

## Graph algorithms for predicting subcellular localization at the pathway level

Chris S Magnano<sup>1,2,3</sup> and Anthony Gitter<sup>1,2,4</sup>

<sup>1</sup>*Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI, USA*

<sup>2</sup>*Morgridge Institute for Research, Madison, WI, USA*

<sup>3</sup>*Center for Computational Biomedicine, Harvard Medical School, Boston, MA, USA*

<sup>4</sup>*Department of Biostatistics and Medical Informatics, University of Wisconsin-Madison, Madison, WI, USA*

Protein subcellular localization is an important factor in normal cellular processes and disease. While many protein localization resources treat it as static, protein localization is dynamic and heavily influenced by biological context. Biological pathways are graphs that represent a specific biological context and can be inferred from large-scale data. We develop graph algorithms to predict the localization of all interactions in a biological pathway as an edge-labeling task. We compare a variety of models including graph neural networks, probabilistic graphical models, and discriminative classifiers for predicting localization annotations from curated pathway databases. We also perform a case study where we construct biological pathways and predict localizations of human fibroblasts undergoing viral infection. Pathway localization prediction is a promising approach for integrating publicly available localization data into the analysis of large-scale biological data.

*Keywords:* Probabilistic graphical model, graph neural network, spatial proteomics

### 1. Introduction

Cellular state is dictated by a wide range of factors from chromatin accessibility to protein abundance to the physical location of proteins within the cell. Cells are compartmentalized into subcellular locations that provide the chemical environment around proteins. That local environment informs proteins' structure and available interaction partners. Protein localization not only dictates protein interactions in normal biological processes,<sup>1</sup> but also is an important factor that can contribute to abnormal cellular behavior. Alzheimer's disease, amyotrophic lateral sclerosis, Wilson disease, and multiple cancers involve abnormal protein localizations.<sup>2</sup>

Although protein localization is dynamic and context-specific,<sup>3</sup> many localization resources present a fixed, static view. Localization databases such as MatrixDB,<sup>4</sup> Organelle DB,<sup>5</sup> Compartments,<sup>6</sup> and ComPPI<sup>7</sup> track primary experimental data, computational predictions, or combinations of multiple information sources. Up to 50% of proteins localize to multiple cellular compartments.<sup>8,9</sup> Databases typically provide multiple possible localizations per protein, but that does not determine the conditions under which subsets of each protein's localizations are relevant. Many tools can predict possible locations of a protein based on its sequence<sup>10–12</sup> using machine learning methods such as logistic regression<sup>13</sup> or deep neural networks.<sup>14</sup> Some

methods incorporate additional information, such as gene expression,<sup>15</sup> Gene Ontology annotations,<sup>16</sup> and network information.<sup>17–20</sup> Methods using network information consider the localizations of neighboring proteins in protein-protein interaction databases to aid in localization prediction and do not attempt to represent any particular biological context. Some predictive methods consider tissue context,<sup>21</sup> but proteins vary in their subcellular localization even between single cells of the same tissue type.<sup>1</sup>

We present graph algorithms for estimating context-specific protein localizations by modeling them in biological pathways<sup>a</sup>. Biological pathways, graphs of biological entities such as proteins, can represent a particular biological process or context. Although traditionally thought of in terms of curated pathway databases, pathway reconstruction graph algorithms<sup>22–24</sup> can generate custom pathway representations of a specific process given a background protein interaction network and condition-specific data such as proteomic measurements as input. However, there is no straightforward way to contextualize and apply available protein localization data to this type of predicted biological pathway. In order to provide context-specific localization information for a particular biological dataset, we develop graph algorithms for the simultaneous prediction of a subcellular localization for all interactions in a reconstructed biological pathway. Computationally, this can be seen as an edge labeling task on an existing graph. This predictive step can be added to existing pathway reconstruction workflows. Estimating localization information at the pathway level enables examining where proteins or other biological entities are when they perform a biological function. Pathway-specific localization annotation can help interpret the predicted pathway and potentially provide additional information to guide followup experiments.

Our strategy to understand context-specific protein localization through graph-based annotations of reconstructed pathways offers advantages over alternative approaches. Some curated pathway databases provide localization information at the interaction level and include information about non-protein biological entities.<sup>25,26</sup> However, many pathway databases contain incomplete or no localization information. For instance, of the 8 pathway databases included in Pathway Commons,<sup>27</sup> 2 are fully labeled with localization information, 5 are partially labeled with localization information, and 1 contains no programmatically available localization information. Additionally, curated pathways often do not line up with experimental data<sup>28–31</sup> and a curated pathway may not be available for a particular biological condition of interest. While condition-specific localization information can be experimentally derived<sup>1</sup> using mass spectrometry or cellular imaging, these methods can be expensive, require experimental expertise, and have incomplete coverage. Predicting localization based on pathways is less precise than acquiring localization data experimentally, but the predictions provide an initial coarse estimate of all proteins' localizations without requiring new specialized data.

We develop and compare three categories of methods for predicting localization for interactions within the context of a biological pathway: graph neural networks, probabilistic graphical models, and classifiers that do not use graph topology. First, we quantitatively evaluate these strategies for pathway-based localization prediction by holding out annotated localizations

---

<sup>a</sup>Supplementary Information and code can be found at <https://github.com/gitter-lab/pathway-localization> and archived at <https://doi.org/10.5281/zenodo.7140733>.

from pathway databases. Then, we demonstrate how our approach can be used in practice with a case study involving human cytomegalovirus (HCMV) infection over time.<sup>32</sup> While there are disparities between localization information in pathway databases and experimentally-derived localization data, pathway-level localization prediction is a promising approach for combining publicly available localization data with the analysis of large-scale biological data.

## 2. Methods

### 2.1. Pathway Localization Prediction Problem Definition

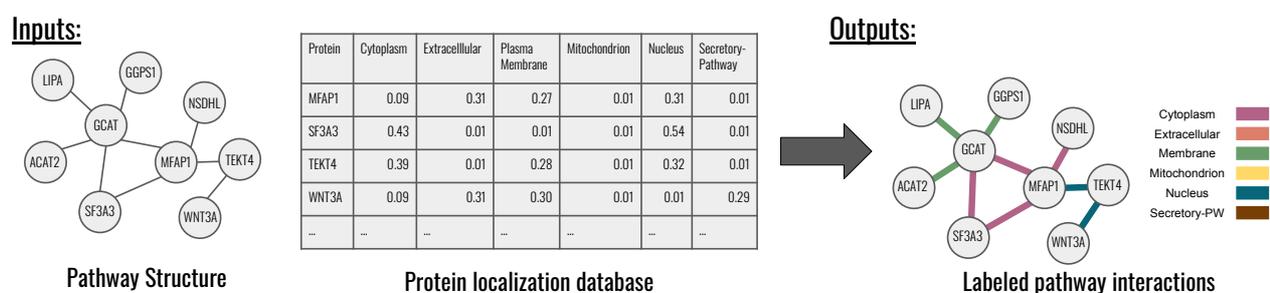


Fig. 1. Overview of the pathway localization prediction experimental workflow.

Given a biological pathway represented as a graph, the goal is to predict one subcellular localization for each edge. The pathway represents some cellular function and can be constructed from large-scale biological datasets using pathway reconstruction.<sup>33</sup> We predict a localization for each edge in the pathway, which can be viewed as a class label assignment for each edge in the graph. Protein-level localization information is used as input to the prediction task as node features. Thus, the pathway-specific subcellular localization task can be defined as:

**Input:** (1) A context-specific pathway graph consisting of nodes and edges  $G = (N, E)$ , and (2) a distribution over possible localizations for each node in the graph. **Output:** A single localization assignment for each interaction  $e \in E$ . See Figure 1.

We chose to assign localizations to edges as opposed to nodes and to assign each interaction a single localization. Pathway databases such as Reactome<sup>25</sup> and popular pathway file formats such as BioPax<sup>34,35</sup> only allow proteins to be in a single subcellular location, creating multiple protein entries if they occur in multiple localizations and assigning them to interactions. While many proteins have multiple localizations, among all Reactome and PathBank pathways less than 5% of total interactions have multiple localizations within the same pathway.

### 2.2. Experimental Setup

#### 2.2.1. Pathway Database Localization Prediction

We investigated how well protein localization databases can be used to predict context-specific localizations in pathway databases, both to examine the feasibility of pathway-specific localization prediction and to elucidate the relationship between node labels in protein localization databases and edge labels in pathway databases. Pathways with interaction localization labels from the Reactome<sup>25</sup> and PathBank<sup>26</sup> databases were each used as ground truth datasets.

The original pathways in both Reactome and PathBank are represented as hypergraphs, where reaction edges can contain more than two nodes. Pathway Commons converts these hypergraphs to graphs using a set of rules<sup>b</sup>. To represent a protein-complex that contains  $n$  proteins, the hypergraph conversions create an edge between every possible pair of nodes, resulting in  $n^2$  edges. For instance, the 4 hyperedges that make up the PathBank pathway Protein Synthesis: Serine are converted to 3,318 edges, of which 3,315 are of type “in-complex-with”. We collapsed protein complexes into single nodes where possible in all pathways. This was done by removing any nodes if all of its edges were redundant with the protein-complex’s edges, leaving a single node for each complex. Though this loses some node information, collapsing protein complexes resulted in pathways that more more closely resembled the original hypergraph in edge distribution, topology, and class balance.

Three different node feature sets were used: the ComPPI database,<sup>7</sup> the Compartments database,<sup>6</sup> and UniProt keyword<sup>36</sup> features. ComPPI and Compartments contain localization scores for each protein, which are used directly as input features. We created a dimensionality reduction-based vectorization of UniProt keyword assignments for all proteins (Section S1.3.3). All 8 predictive models (Section 2.3) were tested on all feature sets with the exception of the NaivePGM model, which could not use the UniProt keyword features as it interprets input features directly as conditional probabilities. All pathways in the 2 pathway databases Reactome and PathBank, which contain interaction-level localization labels, were tested on resulting in a total of 46 runs. Models were trained using 5-fold cross validation, and model selection and hyperparameter selection were performed on a tuning set of the 53 Reactome pathways categorized as developmental and a randomly chosen 10% of all PathBank pathways. Tuning pathways were excluded from cross validation.

### 2.2.2. Human Cytomegalovirus Case Study

To examine how predicting context-specific localization at the pathway level could be used in a realistic setting, we performed a case study with bulk spatial mass spectrometry (MS) data from multi-organelle profiling on primary fibroblasts during HCMV infection.<sup>32</sup> In multi-organelle profiling, gradient centrifugation is used on a bulk sample to partially separate organelles. Protein levels in each subcellular fraction are then measured using tandem mass tags MS, and localization labels are determined by clustering proteins with similar fraction profiles. We investigated whether a predictive model can infer localizations in the context of viral infection, potentially bypassing the need to collect spatial proteomic data.

We performed pathway reconstruction<sup>33</sup> by combining a background protein-protein interaction network<sup>28,37,38</sup> with label-free MS data, which measured protein abundance across the entire fibroblast at 120 hours post infection (hpi) without regards to localization. Measured protein levels were used to create biological networks representing the cell state following infection. The combined top pathways chosen (Section S1.1) contained a total of 386 edges with localization information at 120hpi.

We then trained one of the best performing models from the pathway database prediction

<sup>b</sup><http://www.pathwaycommons.org/pc2/formats>

task, the graph attention network, in three different scenarios. First, we trained a model using data from the PathBank database as described in Section 2.1. Second, we trained a model using a separate dataset that measured protein localization using a similar method on a different cell type and under a different biological condition, HeLa cells undergoing EGF stimulation.<sup>39</sup> Third, we trained a model on the same HCMV experiment at the 24hpi timepoint. This third scenario is unlikely to occur, as it would require a dataset to already exist for an identical cell type and condition, but gives a useful benchmark for best case predictive performance.

### 2.3. Pathway Localization Prediction Models

We evaluated three general categories of models (Section S1.2): general classifiers,<sup>40</sup> probabilistic graphical models, and graph neural networks (Figure 2). The fully-connected neural network (FullyConnectedNN), random forest (RF), and logistic regression (Logit) served as baseline classifiers because they use no topological information from the pathway graph (Figure S1). These models instead concatenate the node features of each interaction’s endpoints as their input. All other models use topological information from the pathway graph to encourage interactions near each other to have similar localizations.

**Graph convolutional network (GCN):** Graph convolutional networks<sup>41</sup> incorporate a series of message-passing convolutional layers before the final fully connected layers. The convolutional layers allow for information to be shared across the topology of the input network, providing a first-order approximation of spectral graph convolutions.<sup>42</sup> All neural network models were implemented using PyTorch Geometric.<sup>43</sup>

**Graph attention network (GAT):** Graph attention networks extend graph convolutional networks by allowing each node to choose which neighbors to pay attention to. As opposed to taking the average of its neighbors, each node computes a weighted average of its neighbors in graph convolutional layers.<sup>44,45</sup> The GAT is multi-headed, where multiple attention weights are computed in parallel for each node. The number of heads is a hyperparameter.

**Graph isomorphism network (GIN):** Graph isomorphism networks<sup>46</sup> take advantage of the similarity between neighbor aggregation in graph neural networks and the Weisfeiler-Lehman (WL) graph isomorphism test.<sup>47</sup> The WL graph isomorphism test is a heuristic algorithm for determining graph isomorphisms. The neighbor aggregation in each graph layer of a graph isomorphism network is formulated to be at least as powerful as the WL isomorphism test; the  $l^{th}$  layer is guaranteed to generate different embeddings of two graphs if those graphs would be found to be non-isomorphic via the WL isomorphism test in  $l$  iterations.

**Probabilistic graphical models:** Given the nature of the label propagation inherent in the pathway level localization prediction task, and that many localization databases provide scores or even probabilities, probabilistic graphical models are a natural choice. However, these models only provide predictions on the nodes of the graph, while we are interested in localization labels on the edges. To convert the input pathway into an appropriate graphical model, each pathway is converted into a bipartite graph, where an additional node is added to that graph for each edge (Figure S2).

Probabilistic graphical models represent a set of  $N$  random variables  $\mathbf{y}$  as nodes and dependencies between them as a set of edges  $E$ . We created two pairwise undirected probabilistic

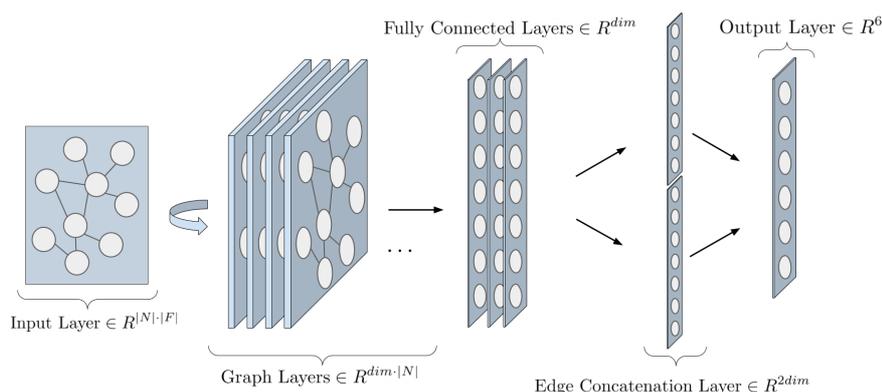


Fig. 2. Overview of neural network architecture for graph neural networks. The number of graph layers (convolutional depth) and number of fully connected layers (linear depth) are hyperparameters.  $|N|$  is the number of nodes in the input pathway.  $|F|$  is the number of input features for each node.

graphical models,<sup>48</sup> which we call NaivePGM and TrainedPGM. In these probabilistic graphical models the random variables obey a local Markov property, such that each random variable is conditionally independent of all others given its neighbors in the graph.

The NaivePGM is a Markov random field, where protein localization database data is used to create conditional probability tables. In the TrainedPGM, input features are treated as observations of additional variables to train potential functions on each node. These potential functions are represented by discriminative classifiers,<sup>49</sup> here random forests. This type of model is referred to as a discriminative random field.<sup>50</sup> This was chosen over a more traditional log linear parameterization due to better performance on the tuning data.

We performed 30 iterations of hyperparameter selection via Bayesian optimization<sup>51</sup> using Ax for neural network models and Scikit-optimize for classifier models<sup>c</sup> (Tables S1 and S2).

### 3. Results

#### 3.1. Comparing Pathway and Localization Databases

To better understand the feasibility of predicting interaction localizations from protein-level localization data, we compared the edge localizations present in biological pathway databases to node localizations in protein localization databases. The Reactome and PathBank pathway databases significantly disagree with both protein localization databases. For instance, among all proteins with an edge localized to the membrane in Reactome, ComPPI scores more as being in the cytosol than in the membrane. In all cases there is a wide distribution when stratifying the ComPPI node scores used as features by the Reactome or PathBank edge localizations used as labels (Figures S3 and S4). Therefore, for any individual protein and interaction there is a significant chance that protein's most likely localization according to ComPPI or Compartments is not the localization Reactome or PathBank assigned it to.

Directly using data from protein localization databases is not sufficient to accurately pre-

<sup>c</sup><https://ax.dev/> and <https://scikit-optimize.github.io/stable/>

dict pathway level localization. Many interactions have at least one contradictory interaction with an identical featurization but a different localization label, over 40% when using ComPPI and over 20% when using Compartments. In addition, many interaction localizations would be considered impossible when using a protein localization database alone. Almost 14% of interactions in Reactome are between proteins that have no protein localizations in common in ComPPI. Even without featurization, for 9.5% and 11.5% of total interactions in Reactome and PathBank, respectively, there exists another interaction between the same unique proteins in another pathway that has a different localization. This indicates that pathway topology or some other form of additional information beyond that of individual proteins is needed to correctly predict localization in context.

### 3.2. Pathway Database Localization Prediction

We used cross-validation to train our models on protein information and some labeled database pathways and evaluate their edge localization predictions for other database pathways given only protein information and graph structure as input. Overall, models were able to achieve better interaction localization prediction performance on PathBank pathways (Figure 3) than Reactome pathways (Figure 4). Generally, models' performance in predicting PathBank interaction localizations was more consistent across pathways. However, on both datasets all models' performance had high variance across pathways. Except for logistic regression, all models got at least some pathways completely correct and some pathways completely wrong across all databases and feature sets. The graph neural network models, GCN, GAT, and GIN, generally outperformed other models in all conditions. However, in Reactome no model was able to achieve a median multiclass F1 score (hereafter called 'F1 score') of over 0.5

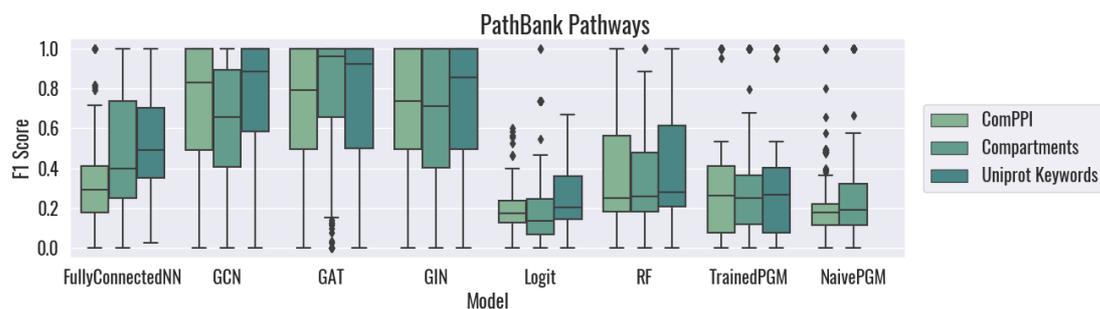


Fig. 3. Multiclass F1 score of predictive performance on PathBank localizations across all 427 considered PathBank pathways. Scores are calculated per pathway; the distribution of scores is shown for each model.

Probabilistic graphical models and models that used no pathway topology had generally comparable performance. The FullyConnectedNN model was able to outperform other models when predicting PathBank localizations using Compartments or UniProt keyword features. It should be noted, however, that when calculating performance by pathway as done in this setting, the size of each pathway is not taken into account. This means that edges in very small pathways can have an outsized effect on total performance.

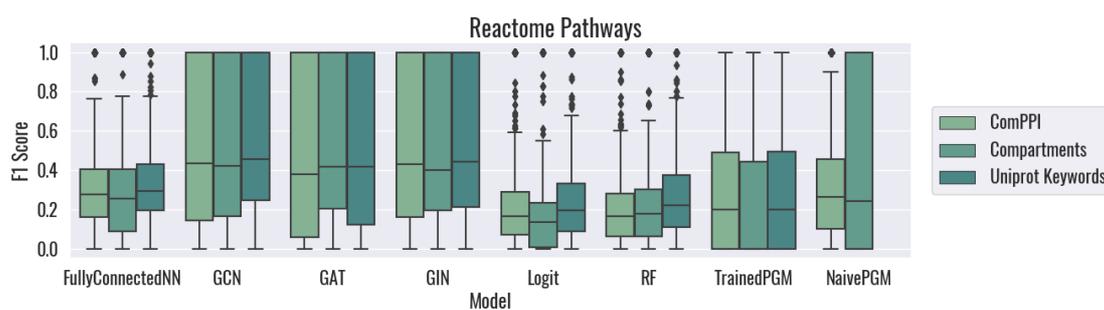


Fig. 4. Multiclass F1 score of predictive performance on Reactome localizations across all 918 considered Reactome pathways. Scores are calculated per pathway; the distribution of scores is shown for each model.

Alternatively, Figures S5 and S6 show F1 scores for each model aggregated from all pathways, where all edges are used for a single performance calculation. When aggregated in this way, all non-neural network models perform comparably. The probabilistic graphical models, and the TrainedPGM model in particular, struggled with small pathways.

The number of real and predicted unique localizations in each pathway also had a large effect on model performance. This can be thought of as the smoothness of the real or predicted localizations in a pathway, or how strong the tendency is for edges nearby in a pathway to have the same localization. Ideally, a model would be able to detect that a pathway exists entirely in a single localization and aggressively smooth its localization predictions over the pathway. Pathways with a single localization had the widest range of performance within each model. More extreme performances, at or nearly at 1.0 or 0.0 for these pathways, indicate that the model correctly predicted that the pathway had only a single localization. Figure S7 shows the distributions of the number of predicted unique localizations by the different models.

### 3.3. HCMV Infection Spatial Proteomics Case Study

We considered three scenarios for evaluating localization prediction in an experimental setting. Here, we examine if localizations can be inferred in the context of a HCMV infection (Section S1.1). We simulate an exploratory workflow by first constructing HCMV infection-specific biological pathways using pathway reconstruction<sup>23</sup> (example pathway topologies can be viewed in Figures S8 and S9). We then use the context provided by these pathways' topologies to predict interaction localizations with the best performing model from pathway database prediction, GAT, using node features from the Compartments database.

In all scenarios, we predict localizations for each interaction of pathways created from protein abundance measurements at 120hpi. Localization data from spatial MS taken at the same timepoint was used as ground truth. Each scenario differs in the labeled training data used: pathways from a pathway database, a different experiment using a different context and cell type, or data from the same experiment at a different timepoint. In all scenarios, all data from the 120hpi timepoint was held out until the final evaluation. We also consider a baseline model that always predicts the most frequent localization among all training set interactions.

While in all scenarios the model substantially outperformed the baseline, there was a large

gap in performance between the model trained using pathway databases versus those trained on a different experiment (Figure 5). Both scenarios using experimental data achieved an F1 score of over 0.8. Although the GAT model predictions do not perfectly recapitulate the spatial proteomics localizations, it is encouraging that the GAT model trained in a plausible setting with data from an unrelated biological context is almost as accurate as the unrealistic, best case GAT model trained on another timepoint from the same HCMV infection experiment.

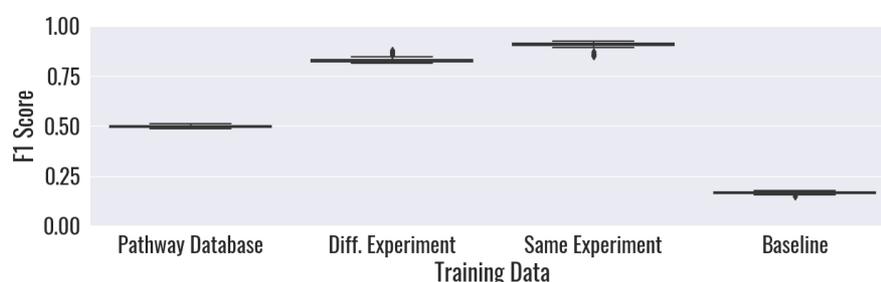


Fig. 5. Multiclass F1 score of the GAT model on spatial MS data of viral infection at 120hpi. Performance is shown in each scenario for the 50 top pathways created from a parameter sweep. The baseline model always predicts the most common localization in the training dataset.

#### 4. Conclusions and Future Work

Although there is some correspondence between protein localization databases and localization data in pathway databases, these two types of localization data generally disagree. Graph neural network models were required to achieve high predictive performance on PathBank localizations, and all models performed poorly in predicting Reactome localizations.

There are a number of possible reasons for this misalignment between localization information in pathway databases and protein localization databases. While the best-performing models include topological information, implying that topology is needed to bring context to protein localization, it is possible that other types of data are needed. Protein features derived from UniProt keywords only slightly improved performance, and tissue- or cell-specific localization may be necessary to fully realize context-specific localization. That type of information may not be available for pathway databases, which are often provided independent of tissue type, but could be for reconstructed pathways. The protein localization databases may also be too noisy and general for context-specific localization prediction. While some signal does exist, the wide range of distributions for ComPPI and Compartments scores across different pathway localizations highlights the imprecise nature of the prediction problem.

While graph neural networks outperformed other methods in predicting pathway localizations, it is unclear how large a role pathway topology played in these methods' performance. It is possible that increased performance over other models comes solely from how graph convolutions share information between nodes, as opposed to the biological information inherent in each pathway's topology aiding localization prediction.

The conversion of pathways from hypergraphs to graphs greatly impacted the class distribution and topology of Reactome and PathBank pathways. Treatment of protein complexes

can lead to orders of magnitude difference in the number of edges in the resultant pathways. We created protein complex nodes to represent complexes, which removes node information but better preserves the edge structure and balance in the pathway. An analysis task focused specifically on nodes may want a conversion that better preserves node information at the possible cost of edge information. Important future work would be to consider these conversions in a more systemic way and quantify the hypergraph properties they alter or keep invariant.

Pathway reconstruction has already proven to be a powerful strategy for interpreting transcriptomic, proteomic, or other data in a network context, and the ability to coarsely approximate interaction localizations could further increase its value. We observed the GAT model may have sufficient accuracy to roughly estimate such pathway localizations as long as it is trained on experimental data instead of pathway databases. Predictions using the model trained on HeLa cells still had an error rate of approximately 17% but could plausibly be used to obtain an estimate of context-specific localization predictions in the absence of other data. Further testing is required to assess how similar the training conditions and assay types must be to the test conditions and assays and what types of pathway reconstruction algorithms are compatible with our GAT localization prediction model.

There are additional biological contexts where localization prediction could prove valuable. Single-cell spatial proteomics experiments have previously found proteins to vary by as much as 16% in either expression or spatial distribution between cells undergoing the same process in the same tissues.<sup>8</sup> Predicted protein localizations for individual cells could add an additional layer of information in single-cell analyses. Additionally, targeted identification of abnormal protein localizations could provide insight in diseases where protein localization is known to play a role.<sup>52</sup> The current predictive method could be expanded to attempt to quantify a localization being unexpected given a constructed pathway representing some cellular state.

## References

1. E. Lundberg and G. H. H. Borner. Spatial proteomics: A powerful discovery tool for cell biology. *Nature Reviews Molecular Cell Biology*, 20(5):285–302, May 2019.
2. M.-C. Hung and W. Link. Protein localization in disease and therapy. *Journal of Cell Science*, 124(20):3381, October 2011.
3. N. C. Bauer et al. Mechanisms regulating protein localization. *Traffic*, 16(10):1039–1061, 2015.
4. E. Chautard et al. MatrixDB, the extracellular matrix interaction database. *Nucleic Acids Research*, 39(suppl\_1):D235–D240, September 2010.
5. N. Wiwatwattana and A. Kumar. Organelle DB: a cross-species database of protein localization and function. *Nucleic Acids Research*, 33(suppl\_1):D598–D604, January 2005.
6. J. X. Binder et al. COMPARTMENTS: unification and visualization of protein subcellular localization evidence. *Database*, 2014(bau012), February 2014.
7. D. V. Veres et al. ComPPI: a cellular compartment-specific database for protein-protein interaction network analysis. *Nucleic acids research*, 43(Database issue):D485–D493, January 2015.
8. P. J. Thul et al. A subcellular map of the human proteome. *Science*, 356(6340):eaal3321, May 2017.
9. S. Zhang et al. DBMLoc: A Database of proteins with multiple subcellular localizations. *BMC Bioinformatics*, 9(1):127, February 2008.
10. J. L. Gardy and F. S. L. Brinkman. Methods for predicting bacterial protein subcellular localization. *Nature Reviews Microbiology*, 4(10):741–751, October 2006.

11. K. Imai and K. Nakai. Prediction of subcellular locations of proteins: Where to proceed? *PROTEOMICS*, 10(22):3970–3983, 2010.
12. A. Alaa et al. Protein Subcellular Localization Prediction Based on Internal Micro-similarities of Markov Chains. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 1355–1358, July 2019.
13. S. Hua and Z. Sun. Support vector machine approach for protein subcellular localization prediction. *Bioinformatics*, 17(8):721–728, August 2001.
14. J. J. Almagro Armenteros et al. DeepLoc: prediction of protein subcellular localization using deep learning. *Bioinformatics*, 33(21):3387–3395, July 2017.
15. A. Drawid and M. Gerstein. A Bayesian system integrating expression data with sequence patterns for localizing proteins: comprehensive application to the yeast genome. *Journal of Molecular Biology*, 301(4):1059–1075, August 2000.
16. A. Fyshe et al. Improving subcellular localization prediction using text classification and the Gene Ontology. *Bioinformatics*, 24(21):2512–2517, August 2008.
17. M. M. Ananda and J. Hu. NetLoc: Network based protein localization prediction using protein-protein interaction and co-expression networks. In *2010 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 142–148. IEEE, December 2010.
18. P. Du and L. Wang. Predicting Human Protein Subcellular Locations by the Ensemble of Multiple Predictors via Protein-Protein Interaction Network with Edge Clustering Coefficients. *PLOS ONE*, 9(1):e86879, January 2014.
19. H. S. Garapati et al. Predicting subcellular localization of proteins using protein-protein interaction data. *Genomics*, 112(3):2361–2368, May 2020.
20. A. Grover and L. Gatto. ProtFinder: finding subcellular locations of proteins using protein interaction networks. *bioRxiv*, 2022.
21. L. Zhu et al. Tissue-Specific Subcellular Localization Prediction Using Multi-Label Markov Random Fields. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(5):1471–1482, September 2019.
22. A. Ritz et al. Pathways on demand: automated reconstruction of human signaling networks. *npj Systems Biology and Applications*, 2(1):1–9, March 2016.
23. N. Tuncbag et al. Network-Based Interpretation of Diverse High-Throughput Datasets through the Omics Integrator Software Package. *PLOS Computational Biology*, 12(4):e1004879, April 2016.
24. E. Cerami et al. Automated Network Analysis Identifies Core Pathways in Glioblastoma. *PLOS ONE*, 5(2):e8918, February 2010.
25. A. Fabregat et al. The Reactome Pathway Knowledgebase. *Nucleic Acids Research*, 46(D1):D649–D655, 2018.
26. D. S. Wishart et al. PathBank: a comprehensive pathway database for model organisms. *Nucleic Acids Research*, 48(D1):D470–D478, 10 2019.
27. I. Rodchenkov et al. Pathway Commons 2019 Update: integration, analysis and exploration of pathway data. *Nucleic Acids Research*, 48(D1):D489–D497, 10 2019.
28. A. S. Köksal et al. Synthesizing Signaling Pathways from Temporal Phosphoproteomic Data. *Cell Reports*, 24(13):3607–3618, September 2018.
29. L. Cao et al. Quantitative Phosphoproteomics Reveals SLP-76 Dependent Regulation of PAG and Src Family Kinases in T Cells. *PLOS ONE*, 7(10):e46725, October 2012.
30. S. J. Humphrey et al. High-throughput phosphoproteomics reveals in vivo insulin signaling dynamics. *Nature Biotechnology*, 33(9):990–995, September 2015.
31. R. C. J. D’Souza et al. Time-resolved dissection of early phosphoproteome and ensuing proteome changes in response to TGF-beta. *Science Signaling*, 7(335):rs5, 2014.
32. P. M. Jean Beltran et al. A Portrait of the Human Organelle Proteome In Space and Time during

- Cytomegalovirus Infection. *Cell Systems*, 3(4):361–373.e6, October 2016.
33. C. S. Magnano and A. Gitter. Automating parameter selection to avoid implausible biological pathway models. *npj Systems Biology and Applications*, 7(1):1–12, 2021.
  34. E. Demir et al. The BioPAX community standard for pathway data sharing. *Nature Biotechnology*, 28(9):935–942, September 2010.
  35. B. M. Gyori and C. T. Hoyt. PyBioPAX: biological pathway exchange in Python. *Journal of Open Source Software*, 7(71):4136, March 2022.
  36. The UniProt Consortium. UniProt: the universal protein knowledgebase in 2021. *Nucleic Acids Research*, 49(D1):D480–D489, 11 2020.
  37. S. Razick et al. iRefIndex: A consolidated protein interaction database with provenance. *BMC Bioinformatics*, 9(1):405, September 2008.
  38. P. V. Hornbeck et al. PhosphoSitePlus, 2014: mutations, PTMs and recalibrations. *Nucleic Acids Research*, 43(D1):D512–D520, December 2014.
  39. D. N. Itzhak et al. Global, quantitative and dynamic mapping of protein subcellular localization. *eLife*, 5:e16950, June 2016.
  40. F. Pedregosa et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
  41. T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
  42. D. K. Hammond et al. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
  43. M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
  44. P. Veličković et al. Graph Attention Networks. *International Conference on Learning Representations*, 2018.
  45. S. Brody et al. How attentive are graph attention networks? *arXiv:2105.14491*, 2021.
  46. K. Xu et al. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
  47. B. Weisfeiler and A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Tekhnicheskaya Informatsia*, 2(9), 1968.
  48. U. B. Gwali and S. T. Monteiro. A tutorial on modelling and inference in undirected graphical models for hyperspectral image analysis. *International Journal of Remote Sensing*, 39(20):7104–7143, 2018.
  49. S. Kosov. *Multi-layer conditional random fields for revealing unobserved entities*. PhD thesis, Universität Siegen, 2018.
  50. S. Kumar and M. Hebert. Discriminative random fields. *International Journal of Computer Vision*, 68(2):179–201, 2006.
  51. M. Balandat et al. Botorch: A framework for efficient Monte-Carlo Bayesian optimization. In H. Larochelle et al., editors, *Advances in Neural Information Processing Systems*, volume 33, pp. 21524–21538. Curran Associates, Inc., 2020.
  52. K. E. Blise et al. Single-cell spatial architectures associated with clinical outcome in head and neck squamous cell carcinoma. *npj Precision Oncology*, 6(1):1–14, February 2022.

## 5. Acknowledgements

This work was supported by NIH award T15LM007359, NSF award DBI 1553206, the Morgridge Institute for Research, and the University of Wisconsin–Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation. We thank Sushmita Roy for her valuable feedback.