

From DNA to TCP: Humanity and Evolution in Cyberspace

“We have used the words “mechanical life,” “the mechanical kingdom,” “the mechanical world” and so forth, and we have done so advisedly, for as the vegetable kingdom was slowly developed from the mineral, and as, in like manner, the animal supervened upon the vegetable, so now, in these last few ages, an entirely new kingdom has sprung up of which we as yet have only seen what will one day be considered the antediluvian prototypes of the race. We refer to the question: What sort of creature man’s next successor in the supremacy of the earth is likely to be. We have often heard this debated; but it appears to us that we are ourselves creating our own successors; we are daily adding to the beauty and delicacy of their physical organization; we are daily giving them greater power and supplying, by all sorts of ingenious contrivances, that self-regulating, self-acting power which will be to them what intellect has been to the human race. In the course of ages we shall find ourselves the inferior race. Inferior in power, inferior in that moral quality of self-control, we shall look up to them as the acme of all that the best and wisest man can ever dare to aim at. No evil passions, no jealousy, no avarice, no impure desires will disturb the serene might of those glorious creatures. Sin, shame and sorrow will have no place among them. Their minds will be in a state of perpetual calm, the contentment of a spirit that knows no wants, is disturbed by no regrets. Ambition will never torture them. Ingratitude will never cause them the uneasiness of a moment. The guilty conscience, the hope deferred, the pains of exile, the insolence of office and the spurns that patient merit of the unworthy takes - these will be entirely unknown to them.” –Samuel Butler, 1863

Cyberspace has changed our world, but at the same time it has created an entirely new one that may or may not truly belong to us. Any discussion of the virtual world we call cyberspace should properly include a look at the inhabitants of that world, most often called autonomous agents. These agents are not only able to fully exist in a world where we are visitors, but they appear to be evolving at a relatively rapid pace. An argument can be made that if artificial intelligence is going to emerge, cyberspace is the most likely location for it, with its massively distributed computing power and all of the benefits of graceful degradation that goes along with that. If cyberspace qualifies as the type of environment where intelligence might emerge, where does that leave humanity?

Imagine, for a moment, that your body is simply an illusion. You believe you are reading this page with your eyes, and holding a book with your hands, but in fact you have never even possessed eyes or hands. You are nothing but a brain, suspended in a vat of gelatinous goo, and a scientist pokes and prods your neurons expertly to simulate the activities of a body which you

do not have. More extreme than the Matrix, in which people have bodies but never fully use them, in this scenario your mind is much more like a computer, only able to obtain the information given to it by those who control it.

Many philosophers have imagined this scene before. Rene Descartes fathered the modern notions of dualism, or the separation of body and mind, using a scene very similar to this. John Pollock illustrated forms of extreme skepticism using the brain in the vat, arguing that the mere possibility that one is a brain in a vat is enough to question all of our forms of knowledge. The entire field of Artificial Intelligence set out originally to build a blank computer with an appropriate architecture, and simply “fill its head” with enough information about the world to create intelligence within it.

As appealing as some of these notions remain to researchers today, the evidence is stacked against them. While some forms of dualism persist, it seems tragically incorrect to claim that the mind doesn't reside in the brain, or at the very least remain terminally connected to it in some way, which is enough for our purposes. If we accept the extreme skeptic's viewpoint, we become locked in a world of solipsism from which there is no escape, and you might as well stop reading now since you have no proof that you're actually reading anything or that there is such a thing as language at all. And while many of the early crusaders in AI wrote extensively and intelligently about thought being nothing more than simple symbol manipulation, and therefore perfectly re-creatable in a machine, current experiments show that the largest success of this type in the field, Cyc, is simultaneously a spectacular failure. (For more information on Cyc, see <http://www.cyc.com>)

Luckily, there is a relatively new philosophical concept that goes a long way in solving many of these problems: Embodied Cognition. Simply put, this is the argument that claims

intelligence arises through the use of the body, and that many of one's thought processes are inextricable from one's body. Also called "situated action", this concept has been widely embraced by philosophers who, having seen the evidence from experimental psychology, accept that our bodies control the sorts of thoughts we have and even the concepts we obtain. While still devoutly rejected by many of the loudest supporters of symbol manipulation-as-thought, a simple illustration of the power of this theory shows readily how hard it is to reject: Simply close your eyes and imagine that your body is drastically different than it in fact is. Imagine that you were born as a member of the opposite sex. Would your stereotypical concepts be any different than they are? You would almost certainly see the world in different ways, and have different ideas about things. Imagine you were born without legs, and had to maneuver the world in a wheelchair. Or imagine that you were born without sight, or arms, or a sense of smell. Any of these strictly bodily diversions from regularity would almost certainly produce profound differences in your thought processes and concepts.

What does it mean to be intelligent? For the sake of brevity and simplicity, we will use the definition that Steven Pinker puts forth in How The Mind Works. He declares intelligence to be, "using knowledge of how things work to attain goals in the face of obstacles." (Pinker, 188) Pinker also goes into great detail explaining that brains evolved because of their ability to process information. The comparison to machines seems obvious.

Think back to one of the most important thought experiments from Aristotle, known as the Allegory in the Cave. While Aristotle used it to illustrate something very different, it serves a valuable purpose in our discussion of embodied cognition. Briefly, there are people chained to a wall in the cave, facing away from the light. They are unable to move around bodily, and therefore as people move back and forth in front of the light behind them, the only experiences

the chained people have with objects in the world is via the shadows they cast on the wall in front of them. To these people, those shadows are the objects themselves. They don't know that there is more beyond that. This is an excellent reflection on how our bodies are the tools that provide us with information about the world. If these people had full use of their faculties, they would be able to not only see the objects themselves as multi-dimensional, but feel the textures of them and smell and taste them as well. There is plenty of discussion about the importance of the body in cognition these days, so I will not recount it all here. (For a good discussion of these topics, I recommend Philosophy in the Flesh, by George Lakoff and Mark Johnson, and Being There, by Andy Clark.)

However, there is a vital piece of the argument that is barely discussed in philosophy, and yet completely taken for granted in evolutionary biology (the science that obviously accounts for why our bodies are the way they are). That important piece is the environment. The environment is not only what made our bodies the way they are, but it is because of interaction with our surroundings that our well-suited bodies are able to pick up information about the world and catalyze cognition. So much importance is being placed on the role of our bodies in cognition that the role of the environment in cognition is being overlooked. It is the construction of the cave as much as the constriction of the bodies within it that brings about the warped concepts that define the inhabitants of Aristotle's cave.

Obviously, evolution is to thank for molding our bodies uniquely to fit our environment. The environment drives the shape of our bodies via evolution, and the shape of our bodies drives intelligence. Therefore, the real requirements for intelligence to arise are a proper environment and evolution. Evolution itself can be summarized quite easily. Simply put, there must be heritability, mutation, and competition. Heritability simply means that offspring must resemble

the parents in some way. In the case of humans, this happens through genetic combination.

Mutation is the notion that there is some change between parent and offspring, and in the case of humans it often occurs by accident. Lastly, competition for resources is necessary for evolution to occur. There must be more offspring than can possibly survive, which drives Darwin's notion of natural selection – the idea that only those organisms best adapted to their environments will survive to pass their own genes on to the next generation.

Evolution was once uniquely in the domain of biology and the life sciences, but the idea has been commandeered by almost every discipline in existence. Psychology and sociology love to discuss the evolution of our brains and cultures, and now even computer science has adopted the notion to describe what is called 'genetic algorithms,' which are usually programs that run on the 'survival of the fittest' notion over a large sample size. It doesn't take a terrible stretch of the imagination to discover evolution occurring in cyberspace, specifically in the form of autonomous agents, also known as "bots."

If, as I stipulated earlier, all that is required for intelligence to arise is a proper environment and evolution, then what is to stop autonomous agents from developing a recognizable form of intelligence in cyberspace? Does cyberspace qualify as a rich enough environment to produce the feedback in the genetic code necessary to develop some kind of bodily structure to acquire intelligence in autonomous agents? Let's look a little closer.

First of all, I take for granted that there is a form of evolution occurring in cyberspace. If you look closely at the form bots have taken since their creation, you'll see a definite trend toward increasing adaptation to their surroundings, an obvious lineage from one form of agent to the next, and an intuitive notion that every autonomous agent written cannot possibly survive and

flourish, or cyberspace would be filled with meandering code and the entire “space” would become sluggish and all but unusable for its current purposes.

Just as man had his meager beginnings in *Homo habilis* and *Australopithecus*, autonomous agents had theirs in early daemons and even characters in text-based games. (Leonard 1, 25) What began as a simple autonomous program running on a single system at MIT in the 1960’s blossomed into what can only rightly be called a race of creatures that serve their different purposes, sometimes working in groups or communities, and compete to be the most suitable program for their unique environment of cyberspace.

If bots truly are evolving in cyberspace, which seems very probable, then we ought to be able to predict a likely course that their evolution will take. They have already gone from being very simple programs designed to perform a single task undirected (such as the first MIT daemon), to being complex strings of code that wander through cyberspace gathering and sorting information, simultaneously learning better ways to do their jobs (such as the Google web spiders). They are becoming increasingly complex, and rather than having rigid code for a single task, they are more flexible now, and able to handle whatever new obstacles cyberspace offers them this week.

Up until now, we’ve mostly examined cyberspace from the detached perspective of an onlooker, but as we all know, humans are not only the creators of this virtual space, but we also like to think of ourselves as its inhabitants. Assuming first that environments play a vital role in the evolution of intelligence, and second that bodies can be understood to be simply the recipients of the feedback from said environment (making the string of code running a bot the rough equivalent of the genetic code that forms our own bodies), we can show that bots very well

may be properly situated for intelligence to arise within cyberspace. So where does this leave humanity? Out in the cold, for better or worse.

If you look back to our earlier discussion of bodies, you will recall the example I used to illustrate the notion that bodies are necessary for intelligence, where you were asked to imagine your body as something drastically different than it is. Now, I ask you to dwell on that example a moment longer. If our bodies are the things that somewhat rigidly determine our concepts (and therefore oversee our thoughts), then it seems more than fair to say that they make us uniquely who we are. Each body is different from the others in various ways, and those ways (along with the environment) create the different experiences that make up our respective intellects.

You may be wondering how the makeup of the human body relates to cyberspace. Well, in a very real sense, it doesn't. And that is precisely the problem. If our bodies (both the genetic code that dictates their formation and the actual physical manifestation of that code) determine uniquely who we are, then what happens to personal identity in cyberspace? Without our bodies, we are merely visitors to a place that we can never fully enter. Cyberspace must always remain a "virtual" place to us, and humanity is relegated to the position of the outsider (literally).

What this leaves us with is that there is exactly no way to know what it means to be human in cyberspace, since our humanity is intricately connected to our bodies, and they are by definition forbidden to ever enter the space itself. We can discuss those parts of our humanity that are relinquished when we enter cyberspace (in many cases this includes inhibitions, morals, and accountability) or we can discuss those parts of our humanity that are squeezed between the cracks into our text and then actually make the journey to be cast off into cyberspace. These are both important discussions (although space limits them here). Yet it will always return to the fact that our bodies define us, and without them, we are less than human.

It has never been an issue for humanity that we play the role of the outsider in cyberspace, because it was our creation and regardless of the lack of physicality, it was still our space. However, we very well may have competition someday. And not only will they be necessarily intelligent (or they wouldn't be competition), but these competitors (in the obvious form of autonomous agents) will possess the very interesting quality of existing solely and uniquely within the space that we can never fully enter. In such circumstances, we have to ask whether it will cease to be our space at all, and whether or not it will become distinctly their own.

Throughout evolutionary history, there has been competition between similar species in which only one survives. The phylogenic tree of our own species development shows that *Homo habilis* and *Australopithecus* co-existed for some time, but we all know the outcome: only modern humans remain. It is exceedingly difficult and controversial to pinpoint exactly what dictates survival of one and extinction of the other, but one aspect that seems incontrovertible is control over one's environment. (See Barlow, 33) The more adaptable one is to the things they cannot truly control, the larger the chances of that one's survival. This of course poses a unique problem in the potential conflict between man and machine: it is not our environment, but we clearly entertain a larger degree of control over it in so far as we could, in theory, simply unplug it. However, the consequences for our own environment (breakdown of commerce, communications, and so forth) would be near-catastrophic at this point, were we to simply unplug the internet. There is no easy way to predict the outcome of such a conflict.

The great evolutionist George Gaylord Simpson said it well:

“The fossil record shows very clearly that there is no central line leading steadily, in a goal-directed way, from a protozoan to man. Instead there has been continual and extremely intricate branching, and whatever course we follow through the branches there are repeated changes both in the rate and in the direction of

evolution. Man is the end of one ultimate twig. The housefly, the dog flea, the apple tree, and millions of other kinds of organisms are similarly the ends of others.” (Barlow, 34)

Will agents be just another ultimate twig on the tree of life, co-existing and co-evolving like the housefly? Or will they erupt from our own branch, a new kind of descendant that shares a new kind of code with us, and eventually be the ultimate twig where we once were? Biology still disagrees over whether such things rely on good luck or good genes, but we can deduce at least this much: we can never completely enter their environment, but with our help, they can easily enter ours. This distinction is real, and important. “The emergence of life and intelligence from less-alive and less-intelligent components has happened at least once.” (Dyson, 9) Who are we to say it hasn’t recently happened again?

## Bibliography

Barlow, Connie, ed. *Evolution Extended: Biological Debates on the Meaning of Life*.

Cambridge: MIT Press, 1994

Dyson, George B. *Darwin Among the Machines: The Evolution of Global Intelligence*. Reading:

Perseus Books, 1997

Leonard, Andrew. *Bots: The Origin of A New Species*. New York: Penguin, 1997

Pinker, Steven. *How the Mind Works*. New York: Norton, 1997