Embodied Cognition in Robots and Human Evolution

by

Conor L. Myhrvold

A.B. Geosciences
Princeton University, 2011

SUBMITTED TO THE PROGRAM IN WRITING AND HUMANISTIC STUDIES IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN SCIENCE WRITING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SEPTEMBER 2012

Signature of Author: ______________________________
Program in Writing and Humanistic Studies
May 10, 2012

Certified by: ______________________________
Alan Lightman
Thesis Supervisor

Accepted by: ______________________________
Thomas Levenson
Professor of Science Writing
Director, Graduate Program in Science Writing

EMBODIED COGNITION IN ROBOTS AND HUMAN EVOLUTION

by

Conor L. Myhrvold

Submitted to the Program in Writing and Humanistic Studies on
May 10, 2012 in Partial fulfillment of the
requirements for the Degree of Master of Science in
Science Writing

ABSTRACT

This thesis investigates the notion of embodied cognition in humans using the research of former University of Washington researcher William Calvin and robots using the research of former MIT professor Rodney Brooks. The idea is that the feedback from the physicality of humans is a precognition to our intelligence. The choice example I use for our physicality is the motion of throwing, particularly the javelin throw. For robotics, I focus on the development of 'eyes' in Brooks' robot Cog and show how it demonstrated behavior we deem to be intelligent using the feedback gleaned from 'seeing'. Altogether, I present evidence for and against the notion that we are who we are, cognitively speaking, because of the sensory feedback of our physical bodies, and what that may mean going forward in the future for our intelligence.

Thesis Supervisor: Alan Lightman

Title: Adjunct Professor Humanities. Creative Writing, Physics.

Introduction

May 15, 1983. A man weighing 217 lbs and 6 ft 2 inches in height throws an 800 gram hollow metal spear. He throws from a run and it takes him 18 steps to do so, 16 plus 2 for the release, using 86 ft out of the 110 ft allowed. He draws back the spear at exactly 46 ft. He knows all of this without thinking about it because he has practiced this same run-up hundreds of thousands of times since the late 1970s.

Just before the release, the man sticks his left leg out to vault over it, putting over 1000 lb per square inch on his block foot, half the jaw force of a crocodile clap. The spear is released at 72 miles per hour, rotating counterclockwise at 1 rotation per second and 14 degrees parallel to the ground. We know this because even more significant than the throw itself was the way it was captured; with high speed video, at 200 frames per second. This technique would revolutionize the way track and field is studied following a publication of the throw's analysis in the *Journal of Applied Biomechanics*. Since then, high speed video in sports has given us an understanding of what our bodies do that has caused us to adapt how we train, akin to what a series of photographs of a galloping horse called "The Horse in Motion" did for a how a horse moves in 1878.

Yet the first high speed video of a biomechanical movement captured something that had also never happened before. At the 1983 Pepsi Invitational Track Meet in UCLA, Tom Petranoff had set the new world record for the javelin throw by just under ten feet, the height of a basketball hoop (exactly 3.00 meters). 327 ft later – further than the length of a football field – Tom Petranoff's javelin lands in the grass, head first. We know this because it was measured.

Throwing the javelin is conceptually simple but physically complex. A long throw looks effortless and fluid on film, yet this observation is literally only skin deep. For all that has been written and recorded about our bodies, we still know little about how we enable ourselves to do a physical motion like throwing. A throw involves millions of muscle parts all timed by an unconscious process that is not understood, making the javelin throw arguably one of the most technically complex athletic events in track and field.

And while some animals may run faster or jump higher than us, no other animal comes close to doing this. Throwing on the run is uniquely human. We share 98% of our genetic material with chimpanzees, so much so that Jared Diamond's book on human origins was called "The Third Chimpanzee", but an alien scientist who landed on Earth would know immediately from our physical abilities that we are a totally different animal.

Throwing: an Extreme Example of Our Physicality

It is well known that humans have unique cognitive abilities, what we call intelligence. Lesser known are our unique physical abilities, such as throwing. The question naturally arises: is there a causal link and if so, how would we know? To explain these behavioral traits and what separates us from other animals in form and function, many anthropologists have spent much time theorizing on how events in history could have fortuitously fallen into place to get us physiologically and cognitively to where we are today.

The standard story of human evolution is one of cleverness and cunning. Bipedalism freed our hands to make tools. Occasional mutations gave us a slightly bigger brain to outthink our opponents. These two traits allowed us to outwit and out-innovate our predecessors such as the australopithecines on the savanna and our Neanderthal neighbors up north. Our direct ancestral lineage spread out of Africa to the rest of the world.

But modern anthropologists realize that, at the very least, details are lacking on exactly how bipedalism and bigger brains fed off each other. Distinct million year gaps exist between when we became bipeds (4 m.y.a.) and the earliest evidence of tool making (2.5 m.y.a. at the beginning of the Oldowan period). Similar gaps exist between when *Homo erectus* developed big brains (tripling in size, and the cortex section quadrupling in a small period of time) and when we made use of those bigger brains, manifested through tool making and other such behavior. Furthermore, there is another million year period of stasis in how we made hand axes (from 2.5 to 1.5 m.y.a.). These gaps are unexplained.

Exactly when these evolutionary progress came about may be in part due to chance. Nevertheless, something tangible, and not coincidental, advanced us in each of those periods. The appearance of these traits may not be dated to closer than a several hundred thousand year range, but the nature of the changes were clear cut cognitive and physical evolutionary advantages. So the question remains: what are these advantages, and can we look to unique features of humans to pinpoint what they are?

Is Physicality the Key to AI?

For several decades, stabs-in-the-dark in computer science attempted to stumble on the holy grail of cognition by creating machine intelligence. United under the banner of "Good Old Fashioned Artificial Intelligence" and the Physical Symbol Systems hypothesis, rules about how the world worked were passed onto devices we built. Intelligence was not built from the ground up but rather dictated from us to our creations. But some Artificial Intelligence researchers like former MIT professor Rodney Brooks did not believe this was the best way to develop the same cognition we take for granted in humans, in machines. They believe the answer to creating the intelligence complexity of a human can be gleaned from understanding the feedback between our physical infrastructure and our brain...and there is increasing anthropological evidence that this is how we became so smart in the first place. From an evolutionary standpoint, the solutions to problems we needed to solve – even those that employed abstract thought – have roots in what we can physically feel and sense.

Brooks believes we needed this physical feedback to create cognition like ours. We needed an environmental context. This overall philosophy, called *embodied cognition*, claims that many behaviors we take for granted of brain-body interplay like throwing were developed over successive sequences rather than in a sudden evolutionary jump. Brooks and other like-minded researchers point out that since our body and brain developed together, our unique cognitive skills could be the direct result of our unique physical skills, or vice versa, and that although we don't currently have the evidence to say that one led to the other, this evidence may be coming in the near future.

Brooks has spent his career using robots to change the way we think about artificial intelligence. In fact, one could even say he spent much of his childhood doing so. Growing up in a lower middle class household in Australia as a teenager, Brooks dove for errant underwater golf balls on a nearby golf course to fund his forays in electronics.

Today Brooks is in his mid-fifties. In his modern Cambridge apartment, a sleek urban kitchen with polished wood panel floors gives way to a fully-fledged machine shop tucked away in the corner, with lathes, screws, wires, metal components, and all the electronic gizmos and gadgets he wishes he had access to as a kid. In fact, in what Brooks described to me beforehand as his "wacky post-MIT hobby", he is using the same components that could have made a computer or robot not only in the 1970s, but well back in the nineteenth century as well: in an era before the transistor, when a mechanical computer could have been possible with wire intersection crossings representing what is nowadays bits. People, says Brooks, have a mystical idea of what constitutes a computer's memory and when you see it for yourself (not possible now given the small size of a digital bit), it becomes more real.

So Brooks has spent an hour or two each evening, every day, for the past several months, working on a project to demonstrate just how you could construct a computer with technology that was available in 1860. Each major memory platter of the computer is composed of rows and rows of switches where information can be stored and relayed to another section. There are only two possible positions for each switch, a binary system represented by a 1 or a 0 depending on the switch orientation. In this regard, information regarding numbers can be mechanically stored using the binary system, as all digital computers still use today. In aggregate the memory platters are equivalent to 64 bytes, "below *anything*" used today in computing, emphasizes Brooks. In order to do 64 kilobytes, a more reasonable amount of memory for today's applications, Brooks says in a characteristic Aussie accent, "I would be here for the rest of my life!" making memory platters.

Brooks, true to physical form, is manually assembling his machine, by hand. (Using tools that weren't available during 1860, says Brooks, would be cheating, although he says double-checking the machine using modern technology is okay so long as the old proof-of-concept is shown, for if you bump just one switch out of the hundreds to a different position, you mess up the representation of information.) The end result is that you can see the switches on the physical bits adjusted to either ‘on’ or ‘off’.

Brooks is just getting started; his mechanical computer will take at least as long to finish as the work he has currently put into it to date (which he characterizes as an understatement). Several more completed platters are needed before the ‘computer’ can be fully functional – though Brooks admits he has scaled back his original ambitions of much more memory, “just because I don’t want to work on this project forever.”

In the interim, more platters means more parts. Propped open on the surface of a nearby countertop is an electronics catalog worn from use, ready to aid in placing another order of more mechanical parts – and other more modern components that Brooks fiddles around with – at a moment’s notice.

Brooks’ machine will show that ultimately all information is physically grounded even if today it occurs – on modern microprocessors and memory – at scales we cannot see with the naked eye. But Brooks emphatically exclaims that, “my ultimate goal is to *have fun*!” What excites him, he says, is that “you can *see* the bits. You can see them!!” Brooks is clear to point out that this latest project “is an art piece, not a science piece.” When you have made enough scientific contributions over the years, he said, you’re allowed to engage in art even if you’re still an amateur at it. (The Nobel laureate physicist Richard Feynman had a similar renaissance with the liberal arts starting when he was middle-aged.)

In this regard, Brooks’ love of the physicality of computer system is similar to the way he champions embodied cognition as a way to give the physical component required for intelligence the respect it deserves. Rodney Brooks’ views on artificial intelligence can summed up in a now-famous 1990 paper, "Elephants Don't Play Chess". In it, Brooks lays out how a physical motion can lead to cognitive evolution that would follow. His premise is that a physical development is a prerequisite to what we deem intelligence. The physical feedback that was the cornerstone of the beginning to intelligence preceded what we call complex cognition – such as inferring human emotion from body language – if cognition did indeed proceed in a stepwise process.

Brooks believes that building a robot is best done under a principle called the physical grounding hypothesis; creating a complex system from scratch requires feedback about the physical environment the system operates in, and the more the merrier. Without this information, Brooks argues, the path to complexity is doomed because the robot's cognition will not be grounded in the reality of the world it operates in. In the dozens of robots Brooks has built to perform robot-like tasks – recognizing and snatching empty soda cans from desks, and scurrying from one dark corner under a table to another – physical grounding is a better means of accomplishing a task and guiding a robot than the actual instruction of the task from the top down. In the case of a soda can, the robot does not need to distinguish its objective as a stand-alone entity but rather only from other non-soda can objects on a desk, which can be inferred by size, shape, weight and a cylindrical orthogonal orientation with respect to the table top, among other qualifications. The robot has no depth perception; it has no idea of depth versus scale. But it does know to look for the a cylindrical shape with the right ratio of height to diameter, and then drive up to the object and reach out with its hand to make sure that it is dealing with the right-sized object at the right distance instead of a larger object viewed from a larger distance in the camera.

Yet the fact that computers can outperform or enhance human analytical thinking has led us to characterize them as an amazing tool. Arithmetic and higher math could be done at first thousands, then millions, then billions times better on a computer than by with humans alone. (And now in combinatory-intensive fields such as cryptography, what humans could not do at all over the entirety of our civilization.) Computers were so much better right from the beginning that computer scientists thought that computers would soon begin to think like people. Artificial intelligence would soon equal, and then trump, our judgment and wit. In short, we would be obsoleted. Yet the computers which are fabulously good at Boolean logic and arithmetic are still, as of 2012, not able to make robots like us. No amount of doing trigonometry or arithmetic well has helped solve the greater problem of mimicking human thought.

The Evolutionary History of Throwing

How did humans learn to perform a complex physical feat like throwing? Millions of muscles fire in an unconscious effort that lasts just milliseconds. Throwing not only requires quick reflexes and hand eye coordination, but also the know-how to set our bodies up in a position to cause an involuntary reaction on the side of the body that is propelled forward to release the throw. One hypothesis is that intelligence was a prerequisite. With self-awareness that resulted from our bigger brains, we got smart, evolved and learned to throw. But given the uncertainty in human evolutionary history, one could equally argue that our progression toward intelligence was initiated by unique physical feats such as throwing.

University of Washington neuroscientist William Calvin has a controversial theory on why throwing, more than other locomotive behaviors like long distance running, could be how we became smart. Calvin started his academic career looking at the jitter of neurons in the late 1960s and found that their resolution limited the ability to precisely perform a several millisecond motion. Throwing a dart, describes Calvin, takes approximately one eighth of a second, the limit of a single neuron, “so if you make a mistake midway then you’ll be off target. The feedback between your brain stem, traveling down your spine and through the nerves on your arms, is not enough time to make a self-correction once you have initiated the throwing motion.” Calvin first laid out his ideas in finished form in a 1982 paper titled “Did Throwing Stones Shape Hominid Brain Evolution?”

In the paper Calvin wrote, “During about 1 sec of throwing time, the desired trajectory of the rock must be translated into a muscle sequence, orchestrating the times at which the various muscles are brought into play to accelerate and then precisely releasing the rock with a very narrow “launch window”.” The paper is well known in anthropology yet almost 30 years later Calvin’s way of thinking has not become mainstream thought.

Calvin today is in his sixties and retired, but still very much engaged in writing and thinking about society’s grand science issues. His mind shows no sign of slowing. In addition to human evolution, he has recently written a book on global warming and how he hypothesizes solutions to anticipated climate change can come about. In the Seattle area, where I am from, he remains a connected and intellectually curious man. His wife, also an academic, worked in the biology department at the University of Washington while I was working there.

In an interview with me last December, Calvin points out that a lot of cognition was required to mentally calculate a throwing trajectory at a moving target. There was little room for error. The fact that these two timescales – throwing and the firing speed of a neuron – are roughly comparable with each other is suggestive of how one could have led to the other (both on the order of several milliseconds; the exact measure is still subject to uncertainty).

The fact that our physical and cognitive capacity evolved hand-in-hand is already indisputable since our brain has always been a part of our bodies. Consequently, the idea that we can only think as fast as we can physically act – because that's what was necessary to survive – is a natural extrapolation of the connectivity that our nervous system provides. This would especially be true in the hunting and gathering era, before the advent of agriculture and we developed sedentary societies.

Calvin speculates that the neuron firing for coordinating our bodies and directing our throws could have become idle at times and thus available for other purposes, like language. Any combination or concert of neurons, an ensemble of sorts, would immediately be selected for as it would create a system that was greater than the sum of its parts. With this new setup came new uses that could only happen with many neurons working together, such as increased motor skills and eventually cognition. Thus, a selection for increased neuron bundling over evolutionary time scales (hundreds of thousands to millions of years) was also a selection for increased intelligence, a scaling phenomenon Calvin argues enabled a critical mass of neurons to lead to complex thought.

Thus, the physical demands for precision timing would select for large brains of many neurons. Calvin argues that we had different motives for throwing, and that hunting prey and warfare may not have been the original reason to develop the ability to throw. At first we could use throwing to cause commotion and distract prey or predators. Calvin calls this first stage "side-of-the-barnyard throwing" because it was inaccurate. Even if throwing could not inflict direct damage, Calvin tells me, "it could knock small prey off balance by the rock momentum against the upper part of body tipping over, or by causing a reflexive muscle reaction, delaying an animal's escape." Then, we could dispatch it by other means. Afterward, theorizes Calvin, we started to throw for accuracy and because it was easier to strike an animal from afar than approaching it up close. Such a progression would select for accuracy because throwing and missing an animal is worse than not throwing at all – animals are intelligent enough to figure out that our intentions are not just to get their attention.

After that, Calvin says, the toolmaker took over. "Throwing specializations would seem a reasonable candidate for an early special-purpose invention whose neural machinery could have important secondary uses such as tool-sharpening and fire-starting." So in fact by learning to throw, maybe we learned to become humans.

Any theoretical aspects of how throwing could give us cognition needs to have a credible evolutionary path. Applied to throwing, until you have a projectile, you do not care much about the projectile trajectories, so there would be no reason to expect that we would have the depth perception and depth of field ability to track a moving object and hit it, without the intermediate prerequisites of the throwing motion and knowing what stones to select for to yield a reliable flight. The key to the evolution of throwing, and other such features of human evolution, says Calvin, is that the progression had to "pay for itself" at each stage. That is, each evolutionary stage had to have enough of an advantage that it would become permanent, "selected for", a trait important enough to be deemed evolutionarily indispensable in the agnostic eyes of nature.

The same sort of manual dexterity and athletic ability that made us able to hit a moving target could then be applied to make us toolmakers in a broader sense. "Once invented, precision throwing would be on a faster-is-better curve for some time…The bigger-is-faster-is-better aspect of a sequencing center suggests that throwing skills could have continued to select for bigger brains from the Pliocene to recent times, aided by such supplementary inventions as the javelin, throwing stick, and sling."

The History of the Javelin

And of all the supplementary throwing inventions we have fashioned, perhaps none is more famous than a spear. The javelin is one of the oldest in a long line of throwing implements, tools that we created as a consequence of our intelligence to help us capture prey and defend ourselves. The earliest spears known to man are about 500,000 years old and found in Germany. They are wooden and have a tapered shaft and fire hardened tip. It is generally accepted that they were used as a throwing projectile. Compared to the modern javelin of today, throwing spears in the prehistoric era were lighter, says throws coach Kevin McGill. They were also made of wood instead of the metal or carbon javelins used for competitions today; in fact, these were only introduced by Dick and Bud Held, two brothers in the second half of the twentieth century. This material density difference would have certainly affected the aerodynamic properties and feel of these ancient, light wooden throwing implements.

Professional javelin throwers such as Duncan Atwood point out that we also ran mostly barefoot, on barren ground or grass, which means we likely ran for a shorter distance before releasing the spear (much less than the 80 to 100+ ft today). Spears predated shoes by a huge margin. Some of the earliest shoes are preserved in peat bogs in what is now known as Denmark, and date to only the Bronze Age, about five thousand years ago at most. Many modern hunter gatherers today still do not use shoes.

Modern javelin throwing appeared in the ancient Greek Olympic Games several thousand BC, and it had a functional purpose: training for warfare, akin to a sharpshooting competition today with firearms. But even this javelin throw was sometimes flung from a sling. We have taken every aspect of this more recent throw and amplified it even more, and the result is one of track and field's most extreme and notoriously injury-prone events.

We throw other things today too, but the javelin is most extreme. Not just one of the most extreme events from a biomechanical standpoint (the forces that are produced from the throw are just below the threshold to snap the ligaments of our elbow in a professional throw), but also in the long time period it takes to train your body to put yourself in the positions necessary to throw far.

The best throwers in the world typically peak in their thirties, unheard of in running and jumping events which rival the javelin in the required timing and technique. And there is ample evidence that even the world's best throwers have not reached what is physiologically possible: the margin between the current world record holder, Jan Železný, and the second longest throw vastly greater than the equivalent percentage differences in other track and field events.

The javelin throw is linear, done from a run and preserves the component of accuracy and distance that made spears a lethal part of hunting and warfare. Other field events such as discus, hammer and shot are rotational and not linear, while baseball pitching is stationary and done from a mound. (An extra step is considered a "balk" and is not permitted in baseball; it is in the game of Cricket.) The javelin throw today has tested the physical limits of how and how far we throw.

In an effort to try to understand what it takes to make the throw, we have made every attempt to slow it down. In electronic devices, by recording the movement. In our minds, by translating different exterior components of it to a number that can be compared to something else: a release speed, an angle of release, a distance where medals are won and lost over a mere centimeters. Tom Petranoff's World Record throw in 1983 was an example of all of the above, and the throw was further broken down in detail – along with the numbers from many of Petranoff's others – in a 239 page academic book titled "Biomechanics of the Javelin Throw". Three years earlier at the 1980 US Olympic Trials, Petranoff himself finished less than an inch behind Duncan Atwood on each of their second longest throws (their longest throws were exactly tied down to what you could measure, and the tiebreaker became their next longest throw.) The top three throwers would make the Olympic team and after the tiebreaker, Duncan Atwood came in third. Tom Petranoff came in fourth.

The History of Throwing

Under what original context did our physical abilities evolve? And if Calvin is right, when did we start to throw? Most experts believe our ancestors started throwing started several million years ago with our ancestors. In prehistory, throwing likely first started out as a means of a defense mechanism and did not have a guided direction, like a chimpanzee throws things today. The fact that close cousins like chimpanzees throw today (yet chimpanzees are not bipedal so cannot throw well upright) supports this notion.

Anthropologist and primatologist Richard Wrangham of Harvard University is a personal witness to the throwing ability of chimpanzees. He studied them for many years in Kigali, Uganda and dodged many an errant throw. Wrangham points out that the drawbacks of a chimpanzee throw compared to a human throw – the haphazard accuracy and the fact that they prefer to throw underhand instead of overhand – does not matter for the primary use of this physical ability: for defense. A commonly observed instinctive defense reaction by primates is to take their own excrement in their hand and fling it at an adversary.

Chimpanzees on occasion, however, had more aggressive motives to throw. One story Wrangham recounts is from a French colleague who witnessed a troop of chimpanzees hunting a mother and baby forest hog. Hogs back up to each other and face their tusks outward when threatened, a defense maneuver which can be lethal to a chimpanzee if they are not careful. When encircled, hogs – which run in the hundreds of pounds, and run fast as well – are hard prey to capture. But as this particular scene unfolded, a chimpanzee took a big stick and heaved it toward one of the hogs, fortuitously hitting the animal square in the rump. A squeal ensued, and the pigs scattered. Separated from its parent, the piglet was soon surrounded and caught. Calvin cites a similar example in *Ethology and Sociobiology*: "A chimpanzee has been seen to throw a melon-sized rock to scare away adult bush pigs who were protecting a piglet from predation."

In neither example, however, did the chimpanzees show capacity to improve their throwing ability through training or aiming. They cannot throw like we throw. But somewhere along the line, a pre-human ancestor threw a little better and more accurately. The evolutionary forces which enabled an improvement in throwing shaped both our minds and bodies through the timing of throw and through the selection of materials and creation of tools to throw, such as shaped stones and whittled wood for spears. Throwing evolved into a primary means of hunting and conflict resolution, no longer done in a stationary manner for defense but from a run, as a primary way to get food. We became throwing machines, perhaps due to or from our bipedalism.

Today, hunter gatherers all over the world throw other objects to get food. Rocks and boomerangs are used by Australian aborigines. (The bow made it all the way to the Torres

Islands but not the thirty mile strait of open water to Australia.) In Papua New Guinea, young men practice their stone-throwing capabilities for many hours a day, every day, to get better. And the Masai tribe in East Africa still hunt by throwing spears, even if the activity is largely ceremonial.

Old AI and Chess

In a 1998 article in *Scientific American*, Calvin points out that part of the definition of intelligence as we humans define it is our ability to be shape-shifters and adapt: "Versatility is another characteristic of intelligence. Most animals are narrow specialists."

Animals are narrow specialists in part because their body physiology does not allow them to get cognitive feedback that can help them improve their thinking abilities. The physiology of animals like elephants enables them to use their bodies to accomplish incredibly dexterous feats which they can use like tools. An elephant's trunk, like other animal features functioning as tools, is a part of its body. Yet although elephants can learn how to paint modern art with their trunk, elephants cannot do with their trunk what we do with our hands…including playing the game of chess.

First and foremost, chess is not a problem that would appeal to an elephant. What an elephant lacks is not the ability to make such a fine motor movement, but rather the recognition that chess pieces move in a very specific manner. To play chess, you need to understand how the pieces move. An elephant has adapted to its natural environment and its physical appendages could not adapt to give the necessary feedback in several thousand years. This is in part due to the fact that chess has been around only thousands of years while elephants, in one form of another, have been around for millions. (Modern elephants are only a surviving remnant of the once-large Proboscidea order.)

On the other hand, the human ability to play chess is a byproduct of our intelligence (and not its source). Our high intelligence is in turn derived from the necessary skills it provided us in our natural environment – the same environment elephants occupied. Our intelligence was our survival strategy. Because we can already do so much more with our hands, our multifaceted use of these parts of our body open up the possibility for yet more uses, such as moving a chess piece. With embodied cognition, this general-purposeness extends to our brain and we figure out yet more uses.

Animal physiology can also do more than just create crazy body part shapes. It can also enable specialized movements that occur on a time scale much beyond what humans can accomplish. The lunge of a frog fish is an amazing physiological feat, able to lunge from a camouflaged position on the seafloor to snatch a fish in only 6 milliseconds, whereas the fastest human voluntary motion is approximately 225 milliseconds. (The fastest period a javelin has been released, created by an involuntary reflex set up by the throw, is around 160 milliseconds, an extremely low number accomplished only by the reigning world record holder Jan Železný).

Yet the impressive lunge of the frog fish is, from a cognitive standpoint, an evolutionary dead end. Its physical prowess has not enabled a frog fish to possess anything we would remotely call intelligence. The frogfish is well adapted to catch unsuspecting fish, vastly better than a person. But that single purpose, one-time adaptation has not given rise to something even better. Animals have fast reflexes but aren't smart; their reaction times have not led to a progression in intelligence. Evolution for animals has meant a physical body change that could not enable further changes that cognition gives us. William Calvin writes that although "the escape reflex of fish and crustacean illustrate one way to orchestrate a rapid multi-muscle sequence", most animals evolved uses of their body that are part software and part hardware, part learning part instinctive, and the physical limitations that resulted from such a combination of (a lot of) brawn and (a little) brain have not given rise to the cognitive abilities of humans. The toolmaker path of humans gave a premium to a general purpose ability, cognition, which in this case would be developing a spear gun to catch a fish.

Humans are unique in that we are predators using artificial weapons that we construct as a tool. Without a tool, we were not fearsome. We developed tools deliberately because it was a necessity. That occurrence created a gateway effect. It opened the possibility of evolving better muscles for better spear use. In turn, we also evolved a better brain which resulted in better spear designs. The utility for using a spear as a weapon is this that the tool itself is adaptable on a time scale much shorter than that of physiological evolution. You can make more progress inventing a new spear, or a spear throwing device such as an atl-atl to throw it better, than evolving your body.

On the inverse side, dolphins possess what we perceive to be high levels of cognition but inadequate physical features to fully utilize those features, which are shaped by echolocation and a lack of hands. Their cognitive complexity has not allowed them to become toolmakers. Ravens, some birds and chimpanzees do use tools but their use is quite limited.

However, Calvin's versatility explanation in itself is not good enough, for it does not differentiate us from other primates like chimpanzees, who are more intelligent than most other animals but fall well short of human accomplishment. To do that, writes Calvin "To understand why humans are so intelligent, we need to understand how our ancestors remodeled the apes' symbolic repertoire and enhanced it by inventing syntax."

How Throwing Influenced Cultural Evolution: The Toolmaker's Arc

William Calvin's stone throwing theory, despite its originality, is not without challenges. Calvin himself acknowledges that human evolution is still very much an uncertain science. Evolutionary psychologist and linguist Steven Pinker argues that since humans derived their advantage from being adaptable and generalists, any one motion such as throwing is too specific to be selected for. Rather, a whole host of behavioral traits arose from collaboration and communication. Throwing and making a spear are results of that process but Pinker believes they are not indicative of their origination.

Steven Pinker's ideas add a caveat to the evolutionary explanation that throwing was used as a means of acquiring or amplifying language. Given the available evidence Pinker believes that throwing is a manifestation of evolutionary advantages, and likely not the cause of them due to its over-specificity. According to Pinker, fundamental physiological changes that differentiated us from our ape ancestors, such as human endurance, were more important to hunting strategy than any one tool such as a spear, or one motion such as throwing. Pinker outlines three key traits he believes are important to human evolution: technological process ("tool making et cetera"), sociality (what he terms "nonrelative cooperativeness" – the ability of humans to be altruistic and have symbiotic relations with strangers, something Pinker says other animal kingdom denizens do not do), and finally language ("talk is cheap, and language is a facilitator of cultural evolution").

The evolution of language, says Pinker, is particularly problematic for Calvin's neuron bundling idea because spoken words in language arises not just from where things are located physiologically. Rather, language originates from a cultural element that throwing in and of itself is unlikely to account for. Sounds in language are either all or nothing. There is no in-between sound that throwing, or any other physical feat, could help facilitate. In this regard, Pinker says that language is "digital" and not analog. As an example, Pinker points out that "there is no sound halfway between "gat" and "bat" at the beginning of the word". So a physical ability like throwing could not have been a catalyst that instigating the ability of humans to make gat from bat or vice versa. The expansion of language would have to have some other origin than throwing.

But however the evolution of phonetic sounds in language came about, throwing in itself was not done in isolation. That language could have developed from activities derivative to and a prerequisite from throwing, such as communication and inherently social activities that the process of stone selection and shaft-shaping required for throwing implements, is an open possibility given the available evidence.

Besides, as Rodney Brooks points out, Berkeley researcher George Lakoff has pointed out that our language is filled with extensive active physical metaphors, such as "we attack a problem", which could mean that the development of our language was a result of our need to describe our physicality. ("I don't remember his name off the top of my head", admitted Brooks, "but Google 'Women fire and dangerous things' and his name will come up".)

Calvin argues that throwing stones could have led us to enlarged brains, but a lesser catch-all could be that the all-encompassing behavior that throwing necessitated for hunting – the preparation of the implement, the setup, the communication – was the catalyst for driving the evolutionary changes that allowed us to proliferate. Because throwing is unique to us and integral as a major protein source before agriculture and animal domestication, it is easier to imagine an activity like throwing than, say, long-distance running, tree climbing, or bush foraging – all of which are done by other animals as well – as ones that were make or break for our cognitive capacity and ultimately our physiology. This is the toolmaker's arc.

Because as Pinker points out, an activity such as tool making also allows cultural evolution. If one inventor can communicate to another it gives them a huge advantage compared to reinventing from scratch. Throwing is part using the tool, which is not part of our body, and part using behavior that enabled the throw and evolved alongside it: tool making. Throwing involved making the spear as well as throwing it. Toolmakers are designers and the path that arose from throwing occurred through tool making.

Besides making the spear itself, other adaptations such as the atl-atl arose. The atl-atl is a hand-held counterweighted arm extension made out of bone or wood. A spear attached to the end of it, allowing a spear thrown with an atl-atl to travel faster and farther. An atl-atl works because it adds weight and arm extension and velocity to the spear. The momentum transfer from a heavier to a lighter object accelerates the latter. You would never invent the atl-atl unless you threw spears. Throwing spears put you on a path where the notion of being an inventor has a tangible result. Without these tangible results there would have been no need to select for innovation, and this includes cognition. If our brains and bodies developed together each could be a limiting factor for tool development. However, given that we could create tools to help fashion other tools, our mind above all else set us apart.

Atl-atls are considered to be so effective – based on tests today – that they have been singled out by scientists as the main culprit for the demise of a long list of Ice Age megafauna: the giant ground sloth, the woolly mammoth and all of the modern day elephant's other ancestors. With atl-atls, humans used throwing to influence their environment. Across Asia and the Americas, the atl-atl shows you what happens when you become a toolmaker.

In this regard, once you become a spear thrower, you become a technologist. This progression of a physical movement leading to tool development (through cognition) can be called the "Toolmaker's Arc": how, once we learned how to aim and throw stones, we progressed to where we are now, enabled by our cognitive ability and ingenuity. If we did not have what we needed, we could think it and make it. Tools led to tool making, which in turn led to tool designing and ultimately cognition. Hunter and gatherers and farmers became inventors and engineers, allowing specialization in cognition and tool making to continue to today through mathematicians, musicians, and novelists.

This progression in some sense sums up the specialization of Western civilization. In the words of John Adams, "I must study politics and war that my sons may have liberty to study mathematics and philosophy. My sons ought to study mathematics and philosophy, geography, natural history, naval architecture, navigation, commerce and agriculture in order to give their children a right to study painting, poetry, music, architecture, statuary, tapestry, and porcelain."

Once evolution is acting on something that includes a tool, the possibility arises that making the tool is as important as using it, leading to the evolution of invention and engineering. We only have one example of toolmakers – humans – but we know throughout archeological history of mankind that there are various breakthroughs in making hand axes, atl-atls, bows; technological innovations that did change our behavior.

Previous Biomechanics Work

Other researchers have used biomechanics to determine what behaviors were most important for human evolution and their results would also seem to downplay the implications of embodied cognition and throwing, for different reasons.

Dan Lieberman, Chair of the Human Evolutionary Biology department at Harvard University, has spent his adult life trying to figure out "why the human body works the way it does." Anthropologists-turned-biologists like Lieberman would say that we already know enough from existing biomechanics experiments to infer what we were like way back when. In Lieberman's view, any effort to understand the current physical structure of humans – our muskoskeletal version of a robot – needs to first focus on behavior. If we can hypothesize the behavior of our human ancestors, the logic goes, then we try to see if the function of the fossils we find match what we would expect for the infrastructure to be consistent with the supposed behavior. The idea is that behavior influences the function of our physical structure. To a certain extent, all fossils were shaped by the same physical forces and physics that we experience today, so if we can determine force relationships and scale based on skeletal structure (through the fossil record), we can get some inference as to what our earlier ancestors were capable of. As Lieberman succinctly sums up, "the laws of physics are the laws of physics". Bones with the same shape are subject to the same physical forces regardless of their ratio difference with bones in other animals.

Yet in a strange way Lieberman's approach, which his mentor David Pilbeam (also a Harvard Human Evolutionary Biology professor) classifies as a "behavior-function-structure" approach, is exactly the opposite of what Brooks and modern day artificial intelligence tells us about the human body. Brooks and others prescribing to the model of embodied cognition believe that our body structure provided feedback to the cognitive side, in turn changing the function of our body through behavior. Starting with the behavior is arguably analogous to the sort of AI symbology of yesteryear, because there is little initial input from the world in which our ancestors lived.

Lieberman has published work arguing that foot strike patterns indicate that humans evolved to be great long distance runners based on measurements of the exterior of the foot done on treadmills while running. This is contrast to Calvin who believed that a specifically coordinated motion – that would not take much time – is responsible for our evolution. One of Lieberman's graduate students, Neil Roach, is now working on comparing and characterizing the way in which humans and chimpanzees throw, and how they differ from each other. Yet the research of Lieberman and the modern anthropologist's school of thought is not really an impediment to embodied cognition or the pioneering neuroscience work of Lieberman's predecessors like Calvin for amongst Lieberman's skeptics are the very scientists whose rules he tries to invoke: the physicists.

How Javelin Flight Can and Can't Help Us

How the throwing motion evolved and its connection to the development of our cognition could be determined by examining exactly how we throw. We would finally observe the smallest resolution of an actin or a myosin (equivalent to determining the molecular structure in chemistry to understand a compound). Knowing throwing down to the smallest detail would also likely reveal the underlying neurological mechanisms responsible for coordinating the muscle movement with our impetus to do so in our brain stem, which while not responsible for our conscious thought is nevertheless an important part of our brain.

While we know javelin flight with a certainty, that same knowledge does not extend to how we throw the javelin. The properties of a javelin in flight are a well understood fluid mechanics calculation ("a javelin does not have a varying internal structure in time"), says physicist Lowell Wood, one of the government's top scientists at Lawrence Livermore National Laboratory in the 1970s through the 1990s.

But because we cannot see the internal degrees of freedom in our body, the throwing of an object like a javelin is too complicated to compartmentalize into pieces we can understand and analyze. We know little about what happens in our body when we throw, because simply put, right now we cannot instrument a human body, says Wood. Muscle fibers, unlike the tools and contraptions we build, do not have gauges on them. In the last decade, the best we have been able to do is measure a single actin fiber. Trillions of fibers in a muscle twitch and react on a time scale of milliseconds. We cannot possibly sense these movements. And our ability to measure with instrumentation is below the resolution to see the process at its smallest scale, which is the necessary unit for us to reliably understand, scientifically, how we throw.

In order to get the understanding that Lieberman purports to have – even if he acknowledges the need for more research in his area – Wood says we need to get inside muscles to see what is going on and eliminate the possibility of counteracting muscles that can obscure what is going on. Wood believes that ultimately, if we want to figure out how throwing works, we have to "build a robot" to do it.

Building Robots to Reveal Human Intelligence

Building robots is exactly what Rodney Brooks has done. His goal was to try to build a robot and see how the physical design of the robot grew with its intelligence, using what he knew about the development of humans, particularly the interaction between mother and child. Brooks' idea is that we can't understand cognition until we understand how the physical structure influences it; it is a mistake to look only at the brain without seeing the body. The two go hand in hand, in evolution.

Brooks has developed a variety of specialized robots designed to do tasks that demonstrate how cognition can arise from modifying the physical infrastructure and "brain" of the robot in tandem. In each case, Brooks has demonstrated that a relatively simple set of rules – governed by the robot's sensing its environment, as opposed to a predefined input/output which he terms the "symbol system hypothesis" – can result in behavior which we deem to be intelligent, (like being able to find empty soda cans on randomly rearranged desks and pick them up. "Graduate student Jon Connell "did all the work", Brooks is quick to point out, "I just supervised, so he deserves the credit.")

The neatly kept clothes that Brooks wears is interrupted by small wild curls of gray hair now tamed to just below his neck – but that in past pictures in the 1980s and 1990s, flowed below to his upper back – and his green eyes stayed focused on me as he animatedly described his research, driven by the phrase *ontogeny recapitulates phylogeny*. There seems to be a continual sparkle in his eye and he pauses for several moments to recall what he says was a different time period in his life.

"A bunch of my grad students and I looked at, almost in a hand-wavy way, how animals had developed over time. What modules could be built on top of other modules? We also looked at child development, to try to figure out what modules there might be to then carve up a modular construction of an intelligent robot, so we could build it. We did that with a robot called Cog, and then another one called Kismet. Cog was more neurally inspired; Kismet more socially inspired. We tried to build pieces on top of pieces that make things work that looked similar to how humans work. And the test was to see how humans interacted with it and measure the response of it."

A key example evident in many of Brooks' robots was demonstrating motor control of the eyes. As an inspiration for his robots, Brooks looked to the development of gaze in humans. Young children, says Brooks, can only focus on their mother and cannot look around elsewhere. ("What's the mother looking at?" Brooks pretends to peer around and look over my shoulder as I sit across from him.) Over time, as the mother looks elsewhere, the child follows. Eventually, not only can the child focus on what is visible in the gaze of the mother, but they can learn to anticipate where the mother will look, including turning their head around when the mother is looking behind the child.

Brooks said that this progression of learning can be broken down into a series of developments, and that modular pieces in a robot could mimic that pattern of progression. In fact, that is what he and his students did. At first, they did not have much available computational power for the robot's "brain", only "16 MHz in the early days of Cog – Motorola 68332s", "so [we] couldn't afford to do much vision, so we had two cameras, one above the other, one with wide angle and one with a narrow lens." "The lenses reduced the number of pixels to see the whole scene and look in detail at a portion", 'zooming in' on an embedded, lower resolution space.

Ironically, Brooks pointed out that this limitation was analogous to how we visually interpret the world. Because we have to pay attention to a lot of things, we time share by physically moving our eyes, only focusing on a small area of our total vision each moment. We trick ourselves into thinking we focus everywhere because our eyes can change focus quickly. "Your view of the world is delusional", says Brooks, because "you've got that *hole* moving around that you're not even aware of!", referring to the uneven concentration of photo receptors in our eye architecture that results in a "blind spot" on our eyes.

This focusing deficiency is apparent when we look at a photograph, because a camera like us only focuses on a given region of the field of frame. It is also apparent when you draw two objects on the back of a piece of paper and focus on one while zooming out, as Brooks demonstrated. Focus exclusively on one of the objects while closing one eye, at the right distance away from you (depending on the distance between the two drawn objects on the piece of paper), and the other one will disappear from the periphery of your vision.

"Because of our hardware we don't have Moore's Law, we only have a certain amount of stuff you can carry around in your head, a certain amount of visual cortex. It would have to be enormously bigger to do high resolution everywhere. Half of our resolution is [contained] in a thumb nail", said Brooks, meaning the angular width you can see in your thumb with your arm outstretched.

Brooks initiated this visual attention mechanism in robots in 1995, continuing through 2003 and 2004. Over that time period, Rodney's robots were able to get people to automatically look in the robot eyes by placing the moving cameras in an arrangement on the head that looked distinctly like a human head, creating 'robot eyeballs' like Kismet had.

One of Brooks' graduate students, Brian Scassellati, extensively documented how the visual viewing system of one of Brooks' robots, Cog, was modeled to mimic a human's. In a paper titled "A Binocular, Foveated Active Vision System", Scassellati lays out the goal for the creation of Cog:

> *The Cog Project at the MIT Artificial Intelligence Laboratory has focused on the construction of an upper torso humanoid robot, called Cog, to explore the hypothesis that human-like intelligence requires human-like interactions with the world (Brooks & Stein 1994)... In designing a visual system for Cog, we desire a system that closely mimics the sensory and sensori-motor capabilities of the human visual system...The system should allow for simple visual-motor behaviors, such as tracking and saccades to salient stimuli, as well as more complex visual tasks such as hand-eye coordination, gesture identification, and motion detection.*

In the interest of time, money, and the available technology in the late 1990s, Brooks had Scassellati focus specifically on two traits of the human visual system. One was "wide field of view", used for perceiving everyday objects in the environment and correcting for 'egomotion', which is camera movement relative to the reference frame. The other trait was "high acuity", used for making fine 'motor movements' and recognizing behavioral cues such as gesturing and the recognition of unique individuals (i.e., 'facial recognition').

Just like the robot construction had sensory limitations due to the technological cost of the robots' parts, so do the human eyes with biological cost and maintenance. In particular, points out Scassellati, the distribution of photoreceptors in our eyes is not even. Instead, they are positioned to give us the greatest visual cues from the fewest amount of receptors possible.

In the case of humans, the receptors for the eyes are classified as rods and cones. Rods can only 'see' in black and white and are concentrated on the edges of our eyes. They are responsible for detecting motion and seeing thing from the corners of our eye. Rods, for example, are concentrated at the edges of our eyes while the more expensive color-seeing cones are concentrated towards the eye's center.

The trick is to make a robot that tries to best model this array of sensors to get as close and complete a picture of human vision as possible – which, by the way, is still flawed vision with a different set of limitations than the design of our eyes, but limitations nevertheless. The idea is that if a robot can see what we can see, it may have a chance to develop traits of cognition similar to our own sense of intelligence. "We're machines with limitations", as Brooks puts it.

Brooks and his graduate students draw upon the research of eye experts who classify human eye movement into five main varieties. Three are voluntary: tracking a fast or slow moving object (saccades, smooth pursuit) across the horizontal plane using your periphery vision, as well as tracking an object with a changing depth of field (vertical movement, called 'vergence'). Two are involuntary: keeping your eyes locked looking in the same place despite head movement ('vestibule-ocular reflex' and the 'optokinetic response').

To create Cog's visual system Brooks' and company used four small black and white cameras called an Elmo ME411E. Each operated at 12 V and 3.2 Watts. Since "color is three bytes per pixel, whereas black and white is one byte per pixel, when you use a 200 MHz CPU it saves computation time. You have to play tricks to operate in real time", says Brooks, otherwise the image feedback the robot gets from its cameras will lag and everything will be going slowly (frustrating the human attempting to interact with it.) The cameras were oriented in a way to focus in on an up-close object as well as spot surroundings from the periphery of our vision. "We tend to have one camera fixed to get the wide angle and the moving cameras would be in the eyeballs."

However, we sacrifice color to see things out of the corners of our eyes and there is some distortion in doing so. The tradeoffs our vision system uses to see far and wide are akin to a fish-eye lens on a camera which makes similar tradeoffs for the benefit of capturing a wider, farther flung field of view. And just as Brooks' cameras have "blind spots" where the robot cannot see some of its immediate surroundings, so do our retinas as a result of a lack of receptors in certain regions. To compensate for the loss of that would-be data we use the surroundings to interpolate what we "think" would be there. Scassellati describes as "filling in" blind spots. These details of these examples of visual limits are still unknown and are not addressed in the works of Brooks ("Until we have a complete understanding of biological (or artificial) machine vision systems, there will always be exceptions to our models.") In fact, these exceptions are well known to us and have been actively exploited for centuries by the numerous visual illusions employed by magicians all around the world.

Brooks was surprised that it did not take much to convince humans that they were talking with something intelligent, even though the only visual cue was human-like eye movement and the intermittent sounds the robot made. Brooks calls this "social interaction without content", and likened it to what happens between humans as well, "The 2 a.m., very drunk in a bar interaction. Who knows what's going on back and forth!?"

But we do know what is going on in the robot's eye architecture which allows its CPU visual attention mechanism to work. The eyes of Rodney's robots, particularly Cog and Kismet, are a good example how embodied cognition manifests robot behavior which human participants interacting with the robot deem to be intelligent.

Graph segmentation algorithms, in particular the ones developed by one of Brooks' graduates students Paul Fitzpatrick, and other relatively recent computer science theory advancements in the 1980s and 1990s, dictate how the robot eyes can perceive movement and initiate a physical response to it. (Which in the case of robot-human interaction, instigates further movement and thus an interaction due to the feedback).

Here's how it works: a person looks at the robot to get its attention, then moves an object, such as a phone, down to the table. The person has to watch robots' eyes. The robot looks for a high motion cue, such as waving back and forth, that a person iterates and the robot tracks. The robot movement it turns triggers our eyes through saccade, eye motion darting back and forth, which forces the person to "latch into what the robot is doing". As person puts phone down, the robots' eyes go to it. When motion stops, robot figures where it last looked, is where the object is. The vision is tracking motion at all times.

The way the algorithms are set up in its CPU, the robot only models an object when it's flat on the table; otherwise, it just detects motion. This shortcut saves lots of computational time and effort and achieves the same result as if the robot knew that the object was in fact an object all along.

Once the object in question reaches the tabletop the robot only knows of its precise location, and what it is, from its previous movement that got it there in the first place. On a wooden table, the color contrast (even in grayscale if the cameras are black and white) can make it difficult to determine what is the object versus what is the table. This limitation is from the architecture of the robot cameras which, as one may recall, mean that the robot lacks depth perception and depth of field. For example, Brooks says that when presented with a "yellow fluffy cube for a kid with different patterns on the side", Cog could not tell the toy from the table because the contrast between the wood table and toy was not high enough.) So instead, the robot pushes along the table until the object moves to detect it. This work-around is as consequence of Cog's ever evolving processing power. "The tech was always changing so were always updating." "Over ten to fifteen years of the Cog project, processors increased by factors of hundreds, so then we could use colors and more sophisticated algorithms. We couldn't push and segment from the table early on. It would have taken *minutes* of processing time."

Early on, in the early 1990s, instead of instructing the robot's algorithms and transferring them via wires or wireless, Brooks and his students burnt read only memories (ROMs), a chip window in the center containing an ultraviolet eraser. Information on these chips could be downloaded, compiled, and transferred to the robot by plugging the chip into the processor board of the robot.

A similar conundrum existed for skin tone versus the plywood in the new building Brooks moved to while at MIT. "One of the things we did is have it look for skin tone; the color of human skin, if you map to HSV [hue saturation value] instead of RGB [the standard color space used]. Humans have a narrow range of skin tones within this space." In Brooks' old building at Tech Square, the robot's cameras "were really good at picking up spaces. But plywood maps plot to the same spot in HSV space, so human face recognition and skin detector picked up lots of human skin because it mistook it for plywood."

In the twenty years since the publication of *Elephants Don't Play Chess* Brooks has developed dozens of robot designs. New ones are still being developed under Heartland Robotics, a company he founded in 2008, where he is Chairman and CTO. But while Brooks' robots have been useful to us in performing perfunctory tasks, the big breakthrough to cognition has not happened yet. To date, Brooks has little to show for his efforts to evolve cognition in robots. So far, the main tangible commercial result produced from the efforts of Brooks' Lab at MIT has been the popular vacuum robot known affectionately as the Roomba.

AI and Chess Again

If the new paradigm in artificial intelligence still has not led to any big breakthroughs, how is the old one faring? In the past two decades computers have been used to completely solve the game of checkers and beat the world's best humans at chess. Six years after Brooks' "Elephants Don't Play Chess', a program called Deep Blue "played" world chess champion Garry Kasparov in a best of five game series in 1996. It lost, but in 1997 was able to controversially win a rematch after Kasparov misremembered a crucial move in the opening game of play.

Deep Blue relied on what was called "brute force"...and humans to move the pieces for it and interpret its analysis. What Deep Blue lacked in direct feedback, which could over time be developed to evolve intuition according to the physical grounding hypothesis, is more than made up for in computational power. Deep Blue is able to sort through millions of moves in a fraction of the time that a human can, evaluating possible futures and comparing them to a rich database of previous high level games between chess masters to determine what move to make next.

If you told artificial intelligence researchers working under the old paradigm in the 1970s and 1980s that in 1997 a computer would beat a human in chess, they would tell you that we would have solved AI in doing so. In other words, the only way for computers to beat humans in chess is if they developed the same cognitive ability that humans used to play chess. But it turns out there was another way to do so, that involved millions of computations per second due to Moore's Law that even the most optimistic futurist (besides Moore) would not have predicted: brute force by analysis.

But that was completely wrong, for it turns out computers cannot play chess either. To be fair, what is going on in Kasparov's head is nothing like what is going on in Deep Blue's. Not only could a computer like Deep Blue not make the moves on the chess board or choose to play on its own accord, but its supposed intelligence was based on a vast database of previous moves in games between grandmasters (chess experts) and their relative outcomes. So Deep Blue could not strategize but it could spot a mistake.

Additionally, computers such as Deep Blue were constructed with an extremely special purpose, so even though its feats were impressive it lacked the alacrity to use its talent for something else, on its own. We know from studies of cognition that the neurons of the brain fire much less quickly than Deep Blue can run through positions. Deep Blue may be able to win, yet that does not mean it is mimicking cognitive processes that humans use to play chess.

After the match was over, Kasparov claimed that the computer had cheated, or rather, the humans controlling the computer had, by updating and changing software architecture behind Deep Blue's "brain". IBM denied the accusation and quickly dismantled the computer, to attempt to re-tool it in the medical field for drug discovery. We have not heard much from it since, reinforcing the notion that the computers designed for a specific task cannot often be successfully repurposed.

This silence, and the lack of a Deep Blue breakthrough for activities other than games such as chess (checkers was solved in 2007), demonstrates that there is a big difference between an isolated brain without physical feedback – ala Deep Blue – and brains connected to bodies, or CPUs to robots, that can learn from interacting with their environment. And it is in the latter case where embodied cognition is so crucial.

Since 1997 numerous other computers have played humans. All have won. Deep Blue considered almost all possible future chess positions by going through millions of positions per second to beat the world champion in 1997, but it cost millions of dollars in custom hardware. Now it is much easier to do so. Fritz 13, a commercial computer program that can be run on an ordinary laptop, can beat nearly any human, by doing the same sorts of computations that Deep Blue did, and is available for purchase on Amazon.com for $43.95. Soon, the same capabilities may be offered on even smaller devices such as a smartphone.

Thus the failure to adopt the physical grounding hypothesis in artificial intelligence has not stopped computer analysis from beating human cognition at chess, but it has stopped chess playing computers from changing the way we think about building robots. The physical context that drove the long process of brain evolution which led humans to play chess the way they do has not been replicated in several decades of curious thinking and electronics tinkering. Physical context is crucially important to why humans can play chess remarkably well, exemplified by the radically different way in which humans and computers play chess.

"Just-So" Stories and the Future

Just like robotics has had difficulties fulfilling their promise, so have explanations for human origins that have been proffered by anthropologists. A skeptic could say that all of these stories about how humans evolved are really complicated just-so stories, like Rudyard Kipling's "How the Leopard Got Its Spots". And in fact whether it's the physicality or throwing, we do not totally understand the mechanisms behind what we are able to do with our bodies. We need this understanding our body today because without it we cannot characterize how our body has changed from the past. Since we only have fossil bones to go on – the hardware – complete knowledge of today's software is our best bet for inferring what sort of software could have been possible before.

A theory in science is something that has predictive power, not just a description of an existing framework. Yet all theories start with what could be criticized as a "just so" story…and an idea or belief that needs to be tested. Charles Darwin didn't know about genes, even though Gregor Mendel was a contemporary, and Mendel didn't know about the exact chromosome advancements that Thomas Hunt Morgan pioneered. James Watson and Francis Crick finally described the structure of DNA, but at the time, in 1953, many thought it was much too complicated to figure out the exact details. Yet in 2008, Watson became one of the first people to have their entire genome sequenced as part of a $1.5 billion dollar project. In 2012, human genome sequencing can be done (in volume) for several thousand dollars apiece.

Thus genomics, says Boston University anthropologist DeSilva, is a legitimate contender for how we would figure out our human origins. Genomics will offer an independent test that can further constrain the wide range of possibilities to how we evolved and the extent to which our physicality played a role in our development of advanced cognition. DeSilva points out that we already know work in genetics that small genetic changes can radically alter the pelvis, which has important implications for how our hips, when modified from our predecessors, enabled us to become bipedal.

A knowledge of how muscles evolved could tell us more about embodied cognition because we simply do not know muscles have evolved to be able to fire, in sequence, the way they do when we move our body today. DeSilva says that genetics work to uncover human muscle-building mutations, such as that of myosin, can help us determine exactly how our muscles are made, in turn helping us figure out how they function. And since muscle movement is the manifestation of our cognitive ability (albeit not conscious thought), if we can figure out more about how muscles operate we can get a better idea about how we "control" motor control.

When it comes to our brain, however, we have much more work to do. William Calvin points out that although foundations such as Paul Allen's Brain Institute have been pouring money into neuroscience genomics, the neuron resolution is still at the "cortex level" and thus unable to dig down deep enough to test his hypotheses. Analogous to computers, our skeletal structure is hardware and our minds software. Ultimately, says Calvin, we are limited to looking at only the hardware of our past, and from that alone it is hard to infer the software; our brain and our cognition. Much work has been done looking at the inside surfaces of brain casing as a way to infer evolution, but by examining even a perfectly preserved case we will not unlock how we are physically and cognitively wired.

When I asked him about Brooks' work, Calvin said he thought that robots like those Brooks has created could offer future insight into our cognition. From an embodied cognition standpoint, this makes sense, as Calvin wrote in *Scientific American*: "In both its phylogeny and ontogeny, human intelligence first solves movement problems and only later graduates to ponder more abstract ones. An artificial or extraterrestrial intelligence freed of the necessity of finding food and avoiding predators might not need to move—and so might lack the "what happens next" orientation of human intelligence." However, Calvin remains slightly skeptical of Brooks' validation of human-mimicking robot behavior because he believes it may reveal more about our intelligence today than what we have been able to replicate in the robots. Robots (such as Cog and Kismet), say Calvin, are solving the problem of human-robot interaction by highlighting the adaptability and intelligence of us more so than what we have created. The fact that a robot can mimic some our behavior "like scratching their head a minute after you scratch yours" goes a long way in appearing to be human to us, is a hallmark of our intelligence and ability to anthropomorphize our man-made creations.

However, the evolution of the eyes of Brooks' robots are in concordance with what Brooks says would be the impetus for our eyes developed as humans. "We have to have a moving object to see", says Brooks, so "smoothly tracking an object" (smooth pursuit) – required for tracking a moving animal you would throw a stone or a spear at – is on the surface a plausible explanation for how we developed the depth and timing sides of our mental acuity. "It would make sense"…we would want to throw at "something we would want to chase or run away from, eat or be eaten", says Brooks, laughing.

Today, we are engineers instead of apes. We do not rely on human strength anymore and hence the importance of our physical prowess has diminished. What was our strength at the beginning, our physical body, has now diminished in importance thanks to using tool using. We use leverage but we do not contort our bodies to gain this leverage, nor do we gain it by running, throwing or using an atl-atl. Modern tools to gain leverage are much more powerful (winches, bulldozers, cranes). If we are serious about it, we do it with a machine.

We have even outsourced elements of cognition that are needed to throw, such as rhythm, used to ensure that the javelin is accelerated upon release. Our timing problems, for the creation of electronic music, drum synths, and even how we shift gears with a car are done for us (and the limitations for all of those now are how fast we react or think instead of a limiting factor of some physical feature).

In warfare, the javelin is now the name of a shoulder-mounted anti-tank missile. It is not primarily known as a primitive projectile in the lexicon.

Thus we have overcome our own biological limitations by becoming toolmakers. We have come up with an amazing set of technologies in a short time: an elaborate calendar to quantify our chronology and atomic clocks that are so accurate as to be irrelevant for most of our daily lives. We could not see super small objects with our unaided eyes until van Leeuwoenhoek invented the modern microscope. We could only see a fraction of the cosmos as stars in the nights sky, until Galileo used his first telescope (and got in trouble for it). The physical perception of our tools have become better than that derived from our body and senses.

Humans somehow went down a general purpose path which later set the stage for cultural evolution. In our modern world, cultural evolution and the evolution of our technology is overtaking physical evolution. But at the same time no matter how abstract and highly evolved we get, the fact is we are still shaped in deep ways by the context of where we came from. Our biological infrastructure ultimately controls how we make our thoughts a reality, whether it's typing, talking, or throwing.

The near future of human evolution is going to be shaped by social factors, yet embodied cognition means we will still be influenced by our physical features. Thousands of years from now, if we don't destroy ourselves, human intelligence will still be shaped by the feedback that was crucial to us millions of years ago…but that today lets us hear high quality music on headphones, allows us drive racecars at 180 mph around a city obstacle course, or compete against one another to see how we run, jump or throw. Our physical feedback still influences the way we think, and have evolved to do so, over the past 160,000 years since the first modern humans walked on Earth. ~

Selected Bibliography

Peer-Reviewed

Robert J. Gregor and Marilyn Pink. “Biomechanical Analysis of a World Record Javelin Throw: A Case Study”, *International Journal of Sport Biomechanics* 1 (1985), 73-77.

Rodney A. Brooks. “Elephants Don't Play Chess”, *Robotics and Autonomous Systems* 6 (1990), 3-15.

William H. Calvin, “A Stone's Throw and its Launch Window: Timing Precision and its Implications for Language and Hominid Brains,” *Journal of Theoretical Biology* 104, 121-135 (1983).

William H. Calvin. “Did Throwing Stones Shape Hominid Brain Evolution?”, *Ethology and Sociobiology* 3 (1982), 115-124.

Other Academic

Brian M. Scassellati. “High-level Perceptual Contours from a Variety of Low-level Physical Features”. MEng and SB thesis, Massachusetts Institute of Technology, May 1995. (Advisor: Rodney Brooks).

Brian M. Scassellati. “A Binocular, Foveated Active Vision System”. MIT Artificial Intelligence Laboratory, AI Memo No. 1628, March 1998.

William H. Calvin. “The Emergence of Intelligence”, *Scientific American* (1998), 44-50.

Books

Alfred W. Crosby, “Throwing Fire: Projectile Technology through History”, Cambridge University Press, 2002.

Ian Tattersall. “Becoming Human: Evolution and Human Uniqueness”, Harcourt Brace and Company, 1998.

Jared Diamond. “The Third Chimpanzee: The Evolution and Future of the Human Animal”, HarperCollins, 1992.

Juris Terauds. “Biomechanics of the Javelin Throw”, Academic Publishers, 1985.

William H. Calvin, “The Throwing Madonna: Essays on the Brain”, McGraw-Hill, 1983.

Additional Sources Used in Thesis

In-person interviews with academics Rodney Brooks (MIT), William Calvin (University of Washington), Richard Wrangham and his graduate student Rachel Carmody (Harvard University), Steven Pinker (Harvard University), Dan Lieberman and his graduate students Neil Roach and Brian Addison (Harvard University), David Pilbeam (Harvard University), Jeremy DeSilva (Boston University), Lowell Wood (Lawrence Livermore National Laboratory) and javelin throwers Duncan Atwood, Tom Petranoff, Garry Calvert, Matt Stack. Email correspondence with throws coach Kevin McGill.

Many thanks to thesis advisor Alan Lightman for extensive feedback on manuscript drafts and conceptual suggestions and advice throughout the thesis process. My father Nathan Myhrvold and brother Cameron Myhrvold also provided feedback and this thesis has been strengthened as a result of their critiques.