

SPRING DAWNS ON ARTIFICIAL INTELLIGENCE

Michael K. Bergman¹, Coralville, Iowa USA

June 2, 2014

AI3::Adaptive Information blog

When I inaugurated this **AI3** blog in 2005 I made this statement in the [about section](#) to clarify that the “three AIs” stood for *adaptive information*, *adaptive innovation*, and *adaptive infrastructure*, and not the AI of [artificial intelligence](#):

. . . I personally believe artificial intelligence to be a lot of hoovey and hype at best, and a misnomer and misdirection at worst. . . ‘Artificial intelligence’ is a misdirection of attention and energy.

Gulp. OK. Time to take my medicine.

I am today formally retracting those statements — probably should have done so some time ago — and want to explain why. As much as anything, it has to do with the changing understanding of what is artificial intelligence, recently affirmed by global-scale applications and technologies, working effectively right now.

Many Winters within AI

Though the idea of automatons and intelligent agents standing in for humans is about as [old as human storytelling](#), the real basic ideas around artificial intelligence became current as part of the World War II effort and were finally given a name in a famous [1956 conference at Dartmouth](#). Initially namers and advocates of artificial intelligence included such founders as [John McCarthy](#), [Herbert Simon](#), [Claude Shannon](#) and [Marvin Minsky](#). Money to support early interest in artificial intelligence came from the part of the US military that eventually became [ARPA](#) (now [DARPA](#)), with the funding going to individual researchers to use as they wished as opposed to specific projects. Along with many futuristic visions of the 1950s to 1970s, the promises for artificial intelligence were bold, including being able to capture and automate most notable basic human capabilities.

Popular movies and books promoted the ideas of autonomous robots that we could speak with and command and that would anticipate our needs and wishes so as to act as [simulacrum](#) agents lessening our burdens and adding to our leisure and capabilities [\[1\]](#). Algorithms would be discovered and codified that would mimic the basis of human thought and intelligence. The idea of the [Turing machine](#) established a defensible basis for foreseeing that any problem of mathematical logic could be captured and taken on by computers.

The predictable failure of this vision to deliver caused a backlash, sufficient that the US Congress prohibited further open-ended funding via the [Mansfield Amendments](#) in 1969 and 1973, such that by 1974 AI funding in the US had largely dried up. Similar restrictions were applied to the British research community. The result of this backlash caused the first of what would prove to be many “winters” of funding and acceptance for AI.

Roughly a decade later, in response to the perceived Japanese threat for [“fifth-generation” computing](#) in the mid-1980s, a number of AI programs were again funded. While hardware developments were proceeding apace, efforts around McCarthy’s AI-oriented language [Lisp](#) and common sense logic frameworks (what are now called

¹Email: mike@mkbergman.com

[ontologies](#) or [knowledge graphs](#)) such as [Cyc](#) began to receive sponsorship again. The mid-1990s were the time of “[expert systems](#),” to be populated by [knowledge engineers](#) charged with interviewing internal subject matter experts (SMEs) to codify their knowledge for later reuse. These efforts, too, disappointed in terms of the lack of practical benefits delivered. More AI winters ensued.

AI (“artificial intelligence”) came to again lose its credibility. Some researchers moved into specific algorithmic disciplines — [Bayesian statistics](#) and [neural networks](#) predominant — while others shifted into such areas as “[hyperlinks](#)” and what became the [semantic Web](#). Today, one could argue, that the lost mojo of AI has affected those in the semantic Web in almost a dialectic way. First, there are those who embrace the idea of [intelligent agents](#) and global knowledge structures, more-or-less in keeping with some sort of vision of artificial intelligence. Second, there are those that have seen the failures of the past, do not want to repeat them, and are more inclined to support “loosely bounded” structure focused on bottoms-up assertions. [OWL](#) modelers and ontologists tend to occupy the first camp; [linked data](#) advocates more the second camp.

The natural community for [knowledge representation](#) and [management](#) has thus tended to bifurcate a bit: global, “visionary” AI types, with history to overcome and challenged by the sheer scale of what emerged from the Internet; and incrementalists, happy to accept a bit of [RDF structured data](#) in the hopes of an ongoing evolution to more structure and interoperability.

Ten years ago, when I made the conscious decision to reject the AI of artificial intelligence as a label for this blog, an algorithmic-vision of AI seemed “wrong” and not in keeping with the general trends of the Web. That was the basis and justification for my then-statements on AI. But a funny thing happened on the way to a cogent forecast: a massive disruption called the Internet came about that — while it took a decade to gestate — changed the whole underlying substrate over which AI could take place. Like so much of history, innovation had presented to us an entirely different reality upon which to “understand” and develop artificial intelligence. It is those changes — plus the fruits from them — that is defining AI in a new light.

Eight AI Megatrends

There are, by my reckoning, at least eight major trends that have been improving AI’s prospects, especially over the past decade (Numbers #3 to #7 below are quite related to AI, the other three are general trends.) Some of the proven wonders we now see in use such as [speech recognition](#), [speech synthesis](#), [language translations](#), [entity recognition](#), [image](#) and [facial recognition](#), [computer vision](#), [question answering](#), [autocompletion](#) and [spell correction](#), [recommendation systems](#), [sentiment analysis](#), [information extraction](#), [document categorization](#), [natural language processing](#), [machine learning](#), [reasoning](#), [optical character recognition](#), [word sense disambiguation](#), [search and information retrieval](#), and [text generation](#) and [summarization](#), with their many additional categories and sub-categories, are proof these trends are making a difference. None individually constitutes what may be called “AI”, but, in combination, they show compellingly that much of AI’s initial vision is indeed being fulfilled to some degree and in some specific aspect today.

Nearly all of these applications correspond to the [Grand Challenges](#) for symbolic computing identified in the 1980s. Until a decade ago, very few of them save search and initial NLP were producing results with sufficient quality and accuracy. Now, all are.

In the past ten years, most evident in the past five, tremendous breakthroughs have occurred across the entire spectrum of artificial intelligence applications. We can point to at least the eight following megatrends enabling these breakthroughs.

#1 Computer Power

A constant river of innovation has fueled the logarithmic power improvements in computers since the first transistor. [Moore’s law](#) has led to massive improvements in hardware cost, numbers of computation cycles, and

amounts of bits stored. Networking capabilities are now truly global and numbers of interconnected devices exceeds billions. Computer software innovations lead to faster and better procedures and methods; as a category, software innovation likely exceeds hardware improvements as a source of computing productivity. What today fits in the palm of our hand thirty years ago required entire rooms, and did not do one billionth of what can be done today.

The rich savanna of computing has itself encouraged a bloom of innovations, many of which contribute to artificial intelligence prospects.

#2 The Internet (and Web)

Though clearly a related function to the general improvements in computing and hardware, the advent of the Internet and its more relevant offspring of the Web has had, I believe, the most fundamental impact on the change in prospects for artificial intelligence. The sheer scale of the Web network has made available crowdsourced innovations like Wikipedia and other crowdsourced data and [knowledge bases](#). More broadly, global content across the entire Web, accessible via a common HTTP protocol, multiplied every individual's access to information — pay close attention — by a factor of a billion or more.

Because the entire Web is interconnected, the sheer raw grist of connected data available to analyze such things as relatedness or similarity is gamechanging. Manual constructs and derived relations from years past can now be multiplied and magnified at Web scale. Any relationship test or validation can be accomplished nearly instantaneously and at (essentially) zero cost. Phenomenal!

#3 Expectations

The discrediting of AI and its holdover smell has itself been a factor working in its favor. By being discredited, it has been possible for multiple possible AI components, many listed herein, to be developed and attended to in relative isolation. Each of today's current pieceparts to AI could be focused upon on their own, without taint from the broader "AI" brush. Because the constituents were recognizable and justifiable on their own, they did not need to fulfill the past overblown visions and expectations for "AI" writ large. The pieceparts could develop in peace.

This observation, if true, means that grand visions like "artificial intelligence" are perhaps rarely (ever?) the result of a grand top-down plan. Rather, like a good stew, it is individual components that need to mature and become available to create the final meal. Since these ingredients need to stand or contribute on their own for their own purposes, the actual resulting stew may vary as to its ultimate ingredients. If one ingredient is not ripe or available, we vary our recipe according to what is available. There is no one single recipe leading to a tasty stew.

Put another way, AI has been flying under the radar for at least the last ten to fifteen years. Portions of the older AI agenda have benefited from specific attention. Better still, the new emergence of the idea of artificial intelligence is also more toned down and practical. Artificial intelligence is now, I believe, understood to be part of a process and not some autonomous embodiment. Human interaction and communication are themselves imprecise and subject to error. Why should not be artificial means to boost those same human capabilities?.

From the standpoint of expectations, artificial intelligence has evolved from science fiction to essentially zero awareness, meanwhile delivering, on a broad scale, focused wonder capabilities such as (nearly) instantaneous translations across 60 leading human languages.

#4 Global Knowledge Bases

How can a system promise useful suggestions or alternatives if it is bereft of information?

At the local or personal level we well understand that we need to describe ourselves via attributes, the more the merrier in terms of a more complete description. A pretty good record for me would include such things as physical description, image, work and economic description, family and life description, education description, text narratives from fun to historical, etc. The more complete description of me requires many sources and many attributes and many perspectives. But, of course, I do not live alone in the world. To describe my world, which constantly changes, I need to describe other thousands of entities I encounter daily. Each of these, too, has many attributes and relationships to other entities. Each of these entities also changes over time (has histories) and place. So, context becomes another critical dimension.

The growth of the Web at scale has resulted in some tremendous knowledge bases of entities and concepts. [Freebase](#) and [Wikipedia](#) are two of the best known, but virtually every domain has its own sources and richness. These knowledge bases, in turn, are often open for use by others. Text mining and digital data mean these data can be combined and made to interoperate. That process is only just beginning.

Though early efforts in artificial intelligence understood that capturing and modeling common sense was both an essential and surprisingly difficult task — the impetus, for example, behind the thirty-year effort of the Cyc knowledge base — what is new in today’s circumstance is how these massive knowledge bases can inform and guide symbolic computing. The literally thousands of research papers regarding use of Wikipedia data alone [\[2\]](#) shows how these massive knowledge bases are providing base knowledge around which AI algorithms can work.

The abiding impression is that the availability of these data sources has fundamentally changed how AI is done. Unlike the early years of mostly algorithms and rules, AI has now evolved to explicitly embrace Web-scale content and data and the statistics that may be derived from global corpora.

#5 Deep Learning

Machine learning is a core AI concept used to determine discriminative characteristics or patterns within source input data. It has been a constant emphasis of AI since the beginning.

Various machine learning algorithms — such as [Markov chains](#), [neural networks](#), [conditional random fields](#), [Bayesian statistics](#), and [many other options](#) — can be characterized among many dimensions. Some are supervised, meaning they need to be trained against a standard corpus in order to estimate parameters; others require little or no training, but may be less accurate as a result. Some are statistically based; others are based on pattern matching of various forms.

A more recent trend has been to combine multiple techniques in what is known as [deep learning](#), where the problem set is modeled as a layered hierarchy of distributed representations, with each layer using (often) neural network techniques for unsupervised learning, followed by supervised feedback (often termed “back-propagation”) to fine-tune parameters. While computationally slower than other techniques, this approach has the advantage of automating the supervised learning phase and is proving generally most effective across a range of AI applications.

More fundamentally, there is a virtuous circle of feedback occurring between AI machine learning algorithms and reference knowledge and statistical bases (see next). This can extend the accuracy, completeness and efficiency of supervised methods. Some notable academic departments have relied on Web-scale corpora (University of Washington and Carnegie Mellon University are two prominent examples in the US). The most dominant player in this realm, however, has been Google (though all of the major search engine and social networking companies have smaller initiatives of similar character).

#6 Big Statistical Data

Using both statistical techniques and results from machine learning, massive datasets of entities, relationships and facts are being extracted from the Web. Some of these efforts, such as the academic [NELL](#) (CMU) or [KnowItAll](#) or [Open IE](#) (UWash) involve extractions from the open Web. Others, such as the terabyte (TB) n-gram listings from Google, are derived from Web-scale [pages](#) or [Google books](#). These examples are but a sampling of various datasets and corpora available.

These various statistical datasets may be used directly for research on their own, or may contribute to further bootstrapping of still further-refined AI techniques. Similar datasets are aiding advertising placements, search term disambiguation and machine (language) translation. In some cases, while the full datasets may not be available, open APIs may be available for areas such as [entity identification](#) or [tabular data](#).

What is important about these trends is that data, statistics and algorithms are all now being combined in various ways with the aim of achieving acceptable AI-backed results at Web scale. It is really via the combination of these techniques that we are seeing the most impressive AI results.

#7 Big Structure

A more nascent area, really in just its first stages of effectiveness, is the application of “big structure” to all of this information. By “big structure” I mean the application of domain and knowledge graphs to help arrange and place the concepts and entities at hand.

At Web scale, the early [Yahoo! directory](#) and [Open Directory](#) were the first examples of structuring domains. Wikipedia next became the most widely used category structure; Freebase, for example, used Wikipedia to initially bootstrap its own structure. A portion of Freebase is now what is used for Google’s own [Knowledge Graph](#). [DBpedia](#) also created its own ontology out of the infobox structure of Wikipedia. The major search engines have also put forward the [schema.org](#) structure as a means of (mostly) organizing entity and attribute information and structured data. [schema.org](#) putatively is an input to the Google Knowledge Graph, but the exact mechanism and ability to trace the results is pretty opaque.

The need for big structure is rapidly emerging as one of the key challenges for Web-scale AI. The Web and crowdsourcing appears well suited to being able to generate entity and attribute data. What remains unclear is how this information can be coherently organized at the scale of the Web. This problem is becoming acute, because the success of “big data” on the Web needs to ultimately find an organized, coherent expression in the aggregate. This is one major AI challenge that remains distinctly unsolved, though promising first steps exist.

#8 Open Source and Content

The major theme of these AI breakthroughs comes from leveraging the global content of the Web. And this enabler, in turn, has been critically dependent on the open source nature of AI algorithms, software code and code infrastructure and architecture, and open content and (generally) open APIs. Open code, algorithms, datasets and knowledge have expanded the pool of *human* intelligence that can be brought to bear on the question of *artificial* intelligence. The positive feedbacks greased through open channels of information, code and data have been absolutely essential to the amazing AI progress of the past few years.

To be sure, open does not mean a level playing field. (See discussion on Google, next.) But, without open source and open content and data, I think no one could argue that progress would have been anywhere near as rapid as it has been. The synergy arising from open source and content has thus been another essential factor in the recent and rapid progress in AI.

The Race to Intelligence

Since innovation is the source of wealth creation, it is also no surprise that the megatrends surrounding AI have also drawn significant investment interest. This interest is in the form of a race to acquire the most innovative AI startups and human expertise (capital) in AI. Since Google has been my common touchstone in this piece — and because Google is the biggest gorilla in the room — we can use them to illustrate the scope and pace of this race. (Though Amazon, Facebook, Microsoft and IBM are also clearly entrants in this race.)

A number of recent articles, notably ones in the [Washington Post](#) and [The Economist](#), have highlighted the total dollars at stake in this AI race. Over the past few years, there have been perhaps more than \$20 billion in AI-related company acquisitions, with [Nest Technologies](#) (Google, \$3.2 B), [Kiva Systems](#) (Amazon, \$775 M), and [DeepMind](#) (Google, \$660 M) some of the largest.

Within Google alone, there has been a buying spree in search improvements (~ \$1.4 B total), robotics (\$80 M), machine synthesis and recognition (\$250 M), machine learning (\$700 M), smart devices (\$3.6 B), compression technologies (\$200 M), natural language processing (\$80 M), and a smattering of others (\$50 M), not to mention its internal efforts in [self-driving cars](#). I don't monitor Google on a constant basis and likely missed some major and relevant acquisitions, but it does appear that Google has perhaps spent over \$6 billion over the past five years or so for AI-related acquisitions [3].

As important as start-up acquisitions has been Google's commitment to hire and partner with many of the leading AI researchers in the world. Besides the strong partnerships Google maintains with such institutions such as the University of Washington, Carnegie Mellon University, MIT, Stanford, UC Berkeley and others, it has also staffed its research ranks with prominent names from those institutions and others.

[Peter Norvig](#), one of the early advocates for combining algorithmic and statistical AI, joined Google in 2001 and is now its Director of Research. Most recently and notably, [Ray Kurzweil](#) joined Google as Director of Engineering in 2012. Other notable AI researchers at Google include [Alon Halevy](#) (FusionTables), [Ramanathan Guha](#) (schema.org), [Geoffrey Hinton](#) (deep learning), [Evgeniy Gabrilovich](#) (search and machine learning), and many others for whom I am not as familiar with their research. There is probably more AI talent combined at Google than has ever been assembled in one institution before.

With [IBM's Watson](#) getting its own division and Facebook funding an AI center to the tune of \$10 B, plus Apple making a similar commitment to robotic manufacturing, it is clear that all of the major players in the computing space are making big bets on AI moving into the future.

AI is Itself But One Beneficiary of These Trends

Since the early winters in artificial intelligence, a phenomenon has developed called the "[AI effect](#)". It really has meant two different things.

First, AI researchers have tended to call their research anything but artificial intelligence. One of the broader and trendy substitutes is known as [cognitive computing](#). Many of the domains and disciplines I noted above got their names and prominent use as substitutes for what used to be labeled as AI. In any case, we can see that AI indeed is a big tent with many components and thrusts.

Second, the "AI effect" also refers to the fact that once an AI technique is embedded in some everyday use, it is no longer perceived as something AI and is taken as a given. [Douglas Hofstadter](#) expressed the AI effect concisely by quoting [Tesler's Theorem](#): "AI is whatever hasn't been done yet."

I was perhaps right to initially reject the algorithmic-centric view of AI from the early years. But now, when matched with big data, big statistics and big structure, all embedded into phenomenal advances in computing power, it is also clear that we are dawning into a new age of AI. One only needs to look at the wondrous progress

on many of what had seemed to be impossible Grand Challenges over the past five years to gain an appreciation of the pace and breadth of new developments to come.

These developments will reify and foster similar emphases in semantic technologies, graph structures and analysis, and [functional programming](#) and [homoiconicity](#) (“data as code, code as data”) that my colleague, Fred Giasson, is [now actively exploring](#). We will find that representational paradigms and the basis of how our tools and algorithms work will increasingly align. There appear to be natural underpinnings to these phenomena, including the pivot of language and meaning, that are closely aligned with the thoughts and writings of that great American pragmatist and logician, [Charles S. Peirce](#). We will increasingly come to see that the wondrous innovations of self-driving cars, talking smartphones, warehouses of fulfillment robots, and computer vision systems can trace their roots back to basic truths of how to see and understand our world.

Understanding these forces will, themselves, help to formulate guidelines and ideas that can foster further innovation. So, in the end, while I still don’t like the term of “artificial” intelligence, it is merely a sign or a term. Adaptive innovations expressed by machines are simply part of the intelligence and structure embodied in the universe, for which we are now gaining the tools and understanding to exploit.

Acknowledgements

This article was originally posted on the *AI3::Adaptive Information* Web site at <http://www.mkbergman.com/1731/spring-dawns-on-artificial-intelligence/>. This version has been edited and reformatted slightly for PDF distribution. We thank Cognonto Corporation for making this content freely available.

-
- [1] [Douglas Adams’ Hyperland](#) is a great exposition on this vision, with my [2007 blog post pointing to the online video](#).
 - [2] Wikipedia maintains its [own page of research](#) that relies on Wikipedia; I have earlier captured about 250 selected sources called [SWEETpedia](#) that relate specifically to semantic technologies and AI.
 - [3] These are merely estimates, and likely quite wrong in many specifics. The estimates were compiled by reviewing a [listing of Google acquisitions](#) (since 2009), supplemented by individual company searches when the acquisition amounts were not listed, followed by analysis of [Google’s SEC Edgar filings](#) in a manner similar to [this analysis](#) (which was also used for the robotics estimate).