



New Sapience 101

Background context for our revolutionary approach to *Artificial General Intelligence*

June 2018

Welcome

Since you are reading this, you have already had some introduction to our extraordinary claim: New Sapience, a start-up founded in 2015, has achieved a break-through in Artificial Intelligence. We have succeeded in creating a technology that has actual comprehension of human language and the potential to achieve Artificial General Intelligence in just a few years.

We are bringing this technology to market at a time when there is much excitement about the world-changing possibilities of Artificial Intelligence with huge on-going investments by the big tech companies, the Venture Capital community, government and academia. AI is much on the minds of the public as well, every day we see new articles touting both the promise and dangers of machines that can think.

There is, however, a huge gap between what the current “Big Data” and “Deep Learning” techniques, that are based on stochastic or statistical methods can achieve, and machines that can think in general terms and converse with people in our native languages. Experts in the main-stream technical community are aware of this gap and when pressed, admit that they don’t know how to cross it, while much of the media is either oblivious or chooses to ignore it.

Our technology, **Machine Knowledge** closes this gap by putting knowledge derived from human understanding of the world directly into the machine. The result is something the world has never seen before: a digital entity with the core knowledge needed to understand the meaning of natural language words and grammar. Each instance of our software becomes a unique individual as it learns, like a person, through reading and conversation. We call these entities **sapiens**.

Over the following pages we will define with precision many terms that are often used loosely in discussions about AI, providing a solid foundation for our further presentations and demonstrations. Finally, we will show you:

- that mainstream AI today is a “one-trick pony” stuck on an approach that has little hope of achieving Artificial General Intelligence in the foreseeable future
- that there is little worry that the big tech companies are pursuing a solution like ours in some secret lab
- why, when our solution is revealed to be the true one, it will be too late for others to overtake us
- why a company like New Sapience discovered the solution while “big tech” has not

We will demonstrate to you that New Sapience technology holds the key to a future rich with possibilities:

“If you invent a breakthrough in artificial intelligence, so machines can learn, that is worth 10 Microsofts.”

Bill Gates (2004)

AI: “The Holy Grail of Computing.”



Since the dawn of computing, people have been fascinated with the question of whether it is possible to create Artificial Intelligence, machines with intelligence comparable to humans. Over the years, the very question has generated quite a bit of confusion and controversy because there is no general agreement about exactly what intelligence is.

We need an essential definition of the term AI, one that is not dependent on any

one approach to creating it (which may not be the same in an organic brain and a machine).

Intelligence

AI is simply intelligence in a machine, but what is intelligence? It is the defining attribute of our species, the one that distinguishes us from all other animals. This is the key to a functional definition because it is really the fact that humans have power to radically alter our environment that makes us different. That power derives from our knowledge of the world.

So, a simple definition: **Intelligence is a set of capabilities that acquire, create and apply knowledge.** This definition applies equally to humans, machines, and any yet undiscovered intelligent alien species. It also explains why the very idea of AI is so exciting. We build machines to extend our own power. The first machines amplified the power of our muscles. Later, with devices like the telescope, we learned to amplify our senses. More recently, we have found ways to amplify some of our brain’s cognitive abilities, such as the ability to do arithmetic, with computers.

A device that amplifies our greatest power, the ability to create and apply knowledge, would be the **ultimate machine: the Holy Grail of computing.**

Knowledge, as opposed to *data*, which is individual facts or assertions, and *information*, which is organized data, is an internal model or representation of the external world.ⁱ Computers

already excel at data and information processing but, until now, no computer program could process knowledge. New Sapience has created a technology based on a computable model of **common-sense** knowledge.

No other AI approach can remotely even conceive of doing this:

*“The crucial piece of science and technology we don’t have is how we get machines to build **models** of the world.”ⁱⁱ*

*“Right now, even the best AI systems are dumb, in the way that they don’t have **common sense**.”ⁱⁱⁱ*

Yann LeCun, director of AI Research, Facebook.

Characteristics of General Intelligence

Intelligence in humans evolved as a set of cognitive capabilities that culminated in the ability to model the world. New Sapience has reversed this process in machines, by first creating a world model and then implementing the intelligence - the logical reasoning routines and algorithms that allow the system to extend and refine the original model in response to incoming information.

We have been guided in our implementation by two major characteristics that, as is widely accepted, *general* intelligence is expected to have:

1. It should be able to comprehend human language. This is key, because understanding language opens the door to unlimited education through reading and comprehension. Key features of language comprehension include:
 - The ability to deduce the meaning of unknown words in context as they are encountered in reading or conversation.
 - The ability to recognize that an assertion is illogical, contradicts what is already known, or simply doesn’t make sense.
 - The ability to disambiguate homonyms from context.
 - The ability to resolve pronoun and common noun references.
 - The ability to recognize actual instances of things as opposed to classes of things.
2. It should be able to explain its reasoning or conclusions.

These major characteristics are emergent from a larger collection of individual skills that may be categorized in steps of increasing sophistication as shown in Table One below.^{iv}

Capabilities of General Intelligence

KNOWLEDGE	UNDERSTAND	APPLY	ANALYZE	SYNTHESIZE	JUDGEMENT
<ul style="list-style-type: none"> ✓ define ✓ repeat ✓ record list ✓ recall ✓ name ✓ state 	<ul style="list-style-type: none"> ✓ classify ✓ discuss ✓ describe express ✓ identify ✓ locate report review ✓ tell 	<ul style="list-style-type: none"> interpret apply employ use demonstrate practice illustrate operate schedule 	<ul style="list-style-type: none"> distinguish analyze differentiate appraise compare contrast criticize inspect debate inventory question 	<ul style="list-style-type: none"> compose plan propose design formulate arrange collect create set up organize manage prepare 	<ul style="list-style-type: none"> judge appraise evaluate rate value revise score select choose assess estimate measure

Table One

The check marks indicate current capabilities of Machine Knowledge as can be demonstrated *today* in an English conversation with a sapiens. The table was not originally compiled to assess AI. It is in fact used by educators to assess the learning sophistication of *human students*. That it has turned out to be equally useful at assessing Machine Knowledge is unprecedented and extraordinary.

No other AI technique can demonstrate even one of these capabilities or, indeed, any comprehension of natural language at all:

“Ultimately, the real challenge is human language understanding – that still doesn’t exist. We are not even close to it...”^v

Satya Nadella, Microsoft CEO

“I’ll be honest with you, I believe that solving language is equivalent to solving general artificial intelligence. I don’t think one goes without the other.”^{vi}

Emmanuel Mogenet, head of Google Research Europe

Thought

The great mathematician and computer scientist, Alan Turing, proposed his now famous test for Artificial Intelligence in 1950 in response to a simple question: “Is it possible that machines



Thought is a process that produces knowledge.

could **think?**” His test was simple: if the judge could not reliably tell which was a machine and a person in a text conversation between the two, then it would be unreasonable to say the machine was not intelligent.

Thought is the process that results in **knowledge** and anything that can think is **intelligent**.

That includes our sapiens technology, but not any of the other current AI approaches.

Learning (Cognitive versus Neural)

Despite having over one million words, the English language sometimes makes do with a single word when it really should have two. “Learning” is a case in point.

When a human “learns” how to ride a bicycle, they do so by practicing until neural pathways coordinating the interaction of the senses and muscles have been sufficiently established to allow one to stay balanced. This “neural learning” is clearly very different than the kind of “cognitive learning” we do in school, which is based on the acquisition and refinement of knowledge. *Neural learning* cannot be explained and cannot be unlearned; no abstract knowledge of the world is produced. Neural learning is the result of *training*.

“Cognitive learning” is the creation of new knowledge by assembling pre-existing concepts in new ways. Most of the knowledge by which humans build and control their world comes through natural language. Language is information that specifies how to construct a more complex concept from simpler ideas. When the assembled “cognition” corresponds to something in the external world that was not previously modeled, we have cognitive learning. Cognitive learning is the result of *education*. No other AI approach can claim cognitive learning:

*“We’re very far from having machines that can **learn** the most basic things about the world in the way humans and animals can do.”^{vii}*

Yann LeCun, Facebook

Understanding and Comprehension

If you asked a four-year-old “What is a cat?” and were answered “It is an animal,” you could be sure the child understood or comprehended the question and the answer. But consider a computer program of the form:

```
if input = “What is a cat?”  
then output = “It is an animal.”
```

Is it fair to say the program understood the question or answer? No, for the program it is pattern in and pattern out - it has no notion that the words signify internal concepts in the rich world models in human minds. Chatbots are programs of this type, no matter how many patterns they can handle or how convincingly they may parrot human speech. Even when their output is useful to a human, they are devoid of knowledge, possess no generality, and cannot learn.

Narrow AI versus Artificial General Intelligence (AGI)

AI became a recognized branch of computer science over the last century. At first the term was meant in the sense of general intelligence as identified by the Turing test. Gradually the term was loosened to include many technologies that clearly lacked even the potential for general intelligence. It became a term applied to all kinds of non-traditional software. Any software that attempted to model any aspect of human brain processes, including ancillary capabilities such as vision and hearing, or structural elements of the brain such as neurons, were thrown in the “AI” bucket.

This unfortunate development of playing fast and loose with the term “intelligence” has resulted in a great deal of confusion. Eventually, those researchers still in pursuit of the genuine article - endowing computers with knowledge generating thought - needed a new term to distinguish their goal from all those other “AIs.”

AI is powerful. AI is cool. AI attracts funding. Naturally everyone wants to be on the inside of it. So, the term “**strong AI**” - for the real thing - was floated. More acceptable, but that still implied the others were “weak.”

Finally, the terms “**artificial general intelligence**” (AGI) or, alternatively, “**real AI**” became widely accepted. (Real AI is less commonly used because it implies narrow AI is not AI at all; which, of course, it isn’t.) In any case, it is only the people in the AGI camp who routinely make the distinction; the people working in the “narrow AI” sector naturally prefer to call their technology simply “AI”, and so the confusion goes on.

New Sapience has the only technology today consistent with what is now called artificial general intelligence; others admit they are not even close:

“We should not claim that artificial general intelligence is just around the corner”^{viii}

Satya Nadella, Microsoft

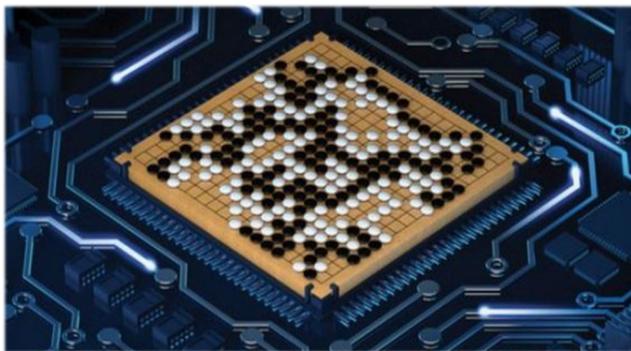
Stochastic versus Deterministic Algorithms

These terms are very significant for our discussion of AI since all approaches are rooted in one or the other technique.

Personal assistants, such as Siri and Artificial Neural Networks (ANNs) and all “big data” systems, use the stochastic approach. The term **stochastic** refers to patterns that may be analyzed statistically but may not be predicted precisely. This makes stochastic algorithms useful for a large class of problems where vast amounts of data are involved and where there may be hidden or incomplete data, common characteristics of many real-world problems.

New Sapience uses a deterministic approach. A **deterministic** algorithm is one which, given an input, will always produce the same output, with the underlying machine always passing through the same sequence of states. Deterministic algorithms are the most familiar kind, as well as one of the most practical, since they can be run on ordinary computers efficiently. Statistical approaches usually require super computers with large data sets to “train” the system.

Big Data and Deep Learning



In 2017 AlphaGo a program created by Google DeepMind beat the human world champion Go player. AlphaGo is now retired. There is nothing else for it to do. There is nothing else it can do.

“deep learning.”

However, being stochastic, ANNs require lots of computing power and large training sets, neither of which were available in sufficient degree for any practical applications at that time. With the advent of supercomputers and massive datasets after the turn of the century, artificial neural networks are finally having their day.

These stochastic techniques excel at certain types of information processing based on pulling patterns or correlations out of large datasets. They are at the foundation of speech recognition and synthesis, facial recognition, and the adaptive behavior of self-driving vehicles, to name a

Articles about AI are published every day. In most of these articles the term “AI” is used in a very narrow sense. It means applications based on training artificial neural networks (ANNs) under the control of sophisticated statistical algorithms to solve very particular problems.

The computational model for neural networks based on mathematics and algorithms called “threshold logic” dates from 1943. But it wasn't until 1986, when Geoff Hinton and colleagues, described a technique called backpropagation, that ANN's become capable of what is now called

few. More and more practical applications are being brought forward every day. Today ANNs are solving many hitherto intractable problems and are generating large revenues and with that much well-deserved excitement. But a key aspect of these ANNs is that they work very much as our muscle training works as described above, and don't generate cognitive knowledge. They work as a result of training.

Stochastic AI and Natural Language

However, one problem domain where great investments are being made in stochastic AI without generating so much excitement is in the processing of natural language. There are good reasons why, despite vast investments by their respective owners, the current crop of “digital personal assistants” are as likely to be made fun of as made use of.

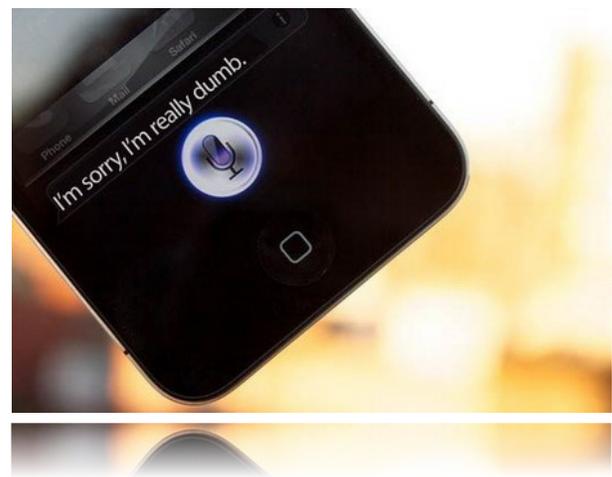
Watson, IBM's question-answering computer system, beat Ken Jennings, the world champion at the game of Jeopardy, in 2011. *Watson* is stochastic AI, it processes the text of questions and answers as patterns to be matched up without any notion of what the words mean.

Watson was against a dataset consisting of 200 million pages of structured and unstructured text content (including the full text of Wikipedia) consuming four terabytes of disk storage.

Although an impressive achievement, IBM shamelessly hypes *Watson* as something it is not with the tagline: “*Watson* – The Power of Knowledge”.

Here, as in many cases of advertising hype in AI, the distinction between genuine comprehension based on knowledge and just outputting what a human recognizes as a correct answer is completely glossed over.

The big tech companies are currently engaged in an expensive arms race to own “conversation as a platform.” They are doing this because they understand that conversational interfaces have the potential to completely alter the landscape with respect to how people interact with their products, and that impacts how these companies generate revenue. For example, where would Google (Alphabet) be if you could ask your digital assistant to “google” things for you and you never see the ads?



Between May 2016 and May 2017, SIRI lost 7.3 million monthly users, or about 15% of its total U.S. user base, [according to data from researcher Verto analytics.](#)

Why are these programs so limited?

The difficulty lies in the much-misunderstood relationship between language and knowledge. Common sense tells us that the Library of Congress is a great repository of knowledge. But consider what is going on when people use language to share knowledge.



*The Library of Congress: A great repository of **Information**.*

Our concepts are composed of simpler concepts. So, to convey a concept from one person to another, we first mentally break it down into its component parts.

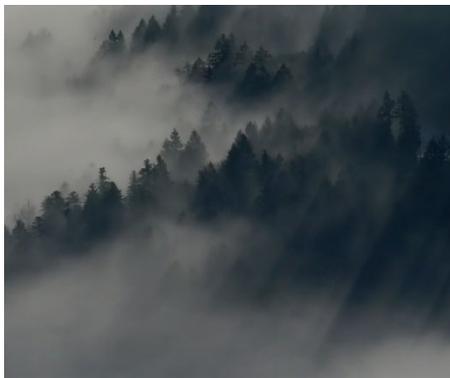
For example, to convey the idea of “cell phone”, we first break it down into component concepts represented here by their English word referents: device, portable, electronic, communication. Then we put those words into a grammatical structure: “A cell phone is a portable, electronic communication device.”

What is conveyed is a parts list which the recipient matches to concepts they already have in their memory. The grammar corresponds to assembly instructions. If all the components are well understood, it may be obvious how they connect to form the new concept in the recipient’s mind regardless of the grammar. However, if one or more parts are missing, full communication is not possible.

The bottom line is that **no stochastic program will ever be able to extract knowledge from text repositories because *the knowledge isn’t there***.

Comprehension of natural language is dependent on pre-existing and commensurate *world models*, independent of language, on both ends of the communication process. New Sapience, and no other approach, puts just such a world model, designed with reference to and compatible with human world models, in the machine, enabling the creation of new knowledge through language comprehension.

Deterministic Approaches to AI



Disappointment over Rule-based systems and Semantic Networks led to a period (1990s) often referred to as the “AI Winter”

Back in the 1980’s, the AI community was focused on deterministic programs that reasoned according to rules and logic, making their inner workings transparent to anyone who cared to examine the code. *Rule-based*, or so-called *expert systems*, and *semantic networks* both generated considerable optimism about AI. But, like artificial neural networks which were just a side note at the time, these approaches also failed to rise above the level of information processing.

While early rule-based applications showed promise, ultimately trying to represent knowledge as rules proved

impractical. Each rule was an individual assertion, that is, an individual datum. Chaining these data together with inference could result in useful information, but there was no way to create a model independent of the logic that reasoned about it. Rule-based applications required progressively more rules to achieve diminishing results; *they did not scale*.

Sematic networks are an information technology inspired by language in which the information is represented as nodes corresponding to natural language words with connections corresponding to phrases. Facts and assertions could be stored by the program and matched to key words in a natural language query. However, there was no appreciation of the difference between the language and knowledge and once again the technology *could not scale* to commercial applications.

Stochastic Approaches to AI: A One Trick Pony



"When the only tool you have is a hammer every problem looks like a nail"

Abraham Maslow

The belief that deterministic approaches cannot be made to scale has become unquestioned within the AI community. That, combined with the success of artificial neural networks in solving problems in narrow domains, has created a monolithic culture. Today, outside of New Sapience, ANNs and AI are considered *one and the same*. The AI curriculum of the universities are filled with courses on statistical algorithms while cognitive science programs focus upon reverse engineering the human brain. Upon graduation (or before) students are snapped up by the big tech companies at enormous salaries. The system re-enforces itself.

This is very good news for New Sapience. The big tech companies are engaged in trying to move neural network approaches closer to AGI an inch at a time at enormous cost. Their approach, by its very nature, requires the creation of something very much like an artificial human brain. Though they have vast resources, they have hitched their star on an approach that is enormously complex and fundamentally obscure. No one knows how the brain works, and, what is worse, the developers don't really even understand how the neural networks they build work:

"When it comes to neural networks, we don't entirely know how they work, and what's amazing is that we're starting to build systems we can't fully understand. The math and the behavior are becoming very complex and my suspicion is that as we create these networks that are ever larger and keep throwing computing power to it, ... (it) creates some interesting methodological problems."^{ix}

Jerome Pesenti, former VP of the Watson team at IBM

“We can build these (statistical) models, but we don’t know how they work.”^x

Joel Dudley, lead, Mount Sinai AI team.

“Deep learning is greedy, brittle, opaque, and shallow.”^{xi}

Gary Marcus, former director, Uber AI Lab

At the end of this paper we list the original references of the quotations presented. These are articles in which the very practitioners of stochastic AI describe what it is they cannot currently do. However, nowhere in any of these articles will you find a hint that, despite the obvious limitations, anyone imagines that there may be a viable approach to AGI that lies outside of their own expertise in statistics and ANNs. They are completely unaware that they are inside a box they cannot think outside of.

Because their financial resources are so vast, one would think that all approaches would be tried, and that somewhere in a basement tucked away at Apple or Google, teams must be working on a deterministic, model-based approach like ours. We have become increasingly confident this is not the case.

The swift grant of our first patent (in only 3 years; impressively fast) provides compelling evidence. The big tech companies patent everything they do but the patent search came up empty. Another telling fact: we are already demonstrating actual natural language comprehension after a relatively miniscule investment compared to the R&D budgets at these companies. If they were pursuing a similar approach it should have borne fruit by now and we would be seeing the results.

It seems very likely that they will continue down the path they are on until, not long from now, someone like Satya Nadella or Tim Cook will have their first conversation with one of our sapiens. They will realize at once that the path to true AGI has been found.

But by that time, New Sapience will have a lead that the others will find difficult to overtake even with vast resources, especially because they would have to staff up new development teams from scratch since different skill sets are required to do what we do. More importantly, we will be armed with a defensive set of patents in one hand and will be extending the offer of a reasonable license opportunity with the other.

Meanwhile we are sailing toward the finish line. While others are narrowly focused on stochastics, we have created the world’s first *deterministic* approach that is inherently *scalable*.

A Contrarian Approach to AGI

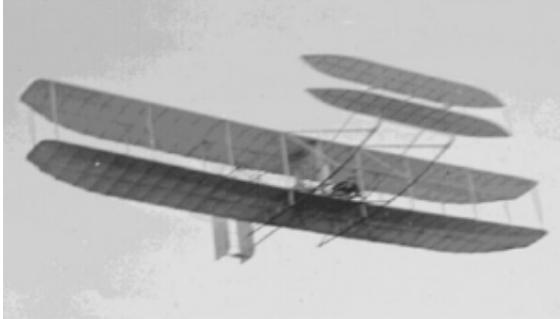
At New Sapience we ask not:

“How can we emulate the cognitive processes of the human brain?”

the question everyone else is asking but, instead:

“Can we put human knowledge of the common-sense world, as we introspect it in our own minds, into a computer?”

As it has turned out, ***we can***.



With the advent of the internal combustion engine, it was only a matter of time before the problem of powered flight was solved.

Many of the inventions that have changed our world happened because it was their time to happen.

The Wright brothers succeeded because they asked themselves the right question, which was not: “How do we build an artificial bird?” but rather: “How can we apply the aerodynamics of bird wings to something that can hold itself together?”

Other inventions, like our Machine Knowledge, come “out of left field,” the result of a series of unforeseeable influences and events, pieces of a puzzle that come together at a certain point in time independent of and often contrary to the technology mainstream.

A contrarian approach is often the result of asking a different question than the mainstream. We are good at that at New Sapience.

Here simply stated, is the New Sapience thesis. It is possible to build a computer model, a software object structure, that can represent core concepts of the common-sense world. Such a model can then serve as the basis for reasoning and language comprehension based on deterministic, logical algorithms.

Our model, comprised of highly interconnected concept-objects reflect the world itself rather than words of any human language, grammar or linguistic construct. The model supports comprehension of any natural language once an appropriate grammar/vocabulary mapping layer has been put in place.

While it provides model referents for a vocabulary list that covers over 80 percent of the words on the Internet, the total number of concepts is manageable. It is roughly equivalent to the world knowledge of a human 5-year-old which corresponds to a vocabulary of about 2500 words.



If the only technology you have is neural networks, your only path to AGI is to build an artificial human brain. An approach that has much in common with those who equated artificial flight with artificial birds.

This core model, the crown jewel of our technology, is being meticulously constructed. Our team has deep roots in traditional AI techniques and understands their capabilities and limitations. Upon that base, we have people with practical expertise in modelling complex real-world systems in software, expertise in modern object-oriented languages and non-SQL databases, linguistics, as well as individuals with broad inter-disciplinary backgrounds spanning mathematics, the humanities, and philosophy, particularly epistemology.

A team of such breadth and depth exists nowhere else in the AI landscape and is the basis of our unique path and extraordinary success.

Our sapiens' view of the world is based on knowledge curated from the best human sources, a model already fully realized when the sapiens first boots and which serves as the foundation for all that it learns afterwards. Its thought processes are formulated using deterministic logic, guaranteed to reliably return the same results from the same inputs. This design permits the sapiens' reasoning process to be open and transparent, a sapiens will be able to explain how and why it reaches any conclusion or takes any action.

This is very good news for humanity. Deterministic design means that a sapiens can only do what it is programmed to do. No matter how far sapiens progress up the scale of intelligence, they will reliably be guided by carefully designed principles of service to their creators, immutably laid down in their core programming.

This happy truth is in stark contrast with stochastic approaches which have none of these built-in safeguards and as has been pointed out, we don't even know how they work. Elon Musk, Steven Hawkins, Bill Gates and others have warned about the potential dangers of AI. (Since, until now, stochastic AI has been the only game in town, people have assumed that it must be the only path to AGI, even if far in the future.) Their concerns are spot on; humanity would be foolish to trust a stochastic AI, if one could ever be built.

What we are creating at New Sapience is AI with a human face. Reliable servants whose world view has been received directly from the hands of their creators. Sapiens of every use and purpose will eventually, when equipped with robot bodies, perform all tasks that humans find dull, dirty or dangerous. Our sapiens will liberate people to follow their individual bliss and concentrate their energies on what uplifts and refines the human spirit.



ⁱ A website with a concise discussion of the difference between data, information and knowledge:

www.infogineering.net/data-information-knowldge.htm

ⁱⁱ “Facebook’s head of AI wants us to stop using the Terminator to talk about AI”

<https://www.theverge.com/2017/10/26/16552056/a-intelligence-terminator-facebook-yann-lecun-interview>

ⁱⁱⁱ “Inside Facebook’s Artificial Intelligence Lab”

<https://www.popsci.com/facebook-ai>

^{iv} The table is in fact **Bloom’s Taxonomy of Learning** used by educators to assess the learning sophistication of human students. https://en.wikipedia.org/wiki/Bloom%27s_taxonomy

^v Microsoft CEO says artificial intelligence is the ‘ultimate breakthrough’

<http://mashable.com/2017/02/20/microsoft-satya-nadella-artificial-intelligence-focus/#T3pkMIHKGiqw>

^{vi} “Why Google can’t tell you if it will be dark by the time you get home – and what it’s doing about it”

<http://www.businessinsider.com/why-google-cant-tell-you-if-it-will-be-dark-by-the-time-you-get-home-2017-6>

^{vii} “Facebook’s head of AI wants us to stop using the Terminator to talk about AI”

<https://www.theverge.com/2017/10/26/16552056/a-intelligence-terminator-facebook-yann-lecun-interview>

^{viii} Microsoft CEO says artificial intelligence is the ‘ultimate breakthrough’

<http://mashable.com/2017/02/20/microsoft-satya-nadella-artificial-intelligence-focus/#T3pkMIHKGiqw>

^{ix} “The Real Trouble with Cognitive Computing” <https://www.nextplatform.com/2015/06/25/the-real-trouble-with-cognitive-computing/>

^x “The Dark Secret at the Heart of AI” <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>

^{xi} “Greedy, Brittle, Opaque, And Shallow: The Downsides to Deep Learning” <https://www.wired.com/story/greedy-brittle-opaque-and-shallow-the-downsides-to-deep-learning/>