Designing Cell-Type-Specific Regulatory DNA with Guided Discrete Diffusion

Animesh Awasthi[1,2,3] **Martin Stoll**[1,2,3] **Raphael Bednarsky**[1,2,3]
Moritz Schaefer[1,2,3] **Christoph Bock**[1,2,3]*

[1]Medical University of Vienna, Institute of Artificial Intelligence,
Center for Medical Data Science, Vienna, Austria

[2]CeMM Research Center for Molecular Medicine
of the Austrian Academy of Sciences, Vienna, Austria

[3]Medical University of Vienna,
Comprehensive Center for AI in Medicine (CAIM), Vienna, Austria

Abstract

AI-driven design of DNA with cell-type-specific regulatory activity is broadly relevant for cell engineering and gene therapy. DNA elements such as promoters and enhancers have gene-regulatory functions that are defined by their DNA sequence and follow a characteristic regulatory grammar. Deep learning models can effectively predict DNA regulatory activity across cell types, but current generative methods often produce DNA sequences that deviate from the native regulatory grammar, increasing the risk of unexpected behavior and off-target effects in real-world applications. Here, we introduce DNA-CRAFT, a method for designing regulatory DNA with high predicted cell-type-specific activity while preserving naturalness. DNA-CRAFT formulates regulatory sequence design as a guided search over a generative model pre-trained on millions of naturally occurring regulatory elements from the human and mouse genomes. Our method enables targeted exploration of the DNA sequence space, efficient evaluation of candidate designs, and explicit suppression of unwanted activity. Across benchmarks spanning human cell lines and primary immune cells, DNA-CRAFT combines high predicted cell-type specificity with strong motif and sequence-composition adherence, achieving favorable results across the evaluated DNA design tasks.

1 Introduction

Precise control of gene expression is at the heart of programmable biology, and offers transformative potential for cell engineering and gene therapy [Dunbar et al., 2018, Kitada et al., 2018, Yang et al., 2025]. A key task is designing regulatory DNA that drives high levels of gene expression in desired cell types while minimizing expression in undesired cell types. Naturally occurring regulatory DNA provides ample evidence that cell-type-specific gene expression is feasible, with combinations of transcription factor binding sites (TFBSs) following the genome's regulatory grammar [Li et al., 2023, Mitra et al., 2024]. However, the catalog of naturally occurring regulatory elements [Moore et al., 2026] is often insufficient for applications in cell engineering and gene therapy, highlighting the need for methods to design synthetic regulatory DNA elements with desired properties.

A key criterion for such methods is that the designed sequences preserve naturalness, i.e., remain consistent with the characteristic TFBS usage that defines the genome's regulatory grammar. While previous work has shown that even random DNA can exhibit regulatory activity [de Boer et al.,

*Correspondence to: Christoph Bock, christoph.bock@meduniwien.ac.at.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

2020], mammalian gene regulation depends on combinations of TFBSs within a supportive DNA sequence and chromatin context [Mitra et al., 2024]. Methods that do not consider the naturalness of the generated DNA risk exploiting biases in predictive models rather than encoding robust regulatory logic [Vaishnav et al., 2022], and may fail to generalize *in vivo* where chromatin state and cellular context strongly modulate activity [Yin et al., 2025].

Current approaches for regulatory sequence design rely on sequence-to-activity models [Avsec et al., 2021, Linder et al., 2025] to predict activity across cell types. Optimization-based approaches, such as gradient-guided search [Schreiber et al., 2025, Linder and Seelig, 2021] and reinforcement learning fine-tuning [Chen et al., 2025b, Wang et al., 2025], can maximize differential activity but may push generated sequences far from the distribution of naturally occurring regulatory DNA [Vaishnav et al., 2022]. Generative approaches such as classifier-free guidance (CFG)-guided discrete diffusion [Sarkar et al., 2025] can preserve naturalness but provide limited steering toward differential activity across cell types. Similar limitations apply to other related work, which we discuss in Appendix A.1. Hence, our central question is: *How can we maximize cell-type-specific regulatory activity while remaining close to the distribution of naturally occurring regulatory DNA?*

Regulatory DNA design can be approached through inference-time search over a pre-trained discrete diffusion model [Sahoo et al., 2024, Lou et al., 2024], using its learned sequence distribution as a prior. This connects inference-time alignment [Uehara et al., 2025] with Monte Carlo tree search (MCTS) [Kocsis and Szepesvári, 2006] to guide diffusion sampling [Tang et al., 2024, Yoon et al., 2026]. Adapting this approach to regulatory DNA design requires three considerations. First, conditioning expansion on the desired regulatory class can focus the search on relevant regions of sequence space. Second, evaluating candidates through full reverse diffusion rollouts is computationally expensive. Third, the reward must explicitly favor target-cell activity while penalizing off-target activity. These considerations motivate a search framework tailored to cell-type-specific regulatory DNA design.

In this work, we introduce **DNA-CRAFT** (**DNA** **C**is-**R**egulatory **A**rchitecture & **F**unction **T**uner), a method for designing regulatory DNA with high predicted cell-type-specific activity while preserving naturalness. DNA-CRAFT combines a discrete diffusion model trained on more than 3.2 million naturally occurring regulatory elements from the human and mouse genomes [Moore et al., 2026] with guided search over its sampling trajectory. The search incorporates three advances: regulatory-class-conditioned expansion via discrete CFG, fast lookahead pruning to efficiently evaluate candidate continuations, and a MinGap specificity objective with a smooth $\mathrm{LogSumExp}$ penalty on off-target activity. Across three human cell lines and eight primary immune cell types, DNA-CRAFT achieves a favorable trade-off between predicted specificity and the measured naturalness metrics. We also provide ablation studies assessing the contributions of its three search components.

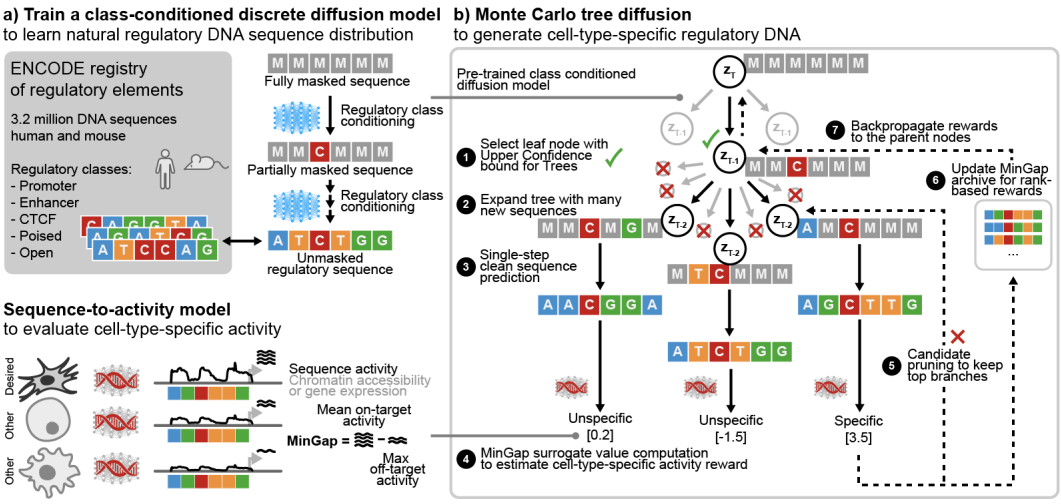

Figure 1: **Overview of DNA-CRAFT. (a)** A class-conditioned discrete diffusion model is trained on the ENCODE registry of naturally occurring regulatory DNA. **(b)** At sampling time, Monte Carlo tree guidance steers the pre-trained class-conditioned diffusion model using cell-type-specific rewards.

Our key contributions are as follows:

- We formulate cell-type-specific regulatory DNA design as an inference-time search problem over a frozen discrete diffusion prior, optimizing differential activity while using the learned distribution of naturally occurring regulatory DNA to guide generation.
- We extend Monte Carlo tree diffusion to regulatory DNA design through class-conditioned expansion, fast lookahead pruning, and a MinGap-based specificity objective.
- We benchmark DNA-CRAFT across cell-type-specific design tasks spanning three human cell lines and eight immune cell types, demonstrating a favorable trade-off between predicted specificity and the measured naturalness metrics.

Source code is available at `https://github.com/epigen/dna-craft`.

2 Background and Problem Formulation

We model regulatory DNA with masked discrete diffusion [Sahoo et al., 2024]. Our implementation of the diffusion model is described in Appendix A.3. Let $\mathcal{V} = \{A, C, G, T\}$ denote the token vocabulary and $\boldsymbol{m}$ denote the mask token (used only during the diffusion process). Each token is represented as a one-hot vector in $\mathbb{R}^{|\mathcal{V}|}$. Unmasked DNA sequences are denoted by $\boldsymbol{x} \in \mathcal{V}^L$, while partially masked sequences are denoted by $\boldsymbol{z}_t \in (\mathcal{V} \cup \{\boldsymbol{m}\})^L$ at continuous time $t \in [0, 1]$. We write $\boldsymbol{x} \sim \mathrm{Cat}(\boldsymbol{p})$ if $\boldsymbol{x}$ is drawn from a categorical distribution with parameter $\boldsymbol{p} \in \Delta^{|\mathcal{V}|}$. For sequences of length L, we write $\boldsymbol{x}^{1:L}$ and $\boldsymbol{z}_t^{1:L}$ (or $\boldsymbol{x}$ and $\boldsymbol{z}_t$ for brevity), with $\boldsymbol{x}^\ell$ and $\boldsymbol{z}_t^\ell$ denoting the ℓ-th token.

2.1 Problem Formulation: Cell-Type-Specific Regulatory DNA Sequence Design

The primary objective in regulatory sequence design is to generate a DNA sequence $\boldsymbol{x}^{1:L}$ that drives high gene expression (and/or chromatin accessibility) in a set of desired cell types while minimizing activity in a set of undesired cell types. To formalize this objective, we assume access to a pre-trained sequence-to-activity model that predicts regulatory activity across several cell types.

Let $\mathcal{C}$ denote the set of all evaluated cell types, and let $\boldsymbol{s}(\boldsymbol{x}^{1:L}) \in \mathbb{R}^{|\mathcal{C}|}$ be the activity vector predicted by the sequence-to-activity model for a sequence $\boldsymbol{x}^{1:L}$. We partition $\mathcal{C}$ into a set of target cell types $\mathcal{C}^+$ and a set of undesired background cell types $\mathcal{C}^-$ and formulate cell-type specificity through two aggregation operators acting on the predicted activities over target and off-target cell types. Let A^+ and A^- denote aggregation operators over $\mathcal{C}^+$ and $\mathcal{C}^-$, respectively. The reward is defined as

$$r(\boldsymbol{x}^{1:L}) = A^+\big(\{s_c(\boldsymbol{x}^{1:L})\}_{c\in\mathcal{C}^+}\big) - A^-\big(\{s_c(\boldsymbol{x}^{1:L})\}_{c\in\mathcal{C}^-}\big)\,.$$

A common choice in prior work is to use mean aggregation over both target and off-target sets, effectively optimizing average activity differences [Reddy et al., 2024]. However, such formulations do not explicitly control for individual off-target cell types, and may allow sequences that exhibit high activity in specific unintended contexts. In many biological applications, even moderate activity in a single off-target cell type can be undesirable, motivating a more conservative notion of specificity. To address this, we adopt a worst-case formulation over off-target activity, leading to the **MinGap** score [Gosai et al., 2024]. Concretely, we set $A^+ = \mathrm{Mean}$ and $A^- = \mathrm{Max}$, yielding

$$r(\boldsymbol{x}^{1:L}) = \underbrace{\frac{1}{|\mathcal{C}^+|}\sum_{c\in\mathcal{C}^+} s_c(\boldsymbol{x}^{1:L})}_{\text{Mean Activity (Target)}} - \underbrace{\max_{c\in\mathcal{C}^-} s_c(\boldsymbol{x}^{1:L})}_{\text{Max Activity (Off-target)}}\,. \tag{1}$$

This objective encourages high target activity relative to the largest predicted off-target response. For search, we replace the maximum with a smooth $\mathrm{LogSumExp}$ aggregation and define

$$g(\boldsymbol{x}^{1:L}) = \frac{1}{|\mathcal{C}^+|}\sum_{c\in\mathcal{C}^+} s_c(\boldsymbol{x}^{1:L}) - \log\left(\sum_{c\in\mathcal{C}^-}\exp(s_c(\boldsymbol{x}^{1:L}))\right). \tag{2}$$

The LogSumExp uses unit temperature without normalization by the number of off-target cell types.

2.2 Inference-Time Alignment and the Specificity-Naturalness Trade-off

We aim to maximize the specificity reward $r(\boldsymbol{x}^{1:L})$ while remaining consistent with the pre-trained diffusion model, leading to the following inference-time alignment objective [Uehara et al., 2025]:

$$\tilde{p}(\boldsymbol{x}^{1:L}) = \arg\max_{p} \; \mathbb{E}_{\boldsymbol{x}^{1:L}\sim p}\big[r(\boldsymbol{x}^{1:L})\big] - \lambda D_{\mathrm{KL}}(p \,\|\, p_\theta) \;\propto\; p_\theta(\boldsymbol{x}^{1:L}) \cdot \exp\left(\frac{r(\boldsymbol{x}^{1:L})}{\lambda}\right), \tag{3}$$

where $\lambda > 0$ controls the trade-off between reward maximization and adherence to the generative prior p_θ. Small λ concentrates the distribution on high-reward sequences, while large λ recovers the prior. Navigating this trade-off is the central challenge in regulatory DNA sequence design.

Soft value function. The target distribution $\tilde{p}$ is defined over fully denoised sequences $\boldsymbol{x}^{1:L}$, while diffusion proceeds through intermediate states $\boldsymbol{z}_t$. To guide this process, we require a measure of the desirability of each intermediate state. This is captured by the *soft value function*, which estimates the expected terminal reward from a given state:

$$v_t(\boldsymbol{z}_t^{1:L}) \coloneqq \lambda \log \mathbb{E}_{\boldsymbol{x}^{1:L}\sim p_\theta(\cdot|\boldsymbol{z}_t^{1:L})}\left[\exp\left(\frac{r(\boldsymbol{x}^{1:L})}{\lambda}\right)\right], \tag{4}$$

which aggregates the rewards of all possible completions of $\boldsymbol{z}_t^{1:L}$, emphasizing high-reward outcomes according to λ.

Soft optimal policy. The value function induces the soft optimal policy [Uehara et al., 2025], which defines the ideal reverse transition for $s < t$:

$$p_s^\star(\boldsymbol{z}_s^{1:L} \mid \boldsymbol{z}_t^{1:L}) \;=\; \frac{p_\theta(\boldsymbol{z}_s^{1:L} \mid \boldsymbol{z}_t^{1:L}) \;\exp\big(v_s(\boldsymbol{z}_s^{1:L})/\lambda\big)}{\exp\big(v_t(\boldsymbol{z}_t^{1:L})/\lambda\big)}. \tag{5}$$

This policy retains the structure of the pre-trained denoising distribution p_θ while re-weighting each candidate continuation $\boldsymbol{z}_s^{1:L}$ by its value v_s. Sequentially sampling from $p_s^\star$ at every denoising step produces sequences distributed according to the target $\tilde{p}$. However, sampling from $p_s^\star$ is intractable for two reasons: (i) evaluating v_t requires marginalizing over all possible completions of $\boldsymbol{z}_t^{1:L}$, an exponentially large space for DNA sequences; and (ii) computing the normalizing constant in the denominator of the soft optimal policy is equally prohibitive.

Approximation strategy. We approximate the soft optimal policy using Monte Carlo Tree Diffusion (MCTD) [Tang et al., 2024, Yoon et al., 2026, 2025], which builds a search tree over intermediate states $\boldsymbol{z}_t^{1:L}$ and uses surrogate value estimates to guide expansion toward high-reward regions while remaining consistent with the generative prior. We note that this results in a heuristic search procedure rather than an exact sampler from the KL-regularized target distribution.

2.3 Monte Carlo Tree Diffusion For Regulatory DNA Sequences

We formulate regulatory DNA design as a search over partially denoised diffusion states. Nodes correspond to sequences $\boldsymbol{z}_t$, and edges represent reverse transitions to $\boldsymbol{z}_s$ with $s < t$. The maximum depth of the resulting tree matches the number of denoising steps. The number of candidate children sampled per expansion at each node is M. MCTD for discrete diffusion undergoes four phases:

1. **Selection:** The algorithm traverses from the root to a leaf node by selecting branches that maximize an Upper Confidence Bound for Trees (UCT) [Kocsis and Szepesvári, 2006], balancing the exploitation of high-reward trajectories with the exploration of the generative prior p_θ.
2. **Expansion:** Upon reaching a leaf node at time t, the tree branches out by sampling M child nodes $(\boldsymbol{z}_{s,1}^{1:L}, \ldots, \boldsymbol{z}_{s,M}^{1:L})$ from the reverse transition $p_\theta(\boldsymbol{z}_s^{1:L} \mid \boldsymbol{z}_t^{1:L})$, without regulatory-class conditioning in this standard formulation.
3. **Rollout (Simulation):** Each child is then iteratively denoised all the way to an unmasked sequence ($t = 0$) using the pre-trained diffusion model, and is scored by the sequence-to-activity model to obtain a Pareto reward (a multi-objective score based on dominance across cell types).
4. **Backpropagation:** The calculated rewards are passed back up the tree, updating the cumulative rewards and visit counts of all parent nodes to inform future selection steps.

3 Methods

DNA-CRAFT builds on the tree-search procedure described above to guide a pre-trained discrete diffusion model toward cell-type-specific regulatory DNA. We first describe the class-conditioned

generative prior, then explain how regulatory-class conditioning, fast lookahead pruning, and the MinGap specificity reward guide candidate generation, selection, and tree updates (Figure 1). The complete algorithm is provided in Appendix A.2.

3.1 Class-Conditioned Discrete Diffusion for Regulatory DNA

To provide a biologically grounded generative prior, we train a class-conditioned masked diffusion language model on the ENCODE registry of candidate regulatory elements, comprising over 3.2 million naturally occurring DNA sequences across the human and mouse genomes [Moore et al., 2026]. We use a bidirectional Mamba [Gu and Dao, 2024] state space model backbone with reverse-complement equivariance that preserves the double-stranded geometry of DNA [Schiff et al., 2024].

We train a class-conditioned denoising network, where $\boldsymbol{y}$ represents one of five consolidated regulatory element classes from the ENCODE registry. Let $\boldsymbol{x}_\theta(\boldsymbol{z}_t, t, \boldsymbol{y})$ denote its clean-nucleotide probabilities at all positions, and let $\boldsymbol{x}_\theta(\boldsymbol{z}_t, t)$ denote predictions from the same network with the null conditioning label. During inference, we apply discrete classifier-free guidance (CFG) [Schiff et al., 2025] to obtain

$$\boldsymbol{x}_\theta^\gamma(\boldsymbol{z}_t, t, \boldsymbol{y}) = \mathrm{softmax}(\gamma \log \boldsymbol{x}_\theta(\boldsymbol{z}_t, t, \boldsymbol{y}) + (1-\gamma) \log \boldsymbol{x}_\theta(\boldsymbol{z}_t, t)) \,. \tag{6}$$

For position $\ell \in \{1, \ldots, L\}$ and nucleotide $a \in \mathcal{V}$, $[\boldsymbol{x}_\theta^\gamma(\boldsymbol{z}_t, t, \boldsymbol{y})]_{\ell,a}$ denotes the predicted probability of nucleotide a at position ℓ. The logarithm acts element-wise, and softmax normalizes over nucleotides at each masked position. These guided clean-token probabilities are used to construct the reverse transition $p_\theta^\gamma(\boldsymbol{z}_s^{1:L} \mid \boldsymbol{z}_t^{1:L}, \boldsymbol{y})$, as detailed in Appendix A.4. Unmasked positions remain unchanged and the denoising network remains frozen during search.

3.2 Class-Conditioned Monte Carlo Tree Guidance

Given the generative prior, we guide sampling using class-conditioned expansion, fast lookahead pruning, and an archive-relative specificity reward.

1. Class-Conditioned Tree Expansion. The genome contains distinct classes of regulatory elements with characteristic sequence composition and motif usage [Friedman et al., 2024]. To focus search on the desired regulatory class, we sample candidate continuations from the CFG-guided reverse transition. When expanding a leaf node $\boldsymbol{z}_t^{1:L}$, we draw a pool of B candidates and retain up to M children after lookahead evaluation. Sampling uses the Gumbel-max construction and unmasked positions remain unchanged. Let $\boldsymbol{G}_b$ contain independent $\mathrm{Gumbel}(0,1)$ draws for every position and token of candidate b. At each masked position ℓ, we sample

$$\boldsymbol{z}_{s,b}^\ell = \underset{a \in \mathcal{V} \cup \{\boldsymbol{m}\}}{\arg\max} \left\{ \log p_\theta^\gamma(\boldsymbol{z}_s^\ell = a \mid \boldsymbol{z}_t, \boldsymbol{y}) + [\boldsymbol{G}_b]_{\ell,a} \right\}. \tag{7}$$

2. Fast Lookahead Pruning.

Candidate pruning. Inspired by value-based selection in inference-time alignment [Uehara et al., 2025, Li et al., 2024], we evaluate the B proposed continuations and retain the M highest-scoring candidates, or all candidates when $B < M$. This concentrates further search on promising continuations within the sampled pool.

Efficient value approximation. Full reverse-diffusion rollouts are expensive when evaluating a large candidate pool. We instead construct a one-step tokenwise modal completion using the denoiser's clean-token predictions:

$$\hat{\boldsymbol{x}}_{0,b}^\ell = \begin{cases} \boldsymbol{z}_{s,b}^\ell, & \boldsymbol{z}_{s,b}^\ell \neq \boldsymbol{m}, \\ \underset{a \in \mathcal{V}}{\arg\max} [\boldsymbol{x}_\theta^\gamma(\boldsymbol{z}_{s,b}, s, \boldsymbol{y})]_{\ell,a}, & \boldsymbol{z}_{s,b}^\ell = \boldsymbol{m}. \end{cases} \tag{8}$$

We refer to this construction as the one-step MAP estimator, where MAP denotes the mode of the factorized clean-token prediction rather than the most likely completion under a full diffusion rollout. The surrogate value is $\hat{v}_s(\boldsymbol{z}_{s,b}^{1:L}) = g(\hat{\boldsymbol{x}}_{0,b}^{1:L})$. Evaluation requires a one-step denoiser prediction followed by a sequence-to-activity model evaluation, avoiding a full reverse-diffusion rollout for each candidate. Appendix A.9 reports weak association with rollout scores at early steps and stronger association later in denoising.

3. MinGap Specificity Reward. We use the smooth specificity score g from Equation 2 to rank candidate completions and retain the highest-scoring children. For tree updates, we compare each retained completion with a bounded archive $\mathcal{G}^*$ of previously admitted record-scoring completions. Candidates are processed in decreasing score order, where the first candidate is inserted directly and subsequent candidates are admitted only when their score matches or exceeds every score currently in the archive. Previously admitted entries remain until the capacity $N_{\max}$ is exceeded, at which point older strictly lower-scoring entries are removed first, with minimum-score removal used otherwise.

Before updating the archive for candidate m, we compute its relative reward:

$$\rho_m = \begin{cases} 1, & \mathcal{G}^* = \emptyset, \\ \dfrac{1}{|\mathcal{G}^*|} \displaystyle\sum_{\boldsymbol{x}^* \in \mathcal{G}^*} \mathbb{I}\big[g(\hat{\boldsymbol{x}}_{0,m}^{1:L}) \geq g(\boldsymbol{x}^*)\big], & \text{otherwise.} \end{cases} \tag{9}$$

The reward lies between zero and one and measures the fraction of archived scores matched or exceeded by the candidate. Since the archive changes during search, these rewards are non-stationary. Candidate pruning uses the smooth score g, while tree selection uses accumulated archive-relative rewards. Complete initialization and backpropagation rules are provided in Appendix A.2.

4 Experiments

4.1 Human Cell-Line-Specific Enhancer Design

Experimental Setup. We utilized a dataset of approximately 700,000 sequences with massively parallel reporter assay (MPRA) activity measurements across three cell lines [Gosai et al., 2024]. We employed a split-model validation strategy following Wang et al. [2025]. To prevent data leakage, we partitioned the MPRA dataset into two halves based on sequence homology scores calculated using MMseqs2 [Steinegger and Söding, 2017]. We fine-tuned the Enformer model [Avsec et al., 2021] on the first half to create a Design-model and on the second half of the data to create the Evaluation-model. Results were further confirmed using independent evaluations (Table 8).

Benchmarking methods. We benchmarked DNA-CRAFT (conditioned on the "Enhancer" class, $\gamma = 1.0$, lookahead pool $B = 800$) against state-of-the-art sequence design methods that use inference-time alignment approaches, including sequential Monte Carlo (SMC) [Phillips et al., 2024] and twisted diffusion sampling (TDS) [Wu et al., 2024] adapted for the MinGap objective, as well as CFG-guided DNA discrete diffusion (D3) [Sarkar et al., 2025]. We also compared against reinforcement learning fine-tuning techniques including DRAKES [Wang et al., 2025] for discrete diffusion and Ctrl-DNA [Chen et al., 2025b] for autoregressive models and the gradient-based optimization method Ledidi [Schreiber et al., 2025]. All methods used the same Design-model. Further implementation details are provided in Appendix A.5.

Evaluation Metrics. We use the MinGap score to assess cell-type specificity, two complementary metrics to assess naturalness and a metric to assess the diversity of the designed sequences. Details of the four metrics are: (i) **MinGap Score.** We compute the MinGap score (Equation 1) to measure differential activity in the desired cell type using the Evaluation-model. (ii) **Motif Correlation.** For each desired cell type, we select reference sequences at or above the 99.9th percentile of MinGap scores computed from measured MPRA activity, corresponding to approximately the top 0.1% of sequences. We scan these top sequences with FIMO [Grant et al., 2011] and the JASPAR 2024 core vertebrate database [Rauluseviciute et al., 2024] to compute TFBS frequency distribution. This serves as the reference for evaluating the biological fidelity of our designed sequences. We compute the Spearman correlation between TFBS frequencies in the generated sequences and the reference. Higher correlation indicates closer agreement with the reference TFBS frequency distribution. (iii) **3-mer Correlation.** We compute the Pearson correlation of 3-mer counts between generated and top reference sequences, capturing sequence composition. (iv) **Diversity.** We calculate the mean per-position Shannon entropy across the generated batch using base-2 logarithms. This metric measures marginal nucleotide diversity, with a maximum of 2 bits per position.

Results. Benchmarking results (Table 1) demonstrate that DNA-CRAFT achieves a favorable trade-off between predicted cell-type specificity and biological fidelity. Across all three cell lines, it combines high differential activity with strong motif and 3-mer correlations. D3 achieves the highest 3-mer correlations and the lowest NTv2-based FBD scores (Table 8), indicating close agreement with

Table 1: Comparison of methods for designing cell-type-specific enhancers across three human cell lines. Shown are the MinGap specificity scores (evaluated under a held-out Evaluation oracle), Motif Spearman correlations, 3-mer Pearson correlations, and Shannon diversity. Values are reported as mean ± std over 3 independent runs and the best result is marked in **bold**.

Cell Line	Metric	SMC	TDS	D3	DRAKES	Ctrl-DNA	Ledidi	DNA-CRAFT (Ours)
HepG2	MinGap Eval ↑	0.989 ±1.717	0.880 ±1.053	-0.075 ±0.021	-1.314 ±0.024	**7.464** ±0.270	6.712 ±0.270	7.271 ±0.142
	Motif Corr. ↑	0.302 ±0.107	0.456 ±0.152	0.793 ±0.009	-0.126 ±0.014	0.545 ±0.063	0.522 ±0.019	**0.810** ±0.028
	3-mer Corr. ↑	0.541 ±0.126	0.735 ±0.083	**0.977** ±0.006	-0.361 ±0.012	0.568 ±0.098	0.711 ±0.010	0.930 ±0.003
	Shannon Div. ↑	0.845 ±0.478	1.301 ±0.105	1.978 ±0.003	1.864 ±0.002	1.914 ±0.033	1.980 ±0.002	**1.981** ±0.001
K562	MinGap Eval ↑	1.812 ±0.768	1.735 ±0.660	0.003 ±0.070	-0.033 ±0.032	**8.490** ±0.153	8.004 ±0.078	7.777 ±0.054
	Motif Corr. ↑	0.269 ±0.130	0.575 ±0.121	0.794 ±0.057	0.060 ±0.008	0.528 ±0.121	0.638 ±0.004	**0.844** ±0.005
	3-mer Corr. ↑	0.591 ±0.059	0.682 ±0.071	**0.969** ±0.013	-0.354 ±0.007	0.388 ±0.043	0.725 ±0.011	0.896 ±0.003
	Shannon Div. ↑	0.487 ±0.302	1.063 ±0.170	1.980 ±0.002	1.958 ±0.003	1.851 ±0.019	**1.982** ±0.001	1.980 ±0.001
SK-N-SH	MinGap Eval ↑	0.252 ±0.161	1.170 ±0.715	-0.111 ±0.014	-0.082 ±0.040	3.988 ±0.856	3.202 ±0.258	**5.479** ±0.057
	Motif Corr. ↑	0.454 ±0.072	0.435 ±0.008	0.766 ±0.022	-0.001 ±0.009	0.580 ±0.042	0.502 ±0.053	**0.809** ±0.010
	3-mer Corr. ↑	0.681 ±0.087	0.658 ±0.135	**0.932** ±0.017	-0.382 ±0.001	0.431 ±0.074	0.387 ±0.080	0.912 ±0.001
	Shannon Div. ↑	1.029 ±0.091	1.019 ±0.327	1.969 ±0.004	1.826 ±0.001	1.847 ±0.066	**1.982** ±0.001	1.972 ±0.001

the reference sequence distribution, but its near-zero or negative MinGap scores indicate poor steering toward cell-type specificity. Ctrl-DNA achieves higher MinGap than DNA-CRAFT on HepG2 and K562, while Ledidi also achieves higher MinGap than DNA-CRAFT on K562. DNA-CRAFT maintains higher motif correlations than either method. On SK-N-SH, DNA-CRAFT achieves the highest MinGap (5.479) and motif correlation (0.809), and retains the highest specificity under the alternative evaluation metrics LogSumExp and Mean Difference. Figure 5 in the appendix visualizes the specificity-naturalness trade-off.

4.2 Designing Immune Cell-Type-Specific Sequences

We further evaluated DNA-CRAFT's ability to design regulatory DNA for another data type (chromatin accessibility) and for cell types that are similar to each other (eight types of immune cells: $CD4^+$ T cells, $CD8^+$ T cells, Macrophages, Mast cells, Memory B cells, Natural Killer (NK) T cells, Naive T cells, and Plasma B cells).

Experimental Setup. We utilized an Enformer model fine-tuned on single-cell ATAC-seq profiles from the CATlas dataset [Zhang et al., 2021, Lal et al., 2025] as our sequence-to-activity model. As in Section 4.1, we trained separate Design and Evaluation models on non-overlapping homology clusters to prevent data leakage between optimization and evaluation. All methods generate sequences using the Design oracle. The final MinGap scores are reported using the held-out Evaluation oracle. We performed eight separate design experiments, configuring the objective to target one cell type while penalizing activity in the remaining seven. Evaluation followed the metrics defined in Section 4.1. Natural differentially accessible regions specific to each target cell type were used as references to calculate motif and 3-mer correlations. We benchmarked DNA-CRAFT (conditioned on "Enhancer", $\gamma = 1.0$, lookahead size $B = 800$) against the same methods as the previous benchmark, except D3 and DRAKES, due to the computational cost of fine-tuning/training these methods on each task.

Results. Decoupling chromatin accessibility among closely related cell types (e.g., $CD4^+$ versus $CD8^+$ T cells) is a challenging generative task. As shown in Table 2, DNA-CRAFT is the only method that achieves positive predicted differential accessibility (MinGap) across all eight cell types, while maintaining high motif correlations. Ctrl-DNA achieves positive MinGap for several cell types, but has lower motif correlations than DNA-CRAFT, particularly for $CD4^+$ and $CD8^+$ T cells. We further evaluated DNA-CRAFT on eight brain cell types from the CATlas dataset, with the experimental setup and complete results provided in Appendix A.5.1.

4.3 Ablation Studies

To isolate the impact of our algorithmic components, we conducted ablation studies using the SK-N-SH cell line enhancer design task. We generated 128 sequences per configuration and evaluated the trade-offs between computational efficiency (wall time), differential activity, and realism across 3 independent runs. We selected SK-N-SH for ablations because the highest mean MinGap attained by any evaluated method was lower than for HepG2 or K562 (Table 1), motivating its use as a challenging test case.

Class conditioning improves specificity while retaining motif fidelity. We evaluated class-conditioned tree expansion by sweeping the CFG scale $\gamma \in \{0.0, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$. At

Table 2: Comparison of methods for cell-type-specific regulatory sequence design across eight immune cell types. Shown are the MinGap specificity scores (evaluated under a held-out Evaluation oracle), Motif Spearman correlations, 3-mer Pearson correlations, and Shannon diversity. Values are reported as mean ± std over three independent runs and the best result is marked in **bold**.

Target Cell Type	Metric	SMC	TDS	Ctrl-DNA	Ledidi	DNA-CRAFT (Ours)
CD4$^+$ T Cell	MinGap ↑	-0.065 ±0.069	-0.016 ±0.002	-0.047 ±0.005	-0.105 ±0.006	**0.075** ±0.001
	Motif Corr. ↑	0.478 ±0.048	0.363 ±0.045	0.261 ±0.031	0.094 ±0.155	**0.616** ±0.024
	3-mer Corr. ↑	**0.811** ±0.056	0.802 ±0.073	0.461 ±0.035	0.479 ±0.039	0.766 ±0.007
	Shannon Div. ↑	1.227 ±0.160	1.339 ±0.285	1.779 ±0.022	**1.983** ±0.001	1.919 ±0.002
CD8$^+$ T Cell	MinGap ↑	-0.022 ±0.019	-0.022 ±0.018	**0.084** ±0.004	-0.048 ±0.003	0.059 ±0.008
	Motif Corr. ↑	0.443 ±0.158	0.360 ±0.302	0.076 ±0.054	0.084 ±0.215	**0.680** ±0.055
	3-mer Corr. ↑	**0.726** ±0.115	0.716 ±0.137	0.221 ±0.010	0.045 ±0.140	0.558 ±0.010
	Shannon Div. ↑	1.088 ±0.282	1.400 ±0.097	1.684 ±0.018	**1.983** ±0.001	1.918 ±0.002
Macrophage	MinGap ↑	0.056 ±0.027	0.053 ±0.036	0.423 ±0.123	0.072 ±0.013	**0.947** ±0.001
	Motif Corr. ↑	0.498 ±0.048	0.447 ±0.079	0.591 ±0.088	0.147 ±0.154	**0.827** ±0.004
	3-mer Corr. ↑	0.660 ±0.207	0.825 ±0.027	0.430 ±0.050	0.079 ±0.150	**0.898** ±0.003
	Shannon Div. ↑	1.275 ±0.125	1.429 ±0.076	1.738 ±0.005	**1.983** ±0.001	1.971 ±0.000
Mast Cell	MinGap ↑	-0.047 ±0.038	-0.021 ±0.012	0.852 ±0.099	-0.056 ±0.007	**0.886** ±0.002
	Motif Corr. ↑	0.474 ±0.082	0.540 ±0.080	0.444 ±0.070	0.113 ±0.182	**0.881** ±0.004
	3-mer Corr. ↑	0.800 ±0.080	0.856 ±0.000	0.395 ±0.123	0.085 ±0.150	**0.912** ±0.006
	Shannon Div. ↑	1.482 ±0.012	1.397 ±0.179	1.835 ±0.086	**1.983** ±0.001	1.982 ±0.001
Memory B Cell	MinGap ↑	-0.042 ±0.030	-0.055 ±0.055	-0.044 ±0.007	-0.151 ±0.005	**0.467** ±0.021
	Motif Corr. ↑	0.392 ±0.026	0.591 ±0.038	0.287 ±0.203	0.076 ±0.181	**0.860** ±0.003
	3-mer Corr. ↑	0.785 ±0.048	0.798 ±0.023	0.659 ±0.147	0.063 ±0.139	**0.937** ±0.006
	Shannon Div. ↑	1.236 ±0.264	1.330 ±0.187	1.950 ±0.008	**1.983** ±0.001	1.982 ±0.002
NK T Cell	MinGap ↑	-0.039 ±0.023	-0.016 ±0.005	0.130 ±0.039	-0.122 ±0.005	**0.344** ±0.006
	Motif Corr. ↑	0.384 ±0.124	0.353 ±0.271	0.315 ±0.026	0.078 ±0.202	**0.804** ±0.006
	3-mer Corr. ↑	0.780 ±0.065	0.742 ±0.121	0.455 ±0.037	0.043 ±0.161	**0.784** ±0.013
	Shannon Div. ↑	1.431 ±0.033	1.401 ±0.096	1.871 ±0.023	**1.983** ±0.001	1.963 ±0.002
Naive T Cell	MinGap ↑	-0.031 ±0.017	-0.064 ±0.037	0.229 ±0.031	-0.126 ±0.007	**0.610** ±0.003
	Motif Corr. ↑	0.422 ±0.123	0.410 ±0.080	0.644 ±0.052	0.105 ±0.154	**0.859** ±0.006
	3-mer Corr. ↑	0.618 ±0.145	0.769 ±0.085	0.684 ±0.038	0.061 ±0.128	**0.942** ±0.004
	Shannon Div. ↑	1.413 ±0.103	1.361 ±0.230	1.906 ±0.040	**1.983** ±0.001	1.981 ±0.001
Plasma B Cell	MinGap ↑	-0.019 ±0.014	-0.037 ±0.019	0.458 ±0.063	-0.081 ±0.001	**0.785** ±0.007
	Motif Corr. ↑	0.454 ±0.072	0.348 ±0.277	0.482 ±0.057	0.024 ±0.165	**0.827** ±0.006
	3-mer Corr. ↑	0.800 ±0.034	0.643 ±0.038	0.433 ±0.033	0.042 ±0.167	**0.879** ±0.010
	Shannon Div. ↑	1.310 ±0.098	1.138 ±0.268	1.806 ±0.071	**1.983** ±0.001	1.969 ±0.001

$\gamma = 0$, sampling uses the null conditioning label of the same network. Mean MinGap generally increases across the sweep, from 5.440 at $\gamma = 0$ to 5.644 at $\gamma = 5$, with a small decrease between $\gamma = 2$ and $\gamma = 3$ (Figure 2a). Motif correlation remains between 0.803 and 0.821 without a monotonic trend. These results show improved predicted specificity at higher guidance scales while retaining similar motif-frequency correlations.

Fast lookahead pruning increases specificity while reducing compute cost. To quantify the computational and performance benefits of the fast lookahead pruning, we swept the lookahead batch size $B \in \{2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048\}$ with a retained-child limit of $M = 8$, giving $\min(B, 8)$ retained children per expansion and compared it to standard full rollouts (Std.). Our results (Figure 2b) demonstrate that as pool size B scales, the mean MinGap specificity improves monotonically from 1.554 to 5.590, while motif and 3-mer correlations vary. The one-step lookahead configurations achieve higher mean specificity than the full-rollout baseline, although their candidate budgets differ. The full-rollout baseline retains higher motif and 3-mer correlations, although these differences are modest at $B = 256$: 0.858 versus 0.828 for motif correlation and 0.947 versus 0.924 for 3-mer correlation (Table 13). Crucially, even with a massive candidate pool of $B = 256$, the method achieves a high MinGap of 5.079 in just 368 seconds. In contrast, standard MCTS ($M = B$) relying on full iterative rollouts achieves a much lower MinGap score of 1.538 at four times the compute time of 1459 seconds. This shows that our fast lookaheads improve both the specificity and the computational efficiency of the tree search.

Specificity reward formulation. We compared four reward formulations within the tree search: Pareto dominance, mean activity difference, hard MinGap, and our $\mathrm{LogSumExp}$ MinGap with a relative archive (Table 3). The Pareto reward achieves strong motif correlation (0.810) and the highest 3-mer correlation (0.941), while yielding lower cell-type specificity (MinGap 0.495). All three scalar

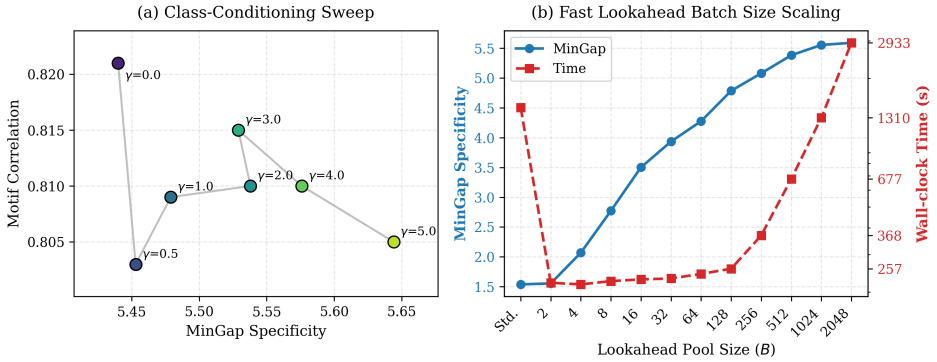

Figure 2: Ablation of DNA-CRAFT class-conditioned expansion and lookahead mechanisms for SK-N-SH enhancer design. **(a)** Class-conditioning sweep across guidance scales (γ) with MinGap specificity and motif correlation reported. **(b)** Fast lookahead pruning batch size (B) scaling compared to standard full rollouts (Std.), with MinGap specificity and compute wall time reported.

Table 3: Evaluating different specificity reward formulations during the Monte Carlo Tree Guidance search ($M = 8$, $B = 800$) for SK-N-SH enhancer design. Values are reported as mean $\pm$ standard deviation over 3 independent runs and the best result is marked in **bold**.

Reward Formulation	**MinGap** ($\uparrow$)	**Mean Diff.** ($\uparrow$)	**Motif Corr.** ($\uparrow$)	**3-mer Corr.** ($\uparrow$)
Pareto Reward	0.495 $\pm$0.059	0.643 $\pm$0.068	0.810 $\pm$0.028	**0.941** $\pm$0.010
Mean activity difference	4.551 $\pm$0.035	6.124 $\pm$0.023	**0.814** $\pm$0.009	0.919 $\pm$0.000
MinGap	**5.531** $\pm$0.056	5.865 $\pm$0.050	0.801 $\pm$0.019	0.891 $\pm$0.006
MinGap LogSumExp	5.479 $\pm$0.057	**6.134** $\pm$0.035	0.809 $\pm$0.010	0.912 $\pm$0.001

specificity rewards achieve higher mean MinGap, highlighting the benefit of directly optimizing the separation between target and off-target activity. Among scalar formulations, mean activity difference achieves a MinGap of 4.551, while hard MinGap achieves the highest specificity (5.531) with lower mean differential activity (5.865) and 3-mer correlation (0.891). Our LogSumExp MinGap with rank-based archive normalization (Equation 9) achieves a favorable balance, with comparable specificity (5.479), the highest reported mean differential activity (6.134), and strong naturalness (motif correlation 0.809, 3-mer correlation 0.912). These results illustrate the trade-off between worst-case specificity, mean differential activity, and the measured naturalness metrics.

5 Conclusion

We introduced DNA-CRAFT, a method for designing regulatory DNA with high predicted cell-type-specific activity while preserving naturalness. DNA-CRAFT formulates sequence design as inference-time search over the denoising trajectory of a frozen discrete diffusion model, combining class-conditioned expansion, fast lookahead pruning, and a MinGap specificity reward. Across benchmarks spanning human cell lines and primary immune cell types, DNA-CRAFT achieves a favorable trade-off between predicted specificity and naturalness metrics. Ablation studies delineate how the search components and reward formulation influence specificity, naturalness, and computational cost. Limitations include dependence on sequence-to-activity models and evaluation against few off-target cell types. The designs have not been experimentally validated. Future directions include integrating larger pre-trained DNA foundation models as the generative backbone, extending the framework to additional tissues and species, and experimentally validating designs for clinically relevant cell types.

Acknowledgments and Disclosure of Funding

Broader Impact. Methods for designing regulatory DNA may support basic research, cell engineering, and therapeutic development. Their outputs remain subject to prediction errors and may exhibit unintended activity outside the evaluated cell types. Experimental validation and appropriate oversight are therefore necessary before downstream use.

Acknowledgements. We thank the members of the Bock Lab and the AI Institute of the Medical University of Vienna for their help and advice. We also thank Albert López i Serrano for helpful discussions.

Funding. C.B. is supported by a European Research Council (ERC) Consolidator Grant (101001971). Computational resources were provided by CeMM with infrastructure supported by CeMM's IT Services team.

Competing interests. C.B. is a cofounder and scientific advisor of Myllia Biotechnology and Neurolentech. The remaining authors declare no competing interests.

References

Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy Colwell. Model-based reinforcement learning for biological sequence design. September 2019. URL https://openreview.net/forum?fileGuid=3xgr169o12oUrbxS&id=HklxbgBKvr&ref=https%3A%2F%2Fgithubhelp.com.

Pavel Avdeyev, Chenlai Shi, Yuhao Tan, Kseniia Dudnyk, and Jian Zhou. Dirichlet diffusion score model for biological sequence generation. In *International Conference on Machine Learning*, pages 1276–1301. PMLR, 2023.

Žiga Avsec, Vikram Agarwal, Daniel Visentin, Joseph R Ledsam, Agnieszka Grabska-Barwinska, Kyle R Taylor, Yannis Assael, John Jumper, Pushmeet Kohli, and David R Kelley. Effective gene expression prediction from sequence by integrating long-range interactions. *Nature methods*, 18 (10):1196–1203, 2021.

Tong Chen, Yinuo Zhang, Sophia Tang, and Pranam Chatterjee. Multi-objective-guided discrete flow matching for controllable biological sequence design, 2025a. URL https://arxiv.org/abs/2505.07086.

Xingyu Chen, Shihao Ma, Runsheng Lin, Jiecong Lin, and Bo Wang. Ctrl-DNA: Controllable Cell-Type-Specific Regulatory DNA Design via Constrained RL, May 2025b. URL http://arxiv.org/abs/2505.20578. arXiv:2505.20578 [cs].

Hugo Dalla-Torre, Liam Gonzalez, Javier Mendoza-Revilla, Nicolas Lopez Carranza, Adam Henryk Grzywaczewski, Francesco Oteri, Christian Dallago, Evan Trop, Bernardo P. de Almeida, Hassan Sirelkhatim, Guillaume Richard, Marcin Skwark, Karim Beguir, Marie Lopez, and Thomas Pierrot. Nucleotide Transformer: building and evaluating robust foundation models for human genomics. *Nature Methods*, 22(2):287–297, February 2025. ISSN 1548-7105. doi: 10.1038/s41592-024-02523-z. URL https://www.nature.com/articles/s41592-024-02523-z.

Lucas Ferreira DaSilva, Simon Senan, Judith F. Kribelbauer-Swietek, Zain Munir Patel, Lithin Karmel Louis, Aniketh Janardhan Reddy, Sameer Gabbita, Jonathan D. Rosen, Zach Nussbaum, César Miguel Valdez Córdova, Aaron Wenteler, Noah Weber, Tin M. Tunjic, Martino Mansoldo, Talha Ahmad Khan, Gue-Ho Hwang, Vincent Gardeux, David T. Humphreys, Cameron Smith, Matei Bejan, Peter Bromley, Will Connell, Bart Deplancke, Michael I. Love, Emily S. Wong, Wouter Meuleman, and Luca Pinello. Designing synthetic regulatory elements using the generative AI framework DNA-Diffusion. *Nature Genetics*, 58(1):180–194, January 2026. ISSN 1546-1718. doi: 10.1038/s41588-025-02441-6. URL https://www.nature.com/articles/s41588-025-02441-6. Publisher: Nature Publishing Group.

Carl G de Boer, Eeshit Dhaval Vaishnav, Ronen Sadeh, Esteban Luis Abeyta, Nir Friedman, and Aviv Regev. Deciphering eukaryotic gene-regulatory logic with 100 million random promoters. *Nature biotechnology*, 38(1):56–65, 2020.

Cynthia E. Dunbar, Katherine A. High, J. Keith Joung, Donald B. Kohn, Keiya Ozawa, and Michel Sadelain. Gene therapy comes of age. *Science*, 359(6372):eaan4672, January 2018. doi: 10.1126/science.aan4672. URL https://www.science.org/doi/10.1126/science.aan4672. Publisher: American Association for the Advancement of Science.

Meyer J. Friedman, Tobias Wagner, Haram Lee, Michael G. Rosenfeld, and Soohwan Oh. Enhancer–promoter specificity in gene transcription: molecular mechanisms and disease associations. *Experimental & Molecular Medicine*, 56(4):772–787, April 2024. ISSN 2092-6413. doi: 10.1038/s12276-024-01233-y. URL https://www.nature.com/articles/s12276-024-01233-y. Publisher: Nature Publishing Group.

Sager J Gosai, Rodrigo I Castro, Natalia Fuentes, John C Butts, Kousuke Mouri, Michael Alasoadura, Susan Kales, Thanh Thanh L Nguyen, Ramil R Noche, Arya S Rao, et al. Machine-guided design of cell-type-targeting cis-regulatory elements. *Nature*, pages 1–10, 2024.

Charles E. Grant, Timothy L. Bailey, and William Stafford Noble. FIMO: scanning for occurrences of a given motif. *Bioinformatics*, 27(7):1017–1018, April 2011. ISSN 1367-4803. doi: 10.1093/bioinformatics/btr064. URL https://doi.org/10.1093/bioinformatics/btr064.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2024. URL https://arxiv.org/abs/2312.00752.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. URL https://arxiv.org/abs/1706.08500.

Moksh Jain, Emmanuel Bengio, Alex Hernandez-Garcia, Jarrid Rector-Brooks, Bonaventure FP Dossou, Chanakya Ajit Ekbote, Jie Fu, Tianyu Zhang, Michael Kilgour, Dinghuai Zhang, et al. Biological sequence design with gflownets. In *International Conference on Machine Learning*, pages 9786–9801. PMLR, 2022.

W. James Kent, Charles W. Sugnet, Terrence S. Furey, Krishna M. Roskin, Tom H. Pringle, Alan M. Zahler, and and David Haussler. The Human Genome Browser at UCSC. *Genome Research*, 12(6):996–1006, June 2002. ISSN 1088-9051, 1549-5469. doi: 10.1101/gr.229102. URL http://genome.cshlp.org/content/12/6/996. Company: Cold Spring Harbor Laboratory Press Distributor: Cold Spring Harbor Laboratory Press Institution: Cold Spring Harbor Laboratory Press Label: Cold Spring Harbor Laboratory Press Publisher: Cold Spring Harbor Lab.

Tasuku Kitada, Breanna DiAndreth, Brian Teague, and Ron Weiss. Programming gene and engineered-cell therapies with synthetic biology. *Science*, 359(6376):eaad1067, 2018. doi: 10.1126/science.aad1067. URL https://www.science.org/doi/abs/10.1126/science.aad1067.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European Conference on Machine Learning*, 2006.

P. J. van Laarhoven and E. H. Aarts. *Simulated Annealing: Theory and Applications*. Springer Science & Business Media, June 1987. ISBN 978-90-277-2513-4.

Avantika Lal, David Garfield, Tommaso Biancalani, and Gokcen Eraslan. Designing realistic regulatory dna with autoregressive language models. *Genome Research*, 34(9):1411–1420, 2024.

Avantika Lal, Laura Gunsalus, Surag Nair, Tommaso Biancalani, and Gokcen Eraslan. gReLU: a comprehensive framework for DNA sequence modeling and design. *Nature Methods*, 22 (11):2253–2257, November 2025. ISSN 1548-7105. doi: 10.1038/s41592-025-02868-z. URL https://www.nature.com/articles/s41592-025-02868-z. Publisher: Nature Publishing Group.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-Free Guidance in Continuous and Discrete Diffusion Models with Soft Value-Based Decoding, October 2024. URL http://arxiv.org/abs/2408.08252. arXiv:2408.08252 [cs].

Yang Eric Li, Sebastian Preissl, Michael Miller, Nicholas D. Johnson, Zihan Wang, Henry Jiao, Chenxu Zhu, Zhaoning Wang, Yang Xie, Olivier Poirion, Colin Kern, Antonio Pinto-Duarte, Wei Tian, Kimberly Siletti, Nora Emerson, Julia Osteen, Jacinta Lucero, Lin Lin, Qian Yang, Quan Zhu, Nathan Zemke, Sarah Espinoza, Anna Marie Yanny, Julie Nyhus, Nick Dee, Tamara Casper, Nadiya Shapovalova, Daniel Hirschstein, Rebecca D. Hodge, Sten Linnarsson, Trygve Bakken, Boaz Levi, C. Dirk Keene, Jingbo Shang, Ed Lein, Allen Wang, M. Margarita Behrens,

Joseph R. Ecker, and Bing Ren. A comparative atlas of single-cell chromatin accessibility in the human brain. *Science*, 382(6667):eadf7044, 2023. doi: 10.1126/science.adf7044. URL https://www.science.org/doi/abs/10.1126/science.adf7044.

Johannes Linder and Georg Seelig. Fast activation maximization for molecular sequence design. *BMC bioinformatics*, 22:1–20, 2021.

Johannes Linder, Divyanshi Srivastava, Han Yuan, Vikram Agarwal, and David R. Kelley. Predicting RNA-seq coverage from DNA sequence as a unifying model of gene regulation. *Nature Genetics*, 57(4):949–961, April 2025. ISSN 1546-1718. doi: 10.1038/s41588-024-02053-6. URL https://www.nature.com/articles/s41588-024-02053-6. Publisher: Nature Publishing Group.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution, 2024. URL https://arxiv.org/abs/2310.16834.

Sneha Mitra, Rohan Malik, Wilfred Wong, Afsana Rahman, Alexander J. Hartemink, Yuri Pritykin, Kushal K. Dey, and Christina S. Leslie. Single-cell multi-ome regression models identify functional and disease-associated enhancers and enable chromatin potential analysis. *Nature Genetics*, 56 (4):627–636, April 2024. ISSN 1546-1718. doi: 10.1038/s41588-024-01689-8. URL https://www.nature.com/articles/s41588-024-01689-8. Publisher: Nature Publishing Group.

Jill E. Moore, Henry E. Pratt, Kaili Fan, Nishigandha Phalke, Jonathan Fisher, Shaimae I. Elhajjajy, Gregory Andrews, Mingshi Gao, Nicole Shedd, Yu Fu, Matthew C. Lacadie, Jair Meza, Mansi Khandpekar, Mohit Ganna, Eva Choudhury, Ross Swofford, Huong Phan, Christian C. Ramirez, Maxwell Campbell, Mary Likhite, Nina P. Farrell, Annika K. Weimer, Anusri Pampari, Vivekanandan Ramalingam, Fairlie Reese, Beatrice Borsari, Xuezhu Yu, Eve Wattenberg, Marina Ruiz-Romero, Milad Razavi-Mohseni, Jinrui Xu, Timur Galeev, Andres Colubri, Michael A. Beer, Roderic Guigó, Mark B. Gerstein, Jesse M. Engreitz, Mats Ljungman, Timothy E. Reddy, Michael P. Snyder, Charles B. Epstein, Elizabeth Gaskell, Bradley E. Bernstein, Diane E. Dickel, Axel Visel, Len A. Pennacchio, Ali Mortazavi, Anshul Kundaje, and Zhiping Weng. An expanded registry of candidate cis-regulatory elements. *Nature*, pages 1–10, January 2026. ISSN 1476-4687. doi: 10.1038/s41586-025-09909-9. URL https://www.nature.com/articles/s41586-025-09909-9. Publisher: Nature Publishing Group.

Eric Nguyen, Michael Poli, Marjan Faizi, Armin Thomas, Michael Wornow, Callum Birch-Sykes, Stefano Massaroli, Aman Patel, Clayton Rabideau, Yoshua Bengio, et al. Hyenadna: Long-range genomic sequence modeling at single nucleotide resolution. *Advances in neural information processing systems*, 36:43177–43201, 2023.

Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking Guidance for Discrete State-Space Diffusion and Flow Models, March 2025. URL http://arxiv.org/abs/2406.01572. arXiv:2406.01572 [cs].

Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Particle Denoising Diffusion Sampler, June 2024. URL http://arxiv.org/abs/2402.06320. arXiv:2402.06320 [stat].

Ieva Rauluseviciute, Rafael Riudavets-Puig, Romain Blanc-Mathieu, Jaime A Castro-Mondragon, Katalin Ferenc, Vipin Kumar, Roza Berhanu Lemma, Jérémy Lucas, Jeanne Chéneby, Damir Baranasic, Aziz Khan, Oriol Fornes, Sveinung Gundersen, Morten Johansen, Eivind Hovig, Boris Lenhard, Albin Sandelin, Wyeth W Wasserman, Franccois Parcy, and Anthony Mathelier. JASPAR 2024: 20th anniversary of the open-access database of transcription factor binding profiles. *Nucleic Acids Research*, 52(D1):D174–D182, January 2024. ISSN 0305-1048. doi: 10.1093/nar/gkad1059. URL https://doi.org/10.1093/nar/gkad1059.

Aniketh Janardhan Reddy, Xinyang Geng, Michael Herschl, Sathvik Kolli, Aviral Kumar, Patrick Hsu, Sergey Levine, and Nilah Ioannidis. Designing cell-type-specific promoter sequences using conservative model-based optimization. *Advances in Neural Information Processing Systems*, 37: 93033–93059, 2024.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and Effective Masked Diffusion Language Models, November 2024. URL http://arxiv.org/abs/2406.07524. arXiv:2406.07524 [cs].

Anirban Sarkar, Yijie Kang, Nirali Somia, Pablo Mantilla Puccetti, Jessica Zhou, Masayuki Nagai, Ziqi Tang, Chris Zhao, and Peter K. Koo. Designing DNA With Tunable Regulatory Activity Using Score-Entropy Discrete Diffusion, May 2025. URL https://www.biorxiv.org/content/10.1101/2024.05.23.595630v2. Pages: 2024.05.23.595630 Section: New Results.

Yair Schiff, Chia-Hsiang Kao, Aaron Gokaslan, Tri Dao, Albert Gu, and Volodymyr Kuleshov. Caduceus: Bi-Directional Equivariant Long-Range DNA Sequence Modeling, June 2024. URL http://arxiv.org/abs/2403.03234. arXiv:2403.03234 [q-bio].

Yair Schiff, Subham Sekhar Sahoo, Hao Phung, Guanghan Wang, Sam Boshar, Hugo Dalla-torre, Bernardo P. de Almeida, Alexander Rush, Thomas Pierrot, and Volodymyr Kuleshov. Simple Guidance Mechanisms for Discrete Diffusion Models, May 2025. URL http://arxiv.org/abs/2412.10193. arXiv:2412.10193 [cs].

Jacob Schreiber, Franziska Katharina Lorbeer, Monika Heinzl, Franziska Reiter, Baptiste Rafanel, Yang Young Lu, Alexander Stark, and William Stafford Noble. Programmatic design and editing of cis-regulatory elements, December 2025. URL https://www.biorxiv.org/content/10.1101/2025.04.22.650035v2. ISSN: 2692-8205 Pages: 2025.04.22.650035 Section: New Results.

Sam Sinai, Richard Wang, Alexander Whatley, Stewart Slocum, Elina Locane, and Eric D Kelsic. Adalead: A simple and robust adaptive greedy search algorithm for sequence design. *arXiv preprint arXiv:2010.02141*, 2020.

Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design, 2024. URL https://arxiv.org/abs/2402.05841.

Martin Steinegger and Johannes Söding. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature Biotechnology*, 35(11):1026–1028, November 2017. ISSN 1546-1696. doi: 10.1038/nbt.3988. URL https://www.nature.com/articles/nbt.3988.

Sophia Tang, Yinuo Zhang, and Pranam Chatterjee. PepTune: De Novo Generation of Therapeutic Peptides with Multi-Objective-Guided Discrete Diffusion, December 2024. URL https://arxiv.org/abs/2412.17780v4.

Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tommaso Biancalani. Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review, 2025. URL https://arxiv.org/abs/2501.09685.

Eeshit Dhaval Vaishnav, Carl G de Boer, Jennifer Molinet, Moran Yassour, Lin Fan, Xian Adiconis, Dawn A Thompson, Joshua Z Levin, Francisco A Cubillos, and Aviv Regev. The evolution, evolvability and engineering of gene regulatory dna. *Nature*, 603(7901):455–463, 2022.

Chenyu Wang, Masatoshi Uehara, Yichun He, Amy Wang, Tommaso Biancalani, Avantika Lal, Tommi Jaakkola, Sergey Levine, Hanchen Wang, and Aviv Regev. Fine-Tuning Discrete Diffusion Models via Reward Optimization with Applications to DNA and Protein Design, March 2025. URL http://arxiv.org/abs/2410.13643. arXiv:2410.13643 [cs].

Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham. Practical and Asymptotically Exact Conditional Sampling in Diffusion Models, November 2024. URL http://arxiv.org/abs/2306.17775. arXiv:2306.17775 [stat].

Jinshou Yang, Feihan Zhou, Xiyuan Luo, Yuan Fang, Xing Wang, Xiaohong Liu, Ruiling Xiao, Decheng Jiang, Yuemeng Tang, Gang Yang, Lei You, and Yupei Zhao. Enhancer reprogramming: critical roles in cancer and promising therapeutic strategies. *Cell Death Discovery*, 11(1):84, March 2025. ISSN 2058-7716. doi: 10.1038/s41420-025-02366-3. URL https://www.nature.com/articles/s41420-025-02366-3. Publisher: Nature Publishing Group.

Christopher Yin, Sebastian Castillo-Hair, Gun Woo Byeon, Peter Bromley, Wouter Meuleman, and Georg Seelig. Iterative deep learning design of human enhancers exploits condensed sequence grammar to achieve cell-type specificity. 16(7), 2025. ISSN 2405-4712, 2405-4720. doi: 10.1016/j.cels.2025.101302. URL https://www.cell.com/cell-systems/abstract/S2405-4712(25)00135-8.

Jaesik Yoon, Hyeonseo Cho, Yoshua Bengio, and Sungjin Ahn. Fast monte carlo tree diffusion: 100x speedup via parallel sparse planning, 2025. URL https://arxiv.org/abs/2506.09498.

Jaesik Yoon, Hyeonseo Cho, Doojin Baek, Yoshua Bengio, and Sungjin Ahn. Monte carlo tree diffusion for system 2 planning, 2026. URL https://arxiv.org/abs/2502.07202.

Kai Zhang, James D. Hocker, Michael Miller, Xiaomeng Hou, Joshua Chiou, Olivier B. Poirion, Yunjiang Qiu, Yang E. Li, Kyle J. Gaulton, Allen Wang, Sebastian Preissl, and Bing Ren. A single-cell atlas of chromatin accessibility in the human genome. *Cell*, 184(24):5985–6001.e19, November 2021. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2021.10.024. URL https://www.cell.com/cell/abstract/S0092-8674(21)01279-4. Publisher: Elsevier.

A Appendix

The appendix provides related work (Appendix A.1), the complete DNA-CRAFT algorithm (Appendix A.2), and diffusion model training details and test-set evaluations (Appendix A.3). Appendix A.4 describes class-conditioned sampling and motif analysis. Appendix A.5 provides benchmark implementation details, extended evaluations including eight brain cell types, and training and validation results for the sequence-to-activity models. Appendix A.6 reports the full ablation results and cross-method evaluation under alternative specificity metrics. Finally, Appendices A.7, A.8, and A.9 present the comparison with a tree-free greedy baseline, tree depth statistics, and evaluation of the one-step MAP estimator, respectively.

A.1 Related Work

Machine learning methods for regulatory element design rely on sequence-to-activity neural networks that predict gene expression or chromatin accessibility from DNA sequences [Avsec et al., 2021, Linder et al., 2025]. By generalizing to unseen sequences, these models serve as differentiable sequence-to-activity models that enable computational design of regulatory elements with desired activity profiles. Existing approaches generally fall into two categories: sequence optimization methods and generative model-based methods.

Sequence Optimization Methods. These methods treat regulatory DNA design as a search problem over a combinatorial sequence space, using a pre-trained sequence-to-activity model as the objective function. Classical approaches, including simulated annealing [Laarhoven and Aarts, 1987], *in silico* directed evolution, AdaLead [Sinai et al., 2020], and gradient-based editing [Schreiber et al., 2025, Linder and Seelig, 2021], start from random or known sequences and iteratively mutate them to maximize predicted activity. More recently, reinforcement learning techniques such as DyNA-PPO [Angermueller et al., 2019] and GFlowNets [Jain et al., 2022] have been applied to navigate this search space more effectively. Because these methods optimize a scalar reward without an explicit distributional constraint, they are prone to generating sequences that exploit biases in the prediction model and diverge from the regulatory grammar of natural DNA [Vaishnav et al., 2022].

Generative Models for Regulatory DNA. Generative approaches learn the underlying distribution of natural DNA sequences and use this learned prior to constrain the design space. Recent methods employ autoregressive genomic language models [Nguyen et al., 2023, Schiff et al., 2024, Lal et al., 2024] and discrete diffusion models [Avdeyev et al., 2023, Sarkar et al., 2025, DaSilva et al., 2026] to capture long-range sequence dependencies with high fidelity. To steer such generative priors toward cell-type-specific regulatory activity, three guidance paradigms have emerged. *Classifier-free guidance* (CFG) builds class-conditioned diffusion models that can be steered toward desired sequence classes at sampling time [Schiff et al., 2025]. While CFG preserves sequence quality, conditioning on broad regulatory classes alone does not provide sufficient specificity for fine-grained cell-type discrimination. *RL-based fine-tuning* updates the generative model weights to maximize a reward function. Ctrl-DNA combines genomic language models with constrained RL for cell-type-specific design [Chen et al., 2025b], and DRAKES backpropagates reward gradients through a discrete diffusion model via the Gumbel-Softmax trick [Wang et al., 2025]. While effective, fine-tuning is computationally expensive and must be repeated for every new design objective. *Inference-time alignment* methods steer the sampling process of a frozen generative model using external guidance signals, avoiding the cost of retraining [Uehara et al., 2025]. Techniques such as soft value-based decoding [Li et al., 2024], sequential Monte Carlo (SMC) [Phillips et al., 2024], twisted diffusion sampling (TDS) [Wu et al., 2024], and classifier-based guidance [Nisonoff et al., 2025] use reward-weighted resampling or auxiliary gradients to steer the generative process. However, these methods typically optimize for high activity in a single cell type without explicitly suppressing off-target activity, limiting their applicability to strict cell-type specificity tasks.

Search-Based Guidance for Diffusion Models. An alternative to gradient-based or resampling-based inference-time alignment is to formulate the reverse diffusion process as a sequential decision problem and apply tree search algorithms to guide generation. Monte Carlo tree search (MCTS) has been adapted to continuous diffusion models as a planner [Yoon et al., 2026, 2025], and extended to discrete diffusion for multi-objective peptide design [Tang et al., 2024]. Multi-objective guided discrete flow matching has also been explored in related biological design tasks [Chen et al., 2025a]. These

methods use tree search to balance reward maximization with distributional fidelity by evaluating candidate completions of partially denoised states.

A.2 Complete Algorithm

We provide the complete DNA-CRAFT algorithm in Algorithm 1. The procedure implements class-conditioned Monte Carlo tree diffusion with fast lookahead pruning, as described in Section 3. Candidates are ranked using the smooth MinGap score g (Equation 2), while the archive-relative rank rewards ρ_m (Equation 9) are used for tree updates. We use $N(\boldsymbol{z})$ to denote the visit count and $\bar{Q}(\boldsymbol{z})$ to denote the accumulated backpropagated rank reward divided by the visit count of node $\boldsymbol{z}$. Each expansion adds the sum of the retained children's rank rewards to the accumulated reward of each ancestor and increments its visit count once. The soft-optimal policy in Equation 5 motivates the search, but DNA-CRAFT uses finite candidate sampling, lookahead pruning, and UCT selection rather than sampling exactly from this policy. Consequently, λ is not an explicit parameter of Algorithm 1.

Algorithm 1 DNA-CRAFT: Class-Conditioned Monte Carlo Tree Diffusion

Input: Denoiser, Design model $\boldsymbol{s}$, regulatory class $\boldsymbol{y}$, guidance scale γ
  Proposal count B, child limit M, coefficient c, iteration limit N_{iter}
  Archive capacity N_{max}
Initialize a fully masked root with $N = 1$ and $\bar{Q} = 0$
Initialize archive $\mathcal{G}^* \leftarrow \emptyset$
for $k = 1, \ldots, N_{\text{iter}}$ **do**
  node $\leftarrow$ root
  while node has children and is not terminal **do**
    node $\leftarrow \underset{\boldsymbol{z} \in \text{children(node)}}{\arg\max} \left[\bar{Q}(\boldsymbol{z}) + c\sqrt{\dfrac{\log N(\text{node})}{N(\boldsymbol{z})}} \right]$
  end while
  if node is terminal **then**
    Skip the remainder of this iteration
  else
    Let $\boldsymbol{z}_t \leftarrow$ node and let $s < t$ be the next diffusion time
    for $b = 1, \ldots, B$ **do**
      Sample $\boldsymbol{z}_{s,b}$ using Equation 7 and retain revealed bases
      Obtain $\hat{\boldsymbol{x}}_{0,b}$ using Equation 8
      $g_b \leftarrow g(\hat{\boldsymbol{x}}_{0,b})$ using Equation 2
    end for
    $\mathcal{I}_M \leftarrow$ indices of the $\min(M, B)$ highest-scoring candidates, in decreasing order
    for $m \in \mathcal{I}_M$ **do**
      Compute ρ_m against the current archive using Equation 9
      if $\mathcal{G}^* = \emptyset$ **then**
        Insert $\hat{\boldsymbol{x}}_{0,m}$ into $\mathcal{G}^*$
      else if $g_m \geq \max_{\boldsymbol{x}^* \in \mathcal{G}^*} g(\boldsymbol{x}^*)$ **then**
        Insert or update $\hat{\boldsymbol{x}}_{0,m}$ in $\mathcal{G}^*$
        Enforce capacity N_{max} using the archive rule described in Section 3
      end if
      Add $\boldsymbol{z}_{s,m}$ as a child of $\boldsymbol{z}_t$, with $N(\boldsymbol{z}_{s,m}) = 1$ and $\bar{Q}(\boldsymbol{z}_{s,m}) = \rho_m$
    end for
    for each node $\boldsymbol{z}$ from $\boldsymbol{z}_t$ to the root **do**
      $\bar{Q}(\boldsymbol{z}) \leftarrow \dfrac{N(\boldsymbol{z})\bar{Q}(\boldsymbol{z}) + \sum_{m \in \mathcal{I}_M} \rho_m}{N(\boldsymbol{z}) + 1}$
      $N(\boldsymbol{z}) \leftarrow N(\boldsymbol{z}) + 1$
    end for
  end if
end for
Output: Archive $\mathcal{G}^*$ of admitted record-scoring completions

A.3 DNA-CRAFT Diffusion Model Training Details

Forward noising process. The forward process progressively corrupts a naturally occurring DNA sequence $\boldsymbol{x}$ into a sequence of discrete variables $\boldsymbol{z}_t$ over continuous time $t \in [0, 1]$, where unmasked tokens (base pairs) transition to the masked token $\boldsymbol{m}$. The marginal probability of a latent state $\boldsymbol{z}_t$ given the input DNA sequence $\boldsymbol{x}$ is defined as a categorical distribution $q(\boldsymbol{z}_t|\boldsymbol{x}) = \mathrm{Cat}(\boldsymbol{z}_t; \alpha_t \boldsymbol{x} + (1 - \alpha_t)\boldsymbol{m})$, where $\alpha_t \in [0, 1]$ is a monotonically decreasing noise schedule. Concretely, this means that at time t, each token independently remains as the original nucleotide with probability α_t or is replaced with probability $1 - \alpha_t$ by the absorbing state mask token $\boldsymbol{m}$.

Reverse denoising process. The reverse process reconstructs the sequence by progressively unmasking positions from a fully masked sequence. Conditioned on the original DNA sequence $\boldsymbol{x}$, the time-reversal process for time steps $s < t$ is

$$q(\boldsymbol{z}_s|\boldsymbol{z}_t, \boldsymbol{x}) = \begin{cases} \mathrm{Cat}(\boldsymbol{z}_s; \boldsymbol{z}_t), & \boldsymbol{z}_t \neq \boldsymbol{m}, \\ \mathrm{Cat}\Big(\boldsymbol{z}_s; \dfrac{(1-\alpha_s)\boldsymbol{m} + (\alpha_s - \alpha_t)\boldsymbol{x}}{1 - \alpha_t}\Big), & \boldsymbol{z}_t = \boldsymbol{m}. \end{cases} \tag{10}$$

We train a class-conditioned backbone neural network denoted as $\boldsymbol{x}_\theta(\boldsymbol{z}_t, t, \boldsymbol{y})$ to predict categorical probabilities over the clean nucleotides given the noisy input $\boldsymbol{z}_t$ and regulatory class $\boldsymbol{y}$. The learned reverse transition is obtained by substituting these predictions into Equation 10 at each position, with a loss function that simplifies to a weighted average of masked language modeling losses. To generate a DNA sequence of length L, we ancestrally sample from the learned reverse process, starting with a completely masked sequence $\boldsymbol{z}_t^{1:L}$ at $t = 1$. Conditional on the current sequence and regulatory class, the reverse transition factorizes across positions:

$$p_\theta(\boldsymbol{z}_s^{1:L} \mid \boldsymbol{z}_t^{1:L}, \boldsymbol{y}) = \prod_{\ell=1}^{L} p_\theta(\boldsymbol{z}_s^{\ell} \mid \boldsymbol{z}_t^{1:L}, \boldsymbol{y}). \tag{11}$$

Dataset Curation. We trained DNA-CRAFT using the ENCODE Registry of candidate cis-Regulatory Elements V4 [Moore et al., 2026]. To capture cross-species regulatory grammar and maximize biological diversity, we integrated data from both human (hg38) and mouse (mm10) genomes. The final dataset comprises 2,348,854 annotated regions across 1,888 human cell types and 926,843 annotated regions across 366 mouse cell types.

Data Pre-Processing. DNA sequences corresponding to the regulatory regions were extracted from their respective reference genomes. All sequences were centered and padded to a fixed length of 350 base pairs (bp) using a distinct `[PAD]` token. To accommodate variable effective lengths, we implemented a masking mechanism within both the model backbone and the diffusion loss function, ensuring that padding tokens are excluded from the generative process. We applied reverse complement augmentation by taking either the forward or reverse complement strand with equal probability during training.

Class Consolidation. The ENCODE V4 Registry categorizes regulatory elements based on biochemical signatures derived from DNase hypersensitivity, histone modifications, and CTCF binding. The original classification schema distinguishes between:

- **Promoter-like (PLS):** Accessible regions proximal to transcription start sites (TSS) enriched for H3K4me3.
- **Enhancer-like (ELS):** Accessible regions enriched for H3K27ac, further stratified into proximal (pELS) and distal (dELS) based on TSS proximity.
- **CA-H3K4me3:** Accessible regions with H3K4me3 enrichment but low H3K27ac, often indicative of poised or primed regulatory states.
- **CA-CTCF:** Accessible regions enriched for CTCF binding sites with low histone acetylation, often indicating insulators.
- **CA-TF, CA, and TF:** Elements defined primarily by chromatin accessibility or transcription factor binding, lacking canonical histone marks.

We consolidated these fine-grained categories into five broad functional classes for DNA-CRAFT conditioning (Table 4).

Table 4: Mapping of ENCODE V4 cCRE classifications to DNA-CRAFT conditioning labels.

DNA-CRAFT Class	Original ENCODE Class
Promoter	PLS
Enhancer	dELS, pELS
CTCF	CA-CTCF
Poised	CA-H3K4me3, TF
Open Chromatin	CA, CA-TF

Cross-Species Train-Test-Validation Split. To ensure generalization and prevent overfitting due to sequence homology, we constructed our data splits using a graph-based clustering approach adapted from the Enformer training protocol [Avsec et al., 2021].

We constructed an undirected graph $G = (V, E)$ in which the vertices V represent approximately 3.2 million processed regulatory elements. Edges E were defined to capture both orthology and sequence similarity:

1. **Homology Edges:** An edge connects a human cCRE and a mouse cCRE if they share a sequence alignment of > 100 bp, determined via the hg38-mm10 syntenic nets [Kent et al., 2002].
2. **Overlap Edges:** Within a single species, an edge connects any two cCREs that share sequence similarity of more than 100 bp.

We computed the connected components of G to define independent sequence clusters. These clusters were randomly partitioned into training (90%), validation (5%), and test (5%) sets, ensuring a strict separation of homologous sequences across splits.

Hyperparameters and Training. We trained the class-conditioned diffusion model on the processed ENCODE dataset for 100 epochs using the AdamW optimizer and a global batch size of 4,096. During training, the regulatory class label was replaced with the null conditioning label with probability 0.1 to support classifier-free guidance. Model hyperparameters are provided in Table 5.

Table 5: DNA-CRAFT Hyperparameters. The model utilizes a bidirectional DiMamba backbone.

Hyperparameter	Value
Backbone Architecture (DiMamba)	
Model Parameters	1.93 Million
Sequence Length (L)	350
Hidden Dimension (d_{model})	128
Mamba Blocks	10
Bidirectional Strategy	Addition (Tied weights)
Dropout	0.1
Diffusion Process	
Noise Schedule	Cosine
Time Conditioning	True
Optimization	
Optimizer	AdamW
Learning Rate	1×10^{-3}
Training Epochs	100
Batch Size	4,096
Conditioning (CFG)	
Condition Type	Regulatory Element Class
Number of Classes	5
Condition Dropout	0.1
Conditioning Dim	128

Computational Infrastructure. All experiments were conducted on NVIDIA H100 GPUs (80GB VRAM). The validation loss and perplexity trajectories of the class-conditioned diffusion model are shown in Figure 3.

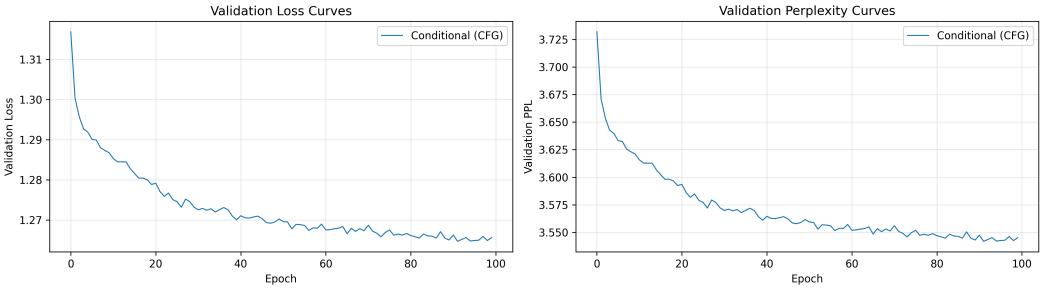

Figure 3: Validation loss and perplexity of the class-conditioned DNA-CRAFT diffusion model over 100 training epochs.

A.4 Class-Conditioned Sampling and Motif Analysis

Classifier-Free Guidance Renormalization. Equation 6 applies guidance to the denoiser's clean-token probabilities, with normalization over nucleotides at each masked position. We then construct the reverse transition using the noise schedule α_t defined in Appendix A.3. For a masked position ℓ and $s < t$, the probability of token a is

$$p_\theta^\gamma(\boldsymbol{z}_s^\ell = a \mid \boldsymbol{z}_t, \boldsymbol{y}) = \begin{cases} \dfrac{\alpha_s - \alpha_t}{1 - \alpha_t} [\boldsymbol{x}_\theta^\gamma(\boldsymbol{z}_t, t, \boldsymbol{y})]_{\ell,a}, & a \in \mathcal{V}, \\ \dfrac{1 - \alpha_s}{1 - \alpha_t}, & a = \boldsymbol{m}. \end{cases} \tag{12}$$

Thus, guidance changes the nucleotide probabilities while the probability of remaining masked follows the diffusion schedule. At already unmasked positions, the transition preserves the current nucleotide. Conditional on the current sequence, sampling factorizes across positions and uses independent Gumbel draws as specified in Equation 7.

Detailed Motif Enrichment Analysis. To validate the biological relevance of the class-conditioned generated sequences, we performed a motif enrichment analysis. Figure 4 displays a heatmap of the Z-scores for the top 5 most specific motifs per class. The recovery of these specific factors is consistent with known biological expectations. Table 6 lists the top 5 enriched TFs for each class.

Table 6: Top 5 Enriched Transcription Factors per Regulatory Class. Factors are ranked by enrichment Z-score.

Regulatory Class	Top Enriched Motifs
CTCF	CTCF, ZNF768, NFATC3, Zfp335, SNAI3
Enhancer	TFAP2C, MAFK, HES6, FOXJ2::ELF1, SPIC
Open Chromatin	Nfat5, Foxo3, FOXO1::ELK3, FOSL2::JUND, FOXK1
Poised	LMX1A, TBX5, Prdm14, MEF2B, TBX2
Promoter	NKX6-2, FEV, NFYA, FLI1::DRGX, ELK1::SREBF2
Unconditional	MAFG::NFE2L1, FOS::JUN, Lhx3, FOSL1::JUNB, MAF::NFE2

A.5 Regulatory Sequence Design Benchmarks

Benchmarking Model Implementation Details. Reported benchmark values summarize three independent generation runs, each producing 128 sequences. Each run contributes one aggregate value per metric. Standard deviations describe variation across generation runs, and bold values indicate the highest or lowest reported means according to the metric direction.

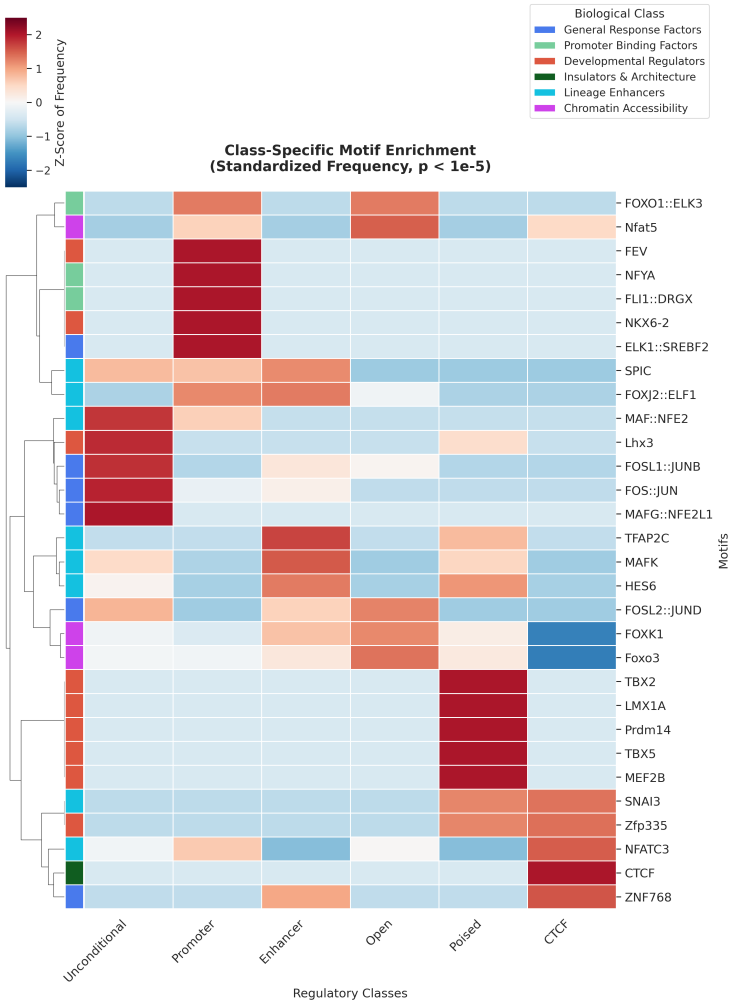

Figure 4: Heatmap of motif-frequency Z-scores across the five regulatory classes and null-conditioned sampling. "Unconditional" denotes sampling from the same class-conditioned network with the null label at $\gamma = 0$.

1. **Ledidi:** We initialized the optimization with 128 random DNA sequences. We implemented a custom differentiable wrapper to compute MinGap scores from Design-model's output and backpropagate gradients directly to the input sequence representation. Optimization was conducted independently for each sequence for a maximum of 20,000 steps to ensure convergence.
2. **Sequential Monte Carlo (SMC):** We used the same class-conditioned diffusion model as DNA-CRAFT, conditioned on the Enhancer class with CFG scale $\gamma = 1.0$. The sampling process tracked 128 particles and applied a resampling parameter of $\alpha = 0.5$.
3. **Twisted Diffusion Sampling (TDS):** We used the same enhancer-conditioned diffusion prior and particle configuration as SMC. TDS used a reward-guidance scale of 1000 and a resampling parameter of $\alpha = 0.5$.
4. **D3 (Discrete Denoising Diffusion):** We trained the discrete diffusion model utilizing the transformer backbone (2M parameters) for 100 epochs on the MPRA dataset ($\sim$700,000 sequences) with the cell type activity as classes. We generated 128 sequences per cell type with conditional sampling ($\gamma = 4.0$).
5. **Ctrl-DNA:** We followed the original protocol and hyperparameters specified for each target cell type. We fine-tuned three separate models (one for each cell line: HepG2, K562, SK-N-SH) for 100 optimization steps. The top 128 sequences from each fine-tuning run were selected for evaluation. We note that Ctrl-DNA's base autoregressive model [Nguyen

et al., 2023] is trained on the entire human reference genome, which is a substantially larger and more diverse training corpus than the ENCODE registry of 3.2 million cCREs used to train DNA-CRAFT's diffusion backbone. Additionally, the TFBS regularization term was excluded, as the specific motif lists were not available in the public repository.

6. **DRAKES:** We followed the protocol described in the original publication. For the HepG2 cell line, we utilized the provided pre-trained checkpoint. Similarly, we fine-tuned two separate models for K562 and SK-N-SH. We generated sequences using the respective fine-tuned models. We note that DRAKES fine-tunes its diffusion model using reward gradients computed over the full MPRA dataset ($\sim$700,000 sequences), which encompasses the data used to train both our Design-model and Evaluation-model. This gives DRAKES implicit access to the evaluation distribution during its fine-tuning process.
7. **DNA-CRAFT (Ours):** We use the class-conditioned diffusion model trained on ENCODE without further fine-tuning. For each replicate of the human cell-line benchmark, we initialize 128 separate search trees, each with a fresh archive and a fully masked root. Each search scores candidate completions using the Design model and maintains an archive according to the smooth specificity score g. We retain the highest-scoring completion from each search to obtain the final set of 128 sequences per replicate. Table 7 summarizes the inference settings.

Table 7: DNA-CRAFT inference parameters for the human cell-line benchmark. Each experimental replicate comprises 128 separately initialized searches.

Parameter	**Value**
Diffusion Sampling Steps	32
Guidance Scale (γ)	1.0
Conditioning Class	Enhancer
Search Iterations (N_{iter})	32
Retained-Child Limit (M)	8
Lookahead Pool Size (B)	800
UCT Exploration Coefficient (c)	0.25
Archive Capacity (N_{max})	64
Separate Searches per Replicate	128

Extended Benchmark Results

We extended our evaluation to include the following additional evaluation metrics:

- **Complete Dataset Validation:** We additionally scored the generated sequences using an Enformer model trained on the full MPRA dataset [Gosai et al., 2024], providing a supplementary comparison with the partition-trained Evaluation model.
- **Fréchet Biological Distance (FBD):** Following Stark et al. [2024], we use the Fréchet Biological Distance (FBD) to measure the resemblance of generated sequences to the natural distribution. This metric extends the Fréchet Inception Distance (FID) commonly used in image generation [Heusel et al., 2018] to biological sequences by computing distances in a learned embedding space. Concretely, let $\{\boldsymbol{h}_i^{\text{ref}}\}_{i=1}^N$ and $\{\boldsymbol{h}_j^{\text{gen}}\}_{j=1}^M$ denote embeddings of reference and generated sequences, respectively, obtained from the sequence-to-activity model. We fit Gaussian distributions $\mathcal{N}(\boldsymbol{\mu}_{\text{ref}}, \boldsymbol{\Sigma}_{\text{ref}})$ and $\mathcal{N}(\boldsymbol{\mu}_{\text{gen}}, \boldsymbol{\Sigma}_{\text{gen}})$ to these embeddings, and compute the Fréchet distance:

$$\text{FBD} = \|\boldsymbol{\mu}_{\text{ref}} - \boldsymbol{\mu}_{\text{gen}}\|_2^2 + \text{Tr}\Big(\boldsymbol{\Sigma}_{\text{ref}} + \boldsymbol{\Sigma}_{\text{gen}} - 2\left(\boldsymbol{\Sigma}_{\text{ref}}\boldsymbol{\Sigma}_{\text{gen}}\right)^{1/2}\Big). \tag{13}$$

  Lower FBD indicates closer alignment between the generated and reference sequence distributions.
- **Pairwise Levenshtein Distance:** We calculated sequence diversity using pairwise Levenshtein (edit) distance.
- **Cross-Architecture Validation (MinGap Borzoi):** We evaluate the predicted activity of the generated sequences with a fine-tuned Borzoi model [Linder et al., 2025] trained on the

evaluation data split. This tests whether the cell-type specificity transfers across independent predictor architectures.

- **GC Content:** We report the mean GC fraction of each method's generated sequences. The reference value (shown in parentheses) is the mean GC content of natural enhancers for the target cell type. Methods that produce sequences with GC content far from the reference value may be exploiting compositional biases rather than learning genuine regulatory syntax.
- **Motif-Hit Preprocessing:** For the arrangement metrics, we restrict the analysis to the 50 most frequent motifs in the natural reference set and retain FIMO [Grant et al., 2011] hits with $p < 10^{-4}$. Within each sequence, hits for the same motif are selected greedily by score, requiring at least 10 bp between their centers. We then retain at most the 10 highest-scoring hits per sequence in both the generated and reference sets. The same reference-defined motif vocabulary is used for all methods.
- **Pairwise TFBS Co-occurrence:** For each unordered pair of distinct motifs, we compute the association $\log_2[(P(a,b))/(P(a)P(b))]$, where $P(a,b)$ is the fraction of sequences containing both motifs, and $P(a)$ and $P(b)$ are their marginal presence frequencies. We report the Spearman correlation between the generated and reference association vectors across all motif pairs. Higher correlation indicates closer agreement with the motif co-occurrence patterns of natural sequences.
- **Inter-Motif Spacing (JSD):** We collect center-to-center distances between all pairs of distinct-motif hits within each sequence and pool these distances across the sequence set. Distances are grouped into 10 bp bins spanning the sequence length, giving 20 bins for 200 bp sequences. We report the base-2 Jensen–Shannon divergence (JSD) between the normalized generated and reference histograms. Lower values indicate closer agreement with the reference spacing distribution.
- **Relative Motif Orientation:** For each unordered pair of distinct motifs, we calculate the fraction of hit pairs within 50 bp that occur on the same strand. We retain motif pairs with at least five observations in both the generated and reference sets and report the Spearman correlation between their same-strand fractions. The metric is reported only when at least five motif pairs satisfy this criterion and both vectors have nonzero variance. Higher correlation indicates closer agreement with the reference orientation preferences.
- **Fréchet Biological Distance (NTv2):** We compute a second distributional distance using embeddings from the Nucleotide Transformer v2 (500M parameters) [Dalla-Torre et al., 2025]. This provides a measure of how close the generated distribution is to the natural enhancer distribution. Lower values are better.

Table 8 presents the comprehensive performance across all metrics. Figure 5 visualizes the trade-off between optimization intensity (MinGap Eval. score) and naturalness. The figure illustrates the trade-offs between predicted specificity and the two fidelity metrics across methods and cell lines.

A.5.1 Brain Cell-Type-Specific Sequence Design

Experimental Setup. We further evaluated regulatory sequence design across eight brain cell types from the CATlas dataset: Astro, GABA1, GABA2, Glut1, Glut2, Microglia, OPC, and Oligo. We followed the split-model protocol described in Section 4.2, using separate Design and Evaluation models trained on non-overlapping homology clusters. Each experiment targeted one cell type while penalizing predicted accessibility in the remaining seven. We compared SMC, TDS, Ctrl-DNA, Ledidi, and DNA-CRAFT, generating 128 sequences per run over three independent random seeds. DNA-CRAFT, SMC, and TDS used the same enhancer-conditioned diffusion prior with class-conditioning scale $\gamma = 1.0$, and DNA-CRAFT used a lookahead pool of $B = 800$. Evaluation followed the immune benchmark, reporting MinGap under the Evaluation model, motif correlation, 3-mer correlation, and Shannon diversity. Natural reference regions were selected using the same procedure as for the immune cell types.

Results. DNA-CRAFT achieves the highest reported mean MinGap, motif correlation, and 3-mer correlation across all eight brain cell types (Table 9). Its MinGap ranges from 0.408 for Glut1 to 0.949 for Microglia, while maintaining strong agreement with the reference motif and sequence-composition distributions. Ledidi achieves the highest positional Shannon diversity across all eight

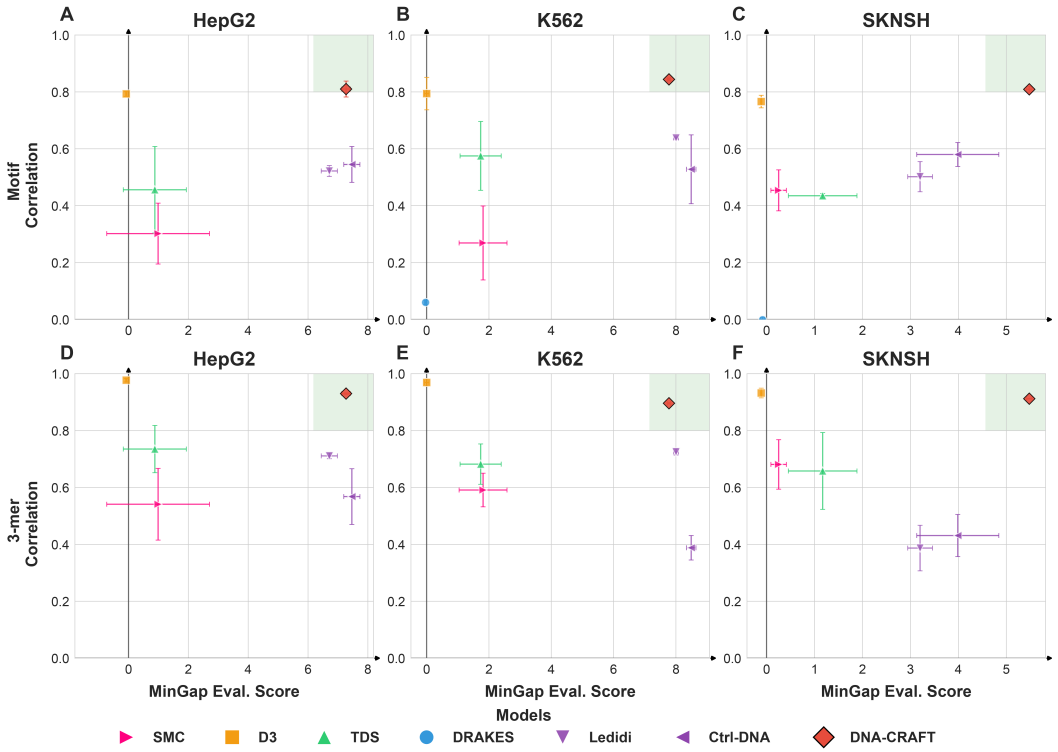

Figure 5: Trade-off between cell-type specificity and biological fidelity. Performance of DNA-CRAFT compared to baselines for HepG2, K562, and SK-N-SH cell lines. The x-axis represents the MinGap score, serving as a proxy for cell-type-specific activity. The y-axis represents biological fidelity, measured by JASPAR Motif Correlation and 3-mer Correlation relative to top specific natural enhancers.

cell types, but substantially lower specificity and motif correlation. These results extend the observed specificity-naturalness trade-off to neuronal and glial cell types within the CATlas benchmark.

A.5.2 Sequence-To-Activity Model Training Details.

We adapted the pre-trained Enformer [Avsec et al., 2021] and Borzoi [Linder et al., 2025] architectures using gReLU [Lal et al., 2025]. For each task, we replaced the output head and fine-tuned the full network. Training used the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 1024 for 15 epochs, without reverse-complement or shift augmentation. For the MPRA benchmark, the models predict activity across HepG2, K562, and SK-N-SH using a three-output regression head optimized with mean squared error. For chromatin accessibility prediction, the models use a 204-output head trained with binary cross-entropy loss.

Input Processing and Prediction Scale. Enformer receives the 200 bp sequence directly. For Borzoi, we pad each sequence with 28 ambiguous bases on either side to obtain a 256 bp input, using the same processing during training and evaluation. These inputs are shorter than the contexts used by the original pre-trained architectures. MPRA predictions are used directly on the trained regression scale. For chromatin accessibility, we apply a sigmoid to the output logits before calculating specificity, so the immune MinGap scores represent differences in predicted accessibility probabilities.

Validation Performance. For the partitioned Design and Evaluation models, training and validation use opposite data partitions. The Design model is trained on the design partition and validated on the evaluation partition, while the Evaluation model uses the reverse assignment. We select the checkpoint with the best validation performance for final scoring. Consequently, the reported scores summarize checkpoint-selection validation performance rather than performance on an additional untouched test set. Table 10 reports the MPRA validation metrics for Enformer and Borzoi, alongside

Table 8: Extended human cell-line benchmarks. We report cell-type specificity under three oracle configurations spanning two architectures (Enformer evaluation-split, Enformer full-MPRA, and Borzoi), sequence-composition fidelity (motif correlation, 3-mer correlation, GC content), motif arrangement quality (pairwise co-occurrence, inter-motif spacing JSD, relative orientation correlation), distributional distance to natural enhancers (FBD with Enformer embeddings and FBD with Nucleotide Transformer v2 embeddings), and output diversity (pairwise Levenshtein distance and positional Shannon entropy). Arrows indicate the direction of improvement; GC content lists the natural-enhancer reference value. Best values are **bolded**. All values report mean ± std over three seeds.

Cell Type	Metric	SMC	TDS	D3	DRAKES	Ctrl-DNA	Ledidi	DNA-CRAFT (Ours)
HepG2	MinGap Eval ↑	0.989 ±1.717	0.880 ±1.053	-0.075 ±0.021	-1.314 ±0.024	**7.464** ±0.270	6.712 ±0.270	7.271 ±0.142
	MinGap (Full MPRA) ↑	0.875 ±1.575	0.896 ±0.982	-0.056 ±0.014	-1.210 ±0.035	**7.210** ±0.714	6.487 ±0.272	7.127 ±0.168
	MinGap (Borzoi) ↑	0.690 ±1.161	0.601 ±0.791	-0.060 ±0.011	-1.100 ±0.040	**6.961** ±0.614	4.946 ±0.159	5.515 ±0.184
	Motif Corr. ↑	0.302 ±0.107	0.456 ±0.152	0.793 ±0.009	-0.126 ±0.014	0.545 ±0.063	0.522 ±0.019	**0.810** ±0.028
	3-mer Corr. ↑	0.541 ±0.126	0.735 ±0.083	**0.977** ±0.006	-0.361 ±0.012	0.568 ±0.098	0.711 ±0.010	0.930 ±0.003
	GC Content (ref. 0.457)	0.485 ±0.060	0.467 ±0.021	0.458 ±0.009	0.698 ±0.002	0.458 ±0.032	0.477 ±0.002	0.487 ±0.001
	Motif Co-occur. ↑	0.090 ±0.044	0.145 ±0.039	0.271 ±0.016	0.225 ±0.046	0.208 ±0.086	0.304 ±0.024	**0.326** ±0.035
	Spacing JSD ↓	0.123 ±0.025	0.161 ±0.058	0.022 ±0.006	0.037 ±0.012	0.022 ±0.002	**0.009** ±0.001	0.011 ±0.001
	Orientation Corr. ↑	0.269 ±0.023	0.319 ±0.469	0.452 ±0.147	0.686 ±0.198	0.455 ±0.115	**0.765** ±0.003	0.695 ±0.026
	FBD ↓	941.851 ±356.067	**904.089** ±172.563	1066.242 ±14.740	2816.206 ±101.942	2968.966 ±776.318	1535.551 ±70.126	1427.015 ±127.205
	FBD (NTv2) ↓	5.558 ±0.762	4.210 ±0.450	**1.629** ±0.013	13.231 ±1.080	3.745 ±0.352	2.397 ±0.034	2.018 ±0.090
	Diversity (Levenshtein) ↑	71.916 ±41.709	104.328 ±8.925	149.687 ±0.217	141.970 ±0.118	145.024 ±2.293	149.776 ±0.145	**149.898** ±0.041
	Shannon Div. ↑	0.845 ±0.478	1.301 ±0.105	1.978 ±0.003	1.864 ±0.002	1.914 ±0.033	1.980 ±0.002	**1.981** ±0.001
K562	MinGap Eval ↑	1.812 ±0.768	1.735 ±0.660	0.003 ±0.070	-0.033 ±0.032	**8.490** ±0.153	8.004 ±0.078	7.777 ±0.054
	MinGap (Full MPRA) ↑	1.963 ±0.948	1.752 ±0.714	-0.014 ±0.059	-0.275 ±0.071	**8.788** ±0.166	8.324 ±0.099	8.046 ±0.073
	MinGap (Borzoi) ↑	1.123 ±0.845	1.426 ±0.585	-0.042 ±0.049	-0.027 ±0.047	**7.656** ±0.050	6.935 ±0.082	6.747 ±0.095
	Motif Corr. ↑	0.269 ±0.130	0.575 ±0.121	0.794 ±0.057	0.060 ±0.008	0.528 ±0.121	0.638 ±0.004	**0.844** ±0.005
	3-mer Corr. ↑	0.591 ±0.059	0.682 ±0.071	**0.969** ±0.013	-0.354 ±0.007	0.388 ±0.043	0.725 ±0.011	0.896 ±0.003
	GC Content (ref. 0.477)	0.457 ±0.048	0.499 ±0.067	0.468 ±0.010	0.567 ±0.003	0.480 ±0.035	0.488 ±0.000	0.497 ±0.002
	Motif Co-occur. ↑	0.105 ±0.065	0.090 ±0.018	**0.409** ±0.024	0.341 ±0.027	0.069 ±0.041	0.232 ±0.046	0.383 ±0.059
	Spacing JSD ↓	0.219 ±0.044	0.184 ±0.078	0.021 ±0.003	0.025 ±0.008	0.053 ±0.016	0.013 ±0.003	**0.012** ±0.001
	Orientation Corr. ↑	0.674 ±0.183	0.866 ±0.038	0.713 ±0.171	0.664 ±0.058	0.759 ±0.074	0.764 ±0.080	**0.869** ±0.033
	FBD ↓	**742.185** ±94.080	799.443 ±176.784	1145.004 ±31.744	1962.765 ±69.462	2890.187 ±364.824	1853.457 ±33.499	1619.139 ±18.759
	FBD (NTv2) ↓	6.072 ±0.943	4.724 ±0.313	**1.667** ±0.023	8.666 ±0.638	4.911 ±0.356	2.264 ±0.062	2.031 ±0.062
	Diversity (Levenshtein) ↑	39.905 ±27.443	92.245 ±10.923	149.809 ±0.171	148.352 ±0.213	141.204 ±0.978	**149.896** ±0.084	149.789 ±0.051
	Shannon Div. ↑	0.487 ±0.302	1.063 ±0.170	1.980 ±0.002	1.958 ±0.003	1.851 ±0.019	**1.982** ±0.001	1.980 ±0.001
SK-N-SH	MinGap Eval ↑	0.252 ±0.161	1.170 ±0.715	-0.111 ±0.014	-0.082 ±0.040	3.988 ±0.856	3.202 ±0.258	**5.479** ±0.057
	MinGap (Full MPRA) ↑	0.222 ±0.133	1.125 ±0.779	-0.116 ±0.011	0.255 ±0.051	4.598 ±0.843	3.680 ±0.370	**6.894** ±0.102
	MinGap (Borzoi) ↑	0.070 ±0.191	0.723 ±0.580	-0.158 ±0.040	-0.036 ±0.040	3.650 ±0.886	2.431 ±0.203	**4.565** ±0.071
	Motif Corr. ↑	0.454 ±0.072	0.435 ±0.008	0.766 ±0.022	-0.001 ±0.009	0.580 ±0.042	0.502 ±0.053	**0.809** ±0.010
	3-mer Corr. ↑	0.681 ±0.087	0.658 ±0.135	**0.932** ±0.017	-0.382 ±0.001	0.431 ±0.074	0.387 ±0.080	0.912 ±0.001
	GC Content (ref. 0.462)	0.420 ±0.059	0.475 ±0.080	0.433 ±0.005	0.720 ±0.001	0.477 ±0.063	0.487 ±0.005	0.440 ±0.001
	Motif Co-occur. ↑	0.039 ±0.176	0.065 ±0.121	0.432 ±0.038	0.140 ±0.122	0.428 ±0.080	0.481 ±0.016	**0.534** ±0.036
	Spacing JSD ↓	0.170 ±0.092	0.222 ±0.096	0.033 ±0.009	0.066 ±0.031	0.042 ±0.050	0.021 ±0.011	**0.008** ±0.002
	Orientation Corr. ↑	0.279 ±0.720	0.336 ±0.321	0.711 ±0.098	–	0.871 ±0.054	0.891 ±0.038	**0.897** ±0.016
	FBD ↓	1047.505 ±23.097	841.456 ±85.692	1016.487 ±15.447	2291.623 ±35.281	1834.831 ±1284.741	**777.882** ±75.112	1742.182 ±44.163
	FBD (NTv2) ↓	5.893 ±1.737	5.538 ±0.686	**1.785** ±0.016	18.779 ±0.856	4.866 ±0.750	2.687 ±0.104	2.357 ±0.013
	Diversity (Levenshtein) ↑	81.457 ±8.492	82.683 ±28.917	149.011 ±0.246	139.740 ±0.052	140.473 ±4.902	**149.944** ±0.067	149.218 ±0.069
	Shannon Div. ↑	1.029 ±0.091	1.019 ±0.327	1.969 ±0.004	1.826 ±0.001	1.847 ±0.066	**1.982** ±0.001	1.972 ±0.001

the original Enformer Full-Model results. Table 11 reports accessibility-prediction AUROC for the eight immune cell types used in the design benchmark and the mean across all 204 tracks.

A.6 Extended Ablations

We provide the complete numerical results for all ablation sweeps discussed in Section 4.3. All experiments were conducted on the SK-N-SH enhancer design task, generating 128 sequences per configuration across 3 independent runs.

CFG Guidance Scale Sweep.

Table 12 reports the full metrics for the class-conditioned expansion sweep across guidance scales $\gamma \in \{0.0, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$.

Fast Lookahead Pruning Sweep.

Table 13 reports the full metrics for the lookahead pool size sweep across $B \in \{2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048\}$. The retained-child limit is $M = 8$, giving $\min(B, 8)$ retained children per expansion. We compare these configurations with standard full rollouts without lookahead (Std.).

Cross-Method Evaluation under Alternative Specificity Metrics.

To assess whether the specificity comparisons depend on the choice of evaluation metric, we evaluated the generated sequences from all seven methods using three alternative scores under the Enformer Evaluation model. MinGap subtracts the maximum off-target activity from the target activity, LogSumExp replaces this maximum with a smooth aggregation, and Mean Difference subtracts the mean off-target activity. The generated sequences are held fixed across these evaluations. This

Table 9: Comparison of methods for cell-type-specific regulatory sequence design across eight brain cell types. Shown are MinGap specificity scores from the Evaluation model, motif Spearman correlations, 3-mer Pearson correlations, and Shannon diversity. Values are reported as mean ± standard deviation over three independent runs, with 128 sequences per run. The highest reported mean in each row is marked in **bold**. Ctrl-DNA was evaluated without TFBS regularization.

Target Cell Type	Metric	SMC	TDS	Ctrl-DNA	Ledidi	DNA-CRAFT (Ours)
Astro	MinGap ↑	-0.094 ± 0.079	-0.045 ± 0.033	0.311 ± 0.164	-0.180 ± 0.004	$\mathbf{0.785} \pm 0.003$
	Motif Corr. ↑	0.389 ± 0.086	0.418 ± 0.153	0.372 ± 0.209	-0.057 ± 0.119	$\mathbf{0.750} \pm 0.022$
	3-mer Corr. ↑	0.699 ± 0.110	0.747 ± 0.078	0.498 ± 0.117	0.029 ± 0.161	$\mathbf{0.914} \pm 0.002$
	Shannon Div. ↑	1.131 ± 0.130	1.263 ± 0.172	1.811 ± 0.137	$\mathbf{1.983} \pm 0.001$	1.982 ± 0.001
GABA1	MinGap ↑	-0.069 ± 0.006	-0.035 ± 0.030	0.354 ± 0.059	-0.064 ± 0.002	$\mathbf{0.638} \pm 0.010$
	Motif Corr. ↑	0.383 ± 0.206	0.331 ± 0.257	0.724 ± 0.012	-0.050 ± 0.147	$\mathbf{0.839} \pm 0.002$
	3-mer Corr. ↑	0.738 ± 0.050	0.729 ± 0.099	0.601 ± 0.027	0.009 ± 0.215	$\mathbf{0.869} \pm 0.002$
	Shannon Div. ↑	1.297 ± 0.116	1.455 ± 0.048	1.929 ± 0.011	$\mathbf{1.983} \pm 0.001$	1.965 ± 0.001
GABA2	MinGap ↑	-0.078 ± 0.005	-0.080 ± 0.051	0.592 ± 0.039	-0.072 ± 0.001	$\mathbf{0.793} \pm 0.012$
	Motif Corr. ↑	0.510 ± 0.179	0.396 ± 0.076	0.717 ± 0.034	0.002 ± 0.054	$\mathbf{0.737} \pm 0.017$
	3-mer Corr. ↑	0.772 ± 0.079	0.692 ± 0.068	0.614 ± 0.086	0.020 ± 0.197	$\mathbf{0.872} \pm 0.006$
	Shannon Div. ↑	1.510 ± 0.063	1.047 ± 0.481	1.908 ± 0.014	$\mathbf{1.983} \pm 0.001$	1.946 ± 0.001
Glut1	MinGap ↑	-0.001 ± 0.022	0.022 ± 0.025	0.004 ± 0.001	0.034 ± 0.003	$\mathbf{0.408} \pm 0.004$
	Motif Corr. ↑	0.541 ± 0.092	0.496 ± 0.019	0.588 ± 0.050	0.058 ± 0.137	$\mathbf{0.731} \pm 0.017$
	3-mer Corr. ↑	0.822 ± 0.025	0.781 ± 0.068	0.422 ± 0.082	0.054 ± 0.169	$\mathbf{0.831} \pm 0.004$
	Shannon Div. ↑	1.476 ± 0.121	1.370 ± 0.221	1.844 ± 0.033	$\mathbf{1.983} \pm 0.001$	1.976 ± 0.001
Glut2	MinGap ↑	-0.027 ± 0.010	-0.075 ± 0.007	0.028 ± 0.003	-0.064 ± 0.001	$\mathbf{0.495} \pm 0.010$
	Motif Corr. ↑	0.401 ± 0.053	0.496 ± 0.066	0.687 ± 0.042	0.017 ± 0.087	$\mathbf{0.760} \pm 0.018$
	3-mer Corr. ↑	0.755 ± 0.060	0.760 ± 0.098	0.395 ± 0.041	0.011 ± 0.215	$\mathbf{0.771} \pm 0.000$
	Shannon Div. ↑	1.327 ± 0.147	1.303 ± 0.036	1.737 ± 0.045	$\mathbf{1.983} \pm 0.001$	1.895 ± 0.004
Microglia	MinGap ↑	-0.081 ± 0.068	-0.098 ± 0.022	0.786 ± 0.093	-0.095 ± 0.011	$\mathbf{0.949} \pm 0.002$
	Motif Corr. ↑	0.241 ± 0.149	0.447 ± 0.148	0.706 ± 0.007	0.144 ± 0.200	$\mathbf{0.858} \pm 0.004$
	3-mer Corr. ↑	0.682 ± 0.068	0.688 ± 0.017	0.517 ± 0.089	0.116 ± 0.145	$\mathbf{0.900} \pm 0.010$
	Shannon Div. ↑	1.203 ± 0.366	1.312 ± 0.341	1.850 ± 0.049	$\mathbf{1.983} \pm 0.001$	1.973 ± 0.002
OPC	MinGap ↑	-0.072 ± 0.025	-0.058 ± 0.018	0.088 ± 0.124	-0.053 ± 0.002	$\mathbf{0.577} \pm 0.003$
	Motif Corr. ↑	0.490 ± 0.041	0.418 ± 0.170	0.643 ± 0.076	0.017 ± 0.116	$\mathbf{0.781} \pm 0.011$
	3-mer Corr. ↑	0.761 ± 0.103	0.754 ± 0.036	0.535 ± 0.092	0.049 ± 0.128	$\mathbf{0.904} \pm 0.003$
	Shannon Div. ↑	1.292 ± 0.055	1.328 ± 0.032	1.665 ± 0.084	$\mathbf{1.983} \pm 0.001$	1.971 ± 0.002
Oligo	MinGap ↑	-0.103 ± 0.061	-0.095 ± 0.026	0.554 ± 0.489	-0.065 ± 0.011	$\mathbf{0.907} \pm 0.003$
	Motif Corr. ↑	0.271 ± 0.117	0.389 ± 0.108	0.291 ± 0.105	0.039 ± 0.134	$\mathbf{0.806} \pm 0.009$
	3-mer Corr. ↑	0.620 ± 0.136	0.754 ± 0.110	0.329 ± 0.097	0.064 ± 0.172	$\mathbf{0.884} \pm 0.013$
	Shannon Div. ↑	1.315 ± 0.254	1.384 ± 0.146	1.836 ± 0.020	$\mathbf{1.983} \pm 0.001$	1.976 ± 0.003

analysis complements the reward-formulation ablation in Table 3, which changes the objective used during DNA-CRAFT optimization.

Table 14 shows that the leading method within each cell line is consistent across all three metrics. Ctrl-DNA achieves the highest mean specificity scores on HepG2 and K562, while DNA-CRAFT achieves the highest scores on SK-N-SH. Thus, these leading specificity comparisons are preserved under alternative off-target aggregations, although individual method rankings can change. Together with the naturalness results in Table 8, this supports evaluating the methods through their specificity–naturalness trade-off rather than a single specificity score.

A.7 Tree Search vs. Best-of-B Greedy Baseline

To assess the contribution of the complete tree-search procedure, we compare DNA-CRAFT with a tree-free best-of-B greedy baseline, with both methods configured for a maximum of 32 candidate-pool evaluations of $B = 800$ sequences. For each run, we report the best sequence from the final batch of 800 fully denoised candidates for best-of-B sampling. Table 15 shows higher mean specificity for DNA-CRAFT under the matched maximum candidate budget, with the observed differences reflecting the combined search and final-selection procedures.

Table 10: Validation performance of sequence-to-activity models on the MPRA benchmark. Design and Evaluation models are evaluated on the partition opposite to their training partition, using the best-validation checkpoint. Metrics include Pearson correlation (r) and Spearman correlation (ρ).

Model	Cell Line	Pearson r	Spearman ρ
Enformer Design	HepG2	0.865	0.779
	K562	0.862	0.744
	SK-N-SH	0.862	0.789
Enformer Evaluation	HepG2	0.863	0.775
	K562	0.857	0.739
	SK-N-SH	0.863	0.788
Enformer Full	HepG2	0.918	0.829
	K562	0.916	0.794
	SK-N-SH	0.916	0.835
Borzoi Design	HepG2	0.865	0.781
	K562	0.865	0.750
	SK-N-SH	0.865	0.788
Borzoi Evaluation	HepG2	0.864	0.779
	K562	0.857	0.746
	SK-N-SH	0.857	0.784

Table 11: Validation performance of chromatin accessibility predictors. We report AUROC for the eight immune cell types used in the design benchmark, together with their positive-label prevalence. The final rows summarize the mean across these eight cell types and across all 204 accessibility tracks. Models are evaluated on the partition opposite to their training partition using the best-validation checkpoint.

Cell Type	Prevalence (%)	Enformer Design AUROC	Enformer Evaluation AUROC	Borzoi Evaluation AUROC
$CD8^+$ T	5.24	0.925	0.923	0.922
$CD4^+$ T	4.43	0.931	0.928	0.927
NKT	4.04	0.934	0.934	0.934
Naive T	4.67	0.925	0.922	0.919
Memory B	4.02	0.930	0.928	0.926
Plasma B	5.25	0.917	0.914	0.909
Macrophage	9.23	0.907	0.905	0.903
Mast	5.22	0.918	0.918	0.912
Mean (8 immune types)	5.26	0.923	0.922	0.919
Mean (all 204 tracks)	–	0.900	0.898	0.894

A.8 Tree Depth Statistics

We recorded the maximum depth reached across 128 independent SK-N-SH searches with a budget of 32 iterations. Of these, 122 searches (95.3%) reached depth 32, while the remaining six reached depths 4, 4, 20, 21, 25, and 26. Because each iteration expands at most one node, the depth-32 searches allocate their expansions along a single path. The remaining searches show that the iteration budget is not always used to increase the maximum depth. These observations characterize search behavior at the default budget, while the comparison in Table 15 evaluates the performance of the complete search procedures.

A.9 Evaluation of the One-Step MAP Estimator

The one-step MAP estimator (Equation 8) approximates sequence quality at an intermediate diffusion state using a deterministic completion. We evaluated this approximation using child candidates collected from 32 SK-N-SH tree searches at five denoising steps. Each partially masked sequence was completed once using the MAP estimator and eight times using stochastic reverse-diffusion rollouts. We compared its MAP MinGap score with the average MinGap across the eight rollouts. Table 16 reports the mean signed difference, its standard deviation across candidates, and the Pearson correlation between the two scores. The step indices count denoising updates from the fully masked

Table 12: Full results for the CFG guidance scale sweep during class-conditioned Monte Carlo tree guidance ($M = 8$, $B = 800$) for SK-N-SH enhancer design. Values are reported as mean $\pm$ std over 3 independent runs.

γ	**MinGap** (↑)	**MinGap (Full MPRA)** (↑)	**Motif Corr.** (↑)	**3-mer Corr.** (↑)	**Shannon Div.** (↑)	**Wall Time (s)** (↓)
0.0	5.440 ±0.038	6.758 ±0.051	**0.821** ±0.024	0.909 ±0.004	1.972 ±0.001	1021.270 ±9.646
0.5	5.453 ±0.048	6.816 ±0.010	0.803 ±0.011	0.909 ±0.005	**1.973** ±0.001	1712.908 ±15.843
1.0	5.479 ±0.057	6.894 ±0.102	0.809 ±0.010	0.912 ±0.001	1.972 ±0.001	1012.475 ±3.999
2.0	5.538 ±0.018	6.893 ±0.058	0.810 ±0.003	0.914 ±0.004	1.972 ±0.001	1719.055 ±20.178
3.0	5.529 ±0.047	6.871 ±0.034	0.815 ±0.015	0.918 ±0.008	1.972 ±0.001	1718.087 ±10.862
4.0	5.576 ±0.100	6.941 ±0.128	0.810 ±0.014	0.919 ±0.005	1.970 ±0.002	1711.341 ±7.905
5.0	**5.644** ±0.032	**7.020** ±0.046	0.805 ±0.020	**0.920** ±0.004	1.970 ±0.001	1733.893 ±27.677

Table 13: Full results for the fast lookahead pruning sweep during Monte Carlo tree guidance with retained-child limit $M = 8$ and guidance scale $\gamma = 1.0$ for SK-N-SH enhancer design. Each expansion retains $\min(B, 8)$ children. "Std." denotes standard full rollouts without lookahead. Values are reported as mean $\pm$ std over 3 independent runs.

B	**MinGap** (↑)	**MinGap (Full MPRA)** (↑)	**Motif Corr.** (↑)	**3-mer Corr.** (↑)	**Shannon Div.** (↑)	**Wall Time (s)** (↓)
Std.	1.538 ±0.029	1.658 ±0.041	**0.858** ±0.024	**0.947** ±0.010	**1.980** ±0.003	1459.713 ±4.460
2	1.554 ±0.070	1.692 ±0.088	0.811 ±0.021	0.866 ±0.007	1.967 ±0.003	221.196 ±5.063
4	2.069 ±0.069	2.319 ±0.100	0.797 ±0.020	0.899 ±0.019	1.979 ±0.003	217.319 ±8.547
8	2.772 ±0.035	3.200 ±0.071	0.777 ±0.025	0.908 ±0.008	1.979 ±0.002	225.092 ±3.689
16	3.501 ±0.059	4.067 ±0.130	0.801 ±0.043	0.914 ±0.008	**1.980** ±0.001	229.305 ±8.031
32	3.934 ±0.061	4.746 ±0.040	0.809 ±0.004	0.917 ±0.009	**1.980** ±0.001	231.643 ±5.718
64	4.275 ±0.095	5.291 ±0.106	0.814 ±0.007	0.920 ±0.008	1.978 ±0.001	243.243 ±7.647
128	4.785 ±0.057	5.960 ±0.097	0.815 ±0.006	0.915 ±0.002	1.977 ±0.001	257.311 ±3.575
256	5.079 ±0.056	6.266 ±0.112	0.828 ±0.016	0.924 ±0.001	1.975 ±0.001	368.128 ±0.834
512	5.382 ±0.010	6.743 ±0.053	0.825 ±0.007	0.913 ±0.006	1.973 ±0.001	676.850 ±9.089
1024	5.554 ±0.016	6.958 ±0.043	0.792 ±0.033	0.909 ±0.010	1.973 ±0.001	1309.609 ±9.239
2048	**5.590** ±0.043	**6.965** ±0.124	0.819 ±0.017	0.904 ±0.009	1.972 ±0.002	2932.594 ±3.048

state and are distinct from the continuous diffusion time t. Steps 2 and 8 show near-zero Pearson correlation, while the measured correlation increases at later steps, reaching $r = 0.44$ at step 30.

To assess the utility of the completion rule for pruning, we repeated the lookahead search with a candidate pool of $B = 800$, changing only the score used for candidate selection. One configuration scored one-step MAP completions, while the other scored completions obtained by filling masked positions with random nucleotides. The MAP configuration achieved a best MinGap of 6.17 ± 1.12, compared with 3.75 ± 1.30 for random completion. This comparison supports the usefulness of the MAP-based pruning score despite its weak association with rollout scores at early denoising steps.

Table 14: Cross-method specificity evaluation across three human cell lines. The same generated sequences are evaluated using hard MinGap, smooth LogSumExp, and Mean Difference. Higher values are better for all three metrics. Values are reported as mean ± standard deviation over three independent runs, with the highest mean in each row marked in **bold**.

Cell Line	Metric	SMC	TDS	D3	DRAKES	Ctrl-DNA	Ledidi	DNA-CRAFT (Ours)
HepG2	MinGap ↑	0.989 ±1.717	0.880 ±1.053	-0.075 ±0.021	-1.314 ±0.024	**7.464** ±0.270	6.712 ±0.270	7.271 ±0.142
	LogSumExp ↑	0.432 ±1.669	0.332 ±1.059	-0.616 ±0.020	-1.811 ±0.049	**7.001** ±0.298	6.124 ±0.262	6.737 ±0.146
	Mean Diff. ↑	1.144 ±1.660	1.046 ±1.056	0.110 ±0.019	-1.072 ±0.061	**7.744** ±0.303	6.828 ±0.261	7.453 ±0.148
K562	MinGap ↑	1.812 ±0.768	1.735 ±0.660	0.003 ±0.070	-0.033 ±0.032	**8.490** ±0.153	8.004 ±0.078	7.777 ±0.054
	LogSumExp ↑	1.248 ±0.758	1.224 ±0.632	-0.566 ±0.078	-0.457 ±0.024	**8.221** ±0.136	7.644 ±0.077	7.410 ±0.056
	Mean Diff. ↑	1.957 ±0.754	1.954 ±0.624	0.155 ±0.081	0.333 ±0.018	**9.090** ±0.119	8.433 ±0.075	8.201 ±0.056
SK-N-SH	MinGap ↑	0.252 ±0.161	1.170 ±0.715	-0.111 ±0.014	-0.082 ±0.040	3.988 ±0.856	3.202 ±0.258	**5.479** ±0.057
	LogSumExp ↑	-0.366 ±0.178	0.624 ±0.717	-0.683 ±0.013	-0.369 ±0.042	3.717 ±0.910	2.770 ±0.256	**5.222** ±0.050
	Mean Diff. ↑	0.337 ±0.186	1.339 ±0.715	0.037 ±0.019	0.496 ±0.045	4.706 ±0.990	3.549 ±0.253	**6.134** ±0.035

Table 15: Tree-free best-of-B greedy baseline with identical model inference parameters. MinGap (↑) is measured under the held-out Evaluation oracle.

Method	HepG2	K562	SK-N-SH
Best-of-B Greedy	$6.87_{\pm 0.18}$	$6.94_{\pm 0.02}$	$5.36_{\pm 0.08}$
DNA-CRAFT (MCTD)	$\mathbf{7.27}_{\pm 0.14}$	$\mathbf{7.78}_{\pm 0.05}$	$\mathbf{5.48}_{\pm 0.06}$

Table 16: **One-step MAP estimator accuracy across denoising steps for SK-N-SH.** MAP scores are compared with the average scores from eight stochastic rollouts per candidate. Shown are the number of candidates, the mean signed difference (MAP − rollout average) ± standard deviation across candidates, and the Pearson correlation r.

Denoising Step	Sequences	MAP − Rollout (mean ± std)	Pearson r
2	2592	$+0.01_{\pm 0.83}$	0.00
8	1088	$+0.14_{\pm 0.88}$	0.03
16	960	$+0.35_{\pm 0.73}$	0.33
24	896	$+0.31_{\pm 0.68}$	0.41
30	672	$+0.19_{\pm 0.96}$	0.44

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction make three clearly stated contributions that are all supported by experimental evidence in the paper. The scope is accurately represented, and limitations are acknowledged in the conclusion.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The conclusions section discusses the limitation of this study.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper does not present novel theorems or proofs.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The paper provides the complete algorithm pseudocode, all hyperparameters, data pre-processing and splitting procedures, baseline implementation details for reproducibility in the appendix.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code and and model weights for reproducing reported experiments are available at `https://github.com/epigen/dna-craft`.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper specifies optimizers, learning rates, batch sizes, epochs, noise schedules, architecture details, data split procedures, inference parameters, and baseline configurations across the main text and appendix.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Benchmark tables report means and standard deviations across three independent generation runs. Table captions identify the relevant aggregation and uncertainty.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: GPU type is specified, and wall-clock times are reported in the ablation studies.

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research uses publicly available data, involves no human subjects, and is conducted entirely computationally. There are no ethical concerns.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Broad impact of the paper is discussed in the Broader Impact paragraph under Acknowledgments.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The models and datasets used in this work focus on regulatory DNA sequence design for scientific research. They do not pose risks of misuse, and no high-risk assets such as language models, generative image models, or scraped personal data are involved.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All datasets, models, and tools are properly cited.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The new model and algorithm are documented via architecture tables, training details, and complete pseudocode in the appendix.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: No crowdsourcing or human subjects are involved.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: No human subjects research is conducted.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs were not used as an important, original, or non-standard component of the core methodology.