

**Diffusion Probabilistic Modeling of Protein  
Backbones in 3D for the Motif-Scaffolding problem**

by

Jason Yim

B.S., Johns Hopkins University (2018)

Submitted to the Department of Electrical Engineering and Computer  
Science

in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2023

© Massachusetts Institute of Technology 2023. All rights reserved.

Author .....  
Department of Electrical Engineering and Computer Science  
January 9, 2023

Certified by.....  
Tommi S. Jaakkola  
Professor of Electrical Engineering and Computer Science  
Thesis Supervisor

Certified by.....  
Regina Barzilay  
Professor of Electrical Engineering and Computer Science  
Thesis Supervisor

Accepted by .....  
Leslie A. Kolodziejcki  
Professor of Electrical Engineering and Computer Science  
Chair, Department Committee on Graduate Students



# Diffusion Probabilistic Modeling of Protein Backbones in 3D for the Motif-Scaffolding problem

by

Jason Yim

Submitted to the Department of Electrical Engineering and Computer Science  
on January 9, 2023, in partial fulfillment of the  
requirements for the degree of  
Master of Science

## Abstract

Construction of a *scaffold* structure that supports a desired *motif*, conferring protein function, shows promise for the design of vaccines and enzymes. But a general solution to this motif-scaffolding problem remains open. Current machine-learning techniques for scaffold design are either limited to unrealistically small scaffolds (up to length 20) or struggle to produce multiple diverse scaffolds. We propose to learn a distribution over diverse and longer protein backbone structures via an  $E(3)$ -equivariant graph neural network. We develop **SMCDiff** to efficiently sample scaffolds from this distribution conditioned on a given motif; our algorithm is the first to theoretically guarantee conditional samples from a diffusion model in the large-compute limit. We evaluate our designed backbones by how well they align with AlphaFold2-predicted structures. We show that our method can (1) sample scaffolds up to 80 residues and (2) achieve structurally diverse scaffolds for a fixed motif.

Thesis Supervisor: Tommi S. Jaakkola

Title: Professor of Electrical Engineering and Computer Science

Thesis Supervisor: Regina Barzilay

Title: Professor of Electrical Engineering and Computer Science



# Acknowledgments

I would like to first thank my advisors Tommi S. Jaakkola and Regina Barzilay, whose guidance was invaluable in pursuing the work stated in this thesis. I am thankful for their patience in exploring ideas, failing, and learning. The research environment and support in their respective groups made every day a joy to work at MIT.

I am thankful to the students in Tommi and Regina’s labs for interesting conversations, sharing of ideas, and a joyful work environment. The office I share with Felix Faltings, Gabriele Corso, and Hannes Stärk is a neverending discussion of machine learning, science, debugging, and Sisyphean discussions of acquiring a couch. The work in this thesis specially benefited from discussion and feedback with Rachel Wu, Jeremy Wohlwend, Wengong Jin, Nitan Shalon, Wenxian Shi, Sean Murphy, Bowen Jing, and John Yang.

I dedicate this work to Octavian Ganea, a talented and kind researcher who left us too soon. He introduced me to Brian Trippe, whom I am the most thankful for being my collaborator during the work in this thesis. Through Brian, I collaborated with David Baker’s group. Without Octavian, I would not have my current collaborators.

I thank David Baker for his insightful comments and infectious energy for always going after big, ambitious problems. His openness to collaborate led to friends and colleagues outside of MIT – Doug Tischer, Justas Dauparas, Jue Wang, David Jurgens, Joseph Watson, Nathaniel Bennett, Woody Ahern, Sidney Lianza, and the rest of members at the Institute for Protein design.

Finally, I would like to thank my loved ones. My mother and father for their sacrifice, love, and support for my career. My partner for her patience and love through my studies.



# Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Introduction</b>  | <b>9</b>  |
| 1.1      | Motif-scaffolding . . . . .  | 9         |
| 1.2      | Conditional sampling with diffusion models . . . . .                           | 10        |
| 1.3      | Related work . . . . .   | 12        |
| <b>2</b> | <b>Preliminaries</b>   | <b>15</b> |
| 2.1      | Problem formulation . . . . .  | 15        |
| 2.2      | Diffusion probabilistic models . . . . .                                       | 17        |
| <b>3</b> | <b>ProtDiff: A diffusion model of protein backbones in 3D</b>                  | <b>19</b> |
| 3.1      | Equivariant neural network of protein backbones . . . . .                      | 19        |
| 3.2      | Feature initialization . . . . .   | 20        |
| <b>4</b> | <b>SMCDiff: Conditional sampling in diffusion models by particle filtering</b> | <b>23</b> |
| 4.1      | Challenge of conditional sampling . . . . .                                    | 24        |
| 4.2      | Conditional sampling is a sequential Monte Carlo problem . . . . .             | 29        |
| <b>5</b> | <b>Experiments</b>   | <b>35</b> |
| 5.1      | ProtDiff training details . . . . .  | 36        |
| 5.2      | <i>In silico</i> evaluation of designed backbones . . . . .                    | 36        |
| 5.3      | Motif-scaffolding via conditional sampling . . . . .                           | 38        |
| 5.4      | Unconditional sampling . . . . .   | 41        |
| 5.4.1    | Sample diversity analysis . . . . .  | 43        |

|          |   |           |
|----------|---|-----------|
| 5.4.2    | Failure mode analysis . . . . .   | 44        |
| 5.4.3    | Noise interpolation analysis . . . . .  | 44        |
| 5.4.4    | Qualitative analysis of <code>scTM</code> in different ranges . . . . .           | 46        |
| 5.5      | Applicability of <code>SMCDiff</code> beyond proteins: MNIST inpainting . . . . . | 49        |
| <b>6</b> | <b>Conclusion</b>   | <b>51</b> |



# Chapter 1

## Introduction

The ability to engineer novel proteins holds promise in developing bio-therapeutics towards global health challenges such as SARS-COV-2 [Arunachalam et al., 2021] and cancer [Quijano-Rubio et al., 2020]. Unfortunately, efforts to engineer proteins requires substantial domain knowledge and laborious trial and error. As reviewed in Kuhlman and Bradley [2019], advancements in deep learning have paralleled improvements in protein engineering by automating knowledge acquisition from data and improving efficiency in designing proteins. The present work addresses a deep learning method for an important task in protein design called *motif-scaffolding*.

### 1.1 Motif-scaffolding

A central task in protein design is creation of a stable *scaffold* to support a target *motif*. Here, motifs are structural protein fragments imparting biological function while scaffolds stabilize the motif’s structure. Vaccines and enzymes have already been designed by solving certain instances of this *motif-scaffolding* problem [Procko et al., 2014, Correia et al., 2014, Jiang et al., 2008, Siegel et al., 2010]. However, successful solutions to this problem in the past have necessitated substantial expert involvement and laborious trial and error. Machine learning (ML) offers the hope to automate, and better direct this search. But existing ML approaches face one of two major roadblocks. First, many method do not build scaffolds longer than about 20

residues. For many motif sizes of interest, the resulting proteins would be smaller than the shortest commonly-studied simple protein folds (35–40 residues) [Gelman and Gruebele, 2014]. Second, while other methods may generate longer scaffolds using stochastic search techniques, they require hours of computation to generate a single plausible scaffold [Wang et al., 2022, Anishchenko et al., 2021, Tischer et al., 2020]. Moreover, when a plausible scaffold is found, it remains to be experimentally validated. Therefore, it is desirable to return not just a single scaffold but rather a set of scaffolds exhibiting diverse sequences and structural variation to increase the likelihood of success in practice.

## 1.2 Conditional sampling with diffusion models

In the present work, we demonstrate the promise of a particular generative modeling approach within ML for efficiently returning a diverse set of motif-supporting scaffolds. Generative models have been shown to capture a distribution over diverse protein structures [Lin et al., 2021]. But it is not clear how to handle conditioning (on the motif) using these approaches. Diffusion probabilistic models (DPMs) offer a potential alternative; not only do they provide a more straightforward path to handling conditioning, but they have also enjoyed success generating small-molecules in 3D [Hoogeboom et al., 2022]. Extending DPMs to protein structures, though, is non-trivial; since proteins are larger than small molecules, modeling proteins requires handling the sequential ordering of residues and long-range interactions. Finally, while existing models often generate distance matrices [Anand and Huang, 2018, Lin et al., 2021], we instead focus on generating a full set of 3D coordinates, which should improve designability in practice. Our resulting model, `ProtDiff`, is similar to concurrent work on  $E(3)$ -equivariant diffusion models for molecules [Hoogeboom et al., 2022], but with modifications specific to protein structure. Moreover, we develop a novel motif-scaffolding procedure based on Sequential Monte Carlo, `SMCDiff`, that repurposes an unconditionally trained DPM for conditional sampling. We prove that if a DPM matches the data distribution, `SMCDiff` is guaranteed to provide *exact*

conditional samples in a large-compute limit; this property contrasts with previous methods [Song et al., 2021, Zhou et al., 2021], which we show introduce non-trivial approximation error that impedes performance. Our final motif-scaffolding generative framework, then, has two steps (Fig. 1-1): first we train `ProtDiff` to learn a distribution over protein backbones, and then we use `SMCDiff` with `ProtDiff` to inpaint arbitrary motifs. Using sequence design and protein folding methods, we evaluate the quality of our backbone samples.

Ours is the first machine-learning method to construct scaffolds longer than 20 residues around motifs — we build up to 80 residues scaffolds on a test case. Beyond our progress on the motif-scaffolding problem, we provide the following technical contributions: (1) we introduce a protein-backbone generative model in 3D – with the ability to generate backbone samples that structurally agree with `AlphaFold2` predictions, and (2) we develop a novel conditional sampling algorithm for inpainting.

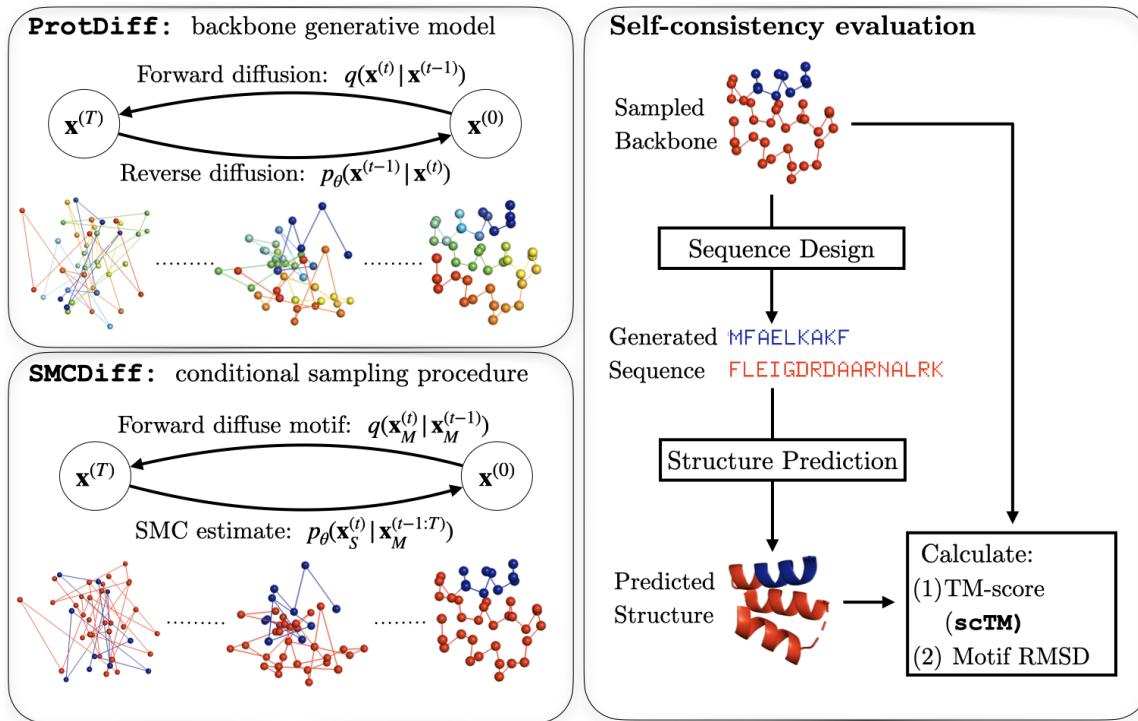


Figure 1-1: Overview of the conditional generative modeling approach to the motif-scaffolding problem. We train our new protein backbone diffusion model, **ProtDiff**, to generate realistic protein backbone structures. Next, we run **SMCDiff**, our conditional sampling algorithm, with **ProtDiff** to generate scaffolds (colored in red) conditioned on the motif (colored in blue). For self-consistency evaluation, we use a pretrained fixed-backbone sequence-design model (**ProteinMPNN** [Dauparas et al., 2022]) to generate the scaffold sequence from a sampled backbone. We then input the sequence to a structure prediction model, in our case **AlphaFold2** (AF2) [Jumper et al., 2021], to generate the full protein structure from the generated sequence. We compare the backbone of the predicted structure with the original backbone structure using TM-score [Xu and Zhang, 2010] and root-mean-square-distance (RMSD) for the motif.

### 1.3 Related work

**Motif-scaffolding.** Past approaches have sought to scaffold a motif with native or prespecified protein fragments, but are limited to finding a suitable match in the Protein Data Bank (PDB) and cannot adapt the scaffold to compensate for slight structural mismatches [Cao et al., 2022, Silva et al., 2016, Yang et al., 2021, Sesterhenn et al., 2020, Linsky et al., 2020]. More recently Wang et al. [2022] used pre-trained protein structure prediction networks to recapitulate native scaffolds, but this method

failed to generate scaffolds longer than 20 residues and can output only a single candidate scaffold rather than a diverse set. By contrast, our goal is to sample diverse, long scaffolds.

**Diffusion models for molecule generation.** Several concurrent works have extended equivariant diffusion models to molecule generation. [Anand and Achim \[2022\]](#) extended diffusion models for generation of protein backbone frames and sequences *conditioned* on secondary structure adjacency matrices. Similarly, [Luo et al. \[2022\]](#) focused on CDR-loop generation using diffusion models conditioned on non-CDR regions of the antibody-antigen. Our method does not require conditioning and is applicable to general proteins. [Lee and Kim \[2022\]](#) approach the same problem as our work but build diffusion models over 2D distances matrices that requires post-processing to produce 3D structures through Rosetta minimization. We demonstrate capability of diffusion models to directly model 3D coordinates of proteins. [Hoogboom et al. \[2022\]](#) developed an equivariant diffusion model (EDM) for generating small molecules in 3D. However, because EDM does not enforce a spatial ordering of the atoms that compose small molecules, as we describe in Chapter 3, it does not learn a coherent chain structure as needed in proteins.

**Inpainting and conditional sampling in diffusion models.** Point-Voxel Diffusion (PVD) [[Zhou et al., 2021](#)] is a 3D diffusion model for generating shapes from the ShapeNet dataset. Though trained to generate shapes unconditionally, PVD completes (or *inpaints*) full shapes when a partial point cloud is fixed during inference. For general diffusion models, [Song et al. \[2021\]](#) proposed an alternative inpainting approach and remarked that this approach produces *approximate* conditional samples. However, these methods do not provide theoretical guarantees, and when we compare them to SMCDiff, we find that their approximation error impedes performance when applied to motif-scaffolding. [Saharia et al. \[2021\]](#) developed an inpainting diffusion model by training a diffusion model to denoise randomly generated masked regions while unmasked regions were unperturbed. However, their approach requires a detailed data augmentation strategy that does not exist for proteins.

**Generative models of proteins.** Following the success of deep language models, [Ferruz et al. \[2022\]](#) developed protein sequence models to generate new proteins, but these models do not allow specification of structural motifs. Another class of methods, referred to as fixed backbone sequence design [[Fleishman et al., 2011](#), [Ingraham et al., 2019](#), [Xiong et al., 2020](#), [McPartlon et al., 2022](#), [Hsu et al., 2022](#)], attempts to solve the problem of identifying a sequence that folds into any given designable backbone structure. In the present work, we utilize a particular sequence design method, ProteinMPNN [[Dauparas et al., 2022](#)], but in principle any other fixed-backbone sequence design method could be used in its place. [Anand and Huang \[2018\]](#), [Lin et al. \[2021\]](#), [Wu et al. \[2021\]](#) propose generative adversarial networks, variational autoencoders, and energy-based models, respectively, on distance matrices, but these approaches (1) do not generate backbones compatible with a specified motif and (2) rely on an unwieldy optimization step to translate the distance matrix into backbone coordinates. Other authors use neural net [[Tischer et al., 2020](#), [Anishchenko et al., 2021](#), [Wang et al., 2022](#), [Huang et al., 2022](#), [Wu et al., 2021](#)], but require a computationally challenging conformational landscape exploration.

# Chapter 2

## Preliminaries

### 2.1 Problem formulation

A protein can be represented by its amino acid sequence and backbone structure. Let  $\mathcal{A}$  be the set of 20 genetically-encoded amino acids. We denote the sequence of an  $N$ -residue protein by  $s \in \mathcal{A}^N$  and its C- $\alpha$  backbone coordinates in 3D by  $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N,3}$ . We describe a protein as having a fixed structure that is a function of its sequence, so we may write  $\mathbf{x}(s)$ . We divide the  $N$  residues into the functional motif  $\mathcal{M}$  and the scaffold  $\mathcal{S}$ , such that  $\mathcal{M} \cup \mathcal{S} = \{1, 2, \dots, N\}$ . The goal is to identify, given the motif structure  $\mathbf{x}_{\mathcal{M}}$ , sequences  $s$  whose structure recapitulates the motif to high precision  $\mathbf{x}(s)_{\mathcal{M}} \approx \mathbf{x}_{\mathcal{M}}$ . Our problem formulation makes some assumptions and heuristics that we describe next.

**Protein sequence-structure relationship.** Generally speaking, a protein’s sequence encodes an ensemble of conformations, populated to different degrees at biological temperatures. Anfinsen’s hypothesis states that the ground state conformation is thermodynamically accessible [Anfinsen, 1973], providing a mapping from sequence to a unique (ground state) structure. In practice, the ground state structures make up the vast majority of experimentally determined protein conformations, as over 95% of structures in the Protein Data Bank (PDB) are collected at cryogenic temperatures [Fraser et al., 2011]. Thus we simplify our problem by saying that a sequence uniquely

maps to a static structure (i.e. the ground state structure). However, violations of this assumption arise in some PDB structures as a result of (1) of context specific determinants of structure such as post-translational modifications and environmental factors including pH, binding partners, and salts, as well as (2) thermodynamic inaccessibility of the ground state.

**Motif sequence and side-chains.** We assume we may represent a functional motif by the coordinates of its C- $\alpha$  atoms. However, the biochemical functions of proteins depend not only on backbone structure, but also on side-chains. For example, the activity of many enzymes is imparted by triplets of residues, known as *catalytic triads*, whose ability to catalyze reactions depends on the spatial organization of side-chain atoms. Our problem statement and subsequent evaluation scheme are agnostic to the amino acid identity of motif residues, let alone side-chain positioning. A more complete representation of a motif would include the side-chain identities (i.e. the amino acid *sequence*) and side-chain atom coordinates.

**Scaffold length and motif placement.** We have additionally assumed that the size of scaffolds and the indices of motif residues within the backbone chain,  $\mathcal{M}$ , are known a priori. However, in practice satisfactory scaffolds could have different lengths and different motif placements, and typically it is not known a priori what lengths and placements will be best. Previous works have addressed this challenge through brute force by sampling multiple lengths and placements, and relied on post-hoc filtering to identify the most promising scaffolds [Wang et al., 2022, Yang et al., 2021]. Subsequent work on ML methods could potentially generalize beyond this assumption to efficiently sample appropriate scaffold lengths and motif placements.

**Sequence and side-chain modeling.** ProtDiff models only the backbone coordinates and leaves sequence design to a subsequent stage, for which we have used ProteinMPNN. A more complete representation of a proteins could include both sequence and structure (where structure can be divided into the backbone and side-chain atom coordinates). To model sequence, we rely on a separately trained neural



network, ProteinMPNN, but this is not ideal. Unless ProtDiff produces perfect backbones, one would expect the backbone samples of ProtDiff to present a substantial domain shift when used as input for ProteinMPNN.

**3D backbone representation.** In this work, we represent a protein structure using the C- $\alpha$  coordinates of every residue along the backbone. However, this representation is coarse-grained and ignores additional backbone atomic coordinates, namely the backbone carbon and nitrogen atoms. Dauparas et al. [2022] observed additionally modeling the heavy atoms of the backbone nitrogen and carbon atoms along with the C- $\beta$  of every residue (to capture side-chain information) improved performance (by sequence recovery) for fixed-backbone sequence design. We hypothesize modeling additional coordinates of every residue would also improve designability performance of ProtDiff. Constraining ProtDiff to place the remaining atoms in the correct orientation could help enforce correct chirality and mitigate chain breaks.

## 2.2 Diffusion probabilistic models

Our approach to the motif-scaffolding problem builds on denoising diffusion probabilistic models (DPMs) [Sohl-Dickstein et al., 2015]. We follow the conventions and notation set by Ho et al. [2020], which we review here. DPMs are a class of generative models based on a reversible, discrete-time diffusion process. The *forward process* starts with a sample  $\mathbf{x}^{(0)}$  from an unknown data distribution  $q$ , with density denoted by  $q(\mathbf{x}^{(0)})$ , and iteratively adds noise at each step  $t$ . By the last step,  $T$ , the distribution of  $\mathbf{x}^{(T)}$  is indistinguishable from an isotropic Gaussian:  $\mathbf{x}^{(T)} \sim \mathcal{N}(\mathbf{x}^{(T)}; 0, \mathbf{I})$ . Specifically, we choose a variance schedule  $\beta^{(1)}, \beta^{(2)}, \dots, \beta^{(T)}$ , and define the transition distribution at step  $t$  as  $q(\mathbf{x}^{(t)} | \mathbf{x}^{(t-1)}) = \mathcal{N}(\mathbf{x}^{(t)}; \sqrt{1 - \beta^{(t)}}\mathbf{x}^{(t-1)}, \beta^{(t)}\mathbf{I})$ .

DPMs approximate  $q$  with a second distribution  $p_\theta$  by learning the transition distribution of the *reverse process* at each  $t$ ,  $p_\theta(\mathbf{x}^{(t-1)} | \mathbf{x}^{(t)})$ . We follow the conventions set by Ho et al. [2020] in our parameterization and choice of objective. In particular, we take  $p_\theta(\mathbf{x}^{(t-1)} | \mathbf{x}^{(t)}) = \mathcal{N}(\mathbf{x}^{(t-1)}; \mu_\theta(\mathbf{x}^{(t)}, t), \beta^{(t)}\mathbf{I})$  with  $\mu_\theta(\mathbf{x}^{(t)}, t) =$

$\frac{1}{\sqrt{\alpha^{(t)}}} \left( \mathbf{x}^{(t)} - \frac{\beta^{(t)}}{\sqrt{1-\bar{\alpha}^{(t)}}} \epsilon_{\theta}(\mathbf{x}^{(t)}, t) \right)$ ,  $\alpha^{(t)} := 1 - \beta^{(t)}$ , and  $\bar{\alpha}^{(t)} := \prod_{s=1}^t \alpha^{(s)}$ . We implement  $\epsilon_{\theta}(\mathbf{x}^{(t)}, t)$  as a neural network. For training, we marginally sample  $\mathbf{x}^{(t)} \sim q(\mathbf{x}^{(t)} | \mathbf{x}^{(0)})$  from the forward process as  $\mathbf{x}^{(t)} = \sqrt{\bar{\alpha}^{(t)}} \mathbf{x}^{(0)} + \sqrt{1 - \bar{\alpha}^{(t)}} \epsilon$  and minimize the objective  $T^{-1} \sum_{t=1}^T \mathbb{E}_{q(\mathbf{x}^{(0)}, \mathbf{x}^{(t)})} [\|\epsilon - \epsilon_{\theta}(\mathbf{x}^{(t)}, t)\|^2]$  by stochastic optimization [Ho et al., 2020, Algorithm 1]. To generate samples from  $p_{\theta}(\mathbf{x}^{(0)})$ , we simulate the reverse process. That is, we sample noise for time  $T$  as  $\mathbf{x}^{(T)} \sim \mathcal{N}(0, \mathbf{I})$ , and then for each  $t = T - 1, \dots, 0$ , we simulate progressively “de-noised” samples as  $\mathbf{x}^{(t)} \sim p_{\theta}(\mathbf{x}^{(t)} | \mathbf{x}^{(t+1)})$ .

# Chapter 3

## ProtDiff: A diffusion model of protein backbones in 3D

Implementation of diffusion probabilistic models requires choosing an architecture for the neural network  $\epsilon_\theta(\mathbf{x}^{(t)}, t)$  introduced abstractly in Section 2.2. In this section we describe ProtDiff, which corresponds to the choice of  $\epsilon_\theta(\mathbf{x}^{(t)}, t)$  as a translation and rotation equivariant graph neural network tailored to modeling protein backbones.

### 3.1 Equivariant neural network of protein backbones

The properties and functions of proteins are dictated by the relative geometry of their residues, and are invariant to the coordinate system chosen to encode them. Recent work on neural network modeling of 3D data has found, both theoretically and empirically, that neural networks constrained to satisfy geometric invariances can provide inductive biases that improve generalization and training efficiency [Batzner et al., 2022]. Motivated by this observation, we parameterize  $\epsilon_\theta$  by an equivariant graph neural network (EGNN) [Satorras et al., 2021], which in 3D is equivariant to transformations in the Euclidean group. Xu et al. [2022] proved that if  $\epsilon_\theta$  is equivariant to a group then  $p_\theta$  is invariant to the same group.

**Tailoring EGNN to protein backbones.** We now describe our EGNN implementation, which we tailor to protein backbones and DPMS through the choice of edge and node features. To model every pairwise residue interaction, we represent backbones by a fully connected graph. Each node in the graph is indexed by  $n = 1, \dots, N$ , and corresponds to a residue. We associate each node with coordinates  $\mathbf{x}_n \in \mathbb{R}^3$  and  $D$  features  $h_n \in \mathbb{R}^D$ . For each pair of nodes  $n, n'$  we define an edge and associate it with edge features.

We construct our EGNN by stacking  $L$  equivariant graph convolutional layers (EGCL). Each layer  $l = 1, \dots, L$  defines an update as  $(\mathbf{x}^l, h^l) = \text{EGCL}[\mathbf{x}^{l-1}, h^{l-1}]$  where for each node  $n$

$$\begin{aligned} \mathbf{x}_n^l &= \mathbf{x}_n^{l-1} + \sum_{n' \neq n} \vec{\omega}_{nn'} \cdot \phi_{\mathbf{x}}(h_n^{l-1}, h_{n'}^{l-1}, d_{nn'}, a_{nn'}) \quad \text{and} \quad h_n^l = \phi_h(h_n^{l-1}, m_n), \quad \text{for} \\ \vec{\omega}_{nn'} &= \frac{\mathbf{x}_n^{l-1} - \mathbf{x}_{n'}^{l-1}}{\sqrt{d_{nn'} + \gamma}}, \quad m_n = \sum_{n' \neq n} \phi_e(h_n^{l-1}, h_{n'}^{l-1}, d_{nn'}, a_{nn'}), \quad \text{and} \quad d_{nn'} = \|\mathbf{x}_n^{l-1} - \mathbf{x}_{n'}^{l-1}\|_2^2. \end{aligned}$$

$\phi_e, \phi_h$ , and  $\phi_{\mathbf{x}}$  are fully connected neural networks, and  $\gamma$  is a small positive constant included for numerical stability. The first EGCL layer takes in initial node embeddings,  $h^0$  while edge embeddings,  $a_{nn'}$ , are kept fixed throughout.

We write the output of EGNN after  $L$  layers as  $\hat{\mathbf{x}} = \text{EGNN}[\mathbf{x}, h]$ . In the context of diffusion models, we predict the noise at time  $t$  with the following parameterization:

$$\epsilon_{\theta}(\mathbf{x}^{(t)}, t) = \hat{\mathbf{x}} - \mathbf{x}^{(t)}, \quad \hat{\mathbf{x}} = \text{EGNN}[\mathbf{x}^{(t)}, h(t)]. \quad (3.1)$$

## 3.2 Feature initialization

We now describe our choice of node and edge features. Our choice is motivated by the linear chain structure of protein backbones; residues close in sequence are necessarily close in 3D space. To allow this chain constraint to be learned more easily, we fix an ordering of nodes in the graph to correspond to sequence order. We include as edge features positional offsets as done in [Ingraham et al. \[2019\]](#), which we represent using

sinusoidal positional encoding features [Vaswani et al., 2017]. For node features, we similarly use a sinusoidal encoding of sequence position as well as of the diffusion time step  $t$  following Kingma et al. [2021]. We additionally process the time encoding to be orthogonal to the positional encoding.

**Initial node and edge embeddings.** Each edge between two residues indexed in the sequence by  $(n, n')$  is featurized with  $D$  features obtained through a sinusoidal encoding of its relative offset:

$$a_{nn'} = \begin{bmatrix} \varphi(n - n', 1) \\ \vdots \\ \varphi(n - n', D) \end{bmatrix}, \text{ where } \varphi(x, k) = \begin{cases} \sin(x \cdot \pi / N^{2 \cdot k / D}), & k \bmod 2 = 0 \\ \cos(x \cdot \pi / N^{2 \cdot (k-1) / D}), & k \bmod 2 = 1. \end{cases}$$

For node features, we similarly use a sinusoidal encoding of sequence position as well as of the diffusion time step  $t$  as

$$h_n(t) = \begin{bmatrix} \varphi(n, 1) \\ \vdots \\ \varphi(n, D) \end{bmatrix} + R \begin{bmatrix} \varphi(t, 1) \\ \vdots \\ \varphi(t, D) \end{bmatrix},$$

where  $R$  is a  $D \times D$  orthogonal matrix chosen uniformly at random. Intuitively, applying  $R$  transforms the time encoding to be orthogonal to the positional encoding.

**Coordinate scaling** While protein structures are typically parameterized in Angstroms, we transform the input protein coordinates to be in nanometers rather by dividing by 10. This scaling brings the backbones to a spatial scale similar to the reference distribution at which the forward noising process is stationary, a unit variance isotropic Gaussian. Importantly, the distribution of the final step  $T$  is indistinguishable from an isotropic Gaussian ( Fig. 3-1.)




Figure 3-1: Distribution of  $x^{(T)}$  after centering and scaling  $x^{(0)}$  to nanometers.

# Chapter 4

## SMCDiff: Conditional sampling in diffusion models by particle filtering

The second stage of our generative modeling approach to the motif-scaffolding problem is to sample scaffolds  $\mathbf{x}_S^{(0)}$  from  $p_\theta(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)})$ . Section 4.1 discusses the intractability of sampling from  $p_\theta(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)})$  exactly and the limitations of a simple approximation introduced by [Song et al., 2021]. In Section 4.2, we then formulate computation of  $p_\theta(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)})$  as a sequential Monte Carlo (SMC) problem [Doucet et al., 2001] and approximate it with a particle filtering algorithm (Algorithm 2).

**Notation.** In the following, we require notation that is more precise than in previous sections. For each  $t = 0, \dots, T$ , we let  $q_t(\cdot)$  and  $p_t(\cdot)$  denote the density functions of  $\mathbf{x}^{(t)}$  according to the forward process and to our neural network approximation of the reverse process, respectively. We denote densities restricted to the motif and scaffold with subscripts  $\mathcal{M}$  and  $\mathcal{S}$ . For example, we here write  $p_{\mathcal{M},t}(\mathbf{x}_M^{(t)})$ , whereas we wrote  $p_\theta(\mathbf{x}_M^{(t)})$  in the main text. We write (random) conditional densities as  $q_{\mathcal{M},t}(\cdot | \mathbf{x}_M^{(t-1)})$  and write the (deterministic) conditional density for an observation  $\mathbf{x}_M^{(t-1)} = x_M$  as  $q_{\mathcal{M},t}(\cdot | \mathbf{x}_M^{(t-1)} = x_M)$ .

An object of interest will be the Kullback-Leibler (KL) divergence. We write  $\text{KL}[q_t(\cdot) || p_t(\cdot)] := \int q_t(x) \log \frac{q_t(x)}{p_t(x)} dx$ , where  $\log(\cdot)$  is the natural (base  $e$ ) logarithm. We will also encounter the expected KL between conditional densities, which we will

write as  $\text{EKL} [q_t(\cdot | \mathbf{x}^{(t-1)}) || p_t(\cdot | \mathbf{x}^{(t-1)})] := \int q_{t-1}(x) \text{KL} [q_t(\cdot | \mathbf{x}^{(t-1)} = x) || p_t(\cdot | \mathbf{x}^{(t-1)} = x)] dx$ , where the outer expectation is taken with respect to the unconditional density associated with first argument of  $\text{EKL} [\cdot || \cdot]$ .

## 4.1 Challenge of conditional sampling

The conditional distributions of a DPM are defined implicitly through the steps of the reverse process. We may write the conditional density explicitly as

$$p_\theta(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)}) \propto p_\theta(\mathbf{x}_S^{(0)}, \mathbf{x}_M^{(0)}) = p_\theta(\mathbf{x}^{(0)}) = \int p_\theta(\mathbf{x}^{(T)}) \prod_{t=0}^{T-1} p_\theta(\mathbf{x}^{(t)} | \mathbf{x}^{(t+1)}) d\mathbf{x}^{(1:T)}.$$

However, the high-dimensional integral on the right-hand side above is intractable (both analytically and numerically) to compute.

To overcome this intractability, we build on the work of [Song et al. \[2021\]](#), who introduced a practical algorithm that generates *approximate* conditional samples. This strategy is to (1) forward diffuse the conditioning variable to obtain  $\mathbf{x}_M^{(1:T)} \sim q(\mathbf{x}_M^{(1:T)} | \mathbf{x}_M^{(0)})$ , and then (2) for each  $t$ , sample  $\mathbf{x}_S^{(t)} \sim p_\theta(\mathbf{x}_S^{(t)} | \mathbf{x}_M^{(t+1)}, \mathbf{x}_S^{(t+1)})$ . We call this approach the *replacement method* (following [Ho et al. \[2022\]](#)).

Although [Song et al. \[2021\]](#) notes that this approach can be understood as *approximate* conditional sampling, they provide no discussion of approximation error. We here show that the replacement method introduces irreducible error that is inherent to the forward process. Algorithm 1 provides an explicit description of the replacement method.

The first return of Algorithm 1,  $\mathbf{x}_S^{(0)}$ , is used as a putative inpainting solution or approximate conditional sample. But Algorithm 1 additionally returns subsequent time steps,  $\mathbf{x}^{(1:T)}$ .



---

**Algorithm 1** Replacement method for approximate conditional sampling

---

```
1: Input:  $\mathbf{x}_{\mathcal{M}}^{(0)}$  (motif)
2: // Forward diffuse motif
3:  $\check{\mathbf{x}}_{\mathcal{M}}^{(1:T)} \sim q(\mathbf{x}_{\mathcal{M}}^{(1:T)} | \mathbf{x}_{\mathcal{M}}^{(0)})$ 
4:
5: // Reverse diffuse scaffold
6:  $\mathbf{x}^{(T)} \sim p_{\theta}(\mathbf{x}^{(T)})$ 
7: for  $t = T, \dots, 1$  do
8:   // Replace with forward diffused motif
9:    $\mathbf{x}^{(t)} \leftarrow [\check{\mathbf{x}}_{\mathcal{M}}^{(t)}, \mathbf{x}_{\mathcal{S}}^{(t)}]$ 
10:
11:  // Propose next step
12:   $\mathbf{x}^{(t-1)} \sim p_{\theta}(\mathbf{x}^{(t-1)} | \mathbf{x}^{(t)})$ 
13: end for
14: Return  $\mathbf{x}_{\mathcal{S}}^{(0)}, \mathbf{x}^{(1:T)}$ 
```

---

We denote the approximation over all steps implied by the generative procedure in Algorithm 1 by  $p_{1:T}^{\text{Repl}}(\cdot | \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})$  and compare it to the exact conditional,  $q_{1:T}(\cdot | \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})$ . We here consider error in KL divergence because it permits an analytically tractable and transparent analysis. We additionally consider the idealized scenario where  $p_{0:T}(\cdot)$  perfectly captures the reverse process. Under this condition, the forward KL takes a surprisingly simple form.

**Proposition 4.1.1.** *Suppose that  $p_{0:T}(\cdot)$  exactly matches the forward diffusion process such that for every  $x$ ,  $p_t(\cdot | \mathbf{x}^{(t+1)} = x) = q_t(\cdot | \mathbf{x}^{(t+1)} = x)$ . Then for any motif  $x_{\mathcal{M}}$ ,*

$$\begin{aligned} & \text{KL} \left[ q_{1:T}(\cdot | \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \| p_{1:T}^{\text{Repl}}(\cdot | \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \right] \\ &= \sum_{t=1}^{T-1} \text{EKL} \left[ q_{\mathcal{S},t}(\cdot | \mathbf{x}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \| q_{\mathcal{S},t}(\cdot | \mathbf{x}^{(t+1)}) \right]. \end{aligned} \tag{4.1}$$

*Proof.* The result obtains from recognizing where the replacement method approximation agrees with the forward process, using conditional independences in both processes, and applying the chain rule for KL divergences. We make this explicit in

the derivation below, with comments explaining the transition to the following line.

$$\begin{aligned} & \text{KL} \left[ q_{1:T}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \parallel p_{1:T}^{\text{Repl}}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \right] \\ &= \int q_{1:T}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \log \frac{q_{1:T}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})}{p_{1:T}^{\text{Repl}}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})} dx^{(1:T)} \end{aligned}$$

// By the chain rule of probability.

$$\begin{aligned} &= \int q_{1:T}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \left[ \log \frac{q_{\mathcal{M},1:T}(x_{\mathcal{M}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})}{p_{1:T}^{\text{Repl}}(x_{\mathcal{M}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}})} + \right. \\ &\quad \left. \log \frac{q_{\mathcal{S},1:T}(x_{\mathcal{S}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})}{p_{\mathcal{S},1:T}^{\text{Repl}}(x_{\mathcal{S}}^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})} \right] dx^{(1:T)} \end{aligned}$$

// By the agreement of  $q$  and  $p_{\text{Repl}}$  on the motif, and the chain rule of probability.

$$\begin{aligned} &= \int q_{1:T}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \left[ \log \frac{q_{\mathcal{S},T}(x_{\mathcal{S}}^{(T)} \mid \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})}{p_T^{\text{Repl}}(x_{\mathcal{S}}^{(T)} \mid \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})} + \right. \\ &\quad \left. \sum_{t=1}^{T-1} \log \frac{q_{\mathcal{S},t}(x_{\mathcal{S}}^{(t)} \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})}{p_{\mathcal{S},t}^{\text{Repl}}(x_{\mathcal{S}}^{(t)} \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0:T)} = x_{\mathcal{M}}^{(0:T)})} \right] dx^{(1:T)} \end{aligned}$$

// Because  $q_{\mathcal{S},T}(\cdot) = p_{\mathcal{S},T}^{\text{Repl}}(\cdot) = \mathcal{N}(\cdot; 0, I)$  and the assumption that  $p_{\theta}$  matches  $q$ .

$$\begin{aligned} &= \int q_{1:T}(x^{(1:T)} \mid \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}) \left[ \sum_{t=1}^{T-1} \log \frac{q_{\mathcal{S},t}(x_{\mathcal{S}}^{(t)} \mid \mathbf{x}^{(t+1)} = x^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}^{(0)})}{q_{\mathcal{S},t}(x_{\mathcal{S}}^{(t)} \mid \mathbf{x}^{(t+1)} = x^{(t+1)})} \right] dx^{(1:T)} \\ &= \sum_{t=1}^{T-1} \text{EKL} \left[ q_{\mathcal{S},t}(\cdot \mid \mathbf{x}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}^{(0)}) \parallel q_{\mathcal{S},t}(\cdot \mid \mathbf{x}^{(t+1)}) \right]. \end{aligned}$$

□

Proposition 4.1.1 reveals that the replacement method introduces approximation error that is intrinsic to the forward process and cannot be eliminated by making  $p_{0:T}(\cdot)$  more expressive. Although the individual terms in the right hand side of Equation (4.1) are not analytically tractable in general, in the following corollary we show that this approximation error can be non-trivial by considering a special case. For this following example, we depart from the earlier assumption that  $\mathbf{x}$  is in 3D, and consider scalar valued  $\mathbf{x}_{\mathcal{M}}$  and  $\mathbf{x}_{\mathcal{S}}$ .

**Corollary 4.1.1.** *Suppose  $[\mathbf{x}_{\mathcal{M}}^{(0)}, \mathbf{x}_{\mathcal{S}}^{(0)}]$  is bivariate normal distributed with mean zero, unit variance, and covariance  $\rho$ . Further suppose that  $q_{\mathcal{S},t}(\cdot \mid \mathbf{x}_{\mathcal{S}}^{(0)}) = \mathcal{N}(\cdot; \sqrt{\bar{\alpha}^{(t)}} \mathbf{x}_{\mathcal{S}}^{(0)}, 1 -$*

$\bar{\alpha}^{(t)}$  and  $q_{S,t+1}(\cdot | \mathbf{x}_S^{(t)}) = \mathcal{N}(\cdot; \sqrt{1 - \beta^{(t+1)}} \mathbf{x}_S^{(t)}, \beta^{(t+1)})$  as in Section 2.2, where  $\beta^{(t+1)}$  and  $\bar{\alpha}^{(t)}$  are between 0 and 1. Then

$$\text{EKL} \left[ q_{S,t}(\cdot | \mathbf{x}_S^{(t+1)}, \mathbf{x}_M^{(0)}) \| q_{S,t}(\cdot | \mathbf{x}_S^{(t+1)}) \right] \geq -\frac{1}{2} \left( \log(1 - \beta^{(t+1)} \bar{\alpha}^{(t)} \rho^2) + \beta^{(t+1)} \bar{\alpha}^{(t)} \rho^2 \right).$$

The proof of the corollary relies on a lemma on the variances of the two relevant conditional distributions. For notational simplicity, we drop the scripts and annotations on  $\bar{\alpha}^{(t)}$  and  $\beta^{(t+1)}$ , and instead write  $\alpha$  and  $\beta$ , respectively.

**Lemma 4.1.1.** *Suppose  $\mathbf{x}_M^{(0)}, \mathbf{x}_S^{(t)}$ , and  $\mathbf{x}_S^{(t+1)}$  are distributed as in Corollary 4.1.1. Then  $\text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}] = \beta$  and  $\text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}, \mathbf{x}_M^{(0)}] \leq \beta(1 - \beta\rho^2\alpha)$ .*

*Proof.* That  $\text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}] = \beta$  follows immediately from that  $[\mathbf{x}_S^{(t)}, \mathbf{x}_S^{(t+1)}]$  is marginally bivariate normal distributed with covariance  $\sqrt{1 - \beta}$ .

The upper bound on  $\text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}, \mathbf{x}_M^{(0)}]$  is trickier. Observe that  $[\mathbf{x}_S^{(t)}, \mathbf{x}_S^{(t+1)}] | \mathbf{x}_M^{(0)}$  is bivariate Gaussian and that

$$\text{Var} \left[ \begin{bmatrix} \mathbf{x}_S^{(t)} \\ \mathbf{x}_S^{(t+1)} \end{bmatrix} \middle| \mathbf{x}_M^{(0)} \right] = \begin{bmatrix} 1 - \rho^2\alpha & \sqrt{1 - \beta}(1 - \rho^2\alpha) \\ \sqrt{1 - \beta}(1 - \rho^2\alpha) & 1 + \beta\rho^2\alpha - \rho^2\alpha \end{bmatrix}.$$

As such, the conditional variance may be computed in closed form as  $\text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}, \mathbf{x}_M^{(0)}] = \beta(1 - \rho^2\alpha) + (1 - \beta)(1 - \rho^2\alpha) \left( 1 - \frac{1 - \rho^2\alpha}{1 - \rho^2\alpha + \beta\rho^2\alpha} \right)$ . But since  $(1 - \rho^2\alpha)/(1 - \rho^2\alpha + \beta\rho^2\alpha) \geq 1 - (\beta\rho^2\alpha)/(1 - \rho^2\alpha)$  and therefore  $1 - (1 - \rho^2\alpha)/(1 - \rho^2\alpha + \beta\rho^2\alpha) \leq (\beta\rho^2\alpha)/(1 - \rho^2\alpha)$  we can write

$$\begin{aligned} \text{Var}[\mathbf{x}_S^{(t)} | \mathbf{x}_S^{(t+1)}, \mathbf{x}_M^{(0)}] &= \beta(1 - \rho^2\alpha) + (1 - \beta)(1 - \rho^2\alpha) \left( 1 - \frac{1 - \rho^2\alpha}{1 - \rho^2\alpha + \beta\rho^2\alpha} \right) \\ &\leq \beta(1 - \rho^2\alpha) + (1 - \beta)(1 - \rho^2\alpha) \frac{\beta\rho^2\alpha}{1 - \rho^2\alpha} \\ &= \beta(1 - \rho^2\alpha) + (1 - \beta)\beta\rho^2\alpha \\ &= \beta(1 - \beta\rho^2\alpha). \end{aligned}$$

□

Now we provide a proof of Corollary 4.1.1.

*Proof.* First recall that

$$\begin{aligned} \text{KL} [\mathcal{N}(\mu_1, \sigma_1^2) \|\mathcal{N}(\mu_2, \sigma_2^2)] &= \frac{1}{2} \left( \log \frac{\sigma_2^2}{\sigma_1^2} + \frac{\sigma_1^2 + (\mu_1 - \mu_2)^2}{\sigma_2^2} - 1 \right) \\ &\geq \frac{1}{2} \left( \log \frac{\sigma_2^2}{\sigma_1^2} + \frac{\sigma_1^2}{\sigma_2^2} - 1 \right) \end{aligned}$$

and observe that this lower bound is monotonically decreasing in  $\sigma_1^2$  for  $\sigma_1^2 \leq \sigma_2^2$ .

Therefore

$$\begin{aligned} &\text{EKL} \left[ q_{\mathcal{S},t}(\cdot \mid \mathbf{x}_{\mathcal{S}}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)}) \|\ q_{\mathcal{S},t}(\cdot \mid \mathbf{x}_{\mathcal{S}}^{(t+1)}) \right] \\ &= \int q_{\mathcal{M},0}(x_{\mathcal{M}}^{(0)}) q_{\mathcal{S},t+1}(x_{\mathcal{S}}^{(t+1)} \mid x_{\mathcal{M}}^{(0)}) \left[ \right. \\ &\quad \left. \text{KL} \left[ q_{\mathcal{S},t}(\cdot \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}^{(0)}) \|\ q_{\mathcal{S},t}(\cdot \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}) \right] \right] \\ &\quad \left. dx_{\mathcal{M}}^{(0)} dx_{\mathcal{S}}^{(t+1)} \right] \\ &\geq \int q_{\mathcal{M},0}(x_{\mathcal{M}}^{(0)}) q_{\mathcal{S},t+1}(x_{\mathcal{S}}^{(t+1)} \mid x_{\mathcal{M}}^{(0)}) \left[ \right. \\ &\quad \left. \text{KL} \left[ \mathcal{N}(0, \text{Var}[\mathbf{x}_{\mathcal{S}}^{(t)} \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}, \mathbf{x}_{\mathcal{M}}^{(0)} = x_{\mathcal{M}}^{(0)}]) \|\ \mathcal{N}(0, \text{Var}[\mathbf{x}_{\mathcal{S}}^{(t)} \mid \mathbf{x}_{\mathcal{S}}^{(t+1)} = x_{\mathcal{S}}^{(t+1)}]) \right] \right] \\ &\quad \left. dx_{\mathcal{M}}^{(0)} dx_{\mathcal{S}}^{(t+1)} \right] \\ &\geq \text{KL} [\mathcal{N}(0, \beta(1 - \beta\rho^2\alpha)) \|\ \mathcal{N}(0, \beta)] \\ &\geq \frac{1}{2} \left( \log \frac{\beta}{\beta(1 - \beta\rho^2\alpha)} + \frac{\beta(1 - \beta\rho^2\alpha)}{\beta} - 1 \right) \\ &= -\frac{1}{2} (\log(1 - \beta\rho^2\alpha) + \beta\rho^2\alpha) \end{aligned}$$

where the second inequality follows from Lemma 4.1.1, and the monotonicity of the KL in  $\sigma_1^2$ .  $\square$

We note two takeaways of Corollary 4.1.1. First, as we might intuitively expect, this error can be large when significant correlation in the target distribution is present. Second, we see that the approximation error can be larger at earlier time steps, when  $\bar{\alpha}^{(t)}$  is closer to 1.

## 4.2 Conditional sampling is a sequential Monte Carlo problem

We next frame approximation of  $q(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)})$  as a sequential Monte Carlo problem that we may solve by particle filtering. Intuitively, particle filtering addresses a limitation of the replacement method: the failure at each time  $t$  to look beyond the current step to the less-noised motif  $\mathbf{x}_M^{(t-1)}$  when sampling  $\mathbf{x}_S^{(t)} \sim p_\theta(\mathbf{x}_S^{(t)} | \mathbf{x}^{(t+1)})$ . Our key insight is that because  $p_\theta(\mathbf{x}_M^{(t-1)} | \mathbf{x}^{(t)})$  provides a mechanism to assess the likelihood of  $\mathbf{x}_M^{(t-1)}$ , we can prioritize noised scaffolds that are more consistent with the motif. Particle filtering leverages this mechanism to provide a sequence of discrete approximations to each  $p_\theta(\mathbf{x}_S^{(t)} | \mathbf{x}_M^{(t-1:T)})$  that look ahead by this extra step. Finally, at  $t = 0$  we have an approximation to  $p_\theta(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0:T)})$ . Then, using Proposition 4.2.1 below, we can obtain an approximate sample from  $q(\mathbf{x}_S^{(0)} | \mathbf{x}_M^{(0)})$ . This framing permits the application of standard particle filtering algorithms [Doucet et al., 2001]. Algorithm 2 summarizes an implementation of this procedure that uses residual resampling [Doucet and Johansen, 2009] to mitigate the collapse of the sequential approximations into point masses.

---

**Algorithm 2** SMCDiff: Particle filtering for conditionally sampling from unconditional diffusion models

---

```

1: Input:  $\mathbf{x}_{\mathcal{M}}^{(0)}$  (motif),  $K$  (# particles)
2: // Forward diffuse motif
3:  $\check{\mathbf{x}}_{\mathcal{M}}^{(1:T)} \sim q(\mathbf{x}_{\mathcal{M}}^{(1:T)} | \mathbf{x}_{\mathcal{M}}^{(0)})$ 
4:
5: // Reverse diffuse particles
6:  $\forall k, \mathbf{x}_k^{(T)} \overset{i.i.d.}{\sim} p_{\theta}(\mathbf{x}^{(T)})$ 
7: for  $t = T, \dots, 1$  do
8:   // Replace motif
9:    $\forall k, \mathbf{x}_k^{(t)} \leftarrow [\check{\mathbf{x}}_{\mathcal{M}}^{(t)}, \mathbf{x}_{S,k}^{(t)}]$ 
10:
11:   // Re-weight based on  $\check{\mathbf{x}}_{\mathcal{M}}^{(t-1)}$ 
12:    $\forall k, w_k^{(t)} \leftarrow p_{\theta}(\check{\mathbf{x}}_{\mathcal{M}}^{(t-1)} | \mathbf{x}_k^{(t)})$ 
13:    $\forall k, \tilde{w}_k^{(t)} \leftarrow w_k^{(t)} / \sum_{k'=1}^K w_{k'}^{(t)}$ 
14:    $\tilde{\mathbf{x}}_{1:K}^{(t)} \sim \text{Resample}(\tilde{w}_{1:K}^{(t)}, \mathbf{x}_{1:K}^{(t)})$ 
15:
16:   // Propose next step
17:    $\forall k, \mathbf{x}_k^{(t-1)} \overset{indep.}{\sim} p_{\theta}(\mathbf{x}^{(t-1)} | \tilde{\mathbf{x}}_k^{(t)})$ 
18: end for
19: Return  $\mathbf{x}_{S,1:K}^{(0)}$ 

```

---

SMCDiff provides a tunable trade-off between computational cost and statistical accuracy through the choice of the number of particles  $K$ . In our next proposition we make this trade-off explicit.

**Proposition 4.2.1.** *Suppose that  $p_{\theta}$  exactly matches the forward diffusion process such that for every  $\mathbf{x}^{(t+1)}$ ,  $p_{\theta}(\mathbf{x}^{(t)} | \mathbf{x}^{(t+1)}) = q(\mathbf{x}^{(t)} | \mathbf{x}^{(t+1)})$  and consider any motif  $\mathbf{x}_{\mathcal{M}}^{(0)}$ . Let  $\mathbf{x}_{S,K}$  be a particle chosen at random from the output of Algorithm 2 with  $K$  particles. Then  $\mathbf{x}_{S,K}$  converges in distribution to  $q(\mathbf{x}_S^{(0)} | \mathbf{x}_{\mathcal{M}}^{(0)})$  as  $K$  goes to infinity.*

The significance of Proposition 4.2.1 is that it guarantees Algorithm 2 can provide arbitrarily accurate conditional samples provided an accurate diffusion model and large enough compute budget (determined by the number of particles). To our knowledge, SMCDiff is the first algorithm for asymptotically exact conditionally sampling from unconditional DPMs.

Our proof of the proposition is obtained from an application of standard asymptotics for particle filtering [Chopin and Papaspiliopoulos, 2020, Proposition 11.4]. The proof will proceed by first discussing technical details of Algorithm 2 then reducing the proof down to proving a lemma.

The idea behind the `SMCDiff` procedure in Algorithm 2 is to break sampling of  $\mathbf{x}_S^{(0)} \sim q_{S,0}(\cdot \mid \mathbf{x}_M^{(0)})$  into three stages:

1. Draw  $\mathbf{x}_M^{(1:T)} \sim q_{M,1:T}(\cdot \mid \mathbf{x}_M^{(0)})$ .
2. Draw  $\mathbf{x}_S^{(1:T)} \sim q_{S,1:T}(\cdot \mid \mathbf{x}_M^{(0:T)})$ .
3. Draw  $\mathbf{x}_S^{(0)} \sim q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:T)}, \mathbf{x}_S^{(1:T)})$

If all three steps were performed exactly, by the law of total probability  $\mathbf{x}_S^{(0)}$  in step (3) would (marginally) be an exact sample from  $q_{S,0}(\cdot \mid \mathbf{x}_M^{(0)})$ . As such, `SMCDiff` aims to perform step (1) and approximate steps (2) and (3). Step (1) corresponds to forward diffusing the motif in lines 2–3 and is exact because we diffuse according to  $q$ .

Step (3) corresponds to line 17 in the last iteration (when  $t = 1$ ). Specifically, to sample from  $q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:T)}, \mathbf{x}_S^{(1:T)})$  we make three observations. (i) The Markov structure of the forward process implies that  $q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:T)}, \mathbf{x}_S^{(1:T)}) = q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:1)}, \mathbf{x}_S^{(1)})$ . (ii) By the assumption that the forward and approximated reverse process agree, we have  $q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:1)}, \mathbf{x}_S^{(1)}) = p_{S,0}(\cdot \mid \mathbf{x}_M^{(0:1)}, \mathbf{x}_S^{(1)})$ . (iii) Finally, because  $p_t(\cdot \mid \mathbf{x}^{(t+1)})$  factorizes across  $\mathcal{M}$  and  $\mathcal{S}$  for each  $t$ ,  $p_{S,0}(\cdot \mid \mathbf{x}_M^{(0:1)}, \mathbf{x}_S^{(1)}) = p_{S,0}(\cdot \mid \mathbf{x}_M^{(1)}, \mathbf{x}_S^{(1)})$ . As a result, under the assumptions of the proposition, we may sample from  $q_{S,0}(\cdot \mid \mathbf{x}_M^{(0:T)}, \mathbf{x}_S^{(1:T)})$ , and perform step (3) exactly as well.

Step (2) is the only non-trivial step, and cannot be performed exactly. The challenge is that although the reverse process approximation,  $p_{S,1:T}(\cdot \mid \mathbf{x}_M^{(0:T)})$ , is well-defined, computing it explicitly involves an intractable, high-dimensional integral.

The sequential Monte Carlo approach of `SMCDiff`, then, is to circumvent this intractability by constructing a sequence of approximations. For each  $t = T, T-1, \dots, 1$ , we approximate  $p_{S,t}(\cdot \mid \mathbf{x}_M^{(t-1:T)})$  (and thereby  $q_{S,t}(\cdot \mid \mathbf{x}^{(t-1:T)})$ ) with  $K$  weighted atoms (the *particles*). We denote these approximations (which are implicit in Algorithm 2)

by  $\mathbb{P}_K^{(t)}(\cdot) := \sum_{k=1}^K \tilde{w}_k^{(t)} \delta(\cdot; \mathbf{x}_{S,k}^{(t)})$ , where each  $\tilde{w}_k^{(t)}$  and  $\mathbf{x}_{S,k}^{(t)}$  are as in Algorithm 2, and  $\delta(\cdot; \mathbf{x})$  denotes a Dirac mass at  $\mathbf{x}$ . In particular,  $\mathbb{P}_K^{(1)}(\cdot)$  is an approximation to  $p_{S,1}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$ .

Proving the proposition amounts to showing that in the limit as  $K$  goes to infinity, each  $\mathbb{P}_K^{(1)}(\cdot)$  converges weakly to  $p_{S,1}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$ , which by assumption is equal to  $q_{S,1}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$ . This weak convergence follows from standard asymptotics for particle filters [Chopin and Papaspiliopoulos, 2020, Proposition 11.4], which we make explicit in Lemma 4.2.1. As a result, if we perform step (3) with  $\mathbf{x}_S^{(1)} \sim \mathbb{P}_K^{(1)}(\cdot)$ , then this lemma implies that  $\mathbf{x}_S^{(0)}$  converges in distribution to  $q_{S,0}(\mathbf{x}_S^{(0)} \mid \mathbf{x}_{\mathcal{M}}^{(0)})$ , since (i)  $q_{S,0}(\mathbf{x}_S^{(0)} \mid \mathbf{x}_{\mathcal{M}}^{(1)}, \mathbf{x}_S^{(1)})$  is continuous in  $\mathbf{x}_S^{(1)}$  and (ii)  $\mathbf{x}_S^{(0)}$  is independent of  $\mathbf{x}_{\mathcal{M}}^{(0)}$  conditional on  $\mathbf{x}^{(1)}$ .

Recall that to show the proposition, it was sufficient to show that  $\mathbb{P}_K^{(1)}$  converged weakly to  $q_{S,1}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$ ; this implied that the  $K$  particle returned by Algorithm 2 would then converge in distribution to  $q_{S,0}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$  which, by the law of total probability, implied that they marginally converge to  $q_{S,0}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0)})$ . However, while the particles return by Algorithm 2 may be treated as exchangeable, they are not independent, because they depend on shared randomness in  $\mathbf{x}_{\mathcal{M}}^{(1:T)}$ . To obtain approximate samples that are independent, it is necessary to run Algorithm 2 multiple times.

**Residual resampling.** Line 14 of Algorithm 2 indicates a `Resample` step. In particle filtering, resampling steps (or *branching mechanisms* [Doucet et al., 2001, Chapter 2]) filter out particles with very small weights, and replace them with additional copies of particles with large weights. Notably, the resampling step is the only point of departure of Algorithm 2 from the replacement method; without resampling, the algorithms behave identically. While a variety of possible branching mechanisms exist, we use *residual resampling* (Algorithm 3) in our implementation for its simplicity.



---

**Algorithm 3** Residual Resample
 

---

- 1: **Input:**  $w_{1:K}$  (weights),  $\mathbf{x}_{1:K}$  (particles)
  - 2:  $\forall k, (c_k, r_k) \leftarrow (\lfloor Kw_k \rfloor, Kw_k - \lfloor Kw_k \rfloor)$
  - 3:  $\tilde{\mathbf{x}}_C = \underbrace{[\mathbf{x}_1, \dots, \mathbf{x}_1]}_{c_1}, \dots, \underbrace{[\mathbf{x}_K, \dots, \mathbf{x}_K]}_{c_K}$
  - 4:  $R \leftarrow K - \sum_{k=1}^K c_k$
  - 5:  $[i_1, \dots, i_R] \sim \text{Multinomial}(r_{1:K}, R)$
  - 6:  $\tilde{\mathbf{x}}_R \leftarrow [\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_R}]$
  - 7:  $\tilde{\mathbf{x}} = \text{concat}(\mathbf{x}_R, \mathbf{x}_C)$
  - 8: Return  $\tilde{\mathbf{x}}$
- 

**Lemma 4.2.1.** Consider  $\mathbb{P}_K^{(1)} := \sum_{k=1}^K \tilde{w}_k \delta(\cdot; \mathbf{x}_{S,k}^{(1)})$ , where  $\tilde{w}_k$  and  $\mathbf{x}_{S,1:K}^{(1)}$  are as constructed in Algorithm 2. Assume the conditions of Proposition 4.2.1. Then  $\mathbb{P}_K^{(1)}$  converges weakly to  $p_{S,1}(\cdot \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$  as  $K$  goes to infinity. That is, for any Borel measurable  $A$ ,  $\lim_{K \rightarrow \infty} \mathbb{P}_K^{(1)}(A) = \int_A p_{S,1}(x \mid \mathbf{x}_{\mathcal{M}}^{(0:T)}) dx$ .

*Proof.* The proof of the lemma follows from an application of standard asymptotics for particle filtering [Chopin and Papaspiliopoulos, 2020, Proposition 11.4]. In particular, to apply Proposition 11.4 we use the formalism of Feynman–Kac (FK) models, following the notation of [Chopin and Papaspiliopoulos, 2020, Chapter 5]. Though typically (and in [Chopin and Papaspiliopoulos, 2020]) FK models are defined via a sequence of approximations at increasing time steps, we consider decreasing time steps because we are approximating the reverse time process. We take the initial distribution as  $\mathbb{M}_T(\mathbf{x}_S^{(T)}) = p_{S,T}(\mathbf{x}_S^{(T)})$ , the transition kernel as  $M_t(\mathbf{x}_S^{(t+1)}, \mathbf{x}_S^{(t)}) = p_{S,t}(\mathbf{x}_S^{(t)} \mid \mathbf{x}^{(t+1)})$ , and the potential functions as  $G_t(\mathbf{x}_S^{(t)}) = p_{\mathcal{M},t-1}(\mathbf{x}_{\mathcal{M}}^{(t-1)} \mid \mathbf{x}^{(t)})$ . The sequence of FK models,  $\mathbb{Q}_t$ , then correspond to

$$\mathbb{Q}_t(\mathbf{x}_S^{(t:T)}) = L_t^{-1} \mathbb{M}_T(\mathbf{x}_S^{(T)}) G_T(\mathbf{x}_S^{(T)}) \prod_{i=T-1}^t M_i(\mathbf{x}_S^{(i+1)}, \mathbf{x}_S^{(i)}) G_i(\mathbf{x}_S^{(i)})$$

for each  $t$ , where  $L_t$  is a normalizing constant.

By substituting in our choices of  $M_t$  and  $G_t$ , we can rewrite and simplify  $\mathbb{Q}_t$  as

$$\begin{aligned}
\mathbb{Q}_t(\mathbf{x}_S^{(t:T)}) &= L_t^{-1} p_{S,T}(\mathbf{x}_S^{(T)}) p_{\mathcal{M},T-1}(\mathbf{x}_{\mathcal{M}}^{(T-1)} \mid \mathbf{x}^{(T)}) \prod_{i=T-1}^t p_{S,i}(\mathbf{x}_S^{(i)} \mid \mathbf{x}^{(i+1)}) p_{\mathcal{M},i-1}(\mathbf{x}_{\mathcal{M}}^{(i-1)} \mid \mathbf{x}^{(i)}) \\
&= L_t^{-1} p_{S,T}(\mathbf{x}_S^{(T)}) p_{t:T-1}(\mathbf{x}^{(t:T-1)} \mid \mathbf{x}^{(T)}) p_{\mathcal{M},t-1}(\mathbf{x}_{\mathcal{M}}^{(t-1)} \mid \mathbf{x}^{(t)}) \\
&\propto p_{t:T}(\mathbf{x}^{(t:T)} \mid \mathbf{x}_{\mathcal{M}}^{(t-1)}) \\
&\propto p_{S,t:T}(\mathbf{x}_S^{(t:T)} \mid \mathbf{x}_{\mathcal{M}}^{(t-1:T)}),
\end{aligned}$$

where lines 3 and 4 drop multiplicative constants that do not depend on  $\mathbf{x}_S^{(t:T)}$ . From the above derivation, we see that each  $\mathbb{Q}_t(\mathbf{x}_S^{(t)}) = p_{S,t}(\mathbf{x}_S^{(t)} \mid \mathbf{x}_{\mathcal{M}}^{(t-1:T)})$ , and in particular that  $\mathbb{Q}_1(\mathbf{x}_S^{(1)}) = p_{S,1}(\mathbf{x}_S^{(1)} \mid \mathbf{x}_{\mathcal{M}}^{(0:T)})$ . As such, the desired convergence in the statement of the lemma is equivalent to that  $\mathbb{P}_K^{(1)}$  converges to  $\mathbb{Q}_1$ .

[Chopin and Papaspiliopoulos \[2020, Proposition 11.4\]](#) provide this result for the generic particle filtering algorithm (see [Chopin and Papaspiliopoulos \[2020, Algorithm 10.1\]](#)], which is written in the FK model form described above). More specifically, Proposition 11.4 proves almost sure convergence of all Borel measurable functions of  $\mathbb{P}_K^{(t)}$ , which implies the desired weak convergence.

Although the proof provided in [Chopin and Papaspiliopoulos \[2020\]](#) is restricted to the simpler, but higher variance, case where the resampling step uses multinomial resampling, the authors note that [Chopin \[2004\]](#) proves it holds in the case of residual resampling (which we use in our experiments) as well.  $\square$

# Chapter 5

## Experiments

We empirically demonstrate the ability of our method to scaffold motifs and sample protein backbone structures. We train a single instance of `ProtDiff` and use it across all of our experiments. For simplicity, we limited our training data to single chain proteins taken from PDB that are no longer than 128 residues. We describe training details in Section 5.1. We describe our procedure for evaluating backbone designs in Section 5.2. We demonstrate the promise of our method for the motif-scaffolding problem in Section 5.3. And we investigate our method’s strengths and weaknesses via experiments in unconditional sampling in Section 5.4.4. Finally, we demonstrate the applicability of `SMCDiff` outside of proteins by applying it to MNIST inpainting in Section 5.5.

**Baselines.** As mentioned in Section 1.3, Wang et al. [2022] is the only prior machine learning work to address the motif-scaffolding problem. We do not compare against this as a baseline because no stable implementation was available at the time of writing. The most closely related method for unconditional sampling with available software is `trDesign` [Anishchenko et al., 2021], but this method does not allow specification of a motif. The ML method most similar to `ProtDiff` is the concurrently developed equivariant diffusion model (EDM) proposed by Hoogeboom et al. [2022]. Like `ProtDiff`, EDM uses a densely connected EGNN architecture but without sequence-distance edge features. Consequently, it does not impose any sequence

order, and therefore does not yield a way to relate generated coordinates to a backbone chain.

## 5.1 ProtDiff training details

ProtDiff uses 4 equivariant graph convolutional layers (EGCL) with 256 dimensions for node and edge embeddings. The training data was restricted to single chain proteins (monomers) found in PDB and lengths in the range [40, 128]. We additionally filtered out PDB with  $>5\text{\AA}$  atomic resolution. This amounted to 4269 training examples. Training was performed using the Adam optimizer with hyperparameters `learning_rate=1e-4`,  $\beta_1 = 0.9$ , and  $\beta_2 = 0.999$ . We trained for 1,000,000 steps using batch size 16. We used a single Nvidia A100 GPU for approximately 24 hours. We implemented all models in PyTorch. We used the same linear noise schedule as Ho et al. [2020] where  $\beta_0 = 0.0001$ ,  $\beta_T = 0.02$ , and  $T = 1024$ . We did not perform hyperparameter tuning.

## 5.2 *In silico* evaluation of designed backbones

While experimental validation via X-ray crystallography remains the gold standard for evaluating computationally designed proteins, recent work [Wang et al., 2022, Lin et al., 2021] has proposed to leverage highly accurate protein structure prediction neural networks as an *in silico* proxy for true structure. More specifically, Wang et al. [2022] jointly design protein sequence and structure, and validate by comparing the design and AlphaFold2 (AF2) [Jumper et al., 2021] predicted structures. Here, our goal is to assess the quality of scaffolds generated independent of a specific sequence, so we treat fixed backbone sequence design as a downstream step as in Lin et al. [2021].

Our evaluation with AF2 is as follows. For each generated scaffold we use a C- $\alpha$  only version of ProteinMPNN [Dauparas et al., 2022] with a temperature of 0.1 to sample 8 amino acid sequences likely to fold to the same backbone structure. We

then run AF2 with the released CASP14<sup>1</sup> weights and 15 recycling iterations. We do not include a multiple sequence alignment as an input to AF2. Our choice of utilizing ProteinMPNN and AF2 (without MSAs) is motivated by their empirical success in various de novo protein design tasks and the ability to recapitulate native proteins [Dauparas et al., 2022, Bennett et al., 2022]. To assess unconditionally sampled scaffolds, we then evaluate the agreement of our backbone sample with the AF2 predicted structures using the maximum TM-score [Zhang and Skolnick, 2005] across all generated sequences which we refer to as **scTM**, for *self-consistency* TM-score. To assess whether prospective scaffolds generated support a motif, we compute the root mean squared distances (RMSD) of the desired and predicted motif coordinates after alignment and refer this metric as the *motif RMSD*. Algorithm 4 outlines the exact steps.

Because a TM-score  $> 0.5$  indicates that two structures have the same fold [Zhang and Skolnick, 2005], we say that a backbone is designable if **scTM**  $> 0.5$ . The ability for AF2 to reproduce the same backbone from an independently designed sequence is evidence a sequence can be found for the starting structure. To verify this is a reasonable cutoff, we analyzed **scTM** over our training set and found 87% to be designable.

---

**Algorithm 4** Self-consistency calculation

---

**Input:**  $\mathbf{x} \in \mathbb{R}^{N,3}$

- 1: **for**  $i \in 1, \dots, 8$  **do**
- 2:      $s_i \leftarrow \text{ProteinMPNN}(\mathbf{x})$
- 3:      $\hat{\mathbf{x}}_i \leftarrow \text{AF2}(s_i)$
- 4: **end for**
- 5: **sc\_tm**  $\leftarrow \max_{i \in 1, \dots, 8} \text{TMscore}(\hat{\mathbf{x}}_i, \mathbf{x})$

**Output:** , **sc\_tm**

---



---

<sup>1</sup>Biannual protein folding competition where AF2 achieved first place. Weights available under Apache License 2.0 license.

### 5.3 Motif-scaffolding via conditional sampling

We evaluated our motif-scaffolding approach (combining `SMCDiff` and `ProtDiff`) on motifs extracted from existing proteins in the PDB and found that our approach can generate long and diverse scaffolds that support these motifs.

We chose to first evaluate on motifs extracted from proteins present in the training set because we knew that at least one stabilizing scaffold exists. We considered 2 examples taken from the PDB with IDs `6exz` and `5trv`, which are 69 and 118 residues long, respectively. We chose these examples due to their high secondary structure composition while being representative of the shortest and longest lengths seen during training. For each structure, we chose a 15–25 residue helical segment as the motif (see Figure 5-1 and Table 5.1). The remainder of each protein is one possible supporting scaffold. We sought to assess if we could recover this and other scaffolds with the same size and motif placement.




Figure 5-1: Structures used for motif-scaffolding test cases. Native structures (grey) and their motifs (orange) that were used for the motif-scaffolding task are shown.

Table 5.1: Motif-scaffolding test case details.

| Origin/ Protein        | Total length | Motif size (residue range) |
|------------------------|--------------|----------------------------|
| 6exz                   | 72           | 15 (30–44)                 |
| 5trv                   | 118          | 21 (42–62)                 |
| RSV (PDB-ID: 5tpn)     | 62           | 19 (16–34)                 |
| EF-hand (PDB-ID: 1PRW) | 53           | 5 (0–4), 13 (31–43)        |

Based on prior work [Wang et al., 2022], we expected that building larger scaffolds around a motif would be more challenging than building smaller scaffolds. To assess this length dependence, we expanded the segment of used as the motif when running `SMCDiff` by including additional residues on each side. In each case, though, we compute the motif RMSD over the minimal motif. In Figure 5-2B, we present motif-scaffolding performance and its dependence on scaffold size for 5trv, the longer of the two test proteins. For the 5trv test case, the lower quartile of the motif RMSD for `SMCDiff` is below 1Å for scaffolds up to 80 residues. Since 1Å is atomic-level resolution, we conclude that our approach can succeed in this length range.

Figure 5-2A provides a visualization of our method’s capacity to generate long and diverse scaffolds. The figure depicts two dissimilar scaffolds of lengths 34 and 54 produced by `SMCDiff` with 64 particles. Both scaffolds are designable and agree with AF2 ( $s_{\text{cTM}} > 0.5$ ). Diversity is particularly evident in the different orderings of secondary structures.

Figure 5-2B compares `SMCDiff` to two naive inpainting methods, `fixed` and `replacement`. In `fixed`, the motif is fixed for every timestep  $t$ , and the reverse diffusion is applied only to the scaffold (as done by Zhou et al. [2021]); `replacement` is the method described in Section 4.1. In contrast to `SMCDiff`, these baselines fail to generate a successful scaffolds longer than 50 residues on 5trv, as determined by the location of their lower quartiles.

We next applied these three inpainting methods to harder targets in order to measure generalization to out-of-distribution and more difficult motifs comprising of dis-contiguous regions and loops. We consider a motif obtained from the respiratory

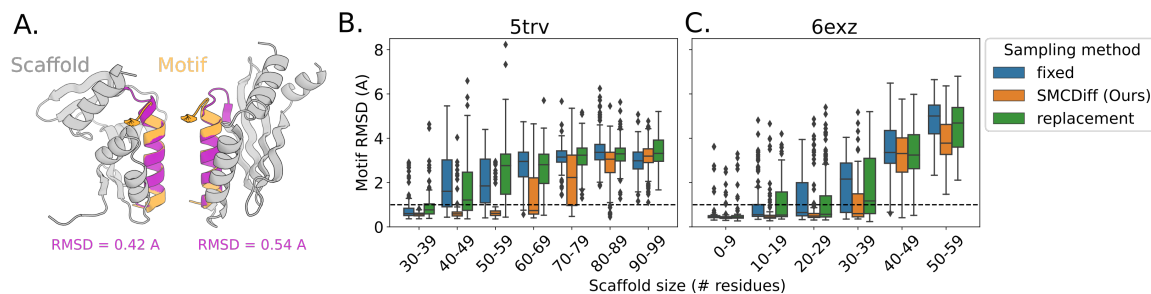


Figure 5-2: Motif-scaffolding case studies. (A) Example of two scaffold structures generated around a segment of 5trv. Orange: desired input motif, Grey: AlphaFold-predicted structure of two scaffolds, with the motif highlighted (purple). Both scaffolds were sampled using SMCDiff with  $\text{sCTM} > 0.5$ . (B,C) Motif RMSD for 5trv and 6exz test cases, its dependence on scaffold size, and comparison of SMCDiff to two naive inpainting methods (fixed, replacement).

syncytial virus (RSV) protein and calcium binding EF-hand motif, both of which are not in the training dataset. RSV is known to be difficult due to its composition of helical, loop, and sheet segments, while EF-hand is a dis-contiguous loop motif found in a calcium binding protein. More details about both motifs can be found in Wang et al. [2022]; there the authors report the only known successful scaffold of these motifs but they attain it with a computationally intensive *hallucination* approach. We found that our method failed to generate scaffolds predicted to recapitulate the motif (Figure 5-3); however, SMCDiff provided smaller median motif RMSDs than the other two inpainting methods.

**Compute cost.** The computation of SMCDiff with 64 particles is approximately 2 minutes per independent sample, while alternative methods fixed and replacement can produce 64 independent samples in the same time. By contrast, the hallucination approach of Wang et al. [2022] involves running a Markov chain for thousands of steps, and has runtime on the order of hours for a single sample [Anishchenko et al., 2021].



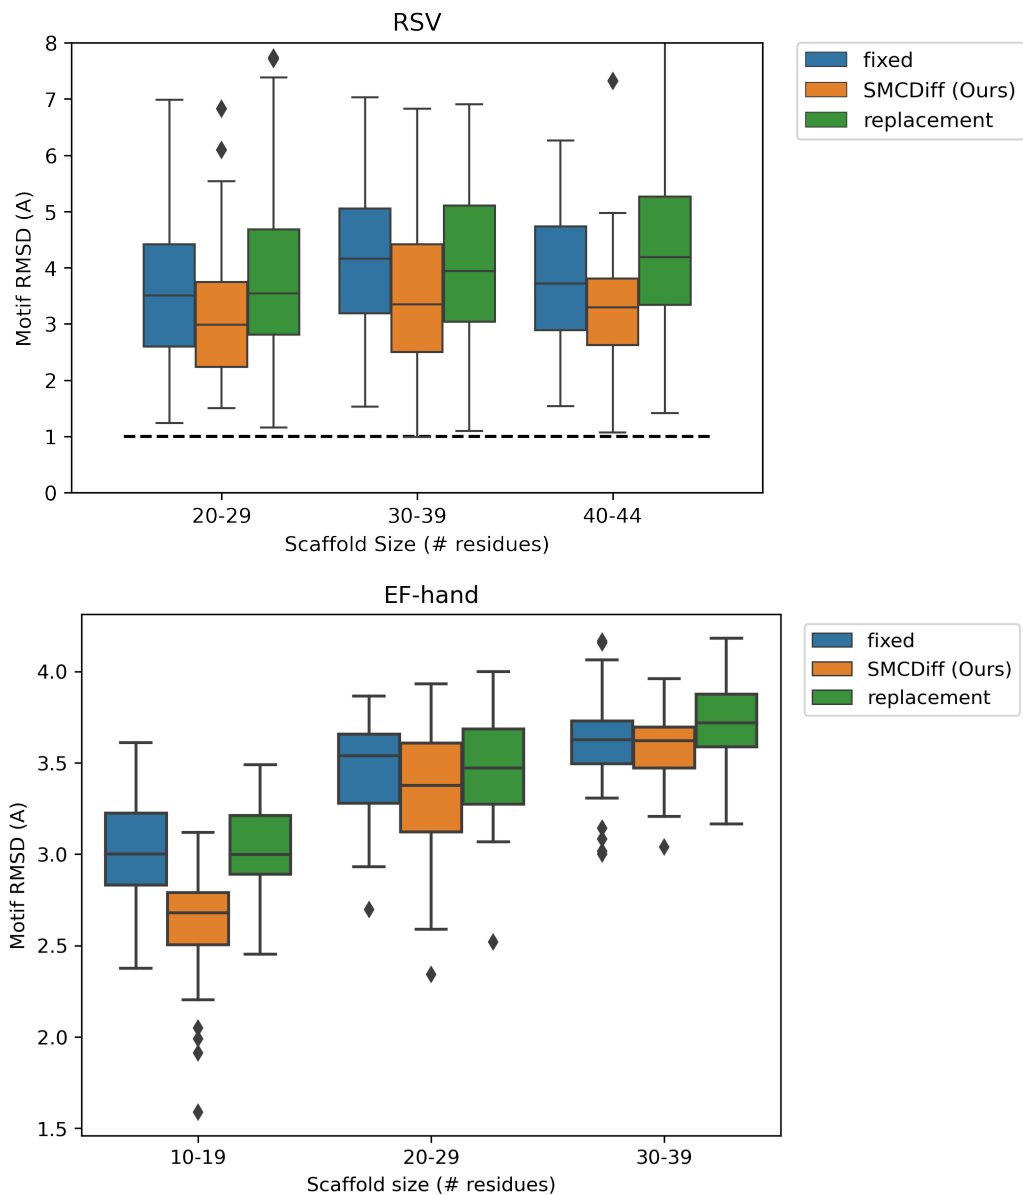


Figure 5-3: Additional inpainting results on a more challenging motif extracted from the respiratory syncytial virus (RSV) and EF-hand motif. The three inpainting methods are evaluated as described in Section 5.3.

## 5.4 Unconditional sampling

We next investigate the origins of the diversity seen in Figure 5-2 by analyzing the diversity and designability of ProtDiff samples without conditioning on a motif.

We first check that ProtDiff produces designable backbones. To do this, we generated 10 backbone samples for each length between 50 and 128 and then calculated

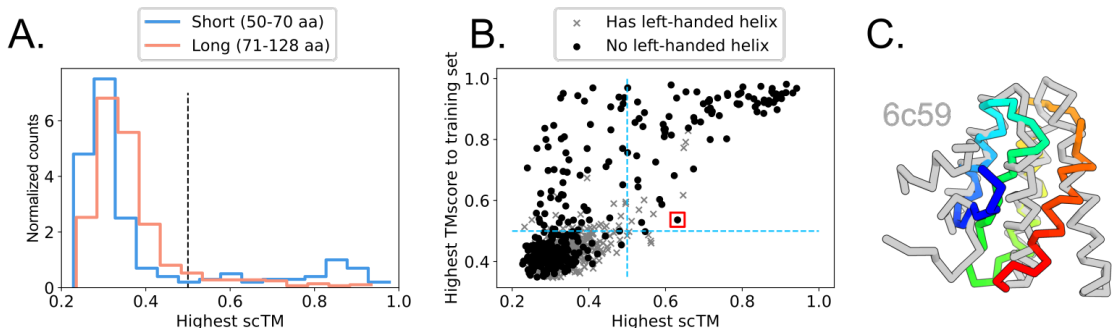


Figure 5-4: Protein backbone samples from ProtDiff. (A) Density plot of  $\text{scTM}$  for different length categories (50–70, 70–128). The dashed line at  $\text{scTM} = 0.5$  indicates the threshold of “designability”, points to the right are considered “designable” (see text). (B) Scatter plot of  $\text{scTM}$  and the highest TM-score of each sample to all of PDB. Points represented as a grey “ $\times$ ” are detected to contain an (invalid) left-handed helix. Dashed lines indicate thresholds  $\text{scTM} = 0.5$ . (C) Example of a designable backbone sample (rainbow) with  $\text{scTM} > 0.5$  (boxed in red in panel B) to its closest PDB example (6c59, grey) with a TM-score of 0.54.

$\text{scTM}$  for each sample. In Fig. 5-4A, we find that 11.8% of samples have  $\text{scTM} > 0.5$ . However, the majority of backbones do not pass this threshold. We also observe designability has strong dependence on length since we expect that longer proteins are harder to model in 3D and design sequences for. We separated the lengths below 128 residues into two categories and refer to them as *short* (50–70) and *long* (70–128). Our results in Figure 5-4A indicate 17% of designs in the short category are designable vs. 9% in the long category.

We next sought to evaluate the ability of ProtDiff to generalize beyond the training set and produce novel backbones. In Figure 5-4B each point represents a backbone sample from ProtDiff. The horizontal coordinate of a point is the  $\text{scTM}$ , and the vertical coordinate is the minimum TM-score across the training set. We found a strong positive correlation between  $\text{scTM}$  and this minimum TM-score, indicating that many of the most designable backbones generated by ProtDiff were a result of training set memorization. However, if the model were only memorizing the training set, we would see TM-scores consistently near 1.0; the range of scores in Figure 5-4B indicate this is not the case – and the model is introducing a degree of variability. Figure 5-4C gives an example of backbone with  $\text{scTM} > 0.5$  that appears

to be novel. Its closest match in the PDB has TM-score = 0.54.

### 5.4.1 Sample diversity analysis

To measure diversity of our designed samples, all 92 samples with  $\text{scTM} > 0.5$  were compared and clustered using MaxCluster [Herbert and Sternberg \[2008\]](#). Structures were compared in a sequence independent manner, using the TM-score of the maximal subset of paired residues. They were subsequently clustered using hierarchical clustering with average linkage,  $1 - \text{TM-score}$  as the distance metric and a TM-score threshold of 0.5 (Figure 5-5).




Figure 5-5: Clustering of self-consistent ProtDiff samples. The distance matrix is  $1 - \text{TM-score}$  between pairs of samples, and ranges from 0 (exact match) to 1 (no match). Dendrograms are from hierarchical clustering using the average distance metric. Designs on the right are cluster centroids. Gray lines connect larger clusters with more than one member to its centroid, while the remaining designs are from a random selection of the remaining single-sample clusters. Protein backbones are colored from blue at the N-terminus to red at the C-terminus.

### 5.4.2 Failure mode analysis

Fig. 5-4B illustrates a limitation of our method: many of our sampled backbones are not designable. One contributing factor is that ProtDiff does not handle chirality. Hence ProtDiff generates backbones with the wrong handedness, which cannot be realized by any sequence. Fig. 5-4B shows that 45% of all backbone samples had at least one incorrect, left-handed helix. Of these, most have  $\text{scTM} < 0.5$ . We additionally find chain breaks to be a common failure case of our samples. Visualizations of our failure cases are seen in Figure 5-6.




Figure 5-6: Failure modes in ProtDiff backbone samples. (A) Backbone clashes and chain breaks. The C- $\alpha$  atoms can be spaced further than the typical 3.8Å between neighbors, resulting in a chain break (dashed lines). Additionally, backbone segments can be too close to each other, resulting in obvious overlaps and clashes. (B) Backbones with a mixture of left (circled in red) and right (circled in green) handed helices. These chirality errors cannot be corrected simply by mirroring the sampled backbone.

### 5.4.3 Noise interpolation analysis

Figure 5-7 and Figure 5-8 illustrates an interpolation between two samples, showing how ProtDiff’s outputs change as a function of the noise used to generate them. To generate these interpolations, we pick two backbone samples that result in different folds. For independent samples generated with noise  $\epsilon^{(0:T)}$  and  $\tilde{\epsilon}^{(0:T)}$  we interpolate with noise set to  $\sqrt{\alpha}\epsilon^{(0:T)} + \sqrt{1-\alpha}\tilde{\epsilon}^{(0:T)}$  for  $\alpha$  between 0 and 1. The depicted values

of  $\alpha$  are chosen to highlight transition points. A future direction is to exploit the latent structure of ProtDiff to control backbone topology.




Figure 5-7: Latent interpolation of length 89 backbone sample from  $\alpha = 0$  to 1.




Figure 5-8: Latent interpolation of length 63 backbone sample from  $\alpha = 0$  to 1.

#### 5.4.4 Qualitative analysis of scTM in different ranges

In this section, we give intuition for backbone designs and AF2 predictions associated with different values of scTM to aid the interpretation of the scTM results. Figure 5-9 examines a possible categorization of scTM in three ranges. The first two rows correspond to backbone designs that achieve scTM > 0.9. We see the backbone designs in the first column closely match the AF2 prediction in the second column. A closely related PDB example can be found when doing a similarity search of the highest PDB chain with the highest TM-score to the AF2 prediction. We showed in Figure 5-4B that scTM > 0.9 is indicative of a close structural match being found in PDB.

The middle two rows correspond to designs that achieve scTM  $\sim$  0.5. These are examples of backbone designs on the edge of what we deemed as designable (scTM > 0.5). In these cases, the AF2 prediction shares the same coarse shape as the backbone design but possibly with different secondary-structure ordering and composition. In the length 69 example, we see the closest PDB chain has a TM-score of only 0.65 to the

AF2 prediction but roughly the same secondary-structure ordering as the backbone design. The length 100 sample is a similar case of AF2 producing a roughly similar shape to the backbone design, but has no matching monomer in PDB.

The final category of  $\text{sctM} < 0.25$  reflects failure cases when  $\text{sctM}$  is low. The AF2 predictions in this case have many disordered regions and bear little structural similarity with the original backbone design. Similar PDB chains are not found. We expect that improved generative models of protein backbones would not produce any samples in this category.

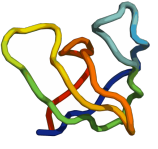
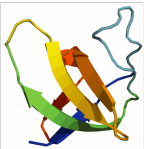
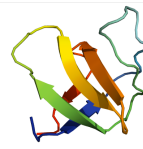
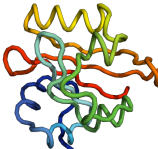
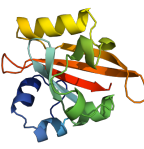
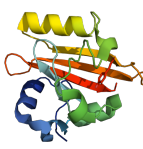
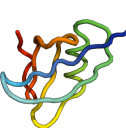
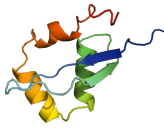
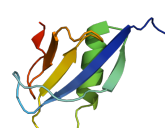
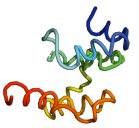
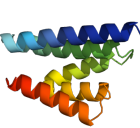
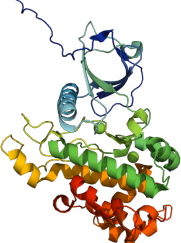
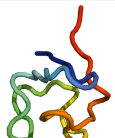
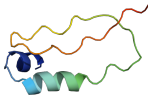
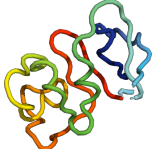
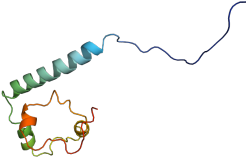
|                    | Backbone design   | AF2 prediction  | PDB parent  |
|--------------------|---|---|---|
| <b>scTM</b> > 0.9  | <br>Length 55    | <br><b>scTM</b> =0.91    | <br>PDB ID=1qkx, TMscore=0.95  |
|                    | <br>Length 125   | <br><b>scTM</b> =0.91    | <br>PDB ID=1f9i, TMscore=0.98  |
| <b>scTM</b> ~ 0.5  | <br>Length 69    | <br><b>scTM</b> =0.5     | <br>PDB ID=7c78, TMscore=0.65  |
|                    | <br>Length 100  | <br><b>scTM</b> =0.51   | <br>PDB ID=4ix3, TMscore=0.56 |
| <b>scTM</b> < 0.25 | <br>Length 54  | <br><b>scTM</b> =0.25  |   |
|                    | <br>Length 120 | <br><b>scTM</b> =0.24 |   |

Figure 5-9: Qualitative analysis of unconditional backbone samples from ProtDiff. The first column displays backbone designs from ProtDiff and their sequence lengths. The second column displays the highest scTM scoring AF2 predictions from the ProteinMPNN sequences of the corresponding backbone design in the first column. The third column displays the closest PDB chain to the AF2 prediction in the second column with the PDB ID and TM-score written below. The third column is blank for the last two rows since no PDB match could be found.



## 5.5 Applicability of SMCDiff beyond proteins: MNIST inpainting

Our goal in this section is to study the applicability of SMCDiff beyond motif-scaffolding, by applying it to inpainting on the MNIST digits dataset. We compare SMCDiff with the replacement method on the task of sampling the remaining half of MNIST digits. We first train DDPM with  $\beta_1 = 10^{-4}, \beta_T = 0.2, T = 1000$  using a small 8-layer CNN on MNIST with batch size 128 and ADAM optimizer for 100 epochs until it is able to generate reasonable MNIST samples (Figure 5-10). We then selected 3 random MNIST images and occluded the right half. The left half would then serve as the conditioning information to the diffusion model (Figure 5-11).




Figure 5-10: Unconditional MNIST samples.



Figure 5-11: Full MNIST images and their occluded halves used for inpainting experiments.

For each occluded image, we fixed a single forward trajectory and sampled 16 images from each method: replacement method and SMCDiff with 16 or 64 particles ( $K$ ). Results are shown in Fig. 5-12. We observe the replacement method can sometimes produce coherent samples as a continuation of the conditioning information, but more often it attempts to produce incoherent digits. SMCDiff on the other hand tends to produce digits that compliment the conditioning information. For more difficult occlusions, such as 5 and 0, SMCDiff can still fail although increasing the number

of particles ( $K = 64$ ) tends to produce samples that are more visually coherent.

It is important to note SMCDiff has additional computation overhead based on the number of particles. It can be more expensive than replacement method but result in higher quality samples. Investigating SMCDiff in more difficult datasets with improved architectures is a direction of future research.




Figure 5-12: MNIST inpainting results for replacement and SMCDiff. See text for explanation.

# Chapter 6

## Conclusion

The motif-scaffolding problem has applications ranging from medicine to material science [King et al., 2012], but remains unsolved for many functional motifs. We have created the first generative modeling approach to motif-scaffolding by developing `ProtDiff`, a diffusion probabilistic model of protein backbones, and `SMCDiff`, a procedure for generating scaffolds conditioned on a motif. Although our experiments were limited to a small set of proteins, our results demonstrate that our procedure is the first capable of generating diverse scaffolds longer than 20 residues with computation time reliably on the order of minutes or less. Our work demonstrates the potential of machine learning methods to be applied in realistic protein design settings.

**General conditional sampling.** `SMCDiff` is applicable to generic DPMs and is not limited to only proteins and motif-scaffolding. While we do not make claims of `SMCDiff` outperforming state-of-the-art conditional diffusion models on other tasks such as image generation, we demonstrate a clear advantage of `SMCDiff` over the replacement method on a toy task of inpainting MNIST images in Section 5.5. Extending `SMCDiff` outside of motif-scaffolding is outside the scope of the present work, but the advantages of a single model for both unconditional and conditional generation warrants additional research.

**Modeling limitations.** Our present results do not indicate our procedure can generalize to motifs that are not present in the training set. We believe improvements in protein modeling could provide better inductive biases for generalization. `ProtDiff`, based on EGNN, is reflection equivariant since it only sees pairwise distances between 3D C- $\alpha$  coordinates. Additionally, `ProtDiff` does not explicitly model primary sequence or side-chains. [Hoogeboom et al. \[2022\]](#) demonstrate the benefits of modeling sequence information in small molecules; joint modeling sequence and structure in a single model could improve the designability of protein scaffolds and backbones as well.

**Data limitations.** We remarked our training set is small due to filtering based on length and oligometry (using only monomeric proteins). Scaling up to longer proteins opens up thousands more examples from the PDB, but in preliminary experiments has proven challenging. Lastly, further development and comparison of methods for motif scaffolding will benefit from standard evaluation benchmarks. Developing a benchmark proved to be difficult since motifs are not labeled in protein databases. It will be important to gather motifs of biological importance in order to guide ML method development towards real-world applications. Because no such benchmarks exist, developing them is a valuable direction for future work.

# Bibliography

- N. Anand and T. Achim. Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*, 2022.
- N. Anand and P. Huang. Generative modeling for protein structures. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- C. B. Anfinsen. Principles that govern the folding of protein chains. *Science*, 181(4096):223–230, 1973.
- I. Anishchenko, S. J. Pellock, T. M. Chidyausiku, T. A. Ramelot, S. Ovchinnikov, J. Hao, K. Bafna, C. Norn, A. Kang, A. K. Bera, F. DiMaio, L. Carter, C. M. Chow, G. T. Montelione, and D. Baker. De novo protein design by deep network hallucination. *Nature*, 600(7889):547–552, 2021.
- P. S. Arunachalam, A. C. Walls, N. Golden, C. Atyeo, S. Fischinger, C. Li, P. Aye, M. J. Navarro, L. Lai, V. V. Edara, et al. Adjuvanting a subunit covid-19 vaccine to induce protective immunity. *Nature*, 594(7862):253–258, 2021.
- S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Kornbluth, N. Molinari, T. E. Smidt, and B. Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13, 2022.
- N. Bennett, B. Coventry, I. Goreshnik, B. Huang, A. Allen, D. Vafeados, Y. P. Peng, J. Dauparas, M. Baek, L. Stewart, , F. DiMaio, S. De Munck, S. N. Savvides, and D. Baker. Improving de novo protein binder design with deep learning. *bioRxiv*, 2022.
- L. Cao, B. Coventry, I. Goreshnik, B. Huang, J. S. Park, K. M. Jude, I. Marković, R. U. Kadam, K. H. G. Verschuere, K. Verstraete, S. T. R. Walsh, N. Bennett, A. Phal, A. Yang, L. Kozodoy, M. DeWitt, L. Picton, L. Miller, E.-M. Strauch, N. D. DeBouver, A. Pires, A. K. Bera, S. Halabiya, B. Hammerson, W. Yang, S. Bernard, L. Stewart, I. A. Wilson, H. Ruohola-Baker, J. Schlessinger, S. Lee, S. N. Savvides, K. C. Garcia, and D. Baker. Design of protein binding proteins from target structure alone. *Nature*, 605(7910):551–560, 2022.
- N. Chopin. Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics*, 32(6):2385–2411, 2004.

- N. Chopin and O. Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*. Springer, 2020.
- B. E. Correia, J. T. Bates, R. J. Loomis, G. Baneyx, C. Carrico, J. G. Jardine, P. Rupert, C. Correnti, O. Kalyuzhniy, V. Vittal, M. J. Connell, E. Stevens, A. Schroeter, M. Chen, S. Macpherson, A. M. Serra, Y. Adachi, M. A. Holmes, Y. Li, R. E. Kleivit, B. S. Graham, R. T. Wyatt, D. Baker, R. K. Strong, J. E. Crowe, Jr, P. R. Johnson, and W. R. Schief. Proof of principle for epitope-focused vaccine design. *Nature*, 507(7491):201–206, 2014.
- J. Dauparas, I. Anishchenko, N. Bennett, H. Bai, R. J. Ragotte, L. F. Milles, B. I. M. Wicky, A. Courbet, R. J. de Haas, N. Bethel, P. J. Y. Leung, T. F. Huddy, S. Pellock, D. Tischer, F. Chan, B. Koepnick, H. Nguyen, A. Kang, B. Sankaran, A. K. Bera, N. P. King, and D. Baker. Robust deep-learning based protein sequence design using ProteinMPNN. *Science*, 378(6615):49–56, 2022.
- A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, 12(656-704):3, 2009.
- A. Doucet, N. De Freitas, and N. J. Gordon. *Sequential Monte Carlo Methods in Practice*, volume 1. Springer, 2001.
- N. Ferruz, S. Schmidt, and B. Höcker. ProtGPT2 is a deep unsupervised language model for protein design. *Nature Communications*, 12(4348), 2022.
- S. J. Fleishman, A. Leaver-Fay, J. E. Corn, E.-M. Strauch, S. D. Khare, N. Koga, J. Ashworth, P. Murphy, F. Richter, G. Lemmon, J. Meiler, and D. Baker. RosettaScripts: a scripting language interface to the Rosetta macromolecular modeling suite. *PLOS ONE*, 6(6):e20161, 2011.
- J. S. Fraser, H. van den Bedem, A. J. Samelson, P. T. Lang, J. M. Holton, N. Echols, and T. Alber. Accessing protein conformational ensembles using room-temperature x-ray crystallography. *Proceedings of the National Academy of Sciences*, 108(39):16247–16252, 2011.
- H. Gelman and M. Gruebele. Fast protein folding kinetics. *Quarterly Reviews of Biophysics*, 47(2):95–142, 2014.
- A. Herbert and M. Sternberg. MaxCluster: a tool for protein structure comparison and clustering. 2008.
- J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet. Video diffusion models. In *Deep Generative Models for Highly Structured Data Workshop, ICLR*, volume 10, 2022.

- E. Hoogeboom, V. G. Satorras, C. Vignac, and M. Welling. Equivariant diffusion for molecule generation in 3D. In *International Conference on Machine Learning*, 2022.
- C. Hsu, R. Verkuil, J. Liu, Z. Lin, B. Hie, T. Sercu, A. Lerer, and A. Rives. Learning inverse folding from millions of predicted structures. In *International Conference on Machine Learning*, 2022.
- B. Huang, Y. Xu, X. Hu, Y. Liu, S. Liao, J. Zhang, C. Huang, J. Hong, Q. Chen, and H. Liu. A backbone-centred energy function of neural networks for protein design. *Nature*, 602(7897):523–528, 2022.
- J. Ingraham, V. Garg, R. Barzilay, and T. Jaakkola. Generative models for graph-based protein design. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- L. Jiang, E. A. Althoff, F. R. Clemente, L. Doyle, D. Rothlisberger, A. Zanghellini, J. L. Gallaher, J. L. Betker, F. Tanaka, C. F. Barbas III, D. Hilvert, K. N. Houk, B. L. Stoddard, and D. Baker. De novo computational design of retro-aldol enzymes. *Science*, 319(5868):1387–1391, 2008.
- J. M. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Zidek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. A. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583 – 589, 2021.
- N. P. King, W. Sheffler, M. R. Sawaya, B. S. Vollmar, J. P. Sumida, I. André, T. Gonen, T. O. Yeates, and D. Baker. Computational design of self-assembling protein nanomaterials with atomic level accuracy. *Science*, 336(6085):1171–1174, 2012.
- D. Kingma, T. Salimans, B. Poole, and J. Ho. Variational diffusion models. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- B. Kuhlman and P. Bradley. Advances in protein structure prediction and design. *Nature Reviews Molecular Cell Biology*, 20(11):681–697, 2019.
- J. S. Lee and P. M. Kim. ProteinSGM: Score-based generative modeling for de novo protein design. *bioRxiv*, 2022.
- Z. Lin, T. Sercu, Y. LeCun, and A. Rives. Deep generative models create new and diverse protein structures. In *Machine Learning for Structural Biology Workshop, NeurIPS*, 2021.
- T. W. Linsky, R. Vergara, N. Codina, J. W. Nelson, M. J. Walker, W. Su, C. O. Barnes, T.-Y. Hsiang, K. Esser-Nobis, K. Yu, Z. B. Reneer, Y. J. Hou, T. Priya,

- M. Mitsumoto, A. Pong, U. Y. Lau, M. L. Mason, J. Chen, A. Chen, T. Berrocal, H. Peng, N. S. Clairmont, J. Castellanos, Y.-R. Lin, A. Josephson-Day, R. S. Baric, D. H. Fuller, C. D. Walkey, T. M. Ross, R. Swanson, P. J. Bjorkman, M. Gale, L. M. Blancas-Mejia, H.-L. Yen, and D.-A. Silva. De novo design of potent and resilient hACE2 decoys to neutralize SARS-CoV-2. *Science*, 370(6521):1208–1214, 2020.
- S. Luo, Y. Su, X. Peng, S. Wang, J. Peng, and J. Ma. Antigen-specific antibody design and optimization with diffusion-based generative models. In *Advances in Neural Information Processing Systems*, 2022.
- M. McPartlon, B. Lai, and J. Xu. A deep SE(3)-equivariant model for learning inverse protein folding. *bioRxiv*, 2022.
- E. Procko, G. Y. Berguig, B. W. Shen, Y. Song, S. Frayo, A. J. Convertine, D. Margineantu, G. Booth, B. E. Correia, Y. Cheng, W. R. Schief, D. M. Hockenbery, O. W. Press, B. L. Stoddard, P. S. Stayton, and D. Baker. A computationally designed inhibitor of an Epstein-Barr viral BCL-2 protein induces apoptosis in infected cells. *Cell*, 157(7):1644–1656, 2014.
- A. Quijano-Rubio, U. Y. Ulge, C. D. Walkey, and D.-A. Silva. The advent of de novo proteins for cancer immunotherapy. *Current Opinion in Chemical Biology*, 56:119–128, 2020.
- C. Saharia, W. Chan, H. Chang, C. A. Lee, J. Ho, T. Salimans, D. J. Fleet, and M. Norouzi. Palette: image-to-image diffusion models. *arXiv preprint arXiv:2111.05826*, 2021.
- V. G. Satorras, E. Hoogeboom, and M. Welling. E(n) equivariant graph neural networks. In *International Conference on Machine Learning*, 2021.
- F. Sesterhenn, C. Yang, J. Bonet, J. T. Cramer, X. Wen, Y. Wang, C.-I. Chiang, L. A. Abriata, I. Kucharska, G. Castoro, S. S. Vollers, M. Galloux, E. Dheilly, S. Rosset, P. Corthésy, S. Georgeon, M. Villard, C.-A. Richard, D. Descamps, T. Delgado, E. Oricchio, M.-A. Rameix-Welti, V. Más, S. Ervin, J.-F. Eléouët, S. Riffault, J. T. Bates, J.-P. Julien, Y. Li, T. Jardetzky, T. Krey, and B. E. Correia. De novo protein design enables the precise induction of RSV-neutralizing antibodies. *Science*, 368(6492), 2020.
- J. B. Siegel, A. Zanghellini, H. M. Lovick, G. Kiss, A. R. Lambert, J. L. StClair, J. L. Gallaher, D. Hilvert, M. H. Gelb, B. L. Stoddard, K. N. Houk, F. E. Michael, and D. Baker. Computational design of an enzyme catalyst for a stereoselective bimolecular Diels-Alder reaction. *Science*, 329(5989):309–313, 2010.
- D.-A. Silva, B. E. Correia, and E. Procko. Motif-driven design of protein-protein interfaces. In *Computational Design of Ligand Binding Proteins*, pages 285–304. 2016.



- J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2015.
- Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- D. Tischer, S. Lisanza, J. Wang, R. Dong, I. Anishchenko, L. F. Milles, S. Ovchinnikov, and D. Baker. Design of proteins presenting discontinuous functional sites using deep learning. *bioRxiv*, 2020.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- J. Wang, S. Lisanza, D. Juergens, D. Tischer, J. L. Watson, K. M. Castro, R. Ragotte, A. Saragovi, L. F. Milles, M. Baek, I. Anishchenko, W. Yang, D. R. Hicks, M. Expòsit, T. Schlichthaerle, J.-H. Chun, J. Dauparas, N. Bennett, B. I. M. Wicky, A. Muenks, F. DiMaio, B. Correia, S. Ovchinnikov, and D. Baker. Scaffolding protein functional sites using deep learning. *Science*, 377(6604):387–394, 2022.
- J. Wu, S. Luo, T. Shen, H. Lan, S. Wang, and J. Huang. EBM-Fold: fully-differentiable protein folding powered by energy-based models. *arXiv preprint arXiv:2105.04771*, 2021.
- P. Xiong, X. Hu, B. Huang, J. Zhang, Q. Chen, and H. Liu. Increasing the efficiency and accuracy of the ABACUS protein sequence design method. *Bioinformatics*, 36(1):136–144, 2020.
- J. Xu and Y. Zhang. How significant is a protein structure similarity with TM-score=0.5? *Bioinformatics*, 26(7):889–895, 2010.
- M. Xu, L. Yu, Y. Song, C. Shi, S. Ermon, and J. Tang. GeoDiff: A Geometric Diffusion Model for Molecular Conformation Generation. In *International Conference on Learning Representations*, 2022.
- C. Yang, F. Sesterhenn, J. Bonet, E. A. van Aalen, L. Scheller, L. A. Abriata, J. T. Cramer, X. Wen, S. Rosset, S. Georgeon, T. Jardetzky, T. Krey, M. Fussenegger, M. Merckx, and B. E. Correia. Bottom-up de novo design of functional proteins with complex structural features. *Nature Chemical Biology*, 17(4):492–500, 2021.
- Y. Zhang and J. Skolnick. TM-align: a protein structure alignment algorithm based on the TM-score. *Nucleic Acids Research*, 33(7):2302–2309, 2005.
- L. Zhou, Y. Du, and J. Wu. 3D shape generation and completion through point-voxel diffusion. In *International Conference on Computer Vision*, pages 5826–5835, 2021.