Sean Maher | sean.maher@entext.com

Rise of the Machines: The AI Arms Race…

'A Northwestern University team has developed a new visual problem-solving computational model that performs in the 75th percentile for American adults on a standard intelligence test. The research is an important step toward making artificial-intelligence systems that see and understand the world as humans do. The researchers tested the AI system on Raven's Progressive Matrices, a nonverbal standardized test that measures abstract reasoning. 'The Raven's test is the best existing predictor of what psychologists call 'fluid intelligence,' or the general ability to think abstractly, reason, identify patterns, solve problems, and discern relationships," Professor of Computer Science Ken Forbus, whose research paper was published in the latest issue of academic journal Psychological Review

'A lot of what we would traditionally think of as human intelligence is built around pattern matching and a lot of what we would think of as learning is having seen these patterns in the past, and being able to realize how they connect to a current situation.' Michael Bowling, professor of computer science at the University of Alberta

'The revolution in AI has been profound, it definitely surprised me, even though I was sitting right there.' Sergey Brin, Google co-founder

- The term "artificial intelligence" was first coined in 1956, but after much premature hype and a series of false starts has finally arrived as a viable technology, thanks to simultaneous advances in parallel processors and machine learning software. **Combined breakthroughs in hardware, scalable cloud computing, "deep learning" algorithms and the availability of unprecedented amounts of tagged data have allowed AI to reach human-level accuracy in image/facial and voice recognition**. The industry has gone through several boom-and-bust cycles, the most recent of which was in the late 1980s when tech companies bet billions on "expert systems" that could supposedly emulate human decision-making ability. They proved too expensive to maintain and the technology was effective only in a narrow set of applications, leading to a collapse in interest in the sector.

- However, as Google's Brin acknowledges, **progress has accelerated dramatically over the past five years and surprised even industry insiders** - the latest applications include cataloguing images, translating documents, interpreting medical scans and legal documents, flying drones (even in intelligent 'swarms', extracting predictive value from consumer behavioural data sets etc. While there is undoubtedly hype (in areas such as the timeline for truly autonomous vehicles, for example) **AI appears to finally be at an inflection point, which will drive very rapid changes in the economic and investing landscape over the next 5-10 years.**

- The jargon can seem confusing but essentially, instead of adding more lines of instruction code to a program, engineers can via a learning model show the software enough examples of correct outcomes for it to draw its own inferences and adapt behaviour accordingly. This 'learning' process will become increasingly autonomous and adaptive over the next few years. Unsurprisingly given the stunning series of headline grabbing breakthroughs (and huge sums being spent to buy promising start-ups, often amounting to tens of millions of dollars per engineer), the brightest computer science, behavioural psychology and neuroscience graduates want to be in AI and help make computer software more like the human brain.

- Registrations to the US flagship Neural Information Processing Systems conference soared almost 50% this year. Amazon's Alexa virtual assistant (while relatively primitive) is demonstrating the potential of voice interaction with computers, which will soon be in many offices as well as homes. Machine learning techniques as well as faster processors have been the key to unlocking AI potential; the basic premise of machine learning is that instead of hardcoding specific sets of instructions, AI **algorithms instead parse huge volumes of raw data, learn from it and make predictions about the future**.

- It's much harder for a computer to learn simple inferences than a human brain, but they **can once 'trained' via this iterative data crunching apply the lessons far more reliably and cheaply than humans,** from recognizing anomalies in potentially cancerous skin moles to fraudulent insurance claims. The most powerful parallel processing chips and the largest possible pool of 'trainers' tagging/using the data being analyzed helps the development of superior AI applications (the huge volume of tagged images being created every day helped AI crack vision, for example).
- The next big target is language – the techniques that have produced spectacular progress in voice and image recognition may also help computers parse language more effectively as voice gradually becomes our default interface with computers. The challenge is huge, given the complexity, subtlety, and power of language but **then even a couple of years ago, few experts would have predicted AI to master complex games like Go or poker (see below) so rapidly.**

Neural Networks Need to Replicate Human Inference, Become 'Data Efficient'

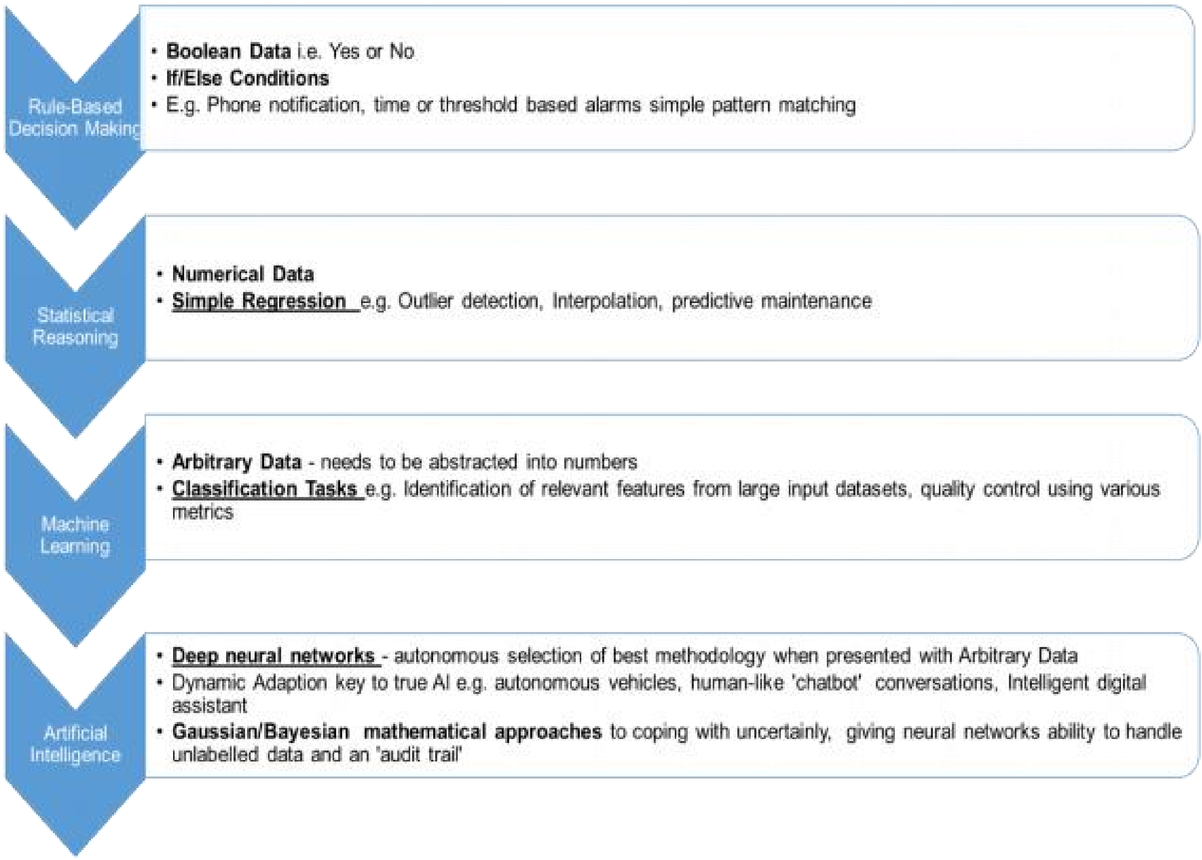

- **Right now, a neural network having been trained can only fulfil that one specific task, unlike an endlessly multi-tasking human brain.** They cannot add new skills without erasing what they already know so an AI trained to recognise buildings in photographs wouldn't cope well with another visual task, such as recognising street signs. Each system would need to be trained for its own limited task - that's one reason why replicating a human driver in a fluid real world scenario will be so difficult as so many different AI systems and associated computing power would have to be in the vehicle, as covered below. Meta-learning algorithms may be a medium-term solution - they learn to perform well on specific tasks, but are designed to discover fundamental rules of learning itself, so that the same algorithm may use the inferred rules to solve other related, but previously unseen tasks.
- **Neural networks have another key limitation – they can't make sense of the world without large amounts of carefully labelled data.** Furthermore, neural networks don't explain why they make particular decisions, which will be critical for an 'audit trail' in applications like healthcare or

driverless vehicles where human safety is an issue. Those constraints have led to the rise of researchers that approach AI by beginning with a hypothesis and then updating this hypothesis based on the data, rather than just relying on brute force computing to drive the conclusions, as neural networks do.

This 'Bayesian' probabilistic approach looks for ways of dealing with uncertainty and feeding new evidence into existing models. Bayesian methods allow AI systems that can learn tasks by inference from just a little data, much like humans do. Such 'small data' systems are essential to building machines that can carry on a conversation or cars that can drive public roads all on their own. Cars will never have enough data or on-board computing power to use brute force iteration the way that deep learning does.