

MRNA-GPT: END-TO-END GENERATIVE DESIGN AND OPTIMIZATION OF FULL-LENGTH MRNA

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce mRNA-GPT, a generative model for end-to-end full-length mRNA sequence design and optimization. Unlike existing approaches that optimize isolated regions, mRNA-GPT jointly optimizes across all three regions (5' UTR, CDS, and 3' UTR) to capture cross-region regulatory interactions critical for therapeutic efficacy. Pre-trained on 10 million full-length natural mRNA sequences across diverse species and organisms, establishing a robust foundation for sequence generation. mRNA-GPT employs an iterative optimization framework with oracle-based rewards to progressively enhance target properties including translation efficiency and half-life. mRNA-GPT supports flexible generation modes: single regions, full-length sequences, or conditional generation given any region. Empirical results demonstrate superior performance over state-of-the-art methods: CDS optimization achieves higher predicted translation rates than LinearDesign and GEMORNA while maintaining structural stability, and full-length design captures critical cross-region interactions yielding enhanced translation efficiency. This unified approach establishes mRNA-GPT as a versatile platform for rational mRNA therapeutics design.

1 INTRODUCTION

Messenger RNA (mRNA) medicines have transitioned from conceptual promise to clinical reality with the success of COVID-19 vaccines (Polack et al., 2020; Baden et al., 2021). However, therapeutic performance remains sensitive to sequence-level determinants distributed across the entire transcript. Coding sequences (CDS) dictate translation kinetics and *in vivo* stability through codon usage and decoding dynamics (Presnyak et al., 2015; Wu et al., 2019); the 5' untranslated region (UTR) modulates ribosomal initiation via structural motifs and upstream open reading frames (Leppek et al., 2018; Karollus et al., 2021); and the 3' UTR governs stability, localization, and translational efficiency through microRNA and RNA-binding protein interactions (Mayr, 2017).

Recent evidence strongly supports joint optimization across all three regions. UTR interactions can reduce protein output by 2-fold (Medjmedj et al., 2025), while UTR combinations that perform well in isolation are incompatible when paired (Ye et al., 2023). Notably, identical UTR combinations produced dramatically different outcomes depending on the CDS (Medjmedj et al., 2025). Dinucleotide interactions at the 5' UTR-CDS boundary drive up to 24.8% differences in translation efficiency (Noderer et al., 2014). These findings demonstrate that full-length mRNA design requires co-optimization across the combinatorial design space.

Recent machine learning advances have accelerated mRNA design optimization. LinearDesign (Zhang et al., 2023) jointly optimize CDS structure and codon usage via dynamic programming. Deep generative models like GEMORNA (Zhang et al., 2025) employ separate transformers for generating each independently, which are then combined to create full-length optimized mRNA sequences, while iDRO (Gong et al., 2023) uses BiLSTM-CRF for CDS optimization coupled with RNA-Bart (a bidirectional auto-regressive transformer) for UTR generation. However, none of these approaches jointly optimize all three regions while capturing long-range cross-region interactions critical to biological function.

To address these limitations, we propose mRNA-GPT, a generative model for *de novo* full-length mRNA sequence design and optimization. The model was pre-trained on 10 million full-length mRNA sequences with annotated regions to capture cross-region interactions (5' UTR, CDS, and

3' UTR) often missed by modular approaches. We employ iterative optimization framework that directly optimize the mRNA-GPT model for enhanced target properties, with flexibility to incorporate any task-specific reward model. In this work, mRNA-GPT demonstrates superior performance over existing methods across multiple design objectives. Through iterative optimization, mRNA-GPT generates CDS sequences with significantly higher predicted translation rates than LinearDesign and GEMORNA while maintaining balanced codon adaptation and structural stability. For full-length sequences, our end-to-end approach captures critical cross-region interactions that modular methods miss, yielding enhanced translation efficiency and substantially higher sequence diversity. These improvements stem from joint optimization across all three mRNA regions rather than treating them as independent modules.

2 METHODS

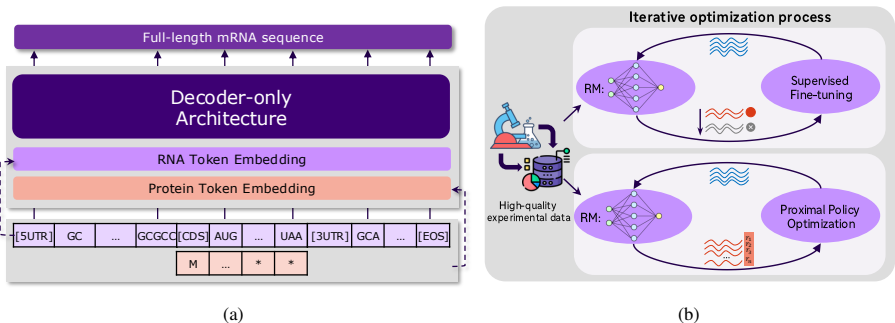


Figure 1: **mRNA-GPT model overview.** (a) **Decoder-only architecture processes full-length mRNA sequences** with region-specific tokens and dual embedding layers for RNA and protein sequences. (b) **Iterative optimization framework** employs reward model scoring followed by either supervised fine-tuning or PPO-based policy updates for progressive property enhancement.

We propose mRNA-GPT, a decoder-only architecture for full-length mRNA generation. The pre-training dataset consists of 10 million natural mRNA sequences collected from RefSeq (Li et al., 2024; 2025). Each sequence contains three regions: 5' UTR, CDS, and 3' UTR. As illustrated in Fig. 1a, the UTR regions are tokenized into segments using pre-trained tokenizers, while coding regions are split into sequential codons. Special tokens [5UTR], [CDS], and [3UTR] are added at the beginning of each region to indicate functional areas, and the special token [EOS] is added after the last region.

Along the coding region, target protein sequences are incorporated to guide design, ensuring that generated CDS sequences can be translated into the target protein. The protein sequences consist of amino acids that are shifted left by one position to align the next amino acid with the generation of the corresponding codon. The architecture employs two input embedding layers: for coding regions, codon embeddings and amino acid embeddings are concatenated and fed into the model, while for UTR regions, special tokens [5UTR] and [3UTR] are encoded as the amino acid and serve as regional indicators. We trained data-driven tokenizers separately for 5' and 3' UTR regions using the SentencePiece framework (Kudo & Richardson, 2018) with a vocabulary size 8000. Compared to fixed-length n-gram tokenizers, pre-trained tokenizers learn variable-length tokens that better capture meaningful regulatory motifs (See Appendix A.1 for details). This compression enables the model to process longer RNA sequences within fixed context window limits.

To optimize target properties, such as mRNA half-life, translation rate, or protein output, we applied an iterative optimization process using two approaches (Fig. 1b): supervised fine-tuning (SFT) or Proximal Policy Optimization (PPO). Both approaches incorporate a static reward model trained on experimental datasets to predict mRNA property values. The two approaches differ in their data usage and the optimization objectives. In the SFT approach, sequences are scored and ranked by the reward model, with only the top-performing sequences retained for fine-tuning using next-token prediction loss. Through this process, the model learns the statistical patterns of high-performing sequences and subsequently generates optimized sequences with improved property values. The

PPO strategy instead utilizes all generated sequences, updating the policy model conservatively while using a frozen reference model to constrain updates and maintain output quality. PPO employs a clipped surrogate objective with KL divergence penalty to prevent excessively large policy shifts, balancing reward maximization with stable optimization (see Appendix A.2).

3 IN SILICO EXPERIMENTS AND RESULTS

3.1 ITERATIVE CDS OPTIMIZATION ENHANCES TRANSLATION RATES




Figure 2: **CDS optimization performance for SARS-CoV-2 membrane protein.** Comparison of synonymous CDS sequences generated by mRNA-GPT over iterations against LinearDesign and GEMORNA showing: (a) predicted translation rates improving with iterations, (b) CAI progression toward human-preferred codons, and (c) MFE values indicating RNA structural stability.

We applied the iterative optimization framework with supervised fine-tuning (SFT) to coding sequence design, training a conditional generator $p(\text{CDS} \mid \text{protein})$ to produce CDS with progressively higher predicted translation rates, as scored by the reward model, mRNA-LM, which is an mRNA language model fine-tuned on experimentally measured human transcripts (Li et al., 2025). We optimized the CDS regions of mRNA sequences corresponding to four target proteins, as listed in Table B.1. At iteration t (with $t=0$ representing the mRNA-GPT pre-trained model), we used a fixed orchestration pipeline to perform a closed-loop sample–score–select–fine-tune procedure independently for each protein. Starting from a common base checkpoint at $t=0$, the generator produces 6000 candidate CDS sequences per iteration for each protein. All candidates were scored by the mRNA-LM, generating a ranked list per protein and iteration. The top $N=1500$ highest-scoring CDS sequences were used as the fine-tuning dataset for that protein at the given iteration.

We focus on a representative target, the SARS-CoV-2 membrane protein (Protein 3 in table B.1), with results for the remaining three proteins provided in the Appendix B. For benchmarking, we included both a classical algorithm (LinearDesign) and the state-of-the-art language model-based GEMORNA as baselines. GEMORNA-CDS, an encoder-decoder module specifically designed for CDS design, treats protein sequences as the source and mRNA CDS as the target. We used GEMORNA-CDS to generate 200 unique CDS sequences for each protein. For LinearDesign, we systematically evaluated a range of λ values (from 0 to 100, increments by 0.1) to balance codon adaptation and structural stability. After deduplication, this yielded 37 unique CDS sequences.

Fig. 2a presents the distribution of predicted translation rates for the eleven iterations, alongside with the LinearDesign and GEMORNA-CDS baselines. As optimization proceeds from $t=0$ to $t=10$, the predicted translation rate distribution shifts upward, with the median approximately doubling and low-scoring outliers nearly eliminated. The final iterations surpass both baselines. The GEMORNA-CDS distribution lies below mRNA-GPT outputs from rounds $t=7-10$ in both median and upper quartiles. LinearDesign performs better than GEMORNA-CDS, but remains slightly below mRNA-GPT at round $t=9-10$, whose median and 75th percentile are higher. Note that the GEMORNA-CDS model uses zero-shot generation, thus its sequences provides a strong starting point but lack further optimization. In contrast, mRNA-GPT’s iterative oracle-guided optimization yields continued performance improvements, surpassing both baselines and demonstrating superior capacity for translation rate enhancement.

Fig. 2b reports the codon adaptation index (CAI) distributions, computed from a fixed human codon-usage reference table Zhang et al. (2023), following the classical definition (see Equation 10) (Sharp & Li, 1987). Over ten optimization rounds, the CAI distribution shifts from ~ 0.70 to ~ 0.87 , indicating improved adaptation to human codon usage. Both LinearDesign and GEMORNA achieve

even higher CAI values than $t=10$ mRNA-GPT, with medians in the $\sim 0.9-0.95$ range. However, this higher CAI does not yield superior predicted translation rates (Fig. 2a), in line with the view that codon optimality contributes to, but does not solely determine, translation efficiency and mRNA stability (Presnyak et al., 2015; Wu et al., 2019). This suggests that mRNA-GPT converges to a more balanced region of sequence space.

We next examine global RNA secondary structure using ViennaRNA (Lorenz et al., 2011) to compute the Minimum Free Energy (MFE) for the designed CDS sequences. Fig. 2c plots the mean MFE (kcal/mol) and its standard deviation over generated sequences. The mean MFE decreases from around -200 to -250 kcal/mol after three iterations, then plateaus in later iterations, achieving improved stability without extreme over-folding. LinearDesign produces much more stable structures, reflecting its explicit energy-minimization objective (Zhang et al., 2023). Combined with the translation-rate results, this indicates a compromise between stability and accessibility: mRNA-GPT transcripts are more structured than the pre-trained baseline yet less over-stabilized than energy-minimizing designs, while achieving higher predicted values.

Nucleotide composition analysis shows that optimization progressively enriches C and G while depleting A and U, matching LinearDesign and GEMORNA (Fig B.2a). Despite this convergence, mRNA-GPT maintains substantially higher diversity than both baselines, as measured by mean pairwise Hamming distance (Fig B.2b). The t-SNE visualization reveals that later iterations drift smoothly across a broad manifold rather than collapsing onto a single codon pattern, indicating that optimization reweighs rather than collapses codon-usage modes.

Analogous analyses for Rabies virus G, Ebola GP, and human DCT/TYRP2 are provided in Appendix B. Qualitatively, they show the same hallmarks as SARS-CoV-2 membrane protein: batch shifts toward higher oracle scores, improved CAI and GC content, moderate compaction but no catastrophic loss of diversity, and a balance between increased structural stability and avoidance of extreme over-folding. We also applied mRNA-GPT to optimize 3' UTR sequences for enhanced stability using the Saluki model as reward model (Agarwal & Kelley, 2022) (see Appendix B). Fine-tuned mRNA-GPT outperformed GEMORNA-UTR and genetic algorithm in predicted half-life, demonstrating the generalization of iterative optimization approach across different mRNA regions.

3.2 END-TO-END FULL-LENGTH MRNA DESIGN OUTPERFORMS REGIONAL OPTIMIZATION

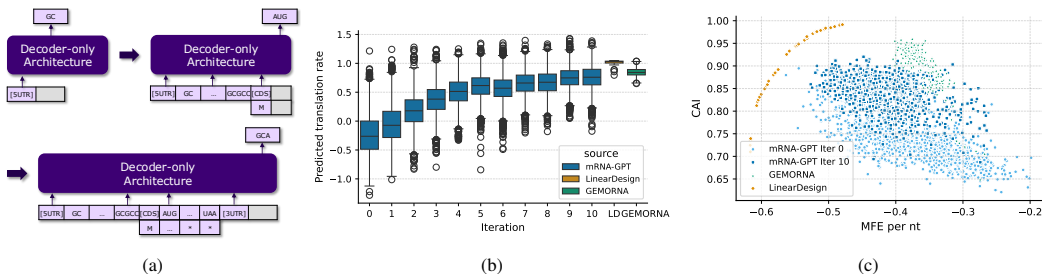


Figure 3: **End-to-end Optimization of Full-length mRNA sequences.** (a) Stepwise autoregressive generation workflow from 5' UTR through CDS to 3' UTR with protein-guided synthesis. (b) Translation efficiency improvements across optimization iterations compared to GEMORNA and LinearDesign. (c) Trade-off between CAI versus MFE across different design methods, showing mRNA-GPT's broader exploration of the sequence design.

Building upon our evaluation of mRNA-GPT's regional optimization capabilities, we next assessed the model's ability to generate complete full-length mRNA sequences. We employed a stepwise auto-regressive approach, as illustrated in Fig. 3a, to design mRNA sequences from the 5' to 3' end. The generation process proceeds in three sequential stages. First, the model generate the 5' UTR by prompting it with the special token [5UTR], feeding each output token back as input, until it encounters a special termination token. Second, we prompt the model to generate the CDS by providing the previously generated 5' UTR as context, extended with the special token [CDS], along with the target protein sequence. Finally, the 5' UTR and CDS are concatenated and extended with the [3UTR] special token. The model then generates the 3' UTR auto-regressively, until reaching the end-of-sequences token ([EOS]).

216 Following the CDS design and optimization, we also utilized mRNA-LM as an oracle to predict
217 translation efficiency given a complete mRNA sequence (Li et al., 2025). We applied iterative
218 supervised fine-tuning using mRNA-GPT to optimize sequences encoding the SARS-CoV-2 mem-
219 brane protein. Fig. 3b demonstrates progressive improvement in predicted translation efficiency
220 over 10 optimization iterations. The average predicted translation efficiency increased from -0.24
221 to 0.76, with convergence observed after approximately five iterations. Mann-Whitney-Wilcoxon
222 tests (Mann & Whitney, 1947) confirmed that each iteration (except iteration 6) yielded sequences
223 with statistically significantly higher translation efficiency than the previous iteration ($p < 0.05$).

224 For GEMORNA, we combined top-scoring UTRs with 200 CDS sequences, yielding 20,000 full-
225 length sequences. The predicted translation rate distribution for GEMORNA ranged from 0.65
226 to 1.04 with an average value of 0.85. Mann-Whitney-Wilcoxon tests revealed that the top 10%
227 sequences from mRNA-GPT iteration 7 to 10 exhibited significantly higher translation efficiency
228 than the top 10% of GEMORNA sequences. GEMORNA builds separate models for UTR and CDS,
229 thereby neglecting interactions across regions. For LinearDesign, the 37 unique CDS sequences
230 were paired with the same UTR sequences used in the LinearDesign publication. The resulting
231 predicted translation rate distribution ranged from 0.79 to 1.05 with an average value of 1.01. In
232 the final iteration (10th), the top 10% of mRNA-GPT sequences demonstrated significantly higher
233 translation efficiency compared to the top 10% of LinearDesign sequences.

234 We further computed the CAI value (only the CDS sequences) and MFE (of full sequences) and
235 visualized the relationship in Fig. 3c. Since mRNA-GPT and GEMORNA generate UTR sequences
236 with variable lengths. We normalized the MFE by the number of nucleotides in the sequence. The
237 strong negative correlation between CAI and MFE per *nt* with a Pearson correlation coefficient
238 value of -0.7 agrees with the fundamental trade-off between MFE and CAI (Zhang et al., 2023).
239 Sequences designed by LinearDesign draw the boundary of searching space because the algorithm
240 directly optimizes CAI and MFE simultaneously. The sequences from the optimized mRNA-GPT
241 shift towards higher CAI and lower MFE compared to the pre-trained mRNA-GPT model indicates
242 that the optimization is in the right direction. The broader distribution of CAI and MFE suggests
243 that mRNA-GPT enables exploring the large space under the boundary curve built by LinearDesign.
244 While sequences from GEMORNA just explore a relatively small space with higher CAI but also
245 higher MFE, indicating the GEMORNA model implicitly focuses more on CAI than MFE, and lacks
246 capturing the interaction across regions.

246 Full-length mRNA sequences consistently yield lower predicted translation rates compared to the
247 CDS region alone, as shown in Fig. 3b and Fig. 2a. This reflects the critical influence of UTR
248 regions on overall translation rate. UTR regions impact ribosome recruitment, initiation, elongation,
249 and transcript stability, often dampening the translation rate relative to optimally designed CDS-
250 only constructs. This context-dependent effect underscores the need for comprehensive design.
251 Considering full-length mRNA, rather than isolated CDS, is essential for realistic and biologically
252 relevant predictions and performance in mRNA therapeutics and synthetic biology applications.

253 254 255 4 CONCLUSION

256
257
258 mRNA-GPT addresses a fundamental limitation in current mRNA design: the inability to jointly
259 optimize all functional regions while capturing cross-region interactions critical for therapeutic effi-
260 cacy. Pre-trained on 10 million full-length natural mRNA sequences, our model learns complex de-
261 pendencies between 5' UTR, CDS, and 3' UTR that modular approaches inherently miss. Through
262 iterative optimization via SFT and PPO with oracle-based rewards, mRNA-GPT directly maximizes
263 target properties such as translation efficiency and stability. Empirical results demonstrate superior
264 performance over state-of-the-art methods: CDS optimization achieves higher predicted translation
265 rates than both LinearDesign and GEMORNA while maintaining balanced codon adaptation and
266 structural stability. Full-length sequence design confirms that optimal UTR-CDS combinations are
267 context-dependent, validating the necessity of end-to-end optimization. The model's flexible gener-
268 ation modes establish mRNA-GPT as a versatile platform for rational mRNA design. This unified
269 approach represents a paradigm shift in nucleic acid therapeutics that could accelerate development
across protein replacement therapies, antibody production, and gene editing applications.

REFERENCES

- 270
271
272 Vikram Agarwal and David R Kelley. The genetic and biochemical determinants of mRNA degra-
273 dation rates in mammals. *Genome biology*, 23(1):245, 2022.
- 274
275 Lindsey R. Baden et al. Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *New England*
276 *Journal of Medicine*, 384(5):403–416, 2021.
- 277
278 Sina Barazandeh, Furkan Ozden, Ahmet Hincer, Urartu Ozgur Safak Seker, and A Ercument Cicek.
279 UTRGAN: learning to generate 5’ UTR sequences for optimized translation efficiency and gene
expression. *Bioinformatics Advances*, 5(1):vbaf134, 2025.
- 280
281 Yanyi Chu, Dan Yu, Yupeng Li, Kaixuan Huang, Yue Shen, Le Cong, Jason Zhang, and Mengdi
282 Wang. A 5’ UTR language model for decoding untranslated regions of mRNA and function
283 predictions. *Nature Machine Intelligence*, 6(4):449–460, 2024.
- 284
285 Peter JA Cock, Tiago Antao, Jeffrey T Chang, Brad A Chapman, Cymon J Cox, Andrew Dalke,
286 Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, et al. Biopython: freely
287 available python tools for computational molecular biology and bioinformatics. *Bioinformatics*,
25(11):1422–1423, 2009.
- 288
289 Haoran Gong, Jianguo Wen, Ruihan Luo, Yuzhou Feng, JingJing Guo, Hongguang Fu, and Xiaobo
290 Zhou. Integrated mRNA sequence optimization using deep learning. *Briefings in bioinformatics*,
291 24(1):bbad001, 2023.
- 292
293 Kaixuan Huang, Yukang Yang, Kaidi Fu, Yanyi Chu, Le Cong, and Mengdi Wang. Latent diffusion
models for controllable rna sequence generation. *arXiv preprint arXiv:2409.09828*, 2024.
- 294
295 A. Karollus et al. Sequence determinants of human gene regulatory elements inferred from mas-
296 sively parallel reporter assays. *Nature Genetics*, 53(3):278–288, 2021.
- 297
298 Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword
tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018.
- 299
300 Kathrin Leppek, Rohit Das, and Maria Barna. Functional 5’ UTR mRNA structures in eukaryotic
301 translation regulation and how to find them. *Nature Reviews Molecular Cell Biology*, 19:158–174,
302 2018.
- 303
304 Sizhen Li, Saeed Moayedpour, Ruijiang Li, Michael Bailey, Saleh Riahi, Lorenzo Kogler-Anele,
305 Milad Miladi, Jacob Miner, Fabien Pertuy, Dinghai Zheng, et al. CodonBERT large language
model for mRNA vaccines. *Genome research*, 34(7):1027–1035, 2024.
- 306
307 Sizhen Li, Shahriar Noroozizadeh, Saeed Moayedpour, Lorenzo Kogler-Anele, Zexin Xue, Dinghai
308 Zheng, Fernando Ulloa Montoya, Vikram Agarwal, Ziv Bar-Joseph, and Sven Jager. mRNA-LM:
309 full-length integrated SLM for mRNA analysis. *Nucleic Acids Research*, 53(3):gkaf044, 2025.
- 310
311 Adam J Litterman, Robin Kageyama, Olivier Le Tonqueze, Wenxue Zhao, John D Gagnon, Hani
312 Goodarzi, David J Erle, and K Mark Ansel. A massively parallel 3’ UTR reporter assay re-
313 veals relationships between nucleotide content, sequence conservation, and mRNA destabiliza-
tion. *Genome research*, 29(6):896–906, 2019.
- 314
315 Ronny Lorenz, Stephan H Bernhart, Christian Höner zu Siederdisen, Hakim Tafer, Christoph
316 Flamm, Peter F Stadler, and Ivo L Hofacker. ViennaRNA Package 2.0. *Algorithms for molecular*
biology, 6(1):26, 2011.
- 317
318 Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochas-
319 tically larger than the other. *The annals of mathematical statistics*, pp. 50–60, 1947.
- 320
321 Christine Mayr. Regulation by 3’-untranslated regions. *Annual Review of Genetics*, 51:171–194,
322 2017.
- 323
Christine Mayr. What are 3’ UTRs doing? *Cold Spring Harbor perspectives in biology*, 11(10):
a034728, 2019.

- 324 Ayoub Medjmedj, Hugo Genon, Dounia Hezili, Albert Ngalle Loth, Rudy Clemençon, Cyril
325 Guimpied, Lucile Mollet, Anne Bigot, Frank Wien, Josef Hamacek, et al. Evaluation of syn-
326 thetic mRNA with selected UTR sequences and alternative poly (A) tail, in vitro and in vivo.
327 *Molecular Therapy Nucleic Acids*, 36(3), 2025.
- 328
329 Melanie Mitchell. *An introduction to genetic algorithms*. MIT press, 1998.
- 330 William L Noderer, Ross J Flockhart, Aparna Bhaduri, Alexander J Diaz de Arce, Jiajing Zhang,
331 Paul A Khavari, and Clifford L Wang. Quantitative analysis of mammalian translation initiation
332 sites by FACS-seq. *Molecular systems biology*, 10(8):748, 2014.
- 333
334 Sawan Patel, Sophia Tang, Yinuo Zhang, Pranam Chatterjee, and Sherwood Yao. Multi-objective-
335 guided generative design of mRNA with therapeutic properties. In *ICML 2025 Generative AI and*
336 *Biology (GenBio) Workshop*, 2025.
- 337 Fernando P. Polack et al. Safety and efficacy of the BNT162b2 mRNA COVID-19 vaccine. *New*
338 *England Journal of Medicine*, 383(27):2603–2615, 2020.
- 339
340 Veronica Presnyak et al. Codon optimality is a major determinant of mRNA stability. *Cell*, 160(6):
341 1111–1124, 2015.
- 342 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language
343 models are unsupervised multitask learners. *OpenAI Blog*, 1(8), 2019.
- 344
345 Paul M. Sharp and Wen-Hsiung Li. The codon adaptation index – a measure of directional synony-
346 mous codon usage bias, and its potential applications. *Nucleic Acids Research*, 15(3):1281–1295,
347 1987.
- 348 Xiaoshan Tang, Miaozhe Huo, Yuting Chen, Hai Huang, Shugang Qin, Jiaqi Luo, Zeyi Qin, Xin
349 Jiang, Yongmei Liu, Xing Duan, et al. A novel deep generative model for mRNA vaccine devel-
350 opment: Designing 5' UTRs with n1-methyl-pseudouridine modification. *Acta Pharmaceutica*
351 *Sinica B*, 14(4):1814–1826, 2024.
- 352
353 Q. Wu, S.G. Medina, G. Kushawah, M.L. DeVore, L.A. Castellano, J.M. Hand, M. Wright, and A.A.
354 Bazzini. Translation affects mRNA stability in a codon-dependent manner in human cells. *eLife*,
355 8:e45396, 2019.
- 356 Zhongfeng Ye, Srinivasa Reddy Bonam, Lindsay GA McKay, Jessica A Plante, Jordyn Walker,
357 Yu Zhao, Changfeng Huang, Jinjin Chen, Chutian Xu, Yamin Li, et al. Monovalent SARS-
358 COV-2 mRNA vaccine using optimal UTRs and Inps is highly immunogenic and broadly pro-
359 tective against omicron variants. *Proceedings of the National Academy of Sciences*, 120(52):
360 e2311752120, 2023.
- 361
362 H. Zhang, H. Liu, Y. Xu, et al. Deep generative models generate mRNA sequences with enhanced
363 translation capacity and stability. *bioRxiv*, pp. 2024–06, 2024.
- 364
365 He Zhang, Liang Zhang, Ang Lin, Congcong Xu, Ziyu Li, Kaibo Liu, Boxiang Liu, Xiaopin Ma,
366 Fanfan Zhao, Huiling Jiang, et al. Algorithm for optimized mRNA design improves stability and
367 immunogenicity. *Nature*, 621(7978):396–403, 2023.
- 368
369 He Zhang, Hailong Liu, Yushan Xu, Haoran Huang, Yiming Liu, Jia Wang, Yan Qin, Haiyan Wang,
370 Lili Ma, Zhiyuan Xun, et al. Deep generative models design mRNA sequences with enhanced
371 translational capacity and stability. *Science*, pp. eadr8470, 2025.
- 372
373
374
375
376
377

Appendices

A METHODS

The core architecture utilizes a GPT-2 model with a maximum length of 1024 tokens, 16 layers, 16 attention heads, and 768-dimensional token embeddings (Radford et al., 2019).

A.1 UTR TOKENIZATION DETAILS

Unlike simple nucleotide-level or fixed n-gram tokenizer for UTR regions, we trained data-driven tokenizers separately for 5' UTR and 3' UTR. Each tokenizer was pre-trained on the same corpus used for mRNA-GPT (10 million sequences) with a fixed vocabulary size of 8000 subword units using the SentencePiece framework (Kudo & Richardson, 2018). Tokenization employed a subword model based on byte-pair encoding (BPE). On average, the 5' UTR is approximately 200 *nt* and the 3' UTR is approximately 800 *nt* in length. The use of pre-trained subword tokenizers significantly reduces the number of tokens representing these sequences, shrinking hundreds of nucleotide tokens to fewer than a hundred subword tokens. This compression enables models to process longer RNA spans within fixed context window limits, thus benefiting performance on full-length mRNA tasks. Compared to fixed-length n-gram tokenizers, pre-trained tokenizers learn variable-length tokens according to sequence frequency; these tokens can better capture meaningful regulatory motifs, whereas rigid n-gram tokenization may break motifs and result in less biologically relevant segmentation.

A.2 PRE-TRAINING DATA ORDER SHUFFLING.

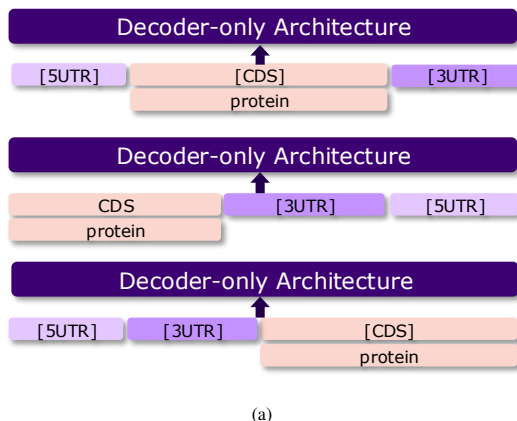


Figure A.1: Pre-training permutes 5' UTR, CDS, and 3' UTR order.

As shown in Fig A.1a, during pre-training, we randomly shuffle the order of the three regions to enhance scalability and enable diverse downstream tasks. Three possible orders are employed: 5' UTR- CDS- 3' UTR, CDS- 3' UTR- 5' UTR, and 5' UTR- 3' UTR- CDS. Such flexible training strategy enables conditional generation capabilities, allowing the model to generate any region conditioned on other regions. For instance, beyond following the natural 5' to 3' order, the model can generate 5' UTR given a 3' UTR, CDS given both UTR regions, or UTR regions based on CDS. Besides the customized two input layers, pre-trained tokenizer, and shuffled order of regions, the core architecture utilizes a GPT-2 model with a maximum length of 1024 tokens, 16 layers, 16 attention heads, and 768-dimensional token embeddings.

432 A.3 MATHEMATICAL FORMULATIONS AND OPTIMIZATION DETAILS

433 A.3.1 SUPERVISED FINE-TUNING LOSS FUNCTION

434 For a sequence of length N with tokens x_1, x_2, \dots, x_N , where the model predicts the probability
435 $p_\theta(x_i|x_{<i})$ of token x_i given its left context, averaged over sequence length (per-token loss):

$$436 \mathcal{L}_{\text{avg}} = -\frac{1}{N} \sum_{i=1}^N \log p_\theta(x_i|x_{<i}) \quad (1)$$

437 where θ represents the set of all trainable model parameters.

438 A.3.2 PPO MATHEMATICAL FRAMEWORK

439 PPO updates the policy model conservatively: smaller parameter changes per update empirically
440 lead to better convergence and stability. To prevent excessively large policy shift, PPO computes a
441 ratio between the current and the reference policy probabilities:

$$442 r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\text{ref}}(a_t|s_t)} \quad (2)$$

443 where $\pi_\theta(a_t|s_t)$ and $\pi_{\text{ref}}(a_t|s_t)$ are the probabilities of taking action a_t in state s_t under the current
444 policy (π_θ) and reference models (π_{ref}), respectively. Here, π_{ref} denotes a frozen reference policy
445 (e.g., the initial supervised model or the policy from the previous iteration), which remains fixed
446 while π_θ is updated. If $r_t(\theta) > 1$, the action a_t is more likely under the current policy than under
447 the reference, quantifying the policy shift for that action. PPO uses the following clipped surrogate
448 objective to ensure conservative update:

$$449 L^{\text{CLIP}}(\theta) = \hat{\mathbb{E}}_t \left[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t) \right] \quad (3)$$

450 where \hat{A}_t is the advantage function, representing how much better an action is compared to the
451 expected value for a given state. In practice, advantage values are often computed from normalized
452 rewards with a baseline value. The "min" operator provides a pessimistic (lower bound) estimate: it
453 takes either the clipped or unclipped objective depending on whether the update would exceed the
454 allowed ϵ -range (typically $\epsilon = 0.2$). In RL PPO, an additional Kullback-Leibler (KL) divergency
455 penalty between the current and reference model distributions is computed:

$$456 L^{\text{KL}}(\theta) = \text{KL} [\pi_{\text{ref}}(\cdot|s_t), \pi_\theta(\cdot|s_t)] \quad (4)$$

457 The complete loss is

$$458 \mathcal{L} = -L^{\text{CLIP}}(\theta) + \beta \cdot L^{\text{KL}}(\theta) \quad (5)$$

459 where β is the KL penalty coefficient. Training minimizes \mathcal{L} , which is equivalent to maximizing the
460 clipped surrogate objective while penalizing KL divergence from the reference policy.

461 B RESULTS

462 B.1 METRIC EQUATIONS

463 The normalized minimum free energy (MFE) is defined as:

$$464 \text{MFE}_{\text{norm}} = \frac{1}{L} \min_S \Delta G(S) \quad (6)$$

465 where L is the length of the RNA sequence (in nucleotides), S represents a possible secondary
466 structure of the RNA, and $\Delta G(S)$ is the total free energy of structure S (in kcal/mol) as computed
467 by ViennaRNA RNAfold.

468 We report length-normalized Tm values ($T_{m,\text{norm}} = T_m/L$). Where T_m is defined as:

$$469 T_m = \frac{\Delta H}{\Delta S + R \ln(C_t)} - 273.15 \quad [^\circ\text{C}] \quad (7)$$

where ΔH and ΔS are the enthalpy and entropy changes from nearest-neighbor parameters. R is the gas constant, and C_t is the strand concentration. R and C_t use consistent units. Here we use a simplified nearest-neighbor melting temperature formula without explicit salt corrections.

The cosine similarity between two sequences s_i and s_j is defined as:

$$\text{sim}(s_i, s_j) = \frac{\mathbf{v}_{s_i} \cdot \mathbf{v}_{s_j}}{\|\mathbf{v}_{s_i}\| \|\mathbf{v}_{s_j}\|} = \frac{\sum_{m=1}^{4^k} f_m g_m}{\sqrt{\sum_{m=1}^{4^k} f_m^2} \sqrt{\sum_{m=1}^{4^k} g_m^2}} \quad (8)$$

where the index $m \in \{1, \dots, 4^k\}$ enumerates all possible k -mers over the nucleotide alphabet (A, C, G, U), and f_m and g_m are the frequencies of the m -th k -mer in sequences s_i and s_j , respectively. The overall diversity of a sequence set is computed as $D = 1 - \bar{s}$, where \bar{s} is the mean pairwise cosine similarity.

For CAI, we use codon weights

$$w_{a,i} = \frac{f_{a,i}}{\max_{1 \leq j \leq n_a} f_{a,j}}, \quad 0 < w_{a,i} \leq 1, \quad (9)$$

for the i -th synonymous codon of amino acid a with usage fraction $f_{a,i}$ among its n_a synonymous codons. The per-sequence CAI for a CDS g of length L codons c_1, \dots, c_L is then

$$\text{CAI}(g) = \left(\prod_{k=1}^L w_{c_k} \right)^{1/L} = \exp \left(\frac{1}{L} \sum_{k=1}^L \ln w_{c_k} \right). \quad (10)$$

B.2 ITERATIVE CDS OPTIMIZATION ENHANCES TRANSLATION RATES

Table B.1: Target proteins used in CDS design and optimization.

ID	Organism	Protein name	Accession	Length (aa)	Brief description
1	Rabies virus (strain Pasteur vaccins / PV) (RABV)	Glycoprotein	P08667	524	Trimeric viral fusion glycoprotein in the rabies virion envelope; main target of neutralizing antibodies.
2	Zaire ebolavirus (strain Eckron-76) (ZEBOV) (Zaire Ebola virus)	Envelope glycoprotein	P87671	676	Heavily glycosylated class I envelope glycoprotein cleaved into GP1/GP2; mediates host-cell attachment and membrane fusion.
3	Severe acute respiratory syndrome coronavirus 2 (2019-nCoV) (SARS-CoV-2)	Membrane protein	P0DTC5	222	Multi-pass structural protein that shapes the coronavirus envelope and scaffolds virion assembly.
4	Homo sapiens (Human)	L-dopachrome tautomerase	P40126	519	Melanosomal enzyme catalyzing the conversion of L-dopachrome to DHICA in the melanogenesis pathway.

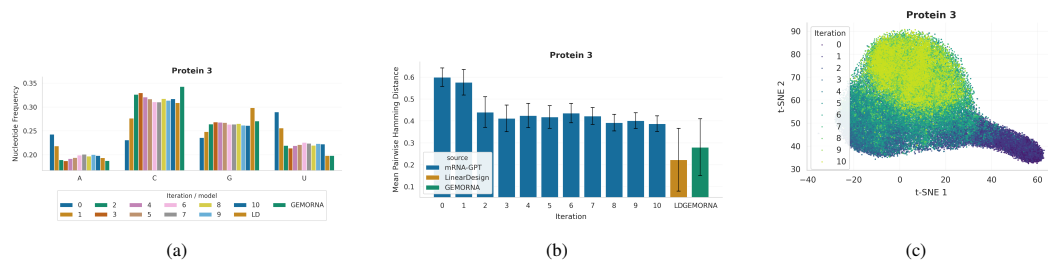


Figure B.2: CDS optimization readouts for SARS-CoV-2 membrane protein. Panels (a)–(f) depict properties of synonymous CDS generated by mRNA-GPT across iterations and by the LinearDesign and GEMORNA baselines: (a) predicted translation rates, (b) CAI, (c) MFE, (d) mean nucleotide composition, (e) sequence diversity, and (f) t-SNE embedding of codon-usage vectors.

Table B.2: Amino-Acid sequences for protein panel used for CDS design and optimization.

ID	Amino-acid sequence
1	MVPQALLFVPLLVPFLCFGKFP IYTIIPDKLGPWSP IDIHHLSCPNNLVEDEGCTNLSGFSYMEKLVGYISAIKM NGFTCTGVVTEAETYTNFVGIVTTTTFKRKHFRPTDACAAYNWKMGADPRYEESLHNPYPDYHWRVTVKTTEKES LVIISPSVADLDPYDRSLHSRVFPPGNC SGVAVSSTYCTNHDTIWMPEPNRPLGMSCDFITNSRGRKASKGSE TCGFVDERGLYKSLKGACKLKLKCGVLGLRLMDGTWVAMQTSNETKWCPPGQLVNLHDFRSEIEHLVVEELVKKR EECLDALESIMTTKSVSFRRLSHLRKLPVGF GKAYTIFNKTLM EADAHYKSVRTWNEIIPSKGCLRVGGRRCHPHV NGVFFNGIILGPDGNVLIPEMQSSLLQHMELLVSSVIPLMHP LADPSTVFKN GDEAEDFVEVHLPDVHERISGV DLGLPNWKGKVVLLSAGALTALMLIIFLMTCWRRVNRSEPTQHNLRG TGREVSVTPQSGKIISSWESYKSGGETGL
2	MGVTGILQLPRDRFKRTSFFLWV IILFQRTFSIPLGVIHNSTLQVNDVDKLVCRDKLSSTNQLRSVGLNLEGNV ATDVP SATKRWFGRSGVPPKVVNYEAGEWAENCYNLEIKKPDGSECLPAAPDGI RGFPRCRYVHKVSGTGPCAGD FAFHKEGAFFLYDRLASTVIYRGTTF AEGVVAFLILPQAKKDFSSHPLREPVNATEDPSSGYSTTIRYQATGF GTNETEYLF EVDNLT YVQLESRFT PQFLLQNETIYTSGRKSNTTGKLIWKVNPEIDTTIGEWAFWETKKNLTKR IRSEELSFTVVSNGAKNISGQSPARTSSDPGTNTT EDHKIMASENSAMVQVHSQGREAAVSHLTLTATISTSP QSLTTKPGPDNSTHNPVYKLDISEATQVEQHRRTDNDSTASDTPSATT AAGPPKAENTNTSKSTDFLDPATTT SPQNHSETAGNNTHHQDTGEE SASSNGKGLITNTIAGVAGLITGGRRTREAI VNAQPKCNPNLHYWTTQDEGA AIGLAWIPYFGPA AEGIYTEGLMHNQGLICGLRQLANETTQALQLFLRATTELRTFSILNRKAIDYFLQRWGGT CHILGDCCIEPHDWTKNITDKIDQIIHDFVDKLTLPQGDNDNWTGWRQWIPAGIGVTGVI IAVIALFCICKFV
3	MADSNGTITVEELKKLLEQWNLVIGFLFTWICLQFAYANRRFLYI IKLIFLWLLWPVTLACFVLAAYRINW ITGGIAIAMA CLVGLMWLSYFIASFRLFARTRSMWSFNPTNILLNVPLHGTILTRP LLESELVIGAVILRGHLR IAGHHLGRCDIKDLPKEITVATSRTLSYKLGASQVAGDSGFAAYSRYRIGNYKLN TDHSSSSDNIALLVQ
4	MSPLWGWFLSCLGCKILPGAQGFPRVCM TVDSL NKECCPRLGAESANVCGSQQGRGQCTEVRADTRPWSGPY ILRNQDDRELWPRKFFHRTCKCTGNFAGYNGDCKFGWTGPN CERKKPPVIRQNIHSLSPQEREQFLGALDLAKK RVHPDYVITTOHWLGLLPNGTQPQFANC SVYDFVWLHYYSVRD TLLGPGRPYRAIDF SHQGP AFVTHRYHLL CLERDLQRLIGNESFALPYWNFATGRNECDVCTDQLFGAARDDPTLISRNSRFS SWETVCDLSDYVNLVTLCN GTYEGLLRNQMG RNSMKLP TLKDIRDCLSLQKFNPPFFQNSTFSFRNALEGF DKADGTLDSQVMSLHNLVHSF LNGTNALPHSAANDP IFVVLHSFTDAIFDEW MKRFNPPADAWPQELAP IGHNRMYNMVFPFPVTNEELFLTSDQ LGYSAIDL PVSVEETPGWPTLLVVMGT LVALVGLFVLLAF LQYRRLRKG YTPLMETHLSSKRYTEEA

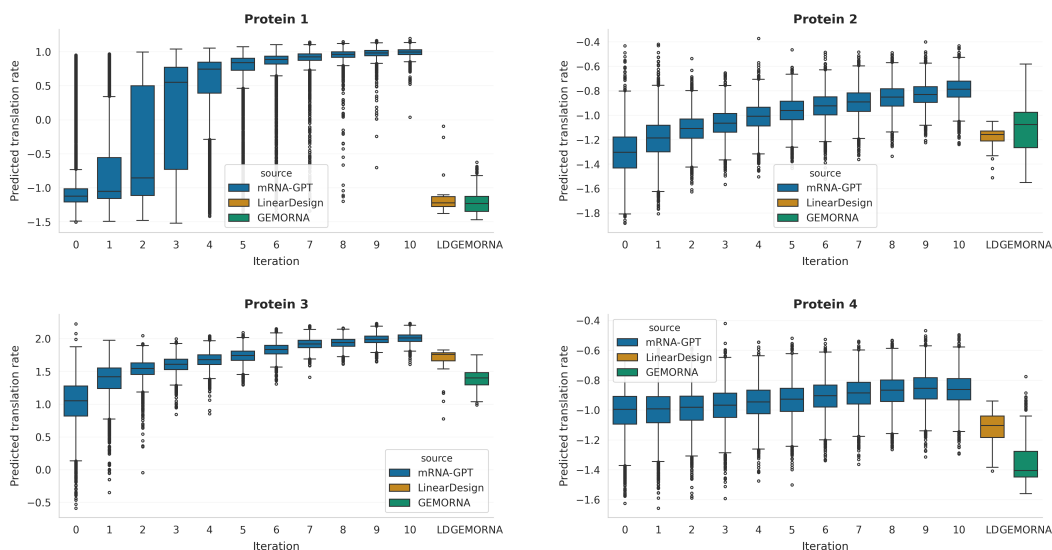


Figure B.3: Oracle-predicted translation rate for CDS-only optimization across all four proteins. Each panel reports the distribution of mRNA-LM oracle scores for 6000 CDS-only sequences generated by mRNA-GPT per optimization round ($t=0-10$; blue boxplots), together with LinearDesign (“LD”, orange) and GEMORNA (green) CDS scored by the same oracle. Across all targets, the median and upper quantiles of mRNA-GPT steadily improve with iteration and low-scoring outliers are progressively suppressed, indicating batch-wide gains rather than improvements restricted to a few top sequences. For all proteins, the final mRNA-GPT iterations clearly outperform both base-lines.

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

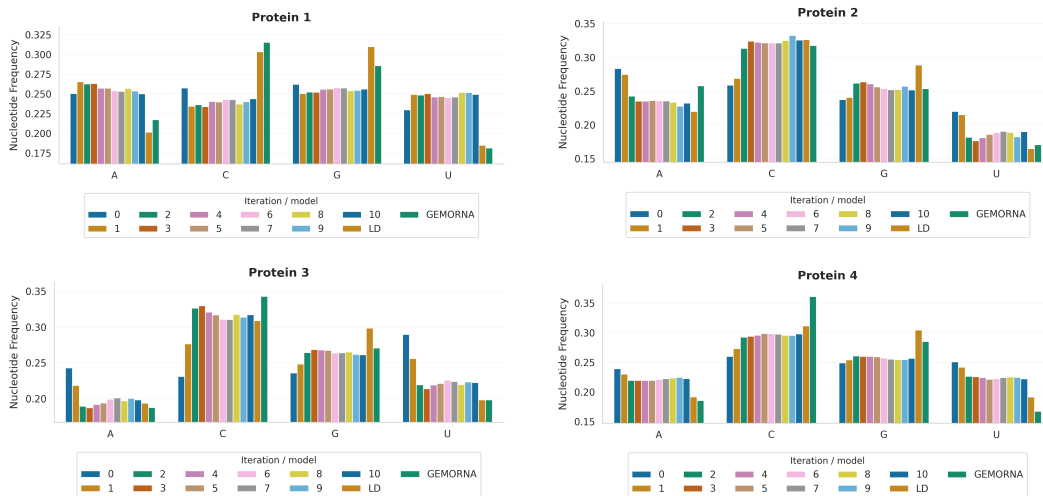


Figure B.4: **Nucleotide composition of CDS-only sequences across proteins and optimization rounds.** Mean nucleotide frequencies (A, C, G, U) in the *coding sequences* optimized by mRNA-GPT are shown for each iteration ($t=0-10$; colored bars) and for the LinearDesign and GEMORNA CDS baselines, separately for the four target proteins. Across proteins, iterative CDS optimization increases GC content (higher C/G fractions with concomitant decreases in A/U) and converges toward the GC-rich regime explored by codon-design baselines, while preserving moderate, protein-specific variability in CDS composition.

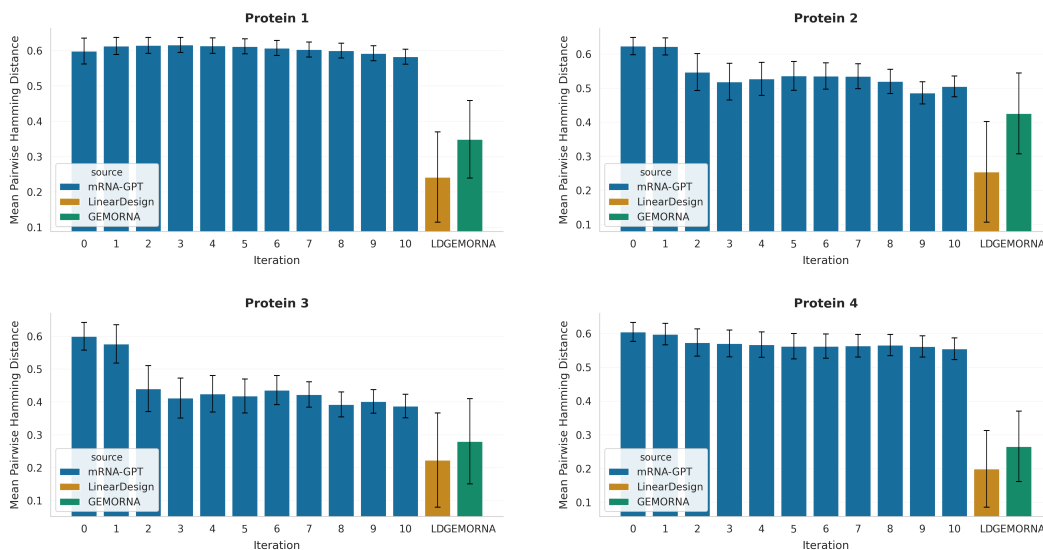


Figure B.5: **Per-protein codon-level diversity for CDS-only optimization (4 proteins).** For each protein, bars report the mean pairwise codon-level Hamming distance between *CDS-only* sequences within a batch, with error bars denoting the corresponding standard deviation. Across all proteins, mRNA-GPT maintains a high degree of intra-batch diversity throughout CDS optimization. In contrast, LinearDesign and GEMORNA produce markedly less diverse CDS sets, indicating that their solutions concentrate in a narrower region of sequence space. These results show that mRNA-GPT improves CDS quality without collapsing to a small number of near-identical coding sequences.

648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

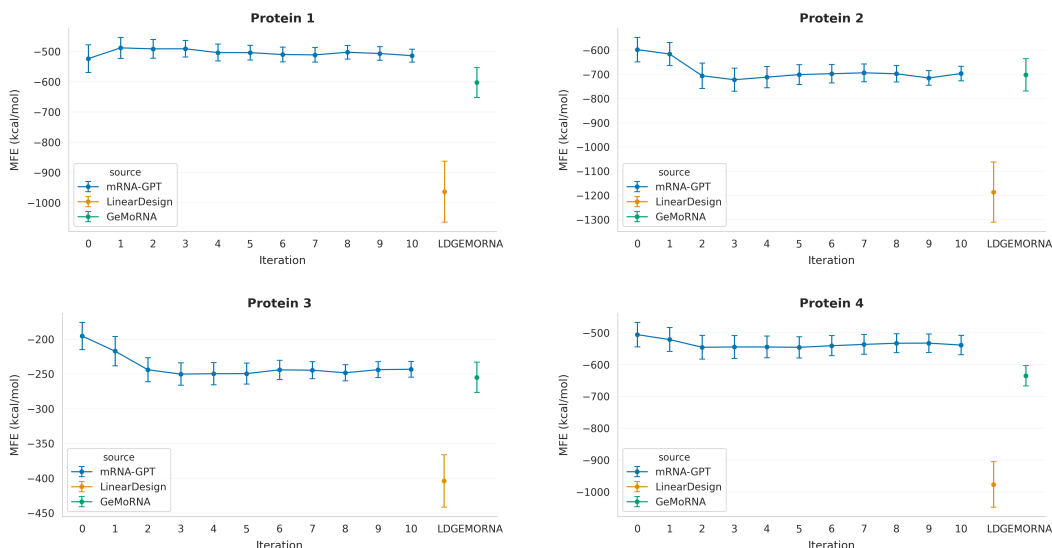


Figure B.6: **Per-protein minimum free energy (MFE) of optimized CDS-only ensembles.** For each protein, blue points trace the mean predicted minimum free energy of mRNA-GPT-generated *CDS-only* sequences as a function of optimization iteration ($t=0$ to $t=10$); vertical bars denote one standard deviation over the corresponding sequence batch. Green and orange markers on the right-hand side summarize the mean \pm s.d. MFE of GEMORNA and LinearDesign CDS, respectively, evaluated with the same RNAFOLD pipeline. Across all four proteins, mRNA-GPT operates in an intermediate MFE regime for CDS-only constructs, remaining systematically less structured (less negative MFE) than LinearDesign while typically matching or modestly improving on GEMORNA, indicating that CDS optimization does not collapse onto overly stable, potentially translation-inhibitory secondary structures.

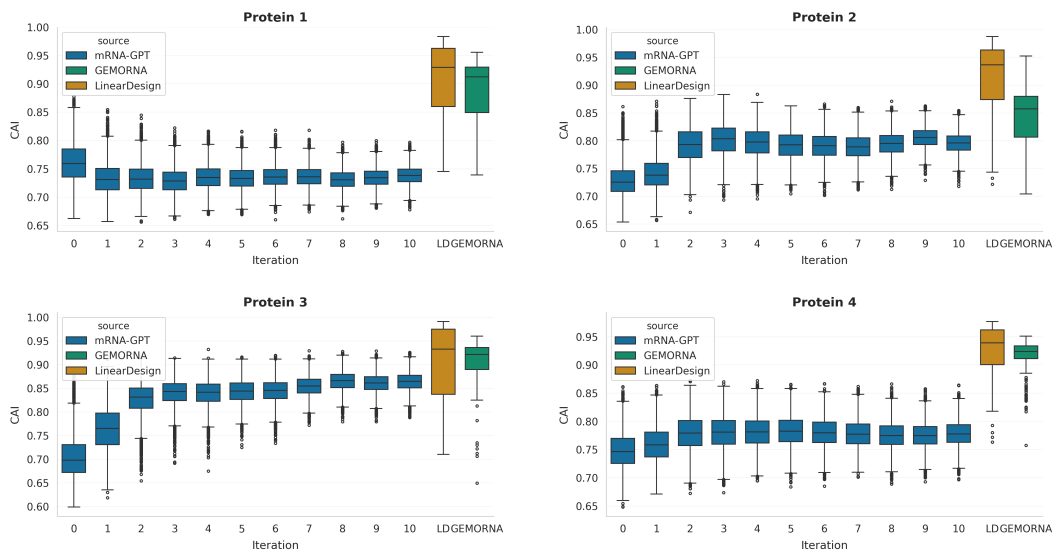


Figure B.7: **Per-protein Codon Adaptation Index (CAI) for CDS-only optimization over eleven rounds.** For each protein, boxplots report the distribution of CAI scores for the mRNA-GPT-generated *CDS-only coding sequences* from the pre-trained model ($t=0$) to iteration $t=10$, with LinearDesign and GEMORNA CDS baselines plotted as separate groups on the right. Across proteins 2, 3 and 4, early optimization rounds ($t \leq 3$) yield a pronounced upward shift and narrowing of the CAI distribution, after which median values plateau and only marginal additional gains are observed. LinearDesign and GEMORNA typically achieve higher absolute CAI values on CDS, but mRNA-GPT closes most of the gap while preserving broad sequence variability (see codon-diversity analysis) for protein 3 for instance.

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

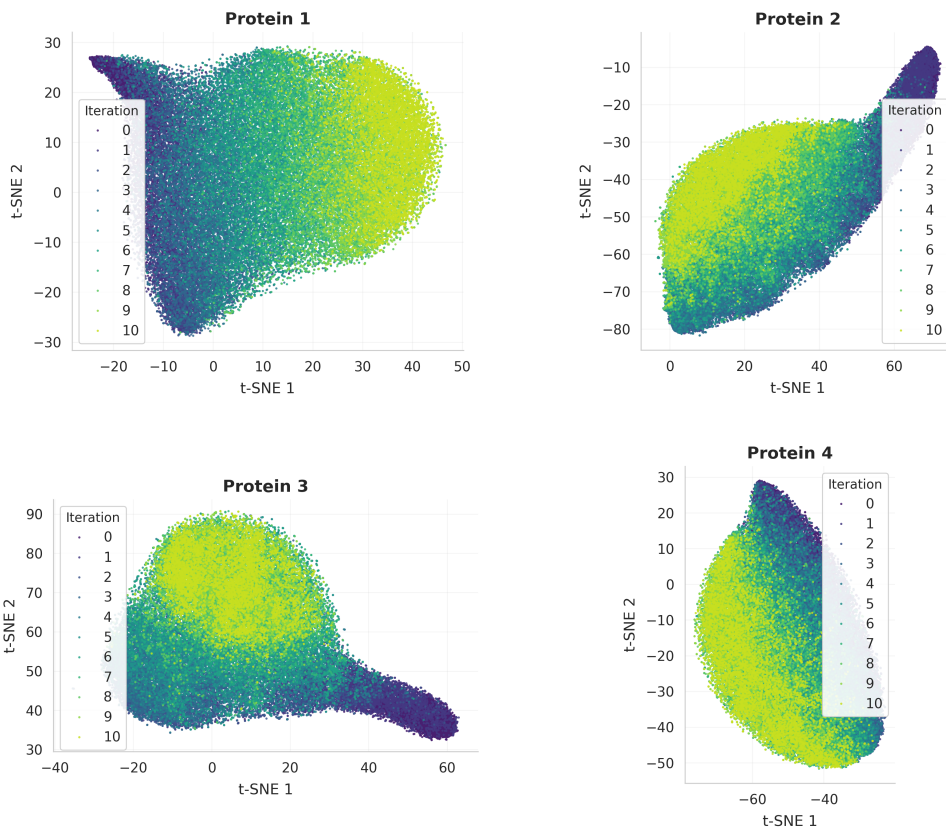


Figure B.8: **Per-protein t-SNE of CDS-only codon-usage profiles across optimization iterations.** Two-dimensional t-SNE embeddings of 64-dimensional codon-usage vectors computed on the *CDS-only* portion of mRNA-GPT-generated sequences for Proteins 1–4. Each point corresponds to one CDS and is colored by optimization iteration ($t=0-10$). Across all targets, later iterations drift smoothly over the manifold explored by the pre-trained model while continuing to occupy a broad region of codon-usage space, consistent with substantial optimization of the CDS-level oracle score without collapse onto a single codon pattern.

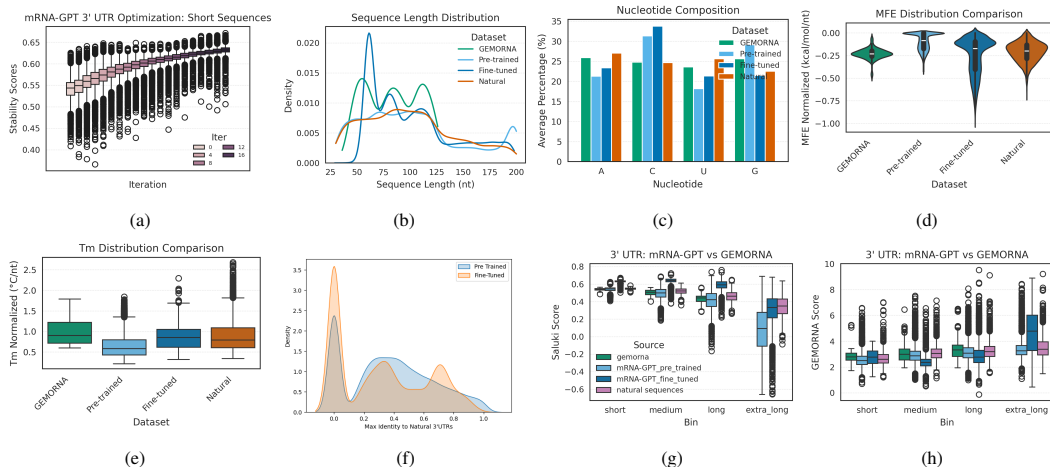


Figure B.9: **3' UTR optimization, analysis, and benchmarking.** (a) Box plot of optimization progression and stability score improvements for short-length bin. (b) Density plots of 3' UTR length distribution. (c) Bar plots of average nucleotide composition. (d) Violin plots of MFE. (e) Box plots of melting temperature. (f) Density plots of mRNA-GPT sequence novelty relative to the pre-training dataset. (g) Box plots of benchmarking using Saluki oracle. (h) Box plots of benchmarking using GEMORNA PRED-3UTR oracle.

B.3 DE NOVO GENERATION OF DIVERSE 3' UTRS WITH ENHANCED STABILITY

3' UTR is an mRNA region that does not translate for protein. However, it is still responsible for a variety of key biological functions including mRNA stability regulation and translation efficiency (Mayr, 2017). The sequence and length of the 3' UTR region influence performance and are context-dependent (Mayr, 2019). Traditional work on the optimization and design of synthetic 3' UTR employs genetic algorithms, whereby preexisting natural or synthetic sequences are iteratively modified via mutations and recombination strategies (Mitchell, 1998). Alternative approaches leveraging generative artificial intelligence have recently emerged for the *de novo* design of UTR sequences. UTRGAN applies a Wasserstein generative adversarial network with gradient penalty to synthesize UTRs, while Smart5UTR uses a multi-task autoencoder for generative reconstruction and optimization (Barazandeh et al., 2025; Tang et al., 2024). More recently, mRNAutilus and RNAdiffusion introduced diffusion-based frameworks for UTR generation (Patel et al., 2025; Huang et al., 2024). The 5' UTR Language Model demonstrated the feasibility of transformer-based generation (Chu et al., 2024).

GEMORNA (Zhang et al., 2024) introduced GEMORNA-UTR, a decoder-only transformer architectures pre-trained on natural UTR sequences (both 5' UTR and 3' UTR). For 3' UTR design, GEMORNA-UTR was further fine-tuned on some top-performing natural sequences, which were scored by a predictive model, PRED-3UTR, trained on 90 000 experimentally labeled 3' UTR sequences with measured degradation rates. This approach is limited by reliance on natural sequences and a single fine-tuning step, increasing the risk of local optima and restricting sequence diversity. In contrast, mRNA-GPT utilizes an iterative optimization strategy with an in-the-loop reward model, enabling direct optimization of desired properties. Our model eliminates the need for seed or natural sequences, evolving the candidate pool entirely through self-generated designs. The following section evaluates the quality of sequences generated by mRNA-GPT, and compares them to both GEMORNA and natural sequences.

Both the publicly available GEMORNA model and the mRNA-GPT model (prompted with the [3UTR] token) were evaluated across predefined length bins. For the GEMORNA model, 10 000 sequences were generated per bin; after removing duplicates, 128, 119, and 121 unique 3' UTR sequences remained in the short (35-65 nt), medium (66-95 nt), and long (96-125 nt) length bins, respectively. Note that the released GEMORNA model produces only about 100 diverse sequences per bin, its sample may not be representative. For both the pre-trained mRNA-GPT and fine-tuned mRNA-GPT, 5000 unique 3' UTR sequences were generated per length bin. Because mRNA-GPT

864 can produce arbitrarily long sequences, an additional extra-long bin (126-200 *nt*) with 5000 se-
865 quences was included. As a natural baseline, 1000 vertebrate 3' UTR sequences from RefSeq were
866 collected per bin, yielding 4000 reference sequences.

867 Comparing sequence properties across samples, we see that the length distribution of natural 3' UTR
868 sequences is broad and relatively uniform, whereas both GEMORNA and fine-tuned mRNA-GPT
869 samples exhibit greater variability, with distinct peaks suggesting preferential generation within
870 certain length ranges. In terms of nucleotide composition, natural and GEMORNA sequences
871 have uniform nucleotide distributions, whereas pre-trained and fine-tuned mRNA-GPT samples are
872 cytosine-rich. Fine-tuning lowers the overall GC content compared to pre-trained mRNA-GPT,
873 which is consistent with the known negative correlation between GC content and stability (Litter-
874 man et al., 2019). Length-normalized minimum free energy (MFE) calculated with the ViennaRNA
875 package (Lorenz et al., 2011) (see Eq. equation 6), shows that fine-tuned mRNA-GPT sequences
876 have lower MFE values, indicating greater thermodynamic stability and closer resemblance to nat-
877 ural profiles (Lorenz et al., 2011; Zhang et al., 2023). Similarly, melting temperature (T_m) analysis
878 computed (see Eq. equation 7) reveals that fine-tuning increases thermal stability to levels compar-
879 able to natural and GEMORNA sequences (Cock et al., 2009).

880 To assess novelty, we converted the pre-trained dataset of over 10 million 3' UTR sequences into
881 a BLAST database using a 15-*nt* threshold for hits, and measured each sequence's longest match.
882 Fine-tuned mRNA-GPT produces a greater number of novel sequences than the pre-trained mRNA-
883 GPT. To assess diversity, we computed 5-mer frequencies for each sequence in the fine-tuned
884 mRNA-GPT and then calculated cosine similarity (see Eq. equation 8). We find that the mean
885 cosine similarity is 0.23, which means there is high diversity and space exploration with minimal
886 overlap in sequence pattern. Collectively, these results suggest improved predicted stability in the
887 fine-tuned mRNA-GPT outputs. The fine-tuned mRNA-GPT performs on par with GEMORNA and
888 natural sequences, occasionally surpassing them. It also produces a higher proportion of novel and
889 diverse 3' UTR sequences, demonstrating its capacity to generate variable sequences with promising
890 characteristics.

891 To apply the iterative optimization with mRNA-GPT for 3' UTR design, we used the Saluki model
892 as the reward model, which is a hybrid convolutional and recurrent deep neural network trained on a
893 large dataset of experimentally measured human mRNA half-lives to predict mRNA stability directly
894 from nucleotide sequence (Agarwal & Kelley, 2022). Optimization was performed via PPO for 20
895 iterations, sampling 5000 candidate 3' UTR sequences per bin each round. Fig. B.9a demonstrates
896 improved predicted half-lives across iterations (shown for the short bin), evidencing the effective-
897 ness of the iterative optimization procedure. Notably, the median predicted half-life increased from
898 0.54 to 0.63, and the maximum value from 0.63 to 0.67. Corresponding results for all four bins
899 are in Fig. B.10. Fig. B.9g displays the predicted half-life distributions for sequences generated
900 by mRNA-GPT, GEMORNA-UTR, and natural samples. The pre-trained mRNA-GPT achieves a
901 comparable median to GEMORNA-UTR and natural sequences, but spans a wider range, especially
902 for longer sequences, thus producing some sequences with higher predictions than GEMORNA-
903 UTR. Fine-tuning mRNA-GPT further improves predicted half-lives beyond all benchmarks. For
904 a direct comparison, sequences from each method were also scored using the PRED-3UTR oracle
905 used in GEMORNA-UTR optimization. In this setting, pre-trained mRNA-GPT has a lower median
906 than GEMORNA-UTR but exhibits many high-value outliers, outperforming GEMORNA-UTR at
907 the distribution tail. Fine-tuned mRNA-GPT is not directly optimized to maximize predictions by
908 the PRED-3UTR oracle and therefore does not achieve high predicted values under this metric,
909 which reflects the specificity of the optimization target. These results demonstrate that mRNA-
910 GPT's iterative, reward-driven optimization enables broader exploration and superior identification
911 of high-stability 3' UTR designs.

912
913
914
915
916
917

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971




Figure B.10: Box plot of 3' UTR optimization progression and stability score improvements for short, medium, long, and extra long bin, respectively.