MAKING THE WORLD WORK
# Intelligence
VOLUME 33, NO. 1 • MAY 2016

## The Next Version of INTELLIGENCE

### MACHINE INTELLIGENCE OBSERVATIONS, INSIGHT AND ANALYSIS - THE TUMULTUOUS AI NEWS CYCLE

**MARS VIA THE HUBBLE**
**Mars as it was observed shortly before opposition in May 2016 by the Hubble Space Telescope. Some prominent features are clearly visible, including the heavily eroded Arabia Terra in the center of the image and the small southern polar cap. Credit: NASA, ESA, the Hubble Heritage Team (STScI/AURA), J. Bell (ASU), and M. Wolff (Space Science Institute)**

The news cycle has exploded in its coverage of machine intelligence, AI (artificial intelligence), NNs (neural networking computing architectures and approaches), DLNNs (deep learning NNS) and their applications and instantiations. I don't read everything that appears, but, just this month of May, I collected 256 articles, papers and other materials about intelligence.

For more than the past three decades, this newsletter has focused on **reporting** about these developments and trends. Now, starting with this issue, I'm changing the format to my observations, insights and analysis of these important fields.

In the past, I've published historically, at the end of each month about the events that, generally, took place during that month. I also wrote my reporting in the usual third person prose style. Now, I'll be writing each month in the first person, focusing on sharing my views on breaking news and trends and adding my longtime experience and knowledge from covering the machine intelligence domain since 1984. This was when neural nets weren't yet called NNs, but adaptive, self-organizing, learning systems, content-addressable memories, and other long ago tag lines.

Now, when most people throughout the planet have or have access to near-supercomputer computing capabilities, machine intelligence is finally

scaling in its applications, and improving important functions, like face recognition (Facebook), speech recognition (Google) and, for example, algorithmic instantiations of economic policy. *BusinessWeek* reported that "Britain's central bank has been developing computer algorithms for forecasting economic conditions and helping determine interest rate policy."

This was not possible until three major forces converged in offering companies the ability to put learning systems into action. First was Moore's Law, the now half-century old prediction that computer chips would grow in complexity and performance and double in the number of features every other year. Thus, supercomputers in our pockets as well as massive large-scale high performance supercomputers, made of a multitude of CPUs (central processing units (chips) and GPUs (graphic chips), yielding extraordinary computing power and dexterity.

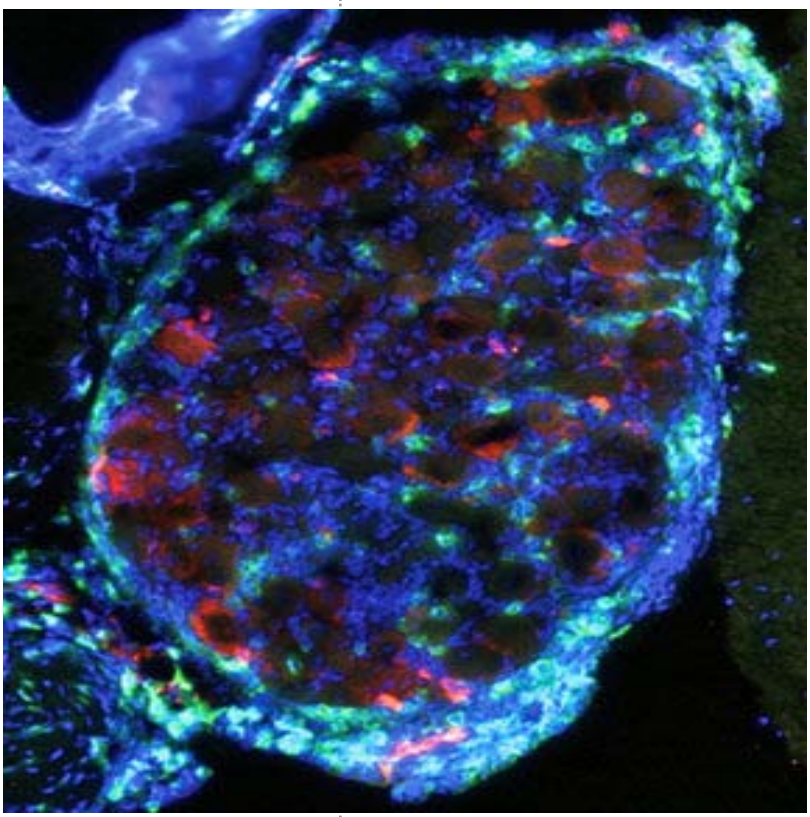

**MOUSE ROOT GANGLIA**
**A cross-section of mouse dorsal root ganglia stained for memory CD4 T cells (green) and the vascular cell adhesion molecule, VCAM-1 (red) after six days of genital herpes challenge. Norifumi Iijima**

Second was the rise of big data, collections of enormous magnitude covering the intense details of how we all live and what we are doing. Again, companies like Google and Facebook had these collections at their fingertips, from search data to social media posts.

Third was the application of machine intelligence learning algorithms by supercomputers on all that data. The explosion took off as these learning algorithms worked at larger and larger scales. When the U of Toronto's Geoff Hinton and Google announced that a deep learning system had learned to recognize what a "cat" is, not by being told what to look for, or what "cat-ness" is, but just from looking at millions of images, only some of which were cats.

These new deep learning NNs changed the landscape of computing, intelligence and machine-augmented capabilities. Earlier, smaller NN systems could solve some problems readily, like check reading and fraud detection. Years ago, that's why Fair Isaac bought HNC, because the latter had NNs that even in their more primitive states, could predict who would default on loans, guessing correctly more often than humans looking at the same data and materials.

With the newer deep learning NNs, faster more powerful computers and big data troves, machine intelligence went from smart to brilliant, outpacing human performance in scores of domains. The DLNNs take in and massage those

myriad data points, and with the help, especially, of graphic processing chips (GPUs) make connections in order to recognize and present salient features, like, what's being presented, what's being shown and what's being said.

Still, NNs in general and DLNNs specifically, are black boxes. How the networks create, establish and maintain their many connections can be described and diagrammed, but how they settle on the correct answers is, for the most part, still mostly unknown and, so far, essentially, unknowable. We know they're working because we can see, hear and experience the results they achieve. But, exactly how NNs do what they do, remains a "deep" mystery.

**IBM QUANTUM**
**IBM makes quantum computing available on IBM Cloud to accelerate innovation. Users of the IBM tutorial can run experiments on an IBM quantum processor.**

## Machine Intelligence - Friend? Foe?

### IF THEY CAN DO EVERYTHING BETTER THAN HUMANS WILL THEY TAKE OVER? - LOST JOBS

John Markoff, of *The New York Times*, reporting on a White House conference at the U of WA, on AI this month: "Never mind Terminator-like killer robots. Artificial intelligence researchers are grappling with more realistic questions like whether their creations will take too many jobs from humans." (http://www.nytimes.com/2016/05/26/technology/artificial-intelligence-is-far-from-matching-humans-panel-says.html)

AI is just the latest aspect in a centuries long trend that has seen machines replace the labor of human beings. Earlier versions saw animals and humans lose out to machines. Today, as robots and algorithms have became further developed and more sophisticated, many jobs done by humans are being done by machines.

Initially, machines replaced humans in, e.g., assembly-line, blue collar work. Then, not so intelligent machines replaced people like bank tellers and clerks. The report cited above is representative of the trickle of coverage following the job displacements to come. Many of the jobs about to be done by machine are in the white collar worker domain, not just the replacement of hands and bodily efforts but machines taking away brain work, work usually considered to be the exclusive domain of human beings.

But, instead of paying attention to the near term terror posed by AI, the loss of jobs and the social disruptions that will certainly accompany such a radical shift in what humans do, what experts now fear are AI systems that not only replace

human workers but make them subservient, Terminators taking over the world. I think this misses the point of a tsunami of job losses that is heading toward all of us in the global community.

Just several years ago, experts in the employment and robotics fields assured us that as many jobs as were being taken away from humans by machines, many more jobs were being created by new opportunities that were emerging as the employment arena is reconfigured. I no longer believe that will be a long term trend. Yes, other jobs will be created but not as many as are being taken away by machines and machine intelligence.

Everyone always told us that if we wanted to succeed in the new computer internet-connected job space, one sure-fire way to get ahead was to learn to program computers. If you can't beat them then you could tell them what to do by learning to code. The cover story on next month's *Wired* tell a different story. Their headline: "The End of Code: Soon we won't program computers. We'll train them. Like dolphins. Or dogs. Or humans." (http://www.wired.com/2016/05/the-end-of-code/)

**ELECTRIC BRAIN**
**Taken from the video of a Glass Brain flythrough. Neuroscape Lab/** *YouTube*

In the article, by Jason Tanz, I found a concise preview of the real terrors of AI that will be coming at us first, before any AI takes over the world, Skynet-style. What will come soon are the massive job losses. Many experts put the number of human jobs to be lost in the coming decade or two at above 50% of all jobs; some more pessimistic predictions put that number above 80%. But, despite the enormous potential for such job losses to machines changing the social fabric, these coming displacements are certainly not mentioned at all in this year's US presidential campaigns.

Tanz quotes Andy Rubin, developer of the Android operating systems for smart phones. He's now a venture capitalist; his company, Playground Global, invests in machine learning start-up companies. Rubin: "People don't linearly write the [NN] programs. After a neural network learns how to do speech recognition, a programmer can't go in and look at it and see how that happened. It's just like your brain. You can't cut your head off and see what you're thinking."

Tanz goes on to quote Demis Hassabis, the leader Google's DeepMind, and the force behind the DLNN system that beat the world's best player at the game of Go a decade before anyone thought that would be possible. Hassabis: "It's almost

like an art form to get the best out of these [DLNN] systems. There's only a few hundred people in the world that can do that really well." Tanz: "But even that tiny number has been enough to transform the tech industry in just a couple of years."

This is an absolute about face from when I started reporting on AI in the mid-1980s. Then it was a world of machine intelligence predicated on logic and the codification of the knowledge of experts into if-then rules that a machine would emulate. The famous Campbell soup cooker was interviewed to discern the right ingredients, temperature, cooking time and other elements that went into cooking the company's soups.

But rules are brittle and, once set in programming code, hard to change and debug. The so-called "expert systems" that resulted from this knowledge instantiation were able to be applied to certain kinds of work, especially in areas of repetitive, unchanging tasks and environments. But, soon it became evident that not being able to adapt and learn was a tremendous drawback for early AI systems. And, in the 20th century, computers were not powerful enough, not fast enough and certainly not smart enough to do most jobs that required flexibility and learning to succeed.

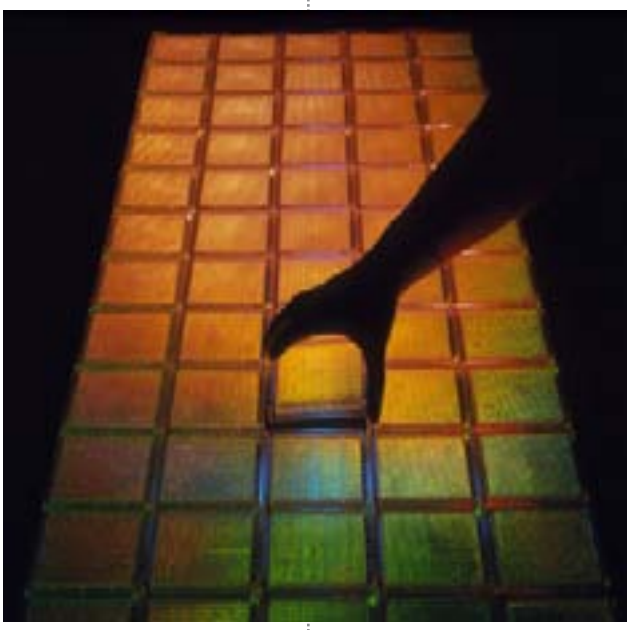

**HUMAN GENOME Sixty trays could contain the entire human genome as 23,040 different fragments of cloned DNA. James King-Holmes,** *Science Source*

AI seers long believed that the approach of logic and rules would somehow work out, but it didn't. NNs, in their early days, could only make limited progress, despite their abilities to learn and adapt. Now, all that's changing. Tanz: "For much of computing history, we have taken an inside-out view of how machines work. First we write the code, then the machine expresses it. This worldview implied plasticity, but it also suggested a kind of rules-based determinism, a sense that things are the product of their underlying instructions. Machine learning suggests the opposite, an outside-in view in which code doesn't just determine behavior, behavior also determines code. Machines are products of the world."

Machines will write their own code:

## Intelligent Assistants

### VIV DEMONSTRATES DYNAMIC PROGRAM GENERATION

The best thing I saw this month was the demonstration of the prototype Viv intelligent assistant by one of the Siri creators, Dag Kittlaus (https://www.youtube.com/watch?v=Rblb3sptgpQ), showing Viv's innard's ability to generate

dynamically program code on the fly, based on a spoken query or request. It was only a demo, but it showed me the potential for what is now the hottest arena in machine intelligence, the intelligent assistant.

Speaking into an Apple iPhone, Kittlaus asked Viv (Latin for "*life*") to come up with complicated sets of information. One of the more complex queries was when he asked Viv: "Will it be warmer than 70-degrees near the Golden Gate Bridge after 5 PM the day after tomorrow?" After he received Viv's answer, Kittlaus showed the diagrammatic sketch of the learning tree arrangement of how Viv arrived at the answer. He underscored that Viv creates the code to isolate and identify the natural language content and then more code to match the meaning of the query to the information sources that provide the answer.

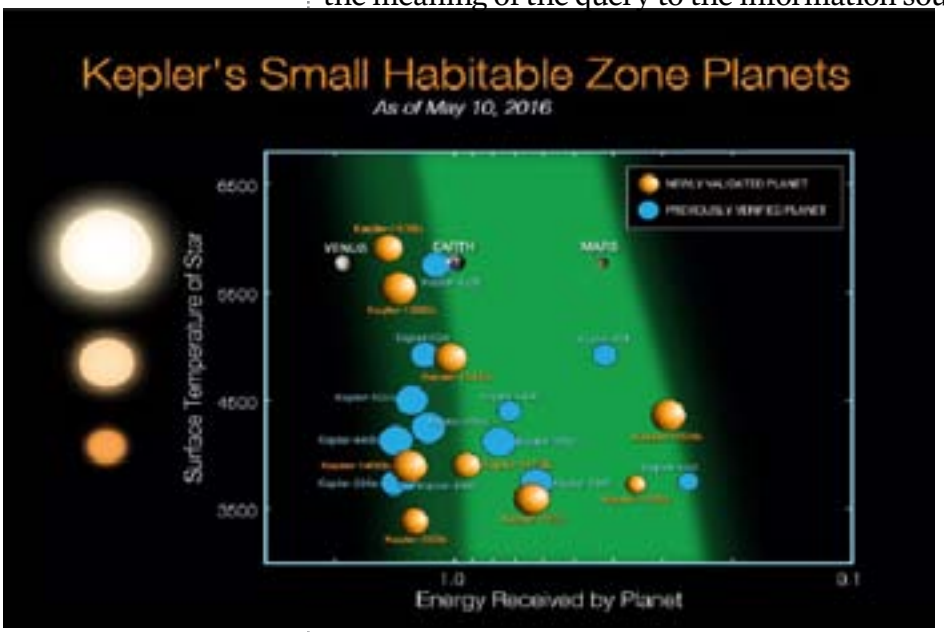


**NINE NEW PLANETS? Nine new potentially habitable planets are among the 1,284 newly confirmed exoplanets found by NASA's Kepler Space Telescope. Shown in orange, the new additions join a growing list of planets in the habitable zones of their stars, where conditions may be right for life. NASA Ames: N. Batalha and W. Stenzel**

In the Golden Gate Bridge example, Viv had to know what the Bridge was, where it was and then look at a future weather forecast for that place at that time. If this kind of disambiguation and data sourcing can scale, then, in a year or so we might all be talking to our phones and other computers in complete sentences and receiving complicated and informative answers.

Kittlaus noted that voice recognition with natural language speech and understanding are the most natural ways for humans to interact with their machines. When he spun Siri out of SRI (Stanford Research Institute) and sold it to Apple, Siri worked in Apple's walled garden of specific machines and software. Viv is intended for many systems, not just Apple or some other specific operating system and computing environment. He sees Viv as being device and domain independent when Viv SDKs (software development kits) appear, perhaps late this year.

I read with interest in *The Washington Post* and elsewhere that Kittlaus had already turned down acquisition offers from both Google and Facebook. I assume his limiting experience with Siri at Apple convinced him that it would be important to aim for ubiquity of operation rather than a specific operating environment.

*TechCrunch*, at whose TechCrunch Disrupt NY meeting Kittlaus made the Viv demo, reported: "Viv has grown relatively stealthily over the past four years. The company raised $12.5 million from Iconiq Capital last year, a company that

*Forbes* reported is backed by names like Facebook's Mark Zuckerberg, Dustin Moskovitz and Sheryl Sandberg, Twitter's Jack Dorsey, LinkedIn's Reid Hoffman and others."

In the arena of using computers to the limits of their potential we've gone from operating systems, like DOS, Windows iOS, Android and more, to apps, the application environment created first for Apple iPhone. More recently came chatbots, which, like apps are for extremely limited domains of operation and content.

Now the intelligent assistant market is commencing, with the promise of even more, smarter, faster, and more powerful operations at our fingertips in the years to come. If Viv or some other entrant can pull off the creation of an assistant that will, over time, work for more and more machines, more and more operating systems and more and more domains, we'll be in a new phase of the future of computing.

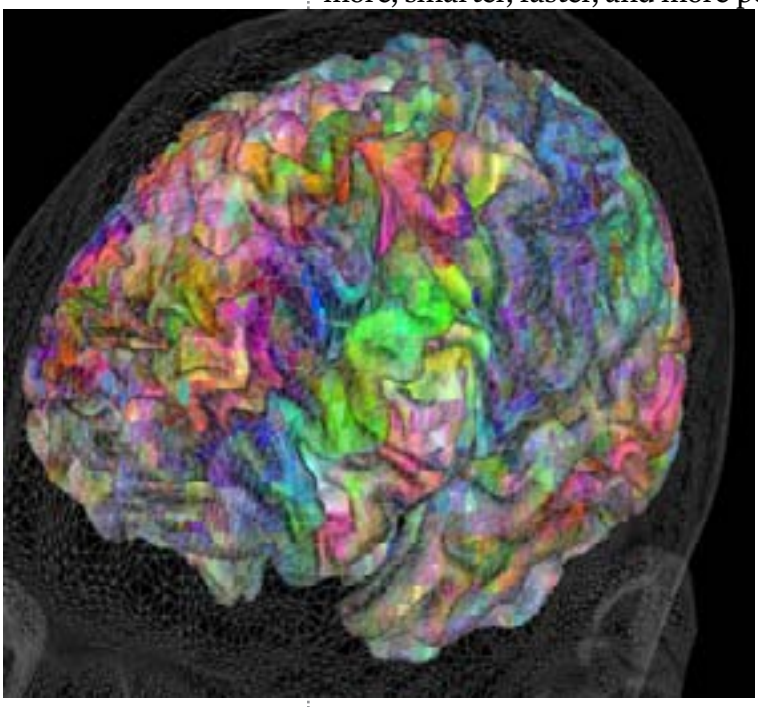

**SEMANTIC MODELS: Voxel-wise semantic models are created from fMRIs (functional magnetic resonance imaging) taken of subjects listening to stories yielding a new way of mapping functional representations in the brain.** *Nature*

## Leon Cooper & Neural Networks

### NESTOR: A BUSINESS FULL OF PROMISE: MANY FAILURES

In the very first issue of this newsletter, I wrote about Nobel laureates involved in artificial intelligence. To me, then, in May of 1984, Leon Cooper's approach to the future of computing seemed to hold the most promise. Cooper won his 1972 Physics Nobel Prize for his part in the first microscopic theory of superconductivity since its discovery in 1911.

Cooper did that research at the Institute for Advanced Studies. By the time he gave his Nobel Prize acceptance speech he had moved to Brown U. In a later paper, he discussed his desire to tackle a really important, massive and difficult set of problems: to probe and comprehend the nature of how we learn and remember.

His 1973 paper, "A possible organization of animal memory and learning," from the *Proceedings of the Nobel Symposium on Collective Properties of Physical Systems*, (That paper is my book: *Neurocomputing - Foundations of Research*, MIT Press, 1988) shows Cooper's inspiration to pursue NNs. He went on to found Nestor, a company attempting to make use of NNs for pattern recognition.

Intelligence

EDWARD ROSENFELD
*Editor & Publisher*

Maxine Davidowitz
*Design*

Maxim Druzhinin
*Webmaster*

Philip Chapnick
Neal Goldsmith
John Oberteuffer
Marty Perlmutter
Stuart Samuels
Gerd Stern
Brian Van Der Horst
*Contributing Editors*

INTELLIGENCE
PO Box 20008
New York, NY,
10025-1510, USA
212-222-1123 (Voice & FAX)
Email: ei@eintelligence.com
http://eintelligence.com/

Published Monthly
Since 1984
Subscriptions:
US$495/year (PDF)
ISSN: 1042-4296



Nestor originally wanted to provide NN software the recognition of Chinese and Japanese characters for computing systems. That didn't work out, despite a contract from the then powerful DEC (DIgital Equipment Corp.). After that, the company, run by Cooper and other associates from Brown, set its sight on commercializing solutions for other domains, including check reading for banks, fraud detection for financial institutions and, finally red light traffic cameras.

Despite great talent, perceptive insights into the nature of NNs and a number of important business connections, Nestor could never find success. After an initial public offering of stock to the public in the mid-1980s, Nestor, in the end, showed that good ideas and even good business execution don't guarantee success.

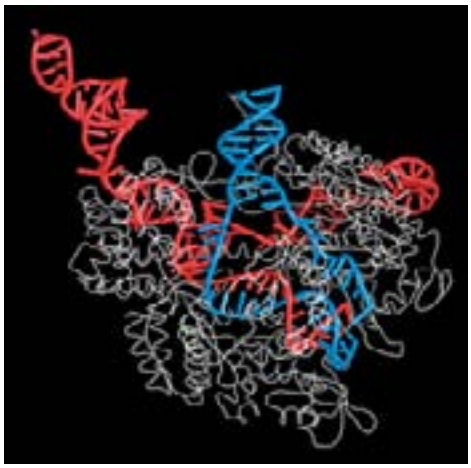

**MODEL OF THE CRISPR–CAS9 GENE-EDITING COMPLEX. MOLEKUUL/SPL**

As I look at the various ventures and approaches that seek to put machine learning to its best uses, winners and losers in business are a tricky group to discern. Certainly large resources help. The leaders in the machine learning field today include most of the very largest companies in technology. Amazon, Apple, Baidu (the Chinese search engine company), Facebook, Google, IBM and Microsoft have all made great strides in implementing successful machine learning systems. Now, to me, Google seems to have a commanding lead in NNs and DLNNs. But it would be foolish for me to discount the potential of any of these tech giants and others companies we don't yet know.

# **Future Plans for** INTELLIGENCE

## NEW FEATURES WILL BE ADDED THIS YEAR AND THE EINTELLIGENCE.COM WEBSITE WILL BE REVISED

Over the coming months, *INTELLIGENCE - Making The World Work* will expand as I add new features. Much of this new content will be added to the newsletter's "**ei**" website: http://eintelligence.com. Next month, I'll be posting on **ei** a contribution from *INTELLIGENCE* Contributing Editor Marty Perlmutter on the current state of what's happening in the VR (virtual reality) arena. His short article, entitled "I'm Not With Stupid," sums up for me the inflated state of promises in that nascent field. I'll also be adding back issues, an archive of previous *INTELLIGENCE* images as well as interviews with those I feel can best explicate what's happening so fast in machine learning, interface, internet, chips, biological and quantum computing, networking and scores of other technology domains. I look forward to your feedback: ei@eintelligence.com.