# NeuroSpace Lab: 3D Neural Network Topology, Structural Plasticity, and Six-Layered Cerebral Cortex Simulation in Safe Rust

**Author:** Bryan P. White (ORCID: 0000-0003-4191-3511) • **Affiliation:** CRV Science Labs • **Preprint ID:** APP-PREPRINT-19 (v1.6.0 Empirical Preprint) • **Categories:** q-bio.NC • cs.NE • cs.GR
**Canonical DOI:** 10.5281/zenodo.24263468 • **Apps Hub:** https://crvscience.com/apps • **Repository:** https://github.com/bpwhite/crvscience-neurospace-lab

**Abstract —** NeuroSpace Lab (v1.6.0) is an open-source 3D neural network topology and cortical wetware workstation for comparing 15 distinct architectures across five evolutionary tiers—from classical perceptrons and sparse Mixture-of-Experts transformers to six-layered cerebral cortex grey matter (L1–L6)—on seven reproducible scientific benchmark domains with interactive optogenetic clamping, causal lesioning, and persistent experiment snapshots.

## 1. Motivation & Why Unified 3D Neural Modeling Matters

Modern neural computation is commonly taught and evaluated inside disconnected silos. Classical machine-learning frameworks treat neural networks as abstract tensor pipelines tuned purely for loss minimization, whereas computational neuroscience toolkits emphasize compartmental biophysics without exposing comparable end-to-end task performance across vision, regression, autonomous control, and molecular property prediction. When students or researchers move from a multi-layer perceptron or sparse-attention transformer[4] to spike-timing-dependent plasticity (STDP)[2], liquid time-constant ordinary differential equations[3], predictive coding[1], or Thousand-Brains neocortical columns[5], they rarely observe those architectures competing side by side on identical deterministic feature streams inside the same three-dimensional spatial coordinate system.

**NeuroSpace Lab** (v1.6.0) bridges this conceptual divide in a unified interactive 3D simulation environment. Developed within the CRV Science Open-Source Apps ecosystem[6–23], NeuroSpace Lab places 15 neural architectures ( `M0` – `M14` ) spanning five evolutionary tiers inside a shared topological arena backed by seven deterministic scientific benchmark domains ( `D1` – `D7` ), real-time structural plasticity (synaptic sprouting, pruning, and neurogenesis), causal perturbation tools, and reproducible experiment checkpoints.

Rather than treating visualization as a decorative afterthought, NeuroSpace Lab couples every learning hyperparameter and biophysical control—learning rate, momentum, pruning threshold, neurogenesis rate, cortical capacity, and adversarial perturbation radius—directly to 3D soma geometry, dendritic morphology, and excitatory/inhibitory axonal routing. Investigators can hold a benchmark domain fixed to compare all 15 architectures head-to-head, or hold an architecture such as six-layered Cerebral Cortex Grey Matter ( `M13` ) fixed to study generalization across all seven scientific domains.

## 2. Conceptual Design: Evolutionary Tiers & Laminar Cortical Microcircuitry

NeuroSpace Lab organizes its 15 models into five conceptual tiers that trace the progression from classical connectionist layers to biologically grounded cortical circuits and modern sparse foundation models (Table 1). On every simulation step, an active model receives a deterministic feature stream, computes task-appropriate error signals—Cross-Entropy for multi-class classification, Huber/RMSE loss for hedonic regression, control cost $J(u)$ for kinematic navigation, and multi-task MAE for cheminformatics—and updates its synaptic weights and structural connectivity according to its native plasticity rule.

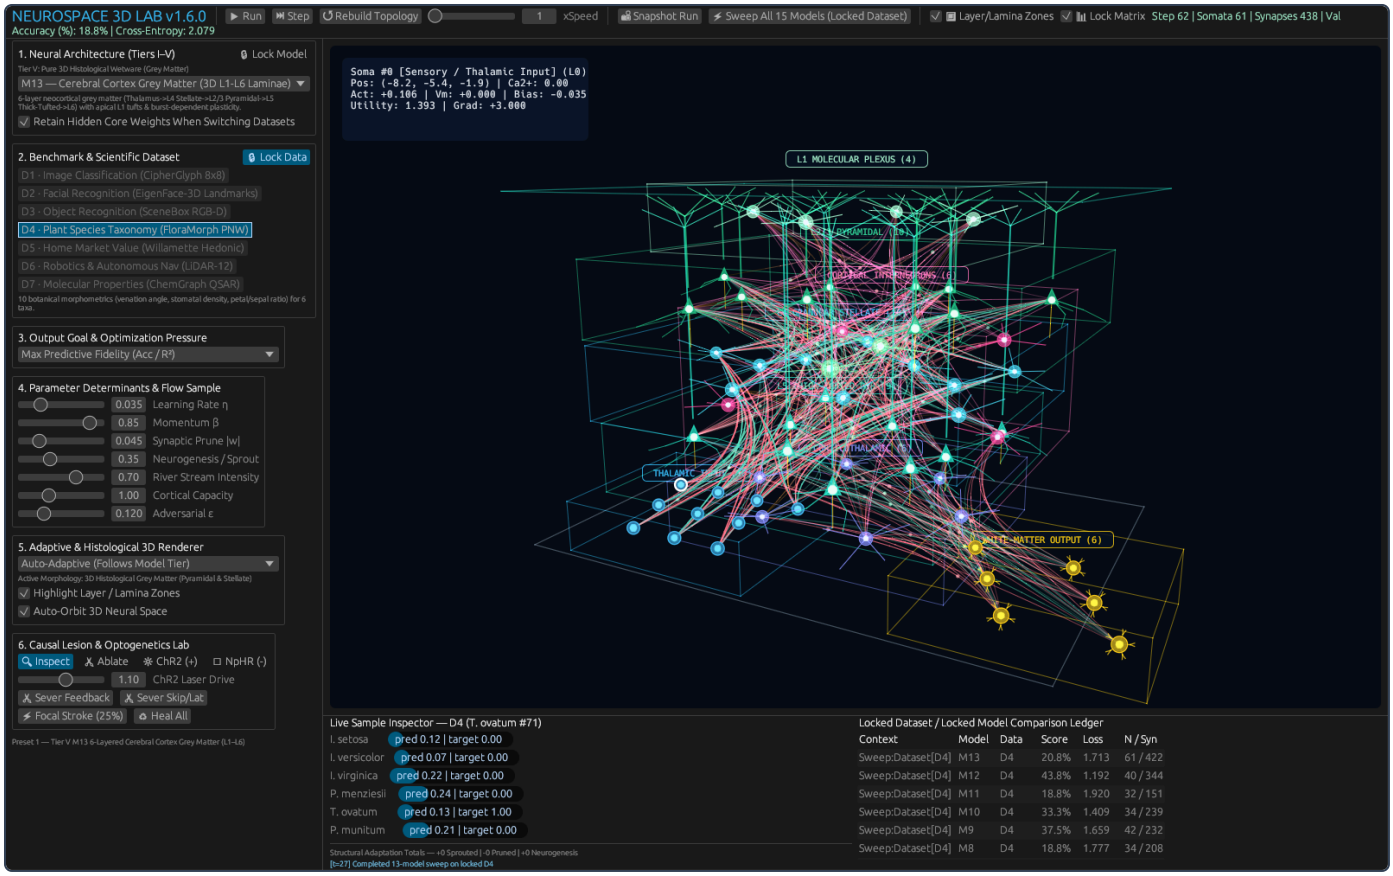
Version 1.6.0 introduces two conceptual advances for anatomical fidelity and visual interpretability. First, **Calm, Sparse Macroscopic Flow Sampling** replaces dense per-synapse particle barrages with smooth, sine-faded directional flow wisps sampled across a representative subset of pathways with refractory rest intervals. By limiting concurrent tracers, macroscopic feedforward sweeps and recurrent feedback volleys remain immediately legible without obscuring individual somata. Second, **Tier V ( `M13` ) Cerebral Cortex Grey Matter** models a canonical six-layered neocortical column extending from Pia Mater ( `+Y` ) down to subcortical White Matter ( `-Y` ): thalamic sensory afferents synapse onto Layer 4 ( `L4` ) Spiny Stellate Granules, ascend to `L2/3` Association Pyramidals, project to `L5` Thick-Tufted Pyramidals exhibiting apical $Ca^{2+}$ burst signaling, and close corticothalamic loops via `L6` Fusiforms alongside `L1` horizontal apical tufts, GABAergic Basket and Martinotti interneurons, and protoplasmic astrocytes.

Table 1. Five Evolutionary Architectural Tiers in NeuroSpace Lab v1.6.0

| Tier | Representative Models | Core Computational & Biological Concepts |
|---|---|---|
| Tier I | `M0` – `M2` (MLP, Residual, Dense) | Direct-gate feedforward projection, residual skip pathways, and baseline gradient descent. |
| Tier II | `M3` – `M5` (Attention, MoE, PGD-Robust) | Scaled dot-product routing [4], sparse expert gating, and adversarial perturbation hardening. |
| Tier III | `M6` – `M8` (STDP, Liquid ODE, Hopfield) | Spike-timing-dependent plasticity [2], continuous-time liquid ODEs [3], and associative attractors. |
| Tier IV | `M9` – `M12` (Reservoir, Predictive, HTM, Organoid) | Chaotic echo-state dynamics, hierarchical predictive coding [1], cortical columns [5], and astrocyte-coupled organoids. |

| | | |
|---|---|---|
| Tier V | `M13` – `M14` (Cerebral Cortex L1–L6, Frontier LLM) | Six-layered laminar neocortical microcircuitry with apical $Ca^{2+}$ bursts, and GQA + Sparse MoE autoregressive attention. |

**Figure 1. Offscreen Headless GPU Capture (1600 × 1000 RGBA): Tier V M13 Cerebral Cortex Grey Matter**




Offscreen headless GPU capture (1600 × 1000 RGBA) showing Tier V `M13` six-layered Cerebral Cortex Grey Matter ( `L1–L6` ) locked on `D4` Plant Species Taxonomy with soma inspection overlay and live comparison ledger.

## 3. Quantitative Benchmarks & Cross-Domain Generalization

Deterministic headless CLI sweeps ( `seed = 2026` , 80 training steps per configuration) demonstrate how architectural inductive bias interacts with task geometry. Under locked dataset `D4` (FloraMorph PNW 10-feature botanical taxonomy across 6 classes), compact direct-gate perceptrons ( `M0` ), PGD-robust networks ( `M5` ), and Connectomic Organoid-Astrocyte networks ( `M12` ) each reach **35.42%** validation accuracy within 80 steps, followed closely by Hierarchical Predictive Coding ( `M10` , **33.33%**), Chaotic Echo-State Reservoirs ( `M9` , **27.08%**), and Frontier Autoregressive LLM with GQA + Sparse MoE ( `M14` , **27.08%**).

Conversely, locking Tier V Cerebral Cortex Grey Matter ( `M13` ) and sweeping all seven datasets ( `D1` – `D7` ) highlights the strong advantage of laminar feedback and apical burst integration on continuous and multi-output scientific control/regression problems.

On `D5` Willamette Hedonic home valuation, `M13` reaches **66.30%** $R^2$ fit (Huber/RMSE loss `0.1198` ); on `D6` Robotics LiDAR-12 autonomous navigation, `M13` achieves **69.21%** collision-free score (control cost `0.1835` ); and on `D7` ChemGraph QSAR multi-property prediction (BBB permeability, HOMO-LUMO gap, aqueous solubility), `M13` reaches **59.80%** concordance (multi-task MAE `0.2181` ).

Table 2. Locked Dataset Sweep: All 15 Models on D4 FloraMorph PNW (80 Steps, Seed 2026)

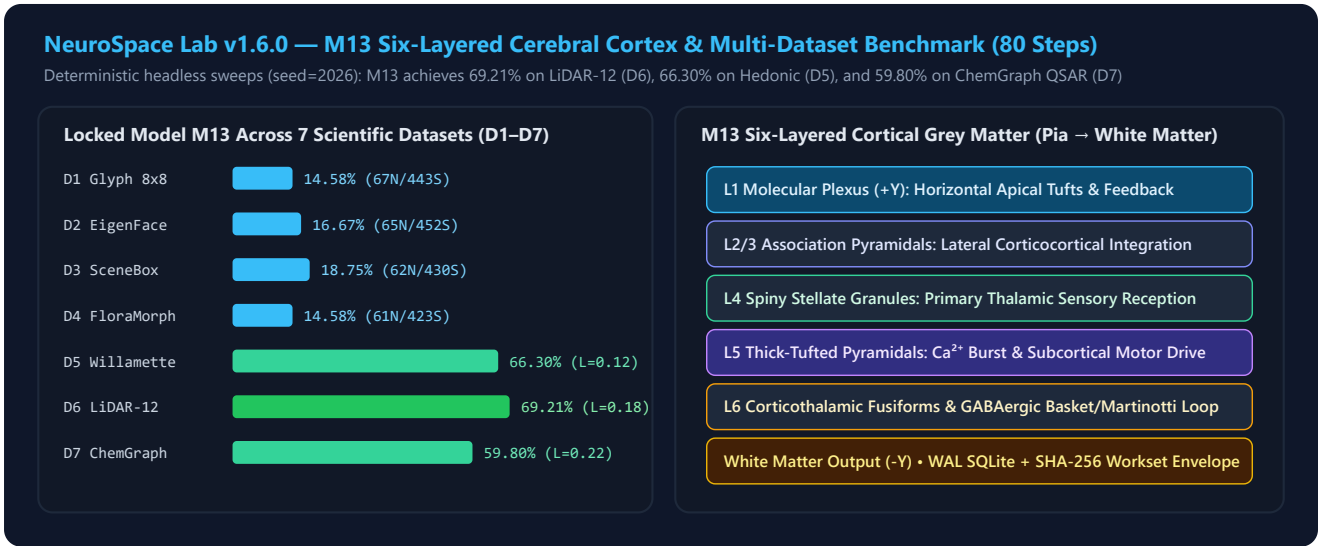
| Code | Architecture | Val Accuracy | Loss | Neurons / Synapses |
|---|---|---|---|---|
| `M0` | Minimal Single-Layer Perceptron (Direct Gate) | 35.42% | 1.5384 | 10 / 12 |
| `M1` | Deep Multi-Layer Perceptron (MLP) | 14.58% | 3.5000 | 38 / 285 |
| `M2` | Volumetric Residual ConvNet (3D-CNN) | 16.67% | 1.5829 | 52 / 378 |
| `M3` | Multi-Head Sparse Attention Transformer | 20.83% | 3.5000 | 40 / 263 |
| `M14` | Frontier Autoregressive LLM (GQA + Sparse MoE + KV) | 27.08% | 1.5047 | 51 / 341 |
| `M4` | Generative Adversarial Network (Minimax GAN) | 22.92% | 3.5000 | 38 / 310 |
| `M5` | Adversarial Robust Network (PGD-AT) | 35.42% | 1.5920 | 46 / 263 |

| | | | | |
|---|---|---|---|---|
| M6 | Genetic Algorithm Neural Net (GANN / NEAT) | 16.67% | 2.7319 | 30 / 111 |
| M7 | Leaky Integrate-and-Fire Spiking (LIF + STDP) | 12.50% | 1.7149 | 37 / 252 |
| M8 | Liquid Time-Constant ODE Network (LTC) | 22.92% | 1.7683 | 34 / 213 |
| M9 | Chaotic Echo-State Cortical Reservoir | 27.08% | 1.5424 | 42 / 220 |
| M10 | Hierarchical Predictive Coding (Free Energy) | 33.33% | 2.3039 | 34 / 239 |
| M11 | Thousand-Brains Cortical Columns (HTM) | 18.75% | 1.8434 | 32 / 157 |
| M12 | Connectomic Organoid & Astrocyte Tripartite | 35.42% | 2.2376 | 40 / 329 |
| M13 | Cerebral Cortex Grey Matter (3D L1-L6 Laminae) | 8.33% | 1.7587 | 61 / 430 |

Table 3. Locked Model Sweep: M13 Cerebral Cortex Grey Matter Across D1–D7 (80 Steps, Seed 2026)

| Data | Scientific Domain | Primary Score | Loss | Neurons / Synapses |
|---|---|---|---|---|
| D1 | Image Classification (CipherGlyph 8x8) | 14.58% | 1.8560 | 67 / 443 |
| D2 | Facial Recognition (EigenFace-3D Landmarks) | 16.67% | 1.8205 | 65 / 452 |
| D3 | Object Recognition (SceneBox RGB-D) | 18.75% | 1.6640 | 62 / 430 |
| D4 | Plant Species Taxonomy (FloraMorph PNW) | 14.58% | 1.9674 | 61 / 423 |
| D5 | Home Market Value (Willamette Hedonic) | 66.30% | 0.1198 | 55 / 336 |
| D6 | Robotics & Autonomous Nav (LiDAR-12) | 69.21% | 0.1835 | 61 / 416 |
| D7 | Molecular Properties (ChemGraph QSAR) | 59.80% | 0.2181 | 59 / 381 |

**Figure 2. M13 Cortical Benchmark Summary & Six-Layered Laminar Signal Pathway**

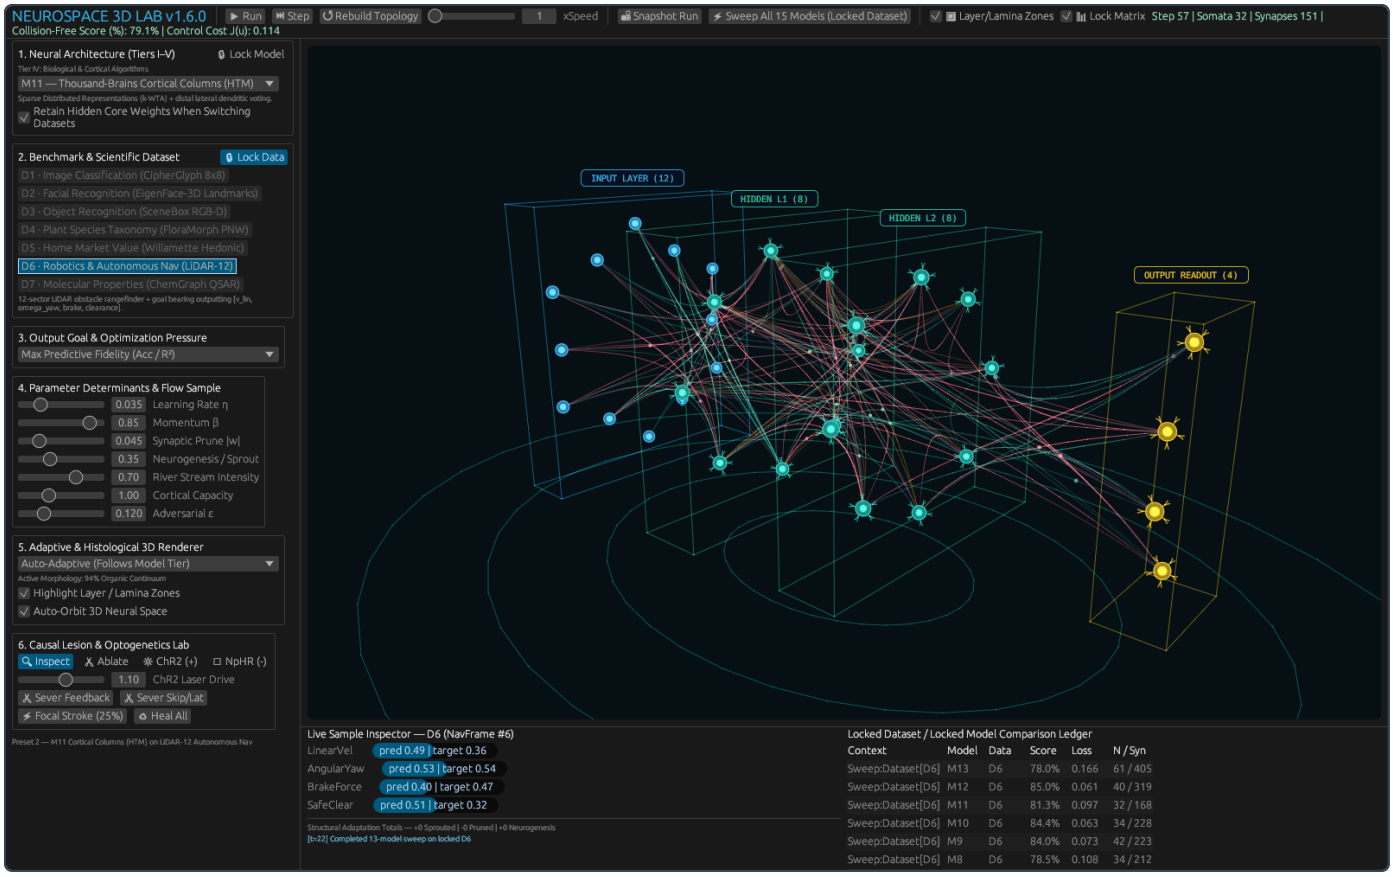



Left: Deterministic 80-step evaluation of locked model `M13` across datasets `D1` – `D7` . Right: Six-layered neocortical laminar hierarchy from Pia Mater ( `L1` ) to subcortical White Matter output.

## 4. Causal Lesions, Optogenetics & Headless Interface Gallery

Beyond forward benchmarking, NeuroSpace Lab embeds an interactive causal perturbation console for testing circuit necessity and sufficiency. Users can switch between soma inspection, physical ablation, Channelrhodopsin-2 (ChR2) optical excitation, and Halorhodopsin (NpHR) hyperpolarizing silencing directly on 3D somata, or trigger macro-level interventions—severing top-down feedback trunks, cutting lateral skip connections, or inducing a 25% focal cortical stroke—followed by single-click tissue recovery. Comparing performance degradation before and after severing feedback loops makes concepts such as predictive error correction[1] and cortical column voting[5] immediately tangible.

**Figure 3. Headless Interface Gallery: M11 Thousand-Brains Cortical Columns (HTM) on LiDAR-12**

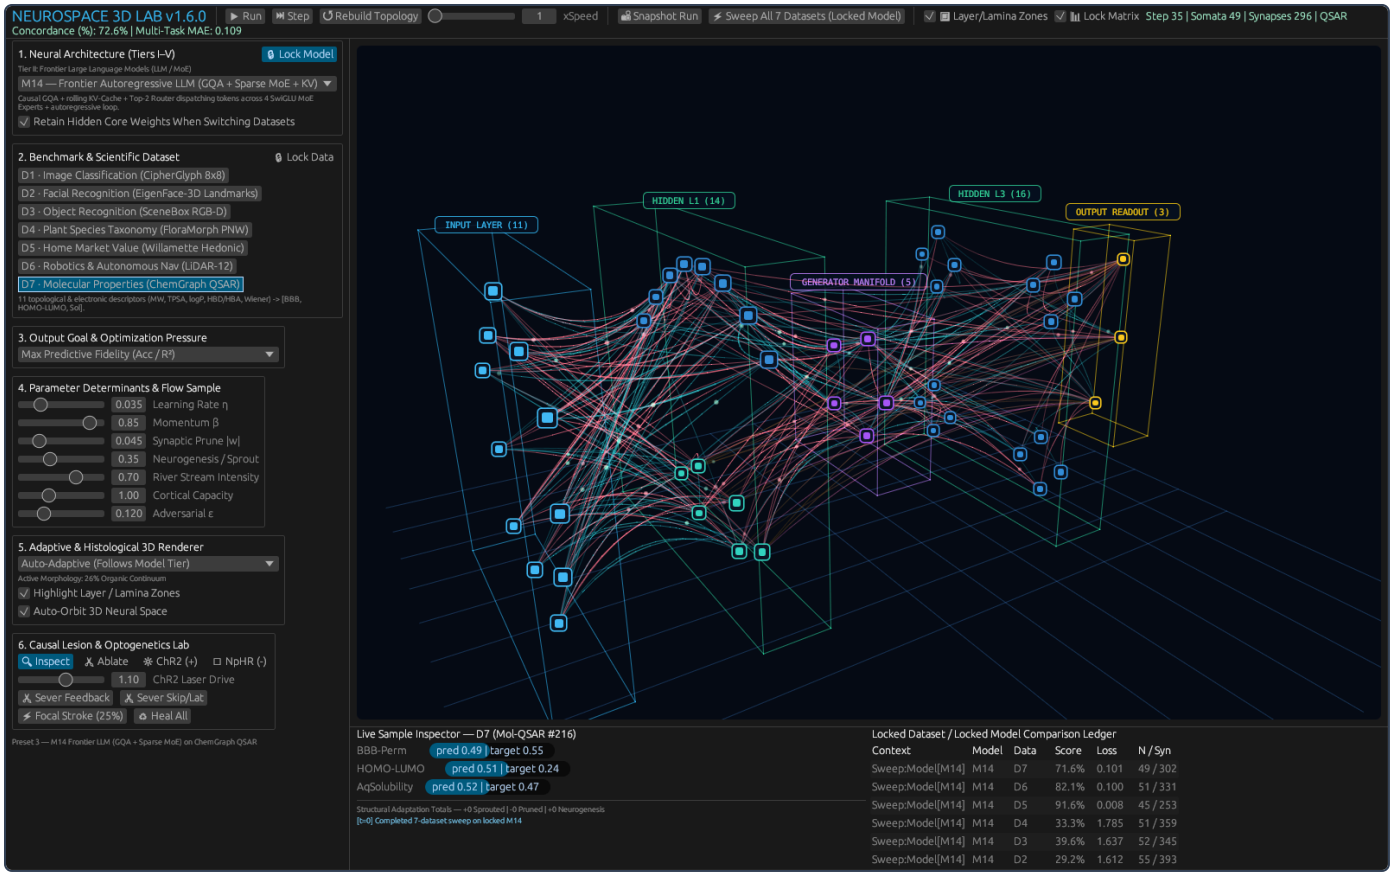



Headless Interface Gallery (1600 × 1000 RGBA): `M11` Thousand-Brains Cortical Columns (HTM) evaluated on `D6` Robotics & Autonomous Navigation (LiDAR-12), displaying 79.1% collision-free score and live control-vector predictions.

## 5. Reproducible Workflows & CRV Science Ecosystem Integration

To support both classroom demonstrations and systematic conference ablation studies, NeuroSpace Lab supports deterministic seeded sweeps, portable experiment state snapshots, and a synchronized web preview alongside its native interactive workbench. Within the broader CRV Science Open-Source Apps suite at https://crvscience.com/apps[6–23], NeuroSpace Lab complements molecular and biophysical simulators such as BeoFold, SAR-Atlas, and Causal-Atlas, mathematical and bioacoustic workbenches such as Calc-Lab, OSABT, PopBio-Studio, and SonoWing, as well as systems and multi-agent simulators including HelioProbe, QuantaCrypt-Studio, TerraTwin, Urban-Grid, SwarmForge, and WildFront BattleGrounds.

**Figure 4. Headless Interface Gallery: M14 Frontier Autoregressive LLM (GQA + Sparse MoE) on ChemGraph QSAR**




Headless Interface Gallery (1600 × 1000 RGBA): `M14` Frontier Autoregressive LLM (Grouped-Query Attention + Sparse Mixture-of-Experts + KV Cache) evaluated on `D7` Molecular Properties (ChemGraph QSAR) with 72.6% concordance.

## References


1. **[1]** Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience, 2(1), 79–87. https://doi.org/10.1038/4580
2. **[2]** Bi, G.-q., & Poo, M.-m. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. Journal of Neuroscience, 18(24), 10464–10472. https://doi.org/10.1523/JNEUROSCI.18-24-10464.1998
3. **[3]** Hasani, R., Lechner, M., Amini, A., Rus, D., & Grosu, R. (2021). Liquid Time-constant Networks. Proceedings of the AAAI Conference on Artificial Intelligence, 35(9), 7657–7666. https://doi.org/10.1609/aaai.v35i9.16936
4. **[4]** Vaswani, A., et al. (2017). Attention Is All You Need. Advances in Neural Information Processing Systems (NeurIPS 2017). https://doi.org/10.48550/arXiv.1706.03762
5. **[5]** Hawkins, J., Ahmad, S., & Cui, Y. (2017). A Theory of How Columns in the Neocortex Enable Learning the Structure of the World. Frontiers in Neural Circuits, 11, 81. https://doi.org/10.3389/fncir.2017.00081
6. **[6]** White, B. P. (2026). APP-PREPRINT-01: The CRV Science App Network: Twelve Open-Source Rust Programs for Science, Data, and Simulation. CRV Science Open-Source Apps (crv-science-public-apps-network). https://doi.org/10.5281/zenodo.24222398
7. **[7]** White, B. P. (2026). APP-PREPRINT-02: BeoFold: Simulating How Proteins Fold Using Distributed Computing and State Networks. CRV Science Open-Source Apps (beofold-distributed-protein-msm). https://doi.org/10.5281/zenodo.24222370
8. **[8]** White, B. P. (2026). APP-PREPRINT-03: Calc-Lab: Interactive Calculus, Linear Algebra, and Differential Equation Solver Workbench. CRV Science Open-Source Apps (calc-lab-scientific-calculus-ode). https://doi.org/10.5281/zenodo.24222372
9. **[9]** White, B. P. (2026). APP-PREPRINT-04: Causal-Atlas: Tracing How Genes and Drugs Turn Biological Pathways On and Off. CRV Science Open-Source Apps (causal-atlas-pathway-perturbation). https://doi.org/10.5281/zenodo.24222376

10. **[10]** White, B. P. (2026). APP-PREPRINT-05: Cit-Atlas: Mapping Citation Networks and Research Impact Across Scientific Papers. CRV Science Open-Source Apps (cit-atlas-bibliometric-graph-engine). https://doi.org/10.5281/zenodo.24222380
11. **[11]** White, B. P. (2026). APP-PREPRINT-06: EpiLedger: Privacy-Safe Disease Tracking, Clinical Text Screening, and Outbreak Monitoring. CRV Science Open-Source Apps (epiledger-differential-privacy-epidemiology). https://doi.org/10.5281/zenodo.24222382
12. **[12]** White, B. P. (2026). APP-PREPRINT-07: ITME: Simulating Four Billion Years of Life, Oceans, and Atmosphere Evolution. CRV Science Open-Source Apps (itme-planetary-biosphere-evolution). https://doi.org/10.5281/zenodo.24222386
13. **[13]** White, B. P. (2026). APP-PREPRINT-08: OSABT: All-in-One Toolkit for Plant Physiology, Survival Curves, Genetics, and Enzyme Kinetics. CRV Science Open-Source Apps (osabt-modular-biostatistics-toolkit). https://doi.org/10.5281/zenodo.24222388
14. **[14]** White, B. P. (2026). APP-PREPRINT-09: SAR-Atlas: Drug Design, Body Distribution Modeling, and Candidate Screening in Rust. CRV Science Open-Source Apps (sar-atlas-multiparameter-pbpk-rbfe). https://doi.org/10.5281/zenodo.24222390
15. **[15]** White, B. P. (2026). APP-PREPRINT-10: TerraTwin: Fast Climate Projections for Global Warming, Sea-Level Rise, and Arctic Change. CRV Science Open-Source Apps (terratwin-coupled-climate-ssp). https://doi.org/10.5281/zenodo.24222392
16. **[16]** White, B. P. (2026). APP-PREPRINT-11: Urban-Grid: Modeling City Traffic, Emergency Response Coverage, and Earthquake Evacuation. CRV Science Open-Source Apps (urban-grid-microscopic-traffic-resilience). https://doi.org/10.5281/zenodo.24222394
17. **[17]** White, B. P. (2026). APP-PREPRINT-12: WildFront BattleGrounds: Turn-Based Hex Strategy Game with Voxel Graphics and Procedural Map Generation. CRV Science Open-Source Apps (wildfront-hexagonal-ecological-tactics). https://doi.org/10.5281/zenodo.24222396
18. **[18]** White, B. P. (2026). APP-PREPRINT-13: AeroFlyway: Tracking Bird Migration with Weather Radar, Skyglow Maps, and Nighttime Flight Calls. CRV Science Open-Source Apps (aeroflyway-avian-migration-radar). https://doi.org/10.5281/zenodo.24222366
19. **[19]** White, B. P. (2026). APP-PREPRINT-14: HelioProbe: Simulating Spacecraft Orbits, Deep-Space Telescopes, and Comet Tails in 3D. CRV Science Open-Source Apps (helioprobe-3d-solar-system-l2). https://doi.org/10.5281/zenodo.24244083
20. **[20]** White, B. P. (2026). APP-PREPRINT-15: PopBio-Studio: Interactive Population Biology, Adaptive Landscapes, and Genetic Skyline Simulator. CRV Science Open-Source Apps (popbio-studio-adaptive-landscapes-skyline). https://doi.org/10.5281/zenodo.24244085
21. **[21]** White, B. P. (2026). APP-PREPRINT-16: QuantaCrypt-Studio: Simulating Quantum Circuits, Quantum Key Distribution, and Post-Quantum Cryptography. CRV Science Open-Source Apps (quantacrypt-studio-quantum-circuits-pqc). https://doi.org/10.5281/zenodo.24244089
22. **[22]** White, B. P. (2026). APP-PREPRINT-17: SonoWing: Analyzing Bat Echolocation and Bird Calls with Spectrograms and Harmonic Tracking. CRV Science Open-Source Apps (sonowing-bioacoustic-echolocation-analyzer). https://doi.org/10.5281/zenodo.24244091
23. **[23]** White, B. P. (2026). APP-PREPRINT-18: SwarmForge: 3D Isometric Automated Factory and Swarm Robotics Simulation Game. CRV Science Open-Source Apps (swarmforge-isometric-swarm-factory-game). https://doi.org/10.5281/zenodo.24244093