AI-Assisted Prioritization of Combination Therapies Based on Metabolic Pathway Interactions and Cellular Network Connectivity

Therese Ginevan
Department of Computer Science, George Mason University, Fairfax, VA, USA.
thereseginevan@gmu.edu

Abstract

Combination therapy design faces a combinatorial search problem in which the number of possible drug pairs or higher-order regimens vastly exceeds the capacity of experimental testing. Computational prioritization can reduce this burden, yet many existing predictors treat drug synergy as a property of molecular similarity or target overlap and do not explicitly represent the metabolic pathway interactions and cellular network context that shape combination response. This manuscript develops an AI-assisted prioritization framework that integrates metabolic pathway cross-talk with cellular network connectivity to rank candidate combinations. The framework represents drugs, targets, pathways, metabolic reactions, and cell-context nodes in a heterogeneous graph; derives pathway-interaction features from shared reactions and metabolite overlap; computes connectivity scores by network propagation from drug targets; and uses a relational graph neural network to learn context-dependent drug embeddings. A bilinear synergy model combines learned embeddings with pathway-interaction and connectivity features, while Monte Carlo dropout supplies predictive uncertainty for risk-adjusted ranking. The central methodological contribution is a coupled representation in which pathway interactions modify the connectivity features that propagate through the cellular network, rather than being treated as independent covariates. Analytical assessment identifies conditions under which pathway-interaction features are informative, such as convergent metabolic dependencies and compensatory flux, and conditions under which connectivity can mislead, such as network rewiring and cell-type heterogeneity. The framework also specifies baselines, ablation logic, and failure modes, including label noise, batch effects, and distribution shift. By making uncertainty explicit and linking predictions to mechanistic pathway and network quantities, the approach supports triage of combination therapies for experimental follow-up while preserving the distinction between model-derived prioritization and validated biological effect.

Keywords

drug combination prediction; network medicine; metabolic pathways; graph neural networks; synergy modeling; uncertainty quantification

1. Introduction

Combination therapies are central to oncology, infectious disease, and complex metabolic disorders because simultaneous perturbation of multiple nodes can produce effects that monotherapy cannot achieve. The pharmacological basis of combination benefit is usually expressed through synergy, additivity, or antagonism, with formal models such as the median-effect principle providing a framework for quantifying departure from expected additive behavior [1]. Yet the space of plausible combinations is enormous, and experimental screens

cover only a small fraction of drug pairs, doses, schedules, and cell contexts. Computational prioritization is therefore essential for directing limited experimental capacity toward combinations with stronger mechanistic support and greater likelihood of reproducible benefit.

Network medicine provides a conceptual foundation for this task by treating disease and drug action as perturbations of interconnected molecular systems rather than isolated targets [2]. Network pharmacology extends this view to multi-target intervention, arguing that combination effects depend on the structure of the target network and on the propagation of perturbation through it [3]. Drug-target networks make the multi-target character of approved drugs explicit and show that drugs often engage several proteins with shared neighborhood relationships [4]. Computational drug-target interaction models have made these relationships learnable at scale, using chemical, genomic, and network features to predict previously unobserved interactions [5]. Network-community approaches have also been proposed to identify synergistic combinations from herbal medicines and complex systems, illustrating how community structure can organize candidate combinations for further assessment [6].

A persistent gap remains between network-level representations and the specific metabolic and cellular mechanisms that determine combination response. Metabolic pathways generate dependencies that are not captured by target overlap alone: two drugs may inhibit distinct enzymes in the same pathway, produce compensatory flux through parallel branches, or alter metabolite pools that change the effective concentration of a third agent. Cellular network connectivity adds a second layer of context because the consequence of inhibiting a target depends on which downstream modules remain connected and which compensatory routes are available. Many AI models incorporate either pathway annotations or network propagation, but they often treat these quantities as parallel features rather than as coupled determinants of combination effect. This separation weakens mechanistic interpretability and can produce rankings that are sensitive to graph artifacts rather than to plausible biology.

This manuscript develops a methodological framework for AI-assisted prioritization of combination therapies that couples metabolic pathway interactions with cellular network connectivity. The framework constructs a heterogeneous graph over drugs, targets, pathways, reactions, and cell-context nodes; derives pathway-interaction features from shared reactions and metabolite overlap; computes target-to-pathway and target-to-module connectivity by network propagation; and learns relational drug embeddings with a graph neural network. A synergy-scoring layer combines these representations, and an uncertainty-aware decision rule ranks combinations by expected benefit relative to predictive risk. The contribution is not an empirical claim about any particular drug pair but a technically specified architecture and analytical rationale for integrating pathway-level and network-level evidence in combination prioritization.

2. Biological and Computational Foundations

2.1 Metabolic Pathway Interactions in Combination Response

Metabolic pathway interactions determine how simultaneous perturbations combine. When two drugs inhibit separate steps in a linear pathway, the combined effect depends on whether flux control is distributed across steps or concentrated at one enzyme. When inhibition diverts flux through a parallel branch, the pathway can buffer the perturbation and reduce synergy. When two drugs converge on complementary metabolic dependencies, such as nucleotide synthesis and one-carbon metabolism, the combination may exhaust compensatory capacity. These mechanisms are pathway-specific and context-dependent, so a

model that represents only drug similarity or target overlap cannot distinguish combinations that share targets from combinations that share pathway-level consequences.

Learned drug representations from similarity integration provide a practical basis for pairwise scoring because they compress chemical, target, and side-effect information into vectors that can be compared across large libraries [7]. Polypharmacy graph models further show that drug combinations can be represented as edges in a multi-relational graph, with side-effect prediction benefiting from explicit modeling of drug-drug relationships [8]. Deep synergy models demonstrate that molecular and genomic features can be combined in nonlinear predictors trained on combination screening data [9]. Target-network ranking approaches add a network-propagation step and show that drug effects can be prioritized by their proximity to disease modules [10]. These developments motivate a representation that keeps pathway-interaction structure explicit rather than absorbing it into a generic drug embedding.

2.2 Cellular Network Connectivity and Drug Action

Cellular network connectivity describes how a perturbation at one node affects distant nodes through paths in a protein interaction, signaling, or regulatory network. A drug that inhibits a target with many short paths to a disease module may produce a stronger functional effect than a drug that inhibits a target with similar annotation but weaker connectivity. Connectivity also determines compensation: if a target sits in a densely connected neighborhood with redundant routes, inhibition may be buffered; if it sits at a bottleneck, the effect may be amplified. Combination prioritization therefore benefits from representations that encode both the local target neighborhood and the global position of the target relative to pathway and disease modules.

Graph neural networks provide a natural formalism for this problem. Graph convolutional networks aggregate information from local neighborhoods and can learn node embeddings that reflect connectivity structure [11]. Graph attention mechanisms allow the model to weight different neighbors according to context, which is useful when some interactions are more informative than others [12]. Relational graph convolutional networks extend message passing to multiple edge types, making them suitable for heterogeneous graphs that contain drug-target, protein-protein, pathway-member, and reaction-participant relations [13]. These architectures do not by themselves encode metabolic pathway interactions, but they supply the connectivity-aware embedding layer into which pathway-interaction features can be injected.

2.3 AI Approaches for Combination Prioritization

AI approaches to combination prioritization differ in their input representations, supervision signals, and treatment of uncertainty. Fingerprint-based models use chemical structure similarity to transfer synergy labels between similar compounds [16]. Feature-based models combine drug descriptors with genomic or transcriptomic context. Graph-based models use network structure to propagate information across drugs, targets, and pathways. Each approach has characteristic failure modes: fingerprint models can overemphasize chemical similarity when mechanism depends on pathway context; feature-based models can overfit to batch-specific genomic patterns; graph models can amplify degree biases and spurious edges. Uncertainty quantification is therefore not a secondary add-on. Monte Carlo dropout provides a practical approximation to Bayesian model uncertainty by keeping dropout active at inference and measuring variance across stochastic forward passes [14].

Distinguishing aleatoric from epistemic uncertainty helps identify whether high predictive variance arises from noisy labels or from limited coverage of the chemical and network space [15]. A prioritization framework should expose this distinction because experimental follow-up is costly and because a ranked list without calibrated uncertainty can be misleading.

3. Methodological Framework

3.1 Formal Problem Definition

The framework represents the biological and pharmacological context as a heterogeneous graph. Nodes correspond to drugs, protein targets, metabolic pathways, metabolic reactions, and cell-context modules. Edges encode drug-target binding, protein-protein interaction, pathway membership, reaction participation, and cell-context expression association. The graph is written as follows.

$$G = (V, E, R), X \in \mathbb{R}^{|V| \times d}$$

Here V is the node set, E is the edge set, R is the set of relation types, and X is the node feature matrix with d features per node. Each drug node has a chemical descriptor vector and an indication vector. Each target node has sequence, functional, and network centrality features. Each pathway node has reaction-composition and annotation features. Each reaction node has substrate-product and enzyme annotations. Cell-context nodes carry expression and dependency summaries. The prediction task is to assign a synergy score to an ordered or unordered drug pair in a specified cell context, using the graph and node features as inputs.

3.2 Heterogeneous Cellular Network Construction

Drug nodes and drug-target edges can be assembled from curated resources such as DrugBank, which provides approved and investigational drug records, target annotations, and interaction information [17]. Metabolic pathway and reaction structure can be assembled from KEGG, which supplies pathway maps, enzyme commissions, and metabolite relationships [18]. Protein-protein association networks can be assembled from STRING, which integrates experimental, database, and text-mined evidence with confidence scores [19]. Cell-context features can be derived from single-cell or bulk expression atlases such as the Human Cell Atlas, which provides cell-type-resolved expression profiles for context-specific network weighting [20]. These sources differ in coverage, granularity, and reliability, so the graph construction step should record provenance and evidence strength for each edge rather than collapsing all relations into a single adjacency matrix.

Two design choices matter for downstream prioritization. First, pathway and reaction nodes should be retained as first-class nodes rather than reduced to drug annotations, because otherwise the model cannot represent shared metabolic dependencies or compensatory flux. Second, cell-context nodes should be connected to targets and pathways through expression-weighted edges, so that the same drug pair can receive different scores in different cellular contexts. This design supports context sensitivity without requiring a separate model for every cell type; the context enters through graph structure and node features.

3.3 Pathway-Interaction Embedding

Pathway interactions are derived from shared reactions and metabolite overlap. For pathways p and q, the interaction strength can be defined as a normalized reaction-sharing score.

$$M_\{pq\} \; = \; \sum_\{k\} \, r_\{pk\} \, r_\{qk\} \, / \, \sqrt{(d_p \; d_q)}$$

Here r_{pk} is an indicator equal to one if pathway p contains reaction k, d_p is the number of reactions in pathway p, and the denominator normalizes for pathway size. The resulting matrix M captures direct reaction sharing and provides a compact representation of pathway cross-talk. Metabolite overlap can be incorporated by replacing the indicator product with a metabolite-sharing weight or by adding a second matrix with the same normalization. The pathway-interaction matrix is then used to compute a pathway-context vector for each drug pair, with larger values indicating that the two drugs perturb pathways that share reaction components.

Connectivity from drug targets to pathways and cell-context modules is computed by propagation on the graph. For a target set T, the propagation score vector is:

$$s_T = (1 - \alpha)(I - \alpha W)^{-1} b_T$$

Here W is the row-normalized adjacency matrix over the target, pathway, and module subgraph, b_T is a seed vector with mass on the targets of the drug or drug pair, α is a restart probability between zero and one, and I is the identity matrix. The score s_T measures how strongly the target set connects to each pathway and module after accounting for multi-step paths. The parameter α controls the balance between local and global propagation; smaller values emphasize immediate neighbors, while larger values allow information to reach distant modules. This propagation step is where pathway interactions and network connectivity become coupled: pathway-interaction features can reweight edges or seed masses before propagation, so that shared metabolic dependencies alter the connectivity scores rather than remaining separate covariates.

3.4 Connectivity-Aware Message Passing

The graph is processed by a relational graph neural network that updates node embeddings through relation-specific message passing. For node v at layer l, the update is:

$$h_v\text{^}\{(l+1)\} = \sigma(\, W_0\, h_v\text{^}\{(l)\} + \sum_\{r\} \sum_\{u \in N_r(v)\} (1/c_\{vu\})\, W_r\, h_u\text{^}\{(l)\}\,)$$

Here h_v^{(l)} is the embedding of node v at layer l, σ is a nonlinear activation, W_0 is the self-connection weight, W_r is the weight matrix for relation type r, N_r(v) is the set of neighbors of v under relation r, and c_{vu} is a normalization factor such as the product of node degrees. The relational structure allows the model to distinguish drug-target edges from protein-protein edges, pathway-membership edges, and reaction-participation edges. After several layers, the drug embedding aggregates chemical information, target neighborhood structure, pathway membership, and context-dependent connectivity. The pathway-interaction matrix from the previous subsection can be used to initialize or regularize the pathway embeddings, which encourages pathways with shared reactions to occupy nearby regions of the embedding space.

3.5 Synergy Scoring and Uncertainty Estimation

The synergy score for a drug pair is computed from the learned drug embeddings and the pathway and connectivity features. A bilinear form provides a flexible pairwise interaction term while remaining interpretable.

$$y_{ij} = \sigma\left(u_i^T W u_j + \beta c_{ij} + \gamma m_{ij}\right)$$

Here y_{ij} is the predicted probability or normalized score for synergy between drugs i and j, u_i and u_j are their learned embeddings, W is a bilinear weight matrix, c_{ij} is the connectivity concordance between the target sets of the two drugs, m_{ij} is the pathway-

interaction score derived from M, and β and γ are coefficients that control the contribution of connectivity and pathway features. The bilinear term captures learned chemical and network compatibility, while c_{ij} and m_{ij} provide explicit mechanistic channels. In practice, c_{ij} can be computed as the cosine similarity between propagated target-set vectors in pathway space, and m_{ij} can be computed as the maximum or weighted average of M over the pathways perturbed by the two drugs. The score is not assumed to be symmetric unless the underlying features are symmetric; asymmetric cases can represent order-dependent effects such as schedule dependence.

Predictive uncertainty is estimated by Monte Carlo dropout. At inference, dropout remains active and multiple forward passes produce a distribution of scores for each drug pair. The mean and variance of this distribution provide a predictive mean and an epistemic uncertainty estimate. For the ranking stage, this uncertainty is used to avoid overconfident placement of poorly covered pairs near the top of the list.

3.6 Training Objective and Decision Rule

The model is trained with a weighted binary cross-entropy loss for synergy labels, with an optional regression loss for continuous synergy scores. The loss for a batch of drug pairs is:

$$L = -\sum_\{i,j\} w_\{ij\} [y_\{ij\} log \hat{y}_\{ij\} + (1 - y_\{ij\}) log(1 - \hat{y}_\{ij\})] + \lambda \|\theta\|\^2$$

Here y_{ij} is the observed label or normalized synergy value, ŷ_{ij} is the model prediction, w_{ij} is a sample weight that can downweight noisy or imbalanced pairs, λ is a regularization strength, and θ denotes all trainable parameters. The weighting term is important because combination screens are unbalanced: most pairs are untested, and tested pairs are enriched for known or plausible interactions. The regularization term controls overfitting in the high-dimensional embedding space.

For prioritization, the uncertainty-adjusted score is:

$$S_{ij} = \mu_{ij} - \kappa\sigma_{ij}$$

Here μ_{ij} is the mean predicted synergy from Monte Carlo dropout, σ_{ij} is the predictive standard deviation, and κ is a risk-aversion parameter. Larger κ penalizes uncertain pairs more strongly, shifting the ranking toward combinations with both high predicted benefit and stable predictions. This decision rule does not replace experimental validation; it orders candidates for follow-up under an explicit risk preference. The parameter κ can be set according to experimental budget, with larger values when false positives are costly and smaller values when broad exploration is desired.

4. Analytical Assessment and Discussion

4.1 Baselines and Ablation Logic

A rigorous assessment of this framework requires baselines that isolate the contribution of pathway interactions and network connectivity. A drug-similarity baseline using chemical fingerprints and target overlap provides a reference for non-network prediction. A target-overlap baseline using the Jaccard index between target sets tests whether the model adds information beyond shared targets. A network-propagation baseline using connectivity scores without pathway interactions tests the value of metabolic coupling. A graph neural network without pathway-interaction features tests whether relational message passing alone captures the relevant structure. The full model adds pathway-interaction features, connectivity-aware propagation, and uncertainty-adjusted ranking. Ablation should be performed in that order so

that performance changes can be attributed to specific components rather than to overall model capacity.

4.2 Failure Modes and Boundary Conditions

Several failure modes are likely. Label noise in combination screens can arise from dose selection, schedule differences, and assay-specific readouts; if labels are noisy, the model may learn batch-specific patterns rather than biological synergy. Degree bias in graph neural networks can cause high-degree targets and pathways to dominate embeddings, which may inflate scores for well-studied drugs and underrepresent less-characterized compounds. Pathway annotation incompleteness can obscure shared reactions, especially for non-model organisms or rare cell types. Context mismatch can occur when a model trained on one tissue or cell line is applied to another with different expression and network rewiring. Finally, metabolic compensation can produce non-additive effects that are not captured by static pathway maps, because enzyme expression and flux change after perturbation. These boundary conditions should be reported alongside rankings, and sensitivity analysis should vary graph construction choices, propagation parameters, and uncertainty thresholds.

4.3 Evaluation Design and Evidence Standards

Evaluation should distinguish between retrospective ranking and prospective validation. Retrospective evaluation can use held-out drug pairs and cross-validation by drug, cell line, or pathway to test generalization. Cross-validation by drug is particularly important because random pair splitting can leak information about the same drug into training and test sets. Metrics should include precision at fixed recall, area under the precision-recall curve, and calibration of predicted probabilities. Because experimental validation is costly, the framework should also report the number of candidates required to achieve a given hit rate under a specified uncertainty threshold. Any empirical claim should be tied to a predefined evaluation protocol and should not generalize beyond the screened contexts. The present manuscript specifies the design logic and expected failure modes; it does not report experimental results.

4.4 Interpretability and Translational Use

Interpretability is a practical requirement for combination prioritization because experimentalists need to know why a pair is ranked highly. The framework supports interpretation at three levels: pathway-interaction scores identify shared reactions and metabolic dependencies; connectivity scores identify target-to-module proximity; and uncertainty estimates indicate whether a ranking is stable. These quantities can be inspected for individual pairs and aggregated across cell contexts. Translational use should treat the output as a triage tool rather than a therapeutic recommendation. Combinations that rank highly but rely on weak or inconsistent pathway evidence should be flagged for mechanistic validation, while combinations with moderate predicted benefit but high confidence may be preferable when experimental resources are limited.

5. Conclusion

This manuscript has developed an AI-assisted prioritization framework for combination therapies that integrates metabolic pathway interactions with cellular network connectivity. The framework represents drugs, targets, pathways, reactions, and cell contexts in a heterogeneous graph, derives pathway-interaction features from shared reactions, computes connectivity by network propagation, learns relational embeddings with a graph neural

network, and ranks combinations with an uncertainty-adjusted decision rule. The central methodological contribution is the coupling of pathway-interaction features with connectivity propagation, which allows shared metabolic dependencies to modify network-level scores rather than remaining independent covariates. Analytical assessment identifies key boundary conditions, including label noise, degree bias, pathway annotation incompleteness, context mismatch, and metabolic compensation. The framework also specifies baselines, ablation logic, and evaluation standards that distinguish retrospective ranking from prospective validation. By making uncertainty explicit and linking predictions to mechanistic pathway and network quantities, the approach provides a disciplined basis for triaging combination therapies for experimental follow-up.

References

1. Chou, T. C. (2006). Theoretical basis, experimental design, and computerized simulation of synergism and antagonism in drug combination studies. Pharmacological Reviews, 58(3), 621-681. https://doi.org/10.1124/pr.58.3.10

2. 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

3. 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

4. Yıldırım, M. A., Goh, K. I., Cusick, M. E., Barabási, A. L., & Vidal, M. (2007). Drug-target network. Nature Biotechnology, 25(10), 1119-1126. https://doi.org/10.1038/nbt1338

5. Chen, X., Yan, C. C., Zhang, X., & You, Z. H. (2016). Drug-target interaction prediction: databases, web servers and computational models. Briefings in Bioinformatics, 17(4), 696-712. https://doi.org/10.1093/bib/bbv066

6. 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

7. Ma, T., Xiao, C., Zhou, J., & Wang, F. (2018). Drug similarity integration through attentive multi-view graph auto-encoders. In Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI) (pp. 3477-3483). AAAI Press.

8. 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

9. Preuer, K., Lewis, R. P., Hochreiter, S., Bender, A., Bulusu, K. C., & Klambauer, G. (2018). DeepSynergy: predicting anti-cancer drug synergy with Deep Learning. Bioinformatics, 34(9), 1538-1546. https://doi.org/10.1093/bioinformatics/btx806

10. Huang, L., Li, F., Sheng, J., Xia, X., Ma, J., Zhan, M., & Wong, S. T. (2014). DrugComboRanker: drug combination discovery based on target network analysis. Bioinformatics, 30(12), i228-i236. https://doi.org/10.1093/bioinformatics/btu278

11. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In International Conference on Learning Representations (ICLR).

12. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018). Graph attention networks. In International Conference on Learning Representations (ICLR).

13. Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., & Welling, M. (2018). Modeling relational data with graph convolutional networks. In European Semantic Web Conference (pp. 593-607). Springer. https://doi.org/10.1007/978-3-319-93417-4_38

14. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Proceedings of the 33rd International Conference on Machine Learning (pp. 1050-1059). PMLR.

15. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In Advances in Neural Information Processing Systems (Vol. 30, pp. 5574-5584).

16. O'Boyle, N. M., & Sayle, R. A. (2016). Comparing structural fingerprints using a literature-based similarity benchmark. Journal of Cheminformatics, 8, 36. https://doi.org/10.1186/s13321-016-0148-0

17. Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., Assempour, N., Iynkkaran, I., Liu, Y., Maciejewski, A., Gale, N., Wilson, A., Chin, L., Cummings, R., Le, D., ... Wilson, M. (2018). DrugBank 5.0: a major update to the DrugBank database for 2018. Nucleic Acids Research, 46(D1), D1074-D1082. https://doi.org/10.1093/nar/gkx1037

18. Kanehisa, M., & Goto, S. (2000). KEGG: Kyoto encyclopedia of genes and genomes. Nucleic Acids Research, 28(1), 27-30. https://doi.org/10.1093/nar/28.1.27

19. Szklarczyk, D., Gable, A. L., Lyon, D., Junge, A., Wyder, S., Huerta-Cepas, J., Simonovic, M., Doncheva, N. T., Morris, J. H., Bork, P., Jensen, L. J., & von Mering, C. (2019). STRING v11: protein-protein association networks with increased coverage. Nucleic Acids Research, 47(D1), D607-D613. https://doi.org/10.1093/nar/gky1131

20. Regev, A., Teichmann, S. A., Lander, E. S., Amit, I., Benoist, C., Birney, E., Bodenmiller, B., Campbell, P., Carninci, P., Clatworthy, M., Clevers, H., Deplancke, B., Dunham, I., Eberwine, J., Eils, R., Enard, W., Farmer, A., Fugger, L., Göttgens, B., ... Human Cell Atlas Meeting Participants. (2017). The human cell atlas. eLife, 6, e27041. https://doi.org/10.7554/eLife.27041