

GenovaX

An Evidence-Gated Workflow for AI-Assisted Molecular Design

ABSTRACT

This paper defines GenovaX as a proposed computational workflow for generating, predicting, filtering, ranking, and handing off molecular candidates for experimental review. The core requirement is evidence gating: a computational score remains a model output, not evidence of binding, biological activity, safety, efficacy, manufacturability, or therapeutic value. The architecture begins with a defined context of use, preserves target and dataset provenance, separates generation from prediction, applies developability and liability filters, records uncertainty and model disagreement, requires expert review, and produces an experimental handoff package with predeclared acceptance criteria. The existing Planckchron lysozyme benchmark demonstration is computational only and is not a therapeutic program. No candidate described by Planckchron has been represented here as experimentally validated, clinically tested, approved, safe, or effective.

PLANCKCHRON RESEARCH

Self-published technical paper / not peer-reviewed / open access

planckchron.com/publications/genovax-evidence-gated-molecular-design

TECHNICAL NOTE

Executive brief

This paper defines GenovaX as a proposed computational workflow for generating, predicting, filtering, ranking, and handing off molecular candidates for experimental review. The core requirement is evidence gating: a computational score remains a model output, not evidence of binding, biological activity, safety, efficacy, manufacturability, or therapeutic value. The architecture begins with a defined context of use, preserves target and dataset provenance, separates generation from prediction, applies developability and liability filters, records uncertainty and model disagreement, requires expert review, and produces an experimental handoff package with predeclared acceptance criteria. The existing Planckchron lysozyme benchmark demonstration is computational only and is not a therapeutic program. No candidate described by Planckchron has been represented here as experimentally validated, clinically tested, approved, safe, or effective.

EVIDENCE BOUNDARY

Computational research workflow only. No diagnostic, therapeutic, safety, efficacy, regulatory, patient-care, or experimentally validated molecular claim is made.

Four operating conclusions

- Structure and interaction predictions generate hypotheses; they do not establish biological truth.
- The context of use determines the evidence required from an AI model.
- Model disagreement and negative results belong in the ranking and handoff record.
- The workflow is complete only when a reviewer can reproduce why a candidate advanced or stopped.

STATUS	Research workflow
PUBLISHER	Planckchron Inc.
CITATION	<i>Planckchron Research (2026). GenovaX: An Evidence-Gated Workflow for AI-Assisted Molecular Design. Version 1.0.</i>

TECHNICAL NOTE

Reference architecture

The architecture below is a proposed control and evidence flow. It is a design artifact, not a representation of deployed capability.



TECHNICAL NOTE

1. Context and objective

Modern structure and interaction models can reduce the cost of forming hypotheses, but their outputs are easy to overinterpret. A high confidence score may reflect confidence within a model's learned distribution, not proof of binding or useful biological effect. A visually plausible complex may still fail expression, stability, specificity, kinetics, function, manufacturability, or safety requirements.

GenovaX is proposed as a research workflow that makes those boundaries explicit. It combines computational generation and prediction with provenance, uncertainty, filtering, expert review, and experimental planning. It is not a medical product and this paper does not establish a therapeutic candidate.

TECHNICAL NOTE

2. Context-of-use contract

The workflow begins by defining how the model output will be used. Examples include prioritizing a small set for an assay, exploring sequence diversity, generating a structural hypothesis, or identifying liabilities for redesign. The same model output requires different credibility when used for exploratory ranking than when used to support a regulated decision.

The contract records the biological question, target identity and version, intended and prohibited interpretations, decision owner, input data rights and provenance, model versions, acceptance criteria, uncertainty plan, experimental follow-up, and stop conditions. If the context changes, the credibility assessment must be repeated.

TECHNICAL NOTE

3. Reference workflow

The target and data layer verifies identifiers, sequences, structures, biological assumptions, and licensing or consent constraints. Candidate generation may use retrieval, mutation, language models, diffusion methods, or other design tools, but generation is separated from scoring to reduce circular evidence. Structure and interaction prediction produces ensembles where feasible rather than one preferred image.

A filtering layer evaluates interpretable liabilities such as problematic motifs, unusual cysteine patterns, charge and hydrophobicity extremes, aggregation risk indicators, and other domain-specific rules. A ranking layer combines model evidence without collapsing it into false certainty. Expert review examines biological plausibility, diversity, novelty, failure modes, and assay value. The output is an experimental handoff package, not a therapeutic claim.

TECHNICAL NOTE

4. Evidence ladder

The workflow uses an explicit evidence ladder. A generated candidate is a design artifact. A predicted structure or interaction is a computational observation. Agreement across models or seeds strengthens confidence in reproducibility of the prediction, not in the underlying biology. A controlled assay can establish a result for that assay under its conditions. Replication, orthogonal methods, functional evidence, and later regulated development add different evidence that cannot be inferred from the computational stage.

Public language should identify the rung. Terms such as designed, predicted, ranked, expressed, bound, functional, safe, and effective must not be treated as synonyms.

TECHNICAL NOTE

5. Evaluation and uncertainty

Evaluation should include holdout and out-of-distribution tasks, target leakage checks, calibration, ranking quality, structural accuracy where reference data exist, model disagreement, sensitivity to input changes, and failure analysis by target class. Benchmark selection and data overlap must be disclosed. A single aggregate result can conceal where the workflow fails.

For candidate ranking, uncertainty can include model confidence, ensemble dispersion, disagreement between structure and property methods, distance from training-like examples, rule-based liabilities, and missing biological knowledge. The ranking should preserve these components so experts can disagree with the weighting.

TECHNICAL NOTE

6. Experimental handoff

A strong handoff includes candidate identity and version, design provenance, predicted structures and confidence, ranking rationale, known liabilities, negative computational results, proposed controls, assay sequence, acceptance and stop criteria, sample and data handling requirements, and a plan for returning results to the model record.

The wet-lab plan should be capable of disproving the computational premise. Selecting only assays likely to confirm the model weakens the evidence. Negative results should remain linked to the design lineage so failed patterns are not repeatedly rediscovered.

TECHNICAL NOTE

7. Benchmark demonstration boundary

Planckchron has separately described a methodology demonstration using hen egg-white lysozyme, a standard benchmark target. The reported output is a computational shortlist produced by an AI-assisted workflow. The molecules were not represented as expressed or assayed, and lysozyme is not represented as a Planckchron therapeutic target.

That demonstration is useful for testing workflow mechanics, provenance, ranking, and disclosure. It cannot establish biological activity, translational value, or performance on a new therapeutic target. Any follow-on report should preserve that distinction and publish experimental methods and negative results if testing occurs.

TECHNICAL NOTE

8. Governance, limitations, and next gates

Access control, data governance, biosafety review, dual-use assessment, reproducible environments, model and dataset documentation, and qualified scientific review are required in proportion to the work. Any regulated use would require a separate context-specific quality and evidence program.

The next evidence gates are reproducible reruns of the computational pipeline; stronger benchmark and data-overlap analysis; independent review of the filtering and ranking logic; and, only with appropriate partners and controls, predefined experimental testing. This paper does not represent completion of those gates.

TECHNICAL NOTE

References and source notes

References are provided for the external standards, public guidance, and primary research that informed this paper. A citation does not imply endorsement, partnership, certification, or validation of Planckchron capability.

01 Nature, 2024**Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3**

<https://www.nature.com/articles/s41586-024-07487-w>

External peer-reviewed research; prediction does not substitute for Planckchron experimental evidence.

02 Boltz-1, 2025**Boltz-1: Democratizing Biomolecular Interaction Modeling**

<https://pmc.ncbi.nlm.nih.gov/articles/PMC11601547/>

External primary research and open-source method reference.

03 FDA Draft Guidance, 2025**Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products**

<https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>

External draft regulatory guidance; no Planckchron approval or regulatory status is implied.

04 FDA and EMA, 2026**Guiding Principles of Good AI Practice in Drug Development**

<https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development>

External public principles; no Planckchron endorsement or compliance claim is implied.

PUBLICATION DISCLOSURE

These self-published technical papers are not peer-reviewed, audited, or independent validation. They describe research methods, proposed architectures, and evaluation plans. Product capability exists only where a dated evidence record expressly reports a result. Computational research workflow only. No diagnostic, therapeutic, safety, efficacy, regulatory, patient-care, or experimentally validated molecular claim is made.