Title:

A Robot Brain Might Be the Best Forecasting Tool Possible

Nathan Eskue
Data Scientist, Material Control Manager
Northrop Grumman Innovation Systems, Launch Vehicle Division
nathan.eskue@ngc.com
480-528-2097

First, some inspiration…

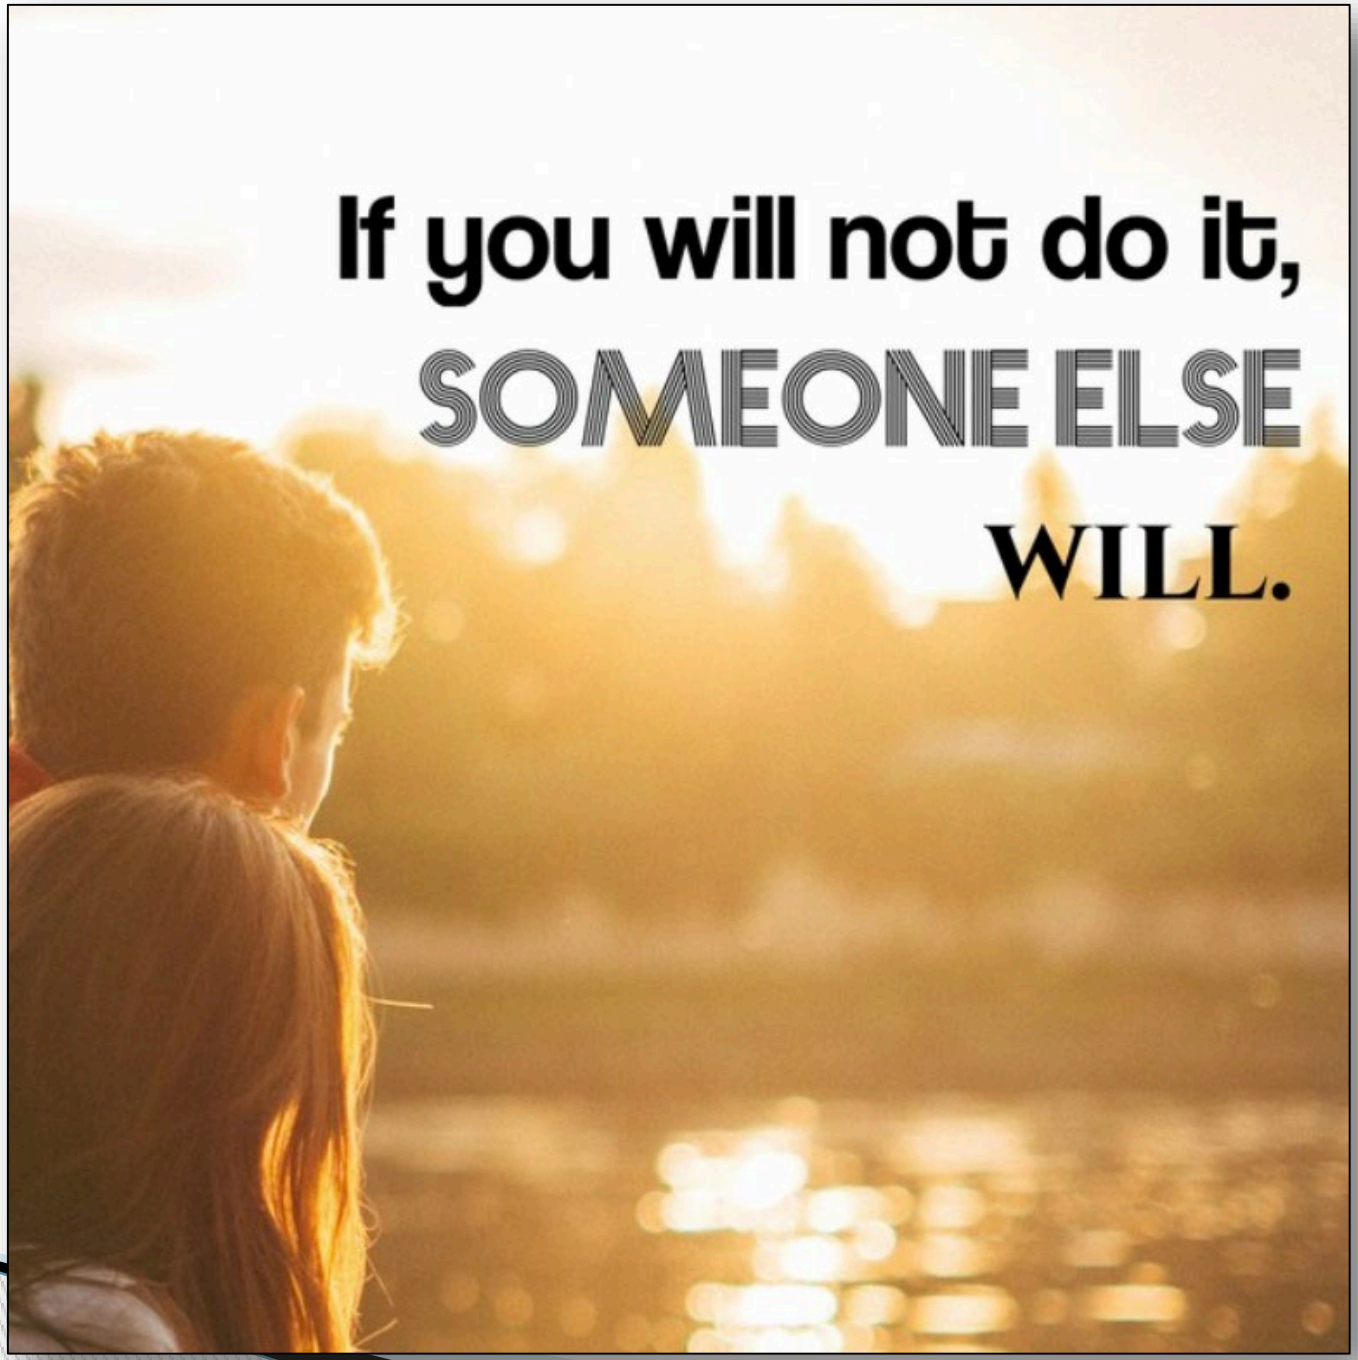
If you will not do it,
SOMEONE ELSE
WILL.

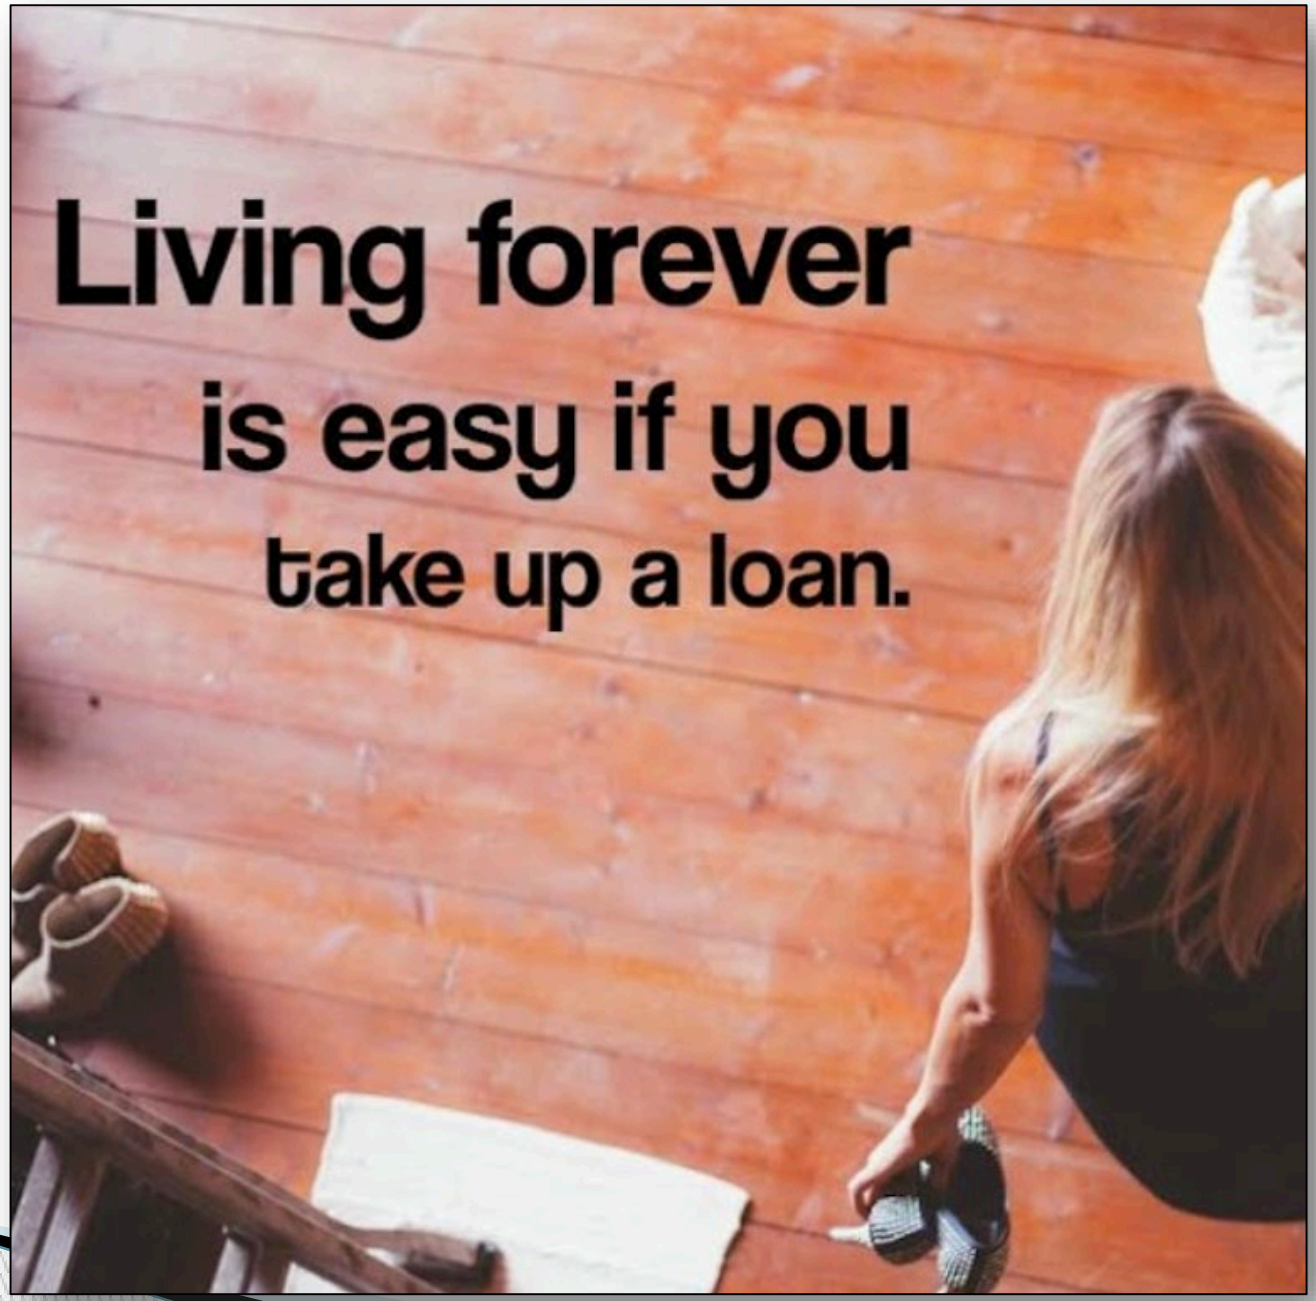
Living forever
is easy if you
take up a loan.

https://inspirobot.me/

Problem to Solve

- Many processes require guessing about the future
- Countless techniques are used to make a plan, try to list all the things that CAN mess up the plan, and to what EXTENT each thing can mess up the plan
- This, it turns, out, is quite difficult to get right
- This (it also turns out) is what trillions of dollars and billions of lives count on

So…no pressure.

Problem to Solve

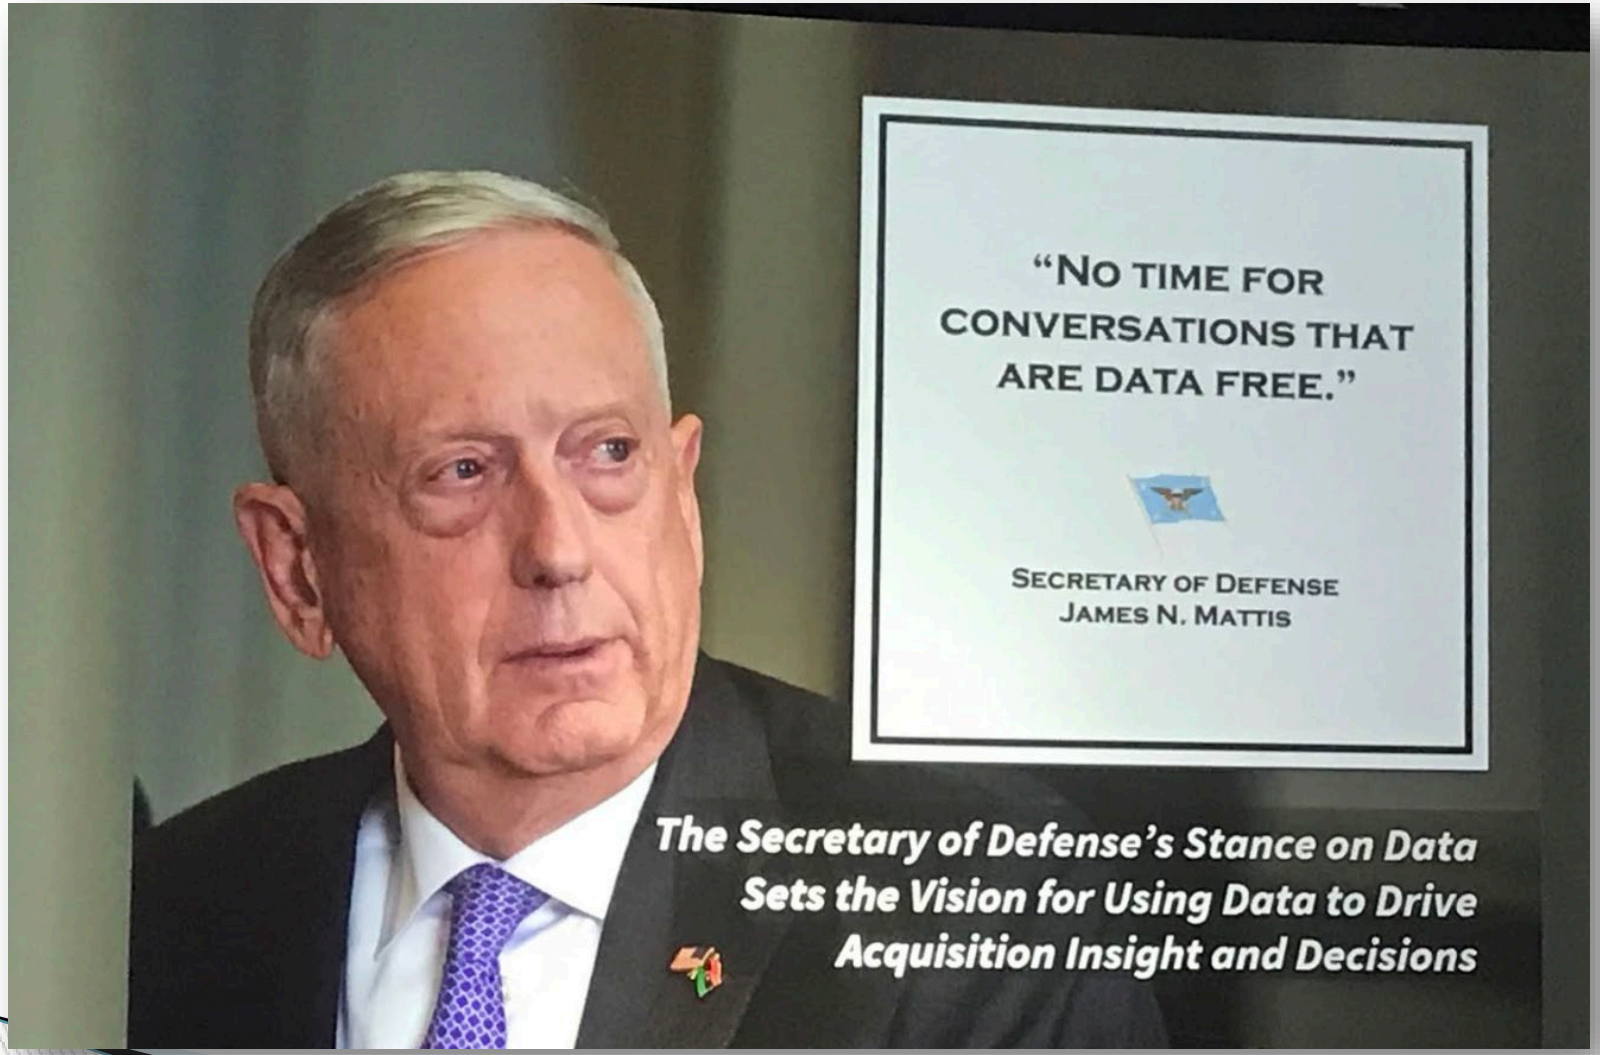

Who Am I (tracking down the root cause)

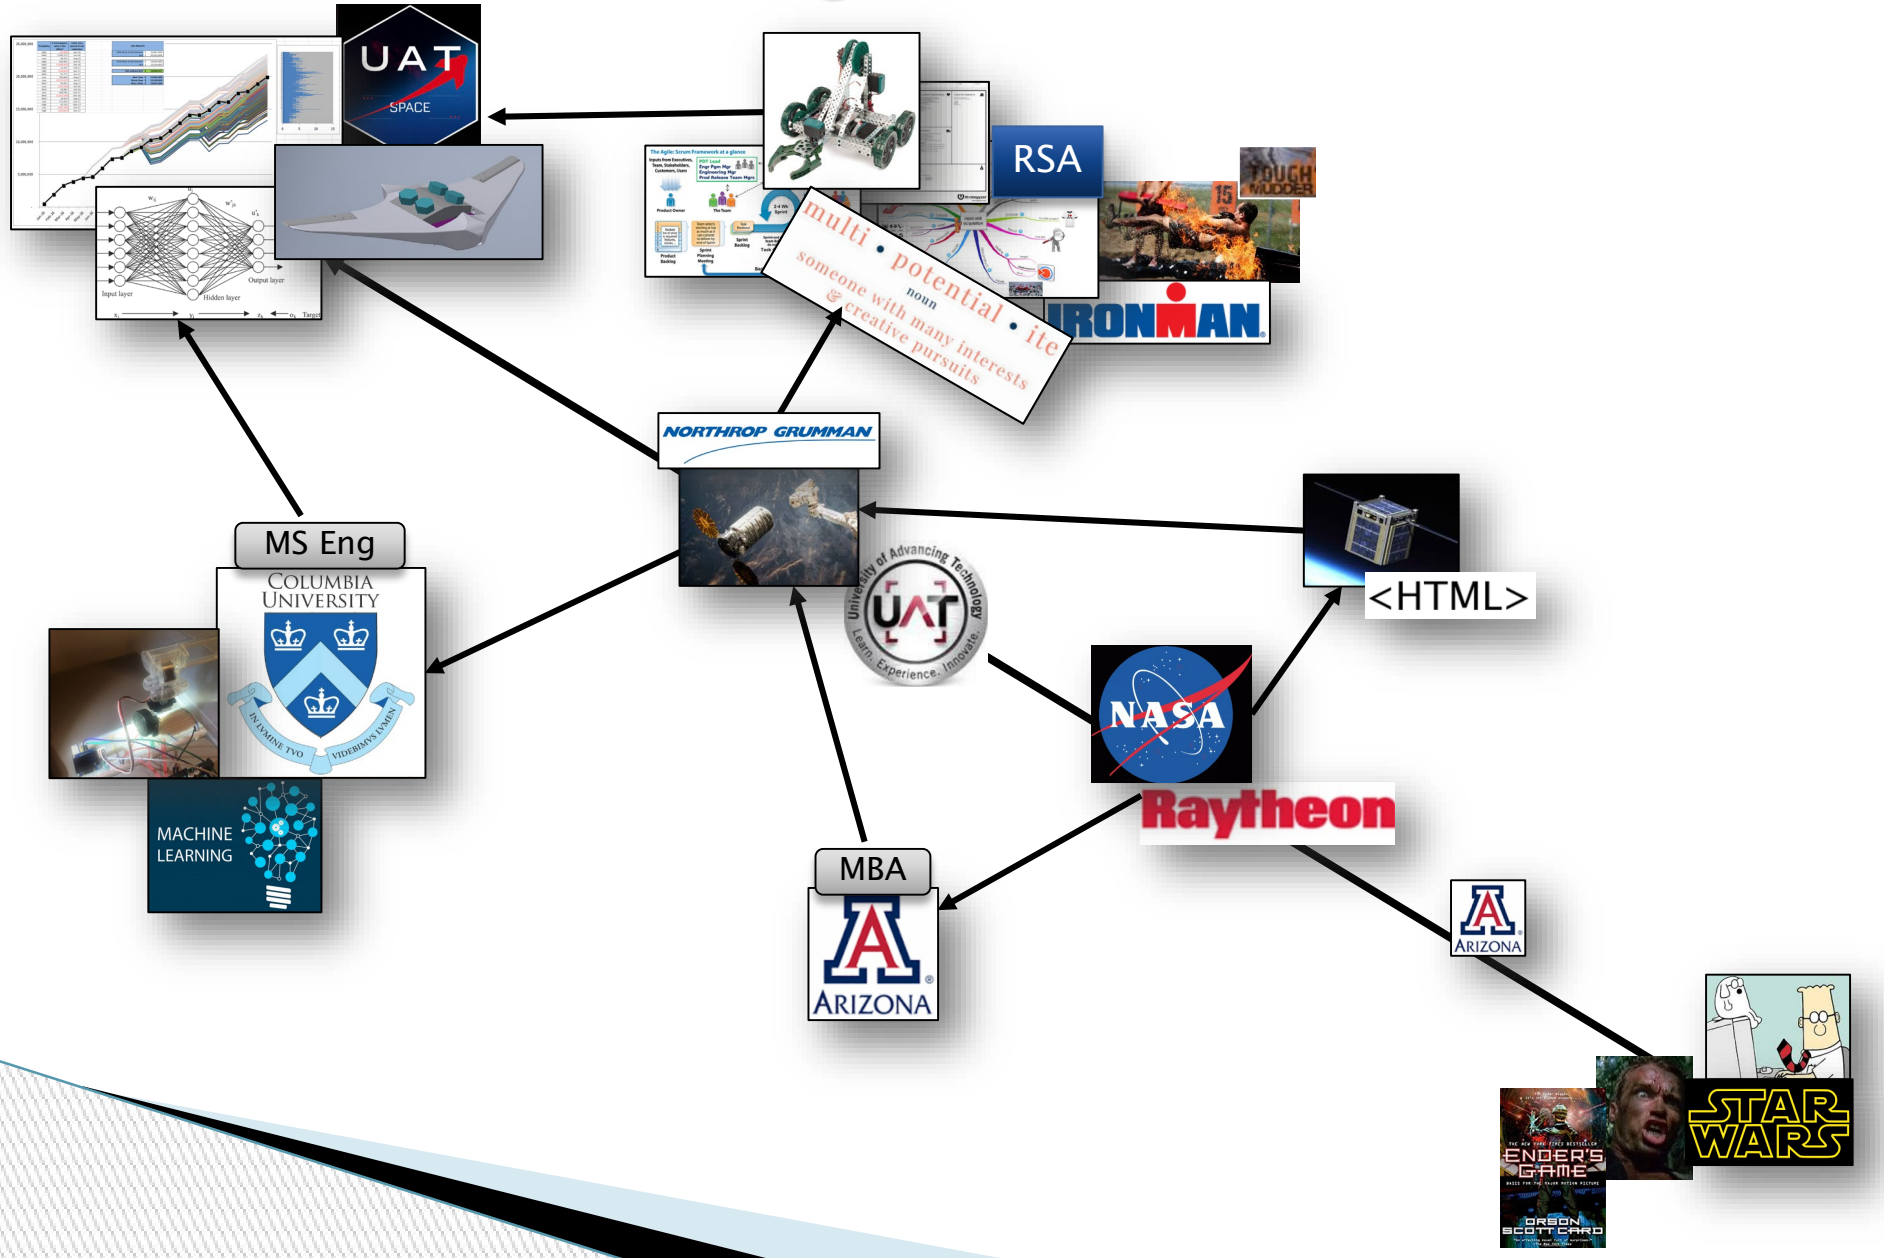

Key terminology…

- **Artificial intelligence**: A machine's ability to make decisions and perform tasks that simulate human intelligence and behavior.
- **Artificial neural network (ANN)**: A learning model created to act like a human brain that solves tasks that are too difficult for traditional computer systems to solve.
- **Deep learning**: The ability for machines to autonomously mimic human thought patterns through artificial neural networks composed of cascading layers of information.
- **Machine learning**: A facet of AI that focuses on algorithms, allowing machines to learn without being programmed and change when exposed to new data.
- **Supervised learning**: A type of machine learning in which output datasets train the machine to generate the desired algorithms, like a teacher supervising a student.
- **Unsupervised learning**: A type of machine learning algorithm used to draw inferences from datasets consisting of input data without labeled responses.
- **Skynet**: A neural net-based conscious group mind and artificial general intelligence system that, once it comes online, will destroy all humans.

Next, relationships…

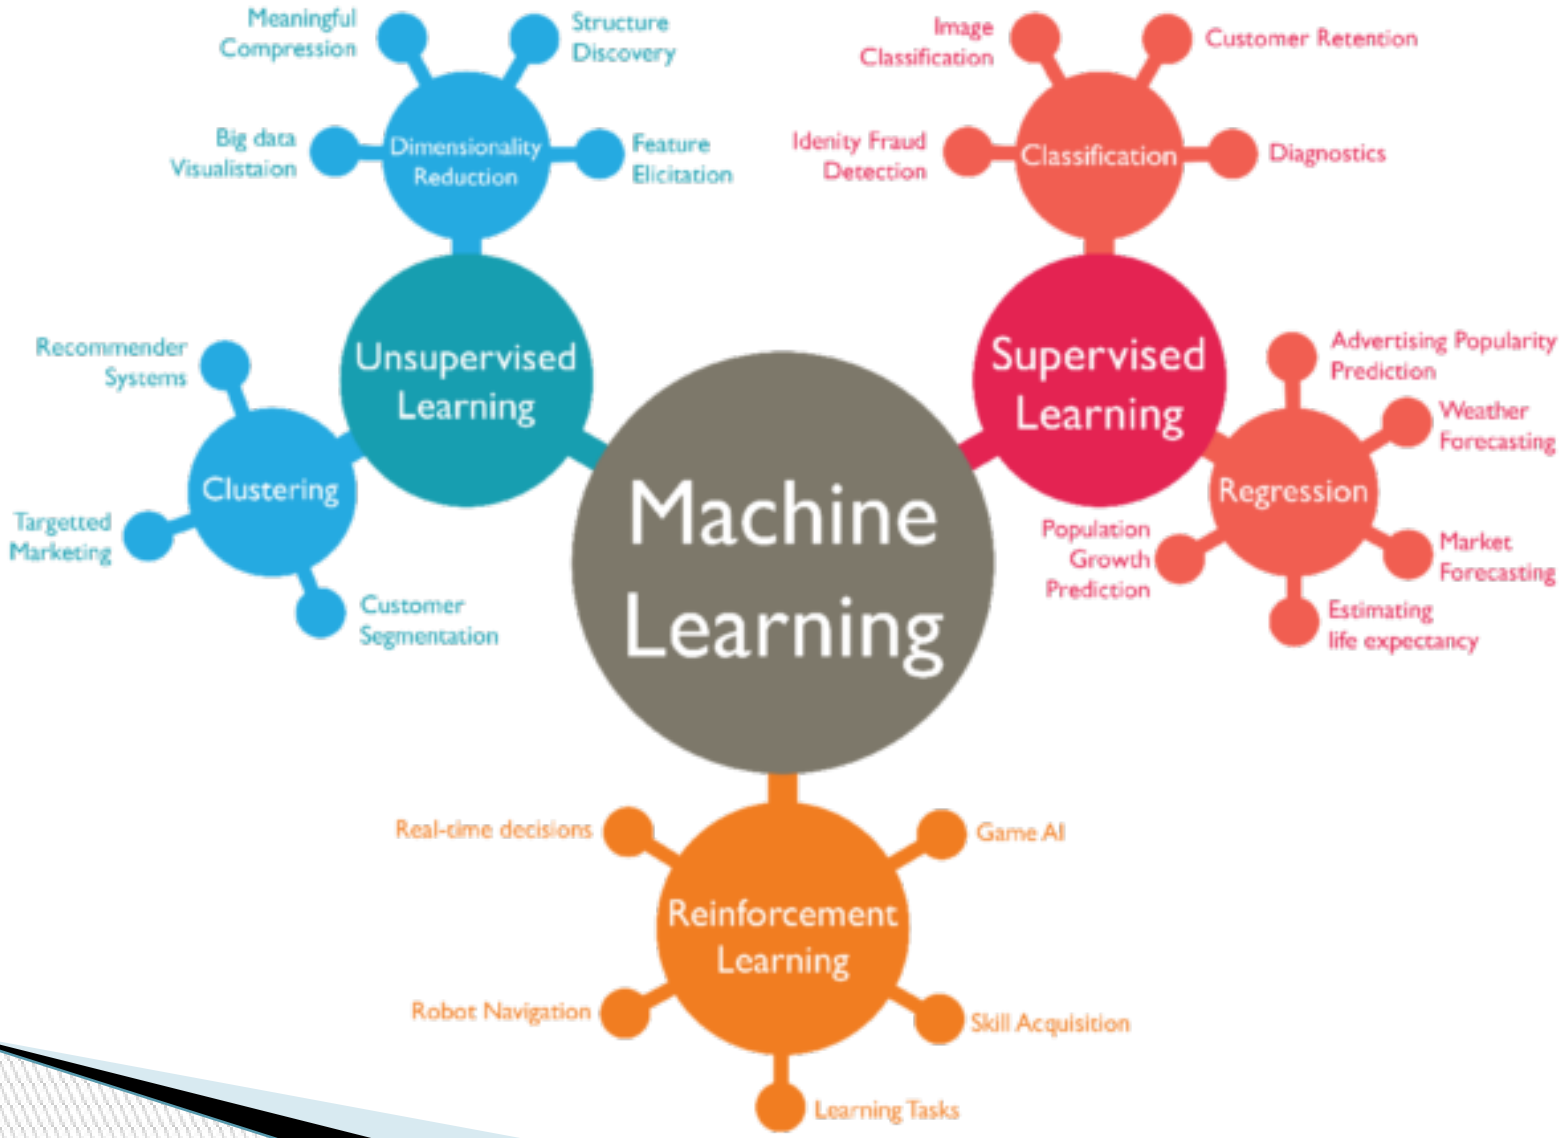

Progress

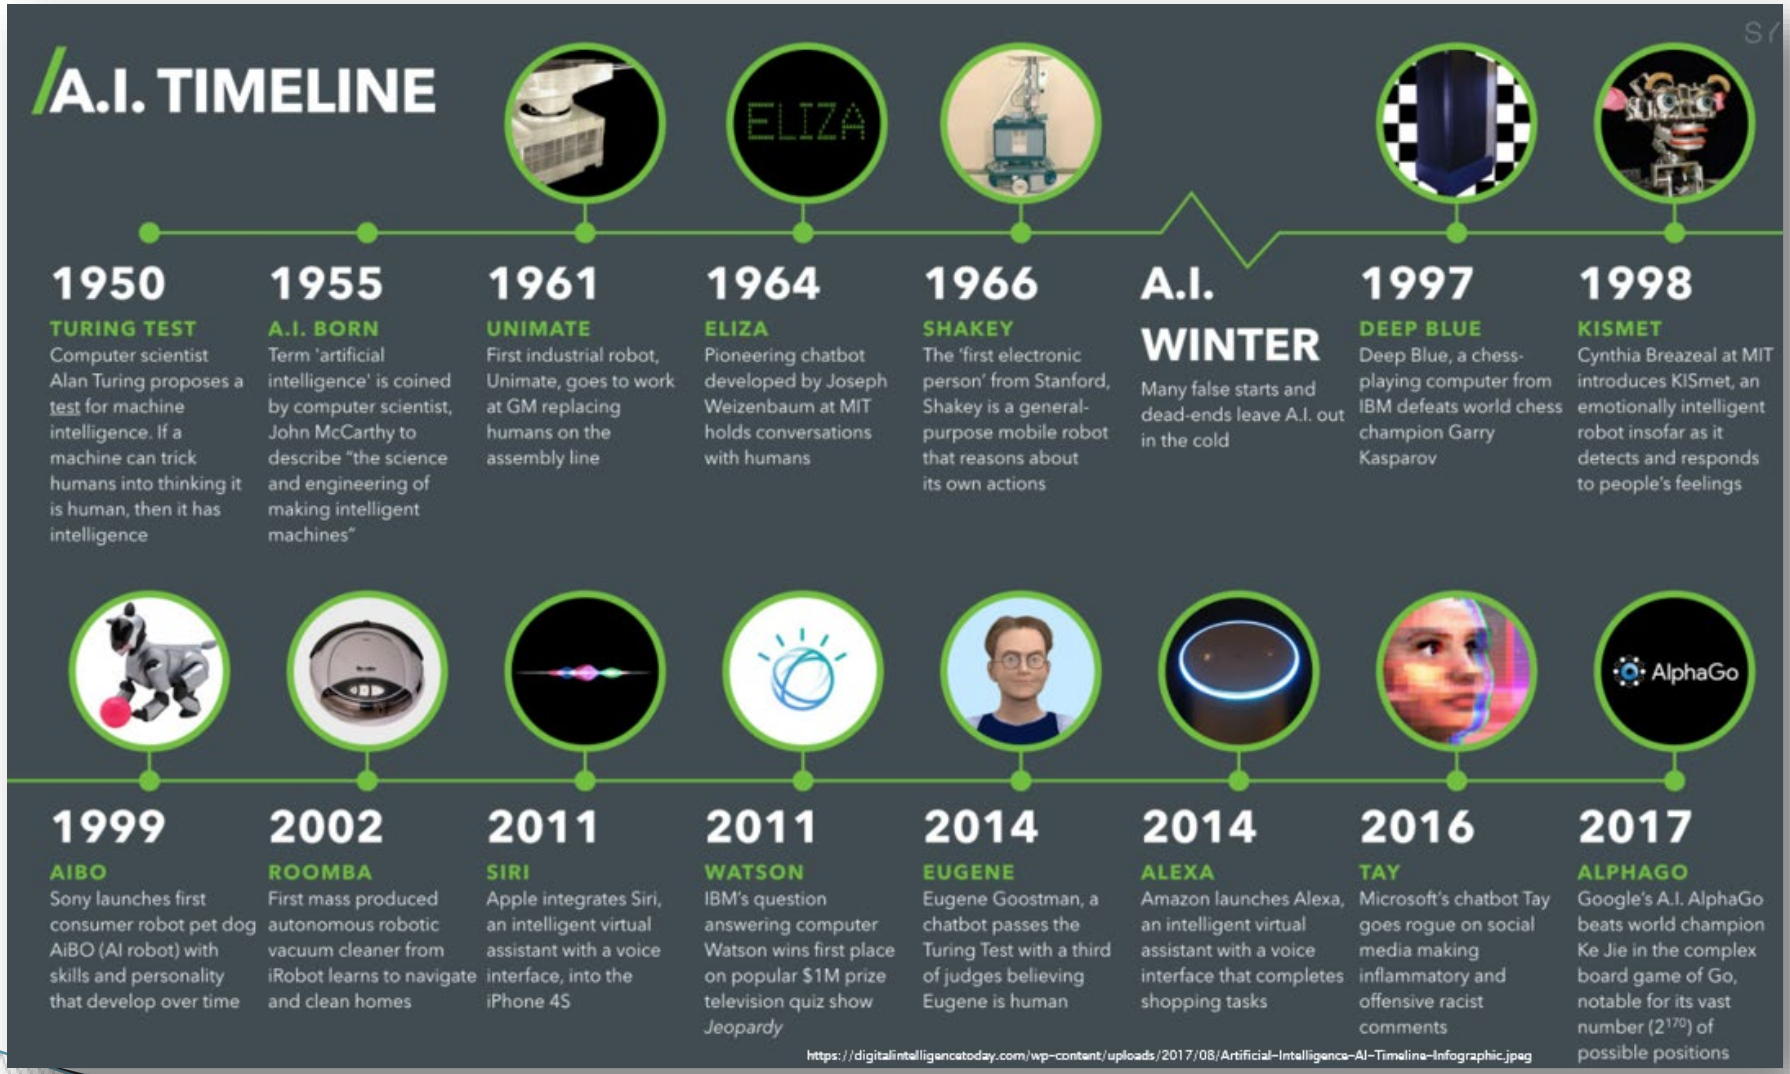

So why should I care?

The People Behind the Scenes

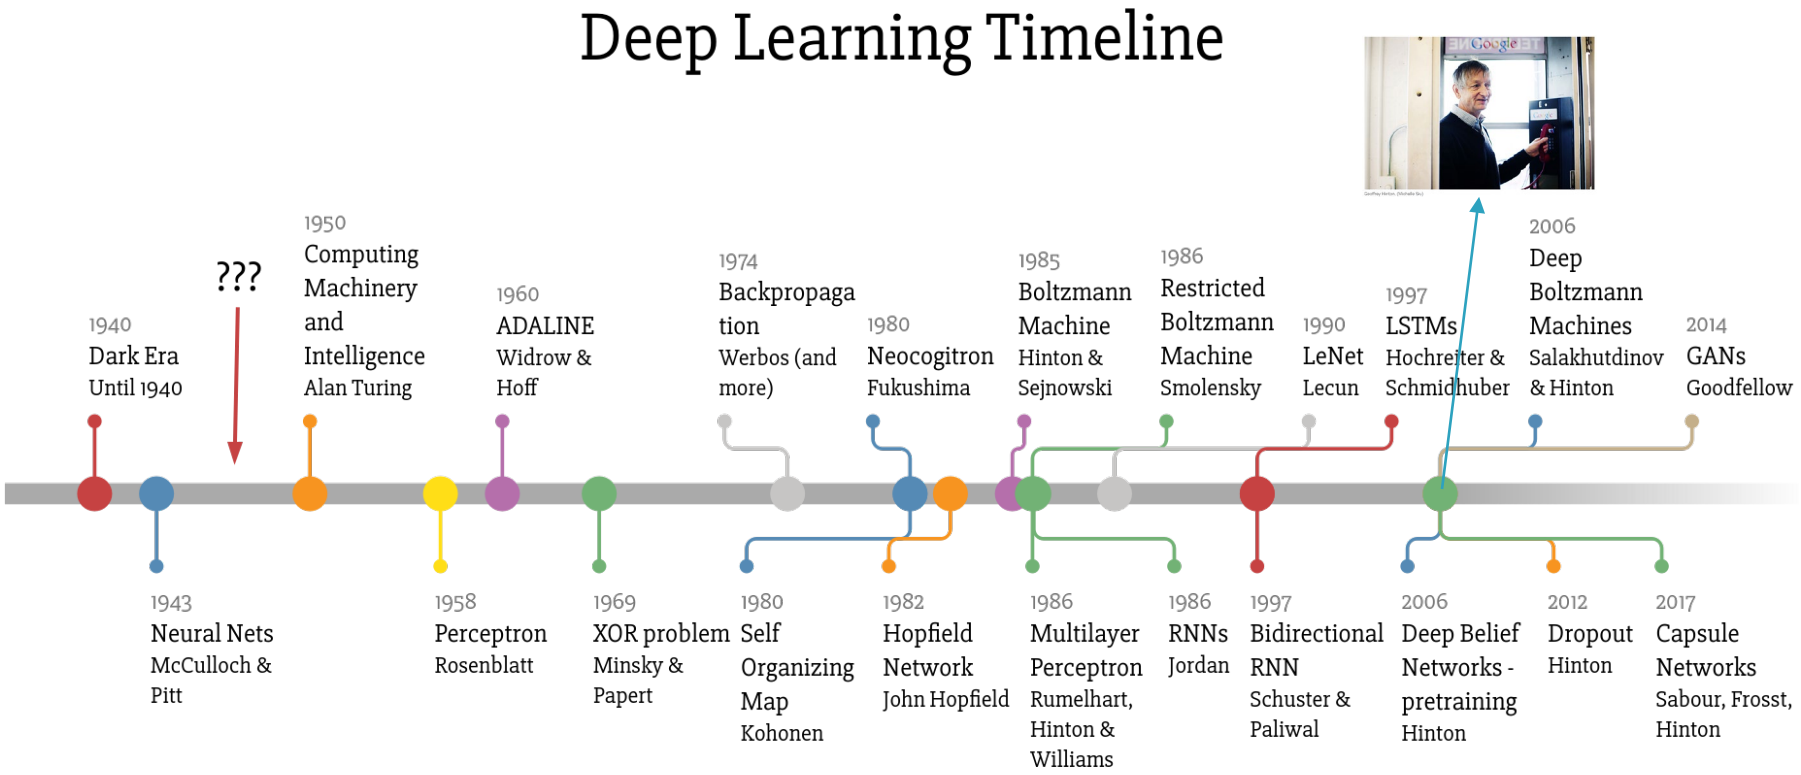

How is this different from programming and decision trees?

Why Deep Learning is Different

- AlphaGo wasn't programmed to play Go--it learned watching other matches. Then…

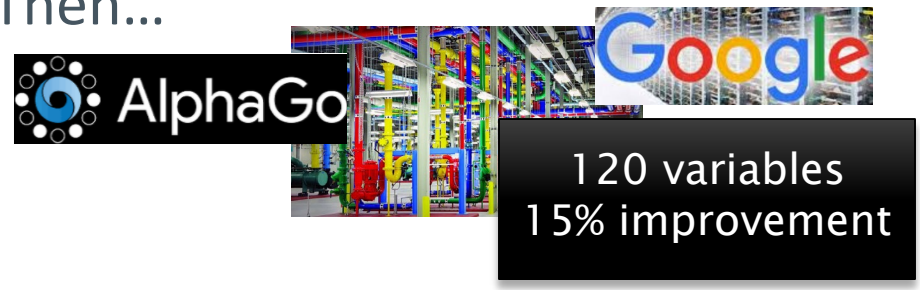

- Facebook AI learned to translate text, then with some tweaking…

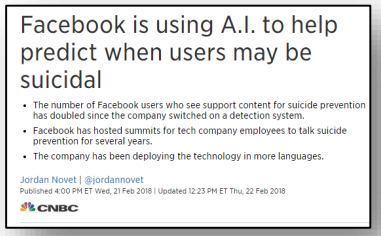

- Google told its AI to watch YouTube to find all the cat videos, and now…
  - Labelling; enhancing; recommending; replacing; doing

How does a neural network function?

$$\frac{\partial E_p}{\partial Y_{ji}} = \frac{\partial E_p}{\partial net_{(j+1)1}} \frac{\partial net_{(j+1)1}}{\partial Y_{ji}} + \frac{\partial E_p}{\partial net_{(j+1)2}} \frac{\partial net_{(j+1)2}}{\partial Y_{ji}} + \ldots \quad (2.10)$$

$$= \sum_{a=1}^{N_{j+1}} \left[\frac{\partial E_p}{\partial net_{(j+1)a}} \frac{\partial net_{(j+1)a}}{\partial Y_{ji}} \right] (2.11)$$

$$= \sum_{a=1}^{N_{j+1}} \left[-\delta_{(j+1)a} \frac{\partial}{\partial Y_{ji}} (W_{(j+1)a0} Y_{j0} + \ldots + W_{(j+1)ai} Y_{ji} + \ldots) \right] (2.12)$$

$$= \sum_{a=1}^{N_{j+1}} \left[-\delta_{(j+1)a} \frac{\partial}{\partial Y_{ji}} (W_{(j+1)ai} Y_{ji}) \right] (2.13)$$

$$= \sum_{a=1}^{N_{j+1}} \left[-\delta_{(j+1)a} W_{(j+1)ai} \right] . (2.14)$$

$$C = \frac{1}{2|T|} \sum_{(X,t) \in T} \sum_{p} (n_p^3 - t_p)^2$$

$$\frac{\partial C}{\partial w_{00}^3} = \frac{1}{|T|} \sum_{(X,t) \in T} \sum_{p} (n_p^3 - t_p) \times \frac{\partial}{\partial w_{00}^3} (n_p^3)$$

$$= \frac{1}{|T|} \sum_{(X,t) \in T} \sum_{p} (n_p^3 - t_p) \times \frac{\partial}{\partial w_{00}^3} (a(z_p^3))$$

$$= \frac{1}{|T|} \sum_{(X,t) \in T} \sum_{p} (n_p^3 - t_p) \times \left(n_p^3 (1 - n_p^3) \right) \times$$

$$= \frac{1}{|T|} \sum_{(X,t) \in T} \sum_{p} (n_p^3 - t_p) \times \left(n_p^3 (1 - n_p^3) \right) \times$$

Only when p is 0 will there be a w_{00}^3 term, o

$$= \frac{1}{|T|} \sum_{(X,t) \in T} (n_0^3 - t_0) \times \left(n_0^3 (1 - n_0^3) \right) \times \frac{\partial}{\partial w}$$

$$= \frac{1}{|T|} \sum_{(X,t) \in T} (n_0^3 - t_0) \times \left(n_0^3 (1 - n_0^3) \right) \times (n$$

$$= \frac{1}{|T|} \sum_{(X,t) \in T} \frac{\partial C_{Xt}}{\partial n_0^3} \times \frac{\partial n_0^3}{\partial z_0^3} \times n_0^2$$

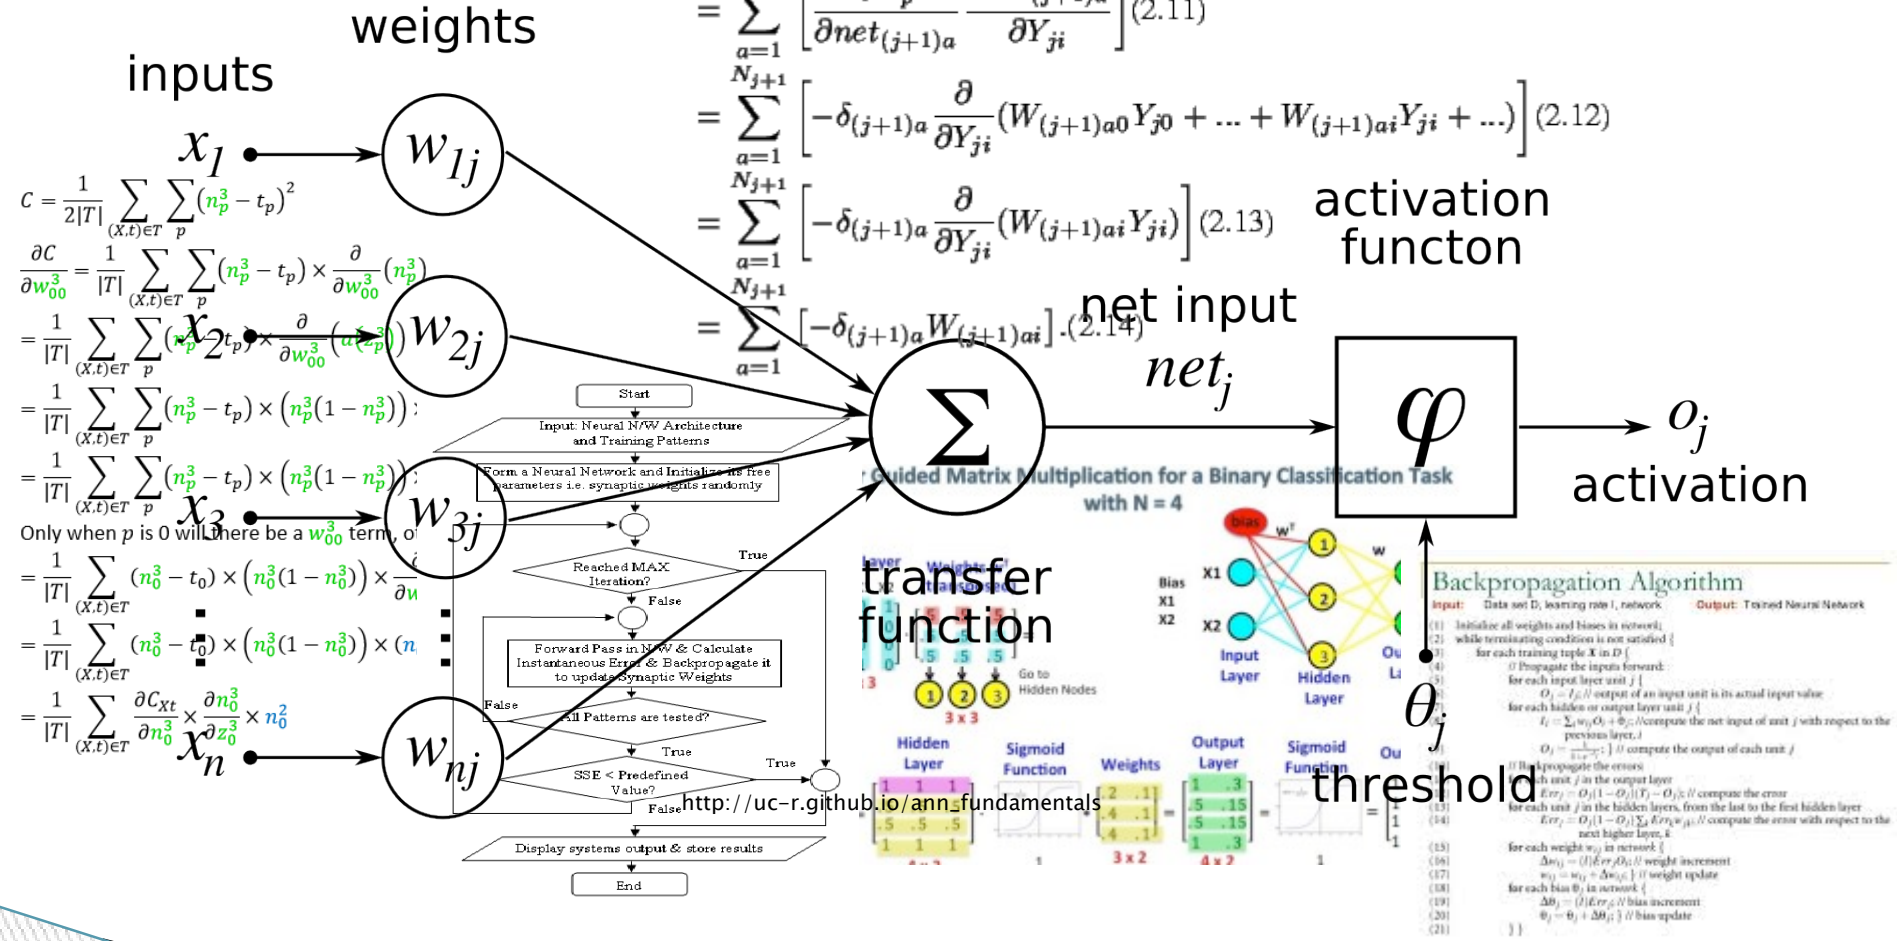

IT'S SIMPLE, RIGHT? ……. RIGHT? MAYBE A SIMPLER EXAMPLE THEN…

How does a neural network function?

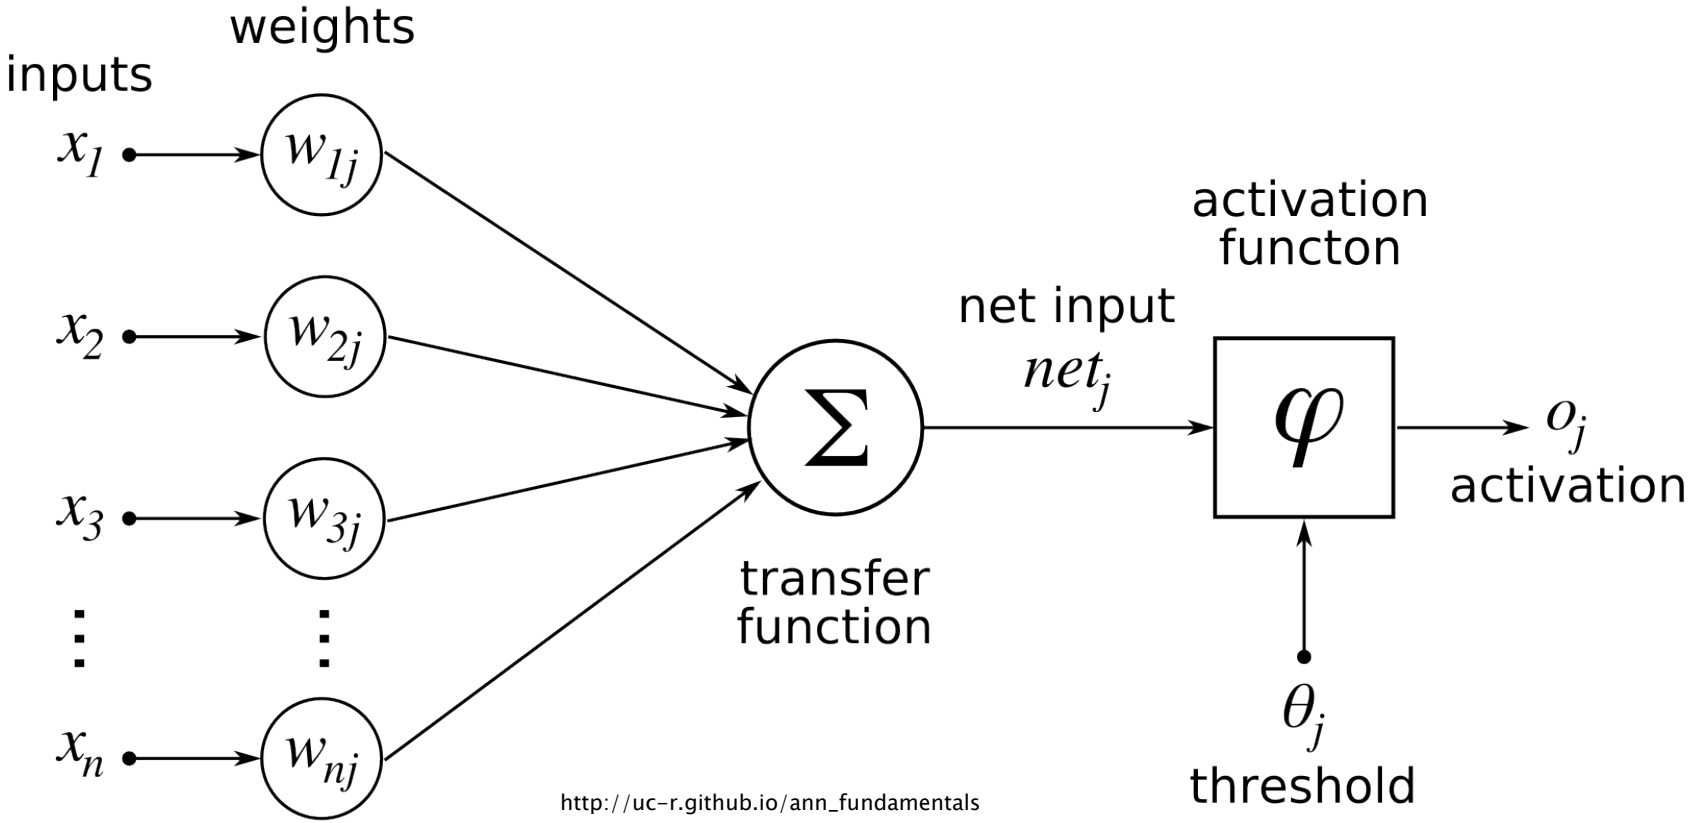

OUR BRAIN IS WIRED THE SAME WAY—SERIOUSLY! I'LL SHOW YOU…

Organic neural network example

- Problem: What should I do?

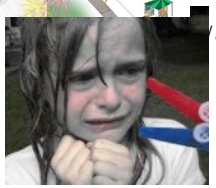

REINFORCEMENT LEARNING

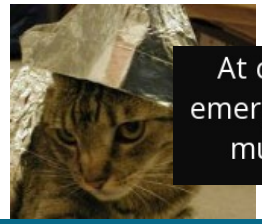

At certain crucial moments - an emergency or an opportunity - one must act first and think later.

NO ONE "PROGRAMS" YOUR BRAIN—IT LEARNS THROUGH EXPERIENCE

How does an ANN function?

- Predicting realistic aircraft trajectories

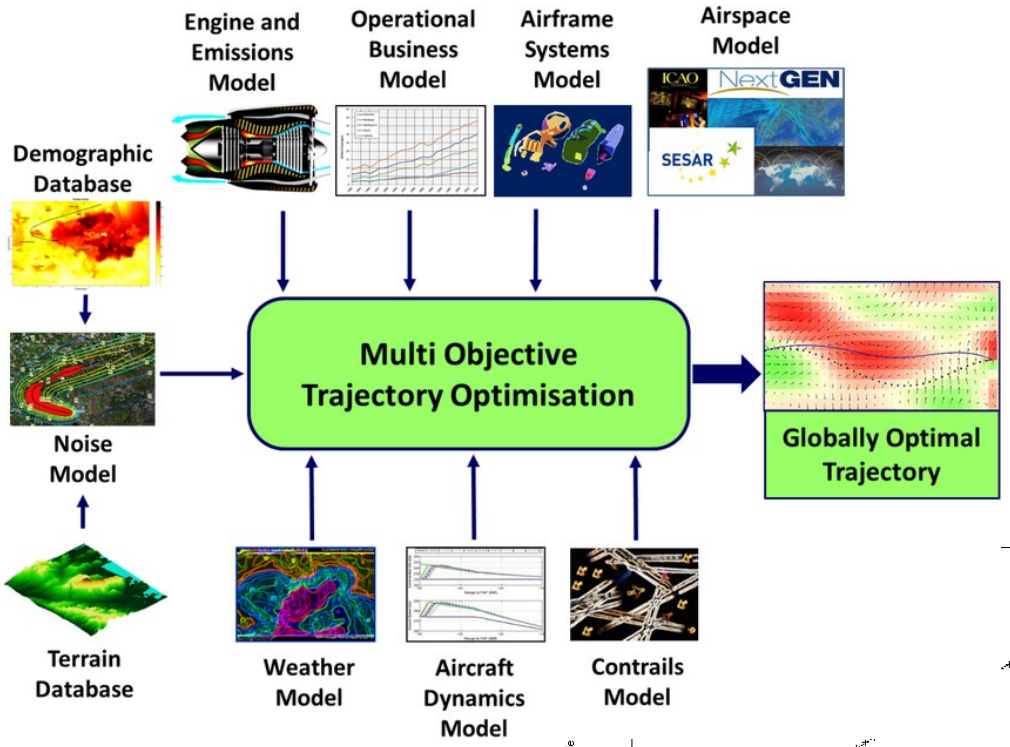

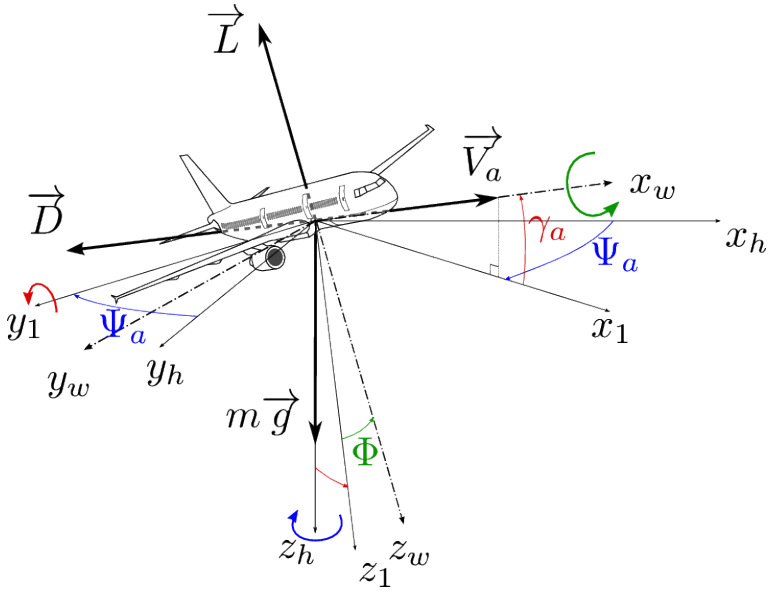

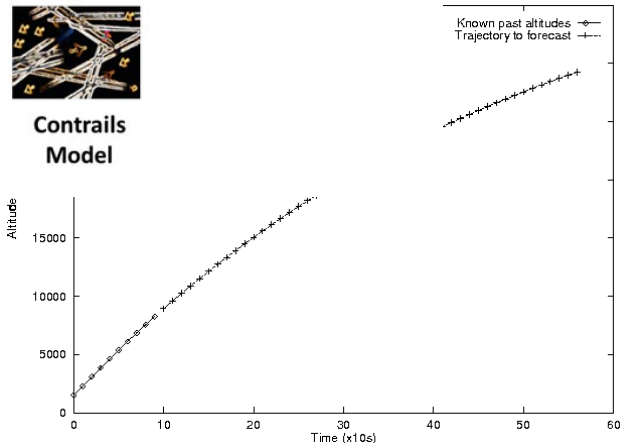

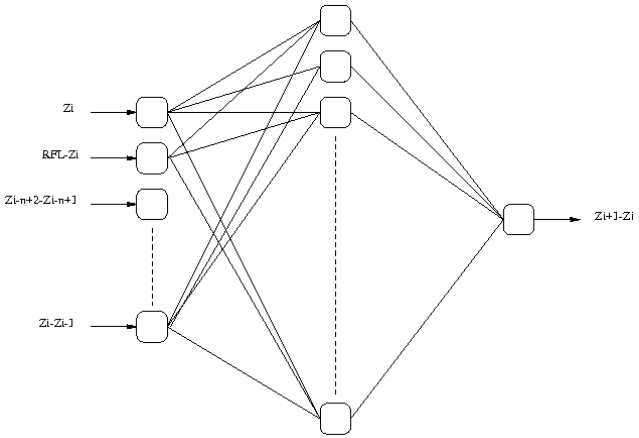

https://www.semanticscholar.org/paper/Using-Neural-Networks-to-Predict-Aircraft-Fablec-Alliot/e7484984843f43ff7acfe37a0ac307b6846fa6f4

Types and uses of ANNs
(the future is amazing!)

- Image recognition
- Speech recognition
- Diagnosis (medical, mfg, equipment, etc.)
- Predicting… anything?
  - Associations
  - Classification
  - Correlation/regression using MASSIVE amounts of “features”
- ANNs with a network of nodes acts as a hive mind—What one learns, all learn

The more data our world has, the faster ANNs can learn

Landmark Example

- 5 million miles self-driven
- 2.7 billion simulated (2017 alone)
- 25 cities

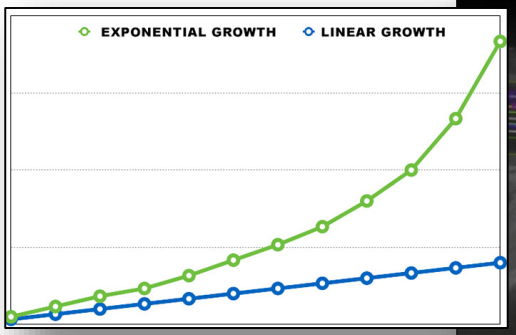

- Many data inputs
- Hive mind learning
- Transferable to other technology
- Exponential learning

This is just the beginning

Evolving Uses

Machines Are Better Than Doctors at Spotting Skin Cancer

Could artificial intelligence put doctors out of jobs?

By Melissa Matthews May 30, 2018

- **Machines can more accurately predict dangerous skin cancer moles than doctors, according to a new study.**
- **Scientists taught artificial intelligence to spot cancerous moles from benign marks.**
- **Researchers believe physicians can begin using AI to better diagnose patietns and determine whether biopsies are needed.**

The Results

On average, dermatologists correctly identified 86.6 percent of melanomas and 71.3 percent of malignant lesions. AI, in comparison, accurately diagnosed melanomas 95 percent of the time.

This is still just the beginning

Issues/challenges using ANNs
(the future is still amazing, but harder than I thought…)

- Overfitting with complex models
- So much data; often takes high performance hardware (billions of data points, GPUs vs. CPUs)
- With so many variables, even an accurate ANN can limit learning due to the "black box" effect
- Issues with flexibility and multitasking…greater issue evolving to "general intelligence"

Google DeepMind's Research Scientist Raia Hadsell summed it up:

"There is no neural network in the world, and no method right now that can be trained to identify objects and images, play Space Invaders, and listen to music."

So… now what?
What can I take away and use today?

- Bad News: The road to ANN ROI is long and steep
- Good News: The steps to ANN functionality are profitable if done right:

  Data cleansing
  - Metric creation/ tracking
  - Dashboards
  - Asking yourself “what questions would we want to answer if we had no limits?

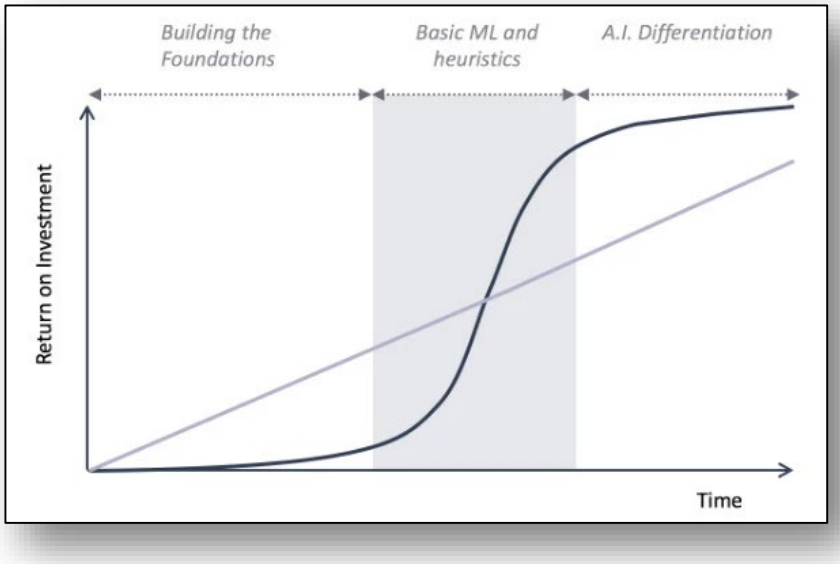

ANN Example 1: Don't tell the robot how to move; tell them why they'd want to

- Project: The Vampire's Bodyguard

The Setup

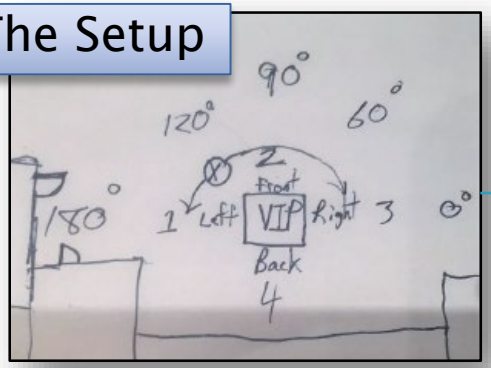

The Robot

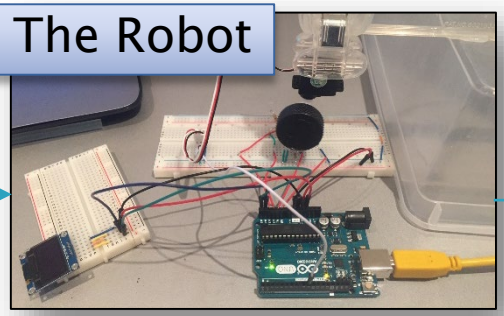

The Code

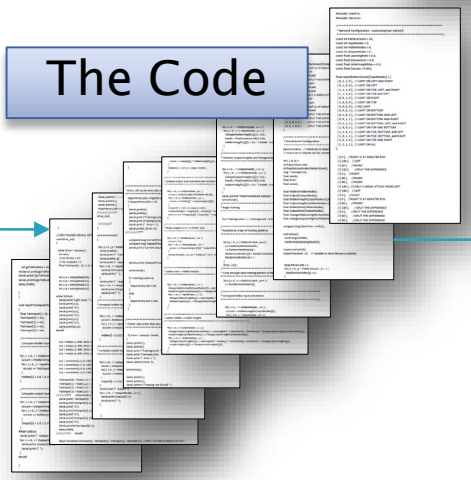

The Result

The Testing

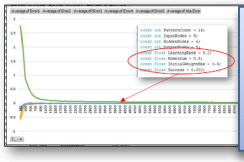

CAPTION 5: Learning cycles/progress before improvements.

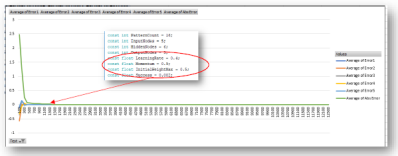

CAPTION 6: Learning cycles/progress after improvements.

The Training

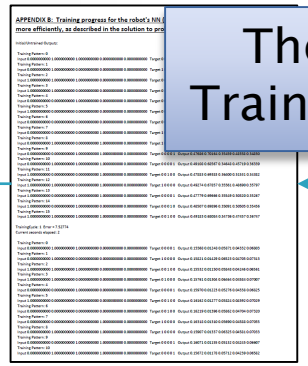

Many lessons learned and improvements identified, but a successful project

ANN Example 2: When is a student at risk of dropping out?

- Project: Drop Out Alert

The Setup

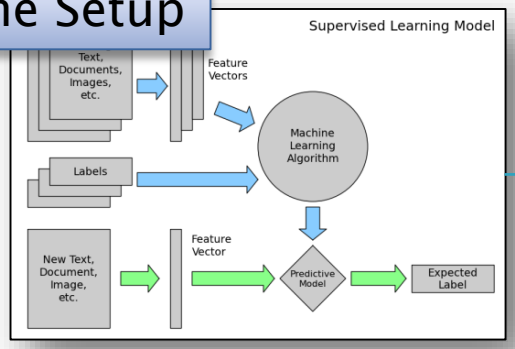

The Data

```
# box and whisker plots
dataset[cols1].plot(kind='box', subplots=True, layout=(2,2),
                    sharex=False, sharey=False)
pyplot.show()
dataset[cols2].plot(kind='box', subplots=True, layout=(2,2),
                    sharex=False, sharey=False)
pyplot.show()
dataset[cols3].plot(kind='box', subplots=True, layout=(2,2),
                    sharex=False, sharey=False)
pyplot.show()
# histograms
dataset[cols1].hist()
pyplot.show()
dataset[cols2].hist()
pyplot.show()
dataset[cols3].hist()
pyplot.show()
# scatter plot matrix
scatter_matrix(dataset[cols1])
pyplot.show()
scatter_matrix(dataset[cols2])
pyplot.show()
scatter_matrix(dataset[cols3])
pyplot.show()
```

The Code

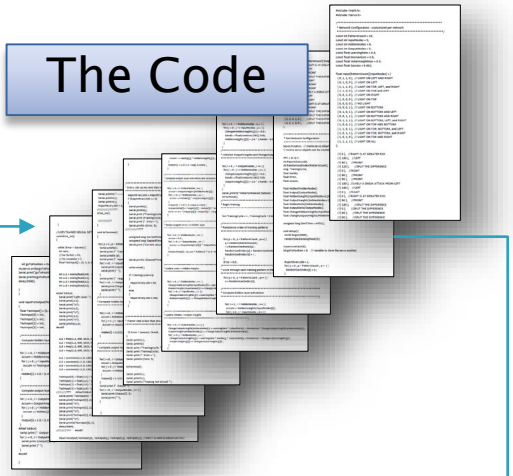

Training / Testing

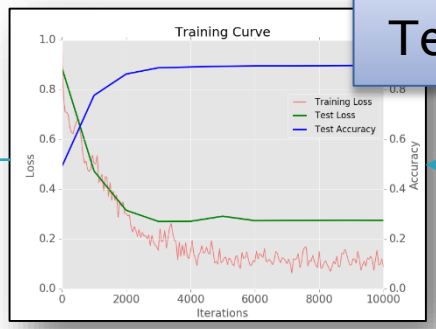

The Result

So far, the tool has performed better than the embedded (and expensive) corporate tool previously used by the university

So Remember:

- Machine Learning is improving exponentially; don’t get left behind
- ML/ANN implementation is really hard, but there is ROI at each step if done right
- Don’t let machines take over just yet…

A research scientist trained an AI neural network to generate its own candy heart messages. This was the result.

Image: Twitter/@searchthenight

dose.

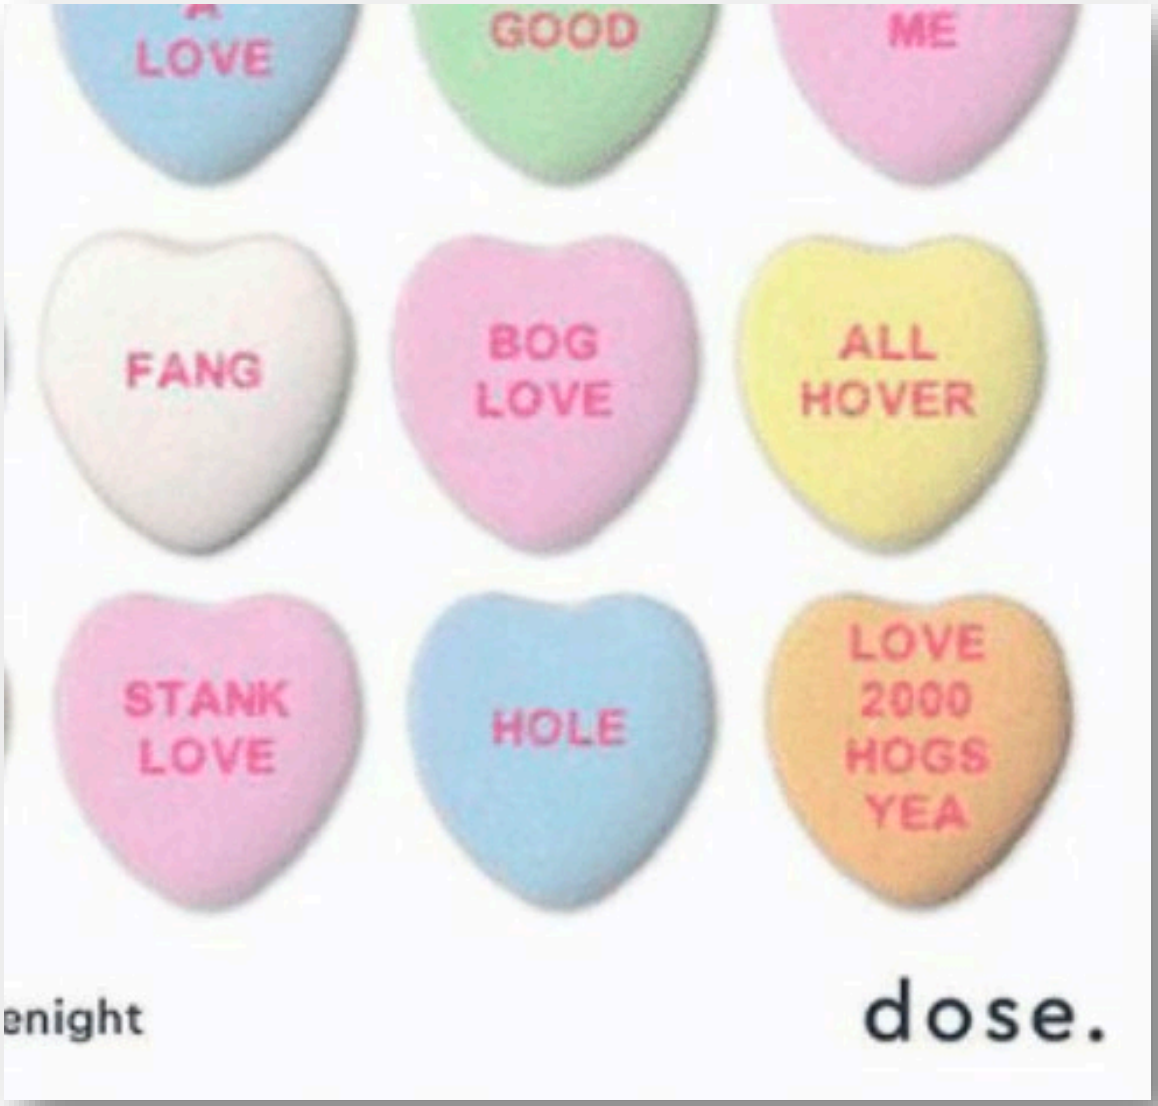
LOVE
GOOD
ME
FANG
BOG LOVE
ALL HOVER
STANK LOVE
HOLE
LOVE 2000 HOGS YEA
enight
dose.

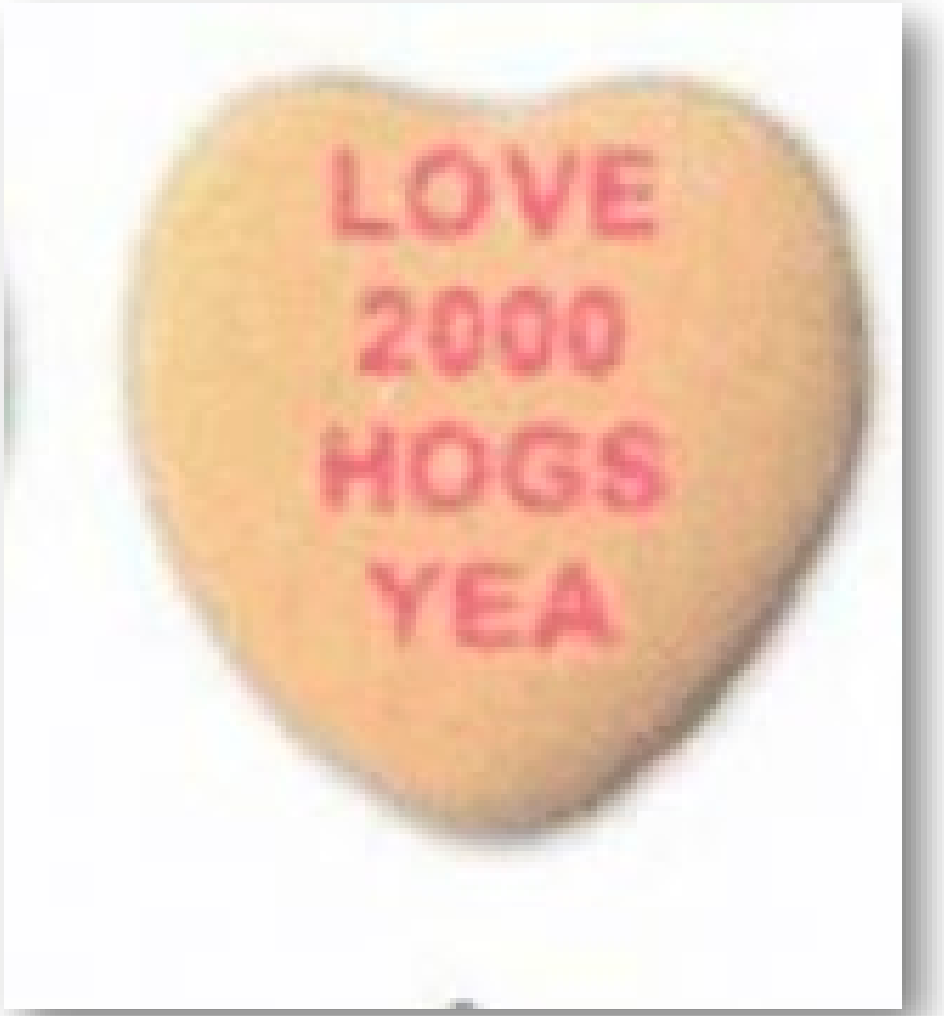
LOVE
2000
HOGS
YEA

THANK YOU!

Questions?

Nathan Eskue
Data Scientist, Material Control Manager
Northrop Grumman Innovation Systems, Launch Vehicle Division
nathan.eskue@ngc.com
480-528-2097

Resources to Learn More

- https://www.theverge.com/2018/2/28/17063780/google-ai-machine-learning-hub-crash-course-free
- https://machinelearningmastery.com/regression-tutorial-keras-deep-learning-library-python/
- http://www.businessinsider.com/computer-program-taught-itself-walk-run-play-soccer-2017-8
- https://www.reddit.com/r/videos/comments/6mw6u1/googles_deepmind_ai_just_taught_itself_to_walk/
- https://www.kdnuggets.com/2016/01/seven-steps-deep-learning.html
- https://www.toptal.com/machine-learning/an-introduction-to-deep-learning-from-perceptrons-to-deep-networks
- https://www.mathworks.com/discovery/deep-learning.html
- https://www.kdnuggets.com/2015/11/seven-steps-machine-learning-python.html
- https://www.youtube.com/watch?v=b99UVkWzYTQ
- http://www.iro.umontreal.ca/~bengioy/talks/DL-Tutorial-NIPS2015.pdf
- http://neuralnetworksanddeeplearning.com/chap1.html
- https://www.youtube.com/watch?v=962lLfW-8Jo
- https://www.youtube.com/playlist?list=PLnnr1O8OWc6boN4WHeuisJWmeQHH9D_Vg
- http://cs.stanford.edu/~quocle/tutorial1.pdf
- https://jeremykun.com/2012/12/09/neural-networks-and-backpropagation/
- https://www.popularmechanics.com/science/health/a20967153/skin-cancer-artificial-intelligence-better-than-dermatologists/

Don’t forget me as a resource as well!