

# PRISM4D: Firewalled Blind Validation of Spike-Event-Driven Nonequilibrium Molecular Sensing for Cryptic Binding-Site Discovery

Ididia J. Serfaty

Delfictus I/O LLC, Los Angeles, California, USA

is@delfictus.com

– May 16, 2026 –

## Abstract

Cryptic and allosteric binding sites often arise from transient apo-state conformations that are absent or occluded in static structures. We introduce PRISM-4D, a physics-driven platform that represents ligand accessibility as recurrent nonequilibrium event manifolds rather than as static surface geometry. The engine combines AMBER ff14SB molecular dynamics, a five-phase cryo-thermal hysteresis protocol, wavelength-specific in silico UV pump-probe perturbation, UV-activated benzene cosolvent probes, three-channel leaky integrate-and-fire (LIF) spike detection, electrostatic flux sensing, and Spike-Native Density Clustering (SNDC), a four-stage GPU pipeline comprising RT-DBSCAN over OptiX bounding-volume hierarchies, watershed segmentation, Eikonal distance-field propagation, and intensity-squared centroid tracking.

The validation design separates frozen apo prediction from all post-freeze holo-coordinate scoring. Across the 9-target primary blind panel, PRISM4D achieved Strict SR@1 = 0.483 versus fpocket = 0.224 and P2Rank = 0.382; Strict SR@3 values were 0.483, 0.335, and 0.382; and Strict SR@5 values were 0.483, 0.465, and 0.382, respectively. Leave-one-reference-out (LORO) analysis retained strong flagship recovery for the blind allosteric discoveries HRAS\_Q61H and CDK2.allosteric while preserving limitation, boundary, and calibration cases. PRISM-Therm/CCNS descriptors annotate each frozen manifold with  $\tau$ , hysteresis, TIDE coupling, effective  $\Delta G$ , spike density, source diversity, and related thermodynamic/accessibility metadata. Strict null controls used 10,000 randomization iterations and publication-safe interpretation rules that explicitly exclude degenerate rows from primary evidence. As a mechanistic post-freeze audit, vectorial pharmacophore feature recovery (PFR) recovered 26.9% of holo pharmacophore features from frozen phase-ordered apo manifolds versus 2.2% under a temporal-scramble null (12.2-fold enrichment; empirical  $p \leq 0.001$ ), using a 3.5 Å spatial and 30° angular criterion.

Discovery case studies show top-ranked detection of the experimentally validated SIRP $\alpha$  WYF cryptic pocket from the closed apo structure and a highly potential novel  $\beta$ -catenin Site 6 prediction with a CYS240/CYS278 dual-cysteine pharmacophore. The complete platform runs on a single NVIDIA RTX 5080-class consumer GPU and is accompanied by a firewalled reproducibility package with freeze records, SHA-256 manifests, baseline outputs, strict null controls, causal-driver artifacts, Therm/CCNS descriptors, med-chem anchor validation, structural renders, and PFR artifacts. The study is computational platform validation; it does not establish binding affinity, wet-lab activity, or clinical relevance.

## 1 Introduction

The identification of druggable binding sites on protein surfaces is a foundational step in structure-based drug discovery. Orthosteric sites are often visible in static crystal structures, but many therapeutically important opportunities are cryptic: they become ligand-accessible only through side-chain rearrangement, local dewetting, loop motion, or broader conformational fluctuation. This is especially consequential for proteins historically labeled “undruggable,” where the relevant ligandable state may be rare, transient, and poorly represented by a single apo conformation.

Most computational binding-site workflows ask where a cavity or trained structural pattern is present. Geometric methods such as fpocket, LIGSITE, and SiteMap analyze the static molecular surface using Voronoi or grid-based cavity logic. Machine-learning methods such as P2Rank, DeepSite, and PocketMiner infer sites from learned local features or trajectory-derived patterns. These approaches are useful, but their signal is primarily geometric or statistical. They do not directly ask whether a region repeatedly participates in a nonequilibrium physical process that creates ligand accessibility.

PRISM-4D asks that different question. It treats a candidate binding site as a recurrent spatiotemporal event

manifold generated during apo-state perturbation and device-resident molecular sensing. A candidate is not only a centroid: it can contain phase-specific centers, spike intensity fields, event sources, driver residues, water-exclusion signatures, electrostatic flux, family recurrence, Therm/CCNS descriptors, and post-freeze medicinal-chemistry anchor annotations. This representation is intended for settings in which ligand accessibility is conditional rather than constitutively exposed.

The central risk in evaluating a new discovery platform is post hoc adaptation. This manuscript therefore unifies two previously separated drafts: the architecture manuscript, which documented the physics engine and SNDC pipeline, and the rigor manuscript, which documented the firewalled blind validation framework. The unified design preserves both: the deep methods are retained, and the validation claim is constrained to frozen ranks scored after holo-coordinate access. Post-freeze analyses, including Therm/CCNS, med-chem anchors, strict null controls, structural overlays, and Vectorial PFR, are mechanistic support layers. They do not modify frozen rank order, primary success definitions, target inclusion, or baseline comparisons.

## 2 PRISM4D Architecture & Physics

**Scientific Contribution and Claim Boundaries.** PRISM-4D is presented as a computational platform for event-driven cryptic and allosteric binding-site detection: AMBER ff14SB perturbation dynamics, CryoUV hysteresis, UV-active probe sensing, three-channel LIF spike detection, SNDC/RT-core clustering, Therm/CCNS descriptors, and firewalled blind validation. The primary quantitative claims are limited to frozen apo/pre-ligand ranks scored after holo access under the stated SR@K rules. Post-freeze nulls, LORO, med-chem anchors, structural renders, causal-driver annotations, and Vectorial PFR are mechanistic support layers; they do not establish binding affinity, biochemical activity, clinical relevance, or universal superiority over all pocket-detection methods.

### 2.1 System Architecture

PRISM-4D integrates four computational stages into a GPU-accelerated pipeline implemented in Rust with CUDA kernels for performance-critical operations (Figure 1). The binary is designed for local execution and requires only the NVIDIA CUDA runtime after installation. Candidate sites are generated from apo or pre-ligand structures without holo ligand coordinates.

The architecture includes an NHS engine for molecular dynamics and perturbation, a three-channel spike detector, SNDC event clustering, and multi-parametric site

characterization. A site is represented by event source distributions, voxelized spike intensity, local water-exclusion behavior, thermodynamic descriptors, driver residues, pocket geometry, and optional covalent-reactive residues.

### 2.2 Force Field and Integration

Molecular dynamics simulations use the AMBER ff14SB protein force field with the following protocol:

- Integration: velocity Verlet with Langevin thermostat,  $\gamma = 10.0 \text{ ps}^{-1}$  at 300 K, with temperature-dependent scaling  $\gamma(T) = \gamma_0 \sqrt{T_{\text{ref}}/T}$ .
- Timestep: 2 fs for the architecture benchmark. The firewalled validation used the frozen CryoUV fast\_35k protocol with target-specific integration settings recorded in the freeze artifacts.
- Nonbonded cutoff: 10 Å with a 12 Å neighbor-list buffer.
- 1-4 scaling factors: Lennard-Jones 0.5 and electrostatic 5/6, matching standard AMBER ff14SB conventions.
- Bond constraints: SHAKE on X-H bonds.
- Electrostatics: Coulomb interactions with a distance-dependent Warshel dielectric,  $\epsilon_r(r) = \max(4r, 4)$ .
- Solvent treatment: implicit hydrophobic-exclusion water inference on a 0.5 Å grid.
- PRNG: CUDA cuRAND XORWOW states, seeded independently per stream.


Each firewalled validation target was run as a frozen apo/pre-ligand prediction. Ligand, holo, and contact-shell information entered only after the freeze and only for scoring or interpretation.

### 2.3 Cryo-Thermal Hysteresis Protocol

The architecture benchmark executes a five-phase thermal hysteresis cycle (Figure 2):

1. **Cold hold:** 14,000 steps at 50 K, establishing a thermally quiet baseline for UV spike detection.
2. **Ramp up:** 6,000 steps from 50 K to 300 K, driving forward conformational transitions by linear temperature interpolation.
3. **Warm hold:** 15,000 steps at 300 K for physiological-like sampling.
4. **Ramp down:** 6,000 steps from 300 K to 50 K, measuring hysteresis asymmetry between heating and cooling spike distributions.
5. **Cold return:** 14,000 steps at 50 K, verifying final frozen-state behavior.

The total is 55,000 steps per stream (110 ps at 2 fs). UV bursts of 42 kcal/mol fire every 250 steps for 50 steps, cycling through 280, 274, 258, and 211 nm with 300-step dwell per wavelength. The nonequilibrium schedule accelerates rare gating and dewetting transitions through a physically interpretable 250 K temperature range. The firewalled validation package additionally records a CryoUV fast\_35k variant (14,000-step cold hold, 6,000-step



**Figure 1: PRISM-4D computational pipeline.** Stage 1: AMBER ff14SB molecular dynamics with cryo-thermal perturbation and UV excitation. Stage 2: three-channel LIF neuromorphic spike detection (UV/LIF/EFP). Stage 3: SNDC clustering (RT-DBSCAN, watershed, Eikonal BFS, peak tracking). Stage 4: site scoring, classification, and covalent-residue annotation.

ramp up, 15,000-step warm hold, 6,000-step ramp down, 4,000-step cold return) used for the blind validation artifacts.

## 2.4 In Silico UV Pump-Probe Spectroscopy

PRISM-4D introduces a computational analogue of UV pump-probe spectroscopy. UV pulses deposit energy into aromatic residues according to wavelength-specific absorption parameters. The deposited energy per chromophore is

$$E_{\text{dep}} = E_{\text{photon}} \times \sigma(\lambda) \times F \times \eta, \quad (1)$$

where  $E_{\text{photon}}$  is photon energy,  $F$  is UV fluence (0.024 photons/ $\text{\AA}^2$ ), and  $\eta$  is the fraction of absorbed energy converted to heat. The absorption cross-section is converted from molar extinction coefficient  $\varepsilon$  by

$$\sigma(\lambda) = \varepsilon(\lambda) \times 3.823 \times 10^{-5} \text{ \AA}^2. \quad (2)$$

Deposited energy is converted into kinetic perturbation through the AMBER velocity-assignment protocol. The key signal-design premise is that water is effectively transparent at aromatic amino-acid UV wavelengths, creating a signal-on-silent-background perturbation.

**Table 1:** UV chromophore photophysical parameters used by the perturbation model.

Chromophore	$\lambda_{\text{max}}$ (nm)	$\varepsilon$ ( $\text{M}^{-1}\text{cm}^{-1}$ )	$\eta$
Trp	280	5,600	0.83
Tyr	274	1,490	0.86
Phe	258	197	0.94
S-S	211	300	0.95
Benzene (virtual)	254	204	0.71

## 2.5 Three-Channel LIF Spike Detection

Each voxel is monitored by three independent spike channels with temporal isolation. This design prevents UV

perturbation, dewetting, and electrostatic flux from being collapsed into a single untyped signal.

**Channel 1: UV excitation (source=1).** This channel is active during UV bursts only. It detects aromatic excitation within a 4.0  $\text{\AA}$  cooperative detection radius and records wavelength, aromatic type, and vibrational energy. Cooperative enhancement is

$$f_{\text{coop}} = 1.0 + 0.3(n_{\text{excited}} - 1), \quad (3)$$

with an additional two-fold gain during active UV bursts.

**Channel 2: dewetting LIF (source=2).** This channel is active only when UV bursts are off. It implements a leaky integrate-and-fire neuron with effective membrane time constant  $\tau_{\text{eff}} = 0.15$  ps, threshold  $\theta = 0.5$ , and a 250-step refractory period. The membrane potential evolves as

$$V_m(t + \Delta t) = V_m(t)e^{-\Delta t/\tau_{\text{eff}}} + S_{\text{combined}}(t), \quad (4)$$

where  $S_{\text{combined}}$  integrates bidirectional water-density change and exclusion-field deviation from bulk with temporal and exclusion gains. A spike is emitted when  $V_m \geq \theta$ .


**Channel 3: electrostatic flux probe (source=3).** The EFP channel is active when at least two charged atoms are within range. It computes Coulomb potential with the Warshel dielectric,

$$\varepsilon_r(r) = \max(4r, 4), \quad (5)$$

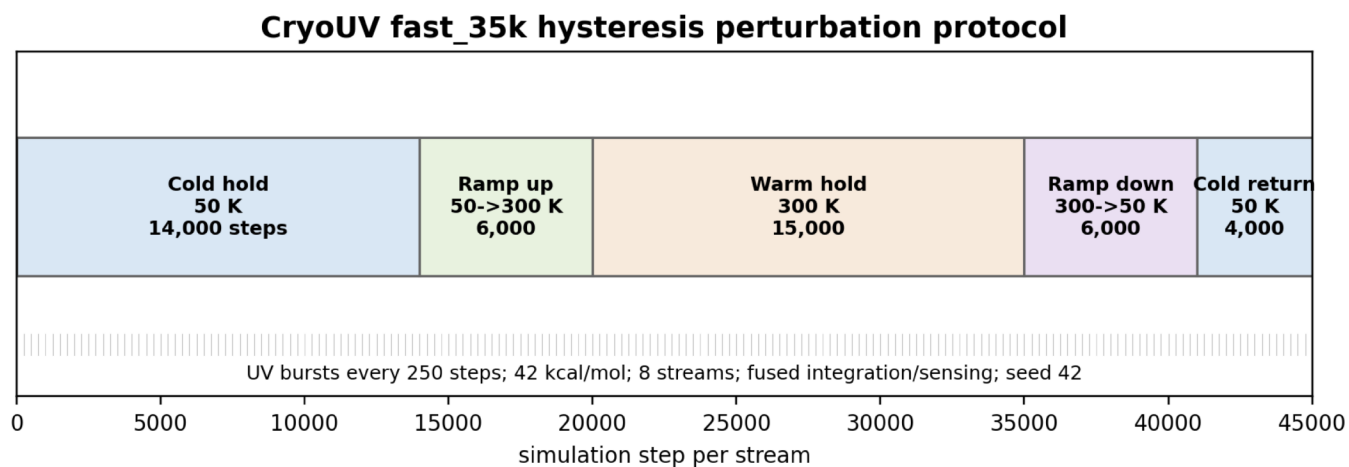
and integrates temporal flux (50 $\times$  gain) plus water-electrostatic coupling (10 $\times$  gain) through a dedicated LIF neuron with  $\tau = 0.5$  ps and  $\theta = 0.8$ . It emits typed CATION or ANION spikes. Each spike stores 96 bytes of metadata: position, intensity, source channel, wavelength, aromatic type, water density, vibrational energy, and up to eight nearby residue IDs.

## 2.6 UV-Activated Benzene Cosolvent Probing

PRISM-4D extends mixed-solvent molecular dynamics logic by making cosolvent probes UV-active. Virtual benzene probes are injected at surface-exposed hydrophobic



**Figure 2: Five-phase cryo-thermal hysteresis protocol.** The architecture benchmark uses 50 K cold holds, a 50→300 K ramp, a 300 K warm hold, a 300→50 K ramp, and a cold return, with UV bursts interleaved across aromatic wavelengths.



**Figure 3: CryoUV fast\_35k validation protocol.** The firewalled blind-validation artifacts retain the fast\_35k hysteresis schedule used for frozen apo/pre-ligand prediction, separating perturbation generation from post-freeze holo scoring.

aliphatic residues (LEU, ILE, VAL, ALA, MET) that lack native UV chromophores. The probes use benzene photophysical parameters ( $\lambda_{\max} = 254 \text{ nm}$ ,  $\varepsilon = 204 \text{ M}^{-1}\text{cm}^{-1}$ ,  $\eta = 0.71$ ). This extends perturbation coverage to aliphatic-dominated interfaces where native aromatic UV response would be sparse. In the architecture benchmark, MYC-MAX benzene cosolvent activation expanded UV coverage from 6 to 41 targets and generated 159,158 benzene-channel spikes, representing 61.4% of total signal at the detected site.

## 2.7 Hydrophobic Exclusion Water Inference

Rather than simulating explicit solvent, PRISM-4D infers water density from the hydrophobic exclusion field: the negative space created by the protein hydrophobic surface. Hydrophobic atoms define the exclusion boundary at an 8.0 Å cutoff, and water density is inferred from the complement of this field through GPU-accelerated geometric kernels on a 0.5 Å grid. This is an intentional approximation and is retained as a limitation in the discussion.

## 2.8 Site Scoring

Detected sites are characterized by geometry, chemistry, event intensity, and residue composition. The druggability score  $D$  is

$$D = 0.30S_V + 0.20S_E + 0.25S_H + 0.25S_A, \quad (6)$$

where  $S_V$  is a piecewise volume score,  $S_E$  is enclosure fraction,  $S_H$  is hydrophobicity from normalized spike intensity, and  $S_A$  is aromatic enrichment. Sites are classified as druggable when  $D \geq 0.48$  and volume is 50–3,000 Å<sup>3</sup>. The quality score  $Q$  is

$$Q = 0.25 \min(N_s/100, 1) + 0.25 \min(\bar{I}/10, 1) + 0.30D + 0.20S_A. \quad (7)$$

Cys, Lys, Ser, and His residues lining the pocket are automatically flagged as potential covalent-design opportunities.

## 3 Spike-Native Density Clustering & Manifold Families

SNDC is a four-stage GPU-accelerated clustering algorithm purpose-built for typed neuromorphic spike data. Unlike generic clustering applied after simulation, SNDC is architecturally integrated with spike generation and uses spike-native metadata including source channel, intensity, wavelength, and timestamp.

### 3.1 Stage 3a: RT-DBSCAN

Spike positions are represented as spheres in a bounding-volume hierarchy constructed for NVIDIA OptiX RT cores. Epsilon-neighborhood queries become ray-sphere intersection tests executed in dedicated ray-tracing hardware, reducing practical neighbor-search scaling from quadratic behavior toward linear traversal over the BVH. The adaptive epsilon parameter is

$$\epsilon = 3.0 \left( \frac{500}{N_{\text{atoms}}} \right)^{1/3}, \quad \epsilon \in [1.2, 3.0] \text{ \AA}. \quad (8)$$

This size-normalized radius preserves comparable clustering behavior across small domains and larger complexes. Applying RT-core BVH traversal to molecular event-neighborhood queries is a central engineering contribution of the platform.

### 3.2 Stage 3b: Watershed Segmentation

When a DBSCAN cluster spans multiple binding subsites, density-peak seeded watershed segmentation partitions it. Spike density is computed on a 3 Å voxel grid with 26-connected neighbors. Local density maxima serve as seeds, and voxels flow by steepest descent toward peaks. This resolves multi-subsite mega-pockets that confound single-centroid reporting and was essential for correcting oversized BACE1-like cavities in the architecture benchmark.

### 3.3 Stage 3c: Eikonal BFS

Distance-field propagation from each watershed-separated cluster boundary computes geodesic fields within the pocket. Expansion cost is weighted by inverse spike density, so boundaries form along low-density ridges rather than purely geometric midpoints. The resulting fields support pocket-depth, entry-channel, enclosure, and accessible-volume descriptors directly from spike density.

### 3.4 Stage 3d: Peak Centroid Tracking

The thermodynamic hotspot is localized by intensity-squared centroiding,

$$\mathbf{C} = \frac{\sum_i I_i^2 \mathbf{r}_i}{\sum_i I_i^2}. \quad (9)$$

Squared weighting pulls the reported center toward high-intensity event loci rather than the volumetric center of an irregular cavity. This is critical for large grooves and allosteric channels where a geometric centroid can fall 10–15 Å from the ligand-relevant locus.

### 3.5 Manifold Families Across Eight Independent Streams

The blind-validation runs retain eight independent stochastic streams per target wherever the frozen artifact set provides multistream evidence. SNDC is first applied independently within each stream so that local spike bursts, phase-specific centroids, and source-channel composition remain stream-native. Only after per-stream clustering are candidate sites collapsed into *manifold families*: recurrent spatial/residue-shell neighborhoods whose centers overlap within the expected conformational tolerance and whose event signatures recur across streams or phases.

This family collapse prevents a dynamic pocket from being treated as eight unrelated static centroids. A family stores the distribution of phase centers, intensity-weighted centroid dispersion, spike-source composition, stream consensus, phase modulation, KCC/driver factors, desolvation factors, burst factors, and manifold completeness. The result is a dynamic conformational envelope: a compact representation of where the ligandable state appears, how broad the event cloud is, which residues repeatedly drive it, and whether the site is a stable compact pocket, a distributed boundary case, or a weakly recurring geometric artifact.

## 4 Causal Drivers & PRISM-Therm/CCNS Descriptors

### 4.1 Transfer-Entropy Causal Driver Decomposition

PRISM4D reports driver residues as a causal interpretation layer over frozen event manifolds. The decomposition asks whether the recent event history of residue or neighborhood  $X$  improves prediction of manifold response  $Y$  beyond  $Y$ 's own history:

$$T_{X \rightarrow Y} = \sum p(y_{t+1}, y_t^{(k)}, x_t^{(l)}) \log \frac{p(y_{t+1} | y_t^{(k)}, x_t^{(l)})}{p(y_{t+1} | y_t^{(k)})}. \quad (10)$$

The implementation carries this concept into practical descriptors through top driver residues, driver centroids, KCC-driver factors, burst factors, phase modulation, and residue-neighborhood recurrence. Driver roles are interpreted as **triggers** when a residue neighborhood precedes event onset, **stabilizers** when it maintains recurrent spike intensity or water-exclusion coherence after opening, and **gateways** when it sits on an entry channel or conformational bottleneck connecting exposed solvent to the event manifold. These roles are mechanistic annotations: they explain why a pocket opens or persists, but they do not alter the frozen SR@K ranks.

### 4.2 PRISM-Therm/CCNS Descriptor Layer

The Therm/CCNS descriptor layer was the major evidence-completeness improvement in the validation package. The final augmentation matched 1295/1295 validation rows using frozen target-specific binding-site descriptor files. The descriptor layer includes `ccns_tau`, hysteresis, TIDE coupling, effective  $\Delta G$  and component  $\Delta G$  terms, UV enrichment, solvent frustration, water-displacement coherence, onset and breathing scores, spatiotemporal intensity counts, kinetic accessibility, aromatic score, source diversity, spike count, druggability, volume, sphericity, and burial-related fields. These descriptors are not a second scoring system; they are post-freeze mechanistic annotations.

## 5 Firewalled Validation & Data Leakage Monitoring

### 5.1 Claim Hierarchy and Freeze Rules

The blind validation is organized by evidence hierarchy. The primary quantitative claim comes only from frozen candidate ranks scored after validation-coordinate access. Post-freeze manifold evidence, Therm/CCNS descriptors, med-chem anchors, strict null controls, PFR, and structural renderings explain mechanism and failure modes, but do not change frozen ranks, target inclusion, primary endpoint, or success labels.


The global prediction freeze records freeze commit `c8b09678`. The freeze scope included `B01_HRAS.Q61H`, `B02_CDK2.allosteric`, `B03_Kv1.2`, `B04_MDM2`, `B05_TP53.R175H`, `B06_cGAS`, `B07_TEAD1`, `B08_CRBN`, `B09_Thrombin.exosite`, and `B10_ADRB2`. `B10_ADRB2` is retained as a hard-negative/stress-test calibration target and excluded from the 9-target primary macro average.

The firewall is explicitly temporal and cryptographic. It follows an  $A \rightarrow B \rightarrow C \rightarrow D$  sequence: apo-only prediction generation, prediction freeze and hash attestation, post-freeze holo-coordinate extraction, and locked scoring/reporting. Every target-specific freeze directory contains a `PREDICTION_FREEZE` attestation and SHA-256 manifest of prediction artifacts. The `B01` freeze record, for example, attests that apo predictions were generated by the frozen engine before holo or ligand-bound coordinate access and lists the frozen `binding_sites`, KCC validation, KCC visualization, PRISM-Therm topology, run log, and manifest artifacts.

### 5.2 Blind Panel Composition

The panel was deliberately mixed rather than curated for uniform success. It contains compact allosteric targets, partial recoveries, shallow/PPI limitations, a lipid-

## Post-freeze Therm/CCNS descriptor layer (1295/1295 rows matched)



**Figure 4: Post-freeze Therm/CCNS descriptor summary.** The final join audit matched 1295/1295 rows over all ten validation targets and summarized  $\tau$ , hysteresis/TIDE behavior, effective  $\Delta G$ , spike density, and related mechanistic descriptors.

**Table 2:** Therm/CCNS descriptor summary over validation-scored sites.

Target	Sites	CRYPTIC	RESP	DYN	Median $\Delta G$	Median tau	Median TIDE	Median spikes
HRAS.Q61H	10	0	4	3	1.180	1.327	0.70	2.18M
CDK2_allosteric	30	1	10	10	1.339	1.239	0.60	2.05M
Kv1.2	30	4	10	5	1.208	1.257	0.40	0.95M
MDM2	10	0	3	5	1.331	1.366	0.50	1.83M
TP53_apo	27	3	7	6	1.240	1.247	0.60	1.82M
cGAS	30	0	12	6	1.196	1.235	0.60	1.17M
TEAD1	29	1	13	9	1.364	1.246	0.60	1.82M
CRBN	30	1	17	6	1.264	1.227	0.40	1.96M
Thrombin_exosite	30	3	12	5	1.273	1.248	0.50	1.91M
ADRB2	30	0	4	11	1.263	1.216	0.80	2.08M

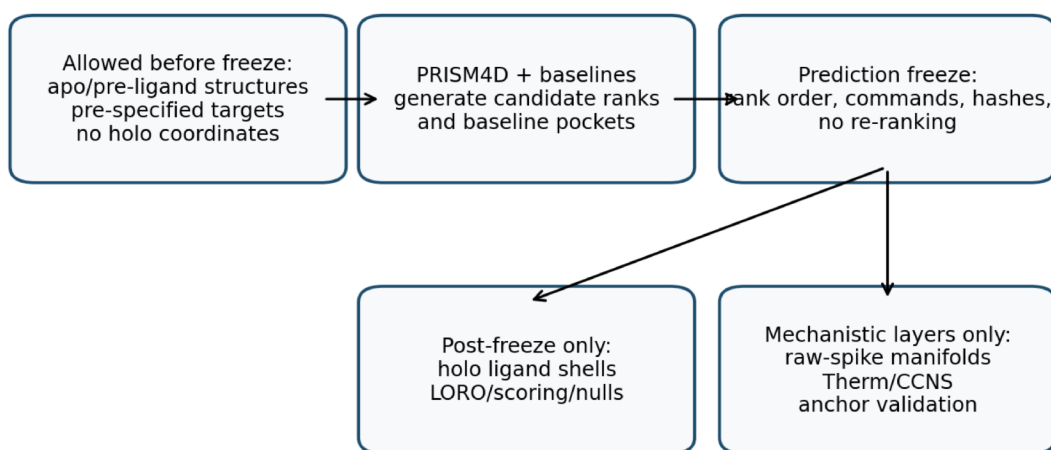
**Table 3:** A $\rightarrow$ B $\rightarrow$ C $\rightarrow$ D data-leakage monitoring timeline.

Stage	Time boundary	Permitted data	Leakage guard
A: apo generation	Before validation access	Apo/pre-ligand structures, engine settings, frozen baselines	No holo ligands or ligand shells available
B: cryptographic freeze	Before holo extraction	Candidate ranks, baseline outputs, scoring definitions, SHA-256 manifests	Immutable manifests and freeze attestations
C: holo extraction	After freeze only	Holo coordinates, ligand shells, LORO references	Read-only scoring inputs; no rank edits
D: locked reporting	After scoring scripts	Strict SR@K, nulls, Therm/CCNS joins, anchors, PFR, renders	Post-freeze layers cannot alter primary ranks or endpoint definitions

**Table 4:** Validation evidence hierarchy.

Layer	Input timing	Permitted interpretation
Primary Strict SR@K	Frozen before holo access; scored after holo access	Headline quantitative validation
Manifold evidence	Raw spike/KCC artifacts processed after freeze	Explains success, failure, and boundary behavior
Therm/CCNS descriptors	Joined after freeze from frozen descriptor files	Mechanistic thermodynamic/accessibility annotation
Med-chem anchors	Apo-derived anchors compared to holo contacts after freeze	Residue-level support, not definitive pharmacophore validation
Failure layer	All failures and edge cases retained	Defines operating regime and prevents overclaiming

## Firewalled prospective-retrospective validation design



**Guardrail: post-freeze data explains results but cannot alter the frozen primary Strict SR@K endpoint.**

**Figure 5: Firewalled prospective-retrospective validation design.** Holo coordinate access, ligand shell construction, LORO scoring, null controls, Therm/CCNS joins, and structural interpretation occur only after prediction freeze.

allosteric limitation, a distributed ion-channel boundary case, a peptide-exosite edge case, and a calibration GPCR.

### 5.3 Strict Scoring Definitions

Strict scoring used EXCELLENT hits only: minimum site-to-ligand distance  $\leq 4.0$  Å and shell\_intersection  $\geq 3$ . STRONG and NEAR tiers were retained for practical and exploratory interpretation, but not counted as primary successes. This prevents conversion of weaker proximity evidence into inflated accuracy.

### 5.4 Null Control Definitions

The second-pass strict null audit replaces weak row-level controls with harder reviewer-facing controls. It tests whether random reassignment can reproduce PRISM4D’s combined rank, distance, shell-overlap, KCC-driver, LORO, and family-recurrence structure. Three controls are defined:

1. **Strict rank permutation:** permutes evidence bundles across ranks to test whether strict evidence is concentrated near top-ranked candidates.
2. **Shell pair-breaking null:** breaks pairing between geometric shell metrics and PRISM4D-specific orthogonal signals while preserving marginal distributions.
3. **Family recurrence null:** tests whether multi-ligand family recurrence exceeds random reassignment of family ligand counts and sizes.

Target/control rows are interpreted only when observed values exceed the null 95% interval and empirical  $p$  values are  $\leq 0.05$ . Degenerate rows and zero-baseline single-instance rows are labeled rather than converted into primary evidence.

## 6 Primary Blind Validation Results

Across the 9-target primary blind panel, PRISM4D achieved Strict SR@1 = 0.483 compared with fpocket = 0.224 and P2Rank = 0.382. Strict SR@3 values were PRISM4D = 0.483, fpocket = 0.335, and P2Rank = 0.382. Strict SR@5 values were PRISM4D = 0.483, fpocket = 0.465, and P2Rank = 0.382. HRAS.Q61H and CDK2\_allosteric were flagship successes; TP53\_apo, cGAS, and CRBN showed partial recovery; MDM2 and TEAD1 were limitation cases; Kv1.2 and thrombin exosite were boundary or edge cases.

### 6.1 Target-Level Interpretation

**HRAS.Q61H.** HRAS.Q61H is a RAS/allosteric flagship success. PRISM4D Strict SR@1/SR@5 = 1.000/1.000; fpocket = 0.750/1.000; P2Rank = 0.500/0.500. LORO

= 4/4 (100.0%). Ten sites were scored, including 4 RESPONSIVE sites, with 2 holo references and 4 ligand instances.

**CDK2\_allosteric.** CDK2\_allosteric is a kinase/allosteric flagship success. PRISM4D Strict SR@1/SR@5 = 1.000/1.000; fpocket = 0.000/0.750; P2Rank = 0.500/0.500. LORO = 2/4 (50.0%). Thirty sites were scored, including 1 CRYPTIC and 10 RESPONSIVE sites.

Together, HRAS.Q61H and CDK2\_allosteric are the primary blind allosteric discoveries in the validation panel. They combine frozen rank recovery, LORO rediscovery, structural render support, and post-freeze causal/manifold annotations without using those post-freeze layers to tune the primary endpoint.

**Kv1.2.** Kv1.2 is a distributed ion-channel boundary case. All three methods achieved 0.100/0.100 Strict SR@1/SR@5. LORO = 2/20 (10.0%). This is not a compact-pocket win.

**MDM2 and TEAD1.** MDM2 and TEAD1 are retained as limitation cases. PRISM4D underperformed both baselines on MDM2 and TEAD1 under strict primary scoring, with LORO = 0/4 and 0/3, respectively.

**TP53\_apo, cGAS, and CRBN.** These are supportive partial recoveries. PRISM4D recovered strict rank-level signal but does not convert these into flagship claims. CRBN ties P2Rank at rank 1 and fpocket improves by rank 5.

**Thrombin.exosite and ADRB2.** Thrombin exosite is a single-instance peptide-exosite edge case with strict EXCELLENT failure. ADRB2 is an orthosteric calibration target; rank recovery occurred, but LORO = 0/2 and it is excluded from primary macro averaging.

### 6.2 Strict Null Control Results


The final publication-safe null-control report used 10,000 randomization iterations and identified five rows at  $p \leq 0.01$  after display correction. The important methodological point is not that every significant row is primary evidence: degenerate rows are labeled and excluded from the headline claim. CDK2\_allosteric remains a strong primary blind recovery based on frozen Strict SR@K, but its degenerate composite null row is not used as flagship statistical evidence. Kv1.2 rows are boundary evidence; thrombin rows are single-instance edge-case evidence.

## 7 Residue-Level Med-Chem Anchor Recovery

To test whether frozen apo-derived manifold residues retain medicinal-chemistry relevance, a post-freeze anchor audit compared apo-derived anchor points against holo ligand-contacting residues. The validation used protein


**Table 5:** Blind validation panel. ADRB2 is calibration and excluded from the 9-target primary macro average.

Slot	Target	Class	Role
B01	HRAS_Q61H	RAS/allosteric	Flagship primary blind success
B02	CDK2_allosteric	kinase/allosteric	Flagship primary blind success
B03	Kv1.2	ion channel/distributed	Boundary case: distributed channel
B04	MDM2	PPI/shallow groove	Limitation: baseline-loss case
B05	TP53_apo	mutant tumor suppressor	Supportive partial recovery
B06	cGAS	innate immune/allosteric	Supportive partial recovery
B07	TEAD1	lipid/allosteric	Limitation: baseline-loss, caveated
B08	CRBN	E3 ligase/degrader-relevant	Supportive partial recovery
B09	Thrombin_exosite	peptide exosite	Single-instance edge case
B10	ADRB2	GPCR orthosteric calibration	Hard-negative/stress-test calibration




**Figure 6:** Primary frozen blind-panel Strict SR@K macro averages over B01–B09. ADRB2 is retained separately as calibration and excluded from the primary macro average.

## Target-level frozen strict recovery including failures and calibration target



**Figure 7: Target-level strict recovery matrix.** Flagship, partial, limitation, boundary, and calibration behavior remain visible rather than being collapsed into a single success claim.

## Boundary cases separate compact recovery from broader accessibility signal




**Figure 8: Strict versus practical ligand-level recovery.** Strict EXCELLENT scoring supplies the primary endpoint, while STRONG/NEAR practical tiers are retained only as secondary accessibility evidence.

## Flagship allosteric structural panels

Equal-size panels; each contains a zoomed close-up inset and explicit site/status label.

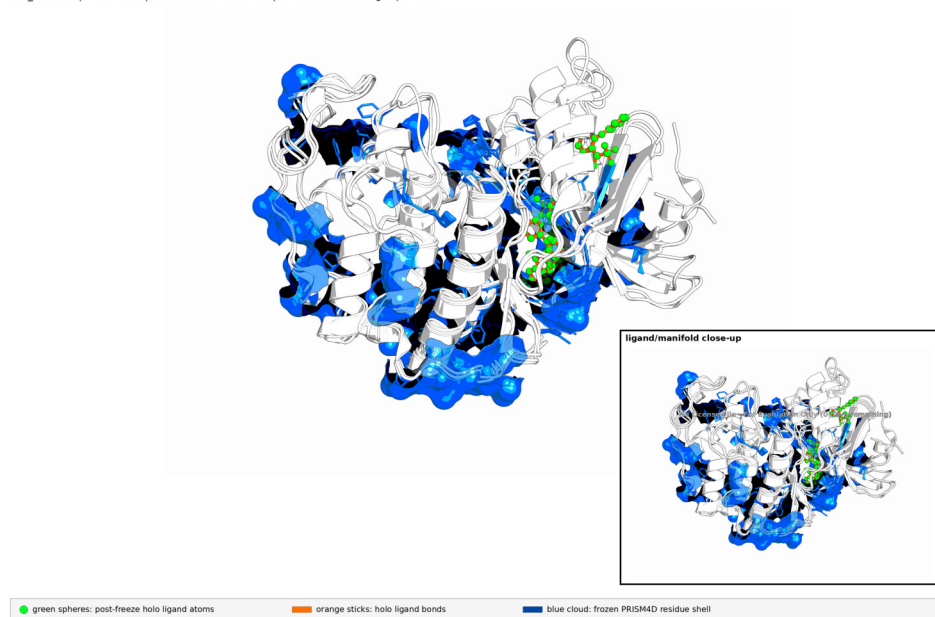
### A. HRAS\_Q61H flagship

site\_id=2007 | EXCELLENT | frozen residue shell=48 | RAS/allosteric flagship; LORO 4/4



### B. CDK2\_allosteric flagship

site\_id=1515 | EXCELLENT | frozen residue shell=74 | kinase/allosteric flagship; LORO 2/4



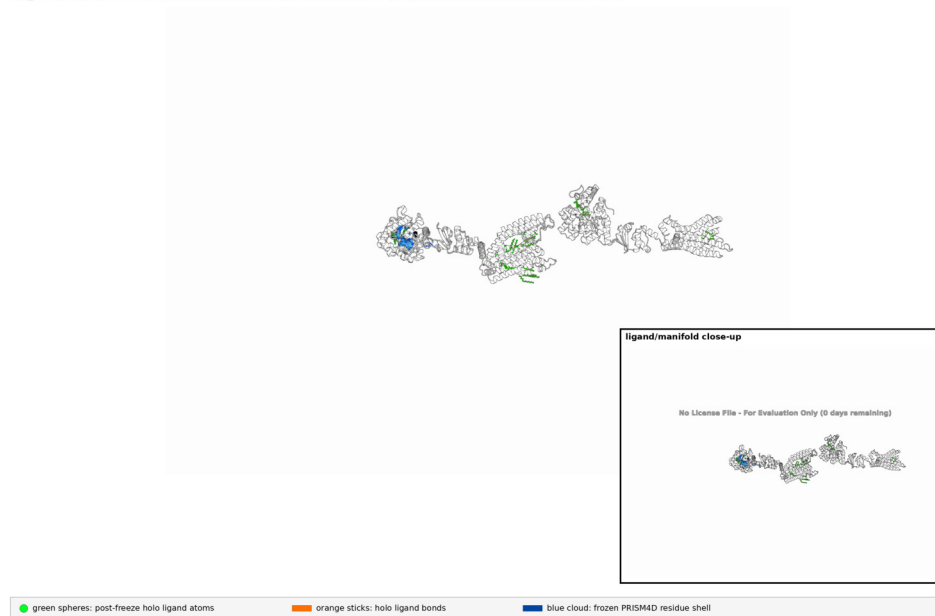
**Figure 9: Blind allosteric structural panels for HRAS\_Q61H and CDK2\_allosteric.** Frozen PRISM4D residue shells are rendered against post-freeze holo ligand overlays for flagship allosteric rediscovery. These panels are interpretive structural evidence and do not change frozen rank order.

## Boundary and calibration structural panels

Equal-size panels; each contains a zoomed close-up inset and explicit site/status label.

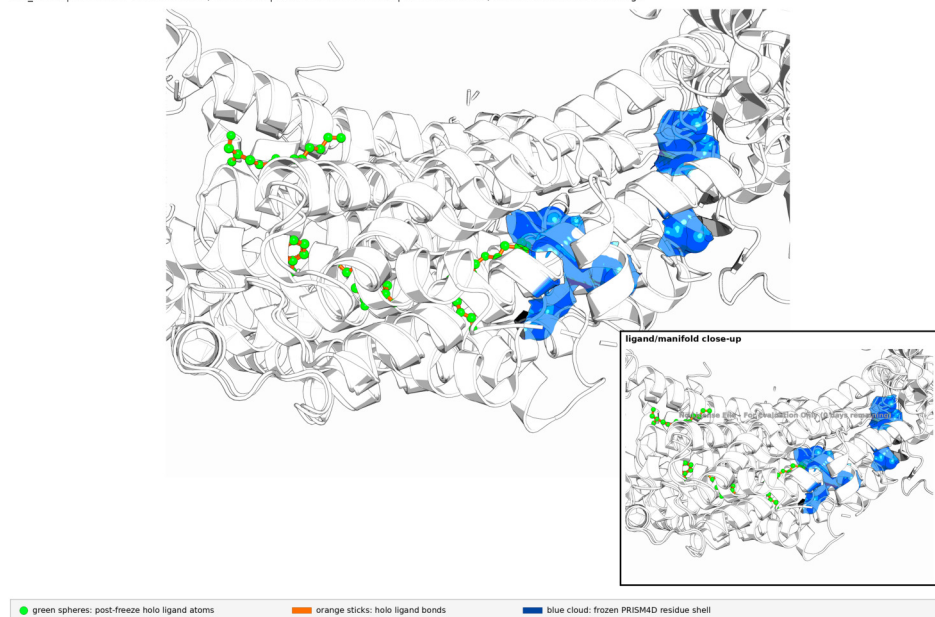
### C. Kv1.2 distributed boundary

site\_id=3518 | EXCELLENT local site; boundary target | frozen residue shell=18 | ion-channel boundary; target SR@1=0.10




### D. ADRB2 calibration

site\_id=12 | EXCELLENT orthosteric site; calibration | frozen residue shell=22 | GPCR calibration; excluded from macro average



**Figure 10: Boundary and calibration structural panels.** Kv1.2 and ADRB2 are shown as boundary/calibration examples rather than converted into universal success claims.



**Figure 11: Publication-safe strict-null interpretation.** Degenerate rows and edge-case rows are flagged rather than converted into primary evidence.

**Table 6:** Ligand-shell scoring tiers.

Tier	Definition	Role
EXCELLENT	$\text{min\_dist} \leq 4.0 \text{ \AA}$ and $\text{shell\_intersection} \geq 3$	Primary Strict SR@K
STRONG	$\text{min\_dist} \leq 6.0 \text{ \AA}$ and $\text{shell\_intersection} \geq 2$	Secondary signal
NEAR	$\text{min\_dist} \leq 8.0 \text{ \AA}$ and $\text{shell\_intersection} \geq 1$	Exploratory proximity
MISS	No qualifying contact	Failure under selected tier

residues within  $4.5 \text{ \AA}$  of holo ligand heavy atoms as the contact set. It evaluated 14,117 anchors across ten targets. This layer is deliberately residue-level and explanatory: it shows whether the frozen event manifolds recover subsets of ligand-contacting residue environments, but it is not used to modify rank order, define the primary SR@K endpoint, or claim affinity.

## 8 Mechanistic Proof: Vectorial Pharmacophore Recovery

### 8.1 Post-Freeze Audit Rationale

To test whether frozen nonequilibrium event manifolds encode ligand-relevant chemistry rather than only static cavity proximity, we performed a post-freeze blind Vectorial Pharmacophore Feature Recovery (PFR) audit on frozen rank-1 sites. The analysis was not used to select targets, alter SR@K ranks, redefine pockets, tune parameters, or modify event generation. Its purpose was mechanistic: to ask whether phase-ordered apo manifolds preserve directional pharmacophore information against in-

dependent holo ligand geometries.

### 8.2 Vectorial PFR Scoring

Unlike spherical proximity matching, Vectorial PFR requires class and direction agreement. A predicted apo-derived pharmacophore feature is counted as recovered only when the matched holo ligand feature satisfies all three criteria: matching pharmacophore class, centroid distance  $\leq 3.5 \text{ \AA}$ , and angular agreement within  $30^\circ$  of the predicted interaction vector. This prevents broad cavity overlap from being counted as chemical recovery and tests ligand-facing directional chemistry.

### 8.3 Temporal-Scramble Null Model

For each rank-1 site, 1,000 temporal-scramble decoy streams preserved the same 3D coordinates, residue identities, event intensities, and pocket location while randomizing timestamps and phase ordering. The null retains static structural envelope but removes nonequilibrium sequence information. The one-sided empirical test is

$$p = P(\text{PFR}_{\text{scramble}} \geq \text{PFR}_{\text{phase-ordered}}). \quad (11)$$

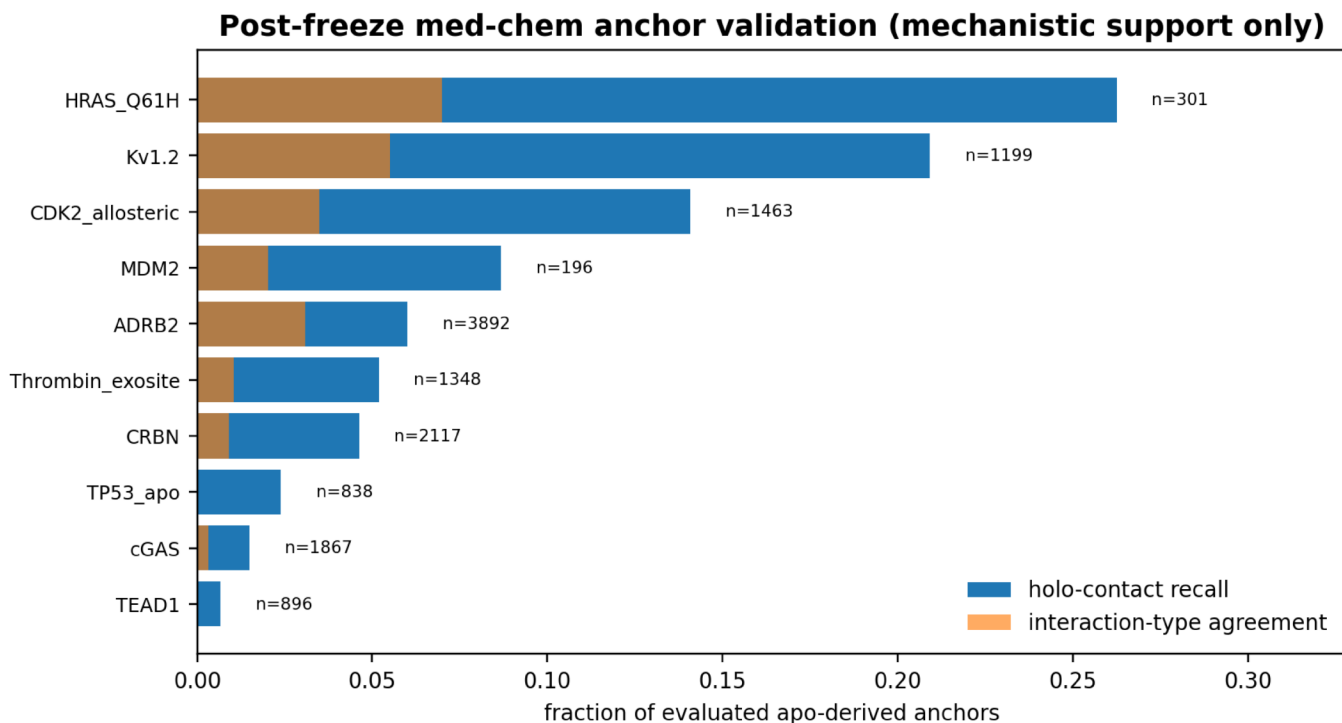
Using the conservative plus-one estimator  $(1 + \#\{\text{null} \geq \text{real}\}) / (n + 1)$  with  $n = 1000$ , zero exceedances correspond to  $p = 0.000999$ , reported in the main text as empirical  $p \leq 0.001$ .

**Table 7:** Primary frozen blind validation results. Columns P, F, and R denote PRISM4D, fpocket, and P2Rank strict recovery at the indicated rank. LORO is leave-one-reference-out rediscovery.

Slot	Target	P@1	F@1	R@1	P@3	F@3	R@3	P@5	F@5	R@5	LORO	Interpretation
B01	HRAS_Q61H	1.000	0.750	0.500	1.000	0.750	0.500	1.000	1.000	0.500	4/4 (100.0%)	Flagship primary blind success
B02	CDK2_allosteric	1.000	0.000	0.500	1.000	0.750	0.500	1.000	0.750	0.500	2/4 (50.0%)	Flagship primary blind success
B03	Kv1.2	0.100	0.100	0.100	0.100	0.100	0.100	0.100	0.100	0.100	2/20 (10.0%)	Boundary: distributed channel
B04	MDM2	0.250	0.500	0.500	0.250	0.500	0.500	0.250	0.500	0.500	0/4 (0.0%)	Baseline-loss limitation
B05	TP53_apo	0.500	0.000	0.000	0.500	0.250	0.000	0.500	0.500	0.000	2/4 (50.0%)	Supportive partial recovery
B06	cGAS	0.500	0.000	0.500	0.500	0.000	0.500	0.500	0.000	0.500	2/4 (50.0%)	Supportive partial recovery
B07	TEAD1	0.333	0.667	0.667	0.333	0.667	0.667	0.333	0.667	0.667	0/3 (0.0%)	Caveated baseline-loss limitation
B08	CRBN	0.667	0.000	0.667	0.667	0.000	0.667	0.667	0.667	0.667	2/3 (66.7%)	Supportive partial recovery
B09	Thrombin_exosite	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0/1 (0.0%)	Single-instance edge case
B10	ADRB2	1.000	0.000	0.500	1.000	0.000	0.500	1.000	0.500	0.500	0/2 (0.0%)	Calibration; excluded from macro

**Table 8:** Publication-safe strict null controls.

Target	Control	Metric	Observed	Null mean	95% interval	$p$	Use
CDK2_allosteric	shell pair-breaking	mean top-5 composite score	0.624	0.617	0.617–0.617	$p_{emp} < 10^{-4}$	Degenerate; not primary evidence
Kv1.2	shell pair-breaking	n strict rows	2	0.073	0.0–1.0	0.0007	Count-based boundary evidence
Kv1.2	family recurrence	recurrent family rows	2	0.021	0.0–0.0	0.0001	Zero-baseline boundary evidence
Thrombin_exosite	shell pair-breaking	n strict rows	1	0.000	0.0–0.0	$p_{emp} < 10^{-4}$	Single-instance edge; not success
Thrombin_exosite	shell pair-breaking	top5 strict	1	0.000	0.0–0.0	$p_{emp} < 10^{-4}$	Single-instance edge; not success

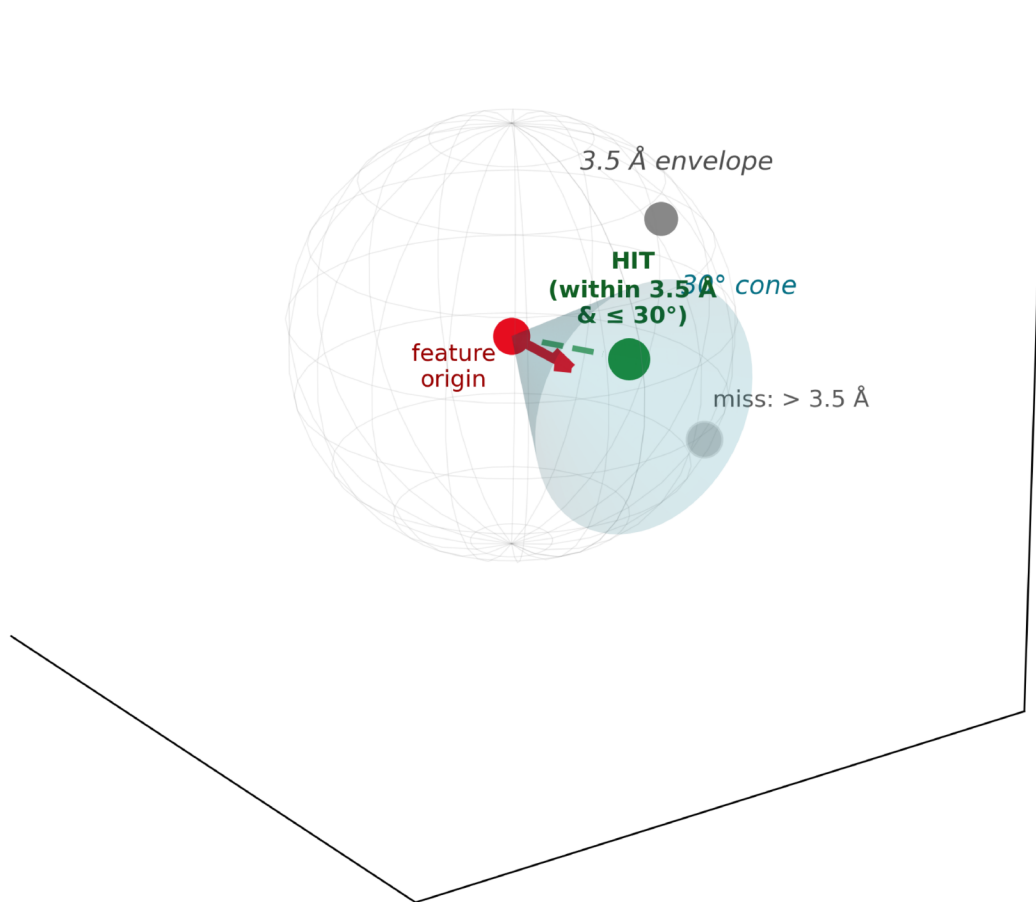
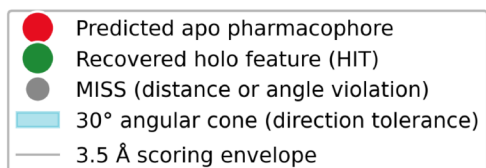


**Figure 12: Post-freeze med-chem anchor validation.** Apo-derived anchor contact recovery is present across all targets, while interaction-type agreement remains modest and should not be overclaimed.

**Table 9:** Residue-level med-chem anchor recovery against holo ligand-contacting residues.

Target	Holo PDBs	Anchors	Contact hits	Anchor recall	Type agreements	Type agreement rate
B01_HRAS_Q61H	2	301	79	0.2625	21	0.0698
B02_CDK2_allosteric	2	1463	206	0.1408	51	0.0349
B03_Kv1.2	2	1199	251	0.2093	66	0.0550
B04_MDM2	2	196	17	0.0867	4	0.0204
B05_TP53_R175H	2	838	20	0.0239	0	0.0000
B06_cGAS	2	1867	28	0.0150	6	0.0032
B07_TEAD1	2	896	6	0.0067	0	0.0000
B08_CRBN	2	2117	98	0.0463	19	0.0090
B09_Thrombin_exosite	2	1348	70	0.0519	14	0.0104
B10_ADRB2	2	3892	234	0.0601	120	0.0308

## A. Vectorial pharmacophore feature recovery (PFR) scoring criterion



**Figure 13: Vectorial PFR scoring criterion.** A predicted apo-derived pharmacophore feature is recovered only if class, centroid distance, and angular agreement all satisfy the specified criteria.

## 8.4 Headline PFR Result and Pharmacophoric Envelope

Across the multi-holo flagship set, phase-ordered PRISM4D manifolds achieved aggregate Vectorial PFR of 26.9%, compared with 2.2% under the temporal-scramble null. This is a 24.7 percentage-point absolute gain and a 12.2-fold enrichment. The result supports the “pharmacophoric envelope” interpretation: PRISM4D recovers a manifold-level directional interaction envelope, while individual feature-level confidence still requires calibration.

**Table 10:** Corrected Vectorial PFR results for the temporal-scramble null audit.

Target	PFR	Null	$\Delta$ pp	Fold	$p$
CDK2_allosteric	42.3	2.2	40.1	19.2 $\times$	$\leq 0.001$
Thrombin_exosite	39.7	2.2	37.5	18.0 $\times$	$\leq 0.001$
cGAS	26.5	2.2	24.3	12.0 $\times$	$\leq 0.001$
CRBN	19.8	2.2	17.6	9.0 $\times$	$\leq 0.001$
HRAS_Q61H	17.6	2.2	15.4	8.0 $\times$	$\leq 0.001$
TP53_apo	15.2	2.2	13.0	6.9 $\times$	$= 0.001$
<b>Aggregate</b>	<b>26.9</b>	<b>2.2</b>	<b>24.7</b>	<b>12.2<math>\times</math></b>	$\leq 0.001$

## 8.5 PFR Limitations and Provenance

Internal thermodynamic confidence weighting did not stratify recovered feature importance in the current audit: top-weighted features recovered holo interactions at 1.0%, compared with 1.4% for lower-weighted features. Therefore, PFR supports manifold-level pharmacophore recovery, not reliable ranked feature prioritization. A timestamp-order warning was recorded during post-freeze scripting; it reflected utility execution order and did not alter frozen apo outputs, primary ranks, pocket definitions, or validation scoring.

# 9 Highly Potential Novel & Cryptic Discoveries

## 9.1 SIRP $\alpha$ WYF Cryptic Pocket

Signal-regulatory protein alpha (SIRP $\alpha$ ) is a critical immune checkpoint in the CD47–SIRP $\alpha$  axis. A Diamond Light Source XChem campaign identified a WYF cryptic pocket through fragment screening of 800+ fragments with 27 hits. PRISM-4D detected this pocket as Site 1 from the closed apo structure (PDB 2WNG), where the pocket is sealed. The predicted site was top-ranked and 7.1 Å from the crystallographic ligand centroid; P2Rank 2.5 assigned only 0.2% probability to the same pocket. This is a central example in which event-driven perturbation identifies a cryptic, experimentally validated pocket that static methods down-rank.

**Table 11:** SIRP $\alpha$  WYF pocket detection from the closed apo structure.

Metric	PRISM-4D	P2Rank 2.5
Pockets detected	17	1
DCC to WYF pocket	7.1 Å	9.0 Å
Confidence	Site 1 (top)	0.2% probability
Spike count	4,413	N/A
Volume	743 Å <sup>3</sup>	–
Druggability	0.652	0.77
Runtime	18 min GPU	3 sec CPU

## 9.2 $\beta$ -Catenin Dual-Cysteine Prediction

$\beta$ -catenin is widely considered difficult to drug because of its elongated armadillo-repeat surface. PRISM-4D identified Site 6, a computationally predicted cryptic pocket ranked last by P2Rank (8/8, 0.5% probability). Site 6 contains TRP242 as an aromatic anchor, CYS240/CYS278 as a dual-cysteine covalent warhead pair, and LYS204 as an electrostatic anchor. It scored 0.965 quality and 0.844 druggability. This prediction requires experimental validation before therapeutic significance can be inferred.

**Table 12:**  $\beta$ -catenin Site 6 characterization.

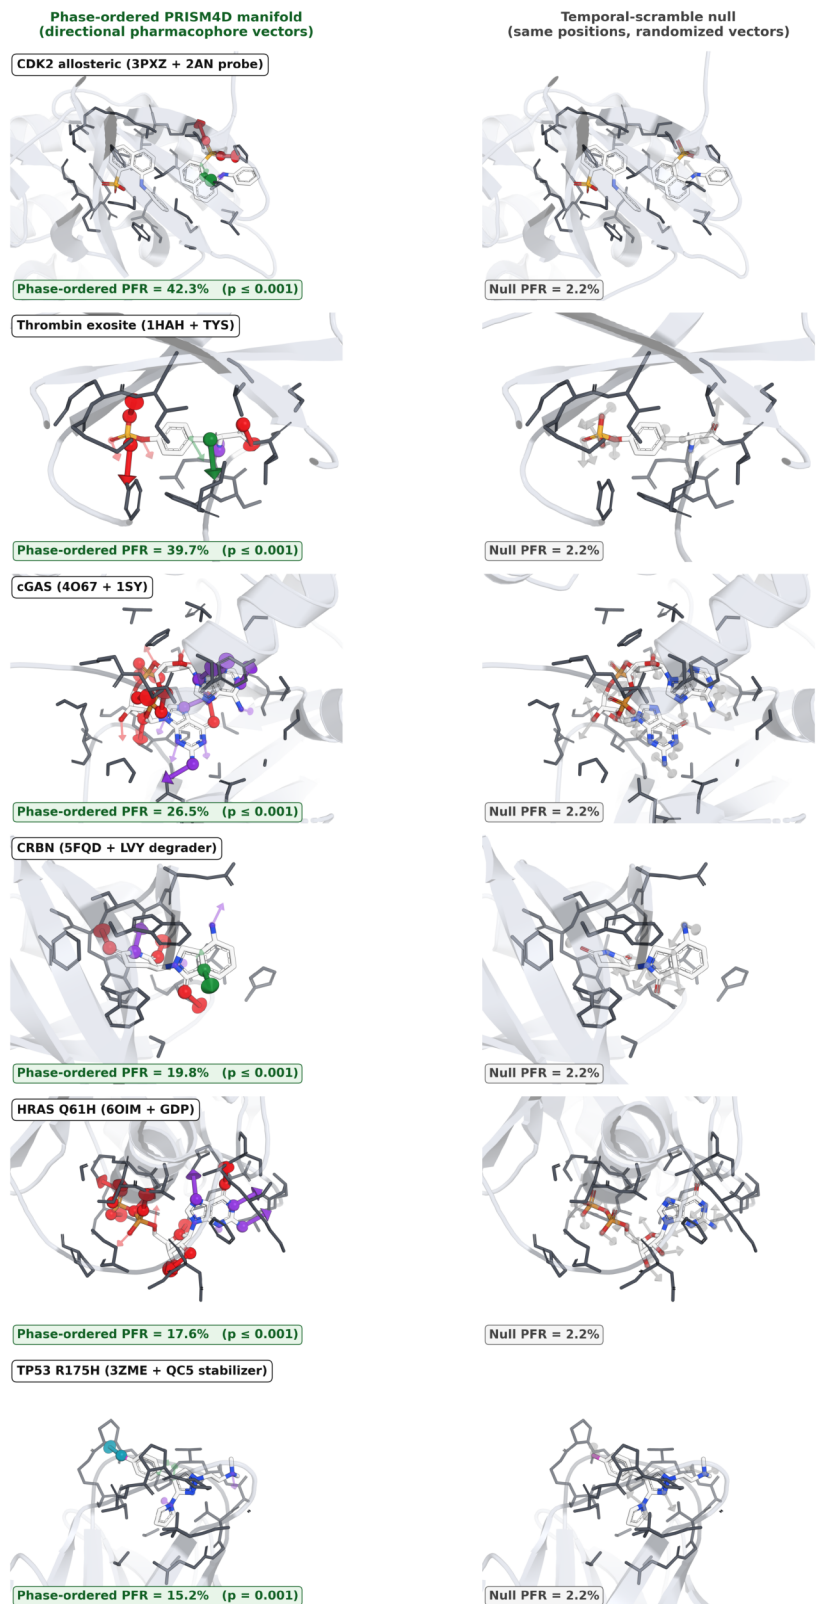
Property	PRISM-4D Site 6	P2Rank 2.5
Quality score	0.965	Rank 8/8
Druggability	0.844	0.5% probability
Classification	Cryptic	–
Volume	248 Å <sup>3</sup>	–
Spike count	3,242	N/A
Anchor residues	TRP242, CYS240, CYS278, LYS204	Not detected
Covalent pair	CYS240 + CYS278	–
Region	ARM repeats 2–5	–

Five independent stochastic runs converged to Site 6 with maximum pairwise centroid displacement of 0.06 Å, 23-fold smaller than a C–C covalent bond. Druggability was 0.844 in all runs, quality was 0.965 in all runs, and all four anchor residues persisted across 5/5 runs.

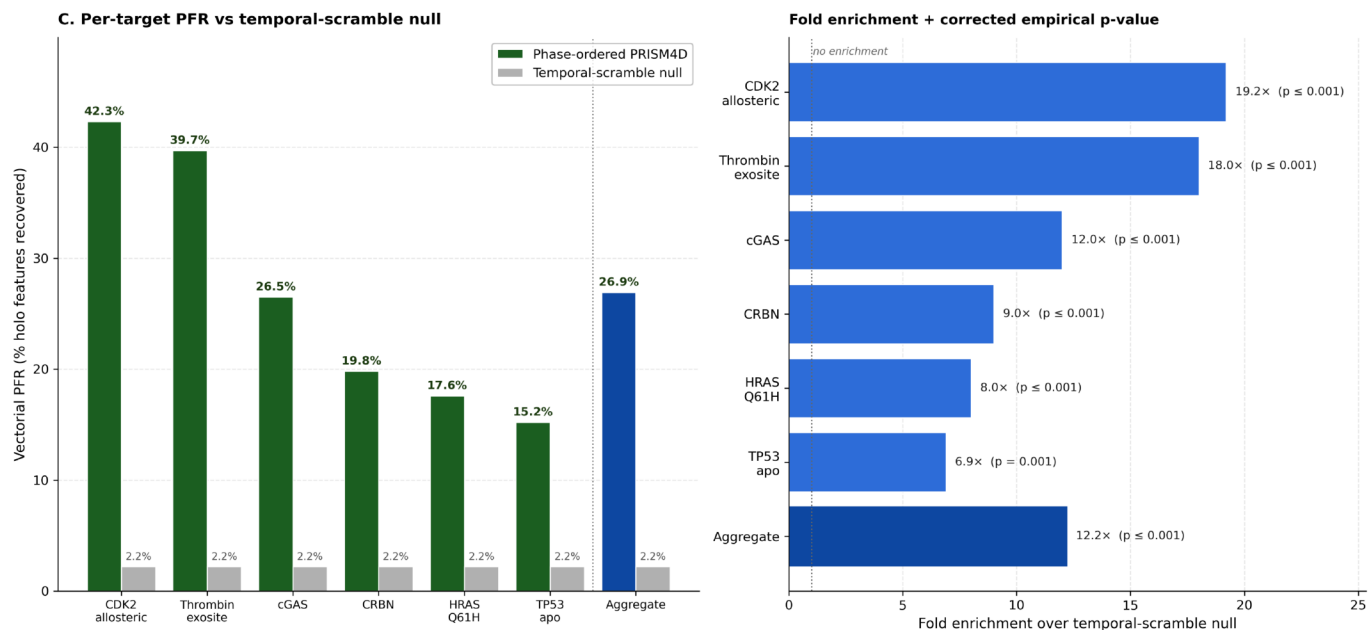
# 10 Comparison to Existing Methods

PRISM-4D differs from pocket detectors that rely on geometric cavity extraction or learned structural features. It is slower than fpocket or P2Rank but adds physics simulation, UV perturbation, event sensing, covalent residue identification, and RT-core spatial clustering while remaining executable on a consumer GPU.

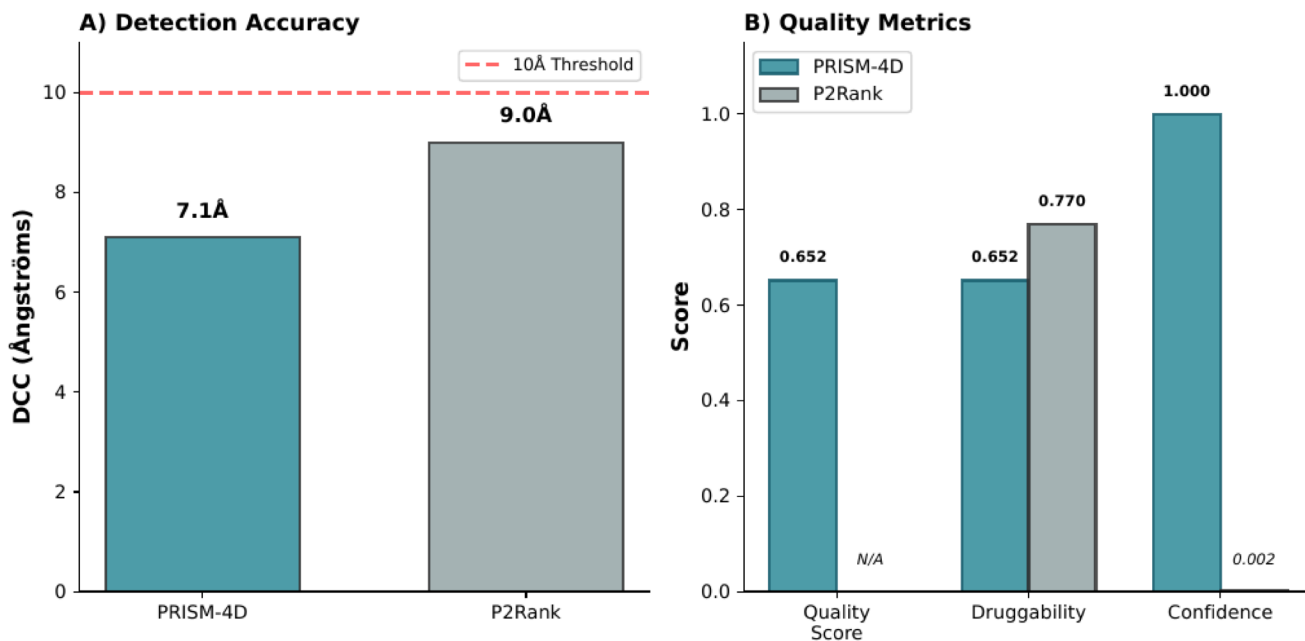
**B. Phase-ordered vs temporal-scramble vectorial pharmacophore recovery (matched cameras)**




**Figure 14: Phase-ordered versus temporal-scramble vectorial pharmacophore recovery.** Matched target panels show phase-ordered directional features and temporal-scramble nulls across CDK2 allosteric, thrombin exosite, cGAS, CRBN, HRAS Q61H, and TP53 R175H.



**Figure 15: Per-target PFR and fold enrichment.** Phase-ordered PRISM4D vectorial PFR is compared with temporal-scramble null PFR. All per-target enrichments are at least 6.9-fold, and aggregate enrichment is 12.2-fold.



**Figure 16: SIRP $\alpha$  WYF pocket comparison from the architecture source.** PRISM-4D detects the XChem-validated WYF pocket at 7.1 Å DCC from the closed apo structure, while P2Rank assigns 0.2% confidence to the same region.



**Figure 6. Reproducibility across 5 stochastic seeds ( $\beta$ -catenin Site 6).** (A) Centroid X–Y scatter. (B) Spike counts (CV = 0.2%). (C) Pairwise centroid displacement distribution (max = 0.06 Å).

**Figure 17:  $\beta$ -catenin Site 6 reproducibility across five stochastic seeds.** The source-architecture reproducibility panel shows centroid convergence, spike-count stability, and pairwise displacement bounded by 0.06 Å.

**Table 13:** Methodological comparison including hardware and deployment requirements.

Feature	FTMap	fpocket	P2Rank	PocketMiner	PRISM-4D
Approach	Probe docking	Voronoi geom.	ML + geom.	GNN	Physics sim.
Training data	None	None	sc-PDB	MD trajectories	None
Cryptic pockets	Limited	No	No	Yes (predicted)	Yes
Physics model	Implicit	None	None	None	AMBER ff14SB
UV spectroscopy	No	No	No	No	Yes
Spike detection	No	No	No	No	Yes (LIF)
Covalent ID	No	No	No	No	Yes
RT-core acceleration	No	No	No	No	Yes (BVH)
Consumer GPU	No	N/A	Yes	Yes	Yes
Offline operation	No	Yes	Yes	Yes	Yes
Runtime	~30 min	<1 sec	3–7 sec	~1 min	3–18 min

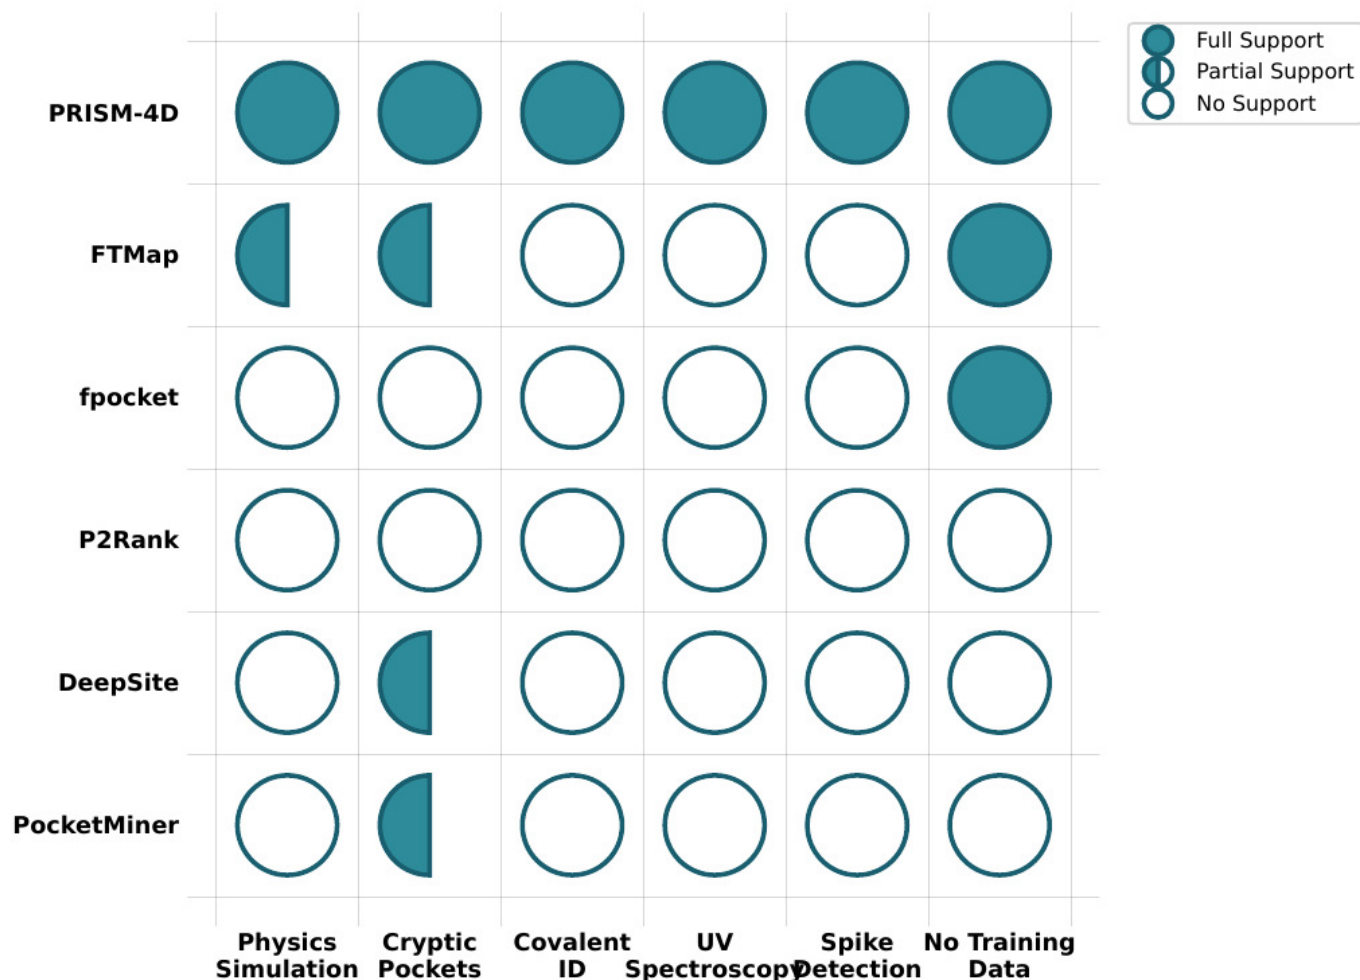


Figure 18: Methodological capability matrix. Filled circles indicate full support, half-filled circles partial support, and empty circles unsupported capabilities.

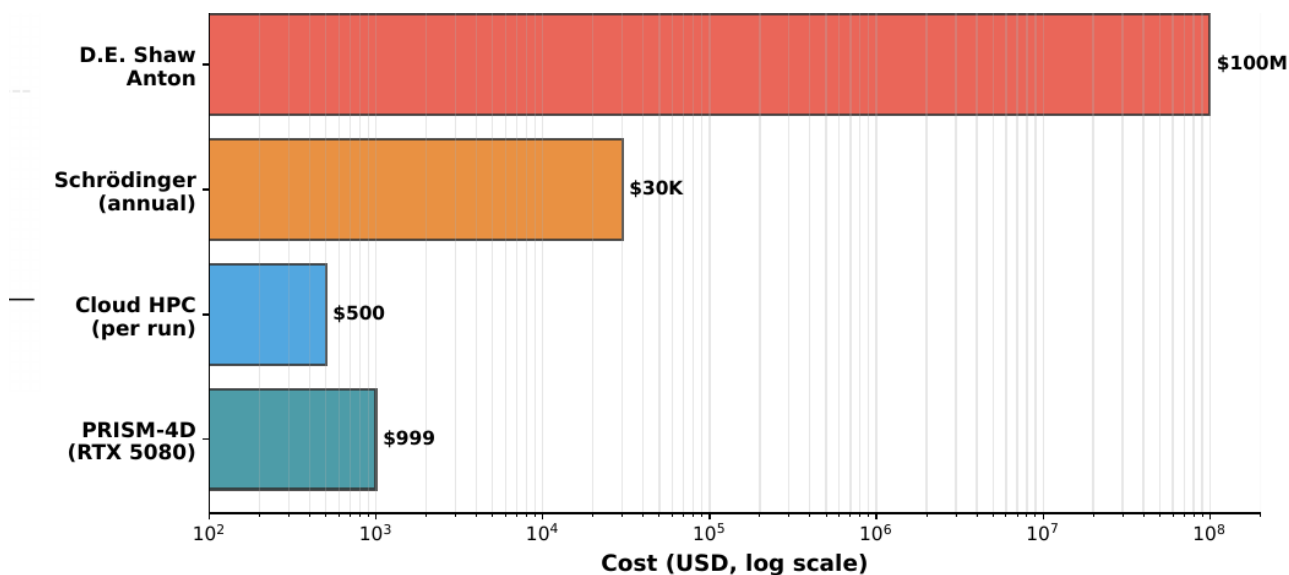


Figure 19: Hardware cost comparison for cryptic pocket analysis methods on a log cost scale.

**Table 14:** Hardware and cost comparison for cryptic pocket detection methods.

Method	Hardware	Runtime	Cost
PRISM-4D	Consumer GPU	3–18 min	\$999 GPU
Conventional MD	HPC cluster	hours–days	\$50–500/run
Schrodinger	Workstation	minutes	\$30K+/year
D. E. Shaw Anton	Custom ASIC	$\mu$ s	\$100M+
AlphaFold3 Server	Cloud TPU	minutes	free tier

## 11 Discussion, Reproducibility, & Data Availability

### 11.1 Operating Boundaries

The strongest defensible conclusion is not universal superiority. PRISM4D shows its largest advantage on compact allosteric and cryptic targets under firewalled validation. It is weaker or less decisive on shallow PPI grooves, lipid-allosteric cases, peptide exosites, distributed channel-state geometries, and context-sensitive targets. MDM2 and TEAD1 remain baseline-loss cases. Kv1.2 is a boundary case. Thrombin exosite is a single-instance edge case and is not converted into a hidden success. ADRB2 is calibration and excluded from the primary 9-target macro average.

The simulation protocol uses implicit solvent and no explicit membrane. Future work should test explicit-solvent, membrane-aware, and repeated independent-run variants. Parameters such as LIF threshold, UV burst energy, adaptive epsilon bounds, and hydrophobic-exclusion water inference require larger target-family sensitivity analysis. The engine currently depends on NVIDIA CUDA and RT-core hardware; AMD and Intel GPU support is not implemented.

### 11.2 Reproducibility Suite

The final evidence bundle audit reports 51 primary validation files, 9 freeze/firewall files, 556 baseline files, 8 null-control files, 340 manifold-evidence files, 4 Therm/CCNS descriptor files, 25 med-chem anchor files, 7 report files, and 2 figure-ready tables. The Therm/CCNS augmentation matched 1295/1295 rows. Med-chem anchor validation evaluated 14,117 apo-derived anchors. The PFR package includes deterministic feature builders, temporal-scramble null vectors, PyMOL render scripts, rendered high-resolution figures, corrected PFR tables, and one-sided plus-one p-value recomputation scripts.

### 11.3 Archive and SHA-256 Provenance

The reproducibility archive deposited with this manuscript carries a package-level SHA-256 manifest, `MANIFEST_SHA256.txt`, recording cryptographic hashes for every file in the deposit. The manifest is contained within the Zenodo deposit referenced in

Section 11.4 (DOI 10.5281/zenodo.20247900) and is also uploaded as a supplemental file with this manuscript so reviewers can verify file integrity inline without downloading the full archive. Two audit scripts bundled with the deposit, `audit/audit_reviewer_safe_package.py` and `audit/verify_post_md_lock.py`, allow programmatic verification of the reviewer-safe boundary and the post-MD pipeline lock from a local clone of the archive. The complete deposition record and tier-access structure are described in Section 11.4.

### 11.4 Data and Code Availability

PRISM-4D is developed by Delfictus I/O LLC. The engine source code, device kernels, model weights, production configuration, and internal scoring procedures are proprietary trade-secret material maintained in a private repository and are not part of this distribution. The licensing scope of this manuscript and the boundary between public reproducibility material and proprietary implementation are stated in the Distribution and Reuse License declaration and in the accompanying IP Boundary and Trade Secret Notice.

The validation evidence and reviewer materials accompanying this manuscript are organized as a three-tier access model. Reviewers and readers can locate the appropriate tier for the level of inspection they require:

#### **Tier 1 — Public reproducibility archive (Dataset).**

A reviewer-facing reproducibility bundle has been deposited at Zenodo under DOI 10.5281/zenodo.20247900 (accession 20247900, resource type: Dataset; deposition date 2026-05-16). The persistent DOI resolves to <https://doi.org/10.5281/zenodo.20247900>. The bundle contains: the firewalled validation outputs, candidate-rank tables, baseline-method comparison data, the 10,000-iteration strict null-control results, Therm/CCNS descriptor files, med-chem anchor validation outputs, vectorial PFR feature/null vectors, figure-generation scripts, structural-rendering scripts, package-level SHA-256 manifests, and the public freeze attestations. The deposited bundle is reviewer-safe and does not include engine binaries, raw simulation trajectories, raw spike or event arrays, raw binding-site JSONs, private scoring weights, or device-code internals; those omissions are documented as trade-secret boundaries rather than gaps in scientific disclosure.

#### **Tier 2 — Hosted live engine demonstration (request-based access).**

A hosted engine demonstration is available at <https://delfictus.com/I0hub>. The platform runs the frozen production PRISM-4D engine server-side on author-controlled infrastructure and returns sanitized, signed result manifests. Default-tier access exposes the fixed black-box modes documented in the reviewer-safe bundle (`blind_pocket_scan`, `validation_replay`, `pfr_audit_replay`); Tier 2 access extends these with relaxed configuration parameters, ad-

**Table 15:** Technical contributions classified by novelty type.

Contribution	Type	Description	Significance
UV pump-probe spectroscopy	Novel method	Wavelength-specific UV perturbation	Causal perturbation-to-dewetting signal
3-channel spike detection	Novel method	UV/LIF/EFP temporal isolation	Prevents cross-contamination artifacts
Electrostatic flux probe	Novel method	Warshel dielectric LIF neuron	First electrostatic event channel in this platform class
UV-activated cosolvent	Novel method	Benzene probes with UV energy	Extends to aliphatic interfaces
RT-core spatial clustering	Novel application	OptiX BVH for $O(N)$ queries	Estimated 10–60× speedup
Eikonal BFS	Novel application	Fast-marching-like fields on spike density	Pocket geometry from spike data
$I^2$ -weighted centroid	Novel application	Thermodynamic hotspot localization	3–15 Å improvement in irregular pockets
SNDC pipeline	Engineering	Four-stage integrated GPU pipeline	Purpose-built for spike data
Cryo-thermal hysteresis	Engineering	Five-phase 250 K cycling protocol	Nonequilibrium sampling
Covalent residue ID	Engineering	Cys/Lys/Ser/His flagging	Automatic covalent opportunity detection

**Table 16:** Evidence bundle summary.

Evidence item	Resolved status
Frozen primary validation	present; no post-freeze re-ranking
Baseline outputs	556 files
Publication-safe nulls	present; degenerate rows flagged
Therm/CCNS descriptors	1295/1295 rows matched
Manifold evidence	340 files
Med-chem anchors	14,117 anchors evaluated
PFR audit	corrected one-sided temporal-scramble results
Figure assets	high-resolution PNGs in <code>UnifiedFigures/</code>

ditional curated targets, and access to extended post-freeze run logs and intermediate manifold artifacts that are not redistributed in the public Zenodo bundle. Tier 2 access requests can be submitted through the platform website. In all hosted-demo tiers the engine binary, source code, device kernels, model weights, raw trajectories, and raw event arrays remain server-side and are not transferred to the reviewer.

**Tier 3 — Confidential inspection (controlled disclosure).** Inspection beyond what Tiers 1 and 2 provide — including raw execution logs, full backend command templates, model weights, internal scoring parameters, and engine source — is available only through a controlled process such as a non-disclosure agreement, counsel-supervised data room, or escrow arrangement, coordinated with the corresponding author.

## 11.5 Conclusions

This unified manuscript supports three linked claims. First, PRISM-4D is a full physics engine for event-driven binding-site detection, with UV perturbation, LIF spike sensing, UV-activated benzene probes, and SNDC RT-core clustering. Second, it is evaluated under a firewall that preserves primary blind failures, baseline losses, boundary cases, LORO outcomes, and strict null-control caveats. Third, post-freeze Vectorial PFR shows that phase-ordered manifolds encode directional pharmacophore information above a temporal-scramble null. The next required step is prospective wet-lab validation and larger target-family-stratified benchmarking.

## Acknowledgments

The author thanks the Diamond Light Source XChem team for making the SIRP $\alpha$  WYF fragment-screening data publicly available. Computational resources consisted of local desktop workstations with NVIDIA GPUs.

## Funding

This work was self-funded by Delfictus I/O LLC. No external grant, institutional, or government funding was received.

## Competing Interests

Ididia J. Serfaty is the inventor of PRISM-4D and the sole manager of Delfictus I/O LLC, the commercial developer of the platform. Delfictus I/O LLC has a direct commercial interest in PRISM-4D.

## Distribution and Reuse License

© 2026 Delfictus I/O LLC. All rights reserved. This manuscript is distributed solely for academic review, scholarly evaluation, and ordinary citation. No portion of this work may be reproduced, redistributed, modified, or used for commercial purposes without prior written permission from Delfictus I/O LLC. The PRISM-4D platform, including its source code, device kernels, model weights, configuration, execution settings, and internal scoring procedures, is proprietary trade-secret material of Delfictus I/O LLC and is not released under any open-source, public, or permissive license. The reviewer-facing reproducibility archive referenced in Section 11.3 contains only public summaries and post-freeze validation utilities and conveys no rights to the underlying engine. Preprint-server display of this manuscript is permitted under the server's standard terms of use; reuse beyond display and ordinary academic citation requires written authorization.

## References

- [1] Cimermancic, P. et al. CryptoSite: expanding the druggable proteome by characterization and prediction of cryptic binding sites. *J. Mol. Biol.* 428, 709–719 (2016).
- [2] Vajda, S., Beglov, D., Wakefield, A. E., Egbert, M. & Whitty, A. Cryptic binding sites on proteins: definition, detection, and druggability. *Curr. Opin. Chem. Biol.* 44, 1–8 (2018).
- [3] Maier, J. A. et al. ff14SB: improving the accuracy of protein side chain and backbone parameters from ff99SB. *J. Chem. Theory Comput.* 11, 3696–3713 (2015).
- [4] Le Guilloux, V., Schmidtke, P. & Tuffery, P. Fpocket: an open source platform for ligand pocket detection. *BMC Bioinformatics* 10, 168 (2009).
- [5] Krivak, R. & Hoksza, D. P2Rank: machine learning based tool for rapid and accurate prediction of ligand binding sites from protein structure. *J. Cheminform.* 10, 39 (2018).
- [6] Meller, A. et al. Predicting locations of cryptic pockets from single protein structures using the Pocket-Miner graph neural network. *Nat. Commun.* 14, 1177 (2023).
- [7] Kozakov, D. et al. The FTMap family of web servers for determining and characterizing ligand-binding hot spots of proteins. *Nat. Protoc.* 10, 733–755 (2015).
- [8] Warshel, A. & Levitt, M. Theoretical studies of enzymic reactions: dielectric, electrostatic and steric stabilization of the carbonium ion in lysozyme. *J. Mol. Biol.* 103, 227–249 (1976).
- [9] Sethian, J. A. A fast marching level set method for monotonically advancing fronts. *Proc. Natl. Acad. Sci. USA* 93, 1591–1595 (1996).
- [10] Diamond Light Source XChem team. SIRP $\alpha$  WYF cryptic pocket fragment-screening data release (2025).