Paper 1341-2017

The Rise of Machines, Analytics and Humans: Artificial Intelligence in the Age of the IIoT

Kathy Ball, Devon Energy; Robert Silverman, Franklin and Marshall College

ABSTRACT

In the next 10 years, the Industrial Internet of Things (IIoT) will dramatically alter nearly all sectors of the industrial economy, which accounts for nearly two-thirds of the global gross domestic product, according to findings at the World Economic Forum in Davos, Switzerland. The IIoT will radically change how humans, machines and our current infrastructure operate to achieve results and compete in the new digital world. Considered a disruptive technology, the IIoT will create new value streams for industries, including automated decisions and reactions in real time, massively improved operational efficiencies, connected infrastructure platforms and much better interaction between machines and humans. Erik Brynjolfsson, director of MIT's Initiative on the Digital Economy, said, "Humans must adapt to collaborate with machines, and when that collaboration happens, the end result is stronger." This session outlines three ways in which data analytics can help bridge the gap between humans and machines to achieve value in the IIoT world: 1) predictive analytics from a case study in the oil and gas industry; 2) an automation case study; and 3) text analysis from chaos theory. The session outlines the people, processes and technologies needed to enable this infrastructure. The session also covers some pitfalls to avoid regarding systems, silos and the human barriers to understanding artificial intelligence.

INTRODUCTION

Amid rapidly expanding curiosity about Machine Learning and Artificial Intelligence (AI), this session will attempt to provide real-world examples of AI in both consumer and industrial applications. It is important to remember that this paper will focus on AI in its current form, as opposed to the more vivid depiction of AI through science fiction. We will examine where AI currently sits along its expected path of development, as well as its broad spectrum of current and future uses. AI is viewed as the future, but commercial products such as Siri and Alexa, self-driving cars and Netflix recommendation engines show us that the future is arriving now.

What are AI, Deep Learning and Machine Learning?

To understand the relationship between the concepts of AI, Machine Learning and Deep Learning, it is helpful to picture three concentric circles, with AI as the largest and Deep Learning as the smallest. Machine Learning has emerged as a subset of AI research, and Deep Learning later emerged as a way of discussing highly complex forms of machine learning.

At its most basic level, work in the field of AI has been about attempting to combine machines' information storage and retrieval systems with the complex logic and adaptive nature of human thought. While there are several definitions of intelligence, they mostly involve information acquisition, comprehension and the application of knowledge toward achieving a set of goals. In cases of adaptive intelligence, these goals can be of either a pre-defined or emergent nature.

AI, in its basic form, is the capacity of a computer to perform operations analogous to learning and decision-making in humans, as by an expert system, or a program for the perception and recognition of shapes in computer vision systems. Modern technological improvements have enabled the growth of AI through the wider availability of GPUs (Graphics Processing Units) to make parallel processing faster, cheaper and more powerful. Leaders in numerous industries, such as Andrew Ng of Baidu Research, believe that AI will fundamentally change the way that work is done in their respective fields. Ng stated his belief that CEOs of S&P 500 companies will wish that they had started earlier on an overall AI strategy.

Artificial Intelligence

When most folks think of AI, they think of popular movie characters C-3P0 or the Terminator, each of whom mimics aspects of human behavior. General AI such as this has existed only in fiction prior to the formative work that is currently being pursued on the subject. Due to the high cost and low certainty surrounding General AI, most companies tend to use less complex forms of AI commonly known as Narrow Intelligence.

This form of AI performs specific tasks better and more efficiently than humans can. Examples of Narrow AI are things such as image classification on Pinterest to facial recognition on Facebook. These technologies exhibit a small piece of human-like intelligence, but do not achieve the more "intelligent" tasks desired. From this stage, analytics has truly blossomed with new computing power combined with improved advanced algorithms to better understand the achievement of logic and reasoning.

Machine Learning

Machine Learning as seen below in Figure 1 is a simple description of analytical techniques separated by topic. The use and advancement of these algorithms will catapult organizations to the more intelligent, Deep Learning and Cognitive Computing previously thought impossible. Machine Learning is a subset of AI that enables machines to improve their performance of tasks through experience. Most applications of AI today involve the use of Supervised Learning where the human explains the pattern to the machine.

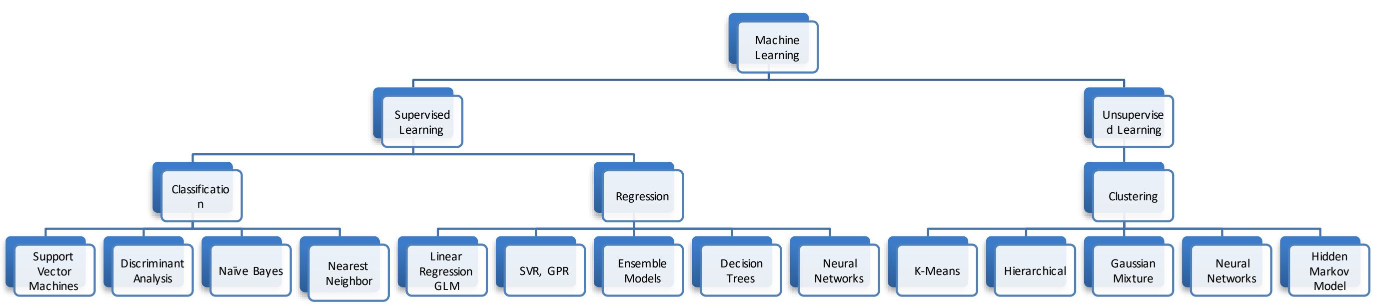

Figure 1. Caption for Machine Learning Algorithm

Source Vincent Granville Data Scientist, LinkedIn 2017

Artificial neural networks (ANN) are a form of machine learning inspired by how the human brain works through interconnections of neurons. Unlike a biological brain where any neuron can connect to any other neuron within a certain physical distance, artificial neural networks have discrete layers, connections and directions of data propagation. (See Figure 2) These networks come up with a probability vector, or a highly educated guess based on our weighting. The networks can analyze millions of images or datasets. Biological neural networks can process roughly a billion neurons connected along pathways throughout these networks. At a high level, neurons interact and communicate with one another through axons connected via synapses. Simply said, a single neuron will pass a message to another neuron if the summation of the weighted signals exceeds a threshold to cause message transmission.

The cost function is what learns the optimal solution for the problem being solved. This determines the best values for all the tunable model parameters such as the learning rate. The optimization techniques used are generally gradient descent or stochastic gradient descent.

The artificial neural networks are trained with labeled data like images with ImageNet. With unsupervised learning, the neural network trains on unlabeled data and is asked to look for recurring patterns. Thus, with the large amounts of data available, the machine would learn on its own by making sense of the world, similar to a human infant.

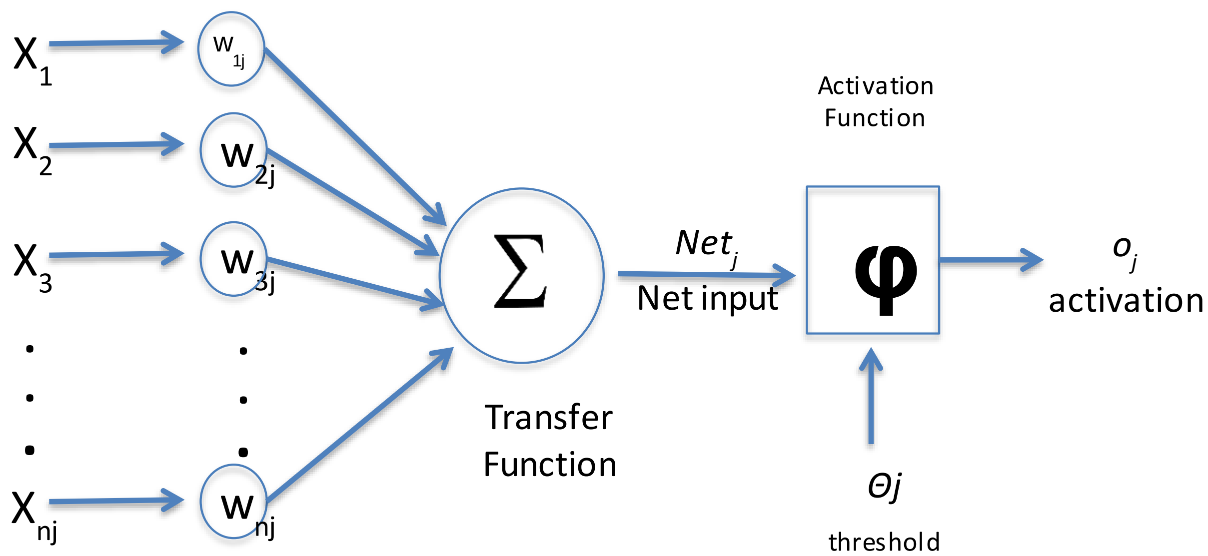

Figure 2: Diagram of Neural Network Activation

As seen in Figure 3, the first layer is the input layer, followed by a hidden layer, and last by an output layer. Models can become extremely complex by increasing the number of hidden layers. The most important thing to remember about ANNs is complex but considered black box algorithms whose inner workings are very difficult to explain to a layperson.

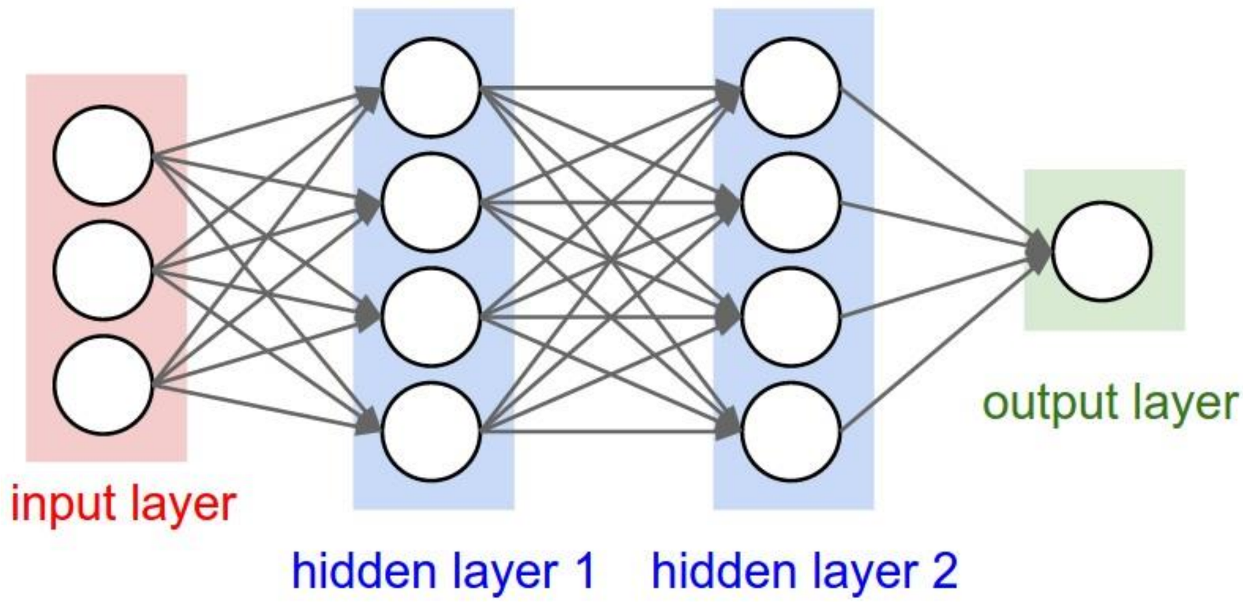

Figure 3: Hidden Layers in a Convolutional Neural Network

Deep Learning

Deep Learning is a subset of Machine Learning and is comprised of algorithms that permit software to train itself to perform tasks, like speech and image recognition, by exposing multilayered neural networks to vast amounts of data. Experts agree that Deep Learning will likely transform all industries. Jeff Dean, who leads the Google Brain project, has said, “There are fundamental changes that will happen now that computer vision really works … now that computers have opened their eyes.” It is important to understand the limitations of AI at this point. Neural networks are good at identifying patterns as well or better than we can, but they cannot reason. Neural networks offer the potential for computers to learn the way children do through trial and error rather than rule-based heuristics. Deep Learning is the use of advanced neural networks that use multiple layers to detect things like corners and edges joining at an angle. As more layers are added, the perceptions can analyze more difficult configurations. The neurons of each higher layer respond to greater levels of complexity and abstraction until the pattern is defined. The neural network must then determine if it is obtaining the right results, so it sends results back through the former layers so that the first neurons can retune their activations to improve the results, thus learning. As to the improvement over earlier Machine-Learning techniques, we can now use volumes of data currently available through sensors and connected devices so that data drives learning versus the other way around. A simple practical application of Deep Learning is the modern Chatbot that can replace customer service interactions through the smart use of technology.

Deep Learning’s greatest impact will be felt when it is integrated into the entire AI toolset such as the combination of Chaos Theory with Advanced Artificial Intelligence. One will define how the human behaves with Chaos Theory. Deep Learning and Machine Learning will define how intelligence can be garnered through machine interaction. AI will bridge the two concepts, similar to how Google’s DeepMind

has created “Reinforced Learning” playing the AlphaGo system. The same techniques applied to gaming theory can be used in Industrial Internet of Things (IIoT) examples where humans must interact with machines and AI helps to gain the operational improvements or intelligence desired. Deep Learning allows a machine to learn for a specific task using well-suited features and also learns from the features themselves. The algorithms virtually learn how to learn. The key point to remember is that due to the relative complexity in Deep Learning and neural network algorithms, the models are prone to over-fitting. Also, due to the increased model and mathematical complexity, additional computational and time are required. The second key point to remember is that Deep Learning models represent local solutions versus global solutions. The models have no obvious ways of performing logical inferences and they are still a long way from integrating abstract knowledge. Deep learning models can also be tricked by random patterns and/or pixels. Using Siri as an example, it appears as if the machine understands us and can tell a few programmed jokes. However, it can only comprehend and react based upon what it’s been taught. Supervised neural networks, like Siri, must have been told the answer to your question at some point for it to learn it.

It is fun to note that “Intelligence is one thing, common sense is another.” Author Unknown.

INDUSTRIAL INTERNET OF THINGS (IIOT)

The Industrial Internet of Things (IIoT) is the use of connected technologies in industrial applications such as oil and gas. The IIot incorporates machine learning and big data technology using sensor data and machine-to-machine communications (M2M). Automation and the increase in usage of sensored equipment in the field create tremendous opportunities in the oil and natural gas industry. The push for IIoT is based on the premise that smart machines are better than humans at accurately and consistently capturing and communicating data. Coupled with advanced analytics, companies can now receive this information, create knowledge and complete the loop by acting upon the discovery. IIoT can help companies discover inefficiencies and problems sooner, saving time and money and supporting continuous improvement efforts.

ARTIFICIAL INTELLIGENCE EXAMPLES

With AI used to solve real world problems, neural networks are designed to mimic our own decision-making. Deep Learning focuses even more narrowly on a subset of Machine Learning tools and techniques and applies them to anything requiring thought, whether human or artificial. One of the most common examples of DL today is seen in the analysis of language comprehension and processing. This is known as Natural Language Processing (NLP) and is fairly well documented. It can be broken down into two components: Content Categorization and Sentiment Analysis. This area of data science has become very popular through open source tools such as those in Figure 4. Our first examples of AI use NLP to help users better understand Chaos Theory. Since AI and Chaos Theory are linked to future success in this field, the example we provide should help you better understand through Machine Learning how to interpret the issues around Chaos Theory.

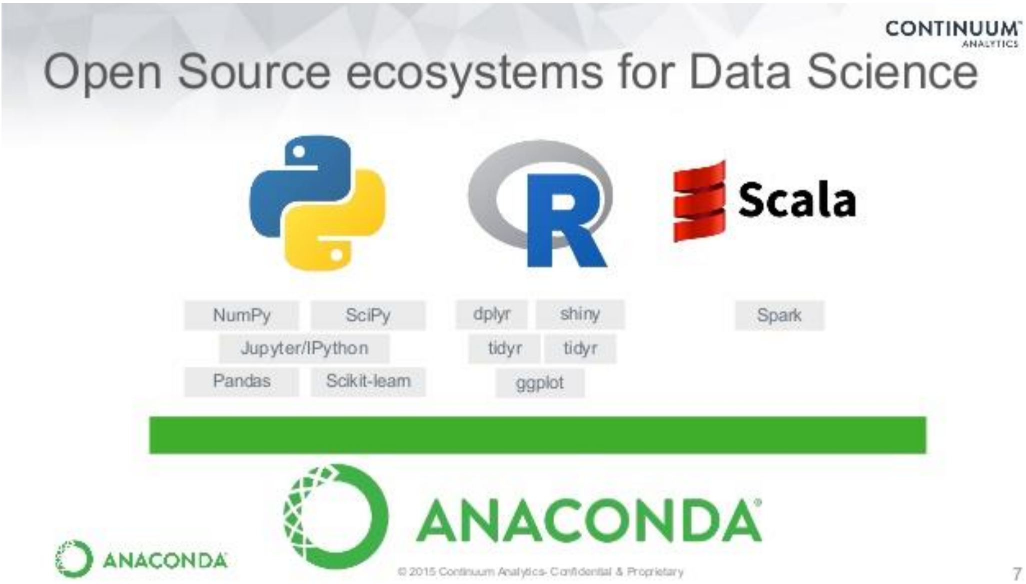

Figure 4: Open Sources Analytics Tools that can be used with SAS®

EXAMPLE 1: TEXT ANALYTICS USING NLP ON CHAOS THEORY

One of the best movies introducing AI to a layperson was *Jurassic Park*. In the film, one character explains Chaos Theory to another. Chaos Theory simply deals with unpredictability in complex systems. It is most commonly known as the butterfly effect. A butterfly can flap its wings in Africa and, as a result, it will rain in the U.S.

Humans, much like the tyrannosaur, do not always obey set patterns or schedules – hence, chaos. The character went on to describe trying to predict which way a drop of water would fall off your hand. There are so many tiny variations that will affect the outcome, leading to unpredictability. The attempts to reach sentient AI will rely on our ability to computationally narrow the unpredictability "cost function."

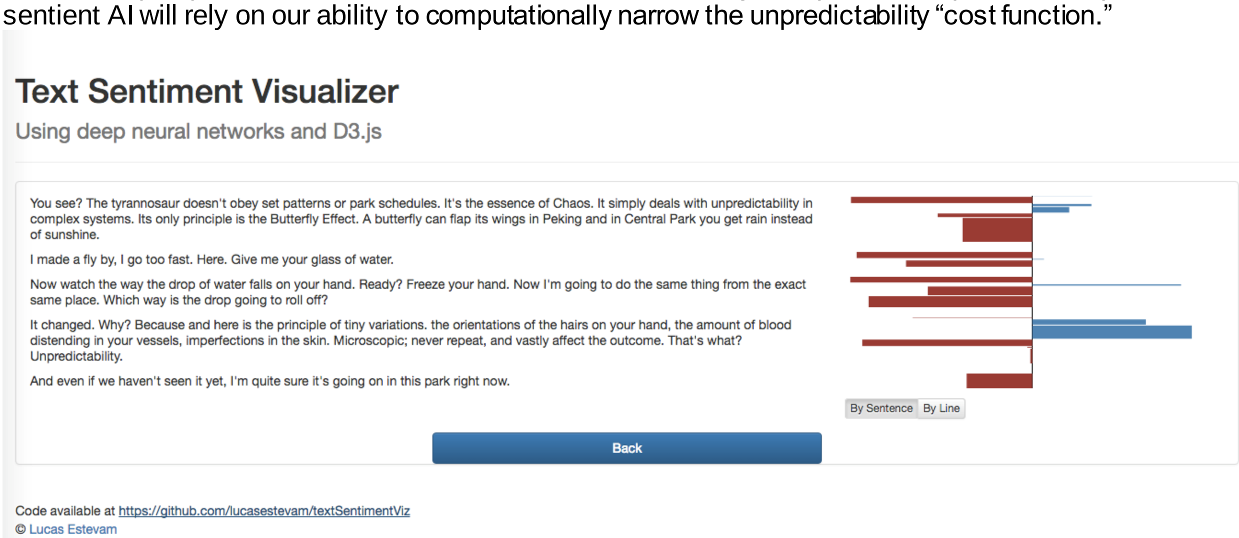

Figure 5: Open Source Text Sentiment Analysis of *Jurassic Park* Script

EXAMPLE 2: DEEP LEARNING USING SUPERVISED MACHINE LEARNING

As machines have tended to learn based on what it has been taught, many examples exist of rule-based learning to optimize or automate a particular task. In this example, the business problem is to train a neural network to discover video images for particular types of downhole video logs. The point of the process is to use technology to improve video analysis for oil, gas and water detection within a well.

The primary challenges for production log videos are time constraints, data volumes are tremendous and text can appear in nonstandard ways. Deep Learning using Supervised Learning allows us to speed time to discovery and creates actionable insights. This type of problem requires a large data solution such as Hadoop and Deep Learning analytics tools not present in the market. In this example, we were able to create a frame-by-frame breakdown with pattern recognition technologies to identify changes in text as well as changes in input in the process of defining how to detect water, oil and gas. Figure 6 displays different events/features that were automated using image recognition techniques for detecting water, gas flow, oil flow and wellbore issues. We were able to determine depths, time and orientations automatically within hydrocarbon wells.

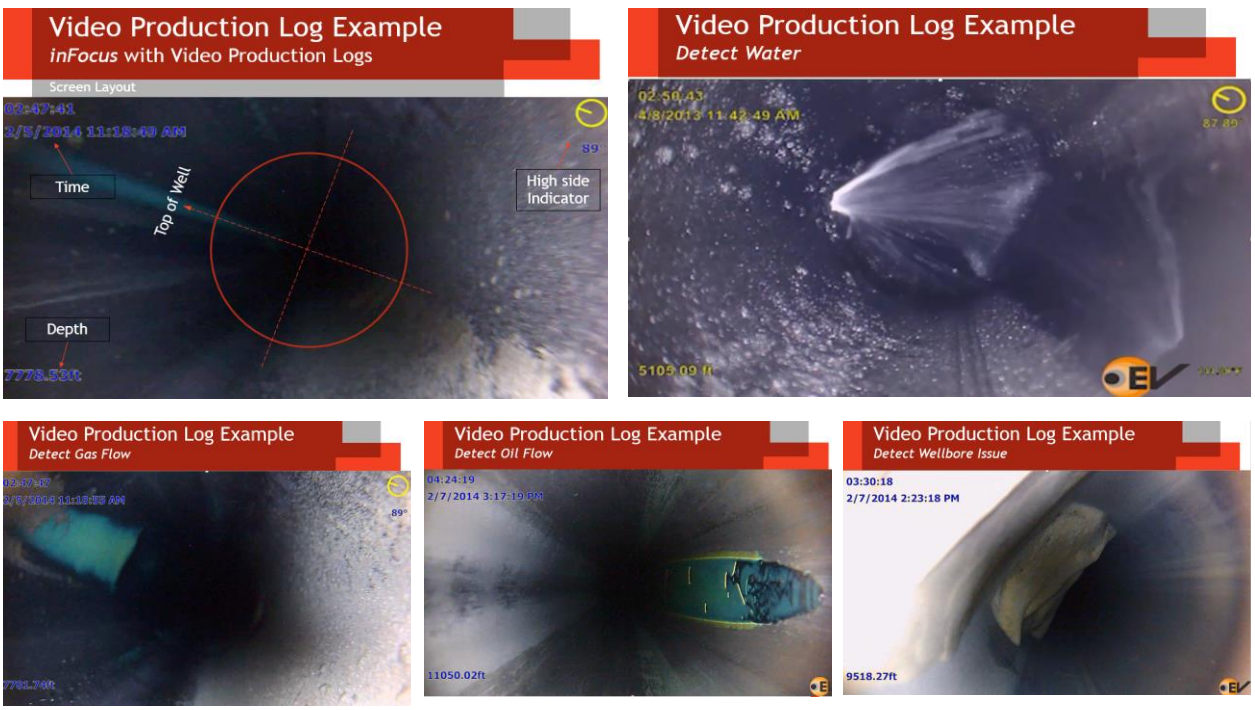

Figure 6: Supervised Deep Learning in Video Image Processing

EXAMPLE 3: DEEP LEARNING USING UNSUPERVISED MACHINE LEARNING

As machines have tended to learn based on what has been taught, many examples exist of rule-based learning to optimize or automate a particular task. In this example, the business problem was to train computers to perform human-like tasks such as finding patterns via image processing. This example attempts to find patterns in microseismic surveys to better determine reservoir properties. The detectable patterns are time consuming and difficult to find when automating for hydrocarbon discovery. Neural network needs to train a five-layer autoencoder with multiple hidden units per layer and then use stacked denoising autoencoders to diminish noise and/or corruption. The results are showing positive results around clustering and groupings that could be of use in further seismic interpretations. Technology is still in its infancy in terms of advanced artificial intelligence.

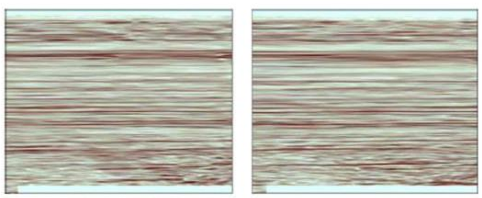

Figure 5. High-resolution seismic images.

The *k*-means algorithm usually results in spherical clusters. However, clustering patches (as opposed to entire images) can result in complex clusters in the original image as can be seen in Figure 7.

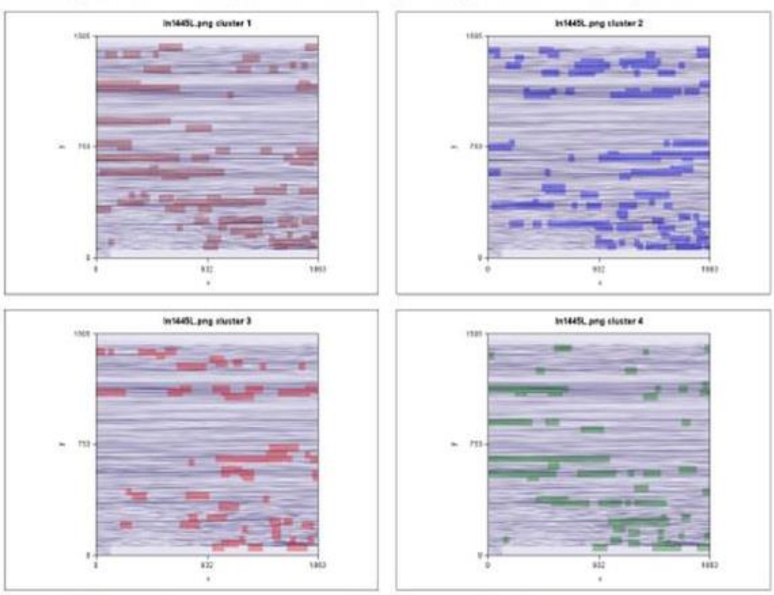

Figure 7. Four clusters of patches overlaid onto a single inline, two-dimensional seismic image.

Figure 7: Unsupervised Deep Learning to Detect Patterns and Trends in Microseismic

CONCLUSION

The IIoT allows objects to be sensed and/or controlled remotely across existing network infrastructure, creating opportunities for more direct integration of the physical world into computer-based systems, and resulting in improved efficiency, accuracy and economic benefit in addition to reduced human intervention. This involves the blending of technology and AI. When the IIoT is augmented with automation such as sensors, the technology becomes an instance of the narrower types of AI. Each thing is uniquely identifiable through its embedded computing system but is able to interoperate within the existing interconnected infrastructure.

Typically, the IIoT is expected to offer advanced connectivity of devices, systems and/or services. The interconnection of these embedded devices will be seen in connected cities, smart grids, self-driving cars and other applications where automation creates efficiencies. "Things" is a mixture of hardware, software, data and services. Devices collect useful data with the help of various existing technologies and then autonomously flow the data among other devices.

The IIoT can be intimidating and it can be difficult to make the most of networked devices. Upgrades to infrastructure, processes and enhanced security are essential. The Industrial Internet of things (IIoT) is considered to be the fifth digital revolution that is changing the way we live and work.

REFERENCES

- Forbes.com, *Risks give pause, but rewards abound with networked devices, Tech Target*
- ComputerWeekly.com, *Big data to unlock value from the industrial internet of things, Nick Millman*
- Fortune.com, *The AI Revolution: Why Deep Learning is Suddenly Changing Your Life*, September 28, 2016
- NVIDIA Blog, *The Difference Between AI, Machine Learning, and Deep Learning*, Michael Copeland, July 29, 2016
- Archi, Artificial Intelligence, *Deep Learning, and Neural Networks Explained*, Alex Castrounis, September 1, 2016
- The Verge, *What counts as artificially intelligent? AI and deep learning, explained*, James Vincent, February 29, 2016

RECOMMENDED READING

- *Base SAS® Procedures Guide*
- *SAS® Enterprise Miner Procedures Guide*
- *SciKit-Learn Documentation (SciKit-Learn.org)*

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Kathy Ball
Devon Energy
Kathy.Ball@dvn.com

Robert Silverman
Franklin and Marshall College
rsilverm@fandm.edu