



FILTER ORDERING - WHICH FILTER GOES WHERE IN THE CHAIN

📅 1 Aug 2026

Summary

The Causal Map app applies filters as a chain, from the top of the pipeline to the bottom. Each filter gets only the links the filter above it kept, and passes on whatever survives. So the same set of filters in a different order is a different analysis, and often answers a different question.

Two orderings go wrong often enough that the app now asks about them. It colours the filter pane, puts the question at the top of the pane, and lets you dismiss it. MapCat follows the same rules when it builds a chain for you.

Neither rule is a restriction. Any order means something, and an analyst with a reason should use the order they want.

The order to build in

1. Filters that choose which sources or links you are looking at: source groups, tags, the Everything filter, a factor filter used to pick a set of factors.
2. Tracing: path tracing, source tracing, pathways, a factor filter set to more than one step.
3. Combining opposites.
4. The rest of the filters that rewrite labels: zoom, collapse, removing brackets, soft relabelling, soft recode, clustering.
5. The frequency filters, link and factor.
6. Exclude self-loops, if you use it.

Rule 1. Put tracing at the top

Tracing exists to find causal paths that are really in the data: chains of links from one factor to another, and in source tracing, chains one person told in one interview. That only works if the tracing gets the links as they were coded.

Two kinds of filter above a tracing filter break it, in different ways.

A filter that rewrites labels merges factors. That is what zoom, collapse, combining opposites and the recoding and clustering filters are for, and it is usually what you want on a map. But two factors merged

into one label become a single point on a path, when in the coding they were two different things that nobody connected. You then trace a route through a junction nobody described.

A filter that drops links by frequency removes the material the paths are made of. A chain from **training** to **income** may depend on a link only two people mentioned. Cut the rare links first and that chain is gone, so you find no route where the data has one.

In both cases you are tracing a map that is no longer the coded data, and the result looks like a finding.

So trace first, and simplify afterwards for presentation. This is the conservative workflow on **Path tracing and source tracing**: source-trace to keep within-source chains, then zoom or collapse for a readable map.

Rule 2. Put the frequency filters at the bottom

A frequency filter counts the links it is given. Nothing else about it changes with position, which is exactly why its position changes the answer: at the top it counts the whole project, half way down it counts what has survived so far.

The clearest case is an ego network, where you want to know what people said about one factor. Filter to the neighbourhood of that factor first, then take the most frequent material within it, and you learn what people said most about that factor. The other way round you keep the most popular material in the whole corpus and then cut it down to the ego, which leaves you with the parts of the corpus-wide top twenty that happen to touch your factor. That is a defensible thing to look at, but it is rarely what the analyst meant, and the map does not show which of the two you did.

Two frequency filters in a row are fine, and often useful: take the top factors, then thin the links between them, as on **Simplification - factor and link frequency**. You can put an exclude-self-loops filter after them.

Rule 3. Put combining opposites above the filters that replace labels

This one fails silently, which is why it is worth a rule of its own.

Combining opposites finds each pair by reading the labels it is given. Either a numeric tag, **[~3]** against **[3]**, or a label starting with **~** against the same label without it. That is the whole mechanism, described on **opposites**.

Several filters replace a label with a different one. Removing brackets deletes the square-bracket tags. Collapse, soft relabelling, soft recode, clustering and auto recode swap the label for a new one, a search term or a magnet, which carries no marker. Put any of them above combining opposites and it finds no pairs at all. The negative and the positive then stay on the map as two separate factors, the map looks reasonable, and nothing tells you the pairing was lost.

So combine the opposites first, and replace labels afterwards for presentation. Zoom is the exception: it truncates the hierarchy rather than replacing the label, and keeps the marker on the part it keeps, so

opposites still pair at each level.

What the app does about it

- The filter pane colours itself and asks a question when a chain breaks either rule. Dismiss the question and it stays away until you next change the chain.
- MapCat puts a filter in the right place when you do not say where, and will explain either rule if you ask why.
- Nothing is blocked. Drag a filter wherever you want it.

More rules

Two other candidates turned out to need no rule of their own. Clustering sources into tribes on a map that has already been thinned by popularity is a real mistake, and so is animating frames over a thinned map, but rule 2 catches both: it flags any filter below a frequency filter. Excluding self-loops above a frequency filter is fine either way, since ranking the links without the self-loops is usually what you want.

We will add rules four onward here as we find them.