Comparative Analysis of De Novo Antibody Design Platforms: JAM-2, Chai-2, Origin-1, and RFAntibody

With Comprehensive Oncology Target Annotations

Comprehensive Review of Generative AI Approaches
for Therapeutic Antibody Discovery

January 2026 — Version 7

Abstract

The field of computational antibody design has undergone a paradigm shift with the emergence of generative artificial intelligence platforms capable of designing therapeutic antibodies de novo. This comprehensive analysis examines five leading approaches: **JAM-2** from Nabla Bio, **Chai-2** from Chai Discovery, **Origin-1** from AbSci, and **RFAntibody** from the Baker Lab. We systematically compare their architectures, training methodologies, target selection strategies, and experimental validation results. JAM-2 demonstrates the highest reported hit rates (39% for VHH-Fc, 18% for mAb formats) with comprehensive developability profiling. Chai-2 achieves remarkable structural accuracy (<1.7 Å RMSD) and first-in-class functional GPCR agonist design. Origin-1 addresses the challenging "zero-prior" epitope problem with all-atom diffusion, while RFAntibody provides atomic-level precision validated by cryo-EM. Special emphasis is placed on oncology applications, including targeting of tumor-specific neoepitopes (KRAS G12V), immune checkpoint molecules (PD-L1), and G protein-coupled receptors implicated in cancer (GPRC5D, CXCR4). This analysis provides a framework for understanding the current landscape and future directions of AI-driven antibody therapeutics.

Keywords: de novo antibody design, generative AI, diffusion models, therapeutic antibodies, oncology, GPCR targeting, structure prediction

Contents

1 Introduction

Therapeutic antibodies represent one of the most successful classes of biopharmaceuticals, with over 100 approved products generating annual revenues exceeding $200 billion [Kaplon et al., 2024]. Traditionally, antibody discovery has relied on immunization campaigns, phage display, or hybridoma technology—approaches that are time-consuming, expensive, and often converge on limited epitope diversity [Bradbury et al., 2011].

The emergence of deep learning has fundamentally transformed protein structure prediction, exemplified by AlphaFold2's breakthrough in the CASP14 competition [Jumper et al., 2021]. This success has catalyzed efforts to extend these capabilities from structure prediction to structure generation, enabling the computational design of proteins with novel functions [Watson et al., 2023].

In 2024-2026, four distinct platforms have emerged as leading approaches for de novo antibody design:

1. **JAM-2** (Joint Atomic Modeling) from Nabla Bio: A general-purpose design system achieving double-digit hit rates across diverse targets [Nabla Bio, 2025].

2. **Chai-2** from Chai Discovery: An all-atom foundation model demonstrating atomically accurate predictions and functional GPCR agonist design [Chai Discovery, 2025].

3. **Origin-1** from AbSci: A platform specifically targeting "zero-prior" epitopes lacking structural precedent [AbSci, 2026].

4. **RFAntibody** from the Baker Lab: An RFdiffusion-based approach with cryo-EM-validated atomic precision [Bennett et al., 2025].

This review provides a systematic comparison of these platforms across multiple dimensions: architectural design, training methodology, target selection strategy, performance metrics, and oncology applications.

2 Antibody Biology Primer

Before examining computational design approaches, we provide an overview of antibody structure, function, and traditional discovery methods. This context is essential for understanding both the challenges these AI platforms address and the metrics by which they are evaluated.

2.1 Antibody Structure and Function

Antibodies (immunoglobulins) are Y-shaped glycoproteins produced by B cells as part of the adaptive immune response. Their primary function is to recognize and bind specific molecular targets (antigens) with high affinity and specificity, thereby neutralizing pathogens or marking them for destruction by other immune cells.

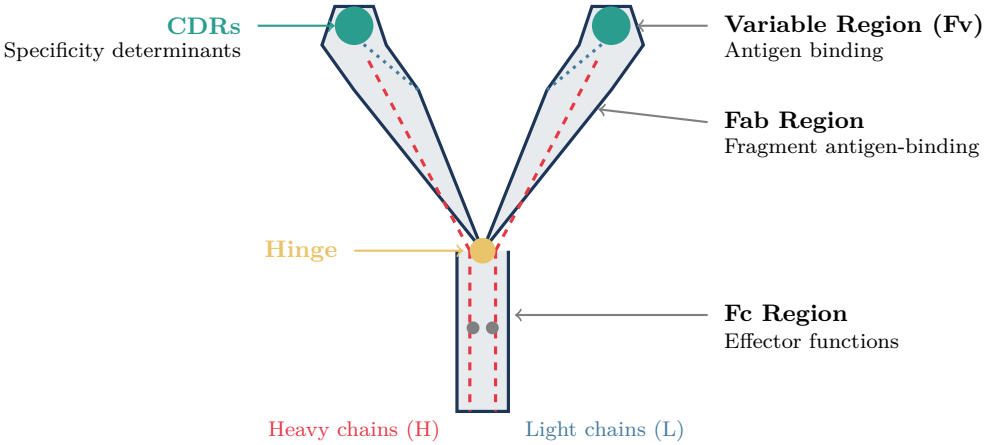

Figure 1: **Immunoglobulin G (IgG) structure.** The canonical antibody consists of two heavy chains (H, red dashed) and two light chains (L, blue dotted) arranged in a Y-shape. The Fab regions contain the variable domains responsible for antigen binding, with specificity determined by the complementarity-determining regions (CDRs, green). The Fc region mediates effector functions including complement activation and Fc receptor binding. Glycosylation sites (gray) in the Fc region influence pharmacokinetics and effector function.

2.1.1 Structural Hierarchy

The IgG antibody (the dominant therapeutic format) comprises:

- **Heavy Chains (H):** Two identical 450 amino acid chains, each containing one variable domain (V_H) and three constant domains (C_H1, C_H2, C_H3).
- **Light Chains (L):** Two identical 220 amino acid chains, each containing one variable domain (V_L) and one constant domain (C_L). Either kappa (κ) or lambda (λ) type.
- **Disulfide Bonds:** Inter-chain disulfides link H-H and H-L chains; intra-chain disulfides stabilize immunoglobulin domains.
- **Hinge Region:** Flexible linker between Fab and Fc, enabling bivalent binding to spatially separated epitopes.

2.1.2 Complementarity-Determining Regions (CDRs)

The antigen-binding site (paratope) is formed by six hypervariable loops—three from V_H (HCDR1, HCDR2, HCDR3) and three from V_L (LCDR1, LCDR2, LCDR3). These CDRs determine binding specificity:

Table 1: CDR characteristics and their roles in antigen recognition.

CDR	Length Range	Variability	Role in Binding
HCDR1	5–7 aa	Moderate	Often contacts antigen; canonical structures
HCDR2	16–19 aa	Moderate	Broad antigen contact; some canonical structures
HCDR3	3–30+ aa	**Highest**	Primary specificity determinant; no canonical structures
LCDR1	10–17 aa	Moderate	Antigen contact; canonical structures
LCDR2	7 aa	Low	Limited antigen contact
LCDR3	7–11 aa	Moderate	Antigen contact; canonical structures

HCDR3 is of particular importance for computational design: it is the most variable in both length and sequence, lacks predictable canonical structures, and typically makes the most extensive contacts with the antigen. Accurate modeling and design of HCDR3 represents a key challenge for all platforms reviewed here.

2.1.3 Framework Regions

The CDRs are embedded within four framework regions (FR1–FR4) that provide structural scaffolding. Framework sequences are more conserved and can be categorized into germline gene families. The platforms reviewed here typically use established frameworks (e.g., VH3-23, VH3-66) to ensure proper folding and reduce immunogenicity.

2.2 Alternative Antibody Formats

Beyond conventional IgG, several alternative formats are relevant to computational design:

Table 2: Antibody formats and their characteristics.

Format	Size (kDa)	Chains	Valency	Key Features
IgG (mAb)	∼150	4 (2H + 2L)	Bivalent	Full effector function; long half-life
Fab	∼50	2 (H + L)	Monovalent	No Fc; shorter half-life
scFv	∼25	1 (linked)	Monovalent	Single-chain; bacterial expression
VHH (Nanobody)	∼15	1	Monovalent	Camelid-derived; high stability
VHH-Fc	∼80	2	Bivalent	Nanobody + Fc; combines advantages

VHH (Nanobodies): Derived from camelid heavy-chain-only antibodies, VHHs consist of a single variable domain (∼110 amino acids). Their advantages include:

- Small size enabling tissue penetration and access to cryptic epitopes
- High thermostability and solubility
- Amenable to bacterial expression (no glycosylation required)
- Extended HCDR3 loops can penetrate concave epitopes (e.g., enzyme active sites)

JAM-2 and Chai-2 both demonstrate strong performance in VHH/VHH-Fc formats, which may reflect the reduced complexity of single-domain design.

2.3 Traditional Antibody Discovery Methods

Understanding traditional methods contextualizes the advances offered by computational design.

2.3.1 Hybridoma Technology

The classical approach, developed by Köhler and Milstein (1975, Nobel Prize 1984):

1. Immunize mice with target antigen over several weeks
2. Harvest splenic B cells producing antigen-specific antibodies
3. Fuse B cells with immortal myeloma cells to create hybridomas
4. Screen thousands of hybridoma clones for binding

5. Subclone positive hits and characterize

6. Humanize murine sequences for therapeutic use

Limitations: 12–24 month timelines; limited epitope diversity (immunodominant epitopes favored); requires humanization (immunogenicity risk); cannot target self-antigens or toxic proteins.

2.3.2 Phage Display

In vitro selection using bacteriophage-displayed antibody libraries:

1. Construct large combinatorial libraries (10^9–10^{11} variants)

2. Display scFv or Fab fragments on phage surface

3. Pan against immobilized antigen through multiple rounds

4. Sequence enriched clones and characterize

Advantages: Access to human germline sequences; can target self-antigens; no animal immunization required.

Limitations: Library diversity constrained by transformation efficiency; bacterial expression may miss some sequences; multiple rounds of panning required.

2.3.3 B Cell Sorting and Single-Cell Sequencing

Modern approaches directly isolate antigen-specific B cells from immunized animals or human donors, followed by single-cell sequencing of paired heavy/light chain genes. This preserves natural pairing but remains dependent on biological immune responses.

2.4 The Case for Computational Design

Traditional methods share fundamental limitations that computational approaches can address:

Table 3: Comparison of traditional versus computational antibody discovery.

Aspect	**Traditional Methods**	**Computational Design**
Timeline	12–24 months	4–8 weeks
Epitope control	Limited (immune response determines)	Precise (user-specified)
Epitope diversity	Often converges on immunodominant sites	Systematic coverage possible
Self-antigen targeting	Challenging (tolerance)	No biological constraints
Toxic/unstable targets	Difficult to immunize	Requires only structure
Developability	Post-hoc optimization needed	Can be designed in
Throughput	Limited by biology	Computationally scalable

2.5 Key Metrics for Evaluating Antibody Designs

The platforms reviewed here report various metrics. Understanding these is essential for interpreting results:

2.5.1 Binding Metrics

- **Hit Rate:** Percentage of designs showing detectable binding. The primary measure of design efficacy.
- **Dissociation Constant (K_D):** Equilibrium binding affinity; lower values indicate tighter binding. Therapeutic antibodies typically require <10 nM; best-in-class achieve pM.
- **On-rate (k_{on}):** Association rate constant. Fast on-rates are important for tissue penetration.
- **Off-rate (k_{off}):** Dissociation rate constant. Slow off-rates correlate with sustained target engagement.

2.5.2 Structural Metrics

- **RMSD (Root Mean Square Deviation):** Measures structural similarity between predicted and experimental structures. Sub-2 Å RMSD indicates high accuracy.
- **Interface RMSD:** RMSD calculated only for residues at the antibody-antigen interface. More relevant for binding prediction.
- **pLDDT:** Per-residue confidence score from AlphaFold-family models (0–100). Scores >70 indicate reliable predictions.

2.5.3 Developability Metrics

Therapeutic antibodies must satisfy multiple biophysical criteria:

- **Thermostability:** Melting temperature (T_m) >60°C preferred; ensures stability during manufacturing and storage.
- **Monomericity:** >95% monomeric by SEC; aggregation indicates instability.
- **Polyreactivity:** Low binding to off-target proteins (DNA, insulin, cardiolipin panel); high polyreactivity predicts fast clearance and toxicity.
- **Hydrophobicity:** Hydrophobic patches promote aggregation; HIC retention time used as proxy.
- **Expression Yield:** Adequate production in mammalian (CHO) or bacterial systems.

2.6 Manufacturing Considerations

Therapeutic antibodies are among the most complex pharmaceuticals to manufacture:

- **Expression System:** Most IgGs require mammalian cells (CHO, HEK293) for proper folding and glycosylation. VHHs and scFvs can be produced in bacteria.
- **Glycosylation:** N-linked glycans at Asn297 in the Fc region influence effector function and pharmacokinetics. Afucosylation enhances ADCC.
- **Scale:** Commercial production requires multi-thousand-liter bioreactors. Global capacity is approximately 30 metric tons annually.
- **Cost:** Manufacturing costs have decreased to $50–150/gram at scale, but drug pricing remains high ($1,000–$10,000+ per dose).

Computational design platforms that optimize for developability from the outset can significantly reduce downstream optimization cycles and manufacturing failures.

3 Architectural Approaches

3.1 Overview of Design Paradigms

The four platforms represent distinct architectural philosophies for antibody generation, as illustrated in Figure 2.

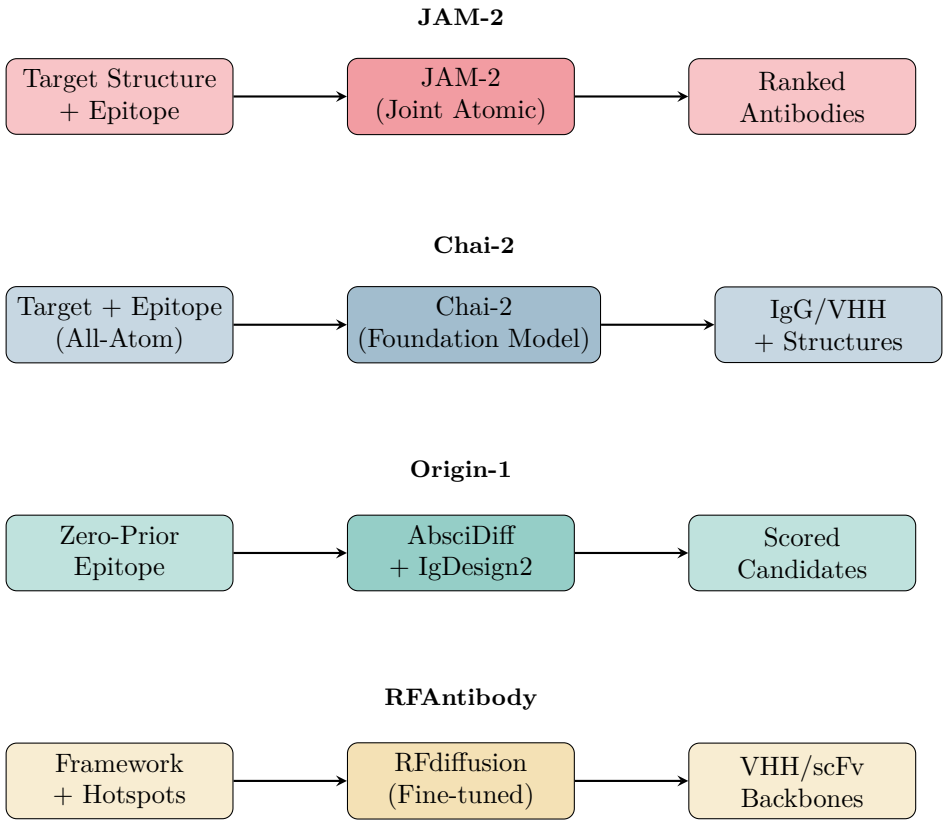

Figure 2: Schematic comparison of the four de novo antibody design pipelines. Each platform takes different input specifications and employs distinct generative architectures.

3.2 JAM-2: Joint Atomic Modeling

JAM-2 represents Nabla Bio's second-generation antibody design system. The system operates as a general-purpose de novo designer capable of generating both VHH-Fc and full-length mAb formats.

Key Architectural Features:

- Zero-shot design capability without target-specific training
- Complete CDR generation (IMGT definition) plus adjacent framework residues
- Generates ranked design sets (typically 45 designs per target per format)
- Design cycle completion in 2–3 days computationally

The system outputs novel sequences with >95% having RMSD >10 Å to the closest structure in SAbDab, demonstrating genuine novelty.

3.3 Chai-2: All-Atom Foundation Model

Chai-2 builds upon Chai-1, Chai Discovery's multimodal foundation model that demonstrated competitive performance with AlphaFold3 on structure prediction benchmarks.

Key Architectural Features:

- All-atom resolution modeling (not just backbone)
- Atomic-level reasoning about binding interfaces
- Direct full-length IgG format design
- Framework selection using well-characterized VH3-23 and VH3-66 germlines

3.4 Origin-1: Zero-Prior Epitope Targeting

Origin-1 addresses a specific challenge: targeting epitopes that lack any structural precedent from antibody-antigen or protein-protein complexes.

Platform Components:

1. **AbsciDiff**: All-atom structure generation via diffusion, fine-tuned from Boltz-1
2. **IgDesign2**: GNN encoder + causal transformer decoder for sequence design
3. **AbsciBind**: Modified AlphaFold-Multimer scoring protocol for design selection

3.5 RFAntibody: Diffusion-Based Backbone Generation

RFAntibody extends the RFdiffusion framework with antibody-specific fine-tuning and a multi-stage pipeline.

Pipeline Components:

1. **RFdiffusion (fine-tuned)**: Backbone generation with framework conditioning
2. **ProteinMPNN**: CDR sequence design with fixed framework sequences
3. **RoseTTAFold2 (fine-tuned)**: Structure prediction and filtering

A distinguishing feature is the MIT license, making it the only fully open-source option.

4 Performance Metrics and Validation

4.1 Hit Rate Comparison

Hit rate—the percentage of computational designs that yield experimental binders—represents the primary metric of design efficacy.

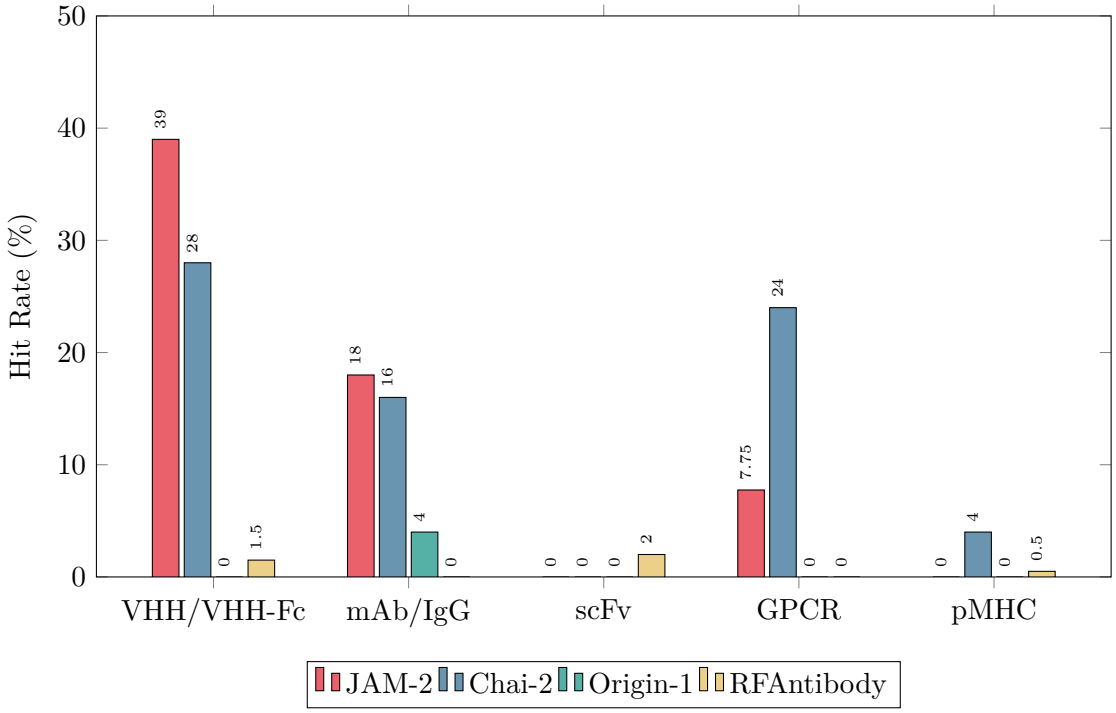

Figure 3: Hit rates across antibody formats and target categories. Values represent approximate averages where multiple targets were tested. Zero indicates format/category not tested or no hits reported.

4.2 How Platforms Define "Hits"

A critical challenge in comparing platform performance is understanding what constitutes a "hit." Each platform employs different experimental assays, thresholds, and validation criteria, complicating direct comparisons. Below we summarize the definitions used by each platform in their reported metrics, with full citations to primary sources.

Table 4: **Hit Rate Definitions by Platform.** All platforms use experimental binding confirmation as the primary hit criterion, but differ in assay methodology, threshold stringency, and orthogonal validation requirements.

Platform	Hit Definition	Validation Criteria
JAM-2	Percentage of designs yielding measurable binding by SPR/BLI (<10μM to pM)	Post-reformatting confirmation (VHH→VHH-Fc or mAb)
Chai-2	Number of designs showing target binding by SPR, BLI, or FACS with measurable affinity	Cell-surface binding confirmed; 93% of scFv hits retained mAb activity
Origin-1	SPR-identified binders confirmed by BLI and aSEC complexation	Three-tier cascade: SPR → BLI → complexation; 1–4 hits per 95 designs
Latent-X2	Target-level success: measurable binding in single-point BLI (>0.1 nm response)	5-point SPR/BLI with $R^2 > 0.95$ for kinetic characterization
RFAntibody	Yeast display enrichment followed by SPR confirmation	No aggregate hit rate reported; focus on structural validation via cryo-EM

4.2.1 Key Differences in Hit Definition

Terminology Distinction. Platforms differ in whether they report "success rates" (JAM-2), "hit rates" (Chai-2), or "target-level success" (Latent-X2). JAM-2 defines success as the percentage of designs that bound experimentally with measurable affinity, achieving 39% for VHH-Fc and 18% for mAb across 16 unseen targets Nabla Bio [2025]. Chai-2 reports per-target hit rates ranging from 4% to 87% (median 24%) depending on target class Chai Discovery [2025]. Latent-X2 achieved 50% target-level success (9 of 18 targets) with 4–24 designs tested per target per modality Latent Labs [2025]. Origin-1 does not report aggregate hit rates but achieved 100% target-level success (4 of 4 targets) with 1–4 hits identified per target from approximately 95 designs each AbSci [2026]. RFAntibody explicitly notes that "experimental success rates remain low" and does not report a specific hit rate, instead emphasizing structural validation Bennett et al. [2025].

Assay Methodology. JAM-2 and Chai-2 primarily use SPR and BLI for primary screening. Chai-2 additionally employs FACS (flow cytometry) for cell-surface targets such as GPCRs Chai Discovery [2025]. Origin-1 uses a tiered approach: SPR for initial identification, followed by BLI confirmation and aSEC complexation experiments AbSci [2026]. Latent-X2 uses single-point BLI screening with a response threshold of >0.1 nm classified as hits Latent Labs [2025]. RFAntibody relies on yeast surface display (YSD) enrichment followed by SPR confirmation Bennett et al. [2025].

Threshold Stringency. JAM-2 reports "average per-target success rates" as the percentage of designs that bound with measurable affinity (ranging from high-nanomolar to picomolar), representing binary confirmation rather than graded thresholds Nabla Bio [2025]. Chai-2 distinguishes between "binding hits" (confirmed target engagement by SPR/BLI/FACS) and "functional hits" (activity in cellular assays such as HEKBlue reporter assays) Chai Discovery [2025]. Latent-X2 applies a quantitative BLI response threshold (>0.1 nm reference-subtracted binding) for initial hit designation, followed by 5-point SPR or BLI with strict quality criteria ($R^2 > 0.95$) for kinetic characterization Latent Labs [2025].

Reformatting Considerations. JAM-2 designs are initially generated as VHHs but are reformatted to VHH-Fc fusions or full-length mAbs for experimental characterization. The reported 39% VHH-Fc success rate reflects post-reformatting binding confirmation Nabla Bio [2025]. Chai-2 demonstrated that 93% of scFv hits retained binding activity when reformatted to full-length IgGs, indicating robust transferability between formats Chai Discovery [2025]. This reformatting retention rate is a critical metric for therapeutic viability, as most clinical antibodies are full-length IgGs.

Orthogonal Validation. Origin-1 employs the most stringent validation pipeline: SPR identification → BLI confirmation → aSEC complexation experiments. This three-tier approach minimizes false positives but may exclude weak binders that could be improved through affinity maturation AbSci [2026]. RFAntibody emphasizes structural validation via cryo-EM, achieving atomic-level agreement between designed and experimentally determined structures (interface RMSD 0.9–1.45 Å) Bennett et al. [2025], but does not report systematic binding confirmation rates.

Benchmark Implication: When comparing hit rates across platforms, it is essential to consider that *a hit is not always a hit.* A platform with lower reported hit rates but stricter validation criteria may actually produce higher-quality leads. DADB-v1.0 addresses this by requiring standardized validation tiers (screening → confirmation → functional) in its scoring system.

4.3 Training Data Composition and Target Similarity

A critical question in evaluating de novo antibody design platforms is whether reported successes reflect genuine generalization to novel targets or merely memorization of training data patterns. This section examines how each platform addresses target antigen similarity to training set epitopes and antigens.

4.3.1 The Data Leakage Challenge

Why This Matters. Machine learning models for protein design are trained on databases of known protein structures (primarily the Protein Data Bank, PDB). If test targets or their epitopes were present in the training data, reported "success rates" may overestimate true generalization capability. This is particularly concerning for antibody design, where:

- Known antibody-antigen complexes in the PDB may include the exact test target or highly similar homologs
- Epitope-specific binding patterns could be memorized rather than learned de novo
- Sequence similarity at the antigen level may enable retrieval rather than design

Table 5: **Training Data Transparency by Platform.** Comparison of how each platform addresses target similarity to training data.

Platform	Training Set Description	Similarity Filtering	Test Target Claims
JAM-2	Not disclosed	Claims "no close homology" to any antibody-antigen complex at training cutoff	16 targets "completely unseen" Nabla Bio [2025]
Chai-2	Not disclosed	No explicit filtering described; one target (GPRC5D) confirmed absent from training	GPRC5D structure "not included" Chai Discovery [2025]
Origin-1	AbData: 3,242 antibody-antigen complexes + 29,835 PPIs	Temporal + sequence homology ($\leq$60% identity) + MMseqs2 filtering AbSci [2026]	4 targets with $\leq$60% identity to any training complex
Latent-X2	Not disclosed; 200 held-out benchmark set	Sequence novelty check: CDR edit distance $\geq$11 to SAbDab; MMseqs2 clustering at 80% identity Latent Labs [2025]	9/18 targets successful; novelty verified
RFAntibody	Fine-tuned RoseTTAFold2 on antibody structures; PDB cutoff Jan 13, 2023	Target similarity analysis shown in Extended Data Fig. 1 Bennett et al. [2025]	28 VHHs confirmed binders show no enhanced similarity to training set

4.3.2 Platform-Specific Approaches

Origin-1: The Most Transparent Approach. Origin-1 provides the most detailed description of training data curation and similarity filtering AbSci [2026]:

- **Temporal filtering**: Training data includes PDB entries released before 2024-01-01; validation (2024-01-01 to 2024-09-30); test (2024-09-01 to 2025-07-10)

- **Sequence homology filtering**: For antibody-antigen complexes, removed entries with antigen >40% sequence similarity to reference set; for protein-protein complexes, removed entries with either chain >40% similar
- **Redundancy removal**: Clustered antigens at 40% identity and Fv sequences at 100% identity; selected single representative per cluster-pair
- **Test target selection**: Explicitly selected targets with "limited sequence overlap (≤60% identity by MMseqs2) with any protein for which a protein-protein complex structure existed in the training data"

The final dataset comprised 3,242 antibody-antigen training entries, 63 validation, and 84 test entries.

RFAntibody: Post-Hoc Similarity Analysis. The RFAntibody study explicitly tested whether their designs showed enhanced similarity to the training dataset Bennett et al. [2025]:

> "Designed VHH sequences are distinct from the training dataset. Blastp was used to find hits against the SAbDab, and the similarity of the CDR loops in the top blast hit were reported for all VHHs experimentally tested in this study. Note also that the 28 VHHs confirmed to bind their targets by SPR do not show enhanced similarity to the training set."

Additionally, they confirmed that designed VHHs for the TcdB Frizzled epitope bound to a site "for which no antibody or VHH structure exists in the PDB," demonstrating true de novo design for a structurally uncharacterized epitope.

Latent-X2: Sequence Novelty Verification. Latent-X2 reports explicit sequence novelty checks Latent Labs [2025]:

- CDR edit distance of at least 11 residues to closest match in SAbDab (most >20)
- BLASTP against non-redundant database: all designs at least 16 residues distant from nearest neighbor
- MMseqs2 clustering at 80% identity to avoid testing closely related variants

JAM-2 and Chai-2: Limited Disclosure. Both JAM-2 and Chai-2 make claims about target novelty but provide less methodological detail:

- **JAM-2**: States that 16 test targets "had no close homology to any antibody–antigen complex available at the model's training cutoff" Nabla Bio [2025]. No specific similarity thresholds or filtering methodology is disclosed.
- **Chai-2**: Notes that GPRC5D structures "were not included in Chai-2's training dataset" Chai Discovery [2025]. No systematic analysis of target similarity for other test cases is provided.

4.3.3 Implications for Benchmarking

The varying levels of transparency in training data composition raise important considerations for comparing platform performance:

1. **Stringency differences**: Platforms with stricter similarity filtering (Origin-1 at ≤60%, Latent-X2 at edit distance ≥11) likely face harder test conditions than those without explicit filtering.

2. **Target selection bias**: Platforms may inadvertently or deliberately select test targets similar to training examples. Without pre-registration of test targets or explicit similarity thresholds, reported success rates may not reflect true generalization.

3. **Epitope vs. antigen similarity**: Even if the overall antigen is dissimilar, specific epitopes may resemble binding sites present in training data. Only Origin-1 explicitly addresses "zero-prior" epitopes lacking structural precedent.

4. **Temporal validation**: Origin-1's use of temporal cutoffs (training on structures released before 2024) provides the strongest guarantee against data leakage, as test targets were genuinely unknown during training.

Critical Assessment: When evaluating platform claims, ask: (1) Were test targets explicitly filtered for similarity to training data? (2) What similarity threshold was used? (3) Was filtering applied *a priori* or justified *post hoc*? Origin-1's temporal + sequence homology approach represents the current gold standard for preventing data leakage in antibody design benchmarks.

4.4 Structural Validation

Cryo-electron microscopy provides the gold standard for validating computational predictions. Table 6 summarizes structural validation results.

Table 6: Cryo-EM structural validation across platforms.

Platform	Complex	Resolution	Global RMSD	Interface RMSD
Chai-2	S1433B	2.9 Å	0.41 Å	0.54 Å
	CSF1	3.3 Å	—	1.9 Å
	EFNA5	3.9 Å	1.7 Å	—
	IL20	3.3 Å	—	—
	EPCR	3.5 Å	—	—
Origin-1	COL6A3	3.0 Å	2.56 Å	0.96 Å
	AZGP1	3.1 Å	1.79 Å	1.35 Å
RFAntibody	Influenza HA	3.0 Å	1.45 Å	—
	TcdB-scFv6	3.6 Å	0.9 Å	—
	TcdB-VHH	4.6 Å	—	—
	SARS-CoV-2 RBD*	3.9 Å	*Failed*	*Failed*

*Design failure: correct epitope but incorrect binding mode.

4.5 Developability Assessment

Drug-like properties are essential for therapeutic development.

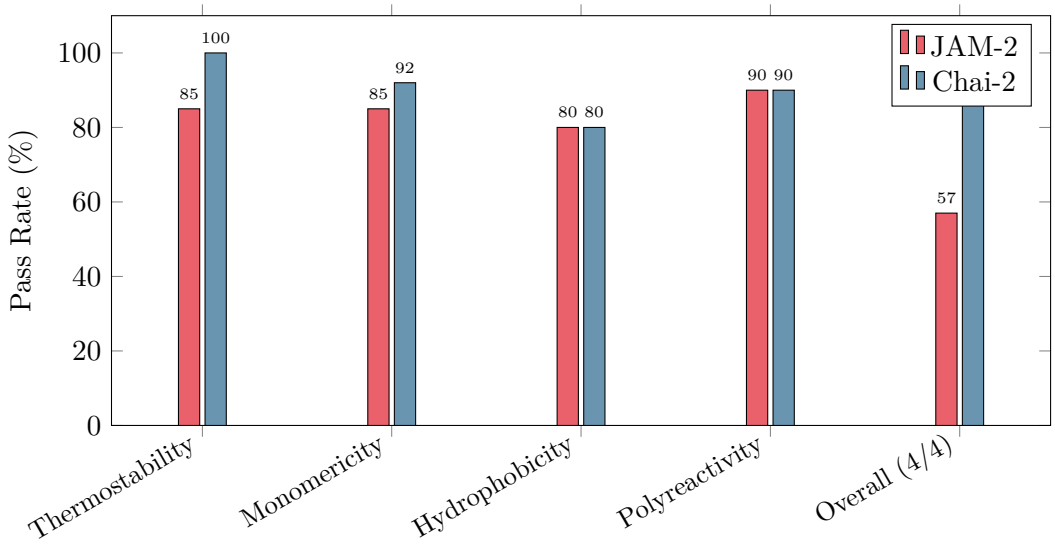

Figure 4: Developability pass rates for JAM-2 and Chai-2. Origin-1 and RFAntibody reported limited developability data.

5 Oncology Applications

Oncology represents the primary application domain for de novo antibody design, encompassing traditional targets and emerging modalities.

5.1 Traditional Oncology Targets

5.1.1 Immune Checkpoints

Checkpoint inhibitors transformed oncology. Keytruda generates over $30 billion annually.

PD-L1 (CD274): JAM-2 demonstrated low-nanomolar VHH binders against PD-L1. With 1,996 disease associations in Open Targets and 13 approved drugs, PD-L1 is maximally validated.

5.1.2 Receptor Tyrosine Kinases

VEGFR2 (KDR): JAM-2 achieved 3.3 nM binding to VEGFR2. With 1,397 disease associations and ramucirumab approved, this represents a validated but competitive space.

TrkA (NTRK1): JAM-2's sub-100 pM affinity (avid format) against TrkA is remarkable. With 1,244 disease associations and tumor-agnostic approval for larotrectinib, TrkA exemplifies precision medicine.

FGFR1: Chai-2 tested FGFR1 as part of their conventional target panel. With 1,653 disease associations and futibatinib approved, FGFR represents an active target class.

5.2 GPCR Targets in Oncology

G protein-coupled receptors represent historically "undruggable" targets with significant oncology relevance.

Table 7: GPCR targeting across platforms with oncology relevance.

Target	Platform	Hit Rate	Best K_D	Oncology Relevance
GPRC5D	Chai-2	48%	189 nM	Multiple myeloma target (talquetamab approved)
CXCR4	JAM-2	11.7%	1.4 nM	Metastasis, tumor microenvironment
CXCR4	Chai-2	11%	164 nM*	First functional GPCR agonist antibody
CXCR7	JAM-2	3.8%	23 nM	Tumor angiogenesis
CCR8	Chai-2	50%	453 pM	Tumor-infiltrating Tregs; 0 approved drugs

*EC50 for partial agonist activity.

5.2.1 GPRC5D: Multiple Myeloma Target

Chai-2 achieved remarkable success against GPRC5D, an orphan GPCR that was not included in their training data:

- 48% hit rate (35/73 VHH-Fc designs)
- 189 nM monovalent affinity (nanodisc SPR)
- Cross-reactivity to cynomolgus (334 nM)—important for preclinical development

5.2.2 CXCR4: Metastasis and Microenvironment

Both JAM-2 and Chai-2 successfully targeted CXCR4:

- JAM-2: 11.7% hit rate, 1.4 nM best affinity
- Chai-2: Designed partial agonist (EC50 = 164 nM)
- First computationally designed GPCR functional antibodies

5.2.3 CCR8: Tumor Treg Depletion

CCR8 is expressed selectively on tumor-infiltrating regulatory T cells. Chai-2's 50% hit rate and 453 pM best affinity represents a potential first-in-class opportunity—zero approved therapeutics exist.

5.3 Tumor-Specific Neoepitopes

Targeting mutant peptide-MHC (pMHC) complexes enables selective killing of cancer cells.

5.3.1 KRAS G12V Neoepitope (Chai-2)

KRAS mutations are among the most common oncogenic drivers:

- Target: HLA-A*03:01 presenting KRAS G12V peptide (VVVGAVGVGK)
- Hit rate: 4% (2/50 designs)
- Best affinity: 1.5 nM (bivalent SPR)
- Specificity: No binding to KRAS WT or G12D (single residue discrimination)
- Allele specificity: No binding to HLA-A*11:01-KRAS G12V

This level of discrimination—valine versus aspartic acid versus glycine at position 12—is precisely what neoepitope therapeutics require.

5.3.2 TP53 R175H Neoepitope

Chai-2 targeted the most commonly mutated tumor suppressor. With 3,277 disease associations, TP53 represents the ultimate "undruggable" oncogene.

5.3.3 PHOX2B Neoepitope (RFAntibody)

RFAntibody targeted the neuroblastoma-associated PHOX2B peptide:

- Target: HLA-C*07:02 presenting PHOX2B peptide
- Affinity: 400 nM (SPR), 1 μM (ITC)
- Peptide specificity confirmed (no R6A mutant binding)
- CAR-T construct tested but showed no detectable killing (likely due to modest affinity)

5.4 Zero-Prior Oncology Targets

Origin-1 demonstrated success against targets lacking structural precedent:

FOLR1 (Folate Receptor Alpha): With mirvetuximab soravtansine approved (2022), FOLR1 represents a validated ADC target. Origin-1's zero-prior approach could identify novel epitopes.

IL36RN: Origin-1 achieved functional antagonist activity (EC50 = 104 nM) after 68-fold affinity maturation. The target modulates inflammatory signaling in the tumor microenvironment.

6 Comprehensive Oncology Target Table

Table 8: Oncology Targets: Platform Performance and Clinical Context. Disease associations from Open Targets Platform.

Target	Platform	Affinity	Diseases	Clinical Context
Immune Checkpoints				
PD-L1	JAM-2	Low nM	1,996	13 approved drugs incl. durvalumab, atezolizumab
Receptor Tyrosine Kinases				
VEGFR2	JAM-2	3.3 nM	1,397	Ramucirumab approved; 69 drugs in development
TrkA	JAM-2	<100 pM	1,244	Larotrectinib tumor-agnostic approval
FGFR1	Chai-2	—	1,653	Futibatinib approved for cholangiocarcinoma
Oncology GPCRs				
CXCR4	JAM-2	1.4 nM	1,500	Metastasis driver; plerixafor approved

Continued on next page

Table 8 – continued

Target	Platform	Affinity	Diseases	Clinical Context
CXCR4	Chai-2	164 nM*	1,500	First-ever computational GPCR agonist antibody
CCR8	Chai-2	453 pM	259	Zero approved drugs; tumor Treg marker
GPRC5D	Chai-2	189 nM	91	Talquetamab approved Aug 2023 (myeloma)
CXCR7	JAM-2	23 nM	—	Tumor angiogenesis and metastasis
Neoepitopes / pMHC				
KRAS G12V	Chai-2	1.5 nM	1,801	Single-residue specificity; no WT/G12D binding
TP53 R175H	Chai-2	—	3,277	Most mutated gene in cancer
PHOX2B	RFAntibody	400 nM	—	Neuroblastoma driver mutation
Tumor Antigens				
CD38	JAM-2	—	990	Daratumumab ($9B+); isatuximab approved
Zero-Prior Targets				
FOLR1	Origin-1	—	402	Mirvetuximab soravtansine approved 2022
IL36RN	Origin-1	89 nM	245	Tumor microenvironment modulation

*EC50 for agonist activity.

7 Comparative Analysis

7.1 Strengths and Limitations

Table 9: Comparative strengths and limitations of each platform.

Platform	Key Strengths	Limitations
JAM-2	• Highest hit rates (39% VHH-Fc) • Comprehensive developability (57% pass all 4) • Systematic epitope coverage • GPCR orthosteric targeting • Large-scale validation (923+ designs)	• Architecture not disclosed • No pMHC targeting demonstrated • Proprietary/not available
Chai-2	• Atomic accuracy (<1 Å HCDR3) • First functional GPCR agonists • Single-residue pMHC discrimination • 100% thermostability • Full-length IgG direct design	• Not publicly available • Variable hit rates (4–87%) target-dependent • Training data undisclosed
Origin-1	• Zero-prior epitope capability • All-atom diffusion generation • Functional antagonist design • AbsciBind scoring innovation	• Low per-design hit rates (1–4 hits per 95 designs) • Initial μM affinities • Requires optimization
RFAntibody	• Open-source (MIT license) • Atomic precision validated • Framework conditioning flexibility • Nature publication	• Low experimental success rates (no aggregate rate reported) • Requires yeast display enrichment • Design failures occur

7.2 Timeline and Resource Requirements

Table 10: Estimated resource requirements for a typical antibody design campaign.

Metric	JAM-2	Chai-2	Origin-1	RFAntibody
Computational time	2–3 days	∼1 week	∼1 week	∼1 week
Designs per target	45–100	50–100	<100	9,000+
Wet-lab validation	<4 weeks	2 weeks	6 weeks	Months
Affinity maturation	Rarely needed	Rarely needed	Often needed	Usually needed
Open source	No	No	Partial	Yes (MIT)

8 Discussion

8.1 Convergent Themes

Despite distinct architectural approaches, several convergent themes emerge:

1. **Diffusion as Foundation**: Three of four platforms employ diffusion-based generation.
2. **All-Atom Modeling**: Moving beyond backbone-only to all-atom resolution appears critical.
3. **Epitope Conditioning**: Explicit epitope specification enables targeted interface design.
4. **Structure-Based Scoring**: Integration of structure prediction significantly improves success rates.

8.2 Implications for Drug Discovery

The demonstrated capabilities have significant implications:

1. **Timeline Compression**: Traditional 12–24 months reduced to 4–8 week cycles.
2. **Target Expansion**: GPCRs and pMHCs now accessible.
3. **Epitope Control**: Systematic epitope coverage from the outset.
4. **Reduced Optimization**: High initial affinities (pM–nM) may eliminate optimization cycles.

8.3 Limitations

Critical gaps remain:

- **In Vivo Validation**: None demonstrated in vivo efficacy or PK.
- **Immunogenicity**: Clinical immunogenicity untested.
- **Manufacturing**: Large-scale production unvalidated.
- **Reproducibility**: Most platforms not publicly available.

9 Conclusions

The emergence of JAM-2, Chai-2, Origin-1, and RFAntibody marks a transformative moment in antibody therapeutics:

1. De novo antibody design has progressed from proof-of-concept to practical utility.
2. Hit rates of 15–40% are achievable, representing 10–100× improvements.
3. Drug-like affinities (single-digit nanomolar) achievable without optimization.
4. Challenging targets—GPCRs, pMHC neoepitopes—are now accessible.
5. Oncology applications span checkpoints, RTKs, GPCRs, and neoepitopes.

The field is transitioning from "can we design antibodies computationally?" to "how do we optimally deploy these capabilities?"

References

Bennett, N.R., Watson, J.L., Ragotte, R.J., et al. (2025). Atomically accurate de novo design of antibodies with RFdiffusion. *Nature*, published online.

Bradbury, A.R., Sidhu, S., Dübel, S., McCafferty, J. (2011). Beyond natural antibodies: the power of in vitro display technologies. *Nature Biotechnology*, 29(3):245–254.

Chai Discovery Team (2025). Drug-like antibody design against challenging targets with atomic precision. *bioRxiv*, doi:10.1101/2025.11.29.691346.

Jumper, J., Evans, R., Pritzel, A., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589.

Kaplon, H., Crescioli, S., Chenoweth, A., et al. (2024). Antibodies to watch in 2024. *mAbs*, 16(1):2297450.

Nabla Bio (2025). JAM-2: Fully computational design of drug-like antibodies with high success rates. Technical Report.

AbSci Corporation (2026). Origin-1: a generative AI platform for de novo antibody design against novel epitopes. *bioRxiv*, doi:10.64898/2026.01.14.699389.

Latent Labs Technologies (2025). Latent-X2: Drug-like antibodies with low immunogenicity in human panels designed with Latent-X2. Technical Report, https://platform.latentlabs.com.

Watson, J.L., Juergens, D., Bennett, N.R., et al. (2023). De novo design of protein structure and function with RFdiffusion. *Nature*, 620(7976):1089–1100.