

Absolute Affinity — Validation Summary

Evidence disclosure of LatticeZero's production affinity engine across held-out and independent panels.

Version 1.0

June 2026

Executive summary

LatticeZero's production affinity engine predicts absolute binding free energy (ΔG) for protein–ligand complexes. This summary reports validation evidence across (1) a held-out blind panel evaluated by a product **candidate** scorer and (2) an **independent, outside-the-loop** challenge panel of 20 public complexes scored by the **production full-feature scoring path** with no model adjustment. Across the independent panel the engine **ranks** binding affinities well (Spearman $\rho \approx 0.8$) with a mean absolute error around 2 kcal/mol. Results are reported in full, including a single documented chemistry limitation.

Panels (public level)

Held-out blind panel — used to evaluate a product-candidate scorer; the candidate is not the production default.

PDBbind-scale scoring panel — standard public-database-scale scoring-power evaluation for context on the production path.

Independent 20-complex challenge — 20 public complexes with measured affinities, sourced so they were not part of any prior LatticeZero validation set; a held-out public-database subset (15) and a fully-independent external subset (5). Structures used quality-controlled preparation; scoring used the production full-feature path.

Results — product candidate [Candidate]

Held-out blind panel. Reported as evidence only — not the production default, not promoted.

Metric	Value
Rank agreement (Spearman ρ)	0.9004
Mean absolute error	0.85 kcal/mol
Cases off by ≥ 2 kcal/mol	2

Results — independent 20-complex challenge [Production path]

Production full-feature path, no model adjustment of any kind.

subset	n	Spearman ρ	MAE	RMSE
Combined	20	0.77	2.11	2.86
Combined, excluding one documented outlier	19	0.80	1.74	2.05
Held-out public-database subset	15	0.75–0.79	1.6–2.1	—
Fully-independent external subset	5	0.90	2.11	2.35

CAVEATS

One carbohydrate-rich case exposed a known chemistry limitation; it is reported **both included and excluded**. The independent panel is a **ranking / sanity result**, not a universal calibrated- ΔG guarantee (mean absolute error out-of-panel around 2 kcal/mol). The product-candidate result is evidence only — not the shipped default.

Status

Production: the production full-feature scoring path is live and reproducible in-browser. **Candidate:** the product-candidate scorer is evaluated but not promoted and not the default.

This document discloses **evidence**, not methodology. It contains no proprietary methods, model internals, feature definitions, formulas, or architecture. © 2026 LatticeZero.