

Genesis Molecular AI - Technical Report

Pearl: A Foundation Model for Placing Every Atom in the Right Location

Genesis Research Team¹, Alejandro Dobles¹.†, Nina Jovic¹.†, Kenneth Leidal¹.†, Pranav Murugan¹.†, David C. Williams¹.†, Drausin Wulsin¹.†, Nate Gruver¹.‡, Christina X. Ji¹.‡, Korrawat Pruegsanusak¹.‡, Gianluca Scarpellini¹.‡, Ansh Sharma¹.‡, Wojciech Swiderski¹.‡, Andrea N. Bootsma¹.††, Richard Strong Bowen¹.††, Charlotte Chen¹.††, Jamin Chen¹.††, Marc André Dämgen¹.††, Benjamin DiFrancesco¹.††, J. D. Fishman¹.††, Alla Ivanova¹.††, Zach Kagin¹.††, David Li-Bland¹.††, Zuli Liu¹.††, Igor Morozov¹.††, Jeffrey Ouyang-Zhang¹.††, Frank C. Pickard IV¹.††, Kushal S. Shah².††, Ben Shor¹.††, Gabriel Monteiro da Silva¹.††, Roy Tal².††, Maxx Tessmer¹.††, Carl Tilbury¹.††, Cyr Vetcher¹.††, Daniel Zeng¹.††, Maruan Al-Shedivat¹.#, Aleksandra Faust¹.#, Evan N. Feinberg¹.#, Michael V. LeVine¹.# and Matteus Pan¹.#

 1 Genesis Molecular AI; 2 NVIDIA; † Lead Contributor; ‡ Core Contributor; † Contributor; $^\sharp$ Senior Contributor; $^\sharp$ Work done during internship Ordered alphabetically within each group

Abstract—Accurately predicting the three-dimensional structures of protein-ligand complexes remains a fundamental challenge in computational drug discovery that limits the pace and success of therapeutic design. Deep learning methods have recently shown strong potential as structural prediction tools, achieving promising accuracy across diverse biomolecular systems. However, their performance and utility are constrained by scarce experimental data, inefficient architectures, physically invalid poses, and the limited ability to exploit auxiliary information available at inference. To address these issues, we introduce PEARL (Placing Every Atom in the Right Location), a foundation model for protein-ligand cofolding at scale. PEARL addresses these challenges with three key innovations: (1) training recipes that include large-scale synthetic data to overcome data scarcity; (2) architectures that incorporate an SO(3)-equivariant diffusion module to inherently respect 3D rotational symmetries, improving generalization and sample efficiency, and (3) controllable inference, including a generalized multi-chain templating system supporting both protein and non-polymeric components as well as dual unconditional/conditional modes. PEARL establishes a new state-of-the-art performance in protein-ligand cofolding. On the key metric of generating accurate (RMSD < 2 Å) and physically valid poses, PEARL surpasses AlphaFold 3 and other open source baselines on the public Runs N' Poses and PoseBusters benchmarks, delivering 14.5% and 14.2% improvements, respectively, over the next best model. In the pocket-conditional cofolding regime, PEARL delivers 3.6x improvement on a proprietary set of challenging, real-world drug targets at the more rigorous RMSD < 1 Å threshold. Finally, we demonstrate that model performance correlates directly with synthetic dataset size used in training.

1. Introduction

Small molecule therapeutics are an essential component of modern medicine, but their discovery remains a slow, expensive, and high-risk endeavor. A common approach involves building *structure-activity relationship* (SAR) models to predict how changes to a ligand's chemical structure will affect its activity. While both qualitative and quantitative SAR models have been successfully used to design drugs, they require large amounts of expensive experimental data, making the process inefficient and often yielding an incomplete understanding of the activity landscape, especially activity cliffs. A key strategy to improve SAR modeling and to accelerate this process is *structure-based drug design* (SBDD), where scientists design new ligands by applying chemical and physical principles to 3D protein-ligand structures. Historically, SBDD has been limited by its reliance on experimentally determined structures, which are still slow and expensive to acquire. The rise of computational methods presents an opportunity to apply this approach at an unprecedented scale, using predicted 3D structures—referred to as *poses*—to triage candidates and to produce more information-rich SAR models.

Traditionally, 3D structure prediction has been performed using physics-based ligand docking methods [1–8]. However, these methods tend to treat the protein (semi-)rigidly, and are unable to account for protein conformational changes that may occur during ligand binding [9–11]. On the other hand, early protein folding

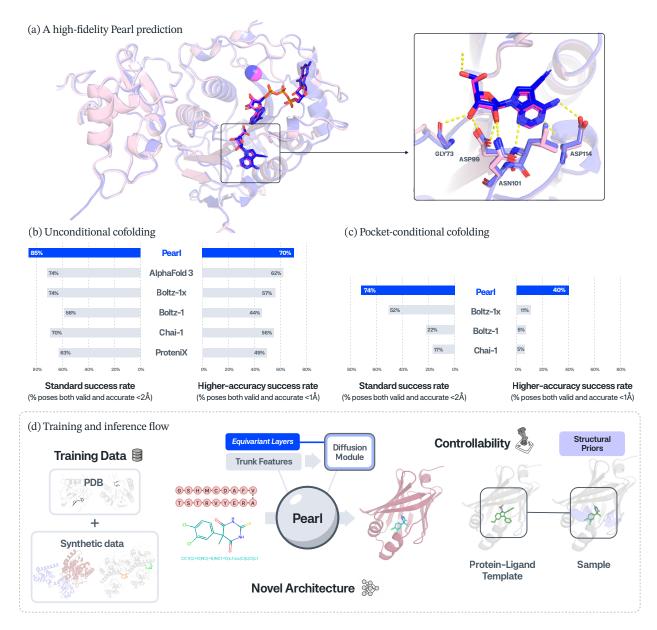


Figure 1. (a) PEARL prediction (magenta) superimposed on the experimental ground truth (blue) for SARS-CoV-2 (PDB: 8S8X, in the Runs N' Poses, released March 2024); PEARL performance in the (b) unconditional (Runs N' Poses) and (c) pocket-conditional (InternalXtals) cofolding modes; (d) PEARL training and inference flow.

models treat the protein as fully flexible, but are unable to jointly predict protein-ligand complexes [12]. Recently, deep learning models have evolved beyond protein-only folding to predict more general biomolecular complexes. Pioneering systems, such as AlphaFold 3 [13] and RoseTTAFold All-Atom [14], have demonstrated that a single model could jointly model proteins, ligands, nucleic acids, and other cofactors, predicting structures directly from their sequences and chemical topologies. The development of these models, and their descendants [15–20], is a major advancement in biomolecular modeling. Crucially, these cofolding models adapt the protein binding pocket geometry to different ligands, accounting for flexibility more fundamentally than previous approaches. Despite this, achieving the reliability needed to advance

drug discovery campaigns remains a significant challenge.

Reliable pose prediction faces three primary challenges. First, cofolding models learn from a relatively limited and biased corpus of experimentally determined structures. The Protein Data Bank (PDB) [21] offers orders of magnitude less data than domains like text or images, and this corpus is skewed toward certain targets and chemotypes due to varying experimental difficulty and scientific interest [22]. This scarcity and bias limit generalizability, leading models to memorize common structures, rather than learn transferable rules [23]. Second, a useful binding pose must satisfy complex physical requirements, such as low ligand strain, shape complementarity, and favorable non-covalent interactions. Subtle violations can produce superficially plausible but physically incorrect "hallucinations" that are not representative of real-world protein-ligand binding behavior. Finally, many existing models offer limited inference-time controllability, preventing them from using auxiliary structural information, such as a homologous ligand-bound structure, which are often available in real-world discovery programs.

To overcome these challenges, we introduce PEARL (Placing Every Atom in the Right Location), a generative foundation model for biomolecular structure prediction (Figure 1). PEARL addresses data scarcity and bias by training on a diverse mixture of experimental and synthetically generated protein-ligand complexes. Its novel architecture and training protocol (Figure 1) are designed to better learn the physical principles of binding, improving both generalization and the physical validity of generated poses. Furthermore, PEARL is explicitly built for adaptability and use during drug discovery campaigns. Its novel multi-chain templating system acts as a flexible conditioning mechanism, allowing scientists to leverage auxiliary structural information about the target protein, cofactors, and related ligands.

This paper makes three key contributions. First, we provide evidence for model performance scaling by training with large-scale synthetic data. Second, we introduce key modeling innovations, including curriculum training and an architecture that encompasses an SO(3)-equivariant diffusion module. Third, we develop a novel multi-chain templating system that incorporates non-polymeric information. This provides two key benefits: it enables conditioning on structural information during inference, and it supplies contextual data during training. Taken together, these contributions establish PEARL as the new state-of-the-art model for protein-ligand cofolding. PEARL achieves 85.2% and 84.7% success rates for generating accurate and physically valid poses on Runs N' Poses and PoseBusters benchmarks, respectively—both significant ($p \le 0.001$) improvements over AlphaFold 3 and all other baselines (RMSD < 2 Å and PB-valid). Notably, this high success rate shows almost no drop when physical validity checks are applied (a mere 0.7% on Runs N' Poses and 0.4% on PoseBusters), highlighting PEARL's ability to generate almost exclusively physically plausible poses. PEARL's performance advantage grows at the stricter, high-accuracy thresholds (RMSD < 1 Å and PB-valid) that are essential for guiding medicinal chemistry efforts, where it exhibits an almost 4× relative improvement over other models on a set of challenging, real-world targets (InternalXtals).

In the following sections, we contextualize PEARL in relation to prior work (Section 2), provide methodological detail—including its data, architecture, and inference strategies (Section 3)—present a comprehensive performance evaluation (Section 4), discuss the implications of our findings (Section 5), and offer concluding remarks (Section 6). This work has a dual aim: first to advance the development of foundation models for drug discovery, providing a path for cofolding systems that are not only state-of-the-art but also robustly generalizable, physically valid, and *practically useful* for accelerating therapeutic design. Second, and more broadly, we intend to push the frontier of generative AI for complex, physics-based problems in a low data regime. We believe the principles and methods outlined here will be of interest to the broader scientific community at the intersection of machine learning, chemistry, and biology.

2. Related Work

The computational prediction of protein-ligand binding poses has long been a cornerstone of SBDD, and it was traditionally carried out by physics-based ligand docking [24] (e.g., FlexX [1], Surflex-Dock [2], FRED [3], GOLD [4], Glide [5], MOE DOCK [6], DOCK [7], Autodock Vina [8]). These methods, which treat the pocket as (semi-)rigid, are either unable to treat induced fit—the protein conformational changes that may occur upon binding [9–11]—or can only account for minor, local rearrangements of the binding pocket. While there exist deep learning approaches for molecular docking [25, 26], including diffusion-based generative models, they do not jointly fold the protein and ligand, nor do they outperform classical docking methods [27].

A paradigm shift occurred with the advent of end-to-end folding models, most notably when AlphaFold 2 achieved near-experimental accuracy on the CASP14 challenge on protein monomers, starting from sequence and evolutionary information alone [12, 28]. RoseTTAFold extended this approach to complexes with multiple protein chains [29]. AlphaFold 3 (AF3) [13] generalized protein folding models to nearly all molecule types in the Protein Data Bank (PDB) [22], inspiring a new generation of *cofolding* models, such as Chai-1 [17], Boltz-1(x) [15], Boltz-2 [16], HelixFold3 [19], NeuralPlexer3 [20], and ProteniX [18]. Modern cofolding architectures typically share a blueprint, consisting of a trunk module (e.g., AF3's Pairformer) and a generative structure module, which are often denoising diffusion probabilistic models [30]. While PEARL is inspired by these approaches, it also altogether reflects a novel approach to cofolding. PEARL incorporates geometric deep learning principles [31] among other architectural, pretraining, and training improvements. These innovations, together with \pearl's unique inference-time controllability, enable state-of-the-art usability for practical small molecule drug discovery, including an enhanced ability to exploit known binding pockets and prior liganded structures.

A persistent bottleneck for all cofolding systems is their reliance upon relatively limited and biased training data from the PDB [32], which can lead to poor generalization to novel structures, instead relying on *memorization* [23]. To mitigate this challenge, models have incorporated strategies like training on "distillation data" [33] from other models [34] or integrating molecular dynamics and biochemical assay data [16]. In contrast, PEARL's training corpus is augmented with a large-scale, dataset of synthetically generated protein-ligand complexes, exposing the model to wider chemical diversity.

Finally, controllability, *e.g.*, the ability of users to incorporate structural priors during inference via templates, can improve the practical utility of cofolding models. AF3 established per-protein-chain templates as the standard for cofolding models. While scalable, this strategy can miss cross-chain interactions [13] and more sophisticated approaches have emerged. For example, later open source models introduced external experimental restraints and multimeric templates for protein complexes [16, 17]. PEARL further generalizes multi-chain templates to additionally encompass non-polymeric components, such as cofactors and related ligands.

3. Methods

PEARL is a generative foundation model for protein-ligand structure prediction. This section outlines its key novel components and their functions.

PEARL architecture PEARL's architecture consists of rotationally- and translationally-invariant representation modules (the "trunk") and diffusion modules responsible for predicting 3D coordinates. The trunk, featuring a lightweight triangle multiplication (trimul) module for efficiency [35], learns a rich,

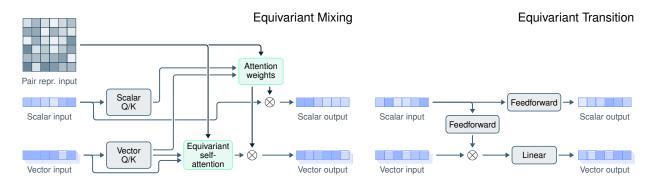


Figure 2. Key components of the equivariant diffusion module, including the equivariant transformer architecture (left) and the equivariant feed-forward layer with a gated nonlinearity for vector components (right).

position-independent pairwise representation that conditions the diffusion module. This design amortizes the trunk's high computational cost over the diffusion process's many sampling steps. A combination of data augmentation and equivariant architecture achieves rotational and translational equivariance of the final 3D structure. A key component is a novel SO(3)-equivariant diffusion module, which is unique among cofolding models. This module is constructed from equivariant transformer (EqT) blocks (Figure 2). By obeying SO(3) symmetry by construction, this module works in synergy with data augmentation to improve training sample efficiency and ensure predictions naturally respect the rotational symmetries of 3D space.

Templates To provide the model with a richer and more complete structural prior, PEARL generalizes the standard templating approach from protein-only templates to templates that also include non-polymeric components. The goal is to supply the model with a coherent, "holo-like" pocket environment—that is, a ligand-bound conformational state, in contrast to the unbound (apo) state. This provides context from similar ligands and their interactions with protein and cofactor components. In prospective settings, this templating form exposes crucial controllability and specificity for drug discovery scientists.

Training PEARL's novel training recipe enhances generalization and structural fidelity through its data mixture, curriculum, and optimization strategies. We train PEARL on a diverse dataset combining curated PDB structures, monomer distillation data (OpenFold [34] and the AlphaFold Database [36]), and a *novel*, *large-scale synthetic dataset*. Derived from public data, this dataset is generated using physics-based methods with diverse virtual ligands. The synthetic data introduces important chemical diversity beyond experimental data limitations.

Training follows a **five-stage curriculum** that progressively increases task complexity and data diversity. Initial stages use smaller crop sizes and simpler data mixtures (non-templated PDB, monomer distillation). Later stages gradually introduce more complex structural priors and templating information across datasets. The training applies structural templates across the datasets to vary the structural context.

To manage computational cost, PEARL employs a **conservative** bfloat16 (bf16) **mixed-precision strategy**. Computationally intensive trunk operations (*e.g.*, triangle ops via NVIDIA cuEquivariance [37], LayerNorm [38]) use bf16. However, numerically sensitive components (losses, coordinate projections), unstable operations (*e.g.*, softmax), and all model weights remain in full fp32 precision to ensure stability. This balanced approach significantly improves efficiency without compromising numerical stability.

Inference-time strategies PEARL uses several in-context learning strategies for high-quality pose generation, ranging from two fundamental operating modes to advanced techniques for fine-grained control. The dual

retrieval capabilities of its versatile multi-chain templating system and pocket conditioning enable PEARL to generate structures in two distinct cofolding modes tailored to different use cases: *unconditional* and *(pocket-aware) conditional*. In the unconditional cofolding mode, PEARL predicts the complex structure using only the protein's amino acid sequence and the ligand's topology. The method may use this sequential information to search for template structures and evolutionary information (*i.e.* MSAs) that the model may also use, similarly to AF3. The unconditional mode is useful for novel targets where the binding pocket is not known *a priori*. The conditional cofolding mode enables a common drug discovery scenario in which a reference apo or holo structure, or a hypothesized binding pocket, is available to guide generation.

Optional guidance and steering techniques can further control PEARL's sampling by modifying the denoising trajectory to enforce specific constraints or physical priors with a configurable parameter governing the strength of the enforcement. These methods are developed for practical drug discovery needs and allow for the additional use of a scientist's contextual knowledge without requiring costly model retraining. PEARL also offers the ability to modulate pose diversity at inference time.

4. Results

This section comprehensively evaluates PEARL against several baselines on public and proprietary benchmarks, addressing its performance, practical utility in drug discovery, and the factors underlying its capabilities.

4.1. Evaluation methodology

All training and evaluations are conducted on a high-performance computing cluster with NVIDIA H100 and H200 GPUs. We use optimized kernels from cuEquivariance (v0.6.0) [37] and a custom CUDA LayerNorm kernel [18] to accelerate key trunk operations. We evaluate PEARL against the following baselines: AlphaFold 3 [13], Boltz-1(x) [15], Boltz-2 [16], Chai-1 [17], and ProteniX [18] on diverse benchmarks. These benchmarks include the public Runs N' Poses [23] (RnP) and PoseBusters [27] datasets, as well as a proprietary InternalXtals dataset. InternalXtals dataset consists of 111 structures from a variety of internal programs. It is designed to test real-world generalization. To ensure consistent benchmarking, the public RnP and PoseBusters datasets were subsetted to structures released after 2023-06-01 and 2021-10-01, respectively (see Appendix A.1). PEARL's training set consisted strictly of publicly available structures from the PDB released on or before 2021-09-30; synthetic complexes were derived only from the pre-2021-09-30 publicly available structures. The PEARL model evaluated in this paper was trained without any proprietary experimental data.

We compare models in both unconditional cofolding and, when available, (pocket-aware) conditional cofolding modes to test for controllability of generation with structural priors (Appendix A.2). The primary accuracy metric is ligand root mean squared deviation (RMSD) [39]. We use a thresholded RMSD success rate ($<2\,\text{Å}$ or $<1\,\text{Å}$ denotes success), supplemented by physical plausibility checks from the PoseBusters [27] validation suite (PB-valid; Appendix A.3). To decouple the evaluation of *pose generation* from *pose selection*, we adopt the best@k protocol common in generative modeling [40], evaluating for k=1,5, and 20 out of 20 samples (see Appendix A.4 for discussion). We report the main results for the best@5 protocol, common in the generative modeling literature; the complete set of results is in Appendix B.1.

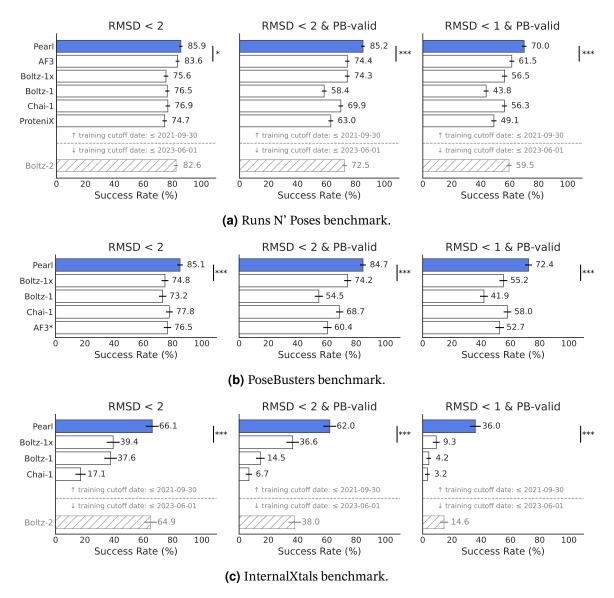


Figure 3. PEARL demonstrates state-of-the-art performance in the unconditional cofolding mode across multiple benchmarks. Shown are best@5 success rates for generating accurate (RMSD < 1 Å and < 2 Å) and physically valid (PB-valid) poses. Dashed lines group models for fair comparison (e.g., based on the training cutoff dates). Note 1(*): AlphaFold 3 results for PoseBusters use officially released metrics for poses selected for max confidence out of 25 samples (not best@5). Note 2: AlphaFold 3 was excluded from the InternalXtals benchmark due to license restrictions. Note 3: Trained with additional public data up to 2023-06-01, Boltz-2 is not directly comparable to the other evaluated models, which use 2021-09-30 or earlier training cutoff. Note 4: Statistical significance (one-sided t-test) is shown as: $\mathfrak{p} \leq 0.05, **: \mathfrak{p} \leq 0.01, ***: \mathfrak{p} \leq 0.001$.

4.2. PEARL is a state-of-the-art cofolding model

Benchmarked against leading cofolding models on public and proprietary datasets, PEARL consistently establishes a new state-of-the-art, delivering significantly improved success rates for accurate and physically valid poses compared to all baselines. The following sections detail these results, focusing on public benchmarks

(Section 4.2.1) and PEARL's generalization capabilities (Section 4.2.2).

4.2.1. PEARL achieves state-of-the-art performance on public benchmarks

Figure 3 shows PEARL's significant and consistent improvement over AlphaFold 3 and other methods across progressively stricter metrics. For the standard RMSD < 2 Å metric (best@5), PEARL attains top-ranking success rates of 85.9% on Runs N' Poses and 85.1% on PoseBusters (Figures 3a and 3b, left). This advantage grows when also evaluating for physical validity (RMSD < 2 Å and PB-valid, best@5). On Runs N' Poses, PEARL's 85.2% success rate is a 14.5% relative improvement over the next best models, AlphaFold 3 (74.4%) and Boltz-1x (74.3%) (Figure 3a, middle). Similarly, on PoseBusters, its 84.7% success rate represents a 14.2% relative improvement over Boltz-1x (74.2%) and a 40% relative gain over AlphaFold 3's officially released max confidence result (60.4%) (Figure 3b, middle). On PoseBusters, PEARL in the best@1 regime (selecting 1 pose at random) still maintains an 8% relative advantage over AlphaFold 3, which selects one max confidence pose from over 25 generated poses. In the PoseBusters best@20 setting PEARL has an 46% relative advantage over AlphaFold 3's max-confidence. For more details, see Appendix B.1.

Notably, PEARL's success rate shows almost no drop when validity checks are applied—a mere 0.7% drop on Runs N' Poses and 0.4% on PoseBusters (Figures 3a and 3b, middle)—highlighting its ability to generate almost exclusively physically plausible poses. A comprehensive breakdown of metrics at different aggregation levels (best@1, best@20) is provided in Appendix B.1 (Tables 1, 2, 3).

4.2.2. PEARL demonstrates strong generalization to novel drug targets

Beyond public benchmarks, we evaluate PEARL's generalization on the InternalXtals dataset, designed to test performance on targets and ligands structurally and chemically dissimilar to the public training corpus (Figure 3c). In the unconditional cofolding mode, PEARL achieves a 66.1% success rate for RMSD < 2 Å (best@5), a 67.8% relative improvement over the next best open source model (Boltz-1x at 39.4%). For physically valid poses (RMSD < 2 Å and PB-valid, best@5), PEARL maintains a strong lead with 62.0% success, a 69.4% relative gain over Boltz-1x (36.6%) (Figure 3c, left and middle). This robust performance on challenging, real-world data demonstrates that PEARL's advantages extend beyond academic benchmarks.

To further probe generalization, Figure 4 analyzes PEARL's performance on the RnP benchmark stratified by similarity to the training data, following [23]. Focusing on the lowest similarity buckets—the most rigorous test of generalization by this metric—reveals PEARL's consistent advantage. PEARL leads when generalizing to novel protein pockets (similarity to the training set < 0.2, as defined in [23, 41]), novel ligands (frequency = 0), and dissimilar chemotypes (Tanimoto similarity < 0.2). This strong performance across different generalization axes indicates PEARL learns transferable rules rather than memorizing training examples, establishing it as a suitable foundation model for diverse drug discovery tasks.

4.3. PEARL's predictions are relevant for drug discovery

This section establishes PEARL's practical relevance for drug discovery. Section 4.3.1 demonstrates its superior performance at the high-accuracy thresholds required for fine-grained molecular design. We also showcase its lead in accuracy and controllability in the presence of structural priors in the conditional cofolding mode (Section 4.3.2). Together, these results show that PEARL generates poses that are not just statistically accurate but better poised to be genuinely useful for guiding real-world therapeutic design.

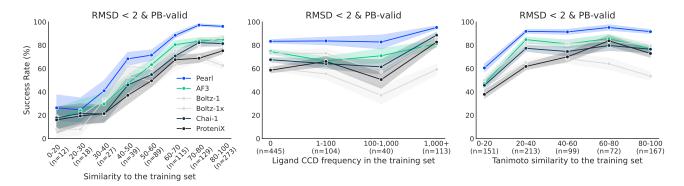


Figure 4. Stratified analysis of PEARL's generalization on the RnP in the unconditional cofolding mode. Shown are the best@5 success rate for generating valid poses (RMSD < 2 Å and PB-valid) when stratified by (left) overall similarity to the training set (product of binding pocket coverage and combined overlap score [SuCOS] of the ligand pose [23]), (middle) ligand frequency, and (right) Tanimoto similarity. PEARL exhibits strong generalization in the most challenging, low-similarity regimes, where it generally leads other models. The number of test examples in each slice is denoted by n. Boltz-2 is excluded as its expanded training set does not conform to the same similarity measures.

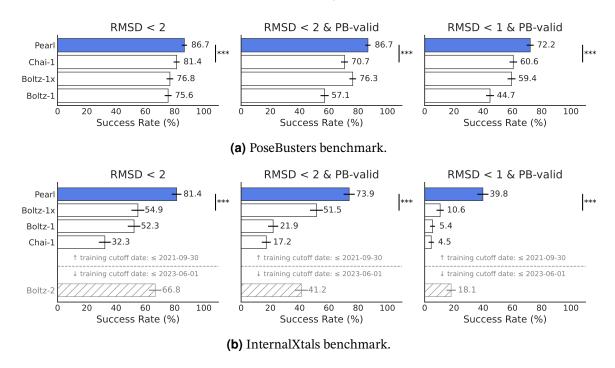


Figure 5. PEARL demonstrates state-of-the-art performance in the conditional cofolding mode across multiple benchmarks. Metrics shown are best@5 (from 20 samples). Dashed lines group models for fair comparision based on the training cutoff dates. Note 1: AlphaFold 3 was excluded due to license restrictions. Note 2: Trained with data up to 2023-06-01, Boltz-2 is not directly comparable to models using earlier cutoffs.

4.3.1. PEARL excels at the high-accuracy thresholds required for drug discovery

While RMSD < 2 Å is a standard benchmark metric, guiding fine-grained molecular design for affinity improvements often requires higher accuracy. We therefore evaluate models with a stricter threshold of RMSD < 1 Å and PB-valid (best@5). At this high-accuracy threshold, PEARL's performance advantage

becomes even more pronounced across all benchmarks (Figure 3, right panels).

On public benchmarks, PEARL reaches success rates of 70.0% on Runs N' Poses and 72.4% on PoseBusters, uniquely maintaining over 70% success while baselines' performance drops significantly (Figures 3a and 3b right). For instance, on Runs N' Poses, the next best model, AlphaFold 3, achieves 61.5%. On PoseBusters, the next best comparable model (Chai-1) achieves 58.0%. This state-of-the-art high-resolution performance is confirmed on the InternalXtals benchmark (Figure 3c, right). Here, PEARL achieves a 36.0% success rate at RMSD < 1 Å and PB-valid. This represents a nearly 4-fold relative improvement over the strongest comparable baseline (Boltz-1x at 9.3%) and more than double the success rate of Boltz-2 (14.6%), which benefited from a later training data cutoff. This demonstrates PEARL's superior ability to generate the high-fidelity poses critical for medicinal chemistry.

4.3.2. PEARL's structure-aware conditioning provides high-fidelity poses for guided research

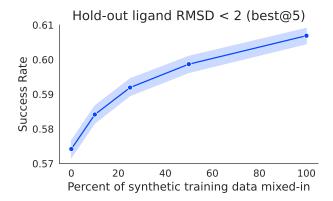
To evaluate PEARL's utility for practical structure-based drug design, we assessed its performance in the (pocket-aware) conditional cofolding mode, where structural priors (*e.g.*, known apo/holo structures) guide pose generation. In this mode, PEARL consistently outperforms all comparable baselines, demonstrating its effectiveness as a controllable tool (Figure 5). On the public PoseBusters benchmark (Figure 5a middle and right), PEARL's success rate is 86.7% (RMSD < 2 Å and PB-valid, best@5), a 14% relative improvement over Boltz-1x (76.3%). At the stricter RMSD < 1 Å and PB-valid threshold, its success rate of **72.2**% is significantly ahead of Chai-1 (60.6%). This guided approach delivers even stronger results on the InternalXtals benchmark (Figure 5b), with a 73.9% success rate (RMSD < 2 Å and PB-valid, best@5), a 43% relative improvement over Boltz-1x (51.5%). At the high-accuracy RMSD < 1 Å threshold on this proprietary set, PEARL's **39.8**% success rate is nearly double that of Boltz-2 (18.1%)—despite Boltz-2 benefiting from a later training data cutoff—and almost 4× better than Boltz-1x (10.6%). All evaluated models are provided the same pocket residues for conditioning (see Appendix A.2). These results confirm PEARL's superior performance and controllability when leveraging structural priors common in drug discovery workflows.

4.4. Key enablers of PEARL's performance

Having established PEARL's state-of-the-art performance, this section highlights two key enablers, presenting evidence for the contributions of large-scale synthetic data and the efficiency gains from architectural choices and hardware acceleration.

Synthetic data improves generalization: A core hypothesis in this work is that large-scale synthetic data can overcome PDB limitations and improve generalization. Our experiments with a smaller variant of PEARL validate this, revealing a clear scaling relationship: PEARL's success rate (RMSD < 2 Å, best@5) on a fixed hold-out set increases monotonically with the proportion of synthetic training data mixed into the training corpus (Figure 6). This data scaling trend confirms that augmenting experimental structures with diverse synthetic examples is a driver of PEARL's robust performance and generalization.

Architectural innovations and hardware acceleration improve efficiency: PEARL's architectural design significantly enhances computational efficiency. The conservative bfloat16 (bf16) mixed-precision strategy alone increases training speed by 22% and reduces memory usage by 11% compared to full fp32 precision, without impacting final model accuracy. Optimized kernels from NVIDIA's cuEquivariance (v0.6.0) [37] library for computationally intensive operations in the model's trunk and diffusion modules yield further acceleration. These kernels provide substantial speedups relative to a vanilla PyTorch baseline: an additional 15% speedup during training and a 10–80% speedup at inference time, with greater inference gains observed



Unique proteins	910
Average number of ligands per protein	640
Tanimoto similarity of ligands vs. reference	0.364
Total synthetic structures	582,065

Figure 6. PEARL's performance scales with the addition of synthetic data. The plot (left) shows a monotonic improvement in success rate (best@5, RMSD < 2 Å) for a smaller variant of PEARL on a fixed hold-out set as the proportion of synthetic data increases. The table (right) summarizes the scaling experiment's dataset, which included 582,065 synthetic structures across 910 proteins. A larger synthetic dataset was used to train the flagship PEARL model.

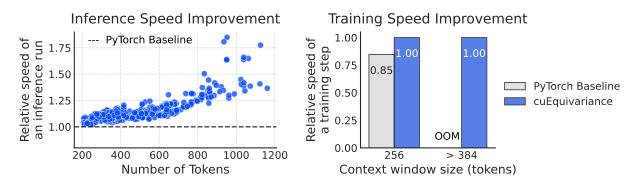


Figure 7. NVIDIA's optimized cuEquivariance kernels provide significant acceleration for PEARL relative to a vanilla PyTorch baseline. The kernels deliver a 10–80% speedup at inference time (left) and a 15% training speedup (right).

for larger inputs (more tokens) (Figure 7). Moreover, the GPU memory savings from cuEquivariance enable training with larger context window sizes. These combined optimizations—mixed precision and hardware-accelerated kernels—were critical for executing the large-scale training and evaluation in this work. Appendix B.2 provides more details on the mixed-precision implementation.

4.5. Case studies

Complementing the quantitative benchmarks, we qualitatively analyze PEARL and Boltz-2 poses drawn from Runs N' Poses in the unconditional cofolding mode. The five case studies highlight common cofolding failure modes, including both those that PEARL overcomes and those that it still struggles to resolve (Figures 8 and 9). Overall, this analysis links key failures to training set memorization. Models often fail to predict novel pockets or conformational changes and exhibit a bias for frequently occupied pockets, even when binding to those pockets is incompatible with the predicted conformation. This behavior, likely sensitive to training data composition and sampling, indicates a reliance on memorization over full generalization. We also show cases that underscore the necessity of strict RMSD < 1 Å thresholds. Many RMSD < 2 Å poses miss critical interactions, limiting their practical utility for potency prediction or compound ideation in drug discovery.

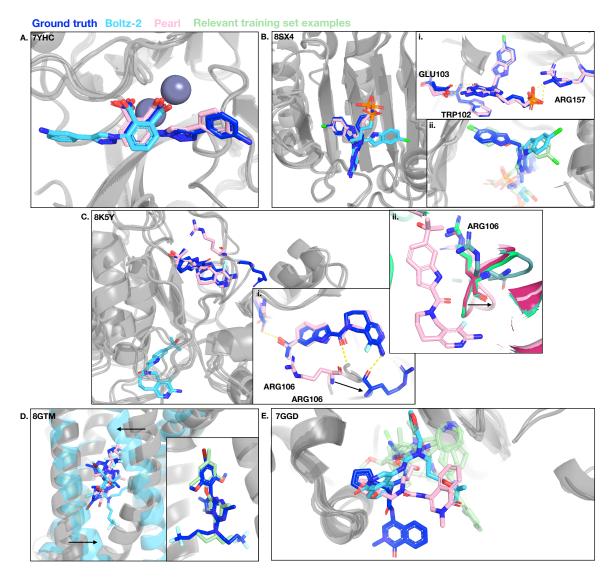


Figure 8. Qualitative analysis of PEARL's successes and common cofolding failure modes. The ground truth is shown in dark blue, PEARL's prediction in pink, Boltz-2's prediction in light blue, and relevant training set examples in green. PEARL *Successes* (A, B, C): PEARL correctly predicts challenging poses where Boltz-2 fails, such as for an inhibitor of VIM-2 MBL (A), an inhibitor of eIF4E (B), and a proMMP-9 inhibitor (C). *Common Failure Modes* (D, E): Both models can place ligands in pockets that are more frequently observed in the training data, leading to incorrect predictions for inhibitors of CRF1R (D) and SARS-CoV-2 main protease (E).

4.5.1. PEARL successes

Figure 8A shows PEARL correctly predicting the pose of a phthalic acid-based VIM-2 MBL inhibitor precursor. Boltz-2 identifies the correct overall pocket and key phthalic acid-Zn²⁺ interactions, but it fails to place the aminophenyl and triazole rings in their correct subpocket. This is despite both models seeing training examples occupying both sides of the pocket (*e.g.*, 1JJT, 1KR3, 3WXC, 5LCF). Similarly, for an inhibitor of the rate-limiting translation factor eIF4E (Figure 8B), PEARL correctly predicts the binding pocket and hydrogen bonds formed with Arg157, Glu103 and Trp102 (Figure 8B i. inset). Boltz-2, however, misplaces the indole

ring and the phosphate tail, placing the indole in an alternative subpocket that is seen in the training set (e.g., 5EHC, 6YLR). This demonstrates a memorization failure, as data for the correct subpocket was also available in the training set (4DUM, Figure 8B ii. inset). Figure 8C shows that PEARL correctly predicts a $< 2 \,\text{Å}$ pose for a proMMP-9 catalytic domain inhibitor, while Boltz-2 predicts the incorrect pocket. In fact, while all the baselines predict a $> 2.0 \,\text{Å}$ pose, only PEARL predicts not only a sub 2 Å pose, but also a sub 1 Å pose (0.96 Å).

While PEARL is sufficiently accurate to be useful, it is not without imperfections. For example, in the case of the proMMP-9 catalytic domain inhibitor, PEARL does not predict the correct loop conformation, and as a consequence the ligand fails to form an important loop backbone interaction. This is likely a memorization artifact, as earlier structures (1GKD, 2OVZ, 2OVX, 6ESM) were crystallized without the loop domain. Additionally, related proteins (proMMP-1, proMMP-7, and proMMP-9desFnII) lack the crucial interacting arginine in their homologous loop and that region is unresolved. Overall, a similar loop conformation is sampled in only a few training structures (*e.g.*, 5TH6, 5UE3, 5UE4), and superimposing this conformation clashes with the inhibitor (Figure 8C ii. inset). Nevertheless, PEARL is the only model that manages to predict a usefully accurate pose, even without capturing the exact conformation of the loop.

4.5.2. Common failure modes

Next, we examine cases where multiple cofolding models fail (Figure 8D and 8E). In the case of the CRF1R antagonist, BMK-C203, both PEARL and Boltz-2 place the ligand in the right pocket, but they fail to predict the correct ligand conformation. While PEARL does predict the correct pocket shape, Boltz-2 predicts a more closed conformation of the pocket. Interestingly, there are many structures of CRF1R and related proteins bound to ligands in other pockets (*e.g.*, 3PDS, 4QKX, 4PHU, etc.), but only one where the ligand, CP-376395, occupies the same pocket as BMK-C203 (4Z9G, Figure 8D inset). It is possible that the abundance of structures with ligands in other pockets lead the cofolding models to struggle when placing ligands into the BML-C203 pocket. Another failure mode is exemplified by the SARS-CoV-2 main protease inhibitor, Mpro-x12350. Both PEARL and Boltz-2 fail to place the inhibitor in the correct subpocket of the catalytic site (Figure 8E). Holo structures of other inhibitors in the training set occupy a similar subpocket as the PEARL and Boltz-2 predicted poses. Having never seen the Mpro-x12350 subpocket during training, both models appear to have failed to generalize.

4.5.3. Failure modes with RMSD < 2 Å

Poses below 2 Å RMSD, the accuracy threshold used in many benchmarks, can still have significant flaws that may mislead a drug discovery scientist. Even under this threshold, unrecoverable mistakes—such as missed key protein-ligand interactions, flipped rings leading to suboptimal interactions, or other subtle errors—can greatly reduce or even fully negate a pose's utility in drug discovery. We highlight some of these cases, as they represent further opportunities improving cofolding models. Figure 9 compares PEARL's poses selected from the 20 generated samples, comparing those that are RMSD < 1 Å with those that are between 1–2 Å RMSD. Figure 9A, Figure 9B and Figure 9C all show poses between 1–2 Å RMSD that contain problematic ring flips. Figure 9A shows the pyrazole ring of CK2-TN01 flipped, leading the pose to miss an important interaction with the backbone of Val116. Similarly, we observe a pose of the inhibitor M689 bound to AKR1C3 in which a phenol flips and as a consequence misses key interactions with the sidechain of Ser87 and the backbone of Met120 (Figure 9B). Sometimes ring flips lead to less obvious errors that would greatly mislead a drug discovery scientist; we find a pose in which the thiazole ring of the PPAR δ ligand lanifibranor is flipped, but no interactions are made or broken (Figure 9C). Finally, Figure 9D shows an example where the the pose is

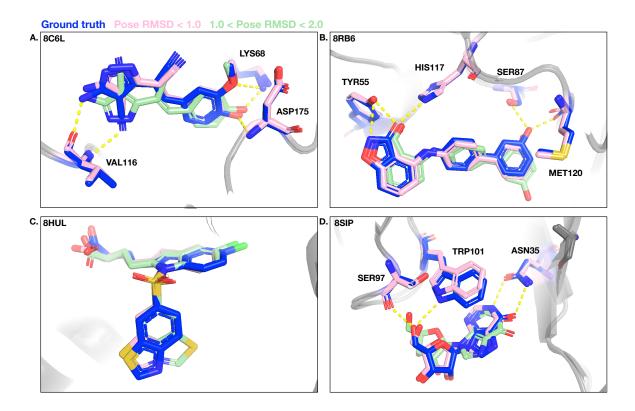


Figure 9. Qualitative analysis of the importance of high-accuracy predictions. The figure compares high-accuracy poses (RMSD < 1 Å, pink) with lower-accuracy poses (1 Å < RMSD < 2 Å, green). The lower-accuracy poses, while under the standard RMSD < 2 Å threshold, exhibit critical errors that limit their utility for drug discovery, such as ring flips (A, B, C) and slight translations (D). (**A**) A 1.99 Å pose of CK2-TN01 in complex with human protein kinase CK2 alpha displays a ring flip and misses a key interaction, which is correctly formed in the 0.42 Å pose. (**B**) A 1.70 Å pose of M689 in complex with AKR1C3 displays a problematic ring flip. Shown in comparison to the 0.36 Å pose. (**C**) A 1.51 Å pose of lanifibranor bound to human PPAR δ displays an incorrect flip of the thiazole ring. Shown in comparison to a 0.58 Å pose. (**D**) A 1.54 Å pose of inosine bound to a mouse IgG fragment displays a slight translation, missing all the key interactions that are correctly predicted in the 0.38 Å pose.

simply "not quite right." The inosine is very slightly translated in the pocket, leading it to miss all the key interactions with the protein.

5. Discussion

Interpretation of the results PEARL's state-of-the-art performance arises not from a single improvement, but from the synergy between its core innovations in data, training recipe, architecture, and inference, particularly in addressing the *low-data regime*, common in scientific ML. We tackle this challenge in several ways. First, we provide strong evidence that leveraging large-scale synthetic data helps overcome PDB limitations and boosts generalization. Crucially, we demonstrate a clear scaling relationship where model performance improves monotonically with the amount of synthetic data used in training (Figure 6). This finding validates synthetic data augmentation as a powerful strategy for building and scaling foundation models for science in data-scarce domains. The effectiveness of this data is further enhanced by a training

curriculum that progressively increases task complexity, aiding generalization. Second, incorporating an SO(3)-equivariant diffusion module introduces a strong geometric inductive bias from the domain of geometric deep learning. By enforcing rotational symmetry architecturally, the model achieves greater sample efficiency. We hypothesize this contributes significantly to PEARL's improved generalization observed in generated poses. Moreover, PEARL's substantially improved physical validity, evidenced by the minimal performance drop when PB-valid checks are applied (Section 4.2), suggests its architecture and training effectively enable both accuracy and physical plausibility with minimal tradeoff between the two. Third, PEARL's multi-chain templating system aids generalization by providing richer contextual data during training and functions as a sophisticated conditioning and in-context learning framework during inference. This allows expert users to provide structural "prompts" (templates) to guide generation effectively, enhancing accuracy in conditional settings (Figure 5).

Implications for drug discovery workflows These findings have direct implications for scientists using computational tools for structure-based drug design. The unconditional cofolding mode suits early-stage hypothesis generation for novel targets where no known pockets or experimental structures exist. The conditional cofolding mode is better suited for hit ID against a known pocket, or for hit-to-lead and lead optimization, where accuracy is paramount and an existing reference structure can be used to guide the model to produce poses of new or prospective analogs.

The need for high accuracy Critically, our results underscore the substantial limitations of the standard RMSD < 2 Å threshold for assessing practical utility in drug discovery. As demonstrated qualitatively (Figure 9) and increasingly recognized, many poses below this threshold exhibit critical flaws like incorrect ring flips or missed interactions that render them misleading for medicinal chemistry or downstream physics-based modeling. Guiding molecular design requires higher accuracy. PEARL's significantly improved success rates at the stricter RMSD < 1 Å threshold across all benchmarks (Figure 3, right panels) represent a key advancement. This superior high-accuracy performance, combined with its high overall physical validity, makes PEARL's predictions substantially more reliable and genuinely useful for structure-based drug design workflows. This practical advantage is particularly evident in demanding guided research scenarios, where PEARL's conditional cofolding mode delivers state-of-the-art accuracy even at the critical RMSD < 1 Å threshold on both the PoseBusters and InternalXtals datasets (Figure 5b).

Implications for AI for science More broadly, this work serves as a case study for building generative *foundation models in scientific domains*. The principles demonstrated—most notably the effectiveness of scaling with synthetic data in low-data regimes (Figure 6), alongside the benefits of architectural symmetry enforcement and flexible conditioning mechanisms—are likely transferable. These techniques hold potential for tackling other complex, generative modeling problems at the intersection of ML, chemistry, physics, and biology.

Limitations and future work While PEARL significantly improves performance, particularly in novel sequence and chemical spaces (Figures 3c, 5b), limitations remain. Like other models, PEARL's accuracy degrades on out-of-distribution (OOD) data, and it is susceptible to "long-tail" failures like mispredicting large induced-fit changes. Memorization artifacts persist, indicating synthetic data mitigates but does not eliminate this challenge. Furthermore, as also found by previous studies [23], existing confidence models are often uninformative for ranking and selecting the highest-accuracy pose, as the top-ranked pose is frequently no better than a random sample (see Figure 11 in the Appendix). Thus, while PEARL excels at generating high-quality poses (high best@k rates), reliably selecting the optimal pose remains a critical challenge for the field. Future work must focus on improving OOD generalization, robustness, and developing reliable pose selection methods. This includes exploring new training/inference protocols, better capturing protein

dynamics, advanced conditioning, alternative geometric architectures, integrating physics-informed ML, and addressing computational bottlenecks. PEARL provides a strong foundation for this next generation of generative models for science.

6. Conclusion

We introduced PEARL, a generative foundation model that establishes a new state-of-the-art for protein-ligand structure prediction by effectively addressing key limitations in existing cofolding systems. We identified primary challenges in data scarcity, learning physical constraints, and limited controllability, and presented PEARL's three-pronged solution: leveraging large-scale *synthetic data* within a curriculum *to improve generalization and overcome PDB biases*; incorporating *SO(3)-equivariant components for enhanced sample efficiency*; and employing a flexible *multi-chain templating* approach *for superior controllability and conditional accuracy*. The results demonstrate that these innovations lead to significant performance gains over existing baselines, including AlphaFold 3, particularly at the strictest drug discovery relevant accuracy metrics for protein-ligand complexes and small molecule tasks.

More importantly, our work goes beyond statistical benchmarks to validate the practical utility of these high-fidelity structures. We show that poses generated by PEARL are significantly more accurate (RMSD) and physically plausible (PB-valid), rendering them tangibly more valuable for *guiding medicinal chemistry decisions* and accelerating structure-based drug design workflows. The principles established here—particularly the effectiveness of scaling with synthetic data and the benefits of equivariant design—provide a clear path forward for developing more powerful and reliable foundation models. Continued development following these principles promises even more powerful tools for AI-driven scientific discovery. We believe these findings are not only significant for drug discovery but also hold broader implications for generative AI at the intersection of chemistry, physics, and biology.

Acknowledgments

The authors would like to thank Olivia Bass, Kathy Benemann, Ken Gong, Stacie Calad-Thomson, Kyle Tretina, and Pat Walters for the helpful feedback on the manuscript. In addition, we thank Michael Atkin, Nicholas Barry, Rebecca Boiarsky, Elizabeth Brown, Fenghong Chen, Kush Desai, Ajinkya Deshpande, Drake Diedrich, Michael Dumont, Gianpaolo Gobbo, Alex Goldberg, Anindit Gopalakrishnan, Victoria Ingman, Rohit Kundu, Dean Latney, Shir Levkowitz, Pavel Mikhalchuk, William McCarthy, Tom Metzger, Nikhil Murthy, Parvin Parineh, Diego Puppin, Davide Sabbadin, Colter Spearsmith, Kendrick Shen, Justin Steinman, Vyom Thakkar, William Wang, Lillian Weng, Jianbo Zhao for the helpful discussion and support.

References

- 1. Rarey, M., Kramer, B., Lengauer, T. & Klebe, G. A Fast Flexible Docking Method using an Incremental Construction Algorithm. *J. Mol. Biol.* **261**, 470–489 (1996).
- 2. Jain, A. N. Surflex: Fully Automatic Flexible Molecular Docking Using a Molecular Similarity-Based Search Engine. *J. Med. Chem.* **46**, 499–511 (2003).
- 3. McGann, M. R., Almond, H. R., Nicholls, A., Grant, J. A. & Brown, F. K. Gaussian docking functions. *Biopolymers* **68**, 76–90 (2003).

- 4. Verdonk, M. L., Cole, J. C., Hartshorn, M. J., Murray, C. W. & Taylor, R. D. Improved protein-ligand docking using GOLD. *Proteins* **52**, 609–623 (2003).
- 5. Friesner, R. A. *et al.* Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy. *J. Med. Chem.* **47**, 1739–1749 (2004).
- 6. Corbeil, C. R., Williams, C. I. & Labute, P. Variability in docking success rates due to dataset preparation. *J. Comput.-Aided Mol. Des.* **26,** 775–786 (2012).
- 7. Venkatraman, V. & Ritchie, D. W. Flexible protein docking refinement using pose-dependent normal mode analysis. *Proteins* **80**, 2262–2274 (2012).
- 8. Eberhardt, J., Santos-Martins, D., Tillack, A. F. & Forli, S. AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings. *J. Chem. Inf. Model.* **61,** 3891–3898 (2021).
- 9. Koshland, D. E. Application of a Theory of Enzyme Specificity to Protein Synthesis. *Proc. Natl. Acad. Sci. USA* **44**, 98–104 (1958).
- 10. Koshland, D. E. Enzyme flexibility and enzyme action. J. Cell. Physiol. 54, 245–258 (1959).
- 11. Miller, E. B. *et al.* Reliable and accurate solution to the induced fit docking problem for protein–ligand binding. *J. Chem. Theory Comput.* **17**, 2630–2639 (2021).
- 12. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
- 13. Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).
- 14. Krishna, R. *et al.* Generalized biomolecular modeling and design with RoseTTAFold All-Atom. *Science* **384**, eadl2528 (2024).
- 15. Wohlwend, J. *et al.* Boltz-1: Democratizing Biomolecular Interaction Modeling. *bioRxiv*, 2024.11.19.624167 (2024).
- 16. Passaro, S. *et al.* Boltz-2: Towards Accurate and Efficient Binding Affinity Prediction. *bioRxiv*, 2025.06.14.659707 (2025).
- 17. Chai Discovery team *et al.* Chai-1: Decoding the molecular interactions of life. *bioRxiv*, 2024.10.10.615955 (2024).
- 18. ByteDance AML AI4Science Team *et al.* Protenix-advancing structure prediction through a comprehensive AlphaFold3 reproduction. *bioRxiv*, 2025.01.08.631967 (2025).
- 19. Liu, L. *et al.* Technical report of HelixFold3 for biomolecular structure prediction. *arXiv*, 2408.16975 (2024).
- 20. Qiao, Z. et al. NeuralPLexer3: Accurate Biomolecular Complex Structure Prediction with Flow Models in Advances in Neural Information Processing Systems 37 (2024).
- 21. Berman, H. M. et al. The Protein Data Bank. Nucleic Acids Res. 28, 235-242 (2000).
- 22. wwPDB Consortium. *Worldwide Protein Data Bank: Deposition Statistics* https://www.wwpdb.org/stats/deposition. 2025-09-20.
- 23. Skrinjar, P., Eberhardt, J., Durairaj, J. & Schwede, T. Have protein-ligand cofolding methods moved beyond memorisation? *bioRxiv*, 2025.02.03.636309 (2025).

- 24. Wang, Z. *et al.* Comprehensive evaluation of ten docking programs on a diverse set of protein–ligand complexes: the prediction accuracy of sampling power and scoring power. *Phys. Chem. Chem. Phys.* **18**, 12964–12975 (2016).
- 25. Corso, G., Stärk, H., Jing, B., Barzilay, R. & Jaakkola, T. DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking in Proceedings of the 11th International Conference on Learning Representations (2023).
- 26. Lee, S. et al. GenMol: A Drug Discovery Generalist with Discrete Diffusion in Proceedings of the 42nd International Conference on Machine Learning (2025).
- 27. Buttenschoen, M., Morris, G. M. & Deane, C. M. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chem. Sci.* **15**, 3130–3139 (2024).
- 28. Kryshtafovych, A., Schwede, T., Topf, M., Fidelis, K. & Moult, J. Critical assessment of methods of protein structure prediction (CASP)-Round XIV. en. *Proteins* **89**, 1607–1617 (Dec. 2021).
- 29. Baek, M. *et al.* Accurate prediction of protein structures and interactions using a three-track neural network. en. *Science* **373**, 871–876 (Aug. 2021).
- 30. Ho, J., Jain, A. & Abbeel, P. Denoising Diffusion Probabilistic Models in Advances in Neural Information Processing Systems 33 (2020), 6840–6851.
- 31. Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A. & Vandergheynst, P. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Process. Mag.* **34**, 18–42 (2017).
- 32. Burley, S. K. *et al.* Updated resources for exploring experimentally-determined PDB structures and Computed Structure Models at the RCSB Protein Data Bank. *Nucleic Acids Res.* **53**, D564–D574 (2025).
- 33. Melnyk, I., Lozano, A., Das, P. & Chenthamarakshan, V. AlphaFold distillation for inverse protein design. *Sci. Rep.* **15**, 21743 (2025).
- 34. Ahdritz, G. *et al.* OpenFold: Retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization. *Nat. Methods* **21**, 1514–1524 (2024).
- 35. Ouyang-Zhang, J. *et al.* Triangle Multiplication Is All You Need For Biomolecular Structure Representations. *arXiv*, 2510.18870 (2025).
- 36. Varadi, M. *et al.* AlphaFold Protein Structure Database in 2024: providing structure coverage for over 214 million protein sequences. *Nucleic Acids Res.* **52**, D368–D375 (2024).
- 37. NVIDIA. *NVIDIA/cuEquivariance: A collective of low-level primitives and tensor ops to accelerate equivariant neural networks* https://github.com/NVIDIA/cuEquivariance. GitHub repository. Accessed: 2025-10-13. 2024.
- 38. Ba, J. L., Kiros, J. R. & Hinton, G. E. Layer Normalization. *arXiv*, 1607.06450 (2016).
- 39. Gilson, M. *et al.* Assessment of Pharmaceutical Protein–Ligand Pose and Affinity Predictions in CASP16. *Proteins* **0**, 1–9 (2025).
- 40. Chen, M. et al. Evaluating large language models trained on code. arXiv, 2107.03374 (2021).
- 41. Leung, S., Bodkin, M., Von Delft, F., Brennan, P. & Morris, G. SuCOS is better than RMSD for Evaluating Fragment Elaboration and Docking Poses. *ChemRxiv*, 8100203 (2019).
- 42. Robin, X. et al. Assessment of protein-ligand complexes in CASP15. Proteins: Structure, Function, and Bioinformatics **91**, 1811–1821 (2023).
- 43. Efron, B. & Tibshirani, R. J. An Introduction to the Bootstrap (Chapman and Hall/CRC, 1994).

44. Mariani, V., Biasini, M., Barbato, A. & Schwede, T. lDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* **29**, 2722–2728 (2013).

A. Detailed Evaluation Methodology

In this section, we provide full details on our evaluation methodology, including details on the datasets and baselines (Section A.1), inference modes (Section A.2), and evaluation protocol and metrics (Section A.3).

A.1. Datasets and baselines

We evaluate PEARL against multiple cofolding baselines, including AlphaFold 3 [13] and open source descendants such as Boltz-1(x) [15], Boltz-2 [16], Chai-1 [17], and ProteniX [18] across three benchmarks: (i) Runs N' Poses [23] (which we refer to as RnP), a recent public dataset for primary head-to-head comparisons; (ii) PoseBusters [27], a standard benchmark to ensure robust performance; and (iii) InternalXtals, a curated proprietary dataset designed to test models' generalization on data representative of real-world small molecule drug discovery programs.

Release date cutoffs for test structures PEARL has been trained on crystal structures from the PDB released prior to 2021-09-30. AlphaFold 3, Boltz-1(x), Chai-1, and ProteniX all have used release date cutoffs on or before 2021-09-30. However, Boltz-2 used a later cutoff date of 2023-06-01, and therefore can only be properly evaluated on proprietary structures or public structures released after that date. For that reason, we exclude Boltz-2 from the comparison on the PoseBusters dataset, which contains only structures released prior to 2023-06-01. In order to compare other methods against Boltz-2 on the RnP dataset, we focus our evaluation on the subset of structures released after 2023-06-01. Analogously, we subset PoseBusters to structures released on or after 2021-10-01 so that we can fairly compare PEARL with the other baselines. Finally, all our models are allowed to use structures from the PDB as templates only if these structures were released prior to 2021-09-30 (for reference, same template cutoff used by AlphaFold 3 [13]), and more than 60 days prior to the release of the test structure.

Comparison with AlphaFold 3 Due to the restrictive license of AlphaFold 3, we only evaluate PEARL against it on the RnP benchmark, for which a third-party academic group has publicly released generated structures and computed metrics for the poses. For illustrative reasons, we include results in Figure 3 and Tables 1–3 for AlphaFold 3 for PoseBusters, but using published results for poses selected for max confidence out of 25 samples (not best@5).

Baselines and protocol RnP provides publicly released metrics for all baselines in which we are interested; we report results for all these baselines based on this data. RnP did not evaluate these models in the pocket conditional regime, so we do not report performance in this setting for this benchmark. To ensure consistency of our results, for PEARL, we followed the exact inference and evaluation protocols prescribed by the authors of the benchmark [23]. For PoseBusters and InternalXtals, inference has been run for Boltz-1, Boltz-1x, Boltz-2, and Chai-1 to establish baselines.

Processing of RnP released metrics for consistency Even though RnP authors aimed to generate 25 poses per test structure (5 seeds \times 5 samples) for each evaluated cofolding models, they reported that different models failed on some of the structures. In our analysis, we noticed that some models not only failed to produce poses entirely for some of the structures, but the released metrics had an inconsistent number of poses per structure.² To ensure consistency of our analysis, in addition to imposing 2023-06-01+ release date

¹Synthetic data used for training PEARL was also derived only from these publicly available structures with the same release date cutoff. No proprietary experimental data was used in training.

²We found cases where Runs N' Poses released data for each model that contained some structures with fewer than 25 poses generated, ranging from as few as 1−5 poses to often between 20 and 24 poses per structure.

Dataset	Structure release date	# structures
Runs N' Poses	2023-06-01+	702
PoseBusters	2021-10-01+	297
Internal Xtals	proprietary	111

		RnP: pocket-shape similarity distribution								
	- 1								273	
ures	250									
ucti	200									
of structures	150						115	129		
er	100					89				
Number	50	12	18	27	39					
	0 1	 0-20	20-30	30-40	40-50	50-60	60-70	70-80	80-100	_
			Po	cket-s	hape	simila	rity bi	ns		

Figure 10. (Left) statistics of the RnP, PoseBusters, and InternalXtals datasets used for analysis and benchmarking. (Right) Histogram of the number of structures in each pocket-shape similarity bin of the RnP dataset.

cutoff, we processed RnP data in the following manner: (i) we only considered a subset of structures for which all methods had least 20 poses, (ii) we selected the first 20 generated poses (4 seeds \times 5 samples) for each structure for each baseline method. This enabled a fully consistent comparison with PEARL for which we generated 20 poses per structure. This resulting RnP test set contained 702 structures.

Size and composition of the datasets The keys statistics of the final benchmarking sets we used in our analysis (after filtering and processing) are summarized in Figure 10 (left). For each single structure in each of the datasets, we ensured that each cofolding model was able to produce 20 poses for which metrics were computed (for RnP, we used metrics released by the authors). Figure 10 (right) shows the distribution over pocket-shape similarity bins of the final RnP dataset used in our analysis.

A.2. Unconditional and conditional evaluation regimes

To assess flexibility and performance in different scenarios, we evaluate PEARL in two distinct experimental setups: (i) the unconditional cofolding and (ii) the pocket-conditional cofolding modes.

The unconditional cofolding mode tests the model's ability to generate a complex structure starting from only the protein's amino acid sequence and the ligand's 2D structure, as well as MSA information. This evaluation assumes that no prior structural information about the target is available and the binding pocket is unknown. Consequently, any template search must rely on sequence and ligand similarity alone. As the publicly available version of AlphaFold 3 operates in this mode, our primary comparison against it is in this regime.

The conditional cofolding mode is designed to emulate a common scenario in drug discovery, where a reference structure (from public sources or proprietary crystal data) for the target protein is available and a specific binding pocket is hypothesized. This setup evaluates PEARL's controllability and its performance in practical, structure-guided scenarios where a scientist provides a structural prior to guide the generation process. Our evaluation in this regime includes PEARL and the open source baselines that support this type of conditioning: Boltz-1(x), Boltz-2, and Chai-1. More concretely, each of the evaluated models is given access to some of the residues that constitute the binding pocket, which are then provided to the models at inference time. For the purposes of automated benchmarking, the pocket residues must be carefully selected to avoid unfairly leaking information. To do so, we select at most two residues. To select the first, the distances to any residue within 6 Å is calculated for each ligand atom, and the residue with the smallest

median distance is selected. The same procedure is used to find the next best residue that is separated in protein sequence from the first by at least eight residues. The goal is to approximately select two residues from opposing sides of the pocket with the target benchmark ligand reasonably centered between them. For each test complex, the identities of the same ≤ 2 pocket residues were provided to all evaluated models.

A.3. Evaluation protocol and metrics

Our protocol is designed to provide a multifaceted view of model performance.

Pose Accuracy: Our primary metric for structural accuracy is ligand RMSD (also known as BiSyRMSD, or Binding-Site Superposed, Symmetry-Corrected Pose RMSD) [42]. We report the industry standard success rates at specific quality thresholds (*e.g.*, RMSD < 2 Å and RMSD < 1 Å). Our main results are focused on ligand RMSD < 2 Å and RMSD < 1 Å success rates. We report additional metrics in Appendix B.

Pose Quality: To ensure poses are physically plausible, we use PoseBusters' plausibility checks in the "redock" mode using the PoseBusters package version 0.2.9 [27]. This includes the following checks. A pose is considered valid (PB-valid) only if it passes all checks.

- mol_pred_loaded
- mol_true_loaded
- mol_cond_loaded
- sanitization
- all_atoms_connected
- molecular_formula
- molecular_bonds
- double_bond_stereochemistry
- tetrahedral_chirality
- bond lengths
- bond_angles
- internal_steric_clash

- aromatic_ring_flatness
- double_bond_flatness
- internal_energy
- protein-ligand_maximum_distance
- minimum_distance_to_protein
- minimum distance to organic cofactors
- minimum_distance_to_inorganic_cofactors
- minimum_distance_to_waters
- volume_overlap_with_protein
- $\bullet \ \ {\tt volume_overlap_with_organic_cofactors}$
- volume_overlap_with_inorganic_cofactors
- volume_overlap_with_waters

Uncertainty estimates via bootstrapping: We bootstrap [43] with 1000 iterations to estimate the mean and standard error over the complexes in each dataset for all metrics computed in our study. Unless specified otherwise, all error bars on the reported charts and tables correspond to standard error of the mean.

A.4. The best@k evaluation protocol vs pose selection with a confidence model

Given the input information about the system of interest, cofolding models can generate multiple pose candidates. Some of the generated poses can be more accurate and higher quality than others. In cofolding, some of the previous benchmarking efforts computed confidence scores and reported metrics on the highest confidence pose [23]. However, we see *pose generation* and *pose selection* as two different problems, which should not be conflated. Instead of evaluating only the model's top-ranked pose by a confidence score, we adopt the common practice from generative modeling literature and compute metrics for the best@k poses, where best@5 is one of the standards in the literature [40]. In essence, best@k means that k random samples are generated from the model and the best one is selected according to the metric of interest.³

best@k metrics have a couple of important properties:

 $^{^{3}}$ In practice, we pre-generate n=20 poses for each test complex and use an unbiased estimator for best@k as proposed in [40].

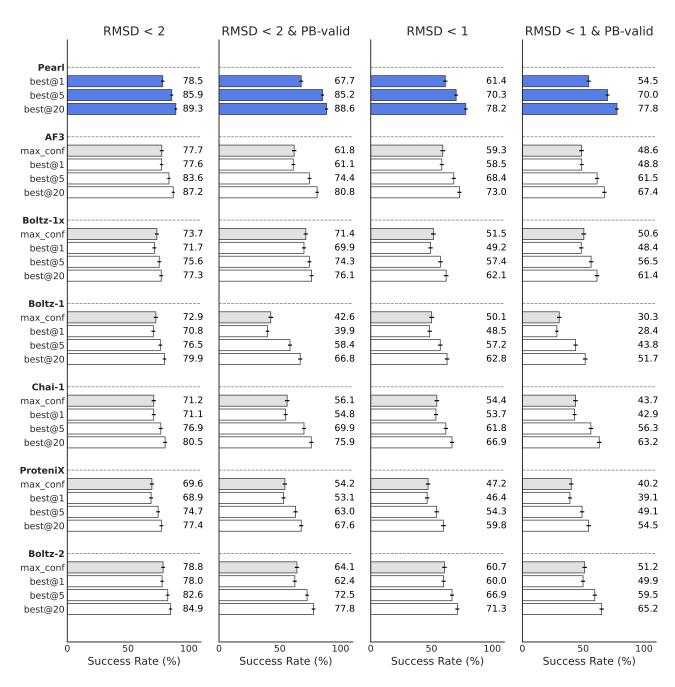


Figure 11. Comparison of the success rates for the best@1, best@5, and best@20 poses versus poses selected with a confidence model (out of 20 sampled) for each of the models on the Runs N' Poses benchmark. The differences between max confidence and randomly selected (best@1) poses are not statistically significant for any of the models. In many cases, PEARL's best@1 pose (*i.e.*, a single random pose) outperforms max confidence poses produced by other models.

- 1. best@1 behaves similar to *precision*. If the metric is a binary indicator of success (e.g., RMSD < 2 Å), best@1 returns the expected probability of success of a random sample generated by the model.
- 2. best@k for larger k values behave similar to recall@k. Again, for a binary indicator of success, best@k returns the expected probability of generating a successful sample after k attempts.

Which k is more important in practice depends on the setting of interest. For example, if a user wanted to guarantee that every pose generated by the model has high probability of success, it's important to optimize for best@1. On the other hand, if a user has access to an appropriate scoring method, and computational cost of inference is not rate-limiting, one should optimize for the probability to "recall" a high quality pose given enough attempts, and use best@k with larger k. Our main results focus on best@5 evaluation, which is a standard metric in the broader generative modeling literature, and it strikes a good balance between best@1 and best@20.

Comparison of best@k with pose selection using confidence model To understand the differences between best@1, best@5, best@20, and top-ranked poses selected by confidence (out of 20 generated), we computed ligand RMSD metrics for each of these pose aggregation methods on the RnP dataset (which includes confidence scores for each pose for each baseline method). The results are presented in Figure 11. Surprisingly, we found that confidence scores are not particularly informative for ranking poses of protein-ligand interfaces: the results show that the top-ranked pose by confidence score is no better than a randomly selected pose on (best@1); the difference between max confidence and best@1 is not statistically significant for any of the models. Our finding is consistent with the observations made in [23, Section 2.5]. We also note that random poses (best@1) generated by PEARL in many cases are superior compared to max confidence poses produced by other models.

B. Extended results

Below, we present additional results for PEARL and baseline cofolding models on the RnP, PoseBusters, and InternalXtals benchmarks and provide further analysis and discussion.

B.1. Comprehensive metrics

Tables 1, 2, and 3 present the extended set of metrics for best@1, best@5, and best@20 pose aggregation, respectively. For each pose aggregation level, we provide detailed results across inference modes (unconditional and pocket-conditional), datasets (Runs N' Poses, PoseBusters, and InternalXtals), and metrics. All metrics are success rate percentages where higher values are better. The highest value across the methods for each dataset, pose aggregation, and inference regime is bolded.

On Runs N' Poses, we use non-PEARL metrics from the structures and poses released by the authors [23], who did not report any metrics for pocket-conditional cofolding.

On PoseBusters, we use AlphaFold 3's own released structures and metrics in the unconditional cofolding mode; we run inference and calculate metrics ourselves for PEARL and other methods. We omit Boltz-2 from PoseBusters because its structures fall within Boltz-2's training window.

For InternalXtals, we run inference and calculate metrics for PEARL and other methods.

We separate Boltz-2 from the other methods and denote it with a † due to its later 2023-06-01 training cutoff date, which gives it an advantage over the other baselines trained only on structures released up until

Table 1. Model performance at best@1 aggregation.

Part			Table 1. Model p	- CITOIIII and C	at bob 061 455	regution.		
Runs N' Poses Boltz-1x 71.7 ± 1.6 69.9 ± 1.7 49.2 ± 1.8 48.4 ± 1.8 76.9 ± 1.0	Mode	Dataset	Method	RMSD < 2		RMSD < 1		lDDT-PLI
Runs N' Poses			PEARL	79.0 ± 1.4	67.9 ± 1.6	61.8 ± 1.7	55.0 ± 1.7	81.9 ± 0.9
Runs N' Poses Boltz-1 70.8 ± 1.6 39.9 ± 1.5 48.5 ± 1.8 28.4 ± 1.4 76.5 ± 0.9 ProteniX 68.9 ± 1.6 53.7 ± 1.7 42.9 ± 1.6 76.5 ± 1.0 Boltz-2 [†] (2023-06) 78.0 ± 1.4 62.4 ± 1.6 60.0 ± 1.7 49.9 ± 1.7 80.6 ± 0.9 PEARL 79.3 ± 2.3 65.4 ± 2.7 62.4 ± 2.7 54.7 ± 2.7 82.1 ± 1.4 AF3 * 76.5 ± 2.5 60.4 ± 2.8 62.1 ± 2.8 52.7 ± 2.9 — PoseBusters Boltz-1x 66.5 ± 2.5 64.1 ± 2.5 41.4 ± 2.5 40.2 ± 2.4 71.9 ± 1.5 Boltz-1 66.6 ± 2.6 41.6 ± 2.5 43.8 ± 2.7 30.3 ± 2.4 72.9 ± 1.5 Chai-1 69.6 ± 2.4 53.7 ± 2.5 51.5 ± 2.5 42.7 ± 2.4 74.7 ± 1.6 PEARL 63.2 ± 4.4 40.2 ± 4.1 29.5 ± 3.5 22.0 ± 3.3 63.8 ± 3.7 Boltz-1x 33.1 ± 4.0 28.6 ± 3.7 3.8 ± 1.2 3.6 ± 1.2 53.8 ± 3.2 Boltz-2 [†] (2023-06) 59.3 ± 4.3 29.2 ± 3.9 12.9 ± 2.6 86.5 ± 2.2 62.3 ± 3.6 Boltz-2 [†] (2023-06) 59.3 ± 4.3 29.2 ± 3.9 12.9 ± 2.6 86.5 ± 2.2 62.3 ± 3.6 PEARL Boltz-1x 69.4 ± 2.5 67.0 ± 2.5 45.1 ± 2.5 43.9 ± 2.4 75.4 ± 1.2 Boltz-1 69.6 ± 2.6 43.9 ± 2.6 47.6 ± 2.7 32.5 ± 2.4 76.4 ± 1.2 Boltz-1 69.6 ± 2.6 43.9 ± 2.6 47.6 ± 2.7 32.5 ± 2.4 76.4 ± 1.2 Chai-1 72.8 ± 2.3 56.1 ± 2.5 55.1 ± 2.5 45.6 ± 2.4 77.9 ± 1.3 PEARL Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1			AF3	77.6 ± 1.4	61.1 ± 1.5	58.5 ± 1.7	48.8 ± 1.6	80.9 ± 0.9
Protein Prot			Boltz-1x	71.7 ± 1.6	69.9 ± 1.7	49.2 ± 1.8	48.4 ± 1.8	76.9 ± 1.0
Protenix 68.9 ± 1.6 53.1 ± 1.6 46.4 ± 1.8 39.1 ± 1.6 76.0 ± 0.9		Runs N' Poses	Boltz-1	70.8 ± 1.6	39.9 ± 1.5	48.5 ± 1.8	28.4 ± 1.4	76.5 ± 0.9
Boltz-2 [†] (2023-06) 78.0 ± 1.4 62.4 ± 1.6 60.0 ± 1.7 49.9 ± 1.7 80.6 ± 0.9			Chai-1	71.1 ± 1.6	54.8 ± 1.6	53.7 ± 1.7	42.9 ± 1.6	76.5 ± 1.0
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$			ProteniX	68.9 ± 1.6	53.1 ± 1.6	46.4 ± 1.8	39.1 ± 1.6	76.0 ± 0.9
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	onal		Boltz-2 [†] (2023-06)	78.0 ± 1.4	62.4 ± 1.6	60.0 ± 1.7	49.9 ± 1.7	80.6 ± 0.9
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	ditio		PEARL	79.3 ± 2.3	65.4 ± 2.7	62.4 ± 2.7	54.7 ± 2.7	82.1 ± 1.4
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	иоз		AF3 *	76.5 ± 2.5	60.4 ± 2.8	62.1 ± 2.8	52.7 ± 2.9	_
$ \begin{array}{ c c c c c c c c c } \hline & Chai-1 & 69.6 \pm 2.4 & 53.7 \pm 2.5 & 51.5 \pm 2.5 & 42.7 \pm 2.4 & 74.7 \pm 1.6 \\ \hline & PEARL & 63.2 \pm 4.4 & 40.2 \pm 4.1 & 29.5 \pm 3.5 & 22.0 \pm 3.3 & 63.8 \pm 3.7 \\ \hline & Boltz-1x & 33.1 \pm 4.0 & 28.6 \pm 3.7 & 3.8 \pm 1.2 & 3.6 \pm 1.2 & 53.8 \pm 3.2 \\ \hline & Boltz-1 & 32.8 \pm 4.1 & 7.5 \pm 1.9 & 4.8 \pm 1.5 & 1.4 \pm 0.6 & 54.4 \pm 3.3 \\ \hline & Chai-1 & 15.5 \pm 3.3 & 5.5 \pm 2.0 & 5.4 \pm 1.7 & 1.6 \pm 0.8 & 38.7 \pm 2.8 \\ \hline & Boltz-2^{\dagger} (2023-06) & 59.3 \pm 4.3 & 29.2 \pm 3.9 & 12.9 \pm 2.6 & 8.6 \pm 2.2 & 62.3 \pm 3.6 \\ \hline & PoseBusters & & PEARL & 80.6 \pm 2.2 & 66.8 \pm 2.6 & 62.9 \pm 2.6 & 55.2 \pm 2.7 & 84.1 \pm 1.1 \\ \hline & Boltz-1x & 69.4 \pm 2.5 & 67.0 \pm 2.5 & 45.1 \pm 2.5 & 43.9 \pm 2.4 & 75.4 \pm 1.2 \\ \hline & Boltz-1 & 69.6 \pm 2.6 & 43.9 \pm 2.6 & 47.6 \pm 2.7 & 32.5 \pm 2.4 & 76.4 \pm 1.2 \\ \hline & Chai-1 & 72.8 \pm 2.3 & 56.1 \pm 2.5 & 55.1 \pm 2.5 & 45.6 \pm 2.4 & 77.9 \pm 1.3 \\ \hline & PEARL & 68.0 \pm 4.1 & 42.6 \pm 4.0 & 30.7 \pm 3.6 & 23.6 \pm 3.4 & 75.1 \pm 2.3 \\ \hline & Boltz-1x & 42.9 \pm 4.0 & 37.8 \pm 3.8 & 5.2 \pm 1.6 & 5.0 \pm 1.6 & 62.1 \pm 3.1 \\ \hline & Boltz-1x & 42.9 \pm 4.0 & 37.8 \pm 3.8 & 5.2 \pm 1.6 & 5.0 \pm 1.6 & 62.1 \pm 3.1 \\ \hline & Boltz-1x & 42.9 \pm 4.0 & 37.8 \pm 3.8 & 5.2 \pm 1.6 & 5.0 \pm 1.6 & 62.1 \pm 3.1 \\ \hline & Chai-1 & 25.0 \pm 3.7 & 9.2 \pm 2.2 & 6.1 \pm 1.7 & 1.9 \pm 0.8 & 54.3 \pm 2.5 \\ \hline \end{array}$	Una	PoseBusters	Boltz-1x	66.5 ± 2.5	64.1 ± 2.5	41.4 ± 2.5	40.2 ± 2.4	71.9 ± 1.5
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$			Boltz-1	66.6 ± 2.6	41.6 ± 2.5	43.8 ± 2.7	30.3 ± 2.4	72.9 ± 1.5
InternalXtals Boltz-1x 33.1 ± 4.0 28.6 ± 3.7 3.8 ± 1.2 3.6 ± 1.2 53.8 ± 3.2 32.8 ± 4.1 7.5 ± 1.9 4.8 ± 1.5 1.4 ± 0.6 54.4 ± 3.3 15.5 ± 3.3 15.5 ± 3.3 15.5 ± 2.0 15.4 ± 1.7 15.6 ± 0.8 15.4 ± 0.8 15.5 ± 3.8 15.5 ± 2.0 15.4 ± 1.7 15.6 ± 0.8 15.5 ± 2.8 Boltz-2† (2023-06) 15.5 ± 2.3 12.9 ± 2.6 15.5 ± 2.6 15.5 ± 2.7 15.5 ± 2.8 Boltz-1x 15.5 ± 2.8 Boltz-1x 15.5 ± 2.8 Boltz-1x 15.5 ± 2.8 Boltz-1			Chai-1	69.6 ± 2.4	53.7 ± 2.5	51.5 ± 2.5	42.7 ± 2.4	74.7 ± 1.6
InternalXtals Boltz-1 32.8 ± 4.1 7.5 ± 1.9 4.8 ± 1.5 1.4 ± 0.6 54.4 ± 3.3 15.5 ± 3.3 5.5 ± 2.0 5.4 ± 1.7 1.6 ± 0.8 38.7 ± 2.8 Boltz-2 [†] (2023-06) 59.3 ± 4.3 29.2 ± 3.9 12.9 ± 2.6 8.6 ± 2.2 62.3 ± 3.6 PEARL Boltz-1x 69.4 ± 2.5 67.0 ± 2.5 45.1 ± 2.5 43.9 ± 2.4 75.4 ± 1.2 Boltz-1 69.6 ± 2.6 43.9 ± 2.6 47.6 ± 2.7 32.5 ± 2.4 76.4 ± 1.2 Chai-1 72.8 ± 2.3 56.1 ± 2.5 55.1 ± 2.5 45.6 ± 2.4 77.9 ± 1.3 PEARL Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 InternalXtals Boltz-1 43.5 ± 4.2 10.0 ± 2.0 6.6 ± 1.8 1.8 ± 0.7 63.8 ± 3.0 Chai-1 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5		InternalXtals	PEARL	63.2 ± 4.4	40.2 ± 4.1	29.5 ± 3.5	22.0 ± 3.3	63.8 ± 3.7
$ \begin{array}{ c c c c c c c c c } \hline & Chai-1 & 15.5 \pm 3.3 & 5.5 \pm 2.0 & 5.4 \pm 1.7 & 1.6 \pm 0.8 & 38.7 \pm 2.8 \\ \hline \hline & Boltz-2^{\dagger} \left(2023\text{-}06\right) & 59.3 \pm 4.3 & 29.2 \pm 3.9 & 12.9 \pm 2.6 & 8.6 \pm 2.2 & 62.3 \pm 3.6 \\ \hline & PEARL & \textbf{80.6} \pm 2.2 & 66.8 \pm 2.6 & \textbf{62.9} \pm 2.6 & \textbf{55.2} \pm 2.7 & \textbf{84.1} \pm 1.1 \\ \hline & Boltz-1x & 69.4 \pm 2.5 & \textbf{67.0} \pm 2.5 & 45.1 \pm 2.5 & 43.9 \pm 2.4 & 75.4 \pm 1.2 \\ \hline & Boltz-1 & 69.6 \pm 2.6 & 43.9 \pm 2.6 & 47.6 \pm 2.7 & 32.5 \pm 2.4 & 76.4 \pm 1.2 \\ \hline & Chai-1 & 72.8 \pm 2.3 & 56.1 \pm 2.5 & 55.1 \pm 2.5 & 45.6 \pm 2.4 & 77.9 \pm 1.3 \\ \hline & PEARL & \textbf{68.0} \pm \textbf{4.1} & \textbf{42.6} \pm \textbf{4.0} & \textbf{30.7} \pm \textbf{3.6} & \textbf{23.6} \pm \textbf{3.4} & \textbf{75.1} \pm \textbf{2.3} \\ \hline & Boltz-1x & 42.9 \pm 4.0 & 37.8 \pm 3.8 & 5.2 \pm 1.6 & 5.0 \pm 1.6 & 62.1 \pm 3.1 \\ \hline & Boltz-1 & 43.5 \pm 4.2 & 10.0 \pm 2.0 & 6.6 \pm 1.8 & 1.8 \pm 0.7 & 63.8 \pm 3.0 \\ \hline & Chai-1 & 25.0 \pm 3.7 & 9.2 \pm 2.2 & 6.1 \pm 1.7 & 1.9 \pm 0.8 & 54.3 \pm 2.5 \\ \hline \end{array}$			Boltz-1x	33.1 ± 4.0	28.6 ± 3.7	3.8 ± 1.2	3.6 ± 1.2	53.8 ± 3.2
			Boltz-1	32.8 ± 4.1	7.5 ± 1.9	4.8 ± 1.5	1.4 ± 0.6	54.4 ± 3.3
PoseBusters PoseBusters $\begin{pmatrix} 80.6 \pm 2.2 & 66.8 \pm 2.6 & 62.9 \pm 2.6 & 55.2 \pm 2.7 & 84.1 \pm 1.1 \\ Boltz-1x & 69.4 \pm 2.5 & 67.0 \pm 2.5 & 45.1 \pm 2.5 & 43.9 \pm 2.4 & 75.4 \pm 1.2 \\ 69.6 \pm 2.6 & 43.9 \pm 2.6 & 47.6 \pm 2.7 & 32.5 \pm 2.4 & 76.4 \pm 1.2 \\ Chai-1 & 72.8 \pm 2.3 & 56.1 \pm 2.5 & 55.1 \pm 2.5 & 45.6 \pm 2.4 & 77.9 \pm 1.3 \\ PEARL & 68.0 \pm 4.1 & 42.6 \pm 4.0 & 30.7 \pm 3.6 & 23.6 \pm 3.4 & 75.1 \pm 2.3 \\ Boltz-1x & 42.9 \pm 4.0 & 37.8 \pm 3.8 & 5.2 \pm 1.6 & 5.0 \pm 1.6 & 62.1 \pm 3.1 \\ Boltz-1 & 43.5 \pm 4.2 & 10.0 \pm 2.0 & 6.6 \pm 1.8 & 1.8 \pm 0.7 & 63.8 \pm 3.0 \\ Chai-1 & 25.0 \pm 3.7 & 9.2 \pm 2.2 & 6.1 \pm 1.7 & 1.9 \pm 0.8 & 54.3 \pm 2.5 \\ \end{pmatrix}$			Chai-1	15.5 ± 3.3	5.5 ± 2.0	5.4 ± 1.7	1.6 ± 0.8	38.7 ± 2.8
PoseBusters			Boltz-2 [†] (2023-06)	59.3 ± 4.3	29.2 ± 3.9	12.9 ± 2.6	8.6 ± 2.2	62.3 ± 3.6
PoseBusters Boltz-1 69.6 ± 2.6 43.9 ± 2.6 47.6 ± 2.7 32.5 ± 2.4 76.4 ± 1.2 2.8 ± 2.3 2.8 ± 2.3 2.8 ± 2.5 2.8			PEARL	80.6 ± 2.2	66.8 ± 2.6	62.9 ± 2.6	55.2 ± 2.7	84.1 ± 1.1
Boltz-1 69.6 ± 2.6 43.9 ± 2.6 47.6 ± 2.7 32.5 ± 2.4 76.4 ± 1.2 72.8 ± 2.3 56.1 ± 2.5 55.1 ± 2.5 45.6 ± 2.4 77.9 ± 1.3 PEARL Boltz-1x 42.9 ± 4.0 37.8 ± 3.8 5.2 ± 1.6 5.0 ± 1.6 62.1 ± 3.1 InternalXtals Boltz-1 43.5 ± 4.2 10.0 ± 2.0 6.6 ± 1.8 1.8 ± 0.7 63.8 ± 3.0 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5		PocoPuctors	Boltz-1x	69.4 ± 2.5	67.0 ± 2.5	45.1 ± 2.5	43.9 ± 2.4	75.4 ± 1.2
InternalXtals Boltz-1 43.5 ± 4.2 10.0 ± 2.0 6.6 ± 1.8 1.8 ± 0.7 63.8 ± 3.0 Chai-1 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5	n	rosebusiers	Boltz-1	69.6 ± 2.6	43.9 ± 2.6	47.6 ± 2.7	32.5 ± 2.4	76.4 ± 1.2
InternalXtals Boltz-1 43.5 ± 4.2 10.0 ± 2.0 6.6 ± 1.8 1.8 ± 0.7 63.8 ± 3.0 Chai-1 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5	ionc		Chai-1	72.8 ± 2.3	56.1 ± 2.5	55.1 ± 2.5	45.6 ± 2.4	77.9 ± 1.3
InternalXtals Boltz-1 43.5 ± 4.2 10.0 ± 2.0 6.6 ± 1.8 1.8 ± 0.7 63.8 ± 3.0 Chai-1 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5	nditi		PEARL	68.0 ± 4.1	42.6 ± 4.0	30.7 ± 3.6	23.6 ± 3.4	75.1 ± 2.3
Chai-1 25.0 ± 3.7 9.2 ± 2.2 6.1 ± 1.7 1.9 ± 0.8 54.3 ± 2.5	C_{OJ}		Boltz-1x	42.9 ± 4.0	37.8 ± 3.8	5.2 ± 1.6	5.0 ± 1.6	62.1 ± 3.1
		InternalXtals		43.5 ± 4.2	10.0 ± 2.0		1.8 ± 0.7	63.8 ± 3.0
Boltz-2 [†] (2023-06) 60.4 ± 4.3 30.9 ± 3.9 14.2 ± 2.7 9.8 ± 2.3 69.7 ± 2.6			Chai-1	25.0 ± 3.7	9.2 ± 2.2	6.1 ± 1.7	1.9 ± 0.8	54.3 ± 2.5
			Boltz-2 [†] (2023-06)	60.4 ± 4.3	30.9 ± 3.9	14.2 ± 2.7	9.8 ± 2.3	69.7 ± 2.6

^{*} AlphaFold 3 results for PoseBusters use officially released metrics for poses selected for max confidence out of 25 samples (not best@1).

[†] Since it is trained with data up to 2023-06-01, Boltz-2 is not directly comparable to the other evaluated models, which in contrast all use earlier training cutoffs of \leq 2021-09-30.

Table 2. Model performance at best@5 aggregation.

		Table 2. Model p			1080010111	DMCD <1	
Mode	Dataset	Method	RMSD < 2	RMSD < 2 & PB-valid	RMSD < 1	RMSD < 1 & PB-valid	lDDT-PLI
		PEARL	85.9 ± 1.2	85.2 ± 1.2	70.3 ± 1.5	70.0 ± 1.6	$\textbf{84.6} \pm \textbf{0.7}$
		AF3	83.6 ± 1.3	74.4 ± 1.5	68.4 ± 1.7	61.5 ± 1.7	84.4 ± 0.8
		Boltz-1x	75.6 ± 1.6	74.3 ± 1.6	57.4 ± 1.8	56.5 ± 1.8	79.5 ± 0.9
	Runs N' Poses	Boltz-1	76.5 ± 1.5	58.4 ± 1.7	57.2 ± 1.8	43.8 ± 1.7	79.9 ± 0.9
		Chai-1	76.9 ± 1.5	69.9 ± 1.6	61.8 ± 1.7	56.3 ± 1.8	80.1 ± 1.0
		ProteniX	74.7 ± 1.6	63.0 ± 1.7	54.3 ± 1.8	49.1 ± 1.8	79.5 ± 0.9
Unconditional		Boltz-2 [†] (2023-06)	82.6 ± 1.3	72.5 ± 1.5	66.9 ± 1.7	59.5 ± 1.7	83.4 ± 0.8
diti		PEARL	85.1 ± 2.0	$\textbf{84.7} \pm \textbf{2.0}$	72.7 ± 2.4	$\textbf{72.4} \pm \textbf{2.4}$	84.9 ± 1.1
иоз		AF3 *	76.5 ± 2.5	60.4 ± 2.8	62.1 ± 2.8	52.7 ± 2.9	_
U_{n}	PoseBusters	Boltz-1x	74.8 ± 2.4	74.2 ± 2.4	55.6 ± 2.7	55.2 ± 2.7	77.7 ± 1.3
		Boltz-1	73.2 ± 2.5	54.5 ± 2.8	54.0 ± 2.8	41.9 ± 2.7	77.5 ± 1.4
		Chai-1	77.8 ± 2.3	68.7 ± 2.5	63.6 ± 2.5	58.0 ± 2.6	79.8 ± 1.4
	InternalXtals	PEARL	66.1 ± 4.4	62.0 ± 4.4	37.2 ± 3.7	36.0 ± 3.7	64.3 ± 3.5
		Boltz-1x	39.4 ± 4.4	36.6 ± 4.3	9.9 ± 2.3	9.3 ± 2.3	55.6 ± 3.3
		Boltz-1	37.6 ± 4.3	14.5 ± 3.0	9.9 ± 2.5	4.2 ± 1.5	55.9 ± 3.3
		Chai-1	17.1 ± 3.5	6.7 ± 2.2	9.5 ± 2.6	3.2 ± 1.6	40.9 ± 2.8
		Boltz-2 [†] (2023-06)	64.9 ± 4.4	38.0 ± 4.3	22.3 ± 3.4	14.6 ± 2.9	63.8 ± 3.6
		PEARL	86.7 ± 1.8	$\textbf{86.7} \pm \textbf{1.8}$	72.6 ± 2.4	$\textbf{72.2} \pm \textbf{2.4}$	86.2 ± 0.9
	PoseBusters	Boltz-1x	76.8 ± 2.3	76.3 ± 2.3	59.5 ± 2.7	59.4 ± 2.7	80.5 ± 1.1
rı	rosedusieis	Boltz-1	75.6 ± 2.4	57.1 ± 2.8	57.8 ± 2.8	44.7 ± 2.7	80.5 ± 1.1
оио		Chai-1	81.4 ± 2.1	70.7 ± 2.4	67.7 ± 2.5	60.6 ± 2.6	82.9 ± 1.2
Conditional		PEARL	81.4 ± 3.3	73.9 ± 3.8	41.1 ± 3.7	39.8 ± 3.8	81.4 ± 1.4
		Boltz-1x	54.9 ± 4.4	51.5 ± 4.4	11.2 ± 2.5	10.6 ± 2.5	66.2 ± 3.0
=	InternalXtals	Boltz-1	52.3 ± 4.3	21.9 ± 3.5	13.0 ± 2.8	5.4 ± 1.7	67.7 ± 2.9
		Chai-1	32.3 ± 4.1	17.2 ± 3.1	11.9 ± 2.9	4.5 ± 1.8	60.5 ± 2.4
		Boltz-2 [†] (2023-06)	66.8 ± 4.3	41.2 ± 4.4	25.0 ± 3.6	18.1 ± 3.2	74.0 ± 2.3

^{*} AlphaFold 3 results for PoseBusters use officially released metrics for poses selected for max confidence out of 25 samples (not best@5).

[†] Since it is trained with data up to 2023-06-01, Boltz-2 is not directly comparable to the other evaluated models, which in contrast all use earlier training cutoffs of \leq 2021-09-30.

Table 3. Model performance at best@20 aggregation.

		Table 3. Woder p			5 8		
Mode	Dataset	Method	RMSD < 2	RMSD < 2 & PB-valid	RMSD < 1	RMSD < 1 & PB-valid	lDDT-PLI
		PEARL	89.3 ± 1.2	88.6 ± 1.2	78.2 ± 1.5	77.8 ± 1.5	87.4 ± 0.6
		AF3	87.2 ± 1.2	80.8 ± 1.5	73.0 ± 1.7	67.4 ± 1.8	86.3 ± 0.8
		Boltz-1x	77.3 ± 1.6	76.1 ± 1.6	62.1 ± 1.9	61.4 ± 1.9	80.8 ± 0.9
	Runs N' Poses	Boltz-1	79.9 ± 1.5	66.8 ± 1.8	62.8 ± 1.9	51.7 ± 1.9	81.8 ± 0.8
		Chai-1	80.5 ± 1.5	75.9 ± 1.6	66.9 ± 1.8	63.2 ± 1.9	82.2 ± 0.9
		ProteniX	77.4 ± 1.6	67.6 ± 1.7	59.8 ± 1.9	54.5 ± 1.9	81.6 ± 0.9
Unconditional		Boltz-2 [†] (2023-06)	84.9 ± 1.3	77.8 ± 1.5	71.3 ± 1.7	65.2 ± 1.8	85.2 ± 0.8
diti		PEARL	88.3 ± 1.9	88.3 ± 1.9	79.8 ± 2.4	79.1 ± 2.4	87.9 ± 1.0
иоз		AF3 *	76.5 ± 2.5	60.4 ± 2.8	62.1 ± 2.8	52.7 ± 2.9	_
Un_0	PoseBusters	Boltz-1x	80.1 ± 2.3	80.1 ± 2.3	65.8 ± 2.8	65.8 ± 2.8	81.1 ± 1.2
		Boltz-1	78.0 ± 2.5	62.0 ± 2.9	61.7 ± 2.9	50.1 ± 3.0	80.5 ± 1.3
		Chai-1	81.1 ± 2.3	76.0 ± 2.5	69.9 ± 2.7	65.8 ± 2.8	82.1 ± 1.4
	InternalXtals	PEARL	67.7 ± 4.4	64.1 ± 4.5	52.4 ± 4.7	50.6 ± 4.7	66.9 ± 3.5
		Boltz-1x	43.5 ± 4.6	40.8 ± 4.6	19.1 ± 3.8	17.3 ± 3.6	56.8 ± 3.4
		Boltz-1	43.5 ± 4.6	21.0 ± 3.8	14.6 ± 3.4	8.30 ± 2.6	57.0 ± 3.4
		Chai-1	20.0 ± 3.8	9.10 ± 2.8	12.8 ± 3.2	4.60 ± 2.0	42.8 ± 2.9
		Boltz-2 [†] (2023-06)	65.9 ± 4.4	46.0 ± 4.7	34.3 ± 4.5	23.4 ± 4.0	64.6 ± 3.7
		PEARL	91.1 ± 1.6	91.1 ± 1.6	79.8 ± 2.4	79.4 ± 2.4	89.1 ± 0.8
	PoseBusters	Boltz-1x	83.5 ± 2.2	83.5 ± 2.2	67.2 ± 2.8	67.2 ± 2.8	83.6 ± 1.0
rı	rosedusieis	Boltz-1	80.8 ± 2.3	64.1 ± 2.8	63.4 ± 2.9	51.8 ± 2.9	82.9 ± 1.1
оио		Chai-1	84.5 ± 2.1	76.8 ± 2.4	73.9 ± 2.5	65.8 ± 2.7	85.1 ± 1.1
Conditional		PEARL	88.3 ± 3.0	79.2 ± 3.9	60.4 ± 4.7	57.7 ± 4.8	85.9 ± 1.1
		Boltz-1x	60.4 ± 4.6	57.7 ± 4.7	19.9 ± 3.8	18.0 ± 3.7	68.8 ± 2.9
	InternalXtals	Boltz-1	59.6 ± 4.6	29.9 ± 4.3	20.0 ± 3.9	10.1 ± 2.9	70.4 ± 2.7
		Chai-1	39.1 ± 4.6	25.4 ± 4.1	15.5 ± 3.5	5.6 ± 2.2	62.7 ± 2.4
		Boltz-2 [†] (2023-06)	69.5 ± 4.3	48.8 ± 4.8	35.3 ± 4.5	26.4 ± 4.2	76.8 ± 2.1

^{*} AlphaFold 3 results for PoseBusters use officially released metrics for poses selected for max confidence out of 25 samples (not best@20).

[†] Since it is trained with data up to 2023-06-01, Boltz-2 is not directly comparable to the other evaluated models, which in contrast all use earlier training cutoffs of \leq 2021-09-30.

2021-09-30. PoseBusters results for AF3 shown in each of the tables are based on the officially released metrics for the top pose selected by the confidence score (*i.e.*, not computed using best@k aggregation).

B.2. Mixed Precision

Training at full precision (fp32) can unnecessarily limit model size due to memory and speed constraints. A common strategy in large-scale language and protein structure models alike is to leverage half-precision arithmetic, most frequently using mixed-precision training with bfloat16 (bf16). Several cofolding models such as ProteniX [18] and OpenFold [34] adopt this approach. Protenix adopts half-precision training aggressively, training nearly the entire model at bf16 (with a custom CUDA LayerNorm) but reverting to fp32 for inference. Boltz-2 instead applies bf16 only to the trunk, keeping the atom attention encoder and diffusion module in fp32. Our approach is to improve scalability through a conservative mixed-precision implementation designed to balance computational efficiency with numerical stability.

For training, PEARL executes most of the computationally heavy trunk module in bfloat16 (bf16) precision. However, numerically sensitive components, such as the loss calculations and final coordinate projections, are kept in fp32 to ensure training stability. This strategy yields significant efficiency gains while avoiding the numerical instability issues faced by more aggressive half-precision approaches.

Specifically, we automatically upcast numerically unstable operations (e.g. softmax, exp, log) to fp32, while computationally heavy operations (e.g. matrix multiplications, linear projections, element-wise activations) are executed in bf16. All model weights are stored in fp32, ensuring compatibility with both bf16 and fp32 inference. PEARL combines these approaches: most of the trunk (triangle operations, via NVIDIA cuEquivariance kernels, and LayerNorm via a custom CUDA kernel) is executed in bf16, while numerically sensitive modules (e.g., conditioned transitions, input and output coordinate projections, losses) are kept in fp32. This conservative strategy avoids instable activations, while still yielding significant efficiency gains.

C. Glossary

Term / Abbreviation	Description
AF2	AlphaFold 2 [12]
AF3	AlphaFold 3 [13]
AFDB	AlphaFold Database [36]
Å	Angstrom, 10^{-10} m
apo	A protein conformation in the absence of a binding cofactor
bf16	Representation standard for floating point numbers using 16 bits
EqT	Equivariant transformer
Equivariance	The property of a function $f(x)$ such that a transformation of its input results in
	an equivalent transformation of its output
fp32	Representation standard for floating point numbers using 32 bits
holo	A ligand-bound protein conformational state
lDDT-PLI	Local difference distance test [44] of the protein-ligand interface
MSA	Multiple Sequence Alignment
PDB	Protein Data Bank [21]
RMSD	Root mean square deviation
SAR	Structure-activity relationship
SBDD	Structure-based drug design
SO(3)	The group of all rotations about the origin in 3D space