
Multi-Objective-Guided Generative Design of mRNA with Therapeutic Properties

Sawan Patel^{*1} Sophia Tang^{*2} Yinuo Zhang³ Pranam Chatterjee^{2,4} Sherwood Yao¹

Abstract

Therapeutic mRNA design requires simultaneous optimization of multiple competing properties, including stability, translation efficiency, and protein expression—a challenge that current single-objective approaches struggle to address systematically. We introduce **mRNA generation via Unrolled Trajectories and Informed Latent UpdateS (mRNAutilus)**, the first multi-objective guided generative model for simultaneous mRNA codon optimization and *de novo* design of untranslated region (UTR) sequences. mRNAutilus combines a Masked Diffusion Model (MDM) trained on 6 million evolutionary-scale mRNA sequences with Monte Carlo Tree Search to guide generation toward Pareto-optimal solutions across multiple therapeutic properties. As objective functions, we use the latent mRNAutilus embeddings to train classifiers for predicting half-life, ribosome profiling, and translation rate. Guided by our classifiers, we demonstrate that mRNAutilus successfully generates diverse, high-quality mRNA sequences with enhanced properties compared to natural sequences while preserving coding regions for therapeutic applications, establishing a new paradigm for rational mRNA design.

1. Introduction

Messenger RNA (mRNA) vaccines have grown in popularity within both the pharmaceutical industry and academic research (Pardi et al., 2018; 2020; Jackson et al., 2020; Zhang et al., 2019; Barbier et al., 2022). As demonstrated during

the COVID-19 pandemic, novel mRNA vaccines were developed, mass-manufactured, and widely distributed in a matter of months with demonstrable success (Park et al., 2021). This technology benefits from the rapid manufacturability of mRNAs and the use of a vaccine-administrable lipid nanoparticle composition for cell-specific antigen production. Ultimately, this evokes a natural immune response without the use of a living pathogen. mRNA vaccine development is ongoing, with several candidates progressing into Phase 3 clinical trials (Soens et al., 2025; Weber et al., 2024; Wilson et al., 2023). Aside from vaccine development, mRNAs are also commonly used in other contexts, such as in protein replacement therapies, gene editing, diagnostics, and for general protein expression.

mRNAs are long, single-stranded ribonucleic acid polymers commonly decomposed into the following segments: the 5' untranslated region (UTR), coding sequence (CDS), 3'UTR, and a poly-A tail. The CDS provides all mRNAs with their identity as it is directly translated into a protein product. Importantly, the choice of codons within the CDS can have a significant impact on the expression and stability of mRNA in an intracellular environment (Mauro & Chappell, 2014). The 5' and 3'UTRs also play vital roles in mRNA efficacy; the former is the site of translation initiation, where the 5' cap (in eukaryotes) is recognized by ribosome-associated scaffolding proteins, which recruit the functional ribosomal subunits. Notably, certain secondary structure motifs in the 5'UTR can prohibit, deter, or delay the translation initiation procedure (Xiao et al., 2020; Leppek et al., 2018). In viral genomes, 5'UTR structural elements have been shown to limit nuclease degradation, while translation is initiated in a cap-independent manner (Ng et al., 2017). The 3'UTR, often longer in sequence length, notably affects mRNA stability and localization. RNA-binding proteins dock to the 3'UTR and recruit different functional proteins. Critically, inter-region interactions have also been deemed influential determinants of mRNA efficacy (Lin & Bundschuh, 2015; Lin et al., 2022a).

Despite their widespread use, mRNAs are generally unstable with high turnover and are prone to nuclease degradation (Ross, 1995). Though this is a feature inherent to mRNA, it often compromises its therapeutic effectiveness as limited

^{*}Equal contribution ¹Atom Bioworks, Cary, NC, United States ²Department of Computer and Information Science, University of Pennsylvania ³Center of Computational Biology, Duke-NUS Medical School ⁴Department of Bioengineering, University of Pennsylvania. Correspondence to: Pranam Chatterjee <pranam@seas.upenn.edu>, Sherwood Yao <s.yao@atombioworks.com>.

Proceedings of the Workshop on Generative AI for Biology at the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

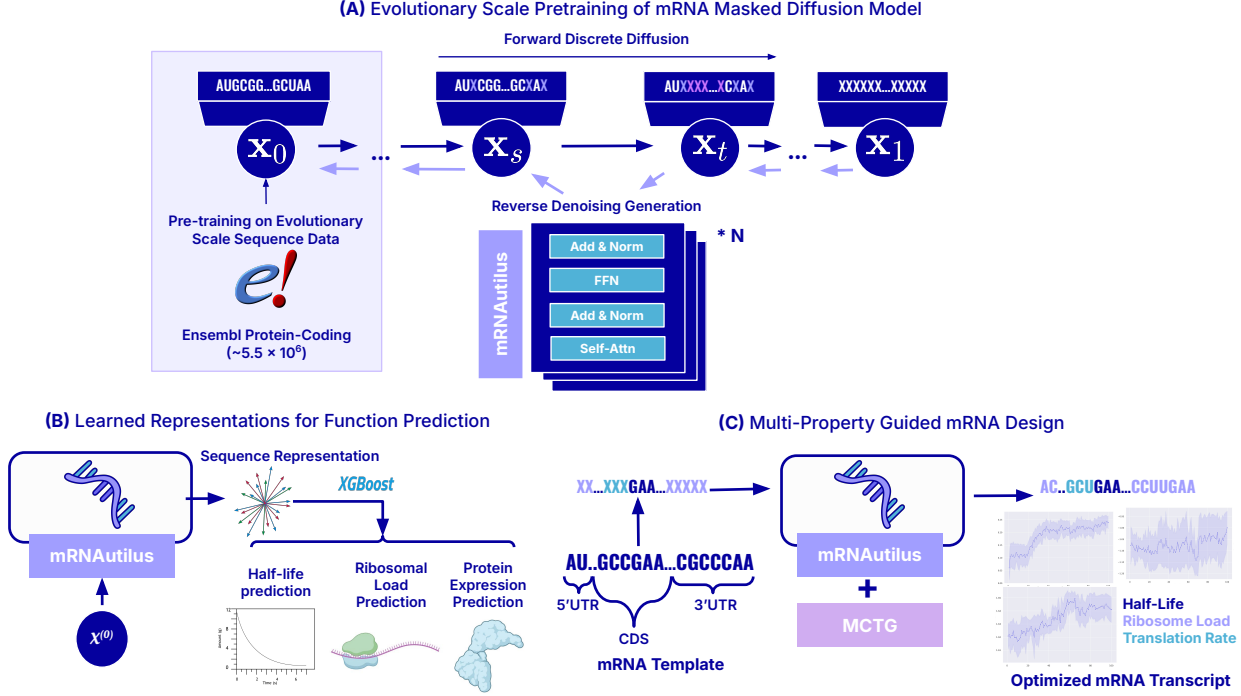


Figure 1. mRNAAutilus overview (A) Masked diffusion model pretraining and *de novo* mRNA sequence generation. (B) mRNA function and property prediction using XGBoost regression analysis from model embeddings. (C) Simultaneous codon optimization and UTR generation via multi-property co-optimization, using PepTune classifier guidance during sampling (Tang et al., 2025).

protein expression ultimately yields unsatisfactory drugability. Designing an mRNA to achieve optimality on a particular property of interest, such as half-life, is crucial. While existing work has explored the optimization of mRNA untranslated regions (UTRs) and coding sequence (CDS) (Mauro & Chappell, 2014; Fu et al., 2020; Leppek et al., 2022; Rosa et al., 2022; Gebre et al., 2022; Kim et al., 2022), effective mRNA therapies possess multiple desirable properties that require optimization without compromising any other property.

Leveraging predictive models to screen mRNA candidates *in silico* can provide quick insights when evaluating the candidacy of an mRNA, with demonstrated experimental success. Prior works are diverse, ranging from graph-based algorithms capable of designing a well-rounded SARS-CoV-2 spike protein mRNA (Zhang et al., 2023) to language models with high predictive capabilities on benchmarks (Li et al., 2025; Yazdani-Jahromi et al., 2024; Zhang et al., 2025; Wood et al., 2025; Patel et al., 2025). However, generative language models for mRNA design remain underexplored. While guided generation of mRNA sequences optimized for 3'UTR-conferred stability has been investigated (Morrow et al.), stability is not solely responsible for mRNA fitness, particularly if it compromises expression. Therefore, there remains a need to steer generative mRNA design toward

optimality across multiple therapeutic properties. Moreover, optimal mRNA design requires consideration beyond the UTRs, with modifications to the CDS seen as crucial to ongoing research.

Recent work in the *de novo* peptide design space has faced similar challenges. In response, Tang et al. (2025) introduced PepTune, which leveraged a Monte-Carlo Tree Guidance strategy to explore the vast design space of non-canonical and cyclic peptides and generate *de novo* sequences optimized across multiple therapeutic properties, including binding affinity, solubility, and non-toxicity.

Here, we present **mRNA generation via Unrolled Trajectories and Informed Latent Updates (mRNAAutilus)**, a multi-objective guided masked diffusion model specific for simultaneous mRNA codon optimization and full-UTR design with Monte-Carlo Tree Guidance (MCTG) (Tang et al., 2025). We demonstrate the unique capability for mRNAAutilus to explore the vast design space of mRNA transcripts with optimal therapeutic properties.

2. Methods

2.1. Masked Discrete Diffusion Framework

Notation Let $x = (x^1, \dots, x^L) \in \mathcal{V}^L$ denote a sequence with length L , where each element occupies one of D possible states. For simplicity of notation, we let $x \in \Delta^V$ denote the distribution of a single token in the sequence. The extension to larger dimensionality follows with the requisite factorization.

We denote $p_1(x)$ as the prior noise distribution, $p_0(x) = p_{\text{data}}(x)$ as the target data distribution, and $p_t(x)$ as the time-dependent marginal distribution for $t \in [0, 1]$ that interpolates between p_0 at $t = 0$ and p_1 at $t = 1$. Given a noisy sequence $x_1 \sim p_1$ and a clean sequence $x_0 \sim p_0$ where $x_0 \in \{0, 1\}^V$ is a one-hot vector with 1 in the true state, we define the *reverse posterior* as the conditional probability of obtaining x_0 given x_1 as $p_{0|1}(x_0|x_1)$. Instead of sampling directly from the prior distribution, we sample iteratively from a marginal reverse posterior $p_{s|t}(x_s|x_t)$ that defines the probability of sampling a slightly less noisy sequence x_s at time $s = t - \frac{1}{T}$, where T is the total number of sampling steps. We define the *conditional reverse posterior* given a clean sequence x_0 as $p_{s|t,0}(x_s|x_t, x_0)$. We train a neural network with parameters θ to predict the distribution $p_{s|t,0}(x_s|x_t, x_0)$, and sample from the parameterized reverse posterior $p_{s|t}^\theta(x_s|x_t)$ at inference.

Discrete Diffusion When generating sequences of *discrete* tokens, we leverage the discrete diffusion framework (Austin et al., 2021), where each state is a V -dimensional vector probability distribution over the vocabulary of V possible states. During the forward noising process, we apply categorical noise to x_0 for continuous times $t \in [0, 1]$ in the form $p_{t|0}(x_t|x_0) = \text{Cat}(x_t; \alpha_t x_0 + (1 - \alpha_t)\pi)$, where α_t is the probability that the token does not change, $\pi \in \Delta^V$ is a noisy distribution, and α_t is a time-dependent noise schedule. We define the marginal distribution p_t as the expectation of conditional distributions given by

$$p_t(x_t) = \mathbb{E}_{p_{\text{data}}} [p_{t|0}(x_t|x_0)], \quad (1)$$

which defines a tractable learning objective by conditioning on sequences in the dataset $x_0 \sim p_{\text{data}}$.

Masked Discrete Diffusion Model For our conditional flow $p_{t|0}(x_t | x_1)$, we leverage the mask discrete diffusion formulation given its proven efficacy in generating high-quality sequences with long-range dependencies (Sahoo et al., 2024; Shi et al., 2024; Ou et al., 2024; Zheng et al., 2024). The mask formulation defines the conditional flow as

$$p_{t|0}^{\text{mask}}(x_t|x_0) = \text{Cat}(x_t; tx_0 + (1 - t)M) \quad (2)$$

where M denotes the one-hot vector of the [MASK] token. At $t = 1$, the distribution converges to only the fully masked sequence $p_{t|0}^{\text{mask}}(x_t|x_0) = \delta(x_t - M)$, and

at $t = 0$, the distribution converges to the clean sequence $p_{t|0}^{\text{mask}}(x_t|x_1) = \delta(x_t - x_0)$. From Bayes' rule, the conditional reverse posterior can be derived

$$p_{s|t,0}(x_s|x_t, x_0) = \begin{cases} (1 - \frac{s}{t})x_0 + \frac{s}{t}M & x_t = M \\ x_s & x_s \neq M \end{cases} \quad (3)$$

However, since x_0 is unknown during inference, we train a parameterized model $p_{s|t}^\theta(x_s|x_t)$ which is trained to predict the mixture of conditional reverse posteriors given only x_t .

Loss Functions We aim to approximate the denoising distribution $p_{0|t}(x_0|x_t)$ that we parameterize with a neural network $p_{0|t}^\theta(x_0|x_t)$. Within each training batch, we apply masking for uniformly sampled time steps $t \sim \mathcal{U}(0, 1)$ and use the model to predict $x_0^\theta \sim p_{0|t}(x_0|x_t)$. Then, we minimize the continuous-time negative evidence lower bound (NELBO) between the predicted distribution x_0^θ and the clean one-hot vector x_0 given by

$$\mathcal{L}_{\text{NELBO}}^\infty = \mathbb{E}_{t, p_{t|0}(x_t|x_0), p_0(x_0)} \left[-\frac{1}{t} \log \langle x_0, x_0^\theta \rangle \right] \quad (4)$$

where the ∞ superscript denotes the continuous-time formulation from (Sahoo et al., 2024). This loss is equivalent to the negative-log loss (NLL), which is 0 when $x_0 = x_0^\theta$.

Unconditional Sampling To sample from the data distribution, we start with a fully masked sequence $x_1 = [M]^L$ of length L . Then for a sequence of discretized time steps $t = \{1, \dots, \frac{1}{T}\}$ where T is the number of sampling steps, we predict $x_0^\theta \sim p_{0|t}^\theta(x_0|x_t)$ with the parameterized reverse posterior. Then, we compute the marginal reverse posterior as $p_{s|t}^\theta$ as

$$p_{s|t}^\theta(x_s|x_t) = \begin{cases} (1 - \frac{s}{t})x_0^\theta + \frac{s}{t}M & x_t = M \\ x_s & x_s \neq M \end{cases} \quad (5)$$

After applying Gumbel noise $g \sim -\log(-\log(u + \epsilon) + \epsilon)$, where $u \sim \mathcal{U}(0, 1)$ and $\epsilon = 10^{-10}$, to integrate stochasticity, we sample the argmax token from the perturbed distribution $x_s = \arg \max(\log p_{s|t}^\theta(x_s|x_t) + g)$. Given our constraint that masked tokens remain unchanged after unmasking, we change x_s only when $x_t = M$. After T unmasking steps, we ensure all tokens are unmasked and return the generated sequence x_0 .

2.2. mRNAutilus

Encoder Architecture mRNAutilus is a BERT-style mRNA generative language model trained on a masked diffusion objective for mRNA sequence design, which parametrizes $p_{0|t}^\theta(x_0|x_t)$ defined above. Given an input sequence $x = (x^1, \dots, x^L)$ of length L , the model maps it

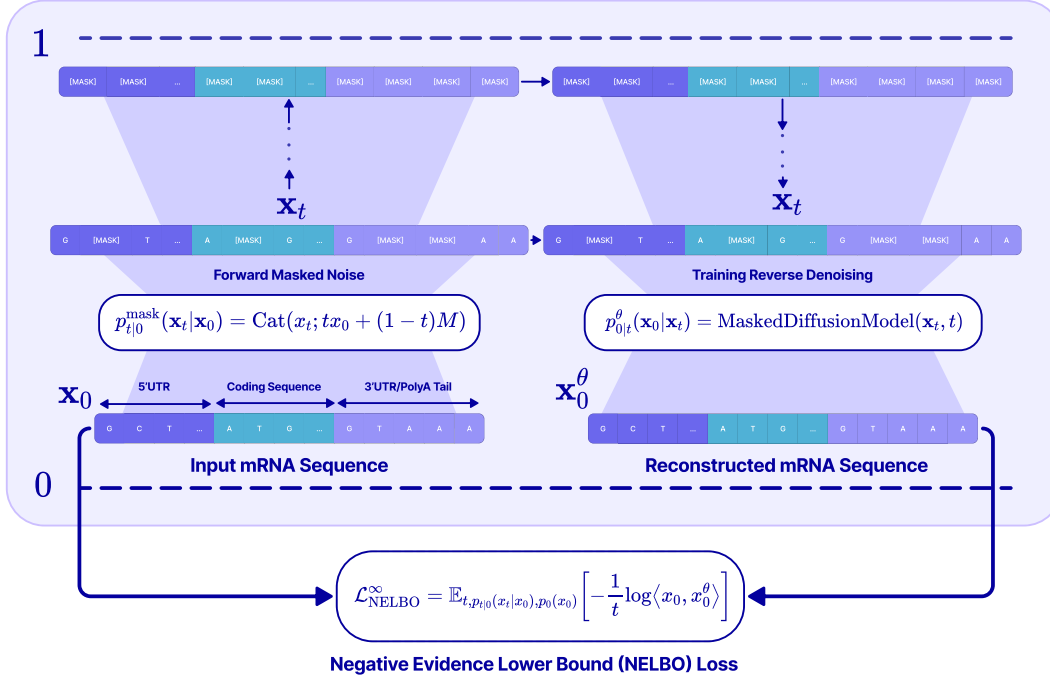


Figure 2. **mRNAutilus discrete diffusion framework.** mRNAutilus is a masked discrete diffusion model pretrained on evolutionary-scale mRNA sequence data.

to a sequence of token embeddings of shape $(L \times d)$. The architecture consists of several transformer blocks, where each block contains a multi-head self-attention module with 20 attention heads, a feed-forward network (FFN) using the SwiGLU activation function (Shazeer, 2020), and Layer Normalization within residual connections to stabilize training dynamics. mRNAutilus adopts Rotary Positional Embeddings (RoPE) (Su et al., 2024) to encode both absolute and relative positional information, thereby improving its ability to model long-range dependencies in mRNA sequences. To enhance computational efficiency, we also utilize FlashAttention-2 (Dao, 2024). In total, mRNAutilus has 150M parameters across all modules.

Data Ensembl is a genome browser containing the largest curated, annotated deposit of vertebrate mRNA (Dyer et al., 2025). Using the Ensembl Perl API, we queried the database for protein-coding mRNA transcripts across 342 vertebrate species and obtained approximately 5.5 million total sequences. Transcripts were filtered to only include sequences with a valid coding sequence length, appropriate start/stop codons, and no missing UTRs. Following this procedure, 5,526,848 sequences remained. An overview of the dataset is shown in Figure A1.

We leverage a hybrid tokenization scheme, encoding the coding sequences (CDS) and UTRs for each mRNA tran-

script separately. We tokenize the coding sequence as 3-mers, effectively by codon. UTRs are separately encoded at single-base resolution. Following this process, the 5'UTR tokens, CDS tokens, and 3'UTR tokens are concatenated in that order. For the CDS, we reference the standard codon table, which spans 64 codons, 1 primary start codon, and 3 stop codons. For the UTR sequences, all uracil (U) bases are replaced with thymine (T) to maintain consistency across representations, resulting in a primary UTR vocabulary of four nucleotides: 'A', 'C', 'T', and 'G'. Additionally, we include IUPAC degenerate nucleotide codes such as 'I', 'R', 'Y', 'K', 'M', 'S', 'W', 'B', 'D', 'H', 'V', 'N', and '-' to account for sequence ambiguities commonly found in biological datasets. We do not discard any sequences.

We include special tokens [CLS], [EOS], [MASK], [UNK], and [PAD], which denote sequence classification, end-of-sequence markers, masked tokens for pretraining, and padding, respectively. Sequences are truncated or padded to a total of 7,500 tokens for batched training. In total, our vocabulary spans 86 tokens.

Pretraining Our model is trained for 100K updates from scratch, with a token batch size of 200K tokens per gradient step. We additionally utilize a batch sampler to construct batches using sequences of similar lengths. We use the aforementioned weighted cross-entropy loss where the weight

is $\frac{1}{mask_ratio}$, where $mask_ratio$ reflects $t \in [0, 1]$ and is clamped to $[\frac{1}{500}, 1]$. During training, tokens are sampled for masking uniformly at random at a rate modulated by the sampled timestep. Additionally, any codon token not coding for the same amino acid as the codon in the true mRNA sequence is rejected during denoising. We find that including this rule slightly improves performance on the training objective (not shown).

2.3. Multi-Objective-Guided Sampling

Pareto Optimality In multi-objective optimization, rather than a single optimal solution, there is a set of *Pareto-optimal solutions*, where not one objective can be improved without sacrificing performance in another objective. Intuitively, this set *maximizes* the trade-offs between solutions such that all objectives are exhausted to their limit.

Given a vector of objective functions $\mathbf{f}(x) = [f_1(x), \dots, f_K(x)]$ where $f_k : \mathcal{V}^L \rightarrow \mathbb{R}$, suppose we aim to identify the set of solutions x that *maximizes* all K objective functions. For two sequences x, y , we say x *dominates* y if for all K objectives, the score for sequence x is strictly greater than the score for sequence y

$$\underbrace{\forall k \in [1, \dots, K] \quad f_k(x) > f_k(y)}_{\text{denoted as } \mathbf{f}(x) \succ \mathbf{f}(y)} \quad (x \text{ dominates } y)$$

A Pareto-optimal sequence x^* is *non-dominated*, such that no other sequence in the solution space $y \in \mathcal{X}$ *dominates* it. This means it is just as optimal or more optimal than any other sequence in the solution space. Given that there are a potentially infinite number of trade-offs between objectives, we can define a *Pareto-optimal set* \mathcal{P}^* where all sequences in the set are non-dominated.

$$\mathcal{P}^* = \{x \mid \forall y \in \mathcal{X}, \mathbf{f}(x) \succeq \mathbf{f}(y)\} \quad (6)$$

Since the size of \mathcal{P}^* is potentially unbounded, the goal of multi-objective optimization is to approximate a finite set of Pareto-optimal solutions in a tractable manner.

Monte-Carlo Tree Guidance To guide the unconditional MDM towards Pareto-optimality across multiple properties relevant to mRNA therapeutic efficacy, we leverage the Monte Carlo Tree Guidance (MCTG) algorithm introduced in Tang et al. (2025) for multiobjective guidance of discrete diffusion across several pre-trained classifier models (Algorithm A1).

Initialization We initialize a tokenized sequence with a partially masked coding sequence (CDS), masked 5'UTR sequence, and a masked 3'UTR sequence. This becomes the root node of the MCTS tree with timestep $t = 0$. Additionally, we construct a vector scoring function $\mathbf{f} : \mathcal{V}^L \rightarrow \mathbb{R}^K$ that returns a score vector given a clean input sequence $x_0 \in$

\mathcal{V}^L and initialize an empty Pareto-optimal set $\mathcal{P}^* = \{\}$. We define the following hyperparameters: number of MCTS iterations N_{iter} and number of children nodes N_{child} . A template CDS is also provided at input, which is referenced during expansion.

Selection To determine the next unmasking step from a parent node with an already expanded set of child sequences $x_i^{\text{child}} \in \text{children}(x)$, we compute a selection score vector $\mathbf{U}(x_i^{\text{child}}) = \mathbf{R}(x_i^{\text{child}}) / N_{\text{visit}}(x_i^{\text{child}})$ that divides the cumulative reward vector $\mathbf{R}(x_i^{\text{child}})$ of a child node to the number of times it was visited $N_{\text{visit}}(x_i^{\text{child}})$. This encourages exploration of child nodes that have not been visited, in addition to high-reward nodes. Then, we select from the Pareto-optimal set of child nodes based on the normalized selection score vector.

$$\begin{aligned} \mathcal{P}_{\text{select}}^* = \{ & x_i^{\text{child}} \mid \nexists x_j^{\text{child}} \in \text{children}(x) \\ & \text{s.t. } \mathbf{U}(x_j^{\text{child}}) \succ \mathbf{U}(x_i^{\text{child}}) \} \end{aligned} \quad (7)$$

If we select a non-leaf node, we restart the selection process with the selected node as the new parent until we reach an expandable leaf node or a fully unmasked sequence. In the case of a fully unmasked sequence, we restart the selection from the root node.

Expansion From the leaf node x_t at time t , we sample N_{child} distinct sequences by applying Gumbel noise to the denoising distribution $p_{0|t}^\theta(x_0|x_t)$

$$x_0^i \sim \log p_{0|t}^\theta(x_0|x_t) + g_i \quad (8)$$

where $g_i \sim \text{Gumbel}(0, 1)$. Then, for each child node x_i^{child} , we randomly select k positions in x_t to unmask according to x_1^i .

During sampling, we employ constraints to ensure that an invalid mRNA sequence is not produced during generation. Specifically, we do not allow codon tokens to be sampled in the position of a UTR token. Additionally, for any masked codon token, we permit only codons encoding for the same amino acid as the corresponding codon in the input template to be sampled. This constraint guarantees that the sampled mRNA sequence does not form an unintended protein product.

Rollout For each child node x_i^{child} , we use self-planned ancestral sampling to obtain a fully unmasked sequence \hat{x}_i^{child} . Then, we feed each of the clean sequences into the vector score function $\mathbf{f}(\hat{x}_i^{\text{child}})$ to obtain a score vector. We then compute the reward vector $\mathbf{r}(x_i^{\text{child}}) \in \mathbb{R}^K$ of each child x_i^{child} where each element is the fraction of sequences in the Pareto-optimal set $x^* \in \mathcal{P}^*$ that x_i^{child} dominates.

$$r_k(x_i^{\text{child}}) = \frac{1}{|\mathcal{P}^*|} \sum_{x^* \in \mathcal{P}^*} \mathbf{1}[f_k(x_i^{\text{child}}) \geq f_k(x^*)] \quad (9)$$

For all N_{child} children sequences, we add x_i^{child} to \mathcal{P}^* if it is *non-dominated* by all sequences $x^* \in \mathcal{P}^*$ and remove all sequences in \mathcal{P}^* dominated by some x_i^{child} .

Backpropagation For each child node x_i^{child} , we use its reward vector $\mathbf{r}(x_i^{\text{child}})$ to initialize the cumulative reward $\mathbf{R}(x_i^{\text{child}}) = \mathbf{r}(x_i^{\text{child}})$ and the number of visits $N_{\text{visit}}(x_i^{\text{child}})$ is set to 1.

We then trace backward from x_i^{child} through its ancestors up to the root node x_0^{root} , updating each node x along the path by accumulating the child rewards and incrementing the visit count:

$$\mathbf{W}(x) \leftarrow \mathbf{W}(x) + \sum_{i=1}^{N_{\text{child}}} \mathbf{r}(x_i^{\text{child}}) \quad (10)$$

$$N_{\text{visit}}(x) \leftarrow N(x) + 1 \quad (11)$$

These cumulative updates influence future selection, prioritizing unmasking paths that lead to higher-reward sequences for further expansion. After N_{iter} iterations of these four steps, we return the set of Pareto-optimal mRNA sequences optimized across the set of properties $x^* \in \mathcal{P}^*$ and their corresponding score vectors $\mathbf{f}(x^*)$.

3. Results

3.1. Unconditional Generation

Though unconditional mRNA generation is not a practical use of mRNAutilus, we compare unconditionally generated mRNAs to vertebrate mRNAs from the Ensembl training corpus to ensure some degree of correspondence. To facilitate this study, we unconditionally generated three mRNA libraries of lengths 1,000, 2,500, and 5,000, respectively, and compared each to mRNAs of the same length in the training corpus. For all populations, we consider libraries of size 200 in total. For all libraries, we characterize the similarity of each population of equal length by investigating the frequency of contained Kozak consensus sequences, GC content, and predicted minimum free energies of each sequence.

Kozak Consensus Sequences The Kozak consensus sequence is a highly-conserved motif found within eukaryotic mRNA transcripts and typically serves as the primary initiation site of protein translation (Kozak, 1999). Originally determined by sequencing a small number of mRNA transcripts across a small number of vertebrate species, it was later found to be a common hallmark of all eukaryotic mRNAs (Kozak, 1987; 1991). As a characteristic feature, we sought to compute the frequency with which these motifs appeared in all mRNA sequences, both generated and from the Ensembl training corpus. For this analysis, we identified the Kozak Consensus sequence using the following regex

string: "(A—G)CCATGG". As shown in Figure 3A, the frequency of Kozak sequences within mRNAutilus-generated sequences reflects the frequency found in nature for the length 1,000 and 2,500 size libraries. However, this similarity tapers off as sequence length grows to 5,000 nucleotides (nt).

GC Content GC content is the percentage of guanine and cytosine nitrogenous bases present in a nucleic acid molecule. For mRNA, high-GC content at the 5'UTR is a hallmark of eukaryotic mRNAs (Leppek et al., 2018). Though we consider unconditionally-generated mRNAs here with no rigid division into a set of UTRs and a CDS, we still compare the GC contents of unconditionally-generated libraries to mRNAs within the Ensembl corpus in Figure 3B. We see that, though the breadth of the GC content distribution in true mRNAs is encompassed by those in their generated library counterparts, the range of GC contents is far wider in those generated by mRNAutilus. This disparity increases with sequence length.

Minimum Free Energy Minimum free energy (MFE) refers to the lowest possible energy state of a system, in this case, the mRNA in intracellular-like buffer conditions. We predict the free energies for all mRNA sequences in our libraries using ViennaRNA (Lorenz et al., 2011) and show their distributions in Figure 3C. Though similarity within the two populations persists at length 1,000, similarity deteriorates as sequence length increases. Notably, though the sequences occupy lower energy states as length increases, the overall distribution of MFEs is highly conserved in both the generated and true populations.

3.2. mRNA Property Classifiers

mRNA efficacy is highly dependent on factors that are largely influenced by sequence. We train classifiers for several key mRNA properties, including half-life, ribosomal load, and translation rate, from sequence embeddings from the learned mRNAutilus latent space.

Half-Life The half-life of a substance is the time required for it to decay to half of its original amount (or concentration within a solution). While more complex in the context of therapeutics, as half-life partially controls protein product abundance, dosage, and overall design flexibility, it is a critical factor to consider and tune when designing an mRNA therapeutic. For example, the mRNA concentration must be controlled to produce an antigen that can produce an immune response, but not induce a severe immunogenic effect.

We trained an XGBoost boosted tree (Chen & Guestrin, 2016) regression model using mRNAutilus embeddings that predicts half-life directly from mRNA sequence (Table A1).

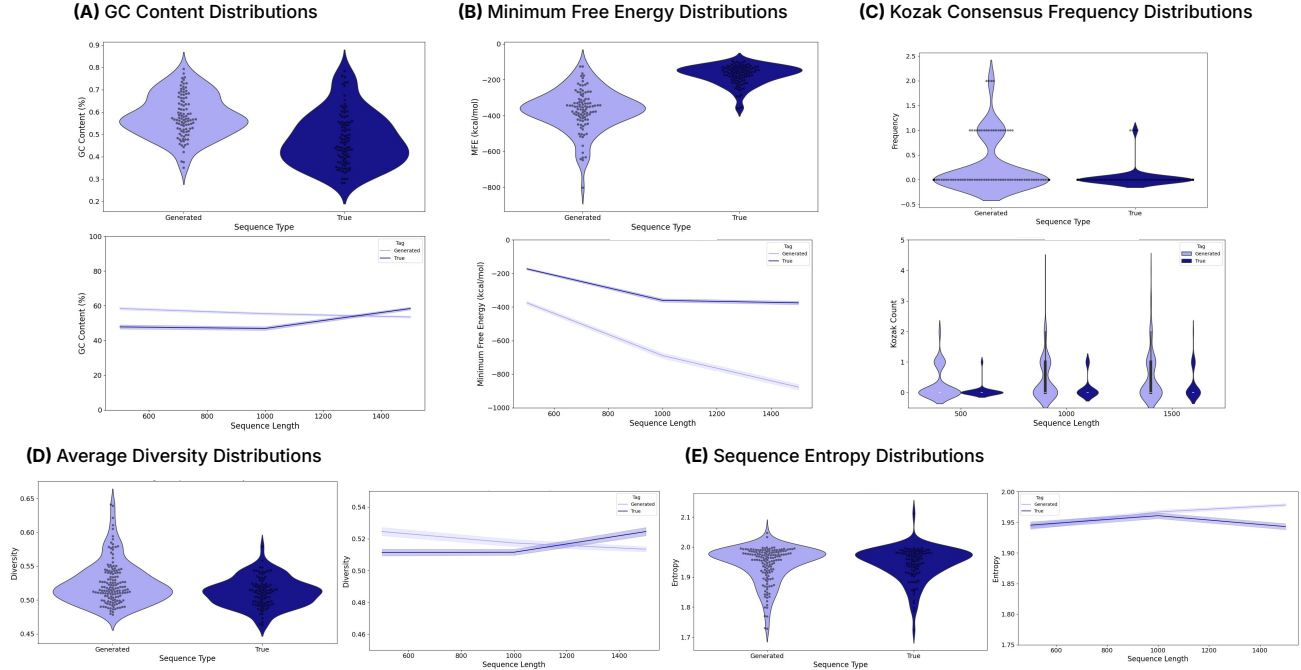


Figure 3. Unconditional mRNA generation results. (A) GC content, (B) minimum free energy, and (C) kozak motif frequency evaluations on unconditionally-generated sequences. (D) Average sequence diversity and (E) Shannon entropy evaluations on unconditionally-generated sequences. Sequences of length 1,000 have their distributions shown (left), while the average values for each property at each length are shown as line plots (right).

Paired data was collected from (Agarwal & Kelley, 2022), which sourced transcriptome-wide half-life data from human samples across 16 publications and deposits. Overall, it contains data for nearly 13,000 transcripts across 39 human samples. Results are shown in Figure 4A.

Ribosome Profiling Following transcription, processing (5'capping, splicing, and 3'polyadenylation), and nuclear export, a mature mRNA transcript will interact with several proteins before ultimately forming a complex with ribosomal subunits to initiate translation. Several ribosomes bound to an mRNA transcript form a polysome, which serves as a proxy for mRNA fitness up until an actual protein product is formed. Polysome formation can be quantified using polysome profiling. Similarly to half-life, we train a regression model on ribosome profiling data collected from (Leppek et al., 2022) that spans 233 engineered mRNA transcripts using a novel PERSIST-Seq assay in HEK293T cells (Figure 4B).

Translation Rate Protein expression is the ultimate purpose of mRNA therapeutics. Following polypeptide release, folding, and trafficking, a mature protein product can perform its function within the cell. Quantifying the projected protein product formed within the cell also fundamentally

determines the efficacy of the originating transcript. In a therapeutic context, a sufficient concentration of antigen must exist in circulation to elicit an immune response. We train a regression model that predicts protein abundance (log-10 scale) from conjoined mRNA sequencing and ribosomal profiling experiments conducted in human B2 and U2OS cells (Figure 4C) (Eichhorn et al., 2014).

3.3. Multi-Objective Guided mRNA Generation

In this section, we evaluate the capability of mRNAutilus to simultaneously generate both the 5' and 3'UTRs and optimize the coding sequences for mRNAs commonly used in engineering and therapeutic contexts using MCTS-guided diffusion. Using the classifiers described in the previous section, we design transcripts for *P. pyralis* luciferase, SARS-CoV-2 spike glycoprotein (S-Protein), and the human mucin 1 (MUC1) transmembrane protein.

***P. Pyralis* luciferase protein** *P. pyralis* luciferase, or Fluc, is an enzyme that catalyzes the oxidation of firefly luciferin, yielding bioluminescence. Due to this natural luminescence mechanism, it is commonly used in research as a reporter for evaluating the transcriptional activity profile in cells transfected with a construct containing the luciferase gene

Table 1. Comparison of property scores for wild-type mRNA transcript and codon optimized mRNA transcript with fully-designed UTR sequences. 24 sequences were generated for the Fluc template and 12 were generated for the S-Protein and MUC1 templates. The scores of a single optimized transcript with the highest overall scores are shown in comparison to its wild-type counterpart.

Template Gene	Half-Life (log-10 hours; \uparrow)		Ribosome Profiling (MRL) (\uparrow)		Translation Rate (log-10 scale; \uparrow)	
	wild-type	designed	wild-type	designed	wild-type	designed
Fluc	-0.112	0.537	1.58	1.62	-1.22	-0.451
SARS-CoV-2-S-Protein	-0.590	0.297	1.61	1.61	-0.741	-0.0562
MUC1	-0.092	0.026	1.59	1.49	-0.630	-0.0389

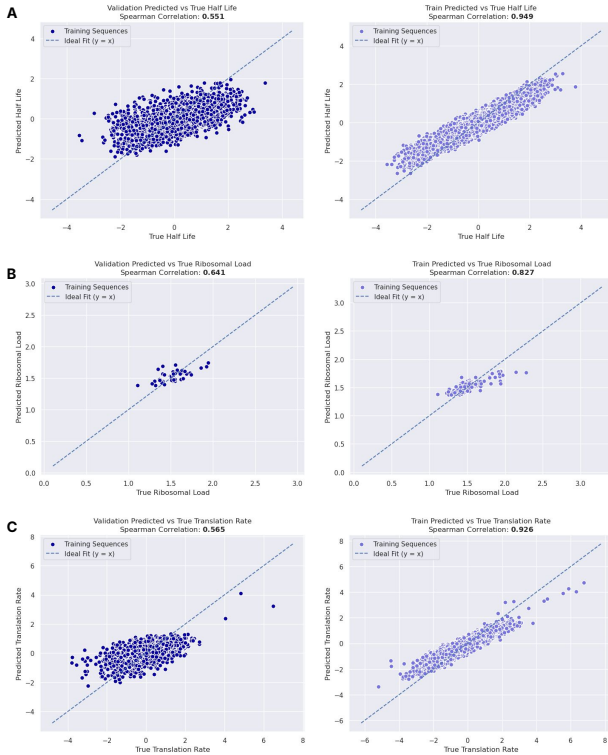


Figure 4. Correlation plots for property classifiers. (A) Half-life, (B) ribosome profiling, and (C) translation rate correlation for validation (left) and training (right) data.

controlled by a promoter of interest. As a commonly used protein in R&D efforts, we sought to design a more optimal firefly luciferase mRNA using mRNAutilus. The template luciferase was constructed referencing the TriLink Clean-Cap Fluc mRNA ORF (BioTechnologies, n.d.). The lengths of the UTRs were arbitrarily set at 41 and 100 nt, respectively.

We performed guided design of Fluc mRNA via two methods: (1) uniform masking of the coding sequence at a 15% masking rate and (2) uniform masking of the coding sequence up until the 50th codon at a 50% masking rate. In

the latter method, we limited model-permitted design to the 50th codon, as recent work has demonstrated that codons at the beginning of the coding sequence primarily affect protein expression (Boël et al., 2016). While prior work highlights that this effect is limited to around the 16th codon, we extend the masking window to 50 since we observe that the model makes fewer changes on sequences that are already optimized, such as the Fluc CleanCap sequence. The Pareto front size was set to 24 sequences, with the number of MCTS iterations set to 200.

SARS-CoV-2 S-Protein The S-Protein is a trimeric class I fusion protein on the surface of the SARS-CoV-2 virus, which directly mediates viral entry into host cells. It contains several subunits, notably the receptor-binding domain (RBD), which binds the human ACE2 receptor. This protein, albeit across different isoforms, is the primary target for all current therapeutic efforts against COVID-19. For the design task, we constructed a template mRNA transcript for the wild-type S-protein around the publicly available CDS (National Center for Biotechnology Information, 2020). Referencing an S-Protein mRNA vaccine patent (Corbett et al., 2021), we set lengths 47 and 119 for the 5' and 3'UTRs, respectively. We performed simultaneous 5'/3' UTR design and partial codon optimization of the CDS, with uniform masking at a 15% mask rate over the entire frame. The Pareto front size was set to 12 sequences, with the number of MCTS iterations set to 128.

Human MUC1 transmembrane protein The MUC1 transmembrane glycoprotein is typically expressed on the apical surface of epithelial cells and is responsible for lubrication, protection, and cell signaling. In epithelial cancers, such as in the lung, pancreas, ovaries, breast, and colon, MUC1 is over-expressed and atypically glycosylated. Its characteristic profile in the context of cancers qualifies it as a tumor-associated antigen. We define a template with the entire wild-type CDS, which we partially optimize, and set lengths to 66 and 311 for the 5' and 3'UTRs, respectively. We performed simultaneous 5'/3' UTR design and partial codon optimization of the CDS, with uniform masking at a mask rate of 15% throughout the frame. The Pareto front

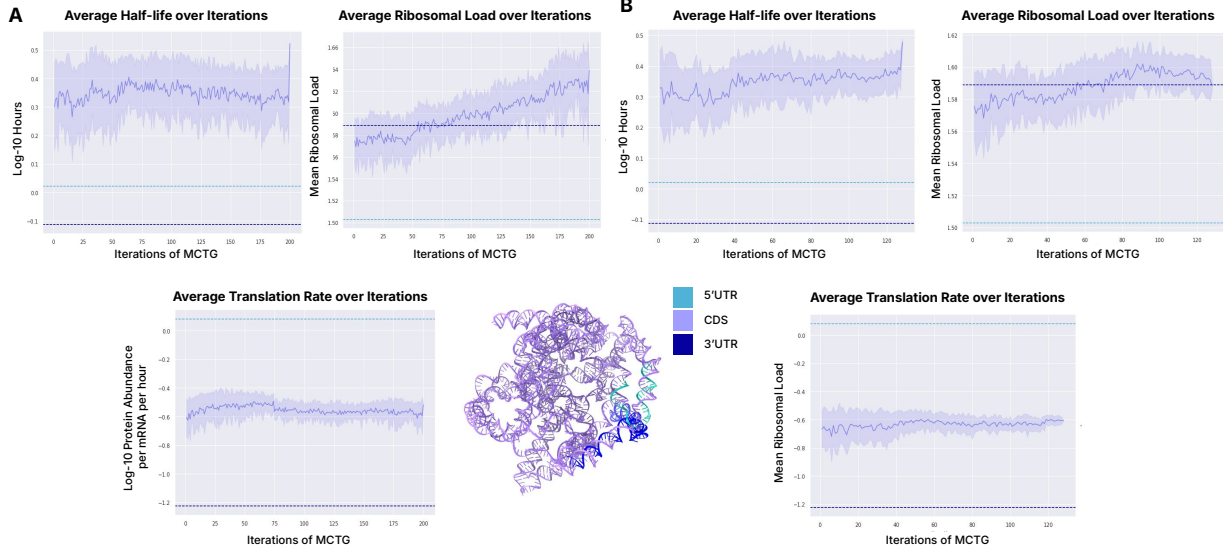


Figure 5. Multi-objective-guidance curves during *P. pyralis* luciferase (Fluc) mRNA design for half-life, ribosome load, and translation rate over iterations of Monte-Carlo Tree Guidance (MCTG). (A) The guidance curves for Fluc mRNA design, where both UTRs are fully designed and the CDS is partially masked at a 15% rate over the entire coding sequence. (B) The guidance curves for Fluc mRNA design, where both UTRs are fully designed and the CDS is partially masked at a 50% rate up until the 50th codon. The average scores are calculated from the unrolled children sequences generated at each iteration from the expanded leaf node sequence. Higher scores are therapeutically desirable for each property. Half-life and translation rate are in the log-10 scale. Both algorithms were run for 200 MCTG iterations, with 24 sequences in the Pareto front. Navy and teal lines correspond to the classifier dataset medians and the WT mRNA classifier scores, respectively. Error bars denote the standard error across all sequences in the Pareto front. Additionally, a predicted structure for one of the Fluc mRNAs is shown, provided by AlphaFold3 (Abramson et al., 2024).

size was set to 12 sequences, with the number of MCTS iterations set to 128.

Conditional Generation The mRNA transcripts designed by mRNAutilus demonstrate enhanced half-life, ribosomal load, and translation rate across all template transcripts through simultaneous codon optimization and UTR design (Table 1), except for MUC1 predicted mean ribosomal load. Furthermore, the average property values of the rolled-out sequences of the expanded child nodes show consistently increasing curves over iterations of MCTS across all properties, demonstrating the efficacy of our method on multi-property optimization. We demonstrate this observation for both Fluc mRNA generation strategies in Figure 5. Optimization curves for SARS-CoV-2 S-Protein and MUC1 can be found in Figure A3.

4. Discussion

The rational design of therapeutic mRNAs requires optimizing multiple competing properties, including stability, translation efficiency, and expression levels, a challenge that traditional single-objective approaches struggle to address systematically. Here, we introduce mRNAutilus, the first

multi-objective guided generative model for mRNA design that combines masked diffusion models with Monte Carlo Tree Search to navigate Pareto-optimal solutions for simultaneous codon optimization and UTR design. Our results demonstrate that unconditionally generated sequences capture key biological features, including appropriate Kozak consensus frequencies and GC content distributions. Importantly, our property classifiers achieve strong predictive performance (Spearman correlation of 0.641 for mean ribosomal load prediction), and MCTS-guided sampling generates diverse mRNA sequences with superior multi-objective performance compared to natural sequences while maintaining biological validity.

Several technical improvements could enhance the framework’s capabilities. Incorporating additional property classifiers for immunogenicity and tissue-specific expression would broaden therapeutic applications. Moreover, integrating manufacturing constraints into our guidance procedure, such as sequence complexity and local GC content, can reduce the need to manually alter sequences following generation. Finally, we envision that integration with experimental high-throughput screening, such as SELEX, will create powerful design-build-test-learn cycles.

Impact Statement

This work introduces the first multi-objective guided generative model for optimization and design of full therapeutic mRNA sequences, addressing a critical need in the rapidly expanding field of mRNA therapeutics. Our approach has the potential to accelerate the development of more effective mRNA therapeutics by enabling systematic optimization of multiple competing objectives simultaneously, potentially leading to improved vaccine efficacy, reduced development timelines, and expanded therapeutic applications. While the effectiveness of designed sequences ultimately requires experimental validation, and there exists potential for misuse as with any powerful design tool, our focus on therapeutic applications and the requirement for specialized expertise in mRNA manufacturing provide natural safeguards. Our work contributes to the responsible development of AI-driven biotechnology by providing transparent, scientifically rigorous methods that can enhance human health while maintaining appropriate oversight and validation requirements inherent to therapeutic development.

Acknowledgements

The authors declare that the reported research was supported by Atom Bioworks.

References

- Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630 (8016):493–500, 2024.
- Agarwal, V. and Kelley, D. R. The genetic and biochemical determinants of mRNA degradation rates in mammals. *Genome biology*, 23(1):245, 2022.
- Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and Van Den Berg, R. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.
- Barbier, A. J., Jiang, A. Y., Zhang, P., Wooster, R., and Anderson, D. G. The clinical progress of mRNA vaccines and immunotherapies. *Nature biotechnology*, 40(6):840–854, 2022.
- BioTechnologies, T. CleanCap® FLuc mRNA (5moU). <https://www.trilinkbiotech.com/cleancap-fluc-mrna-5mo.html>, n.d. Accessed: 2025-06-21.
- Boël, G., Letso, R., Neely, H., Price, W. N., Wong, K.-H., Su, M., Luff, J. D., Valecha, M., Everett, J. K., Acton, T. B., et al. Codon influence on protein expression in e. coli correlates with mRNA levels. *Nature*, 529(7586):358–363, 2016.
- Chen, T. and Guestrin, C. Xgboost: A scalable tree boosting system. In *International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016.
- Consortium, R. Rnacentral 2021: secondary structure integration, improved sequence search and new member databases. *Nucleic acids research*, 49(D1):D212–D220, 2021.
- Corbett, K. S., Seder, R. A., McMahan, K., and Morabito, K. M. Prefusion coronavirus spike proteins and their use. World Intellectual Property Organization (WIPO), August 2021. WO Patent WO2021159040A2.
- Dao, T. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*, 2024.
- Dyer, S. C., Austine-Orimoloye, O., Azov, A. G., Barba, M., Barnes, I., Barrera-Enriquez, V. P., Becker, A., Bennett, R., Beracochea, M., Berry, A., et al. Ensembl 2025. *Nucleic Acids Research*, 53(D1):D948–D957, 2025.
- Eichhorn, S. W., Guo, H., McGeary, S. E., Rodriguez-Mias, R. A., Shin, C., Baek, D., Hsu, S.-h., Ghoshal, K., Villén, J., and Bartel, D. P. mRNA destabilization is the dominant effect of mammalian microRNAs by the time substantial repression ensues. *Molecular cell*, 56(1):104–115, 2014.
- Fu, H., Liang, Y., Zhong, X., Pan, Z., Huang, L., Zhang, H., Xu, Y., Zhou, W., and Liu, Z. Codon optimization with deep learning to enhance protein expression. *Scientific reports*, 10(1):17617, 2020.
- Gebre, M. S., Rauch, S., Roth, N., Yu, J., Chandrashekar, A., Mercado, N. B., He, X., Liu, J., McMahan, K., Martinot, A., et al. Optimization of non-coding regions for a non-modified mRNA COVID-19 vaccine. *Nature*, 601(7893):410–414, 2022.
- Jackson, N. A., Kester, K. E., Casimiro, D., Gurunathan, S., and DeRosa, F. The promise of mRNA vaccines: a biotech and industrial perspective. *npj Vaccines*, 5(1):11, 2020.
- Kim, S. C., Sekhon, S. S., Shin, W.-R., Ahn, G., Cho, B.-K., Ahn, J.-Y., and Kim, Y.-H. Modifications of mRNA vaccine structural elements for improving mRNA stability and translation efficiency. *Molecular & cellular toxicology*, pp. 1–8, 2022.
- Kozak, M. An analysis of 5'-noncoding sequences from 699 vertebrate messenger RNAs. *Nucleic acids research*, 15(20):8125–8148, 1987.

- Kozak, M. An analysis of vertebrate mrna sequences: intimations of translational control. *The Journal of cell biology*, 115(4):887–903, 1991.
- Kozak, M. Initiation of translation in prokaryotes and eukaryotes. *Gene*, 234(2):187–208, 1999.
- Leppek, K., Das, R., and Barna, M. Functional 5 utr mrna structures in eukaryotic translation regulation and how to find them. *Nature reviews Molecular cell biology*, 19(3):158–174, 2018.
- Leppek, K., Byeon, G. W., Kladwang, W., Wayment-Steele, H. K., Kerr, C. H., Xu, A. F., Kim, D. S., Topkar, V. V., Choe, C., Rothschild, D., et al. Combinatorial optimization of mrna structure, stability, and translation for rna-based therapeutics. *Nature communications*, 13(1):1536, 2022.
- Li, S., Noroozizadeh, S., Moayedpour, S., Kogler-Anele, L., Xue, Z., Zheng, D., Montoya, F. U., Agarwal, V., Bar-Joseph, Z., and Jager, S. mrna-lm: full-length integrated slm for mrna analysis. *Nucleic Acids Research*, 53(3):gkaf044, 2025.
- Lin, J., Chen, Y., Zhang, Y., Lin, H., and Ouyang, Z. Deciphering the role of rna structure in translation efficiency. *BMC bioinformatics*, 23(Suppl 3):559, 2022a.
- Lin, Y.-H. and Bundschuh, R. Rna structure generates natural cooperativity between single-stranded rna binding proteins targeting 5 and 3 utrs. *Nucleic acids research*, 43(2):1160–1169, 2015.
- Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., dos Santos Costa, A., Fazel-Zarandi, M., Sercu, T., Candido, S., et al. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv*, 2022b.
- Lorenz, R., Bernhart, S. H., Höner zu Siederdissen, C., Tafer, H., Flamm, C., Stadler, P. F., and Hofacker, I. L. Viennarna package 2.0. *Algorithms for molecular biology*, 6:1–14, 2011.
- Mauro, V. P. and Chappell, S. A. A critical analysis of codon optimization in human therapeutics. *Trends in molecular medicine*, 20(11):604–613, 2014.
- Morrow, A., Jastrzebski, M., Wintermute, J., Zhao, S., Gardin, J., van Niekerk, L., Zulkower, V., Flynn, E., Moller, J., Cadena, P. Q., et al. mdd-0: mrna discrete diffusion for generation of stable mrna sequences.
- National Center for Biotechnology Information. SARS-CoV-2 isolate Wuhan-Hu-1, complete genome, 2020.
- Ng, W. C., Soto-Acosta, R., Bradrick, S. S., Garcia-Blanco, M. A., and Ooi, E. E. The 5 and 3 untranslated regions of the flaviviral genome. *Viruses*, 9(6):137, 2017.
- Ou, J., Nie, S., Xue, K., Zhu, F., Sun, J., Li, Z., and Li, C. Your absorbing discrete diffusion secretly models the conditional distributions of clean data. *arXiv preprint arXiv:2406.03736*, 2024.
- Pardi, N., Hogan, M. J., Porter, F. W., and Weissman, D. mrna vaccines—a new era in vaccinology. *Nature reviews Drug discovery*, 17(4):261–279, 2018.
- Pardi, N., Hogan, M. J., and Weissman, D. Recent advances in mrna vaccine technology. *Current opinion in immunology*, 65:14–20, 2020.
- Park, J. W., Lagniton, P. N., Liu, Y., and Xu, R.-H. mrna vaccines for covid-19: what, why and how. *International journal of biological sciences*, 17(6):1446, 2021.
- Patel, S., Peng, F. Z., Fraser, K., Friedman, A. D., Chatterjee, P., and Yao, S. Evoflow-rna: Generating and representing non-coding rna with a language model. *bioRxiv*, pp. 2025–02, 2025.
- Rosa, S. S., Nunes, D., Antunes, L., Prazeres, D. M., Marques, M. P., and Azevedo, A. M. Maximizing mrna vaccine production with bayesian optimization. *Biotechnology and Bioengineering*, 119(11):3127–3139, 2022.
- Ross, J. mrna stability in mammalian cells. *Microbiological reviews*, 59(3):423–450, 1995.
- Sahoo, S. S., Arriola, M., Schiff, Y., Gokaslan, A., Marroquin, E., Chiu, J. T., Rush, A., and Kuleshov, V. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 2024.
- Shazeer, N. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.
- Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. K. Simplified and generalized masked diffusion for discrete data. *Advances in Neural Information Processing Systems*, 2024.
- Soens, M., Ananworanich, J., Hicks, B., Lucas, K. J., Cardona, J., Sher, L., Livermore, G., Schaefer, K., Henry, C., Choi, A., et al. A phase 3 randomized safety and immunogenicity trial of mrna-1010 seasonal influenza vaccine in adults. *Vaccine*, 50:126847, 2025.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

- Tang, S., Zhang, Y., and Chatterjee, P. Peptune: De novo generation of therapeutic peptides with multi-objective-guided discrete diffusion. *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.
- Weber, J. S., Carlino, M. S., Khattak, A., Meniawy, T., Ansstas, G., Taylor, M. H., Kim, K. B., McKean, M., Long, G. V., Sullivan, R. J., et al. Individualised neoantigen therapy mrna-4157 (v940) plus pembrolizumab versus pembrolizumab monotherapy in resected melanoma (keynote-942): a randomised, phase 2b study. *The Lancet*, 403(10427):632–644, 2024.
- Wilson, E., Goswami, J., Baqui, A. H., Doreski, P. A., Perez-Marc, G., Zaman, K., Monroy, J., Duncan, C. J., Ujiie, M., Rămet, M., et al. Efficacy and safety of an mrna-based rsv pref vaccine in older adults. *New England Journal of Medicine*, 389(24):2233–2244, 2023.
- Wood, M., Klop, M., and Allard, M. Helix-mrna: A hybrid foundation model for full sequence mrna therapeutics. *arXiv preprint arXiv:2502.13785*, 2025.
- Xiao, J., Peng, B., Su, Z., Liu, A., Hu, Y., Nomura, C. T., Chen, S., and Wang, Q. Facilitating protein expression with portable 5-utr secondary structures in bacillus licheniformis. *ACS synthetic biology*, 9(5):1051–1058, 2020.
- Yazdani-Jahromi, M., Prakash, M., Mansi, T., Moskalev, A., and Liao, R. Helm: Hierarchical encoding for mrna language modeling. *arXiv preprint arXiv:2410.12459*, 2024.
- Zhang, C., Maruggi, G., Shan, H., and Li, J. Advances in mrna vaccines for infectious diseases. *Frontiers in immunology*, 10:594, 2019.
- Zhang, H., Zhang, L., Lin, A., Xu, C., Li, Z., Liu, K., Liu, B., Ma, X., Zhao, F., Jiang, H., et al. Algorithm for optimized mrna design improves stability and immunogenicity. *Nature*, 621(7978):396–403, 2023.
- Zhang, H., Gao, X., Zhang, J., and Lai, L. mrna2vec: mrna embedding with language model in the 5’utr-cds for mrna design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 1057–1065, 2025.
- Zheng, K., Chen, Y., Mao, H., Liu, M.-Y., Zhu, J., and Zhang, Q. Masked diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical sampling. *arXiv preprint arXiv:2409.02908*, 2024.

A. Appendix

Overview

In Appendix A.1, we provide additional implementation details including an analysis of the mRNA pertaining dataset (A.1.1) and the hyperparameters for the XGBoost regression classifiers (A.1.2). Appendix A.2 presents additional experimental results, including a principal component analysis of the latent embeddings learned by our unconditional MDM (A.2.1) and the guidance curves for human MUC1 and SARS-CoV-2-S-Protein mRNA generation (A.2.2). We also provide pseudo code for our multi-objective guidance algorithm in Appendix A.3. Finally, we compare our model to existing mRNA language models in Appendix A.4, highlighting the novel contributions of our work.

A.1. Implementation Details

A.1.1. MRNA PRETRAINING DATA

Ensembl pretraining data was sourced by querying the Perl API for all protein-coding vertebrate transcripts in the database. Species with metainfo and associated with the 'EnsemblVertebrates' division were considered, resulting in 342 total vertebrate species whose partial transcriptomes, as maintained by Ensembl, were retrieved. For every gene associated with a species, all protein-coding transcripts were collected in the form of a 5'UTR, CDS, and a 3'UTR. Pretraining data meta-analysis is shown in Figure A1.

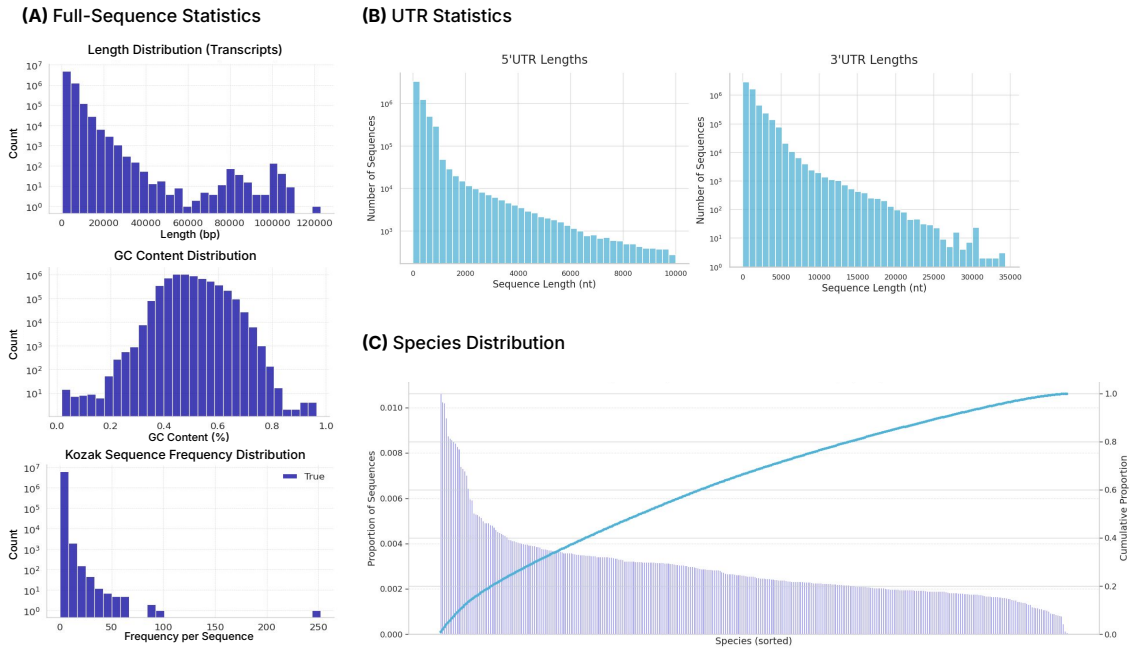


Figure A1. mRNA pretraining data. (A) Distributions for all 5.5 million protein-coding transcripts in the Ensembl dataset for sequence length (top-left), GC content (middle-left), and Kozak consensus sequence frequency (bottom-left). (B) Sequence length distribution distributions for the 5' (left) and 3'UTRs (right) of all sequences used during pretraining. (C) The proportion of protein-coding transcripts attributed to each of the 342 vertebrate species in the pretraining corpus.

A.1.2. REGRESSOR HYPERPARAMETERS

Hyperparameters for the XGBoost regressors are included in Table A1. All classifiers were trained on the data described in the 3.2 section, without any modifications to the sequences or pre-processing. Classifier training/validation sets were 80/20% splits done at random.

Table A1. XGBoost Hyperparameters for Regression

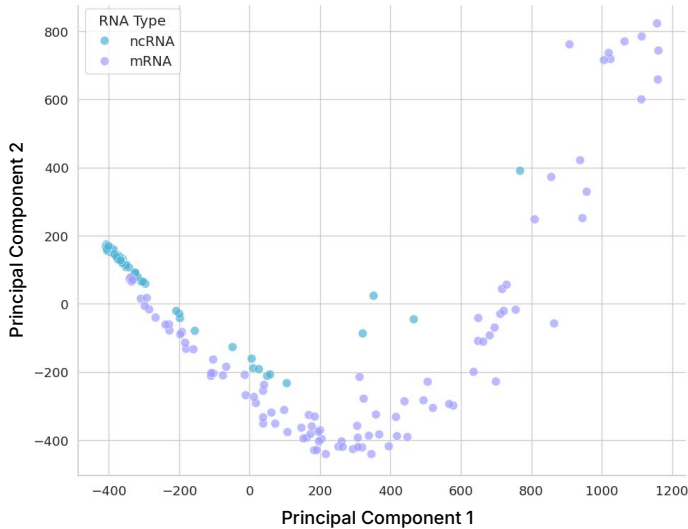
Hyperparameter	Value/Range
Objective	reg:squarederror
Lambda	[0.1, 10.0] (log scale)
Alpha	[0.1, 10.0] (log scale)
Gamma	[0, 5]
Colsample by Tree	[0.5, 1.0]
Subsample	[0.6, 0.9]
Learning Rate	[1e-5, 0.1]
Max Depth	[2, 30]
Min Child Weight	[1, 20]
Tree Method	hist
Scale Pos Weight	[0.5, 10.0] (log scale)

A.2. Additional Results

A.2.1. PCA ANALYSIS

To investigate the representations provided by mRNAutilus, we compared the model-produced embeddings for 100 mRNA sequences and 100 non-coding RNA (ncRNA) sequences. The mRNAs were sourced from the pretraining data validation set, and the non-coding RNAs were taken from RNACentral (Consortium, 2021). All 200 sequences were embedded and flattened, after which they were projected onto the first two principal components of the resulting two-dimensional matrix. Per Figure A2, the mRNAs are separable from the ncRNAs.

(A) PCA Projected Embeddings of RNA Sequences



(B) Explained Variance by PCA Components




Figure A2. **PCA Analysis of mRNAutilus representations.** (A) mRNAutilus embeddings are collected for 100 mRNAs and ncRNAs each, projected onto the two-dimensional vector space defined by the first two principal components of the flattened embeddings. (B) The associated scree plot for the PCA.

A.2.2. CONDITIONAL GUIDANCE

For the *SARS-CoV-2* S-Protein and human MUC1 mRNA generation, we produced 12 pareto-optimal mRNAs using 3 property-centric regressors over 128 MCTS iterations. Results for conditional generation are shown in Figure A3

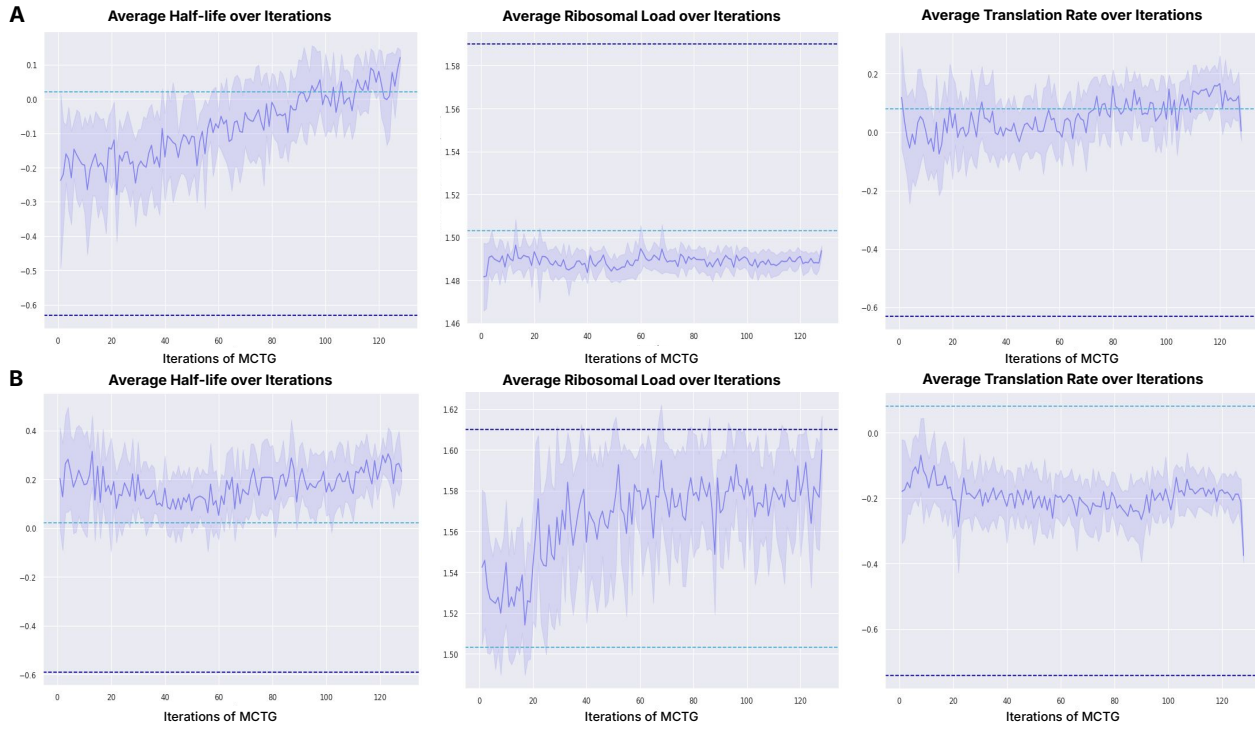


Figure A3. Conditional generation results for (A) Human MUC1 and (B) SARS-CoV-2 S-Protein. Conditional sequence generation spanned 128 iterations, with a Pareto front size set to 12. Navy and teal lines correspond to the classifier dataset medians and the WT mRNA classifier scores, respectively. Error bars denote the standard error across all sequences in the Pareto front.

A.3. Algorithm

Here, we provide the pseudocode for Monte-Carlo Tree Guidance (MCTG) for the mRNAutilus model.

Algorithm A1 mRNAutilus with Monte-Carlo Tree Guidance

```

0: Input: Denoising model  $p_\theta(x_1|x_t, t)$ , score function  $\mathbf{f}(x) : \mathcal{V}^L \rightarrow \mathbb{R}^K$ , number of iterations  $N_{\text{mcts}}$ , number of children  $N_{\text{child}}$ 
0: Output: Set of Pareto-optimal mRNA sequences  $\mathcal{P}^*$ 
0: Initialize:  $x_0 \leftarrow [M]^L$ ,  $\mathcal{P}^* \leftarrow \{\}$ ,  $t \leftarrow 1$ 
0: for  $i = 1$  to  $N_{\text{mcts}}$  do
0:   Selection:
0:    $\triangleright$  Traverse tree to expandable leaf
0:    $x_{\text{leaf}} \leftarrow \text{SELECT}(x_0)$ 
0:   Expansion:
0:   for  $j = 1$  to  $N_{\text{child}}$  do
0:     Sample  $x_1^{(j)} \sim p_\theta(\cdot|x_{\text{leaf}}, t)$  with Gumbel noise
0:     Create child  $x_{\text{child}}^{(j)}$  by unmasking  $k$  positions according to  $x_1^{(j)}$ 
0:     Add  $x_{\text{child}}^{(j)}$  to  $\text{children}(x_{\text{leaf}})$ 
0:   end for
0:   Rollout:
0:   for each  $x_{\text{child}}^{(j)} \in \text{children}(x_{\text{leaf}})$  do
0:      $\triangleright$  Complete sequence using self-planned ancestral sampling
0:      $\tilde{x}^{(j)} \leftarrow \text{SELFPLANNEDSAMPLING}(x_{\text{child}}^{(j)})$ 
0:      $\triangleright$  Evaluate properties using trained classifiers
0:      $\mathbf{s}^{(j)} \leftarrow \mathbf{f}(\tilde{x}^{(j)})$ 
0:     Update  $\mathcal{P}^*$  with  $\tilde{x}^{(j)}$  if non-dominated
0:      $\triangleright$  Compute reward based on Pareto dominance (Equation 12)
0:      $\mathbf{r}^{(j)} \leftarrow \text{COMPUTEREWARD}(\mathbf{s}^{(j)}, \mathcal{P}^*)$ 
0:   end for
0:   Backpropagation:
0:    $\triangleright$  Update ancestor node statistics with child rewards
0:    $\text{BACKPROPAGATE}(x_{\text{leaf}}, \{\mathbf{r}^{(j)}\})$ 
0: end for
0: return  $\mathcal{P}^*$ 
=0
    
```

A.4. Related Works

In this section, we provide a detailed description of three related works, highlighting their relevance to our work and the advancements driven by mRNAutilus.

mRNA-LM mRNA-LM is a language model that utilizes a contrastive loss to relate the latent spaces of three separate language models for accurate mRNA property prediction (Li et al., 2025). With each component language model pretrained on millions of mRNA segments sourced from NCBI, the contrastive loss enforces a learned joint representation for a full-length mRNA. Notably, this model shows accurate performance on the prediction of mRNA half-life, translation rate, and transcript expression from sequence only. The authors also highlight the model’s propensity for identifying mRNA fitness in vaccine-related applications.

In contrast, mRNAutilus is primarily a generative model that is also capable of function prediction. Using the same half-life and translation rate data in mRNA-LM, we demonstrate that mRNAutilus function prediction performance surpasses that of mRNA-LM, though we leverage XGBoost classifiers trained on model embeddings rather than an entirely neural network approach. Moreover, mRNAutilus is an individual model pretrained on whole mRNA sequences, dismissing the need for separating the mRNA sequence into smaller constituent units.

mDD-0 mDD-0 is a large language model trained on millions of protein-coding mRNA sequences sourced from Ensembl (Morrow et al.). Though implementation details are not completely provided, we describe our interpretation of their method here. The model features several component language models, each responsible for embedding a different portion of an mRNA sequence. These embeddings are combined and fed into a transformer which produces a unified set of mRNA logits. The overall model is trained on a masked diffusion modeling objective. Of note, they demonstrate their model’s ability to conditionally generate diverse, functionally plausible 3’UTRs for more stable mRNA design.

While primarily a generative model, mDD-0 relies on the concatenation of disjoint sequence representations. In contrast, mRNAutilus represents an mRNA sequence using a single embedding. Moreover, while not explicitly stated, mDD-0 relies on tokenization at the single-nucleotide level. Additionally, it represents codons as their corresponding amino acids embedded via ESM-2 (Lin et al., 2022b). This approach limits the model’s ability for tailored codon optimization. Additionally, our conditional generation approach features multi-objective guidance in parallel. In doing so, mRNAutilus can produce more well-rounded mRNA transcripts as opposed to mDD-0’s single property objective.

HELM HELM leverages hierarchical encoding for mRNA coding region representation, modulating a causal and masked language modeling loss based on codon hierarchy (Yazdani-Jahromi et al., 2024). Using this objective, they comprehensively compare performance across several model architecture types and tokenization strategies for mRNA fitness prediction. Using their models trained on a causal learning objective, the authors demonstrate the HELM method’s ability to unconditionally generate mRNA coding sequences with functional properties similar to those found in nature.

In comparison, mRNAutilus represents both the UTRs and coding sequences present in an mRNA transcript. Additionally, we leverage the programmability native to masked diffusion modeling for multi-objective conditional design of mRNA.