Website Press Release - Corvallis, Oregon - August 13, 2026

CIPHER-AXIS: Choosing the Right Target Before the Clinic

The problem is not only the molecule.

A large share of the roughly 90% clinical failure rate is decided long before Phase 2 and Phase 3 readouts. When a drug reaches large patient studies, lack of efficacy is among the most common reasons programs stop: the compound may be tolerable, but the biological target is not actually driving the disease in humans. Animal models, narrow hypotheses, and correlation-only gene lists all contribute to that risk.

CIPHER-AXIS (Adaptive eXploration of Intervention Structures) is CIPHERICA's system for ranking the most appropriate target interventions for a given disease.

It does not only ask which genes are associated with a condition. It asks which legal changes to the disease network are most justified so chemistry and assays are aimed at editable mechanisms, not only at correlation strength.

How CIPHER-AXIS works

1. **Disease association graph** Nodes are genes linked to the disease, weighted by association evidence. Edges are physical interactions and signed regulatory relationships among those genes. The disease selects the node set; it is not treated as a single hub vertex.
2. **Legal interventions as graph rewrites** Candidate targets are explicit edits to that graph:
    - **Node loss:** remove a protein and its edges (loss-of-function / degrader-style hypothesis)
    - **Edge block:** interrupt a regulatory link (pathway interruption)
    - **PPI drop:** break a physical interaction (complex disruption)
3. **Physics-informed ranking** Each rewrite is scored by what it removes (association mass and network structure). Energies enter a partition function; temperature sets how sharply probability concentrates on high-utility moves. This is the same statistical-mechanics control layer used in CIPHER-SKI for molecular graph rewrites.

Targets therefore emerge as ranked interventions, not only as a sorted list of associated genes.

ALS pilot (publication comparison only)

We built ALS association graphs of increasing size from public association and interaction data, ranked interventions with Boltzmann scoring, and only afterward compared results to the 28 therapeutic targets published by Pun et al.

(PandaOmics, *Front. Aging Neurosci.* 2022, Table 2). Competitor gene lists were never used to construct the graph.

Association universe	Overlap with Pun et al.
Top 200 genes	2 / 28
800	8 / 28
2,000	18 / 28
4,000	22 / 28
~8,000	**25 / 28**

Raising the association cutoff does not retrain CIPHER-AXIS. It enlarges the board: genes outside the node set cannot appear as interventions. As more disease-linked genes become eligible, overlap with the independent curated list rises. That is expected if the pipeline is working, coverage expands, and scoring still runs on every legal move.

Utility does not flatten with coverage. Within the shared set, VCP (valosin-containing protein / p97) remained the strongest loss-of-function intervention at every scale (high utility throughout). NOS1 was a clear but distant second. Most other Table 2 overlaps entered as low-mass, low-probability interventions: present on the graph, not promoted to the top of the rank. AXIS therefore shows both broad concordance with an external list and concentrated preference on a few high-utility moves, evidence that ranking is discriminating, not only recalling names.

Classic ALS genetics (SOD1, TARDBP, FUS, OPTN, TBK1, and related genes) still dominate the global top of the AXIS rank where association mass is highest. VCP sits inside that core. The two systems optimize different objectives; the overlap is complementary signal, not duplication.

The largest comparison run (~6,000 nodes, tens of thousands of candidate interventions) completed on a local workstation and was completed in under two hours. Smaller association universes rank much faster and are the practical operating range for iterative work.

Why this matters for Phase 2/3 risk

Target risk is not solved by longer gene lists alone. It is reduced when teams can state, before major chemistry and trial spend:

- what node or edge they intend to change
- why that change ranks above alternatives
- how that hypothesis maps to a molecule and a functional test

CIPHER-AXIS is built to make that statement explicit. It does not replace multi-omics platforms, biomarkers, or clinical design and it cannot guarantee late-stage success. It is an intervention layer intended to push weak target hypotheses out earlier and stronger ones into chemistry with a clearer mechanism.

One stack with CIPHER-SKI

- **CIPHER-AXIS** ranks *what to change*
- **CIPHER-SKI** generates *what to make* - scaffold hops, branch optimization, and property-aware molecular rewrites for ligands, warheads, or degraders

Target selection and molecular design share the same rewrite + statistical-mechanics philosophy.

Status

CIPHER-AXIS is under active development within the CIPHERICA platform. The ALS comparison is a retrospective, publication-only benchmark. For disease-focused graphs, assay readouts, or chemistry follow-up, contact gershon@cipherica.com · cipherica.com