2015 IEEE International Conference on Systems, Man, and Cybernetics

More A than I

Why Artificial Intelligence Isn't, But You Are

James Munis, MD, PhD
Division of Neuroanesthesia,
Departments of Anesthesiology,
Physiology and Biomedical Engineering
Mayo Clinic, Rochester, MN, USA
munis.james@mayo.edu

Artificial Intelligence has the potential to change the world. The application of A.I. to robotics, control systems, text recognition, voice recognition, and autonomous vehicles will prove both revolutionary and useful. However, humans and machines process information in fundamentally different ways, and if those differences are not appreciated and exploited, a great deal of time and money will be wasted by expecting the wrong outcomes and capabilities of our machines. This paper reviews the most critical differences between human and machine information processing from the point of view of information theory, and also introduces the concept of "information lensing" that provides a novel framework for understanding patterns of information processing that are currently unique to humans. The addition of information lensing to efforts to mimic human thought processes through a combination of neural networks, distributional semantics, and natural language computation may provide a closer convergence between human and machine information processing.

Keywords-Computational models of human cognition and interaction, Artificial Agent cognitive development, Natural language understanding

I. INTRODUCTION

We have made a magnificent tool. So magnificent that we have forgotten ourselves in the process and asked the tool to do something beyond its nature. We have asked it not just to *help* us, but to *be* us.

Arthur C. Clarke predicted as much when he observed that a sufficiently advanced technology will be indistinguishable from magic [1]. In the case of A.I., it will be indistinguishable from us. And so we have fallen, with great willingness and enthusiasm, under the spell of our own creation.

Does it matter? Isn't it better to ask too much, rather than too little, of our ever more necessary invention? It does matter, and for two reasons. First, we run the risk of inhibiting the potential of computers and robotics by expecting the wrong things from them. As a species, as clever and determined as we are, we do not have infinite resources of time, money, and energy. What is wasted in trying to use a shovel as a pen ultimately robs from the development of the shovel. Shovels are useful, but they are not, and never will be, pens. Secondly, we run the risk of misunderstanding and discounting the greatest creation in the universe: our own minds.

Computers have almost infinite potential as tools: whether for good or for evil is up to us, not to them. The problem with the A.I. misconception is not that it threatens to release some kind of destructive genie, like the "singularity", but instead, that it distracts us from making a better tool. We have a puppy to raise and to train – a puppy with the potential to serve as our companion, protector, and servant. Instead, we are so mindlessly fond of it that we have left aside the training and wasted time pretending that it is the same as us. Even better than us. The danger in this is not the specter of a self-fulfilled prophesy, but only of a disappointing waste of time and potential.

To understand this, we need to understand what computers are and what the nature of their work and capabilities are. And, in turn, to grasp that demands that we know something about "information theory". Coming to grips with the fundamentals of information theory allows us to see clearly the boundary between creature and creator, between man and machine, and to be able to make out the difference between a dead end on the one hand, and a road to adventure and fortune on the other.

II. WHAT CLAUDE SHANNON TRIED TO TEACH US

Computers and humans both deal with information, but in very different ways. This difference was hinted at in the seminal and largely overlooked work of a computer and cryptography genius named Claude Shannon, who like his friend and colleague Alan Turing on the other side of the Atlantic, helped the allies make and break ever more sophisticated codes during World War II. The same eccentric had already made important contributions to the nascent field of computer science by applying Boolean logic to machine computation – a direct outgrowth from his master's thesis at M.I.T. in 1937 [2].

Employed by Bell Labs, Shannon's skills were also lent to the solution of more prosaic problems than code breaking or developing computer science. He took on the task of analyzing the loss of information during transmission across telegraph lines – especially the long distance transfer of signals over the trans-Atlantic cable.

978-1-4799-8697-2/15 $31.00 © 2015 IEEE
DOI 10.1109/SMC.2015.425

It was this interest that would eventually define Claude Shannon and also radically redefine what we know about information. He published a two-part paper in July and October 1948 with the title “A Mathematical Theory of Communication”[3]. In the following year, and along with co-author Warren Weaver, Shannon republished his earlier paper in book form with the more definitive title *The Mathematical Theory of Communication* [4].

The little book was as thin as its initial readership, but once it was discovered and then exploited, that small volume, and the work that surrounded it, including the new field of cybernetics [5,6,7], would change the world. Shannon saw something that no one else had seen before him: information could be treated as yet another aspect of the physical world – subject to physical laws. In particular, it was captive to the Second Law of Thermodynamics, and to the concept of entropy. Shannon even coined the term “information entropy” and quantified the relationship between non-randomness, or predictability, and the capacity to contain information.

Just as Shannon’s book was overlooked initially, his first and seminal observation in that book was equally ignored. Shannon began his treatise by pointing out that he would leave aside the separate problem of “semantics” – or *meaning* in information - and instead deal only with the problem of how information is handled in real world physical systems. What the world of computer science – a world that *The Mathematical Theory of Communication* would eventually help to define - ignored was the distinction that Shannon himself made between “communication” and “information”. Shannon did not invent, as has so often been stated, the science of *information* theory. Instead, he invented the science of *communication* theory. And he was the first to acknowledge that difference in the very title of his book. From the outset, he insisted that his book was about the processes of encoding, storing, transmitting, and decoding information, as well as the inevitable loss of information in the physical systems that handle and transmit it. Together, this is “communication”. “Information”, on the other hand, is different. It is the actual *content*, or *meaning*, that the process of communication moves from Point A to Point B – from sender to receiver.

III. Lower and Higher Information, and Placeholders in Space and Time

Fundamentally, there are two levels of information. The lower level is comprised of what I’ll define as either “instructive” or “descriptive” information, and the higher level contains meaning, or “semantics”. Computers work at the lower level. We work from above.

Before we understand any type of information, though, we first have to come to grips with a fascinating weakness of the two lower forms information: they are dependent for their very existence on some kind of physical medium. This medium can be a magnetic tape, an old-fashioned computer punch card, a printed series of linguistic symbols in the pages of a book, a handwritten note passed between school kids, a photographic film, a series of 0’s and 1’s encoding a JPEG picture and stored on a flash drive, a two dimensional digital picture printed on paper; a three dimensional holographic image displayed by the action of lasers, or even the complex array of biochemical and electrical events that encode the olfactory memory, or aroma, of your grandmother’s oatmeal cookies within the trigger of your brain’s temporal lobes.

Consider just one of those examples. If the handwritten note passed between Christina and Jason eventually meets its end in an incinerator, and if Christina and Jason eventually marry someone else, grow old and forget ever having passed the note, let alone what the message said, does the information content in that note still exist? No. It is gone. While information itself is without mass, substance, energy, or any other physical quality that we can get our hands around, it cannot be created, stored, transmitted, or retrieved without the help of the physical universe. When its storage medium and its transmission vehicles are gone, it ceases to exist and cannot be retrieved or even proven to have existed.

Back to those two flavors of information that make up the lower level: instructional and descriptive. These are important domains of information because they are the only kinds that computers can deal with.

We can refine the hard condition that information depends on physical media for storage and physical vehicles for transmission one step further by understanding that all information is encoded and transmitted as a series of events that each occupies a unique position within a placeholder – a placeholder that is defined in both space and time. By “events”, I mean anything from a mathematical or linguistic symbol, like a number or a letter, to a 0 or 1 in binary code, a colored pixel, or even the more effervescent arrangement of photons of different wavelength moving through space and time in a specific pattern, or of sound waves moving as compression events in the air inside a concert hall. All of these are examples of information content being captured, held, or transmitted by discrete events that each occupies a placeholder in time and space. Importantly, those events can move between placeholders. Information is not necessarily static. It can move and change.

IV. Restriction and Exclusion

At this point, things get a little strange. Each “event” that occupies a particular placeholder is most properly seen as an *exclusionary* event. That is, the letter “T” at the beginning of this sentence occupies a placeholder in space and time, and in that place and in that time, it excludes all other occupants other than the letter “T”.

Why do we look at an event, like a symbol, as exclusionary, rather than inclusionary? Or, to use a concept championed by the brilliant Berkeley computer scientist and mathematician Lotfi Zadeh (the founder of fuzzy mathematics), why can’t we just see lower level information as “restrictive” – that is, as a narrow choice within a broader set of possibilities or probabilities? Isn’t the difference between “exclusionary” and “restrictive”, pun intended, just a matter of semantics? The answer is yes. But forget the “just”. Semantics is the holy grail of information, and as Lotfi Zadeh postulates [8], it is only through the understanding of information as a restrictive entity that we can hope to get a handle on meaning, or semantics. This plays into Dr. Zadeh's more practical goal of developing natural language computation.

While restriction is both conceptually interesting as a possible gateway to semantics, as well as amenable to mathematical analysis, it is not as interesting as exclusion in terms of understanding how humans deal with semantics.

The difference between a restrictive process and an exclusionary process is the difference between "to" and "from". We restrict a choice "to" something, but we exclude *all* options in the universe "from" our one choice. A restrictive process begins with a defined set of possibilities and then makes its choice. It can therefore be approached mathematically. On the other hand, an exclusionary process begins with an undefined, or infinite, array of possibilities, and then pushes that whole infinity of choices away in order to grant exclusive rights to one and only one choice. This is why computers can work with restriction but not with exclusion. Infinity, oddly enough, is something that our relatively slow brains can deal with that much faster, and otherwise more powerful, computers simply can't.

It is critical to understand that, once lower level information is created or encoded, it is now part of a package that computers can work with. Once created, the mechanics of storing, encrypting, transmitting, or decrypting information is indifferent to how that information was created in the first place. My iPhone is just as capable of storing and displaying an affectionate text from my wife as it is an affectionate text from my granddaughter. But the two have entirely different meanings. The mechanics of storing and transmitting are efficient but stupid. The semantics, or meaning, of the information, however is smart. It is human. It depends on a sender and a receiver who are both "sentient", or thinking.

Humans work and trade within the currency of information, just as computers do. But, unlike computers, we apprehend the value of that currency. When we do this, we are operating, as we are wont to do, at the level of semantic information.

V. Intention and Volition at Time T_0

In addition to the fact that semantics operate at a different, and higher, level than instructional or descriptive information, there are two other key differences: First, what happens at the very act of creation, at time t_0 ("t-zero"). I'll call that "*volition*". Second, what the *intention* of the creator, or sender, is. Volition and intention are not only necessary features of semantic information (in this case, the creation of semantic information), but they are also part of what defines us as human beings and separates us from computers.

The handwritten note passed between Christina and Jason cannot be dealt with properly, cannot be interpreted properly, by Jason unless he infers what Christina's intention was in creating it. Jason does not have to be taught to do that. It comes naturally to him because he is a human being. We are very good at searching out and discovering intention. When we are not very good at it, we have a problem in our relationships and in our functioning as humans.

This is part of the burden carried by people with autism, for example, as well as by the people around them who love them and want to communicate with them. They have a hard time interpreting the intentions of others, even if they can decode the plaintext of their communications. Ditto for many of us as spouses. When we fail to discern and act on intention, we sink a little lower than the level of humans (that is, the level of semantics), and become, in the eyes of our spouses, more like computers. We can decrypt the incoming code, and even act on it in the most algorithmic or predictable way, but that is ultimately both unsatisfying and inefficient.

Why is failure, or poverty, at the semantic level so unsatisfying? Computers, after all, are marvelously adept at performing faster than we do. They can even be programmed to substitute randomness for volition: that is, to imitate original and unpredictable thoughts. But is only a fake. Only an imitation, like Turing's famous game that serves to test for what is human and what is machine.

Consider football. We admire the players for their size, speed, and power. So why not save your money and your time and go to the zoo instead? A rhinoceros is bigger and more powerful than any National Football League (NFL) lineman. And a gazelle is faster than any NFL wide receiver. The problem, though, is that rhinos and gazelles, as interesting as they may be in their own right, cannot excite us with either their volition or their intentions. A large part of the uniquely human excitement of watching a sporting event, listening to a concert, or watching a debate, is what we perceive or discover about the creative thought (i.e., t_0 thought – volitional thought) and the intentions of the participants. That's where the real interest lies.

I've spent quite a bit of time at the San Diego Zoo. It makes for a marvelous afternoon, but I have to say that I was never very impressed or entertained by the intentions of the rhinos. Size, power, raw beauty? Yes. Intentions or creative volition? No. Similarly, while I have great admiration for a computer that can beat a human champion at *Jeopardy*, or checkmate a chess grandmaster, I don't want to spend a vacation week with one. It may "know" more than me – in the sense that it stores massive amounts of data, but it has no volition, no intention, and no sense of humor.

I will anticipate the objection that, like Arthur C. Clarke's quote at the beginning of this article, a sufficiently advanced computer will be indistinguishable from magic – or in the present case, indistinguishable from a human. How could that be? The false answer, inevitably, is that an advanced computer will generate feedback loops that sense human actions and physiologic states (for example, emotions), process that afferent (incoming) information, and generate efferent (outgoing) actions that mimic human responses. Wed those feedback loops, the theory goes, to machines that look, sound, and feel like real people and what you get is a robot that wins the imitation game. *Ex Machina.*

What is the problem with that scenario? Aren't we just waiting for the technology to do what we have seen technology do time and again in our lifetime – make the previously impossible possible? The answer is no. There is a qualitative, not a quantitative, wall that separates man from machine: one that no amount of processing power, no long run of Moore's Law [9], can breach. That wall is the boundary between lower and upper information – between instructional and semantic information.

The problem with the concept of a "technological singularity", where computers match and eventually exceed human intelligence, is that the imitation game, or singularity contest, is keeping the wrong score. That game is counting function, not meaning, and function doesn't get you to "human". If we treat computers as black boxes, without caring what's inside, and only assess them based on what goes in and what comes out (i.e., function), they have already achieved the goal of technological singularity. But we don't consider that as intelligence. My smart phone, along with its Wikipedia app, can already answer far more factual questions than I can. But I don't consider it to have an intellect. Even if it is eventually connected to a robot or to some other device within the Internet of Things, so that its massive data processing is wedded to physical actions and responses, I will still know that it lacks volition and intention, and that it doesn't "do" meaning.

How can I know that? Because it is the result of the very nature of information. It is the result, not only of the boundaries that we've already considered, but also of one additional and very important feature of human intelligence: what I'll call "information *lensing*".

VI. Information "Lensing"

An optical lens refracts, or bends, light. A series of lenses can "un-bend" the previously refracted light back into its original form if we construct them just right. Human beings use an analogous process of manipulating information when we communicate with each other. As I mentioned earlier, all information depends for its life on being stored or transported in a physical medium. This is as true of our thoughts, or neurally encoded information, as it is of our writings, paintings, drawings, or video recordings. All information, high or low, descriptive, instructional, or semantic, needs to ride atop or within physical vehicles – the vehicles that establish placeholders in time and space.

It is physically and biologically impossible for me to communicate with another human being without processing and transmitting my thoughts through a conduit that passes through the lower form information, through instructional and/or descriptive information. I will include linguistic information as an interesting subset, but still just a subset, of instructional information.

No emotional nuances get around this requirement. Even a sly wink of an eye cannot connect the human sender and receiver without all of the physiology and physics that attend to the muscles of facial expression, the brainstem nuclei that move those muscles, the light waves that transmit to the receiver's retina the information that a wink has just occurred, the signal transduction from retina to receiver's brain, and finally, the electrical state change that occurs in multiple areas of the winked-at's brain.

But there is more happening here than physics and physiology. The series of events that I just described are necessary, but not sufficient, for the very human exchange. On the sender end, an original volitional act occurred at time t_0, and that act enabled the sender's intention to move to and affect the receiver. At the receiver end, the lower level information (descriptive information, in this case: the visual "picture" of an eye wink) was received and processed to be sure, but the exchange didn't stop there. The important matter of semantics, or meaning, of the event had to be inferred in order for the exchange to be, literally as well as figuratively, meaningful.

This process of intention and volition being "lensed" through lower levels of information and then being reconstituted back into a higher level of information (semantics) within the receiver's brain involves more than function or functional outcome. After all, we can program a robot to wink under certain circumstances and another robot to sense and respond, functionally, to that wink. But nowhere within that robotic chain of events and responses does volition, intention or meaning reside. If the robots are suitably lifelike, can an external observer tell the difference – tell that the exchange just happened between robots and not between humans? The answer is no, but that question is the wrong one to ask. A non-participating external observer perceives function and function only, but is outside of the semantic events, or in this case, outside of the lack of semantic events, between the sender and receiver.

The more interesting question to ask is whether a sufficiently advanced robot can participate with a human in a similar chain of events – leaving semantics to the human and only imitating semantics within the robot, but without the human being able to tell the difference. The answer is yes – a sufficiently advanced robot can get away with this. For a while. Then what happens? The problem occurs when the robot cannot fully reciprocate in the "information lensing" exchange. Remember that the process moves semantics through a lower level information lens, and then reconstitutes that information back into higher-level semantics within the receiver's mind. Intention and volition within the mind of the sender beget information that moves through physical media to, in turn, produce a semantic event within the receiver. The receiver perceives and understands the intention of the sender.

As we discovered before, this one-way exchange, in isolation, can be functionally imitated to the point of winning the imitation game. But what happens as the game continues – as we wait for the table tennis ball to come back to us? Here is where machines drop the ball, and always will. They have no original intention, no original volition at time t_0. And, as thinking, sentient, human beings, we can sense that. We find ourselves playing an information lensing game that reveals, sooner or later, that, on the other end, no one is home.

VII. CAPTCHA's, Codes, and Password Hints

Ironically, computers themselves are asked to administer a sort of reverse Turing test in order to discover what, or who, is on the other end. Is it you (a potential customer) or a bot? CAPTCHA's were designed, by humans, to be used by machines to determine if you are a human. How do they do this? The obvious suite of answers only scratches the surface. It is a good example of how we can sometimes get the right result without understanding why.

For example, one rationale for the way CAPTCHA's work is that they exploit the human ability to perform *invariant recognition* – that is, the ability to deal with an almost infinite number of variations in the shape and features of letters or numbers. Computers, on the other hand, are captive to sifting through a finite set of known possibilities until a match is found. This is a case of restriction (computer) versus exclusion (human), although it is not generally recognized as such.

Another explanation for how CAPTCHA's work is that computers have difficulty separating individual symbols when the two dimensional boundaries between those symbols are distorted – the problem of *segmentation*. Looked at from the point of view of what I described earlier as placeholders in time and space, however, there is more at play here than the conventional explanation offers. That is, humans are better able to deal with distorted placeholders because they have the *expectation* of placeholders occurring in a string of symbols. By "expectation", I do not mean a simple programmed search for the next symbol in a string – the thing that computers do so well – but, instead, the very human ability to search out the *intention* of the sender.

So here is another example of information lensing: the human creator of a CAPTCHA is operating at the higher level of intention, then creating and transmitting his intended information through a lower level conduit of information (in this case, descriptive – the CAPTCHA is essentially a picture, rather than a conventional string of symbols), and, finally, the receiving human is able to reconstitute the original intention of the human sender because he or she is *expecting* intention on the sending end. In other words, the usual explanation of a CAPTCHA being a reverse Turing test is not quite accurate. Understood as a case of information lensing, the whole CAPTCHA process is really a human testing for humanness, rather than a computer testing for humanness. In fact, the test is only passed, in part, because the receiving human has an expectation that the sender can operate at the higher (human) level of information that includes intention.

The expectation of intention is akin to the well-known observation that we find what we are looking for more easily than what we are not looking for. You will discover the face of a friend in a large crowd much faster if you are looking for her than if you aren't. Human beings are also adept at making subconscious comparisons between expectations and observations. The difference between what we expect to observe, or sense, and what we actually observe I will call "alarm". At high levels of alarm, we iterate, while adjusting with each new try, a series of observational loops, sometimes including our own efforts to provoke a response or a change in what we are observing so that the observed thing moves closer and closer to what we expect to observe. When the process succeeds, the observed and expected are reconciled. If you also see that the CAPTCHA involves exclusion, rather than restriction; that it involves the uniquely human ability to exclude an infinite number of wrong answers *from* each placeholder, rather than the simple restriction of a finite set of wrong answers *to* the right answer, you've been paying attention.

Another example of information lensing is the password hint. Here, I am playing the Ping-Pong game with myself. A password hint will not work unless it contains semantic information – unless it means something to me. The hint itself is lensed into a lower level that is comprised of a simple series of linguistic symbols that are stored in a computer and then displayed to me at a later time (an example where the placeholders have moved through time but retained their original spacial relationships) when I have forgotten the password and ask for help. At that point in the game, I reconstitute the higher, semantic, level of information from the lower level string of symbols. The lensing is complete.

This type of information lensing can also occur between two different humans, of course. Ping-Pong is not usually solitary. If I want to transmit a password to someone else without having it discovered by a nefarious "man in the middle", I will lens a semantic hint to my accomplice that he or she can reconstitute from a lower level information transfer into the original higher level semantic information. This type of lensing will defeat a human bad actor in the middle who does not share the key to interpret that semantic information (e.g., "What you ate on our first date"), and it will also defeat an unwanted computer in the middle because the computer in the middle isn't even capable of *trying* to operate at the semantic level.

All of this has obvious application to cryptography. We have passed through the middle era in the science of cryptography – an era where we have discovered the utility of using computers to both make and break codes. That is an example of "I.A.", intelligence augmentation, rather than "A.I.", or artificial intelligence. The cryptography community has never pretended to have no need for humans on either side of the cryptographic chain. In that sense, code makers and code breakers have not wasted time by asking of computers what they cannot give, and they are further along than the A.I., or "singularity" community, by having avoided that false lead. Nonetheless, having refined the speed and power of the lower level information processing in the middle, the next stage in cryptography will take better advantage of the humans on either end. It will be the stage where encryption and decryption is better understood to be a special application of information lensing while exploiting all that high level information (volition, intention, and semantics) offers at either end.

VIII. We Lens, Therefore We Are: The Concept of Degeneracy

The final discrimination between man and machine involves a separate concept borrowed from information science: the concept of "degeneracy". The clearest example of this comes from the genetic code. The three-nucleotide (for example, "AGU") codons in messenger RNA are translated, or code for, specific amino acids. Interestingly, though, the genetic code is not cleanly and simply comprised of a one-for-one mapping between one codon and one amino acid. There are several examples of "redundant" codons that can each code for the same amino acid [10]. The opposite is not true, and for good reason: if one codon were associated with more than one amino acid (rather than the opposite), we would have a mess of non-functioning proteins. This latter situation, where one codon could be associated with multiple different amino acids would be called *degeneracy*, and the code that allowed it would, in turn, be called a "degenerate" code.

Inherent to the distinction between redundancy and degeneracy is the necessary concept of directionality. To tell whether a system is redundant versus degenerate, it makes all the difference which direction you're moving in.

Take the case where you win a high score during a carnival game. With that specific score, you can take your pick of any one of three stuffed animals: a hippo, a lion, or a badger. Moving from stuffed animal to score is a case of redundancy. If you give me a hippo, a lion, or a badger, I'll give you a score – the same score for each. On the other hand, if you give me the score, I can't give you in turn an exact prediction of which animal you will choose to associate with that score – which animal you will redeem for that the score.

How does degeneracy play into human intelligence? If you consider information lensing, that most human of all information processing, it is a case of both redundancy and degeneracy: redundancy during the initial step where the sender encodes complex semantic information into simpler, or lower, instructional or descriptive information; and then degeneracy when the receiver reconstitutes, or expands, that lower level information back into semantic information. Prior to this, I have been using the term "reconstitute", but now we can call it for what it is: degeneracy.

It may seem that degeneracy – expanding a little into a lot – is a violation of some physical conservation law. Perhaps a violation of a "Law of Conservation of Information"? Remember, though, that while information depends on physical media, it is not a physical entity itself. It can be created *ex nihilo* from nothing, and it can be (like Christina's and Jason's note) destroyed back to nothing. Like *A Wrinkle in Time*, the directional sequence of (1) exclusional volition, (2) redundancy, (3) degeneracy, and (4) expectation, which, together, constitute lensing, allows something real and tangible, like semantic information, to move through an intermediate no-man's land only to be reconstituted at the other end. To the extent that artificial systems mimic this sequence and can be engineered to deal effectively with degeneracy, they will mimic human information processing.

IX. Conclusion

What ultimately allows this phenomenon of lensing? The humans on either end. The fact that humans are the only entities in the universe that can operate at the higher level of information, at the level where volition, intention, and semantics play, allows lensing to occur.

What about the application of "distributional semantics" [11,12,13] and neural networks to the development of natural language computation? There are likely to be impressive strides made in this domain, but those efforts are really only an expansion of a system that is self-referencing but not self-aware. It is a case of mimicry that may achieve a powerful level of association, but association can neither apprehend nor make maximal use of *meaning*.

To bring this back to the question of A.I. and robotics is to bring us around to the observation that computers and robots have a glass ceiling but a virtually infinite amount of space in every other direction. To the extent that we expect them to think, rather than operate, we will waste time and money. More to the point, we will waste great opportunities to develop robotics that are truly amazing and truly useful because we are seeing the wrong kind of potential in them. Should robots look like humans and act like humans? Why not? That is not my objection. Instead, we should be exploiting their speed and their power without fear of a singularity for one simple reason: computers and robots can't lens, and without lensing, there will be no singularity. On the other hand, to the extent that we are able to incorporate a closer functional machine equivalent to lensing, A.I. will begin to converge on more I and less A.

References

[1] A.C. Clarke, Profiles of the Future, a Daring Look at Tomorrow's Fantastic World (Revised). New York: Henry Holt & Co., 1984.

[2] C.E. Shannon, A Symbolic Analysis of Relay and Switching Circuits. Unpublished MS Thesis, MIT, August 10, 1937.

[3] C.E. Shannon, "A Mathematical Theory of Communication," Bell Sys Tech J, vol 27, pp. 379-423, 1948.

[4] C.E. Shannon and W. Weaver, The Mathematical Theory of Communication. Champaign: University of Illinois Press, 1949.

[5] N. Wiener, Cybernetics: On Control and Communication in the Animal and the Machine. 2nd ed. Cambridge, MA: MIT Press, 1961.

[6] W.S. McCulloch and J. Pfeiffer, "Of Digital Computers Called Brains," Sci Monthly, vol. 69, pp. 368-376, 1949.

[7] J.R. Munis, "Man, Machine, and Homeostasis," in Just Enough Physiology, J.R. Munis. Oxford: Oxford University Press, 2012, pp. 141-146.

[8] L.A. Zadeh, "The Information Principle," IEEE SMC2014, San Diego, Oct 6, 2014.

[9] G.E. Moore, "Cramming More Components Onto Integrated Circuits," Electronics, vol. 38, p. 114-117, 1965.

[10] D.P. Clark, "DNA, RNA, and Protein," in Molecular Biology, D.P. Clark and N.J. Pazdernik. Amsterdam: Elsevier, 2013, pp. 62-93.

[11] M. Sahlgren, "The Distributional Hypothesis," Rivista di Linguistica, vol. 20, pp. 33-53, 2008.

[12] T. Cohen and D. Widdows, "Empirical Distributional Semantics: Methods and Biomedical Applications," J Biomed Inform, vol. 42, pp. 390-405, 2009.

[13] P.D. Turney and P. Pantel, "From frequency to meaning: Vector space models of semantics," J. Artif Intel Res, vol. 37, pp. 141-188, 2010.