

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CONSTRINED DIFFUSION FOR PROTEIN DESIGN WITH HARD STRUCTURAL CONSTRAINTS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Diffusion models offer a powerful means of capturing the manifold of realistic protein structures, enabling rapid design for protein engineering tasks. However, existing approaches observe critical failure modes when precise constraints are necessary for functional design. To this end, we present a constrained diffusion framework for structure-guided protein design, ensuring strict adherence to functional requirements while maintaining precise stereochemical and geometric feasibility. The approach integrates proximal feasibility updates with ADMM decomposition into the generative process, scaling effectively to the complex constraint sets of this domain. We evaluate on challenging protein design tasks, including motif scaffolding and vacancy-constrained pocket design, while introducing a novel curated benchmark dataset for motif scaffolding in the PDZ domain. Our approach achieves state-of-the-art, providing perfect satisfaction of bonding and geometric constraints with no degradation in structural diversity.

## 1 INTRODUCTION

Diffusion models have revolutionized protein engineering with notable successes demonstrated in the design of protein monomers, assemblies, and protein binders against biomolecular targets (Watson et al., 2023). In many cases, predefined binding or catalytic motifs are introduced into designed proteins via motif scaffolding but there are no guarantees that the generated backbones will accurately include the motif (Trippe et al., 2022; Didi et al., 2023). Furthermore, the motifs are typically pre-defined as structural fragments, rather than more physically-based (e.g. hydrogen bonds to chosen target residues), which narrows the accessible design space (Song et al., 2024). Negative space constraints (e.g. tunnels for substrate access and product egress), while a ubiquitous feature of naturally evolved proteins such as enzymes, are not readily incorporated in current generative protein models. These obstacles restrict the scope of design goals accessible to current methods.

These limitations highlight the broader challenge of designing structured objects under strict feasibility. Many existing approaches augment diffusion models through constraint guidance (Ho et al., 2020; Ho & Salimans, 2022). For example, Gruver et al. (2023) inject soft constraints into discrete sequence prediction for protein design through gradient-based guidance; however, while guidance approaches result in increased feasibility rates, they fail to consistently provide constraint adherent outputs. Others adopt post-processing optimizations which more rigorously target the constraint set. However, these methods rely on either simplifications of the highly nonconvex constraint set (e.g., matching an existing ligand template) or result in samples falling outside the data manifold (Bergues et al., 2025; Christopher et al., 2024). A more effective approach is to inject constraint into the generative process, e.g., by projecting intermediate states back to the feasible set (Christopher et al., 2024). Yet projecting noisy states early in the sampling process has been shown to disrupt the diffusion trajectory, potentially biasing samples (Blanke et al., 2025).

To resolve this tension, this paper proposes to view constrained diffusion through the lens of stochastic proximal methods. To enable strict constraint enforcement throughout the diffusion process, while removing the need to project at earlier noisy states, we introduce a framework which applies *final-state corrections*. Proximal steps are applied to a predicted clean posterior, rather than on a noisy intermediate state, and the feasible clean state is then renoised to steer the sampling trajectory along the data manifold, while converging to exact feasibility at the terminal state. A schematic illustration of the proposed scheme is provided in Figure 1.

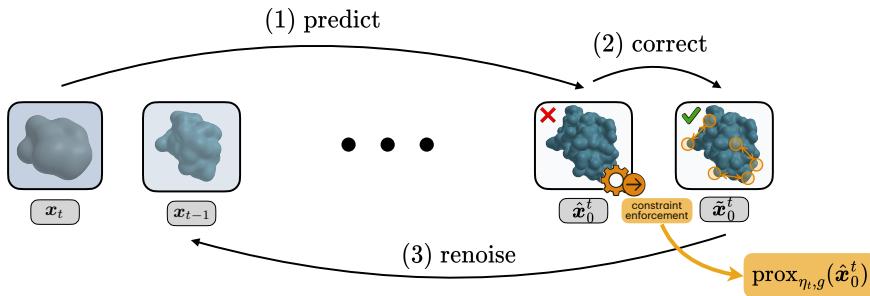


Figure 1: Illustration of our stochastic proximal sampling for structured-constrained protein design.

**Contributions.** This paper makes five key contributions: **(1)** It introduces a stochastic proximal method for constrained diffusion feasibility; **(2)** based on this view, it proposes a consensus-based ADMM scheme that separates local stereochemistry from sparse global couplings; **(3)** it provides a theoretical analysis characterizing the convergence to the constraint set and provides arguments for why final-state projection is preferred over per-step projections; **(4)** it demonstrates the efficacy of the approach on challenging protein design tasks with nonconvex constraints, including global topology (e.g., chain closure, ligand binding feasibility) and local stereochemistry (bond lengths, angles, and chirality), achieving perfect constraint satisfaction and providing state-of-the-art performance; **(5)** it introduces a novel curated benchmark for protein motif scaffolding in PDZ domains, providing the first systematic standard for constrained diffusion methods in modular domain engineering.

## 2 SETTINGS AND BACKGROUND

The goal of *de novo* protein design is to generate three-dimensional representations which satisfy physical plausibility and functional requirements, such as protein motif scaffolding, binder design, and monomer generation. This task is fundamentally constrained by the physics and chemistry of proteins, where bond geometries must be preserved, chains must remain connected, and higher-level properties such as specific inter-chain interactions or interface complementarity must be realized.

**Diffusion-based protein backbone generation.** Let  $p_{\text{data}}$  denote the unknown distribution of clean molecular structure representations  $\mathbf{x}_0 \in \mathbb{R}^d$  (e.g., atomic coordinates). Consider a constrained feasibility set  $\mathcal{C} \subset \mathbb{R}^d$ , encoding physical and geometric constraints, seeking to sample from the target distribution

$$p_{\mathcal{C}}(\mathbf{x}_0) \propto p_{\text{data}}(\mathbf{x}_0) \mathbb{1}\{\mathbf{x}_0 \in \mathcal{C}\}, \quad (\star)$$

where  $\mathbb{1}\{\mathbf{x}_0 \in \mathcal{C}\}$  is an indicator function on  $\mathcal{C}$ . To achieve this, a diffusion model learns to reconstruct samples from the data distribution  $p_{\text{data}}$  by coupling a *forward* noising Markov chain with a learned *reverse* denoising chain (Song et al., 2020; Ho et al., 2020).

**Forward process.** Diffusion models construct a Markov chain  $\{\mathbf{x}_t\}_{t=0}^T$  starting from  $\mathbf{x}_0 \sim p_{\text{data}}$ . At each step, Gaussian noise is added according to a fixed variance schedule  $\{\alpha_t\}_{t=0}^T$ . This process admits a closed-form marginal distribution,  $q(\mathbf{x}_t \mid \mathbf{x}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t)I)$ , where  $\bar{\alpha}_t = \prod_{s=1}^t (1 - \alpha_s)$ . As  $t \rightarrow T$ , the distribution converges to an isotropic Gaussian  $q(\mathbf{x}_T) \approx \mathcal{N}(\mathbf{0}, I)$ .

**Reverse process.** A generative model learns to approximate the reverse dynamics, sampling  $\mathbf{x}_{t-1}$  given  $\mathbf{x}_t$ . Since the true reverse kernel  $q(\mathbf{x}_{t-1} \mid \mathbf{x}_t)$  is intractable, a neural network  $\mathbf{x}_\theta(\mathbf{x}_t, t)$  is introduced to parameterize this transition. At inference, the learned reverse transitions are applied iteratively, gradually denoising a random Gaussian vector  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, I)$  into a clean structure  $\mathbf{x}_0$ . Standard sampling processes will result in outputs distributed approximately as  $p_{\text{data}}$ , while our underlying goal is instead to generate outputs distributed in  $p_{\mathcal{C}}$ . Generating from  $p_{\mathcal{C}}$ , requires the sampling procedure to be modified to incorporate constraints, as discussed in the next section.

## 3 CHALLENGES WITH CONSTRAINED DIFFUSION

As previously noted, diffusion models natively learn to reconstruct samples from an unconstrained data distribution  $p_{\text{data}}$ , which is misaligned with the true goal for constrained generation described

108 by Equation (\*). In other domains, this gap is often addressed by formulating constrained sampling  
 109 as an optimization problem, and recent work has extended this perspective to diffusion models by  
 110 proposing a general framework for inference-time constrained generation (Christopher et al., 2024).  
 111 The approach can be viewed as a sequential optimization problem:

$$112 \quad \min_{\{\mathbf{x}_t\}_{t=0}^T} \sum_{t=0}^T \ell_t(\mathbf{x}_t, \mathbf{x}_0) \quad \text{s.t. } \mathbf{x}_t \in \mathcal{C} \quad \forall t \quad (1)$$

115 where the single stage cost  $\ell_t(\mathbf{x}_t, \mathbf{x}_0) := -\log p(\mathbf{x}_t \mid \mathbf{x}_0)$ , and the constraint set  $\mathcal{C}$  captures geo-  
 116 metric or structural feasibility. Enforcing these constraints typically involves applying a projection  
 117 operator  $\Pi_{\mathcal{C}}(\mathbf{x}_t) = \arg \min_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{x}_t\|_2^2$  after each reverse diffusion step. This is appealing be-  
 118 cause it embeds feasibility directly into the generative process; however, enforcing strict projections  
 119 at every step raises two domain-specific challenges:

- 120 1. As observed by Blanke et al. (2025), projections on intermediate states introduces statistical  
 121 biases. This arises as intermediate samples concentrate near constraint boundaries. This issue  
 122 has also been reported for soft guidance schemes, where increased weight on the guidance terms  
 123 tends to disrupt the diffusion trajectory and degrading performance (Dhariwal & Nichol, 2021;  
 124 Nichol & Dhariwal, 2021; Ho & Salimans, 2022).
- 125 2. Additionally, intermediate feasibility requires projecting highly noisy states onto complex, non-  
 126 convex constraints, which can result in solutions trapped in local minima (Pardalos & Vavasis,  
 127 1991; Boyd & Vandenberghe, 2004), disrupting the diffusion performance while simultaneously  
 128 limiting the efficacy of feasibility updates.

129 Imposing constraints precisely on noisy samples may be reasonable under convexity assumptions,  
 130 where the theoretical provisions apply, but within protein generation settings characterized by non-  
 131 convex constraints, it is necessary to develop inference-time strategies which do not fundamentally  
 132 rely on  $\mathbf{x}_t \in \mathcal{C}$ . In the next section, we present our proposed method, designed intentionally with  
 133 this principle in mind.

## 135 4 REVERSE DIFFUSION AS PROXIMAL OPTIMIZATION

137 To effectively sample constraint compliant outputs  $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x}_0) \mathbb{1}\{\mathbf{x}_0 \in \mathcal{C}\}$  (Equation (\*)), we  
 138 design an inference-time method that converges to  $\mathcal{C}$ , where the terminal distribution  $\pi_0$  minimizes  
 139  $\text{KL}(\pi_0 \mid p_{\mathcal{C}})$ . Under this framing, a single reverse step of a diffusion sampler is viewed as an optimi-  
 140 zation problem in which the denoiser provides a data-driven “anchor”, and feasibility is enforced  
 141 by penalizing the distance to the constraint set.

142 Since  $p_{\text{data}}$  is only available through a denoiser, the reverse process is realized incrementally. Each  
 143 state  $t$  of the reverse diffusion is composed of three stages: (1) **predict** the clean structure  $\hat{\mathbf{x}}_0^t$  from  
 144 the current noisy state  $\mathbf{x}_t$ , (2) **correct** this prediction by a proximal operator  $\text{prox}_{\eta_t, g}$  that enforces  
 145 feasibility, and (3) **renoise** the corrected clean structure with the forward kernel, denoted  $\text{FWD}(\cdot, \varepsilon)$ ,  
 146 to obtain the next noisy sample. The procedure is outlined in Figure 1 and summarized as:

$$147 \quad \mathbf{x}_t \xrightarrow{\text{predict}} \hat{\mathbf{x}}_0^t \xrightarrow{\text{prox}} \tilde{\mathbf{x}}_0^t \xrightarrow{\text{FWD}} \mathbf{x}_{t-1}.$$

148 This modus operandi has strong theoretical properties as discussed in Section 6. First, a description  
 149 of each step is detailed.

150 **1. Clean state prediction.** In protein design applications, it is common to employ an  $\mathbf{x}_0$ -prediction  
 151 parameterization. This design is intentional, as it enables adaptation from pretrained folding  
 152 models like RoseTTFold, reusing their architectures and learned weights for initialization (Baek et al.,  
 153 2021). At reverse time  $t \in T, \dots, 0$  the model takes a noisy latent  $\mathbf{x}_t$  and predicts a clean structure,  
 154

$$155 \quad \mathbf{x}_{\theta}(\mathbf{x}_t, t) = \hat{\mathbf{x}}_0^t$$

156 providing an approximation of the final state. As  $t \rightarrow 0$ , predictions improve in accuracy as noise  
 157 signal reduces. The availability of the predictor  $\mathbf{x}_{\theta}$  is convenient as its output can be leveraged in  
 158 our next step to restore feasibility.

159 **2. Feasibility step (proximal projection).** Next, feasibility requirements are applied on the pre-  
 160 dicted clean state. A feasible estimate is produced through a proximal map:

$$161 \quad \tilde{\mathbf{x}}_0^t = \text{prox}_{\eta_t, g}(\hat{\mathbf{x}}_0^t) := \arg \min_{\mathbf{x}} \frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2 + g(\mathbf{x}) \quad (2)$$

162 where  $\eta_t > 0$  is a step size determined by the degree of trust in the denoiser’s prediction at step  $t$   
 163 and  $g : \mathbb{R}^{3 \times d} \rightarrow \mathbb{R}^1$  is a feasibility potential.  
 164

165 A natural first choice would be to take  $g$  as the indicator function of the feasible set. However, this  
 166 would require exact projections onto a potentially nonconvex set, which can be ill-defined when  $\hat{x}_0^t$   
 167 lies far from  $\mathcal{C}$ . To avoid instability, the hard indicator is replaced by its Moreau envelope (Boyd &  
 168 Vandenberghe, 2004), yielding the smooth penalty

$$169 \quad g(\mathbf{x}) = \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\mathbf{x})^2 \quad \text{with } \lambda_t > 0 \quad (3)$$

170 where  $\text{dist}_{\mathcal{C}}$  is a distance metric from the feasible set (e.g., in SE(3), it could be  $\inf_{\mathbf{y} \in \mathcal{C}} \|\mathbf{x} - \mathbf{y}\|$ ).  
 171 The parameter  $\lambda_t > 0$  plays the role of an inverse smoothing radius: as  $\lambda_t \rightarrow \infty$  the penalty enforces  
 172 exact feasibility, and for finite  $\lambda_t$  it softly biases toward  $\mathcal{C}$ .  
 173

174 **3. Forward renoising.** Having obtained a feasible estimate, the next step reintroduces noise by  
 175 sampling from the forward marginal at  $t - 1$  conditioned on the corrected clean sample  $\tilde{x}_0^t$ :

$$176 \quad \mathbf{x}_{t-1} = \text{FWD}(\tilde{x}_0^t, \varepsilon) = \sqrt{\bar{\alpha}_{t-1}} \tilde{x}_0^t + \sigma_{t-1} \varepsilon \quad (4)$$

177 with  $\varepsilon \sim \mathcal{N}(0, I)$ . This guarantees that, conditioned on  $\tilde{x}_0^t$ , the marginal of  $\mathbf{x}_{t-1}$  matches the  
 178 forward diffusion at time step  $t - 1$ . Note that as  $\sigma_t \rightarrow 0$  and  $\bar{\alpha}_t \rightarrow 1$ , the Markov chain terminates  
 179 at a clean  $\mathbf{x}_0 = \tilde{x}_0^t$ . If  $g$  acts as in indicator function, exact feasibility is recovered, while otherwise  
 180  $\mathbf{x}_0$  becomes arbitrarily close to  $\mathcal{C}$  as  $\lambda_t$  increases.  
 181

182 **Selecting the schedule.** It is instructive to connect the proximal subproblem (Equation (2)) to prob-  
 183 abilistic reasoning. By modeling the network’s clean error at step  $t$  as Gaussian with variance  $\eta_t I$ :

$$184 \quad p(\mathbf{x}_0 | \mathbf{x}_t) \propto \exp\left(-\frac{1}{2\eta_t} \|\mathbf{x}_0 - \hat{x}_0^t\|^2\right)$$

185 then interpreting the penalty  $g$  as a soft prior of the form  $\propto \exp\{-g(\mathbf{x}_0)\}$ , Equation (2) computes  
 186 the *per-step MAP estimate*  $\hat{x}_0^t$  of the clean state. The subsequent renoising step, Equation (4),  
 187 reinstates the correct stochasticity for the reverse chain while anchoring it to constraint set  $\mathcal{C}$ .  
 188

189 Because  $\sigma_t^2$  shrinks over time, it is natural to schedule  $\lambda_t$  to grow, so that feasibility becomes domi-  
 190 nant only when the model’s  $\hat{x}_0^t$  is accurate. Similarly, if  $\eta_t = \sigma_{t-1}^2$  the trust weight can be directly  
 191 connected to the diffusion variance, and the clean proximal problem remains on the same scale.  
 192

193 This predict-prox-renoise step is the stochastic analogue of a proximal gradient step. As elaborated  
 194 in Section 6, the predict-prox-renoise cycle both respects the diffusion dynamics and guarantees  
 195 convergence to feasible terminal states.  
 196

## 5 DECOUPLING GLOBAL TOPOLOGY FROM LOCAL GEOMETRY VIA ADMM

197 The constraint set  $\mathcal{C}$  captures a strong coupling between *local* stereochemical variables and *global*  
 198 variables governing topology and long-range residue interactions. Because residues that are far apart  
 199 in sequence may lie adjacent in the folded structure, enforcing global constraints thus necessitates  
 200 coordinated updates that can significantly impact nearby stereochemistry. For instance, we observe  
 201 that applying non-covalent bond constraints on a  $\beta$ -strand can cause the associated residues to shift  
 202 substantially, degrading the fidelity of this local geometry (Budyak et al., 2024). These interdepen-  
 203 dencies make the proximal step computationally complex. However, the presence of these separable  
 204 local and global constraints confers structure to the problem and thus presents an opportunity to  
 205 exploit it, enabling the use of decomposition approaches.  
 206

207 Consider that a feasible point can be equivalently represented as  $\mathbf{x} \in \mathcal{C}_{\text{local}} \cap \mathcal{C}_{\text{global}}$ . The local  
 208 constraints  $\mathcal{C}_{\text{local}}$  capture properties that are applicable in all backbone design tasks (e.g., adherence  
 209 to stereochemical bond lengths and angles between consecutive atoms and residues). The global  
 210 constraints  $\mathcal{C}_{\text{global}}$  are problem-specific functionals: for example, in our first experiment, this  
 211 constraint set defines bond lengths and angles between specific non-neighboring residues, characterizing  
 212 non-covalent bonds which are necessary for protein-ligand pocket design.  
 213

214 Following this intuition, the feasibility potential is decomposed as  $g(\mathbf{x}) = g_{\text{local}}(\mathbf{x}) + g_{\text{global}}(\mathbf{x})$ ,  
 215 where

$$g_{\text{local}}(\mathbf{x}) = \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}_{\text{local}}}(\mathbf{x})^2, \quad g_{\text{global}}(\mathbf{x}) = \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}_{\text{global}}}(\mathbf{x})^2$$

and  $g_{\text{local}}, g_{\text{global}} : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ . These functions are proximable; squared distance penalties are adopted, which are treated as indicator functions when  $\lambda_t \rightarrow \infty$ . Hence, our proximal update is reframed as:

$$\Pi(\hat{\mathbf{x}}_0^t) := \arg \min_{\mathbf{x}} \underbrace{\frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2}_{=:F(\mathbf{x})} + \underbrace{g_{\text{local}}(\mathbf{x}) + g_{\text{global}}(\mathbf{x})}_{=:G(\mathbf{x})}.$$

Crucially, we *define* the local block  $F$  to include the distance-to-denoiser term so that the local step both repairs stereochemistry and stays close to  $\hat{\mathbf{x}}_0^t$ , while the global block  $G$  focuses on long-range feasibility. We solve this by a consensus ADMM on

$$\min_{\mathbf{y}, \mathbf{z}} F(\mathbf{y}) + G(\mathbf{z}) \quad \text{s.t.} \quad \mathbf{y} = \mathbf{z}, \quad (5)$$

with scaled dual variable  $\mathbf{u}$  and penalty  $\rho > 0$ . This leads to the proximal splitting form of ADMM (Douglas–Rachford) (Parikh et al., 2014), with the update

$$\mathbf{y}^{k+1} := \text{prox}_{\rho^k, F}(\mathbf{y}^k - \mathbf{u}^k), \quad (6a)$$

$$\mathbf{z}^{k+1} := \text{prox}_{\rho^k, G}(\mathbf{z}^k + \mathbf{u}^k), \quad (6b)$$

$$\mathbf{u}^{k+1} := \mathbf{u}^k + \mathbf{y}^{k+1} - \mathbf{z}^{k+1} \quad (6c)$$

where  $k$  is an iteration counter and  $\mathbf{y}^0 = \mathbf{z}^0 = \hat{\mathbf{x}}_0^t \in \mathbb{R}^{3 \times d}$ . Here  $\mathbf{y}$  and  $\mathbf{z}$  are two copies of the backbone, associated with  $F$  and  $G$  respectively, and the dual variable  $\mathbf{u}$  accumulates their mismatch. At convergence the iterates satisfy  $\mathbf{y} = \mathbf{z}$ , recovering the minimizer of  $F + G$ ; in practice it is only necessary to take a single sweep per diffusion step, but warm-starting across steps ensures the two copies remain close.

These updates can, thus, be interpreted as applying ADMM to the consensus problem, where the dual variable  $\mathbf{u}$  carries forward residuals between local and global feasibility corrections. In practice, this is implemented by minimizing the associated augmented Lagrangian. For clarity, we present only the proximal form here but detail the explicit augmented Lagrangian in Appendix G.

## 6 THEORETICAL ANALYSIS

Now, we show that the samples generated by the proposed stochastic proximal method come with feasibility guarantees. In the following we assume that the constraint set  $\mathcal{C}$  is prox-regular and defer all proofs in Appendix H.

We start by providing a bound on the feasibility guarantees attained by the generated final sample.

**Theorem 6.1.** *Consider a feasibility potential  $g(\mathbf{x}) = \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\mathbf{x})^2$  defined as in Equation 2. Then, the proximal minimizer  $\tilde{\mathbf{x}}_0$  satisfies:*

$$\underbrace{\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0)}_{\text{feasibility}} \leq \frac{1}{\sqrt{2\lambda_t\eta_t}} \text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0) \quad (7)$$

The inequality shows that the proximal step contracts the violation by  $(2\lambda_t\eta_t)^{-1/2}$ , guiding the reverse process towards the constraint set. Then, as  $t \rightarrow 0$  and  $\lambda_t\eta_t \rightarrow \infty$ , the corrected iterate converges arbitrarily close to the constraint set.

The following result provides rationale for how to schedule  $\lambda_t$ .

**Theorem 6.2.** *Let  $K$  be a finite number such that  $\mathbb{E}[\text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_t, t)^2] \leq K$  for all  $t$ , then choosing  $\lambda_t = \frac{c_t}{\eta_t}$  with a non decreasing  $c_t$  yields:*

$$\mathbb{E}[\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2] \leq \frac{K}{2c_t} \quad \text{and hence} \quad \mathbb{E}[\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0)^2] \leq \frac{K}{2c_1} \quad (8)$$

As a consequence, tightening  $c_t$  towards the end guarantees decreasing expectation over the violations, leading to arbitrary small terminal violations. Note also that taking  $g_1$  as the identity function over the constraint set  $\mathcal{C}$  gives *exact feasibility*.

Beyond these quantitative feasibility bounds, the per-step proximal subproblem needs to be well-posed. In particular, the question becomes whether a (possibly local) minimizer of the proximal mapping exists under the modeling assumptions.

270 **Theorem 6.3.** Consider the proximal subproblem  
 271

272 
$$\text{prox}_{\eta_t, g}(\mathbf{x}) = \frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2 + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\mathbf{x})^2,$$
  
 273

274 with  $\eta_t, \lambda_t > 0$  and  $\mathcal{C} \subset \mathbb{R}^{3 \times n}$  nonempty and closed. Then:  
 275

276 1. **Existence.**  $\text{prox}_{\eta_t, g}$  is continuous and coercive, hence attains a global minimizer for every  $\hat{\mathbf{x}}_0^t$ .  
 277 In particular,  $\arg \min \text{prox}_{\eta_t, g} \neq \emptyset$ .  
 278 2. **Local uniqueness.** If  $\hat{\mathbf{x}}_0^t$  lies within the prox-regularity neighborhood of  $\mathcal{C}$ , then the projection  
 279  $\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$  is single-valued. Moreover, if  $\text{prox}_{\eta_t, g}$  is strongly convex in a neighborhood of  $\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$ ,  
 280 then the proximal minimizer is unique within that neighborhood.

281 The analysis shows that feasibility improves monotonically as  $\lambda_t$  tightens relative to  $\eta_t$ , while the  
 282 proximal subproblem remains well-posed due to the quadratic distance term. In practice, this pro-  
 283 vides principled guidelines for selecting schedules which balance denoiser trust against constraint  
 284 enforcement. The next section demonstrates that the theoretical foundations translate to tangible  
 285 performance improvements in protein backbone design.

287 **7 EXPERIMENTS**

290 We conduct our evaluation on two key tasks for protein motif scaffolding in the PDZ domain and  
 291 vacancy-constrained pocket design for unconditional generations, described in more details in Sec-  
 292 tions 7.1 and 7.2.

293 **Baselines.** The comparison includes state-of-the-art structure-based design methods incorporating  
 294 constraint conditioning. We use *RFDiffusion* (Watson et al., 2023) as the underlying backbone  
 295 across all baselines and our method, although we note that any method evaluated could be extended  
 296 to other backbone structure diffusion models (Yim et al., 2023; Cutting et al., 2025). *RFDiffusion* is  
 297 selected, as it's considered a state-of-the-art approach for the tasks considered in this work:

298 1. **Standard Diffusion:** *RFDiffusion* conditioned on relevant motifs or structures, as it is often used  
 299 in practice.  
 300 2. **Recentering of Mass Guidance:** *RFDiffusion* with conditioning to bias the generation towards  
 301 a particular Cartesian coordinate where bond interactions or cavity-defining residues enforce  
 302 constraint satisfying formations (Braun et al., 2024). Conditioning is fine-tuned to find the lowest  
 303 average constraint violation per sample prior to beginning experiments (see Appendix B).  
 304 3. **Constraint-Guided Diffusion (CGD):** *RFDiffusion* with constraint-based guidance Sequential  
 305 Monte Carlo sampling (Lee et al., 2025). This SOTA method uses importance sampling with a  
 306 guided rate matrix to define weights and periodic resampling based on constraint violations.

307 **Metrics.** The following metrics are adopted to evaluate the performance of our method:

308 1. **Constraint Satisfaction:** The percentage of samples satisfying all domain-specific constraints.  
 309 For Section 7.1, we verify presence of the described [hydrogen-bond pairing of the backbone](#)  
 310 through DSSP (Kabsch & Sander, 1983), while Section 7.2 computes constraints via Cartesian  
 311 coordinate bounds.  
 312 2. **Structure Realism:** Percentage of samples containing secondary structures within constrained  
 313 regions while maintaining general backbone realism (e.g.,  $\beta$ -sheets remain less than ten residues,  
 314 inter-residue distances and angles are preserved).  
 315 3. **Usable Percentage:** Percentage of samples passing above conditions for Structure Realism  
 316 and Constraint Violation. Indicates frequency of generating a structure that satisfies physical  
 317 plausibility and functional requirements.  
 318 4. **Radius of Gyration:** Average radius of gyration across generated backbones, measuring the  
 319 overall spatial compactness of the structure. Lower values typically correspond to more compact,  
 320 globular folds, while higher values indicate extended or unfolded conformations.  
 321 5. **Diversity:** Percentage of samples which are both useable and satisfy a minimum root mean  
 322 squared error between all other samples (2 Å). Higher diversity indicates better coverage of pos-  
 323 sible structure.

324 Additionally, we note that further details on all experimental setups are provided in Appendix B.

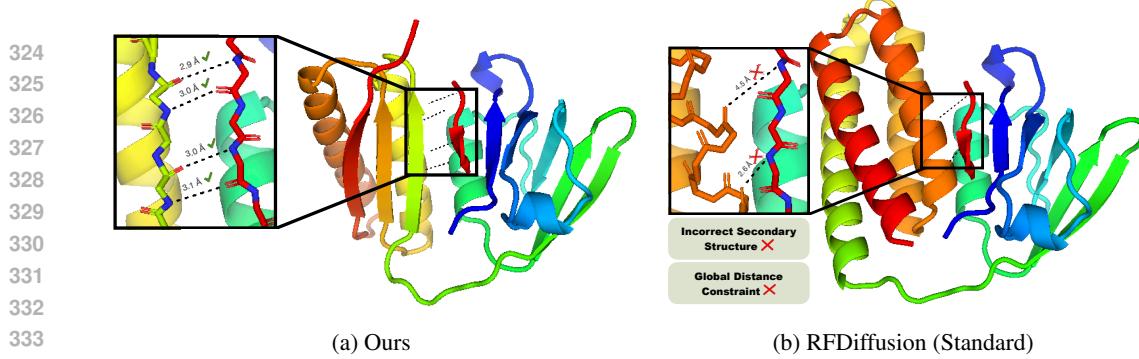


Figure 2: Visualization of randomly selected samples generated by (a) our proximal method and by (b) Standard RFDiffusion on our introduced PDZ domain benchmark.

### 7.1 PDZ DOMAIN: NON-COVALENT BONDS

Antibodies are widely used in molecular biology to target specific proteins, but their large size and extracellular restriction limit their utility, particularly in intracellular contexts. In addition, their binding interfaces, often mediated by flexible loop regions, pose challenges for computational design. As an alternative, biology employs small modular protein domains (e.g., WW, SH2, SH3, PH, and PDZ), which provide more designable binding modes. PDZ domains, for instance, recognize unstructured C-terminal motifs of partner proteins, typically through  $\beta$ -sheet-like hydrogen bond contacts. These interactions are generally weak and promiscuous, serving primarily in protein localization. To enhance affinity and specificity, Huang et al. (2009) engineered concatenated PDZ fusions with other small domains, demonstrating improved performance. Protein diffusion models extend this concept by enabling conditioning on both PDZ domains and their peptide ligands, thereby facilitating the *de novo* design of concatenated architectures.

**Dataset construction.** To benchmark constrained diffusion for PDZ engineering, we collected all resolved PDZ/PDZ binding motif (PBM) complexes from the RCSB Protein Data Bank (Berman et al., 2000). Seventy-two structures were initially retrieved and manually curated to remove entries with unresolved regions or peptides too short for recognition, yielding fifty-two usable complexes. **To expand the existing PDZ domain in a reliable way to make additional contacts with the target PBM, the N- and C-termini had to be rearranged.** Each structure was processed to reposition termini closer to the bound peptide by introducing a cut in a loop adjacent to the ligand, trimming the original termini, and applying vanilla RFDiffusion to in-paint the resulting gap. Candidate backbones were filtered to exclude chain breaks and non- $\beta$ -sheet pairings. Sequences for the redesigned regions were generated using ProteinMPNN, followed by structure prediction with AlphaFold2. Predicted models were retained only if they satisfied stringent criteria: (i) self-consistency RMSD to the RFDiffusion backbone  $<2.5$  Å, (ii) mean pLDDT  $>90$ , and (iii) peptide RMSD  $<2.0$  Å. After filtering, 31 high-confidence PDZ designs remained for benchmarking. **While all 31 structures are valid targets for this approach, six of them (PDB IDs: 2AWW, 2G2L, 4JOG, 6UBH, 6Y9O, 8CN3) contain short peptide ligands or peptides with geometrically restrictive residues such as prolines. These structures remain in the benchmark set, but they are distinguished in the results as having poor-posed ligands.** More details are provided in Appendix C. This benchmark for constrained diffusion is also a novel contribution of this work.

**Task description.** Given a target PDZ domain and its peptide ligand, the objective is to design an additional domain concatenated to the PDZ that contributes new stabilizing contacts with the peptide. This requires satisfying global inter-chain constraints. Specifically, because the peptide forms  $\beta$ -sheet-like contacts with the PDZ domain, we aim to create a complementary set of  $\beta$ -sheet-like interactions on the opposite face of the peptide using the engineered concatenated domain. Constraints are enforced to ensure valid bond lengths and angles, while also ensuring local stereochemistry is preserved. Success is evaluated by the metrics aforementioned, where global feasibility is defined by ideal bond lengths ( $2.9 \pm 0.2$  Å), C=O  $\cdots$  N angles ( $155 \pm 10^\circ$ ), and C<sub>α</sub>-N  $\cdots$  O ( $120 \pm 10^\circ$ ).

**Results.** Table 1 provides results on the PDZ benchmark comparing our constrained diffusion approach (visualized in Figure 2 (a)) to the baseline methods. Notably, across nearly one hundred thousand samples generated for the three baselines, not one sample perfectly satisfied the bonding distance and angle constraints. The baselines frequently generate incorrect secondary structures, as

	RFDiffusion			Ours
	Standard	Recenter	CGD	
Constraint Satisfaction (%) [ $\uparrow$ ]	0.0	0.0	0.0	<b>100.0</b>
Structure Realism (%) [ $\uparrow$ ]	(32.0)	(18.7)	(38.2)	<b>21.0</b>
Usable Percentage (%) [ $\uparrow$ ]	0.0	0.0	0.0	<b>21.0</b>
Radius of Gyration ( $\text{\AA}$ ) [ $\downarrow$ ]	(13.6)	(13.2)	(16.2)	<b>12.4</b>
Diversity (%) [ $\uparrow$ ]	N/A	N/A	N/A	<b>18.8</b>

Table 1: Comparison to structure-based design baselines, Standard (Watson et al., 2023), Recenter (Braun et al., 2024), and CGD (Lee et al., 2025), and ours, for the *PDZ domain*. Results reported across 31,000 samples for each baseline, highlighting **best** and second best results. Parentheses indicate when statistics are computed over **unusable** structures.

illustrated in Figure 2 (b), making it implausible that generations will effectively bind with the peptide ligand. While recentering and CGD perform well in terms of the local geometric requirements captured by the structure realism measurement, they are unable to cope with the global requirement of non-covalent bonding between the PDZ backbone continuation and the peptide ligand, as reflected by the constraint satisfaction rates. While constraint-guided diffusion satisfies the bond distance constraints for some generations, it is never able to generate residues which appropriately meet the angle requirements. Importantly, the *baselines yield no usable generations* for any of the 31 structures in the benchmark. In contrast, our method achieves state-of-the-art results, generating usable structures in **21.0%** of total generations (and up to **83.0%** for well-posed ligands), markedly outperforming the existing baselines. In addition to observing *perfect constraint satisfaction*, we outperform all methods substantially in radius of gyration and diversity metrics. These results highlight the unique ability of our approach to handle both local stereochemical properties while enforcing global functional constraints, providing a vastly more viable approach to protein engineering under specific property requirements and design constraints.

## 7.2 MOLECULE ENCAPSULATION: VACANCY CONSTRAINTS

Several recent approaches have improved the usefulness of diffusion-based protein generation with the integration of all-atom models and catalytic site scaffolding (Krishna et al., 2024; Ahern et al., 2025; Braun et al., 2024). A key remaining limitation is precise spatial control over where new structure is placed. Adjusting the diffusion origin can help, but explicit user control over inclusion/exclusion volumes would better enable tasks such as shaping small-molecule pockets, peptide-binding grooves, or membrane-embedded features. For example, Braun et al. (2024) approximated pocket formation by inserting a placeholder  $\alpha$ -helix to occupy volume during generation and deleting it afterward. Generalizing this idea to geometric volume constraints, hard inclusion masks and forbidden regions, could provide finer control and higher success rates in targeted protein design.

**Task description.** Given a fixed spatial environment defined by a rectangular box with an internal conical exclusion zone, the goal is to design protein backbones which fall exclusively in this nonconvex region while preserving the local geometries and secondary structures. Feasible structures are characterized by all atoms falling within the defined box ( $20 \text{\AA} \times 40 \text{\AA} \times 40 \text{\AA}$ ), while simultaneously avoiding the exclusion zone introduced by the displacement (visualized in Figure 3).

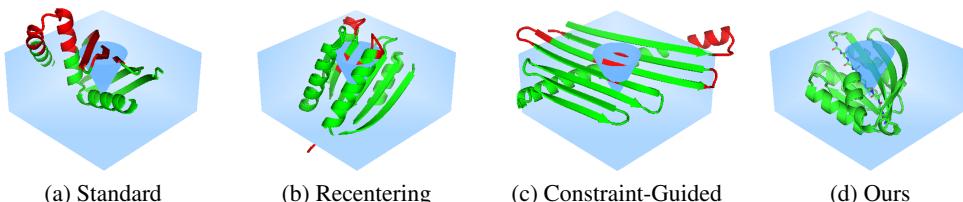


Figure 3: Visualization of randomly selected samples for the molecule encapsulation experiment; green parts of the structure fall within **feasible regions**, while the red parts **violate the constraints**.

**Results.** Figure 3 provides a visualization of representative samples from each baseline. Standard diffusion and recentering of mass guidance perform similarly in this domain, yielding realistic struc-

		RFDiffusion			
	Standard	Recenter	CGD	Recenter + CGD	Ours
Constraint Satisfaction (%) $\uparrow\uparrow$	0.0	0.0	21.6	<u>27.4</u>	<b>100.0</b>
Structure Realism (%) $\uparrow\uparrow$	(100.0)	(100.0)	96.1	<u>93.8</u>	<b>97.8</b>
Usable Percentage (%) $\uparrow\uparrow$	0.0	0.0	20.5	<u>24.2</u>	<b>97.8</b>
Radius of Gyration (Å) $\downarrow\downarrow$	(15.2)	(14.3)	<u>23.9</u>	<u>26.6</u>	<b>14.8</b>
Diversity (%) $\uparrow\uparrow$	N/A	N/A	20.5	<u>24.2</u>	<b>97.8</b>

Table 2: Comparison to structure-based design baselines: Standard (Watson et al., 2023), Recenter (Braun et al., 2024), and CGD (Lee et al., 2025), and ours, for *molecule encapsulation*. Results reported across 4000 samples, highlighting **best** and second best results. Parentheses indicate when statistics are computed over **unusable** structures.

tures but failing to satisfy the functional requirements in any out of 1000 samples generated for each. We observe slightly different failure modes: standard diffusion generally violates the box constraint, while recentering more often violates the vacancy constraint. The recentering is effective at keeping the structures inside the box, but it cannot capture the exclusion zone. **Table 2 reports much stronger performance for constraint-guided diffusion, especially when augmented with recentering, which generates feasible samples 27.4% of the time, resulting in 24.2% usable samples.** However, it is worth noting that the qualitative performance suffers, as the radius of gyration is much higher than other baselines. This is often indicative of structures which contain unfolded conformations, ultimately undermining structural stability and realism. In comparison, our method reports *perfect constraint satisfaction*, producing **4× as many usable samples** as the nearest baseline, with an impressive **97.8%** success rate. Furthermore, it maintains radii of gyration comparable to standard diffusion, indicating the generated samples combine structural plausibility, compactness, and fold coverage.

## 8 RELATED WORK

While existing *de novo* protein structure design models produce plausible generations, as we have shown in this paper, sampled protein backbones frequently violate inter-atomic bond lengths, angles, or chain closure requirements, often necessitating the generation of tens of thousands of candidates to obtain a handful of viable designs (e.g., Watson et al. 2023; Sappington et al. 2024). Although backbone generators such as RFDiffusion (Watson et al., 2023; Ahern et al., 2025) provide major advances in functional conditioning, outputs still require post hoc filtering to ensure stereochemical correctness, and current pipelines continue to rely heavily on rejection sampling. **Recent models such as Genie 2 and OriginFlow further illustrate the growing interest in diffusion and flow-based approaches for backbone generation (Lin et al., 2024; Yan et al., 2025); however, these methods exhibit similar limitations, with generated backbones often violating global constraints and requiring substantial downstream filtering.**

Training-time methods have been proposed to address these issues by embedding structural constraints into generative models (Eguchi et al., 2022; Lutz et al., 2023). However, because protein design requires task-specific constraints, a model trained on one constraint set does not generalize, making broad applicability impractical without retraining. Furthermore, training-time approaches typically provide only distributional guarantees, biasing samples on average rather than ensuring per-sample feasibility. Models such as ReQFlow (Yue et al., 2025) and FoldFlow-2 (Huguet et al., 2024) provide valuable tools, but likewise do not directly enforce hard geometric constraints, instead incorporating them as soft biases.

Inference-time approaches have emerged as a strong alternative, enforcing per-sample compliance and removing the need for model retraining. Diffusion guidance methods were first introduced for soft constraint imposition but are fundamentally limited, offering only probabilistic bias rather than guaranteed adherence to the constraint set (Ho & Salimans, 2022). While these techniques have improved performance for protein backbone generation, with models such as Chroma (Ingraham et al., 2023) leveraging conditioning on specific substructures, as we have shown in our experiments, they are often ineffective in providing consistent constraint satisfaction. Overcoming these limitations requires generative models which can effectively integrate these constraints into the design process, as presented in this work.

---

**486 9 CONCLUSION**  
487

488 Motivated by the significant challenge of integrating functional design constraints into protein en-  
489 gineering tasks, this paper present a constrained diffusion framework for structure-guided design.  
490 By applying proximal feasibility updates with ADMM decomposition, the approach couples local  
491 stereochemical property enforcement with global utility requirements. To assess the quality of ex-  
492 isting solutions as compared to the methodology presented in this paper, the work introduces a novel  
493 curated benchmark for protein motif scaffolding in PDZ domains, providing the first standard for  
494 constrained diffusion methods in modular domain engineering. Evaluation reports state-of-the-art  
495 results across motif scaffolding and vacancy-constrained pocket design, illustrating the ability of  
496 this approach to generate high quality proteins which adhere to precise domain-centric constraints.  
497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540  
541  
**ETHICS STATEMENT**542  
543  
544  
545  
546  
547  
This work develops methods for constrained generative modeling in protein design. Our method im-  
proves feasibility in backbone design tasks, which holds significant potential to accelerate existing  
protein engineering pipelines. To mitigate risks of potential misuse, this paper restricts evaluation to  
safe, publicly available structural benchmarks, curated from the Protein Data Bank. All implemen-  
tations follow standard open science practices, ensuring safe and transparent research548  
549  
**REPRODUCIBILITY STATEMENT**550  
551  
552  
553  
554  
The code and benchmark are released in the supplementary material along with instructions to guide  
the reproduction of the results presented in this paper. Methodological details are described exten-  
sively in the paper and accompanying appendix. Section 7, Appendix B, Appendix F describe the  
specific constraints, hyperparameters, models, and hardware used for evaluation pipelines.555  
556  
**REFERENCES**557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
Woody Ahern, Jason Yim, Doug Tischer, Saman Salike, Seth M Woodbury, Donghyo Kim, Indrek  
Kalvet, Yakov Kipnis, Brian Coventry, Han Raut Altae-Tran, et al. Atom level enzyme active site  
scaffolding using rfdiffusion2. *bioRxiv*, pp. 2025–04, 2025.  
Minkyung Baek, Frank DiMaio, Ivan Anishchenko, Justas Dauparas, Sergey Ovchinnikov, Gyu Rie  
Lee, Jue Wang, Qian Cong, Lisa N Kinch, R Dustin Schaeffer, et al. Accurate prediction of  
protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–  
876, 2021.  
Noémie Bergues, Arthur Carré, Paul Join-Lambert, Brice Hoffmann, Arnaud Blondel, and Hamza  
Tajmouati. Template-guided 3d molecular pose generation via flow matching and differentiable  
optimization. *arXiv preprint arXiv:2506.06305*, 2025.  
Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig,  
Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1):  
235–242, 2000.  
Matthieu Blanke, Yongquan Qu, Sara Shamekh, and Pierre Gentine. Strictly constrained generative  
modeling via split augmented langevin sampling. *arXiv preprint arXiv:2505.18017*, 2025.  
Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.  
Markus Braun, Adrian Tripp, Morakot Chakatok, Sigrid Kaltenbrunner, Celina Fischer, David Stoll,  
Aleksandar Bijelic, Wael Elaily, Massimo G Totaro, Melanie Moser, et al. Computational enzyme  
design by catalytic motif scaffolding. *bioRxiv*, pp. 2024–08, 2024.  
Ivan L Budyak, Lihua Huang, and Rina K Dukor. Higher order structure characterization of two in-  
terdomain disulfide bond variants of a bispecific monoclonal antibody. *Journal of Pharmaceutical  
Sciences*, 113(8):2651–2655, 2024.  
Sidhartha Chaudhury, Sergey Lyskov, and Jeffrey J Gray. Pyrosetta: a script-based interface for im-  
plementing molecular modeling algorithms using rosetta. *Bioinformatics*, 26(5):689–691, 2010.  
Jacob K Christopher, Stephen Baek, and Nando Fioretto. Constrained synthesis with projected  
diffusion models. *Advances in Neural Information Processing Systems*, 37:89307–89333, 2024.  
Daniel Cutting, Frédéric A Dreyer, David Errington, Constantin Schneider, and Charlotte M Deane.  
De novo antibody design with se (3) diffusion. *Journal of Computational Biology*, 32(4):351–361,  
2025.  
Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances  
in neural information processing systems*, 34:8780–8794, 2021.

594 Kieran Didi, Francisco Vargas, Simon V Mathis, Vincent Dutordoir, Emile Mathieu, Urszula J Ko-  
 595 morowska, and Pietro Lio. A framework for conditional diffusion modelling with applications in  
 596 motif scaffolding for protein design. *arXiv preprint arXiv:2312.09236*, 2023.

597

598 Raphael R Eguchi, Christian A Choe, and Po-Ssu Huang. Ig-vae: Generative modeling of protein  
 599 structure by direct 3d coordinate generation. *PLoS computational biology*, 18(6):e1010271, 2022.

600 Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanassee,  
 601 Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete  
 602 diffusion. *Advances in neural information processing systems*, 36:12489–12517, 2023.

603

604 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*  
 605 *arXiv:2207.12598*, 2022.

606

607 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
 608 *neural information processing systems*, 33:6840–6851, 2020.

609 Jin Huang, Koki Makabe, Matthew Biancalana, Akiko Koide, and Shohei Koide. Structural basis  
 610 for exquisite specificity of affinity clamps, synthetic binding proteins generated through directed  
 611 domain-interface evolution. *Journal of molecular biology*, 392(5):1221–1231, 2009.

612

613 Guillaume Huguet, James Vuckovic, Kilian Fatras, Eric Thibodeau-Laufer, Pablo Lemos, Riashat  
 614 Islam, Chenghao Liu, Jarrid Rector-Brooks, Tara Akhoud-Sadegh, Michael Bronstein, et al.  
 615 Sequence-augmented se (3)-flow matching for conditional protein generation. *Advances in neural*  
 616 *information processing systems*, 37:33007–33036, 2024.

617

618 John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent  
 619 Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein  
 620 space with a programmable generative model. *Nature*, 623(7989):1070–1078, 2023.

621

622 Wolfgang Kabsch and Christian Sander. Dictionary of protein secondary structure: pattern recog-  
 623 nition of hydrogen-bonded and geometrical features. *Biopolymers: Original Research on*  
*Biomolecules*, 22(12):2577–2637, 1983.

624

625 Rohith Krishna, Jue Wang, Woody Ahern, Pascal Sturmels, Preetham Venkatesh, Indrek Kalvet,  
 626 Gyu Rie Lee, Felix S Morey-Burrows, Ivan Anishchenko, Ian R Humphreys, et al. Generalized  
 627 biomolecular modeling and design with rosettafold all-atom. *Science*, 384(6693):eadl2528, 2024.

628

629 Cheuk Kit Lee, Paul Jeha, Jes Frellsen, Pietro Lio, Michael Samuel Albergo, and Francisco Var-  
 630 gas. Debiasing guidance for discrete diffusion with sequential monte carlo. *arXiv preprint*  
*arXiv:2502.06079*, 2025.

631

632 Yeqing Lin, Minji Lee, Zhao Zhang, and Mohammed AlQuraishi. Out of many, one: Design-  
 633 ing and scaffolding proteins at the scale of the structural universe with genie 2. *arXiv preprint*  
*arXiv:2405.15489*, 2024.

634

635 Isaac D Lutz, Shunzhi Wang, Christoffer Norn, Alexis Courbet, Andrew J Borst, Yan Ting Zhao,  
 636 Annie Dosey, Longxing Cao, Jinwei Xu, Elizabeth M Leaf, et al. Top-down design of protein  
 637 architectures with reinforcement learning. *Science*, 380(6642):266–273, 2023.

638

639 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.  
 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

640

641 Panos M Pardalos and Stephen A Vavasis. Quadratic programming with one negative eigenvalue is  
 642 np-hard. *Journal of Global optimization*, 1(1):15–22, 1991.

643

644 Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and trends® in Optimization*,  
 1(3):127–239, 2014.

645

646 Isaac Sappington, Martin Toul, David S Lee, Stephanie A Robinson, Inna Goreshnik, Clara Mc-  
 647 Curdy, Tung Ching Chan, Nic Buchholz, Buwei Huang, Dionne Vafeados, et al. Improved protein  
 648 binder design using beta-pairing targeted rfdiffusion. *bioRxiv*, pp. 2024–10, 2024.

648 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 649 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*  
 650 *arXiv:2011.13456*, 2020.

651

652 Zhenqiao Song, Yunlong Zhao, Wenxian Shi, Wengong Jin, Yang Yang, and Lei Li. Generative en-  
 653 zyme design guided by functionally important sites and small-molecule substrates. *arXiv preprint*  
 654 *arXiv:2405.08205*, 2024.

655

656 Brian L Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and  
 657 Tommi Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-  
 658 scaffolding problem. *arXiv preprint arXiv:2206.04119*, 2022.

659

660 Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eise-  
 661 nach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of  
 662 protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

663

664 Junyu Yan, Zibo Cui, Wenqing Yan, Yuhang Chen, Mengchen Pu, Shuai Li, and Sheng Ye. Robust  
 665 and reliable de novo protein design: A flow-matching-based protein generative model achieves  
 666 remarkably high success rates. *bioRxiv*, pp. 2025–04, 2025.

667

668 Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay,  
 669 and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation.  
 670 *arXiv preprint arXiv:2302.02277*, 2023.

671

672 Angxiao Yue, Zichong Wang, and Hongteng Xu. Reqflow: rectified quaternion flow for efficient  
 673 and high-quality protein backbone generation. *arXiv preprint arXiv:2502.14637*, 2025.

674

## A LIMITATIONS AND FUTURE WORK

675 In scientific domains such as protein engineering, exact constraint satisfaction is essential for both  
 676 physical realism (local constraints) and functional utility (global constraints). However, it is valua-  
 677 ble to acknowledge several important trade-offs associated with the adoption of constrained meth-  
 678 ods. First, enforcing feasibility inherently introduces increased computational complexity. Sampling  
 679 from  $p_C$  is a fundamentally harder task than sampling from  $p_{\text{data}}$ , especially under the assumption of  
 680 hard constraint satisfaction, as reflected in reported runtimes (Appendix F). While runtime is worth  
 681 considering, especially as the method scales, it is important to note that when sampling from  $p_C$ , our  
 682 approach offers the most efficient runtime, providing much better performance than the current re-  
 683 jection sampling paradigm. Although not applicable in our experimental settings, if sampling from  
 684  $p_{\text{data}}$  is sufficient, it is not computationally tractable to adopt constrained diffusion methods.

685 As second caveat concerns theoretical assumptions. Our analysis assumes prox-regularity of the  
 686 constraint set to provide formal guarantees. This condition is standard in proximal theory, but it may  
 687 not apply for highly nonconvex constraints encountered in protein design tasks. In such cases, con-  
 688 vergence guarantees do not explicitly apply, and solutions are instead justified empirically. However,  
 689 the development of formal guarantees for nonconvex constraint sets remains an open and largely un-  
 690 explored direction, as existing optimization theory cannot provide guarantees for general nonconvex  
 691 sets.

692 Finally, while our experimental evaluation illustrates state-of-the-art performance on two highly  
 693 nontrivial and practically significant protein engineering problems, it necessarily leaves many other  
 694 design challenges open for future exploration. The problems we consider already capture central  
 695 difficulties in enforcing exact feasibility under realistic structural and functional constraints, making  
 696 them representative of some of the hardest settings encountered in practice. Extensions such as plac-  
 697 ing global constraints on higher-level properties (e.g., polarity of specific regions in the structure)  
 698 or imposing specific formation constraints on secondary structures (e.g., controlling the radius of a  
 699  $\beta$ -barrel) are complimentary next steps. We view these as exciting opportunities for future work,  
 700 with the present study establishing a critical foundation for handling such broader classes of design  
 701 challenges.

Task	$u^0$	$c_T$	$c_1$	$T$	Diffusion schedule
PDZ Domain	0.0	finite ( $\approx$ tol. 3.0 Å)	$\infty$	45	linear in $t$ , from $\bar{\alpha}_0 \approx 1$ to $\bar{\alpha}_T \approx \prod_{t=0}^{45} (1 - 7e^{-2})$ ; SO(3): 1.5 → 2.5; $\sigma \in [0.02, 1.5]$
Molecule Encapsulation	0.0	$\infty$	$\infty$	50	linear in $t$ , from $\bar{\alpha}_0 \approx 1$ to $\bar{\alpha}_T \approx \prod_{t=0}^{50} (1 - 7e^{-2})$ ; SO(3): 1.5 → 2.5; $\sigma \in [0.02, 1.5]$

Table 3: Experiment hyperparameters.

## B EXPERIMENTAL DETAILS

All hyperparameters are reported in Table 3. Both tasks use RFDiffusion defaults:  $T = 50$  steps, linear  $\beta$  schedule from  $10^{-2}$  to 0.07, SO(3) schedule 1.5–2.5, coordinate scaling 0.25, and Gaussian noise  $\sigma \in [0.02, 1.5]$ . Proximal multipliers  $c_t$  tighten from  $c_T$  (finite) to  $c_1 = \infty$  across the trajectory. Other hyperparameters  $\lambda$  and  $\eta$  can be derived by Theorem 6.2.

### B.1 PDZ DOMAIN: NON-COVALENT BONDS

The evaluation is conducted across our newly introduced dataset, including 31 distinct PDZ structures. For each structure, we generate 1000 samples with each of the baselines, yielding a total of 31,000 designed proteins. From these generations, we assess the overall performance of each method, leveraging PyRosetta to conduct the final evaluation of the generations (Chaudhury et al., 2010).

**Recenter of Mass Guidance.** Prior to running the evaluation, we search the local space surrounding the peptide ligand to empirically optimize the selected center of mass. We generate ten samples for each selected point in space, providing representative data for the performance on different conditioning. Selection is determined based on the lowest average constraint violation (capturing angles and bond lengths between the peptide ligand and the sampled generations). After completing this search, this center of mass is used for all runs on the particular motif.

**Constraint-Guided Diffusion.** Adopting the Sequential Monte Carlo conditioning proposed by Lee et al. (2025), we set the sample weighting via the constraint violation of the clean predicted state  $\hat{x}_0^t$ . For the evaluation, we fix the number of particles  $P = 200$ , as this seems to balance potential performance and overall runtime. We adopt a multinomial resampling function and introduce an inverse temperature parameter,  $\beta$ , which controls the sharpness of the resampling distribution. Candidate weights are computed as

$$w_i \propto \exp\left(-\beta (c_i - \min_j s_j)\right)$$

where  $c_i$  denotes the score of constraint violation  $i$ . Earlier in the diffusion process,  $\beta = 30$  remains high, keeping the resampling fairly stochastic, but we lower this  $\beta = 1$  as  $t \rightarrow 0$  to improve selection for the final state and increase the likelihood of non-covalent bonds forming.

**Constrained Diffusion (Ours).** As described in Section 5, our projection is implemented through an ADMM derived decomposition. We assign  $\mathbf{y}$  to be consecutive residues which fall between the start and end of the bonding points on the peptide ligand. For instance, if the peptide has six covalent bonding sites,  $\mathbf{y}$  is composed of five residues. In this case, O–N and N–O bonds are placed on the first, third, and fifth residues. Then,  $\mathbf{z}$  is composed of the linker preceding  $\mathbf{y}$  (typically six residues), and the chain continuation following  $\mathbf{y}$  (typically 60–100 residues). Additionally, we note that our approach leverages constraint guidance for this experiment, as our approach is complementary to these guidance methods.

### B.2 MOLECULE ENCAPSULATION: VACANCY CONSTRAINTS

For this setting, we evaluate across 1000 samples for each of the baselines. We use a fixed box of size  $20 \text{ \AA} \times 40 \text{ \AA} \times 40 \text{ \AA}$ , centering the apex of the cone  $5 \text{ \AA}$  above the bottom face of the box, with a half-angle of  $25^\circ$ . As we identify that many *baseline* samples violate only the width constraints of the box ( $40 \text{ \AA} \times 40 \text{ \AA}$ ), which is not integral to the application, we do not report these violations in our evaluation provided in Figure 2.

756 **Recenter of Mass Guidance.** The center of mass is positioned at the center of the box, encouraging  
 757 the sample to (1) remain inside of the box and (2) maximize contact with the cone shaped  
 758 vacancy. Similar to the previous setting, the guidance is tuned on a small set of sample prior to  
 759 generating outputs for evaluation; however, for this experiment the center of mass is fixed during  
 760 tuning, only searching over the *guidance strength* parameter.

761 **Constraint-Guided Diffusion.** We adopt an identical setup for the constraint guidance as de-  
 762 scribed for the PDZ domain, but leveraging an identical constraint violation measure:

$$g(x) = \inf_{y \in \mathcal{C}} \|x - y\|^2$$

763 which captures the sum of the distance of each atom from the feasible region. We similarly adopt  
 764 a particle count of  $P = 200$ , and follow an identical schedule to the previous experiment for the  
 765 inverse temperature.

766 **Constrained Diffusion (Ours).** In this setting, all atomic positions are subject to the global con-  
 767 straints, making  $y = x$ , removing the need for explicit variable splitting in this setting. While we  
 768 note that the vacancy constraint is indeed nonconvex, the projection operator can be represented in  
 769 closed form, as it is composed of simple convex shapes; this alleviates concerns surrounding the  
 770 tractability of the projection operator, which is the primary motivation for decoupling global and  
 771 local requirements. To preserve local geometric properties and secondary structures, we extend  
 772 this operator by enforcing rigid-body consistency on secondary structure segments, while adjoining  
 773 linker regions are adjusted to absorb residual displacements. This construction ensures that the pro-  
 774 jection both satisfies global constraints and maintains local stereochemical integrity. Additionally,  
 775 we note that our approach leverages constraint guidance and recentering of mass for this experiment,  
 776 as our approach is complementary to these guidance methods.

777 **Additional Results.** To supplement the results reported in the paper, we include additional anal-  
 778 ysis of the secondary structure preservation rates for samples generated by each baseline and by  
 779 our method. This metric quantifies the percentage of residues assigned to a canonical secondary  
 780 structure state in the relaxed models. Higher values indicate improved preservation of fold-level  
 781 features.

	RFDiffusion			Ours
	Standard	Recenter	CGD	
Secondary Structure (%)	80.8	82.5	83.5	<b>85.0</b>

782 Table 4: Secondary structure preservation across design methods.

783 In addition to the primary metrics reported in the main text, we assess the structural plausibility of  
 784 generated designs using a secondary structure preservation metric. This measure is assessed using  
 785 PyRosetta to determine the percentage of residues which fall in helix, sheet, or turn conformations in  
 786 the relaxed structures (Chaudhury et al., 2010). As shown in Table 4, our method achieves the high-  
 787 est secondary structure frequency, suggesting that the model maintains global fold characteristics  
 788 more reliably than the baselines.

## C DATASET CONSTRUCTION

801 Details for the filtering are included in Figure 4. For every candidate structure per target PDZ-PBM  
 802 loop closure, AlphaFold models were constructed. To determine the confidence of each model,  
 803 predicted Local Distance Difference Test (pLDDT) was collect. The cutoff for pLDDT for a viable  
 804 loop closure was set to  $>90$  (top left). The former N- and C- termini form a  $\beta$ -sheet pairing, to  
 805 conserve this, a filter of  $>60\%$  of the residues being  $\beta$ -sheet was used (top right). ProteinMPNN  
 806 was used to generate potential sequences that conform to generated backbones for the in-painted  
 807 region. After having AlphaFold predict the structure given the new sequence, RMSD was measured  
 808 between the generated backbone and the predicted structure. A cutoff of  $<2.5\text{\AA}$  was used (bottom  
 809 left). To ensure the peptide was not interfered with after predicting the new structure with AlphaFold,  
 a stringent cutoff of  $<1\text{\AA}$  RMSD was used (bottom right)

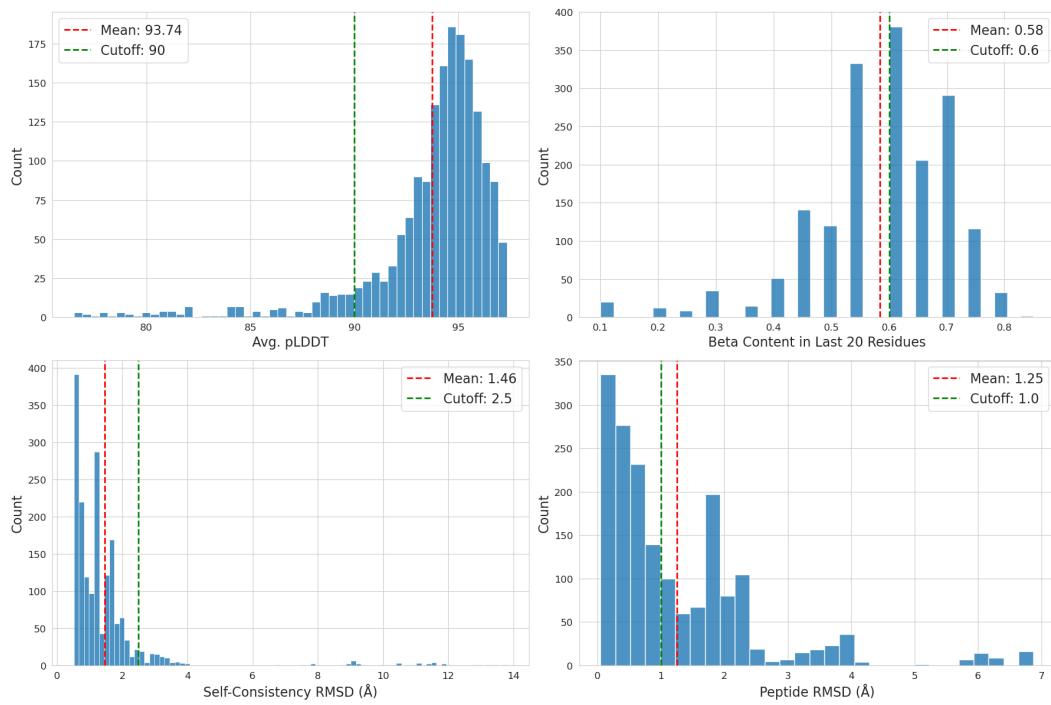


Figure 4: Dataset filtering details expanded.

## D COMPARISON TO PROJECTED DIFFUSION MODELS

As discussed in Section 3, prior work on constrained generative sampling has proposed Projected Diffusion Models (Christopher et al., 2024). These models augment the sampling process by imposing a projection onto the feasible set after each diffusion step. While theoretical guarantees have been derived for convex constraint sets, and empirical observations have illustrated that these properties hold for certain nonconvex sets, the paper has argued critical limitations for step-then-project approaches when applied to challenging settings as explored in this work. This section presents empirical justification of these aforementioned claims.

	<i>Projected Diffusion Model</i>	<i>Ours</i>
	Local + Global	Global
Constraint Satisfaction (%)	0.0 (Failed to Converge)	<b>100.0</b>
Structure Realism (%)	N/A	0.0
Usable Percentage (%)	0.0	0.0
Radius of Gyration (Å)	N/A	(13.1)
Diversity (%)	N/A	N/A

Table 5: Comparison to Projected Diffusion Models on the Molecule Encapsulation setting.

The results in Table 5 reaffirm the limitations of integrating intermediate projection steps when working with high-dimensional data subject to complex nonconvex constraints. While in this setting it was possible to exactly satisfy the global constraints, the resulting structures are unusable, with high concentrations of atoms placed on the edges of the box and cone boundaries. When also enforcing the local constraints explicitly, the projection operator never converged on any samples or timesteps.

864 

## E HYPERPARAMETER SENSITIVITY ANALYSIS

866 To assess the robustness of our method to hyperparameter choices, we performed two ablation studies:  
 867 (1) varying constraint weight schedule  $\lambda$ , and (2) varying the number of ADMM iterations per  
 868 diffusion step. All experiments were conducted on the PDZ domain design target 6ubh, a well-posed  
 869 example from our benchmark set suited to constrained generation.

	$\lambda = \frac{10}{t}$	$\lambda = \frac{100}{t}$	$\lambda = \frac{1000}{t}$	$\lambda = \text{inf}$
Constraint Satisfaction (%)	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>
Structure Realism (%)	80.0	<b>100.0</b>	<b>100.0</b>	90.0
Usable Percentage (%)	80.0	<b>100.0</b>	<b>100.0</b>	90.0
Radius of Gyration (Å)	10.58	10.53	10.56	<b>10.34</b>
Diversity (%)	80.0	<b>100.0</b>	<b>100.0</b>	90.0

879 Table 6: *Ablation over constraint-weight schedules  $\lambda$  on PDZ domain sample 6ubh.*

880 As illustrated by Table 6, increasing the initial weight of  $\lambda_T$  yields improvements in structural real-  
 881 ism and usable sample percentage. This trend aligns with the interpretation that stricter constraint  
 882 enforcement at early timesteps guides the diffusion trajectory toward the feasible manifold, improv-  
 883 ing geometric quality. However, increasing this too strictly, results in over constraining early states,  
 884 analogous to placing an overly large weight on the diffusion-guidance terms as discussed in Section  
 885 3.

	ADMM = 1	ADMM = 2	ADMM = 3
Constraint Satisfaction (%)	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>
Structure Realism (%)	<b>80.0</b>	70.0	40.7
Usable Percentage (%)	<b>80.0</b>	70.0	40.7
Radius of Gyration (Å)	<b>10.58</b>	10.72	11.24
Diversity (%)	<b>80.0</b>	70.0	40.7

894 Table 7: *Ablation on the number of ADMM iterations per diffusion step.*

895 The next analysis centers on determining the optimal number of ADMM iterations per diffuson  
 896 step. As described in Section 5, a single ADMM sweep per diffusion step is sufficient in practice.  
 897 In fact, the ablation suggests that multiple iterations degrade quality. The failure case that emerges  
 898 over additional ADMM iterations is actually always tied to the formation of a  $\beta$ -sheet across the  
 899 binding sites; additional iterations do not result in local stereochemical properties breaking, but the  
 900 analysis suggests that they do lower the liklihood of secondary structure conformations. This is  
 901 because a higher wieght is placed on the ADMM optimization than the diffusion updates, leading to  
 902 feasible generations but not ones which fall closely on the data manifold. Furthermore, we note that  
 903 additional ADMM steps also become impractical because adding these iterations scales the runtime  
 904 nearly linearly.

905 

## F RUNTIME

906 Benchmarks were run on a single slice of an NVIDIA A100 Multi-Instance GPU (MIG) to effi-  
 907 ciently parallelize sampling across available hardware. We posit that overall runtimes could be  
 908 improved by scaling the available resources, but this is beyond the scope of this study; for instance,  
 909 both constraint-guided diffusion and our approach could be significantly parallelized to increase  
 910 efficiency, as they operate over up to  $P = 200$  particles. Note that reducing  $P$  will accelerate per-  
 911 formance, providing a tunable trade-off between speed and performance. Hence, we expect that  
 912 if resource constraints could be ignored, these methods could be accelerated by multiple orders of  
 913 magnitude.

914 Our method provides the best runtimes *per usable sample*. In protein design experiments, this is a  
 915 more important metric than raw runtime, as only usable samples are relevant when considering over-  
 916 all performance. Hence, under the assumption that functional requirements are necessary, which is

	Standard	Recenter	CGD	<i>Ours</i>
<i>PDZ Domain</i>				
Time Per Usable Sample (min)	$\infty$	$\infty$	$\infty$	<b>935.2</b>
<i>Molecule Encapsulation</i>				
Time Per Usable Sample (min)	$\infty$	$\infty$	142.4	<b>98.4</b>

Table 8: Reported runtimes on both settings. Analysis considers the average time per usable sample.

specifically the applications this work targets, **the proposed method is indeed the most computationally efficient!**

## G ADMM

This section expands the discussion of the ADMM formulation introduced in Section 5. Begin by recalling the consensus formulation of our proximal update:

$$\text{prox}_{F+G}(\mathbf{y}, \mathbf{z}) := \arg \min_{\mathbf{y}, \mathbf{z}} F(\mathbf{y}) + G(\mathbf{z}) \quad \text{s.t. } \mathbf{y} = \mathbf{z}$$

where  $F$  includes the data anchor term and local feasibility requirements, and  $G$  encodes global feasibility.

$$\begin{aligned} F(\mathbf{y}) &= \frac{1}{2\eta_t} \|\mathbf{y} - \hat{\mathbf{x}}_0^t\|^2 + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}_{\text{local}}}(\mathbf{y})^2 \\ G(\mathbf{z}) &= \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}_{\text{global}}}(\mathbf{z})^2 \end{aligned}$$

Recall that the consensus initialization  $\mathbf{y}^0 = \mathbf{z}^0 = \hat{\mathbf{x}}_0^t \in \mathbb{R}^{3 \times n}$ , which aligns the start of the iterates with the denoiser's prediction.

Classical ADMM results (see Parikh et al. 2014) guarantee convergence to a minimizer of  $F + G$  under convexity assumptions. In our case, convexity does not strictly hold, but in practice this algorithm is often effectively applied to nonconvex settings, achieving feasible solutions.

The formulation yields the augmented Lagrangian, introducing the dual variable  $\mathbf{u}$  and a penalty parameter  $\rho$ :

$$\mathcal{L}_\rho(\mathbf{y}, \mathbf{z}, \mathbf{u}) = F(\mathbf{y}) + G(\mathbf{z}) + \mathbf{u}^\top (\mathbf{y} - \mathbf{z}) + \frac{\rho}{2} \|\mathbf{y} - \mathbf{z}\|^2$$

Then, expanding the ADMM updates introduced in Equation (6) yields:

$$\begin{aligned} \mathbf{y}^{k+1} &= \arg \min_{\mathbf{y}} F(\mathbf{y}) + \frac{\rho^k}{2} \|\mathbf{y} - \mathbf{z}^k + \mathbf{u}^k\|^2 \\ \mathbf{z}^{k+1} &= \arg \min_{\mathbf{z}} G(\mathbf{z}) + \frac{\rho^k}{2} \|\mathbf{y}^{k+1} - \mathbf{z} + \mathbf{u}^k\|^2 \\ \mathbf{u}^{k+1} &= \mathbf{u}^k + \mathbf{y}^{k+1} - \mathbf{z}^{k+1} \end{aligned}$$

These updates are equivalent to the proximal splitting form presented in Section 5, with  $\text{prox}_F$  realized by the  $\mathbf{y}$ -update and  $\text{prox}_G$  by the  $\mathbf{z}$ -update. The dual variable  $\mathbf{u}$  accumulates the residual mismatch between the two copies, driving them toward consensus. At convergence,  $\mathbf{y} = \mathbf{z}$ , recovering the minimizer of  $F + G$ .

## H COMPLETE PROOFS

**Proof of Theorem 6.1.** *Claim:*  $\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t) \leq \frac{1}{\sqrt{\lambda_t \eta_t}} \text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$

**Lemma H.1.** *Let  $\phi(\mathbf{z}) = \frac{1}{2\eta_t} \|\mathbf{z} - \hat{\mathbf{x}}_0^t\|^2 + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\mathbf{z})^2$ . By optimality of  $\hat{\mathbf{x}}_0^t$ ,  $\phi(\tilde{\mathbf{x}}_0^t) \leq \phi(\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t))$ .*

Lemma H.1 holds as the minimizer's objective value can't exceed the value at any feasible point, in particular at the projection of  $\hat{\mathbf{x}}_0^t$ .

972 Then, expanding from the lemma:  
 973

$$\frac{1}{2\eta_t} \|\tilde{\mathbf{x}}_0^t - \hat{\mathbf{x}}_0^t\| + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2 \leq \frac{1}{2\eta_t} \|\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t) - \hat{\mathbf{x}}_0^t\| + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t))^2$$

974 Since  $\frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t))^2 = 0$ , this term can then be omitted.  
 975

$$\frac{1}{2\eta_t} \|\tilde{\mathbf{x}}_0^t - \hat{\mathbf{x}}_0^t\| + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2 \leq \frac{1}{2\eta_t} \|\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t) - \hat{\mathbf{x}}_0^t\|$$

976 Dropping the non-negative first term gives:  
 977

$$\frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2 \leq \frac{1}{2\eta_t} \|\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t) - \hat{\mathbf{x}}_0^t\|$$

978 or equivalently  
 979

$$\left[ \text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t) \leq \frac{1}{\sqrt{\lambda_t \eta_t}} \text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0^t) \right]$$

980  $\square$   
 981

982 **Proof of Theorem 6.2.** *Claim:*  $\mathbb{E}[\text{dist}_{\mathcal{C}}(\mathbf{x}_0)^2] \leq \frac{K}{2c_1}$   
 983

984 Begin by assuming  $\mathbb{E}[\text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)^2] \leq K$ , and the schedule is defined such that  $\lambda_t = \frac{c_t}{\eta_t}$ .  
 985

986 By Theorem 6.1, it is known that  
 987

$$\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t) \leq \frac{1}{\sqrt{2\lambda_t \eta_t}} \text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$$

988 and thus, on the squared expectation  
 989

$$\mathbb{E}[\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2] \leq \frac{1}{2\lambda_t \eta_t} \mathbb{E}[\text{dist}_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)^2].$$

990 Then, by our assumption,  
 991

$$\mathbb{E}[\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2] \leq \frac{1}{2\lambda_t \eta_t} K = \frac{K}{2c_t}$$

992 Substituting our schedule  $\lambda_t = \frac{c_t}{\eta_t}$  yields:  
 993

$$\mathbb{E}[\text{dist}_{\mathcal{C}}(\tilde{\mathbf{x}}_0^t)^2] \leq \frac{K}{2c_t}$$

994 By applying Corollary 6.2 at time  $t = 1$ :  
 995

$$\left[ \mathbb{E}[\text{dist}_{\mathcal{C}}(\mathbf{x}_0)^2] \leq \frac{K}{2c_1} \right]$$

996 Thus, if the schedule of  $c_t$  is decreasing, the expected violation converges as  $t \rightarrow 0$ .  
 997

998  $\square$

1000 **Proof of Theorem 6.3.** *Claim:* Existence of a local solution for the proximal mapping.  
 1001

1002 *Existence:*  $\text{prox}_{\eta_t, g}$  is continuous and coercive, hence attains a global minimizer for every  $\hat{\mathbf{x}}_0^t$ . In  
 1003 particular,  $\arg \min \text{prox}_{\eta_t, g} \neq \emptyset$ .

1004 First, assume  $\eta_t, \lambda_t > 0$ ,  $\mathcal{C} \subset \mathbb{R}^{3 \times n}$  nonempty, closed. For ease of notation:  
 1005

$$\Phi_t(\mathbf{x}) = \frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2 + \frac{\lambda_t}{2} \text{dist}_{\mathcal{C}}(\mathbf{x})^2$$

1006 Since the distance from the constraint set is strictly non-negative,  
 1007

$$\Phi_t(\mathbf{x}) \geq \frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2 = \frac{1}{2\eta_t} \|\mathbf{x}\|^2 + \frac{1}{2\eta_t} \langle \mathbf{x}, \hat{\mathbf{x}}_0^t \rangle + \frac{1}{2\eta_t} \|\hat{\mathbf{x}}_0^t\|^2.$$

1008 Since the term  $\|\mathbf{x}\| \rightarrow \infty$  as  $\mathbf{x} \rightarrow \infty$ ,  $\Phi$  is  $\Phi$ .  
 1009

1010  $\Phi$  is also lower semicontinuous since (i) the quadratic anchor term  $\frac{1}{2\eta_t} \|\mathbf{x} - \hat{\mathbf{x}}_0^t\|^2$  is continuous, and  
 1011 (ii) the squared distance to the closed set is lower semicontinuous and finite everywhere.  
 1012

1026 Since  $\Phi$  is coercive and lower semicontinuous, it obtains a global minimum.  
 1027

$$\left[ \arg \min \Phi_t \neq \emptyset \right]$$

1031 *Local uniqueness.* If  $\hat{\mathbf{x}}_0^t$  lies within the prox-regularity neighborhood of  $\mathcal{C}$ , then the projection  
 1032  $\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$  is single-valued. Moreover, if  $\text{prox}_{\eta_t, g}$  is strongly convex in a neighborhood of  $\Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$ ,  
 1033 then the proximal minimizer is unique within that neighborhood.

1034 Now, assume  $\mathcal{C}$  is prox-regular at all points in a neighborhood  $\mathcal{X}$  of  $\mathbf{x}^* = \Pi_{\mathcal{C}}(\hat{\mathbf{x}}_0^t)$ . Prox-regularity  
 1035 then implies:

1036

- 1037 • **Single valued projection.** The projection  $\Pi_{\mathcal{C}}(\cdot)$  is single valued and Lipschitz on the  
 neighborhood  $\mathcal{U} \subset \mathcal{X}$ .
- 1039 • **Smooth distance function.** Thus,  $\text{dist}_{\mathcal{C}}(\cdot)^2$  is  $C^1$  on  $\mathcal{U}$  with gradient:

$$\nabla \left( \frac{1}{2} \text{dist}_{\mathcal{C}}(\mathbf{x})^2 \right) = \mathbf{x} - \Pi_{\mathcal{C}}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{U}.$$

1043 Let  $\mathcal{U}$  be small enough that the above holds and  $\hat{\mathbf{x}}_0^t \in \mathcal{U}$ . Consider  $\Phi_t$  restricted to  $\mathcal{U}$ . Its gradient is  
 1044

$$\nabla \Phi_t(\mathbf{x}) = \frac{1}{\eta_t}(\mathbf{x} - \hat{\mathbf{x}}_0^t) + \lambda_t(\mathbf{x} - \Pi_{\mathcal{C}}(\mathbf{x})), \quad \mathbf{x} \in \mathcal{U},$$

1045 and its (generalized) Hessian can be bounded below using the Jacobian of  $\Pi_{\mathcal{C}}$ . In particular, since  
 1046  $\Pi_{\mathcal{C}}$  is nonexpansive in  $\mathcal{U}$ ,

$$\langle \nabla^2 \Phi_t(\mathbf{x}) v, v \rangle \geq \frac{1}{\eta_t} \|v\|^2 + \lambda_t \langle (I - D\Pi_{\mathcal{C}}(\mathbf{x}))v, v \rangle \geq \frac{1}{\eta_t} \|v\|^2,$$

1053 for all  $v$  and almost every  $\mathbf{x} \in \mathcal{U}$ . Thus  $\Phi_t$  is locally  $\mu$ -strongly convex on  $\mathcal{U}$  with  $\mu = 1/\eta_t > 0$ .  
 1054

1055 Local strong convexity implies a unique critical point in  $\mathcal{U}$ , hence a unique local minimizer of  $\Phi_t$   
 1056 in  $\mathcal{U}$ . **Since  $\Phi_t$  is coercive with at least one global minimizer, and  $\hat{\mathbf{x}}_0^t$  is in the prox-regularity  
 1057 neighborhood so the minimizer lies in  $\mathcal{U}$ , the proximal minimizer is unique in that neighbor-  
 1058 hood.**

1059  $\square$

1060  
 1061  
 1062  
 1063  
 1064  
 1065  
 1066  
 1067  
 1068  
 1069  
 1070  
 1071  
 1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079