Synthese (2010) 177:213–245
DOI 10.1007/s11229-010-9842-z

Topological explanations and robustness in biological sciences

Philippe Huneman

Received: 9 June 2010 / Accepted: 1 September 2010 / Published online: 3 December 2010

Abstract This paper argues that besides mechanistic explanations, there 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. I first investigate topological explanations in the case of ecological research on the stability of ecosystems. Then I contrast them with mechanistic explanations, thereby distinguishing the kind of realization they involve from the realization relations entailed by mechanistic explanations, and explain how both kinds of explanations may be articulated in practice. The second section, expanding on the case of ecological stability, considers the phenomenon of robustness at all levels of the biological hierarchy in order to show that topological explanations are indeed pervasive there. Reasons are suggested for this, in which "neutral network" explanations are singled out as a form of topological explanation that spans across many levels. Finally, I appeal to the distinction of explanatory regimes to cast light on a controversy in philosophy of biology, the issue of contingence in evolution, which is shown to essentially involve issues about realization.

Keywords Explanation · Topology · Mechanisms · Evolutionary contingency · Robustness · Realization

It has often been argued that to explain means to point out the causes of some event, phenomenon, state of affairs, etc. (e.g. Salmon 1984). Philosophers of science often debated this claim, and classically the view that explanations are causal is contrasted with the view that being explanatory means being unifying (e.g. Kitcher 1989; Friedman 1974). In this paper, I will focus on a kind of explanation that is not

P. Huneman (✉)
IHPST, CNRS/Université Paris I Sorbonne, Paris, France
e-mail: philippe.huneman@gmail.com

straightforwardly causal,[1] and which becomes more and more pervasive across many disciplines, i.e. ecology, social sciences, molecular biology, evolutionary biology or cell biology. I remain neutral regarding the stance of unifying explanation; however I will argue that this kind of explanation displays specific features which preclude its entering into one of the classical stances. The kind of explanation is "topological explanation", which is a kind of explanation that abstracts away from causal relations and interactions in a system, in order to pick up some sort of "topological" properties of that system and draw from those properties mathematical consequences that explain the features of the system they target. They are pervasive in recent important styles of modeling: the idea of neutral systems (neutral theory of ecology (Hubbell 2001) is a major theme here, but not the only one); the ever-increasing use of networks and graphs salient at all levels of investigation. The present investigation characterizes this type of explanation by specifying the entailment relation between some specific properties of systems and the explanandum of interest, shows by appeal to examples from ecology how topological explanations are elaborated and why they are pervasive, contrasts them with other explanatory regimes and explains how taking them into account has important consequences for issues in the philosophy of science.

I will first give an ecological example of topological explanation—about the diversity-stability hypothesis-, and then show its pervasiveness across many sciences. This will allow me to illuminate the differences between what I will call "mechanistic" and "topological" explanations, and sketch their possible articulations in practice, as well as their correlation with specific types of realization. The second section considers a major set of topological explanations, namely the explanation of robustness of various kinds of living systems: it will enable me to demonstrate the necessity of topological explanations in the biological sciences, as well as their pervasiveness. I show then that considering this explanatory regime is relevant if one wants to solve several philosophical problems about evolution, such as the problem of contingency in long-term evolution.

1 Topological vs. mechanistic explanations

1.1 Topological properties as explanatory

When we explain a phenomenon, we generally consider how the system from which it is a property, an outcome, a characteristic or a consequence behaves. We may try to determine causal relations, or to subsume several features of the system and its functioning under some laws of nature—those two concerns not being mutually exclusive. The system considered may have many properties. Some of them concern how, to put it vaguely, it fills a space; how parts of the system are located regarding one another, and whether those relations can still hold under some continuous deformations of the system (and which ones). Those are what I call the topological properties of the system. But these "parts" are not restricted to the physical parts of the system as a

[1] In the sense that their being «causal» depends upon the view of causation one adopts, unlike some other types of causal explanations.

material entity; they can be parts of a more abstract space, possibly a mathematical hyperspace, and then be involved in such topological properties. For example, an ecological community is a set of individuals of various species: prima facie its parts are the organisms; but on a more abstract level, one can see the species as parts. Those species have relations, for instance they prey on one another. The schema of their predation relationships is therefore a feature of the system; it takes place in a space, and can display topological properties. Let's write (A, B, C, D) a community, (A-> B) meaning that A preys on B. We could have the structure of predation relations A->B->C->D (type 1a), or we could have a closed structure A->B->C->D**->A** (type 1b), and, besides those structures where each prey has one predator, a structure with iterations A->B->C->**A->D**->C (type 2). Those schemes are different because closeness and openness, as well as iterated and non-iterated schemes, are topologically different. Larger communities can be equivalent to type 1a or type 1b, or type 2 communities. And different communities may be equivalent with respect to their displaying one of those three kinds of structures. Clearly the difference between those structures may have consequences upon the evolution of frequencies of the species—for example, the iterated scheme implies that type 2 communities will not react in the same way as type 1 communities in the face of the extinction of species D. When, among the consequences of some topological properties, stands the behavior, property or outcome we want to explain, then I say that we have given a *topological explanation* of our explanandum.

Here is a small example, taken from evolutionary theory (Wilke et al. 2001), of an actual topological explanation. Richard Lenski and colleagues studied evolution of strains of bacteria with specific fitness landscapes under various mutations rates, in order to capture various schemes of evolution by natural selection in different settings. It appeared that in changing mutations rates, a specific distribution of phenotypes is surprisingly selected against the fittest. The idea is that two types of topological structure of populations in fitness space can be distinguished: a population A with the fittest individuals, which could be represented by a peak at the optimum; and a population B whose fitness curve has a flat shape, where the large majority of individuals are very close in fitness, which could be represented by a kind of plateau. If two strains of bacteria with those profiles are put in competition, survival of the flattest (against the fittest) is common under some conditions because of the compared properties of the fittest and the flattest. In the A type, a random mutation has many chances to occur in the steep of the fitness peak, and to decrease the fitness; then such population can only give rise to mutations which will fall off sharply from the fitness peak. Concerning the flattest, a mutation has many chances to still be on the plateau. Thereby in identical environments but *at high mutation rates* μ, this flatness gives an advantage over the sharpness, since the detrimental mutations will accumulate faster in the sharp type A. It is as if mutation rate defined the selective pressures.[2]

In this example, the properties considered are stricto sensu topological properties: flatness vs. sharpness. They are topologically different because one is continuous while

[2] "Selection in an asexual population should maximize the overall replication rate of a cloud of genotypes connected by mutation, rather than any one genotype that has the highest replication." (Wilke et al. 2001, p. 331).

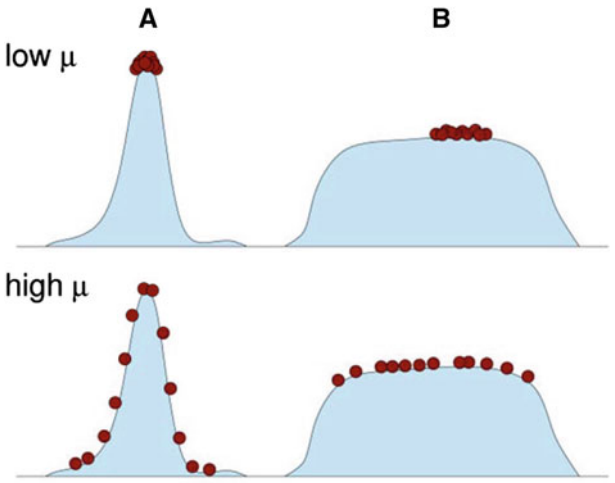

Fig. 1 Survival of the flattest. μ is the mutation rate, vertical axis is fitness (after Wilke et al. 2001)

the other displays a singularity, so no continuous transformation can change one into the other. They are properties of the system, in the sense that the system considered (the strains of bacteria) also incorporates an abstract feature, the fitness distribution of possible mutations, located not in the actual space, but in a mathematical one. Flatness and peaks are two properties possessed by shapes in this abstract space. The probability that a mutation displays a large drop in fitness as compared to the current alleles is clearly a consequence of the shape of the curve, because it only depends on the fitness distribution of possible mutations. And this probability is what determines the evolutionary fate of each strain. So in this case topological properties of the two strains explain the different evolutionary outcomes (Fig. 1).

More precisely now, I call a topological explanation, an explanation in which a feature, a trait, a property or an outcome X of a system S is explained by the fact that it possesses specific topological properties T_i. What "possess topological properties" means is the following: S has elements, parts, features or capacities, moments of its regular behavior, or of the set of its possible behaviors, which are likely to be represented in a graph, a network or a variety S′ in a space E. (For instance, the phase space of the system can be such a space.) In the above example, the fitness landscape of possible mutations is S′. In the imaginary ecological community above, it was the graph of predation relationships.

You can define a topological space on E. Once this topology is defined, S′ will have topological properties T_i, namely, properties which specify its invariance under some continuous transformations, and which will determine equivalence classes between all structures S″ homotopic to S′. Or, if S′ is a graph, you can specify some properties of S′ (e.g. connexity; cyclicity, etc.[3]) which will define an equivalence class, and distinguish S′ from other graphs S* not having those properties. Topology or graph theory

[3] I provide more examples in the next paragraph.

thus provides tools to classify S′ within specific categories defined by the properties they have. From now on, I call "topological properties", those properties that are either proper to subsets in a topological space or to some graphs and networks.

Topology and graph theory stricto sensu are not the same mathematic subfields. However, one knows that they have the same origin, the "problem of the seven bridges of Königsberg" set by Euler, a proximity which is quite meaningful; more substantially, they study the properties of structures of sets of points in an abstract space, and allow one to classify those structures into distinct equivalence classes. Here the relations between parts (points, nodes vertices, lines, open intervals) of S′ define such a structure of a set of points, and determine the relevant properties T_i. Those two fields, topology and graph theory, have indeed been joined in a topological graph theory, currently developed for its own sake (Gross and Tucker 1987). In the following, I consider both fields together because they are working in the same way with respect to what I am calling topological explanations, namely they specify the nature of the properties whose existence entail the fact that the explanandum happens[4]. S has topological properties in virtue of its relation to S′ and its elements and relations—exactly like Lenski's bacteria strains have the properties of "flatness" and "sharpness" in virtue of their essential relations to their mutations distribution.

In a sense, S realizes the topological properties T_i, and many S_j can be said to be equivalent because they realize the same T_i (for example, belonging to the same equivalence class defined by homotopic paths in some space), and their equivalence may play an explanatory role in some explanations, as we will see. However, the realization here is not the same as the realization involved when we say that atoms of carbon realize the properties of hardness. Gilllett (2010) usefully distinguished two senses of realization—the M-realization, of which hardness is a typical example, and the A-realization, of which we are speaking here.[5] A-realization occurs because of an isomorphism between entities in the world, here the parts of the system S, and possibly states of its behavior, etc., and a mathematical entity, here S′. Bacteria strain A realizes flatness in this sense. Interestingly, the chemical and physical properties of the bacteria are not involved in this relation, unlike cases of M-realization where it is the properties of carbon atoms (covalent liaisons, etc.) that underlie the realization of hardness. Because there are two very different senses of realization at stake here, one understands that the realized properties may not be explanatory in the same way, when it's about hardness and when it's A-realization of topological properties. I will elaborate on this difference below.

These T_i have many consequences, especially, they may constrain the possible transformations of S′, for example because all continuous transformations should lead

[4] Some of the properties I am considering are not topological but geometrical and instead of calling them "topological" I could have called them "formal" explanations (from encompassing both topological and geometrical explanations). Although this would be accurate, it would be very confusing given the extant literature on formal causation, explanation, etc.

[5] "X is taken to A-realize Y if the elements of X map onto, or are isomorphic with, the elements of Y. This notion of 'realization' is commonly utilized with formal models and the relata of such 'realization' relations are largely unconstrained because A-realization simply holds in virtue of a mathematical mapping or isomorphism" (Gilllett 2010) This paper largely develops A-realization, whereas Gillett (2007); Gilllett (2010) are essentially about M-realization. I thank Carl Gillett for his help on those issues.

to some homotopic structure S″. Among the consequences, is also the fact that the behavior of S, because of the topological properties of S′, will in many respects be different from the behavior of systems S* such that their associated shape in space E, S'_2, belongs to another topological equivalence class. Hence in some cases, the explanandum will be wholly explained by S′ having T_i, for example if one wants to explain the difference between two outcome behaviors, or the fact that some systems seem to behave in analogous ways, or that several states of a system are equivalent from some viewpoint of interest.

The topological explanatory relation implies that whatever possible process B_j occurs to S that involves some or all elements or parts of S, no B_j is sufficient to account for X, but the simple fact that S realizes T_i entails as a consequence the fact that S has X. We therefore have a topological explanation when none of the B_j is needed to explain. It may be that a given X causally results from some B_j, but that what explains X is not the B_j itself; the reason would be that another causal process B_k also produces X, but that the topological properties T_i which constrain in the same way all processes B_i, entail that whatever the process B_j, the outcome will be of type X, so that the causal process B_j itself does not make any difference (conditionally on the fact of T_i) to the outcome and is therefore not explanatory. And unlike a mechanism or a process, a topology is not something that takes place in time: so those kinds of explanations are of very different kind than the uncovering of mechanisms or the understanding of specific processes. Explanation of X goes like a relation of entailment between topological properties T_i and X or features of X, and not like the display of a mechanism from which X would be a temporal outcome.

I draw this contrast between B_j and T_i to emphasize the fact that topological explanations are exclusive of explanations which pinpoint some mechanisms, or some specific causal interactions between elements of S underlying mechanisms, as explanatory relevant. This difference parallels the one between the A-realization in topological explanations and the M-realizations in mechanistic explanations according to Gillett. What are at stake when someone unravels a mechanism or process are those specific causal relations uncovered between elements. However, in the explanations considered here, those specific causal relations are irrelevant; what counts is the fact that there are relations in S′, the associated shape of S in space E, because the patterns of relation in S′ can define a topological structure with its properties realized by S and then be explanatory of X.

Unlike in mathematics (e.g. Randrup and Rogen 1997), however, the term "topological explanation" is not widely used in natural sciences except in some regions of physics. For example, authors may design what they called a topological explanation of charge and mass of particles (Arcos and Pereira 2007), which means that those properties are accounted for by the topological properties they display in the model; however, given that they are mathematically defined from the beginning it is quite natural to talk about their topology. I found an instance in the scientific literature in which the use of the term "topological explanation" was put forward; namely a paper concerning some formal features of the internet (Park and Newman 2003).[6] I

[6] "Rather than supposing the anticorrelation of vertex degrees to be the result of some specific social or engineering constraints on the construction of data networks, they suggest instead a topological explanation.

claim here that many explanations in biology, cell biology, and the social sciences are topological explanations in this same sense, i.e. the properties that are *explananda* are explained by reference to the topological properties of the system, no matter the processes and mechanisms happening in those systems. I will now give an example of such explanation, which is pervasive in community ecology, and then refine the notion of topological explanation with respect to mechanistic explanations.

1.2 The topological explanations of ecological stability

The relationship between the diversity of species in an ecological community and the stability of this community is a longstanding debated topic in ecology (McCann et al. 1998). Although often supported by ecologists, the idea that diversity would yield stability was not demonstrated until the 1970s. May (1974) formally demonstrated that in fact, if you consider that species are randomly connected, more diversity would imply less stability in terms of the constancy of frequencies of individuals. This prompted ecologists to realize that the exact meanings of stability and diversity were central issues. Functional diversity, for example, cannot be equated with *number* of species, since one could have many species preying on the same preys, being therefore functionally equivalent. And as it turned out, if stability is understood as the constancy of some property like the biomass (Tilman 1996), then diversity enforces it. So many varieties of diversity-stability hypotheses, often involving complexity (Pimm 1984), flourished during the past two decades (Pimm 2001). Some of those (Solé and Montoya 2001; Dunne et al. 2002a; McCann et al. 1998; Montoya and Solé 2002; Montoya et al. 2006—after Yodzis 1989; Dunne 2006) considered only three parameters: the number of species and their average number of links, the number of connections realized between species as compared to the number of possible connections (connectance), and the distribution of the connections between species. These hypotheses provide examples of topological explanations.

The general aim of these studies consists in specifying for a community S the network S′ of relationships between species, each interaction being represented by a link, and then inferring from properties of this network some properties concerning stability, resilience, etc. In ecology, two species can have several kinds of causal relations: A can prey on B, be the prey of B, compete with B, be mutualistic or commensal with B, or parasitic on B, for example. However, in those networks, 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 (Montoya et al. 2006).

If we want to explain why an ecological community is stable, or which communities are more likely to be stable, we can consider how they react to invading species, or extinction of some of their species. The stable ones are the communities where those events don't have many consequences upon the repartition of species or the frequen-

Footnote 6 continued
Using computer simulations, they show for a network of the size and degree sequence of the Internet that the requirement that there is at most one edge between any pair of vertex induces degree anticorrelations very similar to those observed."

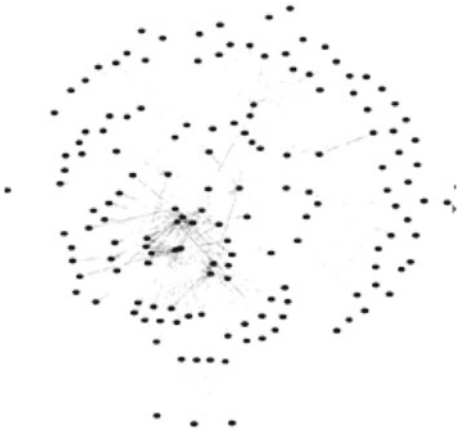

Fig. 2 An ecological network in a community (Silkwood Park, after Solé et al. 2002)

cies of individuals in those species; the less stable are the ones that undergo dramatic changes in composition. There can be a simple ecological explanation for that: suppose that an ecological network has a few hubs that are vastly connected, and many species that are not very connected. Concretely, you can identify an omnivore species, therefore being a hub connected to many species, and then species specialized on only one prey, which will comprise isolated nodes. Now a random deletion of species will be more likely to touch a weakly connected node: suppose that each species is equiprobably likely to be deleted, and that you have two hubs and 464 simple nodes; then of course the deletion has 232 more chances to delete only an isolated node. So the whole structure of links won't be altered, therefore the whole community will have some stability, in terms of response to species extinction or species invasion (Solé and Montoya 2001; Dunne et al. 2002a,b). In this example, the community S has a network S′ and then realizes the property T of being "two-hubbed", which entails this fact about probability and then its stability (Fig. 2).

A network with such pattern is not a random network, like ones mostly considered by May (1974), who contested that diversity enforces stability. Most of the research about diversity and stability consisted therefore in exploring cases where networks of links depart from May's assumption that connections are random. So let's consider now actual explanations of stability in communities. The mathematical study of networks highlighted some precise kinds of graphs, for instance scale free networks, where the repartition of links on each node follows a power-law[7], and small worlds, where two random nodes are always never further than a specific number of links, and where the degree of clustering is high[8]. Some processes to get scale free networks and small worlds are known: the construction rule known as preferential attachment or "rich get richer" (i.e., the probability of having a new link is proportional to the

[7] A scale free network is such that if the number of connections is $x = a^{M-n}$, the number of nodes having x connections is a^n. This means that the network is invariant if you zoom in or zoom out—hence its invariance of scale.

[8] A cluster is a subset of nodes connected with each other to a degree far higher than they are connected to other nodes.

number of existing connections) yields power laws (Albert and Barabasi 2002), and then scale-free networks; adding some random links to a graph where nodes are connected only between neighbors yields small worlds (Watts 2003; Watts and Strogatz 1998).

A community whose associated network S′ is an ecological scale-free network will have the property of stability against random extinction mentioned above. In this sense proving that a community is so structured triggers a topological explanation of its stability. This property is explanatory sufficient and one does not need to delve into the details of who is parasitic on whom, who preys on whom, etc., to establish the causes of stability. The details of these processes here do not make a difference, and the stability is not what eventually results from a sequence of causal interactions involving the species[9]. Interesting parallels have been drawn between those networks and the Internet, as well as cell metabolic pathways and networks of words in texts (Solé et al. 2002), or between ecological networks and financial networks (Levin et al. 2008), especially regarding the stability requirements: all those systems A-realize the same topological properties, thus their stability is identically explained.

Considering small worlds, Montoya and Solé (2002) have shown that their structure happens in some communities and also yields forms of stability, because if you break some links between species, the structure allows for alternative links and therefore preserves the general stability. For example, suppose that a trophic web is a small world: between two given species you always have few nodes; if you delete one (e.g., you suppress the predator of A preyed on by B), in the cluster near the two unrelated nodes are many connections to be found so you can still find an alternative pathway. Hence, the property of the interaction network S′ proper to a community S, of being a small world, explains why it is resistant to species deletion and invasion, thus why it is stable.

In the same perspective, ecologists elaborated the idea, ancient in social sciences (Granovetter 1973), that weak links between individuals weakly connected (which means, link between individuals belonging to distant clusters) are likely to make everyone connected in a kind of small world. To this extent, in the case of food webs, weak links proved to be important in stabilizing the communities (McCann et al. 1998; Berlow 1999) because they inhibit the oscillatory relations which happen with strong relations of competition or mutualism or predation between few species; hence many weak interactions in food webs make chaotic dynamics unlikely, so that "weak interactions may be the glue that binds natural communities together" (p. 797). In such a case we see that the direction and the strength of interactions are explanatorily relevant, but not their nature (e.g., whether species are predators, preys, mutualists, competitors etc.).

What accounts for the stability of the community, under many senses of stability, including "resilience" (Holling 1973)—the ability to come back to a state of equilibrium after some perturbations—are the properties of this net of interactions weighted

[9] In fact, ecological networks are not exactly scale-free networks or small world networks, but many of them, fulfilling the rough characterization just given, resemble scale-free networks (Dunne 2006) or are truncated scale-free networks, which is enough to explain topologically their stability.

by their strength (see Tylianakis 2008 for a recent overview)[10]. Many causal interactions are possible between all the individuals in the community, but this web of relations gives rise in every case to the same property of stability, affecting sometimes the community, sometimes a metacommunity (Maser et al. 2007). Suppose indeed that many different communities S_i of different species x_k have the same interaction network S', where a link is defined both by its relata and its strength. This network is a specific graph in an abstract space, and specific properties regarding weak links and high degree of clustering will entail the stability of the community (because of their anti-oscillatory consequences mentioned above). No matter which species x_k in a given S_i are connected through which links in the real web of relations realizing S', what matters for the outcome is only the general shape S' of the web of links, and the extent to which it has high degree of clustering and weak links, and this holds for those different communities S_j. In this sense, explaining by considering the topology of food webs, or more generally ecological links, abstracts away from the nature of causal relations, in the sense of the genuine ecological relations between definite species and the causal mechanisms at work within the community. Suppose that in S_1 half of the links are of predation, and in S_2 half of those links are of parasitism, and that the species implied in S_1 and S_2 are differing by 30%. If the interaction network S' is quite the same (more precisely, S'_1 and S'_2 belong to a same equivalence class regarding a property of having weak links and high degree of clustering), then stability in S_1 and S_2 is explained; a detailed understanding of the processes B_1 in or B_2, which are different in S_1 and S_2, will not add anything to the understanding of why S_1 or S_2 are stable. In this sense, mechanistic explanations are superfluous relative to this explanandum.

In the following Section, I detail the relationship between those topological explanations and mechanistic explanations in general, by considering what is included in this idea of being sufficient relative to one explanandum.

1.3 Clarifications about the difference between explanatory regimes

I said in the beginning that topological explanations differed from explanations that are straightforwardly causal. Yet someone could still say that the topological properties are in this sense a cause, because obviously in a counterfactual sense of cause, they cause the explanandum. Indeed, to some extent the discussion here is merely semantic. Regarding those people who say all explanations are causal explanations, I specify that topological explanations compare indeed to explanations which *describe actual causal relationships* between entities, and therefore describe what Machamer et al. (2000) call mechanisms. If one says that the topology of food webs causes some resilience properties, this however contrasts with a mechanistic explanation that would specify how the disturbance of some interaction between two species would provoke another kind of interaction that in the end would restore the same pattern of community. In the

[10] Strength is sometimes strength being "measured roughly as the proportion of individuals of a species at the lower trophic level fed upon by, or interacting with, a species at the higher trophic level" (Tylianakis 2008, p. 224). Concerning in general what the strength of an ecological interaction would be and how it can be measured, see Ulanowicz (2002) and Huneman (2010)

topological explanation you do not have to consider the mechanism triggered by the perturbation, because for example you already know that whatever happens, the preys and the predators, the parasites and the hosts, etc., will manage to stay connected. I will now state this sharp contrast between these two kinds of explanations, and then show how they can be articulated.

In the terms of Machamer et al. (2000), a mechanism is defined by some entities which have proper activities and some regular interactions, and which results in a general activity of the whole that regularly yields definite outcomes from definite inputs. So, on their account, "mechanisms are composed of both entities (with their properties) and activities"; "activities are the producers of change"; "entities are the things that engage in activities"; "activities usually require that entities have specific types of properties" and "the organization of these entities and activities determines the ways in which they produce the phenomenon" (2000, p. 3). They claim that describing such a mechanism explains the phenomenon.[11] Identifying activities, in their view, is what scientists do instead of discovering laws.[12]

The contrast between kinds of explanation is easy to grasp then. What is explanatory here in mechanisms is the essential link between an entity and what it does; it is in virtue of X's Phi-ing, that Y and Z will for example be triggered to do J and K, and then make the system S likely to be Psi-ing as a result of combined J and K. As Machamer et al. state, "mechanisms are identified and individuated by the activities and entities that constitute them, by their start and finish conditions, and by their functional roles" (2000, p. 4). However, in topological explanations, there are no specific activities listed: no matter what the species do, whether they prey or not, whether they Psi or they Phi, and on whom etc. Stability rather occurs because of some network property of S'. Suppose that two systems S_1 and S_2 have the same associated shape S' in abstract space but that the relata of links in S_1 are J and K, and in S_2 are J_2 and K_2, with their distinct associated activities. Yet the topological properties of S', realized identically by S_1 and S_2, will not be affected by this difference of activities and properties; therefore the identification of mechanisms in S_1 and S_2 is not relevant for explaining their outcomes. In my first example, the probability of a mutation staying on the plateau is not a mechanism of mutation or variation; it is simply a constant property of the system due to the topological property of S', which is namely the distribution of mutations, and this is true irrespective of the reasons for the mutations, i.e. the exact mechanisms that produce the mutations and their specific distributions of fitness.

Another point of contrast is that the explanandum of a mechanistic explanation is determined by stating the "set-up" and "termination" conditions. For instance, in neuroscience explaining the NDMA receptor–ionophore complex: "activation of the NMDA receptor is a means of transforming an extra-cellular chemical signal (born by neurotransmitters) and an intracellular electrical signal (born by ion fluxes in the

[11] "To give a description of a mechanism for a phenomenon is to explain that phenomenon, i.e., to explain how it was produced." (Machamer et al. 2000, p. 4)

[12] "A mechanism is the series of activities of entities that bring about the finish or termination conditions in a regular way. These regularities are non-accidental and support counterfactuals to the extent that they describe activities." (ibid. p. 6)

cell)", which are the set-up conditions, "into an intracellular chemical signal (born by intracellular ions and molecules)" (Craver, forth.), which is the termination condition. The chronology of sequences through which entities act is crucial to explain the production of this termination condition—for example here, "the depolarization precedes the release of the Mg+ ions" (ibid.) To this extent, another difference between topological and mechanistic explanations is made manifest: the temporal sequence between events in topological explanations is irrelevant, since switching the order of events does not affect the explanandum.

Now Craver explains that mechanistic explanations consider a system along two dimensions, the interlevel and the intralevel ones. If M_i are entities at a same level of a system S, hence components of S[13], intralevel explanations consider how causal relations account for the production of the termination condition of a system S by activities of entities such as M_i. Yet S often enters into higher-level mechanisms, and hence the outcome of a given M_i will enter into an explanation spanning across levels. So, as Craver and Bechtel (2007) indicate "each new decomposition of a mechanism into its component parts reveals another lower-level mechanism until the mechanism bottoms out in items for which mechanistic decomposition is no longer possible". And reciprocally, "higher levels of mechanisms are aggregated (i.e., built up from) or composed from parts that are organized into more complex spatial, temporal, and causal relations" (ibid).

Craver and Bechtel argue that the only *causal* relations are intra-level, but interlevel explanations display relations of *constitution*: "there is a temptation to say that the activation of cyclic GMP phosphodiesterase, which catalyzes the conversion of cyclic GMP to 5¢-GMP, causes rod cells to hyperpolarize, which in turn causes the eye to transduce light into neural activity. But the activation of cyclic GMP phosphodiesterase is part of the activity of depolarization, which is part of the eye's transduction of light" (2006, p. 15). In general, M_i causally contributes to the outcome of S but it constitutes a high-level activity within which S itself is a contributing entity.

A full mechanistic explanation therefore embodies both interlevel constitutional and intralevel causal explanation. The building-up relation between two levels, for example the way activation of cyclic GMP phosphodiesterase *constitutes* and does not cause transduction of light, is a kind of realization. But it is rather an M-realization, an asymmetric "ontological determination relation between causally individuated property instances and/or properties" (Gilllett 2010) Until now, when contrasting them with topological explanations this section considered intra-level mechanistic explanations. Yet the complete picture of mechanistic explanations actually includes interlevel mechanistic explanations, which envelop realization relations—namely, M-realizations, which contrast with A-realizations enveloped in topological explanations, as emphasized above. The mechanistic explanation is driven by "constitution", while the other uses features of abstract descriptions.

However, in a full mechanistic explanation it could be possible for a realization relation to be an A-realization, if the activity of a component of a mechanism has a topological explanation. For example, suppose hypothetically that a sort of Gaia

[13] "X is a component in a mechanism if and only if it is one of the entities or activities organized such that S Psi-s." (Craver and Bechtel 2007)

hypothesis holds; ecological communities would compose the system Earth, and they will contribute to its functioning and homeostasis. The stability of a community could have an important role in the global mechanisms; and such stability may have a topological explanation, as we know. Therefore here an A-realization will be involved in the general interlevel explanation why a Gaia hypothesis holds.

1.4 Nuances in the picture

Notwithstanding this contrast between causal and topological explanations, in general one could also say that there is a continuum between two poles, one that consists in unraveling mechanisms without regard for any topological properties and one that is purely topological. More precisely, when I said that the nature of causal relations is not explanatorily relevant in the network explanations of ecological stability, of course it did not mean that *just any relation* is relevant. So there is, among all the infinitely possible relations, and moreover, among all the causal relations, a set of some (e.g. in the community ecology example: ecological interactions) that are equivalent regarding the explanation—in the sense that no matter how they interact, what is explanatory is that interactors are related and that their relation maps onto a link in the abstract space E, and that they are the ones that map onto the associated shape S′ in the abstract space E where a topology can be defined, or where graphs are considered (see Sect. 1.2). Once the set of relations that is, as a whole, explanatory relevant, is identified, those relations are the ones which enter into the network, the graph, or any structure which will possess topological properties.[14]

One could say that when *all* the relations are explanatorily equivalent and enter into S′ as nodes, vertices, points or sides, we have a pure topological explanation; yet, in my ecological diversity-stability example, we have already identified a set of causal interactions, in order to elaborate the topological explanation. A pure mechanistic explanation would consist in considering that all differences between causal interactions are relevant. So it is plausible that a more realistic picture of explanation would see a continuum where I emphasize a contrast between two types.

Finally, how are mechanistic and topological explanations related to one another, and how can they be articulated in actual scientific practice? I am not claiming that topological explanations of the stability of communities are the only types of explanation that may be offered to explain the property, nor am I claiming that such stability is not subject to alternative kinds of explanation. Topological explanations are different from mechanistic explanations, but not independent from any mechanisms. I considered here explanations of stability of some communities by an appeal to the properties of interaction webs. They answer the question: “why is this community stable?” But suppose I want to know why this community has precisely this or that stability-enhanc-

[14] So clearly, the properties are first predicated of this representation, which captures an aspect of the system. They are not immediately grasped with respect to the system as such, but their determination depends on the specification of one aspect of the system that is represented. However, there is nothing more problematic here than in any case where we consider a representation of a system and ascribe properties to the system from its consideration: for example when the solutions of differential equations of motion are used to compute the position of a particle.

ing network of interactions—e.g. a scale free network? This question may warrant a mechanistic explanation of, for example, preferential attachment, which is the fact that the probability of having a new connection is proportional to the already extant connections. The mechanisms through which individual organisms forage, look for preys, compete with each other etc., all contribute to explain why preferential attachment as an outcome of those aggregated processes will hold. So the explanatory topological property itself, which is a property of the community, may in turn be the object of a mechanistic explanation.[15] And even more generally suppose my question is: why are scale-free networks and truncated scale-free networks so common across ecological networks? Then the answer will have to take into consideration the most general properties of ecological interactions—the ones that may yield phenomena like preferential attachment—and once again possibly invoke mechanistic explanations. So a first kind of relation between mechanistic and topological explanations is that they can be two stages of a complete explanation of the same phenomenon, related in a diachronic way: why does phenomenon X exists? Because of some topological properties T_i of the system X; and why does S have T_i? Because of some mechanism proper to S, its surroundings and parts, etc.

A second type of distinction between the two types of explanations may also be drawn. What I call topological explanations here may be of two sorts. Some of them, like the explanation of stability through properties of scale-free-style networks, are explanations per se. Others pinpoint topological properties as explanatorily relevant, in the sense that they put constraints on whatever is likely to occur in a system, but the process itself still plays a role in the explanation, so that explaining means uncovering topologically constrained mechanisms. Different topologies may result in differences with respect to what the same mechanism will do. Suppose that we have a family of systems S_i, which have an associated shape S_i' in a space E. There is a mechanism M that explains a property or an outcome etc. in which we are interested, and which works in the same way in all S_i (because those S_i have the same entities and activities). Now, suppose that the outcome of the mechanisms will be X_j if some topological properties of S_i are T_j, and X_k if they are T_k—I call it a constraining topological explanation, where mechanistic and topological explanations cooperate to produce the full explanation.[16]

Constraining topological explanations have been exemplified by recent studies on the evolution of cooperation. Evolutionary biologists have indeed been interested since the 1960s in the processes which may yield cooperation in populations, given that evolution prima facie seems to increase the frequency of behaviors which benefit that

[15] Clearly, the same reasoning holds about the "survival of the flattest" case—there should be mechanistic explanations of why the distribution of mutation is flat in the case of one strain, given in terms of properties of nucleotides, chemical reactions, etc.

[16] Indeed, Machamer et al. in passing acknowledged such possibility when they mention the causes of the activities proper to entities. They say: "The neurotransmitter and receptor, two entities, bind, an activity, by virtue of their structural properties and charge distributions. A DNA base and a complementary base hydrogen bond because of their geometric structures and weak charges." (2000, p. 3) But how is the "geometric structure" or the "structural properties" likely to explain why neurotransmitter or DNA and complementary base bind, given the charges hence the electric processes, if it's not by constraining the electric mechanisms occurring at the atomic level?

population that exhibits the increase rather than some other population. Cooperation by definition benefits the others, and indeed altruism, inter and intra-species, is frequent in nature (West et al. 2007). The key consideration is that if altruists are, by one reason or another, interacting with individuals prone to be altruist with them in some manner, then altruism is likely to evolve. Recent studies emphasize the fact that the topology of networks in which individuals interact is able to ensure such a feature, namely altruists being more likely to meet other altruists (Nowak 2005). Many biologists studied this fact in terms of constraints given by population structure: often the reason for this is limited dispersal, which compels kin to stay together, and kin are likely to be altruistic with one another in proportion of their relatedness, as is stated by the famous Hamilton's rule,[17] which dominated cooperation studies for 3 decades.[18] The point for us is that when you model those phenomena in a network in which individuals are interacting, the topological properties of the network are exactly what constrains interactions to be more cooperation-yielding. As Santos and Pacheco (2006) indicate, "the diversity of connectivity patterns in a population, which translates into a heterogeneous Network Of Contacts (NOC), is efficiently explored by cooperators to outperform defectors, leading to evolutionary outcomes in which cooperators easily survive and may even dominate. NOCs exhibiting strong heterogeneities and tight connections between the few most connected agents favor the dominance of cooperation." In this example, part of the explanation of cooperation is natural selection, which, through the differential replication of individuals featuring various strategies, and through either Hamilton rule and kin selection, or preferential assortment between altruists, leads to cooperation[19]. This is the mechanism in the explanation[20]; but often non-random assortment is explained by the topology of the network: what Pacheco and Santos call "Network of Contacts". Some topologies will confer an individual higher chances to meet altruists, and that is why in such instances selection will lead to cooperation, whereas in another network it would lead to the fixation of selfishness. Topological properties are clearly explanatory relevant, because they constrain selection to work in one sense rather than in another[21]. In the evolution of cooperation, the topology of connections should be part of the explanation because it provides the context within which natural selection yields different cooperative outcomes.

Another example of constraining topological explanation, this time in ecology, is found in "neutral ecology theory" (Hubbell 2001). This theory gives the crucial role to stochastic processes (ecologically analogous to genetic drift) against selection in the creation of biodiversity patterns. Many of these studies have been done by simula-

[17] A disposition to a behavior with cost c and benefit b, directed unto an individual related at degree r to the focal individual, will evolve iff c = br. Proper interpretations of relatedness r are extremely difficult, but the intuition is that it often correlates to the degree of kinship.

[18] Hereafter cooperation and altruism are equivalent, technically they are not, see West et al. (2007) for semantic clarifications.

[19] Some argue that those two things are varieties of kin selection (West et al. 2007) but this is not our concern here.

[20] It is controversial that "natural selection" is a mechanism sensu Machamer et al. (2000), or in another sense. Here I use it in a very general way in order to contrast it with the topological explanations.

[21] When including small worlds, they can even lead otherwise divergent models (such as iterated prisoner dilemma and snowdrift games) to a same cooperative outcome (Santos and Pacheco 2005).

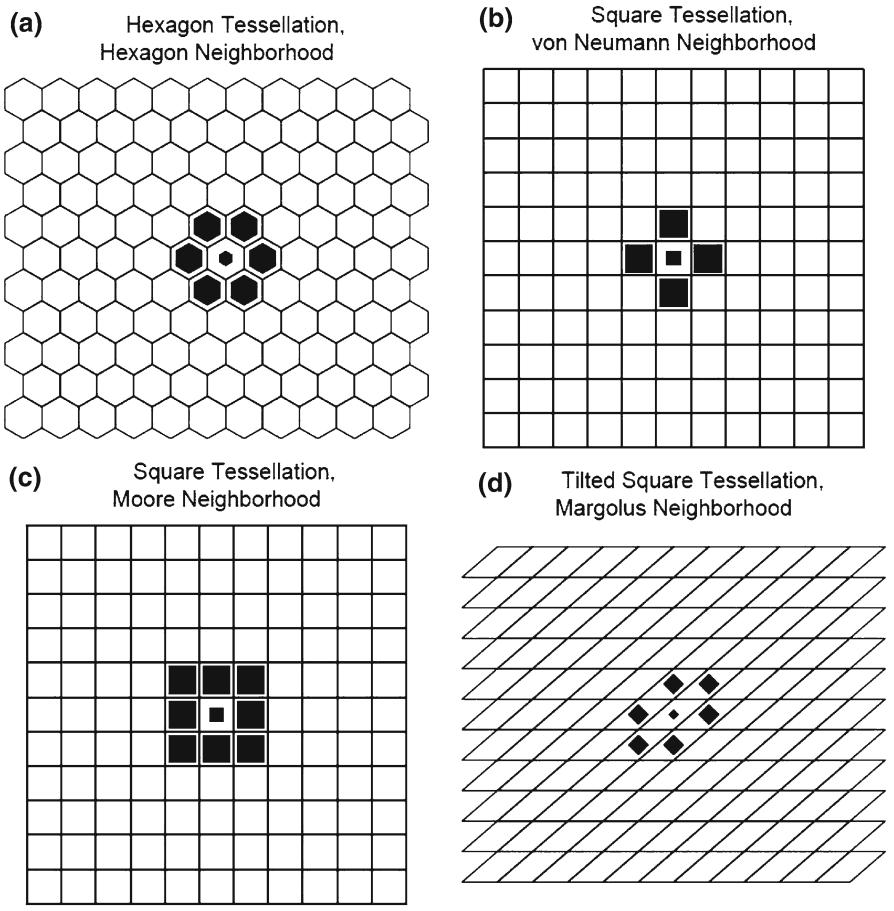

Fig. 3 Topologies of neighborhood used to simulate neutral ecological models (after White and Kiester 2008)

tion. Recently, White and Kiester 2008 paid attention to the topology of the grid onto which one designs the simulation by cellular automata. These grids are often square grids (Moore or Von Neumann neighborhoods) (Fig. 3)—which differ from each other with respect to the number of neighbors for each cell and the possibility of diagonal neighbors. The results of the same simulation, on the basis of the neutral theory, will differ with respect to the size of communities and diversity, depending upon which tessellation is used.

In this case, the many communities S_i embody one mechanism, the ecological drift (M), is modeled by stochastic algorithms. However, those S_i may have in their associated shape (here, the grid of the simulation) four topologies a, b, c, d, and the outcome of mechanism M in terms of biodiversity will be different in accordance with which topology is used. Therefore we clearly see how topology constrains the mechanistic explanation here.

Probably, to explain and to constrain explanations also define two poles—pure topological, and constraining topological, explanations—between which the topological explanations can be distributed. Having contrasted mechanistic and topological explanations and sketched how they can be articulated, I have to show on a specific set of cases how and why they differ, and how they are compatible. In the next section I will consider one main explanandum for topological explanations, taking place at all levels of biological hierarchies namely, *robustness*. Stability of ecosystems and communities in the face of disturbances and invasions is indeed one kind of robustness, occurring at the highest level of the biological hierarchy, and I have shown how pervasive were topological explanations in that context. The question is whether this is only specific to ecology, or whether it is true of the biological sciences in general. I will thereby investigate how robustness at many levels is explained, and see how mechanistic and explanations are distributed with respect to such an explanandum. First, I demonstrate that the kinds of claims I have made with respect to the ecological case study that pertain to the relationship between topological and mechanistic explanations may be generalized to this more general case. Then I consider why topological explanations are so pervasive when it comes to explaining robustness, and explain that it has something to do with the very nature of topological properties.

2 Robustness: mechanistic and topological explanations reconsidered

The stability of ecosystems provided us with a family of examples of powerful topological explanations. The explanandum, here, ranges under the rubric "stability" in the sense of "disposition to stay the same across changes, under some important ecological aspects", and the nature of those aspects define the various stability properties (e.g. resilience, constancy of biomass, etc.) that communities may possess. Given this characterization, ecological stability is however not the only kind of biological stability; the robustness of biological systems in general means something very similar. Recently we witnessed an inflation of research programs dealing with various kinds of *robustness* of systems *at all levels* of the biological hierarchy. I wish to show here that in this large field of biology the two kinds of explanations can be used side by side, and that at each level in the biological hierarchy topological explanations are used. This latter fact is quite meaningful with respect to the nature and capacities of topological explanations, and I will explicate it. I will emphasize a specific kind of explanation among them, called *neutral spaces*, which will prove relevant for understanding varieties of explanation in general.

By "robustness" here, and in order to encompass many different researches in the field, I mean keeping some parameters stable in the face of changing values of variables. All modes of stability mentioned above in ecology satisfy this characterization, the variable concerning in general the number of species, or of individuals in species. More precisely one can parse this into two kinds: a "level" definition, in the sense that change in low level variables doesn't involve change in high level variables—for instance changing the identity of several species does not affect some general properties like biomass or abundance pattern; and a "functional" definition (Wagner 2005a): robust are those systems that are able to maintain (some of) their functions in the

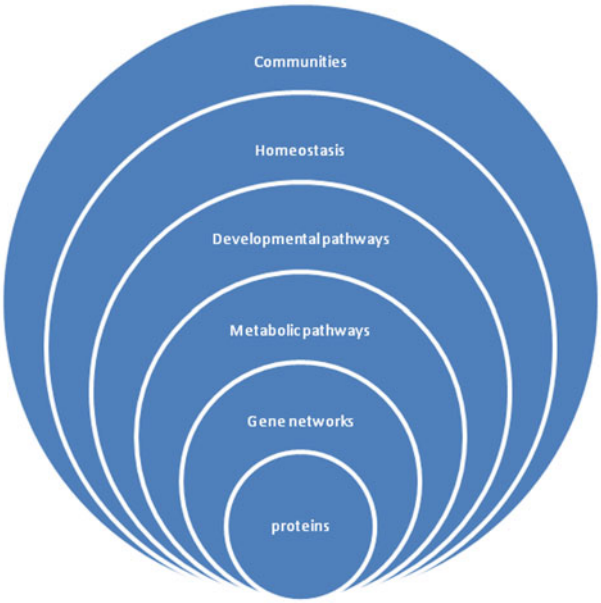

Fig. 4 Successive spheres circles of robustness

face of perturbations. Plausibly, the latter definition works better with physiological or cellular systems.

Canalization and homeostasis are time-honored well-known cases of robustness in living systems. Homeostasis in physiology falls under this second definition, because it means that organisms are able to keep vital functions efficient in the face of large environmental changes; concerning development, organisms display what Waddington called "canalization" of their processes, which also satisfies the definition. "Canalization" means that perturbations in the genetic makeup of an organism are not always likely to change the outcome of development; inversely, development shows also some robustness *vis-à-vis* the environmental perturbations, in the sense that the final product may be constant across a large range of environmental conditions.

However, properties of robustness have been found at many lower levels: within the metabolic networks of cells, at the level of the genome and its expression, and at the level of the genetic code and proteins. Very generally defined robustness properties range from nucleotides to ecosystems (Fig. 4), and they may require analogous or identical explanations in several cases. Wagner (2005a) provides a groundbreaking synthesis of those findings. I here after give two examples of robustness at low level, and in the next paragraph provide a more systematic view of robustness research according to the explanatory distinction sketched above. The last section will draw consequences about the nature of topological explanations and highlight a specific case of it.

Concerning gene networks, a striking study compared two networks of gene expressions for gene Endo 16 in sea urchins. In *Strongylocentrotus purpuratus* and *Lytechinus variegatus*, you find this gene, whose product is a protein involved in cell adhesion,

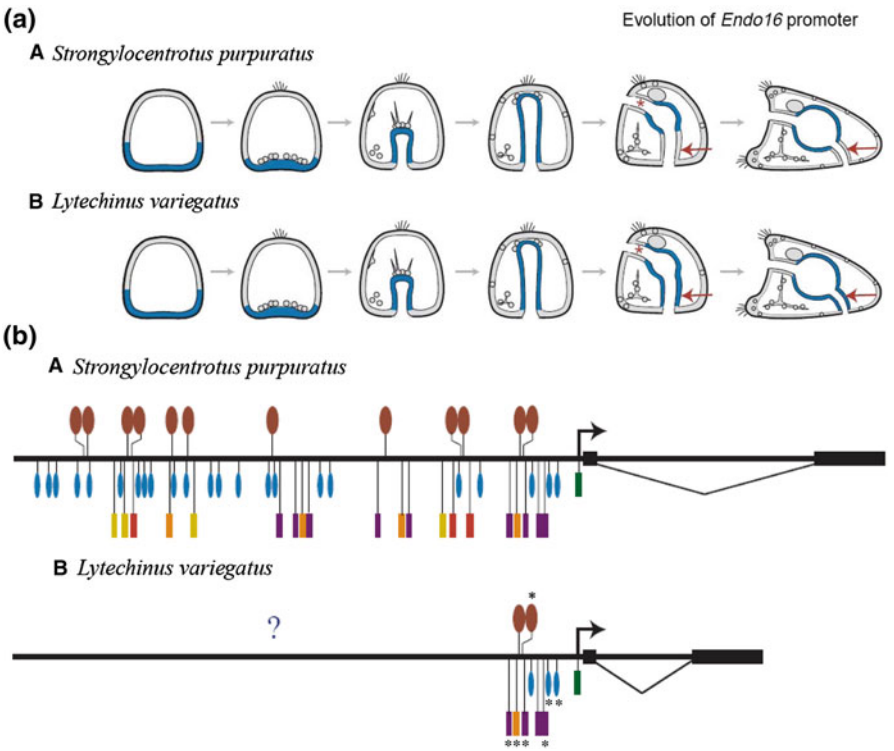

Fig. 5 **a** (*above*) Representation of the product of Endo 16 in two species of sea urchins along development; **b** (*below*) Regulatory regions for Endo 16 in both species. Reproduced with permission of *Development* from Romano and Gray (2003)

crucial in the development of several tissues. Those two species diverged 16 Million years ago. The regulation networks for the expression of this gene in those two species are therefore very different; however the small regulatory region in the most recent species corresponds to a small part of the regulatory region in the oldest species (Fig. 5b). Yet, the final product (Fig. 5a) is the same in both cases (Revilla et al. 2003). Hence everything occurs as if altering the Endo 16 regulation network in *Strongylocentrotus purpuratus* at a wide scale entails very slight perturbations of the end product. Evolution here provides an analogue of perturbation experiments on regulatory networks.

Metabolic pathways in cells display analogous features. *E. Coli* have seven central reactions involved in their metabolism; each of them involving a given flux of chemical substances. However, with mutations reducing drastically the fluxes in one of those reactions, the transketolase, up to 85%, the global outcome of the metabolism remains remarkably constant.[22]

[22] "Although complete elimination of flux through the seven essential reactions of *E. Coli* central metabolism is lethal, substantial quantitative reductions in flux may be neutral. A case in point is the essential transketolase reaction in the pentose phosphate shunt. As long as the mutation preserves more than 15% of the wild type flux through this reaction, the resulting growth rate is no less than 99.2% of the wild type growth rate." (Wagner 2005a, p. 131)

In the previous case of gene regulation networks, granted, we saw that robustness meant that taking away a very large specific part of the network is not so consequential, whereas here, it means the fact that withdrawing a non-specific large quantitative part of the system is not consequential: which 15% should be kept is not at stake here. Yet in both cases robustness needs an investigation into how the system manages to respond in such a way to its dramatic alterations. In the genome itself, changing nucleotides will alter genes but in fact often leave intact the phenotypic result. Redundancy is a potential cause of this phenomenon: if a gene exists in two copies, altering one will leave intact the other, and then the phenotypic outcome will be unchanged. Many authors also have shown that redundancy plays an important role in evolution since it allows variation on one copy whereas its function is still fulfilled by the other copies—in other words, thanks to redundancies, the system can undergo variation in its nucleotides and therefore benefit from the rare possibly favorable new outcomes, without losing fitness (since a copy of the varying nucleotides is still there and functions) (Kitano 2004). Andreas Wagner however established that an important amount of genomic robustness is *not* due to redundancy but to what he calls "distributed robustness", namely: two alternative pathways exist, mobilizing two sets of connections in the same genomic network, which buffers the outcome against change in one of the genes (Wagner 2005b). Distributed robustness may be involved in the Endo 16 case.

The constant existence of robustness phenomena at all levels call for explanations, and many have been given which can be classified according to the categories proposed in this paper. Here is a more systematic overview.

2.2 Systematic view of robustness explanations

In general, researchers ask two types of questions about robustness at all levels: (a) *proximate*: how is it ensured? and (b) *evolutionary*—which breaks down into (b1), why has it evolved? And (b2) what role may it play in evolution?

About (b1), if one considers for instance canalization, the evolution of robustness, raises the issue of the relationships between environmental and genetic robustness, and some researchers argue that the latter is a byproduct of selection for the former: " Genetic robustness may thus evolve as a correlated side effect of the evolution for environmental robustness. Since environmental perturbations often have a higher frequency and impact on fitness, they serve as the driving force. (. . .) Character robustness can be separated from, and is secondary to, character adaptation itself, but… genetic robustness does not evolve for its own sake." (de Visser et al. 2003). Nevertheless I will mostly consider here (a) and (b2).

About (a), I claim that one can partition robustness into several classes, according to the explications for robustness, and whether they are topological or mechanistic (Table 1). And it may turn out that those classes are not concerned in the same manner by the two evolutionary questions (b). To show why those two kinds of explanations are held, let's consider two examples, which concern explaining the robustness of genetic sequences against mutations. If one indentifies a feedback process that is triggered by the alteration of some nucleotide and leads to a restoration of some other genomic circuit, which then leads to the initial phenotypic state, this would clearly

be a mechanistic explanation. On the other hand, an explanation, instead of appealing to any mechanism, may simply cite some properties of the genetic code, namely, the redundancy of some triplets of nucleotides coding for amino acids. What is at stake here are the properties of the mapping between DNA sequences and proteins: due to redundancy, which is a topological property of this mapping, many random changes in nucleotides will not change anything because they will turn a codon into a synonymous codon—for example, given that AGT and ACT code for the same amino acid, mutating C to G or G to C in the second position is of no consequence upon the final protein. This explanation works in the same way as the quasi-scale-free network explanation of stability of some ecosystems, and does not need the unraveling of any specific causal interaction.

At many levels of the biological hierarchies, then, you can find these two kinds of explanations of robustness coexisting. I now detail some of them as they appear in Table 1 and then emphasize a specific style of topological explanation of robustness using a tool called neutral spaces, which can be developed at several biological levels.

2.2.1 Some mechanisms

Feedbacks. In physiology, robustness of organisms against environmental perturbations has been for a long time acknowledged and studied and researchers discovered many mechanisms producing a general homeostasis—*negative feedbacks* controlled by specific parameters (pressure, pulse, etc.). Here, robustness investigations of course overlap with the study of *regulations,* which occur at all levels. For instance, genic regulation has been discovered by the Jacob, Wolff and Monod in their work on the lactose operon, which is a system that allows a stable level of lactose in the cell through mechanisms of positive feedback involving both a repressor gene and a lactose-producing gene (Morange 2000).

DNA repair. Research on the evolutionary origins of sex—which is a major area of investigation for evolutionary biologists—led to the hypothesis that sexual reproduction could be a byproduct of a selection for *DNA repair* (Bernstein et al. 1985). The fact of diploidy and the mechanisms of recombination are such that if DNA is altered, another copy is intact, which can be a template for DNA repair. In this sense the robustness of DNA against alterations of nucleotides is produced by this ability to copy on the basis of a second haplotype as a template. As an explanation of the robustness of the genome against environmental (intracellular) perturbations or

Table 1 Kinds of explanations of biological robustness

Mechanisms	Topology
Feedback control physiological homeostasis; lactose operon)	Modularity
	Redundancy
	Small world topologies
DNA repair systems (→ fct of sex (Bernstein et al. 1985))	*Neutral spaces* (Stadler et al. 2001; Schuster et al. 1994; Wagner 2005a)

cellular noise (both altering fidelity of replication), the DNA-repair hypothesis provides a clearly mechanistic account.

2.2.2 Some topologies

Modularity reduces the effect of perturbations because they will not break the whole systems, since they have chances to concentrate upon only one module, the others being likely to still function, at least under some conditions on modularity (e.g. avoiding serial architectures). It could be argued that this is a kind of topological explanation of robustness to the extent that modularity as a property of a system made of elements that may be grouped into modular subsystems can be represented by a network with a high degree of clustering. This is clearly a topological property. Through explaining robustness, modularity accounts also for "evolvability", the ability of a system to evolve due to its internal properties, since modularity is for this precise reason crucially involved in this property (Wagner and Altenberg 1996). Therefore here the topological answer to the proximate question (a) also impinges on our answer to evolutionary questions (b).

Another striking kind of robustness is the robustness of proteins in the face of recombination. Given that sexual reproduction breaks and recombines two different genotypes, and given that many traits are yielded by several genes with relations of epistasis, one should expect that recombination, by destroying epistasis, would make many gene complexes non-functional, thereby reducing the ability to form functional proteins, then dramatically reducing fitness and in the end making sexual reproduction counterselected. Given that this is obviously not the case (sexual reproduction is pervasive among many phyla), proteins should display some robustness regarding recombination. One of the explanations thereof is made in terms of topological properties. The main idea, here, is that not all topologies of epistasis relations are equally likely to be altered by random recombination.

If you consider proteins, amino acids in them are not related in whatever fashion. The frequency at which some sets of amino acids are functionally related to other sets of amino acids is what will provide the robustness of proteins regarding recombination. The fact is that given the topology of connections between subsequences of amino acids, the chances are very high that when you break a protein into two subsequences of amino acids and recombine, you still preserve the connection between some of them by reassembling them into identical subsequences (Fig. 6).

I tried to show that at many levels of biological hierarchies topological explanations stand side by side with mechanistic explanations when it comes to explanations of robustness. They have many forms. This raises the question why topological explanation is *in general* so well suited to explain robustness. The answer amounts to the very nature of topological properties. In the strict sense, if a continuous deformation leads from a to b they are topologically equivalent (homotopic). A topological property of a biological system S is defined by belonging to such equivalence class of homotopic shapes S′. Topological properties define features of shapes in space that are constant across deformations. Therefore, if a system has such properties this means that under this perspective it will stay constant across the processes realizing the deformations to which the topological properties are immune. There is thereby a natural

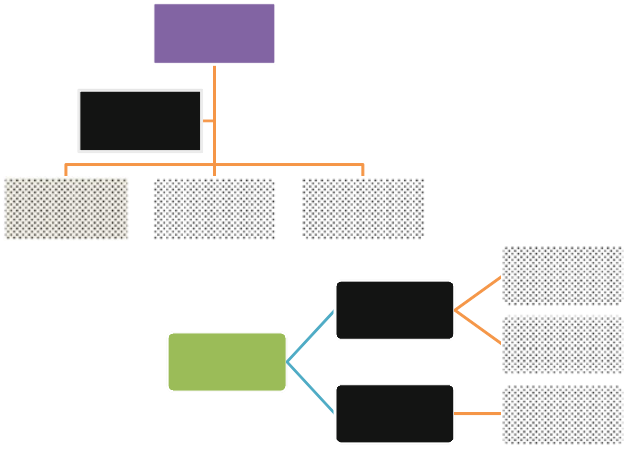

Fig. 6 Functional connections between subsequences of amino acids in proteins. The topology entails that randomly breaking a sequence (e.g. the *dotted squares*) still has some chances to keep functional connections (*dotted together*; *dark–dotted*.)

link between having some topological properties and displaying a sort of stability. In this sense, robustness, as constancy across changes, is naturally suited to be explained by means of topological explanations. Now, many of the explanations here mentioned do not pertain to topology stricto sensu but to graph theory (granted, their affinities have been emphasized above): then many properties there concern the equivalence between graphs in terms of paths between nodes; graphs can change whereas the nodes are still connected through equivalent pathways. Changing a graph while keeping the possibility of connections between the same nodes is exactly what underlies the property of robustness as constancy across change.

2.3 Neutral spaces

Finally a most interesting and encompassing recent approach to robustness is provided by the concept of *neutral spaces* developed by Wagner (2005a), after Stadler, Shuster, Fontana and others (Stadler and Stadler 2004; Schuster 2002; Cupal et al. 1999; Attolini and Stadler 2005; Van Nimwegen et al. 1999; Stadler et al. 2001; Fontana et al. 1999; Schuster et al. 1994). As compared to feedback mechanisms, this is clearly a topological explanation, ranging together with the scale-free networks explanations and the redundancy explanations of the robustness of the genetic code. The idea of neutral spaces comes from the construction of a genotype-phenotype mapping and the modeling of RNA folding, notions that I will explicate now.

First, comes the concept of "neutral network". RNA sequences are, like DNA sequences, linear entities. However, they owe their functionality to the spatial structure made by this sequence, because it determines the receptor sites, bonding sites, etc., in a word, all that is operational in the chemical working of RNA. So RNAs can be seen both as *genotypes*—the sequence—and as *phenotypes*—the structure, in

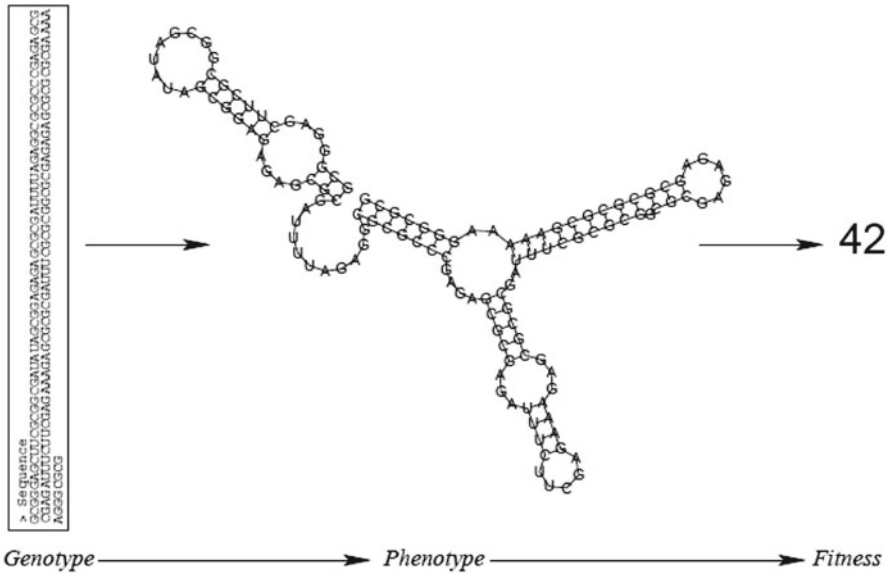

Fig. 7 Mapping: RNA sequence—RNA structure—fitness (after Stadler and Stadler 2004)

the sense that RNA interacts with other molecules and performs some functions of catalysis, bonding, etc. Therefore one can construct on this basis a kind of genotype (sequence)-phenotype (structure) mapping, and therein construct a fitness landscape by ascribing a fitness value to each structure (Fig. 7). The fitness does not need to be absolute; the point is that given all possible RNA structures, they can be ranked according to their differential performances in terms of some biochemical parameters (ability to synthesize proteins, chemical kinetics, catalyses, etc.). All structures can therefore be ascribed a fitness value in a scale. Of course, many structures can have the same fitnesses—the properties of many-to-one mappings will be crucial in the use of those maps.

Researchers noticed that often many different sequences, due to the laws of chemistry and the requirements of minimizing free energy, fold into the same spatial structure (i.e., with the same molecules at surface sites, hence functionally behaving). In fact, the space of structures is such that few structures are reached by several large sets of RNA sequences, whereas many structures are reached only by a few sequences. If you consider the sets of RNA sequences of a given length n realizing the same structure, these have interesting topological properties.

You can view the space of sequences as an n-dimensional space, each axis being likely to have 4 values (A, G, C, and U). Then you can define the distance between two sequences as the number of mutations required to get from one to another. Suppose a structure S onto which many different sequences map, noted as set K, and among them a sequence X. You can then define the set of all phenotypically equivalent sequences X_1 accessible from X by one mutation (the set of sequences mapping onto S, in K, and at a distance 1 from X). Then let us reiterate this operation on each of the X_1 and so on, and finally you have a set K_X of sequences which all give the same structure,

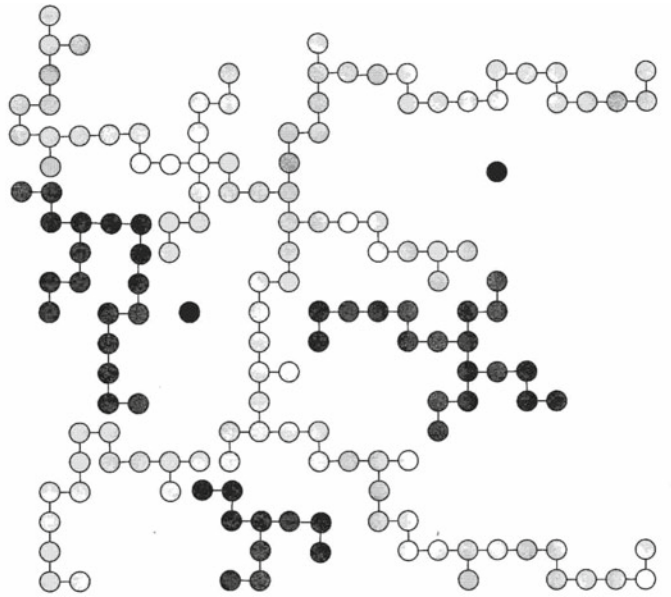

Fig. 8 *Neutral networks in RNA sequence-structure mappings* (after Wagner 2005a). Schematic illustration of three different classes of RNA or protein structures. The *rectangular area*, symbolizes sequence space. *Circles* correspond 10 individual sequences in this space, and *circles* with the same shading correspond to sequences folding into the same (secondary or tertiary) structure, The network of *circles shaded* in *light gray* corresponds to a highly frequent structure, a structure: realized by many sequences. All sequences folding into this highly frequent structure form a Connected, neural network. The three groups of *circles shaded in dark gray* correspond to sequences folding into the same moderately frequent structure. The sequences folding into this structure do not form a connected network, but instead form three disjoint sets of sequences. Finally, the *two black circles* correspond to a rare structure, a structure realized by only two sequences that occur at different points in sequence space. The image is misleading in that the actual sequence space is high-dimensional, not two dimensional as suggested by *box*

and which are such that each of them are connected by a series of mutations which never leaves set K. Producing the same structure, all the elements in K_X have the same fitness. Now, if X is somehow deep in K_X—meaning that at a distance 1 all its neighbors are in the same set K_X—then all mutations on X will be neutral. K_X defines a neutral network, in the sense that circulating by mutations across the network does not change the fitness (Fig. 8). Of course nothing proves that $K = K_X$, which means that nothing guarantees that you can circulate between all sequences in K, realizing the same structure by a series of mutations without getting out of K (K is not necessarily connex). What is important about robustness is that for X deeply in K_X, mutations will not change the structure and the fitness, so X will be robust to mutations. Mutational robustness is precisely the very interesting biological property of genotypes that keeps their fitness stable against many mutations of genes. Here, given that the topology of the set K_X as well as the position of X within it account by themselves for the fact that X is robust against many mutations, we can say that the structure of neutral networks provides a topological explanation of mutational robustness of RNA sequences. Moreover, two distant structures X_1 and X_2, rarely realized and with different fitnesses, can be linked through a series of mutations if between

them there is a neutral network K which has boundaries with each of them—so you can increase fitness of RNA sequences by drifting through neutral networks, a crucial evolutionary property of those.

Andreas Wagner (2005a) generalized this idea into the concept of "neutral space", supposed to account for the fact of robustness at many levels, not only RNA. It is easy to see this with DNA. Each genotype in the Genotype Space (G space) is correlated to a functional protein. However, several genes may lead to the same protein in Phenotype Space. The G space is a n-dimensional grid, and each cell on the grid is a nucleotide; distance between G and G' is measured by the number of changes you have to perform to go from G to G′. So, you can define, with a given X, the set of genotypes K_1 which are at distance 1 of X, and which lead to the same proteins. All the genotypes in this set can also define in the same way an extension K_2 of K_1, which leads to the same proteins, and so on until having a complete set K. In the end all the genes in K are such that there is always a pathway from one to another that involves only one step changes, and all of them relate to the same proteins, hence have the same fitness. K is therefore a neutral space, in the sense that switching from one genotype to another in K is neutral in fitness. In this sense if a genotype x is in K, and rather in the center (Fig. 8), x will be robust because most of the changes in x will lead to another genotype in this neutral space. Now, if two boundaries of the neutral set K are very different in fitness, one sees how the genotypes may evolve new phenotypes with higher fitness, only by drifting in the G space from one boundary of K to another.

So here, what explains the genetic robustness of some genotypes is the topological structure of the GP map, and the specific position of the genotype X in this structure. We can think of many possible mechanisms of gene-protein interaction underlying the neutral spaces but they will make no difference to the topological property of having neutral spaces, which explains robustness. Granted, the topology here represents underlying relations between genotypes—for example that chemical properties of nucleotides determine how substitutions actually occur between nucleotides of two genotypes. Yet those relations—for example this specific process—are not as such explanatory: in my example, one could have the same process leading from X to Y but no robustness because X and Y would not belong to the same neutral space. So the topological structure is what explains the robustness.

Neutral spaces are also likely to explain the robustness of the genetic code (by devising a space of alternative genetic codes (Wagner 2005a, p. 198)), as well as the robustness of gene expression (like in the case of Endo 16). This latter case is easily described by a space of DNA, which maps each regulatory DNA region onto a specific expression. Some genes like Endo 16 will be regulated by many alternative regulatory regions, and one can define the set K of regulatory DNA regions expressing the same gene. Very likely, K can be partitioned onto several neutral spaces, and a robust gene (regarding its expression) like Endo 16, will have a large neutral regulatory space associated.

Finally the models of neutral spaces are quite encompassing and pertain to many levels (as soon as one can define fitness, indeed), and those neutral spaces allow biologists to formulate issues and conclusions about the evolution of robustness. For instance, as is easily seen, the more robust a system, the larger are the neutral spaces in it. So if a system—a population of a species, let us say—has such a neutral space,

then the genotypes in them are likely to vary a lot and explore many regions, and may evolve towards higher fitness individuals. Without such a space, the variations would be very likely to decrease fitness, and then the species would be less likely to evolve (see Huynen et al. (1996) for a model). Robustness therefore has a key role in promoting the evolvability of species, and the topological properties of neutral spaces are clearly explanatory here. So the neutral spaces are providing proximate explanations of mutational robustness (question a), but also, contribute evolutionary explanations for the role robustness plays in evolution (question b2). They exemplify a kind of topology that may be explanatory of robustness at many biological levels (where relations analogous to GP and fitness can be defined). To this extent they also account for the pervasiveness of topological explanations of robustness across those levels.

The last section will draw consequences of the acknowledgement of topological explanations and their role in explaining robustness at all biological levels. The distinction between topological explanations and mechanistic explanations can shed light on a specific issue in evolutionary biology, about the contingency of evolution—the connection between kinds of explanations and kinds of realization being decisive in such regard.

3 Contingency, realizations and the scope of topological explanations

A longstanding metaphysical issue raised by evolutionary biology is whether evolution is mainly contingent, or not so contingent. Gould (1989) famously argued that if we "replay the tape of life", we would clearly not get the same kind of families and species—because even if natural selection is somehow deterministic (as compared to genetic drift, a stochastic process), large scale events in the history of life may involve mass extinctions, even if the species are well adapted by selection to their current environment. Many philosophers (e.g. Beatty 1995) subscribe to this contingency thesis. However other authors (e.g. Dennett 1995; Dawkins 1982) claim that evolution by natural selection hits on some "good tricks", like eyes, or limbs, or social life, which in any contexts are likely to be reached by evolving species. The pathway to them might be contingent, but themselves as endpoints are not. A good argument for this thesis is evolutionary convergences: eyes have been invented 22 times, multicellularity and sexual reproduction also several times. The frequency of such convergences indicates that it is not so contingent to meet those outcomes.

The debates often focus on issues close to adaptationism: if one is "adaptationist", namely one thinks that most of the most important traits in living beings arose through natural selection (Sober 1998), then one may also think that optimal solutions are likely to be reached in any possible world, no matter the initial conditions. The compared weight of drift, or in general contingent factors, and selection, often determines one's position with respect to the contingency issue.

The whole debates have a simple solution if one considers the grain of description. At a fine grain, the features of traits are mostly shaped by contingent factors: it is very unlikely for example that in another possible world we would witness creatures exactly like the ones we know. Too many contingencies in their phylogenies and even their recent evolutionary history (by drift) entered into the fine-grain shaping of those traits. Even the fact that life is carbon based, i.e. it is made of C, H, O and N atoms

which form DNA and RNA, the replicable molecules, is contingent—and the code linking those DNA molecules to amino-acids is also contingent in the sense that some others might have been possible.[23] However, at a very coarse grain, clearly one could expect some similarities in another version of evolution in another world. One could argue that if the conditions for selection are fulfilled, i.e. inheritance, variation and fitness (causal correlation between reproductive chances and heritable variable traits) (Lewontin 1970) then we would get: entities which consume other entities as resources; entities which evolve detectors in order to track environmental changes (because this gives a selective advantage against non-tracking entities in changing environments); entities which, amongst those ones, evolve detectors tracking variations of light (which is a good predictor of environmental changes); entities which finally could detect light modulations, and eventually something not so far from what we call eyes. Moreover the competitive exclusion principles in ecology (Gause), directly deriving from natural selection,[24] stating that no two species can live in the same niche, would imply that an increasing variety of environments should be colonized, so other worlds should witness forms of avian life, aquatic life etc.

In this sense, at a very coarse grain, there is something necessary in an evolutionary sequence. Experiments in Artificial Life, where researchers designed silicon-based life likely to undergo natural selection, attested that we have robust outcomes of natural selection—for instance Thomas Ray's Tierra experiment have shown that in many cases parasites, antiparasites, predators, etc., are regularly occurring. The selected effects which drive evolution, in other words the functions (Neander 1991) are somehow recurrent in the possible histories of life, but the details of the structures realizing those functions depend upon the kind of chemistry underlying forms of life (first of all carbon-based vs. silicon based life), so are plainly contingent.

The whole debate can then be phrased in terms of realization. The solution sketched here can be formulated in the following way: coarse grained features like being a detector, being a light detector, being a predator, or being a parasite, can be thought as *realized* properties. Organisms of given species are carrying *realizers* of those properties (e.g. eyes), or are *realizers* themselves (e.g. some species of fungi as parasites). So contingency pertains to the level of realizers, necessity to the level of (multi)realized properties.

However this is not the whole controversy. There is another reason for which one could argue for some necessity across evolution-hosting possible worlds. D'Arcy Thompson was one of the most influential thinkers—and the first—to highlight the role of non-biological laws in the shaping of all species. Often we call "Structuralists" (Crutchfield 1994; Amundson 2005; etc.) those biologists who think that such laws of structures are the principal laws. According to this position, natural selection fine-tunes organisms to their environments (for example, it determines the color of the wings of the peppered moths, if we consider the famous case of industrial melanism

[23] In this sense Crick proposed the theory that the current code is a frozen accident, in the sense that it's a contingent result on which all the evolutions of living beings built thereafter so it's kept as "frozen" (Crick 1968).

[24] Although its universal validity is contested by neutral ecologists—yet a limited validity here is enough for my argument.

analyzed by Kettlewell 1955); but structural laws determine the essential requirements constraining the make-up of organisms (e.g. they explain the wings themselves, and then explain the fact that moths, and generally insects, living in very diverse environments, have the same shape of wings). Those laws of structure define what Darwin already called "Unity of Type" (after Geoffroy Saint-Hilaire)—meaning that for example several insects are very different but have organs which in fact are somehow the same, like wings of bats and fins of fish.[25]

Interestingly, many of those laws of unity of type are (indeed) topological. They pertain to morphology, but also to the lowest level of biological organization—for example protein folding, as explained by molecular biology (Dokholyan et al. 2002), where especially small-worlds networks entail important properties (Vendruscolo et al. 2002). Sometimes commonalities of form that have prima facie seemed to result from natural selection, in fact came out from pure topological properties. For instance, the fact that almost all cell metabolic networks are modular seems to be due to natural selection, which is known to favor modularity in general (because of its robustness, as we have seen). A default state of a possible metabolic network picked at random would be non-modularity, and then natural selection will favor and fix modular networks. However, Sole and Valverde (2008) have shown that properties of graphs in general, applied to the cell networks, makes it very likely that when you randomly build a metabolic network it will be modular, so that here modularity comes for free and does not need natural selection. This fact could explain why we witness so many modular cell networks. Now, evolution in another possible world would be subject to the same topological graph regularities and therefore also display the same kind of modularity patterns. Because they are topological, hence not wedded to a specific kind of material constitution, the properties explaining modularity in cell network can hold whatever has been contingently evolved to form the chemistry of life and the substrates of replication. The argument here, in favor of the necessity thesis, would be: many biological organizations can evolve, in some other worlds, anyway they will display only a few topological features, even if the details of the way they realize those features may vary. So here, realizers are contingent, once again; but the realized properties are necessary, explaining commonalities in this world as well as general features of life, whatever the chemical substrate of life is made of elsewhere. And those realized properties are indeed topological properties—unlike what they were in the adaptationist version of the non-contingency thesis, i.e. functional properties.

So we have two versions of the necessity thesis in evolutionary biology, both of them emphasizing the universality of realized properties, against the contingency of the realizers that hang on many historical events in actual evolution (Table 2). But in the adaptationist version, the realized properties are functional ones; so basically they are realized by mechanisms. Shapiro (2000) clearly related realizations to explanations and functions: pump mechanisms and screw-driver mechanisms in corkscrews together distinctly allow for the same property of cork extracting, because these different mechanisms explain why they fulfill their cork-screw function. This is the M-realization proper to such kind of explanation. The camera eye and the compound

[25] This amounts to the issue of what biologists call *homologies*, i.e. similarities that are not caused by adaptation to identical selective pressures, but by common descent (see Amundson 2005).

Table 2 Summary of the contingency debates

Necessity thesis 1	Necessity thesis 2
Structuralist version	*Adaptationist version*
Realizers are contingent across possible worlds	*Realizers* are contingent across possible worlds
Realized features and entities are necessary; they are consequences of topological properties	*Realized* features and entities are necessary; they are consequences of selection, hitting regularly on the same «good tricks» (i.e., constant causes of selective advantages)

eye of insects, for example, are two mechanisms for realizing the function "seeing". Explaining why a type of eye allows its bearer to see would require a mechanistic explanation.

In the structuralist version of the necessity thesis, the realized properties are variously realized by material structures. And I emphasized that their universality would often require topological explanations. Yet we have seen that topological explanations and mechanistic explanations embody different kinds of realization, A-realizations and M-realizations *sensu* Gillett. Therefore it appears that both necessity theses in evolutionary biology are emphasizing the realized properties rather than their realizers, in evolution, but that their distinction pertains to a distinction internal to realization in general, the one between M- and A-realization, correlated to the difference between types of explanations.

4 Conclusions

This paper intended to make sense of a difference between two kinds of explanation: mechanistic ones, absolutely pervasive in natural sciences and intending to display causal processes responsible for an explanandum—and topological ones, which do not consider mechanisms at stake. The latter kind is more and more frequent because of both the science of networks, which emerged in the wake of the general focus on complexity—and the "neutrality" theme in ecology and biology. An important part of the explanations of robustness of many kinds (resilience, stability, etc.) in many systems relies on topological considerations. I argued that paying attention to the structure of those explanations may both cast new light on the idea of multiple realizability as it is met in the special sciences, and contribute to important debates in the philosophy of biology (e.g. contingency in evolution).

Acknowledgements The author warmly thanks Isabelle Drouet, Paul Humphreys, Carl Gillett, Michael Strevens and Jackie Sullivan, for their precious and precise critiques and comments, very helpful in rewriting the manuscript.

References

Albert, R., & Barabasi, A.-L. (2002). Statistical mechanics of complex networks. *Review Modern Physics, 74*, 47–97.

Amundson, R. (2005). *The changing role of the embryo in evolutionary thought*. Cambridge: Cambridge University Press.

Arcos, H., & Pereira, J. (2007). Spacetime: Arena or reality?. In V. Petkov (Ed.), *Relativity and the dimensionality of the world. Fundamental theories of physics* (Vol. 153, pp. 217–227). Berlin: Springer.

Attolini, C. S. O., & Stadler, P. F. (2005). Neutral Networks of interacting RNA secondary structures. *Advances in Complex Systems, 8*(2–3), 275–283.

Beatty, J. (1995). The evolutionary contingency thesis. In G. Wolters & J. Lennox (Eds.), *Concepts, theories, and rationality in the biological sciences* (pp. 45–81). Pittsburgh: University of Pittsburgh Press.

Berlow, E. L. (1999). Strong effects of weak interactions in ecological communities. *Nature, 398*, 330–334.

Bernstein, H., Byerly, H. C., Hopf, F. A., & Michod, R. E. (1985). Genetic damage, mutation, and the evolution of sex. *Science, 229*(4719), 1277–1281.

Craver, C., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy, 22*, 547–563.

Craver, C. (forth.). Functions and mechanisms in contemporary neuroscience. In P. Huneman, (Ed.), *Functions: Selection and mechanisms*. Berlin: Synthese Library, Springer.

Crick, F. H. C. (1968). The origin of the genetic code. *Journal Molecular Biology, 38*, 367–379.

Crutchfield, J. (1994). Is anything ever new? Integrative themes. In G. Cowan, D. Pines, & D. Melzner (Eds.), *Santa Fe Institute studies in the sciences of complexity XIX*. Reading, MA: Addison-Wesley.

Cupal, J., Stadler, P., & Schuster, P. (1999). Topology in phenotype space. In J. Giegerich (Ed.), *Computer science in biology* (pp. 9–15). Berlin: Springer.

Dawkins, R. (1982). *The extended phenotype*. Oxford: Oxford University Press.

de Visser, J. A. G. M., Hermisson, J., Wagner, G. P., Ancel Meyers, L., Bagheri-Chaichian, H., Blanchard, J. L., & Chao, L. (2003). Evolution and detection of genetic robustness. *Evolution, 57*, 1959–1972.

Dennett, D. (1995). *Darwin's dangerous idea*. New York: Simons & Shuster.

Dokholyan, N., Li, L., Ding, F., & Shakhnovich, E. I. (2002). Topological determinants of protein folding. *PNAS, 99*(13), 8637–8641.

Dunne, J. (2006). The network structure of food webs. In M. Pascual & J. Dunne (Eds.), *Ecological networks: Linking structure to dynamics in food webs*. Oxford: Oxford University Press.

Dunne, J. A., Williams, R. J., & Martinez, N. D. (2002a). Network structure and biodiversity loss in food webs: Robustness increases with connectance. *Ecology Letters, 5*, 558–567.

Dunne, J. E., Williams, R. J., & Martinez, N. D. (2002b). Food web structure and network theory: The role of connectance and size. *PNAS, 99*, 12917–12922.

Fontana, W., Stadler, P. F., Bornberg-Bauer, E., Griesmacher, T., Hofacker, I. L., Tacker M., et al. (1999) RNA folding and combinatory landscapes. *Physics Review E, 47*, 2083–2099.

Friedman, M. (1974). Explanation and scientific understanding. *Journal of Philosophy, 71*, 5–19.

Gillett (2007). Understanding the new reductionism: The metaphysics of science and compositional reduction. *The Journal of Philosophy*, 193–216.

Gilllett (2010). Moving beyond the subset model of realization: The problem of qualitative distinctness in the metaphysics of science. *Synthese*. doi:10.1007/s11229-010-9840-1.

Gould, S. J. (1989). *Wonderful life. The Burgess shale and the nature of history*. New-York: Norton.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology, 78*(6), 1360–1380.

Gross, J. L., & Tucker, T. W. (1987). *Topological graph theory*. Reading, MA: Wiley Interscience.

Holling, G. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics, 4*, 1–23.

Hubbell, S. P. (2001). *The unified neutral theory of biodiversity and biogeography*. Princeton: Princeton University Press.

Huneman P. (2010). The conceptual foundations of ecological engineering. Stability, individuation, ethics. In S. Barot (Ed.), *Ecological engineering, from concepts to applications*. Physics Procedia. Amsterdam: Elsevier.

Huynen, M., Stadler, P. F., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *PNAS, 93*, 397–401.

Kettlewell, H. B. D. (1955). Selection experiments on industrial melanism in the Lepidoptera. *Heredity, 9*, 323–342.

Kitano, H. (2004). Biological robustness. *Nature Review Genetics, 5*, 826–837.

Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. Salmon (Eds.), *Scientific explanation* (pp. 410–505). Minneapolis: University of Minnesota Press.

Lehmann, L., & Keller, L. (2006). The evolution of cooperation and altruism—A general framework and a classification of models. *Journal of Evolutionary Biology, 19*, 1365–1376.

Levin, R., May, R., & Sugihara, C. (2008). Complex systems: Ecology for bankers. *Nature, 451*, 893–895.

Lewontin, R. C. (1970). Units of selection. *Annual Review of Ecology and Systematics, 1*, 1–18.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science, 67*(1), 1–25.

Maser, G. L., Guichard, F., & McCann, K. S. (2007). Weak trophic interactions and the balance of enriched metacommunities. *Journal of Theoretical Biology, 247*, 337–345.

May, R. M. (1974). *Stability and complexity in model ecosystems*. Princeton: Princeton University Press.

McCann, K., Hastings, J., & Huxel, R. (1998). Weak trophic interactions and the balance of nature. *Nature, 395*, 794–798.

Montoya, J., Pimm, S., & Solé, R. (2006). Ecological networks and their fragility. *Nature, 442*, 259–267.

Montoya, J. M., & Solé, R. V. (2002). Small world patterns in food webs. *Journal of Theoretical Biology, 214*, 405–412.

Morange, M. (2000). *A history of molecular biology*. New Haven: Harvard University Press.

Neander, K. (1991). Functions as selected effects: The conceptual analysts defense. *Philosophy of Science, 58*, 168–184.

Nowak, M. (2005). Five rules to the evolution of cooperation. *Science, 314*(5805), 1560–1563.

Park, J., & Newman, M. (2003). The origin of degree correlations in the Internet and other networks. *Physical Review E, 68*(2), 026112.

Pimm, S. L. (1984). The complexity and stability of ecosystems. *Nature, 307*, 321–326.

Pimm, S. L. (2001). *The balance of nature? Ecological issues in the conservation of species and communities*. Chicago: University of Chicago Press.

Randrup, T., & Rogen, T. (1997) How to twist a knot. *Archiv der Mathematik, 68*(3), 252–264.

Revilla, I., Domingo, R., & Davidson, E. (2003). Developmental gene network analysis. *International Journal of Developmental Biology, 47*, 695–703.

Romano, L. A., & Gray, G. A. (2003). Conservation of endo 16 expression in sea urchins despite evolutionary divergence in both cis and trans-acting components of transcriptional regulation. *Development, 130*(17), 4187–4199.

Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.

Santos, F. C., & Pacheco, J. M. (2005). Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letter, 95*(098104), 1–5.

Santos, F. C., & Pacheco, J. M. (2006). A new route to the evolution of cooperation. *Journal of Evolutionary Biology, 19*(3), 726–733.

Schuster, P. (2002). A testable genotype-phenotype map: Modeling evolution of RNA molecules. In M. Lässig & A. Valleriani (Eds.), *Lecture notes in physics, 585* (pp. 55–81). Berlin: Springer.

Schuster, P., Fontana, W., Stadler, P. F., & Hofacker, I. (1994). From sequences to shapes and back: A case study in RNA secondary structures. *Proceedings of the Royal Society of London Series B, 255*, 279–284.

Shapiro, L. (2000). Multiple realizations. *Journal of Philosophy, 97*, 635–654.

Sober, E. (1998). Six sayings about adaptationism. In D. Hull & M. Ruse (Eds.), *Oxford readings in philosophy of biology* (pp. 72–86). Oxford: Oxford University Press.

Solé, R. V., & Montoya, J. M. (2001). Complexity and fragility in ecological networks. *Proceedings of the Royal Society of London Series B, 268*, 2039–2045.

Sole, R., & Valverde, S. (2008). Spontaneous emergence of modularity in cellular networks. *Journal of the Royal Society Interface, 5*, 129–133.

Solé, R. V., Ferrer, R., Montoya, J. M., & Valverde, S. (2002). Selection, tinkering and emergence in complex networks. *Complexity, 8*, 20–33.

Stadler, P., Stadler, P., Wagner, G., & Fontana, W. (2001). The topology of the possible: formal spaces underlying patterns of evolutionary change. *Journal Theoretical Biology, 213*(2), 241–274.
Stadler, B., & Stadler, P. (2004). The topology of evolutionary biology. In C. Ciobanu (Ed.), *Modeling in molecular biology. Natural computing series* (pp. 267–286). Berlin: Springer.
Strogatz, S. (2001). Exploring complex networks. *Nature, 410*, 268–276.
Tilman, D. (1996). Biodiversity: Population versus ecosystem stability. *Ecology, 77*(2):350–363.
Tylianakis, J. (2008). Understanding the web of life: The birds, the bees, and sex with aliens. *PLoS Biology, 6*(2), 224–228.
Ulanowicz, R. (2002). The balance between adaptability and adaptation. *Biosystems, 64*(1–3), 13–22.
Van Nimwegen, E., Crutchfield, J., & Huynen, M. (1999). Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences of the United States of America, 96*(17), 9716–9720.
Vendruscolo, M., Dokholyan, N., Paci, E., & Karplus, N. (2002). Small-world view of the amino acids that play a key role in protein folding. *Physical Review E, 65*(061910), 1–5.
Wagner, G., & Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution, 50*(3), 967–976.
Wagner, A. (2005a). *Robustness and evolvability in living systems*. Princeton: Princeton University Press.
Wagner, A. (2005b). Distributed robustness versus redundancy as causes of mutational robustness. *Bioessays, 27*, 176–188.
Watts, D. (2003). *Six degrees of separation*. New York: Norton.
Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics in "small-world" networks. *Nature, 393*, 440–442.
West, S. A., Griffin, A. S., & Gardner, A. (2007). Social semantics: Altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology, 20*, 415–432.
White, S., & Kiester, D. (2008). Topology matters: Network topology affects outcomes from community ecology neutral models. *Computers, Environment and Urban Systems, 32*(2), 165–171.
Wilke, C. O., Wang, J., Ofria, C., Lenski, R. E., & Adami, C. (2001). Evolution of digital organisms at high mutation rate leads to survival of the flattest. *Nature, 412*, 331–333.
Yodzis, P. (1989). *Introduction to theoretical ecology*. New York: Harper & Row.