ROBOTS BEHAVING BADLY: SIMULATION AND PARTICIPATION IN THE STUDY OF LIFE

Christopher KELTY

A 2011 paper in the journal *PLOS Biology* described an experimental test of Hamilton's Rule; the authors report success.[1] It is, they claim, 'the first quantitative test of Hamilton's rule in a system with a complex mapping between genotype and phenotype' and they 'demonstrate the wide applicability of kin selection theory.'

This is surprising for many reasons. First, Hamilton's rule is one of the more controversial features of evolutionary biology in the last half-century.[2] Second, the authors refer to their experiment as one in 'artificial evolution.' This is obvious enough given that experimenting on evolution is not easy to do with any but the most short-lived organisms like fruit flies and microbes, except that, third, this experiment did not use social organisms of any predictable kind such as wasp, ants, or even microbes, but tiny mobile robots named *ALICE*.

These robots raise a set of confusing questions that this article will address. The first concerns the role of the robots in biological research: do they *simulate* something (life, evolution, sociality) or do they *participate* in something? The second question concerns the physicality of the robots: what difference does embodiment make to the role of the robot in these experiments, where as we will see, there are some subtle distinctions between what is abstract, what is digital, what is simulated and what is physical? Thirdly, how do life, embodiment and social behavior relate in contemporary biology and why is it possible for robots to illuminate this relation? These questions are provoked by a strange similarity that has not been noted before: between the problem of simulation in philosophy of science, and Deleuze's reading of Plato on the relationship of ideas, copies and simulacra. Whereas robot scientists, biologists and some philosophers of science may argue that robots are an 'object' on which one can do precise experiments about a 'target' (life, social behavior, evolution), Deleuze might counsel instead that we do not treat robots enough like robots, nor take our curiosity about them quite seriously enough – for they themselves might be the object we are studying, or should be.

The Robot Invasion Begins...

Using robots to study evolution might seem at first extremely unlikely. The practice has steadily grown more common over the last twenty years, owing in large part to the manifest enthusiasm that computer scientists and engineers have shown for evolutionary theory as both a theory and as an engineering principle. Recently, however, biologists have also started to take such work a bit more seriously. Several reviews have appeared recently aimed at enlightening biologists about this work.[3] And indeed, one of the authors of the paper about Hamilton's rule, Laurent Keller (University of Lausanne), is a widely lauded biologist of social evolution (primarily in real ants) who is thus lending a modicum of legitimacy to a style of experiment that might otherwise be dismissed by his colleagues.

Since about 2005, Keller and collaborator Dario Floreano, a robot scientist at École Polytechnique Fédérale de Lausanne (EPFL), have done a number of experiments with robots that explore different aspects of evolution. In most of them, tiny robots move around on motorized wheels in an enclosed area and interact with each other. Often there is 'food' or 'poison' (objects with blinking LEDs for instance that can be detected by a sensor), which the robots are programmed to forage for (or to avoid). In some, the robots are predator or prey. Properly programmed and set to work, the robots have demonstrated the evolution of communication, the evolution of altruism, the evolution of information suppression, the relationship of signal reliability and relatedness; predator/prey co-evolution; and various aspects of the morphology of a robot body.[4]

In the test of Hamilton's rule, the Alice robots had infrared sensors that allowed them to differentiate the other robots from the 'food items' deposited amongst them, and two vision sensors that allow them to 'see' the walls (three black and one white) of the enclosure. The robots were instructed to gather up the food (i.e. push it towards the white wall) and then decide whether or not to 'share' it with the other robots. The robots do not consume the food, nor metabolize it – they are designed only to be highly simplified phenotypic vehicles for a limited set of genes that govern their behavior. They do not reproduce, sexually or asexually, but are manually regenerated by a combination of precise but capricious software (the source of mutation) and beneficent human intervention (the robots need to be plugged in or otherwise connected to a computer in order to download the new genome and become their own offspring).

Since these robots have neither brains nor DNA, they are equipped with a software-based 'neural network.' This simple set-up allows the experimenters to precisely vary the measure of relatedness in Hamilton's rule (*r*), and watch as the fitness of the robots changes over a series of generations. In each generation, the inclusive fitness of each individual was determined by how many times a robot shared the food, and how many times food was shared by another robot in the group. The robots were then 'selected' based on inclusive fitness, subjected to mutations, and then the next generation was put back into the arena with new food items. After hundreds of generations, the gene frequencies of the resulting robots confirmed just what Hamilton's rule would predict. As the experimenters put it:

> 'Because the 33 genes were initially set to random values, the robots' behaviors were completely arbitrary in the first generation. However, the robots' performance rapidly increased over the 500 generations of selection. The level of altruism also rapidly changed over generations with the final stable level of altruism varying greatly depending on within group relatedness and c/b ratio'.[5]

A careful reading of the article reveals that the experiment was conducted with 200 groups of 8 robots over 500 generations and 25 different treatments. That's a lot of robots. However, as the authors explain: 'because of the impossibility to conduct hundreds of generations of selection with real robots, we used physics-based simulations that precisely model the dynamical and physical properties of the robots'. So, to be clear: there are at least two kinds of robots in these experiments: 'real' (physical) robots and simulated robots.[6] But both kinds do 'exactly' the same thing, which raises the interesting question: why use physical robots at all? What exactly is the difference between a physical robot and a simulated robot? Do robots *simulate* nature or *participate* in it?

Robot Trouble

Certain aspects of the distinction between the natural and the artificial are reassuringly intransigent. And yet, as with anyone who studies animals, robot scientists love to chip away at those divisions. In the 1990s , Artificial Life researchers like Christopher Langton insisted that the pixelated creatures on their screens were alive. Other more jocular researchers, like Rodney

Brooks, weren't so much convinced that robots or their simulations were alive – only that they can no longer be ignored, or maybe even controlled. For most participants in the fields of robotics and artificial life, the processes of biology, evolution or the meaning of 'life' were useful primarily in order to build better robots. John Holland and David Goldberg applied the ideas of evolutionary theory to the design of algorithms; researchers in robotics initiated projects in evolutionary, cognitive, developmental, affective and epigenetic computing; the 'inspiration' of nature was oft cited amongst engineers as 'biomimesis' or 'bioinspired' technology.[7] Robots have obvious power in understanding issues of locomotion, perception and communication, as well as more recently, emotion and human health.

At roughly the same time, and often with very little interaction, evolutionary theory itself has seen a series of changes so that there is now a marked split between those who study *evolutionary theory*, and those who study living organisms (now or in the past). The link between evolutionary theory and the study of life has steadily become less obvious with generations of thinkers applying the concepts of evolution to everything from economic growth to Epistemology to universes.[8] As 'evolutionary theory' has expanded it has come to be regarded less as a theory of life and its organization and more as an incredibly powerful theory of change and diversity in any system.[9] It is clear that evolutionary theory has taken on not just a life but an ecology of its own.

Classical biologists might be justifiably skeptical: driven to extremes, evolutionary theory of this kind loses touch with the empirical – as well as with other domains of theory like those of ecology, development and physiology. They demand exploration, experimentation and the observation of living things.

But at the same time, computer simulation has risen in respectability: not only has there been a marked shift away from 'law-governed' theories to 'model-based' theories over the course of the 20th century, but also a recognition that computer simulation might be scientifically and philosophically novel.[10] Debates have emerged about the difference between simulation and experiment[11], about the rise of 'exploratory' experiments[12], about the role of surprise generated by simulations[13], and about the epistemological and ontological consequences of understanding based on the radically advancing computational power available to scientists.[14] Robots, therefore, seem to fall into this niche of respectability as they too become more powerful, diverse and tractable as tools of exploration and experiment.

Robots can be treated like model organisms responding to an experimental set-up, or they can be used as traps, lures or decoys that provoke behavior or reaction from an animal or human. The distinction between a robot that *simulates* something else (stands in for) and one that *participates* in something is not at all clear, whether to the biologists using robots or for those observing the scientists.[15]

Using robots (or computer simulations) to study *evolution* is apposite though because experiments on evolution are difficult to design in the first place. The time-scales involved require experimental set-ups that accelerate time relative to that of humans and other animals, as well as a significant degree of inference and assumption.. Today the mathematical theory of evolution allows considerable analytic and predictive power, but often requires an unsettling degree of simplification and only a tentative generality. Computer-based simulations offer a way to add ‘complexity’ back in, and robots, therefore, seem to be the next obvious step in such an exploration. Still they are just computer simulations. For many scientists the question remains: how do these simulations relate to the nature they describe? Philosophers of simulation pose this question routinely as well. Eric Winsberg (among others) offers a helpful distinction by pointing out that both in simulations and in experiments, there is a difference between the *object* of explanation, and the *target* of explanation – a mouse model can be an object while the target is humans or human cancer, just as a simulation of fluid dynamics in a computer can be an object while the target is the center of a black hole – something that manifestly cannot be accessed by humans.[16] Target and object can be distinguished without depending on other spurious distinctions like digital/analog or real/virtual and can be distinguished equally in bench experiments in a lab, field experiments, or experiments using software models of phenomena. Such a distinction would seem to appeal to the experimental biologists who are using robots, as it implies the possibility of a continuum of relations between object and target. Indeed, in a recent review by Mitri, Floreano and Keller, the uses of robots in biological experiments concerning social behavior are surveyed; they lay out just such a continuum that they call the ‘scale of situatedness' (figure 1).

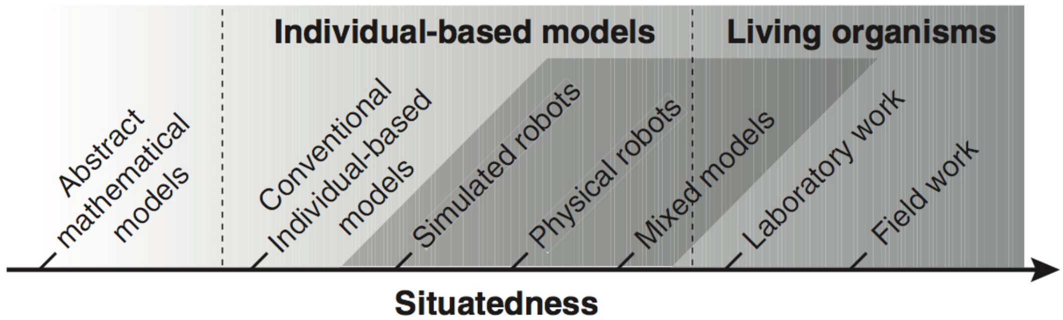

Fig. 1. Different approaches to studying social behaviour on a scale of situatedness, i.e. the extent to which individuals are embedded in an environment that they can sense and modify. The shaded box represents the robotic models covered herein.

Figure 1: **The Scale of Situateness.**

On the right-hand end of the scale are versions of 'situatedness' such as 'field work' and 'laboratory work' which 'include the whole complexity of the organisms and their environment... but rarely permit the unambiguous demonstration of causations',[17] especially concerning social behavior – pure 'target', to be sure, but seemingly as complex and inaccessible as a black hole. On the other end they put 'abstract mathematical models' that 'boil down collective systems to their minimal components'.[18] They are abstractions and formulas that model populations as a whole, rather than individuals in populations. Next on the scale are so-called 'individual' or 'agent-based models' that can take into account the varying behavior of a potentially large population of individuals. These models necessarily inhabit a computer simulation (rather than a formula on a page or screen in the case of 'abstract mathematical models'), usually represented not graphically but as something like a database or spreadsheet of changing parameters over time. Such models depend on the ability to both compute the complexity of interacting individuals (each of which might have its own more or less complex genotype and phenotype) in a population and in most cases to display that computation in some form (though not always as a graphical visualization).

Further down the scale are robots (both physical and simulated), which are essentially agent-based models in physical, programmed robots that interact in real (or simulated) space. Physical robots are better than agent-based model simulations because they do not need to make simplifying assumptions about the physical environment or the ‘physiology’ of the agent/individual (i.e. physical organization of the robot). As the authors put it, realism in an agent-based model is *costly* and *complex* whereas ‘the laws of physics are included “for free” in robotic models.’[19]

The scale is implicitly one of the complexity of bodies and their environments. Robots straddle a boundary: they share some of bodily complexity with organisms at the same time that they share putatively simpler (or just more controllable) computational and digital modes of existence with simulated organisms. That robots have bodies is not interpreted as a metaphor or a supplement, but instead as a crucial determinant of what robots are and can be. Embodiment is central to cognition not just in the sense that the robot needs to perceive and sense the world, but also implies that the kind of body it has will transform the kind of cognition of which it is capable.[20] When the authors say that the ‘laws of physics are included for free’ they mean that all of the complexity of the robot body – it’s weight, orientation, speed, the distance between its eyes and its wheels, etc. do not need to be simulated by a digital computer – they are just there, as part of what a robot is.

I’m not a Robot, But I Play One on TV…

> In Plato’s world variation is accidental, while essences record a higher reality; in Darwin’s reversal, we value variation as a defining (and concrete earthly) reality, while averages (our closest operational approach to essences) become mental abstractions.
>
> Stephen Jay Gould[21]

A key moment in the work of philosopher Gilles Deleuze is his confrontation with Nietzsche’s claim to be ‘overturning’ Plato’s philosophy.[22] For most interpreters, and arguably for Nietzsche, such an inversion concerns the opposition between the transcendent realm of the Ideas, and the fallen or immanent realm of copies. To invert Platonism would mean that the material or the concrete is elevated over the realm of Ideas.

However, for Deleuze it is not the case that Plato is distinguishing between a transcendent realm of Ideas and a fallen realm of copies or images, but offers instead a tripartite relation amongst idea, true copy and false copy, or among idea, copy and *simulacrum*. Copies in Plato are not all equally subordinate to the Ideas they mimic – indeed they do not mimic at all, but instead *participate* in the Idea, some much better and more authentically than others. The sculptor captures the image of a man when she understands the internal and external resemblance of the copy she creates to the model she observes; but the 'simulacrum' or false copy bears only a semblance, it looks like the real thing, but isn't the same 'on the inside' as it were. According to Deleuze, Plato observes the problem of the *rivalry* amongst copies, and the problem of deciding amongst them.[23] Thus, the question of whether robots simulate evolution, or actually participate in it (i.e. actually evolve) might be put the test of Deleuze's reading, by asking how the scientists themselves conceive of and practice the exploration of evolution using robots.

Deleuze's reading of Plato bears a striking similarity to the problem of simulated and physical robots discussed here. Rather than a simple opposition between real (living things) and fake (robots), the 'scale of situatedness' seems to set up a scale of rivalry along which some kinds of copies are better than others. But as in the case of Plato's philosophy, the question remains: on what basis does one choose one rival over others? What troubles some biologists about using robotsis that *they simply are not alive.*

It would seem therefore that the question is not 'what's the difference between an animal and a robot?' (model and copy), but 'what is the difference between a physical robot and a simulated robot?' (copy and simulacrum)? For Plato, the danger of the simulacrum is that it resembles the original on the outside, but lacks internal similarity (it is not a true copy). This would seem to describe any robot or especially classical automata that have the appearance of being what they copy but a radically different internal organization (the story of the mechanical Turk notwithstanding). Robots with fleshy plastic masks that simulate emotionally specific expressions are troubling and uncanny because they copy (increasingly well) the outside without being the same on the inside. Whatever we think it is that 'makes us human', it is present only in the *behavior* of the robot, not in its 'essence'. So what then is it that biologists think they are seeing when they decide to use a physical, embodied robot instead of a simulated one? That is to say, on what grounds do they distinguish these claimants to participation, and in turn, place them in order on the scale of situadedness?

One Robot May Hide Another...

The difference between agent-based models and robot models is frequently mentioned in experimental studies using robots. But is there also a difference between *physical* robots and *simulated* robots. 'Physics-based' robot simulations – simulations of the robots in question include more of the physical parameters of the robot than an agent-based simulation would, but presumably less that that of the physical robot. In these physics-based simulations, the price of the laws of physics starts to fluctuate: how detailed must a physics-based simulated robot be to be the same as a physical robot?

There is obviously a lot of complexity to a robot simulation, and not much information is offered by the experimenters when it comes to understanding what a simulated robot is or does. The materials and methods sections in the work of Floreano and Keller, for instance, often describe in detail the robot's configuration, precise dimensions and speed, programming, etc. as if all the robots in the study were physically real, and they rarely refer readers to the simulation tools used. They are modular, extensible software environments of a standard sort that can be used to cobble together simulated robots and run experiments using these robots. Physical and simulated robots are indistinguishable because they use the same operating systems, programming environments and software tools and libraries. It is a feature of all robots that they are first simulated in software and then, essentially, 'printed out' in physical form and set to work – and it is always the differences, breakdowns, or hardware-dependent surprises that form the core of an iterative process of inquiry in the robot sciences. The assumption that there is no difference between a simulated and a physical robot is therefore both deeply ingrained in the practice of creating robots, and warranted by an understanding of what gives the robot it's being: its software.

To return to Deleuze, the problem of distinguishing between false claimants is that they resemble the copy on the outside but not the inside –they do not really *participate* in the idea they copy. The simulated robots, however, offer a twist on this: they may have a visual representation (such as a 3D rendering on a screen) that resembles a physical robot, but what is indistinguishable is actually the internal aspect – the program or software governing the robot. Simulated robots have simulated bodies but their 'brains' and 'genes' are identical to those of physical robots: they are identical on the inside, but share no resemblance on the outside.

But if so, why not just run simulated robots through the motions and report on that? Implicit in the design of this experiment is that physical robots are necessary to produce that minimum of difference – that surprise – which can only come by running an experiment.[24] Reseachers expect, perhaps, that physical robots will confirm what the simulation demonstrates, but they must be included, observed, and the hypotheses thereby tested in order to assure them that the simulation acts as the physical robots would. In every case reported by Floreano and Keller so far, they use physical robots in the experiment – but only a few – and they do not seem to do anything other than confirm the results of the simulation. They produce no surprises, but are nonetheless essential to confirming that they are, as yet, unnecessary.

These Aren't the Droids You Are Looking for...

There are some obvious reasons why scientists might prefer robots (whether physical or simulated) to animals. They can do all kinds of ethically suspect things with robots that they cannot do with animals. There are also obvious reasons why an experimenter would prefer *simulated* robots over physical ones: they can reproduce and mutate dramatically faster than the embodied robots, they don't break down as often, and they can be reprogrammed and reconfigured much more quickly. The trade-off is that robots are not animals, and simulated robots are not robots. On the 'scale' of situatedness, there is a clear hierarchy.

On the one hand, the troubling implication of such a spectrum is that mathematical abstractions are figured as the least appropriate object for understanding a target – as opposed to the traditional assumption that they bear a kind of representational accuracy, if not identity with the target (mathematics as the 'language of nature'). But on the other, it seems to put mathematical abstractions, models, robots, model organisms and living organisms all on the same ontological plane – formulas are simpler versions of agent based models, which are simpler forms of robots, which are simpler forms of organisms, all of which *participate* in something not quite specified: life, evolution or social behavior. In other terms, it collapses the object/target distinction.

There appears to be an epistemological tension here: on the one hand robots are merely physical entities which can be simulated as other physical entities can (assuming we get the physical theories 'right') and what they do is ontologically indistinguishable from what mathematical formulae do or from what organisms do. But on the other hand is the assertion that *robots are different*: robots are *embodied* and so they will exhibit forms of behavior and/or cognition that can only come from being embodied – i.e. something that cannot be or does not reduce to their program, something that comes of having a 'body' not just a 'mind' (or program).

Thus the choice to work with robots appears strictly pragmatic: robots give some kind of partial access to (more complex) nature in return for being more tractable (simpler).[25] It sets up a *rivalry* between different kinds of objects which are also targets, and at the same time implicitly orders those that are most authentic expressions of some unstated essence: life, evolution, behavior, etc.

The implicit definition of embodiment here, however, is a physical one. 'Having a body' means 'including the laws of physics for free'. By doing so we gain purchase on the physics of bodies and environments that we would otherwise have to painstakingly model in a software system. But in the case at hand, what is being simulated is something else: *social behavior*, and in particular motivated social behavior that can be labeled as either selfish or altruistic. The scientists involved are no doubt committed to the epistemological claim that all such behaviors are fundamentally physical – either in a reductionist or an emergentist sense. There are genes which are biochemical entities that interact with the physical features of an environment which are then translated into proteins and on up the chain into organisms whose metabolically self-sustaining whole is liable to behave in predictable ways – such as by choosing to share food with another organism (QED, Hamilton's rule).

The question the scientists do not pose, however, is whether these robots have a 'social' body (real or simulated) as well as a 'physical' one, and whether there is any difference. Is social behavior a feature of the environment or is there a kind of social behavior that is not influenced by physical environments?

Is 'social behavior' simply 'physical behavior' governed by brains and genes? Could it be something more, and could the robots reveal it? Do the scientists expect – or perhaps hope – that the robots will surprise them? Unfortunately, to make things even worse, the very concept of 'social behavior' (to say nothing of 'altruism' or 'self-interest') is so elastic and vague that it's hard to know what most scientists mean even when humans and animals are both object and target, much less robots.[26] One cynical answer to the question 'Why use physical robots?' is simply that it's cool, or that it brings in funding, or that interdisciplinarity is (over)valued. However, the larger question concerning simulation and the 'style of reasoning' that is being developed does not thereby disappear – robots and simulations have entered the practice of science to such an extent that it is no longer possible to continue to treat them as gimmicks or illustrations.

But neither is it possible to treat these robots as less complex versions of living organisms. To do so is to misrecognize what they are, to mistake the false for the true claimant, and even more importantly, to follow Deleuze, to fail to give the false claimant a positivity of its own. To investigate, forthrightly, the internal difference that these robots harbor vis-à-vis animals and humans, and even, ultimately to argue against the existence of any model or original.

Do Robots Have a Life of Their Own?

When we learn something from an experiment using robots, are we learning about robots or are we learning about something else? Floreano and Keller tell us that we are learning something about ('confirming') Hamilton's rule when we watch these robots do what they have been programmed to do. The algorithms of the robots instantiated in their physical and simulated bodies unfold in accordance with our expectations about how ants or bees or humans would do the same. We are then comparing expectations to expectations and confirming their identity – this tells us that theory is correct for robots, although does not necessarily tell us that it is true for animals.

Given these robots and their algorithms, we can assert that simulations (including robots simulating animals) *animate* theory: they bring theories to life in time. In the case at hand the robots are animating Hamilton's Rule. Eric Winsberg's term for this is that simulations are 'downward': they draw on theory and perform it in a computer rather than being the kind of thing from which one makes observations and builds theories ('upward').[27] However, the spatial metaphor is misleading, as it relies too much on a hierarchical relationship and separation between theory and observation. Yet what simulations do is less about higher and lower, and more about static and dynamic. A simulation is preferable to an equation because it gives the equation life. .Confirmation becomes something the simulation does for us by enacting the theories we devise.

Further, if equations can be observed and experienced, then they have the capacity to generate insight, surprise and difference in an experimental system. They do in fact possess the potential to be 'upward' in Winsberg's sense, but only if one approaches them differently – as that which must be observed in its own right: not a true copy but a false one, a simulacrum with its own positivity, its own internal difference as a motor of change and exploration. There are many examples of such simulations generating both surprise and new difference and understanding. Simulations therefore are obviously 'downward' when they animate a theory, but they can also be upward, with a swerve, when observed. They generate surprises and insights and when this happens, there is no obvious commitment to realism, but there is a commitment to curiosity – the simulation does not represent something, it literally *becomes* something else. When this happens, a simulation is neither an attenuated nor a false copy of the real, but an object and process with its own complexity, an *assemblage*. Simulations therefore – and one must agree with Winsberg here, even if the metaphor is wrong – take on a life of their own; some of them even escape and find new niches in neighboring disciplines. To say that the robots *participate* in the experiment, however, implies that the robot (physical or simulated) is an image of something else – an instance of some other ideal form in which it participates. A mouse, as a model organism, clearly participates in 'life' insofar as it is a 'copy' of human life.

To view robots as participating in life or evolution is exactly what generates the skepticism among biologists – they are so obviously false claimants to the idea of life that it is exceedingly hard to treat any experimental result as having even the faintest relevance to understanding life. These robots have 'genes' and 'brains' that can only be used in scare quotes: they lack that resemblance (merely homologous thought it be) that we comfortably attribute to mice and men.. However, this suspicion of robots requires a certain Platonism amongst the biologists – or perhaps not enough Deleuze. The robots in these studies are not model organisms – they should not be taken as simply lower down on the great chain of being, less representative models – less realistic – but still useful. They are what they are – and we ought to be honest about it: they are *robots*. They are simulacra with a power of their own: assemblages of software, hardware, theories of computation and embodiment, theories of evolution and neural processing, vision sensors and motors, human operators and software environments for testing, all cobbled together for a range of reason like understanding, human companionship, capitalist efficiency as well as insane longings for other kinds of creation and control. Thus they are not *copies* of a 'wild type' of robots against which we might compare our laboratory results. Neither alive nor not alive, robots give us a glimpse into the complex imbrications of knowledge, living substance, existence in time, and the ability to affect or control any of these things. The more robots that are inserted into our environments and bodies, the more our sense of what it means to live will be transformed; the lines that so clearly seem to be both warrant for an experiment using robots, and argument against it, will fade.

---

[1] M. Waibel, D. Floreano, and L. Keller, 'A Quantitative Test of Hamilton's Rule for the Evolution of Altruism', *PLoS Biology*, 9:5 (2011), e1000615.

[2] Hamilton's Law (often written $rB > C$) asserts that a behavior that otherwise seems to be unselfish, perhaps even altruistic, might actually confer a fitness benefit (i.e. more success in reproduction) because it increases the chance that not only an individual's genes, but those of his

close kin, will succeed. Hamilton gave precise form to this intuition by introducing a formalism for what is now called "kin selection": if the organism benefiting from the altruist is sufficiently genetically related to it, then it might cancel out the reproductive cost of that act to the altruist; see W. D. Hamilton, 'The Genetical Evolution of Social Behaviour. I', *Journal of Theoretical Biology*, 7:1 (1964), pp. 1–16. The concept has been controversial ever since – whether associated with sociobiology, evolutionary psychology or the biology of social behavior; see J. A. Kurland, 'Kin Selection Theory: A Review and Selective Bibliography', *Ethology and Sociobiology*, 1:4 (1980) pp. 255–74; S. Okasha, 'Why Won't the Group Selection Controversy Go Away?', *British Journal for the Philosophy of Science*, 52:1 (2001), pp. 25–50; E. G. Leigh, 'The Group Selection Controversy', *Journal of Evolutionary Biology*, 23:1 (2010), pp. 6–19.

[3] J. Krause, F. T. A.Winfield, and J.-L. Deneubourg, 'Interactive Robots in Experimental Biology', *Trends in Ecology & Evolution*, 26:7 (2011), pp. 369–75; S. Garnier, 'From Ants to Robots and Back: How Robotics Can Contribute to the Study of Collective Animal Behavior', in Y. Meng and Y. Jin (eds), *Bio-Inspired Self-Organizing Robotic Systems* (Berlin, Heidelberg: Springer, 2011), pp.105–20; S. Mitri, S. Wischmann, D. Floreano, and L. Keller, 'Using Robots to Understand Social Behaviour', *Biological Reviews of the Cambridge Philosophical Society*, 88:1 (2013), pp 31–9.

[4] Mitri et al. 'Using Robots to Understand Social Behaviour'; D. Floreano, S. Mitri, S. Magnenat, and L. Keller, 'Evolutionary conditions for the emergence of communication in robots', *Current Biology,* 17:6 (2007), pp. 514–19; D. Floreano, S. Mitri, and L. Keller, 'Evolution of adaptive behaviour in robots by means of Darwinian selection', *PLoS biology* 8:1 (2010), e1000292.

[5] Waibel et al., 'A Quantitative Test of Hamilton's Rule for the Evolution of Altruism', on p. 5.

[6] There is some slippage between the terms 'real' and 'physical' – often the scientific publications in question use the two terms interchangeably. I follow their usage where possible, otherwise default to 'physical' to refer to robots that are extended in space, use electricity and are built out of plastic, metal and other materials.

[7] See e.g. J. H. Holland, 'Genetic algorithms and the optimal allocation of trials', *SIAM Journal on Computing*, 2:2 (1973), pp. 88–105; D. Goldberg & J. H. Holland, 'Genetic Algorithms and Machine Learning', *Machine Learning*, 3:2–3 (1988), pp. 95–9. For history and critical readings see S. Helmreich, 'Recombination, Rationality, Reductionism and Romantic Reactions: Culture, Computers, and the Genetic Algorithm', *Social Studies of Science*, 28:1 (1998), pp. 39–71; C. Emmeche, *The Garden in the Machine: The Emerging Science of Artificial Life* (Princeton, NJ: Princeton University Press, 1996); B. Bensaude-Vincent, 'Reconfiguring nature through syntheses: From plastics to biomimetics', in B. Bensaude-Vincent and W. R. Newman (eds), *The Artificial and the Natural. An Evolving Polarity* (Cambridge, MA: MIT Press, 2007), pp. 293-312; J. Riskin, *Genesis Redux: Essays in the History and Philosophy of Artificial Life* (Chicago, Mi: University of Chicago Press, 2007).

[8] R. R. Nelson and S. G. Winter, *An Evolutionary Theory of Economic Change* (Harvard, MA: Harvard University Press, 1982); S. E. Toulmin, *Human Understanding : Vol. 1.* (Princeton, NJ: Princeton University Press, 1972). L. Smolin, *The Life of the Cosmos* (New York: Oxford University Press, 1997).

[9] See especially Daniel Dennett's claim that evolutionary theory is "substrate-neutral": D. Dennett, *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon & Schuster, 1995).

[10] A. N. H. Creager, E. Lunbeck and N. M. Wise, *Science Without Laws: Model Systems, Cases, Exemplary Narratives* (Durham: Duke University Press, 2007); M. S. Morgan and M. Morrison, *Models as Mediators: Perspectives on Natural and Social Science* (Cambridge University Press, 1999). On simulation, see S. Schweber and M. Wächter, 'Complex Systems, Modelling and Simulation', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 31:4 (2000), pp. 583–609; P. Humphreys, P., 'Extending Ourselves Computfational Science, Empiricism, and Scientific Method' (New York : Oxford University Press, 2004); J. Lenhard, G. Küppers and T. Shinn (eds.), *Simulation: Pragmatic Constructions of Reality* (Dordrecht: Springer, 2006); T. Grune-Yanoff and P. Weirich, 'The Philosophy and Epistemology of Simulation: A Review', *Simulation & Gaming*, 41:1 (2010), pp. 20–50.

[11] Rohrlich, F., 'Computer Simulation in the Physical Sciences', *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1990 (1990), pp. 507–18; Winsberg, E., 'A tale of two methods' *Synthese*, 169:3 (2009), pp. 575–92.

[12] F. Steinle, 'Entering New Fields: Exploratory Uses of Experimentation', *Philosophy of Science*, 64:4 (1997), pp. S65–74; L. R. Franklin, 'Exploratory Experiments', *Philosophy of Science*, 72:5 (2005), 888–99.

[13] J. Lenhard, 'Surprised by a Nanowire: Simulation, Control, and Understanding', *Philosophy of Science*, 73:5 (2006), pp. 605–16.

[14] E. Fox Keller, 'Models, Simulation, and "Computer Experiments"', in Hans Radder (ed), *The Philosophy of Scientific Experimentation* (Pittsburgh, PA: University of Pittsburgh Press, 2003), pp. 198–215; W. Parker, 'Does Matter Really Matter? Computer Simulations, Experiments, and Materiality', *Synthese*, 169:3 (2009), pp. 483–96; P. Humphreys, 'Computational and Conceptual Emergence', *Philosophy of Science*, 75:5 (2008), 584–94.

[16] E. Winsberg, 'A tale of two methods'.

[20] See e.g. R. Chrisley and T. Ziemke, 'Embodiment', *Encyclopedia of Cognitive Science* (John Wiley & Sons, Ltd, 2006). Mitri et.al. cite two influential books in cognitive science: *The Embodied Mind*, by Varela et al. (Cambridge, MA: MIT press, 1991) and Andy Clark's *Being There: Putting Brain, Body and World together again* (Cambridge, MA: MIT Press, 1997). See also R. Pfeifer, J. Bongard and S. Grand, *How the body shapes the way we think: a new view of intelligence* (Cambridge, MA: MIT press, 2007).

[21] S. J. Gould, *Full House: The Spread of Excellence from Plato to Darwin* (New York: Random House, 1996), on p. 41.

[22] Key locations in the corpus include G. Deleuze, *The Logic of Sense* (New York: Columbia University Press, 1990), Appendix 1 'The Simulacrum and Ancient Philosophy'; *Nietzsche and Philosophy* (New York: Columbia University Press, 1983), and the conclusion to *Difference and Repetition* (New York: Columbia University Press, 1994).

[23] I rely here on Daniel Smith's reading of Deleuze: D. W. Smith, 'The Concept of the Simulacrum: Deleuze and the Overturning of Platonism', *Continental Philosophy Review*, 38:1–2 (2006), pp. 89–123.
[24] H. J. Rheinberger, *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube* (Stanford, CA: Stanford University Press, 1997).
[25] A stronger assertion might be that there is 'ontological indifference' at work here (See above Galison, this volume); perhaps the scientists are no longer concerned about the fundamental features of life, but only about the tractability or manipulability of an experimental set-up and the controllability of its components – how to make life (evolutionary theory, social behavior) *do something*.
[26] J. A. Piliavin and H. W. Charng, 'Altruism: A review of recent theory and research', *Annual review of sociology*, 16 (1900), pp. 27–65; D. A. Levitis, W. Z. Lidicker and G. Freund, 'Behavioural Biologists Don't Agree on What Constitutes Behaviour', *Animal Behaviour*, 78:1 (2009), pp. 103–10; Bergner, R. M., 'What Is Behavior? And so What?', *New Ideas in Psychology*, 29:2 (2011), pp. 147–55.
[27] E. Winsberg, 'A tale of two methods'.

Index terms:
Assemblage
complexity
Deleuze
Evolutionary theory
Hamilton's Rule
Kin selection
Participation
Philosophy of experiment
Plato's theory of forms
Robots
"scale of situatedness"
Simulacrum
Simulated robots
Simulation
Social behavior