

Foundation Models for AI-enabled Biological Design

Asher Moldwin¹, Amarda Shehu¹

¹Department of Computer Science, George Mason University, Fairfax, VA, USA
amoldwin@gmu.edu, amarda@gmu.edu

Abstract

This paper surveys foundation models for AI-enabled biological design, focusing on recent developments in applying large-scale, self-supervised models to tasks such as protein engineering, small molecule design, and genomic sequence design. Though this domain is evolving rapidly, this survey presents and discusses a taxonomy of current models and methods. The focus is on challenges and solutions in adapting these models for biological applications, including biological sequence modeling architectures, controllability in generation, and multi-modal integration. The survey concludes with a discussion of open problems and future directions, offering concrete next-steps to improve the quality of biological sequence generation.

Introduction

Natural language processing (NLP) has undergone a recent revolution driven by Transformer-based neural networks and the attention mechanism. These architectures have enabled large language models (LLMs) to achieve remarkable performance on natural language understanding and generation tasks, many of which were once thought to be the exclusive domain of human intelligence. While this transformation in NLP has captured widespread attention, a quieter but equally significant revolution has been unfolding in the biological sciences.

Advances in molecular biology, particularly in large-scale data collection initiatives such as the Protein Structure Initiative[1], the Human Genome Project[2] and structural genomics efforts[3, 4, 5], have paved the way for a new era of biological research. The rapid development of next-generation sequencing technologies has led to a wealth of widely available genomics, proteomics, and metabolomics data. This multi-omics view of organisms has driven breakthroughs such as AlphaFold’s[6] success in solving protein structure prediction, earning two of its main authors, Demis Hassabis and John Jumper, the 2024 Nobel Prize in Chemistry. Perhaps more importantly, these rapid developments are catalyzing a boom in bioinformatics methodologies for various prediction and generation tasks.

The aim of this survey is to provide an overview of recent contributions to AI-enabled biological design, where the goal is to leverage biological data and new methods to design and engineer biological entities with impactful ap-

plications in health, drug discovery[7], synthetic biology[8], and material sciences[9]. Notably, David Baker, the third recipient of the 2024 Nobel Prize in Chemistry, was recognized for pioneering work in protein design.

While many neural network architectures and machine learning methods are being developed for biological design, this survey focuses on a particularly promising frontier: Foundation Models (FMs). These models, capable of learning general representations from vast datasets in a task-agnostic setting, offer exciting opportunities for biological applications. Our understanding of FMs has evolved in recent years, and we adopt a broad definition that aligns with the Stanford University Human-Centered AI group’s understanding[10] in coining the term: FMs are architecture-agnostic and defined by their ability to learn from task-agnostic pre-training, making them applicable across a range of tasks. Although some researchers narrowly define FMs as sequential models or strictly equate them with LLMs, we take the position that FMs should be considered more broadly for their capacity to learn deep representations that can be leveraged for downstream tasks.

Given the rapid progress in FMs, it is timely to survey this expanding landscape. These models are quickly transforming AI research, with new developments appearing almost weekly, particularly in biological design applications. While this survey cannot be fully comprehensive, it focuses on key areas of concentrated activity or where significant challenges are being articulated and addressed. Specifically, we examine FMs that, leveraging the analogies between natural language and biological sequences, directly support biological sequence-design tasks; though this focus may be somewhat narrow, we believe it encompasses some of the most exciting developments in the field.

The survey highlights sequence-based FMs and focuses on commonly-applied architectures, such as the Transformer, Diffusion, and State Space Model (SSM) architectures. While FMs can encompass a broader range of architectures, we concentrate on these three due to their demonstrated success in recent biological design literature. Moreover, while FMs can support a variety of prediction and downstream tasks, our focus on biological design narrows our attention to generative tasks, where models not only analyze existing biological data but also aim to create novel biological entities with desired properties. As we will discuss,

this type of biological design includes *de-novo* drug design, protein engineering, and DNA design through gene-editing.

Background and Preliminaries

FMs

While FMs do not imply a specific neural architecture, certain architectures are more naturally suited to the paradigm of representation learning, which is a critical component in modern FMs. We begin by formalizing the concept of implicit representation learning and illustrate it using a decoder-only text-based Transformer, before introducing the diffusion and SSM-based paradigms.

Given an input sequence $X_{input} = \{x_1, x_2, \dots, x_T\}$, where each token x_i belongs to a vocabulary V , a foundation-model must learn a function to compute an internal representation vector h_{final} such that h_{final} encodes the necessary information to reliably optimize the pre-training objective for different values of X_{input} .

In deep learning, this function is typically parameterized such that it can be represented as a series of matrix multiplications and nonlinear transformations. Each network layer computes an intermediate representation $h_l(x) = \sigma(W_l h_{l-1}(x) + b_l)$ where W_l is the trainable weight matrix layer l , b_l is the bias vector, and σ is a nonlinear activation function. The final output layer maps h_{final} to the appropriate form to compute the loss with respect to ground truth.

Transformer Architecture In Transformer architectures[11], sequences are handled by maintaining separate representations for each token at each layer, giving: $h_i^{(l)} = \text{TransformerLayer}^{(l)}(h_1^{(l-1)}, h_2^{(l-1)}, \dots, h_T^{(l-1)})$. This allows the model to create contextualized representations through the attention mechanism, where each token’s representation at layer l depends on the representations of all tokens from the previous layer. Specifically, for each token embedding $h_i^{(l-1)}$, Transformer attention computes attention scores using the query, key, and value projections: $Q_i = W_q h_i^{(l-1)}$, $K_j = W_k h_j^{(l-1)}$, and $V_j = W_v h_j^{(l-1)}$ for all j , where W_q , W_k , and W_v are trainable weight matrices. The layer representation $h_i^{(l)}$ is then computed as a weighted sum of the values: $h_i^{(l)} = \sum_{j=1}^T \text{softmax}(Q_i \cdot K_j) V_j$.

In this formulation, attention incurs a computational cost during training proportional to the number of input tokens squared, rendering its use for long sequences impractical. Memory usage during inference scales linearly with the sequence length, as key and value representations for all tokens need to be stored, further limiting its efficiency.

To effectively leverage the representational power added by the inclusion of context, the use of general, self-supervised pre-training objectives has proven crucial for learning representations that can be repurposed for other downstream tasks. For decoder-only models, the pre-training objective is typically autoregressive language modeling, where the model predicts the next token in a sequence given all previous tokens. Formally, the model learns to approximate the probability distribution $P(x_i | x_1, x_2, \dots, x_{i-1})$. The attention mechanism is

thus constrained to left-to-right context, ensuring that each token attends only to previous tokens: $h_i^{(l)} = \text{Attention}^{(l)}(h_1^{(l-1)}, h_2^{(l-1)}, \dots, h_{i-1}^{(l-1)})$.

Finally, while the attention mechanism handles dependencies between tokens, applying attention multiple times in parallel at each layer (multihead attention), and stacking many Transformer layers in a deep network builds the capacity to learn complex and hierarchical relationships in data.

State-Space Models (SSMs) State-Space Models (SSMs) provide an efficient alternative to Transformers for sequence modeling, particularly for tasks requiring long-range dependencies. While Transformers rely on quadratic self-attention and increasing memory during inference, SSMs achieve sub-quadratic scaling and constant memory by employing a Recurrent Neural Network architecture and using fixed-size hidden states and linear state-space equations. These equations model state evolution over time as $h_{t+1} = Ah_t + Bu_{t+1}$, with outputs $y_t = Ch_t + \Delta u_t$. Key architectures like S4[12] enhance SSMs with optimizations for long sequences, while Mamba[13] employs Selective State Spaces to learn dynamic versions the A, B and C matrices whose values depend on the input token.

Diffusion Modeling Models can forgo attention by instead leveraging diffusion processes wherein noise is iteratively introduced to the inputs and then a learned reverse diffusion process is performed to reconstruct realistic inputs. This allows global contextual relationships to gradually emerge in the outputs of diffusion models over many iterations of denoising. Formally, diffusion models progressively add random noise to input data over T timesteps, following a Markov process in which each timestep’s result depends only on the previous timestep and the added noise. The probability of transforming the input x_0 into the noisy state x_T after T steps is given by: $q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1})$, where each transition is defined as: $q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbf{I})$ with β_t controlling the variance of the added noise at each step. The reverse process, or “denoising,” mirrors this forward process, allowing the model to reconstruct the input by removing noise over time. For reverse diffusion, a neural network ϵ_θ is trained to progressively denoise x_T by reversing the diffusion process at each timestep t . The standard training objective is derived from a variational bound on the data likelihood, giving the loss function $\mathcal{L}(\theta) = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2]$.

While diffusion models were originally developed for continuous data, like images, alternative approaches have adapted this framework to discrete data, such as text. In one approach, tokens are mapped into a continuous latent space, where noise is progressively added over T timesteps. The reverse process denoises the embeddings step by step, and a decoder maps the final latent embeddings back to discrete tokens. This allows diffusion models to maintain local coherence while capturing global structure over multiple denoising steps. Another approach directly applies noise to discrete tokens[14] by either replacing tokens with others sampled from the vocabulary[15] or by masking them[16].

In the reverse process, a neural network restores the correct token sequence from its noisy version by minimizing the difference between the predicted tokens and the ground truth at each timestep. This iterative denoising process enables diffusion models to generate coherent text, effectively adapting the diffusion framework to discrete data.

Biological Design

Biological design encompasses the creation or modification of biological entities for specific functions or properties. Researchers focus on three primary classes of biological or chemical objects: proteins, DNA/RNA, and small molecules. Different representations of these objects typically enable different neural network architectures, with the two main representations being sequence-based and graph-based. Sequence-based representations are derived from the chemical formulas, such as representing proteins as sequences of characters corresponding to the twenty naturally occurring amino acids or representing genomic sequences as strings of the four nucleotide bases. The SMILES representation for small molecules captures atoms as well as bonds and branches. Graph-based representations, on the other hand, model the bonds and interactions (e.g., hydrogen bonds, van der Waals interactions) between atoms or molecular units as edges connecting vertices.

Biological design combines principles from biology, chemistry, and computational science to engineer biological systems. The primary areas of greatest activity are: (1) Protein engineering: Designing novel proteins or modifying existing ones for enhanced function, stability, or specificity; (2) Small molecule design: Creating new drug candidates or optimizing existing compounds for improved efficacy and reduced side effects. Advances toward generating biologically functional small molecules or designing novel proteins with precise control over properties—such as binding affinity, solubility, or toxicity—could significantly accelerate therapeutic development by streamlining the identification of drug candidates and (3) Genomic sequence design: Engineering DNA sequences for applications in synthetic biology, gene therapy, or CRISPR-based genome editing. In this survey we focus on how text-based FMs are adapted for biological design due to their generative capabilities. We will refer to FMs for these three topics as Chemical Language Models (CLMs), Protein Language Models (PLMs), and Genomic Language Models (GLMs).

Taxonomy and Survey

Figure 1 shows our taxonomy highlighting recent areas of focus in FMs for Biological Sequence Design. We include references to papers that address each topic in the leaf nodes of the taxonomy tree, color-coded by the biological domain that they discuss.

Methods aligned with these categories are summarized in Table 1.

Architectures

Transformer Models

Classic Transformer Models

The SMILES representation facilitates using standard Transformer architectures like GPT and XLNet [41] for small-molecule design, often with minimal modifications and NLP-optimized hyperparameters. MolGPT [17], MCMG [18], Regression Transformer [42], and Taiga [19] employ GPT-like architectures but differ in training and controllability strategies. PLMs like Progen [21], Progen2 [22], and ProtGPT2 [23] also use standard architectures with minor enhancements. For instance, Progen2 incorporates Rotary Position Embeddings (RoPE [39]) and improved parallelization. In contrast, most DNA FMs adopt highly customized attention mechanisms or alternative architectures.

Specialized for Biological Design

Transformer-based models often adapt attention mechanisms to suit biological tasks. Examples include RFDiffusionAllAtom [24], NOS [25], and Evo [26]. RFDiffusionAllAtom extends RoseTTAFold2’s multi-track attention with atomic-level details for protein-molecule complex design. NOS employs an encoder-decoder Transformer for forward and backward diffusion, optimized for antibody design. Evo uses a hybrid StripedHyena architecture, combining RoPE-based attention with convolutional layers [43], enabling efficient processing of long genomic sequences (up to 131 kilobases) for tasks like CRISPR DNA design.

State Space Models

Classic SSMs

Many SSM architectures struggle to effectively capture relevant information between tokens separated by large distances in the input. This challenge is significant for protein data, where distant amino acids in the sequence may interact through hydrogen bonds, disulfide bridges, or hydrophobic interactions. Likewise, genomic data contain regulatory elements located hundreds of thousands of tokens away from the genes they regulate. A similar challenge exists in modeling ring closure and branching in small molecules [27], though these sequences are considerably shorter.

Several studies have applied Mamba and S4 architectures directly to biological tasks and report superior performance. However, what qualifies as a fair comparison remains unclear (e.g., should models be compared based on the same number of parameters, training time, or inference time?). The CLM introduced by Ozcelik et al. [27] was the first to apply S4 models to SMILES chemical sequences, benchmarking them against other CLMs such as LSTMs and GPT-type models for *de-novo* molecule generation. They demonstrate SSM advantages in generating valid, unique, and novel molecules.

Various SSM-based protein sequence models have emerged, including ProtMAMBA[28], ProteinMamba [44], and PTMMamba [45]. ProtMAMBA can autoregressively generate new protein sequences, with or without conditional information from homologous sequences. Motivated by long-range dependencies in proteins, it employs a variant of the “fill-in-the-middle” pre-training objective [46],




Figure 1: Taxonomy of recent contributions in FMs for Biological Design. Leaf nodes contain references to the relevant papers, color coded by the type of biological sequence that they discuss.

Table 1: Summary of current FM models for Biological Design.

Model Name	Domain	Architecture	Design Goal	Generation Method
MolGPT [17]	Small Molecule	Decoder-only Transformer	Control TPSA, QED, SAS, LogP and molecular scaffolds	Conditional Autoregressive Generation, pre-trained with prepended property condition embeddings
MCMG [18]	Small Molecule	Decoder-only Transformer Distilled into RNN	Control Bioactivity, QED, SAS	Same as above, with added reinforcement learning rewarding low property error
Taiga [19]	Small Molecule	Decoder-only Transformer	QED, inhibitory potency (pIC50)	Policy gradient reinforcement learning rewarding high property values
S4 CLM [27]	Small Molecule	S4-based SSM	MAPK1 kinase inhibitors	fine-tuning for transfer to desired class
DiffuMol [32]	Small Molecule	Diffusion on embeddings, Transformer decoder for denoising	LogP, QED, TPSA, SAS, and scaffolds	Noiseless conditional token anchors
Regression Transformer[20]	Small Molecule and Protein	Decoder-only Transformer (XLNet)	Control QED, solubility, lipophilicity for small molecules, fluorescence, stability and Bowan index for proteins	Alternating training scheme that switches between property prediction and conditional sequence generation
Progen [21]	Protein	Decoder-only Transformer	Proteins associated with specific families, biological functions	Conditional autoregressive generation and fine-tuning
ProGen2 [22]	Protein	Decoder-only Transformer with RoPE[39]	Structurally valid proteins, specific folds, antibodies	Fine-tuning, and three-residue motif prompts for antibodies
ProtGPT2 [23]	Protein	Decoder-only Transformer	Novel proteins that are also plausible, mostly globular proteins	Autoregressive sampling with modified sampling schemes
ProtMamba [28]	Protein	Mamba-based SSM	Homologous proteins, valid inpainting-based protein modifications	Autoregressive with homologs for context, or inpainting using Fill-in-the-Middle
EvoDiff [33]	Protein	Diffusion with Dilated CNN	Proteins from specific families, inpainting functional domains, generating around motifs	MSA-conditioned for family-based, motif-based conditioning and inpainting for scaffolds
NOS [25]	Protein	BERT-based Encoder-Decoder	Antibodies with high expression yield and binding affinity	Diffusion, multi property value function, Latent Multi-Objective Bayesian Optimization
RFDiffusionAA [24]	Protein complexes	Specialized Multi-track attention based architecture	Protein binders for small molecules, including nucleic acids, proteins, and ligand complexes	Condition on ligand as fixed noiseless anchor
MMDIFF [34]	DNA and Protein	One-hot vector for discrete sequences, FrameDiff architecture for 3d structure	Macromolecular Complexes	Joint reverse diffusion with separate loss components for sequence and structure
regLM [29]	DNA	Based on Hyena-DNA[40], SSM-like	Cis-regulatory elements with specified levels of activity or cell-type specificity	Special condition tokens for activity level or cell type, outputs filtered using regression model
Evo [26]	DNA	StripedHyena using attention and SSM-like blocks	Nucleotide sequences for CRISPR systems, function-conserving transposable elements	Fine-tuning to enable conditioning on special tokens
DiscDiff [30]	DNA	2-stage VAE, with U-net based denoising networks	DNA for specific species, gene expression levels, or cell types	Conditional generation with absorb-escape algorithm
DNA-diffusion [31]	DNA	U-net backbone	Regulatory sequences to control chromatin accessibility	Conditioning on cell-type labels

which trains the model to predict masked segments of a sequence using both preceding and succeeding amino acids. This enables more accurate sequence inpainting and generation tasks designed to capture long-range dependencies. ProtMAMBAFoundation is proposed for general tasks, such as sequence generation and inpainting, while ProtMAMBAFine-Tuned is designed specifically for inpainting.

While Mamba and S4 models are well-established in the State-Space Modeling paradigm, the precise definition of SSMs can be ambiguous. Following Benegas et al. [47], models based on the Hyena architecture[43] can also be classified as SSMs due to shared features like data-dependent gating and long convolutions. Hyena, which generalizes techniques from H3[48], excels in long-range sequence comparisons, making it highly effective for DNA modeling. In the realm of DNA-sequence generation, models like HyenaDNA [40] serve as backbones. This efficient long-sequence architecture is critical, as HyenaDNA is trained on sequences of up to 1 million nucleotides. It allows the model to capture long-range interactions at single-nucleotide resolution, which is crucial for tasks like detecting regulatory elements, mutations, and other genomic features spanning large distances. RegLM[29] is a generative model that uses a Hyena-based backbone, and Evo [26] relies on a variant called StripedHyena, combining both attention and Hyena Layers.

Diffusion Models

Classic Diffusion Models

Diffusion models have become popular for generating continuous data, such as images or molecular graphs. Recently, these models have also been extended to biological sequences by embedding discrete token inputs. For example, DNA-Diffusion [31] generates regulatory sequences controlling chromatin accessibility and gene expression. It uses a U-net convolutional architecture, similar to image diffusion networks like DALL-E [49]. By applying a transformation similar to one-hot encoding—where each nucleotide is represented within a continuous range of $[-1, 1]$ —DNA-Diffusion is able to introduce Gaussian noise to discrete biological sequences, facilitating the diffusion process. Building on similar principles, DiscDiff [30] leverages a U-net with ResNet blocks for DNA generation, using a Variational Autoencoder to encode nucleotide sequences into a continuous latent space for diffusion, before decoding them back into discrete form after the reverse diffusion process.

Specialized for Biological Design

EvoDiff [33] explores two distinct forward diffusion processes tailored for protein modeling: Order Agnostic Autoregressive Diffusion (OADM), which involves randomly masking tokens, and the Discrete Denoising Diffusion Probabilistic Model (D3PM), which applies mutation-based noise by replacing amino acids according to a transition matrix derived from natural mutation probabilities. This mutation-driven approach allows EvoDiff to navigate biologically meaningful regions of evolutionary space, increasing the likelihood that the generated proteins will retain essential properties such as functionality, stability, and activ-

ity—features typically conserved in naturally-evolved proteins.

We observe hybridizations of classic diffusion and transformer architectures to capture broader biological contexts. NOS[25] adapts a BERT-small model encoder-decoder model to better integrate contextual information in protein sequences during diffusion, while DiffuMol [32] targets small-molecule design, incorporating a Transformer decoder to generate molecules with desired properties like QED and LogP.

In multi-modal data contexts, RFDiffusionAA extends RoseTTAFold All-Atom for diffusion-based generation of protein-molecule complexes, utilizing both sequence and structure. Similarly, MMDIFF [34] introduces joint diffusion across DNA and protein sequences, incorporating distinct loss components for sequence and structural modalities to enable coordinated generation of biologically coherent macromolecular complexes.

Controllability in Generation

A key challenge in advancing FMs for biological design is achieving fine-grained control over generated data. We note that biological design is inherently an engineering endeavor. The designed entities are intended to be ultimately synthesized in wet laboratories and then operationalized for particular outcomes. So, the issue of control is inherent to successful, synthetically-accessible and operationalizable design.

For small molecules, continuous numerical properties may need to be controlled, including Drug-likeness, LogP, Molecular Weight, Synthetic Accessibility, Toxicity, and Topological Surface Area. One may also require particular classes of molecular compounds, such as “kinase inhibitors” or “Quaternary Ammonium Compounds” depending on the downstream task/application. For proteins, one typically controls for specific functions or similarity to other known proteins. For DNA, it may be necessary to ensure specific regulatory properties, such as promoters or enhancers, or design guide RNAs (gRNAs) to facilitate CRISPR-based gene editing.

Fine-Tuning

Many of the models discussed above rely on some form of fine-tuning to adapt a pre-trained FM for controlled generation. In its simplest form, this involves continued training with the same pre-training objective but on a smaller subset of data that matches the desired properties/criteria. Ozcelik et al[27] pre-train- their S4-based CLM on 1.9 million SMILES, then fine-tune it on 68 manually-annotated molecules known to inhibit MAPK1, generating new candidates likely to have high binding affinity to MAPK1. Madani et al. [21] fine-tune the ProGen model on lysozyme families after pre-training on 280 million protein sequences, then generating proteins with catalytic efficiencies comparable to natural lysozymes despite low sequence identity.

Fine-tuning can be computationally expensive, especially for large models like ProGen, with 1.2 billion parameters, or insufficient when more granular control is required. Nevertheless, fine-tuning is often a necessary step before applying more controlled generation techniques, such as conditional

generation. It is worth noting that fine tuning may not be effective if the dataset with properties of interest is particularly small relative to the pre-training dataset.

Conditional Generation

Conditional generation offers another method for controlling model outputs by embedding "prompts" or "control tags" into the input, which guide the generation toward desired characteristics. RegLM [29] employs conditional generation to design synthetic cis-regulatory elements (CREs; e.g., promoters, enhancers), where autoregressive generation is conditioned on a starter fragment that guides the prediction of subsequent tokens. Similarly, Evo [26] fine-tune a model on 82,430 loci containing CRISPR-Cas sequences, using prompt tokens (e.g., "cas9," "cas12," "cas13") to optimize gRNA sequences for improved efficiency and reduced off-target effects.

Progen2 conditions the generation of antibody (protein) sequences on three-residue motif prompts, while ProtMamba enables homologous protein generation by conditioning on proteins from the desired family as context. RFDiffusionAA [24] ensures accurate protein-ligand complex generation by conditioning the denoising process on the presence of an unaltered ligand during diffusion. In contrast, MolGPT [17] embeds control tokens directly in the pre-training inputs. While effective, this approach re-trains the model from scratch for each new property or combination of properties, diverging from the more general FM paradigm.

While conditional generation enables control over fine-grained and numerical properties, it presents several limitations. Chief among these is the reliance on examples in the training or fine-tuning data, which can constrain the model's ability to discover novel ways of achieving desired properties. Additionally, the model may develop a bias toward retaining features correlated with the target properties from the training data, potentially hampering diversity in generated data.

Reinforcement Learning

Reinforcement learning (RL) offers a promising alternative by guiding the generation with a reward function. In RL, an agent learns to make decisions through feedback from an environment, optimizing for specific objectives. In sequence generation, each token is treated as an action, and the model is fine-tuned based on how well the generated sequence meets the desired conditions. RL's flexibility in handling complex, non-linear reward landscapes makes it ideal for multi-objective optimization, where properties like bioactivity and molecular diversity must be balanced.

The MCMG model [18] employs RL to control multiple properties in small-molecule generation. Initially, a conditional Transformer generates molecules based on property embeddings. To improve the model's efficiency and balance diversity and novelty, the trained model is distilled into a simpler recurrent model, which is then fine-tuned with RL. During RL training, the model generates molecules autoregressively, with a reward function evaluating properties such

as bioactivity, drug-likeness, and synthetic accessibility. Invalid molecules receive a reward of zero, and the reward is incorporated into an augmented likelihood loss function that balances property optimization with sequence validity.

Similarly, Mazuz et al. [19] combine a Transformer with RL, using a policy gradient approach and the REINFORCE algorithm [50]. After pre-training on SMILES strings, the model is fine-tuned to optimize molecular properties like QED and inhibitory concentrations (pIC50). A reward function proportional to the property values is applied for fully generated, valid molecules, while intermediate generation steps receive zero reward. The model optimizes the expected return, with a high discount factor ($\gamma = 0.99$) to prioritize near-term rewards during sequence generation. Despite these successes, we note that RL is often impractical due to the challenges of designing effective reward functions, which demand significant biological and computational insight, especially given the high-dimensional action spaces in biological design.

Custom Objective Functions

Non-RL methods also employ custom loss functions to directly evaluate sequence properties, allowing for property-specific optimization during generation. NOS[25] uses custom objective functions to optimize properties such as binding affinity and expression yield. This is achieved by incorporating a value function, $v_\theta(w)$, trained via a separate discriminator (or value network), which shares hidden layers with the generative diffusion model up to a certain depth. The discriminator is trained to predict the target objective for proteins in the training set. Once trained, the value function is applied to the hidden states of the sequence representation in the denoising network, and the model uses gradient ascent on these hidden states to maximize the objective value.

The Regression Transformer takes a similar approach by alternating between sequence generation and numerical-property regression tasks. It introduces a custom "property prediction objective" that masks numerical tokens representing property values, requiring the model to predict these values from the associated textual sequence: $J_P = \max_{\theta} \mathbb{E}_{z \sim Z_T^p} [\log p_{\theta}(x^p | x^t)]$, where x^p are the property tokens, x^t are the associated textual tokens, and $z \sim Z_T^p$ denotes the constrained permutation order ensuring that only property tokens are masked. Additionally, a "self-consistency loss" is introduced to ensure the generated sequence adheres to the primed property by using the model's own property prediction as a target during generation: $J_{SC} = J_G(x) + \alpha \cdot J_p(\hat{x})$, where J_G is the conditional text generation objective, J_p is the property prediction objective. By alternating between these custom loss functions, the model can effectively optimize conditional generation for fine-grained numerical properties.

Post-Generation Filtering

A key measure of success for FMs in biological sequence design is the ability to reliably generate sequences worthwhile for practical validation. If *in-silico* validation is cost-effective, high-throughput filtering can screen large numbers

of candidates, discarding those that fail to meet desired conditions before advancing to *in-vitro* validation.

An example of this is the SSM-based CLM proposed by Ozcelik et al[27], where, after fine-tuning on MAPK1 inhibitors as mentioned above, they next generated 256,000 molecules. The molecules are ranked based on log-likelihood scores, computed by subtracting pre-training scores from fine-tuning scores to focus on newly learned bioactivity features. From the top 5,000 molecules, scaffold similarity filtering divided them into high and low similarity groups. The highest-scoring molecules from each group were further evaluated using molecular dynamics simulations to assess binding affinity.

Similarly, RegLM [29] apply post-generation filtering to synthetic CREs by using a sequence-to-function regression model to predict activity levels. Sequences outside two standard deviations from the mean activity were discarded. The model also filtered sequences for cell type specificity, selecting those with the highest specificity based on predicted activity in target and off-target cell lines. Biological realism was ensured through metrics like GC content and k-mer frequency, ensuring the synthetic sequences mimicked natural CREs.

Numerical Property Optimization Strategies

In biological design, many properties of interest are numerical, such as synthetic accessibility, drug-likeness, lipophilicity, and molecular weight. These molecular properties, evaluated using tools like RDKit, were considered in models like MolGPT[17] and DiffuMol[32]. Additional numerical properties include bioactivity[18], solubility[20], and inhibitory potency[19]. For proteins, numerical properties like interaction potential (e.g., the Boman index), fluorescence, and stability (from datasets like TAPE [51]) are of interest [20]. In DNA design, properties like gene expression levels are key, as seen in DiscDiff [30], which optimizes regulatory elements to control expression.

Several models incorporate these numerical targets directly into training. For instance, RL-based methods [19, 18] optimize numerical properties as part of their reward functions. One specialized approach is the Regression Transformer [20], which encodes numerical properties as tokens within a decoder-only architecture. To handle numerical tokenization, the authors introduced a tokenization scheme for floating-point numbers and developed numerical encodings, similar to positional encodings, to help the model understand numerical relationships. MolGPT [17], by contrast, reserves specific dimensions in the input embeddings for numerical values, avoiding explicit tokenization. This differs from the common NLP approach used in models like T5 [52] and GPTs [53], which apply word-piece tokenization to numbers.

Multi-Condition Generation and Tradeoff Handling

Generating molecules with multiple specified properties can be challenging, especially when the desired combinations are rare in the training data or constrained by physical and

chemical limitations. A key challenge is managing tradeoffs, where the model must balance generating valid and diverse molecules that meet all specified conditions.

In molecule generation, models are often evaluated based on criteria like validity, uniqueness, and novelty. A tradeoff exists between validity and uniqueness: the more distinct or novel a generated molecule is, the harder it is to ensure its validity. This tradeoff is especially important in applications like antibacterial drug design, where novelty is critical to address bacterial resistance. Users can adjust priorities by tuning parameters such as the temperature in autoregressive Transformer models, emphasizing either validity or uniqueness depending on the task. Several studies explore this tradeoff. For instance, MCMG [18] uses reinforcement learning to optimize multiple molecular properties while maintaining both validity and diversity. DiffuMol [32] leverages a diffusion framework that provides stepwise control over multiple objectives, making it easier to generate novel molecules while optimizing specific properties. DiffuMol’s use of the Transformer’s attention mechanism helps balance diversity and specificity by focusing on key regions of the molecule that must meet certain property or scaffold requirements, while allowing variation in other parts to promote molecular diversity.

Sampling Algorithms

One part of autoregressive generation process that can affect the natural plausibility of the output sequences is the sampling algorithm used to predict tokens based on the model’s probability distribution. ProtGPT2[23] explored various sampling strategies to generate *de-novo* protein sequences, such as greedy search, beam search, and random sampling. They motivate this by noting that greedy and beam search often lead to repetitive and deterministic sequences, which do not reflect the natural variability found in proteins. On the other hand, random sampling from the top “k” tokens, especially when using a high value of “k” (e.g., $k = 950$), was found to produce more natural-like sequences by introducing variability and capturing natural amino acid propensities.

Multi-Modal FMs

While FMs often excel at learning from a single data modality, biological systems are inherently multi-modal, involving interactions across sequence, structure, and function at various scales. Effectively integrating these diverse data types remains an open challenge in biological modeling. Often, biological data sources are treated independently, overlooking the interactions between systems. DNA sequences are often viewed in isolation, but DNA alone does not fully describe an organism’s phenotype. As Denis Noble[54, 55] and others emphasize, biological processes result from interactions across molecular, cellular, physiological and organismal scales[56].

Combining Multiple Forms of Biological Sequences

Multimodal approaches enhance GLMs’ understanding by incorporating epigenetic and transcriptomic data. The gLM2

model [57], trained on the Open MetaGenomic corpus, uses both nucleotide sequences and corresponding amino acid sequences from coding regions, along with strand direction information, helping process gene regulation and reading frames.

Hybrid approaches that combine multiple input modalities or architectures are increasingly common in biological design. For example, MMDiff[34] integrates both structural (using rigid body frames for rotation and translation) and sequence data to jointly generate nucleic acid-protein complexes. RFAA[24] tackles the complexity of modeling biomolecular systems, including proteins, nucleic acids, small molecules, and metals. RFAA uses a three-track network architecture inspired by RoseTTAFold2 [58] with (1) 1D sequence data for amino acids, nucleic acids, and atom types, (2) pairwise relationships like residue distances or bond types, and (3) 3D coordinate information, including atom positions and stereochemistry. RFAA also incorporates a conditional diffusion fine-tuning approach to generate binding proteins, starting from substructures like ligands.

Combining Sequence and Structure

Integrating structural information with sequence data can significantly improve model performance, especially for tasks where geometry and physical constraints are crucial. A common approach is to combine molecular graphs with sequence data, providing geometric context that complements raw sequences. Protein structures—secondary, tertiary, and quaternary—also offer essential spatial information that sequence data alone may not reveal. Additional inputs can include biological environment details, target receptors, or transcriptomic data linking RNA, DNA, and amino acid sequences.

For example, models like MMDiff [34], which generate nucleic acid-protein complexes, use both sequence and structural data for more accurate macromolecule modeling. This is intended to help with handling Intrinsically Disordered Regions is challenging since their structure cannot be fully inferred from sequence alone, underscoring the need for both representations. ProtMamba [28] enhances protein inpainting and de novo protein generation by using homologous sequences, without relying on multiple sequence alignments (MSAs). This approach captures evolutionary conservation and variability, providing valuable context for tasks like protein editing. The Gene Ontology (GO) [59] also provides a structured framework for annotating biological processes, molecular functions, and cellular components. Although models like Progen2 [22] partially leverage GO terms for conditioning, the full graph structure of the ontology, which allows for rich parent-child relationships, remains underutilized for generation.

Incorporating Natural Language and Domain Knowledge

One of the most appealing aspects of the FM paradigm, particularly with models like GPT-4, is their ability to interact via natural language, offering flexibility that is largely absent in biology-specific models. To bridge this gap, recent works have focused on integrating natural language

with biological sequence models, enabling tasks to be performed through natural language prompts. Notable examples include ChatNT [38] for genomics and Nach0 [36] and InstructBioMol [37] for chemical sequences.

ChatNT is a multimodal agent that combines biological sequence processing with NLP, handling DNA, RNA, and proteins through conversational input. It integrates the Nucleotide Transformer [60] with the Vicuna-7B language model using a Perceiver-based projection layer, allowing for seamless interaction in biological tasks through natural language. Nach0 extends this concept to SMILES molecular sequences, using training data from PubMed abstracts and patents. It has been applied to tasks like molecular property prediction, generation, and reaction prediction. In a case study, Nach0 was prompted to “Generate a random drug-like small inhibitor molecule for Janus Kinase 3 (JAK3) that contains a classic kinase hinge binding motif.” Despite lacking 3D structural information, Nach0 successfully generated eight valid molecules, with a discovery rate of 0.11%, compared to 1.53% from a structure-aware baseline. InstructBioMol takes this further by training on natural language, 2D molecular graphs, protein sequences, and 3D structures. Its Motif-Guided Multimodal Feature Extraction Module allows the model to handle tasks like molecule captioning, description-based molecule and protein generation, and protein property question answering.

Open Problems

Which Architectures for Which Biological Task:

Many studies introduce specific approaches (e.g., Transformers, diffusion, SSMs) for distinct applications, but there is little consensus on which architecture is best across biological tasks. Direct comparisons are challenging due to inconsistent evaluations and benchmarks. We propose exploring architectures that are underused in certain domains, such as diffusion models for textual molecular representations. Additionally, we would compare the effects of different pre-training objectives, especially since many current objectives are borrowed from NLP, where sequence behavior differs from biological data. Developing domain-specific, self-supervised objectives—both multi-modal and sequence-only—would help create better representations for biological tasks where additional modalities are unavailable.

Innovative Architectures to handle Data Limitations:

While FMs have shown strong performance in specific biological domains, scalability and generalization remain difficult due to the relatively small and fragmented datasets in biology. In NLP, large datasets have driven breakthroughs in models like GPT-3 [53], with scaling laws analyzed by Kaplan et al. [61] and Bahri et al. [62]. Similar trends have been observed in biological models [21, 63, 64], but addressing data limitations may require lighter architectures or novel methods for obtaining generalizable representations without massive datasets. This demands new strategies for scalable biological models.

Enhancing Generalization and Transfer to Low-Data Regimes: Finally, we would like to explore the important application transfer learning to help CLMs generalize to

classes of molecules that do not have a lot of available training data. Specifically, Quaternary Ammonium Compounds are a group of chemicals with important antibiotic properties but currently the number of available examples of molecules in this class is small. For example, we may experiment using techniques similar to Progen2’s [22] antibody generation approach combining fine-tuning with conditional generation to overcome the bias from pre-training. We would like to explore which architectures best enable transfer learning in this setting and additional techniques to guide generation toward the correct region of chemical space.

Better Integration of Biological Modalities:

A promising direction is the integration of various biological modalities, such as combining chemical graphs with SMILES representations in CLMs. Techniques like Adapters, used by Liu et al. [65] for molecular captioning, may also improve molecule generation. We would begin by evaluating the impact of these approaches by testing generated sequence validity and property control. Furthermore, we aim to explore pre-training models with Gene Ontology graph structures, extending Progen’s use of GO-derived keywords, to enhance protein generation.

Enabling Better Control in Generation:

We would like to push the boundaries of FMs’ capabilities to generalize to rare or physically unlikely property combinations in the biological sequence space. This would involve studying whether the difficulty in generating these combinations is due to their absence from training data or intrinsic physical constraints. This investigation connects to Ferruz et al.’s [23] concept of ”dark” regions in protein space, and could provide insight into improving generation for uncommon combinations.

Conclusion

The field is moving ahead rapidly; a large number of the cited papers here are from preprint servers such as arXiv and BioRxiv, highlighting the need for standardization of datasets and metrics. Furthermore, looking ahead, it is crucial for research to prioritize rigorous comparative evaluations of different model architectures for biological generation, more sophisticated integration of biological modalities, and improved strategies for generalization across diverse biological systems. This is especially timely, as the United States National Academies of Sciences, Engineering, and Medicine are founding a committee on Foundation Models for Scientific Discovery and Innovation, underscoring the growing importance of this area.[66] By addressing these challenges, FMs could truly revolutionize fields such as drug discovery, synthetic biology, and genetic engineering, accelerating breakthroughs and moving us closer to practical, AI-driven biological innovations that can meaningfully impact our understanding and manipulation of life.

Acknowledgements

This work was supported in part by the National Science Foundation Grant No. 2411529 and Grant No. 2310113.

References

- [1] Protein structure initiative. National Institute of General Medical Sciences, U.S. National Institutes of Health, 2000–2015. Available from <https://www.nigms.nih.gov/Research/specificareas/PSI>.
- [2] Francis S. Collins and Leslie Fink. The human genome project. *Alcohol Health and Research World*, 19(3):190–195, 1995.
- [3] Natalie L Dawson, Ian Sillitoe, Jonathan G Lees, Su Datt Lam, and Christine A Orengo. Cath-gene3d: generation of the resource and its use in obtaining structural and functional annotations for protein sequences. *Protein Bioinformatics: From Protein Modifications and Networks to Proteomics*, pages 79–110, 2017.
- [4] Molly Morgan Jones, Sophie Castle-Clarke, Daniel Brooker, Edward Nason, Farah Huzair, and Joanna Chataway. The structural genomics consortium: a knowledge platform for drug discovery: a summary. *Rand health quarterly*, 4(3), 2014.
- [5] Protein Data Bank. Protein data bank. *Nature New Biol*, 233(223):10–1038, 1971.
- [6] John Jumper, Richard Evans, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.
- [7] Alexandre Blanco-Gonzalez, Alfonso Cabezon, Alejandro Seco-Gonzalez, Daniel Conde-Torres, Paula Antelo-Riveiro, Angel Pineiro, and Rebeca Garcia-Fandino. The role of ai in drug discovery: challenges, opportunities, and strategies. *Pharmaceuticals*, 16(6):891, 2023.
- [8] Christopher A Voigt. Synthetic biology 2020–2030: six commercially-available products that are changing our world. *Nature Communications*, 11(1):1–6, 2020.
- [9] Tzu-Chieh Tang, Bolin An, Yuanyuan Huang, Sangita Vasikaran, Yanyi Wang, Xiaoyu Jiang, Timothy K Lu, and Chao Zhong. Materials design by synthetic biology. *Nature Reviews Materials*, 6(4):332–350, 2021.
- [10] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv e-prints*, pages arXiv–2108, 2021.
- [11] Ashish Vaswani. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [12] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021.
- [13] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.
- [14] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343, 2022.

- [15] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021.
- [16] Emiel Hoogeboom, Alexey A Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and Tim Salimans. Autoregressive diffusion models. In *International Conference on Learning Representations*, 2021.
- [17] Viraj Bagal, Rishal Aggarwal, PK Vinod, and U Deva Priyakumar. Molgpt: molecular generation using a transformer-decoder model. *Journal of Chemical Information and Modeling*, 62(9):2064–2076, 2021.
- [18] Jike Wang, Chang-Yu Hsieh, Mingyang Wang, Xiaorui Wang, Zhenxing Wu, Dejun Jiang, Benben Liao, Xujun Zhang, Bo Yang, Qiaojun He, et al. Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning. *Nature Machine Intelligence*, 3(10):914–922, 2021.
- [19] Eyal Mazuz, Guy Shtar, Bracha Shapira, and Lior Rokach. Molecule generation using transformers and policy gradient reinforcement learning. *Scientific Reports*, 13(1):8799, 2023.
- [20] Jannis Born and Matteo Manica. Regression transformer enables concurrent sequence regression and generation for molecular language modelling. *Nature Machine Intelligence*, 5(4):432–444, 2023.
- [21] Ali Madani, Ben Krause, Eric R Greene, Subu Subramanian, Benjamin P Mohr, James M Holton, Jose Luis Olmos, Caiming Xiong, Zachary Z Sun, Richard Socher, et al. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, 41(8):1099–1106, 2023.
- [22] Erik Nijkamp, Jeffrey A Ruffolo, Eli N Weinstein, Nikhil Naik, and Ali Madani. Progen2: exploring the boundaries of protein language models. *Cell systems*, 14(11):968–978, 2023.
- [23] Noelia Ferruz, Steffen Schmidt, and Birte Höcker. Protgpt2 is a deep unsupervised language model for protein design. *Nature communications*, 13(1):4348, 2022.
- [24] Rohith Krishna, Jue Wang, Woody Ahern, Pascal Sturmfels, Preetham Venkatesh, Indrek Kalvet, Gyu Rie Lee, Felix S Morey-Burrows, Ivan Anishchenko, Ian R Humphreys, et al. Generalized biomolecular modeling and design with rosettafold all-atom. *Science*, 384(6693):ead12528, 2024.
- [25] Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanasse, Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete diffusion. *Advances in neural information processing systems*, 36, 2024.
- [26] E Nguyen, M Poli, MG Durrant, AW Thomas, B Kang, J Sullivan, MY Ng, A Lewis, A Patel, A Lou, et al. Sequence modeling and design from molecular to genome scale with evo. *Science*, 2024.
- [27] Rıza Özçelik, Sarah de Ruiter, Emanuele Criscuolo, and Francesca Grisoni. Chemical language modeling with structured state space sequence models. *Nature Communications*, 15(1):6176, 2024.
- [28] Damiano Sgarbossa, Cyril Malbranke, and Anne-Florence Bitbol. Protmamba: a homology-aware but alignment-free protein state space model. *bioRxiv*, pages 2024–05, 2024.
- [29] Avantika Lal, David Garfield, Tommaso Biancalani, and Gokcen Eraslan. reglm: Designing realistic regulatory dna with autoregressive language models. In *International Conference on Research in Computational Molecular Biology*, pages 332–335. Springer, 2024.
- [30] Zehui Li, Yuhao Ni, William AV Beardall, Guoxuan Xia, Akashaditya Das, Guy-Bart Stan, and Yiren Zhao. Discdiff: Latent diffusion model for dna sequence generation. *CoRR*, 2024.
- [31] Simon Senan, Aniketh Janardhan Reddy, Zach Nussbaum, Aaron Wenteler, Matei Bejan, Michael I Love, Wouter Meuleman, and Luca Pinello. Dna-diffusion: Leveraging generative models for controlling chromatin accessibility and gene expression via synthetic regulatory elements. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*, 2024.
- [32] Xinmiao Peng and Fei Zhu. Hitting stride by degrees: Fine grained molecular generation via diffusion model. *Expert Systems with Applications*, 244:122949, 2024.
- [33] Sarah Alamdari, Nitya Thakkar, Rianne van den Berg, Alex Lu, Nicolo Fusi, Ava Amini, and Kevin Yang. Protein generation with evolutionary diffusion: sequence is all you need. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
- [34] Alex Morehead, Jeffrey Ruffolo, Aadyot Bhatnagar, and Ali Madani. Towards joint sequence-structure generation of nucleic acid and protein complexes with se (3)-discrete diffusion. In *Proceedings of the NeurIPS 2023 Workshop on Machine Learning in Structural Biology*, 2023.
- [35] Christof Angermueller, David Dohan, David Belanger, Ramya Deshpande, Kevin Murphy, and Lucy Colwell. Model-based reinforcement learning for biological sequence design. In *International conference on learning representations*, 2019.
- [36] Micha Livne, Zulfat Miftahutdinov, Elena Tutubalina, Maksim Kuznetsov, Daniil Polykovskiy, Annika Brundyn, Aastha Jhunjhunwala, Anthony Costa, Alex Aliper, Alán Aspuru-Guzik, et al. nach0: Multimodal natural and chemical languages foundation model. *Chemical Science*, 15(22):8380–8389, 2024.
- [37] Xiang Zhuang, Keyan Ding, Tianwen Lyu, Yinuo Jiang, Xiaotong Li, Zhuoyi Xiang, Zeyuan Wang, Ming Qin, Kehua Feng, Jike Wang, et al. Instruct-biomol: Advancing biomolecule understanding and design following human instructions. *arXiv preprint arXiv:2410.07919*, 2024.

- [38] Guillaume Richard, Bernardo P de Almeida, Hugo Dalla-Torre, Christopher Blum, Lorenz Hexemer, Priyanka Pandey, Stefan Laurent, Marie P Lopez, Alexander Laterre, Maren Lang, et al. Chatnt: A multimodal conversational agent for dna, rna and protein tasks. *bioRxiv*, pages 2024–04, 2024.
- [39] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- [40] Eric Nguyen, Michael Poli, Marjan Faizi, Armin Thomas, Michael Wornow, Callum Birch-Sykes, Stefano Massaro, Aman Patel, Clayton Rabideau, Yoshua Bengio, et al. Hyenadna: Long-range genomic sequence modeling at single nucleotide resolution. *Advances in neural information processing systems*, 36, 2024.
- [41] Zhilin Yang. Xlnet: Generalized autoregressive pre-training for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
- [42] Michael Moret, Lukas Friedrich, Francesca Grisoni, Daniel Merk, and Gisbert Schneider. Generative molecular design in low data regimes. *Nature Machine Intelligence*, 2(3):171–180, 2020.
- [43] Michael Poli, Stefano Massaro, Eric Nguyen, Daniel Y Fu, Tri Dao, Stephen Baccus, Yoshua Bengio, Stefano Ermon, and Christopher Ré. Hyena hierarchy: Towards larger convolutional language models. In *International Conference on Machine Learning*, pages 28043–28078. PMLR, 2023.
- [44] Bohao Xu, Yingzhou Lu, Yoshitaka Inoue, Namkyeong Lee, Tianfan Fu, and Jintai Chen. Protein-mamba: Biological mamba models for protein function prediction. *arXiv preprint arXiv:2409.14617*, 2024.
- [45] Zhangzhi Peng, Benjamin Schussheim, and Pranam Chatterjee. Ptm-mamba: A ptm-aware protein language model with bidirectional gated mamba blocks. *bioRxiv*, 2024.
- [46] Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. Efficient training of language models to fill in the middle. *arXiv preprint arXiv:2207.14255*, 2022.
- [47] Gonzalo Benegas, Chengzhong Ye, Carlos Albors, Jianan Canal Li, and Yun S Song. Genomic language models: Opportunities and challenges. *arXiv preprint arXiv:2407.11435*, 2024.
- [48] Daniel Y Fu, Tri Dao, Khaled Kamal Saab, Armin W Thomas, Atri Rudra, and Christopher Re. Hungry hungry hippos: Towards language modeling with state space models. In *The Eleventh International Conference on Learning Representations*, 2022.
- [49] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021.
- [50] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.
- [51] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Peter Chen, John Canny, Pieter Abbeel, and Yun Song. Evaluating protein transfer learning with tape. *Advances in neural information processing systems*, 32, 2019.
- [52] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [53] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [54] Denis Noble. A theory of biological relativity: no privileged level of causation. *Interface focus*, 2(1):55–64, 2012.
- [55] Denis Noble. It’s time to admit that genes are not the blueprint for life. *Nature*, 626(7998):254–255, 2024.
- [56] Jeremy Ramsden. *Bioinformatics: an introduction*. Springer Nature, 2023.
- [57] Andre Cornman, Jacob West-Roberts, Antonio Pedro Camargo, Simon Roux, Martin Beracocha, Milot Mirdita, Sergey Ovchinnikov, and Yunha Hwang. The omg dataset: An open metagenomic corpus for mixed-modality genomic language modeling. *bioRxiv*, pages 2024–08, 2024.
- [58] Minkyung Baek, Ivan Anishchenko, Ian R Humphreys, Qian Cong, David Baker, and Frank DiMaio. Efficient and accurate prediction of protein structure using rosettafold2. *BioRxiv*, pages 2023–05, 2023.
- [59] Michael Ashburner, Catherine A Ball, Judith A Blake, David Botstein, Heather Butler, J Michael Cherry, Allan P Davis, Kara Dolinski, Selina S Dwight, Janan T Eppig, et al. Gene ontology: tool for the unification of biology. *Nature genetics*, 25(1):25–29, 2000.
- [60] Hugo Dalla-Torre, Liam Gonzalez, Javier Mendoza-Revilla, Nicolas Lopez Carranza, Adam Henryk Grzywaczewski, Francesco Oteri, Christian Dallago, Evan Trop, Bernardo P de Almeida, Hassan Sirelkhatim, et al. The nucleotide transformer: Building and evaluating robust foundation models for human genomics. *BioRxiv*, pages 2023–01, 2023.
- [61] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [62] Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *Proceedings of the National Academy of Sciences*, 121(27):e2311878121, 2024.

- [63] Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C Lawrence Zitnick, Jerry Ma, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021.
- [64] Nathan C Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor W Coley, and Vijay Gadepally. Neural scaling of deep chemical models. *Nature Machine Intelligence*, 5(11):1297–1305, 2023.
- [65] Zhiyuan Liu, Sihang Li, Yanchen Luo, Hao Fei, Yixin Cao, Kenji Kawaguchi, Xiang Wang, and Tat-Seng Chua. Molca: Molecular graph-language modeling with cross-modal projector and uni-modal adapter. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15623–15638, 2023.
- [66] National Academies of Sciences Engineering and Medicine. Foundation models for scientific discovery and innovation: Opportunities across the department of energy, 2024. Accessed: 2024-10-22.