

Distinguishing topological and causal explanation

Lauren N. Ross

Recent philosophical work has explored the distinction between causal and non-causal forms of explanation. In this literature, topological explanation is viewed as a clear example of the non-causal variety—it is claimed that topology lacks temporal information, which is necessary for causal structure (Pincock 2012; Huneman 2010). This paper explores the distinction between topological and causal forms of explanation and argues that this distinction is not as clear cut as the literature suggests. One reason for this is that some explanations involve both topological and causal information. In these “borderline” cases scientists explain some outcome by appealing to the causal topology of the system of interest. These cases help clarify a type of topological explanation that is genuinely causal, but that differs from standard topological and interventionist accounts of explanation (Woodward 2003).

1 Introduction. Modern discussions of scientific explanation often start with the deductive-nomological model, put forward by Hempel and others in the mid-to-late twentieth century (Hempel 1965). The deductive-nomological model helped popularize the topic of scientific explanation and it eventually inspired analyses of different forms of explanation. These forms included statistical, unification-style, and causal explanation, among many others. Although various types of scientific explanation have been discussed in this literature, no single type has received as much attention as causal explanation. Some have even suggested that this is the only form (or the main form) of explanation in all of science (Skow 2014). After enjoying this attention for the last few decades, assumptions about the singular importance of causal explanation in science have drawn criticism in recent work. Much of this criticism is driven by the view that non-causal explanation is both common and legitimate. One form of non-causal explanation that is discussed in this literature is topological explanation. While causal explanations appeal to the causes of some event, topological explanations appeal to the topological properties of some system. These topological properties are said to capture some “formal” or mathematical structure of a system, in which this structure is non-causal and abstracts from lower-level details (Pincock 2012, 53). As these topological features are viewed as explanatory while lacking causal information, they are considered straightforward examples of non-causal explanation.

While much of this literature assumes a clear divide between causal and topological forms of explanation, there are a number of complications for this view. One complication, is that various “borderline” cases appear to involve aspects of both types of explanation. In particular, scientists sometimes provide explanations by appealing to topological properties that are causal in nature—that is to say, by appealing to the causal topology of a system or a higher level “causal pattern” that

*I would like to thank Nicholas Jones, Philippe Huneman, Daniel Kostic, two anonymous reviewers, and audience members at the “Scientific Explanations, Competing and Conjunctive” conference, held by Philosophy Department at the University of Utah, for helpful feedback on this paper.

†To contact the author, please write to: Lauren N. Ross, Department of Logic and Philosophy of Science, University of California, Irvine, 3151 Social Science Plaza A, University of California, Irvine 92697-5100; email: rossl@uci.edu.

is present in some domain. If a strict, non-overlapping boundary exists between explanations that appeal to causal and topological properties, how should we understand explanations that appeal to both of these? What does it mean to say that an explanation appeals to causal topology and how should we understand these cases? This paper examines a variety of scientific examples in order to better understand how topology figures in scientific explanation, which explanatory-why questions it can answer, and whether such explanations are ever causal in nature.

2 Topological explanation. In the philosophical and mathematical literature, one of the more frequently discussed examples of topological explanation is the case of the Königsberg bridges. In this example, the Pregel River runs through the eighteenth-century city of Königsberg and seven bridges connect up two central islands to nearby landmasses.¹ Given the layout of these bridges there was interest in determining whether one could walk a path that crossed each of the bridges exactly and only once. Failed attempts to find such a route led to two related explanatory-why questions: First, does the Königsberg bridge system have a single route that traverses each of its bridges exactly and only once? And second, what explains whether a bridge system has or lacks such a route?

An answer to these questions was provided by Euler, who represented the bridges and landmasses graphically, in what is now considered the earliest work in graph theory and topology (Wilson 1999). Suppose that this bridge system is represented graphically such that the bridges are represented as edges, landmasses are nodes (or vertices), and the “degree” of a node refers to the number of edges that are connected to it, as shown in figure 1. Euler proved that in order for there to be a path that traverses each bridge only once—something we now call an “Eulerian path”—a bridge system needs to meet two conditions: all nodes should be connected to each other and there should be “either zero or two nodes of odd degree” (Euler 1956; Woodward 2019; Adams and Franzosa 2008). As the Königsberg bridge system fails to meet both of these conditions, it lacks such a traversable path.² In this sense, Euler’s work provided a “mathematical solution” to the problem of determining and explaining why a system has or lacks an Eulerian path (Carlson 2001, 104).




Figure 1: Introduction to Topology (Adams and Franzosa 2008, p. 414)

Euler’s solution to this problem contains at least three features that are viewed as characteristic of topological explanation. First, the explanation in this case appeals to the topology of the

¹Königsberg is a former German city, which is now Kalingrad, Russia.

²The Königsberg bridge system fails to meet the second condition, as all of its four nodes are of odd degree.

system, which captures a higher-level structure that abstracts away from various lower-level details. For Euler, this higher-level structure is represented graphically by connections between entities in the system, represented abstractly by edges and nodes. He takes the Königsberg explanatory-why questions to be concerned with “geometry of position,” which involves the relative position, organization, and structure of connections among entities in some domain—this is later referred to as the “topology” of a system. To say that this topology captures a “higher-level” structure means that this structure can be instantiated or realized by a variety of different physical or microstructural details (Kostić 2018; Huneman 2010).³ For example, the same Königsberg bridge topology can be found across bridges that are composed of different materials, such as wood, steel, brick and so on.

Second, philosophers interpret this topological information as capturing some structure of the system that is explicitly non-causal (Pincock 2012, 53). This is supported by claims that Euler’s representation lacks any time or temporal element, which causal structure necessarily contains (Huneman 2010; Pincock 2012). For example, the Königsberg bridge topology is said to lack at least two types of causal information. It lacks information about the causal process of individuals walking across bridges and it lacks information about the “material causes” that make up the bridge’s microstructure (Pincock 2012, 14,53). These “material causes” can be interpreted as lower-level *causal interactions* among molecular (and other) entities in a system, as opposed to physical properties that *realize* the topological structure.⁴ Both forms of causal information are absent from the graphical representation of the Königsberg bridge topology and they are unnecessary for providing the explanation of interest. With respect to whether a bridge system has a Eulerian path or not, all that matters is the topology of the system—its higher-level connections—and not its lower-level constituents or causal relations.

A third feature of the Königsberg example has to do with dependency relationships, which concern the relation between the explanans and explanandum. Various accounts of scientific explanation suggest that all (or many) forms of explanation involve dependency relations that specify how the explanandum is dependent (in some way) on the explanans (Woodward 2003; Jansson and Saatsi 2017; Reutlinger 2016). In the case of causal explanation, this dependency is often cashed out in terms of “empirical” and “causal” difference-making relationships Woodward (2003). These relationships capture the fact that manipulations of a cause produce changes in an effect and that changes in an effect depend on changes in a cause. In the case of causal explanation, these dependency relations are identified and verified empirically. However, in the case of topological explanation the dependency relation is provided by mathematical derivation as opposed to empirical study. For example, once we know the physical, non-causal structure of the Königsberg bridge system we can apply mathematical understanding to answer the explanatory-why questions. This explanatory work is facilitated by mathematics, as opposed to empirical investigations of causal relations in the world. In this sense, topological explanations are said to have dependency relations

³As best I know, this interpretation of “level” and “higher-level” originates in the multiple-realizability literature. In this literature, when some property A can be realized or instantiated by a variety of different physical details, the multiply realized property is said to be as a “higher-level” than its “lower-level” realizers (Putnam 1975). This characterization of “level” helps clarify one form of “abstraction” in these cases—namely, a higher-level property can “abstract” from the lower-level details that instantiate it.

⁴If the bridge’s higher-level topology is realized by some microstructure—and realization relations are not causal—how does this capture abstraction from causal detail? Although Pincock and others are not explicit about this, one charitable interpretation is that the bridge’s lower-level details consist of molecular interactions, which are causal in nature. These lower-level causal interactions differ across bridges made of different materials and they are omitted from representations of the system’s topology.

that are mathematical, while causal explanations have dependency relations that are empirical Woodward (2019).

Other examples of topological explanation that are discussed in this literature originate in ecology and molecular biology. In various ecological examples, mainly discussed by Huneman, networks and graphs are used to represent connections among distinct species in an ecosystem (Huneman 2010). In these cases, species are represented by nodes and the edges that connect them represent prey-predator relationships. Topological features of these systems are said to explain some of the properties that they exhibit. For example, one type of topological feature that these systems can exhibit is the property of being scale-free. A scale-free system is one in which few nodes are highly-connected and many nodes are only minimally connected. This scale-free system has a very notable feature—if single nodes are randomly deleted from this system (in which these deletions represent the extinction of an entire species), the system is more likely to remain stable than if it lacked this scale-free character. This is simply because a random deletion in a scale-free system is more likely to target a minimally-connected node, which will do less damage than the deletion of a highly-connected node. Alternatively, if the system is not scale-free, the random deletion is more likely to target a highly-connected node, which increases chances of collapsing the system. This explains how scale-free systems are stable or robust in situations of random deletions to the network.

This ecological case exhibits the three characteristics of topological explanation that are present in the Königsberg example. In this ecological case, the explanation is provided by some topological structure, this structure is taken to be non-causal, and the dependency relation between the explanandum and explanans is mathematical (as opposed to empirical). As the topological structure is “sufficient” for this explanation and devoid of causal information, the explanation is viewed as topological and non-causal. As the explanatorily relevant topology of the system abstracts from lower-level causal details, these details are viewed as “irrelevant” to the explanation and unable to “add anything to the understanding” of the explanandum (Huneman 2010, 222).

The non-causal nature of this topological explanation is supported in two main ways. First, similar to the Königsberg case, it is claimed that the relevant topology lacks any temporal or time-like feature, as “topology is not something that takes place in time” (Huneman 2010, 218). As time is an essential feature of causation, the topological structure is viewed as clearly non-causal. Second, the non-causal nature of this case is also supported through Huneman’s response to an objection. Suppose someone objects to the “non-causal” characterization of the ecological explanation, on the grounds that the prey-predator relationships (the edges connecting nodes in the graph) do contain causal information. This causal information might be represented as the flow of energy along prey-predator relationships, which is depicted in ecological pathways. Huneman responds to this objection by indicating that the directionality of these connections is irrelevant to the explanandum. In order to explain the stable or fragile character of the system, all you know to know is who is connected to who, not the directionality of these connections. As Huneman states, while “two species can have several kinds of causal relations” and “the nature of interactions between species—whether A preys on B, or is parasitic on B, or is preyed on by B, etc.—is not relevant, but only their number and the global shape of the connections between them as represented by a graph” (Huneman 2010, 219). The suggestion is that, while some causal information and directionality is present in lower-level characterizations of the system, this information is not necessary for the explanation and it is not present in the topology.

3 Borderline cases: Causal topology. While mainstream philosophical work suggests a clear division between causal and topological forms of explanation, some scientific examples threaten this view. Various “borderline” cases appear to involve aspects of both types of explanation. In particular, scientists sometimes explain by appealing to topological properties that are causal in nature—that is to say, by appealing to the *causal* topology of some system of interest. What can an analysis of these scientific cases contribute to philosophical discussions of topological explanation? What do these borderline cases reveal?

First, despite conflicting claims in the literature, we have good reason to think that topology can be causal. This is seen in various forms of topological analysis, including network and pathway approaches. These approaches represent biological systems graphically, in which nodes correspond to entities, while edges (or arrows, arcs, etc.) correspond to causal relations among these entities. Scientists explicitly state that these network topologies capture causal information. In describing the components of these graphs, they claim that “the edges denote interactions among...elements,” in which these interactions include things “such as ‘is transformed into...’ or ‘...binds to...’, ‘co-operates with...’ ” and so on (Palumbo et al. 2006, 220). These directed edges string together causally related entities in sequences that are often referred to as “pathways” or “causal pathways.” Examples of these include gene expression pathways, cell signaling pathways, metabolic pathways, developmental pathways, and ecological pathways (Palumbo et al. 2006, 220). When many pathways are integrated together—in a way that represents the causal connections among them—this forms a network, which captures a more extensive “map,” “chart,” or “web” of causal connections in some domain. These networks capture the topology of causal connections in some system and they are sometimes referred to as a “wiring topology,” as they bear similarity to wiring diagrams that trace causal connections in an electrical circuit (Tun et al. 2006, 5).

Does it really make sense to say that a graphical structure has *causal* topology? Yes—these network and pathway examples involve topology in the sense that the “topology of a graph defines how the links between system elements are organized” (Fornito et al. 2016, 6). It just so happens that the links in these cases are causal, as opposed to being undefined or correlational. Furthermore, this organization of causal connections captures a property that is preserved under continuous deformations of a system. If we take graphs from the previous examples and change their physical scale, or rotate, stretch, or reflect them, none of this will change the organization of connections that they exhibit. While this topological assessment was first applied to undirected graphs, as seen in Euler’s representation of the Königsberg bridges, it was later applied to directed graphs, which contain causal information. As Fornito et al. state, “[t]he principles of topological analysis have since been extended to more sophisticated graphs that include both weighted and directed connectivity” (Fornito et al. 2016, 6). This is seen in various network analyses in biology, in which scientists frequently refer to the “causal topology” of some system (Cvijovic et al. 2014; Zamir 2016; Shipley et al. 2005). Examples of different causal topologies include causal chains that are linear, branching, and cyclic, and ones that make up more complex structures such as bow tie and final common pathway configurations.

Second, a number of these cases reveal that causal topology is cited in scientific explanations. As an example of this, consider the bow tie configuration of T cell mediated immunity, which is discussed by Jones and Huneman (Jones 2014; Huneman 2018). In this case, the fragility of T cell mediated immunity is explained by the bow tie topology of immune cell interactions. As seen in figure 2, the “bow tie” terminology refers to a “fan-in, fan-out” structure, in which a variety of different inputs all converge on and operate through a single node, which then produces a variety

of different outputs. In this example, CD4+ T cells (a particular type of immune cell) are located at the single node or “knot” position of the bow tie. The bow tie structure of this system captures a particular topology, with unique consequences. If there is an insult that targets and immobilizes the CD4+ T cells—at the “knot” of the tie—the entire system will shut down. This has to do with the fact that all of the upstream signaling converges on and is mediated through these T cells, making them a point that—if manipulated or disrupted—will prevent any type of proper downstream response. In other words, the system-wide shutdown that occurs after T cell attack is explained by the “location of CD4+ T cells in this structure”—the fact that these cells are located at the knot of the bow tie, as opposed to being located at some other part of the topology or as opposed to the topology being different (Jones 2014, 1139).

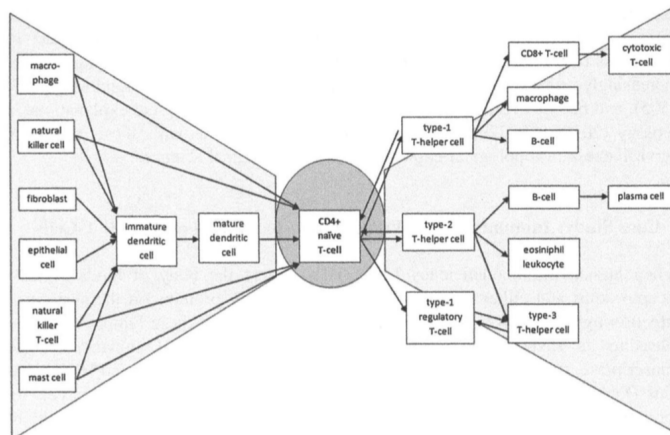


Figure 2: Bow tie structure of T cell mediated immunity (Jones 2014, p. 1138).

Alternatively, if the topology was different and there was a bypass or an additional route that connected the upstream signals to the proper downstream effect (without going through T cells at the knot), T cell malfunction would not lead to the collapse of the system, as this alternative route could take over. Thus, as differences in the topology of the system make a difference to the outcome of interest, these cases “show their explananda to be consequences of a system’s topological properties” (Jones 2014, 1136).

Interestingly, Jones and Huneman interpret this bow tie case as a *non-causal* topological explanation (Jones 2014, 1139-1140)(Huneman 2018, 116). There are a number of problems with this view. One significant problem is that the relevant topology in this case is causal, as opposed to the “formal,” “acausal” structure present in the Königsberg and ecological cases. One indication of this is that the Königsberg and ecological cases involve graphs with undirected edges, while the T cell case involves a graph with directed edges that are causal. This is evident in the figures that scientists use to represent this case, as arrows capture the directional flow of causal influence through the system, as seen in figure 2.⁵ This is not to suggest that “arrows” or “edges” always represent causal relationships—these symbols can be used to represent various types of non-causal relations.⁶ The use of an arrow (or edge) only indicates the potential for representation of causality—

⁵Figure 2 is from Jones (2014), who has adapted it from Kitano and Oda’s (2006) publication on T cell mediated immunity.

⁶One example of this is the use of arrows in conveying relationships of priority, such as trail signs repre-

determining whether it does or not requires assessing other information. This other information can include verbal descriptions of the represented system, intentions of the representer, and so on. When we look to this other information in the bow tie case we find that scientists describe these edges with causal terminology. Scientists describe this system as a “network of molecular interactions” and an “interaction pathway,” which involves the “transmission” of signals to T cells that “stimulate” changes in downstream entities Kitano and Oda (2006).⁷ Furthermore, these causal characterizations of the system explicitly figure in explanations of its fragile character (Kitano and Oda 2006, 1). This explanation is derived from the T cell’s “signal coordinator” role—the fact that T cells receive upstream information, which is delivered downstream. As T cells are a kind of causal intermediate along this pathway, disrupting their functionality severs a link between the upstream and downstream portions of the causal chain, effectively “uncoupling” them (Prochazka et al. 2014, 1). However, it is not just that T cells are a causal intermediate in this system, but that they are the sole causal intermediate linking all inputs to all outputs, which gives the system this fragile property.⁸ If all signals converge on and operate through T cells, then disrupting these cells can completely shutdown this transmission.

Suppose we agree that this topology contains causal information. How do we know that this causal information is *really* doing the explanatory work? Similar to Huneman’s ecological cases, it might be suggested that a non-causal bow tie topology is sufficient to provide this explanation. Suppose we remove all causal information from the bow tie by replacing its arrows with undirected edges. Can this structure explain the fragility of T cell mediated immunity? Perhaps this explanation is provided by the fact that the knot is highly-connected (or has a high degree of centrality) and that disrupting a highly-connected node, no matter what direction these connections take, is what explains collapse of the system? There are a number of issues with this objection. In particular, it is not just the high-connectivity (or centrality) of the bow tie’s knot that explains the fragility of this system. If an upstream or downstream node had a larger number of connections, the knot would still be the weak point.⁹ This has to do with the knot’s location in the network and the fact that it separates all incoming causal signals from outgoing effects—it captures a kind of bottleneck in this causal process. Identifying this location of the knot requires causal information about the network and it explains why intervening on the knot has implications for the entire system. As all causal information flows through T cells, represented by the knot and bottleneck, shutting them down collapses all cellular communication in the system. To the extent that these bow tie explanations depend on identifying causal bottlenecks, they cannot be provided with network topologies that abstract from causal information.

sending “user hierarchy” (arrows capture who should yield to who among pedestrians, cyclists and cars). Alternatively, a figure or diagram can represent causality without using either arrows or edges or all. For example, it has been argued that diagrams such as the periodic table represent causal information, despite lacking these symbols (Ross 2018a).

⁷Even Jones’ description reveals the causal nature of this case. He states that “[t]he directionality of cellular interactions within the immune system pathway determines the pathway’s bowtie structure: various stimuli activate pathways that converge to activate naive CD4+ T cells, which in turn activate a variety of responses” (Jones 2014, 1139). It should be clear that terms such as “interaction,” “activation,” and “directionality” are referring to causal relationships in this system.

⁸This captures a way that the connections to the T cell variable differs from connections to other variables in the system.

⁹I would like to thank Carlos Santana for insightful comments regarding this objection. For helpful discussion of the centrality feature, see (Jackson 2008, 61-65).

Third, aside from the causal nature of the bow tie topology, this case is not entirely different from standard formulations of topological explanation. Similar to these standard formulations, the explanatorily relevant topology abstracts from lower-level causal details and generalizes to various contexts. As in the Königsberg and ecological cases, explanation in the bow tie case does not appeal to (or depend on) lower-level causal details that instantiate the bow tie structure. The intricate details of how exactly different upstream factors all trigger T cells is irrelevant in this case—we only need to know that the system exhibits causal connections in the bow tie configuration and not how these causal connections occur or what instantiates them.¹⁰ Compare this to Pincock’s discussions of the Königsberg example, in which he states “[w]ith the bridges, we have abstracted away from the material causes of the bridges being in the shape that they are in, and represented them as they are independently of the details of their construction” [14]Pincock (2012). We can say something similar about systems that exhibit the bow tie structure. The bow tie structure can be found across systems that greatly differ in lower-level detail. None of these systems needs to share lower-level “material causes” in order to exhibit the same higher-level causal topology. Part of what this shows is that, abstracting from lower-level detail does not automatically produce the “abstract acausal” representations that Pincock is interested in—this process can produce abstract causal representations too. Abstracting from lower-level causal detail does not make a structure or representation acausal, as there may be relevant causal patterns at higher-levels that are shared across systems with different lower-level details (Reutlinger and Andersen 2017).

The abstract nature of the bow tie topology is related to the generality of this explanatory pattern—the same explanation can be provided in different systems with the same causal topology. This is seen in various subfields of biology, such as those focused on signaling, metabolic, and protein networks, in which the bow tie structure is a “well-known network topological motif” (Niss et al. 2018). The bow tie topology is also found in various ordinary life contexts, such as roadways that channel traffic. Consider a situation in which various one-way roads all converge on a single checkpoint, which then leads to a set of diverging roads. Similar to the biological cases, if the checkpoint is shut down all movement of traffic through the system will be stopped—traffic will build up along the incoming roads and vehicles will not pass the checkpoint. However, if the topology of the system were altered, for example, by adding a bypass around this convergence, then disrupting the checkpoint would no longer lead to collapse. An ordinary life example of this is a bow tie-shaped freight marshaling yard in Bologna, which scientists compare their biological bow tie cases to (Tierl et al. 2010, 7). The bow tie topology generalizes to a wide variety of systems and it explains why these systems exhibit similar behaviors, despite having different lower-level details.

This bow tie case uncovers a set of examples that have been overlooked in this literature. These are cases in which the explanatorily relevant topology is causal and captures a kind of higher-level “causal pattern” or causal topology that reoccurs in various contexts. In these cases, the explanation is not provided by a particular set of causal factors, but by the way in which these factors are connected up with each other. Here we see something similar to Euler’s interest in the “geometry of position”—in these cases, we are interested in the position, organization, and configuration of causal connections across nodes and how changes in these features explain changes in some system-wide outcome of interest.

In order to gain a better understanding of these causal topological cases, let us consider two

¹⁰For example, we do not need to know if T cells are triggered by the upstream cell (i) sending a chemical signal, (ii) physically manipulating an extracellular receptor, or any other causal process. It does not matter how these upstream cells trigger T cells, it just matters that they trigger them.

other examples. A second example comes from biochemistry and involves the concept of “chokepoint enzymes.” Enzymes are proteins that chemically convert one kind of metabolic substance into another. In the network models shown in figure 3, nodes represent different metabolic materials and the directed edges capture the enzymatic links between them. These enzymatic links represent the conversion of an upstream substrate into a downstream product.¹¹ In this way, edges capture the directional and causal process in which one metabolite is converted into another.¹² While this description captures a metabolite’s participation in a single enzymatic reaction, any given metabolite is typically “highly-connected” in the sense that it participates in a multitude of enzyme reactions in a living organism. In other words, when represented in the graphical framework shown in figure 3, any given metabolite is typically connected to many enzymes—there are many enzymes that produce this metabolite and many enzymes that use it as a substrate. However, chokepoint enzymes are an exception to this “highly-connected” norm. Chokepoint enzymes are enzymes that either uniquely consume or uniquely produce a particular metabolic material (Taylor et al. 2013). In other words, they are the sole enzyme in the organism that produces a particular product or that uses it as a substrate. This feature gives these enzymes their name as they are a “chokepoint,” or the sole connection to a metabolite, among an otherwise highly-connected network. Figure 3 represents a chokepoint enzyme that (A) uniquely consumes a metabolite and a chokepoint enzyme that (B) uniquely produces one.




Figure 3: In this diagram, nodes represent metabolites and arrows represent enzymes. The direction of the arrow captures the enzyme’s role in converting an upstream metabolic substrate into a downstream product. In (A) the yellow node/metabolite has only one arrow leaving it—this means that it has only one enzyme that consumes it. In (B) the yellow node/metabolite has only one arrow that enters it—this means that it has only one enzyme that produces it. These enzymes that either uniquely consume or uniquely produce a given metabolite are chokepoint enzymes (Taylor 2013, 2).

Chokepoint enzymes interest scientists because they are effective drug targets in disease caus-

¹¹For example, in the first step of glycolysis the substrate glucose is converted into glucose-6-phosphate (G-6P) by a particular enzyme (a variety of hexokinase isozymes catalyze this reaction). The framework in figure 3 would represent this in the following way: glucose would be the upstream node, G-6P would be the downstream node, and the enzyme would be the directed edge between them.

¹²For a detailed analysis of the relevant causal factors in these biochemical processes and their role in explanation, see (Ross 2018b).

ing organisms (such as pathogenic bacteria). Destroying an enzyme that is the sole producer of a metabolite can cause death by starvation, while destroying an enzyme that is the sole consumer of a metabolite can cause death by over-accumulation. Thus, as the only enzymatic link to a metabolic material, chokepoint enzymes can be targeted to potentially destroy the organism. While disrupting any random enzyme in an organism is unlikely to lead to noticeable or serious consequences, disrupting a chokepoint enzyme can be fatal. This is a surprising outcome that deserves explanation. Why does the disruption of a chokepoint enzyme have this deadly consequence, while disrupting any other enzyme is harmless? The answer to this question is provided by the causal topology of the metabolic network and the location of chokepoint enzymes in this structure. The answer has to do with different ways that these enzymes are causally connected in the system. The chokepoint enzyme’s single causal connection to a metabolic material represents a weak link that, if eliminated, shuts down the only pathway to producing or consuming the material. Depending on the directionality of this chokepoint enzyme, disrupting it leads to a lethal abundance or lethal deprivation of some metabolite.

As a third example, consider an ecological food web that is represented by figure 4. In this figure, the letters and nodes represent changes in energy levels of species in an ecosystem, while arrows capture prey-predator relationships between these species. Each arrow connects up some upstream prey to a downstream predator, in which the arrows capture the flow of energy through the ecosystem. In this case, C, D, and E are all different species of fish and A and B are two different species of prey (e.g. a species of crustacean and clam, respectively).¹³ Although these species are all located in the same body of water, scientists identify that fish species C has extremely high levels of Selenium, which is toxic to this fish and causes various pathologies. Meanwhile, fish species D and E have normal levels of Selenium and lack these issues (Ross 2019). Scientists want to know what explains these differences in toxic levels of Selenium across different species of fish. The fact that they are located in the same body of water, yet have different Selenium levels, is surprising and deserves explanation.

Scientists explain this fact by appealing to differences in the way that these predators are connected up in the larger food web. The explanation they provide involves citing the fact that species A (which is an upstream prey) has high levels of Selenium and that fish species C is causally connected to species A, while fish species D and E are not. In this manner, “exposures of top predators can be explained by food web relationships” and the fact that “predators feed differently” (Stewart et al. 2004, 4519).

In this case, the pattern of connections between nodes in the graph helps explain how a toxin accumulates in an ecosystem. The ecological pathways in figure 4 outline chains of causally connected species along which energy and toxic materials flow. If toxic materials enter into an upstream prey in the ecosystem, these materials can collect or biomagnify in downstream species in the food chain.¹⁴ The contours of these food chains—which species they connect up and how they link up the ecosystem—dictate where energy and toxic materials can flow. In the case above, it is the presence or absence of connections to the upstream prey that explains whether a downstream predator con-

¹³Technically, these letters do not represent the species themselves, but changes in their energy levels. This allows them to be interpreted with an interventionist framework. For more on this see (Ross 2019).

¹⁴It is easy to imagine ordinary life examples that are similar to this case. Suppose figure 4 represents a set of flowing rivers, with the letters and nodes as locations and arrows as segments of river. If one drops a basket of bread into the river at location A it will ultimately end up at location C, as opposed to locations D or E. This, of course, is because C is causally connected to A, while D and E are not.




Figure 4: Ecological food web (Ross 2019).

tains Selenium or not. Changing these causal connections will produce changes in the explanatory outcome of interest. If fish species C were no longer causally connected to A (or any primary producer with high Selenium) this species would no longer have high levels of this compound. If D and E were connected to A, they would have high levels of this compound. These differences are not captured with lower-level causal information about “how” a predator captures and eats its prey or how energy and Selenium move through these pathways. All that matters is that these materials move through the ecosystem in the ways captured by the causal topology of this food web diagram.

This section describes three examples of scientific explanation, in which an outcome is explained by appealing to some causal topology of the system of interest. These cases include: (1) biological bow ties, (2) chokepoint enzymes in biochemistry, and (3) bioaccumulation along ecological pathways. With respect to current philosophical accounts of scientific explanation these examples pose an interesting problem. These examples contain features of both topological and causal explanation, despite the fact that these types of explanation are viewed as mutually exclusive in the literature. How should we reconcile the features of these cases with analyses of these two types of explanation?

4 Examining causal topological explanation. These causal topological cases raise a number of questions. First, what does it mean to say that an explanation is topological? Second, are these explanations really topological and causal? If so, how do we know?

First, how do we know whether an explanation is topological or not? Consider Huneman’s answer to this question. Huneman states that topological explanation is “a kind of explanation that relies upon ‘topological’ properties of systems in order to derive the explanandum as a consequence, and which does not consider mechanisms or causal processes” (Huneman 2010, 213). I will suggest some modifications to these claims. To begin, we should avoid defining topological explanation in contrast to causal or mechanistic explanation. It is better to start with the view that explanations involve dependency relations that specify how the explanandum depends on the explanans (Woodward 2003; Reutlinger 2016; Jansson and Saatsi 2017). These dependencies can be understood as “difference-making” relations, which indicate how the explanans “makes a difference” to the explanandum and, relatedly, how the explanandum depends on the explanans. This framework specifies three main components of an explanation: explanans, explanandum, and

dependency relation.

With this framework in mind, topological explanation can be defined as any explanation in which topology does the explanatory work. This definition accommodates both the Königsberg case and Huneman’s ecological example. In the former, whether a system has an Eulerian path or not depends on changes in topology, namely, changes in connections among nodes (or links between bridges). In the later, whether a system is robust or not, depends on changes in connections among nodes (or relations among species in an ecosystem), which determine whether the system is scale-free or not. This suggests that an answer to whether an explanation is topological or not involves assessing the *explanans* component. If the explanans involves topological information, the explanation is topological. If it lacks this information, it is not.

Second, how do we know whether an explanation is causal or not? A helpful way to address this question involves relying on the same three-component assessment above. In recent work, a number of philosophers have suggested that one way to understand the difference between causal and non-causal (or mathematical) explanation is in terms of dependency relations. It has been claimed that non-causal (mathematical) explanations involve dependency relations that are mathematical, while causal explanations involve dependency relations that are empirical (Woodward 2003; Jansson and Saatsi 2017; Reutlinger 2016). This is argued for in analyses of the Königsberg and ecological cases. Both of these cases involve starting with some non-causal topological structure and then using mathematics alone to answer the explanatory-why question. Causal explanations do not follow this pattern. Suppose we want to know whether a gene causes some disease. In this case, we cannot start with some non-causal genetic information and then apply mathematics to get answer. We need empirical information about the how the gene acts in the world—about what happens to the disease outcome when the gene is manipulated, as we find in animal model experiments. In this case, the genetic factor and disease outcome are connected by an empirical relationship. Alternatively, the Königsberg topological structure and availability of an Eulerian path are connected by a mathematical one. What this shows is that one way to determine whether an explanation is causal or not involves assessing its dependency-relation component. If the dependency-relation is empirical then the explanation is causal, and if it is mathematical then the explanation is non-causal.

How do we apply these notions of topological and causal explanation to the causal topological cases examined in this paper? One helpful way to do this involves comparing these causal topological cases to standard interventionist examples of causal explanation Woodward (2003). This highlights differences between these cases and important features of causal topological explanation.

In starting with the interventionist framework, to say that X is a cause of Y means that an ideal intervention that changes the values of X, in background circumstances B, produces changes in the values of Y Woodward (2003). On this framework, causal relationships between variables are naturally represented with directed acyclic graphs and are compatible with interpretations of causal structure in network models. Suppose that we are working within this interventionist framework and someone provides us with the causal diagram shown in figure 4. Furthermore, suppose that this diagram captures something different than the original ecological case. In this new scenario, C, D, and E represent distinct disease outcomes, while A and B capture the causal factors that lead up to them (such as a particular virus, gene, and so on). How does causal explanation work within this interventionist picture? We can address this with our three-component assessment of explanation, in which we examine the (a) explanandum, (b) explanans, and (c) dependency relation. The first step involves specifying an explanatory target or explanandum of interest. In the interventionist framework, the explanatory target is represented with an effect variable that can take on a variety

of different values. Suppose we are interested in variable “D” in figure 4, which represents the occurrence and nonoccurrence of disease D. With this specification of the explanatory target, and information provided in the causal graph, we can now consider what explains this target. One way of identifying explanatorily relevant factors involves locating factors in the causal history of the effect. In the interventionist framework, these causal factors are represented by upstream cause variables connected to the effect by edges (or arrows). If we start at variable D and trace backwards, this reveals factors that lead up to the production of this disease and can be cited in explaining it. This takes care of the explanandum and the explanans, but we still have the dependency relation to consider. In this case, the dependency relations between variables are empirical—they contain information about how changes in one property depend on changes in another, where these dependencies are obtained from empirical investigations of the world. This is what we should expect from the interventionist account, as this provides a causal explanation (as opposed to a non-causal one).

Now we can examine the causal topological cases with the same three-component criteria. As mentioned above, explanations typically start with the specification of some (a) explanatory target of interest. However, while the interventionist-explanandum is well-captured by a single variable in the causal diagram, there is no variable or node in the network model that captures the explanandum in the topological cases. Instead, the causal topological cases involve “system-level” explananda that involve differences across variables in the causal system or some other property of the system that is not captured by a single variable. As an example of this, consider the ecological-bioaccumulation case in which the explanatory target consists of *differences* across variables C, D, and E, namely differences in bioaccumulation across different fish species. As another example, in the bow tie and chokepoint enzyme cases the explanatory targets are system-level properties, such as fragility and collapse, which are not captured by any variable in the corresponding graphs. Thus, one difference between the interventionist and causal topological cases is that the latter involve system-level explanatory targets that are not directed represented in the causal or network models.

Now we can examine the nature of the explanans in the causal topological cases. Recall that within the interventionist account the explanans is represented by the upstream causal history of some effect. Notice how this feature is not present in the causal topological cases. This is, at least, because the effect is not represented by a single variable so there is no identifiable causal chain leading up to it. In these causal topological cases, the explanatory work is not done by causal variables in the history of an outcome, but by a system-level pattern of causal connections among these variables (and ways in which different connections would lead to differences in the explanatory target). This marks a serious difference between the causal topological and interventionist frameworks—causal topological cases appeal to system-level patterns of causal *connections*, as opposed to particular causal *factors*. Causal connections have a more system-level character than causal factors, because examining causal connections requires assessing global features of a causal system, as opposed to limiting one’s focus on only those factors in the direct causal history of an effect. As mentioned earlier, assessing the nature of the explanans helps us determine whether these cases should be viewed as topological or not. As these cases involve appealing to topology—the higher-level causal connections present in some system—they should be viewed as instances of topological explanation.

Finally, we should consider the dependency relation in the causal topological cases. Getting clear on the nature of this relation is important, because this determines whether an explanation

is causal or non-causal. In causal topological cases, the dependency relation connects changes in causal topology to changes in a higher-level explanatory target of interest. Is this relation purely mathematical or is it empirical? While a complete response to this question requires more detail, we have good reason to think that this relation is causal. The main reason for this is that this relation involves a significant amount of empirical information. Given that mathematical dependency relations capture how the explanans leads to the explanandum in a way that is derived from mathematical understanding, causal topological cases appear different. In these cases, changes in causal topological structure are said to lead to the explanatory outcome. These topological structures are explicitly causal—in considering different connections, we imagine how these differences matter for the causal processes that propagate along these connections and, ultimately, for the explanatory target of interest. The causal topological structures that figure in the explanans contain significant empirical information—each causal link specifies information about a difference-making relationship between properties in the world. We have good reason to view these causal topological cases as involving empirical dependency relation, which support causal explanation. However, gaining a more complete understanding of these relations requires further analysis and is a promising area for future work.

In this section, I have suggested that one important difference between causal topological and interventionist explanation is that the former has more of a system-level character than the latter—that the explanandum and explanans refer to system-level properties that are not captured by discrete variables in a graph. Consider an objection to this claim. It might be suggested that the causal topological cases can be interpreted in a way in which the explananda and explanata are captured by single variables. Suppose it is suggested that in understanding the theoretical structure of the explanation, whatever the explanandum is, it is represented by a single effect variable and whatever the explanans is, it is represented by a set of causal variables. So if there is interest in explaining (a) differences in Selenium levels across species or (b) varying degrees of systemic fragility that—although there are causal graphs that split these phenomena into many variables—for the purposes of understanding a scientific explanation, these explanatory targets can be represented by single effect variables. The same might be said for explanatory factors that make up the explanans—that these can be understood as discrete, single variables with difference-making connections to the explanatory target. This objection suggests that the causal topological explanations are not really “system-level” compared to the the interventionist cases, because their explananda and explanata can be captured with discrete variables in a way compatible with the interventionist framework. Maybe these cases are not really different from the standard interventionist examples, after all?

In considering scientific explanation from a theoretical perspective any explanandum could be represented as a single effect variable with one or more variables representing the explanans. The point here is that, in the graphical models that scientists use to provide these explanations, the explanandum and explanans are system-level properties that are not represented by single variables. A causal topological structure will necessarily capture some higher-level property of many causal connections, which is not the case for most explanatory factors in the interventionist framework.¹⁵ The point is not that you cannot represent the explanandum and explanans as single variables in the causal topological cases—it is that the system-level nature of these properties make them different from the properties in the explanandum and explanans of standard interventionist cases.

¹⁵For example, a “gene” that explains a disease does not involve sets of causal connections in the way that the “bow tie topology” that explains some systemic property does.

5 Conclusion. Although causal explanation has received significant attention in the philosophical literature, many now accept that there are legitimate forms of non-causal explanation. While appreciating the nature of non-causal explanation is important, some of this literature erroneously suggests that all topological explanations are non-causal. This paper suggests that there is no clear-cut distinction between topological and causal explanation because they are not mutually exclusive. Some topological explanations are causal, while others are not. We can distinguish causal from non-causal, and topological from non-topological, but the causal and topological categories overlap. The approach taken in this paper involves separating explanation into three components: explanandum, explanans, and dependency relation. One can then compare different examples or types of scientific explanation by assessing how they differ with respect to these components. This provides a way to identify novel forms of explanation and compare them to commonly discussed types of explanation. This approach has promise for addressing further questions that arise in this literature. For example, how should we revise our taxonomy of scientific explanation in light of causal topological cases? Are “causal” and “non-causal” the two main and mutually exclusive categories that all scientific explanations fall into? Relatedly, do causal topological explanations fall into the “causal” category or are they in a category of their own? Addressing these questions will require further attention to the diversity of explanatory patterns in science and ways in which scientific understanding is generated.

References

- Adams, C. and Franzosa, R. (2008). *Introduction to topology*. Pure and applied. Pearson Prentice Hall, New Delhi.
- Carlson, S. (2001). *Topology of Surfaces, Knots, and Manifolds*. John Wiley & Sons, Inc.
- Cvijovic, M., Almquist, J., Hagmar, J., Hohmann, S., Kaltenbach, H.-M., Klipp, E., Krantz, M., Mendes, P., Nelander, S., Nielsen, J., Pagnani, A., Przulj, N., Raue, A., Stelling, J., Stoma, S., Tobin, F., Wodke, J. A. H., Zecchina, R., and Jirstrand, M. (2014). Bridging the gaps in systems biology. *Molecular Genetics and Genomics*, 289(5):727–734.
- Euler, L. (1956). *The seven bridges of Königsberg*, volume 1 of *The world of mathematics*. Simon and Schuster.
- Fornito, A., Zalesky, A., and Bullmore, E. (2016). *Fundamentals of Brain Network Analysis*. Academic Press.
- Hempel, C. (1965). *Aspects of scientific explanation*. And Other Essays in The Philosophy of Science. The Free Press.
- Huneman, P. (2010). Topological explanations and robustness in biological sciences. *Synthese*, 177(2):213–245.
- Huneman, P. (2018). Diversifying the Picture of Explanations in Biological Sciences: Ways of Combining Topology with Mechanisms. *Synthese*, 195(1):115–146.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jansson, L. and Saatsi, J. (2017). Explanatory abstractions. *The British Journal for the Philosophy of Science*.
- Jones, N. (2014). Bowtie Structures, Pathway Diagrams, and Topological Explanation. *Erkenntnis*, 79:1135–1155.
- Kitano, H. and Oda, K. (2006). Robustness trade-offs and host–microbial symbiosis in the immune system. *Molecular Systems Biology*, 2:787.
- Kostić, D. (2018). The topological realization. *Synthese*, 195(1):79–98.
- Niss, K., Hu, J. X., Gomez-Casado, C., Joeris, T., Agace, W. W., Belling, K. G., and Brunak, S. (2018). Bow-tie motifs enable protein multifunctionality by connecting cellular processes in the interactome. *bioRxiv*.
- Palumbo, M. C., Farina, L., Colosimo, A., Giuliani, A., Tun, K., and Dhar, P. K. (2006). Networks Everywhere? Some General Implications of an Emergent Metaphor. *Current Bioinformatics*, pages 219–234.
- Pincock, C. (2012). Mathematics and scientific representation. *Oxford University Press*.

- Prochazka, L., Angelici, B., Haefliger, B., and Benenson, Y. (2014). Highly modular bow-tie gene circuits with programmable dynamic behaviour. *Nature Communications*, 5:1–12.
- Putnam, H. (1975). *Philosophy and our Mental Life*, volume 2 of *Mind, Language, and Reality*. Cambridge University Press, Cambridge.
- Reutlinger, A. (2016). Is there a monist theory of causal and noncausal explanations? The counterfactual theory of scientific explanation. *Philosophy of Science*, 83:733–745.
- Reutlinger, A. and Andersen, H. (2017). Are there non-causal explanations (of particular events)? *International Studies in the Philosophy of Science*, 30(2):129–146.
- Ross, L. N. (2018a). Causal explanation and the periodic table. *Synthese*.
- Ross, L. N. (2018b). Causal selection and the pathway concept. *Philosophy of Science*, 85:551–572.
- Ross, L. N. (2019). Causal concepts in biology: How pathways differ from mechanisms and why it matters (Forthcoming). *The British Journal for the Philosophy of Science*.
- Shiple, B., Vile, D., Garnier, E., Wright, I. J., and Poorter, H. (2005). Functional linkages between leaf traits and net photosynthetic rate: reconciling empirical and mechanistic models. *Functional Ecology*, 19(4):602–615.
- Skow, B. (2014). Are There Non-Causal Explanations (of Particular Events)? *The British Journal for the Philosophy of Science*, 65(3):445–467.
- Stewart, A. R., Luoma, S. N., Schlekot, C. E., Doblin, M. A., and Hieb, K. A. (2004). Food Web Pathway Determines How Selenium Affects Aquatic Ecosystems: A San Francisco Bay Case Study. *Environmental Science & Technology*, 38(17):4519–4526.
- Taylor, C. M., Wang, Q., Rosa, B. A., Huang, S. C.-C., Powell, K., Schedl, T., Pearce, E. J., Abubucker, S., and Mitreva, M. (2013). Discovery of Anthelmintic Drug Targets and Drugs Using Chokeypoints in Nematode Metabolic Pathways. *PLoS Pathogens*, 9(8):1–17.
- Tieri, P., Grignolio, A., Zaikin, A., Mishto, M., Franceschi, C., Remondini, D., and Castellani, G. (2010). Network, degeneracy and bow tie. Integrating paradigms and architectures to grasp the complexity of the immune system. *Theoretical Biology and Medical Modelling*, pages 1–16.
- Tun, K., Dhar, P., Palumbo, M., and Giuliani, A. (2006). Metabolic pathways variability and sequence/networks comparisons. *BMC Bioinformatics*, 7(1):24.
- Wilson, R. J. (1999). Graph Theory. In James, I. M., editor, *History of Topology*. Elsevier, New York.
- Woodward, J. (2003). *Making things happen*. Oxford University Press, Oxford.
- Woodward, J. (2019). Some varieties of non-causal explanation. In Reutlinger, A. and Saatsi, J., editors, *Explanation beyond causation: Philosophical perspectives on non-causal explanation*. Oxford University Press.
- Zamir, E. (2016). Integrative systems and synthetic biology of cell- matrix adhesion sites. *Cell Adhesion & Migration*, 10(5):451–460.