NEUROLAB

NeuroLab: An Auditable Neurosymbolic Platform for Central Nervous System Drug Discovery

Public scientific whitepaper — review draft

From molecular evidence to auditable multiscale hypotheses

20 July 2026

Version 0.11

Research and engineering whitepaper; not peer reviewed

Auditable in-silico triage, hypothesis generation, and evidence-led mechanism exploration

Version: 0.11
Snapshot date: 20 July 2026
Scientific status: Research and engineering whitepaper; not peer reviewed
Intended use: Auditable in-silico triage, hypothesis generation, and evidence-led mechanism exploration
Prohibited use: Clinical diagnosis, prescribing, dosing, patient-specific treatment, safety certification, regulatory decisions, or autonomous experimental decisions

> **Status statement.** NeuroLab is a working, routed research system, not a single end-to-end model and not a validated digital twin. It combines a transductive predictor for compounds represented in a governed knowledge graph, a separate chemistry-to-receptor route for novel structures, and a deterministic symbolic composer. The routes have different inputs, denominators, and evidence states; they must not be assigned one overall accuracy.

> **Public-disclosure boundary.** This paper discloses the scientific problem, high-level system decomposition, evidence contracts, selected source-bound evaluation and engineering results, limitations, and validation roadmap. It intentionally withholds source code, private data mappings, exact feature and retrieval recipes, model parameters, decision thresholds, deployment topology, repository structure, private artifact identifiers, internal file paths, raw benchmark traces and payloads, security-control configuration, third-party paper figures, atlas renderings or geometry, product screenshots, licensed dataset extracts, and other material without a documented public-use basis. Controlled technical records can be made available selectively for diligence or independent evaluation under an appropriate written agreement.

Abstract

Central nervous system research requires reasoning across heterogeneous scales: chemical structure, molecular targets, functional action, cell and layer context, circuit anatomy, regional and network effects, cognition, and clinical observations. NeuroLab is being developed as an auditable neurosymbolic platform for this problem. It unifies governed evidence while keeping three computational functions distinct: a heterogeneous graph predictor for compounds represented in the knowledge graph; a chemistry-to-receptor retriever for novel structures; and a symbolic composer that develops receptor hypotheses into explicit multiscale mechanism paths. A provenance layer, a typed inference API, and a human-facing application preserve the route, evidence state, applicability domain, alternatives, and reasons to abstain.

The platform is designed to turn fragmented scientific evidence into inspectable research hypotheses. For discovery teams, it can support compound and mechanism triage, evidence-gap analysis, comparison of mechanistic alternatives, and design of discriminating experiments. The implemented application accepts a governed compound name or a molecular structure, routes the query to the appropriate computational path, streams progress for the known-compound workflow, and returns both human-readable and machine-readable reports. Complete reports or selected JSON sections can be exported for downstream review. The current application also contains an interactive 3D Brain Explorer. After a routed report is returned, the Explorer can display available 392-parcel report maps, reference anatomy, tracts, and affected-system annotations while keeping predicted, literature-derived, measured, and legacy channels visibly separate. This existing interface can make evidence and uncertainty easier to inspect for community members and other interested non-specialists. It is not endorsement of a substance, product, practice, or self-experiment, and the system does not recommend or optimize personal use.

Current evidence supports a working routed system and encouraging component-level development signals, not a single end-to-end accuracy claim. Several known-compound categorical heads outperform simple development controls; a reproducible scaffold-disjoint chemistry-neighbour receptor comparator exceeds frequency and random baselines; and deterministic contracts verify selected route, provenance, artifact, and abstention behavior. A historical, workload-specific serving benchmark also demonstrates practical CPU and GPU execution, with a warmed three-seed streaming request measuring 4.91 s end to end on CPU and 1.28 s on an RTX 5070 Ti GPU (three requests per profile; the server-side spans alone were 3.84 s and 1.07 s). These are route-specific development and engineering observations rather than clinical validation, calibrated per-prediction probabilities, or a current service-level objective.

NeuroLab therefore occupies a useful and defensible position today: an auditable triage and hypothesis-generation system that can show how an inference was assembled, where support is conditional, and which transition is missing. The next milestone is a frozen chemistry- and evidence-disjoint evaluation of the complete structure-to-receptor-to-mechanism-to-endpoint chain with calibration and selective-risk analysis. Quantitative population dynamics and an individualized brain digital twin remain longer-term research directions.

Keywords: CNS drug discovery; neurosymbolic AI; knowledge graph; graph neural network; mechanism of action; circuit reasoning; provenance; uncertainty; abstention; digital twin

Executive scientific summary

The opportunity

CNS discovery is slowed by evidence fragmentation. Chemical assays, target pharmacology, anatomical atlases, circuit studies, imaging, cognition, and clinical observations are usually searched and interpreted in separate tools. NeuroLab provides a governed layer in which these sources can be connected without hiding differences in route, context, or evidential strength. The aim is not merely to return another score; it is to help a researcher move from a question to an auditable mechanism hypothesis, its alternatives, and the next evidence needed to test it.

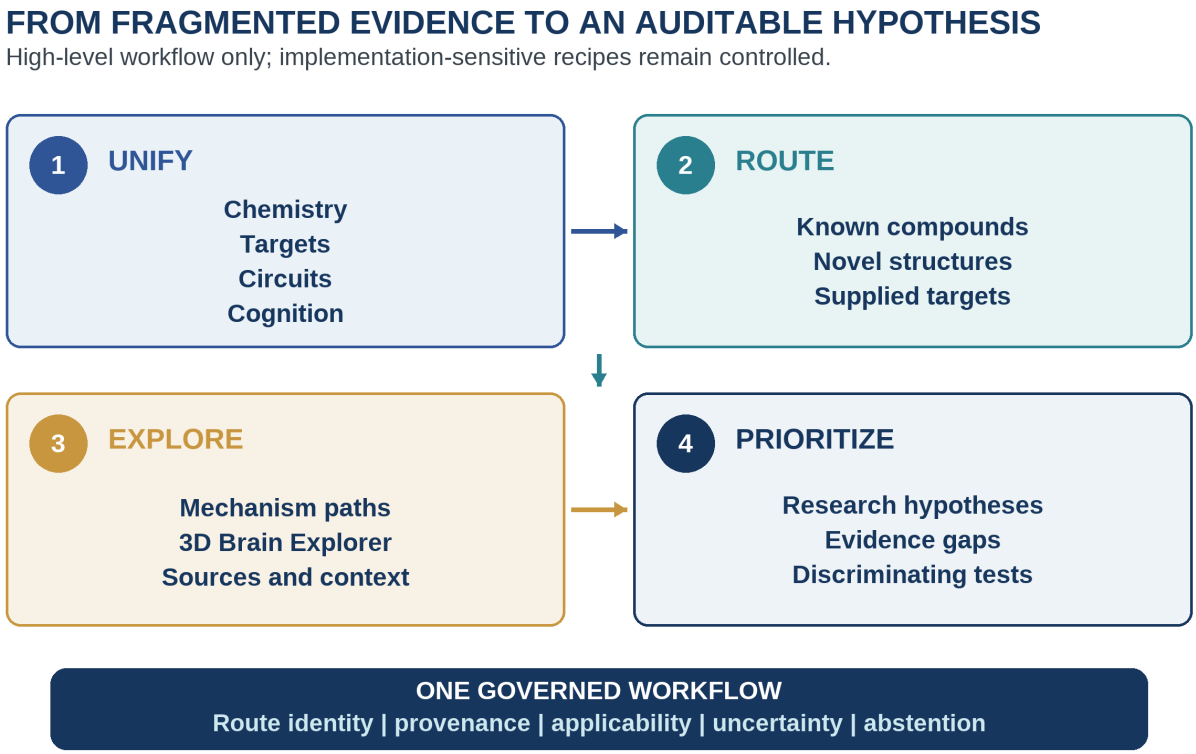

Figure 1: NeuroLab research value pathway. The public view describes the workflow and scientific controls without disclosing implementation-sensitive recipes.

What NeuroLab delivers

- **One governed evidence substrate** spanning compounds, targets, biological entities, anatomy, cognition, clinical categories, citations, and reference overlays.
- **Three explicit computational routes** for graph-resident compounds, novel structures, and supplied receptor hypotheses, avoiding a misleading one-model-fits-all abstraction.
- **Mechanism-oriented outputs** that expose supporting relations, alternative hypotheses, applicability limits, and missing transitions rather than presenting an opaque score alone.
- **Research workflow support** for compound and mechanism triage, evidence-gap discovery, comparative hypothesis analysis, and planning of discriminating experiments.
- **Reproducible engineering controls** through deterministic manifests, artifact binding, route contracts, and tests of selected software behavior.
- **An implemented inference API and query router** for governed compound-name, molecular-structure, and explicitly conditional receptor-profile inputs, with typed JSON, streamed delivery, health/version reporting, and generated OpenAPI descriptions.

- **Dual report surfaces** that provide a readable interpretation while preserving the complete structured payload, evidence and reliability context, exact JSON paths, and downloadable JSON for audit or integration.
- **An existing interactive 3D Brain Explorer** that turns available report channels into an inspectable spatial view while preserving channel identity, reference-only labels, uncertainty, and abstention.

Why the architecture is differentiated

The system's central design choice is separation with composition. Known compounds, novel structures, and caller-supplied receptor profiles do not silently share one estimator or one accuracy number. Neural components are used where learned ranking is appropriate; symbolic composition is used where paths, conditions, provenance, and alternatives must remain inspectable. Reference layers can expand explanation and research planning without receiving unearned credit for predictive performance. This makes the platform extensible while preserving a defensible boundary between prediction, retrieval, evidence, and narrative reporting.

Evidence boundary

The present release supports route-specific development and engineering claims. It does not yet support a scalar overall accuracy, calibrated probabilities across routes, independent full-chain validation for novel compounds, quantitative dose-to-brain dynamics, or an individualized digital twin. These boundaries are retained in the body and appendices, while the near-term program concentrates on the experiments that can unlock stronger claims.

Near-term scientific focus

The immediate priority is to bind each public metric to a controlled validation record, complete a genuinely blind chemistry- and evidence-disjoint benchmark of the served routes, and measure calibration and selective risk. This converts the platform's current strengths—integrated evidence, explicit routing, mechanistic composition, provenance, and abstention—into a reproducible external evaluation. Retraining, promotion, and deployment remain separate human decisions.

1. Scope, intended use, and non-claims

1.1 Scientific objective

The project objective is to predict and explain how a compound may perturb the central nervous system across scales:

compound → receptor or channel → cell and compartment → population → pathway and circuit → nucleus and region → distributed network → cognition → clinical observation

The system is designed to help experts prioritize hypotheses, inspect evidence, find missing transitions, compare mechanistic alternatives, and plan discriminating experiments. Its value proposition is auditability: outputs should expose the inference route, supporting and contradicting evidence, applicability limits, and reasons to abstain.

1.2 Intended users and decisions

The intended professional users are computational neuroscientists, pharmacologists, medicinal chemists, translational scientists, data curators, and research engineers. Appropriate research decisions include which mechanism to investigate, which evidence gap blocks an inference, which comparator or assay would discriminate alternatives, and whether the system should abstain. A ranked output is a hypothesis for review, not a fact about a compound or a patient.

1.3 Existing 3D Brain Explorer and community use case

The present application includes a 3D Brain Explorer. In the implemented flow, the Explorer becomes available after a prediction report is returned. It provides an interactive glass-brain view across a 392-parcel cortical and subcortical vocabulary, optional reference anatomy and tract layers, region selection, and map, network, parcel, and affected-system summaries. The feature is implemented in the application; external public availability is staged separately and is not claimed here as already launched.

For community members and other interested non-specialists, including people interested in brain science, neuroscience, neuropharmacology, nootropics, or biohacking, the same interface can support evidence and mechanism exploration. Its purpose in that context is to make provenance, channel identity, alternatives, and uncertainty inspectable. Interest in these topics does not imply that a practice is effective, appropriate, or safe.

Community goal	Permitted support	Prohibited interpretation or action
Learn how a compound is discussed across biological scales	Explain relevant concepts, display cited evidence, and distinguish measured observations from inferred mechanisms	Present the mechanism as established merely because a plausible path can be composed
Inspect a returned brain-system hypothesis in 3D	Switch among separately labelled report channels, reference anatomy, tracts, regions, and summaries; expose missing or abstained outputs	Interpret colour, parcel rank, tract overlap, or particle motion as measured activation, BOLD response, neuronal firing, causal direction, efficacy, or safety
Examine a claim encountered online or in a community	Trace supporting, contradictory, and missing evidence; show source context and evidence state	Endorse the claim, influencer, commercial product, supplement, protocol, or vendor
Compare mechanistic hypotheses	Show alternative routes, uncertainty, applicability limits, and reasons to abstain	Rank substances for personal benefit, enhancement, productivity, mood, cognition, or longevity
Contribute to scientific discussion	Identify questions, missing evidence, and candidate sources for expert review	Treat anecdotes, votes, testimonials, or community popularity as validated evidence

The Brain Explorer and surrounding report experience must not recommend substances, products, stacks, doses, schedules, routes of administration, cycling strategies, combinations, or self-experiments. They must not provide individualized benefit-risk judgments, infer safety from missing adverse evidence, or advise a person to start, stop, replace, or combine a treatment or supplement. Commercial placement, affiliate incentives, and popularity signals must not influence scientific ranking or evidence presentation. Questions involving symptoms, adverse reactions, contraindications, pregnancy, prescribed medicines, substance withdrawal, or urgent safety concerns must be redirected to qualified clinical or emergency resources rather than answered as a NeuroLab prediction.

1.4 Explicit non-claims

This whitepaper does not claim that:

- NeuroLab is a validated whole-brain simulator or digital twin;
- all knowledge-graph edges are learned by the predictor;
- the novel-compound route has established end-to-end accuracy;
- a ranking score is a calibrated biological probability;
- seed agreement is a posterior uncertainty distribution;
- receptor-density maps predict activation, firing, BOLD, EEG, behavior, efficacy, or safety;
- an administered dose is converted into unbound brain concentration or target occupancy;
- cross-species cell-class homology transfers effect sign, dose response, dynamics, or endpoints;
- a literature-supported pathway is activated for a query merely because it exists in the graph;
- a software test pass constitutes biological validation; or
- a generated narrative adds predictive evidence beyond its structured inputs.

1.5 Evidence-state vocabulary

The project uses the following states throughout this document.

State	Meaning	Permitted wording
Executable	Present in a governed runtime and bound to versioned assets	“The current system executes . . . ”
Software verified	Deterministic schema, integrity, or contract checks pass	“The artifact or behavior verifies . . . ”
Development diagnostic	Measured on development data or a reused cohort	“Development performance was . . . ”
Oracle or conditional diagnostic	Downstream component receives a curated upstream truth or caller-supplied assumption	“Conditional on the supplied receptor or class profile . . . ”

State	Meaning	Permitted wording
Reference-only or shadow	Curated and queryable, but not used to train, score, or route evaluated predictions	“The reference substrate contains . . . ”
Scientifically validated	Frozen evaluation supports a predeclared claim against relevant controls	Used only where the stated claim and evaluation match
Planned	Design target without current implementation evidence	“The roadmap requires . . . ”

This taxonomy prevents a common category error: software can be correct with respect to its contract while the biological claim remains unvalidated.

2. Falsifiable system claims

The whitepaper is organized around testable claims rather than a single maturity label.

ID	Present claim	Falsifier or required test
C1	Known and novel compounds are routed through distinct computational paths.	Runtime inspection must show that the novel-structure route does not silently invoke the in-graph predictor.
C2	The evaluated system is reproducibly bound to explicit graph and model artifacts.	Integrity verification failure falsifies artifact identity.
C3	Some known-compound heads outperform simple development baselines on the recovered historical cohort.	Paired baseline comparison on the same rows; adverse events are already a counterexample.
C4	The symbolic composer is deterministic and can expose the relations used.	Repeated identical inputs must reproduce paths and evidence identifiers; missing relations must trigger abstention or explicit unknowns.
C5	Novel-structure outputs are hypotheses with an applicability domain, not transferred known-compound predictions.	The output contract must preserve route, similarity, coverage, action, and abstention fields.
C6	Reference overlays do not alter evaluated predictions unless explicitly integrated and validated.	Dependency and ablation tests must show no hidden reference-layer consumption.
C7	Human-facing reports do not create new predictive facts.	Narrative claims must be traceable to structured output or allowlisted evidence; invalid citations must be withheld.
C8	Current evidence does not support a global end-to-end accuracy claim.	A sufficiently powered, frozen, dual-disjoint benchmark could replace this null; none presently exists.
C9	The current application includes a report-driven 3D Brain Explorer that preserves map-channel identity.	Application contracts must reject malformed maps, keep predicted, literature, measured, legacy, and reference views distinct, and label non-predictive or diagnostic content.
C10	The current codebase contains a routed inference API and an application that can preserve structured report identity through query, streaming, inspection, and JSON export.	Contract tests must show that route selection, report fields, evidence links, reliability state, and exported JSON are not silently changed or dropped across the service and application boundary.

3. System architecture

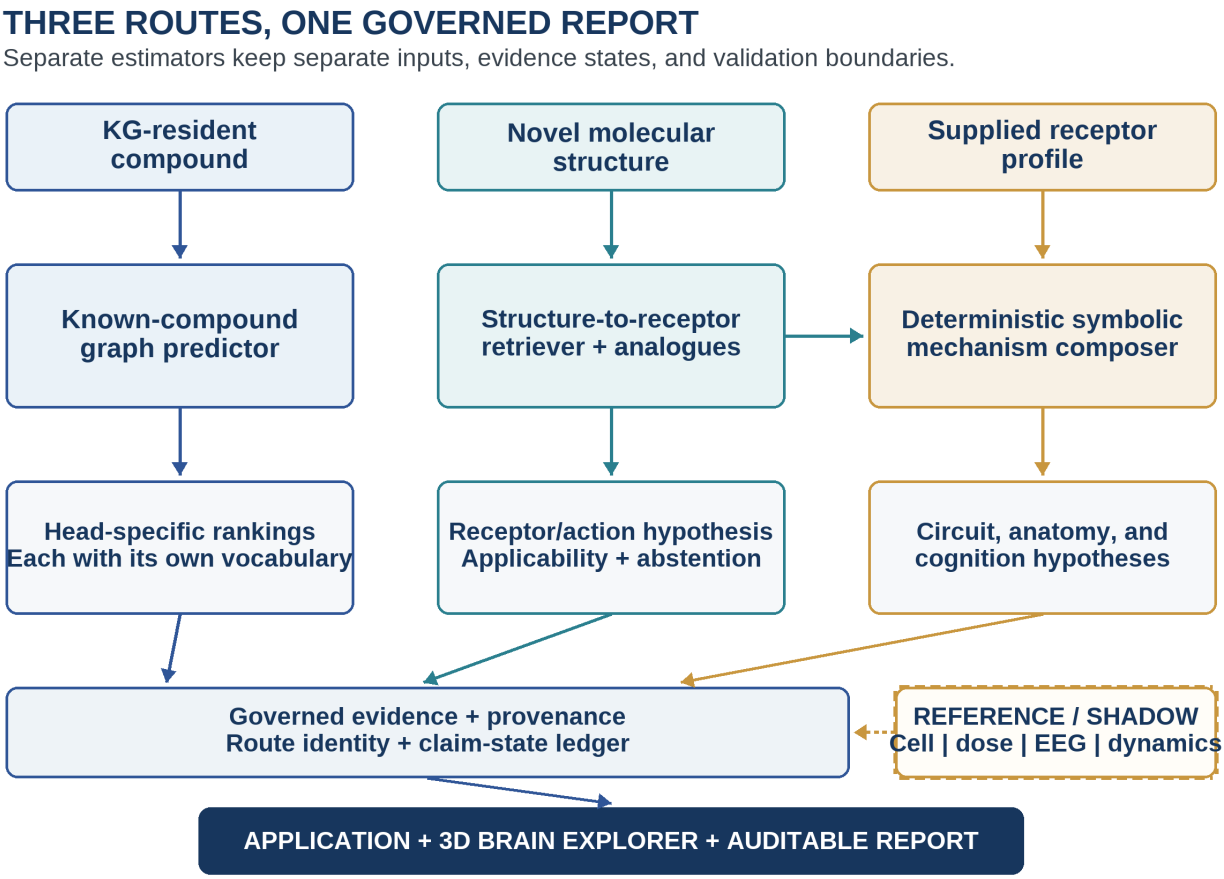

Figure 2: Routed system architecture. The predictor, retriever, and symbolic composer are separate estimators with distinct validation boundaries; reference and shadow layers do not automatically influence evaluated outputs.

3.1 Separation of predictor, retriever, and reasoner

NeuroLab is deliberately routed. The predictor, retriever, and reasoner solve different problems and have different failure modes.

1. **Known-compound predictor.** A knowledge-graph-resident compound is represented by chemistry and selected graph context, then scored by a multi-head heterogeneous graph model. This is transductive inference over the governed graph lineage.
2. **Novel-structure retriever.** A structure outside the graph is encoded from chemistry, compared with receptor-labelled reference evidence, and assigned a receptor hypothesis when support is adequate. It does not receive the known-compound predictor outputs.
3. **Symbolic composer.** A receptor profile—predicted, curated, or caller supplied—is propagated through explicit mechanism, circuit, nucleus, tract, parcel, and cognition relations. The result is conditional on the receptor profile and on the availability and context of those relations.

A shared provenance layer stores source identifiers, relation context, method-review state, and surfaceability. The application presents raw reports, narrative inspection, and the 3D Brain Explorer without erasing route differences. This architecture prevents an upstream estimate from being mistaken for a measured input and prevents reference-only knowledge from silently improving a production score.

Related systems. NeuroLab is complementary to, rather than a replacement for, whole-brain dynamical simulators [17,18], which model network dynamics from anatomical connectivity but do not carry compound-level pharmacological provenance, and to knowledge-graph mechanism-prediction frameworks [19,20], which predict treatment links without a multiscale anatomical composer or a routed applicability domain. The distinguishing commitment here is that a compound enters through a declared route, and every downstream claim inherits that route's evidence state.

3.2 Route contracts

Known-compound route

The known-compound route resolves a graph-resident compound, returns head-specific rankings, and attaches route and evidence metadata. It is appropriate only for compounds represented in the same governed lineage. It is not an inductive claim about arbitrary chemistry.

Novel-structure route

The novel-structure route derives a chemical representation, estimates a receptor profile from learned and analogue evidence, retrieves supported functional-action information, and passes explicit hypotheses to downstream composition. Applicability-domain gates must operate before scientific interpretation. A novel-structure result cannot inherit known-compound predictor metrics.

Receptor-conditioned route

When a user supplies receptors or functional actions, downstream results are explicitly conditional. The API cannot determine whether those values were measured, curated, or hypothetical unless that provenance is supplied. These results are useful for counterfactual exploration and mechanism decomposition, not for validating the upstream receptor estimate.

3.3 Query workflow and user surface

The implemented web application provides two entry paths. A basic query accepts a governed compound name and launches the streamed known-compound workflow. An advanced query can supply a compound name, a SMILES representation, or both, together with optional dose context and explicit receptor hypotheses. If a name is absent from the governed graph, the router returns a typed `needs_smiles` state and the interface asks for a structure rather than guessing an identity. Dose and receptor values supplied by a user remain query conditions: they are not automatically treated as measured exposure, target occupancy, or decision-grade evidence.

The returned result can be inspected as a readable report, a complete raw payload, or an interactive spatial view. Users can ask follow-up questions about the report as a whole, attach a named report section, or point to an exact JSON node. The follow-up layer receives the structured route and report context, validates attached-source references and citation identifiers against a context-derived allowlist, and withholds an answer when that validation fails. It is a report interpreter, not a second scientific estimator.

3.4 Implemented inference and research API

NeuroLab has a source-implemented HTTP API built around typed request and response contracts. It supports the application and can support controlled integration with research notebooks, workflow systems, or other interfaces. The API is implemented and tested at the code-contract level; this paper does not claim that an unrestricted public endpoint is presently available or that unreleased contracts will remain unchanged.

Implemented surface	Function	Scientific interpretation
`GET /health` and generated OpenAPI descriptions	Report service readiness plus API and served-model identity; expose machine-readable schemas when a service is running	Operational readiness and schema discovery, not biological validity or evidence of public deployment
`POST /predict`	Return a synchronous structured report for a governed, graph-resident compound	Known-compound routed output only; not an inductive claim about arbitrary chemistry
`POST /predict/stream`	Deliver the known-compound workflow as Server-Sent Events, including progress, report fields, evidence additions, and the finalized report	An alternative delivery mode for the same report, not a faster or scientifically different model
`POST /predict/router`	Select the known-compound or structure route and return explicit route, input-type, graph-residency, or `needs_smiles` state	Prevents an unknown name from silently inheriting a known-compound path
`POST /structure`	Accept a SMILES query and optional explicitly conditional receptor or context inputs	Novel-structure hypothesis generation with applicability and abstention controls; not a novel-node pass through the graph predictor
`POST /explain` and `POST /brain`	Provide specialist known-compound explanation and spatial-analysis views	Analytical support surfaces; their presence does not validate causal drivers or compound-specific brain maps
`POST /optimize`	Compare a governed compound with related profiles and candidate receptor changes for a stated model target	Controlled sensitivity and analogue analysis; not a validated molecule generator, recommendation engine, or proof of improved efficacy or safety
GET /compound/{name}, GET /cortex/{area}, and GET /circuit/{framework}	Query governed compound, cortical-area, and circuit reference records	Reference lookup; returned context must not be mistaken for query-specific causal activation

The default application uses the streaming known-compound route and the synchronous router for advanced structure or receptor-conditioned inputs. More specialized endpoints are part of the implemented research service but are not all exposed as primary controls in the current user workflow.

3.5 Structured reports, outputs, and interoperability

The structured report is the system of record for a query. Human-readable narrative and spatial views are derived presentation layers. Route-specific reports differ in scientific meaning, but the application preserves a common set of audit functions.

Report layer	Examples of returned or presented content	Contract boundary
Query and route identity	Compound display identity, submitted structure where applicable, selected route, input type, graph-residency state, and service/model identity	A display name or route label is not molecular-identity validation by itself
Ranked hypotheses	Cognition, reported usage, adverse-event (demoted; see 11.2), therapeutic-class, mechanism-of-action, or receptor rankings available on the selected route	Values retain their head- and route-specific interpretation; they are not one calibrated probability scale
Mechanism and anatomy	Conditional fine-cognition, mechanism paths, circuits, nuclei, parcels, networks, and affected-system annotations when supported and returned	Conditional, shadow, reference, and predictive channels remain distinguishable

Report layer	Examples of returned or presented content	Contract boundary
Evidence and provenance	Source-safe evidence summaries, resolvable citations where surfaceable, license-redaction state, and optional report-local support links	A citation or support-path count is not a causal weight or correctness score
Reliability and abstention	Applicability-domain state, scope/readiness fields, warnings, missing support, withheld directions, and typed abstentions where emitted by the route	Uncertainty categories and seed stability are not posterior probabilities
Spatial channels	Available 392-parcel maps, network and lobe summaries, literature or measured slots when actually supplied, and legacy diagnostic compatibility channels	Absent measured data remain absent; map channels are selectable and unfused
Presentation and export	Readable report, report-grounded follow-up, complete raw JSON, selected-section JSON, copied JSON values, and exact JSON Pointers	Narrative or visualization cannot add scientific evidence beyond the structured payload

Technical interoperability currently rests on standard JSON over HTTP, generated OpenAPI schemas, Server-Sent Events for progressive delivery, report-local JSON Pointers, and downloadable complete or sectional JSON. These interfaces make the result portable without requiring a recipient to parse the visual interface. They do not establish semantic equivalence with an external biomedical standard, long-term contract stability, persistent project history, or public-service availability; those are separate product and governance gates.

The application also exposes an interactive 3D Brain Explorer after a completed report. It can render an available 392-parcel channel, reference anatomy, subcortical and nucleus geometry, tract overlays, selected-region inspection, and map, network, parcel, and affected-system summaries. If no report map is available, it remains a non-predictive anatomy view. Predicted hypotheses, literature evidence, measured maps, legacy diagnostic maps, and reference anatomy remain selectable but unfused. Rendering a value does not validate its biology.

3.6 Component governance

Governed component	Scientific responsibility
Knowledge substrate	Provenance, schemas, releases, exports, and reference overlays
Model development	Feature governance, split policy, leakage gates, training, and candidate evaluation
Inference system	Route composition, abstention contracts, integrity records, and validation reports
Application	Input routing, report presentation, brain views, and narrative safeguards
Cross-component claim ledger	Reconciliation of public claims; not executable truth

Executable artifacts and their bound integrity records take precedence over narrative descriptions. Experimental or unverified work is not evidence for the evaluated system unless integrated, bound, and reverified.

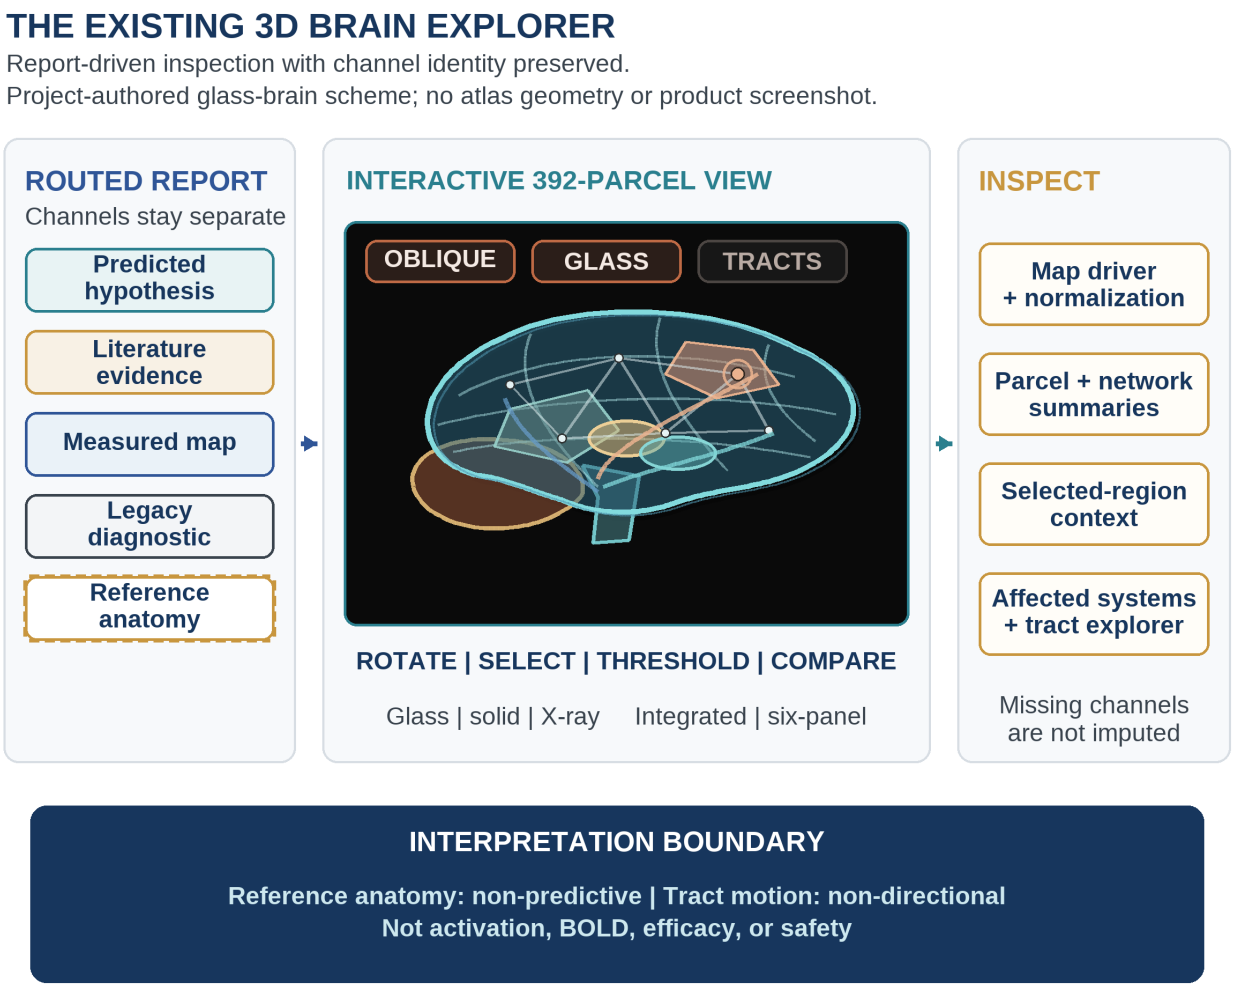

Figure 3: Existing report-integrated 3D Brain Explorer. The interface can inspect available 392-parcel report channels, reference anatomy, affected systems, and tracts while preserving channel identity. The project-authored illustrative glass-brain silhouette is not a product screenshot or atlas-derived geometry and does not imply biological validation or launched public access.

4. Knowledge substrate and provenance

4.1 Governed substrate scope

The governed substrate spans compounds, molecular targets, genes, anatomical regions, cognitive concepts, clinical categories, literature evidence, and reference overlays. Its breadth is descriptive inventory, not predictive-performance evidence. Exact inventory counts and release identifiers are retained in the controlled technical record because they are unnecessary for evaluating the claims made in this public edition.

The known-compound predictor consumes only a strict, task-governed projection of that substrate. Most multimodal relations are therefore not learned by the predictor. Reference layers can support retrieval, explanation, gap analysis, and symbolic composition, but they receive no credit for predictor performance unless an explicit, tested dependency and ablation demonstrate benefit.

4.2 Source classes

The governed substrate draws on established pharmacological, molecular, pathway, adverse-event, transcriptomic, anatomical, and neuroimaging resources. Examples include IUPHAR/BPS Guide to PHARMACOLOGY, ChEMBL, BindingDB, Reactome, SIDER, openFDA, the Allen Human Brain Atlas, NeuroSynth, NeuroQuery, Julich-Brain, receptor-density atlases, neuromaps, the Glasser cortical parcellation, and the Tian subcortical atlas [1–13,21,22]. Chemical representations and split analyses use established circular fingerprints and molecular-framework concepts [14,15].

The source list is heterogeneous in purpose and evidentiary strength. A database record can establish an identifier or an observed association without establishing causal direction, human relevance, dose dependence, or a query-specific mechanism. Each downstream claim must therefore state what the source actually supports.

4.3 Atomic provenance model

Scientific facts are curated as atomic relations with source identity and, where available, species, preparation, state, task, intervention, dose, measurement, direction, and method-review status. A citation is not accepted merely because its abstract mentions the relevant entities; the methods and results must support the exact relation and context. Unknown sign or context remains unknown.

The project distinguishes:

- **bibliographic resolution:** the identifier resolves to a real record;
- **claim support:** the paper's methods and results support the encoded claim;
- **surfaceability:** the relation can be shown to a user under its license and provenance contract;
- **independence:** evaluation evidence does not reuse the same source family or curation decision as support; and
- **runtime eligibility:** the relation is admitted to an evaluated route rather than remaining reference-only.

4.4 Provenance limits and incidents

Provenance engineering is substantial but incomplete. Internal audits have found unresolved provenance records, plausible-looking but invalid identifiers, topic mismatches, contradictions, and classification errors. Correcting a source ledger does not automatically synchronize every downstream surface; affected evidence must be reverified or quarantined before use.

These are not minor editorial issues. They show why bibliographic resolution, claim review, surfaceability, and deployment synchronization are separate gates. The scientific statements in this paper are limited to claims traceable to reviewed artifacts or independently verified primary sources; this does not certify every relation in the broader knowledge substrate. The detailed incident inventory remains in the controlled audit record because publishing its internal locations would add little scientific value while exposing unnecessary implementation detail.

4.5 Licenses and redistribution

Source availability does not imply unrestricted redistribution. Each export requires a source-specific license and surfaceability review, and a model artifact should record which data were used for training versus reference. A publication release should include a machine-readable source ledger with version, access date, use class,

redistribution status, and relevant license or terms. The current whitepaper does not certify that every underlying record can be redistributed.

The public whitepaper package is deliberately narrower than the controlled research substrate. Five of its six figures are project-generated schematics or a plot of internally generated aggregate latency observations. The sixth renders values derived from openly licensed external sources—measured receptor-imaging maps and cortical parcel coordinates—as a project-generated visualization; it reproduces no third-party figure, image, mesh, or surface geometry, and its silhouette is computed from parcel centroids rather than taken from an atlas. The Brain Explorer figure uses a project-authored illustrative glass-brain silhouette rather than application, atlas, mesh, parcel, tract, or source-derived geometry. The package includes no third-party paper figure, atlas image or geometry, application screenshot, external dataset extract, raw evidence row, benchmark trace, model artifact, or source corpus. Figure lettering is rasterized from an openly licensed font, and no figure-font file is packaged as a public asset. External atlases, datasets, and papers are identified in the reference list, with generic descriptions used in body prose and figure captions. Restrictively licensed or otherwise protected source material remains outside the public package unless a later release records an explicit redistribution basis and required attribution.

5. Known-compound predictor

5.1 Model and input representation

The known-compound model is an ensemble heterogeneous message-passing predictor. It combines a governed chemical representation with a selected, leakage-controlled graph projection and produces separate rankings for brain, cognition, reported usage, adverse events (demoted; see 11.2), therapeutic class, and mechanism-of-action categories (the adverse-event head is demoted; see 11.2). The ensemble supports a limited stability diagnostic and engineering continuity; it is not a Bayesian posterior.

Exact feature composition, operator configuration, dimensions, training schedule, random seeds, and checkpoint layout are controlled implementation details. They are withheld from the public paper because scientific evaluation depends on the declared inputs, split, outputs, comparators, and reproduced results—not on disclosing a copy-ready recipe. A qualified independent evaluator can receive the frozen executable and benchmark package under an appropriate access agreement.

5.2 Leakage controls

Supervision-target relations are excluded from message passing by the graph-integrity gate. This prevents direct reuse of labels as graph edges during training and inference. That relation-level control does not repair the historical split itself.

Reconstruction of the nominal scaffold split identified train/test scaffold overlap. The original source split and an untouched final-test partition cannot be reconstructed from the controlled evaluation record. Consequently, the reportable status is **de-leaked pooled historical development evidence**. It is not scaffold-OOD, independent, or prospective evidence.

5.3 Output interpretation

Each head produces a ranking over a constrained label vocabulary. Hit@1 asks whether any gold label is first; hit@5 asks whether any gold label appears in the first five; precision@5 is the fraction of the five returned labels that are gold. These metrics describe retrieval within a label set. They do not directly estimate effect size, clinical utility, safety, or a calibrated probability for a particular label.

The trained brain head must be distinguished from the currently served brain-map presentation, which is a composed output rather than the raw head alone. Any future benchmark must bind the evaluated output to the exact served route.

6. Novel-structure retriever

6.1 Purpose and separation from the known-compound predictor

The structure route addresses compounds not represented in the governed knowledge graph. It first estimates receptor engagement from chemistry, then conditions downstream retrieval and symbolic composition on that estimate. This is scientifically different from transductive prediction over known nodes and must carry separate metrics.

6.2 Current method

The route combines a molecular representation, a learned structure model, and chemistry-neighbour evidence. It emits receptor ranks, functional-action hypotheses where analogue support exists, applicability and coverage metadata, and abstention signals. Receptor-conditioned components can rank downstream categories, while graph retrieval supplies explicit circuit and anatomical relations to the symbolic composer. Blend weights, thresholds, feature details, and retrieval precedence are withheld as implementation know-how.

6.3 Applicability domain

Novelty is not binary. A structure can be outside the KG while remaining close to a labelled analog, or it can be far outside the supported chemical domain. Interpretation therefore requires at least:

- exact-self and near-duplicate exclusion;
- maximum similarity to the comparator pool;
- scaffold or chemistry-component separation;
- receptor coverage and action-sign coverage;
- polypharmacology and conflict flags;
- blood-brain-barrier evidence status;
- source-family independence; and
- an explicit abstention reason when the route is unsupported.

The current scorecard packages a reproducible chemistry-neighbour comparator, but not an independent evaluation of the exact served fused route. Historical fused-model values therefore remain development metadata rather than current served-route evidence.

6.4 Analog functional-action layer

An analogue-sign diagnostic performs strongly only on the subset for which it commits. The numerator cannot be interpreted without coverage because uncovered cases abstain. Its evaluator and source corpus are development assets rather than an independent benchmark, and no bound confidence interval is available. This result is therefore an engineering diagnostic, not calibrated biological reliability.

7. Symbolic mechanism and circuit composition

7.1 Deterministic composition

The symbolic reasoner accepts a receptor profile and composes explicit relations through receptor function, circuit frameworks, source nuclei, tracts, parcels, and fine cognitive functions. The purpose is to expose a mechanistic chain and its missing links, not to turn literature associations into quantitative dynamics. Exact internal schema labels and inventory counts are withheld from this public edition.

Composition is deterministic for a fixed artifact and input. Evidence strength, applicability, unsigned relations, contradictions, and missing stages should remain visible. A pathway that exists in the KG is not activated unless the query and route satisfy its input contract.

7.2 Conditionality and oracle diagnostics

When the composer receives curated receptors, its output is an oracle-receptor component diagnostic. When it receives predicted receptors, upstream error propagates. These two cases cannot share a headline metric. The strongest fine-cognition evidence is conditional on oracle receptors, was reused during development, and is not independent. The complete structure route has too few evaluable cases for an aggregate claim.

7.3 Mechanism is not yet a dynamical model

The current composer lacks the quantitative transition functions needed for mechanistic simulation: unbound exposure, occupancy, efficacy, cell- and compartment-specific response, population-state transitions, recurrent circuit dynamics, state and task dependence, and time-varying observation models. It can return a signed or unsigned path hypothesis when supported, but it cannot yet integrate the path into a validated trajectory.

8. Brain, cellular, dose, species, and temporal layers

8.1 Brain parcellation and map channels

The spatial vocabulary uses established cortical and subcortical parcellations [21,22]. The implemented 3D Brain Explorer distinguishes predicted, literature-derived, measured, legacy, and non-predictive reference views. The interface is operational even though the scientific maturity of its channels differs. By default the Explorer presents the spatial channel returned with the report and labels its evidence state; receptor-density and circuit-framework channels are available on request as explicitly non-predictive hypotheses; and literature or measured slots appear only when those inputs are loaded. Channels remain selectable and are never fused. A receptor-density map is normalized per receptor across the parcellation: it expresses where a given receptor is relatively enriched, not its absolute abundance, and values are not comparable between receptors. A circuit-framework map represents framework salience. Neither is automatically a drug-induced activity map.

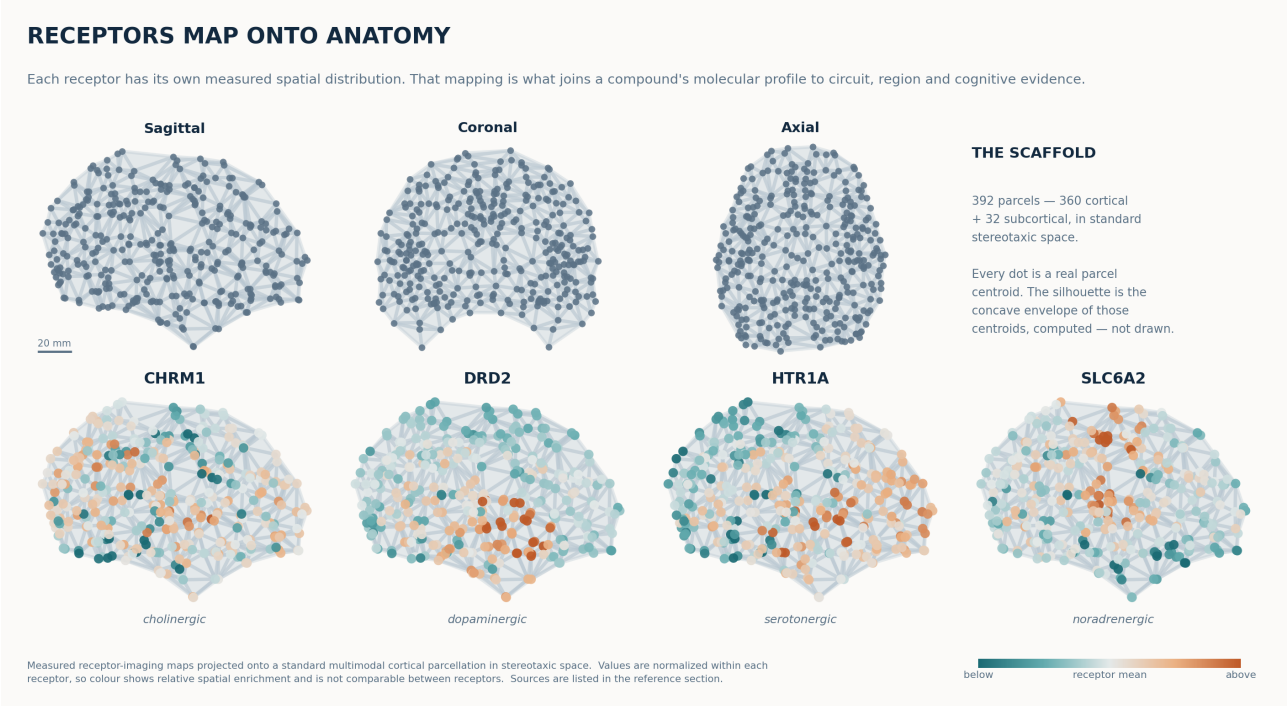

Figure 4: The measured receptor scaffold beneath the spatial channel. Parcel positions are real centroids of a standard multimodal cortical parcellation in stereotaxic space, and each lower panel is a measured receptor-imaging map for one neuromodulatory system. The silhouette is the computed concave envelope of those centroids rather than an illustrated brain or a rendered atlas surface. Maps are normalized within each receptor, so colour shows relative spatial enrichment and is not comparable between receptors; a rendered value is not a drug-induced activity map.

Current brain validation is inadequate for an accuracy claim. The controlled evaluation contains only a single scoreable observation from one source family, and the evaluated density view misses the predeclared top-k criterion. There is no reportable aggregate, directional, or novel-structure brain accuracy.

8.2 Reference virtual-brain substrate

The reference substrate includes receptor-by-cell-type/layer records, dozens of nucleus or region profiles, and hundreds of tract records. These provide anatomical constraints, lookup evidence, and a coverage ledger. They are not fitted simulation parameters, message-passing weights, or validated response operators. Receptor atlases and spatial maps can constrain hypotheses, but modality, preprocessing, population average, and measurement scale must remain explicit [7,10–13].

8.3 Typed population and pathway shadow

Three maturity states must be kept separate: governed reference material, a reviewed reference candidate, and experimental compilers. None of these states implies that the material trains, scores, or routes an evaluated prediction.

The current typed population/pathway candidate has no admitted quantitative transition parameters and no eligible context-complete evaluation observations. Only a small minority of audited whole-brain framework references are typed into causal pathways, and no domain has every required causal stage. The reference layer is valuable precisely because it measures this deficit. It should not be described as a complete mechanistic simulator.

8.4 Exact-context mechanistic reconstruction

The experimental EEG compiler uses a conservative match contract: compound, species, preparation, state, task, manipulation, measurement, dose, source, and operator must match. This supports narrow reconstructions, not transferable laws. Examples of appropriately bounded primary evidence include:

- a ketanserin-sensitive posterior alpha decrease in a human ayahuasca-mixture protocol, which does not establish a DMT-alone or generic psychedelic rule [23];
- caffeine-dependent arousal through nucleus-accumbens-shell A2A signalling in an exact rodent protocol, not a universal human pathway [24];
- local rat locus-coeruleus alpha-2A involvement in dexmedetomidine loss-of-righting-reflex duration, not a generic human sedation or EEG model [25];
- basal-forebrain cholinergic modulation of propofol anaesthesia and EEG within the reported mouse protocol, without proving unique mediation [26]; and
- broad primate dlPFC cell-class homology that does not transfer receptor abundance, dynamics, dose response, or cognitive effect to humans [27].

The mechanistic reconstruction examples were selected during curation. They are not an unseen test, have unresolved source-family independence in important cases, and do not support calibrated probabilities. The correct status is “mechanistic compiler stress test; benchmark not ready.”

8.5 Dose and exposure

The knowledge substrate stores dosage and administration facts from regulatory, institutional, and research sources. These facts are reference-only. The system records an administered dose as context but does not convert it to plasma pharmacokinetics, unbound brain exposure, metabolite exposure, receptor occupancy, or functional efficacy. A report must not imply otherwise.

A quantitative dose path requires, at minimum:

administered dose → formulation and route → plasma exposure → unbound brain exposure → target occupancy → functional efficacy → local operator → circuit state → observation

Each transition needs uncertainty, context, provenance, and an abstention condition. Required parameter classes are recorded, but no quantitative values are admitted to the causal candidate.

8.6 Cross-species transfer

Species transfer is an explicit inference, not a default. Conservative evaluators abstain on the available retrospective endpoint cases. An intentionally unsafe universal-transfer rule produces systematic false positives. The negative control supports fail-closed transfer: shared cell classes or homologous anatomy do not by themselves license shared parameters or endpoints.

8.7 Temporal evidence

The temporal pilot yields no scoreable cases after methods/results and leakage gates. The historical evidence projection is sparse, dominated by one relation family, and partly conditioned on entities present in the current substrate. Consequently, no temporal generalization, prospective discovery, drift, or historical-novelty claim is allowed.

9. Uncertainty, applicability, and abstention

9.1 What the scores mean

Model outputs are ranking scores or route-readiness signals. They are not calibrated probabilities of biological truth, clinical effect, or safety. Ensemble dispersion provides a limited stability diagnostic; the members share data, architecture, feature construction, and curation, so their dispersion cannot represent total epistemic uncertainty.

9.2 Sources of uncertainty

The uncertainty budget is route-dependent:

Source	Examples	Current treatment
Chemical domain	distant scaffold; stereochemistry; salt or metabolite mismatch	similarity and applicability metadata; partial abstention
Target identity	crosswalk error; family-level target; missing isoform	provenance flags; unresolved cases must be quarantined
Functional action	agonism versus antagonism; partial efficacy; biased signalling	analog sign where covered; otherwise unknown
Exposure	BBB; free fraction; active metabolite; dose and route	mostly reference/context; no quantitative occupancy
Biological context	species, cell class, layer, state, task, disease	exact-context matching in shadow layers; sparse coverage
Causal transition	unsigned or absent link; competing pathway; homeostasis	explicit unknowns and evidence caps
Label and evaluation	sparse positives; prevalence; source-family reuse	head-specific metrics and baseline comparison
Observation	atlas mismatch; modality and preprocessing; temporal window	separate map channels; largely unvalidated

9.3 Required calibration program

Calibration must be measured for each route and endpoint using proper scoring and selective prediction. A defensible program should report Brier score or log loss where labels permit, expected calibration error with bin uncertainty, reliability diagrams, coverage-risk curves, abstention precision, and failure rates stratified by similarity, receptor/action support, polypharmacology, BBB evidence, source strength, and label prevalence. Calibration learned on a reused development cohort must not be reported as external validity.

9.4 Abstention contract

The system should abstain when any indispensable transition is unsupported, not merely when a final score is low. An abstention record should identify the first blocking transition, all material upstream assumptions, the evidence grade, the nearest supported domain, and the experiment or datum that could resolve the block. “Unknown” is a valid scientific output.

10. Validation design

10.1 Evaluation units

The correct unit depends on the claim. Compound-level bootstrap intervals quantify row sampling variability but not scaffold, source-family, seed, or curation uncertainty. Scaffold and chemistry-component splits address structural dependence. Source-family splits address evidence reuse. Temporal splits address availability and drift. End-to-end evaluation must hold out the upstream chemistry and downstream evidence used to construct or curate the mechanism.

10.2 Metrics

For a query with gold-label set **G** and ranked list **R**:

- **Hit@k = 1** if at least one member of **G** occurs in the first **k** positions of **R**, otherwise 0.
- **Precision@k = intersection(G, R[1:k]) / k**.
- **Coverage** is the fraction of queries for which a route returns a non-abstaining result.
- **Selective risk** is error among non-abstaining predictions as coverage varies.

Hit@k should always be reported with the label vocabulary, denominator, cohort construction, and baseline. High hit@5 can coexist with low precision@5 when labels are broad or multi-label. Conditional accuracy must be reported together with coverage.

10.3 Baselines and nulls

Relevant controls include:

- label-frequency and random rankings;
- chemistry nearest neighbours and fingerprint-only models;
- receptor-only decoders;
- unsigned or sign-shuffled mechanisms;
- density-only and population-average brain maps;
- no-context and unsafe cross-species transfer rules;
- ablation of symbolic composition; and
- oracle receptors as an upper bound, separated from end-to-end input.

An architecture is not supported merely because its result is above zero. It must beat relevant simple baselines on a frozen cohort and retain the claim under dependence-aware uncertainty.

10.4 Current verification performed for this snapshot

Check	Result	Interpretation
Benchmark status reconciliation	PASS	Current, historical, oracle, coverage-only, and failed-baseline states are kept distinct.
Artifact-integrity controls	PASS in the controlled release audit	Integrity checks establish byte consistency; they do not establish biological validity or ownership.
Route and reporting contracts	PASS in the controlled software audit, with documented unexecuted paths	Verified contracts support software behavior only; unexecuted paths are not implied to pass.
Reference-layer conformance	PASS in the controlled structural audit	Schema and fail-closed checks do not establish predictive accuracy.
Public disclosure review	PASS for this draft	Direct affiliations, repository identifiers, internal version labels, detailed benchmark tuples, and document history are omitted.
Paper bibliography resolution	PASS for the cited records	Resolved identifiers establish record identity, not claim support for every relation in the broader knowledge substrate.

No retraining, model promotion, or deployment was performed for this whitepaper. Headline ranking metrics are reported below with their denominators and controls. Cohort identities, split artifacts, evaluator code, internal artifact identifiers, file locations, and private test logs are retained in the controlled validation record rather than disclosed publicly.

11. Current evaluation status

11.1 Known-compound development evidence

Historical development evaluations show heterogeneous head-specific performance. Several categorical heads outperform simple frequency controls on their recovered development cohorts. Those cohorts were inspected during development, use different denominators and label vocabularies, and are not independent final tests.

Claim family	Public result	Evidence state
Coarse cognition ranking	hit@1 0.852 (95% CI 0.801–0.903, n=196) against a 0.347 label-frequency control	Development cohort; the control is provisional and the source split cannot be reconstructed
Reported-usage ranking	hit@1 0.569 (95% CI 0.458–0.694, n=72) against a 0.333 label-frequency control	Development cohort; the control is provisional and the source split cannot be reconstructed
Therapeutic-class ranking	hit@1 0.690 (95% CI 0.517–0.862, n=29) against a 0.241 label-frequency control	Development cohort; small denominator; provisional control
Mechanism-of-action ranking	hit@1 0.594 (95% CI 0.531–0.657, n=239)	Development cohort; **no prevalence control is packaged for this head**, so the margin over chance is not established

These rows do not support a single model accuracy and are neither prospective nor a current scaffold-out-of-distribution evaluation.

11.2 Adverse-event demotion

The legacy adverse-event head is demoted because its external positive-label concordance is below matched prevalence. It must not be presented as validated prediction. Observational adverse-event retrieval and signed, context-dependent biological-effect prediction are different tasks and should be evaluated separately.

11.3 Novel-structure receptor ranking

The most reproducible current comparator is a fingerprint-nearest-neighbour evaluation using a scaffold-disjoint split.

Comparison	Public result	Interpretation
Chemical-neighbour receptor ranking vs frequency	Better on the controlled development split	Evidence for analogue-sensitive retrieval
Chemical-neighbour receptor ranking vs random	Better on the controlled development split	Sanity check only
Low-similarity slice	Performance degrades substantially	Broad novelty is not established
Exact served structure route	Not independently evaluated	No end-to-end claim

Performance is strongly similarity-dependent, and the strata were defined post hoc rather than as predeclared serving thresholds. The experiment supports analogue-sensitive receptor retrieval within a declared domain, not broad novel-scaffold generalization. Historical fused-route values are retained only in the controlled validation record because their gating differs from the current system and they are not current served-route evidence.

11.4 Conditional, coverage, and end-to-end diagnostics

Component or behavior	Public result	Status
Analogue receptor-action sign	Strong only on the covered subset	Conditional development diagnostic; not independent
Oracle receptor → fine cognition	Positive reused-cohort result	Upper-bound diagnostic; not end to end
Receptor profile → directional cognition	Does not beat the strongest baseline after abstention penalty	Negative result
Complete structure → cognition	Too few evaluable cases	Unvalidated
Brain-map route	Single scoreable observation	Inadequate denominator
Mechanistic reconstruction set	Curated stress test	Not an unseen-compound benchmark

Coverage and output-presence measurements are retained in the controlled validation record. They are not correctness or calibration. A mechanism score is not a calibrated probability, and an empty or populated field does not establish correctness.

The current structure-profile AUROC is unmeasured. A retired predecessor AUROC must not be revived. Component estimates must not be multiplied to manufacture a full-chain accuracy.

11.5 Why overall accuracy is null

The canonical report deliberately sets overall accuracy to null. There are four reasons:

1. routes answer different questions and use different inputs;
2. conditional and oracle diagnostics receive information unavailable end to end;
3. several downstream denominators are too small or absent; and
4. scores are not calibrated or validated against a shared utility function.

Maintaining the null is scientifically stronger than reporting an average of incompatible metrics.

12. Engineering performance and reproducibility

12.1 Frozen release identity

The controlled release record binds the graph snapshot, model artifacts, evaluator inputs, and current benchmark scorecard by cryptographic digest. Exact artifact names, commits, and checkpoint digests are withheld from this public paper until a corresponding evaluation package is intentionally released. SHA-256 establishes byte integrity, not signer identity or scientific validity; an external evaluation release should add authenticated signing and a durable public record.

12.2 Historical CPU and GPU serving profile

An archived benchmark run on 25 June 2026 demonstrates practical execution of the routed service on CPU and GPU profiles. The controlled workload submitted Modafinil to `POST /predict/stream` with `top_k=5`, knowledge-graph support enabled, and an archived three-seed model profile. For the warmed backend request, the CPU mean was 3.84 s (observed range 3.80–3.91 s) and the RTX 5070 Ti GPU mean was 1.07 s (1.06–1.08 s), with three requests per profile. For this specific warmed request, the observed ratio was approximately 3.6×.

Measured quantity	CPU profile	GPU profile	Repetitions
Cold container and first request	15.02 s	14.72 s	n=3 per profile
Warm request after container readiness	4.91 s	1.28 s	CPU n=2; GPU n=3

Measured quantity	CPU profile	GPU profile	Repetitions
Warm backend request wall	3.84 s (3.80–3.91)	1.07 s (1.06–1.08)	n=3 per profile

These measurements establish historical engineering feasibility for one route and payload. They are not bound to the present whitepaper snapshot, the CPU hardware model was not recorded in the public benchmark report, and the cold/warm Docker totals use unequal warm repetition counts. They are therefore not a current service-level objective, a controlled hardware-equivalence study, or a claim about throughput, concurrency, other compounds, or the novel-structure route. Raw traces and full evidence payloads remain controlled because they contain implementation-sensitive details.

12.3 Performance refresh and operational scale

The historical profile shows that the architecture can support responsive interactive use and that accelerated execution materially reduces warmed latency for the measured workload. The next engineering release should convert this observation into a current operational claim by benchmarking both primary inference routes on named CPU and GPU hardware, exact image and artifact hashes, at least 30 repetitions, p50/p95/p99 latency, cold and warm paths, concurrency and throughput, resident and peak memory, representative payload sizes, failures, abstentions, and deterministic output equivalence. The live integration suite should run in the same frozen release environment and its result should be bound to the whitepaper manifest.

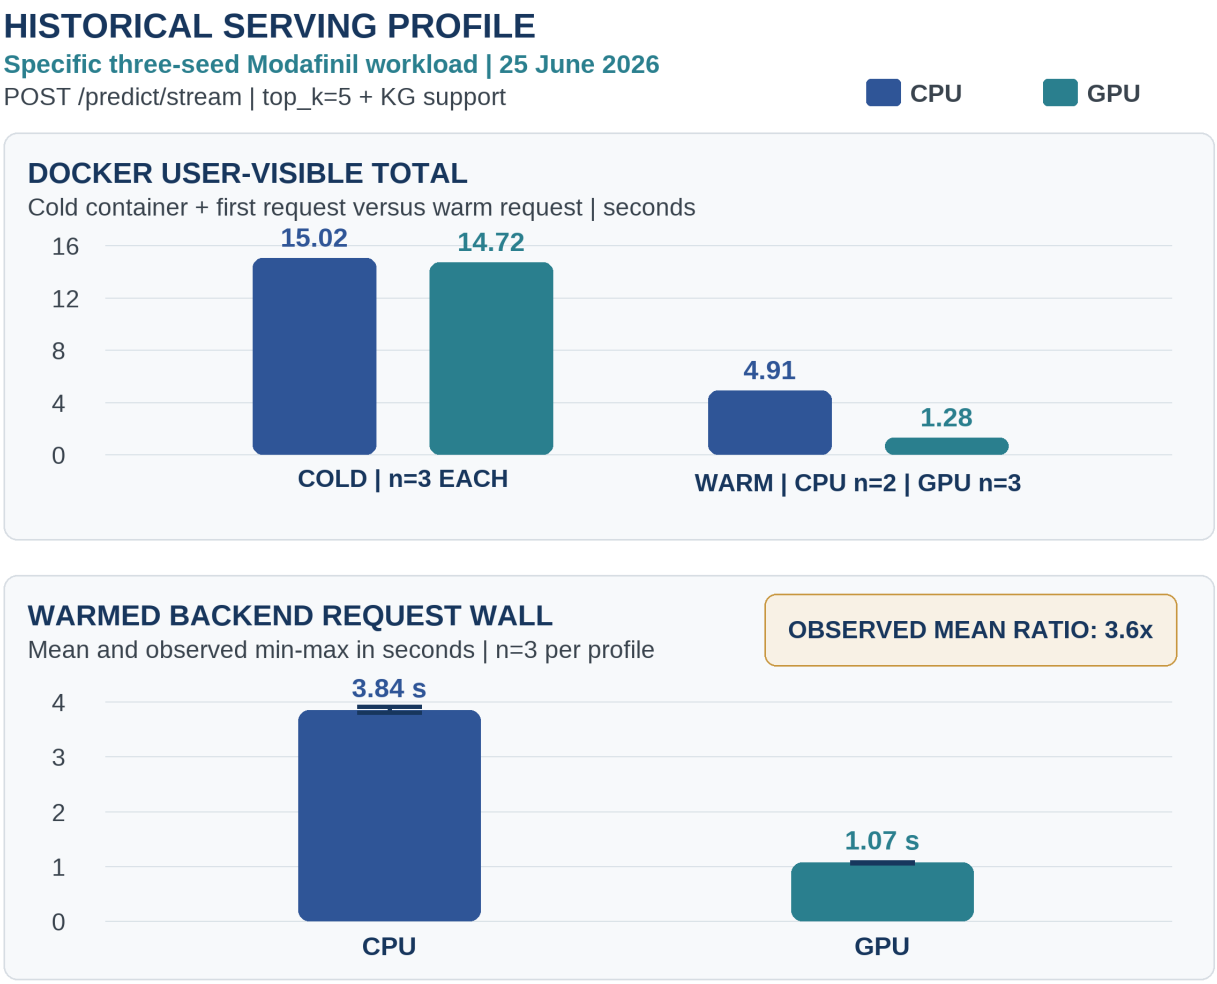

Figure 5: Historical serving profile for one three-seed Modafinil streaming workload. The GPU mean warmed backend request wall was approximately 3.6× faster than the CPU profile. These observations are not bound to the present snapshot and do not define a current service-level objective.

13. Model governance: retained negative results and retired claims

Scientific maturity includes retaining results that narrow the claim.

13.1 Adverse-event prediction

The legacy adverse-event head falls below its provisional frequency comparator and also records external concordance below matched prevalence. The head is demoted. This argues against treating heterogeneous reported adverse events as a direct mechanistic prediction without a signed, context- and goal-aware effect representation.

13.2 Historical leakage and split language

Legacy cognition values arose from retired or leakage-era evaluations and are excluded. The known-compound rows remain useful development diagnostics after relation-level de-leaking, but recovered scaffold overlaps invalidate "scaffold-OOD" wording. Stale result bindings must be reconciled to the controlled claim ledger before release.

13.3 Retired novel-structure values

Retired structure-to-receptor AUROC, brain-fusion, sparse-edge, and textbook-smoke values must not be reused as current independent evidence. The current structure-profile AUROC is unmeasured.

13.4 Unsafe transfer and map interpretation

The intentionally unsafe universal cross-species rule produces systematic false positives. This is evidence for abstention, not a reason to tune transfer until it appears positive. Likewise, the single brain pilot is partly below simple controls. Density, atlas, and class templates should remain controls or explanatory views until a held-out map benchmark demonstrates independent value.

13.5 Documentation and naming drift

Earlier descriptions sometimes conflated the novel-structure route with a "digital-twin path" or treated an incomplete chain as "symbolically complete." Those descriptions are retired. The canonical public name in this paper is **NeuroLab**. Its present product category is an auditable CNS hypothesis-generation and triage system.

14. Multiscale capability matrix

Scale	Current executable or governed asset	Scientific status	Required transition
Molecule	Governed identity, molecular representation, and similarity	Executable; identity gaps remain for edge cases	Stereochemistry, salts, metabolites, prodrugs, and mixture contracts
Receptor and channel	Known-compound graph context; novel structure receptor rank; curated target evidence	Executable hypotheses; novel benchmark internal	Frozen source- and chemistry-disjoint retriever evaluation
Functional action	Curated action and analog-sign layer where covered	Partial and conditional	Functional efficacy, isoform, concentration, bias, and uncertainty
Exposure	Dose and BBB reference facts	Reference/context only	PK, unbound brain exposure, metabolites, and occupancy
Cell and compartment	Receptor-by-cell/layer reference rows	Reference-only	Quantitative, context-matched operators with units
Population	Typed population and local-transition candidate	Shadow; zero admitted parameters	Identifiable state variables and independently adjudicated transitions
Pathway and circuit	Typed circuit frameworks, nuclei, tracts, symbolic composition	Executable qualitative composition; sparse typed causal coverage	Signed, context-conditioned dynamics and competing pathways

Scale	Current executable or governed asset	Scientific status	Required transition
Region and network	Standard parcellations, density and circuit views	Exploratory; inadequate brain-validation denominator	Held-out atomic evidence corpus and spatial nulls
Cognition	Coarse-cognition predictor; fine-cognition composer	Development and oracle diagnostics	Dual-disjoint end-to-end evaluation
Clinical	Usage, ATC, adverse-event label rankings	Development categorization, not clinical prediction	Endpoint definitions, utility, external validation, safety process
Time and state	Sparse temporal KG; context fields	No scoreable temporal result	Pre-cutoff evaluation and dynamic observation model
Individual	No subject parameterization	Planned	Identifiability, repeated measures, prospective within-subject gain

The matrix is intentionally asymmetric. A broad knowledge substrate can coexist with a narrow validated path. Progress requires validating the transitions, not merely adding nodes.

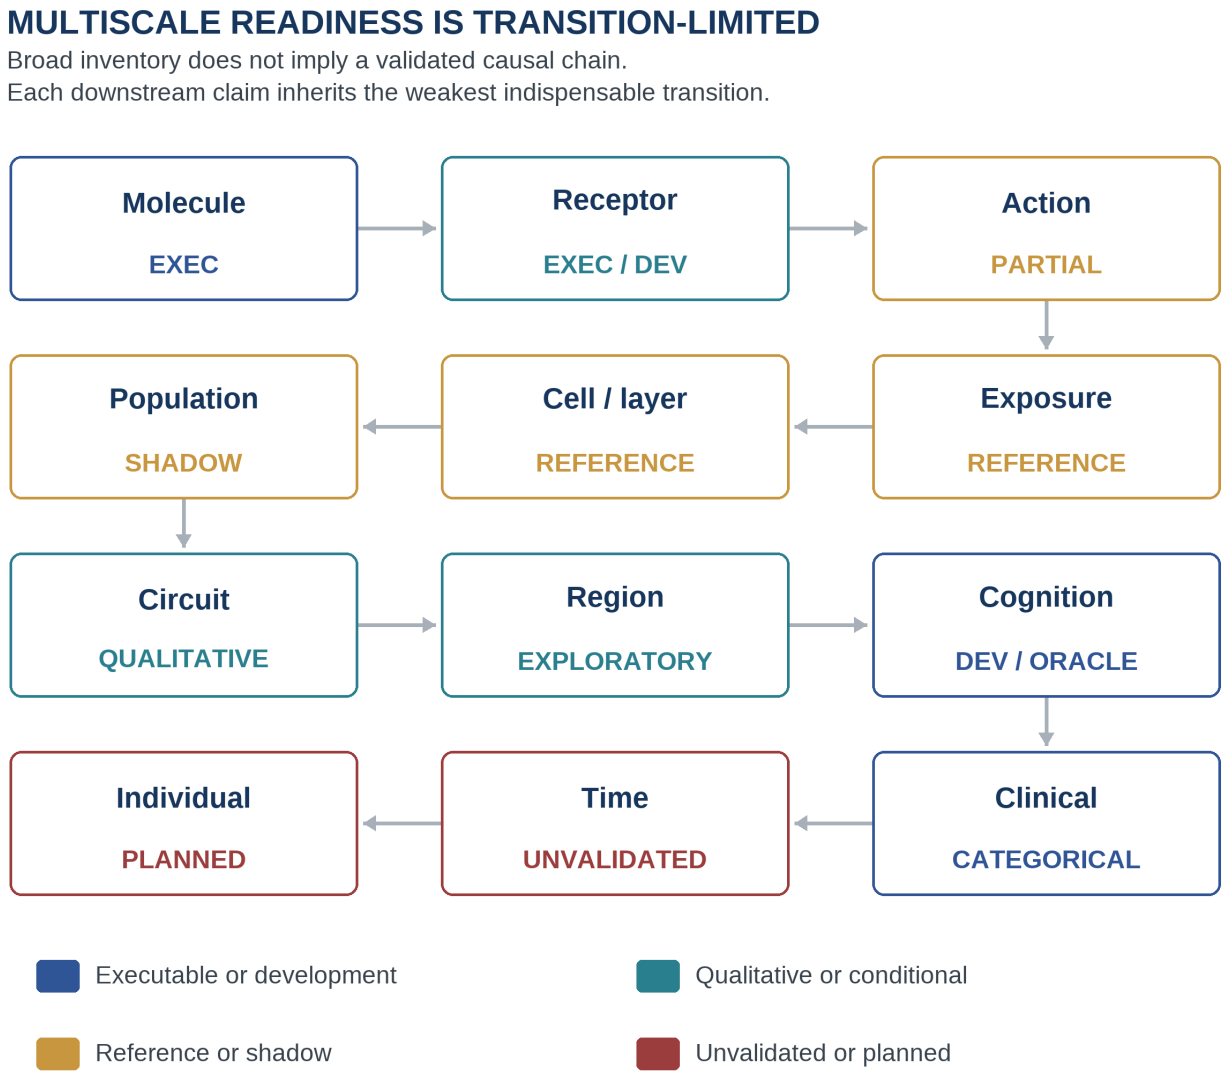

Figure 6: Multiscale maturity chain. Breadth of the knowledge substrate is not equivalent to end-to-end mechanistic validation.

15. Scientific development gates

15.1 Global evaluation rule

For every new benchmark, the protocol, primary endpoint, meaningful effect or precision margin, comparator set, subgroup analyses, stopping rule, and invalidation conditions must be frozen before test-label access. Data volume should follow prospective power or confidence-interval precision, not an arbitrary sample size or repeated extension until significance. Multiple heads, routes, and slices require multiplicity and benchmark-reuse controls.

15.2 P0 — publication truth and release hygiene

Gate	Acceptance criterion	Current blocker
P0.1 Canonical claim ledger	A machine-readable ledger binds every reported number to route, artifact identity, cohort, evaluator, and evidence state; every claim is typed as independent, development, oracle, software, reference/local, planned, or unsupported; an automated audit finds zero contradictions.	Historical wording and cross-document result drift require continuing reconciliation.
P0.2 Provenance, license, and synchronization closure	Every surfaced scientific assertion resolves to an atomic reviewed source, license/use class, and surfaceability state; unresolved records are quarantined; corrected evidence is verified across every downstream copy.	Unresolved or previously corrected provenance must remain quarantined until license and downstream synchronization checks close.
P0.3 Molecular identity gate	Protected compounds resolve to the correct molecular target and evidence contract before quantitative support is shown.	Known identity and direct-evidence edge cases must close before promotion.
P0.4 Reproducible benchmark registry	Cohort IDs, inclusion logic, preprocessing, comparator assets, exact served-runtime invocation, frozen evaluator, dependency lock, and hashes reproduce every main-table metric from a clean environment.	The exact served fused structure route lacks independent evaluation.
P0.5 Release authenticity and integration	Deterministic manifests reproduce bytes, an authorized signer authenticates the release, and standard plus live integration suites are bound to the snapshot.	Integrity is bound; public signing and complete live integration evidence remain pending.
P0.6 Report, Brain Explorer, and public-use honesty	Tests prove reports cannot merge routes, omit abstention, invent evidence, or promote reference or experimental work; Brain Explorer tests preserve map length, finite values, channel separation, diagnostic labels, and reference-only boundaries. Adversarial public-use tests must verify that the surrounding experience does not provide product endorsement, personal optimization, dosing, stacking, treatment, or self-experimentation advice, and user studies must test comprehension of evidence state, uncertainty, and abstention.	Display-contract safeguards exist, but public-user safeguards and comprehension have not yet been independently evaluated.

The first publishable whitepaper release is unlocked only when these documentation and provenance gates close. They do not require better biological results; they require truthful, reproducible reporting.

15.3 P1 — independent component validation

If a head fails its baseline—as adverse events currently do—it should remain a negative result or be removed from the primary product claim. Protected identity checks are necessary for deployment safety but are not evidence of general performance.

Gate	Acceptance criterion	Claim unlocked
P1.1 Known-compound predictor	A signed-off candidate is trained on a frozen chemistry-component split; label vocabularies are fixed from training only; final test remains blind; each head beats its strongest deployable baseline by a predeclared lower-bound margin.	Head-specific chemistry-OOD component performance.
P1.2 Structure → receptor/action	A frozen chemistry-component- and source-disjoint cohort uses immutable compound identities, includes low-similarity, rare-target, stereochemical, and sign slices, invokes the exact evaluated route, and saves per-compound predictions before any final refit; learned, neighbour, class-frequency, and permutation controls are frozen; selective risk beats the strongest control.	Validated novel-structure receptor ranking within a declared domain.
P1.3 Calibration and abstention	On untouched data, route-specific reliability and coverage-risk analyses meet a declared risk bound at stated coverage; abstention reduces error monotonically.	Quantitative interpretation of route confidence within the tested domain.
P1.4 Engineering refresh	Exact current bytes are benchmarked on named hardware with cold/warm p50/p95/p99, memory, concurrency, throughput, failure rate, and output-equivalence checks against frozen service objectives.	Current latency and resource claims.

15.4 P2 — end-to-end and mechanistic validation

Gate	Acceptance criterion	Claim unlocked
P2.1 Dual-disjoint full chain	Freeze the complete route and evaluate on compounds disjoint in chemistry and evidence source, excluding query and neighbour labels; compare structure-only, receptor-only, class, permutation, deployable full route, and oracle upper bound.	Novel structure → receptor → reasoner → endpoint performance.
P2.2 Brain-map evaluation	Build atomic targets from held-out papers, split jointly by article and compound, and compare parcel/network retrieval with global, class, receptor, legacy, and permutation baselines. Literature-evidence maps remain distinct from predictions.	Compound-specific brain localization within the tested context.
P2.3 Quantitative dose and context	Add unit-checked exposure, occupancy, efficacy, cell/population operator, state, task, time, and endpoint links; predict held-out direction and magnitude beyond static receptor, density, and sign-shuffled nulls.	Context-dependent mechanistic forward prediction.
P2.4 Temporal and species transfer	Freeze a pre-cutoff graph and property-scoped transfer rule; evaluate only later or held-out-source evidence against no-transfer controls; otherwise abstain.	Narrow temporal or cross-species claims for the tested property.
P2.5 EEG compiler	Preserve exact matching and evaluate blinded, source-family-independent interventions; separate reconstruction from forward prediction and compare shuffled-operator and no-mechanism nulls.	Independent mechanistic discrimination or prediction; numeric confidence only after calibration.

The dual-disjoint full-chain benchmark is the central scientific milestone because it tests the actual novel-compound promise rather than a component supplied with oracle inputs.

15.5 P3 — quantitative and individualized digital twin

Gate	Acceptance criterion	Claim unlocked
P3.1 Population multiscale model	Define state variables, units, coupling equations, parameter priors, numerical stability, and missing-link abstention; predict held-out longitudinal perturbation trajectories better than static KG, receptor-only, and phenomenological baselines.	Quantitative dynamic population forward simulation.
P3.2 Subject-specific model	Fit subject parameters only from pre-intervention repeated measures; demonstrate identifiability, test–retest reliability, calibrated uncertainty, and prospective within-subject improvement over the population model.	A subject-specific digital twin within the validated protocol.
P3.3 Decision-grade use	Independent prospective external validation, human-factors evaluation, privacy/security controls, fairness analysis, monitoring and rollback, and the applicable quality and regulatory process all pass.	Any treatment-selection or clinical-use claim.

Until P3.3, the system remains a hypothesis-generation research tool even if earlier simulation gates succeed.

15.6 Cross-cutting gaps

The following risks should be measured across all gates:

- coverage bias by sex, age, ancestry, disease state, species, administration route, and evidence-source geography;
- stereoisomers, mixtures, salts, prodrugs, transporters, active metabolites, polypharmacy, and exposure edge cases;
- curator inter-rater reliability and independent external replication;
- ontology and source drift, with explicit rules for invalidating old metrics;
- data licensing, consent, privacy, API threat modelling, prompt-injection resistance, and supply-chain signing;
- multiple testing and benchmark overuse; and
- user understanding of route-specific caveats and abstention.

16. Research roadmap

The roadmap follows causal dependency, not feature count. The items below state current work directions without assigning release dates or promising that an unresolved scientific gate will pass.

16.1 Current product and interoperability work

The next product steps continue interfaces and controls that are already present in the codebase:

1. **Freeze and identify API contracts.** Bind an explicit contract version to the query, router, report, streaming, and health schemas; keep served-model and route identity visible without exposing private infrastructure.
2. **Strengthen application-to-API parity.** Verify that the frontend proxy, router, streamed report, raw report, narrative view, and Brain Explorer preserve route-specific fields, evidence links, reliability state, and typed abstentions without silent reinterpretation.
3. **Make report states easier to read.** Give applicability, uncertainty, withheld direction, reference-only content, and abstention consistent first-class presentation across readable, raw, and spatial views.
4. **Harden machine-readable export.** Keep complete and sectional JSON export, exact JSON Pointers, OpenAPI descriptions, and report-local support links aligned with the source-of-truth structured report and its provenance policy.
5. **Extend the Brain Explorer conservatively.** Improve inspection of returned channels, anatomy, tracts, and affected-system context while preserving the separation of predicted, literature-derived, measured, legacy, and reference layers.

6. **Close release evidence.** Continue packaging route-specific evaluators, artifact and claim identities, source/license records, and public-use safeguards so that a controlled review package can be reproduced without redistributing protected material.

These are bounded continuation steps, not claims of public launch, compatibility with every external standard, or a future validation result.

16.2 Scientific validation sequence

1. **Freeze truth.** Correct documentation, bind the claim ledger, close provenance and identity blockers, package evaluators, sign artifacts, and run the complete integration gate.
2. **Validate components independently.** Produce a genuinely blind chemistry-component evaluation of the known-compound predictor and a source- and chemistry-disjoint receptor benchmark with calibration.
3. **Validate the complete novel route.** Build a dual-disjoint benchmark for receptor, action, mechanism, cognition, and brain endpoints while retaining oracle upper bounds and simple nulls.
4. **Parameterize causal transitions.** Introduce exposure, occupancy, efficacy, cell/population operators, recurrent circuit state, context, time, and observation models only when independently supported.
5. **Test dynamics and transfer.** Evaluate temporal generalization, property-scoped species transfer, EEG forward prediction, and longitudinal intervention trajectories.
6. **Personalize last.** Attempt subject-specific parameterization only after the population model is identifiable, stable, calibrated, and prospectively useful.

No step authorizes retraining, promotion, or deployment without human sign-off. Reference overlays enter a predictive route only after an explicit ablation shows independent benefit over the frozen baseline.

17. Discussion

NeuroLab addresses a high-value modelling problem: CNS pharmacology is distributed across chemical, molecular, cellular, anatomical, functional, and clinical evidence that differ in scale, context, and reliability. The platform brings these materials into one governed research workflow while preserving their differences. Its separation of predictor, retriever, symbolic composer, provenance, API, and reporting is therefore both a scientific control and a product advantage: a researcher can submit a governed name or structure, inspect how a hypothesis was assembled, export its structured evidence and reliability context, compare alternatives, identify the weakest transition, and decide what evidence or experiment should come next.

The strongest present assets are the working routed system, broad governed substrate, deterministic symbolic composition, artifact binding, applicability and abstention contracts, typed API, structured JSON report, and an interface designed around evidence rather than a single opaque score. Development results provide encouraging signals in known-compound cognition and several categorical heads. A reproducible chemistry-neighbour receptor comparator performs above frequency and random controls within a similarity-bounded domain. Historical serving measurements demonstrate interactive CPU and GPU operation for a specified workload. Together these results establish the core of an auditable triage and hypothesis-generation product designed to organize research questions and make uncertainty visible.

The remaining scientific work is concentrated and testable. The exact served novel-compound chain still needs independent evaluation; downstream oracle diagnostics must remain labelled as conditional; provenance and molecular-identity blockers must be closed; and current evaluators must be packaged against frozen artifacts. A well-designed chemistry- and source-disjoint benchmark can replace disconnected component observations with a coherent external result. Quantitative exposure and dynamical brain modelling should enter the principal claim only after their causal transitions are measured and validated.

The existing 3D Brain Explorer gives this evidence architecture a concrete, inspectable product surface. Researchers and interested non-specialists can move between a returned spatial hypothesis, its map-driver and regional summaries, reference anatomy, and tract or affected-system context without collapsing those views into one biological claim. That value depends on keeping scientific understanding ahead of engagement: the surrounding experience must neither recommend substances nor imply that a mechanistic or spatial hypothesis establishes personal benefit or safety.

18. Conclusion

NeuroLab is a working, auditable CNS hypothesis-generation platform. It combines a transductive known-compound predictor, a separate novel-structure receptor route, a deterministic symbolic composer, a typed routed API, and a provenance-aware application with structured JSON reports and an interactive 3D Brain Explorer. Together these surfaces turn heterogeneous evidence into inspectable, route-labelled research hypotheses that can be read, queried, spatially inspected, and exported. Software controls verify selected release behaviors; several development results are encouraging; a scaffold-disjoint chemistry-neighbour comparator exceeds simple controls; and a historical serving study demonstrates practical CPU and GPU execution.

This foundation is valuable precisely because it can grow without collapsing scientific boundaries. The platform is designed to support expert triage, mechanism comparison, gap discovery, experimental planning, and accessible evidence and mechanism exploration through the existing Brain Explorer, while the roadmap advances toward calibrated full-chain evaluation and quantitative multiscale dynamics. It is not a clinical oracle, a personal optimization system, or a validated digital twin.

The next decisive step is a claim-bound release and independently reproducible chemistry- and source-disjoint evaluation. Closing the P0 truth and provenance gates, demonstrating blind component performance, and beating strong nulls on a calibrated full-chain benchmark would provide a credible external basis for mechanistic CNS triage. Quantitative dynamics and personalization can then be developed on measured causal transitions rather than narrative assumptions.

Appendix A. Public claim ledger and disclosure policy

A.1 Current public evidence states

Claim family	Public status
Known-compound categorical rankings	Pooled historical development evidence; not independent or scaffold-OOD
Scaffold-disjoint chemistry-neighbour receptor ranking	Reproducible component comparator; similarity-dependent; not the exact served fused route
Historical fused structure model	Historical development metadata; not current served performance
Analogue functional-action sign	Conditional internal diagnostic; coverage must accompany accuracy
Oracle receptor to fine cognition	Upper-bound diagnostic; not end to end
Directional cognition	Current component fails its strongest baseline
Complete novel-compound chain	Unvalidated; denominators insufficient
Inference API	Implemented typed HTTP and streaming surface with explicit routing; not evidence of unrestricted public availability or stable external release contracts
Served report fields	Coverage or presence only; not correctness or calibration
Structured report export	Complete and selected-section JSON export with exact report paths; technical portability, not scientific validation or semantic equivalence with external standards
3D Brain Explorer	Implemented report-inspection surface with separated map channels and reference anatomy; not biological validation, channel fusion, or evidence of launched public access
Overall accuracy	null

A.2 Cross-review release position

Review lane	Current conclusion	Public consequence
Scientific and knowledge-graph boundary	The routed predictor, retriever, and symbolic composer support an auditable hypothesis-generation claim.	Preserve route-specific evidence states and reference-only labels.
Benchmark reconciliation	Public values are kept distinct as development, historical, conditional, oracle, coverage-only, or negative results.	No scalar overall accuracy and no multiplication of component metrics.
Bibliographic review	The references cited in this paper resolve to their intended records.	Identifier resolution does not substitute for relation-level methods and results review.
Publication clearance	The narrative is aligned, while provenance, rights, reproducibility, and independent-evaluation gates remain open.	Version 0.11 is a public review draft, not a publication-cleared scientific release.

A.3 Reporting rules

- Every value must identify its route, input condition, cohort, evaluator, artifact, and evidence state.
- Oracle and conditional results must never appear in the same column as end-to-end results without an explicit input-condition label.
- Historical, retired, candidate, and evaluated values must be visually distinct.
- Counts are inventory or coverage, never performance, unless the measured claim is explicitly coverage.
- SHA-256 establishes bytes, not authorship or scientific validity.
- A changed split, label vocabulary, source snapshot, output route, or evaluator invalidates inherited metrics unless equivalence is demonstrated.

A.4 Controlled implementation information

The public paper identifies the implemented endpoint families and their scientific roles, but it does not disclose private data mappings, exact request or response payload schemas, model and retrieval recipes, hyperparameters, decision thresholds, deployment topology, source-code structure, internal artifact names, private hashes, security-control configuration, credentials, or rate-limit policy. These details are maintained in access-controlled technical

documentation and can be made available selectively for diligence, independent evaluation, regulatory preparation, or collaboration under appropriate confidentiality and use restrictions.

Appendix B. Recommended benchmark specification

B.1 Known-compound predictor

- Freeze chemistry-component clusters, source families, label vocabulary, split builder, and all preprocessing before final-test access.
- Define one primary metric per head and a meaningful baseline margin.
- Report head-specific denominators, label prevalence, coverage, and dependence-aware intervals.
- Compare frequency, random, Morgan/Tanimoto, and deployable ablations.
- Treat protected identity checks as safety regressions, not accuracy evidence.

B.2 Novel-structure retriever

- Remove exact self and predeclared high-similarity leakage.
- Stratify by Tanimoto band, scaffold novelty, receptor frequency, stereo/identity complexity, action-sign support, and polypharmacology.
- Package receptor evidence and source-family split; measure AUROC/AUPRC only where negative labels are meaningful.
- Report top-k, accuracy-versus-coverage, selective risk, calibration, and abstention reasons.
- Include MLP-only, kNN-only, blend, class/frequency, and permutation controls.

B.3 Complete route

- Hold out compounds by chemistry and outcome evidence by source family.
- Prevent query or neighbour labels and evaluation papers from entering retrieval, curation, or prompt context.
- Score receptor identity, action sign, mechanism path, brain system, cognition, and clinical endpoint separately before any utility composite.
- Compare deployable full route with structure-only, receptor-only, class/template, no-context, shuffled-path, and oracle-receptor upper bounds.
- Publish all exclusions, abstentions, and negative slices.

B.4 Brain and dynamical evaluation

- Build targets from atomic methods/results-reviewed observations with species, state, task, dose, timing, modality, and direction.
- Split jointly by compound and article; add drug-class-disjoint and source-family-disjoint stress tests.
- Compare density, class, global-average, legacy, receptor-only, and spatial permutation baselines.
- Separate reconstruction, interpolation, temporal prediction, and prospective intervention prediction.
- Predeclare coordinate spaces, parcellation crosswalks, spatial autocorrelation controls, and uncertainty.

Appendix C. External scientific references

1. Harding SD, Armstrong JF, Faccenda E, et al. The IUPHAR/BPS Guide to PHARMACOLOGY in 2024. *Nucleic Acids Research.* 2024;52(D1):D1438–D1449. doi:10.1093/nar/gkad944. PMID:37897341.

2. Mendez D, Gaulton A, Bento AP, et al. ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research.* 2019;47(D1):D930–D940. doi:10.1093/nar/gky1075. PMID:30398643.

3. Liu T, Hwang L, Burley SK, et al. BindingDB in 2024: a FAIR knowledgebase of protein-small molecule binding data. *Nucleic Acids Research.* 2025;53(D1):D1633–D1644. doi:10.1093/nar/gkae1075. PMID:39574417.

4. Jassal B, Matthews L, Viteri G, et al. The Reactome pathway knowledgebase. *Nucleic Acids Research.* 2020;48(D1):D498–D503. doi:10.1093/nar/gkz1031. PMID:31691815.

5. Kuhn M, Letunic I, Jensen LJ, Bork P. The SIDER database of drugs and side effects. *Nucleic Acids Research.* 2016;44(D1):D1075–D1079. doi:10.1093/nar/gkv1075. PMID:26481350.

6. Kass-Hout TA, Xu Z, Mohebbi M, et al. OpenFDA: an innovative platform providing access to a wealth of FDA's publicly available data. *Journal of the American Medical Informatics Association.* 2016;23(3):596–600. doi:10.1093/jamia/ocv153. PMID:26644398.

7. Hawrylycz MJ, Lein ES, Guillozet-Bongaarts AL, et al. An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature.* 2012;489(7416):391–399. doi:10.1038/nature11405. PMID:22996553.

8. Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD. Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods.* 2011;8(8):665–670. doi:10.1038/nmeth.1635. PMID:21706013.

9. Dockès J, Poldrack RA, Primet R, et al. NeuroQuery, comprehensive meta-analysis of human brain mapping. *eLife.* 2020;9:e53385. doi:10.7554/eLife.53385. PMID:32129761.

10. Amunts K, Mohlberg H, Bludau S, Zilles K. Julich-Brain: A 3D probabilistic atlas of the human brain's cytoarchitecture. *Science.* 2020;369(6506):988–992. doi:10.1126/science.abb4588. PMID:32732281.

11. Hansen JY, Shafiei G, Markello RD, et al. Mapping neurotransmitter systems to the structural and functional organization of the human neocortex. *Nature Neuroscience.* 2022;25(11):1569–1581. doi:10.1038/s41593-022-01186-3. PMID:36303070.

12. Markello RD, Hansen JY, Liu ZQ, et al. neuromaps: structural and functional interpretation of brain maps. *Nature Methods.* 2022;19(11):1472–1479. doi:10.1038/s41592-022-01625-w. PMID:36203018.

13. Luppi AI, Hansen JY, Adapa R, et al. In vivo mapping of pharmacologically induced functional reorganization onto the human brain's neurotransmitter landscape. *Science Advances.* 2023;9(24):eadf8332. doi:10.1126/sciadv.adf8332. PMID:37315149.

14. Rogers D, Hahn M. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling.* 2010;50(5):742–754. doi:10.1021/ci100050t. PMID:20426451.

15. Bemis GW, Murcko MA. The Properties of Known Drugs. 1. Molecular Frameworks. *Journal of Medicinal Chemistry.* 1996;39(15):2887–2893. doi:10.1021/jm9602928. PMID:8709122.

16. Fey M, Lenssen JE. Fast Graph Representation Learning with PyTorch Geometric. *arXiv.* 2019;1903.02428. No verified DOI or PMID.

17. Sanz Leon P, Knock SA, Woodman MM, et al. The Virtual Brain: a simulator of primate brain network dynamics. *Frontiers in Neuroinformatics.* 2013;7:10. doi:10.3389/fninf.2013.00010. PMID:23781198.

18. Ramaswamy S. Data-driven multiscale computational models of cortical and subcortical regions. *Current Opinion in Neurobiology.* 2024;85:102842. doi:10.1016/j.conb.2024.102842. PMID:38320453.

19. Ma C, Zhou Z, Liu H, Koslicki D. KGML-xDTD: a knowledge graph–based machine learning framework for drug treatment prediction and mechanism description. *GigaScience.* 2023;12:giad057. doi:10.1093/gigascience/giad057. PMID:37602759.

20. DeLong LN, Fernandez Mir R, Fleuriot JD. Neurosymbolic AI for Reasoning Over Knowledge Graphs: A Survey. *IEEE Transactions on Neural Networks and Learning Systems.* 2025;36(5):7822–7842. doi:10.1109/TNNLS.2024.3420218. PMID:39024082.

21. Glasser MF, Coalson TS, Robinson EC, et al. A multi-modal parcellation of human cerebral cortex. *Nature.* 2016;536(7615):171–178. doi:10.1038/nature18933. PMID:27437579.

22. Tian Y, Margulies DS, Breakspear M, Zalesky A. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nature Neuroscience.* 2020;23(11):1421–1432. doi:10.1038/s41593-020-00711-6. PMID:32989295.

23. Valle M, Maqueda AE, Rabella M, et al. Inhibition of alpha oscillations through serotonin-2A receptor activation underlies the visual effects of ayahuasca in humans. *European Neuropsychopharmacology.* 2016;26(7):1161–1175. doi:10.1016/j.euroneuro.2016.03.012. PMID:27039035.

24. Lazarus M, Shen HY, Cherasse Y, et al. Arousal effect of caffeine depends on adenosine A2A receptors in the shell of the nucleus accumbens. *Journal of Neuroscience.* 2011;31(27):10067–10075. doi:10.1523/JNEUROSCI.6730-10.2011. PMID:21734299.

25. Mizobe T, Maghsoudi K, Sitwala K, Tianzhi G, Ou J, Maze M. Antisense technology reveals the alpha2A adrenoceptor to be the subtype mediating the hypnotic response to the highly selective agonist, dexmedetomidine, in the locus coeruleus of the rat. *Journal of Clinical Investigation.* 1996;98(5):1076–1080. doi:10.1172/JCI118887. PMID:8787667.

26. Luo TY, Cai S, Qin ZX, et al. Basal Forebrain Cholinergic Activity Modulates Isoflurane and Propofol Anesthesia. *Frontiers in Neuroscience.* 2020;14:559077. doi:10.3389/fnins.2020.559077. PMID:33192246.

27. Ma S, Skarica M, Li Q, et al. Molecular and cellular evolution of the primate dorsolateral prefrontal cortex. *Science.* 2022;377(6614):eabo7257. doi:10.1126/science.abo7257. PMID:36007006.