

## Understanding non-modular functionality – lessons from genetic algorithms

Jaakko Kuorikoski & Samuli Pöyhönen  
University of Helsinki

Draft, March 2012.  
Forthcoming in PSA 2012 proceedings

### Abstract

Evolution is often characterized as a tinkerer that creates efficient but messy solutions to problems. We analyze the nature of the problems that arise when we try to explain and understand cognitive phenomena created by this haphazard design process. We present a theory of explanation and understanding and apply it to a case problem – solutions generated by genetic algorithms. By analyzing the nature of solutions that genetic algorithms present to computational problems, we show that the reason for why evolutionary designs are often hard to understand is that they often exhibit non-modular functionality, and that breaches of modularity wreak havoc on our strategies of causal and constitutive explanation.

### 1 Introduction

The once dominant classical paradigm of cognitive science has been under attack for several decades. Connectionism, cognitive neuroscience, dynamical systems theory, and new robotics have all questioned whether the classical AI approach to cognition can credibly describe biologically evolved cognitive systems such as human minds. Whereas classical AI tends to approach computational problems with functional decompositions inspired directly by the programmer's intuitions about possible efficient subroutines, the alternative research programs often emphasize that biological evolution is more likely to produce far more complex and messy designs.

In our paper we analyze the nature of the problems that these messy solutions raise to the understanding of cognitive phenomena. In general, the problem of understanding non-intuitive designs produced by natural selection is well-known in philosophy of psychology (e.g., Clark 1997, Ch. 5), philosophy of biology (Wimsatt 2007), and now even in popular psychology (Marcus 2008), but the problem has proved to be difficult to articulate without a clear idea of what exactly it is that evolutionary tinkering is supposed to hinder. The main challenge for understanding is often framed and explained by pointing to the path-dependent nature and the resulting unfamiliarity of the evolved design (Jacob 1977). We argue that this is not the whole story. We hope that providing an explicit theory of explanation and understanding will move us beyond intuitions towards a more

systematic analysis and, ultimately, concrete solutions. We also combine our theory of explanatory understanding to a computational application of evolutionary design: problem-solutions generated by genetic algorithms. By analyzing the nature of solutions that genetic algorithms offer to computational problems, we suggest that an important reason for why evolutionary designs are often hard to understand is that they can exhibit *non-modular functionality*, and that breaches of modularity wreak havoc on our strategies of causal and constitutive explanation.

## 2 Explanation and understanding

The ultimate goal of cognitive neurosciences is to provide mechanistic (constitutive) understanding of system-level properties of the cognitive system in terms of the properties of its parts and their organization. Probably the most developed account of general strategies for reaching such mechanistic understanding is William Bechtel's and Robert Richardson's (2010) study of the *heuristics of decomposition and localization* (DL). The DL procedure goes roughly as follows. First, the different phenomena that the system of interest exhibits are differentiated. Then the phenomenon of interest is *functionally decomposed* in the sense of being analyzed into a set of possible component operations that would be sufficient to produce the phenomenon. One can think of this step as a formulation of a preliminary set of simple functions that taken together would constitute the more complex input-output relation (the system-level phenomenon). The system is also *structurally decomposed* into a set of component parts. The final step is to try to *localize* the component operations by mapping the operations onto appropriate structural component parts. The idea is thus to first think what kinds of more basic properties or behaviors could, taken together, possibly result in the *explanandum* behavior, and then try to find out whether the system is in fact made of such entities that can perform the required tasks. If this cannot be done, the fault may lie with the functional and structural decompositions or with the very identification of the phenomenon, and these may then have to be rethought. The identification and decomposition procedures will in the beginning be guided by earlier theories and common sense, but empirical evidence can always suggest that a thorough reworking of the basic ontology and the form of the possible *explananda* may be in order.

According to Bechtel and Richardson, decomposability is a regulative ideal in such model construction because complex systems are psychologically unmanageable for humans. Decomposition allows the explanatory task to be divided into parts that are manageable for cognitively limited beings, thereby rendering the system intelligible (Bechtel and Richardson 2010). The idea comes originally from Herbert Simon (1962), who claimed that the property of *near-*

*decomposability* is a necessary condition of understandability to any finite cognitive agent. Near-decomposability means that the system can be decomposed into parts in such a way that the intrinsic causal properties of the parts are more important for the behavior of the system than the relational causal properties of the components that are constituted also by their environment and interaction. Near-decomposable systems are thus hierarchical in the sense that the complex whole can be seen as made from a limited set of simpler parts and interactions. Hierarchical systems are manageable for cognitively limited beings because their 'complete description' includes recurring or irrelevant elements describing recurring parts and non-important interactions, the removal of which does not hamper our understanding of the system.

Although there are a number of arguments that conclusively show that such informational economy (redundancy) by itself is not *constitutive* of understanding<sup>1</sup>, we agree with Simon in that a property closely related to near-decomposability, namely *modularity*, is a necessary condition for understanding. As a conceptual starting point for our argument, we follow Jaakko Kuorikoski and Petri Ylikoski in conceiving understanding not as a special or privileged mental state or act, but as a regulative label attributed according to manifest abilities in action and correctness of reasoning. Understanding is a public, behavioral concept. Cognitive processes (comprehension) taking place in the privacy of individual minds are a *causal* prerequisite for possible fulfillment of these criteria, but the processes themselves are not the facts in virtue of which somebody understands or not. They are not the criteria of understanding in the sense that we would have to know them in order to say whether somebody *really* understands something. (Ylikoski 2009; Ylikoski and Kuorikoski 2010)

We take the criterion of understanding to be inferential performance: whether someone understands a concept is evaluated according to whether he or she can make the right inferential connections to other concepts. Likewise, whether someone understands a phenomenon is assessed based on whether he or she can make correct inferences related to it. This view can be further developed by linking it to James Woodward's account of scientific explanation in the following way: Woodward's theory of explanation tells us more specifically what *kinds of inferences* are constitutive of specifically explanatory understanding. According to Woodward (2003), explanation consists in exhibiting functional dependency relations between variables. Knowledge of explanatory

---

<sup>1</sup> First, nobody has actually succeeded in giving a positive argument for equating understanding with increased informational economy (Barnes 1992). Second, successful classification schemes compress information by facilitating inferences to properties probably possessed by individuals on the basis of belonging to a certain known class. However, classification schemes by themselves are usually taken to be merely descriptive and not explanatory. The same general point can be drawn from standard statistical procedures, which by themselves only summarize the data, but do not explain it. (Woodward 2003, 362-364).

relationships facilitates understanding by implying answers to what-if-things-had-been-different questions concerning the consequences of counterfactual or hypothetical changes in the values of the *explanans* variable. Whether someone understands a phenomenon is evaluated according to whether he or she can make inferences not only about its actual state, but also about possible states of the phenomenon or system in question. In the case of causal explanations, these explanatory dependencies concern the effects of interventions and knowledge of causal dependencies thus enables the possessor of this knowledge to act and possibly manipulate the object of explanation. These answers are the basis of the inferential performance constitutive of understanding.

Of course, the limits of inferential performance depend causally on contingencies related to the reasoning processes of the agents whose understanding is being evaluated. Thus the limits of understanding are partly dependent on the cognitive make-up of agents and can certainly be investigated psychologically. For example, if the space limit of our working memory is indeed roughly seven items, then this constitutes an upper boundary for the complexity of our inferences and, consequently, for our understanding.

In order for answers to what-if questions to be well defined, the dependencies grounding the answers have to possess some form and degree of independence in that *a local change in an aspect of the phenomenon under study cannot ramify uncontrollably or intractably*. If local changes in a part of a system disrupt other parts (dependencies) in a way that is not explicitly specified (endogenized) in the (internal or external) representation of the system according to which inferences are made, the consequences of these changes are impossible to predict and counterfactual assertions impossible to evaluate. Things participating in the dependency relations also have to be somewhat localized (physically and/or conceptually) in order for the contemplated changes to be well defined in the first place. (Woodward 2003, 333.) Therefore a necessary condition for a representation to provide understanding of a phenomenon is that the modularity in the representation matches the modularity in the phenomenon.

Let us first discuss the case of causal understanding. If an intervention on a causal system actually changes the system in a way that is not represented in the model of the system, the model as it stands does not give correct answers to what-if-things-had-been-different questions concerning the state of the system after the intervention. If we intervened on a causal input corresponding to variable  $X_i$  in a model and the intervention, no matter how surgical, also changed the dependencies within the system or other variables themselves affecting variables causally downstream of  $X_i$ , the model would give incorrect predictions about the consequences of the intervention. Hence, the model would not provide correct causal understanding of the workings of the system and the causal

role of the variable in it. If the system cannot be correctly modeled on any level of description or decomposition so that it is modular in such a way - if the system itself is not causally modular - no what-if-things-had-been-different-questions concerning interventions in the system can be answered and there is no causal understanding of the system to be had. If the system is in fact such that every local change brings about intractable changes elsewhere in the system to such an extent that there can be no representation that would enable a cognitively finite being to track these changes and make correct inferences about their consequences, then the system is beyond the limits of understanding.

The problem of understanding causally non-modular systems has received some attention in the philosophy of science literature (e.g. Bechtel and Richardson 2010, Ch. 9). However, according to the schema of Bechtel and Richardson, before we can even start thinking about the causal-mechanical understanding of the system realizing the complex behavior to be understood, we need to formulate hypotheses about the possible functional decompositions of the behavior (see also Cummins 1983). For example, what kind of simpler subtasks could possibly produce complex cognitive capacities such as language production and comprehension, long-term memory, and three-dimensional vision? Importantly, this task is separate, though not independent, from hypotheses concerning the implementation of the capacity. Although the understanding offered by the functional decomposition is not strictly speaking causal – component operations do not cause the whole behavior to have certain properties because they are not causal precedents but constitutive parts of the complex whole<sup>2</sup> – the modularity constraint on understandability still applies in the following way. We can only understand the complex behavior by having knowledge of the component operations if we can make reliable what-if inferences concerning the possible consequences of changes in the component operations for the properties of the *explanandum* capacity. We provisionally understand working memory if we can infer from possible changes in its hypothesized component operations (such as differences in the postulated phonological loop or episodic buffer) to changes in the properties of the capacity. These inferences are only possible if the functional decomposition itself is suitably modular, i.e., the consequences of “local” changes in component operations do not ramify in an intractable way making the behavior of the whole completely holistic. We now argue that genetic algorithms demonstrate that design by selection can lead to such non-modular complex behavior.


---

<sup>2</sup> Although we fully agree with Piccinini and Craver (2011) in that insofar as functional decompositions are explanatory, they are to be thought of as mechanism sketches and that the functional hypotheses are not independent of the question of mechanistic implementation.

### 3 Genetic algorithms

Since the 1960s, there have been attempts to apply insights from evolutionary thinking to computer programming. Here we discuss one genre of evolutionary programming: genetic algorithms (cf. Holland 1975; Goldberg 1989; Mitchell 1996). In a nutshell, the idea of the genetic-algorithms approach is to “breed” randomly generated solutions to computational problems. This is done by mimicking the evolutionary mechanisms of inheritance, mutation, selection and crossover in a computer simulation. Although genetic algorithms (henceforth GAs) are not the only strand of evolutionary programming, they serve our purpose well because their basic principles are easy to understand and they are the most well-known kind of evolutionary programming outside computer science (Clark 1997, 2001; Mitchell 2009).

From the point of view of AI, genetic algorithms are a form of non-exhaustive but massively parallel search in the search space of a problem. They can be used for a number of different purposes: for evolving behavioral strategies for simulated agents, for finding weights for a connectionist network, or for evolving cellular automata to perform computations. We illustrate the nature of GAs by presenting a simple example from Melanie Mitchell (2009, ch. 9). Mitchell’s original simulation showed how GAs can be used to evolve a controlling program for a simulated robot picking up soda cans in a 10x10 grid. Robby the robot can only see squares that are adjacent to its location (center, North, South, East, West), and each turn it can either move one step to a particular direction, move at random, try to pick up a can, or do nothing. Each simulation run lasts for a predetermined amount of time steps (originally 200), and Robby’s task is to pick up as many randomly situated soda cans as possible.



Strategy G

Genome G:

```
254355153256235251056355461151336154151034156110550150052030256256132252350325112
052333054055231255051336154150665264150266506012264453605631520256431054354632404
350334153250253251352352045150130156213436252353223135051260513356201524514343432
```

Figure 1. (taken from Mitchell 2009, 137). Each “locus” in the genome  $G$  corresponds to one of the possible immediate environmental states of Robby and each digit (the allele) to a move in that situation (e.g. ‘0’  $\rightarrow$  ‘move north’, ‘5’  $\rightarrow$  ‘pick up’).

Initially a random population of software individuals is generated, each with a “genome” consisting of 243 random numbers. Each locus in the genome guides Robby’s behavior in a particular situation (Fig 1). The fitness score of each candidate program in the population is calculated by running several simulation trials: crudely, the more cans the robot is able to pick up by average, the higher its fitness. Programs with the highest fitness scores are then used to form the next generation of programs: they are paired randomly, and the genomes of the two parents are crossed over at a randomly chosen point to create the genomes of new individuals. Finally, for each descendant, there is a small probability (.05) that a mutation occurs in its chromosome. As a result, the new generation is based on the most successful variants among the previous generation and the process loops back to the fitness-calculation phase. Thus the GA continues searching for efficient solutions to the problem by investigating the surrounding areas in the search space.

After a few hundred generations, the evolved strategies start to achieve impressing results in the simulated task. As we replicated Mitchell’s simulation, we observed that after the 800<sup>th</sup> generation, the best strategies among evolved Robbys started to have higher fitness scores than a simple “rational” solution programmed by a human designer (ultimately 480 vs. 440 points). However, although solutions found with GAs are efficient, their behavior is often hard to understand. The ingenious heuristics that the programs employ cannot usually be deciphered by simply looking at individual genes or sets of genes. Instead, looking holistically at the broad phenotypic behavior of the robot is necessary. A nice illustration of this impenetrability of such evolved solutions is the fact that in some cases when a highly evolved Robby is in the same square with a can, it decides not to pick it up, but rather chooses to move away from the square. While this behavior seems *prima facie* irrational, looking at the total behavioral profile of the robot uncovers a cunning strategy: Robby uses cans as markers to remember that there are cans on its side and explores the adjacent squares for extra cans before picking up the marker can. Thus by not treating cans only as targets but also as navigational tools, Robby uses its environment to extend its severely limited visual capacities and to compensate for its total lack of memory.

Moreover, by examining the behavior of a 1500<sup>th</sup> generation Robby with the highest fitness score in its population, it can be seen that the marker strategy manifests in slightly different ways in different

environmental situations. It is therefore not a discrete adaptation, but rather a collection of independently evolved sub-strategies. Furthermore, the marker strategy appears to tightly intertwine with other environment-employing “hacks” that the sophisticated Robby uses: when there is already a lot of empty space on the grid, Robby employs a “vacuum-cleaner” movement strategy. It follows the walls of the board, departing toward the center when it detects a can, employs the marker strategy if possible, and immediately after cleaning up its local environment, returns directly to the south wall to continue its round around the board.

Such kluges are common to designs created by GAs. Like biological evolution, GAs can come up with solutions that a human designer would not usually think of. These solutions often offload parts of problem solving to the environment, and thus rely on a tight coupling between the system and its environment. And as pointed out by Clark (1997, 2001), recurrent circuitry and complex feedback loops between different levels of processing often feature in systems designed by GAs. Such designs are often difficult to understand. We claim that such difficulties in understanding are often created by the lack of modularity in the functional decomposition of the behavior. This point can be illustrated by looking again at the genome of our most successful Robby (genome G in Fig 1). Robby is leaving cans as markers only in specific situations and only the totality of this selective marking strategy, together with navigational strategies utilizing cans and walls, constitutes the effectiveness of the search procedure. Looking at isolated genes in Robby’s genome only reveals trivially modular elements corresponding to elementary subtasks in Robby’s behavior: one gene corresponds to an elementary move in a specific environmental situation. But we cannot make inferences from local hypothetical changes in these elemental behaviors to consequent effects on fitness. The connection between any single elementary behavioral rule and the strategy is simply too complex and context dependent. A change in a single rule (in situation B and a can present, whether to pick or not to pick the can up) has consequences for the effectiveness of the other elementary behavioral rules constituting the navigational strategy. Such inferences would require an extra “level” of modular sub-operations between the individual movements and the strategy as a whole. The marker and vacuum cleaner strategies mentioned above are examples of such middle-level sub-operations, but they are by themselves insufficient to yield understanding of the whole behavior of our most successful Robby, since the effectiveness of leaving a can is a result of the evolved match between the specific situations in which Robby leaves a can and the rest of the navigation behavior. And genetic algorithms do not, in general, produce such easily discernible designs. Rather, only by simultaneously looking at constellations of different genes, and eventually the whole genome, the interesting heuristics in the system’s behavior can be revealed - if at all.



To recapitulate, our example exhibits several distinct (yet related) challenges to understanding:

1. The discernible middle-level strategies (marker, vacuum-cleaner) do not have a dedicated structural basis. Instead, the nature of the design process leaves all atomic structural elements (the 243 dna elements) open for exploitation by all capacities serving the main goal. In consequence, the system is not structurally nearly-decomposable, but instead has "a flat hierarchy." Strategies are implemented in highly distributed structures, and as pointed out in section 2, this raises a challenge for human cognitive capacities.
2. Challenge 1 above means that the interactions between subtasks tend to be strong: a change in one subtask responsible for the marker-behavior affects also the functioning of the vacuum-cleaner navigation. In general the middle-level strategies can only be discerned and defined in a very abstract way and the interaction-effect in their contribution to the overall fitness is so large as to make any inferences about the consequence of partial changes in one strategy next to impossible.
3. The way in which operations contribute to the fitness of the individual is highly context-dependent and depends on the properties of the environment as well as the dna of the agent. For instance, merely detecting the existence of the marker strategy requires that there are suitable clusters of cans in the environment. Moreover, even small modifications the environment can lead to drastic changes in the performance of a strategy. For instance, adding only a few randomly placed extra walls on the grid radically collapses the average score of the successful Robby described above.

Extrapolating from this very simple case, GAs may yield functional decompositions of the problem that do not follow a tidy hierarchical decomposition into modular subtasks, whose individual contributions would be easy to understand (how a change in a sub-routine would affect the behavior of the mother-task). Instead, feedback, many tasks using same subtasks as resources, and environment couplings lead to holistic design where almost "everything is relevant for everything". The evolved functional architecture is flat in that there are few discernible levels of order between the elementary operations and the complex whole. The counter-intuitiveness of such flat architectures is apparent in the deep mistrust faced by connectionist suggestions for non-hierarchical design of cognitive capacities (see e.g. Rumelhart and McClelland 1986 vs. Pinker and Prince 1988).

Furthermore, GAs underscore the path dependence of evolutionary problem solving. For sufficiently complex computational problems there are often several local maxima in the fitness landscape of the problem, and on different runs of the simulation the population can converge to different maxima. The functional decomposition that a human designer comes up with is just one possible solution among several others. Perhaps our biological evolution actually ended up with a radically different one.

#### **4 Lessons for the study of mind**

Genetic algorithms seem to demonstrate that evolution can in principle lead to non-modular functionality. This imposes a limit on our ability to understand such behavior: if we cannot trace the consequences of changes in the sub-operations, we cannot answer what-if questions concerning the complex behavior. Such behavior also constitutes a thorny problem for mechanistic understanding of the implementation of the said behavioral capacities, since the DL heuristic cannot even get off the ground. We can now ask two questions: should we expect to find such non-modular functionality in nature, especially in human cognition, and if so, what attitude should we adopt with respect to this problem. Should the aim of causal-mechanistic understanding of the brain be given up and replaced with a program of instrumentally interpreted dynamical models and modeling the dynamics of the mind with a few macro-variables?

There are important disanalogies between GAs and biological evolution. (1) in GAs, there usually is no genotype–phenotype distinction. In biological evolution, however, genes do not directly cause properties of the phenotype, but rather participate in guiding ontogenesis. There have been suggestions that ontogenesis itself favors modular design. GAs may also seem a problematic platform for exploring the possibilities of DL heuristics, since the lowest level of functional organization and the level of implementation are the same (i.e., the genome). However, we see no reasons why this would affect our argument. Moreover, the argument developed here is about selection in general, and failures of functional modularity may in principle also arise in the course of development – at least if the idea of neuronal group selection or “neural Darwinism” is taken seriously.

(2) Most studies on genetic algorithms are carried out by using a single fixed goal or a fixed task type. In the Robby example, although the distribution of the cans was generated at random, the task

itself remained essentially the same from generation to generation. However, Nadav Kashtan and Uri Alon (2005, see also Kashtan et al. 2007) have demonstrated that when the goals themselves are composed of modularly varying sub-goals, evolution produces modular functionality. It is easy to see why this is the case. If the tasks to which the system has to adapt to remains the same, the selection environment is stable and the peaks in the fitness landscape are immovable, then selection favors strategies which offload problem solving to that particular environment as much as possible. But if the task itself is composed of changing subtasks, it makes sense to design the adaptive response in such a way that a particular sub-operation can locally adapt to a local change in a subtask without altering the totality of the otherwise well functioning behavior.

It seems likely that cognition has evolved in such a modularly changing selection environment, but the extent to which we should expect to find modular functionality in cognition is hard to estimate and most probably a purely empirical matter. Moreover, as a response to Simon's (1962) *Tempus and Hora* argument, it has been argued that componential specialization in complex systems is a force that works against the development of strictly modular structures (e.g., Levins 1973, Wimsatt 2007, 186-192). Therefore, the extent to which we should expect to find modular functionality in human cognition is hard to estimate and most probably a purely empirical matter. Nonetheless, these arguments as such give us no reason to believe that the produced functional decomposition should respect any intuitive constraints, such as those derived from introspection on our thought processes or the way in which we would program a strategy to tackle similar cognitive challenges.

Genetic algorithms demonstrate that evolution can create designs which are in principle beyond the understanding of unaided cognitive beings such as us. Yet there is nothing mysterious in such designs. Simon pondered whether the relative abundance of hierarchical nearly decomposable complexity was due to our selective attention to precisely such systems, but we believe this to be a somewhat hasty conjecture. We have no trouble finding and delineating systems, such as Robby or possibly ourselves, with behaviors which are functionally non-decomposable and constituted by a flat architecture. However, there certainly might be a psychological bias that makes us see hierarchical design also where there is none. One way of coping with this impasse is to realize that there are no fundamental reasons to limit the relevant understanding *epistemic* agent to be an *unaided* human. Although only a human agent can experience a sense of understanding, this feeling should not be confused with understanding itself. Therefore brute computational approaches can produce understanding as long as the understanding subject, the cognitive unit whose inferential abilities are to be evaluated, is conceived as the human-computer pair.



## References

- Barnes, Eric. 1992. Explanatory Unification and Scientific Understanding. *PSA 1992*, 3-12.
- Bechtel, William and Robert C. Richardson. 2010. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, The MIT Press.
- Clark, Andy. 1997. *Being There: Putting Brain, Body and World Together Again*. The MIT Press.
- Clark, Andy. 2001. *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford University Press.
- Cummins, Robert. (1983). *The Nature of Psychological Explanation*. Cambridge: MIT Press.
- Goldberg, David E. 1989. *Genetic Algorithms in Search Optimization and Machine Learning*. Addison Wesley.
- Holland, J. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- Jacob, Francois. 1977. Evolution and Tinkering. *Science* 196 (4295): 1161-1166.
- Kashtan, Nadav and Uri Alon. 2005. Spontaneous evolution of modularity and network motifs. *PNAS* 102 (39), 13773-13778.
- Kashtan, Nadav, Elad Noor and Uri Alon. 2007. Varying environments can speed up evolution. *PNAS* 104 (34), 13711-13716.
- Levins, Richard. 1973. The Limits of Complexity. Pattee, H. (ed.) *Hierarchy Theory: The Challenge of Complex Systems*. London: Braziller: 73-88.
- Marcus, Gary. 2008. *Kluge: The Haphazard Construction of the Human Mind*. Boston and New York: Houghton Mifflin.
- Mitchell, Melanie. 1996. *An Introduction to Genetic Algorithms*. Cambridge: MIT Press.
- Mitchell, Melanie. 2009. *Complexity. A guided tour*. Oxford: Oxford University Press.
- Piccinini, Gualtiero and Carl Craver. 2011. Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches. *Synthese* 183 (3):283-311.
-

Pinker, S. and Prince, A. 1988. On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition. *Cognition* 23: 73–193.

Rumelhart, D. and McClelland J. 1986. On Learning the Past Tenses of English Verbs. in McClelland and Rumelhart *et al.* 1986. *Parallel Distributed Processing*, vol. I, Cambridge, Mass.: MIT Press: 216–271.

---

Simon, Herbert. 1962. The Architecture of Complexity. *Proceedings of the American Philosophical Society* 106. 476-482.

Wimsatt, William. 2007. *Re-Engineering Philosophy for Limited Beings*. Cambridge MA: Harvard UP.

Woodward, James 2003. *Making Things Happen*. Oxford University Press.

Ylikoski, Petri. 2009. The Illusion of Depth of Understanding in Science. in *Scientific Understanding: Philosophical Perspectives* (edited by H. De Regt, S. Leonelli & K. Eigner), Pittsburgh University Press: 100-119.

---

Ylikoski, Petri and Jaakko Kuorikoski. 2010. Dissecting Explanatory Power. *Philosophical Studies* 148, 201-219.