

# SESAME: OPENING THE DOOR TO PROTEIN POCKETS

Raúl Miñán<sup>1\*</sup>, Carles Perez-Lopez<sup>2</sup>, Javier Iglesias-Fernandez<sup>3</sup>, Álvaro Ciudad<sup>1,4\*</sup>, Alexis Molina<sup>1,4\*</sup>

<sup>1</sup>Department of Artificial Intelligence, Nostrum Biodiscovery, Barcelona, Spain

<sup>2</sup>Department of Drug Discovery, Nostrum Biodiscovery, Barcelona, Spain

<sup>3</sup>Department of R&D, Nostrum Biodiscovery, Barcelona, Spain

<sup>4</sup>Atlas Labs, Barcelona, Spain

\*{raul.minan, alvaro.ciudad, alexis.molina}@nostrumbiodiscovery.com

## ABSTRACT

Molecular docking is a cornerstone of drug discovery, relying on high-resolution ligand-bound structures to achieve accurate predictions. However, obtaining these structures is often costly and time-intensive, limiting their availability. In contrast, ligand-free structures are more accessible but suffer from reduced docking performance due to pocket geometries being less suited for ligand accommodation in apo structures. Traditional methods for artificially inducing these conformations, such as molecular dynamics simulations, are computationally expensive. In this work, we introduce Sesame, a generative model designed to predict this conformational change efficiently. By generating geometries better suited for ligand accommodation at a fraction of the computational cost, Sesame aims to provide a scalable solution for improving virtual screening workflows.

## 1 INTRODUCTION

Within the realm of drug discovery, molecular docking serves as a pivotal methodology for exploring and characterizing the intricate interactions between molecular entities and their corresponding protein binding cavities (Meng et al., 2011). Traditionally, high-resolution ligand-bound, or holo structures, are needed as the starting point for this computational workflow (Bender et al., 2021). However, they might not be always available, as they are produced by costly experimental protocols (Turnbull & Emsley, 2013).

Ligand-free, or apo structures, are more readily available. They can be derived from experimental methods, similar to those used for holo structures, or predicted employing models like AlphaFold2 (Jumper et al., 2021). Despite their availability, when used as input for virtual screening campaigns, they suffer from lower performance compared to holo structures, both in traditional docking (Díaz-Rovira et al., 2023) and deep learning approaches (Corso et al., 2023). This performance drop is attributed to pocket geometries being better suited for ligand accommodation in holo structures. This problem is particularly relevant when transient binding sites, such as allosteric and cryptic pockets, are of interest. These unique pockets are binding sites that only become apparent under an induced conformational change, e.g. drug binding, providing a promising alternative to classical binding sites for drug development (Basciu et al., 2019). Moreover, they can be involved in mechanisms that allow for additional specificity in modulating certain diseases (Hollingsworth et al., 2019).

Traditional approaches for generating these holo-like geometries consist in computationally expensive molecular dynamics (MD) protocols that can take up to several GPU-days (Sabanés Zariquiey et al., 2022). In this work we present Sesame, a generative model that can generate these desired geometries for a fraction of the cost by predicting the apo-holo conformational changes in a ligand-agnostic way.

## 2 RELATED WORK

**Traditional Molecular Dynamics Approaches.** MD simulations are computational methods that calculate atomic and molecular interactions and physical movements over time by solving Newton’s

equations of motion. These simulations can be used to characterize the formation of cryptic pockets through numerous approaches: standard MDs with organic probe molecules to measure binding propensity of fragments (Guvench & MacKerell Jr, 2009), mixed-solvent MDs to characterize hydrophobic regions (Sabanés Zariquiey et al., 2022) and enhanced sampling methods to accelerate the sampling of the conformations of interest (Oleinikovas et al., 2016; Basciu et al., 2019).

**Deep Learning Approaches.** Molecular docking has been a field of interest in deep learning research during the last few years. The seminal work of DiffDock (Corso et al., 2023) introduced generative models to docking, achieving state-of-the-art performance. However, this work had already observed a lower performance in apo predicted structures compared to holo conformations. This finding prompted multiple studies aimed at generating improved protein conformations to use as inputs in computational workflows (Bose et al., 2024; Yim et al., 2024) or at directly modeling the joint distribution of protein and ligand conformations (Abramson et al., 2024).

Closer to our work are ApolloDiff<sup>1</sup> (Zhang et al., 2023) and SBAlign (Somnath et al., 2023). ApolloDiff focuses on the generation of holo conformations using an equivariant diffusion model conditioned on the apo conformation and sequence. SBAlign, meanwhile, formulates an alternative approach to solving the Schrödinger Bridge (Léonard, 2013) problem under the assumption that samples from the source and target distributions are available and paired and recovers a stochastic trajectory between them.

### 3 METHODS

In this work we seek to learn a mapping between two explicit data distributions: the source distribution of apo structures, and the target distribution of holo structures. In this context, flow matching (FM) (Lipman et al., 2023) is a generative modeling paradigm which provides a simulation-free way of learning continuous normalizing flows (CNFs) that generate data over a learned vector field (Chen et al., 2018). Given two distributions  $q_0$  and  $q_1$ , one can then learn the CNF  $\psi_t(x)$  that transports  $q_0$  to  $q_1$  by optimizing the conditional flow matching (CFM) objective:

$$\mathcal{L}_{\text{CFM}} := \mathbb{E}_{t, x_0 \sim q_0, x_1 \sim q_1} \left[ \|v_t(x) - u_t(x|x_1)\|^2 \right] \quad (1)$$

where  $u_t(x|x_1)$  is the conditional vector field associated with the interpolation between  $x_0$  and  $x_1$  and  $v_t(x)$  can be parameterized with a neural network.

Previous works have extended FM to Riemannian Manifolds (Chen & Lipman, 2024) and generalized it to use arbitrary source and target distributions (Tong et al., 2023; Albergo et al., 2023), postulating FM as a suitable framework for a generative model of holo structures given apo samples and allowing us to take advantage of several novel developments in the protein structure generation field (Bose et al., 2024; Jing et al., 2024).

Namely, we follow FoldFlow and parameterize backbones as SE(3)-equivariant frames that represent rigid transformations  $T = (r, x) \in SE(3)$  that consist of a rotation  $r \in SO(3)$  and a translation  $x \in \mathbb{R}^3$  (Jumper et al., 2021). This formulation allows us to decompose the CFM process independently in  $SO(3)$  and  $\mathbb{R}^3$ . In each space, we can define conditional flows using the geodesic (2) and linear (3) interpolants, respectively, and use their associated vector fields to optimize the following simplified objectives:

$$\mathcal{L}_{SO(3)}(\theta) = \mathbb{E}_{t, q(r_0, r_1)} \|v_\theta(t, r_t) - u_t(r_t|r_0, r_1)\|_{SO(3)}^2, \quad \text{with } r_t = \exp_{r_0}(t \log_{r_0}(r_1)) \quad (2)$$

$$\mathcal{L}_{\mathbb{R}^3}(\theta) = \mathbb{E}_{t, q(x_0, x_1)} \|v_\theta(t, x_t) - u_t(x_t|x_0, x_1)\|^2, \quad \text{with } x_t = (1-t)x_0 + tx_1 \quad (3)$$

We optimize these objectives with the loss function used by FoldFlow, and additionally find increased performance with the addition of the Frame Aligned Point Error loss from AlphaFold2 (Jumper et al., 2021). Moreover, we also include the auxiliary losses from FrameFlow (Yim et al.,

<sup>1</sup>We omit ApolloDiff as a baseline in our experiments due to the lack of weights and code availability.

2023), which encourage a better representation of the backbone atoms. Therefore, the final loss is formulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{\text{FOLDFLOW}} + \mathcal{L}_{\text{FAPE}} + \mathcal{L}_{\text{aux}} \quad (4)$$

Further details on FM and the used loss functions can be found in Appendix F.

## 4 EXPERIMENTS

### 4.1 PREDICTING LARGE CONFORMATIONAL CHANGES

We first evaluate our method in the dataset of overall protein motions of the D3PM database (Peng et al., 2022), a dataset comprised of apo-holo pairs where the Root Mean Squared Deviation (RMSD) of the  $C_{\alpha}$  carbon atoms between bound and unbound states is  $> 2.0\text{\AA}$ , which we call D3PM-Large (see Appendix A for details).

In Table 1 we report the statistics of the RMSD between the  $C_{\alpha}$  carbon atoms of the generated structures and the holo conformations. Additionally, we also compute  $\Delta RMSD = RMSD_{\text{apo-holo}} - RMSD_{\text{gen-holo}}$  to assess whether the generated structures are closer to the holo conformation than the respective apo one. Here, positive values are indicative of good performance (Zhang et al., 2023).

Table 1: Conformational changes results in the D3PM-Large Test Set. RMSD between generated conformations and true holo structures. Results marked with an asterisk (\*) were obtained from Somnath et al. (2023).

Methods	RMSD ( $\text{\AA}$ ) $\downarrow$			$\Delta$ RMSD ( $\text{\AA}$ ) $\uparrow$			% RMSD ( $\text{\AA}$ ) $< \tau$ $\uparrow$		
	Med.	Mean	Std	Med.	Mean	Std	$\tau = 2$	$\tau = 5$	$\tau = 10$
EGNN*	19.99	21.37	8.21	—	—	—	1%	1%	3%
SBAlign	3.67	4.82	3.93	1.30	1.92	2.59	0%	71%	93%
Sesame	2.87	3.65	2.95	2.15	3.11	4.26	38%	82%	96%

Sesame outperforms previous baselines, achieving 38% of predictions with an RMSD  $< 2.0\text{\AA}$  with respect to the reference holo structure. Moreover, it also showcases better performance in generating structures that are more similar to the holo conformations than their apo counterparts, aligning more closely with our objectives. This distinction is crucial, since a generated conformation may be similar to the holo structure yet still remain closer to the apo state.

While our model demonstrates strong performance in modeling large conformational changes, most ligand-induced conformational changes occur on a smaller scale, with 88% of the cases having a RMSD between corresponding  $C_{\alpha}$  atoms of less than  $2\text{\AA}$ , and with 75% of the cases where it falls below  $1\text{\AA}$  (Frimurer et al., 2003).

### 4.2 PREDICTING POCKET CONFORMATIONAL CHANGES

The task of predicting small conformational changes can prove to be more challenging due to the even greater scarcity of high-quality data, as these subtle changes are more difficult to capture (Turnbull & Emsley, 2013). Although the D3PM database also has a dataset for pocket motions where the RMSD between  $C_{\alpha}$  atoms is lower than  $2\text{\AA}$  (see Appendix A.1.2), D3PM-Pocket only contains 865 apo-holo pairs, which represents a challenge to build generalizable predictive models. To overcome this, we construct a dataset of apo-holo pairs by generating apo conformations from holo structures after collapsing pockets using MD simulations. Further details in Appendix A.2.

For a more complete comparison with previous methods, we also retrain SBAlign on this new dataset. More details on the retraining procedure can be found in Appendix B. We additionally use the D3PM-Pocket dataset to evaluate the performance of both models in an established benchmark.

Table 2: Conformational changes results in the PDBBind-MD Test Set. RMSD between generated conformations and true holo structures. SBAlign was retrained in the PDBBind-MD dataset and inference was performed for the best model hyperparameters.

Methods	RMSD (Å) ↓			ΔRMSD (Å) ↑			% RMSD (Å) < $\tau$ ↑		
	Med.	Mean	Std	Med.	Mean	Std	$\tau = 0.15$	$\tau = 0.2$	$\tau = 0.5$
SBAlign	0.52	0.55	0.18	0.05	0.03	0.11	0 %	0 %	24.46 %
Sesame	0.18	0.18	0.03	0.39	0.39	0.03	7.57 %	82.16 %	100%

Table 3: Conformational changes results in the D3PM-Pocket set. RMSD between generated conformations and true holo structures. SBAlign was retrained in the PDBBind-MD dataset and inference was performed for the best model hyperparameters.

Methods	RMSD (Å) ↓			ΔRMSD (Å) ↑			% RMSD (Å) < $\tau$ ↑		
	Med.	Mean	Std	Med.	Mean	Std	$\tau = 0.15$	$\tau = 0.2$	$\tau = 0.5$
SBAlign	0.77	1.21	1.08	0.001	0.001	0.002	1.16 %	2.33 %	18.60 %
Sesame	0.49	0.91	1.22	0.18	0.19	0.1	5.59 %	13.61 %	50.67 %

As we can observe on Tables 2 and 3, we outperform the existing baseline in  $\Delta$ RMSD both in the PDBBind-MD and D3PM-Pocket datasets, highlighting the potential of Sesame in capturing small conformational changes.

Summary statistics for each dataset are provided in Appendix A.4 to better contextualize the reported results. Additionally, we report results for cross-inferences, where we evaluate the models trained on large conformational changes on small ones and vice versa to assess generalization capabilities across different data distributions.

### 4.3 SESAME FACILITATES CRYPTIC POCKET IDENTIFICATION AND CHARACTERIZATION

Due to the increasing interest in drug discovery in targeting cryptic and transient pockets, multiple algorithms for identifying and characterizing them have been designed. One example is PocketMiner, a recent deep learning method which predicts the probability of a residue belonging to a cryptic site (Meller et al., 2023). Testing Sesame alongside pocket prediction algorithms can help assess a potential synergy, since Sesame is trained to generate proteins closer to holo states from apo structures. This could improve the prediction of cryptic and transient pockets and allow for better identification of the residues involved in them. Additionally, if the relevant residues are detected more effectively, it would further confirm that the generated conformations are more holo-like.

To evaluate how Sesame could aid in this endeavor, we generate holo-like structures for the validation and test sets used in Meller et al. (2023), a set of known cryptic pockets, labeling as positive those residues closer than 8 Å to the ligand. We then run PocketMiner on the apo, holo and predicted conformations and obtain the precision-recall and receiver operating characteristic (ROC) curves for each.

In Figure 1, we can observe an increase in performance in PocketMiner predictions on Sesame’s generated structures versus their apo counterparts. This can indicate that the generated conformations are more geometrically similar to holo structures, increasing the reliability of pocket detection algorithms. For further discussion and visualizations of results obtained see Appendix D.

### 4.4 SESAME GENERATES VIABLE STRUCTURES FOR MOLECULAR DOCKING

In this section we seek to evaluate the performance of Sesame generated structures using them as input for a standard docking protocol. A common way of measuring a docking software’s performance is reproducing a crystallographically resolved structure, with the degree of geometric similarity serving as said measurement of the performance (Salmaso et al., 2018). In this context, we can evaluate

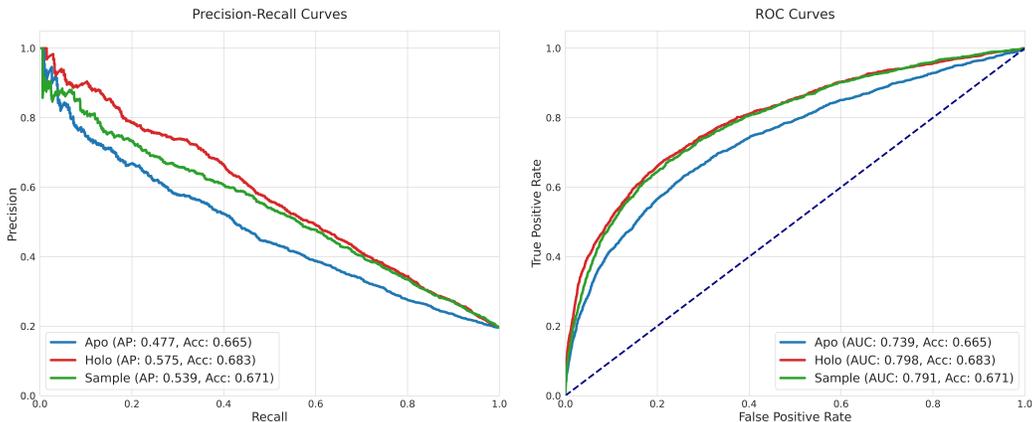


Figure 1: Precision-Recall (left) and ROC (right) curves for PocketMiner predictions.

the ability of Sesame generated structures to accept compound geometries that resemble those in the crystal holo conformations.

It is important to note that Sesame employs only backbone atoms, meaning side chains are not explicitly determined. Since most docking software requires side chains as input, and they are vital for determining the ligand’s correct position, their positions were predicted using Prime (Jacobson et al., 2004). To ensure a fair comparison, the side chains of both holo and apo structures were also reconstructed following the same protocol and included as baselines. Docking performance was then evaluated using RMSD, with all calculations performed against the ligand structure in the holo X-ray.

The two PDB complexes were chosen based on the availability of both apo and holo structures with high resolution and their clinical significance (Wu et al., 2022; Hiron et al., 2010).

Table 4: RMSD and corresponding docking poses for the Holo versus Reconstructed Holo, Apo and Model. (PDBs: 4ZZI, 4XKQ)

	RMSD (Å) ↓	Docking Pose
Holo X-ray vs Reconstructed Holo	1.622	Figure 4.A
Holo X-ray vs Reconstructed Apo	5.704	Figure 4.B
Holo X-ray vs Model	3.724	Figure 4.C

In the results corresponding to the complex 4ZZI in Table 4, we obtain an increased performance when compared to the apo structure, indicating that the obtained backbone structure is better suited for docking protocols. The low RMSD in the X-ray versus reconstruction comparison confirms the validity of the side chain prediction protocol. Further discussion and additional results of the 4LVT complex can be found in Appendix E.2.

For full details of the corresponding docking methods, protocol and visualization of the obtained poses that were used for the RMSD calculations see Appendix E.1 and Appendix E.3, respectively.

## 5 CONCLUSIONS AND FUTURE WORK

In this work, we introduce Sesame, a novel generative model that leverages the flow matching framework to construct a generative process between two data distributions, learning how to transform apo-like protein conformations into holo-like ones. Our model surpasses existing baselines in capturing both large and small conformational changes. Furthermore, we demonstrate that the generated conformations aid in identifying cryptic binding pockets and serve as effective input structures for molecular docking, leading to improved results compared to their original apo counterparts. These

findings indicate that Sesame offers a scalable and efficient approach for enhancing virtual screening workflows and streamlining the drug discovery pipeline.

Future work aims to address the main limitations of the model. Namely, this includes extending the framework to incorporate side chain modeling, as they play a fundamental role in protein-ligand interactions, binding affinity and docking accuracy. Additionally, we plan to expand our data generation efforts using molecular dynamics simulations that capture transitions across energetic barriers, thereby increasing the diversity of targetable binding pockets and providing richer datasets that encompass both subtle and large-scale movements and allow for an end-to-end model capable of accurately modeling both.

## REFERENCES

- Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, Alexander I. Cowen-Rivers, Andrew Cowie, Michael Figurnov, Fabian B. Fuchs, Hannah Gladman, Rishub Jain, Yousuf A. Khan, Caroline M. R. Low, Kuba Perlin, Anna Potapenko, Pascal Savy, Sukhdeep Singh, Adrian Stecula, Ashok Thillaisundaram, Catherine Tong, Sergei Yakneen, Ellen D. Zhong, Michal Zielinski, Augustin Židek, Victor Bapst, Pushmeet Kohli, Max Jaderberg, Demis Hassabis, and John M. Jumper. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, 2024.
- Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- Delphine C Bas, David M Rogers, and Jan H Jensen. Very fast prediction and rationalization of pKa values for protein-ligand complexes. *Proteins: Structure, Function, and Bioinformatics*, 73(3):765–783, 2008.
- Andrea Basciu, Giuliano Malloci, Fabio Pietrucci, Alexandre MJJ Bonvin, and Attilio V Vargiu. Holo-like and druggable protein conformations from enhanced sampling of binding pocket volume and shape. *Journal of Chemical Information and Modeling*, 59(4):1515–1528, 2019.
- Brian J Bender, Stefan Gahbauer, Andreas Lutten, Jiankun Lyu, Chase M Webb, Reed M Stein, Elissa A Fink, Trent E Balius, Jens Carlsson, John J Irwin, et al. A practical guide to large-scale docking. *Nature Protocols*, 16(10):4799–4832, 2021.
- Joey Bose, Tara Akhound-Sadegh, Guillaume Hugué, Kilian FATRAS, Jarrid Rector-Brooks, Cheng-Hao Liu, Andrei Cristian Nica, Maksym Korablyov, Michael M Bronstein, and Alexander Tong. SE(3)-stochastic flow matching for protein backbone generation. In *The Twelfth International Conference on Learning Representations*, 2024.
- David A Case, Hasan Metin Aktulga, Kellon Belfon, David S Cerutti, G Andrés Cisneros, Vinícius Wilian D Cruzeiro, Negin Forouzesh, Timothy J Giese, Andreas W Götz, Holger Gohlke, et al. AmberTools. *Journal of Chemical Information and Modeling*, 63(20):6183–6191, 2023.
- Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. *arXiv preprint arXiv:2302.03660*, 2024.
- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi S. Jaakkola. DiffDock: Diffusion steps, twists, and turns for molecular docking. In *The Eleventh International Conference on Learning Representations*, 2023.
- Warren L DeLano. Pymol: An open-source molecular graphics tool. *CCP4 Newsl. Protein Crystallogr.*, 40(1):82–92, 2002.

- Anna M Díaz-Rovira, Helena Martín, Thijs Beuming, Lucía Díaz, Victor Guallar, and Soumya S Ray. Are deep learning structural models sufficiently accurate for virtual screening? Application of docking algorithms to AlphaFold2 predicted structures. *Journal of Chemical Information and Modeling*, 63(6):1668–1674, 2023.
- R. A. Engh and R. Huber. *Structure quality and target parameters*, chapter 18.3, pp. 474–484. John Wiley & Sons, Ltd, 2012.
- Thomas M Frimurer, Günther H Peters, Lars F Iversen, Henrik S Andersen, Niels Peter H Møller, and Ole H Olsen. Ligand-induced conformational changes: improved predictions of ligand binding conformations and affinities. *Biophysical journal*, 84(4):2273–2281, 2003.
- Olgun Guvench and Alexander D MacKerell Jr. Computational fragment-based binding site identification by ligand competitive saturation. *PLoS computational biology*, 5(7):e1000435, 2009.
- Thomas A Halgren, Robert B Murphy, Richard A Friesner, Hege S Beard, Leah L Frye, W Thomas Pollard, and Jay L Banks. Glide: a new approach for rapid, accurate docking and scoring. 2. enrichment factors in database screening. *Journal of medicinal chemistry*, 47(7):1750–1759, 2004.
- Aurelia Hiron, Brunella Posteraro, Marie Carrière, Laetitia Remy, Cécile Delporte, Marilena La Sorda, Maurizio Sanguinetti, Vincent Juillard, and Elise Borezée-Durant. A nickel abc-transporter of staphylococcus aureus is involved in urinary tract infection. *Molecular microbiology*, 77(5):1246–1260, 2010.
- Scott A Hollingsworth, Brendan Kelly, Celine Valant, Jordan Arthur Michaelis, Olivia Mastromihalis, Geoff Thompson, AJ Venkatakrishnan, Samuel Hertig, Peter J Scammells, Patrick M Sexton, et al. Cryptic pocket formation underlies allosteric modulator selectivity at muscarinic gpcrs. *Nature Communications*, 10(1):3289, 2019.
- Matthew P Jacobson, David L Pincus, Chaya S Rapp, Tyler JF Day, Barry Honig, David E Shaw, and Richard A Friesner. A hierarchical approach to all-atom protein loop prediction. *Proteins: Structure, Function, and Bioinformatics*, 55(2):351–367, 2004.
- Bowen Jing, Bonnie Berger, and Tommi Jaakkola. AlphaFold meets flow matching for generating protein ensembles. *arXiv preprint arXiv:2402.04845*, 2024.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.
- Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976.
- Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- Jianing Li, Robert Abel, Kai Zhu, Yixiang Cao, Suwen Zhao, and Richard A Friesner. The VSGB 2.0 model: a next generation energy model for high resolution protein structure modeling. *Proteins: Structure, Function, and Bioinformatics*, 79(10):2794–2812, 2011.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2023.
- Zhihai Liu, Minyi Su, Li Han, Jie Liu, Qifan Yang, Yan Li, and Renxiao Wang. Forging the basis for developing protein–ligand interaction scoring functions. *Accounts of Chemical Research*, 50(2):302–309, 2017.
- Chao Lu, Chuanjie Wu, Delaram Ghoreishi, Wei Chen, Lingle Wang, Wolfgang Damm, Gregory A Ross, Markus K Dahlgren, Ellery Russell, Christopher D Von Bargen, et al. OPLS4: Improving force field accuracy on challenging regimes of chemical space. *Journal of chemical theory and computation*, 17(7):4291–4300, 2021.

- James A Maier, Carmenza Martinez, Koushik Kasavajhala, Lauren Wickstrom, Kevin E Hauser, and Carlos Simmerling. ff14SB: improving the accuracy of protein side chain and backbone parameters from ff99SB. *Journal of Chemical Theory and Computation*, 11(8):3696–3713, 2015.
- Artur Meller, Michael Ward, Jonathan Borowsky, Meghana Kshirsagar, Jeffrey M. Lotthammer, Felipe Oviedo, Juan Lavista Ferres, and Gregory R. Bowman. Predicting locations of cryptic pockets from single protein structures using the pocketminer graph neural network. *Nature Communications*, 14(1):1177, Mar 2023.
- Xuan-Yu Meng, Hong-Xing Zhang, Mihaly Mezei, and Meng Cui. Molecular docking: a powerful approach for structure-based drug discovery. *Curr. Comput. Aided Drug Des.*, 7(2):146–157, June 2011.
- Vladimiras Oleinikovas, Giorgio Saladino, Benjamin P Cossins, and Francesco L Gervasio. Understanding cryptic pocket formation in protein targets by enhanced sampling simulations. *Journal of the American Chemical Society*, 138(43):14257–14263, 2016.
- Cheng Peng, Xinben Zhang, Zhijian Xu, Zhaoqiang Chen, Yanqing Yang, Tingting Cai, and Weiliang Zhu. D3PM: a comprehensive database for protein motions ranging from residue to domain. *BMC Bioinformatics*, 23(1):70, 2022.
- Francisc Sabanés Zariquiey, Edgar Jacoby, Ann Vos, Herman WT van Vlijmen, Gary Tresadern, and Jeremy Harvey. Divide and conquer. pocket-opening mixed-solvent simulations in the perspective of docking virtual screening applications for drug discovery. *Journal of Chemical Information and Modeling*, 62(3):533–543, 2022.
- Veronica Salmaso, Mattia Sturlese, Alberto Cuzzolin, and Stefano Moro. Combining self-and cross-docking as benchmark tools: the performance of DockBench in the D3R Grand Challenge 2. *Journal of Computer-aided Molecular Design*, 32:251–264, 2018.
- G Madhavi Sastry, Matvey Adzhigirey, Tyler Day, Ramakrishna Annabhimoju, and Woody Sherman. Protein and ligand preparation: parameters, protocols, and influence on virtual screening enrichments. *J. Comput. Aided Mol. Des.*, 27(3):221–234, March 2013.
- Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause, and Charlotte Bunne. Aligned diffusion schrödinger bridges. In *The 39th Conference on Uncertainty in Artificial Intelligence*, 2023.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- Andrew P Turnbull and Paul Emsley. Studying protein-ligand interactions using X-ray crystallography. *Methods Mol. Biol.*, 1008:457–477, 2013.
- Qi-Jun Wu, Tie-Ning Zhang, Huan-Huan Chen, Xue-Fei Yu, Jia-Le Lv, Yu-Yang Liu, Ya-Shu Liu, Gang Zheng, Jun-Qi Zhao, Yi-Fan Wei, et al. The sirtuin family in health and disease. *Signal Transduction and Targeted Therapy*, 7(1):402, 2022.
- Jason Yim, Andrew Campbell, Andrew YK Foong, Michael Gastegger, José Jiménez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S Veeling, Regina Barzilay, Tommi Jaakkola, et al. Fast protein backbone generation with SE(3)-flow matching. *arXiv preprint arXiv:2310.05297*, 2023.
- Jason Yim, Andrew Campbell, Emile Mathieu, Andrew Y. K. Foong, Michael Gastegger, Jose Jimenez-Luna, Sarah Lewis, Victor Garcia Satorras, Bastiaan S. Veeling, Frank Noe, Regina Barzilay, and Tommi Jaakkola. Improved motif-scaffolding with SE(3)-flow matching. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.
- Xuejin Zhang, Tomas Geffner, Matt McPartlon, Mehmet Akdel, Dylan Abramson, Graham Holt, Alexander Goncarenco, Luca Naef, and Michael Bronstein. Bending and binding: Predicting protein flexibility upon ligand interaction using diffusion models. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.

## A DATASETS

### A.1 D3PM DATABASE

The D3PM database collects all kinds of protein motions, including overall structural changes upon ligand binding and the inherently flexibility of protein, which we refer to as D3PM-Large, as well as flap movements of residues within binding pocket; which we do so as D3PM-Pocket (Peng et al., 2022).

#### A.1.1 D3PM-LARGE

The initial dataset is sourced from the D3PM-Large, where the RMSD of the  $C_\alpha$  carbon atoms between bound and unbound states is  $> 2.0\text{\AA}$ , following the protocol outlined in SBAAlign (Somnath et al., 2023). First, apo-holo pairs are selected and filtered to ensure that the reported RMSD of  $C_\alpha$  exceeds  $3\text{\AA}$ . The selected pairs are then aligned using the Kabsch (Kabsch, 1976) algorithm, and the RMSD of these  $C_\alpha$  is recalculated. Pairs are retained if the recalculated RMSD falls within a small margin of error, indicating minimal reconstruction error. Lastly, proteins with more than one chain are removed from the set. Then the structures are randomly split in training, validation and tests sets.

Table 5: D3PM-Large training, validation, and test sets.

	# Unique chains	# Apo chains	# Holo chains	# Unique Proteins
Train	1291	799	1133	443
Val	150	136	147	73
Test	150	141	144	74

#### A.1.2 D3PM-POCKET

This dataset contains pairs of apo-holo proteins where the overall RMSD between  $C_\alpha$  atoms is  $< 2.0\text{\AA}$ . We use the entire dataset for evaluation, with no other data processing than the one outlined in A.3

Table 6: D3PM-Pocket training, validation, and test sets.

	# Unique chains	# Apo chains	# Holo chains	# Unique Proteins
Test	1723	861	862	848

### A.2 PDDBIND-MD

However informative, the D3PM-Pocket dataset has quite a small size. Due to the lack of additional holo-apo pairing data, we decided to artificially generate more data using molecular dynamics. While it is quite complicated to generate holo conformations from apo structures, the reverse procedure is a simpler problem. In this dataset we take advantage of the PDDBind dataset (Liu et al., 2017) split in temporal splits following DiffDock (Corso et al., 2023). This set contains bound holo structures, and employing molecular dynamics, we collapse the pockets to generate an artificial apo structure, that is naturally paired with the original holo one.

First a quality filter is passed on the protein structures, which consists in the removal of structures with backbone breaks. Afterwards, for each remaining entry, the ligand was removed, and the protein structure was prepared for MD simulations. The proteins were parameterized using the AMBER ff14SB (Case et al., 2023; Maier et al., 2015) force field in the gas phase to ensure consistency of force field parameters across all systems.

Each system underwent an initial heating phase followed by an equilibration before trajectory analysis. The heating and equilibration steps were conducted using implicit solvent conditions, with the generalized Born model, to minimize computational cost while maintaining biophysical relevance.

All systems were initially heated from 0 K to 550 K over 1000 steps. This heating phase utilized Langevin dynamics (ntt=3) with a collision frequency of  $100.0 \text{ ps}^{-1}$  to maintain thermal stability. The timestep was set to 0.0005 ps, ensuring controlled energy fluctuations during heating. After heating, each system was equilibrated for 4 ps at a constant temperature of 550 K. A higher Langevin collision frequency ( $200.0 \text{ ps}^{-1}$ ) was used during equilibration to stabilize the system at 550 K while preventing excessive energy fluctuations. The timestep was increased to 0.001 ps to allow for a more efficient sampling of the conformational space.

To analyze the stability and dynamic behavior of each protein, residues within a  $5\text{\AA}$  radius of the removed ligand were extracted. The RMSD of these residues was calculated throughout the MD trajectory to evaluate conformational changes in the binding pocket. This protocol was used to determine which systems exhibited the most significant conformational changes in the binding site after ligand removal. Afterwards, proteins having a minimum RMSD change of  $0.5\text{\AA}$  were kept.

Table 7: PDBBind-MD training, validation, and test sets.

	# Unique Chains	# Apo chains	# Holo chains	# Unique Proteins
Train	6596	6596	6596	2334
Val	395	395	395	353
Test	185	185	185	80

### A.3 DATA PROCESSING

Apo and holo chains are featurized following AlphaFold2 (Jumper et al., 2021), where each residue is associated with SE(3)-equivariant frames that map rigid transformations onto an idealized backbone configuration. Each frame is defined by the atomic coordinates of the nitrogen (N), alpha carbon ( $C_\alpha$ ), carbonyl carbon (C), and oxygen (O) atoms, with transformations incorporating experimental bond angles and lengths (Engl & Huber, 2012). The backbone oxygen position is determined by an additional rotation around the  $C_\alpha$ -C bond. The final representation encodes all heavy atom coordinates as structured transformations in  $\mathbb{R}^{N \times 4 \times 3}$ , ensuring consistency across conformational states.

In this step, we additionally add a mask in the holo structure for pocket residues, defined as those within a  $8\text{\AA}$  distance cutoff of the ligand. After this, apo-holo pairs undergo local sequence alignment and only matching positions are kept. Then apo structures are centered by subtracting the center of mass from the  $C_\alpha$  positions and holo structures are superimposed onto them using Kabsch alignment of the  $C_\alpha$  atoms. During apo-holo alignment, we transfer the holo-defined pocket mask to the apo and retain up to 512 residues, cropping the excess when necessary.

### A.4 DATASET STATISTICS

We report summary statistics for each dataset, as they help interpret results, putting in context the obtained metrics.

Table 8: Summary statistics for the datasets.

Dataset	Length			$C_\alpha$ RMSD			Seq. Id. (%)		
	Med.	Mean	Std	Med.	Mean	Std	Med.	Mean	Std
PDBBind-MD	223	244	127.48	0.58	0.61	0.31	100.00	99.98	1.35
D3PM-Pocket	267	290	121.49	0.72	1.10	1.19	97.67	94.70	10.28
D3PM-Large	270	285	136.35	3.07	4.55	4.27	95.20	87.02	19.88

## B TRAINING SBALIGN

To retrain the SBAlign model on the PDBBind-MD dataset, we modified some hyperparameters to better suit the model to the task of predicting small conformational changes. Primarily, we use a lower diffusivity value, which controls the variance of the stochastic process.

We hypothesize that this led to some instabilities in the loss function, defined as:

$$L(\theta, \phi) := \mathbb{E} \left[ \int_0^1 \left\| \mathbb{G} - (b_t^\theta + m^\phi(X_t)) \right\|^2 + \lambda_t \|m^\phi(X_t)\|^2 dt \right], \quad (5)$$

with  $\mathbb{G}$ :

$$\frac{\mathbf{x}_1 - X_t}{\beta_1 - \beta_t} \quad (6)$$

where  $\beta_t$  is given by:

$$\beta_t := \int_0^t g_s^2 ds. \quad (7)$$

As  $g_s^2 \rightarrow 0$ ,  $\beta_t \rightarrow 0$ , and consequently, the term in Equation 6 grows larger, causing large gradients. These gradients can destabilize the optimization, leading to divergence and therefore to increasingly larger loss values during the training. This hypothesis aligns with the observed instability when training with low diffusivity.

We also note this affects inference<sup>2</sup> equally. In Table 9 we report metrics for several values of  $g$  with the model trained on PDBBind-MD, and in Table 10 we present similar results for the original SBAlign model. These results indicate that SBAlign is very sensitive to this parameter and struggles to generalize across a broader range of motions.

Table 9: Conformational changes results in the PDBBind-MD Test Set. RMSD between generated conformations and true holo structures after doing inference using SBAlign at different  $g$  values.

$g$ value	RMSD (Å) ↓			ΔRMSD (Å) ↑		
	Median	Mean	Std	Median	Mean	Std
0.001	0.57	0.59	0.15	0.0009	0.0009	0.0003
0.01	0.52	0.55	0.18	0.05	0.03	0.1
0.02	0.77	0.87	0.32	-0.28	-0.21	0.28

Table 10: Conformational changes results in the D3PM-Large Test Set. RMSD between generated conformations and true holo structures after doing inference using SBAlign at different  $g$  values.

$g$ value	RMSD (Å) ↓			ΔRMSD (Å) ↑		
	Median	Mean	Std	Median	Mean	Std
0.1	4.43	6.70	5.16	0.036	0.029	0.003
1.0	3.80	4.98	3.95	1.18	1.76	2.52
2.0	7.69	8.87	4.14	-2.60	-2.11	3.06

<sup>2</sup>In all inferences for SBAlign across this work, we generate 10 structures in 10 sampling steps and report metrics for the best one.

## C CROSS-INFERENCES RESULTS

In this section, we assess the capabilities of each model to generalize across different data distributions. To this end, we perform inferences with the models trained on large conformational changes to predict small ones and vice versa.

As we can see in Table 11, both models suffer from reduced performance when predicting conformational changes of a magnitude different that those it was trained on. Intuitively, we observe that the models trained on large conformational changes fail by producing movements that are too large and even lead corrupt structures, as evidenced by RMSD values much higher than the range found in the apo-holo pairs in the dataset. Conversely, the ones trained on small ones are able to capture only a small part of movement and fail to capture the full range of it, reflected in the combination of low  $\Delta$ RMSD values and high RMSD ones, but that fall within the distribution of apo-holo pairs in the dataset.

While these results are perhaps to be expected, it highlights a key limitation of the models, which ultimately stems from a lack of diverse, high-quality data.

Table 11: Conformational changes results for cross-inferences. RMSD between generated conformations and true holo structures after doing inference in the test sets of the corresponding datasets.

Model	Dataset Used		RMSD (Å) ↓			$\Delta$ RMSD (Å) ↑		
	Train	Inference	Med.	Mean	Std	Med.	Mean	Std
SBAAlign	D3PM-Large	PDBBind-MD	2.23	2.36	0.45	-1.66	-1.77	0.45
Sesame	D3PM-Large	PDBBind-MD	2.38	2.50	0.90	-1.79	-1.93	0.90
SBAAlign	PDBind-MD	D3PM-Large	4.46	6.74	5.16	0.025	0.003	0.31
Sesame	PDBind-MD	D3PM-Large	4.22	6.31	5.04	0.12	0.12	0.16

## D VISUALIZATIONS OF CRYPTIC POCKET PREDICTION

Here, we provide two example of PocketMiner’s predictions for apo, holo and generated structures, showing comparisons of the structures and PocketMiner’s per-residue predictions, where residues predicted to belong to a cryptic binding pocket with a probability  $p > 0.5$  are highlighted in red.

As seen in Figures 2 and 3, the generated structures are closer to the holo conformations for some regions of the proteins, and PocketMiner’s per-residue predictions are more closely aligned to those of the holo than the apo as well.

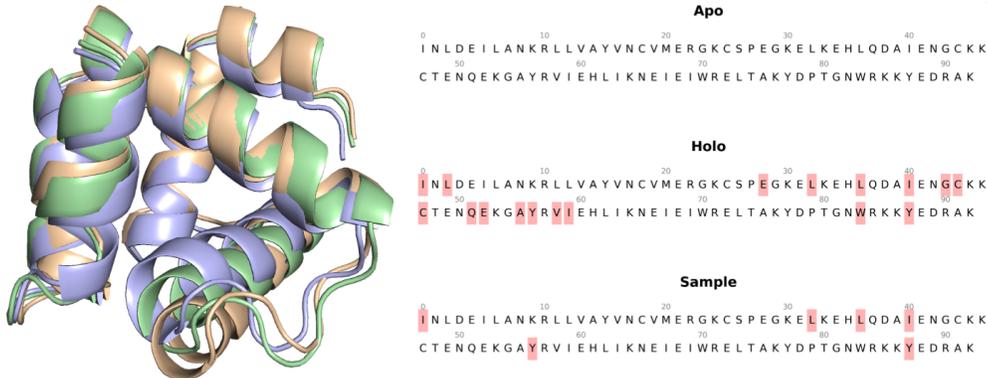


Figure 2: Left: Apo (blue), Holo (yellow) and Generated Sample (green). Right: PocketMiner per-residue predictions for each structure (PDBs: 1KX9-B, 1N8V-B)

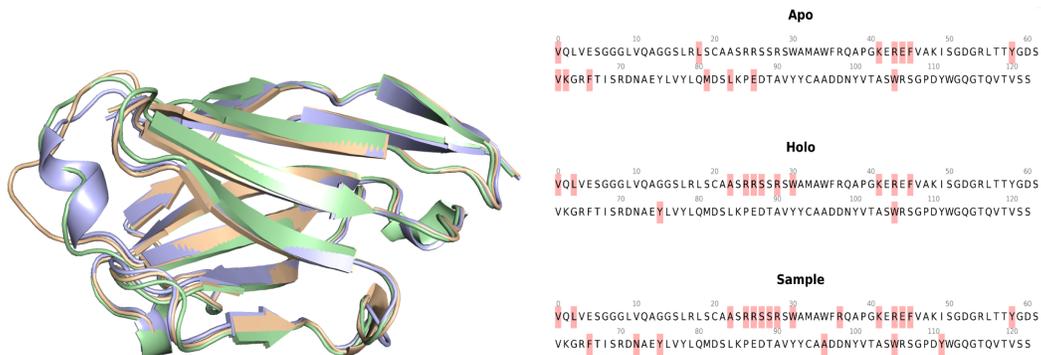


Figure 3: Left: Apo (blue), Holo (yellow) and Generated Sample (green). Right: PocketMiner per-residue predictions for each structure (PDBs: 3QXW-B, 3QVX-A)

## E DOCKING METHODS AND VISUALIZATIONS

### E.1 METHODS

Protein Data Bank (PDB) was the source for the apo and holo X-ray structures. To ensure a fair comparison with the model, side chains (except beta carbon), water, ions, and heteroatom molecules were removed from the crystal structures. Missing side chains in both X-rays and models were reincorporated using Prime (Jacobson et al., 2004). The Protein Preparation Wizard (Sastry et al., 2013) in Schrödinger (Schrödinger Release 2024-4) was used to add hydrogens and assign bond orders. Hydrogen bond and side-chain orientations were optimized at pH 7.4 with PROPKA (Bas et al., 2008). Finally, in those cases where side-chains were reincorporated they were minimized with Prime, using OPLS4 force field (Lu et al., 2021) and solvation VSGB (Li et al., 2011). Otherwise, we applied a restrained minimization of 0.3Å RMSD. Grid-based receptor preparation was performed by centering in the ligand-binding site identified from the X-ray structure. Ligand docking was conducted using the standard Glide SP (Halgren et al., 2004) protocol. Figures were generated with PyMOL (DeLano, 2002).

### E.2 SESAME HELPS CORRECT SIDE CHAIN PREDICTION ERRORS

We further evaluate Sesame using the 4LVT complex. Interestingly, in this case, our side chain reconstruction protocol performs significantly worse. Comparing the RMSD results of the reconstructed side chains to those obtained from self-docking using the X-ray side chains, we observe a substantially lower geometric similarity. We hypothesize that this is due to the apo conformation having a more closed pocket than the holo conformation, leading to a worse placement of side chains.

However, applying the same protocol to Sesame’s generated structure yields improved results. This further reinforces the potential of extending the method to include side chains as a promising direction for future research.

Table 12: RMSD and corresponding docking poses for the X-ray Holo versus X-ray Holo, Reconstructed Holo, X-ray Apo and Model. (PDBs: 4LVT, 1GJH)

	RMSD (Å) ↓	Docking Pose
Holo X-ray vs Holo X-Ray	3.063	Figure 4.D
Holo X-ray vs Reconstructed Holo	8.752	Figure 4.E
Holo X-ray vs Apo X-Ray	14.886	Figure 4.F
Holo X-ray vs Model	6.081	Figure 4.G

### E.3 VISUALIZATIONS

Figure 4 displays the highest-ranked docking pose for each structure based on Glide score, which was subsequently used for RMSD calculations.

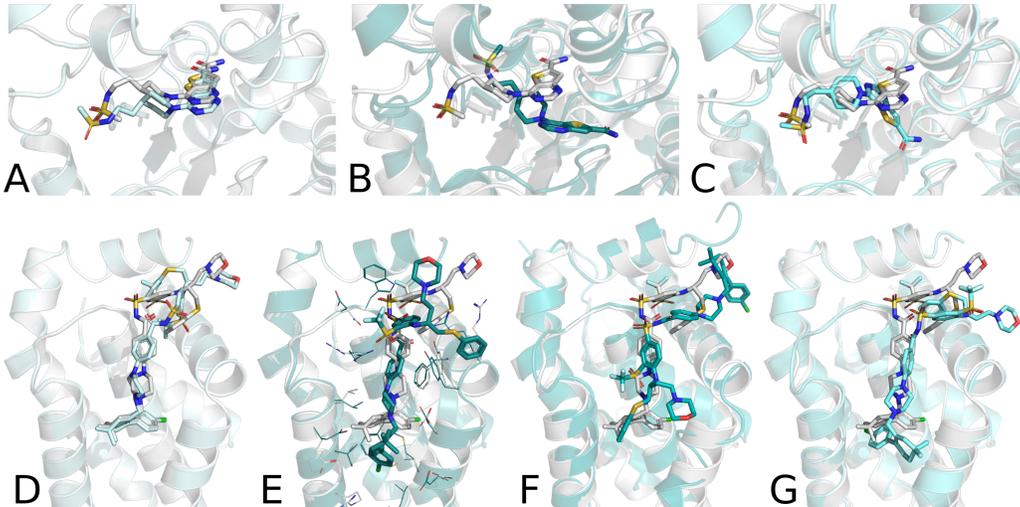


Figure 4: Docking Poses used for RMSD calculations in Tables 4 and 12. Alphabetical ordering follows order of appearance in the respective tables.

### F DETAILED METHODS

In this section, we aim to provide a more detailed theoretical background on flow matching and in the used loss functions in our work.

**Flow Matching.** Flow matching is generative modeling paradigm that was introduced as a efficient and simulation-free way to learn continuous normalizing flows (CNFs) (Lipman et al., 2023). CNFs are a class of deep generative models that generate data by integrating an ordinary differential equation over a learned vector field (Chen et al., 2018).

Let  $q_0$  and  $q_1$  be two distributions, one can learn a vector field  $v_t$  that induces a CNF  $\psi_t(x)$  that transports  $q_0$  to  $q_1$ :

$$q_1(x) = [\psi_1]_{\#} q_0(x) \quad (8)$$

where  $\#$  denotes the pushforward operator.

The marginal vector field  $v_t(x)$  can then be parametrized with a neural network  $v_t(x_t; \theta)$  and learned by defining conditional flows  $\psi_t(x_0|x_1)$  interpolating between  $x_0 \sim q_0$  and  $x_1 \sim q_1$  and their associated vector field  $u_t(x_t|x_1)$ , and regressing against the conditional vector field with the CFM objective:

$$\mathcal{L}_{\text{CFM}} := \mathbb{E}_{t, x_0 \sim q_0, x_1 \sim q_1} \left[ \|v_t(x) - u_t(x|x_1)\|^2 \right], \quad (9)$$

Although originally constrained to a Gaussian starting distribution, flow matching was generalized, relaxing said constraint and allowing arbitrary couplings by taking a joint distribution  $q(z)$ , over several possible choices (Tong et al., 2023; Albergo et al., 2023).

Flow matching has additionally been extended to Riemannian Manifolds (Chen & Lipman, 2024). On a manifold  $\mathcal{M}$ , the CNF  $\psi_t(x)$  can be defined by integrating along a time-dependent vector field  $v_t(x) \in \mathcal{T}_x \mathcal{M}$ , where  $\mathcal{T}_x \mathcal{M}$  is the tangent space of the manifold at  $x \in \mathcal{M}$ .

As such, we can similarly learn the marginal vector field with a similar objective:

$$\mathcal{L}_{\text{CFM}} := \mathbb{E}_{t, x_0 \sim q_0, x_1 \sim q_1} \left[ \|v_t(x) - u_t(x|x_1)\|_g^2 \right], \quad (10)$$

Where  $\|\cdot\|_g^2$  is the norm induced by the Riemannian metric  $g$ , and the conditional flows  $\psi_t(x) = \psi_t(x_0|x_1)$  can be defined by interpolating along the geodesic paths for simple manifolds (Yim et al., 2023).

The extension of flow matching to Riemannian Manifolds makes this framework readily applicable to the protein modeling field and to our goal of a generative model of holo structures given apo ones. More specifically, we follow FoldFlow and parameterize backbones as SE(3)-equivariant frames that represent rigid transformations  $T = (r, x) \in SE(3)$  that consist of a rotation  $r \in SO(3)$  and a translation  $x \in \mathbb{R}^3$  (Jumper et al., 2021). This formulation allows us to decompose the CFM process independently in  $SO(3)$  and  $\mathbb{R}^3$ . In each space, we can define conditional flows using the geodesic (2) and linear (3) interpolants, respectively, and use their associated vector fields to optimize the following simplified objectives:

$$\mathcal{L}_{SO(3)}(\theta) = \mathbb{E}_{t, q(r_0, r_1)} \|v_\theta(t, r_t) - u_t(r_t|r_0, r_1)\|_{SO(3)}^2, \quad \text{with } r_t = \exp_{r_0}(t \log_{r_0}(r_1)) \quad (11)$$

$$\mathcal{L}_{\mathbb{R}^3}(\theta) = \mathbb{E}_{t, q(x_0, x_1)} \|v_\theta(t, x_t) - u_t(x_t|x_0, x_1)\|^2, \quad \text{with } x_t = (1-t)x_0 + tx_1 \quad (12)$$

Where the conditional vector fields for rotation and translation take the following expressions, respectively:

$$u_t(x_t|x_0, x_1) = x_1 - x_0 \quad (13)$$

$$u_t(r_t|r_0, r_1) = \log_{r_t}(r_0)/t \quad (14)$$

Following previous works (Yim et al., 2023), we perform all modelling within the zero center of mass (CoM) subspace, which guarantees that the distribution of the sampled frames is SE(3)-invariant, and pre-align with the Kabsch algorithm to remove global rotations that result in increase kinetic energy in the ODE, and parameterize the velocity predictions  $v_\theta(t, x_t)$  as a prediction of the starting point  $\hat{x}_0$  given  $x_t$  (Yim et al., 2023; Bose et al., 2024; Jing et al., 2024).

**Final loss function.** Finally, the complete loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{FOLDFLOW}} + \mathcal{L}_{\text{FAPE}} + \mathbb{1}_{\{t < 0.25\}} \lambda_{\text{aux}} \mathcal{L}_{\text{aux}}, \quad (15)$$

with

$$\mathcal{L}_{\text{FOLDFLOW}} = \mathcal{L}_{\text{FOLDFLOW-SO(3)}} + \mathcal{L}_{\text{FOLDFLOW-R}^3}, \quad (16)$$

and:

$$\mathcal{L}_{\text{FOLDFLOW-R}^3} = \mathbb{E}_{t \sim \mathcal{U}([0,1]), q(x_0, x_1), \rho_t(x_t|x_0, x_1)} \|v_\theta(t, x_t) - (x_1 - x_0)\|^2 \quad (17)$$

$$\mathcal{L}_{\text{FOLDFLOW-SO(3)}} = \mathbb{E}_{t \sim \mathcal{U}([0,1]), q(r_0, r_1), \rho_t(r_t|r_0, r_1)} \|v_\theta(t, r_t) - \log_{r_t}(r_0)/t\|_{SO(3)}^2 \quad (18)$$

**FrameFlow Auxiliary Losses.** In addition to the flow matching and FAPE losses, we also include the auxiliary losses from Yim et al. (2023). Following Bose et al. (2024), we apply these atomic constraints for  $t < 0.25$ , to encourage a better local neighbour representation during the last steps of the generation procedure. These constraints consist a regression on the positions of the backbone atoms  $\mathcal{L}_{\text{bb}}$  and a pairwise distance loss on the local neighborhood  $\mathcal{L}_{2D}$ .

$$\mathcal{L}_{\text{aux}} = \mathbb{E}_{\mathbb{Q}}[\mathcal{L}_{\text{bb}} + \mathcal{L}_{2\text{D}}], \quad \mathcal{L}_{\text{bb}} = \frac{1}{4N} \sum \|A_0 - \hat{A}_0\|^2, \quad \mathcal{L}_{2\text{D}} = \frac{\left\| \mathbb{1}_{\{D < 6\text{\AA}\}}(D - \hat{D}) \right\|^2}{\sum \mathbb{1}_{D < 6\text{\AA}} - N} \quad (19)$$

where  $\mathbb{Q}$  is defined as the factorized joint distribution  $\mathbb{Q}(t, x_0, x_1, \tilde{x}_t) := \mathcal{U}(0, 1) \otimes \bar{\pi}(x_0, x_1) \otimes \rho_t(\tilde{x}_t | x_0, x_1)$ ,  $\mathbb{1}$  is the indicator function that indicates the membership to the local neighborhood according to atomic distances in Angstroms ( $\text{\AA}$ ) and  $D$  is a multidimensional array constructed from the pairwise distances between the four heavy atoms belonging to the backbone  $D_{ijab} = \|A_{ia} - A_{jb}\|$ .