



BEACON: Benchmark for Comprehensive RNA Tasks and Language Models

Yuchen Ren^{1,2*} Zhiyuan Chen^{1,3*} Lifeng Qiao^{1,4*}
Hongtai Jing⁵ Yuchen Cai^{1,5} Sheng Xu^{1,5} Peng Ye^{1,5†} Xinzhu Ma^{1,6†}
Siqi Sun^{1,5} Hongliang Yan¹ Dong Yuan² Wanli Ouyang¹ Xihui Liu^{1,3}

*equal contribution †corresponding author

¹Shanghai AI Lab ²USYD ³HKU ⁴SJTU ⁵FDU ⁶CUHK

{renyuchen, chenzhiyuan, qiaolifeng, caiyuchen, xusheng, yepeng, maxinzhu, sunsiqi, yanHongliang, ouyangwanli, liuxihui}@pjlab.org.cn, htjing21@m.fudan.edu.cn, dong.yuan@sydney.edu.au

Abstract

RNA plays a pivotal role in translating genetic instructions into functional outcomes, underscoring its importance in biological processes and disease mechanisms. Despite the emergence of numerous deep learning approaches for RNA, particularly universal RNA language models, there remains a significant lack of standardized benchmarks to assess the effectiveness of these methods. In this study, we introduce the first comprehensive RNA benchmark BEACON (BEnchmArk for COmprehensive RNA Task and Language Models). First, BEACON comprises 13 distinct tasks derived from extensive previous work covering structural analysis, functional studies, and engineering applications, enabling a comprehensive assessment of the performance of methods on various RNA understanding tasks. Second, we examine a range of models, including traditional approaches like CNNs, as well as advanced RNA foundation models based on language models, offering valuable insights into the task-specific performances of these models. Third, we investigate the vital RNA language model components from the tokenizer and positional encoding aspects. Notably, our findings emphasize the superiority of single nucleotide tokenization and the effectiveness of Attention with Linear Biases (ALiBi) over traditional positional encoding methods. Based on these insights, a simple yet strong baseline called BEACON-B is proposed, which can achieve outstanding performance with limited data and computational resources. The datasets and source code of our benchmark are available at <https://github.com/terry-r123/RNABenchmark>.

1 Introduction

RNA plays a vital role in numerous biological processes, including protein synthesis, enzymatic activities, and gene regulations [16, 47, 73, 56]. Unlike its more famous counterpart DNA, RNA is not restricted to information storage but actively participates in translating genetic instructions into functional proteins, modulating gene expression through various mechanisms, and regulating cellular responses to internal and external stimuli [27]. As a dynamic intermediary between DNA and protein, RNA governs crucial biological processes, making it a focal point of research in molecular biology and biomedicine. Consequently, understanding the diverse functions of RNA is crucial to unraveling the complexities of cellular processes and deciphering the underlying mechanisms of diseases.

Despite its critical importance, understanding the functional roles of RNA poses significant challenges. Inspired by the success of machine learning in various fields, there have been extensive research efforts in recent years to apply machine learning approaches to RNA tasks. Initially, traditional machine learning algorithms such as support vector machine and random forest paved the way for

predictive modeling in RNA studies [38, 66, 41]. The evolution of deep learning, especially through Convolutional Neural Networks (CNNs), has enabled more nuanced analyses of RNA sequences and structures [59, 6, 35]. More recently, pre-trained language models (LM) have revolutionized RNA research, facilitating more accurate predictions of RNA function and interactions [8, 12, 76]. These advancements significantly deepen our understanding of RNA's regulatory roles in cellular processes.

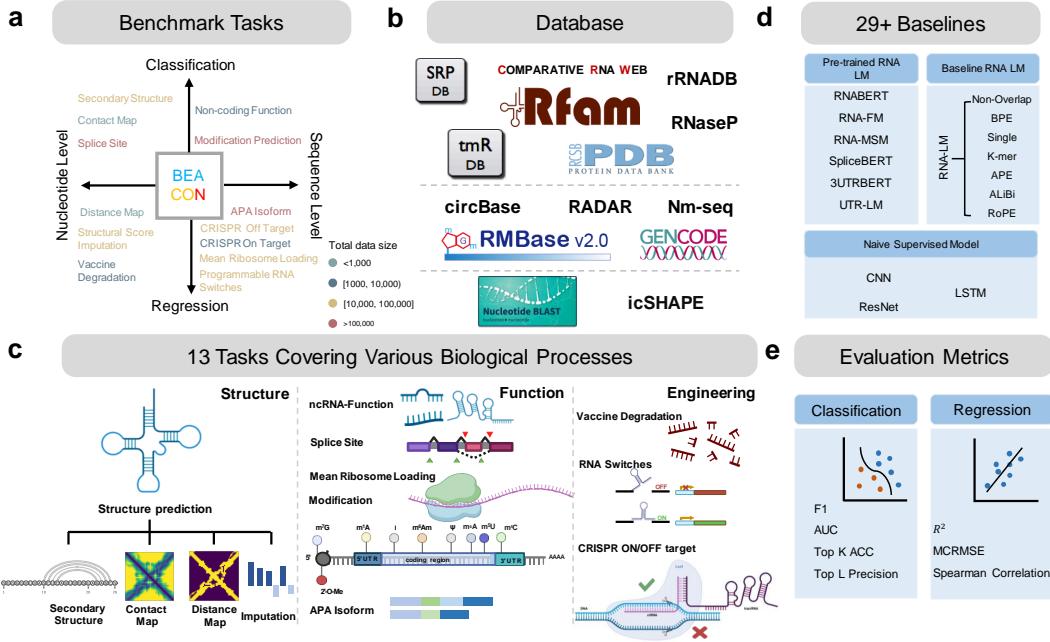


Figure 1: Overview of BEACON: **a:** Categorization of the 13 benchmark tasks into classification and regression at both nucleotide and sequence levels. **b:** Diverse database distinguished by data size and source type. **c:** Visual representations of tasks across Structure, Function, and Engineering. **d:** List of baseline models, including naive supervised deep models and advanced RNA language models. **e:** Metrics for evaluating model performance in classification and regression tasks, tailored to RNA analysis specifics.

According to the central dogma of molecular biology [14], genetic information flows unidirectionally from DNA to RNA and then to protein, or directly from RNA to protein. While established benchmarks for DNA [29, 44] and protein [51, 75] have significantly aided research in these areas, RNA, a crucial component of the central dogma, lacks such standardized benchmarks. Therefore existing RNA models are often evaluated using disparate individual datasets, making it difficult to conduct fair comparisons between different methods and hindering the development of the field.

To address this gap, we present the first comprehensive RNA benchmark called BEACON (**B**enchm**A**rk for **C**Omprehensive RNA Task and Language Models). As shown in Fig. 1, BEACON contains a curated collection of 13 important RNA-related tasks derived from a comprehensive review of RNA-related research papers [18, 35, 13], containing 967k sequences with lengths ranging from 23 to 1182. These tasks cover sequence-level and nucleotide-level analyses across three main fields: **Structural Analysis** focuses on deciphering RNA’s secondary structures and three-dimensional configurations, essential for its interactions with other molecules and for therapeutic design. **Functional Studies** investigate RNA’s roles in gene regulation and its implications for disease, which are vital for protein translation and treatment of disease. **Engineering Applications** explore RNA’s potential in synthetic biology to enhance its utility in biotechnology and medicine, exploring how RNA can be utilized to solve complex biological challenges.

In addition, we evaluate a diverse range of models using our benchmark, including traditional models like CNNs, ResNets, and LSTMs, as well as advanced RNA foundation models like RNA-FM [8] and UTR-LM [12]. Surprisingly, ResNet and LSTM are proved to be strong baselines, managing to outperform language models on several tasks. Additionally, pretrained RNA language models surpassed previous task-specific state-of-the-art (SOTA) performances on 8 out of the 13 tasks, demonstrating significant potential. Next, we explore the impact of various components in RNA

language models for the community, with a particular focus on tokenization methods and positional encodings. We conclude some experimental findings, based on which we further propose a robust yet efficient baseline, BEACON-B, that incorporates Attention with Linear Biases (ALiBi) and single nucleotide tokenization, providing an extremely fast and easy-to-use open-source pre-training model for the community.

Overall, our contributions can be summarized as follows:

- We establish the first comprehensive benchmark for RNA research with 13 diverse tasks, covering structure, function, and engineering aspects.
- We conduct a thorough evaluation of pre-trained RNA language models, providing insights into their strengths and limitations across different tasks.
- We investigate the component impacts of RNA language models in depth, and propose BEACON-B, a simple yet strong baseline, that benefits subsequent research in the field.

2 Related Works

RNA tasks. RNA research is categorized into three primary areas: structure, function, and engineering. Structural tasks, such as predicting secondary structures [18] and contact maps [67], aim to understand RNA configurations. Functional tasks focus on the biological roles of RNA, including splice site prediction [35] and non-coding RNA function classification [2]. Engineering tasks involve designing RNA molecules with specific properties for applications in synthetic biology, such as discriminating programmable RNA switches [4].

Deep learning methods in RNA tasks. Deep learning has been pivotal in addressing these tasks. For structural predictions, U-Net [54] has been employed to model secondary structures [25]. In functional studies, methods like SpliceAI utilize dilated convolutions for effective splice detection [35]. For engineering challenges, LSTMs have been used to design programmable RNA switches, demonstrating their versatility in handling complex sequence data [4]. The development of foundational RNA models like RNA-FM [8], RNA-BERT [1], RNA-MSM [78], SpliceBERT [10], UTR-LM [12] and 3UTRBERT [76] represents a significant advancement, capable of tackling multiple RNA-related tasks by leveraging advanced language modeling techniques. These models promise a broader understanding of RNA biology. However, they often lack thorough evaluations across different tasks, highlighting a gap in the systematic assessment of their capabilities. For instance, UTR-LM [12] only focuses on 5' UTR function-related tasks. This limitation underscores the need for robust, cross-disciplinary evaluation approaches to fully explore and utilize the potential of these models in RNA research.

Benchmarks in molecular biology. While AI for RNA research is a relatively new field and lacks comprehensive benchmarks, numerous benchmarks have been established for DNA [29, 17, 45] and protein [51, 75, 48, 26] studies. Grešová et al. proposed Genomic Benchmarks [29], which includes a collection of genomic sequence classification tasks. Marin et al. constructed BEND [45], a comprehensive benchmark for DNA, encompassing tasks such as gene finding, enhancer annotation, histone modification, CpG methylation, etc. Notin et al. introduced ProteinGym [48], a benchmark specifically designed for protein fitness prediction and design, while Gao et al. proposed ProteinBench [26], focusing on protein design. Additionally, Xu et al. developed PEER [75], a comprehensive protein benchmark involving function, localization, structure, etc. As AI for RNA research develops, some preliminary benchmarks in RNA have emerged. RnaBench [57] is the benchmark for computational RNA modeling, though it only involves tasks related to RNA secondary structure and RNA design. Many real-life applications, such as RNA-based therapeutics, require a comprehensive understanding of the functions of ncRNA and mRNA [80]. To address this gap, we developed BEACON, a comprehensive benchmark for RNA that covers a wide range of tasks covering structural analysis, functional studies, and engineering applications.

3 Benchmark Tasks

BEACON comprises 13 tasks designed to evaluate RNA models comprehensively, covering structural analysis, functional studies, and engineering applications. The following sections provide detailed information, including data statistics, evaluation metrics, and data sources as shown in Tab 1.

3.1 Structure Prediction

Secondary Structure Prediction (SSP) identifies paired regions (stems) and unpaired regions (loops, bulges, junctions) within RNA molecules. The target matrix is $y \in \mathbb{R}^{l \times l}$ indicating whether each nucleotide pair forms a base pair as part of the RNA's secondary structure. We adopt the bpRNA-1m database [18], which contains detailed annotations of over 100,000 single-molecule RNA structures. The evaluation metric is the F1 score.

Impact: Accurate secondary structure prediction is pivotal for elucidating the structural and functional dynamics of RNA. By precisely mapping these structures, researchers gain insights into functional regions and interaction sites, contributing significantly to areas such as drug discovery and genetic research.

Contact Map Prediction (CMP) identifies pairs of nucleotides in RNA that are in close proximity in their three-dimensional structures. Each nucleotide pair is associated with a binary label $y \in \{0, 1\}$ indicating whether they contact (within a distance threshold of 8 Å). Following [67], we utilize a dataset derived from non-redundant RNA 3D structures documented by [39], and evaluate predictions using the Top-L precision metric.

Impact: Accurately identifying nucleotide interactions is pivotal for inferring the tertiary structure of RNA molecules. These predictions enhance our understanding of RNA folding and function, contributing to advancements in RNA-based therapeutics and biotechnology.

Distance Map Prediction (DMP) estimates the physical distances between pairs of nucleotides within an RNA molecule. The target distance matrix $y \in \mathbb{R}^{l \times l}$ records the distance between every pair of nucleotides within the sequence. The same dataset as the contact map prediction task is used, with R^2 serving as the evaluation metric.

Impact: Inter-nucleotide distance prediction offers detailed spatial information, facilitating the construction of accurate three-dimensional models by providing distance restraints.

Structural Score Imputation (SSI) predicts missing structural information within RNA molecules, with each nucleotide assigned with an experimentally derived structural score $y \in R$. The dataset [28] is derived from icSHAPE sequencing data of the HEK293 cell line, with 30% of nucleotides randomly masked as null in the training set, and downsampling leading to missing values in 3,095 fragments in the testing set. The evaluation metric is R^2 .

Impact: Accurate structural score imputation provides enhanced and comprehensive structural information, crucial for the development of RNA-based therapeutics and diagnostics. Improved structural data enable precise targeting of RNA molecules in disease treatment, potentially leading to more effective interventions.

3.2 Function Prediction

Splice Site Prediction (SPL) classifies each base within a sequence into one of three categories: acceptor (a), donor (d), or neither (n), with the categorical label $y \in \{0, 1, 2\}$. The task uses Jaganathan's dataset [35] with Top-k accuracy as the evaluation metric.

Impact: Splice site prediction is crucial for studying gene expression and regulation within biological systems. Accurate prediction of splice sites helps determine the precise locations within the genome where splicing occurs, enabling the detection of non-coding genomic variations that could impact protein synthesis, particularly those resulting in cryptic splicing.

APA Isoform Prediction (APA) predicts the usage ratio of the proximal polyadenylation site (PAS) in the 3' untranslated region (3' UTR) for each variant, recorded in target $y \in \mathbb{R}$. We filter 228k sequences from over 3 million APA reporter gene data from Bogard's dataset [6], this regression task assesses the proportion of proximal APA isoforms. The evaluation metric is the R^2 value.

Impact: APA is a common gene expression regulation mechanism that generates different RNA transcripts and protein isoforms by modulating RNA 3' UTR processing [77]. This regulatory method can affect gene expression levels and functions, playing a crucial role in cellular biological processes and development.

Non-coding RNA Function Classification (ncRNA) classifies ncRNA molecules into categories like microRNAs (miRNAs), long non-coding RNAs (lncRNAs), and small interfering RNAs (siRNAs).

Table 1: Overview of the 13 benchmark tasks across three major RNA task groups. Nucleotide-level tasks require labels to have the same length as input sequences. Sequence-level tasks require nucleotides from one input sequence to share one label. Cls and Reg denote classification and regression, respectively. MCRMSE means Mean Columnwise Root Mean Squared Error.

RNA Task	#Train/Validation/Test	Metric	Task Type	RNA Level	Max/Mean Length	Source/Venue
Structure						
SSP	10,814/1,300/1,305	F1	Multi-label Cls	Nucleotide	499/133.8	bpRNA/NAR [18]
CMP	188/23/80	Top L Precision	Multi-label Cls	Nucleotide	960/110.3	RNAcontact/BIOINF [39]
DMP	188/23/80	R^2	Reg	Nucleotide	960/110.3	RNAcontact/BIOINF [39]
SSI	14,049/1,756/3,095	R^2	Reg	Nucleotide	100/100	StructureImpute/NMI [28]
Function						
SPL	144,628/18,078/16,505	Top-k ACC	Multi-class Cls	Nucleotide	100/100	SpliceAI/Cell [35]
APA	145,463/33,170/49,755	R^2	Reg	Sequence	186/186	APARENT/Cell [6]
ncRNA	5,679/650/2,400	ACC	Multi-class Cls	Sequence	1182/158.4	Noorul's/NMI [2]
Modif	304,661/3,599/1,200	AUC	Multi-label Cls	Sequence	101/101	MultiRM/NC [63]
MRL	76,319/7,600/7,600	R^2	Reg	Sequence	100/61.5	Optimus/NBT [59]
Engineering						
VDP	2,155/245/629	MCRMSE	Multi-label Reg	Nucleotide	130/118.5	OpenVaccine/NMI [71]
PRS	73,227/9,153/9,154	R^2	Multi-label Reg	Sequence	148/148	Angenent-Mari's/NC [3]
CRI-On	1,453/207/416	Spearman Corr	Reg	Sequence	23/23	DeepCRISPR/GB [13]
CRI-Off	14,223/2,032/4,064	Spearman Corr	Reg	Sequence	23/23	DeepCRISPR/GB [13]

Each molecule is assigned a categorical label $y \in \{0, 1, \dots, 12\}$ to denote its function. The dataset [2, 24] comprises contributions from GENCODE, circBase, and Rfam, encompassing various ncRNAs. Accuracy (ACC) at the sequence level is the evaluation metric.

Impact: Classifying ncRNA functions is crucial for understanding their diverse roles in gene regulation and cellular processes. Accurate classification enhances our knowledge of regulatory networks and aids in elucidating disease mechanisms. This contributes to identifying new biomarkers and therapeutic targets, advancing molecular biology research, and improving disease diagnosis and treatment.

Modification Prediction (Modif) predicts twelve widely occurring types of RNA modifications from a given RNA sequence, indicated by a categorical label $y \in \{0, 1, \dots, 11\}$. We adopt Song's dataset that contains 20 epi-transcriptome profiles for 12 different types of RNA modifications obtained from 15 base-resolution technologies, where over 300,000 sites were collected and divided into training, validation, and test sets. We use AUC as the metric.

Impact: Post-transcriptional RNA modifications enhance the structural and functional diversity of RNA molecules, impacting all stages of RNA life [22]. Due to the complex and diverse characteristics of RNA sequences, different modifications may correspond to distinct sequence features. Accurately identifying RNA modification sites is crucial for understanding the functions and regulatory mechanisms of various RNAs.

Mean Ribosome Loading (MRL) predicts the MRL value for a given sequence, with target $y \in \mathbb{R}$ representing the level of mRNA translation activity into proteins. Data from Reid's dataset [59] of 91,519 5' UTR sequences and their variants are used to calculate the MRL for each sequence. The model's performance is evaluated using the R^2 value.

Impact: MRL refers to the average ribosome load on a specific mRNA sequence under given conditions, indicating the translation efficiency of ribosomes on that mRNA. Modulating the features and structures of the 5' UTR sequence can influence ribosome loading on mRNA, thereby regulating protein expression levels [40, 7].

3.3 Engineering Prediction

Vaccine Degradation Prediction (VDP) forecasts the stability and shelf life of vaccines under different environmental conditions. For each nucleotide, the three properties are recorded in target $y \in \mathbb{R}^3$. We use data from the "Stanford OpenVaccine" [71] competition on Kaggle and the RNA design platform Eterna, which includes detailed measurements for 6,043 diverse RNA constructs. The evaluation metric is the Mean Columnwise Root Mean Squared Error (MCRMSE).

Impact: Accurate predictions of vaccine degradation under different environmental conditions are crucial for optimizing storage and transportation protocols, ensuring vaccines remain potent until administration. Enhanced degradation predictions are particularly beneficial for distributing

vaccines in challenging environments, such as resource-limited settings, by providing guidelines to maintain vaccine stability and efficacy.

Programmable RNA Switches (PRS) involves identifying synthetic RNA molecules that can alter their conformation and function in response to specific signals. The target $y \in \mathbb{R}^3$ records the ON, OFF and ON/OFF states activity given an RNA sequence. The dataset, analyzed by Angenent-Mari [3], includes 91,534 toehold switches *in vivo*, covering 23 viral genomes and 906 human transcription factors, with GFP signal intensity measurements indicating ON and OFF states activity levels [4]. The R^2 metric evaluates the effectiveness of these switches.

Impact: Programmable RNA switches provide precise control of gene expression and cellular functions, serving as powerful tools for investigating biological processes [64, 46]. In therapeutic applications, these switches hold promise for developing targeted and personalized treatments by responding to disease-specific signals, offering innovative approaches to medical intervention.

CRISPR On-Target Prediction (CRI-On) evaluates the efficiency of single-guide RNAs (sgRNAs) directed by Cas proteins in gene editing within specific target sites. Each sgRNA's knockout efficacy is quantified and presented as target $y \in \mathbb{R}$. The dataset [13] comprises approximately 15,000 sgRNAs targeting 1,071 genes across four different cell lines, with performance evaluated using the Weighted Spearman correlation coefficient.

Impact: CRISPR-Cas technology has transformed genetic engineering with significant enhancements in genome editing accuracy and safety. Effective on-target predictions are essential for designing sgRNAs that precisely modify genetic sequences without affecting unintended regions, thus improving therapeutic outcomes and research accuracy [74, 33].

CRISPR Off-Target Prediction (CRI-Off) assesses the likelihood and frequency of CRISPR-induced mutations at unintended genomic locations. The efficacy of sgRNA specificity is quantified using a target $y \in \mathbb{R}$, capturing the frequency of off-target cleavage. The evaluation dataset [13] contains data for about 160,000 potential off-target sites across 30 sgRNAs in various cell types, with the Weighted Spearman correlation coefficient serving as the metric.

Impact: Precision in off-target predictions is critical for advancing CRISPR technology by reducing unintended genetic modifications, which can lead to harmful effects. Accurate off-target analysis helps refine sgRNA designs, enhancing the safety and efficacy of CRISPR applications in clinical settings and research.

4 Models

We consider three types of baseline models in our benchmark, including naive supervised models, pre-trained language models, and the proposed BEACON-B. We give the details in the following part and summarize them in Tab. 2.

Naive Supervised Models. We utilize three widely-used sequence encoders: CNN [62], ResNet [51], and LSTM [51]. We mainly follow the design choices described in [75], employing 2 layers for CNN, 8 resblocks for ResNet, and 3 Bi-LSTM layers for LSTM, with 5.4M, 11M, and 26.7M parameters, respectively.

Pre-trained Language Models. We evaluate the performance of several language models, including RNA-FM [8], RNABERT [1], RNA-MSM [78], SpliceBERT [10], 3UTRBERT [76], and UTR-LM [12]. These models vary significantly in size, ranging from 0.48M to 99.52M parameters, and are pre-trained on diverse RNA data sources including ncRNA, pre-mRNA, mRNA-3'UTR, and mRNA-5'UTR. For consistency, we choose to fine-tune them using identical settings.

Baseline RNA LM and BEACON-B. We conduct ablation studies on two key aspects of RNA LM: 1) tokenization methods including Single Nucleotide (Single), Byte-Pair Encodings (BPE) [61, 79], Overlapping K-mer (K-mer, we use 6mer for experiments) [76, 36] and Non-overlapping K-mer (Non-overlap) [17] 2) positional encodings including Absolute Positional Encodings (APE) [21], Attention with Linear Biases (ALiBi) [50] and Rotary Positional Encodings (RoPE) [65]. The findings indicate that single nucleotide tokenization outperforms both K-mer, BPE, and Non-overlap, and ALiBi shows advantages over both RoPE and APE. Consequently, we propose a robust yet efficient BEACON baseline (BEACON-B) that incorporates single nucleotide tokenization and ALiBi as positional encodings, based on the BERT backbone.

Table 2: Detailed specifications and pre-training data of RNA language models analyzed in the study.

RNA Foundation Model	Number of Parameters (M)	Max Token length	Pre-trained Data	Tokenizer	Positional Encoding
RNA-FM [8]	99.52	1024	Multispecies ncRNA [68]	Single	APE
RNABERT [1]	0.48	440	Human ncRNA [68]	Single	APE
RNA-MSM [78]	95.92	1024	Homologous sequences [37, 9]	Single	APE
SpliceBERT-H510 [10]	19.45	510	Human pre-mRNA [30]	Single	APE
SpliceBERT-MS510 [10]	19.45	510	Multispecies pre-mRNA [30]	Single	APE
SpliceBERT-MS1024 [10]	19.72	1024	Multispecies pre-mRNA [30]	Single	APE
UTR-LM-MRL [12]	1.21	1026	Multispecies 5'UTR [15, 58, 7]	Single	RoPE
UTR-LM-TE&EL [12]	1.21	1026	Multispecies 5'UTR [15, 58, 7]	Single	RoPE
3UTRBERT-3mer [76]	86.14	512	Human 3'UTR [32]	K-mer	APE
3UTRBERT-4mer [76]	86.53	512	Human 3'UTR [32]	K-mer	APE
3UTRBERT-5mer [76]	88.45	512	Human 3'UTR [32]	K-mer	APE
3UTRBERT-6mer [76]	98.05	512	Human 3'UTR [32]	K-mer	APE

5 Results

5.1 Training setups

To ensure a fair comparison, we fully fine-tune all the BERT-like RNA foundation models including RNA-FM, RNABERT, RNA-MSM, SpliceBERT, 3UTRBERT, UTR-LM and BEACON-B under the same training settings. For simple supervised methods (CNN, ResNet and LSTM) and baseline RNA LM, we train them from scratch using similar training settings. For each model, we search for its learning rate from 1e-5 to 5e-3. All experiments are repeated with three random seeds, and we report the average performance alongside sample standard deviations. More details are in Appendix A.1.

5.2 Task Pipeline

Our approach incorporates three pipelines for different types of tasks in the BEACON. In nucleotide-level tasks, due to the complexity of outputs in structural tasks, we further categorize the tasks of Secondary Structure, Contact Map, and Distance Map into a more detailed nucleotide-nucleotide level prediction.

Sequence Level Prediction For sequence-level tasks, we apply an attentive weighted sum of all nucleotides for naive supervised models and use the [CLS] token from language models. Both representations are processed through an MLP layer to derive the sequence-level predictions.

Nucleotide Level Prediction For tasks requiring resolution at the nucleotide level, individual representations for each nucleotide are processed through a Multilayer Perceptron (MLP) to generate nucleotide-level predictions. Specifically, the representation for a nucleotide are calculated by averaging the representations of all tokens that cover it, as illustrated in Appendix Fig. 2.

Nucleotide-Nucleotide Relation Prediction To analyze relationships between nucleotides, we compute a self outer product of the nucleotide representations to form a matrix that cotains the pairwise interactions between nucleotides. This matrix is then passed through a simple Resnet to get the final output.

5.3 Benchmark results

In Tab 3, we report the benchmark results for popular and opensource methods, including literature SOTAs, naive supervised models and existing RNA language models.

ResNet and LSTM are strong naive supervised models. ResNet, which has only been trained on downstream tasks, can outperform most if not all language models on some tasks. LSTM outperforms the other Naive supervised Models on 9 out of 13 tasks, and the performance is better by a large margin on many tasks.

Pre-trained RNA language models have good potential for RNA understanding. It outperforms the previous task-specific SOTA on 8 out of 13 tasks, demonstrating that the additional unsupervised pre-training brings a lot of gains. However, there is still a long way to go on individual tasks, such as contact map prediction and distance map prediction in the structural task, vaccine degradation rate prediction in the engineering task, and CRISPR on- and off-target prediction. Of course, the previous SOTA method used additional features such as secondary structure, but it shows that there is still a lot of room for improvement in the RNA language model.

Table 3: Benchmark results across various 13 RNA tasks. We use four color scales of blue to denote the **first**, **second**, **third** and **fourth** best performance among naive supervised models and pre-trained RNA LMs. Mean (std) is reported for each experiment.

Task Metric	SSP F1 (%)	CMP P@L (%)	DMP R^2 (%)	SSI R^2 (%)	SPL ACC@K (%)	APA R^2 (%)	NcRNA ACC (%)	Modif AUC (%)	MRL R^2 (%)	VDP MCRMSE \downarrow	PRS R^2 (%)	CRI-On SC (%)	CRI-Off SC (%)
Literature SOTA													
UFold [25]	65.4	RNACon [67]	SS+Seq [8]	StructImp [28]	SpliceAI [35]	APARENT [6]	GCN [55]	MultiRM [63]	Optimus [59]	NAttn [53]	MLP-O [3]	SSC [74]	DeepCRI [13]
CNN	49.95 _(0.82)	43.89 _(5.53)	27.76 _(5.00)	34.36 _(0.12)	8.43 _(0.38)	50.93 _(0.17)	88.62 _(0.71)	70.87 _(0.40)	74.13 _(0.58)	0.361 _(0.000)	45.40 _(0.66)	29.69 _(2.52)	11.40 _(0.10)
ResNet	57.26 _(3.14)	59.59 _(0.68)	30.26 _(1.81)	37.74 _(0.16)	21.15 _(1.56)	56.45 _(0.94)	88.33 _(1.22)	71.03 _(0.32)	74.34 _(0.22)	0.349 _(0.000)	55.21 _(0.28)	28.55 _(2.42)	11.50 _(0.22)
LSTM	58.61 _(0.21)	40.41 _(1.67)	44.77 _(0.47)	35.44 _(1.13)	36.66 _(1.83)	67.03 _(0.86)	88.78 _(0.10)	94.83 _(0.31)	83.94 _(0.08)	0.329 _(0.002)	55.45 _(0.71)	26.83 _(1.32)	8.60 _(0.13)
Naive supervised Model													
RNA-FM	68.50 _(0.54)	47.56 _(0.73)	51.45 _(0.51)	42.36 _(0.24)	34.84 _(0.87)	70.32 _(0.97)	96.84 _(0.06)	94.98 _(0.042)	79.47 _(0.47)	0.347 _(0.003)	55.98 _(0.09)	31.62 _(1.16)	2.49 _(1.56)
RNABERT	57.27 _(0.30)	45.21 _(10.87)	48.19 _(0.64)	31.62 _(0.64)	0.18 _(0.18)	57.66 _(2.11)	68.95 _(7.285)	82.82 _(19.09)	29.79 _(20.15)	0.378 _(0.009)	54.60 _(0.23)	29.77 _(3.98)	4.27 _(0.05)
RNA-MSM	57.98 _(0.47)	57.26 _(15.38)	37.49 _(1.10)	39.22 _(0.23)	38.33 _(0.76)	70.40 _(1.12)	84.85 _(0.266)	94.89 _(0.14)	83.48 _(0.18)	0.330 _(0.001)	56.94 _(0.38)	34.92 _(0.99)	3.85 _(0.99)
Splice-H510	64.93 _(0.84)	45.80 _(6.03)	55.56 _(1.00)	38.91 _(0.07)	44.80 _(0.93)	58.65 _(2.34)	95.92 _(0.666)	62.57 _(1.92)	83.49 _(0.47)	0.321 _(0.000)	54.90 _(3.45)	26.61 _(1.30)	4.00 _(1.13)
Splice-MS510	43.24 _(28.64)	52.64 _(7.56)	10.27 _(0.20)	38.58 _(0.50)	50.55 _(0.49)	52.46 _(17.36)	95.87 _(0.364)	55.87 _(5.25)	84.98 _(0.28)	0.315 _(0.003)	50.98 _(7.46)	27.13 _(0.27)	3.49 _(2.12)
Splice-MS1024	68.26 _(0.20)	47.32 _(3.16)	55.89 _(0.48)	39.22 _(0.02)	48.52 _(0.49)	60.03 _(1.32)	96.05 _(0.777)	53.45 _(6.25)	67.15 _(10.34)	0.313 _(0.000)	57.72 _(0.45)	27.59 _(4.61)	5.04 _(0.71)
UTR-LM-MRIL	59.71 _(0.30)	45.51 _(23.51)	55.21 _(1.91)	39.52 _(0.36)	36.20 _(1.84)	64.99 _(0.40)	89.97 _(0.617)	56.41 _(1.90)	77.78 _(6.03)	0.325 _(0.002)	57.28 _(0.10)	28.49 _(1.37)	4.28 _(0.15)
UTR-LM-TE&EL	59.57 _(0.20)	60.32 _(7.27)	54.94 _(2.54)	40.15 _(0.11)	37.35 _(5.48)	72.09 _(0.92)	81.33 _(8.551)	59.70 _(0.52)	82.50 _(1.45)	0.319 _(0.001)	53.37 _(3.34)	32.49 _(4.14)	2.91 _(1.18)
UTRBERT-3mer	60.37 _(0.47)	51.03 _(21.48)	50.95 _(0.44)	34.31 _(0.00)	44.24 _(0.53)	69.52 _(4.56)	92.88 _(0.379)	95.14 _(0.11)	83.89 _(0.13)	0.337 _(0.002)	56.83 _(0.26)	29.92 _(1.95)	4.48 _(0.12)
UTRBERT-4mer	59.41 _(0.45)	44.91 _(27.56)	47.77 _(7.08)	33.22 _(0.00)	42.04 _(0.53)	72.71 _(0.85)	94.32 _(0.346)	95.10 _(0.12)	82.90 _(0.75)	0.341 _(0.002)	56.43 _(0.67)	23.20 _(1.10)	3.11 _(1.10)
UTRBERT-5mer	47.92 _(8.75)	44.71 _(7.64)	48.67 _(1.70)	31.27 _(0.00)	39.19 _(0.37)	72.70 _(1.77)	93.04 _(0.367)	94.78 _(0.07)	75.64 _(4.70)	0.343 _(0.001)	57.16 _(0.08)	25.74 _(0.00)	3.93 _(0.24)
UTRBERT-6mer	38.56 _(28.76)	51.56 _(20.30)	50.02 _(0.05)	29.93 _(0.17)	38.58 _(2.72)	71.17 _(2.30)	93.12 _(0.168)	95.08 _(0.17)	83.60 _(0.39)	0.340 _(0.001)	57.14 _(0.12)	28.60 _(1.55)	4.90 _(0.57)
Our BEACON-B													
BEACON-B	64.18 _(0.44)	60.81 _(1.70)	56.28 _(0.41)	38.78 _(0.18)	37.43 _(1.43)	70.59 _(0.91)	94.63 _(0.16)	94.74 _(0.20)	72.29 _(0.28)	0.320 _(0.001)	54.67 _(0.36)	26.01 _(1.81)	4.42 _(0.33)
BEACON-B512	58.75 _(3.72)	61.20 _(2.11)	56.82 _(0.63)	39.13 _(0.08)	37.24 _(1.09)	72.00 _(0.17)	94.99 _(0.21)	94.92 _(0.07)	72.35 _(0.28)	0.320 _(0.001)	55.20 _(0.26)	28.17 _(1.81)	3.82 _(1.04)

SpliceBERT and RNA-FM are superior models for various tasks. Both SpliceBERT-MS1024 and RNA-FM got first place in 3 out of 13 tasks, and had top performances in other tasks as well, showing they have learned rich patterns and evolution knowledge from multi-species RNA sequences.

Pre-training of specific RNA attributes will result in a gain on tasks with corresponding attributes. First, when specific RNA attributes are included in the pre-training it brings gains to the downstream tasks corresponding to the attributes. For example, RNA-FM pre-trained with non-coding RNA achieves the best performance in non-coding RNA family prediction, SpliceBERT pre-trained on pre-mRNA learns information about potential shear mRNAs for the best shear site prediction, and 3UTRBERT uses sequences from the 3'UTR region to learn 3'UTR function worked best in the prediction of APA isoforms in the 3'UTR functional region, and similarly, the pre-trained UTR-LM in the 5'UTR region worked well in the prediction of ribosome loading in the 5'UTR association. Second, specific attributes also give gains for having other RNA attributes, for example, 3UTRBERT, although pre-trained on 3'UTR sequences, also gained on the prediction of 5'UTR function.

Table 4: Performance of baseline RNA LMs with different tokenizers and positional encodings.

Task Metric	SSP F1 (%)	CMP P@L (%)	DMP R^2 (%)	SSI R^2 (%)	SPL ACC@K (%)	APA R^2 (%)	NcRNA ACC (%)	Modif AUC (%)	MRL R^2 (%)	VDP MCRMSE \downarrow	PRS R^2 (%)	CRI-On SC (%)	CRI-Off SC (%)
Baseline RNA LM Analysis													
Non-overlap-APE	12.58 _(0.08)	41.60 _(5.34)	44.66 _(0.59)	10.36 _(0.57)	0.00 _(0.00)	58.49 _(0.75)	82.04 _(0.18)	79.07 _(20.82)	34.56 _(2.79)	0.640 _(0.000)	51.95 _(0.15)	21.64 _(3.85)	8.33 _(0.90)
Non-overlap-ALiBi	5.58 _(0.49)	57.49 _(32.45)	37.44 _(2.69)	10.84 _(1.13)	0.00 _(0.00)	57.92 _(0.36)	80.55 _(1.38)	60.56 _(9.00)	37.75 _(1.07)	0.640 _(0.000)	49.17 _(0.17)	14.85 _(4.09)	7.95 _(0.20)
Non-overlap-RoPE	4.50 _(0.16)	27.47 _(22.17)	38.68 _(0.39)	10.38 _(0.94)	0.00 _(0.00)	46.14 _(0.62)	76.76 _(0.44)	67.12 _(0.25)	24.13 _(18.27)	0.640 _(0.000)	31.76 _(0.06)	15.76 _(2.07)	8.43 _(0.22)
BPE-APE	6.30 _(0.51)	49.48 _(3.75)	41.56 _(0.16)	20.79 _(0.95)	0.00 _(0.00)	65.75 _(1.16)	80.76 _(0.95)	67.67 _(0.28)	48.67 _(1.45)	0.641 _(0.001)	29.75 _(0.08)	16.16 _(2.77)	5.78 _(1.50)
BPE-ALiBi	6.34 _(0.85)	59.17 _(18.04)	37.58 _(1.10)	25.23 _(1.67)	0.00 _(0.00)	69.03 _(1.05)	81.36 _(0.43)	63.95 _(4.77)	45.78 _(5.49)	0.642 _(0.001)	31.45 _(0.79)	16.13 _(3.84)	6.63 _(0.73)
BPE-RoPE	6.32 _(0.38)	51.31 _(23.54)	37.01 _(0.47)	21.96 _(1.13)	0.00 _(0.00)	50.91 _(1.27)	75.74 _(1.10)	62.89 _(4.01)	45.87 _(1.53)	0.642 _(0.001)	19.69 _(0.10)	16.46 _(1.75)	6.63 _(0.73)
Single-APE	48.23 _(0.26)	70.85 _(0.03)	48.22 _(0.69)	24.38 _(12.80)	0.18 _(0.00)	56.35 _(2.65)	84.81 _(0.74)	93.51 _(0.24)	14.45 _(0.22)	0.370 _(0.003)	55.22 _(0.38)	35.51 _(0.63)	5.66 _(0.20)
Single-ALiBi	49.78 _(0.34)	49.70 _(1.05)	42.62 _(6.15)	38.84 _(1.01)	28.27 _(0.91)	66.15 _(2.92)	88.18 _(0.57)	73.62 _(14.68)	69.04 _(3.30)	0.347 _(0.002)	51.68 _(0.50)	22.27 _(0.22)	5.66 _(0.87)
Single-RoPE	39.20 _(0.82)	51.64 _(0.38)	15.72 _(3.01)	10.15 _(0.07)	0.00 _(0.00)	33.34 _(1.17)	38.71 _(2.11)	65.05 _(6.64)	1.36 _(22.24)	0.462 _(0.000)	14.59 _(0.20)	21.11 _(0.08)	4.89 _(0.09)
6mer-APE	16.24 _(0.54)	55.65 _(23.35)	43.21 _(0.83)	12.24 _(0.90)	13.92 _(0.86)	52.59 _(6.66)	87.63 _(0.91)	91.71 _(0.93)	67.75 _(0.89)	0.420 _(0.009)	51.39 _(0.46)	9.99 _(2.58)	4.21 _(1.08)
6mer-ALiBi	13.99 _(0.35)	28.45 _(0.44)	39.48 _(2.99)	11.49 _(0.39)	22.82 _(0.83)	58.93 _(0.09)	87.49 _(0.86)	60.82 _(2.72)	69.01 _(1.04)	0.417 _(0.001)	48.26 _(0.44)	9.70 _(4.31)	4.64 _(0.88)
6mer-RoPE	22.18 _(0.90)	37.95 _(0.18)	36.93 _(1.02)	12.89 _(0.41)	7.46 _(0.65)	45.71 _(0.30)	86.55 _(0.62)	64.41 _(5.70)	66.38 _(0.90)	0.435 _(0.005)	35.17 _(0.46)	9.91 _(2.12)	5.53 _(1.03)

5.4 Component Analysis of RNA language models

In Tab 4 , we study the language model component effect from tokenizer and positional encoding.

The single nucleotide tokenizer is a powerful tool for RNA language models. As shown in Tab 11, it achieves the best performance on 11 out of 13 tasks, significantly outperforming other tokenizers. BPE and non-overlapping tokenizers are generally ineffective at the nucleotide level, as they lose precision from overlapping. Similarly, the 6mer approach adds local information before the individual tokens, potentially introducing redundancy. We argue that the single nucleotide tokenizer can learn global information, including surrounding context, through self-attention mechanisms. Thus, using

the single nucleotide tokenizer is sufficient, and future work should focus on designing interpretable tokenizers based on single nucleotide units [69, 42].

ALiBi is better for RNA sequences understanding. For tasks involving shorter sequences, the specific advantages of RoPE or other complex encoding schemes may not be fully realized. RoPE, which is highly effective in long sequences due to its rotational component that maintains relative positioning across long distances, might not provide significant benefits over simpler methods like ALiBi in shorter sequences. Moreover, ALiBi linearly biases the attention scores based on relative positions, which helps the model better generalize across different sequence lengths.

Table 5: Comparison of GPU days and the number of tasks (total: 13) where BEACON-B demonstrates performance superiority over other methods.

Model	BEACON -B	BEACON -B512	RNA-FM	SpliceBERT -H510	SpliceBERT -MS510	SpliceBERT -MS1024	UTRBERT -3mer	UTRBERT -4mer	UTRBERT -5mer	UTRBERT -6mer
# GPU Days	1.3*8=10.4	0.895*4=3.58	30*8=240	7*8=56	7*8=56	7*8+3*4=68	38*4=152	38*4=152	38*4=152	38*4=152
# Tasks where BEACON-B Performs Better	-	-	6	6	8	5	7	8	8	6
# Tasks where BEACON-B512 Performs Better	-	-	6	8	9	6	6	7	8	7

5.5 BEACON-B: an Efficient Baseline for RNA Language Models

Based on the above analysis of the different components of the RNA language model, combined with the Tab 4, we use the single nucleotide tokenizer, ALiBi as the positional encoding, and pre-train on filtered human ncRNA sequences from RNACentral [68]. We propose the low-resource and cost-effective BEACON-B (pre-trained on 1026 length seqs) and BEACON-B512 (pre-trained on 512 length seqs and FlashAttn [20]) as a baseline to provide an extremely fast and easy to use open-source pre-training model for subsequent researchers.

Although with very small GPU days (days * GPUs) as the cost of pre-training, BEACON-B can even outperform SOTA pre-trained RNA LMs on some tasks such as **contact map prediction** and **distance map prediction**. Compared with other models that also report pre-training resources, BEACON-B and BEACON-B512 can match or even surpass existing RNA language models in one-to-one comparisons on **almost half of the tasks** listed in Tab 3 and Tab 5, despite being pre-trained with significantly fewer resources. This demonstrates that the insights we obtain from the important components in analysing the RNA language model are vital and that biological motifs and configurations on limited RNA data can be fully explored by utilising such a combination of components in a good way.

6 Conclusions

Summary. In this work, we present BEACON, the first comprehensive RNA benchmark, which encompasses 13 diverse tasks spanning structural analysis, functional studies, and engineering applications. BEACON aims to address the critical gap in standardized evaluation for RNA models. We assess various models, from traditional approaches like CNNs to advanced RNA foundation models, providing insights into their task-specific performances. Additionally, we analyse the vital components of RNA LM from tokenization and positional encoding. Building upon this, we propose BEACON-B, an efficient baseline that incorporates single nucleotide tokenization and ALiBi. BEACON's standardized evaluation framework and the insights provided into RNA modeling components are expected to significantly advance RNA research, facilitating the development of more sophisticated models and enhancing our understanding of RNA's diverse roles in biology.

Limitation & future work. Despite the comprehensiveness of BEACON, it has some limitations for future work. While BEACON includes 13 diverse RNA-related tasks, it may not cover all aspects of RNA biology, necessitating the inclusion of additional tasks and datasets in future versions. The influence of pre-training datasets and hyperparameters on model performance also needs further systematic exploration to optimize configurations for specific RNA tasks. Although BEACON-B serves as an efficient baseline, there is potential for developing more advanced models that leverage RNA's unique structural characteristics. Additionally, BEACON primarily evaluates predictive accuracy, suggesting the need to incorporate metrics like interpretability, computational efficiency, and robustness for a more holistic assessment. Addressing these limitations and exploring new directions will not only advance RNA research but also deepen our understanding of its indispensable roles in genetic regulation, disease pathogenesis, and therapeutic development.

References

- [1] M. Akiyama and Y. Sakakibara. Informative rna base embedding for rna structural alignment and clustering by deep representation learning. *NAR genomics and bioinformatics*, 4(1):lqac012, 2022.
- [2] N. Amin, A. McGrath, and Y.-P. P. Chen. Evaluation of deep learning in non-coding RNA classification. *Nature Machine Intelligence*, 1(5):246–256, May 2019.
- [3] N. M. Angenent-Mari, A. S. Garruss, L. R. Soenksen, G. Church, and J. J. Collins. A deep learning approach to programmable rna switches. *Nature Communications*, 11(1), 10 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18677-1. URL <https://www.osti.gov/biblio/1816589>.
- [4] N. M. Angenent-Mari, A. S. Garruss, L. R. Soenksen, G. Church, and J. J. Collins. A deep learning approach to programmable rna switches. *Nature communications*, 11(1):5057, 2020.
- [5] J. Ansel, E. Yang, H. He, N. Gimelshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski, G. Chauhan, A. Chourdia, W. Constable, A. Desmaison, Z. DeVito, E. Ellison, W. Feng, J. Gong, M. Gschwind, B. Hirsh, S. Huang, K. Kalambarkar, L. Kirsch, M. Lazos, M. Lezcano, Y. Liang, J. Liang, Y. Lu, C. Luk, B. Maher, Y. Pan, C. Pührsch, M. Reso, M. Saroufim, M. Y. Siraichi, H. Suk, M. Suo, P. Tillet, E. Wang, X. Wang, W. Wen, S. Zhang, X. Zhao, K. Zhou, R. Zou, A. Mathews, G. Chanan, P. Wu, and S. Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)*. ACM, Apr. 2024. doi: 10.1145/3620665.3640366. URL <https://pytorch.org/assets/pytorch2-2.pdf>.
- [6] N. Bogard, J. Linder, A. B. Rosenberg, and G. Seelig. A deep neural network for predicting and engineering alternative polyadenylation. *Cell*, 178(1):91–106, 2019.
- [7] J. Cao, E. M. Novoa, Z. Zhang, W. C. W. Chen, D. Liu, G. C. G. Choi, A. S. L. Wong, C. Wehrspau, M. Kellis, and T. K. Lu. High-throughput 5' UTR engineering for enhanced protein production in non-viral gene therapies. *Nat. Commun.*, 12(1):4138, July 2021.
- [8] J. Chen, Z. Hu, S. Sun, Q. Tan, Y. Wang, Q. Yu, L. Zong, L. Hong, J. Xiao, T. Shen, et al. Interpretable rna foundation model from unannotated data for highly accurate rna structure and function predictions. *bioRxiv*, pages 2022–08, 2022.
- [9] K. Chen, T. Litfin, J. Singh, J. Zhan, and Y. Zhou. The master database of all possible rna sequences and its integration with rnacmap for rna homology search. *bioRxiv*, pages 2023–02, 2023.
- [10] K. Chen, Y. Zhou, M. Ding, Y. Wang, Z. Ren, and Y. Yang. Self-supervised learning on millions of pre-mrna sequences improves sequence-based rna splicing prediction. *bioRxiv*, 2023. doi: 10.1101/2023.01.31.526427. URL <https://www.biorxiv.org/content/early/2023/05/09/2023.01.31.526427>.
- [11] K. Chen, T. Litfin, J. Singh, J. Zhan, and Y. Zhou. MARS and RNACmap3: The Master Database of All Possible RNA Sequences Integrated with RNACmap for RNA Homology Search. *Genomics, Proteomics & Bioinformatics*, page qzae018, 03 2024. ISSN 1672-0229. doi: 10.1093/gpbjnl/qzae018. URL <https://doi.org/10.1093/gpbjnl/qzae018>.
- [12] Y. Chu, D. Yu, Y. Li, K. Huang, Y. Shen, L. Cong, J. Zhang, and M. Wang. A 5' utr language model for decoding untranslated regions of mrna and function predictions. *Nature Machine Intelligence*, pages 1–12, 2024.
- [13] G. Chuai, H. Ma, J. Yan, M. Chen, N. Hong, D. Xue, C. Zhou, C. Zhu, K. Chen, B. Duan, et al. Deepcrispr: optimized crispr guide rna design by deep learning. *Genome biology*, 19:1–18, 2018.
- [14] F. Crick. Central dogma of molecular biology. *Nature*, 227(5258):561–563, 1970.
- [15] F. Cunningham, J. E. Allen, J. Allen, J. Alvarez-Jarreta, M. R. Amode, I. M. Armean, O. Austine-Orimoloye, A. G. Azov, I. Barnes, R. Bennett, et al. Ensembl 2022. *Nucleic acids research*, 50(D1):D988–D995, 2022.
- [16] A. E. Dahlberg. The functional role of ribosomal rna in protein synthesis. *Cell*, 57(4):525–529, 1989.
- [17] H. Dalla-Torre, L. Gonzalez, J. Mendoza-Revilla, N. L. Carranza, A. H. Grzywaczewski, F. Oteri, C. Dalago, E. Trop, B. P. de Almeida, H. Sirelkhatim, et al. The nucleotide transformer: Building and evaluating robust foundation models for human genomics. *bioRxiv*, pages 2023–01, 2023.
- [18] P. Danaee, M. Rouches, M. Wiley, D. Deng, L. Huang, and D. Hendrix. bprna: large-scale automated annotation and analysis of rna secondary structure. *Nucleic acids research*, 46(11):5381–5394, 2018.

- [19] T. Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023.
- [20] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022.
- [21] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [22] H.-C. Duan, Y. Wang, and G. Jia. Dynamic and reversible rna n6-methyladenosine methylation. *Wiley Interdisciplinary Reviews: RNA*, 10(1):e1507, 2019.
- [23] W. Falcon and The PyTorch Lightning team. PyTorch Lightning, Mar. 2019. URL <https://github.com/Lightning-AI/lightning>.
- [24] A. Fiannaca, M. La Rosa, L. La Paglia, R. Rizzo, and A. Urso. nrc: non-coding rna classifier based on structural features. *BioData mining*, 10:1–18, 2017.
- [25] L. Fu, Y. Cao, J. Wu, Q. Peng, Q. Nie, and X. Xie. Ufold: fast and accurate rna secondary structure prediction with deep learning. *Nucleic acids research*, 50(3):e14–e14, 2022.
- [26] Z. Gao, C. Tan, Y. Zhang, X. Chen, L. Wu, and S. Z. Li. Proteininvbench: Benchmarking protein inverse folding on diverse tasks, models, and metrics. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=bqXduvuW5E>.
- [27] W. Gilbert. Origin of life: The RNA world. *Nature*, 319(6055):618–618, Feb. 1986.
- [28] J. Gong, K. Xu, Z. Ma, Z. J. Lu, and Q. C. Zhang. A deep learning method for recovering missing signals in transcriptome-wide rna structure profiles from probing experiments. *Nature Machine Intelligence*, 3(11):995–1006, 2021.
- [29] K. Grešová, V. Martinek, D. Čechák, P. Šimeček, and P. Alexiou. Genomic benchmarks: a collection of datasets for genomic sequence classification. *BMC Genom. Data*, 24(1):25, May 2023.
- [30] M. Haeussler, A. S. Zweig, C. Tyner, M. L. Speir, K. R. Rosenbloom, B. J. Raney, C. M. Lee, B. T. Lee, A. S. Hinrichs, J. N. Gonzalez, et al. The ucsc genome browser database: 2019 update. *Nucleic acids research*, 47(D1):D853–D858, 2019.
- [31] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, Sept. 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [32] J. Harrow, A. Frankish, J. M. Gonzalez, E. Tapanari, M. Diekhans, F. Kokocinski, B. L. Aken, D. Barrell, A. Zadissa, S. Searle, et al. Gencode: the reference human genome annotation for the encode project. *Genome research*, 22(9):1760–1774, 2012.
- [33] P. D. Hsu, D. A. Scott, J. A. Weinstein, F. A. Ran, S. Konermann, V. Agarwala, Y. Li, E. J. Fine, X. Wu, O. Shalem, et al. Dna targeting specificity of rna-guided cas9 nucleases. *Nature biotechnology*, 31(9):827–832, 2013.
- [34] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- [35] K. Jaganathan, S. Kyriazopoulou Panagiotopoulou, J. F. McRae, S. F. Darbandi, D. Knowles, Y. I. Li, J. A. Kosmicki, J. Arbelaez, W. Cui, G. B. Schwartz, E. D. Chow, E. Kanterakis, H. Gao, A. Kia, S. Batzoglou, S. J. Sanders, and K. K.-H. Farh. Predicting splicing from primary sequence with deep learning. *Cell*, 176(3):535–548.e24, Jan. 2019.
- [36] Y. Ji, Z. Zhou, H. Liu, and R. V. Davuluri. Dnabert: pre-trained bidirectional encoder representations from transformers model for dna-language in genome. *Bioinformatics*, 37(15):2112–2120, 2021.
- [37] I. Kalvari, E. P. Nawrocki, N. Ontiveros-Palacios, J. Argasinska, K. Lamkiewicz, M. Marz, S. Griffiths-Jones, C. Toffano-Nioche, D. Gautheret, Z. Weinberg, et al. Rfam 14: expanded coverage of metagenomic, viral and microrna families. *Nucleic Acids Research*, 49(D1):D192–D200, 2021.

- [38] A. Karollus, Ž. Avsec, and J. Gagneur. Predicting mean ribosome load for 5'utr of any length using deep learning. *PLoS computational biology*, 17(5):e1008982, 2021.
- [39] N. B. Leontis and C. L. Zirbel. Nonredundant 3d structure datasets for rna knowledge extraction and benchmarking. *RNA 3D structure analysis and prediction*, pages 281–298, 2012.
- [40] K. Leppek, R. Das, and M. Barna. Functional 5' utr mrna structures in eukaryotic translation regulation and how to find them. *Nature reviews Molecular cell biology*, 19(3):158–174, 2018.
- [41] J. Li, Y. Huang, X. Yang, Y. Zhou, and Y. Zhou. Rnam5cfinder: a web-server for predicting rna 5-methylcytosine (m5c) sites based on random forest. *Scientific Reports*, 8(1):17299, 2018.
- [42] C. Liang, W. Bai, L. Qiao, Y. Ren, J. Sun, P. Ye, H. Yan, X. Ma, W. Zuo, and W. Ouyang. Rethinking the bert-like pretraining for dna sequences. *arXiv preprint arXiv:2310.07644*, 2023.
- [43] R. Lorenz, S. H. Bernhart, C. H. Zu Siederdissen, H. Tafer, C. Flamm, P. F. Stadler, and I. L. Hofacker. ViennaRNA package 2.0. *Algorithms for molecular biology*, 6(1):26, 2011.
- [44] F. I. Marin, F. Teufel, M. Horrender, D. Madsen, D. Pultz, O. Winther, and W. Boomsma. Bend: Benchmarking dna language models on biologically meaningful tasks. *arXiv preprint arXiv:2311.12570*, 2023.
- [45] F. I. Marin, F. Teufel, M. Horlacher, D. Madsen, D. Pultz, O. Winther, and W. Boomsma. BEND: Benchmarking DNA language models on biologically meaningful tasks. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=uKB4cFNQFg>.
- [46] H. Ning, G. Liu, L. Li, Q. Liu, H. Huang, and Z. Xie. Rational design of microrna-responsive switch for programmable translational control in mammalian cells. *Nature Communications*, 14(1):7193, 2023.
- [47] M. Nirenberg and P. Leder. Rna codewords and protein synthesis: The effect of trinucleotides upon the binding of srna to ribosomes. *Science*, 145(3639):1399–1407, 1964.
- [48] P. Notin, A. W. Kollasch, D. Ritter, L. V. Niekerk, S. Paul, H. Spinner, N. J. Rollins, A. Shaw, R. Orenbuch, R. Weitzman, J. Frazer, M. Dias, D. Franceschi, Y. Gal, and D. S. Marks. Proteingym: Large-scale benchmarks for protein fitness prediction and design. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=URoZHqAohf>.
- [49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [50] O. Press, N. A. Smith, and M. Lewis. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021.
- [51] R. Rao, N. Bhattacharya, N. Thomas, Y. Duan, P. Chen, J. Canny, P. Abbeel, and Y. Song. Evaluating protein transfer learning with tape. *Advances in neural information processing systems*, 32, 2019.
- [52] R. M. Rao, J. Liu, R. Verkuil, J. Meier, J. Canny, P. Abbeel, T. Sercu, and A. Rives. Msa transformer. In *International Conference on Machine Learning*, pages 8844–8856. PMLR, 2021.
- [53] Reppic. Neighborhood attention architecture. <https://www.kaggle.com/code/reppic/neighborhood-attention-architecture/notebook>, 2020. Accessed: 2024-05-26.
- [54] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [55] E. Rossi, F. Monti, M. Bronstein, and P. Liò. ncrna classification with graph convolutional networks. *arXiv preprint arXiv:1905.06515*, 2019.
- [56] I. A. Roundtree, M. E. Evans, T. Pan, and C. He. Dynamic RNA modifications in gene expression regulation. *Cell*, 169(7):1187–1200, June 2017.
- [57] F. Runge, K. Farid, J. K. Franke, and F. Hutter. Rnabench: A comprehensive library for in silico rna modelling. *bioRxiv*, 2024. doi: 10.1101/2024.01.09.574794. URL <https://www.biorxiv.org/cont/ent/early/2024/01/09/2024.01.09.574794>.

- [58] P. J. Sample, B. Wang, D. W. Reid, V. Presnyak, I. J. McFadyen, D. R. Morris, and G. Seelig. Human 5'utr design and variant effect prediction from a massively parallel translation assay. *Nature biotechnology*, 37(7):803–809, 2019.
- [59] P. J. Sample, B. Wang, D. W. Reid, V. Presnyak, I. J. McFadyen, D. R. Morris, and G. Seelig. Human 5'utr design and variant effect prediction from a massively parallel translation assay. *Nature biotechnology*, 37(7):803–809, 2019.
- [60] P. J. Sample, B. Wang, D. W. Reid, V. Presnyak, I. J. McFadyen, D. R. Morris, and G. Seelig. Human 5'UTR design and variant effect prediction from a massively parallel translation assay. *Nat. Biotechnol.*, 37(7):803–809, July 2019.
- [61] R. Sennrich, B. Haddow, and A. Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.
- [62] A. Shanehsazzadeh, D. Belanger, and D. Dohan. Is transfer learning necessary for protein landscape prediction? *arXiv preprint arXiv:2011.03443*, 2020.
- [63] Z. Song, D. Huang, B. Song, K. Chen, Y. Song, G. Liu, J. Su, J. P. d. Magalhães, D. J. Rigden, and J. Meng. Attention-based multi-label neural networks for integrated prediction and interpretation of twelve widely occurring rna modifications. *Nature communications*, 12(1):4011, 2021.
- [64] B. Strobel, M. Spöring, H. Klein, D. Blazevic, W. Rust, S. Sayols, J. S. Hartig, and S. Kreuz. High-throughput identification of synthetic riboswitches by barcode-free amplicon-sequencing in human cells. *Nature Communications*, 11(1):714, 2020.
- [65] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, page 127063, 2023.
- [66] L. Sun, H. Liu, L. Zhang, and J. Meng. Incrscan-svm: a tool for predicting long non-coding rnas using support vector machine. *PloS one*, 10(10):e0139654, 2015.
- [67] S. Sun, W. Wang, Z. Peng, and J. Yang. Rna inter-nucleotide 3d closeness prediction by deep residual neural networks. *Bioinformatics*, 37(8):1093–1098, 2021.
- [68] B. A. Sweeney, A. I. Petrov, C. E. Ribas, R. D. Finn, A. Bateman, M. Szymanski, W. M. Karlowski, S. E. Seemann, J. Gorodkin, J. J. Cannone, et al. Rnacentral 2021: secondary structure integration, improved sequence search and new member databases. *Nucleic Acids Research*, 49(D1), 2020.
- [69] M. H. Vu, R. Akbar, P. A. Robert, B. Swiatczak, G. K. Sandve, V. Greiff, and D. T. T. Haug. Linguistically inspired roadmap for building biologically reliable protein language models. *Nature Machine Intelligence*, 5(5):485–496, 2023.
- [70] M. L. Waskom. Seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021. doi: 10.21105/joss.03021. URL <https://doi.org/10.21105/joss.03021>.
- [71] H. K. Wayment-Steele, W. Kladwang, A. M. Watkins, D. S. Kim, B. Tunguz, W. Reade, M. Demkin, J. Romano, R. Wellington-Oguri, J. J. Nicol, et al. Deep learning models for predicting rna degradation via dual crowdsourcing. *Nature Machine Intelligence*, 4(12):1174–1184, 2022.
- [72] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [73] T. W. Wong and D. A. Clayton. Dna primase of human mitochondria is associated with structural rna that is essential for enzymatic activity. *Cell*, 45(6):817–825, 1986.
- [74] H. Xu, T. Xiao, C.-H. Chen, W. Li, C. A. Meyer, Q. Wu, D. Wu, L. Cong, F. Zhang, J. S. Liu, et al. Sequence determinants of improved crispr sgrna design. *Genome research*, 25(8):1147–1157, 2015.
- [75] M. Xu, Z. Zhang, J. Lu, Z. Zhu, Y. Zhang, M. Chang, R. Liu, and J. Tang. Peer: a comprehensive and multi-task benchmark for protein sequence understanding. *Advances in Neural Information Processing Systems*, 35:35156–35173, 2022.
- [76] Y. Yang, G. Li, K. Pang, W. Cao, X. Li, and Z. Zhang. Deciphering 3' utr mediated gene regulation using interpretable deep representation learning. *bioRxiv*, 2023. doi: 10.1101/2023.09.08.556883. URL <https://www.biorxiv.org/content/early/2023/09/12/2023.09.08.556883>.

- [77] F. Yuan, W. Hankey, E. J. Wagner, W. Li, and Q. Wang. Alternative polyadenylation of mrna and its role in cancer. *Genes & Diseases*, 8(1):61–72, 2021.
- [78] Y. Zhang, M. Lang, J. Jiang, Z. Gao, F. Xu, T. Litfin, K. Chen, J. Singh, X. Huang, G. Song, Y. Tian, J. Zhan, J. Chen, and Y. Zhou. Multiple sequence alignment-based RNA language model and its application to structural inference. *Nucleic Acids Research*, 52(1):e3–e3, 11 2023. ISSN 0305-1048. doi: 10.1093/nar/gkad1031. URL <https://doi.org/10.1093/nar/gkad1031>.
- [79] Z. Zhou, Y. Ji, W. Li, P. Dutta, R. Davuluri, and H. Liu. Dnabert-2: Efficient foundation model and benchmark for multi-species genome. *arXiv preprint arXiv:2306.15006*, 2023.
- [80] Y. Zhu, L. Zhu, X. Wang, and H. Jin. RNA-based therapeutics: an overview and prospectus. *Cell Death Dis.*, 13(7):644, July 2022.

A Appendix / supplemental material

Optionally include supplemental material (complete proofs, additional experiments and plots) in the appendix. All such materials **SHOULD be included in the main submission**.

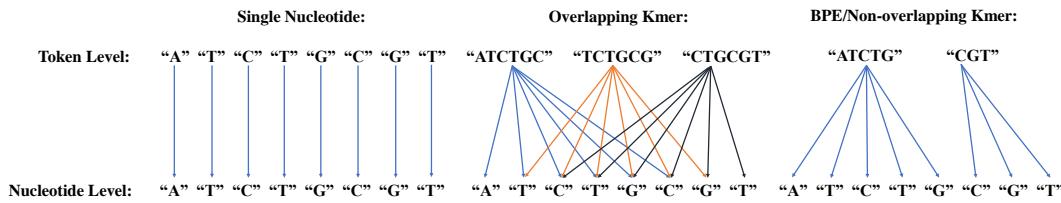


Figure 2: Derivation of nucleotide-level representations. In single nucleotide tokenization, a token directly corresponds to a nucleotide, thus the representations are identical. For overlapping Kmer tokenization, the nucleotide representation is the averaged representation of tokens covering it. In Byte-Pair Encoding (BPE) and non-overlapping K-mer tokenization, the representation is derived from the token covering it.

A.1 Experimental settings for tasks

A.1.1 Most of the tasks

Most of the tasks are trained using the same training settings shown in the Tab 6. These tasks include structural score imputation, splice site prediction, APA isoform prediction, non-coding RNA function classification, modification prediction, programmable RNA switches, CRISPR on-target prediction and CRISPR off-target prediction.

Table 6: Configuration settings for most of the tasks training

config	value
optimizer	AdamW
optimizer epsilon	1e-8
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.999$
weight decay	0.01
learning rate sch.	linear decay
learning rate	[1e-5, 5e-3]
warmup steps	50
epochs	30
batch size	32
gradient accumulation	1
dtype	float16

In particular, for the structural score imputation task, the input is the sequence accompanied by structural scores. The sequence is fed to the model and undergoes the transformation to the nucleotide-level representation. We concatenate it with the MLP-passed structural scores, and then use the regression header to get imputation scores.

For CRISPR off-target, the input is two sequences including sgRNA and target sequences. We feed them through the same model separately, then concat them together. Finally, we use the regression header to get the predicted value of off-target.

A.1.2 Vaccine degradation prediction task

Vaccine degradation prediction is trained using the settings shown in the Tab 7.

Table 7: Configuration settings for vaccine degradation prediction training

config	value
optimizer	AdamW
optimizer epsilon	1e-8
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.999$
weight decay	0.01
learning rate sch.	linear decay
learning rate	[1e-5, 5e-3]
warmup steps	50
epochs	100
batch size	32
gradient accumulation	1
dtype	float16

A.1.3 Nucleotide-nucleotide level tasks

Nucleotide-nucleotide level tasks are trained using the settings shown in the Tab 8. In addition, the representation also follows Fig 2

Table 8: Configuration settings for nucleotide-nucleotide level tasks training

config	value
optimizer	Adam
optimizer epsilon	1e-8
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.999$
learning rate sch.	cosine decay
learning rate	[1e-5, 5e-3]
warmup epochs	1
epochs	100
batch size	1
gradient accumulation	8
dtype	float16

A.2 Methods in Benchmark

All pre-trained benchmarked methods use a Masked Language Modeling (MLM) objective.

In the MLM task, a sequence is provided as input, with 15% of its tokens randomly masked. The entire masked sequence is then processed by the model, which is tasked with predicting the original tokens. This approach is analogous to the Cloze task in traditional language modeling.

- 15% of the tokens in the sequence are masked.
- In 80% of the cases, the masked tokens are replaced by a special <mask> token.
- In 10% of the cases, the masked tokens are substituted with a random token different from the original.
- In the remaining 10% of cases, the masked tokens remain unchanged.

A.2.1 RNABERT

Training Objectives RNABERT was pre-trained with two objectives: masked language modeling (MLM) and structural alignment learning (SAL).

For SAL, the model learns to predict the structural alignment between two RNA sequences. It achieves this by being trained to predict the alignment score of RNA sequence pairs using the Needleman-Wunsch algorithm.

Training Data The RNABERT model was pre-trained using a subset of 76,237 human ncRNA sequences from RNACentral. The dataset was preprocessed by applying 10 different masking patterns to the 76,237 sequences, resulting in a final dataset comprising 762,370 sequences.

A.2.2 RNA-FM

Training Data The RNA-FM model was pre-trained using data from RNACentral. To ensure the dataset was non-redundant, RNA-FM applied CD-HIT (CD-HIT-EST) with a cut-off at 100% sequence identity, resulting in a final dataset containing 23.7 million unique RNA sequences.

A.2.3 RNA-MSM

Unlike other methods, RNA-MSM utilizes homologous sequences as input to provide additional evolutionary information, similar to MSATransformer [52].

To ensure fairness, homologous sequences were not included in the input during evaluation.

Training Data RNA-MSM was pre-trained using data from Rfam, which includes homologous sequences. To prevent potential overfitting in structural inference, RNA-MSM excluded families with experimentally determined structures, such as ribosomal RNAs, transfer RNAs, and small nuclear RNAs. The final dataset comprises 3,932 RNA families, with a median of 2,184 MSA sequences per family. To augment the number of homologous sequences, RNA-MSM employed an automated pipeline, RNACmap3 [11], for homolog search and sequence alignment.

A.2.4 SpliceBERT

Training Data The SpliceBERT model was pre-trained using messenger RNA precursor sequences obtained from the UCSC Genome Browser.

SpliceBERT gathered reference genomes and gene annotations from the UCSC Genome Browser for 72 vertebrate species. Bedtools getfasta was used to extract pre-mRNA sequences from the reference genomes based on these gene annotations. The resulting pre-mRNA sequences were then utilized for pre-training SpliceBERT. The pre-training dataset comprises 2 million pre-mRNA sequences, with a total length of 65 billion nucleotides.

A.2.5 3UTRBERT

Training Data The 3UTRBERT model was pre-trained using human mRNA transcript sequences obtained from GENCODE.

3UTRBERT collected 108,573 unique human mRNA transcripts from GENCODE, utilizing only the longest transcript for each gene in the pre-training process. To avoid codon constraints in the CDS region and to reduce the complexity of the full mRNA transcripts, only the 3' untranslated regions (3'UTRs) of the mRNA transcripts were used. The average length of the 3'UTRs was 1,227 nucleotides, with a median length of 631 nucleotides. Each 3'UTR sequence was divided into non-overlapping patches of 510 nucleotides, with the remaining sequences padded to the same length.

A.2.6 UTR-LM

Training Objectives In addition to MLM pre-training, UTR-LM employs two additional supervised objectives: Secondary Structure (SS) and Minimum Free Energy (MFE).

Both secondary structure and the MFE value are calculated using ViennaRNA [43]. To prevent information leakage, UTR-LM calculates the secondary structure loss only on the masked positions. The output embedding of the cls token is used by UTR-LM to regress the MFE value.

Training Data The UTR-LM model was pre-trained using 5' UTR sequences sourced from three origins: the Ensembl database, synthetic libraries from Sample et al. [60], and endogenous human 5' UTR data analyzed by Cao et al. [7].

The preprocessing of 5' UTR sequences for UTR-LM involved a 4-step pipeline: First, all coding sequence (CDS) and non-5' UTR fragments were removed from the raw sequences. Second, duplicate

sequences were identified and removed. Third, the sequences were truncated to fit within a range of 30 to 1022 base pairs. Finally, incorrect and low-quality sequences were filtered out.

A.2.7 BEACON-B

Training Data We filter 523,934 human ncRNA sequences from the total ncRNA in the RNACentral database [68] as pre-training data. BEACON-B and BEACON-B512 use normal BERT-base [21] architecture with 12 layers.

The pre-training configs of BEACON-B and BEACON-B512 are shown as Tab 9 and Tab 10

Table 9: Configuration settings for the BEACON-B pre-training

config	value
optimizer	AdamW
optimizer epsilon	1e-6
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.98$
weight decay	0.01
learning rate sch.	linear decay
learning rate	2e-4
warmup steps	10000
steps	80000
batch size	256
gradient accumulation	2
dtype	float16
length	1026
pertaining data	RNACentral Human ncRNA

Table 10: Configuration settings for the BEACON-B512 pre-training

config	value
optimizer	AdamW
optimizer epsilon	1e-6
optimizer momentum	$\beta_1, \beta_2 = 0.9, 0.98$
weight decay	0.01
learning rate sch.	linear decay
learning rate	2e-4
warmup steps	10000
steps	80000
batch size	512
gradient accumulation	1
dtype	float16
length	512
pertaining data	RNACentral Human ncRNA
attention	FlashAttention

A.3 Computational Resources

We fine-tune or train each model from scratch on one task using one NVIDIA A100 40g GPU. We pre-train the simple BEACON-B on 8 A100 GPUs of 80GB for 1.3 days and BEACON-B512 on 4 A100 GPUs of 80GB for 0.895 days.

A.4 Additional results

For the experiments 4 on component analysis of the baseline RNA language model, we further counted the number of top performances for different tokenizers and positional encoding as shown in Tab 11. The effectiveness of Single nucleotide tokenizer and ALiBi can be demonstrated directly.

Table 11: The number of Top2 performance among different tokenizers and positional encodings

Rank	Tokenizer			Positional encoding			
	Single	6mer	Non-overlap	BPE	APE	ALiBi	RoPE
1st	11	0	1	1	5	7	1
2nd	1	6	3	3	8	5	0

A.5 Broader Societary Impacts

This work is dedicated to establishing a robust and versatile benchmark for RNA-related tasks, enhancing the understanding of RNA sequences across diverse applications. Our benchmark, encompassing a variety of RNA tasks, aims to rigorously evaluate the efficacy of different RNA representation covering structural analysis, functional studies, and engineering applications. By doing so, it provides a critical assessment of their potential utility in real-world scenarios, thereby laying a foundational framework for applying deep learning in fields such as medical research and genetics.

However, it is also important to acknowledge the dual-use nature of any powerful technology, including those developed from our benchmark. For instance, the enhanced ability to manipulate RNA sequences might be misused, such as in the creation of adverse viral agents. Moving forward, it is crucial to address these risks. We will develop and implement guidelines for the ethical and safe use of our benchmark in the future, ensuring that it contributes positively to society and does not enable harmful applications.

A.6 Assets

A.6.1 Software and Libraries

The open-source software, and corresponding licenses are presented in Tab. 12. The data, licenses and corresponding URL are presented in Tab. 13.

Table 12: Software used in this work

Asset	License
FlashAttention [20, 19]	BSD-3-Clause
Pytorch [5]	BSD-3-Clause
Pytorch Lightning [23]	Apache-2.0
Huggingface [72]	Apache-2.0
Scikit-Learn [49]	BSD-3-Clause
Numpy [31]	BSD-3-Clause
Matplotlib [34]	Matplotlib License
Seaborn [70]	Apache-2.0

Table 13: Dataset used in this work

Dataset	Sub-dataset	License	URL
bpRNA-1M	CRW	-	https://bprna.cgrb.oregonstate.edu/about.php https://crw-site.chemistry.gatech.edu/
	tmRDB	Research Purpose Only	https://rth.dk/resources/rnp/tmRDB/
	SRPDB	Research Purpose Only	https://rth.dk/resources/rnp/SRPDB/
	tRNADB	Public Domain	
RNAcontact	Rnase P	CC0 1.0	https://rfam.org
	RFam	CC0 1.0	https://www.rcsb.org
	PDB	CC0 1.0	
		Public Domain	https://yanglab.qd.sdu.edu.cn/RNAcontact/
StructImpute		MIT & Non Commercial	https://figshare.com/articles/dataset/A_deep_learning_method_for_recovering_missing_signals_in_transcriptome-wide_RNA_structure_profiles_from_probing_experiments/16606850
SpliceAI		GPLv3	https://github.com/illumina/SpliceAI
APARNET		MIT	https://github.com/johli/aparent
ncRNA		Apache 2.0	https://github.com/bioinformatics-sannio/ncrna-deep
MultiRM	MIT	MIT	https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE114002
Optimus	GEO		
OpenVaccine		Non Commercial	https://www.kaggle.com/competitions/stanford-covid-vaccine/data
ProgrammableRNAswitches		GEO	https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE149225
DeepCRISPR		Apache 2.0	