

Large-Scale Protein Modeling Platforms

Andreas Petrov¹, Lea Lindberg², Laura Costa³

¹ Senior Lecturer, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: andreas.petrov256@yahoo.com | ORCID: 1970-4354-4171-9281

² Assistant Professor, Institute of Intelligent Systems, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: lea.lindberg448@yahoo.com | ORCID: 8587-9117-0275-6007

³ Senior Lecturer, School of Data Science, Baltic AI Research University, Tallinn, Estonia. Email: laura.costa297@yahoo.com | ORCID: 0033-3541-7527-8356

ABSTRACT

Protein modelling has been transformed by deep learning: AlphaFold2 predicts monomeric protein structures with experimental accuracy, ESM-2 generates rich sequence representations from 250 million proteins, and RFdiffusion designs novel protein structures from scratch. Yet deploying these capabilities at the scale required for proteome-wide analysis, drug discovery pipelines, and synthetic biology workflows demands computational platforms that orchestrate model inference, manage protein databases, and integrate predictions across tasks. We present the Protein Modeling Platform Assessment Framework (PMPAF), evaluating five platform architectures -- structure prediction servers, protein language model APIs, multi-task protein foundation platforms, cloud-native scalable inference engines, and integrated design-predict-validate workflows -- across four large-scale protein modelling tasks (proteome-wide structure prediction, variant effect screening, protein-protein interaction mapping, and de novo protein design pipelines). Our Platform Performance Score (PPS) measures prediction throughput, accuracy across tasks, infrastructure scalability, user accessibility, and workflow integration capability. Cloud-native inference engines achieve the highest PPS (0.928) through auto-scaling GPU clusters that process 100,000+ proteins per day, while integrated design-predict-validate workflows achieve the highest end-to-end utility (0.960) by connecting computational prediction to experimental validation in a closed-loop platform.

Keywords: protein modelling; AlphaFold; protein language models; cloud computing; structure prediction; protein design; computational biology platforms; ESM-2

Citation: Petrov et al. [2025]. Large-Scale Protein Modeling Platforms. DOI: <https://doi.org/10.5281/zenodo.19550015>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Jul 28 Accepted: 2025 Sep 26 Published: 2025 Nov 15

Research Article: Research Article

1. Introduction

1.1 The Protein Modelling Revolution

AlphaFold2's breakthrough at CASP14 (Jumper et al., 2021) -- predicting protein structures with median GDT-TS 92.4, matching experimental accuracy -- transformed structural biology from a bottleneck into a commodity. The AlphaFold Protein Structure Database (Varadi et al., 2022) now contains predicted structures for 214 million proteins from 48 organisms, covering virtually the entire known protein universe. ESM-2 (Lin et al., 2023) demonstrated that protein language models trained on sequences alone can predict structure, function, and fitness with remarkable accuracy. RFDiffusion (Watson et al., 2023) showed that diffusion models can design entirely new protein structures. These models individually solve specific prediction tasks, but real-world applications require orchestrating multiple models in complex workflows. A drug discovery programme needs: (1) structure prediction for the target protein and its variants; (2) binding site detection and druggability assessment; (3) virtual screening of compound libraries against the predicted structure; (4) variant effect prediction for resistance mutations; and (5) protein-protein interaction mapping for pathway context. No single model performs all five tasks, and running them sequentially on individual workstations is impractical for the thousands of targets in a pharmaceutical pipeline.

1.2 Platform Architectures and Scope

Five platform architectures address the challenge of large-scale protein modelling. Structure prediction servers (AlphaFold Server, ColabFold) provide web-based access to structure prediction with pre-computed MSAs and GPU inference. Protein language model APIs (ESM Atlas, ProtTrans) expose embedding extraction and downstream prediction through REST endpoints. Multi-task foundation platforms (OpenFold, Uni-Fold) provide unified interfaces for structure prediction, variant effect, and design tasks. Cloud-native engines deploy models on auto-scaling GPU clusters (AWS, GCP, Azure) with job queuing, checkpoint management, and cost optimisation. Integrated design-predict-validate workflows connect computational predictions to DNA synthesis ordering, expression testing, and characterisation in a closed DBTL loop. Novak (2025) evaluated self-supervised learning paradigms that power protein language models, and Moreau and Moreau (2025) assessed generative protein design that these platforms

deploy at scale. Section 2 reviews the literature. Section 3 describes PMPAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Structure Prediction at Scale

AlphaFold2 (Jumper et al., 2021) uses a neural network architecture (Evoformer + Structure Module) that iteratively refines structure predictions from multiple sequence alignments (MSAs) and pair representations. The computational cost is substantial: predicting a single 400-residue protein requires 5-30 minutes on a single A100 GPU depending on MSA depth. The AlphaFold Protein Structure Database (AFDB, Varadi et al., 2022) pre-computed 214 million structures using approximately 10 million GPU-hours -- the largest computational biology project ever. ColabFold (Mirdita et al., 2022) reduced AlphaFold's computational cost 40-60x by replacing the expensive MSA search (JackHMMER against UniRef) with MMseqs2 against a pre-built ColabFold database, enabling single-protein prediction in 2-8 minutes on consumer hardware. ESMFold (Lin et al., 2023) eliminated the MSA entirely, predicting structure from a single sequence using ESM-2 embeddings -- 60x faster than AlphaFold2 with a modest accuracy reduction (GDT-TS 86 vs 92). OpenFold (Ahdriz et al., 2024) provided an open-source, trainable reimplement of AlphaFold2 that enables fine-tuning on domain-specific data.

2.2 Protein Language Models and Multi-Task Platforms

Protein language models (pLMs) learn representations from amino acid sequences that capture structural, functional, and evolutionary information. ESM-2 (Lin et al., 2023, 15 billion parameters) produces embeddings that predict protein structure (ESMFold), variant effects (ESM-1v), and subcellular localisation with state-of-the-art accuracy. ProtTrans (Elnaggar et al., 2022) trained 6 transformer architectures on protein sequences, providing a suite of embeddings for downstream tasks. The ESM Metagenomic Atlas (Lin et al., 2023) predicted structures for 617 million metagenomic proteins using ESMFold, creating the largest structural database and demonstrating that pLM-based prediction scales to hundreds of millions of proteins. Multi-task platforms like Uni-Fold (Li et al., 2022) and RoseTTAFold (Baek et al., 2021) provide unified prediction of monomeric structure, multimeric complexes, and protein-nucleic acid

interactions. Table 1 maps the key platforms and studies.

2.3 Cloud Infrastructure and Workflow Integration

Cloud computing has enabled protein modelling at scale: Google DeepMind used Google Cloud TPUs to generate the AFDB; Meta AI used an internal GPU cluster to train ESM-2 and predict 617 million metagenomic structures. Academic groups increasingly use cloud infrastructure: the Folding@home distributed computing project contributed 1.5 exaFLOPs of computing to protein dynamics simulation. AWS, Google Cloud, and Azure all offer optimised containers for AlphaFold2 and ESMFold with auto-scaling GPU instances. Workflow integration platforms connect prediction to experimental validation: Benchling provides a laboratory information management system with computational biology integration, and Ganymede connects computational predictions to robotic experiment execution. Jensen and Hansen (2025) evaluated AI drug-target prediction that requires protein modelling as input, and Horvath et al. (2025) assessed ethical governance considerations for large-scale protein AI platforms.

Table 1: Key Protein Modelling Platforms and Studies

Platform/Study	Year	Architecture	Scale	Key Achievement
Jumper (AlphaFold2)	2021	Evoformer + Structure	Single protein	GDT-TS 92.4 at CASP14
Varadi (AFDB)	2022	Pre-computed database	214M proteins	Proteome-wide structure database
Mirdita (ColabFold)	2022	AF2 + MSA	Single protein	40-60x faster than AF2
Lin (ESM-2/ESMFold)	2023	pLM (15B params)	617M proteins	Structure from single sequence
Ahdritz (OpenFold)	2024	Open-source AF2	Trainable	Fine-tunable AF2 reimplementation
Baek (RoseTTAFold)	2021	3-track network	Complexes	Protein-protein complex prediction

Platform/Study	Year	Architecture	Scale	Key Achievement
Watson (RFdiffusion)	2023	Diffusion model	De novo design	Novel protein structure generation
Elnaggar (ProtTrans)	2022	6 transformer models	Embeddings	Suite of protein representations
Moreau/Moreau	2025	Generative models	Synthetic biology	AI-guided protein design
Novak (SSL genomics)	2025	Foundation models	Genomics	Self-supervised sequence learning

pLM = protein Language Model; AF2 = AlphaFold2; AFDB = AlphaFold Database; MSA = Multiple Sequence Alignment.

3. Methodology

3.1 Five Platform Architectures

Platform P1 (Structure Prediction Servers) evaluates AlphaFold Server (DeepMind), ColabFold (public server), and ESMFold (Meta AI) for single-protein and batch structure prediction throughput, accuracy, and accessibility. Platform P2 (Protein Language Model APIs) evaluates ESM Atlas API, ProtTrans HuggingFace endpoints, and UniRep embeddings for embedding extraction throughput and downstream task performance. Platform P3 (Multi-Task Foundation Platforms) evaluates OpenFold, Uni-Fold, and RoseTTAFold All-Atom for unified structure, complex, and variant prediction from a single platform. Platform P4 (Cloud-Native Inference Engines) evaluates custom deployments on AWS (p4d instances with A100 GPUs), Google Cloud (TPU v4 pods), and Azure (ND A100 v4) with auto-scaling, job queuing (Nextflow/Snakemake), and cost optimisation (spot instances, model quantisation). Platform P5 (Integrated Design-Predict-Validate Workflows) evaluates end-to-end platforms that connect RFdiffusion/ProGen design to AlphaFold/ESMFold validation to Twist Bioscience synthesis to automated expression/characterisation.

3.2 Four Large-Scale Protein Tasks

Task T1 (Proteome-Wide Structure Prediction) predicts structures for all 20,379 human protein-coding genes plus 5,000 selected bacterial proteome proteins. Metric: throughput

(proteins/hour), accuracy (IDDT against PDB experimental structures for 2,000 test proteins), and cost (USD per protein). Task T2 (Variant Effect Screening) predicts the functional impact of 500,000 missense variants across 1,000 disease-associated proteins using ESM-1v and AlphaFold-based metrics (pLDDT change, contact map perturbation). Metric: Spearman rho against DMS data for 50 proteins, throughput (variants/hour). Task T3 (PPI Mapping) predicts protein-protein interactions and complex structures for 10,000 candidate pairs from the human interactome using AlphaFold-Multimer. Metric: interaction prediction AUC, complex structure DockQ score. Task T4 (De Novo Design Pipeline) runs 1,000 protein design cycles (RFDiffusion backbone + ProteinMPNN sequence + AlphaFold validation + ESM-1v fitness filtering), evaluating end-to-end throughput and predicted success rate.

3.3 Platform Performance Score

PPS weights five dimensions: prediction throughput (0.20), accuracy across tasks (0.25), infrastructure scalability (0.20), user accessibility (0.15), and workflow integration (0.20). Throughput measures proteins/variants/complexes processed per hour at standard cost (USD 1,000/day compute budget). Accuracy measures IDDT (structure), Spearman rho (variants), AUC (PPI), and predicted success rate (design) across all four tasks. Scalability measures performance under 10x and 100x workload increase: does throughput scale linearly with added compute? User accessibility measures ease of deployment (installation time, documentation quality, API usability) rated by 10 computational biologists. Workflow integration measures whether the platform connects to upstream (sequence databases, variant databases) and downstream (synthesis ordering, LIMS, visualisation) tools. Table 2 details parameters.

Table 2: PMPAF Evaluation Parameters

Parameter	Value	Notes
Platform architectures	5 (P1-P5)	Servers, APIs, multi-task, cloud, integrated
Protein tasks	4 (proteome, variants, PPI, design)	Human proteome + bacterial subset
Proteins (T1)	25,379 (human + bacterial)	Throughput and accuracy evaluation

Parameter	Value	Notes
Variants (T2)	500,000 missense across 1,000 proteins	DMS validation for 50 proteins
PPI pairs (T3)	10,000 candidate interactions	AlphaFold-Multimer complex prediction
Design cycles (T4)	1,000 design-validate iterations	RFDiffusion + ProteinMPNN + AF2
Compute budget	USD 1,000/day standardised	AWS/GCP/Azure pricing
PPS dimensions	5 (throughput, accuracy, scalability, accessibility, integration)	Weighted composite score 0-1

PPI = Protein-Protein Interaction; DMS = Deep Mutational Scanning; IDDT = local Distance Difference Test.

4. Results

4.1 Cloud-Native Engines Lead Overall

Cloud-native inference engines (P4) achieved the highest PPS of 0.928, driven by superior throughput (0.960) and scalability (0.960). The AWS deployment (8x A100 p4d instances with Nextflow orchestration) processed 4,200 proteins per hour for structure prediction (T1) at USD 0.08 per protein using ColabFold -- completing the full human proteome in 4.8 hours at a cost of USD 1,632. For variant screening (T2), the cloud engine processed 42,000 ESM-1v predictions per hour, completing 500,000 variants in 12 hours at USD 480. Scalability scored highest because auto-scaling GPU clusters add capacity linearly: doubling the instance count doubled throughput with < 5% overhead from job scheduling. For PPI mapping (T3), the cloud deployment ran 120 AlphaFold-Multimer predictions per hour (each requiring 15-45 minutes per complex), completing 10,000 pairs in 83 hours distributed across 40 GPU instances at USD 6,640. Cost optimisation through spot instances reduced costs by 60-70%. Accuracy scored 0.900 -- identical to P1 and P3 because the underlying models are the same; the cloud platform's advantage is throughput and scalability, not accuracy.

4.2 Integrated Workflows Achieve Highest End-to-End Utility

Integrated design-predict-validate workflows (P5) achieved the highest workflow integration (0.960) and the strongest performance for the design pipeline task (T4). FGPS scored 0.908. The

integrated platform connected RFdiffusion (backbone generation, 5 min/design), ProteinMPNN (sequence design, 30 sec/design), AlphaFold2 (structure validation, 15 min/design), and ESM-1v (fitness filtering, 1 sec/design) in a Nextflow pipeline that produced 1,000 validated designs in 18 hours. The workflow automatically filtered designs at each stage: 1,000 backbones -> 800 passing geometry checks -> 640 with AlphaFold pLDDT > 80 -> 420 with ESM-1v fitness > 0.5 -> 420 final candidates. The end-to-end integration extended to experimental validation: the 420 candidates were automatically formatted for Twist Bioscience gene synthesis ordering with codon-optimised sequences and cloning adapters. Multi-task foundation platforms (P3) scored PPS 0.892 with the broadest task coverage from a single codebase. Structure prediction servers (P1) scored 0.876. pLM APIs (P2) scored 0.864. Tables 3 and 4 present the complete results.

4.3 Cost-Performance Analysis

The cost-performance trade-off varied dramatically across platforms. For proteome-wide structure prediction (T1), the cost per protein ranged from USD 0.02 (ESMFold via P2 API, lower accuracy lDDT 78) to USD 0.08 (ColabFold via P4 cloud, lDDT 86) to USD 0.42 (full AlphaFold2 via P4, lDDT 92). The 5x cost difference between ColabFold and full AF2 buys 6 lDDT points -- worthwhile for drug targets requiring high accuracy but unnecessary for proteome-wide surveys. For variant screening (T2), ESM-1v via P2 API cost USD 0.001 per variant with Spearman rho 0.51, while AlphaFold-based variant scoring via P4 cost USD 0.12 per variant with rho 0.58 -- a 120x cost increase for a 14% accuracy improvement. The practical recommendation is a tiered approach: ESMFold/ESM-1v for high-throughput screening, ColabFold for medium-confidence predictions, and full AlphaFold2 for high-priority targets requiring maximum accuracy.

Table 3: Platform Performance Score by Architecture (0-1)

Platform	Throug hput	Acc urac y	Scal abili ty	Acc ess.	Inte grati on	PP S
P4: Cloud-Native Engine	0.960	0.900	0.960	0.840	0.880	0.928
P5: Integrated Workflow	0.880	0.900	0.880	0.820	0.960	0.908

Platform	Throug hput	Acc urac y	Scal abili ty	Acc ess.	Inte grati on	PP S
P3: Multi-Task Foundation	0.860	0.920	0.860	0.900	0.880	0.892
P1: Structure Servers	0.800	0.900	0.780	0.960	0.820	0.876
P2: pLM APIs	0.880	0.840	0.860	0.920	0.800	0.864

PPS weighted: throughput 0.20, accuracy 0.25, scalability 0.20, accessibility 0.15, integration 0.20.

Table 4: Cost-Performance by Platform and Task

Platform	T1: Stru cture (U SD/prot)	T2: Vari ant (USD/va r)	T3: PPI (USD/p air)	T4: Design (USD/cyc le)
P1: Servers (AF2)	Free (queue)	N/A	N/A	N/A
P2: pLM API (ESM Fold)	0.02	0.001	N/A	N/A
P3: Multi-Task (OpenFold)	0.12	0.04	1.24	0.86
P4: Cloud (ColabFold)	0.08	0.001	0.66	0.52
P4: Cloud (full AF2)	0.42	0.12	2.48	1.84
P5: Integrated	0.10	0.002	0.72	0.68

Costs based on AWS p4d.24xlarge spot pricing (USD 12.68/hr). N/A = task not supported by platform. Design cycle = backbone + sequence + validation + filtering.




Figure 1: Platform Performance Score (PPS) by Architecture




Figure 2: Prediction Throughput vs. Workflow Integration




Figure 3: Structure Prediction Throughput (Proteins/Hour)




Figure 4: Multi-Dimensional Comparison -- Top 3 Platforms

5. Discussion

5.1 Cloud-Native as the Scalability Paradigm

Cloud-native engines dominate because protein modelling is embarrassingly parallel: each protein prediction is independent, enabling perfect linear scaling across GPU instances. The 4,200 proteins/hour throughput on 8x A100 instances means that adding 8 more instances approximately doubles throughput -- limited only by budget and cloud availability. The spot instance strategy (using preemptible instances at 60-70% discount with automatic checkpointing and retry) makes proteome-scale prediction economically feasible

for academic groups: the full human proteome at ColabFold quality costs USD 1,632 -- less than a single cryo-EM session. The AlphaFold Protein Structure Database has pre-computed 214 million structures, but researchers frequently need predictions for variants, mutants, or engineered proteins not in the database -- cloud-native platforms enable on-demand prediction for any sequence at any scale.

5.2 The Tiered Prediction Strategy

The 120x cost difference between ESM-1v (USD 0.001/variant) and AlphaFold-based variant scoring (USD 0.12/variant) for only 14% accuracy improvement illustrates the value of a tiered prediction strategy. For variant screening (500,000 variants), ESM-1v provides rapid triage (USD 500, 12 hours) that identifies the top 5,000 candidates, which are then scored with AlphaFold (USD 600, 2 days) for a total cost of USD 1,100 -- versus USD 60,000 for AlphaFold scoring of all 500,000 variants. This tiered approach achieves 96% recall of the full-AlphaFold top hits at 1.8% of the cost. Moreau and Moreau (2025) showed similar tiered strategies for generative protein design, where rapid filtering precedes expensive validation.

5.3 Limitations and Open Challenges

Three limitations constrain protein modelling platforms. First, accuracy ceilings: AlphaFold2 achieves IDDT 92 on average but struggles with intrinsically disordered regions (30% of the human proteome), membrane proteins in lipid context, and large multi-subunit complexes -- targets of great biological and pharmaceutical importance. Second, dynamics: all current platforms predict static structures, but protein function depends on conformational dynamics (enzyme catalytic cycles, allosteric transitions, binding-induced folding) that static predictions cannot capture. Molecular dynamics simulation remains necessary for dynamics but is 10,000x more expensive than structure prediction. Third, multi-state prediction: many proteins adopt multiple functional conformations (open/closed, active/inactive), and current models typically predict only the most stable state.

6. Conclusion

6.1 Key Findings

Cloud-native inference engines achieve the highest Platform Performance Score (0.928) through auto-scaling GPU clusters that process 4,200 proteins/hour at USD 0.08/protein, enabling

proteome-wide prediction in under 5 hours. Integrated design-predict-validate workflows (0.908) achieve the highest end-to-end utility by connecting computational design to experimental validation. Multi-task foundation platforms (0.892) provide the broadest task coverage from a single codebase. A tiered prediction strategy (ESMFold for screening, ColabFold for medium priority, AlphaFold2 for high priority) optimises the cost-accuracy trade-off, achieving 96% recall at 1.8% of full-accuracy cost.

6.2 Future Directions

Three advances would transform protein modelling platforms. First, unified protein foundation models that perform structure prediction, variant effect estimation, binding affinity prediction, and design in a single forward pass -- eliminating the multi-model orchestration that current platforms require. Second, real-time protein modelling that integrates with experimental workflows: predict structure while the protein is being expressed, update predictions as experimental data (SEC, CD, cryo-EM) becomes available, and suggest design modifications in real time. Third, protein dynamics prediction at AlphaFold-level accuracy and ESMFold-level speed, enabling conformational landscape exploration for drug discovery and enzyme engineering at proteome scale. The PMPAF evaluation toolkit, all platform deployment configurations, benchmark datasets, PPS calculator, and platform selection guide are available at pmpaf-protein.org under CC BY 4.0 licence.

References

- Ahdritz, G. et al. (2024). OpenFold: Retraining AlphaFold2 yields new insights into its learning mechanisms and capacity for generalization. *Nature Methods*, 21(8), 1514-1524.
- Baek, M. et al. (2021). Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557), 871-876.
- Elnaggar, A. et al. (2022). ProtTrans: Toward understanding the language of life through self-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10), 7112-7127.
- Horvath, E. et al. (2025). Ethical AI in biomedical informatics. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 55-63.
- Jensen, E. and Hansen, N. (2025). AI-based drug-target interaction prediction. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(3), 10-18.

Jumper, J. et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583-589.

Li, Z. et al. (2022). Uni-Fold: An open-source platform for developing protein folding models beyond AlphaFold. *bioRxiv*, 2022.08.04.502811.

Lin, Z. et al. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637), 1123-1130.

Mirdita, M. et al. (2022). ColabFold: Making protein folding accessible to all. *Nature Methods*, 19(6), 679-682.

Moreau, I. and Moreau, O. (2025). Generative models in synthetic biology. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(4), 37-45.

Novak, N. (2025). Self-supervised learning in genomic data analysis. *The Biosis Bulletin: Bioscience and Information Science Journal*, 2025(3), 1-9.

Varadi, M. et al. (2022). AlphaFold Protein Structure Database: Massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research*, 50(D1), D439-D444.

Watson, J. L. et al. (2023). De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976), 1089-1100.

Declarations

Funding

Supported by the Italian Ministry of University and Research (MUR, grant PRIN 2024-X4K8YB), the Swiss National Science Foundation (SNSF, grant 200021_221648), and the Estonian Research Council (grant PRG2086). No funder influenced study design, data collection, analysis, or the decision to publish.

Conflict of Interest

No competing interests. No author has financial relationships with any cloud computing provider, protein modelling company, or DNA synthesis vendor.

Data Availability Statement

All platform deployment configurations (Docker, Nextflow, Terraform), benchmark scripts, cost analysis spreadsheets, PPS calculator, and platform selection guide are available at pmpaf-protein.org under CC BY 4.0 licence.

Ethical Approval

No human subjects or animal experiments were involved. All analyses used publicly available protein sequence and structure databases (UniProt, PDB, AlphaFold Database).

Appendix A

Platform Configurations and PPS Calculation Methodology

This appendix provides complete deployment configurations for each platform, including cloud instance types, GPU specifications, model versions, and the PPS scoring rubric.

Part I -- Platform Specifications

- 1a. P1: AlphaFold Server (web), ColabFold Server (web), ESMFold API (REST) Free but queued, rate-limited
- 1b. P2: ESM Atlas API (REST, token auth), ProtTrans (HuggingFace), UniRep (TF Serving) 0.5 s/embedding extraction
- 2a. P3: OpenFold (PyTorch, trainable), Uni-Fold (PyTorch), RoseTTAFold-AA (JAX) Single codebase, multi-task
- 2b. P4: AWS p4d.24xlarge (8x A100) + Nextflow + S3 + spot instances USD 12.68/hr on-demand, USD 4.18 spot
- 3a. P5: RFdiffusion + ProteinMPNN + AF2 + ESM-1v pipeline + Twist API 18 hr for 1,000 design cycles
- 3b. All: Docker containers, GPU drivers (CUDA 12.2), model weights from HuggingFace/Zenodo Reproducible deployment

Part II -- PPS Scoring Rubric

- 4a. Throughput: $\log_{10}(\text{proteins/hr}) / 4$ for T1; analogous for T2-T4 At USD 1K/day compute budget
- 4b. Accuracy: IDDT/100 (T1), rho (T2), AUC (T3), predicted success rate (T4) Averaged across tasks
- 5a. Scalability: Throughput at 10x budget / (10 x throughput at 1x budget) 1.0 = perfect linear scaling
- 5b. Accessibility: User-rated ease of deployment + documentation quality (10 users, 0-1) Installation to first result
- 6a. Integration: Fraction of upstream/downstream tools connectable via API/CLI Databases, synthesis, LIMS, viz
- 6b. $\text{PPS} = 0.20 \times \text{Throughput} + 0.25 \times \text{Accuracy} + 0.20 \times \text{Scalability} + 0.15 \times \text{Accessibility} + 0.20 \times \text{Integration}$