ARENA

Adaptive Reseeded Elimination for Novel binder Architectures

De Novo RBX1 Binder Design via Multi-Epitope Protein Hunter Tournaments

Tanay Lohia, Sourab Kulkarni • Mandrake Bio, Bangalore, India
tanay@mandrake.bio • sourab.kulkarni@mandrakebio.com

🎯 Target challenge: RBX1's N-terminus (res 1–39) is an intrinsically disordered region, while the C-terminal RING-H2 is stabilized by three Zn^{2+} ions through 12 load-bearing residues (Cys42/45/53/56/68/75/83/94, His77/80/82, Asp97). The IDR makes AI confidence metrics (ipTM, pLDDT) unreliable as sole quality indicators — we supplement with biophysical property assessment and structural biology reasoning (zinc-aware contact exclusion, polar-only hotspot selection).

① Protein Hunter: Strengths & the Seed Problem

Protein Hunter (Nazari et al., 2025) is a compute-light iterative co-design pipeline coupling **Boltz-2** (structure prediction) and **SolubleMPNN** (sequence design). Each cycle folds a binder in complex, scores it, and re-optimizes the sequence on the predicted geometry. Over 5 cycles, sequence and structure co-evolve toward higher predicted affinity — using only forward passes, no backpropagation through the folding network.

The limitation: Protein Hunter is strongly path-dependent. The initial random seed determines the structural basin Boltz-2 predicts, constraining all subsequent redesigns. Seeds producing low ipTM at cycle 1 rarely recover. The pipeline is highly efficient *given a good seed*, but seed quality is unpredictable a priori.

② ARENA: Adaptive Reseeded Elimination at Low Cost

We treat seed sensitivity as a **multi-armed bandit**: run many seeds cheaply, observe early signal, concentrate compute on winners. We call this **ARENA** — progressive elimination applied *during* generation, not after:

R1: BROAD SWEEP	▸	R2: TRAJECTORY CUT	▸	R3: FULL REFINEMENT
100 seeds → 40		**40 → 20**		**20 → elite**
1 cycle • Keep top 40%		2 cycles • Improving + high ipTM		2 cycles • 5-cycle designs

220	100	2.3×	880
cycle-equiv / combo	starting seeds / combo	diversity vs full traj.	total designs (4 combos)

We use **SolubleMPNN** (not LigandMPNN) — consistent with the EGFR competition's winning formula. LigandMPNN produced charged-helix artifacts in our preliminary experiments.

③ Multi-Epitope Hedging with Zinc-Aware Contact Selection

We hedge against the #1 binder design failure mode — *choosing the wrong surface* — by running **four mechanistically distinct epitope strategies**, each anchored on mutagenesis-validated Trp87 and Arg91:

ID	Strategy	Key Rationale
C1	GLMN-competitive	Directly occludes GLMN-validated E2-recruiting surface (PDB: 4F52)
C2	CRL state-selective	Targets Cullin-facing surface accessible only in assembled CRL complex
C3	Extended E2 face	Broader footprint extending into free accessible groove (res 60–69)
C4	Full E2 sweep	Most permissive hotspot set — Boltz-2 discovers diverse geometries across entire E2 face

Zinc-aware contact selection: All 12 Zn-coordinating residues excluded from every contact list. All selected contacts (8–10 per combo) are strictly polar/charged — steering Boltz-2 through directional H-bonds while hydrophobic anchors (Trp87, Pro95, Leu96) engage naturally via packing.

④ Design Constraints

No cysteines in any binder: (i) Free Cys thiol competes for RBX1's structural Zn^{2+}, risking target destabilization. (ii) E. coli cytoplasm is reducing — unpaired Cys causes disulfide-mediated aggregation on purification. (iii) Field standard (BoltzGen excludes Cys by default). **Wide size range 70–150 AA** for both compact mini-binders and extended-interface designs. **SolubleMPNN** for all sequence design. **Deliberately relaxed in-silico thresholds:** Soft targets of ipTM ≥ 0.85, pLDDT ≥ 0.80, iPLDDT ≥ 0.75 guided initial ranking, but were not applied as hard cutoffs. RBX1's IDR (res 1–39) participates in binder interactions but is inherently disordered - this artificially suppresses confidence metrics even for structurally sound interfaces. Several top candidates by Rosetta energetics and manual inspection score below these soft targets (ipTM as low as 0.69), validating our decision to rely on multi-oracle consensus rather than single-metric filtering. Ranked by harmonic mean of ipTM and pLDDT. One best candidate per seed for diversity.

⑤ What Makes This Novel

🎯 ARENA Tournament	🧬 Multi-Epitope Hedging	⚡ Compute Democratization
Selection pressure during generation via successive halving. 100 seeds per combo explore 2.3× more structural basins than equivalent full-trajectory compute.	4 distinct epitope strategies with zinc-aware polar contact selection. Hedges against #1 failure mode: wrong surface.	Full campaign: 4 epitopes × 100 seeds, 70–150 AA range, 880 designs on a single A100. No HPC required.

⑥ In-Silico Results

Metric	Range	Best	Notes
ipTM	0.691 – 0.932	MBL_020 (C1)	26 submitted (from 880)
pLDDT	0.722 – 0.874	—	Size: 80–131 AA
iPLDDT	0.683 – 0.910	MBL_013 (C2)	4 epitope strategies

⑥ Ranking & Selection (880 → top 50 → 26 submitted)

Funnel: 880 designs generated across 4 ARENA tournaments (100 seeds × ~5 cycles × 4 epitope combos, with progressive elimination). The **top 50** were selected by harmonic mean ranking for manual evaluation, then quality-filtered and curated to **26 final submissions**.

Step 1 — Harmonic mean ranking (880 → 50): From the 880 designs, candidates were ranked by HMean(ipTM, pLDDT) = 2·ipTM·pLDDT / (ipTM + pLDDT). The harmonic mean penalizes imbalance — a design with ipTM=0.95, pLDDT=0.50 scores 0.648 (not 0.725), forcing both metrics high simultaneously. Top 50 selected, ranked within each cohort then interleaved in tier order (C1 > C4 > C2 > C3) so all 4 epitopes are represented before any cohort repeats.

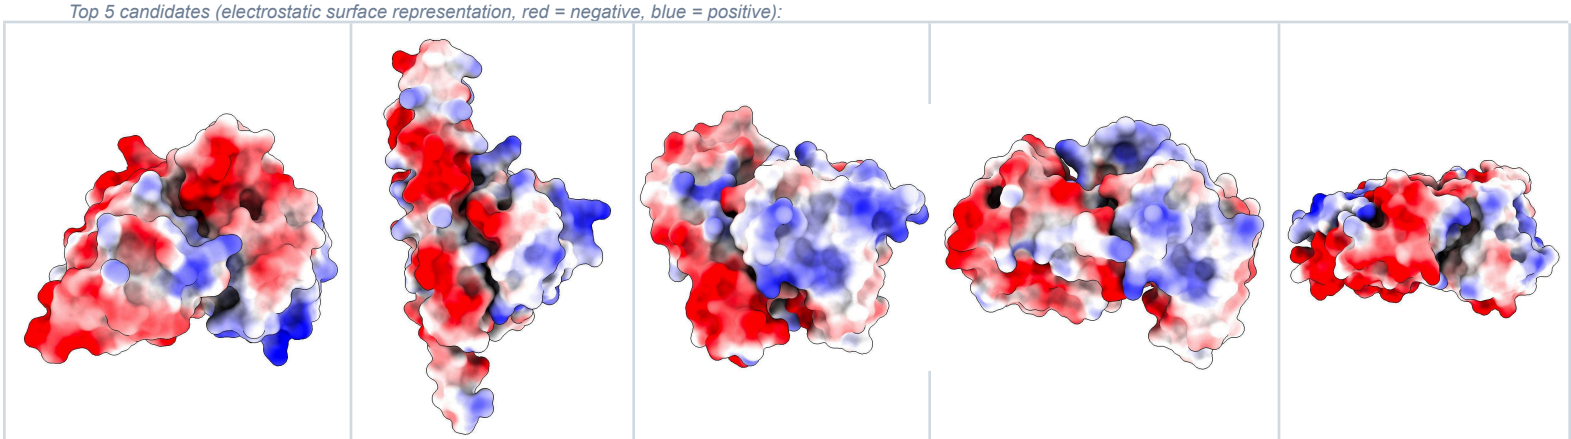
Step 2 — Quality filters + manual curation (50 → 26): Hard cutoffs: (i) ipSAE (interface PAE) < 0.20 → rejected (low contact-level confidence). (ii) Rosetta dG_separated > 100 kcal/mol → rejected (non-physical = misfolded interface). Remaining candidates manually inspected: Boltz-2 complex visualization, Rosetta dG_separated, H-bond count, buried surface area, shape complementarity (SC > 0.60), and E2-face coverage.

Step 3 — Priority ranking (top 13 slots): The top 13 positions (MBL_001–013) were manually assigned based on qualitative structural review. MBL_004 (ipSAE=0.186) force-included for strong dG=−21.4, 50% E2 coverage, SC=0.581 — a functionally relevant CRL-state-selective interface. ***This is because the past Adaptyv competitions have shown low correlation between in-silico metrics and actual experimental results***.

Code	Strategy	#	Priority
C1	GLMN-competitive	10	1st
C4	Full E2 sweep	4	2nd
C2	CRL state-selective	5	3rd
C3	Extended E2 face	7	4th

⑦ Representative Structures (Top Candidates Across C1–C4)

Top 5 candidates (electrostatic surface representation, red = negative, blue = positive):



MBL_001	MBL_002	MBL_003	MBL_004	MBL_005

References

[1] Nazari et al. (2025). Protein Hunter. github.com/adaptyvbio/protein_hunter
[2] Pacesa et al. (2024). BindCraft. bioRxiv 2024.09.30.615802. [3] Goverde et al. (2024). SolubleMPNN. Nature 631, 449–458.
[4] Dauparas et al. (2022). ProteinMPNN. Science 378, 49–56. [5] PDB: 4F52 (GLMN–RBX1). [6] PDB: 2LGV (apo RBX1).