IEEE Buenaventura Section RAS/IAS

ARE THE MACHINES LEARNING?

KARL GEIGER

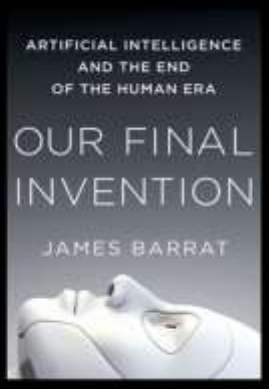

Popular Press:

OH NO, WE ARE DOOMED

If not doomed, machines will take our jobs and we will all starve.

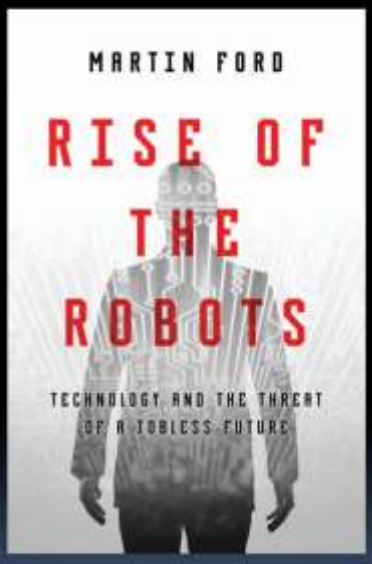

Experts tell us …

"The development of full artificial intelligence could spell the end of the human race."

-- Stephen Hawking, December 2014

… beware summoning demons.

"I think we should be very careful about artificial intelligence. If I had to guess at what our biggest existential threat is, it's probably that. So we need to be very careful."

"With artificial intelligence we are summoning the demon. In all those stories where there's the guy with the pentagram and the holy water, it's like – yeah, he's sure he can control the demon. Doesn't work out."

-- Elon Musk, October 2014

Why the Hubbub?

- Machines-taking-over is an old trope.
- It comes around with a vengeance every 20 years or so …

1990s / Early 2000s

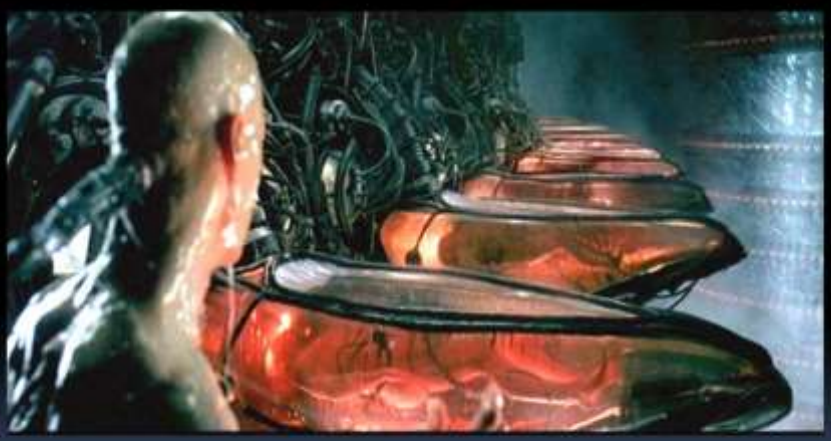

WHOA – THE MATRIX!!

1980s

Johnny Five

Terminator T800

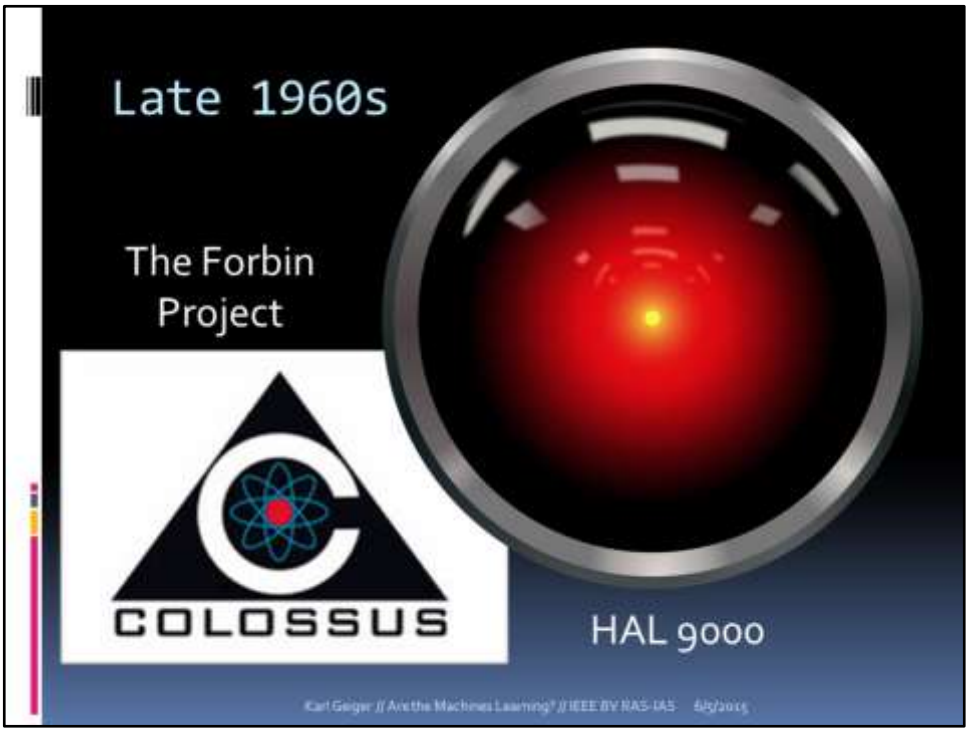
Late 1960s
The Forbin Project
COLOSSUS
HAL 9000
Karl Geiger // Are the Machines Learning? // IEEE RV RAS-IAS 6/5/2015

Interestingly, the scary machines from the 1950s and 1960s were faceless mainframes.

Then we had robots coming to kill us.

Now we're back to the faceless mainframes "in the cloud", eg, the Googlebeast, Amazon Web Services, MS Azure, etc.

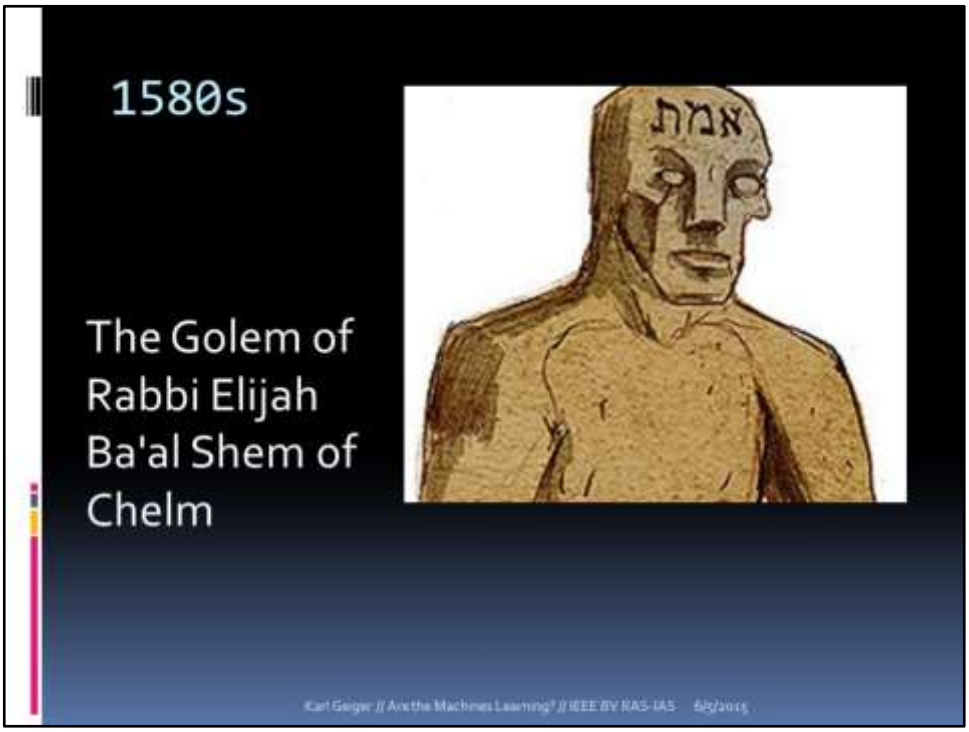

The rabbi created the golem from clay to perform drudge work. The golem began to grow and dangerously out of control. Fearing his destruction, the rabbi removed the holy inscription from the golem's forehead, whereupon it returned to dust and clay.

What we really fear is a robot slave revolt.

No one wants to be enslaved, a la *The Matrix*.

It's just like monster movies: no one wants to be turned into food and eaten, either.

Technical Background

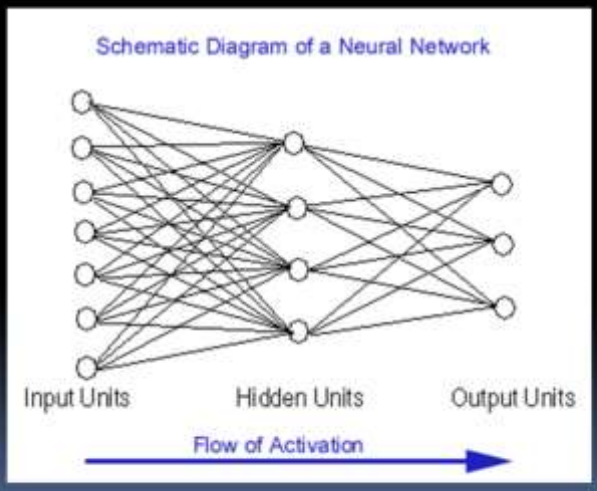

STANDARD DISCLAIMER

- I am *not*
  - a Machine Learning Expert,
  - an Artificial Intelligence Expert, or
  - a Neurophysiologist / Brain Researcher.
- I am
  - a practitioner at a company that sells a product based on machine learning,
  - trained as a computer scientist, and
  - read avidly and broadly (ACM, AAAS, IEEE).
- Your mileage may vary.

What are we trying to do with ML? One simple thing:

What People Do

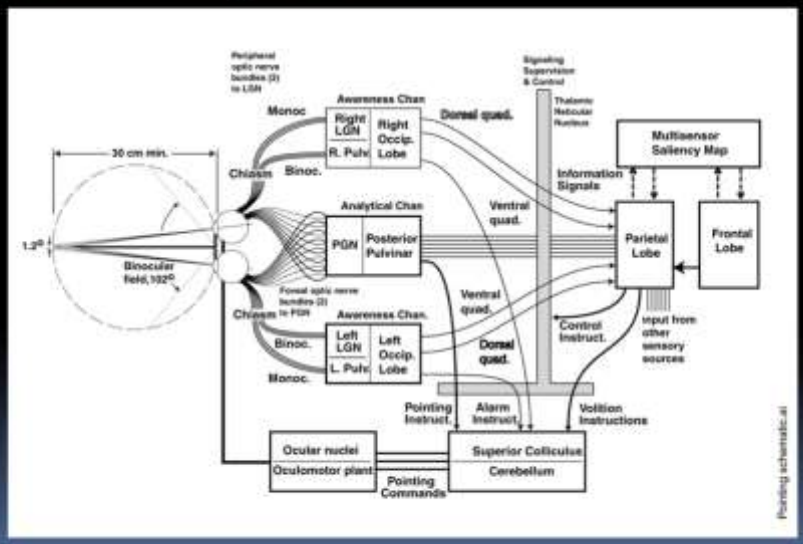

Francis Crick (of DNA structure fame) laid out a research program in his 1995 book *The Astonishing Hypothesis*: let's find out how a human being takes visual information and reduces it to language. The visual system of the brain is very complex, but more mysterious is how it links to our ability to communicate in small, highly compact, and very information rich symbols (words).

Expert systems and "big data" analytics make predictions on data that are already reduced to language, that is, words and numbers.

The data come from various sources: web logs, cookie trails, invoices and order records, historical medical records from health trends in a patient, phone logs, credit card purchases, and so on.

Very few inputs come from speech to text, signal, or image analysis because these data streams are complicated and hard to match to patterns and hence reduce to the language of data.

Active research is taking place, and the use of image analysis applications is improving the ability of machine systems to pick out and recognize details of interest from the noise of the natural world.

ML Techniques – Categorizers

- K Nearest Neighbors (KNN)
- Decision Trees
- Neural Networks / Back Propagation
- Support Vector Machines (sheep vs goats)
- Statistical Classification
  - Linear, Logistic Regression
  - Probabilistic
- Clustering (unsupervised) eg K-Means
- et cetera

ML Techniques - Predictors

- Event dependencies
  - Bayes
  - Frequentist
- Event-driven
  - Shopping cart, web log
  - Typically algorithmic
- et cetera

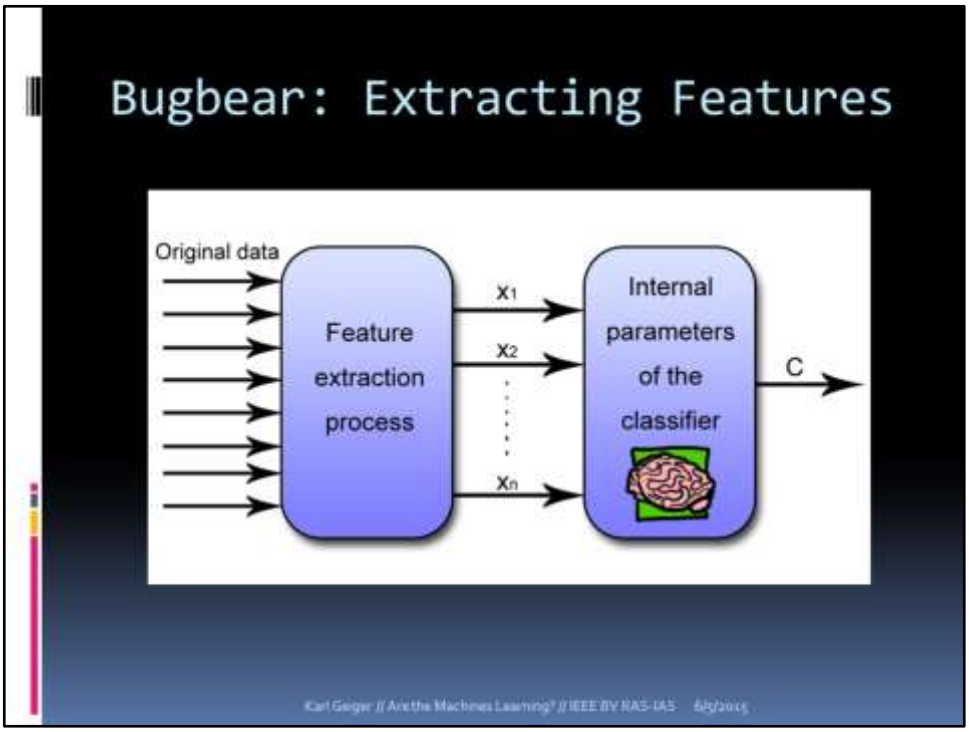

Input data to ML systems is usually already reduced in some way. All the previously listed techniques assume we have already extracted features and reduced them to data (language).

Feature extraction techniques are highly dependent on the choice of which features are selected. This normally has been left up to the human designer of the FE method.

Stanford's Unsupervised Feature Learning / Deep Learning (UFLDL) group is working very hard on developing techniques for automated recognition of patterns in images and other raw data streams.

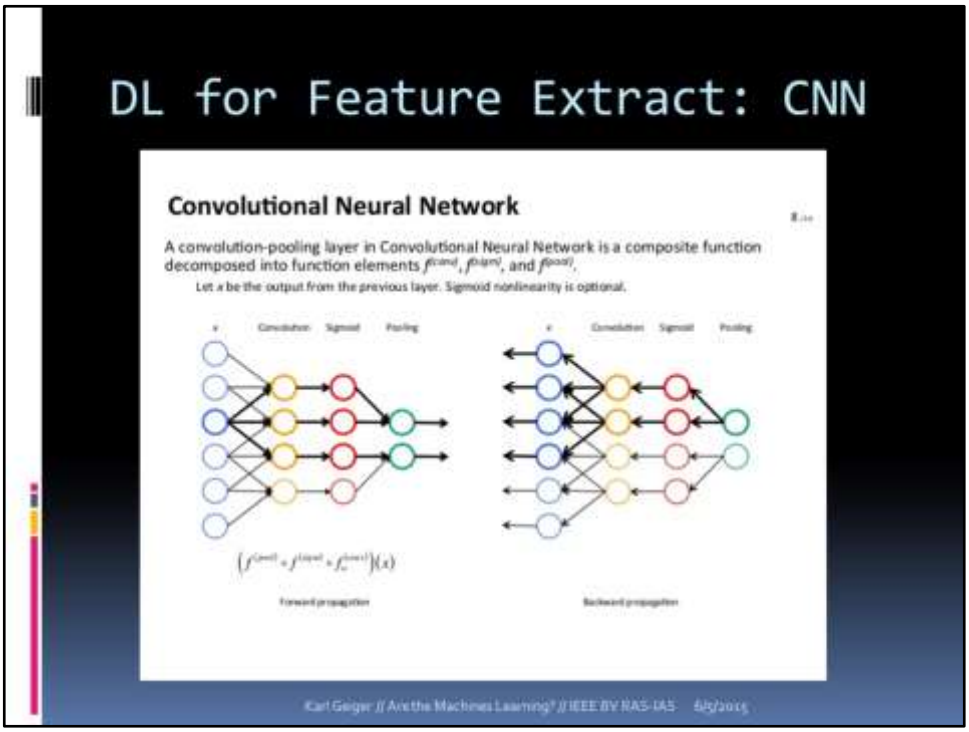

UFLDL techniques include adding additional layers of recognition neurons that assist in sharpening the detection capabilities of artificial neural networks (ANNs). The additional layers act somewhat like the modular structure of human brains. The human visual system passes image information through multiple modules for edge detection, motion detection, etc., before the information ever reaches the conscious level.

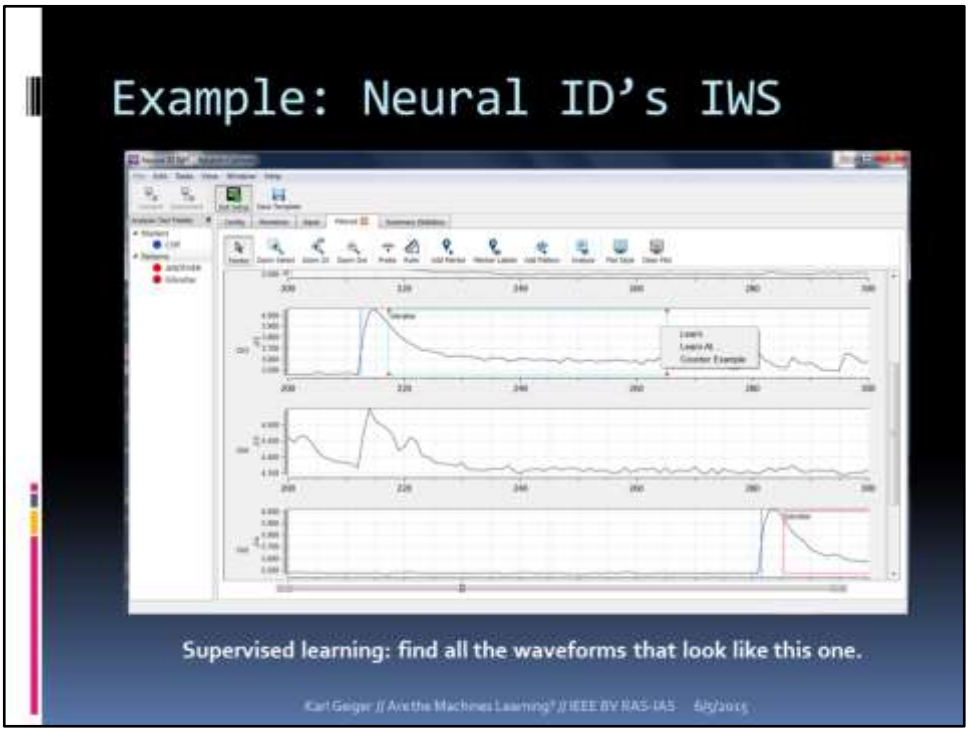

Neural ID's IWS (Intelligent Waveform Service) does one simple thing: it finds amplitude (shape-based) patterns in signal data and turns them into langauge.

To do so, a human shows it what to find. IWS builds a n-dimension model of a waveform shape, then scans the input stream for the pattern. Found patterns are labelled with a name (reduced to language) and noted where they are found (data point == language).

The system is sufficiently powerful enough to permit patterns learned on one channel to be reapplied to other channels of data … or other files. The "tolerance" setting permits the user to control how closely [exact match … very loose match] the input must compare to the model.

Pre-computation of the learning vector set makes the machine very fast at recognizing pattners. Unlike other systems, it's not doing heavy math (Fourier Transforms, polynomial curve fitting) on the input data stream. It's just comparing its pattern model to the incoming signal stream.

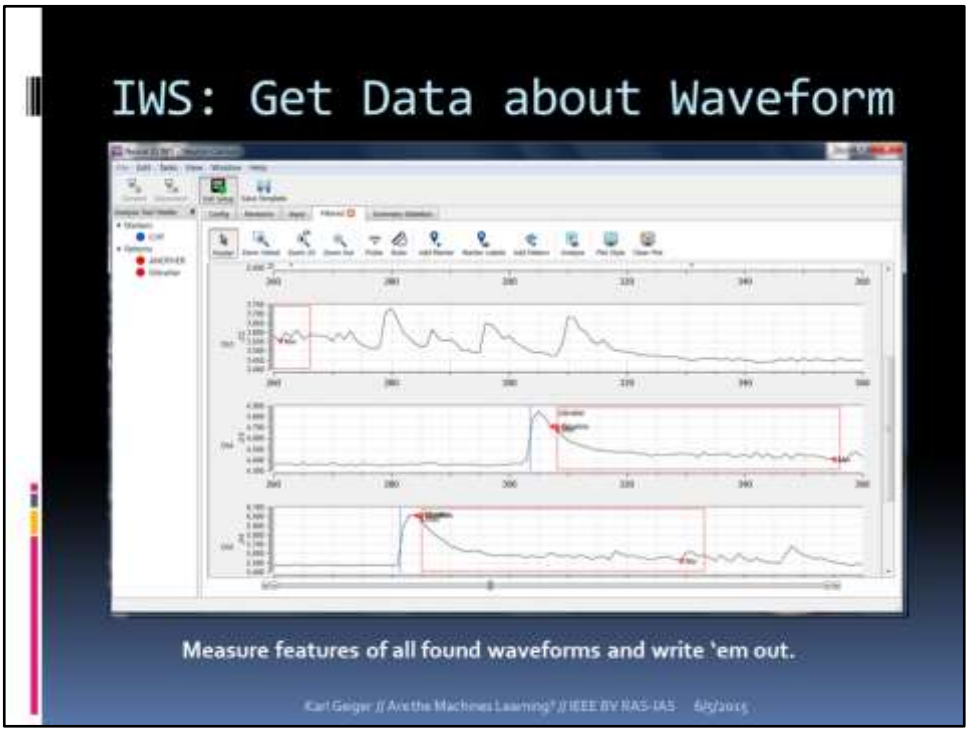

In a broader context, IWS acts as a feature-extraction layer, directed by human-supervised learning. If a human can pick out the interesting bits of the waveform, the machine can "see" what the human sees.

Once the machine has detected the waveform, it then applies mathematical methods to measure waveform features such as pattern width, maxima, minima, curtosis, etc. IWS defines these as "metrology" methods.

Ultimately the data are reduced to named and measured data points suitable for further downstream processing by classic ML systems. We think of its as taking "deep data" (signals) and turning them into "big data".

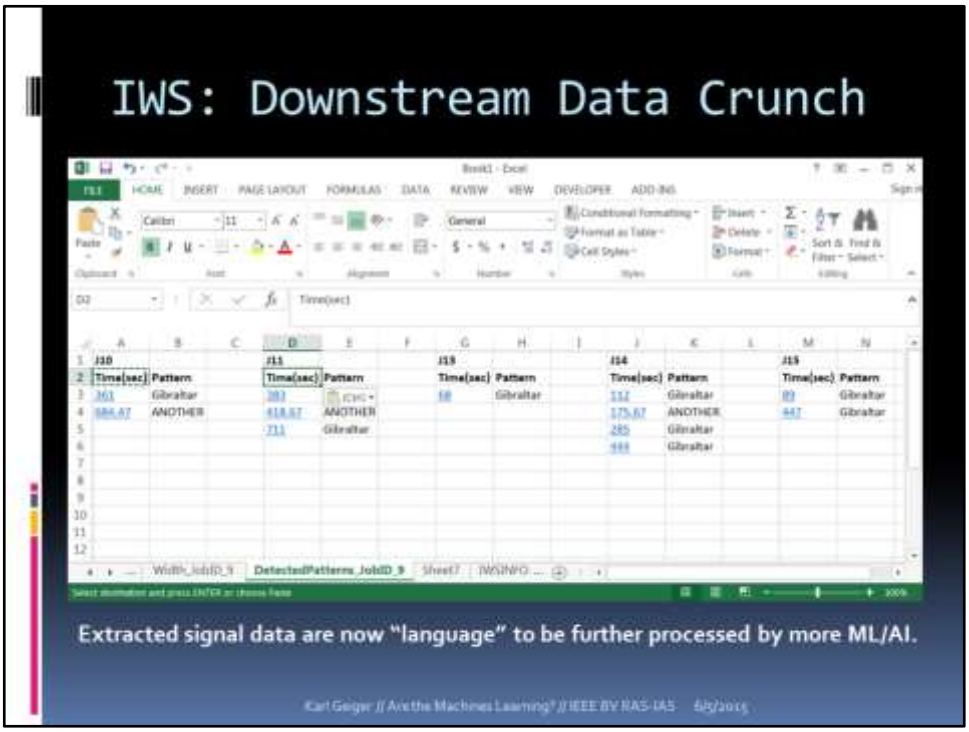

IWS comes with simple SQL tools to extract and process extracted data.

The tools run in Excel and plug into other systems such as IDBS eWorkbook and TIBCO Spotfire.

Key feature is that the downstream datapoints can be extracted with URLs that permit the user to link back to the underlying data signal and view it again in IWS. This allows the "human layer" to inspect output at any time.

Point solutions will reign for the forseeable future. In very narrow domains, we'll see things like self-driving trucks, face recognition for security and fraud detection, and so on.

Important to think about the business context … self-driving trucks turns trucks into trains. Expensive capital equipment can then be used more than 11 hours/day. Drivers are limited by fatigue rules to prevent crashes. If the driver can be on board sleeping all night while the truck trundles along the a dark interstate highway, it will cut delivery times in half AND employ the machinery more cost-effectively.

More point solutions …

"Saleh and Elgammal begin with a database of images of more than 80,000 paintings by more than a 1,000 artists spanning 15 centuries. These paintings cover 27 different styles, each with more than 1,500 examples. The researchers also classify the works by genre, such as interior, cityscape, landscape, and so on.

The researchers then test the algorithm on a set of paintings it has not yet seen. And the results are impressive. Their new approach can accurately identify the artist in over 60 percent of the paintings it sees and identify the style in 45 percent of them."

This is a perfect example of what machines are good at: taking many dimensions and making detailed matches across a broad swath.

There's always sales. Systems to review customer trends and buy patterns help smooth the supply chains for retailers.

(Tell Target pregnant daughter story here.)

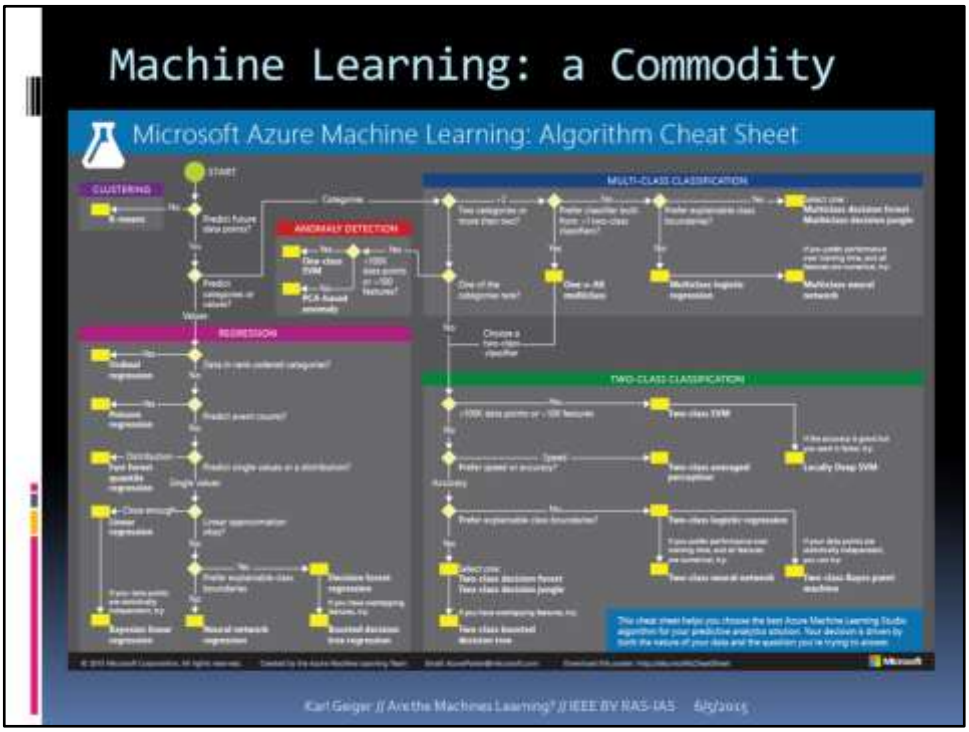

ML techniques are commoditized to the point of being APIs available to any programmer.

The question is will the people apply them “correctly” and “appropriately” or have we a new way to make data smog?

It's been 16 years since the Y2K debacle in 2000. What do we have? Really good phones and smart thermostats. I think the leap over the 16 years from 2029 to 2045 will also be gentler.

"With face recognition, it's been clear for a while now that it can be solved. Beyond faces, you can also talk about other categories of objects: 'there's a cup in the scene'; 'there's a dog in the scene'. But it's still a hard problem to talk about many kinds of different objects in the same scene and how they relate to each other, or how a person or a robot would interact with that scene. There are many, many hard problems that are far from solved."

Michael Jordan, EECS @ UCB

"The fears of runaway AI systems either conquering humans or making them irrelevant are not even remotely well-grounded. Misled by suitcase words, people are making category errors in fungibility of capabilities. These category errors are comparable to seeing more efficient internal combustion engines appearing and jumping to the conclusion that warp drives are just around the corner."

"Artificial intelligence is a tool, not a threat."

Rodney Brooks

Summary

- ML and AI are essentially data processing techniques used to extract patterns and relationships that are otherwise cumbersome and time consuming to find.
- ML relies on methods that operate on input that is already reduced to language, hence, manipulatable with computers.
- Feature extraction (noisy world → data) is still a research problem. We're getting better at it but have a very long way to go.

Summary

- Humans must interpret output to create "meaning", that is action that changes the external world's state. Its whole-system cybernetics.
- The bigger threat is human misapplication of data mining and machine learning. A false-positive determination that you're a bad credit risk or terrorist means you're going to have a very, very bad day.
- Point solutions will help make the world around us accommodating, but won't take over the planet.
- Trust Rodney Brooks: machines won't be our "robot overlords" for quite some time, if ever.

Selected References

- Steven Hawking's AI Alarm http://www.bbc.com/news/technology-30290540
- Elon Musk AI Alarm http://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat

Selected References

- "Our Fear of A.I.", *MIT Tech Review* http://www.technologyreview.com/review/534871/our-fear-of-artificial-intelligence/
- "Artificial intelligence is a tool, not a threat." http://www.rethinkrobotics.com/artificial-intelligence-tool-threat/
- "Mistaking Performance For Competence Misleads Estimates Of AI's 21st Century Promise And Danger" http://edge.org/response-detail/26057

Selected References

- "Machine-Learning Maestro Michael Jordan on the Delusions of Big Data and Other Huge Engineering Efforts", *IEEE Spectrum* http://spectrum.ieee.org/robotics/artificial-intelligence/machinelearning-maestro-michael-jordan-on-the-delusions-of-big-data-and-other-huge-engineering-efforts
- Francis Crick, *The Astonishing Hypothesis: The Scientific Search for the Soul* http://www.amazon.com/Astonishing-Hypothesis-Scientific-Search-Soul/dp/0684801582

Selected References

- "How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did", *Forbes* 2012 http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/
- "16 Analytic Disciplines Compared to Data Science" http://www.datasciencecentral.com/profiles/blogs/17-analytic-disciplines-compared

Selected References

- Harrington, *Machine Learning in Action* http://www.amazon.com/Machine-Learning-Action-Peter-Harrington/dp/1617290181
- Battiti, Brunato, *The LION Way: Learning plus Intelligent Optimization* http://www.amazon.com/The-LION-Way-Intelligent-Optimization/dp/1496034023
- Frankel, "Data Scientists Don't Scale", *Harvard Business Review* https://hbr.org/2015/05/data-scientists-dont-scale

Selected References

- Prediction IO Open Source Engine https://prediction.io/
- Deep Learning Tutorial – Feature Extraction http://ufldl.stanford.edu/tutorial/supervised/FeatureExtractionUsingConvolution/

Contact Information

- Karl Geiger
  - karl@neuralid.com
  - kgeiger@computer.org
- Neural ID's Intelligent Waveform Service
  - http://www.neuralid.com/technology.html
  - https://www.youtube.com/watch?v=NLvWb_-hqjs