

Controllable protein design through Feynman–Kac steering

Erik Hartman^{1,*}, Jonas Wallin², Johan Malmström¹, Jimmy Olsson³

¹Division of Infection Medicine, Faculty of Medicine, Lund University, Sweden

²Department of Statistics, Lund University, Sweden

³Department of Mathematics, KTH Royal Institute of Technology, Sweden

*Correspondence: erik.hartman@med.lu.se

Abstract

Diffusion-based models have recently enabled the generation of realistic and diverse protein structures, yet they remain limited in their ability to steer outcomes toward specific functional or biochemical objectives, such as binding affinity or sequence composition. Here we extend the Feynman–Kac (FK) steering framework, an inference-time control approach, to diffusion-based protein design. By coupling FK steering with structure generation, the method guides sampling toward desirable structural or energetic features while maintaining the diversity of the underlying diffusion process. To enable simultaneous generation of both sequence and structure properties, rewards are computed on models refined through ProteinMPNN and all-atom relaxation. Applied to binder design, FK steering consistently improves predicted interface energetics across diverse targets with minimal computational overhead. More broadly, this work demonstrates that inference-time FK control generalizes diffusion-based protein design to arbitrary, non-differentiable, and reward-agnostic objectives, providing a unified and model-independent framework for guided molecular generation.

Introduction

Diffusion models have enabled *de novo* protein design by learning to iteratively transform random noise into structured molecular configurations [1–4]. Trained on large structural datasets, these models learn to generate realistic protein backbones with diverse structural and functional properties and have been used to successfully design enzymes [1], toxin inhibitors [5], vaccine components [6], and peptide therapeutics [7]. Typically, diffusion-based protein design is largely unguided, with limited means to direct the generation toward specific functional or biophysical objectives. This limits the efficiency of the design process and hinders applications that require targeted sampling of proteins with certain properties, such as the design of protein binders with high affinity and favorable features.

However, introducing control into the denoising process remains challenging. Architectural modifications or fine-tuning approaches can steer generation toward specific targets, but they are cumbersome and often require specific architectures or retraining for each new task [8–10]. A more general approach would introduce a principled mechanism to steer diffusion trajectories toward desired structural or biochemical features. Such control would enable directed exploration of protein conformational space and improve the likelihood of sampling functionally relevant candidates without retraining the underlying model or modifying its architecture.

The Feynman–Kac (FK) framework provides a theoretical foundation for inference-time control by

enabling stochastic trajectories to be reweighted according to user-defined reward functions [11–14]. Integrating FK steering, also formulated as *twisting* [15], into diffusion-based protein generation offers a means to steer the sampling process toward target properties without altering the underlying model architecture. In principle, this allows the diffusion dynamics to be guided in real time by rewards reflecting structural, energetic or functional constraints.

Here, we integrate FK steering with RFDiffusion [1] to achieve controllable and property-aware protein generation. We show that FK steering allows diffusion trajectories to be selectively guided toward target properties and improves the design quality of protein binders relative to unguided diffusion. This integration establishes a general framework for adaptive, reward-driven control of generative protein models, advancing *de novo* protein design from passive sampling toward directed functional generation. Our implementation of FK-steered RFDiffusion is available at [GitHub](#).

Main

To introduce guiding into diffusion-based protein generation, we extend the framework of RFDiffusion with a particle-based implementation of Feynman–Kac (FK) steering. This approach reframes protein design as a guided stochastic process, in which diffusion trajectories are continuously reweighted by user-defined rewards that reflect desired structural or biochemical properties. FK steering operates entirely at inference time, dynamically tilting sampling toward high-reward configurations while preserving stochastic diversity. We first formalize the FK formulation for protein diffusion and describe how it can be integrated with RFDiffusion to yield controllable generation. Using charge and secondary structure as illustrative examples, we show that steering enables fine-grained modulation of both sequence and structure. We then analyze how different potential functions and hyperparameters influence denoising trajectories, and finally demonstrate that FK steering enhances the quality and functional relevance of designed binders across multiple pathogenic targets.

Feynman–Kac steering of protein diffusion models

We model protein generation as a discrete denoising diffusion process that transforms random noise into structured molecular configurations. The process consists of two components: a forward diffusion process that progressively adds noise to native protein structures, and a reverse process that learns to recover structure from noise. Each protein conformation $\mathbf{x}_t$ is represented in a rigid-frame coordinate system, where residues are described by their C α positions together with local orientation frames defined by the N-C α -C atoms [1–4]. The forward process defines a sequence of latent variables $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T$, where $\mathbf{x}_0$ represents the native protein backbone and $\mathbf{x}_T$ approaches a standard Gaussian distribution through successive noising steps $q(\mathbf{x}_{t+1} | \mathbf{x}_t)$. The reverse, or generative, process is parameterized by a neural network p_θ with parameters θ that iteratively denoises $\mathbf{x}_t$ to generate artificial data $\mathbf{x}_0$ according to

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=0}^{T-1} p_\theta(\mathbf{x}_t | \mathbf{x}_{t+1}),$$

where $p(\mathbf{x}_T)$ denotes the distribution over noisy structures. During generation, the model begins from $\mathbf{x}_T \sim p(\mathbf{x}_T)$ and successively generates $\mathbf{x}_0$, corresponding to a valid protein backbone.

To introduce control over the generative process, we extend diffusion-based protein modeling with a sequential Monte Carlo formulation inspired by the FK formula [13, 14]. In this framework, an ensemble of particles, each representing a diffusion trajectory $\mathbf{x}_{0:T}$, evolves under the influence of guiding potentials that tilt sampling toward configurations with desired properties. During the generative process, particles are reweighted according to a scalar reward that quantifies the extent to which the intermediate structure at step t satisfies a specified objective, thereby steering the ensemble toward high-reward regions of structural space.

The objective of FK steering is to reweight the diffusion process so that samples are drawn from a *tilted distribution*

$$p^*(\mathbf{x}_0) = \frac{1}{Z} \int p_\theta(\mathbf{x}_{0:T}) e^{r(\mathbf{x}_0)} d\mathbf{x}_{1:T},$$

where Z is a normalizing constant and $r(\mathbf{x}_0)$ is a reward function. To construct a sequential approximation, we define stepwise potentials $G_t(\mathbf{x}_{t:T})$ such that

$$\prod_{t=0}^T G_t(\mathbf{x}_{t:T}) = e^{r(\mathbf{x}_0)}.$$

This allows us to express the tilted target distribution $p^*(\mathbf{x}_0)$ as the time-zero marginal of the *tilted path measure*

$$p^*(\mathbf{x}_{0:T}) \propto p_\theta(\mathbf{x}_{0:T}) \prod_{t=0}^T G_t(\mathbf{x}_{t:T}),$$

which defines the FK-steered diffusion process. During generation, we maintain an ensemble $\mathbf{X}_{t:T} = \{\mathbf{x}_{t:T}^1, \dots, \mathbf{x}_{t:T}^N\}$ of trajectories, referred to as particles, $\mathbf{x}_{t:T}^i$ denoting the i th particle. At each reverse step t , sampling proceeds through a resampling–propagation update. More precisely, the marginal distribution of $\mathbf{x}_{t-1}$ given the current ensemble $\mathbf{X}_{t:T}$ is

$$q_{\text{FK}}(\mathbf{x}_{t-1} \mid \mathbf{X}_{t:T}) = \sum_j w(\mathbf{x}_{t:T}^j) p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t^j),$$

where the weights $w(\mathbf{x}_{t:T}^j) = G_t(\mathbf{x}_{t:T}^j) / \sum_i G_t(\mathbf{x}_{t:T}^i)$ reflect each trajectory’s normalized potential. Algorithmically, each iteration consists of resampling the current particles, with replacement, according to their weights. This duplicates high-potential particles while eliminating low-potential ones, after which each resampled particle is propagated by extending it with a draw from the reverse transition $p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t)$. This sequential Monte Carlo update tilts the ensemble toward high-reward trajectories and enables controllable protein generation (**Fig. 1a**).

Evaluating relevant rewards such as geometries or energies directly on noisy diffusion states $\mathbf{x}_t$ is infeasible, as intermediate structures lack meaningful geometry or sequence context. We therefore estimate rewards on a denoised proxy of the terminal structure, $\hat{\mathbf{x}}_{0|t} = f_\theta(\mathbf{x}_t, t)$, where f_θ denotes the denoising network of the trained diffusion model, which predicts the final structure $\hat{\mathbf{x}}_{0|t}$ from the noisy intermediate state $\mathbf{x}_t$. To obtain a physically coherent and sequence-specific model suitable for reward evaluation, $\hat{\mathbf{x}}_{0|t}$ is passed through ProteinMPNN [16] to generate a sequence $\mathbf{s}_t$, followed by side-chain packing and local relaxation in PyRosetta to produce a refined structure $\tilde{\mathbf{x}}_{0|t}$ [17]. The resulting pair $(\tilde{\mathbf{x}}_{0|t}, \mathbf{s}_t)$ defines the refined state on which the reward function operates (**Fig. 1a**) and we define these refinement steps as part of the reward evaluation pipeline (see **Methods: Rewards**).

To illustrate this approach, we define two reward functions: one based on sequence-level charge (**Fig. 1b**), and another on secondary structure content (**Fig. 1c**). In both cases, steering guides the diffusion process to maximize the objective reward.




Fig. 1 | Feynman–Kac steering enables controllable protein diffusion. **a** Schematic overview of FK steering applied to a protein diffusion model. Starting from Gaussian noise, multiple diffusion trajectories are propagated and resampled at each denoising step according to *steering potentials* G_t , which tilt sampling toward high-reward configurations. Low-reward trajectories (grey) are progressively abandoned, while high-reward trajectories are retained and further denoised toward $\mathbf{x}_0$. This process yields samples from a *tilted distribution* proportional to $p_\theta(\mathbf{x}_0)e^{r(\mathbf{x}_0)}$, where $r(\mathbf{x}_0)$ defines the reward function corresponding to a desired property. On the right, the FK steering loop integrates with the RFdiffusion generative backbone: RFdiffusion is used to denoise $\mathbf{x}_t$ and the predicted denoised proxy $\hat{\mathbf{x}}_{0|t}$ is fed to the reward function. In the reward function, the denoised proxy is used to sample an accompanying sequence via ProteinMPNN, and the structure is refined with packing and relaxation, followed by property computation. The reward is used to compute potentials, which guide the sampling, closing the FK steering loop. **b** Steering by sequence charge reward functions. Lines show the mean net charge over denoising steps for designs guided toward negative charge (orange), positive charge (purple), or unguided (teal) trajectories. Right: UMAP embeddings of generated sequences colored by charge. **c** Steering secondary structure formation. The FK reward function steers sampling toward β -sheet (orange), loop (red), or α -helix (purple) secondary structures. Left: mean secondary structure reward progression over denoising steps. Right: representative examples of resulting structures, along with a triplot where each point is positioned according to its fractional composition of β -sheet, loop, and α -helix secondary structures. Shaded bands indicate the mean ± 1 standard deviation.

Steering is particularly relevant in peptide and protein binder design, where favorable binding energies and geometries are essential for high-affinity interactions. In this setting, we condition the diffusion process on a context $\mathbf{c}$ encoding the target structure and binding hotspots, sampling from $p_\theta(\mathbf{x}_{0:T} | \mathbf{c})$. The reward function is defined over the full complex, enabling FK steering to guide sampling toward conformations that improve predicted binding affinity (ΔG , **Fig. 2a**) and towards the desired secondary structure when binding (**Fig. 2b**).




Fig. 2 | Applying Feynman–Kac steering to binder design. **a** Comparison of binder design trajectories generated using standard RFDiffusion (teal) and FK-steered RFDiffusion (purple) against pneumolysin (PDB: 5CR6). The y-axis shows $-\Delta G$ (higher is better) across denoising steps. Each dot represents a particle at the given timepoint, with lines connecting its immediate predecessors and successors. **b** FK steering enables control over secondary structure composition in binder designs. Lines indicate the mean secondary structure reward for designs guided toward β -sheet (orange), loop (red), or α -helix (purple) motifs. Right: representative designed binders, illustrating successful steering of secondary structure content. Triplot where each point is positioned according to its fractional composition of β -sheet, loop, and α -helix content, showing that FK steering produces distinct populations enriched for the targeted secondary structure class. Shaded bands indicate the mean ± 1 standard deviation.

The choice of potential G_t determines how rewards influence trajectory resampling during the FK-steered diffusion process. In practice, different functional forms of G_t yield distinct trade-offs between responsiveness and stability.

Table 1: Steering potential domain, form and boundary potentials. r_t is used as shorthand for $r_t(\mathbf{x}_t)$.

Potential	Domain	Functional form	Boundary potentials
Immediate	$\mathbf{x}_t$	e^{r_t}	$G_0 = e^{r_0} / \prod_{t=1}^T G_t$
Difference	$(\mathbf{x}_t, \mathbf{x}_{t+1})$	$e^{r_t - r_{t+1}}$	$G_T = 1$
Max	$\mathbf{x}_{t:T}$	$e^{\max_{s \geq t} r_s}$	$G_0 = e^{r_0} / \prod_{t=1}^T G_t$
Sum	$\mathbf{x}_{t:T}$	$e^{\sum_{s=t}^T r_s}$	$G_0 = e^{r_0} / \prod_{t=1}^T G_t$

Among these, the *difference* potential satisfies the exact FK decomposition, while the *immediate*, *max*, and *sum* variants serve as heuristic alternatives that distribute the reward differently along the diffusion trajectory.

For all formulations, the reward includes a guidance scale $\lambda = 1/\tau$,

$$r_t(\mathbf{x}_t) := \lambda r(\hat{\mathbf{x}}_{0|t}),$$

where $\hat{\mathbf{x}}_{0|t} = f_\theta(\mathbf{x}_t, t)$ is the denoised structural estimate of the terminal state $\mathbf{x}_0$. The function $r(\cdot)$ operates on the refined sequence–structure pair derived from $\hat{\mathbf{x}}_{0|t}$ as described above, such that $r_0(\mathbf{x}_0) = r(\mathbf{x}_0)$. The scale τ controls the sharpness of the resampling distribution: lower τ (higher λ) enforces stronger selection toward high-reward configurations but reduces particle diversity, whereas higher τ produces broader, more exploratory ensembles.

Dynamics of FK steering and control/diversity trade-offs

Different potential formulations shape the trade-off between reward maximization and sequence diversity in distinct ways. To evaluate their impact, we designed peptides to bind streptolysin O (PDB: 4HSC), a bacterial cytolytic toxin part of the large superfamily of cholesterol-dependent cytolysins, under identical conditions while varying the potential functions. Strong guidance, especially under *sum*, *max* and *immediate*, yields high terminal rewards (**Fig. 3a**) but quickly collapses diversity, whereas the *difference* potential maintains a steadier balance between exploration and exploitation (**Fig. 3b**). The UMAP projection shows that guided samples remain largely within the distribution of the unguided model, consistent with guidance acting as a form of importance weighting rather than fully re-training the generative prior (**Fig. 3c**). Sequence diversity was quantified as one minus the mean pairwise identity across generated sequences. Shaded bands indicate the mean ± 1 standard deviation.

Varying FK hyperparameters further modulates this balance: lowering the temperature τ enhances rewards but reduces diversity, while resampling less frequently preserves exploration at the cost of control. Increasing the number of particles improved both reward and diversity, and delaying the onset of guidance until the diffusion exited the high-noise regime led to higher final rewards, likely because early rewards are unreliable and can drive premature convergence (**Fig. 3d**).




Fig. 3 | Effect of steering potentials and parameters on FK-guided binder design. **a** Comparison of terminal rewards across different potential formulations. Violin plots show the reward distribution at $t = 0$ for 3 rounds of unguided (blind) diffusion and four FK steering potentials. **b** Sequence diversity trajectories for each potential. Lines indicate mean sequence diversity over denoising steps. **c** UMAP visualization of generated sequences colored by potential type (left) and denoising step (right). **d** Sensitivity of FK steering to key hyperparameters. Top to bottom, left to right: reward and sequence diversity at $t = 0$ as a function of (i) number of particles, (ii) guidance temperature τ , (iii) guidance start step t , and (iv) resampling frequency Δt . Increasing particle count or lowering τ improves mean reward but reduces sequence diversity. Streptolysin O (PDB: 4HSC) was consistently used as the target. The parameters deemed optimal and which were later used for binder design are highlighted with a star (*).

Because the reward function depends on stochastic sequence generation by ProteinMPNN (see **Methods: Rewards** for details), we examined how repeated sequence sampling influenced the performance of FK steering (**Fig. 4a**). At each timestep, multiple sequence-structure pairs ($\mathbf{s}_t, \tilde{\mathbf{x}}_{0|t}$) were generated for each trajectory and averaged to obtain a more reliable estimate of the reward. Increasing the number of sampled pairs per step consistently raised the mean terminal reward (**Fig. 4b**) yet left population diversity largely unchanged (**Fig. 4c**). This demonstrates that more robust reward estimation improves the performance of FK steering without affecting population diversity, in contrast to earlier experiments where higher rewards often came at the cost of reduced diversity (**Fig. 3d**). The standard deviation of rewards decreased along the denoising trajectory (**Fig. 4d**), with the largest uncertainty occurring before $t = 20$. This is consistent with the observed benefit of delaying guidance until later diffusion steps (**Fig. 3d**), when denoised structures provide sufficient geometric and sequence context for reliable evaluation.




Fig. 4 | Effect of repeated reward evaluations on steering. **a** At every time step t , N independent evaluations are generated by applying ProteinMPNN followed by side-chain packing and relaxation, and lastly property evaluation. The used reward function is then formed as the average of r_t across these evaluations. **b** Distribution of rewards at $t = 0$ as a function of the number of repeated evaluations n . **c** Sequence diversity trajectories over denoising steps for different repetition counts. **d** Standard deviation of r_t across repeated runs as a function of diffusion time t . Variability decreases as diffusion progresses toward $t = 0$. Triplicate runs were used for evaluation. Shaded bands indicate the mean ± 1 standard deviation.

Assessment of FK-steered designs by structural prediction with Boltz-2

To evaluate whether FK steering improves structural quality, we compared FK-steered and unguided diffusion designs after co-folding with Boltz-2 [18] (**Fig. 5a**). *Post hoc* structure prediction is commonly used to filter generative outputs, as designs with higher structural agreement to their targets are more likely to yield functional binders. The final designs generated during potential evaluation in *Fig. 4a* were co-folded with the target using Boltz-2. FK steering markedly increased the proportion of co-folded complexes that aligned with the intended diffused structures (**Fig. 5b**), indicating improved sampling of physically plausible configurations. Among the tested potential functions, the *immediate* potential achieved the highest structural concordance, followed by the *sum* potential (**Fig. 5b,c**).




Fig. 5 | Evaluating FK-steered designs using Boltz-2 structural metrics. **a** Structural alignment between FK-steered RFdiffusion (teal) and Boltz-2 (purple) peptide-protein complex (Streptolysin O, PDB: 4HSC). The target receptor is shown in grey. **b** Quantitative comparison of model outputs across evaluation metrics. The number of peptides passing a 5 Å RMSD-cutoff. **c** The swarmplots shows distributions for the peptide distance score, peptide predicted alignment error (pAE), interface ΔG and interface pLDDT. The designs used in this evaluation were the same as the ones presented in *Fig. 4a*.

Designing peptide binders to diverse targets

To assess the generality of FK steering for binder design, we applied FK-steered RFdiffusion to a panel of clinically relevant virulence factors from important gram positive bacterial pathogens, *Streptococcus pyogenes* and *Streptococcus pneumoniae*, including pneumolysin (PLY), streptolysin O (SLO), EndoS, IdeS, and C5a peptidase. These proteins are key pathogenic virulence factors with characterized catalytic or antibody-binding regions which suggests that they accommodate binding. However, to our knowledge, they have not previously been explored as design targets for *de novo* protein binders and are generally structurally diverse. FK steering was applied using the interface free energy reward. The optimal parameters identified during benchmarking were used: the *immediate* potential, $\Delta t = 2$, $n_{\text{particles}} = 50$, $t_{\text{start}} = 20$, and $\tau = 10$. Across all targets, FK-steered RFdiffusion increased the binding reward compared to unguided diffusion, indicating that FK control improves predicted interface affinity across diverse targets (**Fig. 6**).




Fig. 6 | FK steering improves binder design across diverse bacterial virulence targets. FK-steered RFdiffusion applied to the *Streptococcus* virulence factors streptolysin O (SLO), pneumolysin (PLY), EndoS, IdeS, and C5a peptidase. For each target, the left panels show the receptor (grey) and its known antibody (pink) or active-site region, which define the design interface. Middle panels show reward trajectories over diffusion timesteps comparing standard RFdiffusion (teal) and FK-steered RFdiffusion (purple). Each dot represents a particle at the given timepoint, with lines connecting its immediate predecessors and successors. Shaded bands indicate the mean ± 1 standard deviation. Right panels show representative FK-designed peptides (purple) bound to their corresponding target domains (grey).

Together, these results demonstrate that FK steering provides an effective and general mechanism for introducing controllability into diffusion-based protein generation. Across multiple property classes and biological targets, FK steering enhances model performance without modifying or retraining the underlying diffusion network, establishing it as a flexible inference-time framework for directed protein design.

Discussion

The FK framework provides a principled means to control diffusion-based protein design without retraining or modifying the generative model. By reweighting trajectories according to user-defined rewards, FK steering enables targeted exploration of structural space while maintaining stochastic diversity. When applied to RFdiffusion and coupled with sequence recovery and relaxation, we show that FK steering effectively directs diffusion models toward desired biochemical and structural properties, including charge distribution, secondary structure composition, and predicted binding affinity. The method operates entirely at inference time and supports flexible, property-specific rewards defined over both structure and sequence.

A central challenge in applying FK steering for protein design lies in defining a potential function that consistently guides the diffusion process toward regions of structural space associated with desired molecular properties. The first difficulty lies in recovering a denoised structure $\hat{\mathbf{x}}_{0|t}$ from a noisy sample $\mathbf{x}_t$, which is fundamentally approximate since many plausible structures can correspond to the same noisy representation. The latter represents a broader limitation of protein design itself, where accurate prediction of properties such as binding affinity or stability remains inherently difficult. In this work, we address these challenges by using the diffusion network to predict $\hat{\mathbf{x}}_{0|t}$, followed by sequence generation with ProteinMPNN and structural relaxation in PyRosetta to obtain physically coherent sequence–structure pairs $(\mathbf{s}_t, \hat{\mathbf{x}}_{0|t})$ suitable for evaluation. Both steps alter the potential landscape and impact steering performance. Continued improvements in denoising models and algorithms, and in biophysical property predictors, are therefore likely to enhance the accuracy and stability of FK steering.

Another practical consideration concerns the selection and optimization of steering parameters. The performance of FK control depends on maintaining an appropriate balance between exploration and exploitation, which is influenced by the choice of potential as well as by parameters such as resampling frequency and the time of guidance onset. In this study, we characterized the general behaviour of steering across these parameters, but the optimal configuration is likely dependent on the specific reward formulation and design objective. Developing adaptive schemes that automatically tune these parameters may further improve robustness and generality across different design objectives.

Despite these considerations, FK steering remains computationally efficient, as it operates entirely at inference time and requires no gradient-based optimization or retraining of the generative model. The method adds only a minor overhead relative to standard sampling, making it practical for large-scale or iterative design tasks. Moreover, its modular formulation allows straightforward extension to new reward functions, enabling users to incorporate additional biochemical or structural objectives without modification of the underlying framework. This flexibility makes FK steering well suited for integration with emerging predictive models and experimental feedback, providing a general foundation for controllable protein generation.

Method

Feynman–Kac implementation

FK steering was implemented as a wrapper around RFdiffusion [1] (commit e220924) and uses its pretrained denoising network (p_θ) and diffusion schedule to generate reverse trajectories. The FK layer introduces particle-based guidance, where an ensemble of $n_{\text{particles}}$ particle trajectories evolve randomly through reweighting and resampling operations every Δt timestep (starting at t_{start}) according to user-defined reward functions.

At each guidance step, a reward function $r_t(\mathbf{x}_t)$ is evaluated for every particle, and potentials G_t are computed according to the chosen formulation. To maintain numerical stability, exponentiation is performed after subtracting the maximum value and clipping ($r_t - r_{\text{max}} \geq -10^3$). Particle weights are

then assigned as normalized potentials

$$w(\mathbf{x}_{t:T}^j) = \frac{G_t(\mathbf{x}_{t:T}^j)}{\sum_i G_t(\mathbf{x}_{t:T}^i)}$$

Guided diffusion proceeds from the initial noisy state toward $t = 1$, alternating between denoising by p_θ and resampling by the FK update.

The computational cost of FK steering scales linearly with the number of diffusion particles and reward evaluations. Each guidance step requires $n_{\text{particles}}$ forward passes through RFDiffusion and corresponding reward computations, resulting in a total complexity of $\mathcal{O}(n_{\text{particles}} \cdot (C_{p_\theta} + C_r))$ where C_{p_θ} is the cost of inference of the denoising network and C_r the cost of computing the reward. In practice, diffusion steps are efficiently parallelized across GPU devices, while reward computations, particularly those involving ProteinMPNN inference and PyRosetta relaxation, can be executed in parallel on CPUs, substantially reducing wall time.

Sequence recovery and packing

To obtain a physically meaningful estimate of $\mathbf{x}_0$ for reward evaluation, sequence design from ProteinMPNN (commit 8907e66) [16] was integrated with all-atom refinement in PyRosetta as part of the reward pipeline. This procedure transforms a coarse, denoised backbone representation into a fully atomic, sequence-specific structure better suited for downstream evaluation.

Let $\mathbf{x}_0 \in \mathbb{R}^{3L}$ denote the residue-frame backbone coordinates, and let $\mathcal{A}^L$ be the discrete amino acid sequence space of length L . ProteinMPNN defines a conditional distribution $P_{\text{MPNN}}(\mathbf{s} \mid \mathbf{x}_0)$, from which sequences $\mathbf{s}$ are sampled at a temperature of 0.2 using solubility-optimized model weights. Each sequence $\mathbf{s}$ is threaded onto the backbone to generate a full-atom model $\mathbf{x}_{\text{atom}}(\mathbf{x}_0, \mathbf{s}) \in \mathbb{R}^{3N}$, where N is the number of heavy atoms.

We define the transformation $(\mathbf{x}_0, \mathbf{s}) \mapsto (\tilde{\mathbf{x}}_0, \mathbf{s})$, which represents the combined process of sampling a sequence and reconstructing the corresponding all-atom model through side-chain packing and local energy minimization. The refined coordinates $\tilde{\mathbf{x}}_0$ are obtained by discrete rotamer sampling followed by continuous minimization of torsional angles $\chi_1, \chi_2, \dots$ under the Rosetta energy function, resolving steric clashes and optimizing the side chain poses. This defines a conditional distribution

$$(\tilde{\mathbf{x}}_0, \mathbf{s}) \sim P_{\text{ref}}(\cdot \mid \mathbf{x}_0),$$

which maps each denoised backbone to an ensemble of relaxed, sequence-specific structures. Reward evaluation is thus performed on physically and chemically coherent sequence–structure pairs.

Rewards

In the FK framework, the reward function is a deterministic function. In our case, evaluating such a reward requires intermediate sampling and refinement steps that introduce randomness. Specifically, the combination of ProteinMPNN sequence sampling and PyRosetta relaxation defines a conditional distribution $P_{\text{ref}}(\cdot \mid \mathbf{x}_0)$, from which realizations $(\tilde{\mathbf{x}}_0, \mathbf{s})$ are drawn. These steps are best regarded as part of the FK update mechanism rather than as intrinsic randomness in the reward function itself. However, for clarity, we treat the entire procedure as an evaluation pipeline that provides an approximation to a deterministic reward defined on the expected refined structure, allowing us to tailor the reward pipeline based on the objective.

To obtain a stable estimate of the expected reward, each denoised backbone $\mathbf{x}_0$ is evaluated multiple times. Each realization

$$(\tilde{\mathbf{x}}_0^{(i)}, \mathbf{s}_t^{(i)}) \sim P_{\text{ref}}(\cdot \mid \mathbf{x}_0), \quad i = 1, \dots, n,$$

produces an independent structure–sequence pair. A scalar reward can then be aggregated as either

the mean or maximum across samples,

$$r^{\text{mean}}(\mathbf{x}_0) = \frac{1}{n} \sum_{i=1}^n r(\tilde{\mathbf{x}}_0^{(i)}, \mathbf{s}_t^{(i)}), \quad r^{\text{max}}(\mathbf{x}_0) = \max_i r(\tilde{\mathbf{x}}_0^{(i)}, \mathbf{s}_t^{(i)}).$$

Because r operates on coherent sequence–structure pairs, any scalar property derived from either the structure or sequence can be incorporated as a reward.

For binder design, the interface quality is quantified by the computed binding free energy ΔG between the designed and target chains, with

$$r_{\text{bind}} = -\Delta G.$$

Sequence-level objectives can also be introduced. To control the electrostatic character, a charge reward penalizes deviation from a target net charge Q^* at pH 7,

$$r_{\text{charge}} = -|Q - Q^*|.$$

Secondary structure composition is constrained through a combined geometry- and sequence-based reward. Structural content is estimated using DSSP on the refined backbone and residue-based propensities from the sequence, combined with fixed weights (0.8 for DSSP and 0.2 for sequence) to yield the helix, β -sheet, and loop fractions (α, β, ℓ) . The reward penalizes deviation from user-specified targets $(\alpha^*, \beta^*, \ell^*)$,

$$r_{\text{SS}} = w_\alpha(1 - |\alpha - \alpha^*|) + w_\beta(1 - |\beta - \beta^*|) + w_\ell(1 - |\ell - \ell^*|),$$

where w_α , w_β , and w_ℓ control the relative contribution of each term. When steering toward a specific secondary structure, that term is weighted fourfold relative to the others.

Steering configuration evaluation

To characterize the behavior of FK steering during binder design, we performed a systematic configuration sweep targeting streptolysin O, with hotspot residues 110, 115, and 117. Each condition was repeated three times.

To examine the effect of the steering potential, three independent FK-steered runs were carried out using 50 particles and a binder length ($L_{\text{binder}} = 24$). To investigate the influence of the number of particles ($n_{\text{particles}}$), guidance temperature (τ), resampling interval (Δt), and guidance onset (t_{start}), we performed three repeats of guided diffusion per parameter, keeping all others fixed at the default values presented in **Table 2**.

Table 2. Default parameters during the parameter sweep.

Parameter	Value
L_{binder}	15
τ	10
$n_{\text{particles}}$	20
t_{start}	50
$\Delta t = 2$	2
G_t	Immediate

In each case, only the parameter under evaluation was varied.

Docking and scoring

To evaluate the structural quality of FK-steered designs, we performed peptide-protein docking and scoring using the Boltz-2 model [18] implemented in the BoPep framework [19]. The final designs from

the potential evaluation were used. Each designed peptide was docked to the target with ten recycling steps and five diffusion samples per model, using the target structure as the template for that chain. The BoPep framework was also used to score the complexes.

Campaigns

We applied FK-steered RFDiffusion to design peptide binders against clinically relevant virulence factors from *Streptococcus pyogenes* and *Streptococcus pneumoniae*. The targets: EndoS (PDB: 8A49), IdeS (PDB: 8A47), pneumolysin (PLY, PDB: 5CR6), C5a peptidase (PDB: 3EIF), and streptolysin O (SLO, PDB: 4HSC), are factors for which neutralizing antibodies have been described and/or have known catalytic sites [20, 21]. These guided the selection of receptor hotspot residues defining the design interfaces.

For SLO, the targeted epitope comprised residues A110, A115, and A117. EndoS was targeted at two antibody-binding regions on chain C: interface 1 spanning residues 765-916 with hotspots C798, C835, and C909, and interface 2 spanning residues 295-421 with hotspots C314, C315, and C316. IdeS was targeted at three distinct regions of chain C: an antibody-binding interface near residues C185-C187, a second interface near C255, C258, and C322, and the catalytic triad formed by C93, C264, and C286. PLY was targeted at an inhouse defined epitope (unpublished). C5a peptidase was targeted at the catalytic region, with hotspots A423, A360, A356 and A426.

All campaigns used 50 diffusion particles, resampling every two steps, guidance from timestep 20, and the *immediate* FK potential acting on interface ΔG reward. Guided and unguided runs were performed for each target. The beta sheet checkpoint was used for RFDiffusion weights.

Code availability

FK-steered RFDiffusion is available at [GitHub](#)

Acknowledgements

We thank Di Tang and Alejandro Gomez Toledo for providing their insights on target selection.

References

- [1] Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with RFDiffusion. *Nature*, 620(7976):1089–1100, July 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06415-8. URL <http://dx.doi.org/10.1038/s41586-023-06415-8>.
- [2] Jason Yim, Brian L. Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. SE(3) diffusion model with application to protein backbone generation. 2023. doi: 10.48550/ARXIV.2302.02277. URL <https://arxiv.org/abs/2302.02277>.
- [3] Kevin E. Wu, Kevin K. Yang, Rianne van den Berg, Sarah Alamdari, James Y. Zou, Alex X. Lu, and Ava P. Amini. Protein structure generation via folding diffusion. *Nature Communications*, 15(1), February 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-45051-2. URL <http://dx.doi.org/10.1038/s41467-024-45051-2>.

- [4] Brian L. Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and Tommi Jaakkola. Diffusion probabilistic modeling of protein backbones in 3D for the motif-scaffolding problem, 2022. URL <https://arxiv.org/abs/2206.04119>.
- [5] Susana Vázquez Torres, Melisa Benard Valle, Stephen P. Mackessy, Stefanie K. Menzies, Nicholas R. Casewell, Shirin Ahmadi, Nick J. Burlet, Edin Muratspahić, Isaac Sappington, Max D. Overath, Esperanza Rivera-de Torre, Jann Ledergerber, Andreas H. Laustsen, Kim Boddum, Asim K. Bera, Alex Kang, Evans Brackenbrough, Iara A. Cardoso, Edouard P. Crittenden, Rebecca J. Edge, Justin Decarreau, Robert J. Ragotte, Arvind S. Pillai, Mohamad Abedi, Hannah L. Han, Stacey R. Gerben, Analisa Murray, Rebecca Skotheim, Lynda Stuart, Lance Stewart, Thomas J. A. Fryer, Timothy P. Jenkins, and David Baker. De novo designed proteins neutralize lethal snake venom toxins. *Nature*, 639(8053):225–231, January 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08393-x. URL <http://dx.doi.org/10.1038/s41586-024-08393-x>.
- [6] Sanela Rankovic, Kenneth D. Carr, Justin Decarreau, Rebecca Skotheim, Ryan D. Kibler, Sebastian Ols, Sangmin Lee, Jung-Ho Chun, Marti R. Tooley, Justas Dauparas, Helen E. Eisenach, Matthias Glögl, Connor Weidle, Andrew J. Borst, David Baker, and Neil P. King. Computational design of bifaceted protein nanomaterials. *Nature Materials*, 24(10):1635–1643, July 2025. ISSN 1476-4660. doi: 10.1038/s41563-025-02295-7. URL <http://dx.doi.org/10.1038/s41563-025-02295-7>.
- [7] Stephen A. Rettie, David Juergens, Victor Adebomi, Yensi Flores Bueso, Qinqin Zhao, Alexandria N. Leveille, Andi Liu, Asim K. Bera, Joana A. Wilms, Alina Üffing, Alex Kang, Evans Brackenbrough, Mila Lamb, Stacey R. Gerben, Analisa Murray, Paul M. Levine, Maika Schneider, Vibha Vasireddy, Sergey Ovchinnikov, Oliver H. Weiergräber, Dieter Willbold, Joshua A. Kritzer, Joseph D. Mougous, David Baker, Frank DiMaio, and Gaurav Bhardwaj. Accurate de novo design of high-affinity protein-binding macrocycles using deep learning. *Nature Chemical Biology*, June 2025. ISSN 1552-4469. doi: 10.1038/s41589-025-01929-w. URL <http://dx.doi.org/10.1038/s41589-025-01929-w>.
- [8] Dongjun Kim, Yeongmin Kim, Se Jung Kwon, Wanmo Kang, and Il-Chul Moon. Refining Generative Process with Discriminator Guidance in Score-based Diffusion Models, 2022. URL <https://arxiv.org/abs/2211.17091>.
- [9] Susung Hong, Gyuseong Lee, Wooseok Jang, and Seungryong Kim. Improving Sample Quality of Diffusion Models Using Self-Attention Guidance, 2022. URL <https://arxiv.org/abs/2210.00939>.
- [10] Sidney Lyayuga Lisanza, Jacob Merle Gershon, Samuel W. K. Tipps, Jeremiah Nelson Sims, Lucas Arnoldt, Samuel J. Hendel, Miriam K. Simma, Ge Liu, Muna Yase, Hongwei Wu, Claire D. Tharp, Xinting Li, Alex Kang, Evans Brackenbrough, Asim K. Bera, Stacey Gerben, Bruce J. Wittmann, Andrew C. McShan, and David Baker. Multistate and functional protein design using RoseTTAFold sequence space diffusion. *Nature Biotechnology*, 43(8):1288–1298, September 2024. ISSN 1546-1696. doi: 10.1038/s41587-024-02395-w. URL <http://dx.doi.org/10.1038/s41587-024-02395-w>.
- [11] Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A General Framework for Inference-time Scaling and Steering of Diffusion Models, 2025. URL <https://arxiv.org/abs/2501.06848>.
- [12] Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alán Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-Kac Correctors in Diffusion: Annealing, Guidance, and Product of Experts, 2025. URL <https://arxiv.org/abs/2503.02819>.
- [13] Pierre Del Moral. *Feynman-Kac Formulae*. Springer New York, 2004. ISBN 9781468493931. doi: 10.1007/978-1-4684-9393-1. URL <http://dx.doi.org/10.1007/978-1-4684-9393-1>.
- [14] Nicolas Chopin and Omiros Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*. Springer International Publishing, 2020. ISBN 9783030478452. doi: 10.1007/978-3-030-47845-2. URL <http://dx.doi.org/10.1007/978-3-030-47845-2>.

- [15] Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham. Practical and Asymptotically Exact Conditional Sampling in Diffusion Models. 2023. doi: 10.48550/ARXIV.2306.17775. URL <https://arxiv.org/abs/2306.17775>.
- [16] 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, October 2022. ISSN 1095-9203. doi: 10.1126/science.add2187. URL <http://dx.doi.org/10.1126/science.add2187>.
- [17] Rebecca F. Alford, Andrew Leaver-Fay, Jeliasko R. Jeliaskov, Matthew J. O’Meara, Frank P. DiMaio, Hahnbeom Park, Maxim V. Shapovalov, P. Douglas Renfrew, Vikram K. Mulligan, Kalli Kappel, Jason W. Labonte, Michael S. Pacella, Richard Bonneau, Philip Bradley, Roland L. Dunbrack, Rhiju Das, David Baker, Brian Kuhlman, Tanja Kortemme, and Jeffrey J. Gray. The Rosetta All-Atom Energy Function for Macromolecular Modeling and Design. *Journal of Chemical Theory and Computation*, 13(6):3031–3048, May 2017. ISSN 1549-9626. doi: 10.1021/acs.jctc.7b0125. URL <http://dx.doi.org/10.1021/acs.jctc.7b0125>.
- [18] Saro Passaro, Gabriele Corso, Jeremy Wohlwend, Mateo Reveiz, Stephan Thaler, Vignesh Ram Somnath, Noah Getz, Tally Portnoi, Julien Roy, Hannes Stark, David Kwabi-Addo, Dominique Beaini, Tommi Jaakkola, and Regina Barzilay. Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. June 2025. doi: 10.1101/2025.06.14.659707. URL <http://dx.doi.org/10.1101/2025.06.14.659707>.
- [19] Erik Hartman, Firdaus Samsudin, Peter J. Bond, Artur Schmidtchen, and Johan Malmström. Mining the endogenous peptidome for peptide binders with deep learning-driven optimization and molecular simulations. January 2025. doi: 10.1101/2025.01.20.633551. URL <http://dx.doi.org/10.1101/2025.01.20.633551>.
- [20] Abigail S. L. Sudol, John Butler, Dylan P. Ivory, Ivo Tews, and Max Crispin. Extensive substrate recognition by the streptococcal antibody-degrading enzymes IdeS and EndoS. *Nature Communications*, 13(1), December 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-35340-z. URL <http://dx.doi.org/10.1038/s41467-022-35340-z>.
- [21] Todd F. Kagawa, Maurice R. O’Connell, Pania Mouat, Max Paoli, Paul W. O’Toole, and Jakki C. Cooney. Model for Substrate Interactions in C5a Peptidase from *Streptococcus pyogenes*: A 1.9 Å Crystal Structure of the Active Form of ScpA. *Journal of Molecular Biology*, 386(3):754–772, February 2009. ISSN 0022-2836. doi: 10.1016/j.jmb.2008.12.074. URL <http://dx.doi.org/10.1016/j.jmb.2008.12.074>.