

Strategies of Explanatory Abstraction in Molecular Systems Biology[†]

Nicholaos Jones[‡]

Abstract

I consider three explanatory strategies from recent systems biology that are driven by mathematics as much as mechanistic detail. Analysis of differential equations drives the first strategy; topological analysis of network motifs drives the second; mathematical theorems from control engineering drive the third. I also distinguish three abstraction types: aggregations, which simplify by condensing information; generalizations, which simplify by generalizing information; and structurations, which simplify by contextualizing information. Using a common explanandum as reference point—namely, the robust perfect adaptation of chemotaxis in *Escherichia coli*—I argue that each strategy invokes a different combination of abstraction types and that each targets its abstractions to different mechanistic details.

1 Introductory Remarks

The currently dominant paradigm for understanding explanation in biology puts mechanism at center stage (Nicholson 2012; Levy 2013). Leading accounts of mechanistic explanation, while differing in the particulars of their analysis of *mechanism*, agree that mechanistic explanations explain by alluding to mechanisms or models thereof (Machamer, Darden, Craver 2000; Bechtel and Abrahamsen 2005).

There is a small publishing industry devoted to discerning the scope of mechanistic explanation in scientific practice. Some claim to identify biological explanations that do not allude to mechanisms (Wouters 2007; Huneman 2010; Rice 2015). Fans of mechanistic explanation tend to resist making scope concessions, preferring instead to accommodate the putative explanations as mechanistic despite initial appearances, to broaden the scope of mechanistic explanation or the analysis of *mechanism*, or else to

[†] Draft. For symposium on *Integrating Explanatory Strategies Across the Life Sciences* at the 2016 meeting of Philosophy of Science Association, Atlanta, GA. I thank audiences at Mississippi State University, the Alabama Philosophical Society, and the Society for Philosophy of Science in Practice for comments on earlier drafts.

[‡] Department of Philosophy, University of Alabama in Huntsville, Huntsville AL 35899, nick.jones@uah.edu

deny that the putative explanations are explanations at all (Craver 2006; Bechtel and Abrahamsen 2010; Brigandt 2013; Levy and Bechtel 2013).

I set aside questions about what qualifies as an explanation as well as questions about whether only mechanisms—or models thereof—carry explanatory power. I focus, instead, on *explanatory strategies*, understood as patterns of reasoning directed toward providing explanations. I consider three explanatory strategies from recent systems biology that are driven by mathematics as much as, if not more than, mechanistic detail. Analysis of differential equations drives the first strategy; topological analysis of network motifs drives the second; mathematical theorems from control engineering drive the third.

Systems biologists use these strategies to supplement the explanatory power of traditional molecular mechanisms (see Brigandt et al *forthcoming*). My aim is to identify how the strategies differ from each other, rather than how they differ from standard mechanistic explanations or what might unify them in those differences (for which see Green and Jones 2016). Doing so helps with understanding relations among the strategies, their tactics for integrating mechanistic detail, and explanatory affordances of their mathematical elements.

The key to my analysis is a distinction among three abstraction types: aggregations, which simplify by condensing information; generalizations, which simplify by generalizing information; and structurations, which simplify by contextualizing information. Using a common explanandum as reference point—namely, the robust perfect adaptation of chemotaxis in *Escherichia coli* (Barkai and Leibler 1997; Ma et al 2009; Yi et al 2000)—I argue that each strategy invokes a different combination of abstraction types and that each targets its abstractions to different mechanistic details. I begin with the typology of abstraction.

2 Abstraction Typology

I am interested in abstractions as representational rather than metaphysical. Abstractions, as I understand them, are ontologically innocent, so that characterizing features of representations as abstractions over some parts of reality carries no implication that features correspond to abstract objects (see also Cartwright 1989, 353-354; Levy and Bechtel 2013, 243). So, for example, representing the relation between a person, a hotel, and a date range as a reservation does not entail that some abstract object, a *reservation*, exists; nor does representing the motions of an object's constituents as the motion of the object's center of mass entail that some abstract object, a *center of mass*, exists.

Levy and Bechtel characterize a representation as abstract insofar as a more concrete representation is possible (2013, 242). Brigandt and colleagues suggest that biologists use abstractions to “elucidate system-level patterns of organization that may not be visible at the level of molecular details” (*forthcoming*). I concur. I understand abstractions as representing only some of the many elements—objects, relations, parameters—associated with their targets, thereby making apparent patterns obscured by more detailed representations. I add to these insights that biologists produce (at least) three types of abstraction.

Following Ordorica, I call the first *aggregation* (2015, 163-164). An aggregation represents some relationship among multiple elements of a representational target as a higher-level object, or multiple elements of the target as a single, composite object. (See Figure 1a.) Paradigm cases of aggregations include representations of person-hotel-date relations as *reservations*; of costs of services and costs of goods as *costs*; and of the motions of an object’s parts as the *motion of a center of mass* (from Ordorica 2015, 164). Aggregations abstract from plurality to individual, ignoring differences among many in order to make salient some integrated unity among the elements of a representational target. They thereby simplify representations by condensing information about representational targets.

Following Pincock, I call the second abstraction type *generalization* (2015, 864). A generalization represents some element of a representational target as a class of elements, where potential instances of the class might include elements not present in the target. (See Figure 1b.) For example, because the class of solution measures includes all soap-bubble-like surfaces, such as the cellular froth surrounding radiolarian protozoa, representing a soap-bubble surface as a “solution measure” is a generalization (Pincock 2015, 864). Generalizations abstract from an instance to a class thereof, ignoring differences between instances of the class in order to make salient some more general unity. They thereby simplify representations by generalizing from information about representational targets.

I call the third abstraction type *structuration*. A structuration represents some element of a representational target as a position in a structure, such that potential occupants of the position might include elements not present in the target. (See Figure 1c.) I follow Haslanger in understanding structures as “complex entities with parts whose behavior is constrained by their relation to other parts” (2016, 118). Paradigm cases of structurations include representing Barack Obama as President of the United States of America, or representing Alneias as son of Anchises and Aphrodite. Structurations abstract to a position in a structure, from an occupant of the position, ignoring intrinsic features of the occupant unrelated to its position in order to make salient the

occupant's role relative to occupants of other positions in the same structure. They thereby simplify by contextualizing information about representational targets.




Figure 1: Visualizing Abstraction Types. (a) Aggregation A represents elements e_1 , e_2 , and e_3 (and relations therein) as a single object. (b) Generalization C represents $I_1(c)$ as a class, instances of which also include $I_2(c)$ and $I_3(c)$. (c) Structuration p_1 represents element o_1 as a position in larger structure that also includes p_2 , p_3 , and p_4 .

I understand aggregations as distinct from both generalizations and structurations, by virtue of being many-to-one, rather than one-to-one, simplifications. I also understand being a generalization as insufficient for being a structuration. For representations of positions carry information about functional relationships between their occupants and other positions in the same structure; but representations of classes do not. Finally, insofar as classes are sets, I understand being a structuration as insufficient for being a generalization. For, sometimes, representing target elements as classes carries some information about intrinsic features of those elements apart from their functional relations to elements occupying other positions in the same structure; but representing target elements as positions in structures never carries such information.

3 Robust Perfect Adaptation of *E.coli* Chemotaxis

My central claim is that different explanatory strategies from recent systems biology differ from each other, at least in part, by virtue of appealing to different abstraction types. I support this claim by considering a case in which multiple strategies target the same explanandum. Doing so minimizes confounds that confuse differences due to the nature of each explanatory strategy with differences due to the nature of each

explanatory target. I focus on a particular explanandum known as robust perfect adaptation of bacterial chemotaxis, following others who consider this a paradigmatic target for non-mechanistic explanation (Braillard 2010; Brigandt, Green, and O'Malley *forthcoming*; Matthiessen *forthcoming*).

3.1 Explanandum Context

Escherichia coli (*E.coli*) is popular model organism in biological research. It is very sensitive to small chemical changes over a very large range of background concentrations. It also has a simple and well-understood signal transduction network (Wadhams and Armitage 2004).

E.coli manages two kinds of motion (Berg 2003). It *runs* by rotating its flagellar motor counterclockwise. This aligns all of its flagella into a synchronized bundle, resulting in movement in a straight line for about 1 second. *E.coli* also *tumbles* by rotating its flagellar motor clockwise. This breaks flagellar alignment, and the asynchronized flagella produce stationary changes of direction lasting for about 0.1 second. *E.coli* are randomly reoriented after each tumble. Moreover, while these tumbles occur with regular frequency, *E.coli* with higher concentrations of CheR protein tumble more frequently (Spudich and Kochland 1975).

E.coli's motion in a uniform external environment resembles a random walk. *E.coli* has no ability to control or select its direction of motion, and its straight runs are subject to Brownian motion because of eddies. However, in the presence of a chemical attractant—amino acids such as serine or aspartic acid, or sugars such as maltose or glucose—*E.coli taxis* toward the attractant. This taxi behavior involves less frequent tumbles, leading to longer runs and so gradual motion toward the attractant. (There is an opposite behavior for repellants such as metal ions or leucine.)

The biomolecular mechanism for *E.coli* chemotaxis is well-understood. When an environmental attractant attaches to a receptor, the receptor lowers the activity of the CheW-CheA protein complex. Less activity from this complex reduces the rate of CheY phosphorylation, which results in less phosphorylated CheY diffusing to the flagella. Because CheY induces clockwise rotation of the flagellar motor, the outcome is less frequent tumbling.

3.2 Explanandum Question

Alon and colleagues have experimental verification that, in the presence of a chemical attractant mixed uniformly into the environment at a constant concentration, *E.coli* chemotaxis *perfectly adaptive* (Alon et al 2009). After a brief period of decreased tumbling frequency, the frequency of *E.coli* tumbles increases toward and returns to the

exact frequency prior to the introduction of the attractant. The effect of the attractant, accordingly, becomes entirely forgotten despite its continuing presence.

The biomolecular mechanism for the adaptiveness of chemotaxis for *E.coli* is also well-understood. Some time after a new attractant has been detected by receptors, the lower activity of the CheW-CheA complex induces less CheB activity. This reduces the rate for removing methyl groups from the CheW-CheA complex and, together with continual methylation of the CheR receptor, CheW-CheA methylation increases. More methylation means more CheW-CheA activity, which in turn induces more CheY phosphorylation. This eventually results in more phosphorylated CheY diffusing to the flagellar motor, which increases clockwise motor rotation and thereby raises tumbling frequency.

Alon and colleagues have further experimental verification that this perfectly adaptive chemotaxis of *E.coli* is *robust* across ranges of CheR concentrations 0.5 to 50 times higher than concentration levels in “wild type” *E.coli* (Alon et al 2009). (By contrast, *E.coli*’s adaptation time—the time to return to 50% of its pre-stimulus tumbling frequency—is not robust to different CheR concentrations, because more CheR entails longer adaptation times.) This is the explanandum of interest: why is the perfect adaptation of *E.coli* chemotaxis, in the presence of a well-distributed chemical attractant, robust to CheR protein concentrations?

There are (at least) three strategies for answering this question in recent systems biology literature. (For a fourth, see Kollman et al 2005.) I consider each in turn, first sketching the general strategy and then making explicit the abstractions at work.

4 Distinguishing Explanatory Strategies through Abstraction Types

4.1 Dynamical Modeling

I call the first strategy *dynamical modeling*. This strategy begins by constructing a chemotaxis network for *E.coli*. This network represents the mechanism for *E.coli* chemotaxis, including specific biochemical details about when and how relevant proteins affect each other. (See Figure 2.) For example, Barkai and Leibler (1997) construct a model according to which, among many other specifics, CheB demethylates only the active form of the CheW-CheA complex and CheR works only at saturation.

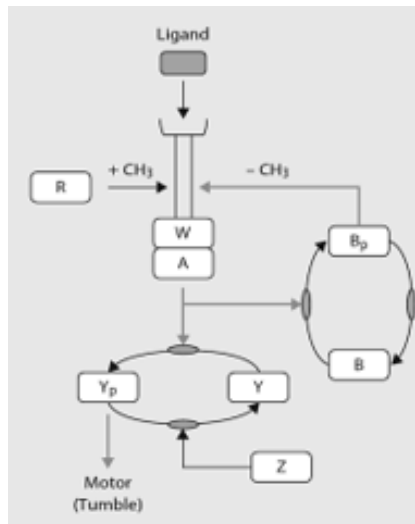


Figure 2. Mechanistic network for *E.coli* chemotaxis (Rao and Ordal 2009).

The dynamical modeling strategy proceeds by constructing a dynamical model—typically a set of differential equations—from the network (see Jones and Wolkenhauer 2012). One then demonstrates, via mathematical proof or simulation, that this model predicts perfect adaptation in the presence of a well-distributed chemical attractant for CheR concentration values varying over several orders of magnitude. (Raerinne 2013 calls this *sensitivity analysis*.) The demonstration supports the inference that *E.coli* chemotaxis exhibits robust perfect adaptation *because of its biochemical specifics*.

Bechtel and Abrahamsen (2010) call the product of this strategy a *dynamical mechanistic explanation*. I set aside the issue of whether the dynamical modeling strategy produces explanations. But I endorse Bechtel and Abrahamsen’s insight that the dynamical modeling strategy produces accounts that are mechanistic, by virtue of depending upon mechanistic details, as well as dynamical, by virtue of analyzing mathematical models built upon those details. For example, Barkai and Leibler’s (1997) mathematical analysis is relevant to *E.coli* chemotaxis only insofar as their network details are relevant; and analysis of the network apart from the model cannot produce an inference about the *robustness* of *E.coli*’s perfectly adaptive chemotaxis.

Let’s treat the dynamical model driving this explanatory strategy as an initial baseline for evaluating the number and severity of abstraction in various explanatory strategies. The model is abstract in various ways. But we shall treat it as a recipient of further abstractions, in the way a vehicle receives freight. Just as we can determine the weight of the freight indirectly by subtracting the gross weight of vehicle and freight from the “tare weight” (the weight of vehicle alone), we shall determine abstraction variety and

severity/extent for models driving other explanatory strategies by “subtracting” their total abstraction variety and severity from the “tare” abstraction.

4.2 Topological Analysis

I call the second explanatory strategy *topological analysis*. This strategy begins by identifying all possible minimal adaptation networks capable of predicting robust perfect adaptation for *E.coli* chemotaxis. These networks, like the networks for dynamical modeling, represent mechanisms for *E.coli* chemotaxis. Yet, unlike the networks for dynamical modeling, these networks are minimal: they contain the fewest possible nodes and links that suffice for robustly perfectly adaptive chemotaxis. The procedure for identifying all possible minimal networks of this sort is brute computational search. It turns out that there are exactly three, each of which has exactly three nodes and no more than three links (Ma et al 2009).

The topological analysis strategy proceeds by identifying a chemotaxis network known to predict robust perfect adaptation. This strategy thereby relies upon the dynamical modeling strategy, but only for mathematical results. The biochemical details of the chosen chemotaxis network turn out to be largely irrelevant, because the topological analysis strategy proceeds by demonstrating that a *reduced form* of the chosen network is topological equivalent to one of the minimal adaptation models. Reduced forms for mechanistic networks functional equivalents for node groups, group nodes or equivalents into modules, and ignore links within modules in favor of links between modules.

Consider, for example, one of the three minimal adaptation networks Ma and colleagues (2009) discover for *E.coli* chemotaxis. (See Figure 3.) The network has an input activating node A, A inhibiting being activated by B, A also activating C, and C activating some output. Ma and colleagues show that Barkai and Leibler’s (1997) model for *E.coli* chemotaxis reduces to this minimal network. Barkai and Leibler have an input and CheR activating, and CheB inhibiting, receptors; these receptors activating the CheW-CheA complex; the complex activating CheB and CheY; and CheY activating some output. Ma and colleagues reconceptualize Barkai and Leibler’s network into one where the input activates a *receptor complex*; this complex activates CheY, which activates the output; the complex also activates CheB, which inhibits a *methylation level* also activated by CheR.; and this methylation level activates the receptor complex. Then, in a second reconceptualization that produces one of their minimal adaptation networks, they group the receptor complex and CheB into module A, group CheR and the methylation level into module B, and rename CheY module C.

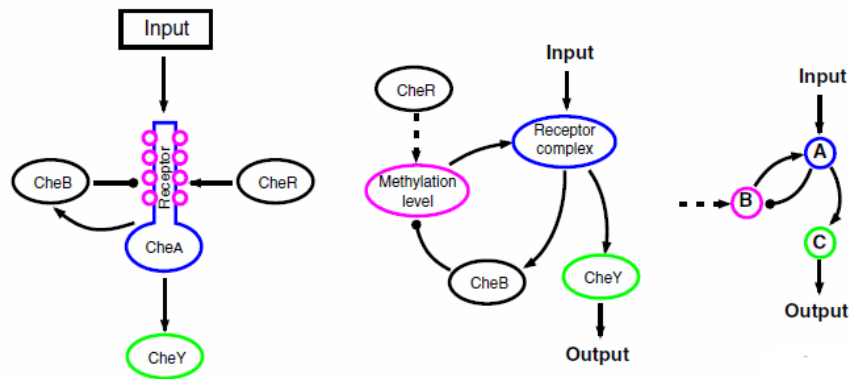


Figure 7. The Network of Perfect Adaptation in *E. coli* Chemotaxis Belongs to the NFBLB Class of Adaptive Circuits
 Left: the original network in *E. coli*. Middle: the redrawn network to highlight the role and the control of the key node "Methylation Level." Right: one of the minimal adaptation networks in our study.

Figure 3: Network topology for *E. coli* chemotaxis (Ma et al 2009).

The topological analysis strategy infers, from the topological equivalence between a minimal adaptation network and the reduced form of a network known to predict robust perfect adaptation for chemotaxis, that *E. coli* chemotaxis exhibits robust perfect adaptation *because of the topology of its chemotaxis network*. Huneman (2010) calls the product of this strategy a *topological explanation*. Regardless of whether analyses such as Ma and colleagues's are explanatory, they are topological by virtue of demonstrating some consequence about the topological properties of a network. This means that, even if the mechanistic details of *E. coli*'s chemotaxis network were different, and even if the biochemical specifics of the network chosen for reduction were different, the product of the topological analysis strategy would remain the same provided that the alternative networks preserve topological equivalence with the originals (see also Jones 2014).

The topological model driving this second explanatory strategy is more abstract than the dynamical model driving our initial ("tare") strategy. The topological model contains more aggregations. For example, it represents CheY and CheZ as "the motor rotation group;" it represents CheA and CheW as "the receptor complex;" and it represents the receptor complex and CheB as "the phosphorylation group." The topological model also contains more structurations. For example, it represents the phosphorylation group as "A" and the motor rotation group as "C." These representations abstract entirely from any intrinsic marks that might distinguish instances of "A" from instances of "C," relying instead upon extrinsic relations to distinguish the nodes from each other. So, for example, "A" but not "C" inhibits "B," "A" activates "C," and so on.

4.3 Organizational Design

I call the third explanatory strategy *organization design*. This strategy begins with a proof to the effect that systems exhibit robust perfect adaptation if and only if they

satisfy the characteristic equation for Integral Feedback Control (IFC). The proof is purely mathematical, well-known from control engineering theory in contexts involving mechanical systems that exhibit IFC such as thermostats. I am not aware of a complete and published version of this proof, but Yi and colleagues (2000) provide a sketch with relevant details. The organizational design strategy proceeds by inferring that *E.coli* chemotaxis exhibits robust perfect adaptation if and only if it satisfies the characteristic equation for IFC, and further inferring that *E.coli* chemotaxis exhibits robust perfect adaptation *because it satisfies the characteristic equation for IFC*. (For better explanatory details regarding this specific case, Braillard 2010; Green and Jones 2016.)

The organizational design strategy invokes neither mechanistic specifics about the chemotaxis network for *E.coli* nor topological details about the structure of that network. The strategy takes the explanandum phenomenon as given, using a mathematical equivalence result to identify a principle both necessary and sufficient for the phenomenon. The strategy thereby has affinities with explanatory strategies that appeal to organizing principles (Green and Wolkenhauer 2013) and design principles (Green 2015).

For simplicity, let's "reset" our abstraction "tare" to the topological model, because the model driving the organizational design strategy—call it the design model—is abstract in all the ways the topological model is abstract and more besides. The simplification thereby focuses attention on ways in which the design model differs from the topological model—and, by extension, from the initial dynamical model.

Compared to the topological model, the design model contains more aggregations. For example, the design model represents CheY phosphorylation and CheB activation as "*k*-box output." This aggregation is, at the same time, a generalization and a structuration. For example, "*k*-box output" is a class, with instances biological as well as mechanical. The standard example of a mechanical instance is heater activation in a thermostat. The *k*-box representation is also a structuration, akin to the "A", "B," and "C" representations from the topological model. For the *k*-box represents whatever has such-and-such input and output (a position in a structure). (See Figure 4.)




Fig. 2. A block diagram of integral feedback control. The variable u is the input for a process with gain k . The difference between the actual output y_1 and the steady-state output y_0 represents the normalized output or error, y . Integral control arises through the feedback loop in which the time integral of y , x , is fed back into the system. As a result, we have $x = y$ and $y = 0$ at steady-state for all u . In the Barkai-Leibler model of the bacterial chemotaxis signaling system, the chemoattractant is the input, receptor activity is the output, and $-x$ approximates the methylation level of the receptors.

Figure 4 Organizational design for bacterial chemotaxis...and thermostats (Yi et al 2000).

The topological model is more abstract than the dynamical model, by virtue of containing various abstractions over protein identities. The design model, in turn, is more abstract than the topological and dynamical models, by virtue of also containing various abstractions over protein interactions. We can, therefore, arrange the various explanatory strategies along a continuum of abstraction type and severity. The dynamical modeling strategy, as our baseline, occupies the “low” end of our continuum. Next is topological analysis, which involves aggregations of and structurations from protein identities (or aggregations thereof). Then there is organizational design, which also involves aggregation of protein interactions as well as generalization and structuration of protein identities (or aggregations thereof).

5 Confirming the Analysis

I consider the foregoing to establish that each explanatory strategy invokes a different combination of abstraction types and that each targets its abstractions to different mechanistic details. Whether this result generalizes beyond my chosen case study awaits future research. There is some reason to expect an affirmative result. For if dynamical, topological, and design explanatory strategies differ as I claim—specifically, along dimensions of number and severity of generalizations and structurations—then we should expect the *more abstract* strategies to have *wider scope*. For the more general models likely have more instances, and the more structural models likely have more position occupants.

We find confirmation of this prediction for the case of robust perfect adaptation of *Bacillus subtilis* (*B.subtilis*) chemotaxis. Details of the organization design strategy for explaining why *E.coli* chemotaxis exhibits robust perfect adaptation *also* apply for explaining why *B.subtilis* chemotaxis exhibits robust perfect adaptation. But details of the corresponding dynamical mechanistic strategy do not. The organization design strategy, as we know, involves more generalization and structuration than the dynamical mechanistic strategy. This confirms our prediction.

Allow me to be brief with the details. Rao and Ordal (2007) develop a dynamic mechanistic explanation for the perfect robustness of chemotaxis for *B.subtilis*. Their explanatory strategy follows the same pattern as Barkai and Leibler's in the case of *E.coli*. But details differ. For example, according to Barkai and Leibler's model, CheB in *E.coli* demethylates only active receptor complexes; according to Rao and Ordal, CheB in *B.subtilis* demethylates inactive ones too. Again, according to Barkai and Leibler's model, without CheY *E.coli* runs but does not tumble; according to Rao and Ordal, without CheY *B.subtilis* tumbles but does not run. One more: according to Barkai and Leibler's model, *E.coli* without CheB cannot run; according to Rao and Ordal, *B.subtilis* without CheB can run. See Figure 5.

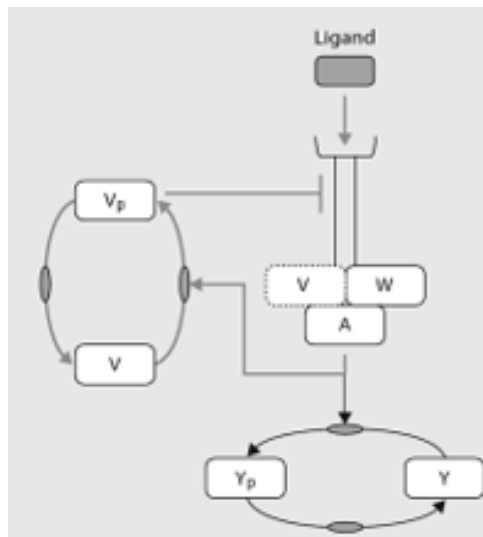


Figure 5: Chemotaxis network for *B.subtilis* (Rao and Ordal 2009).

So Barkai and Leibler's dynamical mechanistic explanation does not apply for the case of *B.subtilis*. But Yi and colleague's organizational design strategy does. For *B.subtilis*, like *E.coli*, exhibits robust perfect adaptation for chemotaxis if and only if it satisfies the characteristic equation for integral feedback control.

6 Toward Abstractive Mechanistic Explanation and its Affordances

Systems biological strategies for explaining the robust perfect adaptation of bacterial chemotaxis (in *E.coli*, *B.subtilis*, etc) apply mathematical techniques to network models. Dynamical, topological, and design strategies apply different techniques to explain the same phenomenon. Each explanatory strategy, moreover, applies its mathematical techniques to network models that embody different kinds and severities of these abstractions such as aggregations, generalizations, structurations. These abstraction types, accordingly, help to explain how these systems biological explanatory strategies differ from each other.

These abstraction types also provide a foundation for unifying various explanatory strategies from systems biology under the banner of mechanistic explanation. Let's consider well known kinds of mechanistic explanation as *standard*. Let's also follow Bechtel and Abrahamsen (2010) by considering *dynamical* mechanistic explanation as a mathematized species of standard mechanistic explanation.

Then let an *abstract network* be any network representation obtained by aggregating, generalizing, or structuring mechanistic details of the sort familiar in standard mechanistic explanation. Also let an *abstractive* mechanistic explanation be any explanation driven by applying mathematical techniques to an abstract network. See Figure 6.

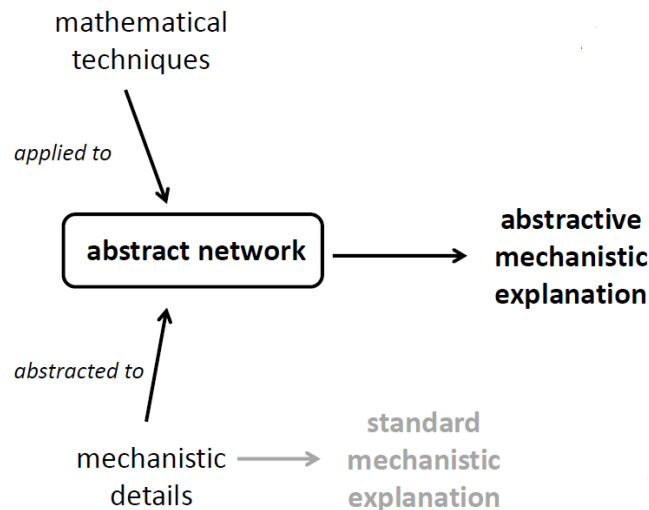


Figure 6. Relating standard and abstractive mechanistic explanation.

Then topological and organizational design explanatory strategies are mechanistic strategies—albeit abstractive ones. Topological explanations apply topological analysis

to aggregated and generalized mechanism networks. Organizational design explanations apply control systems engineering to aggregated, generalized, and structured mechanism networks.

Both kinds of explanation are mechanistic, by virtue of being grounded upon mechanistic details. But both also provide explanatory affordances unavailable through standard mechanistic explanations, by virtue of being abstract. For example, by virtue of using generalizations, topological explanations should have a greater scope than their standard mechanistic counterparts. By virtue of using generalizations and structurations, organizational design explanations should have still greater scope.

That these abstractive mechanistic strategies use novel mathematical techniques is a side effect of their using novel abstractions (in comparison with standard mechanistic explanations and their dynamical cousins). These techniques, of course, support more general conclusions, with wider scope, than the kind of differential equation analysis available for dynamical mechanistic explanations. But the techniques do not explain why the strategies have broader scope.

References

- U.Alon, M.G.Surette, N.Barkai, and S.Leibler, "Robustness in Bacterial Chemotaxis," *Nature* 397 (2009), 168-171.
- N.Barkai and S. Leibler. "Robustness in simple biochemical networks," *Nature* 387 (1997), 913-917.
- W.Bechtel and A.Abrahamsen, "Explanation: A Mechanistic Alternative," *Studies in History and Philosophy of Biological and Biomedical Science* 36 (2005), 421-441.
- W.Bechtel and A.Abrahamsen, "Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science," *Studies in History and Philosophy of Science* 41 (2010), 321-333.
- H.C.Berg, *E.coli in Motion* (Springer, 2003).
- P.A. Braillard, "Systems Biology and the Mechanistic Framework," *History and Philosophy of the Life Sciences* 32.1 (2010), 43-62.
- I.Brigandt, "Systems Biology and the Integration of Mechanistic Explanation and Mathematical Explanation," *Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (2013), 477-492.
- I.Brigandt, S.Green, and M.O'Malley, "Systems Biology and Mechanistic Explanation," in S. Glennan and P. Illari (eds.), *The Routledge Handbook of Mechanisms and Mechanical Philosophy* (forthcoming).
- N.Cartwright, "Capacities and Abstraction," in P. Kitcher and W. Salmon (eds.), *Scientific Explanation* (University of Minnesota Press, 1989), 349-356.
- C.F.Craver, "When Mechanistic Models Explain," *Synthese* 153 (2006), 355-376.

- S.Green, "Revisiting Generality in Biology: Systems Biology and the Quest for Design Principles," *Biology and Philosophy* 30.5 (2015), 629-652.
- S.Green and N.Jones, "Constraint-Based Reasoning for Search and Explanation: Strategies for Understanding Variation and Patterns in Biology," *dialectica* 70.3 (2006), 343-374.
- S.Green and O.Wolkenhauer, "Tracing Organizing Principles: Learning from the History of Systems Biology," *History and Philosophy of the Life Sciences* 35 (2013), 553-576.
- S. Haslinger, "What is a (Social) Structural Explanation?" *Philosophical Studies* 173 (2016), 113-130.
- P.Huneman, "Topological Explanation and Robustness in Biological Sciences," *Synthese* 177 (2010), 213-245.
- N.Jones, "Bowtie Structures, Pathway Diagrams, and Topological Explanation," *Erkenntnis* 79.5 (2014), 1135-1155.
- N.Jones and O.Wolkenhauer, "Diagrams as Locality Aids for Search and Explanation in Molecular Cell Biology," *Biology and Philosophy* 27 (2012), 1135-1155.
- M.Kollman, L.Løvdo, K.Bartholome, J.Timmer, and V.Sourjik, "Design Principles of a Bacterial Signalling Network," *Nature Letters* 438.24 (2005), 504-507.
- A.Levy, "Three New Kinds of Mechanism," *Biology and Philosophy* 28.1 (2013), 99-114.
- A.Levy and W.Bechtel, "Abstraction and the Organization of Mechanisms," *Philosophy of Science* 80 (2013), 241-261.
- W.Ma, A.Trusina, H.El-Samad, W.A.Lim, and C.Tang, "Defining Network Topologies that Can Achieve Biochemical Adaptation," *Cell* 138 (2009), 760-773.
- P.Machamer, Lindley Darden, and C.F.Craver, "Thinking about Mechanisms," *Philosophy of Science* 67 (2000), 1-25.
- D.Matthiessen, "Mechanistic Explanation in Systems Biology: Cellular Networks," *British Journal for Philosophy of Science* (forthcoming).
- D.J.Nicholson, "The Concept of Mechanism in Biology," *Studies in History and Philosophy of Biological and Biomedical Sciences* 43.1 (2012), 152-163.
- S.A.G.Ordorica, "The Explanatory Role of Abstraction Processes in Models: The Case of Aggregations," *Studies in History and Philosophy of Science* (2015), 161-167.
- C. Pincock, "Abstract Explanations in Science," *British Journal for the Philosophy of Science* 66 (2015), 857-878
- J.Raerinne, "Robustness and Sensitivity of Biological Models," *Philosophical Studies* 166 (2013), 285-303.
- C.V.Rao and G.W.Ordal, "The Molecular Basis of Excitation and Adaptation during Chemotactic Sensory Transduction in Bacteria," in M. Collin and R. Schuch (eds.), *Bacterial Sensing and Signaling* (Karger: Basel, Switzerland, 2009), 33-64.
- C.Rice, "Moving Beyond Causes: Optimality Models and Scientific Explanation," *Nous* 49 (2015), 589-615.

- J.L.Spudich and D.E.Kochland, "Non-Genetic Individuality: Chance in the Single Cell," *Nature* 262 (1976), 467-471.
- G.H.Wadhams and J.P.Armitage, "Making Sense of It All: Bacterial Chemotaxis," *Nature Reviews: Molecular Cell Biology* 5 (2004), 1024-1037.
- A.G.Wouters, "Design Explanation: Determining the Constraints on What Can Be Alive," *Erkenntnis* 67 (2007), 65-80.
- T.-M.Yi, Y.Huang, M.I.Simon, and J.Doyle, "Robust Perfect Adaptation in Bacterial Chemotaxis through Integral Feedback Control," *PNAS* 97 (2000), 4649-4653.