MAKING THE WORLD WORK
Intelligence
VOLUME 34, NO. 8 • DECEMBER 2017

The Final Issue of the INTELLIGENCE Newsletter

THE HISTORY AND DEVELOPMENT OF NEURAL NETWORKS AND MACHINE INTELLIGENCE

NGC 362
NGC 362 is one of only about 170 globular clusters of stars that exist in our Milky Way Galaxy. This star cluster is one of the younger globulars, forming likely well after our Galaxy. NGC 362 can be found with the unaided eye nearly in front of the Small Magellanic Cloud, and angularly close to the second brightest globular cluster known, 47 Tucanae. Hubble Space Telescope

This is the 404th and last issue of the **INTELLIGENCE** newsletter. I began this newsletter, in planning, preparation and marketing in 1982 and 1983, with the first issue published in May, 1984 as **INTELLIGENCE - The Future of Computing**. The title changed in 2014 to **INTELLIGENCE - Making the World Work**.

In that very first issue, in an article on Nobel laureates in artificial intelligence (AI), I reported on what would come to be known as neural networks (NNs). I was intrigued by this more neurobiological approach to computing and continued to cover the development of NNs and related approaches to machine intelligence (MI) for the life of this newsletter: 1984 - 2017.

Over the past five years, the NN field, often called AI, has exploded and has found its applications stretching to thousands of different domains and uses, mostly taking advantage of the capability of NNs to recognize patterns of all sorts. Tens to hundreds of articles are published each day, around the globe on NNs, AI, MI, etc. My original goal, to contribute to these fields in any ways I could, has reached a transition point. NNs have become a most powerful force in contemporary technology. I am pleased and can now move on.

How did we get here? Why are NNs and MI so important? Over the years I've done my best to provide answers, in the pages of this newsletter as well as in my three books on NNs: *Neurocomputing - Foundations of Research*, *Neurocomputing - Directions for Research* and *Talking Nets - An Oral History of Neural Networks*, all published by the MIT Press. A brief history of MI as well as my involvement and motivations, is in order here for this parting issue.

Computing and intelligence owes so much to Alan Turing. His work from the 1930s forward until his death in 1954, established the foundations of many of the most fundamental aspects and foci for contemporary computer science and technology. His 1950 paper on "Computing machinery and intelligence," and his creation of the concept of a Turing Test, to distinguish between humans and machines, are all pioneering and far reaching works.

My introduction to the concepts and theories inherent in machine intelligence came by way of reading science fiction. Here I learned about automation and robots as well as about "intelligent" systems that could do some things better than humans. The most important work I read in science fiction was *Last and First Men* by Olaf Stapledon, originally published in 1930. I learned over the years that Stapledon had anticipated most of the themes of science fiction, including the global brain.

NGC 1365
Barred spiral galaxy NGC 1365 is 200,000 light-years across. Located 60 million light-years away toward the chemical constellation Fornax. This image shows star forming regions at the ends of the bar and along the spiral arms; details of dust lanes cut across the galaxy's bright core. At the core lies a supermassive black hole. Dietmar Hager, Eric Benson, Torsten Grossmann

As I read more widely about automation and computing, I was especially fascinated by cybernetics. Even though I couldn't do the partial differential equations, I became intrigued by the prospect of self-organizing systems. As a founding editor of *Omni* magazine, I edited articles on robotics, automation, computing and AI.

After *Omni*, when I was on the editorial board of *Psychology Today* magazine, another board member, Joe Bianco, an attorney and financier, told me about two new companies in the AI space. One was Cognitive Systems, based around the work in natural language by Roger Schank (then at Yale U). The other was Nestor, based around the work of Nobel prize winner Leon Cooper (Brown U); the company I wrote up in my first issue.

In 1984, when I started **INTELLIGENCE**, I went to my first conference on AI. What I saw and heard there convinced me that AI was too brittle, too arcane and didn't really do much. To me, it seemed that specialized hardware and programming (LISP: list processing) was used to instantiate a very tight structure of if:then rules. I didn't see any future in AI and focused my newsletter on NNs. Even today, AI experts decry the advance of NNs and claim that AI methods are worthy. Their results do not support this conclusion.

It was when I was first attempting to read the NN literature that I found, in the days before the Internet, popular scientific and technical papers were hard to come by. When I located where Bernard Widrow's paper, "Adaptive Switching Circuits," (1960) was supposed to be, I discovered only razor marks where the paper had been cut out of the conference proceedings. This motivated me to suggest a book, collecting historical NN papers, to a subscriber, the late Frank Urbanowski of the MIT Press; this led to the two *Neurocomputing* volumes.

Widrow's work inspired a pioneering approach to NNs: back propagation. As I got to know people in the NN field, I heard so many different claims about who had "really discovered" back prop, that I was motivated to do an oral history of the field. *Talking Nets* has many accounts of how back prop came to be. That with other stories of shared research triumphs and tragedies, gave the book a flavor like Akira Kurosawa's insightful and great film, *Rashomon*.

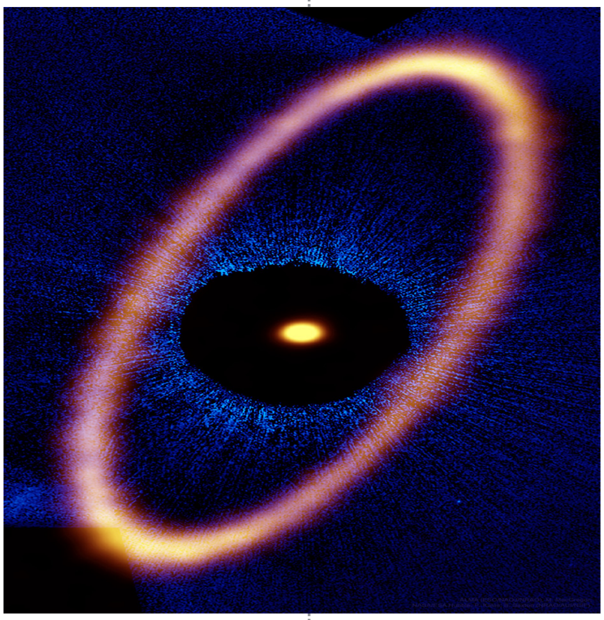

FOMALHAUT ICE RING This image by the Atacama Large Millimeter Array shows this outer ring with complete and unprecedented detail -- in pink -- superposed on a Hubble image of the Fomalhaut system in blue. A theory holds that this ring resulted from numerous violent collisions involving icy comets and planetesimals, while the ring boundaries are caused by the gravity of yet unseen planets. ALMA (ESO/NAOJ/NRAO), M. MacGregor; NASA/ESA Hubble, P. Kalas; B. Saxton (NRAO/AUI/NSF)

In that film, we see four people give differing accounts of the same story, about a man's murder and the rape of his wife. But, as I tried to demonstrate in *Talking Nets*, too, it is the very nature of truth that needs to be carefully considered and is always in question, whether about our stories or computer science discoveries. Back prop has gone on to become a great part of the foundation of the current success and scaling of NNs and their multitude of contemporary applications and systems.

But before back prop's first ascendancy in NNs, the field had both promising and difficult times. The original theoretical breakthroughs in NNs are presented in the opening papers in the first volume of *Neurocomputing*; works from: McCulloch and Pitts, Hebb, Lashley, Von Neumann, Rosenblatt, Selfridge, et al. And it was Frank Rosenblatt's perceptron that initially ignited public interest in the NN field and eventually led to the first of several "neural net winters."

Rosenblatt made very extravagant claims for what his patter recognition system could do. It was the first time the general public heard about something called a brain machine, what he referred to as "a model for information storage and organization in the brain." "*The New York Times* reported the perceptron to be '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.'" Perceptrons could not, in fact, do all these things promised. And then, Rosenblatt died in a mysterious boating accident.

Despite the fact that the famous AI scientist Marvin Minsky's PhD thesis was on NNs. he and Seymour Papert wrote a book, titled *Perceptrons*, purporting to show that NNs, like perceptrons, could only do limited pattern recognition.

Despite errors, including errors that the authors knew about, the book had a chilling effect on government and corporate funding for NN research that lasted for more than a decade, bringing on the first and longest of the NN winters.

But after that time, in the early 1980s, thanks to the research and theoretical work of people like Marr, Fukushima, Grossberg, Hopfield, Kohonen, McClelland and Rumelhart and many others, NNs once again blossomed and showed real promise. During this period, Jasper Lupo, then the head of the US Defense Advanced Research Projects Agency (DARPA), said at a keynote address at an early NN conference: “I believe that this technology [NNs] which we are about to embark upon is more important than the atom bomb.”

With venture capital as well as government funding pouring into the NN field, expectations were high. Once again, as had happened with Rosenblatt’s promotion of the perceptron, the hype about what NNs could do exceeded what was possible. It turned out that NNs didn’t really scale to many problems. There were successes, as when Robert Hecht-Nielsen’s HNC beat out Nestor in financial services applications. HNC was one of the first NN start-ups to be purchased by a major company: Fair Isaac.

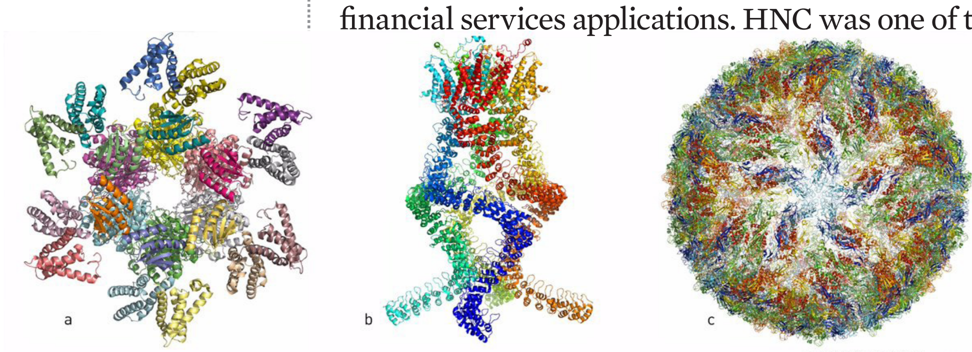

PROTEIN, SENSOR & VIRUS
Atomic structures of a) protein complex that governs circadian rhythm b) pressure sensor of the type that allows us to hear c) Zika virus. The Royal Swedish Academy of Sciences

For a while in the 1990s, it was thought that what was needed to make NNs scale and find great applicability would be the development of dedicated processors, NN chips. They were produced and, for one, the US government was among the first to want to invest and test them out.

These NN chips were the subject of one of my favorite stories that I heard that I couldn’t publish, then, in **INTELLIGENCE**. I was at a NN conference and I was talking about a “black” [secret] project in object recognition. “Can I speak to you off the record?” I was asked. I assured that our conversation would go no further. I was told that the NN chip being used to, ostensibly ”read license plate numbers in Kremlin parking lots” (this was still during the cold war) was very expensive.

“We’ve got to create several different working models and we didn’t get enough money to buy all the NN chips we need. So, I figured out a way we could achieve the same results needed in our models using DSP chips,” my source told me, referring to digital signal processing chips. DSPs then cost several dollars while the NN chips in question cost several thousand dollars each. “Of course, when we produce the final working system, we’ll use the NN chips,” he assured me. That was when I realized it would be Moore’s Law that would advance NNs.

Moore’s Law, stated by Gordon Moore in 1965 and holding steady until this year (and more?), said that, roughly, every couple of years, the number of devices that could be placed on a single chip would double yet the price and size of these chips would remain fairly stable and constant. That Law was the most pertinent reflection of the acceleration taking place over the past few decades in the broad development of technology. And, as more features could be placed on chips, more and more things could be done with those chips.

Chip devices have now become so large in capacity and so small in physical instantiation that their operations are pushing up against physical limits, with excessive heat and cross talk between chip elements that cannot be easily controlled. Exotic metals, and combinations of metals, three-dimensional chip architectures and, soon, quantum computers are all pushing to keep Moore's Law going strong.

PHOTOREALISTIC FACES
Everyone of these faces are machine-rendered... Progressive Growing of GANs for Improved Quality, Stability, and Variation: We describe a new training methodology for generative adversarial networks (GANs). The key idea is to grow both the generator and discriminator progressively: starting from a low resolution, we add new layers that model increasingly fine details as training progresses. This both speeds the training up and greatly stabilizes it, allowing us to produce images of unprecedented quality, e.g., CelebA images at 1024. Nvidia

One of the most revolutionary technological developments came with desire of the US defense department to develop autonomous vehicles. What DoD and DARPA wanted was to take soldiers out of tanks and other vehicles. When I first started writing about the initial field trials under DARPA funding, the most promising system, from then Martin Marietta, couldn't find the side of the road, even at three miles per hour.

DARPA persisted and established its series of Grand Challenges. In 2005, a team from Stanford U, led by Sebastian Thrun, completed the Mojave Desert course in under seven hours, besting two teams from Carnegie Mellon U, where Thrun had previously led robotic driving efforts. According to the DARPA archives: "Vehicles in the 2005 race [over seven miles] passed through three narrow tunnels and negotiated more than 100 sharp left and right turns. The race concluded through Beer Bottle Pass, a winding mountain pass with a sheer drop-off on one side and a rock face on the other. ... there were more curves and narrower roads than in the 2004 Challenge."

Thrun's DARPA Challenge results, and his subsequent move to work at Google, led to the serious development of autonomous vehicles. Most all of the selfdriving systems now being developed, tested and "on the road" use NNs, in part to see where they're going. Since for decades, vehicles have been one of the most important parts of the global economy, the changes that autonomous vehicular systems are bringing to the transportation domain are quite revolutionary and are poised to change society and the world in profound, disruptive and unsettling ways, and soon. I have written in earlier issues repeatedly about the prospect for millions of job losses around the planet in the coming years, as "drivers" are replaced. Yet, governments worldwide seem oblivious to the threat.

After another AI winter in the 1990s, Google was responsible for a major advance in the current development and deployment of new NN systems. Geoff

Hinton's group at the U of Toronto, did better on standardized object recognition tests than ever before. As a result, Hinton did a series of visits and talks at Google for Andrew Ng and Jeff Dean's group there. Google had what Hinton and the entire NN field needed: big iron computing in profusion and some of the largest data sets in the world.

It was this combination of copious computing resources and enormous amounts of data, combined with algorithms developed primarily by Hinton, as well as Yann LeCun (NYU, Facebook) and Yoshua Bengio, (U of Montreal, Element AI, Microsoft), as well as others, that broke open the log jam to release thousands of successful NN applications. The new NN field now is often called by the name "deep learning" after a paper by those three NN scientists (LeCun, Y., Bengio, Y. and Hinton, G. E. Deep Learning. *Nature*, Vol. 521, pp 436-444.).

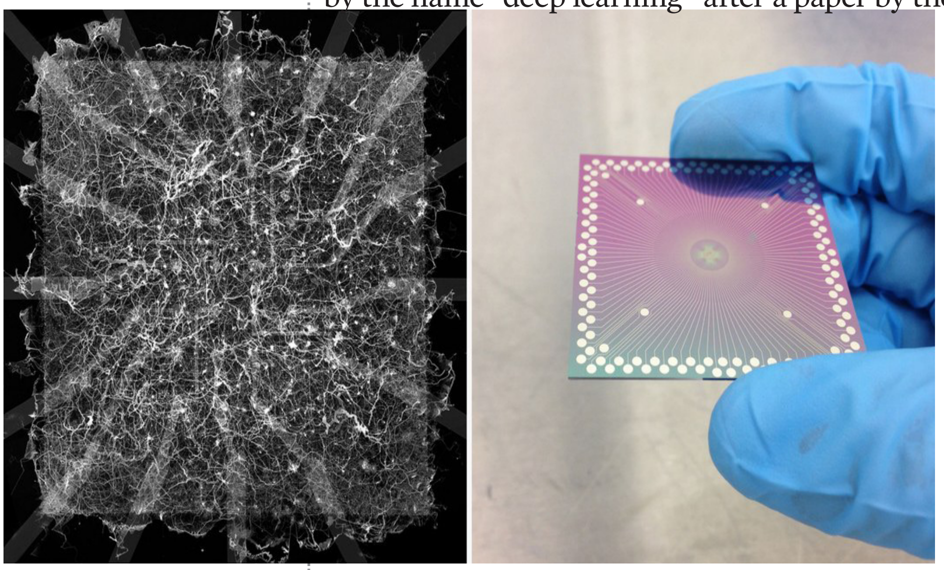

ATOMIC SWITCH BRAIN
The silver nanowire network (left) takes the form of a tiny square of mesh at the center of the device (right). The housing that holds the square mesh allows users to introduce signals as inputs and to measure the output results.Eleanor Demis (SEM image)/ Henry Sillin (Hand with device). A tiny self-organized mesh full of artificial synapses recalls its experiences and can solve simple problems. Its inventors hope it points the way to devices that match the brain's energy-efficient computing prowess.
Quanta Magazine

As a result of these and other breakthroughs in NNs, MI is now being put to work in many different fields. Robotics and automation originally followed the path of assembly line industrial processes, though computing machines did hollow out some employment in non-industrial ranks. Now, however, MI is reporting the news and sports scores, making medical diagnoses and decisions, and even writing computer code.

As this issue was going to press, two headlines were especially indicative of what's to come in the world of MI and automation: "Google's New AI Is Better at Creating AI Than the Company's Engineers," (https://futurism.com/googles-new-ai-is-better-at-creating-ai-than-the-companys-engineers/) and "Robot automation will 'take 800 million jobs by 2030' - report." (http://www.bbc.com/news/world-us-canada-42170100).

The first headline just underscores that it is not just computer programmers who will be replaced by MI systems but those at the very top of the AI jobs pyramid. Several recent reports have underscored the dearth of people available and experienced enough to fill top level AI jobs, those with starting salaries reaching >$100,000. Now, as MI systems gain in power and complexity, they will even replace the people who are building the AI systems!

The second headline is based on a "McKinsey Global Institute study of 46 countries and 800 occupations that found that up to one-fifth of the global work force will be affected. It said one-third of the workforce in richer nations like

Germany and the US may need to retrain for other jobs. Machine operators and food workers will be hit hardest, the report says." The report speculates that new jobs will be created as older jobs for humans are done by robots and other automation systems. "... new technology will yield new types of jobs, similar to the introduction of the personal computer in the 1980s which led to technology support work, and online business. The report's authors urge governments to enact plans to retrain their citizens."

But what of the transitional period between when the old jobs are finished and the new jobs are created? Does the current embrace of populist governments by those who feel left out of the modern world only represent the calm before the storm? Will the hundreds of millions of people who lose jobs, soon, even among a population of seven billion, sit by and quietly lament their fate? Probably not.

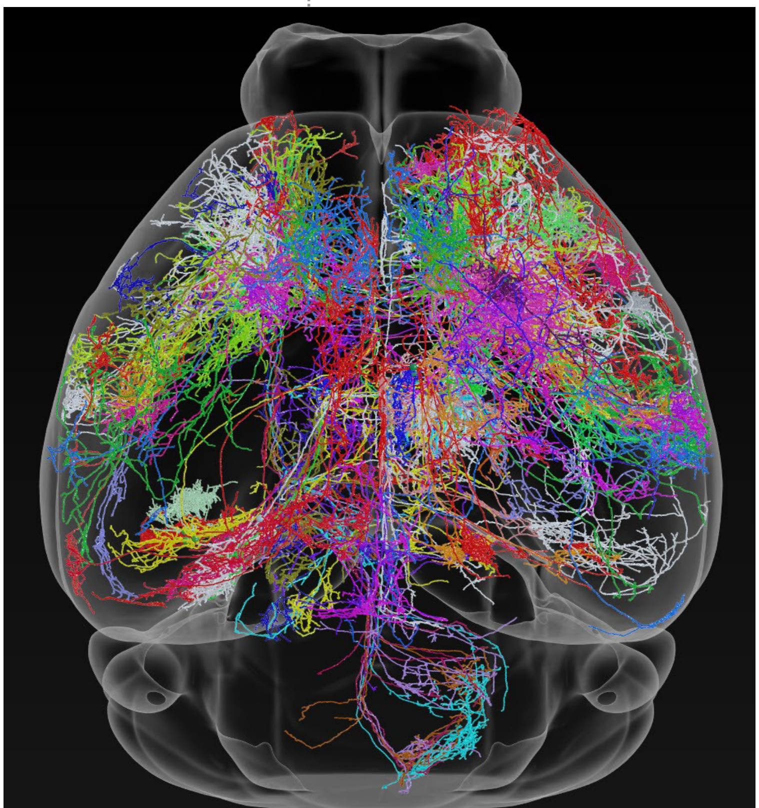

MOUSE LIGHT
MouseLight generates complete morphological reconstructions of individual neurons from datasets of whole mouse brains imaged at sub-micron resolution. We provide an interactive web platform called NeuronBrowser for anyone to explore, search, filter and visualize the single neuron reconstructions. http://mouselight.janelia.org

Yet, government inaction, on a global scale, is rampant and those who govern are in a state of perpetual denial about these issues and threats. These problems will not take care of themselves. Governments will have to take care of the many, or foster new ways for us to all take care of each other. If the inaction continues our governments will be destroyed by the social upheavals that are seemingly close at hand. Problems caused by social and financial inequality are pressing now, yet solutions seem far away, nonexistent.

Pondering and thinking about these global problems was what lead me to change the sub-headline of this newsletter. By 2014, I realized that intelligence had already brought on "the future of computing" and that it was time to focus more on larger issues from a different perspective. Seeing how rapidly MI has grown and the speed of its nearly universal acceptance was a big influence in deciding to stop publishing this newsletter and move on.

What's next for me is stated in the new sub-head I adopted in 2014: Making the world work. We will have to take advantage of intelligence, both human and MI, in order to move forward and deal with potential crises that lurk in the not too distant future. I cannot conceive of a more worthy goal than doing what I can to make the world work. I'm happy to look forward to devoting the rest of my life toward that end, which, if realized, would be truly a new beginning.

In the 1960s, I had the privilege of meeting and working with R. Buckminster Fuller, famous for his geodesic dome designs and so much more. Bucky Fuller

Intelligence

EDWARD ROSENFELD
Editor & Publisher

Maxine Davidowitz
Design

Maxim Druzhinin
CTO

Sankai Lemmens
Webmaster

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

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

Published Monthly:
1984 - 2017
Subscriptions:
US$495/year (PDF)
ISSN: 1042-4296

emphasized that we live on spaceship earth. We are all here, on this planet, and we are lucky to have such a rich and supportive planetary home. To me, one of Bucky's great insights was to observe and communicate that the Earth can support us all: there is enough to go around. He underscored that design and organization were ways to potentially overcome the distribution and other problems that inhibit the world from working.

Fuller took many steps to move toward a world that would work. Most specifically, he devised the World Design Science Decade 1965 - 1975: a program to gather information and ideas about how to surmount the many global problems that interfere with world working. As he stated it: "applying the principles of science to solving the problems of humanity."

Bucky Fuller's Dymaxion map shows planet Earth's land masses as "one island".

Around the same time, Fuller developed a simulation experiment for solving world problems that he called the World Game. Played on a 70-by-35-foot version of Bucky's Dymaxion map, the aim of the game, again in Fuller's words, to: "make the world work, for 100% of humanity, in the shortest possible time, through spontaneous cooperation, without ecological offense or the disadvantage of anyone."

When I've mentioned to friends and colleagues that I'm ending publication of **INTELLIGENCE** and moving on to "making the world work," the response I've been receiving most is: "How are you going to do that?" My answer is that I don't know how to do it, but I'm comfortable with my uncertainty as I go forward to discover just what roles I might play in what I think is the most important effort of our times. If we're not the ones who start NOW to do what needs to be done to make the world work, then who will take on that goal and when?

To begin my new endeavors for world working I've started developing a new website: *worldworking.net*. The site is now under construction but I'll be working up materials to add to it over the coming months. I plan to officially launch the site to the global general public on Earth Day, 2018: Sunday, the 22nd of April. This will be two days after the publication of *The Book of Highs - 255 Ways to Alter Your Consciousness Without Drugs* (Workman) on 4/20. This is a new, revised edition of my book, written in the late 1960s and originally published in 1973.

Over the coming weeks, I'll continue working on the *eintelligence.com* website, adding new materials and completing the archiving of all of the 404 issues of the **INTELLIGENCE** newsletters. I want to make that website into a resource for those seeking to know more about the history, research and development of NNs and the many other related MI technologies that have been reported on over the years in these pages. Thanks for your continued interest and support.