



Contents lists available at SciVerse ScienceDirect

# Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: [www.elsevier.com/locate/shpsc](http://www.elsevier.com/locate/shpsc)

## The mismeasure of machine: Synthetic biology and the trouble with engineering metaphors

Maarten Boudry<sup>a</sup>, Massimo Pigliucci<sup>b</sup><sup>a</sup> Ghent University, Department of Philosophy and Moral Sciences, Belgium<sup>b</sup> City University of New York, Graduate Center, Philosophy Program, United States

## ARTICLE INFO

Article history:  
Available online xxx

Keywords:  
Synthetic biology  
Adaptationism  
Reverse engineering  
Organism-machine metaphor  
Analogical thinking

## ABSTRACT

The scientific study of living organisms is permeated by machine and design metaphors. Genes are thought of as the “blueprint” of an organism, organisms are “reverse engineered” to discover their functionality, and living cells are compared to biochemical factories, complete with assembly lines, transport systems, messenger circuits, etc. Although the notion of design is indispensable to think about adaptations, and engineering analogies have considerable heuristic value (e.g., optimality assumptions), we argue they are limited in several important respects. In particular, the analogy with human-made machines falters when we move down to the level of molecular biology and genetics. Living organisms are far more messy and less transparent than human-made machines. Notoriously, evolution is an opportunistic tinkerer, blindly stumbling on “designs” that no sensible engineer would come up with. Despite impressive technological innovation, the prospect of artificially designing new life forms from scratch has proven more difficult than the superficial analogy with “programming” the right “software” would suggest. The idea of applying straightforward engineering approaches to living systems and their genomes— isolating functional components, designing new parts from scratch, recombining and assembling them into novel life forms—pushes the analogy with human artifacts beyond its limits. In the absence of a one-to-one correspondence between genotype and phenotype, there is no straightforward way to implement novel biological functions and design new life forms. Both the developmental complexity of gene expression and the multifarious interactions of genes and environments are serious obstacles for “engineering” a particular phenotype. The problem of reverse-engineering a desired phenotype to its genetic “instructions” is probably intractable for any but the most simple phenotypes. Recent developments in the field of bio-engineering and synthetic biology reflect these limitations. Instead of genetically engineering a desired trait from scratch, as the machine/engineering metaphor promises, researchers are making greater strides by co-opting natural selection to “search” for a suitable genotype, or by borrowing and recombining genetic material from extant life forms.

© 2013 Elsevier Ltd. All rights reserved.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Biological and Biomedical Sciences*

*We improve our favourite plants and animals—and how few they are—gradually by selective breeding; now a new and better peach, now a seedless grape, now a sweeter and larger flower, now a more convenient breed of cattle. We improve them gradually, because our ideals are vague and tentative, and our knowledge is very limited; because Nature, too, is shy and slow in our clumsy hands. Some day all this will be better organized, and still better. That is the drift of the current in spite of the eddies. The whole world will*

*be intelligent, educated, and co-operating; things will move faster and faster towards the subjugation of Nature. In the end, wisely and carefully we shall readjust the balance of animal and vegetable life to suit our human needs.* (H.G. Wells, *The Time Machine*).

### 1. Introduction

Ever since Descartes (1637/2010), the idea that living beings are a type of complex machine<sup>1</sup> has fascinated philosophers and scien-

E-mail addresses: [maartenboudry@gmail.com](mailto:maartenboudry@gmail.com) (M. Boudry), [massimo@platofootnote.org](mailto:massimo@platofootnote.org) (M. Pigliucci)

<sup>1</sup> Initially, of course, with the partial exception of humans, when it came to consciousness.

tists alike, spurring much research into physiology and molecular biology. It has also engendered much misunderstanding among the general public, because of the popular appeal of the concept of “Intelligent Design” proposed by a number of creationists (Dembski, 2002; Pigliucci, 2002). A large section especially of the American public makes the intuitive (if logically flawed) inference that, if organisms resemble machines, there must have been an engineer who put them together in the first place.

Contemporary biology too is steeped in engineering metaphors (Pigliucci, 2010; Pigliucci & Boudry, 2011): genes are said to carry information that constitutes a “blueprint” for the organism; cells and sub-cellular organelles are talked about in terms of factories full of molecular machines. When the human genome project got started, the hope was to be able to have the information necessary to “read off” a human being stored on a CD or hard drive, thereby allowing biologists to pinpoint with precision where to intervene to cure a number of diseases. But the view of living systems as machines has deeper roots. Ever since the program of natural theology in the 18th and 19th centuries, which made a deep impression on Darwin, discussions of biological functionality have been wedded to the notion of design. Many biologists would argue that the latter is indispensable in their discipline, and that many biological systems are most simply viewed as design solutions to specific adaptive problems.

Scientists are often attracted by the possibility of engineering and “reverse engineering” living organisms (Dennett, 1996; Pinker, 1997), both for practical purposes (Deplazes-Zemp, 2011) and in order to better understand life itself (e.g., Carr & Church, 2009). Indeed, the strategy of “reverse engineering” is a pivotal principle shaping the program of adaptationism in biology. Though engineering analogies have some heuristic value and applicability, we will see that they are also misleading and break down on the level of molecular biology and biochemistry. The recent emergence of “synthetic” biology has upped the ante of the engineering approach in biology, so to speak. The machine metaphor is a guiding principle of the basic idea behind synthetic biology, and it is a logical extension of the adaptationist credo of reverse engineering: again, living organisms are intricate machines (though we now know that they are the result of evolution), and as such we can understand them by taking them apart, isolating their various functional components, and then using rigorous engineering principles to put together a novel life form.

## 2. David Hume and the limits of the machine metaphor

This ‘engineering’ approach in biology, based on analogies with man-made artifacts, is in tension with a different tradition in both philosophy and biology (and which, ironically, is also entwined with the debate over intelligent design). Hume (1779/2012) first cautioned against analogies between the world (and living beings in particular) and products of intelligent engineering.<sup>2</sup> Hume proposed that the universe, if anything, bears more resemblance to an animal or a vegetable. Here is how his stand-in character Philo developed the analogy:

A continual circulation of matter in [the universe] produces no disorder; a continual waste in every part is incessantly repaired: The closest sympathy is perceived throughout the entire system: And each part or member, in performing its proper offices,

operates both to its own preservation and to that of the whole. The world, therefore, I infer, is an animal. (Hume, 1779/2012, p. 39)

Philo further speculates that the world even more resembles a plant, and that it could have come into existence by a process analogous to reproduction or vegetation. Although this still leaves us with the question of what initiated the process of vegetation, Hume’s main concern lies with the limits of analogical reasoning: if we single out some features of the world while ignoring others, we can dream up any number of facile analogies: “[I]n such questions as the present, a hundred contradictory views may preserve a kind of imperfect analogy, and invention has here the full scope to exert itself” (Hume, 1779/2012, p. 49). And if an analogy is not perfect, any conclusions transferred from the source (artificial design) to the target (organisms) are premature.

Hume’s point about engineering analogies as applied to the universe as a whole logically extends to individual living organisms (indeed, the latter issue is subsumed under the former). In the case of living beings, Hume’s critique had only partial force because he was missing a powerful alternative explanation for their apparent design, an explanation that was offered by Darwin almost a century later. He surely would have appreciated Darwin’s mechanism of evolution, as it vindicates Hume’s point that design inferences on the basis of superficial analogies are premature. Interestingly, although the advent of Darwinian biology permanently excluded intelligent designers from any reasonable form of discourse about living organisms, it still seemed to legitimize design talk about the living world.<sup>3</sup> The biological adaptations that fascinated natural theologians before Darwin turned out to be products of natural selection working slowly over countless generations. As we will see in the next section, the study of adaptation through natural selection is one of two major reasons why the engineering approach championed by natural theology has been inherited by modern structural biology.

The other reason why the machine metaphor survived the expulsion of intelligent designers from science is the success of methodological reductionism in some fields of biology throughout the 20th century, leading to a series of spectacular discoveries in genetics and molecular biology. Treating living organisms as complicated pieces of machinery, with cells as miniature factories and genes as design blueprints, clearly paid off—up to a point. That point became increasingly apparent with the mid-century advent of the Modern Synthesis in evolutionary biology, which successfully integrated gradual evolutionary change induced by natural selection with the idea of discrete Mendelian inheritance, in the process creating a rift in the biological community. On the one hand geneticists (and, eventually, the molecular biologists who followed them) increasingly adopted the above mentioned reductionist and mechanistic stance, which served them well into the 21st century era of genomics and beyond. On the other hand, developmental biology was famously left out of the Modern Synthesis, and along with it all its rich scholarship concerning the embryology of complex organisms, precisely because the new emphasis on reduction, mechanism and adaptation did not mesh well with the messy complexity of most developmental systems. Indeed, much of the impetus behind recent and repeated calls for an Extended Synthesis in evolutionary biology (Carroll, 2008; Pigliucci & Muller, 2010) can be traced back to the lingering dissatisfaction of many (though by no means all) organismal biologists with what they see as an

<sup>2</sup> See William Paley’s famous watch-watchmaker analogy, which was published after Hume, but represented a popular way of thinking about natural theology even in Hume’s time.

<sup>3</sup> By the turn of the 20th century biology and Darwinism itself were facing a crisis (Bowler, 1992), catalyzed by the apparent difficulty of reconciling the idea of gradual evolutionary change induced by natural selection with the discrete inheritance implied by the rediscovery of Mendel’s work. The resolution, of course, was the so-called Modern Synthesis of the 1920s–1940s, a crucial aspect of which was the development of population and quantitative genetic theory (Provine, 2001).

excessively reductive approach that does not fit well with the complex realities of the living world.

### 3. Adaptationism, reductionism & reverse engineering

There is a long history of analogies from the realm of human artifices and design to the study of biological adaptation. In a sense, Darwin modernized teleological reasoning in the science of biology, making it independent of theological speculation, but he certainly did not dispense with the design analogy altogether. In his famous discussion of the eye, he marvels at the “extreme perfection” of this optical “contrivance,” and he notes that it is “scarcely possible to avoid comparing [it] to a telescope” (Darwin, 1859/2011, p. 145). In his influential study of adaptation and natural selection, Williams (1966/1996, p. 10) writes that adaptations in biological systems have a “clear analogy with human implements.” He even uses design terminology as a way of discriminating between real adaptations and mere causal relations:

Whenever I believe that an effect is produced as the function of an adaptation perfected by natural selection to serve that function, I will use terms appropriate to human artifice and conscious design. The designation of something as the means or mechanism for a certain goal or function or purpose will imply that the machinery involved was fashioned by natural selection for the goal attributed to it. (Williams, 1966/1996, p. 9)

The research program of adaptationism takes adaptation as the center of focus in biology, although there has been some confusion (not to say, heated controversy) over what this centrality exactly amounts to (Godfrey-Smith, 2001; Sober, 1998). Evolutionary biologists in the tradition of adaptationism have analyzed all kinds of biological systems, including the human brain, by assuming that costly and conspicuous features have been “designed” by natural selection to solve some adaptive problem, similar to how an engineer would tackle a design task with a preset goal. Notably, the Santa Barbara school of evolutionary psychology (Laland & Brown, 2011, chap. 5; Pigliucci & Kaplan, 2006, chap. 7) treats the human mind as analogous to a Swiss army knife (Cosmides & Tooby, 1994), equipped with a neat set of largely independent “modules,” each evolved for a specific adaptive purpose, and each ready to be analyzed in terms of a well-defined function that evolved during the Pleistocene or at other “evolutionarily relevant” times for the human lineage. This massive modularity thesis of evolutionary psychology, with its view of the human mind as a collection of “isolable function-specific processing systems” (Carruthers, 2006, p. 12), goes hand in hand with the program of methodological reductionism.

One of the guiding principle of adaptationism has been the idea of reverse engineering. Fully in line with the machine analogy, proponents of adaptationism have advocated the use of engineering principles as a methodological strategy for figuring out the function of biological systems. For instance, Steven Pinker, a staunch defender of adaptationism in psychology, has no reservations about adopting engineering approaches to living systems, including the human mind: “The mind, like the Apollo spacecraft, is designed to solve many engineering problems, and thus is packed with high-tech systems each contrived to overcome its own obstacles” (Pinker, 1997, p. 4). Daniel Dennett, another advocate of adaptationism, notes that “we are artifacts, after all, designed by natural selection to provide reliable survival

vehicles for our genes into the indefinite future” (Dennett, 1990, p. 187).

The analogy with human artifices is clear. If we want to figure out the function of an unfamiliar man-made artifice, we have to try to get inside the mind of its maker. To distinguish intended from accidental features, we try to see which possible functionality would make sense of the whole device (Pinker, 1997, p. 21). By assuming reasonable principles of engineering, we can reconstruct the intended purpose of the device. If the designer had a clear goal in mind and proceeded in a rational fashion, being little constrained by external limitations, this strategy will usually be successful. It will be less successful if the designer did a poor job, or if (s)he was (heavily) constrained by material resources and the errors of other designers before him/her. As Daniel Dennett himself admits: “If the reverse engineer can’t assume that there is a good rationale for the features they observe, they can’t even begin their analysis.” (Dennett, 1996, p. 213)<sup>4</sup>

As any student of evolution knows, however, foresight or intentionality are completely alien to natural selection. At most, we can think about the selection pressures that organisms have to cope with, or the adaptive problem that they have to “solve.” This means that the reverse engineering project in biology, or what Dennett termed “artifact hermeneutics” (Dennett, 1990) is plagued by a problem of underdetermination, more so than with man-made artifacts. In the case of the latter the designer is accountable to standards of rational engineering, which facilitates the hermeneutic process, allowing the interpreter to make sense of different components within the system. In the case of living systems, however, there is no prior guarantee that everything will make sense. Gould and Lewontin (1979) have warned about the adaptationist temptation to overreach, ascribing functionality to accidental features, historical left-overs, and useless by-products. It is usually not very difficult to invent a spurious adaptationist explanation for a given biological system. As Williams noted, adaptation is a costly explanation, and should not be resorted to lightly or casually (Williams, 1966/1996).

The design stance has undoubtedly proven useful to biologists, when taken as a heuristic for reasoning about biological functionality. In its modest form, adaptationism does not make a blanket assumption of optimality, but uses optimality considerations as a starting point for research. As Sober writes, “optimality models are important even if they turn out to be false” (Sober, 1998, p. 83). By finding out to what extent biological systems fail to match optimal solutions, measured against some engineering standard, biologists can generate hypotheses about the elements missing from their analysis: maybe the optimal solution is inaccessible in the fitness landscape? Maybe there is a trade-off between different selection pressures, for example the cost of implementing a better solution? Maybe selection pressures were not consistent and strong enough, or optimal solutions were out of reach because of physical constraints? Ernst Mayr, while acknowledging the pitfalls of adaptationist reasoning, praises its heuristic value, stating that the adaptationist question “What is the function of a given structure or organ?” has been the “basis for every advance in physiology” (Mayr, 1983, p. 328), even well before Darwin (and Wallace) hit upon the right explanation of apparent design. In any case, this means that, even to the admission of reasoned adaptationists, engineering analogies have a mainly heuristic value in biology, rather than providing an accurate model of living organisms.

<sup>4</sup> Vaesen & van Amerongen (2008) have argued that Dennett places too much emphasis on optimality considerations, and that, even in the case of human artifacts, reverse engineering can get off the ground even when the artifact is damaged, flawed or bears the mark of physical and other constraints. Archeologists, when confronted with an inscrutable device, as well as laymen in everyday situations, consider plausible intentions on the designer’s part to make sense of the artifact, and do not rely solely on optimality considerations.



#### 4. What is the problem with applying machine metaphors to the living world?

We discuss three (related) limitations impinging on the project of synthetic biology.

(1) First, even when natural selection has succeeded in crafting organs of “extreme perfection,” as Darwin called it, which are optimal in the context of some well-defined adaptive problem, the way evolution achieves this result may not resemble anything that a sensible designer would ever come up with. Evolution is a satisficer and tinkerer (Jacob, 1977), it can never go back to the drawing board and start from scratch, as human engineers are free to do. Not only is it heavily constrained by the material at hand, it always operates to maximize fitness in a specific environment, against local competitors, rather than striving for some general standard of optimality. Evolution will typically co-opt and modify existing structures, resulting in a patched-up network of interlocking and partly redundant components. As a result, it often settles for solutions that are opportunistic, cumbersome and ultimately ill-conceived.

To take the prototypical instance of adaptation, the mammalian eye is surely beautifully designed, and even before William Paley people have marveled at the way the eye muscles and lens achieve focus, the pupils dilate and contract to adjust for light intensity. The different parts of the eye are exquisitely arranged to achieve depth of vision and to adjust for different chromatic aberrations, etc. If we have a closer look at how this has been pulled off, however, we should not be surprised to find unwieldy or even botched design. Pinker has claimed that nature, being red in tooth and claw, does “not forgive slapdash engineering” (Pinker, 1997, p. 167), but surely that claim is not warranted by the biological facts. Is the blind spot on our retina, resulting from a clumsy and perfectly avoidable design flaw (the octopus eye is properly wired, without blind spot), not a prime example of slapdash engineering? The distinguished 19th century physiologist and physicist Herman von Helmholtz wrote that the eye was littered with design flaws, such as the blind spot, astigmatism (a failure to focus on all lines of an object) chromatic aberrations (a failure to focus on all colors), the cornea’s imperfect transparency, the shadow image of our veins, etc. If such a design were submitted by an engineer, thought Helmholtz, he would send it back immediately, “blaming his carelessness in the strongest terms” (Helmholtz, Cahan, & Williams, 1995, p. 181):

For the eye has every possible defect that can be found in an optical instrument, and even some which are peculiar to itself but they are all so counteracted, that the inexactness of the image which results from their presence very little exceeds, under ordinary circumstances of illumination, the limits which are set to the delicacy of sensation by the dimensions of the retinal cones....” (Helmholtz et al., 1995, p. 147)

And what about the laryngeal nerve in the giraffe, which takes an improbable 15 feet detour down its neck, looping around the aorta and traveling back up the neck? Even arch-adaptationist Richard Dawkins writes that he got something of an epiphany while watching the dissection a giraffe’s neck:

Not only would a designer never have made a mistake like that nervous detour; a decent designer would never have perpetrated anything of the shambles that is the criss-crossing maze of arteries, veins, nerves, intestines, wads of fat and muscle, mesenteries and more. (Dawkins, 2009, p. 371)

Quoting the biologist Colin Pittendrigh in a paper from 1958, Dawkins describes biological organisms as a “patchwork of makeshifts

pieced together, as it were, from what was available when opportunity knocked.” The problem is thus far from novel, and was also noted by Jacob (1977), who famously talked of evolution as “bricolage,” i.e. tinkering and improvising with what materials happen to be at hand. Jacob went further and directly contrasted the idea of bricolage to the idea of engineering, rejecting the latter as a useful analogy in biology. In his 1977 paper, Jacob wrote: “For the engineer, the realization of his task depends on his having the raw materials and the tools that exactly fit his project. The tinkerer, in contrast, always manages with odds and ends” (Jacob, 1977, pp. 1063–1064). Jacob traces back the objection to seeing living organisms akin to engineering products to Darwin himself, who in *The Origin* wrote:

Nor ought we to marvel if all the contrivances in nature be not, as far as we can judge, absolutely perfect. We need not marvel at the sting of the bee causing the bee’s own death; at drones being produced in such vast numbers for one single act, and being then slaughtered by their sterile sisters; at the astonishing waste of pollen by our firtrees; at the instinctive hatred of the queen bee for her own fertile daughters; at ichneumonidae feeding within the live bodies of caterpillars; and at other such cases. The wonder indeed is, on the theory of natural selection, that more cases of the want of absolute perfection have not been observed. (Darwin, 1859/2011, p. 472)

Of course Darwin was reacting against the then prevalent concept of intelligently designed organisms, but the point remains valid today.

The products of evolution, with their unwieldy complexities and clumsy ways, bear a greater likeness to Rube Goldberg machines than to the products of standard engineering. And even then, two major differences come to mind. First, at least Rube Goldberg contraptions are (baroquely, comically) designed with a specific and unique function in mind. Not so for living organisms, where reproductive fitness is all that matters, and no foresight at all is involved. Second, Rube Goldberg contraptions are not characterized by redundancy. Whereas a typical biological structure is robust and quite resistant to disturbances, owing to its redundancies and overlapping parts, a typical Rube Goldberg machine is extremely “brittle” and liable to errors: if one step in the machine does not work properly, or is not in lockstep with other parts, the whole apparatus breaks down.

(2) Second (though relatedly), even if engineering principles may be relevant for analyzing some biological systems (e.g., the concepts of lift, weight and air resistance for the flight of birds), this does not mean that evolution has somehow cracked the physical principles and equations underlying flight. Evolution typically proceeds by simple rules of thumb and rough-and-ready methods that are (more or less) functionally equivalent to engineering solutions in a well-specified ecological environment. The idea that natural selection is capable of “solving a near intractable physics-problem” (Pinker, 1997, p. 12), though having a kernel of truth, is also profoundly misleading (Pinker was talking about the smooth movement of your limbs). Animals don’t use algebraic fractions to calculate the level of altruism they should extend to their kin (not even unconsciously), any more than birds use latitude and trigonometry to navigate to their brooding places, or dogs compute parabolic trajectories when they’re catching a ball in flight. All these animals use surprisingly simple rules of thumb which, in their specific ecological environments, produce behaviors that more or less track engineering solutions. This may sound trivial, but it is a common fallacy to assume that, when animals engage in sophisticated behavior with an adaptive rationale, there is something similarly complex going on at a subconscious level. For example, with regard to catching balls in flight, Richard Dawkins

(in 1976, well before his above mentioned giraffe-induced epiphany) noted:

When a man throws a ball high in the air and catches it again, he behaves as if he had solved a set of differential equations in predicting the trajectory of the ball. He may neither know nor care what a differential equation is, but this does not affect his skill with the ball. At some subconscious level, something functionally equivalent to the mathematical calculations is going on. (Dawkins, 1976, p. 96)

In fact, experiments show that humans (and dogs) use a deceptively simple heuristic to catch a ball: keep your gaze fixed at the ball, and adjust your running speed such that the angle of the ball remains constant (Gigerenzer & Todd, 1999). When you follow this heuristic, you will be there when the ball hits the ground. As it happens, baseball players are very poor at predicting where a ball is going to hit the ground when they are asked not to run towards it (Babler & Dannemiller, 1993). They just manage to get there when the ball does. This is little surprising since computing the trajectory of a ball is a very complicated physical problem: one has to take into account initial velocity, angle, direction, spin, as well as the air current and the distance from the player. Gerd Gigerenzer and his colleagues (Gigerenzer & Selten, 2002; Gigerenzer & Todd, 1999) have documented several other simple heuristics in animal and human behavior. Again and again, complex behavior that seems suggestive of some mathematical or technical competence, turns out to emerge from deceptively simple rules-of-thumb. This is not to say that engineers can learn nothing from biological adaptations (e.g. the sonar of bats, the infrared pinhole camera of snakes). Adaptationists Cosmides and Tooby (1994, p. 43) have written that “natural selection is known to produce cognitive machinery of an intricate functionality as yet unmatched by the deliberate application of modern engineering.” This is right in one respect (face and speech recognition come to mind), but misleading in another. Naïve metaphors derived from engineering may fool one into believing that, if some biological feat can impress an engineer, it must be based on engineering principles that are at least as (mathematically) sophisticated, while nothing like that is actually the case.

(3) Third, even though optimality assumptions and engineering principles may be fruitfully applied to some macro-properties of biological adaptations, such analogies are a lot more problematic when moving down to the biochemical level. To illustrate this point, let's have a look at the genome, the so-called “blueprint” of an organism, to borrow another popular metaphor inspired by engineering analogies. The machine metaphor implicitly (and sometimes explicitly) assumes a relatively straightforward Genotype => Phenotype (G => P) mapping function, because this is how an engineer would go about constructing an artifact: the blueprint or design plan of a machine should be transparent and easily translatable to its functional components. This allows the engineer to go back to the drawing board to replace components, isolate unwanted parts without compromising the whole device, etc.

The concept of the G => P mapping function was advanced by Alberch (1991), primarily to illustrate why one is highly unlikely to accomplish in biology what the machine metaphor promises. The G => P has actually been extensively studied in a small number of simple cases, primarily the adaptive landscapes that describe the evolution of RNA molecules and of some protein functions (Copperthwaite & Meyers, 2007; Stumpf, Kelly, Thorne, & Wiuf, 2007). Even in these simple cases the degree of redundancy and non-linearity of G => P is such that the only way to make sense of it currently available to researchers is to exhaustively map the function empirically. The problem is that, for the foreseeable fu-

ture, this seems feasible only for simple instances of G => P, like the ones just mentioned, and—to a point, for very simple model organisms like *C. elegans*. Once we start talking about complex organisms, particularly those characterized by flexible developmental trajectories, all bets are off.

As the study of RNA and protein function G => P makes clear, the issue is not just complexity, but historically-induced redundancy and lack of one-to-one correspondence (what biologists call pleiotropy and epistasis). Modern living organisms are the result of a very long and tortuous evolutionary history, and—as mentioned above—natural selection is notoriously a satisficing, not optimizing, process (Ward, 1992).

It is telling that François Jacob, when he described evolution as a “bricoleur,” focused his attention where he saw tinkering most obviously, at the molecular level, pointing out that “the appearance of new molecular structures during much of evolution must, therefore, have rested on alteration of preexisting ones” (Jacob, 1977, p. 196). Famously, Gould and Vrba (1982) gave a name to this phenomenon: exaptation, the continuous modification of previous structures to be molded by natural selection to perform new functions—often in a sub-optimal fashion, when compared with the desiderata of an engineer. This is the same line of reasoning that brought Gould and Lewontin (1979) to criticize the adaptationist program, precisely because they felt it took too literally the idea of natural selection as an optimizing engineering mechanism.

Lynch (2007) has recently argued—following a completely independent line of thought originating from inside population genetic theory—that, contra popular assumption, much genomic architecture is not the result of selection, but of stochastic events. Molecular evolution may therefore be even more affected by chance than Jacob was willing to concede, and chance produces even more baroque outcomes than Rube Goldberg did, with the consequent difficulty in reverse engineering the resulting mess. Proponents of Intelligent Design have been notoriously misled by the huge complexity of the living cell exposed under the microscope, over-attributing functionality to parts that are actually redundant, functionally overlapping or even useless. In fact, ID creationists ironically take a hyper-adaptationist stance, portraying every component as an essential and irreplaceable cogwheel in the clockwork design, fully in the tradition of natural theology. In a previous paper (Pigliucci & Boudry, 2011), we documented how ID creationists have eagerly exploited the wanton use of machine/information metaphors on the part of some biologists. A striking illustration of this engineering rhetoric is the Discovery Institute's promotion clip *Journey Inside the Cell*, which features highly stylized computer-simulations and depicts the cell as a high-tech feat of nano-engineering. In the voice-over and accompanying text, director Stephen C. Meyer describes the biochemical processes in terms of “molecular machines and nano-processors,” “information-recognition devices” and “mechanical assembly lines.”<sup>5</sup>

Evolutionary adaptationists, of course, adamantly resist the inference of intelligent design. They may even grant the pervasiveness of phylogenetic constraints and stochastic factors, but insist that these problems can be safely ignored when engaging in adaptationist reasoning. Elliott Sober, in a balanced review of adaptationism, writes that proponents of the research program are committed to the irrelevance, not the non-existence, of genetic and developmental constraints:

The supposition is that a simplifying assumption about heredity—that like phenotype produces like phenotype—is usually close enough to the truth; the details of the underlying genetics would not materially alter one's predictions about which phenotypes will evolve. (Sober, 1998, p. 76)

<sup>5</sup> <http://www.journeyinsidethecell.com/>.

Indeed, Sober writes that the existence of phylogenetic constraints—evolution can only work with what is available—“is not inconsistent with the claim that the optimal available phenotype evolves” (Sober, 1998, p. 80). No matter how tight the constraints are, adaptationists would argue, it seems that evolution has a knack for homing in on adaptations that pass engineering requirements with flying colors. In many cases of anatomical design (e.g., the bat’s sonar, the hummingbird’s wings, the tuna’s body shape) natural selection has come up with a design that has been independently established as optimal from an engineering point of view.

But this may be exactly the point. Though natural selection may deliver the goods at the end of the day, it will not exactly win an engineering prize for its production process, which may often be described as a wasteful, clumsy and patched-up mess. Though this problem may be ignored to some extent in adaptationist reasoning, as Sober points out, it is highly relevant for synthetic biologists, who want to get down to the nitty-gritty of genes and biochemistry to “re-engineer” natural organisms.

### 5. Developmental encoding and software engineering

Jacob’s (and Darwin’s) point has been confirmed indirectly by research on the structure of artificial neural nets (e.g., Grossberg, 1988; Hornik, Stinchcombe, & White, 1989). These are a type of software that self-develop by trial and error in order for the program to “learn” a certain task, for instance to distinguish artificial objects from living organisms as distinct ontological categories. Different neural nets can be developed to accomplish the same task, and when researchers “open up” the program to see how exactly the strengthening and weakening of input-output links has achieved the desired result, they cannot easily (or, in many cases, at all) reconstruct the logic of the neural net. That is because there is no transparent “logic” to speak of; there is only a particular historical sequence of trials and errors, whose outcome depends on a number of contingencies, including the (usually random) starting point.

It is somewhat ironic that while biologists pursue the idea of reverse engineering, software engineers are more and more turning to biological systems as inspiration for their own work. This started a number of years ago with the idea of genetic algorithms, a now popular approach to develop software by direct analogy with evolutionary processes: software equivalents of mutation, recombination and selection are put to use to develop programs that would be too complicated for a human programmer to produce (e.g., Goldberg, 1989; Mühlenbein, 1991). More recently, software engineers have expanded the scope of their inspiration, moving from straightforward population genetics (on whose principles genetic algorithms are based, though some of those principles were rediscovered independently by engineers without previous knowledge of biological theory) to developmental biology, the very same field that was left out of the Modern Synthesis because it did not fit the reductionist program of the emerging science of genetics.

According to Hartmann, Haddow, and Lehre (2007), for instance, so-called “artificial development” is used to solve complex computational problems by way of replacing direct genetic encoding with indirect “developmental encoding.” Developmental encoding is a way to build electronic circuits by adapting two processes that are common in biological developmental systems (Roggen, Federici, & Floreano, 2007): the computation proceeds by using a signaling phase, where information is communicated locally within a given circuit (just as cells in an organism rely on internal monitoring to maintain homeostasis and functionality in response to their immediate surroundings); this is complemented by an expression phase, where local components (“cells”) of a circuit adopt a particular functional state depending on the signal

that they have received in the previous phase (the rough equivalent of adaptive cell memory in biological organisms).

The interesting point as far as we are concerned is that shifting from genetic to developmental encoding reduces the search space for evolutionary algorithms and produces systems that are less complex and more robust (“fault-tolerant”) than those obtained from standard genetic algorithms. This is in part because direct genetic encoding is limited by the fact that the length of the genetic string grows in roughly linear proportion to the complexity of the “phenotype,” which means that only relatively simple  $G \Rightarrow P$  combinations can evolve before the genome becomes too cumbersome and prone to systemic failures (Roggen et al., 2007). With developmental encoding, instead, the evolving system can take advantage of a small number of genetic instructions mapping to a large number of phenotypic outcomes, because those outcomes emerge from the local interactions among parts of the system and from interactions of the system as a whole with its environment. This approach requires much less explicit encoding of information in the genome itself. The resulting highly non-linear  $G \Rightarrow P$  map is reminiscent of those of living organisms, and yields more flexible and less fragile developmental systems.

From an empirical standpoint, simulations comparing the evolution of standard genetic systems of information encoding with systems based on developmental encoding clearly show that the former reach a maximum level of fitness for low levels of complexity. At higher levels of complexity developmental encoding scales much better, with developmental systems being capable of achieving high fitness more quickly and efficiently. Developmental encoding also leads to the artificial evolution of systems that are both significantly more robust to internal disruptions and significantly more flexible in response to external environmental conditions than standard genetic ones. These results clearly show a direct parallel between living organisms and developing software, but the parallel undermines, rather than strengthens, the machine analogy.

### 6. Different approaches to synthetic biology

How is all of the above relevant to the research program in synthetic biology? To begin with, there is no such thing as a single research program in this emerging area. Deplazes-Zemp (2011) distinguishes at least five conceptually distinct, if somewhat overlapping programs associated with synthetic biology:

- (1) Bioengineering. Uses standard biotechnology tools to build novel biochemical pathways in host organisms.
- (2) In silico synthetic biology. Similar to bioengineering, but carried out using computer simulations of novel metabolic pathways, rather than by experimentation with living organisms.
- (3) Synthetic genomics. As the name plainly implies, this is a much broader scale of bioengineering intervention, at the level of whole genomes—rather than individual pathways—being slated into a (de-genomicized) host cell.
- (4) Protocell synthetic biology. Here the aim is somewhat complementary to that of synthetic genomics: to bioengineer “living” cells that could then be used as entirely artificial hosts for other bioengineering projects.
- (5) Unnatural molecular biology. This approach is arguably the most ambitious, as researchers in this area pursue the goal of producing entirely new molecular biologies, for instance using expanded genetic codes, capable of incorporating more and different amino acids from those used by the natural code.



It is interesting to ask to what extent the above mentioned research programs may be affected by the breakdown of the machine metaphor as a guide to understanding living organisms on the biochemical level, as we think the adverse impact will be different in the five cases in question. Broadly speaking, one of the key features of the different branches of synthetic biology is the systematic “application of rigorous engineering principles to biological system design and development” (Royal Academy of Engineering, 2009, p. 5). Biological systems are functionally analyzed and reduced to individual parts, with specified functionality and input/output characteristics. Some synthetic biologists explicitly use the model of a computer, taking inspiration from its hierarchical and modular structure (Andrianantoandro, Basu, Karig, & Weiss, 2006; Deplazes & Huppenbauer, 2009). These bio-modules and their specifications are compiled into an inventory, in order to allow an engineer to recombine them into new biological systems. Detailed computer modeling is used (or will be used) to simulate the expected behavior of the new system. In the implementation stage, the validity of the model is tested, after which the model can be adjusted. This so-called classical “engineering cycle” can be iterated several times over, with new empirical tests providing input for revised models.

We suggest that the general inadequacy of engineering metaphors for biological systems (particularly at the biochemical level) impacts all five approaches listed above, but to different degrees and in different manners. We should immediately point out that we are not making predictions about the likely future successes and failures of any of these research programs per se, as science is a very pragmatic enterprise, with scientists often able to find a way around a particular problem by previously unforeseen paths. But we do think that whatever successes these researchers will be able to achieve will be in spite, and not because of, the inspiration provided by the machine metaphor. Indeed, we suggest that the more bioengineers will adhere to a straightforward “engineering” perspective on living organisms, the more obstacles they will throw in the way of their own progress.

Consider, for instance, the first two on the list, bioengineering and *in silico* synthetic biology. As pointed out by Deplazes-Zemp (2011) the major difference between them lies not in what the researchers are hoping to achieve, but in how they plan on doing so: direct experimentation with living organisms vs. computer simulation of metabolic pathways. The latter—we think—has to push the engineering approach further along than the former, since straightforward bioengineering can afford, to a considerable advantage, to co-opt naturally evolved systems and treat them as background conditions, which means that the scientists do not need to explicitly model them. Not so in the case of *in silico* approaches, where much more needs to either be assumed without evidence or modeled directly with relatively scant knowledge of the analogous biological systems. To the extent that *in silico* researchers will make “reasonable” simplifying assumptions, such as ignoring historically generated partial redundancies, they will depart from both understanding and empirical replication of actual biological systems.

The problem is exacerbated when we move to the next two projects, synthetic genomics and—even more so—protocell synthetic biology. Cells are much more complex objects than genomes, and of course depend on an interplay between genomes and environments in order to develop and function properly. But the crucial issue is that if individual pathways can hardly be reversed engineered due to their historically contingent redundancies, a fortiori genomes present the same challenge, only orders of magnitude larger. And finally, the idea of developing entirely new genetic codes is tantalizing, and certainly physically possible. But, again, it does not seem to take seriously the fact that the current natural system evolved by trial and error of a kind that would be extremely hard to bypass in the quest for an entirely novel system

that has much more flexibility and scope (and therefore complexity, and number of things that can go wrong) than the natural one.

None of the above is meant as a council for despair, but rather as a caution to possibly over-enthusiastic would be bioengineers, who may not take seriously enough the fact that living organisms are nothing like machines. There are at least two alternative approaches for getting around this problem, and we suggest that they will be the ones followed in practice.

First, the synthetic life program could co-opt the process of random mutation and selection, in the way that has been so useful for both genetic algorithms and neural nets (Andrianantoandro et al., 2006). One notable example is the practice of mutagenesis in biotechnology, in which the genome of some organism (or some targeted part of interest) is bombarded with radioactive rays (or exposed to other mutagens) to increase mutation rates. Instead of designing some desired trait from scratch, researchers produce a sample of random variations on an existing genome and select whatever interesting features arise as a result. It seems that evolution sometimes “discovers” solutions more quickly than we can, in part because it preserves whatever works, not aiming at optimality; and in part because organisms are a product of historical processes that resulted in far more messy and less transparent systems than man-made machines. As mentioned earlier, however, co-opting natural selection would probably carry some limitations on our ability to explicitly understand every detail of exactly why a given outcome evolved, precisely because of the vagaries and redundancies of stochastic, non-optimizing natural processes.

A second way to overcome this problem is to borrow genetic material from existing organisms, where the phenotypic effects are already known and associated with specific DNA sequences. This seems to work for relatively simple problems, such as the production of a single enzyme (for example, inserting anti-freeze genes from polar fish into tomatoes: see Hightower, Baden, Penzes, Lund, and Dunsmuir, 1991).

Notwithstanding the enthusiasm of some researchers at the idea of designing life from scratch, the relative success of these two strategies (co-opting selection and borrowing genes) highlights the limitations of engineering approaches and machine metaphors. Increasingly, researchers apply directed evolution as a designing aid for bio-engineering, a development which Deplazes-Zemp (2011) has even called a “withdrawal from the machine-analogy.” If synthetic biologists take a pragmatic stance, we predict that these strategies will be the ones further pursued, rather than the more radical design-from-scratch project inspired by straightforward machine metaphors. The adage that “evolution is cleverer than you are,” attributed to the biologist Leslie Orgel, may also be good advice to synthetic biologists.

Indeed, the obstacles encountered by synthetic biologists may further our understanding of what is distinctive about life, and where the analogy with human artifacts breaks down. For example, one of the problems currently confronting the program of synthetic genomics is the inherent complexity and messiness of  $G \Rightarrow P$  mapping. Synthesizing and inserting new genetic materials may cause unintended side-effects that are not necessarily predictable via previous knowledge or computer modeling (Giddings, 2008; Regal, 2008). One of the strategies to overcome this challenge is to produce a simplified “chassis,” for example by reducing the genome of *E. coli* bacteria. The hope is to minimize side-effects and unwanted interactions with the cell’s natural metabolic pathways, in other words, to make it more transparent and workable for a human engineer. A related approach, pursued by protocell synthetic biologists, consists of constructing a “minimal cell” containing merely the equipment necessary to support the synthesis of artificial DNA sequences. A report of the British Royal Society of Engineering on Synthetic Biology acknowledges the problem of “the interdependence of biological networks and possibility of non-spe-

cific interactions between the new circuit and existing (natural) circuits” (Royal Academy of Engineering, 2009, p. 30).<sup>6</sup> The problem is that even simplified cells and bacterial chassis are still quite complex, and that further simplification would preclude the implementation of synthetic DNA technology.

## 7. Artificial selection and engineering

In his defense of reverse engineering and artifact hermeneutics, Daniel Dennett writes that there might be no foolproof way to distinguish the products of blind selection and the handiwork of rational designers. Suppose we present Martian scientists with a laying hen, a Pekingese dog, a barn swallow, a cheetah, and a greyhound, challenging them to identify the products of foresighted and conscious design. On the one hand, they might note that the laying hen does not care properly for her eggs, and that the Pekingese is “pathetically ill-suited for fending for itself in any demanding environment” (Dennett, 1990, p. 192). On the other hand:

The barn swallow’s fondness for carpentered nest sites might fool them into the view that it was some sort of pet, and whatever features of the cheetah convinced them that it was a creature of the wild might also be found in greyhounds, and have been patiently encouraged by breeders. Artificial environments are themselves a part of nature, after all. (Dennett, 1990, p. 192)

The last sentence is crucial. Greyhounds were designed by co-opting selection, not by anything resembling the technology envisaged by synthetic biologists. The selection pressures of a natural environment and those consciously enacted by a human breeder may well resemble each other, as both breeders and mother nature may be “interested” in the same phenotypes (e.g., running speed). In such cases, the resulting organisms may be difficult to distinguish. As Darwin recognized, the two mechanisms fade into each other: in the case of sexual selection—itsself nothing but a special form of natural selection—the evolution of conspicuous male features is driven by female choice, much like the breeder’s fancy creates new variety of animals. Indeed, artificial selection can take place without conscious awareness on the selectors’ part, for example when humans unconsciously choose among puppies with the longest ears and largest eyes. Darwin described this as “a kind of Selection, which may be called Unconscious, and which results from every one trying to possess and breed from the best individual animals” (Darwin, 1859/2011, p. 29). No wonder the Martians are confused.

In any case, human breeders start off with a naturally evolved and reasonably functioning organism. They never fundamentally alter the biological constitution of an animal or plant, and in many cases they either exploit existing variation in the gene pool, or bide their time before new mutations arise. Human breeders, no less than nature herself, act on phenotypic variations, and remain blind to the underlying genetic and developmental processes.

All in all, what Dennett’s thought experiment shows is that current methods for designing living organisms still closely mimic the process of natural selection. Given the messiness of living systems, especially with regard to G=>P mapping, this is not a bad strategy at all for breeders (and bio-engineers). H. G. Wells notes that nature is still “*shy and slow in our clumsy hands*,” envisaging a future in which all of this will be done more thoroughly and efficiently. But consider some of the remarkable achievements of artificial selection, such as the speed of the greyhound or the herding qualities of the border collie. Would it have been possible to design such creatures by reverse-engineering the desired characteristics and synthesizing the genome from scratch? There is no way to know

for sure, but we bet that, at least for a long time in the future, some process of directed or co-opted evolution will be a much more viable way to achieve this.

This is not to say that the difference between synthetic life “from scratch” and artificial selection is always clear-cut. For example, synthetic biologists could synthesize a new gene sequence for a simple phenotypic trait, partially borrowing from existing organisms, and then go on to fine-tune the results using artificial selection (Andrianantoandro et al., 2006). Alternatively, they could simply take the living world as a general source of “inspiration,” so to speak, but then proceed from the bottom up, designing systems that would not be bound by any direct analogy with living organisms. Even this last path, however, is not guaranteed to yield to the optimist’s prediction because of the sort of reasons already brought to bear above, particularly the brittleness and relative inflexibility of the type of simple G =>P mapping we might be capable to engineer from scratch in the foreseeable future.

In essence, this is the same conundrum faced by Artificial Intelligence researchers (Ekbia, 2012): while the “weak” AI program can claim some spectacular successes (e.g., Deep Blue and Watson, the computers built by IBM to challenge world champions in chess and Jeopardy respectively), this has come at the price of giving up on the strong AI objective of producing a truly artificial, conscious, intelligence. The latter research program, while certainly possible in principle, seems at the moment to have ground to a screeching halt. One of the possible explanations for the failure of strong AI is precisely that its attempts at “reverse engineering” the brain are too confidently based on the idea that the brain is analogous to a machine (to be precise, to an electronic computer), as opposed to an organic product of blind evolution.

## 8. In search of new metaphors?

Pigliucci and Boudry (2011) have briefly examined some of the alternative metaphors or analogies that have been proposed to help us build a more proper mental model of living organisms, replacing the one generated by too strict engineering talk. Some alternative metaphors for thinking about the relationship between genomes and phenomes, for instance, include the idea of a recipe, where DNA contributes the equivalent of the instructions for cooking, but does not specify all of the details of the process, which are left to a continuous interaction between the recipe itself and the environment and ingredients that are being used.

Although the recipe metaphor does get us away from a straightforward talk of “blueprints,” and particularly from a simplistic, near one-to-one G=>P mapping function, it is of mostly educational use and is unlikely to generate novel insights to guide professional researchers. The same holds for the origami metaphor, due to Lewis Wolpert, which captures some important elements of embryological development (the circuitous step-by-step folding), but which obviously will not work as a new master metaphor for thinking about living organisms (nor was it intended as such). While we acknowledge that metaphorical and analogical thinking are part and parcel of the way human beings make sense of the world, in some highly specialized areas of human endeavor it may simply be the case that the object of study becomes so remote from everyday experience that analogies begin to do more harm than good. In particular, the systematic application of engineering metaphors to a domain that is fundamentally different from the world of human artifacts may send scientists on a wild goose chase. Wittgenstein (1953/2009, p. §109) famously said that “Philosophy is a battle against the bewitchment of our intelligence by

<sup>6</sup> One ambitious solution would be to extend or revise the natural genetic code, in order to design “orthogonal DNA circuits” that are biochemically isolated from natural systems (Royal Academy of Engineering, 2009, p. 30).



means of our language.” Perhaps a contribution of philosophy of biology to the field of synthetic biology is to help free the scientists from the bewitching effects of misleading metaphors, so that they can simply get on with the difficult and unpredictably creative work lying ahead.

## References

- Alberch, P. (1991). From genes to phenotype dynamical systems and evolvability. *Genetica*, 84, 5–11.
- Andrianantoandro, E., Basu, S., Karig, D. K., & Weiss, R. (2006). Synthetic biology: New engineering rules for an emerging discipline. *Molecular Systems Biology*, 2(1).
- Babler, T. G., & Dannemiller, J. L. (1993). Role of image acceleration in judging landing location of free-falling projectiles. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 15.
- Bowler, P. J. (1992). *The eclipse of darwinism: Anti-darwinian evolution theories in the decades around 1900*. Johns Hopkins University Press.
- Carr, P. A., & Church, G. M. (2009). Genome engineering. *Nature Biotechnology*, 27, 1151–1162.
- Carroll, S. B. (2008). Evo-devo and an Expanding evolutionary synthesis: A genetic theory of morphological evolution. *Cell*, 134, 25–36.
- Carruthers, P. (2006). *The architecture of the mind: Massive modularity and the flexibility of thought: Massive modularity and the flexibility of thought*. Clarendon Press.
- Cosmides, L., & Tooby, J. (1994). Beyond intuition and instinct blindness: Toward an evolutionarily rigorous cognitive science. *Cognition*, 50, 41–77.
- Cowperthwaite, M. C., & Meyers, L. A. (2007). How mutational networks shape evolution: Lessons from RNA models. *Annual Review of Ecology Evolution and Systematics*, 38, 203–230.
- Darwin, C. (1859/2011). *On the origin of species*. Empire Books.
- Dawkins, R. (1976). *The selfish gene*. Oxford University Press.
- Dawkins, R. (2009). *The greatest show on earth: The evidence for evolution*. Bantam.
- Dembksi, W. (2002). *Intelligent design: The bridge between science & theology*. IVP Academic.
- Dennett, D. C. (1990). The interpretation of texts, people and other artifacts. *Philosophy and Phenomenological Research*, 50, 177–194.
- Dennett, D. C. (1996). *Darwin's dangerous idea: Evolution and the meanings of life*. Simon and Schuster.
- Deplazes, A., & Huppenbauer, M. (2009). Synthetic organisms and living machines. *Systems and Synthetic Biology*, 3(1), 55–63.
- Deplazes-Zemp, A. (2011). The conception of life in synthetic biology. *Science and Engineering Ethics*, 18(4), 757–774.
- Descartes, R. (1637/2010). *Discourse on the method of rightly conducting the reason*. FQ Books.
- Ekbia, H. R. (2012). Fifty years of research in artificial intelligence. *Annual Review of Information Science and Technology*, 44, 201–242.
- Giddings, G. (2008). The release of genetically engineered micro-organisms and viruses into the environment. *New Phytologist*, 140, 173–184.
- Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. MIT Press.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. Oxford University Press.
- Godfrey-Smith, P. (2001). Three kinds of adaptationism. In S. H. Orzack & E. Sober (Eds.), *Adaptationism and optimality* (pp. 335–357). Cambridge University Press.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley Professional.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London B*, 205, 581–598.
- Gould, S. J., & Vrba, E. S. (1982). Exaptation, a missing term in the science of form. *Paleobiology*, 8, 4–15.
- Grossberg, S. (1988). Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, 1, 17–61.
- Hartmann, M., Haddow, P. C., & Lehre, P. K. (2007). The genotypic complexity of evolved fault-tolerant and noise-robust circuits. *Biosystems*, 87, 224–232.
- Helmholtz, H. V., Cahan, D., & Williams, L. P. (1995). Science and culture: Popular and philosophical essays. *Nature*, 378, 345–346.
- Hightower, R., Baden, C., Penzes, E., Lund, P., & Dunsmuir, P. (1991). Expression of antifreeze proteins in transgenic plants. *Plant Molecular Biology*, 17, 1013–1021.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2, 359–366.
- Hume, D. (1779/2012). *Dialogues concerning natural religion*. Empire Books.
- Jacob, F. (1977). Evolution and tinkering. *Science*, 196, 1161–1166.
- Laland, K. N., & Brown, G. R. (2011). *Sense and nonsense: Evolutionary perspectives on human behaviour*. Oxford University Press.
- Lynch, M. (2007). *The origins of genome architecture*. Sinauer.
- Mayr, E. (1983). How to carry out the adaptationist program? *American Naturalist*, 122, 324–334.
- Mühlenbein, H. (1991). Parallel genetic algorithms, population genetics and combinatorial optimization. *Lecture Notes in Computer Science*, 565, 398–406.
- Pigliucci, M. (2002). *Denying evolution: Creationism, scientism, and the nature of science*. Sinauer.
- Pigliucci, M., & Kaplan, J. (2006). *Making sense of evolution: The conceptual foundations of evolutionary biology*. University of Chicago Press.
- Pigliucci, M. (2010). Genotype phenotype mapping and the end of the genes as blueprint metaphor. *Philosophical Transactions Royal Society B*, 365, 557–566.
- Pigliucci, M., & Muller, G. B. (Eds.). (2010). *Evolution—The extended synthesis*. MIT Press.
- Pigliucci, M., & Boudry, M. (2011). Why machine information metaphors are bad for science and science education. *Science & Education*. <http://dx.doi.org/10.1007/s11191-010-9267-6>.
- Pinker, S. (1997). *How the mind works*. Norton.
- Provine, W. B. (2001). *The origins of theoretical population genetics*. University of Chicago Press.
- Regal, P. J. (2008). Scientific principles for ecologically based risk assessment of transgenic organisms. *Molecular Ecology*, 3, 5–13.
- Roggen, D., Federici, D., & Floreano, D. (2007). Evolutionary morphogenesis for multi-cellular systems. *Genetic Programs and Evolvable Machines*, 8, 61–96.
- Royal Academy of Engineering (2009). *Synthetic biology: Scope, applications and implications*. Royal Academy of Engineering.
- Sober, E. (1998). Six sayings about adaptationism. In *The philosophy of biology*. D. Hull &.
- Stumpf, M. P., Kelly, W. P., Thorne, T., & Wiuf, C. (2007). Evolution at the system level: The natural history of protein interaction networks. *Trends in Ecology and Evolution*, 22, 366–373.
- Vaesens, K., & van Amerongen, M. (2008). Optimality vs. intent: Limitations of Dennett's artifact hermeneutics. *Philosophical Psychology*, 21, 779–797.
- Ward, D. (1992). The role of satisficing in foraging theory. *Oikos*, 63, 312–317.
- Williams, G. C. (1996/1996). *Adaptation and natural selection: A critique of some current evolutionary thought*. Princeton University Press.
- Wittgenstein, L. (1953/2009). *Philosophical investigations*. Wiley-Blackwell.