Multi-Layer Biological Network Analysis for Predicting Synergistic Interactions Among Bioactive Compounds in Neurodegenerative Diseases

Eliza Reijers
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
elizareijers@uab.edu

Niek de Gruijter
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
niek.degruijter@missouri.edu

Abstract

Neurodegenerative diseases such as Alzheimer's disease and Parkinson's disease involve interacting proteinopathies, mitochondrial dysfunction, neuroinflammation, impaired proteostasis, and synaptic failure, yet most therapeutic discovery pipelines still prioritize single targets. Bioactive compounds with multi-target profiles are plausible combination partners, but predicting which compound pairs act synergistically remains difficult because compound-target, pathway-disease, and disease-disease relations are heterogeneous and incompletely observed. This manuscript develops a multi-layer biological network framework for predicting synergistic interactions among bioactive compounds in neurodegenerative disease contexts. The framework represents compounds, proteins, pathways, disease modules, and phenotypes across relation-specific layers; aligns them through shared biological entities; and combines chemical similarity, target engagement, pathway crosstalk, network propagation, and disease-module proximity. A relation-aware graph neural network estimates compound-pair and disease-specific synergy probabilities, while a multi-task objective enforces prediction symmetry, layer regularization, and class-imbalance correction. Uncertainty is quantified through ensemble variance, and candidate prioritization requires both a lower confidence bound and a mechanistic concordance criterion. The paper specifies label construction from Loewe or Bliss reference models, calibration procedures, scaffold and disease-held-out validation, negative-sampling controls, and domain-specific failure modes. The contribution is a falsifiable and auditable integration of multilayer network biology, synergy theory, and uncertainty-aware machine learning that can guide experimental combination screening in neurodegenerative disease research.

Keywords

multilayer networks; network pharmacology; synergy prediction; graph neural networks; neurodegenerative disease; uncertainty calibration

1. Introduction

Neurodegenerative diseases such as Alzheimer's disease and Parkinson's disease present a persistent therapeutic problem: core molecular lesions are identifiable, but single-target interventions have not produced durable disease modification across heterogeneous patient

populations. The amyloid hypothesis framed aggregation as a principal driver of Alzheimer's disease [1], while neuropathological evidence shows that disease states typically combine proteinopathy, glial activation, synaptic loss, and regional vulnerability rather than a single lesion [2]. Aging amplifies these processes through proteostasis failure, mitochondrial dysfunction, and chronic inflammation [3], and Parkinson's disease similarly involves α-synuclein pathology, dopaminergic loss, and widespread non-motor network dysfunction [4]. This mechanistic plurality creates a rationale for combination strategies that act on several disease-associated processes at once.

Network pharmacology and network medicine provide a useful language for this problem because they treat drug action as a perturbation of a biological system rather than as a single receptor event [5,6]. Multilayer network formalisms extend this view by distinguishing relation types that single graphs collapse, such as compound-target binding, protein-protein interaction, pathway membership, disease-module association, and phenotype similarity [7,8]. Herbal and natural-product network pharmacology has already used compound-target-pathway representations to interpret complex mixtures and multi-component interventions [9]. Such representations are especially relevant to bioactive compounds, whose effects are often distributed across several targets and whose combination behavior may depend on pathway crosstalk rather than on one dominant mechanism.

Synergy itself is not a single biological quantity. Loewe additivity defines a reference surface for dose combinations [10], Bliss independence defines an expectation under probabilistic independence [11], and combination-index methods formalize departure from a chosen additivity model [12]. These definitions answer different questions and can disagree when agents have overlapping dose-response ranges, nonlinear pharmacokinetics, or context-dependent mechanisms. A prediction framework therefore needs to specify which reference model generated each label, how uncertainty in dose-response estimation enters the target, and whether predictions are intended for in vitro, ex vivo, or organism-level contexts.

Machine learning has improved related prediction tasks, including drug-drug interaction modeling and polypharmacy side-effect prediction with graph convolutional networks [13]. For herbal medicines, community-based network models have been proposed to identify synergistic combinations from complex systems [14]. These approaches establish feasibility but leave an integration gap: compound-level chemical information, target-level engagement, pathway-level crosstalk, disease-module proximity, synergy labels, and predictive uncertainty are rarely treated within one coherent analytical design. The central contribution here is a multi-layer biological network framework that links these elements and specifies how predictions should be evaluated before experimental follow-up.

2. Biological and Network Rationale

The framework treats neurodegeneration as a module-level perturbation problem. A compound may engage several proteins, some of which lie near a disease module and others of which lie in adjacent pathways that modify cellular vulnerability. Network proximity is useful because disease proteins are often not randomly distributed in interactomes; they tend to form connected or nearby modules [6]. However, proximity alone does not establish efficacy. It is a mechanistic feature that must be combined with chemical realism, target engagement evidence, pathway directionality, and experimental synergy labels.

Let the multilayer network contain L relation-specific layers over a shared entity set. Nodes may represent compounds, proteins, pathways, disease modules, and phenotypes, while each layer encodes only one relation type. The layer-specific adjacency is defined as:

$$A(\ell, u, v) = w(\ell, u, v) \cdot I[(u, v) \in \mathrm{E}(\ell)]$$

Here ℓ indexes the biological layer, u and v are nodes, w(ℓ,u,v) is a nonnegative edge weight, E(ℓ) is the edge set for layer ℓ, and I[·] is an indicator function equal to one when the relation is present. Weights can reflect confidence, assay support, or literature curation strength, but they should be normalized within layer before fusion so that a densely annotated layer does not dominate a sparsely annotated one. Shared nodes align compound, target, pathway, and disease views without forcing all relations into a single adjacency matrix.

Propagation converts compound-target evidence into a distributed perturbation score over the network. For a seed vector y representing compound-engaged targets and a row-normalized transition matrix S, the steady-state propagation score is:

$$f = (1 - \alpha)(I - \alpha S)^{\wedge} - 1\, y$$

In this expression, f is the propagated score, I is the identity matrix, α is a restart parameter with $0 \leq \alpha < 1$, and y encodes initial target weights. The expression is a standard personalized propagation model; it rewards targets that are reachable through several short paths from the compound's direct targets. The condition $\alpha < 1$ ensures that $I - \alpha S$ is invertible for a normalized transition matrix. The resulting vector is a feature, not a validated efficacy estimate.

Disease proximity is computed between a compound target set T and a disease module D. A scale-sensitive proximity score is:

$$s(T, D) = \left(\frac{1}{(|T||D|)}\right) \sum_{t \in T} \sum_{d \in D} \exp\left(- d_G \frac{(t, d)}{\tau}\right)$$

Here |T| and |D| are set sizes, d_G(t,d) is a network distance between target t and disease node d, and τ controls how quickly distant pairs stop contributing. The score is high when many compound targets lie close to many disease-module nodes. Shortest-path distance can be replaced by diffusion distance when connectivity is noisy, but the choice must be fixed before evaluation to avoid tuning the feature definition to the test set.

3. Methodological Framework

3.1 Representation Learning over Relation-Specific Layers

Relation-aware representation learning is needed because the same node may have different roles across layers. Graph attention networks allow a node to weight its neighbors differently according to learned coefficients [15], message-passing architectures provide a general mechanism for exchanging information across graph edges [16], and standard graph convolutional networks offer a simpler baseline [17]. The proposed encoder updates node states by aggregating relation-specific neighborhoods:

$$h_u^{\wedge}(\ell + 1) = \sigma(\sum_\{r \in R\} \sum_\{v \in N_r(u)\}\, a_uv^{\wedge} r\, W_r\, h_v^{\wedge}(\ell))$$

Here h_u^(ℓ+1) is the updated embedding for node u, σ is a nonlinear activation, R is the set of relation types, N_r(u) is the neighborhood of u under relation r, a_uv^r is a normalized attention coefficient, W_r is a relation-specific trainable matrix, and h_v^(ℓ) is the current embedding of neighbor v. Layer-specific transformations prevent compound-target evidence

from being mixed with pathway-disease evidence before the model has learned which relations matter for the disease context. Attention coefficients also provide a partial interpretability mechanism, although attention weights should not be treated as causal explanations without external validation.

The model outputs a compound embedding h_i, a second compound embedding h_j, and a disease-context embedding h_d. A pair representation combines shared and contrastive information, because synergy may depend both on overlapping targets and on complementary mechanisms. Chemical similarity, dose ratio, pathway overlap, and disease-module proximity are appended as explicit features rather than left implicit in the graph structure. This design allows the model to learn nonlinear interactions while retaining mechanistic descriptors that can be audited.

3.2 Biological Layer Construction and Mechanistic Features

The network is assembled from public and study-specific resources in a curated, versioned pipeline. Compound-target edges can be derived from binding, bioactivity, and perturbation evidence; protein-protein interaction edges from curated interactomes; target-pathway edges from pathway databases; pathway-disease edges from disease-gene and module enrichment resources; and compound-compound edges from chemical structure, scaffold, or target-profile similarity. Hetionet-style integrative schemas demonstrate how heterogeneous biomedical knowledge can be represented in a typed network for drug repurposing [18]. Network-proximity approaches further show how target sets can be compared with disease modules for in silico efficacy screening [19]. These precedents motivate the layer architecture but do not remove the need for study-specific curation rules.

Node features should include degree, betweenness centrality, tissue expression breadth, disease specificity, pathway entropy, and compound physicochemical descriptors. Edges should carry provenance labels, such as experimentally measured, curated, inferred, or predicted. Provenance matters because literature bias can make heavily studied proteins appear more central than they are biologically. A compound that has been assayed against many kinases may accumulate edges for reasons of experimental convenience rather than disease relevance. The framework therefore includes edge-confidence weights and optional layer dropout during training so that the model does not rely on one over-annotated layer.

3.3 Synergy Label Definition and Predictive Objective

When dose-response data are available, a compound pair is labeled positive for a disease context if the observed combination effect exceeds the chosen additivity reference by a pre-specified margin with adequate confidence. Loewe additivity [10] and Bliss independence [11] should be recorded separately rather than merged into a single label, because a pair can depart from one reference model and not the other. Combination-index estimates [12] can provide an additional label source when the experimental design supports them. Pairs without sufficient dose-response coverage should be treated as unlabeled rather than forced into the negative class.

The predictive model estimates the probability of synergy for compound pair (i,j) in disease context d:

$$p_ijd = \sigma\left(w^T \psi\left(h_i, h_j, h_d, \varphi_ij\right) + b\right)$$

Here p_ijd is the predicted probability, σ is the logistic function, w and b are trainable parameters, ψ is a feature-fusion function, h_i and h_j are compound embeddings, h_d is the

disease-context embedding, and φ_ij contains explicit pair features such as chemical similarity, target overlap, pathway crosstalk, and dose ratio. The fusion function can be a multilayer perceptron, but its inputs should remain separable so that ablation analyses can identify whether graph structure, chemical similarity, or mechanistic proximity drives predictions.

Training uses a multi-task objective with class-imbalance weighting, regularization, and a symmetry penalty:

$$L = \sum w_ijd\, BCE(p_ijd, y_ijd) + \lambda 1\|\Theta\|^2 + \lambda 2\sum(p_ijd - p_jid)^2$$

In this expression, L is the total loss, w_ijd weights informative or uncertain labels, BCE is binary cross-entropy, y_ijd is the synergy label, Θ denotes trainable parameters, λ1 controls weight decay, and λ2 controls the penalty for asymmetric predictions under compound-order reversal. The symmetry term is a modeling prior, not a biological law; it prevents arbitrary ordering effects when the biological features are otherwise symmetric. Disease-specific tasks can share graph encoder parameters while maintaining separate output heads when the evidence base differs across Alzheimer's disease, Parkinson's disease, and other neurodegenerative contexts.

3.4 Uncertainty, Calibration, and Decision Rule

Predictive probability alone is insufficient for prioritizing expensive experiments. Ensemble variance provides a simple uncertainty estimate:

$$U_ijd = (1/M)\sum_\{m = 1\}\text{^}M\ (p_m(ijd) - \bar{p}_ijd)^2$$

Here M is the number of ensemble members, p_m(ijd) is the synergy probability from member m, $\bar{p}$_ijd is the ensemble mean, and U_ijd is the predictive variance. Variance captures model disagreement but not necessarily epistemic correctness; calibration must therefore be assessed with reliability diagrams, Brier scores, and held-out disease contexts. A candidate is prioritized only when the lower confidence bound exceeds a probability threshold and a mechanistic concordance score exceeds a separate threshold:

$$predict\ synergy\ if\ p_ijd - \kappa\sqrt{U_ijd} \geq \theta\ and\ q_ijd \geq \tau$$

In this decision rule, κ controls risk aversion, θ is the probability threshold, q_ijd is a mechanistic concordance score, and τ is its threshold. The concordance score can combine pathway overlap, directionality agreement, disease-module proximity, and absence of known antagonistic interactions. This two-stage rule reduces the chance that a high-probability but mechanistically incoherent prediction is advanced to experimental testing.

4. Analytical Assessment and Design Trade-offs

The framework should be compared against credible baselines rather than only against itself. These include single-layer compound-target models, matrix factorization over compound-disease associations, random forests or gradient boosting on concatenated similarity features, a standard graph convolutional network [17], a graph attention model without propagation features [15], and an ablation that removes chemical or pathway layers. Machine learning in drug discovery has repeatedly shown that performance depends on representation, label quality, and validation design rather than on model depth alone [20]. Deep learning for drug-drug and drug-food interaction prediction further illustrates the value of combining chemical and biological representations, while also exposing the difficulty of

generalizing across interaction types [21]. The proposed model should therefore report performance by disease context and by evidence type, not only as a pooled score.

Validation must address several sources of optimistic bias. Random pair splits can place the same compound or scaffold in both training and test sets, inflating performance. Scaffold splits, compound-held-out splits, disease-held-out splits, and temporal splits provide stricter tests. Negative sampling is another failure point: randomly sampled non-synergistic pairs are often trivially separable from positives because they share no targets or pathways. Hard negatives should include pairs with overlapping targets but no observed synergy, pairs with similar chemical structures but different disease contexts, and pairs that are individually active but experimentally antagonistic. The model should also be evaluated under missing-edge stress, because incomplete interactomes can make true disease-module connections appear absent [22].

Domain-specific error modes deserve explicit treatment. Compound promiscuity can create many weak edges that mimic multi-target action without producing useful synergy. Dose mismatch can make an in vitro synergistic pair irrelevant at achievable tissue concentrations. Cell-type mismatch can invert pathway effects between neurons, microglia, and astrocytes. Literature bias can overrepresent well-studied targets and underrepresent understudied but disease-relevant pathways. The model should therefore include sensitivity analyses that perturb edge weights, remove high-degree nodes, and vary the propagation restart parameter α. If predictions are unstable under small perturbations, they should not be prioritized for experimental follow-up.

5. Technical Discussion

The main value of a multi-layer formulation is not that it adds complexity but that it preserves distinctions that affect inference. Compound-target edges answer whether a compound can engage a protein; pathway edges answer whether that engagement propagates to a disease-relevant process; disease-module edges answer whether the process lies near the disease's molecular neighborhood; and chemical edges answer whether two compounds are plausibly comparable or complementary. Collapsing these relations into one graph loses the ability to audit which relation drives a prediction. The proposed framework keeps those relations separable and allows layer ablations to test their contributions.

Mechanistic interpretability should be treated cautiously. Attention weights and path enrichments can suggest why a pair is predicted to be synergistic, but they do not demonstrate synergy. The output of the model is a prioritized hypothesis with an uncertainty estimate. Experimental validation remains necessary, ideally using dose-response designs that can distinguish Loewe, Bliss, and combination-index behavior. Community-based screening for herbal medicines [14] can generate candidate groups, while the present framework can refine pair-level predictions and attach mechanistic and uncertainty annotations. Disease-network analyses [22] can inform module boundaries, but incomplete interactomes can distort those boundaries and should be represented as uncertainty rather than ignored.

The framework also has limits that follow from its assumptions. It assumes that curated edges are sufficiently accurate to support propagation, that disease modules can be defined at a useful granularity, and that synergy labels are comparable enough within a disease context to support supervised learning. These assumptions are strong in neurodegeneration, where human tissue is scarce, model systems differ substantially, and clinical endpoints are slow and noisy. A defensible use of the model is therefore triage: ranking compound pairs for

experimental testing under explicitly stated evidence conditions, not replacing mechanistic studies or clinical trials.

6. Conclusion

This manuscript has developed a multi-layer biological network framework for predicting synergistic interactions among bioactive compounds in neurodegenerative diseases. The framework integrates relation-specific biological layers, propagation-based disease-module features, chemical and target-pair descriptors, a relation-aware graph neural network, a symmetry-regularized synergy objective, ensemble uncertainty, and a two-stage decision rule. It also specifies label construction from established synergy models, calibration and validation procedures, and domain-specific failure modes. The central contribution is an auditable analytical design that connects network biology with synergy prediction while distinguishing mechanistic reasoning from observed evidence. Its practical value lies in prioritizing compound pairs for experimental follow-up under transparent uncertainty and validation conditions, particularly for diseases where combination therapies must address several interacting pathological processes.

References

1. Hardy, J., & Selkoe, D. J. (2002). The amyloid hypothesis of Alzheimer's disease: Progress and problems on the road to therapeutics. Science, 297(5580), 353–356. https://doi.org/10.1126/science.1072994

2. Dugger, B. N., & Dickson, D. W. (2017). Pathology of neurodegenerative diseases. Cold Spring Harbor Perspectives in Biology, 9(7), a028035. https://doi.org/10.1101/cshperspect.a028035

3. Hou, Y., Dan, X., Babbar, M., Wei, Y., Hasselbalch, S. G., Croteau, D. L., & Bohr, V. A. (2019). Ageing as a risk factor for neurodegenerative disease. Nature Reviews Neurology, 15(10), 565–581. https://doi.org/10.1038/s41582-019-0244-7

4. Poewe, W., Seppi, K., Tanner, C. M., Halliday, G. M., Brundin, P., Volkmann, J., Schrag, A.-E., & Lang, A. E. (2017). Parkinson disease. Nature Reviews Disease Primers, 3, 17013. https://doi.org/10.1038/nrdp.2017.13

5. Hopkins, A. L. (2008). Network pharmacology: The next paradigm in drug discovery. Nature Chemical Biology, 4(11), 682–690. https://doi.org/10.1038/nchembio.118

6. Barabási, A.-L., Gulbahce, N., & Loscalzo, J. (2011). Network medicine: A network-based approach to human disease. Nature Reviews Genetics, 12(1), 56–68. https://doi.org/10.1038/nrg2918

7. De Domenico, M., Solé-Ribalta, A., Cozzo, E., Kivelä, M., Moreno, Y., Porter, M. A., Gómez, S., & Arenas, A. (2013). Mathematical formulation of multilayer networks. Physical Review X, 3(4), 041022. https://doi.org/10.1103/PhysRevX.3.041022

8. Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. Journal of Complex Networks, 2(3), 203–271. https://doi.org/10.1093/comnet/cnu016

9. Li, S., & Zhang, B. (2013). Traditional Chinese medicine network pharmacology: Theory, methodology and application. Chinese Journal of Natural Medicines, 11(2), 110–120. https://doi.org/10.1016/S1875-5364(13)60037-0

10. Loewe, S. (1953). The problem of synergism and antagonism of combined drugs. Arzneimittelforschung, 3(6), 285–290.

11. Bliss, C. I. (1939). The toxicity of poisons applied jointly. Annals of Applied Biology, 26(3), 585–615. https://doi.org/10.1111/j.1744-7348.1939.tb06990.x

12. Chou, T.-C., & Talalay, P. (1984). Quantitative analysis of dose-effect relationships: The combined effects of multiple drugs or enzyme inhibitors. Advances in Enzyme Regulation, 22, 27–55. https://doi.org/10.1016/0065-2571(84)90007-4

13. Zitnik, M., Agrawal, M., & Leskovec, J. (2018). Modeling polypharmacy side effects with graph convolutional networks. Bioinformatics, 34(13), i457–i466. https://doi.org/10.1093/bioinformatics/bty294

14. HerbSyner_Finder: a network community-based model for identifying synergistic combinations from herbal medicines and complex systems. Targetome. 2026;2(2):0-0. doi: 10.48130/targetome-0026-001

15. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2017). Graph attention networks. arXiv:1710.10903. https://arxiv.org/abs/1710.10903

16. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. arXiv:1704.01212. https://arxiv.org/abs/1704.01212

17. Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. arXiv:1609.02907. https://arxiv.org/abs/1609.02907

18. Himmelstein, D. S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S. L., Hadley, D., Green, A. J., Khankhanian, P., & Baranzini, S. E. (2017). Systematic integration of biomedical knowledge prioritizes drugs for repurposing. eLife, 6, e26726. https://doi.org/10.7554/eLife.26726

19. Guney, E., Menche, J., Vidal, M., & Barabási, A.-L. (2016). Network-based in silico drug efficacy screening. Nature Communications, 7, 10331. https://doi.org/10.1038/ncomms10331

20. Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., Spitzer, M., & Zhao, S. (2019). Applications of machine learning in drug discovery and development. Nature Reviews Drug Discovery, 18(6), 463–477. https://doi.org/10.1038/s41573-019-0024-5

21. Ryu, J. Y., Kim, H. U., & Lee, S. Y. (2018). Deep learning improves prediction of drug-drug and drug-food interactions. Proceedings of the National Academy of Sciences of the United States of America, 115(18), E4304–E4311. https://doi.org/10.1073/pnas.1803294115

22. Menche, J., Sharma, A., Kitsak, M., Ghiassian, S. D., Vidal, M., Loscalzo, J., & Barabási, A.-L. (2015). Disease networks. Uncovering disease-disease relationships through the incomplete interactome. Science, 347(6224), 1257601. https://doi.org/10.1126/science.1257601