

Inferring the size of the causal universe: features and fusion of causal attribution networks

Daniel Berenberg^{1,2} and James P. Bagrow^{3,2,*}

¹Department of Computer Science, University of Vermont, Burlington, VT, United States

²Vermont Complex Systems Center, University of Vermont, Burlington, VT, United States

³Department of Mathematics & Statistics, University of Vermont, Burlington, VT, United States

*Corresponding author. Email: james.bagrow@uvm.edu, Homepage: bagrow.com

December 14, 2018

Abstract Cause-and-effect reasoning, the attribution of effects to causes, is one of the most powerful and unique skills humans possess. Multiple surveys are mapping out causal attributions as networks, but it is unclear how well these efforts can be combined. Further, the total size of the collective causal attribution network held by humans is currently unknown, making it challenging to assess the progress of these surveys. Here we study three causal attribution networks to determine how well they can be combined into a single network. Combining these networks requires dealing with ambiguous nodes, as nodes represent written descriptions of causes and effects and different descriptions may exist for the same concept. We introduce NetFUSES, a method for combining networks with ambiguous nodes. Crucially, treating the different causal attributions networks as independent samples allows us to use their overlap to estimate the total size of the collective causal attribution network. We find that existing surveys capture $5.77\% \pm 0.781\%$ of the $\approx 293\,000$ causes and effects estimated to exist, and $0.198\% \pm 0.174\%$ of the $\approx 10\,200\,000$ attributed cause-effect relationships.

Keywords— causality; knowledge graphs; graph alignment; natural language processing; word embeddings; capture-recapture estimators

Causality and causal reasoning are central questions of statistics, computer science, philosophy, and the cognitive sciences [1, 2, 3, 4]. Recently, our understanding of causality has been revolutionized by new insights and large volumes of data [5, 6]. Large-scale data collection and surveys of large numbers of individuals are now possible at an unprecedented scale. One class of new data enabled by the internet is very large-scale knowledge graphs, annotated semantic networks codifying large numbers of factual statements, events, and interrelationships between concepts [7, 8, 9]. Knowledge graphs allow generalization to new relationships using graph algorithms, and these algorithms have been applied to causal predictions [10].

In this work we study causal relationships encoded into networks. Nodes in these networks represent causes and effects, and directed links indicate cause-effect relationships. Generally, these relationships are gathered by large-

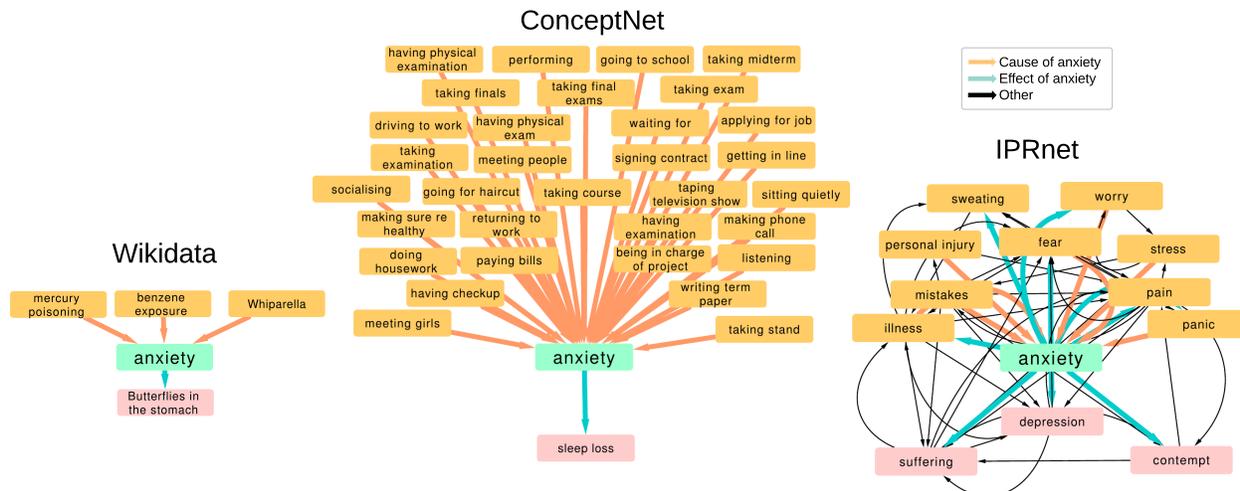


Figure 1: Causes and effects around ‘anxiety,’ a term common to all three networks studied here. This example illustrates similarities and differences of these networks, in particular the sparse, treelike structure of Wikidata and ConceptNet compared with the denser interlinking present in IPRnet.

scale surveying of individuals, often as part of larger efforts to build general-purpose semantic networks [11, 12], although dedicated experiments have also been conducted [13]. As these relationships are contributed by individuals or groups of individuals, we refer to these networks as causal attribution networks. Attribution theory, the study of how individuals perceive causality and attribute causes to effects, has long explored the cognitive biases that affect causal attribution [14, 15, 16].

We study causal attribution networks extracted from three sources, the collaboratively constructed knowledge graph “Wikidata”, the long-running project “ConceptNet,” and “IPRnet,” a network built by members of a crowdsourcing platform to test a network data collection method called “Iterative Pathway Refinement” [13]. Figure 1 shows examples from all three networks, centered on “anxiety” a term common to all three. Wikidata and ConceptNet encode other relationships, but we focus on causal relationships. All three networks represent different efforts to explore the larger causal attribution network, and our goal is to understand similarities and difference between these networks, and whether or not they can be fruitfully combined into a single, larger network. A key challenge when combining these data is resolving ambiguity between entities or concepts: most causes and effects in these networks are originally identified only by short written descriptions, and it is possible to describe the same entity in many different ways. Yet, overcoming this challenge to fuse multiple causal attribution networks together brings multiple benefits: it provides a common network dataset for researchers to study causality and attribution, and by measuring the overlap of different network samples we can estimate the size of the single underlying, incompletely observed causal attribution network.

The rest of this paper is organized as follows. Section 1 describes the data collection and network and text analysis procedures, introduces a method called NetFUSES for fusing graphs with potentially ambiguous node identities, and

discusses capture-recapture estimators for inferring the size of partially observed populations. Section 2 compares the three causal networks using both network and text statistics (Sec. 2.1), then investigates how well these networks may be combined (Sec. 2.2). Section 2.3 leverages these results to estimate the size of the underlying causal attribution network using capture-recapture methods. Finally, we discuss our results in Sec. 3.

1 Data and Methods

1.1 Causal attribution datasets

In this work we compare causal attribution networks derived from three datasets. A causal attribution dataset is a collection of text pairs that reflect cause-effect relationships proposed by humans (for example, “virus *causes* sickness”). These written statements identify the nodes of the network (see also our graph fusion algorithm for dealing with semantically equivalent statements) while cause-effect relationships form the directed edges (“virus” → “sickness”) of the causal attribution network.

We collected causal attribution networks from three sources of data: English Wikidata [12], English ConceptNet [11], and IPRnet [13]. Wikidata and ConceptNet, are large knowledge graphs that contain semantic links denoting many types of interactions, one of which is causal attribution, while IPRnet comes from an Amazon Mechanical Turk study in which crowd workers were prompted to provide causal relationships. Wikidata relations were gathered by running four search queries on the Wikidata API (query.wikidata.org). These queries searched for relations with the properties: "has immediate cause", "has effect", "has cause", or "immediate cause of". The first and third searches reverse the order of the cause and effect which we reversed back. We discarded any Wikidata relations where the cause or effect were blank, as well as one ambiguous relation where the cause was "NaN". ConceptNet attributions were gathered by searching the English ConceptNet version 5.6.0 assertions for “/r/Causes/” relations. Lastly, IPRnet was developed in [13] which we use directly.

The three networks together contain 23 239 causal links and 19 096 unique terms, of which there are 4 265 and 14 831 unique causes and effects, respectively.

1.2 Text processing and analysis

Analyzing causal sentences

Each node in our causal attribution networks consists of an English sentence, a short written description of an associated cause and/or effect. Text analysis of these sentences was performed using CoreNLP v3.9.2 and NLTK v3.2.2 [17, 18].

We computed Part-of-Speech (POS) tags and identified (but did not remove) stop words for these sentences. We used the standard Brown corpus as a text baseline for comparison. Text processing procedures such as lemmatization or removal of casing were not performed in order to retain information for subsequent operations. A small number of ConceptNet sentences contained ‘/n’ and ‘/v’ codes within the text denoting parts-of-speech tags; we removed these before applying our own POS tagger. POS tagging of the causal sentences and the baseline dataset was performed using CoreNLP by tokenizing each input using the Penn Treebank tokenizer then applying the Stanford POS tagger. This tagger uses Penn Treebank tags. We aggregated these 36 tags into NLTK’s universal tagset which consists of a simpler set of 12 tags including NOUN, VERB, ADJ, and more. To simplify presentation, we chose to further collect all non-verb, non-noun, and non-adjective tags into an “Other” tag. Stop words were identified using NLTK’s English stop words corpus.

Vector representations for words and sentences

Word vectors, or embeddings, are modern computational linguistics tools that project words into a learned vector space where context-based semantics of text are preserved, enabling computational understanding of text via mathematical operations on the corresponding vectors [19]. Many different procedures exist for learning these vector spaces from text corpora [19, 20, 21, 22]. Document embeddings, or “sentence vectors,” extend word vectors, representing more complex multi-word expressions in a vector space of their own [23]. Given two nodes i and j with corresponding sentences s_i and s_j and sentence vector representations \mathbf{v}_i and \mathbf{v}_j , respectively, the vector cosine similarity $\frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|}$ is a useful metric for estimating the semantic association between the nodes. High vector similarity implies that textual pairs are approximately semantically equivalent and sentence vectors can better compare nodes at a semantic level than more basic approaches such as measuring shared words or n-grams.

We computed sentence vectors using TensorFlow [24] v1.8.0 using the Universal Sentence Encoder v2, a recently developed embedding model that maps English text into a 512-dimensional vector space and achieves competitive performance at a number of natural language tasks [25]. This model was pretrained on a variety of text corpora [25]. The Universal Sentence Encoder was tested on several baseline NLP tasks including sentiment classification and semantic textual similarity, for each of which it performs with the highest accuracy. Given the higher performance of the Universal Sentence Encoder with respect to textual similarity tasks, we elected to utilize it instead of other sentence encoding models including the character level CNN architecture used in Google’s billion word baseline [26], and weighted averaging of word vector representations [27].

1.3 Graph fusion

Graph fusion takes two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ and computes a *fused* graph $G = (V, E)$ by identifying and combining semantically equivalent nodes (according to some measure of similarity) within and between V_1 and V_2 . Graph fusion is closely related to graph alignment and (inexact) graph matching [28], although fusion assumes the need to identify node equivalents both within and between the networks being fused, unlike alignment and matching which generally focus on uncovering relations *between* V_1 and V_2 . Graph fusion is particularly important when a canonical representation for nodes, such as an ID number, is lacking, and thus equivalent nodes may appear and need to be combined. This is exactly the case in this work, where each node is a written description of a concept, and the same concept can be equivalently described in many different ways.

Here we describe Network FUSion with SEMantic Similarity (NetFUSES). This algorithm computes the fused graph G given a node similarity function $f : V \times V \rightarrow \{0, 1\}$. This f should encode the semantic closeness between nodes u and v , with $f(u, v) = 1$ for semantically equivalent u and v and $f(u, v) = 0$ for semantically non-equivalent u and v . We assume $f(u, v) = f(v, u)$ and $f(u, u) = 1$.

To fuse G_1 and G_2 into G , first compute $F = \{f(u, v) \mid u, v \in V_1 \cup V_2\}$. One can interpret F as (the edges of) a *fusion indicator graph* defined over the combined node sets of G_1 and G_2 . Each connected component in F then corresponds to a subset of $V_1 \cup V_2$ that should be combined into a single node in V . (One can also take a stricter view and combine nodes corresponding to completely dense connected components of F instead of any connected components, but this strictness can also be incorporated by making f more strict.) Let F_i indicate the connected component of F containing node i . Abusing notation, one can also consider F_i as representing the node in G that the unfused node i maps onto. Lastly, we define the edges E of the fused graph based on the neighborhoods of nodes in V . The neighborhood $N_G(F_i)$ of each node $F_i \in V$ in the fused graph is the union of the neighborhoods of the nodes connected to i in F : for any node $i \in V_1 \cup V_2$, let $F_i = \{j \in V_1 \cup V_2 \mid i, j \text{ connected in } F\}$ and $K(i; H) = \bigcup_{u \in F_i \cap H} N_H(u)$. Then the neighborhood $N_G(F_i) = \{F_j \mid j \in K(i; G_1) \cup K(i; G_2)\}$ defines the edges incident on F_i in the fused graph and G may now be computed. Notice by this procedure that if an edge already exists in G_1 and/or G_2 between two nodes u and v that share a connected component in F , then a self-loop is created in G when u and v are combined. For our purposes these self-loops are meaningful, but otherwise they can be discarded.

Semantic similarity In this work, each node i is represented only by a short written sentence s_i , and two sentences $s_i \neq s_j$ may in fact be different descriptions of the same underlying concept. Hence the need for NetFUSES. To relate two sentences s_i and s_j semantically, we rely upon recent advances in natural language processing that can embed words and multiword expressions into a semantically-meaningful vector space (see Sec. 1.2). Let \mathbf{v}_i be the “sentence vector” corresponding to s_i . Then define $f(i, j) = 1$ if $\frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|} > t$ and zero otherwise, for some parameter t . In

other words, we consider nodes i and j to be semantically equivalent when the cosine similarity between their vectors exceeds a given threshold t . Our procedure in the main text determined $t = 0.95$ as an approach threshold.

1.4 Capture-recapture

Capture-recapture (also known as mark-and-recapture and recapture sampling) methods are statistical techniques for estimating the size of an unobserved population by examining the intersection of two or more independent samples of that population [29, 30]. For example, biologists wishing to understand how many individuals of a species exist in an environment may capture n_1 individuals, tag and release them, then later gather another sample by capturing n_2 individuals. The more individuals in the second sample that carry tags, the more likely it is that the overall population N is small; conversely, if the overlap in the samples is small, then it is likely that N is large. Capture-recapture is commonly used by biologists and ecologists for exactly this purpose, but it has been applied to many other problems as well, including estimating the number of software faults in a large codebase [29] and estimating the number of relevant academic articles covering a specific topic of interest [31].

The simplest estimator for the unknown population size N is the Lincoln-Petersen estimator. Assuming the samples generated are unbiased, meaning that each member of the population is equally likely to be captured, then the proportion of captured individuals in the second sample who were tagged should be approximately equal to the overall capture probability for the first sample, $n_1/N \approx n_{12}/n_2$. Solving for N gives the intuitive Lincoln-Petersen estimator $\hat{N} = n_1 n_2 / n_{12}$, for $n_{12} > 0$. While a good starting point, this estimator is known to be biased for small samples [30], and much work has been performed to determine improved estimators, such as the well-known Chapman estimator [32].

In this work we use the recently developed Webster-Kemp estimator [31]:

$$\hat{N} = \frac{(n_1 - n_{12} + 1)(n_2 - n_{12} + 1)}{n_{12}} + n_1 + n_2 - n_{12}, \quad (1)$$

which assumes (i) that one tried to capture as many items as possible (as opposed to predetermining n_1 and n_2 and capturing until reaching those numbers) and (ii) the total number of items found $n_1 + n_2 - n_{12} \gg 1$. Webster and Kemp also derive the variance of this estimator:

$$\sigma_{\hat{N}}^2 = \frac{(n_1 - n_{12} + 1)(n_2 - n_{12} + 1)(n_1 + 1)(n_2 + 1)}{n_{12}^2(n_{12} - 1)}, \quad (2)$$

with $n_{12} > 1$, allowing us to assess our estimate uncertainty. Equations (1) and (2) are approximations when assuming a flat prior on N but are exact when assuming an almost-flat prior on N that slightly favors larger populations N over

Table 1: Network statistics across each dataset. Abbreviations: d-directed, u-undirected, w-weak, s-strong.

	Wikidata	ConceptNet	IPRnet
Nodes	5 316	12 741	394
Edges	3 818	16 796	1 329
Causes	2 558	1 497	302
Effects	3 022	11 687	361
Self-loops	7	126	1
Reciprocated edges	7	7	214
Feedback loops	0	1	986
Feedforward loops	0	87	3541
Average degree	1.436	2.637	6.746
Clustering (u)	0.004194	0.00568	0.3804
Assortativity	-0.06101	-0.0569	-0.1627
Longest shortest path (d)	7	8	15
Longest shortest path (u)	25	12	11
Connected components (w)	1 780	415	1
Connected components (s)	5 309	12 732	153
Giant component size (w)	1 069	11 940	394
Giant component size (s)	3	3	242

smaller [31].

2 Results

Here we use network and text analysis tools to compare causal attribution networks (Sec. 2.1). Crucially, nodes in these networks are defined only by their written descriptions, and multiple written descriptions can represent the same conceptual entity. Thus, to understand how causal attribution networks can be combined, we introduce and analyze a method for fusing networks (Sec. 2.2) that builds off both the network structure and associated text information and explicitly incorporates conceptual equivalencies. Lastly, in Sec. 2.3 we use the degree of overlap in these networks as a means to infer the total size of the one underlying causal attribution network being explored by these data collection efforts, allowing us to better understand the size of collective space of cause-effect relationships held by humans.

2.1 Comparing causal networks

We perform a descriptive analysis of the three datasets, comparing and contrasting their features and properties. We focus on two aspects, the network structure and the text information (the written descriptions associated with each node in the network). Understanding these data at these levels can inform efforts to combine different causal attribution networks (Sec. 2.2).

Network Characteristics

Table 1 and Fig. 2 summarize network characteristics for the three causal attribution networks. We focus on standard measures of network structure, measuring the sizes, densities, motif structure, and connectedness of the three networks. Both Wikidata and ConceptNet, the two larger networks, are highly disconnected, amounting to collections of small components with low density. In contrast, IPRnet is smaller but comparatively more dense and connected, with higher average degree, fewer disconnected components, and more clustering (Table 1). All three networks are degree assortative, meaning that high-degree nodes are more likely to connect to low-degree nodes. For connectedness and path lengths, we consider both directed and undirected versions of the network allowing us to measure strong and weak connectivity, respectively. All three networks are well connected when ignoring link directionality, but few directed paths exist between disparate nodes in Wikidata and ConceptNet, as shown by the large number of strong connected components and small size of the strong giant components for those networks.

To examine motifs, we focus on feedback loops and feedforward loops, both of which play important roles in causal relationships [33, 34]. The sparse Wikidata network has neither loops, while ConceptNet has 87 feedforward loops and 1 feedback loop (Table 1). In contrast, IPRnet has far more loops, 986 feedback and 3541 feedforward loops.

Complementing the statistics shown in Table 1, Fig. 2 shows the degree distributions (2A), distributions of component sizes (2B), and distributions of two centrality measures (2C). All three networks display a skewed or heavy-tailed degree distribution. We again see that Wikidata and ConceptNet appear similar to one another while IPRnet is quite different, especially in terms of centrality. One difference between ConceptNet and Wikidata visible in 2A is a mode of nodes with degree ~ 30 within ConceptNet that is not present in Wikidata.

Text Characteristics

Understanding the network structure of each dataset only accounts for part of the information. Each node i in these networks is associated with a sentence s_i , a written word or phrase that describes the cause or effect that i represents. Investigating the textual characteristics of these sentences can then reveal important similarities and differences between the networks.

To study these sentences, we apply standard tools from natural language processing and computational linguistics (see Sec. 1). In Table 2 and Fig. 3 we present summary statistics including the total text size, average length of sentences, and so forth, across the three networks. We identify several interesting features. One, IPRnet, the smallest and densest network, has the shortest sentences on average, while ConceptNet has the longest sentences (Table 2 and Fig. 3A). Two, ConceptNet sentences often contain stop words ('the,' 'that,' 'which,' etc.; see Sec. 1) which are less likely to carry semantic information (Fig. 3B). Three, Wikidata contains a large number of capitalized sentences and

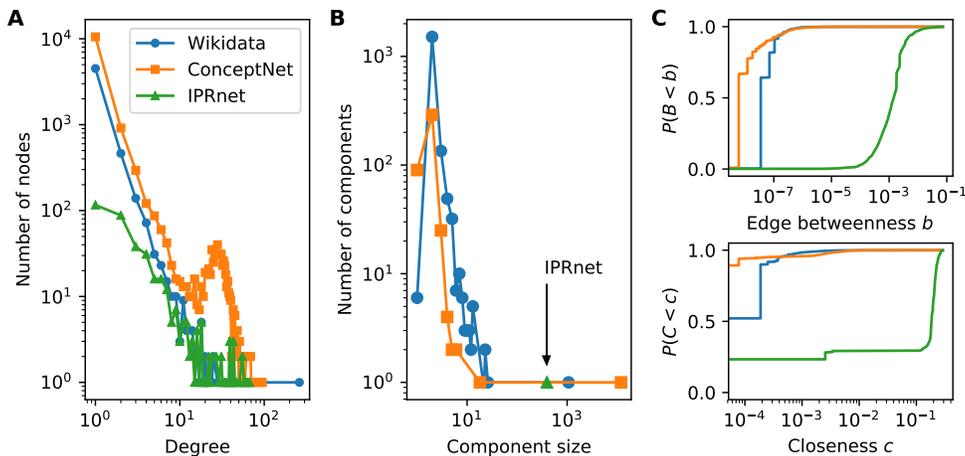


Figure 2: Degree (A) and weakly connected component size (B) distributions for each data set as well as cumulative centrality distributions (C) for edge betweenness (top) and closeness (bottom) centrality measures. Note the interesting modality of high degree nodes in ConceptNet.

Table 2: Text statistics across each dataset.

	Wikidata	ConceptNet	IPRnet
Sentences	5 316	12 741	394
Vocabulary size (unique words)	5 999	6 822	399
Text size (total words)	11 960	33 777	504
Average sentence length	2.25	2.651	1.279
Sentences with stop words	418	5 296	16
Sentences with numerical digits	533	8	0
Sentences with punctuation	571	163	0
Capitalized sentences	1 849	0	0

sentences containing numerical digits. This is likely due to an abundance of proper nouns, names of chemicals, events, and so forth. These textual differences may make it challenging to combine these data into a single causal attribution network.

We next applied a Part-of-Speech (POS) tagger to the sentences (Sec. 1). POS tags allow us to better understand and compare the grammatical features of causal sentences across the three networks, for example, if one network’s text is more heavily focused on nouns while another network’s text contains more verbs. Additionally, POS tagging provides insight into the general language of causal attribution and its characteristics. As a baseline for comparison, we also present in Fig. 3C the POS frequencies for a standard text corpus (Sec. 1). As causal sentences tend to be short, often incomplete statements, it is plausible for grammatical differences to exist compared with formally written statements as in the baseline corpus. For conciseness, we focus on nouns, verbs, and adjectives (Sec. 1). Nouns are the most common Part-of-Speech in these data, especially for Wikidata and IPRnet that have a higher proportion of nouns than the baseline corpus (Fig. 3C). Wikidata and IPRnet have correspondingly lower proportions of verbs than

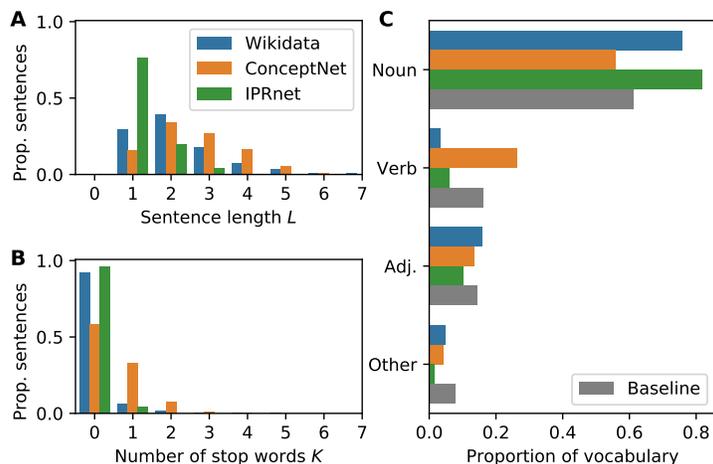


Figure 3: Properties of causal attribution sentences across the networks. (A) Distributions of total sentence lengths L (number of words) for all unique sentences. (B) Numbers of stop words per sentence. (C) Part of speech tags across words for each dataset, compared against a baseline POS tag distribution provided by the Brown corpus (grey). Note that the horizontal axis in panel A has been truncated for clarity: 0.5%, 0.02%, and 0% of Wikidata, ConceptNet, and IPRnet sentences, respectively, have $L > 7$.

the baseline. These proportions imply that causal attributions contain a higher frequency of objects committing actions than general speech. However, ConceptNet differs, with proportions of nouns and verbs closer to the baseline. The baseline also contains more adjectives than ConceptNet and IPRnet. Overall, shorter, noun-heavy sentences may either help or harm the ability to combine causal attribution networks, depending on their ambiguity relative to longer, typical written statements.

2.2 Fusing causal networks

These causal attributions networks are separate efforts to map out the underlying or latent causal attribution network held collectively by humans. It is natural to then ask if these different efforts can be combined in an effective way. Fusing these networks together can provide a single causal attribution network for researchers to study.

At the most basic level, one can fuse these networks together simply by taking their union, defining a single network containing all the unique nodes and edges of the original networks. Unfortunately, nodes in these networks are identified by their sentences, and this graph union assumes that two nodes i and j are equivalent iff $s_i = s_j$. This is overly restrictive as these sentences serve as descriptions of associated concepts, and we ideally want to combine nodes that represent the same concept even when their written descriptions differ. Indeed, even within a single network it can be necessary to identify and combine nodes in this way. We identify this problem as *graph fusion*. Graph fusion is a type of record linkage problem and is closely related to graph alignment and (inexact) graph matching [28], but unlike those problems, graph fusion assumes the need to identify node equivalencies both within and between the networks

being fused.

We introduce a fusion algorithm, NetFUSES (Network FUSion with SEmantic Similarity) that allows us to combine networks using a measure of similarity between nodes (Sec. 1). Crucially, NetFUSES can handle networks where nodes may need to be combined even within a single network. Here we compare nodes by focusing on the corresponding sentences s_i and s_j of the nodes i and j , respectively, in two networks. We use recent advances in computational linguistics to define a semantic similarity $S(s_i, s_j)$ between s_i and s_j and consider i and j as equivalent when $S(s_i, s_j) \geq t$ for some semantic threshold t . See Sec. 1 for details.

To apply NetFUSES with our semantic similarity function (Sec. 1) requires determining a single parameter, the similarity threshold t . One can identify a value of t using an independent analysis of text, but we argue for a simple indicator of its value given the networks: growth in the number of self-loops as t is varied. If two nodes i and j that are connected before fusion are combined into a single node u by NetFUSES, then the edge $i \rightarrow j$ becomes the self-loop $u \rightarrow u$. Yet the presence of the original edge $i \rightarrow j$ generally implies that those nodes are not equivalent, and so it is more plausible that combining them is a case of over-fusion than it would have been if i and j were not connected. Of course, in networks such as the causal attribution networks we study, a self-loop is potentially meaningful, representing a positive feedback where a cause is its own effect. But these self-loops are quite rare (Table 1) and we argue that creating additional self-loops via NetFUSES is more likely to be over-fusion than the identification of such feedback. Thus we can study the growth in the number of self-loops as we vary the threshold t to determine as an approximate value for t the point at which new self-loops start to form.

Figure 4 identifies a clear value of the similarity threshold $t \approx 0.95$. We track as a function of threshold the number of nodes, edges, and self-loops of the fusion of Wikidata and ConceptNet, the two largest and most similar networks we study. The number of self-loops remains nearly unchanged until the level of $t = 0.95$, indicating that as the likely onset point of over-fusion. Further lowering the similarity threshold leads to growth in the number of self-loops, until eventually the number of self-loops begins to decrease as nodes that each have self-loops are themselves combined. Thus, with a clear onset of self-loop creation, we identify $t = 0.95$ to fuse these two networks together.

2.3 Inferring the size of the causal attribution network

These three networks represent separate attempts to map out and record the collective causal attribution network held by humans. Of the three, IPRnet is most distinct from the other two, being smaller in size, denser, and generated by a unique experimental protocol. In contrast, Wikidata and ConceptNet networks are more similar in terms of how they were constructed and their overall sizes and densities.

Treating Wikidata and ConceptNet as two independent “draws” from a single underlying network allows us to

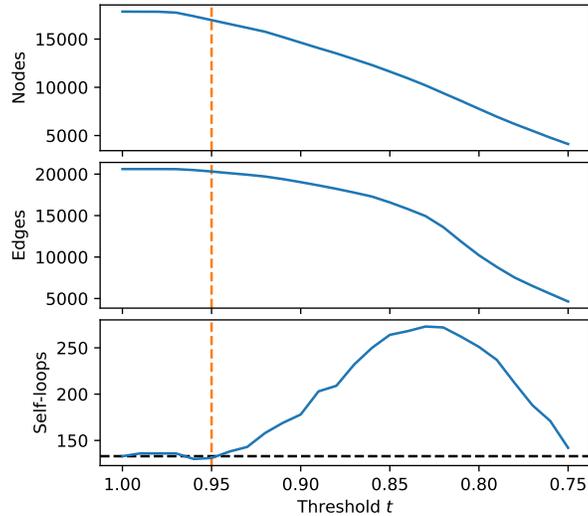


Figure 4: Statistics of fused Wikidata–ConceptNet networks across semantic similarity threshold values. Monitoring the number of self-loops, we observe a relatively clear onset of over-fusion at a threshold of $t \approx 0.95$. At this threshold, we observe a 4.95% reduction in the number of nodes and a 1.43% reduction in the number of edges compared with $t \geq 1$.

estimate the total size of this latent network based on their overlap. (We exclude IPRnet as this network is generated using a very different mechanism than the others.) High overlap between these samples implies a smaller total size than low overlap. This estimation technique of comparing overlapping samples is commonly used in wildlife ecology and is known as capture-recapture or mark-and-recapture (see Sec. 1.4). Here we use the Webster-Kemp estimator (Eqs. (1) and (2)), but given the size of the samples this estimator will be in close agreement with the simpler Lincoln-Petersen estimator.

We first begin with the strictest measure of overlap, exact matching of sentences: node i in one network overlaps with node j in the other network only when $s_i = s_j$. We then relax this strict assumption by applying NetFUSES as presented in Sec. 2.2.

Wikidata and ConceptNet contain 12 741 and 5 316 nodes, respectively, and the overlap in these sets (when strictly equating sentences) is 208. Substituting these quantities into the Webster-Kemp estimator gives a total number of nodes of the underlying causal attribution network of $\hat{N} = 325\,715.4 \pm 43\,139.2$ ($\pm 95\%$ CI). Comparing \hat{N} to the size of the union of Wikidata and ConceptNet indicates that these two experiments have explored approximately $5.48\% \pm 0.726\%$ of causes and effects.

However, this estimate is overly strict in that it assumes any difference in the written descriptions of two nodes means the nodes are different. Yet, written descriptions can easily represent the same conceptual entity in a variety of ways, leading to equivalent nodes that do not have equal written descriptions. Therefore we repeated the above estimation procedure using Wikidata and ConceptNet networks after applying NetFUSES (Sec. 2.2). NetFUSES incorporates

natural language information directly into the semantic similarity, allowing us to incorporate, to some extent, natural language information into our node comparison.

Applying the fusion analysis of Sec. 2.2 and combining equivalent nodes within the fused Wikidata and ConceptNet networks, then determining whether fused nodes contain nodes from both original networks to compute the overlap in the two networks, we obtain a new estimate of the underlying causal attribution network size of $\hat{N} = 293\,819.0 \pm 39\,727.3$. This estimate is smaller than our previous, stricter estimate, as expected due to the fusion procedure, but within the previous estimate’s margin of error. Again, comparing this estimate to the size of the union of the fused Wikidata and ConceptNet networks implies that the experiments have explored approximately $5.77\% \pm 0.781\%$ of the underlying or latent causal attribution network.

Finally, capture-recapture can also be used to measure the number of links in the underlying causal attribution network by determining if link $i \rightarrow j$ appears in two networks. Performing the same analysis as above, after incorporating NetFUSES, provides an estimate of $\hat{M} = 10\,235\,150 \pm 8\,962\,595.9$ links. This estimate possesses a relatively large confidence interval due to low observed overlap in the sets of edges. According to this estimate, $0.198\% \pm 0.174\%$ of links have been explored.

3 Discussion

The construction of causal attribution networks generates important knowledge networks that may inform causal inference research and even help future AI systems to perform causal reasoning, but these networks are time-consuming and costly to generate, and to date no efforts have been made to combine different networks. Our work not only studies the potential for fusing different networks together, but also infers the overall size of the total causal attribution network being explored.

We used capture-recapture estimators to infer the number of nodes and links in the underlying causal attribution network, given the Wikidata and ConceptNet networks and using NetFUSES and a semantic similarity function to help account for semantically equivalent nodes within and between Wikidata and ConceptNet. The validity of these estimates depends on Wikidata and ConceptNet being independent samples of the underlying network. As with many practical applications of capture-recapture in wildlife ecology and other areas, here we must question how well this independence assumption holds. The best way to sharpen these estimates is to introduce a new causal attribution survey specifically designed to capture either nodes or links independently (it is unlikely that a single survey protocol can sample independently both nodes and links), and then perform this same survey multiple times.

NetFUSES is a simple approach to graph fusion, in this case building off advances made in semantic representa-

tions of natural language, although any similarity function can be used to identify semantically equivalent nodes as appropriate. We anticipate that more accurate and more computationally efficient methods for graph fusion can be developed, but even the current method may be useful in a number of other problem domains.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1447634.

References

- [1] J. Pearl, *Causality: Models, Reasoning, and Inference*. Cambridge university press, 2009. [1](#)
- [2] M. Bunge, “Causality—the place of the causal principle in modern science,” 1960. [1](#)
- [3] D. Hume, *A Treatise of Human Nature*. Oxford University Press, 1738. [1](#)
- [4] I. Kant, *Critique of pure reason*. Cambridge university press, 1999. [1](#)
- [5] J. Pearl *et al.*, “Causal inference in statistics: An overview,” *Statistics surveys*, vol. 3, pp. 96–146, 2009. [1](#)
- [6] R. M. Shiffrin, “Drawing causal inference from big data,” *Proceedings of the National Academy of Sciences*, 2016. [1](#)
- [7] T. Berners-Lee, J. Hendler, and O. Lassila, “The semantic web,” *Scientific american*, vol. 284, no. 5, pp. 34–43, 2001. [1](#)
- [8] F. M. Suchanek, G. Kasneci, and G. Weikum, “Yago: a core of semantic knowledge,” in *Proceedings of the 16th international conference on World Wide Web*, pp. 697–706, ACM, 2007. [1](#)
- [9] T. Zesch, C. Müller, and I. Gurevych, “Extracting lexical semantic knowledge from wikipedia and wiktionary,” in *LREC*, vol. 8, pp. 1646–1652, 2008. [1](#)
- [10] K. Radinsky and E. Horvitz, “Mining the web to predict future events,” in *Proceedings of the sixth ACM international conference on Web search and data mining*, pp. 255–264, ACM, 2013. [1](#)
- [11] H. Liu and P. Singh, “Conceptnet—a practical commonsense reasoning tool-kit,” *BT technology journal*, vol. 22, no. 4, pp. 211–226, 2004. [2](#), [3](#)
- [12] D. Vrandečić and M. Krötzsch, “Wikidata: a free collaborative knowledgebase,” *Communications of the ACM*, vol. 57, no. 10, pp. 78–85, 2014. [2](#), [3](#)
- [13] D. Berenberg and J. P. Bagrow, “Efficient crowd exploration of large networks: The case of causal attribution,” *PACM on Human-Computer Interaction, Vol. 2, No. CSCW, Article 24. Publication date: November 2018*, vol. abs/1810.03163, 2018. [2](#), [3](#)
- [14] H. H. Kelley, “Attribution theory in social psychology,” in *Nebraska symposium on motivation*, University of Nebraska Press, 1967. [2](#)
- [15] S. E. Taylor and S. T. Fiske, “Point of view and perceptions of causality,” *Journal of Personality and Social Psychology*, vol. 32, no. 3, p. 439, 1975. [2](#)
- [16] H. H. Kelley and J. L. Michela, “Attribution theory and research,” *Annual review of psychology*, vol. 31, no. 1, pp. 457–501, 1980. [2](#)

- [17] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky, “The Stanford CoreNLP natural language processing toolkit,” in *Association for Computational Linguistics (ACL) System Demonstrations*, pp. 55–60, 2014. 3
- [18] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc., 2009. 3
- [19] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, “A neural probabilistic language model,” *Journal of machine learning research*, vol. 3, no. Feb, pp. 1137–1155, 2003. 4
- [20] J. Turian, L. Ratinov, and Y. Bengio, “Word representations: a simple and general method for semi-supervised learning,” in *Proceedings of the 48th annual meeting of the association for computational linguistics*, pp. 384–394, Association for Computational Linguistics, 2010. 4
- [21] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Advances in neural information processing systems*, pp. 3111–3119, 2013. 4
- [22] J. Pennington, R. Socher, and C. Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014. 4
- [23] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, “Supervised learning of universal sentence representations from natural language inference data,” *arXiv preprint arXiv:1705.02364*, 2017. 4
- [24] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, *et al.*, “TensorFlow: a system for large-scale machine learning,” in *OSDI*, vol. 16, pp. 265–283, 2016. 4
- [25] D. Cer, Y. Yang, S. Kong, N. Hua, N. Limtiaco, R. S. John, N. Constant, M. Guajardo-Cespedes, S. Yuan, C. Tar, Y. Sung, B. Strope, and R. Kurzweil, “Universal sentence encoder,” *CoRR*, vol. abs/1803.11175, 2018. 4
- [26] R. Jozefowicz, O. Vinyals, M. Schuster, N. Shazeer, and Y. Wu, “Exploring the limits of language modeling,” *arXiv preprint arXiv:1602.02410*, 2016. 4
- [27] S. Arora, Y. Liang, and T. Ma, “A simple but tough-to-beat baseline for sentence embeddings,” 2016. 4
- [28] F. Emmert-Streib, M. Dehmer, and Y. Shi, “Fifty years of graph matching, network alignment and network comparison,” *Information Sciences*, vol. 346-347, pp. 180 – 197, 2016. 5, 10
- [29] T. K. Nayak, “Estimating population size by recapture sampling,” *Biometrika*, vol. 75, no. 1, pp. 113–120, 1988. 6
- [30] K. H. Pollock, J. D. Nichols, C. Brownie, and J. E. Hines, “Statistical inference for capture-recapture experiments,” *Wildlife monographs*, pp. 3–97, 1990. 6
- [31] A. J. Webster and R. Kemp, “Estimating omissions from searches,” *The American Statistician*, vol. 67, no. 2, pp. 82–89, 2013. 6, 7
- [32] D. G. Chapman, “Some properties of hypergeometric distribution with application to zoological census,” *University of California Publications Statistics*, vol. 1, pp. 131–160, 1951. 6
- [33] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, “Network motifs: simple building blocks of complex networks,” *Science*, vol. 298, no. 5594, pp. 824–827, 2002. 8
- [34] F. C. Keil, “Explanation and understanding,” *Annu. Rev. Psychol.*, vol. 57, pp. 227–254, 2006. 8