DynCodon: Multi-objective Dynamic Learning Strategies for Robust Codon Optimization

Anonymous authors
Paper under double-blind review

Abstract

Existing codon optimization methods often blindly pursue theoretical extreme values for evaluation metrics. However, these extreme sequences at the upper boundary of the Top 10% distribution are like unreplicable "geniuses", lacking universal biological adaptability; in contrast, the robust regions of the Top 10% distribution are more like high-performing "high achievers", containing generalizable evolutionary patterns. This paper proposes a codon optimization model guided by a multi-objective dynamic learning strategy, called DynCodon, aiming to achieve robustness comparable to the expression levels of "high achievers". This model is based on a dual BERT encoder-decoder architecture and employs a two-stage learning paradigm: pre-training through large-scale mixed species to capture universal amino acid-codon correspondences; fine-tuning using this multi-objective dynamic learning strategy for target species to achieve synergistic optimization across multiple key metrics. Experimental results show that DynCodon effectively avoids the pitfall of over-optimization towards extreme values, achieving high-precision reconstruction of natural sequence features within the Top 10% robust regions.

1 Introduction

The advent of mRNA-based therapies has revolutionized modern medicine, providing rapid and programmable solutions for vaccines and protein replacement therapies (Pardi et al., 2018; Polack et al., 2020). The efficacy of these therapies hinges on the design of the coding sequences (CDS). Due to the degeneracy of the genetic code, a protein can be encoded by an extremely large number of synonymous mRNA sequences (approximately 3^L synonymous sequences for a protein of length L) (Plotkin & Kudla, 2011). Although these synonymous variants encode the same polypeptide, they exhibit drastically different behaviors in terms of translation efficiency, stability, and immunogenicity (Quax et al., 2015; Presnyak et al., 2015). Therefore, codon optimization, the computational task of selecting the optimal mRNA sequence for a target protein, has become a key challenge in computational biology.

Traditional codon optimization methods primarily rely on heuristic algorithms or linear programming to maximize specific evaluation metrics, such as the Codon Adaptation Index (CAI) (Sharp & Li, 1987), Minimum Free Energy (MFE) (Zuker & Stiegler, 1981), or GC content (Kudla et al., 2006). In recent years, advancements in deep learning have further enhanced these methods by generating sequences using Recurrent Neural Networks (RNNs) or Transformer models, thereby pushing these metrics to theoretical limits (Jain et al., 2023; Fallahpour et al., 2025).

However, we find a key "over-optimization trap" in current methodologies. Existing models often blindly pursue the absolute maximization of evaluation metrics, aiming to reach the theoretical upper bound of the natural sequence distribution (e.g., the Top 1%). We believe that these extreme sequences are similar to "geniuses" in nature that cannot be replicated, rare outliers whose high performance may depend on special, specific conditions that are difficult to generalize. In biological systems, pushing sequences to such extremes often disrupts the delicate balance of translation dynamics, leading to protein misfolding or ribosome arrest (Tuller et al., 2010). In contrast, sequences located in the Top 10% robust regions are analogous to well-behaved "high achievers". These sequences may not necessarily reach the theoretical upper bound of various metrics, but based on generalizable evolutionary patterns, they consistently exhibit superior performance. Therefore, the goal of robust optimization models should shift from creating "geniuses" (theoretical extremes) to learning the

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

distribution of "high achievers" (robust regions) to ensure high expression potential and biological rationality.

To address this challenge, we propose DynCodon, a novel sequence generation framework based on a multi-objective dynamic learning strategy. DynCodon aims to balance conflicting biological objectives, ensuring that the generated sequences reside within robust, high-performance regions, rather than unstable theoretical boundaries. DynCodon employs a dual BERT encoder-decoder architecture to capture bidirectional contextual dependencies in gene sequences. To achieve robust optimization, we introduce a two-stage learning paradigm: **General Knowledge Acquisition:** We pre-train the model using a large-scale mixed-species dataset to capture the general "grammar" of amino acid-codon correspondences. **Target Species-specific Optimization:** We fine-tune the model on the target species using a multi-objective dynamic learning strategy; Unlike static loss functions, this strategy dynamically balances gradients from multiple objectives (e.g., CAI, GC content, etc.) during training; This enables the model to achieve co-optimization, effectively avoiding extreme cases caused by a single metric.

Experimental results demonstrate that DynCodon effectively avoids the drawbacks of over-optimization. By focusing on the Top 10% of robust regions in the natural sequence distribution, our model achieves high-precision reconstruction of natural sequence features, providing a reliable and stable solution for mRNA sequence design. In summary, our main contributions are as follows:

- We point out the limitations of pursuing theoretical extreme values in codon optimization and propose a paradigm shift by modeling the Top 10% robust regions ("high-performance regions") to improve biological adaptability.
- We propose DynCodon, a Dual-BERT-based framework that integrates a multi-objective dynamic learning strategy to collaboratively optimize multiple key metrics under dynamic constraints.
- Extensive experiments demonstrate that DynCodon outperforms existing state-of-the-art methods in generating robust mRNA sequences, avoiding the pitfalls of over-optimization while maintaining high predictive expression efficiency.

2 Related Work

2.1 Heuristic and Statistical Methods

Early research focused primarily on simulating Codon Usage Bias (CUB) of highly expressed genes in host organisms. The most representative metric is CAI proposed by Sharp et al. (Sharp & Li, 1987), which quantifies the similarity between gene synonym codon usage and the reference set. Based on this, many heuristic algorithms have been developed to maximize CAI or Optimal Codon Frequency. Tools such as JCat (Grote et al., 2005) and GenScript (Burgess-Brown et al., 2008) employ greedy or genetic algorithms to replace rare codons with common codons. However, these methods typically treat codons as independent units or restrict the context to a short window (e.g., codon pairs), largely ignoring the overall structural stability of mRNA and long-range epistatic interactions (Mauger et al., 2019). Furthermore, maximizing CAI often results in extremely high GC content, which can lead to ribosome arrest or shortened mRNA half-life due to rigid secondary structures (Kudla et al., 2006).

2.2 Deep Learning for Sequence Generation

With the development of natural language processing, codon optimization is increasingly being viewed as a sequence-to-sequence translation task. For example, ICOR (Jain et al., 2023) utilizes a bidirectional Long Short-Term Memory network (LSTM) to learn codon context from large-scale genomic data, demonstrating superior performance to heuristic methods in predicting expression levels. Similarly, Goldman et al. (Yang et al., 2019) proposed an LSTM-based converter for mapping protein sequences to DNA. Despite improvements in RNNs, they still face the vanishing gradient problem, making it difficult to capture long-range dependencies across sequences spanning thousands of nucleotides, which is crucial for modeling mRNA folding structures. In recent years, Transformer-based models, such as CodonTransformer (Fallahpour et al., 2025), and GEMORNA (Zhang et al., 2025), have emerged. CodonTransformer utilizes a taxonomically conscious Transformer architecture, pre-trained on millions of coding sequences, to generate host-specific DNA with natural codon usage

patterns by capturing long-range contextual dependencies. GEMORNA treats mRNA design as a complex language generation task, using deep generative models to create coding sequences and untranslated regions that maximize translation capability and stability through zero-shot inference. While these models excel in feature extraction, they often fall into the "over-optimization trap" mentioned in the introduction. By blindly maximizing the likelihood of the training distribution or a single structural metric, they tend to generate sequences at theoretical extremes that may lack the robustness required for practical translation (Tuller et al., 2010).

2.3 Multi-objective Optimization in Bioinformatics

Biological systems rarely optimize for a single trait; they evolve under multiple conflicting objective constraints (e.g., translation speed, accuracy, and stability). Traditional multi-objective optimization in this field often relies on evolutionary algorithms, such as NSGA-II (Deb et al., 2002), which are computationally expensive and difficult to scale to long sequences. In the field of neural network generation, Reinforcement Learning (RL) has been used to optimize discrete metrics (e.g., QED in drug discovery) (Popova et al., 2018). However, applying RL to codon optimization faces challenges of reward sparsity and convergence instability. DynCodon differs from these approaches by introducing a dynamic learning strategy. Instead of using static scalarization or complex RL loops, we dynamically balance the gradients of multiple objectives during the fine-tuning phase. This allows the model to move efficiently along the Pareto front and find "robust regions" in the solution space, rather than drifting to unstable extremes.

3 Methodology

This section introduces ***DynCodon***, a sequence generation framework with a ***Multi-objective Dynamic Learning Strategy*** designed to map amino acid sequences to optimized mRNA sequences. As shown in Fig. 1, DynCodon consists of a ***Dual-BERT Encoder-decoder Architecture*** based on a ***Two-stage Learning Paradigm***.

3.1 Problem Formulation

Let $\mathbf{A} = (a_1, a_2, \ldots, a_L)$ represent the input amino acid sequence of length L, and $\mathbf{C} = (c_1, c_2, \ldots, c_L)$ represent the corresponding synonymous mRNA sequence, where each c_i is a codon encoding a_i. By convention, a special symbol, such as "*" or "_", is usually appended to the end of the amino acid sequence to correspond to the stop codon of the mRNA sequence. Thus, if the length of the amino acid sequence with the appended special symbol is L, then the length of the corresponding mRNA codon sequence containing the stop codon is also L. In the model, our goal is not to maximize biological metrics to theoretical limits, but to generate sequences $\mathbf{C}'$ whose properties are consistent with the "gold standard" derived from naturally evolutionary robust regions. We define a set of metric functions $\mathcal{M} = \{f_{CAI}, f_{GC}, f_{CFD}\}$ and their corresponding gold standard values or intervals $\mathcal{G} = \{G_{CAI}, G_{GC}, G_{CFD}\}$. The model employs an error-driven adaptive learning rate strategy to optimize the model, which is the ***Multi-objective Dynamic Learning Strategy***. A high learning rate is used in the large error range to accelerate parameter space traversal and localization; a low learning rate is used in the small error range to achieve refined exploration within the target domain. This significantly reduces the metric bias between the generated sequence and the standard value while maintaining translation accuracy.

3.2 Dual-BERT Architecture

To capture long-range dependencies in gene sequences, DynCodon employs a dual-BERT encoder-decoder architecture: ***Protein-BERT Encoder:*** Processes the amino acid sequence $\mathbf{A}$ to generate a contextualized latent representation matrix $\bar{\mathbf{A}}$; ***Codon-BERT Decoder***: Predicts the conditional probability of codon c_i based on the protein contextualized latent representation $\bar{\mathbf{A}}$ and the known preceding codons $\mathbf{C}_{<i}$. The Protein-BERT Encoder and Codon-BERT Decoder are connected using cross-attention mechanism. Both the encoder and decoder have a maximum context window of 128. For gene sequences longer than 128, a sliding window approach can be used. Setting the context window to 128 significantly reduces the number of model parameters, directly improving the difficulty of model optimization and reducing the amount of training data.

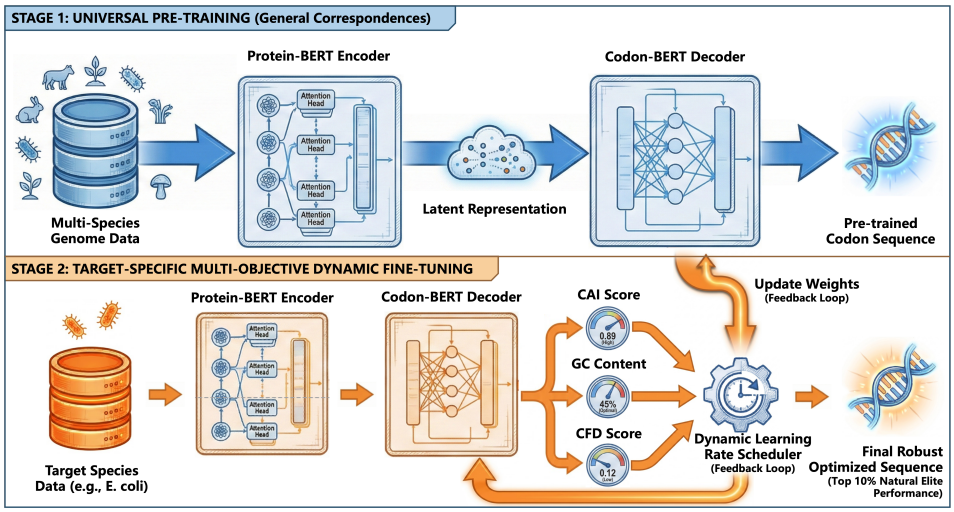

Figure 1: DynCodon: A codon optimization model with multi-objective dynamic learning strategies under a two-stage learning paradigm. The model consists of a protein encoder that encodes amino acid sequences and a codon decoder that decodes codon sequences. They are connected using cross-attention to achieve sequence-to-sequence learning. Both the protein encoder and codon decoder are implemented using the BERT model, named Protein-BERT Encoder and Codon-BERT Decoder, respectively. The model training process is divided into two stages: a first stage of multi-species hybrid pre-training and a second stage of multi-objective dynamic fine-tuning targeting the specific species. Pre-training aims to learn the universal relationships between amino acids and codons, while fine-tuning adjusts the model to align with the unique preferences of the target species. The multi-objective dynamic learning strategy proposed in this paper is used in the multi-objective dynamic fine-tuning stage to enable the model to learn and possess robust codon optimization capabilities.

3.3 Two-stage Learning Paradigm

Stage 1: Universal Pre-training. We first pre-train DynCodon on a large-scale mixed-species dataset. The goal is to capture the general "grammar" of the genetic code. We minimize the standard cross-entropy loss to learn the fundamental mapping between amino acids and codons:

$$\mathcal{L}_{pre} = -\sum_{i=1}^{L} \log P(c_i | \mathbf{C}_{<i}, \bar{\mathbf{A}}_{\leq i}; \theta), \tag{1}$$

where θ represents the parameter to be optimized in the model. We use an autoregressive language model training strategy to optimize the model, which is to generate the i-th codon c_i based on the first $i-1$ predicted codons $\mathbf{C}_{<i}$ and the first i amino acids $\bar{\mathbf{A}}_{\leq i}$.

Stage 2: Multi-objective Dynamic Fine-tuning. At this stage, we fine-tune the model for the target species. To avoid the "over-optimization trap", we introduce a metric-based adaptive learning rate strategy, namely a multi-objective dynamic learning strategy. Instead of using a fixed learning rate, we dynamically adjust the learning rate η based on the real-time deviation between the multi-metric evaluation values of the generated sequence and their gold standards, and then combine this with the error between the generated sequence and the target sequence to achieve directional dynamic optimization of the model parameters. The **Multi-objective Dynamic Learning Strategy** specifically includes the following steps:

1. Metric Divergence Calculation: In each training step t, the model generates a prediction sequence $\mathbf{C}'_t$. We calculate the total metric deviation (Δ_t) between the predicted metrics and the gold standard:

$$\Delta_t = \sum_{k \in \{CAI, GC, CFD\}} \lambda_k \cdot |f_k(\mathbf{C}'_t) - G_k| \tag{2}$$

where G_k represents the average of the k-th metric in the Top 10% robust region ("high achiever" standard), and λ_k is the normalization factor for each metric.

2. Dynamic Learning Rate Modulation: We believe the optimization speed should be adjusted based on the "distance" from the robust region. If the generated sequence is far from the golden ratio (large Δ_t), the model requires a larger learning rate to escape local optima; as the model approaches the robust region (smaller Δ_t), the learning rate should gradually decrease to ensure model stability and prevent overshooting to extreme values. We denote the dynamic learning rate η_t as:

$$\eta_t = \eta_{base} \cdot \Psi(\Delta_t) \tag{3}$$

where η_{base} is the initial learning rate, and $\Psi(\cdot)$ is a non-linear modulation function defined as:

$$\Psi(\Delta_t) = 1 + \alpha \cdot \tanh(\beta \cdot \Delta_t) \tag{4}$$

where α and β are hyperparameters controlling the adjustment sensitivity.

3. Optimization Objective: The final optimization minimizes the combinatorial loss by dynamically adjusting η_t:

$$\theta_{t+1} \leftarrow \theta_t - \eta_t \cdot \nabla_\theta(\mathcal{L}_{fine} + \gamma \cdot \Delta_t) \tag{5}$$

where γ represents the balance coefficient or regularization weight of the loss term, and

$$\mathcal{L}_{fine} = -\sum_{i=1}^{L} \log P(c_i|\mathbf{C}_{<i}, \bar{\mathbf{A}}_{\leq i}; \theta). \tag{6}$$

By directly relating the learning rate to the metric of divergence, DynCodon can dynamically accelerate or decelerate its learning process. This mechanism ensures that the model converges precisely to a robust "high achiever" region, effectively avoiding the generation of unstable, theoretically extreme sequences.

4 Experimental Results & Analysis

4.1 Experimental Setup:

Datasets: We utilize the comprehensive dataset published by CodonTransformer (Fallahpour et al., 2025), which covers a wide range of biological species.

Pre-training dataset: To equip the model with general codon knowledge, we used a complete training set containing 164 organisms. This large-scale corpus enabled the encoder to capture universal amino acid-to-codon mapping rules.

Fine-tuning and test datasets: We selected eight representative organisms from the pre-training set to fine-tune the model, enabling it to learn codon preferences unique to the target species. Correspondingly, we used data from eight organisms in the test set used by CodonTransformer as our test set. These organisms cover different biological fields and have different GC content preferences: ***Mus musculus***, ***Drosophila melanogaster***, ***Escherichia coli general***, ***Nicotiana tabacum***, ***Nicotiana tabacum chloroplast***, ***Arabidopsis thaliana***, ***Bacillus subtilis***, ***Thermococcus barophilus MPT***.

We strictly followed the data partitioning protocol in the original CodonTransformer paper, using their defined training, validation, and test sets to partition these 8 organisms to ensure a fair comparison.

Baselines: We compare DynCodon with five baseline methods, ranging from heuristics to state-of-the-art deep learning models: ***URC*** (Uniform Random Choice): Randomly selects synonymous codons. Used as a lower bound. ***BFC*** (Background Frequency Choice): Selects codons based on their global usage frequency within the host organism. Represents a traditional statistical method. ***CodonTransformer:*** This is a Transformer-based model, pre-trained on the same 164 organisms (Fallahpour et al., 2025). It represents a "theoretically extreme" approach ("genius") and is **the baseline model for the primary comparison**. ***LinearDesign:*** A dynamic programming algorithm that globally optimizes mRNA stability (MFE) and computational efficiency (CAI) (Zhang et al., 2023). ***GEMORNA:*** A deep generative model designed to optimize translational ability and stability (Zhang et al., 2025).

Implementation Details: DynCodon is implemented using PyTorch. The dual BERT architecture consists of a 6-layer BERT model as the encoder and a 6-layer BERT model as the decoder, both

with 8 attention heads and 512-dimensional hidden layers. ***Other unmentioned model parameter settings will be provided in the code that will be released later.***

Hyperparameters: We use the AdamW optimizer with default values for the optimizer parameters. The base learning rate η_{base} is set to $1e^{-4}$. For dynamic learning rate calculation, the gold standard G_k is calculated by averaging the natural sequences of the Top 10% robust regions in the training set according to the CSI metric. The balance coefficient γ is empirically set to 0.5.

Hardware: All models were trained on 8 NVIDIA A100 GPUs with a batch size of 16.

Evaluation Metrics: The evaluation metrics and calculation methods used in this paper are identical to those described in the CodonTransformer paper. The ***Codon Similarity Index (CSI)***, ***Codon Frequency Distribution (CFD)***, and ***GC Content (%GC)*** metrics all utilize the Python toolkit provided with the original paper. ***Sequence Similarity Analysis***, ***Jaccard Index***, ***Dynamic Time Warping (DTW)***, and ***%MinMax*** were implemented by coding ourselves based on the formal definitions given in the original paper. ***Minimum Free Energy (MFE)*** calculation for the mRNA sequence was performed using an open-source Python toolkit.

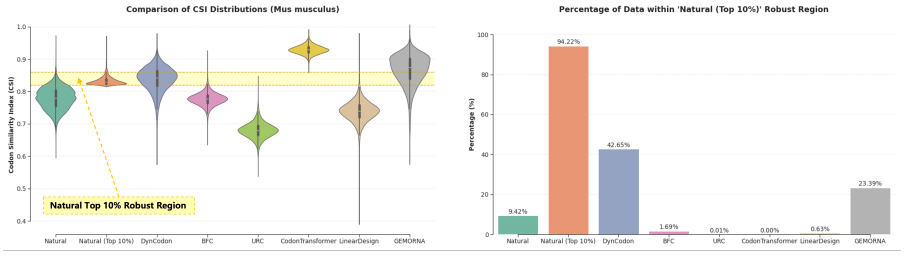

Figure 2: Experimental results for a typical species, ***Mus musculus***, under the CSI evaluation metric. Subfigure 1 shows the violin plots of the comparative models, where ***Natural*** represents the wild-type gene sequence, and ***Natural (Top 10%)*** represents the Top 10% of wild-type gene sequences according to the CSI evaluation metric; ***BFC*** and ***URC*** represent heuristic methods; ***LinearDesign*** and ***GEMORNA*** are the current state-of-the-art (SOTA) models; ***CodonTransformer*** represents an overfitted model and is the main comparative model. Subfigure 2 shows the proportion of the comparative models in the ***Natural Top 10% Robust Region***.

4.2 General Evaluation Metrics Analysis:

Codon Similarity Index (CSI), Codon Frequency Distribution (CFD), and GC Content (%GC) are commonly used evaluation metrics in codon optimization tasks. Based on these metrics, we can comprehensively and quantitatively analyze the performance of the optimization algorithm across different dimensions:

CSI is a key evaluation metric measuring the consistency between the codon usage pattern of exogenous genes and the host's preference for highly expressed genes. The CSI performance shows that the codon sequence optimized by DynCodon achieves higher consistency with the natural sequences of the Top 10% robust regions compared to those generated by the comparative models. As shown in Fig. 2, the CSI distribution in a typical species (Mus musculus) shows that the codon sequences optimized by DynCodon account for the largest proportion (42.65%) in the ***Natural Top 10% Robust Region***; while heuristic methods such as BFC and URC account for only about 1.00%, and the current state-of-the-art models CodonTransformer (0.00%), LinearDesign (0.63%), and GEMORNA (23.39%). The experimental results for the other seven species are shown in Fig. 5 and 6 in the Appendix, all exhibiting a high proportion of the natural Top 10% robust regions. ***This indicates that the codon sequence optimized by DynCodon is more consistent with the host's preference for highly expressed genes, and that DynCodon has learned the most robust and generalizable evolutionary pattern among the host's highly expressed genes, ensuring high expression potential and biological rationality.***

If CSI focuses on "optimization", then CFD emphasizes "balance". This evaluation metric measures whether the frequency of specific codons in an optimized sequence matches the abundance of the

available tRNA library within the host cell. As shown in Fig. 7 and 8 in the Appendix, CFD performance demonstrates that DynCodon, like the wild-type sequences in the natural Top 10% robust regions, effectively maintains a global balance in tRNA utilization, which can avoid the folding problem caused by simply pursuing high-frequency codons. In contrast, current state-of-the-art models LinearDesign and GEMORNA show a higher abundance matching rate than the available tRNA library in the host cell; while CodonTransformer has a low abundance matching rate, close to 0.

GC content is a fundamental physicochemical parameter affecting mRNA structural stability and transcription efficiency. GC content reflects that optimized sequences possess better mRNA structural stability and host genome compatibility. The performance of GC Content, as shown in Fig. 7 and 8 in the Aappendix, reveals that the sequences optimized by DynCodon are most similar in distribution to wild-type sequences in the natural Top 10% robust regions on this evaluation metric. In contrast, current state-of-the-art models LinearDesign and GEMORNA show generally high scores across most species; CodonTransformer performs well across most species, but is slightly high in a few; and heuristic methods BFC and URC perform well across most species, but are slightly low in a few.

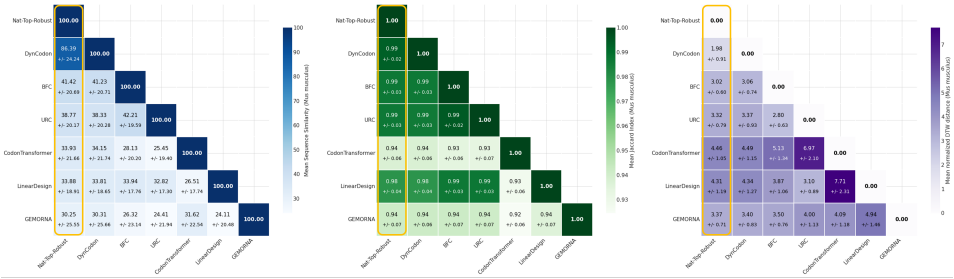

Figure 3: Similarity analysis results for the typical species ***Mus musculus*** under the ***Sequence Similarity Analysis***, ***Jaccard Index***, and ***Dynamic Time Warping (DTW)*** indicators. Subfigure 1 shows the degree of similarity in a global pattern among sequences from optimized sequences generated by ***DynCodon***, ***BFC***, ***URC***, ***CodonTransformer***, ***LinearDesign***, and ***GEMORNA***, and wild-type sequences from the natural Top 10% robust region (***Nat-Top-Robust***); Subfigure 2 shows the degree of similarity in a local pattern among sequences; Subfigure 3 shows the degree of consistency among sequence trends.

4.3 Generated Sequence Similarity Analysis:

To further verify the biological validity of the generated sequences, in addition to examining codon-level metrics, we also introduced three metrics, Sequence Similarity Analysis, Jaccard Index, and Dynamic Time Warping (DTW), to evaluate the consistency between the generated sequences and wild-type sequences in the natural Top 10% robust regions (***Nat-Top-Robust***). This aims to provide an in-depth quantitative evaluation of the generated results from three levels: global pattern, local pattern, and sequence trend.

Sequence Similarity Analysis measures the percentage of matching codons between two sequences. This indicator is crucial for assessing the functional and evolutionary relationships between sequences. As shown in subfigure 1 of Fig. 3, experimental results on a typical species (Mus musculus) under this indicator show that the optimized sequence generated by DynCodon has the highest similarity to wild-type sequences in the natural Top 10% robust regions, compared to the contrasting models. Experimental results for the other seven species are shown in Fig. 9 and 10 in the Appendix, all exhibiting a high proportion of similarity to wild-type sequences in the natural Top 10% robust regions. ***This indicates that the codon sequence optimized by DynCodon better matches the host's preference for highly expressed genes, and that DynCodon has learned the most robust and universal evolutionary patterns from the host's highly expressed genes, thus ensuring high expression potential and biological validity.***

Unlike Sequence Similarity Analysis, the Jaccard Index focuses on assessing the degree of overlap between two sets. In this experiment, we transformed the sequences into 3-mer sets and calculated

the intersection-union ratio of the generated sequences with wild-type sequences from the natural Top 10% robust regions. As shown in Subfigure 2 of Fig. 3, and Fig. 9 and 10 in the Appendix, the optimized sequences generated by DynCodon have the highest similarity scores between them and wild-type sequences from the natural Top 10% robust regions, compared to the contrastive models.

DTW is a classic algorithm for measuring the similarity between two time series or waveform shapes, effectively handling cases where sequence lengths are not perfectly consistent or local phase deviations exist. In this experiment, we utilize the DTW algorithm to quantify the distance between the generated and target sequences in the %MinMax codon usage pattern. %MinMax is a key metric for evaluating local codon bias as it slides along a sequence, reflecting the frequency distribution between high-frequency (common) and low-frequency (rare) codons within a specific length window (18 codons in this experiment). As shown in Subfigure 3 of Fig. 3, and Fig. 9 and 10 in the Appendix, the optimized sequence generated by DynCodon shows good similarity scores with wild-type sequences in the natural Top 10% robust regions, compared to the contrastive models.

Furthermore, we statistically analyzed the codon distribution across sequences and the number of sequences containing the corresponding codons, as shown in Fig. 11 and 12 in the Appendix. The statistical results from eight species also demonstrate that the optimized sequences generated by DynCodon exhibit high distributional consistency with wild-type sequences in the natural Top 10% robust regions.

4.4 Sequence Structure Stability Analysis:

The secondary structure stability of mRNA is a key factor determining its half-life and translation efficiency within the cell. To quantify the thermodynamic stability of the generated sequence, we used Minimum Free Energy (MFE) as the core evaluation metric. Generally, a lower MFE value indicates a more stable secondary structure formed by mRNA folding, and greater resistance to degradation by exonucleases.

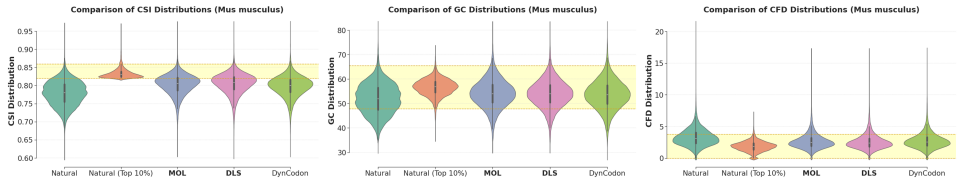

Figure 4: Ablation analysis experimental results for a typical species ***Mus musculus*** under the general evaluation metrics (***CSI***, ***CFD***, ***GC content***), where ***MOL*** represents the multi-objective optimization ablation model, ***DLS*** represents the dynamic learning strategy ablation model, and ***DynCodon*** represents the complete model.

To verify the biological rationale of the optimized sequence generated by DynCodon, we compared it with wild-type sequences from the natural Top 10% robust regions, as shown in Fig. 13 in the Appendix. Analysis revealed that the optimized sequences generated by DynCodon highly overlapped with the wild-type sequences from the natural Top 10% robust regions in the MFE distribution. In contrast, some over-optimized algorithms (such as those simply pursuing maximum GC content) often generated sequences with abnormally low MFE. This "hyperstability" may lead to difficulties in ribosome unwinding during translation, resulting in translational arrest. ***Our data indicate that the optimized sequences generated by DynCodon not only possess extremely high stability, but this stability is also within biologically permissible limits. This precise simulation of the thermodynamic characteristics of natural robust sequences means that the generated mRNA molecules can effectively resist degradation in the cytoplasm while maintaining suitable translational elongation kinetics.***

5 Model Ablation & Analysis

5.1 Model Setup:

Complete Model (DynCodon): Includes the complete network architecture and integrates both multi-objective optimization mode and dynamic learning strategies.

Multi-Objective Optimization Ablation Model (w/o MOL): Includes the complete network architecture, removes the multi-objective optimization mode, and retains the dynamic learning strategy. Specifically, multi-objective optimization ablation refers to using all pre-trained data to further refine the model under the dynamic learning strategy during the fine-tuning phase, instead of fine-tuning the model solely for the target species.

Dynamic Learning Strategy Ablation Model (w/o DLS): Includes the complete network architecture, removes the dynamic learning strategy, and retains the multi-objective optimization mode. Specifically, dynamic learning strategy ablation refers to using traditional parameter optimization methods to refine the model under the target species during the fine-tuning phase, instead of using the proposed dynamic learning optimization strategy.

The hyperparameters and hardware/software environment involved in the ablation models are consistent with those in the main experimental section. ***Other unmentioned model parameter settings will be provided in the code that will be released later.***

5.2 Ablation Analysis:

The ablation analysis experimental results under the general evaluation metrics (CSI, CFD, GC content) are shown in Fig. 4 for the typical species Mus musculus and Fig. 14 in the Appendix. Under the CSI metric, the ablation models MOL and DLS are slightly inferior to the complete model DynCodon for most species, and only slightly superior to DynCodon for a few species. Similar performance was observed under CFD and GC content evaluation metrics. ***This indicates that both the multi-objective optimization mode and the dynamic learning strategy contributed to DynCodon's final performance, but the success of DynCodon lies in the combined efforts of both.***

6 Conclusion

This paper identifies and addresses a key limitation of contemporary codon optimization methods: the "overoptimization trap", where models blindly pursue theoretically extreme values at the expense of biological adaptability. We argue that sequences at theoretical boundaries ("genius sequences") often lack the stability required for stable translation, while robust regions in natural distributions ("high-achiever sequences in the natural Top 10% robust regions") provide a more reliable blueprint for high-performance gene expression. To put this paradigm shift into practice, we propose DynCodon, a novel sequence generation framework driven by a multi-objective dynamic learning strategy. Our extensive experimental results demonstrate that DynCodon successfully avoids the pitfalls of overoptimization. Instead of generating unstable sequences with extreme metric values, the model precisely reconstructs the feature distribution from the natural Top 10% robust regions. This balance ensures that the generated sequences maintain high translation efficiency while preserving the structural stability and evolutionary patterns of "high-efficiency sequences" in nature.

7 Limitations

Although the model performs well in computational metrics (such as CAI, GC content, and reconstruction accuracy), it has not been validated by real in vitro or in vivo biological experiments. The translation efficiency, stability, and immunogenicity of mRNA in a real cellular environment are very complex, and there may still be a gap between pure computational results and actual biological performance. Due to the limitations of our research group, we welcome capable researchers to conduct subsequent wet experimental analyses.

ETHICS STATEMENT

The authors declare that this study adheres to ethical standards. All data used for quantitative and qualitative analysis were obtained from publicly available online sources. The models used are based on authors' open-source code and are appropriately cited. Additionally, Google Gemini and Google Translate were employed solely for language polishing and translation to enhance the manuscript's clarity.

REFERENCES

Nicola A Burgess-Brown, Sujata Sharma, Frank Sobott, Christoph Loenarz, Udo Oppermann, and Opher Gileadi. Codon optimization can improve expression of human genes in escherichia coli: A multi-gene study. *Protein expression and purification*, 59(1):94–102, 2008.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

Adibvafa Fallahpour, Vincent Gureghian, Guillaume J Filion, Ariel B Lindner, and Amir Pandi. Codontransformer: a multispecies codon optimizer using context-aware neural networks. *Nature Communications*, 16(1):3205, 2025.

Andreas Grote, Karsten Hiller, Maurice Scheer, Richard Münch, Bernd Nörtemann, Dietmar C Hempel, and Dieter Jahn. Jcat: a novel tool to adapt codon usage of a target gene to its potential expression host. *Nucleic acids research*, 33(suppl_2):W526–W531, 2005.

Rishab Jain, Aditya Jain, Elizabeth Mauro, Kevin LeShane, and Douglas Densmore. Icor: improving codon optimization with recurrent neural networks. *BMC bioinformatics*, 24(1):132, 2023.

Grzegorz Kudla, Leszek Lipinski, Fanny Caffin, Aleksandra Helwak, and Maciej Zylicz. High guanine and cytosine content increases mrna levels in mammalian cells. *PLoS biology*, 4(6):e180, 2006.

David M Mauger, B Joseph Cabral, Vladimir Presnyak, Stephen V Su, David W Reid, Brooke Goodman, Kristian Link, Nikhil Khatwani, John Reynders, Melissa J Moore, et al. mrna structure regulates protein expression through changes in functional half-life. *Proceedings of the National Academy of Sciences*, 116(48):24075–24083, 2019.

Norbert Pardi, Michael J Hogan, Frederick W Porter, and Drew Weissman. mrna vaccines—a new era in vaccinology. *Nature reviews Drug discovery*, 17(4):261–279, 2018.

Joshua B Plotkin and Grzegorz Kudla. Synonymous but not the same: the causes and consequences of codon bias. *Nature Reviews Genetics*, 12(1):32–42, 2011.

Fernando P Polack, Stephen J Thomas, Nicholas Kitchin, Judith Absalon, Alejandra Gurtman, Stephen Lockhart, John L Perez, Gonzalo Pérez Marc, Edson D Moreira, Cristiano Zerbini, et al. Safety and efficacy of the bnt162b2 mrna covid-19 vaccine. *New England journal of medicine*, 383(27):2603–2615, 2020.

Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science advances*, 4(7):eaap7885, 2018.

Vladimir Presnyak, Najwa Alhusaini, Ying-Hsin Chen, Sophie Martin, Nathan Morris, Nicholas Kline, Sara Olson, David Weinberg, Kristian E Baker, Brenton R Graveley, et al. Codon optimality is a major determinant of mrna stability. *Cell*, 160(6):1111–1124, 2015.

Tessa EF Quax, Nico J Claassens, Dieter Söll, and John van der Oost. Codon bias as a means to fine-tune gene expression. *Molecular cell*, 59(2):149–161, 2015.

Paul M Sharp and Wen-Hsiung Li. The codon adaptation index-a measure of directional synonymous codon usage bias, and its potential applications. *Nucleic acids research*, 15(3):1281–1295, 1987.

Tamir Tuller, Yedael Y Waldman, Martin Kupiec, and Eytan Ruppin. Translation efficiency is determined by both codon bias and folding energy. *Proceedings of the national academy of sciences*, 107(8):3645–3650, 2010.

David K Yang, Samuel L Goldman, Eli Weinstein, and Debora Marks. Generative models for codon prediction and optimization. *Machine Learning in Computational Biology*, 2019.

He Zhang, Liang Zhang, Ang Lin, Congcong Xu, Ziyu Li, Kaibo Liu, Boxiang Liu, Xiaopin Ma, Fanfan Zhao, Huiling Jiang, et al. Algorithm for optimized mrna design improves stability and immunogenicity. *Nature*, 621(7978):396–403, 2023.

He Zhang, Hailong Liu, Yushan Xu, Haoran Huang, Yiming Liu, Jia Wang, Yan Qin, Haiyan Wang, Lili Ma, Zhiyuan Xun, et al. Deep generative models design mrna sequences with enhanced translational capacity and stability. *Science*, 390(6773):eadr8470, 2025.

Michael Zuker and Patrick Stiegler. Optimal computer folding of large rna sequences using thermodynamics and auxiliary information. *Nucleic acids research*, 9(1):133–148, 1981.

A APPENDIX

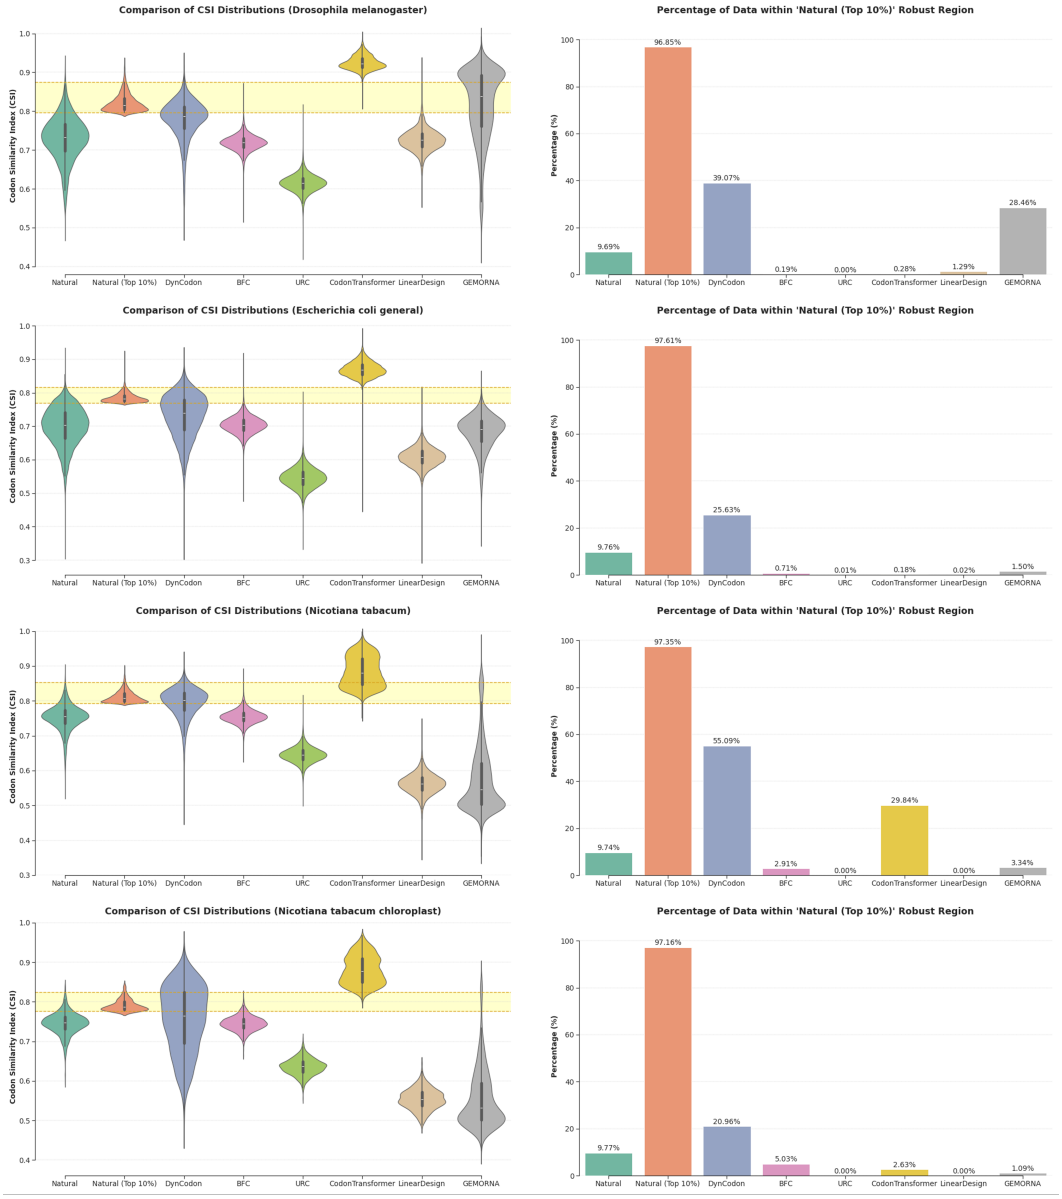

Figure 5: Experimental results for typical species ***Drosophila melanogaster***, ***Escherichia coli general***, ***Nicotiana tabacum***, and ***Nicotiana tabacum chloroplast*** under the CSI evaluation metric.

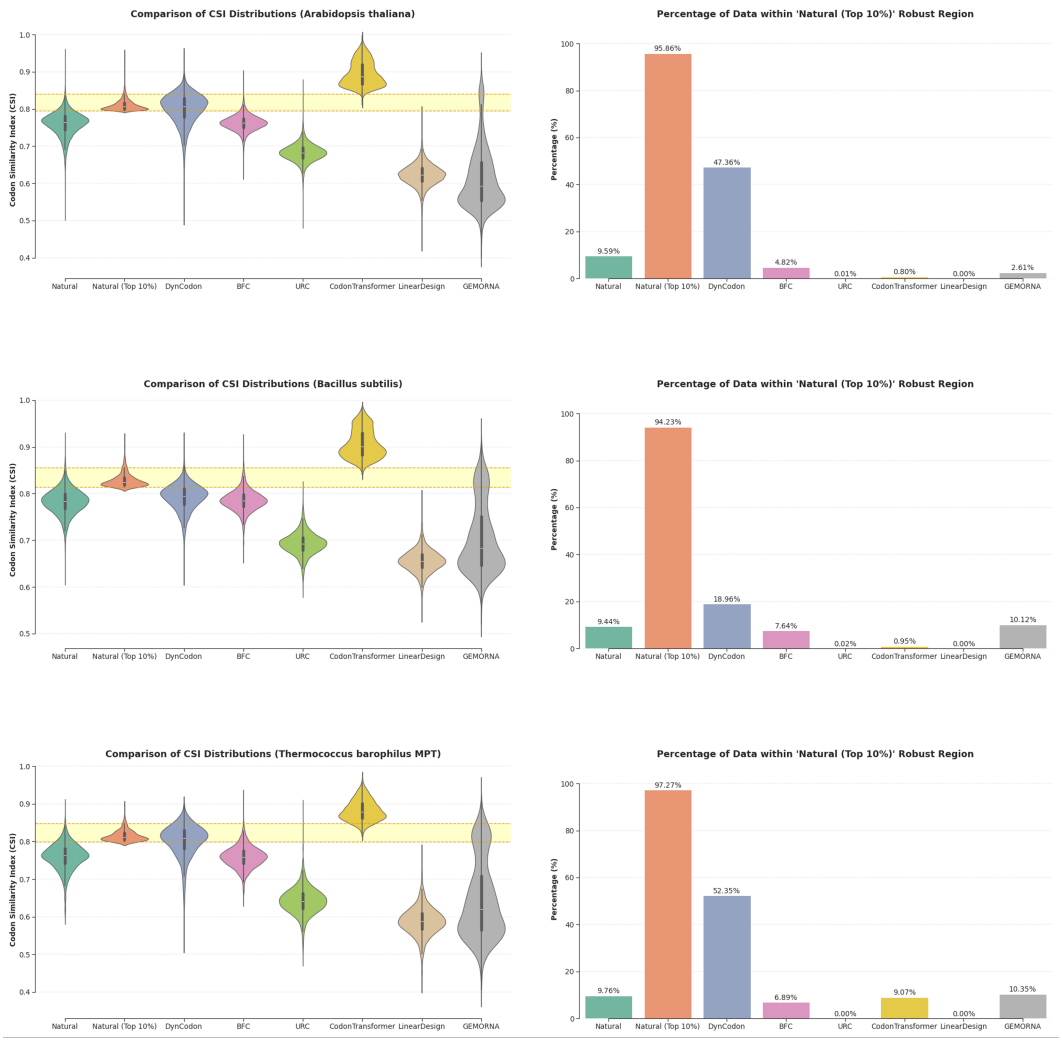

Figure 6: Experimental results for typical species ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT*** under the CSI evaluation metric.

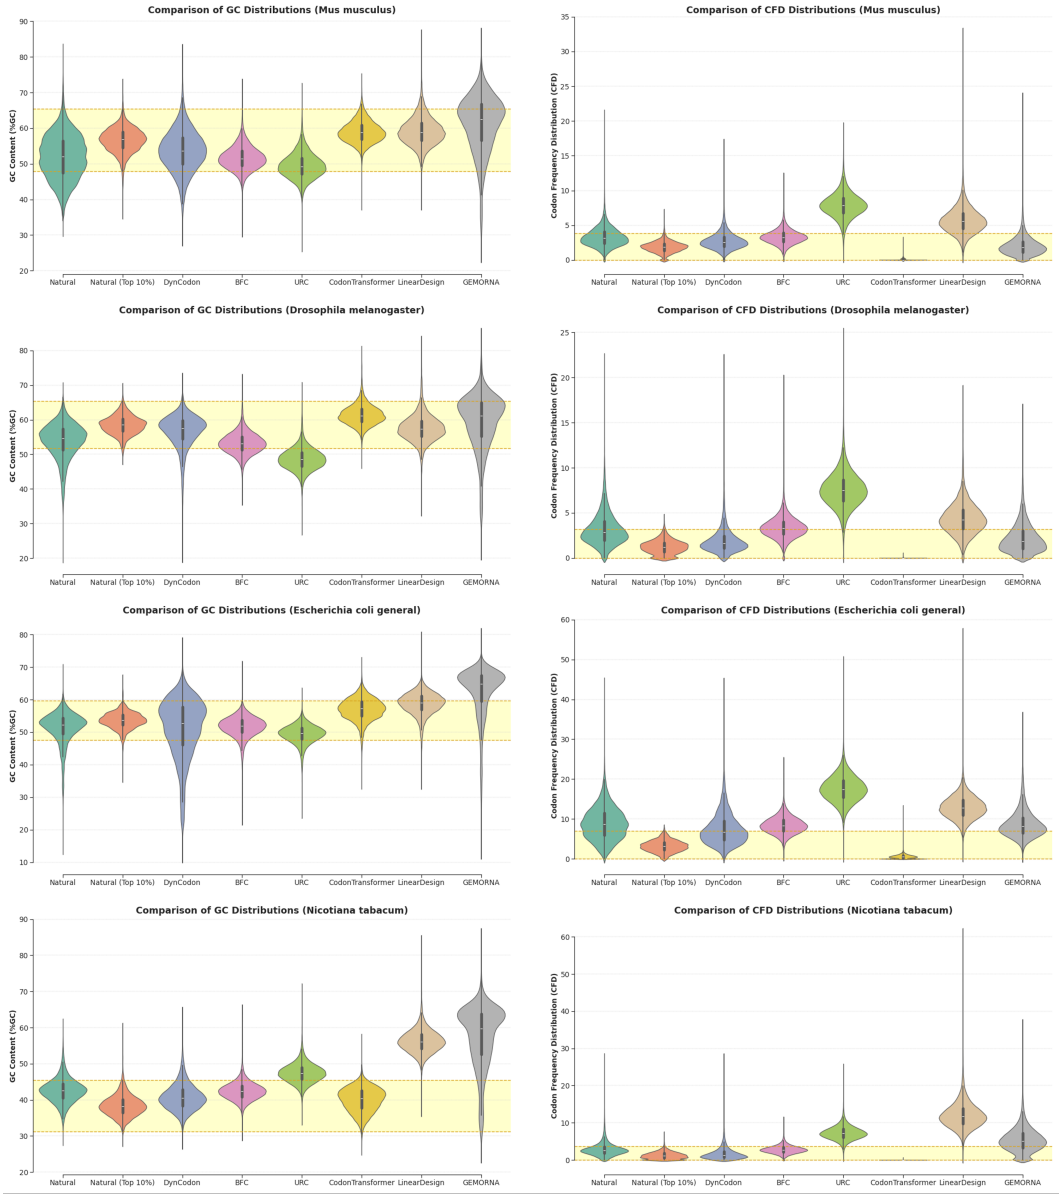

Figure 7: Experimental results for typical species ***Mus musculus***, ***Drosophila melanogaster***, ***Escherichia coli general***, and ***Nicotiana tabacum*** under GC content and CFD evaluation metrics.

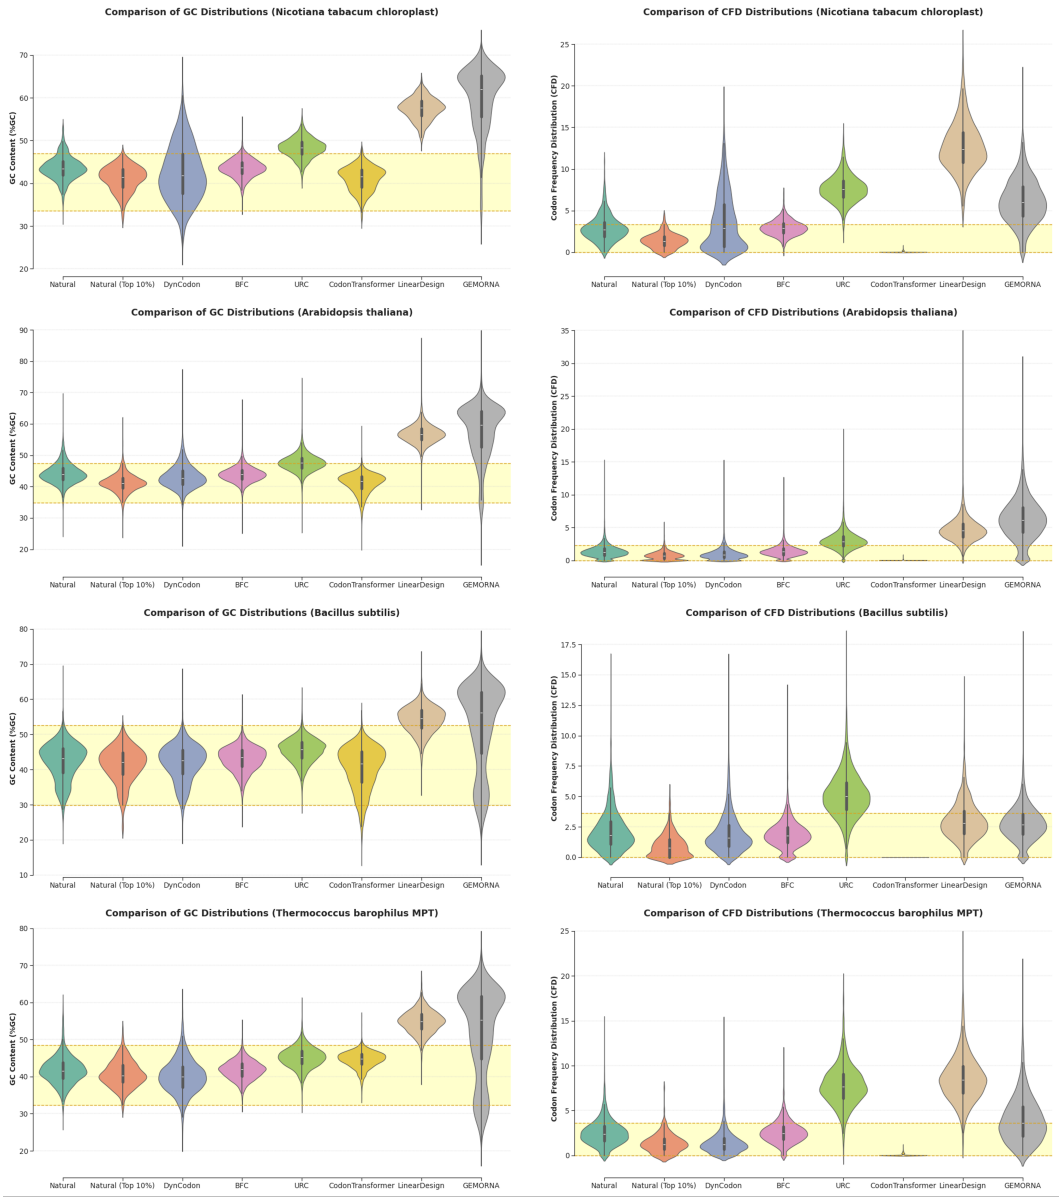

Figure 8: Experimental results for typical species ***Nicotiana tabacum chloroplast***, ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT*** under GC content and CFD evaluation metrics.

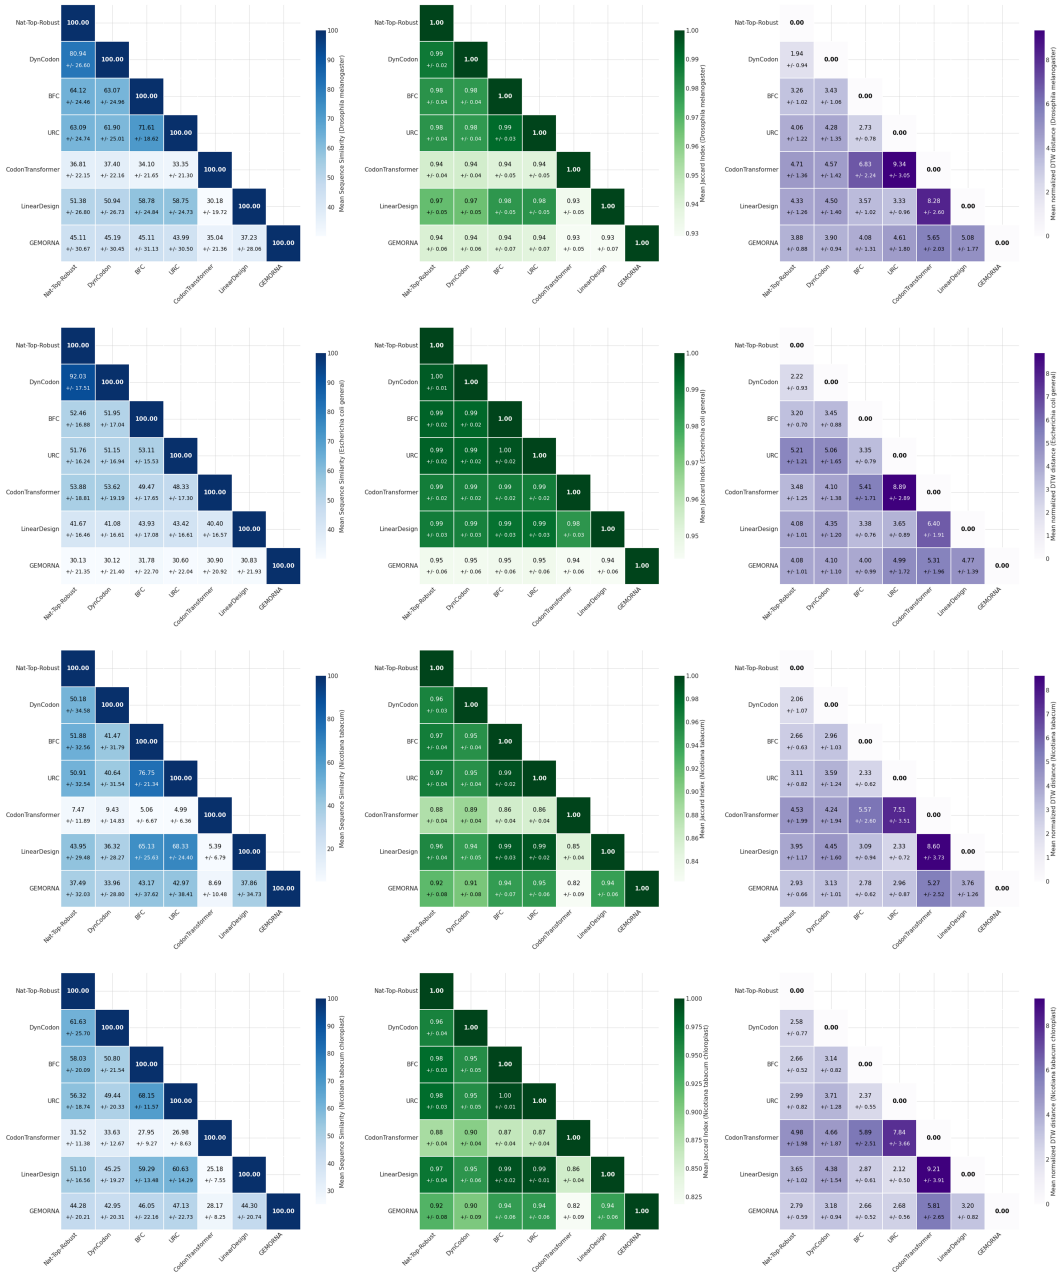

Figure 9: Similarity analysis results for typical species ***Drosophila melanogaster***, ***Escherichia coli general***, ***Nicotiana tabacum***, and ***Nicotiana tabacum chloroplast*** under the ***Sequence Similarity Analysis***, ***Jaccard Index***, and ***Dynamic Time Warping (DTW)*** indicators.

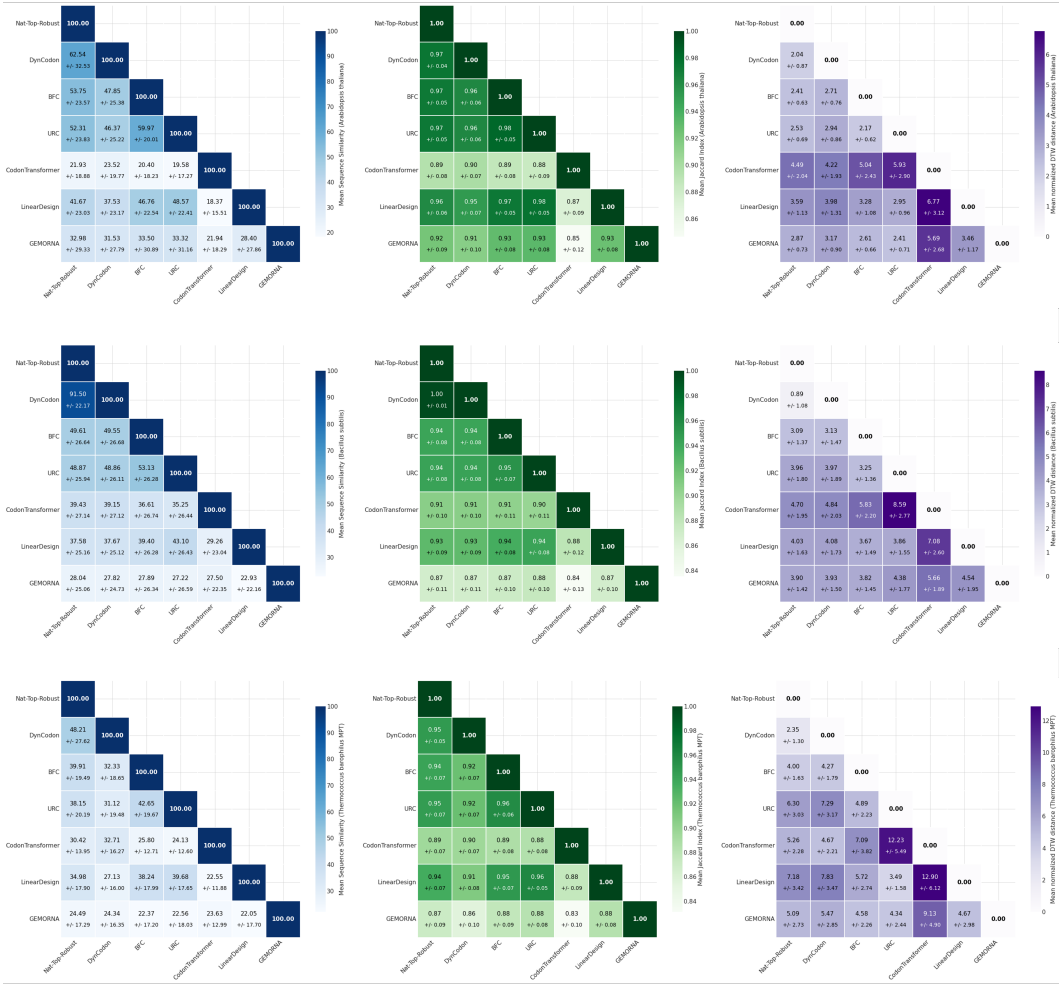

Figure 10: Similarity analysis results for typical species ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT*** under the ***Sequence Similarity Analysis***, ***Jaccard Index***, and ***Dynamic Time Warping (DTW)*** indicators.



Figure 11: Statistical results of codon distribution across sequences and the number of sequences containing the corresponding codons under typical species ***Mus musculus***, ***Drosophila melanogaster***, ***Escherichia coli general***, and ***Nicotiana tabacum***.

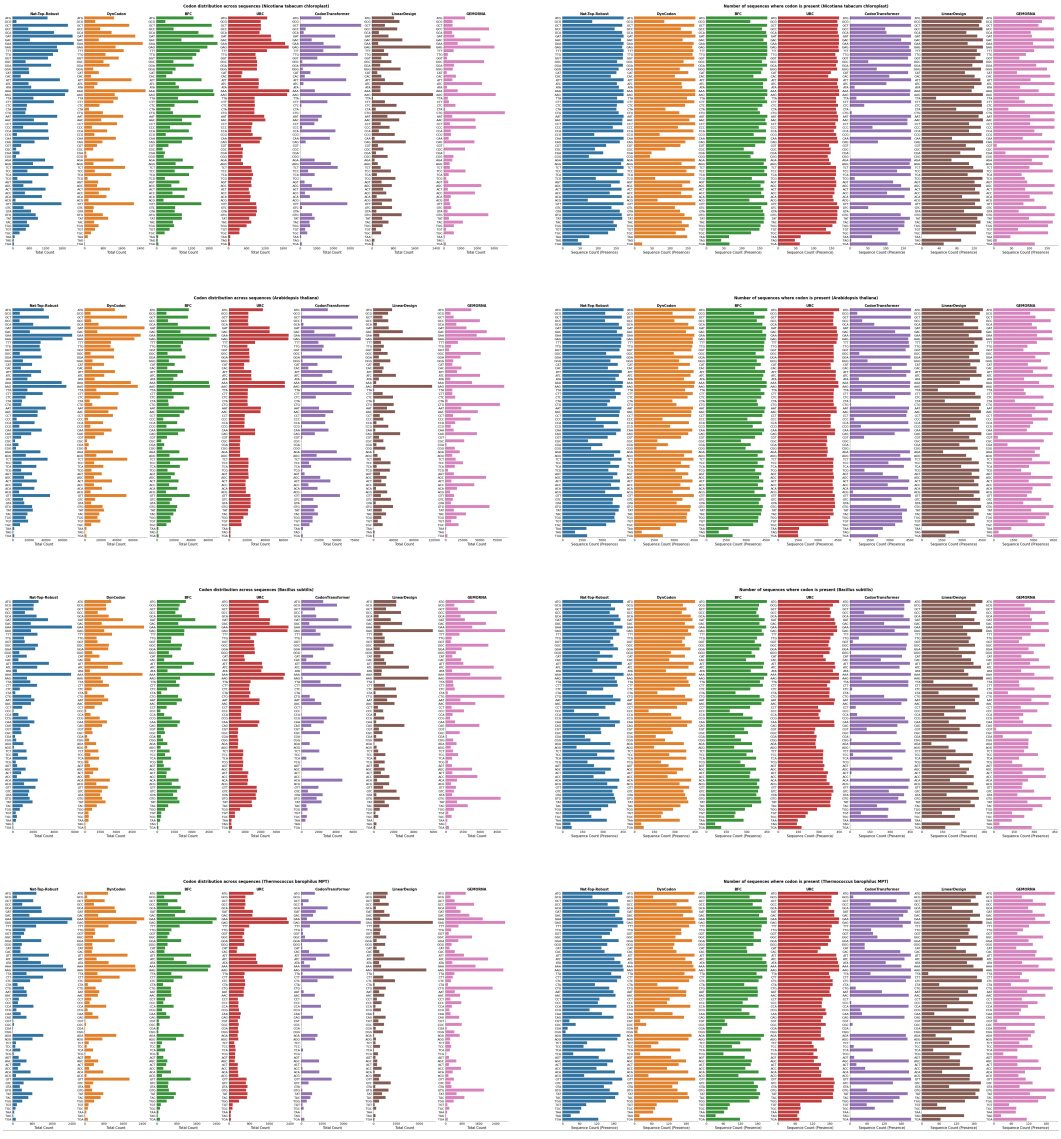

Figure 12: Statistical results of codon distribution across sequences and the number of sequences containing the corresponding codons under typical species ***Nicotiana tabacum chloroplast***, ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT***.



Figure 13: The normalized Minimum Free Energy (MFE) distribution of wild-type sequences from the natural Top 10% robust regions (***Nat-Top-Robust***) and optimized sequences generated by ***DynCodon***, ***BFC***, ***URC***, ***CodonTransformer***, ***LinearDesign***, and ***GEMORNA*** in the typical species ***Drosophila melanogaster***, ***Escherichia coli general***, ***Nicotiana tabacum***, ***Nicotiana tabacum chloroplast***, ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT***.

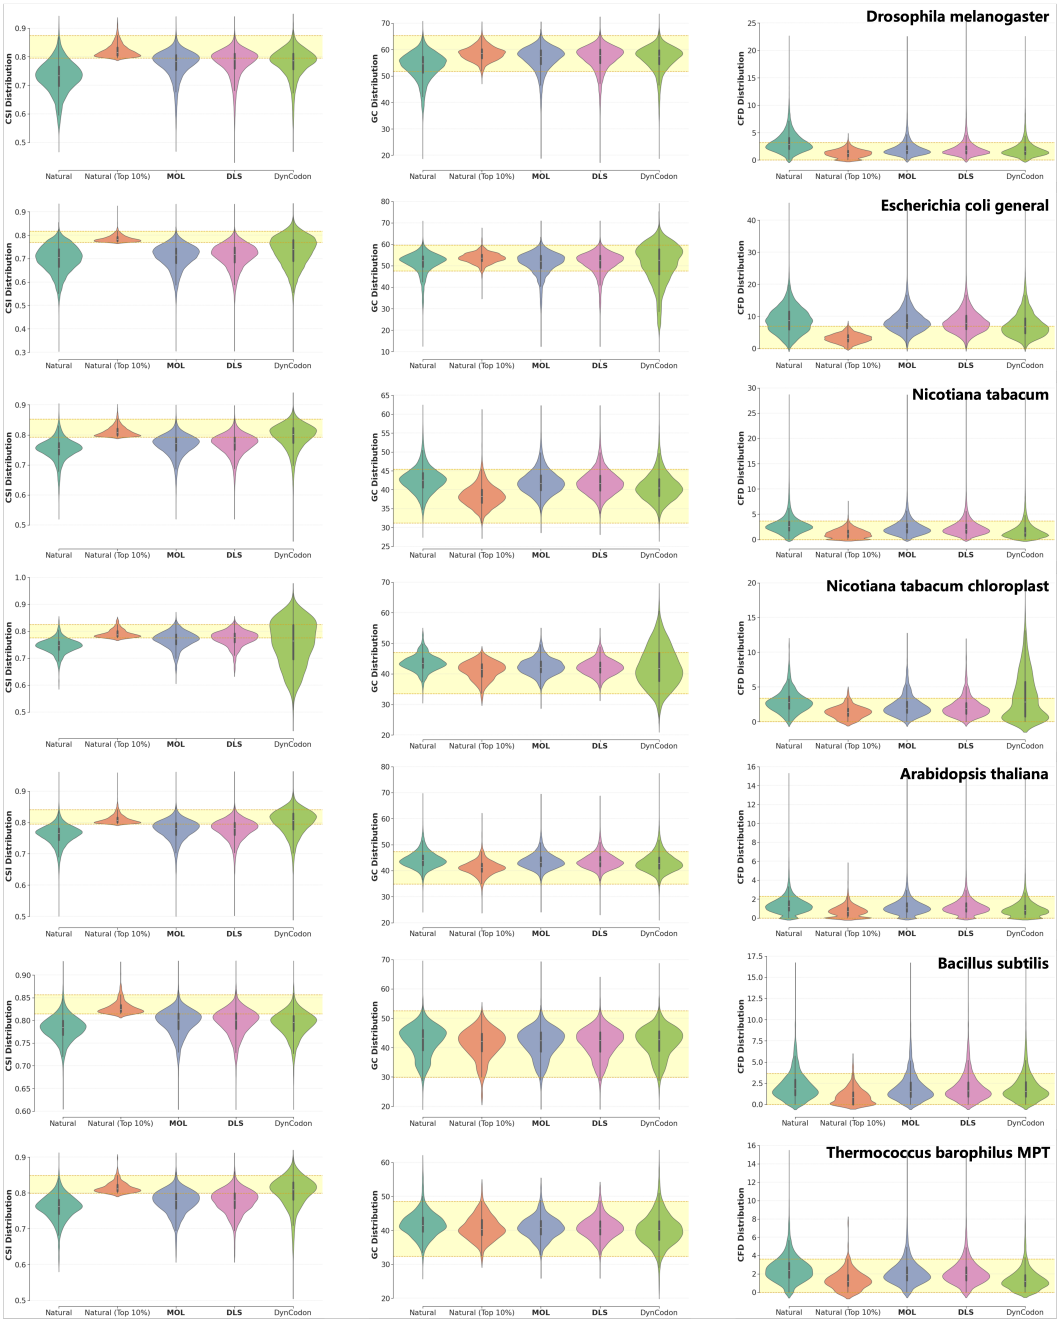

Figure 14: Ablation analysis experimental results for typical species ***Drosophila melanogaster***, ***Escherichia coli general***, ***Nicotiana tabacum***, ***Nicotiana tabacum chloroplast***, ***Arabidopsis thaliana***, ***Bacillus subtilis***, and ***Thermococcus barophilus MPT*** under general evaluation metrics (***CSI***, ***CFD***, ***GC content***), where ***MOL*** represents the multi-objective optimization ablation model, ***DLS*** represents the dynamic learning strategy ablation model, and ***DynCodon*** represents the complete model.