The Ghost in the Machine: Structural Metaphors in the ‘Golden Age’ of Artificial Intelligence Research, 1956-1976

By Joseph Burgess Wilson

A Masters Research Project submitted in conformity with the requirements for the degree of Masters of Arts, Linguistic and Semiotic Anthropology, Department of Anthropology, University of Toronto

The Ghost in the Machine: Structural Metaphors in the 'Golden Age' of Artificial Intelligence Research, 1956-1976

by Joseph Burgess Wilson

Masters of Arts, Linguistic and Semiotic Anthropology,

Department of Anthropology, University of Toronto

2019

Abstract

This research project analyzes the language used in a corpus of academic papers from 1956-1976 generally considered to be the foundational documents of the field of artificial intelligence. Thirty-one papers were assembled and examined for evidence of the use of structural metaphors (Lakoff & Johnson, 1980), first manually with an adapted version of the Metaphor Identification Procedure (Pragglejaz, 2007), and then through key-word-in-context searches (Deignan, 2008) with online corpus analysis tool Sketch Engine. Concordance data shows that the scientists frequently used metaphors to make sense of their work. Some structural metaphors used imagery from the same source domain, suggesting underlying root metaphors (Pepper, 1972), evidence of particular perspectives that comes to constitute the academic field. Root metaphors such as A MACHINE IS A BRAIN or RESEARCH IS A JOURNEY were extremely successful in communicating non-observable phenomena between scientists. Other structural metaphors appeared briefly in the literature but soon disappeared from discourse.

Table of Contents

List of Figures

List of Tables

“Unless you are at home in metaphor, unless you have had your proper poetical education in the metaphor, you are not safe anywhere. Because you are not at ease with figurative values: you don’t know the metaphor in its strength and its weakness. You don’t know how far you may expect to ride it and when it may break down with you. You are not safe in science; you are not safe in history.”

– Robert Frost (Cox and Lathem, 1966, 39)

Introduction

There is a deep-seated suspicion by many scientists, engineers and even amongst many writers, that metaphors can be dangerous things. It is thought that they are too indulgent; that their purpose is to manipulate, to convince or to deceive. There is trickery afoot when figurative language is used. In the words of the Enlightenment scholar Thomas Hobbes, one of the greatest errors in language use is to “use words metaphorically; that is, in other sense than that they are ordained for, and thereby deceive others.… such speeches are not to be admitted” (1651, 20). For Enlightenment thinkers, devoted as they were to rationality and objectivity, metaphor could have no place in scientific communication. Philosopher John Locke similarly counselled his readers to avoid using metaphors. “All the artificial and figurative application of words eloquence hath invented, are for nothing else but to insinuate wrong ideas, move the passions, and thereby mislead the judgment,” he wrote, “and so indeed are perfect cheats” (1998, 677).

One of the curious features of these arguments, however, is in the authors’ use of metaphor to structure their arguments. Hobbes goes on to say that “metaphors, and senseless and ambiguous words, are like *ignes fatui;* and reasoning upon them is wandering amongst innumerable absurdities;

and their <u>end</u>, contention and sedition, or contempt" (1651, 30)[1]. The term *ignes fatui* is usually translated as "will-o-the-wisp," a bright, flitting light designed to distract and deceive travelers (Forrester, 2010, 612). Hobbes' argument from above appears in *Leviathan* (1651), a book structured around a metaphor comparing the State in a democracy to the great sea beast from the Bible of the same name. The cover of the 1651 edition of Leviathan has what modern scholars would call a *visual metaphor*, with a monarch physically comprised of the people from whom he gets his power. In *An Essay Concerning Human Understanding* (1998) Locke uses a metaphor of an empty cabinet to help him explain his model of the mind (Forrester, 2010, 612).

Locke's and Hobbes' lack of self-awareness is not surprising. "Our skill with metaphor, with thought, is one thing— prodigious and inexplicable," writes philosopher I. A. Richards, "our reflective awareness of that skill is quite another thing—very incomplete, distorted, fallacious, over-simplifying" (1936, 116). A more modern understanding of metaphor, based on research from such diverse fields such as psychological linguistics, discourse analysis, cognitive science and comparative literature, suggests that metaphor is much more common in human speech than is often supposed (Deignan, 2008). Metaphor is not only essential for effective communication, but also quite possibly, a necessary component of cognition itself (Lakoff and Johnson, 1980). Instead of obfuscating the truth, often metaphorical statements help to clarify inchoate concepts or new ideas. "Metaphor has been treated as a sort of happy extra trick with words, an opportunity to exploit the accidents of their versatility, something in place occasionally but requiring unusual skill and caution. In brief, a grace or ornament or *added* power of language, not its constitutive form," writes Richards (1936, 90). Instead, "we cannot get through three sentences of ordinary fluid discourse without it," he concludes (92).

[1] In this paper, when analyzing text for metaphor use, words used in a metaphorical sense will be <u>underlined</u>, in accordance with the norms of the literature in Critical Metaphor Analysis.

In keeping with Hobbes' and Locke's core philosophy (if not their opinion on metaphor), we can move towards a greater understanding of the human condition and the world in which people live by carefully examining how they use language. In science, which is commonly supposed to be immune to the manipulative power of language, metaphor analysis is especially crucial. By unveiling the metaphorical roots of arguments laid out in scientific papers and textbooks we can trace the relationships between our cultural models and our scientific models. We can also uncover any ideological biases that might have been unknowingly smuggled into our scientific discourse under the guise of metaphorical reasoning.

This paper will analysis the use of metaphor at the hands of scientists and computer engineers working during the so-called "golden age" of the field of artificial intelligence from 1956 to 1976. Much of the computer technology that forms such an important part of the modern world can trace its roots to this fertile period of discovery which leveraged work being conducted in a diverse range of fields including computer science, mathematics, cognitive science, and psychology. In recent years the field of 'AI' has grown exponentially to command prestigious appointments at Universities (Farrow, 2019), venture capital dollars (Silcoff, 2019), and cause disruptions to traditional industries and technology (Venkatesh, 2018; Fenwick, 2018). Examining the language of the scientists involved allows us to track the direction of this growth over time, and how the metaphors used (or the metaphors that were discarded) shaped this evolution.

In scientists' attempts to describe the physical world, with mathematical formulas, with models, or with words, the use of metaphor becomes a cognitive tool, a means of reasoning through the consequences of a proposed theory. Indeed, there may not be any other means of transmitting science we cannot see, or communicating experiences that are inchoate, than by relying on metaphor (Fernandez, 1974). In the literature published between 1956-1976 we see wide-spread and systematic use of metaphors as tools of reasoning and of explanation. Some metaphors became more prevalent and more powerful due to their ability to elucidate theoretical predictions, such as speaking of machines as if they were brains, both anatomically and functionally. Some metaphors fizzle out, remaining only as anachronistic footnotes in industry textbooks, such as the short-lived attempt to understand the inner workings of early computers by using John Milton's depiction of hell from *Paradise Lost*. The philosopher Max Back wrote that "perhaps every science must start with metaphor and end with algebra," (1962, 242); the goal with the present study is to track this transformation, from metaphor to algebra, in the relevant literature.

First, I will summarize much of the background literature on the role of metaphor in science and its treatment in linguistic anthropology. Second, I will outline the methodology used to analyze the linguistic corpus and pin down some definitions for further use. Third, I will examine a specialized corpus of scientific texts taken from the period in question, hunting for metaphors and untangling their entailments.

Theories of Metaphor

"An apt metaphor suggests directions for experiment. The results of experiments in turn are interpreted in terms of an elaborated, improved metaphor or even a new one. At some stage in this evolutionary process the initial metaphor has acquired sufficient complexity to be considered a model."

- Theodore Brown (2003, 26)

The popular understanding of metaphor is essentially the version that is espoused in a typical North American high school English class: metaphor is a decorative trope, similar in form to simile, that springs from the minds of talented poets and authors. This explanation has its roots in the writing of Aristotle, who described metaphor and simile as being two versions of the same phenomenon, differing only in the "form of expression" (Garrett, 1406b). Although metaphor might be aesthetically pleasing, it is essentially superfluous to goal of communication, an ornamental feature superimposed onto language. "Metaphor consists in giving the thing a name that belongs to something else; the transference being either from genus to species, or from species to genus, or from species to species, on the grounds of analogy," Aristotle writes in *Poetics* (Garett, 1457). The word itself is based on the Greek word *metaphorà* meaning *moving* or *transfer*. "Transfer is implied, perhaps unfortunately, by the etymology of the term metaphor, and it is a study of transfer or comparison that traditional studies of metaphor have developed," writes David Miall from the perspective of English literature and film studies (1982, 89).

This model of metaphor is often called the *substitution model*, which replaces the name for one object with that of another object that is similar in some way (Black, 1962, 33). The *comparison model* goes one step further and suggests that a reader may hold two concepts in their mind simultaneously and will seek to try and compare their attributes to see if there is a congruence of properties (Black, 1962, 37). "The standard approach to metaphor comprehension treats metaphors as comparisons that highlight preexisting but potentially obscure similarities between the target and base concepts," write cognitive psychologists Brian Bowdle and Dedre Gentner (2005, 194). This model assumes a 1:1 isomorphism between features of both sets of concepts in the metaphor.

I. A. Richards explains that a metaphor is composed of two component pieces. The *tenor* is the object being discussed, and the *vehicle* is the figurative term that is borrowed to describe the tenor (1936, 100). In the metaphor "Juliet is the <u>Sun</u>," "Juliet" is the tenor and the "<u>Sun</u>" is the vehicle. Richards felt that Aristotle's model was far too simple and thought that the tenor and vehicle combined in the mind of the reader to create a third, emergent meaning that combined the meanings of both domains. For Richards, the tenor not only takes on some of the meaning of the vehicle, but the reverse is also true. "In many of the most important uses of metaphor," writes Richards, "the co-presence of the vehicle and tenor results in a meaning (to be clearly distinguished from the tenor) which is not attainable without their interaction" (1936, 100).

Max Black uses a different set of terms to explain metaphor, but his focus on the interaction between the two domains is similar. Instead of *tenor* and *vehicle*, Black introduced the terms *focus* (Juliet) and *frame* (the Sun). ("Are *we* now using metaphor? And mixed ones as that?" Black asks as an aside (1962, 28)). He also used the terms *principle* and *subsidiary* subject to suggest the relative importance of the two terms to the speaker's point (or, in a later paper, *primary* and *secondary* domains (Black, 1979, 28)). The meaning of a metaphor is determined not by syntax, which might point to a precedingly existing similarity, but through semantics and pragmatics and the context in which the metaphor is used. "It would be more illuminating in some of these cases to say that the metaphor creates the similarity than to say that it formulates some similarity antecedently existing," writes Black in what has now become known as the *interaction view* of metaphor (1962, 37). "The principle and subsidiary subject take on new meanings by virtue of the metaphorical phrase," he writes (1962, 39). In the example "Man is a wolf," both concepts "Man" and “wolf” are linked to a “system of commonplaces” associated with them, which “organizes” our understanding (1962, 41). Modern scholars often use the terms *target domain* (Man) and *source domain* (wolf) to remind readers that there is an entire system of concepts related to the words used in the metaphor. These are the terms I will be using for the rest of this paper.

The View from Anthropology

For the anthropologist, it is obvious that the “commonplaces” in Black’s definition of metaphor are largely culturally determined. For an anthropologist such as Keith Basso, working amongst the Western Apache in Cibecue, Arizona, the use of metaphor is not merely the purview of grammarians. Instead, metaphors can only be properly decoded by understanding the network of symbols that are meaningful to a particular culture (1976, 95). Even within a community, metaphors can be interpreted in very different ways within a particular speech community (1976, 97). As an example, Basso describes the “wise words” spoken by Apache elders to impart advice such as “butterflies are girls” or “ravens are widows” (1976, 98). When Basso tries to determine the meaning of the metaphors in conversations with interlocuters, his attempts are met with laughter. Basso, perhaps unduly burdened with baggage form Aristotle’s substitution model, draws a parallel between a raven’s black plumage and the widows’ black clothing. His connection is dismissed by the elders. “It doesn’t mean anything because it doesn’t tell us what they *do*,” says one. “You have to think about how they are the same in what they *do*—not what they look like,” he says (1976, 105), for example, by focusing on the tendency of ravens to loiter on the periphery of human activity much like how widows remain detached from community life during their period of mourning (1976, 101).

Amongst the Coeur d’Alene Apache in Idaho, Basso describes another technique of using metaphor to personify inanimate objects that are a product of local culture, a trope common in the field of artificial intelligence, as we will soon see. In the local language, Basso notes that the word for tire-tracks in the dust translates as “wrinkled feet” (Palmer, 1996, 224). The rest of the parts of a truck correspond to the parts of a living being with the following mapping:

Root metaphor: A Truck is a Human	
Target domain: parts of a truck	**Source domain:** parts of an animal
hood	nose
headlights	eyes
windshield	forehead
front wheel	hands/arms
back wheels	feet
- tire tracks	- wrinkled feet
under the hood	innards
- battery	- liver
- electrical wiring	- veins
- gas tank	- stomach
- distributor	- heart
- radiator	- lung
- hoses	- intestines

Table 1: Example of structural metaphors from the Coeur d’Alene Apache describing the parts of a truck in terms of human body parts (Basso, 1967).

Here, the underlying metaphor of “a truck is a human” is made useful by its “entailments,” that is, by the consequences of the original comparison. The use of a metaphor like this is productive in the sense that as new parts are added to the truck, the metaphor can be extended while keeping the relationship between target and source domains coherent. One could imagine members of the community speaking of gasoline as “food” or exhaust as “excrement,” or speaking of a mechanic as a “doctor”. To use Lévis-Strauss’ famous term, there is a *homology* in the structure of the metaphor. “Cognitive topology is preserved” between the relationships of the component parts in both domains (Palmer, 1996, 234).

Metaphors We Live By

In linguistic anthropology, metaphor turned into an object of study in its own right with the publication of George Lakoff and Mark Johnson's *Metaphors We Live By* (1980). Their core argument, echoing Richards, is that "metaphor is pervasive in everyday life, not just in language but in thought and action" (1980, 3). Their argument is one borne of empirical observation of speech and text analysis. "Our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature," they claim (op cit.). The language they use in this book form the basis of Conceptual Metaphor Theory (CMT), the dominant theory used to explain the ubiquity of metaphor in human speech. Any metaphors that are identified in speech or in written text, so-called *linguistic metaphors*, can be said to represent a limited number of core cognitive concepts called *conceptual metaphors*. "Linguistic metaphors *realize* conceptual metaphors," writes Alice Deignan (2008, 14).

Lakoff and Johnson describe three types of conceptual metaphors from which linguistic metaphors derive, all three of which we will see in the artificial intelligence corpus. *Orientational metaphors* give abstract concepts a spatial orientation, a universal human tendency based on the fact that our experiences of the world are embodied. A common conceptual example is HAPPY IS UP[2] which gives rise to linguistic metaphors such as "that <u>boosted</u> my spirits," or "thinking about her gives me a <u>lift</u>," (or conversely, "stop bringing me <u>down</u>") (1980, 14). The second type, the *ontological metaphor*, "allows us to pick out parts of our experience and treat them as discrete entities or substances… [which means] we can refer to them, categorize them, group them, and quantify them—and, by this means, reason about them." This allows us to talk about "<u>grasping</u> an

[2] Conceptual metaphors will be written in ALL CAPS in accordance with the norms of the literature in Conceptual Metaphor Theory, as distinct from linguistic metaphors that will appear in quotes with references to the text in which they were originally found (with words used metaphorically within the phrase <u>underlined</u>).

idea," or "replacing a concept," in our minds. In science, (and as we will see, in artificial intelligence), a common conceptual metaphor is IDEAS ARE BUILDINGS seen in expressions like "constructing a theory" or "foundations of a model" (1980, 46). Lastly, *structural metaphors* are cases where one concept is metaphorically structured in terms of another, seen in the above example A TRUCK IS A HUMAN. Structural metaphors have depth and the comparison goes beyond a singular, original insight (1980, 14). These are also very common in scientific literature, and, as we will see, allow scientists to think through the theoretical implications of new theories by relating them to structures they already understand. Structural metaphors can also be orientationally and/or ontological in nature.

For Lakoff and Johnson, part of what creates and constitutes conceptual metaphors at a cognitive level, is culture. "Cultural assumptions, values, and attitudes are not a conceptual overlay which we may or may not place upon experiences as we choose," they write, "It would be more correct to say that all experience is cultural through and through, that we experience our 'world' in such a way that our culture is already present in the very experience itself… *every* experience takes place within a vast background of cultural presuppositions" (1980, 57). Furthermore, as we saw with Locke and Hobbes, metaphors are used by members of a particular culture unconsciously, leading to a reinforcement of cultural tropes. According to anthropologist Naomi Quinn, the creation of new metaphors is fairly rare in everyday speech because people who share cultural values already have a fairly comprehensive set of metaphors with which to describe their experiences (1991, 57). Quinn points to the data she collected by analyzing hundreds of hours of transcripts of Western married couples reflecting on their understanding of marriage. For Quinn, culture consists of a set of "shared understandings that people hold and that are sometimes, but not always, realized, stored, and transmitted in their language" (op cit.). She found a huge variety of metaphors which could be reduced to several underlying conceptual metaphors, but these metaphors were often mixed, and during speech, speakers would flip between metaphors sometimes within the same sentence (Quinn, 1991, 66). “Metaphors are reintroduced over and over again because they are satisfying instantiations of a ‘conventional’ or culturally shared model,” writes Quinn, “capturing multiple elements of that model” (1991, 79). Good structural metaphors, writes Quinn, are culturally bound and “do not merely map two or more elements of the source domain onto the domain of marriage,” she writes, echoing Basso, “in doing so they map the relationship between or among elements as well” (1991, 80).

Since *Metaphors We Live By* (1980) there has been a tension between scholars looking for metaphors that are culturally specific and those that might be universal. Zoltan Kovecses claims HAPPY IS UP as an example of a universal orientational metaphor because when humans are happy and healthy they are upright with their heads above their bodies. Similarly, most cultures speak of the future as “<u>in</u> <u>front</u>” of a person (i.e. “we’re <u>moving</u> <u>forward</u>” or “<u>looking</u> <u>ahead</u> to the future”) and the past as “<u>behind</u>” (i.e. “don’t <u>look</u> <u>back</u>” or “that event is <u>behind</u> us”) (Kovecses, 2005, 47).

Entailments

Conceptual Metaphor Theory (CMT) has come under some criticism since its inception due to the lack of empirical evidence that linguistic metaphors necessarily represent metaphorical *thinking*. Evidence from discourse analysis shows that, although people use metaphors as a way of organizing their communication, and perhaps their thoughts, they have a fair amount of choice when it comes to using metaphors to express themselves (Kovecses, 2005). Psycholinguistics professor Raymond Gibbs Jr. claims there is now lots of evidence to show that people can ignore conceptual metaphors in discourse and can actively chose alternatives that fit the pragmatic purpose of their communication (Gibbs Jr., 2017, 148). He also points out some circular reasoning in CMT: first, by analyzing language, conceptual metaphors are inferred, and then are reconfirmed by referring to examples in language. It is claimed that speakers are unaware of their adherence to conceptual metaphors yet are cognitively bound to them (2017, 114). Instead of making 'strong' cognitive claims of metaphor in the mind, then, Earl Mac Cormac suggests a compromise, that "the metaphorical process not only involves the mind and the brain but it also presumes the external world with its wealth of symbols and culture," he writes, framing "metaphor as a mediating device among the mind, brain, and the external world" (Mac Cormac, 1985, 21).

As such, although I will continue to use ALL CAPS to denote underlying metaphors that unite a group of linguistic metaphors that pull imagery from the same source domain, without any evidence confirming cognitive processes, I will use the term *root metaphor* instead of conceptual metaphor. They are similar in that they both refer to broad metaphorical statements that underlie different linguistic metaphors, yet root metaphors make no cognitive claims. The term was coined by philosopher Stephen Pepper in *World Hypotheses* (1972) to refer to a basic set of *heuristic aids* that allow people to think about complex ideas. Although they are often applied unconsciously through the social conventions of language use, they can also be applied deliberately and creatively, as we will see in the case of artificial intelligence research. Explaining how root metaphors might come about, Pepper suggests that a person "pitches upon some area of common-sense facts and tries if he cannot understand other areas in terms of this one. This original area becomes then his basic analogy or root metaphor" (1972, 91). The use of this term is a way of avoiding any claims about cognitive determinism that I am not prepared to argue from the perspective of anthropology.

In *Fields, Dramas and Metaphors* (1974), Victor Turner uses Pepper's concept of root metaphor to explore how metaphors are used to structure ritual and social interaction. In speaking of the analysis of metaphor, though, to elucidate cultural norms, Turner warns us that "…they may be misleading; even though they draw our attention to some important properties of social existence, they may and do block our perception of others" (1974, 25). If metaphor can be understood as a filter on the world, it is just as useful to ask what the metaphors are leaving out as well as what they are including in their comparison. "One must pick one's root metaphors carefully," writes Turner, "for appropriateness and potential fruitfulness" (op cit.). Taking this warning seriously, we can be wary of the process by which metaphors become commonplace and lose their metaphorical roots. "The danger is, of course, that the more persuasive the root metaphor or archetype, the more chance it has of becoming a self-certifying myth, sealed off from empirical disproof," writes Turner (1974, 29). Either way, good metaphors have consequences; they suggest derivative metaphors that can be extended into new situations and suggest commonalities, known as *entailments*.

Over the years, different taxonomies have developed to describe the level of *metaphoricity* in a given metaphor and the process of linguistic evolution they go through to become "sealed off" from conscious access. New metaphors are often called *novel metaphors* or *innovative metaphors* and rely on context to make them intelligible (i.e. "his car was a <u>gorilla</u>" or "her love was <u>napalm</u>"). These are metaphors that make a listener pause, or make a reader reflect on the entailments of a certain metaphorical comparison. As we will see, novel metaphors are quite common in scientific literature as scientists strive to make sense of their models and communicate them to other people, but they are fairly rare in everyday speech (Gibbs Jr., 2017). *Conventional metaphors* are those that pepper our speech and, although the words still retain their literal meanings, the metaphors have become 'stock' and are not noticed as novel expressions (i.e. "America is a <u>melting pot</u>" or "the wind <u>whispered</u> through the trees" (Deignan, 2008, 39)). *Dead metaphors* are metaphors that are considered by standard speakers to be synonyms of the literal meaning, although with some prompting, they can see the metaphor underlying the term (i.e. "it was a very <u>deep</u> blue" or "they used a <u>crane</u> to move the bricks" (op cit.)). Sometimes dead metaphors are only traceable through etymology and, as such, can be called *historical metaphors* (i.e. the word *coincidence* comes from the roots *con* (together) and *incidere* (falling), an ontological metaphor of two things physically falling together (Hyde, 1998, 97)). In general, Conceptual Metaphor Theory is concerned with conventional and dead metaphors, and their relative frequency in everyday speech, as a way of inferring cognitive patterns that lie below the threshold of conscious perception of speakers. Moving forward, this study will be more interested in tracing the *trajectory* of novel metaphors as they become conventional.

The so-called Career Model for metaphor gives us a hypothetical structure with which to work. "The career of metaphor hypothesis suggests that a computational distinction can be drawn between novel and conventional metaphors," write Bowdle and Gentner (2005, 199). Conventional metaphors have been used so much in everyday discourse that a listener does not process them as metaphors; the literal and metaphoric meanings are on their way to becoming synonyms. The strength of this theory is that it predicts "that as metaphors become increasingly conventional, there will be a shift in mode of alignment from comparison to categorization" (Bowdle and Gentner, 2005, 208). In the last twenty years, studies using fMRI analysis have shown that "novel metaphors invite sense creation but conventional metaphors invite sense retrieval" (op cit.). Conventional metaphors are understood by the human brain faster than are novel metaphors and even more so than are similes. In this framework, similes include the word 'like' or 'as' to signal to a reader or listener that the speaker/writer is about to introduce a novel comparison into discourse. Experiments show that readers find conventional metaphors written as similes jarring and unnatural when reading (i.e. "America is like a melting pot" or "I understood the problem as if I was grasping it") (Bowdle and Gentner, 2005, 202). But this model of metaphor relies on the reader to understand that a metaphor is being used, which is perhaps why metaphors are difficult for non-native speakers to fully comprehend (Mac Cormac, 2005, 28).

Metaphor in Science

Novel metaphors are of most interest to scholars of scientific language because they are so often used to explore new discoveries and to explain concepts to a general audience. One of Aristotle's first examples, in fact, to illustrate his formula of analogy found in *Rhetoric*, was to draw a parallel between the propagation of light and the propagation of sounds through different densities of media (Leatherdale, 1974, 31). When structural metaphors like this are used, the metaphor can be investigated and the logical consequences of the metaphor can be analyzed for their coherence with observed phenomenon. In short, the analogy itself suggests experiments that can be conducted to test the strength of the model. "New metaphors, by virtue of their entailments, pick out a range of experiences by highlighting, downplaying, and hiding," write Lakoff and Johnson (1980, 152). They conclude that, instead of being opposed to objective and rational descriptions of reality, "metaphor is thus *imaginative rationality*" (1980, 193).

Experimental evidence suggests that people often use metaphors as a way of thinking through complex situations. But not just any metaphor will do. "Even the subtlest instantiation of a metaphor (via a single word) can have a powerful influence over how people attempt to solve social problems like crime and how they gather information to make 'well-informed' decisions," writes psychologist Paul Thibodeau (2011, 1). "Interestingly, we find that the influence of the metaphorical framing effect is covert: people do not recognize metaphors as influential in their decisions; instead they point to more 'substantive' (often numerical) information as the motivation for their problem-solving decision," he writes, suggesting the importance of making visible the use of metaphor in scientific reasoning if objectivity and accuracy are the goals (op cit.).

One of the dangers of relying on metaphor for reasoning in science is that, through their use, metaphors often become more conventional over time and are eventually assumed to be literal. "Forgetting that the foundational theory is a metaphor can allow us to accept it through familiarity rather than evidence," writes Mac Cormac (1985, 34). One of the strategies Mac Cormac suggests is to keep metaphors top of mind is to expose the metaphors in our work by periodically trying to replace key words in the structural metaphor with those of another domain. The exercise might suggest new insights or directions for investigation but will also serve to remind scientists that the model is a metaphorical, and not a literal account of the structure of nature (1985).

Because metaphor can blind us to phenomenon that has been filtered out, scientific models built upon analogy can hinder progress. In the nineteenth century, one of the great puzzles of physics was in determining the medium through which light travelled. The structural metaphor that guided the thinking at the time was in seeing light waves as analogous to sound waves and, the literal originator, to water waves. The initial metaphor was remarkably productive and gave rise to similar vocabulary for both phenomena: propagation, frequency, wavelength, speed, period etc. To maintain metaphorical congruence, scientists assumed that light needed a medium through which to travel, as did sound and water waves. They dubbed this material "ether" even though it had not yet been detected, such was their confidence that it would soon be discovered. 19th century scientist Lord Kelvin talked about ether as a real substance, as having density, elasticity and other properties of physical objects (Black, 1962, 228). Now we know, thanks to Einstein, that ether does not exist and that light travels through a vacuum unlike sound waves, a discovery that was hidden in plain sight behind the trappings of the wave metaphor. Black summarizes: "we reap the advantages of an explanation but are exposed to the dangers of self-deception by myths" (op cit.).

The parts of an analogy that don't fit the model are often called *disanalogies*. In a speech to the American Psychological Association in 1955, physicist Robert Oppenheimer said, "at each point the first scientists have tried to make a theory like the earlier theories: Light, like sound, as a material wave; matter waves like light waves, like a real, physical wave; and in each case it has been found one had to widen the framework a little, and find the disanalogy which enabled one to preserve what was right about the analogy" (1955, 131). Philosopher and historian of science Mary Hesse structures it this way: water waves/water particles :: sound waves/air particles :: light waves/ether particles (Hesse, 1966, 26). The term "ether" must be recognized, in Hesse's words, as a *negative analogy*; an entailment that needs to be expunged, as opposed to a *positive analogy*, an entailment from a metaphor that has productive value in expanding the theory (1966, 8).

Metaphorical models of phenomena are "meant to be exploited energetically and often in extreme quantitative detail in quite novel observational domains; they are meant to be internally tightly knit by logical and causal interrelations" (Hesse, 1980, 119). Any hypothesis in a "theory in the process of growth" (Hesse, 1966, 10) uses metaphor to state a *neutral analogy*, a hypothesis by a scientist who is not sure how deep the homologies (1966, 81). The hypothesis, in short, is open to testing. Hesse calls directly for a modification of the inductive model of science to include an explicit step devoted to reasoning through metaphor. Over time, the metaphorical description will be refined so as to replace any original, literal attempts at explanation in the same way that is described by the 'career model' of metaphor by Bowdle and Gerntner. One way of thinking about this process is that as "metaphoricity" dimishes, "facticity" is increased, to use Latour's term from *Laboratory Life* (1986). Latour explains that the process by which "the statement achieves any degree of facticity" (1986, 80) slowly removes any "trace of authorship" or of contingency on lab equipment (82).

The importance of using metaphor to explain scientific phenomenon is especially important when scientists cannot see or experience first-hand what they are studying. There is a need, essentially, for an ontological metaphor to 'make sense' of what is insensible. Hesse calls this the need for models to be *picturable* (1966, 19). "Scientists understand nature largely in terms of metaphorical concepts, based on embodied understandings of how nature works," argues Brown (2003, 11). Einstein's gedankenexperiments (thought experiments), perhaps the quintessential example of a ‘purely cognitive’ model, were all “deeply embodied,” he argues. Even though they lead to some deeply unintuitive and non-classical conclusions, were based on material world experience such as “riding a beam of light” (op cit.).

Methodology

> *"I believe it would be an interesting exercise to study the key words and expressions of major conceptual archetypes or foundational metaphors, both in the periods during which they first appeared in their full social and cultural setting and in their subsequent expansion and modification in changing fields of social relations"*
>
> - Victor Turner (1974, 28)

Now that we are more "at home in metaphor" (Cox and Lathem 1966, 39), we can turn our attention towards the corpus in question. Written texts are cultural artifacts and as such, should be considered a robust source of anthropological data. Basso considers written texts a part of what he calls "the ethnography of communication," and frames writing as a communicative *activity*, and not as a static set of sentences waiting to be parsed by grammarians (Basso 1973, 426). Writing is one of many channels of communication open to members of a (literate) community, so what factors determine their choice to write their thoughts down? In science, the tradition of summarizing one's findings in a scientific journal and then letting the community critique and, if they can, reproduce experiments, is well established. Scientists write to gain credibility and to establish their findings as nodes in the growing network of knowledge producers. Written texts, in science, gain their power by reference to other written texts (Latour 1987).

In examining scientific research papers from the historical record, we can trace this network of influence back to its roots and interrogate the language used for clues as to the cultural context surrounding the authors. "The anthropologist respects history, but he does not accord it a special value," writes Lévis-Strauss. "He conceives it as a study complementary to his own: one of them unfurls the range of human societies in time, the other in space," he continues (1966, 256). In combination with ethnographic studies of scientists working in their laboratories, anthropologists can put together a full picture of science as a cultural activity.

The problem with hunting metaphors in text is that they are often hard to find. "The first thing to be noted about metaphors is disagreeably paradoxically," writes W. H. Leatherdale. "There is no explicit form which metaphors can be said to take. … Nor is it always clear where a metaphor begins or ends or what words in a sentence constitute the metaphor it contains," he writes (1974, 95). As such, the methodology used here is a unique blend of two methodologies that are widely used to analyze metaphor: the Metaphor Identification Procedure (MIP) described by the Pragglejaz group (2007); and the 'key-word-in-context' method of corpus analysis described by linguist Alice Deignan (2008). Both methods have shortcomings but work well together. Corpus analysis identifies the frequency and form of known metaphors but cannot identify new ones, and as such, is insufficient to both identify metaphors and examine their frequency. As such, before corpus analysis was undertaken, an adapted form of MIP was employed in order to identify metaphors manually in the corpus of scientific papers.

Adapted Metaphor Identification Procedure

In 2007 ten leading researchers in metaphor studies collaborated to standardize their research methodology under the name Metaphor Identification Procedure (MIP). "Variability in intuitions, and lack of precision about what counts as a metaphor, makes it quite difficult to compare different empirical analyses," they write. "More important, the lack of agreed criteria for metaphor identification complicates any evaluation of theoretical claims about the frequency of metaphor, its organization in discourse, and possible relations be-tween metaphoric language and metaphoric thought," they continue[3] (Pragglejaz 2007, 2).

[3] The bizarre name of the author of this paper comes from the first names of the ten collaborators, many of whom are referenced in this paper for their individual contributions: **P**eter Crisp; **R**aymond W. Gibbs Jr.; **A**lice Deignan; **G**raham Low; **G**erard Steen; **L**ynne Cameron; **E**lena Semino; **J**oe Grady; **A**lan Cienki; **Z**oltan Kovecses

The original procedure as described in the 2007 paper demands a close reading of every lexical unit in a particular text to determine its level of metaphoricity. The thirty-one papers in the present corpus are comprised of almost 386, 000 words, so applying the MIP to each word would be impractical. Instead, I applied the MIP in its original form to only one document, the ‘foundational text’ written in 1956 that first used the term “artificial intelligence” (McCarthy, 1956), and instead scanned the rest of the papers manually for metaphors first used in 1956 or for other ‘novel metaphors’ that went beyond the level of a single lexical unit. Analysis was thus conducted at several levels, as summarized below in two sections: identification and analysis. References are included for methodologies I used that are not found in the original Pragglejaz paper (2007).

1. **Metaphor Identification**
   a. **Meta**: Analyze the structure of the text, determine who the author is, who the intended readership is, the genre of the text and the medium of transmission (Basso, 1973).
   b. **General**: Read the entire text to establish a general understanding of the text’s content and its purpose.
   c. **Lexical**: Determine the *lexical units* in the text (if necessary, divide by forward slashes to avoid ambiguity on what constitutes a lexical unit). This level of detail was only conducted on the “foundational text.”
   d. **Metaphoricity**: Establish the unit’s meaning in context and compare it to the standard meaning of the unit.[4] Determine the “level of metaphoricity” of the unit. If the unit is categorized at any of the three non-historical levels, <u>underline</u> the unit.
   e. **Text**: Widen the units of analysis beyond the lexical level to include phrases and recurring images of a source domain spread throughout a text (Steen, 2007, 84). Sustained or *systematic metaphors* are “a particular set of linguistic metaphor vehicles

[4] Standard meaning is defined as the most commonly used meaning of the word. To verify my intuition on this, several sources were used as corroboration: Online Etymology Dictionary; Dictionary.com; Google Ngram.

in talking about a particular topic, or closely connected topics… an emergent grouping of closely connected metaphors," (Steen, 2007, 91) that are used as a particular text unfolds.

2. **Metaphor Analysis**

   a. **Metaphor Markers**: Look for linguistic markers that deliberate metaphors are being used including similes ("like", "as"), hedging ("imagine…", "a good comparison might be…"), quotation marks, modifiers, or "alternate representations" of the same term (Basso 1973). In contrast to the conceptual metaphors described by Lakoff and Johnson which lie below the level of our conscious perception, markers like this denote the use of a "deliberate metaphor" (Steen, 2015).

   b. **Root metaphors**: Determine whether the individual linguistic metaphors suggest one or more root metaphors being used as a heuristic tool. Root metaphors can be categorized in the same way Lakoff and Johnson (1980) categorize their conceptual metaphors: ontological, orientational, or structural (often overlapping) but without the same commitment to cognitive determinism.

   c. **Heuristics**: Look for evidence that the root metaphors are being used to think through the consequences (or entailments) of a particular scientific model (Hesse, 1966).

   d. **Collocation:** Pay particular attention to words that always appear together in text (Deignan, 2008).

Once metaphors were identified and analyzed, key words representing those metaphors were fed into Sketch Engine to determine their frequency and context. Historical metaphors were not included in the corpus analysis because of their ubiquity and their invisibility to the authors. Metaphors that were used to explain a structural feature of artificial intelligence in terms of something else were the main focus of analysis. If several structural metaphors were found to use the same source domain to explain something in the target domain, a root metaphor was hypothesized that might explain them all. For example, different structural metaphors such as the words "<u>neuron</u>" or "<u>memory</u>" can be said to represent the root metaphor A MACHINE IS A BRAIN even though the phrase "a machine is a brain" might never come up in the corpus.

Corpus Analysis

Corpus linguistics goes a long way in providing data to buttress theories in linguistic anthropology because the corpora are large enough to be representative of certain domains of discourse (Deignan, 2008). Corpora can be general (composed of texts assumed to represent typical language use) or specialized (composed of texts written in a particular domain or medium). Corpora can also differ by being closed (where nothing new is added) or open (texts are constantly being added). This difference becomes relevant when considering the corpora that comprise scientific journals that are defunct, or those that are still productive. Corpora can also be complete (i.e. every issue of the *New York Times*) or a sample (either a random sampling or a specific cross-section like the years 1956-1976) (Deignan, 2008, 77).

Key words[5] can be searched for and are presented as a *concordance*, a list of places where the key-word has appeared with the key word lined up in the middle of the screen as a *node*. This method is known as "key-word-in-context". Researchers can examine the semantic or pragmatic context of a phrase, as well as grammatical variations. Instances of *collocation* can also be revealed, showing occurrences of words that typically (sometimes exclusively) occur together. In short, according to Alice Deignan, "concordance data … reveal linguistic patterns that demand

[5] More specifically, the *lemma* is the search term, sometimes called the *head word* in a dictionary, and search results include all the inflections of that word that represent the same *lexeme*.

explanation” (2008, 9). Corpus analysis is ideal for searching for systematic metaphors that appear frequently in speech, possible evidence for the presence of a “theory-constitutive” root metaphor.

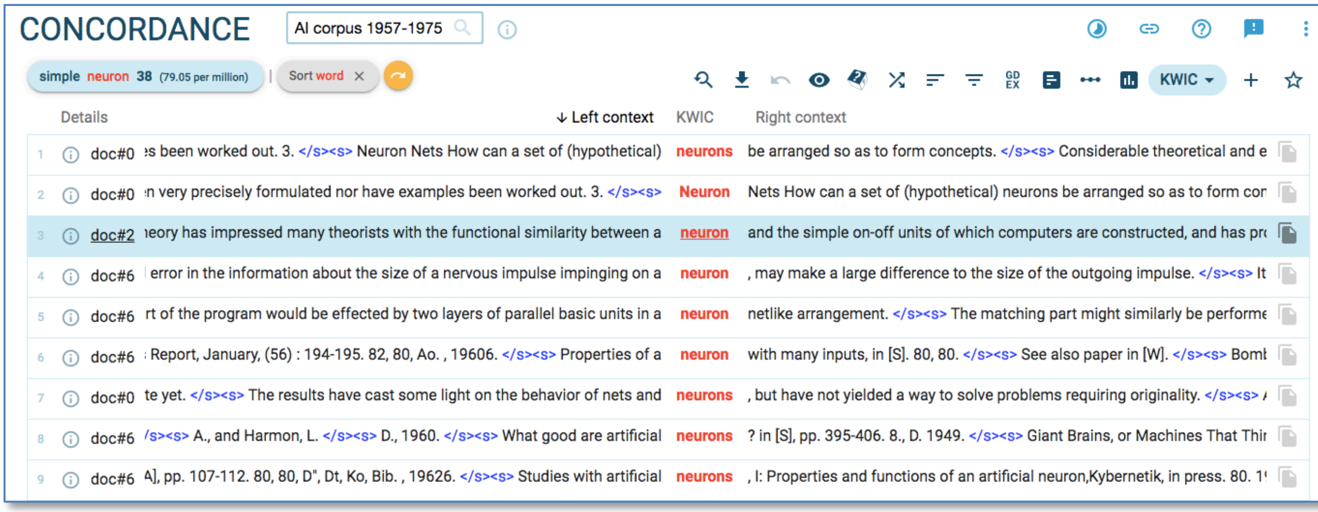

Figure 1: An example of concordance data as presented by Sketch Engine. The word “neuron” is the key word, placed in a ‘node’ in the middle of the screen, showing the context both to the left and right of the node.

For this study, the thirty-one papers[6] assembled into a specialized corpus are widely considered to be the most important documents of the first wave of artificial intelligence research, papers that set the direction of future research. We will refer to it as the “Early AI Corpus” for the remainder of the text. The papers were uploaded to Sketch Engine, an online corpus analysis tool that has been used widely in linguistics and related fields (Kilgarriff, 2014). Before they were uploaded, it was determined that the papers were a) machine-readable; and b) accurate copies of the

[6] A full list of the papers in the Early AI Corpus is included in Appendix A. If any of those papers are referenced directly with in-text citations they are also included in the References list.

original papers.

Tables are included below that summarize the concordance data. The left-hand column includes keywords (in **bold**) that represent structural metaphors found by using the MIP. Keywords in regular font below the bold words show collocation frequency. The "+" sign indicates where the bold term above it was found. For example, the word "intelligence" appeared 272 times in the entire corpus.

Root Metaphor:			A MACHINE IS A PERSON	
	count	literal	metaphorical	marked
intelligence	272	66	206	quotes (x7)
artificial +	96	0	96	modifier quotes (x5)
human +	13	13	0	
Level:			Novel	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

Table 2: An example of a table summarizing concordance data from Sketch Engine.

Concordance data showed that 66 of the appearances were literal (referring to human intelligence) and 206 of them were metaphorical (referring to a machine). The collocation "artificial intelligence" (written here as "artificial +") appeared 96 times and all of them were metaphorical. The term "human intelligence" ("human +") appeared 13 times and was, obviously, not metaphorical. The last column shows if any of the instances of the metaphorical term were "marked" as being deliberately metaphorical, such as through the use of single quotes, modifiers or hedges (see point 2. a. in the previous section). For example, the word "intelligence" was marked as being metaphorical with the use of single quotes only seven times (out of 206 possibilities), and the lexical unit "artificial intelligence" was enclosed in quotes five times. The term "artificial" is also a modifier when collocated with "intelligence". Below the quantitative data the level of the metaphors are classified as being either novel, conventional or dead. The metaphors are also classified as either bring ontological or orientational (see section “Metaphors We Live By” above) or as systematic (see point 1. e. above), appearing many times throughout the corpus.

Scope

Before we dive into the corpus itself, a note on terminology is necessary. Although the two types of literary tropes *metaphor* and *metonym* are different kinds of comparisons, they have a similar form and serve a similar purpose (having one thing stand in for another). Metonym has a “referential function” between objects that are in the same domain (i.e. the ‘White House started a trade war with China’ with the ‘White House’ standing in for the government officials who made the decision). Metaphor, as discussed above, compares two things from different domains. A special kind of metonym is *synecdoche*, where a part of an object stands in for the whole (i.e. a face painted in a portrait stands in for the entire person (Lakoff and Johnson 1980, 37). We will include metonym, synecdoche, along with simile and metaphorical idioms in our analysis in order to gain a holistic understanding of the literary tropes used to convey meaning in the corpus. “Metaphor is the genus of which all the rest are species,” says Colin Thurbayne (1970, 19). Analogies will be treated as *extended metaphors*. The diagram included above, provides a good summary of the relationship between the three core tropes (Maranda, 1971). We’re interested in the entire relationship complex, not just in individual linguistic terms.

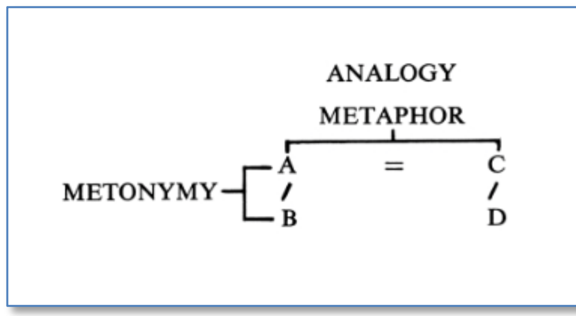

Figure 2: The relationship between metonym, metaphor and analogy (Maranda, 1971)

Lastly, there is a process whereby new terms are introduced into a lexicon through *catechresis,* a situation where a new word emerges from a metaphorical comparison. “The theoretical sciences experience crises of vocabulary,” writes David Miall (1982, 96). Making up new words arbitrarily rarely works because for people to find terms useful they need to be embedded in a network of meaning, with metaphorical relationships to a model or root metaphor suggested by the name (Miall, 1982, 97). Kathryn English gives an example from the 19th century when experiments first revealed evidence of electricity. Terms from fluid dynamics were used to describe the structure of this new phenomenon such as "current" "flowing" and "resistance". These catechreses have survived and these terms are now synonyms under a head word in the dictionary.

Structural Metaphors

"Once the choice of metaphor was made, its use began to create similarity. Those hearing the metaphor were led to think along certain lines as they conceptualized the observations"

- Theodore Brown (2003, 19).

The Dartmouth Conference

In 1955, computer scientist John McCarthy wrote "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," (1956) a bid to the Rockefeller Foundation to financially support a summer-long meeting at Dartmouth College for scientists and engineers to collaborate. This is the "foundational text" of the field of AI and is the first time the term "artificial intelligence" appears to describe the burgeoning scientific field. Chronologically, it is the first in the Early AI Corpus (see Appendix A for full list of titles). It can be analyzed as a rhetorical text, one designed to convince potential funders of the value of such a research project. The readers of the text can be assumed to be members of a review committee at the Rockefeller Foundation and, as such, the proposal is not overtly technical in an effort to connect with readers. The document is structured more as a memo than as a formal scientific paper, including seven short, digestible summaries of different "aspect of the artificial intelligence problem" for the lay readers (McCarthy, 1956, 1). There are short biographies of the four collaborators on the proposal and summaries of what expertise they bring to bear on the seven aspects identified (although it is widely acknowledged that McCarthy himself wrote the memo and he uses "I" liberally throughout). McCarthy spends the rest of the proposal explaining in more detail how his own work would benefit from collaboration with the other scientists and engineers. This section contains footnotes and in-text citations to published scientific papers, situating the proposal in a network of associations (Latour 1987). References like "suggested by Craik" (McCarthy, 1956, 6) are metonyms that rely on a sort of "root metonym" substitution under the category of PRODUCER FOR PRODUCT (Lakoff and Johnson, 1980, 38). These references are made "to mark affiliation and show with which group of scientists he identifies" (Latour 1987, 34).

"Dartmouth," as it is metonymically referred to under the category PLACE FOR EVENT (Deignan, 2008, 56) is "generally recognized as the official birth date of the new science of artificial intelligence" (Crevier, 1993, 49). The conflation of the use of the term "artificial intelligence" with the "birth" of the field of study suggests that the event can be understood as a type of *naming ceremony*. Dartmouth is described as the "christening of the new discipline" (Crevier, 1993, 50). Such a ceremony 'calls into being' the entity being named and is often cast as a birth or re-birth. A naming ceremony is necessary "to label a discipline is to define its boundaries and identity" (op cit.). When looking back on the conference, the standard narrative is that the rapidly advancing fields of computer science, cybernetics, game theory and psychology had started to converge, and a new name was needed to cement the identity of this new field. "The term was chosen to nail the flag to the mast, because I (at least) was disappointed at how few of the papers in *Automata Studies* dealt with making machines behave intelligently," says John McCarthy (Moor 2006, 87).

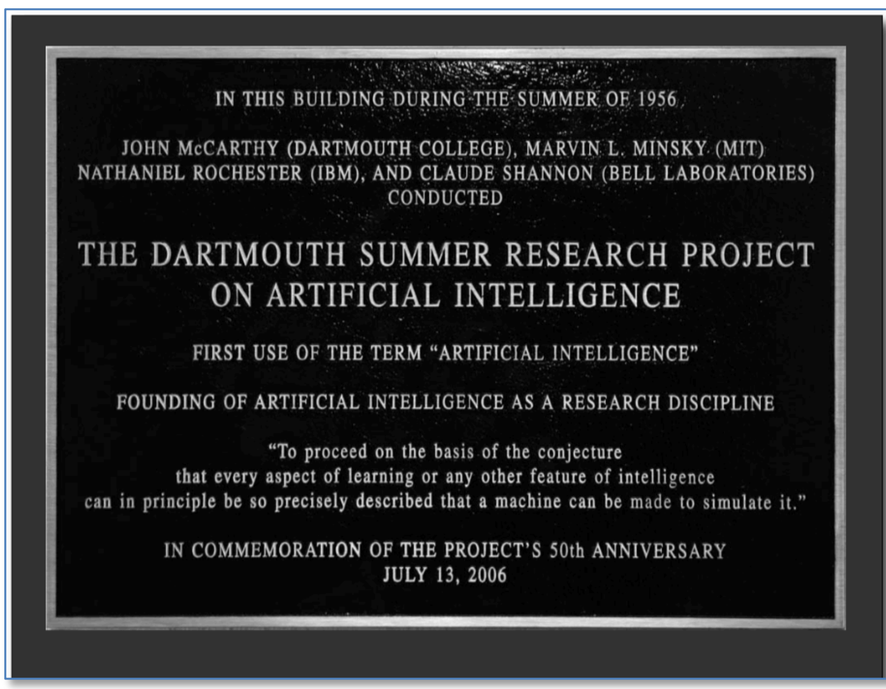

Figure 3: Commemorative plaque unveiled at the 50th Anniversary of the Dartmouth Conference (Moor, 2016, 90)

There is a sense of striving in a good metaphor, a momentum that suggests future actions and suggests ideas for expansion. Once the metaphor of "artificial intelligence" was 'activated' in 1956, it defined the field.

One of the key differences between the Dartmouth Conference and a more traditional naming ceremony is the fact that its importance was only evident to researchers after the fact. "Looking back, that was the start of the community," said AI pioneer Marv Minsky 40 years later (Crevier, 1993, 49). "Dartmouth indeed defined the establishment," he explains, "for almost two decades afterwards, all significant AI advances were made by the original group members or their students" (op cit.). There is a ritualistic structure to academic conferences, and particular meetings grow into the pantheon of legend with the benefit of hindsight. But according to reports at the time, the event was underwhelming. Only six people showed up in the summer of 1956 (plus the four organizers), and even then, not all at the same time. Two participants, frequent collaborators Allen Newell and Herbert Simon, reportedly didn't like the name "artificial intelligence" and persisted in using "Information Processing Language" to describe the field (Crevier, 1993, 51). Still, to mark the 50th anniversary, a plaque was unveiled at Dartmouth with Minsky, McCarthy and several other original attendees looking on.

Artificial Intelligence

The word "intelligence" can be classified as having metaphorical meaning as it is applied to a machine, although it is modified by the word "artificial" to make the meaning of the phrase obvious. "Artificial intelligence" was treated in the corpus as a single lexical unit based on the common collocation of the two words. In the Dartmouth proposal (1956), its first appearance, the term is used only four times (excluding the title and any page headers), while the word "intelligence" appears alone twice (once modified with the word "mechanized"). The term certainly qualifies as a novel metaphor, although its rapid adoption and use throughout the rest of the corpus after 1956 rapidly rendered it conventional. The word “intelligence” appears 272 times in the Early AI Corpus and was modified over one-third of the time with the term “artificial” and 13 times with the word “human” (as a marker for the literal usage of the term).

Root Metaphor:			A MACHINE IS A PERSON	
	count	literal	metaphorical	marked
intelligence	272	66	206	quotes (x7)
artificial +	96	0	96	modifier quotes (x5)
human +	13	13	0	
Level:			Novel	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

Table 3: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A PERSON: ***intelligence****.*

The presence of these linguistic metaphors realizes the root metaphor on which the entire field of AI is built: A MACHINE IS A PERSON. Science historian Richard Boyd calls this a *theory constitutive metaphor* (1979, 361), one upon which the entire field relies to make sense of observed phenomena. The structure and the entailments of the metaphor have largely determined the direction of the research. “The hypothetical, or exploratory, role of metaphor is central to theory development and supports the view that it can provide a way to introduce terminology for features of the world whose existence seems probable but many of whose fundamental properties have yet to be discovered,” he writes (1979, 357).

Machines are People

The two tables in this section contains frequency data on linguistic metaphors around other ‘human-like’ characteristics that support the root metaphor A MACHINE IS A PERSON.

The first type of personification is in using the word "computer" itself, which historically referred to a person who performed calculations. The term "computer" began to be used to refer to machines that aided in such calculations at the end of the 19th century (OED, 2019). Frequency of use increased with the "mechanical computers" built to crack code during WW2, and by

McCarthy's time the word "computer" could be considered a conventional metaphor. The modifier "human" was even employed six times implying the metaphorical usage was so common that it was necessary to distinguish the literal meaning. That said, when looking at the data from the corpus, "machine" was still the preferred terminology.[7] Throughout the text the noun phrase "Turing machine" is the most common, a metonym that points the reader to Alan Turing's famous paper from 1950 in which he described a hypothetical machine that might 'pass' as a human when judged by its answers to questions from a panel of inquisitors (Turing, 1950).

Root Metaphor:			A MACHINE IS A PERSON	
	count	literal	metaphorical	marked
computer	802	6	796	
human +	6	6	0	modifier
digital +	112	0	112	modifier
analog(ue) +	3	0	3	modifier
machine	1,568	1,533	35	
Turing +	47	0	0	metonym
intelligent +	35	0	35	
+ code	20	20	0	
human +	6	6	0	modifier
Level:			Conventional	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

*Table 4: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A PERSON: **computer, machine**.*

[7] In Alan Turing's famous pre-war paper "On Computable Numbers" (1937), he is much more comfortable using "computable" (113) or "computing" (12) as an adjective, or "compute" as a verb (31), than as a noun. The word "computer" appears 16 times in that document as compared to "machine", as a noun, 160 times.

If the machines in question are cast as people, what kind of people are they? A closer analysis of McCarthy's text can give us clues. The machines are described as being obedient people, faithfully executing the commands of the scientists who program them. "One expects the machine to follow this set of rules slavishly and to exhibit no originality or common sense," McCarthy writes in his proposal (1956, 6). While the verb make/made is a verb applied to the machine as a subject four times in McCarthy's text (i.e. "the machine… could make reasonable guesses"), the word appeared twenty times in the text as applied to the machine as the object, in the sense of the scientist forcing the computer to do something, (i.e. "how to make machines use language" or "make the machine form and manipulate concepts") further emphasizing the passive and obedient nature of the machines. The scientists themselves are described as agents who control the "behavior" of the machines. A verb like "try" appears seven times in the text, but only one of them applies to the machine. "Trying" implies flexibility of approach, which implies ambiguity, and perhaps failure, whereas it is assumed the machines will execute tasks provided to them accurately and successfully.

When the word "machine" is placed in a sentence as the subject, verbs are often used to give the machine human-like agency, the most common verbs being "think", "do", "learn" and "have." When the word "machine" is the object of the sentence, the most common verbs used are "build", "construct", and "design," where the scientists are the subjects doing the building. The breadth of the agency given to machines is given here with a complete list of the verbs found in McCarthy's proposal applied to "machine" or "computer" as the subject of the sentence:

MACHINES ARE PEOPLE that can:

Copy, execute (x3), form (x8), formulate (x7), make (x4), do (x8), work (x3), try (x2), manipulate, operate (x3), acquire, respond, improve (x4), self-improve (x2), find, solve (x7), guess (x3), simulate (x8), predict (x3), transmit, learn (x4), develop, be imaginative (x3), acquire, exhibit (x4), be trained (x2), abstract (verb), perform, assemble, explore, get confused, behave (x3).

MACHINES ARE PEOPLE that have (or could have):

Behavior (8), memory (2), capacity (3), a laborious manner, internal habits, language, a sophisticated manner, symptoms, character(istically), speed (2), tend(ency), higher function, originality (8), common sense, intuition (2), strategy.

Agency is given to the machines to undertake a wide variety of actions. This is a common rhetorical trope to get readers to identify with the object of study. In corpus analysis of different types of scientific texts, agency is found to have been given to "organs, proteins, bacteria, drugs etc." (Brown, 2003, 49). This allows for simpler sentences in a syntactic sense, but also allows readers to imagine the subject as the protagonist of a narrative. At several points in the Early AI Corpus machines are described as having "internal habits," in contrast to "external experiments," suggesting a core inside/outside dichotomy constructed from the perspective of the machine. Rhetorically, the root metaphor A MACHINE IS

Root Metaphor:			A MACHINE IS A PERSON	
	count	literal	metaphorical	marked
memory	332	80	252	quotes (x5)
human/his +	15	15	0	modifier
in/into/internal +	68	0	68	
behavior(iour)	602	162	440	
human +	45	45	0	modifier
intelligent +	22	0	22	modifier
social +	10	10	0	modifier
machine +	8	0	8	modifier
computer +	0	0	0	modifier
intuition	14	9	5	
common-sense	15	5	10	
Level:			Conventional	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

Table 5: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A PERSON: ***memory, behavior, intuition, common-sense****.*

A PERSON is strengthened when the scientist supresses their own viewpoint (from which *everything* in the experiment occurs outside their own body) in favour of imagining the machine as the ego in the text and describing the experiment from the machine's perspective.

At one point in McCarthy's text, a curious alternate representation appears as McCarthy puts quotation marks around two phrases usually reserved for humans: a machine "may be "trained" by a "trial and error" process". The presence of the quotes here around one verb ("trained") and one idiom ("trail-and-error") suggests that McCarthy is 'marking' his metaphors as novel. It suggests that McCarthy was well aware of the metaphorical innovations of his text and chose, at this particular point, to flag his use of the terms as they are being used in their new target domain.

All of McCarthy's tendencies to personify are present in the full Early AI Corpus, summarized with the data in the table above. "Intuition" and "common sense," two words reserved for describing only human behaviour by McCarthy, are applied to both humans (literal) and machines (metaphorical). Machines are described often as having "behavior" or "behaving" in a certain way. Various modifiers are used to guide the reader to a certain domain such as "intelligent," "human," or "social" for the human domain, and "computer" or "machine" for the machine domain. Yet these modifiers are used sparingly (in only 14% of all cases), so readers need to rely on the context of the sentence to determine the subject: machine or human. Often, nouns like "behavior" are applied to machines and humans within the same paragraph, resulting in a blurring of the lines between the meanings, which is why such metaphors are conventional in nature and are moving quickly in the direction of being "dead". Note that there are no instances of referring to "computer behavior" in the corpus, perhaps through fear that the metaphorical meaning of the term "computer" might be compromised if they started talking about its "behavior," pointing people back to a literal interpretation.

The use of the term "memory" to describe the computer's ability to store information was fairly new in mid-century computer science, but even by McCarthy's day the term was conventional. The term "memory" was first used by American mathematician John Von Neumann in 1945 to describe a component of the first modern computer, the ENIAC, at Los Alamos (Von Neumann, 1945). In the paper, Von Neumann made explicit comparisons between the ENIAC machine and the human brain, one of the earliest attempts to cement the root metaphor in scientific literature. In the Early AI Corpus, instances show up in the corpus of the word enclosed by quotations suggesting the writers knew it was a metaphor taken from the human domain and should not be taken literally. The use of "human" as a modifier occurred only 15 times in the same way that it is used with "behaviour," relying on context to determine the meaning and to whom/what it applies. Near the end of the period represented by the corpus, though, such semantic confusion was soon gone. "Terms like "memory" and "symbol structure" are now pinned down and defined in sufficient detail to embody their referents in precisely stated programs and data structures," write Simon and Newell by 1970 (148).

Root Metaphor:			A MACHINE IS A BRAIN	
	count	literal	metaphorical	marked
brain	331	321	10	
giant +	5	0	5	
artificial +	0	0	0	modifier
mechanical +	5	0	5	modifier
human +	32	32	0	modifier
Level:			Conventional	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

Table 6: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A BRAIN: ***brain****.*

Machines are Brains

The root metaphor A MACHINE IS A PERSON has a precursor in the form of a root metaphor that was common during the Second World War, one that inspired Von Neumann to use the word “memory”: that A MACHINE IS A BRAIN (Buchanan, 2005, 54). The two are connected in a metonymic sense. Strictly speaking, BRAINS STAND IN FOR PEOPLE is a *synecdoche* because the brain is a physical part of the whole person. By “whole person” we mean to include their behavior, personality, language and thoughts. This leads to a paradox, because although the synecdoche stands for a generic ‘person’, not any particular individual, it is still spoken of as having behavior, character and language which are expressed, in human culture, at an individual level.

Root Metaphor:			A MACHINE IS A PERSON	
	count	literal	metaphorical	marked
language	442	0	0	quotes (x5)
computer +	14	0	14	modifier
machine +	13	0	13	modifier
processing +	31	0	31	modifier
programing +	14	0	14	modifier
natural +	49	49	0	
human +	4	3	0	
ordinary +	29	29	0	
physical +	7	7	0	modifier
Level:			Conventional	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

*Table 7: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A PERSON: **language**.*

The historical context of the phrase comes from the giant machines that were built to decode messages intercepted from enemy lines and to help the military plan logistics for the battlefield (Buchanan, 2005, 54). The speed and power of their calculations lead to them being called “giant brains,” a metaphor that was firmly established in 1949 with the publication of *Giant Brains, Or Machines that Think* (1961) by computer scientist Edmund C. Berkeley. The use of terms such as neurons/neural/nerve ten years later to describe the inner workings of these machines was a natural extension of this metaphor and followed naturally from references to “memory”. “These new machines are called sometimes mechanical <u>brains</u> and sometimes sequence-controlled <u>calculators</u> and sometimes by other names,” writes Berkeley, in a sentence that is notable for the absence of the word “computer” (1961, vii).

In the Early AI Corpus we can see the metaphor of memory being combined with others in a way that allows the scientists to generate new theories as to how these “giant <u>brains</u>” might be made to perform more and more complex tasks. “The basic idea is that, whenever a <u>piece</u> of information is <u>stored</u> in <u>memory</u>, additional information should be <u>stored</u> with it <u>telling</u> <u>where</u> to <u>find</u> the <u>next</u> (associated) <u>piece</u> of information. In this way the entire <u>memory</u> could be organized like a long <u>string</u> of <u>beads</u>, but with the <u>individual beads of the string stored</u> in arbitrary <u>locations,</u>” writes Simon in 1971 (Crevier, 1993, 47). In addition to the mixed metaphors of “<u>memory</u>” and “a long <u>string</u> of <u>beads</u>,” information is spoken of as an ontological reality, as an object that can be “<u>stored</u>” in a “<u>location</u>” in “<u>piece</u>(s)”. The metaphor is written as a simile, with “like a” acting as a marker of a novel metaphor. Simon is well aware this is a new way to speak about memory and draws the readers’ attention to that fact.

By 1963, in the introduction to *Computers and Thought* (1963), the first de-facto textbook for scientists interested in artificial intelligence (and part of the Early AI Corpus; see Appendix A), editors Feigenbaum and Feldman describe a “<u>computer</u>” as any machine that has an “input” device to transform symbolic information external to the machine to internal language and an “output” device to perform the opposite. These two terms are not meant metaphorically. At the time, punch cards were literally "put in" to the machines and a printer would "put out" a series of conclusions derived from the computer's programming. Earlier we looked at examples of *catechresis*, the process of a metaphorical term becoming a new word, and in the process losing its metaphoricity. This appears to be the opposite process, a literal process becoming metaphorical when technology changes. The editors list the other requirements for a computer: a "memory" device, an "arithmetic unit", and a "control unit," which is described as the "executive of the computer organization…calling the other units into action when necessary" (1963, 2).

The use here of "language" to describe both the medium of input and the means of computation occurring inside the computer, is metaphorical in several stages. "Without stretching the ordinary usage of the term too far, we may say that a computer speaks a language, more specifically, a language-system with a specified primitive vocabulary, axioms, and rules of inference," writes pragmatic philosopher Sidney Hook (1960, 190). It is difficult to determine whether the instances of "language" in the text are strictly metaphorical (referring to computers) or literal (referring to humans) because often the reference is to "language" as theoretical concept of communication. The word is often paired with modifiers such as "human," "ordinary" or "physical" to refer to human language, and with "machine," "computer," "programming," or "processing" to refer to machine language. One of the key terms that was introduced around this time is "natural language," which generally refers to the type of language humans use in everyday discourse (as opposed to formal grammar or syntax), but it is often presented as the most difficult information for AI systems to parse because it is so idiosyncratic.

It is worth recalling Max Black's interaction view of metaphor here, because it suggests that through repeated use, not only will machines be thought of as having similar properties to brains, but brains will take on some of the properties of machines. The source domain (the human body) and the target domain (the machine) both have associated with them a set of "commonplaces" that should blur together (Black, 1962, 41). The metaphor functions in both directions; a "reciprocity of perspectives," in the words of Claude Lévis-Strauss. "The reciprocity of perspectives, in which man and the world mirror each other and which seems to us the only possible explanation of the properties and capacities of the savage mind, we thus find transposed to the plane of mechanized civilization," he writes in his seminal book *The Savage Mind* (1966, 222).

The root metaphors that cast machine behaviour and human behaviour as homologous seem to be inherent properties of the systems themselves for many of these early scientists working in artificial intelligence. "Metaphors can be dangerous not only in bewitching us into thinking that what they suggest really does exist but also in leading us to believe that the attributes normally possessed by any of the referents in the metaphor are possessed by the others," writes Mac Cormac (1985, 17). He asks us to consider Lévis-Strauss' "reciprocal perspective" seriously. "From the computer's perspective, perhaps humans do not do much thinking," writes Mac Cormac, with a "haphazard associations of images, emotions, shortcuts etc." (1985, 18).

Neural Nets

In the Dartmouth proposal, references to “neural nets”, “nerve nets” and “neuron nets” are equal (two each). Searches through the Early AI Corpus reveal that these were used as synonyms during this period. The key-word “neural” appears in the corpus 85 times and modifies the noun “net” or “network” only 14 times. It is used much more often in a literal sense about human brains to modify nouns like “process,” “activity,” or “mechanism”. The only time “neural” is used metaphorically is when it modifies the noun “net” or “network”. Collocation of “nerve” with “net” or “network” occurs eight times and “neuron net” only occurs three times. A fourth synonym is introduced in McCarthy’s paper as “nervous nets” but disappears after three uses, perhaps because the semantics were ambiguous, suggesting networks under emotional strain. Early versions of these terms were evident in the 1940s in what were called “neuro-logical networks,” referencing not biological neurological networks as we understand them, but artificially created “logical networks” that were modelled after neuron structure (Minsky and Papert, 1969, ix). The prefix ‘neuro’ is remarkably versatile and can be considered a *metaphorical morpheme*, a word part that can be

Root Metaphor:			A MACHINE IS A BRAIN	
	count	literal	metaphorical	marked
neuron	38	29	9	
+ net(work)	3	0	3	one instance of “netlike”
artificial +	5	0	5	modifier
neural	85	77	8	
+ activity	19	19	0	
+ net(work)	14	6	8	
nerve net(work)	8	0	8	
nervous net(work)	3	3	0	
Level:			Novel	
Ontological?			N	
Orientational?			N	
Systematic?			Y	

Table 8: Concordance data from 'Early AI Corpus' for the root metaphor A MACHINE IS A BRAIN: ***neuron, neural, nerve network, nervous network.***

mixed and matched with other terms to create new meaning, all under the suggestion of the overarching root metaphor A MACHINE IS A BRAIN.

The term itself was used by Marvin Minsky as early as 1954 for his PhD thesis entitled "Theory of Neural-Analog Reinforcement Systems and its Application to the Brain-Model Problem." (Minsky, 1954). In the text the term "net" is defined first as "simple sets of interconnected neurons" (Minsky 1954, 8) and then neural nets are described: "Each 'brain model' is formed of a small number of very large 'random' neural nets with a small number of channels connecting these nets" (op cit.).[8]

Minsky describes the metaphor in his own words in an interview in 2011: "There's one place [in the thesis] where I try to calculate what could be accomplished by loops of neurons that are arranged in circular pathways so that if you put a certain pattern in it will sort of echo around and I showed that mathematically that under some conditions the information that you originally put into such a loop will be gradually destroyed and the pulses will become equally spaced." (Web of Stories, 2011). The underlined metaphorical language reveals the metaphor as a productive one. In the thesis, "circular pathways" are mathematical abstractions, yet the consequences of "neurons" arranged in such an array lead to some conclusions that are described in physical or visual terms with words like "destroyed", "pulses" and "spaced". For proving this theorem, Minsky was awarded his PhD. He went on to experiment with the consequences of neural nets.

The term "neural net" places the reference to neurons in the phrase as an adjective, whereas the other two terms, "nerve nets" and "neuron nets" are noun phrases with no adjective. This suggests that as the theory of modelling brain processes was becoming more sophisticated, Minsky

[8] It is worth noting that in this period, Minsky physically made these networks. Artificial neurons were physical, analogue relays, perhaps another reason the word "machine" remained popular over the term "computer." When computers went digital in the 60s, the approach was to use symbolic Boolean functions instead of physical parts. Perhaps the term "machine" was no longer appropriate (Crevier, 1993, 26).

and others realized that their artificial neurons did not function in the same way as did biological neurons. They were, essentially, 'neuron-like' and not identical to real neurons. In the adjectival phrase the "net(work)" is the object of interest and it is modified by an adjective that suggests behavior but does not make any undue claims concerning biological equivalence.

Research is a Journey

The Early AI Corpus includes a series of *orientational metaphors*, the most basic of which is to treat the passage of time as extending forward, ahead of an observer, with past occurrences extending behind. The position of the observer is defined as the present. This is a fairly universal metaphor to describe the passage of time (Kovacses, 2005; Lakoff and Johnson, 1980). It conceptualizes the observer as physically moving forwards through physical medium, "time," on a journey.

Root Metaphor:			RESEARCH IS A JOURNEY	
	count	literal	metaphorical	marked
steps	89	1	88	
quest	4	0	4	
Level:			Dead	
Ontological?			Y	
Orientational?			Y	
Systematic?			Y	

Table 9: Concordance data from 'Early AI Corpus' for the root metaphor RESEARCH IS A JOURNEY: ***steps, quest.***

The opening sentence of McCarthy's paper uses the preposition "in" to refer to the physical place where this metaphorical journey will begin ("in Dartmouth"). Prepositions like *in, with*, and *on* are known as delexicalized words and do not add much to our analysis of meaning in the context of artificial intelligence as they tend to be fairly arbitrarily assigned (Deignan, 2008, 36). The use of "in" is a thoroughly dead metaphor and as such, prepositions will be largely ignored for the purpose of analysis.[9]

The journey metaphor crops up through the use of plenty of other words in McCarthy's prose, including the following:

RESEARCH IS A JOURNEY that can:

Lead/led to (x3), be "achieved in the future", develop in "parallel", proceed (x2), advance, have "goals" (x3), previous, preceding, further, direction/direct (x5), so far (x4), along these lines, way (x9), field, area, range, course of action, progress (verb: x3), obstacle, place (x3), trial and error (x4), guide (x3), surmount, foresight, approach (x8).

The scientists are positioned as the protagonists in a hero's tale, advancing through the predictable stages of a journey. The mythical quality of such a journey can be seen in words such as "<u>devote</u>" (my research to…), or "they are engaged in a <u>quest</u> for an information processing theory" (Feigenbaum and Feldman 1963, 205). If RESEARCH IS A JOURNEY and the scientist is the journeyer, we might expect there to be a series of obstacles for the protagonist to overcome in keeping with the archetype of the 'Hero's Journey' (Campbell, 2004). In McCarthy's text, these obstacles are either self-imposed by the scientist in the form of computational "problems," or are a function of a "<u>hostile</u> environment" in which the computer needs to navigate a solution. In this way, the computer is like a proxy fighting on behalf of the scientist. The machine's success in the face of obstacles represent the scientist's success on the journey to acquire knowledge. The two work together, as a team, like a knight and his squire.

---

[9] In French it is common to say "Je suis <u>à</u> Paris" which translates literally as (I am <u>at</u> Paris). Even in English it is common to say "I am <u>at</u> the train station" even though one might be "in" a building. The preposition "on" is often also used, such as "I am <u>on</u> the train," even though one might be "in" the vehicle.

In this sense, the machine is sometimes framed as the protagonist. Readers empathize with the machine as, "it can <u>progress</u> <u>through</u> a complicated environment only <u>through</u> <u>painfully</u> <u>slow</u> <u>steps</u>, and in general will not <u>reach</u> a <u>high level</u> of behavior," (McCarthy, 1956, 5) or as "<u>taking</u> 10,000 steps" (8). McCarthy is asking for support for the journey; money for better armour, a better line of defense. We can read the proposal through Latour's theory, that "the text builds a little story in which something incredible (the hero) becomes gradually more credible because it withstands more and more credible trials" (1987, 53).

Problems are Adversaries

Any problems the machine is challenged to solve are understood to be inherently resistant to being solved, giving them agency as adversaries. The engineer "<u>provides</u> the machine with a <u>set</u> of rules to <u>cover</u> each contingency which may <u>arise</u> and <u>confront</u> the machine." (McCarthy, 1956, 6). Several times the problem is cast as a stealthy adversary that remains hidden until the "<u>underlying</u>" problem is revealed by the computer. The role of the scientist is to "<u>give</u>" the computer the knowledge it needs (in the form of instructions or rules) to conquer the problem. "The 'things' behind the scientific text are thus similar to the heroes of the stories," continues Latour, "they are all defined by their performances. Some in fairy tales defeat the ugliest seven-headed dragons or against all odds they save the king's daughter; others inside laboratories resist precipitation or they triumph over bismuth" (1987, 89).

A special case of the PROBLEMS ARE ADVERSARIES root metaphor is THE ENVIRONMENT IS AN ADVERSARY. The machine in question is spoken of as being in conflict with either arbitrarily created problems, or with the environment itself. "A still more complex set of rules might <u>provide</u> for uncertainty about the environment, as for example in playing tic tac toe one must not only consider his next move but the various possible moves of the environment (his opponent)" (McCarthy, 1956, 7). The environment is spoken of in 161 different instances in the Early AI Corpus, often blurring the lines between the physical environment in which the computer is placed and the virtual environment the engineer controls. Often, the writers do not distinguish between the two. An environment is often an agent that is "providing" data to the machine, or described as being "hostile" or "indifferent" to the solving of problems. Sometimes the failings of the scientist are included in what constitutes the hostile environment, "For the machine, randomness will probably be needed to overcome the shortsightedness and prejudices of the programmer" (McCarthy, 1956, 8).

Hills and Trees

As a journey unfolds in a "problem space," metaphors that evoke physical landscape can be used to describe particular features of the machine's process. In 1950 mathematician Claude

Root Metaphor:		RESEARCH IS A JOURNEY	
	count	literal	metaphorical	marked
problem	1,666			
+ space	85	0	85	
attack	39	1	38	
+ [a] problems	19	0	19	
resist	8	0	8	
environment	161	0	0	
hostile +	2	0	2	
task +	36	0	0	
solution	301	0	0	
found [a] +	11	0	11	
seek [a] +	6	0	6	
Level:		Dead		
Ontological?		Y		
Orientational?		Y		
Systematic?		Y		

*Table 10: Concordance data from 'Early AI Corpus' for the root metaphor RESEARCH IS A JOURNEY: **problem, attack, resist, environment, solution**.*

Shannon described a "branching tree" for each of a series of chess moves, one that resulted in 10^{120} possibilities for game play (1950). Trees are ontologically real objects whose form can be used to map a decision-making progress. This metaphor proved to be enormously productive for researchers once it was established. A number of entailments of this metaphor from the domain of "landscape" or "nature" were used to refine the core theory of how artificial intelligence might be achieved. Quotes like this are typical: "Figure 1 shows a 'tree' of moves which might be investigated …. The actual branchings are much more numerous than those shown, and the 'tree' is apt to extend to as many as 20 levels" (Samuel, 1959, 74). Note the metaphors marking one of the first uses of this metaphor. In comparison, in the entire Early AI Corpus, out of 127 instances, there are only seven instances of "tree" marked with quotes.

Root Metaphor:			RESEARCH IS A JOURNEY	
	count	literal	metaphorical	marked
hill	16	7	9	quotes (x1)
climb	8	0	8	
peak	14	0	14	
hill-climbing	30	0	30	quotes (x3)
explore	66	0	0	
tree	127	11	116	quotes (x7)
branching +	4	0	4	
decision +	3	0	3	
branch	74	0	74	
prune	9	0	9	
Level:			Novel	
Ontological?			Y	
Orientational?			Y	
Systematic?			Y	

Table 11: Concordance data from 'Early AI Corpus' for the root metaphor RESEARCH IS A JOURNEY: ***hill, climb, peak, explore, tree, branch, prune****.*

The entailments of the tree metaphor are used by many of the writers of this period to include terms such as "culling branches" (Samuel, 1959, 80); "pruning of the tree" (Slagle, 1963, 197); "fertile"; "sterile" (Feigenbaum and

Feldman, 1963, 6) "grows" (Newell et al. 1958, 50); and "fruitful paths" (op cit.). Newell, Simon and Shaw, the other member of their team at Carnegie-Mellon, were especially fond of this metaphor in the corpus. "We are more accustomed to thinking of problem-solving search as generating lushly branched trees or partial solution possibilities which may grow to thousands, or even millions, of branches, before they yield a solution," they write in 1976 (122). They ask, rhetorically, "from what node in the tree shall we search next? and what direction shall we take from that node?" (1976, 123). In a remarkable sentence rich with imagery, Newell and Simon write of a colleague's work in psychology, that "De Groot found that the tree of move sequences explored by players did not originate as a bushy growth, but was generated, instead, as a bundle of spindly explorations, each of them very little branched. After each branch had been explored to a position that could be evaluated, the player returned to the base position to pick up a new branch for exploration" (Simon and Newell, 1970, 153). The "tree" metaphor is enormously productive as a "probative tool" (Miall, 1982, 101) for scientists to communicate with each other and explore the consequences and limitations of their theories.

Another topographical metaphor that appeared often in the literature is to describe the solving of problems by machines as "hill-climbing." Simon and Newell describe this as a "progressive test" of a suggested solution to a given problem set. "In climbing a (not too precipitous) hill, a good heuristic rule is always to go upward. If a particular spot is higher, reaching it probably represents progress toward the top. The time it takes to reach the top will depend on the height of the hill and its steepness, but not on its circumference or area—not on the size of the total problem space," they write (Simon and Newell, 1970, 152). Metaphorical words are remarkably frequent in this passage and accord well with the root metaphor RESEARCH IS A JOURNEY.

Psychologist Oliver Selfridge describes this in terms that casts the hero as a "blind man trying to climb a hill. There may be, of course, many false peaks on which one may find oneself trapped," he writes (1960, 520). "The main peak may be very prominent," he continues, "but unless it has wide-spread foot-hills it may take a very long time before we ever begin to gain altitude" (1960, 521). Minsky's hill-climber is not blind but is "standing on the hill in a fog so dense that only the immediate vicinity is visible" (Minsky and Papert, 1969, 257)[10]. Minsky is particularly troubled by the "false peak" phenomenon, the lighting upon a "local peak" which results in "much aimless wandering" (Minsky, 1961, 411). He introduces specific mathematical tools to counter this tendency including the "steep ascent" method designed to "get around obstacles" (Minsky and Papert, 1969, 266) or to follow the "contours" of a particular problem (Minsky and Papert, 1969, 180).

The Ghost in the Machine

If neurons are used to model the internal structure of computers as part of the root metaphor A MACHINE IS A PERSON (or, its synecdochal form, A MACHINE IS A BRAIN) then the interaction view of metaphor would suggest its opposite: that A PERSON IS A MACHINE or A BRAIN IS A MACHINE. This is a mechanistic view of the internal structure of the human body that has long roots, tracing back to the structural metaphors of Descarte's clock (Lauden 1966) or Julien Offray de La Mettrie's *Man a Machine* (1750). Of his own analysis of the different root metaphors around the world, Stephen Pepper claimed that "one of the seven or eight root metaphors was "__________ is a machine," where the blank space would often be "human" (1972, 329). Canadian AI researcher Daniel Crevier notes that, "to an engineer's eye, the cortex presents striking similarities with a structure universally present in computers: the printed circuit board" (1993, 284).

[10] This is a good example of an extended metaphor that doesn't parse well at the lexical level.

The key phrase here is “to an engineer’s eye,” implying that under the scrutiny of a different professional, similarities would be mapped to a different domain.

One of the texts included in the Early AI Corpus contains an innovative metaphor coined to derogate Descarte’s instance on dualism between a mechanistic body and an immaterial soul. In 1964, Hilary Putnam calls attention to “a ghost in a machine,” a slight divergence from the original usage introduced by Gilbert Ryle in 1949 of “the dogma of the Ghost in The Machine” (Ryle 1949, ---). Putnam’s use can still be considered novel based on the quotation marks around the term, as does Ryle’s use of capital letters: “a robot could be psychologically isomorphic to a disembodied spirit or to a “ghost in a machine” just as well,” he writes, showing more sympathy for the Cartesian view than does Ryle (Putnam, 1964, 678). The metaphor became more conventional upon the publishing of Arthur Koestler’s book of the same name in 1967.

An artful critique of the emerging field of artificial intelligence was published in 1958 in a paper entitled “How to Make a Computer Appear Intelligent” (1961) by computer scientist Joseph Weizenbaum. He created a program, ELIZA, that could hold conversations with a user through a text interface. His most famous script, DOCTOR, cast the computer as a psychotherapist, responding to statements posed by the user with questions that simply parroted back the main clause of the statement in question form (Crevier, 135). He meant it as a parody of the current fashion for ‘intelligent seeming’ programs, but many reporters and scientists didn’t catch the meaning of the word “appear” in the title.

In 1972, near the end of the first phase of growth of artificial intelligence research, philosopher Hubert Dreyfus extended the critiques from Ryle, Putnam and Weizenbaum in his book *What Computers Can't Do: A Critique of Artificial Reason* (1972). Dreyfus takes aim at the hubris of scientists in the 50s and 60s who made bold claims about how closely computers could perform human-like tasks of intelligence. “<u>Intuition</u>, <u>insight</u>, and <u>learning</u> are no longer exclusive possessions of humans: any large high-speed computer can be programmed to exhibit them also," claimed Simon and Newell as early as 1958 (6). In the same paper the authors claimed that within ten years, a computer would be the world chess champion, be able to write music with "aesthetic value," prove an "important new" mathematical theorem, and render the field of psychology obsolete (1958, 8).

By 1972, none of these things had happened, and Dreyfus criticized the researcher's blindness to their own assumptions. The growth of the "giant brain" metaphor, and the assumption that its entailments must also be true (neurons, memory, learning, self-improvement) was determining the direction of their research. Dreyfus puts the researcher's optimism into historical context by showing that "the brain [was] always understood in terms of the latest technological innovations." In the first half of the 20th century, the guiding metaphor was that of the brain as a telephone switchboard, then after WW2 the metaphor shifted to the brain being akin to computer hardware, then software, and finally interconnected like nodes on the world wide web (Brooks, 2015, 295). "Metaphors for the brain will continue to evolve as technology evolves," writes robotics engineer Rodney Brooks, "with the brain always corresponding to the most complex technology we currently possess" (2015, 296). This metaphorical fallacy lead to the "the naive assumption that man is a walking example of a successful digital computer program," says Dreyfus (1972, 71). There is an "a priori assumption that the mind must work like a heuristically programmed digital computer," he writes, a view that is not supported by evidence according to 20th century phenomenologists (Dreyfus, 1972, 99).

To come full circle with our analysis, Lakoff and Johnson make a similar point in *Metaphors we Live By* (1980) which came out only a few years later than Dreyfus' book, "There is a corollary of this that has to do with the issue of whether a computer could ever understand things the way people do. The answer we give is no—simply because understanding requires experience, and computers don't have bodies and don't have human experiences" (1980, 221).

These "metaphor wars" can be considered as an example of a "social drama" as described by Turner (1974, 17). The arena in which this drama is played out is the halls of academia, and the medium of the conflict is scientific articles and books, as is tradition in academia. "Arenas are the concrete settings in which paradigms become more transformed into metaphors and symbols," writes Turner about the context of social dramas. We can consider Dreyfus' arguments as imposing "a trial of strength between influential paradigm-bearers…. Social dramas represent the phased process of their contestation," (op cit.).

Dreyfus ends his devastating critique of artificial intelligence with yet another metaphor, comparing the scientists to medieval alchemists: "Alchemists were so successful in distilling quicksilver from what seemed to be dirt that, after several hundred years of fruitless efforts to convert lead into gold, they still refused to believe that on the chemical level one cannot transmute metals," he writes (1972, 215). It's hard to imagine a more devastating critique to a practicing scientist than a comparison to a pre-scientific wizard. In the arena of contested metaphors, artificial intelligence was dealt a devastating blow by Dreyfus' book.

To take Turner's "social drama" structure one step further, Dreyfus' book was a *breach* in the commonplace understanding of artificial intelligence and initiated a *crisis* in the field that only calmed down when Dreyfus agreed to play chess against one of Marv Minsky's student's computer programs in 1967 (Turner, 1976). Dreyfus lost the game. As an opportunity for *redressive action*, the chess event failed, causing Dreyfus to double-down on his criticisms, publishing a series of books and papers culminating in the aptly-titled *What Computers* Still *Can't Do* (1992).

At this point it becomes clear that the machines are not the "heroes" of the RESEARCH IS A JOURNEY metaphor; it is the scientists themselves who need to withstand the trials, acting as they do as “spokesmen” for particular theories. The machines are the knights and the scientists are their squires. “Some of these trials are imposed… by the scientific objector” or “dissenter,” writes Latour (1987, 89), a moniker one imagines Dreyfus would wear with pride. The chess game can be imagined as a ‘trial of strength’ for the theory of artificial intelligence, cast as a contest between Marv Minsky, spokesman for neural networks, and Hubert Dreyfus, spokesman for phenomenology.

Knowledge is an Object

Many of the ontological metaphors we have thus-far been considering can be summarized with the common root metaphor KNOWLEDGE IS AN OBJECT. As Alice Deignan states, “it is very difficult to find words that are *not* metaphorical to describe certain abstract things” (2008, 17). In McCarthy’s proposal, he writes of knowledge as something that can be “carried out,” “obtained,” “acquired” or “circulated.” Furthermore, the *knowledge object* can be perceived visually, suggesting that SEEING IS UNDERSTANDING. Results from specific experiments can be “shown” to exhibit a desired principle in accordance with a hypothesis. One sentence, rhetorically designed to juxtapose the appeal of a certain theory with its newness in the literature, is described as being “attractive and clearly incomplete” (McCarthy, 1956, 2). This suggests that the view of the knowledge object can be either obstructed or unobstructed (and that this state is under the control of the scientists, who are able to “clarify” a view), and that it can exhibit aesthetically pleasing characteristics, one of which is its state of completeness. The knowledge in the heads of the individual scientists is a physical object that “lies” in certain locations and is “connected” to other ideas. Again, using McCarthy’s text as typical of the corpus, knowledge has the following characteristics:

KNOWLEDGE IS AN OBJECT that can be:

Carried out (x3), obtained (x3), acquired, circulated, (in)complete (x2), attractive, clear (x6), connected, valuable, accumulated, worthwhile, looked for, shown, seen (from a “point of view”), “cast some light on”, “appear” (x3), given, provided, stored (x2), used (widely), extended, approached, made to appear, near, seen, handled, injected, supported (x4), set up, constructed (x3), built (x3), linked, seen, had on hand (or the other hand) (x2), worked on/out (x10), orderly.

Root Metaphor:			KNOWLEDGE IS AN OBJECT	
	count	literal	metaphorical	marked
knowledge	213	213	0	quotes (x5)
stored +	4	0	4	
+ base	13	0	13	
general +	3	0	0	
specifc +	1	0	0	
Level:			Dead	
Ontological?			Y	
Orientational?			Y	
Systematic?			Y	

Table 12: Concordance data from 'Early AI Corpus' for the root metaphor KNOWLEDGE IS AN OBJECT

KNOWLEDGE IS AN OBJECT that has:

Size (of a calculation), worth, upper and lower bounds, scope, basis (x3), structure (x2), matter, sets (of rules, of experiments) (x7), higher and lower levels/functions (x2), stages, flow, high order.

Chess

The game of chess played an enormous symbolic role in the social drama that unfolded with Dreyfus. It was the 'trial of strength' par excellence and Minsky and his colleagues were tremendously pleased when their machine defeated Dreyfus. Since 1950, chess had been treated in the artificial intelligence community as a metonym that stood for human intelligence. Moreover, it is a physical metonym, a game traditionally played on a flat board with three-dimensional wooden or plastic pieces. We can consider this metonym as further evidence of the ontological metaphor KNOWLEDGE IS AN OBJECT, and more specifically, that the spectacle of the game being played is a ritual enactment of that knowledge. The underlying assumption was that chess-playing represented the penultimate expression of human intelligence. Thus, if a computer could be programmed to play chess as well as, or better than, the best human players, the programmers would have removed the metaphorical veil from the term "artificial <u>intelligence</u>" and actually created something they considered intelligent. "Chess is generally considered to require 'thinking' for skillful play," writes mathematician Claude Shannon in 1950 (4). In this paper, the

Root Metonym:			CHESS REPRESNTS KNOWLEDGE	
	count	literal	metaphorical	marked
chess	215	215	0	metonym
+ player	19	10	9	
human +	2	2	0	modifier
machine +	7	0	7	modifier
rule of thumb	8	0	8	idiom
heuristic	290			-
Level:			Conventional	
Ontological?			Y	
Orientational?			N	
Systematic?			Y	

Table 13: Concordance data from 'Early AI Corpus' for the root metonym CHESS REPRESENTS KNOWLEDGE.

words "artificial" and "intelligence" do not appear, in preference for the term "thinking". "A solution of this problem will force us… to admit the possibility of a mechanized thinking," Shannon writes (1950, 4).

"To code chess, a complete 'chess vocabulary' is built up…. This vocabulary consists of a set of processes for expressing basic concepts in chess," writes Newell et al. (1958, 64), using words that reinforce the idea that machines have language. For the machine, this vocabulary consists of a set of mathematical instructions, but in the context of the paper reporting on the results, Newell uses English chess terminology to explain the machine's logic: blind spots, retreat, attack, defend, take, threat, fork, exchange, sacrifice, capture, pin, opponent, block, en passant, drive, move, interpose, counterattack, and dead position. These are all ontological terms, either referring to the physical movement of the chess pieces or of the human figures the pieces are meant to represent.

"These programs, especially the Los Alamos one, provide real anchor points," writes Newell et al. "They show that, with very little in the way of complexity, we have at least entered the arena of human play—we can beat a beginner" (1958, 48), they conclude, using ontological terminology that again evoke Turner's concept of social drama.

Understandably, Hubert Dreyfus had something to say about chess. He accused researchers of exaggerating the minimal success they had with their early chess programs to "launch the chess machine into the realm of scientific mythology" (1972, xxxi). He traces their language throughout the 1950s to show how the chess-playing computer became a 'fact'. Originally, the chess programs were described as "good in the opening," but that they fell apart towards the middle of the game. Shortly afterwards, this turned into the claim that "machines are already capable of a good game," by cybernetics pioneer Norbert Weiner, and then even later, that machines could "counter the moves of a master" and eventually pull off "combinations as difficult as any that have been recorded in chess history" (1972, xxxi)

In 1959, computer scientist Arthur Samuel tweaked the metaphor slightly and used checkers as his model for human intelligence. He explains his methodology as a "neural-net approach" (1959, 71) to program computers. There is a puzzle implied by this computational metaphor, though. If intelligence emerges from the firing of neurons, then machines should be able to compute all the possible moves in a checkers decision "tree". Mathematically, though, this is impossible. On an 8 X 8 checkers board, there are "10^{40} possible moves. At 3 choices per millimicrosecond," it would take someone 10^{21} centuries to play a game (which is many times longer than the age of the Universe) (1959, 72). This method of computing is known as the "brute force" method suggesting an acknowledgment that the above calculation is not only not reflective of reality, but also inelegant and boorish (Feigenbaum and Feldman, 1963, 5). The assumption is that true human thought should be much more sophisticated. We see again that researchers use metonyms (here, a checkers board) as a platform on which to calculate the consequences of their theory and the results of the thought experiment gave them pause.

Acknowledging there was something missing from the "neural net" model, researchers introduced a new metaphor to describe human decision-making as based on "rules of thumb," an idiom used to suggest that humans don't calculate every move in a game of checkers or chess through "brute force" but are engaged in a "highly selective search" based on past experience (Newell and Simon, 1976). This method was known as *heuristics*, a term that itself is a historical metaphor, based on the Greek root *heuriskein* meaning "to find" (OED, 2019). Newell, Simon and Shaw conclude that "we believe that any information processing system—a human, a computer, or any other — that plays chess successfully will use heuristics generically similar to those used by humans" (1958, 65). One of the inventions of the heuristic method was the use of the now ubiquitous IF…THEN statements to encode these "rules of thumb" (Crevier, 1993, 156).

Knowledge is a Building

Ontological metaphors serve to clarify an object of study in the minds of scientists. Knowledge and strategies are often described as things to be "found," but if they prove to be hidden from view, they can also be "constructed." The root metaphor KNOWLEDGE IS A BUILDING allows for knowledge to be described in McCarthy's paper as having a "basis" (x3), as needing financial "support" (x4), as having "stages" (x3), "levels" (4) and "structure". It can be "built," and "constructed," and distinctions can be made between "higher" and "lower" levels of knowledge (with the higher levels previously reserved for uniquely human thought). In the Early AI Corpus the "objects" being constructed are often immaterial things such as "program," "model," or "problem space." Physical objects that are actually built by the scientists (or their eager grad students) are referred to as "constructed" less often than are the immaterial objects.

Root Metaphor:			KNOWLEDGE IS A BUILDING	
	count	literal	metaphorical	marked
construct	213	38	251	Y – quotes in 3 instances
+ program	16	0	16	
+ model	10	0	10	
+ problem space	7	0	7	
+ robot	4	4	0	
+ machine	11	11	0	
+ computer	3	3	0	
frame	255	0	255	
global +	3	0	3	
room +	6	0	6	
subframe	35	0	35	
Level:			Dead	
Ontological?			Y	
Orientational?			Y	
Systematic?			Y	

Table 14: Concordance data from 'Early AI Corpus' for the root metaphor KNOWLEDGE IS A BUILDING

A visual metaphor was used by Minsky to try and explain how heuristics worked. He suggested "frames" as a "data-structure" for representing a stereotyped situation like "being in a certain kind of living room, or going to a child's birthday party" (1974, 33). It is included in this section because Minsky so often referred to frames in structural terms in keeping with KNOWLEDGE IS A BUILDING and specifically as a window-frame (as opposed to a frame for a picture) (Crevier, 1993, 173). Viewers can organize "chunks" of experience that fit into "slots" of the "frame". One can see through the frame in three dimensions: "For visual scene analysis, the different frames of a system describe the scene from different viewpoints, and the transformations between one frame and another represent the effects of moving from place to place" (Minsky, 1974, 2)

The entailments here suggest morphemic variations of the word including "framework," found in the text as "frame-systems," "inter-frame structures," "niche-frame," "space-frame," "conventional frame," "frame-oriented scenario," "super-frame," and terms that cross with other metaphors previously mentioned such as chess ("fork frame") and A MACHINE IS A BRAIN ("a great collection of frame systems is stored in permanent memory" (Minsky, 1974, 8)).

Perceptrons

There are several novel structural metaphors that are introduced in the Early AI Corpus that disappeared as the discourse unfolded over the years. Instead of systematic metaphors, which move from novel to conventional, here we can speak of *one-shot metaphors* that die out before their level of metaphoricity fades through repeated use.

McCarthy makes a reference to the theory of a colleague, in the following passage: "He [Craik] suggests that mental action consists basically of constructing little engines inside the brain which can simulate and thus predict abstractions relating to environment". He goes on to explain "then the engine operates…" (1956, 6) Metaphorical use of the term engine appear five times throughout McCarthy's propsal, and never literally. But in the Early AI Corpus the "little engines" is not to be found, suggesting other scientists didn't find the metaphor 'good to think with.'

In 1958, just after the Dartmouth project, a psychologist at Cornell, Frank Rosenblatt, coined a new term to refer to neuron-like nodes, something he called "perceptrons." In his famous paper, he described hundreds of little cameras that are connected to "neurons" that are programmed to interpret signals in order to identify visual patterns (Rosenblatt, 1958). The name itself is a curious mash-up of two terms: *perceiving* + *automaton*, both of which are historical metaphors. "Perceive" has an ontological root in the Latin, *per* (thoroughly) + *capere* (to grasp), while "automaton" has its roots in the Greek, *autos* (self) + *matos* (thinking) (OED, 2019). The combination of these two words is thus a *neologism*, as the word itself does not have the metaphorical meaning suggested by its roots (thus, not a strict example of catechresis). Yet creating a machine that "perceives" like a human being is a metaphorical conceit, one that Rosenblatt famously embellished for the *New York Times* in a 1958 article that stated that the perceptron was "the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence" (Olazaran, 1996, 621).

Root Metaphor:			N/A	
	count	literal	metaphorical	marked
engines	18	0	9	
perceptrons	69			catechresis
Level:			Novel	
Ontological?			Y	
Orientational?			N	
Systematic?			N	

Table 15: Concordance data from 'Early AI Corpus' for the novel metaphor COMPUTER 'NEURONS' ARE 'PERCEPTRONS'

Rosenblatt describes the perceptron system as a self-evident extension of the similarity between "the simple on-off units" in computers, and neurons that fire in the brain. "The analogy between the perceptron and biological systems should be readily apparent to the reader," he writes (1958, 387). The term perceptron is fairly malleable in the text and can be modified as a "photo-perceptron" (1958, 390) or a "simplified perceptron" (1958, 391). The perceptron itself "appears to be capable of trial-and-error learning," he claims (1958, 402).

The excitement over perceptrons was short-lived. In 1969, Marvin Minsky and Seymour Papert wrote a famously scathing critique of Rosenblatt's theory, performing calculations that showed perceptrons could not perceive certain types of patterns (Olazaran, 1996). In this scenario, Minsky is now in the position of Latour's "dissenter," creating trials of mathematics too onerous for the hero-figure of the perceptron, represented by its spokesperson Rosenblatt, to overcome. Minsky admits as such in the introduction to the new edition of *Perceptrons*, "intellectual battle lines began to form along such conceptual fronts" (1988, xi).

Minsky and Papert's criticism unfolds by tacitly accepting the root metaphor A MACHINE IS A PERSON, even while critiquing Rosenblatt's claims. "Perceptrons make decisions… by adding up evidence obtained from many small experiments," they argue (1988, 4), however, "usually they are treated so loosely that the species of computing machine defined by them belongs to mythology rather than science" (1988, 2). This is another version of Dreyfus' dismissal of science as pseudo-science. The name "perceptron" was perhaps too evocative, as it gave researchers a false sense that computers were actually perceiving the world around them. "The perceptron seemed powerful enough in function, suggestive enough in architecture, and simple enough in its mathematical definition," write Minsky and Papert, "yet understanding the range and character of its capabilities presented challenging puzzles" (1988, 249).

Pandemonium

Another attempt at creating a 'theory constitutive metaphor' had as its inspiration John Milton's *Paradise Lost.* In 1959, computer scientist Oliver Selfridge created his version of an AI system he called "Pandemonium," based on Milton's description of hell (Crevier, 1993, 40) suggesting a novel metaphor that received knowing nods from colleagues in artificial intelligence: COMPUTER PROCESSING TAKES PLACE IN HELL. Selfridge called his version of neurons "demons," in keeping with the metaphor. They came in four types: image demons, feature demons,

Root Metaphor:			N/A	
	count	literal	metaphorical	marked
Pandemonium	39	0	39	quotes (x3)
demon	73	0	73	
shriek	10	0	10	
shout	7	0	7	
Level:			Novel	
Ontological?			N	
Orientational?			N	
Systematic?			N	

Table 16: Concordance data from 'Early AI Corpus' for the novel metaphor COMPUTER PROCESSING TAKES PLACE IN HELL

cognitive demons and decision demons and would work in parallel within their groups to identify patterns. If a demon detected something it "recognized" it would "yell," "shout" or "shriek". The "decision demon" is the most powerful (sometimes called "Lucifer" by other researchers) and it "listens" for the "loudest shriek" and attempts to discern what the pattern is (Selfridge, 1959).

Selfridge himself prefaces his paper with a moment of self-awaress. “We are not going to apologize for a frequent use of anthropomorphic or biamorphic terminology. They seem to be useful words to describe our notions,” he writes (1959, 513. The Pandemonium metaphor, with the neuron/perceptron/node role filled by demons was also fairly productive in how it led Selfridge’s thinking. Demons can be thought of as contained in a “box” (1959, 514), and modified, much like neurons and perceptrons, with prefixes and adjectives resulting in “worthy demons” (1959, 521), “useful demons” (1959, 522), and “subdemons” (1959, 522).

One of the consequences of this particular metaphor is that Selfridge writes of his “demons” as if they were a colony in a petri dish to be manipulated through the laws of biology. He speaks of his attempts to “control the mutations in subdemon selection” (1959, 523). “If they serve a useful function they survive, and perhaps are even the source for other subdemons who are themselves judged on their merits,” he writes (1959, 523). Strictly speaking, this second quote accords closer to the model of intelligent design that it does natural selection, but the biological parallels are clear. In an appendix to Selfridge’s paper that appears in the Early AI corpus, there is QandA period that shows how the metaphor is used as a scaffold to think about human intelligence by Selfridge’s colleagues. John McCarthy makes an interesting point when he suggests that “If one conceives of the brain as a pandemonium - a collection of demons - perhaps what is going on within the demons can be regarded as the unconscious part of thought, and what the demons are publicly shouting for each other to hear, as the conscious part of thought” (1959, 527).

That said, the Pandemonium model did not ‘stick’ and its only legacy is the presence, in modern computing vocabulary, of “daemons,” describing computer programs that run in the background of a user’s desktop; no shrieking or yelling required.

Revolution

If we look past the linguistic metaphors used to make sense of the "ghost" inside the computing machine, we also notice examples of researchers using metaphors to explain their own position in the field of AI. Scientists often compared their work to other well-known moments of paradigm shift in science, suggesting implicitly (and sometimes explicitly) that their current research deserved as much attention as these historical moments. In 1976, as criticisms started mounting on AI researchers who were missing their predicted targets, Newell and Simon explicitly place their "Physical Symbol System Hypothesis" as "the underlying qualitative theory" of AI next to several other well-known theories. "All sciences characterize the essential nature of the systems they study," they write, "they set the terms within which more detailed knowledge can be developed". They compared their theory to the cell doctrine in biology, to plate tectonics, to the germ theory of disease, and to the "doctrine of the atom" (1976, 116). These are powerful analogies to use because they suggest a whole series of "associated commonplaces," especially for scientists who know how crucial these theories are to the functioning of their respective disciplines.

"We are in the position of nineteenth-century chemistry which postulated atoms on the basis of observations of chemical reactions among molecules, and without any direct evidence for their existence," write Newell and Simon, "or in the position of classical genetics, which postulated the gene before it could be identified with any observed microscopic structures in the cell" (1970, 158). The germ theory, in particular, is used as an analogy to suggest how to test implications of their theory. The germ theory suggests that microscopic organisms cause diseases, so once as a disease has been identified, researchers work backwards to find the pathogen that causes it. Similarly, in AI, researchers identify a human task that calls for intelligence and then "work backwards" to recreate the symbolic-level logic needed to recreate it (1976, 118).

McCarthy worked his way through a puzzle by using genetics as a heuristic guide. He wondered how a simple set of coded rules could give rise to complex and interesting behaviour. “It is as though the human genetic structure were represented by a set of blue-prints,” he writes with a hedge at the beginning of the sentence marking novelty. “Then a mutation would usually result in a wart or a failure of parts to meet, or even an un-grammatical blue-print which could not be translated into an animal at all. It is very difficult to see how the genetic representation scheme manages to be general enough to represent the great variety of animals observed and yet be such that so many interesting changes in the organism are represented by small genetic changes. The problem of how such a representation controls the development of a fertilized egg into a mature animal is even more difficult,” he continues (1959, 3).

Galileo seems to be a favorite referent amongst the researchers, understood as *the* archetypical paradigm-shifter, the patron saint of empirical science. With regard to their use of computers, Simon and Newell admit their reliance on computers just as, “the telescope made sunspots and Jupiter's moons a part of Galileo’s science” (1958, 6). As electrical engineer Kenneth Forbus notes, though, such grandiose comparisons could also backfire. “Unfortunately, one of the best strategies for getting noticed is to declare a revolution, and that everything earlier must now be rejected,” he writes (2010, 347). “Artificial intelligence, plus its sub-fields connectionism, situated cognition, embodied cognition, and dynamical systems have all used this tactic,” he writes (op cit.).

The AI Winter

1976 is chosen as the end date of the Early AI Corpus because it coincides with the first big wave of funding cuts to strike the field of AI. The power of metaphors to suggest worthwhile experiments and produce useful software applications for funding governments was starting to wane. The researchers were “deluded by false expectations” writes Daniel Crevier and remained blind to the "disanologies" of their extended metaphors, focusing on the positive entailments, resulting in exaggerated claims and outlandish predictions, some of which are mentioned above.

In 1973, the infamous Lighthill Report was released by the British government as a summary of the past 20 years of government-funded AI research. The conclusions were bleak for the future of AI. "Enormous sums have been spent with very little useful result," wrote member of parliament Sir James Lighthill (1973, 10). "Most workers who entered the field around ten years ago confess that they then felt a degree of *naive* optimism which they now recognise as having been misplaced… In no part of the field have the discoveries made so far produced the major impact that was then promised," he concluded (1973, 9). Critics started talking of the logic used in AI as "brittle," meaning that it started to break down at the "edge" of a "problem space" (Moor, 2006, 88). Ironically, metaphor was only recognized at the end of this period of innovation in AI research as a potential tool to help AI programmers imbue creativity and flexibility into their systems. "Sometimes, in 'problem-solving' we [humans] use two or more descriptions in a more complex way to construct an analogy or to apply two radically *different* kinds of analysis to the same situation," writes Minsky. "For hard problems, one 'problem space' is usually not enough!" he concludes, the exclamation point emphasizing the counter-intuitiveness of the conclusion to his colleagues (1974).

Curiously, Lighthill claims that the researcher's use of language is one of the reasons for the field's failure, but not in the way we have been analysing here. "The style in which most papers on AI research are written is depressingly turgid or jargon-dominated and almost makes the authors appear antagonistic to the special human gift for relating to, and communicating with, other people in an imaginative way," he writes (1973, 17). A letter written to the RAND Corporation, lays out some of these feelings of alienation:

> "Semantics may have a lot to do with the degree of enthusiasm for supporting research in this area (artificial intelligence). Subjectively, the terms "intelligent

> machine" or "thinking machine" disturb me and even seem a bit threatening: I am a human being, and therefore "intelligent" and these inhuman devices are going to compete with me and may even beat me out. On the other hand, if the very same black boxes were labelled "problem solver," or even "adaptive problem solver," they would seem much more friendly, capable of helping me in the most effective way to do things that I want to do better, but, best of all, I'd still be the boss. This observation is wholly subjective and emotional..." (Armer, 1960, 397).

The letter-writer is perceptive. The battle over root metaphors is, in essence, a battle over power, and who has the right to determine the direction of research, and under what assumptions. "The important thing for those who use metaphorical means is to build up as elaborately as they may a structure of ideas, embodied in symbols, and a structure of social positions, symbolically expressed, which will keep chaos at bay and create a mapped area of security," wrote Victor Turner, from the perspective of analysing metaphor use in ritual life (1974, 297).

After 1974, most of the university research labs had their funding cut almost completely, so AI researchers were forced to look to the corporate world for funding. Big companies like Xerox, DEC, Compaq, IBM, and eventually Apple, funded programs designed to focus on specific applications such as hunting for chemical structures or more efficient computer compilers (Crevier, 1993, 156). The corporate funding eventually dried up, too, leading researchers to declare the onset of an "AI <u>Winter</u>" in 1984 that continued up until the late 1990s, a rather dramatic comparison to a common fear in the Western world of a "nuclear <u>winter</u>". The term "artificial <u>intelligence</u>" itself, and the root metaphors that gave it such resonance, started to dwindle in 1980s. Research continued in what were, in essence, neural networks, but the researchers called themselves "connectionists," and they considered themselves a completely new field, distinct from AI. "Connectionists meet in different conventions, express themselves in different journals, and speak a technical dialect different from that of AI researchers," writes Crevier (1993, 215).

Reflections

"The truth … is a mobile army of metaphors, metonyms, anthropomorphisms, in short, a sum of human relations which were poetically rhetorically heightened, transferred, and adorned, and after long use seem solid, canonical, and binding to a nation. Truths are illusions about which it has been forgotten they are illusions"

– Frederich Nietzsche (Hyde, 1998, 77)

The opinions of most of the scientists whose written work is represented in the Early AI Corpus would likely agree with Hobbes and Locke that communication in science is best achieved soberly and without rhetorical flourishes. In one of their papers, Simon and Newell claim that "our account today will be framed in ordinary language" (1970, 148). But this is followed up, only one page later, with an innovative metaphor not at all necessary for the communication of their content: "These are the bones of the theory. In the next pages, we will undertake to clothe them in some flesh" (1970, 149). The importance of human reasoning through ontological metaphors remained largely invisible to theorists trying to copy human intelligence in the period 1956-1976. In one passage, Simon and Newell puzzle through some of the differences between their human test subjects and the behaviour of the computer, observing that "the human subjects appeared able to move back and forth between concrete and abstract objects without treating the latter as belonging to a separate problem space," seemingly unaware that they are using ontological metaphors themselves to conceptualize abstract concepts as concrete entities (1970, 156).

In fact, without competence in decoding figurative language, many sentences in the literature would be unintelligible. "Memory management, for example, was the programmer's problem until the invention of garbage collection," is a sentence that uses Standard English yet the literal meaning diverges profoundly from the intended meaning (Buchanan, 2005, 57). Daniel Crevier helpfully explains to lay readers that "garbage collection" refers to the practice of "returning unused items to the end of a list/string," but in doing so activates yet another metaphorical domain of "strings" and "items" that in turn needs to be decoded (1993, 60). The use of the passive voice (as in "until the invention of garbage collection") is another rhetorical tool used widely in these papers to suggest that the subject lacks personal agency and that any characteristics applied to it emerge spontaneously from nature. This confuses the boundary between what the scientist does and what nature does, implying they are one and the same (Gopnik, 1972, 41).

The ubiquity of the root metaphor A MACHINES IS A BRAIN suggests to Earl Mac Cormac that the entire field of AI grew out of this one insight, a logical leap unique to Western science. "The discipline of AI was developed by those computer scientists, philosophers, and psychologists who accepted the metaphorical suggestion that computers engage in mental activities similar to those of human minds," he writes (Mac Cormac, 1985, 10). The reverse is also evident; that the focus on 'machines as minds' solidified the "reciprocal perspective" that 'minds are machines'. Philosopher of science Richard Boyd claims that, "a concern with exploring analogies, or similarities, between men and computational devices has been the *most important single factor* influencing post-behaviourist cognitive psychology" (emphasis added; Boyd 1979, 360). One of the consequences of this dogmatic adherence to the "computational metaphor" was that it led researchers to ignore the role of hormones on the function of the brain, and how sensitive brain function could be on environmental factors and hormone levels (Brooks, 2015). Trying to force newly discovered features of human cognition like dependence on hormone levels onto the computational metaphor gets unwieldy, a sign that perhaps the root metaphor is starting to shift and is reaching the limits of its explanatory powers.

Thoughts collected from some of the leading researchers in AI in recent years suggest that many of them were not aware of their dependence on metaphor. "It was always a terrible name, but it was also a bad idea," writes Roger Schank in a cantankerous essay from 2015. "It's not that artificial intelligence has failed; no one actually ever tried. (There, I said it!)" (Schank, 2015, 268). "The fact is that the name 'AI' made outsiders to AI imagine goals for AI that AI never had," he claims. His solution? “The field should be renamed ‘the attempt to get computers to do really cool stuff’ (2015, 269). At the 50th anniversary conference of the historical Dartmouth project, James Moor, the President of the AAAI, lamented that “there is still no general theory of intelligence or learning that unites the discipline” (Moor, 2006, 88). Perhaps stung by previous over-confident predictions, John McCarthy claimed at the conference that human-level AI would not be achieved until 2056, and even then it was not a sure thing (Moor, 2006, 90). At a recent conference in which scientists could submit programs for a ‘Turing test’, “it took maybe 30 seconds to figure out which was a human and which was a computer” (Schank, 2015, 269).

But what if we don’t accept the metaphor? Are we left only with Dreyfus’ phenomenology, and none of the advances in science and technology that have spun out of the AI industry? The Early AI Corpus used for this paper was analyzed with the use of a powerful <u>cloud</u> software based on AI algorithms designed to <u>clean-up</u> text and make it machine-<u>readable</u> for analysis. These words were typed on a computer with AI-developed voice-recognition technology, predictive text technology and many other tools that trace their roots back to AI labs. Metaphor use leads to an irresolvable paradox, an aporia: metaphors aid the process of reasoning and development in science in very tangible ways, yet they can over-simplify or even mis-characterize the complexity of the very systems being observed.

Conclusion

By using a unique combination of manual identification of metaphors through the Metaphor Identification Procedure (Pragglejaz, 2007) and automated key-word-in-context searches (Deignan, 2008), some structural metaphors appeared so thoroughly throughout the Early AI Corpus they have become *doxa* in the field of artificial intelligence. Metaphors that describe machines (computers) metaphorically suggest root metaphors such as A MACHINE IS A PERSON or A MACHINE IS A BRAIN. These root metaphors quickly became theory-constitutive metaphors and today their linguistic realizations can be considered dead. Not all metaphors describing machines (computers) were as equally successful, though. The metaphor COMPUTER PROCESSING TAKES PLACE IN HELL based on John Milton's *Paradise Lost* was not taken-up by colleagues as a useful way to think about their work, and its use was sporadic throughout the corpus. Similarly, the catechresis "perceptrons" did not catch on to describe computer "neurons," perhaps because the term Rosenblatt was trying to replace was already a metaphorical term taken from anatomy and was used widely by scientists to conceptualize the innards of their computers. Evidence of root metaphors that described the research process as an ontological reality such as RESEARCH IS A JOURNEY and KNOWLEDGE IS AN OBJECT were also ubiquitous throughout the corpus. The strength of these metaphorical foundations gave rise to productive metaphorical concepts such as "hills", "trees" or "frames" that allowed for real scientific advancement.

Previously collected data in the field of metaphor studies supports the conclusion that metaphors are widely used in scientific writing. According to a recent study by Raymond Gibbs Jr. (2017), the use of metaphor is highest in academic discourse including science with 18% of all written words being used in a metaphorical sense (excluding historical metaphors). This was followed, in order of decreasing frequency, by news stories (15%), fiction (11%), and finally everyday conversation (7%) (Gibbs Jr., 2017). In short, the use of metaphor *increases* if the topic of discourse is perceived to be rational and objective, suggesting patterns of communication opposite to those mandated by Hobbes and Locke.

If we accept the fact that the only way humans can coherently discuss non-material phenomena is through ontological and orientational language, then we can ask what accounts for one choice of conceptual metaphor over another and why they might have particular resonance with a particular audience (Charteris-Black, 2004, 10). This is the stated aim of the theory of Critical Metaphor Analysis (CMA) and would provide a good theoretical foundation for further investigations into metaphor use beyond what is covered in this paper. Instead of written texts, CMA would take as its object of study real-time communication between scientists. As a subset of Critical Discourse Analysis, CMA focuses on identifying novel metaphors and evaluates how they reinforce existing ideology or how they are used pragmatically by speakers (Charteris-Black, 2004, 34). In discourse analysis the "main interest is in how speakers use language to create meaning, metaphor being one tool in this task" (Deignan, 2008, 123). In scientific discourse one of these pragmatic goals is rhetorical, either in convincing others of the relative merits of a theory or applying for project funding. CMA theories make no claims on cognitive determinism but focus more on ideology and socio-cultural context of metaphor use. "CMA is an approach to metaphor analysis that… aims to reveal the covert (and possibly unconscious) intentions of language users" says Charteris-Black (2004, 34). He continues, "metaphorical interpretation is concerned with textual meaning—that is, identifying the type of social relations that are constructed through them" (2005, 35). CMA departs theoretically from Lakoff and Johnson's Conceptual Metaphor Theory by being more concerned with pragmatics than with semantics and its cognitive consequences (Cameron and Maslow, 2010, 79). As such, CMA is perhaps more relevant to anthropology than is

CMT as it is focused less on cognition than it is in making meaning of scientists' discourse as a social practice.

In general, the writing in the papers studied in the AI corpus from 1956-1976 is creative, illuminating, and clear. But this feat of communicative competence is not accomplished by banishing metaphor and other rhetorical tropes to the dustbin. On the contrary, it is only by embracing figurative language and the importance it holds for our thinking that we can make progress with our scientific theories and make them coherent with our other ways of knowing about the world.

References

Armer, P. 1960. “Attitudes toward intelligent machines.” In *Computers and Thought* (1963) edited by Edward A. Feigenbaum and Julian Feldman. New York: McGraw-Hill.

Basso, Keith. 1967. “Semantic Aspects of Linguistic Acculturation.” In *American Anthropologist*. 69 (5): 471-477.

Basso, Keith. 1973. “The Ethnography of Writing.” In *Explorations in the Ethnography of Speaking*, edited by Bauman and Sherzer. Cambridge: Cambridge University Press.

Basso, Keith. 1976. “Wise Words of the Western Apache" in *Meaning in Anthropology*, edited by Keith Basso and Henry A. Selby. Albuquerque: University of New Mexico.

Berkeley, Edmund C. 1961. *Giant Brains or Machines That Think*. New York: Science Editions Inc.

Black, Max. 1962. *Models and Metaphors: Studies in Language and Philosophy*. Ithaca: Cornell University Press

Black, Max. 1979. “More about metaphor.” In *Metaphor and Thought*, edited by Andrew Ortony. Cambridge: Cambridge U Press.

Bowdle, Brian F. and Dedre Gentner. 2005. “The Career of Metaphor.” In *Psychological Review*. 112 (1): 193–216.

Boyd, Richard. 1979. “Metaphor and Theory Change: What is ‘Metaphor’ a Metaphor for?” In *Metaphor and Thought*, edited by Andrew Ortony. Cambridge: Cambridge U Press.

Brooks, Rodney A. 2015. “The Computational Metaphor.” In *This Idea Must Die: Scientific Theories That are Blocking Progress*, edited by John Brockman. New York: HarperCollins.

Brown, Theodore L. 2003. *Making Truth: Metaphor in Science*. Chicago: University of Illinois.

Buchanan, Bruce G. 2005. “A (Very) Brief History of Artificial Intelligence.” In *AI Magazine*. 26 (4): 53-60.

Bush, Vannevar. 1945 “As We May Think.” In *The Atlantic*, July.

Cameron, Lynne and Robert Maslen. 2010. *Metaphor Analysis: Research Practice in Applied Linguistics, Social Sciences and the Humanities*. Studies in Applied Linguistics. London: Equinox.

Campbell, Joseph. 2004. *The hero with a thousand faces*. Princeton, N.J.: Princeton University Press.

Crevier, Daniel. 1993. *AI: The Tumultuous History of the Search for Artificial Intelligence*. New York: Basic Books.

Charteris-Black, Jonathan. 2004. *Corpus Approaches to Critical Metaphor Analysis*. London: Palgrave Macmillan.

Cox, Hyde and Edward Connery Lathem. 1966. *Selected Prose of Robert Frost*. New York: Holt, Rinehart and Winston.

Deignan, Alice. 2008. *Metaphor and Corpus Linguistics*. Philadelphia: John Benjamins Publishing Company.

Dreyfus, Hubert L. 1972. *What Computers Can't Do: A Critique of Artificial Reason*. New York: Harper and Row

English, Kathryn. 1998. "Understanding science: when metaphors become terms." In *ASp: Théorie et pratique des discours spécialisés.* 19 (22): 151-163.

Farrow, Jon. 2019. "CIFAR announces new research programs and AI Chairs." In CIFAR News. April 8. Accessed July 23, 2019. https://www.cifar.ca/cifarnews/2019/04/08/cifar-announces-new-research-programs-and-ai-chairs

Feigenbaum, Edward A. and Julian Feldman, Eds. 1963. *Computers and Thought.* New York: McGraw-Hill.

Fenwick M., E.P.M. Vermeulen, and M. Corrales 2018. "Business and Regulatory Responses to Artificial Intelligence: Dynamic Regulation, Innovation Ecosystems and the Strategic Management of Disruptive Technology." In *Robotics, AI and the Future of Law. Perspectives in Law, Business and Innovation*, edited by M. Corrales, M. Fenwick, and N. Forgó. Singapore: Springer.

Fernandez, James. 1974. "The Mission of Metaphor in Expressive Culture." In *Current Anthropology*, 15 (2): 119-145.

Fernandez, James W. 1977. "The Performance of Ritual Metaphor." In *The Social Use of Metaphor*, edited by J. David Sapir and J. Christopher Crocker. Philadelphia: University of Pennsylvania Press

Forbus, Kenneth D. 2010. "AI and Cognitive Science: The Past and Next 30 Years." In *Topics in Cognitive Science*. 2 (3): 345-356.

Forrester, Stefan. 2010. "Theories of Metaphor in Seventeenth and Eighteenth Century British Philosophy." In *Literature Compass*. 7 (8): 610-625.

Garett, Jan. 2007. "Aristotle on Metaphor." Accessed July 23, 2019. http://people.wku.edu/jan.garrett/401s07/arismeta.htm

Gibbs Jr., Raymond W. 2017. *Metaphor Wars: Conceptual Metaphors in Human Life*. Cambridge: Cambridge University Press.

Gopnik, Myrna. 1972. *Linguistic Structure in Scientific Texts*. The Hague: Mouton and Co. N.V.

Harré, Ron, J. L. Aronson and Eileen Way. 2000. "Apparatus as Models of Nature." In *Metaphor and Analogy in Science*, edited by Fernand Hallyn. Boston: Kluwer Academic Publishers.

Hesse, Mary. 1966. *Models and Analogies in Science.* Notre Dame: University of Notre Dame Press

Hesse, Mary. 1980. *Revolutions and Reconstructions in the Philosophy of Science*. Sussex: The Harvester Press Ltd.

Hobbes, Thomas. 1651. *Leviathan*. London: Andrew Crooke. Accessed July 22, 2019: https://socialsciences.mcmaster.ca/econ/ugcm/3ll3/hobbes/Leviathan.pdf

Hook, Sidney, Ed. 1960. *Dimensions of Mind: A Symposium*. New York: New York University Press.

Hyde, Lewis. 1998. *Trickster Makes This World: Mischief, Myth, and Art*. New York: Farrar, Straus and Giroux.

Kilgarriff, A. et al. 2014. "Sketch Engine ten years on." In *Lexicography*. 1 (1): 7-36.

Kovecses, Zoltan. 2005. *Metaphor in Culture: Universality and Variation*. Cambridge: Cambridge University Press.

Lakoff, George and Mark Johnson. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.

Latour, Bruno. 1987. *Science in Action*. Cambridge: Harvard University Press

Latour, Bruno and Steve Woolgar. 1986. *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press

Laudan, Laurens. 1966. "The clock metaphor and probabilism: The impact of Descartes on English methodological thought 1650–65." In *Annals of Science*. 22 (2): 73-104.

Leatherdale, W. H. 1974. *The Role of Analogy, Model and Metaphor in Science*. New York: American Elsevier Publishing Co. Inc.

Lévi-Strauss, Claude. 1966. *The Savage Mind*. London: Weindefeld and Nicolson.

Lighthill, Sir James. 1973. "Artificial Intelligence: A General Survey," from Cambridge University.

Locke, John. 1998. *An Essay Concerning Human Understanding*. London: Penguin Classics

Mac Cormac, Earl R. 1985. *A Cognitive Theory of Metaphor*. Cambridge: MIT Press.

Maranda, Elli Köngäs and Pierre Maranda. 1971. *Structural Analysis of Oral Tradition*. Philadelphia: University of Philadelphia Press.

McCarthy, J., M. Minksy, Rochester, L., and C. E. Shannon. 1956. “A proposal for the Dartmouth summer research project on artificial intelligence.” Accessed August 30, 2019. http://www-formal.stanford.edu/jmc/history/dartmouth.pdf

McCarthy, John. 1959. “Programs With Common Sense” in *Semantic Information Processing*, edited by M. Minsky. Cambridge: The MIT Press. 403-418

Mettrie, Julien Offray de La. 1750. *Man a Machine*. La Salle, IL: Open Court.

Miall, David S. Ed. 1982. *Metaphor: Problems and Perspectives*. New Jersey: Humanities Press

Minsky, Marvin. 1954. “Theory of Neural-Analog Reinforcement Systems and its Application to the Brain-Model Problem.” PhD Thesis. MIT.

Minsky, M. 1961. “Steps towards artificial intelligence.” In *Computers and Thought* (1963) edited by Edward A. Feigenbaum and Julian Feldman. New York: McGraw-Hill.

Minsky, Marvin and Seymour Papert. 1988. *Perceptrons*. Cambridge: MIT University Press

Minsky, Marvin. 1974. “A Framework for Representing Knowledge.” in *The Psychology of Computer Vision,* edited by P. Winston. New York: McGraw-Hill.

Moor, James. 2006. “The Dartmouth College Artificial Intelligence Conference: The Next 50 Years.” In *AI Magazine*. 27 (4).

Newell, Allen and Herbert A. Simon. 1976. “Computer Science as Empirical Inquiry: Symbols and Search.” In *Communications of the ACM*. 19 (3): 113-126.

Newell, A., J. C. Shaw, and H. Simon. 1958. “Chess playing programs and the problem of complexity.” In *IBM Journal of Research and Development*. 2: 320-335.

Olazaran, Mikel. 1996. “A Sociological Study of the Official History of the Perceptrons Controversy.” In *Social Studies of Science*. 26 (3): 611–659.

OED: Online Etymology Dictionary. 2019. Accessed May-July 2019. https://www.etymonline.com/

Oppenheimer, Robert. 1955. “Analogy in Science.” In *Sixty-Third Annual Meeting of the American Psychological Association*. San Francisco, California.

Palmer, Gary B. 1996. *Towards a Theory of Cultural Linguistics*. Austin: University of Texas Press

Pepper, Stephen C. 1972. *World Hypotheses*. Berkley: University of California Press

Pragglejaz Group. 2007. “MIP: A Method for Identifying Metaphorically Used Words in Discourse.” In *Metaphor and Symbol*, 22 (1): 1-39.

Putnam, Hilary. 1964. “Robots: Machines or Artificially Created Life.” In *The Journal of Philosophy*. 61 (21): 668-691.

Quinn, Naomi. 1991. “The Cultural Basis of Metaphor.” In *Beyond Metaphor: The Theory of Tropes in Anthropology*, edited by James W. Fernandez. Stanford: Stanford University Press.

Richards, Ivor Armstrong. 1936. *Philosophy of Rhetoric*. Oxford: Oxford University Press.

Rosenblatt, F. 1958. “The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain.” In *Psychological Review*. 65 (6): 386-408.

Ryle, Gilbert. 1949. *The Concept of Mind*. London: Hutchinson

Samuel, A. L. 1959. “Some studies in machine learning using the game of checkers,” in *IBM Journal of Research and Development*. 3: 211-229.

Schank, Roger. 2015. “Artificial Intelligence.” In *This Idea Must Die: Scientific Theories That are Blocking Progress,* edited by John Brockman. New York: HarperCollins.

Selfridge, Oliver. 1959. *Pandemonium: A Paradigm for Learning*. London: H.M.S.O.

Shannon, Claude E. 1950 “Programming a Computer for playing chess.” In *Philosophical Magazine*, Series 7, 41: 314.

Silcoff, Sean. 2019. “Canada’s Element AI is making a high-stakes gamble it can be a world-beater.” In *The Globe and Mail*. July 19, Report on Business, B6.

Simon, Herbert A. and Allen Newell. 1958. “Heuristic Problem Solving: The Next Advance in Operations Research.” In *Operations Research*. 6 (1): 1-10.

Simon, Herbert A. and Allen Newell. 1970. “Human Problem Solving: The State of the Theory in 1970.” In *American Psychology*. 26 (2): 145-159.

Slagle, J. 1963. *A heuristic program that solves symbolic integration problems in freshman calculus*. Cambridge: MIT Lincoln Laboratory.

Steen, Gerard. 2015. “Developing, testing and interpreting Deliberate Metaphor Theory.” In *Journal of Pragmatics*, 90: 67-72.

Steen, Gerard J. 2007. *Finding Metaphor in Grammar and Usage*. Philadelphia: John Benjamins Publishing Company

Thibodeau P.H. and L. Boroditsky. 2011. “Metaphors We Think With: The Role of Metaphor in Reasoning.” In *PLoS ONE*. 6 (2):

Thurbayne, Colin Murray. 1970. *The Myth of Metaphor*. Columbia: University of South California Press

Traweek, Sharon. 1988. *Beamtimes and Lifetimes: The World of High Energy Physicists*. Cambridge: Harvard University Press.

Turing, A. M. 1937. “On Computable Numbers, With an Application to the Entscheidungsproblem.” In *Proceedings of the London Mathematical Society*, Series 2, 42 (1): 230–265.

Turing, A. M. 1950. “Computing machinery and intelligence.” In *Mind*. 59: 433-460.

Turner, Victor. 1974. *Dramas, Fields, and Metaphors*. Ithaca: Cornell University Press.

Venkatesh, Narasima A. 2018. “Industry 4.0: Reimagining the Future of Workplace (Five Business Case Applications of Artificial Intelligence, Machine Learning, Robots, Virtual Reality in Five Different Industries).” In *International Journal of Engineering, Business and Enterprise Applications*. 26 (1): 5-8.

Von Neumann, John. 1945. *First Draft of a Report on the EDVAC*. University of Pennsylvania: Moore School of Electrical Engineering. Accessed July 21, 2019. http://abelgo.cn/cs101/papers/Neumann.pdf

Web of Stories. 2011. *Marvin Minsky*. Clip 25 of 151. Accessed July 22, 2019: https://www.webofstories.com/playAll/marvin.minsky

Weizenbaum J. 1961. “How to make a computer appear intelligent.” In *Datamation* 7: 24–26.

Appendix A: Early AI Corpus—list of included papers

Feigenbaum, Edward A. and Julian Feldman, Eds. 1963. *Computers and Thought.* New York: McGraw-Hill.

In the above volume:

Armer, P. 1960. "Attitudes toward intelligent machines." In *Symposium on Bionics, Wadd Technical Report.* 60 (600): 13-19

Clarkson, G. P. E. 1961. "A model of the trust investment process." In *A Simulation of Trust Investment*. Englewood Cliffs: Prentice-Hall.

Feigenbaum, E. A. 1961. "The simulation of verbal learning behavior." In *Proceedings of the Western Joint Computer Conference*. 19: 121-132.

Feldman, J. 1961. "Simulation of behavior in the binary choice experiment." In *Proceedings of the Western Joint Computer Conference*. 19: 133-142.

Gelernter, H. 1959. "Realization of a geometry-theorem proving machine." In *Proceedings of an International Conference on Information Processing*. Paris: UNESCO House, 273-282.

Gelernter, H., J. R. Hansen and D. W. Loveland. 1960. "Empirical explorations of the geometry-theorem proving machine." In *Proceedings of the Western Joint Computer Conference*. 19: 143-147.

Green, B. F., A. K. Wolf, C. Chomsky, and K. Laughery. 1961. "Baseball: An automatic question answerer." In *Proceedings of the Western Joint Computer Conference*. 19: 143-147.

Gullahorn, J. and J. Gullahorn. 1963. *A computer model of elementary social behavior*. University of California at Berkley: School of Business Administration.

Hunt, E. B. and C. I. Hovland. 1960. "Programming a model of human concept formulation," In *Proceedings of the Western Joint Computer Conference*. 19: 145-155.

Minsky, M. 1961. "Steps towards artificial intelligence." In *Proceedings of the Institute of Radio Engineers*. 49: 8-30.

Minsky, M. 1961. "A selected descriptor-indexed bibliography to the literature on artificial intelligence." In *IRE Transactions on Human Factors in Electronics*. 2: 39-55

Newell, A. and H. Simon. 1961. "GPS, a program that simulates human thought." In *Lernende Automaten*. Munich: R. Oldenbourg KG.

Newell, A., J. C. Shaw, and H. Simon. 1957. "Empirical explorations with the Logic Theory Machine." In *Proceedings of the Western Joint Computer Conference*. 15: 218- 239.

Newell, A., J. C. Shaw, and H. Simon. 1958. "Chess playing programs and the problem of complexity." In *IBM Journal of Research and Development*. 2: 320-335.

Samuel, A. L. 1959. "Some studies in machine learning using the game of checkers." In *IBM Journal of Research and Development*. 3: 211-229.

Selfridge, O. and U. Neisser. 1960. "Pattern recognition by machine." In *Scientific American*. 203: 60-68.

Slagle, J. 1963. *A heuristic program that solves symbolic integration problems in freshman calculus*. Cambridge: MIT Lincoln Laboratory.

Tonge, F. M. 1960. "Summary of a heuristic line balancing procedure." In *Management Science*. 7: 21-42.

Turing, A. M. 1950. "Computing machinery and intelligence." In *Mind*. 59: 433-460.

Uhr, L. and C. Vossler. 1961. "A pattern-recognition program that generates, evaluates, and adjusts its own operators." In *Proceedings of the Western Joint Computer Conference*. 19: 555-570.

Hook, Sidney. Ed. 1960. *Dimensions of Mind: A Symposium*. New York: New York University Press.

Lighthill, Sir James. 1973. "Artificial Intelligence: A General Survey," from Cambridge University.

McCarthy, J., M. Minksy, Rochester, L., and C. E. Shannon. 1956. "A proposal for the Dartmouth summer research project on artificial intelligence." Accessed August 30, 2019. http://www-formal.stanford.edu/jmc/history/dartmouth.pdf

McCarthy, John. 1959. "Programs With Common Sense" in *Semantic Information Processing*, edited by M. Minsky. Cambridge: The MIT Press. 403-418.

Minsky, Marvin. 1974. "A Framework for Representing Knowledge." In *The Psychology of Computer Vision*, edited by P. Winston. New York: McGraw-Hill.

Newell, Allen and Herbert A. Simon. 1976. "Computer Science as Empirical Inquiry: Symbols and Search." In *Communications of the ACM*. 19 (3): 113-126.

Putnam, Hilary. 1964. "Robots: Machines or Artificially Created Life." in *The Journal of Philosophy*. 61 (21): 668-691.

Rosenblatt, F. 1958. "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." In *Psychological Review*. 65 (6): 386-408.

Selfridge, Oliver. 1959. *Pandemonium: A Paradigm for Learning*. London: H.M.S.O.

Simon, Herbert A. and Allen Newell. 1958. “Heuristic Problem Solving: The Next Advance in Operations Research.” In *Operations Research*. 6 (1): 1-10.

Simon, Herbert A. and Allen Newell. 1970. “Human Problem Solving: The State of the Theory in 1970.” In *American Psychology*. 26 (2): 145-159.