

Network Analyses in Systems Biology: New Strategies for Dealing with Biological Complexity

Sara Green,ⁱ Maria Şerban,ⁱⁱ Raphael Scholl,ⁱⁱⁱ Nicholaos Jones,^{iv} Ingo Brigandt,^v
and William Bechtel^{vi,1}

Abstract: The increasing application of network models to interpret biological systems raises a number of important methodological and epistemological questions. What novel insights can network analysis provide in biology? Are network approaches an extension of or in conflict with mechanistic research strategies? When and how can network and mechanistic approaches interact in productive ways? In this paper we address these questions by focusing on how biological networks are represented and analyzed in a diverse class of case studies. Our examples span from the investigation of organizational properties of biological networks using tools from graph theory to the application of dynamical systems theory to understand the behavior of complex biological systems. We show how network approaches support and extend traditional mechanistic strategies but also offer novel strategies for dealing with biological complexity.

1 Revisiting research strategies in biology

Life scientists draw on a variety of research strategies to deal with biological complexity. The success of molecular biology stems from the development of powerful experimental strategies to manipulate molecular interactions underpinning biological functions. Drawing on a variety of case studies, in *Discovering Complexity* Bechtel and Richardson (1993) described how biologists use the strategies of *decomposition* and *localization* to reduce the problem space for biological analysis. These strategies imply the provisional assumption that biological systems can be subdivided into localizable operations of interrelated parts, organized in modules, from which the workings of larger (sub)systems can be recomposed. The result of a successful use of these discovery strategies, a mechanistic explanation, cites

¹ The authors contributed equally to this work.

ⁱ⁾ University of Copenhagen, Denmark, Email: saraehrenreichgreen@gmail.com.

ⁱⁱ⁾ University of Copenhagen, Denmark.

ⁱⁱⁱ⁾ University of Geneva, Switzerland.

^{iv)} University of Alabama in Huntsville, USA.

^{v)} University of Alberta, Canada.

^{vi)} University of California, San Diego, USA.

how interacting and hierarchically organized parts causally produce the phenomenon (Bechtel and Abrahamsen, 2005; Bechtel, 2011; Machamer, Darden, and Craver, 2000).

Mechanistic research strategies are very productive if the system is organized in such a way that functions can be localized to parts of a mechanism, and if the functioning of such a part is due to its internal organization and largely unaffected by its context, so that parts can be investigated in isolation (via decomposition) and their joint operation is relatively easy to understand (via recomposition). However, analyses of complex and highly integrated systems seem to defy these general assumptions, raising challenges for traditional mechanistic strategies. Indeed, Bechtel and Richardson have viewed decomposition and localization as heuristics because they fail in some cases. In the new introduction to the 2010 edition of *Discovering Complexity*, they point to recent biological research on complex networks (e.g., with extensive feedback loops) as cases where decomposability may not hold—while leaving the philosophical investigation of how such network research works as a project for the future.

Additional strategies of mechanistic research have been philosophically articulated by Craver and Darden (2013). They cover not only experimental strategies for investigating mechanisms and validating mechanism schemas, but also several conceptual strategies that guide the formulation of hypothesized mechanism schemas, such as *modular subassembly*, *forward chaining*, and *backward chaining*. The recent accumulation of big quantitative datasets has motivated some life scientists to utilize network modeling as a research strategy for the organization and interpretation of data. Unlike mechanism diagrams, network models typically do not contain details about the molecular properties of the components. By representing interactions among a vast number of molecular species they enable an analysis of the organizational structure of larger systems, sometimes involving automated pattern-detection. Geared toward the analysis of the complex structure and dynamic operation of large-scale networks, these representational, computational, and analytical strategies appear to go beyond the conceptual strategies that have been listed by Craver and Darden.

This calls for philosophers to revisit the question about prominent research strategies for dealing with biological complexity. How do these new strategies relate to the mechanistic research strategies described by Bechtel and Richardson (1993) and Craver and Darden (2013)? Are network approaches an extension of or at least complementary to mechanistic strategies? And what, more generally, can network analysis afford to researchers in the life sciences? In this paper we address these questions by focusing on how various types of biological networks are represented and analyzed. We investigate a variety of network analyses used in systems biology with the aim of identifying some of the distinctive aspects of this new research approach. Focusing on a set of *representational* and *analytical* strategies relying on network modeling, we also show that there are different ways of using network approaches that have different philosophical implications.

For this purpose, we draw on a set of case studies from systems biology. Systems biology differs from most other fields of biology by the lack of an explicit focus on a specific domain of phenomena (in contrast to cell biology, molecular biology, genetics, and developmental biology). Instead, the field distinguishes itself by new research strategies such as analyses of large quantitative datasets on gene regulation and protein interactions through computational simulations and network analysis. Although network analysis precedes systems biology historically and has been widely applied in a variety of other fields such as neuroscience, anthropology, and sociology (see Newman, 2010), the analysis of data-rich network models in systems biology presents a new and interesting direction in need of philosophical analysis. The cases we examine all target the organization and behavior of biological networks but also illustrate the diversity of representational and analytical strategies used in systems biology research.

We begin with some examples in which graph theoretic tools have been successfully applied to investigate the organization of biological networks (Sections 2.1–2.4). These approaches have revealed important principles governing both the structural organization of biological networks and the timing of activities involved in these networks. We then examine cases where network approaches help scientists to explore and understand the dynamic features of biological system activity (Sections 3.1–3.2). Over the course of these different case studies, we gradually move from more local analyses of specific components of complex systems to ones that address their global operation (cf. Huang, 2004). Roughly speaking, the localist approaches are more in continuity with traditional mechanist research strategies than the globalist approaches. However, given the diversity of the network research discussed and the absence of a unique notion of ‘mechanistic’ research (e.g., there are mechanistic modes of explanation as well as mechanistic tools of experimental discovery), for each case of network research we will lay out in what aspects it is mechanistic, complements features of mechanistic research, or departs from classical mechanistic strategies.

2 Graph-theoretic analyses of networks

Large data-collection procedures have enabled biologists not only to identify many different genes, proteins, and metabolites, but also to explore how they interact. To filter, organize, and interpret this data, some biologists employ network modeling strategies from graph theory, a field of mathematics which uses combinatorial and geometrical principles to analyze relational structures (i.e., pairwise relationships between objects). A graph is a collection of objects called nodes or vertices, and a collection of edges or links, which represent connections between these objects. One distinctive feature of this type of approach is that systems biologists often start with large numbers of components, some interacting with a large number of other components.

To construct a graph representation of a biological network, researchers often turn to databases such as RegulonDB, BioGraph, Bioinformatic Harvester, or UniProt, from which large datasets can be downloaded. In network models, entities such as genes and proteins are represented as nodes and the relations between them as edges. Edges may be directed or undirected depending on whether, for example, one has causal or merely correlational information about their interactions. The goal of applying the tools of graph theory to the study of these biological networks is to identify organizational patterns that might be useful for predicting and explaining interesting features of biological organisms. We begin with efforts to identify small-scale patterns in biological networks, and then consider strategies for analyzing the large-scale organization of whole networks.

2.1 Motifs within networks

One strategy for extracting relevant patterns from large graphs is to focus on *motifs*—patterns of connectivity between a small number of nodes (see Figure 1 for an example). Uri Alon, who pioneered this type of analysis, characterized these motifs as basic computational elements (Shen-Orr, Milo, Mangan, and Alon, 2002; Alon, 2007). As the word ‘computational element’ suggests, such a node cluster is deemed to be a module capable of performing a specific function. Modules are connected to others to make the whole functional system. Researchers who seek out computational elements often characterize themselves as engaging in *reverse engineering*. This practice typically involves taking an already made product and exploring its behavior and component functions in systematic ways at many levels of description in order to build (synthesize) a similar product. By using these methods, biologists try to identify the functional organization of the system components, which gives rise to their global behaviors. This approach is motivated by the idea that biological functions, like engineered systems, may be characterized by basic organizational or ‘design principles’ (Green, Levy, and Bechtel, 2014), as we will see below.

Alon developed the network motif approach while exploring the organizational structures of gene regulatory networks, based on data on transcriptional regulation in bacteria. Reverse engineering in this context was highly exploratory in the sense that little was known about the function and structure of the network. The group compiled data on transcriptional regulation from RegulonDB and additional sources and represented the transcriptional interactions as a directed graph. He and his collaborators noticed that particular subgraphs involving three or four nodes occurred with a surprisingly high frequency. To elaborate this finding they developed and applied algorithms comparing the frequency of all the subgraphs of a particular type in a larger graph to randomly constructed graphs with the same degree of connectivity. The idea is thus that automatic pattern-detection (here via frequency of subgraphs) can be used to identify organizational patterns of possible functional significance. Hypothesizing that those subgraphs that appeared especially often were likely to display distinct functions (like circuits in

engineered systems), Alon and colleagues represented them as in Figure 1 and then used computational modeling to analyze their functionality.




Figure 1. Example of a motif. A coherent feedforward loop of the sort investigated by Alon.

Figure 1 shows one of the three-node subgraphs Alon found to be especially frequent, a motif which he labeled a *coherent feedforward loop*. In it, the node X (which receives the external signal from S) is connected to node Z both directly and indirectly through node Y. By analyzing the motif from an engineering perspective, Alon showed that when Z requires both of its inputs to respond, the network constitutes a computational element that will function as a *persistence detector*—Z will be active only with a sustained stimulation from S. This design is particularly useful in contexts in which chemical signals may be generated spuriously and where engaging in an activity such as synthesizing a protein unnecessarily can be particularly wasteful. Using simple Boolean modeling, Alon demonstrated that the coherent feedforward loop would perform this elementary computation as long as it takes time for a node to become active (see Alon, 2007). Since each node ceases to be active once input ceases and the activation through the indirect pathway takes longer than through the direct pathway, Z will be activated only when X remains active long enough for Y to become active so that both X and Y will be simultaneously sending input to Z. He also showed that changing Z to operate as an OR-gate (Z is active when at least one of its two potential inputs obtains) will yield a different behavior. Now Z continues to function even in the face of brief interruptions in the activity of X. This *sign-sensitive accelerator* design can be important in noisy environments in which it takes time to complete a process.

Following the initial analyses, software for automated motif-detection was developed. This enabled a more extensive search for network motifs in different biological systems and at different levels of biological organization. Some critics have questioned Alon's methods of evaluating whether the frequency of subgraphs is beyond what would be expected by chance (cf. Beber, Fretter, Jain, Sonnenschein, Muller-Hannemann, and Hutt, 2012). Searching for overabundant circuit patterns in networks can, however, be useful as a heuristic regardless of the soundness of the null hypothesis used in the comparison, provided that this the network analysis is combined with other strategies that account for the details that network analyses neglect or distort.² Moreover, the research sparked a debate on whether functions are determined by the organizational structure of the regulatory circuits and can be generalized from an abstract analysis (Green, 2015; Levy and Bechtel, 2014). A requirement for such generalizations is that the mathematically computed motif functions can be isolated from the network as a whole in living cells. Initial experiments on extensively studied regulatory systems in *E. coli* provided an affirmative answer (Mangan, Zaslaver, and Alon, 2003). But some have cautioned about the situation for more complex organisms where the same motifs are also present (Isalan, Lemerle, Michalodimitrakis, Horn, Beltrao, Raineri, Garriga-Canut, and Serrano, 2008). In response, more recent research has focused on specifying the conditions (parameter spaces) for which the identified motif functions hold (Tyson and Novak, 2010).

The network motif approach shares with mechanistic strategies the heuristic assumption of near-decomposability of biological systems, and network approaches and mechanistic research strategies are intertwined in the experimental investigation of the computed functions. Individual motifs that have been identified function similarly to what Machamer, Darden, and Craver (2000; see also Craver and Darden 2013) call *mechanism schemas*. The computational analysis of individual motifs is carried out at a higher level of abstraction than typical research in molecular biology, insofar as motifs abstract away from a wide range of details about the parts and operations of specific molecular mechanisms (cf. Levy and Bechtel, 2013). But the very search for motifs also provides a novel heuristic strategy for detecting regulatory units—the computational screening of vast networks—thus extending mechanistic research. This goes beyond the conceptual as well as experimental strategies of mechanism discovery that Craver and Darden (2013) have compiled: Alon's approach allows for motif-detection from a representation of the large-scale structure of

²One limitation of this kind of network analysis is that it represents regulatory interactions as pairwise relations between discrete and static objects (nodes in the network) and assumes that these are uniformly distributed throughout the cell. However, in many regulatory systems (e.g. in development) the spatial distribution of molecules is crucial and needs to be accounted for. Mechanistic research providing such details can therefore provide a useful corrective to network analysis. (We would like to thank an anonymous reviewer for stressing this point.)

regulatory connections, *without* background information on specific functions of the subsystems. The detection of frequent organizational patterns allows for a data-intensive exploratory analysis, revealing motifs which are amenable to further investigation.

We now turn to network analyses that reveal organizational patterns in networks at a scale above that of small subgraphs. As we will see, within these network approaches, researchers often differentiate modules that correspond to particular mechanisms, but, unlike in the case of more traditional mechanistic approaches, the analytical emphasis falls on how these modules are situated in larger networks.

2.2 Tools for Analyzing Large-Scale Networks

Graph theorists have developed various measures to characterize and analyze the patterns of interactions among the many components of graphs and to make inferences about how they function together. These take a somewhat different form if the edges in the graph are directed or undirected; we will consider only undirected graphs and focus on three quantitative measures that have played an especially prominent role in mathematical analyses of networks.

As a global property of a whole network, the mean shortest path length (henceforth, MSPL) measures the average, for all pairs of nodes, of the minimum number of edges that must be traversed to get from one node to the other.³ Shortest paths are of interest in understanding how quickly or slowly activity at one node may affect activity at another; MSPL is then a measure of how quickly signals can be passed between different components of a network. Second, the clustering coefficient C is the degree to which a given node's neighbors are connected among each other. When C for a given node is high, there will be high connectivity in the neighborhood around that node; these connected nodes are often viewed as modules. Below we will illustrate why the average clustering coefficient across all nodes (as a global network measure) is of analytical value. Finally, the degree k of an individual node is the number of other nodes to which it is connected; and the network-wide degree distribution $p(k)$ is the probability that a given node will have degree k .

Most mathematical analyses of graphs in the mid-20th century focused on either randomly connected networks, which are characterized by a short MSPL and a low average clustering coefficient, or lattice structures, which have a long MSPL and high average clustering. Watts and Strogatz (1998) attracted much attention when they theoretically introduced the class of *small-world* networks that exhibit both short MSPL and high average clustering. They

³ In this context, 'path length' does not measure physical distance (only number of edges) and 'organization' does not track spatial organization of the target system (but rather its functional organization).

argued that these characteristics are important as they allow for a high degree of coordination throughout the network yet also specialization of function in clustered units. This illustrates how some representational notions and analytical strategies of network research are not so much about the specific functioning of a network segment—as in the analysis of individual motifs discussed above—but about a network’s overall organization. Network researchers focusing on gene-regulatory, protein interaction, and metabolic networks in a variety of species found that these biological networks indeed exhibit small-world structure.

However, it turns out that, as a matter of statistical probability, nearly every network with a large enough number of nodes and connections will be a small-world network (Newman, 2010). This motivated the search for a more precise description of biological networks using the tools (and measures) of graph theory. Historically, graph theorists assumed that the degree distribution $p(k)$ would be normal (e.g., Gaussian), but Barabási and Albert (1999) found that in many real world networks, node degree is distributed according to a power law of the form $p(k)=c \cdot k^{-a}$ (where the parameter a is often between 2 and 3). This exponentially declining distribution means that there is great variation in the number of connections for individual nodes. Most nodes have a very low degree (e.g., only 1 or 2 edges to other nodes) but a small number of nodes—called *hubs*—have a many connections. Since the tail of a power-law distribution extends over many orders of magnitude and the node degree across the system cannot be characterized using a unique scale, these networks have been called *scale-free networks*.

An analytical implication of a scale-free network organization is that it tends to exhibit *robustness*, i.e., its functioning is usually not seriously affected by perturbations. For while the elimination of a highly connected hub may detrimentally impact functioning, most nodes in a scale-free network are not hubs and have only small impact on the overall network. Given that robustness is a biologically important feature for a system to have, this highlights the usefulness also of global measures of network structure. In the specific philosophical context of explanation, Huneman (2010) has argued that the explanation of a network’s robustness in terms of its scale-free structure is an instance of what he calls topological explanation, which he deems to be non-mechanistic, on the grounds that an account purely in terms of structural organization does not explain by tracking activities during a mechanism’s operation (see also Brigandt, Green, and O’Malley, in press; Jones, 2014). A discussion of the scope of specific accounts of explanation is beyond the scope of this paper, but we mention these aspects to highlight how representing abstract system organization and topology is a strength of network approaches that extends the analytical toolkit in biological research.

As concrete illustrations of the potential of using graph theoretical methods and measures in the study of biological organisms, we turn in the next two subsections to discussing two network analyses; one revealing hierarchical modular organization within metabolic networks, the other temporal dynamics of protein-protein interactions.

2.3 Hierarchical modular networks

Traditional mechanistic approaches identify modules through research strategies that start from a specific phenomenon (e.g., a concrete function), identify a system taken to be responsible for it, and decompose that system into functionally relevant components (Bechtel and Richardson, 1993). In contrast, graph-theoretic analyses that abstractly represent the organization of large numbers of components and their interactions identify modules as classes of nodes that are highly interconnected (compared to the degree of connectedness across modules). Hence, graph-theoretic approaches offer a new perspective on modules, highlighting how they are situated in the broader organization of the network. One type of analytical strategy that has gained considerable popularity focuses on the hierarchical organization of many modules in larger networks.

The characterization of hierarchical modular networks was introduced by Ravasz, Somera, Mongru, Oltvai, and Barabási (2002) to resolve a tension they had identified between two network measures that were concomitantly exhibited by a wide class of biological networks. In analyzing metabolic networks in *E. coli*, Wagner and Fell (2001) had shown that they both (1) have high average clustering—standardly associated with the presence of fairly isolated modules—and (2) are scale-free, having some hub nodes with many connections across the network. So, is the metabolism of *E. coli* organized as a collection of specialized modules or as a highly integrated network (Ravasz et al., 2002)?

The combination of modules and nodes with high degree (which Ravasz et al. found in 43 other species) seems contradictory—modules are relatively isolated, whereas nodes with high degree interconnect the whole system. Ravasz and colleagues proposed to resolve this tension by constructing a model of a network that would exhibit both properties. They began with a module of four completely interconnected nodes arranged as in Figure 2A. They then created three replicas, connecting the peripheral nodes of each replica to the central node of the original module (Figure 2B). This procedure was then iterated (the next iteration is shown in Figure 2C). The network that resulted after several iterations exhibits high average clustering ($C \cong 0.6$), due to it being built by replicating highly-clustered units, and is scale-free (the parameter α of the power law is 2.26), as a consequence of connecting the peripheral units in a replica only to a common unit in the original.




Figure 2. The process Ravasz et al. employ to construct hierarchical modular networks that have high clustering and are scale-free. The module shown in A is replicated three times, where in B the peripheral nodes of each replica are connected to the central node of the original module. This process is iterated one more time in C, but can be iterated indefinitely. From Ravasz et al. (2002).

Ravasz and colleagues observe that nodes that appear in the center of the network model have the highest degree (i.e., are hubs) and the lowest clustering coefficient, while the nodes at the periphery of each of the super-modules have low degree and the largest clustering coefficient. In between are nodes with moderate degree and moderate clustering; these serve to link the smallest modules into larger modules. This creates a hierarchy in which the nodes generated at each replication step are connected not only to their own central node but also to the central node of the entire structure. Ravasz et al. refer to this network structure as a *hierarchical modular network* and propose that networks organized in this way might account for the combination of high clustering and scale-freeness in the metabolic networks of the 43 species they examined.

In particular, Ravasz et al. focused on showing that their theoretical model can be used to represent the organization of the metabolic network of *E. coli*. They considered each substrate as a node and chemical interactions between substrates as edges and executed several simplifying procedures that did not distort the topology of the network. To identify modules, they constructed an overlap matrix (Figure 3) that shows, for each pair of substrates, the probability that any two substrates i and j are connected to the same other substrate (a *white* empty spot in the matrix indicates that the substrates are not connected

to any common substrate, *green* that the two substrates are connected to very few of the same substrates and *dark red* indicates that they are connected mostly to the same substrates). From the overlap matrix representation as the first step in the system analysis, one can visually identify regions with high overlap and regions with less substrate overlap. The high overlap regions were taken to correspond to the component modules in the proposed type of hierarchical modular organization.

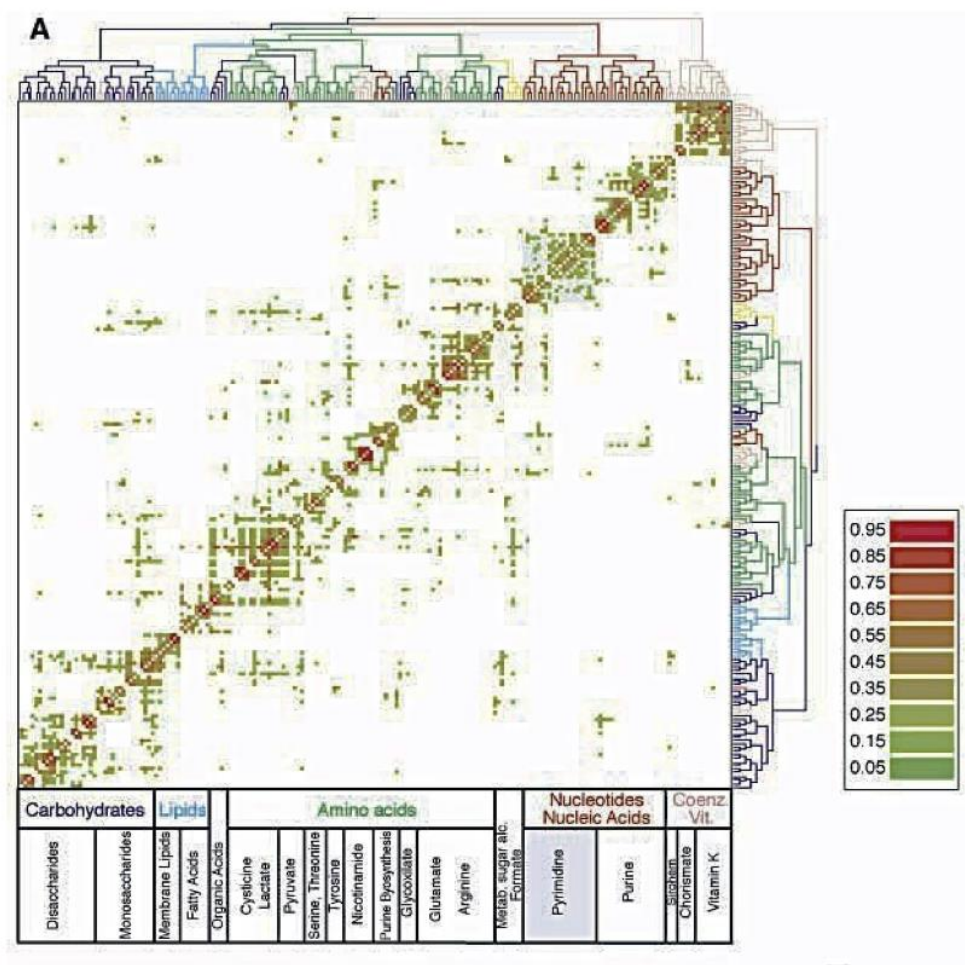


Figure 3. Overlap matrix in which colors indicate degree of overlap. At the bottom the chemical classes of the molecules is indicated and on the top and right the resulting of a clustering algorithm are shown. From Ravasz et al. (2002).

Since the network was assumed to be hierarchical, Ravasz et al. then used a clustering algorithm to group substrates based on their overlap; the results of this analysis are shown along the top and right-hand edges of the matrix square in Figure 3. Because these clusters closely correspond to known chemical groups, they were interpreted as picking out modules that correspond to biochemical functions. Ravasz et al. then focused in more detail on the module corresponding to pyrimidine metabolism and linked each sub-module to particular pathways in the module.

Ravasz et al.'s analysis showed that modules closely corresponding to traditionally characterized biochemical pathways or mechanisms can be recovered from large-scale networks. Noting that the fit between the modules derived via their graph-theoretical analysis and previously established metabolic pathways is not perfect, they suggested that “further experimental and theoretical analyses will be needed to understand the relation between the decomposition of *E. coli* metabolism offered by our topology-based approach and the biologically relevant subnetworks” (Ravasz et al. 2002, p. 1555).

We conclude that while Ravasz et al. succeed in recovering classically defined mechanisms, their discovery strategy differs from more traditional methods of investigating biological modules in the following way. It starts with the whole network and only arrives at modules by identifying clusters in that network—in contrast to starting with an isolated mechanism and looking out to other mechanisms. This hierarchy, achieved by nodes with low clustering but high degree that serve as hubs that connect modules, provides the basis for coordinated regulation of flux in different biochemical pathways. By invoking the graph-theoretic measures of clustering and node degree and characterizing a network architecture compatible with high values on both, Ravasz et al. take advantage of the representational power of graph theory. They propose a new way of thinking about biological mechanisms, according to which they can both be specialized and closely integrated into a larger network—where the latter goes beyond the assumption of traditional mechanistic research that systems can be treated as decomposable. And unlike standard approaches in molecular biology, these network analyses abstract from the molecular details and focus instead on the extracted graph structure which is then shown to represent a type of organization that is biologically relevant.

2.4 Temporal dynamics of activity in networks

Most graph-theoretic studies of biological networks offer a static view of the systems under investigation—edges between nodes represent activity that occurs at some time, but the timing of activities at nodes is not differentiated. In real cells, different metabolic or gene expression activities occur at different time-scales and cellular functioning requires the precise coordination of a large number of events. Accordingly, a proper understanding of cellular function requires the identification of the temporal execution of the interactions

characterized by edges. How can network approaches and graph-theoretical representations incorporate the temporal aspects of biological phenomena? We address this question by analyzing a study by de Lichtenberg, Jensen, Brunak, and Bork (2005) of protein-protein interactions during the cell cycle of yeast.

Protein-protein interactions (PPIs) occur when two or more proteins are shown to be able to bind, which is taken to indicate that they might interact in living cells. High-throughput proteomics technologies have allowed the collection of numerous PPI datasets for many organisms that have led in turn to the development of PPI network analyses investigating principles of cell organization and function. The general goal of such network approaches is to identify complexes of proteins that are co-localized and co-expressed, and to establish how they are organized as modules to perform specific functions within cells. Although the interpretation of PPI data raises a series of methodological concerns (Mackay, Sunde, Lowry, Crossley, and Matthews, 2007; Chatranyamonti, Ceol, Licata, and Cesareni, 2008), de Lichtenberg and colleagues are representative of a larger community of scientists who have argued that robust biologically relevant conclusions can be derived from network analyses by integrating various types of data. Working on budding yeast, *Saccharomyces cerevisiae*, the researchers integrated information about protein complexes from large-scale static PPI datasets with time series data about gene expression data in order to investigate the functional organization of protein complexes during the yeast cell cycle.

Gene expression data allowed de Lichtenberg et al. to distinguish so-called dynamic proteins, which are expressed only at specific steps in the cell cycle, from static proteins, which are constitutively expressed. They then used the PPI data to link the dynamic proteins to those expressed constitutively. Out of the 595 dynamic proteins identified in the study, the researchers focused on the interactions between 184 dynamic proteins and 116 static proteins that are represented by the graphs in the center of Figure 4. Unlike the metabolic network in the previous example, these protein interaction networks do not form a completely connected graph but a collection of independent graphs. The 412 dynamic proteins for which no interactions could be identified are shown as unconnected circles around the outside. The researchers used color to represent the phase of the cell cycle during which these proteins are present (e.g., green for the S phase). The same coloring scheme is used to represent the phase of the cell cycle at which the proteins shown in the center are expressed, with white circles indicating constitutively expressed proteins.




Figure 4. Protein-protein interaction networks in yeast identified by de Lichtenberg et al. The color-coding represents the period in the cell cycle in which specific proteins are synthesized. From de Lichtenberg et al. (2005).

The center of Figure 4 shows 29 graphs of three or more interacting proteins, most of which are associated with particular well-characterized cell-cycle functions. Some of these graphs include only dynamic proteins that are expressed at the same phase in the cell cycle, but most present a much more complex pattern involving proteins expressed constitutively or at different phases in the cell cycle. For instance, the Sister chromatid cohesion graph contains a sizable subgraph of dynamic proteins expressed at the same phase of the cell cycle, while the Cdc-cyclin graph captures a more complex patterns of interacting proteins. The coloring scheme introduced by Lichtenberg et al facilitates the representation of these temporally heterogeneous modules. Focusing on three of these modules (additional modules are discussed in the supplemental material) the authors show how their

integrative approach can be used for further analytical steps: (i) to predict new functional roles for known protein complexes, (ii) to formulate a general principle that governs the construction of protein complexes, and (iii) to provide a dynamic characterization of the interaction between static and dynamic protein complexes. We discuss each of these in turn.

First, some of the modules represented in Figure 4 were newly discovered through analysis of the temporal expression data. An example is the Nucleosome/Bud formation module shown in detail in Figure 5A. The newly identified module includes two proteins, Nis1p and Yol070p, whose function had not been previously identified. By linking them with other proteins in a temporally expressed module whose functional role was already known, and connecting this module to two others, the Sister Chromatid Cohesion complex and the Nucleosome complex (again, whose functions were known), the researchers were able to suggest functional roles for these proteins. For instance, the previously uncharacterized phospho-protein Mmr1p has been shown to have a cell-cycle role in activating transcription at the G₂/M cell phase transition.




Figure 5. Detailed representation of three of the networks shown in Figure 4. A: a network only identified through the use of information about temporal expression of proteins. B: a network illustrating just-in-time assembly. C: a network illustrating a node that serially interacts with proteins synthesized at different times. From de Lichtenberg et al. (2005).

Second, from the analysis of temporal expression data, de Lichtenberg et al. infer a general principle of construction for yeast protein complexes: they are assembled just in time for their use. Many of the proteins in a complex are constitutively expressed and already available, but only assembled into the requisite complex when one or a few additional proteins are synthesized. This process contrasts with the pattern McAdams and Shapiro (2003) found in bacteria—there all the proteins needed to form a complex are synthesized just when the complex is formed. De Lichtenberg et al. illustrate *just-in-time assembly* in Figure 5B. The members of the pre-replication complex, Orc1p to Orc6p, are all constitutively expressed and bound to DNA throughout the cell division

cycle. But their analysis reveals that at different phases in the cycle different proteins are synthesized and bind with the Orc proteins.⁴

Third, de Lichtenberg et al. advance a dynamic perspective on the previously identified module consisting of Cdc28p and nine associated cyclins (which are proteins regulating the cell cycle progression). Previously all the cyclins had been viewed as active parts of the module at all times even though Cdc28p can only interact with a single cyclin at a time. De Lichtenberg et al. established, as shown in Figure 5C, that eight of the cyclins are expressed at different phases of the cell cycle. Thus, although qualifying as a module in the non-dynamic network (which shows any interactions occurring at some time), a dynamic analysis reveals Cdc28p to be a dynamically changing collection of entities: different cyclins constitute it at different phases in the cell cycle. The dynamic study also identified Ypl014p, whose function had previously been unidentified, as a member of Cdc28p complex.

Overall, while proving to be a fruitful means of discovering various molecular features, this type of investigation doesn't start with particular functions (as in classical mechanistic research). Rather than beginning with a mechanism and determining its temporal patterns, as in more traditional mechanistic research, the strategy consists in building a system-wide representation of the time of expression across the whole network which reveals how nodes in a cluster relate temporally throughout the cell cycle. Moreover, this focus on the temporal dimension of the protein interaction network gives rise to a research strategy for discovering new mechanisms and their organization.

3 Dynamical Systems Analyses of Biological Networks

Although often represented as operating at steady-state, living systems are highly dynamic, transitioning between different operating states. The behavior of components of a biological system is dependent on such contexts and in this section we focus on attempts to deal with the context-dependency of causal effects as biological networks enter into different regimes over time. We outline important characteristics of these system dynamic approaches and discuss their philosophical implications by drawing on similarities and differences between two related, yet different, approaches. Both are illustrated in the context of cancer research, a field where analyses are particularly important since cancer involves altered dynamical behavior of cells. The first uses systems analysis as a stepping-stone to identifying specific molecular components that are causal difference-makers, and

⁴ Mcm2p and Mcm3p, along with Cdc46p, Cdc47p, and Cdc54p are all synthesized during the M/G1 cell phase transition and form a complex with constitutively expressed Mcm6p. Next Cdc6p (also expressed at the M/G1 transition) recruits this complex to the Orc complex during the G1 phase. Finally, Cdc45p is expressed early in the S phase and is proposed to then recruit the whole complex to the site where replication originates.

as such retains strong links to mechanistic approaches. In contrast, other system theoreticians aim to go beyond the focus on properties of specific molecular components. To clarify this, Section 3.2 analyzes the philosophical implications of an analysis inspired by Dynamical Systems Theory where network-wide activity is captured by mathematical models and investigated via state-space analysis.

3.1 Difference-making in dynamical networks

One of the cornerstones of molecular biology is reliance on difference-making experiments to identify the causal roles of different components of a system (Woodward, 2003; Craver and Darden, 2013, Ch. 8). Complex systems, however, pose two challenges for this type of experimental practice. First, high-throughput methodologies offer a vast number of potential targets for experimental intervention, and testing all of them individually in various combinations would be prohibitively time and resource demanding. How do scientists narrow the space of promising candidates for intervention? Second, the problem of context-sensitivity: as complex systems change their state over time it is possible for one and the same intervention on a component to have different, or even opposite, effects. How is the causal structure to be investigated when effects are context-sensitive? The present section looks in detail at two cases involving apoptosis, or programmed cell death, in which network strategies helped scientists to elucidate the behavior of the systems—leading to the identification of a partial mechanism underlying cellular context-sensitivity. Our discussion shows that thinking in terms of large-scale networks of interactions is both analytically useful and experimentally productive.

Janes, Albeck, Gaudet, Sorger, Lauffenburger, and Yaffe (2005) investigated the cellular changes that determine whether or not human colorectal adenocarcinoma cells enter into apoptosis. Previous research had generated conflicting data about the role of c-Jun N-terminal kinase (JNK): one study found it to be pro-apoptotic; another found it to be anti-apoptotic; and yet a third found it to be unrelated to apoptosis. The authors hypothesized that changing cellular network states over time might explain these results—JNK as an individual difference-making cause would have different effects depending on the *state of the surrounding signaling network*. To test the hypothesis, they systematically exposed cells to different regimes of apoptosis-related cytokines such as tumor necrosis factor (TNF), epidermal growth factor (EGF), and insulin. Levels of JNK and apoptosis were measured for each regime. The findings of this combinatorial experiment supported the context-sensitivity hypothesis: depending on the levels and combinations of cytokines, increases in JNK were seen to be correlated with increases in apoptosis, decreases in apoptosis, or no change (see especially Fig. 1B-E in Janes et al., not reprinted here).

Reasoning in terms of states of cellular networks provides a conceptual strategy for thinking about context-sensitive effects such as those of the protein JNK. Janes et al. focused on the downstream effects of TNG, EGF and insulin: they analyzed data on a large number of quantitative measures (maximum signal, mean signal, area under the curve, activation

slope, and decay rate) of cellular proteins that were known to be downstream effects of these signal molecules. Applying the statistical tool of principal components analysis, the authors showed that the different cytokines activated different parts of the overall network: while TNF activated a pro-apoptosis group of components, EGF and insulin activated a pro-survival group of components. Interestingly, the reaction to combinations of TNF, EGF, and insulin was not predictable by merely summing the individual responses. Instead, the statistically identified groups exhibited specific interactions when TNF+EGF or TNF+insulin were paired. While EGF mainly antagonized TNF's apoptotic response, insulin exerted its antagonistic effect to TNF by activating the pro-survival components of the network. This indicated the existence of different underlying mechanisms of signal interaction.

Janes et al. thus provided evidence for the existence of both a pro-apoptosis cluster of cellular signals and a pro-survival cluster of signals. Various cytokines were shown to activate one and antagonize the other when combined in different ways. Thus, thinking in terms of changing network states could in principle explain the inconsistent results that had been obtained concerning the JNK protein. Depending on network states, the protein was in fact pro-apoptotic, anti-apoptotic, or neutral.

A subsequent study from the same laboratory, Lee, Ye, Gardino, Heijink, Sorger, MacBeath, and Yaffe (2012), applied the strategy of thinking in terms of changing network states to the treatment of cancer. In the spirit of the conclusion from Janes et al. they proposed that the causal effectiveness of chemotherapeutic agents might depend on the state of cellular networks. Recognizing that this undermines the "component-by-component approaches to understanding human disease" (p. 780), Lee et al. posed new epistemic questions: How can we detect coordinated changes in large numbers of cellular components that serve similar functions? How can the most relevant of these components be identified for further experimental investigation?

Like Janes et al., they began with a systematic combinatorial experiment. They exposed a particular line of cancer cells—triple-negative breast cancer (TNBC) cells—to combinations of signaling inhibitors and DNA-damaging agents (to which TNBC cells are particularly resistant). As in the previous study, apoptosis was the measured effect. Importantly, the authors not only tested different combinations of signal inhibitors and chemotherapeutics. They also studied time-staggered presentations in which cells were given signal inhibitors a variable number of hours before or after chemotherapeutic agents. The results of this systematic approach were striking: giving erlotinib (an inhibitor of EGF receptors) several hours before doxorubicin (a DNA-damaging compound) resulted in a marked, several-fold increase in levels of apoptosis compared to giving erlotinib alone, doxorubicin alone, or erlotinib after doxorubicin. Importantly, this increase in the apoptosis response was found only in TNBC cells and not in other cancers.

As an interpretation of these findings, the authors proposed that giving erlotinib changed the state of one or several cellular signaling networks in TNBC cells, thereby rendering the cells more susceptible to DNA damage induced by doxorubicin. The time delay required for the intervention to work (a delay of roughly 24 hours was found to be optimal) was seen as consistent with the time required for network state changes to take effect, since this may require changes in translation and transcription. To investigate this hypothesis, Lee et al. employed high-throughput methodologies on both TNBC and other cancer cells to identify thousands of genes that are differentially expressed under erlotinib treatment. Using the GeneGO pathway annotation software, they linked these genes to functions such as DNA damage response, apoptosis, and inflammation.

To gain a more precise understanding of the signaling network changes underlying the increased sensitivity to doxorubicin in erlotinib-treated cells, the authors identified 35 candidate genes for a detailed high-throughput analysis (including some previously known genes and some that they found to be differentially expressed under erlotinib treatment). Using principal component analysis combined with partial least-squared regression they statistically identified a cluster of differentially expressed genes in TNBC cells that were not found in other cancer cells. They proposed that these differentially expressed genes were responsible for the increased sensitivity of TNBC cells to the time-staggered treatment with erlotinib and doxorubicin. In particular, they targeted four cellular components in TNBC cells that were strongly correlated with increased doxorubicin efficacy of which the most strongly correlated was cleaved caspase-8. To determine whether caspase-8 really is a cause influencing increased sensitivity to doxorubicin treatment they created small interfering RNA (siRNA) to knock out caspase-8 in TNBC cells. The results of this difference-making experiment confirmed the causal interpretation of the correlation found previously: The intervention reduced the sensitivity of TNBC cells to doxorubicin to the same level as found in other cancer cells. Thus, the success of the time-staggered erlotinib and doxorubicin treatment in TNBC cells could be assumed to involve a caspase-8 pathway specific to these cells.

In this section we have discussed two challenges for difference-making methodologies in the case of complex systems: first, the effects of particular interventions can vary with context; and second, the systems are generally made up of too many components for component-by-component analysis to be feasible. We have seen that network strategies help in dealing with context-sensitivity as well as system size. In one of the cases discussed, investigators found that the same chemotherapeutic agent had different causal effects depending on the previous treatment of cells with cytokines. To elucidate this context sensitivity, the investigators first used combinatorial experiments to determine which combinations of cytokines (and which order of administration) made cells particularly vulnerable to chemotherapy. High-throughput methodologies then revealed the large-scale changes of cellular gene expression under cytokine treatment. The investigators found that individual cytokines caused coordinated changes of whole sets of cellular components,

which suggested changes in the functional state of entire regulatory networks. In a classical difference-making approach, the many affected components would have been intervened upon one-by-one in order to test their relevance to the cell's susceptibility to chemotherapy. Using statistical techniques for network analysis, however, the investigators were able to pinpoint components whose activity seemed to be particularly closely correlated with susceptibility to chemotherapy. The most promising candidates were then subjected to quite classical difference-making experiments in order to determine their actual causal roles.

We here see an extension of classical causal and mechanistic methods by network approaches. In attempting to identify molecular components of interest, the research maintains a tight connection to mechanistic approaches. However, it also considerably expands the set of conceptual strategies for analyzing molecular mechanisms that philosophers have previously articulated to include both statistical methods and analytical strategies for dealing with interconnected complex systems. In the next section we will encounter network researchers who propose a more radical reconceptualization of biological systems.

3.2 Cancer as Dynamic Attractor States

Some systems biologists have proposed a shift of focus from particular mechanisms that are involved in individual cases of cancer to the multitude of ways the organization of cells and tissues can be disrupted and enter specific *dynamic states*. The focus on activity states of a regulatory network rather than specific mutations is motivated by recent empirical studies showing that there is high heterogeneity of mutations in sequenced tumors for non-hereditary cancers (Creixell, Schoof, Erler, and Linding, 2012). The importance of context and reversible dynamics of cell states is further supported by experimental interventions showing that normal cells can turn cancerous when placed next to neoplastic tissues and that cancer cells can be *normalized* if transplanted from tumors to a location next to normal stroma (Lang, Shi, and Chin, 2013).

To account for such experimental results, some researchers appeal to the mathematical framework of *Dynamical Systems Theory* (DST). This approach represents the dynamic state of a system in a state space in which each point represents a possible state of the whole system (e.g., expression rates of genes or concentrations of various proteins). Activity of the system is represented as a trajectory through state space. By mapping the trajectories that biological systems do or can follow, investigators can identify *attractors* to which the system will evolve and where it will remain unless perturbed. A distinctive feature of this modeling strategy, which we will illustrate below, is that the account is based on concrete biological interaction networks, yet—unlike the studies discussed in Section 3.1—the theoretical analysis also captures and visualizes different *dynamic*

possibilities, i.e., system state trajectories that are biologically possible or impossible (beyond a single trajectory taken by an experimental system).

Kauffman (1971) pioneered the idea of describing cancer in terms of attractor states over 40 years ago, building on Waddington's (1940, 1953) much older ideas of epigenetic landscapes. While this approach initially appeared rather speculative, new experimental technologies are providing data that can be used for the development of models and to empirically test the predictions they make. Proponents argue that viewing cancer in terms of dynamic attractor states may resolve some of the problems facing the traditional view that cancer is caused by somatic mutations or specific molecular pathway-interfering causes (e.g., Huang 2011). The view of cancers as attractors aligns with the 'cancer stem cell hypothesis' that highlights similarities between carcinogenesis and developmental processes, such as clonal expansion, fast proliferation, sustained angiogenesis, and tissue invasion. To make sense of how cells with the same genome can differentiate into qualitatively distinct cell types, some researchers investigate the conditions that enable both transformations and stable dynamic states of gene regulatory networks. An important question in this context is why living systems often display stable discrete (or discontinuous) phenotypes, rather than continuous ones that would follow if genetic changes directly caused phenotypic changes (Jaeger and Crombach, 2012).

One approach to understanding the lack of continuity between phenotypic states, including normal and cancer cell states, is to investigate whether there are constraints that make intermediate states unstable so that they follow trajectories into one or another attractor. Kauffman (1969) simulated hypothetical networks in which genes, depending on their states, turned other genes on or off. By starting with different random inputs, he showed that they settled into stable states that were largely independent of the initial states of specific network nodes. Kauffman argued more generally that even highly complex networks often converge to a limited set of stable states because of internal constraints.

New experimental technologies in genomics have started to provide insights into how the organizational structure of gene regulatory networks, through inhibitory and activating connections, may constrain, but also stabilize, the transcriptome signatures of different cell types. We illustrate this approach for a very simple network employing a double negative feedback loop (Figure 6, see also Tyson and Novak, 2010). Its two nodes might represent genes whose products inhibit each other's expression (A and B in Figure 6). The possible dynamic states and trajectories of the two-node network are represented on as a state space in which the expression levels of genes A and B are represented on the corresponding dimensions. The third dimension, labeled 'potential' (or other times, 'energy') indicates how stable (low value) or unstable a particular combination of expression levels is. Trajectories are paths from less stable states to more stable ones, ultimately terminating in attractors. The dynamical landscape shows the state space trajectories the system would follow from various possible states to one of the two attractors. In this case, each of the two attractors consists in high expression of one gene

and low expression of the other. The simple network illustrates a common dynamics that is found in many biological systems. For instance, the situation in Figure 6 corresponds to a simplified model used to describe the development of different segments in *Drosophila* (an instance of discontinuous phenotypic patterns). Protein expression of the gap genes *giant* and *Krüppel* never co-occurs in the same nucleus because the stable expression patterns are mutually exclusive (Jaeger and Crombach, 2012).⁵

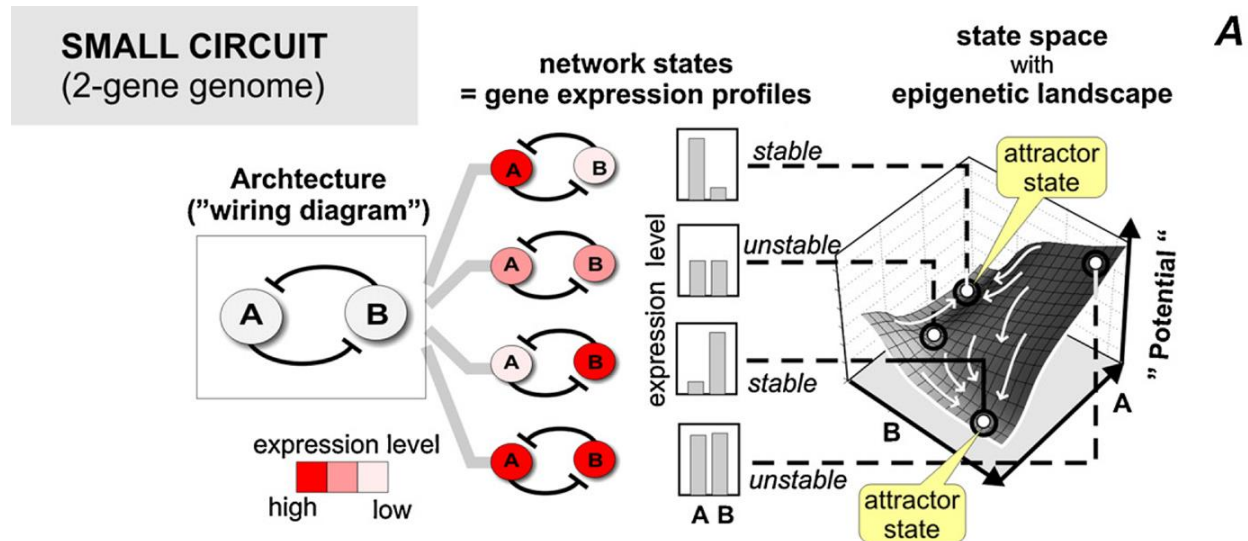


Figure 6. The possible trajectories of a double-negative feedback network between two genes are plotted on a landscape. It has two attractors (stable states of network activity), where some unstable states are also shown. From Huang et al. (2009).

Some researchers have suggested that this visual representation strategy can be also used to indicate what happens in complex networks with thousands of genes. It would require a high-dimensional state space to capture the multitude of stable states, but the number of attractors need not increase proportionately. Because many molecular changes can lead to the same stable cell state, and many lower-level changes will have no effect on the overall state of the system, Huang (2011) argues for a shift of focus from linear molecular pathways (which have been the focus of mechanistic research) to dynamic states of whole networks. Due to the high dimensionality of state spaces for complex networks, this kind of

⁵ The expression patterns change in response to initial conditions, in this context maternally expressed transcription factors that can be measured experimentally and manipulated, and the state space analysis can be based on systems parameters fitted to gene expression data.

analysis is dependent on advanced representational strategies that transform causal interactions into dynamic states of whole networks.




Figure 7. Illustration of multi-stable dynamical profile of a complex network. From Huang et al. (2009).

A network with many nodes but whose dynamical profile (due to constraints within the whole network) has a limited number of attractors is shown in Figure 7. Note that in addition to the two trajectories to stable, mature cell types (indicated by solid arrows), there is another trajectory indicated by a dotted arrow that terminates in an attractor corresponding to cancer. Along the trajectory is a hill denoted as an *epigenetic barrier*. Under normal conditions the system will not cross this barrier, but perturbations, such as somatic mutations, can have the effect of lowering the barrier, allowing the system to settle into the cancer attractor. This landscape representation thereby reveals that the *available trajectories* of network activity can *change* as either the network organization is altered

(e.g., mutations in a cell's gene regulatory network) or the conditions in which the system is operating is altered (e.g., epigenetic and pharmaceutical influences).

Thinking in terms of cancer attractors leads to a shift of focus from how mutations interfere with specific activities in cells to how they reshape the landscape, alter the possible trajectories, and allow access to new cancer attractors (Huang, Ernberg, and Kauffman, 2009). This representational strategy using landscapes also provides the analytical insight of how *the same* mutation can be associated with both carcinogenesis and tumor suppression, depending on the current state of the network as a whole. If the cell is in the normal state, a mutation lowering the barrier between this state and a cancer attractor can allow the system to enter the attractor. If, however, the cell is already trapped in the cancer attractor, the same mutation may allow the system to return to the normal state. Thus, in the analysis of gene regulatory networks through mathematical state space representations, epigenetic landscapes offer an intuitive depiction of how the effects of genetic changes are dependent on the dynamic context of the cell. Huang (2013) appeals to the evidence of normalizing cancer cells via transfer to normal tissue environments to suggest that there may be differentiation therapies that can push cancer cells back to the normal state. Huang and Kauffman (2013) also interpret Lee et al.'s (2012) discovery of the need for sequential application of drugs that we discussed in the previous section from this perspective: the first treatment is needed to change network dynamics to a landscape profile in which the second drug can push the cell back into a normal trajectory.

Although intuitively appealing, it is less clear how the 'normalizing' of network dynamics by means of interventions would work in practice. But an important therapeutic implication suggested by this network perspective is that the most efficient treatment may target a dynamic state of the network as a whole, rather than specific disease-inducing causal pathways. Thus, the analysis directs focus away from 'broken mechanisms' that can be easily delineated towards global dynamical states (Gross, 2011).

Although the DST approach emphasizes that the existence of robust macroscopic dynamics renders many molecular details explanatorily irrelevant for studies of many general features of living systems, the empirical applicability of the models is dependent on extensive gene expression data to bridge the gap between abstract state representations and empirical observables (Huang et al., 2005, 2009). Although researchers have succeeded in building dynamical models from experimental analysis of gene expression profiling for simple regulatory circuits in studies of insect development (Jaeger and Crombach, 2012; Jaeger and Sharpe, 2014), the research is still far from developing multi-dimensional models capturing the dynamic states of cells involved in multiple coordinated processes. The impact of the DST approach on future cancer therapy is therefore an open question. In any case, Huang and Kauffman argue that the DST framework still can offer a fruitful alternative for guiding research and therapy, which is an improvement over what they describe as "*ad hoc*, direct interpretation of gene network topologies" (2013, p. 276).

Among the different network research approaches we have discussed, the Dynamical Systems Theory framework is the one that most decisively differs from traditional mechanistic strategies. Not only does it call for the experimental investigation of larger networks as opposed to individual pathways (a perspective all network approaches share), but in terms of the conceptual analysis of systems, it questions the utility of the heuristic strategy to approach a system as nearly-decomposable (Bechtel and Richardson, 1993; see also Kaiser, 2015). Paradigmatic mechanistic procedures start by assuming that the behavior of a whole can be understood from the individual component functions arranged in a fairly linear sequence (Machamer, Darden, and Craver, 2000), and move toward a more complex organization only once a linear structure has been shown to be inadequate. Yet on the dynamical systems approach, since many different mechanisms can lead to the same attractor state, rather than mapping specific pathways, the proposal is to focus on attractors (which entail possible system trajectories). Even in the case of the existence of only a small number of attractor states (e.g., corresponding to different cell types), these attractors emerge from the activity and dynamical features of the whole network. The two-node network from Figure 6 notwithstanding—which we merely used to illustrate the basic approach—the DST approach presumes that a more *global* representation and analysis of large networks is required, where important characteristics of network functioning cannot be recovered from combining the activities of individual network components (component activities outside the systemic context of the larger network), as a reductionist mechanistic approach assumes.

4 Discussion

We have presented several examples of network analyses that were published in systems biology in recent years. To bring out what is distinctive of these approaches, we have discussed in what respects they align with, extend, or depart from more traditional mechanistic strategies. Our presentation of examples of network approaches started with localist analyses of subgraphs and modules within larger networks and proceeded to more globalist analyses of dynamical activity in whole networks. To a first approximation, we argued that localist analyses directly enhance mechanistic approaches through identification of the ways that mechanisms are organized. Some globalist approaches, on the other hand, either connect indirectly or induce a departure from the previously articulated strategies for mechanism discovery. In the following, we outline the types of relations that can obtain between network approaches and the mechanistic framework by highlighting distinctive contributions made by network approaches in each of the cases examined.

Alon's identification of individual motifs began with the computational screening of a whole network consisting only of gene regulatory connections to identify unusually frequent motifs (Section 2.1). We argued that his procedure differs from more typical

mechanistic strategies which begin by considering a specific function and then reason *locally* about the pathways that realize it. The localist perspective only comes in later in Alon's functional analysis of specific motifs, i.e. after the global search for frequent patterns in the network. Once identified, the motifs were modeled and manipulated experimentally by relying on the mechanistic strategy of decomposition, assuming that a motif's function can be explained in isolation from the rest of the network. This approach adds to mechanistic methods by providing an alternative route to the same type of explanation. At the same time, however, it complements the focus on concrete causal mechanisms with a focus on generalizable design principles that are independent of specific contexts of implementation.

To characterize larger networks, investigators sometimes invoke graph theoretical representations including global measures, such as the node degree distribution across a whole network, which go beyond the arsenal of mechanistic strategies. Discovering clustered nodes (modules) in complex networks using algorithms is relevant for mechanistic research insofar as they can be identified with specific biological functions. But the definition and thus identification of modules is novel. From a traditional mechanistic perspective, a module or part of a mechanism is defined by its concrete internal structure (which underlies its functioning), whereas module identification within large-scale networks initially does not characterize the specific *internal* organization of modules (as the representation of a motif would), but instead identifies a module in terms of its relatively low number of connections *to other* modules of the overall network.⁶ This strategy enables the *identification* of modules without any knowledge about the functions of different system parts or the specific causal nature of their connections. Instead of properties of specific parts or operations, the focus is on the *relational* structures and patterns and their implications for system behavior. In some cases these module clusters were taken to reveal some global features of the network. As we saw in Section 2.3, the motivation for Ravasz et al.'s (2002) research was the challenge of combining a network's functional modularity (segregation into internally highly clustered modules) with the existence of a few highly connected nodes, indicating functional dependency across modules. As a possible solution to this tension they introduced the notion of *hierarchical modular networks* and based on experimental evidence about the metabolic network of *E. coli* showed how various smaller networks might be integrated into the larger metabolic network.

Graph representations of networks are usually limited to static portrayals of biological phenomena. De Lichtenberg et al. (2005) provided an interesting way of incorporating timing information (specifically, timing within the cell cycle) into a graph representation (Section 2.4), and using timing to identify possible mechanisms. The resulting graph revealed 29 interconnected modules in which one or more protein is represented as being

⁶We would like to thank an anonymous reviewer for making this point.

expressed at a specific time. This enabled the researchers to establish that some modules are assembled just at the time they are needed while in other modules a central protein binds with different proteins at different phases in the cell cycle. De Lichtenberg et al. linked these 29 modules to different activities involved in the cell cycle, thereby treating them as mechanisms. The goal, however, was not to map the temporal sequence of complete pathways but to relate the timing of different cellular components so as shed light on their synchronization at the cell (cycle) level. This instance of research employed conceptual strategies to analyze the dynamic organization of networks—computational tools as well as means of visual representation (e.g., Figure 4)—in a way that motivates subsequent analytical exploration and discovery. By capturing some aspects of activity timing these tools enhance the value of standard graph-theoretic representations.

Unlike graph theory, dynamical approaches to networks are explicitly geared at characterizing the temporal character of biological phenomena. Biological systems constantly change their state either due to activities within the system or in response to external perturbations. The studies by Janes et al. (2005) and Lee et al. (2012) showed that the efficacy of drug therapies could be very different depending on the state the network is in at a given time, so that sequential administration of drugs could be effective where simultaneous administration would not be (Section 3.1). To interpret and discover such results in the first place, it was useful for researchers to conceptualize and analyze the system in terms of dynamically changing networks states, where system components may respond differently when in different states. The use of statistical tools on high-throughput data allowed investigators to hone in on putative mechanism components and interactions that warranted further investigation by classical difference-making methodologies.

The final case analyzed in section 3.2. marks a significant departure from a traditional mechanistic perspective by rejecting the core strategy of decomposition, i.e., the assumption that systems can be understood by investigating systems parts and their operation in isolation, and that the behavior of the whole can be recomposed from the parts in a largely sequential fashion. Rather than thinking in terms of specific pathways and mechanisms, such global systems approaches push for the analysis of the dynamics of interconnected networks, revealing possible system trajectories, where the most efficient treatment for a specific disease may be an intervention on the state of the whole network (so as to change the system's dynamical profile consisting in attractors and epigenetic barriers). Representing biological processes as attractors in state space rather than connections between specific molecular entities in a mechanism diagram reveals significant differences in the conceptual approach. A specific attractor state can be reached by many different mechanistic pathways. Therefore, the proposed therapeutic intervention need not target any specific mechanism but the dynamics of the network as a whole. The predictive potential of these newer approaches is still dependent on grounding the dynamic analysis in details of actual biological systems (such as patterns of gene expression or molecular concentrations in cells). Yet, conceiving of cancer in terms of

attractor states has important implications for the view of what cancer is, as it presents a therapeutic alternative to the focus on specific somatic mutations or specific molecular pathway-interfering causes.

Apart from the important, although generic mechanistic strategies of decomposition and localization (originally articulated by Bechtel and Richardson 1993), Craver and Darden (2013) have recently laid out further strategies for mechanism discovery. However, their conceptual strategies, such as modular subassembly, forward chaining, and backward chaining, come from philosophically investigating traditional molecular biology and pertain to reasoning about individual pathways. Our discussion of different instances of network research has covered additional strategies for analyzing systems, which are geared toward the characterization of complex molecular networks. This research often starts out with experimentally acquired data about large-scale networks, so that (although sometimes further experimental steps follow) a good deal of the intellectual burden is on employing—and even developing—strategies for analyzing some aspects of the structure and function of complex networks, so as to provide guidance for further discovery steps. In addition to the use of mathematical measures and statistical and computational tools, these novel network strategies include means of visually representing system aspects. Many of such representations are neither mechanism diagrams nor depictions of the whole network. A good example is the overlap matrix depicted in Figure 3 (from Ravasz et al., 2002), which zeroes in on some aspects of system activity, and not only summarizes previous information, but has the crucial intellectual function of guiding further analysis and discovery. Dynamic landscape representations like the ones in Figures 6 and 7 also go beyond mechanism and network diagrams because the landscapes are derived from a theoretical analysis of network functioning, and point to system potentialities such as possible therapeutic targets that can be explored in future discovery.

Our aim in this paper has been to motivate a philosophical analysis of network approaches as a means increasingly used by life scientists to deal with biological complexity. Throughout the paper we have emphasized the benefits of network approaches for biology, which use the tools of graph theory and dynamical systems theory. Biologists have been drawn to these tools as they confronted the challenge of coping with the complexity of highly interconnected and non-linear biological systems. We have emphasized both how the various network strategies can be (and have been) used to complement mechanistic approaches but also how there is already a wide spectrum of network approaches where some differ from traditional mechanistic strategies. As these methods and their applications are in constant development, we think that future careful philosophical analyses can further clarify the types of inferences that these methods support in the investigation of biological phenomena.

Acknowledgements

This paper was initiated while the authors were fellows at the Center for Philosophy of Science, University of Pittsburgh. We would like to thank the Center for providing us such a rich environment that was conducive to establishing this collaboration. Ingo Brigandt's work is also supported by the Social Sciences and Humanities Research Council of Canada (Insight Grant 435-2016-0500).

References

- Alon, U. (2007). *An introduction to systems biology: Design principles of biological circuits*. Boca Raton, FL: Chapman & Hall/CRC.
- Barabási, A.-L., & Albert, R. (1999). 'Emergence of scaling in random networks'. *Science*, 286, 509-512.
- Beber, M. E., Fretter, C., Jain, S., Sonnenschein, N., Muller-Hannemann, M., & Hutt, M. T. (2012). 'Artefacts in statistical analyses of network motifs: general framework and application to metabolic networks'. *Journal of the Royal Society Interface*, 9, 3426-3435.
- Bechtel, W. (2011). 'Mechanism and biological explanation'. *Philosophy of Science*, 78, 533-557.
- Bechtel, W., & Abrahamsen, A. (2005). 'Explanation: A mechanist alternative'. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 421-441.
- Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.
- Brigandt, I., Green, S., & O'Malley, M. A. (in press) 'Systems biology and mechanistic explanation'. In S. Glennan and P. Illari (Eds), *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. New York: Routledge.
- Chatr-aryamontri, A., Ceol, A., Licata, L., & Cesareni, G. (2008). 'Protein interactions: integration leads to belief'. *Trends in Biochemical Sciences*, 33, 241-242.
- Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. Chicago: University of Chicago Press.
- Creixell, P., Schoof, E. M., Erler, J. T., & Lindig, R. (2012). 'Navigating cancer network attractors for tumor-specific therapy'. *Nature Biotechnology*, 30, 842-848.
- de Lichtenberg, U., Jensen, L. J., Brunak, S., & Bork, P. (2005). 'Dynamic complex formation during the yeast cell cycle'. *Science*, 307, 724-727.
- Green, S. (2015). Revisiting generality in biology: systems biology and the quest for design principles. *Biology & Philosophy*, 30(5), 629-652.

- Green, S., Levy, A., & Bechtel, W. (2014). 'Design sans adaptation'. *European Journal for Philosophy of Science*, 5, 15-29.
- Gross, F. (2011). 'What systems biology can tell us about disease'. *History and Philosophy of the Life Sciences*, 33, 477-496.
- Huang, S. (2004). 'Back to the biology in systems biology: what can we learn from biomolecular networks?' *Briefings in Functional Genomics and Proteomics*, 2, 279-297.
- Huang, S. (2011). 'Systems biology of stem cells: three useful perspectives to help overcome the paradigm of linear pathways'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366, 2247-2259.
- Huang, S. (2013). 'Genetic and non-genetic instability in tumor progression: link between the fitness landscape and the epigenetic landscape of cancer cells'. *Cancer and Metastasis Reviews*, 32, 423-448.
- Huang, S., Eichler, G., Bar-Yam, Y., & Ingber, D. E. (2005). 'Cell fates as high-dimensional attractor states of a complex gene regulatory network'. *Physical Review Letters*, 94, 128701.
- Huang, S., Ernberg, I., & Kauffman, S. A. (2009). 'Cancer attractors: a systems view of tumors from a gene network dynamics and developmental perspective'. *Seminars in Cell & Developmental Biology*, 20, 869-876.
- Huang, S., & Kauffman, S. (2013). 'How to escape the cancer attractor: Rationale and limitations of multi-target drugs'. *Seminars in Cancer Biology*, 23, 270-278.
- Huneman, P. (2010). 'Topological explanations and robustness in biological sciences'. *Synthese*, 177, 213-245.
- Isalan, M., Lemerle, C., Michalodimitrakis, K., Horn, C., Beltrao, P., Raineri, E., Garriga-Canut, M., & Serrano, L. (2008). 'Evolvability and hierarchy in rewired bacterial gene networks'. *Nature*, 452, 840-845.
- Jaeger, J., & Crombach, A. (2012). 'Life's attractors : understanding developmental systems through reverse engineering and in silico evolution'. *Advances in Experimental Medicine and Biology*, 751, 93-119.
- Jaeger, J., & Sharpe, J. (2014). 'On the concept of mechanism in development'. In A. Minelli & T. Pradeu (Eds.), *Towards a theory of development* (pp. 56-78). Oxford: Oxford University Press.
- Janes, K. A., Albeck, J. G., Gaudet, S., Sorger, P. K., Lauffenburger, D. A., & Yaffe, M. B. (2005). 'A systems model of signaling identifies a molecular basis set for cytokine-induced apoptosis'. *Science*, 310, 1646-1653.
- Janes, K. A., & Lauffenburger, D. A. (2013). 'Models of signalling networks - what cell biologists can gain from them and give to them'. *Journal of Cell Science*, 126, 1913-1921.
- Jones, N. (2014). 'Bowtie structures, pathway diagrams, and topological explanation'. *Erkenntnis*, 79, 1135-1155.

- Kaiser, M. I. (2015). *Reductive explanation in the biological sciences*. Cham: Springer.
- Kauffman, S. A. (1969). 'Metabolic stability and epigenesis in randomly constructed genetic nets'. *Journal of Theoretical Biology*, 22, 437-467.
- Kauffman, S. A. (1971). 'Differentiation of malignant to benign cells'. *Journal of Theoretical Biology*, 31, 429-451.
- Lang, J. Y., Shi, Y., & Chin, Y. E. (2013). 'Reprogramming cancer cells: back to the future'. *Oncogene*, 32, 2247-2248.
- Lee, M. J., Ye, A. S., Gardino, A. K., Heijink, A. M., Sorger, P. K., MacBeath, G., & Yaffe, M. B. (2012). 'Sequential application of anti-cancer drugs enhances cell death by re-wiring apoptotic signaling networks'. *Cell*, 149, 780-794.
- Levy, A., & Bechtel, W. (2013). 'Abstraction and the organization of mechanisms'. *Philosophy of Science*, 80, 241-261.
- Machamer, P., Darden, L., & Craver, C. F. (2000). 'Thinking about mechanisms'. *Philosophy of Science*, 67, 1-25.
- Mackay, J. P., Sunde, M., Lowry, J. A., Crossley, M., & Matthews, J. M. (2007). 'Protein interactions: is seeing believing?' *Trends in Biochemical Sciences*, 32, 530-531.
- Mangan, S., Zaslaver, A., & Alon, U. (2003). 'The coherent feedforward loop serves as a sign-sensitive delay element in transcription networks'. *Journal of Molecular Biology*, 334, 197-204.
- McAdams, H. H., & Shapiro, L. (2003). 'A bacterial cell-cycle regulatory network operating in time and space'. *Science*, 301, 1874-1877.
- Newman, M. E. J. (2010). *Networks : An introduction*. New York: Oxford University Press.
- Ravasz, E., Somera, A. L., Mongru, D. A., Oltvai, Z. N., & Barabási, A. L. (2002). 'Hierarchical organization of modularity in metabolic networks'. *Science*, 297, 1551-1555.
- Shen-Orr, S. S., Milo, R., Mangan, S., & Alon, U. (2002). 'Network motifs in the transcriptional regulation network of *Escherichia coli*'. *Nature Genetics*, 31, 64-68.
- Tyson, J. J., & Novak, B. (2010). 'Functional motifs in biochemical reaction networks'. *Annual Review of Physical Chemistry*, 61, 219-240.
- Waddington, C. H. (1940). *Organisers and genes*. Cambridge: Cambridge University Press.
- Waddington, C. H. (1953). 'The interactions of some morphogenetic genes in *Drosophila melanogaster*'. *Journal of Genetics*, 51, 243-258.
- Wagner, A., & Fell, D. (2001). 'The small world inside large metabolic networks'. *Proceedings of the Royal Society of London B: Biological Sciences*, 280.
- Watts, D., & Strogatz, S. (1998). 'Collective dynamics of small worlds'. *Nature*, 393, 440-442.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

