
PathLinker Cytoscape Documentation

Release 1.3

Jeff Law

Jan 19, 2018

1	PathLinker Cytoscape	3
1.1	Description	3
1.2	Overview	3
1.3	Running the App	5
1.4	Inputs	6
1.4.1	Network	6
1.4.2	Sources/Targets	6
1.4.3	Algorithm	6
1.4.4	Edge weights	6
1.5	Outputs	7
1.6	Example	7
1.6.1	Load Network	7
1.6.2	Sources and Targets	9
1.6.3	Edge weight options	13
1.6.4	Results	13
1.6.5	Edge penalty	15
1.6.6	Include tied paths	16
1.7	Errors and Warnings	17
1.7.1	Sources or Targets not in network	18
1.7.2	Invalid input for parameters	19
1.7.3	Not all edges have weight	19
1.8	Cytoscape Automation / cyREST API	20
1.9	How to Cite PathLinker	20
1.10	Contact Information	20

Table of Contents

PathLinker Cytoscape

Welcome to the documentation for the PathLinker Cytoscape App.

Helpful links:

- [PathLinker app on the Cytoscape App Store](#)
- [PathLinker-Cytoscape GitHub repository](#)
- [PathLinker Cytoscape Automation example Jupyter Notebook](#)

1.1 Description

This app implements the PathLinker algorithm for reconstructing signalling pathways from protein interaction networks².

1.2 Overview

PathLinker is a method for reconstructing signalling pathways from protein interaction and regulatory networks. Consider a protein-protein interaction network represented as a directed graph. The network may contain physical, signalling, and regulatory interactions between pairs of proteins. Given a query composed of a set of receptors and a set of transcription factors in this network that represent the "sources" and "targets" of some specific signalling pathway (e.g., the Wnt pathway), PathLinker attempts to compute a high quality reconstruction of this pathway. Specifically, PathLinker computes a sub-network that connects the receptors in the query to the transcription factors in the query. More generally, PathLinker will be useful for connecting any set of sources to any set of targets in an interaction network.

The primary algorithmic component of PathLinker is computing the k -shortest simple paths in the network from any source to any target. PathLinker accomplishes this task through a novel integration of Yen's algorithm with the A* heuristic, which allows very efficient computation for very large k values, e.g., 10,000, on networks with hundreds of thousands of edges. PathLinker outputs the sub-network composed of the k shortest paths.

PathLinker supports several different options for weighting the edges in the network. If the edges in the input graph do not have weights, PathLinker computes the k lowest cost paths, where the cost of a path is the number of edges in it. If the input graph has edge weights, they can be treated additively or as probabilities. If the edge weights are additive, PathLinker defines the cost of a path as the sum of the edge weights and computes the k paths of lowest cost. In the case of probabilities, PathLinker treats the edge weights as multiplicative and computes the k highest cost paths, where the cost of a path is the product of the edge weights.

See the publications referenced below for a formal description of the method, comprehensive evaluations and comparisons to other approaches, and experimental validation of results.

1.3 Running the App

Control Panel

Network Style Select PathLinker

PathLinker

Help About

Sources/Targets

Select network:

Sources separated by spaces (e.g., S1 S2 S3)
Must match the 'name' column in the Node Table

Select sources in the network or enter text manually

Add selected source(s)

Targets separated by spaces (e.g., T1 T2 T3)

Select targets in the network or enter text manually

Add selected target(s)

Allow sources and targets in paths

Connect sources to each other

Clear

Algorithm

k (# of paths):

Edge penalty:

Include tied paths

Edge Weights

Unweighted

Weights are additive

Weights are probabilities

Edge weight column:

Submit Close

To run the app, fill in the inputs (as described below) and press the submit button. Required inputs: 1) a selected network, 2) at least 1 source and target, 3) a value for k.

1.4 Inputs

1.4.1 Network

PathLinker runs on the currently selected network. The *Select network* dropdown box shows the available networks in Cytoscape and can be used to select a network.

1.4.2 Sources/Targets

The user can select nodes directly from the network and then use the buttons 'Add selected sources' and 'Add selected targets' to add the selected nodes names to the corresponding text field. Node names can also be entered manually, separated by spaces, but they must match the 'name' column of the Node table. If there are sources or targets named that do not exist in the network, PathLinker will warn the user.

There are two options available options here:

- **Allow sources and targets in paths:** Normally, PathLinker removes incoming edges to sources and outgoing edges from targets before computing paths. If the user selects this option, PathLinker will not remove these edges. Therefore, source and target nodes can appear as intermediate nodes in paths computed by PathLinker.
- **Connect sources to each other:** If the user selects this option, PathLinker will copy the sources to the targets field. This option allows the user to compute a sub-network of paths connecting any of the sources/targets to each other. If the user selects this option, then PathLinker will automatically allow sources and targets to appear in paths, i.e., the previous option is also selected. Note that since PathLinker computes loopless paths, if the user inputs only a single node and selects this option, PathLinker will not compute any paths at all.

1.4.3 Algorithm

There are three parameters here:

- **k :** PathLinker takes in the number of paths the user wants. If an invalid value is input for k (e.g., a negative number or a non-integer), PathLinker will compute the default $k = 200$ paths.
- **Edge penalty:** PathLinker can apply an edge penalty when using the additive or multiplicative edge weight options. The larger the value of edge penalty, the less likely it is that longer paths will appear in the results before low-weight shorter paths.
 - **Weights are additive:** penalize each path (i.e. add to the total path score) by a factor of (number of edges in the path)*(edge penalty). Default is 0. Numbers > 0 are allowed.
 - **Weights are probabilities:** penalize each path (i.e. add to the total path score) by a factor of (number of edges in the path)^(edge penalty). Default is 1. Numbers ≥ 1 are allowed.
- **Include tied paths:** If this option is selected, PathLinker will output the k lowest cost paths, as well as all paths "tied" (i.e., equal path cost) with the k th path.

1.4.4 Edge weights

There are three options for the edge weights:

- **Unweighted:** PathLinker will compute the k lowest cost paths, where the cost is the number of edges in the path
- **Weights are additive:** PathLinker will compute the k lowest cost paths, where the cost is the sum of the edge weights

- **Weights are probabilities:** PathLinker will compute the k highest cost paths, where the cost is the product of the edge weights

When a weighted option is selected, the user must select which column to use as the edge weight using Edge weight column drop-down. All columns in the Edge Table that are of the type Double, Integer, Float, and Long can be selected.

1.5 Outputs

PathLinker generates three outputs:

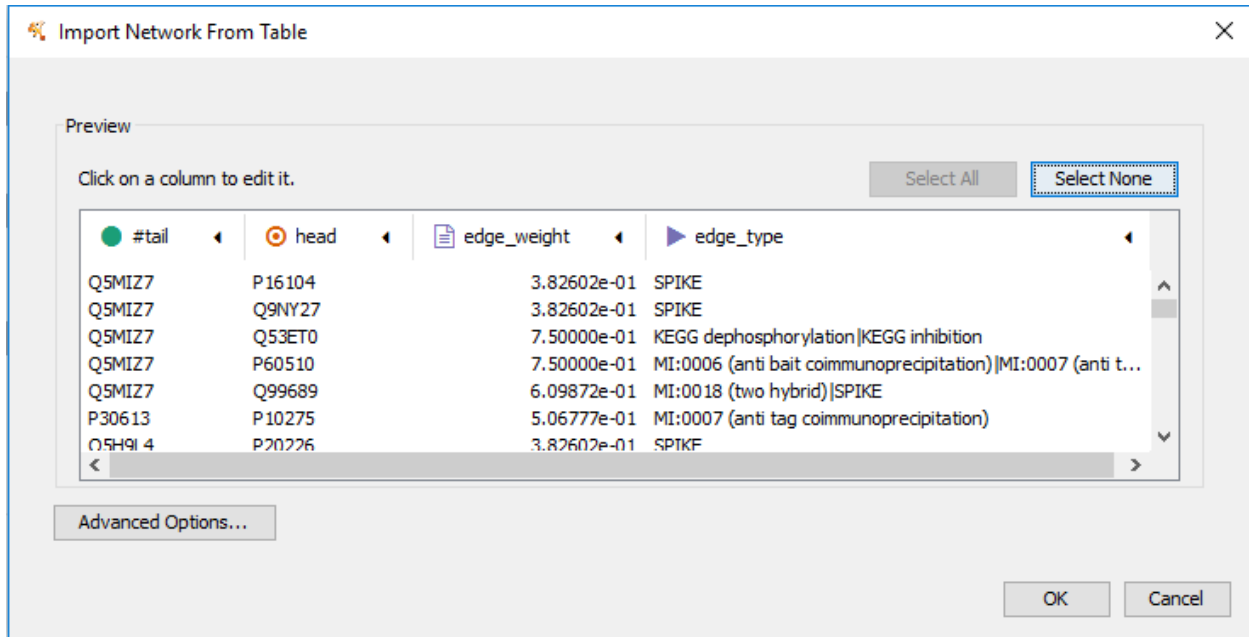
1. **PathLinker-results-X-k-paths:** A sub-network consisting of the nodes and edges involved in the k paths.
2. **PathLinker Results X:** A results panel that displays a table of the rank, score, and nodes of the k paths.
 - Selecting one or more paths in results table will select or highlight those paths in the sub-network.
 - The results can be exported to a TSV (tab separated file).
 - All three of these outputs can be deleted using the Discard button.
3. **Path Rank X:** A column in the Edge Table whose value is the rank of the first path in which a given edge appears.

1.6 Example

Here I go through an example of how to load a network, run PathLinker on it using a set of sources and targets, and look at the results.

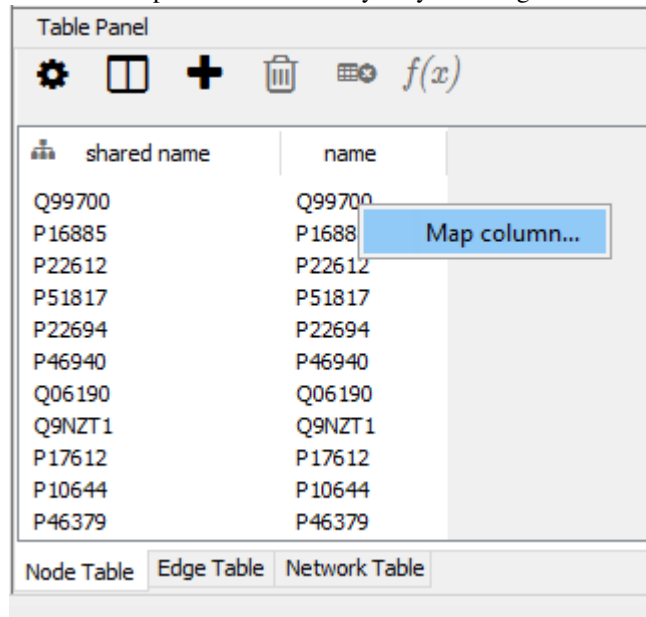
1.6.1 Load Network

To load a network, go to file->import->network->from file. In this example, I downloaded the protein-protein interactome used by Ritz *et al.* [2] from here: <http://bioinformatics.cs.vt.edu/~murali/supplements/2016-sys-bio-applications-pathlinker/>.

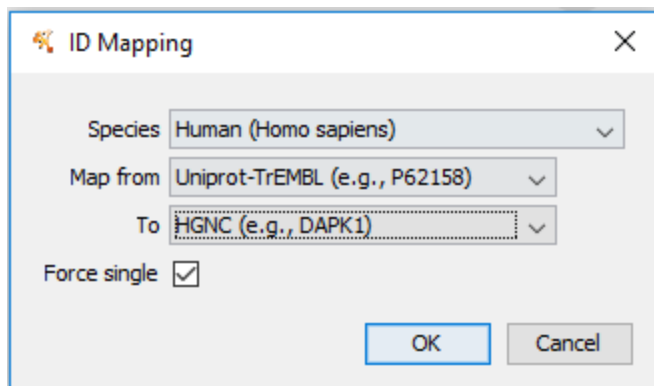


Specify the source and target columns (tail and head columns). Note that Cytoscape assumes the edges are directed, which they are in our case.

I don't know about you, but I don't think using gene names not UniProt IDs! Fortunately cytoscape has various methods to add different namespaces such as gene symbols for uniprot IDs. One easy way is to right click on



the top of a column and use the "Map column" tool.



This has options to automatically map from uniprot to HGNC (Gene Symbol). The screenshots below show the results of the mapping.

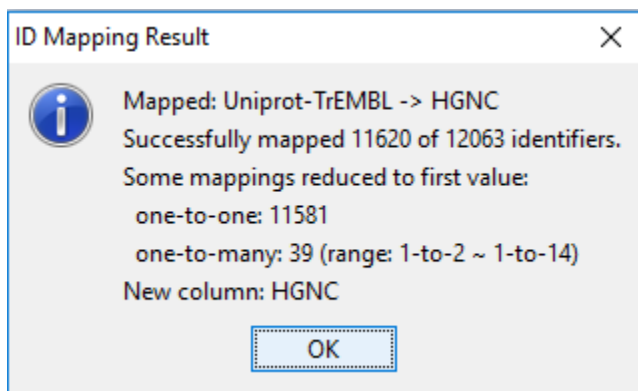
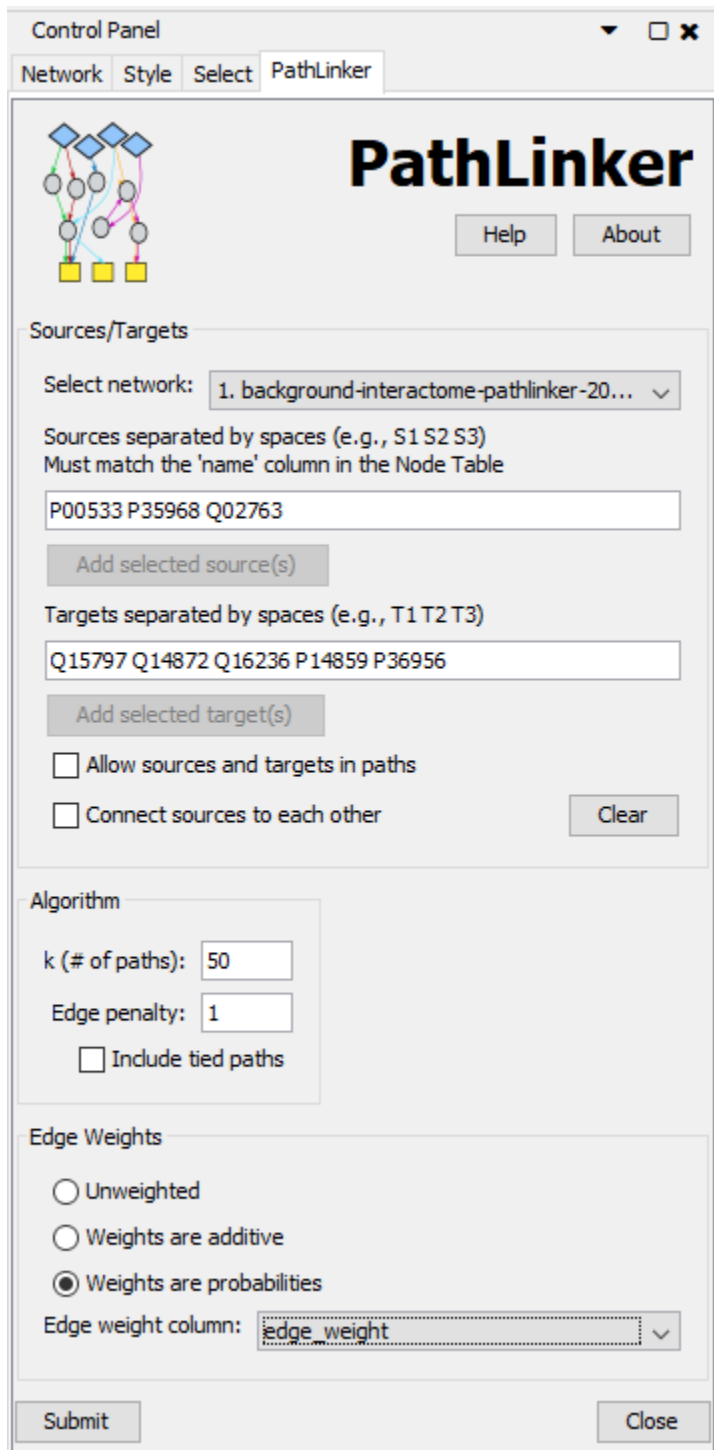


Table Panel		
shared name	name	HGNC
Q99700	Q99700	ATXN2
P16885	P16885	PLCG2
P22612	P22612	PRKACG
P51817	P51817	PRKX
P22694	P22694	PRKACB
P46940	P46940	IQGAP1
Q06190	Q06190	PPP2R3A
Q9NZT1	Q9NZT1	CALML5
P17612	P17612	PRKACA
P10644	P10644	PRKAR1A
P46379	P46379	BAG6

Node Table Edge Table Network Table

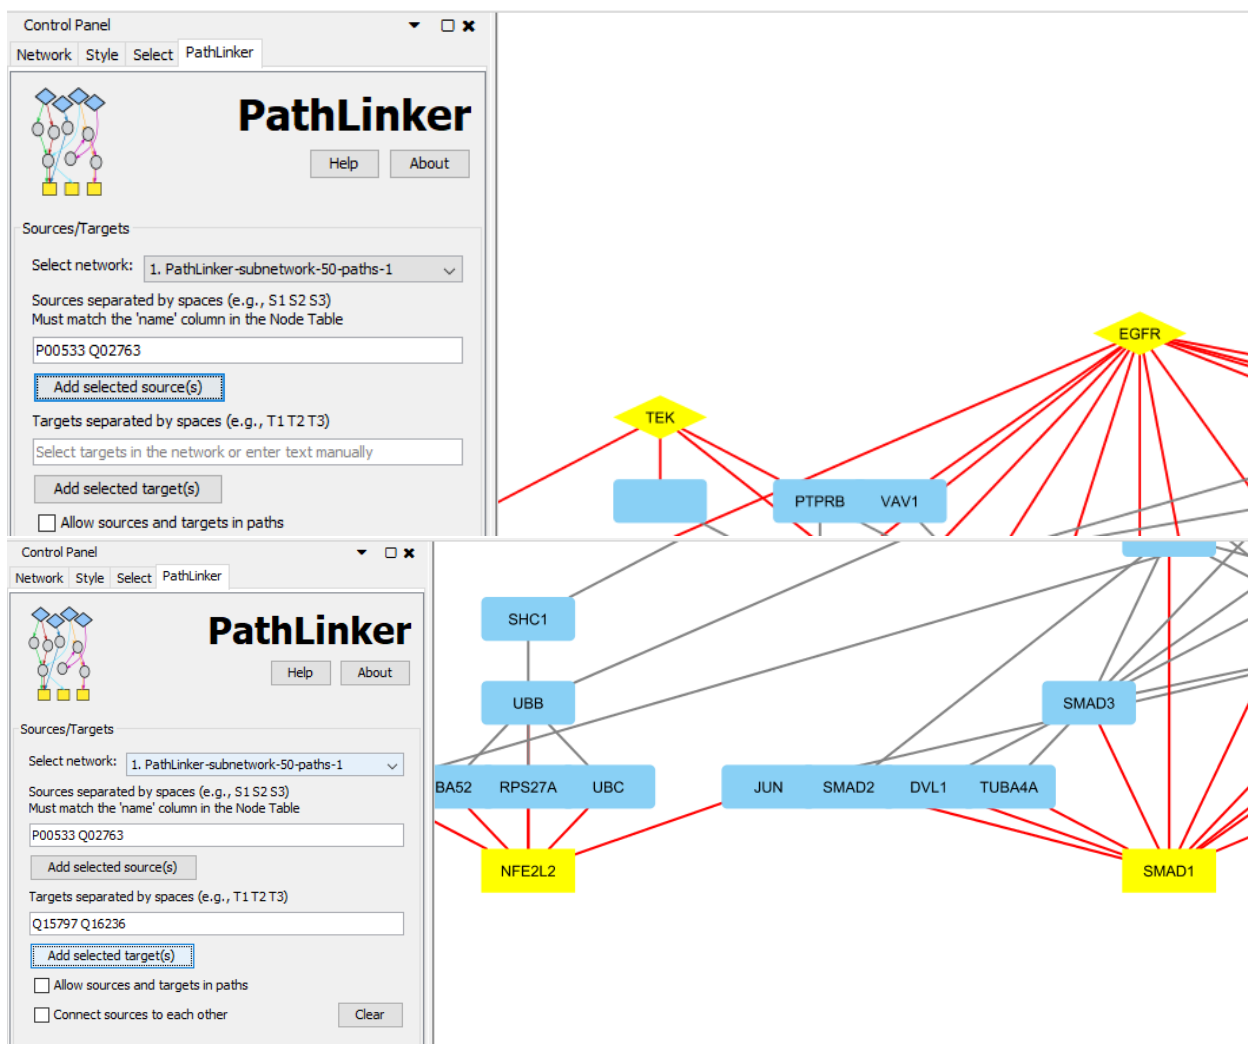
1.6.2 Sources and Targets

PathLinker computes the k shortest or highest scoring paths from a set of sources to a set of targets in an underlying network. In our app paper ([The PathLinker app: Connect the dots in protein interaction networks](#)) as an example use case for PathLinker, we use the receptors and TFs perturbed by drug lovastatin in the ToxCast dataset as the sources and targets. In the ToxCast data, lovastatin perturbed three receptors (EGFR, KDR and TEK) and five TFs (MTF1, NFE2L2, POU2F1, SMAD1 and SREBF1). Because our network uses UniProt IDs as the name of the node, we have to use the UniProt IDs which are: P35968 P00533 Q02763 and Q15797 Q14872 Q16236 P14859 P36956.



Select Sources and Targets from Network

If you have a network view open, instead of manually typing node names into the sources and targets field, you can select nodes from the network and use the "Add selected source(s)" and "Add selected target(s)" buttons to enter the sources and targets. Below I have screenshots of selecting and adding the EGFR and TEK nodes to the sources field, and NFE2L2 and SMAD1 to the targets field. Notice the UniProt IDs are added because, those are the IDs in the 'name' column of the Node table.



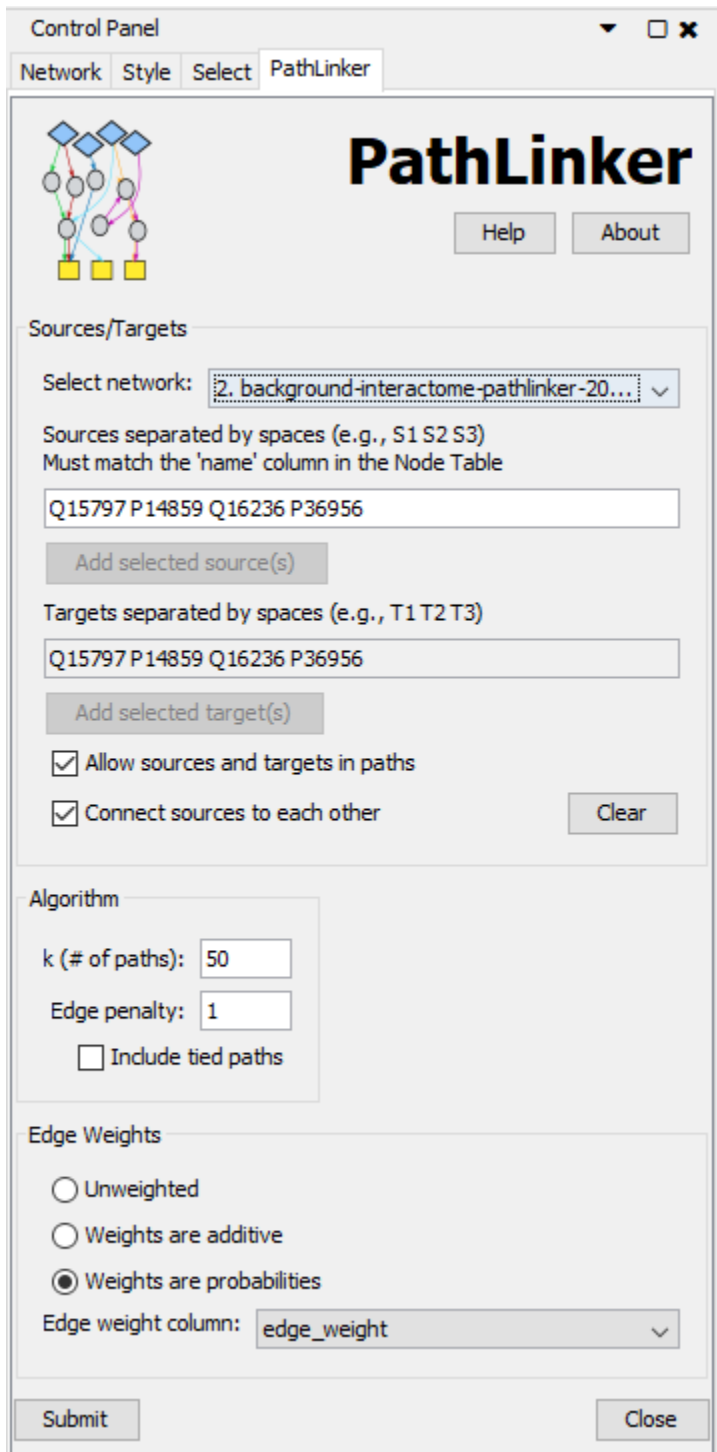
Allow sources and targets in paths

Normally PathLinker removes all incoming edges to sources and outgoing edges to targets. If you want to allow sources and targets to also be intermediate nodes along source->target paths, select this option. Also, if a user wants to include nodes as both a source and a target, this option should be selected, otherwise all incoming and outgoing edges to the nodes in both will be removed.

Connect sources to each other

In some cases, a user may want to find paths or build a subnetwork connecting a set of nodes to each other where any node can be a source and any can be a target. Using the "Connect sources to each other" option will copy the sources to the targets field, select the "Allow sources and targets in paths" option, and disable the targets text field.

In the example below, I actually wanted to connect the lovastatin targets to each other, so I copied them to the sources field.



Here is the resulting subnetwork with the first couple of paths highlighted.

The screenshot displays the PathLinker Cytoscape interface. The top-left panel shows a list of networks, with 'PathLinker-subnetwork-50-paths-2' selected. The top-center panel shows a network graph with nodes and edges. The top-right panel, titled 'Results Panel', shows a table of results for 'PathLinker Result 2'. The bottom panel shows a table of network data.

Path index	Path score	Path
1	0.5625	P14859 P10275 Q15797
2	0.5625	Q15797 P10275 P14859
3	0.502197	P36956 P27361 Q15797
4	0.502197	P36956 P28482 Q15797
5	0.472237	Q15797 P62979 Q16236
6	0.472237	Q15797 P62987 Q16236
7	0.472237	Q16236 P62979 Q15797
8	0.472237	Q16236 P62987 Q15797
9	0.421875	P14859 P08047 P05412 Q16236
10	0.421875	P14859 P08047 P10275 Q15797
11	0.421875	P14859 P10275 P05412 Q16236
12	0.421875	P14859 P10275 P084022 Q15797
13	0.421875	P14859 P10275 Q92831 Q15797
14	0.421875	P14859 P10275 Q94NE7 Q15797
15	0.421875	P14859 P20226 P05412 Q16236
16	0.421875	P14859 P20226 Q09472 Q15797
17	0.421875	P14859 P24385 P10275 Q15797
18	0.421875	P14859 P42229 P10275 Q16236

shared name	shared interaction	edge_weight	path index 1	name	interaction	path index 2
P10275 (NetPath SPIKE)	NetPath SPIKE	0.75	2	P10275 (N...	NetPath SPIKE	2
P14859 (NetPath SPIKE)	NetPath SPIKE	0.75	1	P14859 (N...	NetPath SPIKE	1
P10275 (MI:0090 (protein complem...))	MI:0090 (protein complem...	0.75	3	P10275 (M...	MI:0090 (protei...	1
Q15797 (MI:0090 (pr...	MI:0090 (protein complem...	0.75	1	Q15797 (M...	MI:0090 (protei...	2

1.6.3 Edge weight options

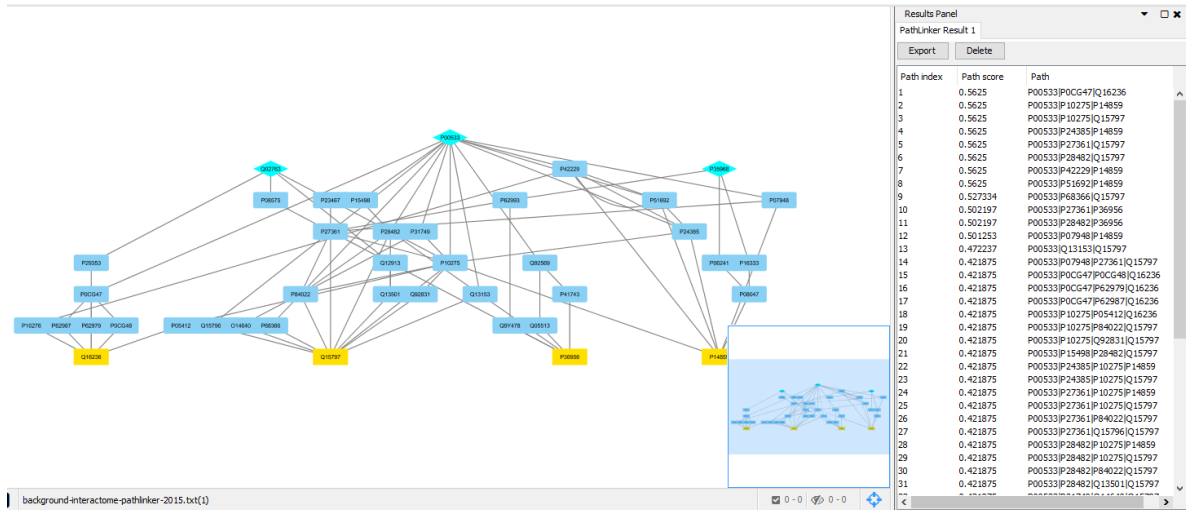
The interactome we imported also has probabilities for edge weights corresponding to the likelihood an interaction is real. We want PathLinker to use those weights, so select the "Weights are probabilities" option, and the "edge_weight" column. PathLinker transforms the probabilities where higher is better into a cost where lower is better by taking the $-\log$ (base 10) of the weight. The total cost of a path found by PathLinker is the sum of the $-\log$ transformed weights. That cost is transformed back into a score by taking e^{cost} which is equivalent to taking the product of the edge weights. The higher the path score, the better.

If your network has weights or costs where lower is better, use the "Weights are additive option". No transformation is necessary there.

1.6.4 Results

PathLinker will run after clicking submit and generate 3 outputs:

1. A subnetwork containing the k paths. The app will apply the hierarchical layout if < 2000 paths are in the results.
2. An interactive results table showing the paths along with the index of each path and the score of each path.
3. A *path index* X column in the edge table whose value is the index of the first path in which a given edge appears.



Control Panel

Network Style Select PathLinker

1 of 2 Networks selected

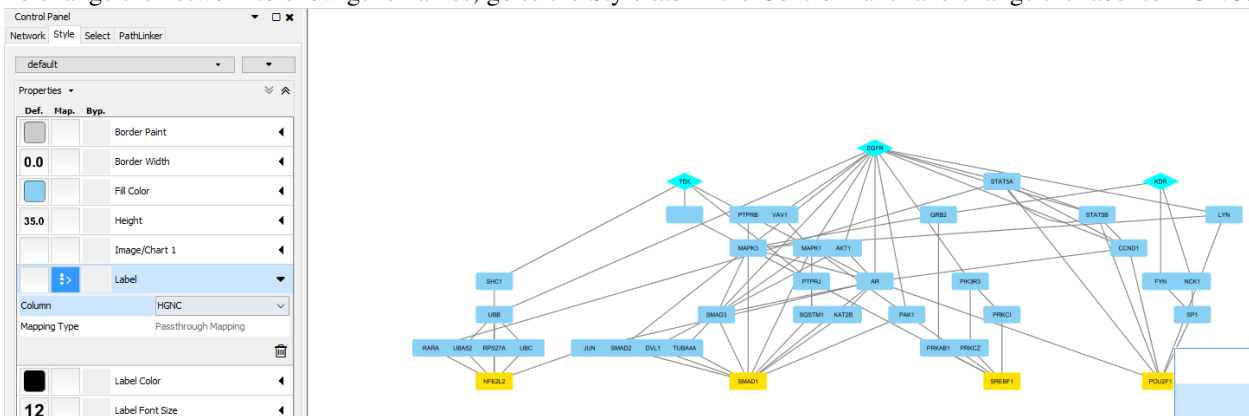
Network Name	Nodes	Edges
background-interactome-pathlinker-2015.txt	12063	159863
PathLinker-subnetwork-50-paths-1	41	80

Table Panel

shared name	shared interaction	edge
P00533 (NetPath) P0...	NetPath	
POCG47 (MI:0808 (c...	MI:0808 (comigration in sd...	
P00533 (NetPath) P1...	NetPath	
P10275 (NetPath SPI...	NetPath SPIKE	
P10275 (MI:0090 (pr...	MI:0090 (protein complem...	
P24385 (MI:0412 (el...	MI:0412 (electrophoretic ...	
P00533 (MI:0007 (an...	MI:0007 (anti tag coimmu...	
P27361 (KEGG phosp...	KEGG phosphorylation KE...	
P00533 (KEGG indire...	KEGG indirect-effect KEGG...	
P00533 (KEGG indire...	KEGG indirect-effect NetP...	
P28482 (KEGG phosp...	KEGG phosphorylation KE...	

Node Table Edge Table Network Table

To change the network to show gene names, go to the Style tab in the Control Panel and change the label to HGNC.



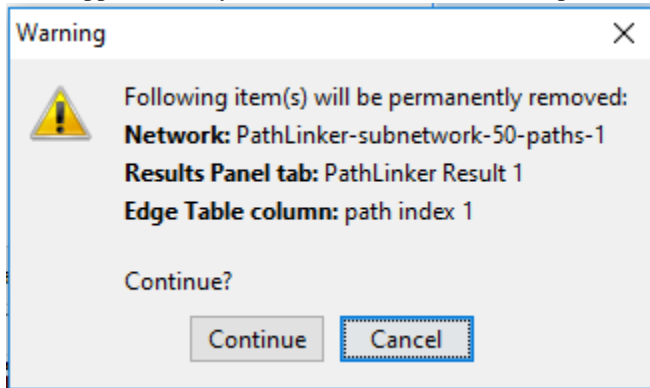
You can select rows in the results table to highlight/select the corresponding nodes and edges in the network.

The screenshot shows the PathLinker Cytoscape interface. On the left is a network diagram with nodes and edges. On the right is the Results Panel, which contains a table of paths. The table has three columns: Path index, Path score, and Path. The paths are listed in order of their path index, with the first path having a score of 0.5625 and the last path having a score of 0.421875. The Results Panel also has an 'Export' button and a 'Delete' button.

Path index	Path score	Path
1	0.5625	P00533 P0CG47 Q16236
2	0.5625	P00533 P10275 P14859
3	0.5625	P00533 P10275 Q15797
4	0.5625	P00533 P24385 P14859
5	0.5625	P00533 P27361 Q15797
6	0.5625	P00533 P28482 Q15797
7	0.5625	P00533 P42229 P14859
8	0.5625	P00533 P51693 P14859
9	0.527334	P00533 P68366 Q15797
10	0.502197	P00533 P27361 P36956
11	0.502197	P00533 P28482 P36956
12	0.501253	P00533 P07948 P14859
13	0.472237	P00533 Q13153 Q15797
14	0.421875	P00533 P07948 P27361 Q15797
15	0.421875	P00533 P0CG47 P0CG48 Q16236
16	0.421875	P00533 P0CG47 P62979 Q16236
17	0.421875	P00533 P0CG47 P62987 Q16236
18	0.421875	P00533 P10275 P05412 Q16236
19	0.421875	P00533 P10275 P84022 Q15797
20	0.421875	P00533 P10275 Q92831 Q15797
21	0.421875	P00533 P15498 P28482 Q15797
22	0.421875	P00533 P24385 P10275 P14859
23	0.421875	P00533 P24385 P10275 Q15797
24	0.421875	P00533 P27361 P10275 P14859
25	0.421875	P00533 P27361 Q10275 Q15797
26	0.421875	P00533 P27361 P84022 Q15797
27	0.421875	P00533 P27361 Q15796 Q15797
28	0.421875	P00533 P28482 P10275 P14859
29	0.421875	P00533 P28482 P10275 Q15797
30	0.421875	P00533 P28482 P84022 Q15797
31	0.421875	P00533 P28482 Q13501 Q15797

Click the export button to save the table to a tab-separated file. Paths are represented as nodes in order of the path with a vertical bar '|' between them

If you have PathLinker results you want to discard, use the *delete* button in the results panel. The app will try to delete all three outputs it generated with the selected Result Panel tab.



1.6.5 Edge penalty

Weights are Additive

If you have a large distribution of edge weights and some edges with very small costs (close to 0), your results could have many source->target paths that are very long (have many small-cost edges). Including an edge penalty would add the given penalty to each edge weight, which effectively penalizes each path (i.e. add to the total path score) by a factor of (number of edges in the path)^{*}(edge penalty). Default is 0 (no penalty) and numbers > 0 are allowed.

Weights are Probabilities

If your weights are probabilities and many are close to 1, they can become very small when taking the -log to transform to costs. This could lead to source->target paths that are very long (have many edges). To favor shorter paths with less edges, you can add an edge penalty which penalizes each path (i.e. add to the total path score) by a factor of (number of edges in the path)[^](edge penalty). Default is 1 (no penalty) and numbers >= 1 are allowed.

1.6.6 Include tied paths

Often paths can have the same score as other paths because the edges in them have the same weight. If there are multiple paths with the same score as the k th path, normally PathLinker ignores this and shows only up to k paths. To include all paths with the same weight as the k th path, use the "Include tied paths" option.

Control Panel

Network Style Select PathLinker

PathLinker

Help About

Sources/Targets

Select network: 2. background-interactome-pathlinker-20...

Sources separated by spaces (e.g., S1 S2 S3)
Must match the 'name' column in the Node Table

P00533 P35968 Q02763

Add selected source(s)

Targets separated by spaces (e.g., T1 T2 T3)

Q15797 P14859 Q16236 P36956

Add selected target(s)

Allow sources and targets in paths

Connect sources to each other

Clear

Algorithm

k (# of paths): 50

Edge penalty: 1

Include tied paths

Edge Weights

Unweighted

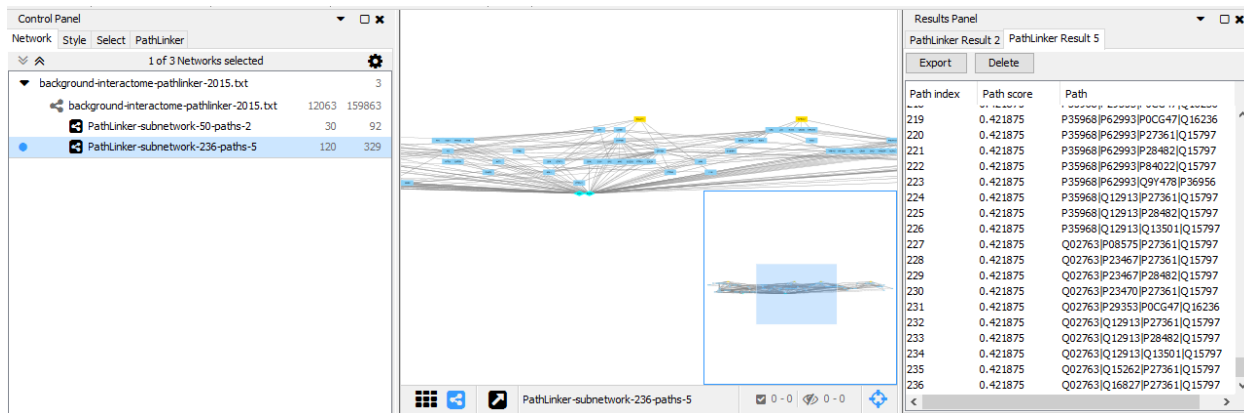
Weights are additive

Weights are probabilities

Edge weight column: edge_weight

Submit Close

In our example, including all paths for the set of receptors and TFs for lovastatin with an input k of 50 gives a total of 236 paths.



1.7 Errors and Warnings

Here I go over some example errors and warnings you may run into.

1.7.1 Sources or Targets not in network

Control Panel

Network Style Select PathLinker

PathLinker

Help About

Sources/Targets

Select network: 2. background-interactome-pathlinker-20...

Sources separated by spaces (e.g., S1 S2 S3)
Must match the 'name' column in the Node Table

EGFR KDR TEK

Add selected source(s)

Targets separated by spaces (e.g., T1 T2 T3)

MTF1 NFE2L2 POU2F1 SMAD1 SREBF1

Add selected target(s)

Allow sources and targets in paths

Connect sources to each other

Clear

Algorithm

k (# of paths): 50

Edge penalty: 1

Include tied paths

Edge Weights

Unweighted

Weights are additive

Weights are probabilities

Edge weight column: edge_weight

Submit Generating subnetwork... Close

Error Message

0 out of 3 sources are found in the network.

- Sources not found: [KDR, TEK, EGFR]
- Please ensure the entered node names match the 'name' column of the Node Table.
- Sources are required to run PathLinker.

0 out of 5 targets are found in the network.

- Targets not found: [SMAD1, SREBF1, POU2F1, MTF1, NFE2L2]
- Please ensure the entered node names match the 'name' column of the Node Table.
- Targets are required to run PathLinker.

OK

background-interactome-pathlinker-2015.txt

Table Panel

shared name	shared interaction	name	interaction	edge
Q92831 (NetPath SP... D10644 (NetPath) D1	NetPath SPIKE NetPath	Q92831 (N... D10644 (N)	NetPath SPIKE NetPath	

If you're getting an error that sources or targets are not found in the network, make sure you're using the right names for the nodes. PathLinker uses the 'name' column of the node table which contains UniProt IDs, so the gene names are not found.

The screenshot shows the PathLinker application interface. On the left is the 'Control Panel' with tabs for 'Network', 'Style', and 'Select'. The 'PathLinker' tab is active. The main area displays a network diagram with nodes and edges. A 'Warning' dialog box is overlaid on the network, indicating that 4 out of 5 targets are found in the network. The targets not found are listed as [Q14872]. The dialog asks if the user wants to cancel and correct the inputs or continue with the current settings (3 sources, 4 targets, k = 50, and edge penalty = 1.0). Below the dialog is a 'Table Panel' showing a table with columns for 'shared name', 'HGNC', and 'name'.

shared name	HGNC	name
P00533	EGFR	P00533
P00533	PTPRB	P00533

Another problem you could run into is all of your sources and targets may have been in the original network, but if you don't re-select the original network when running the app a second time, then PathLinker will use the subnetwork.

It could be that not all of the sources and targets you entered were used in the top k paths. In the screenshot, only 4 of the 5 targets were in the top 50 paths. In this case, you can continue and run PathLinker using the sources and targets that are in the current network.

1.7.2 Invalid input for parameters

k only takes positive integers and the **Edge penalty** only takes positive floats.

1.7.3 Not all edges have weight

In order for PathLinker to work properly, all edges must have weights when using a weighted option. You can either fix the empty edge weight rows, add a column by uploading another file with edge weights, or just use the unweighted option.

1.8 Cytoscape Automation / cyREST API

We added cyREST API functionality to the PathLinker app! Cytoscape automation is a new feature as of Cytoscape 3.6. Check out our [example jupyter notebook](#) to see how to use it. More details coming soon.

1.9 How to Cite PathLinker

We will be very glad to hear from you if you use PathLinker in your work. If you publish a paper that uses PathLinker, please cite:

1. [The PathLinker app: Connect the dots in protein interaction networks](#). Daniel Gil, Jeffrey Law, Li Huang and T. M. Murali. F1000Research 2017, 6:58
2. [Pathways on Demand: Automated Reconstruction of Human Signaling Networks](#). Anna Ritz, Christopher L. Poirel, Allison N. Tegge, Nicholas Sharp, Allison Powell, Kelsey Simmons, Shiv D. Kale, and T. M. Murali, Systems Biology and Applications, a Nature partner journal, 2, Article number 16002, 2016.

1.10 Contact Information

If you have any problems using PathLinker or any suggestions for improvement, please contact us by email at jeffl@vt.edu or murali@cs.vt.edu