🌻 Extension - co-terminal link bundles

8 Nov 2025

In most projects, the data contains many **co-terminal links**: multiple coded claims with the **same cause and the same effect**. We call these **bundles of links**.

This extension is best thought of as two pieces:

1. **A filter (data transformation)**
2. **An interpretation rule (how to read a "link" on a map)**

1) The filter: bundle the links table

The filter operates on the current **links table** (one row per coded claim / citation) and produces a **derived table** with fewer rows by **grouping** on the (current, possibly transformed) factor labels:

- Bundle key: (cause label, effect label)
- One output row = one **bundle** (one unique cause→effect pair)

The bundled output adds aggregate columns such as:

- `bundle`: a readable key like `cause >> effect`
- `citation_count`: number of underlying link rows in the bundle (how many claims / coded citations)
- `source_count`: number of distinct sources contributing at least one link to the bundle (how many sources said it)

Optionally, further bundle-level summaries can be computed from the underlying rows, for example:

- `mean_sentiment`: mean effect sentiment across the bundle (if sentiment is present)
- per-tag or per-group counts (e.g. "how many links in this bundle have tag X?"), when you want enriched breakdowns

Important: bundling uses the **current filtered labels**, i.e. after any upstream label transforms (collapse, zoom hierarchies, bracket removal, soft recode, combine opposites, etc.). So the bundle

definition reflects the *conceptual normalisations* you have chosen.

Most causal mapping approaches which have recorded data from more than one source or context has done this.

2) The interpretation rule: what “a link” means on a map

Strictly, maps show **bundles**, not individual citations. But in practice we often still say “link” to mean:

> the **bundle** representing “many similar claims that x influences y”.

So a map statement like:

> “n sources claimed that x influences y (in m citations)”

is shorthand for: there are m underlying coded link rows whose (possibly transformed) labels share the same cause and effect, contributed by n unique sources.

Showing bundle statistics on maps

Bundling also supports a family of conventions for displaying **bundle-level statistics** visually, for example:

- edge width proportional to `citation_count`
- link label showing `source_count` or `citation_count`
- link label showing bundle summaries like mean sentiment (where available)

These are display choices layered on top of the same underlying bundle definition.