

# High Level Introduction to AI



Nathan Eskue  
Associate Professor, AI in Manufacturing  
Aerospace Faculty, TU Delft

# AGENDA

- Intro
- Story Time
- Activity—What is a Cat?
- AI History and Overview
- Application: Visual inspection for quality control
- Application: Combined AI techniques for complex problems
- Conclusion


STORY TIME: “BUT I WORE THE JUICE!”

# STORY TIME: “BUT I WORE THE JUICE!”


- 1995, Pittsburgh (US), a man name Wheeler had an idea...



+



=



*THEREFORE.....*

## STORY TIME: “BUT I WORE THE JUICE!”


- But they saw this—why???



- When they caught him, he kept yelling “But I wore the juice!!”

## STORY TIME: “BUT I WORE THE JUICE!”

- Dunning-Krueger Effect—“You don’t know what you don’t know, but you really feel like you do.”



## DISCUSSION QUESTION 1

How might this story relate to the field of AI?

What implications/lessons are there for us?

## ACTIVITY: WHAT IS A CAT?

### INSTRUCTIONS

- FIRST—forget everything you know about the concept of a “cat”
- Look through the first ten photos, and try to “learn” what you believe a cat should look like based on the labelled images. (2 minutes)



## DISCUSSION QUESTION 2

In your notes, write down each new detail you notice about this concept called “cat” (remember, you wiped your memory so all this is COMPLETELY new).


IMAGE

CAT? (Y/N)

IMAGE

CAT? (Y/N)

TRAINING DATA (ANSWERS GIVEN)



Y



Y




Y



N



Y



N



Y



Y

IMAGE

CAT? (Y/N)

IMAGE

CAT? (Y/N)

TRAINING DATA (ANSWERS GIVEN)



# ACTIVITY: WHAT IS A CAT?

## INSTRUCTIONS

- Look through remaining photos
  - Label each photo you believe has a cat as “Y”
  - Label photos without a cat as “N”.

(3 minutes)

## DISCUSSION QUESTION 3

In your notes, for each numbered photo, write down both your answer and a few details from your discussion question 2 (the training set) that led you to this answer.

IMAGE

CAT? (Y/N)

IMAGE

CAT? (Y/N)

ANSWER "Y" IF THE IMAGE CONTAINS A CAT, AND "N" IF IT DOES NOT

1



5



2



6



3



7



4



8





IMAGE

CAT? (Y/N)

IMAGE

CAT? (Y/N)

ANSWER "Y" IF THE IMAGE CONTAINS A CAT, AND "N" IF IT DOES NOT

9



13



10



14



11



15



12



16



IMAGE

CAT? (Y/N)

IMAGE

CAT? (Y/N)

ANSWER "Y" IF THE IMAGE CONTAINS A CAT, AND "N" IF IT DOES NOT

17



21



18



22



19



23



20



24





## DISCUSSION QUESTION 4

Was it easier or harder than you first thought to describe what makes a cat a cat?

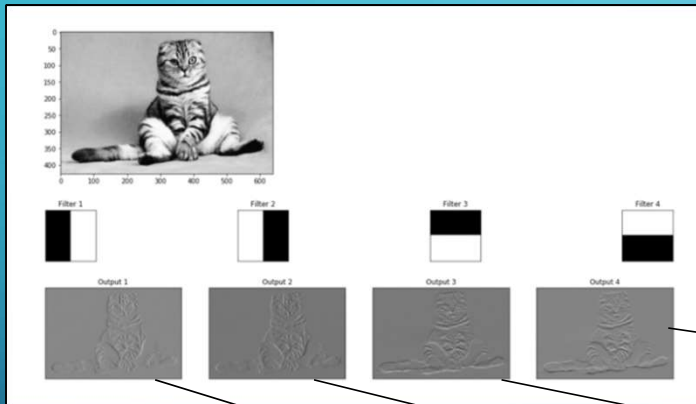
# POST-EXERCISE THOUGHTS

- Consider whether the training set might contain some bias for how to differentiate photos of cats. Is it possible that the training set is causing some patterns to be picked up as false positives?
  - Paw photo (hint—look at the training set)
  - Glowing eyes
  - Cat food bag
  - What about the fish?! Why might the AI call them fish?

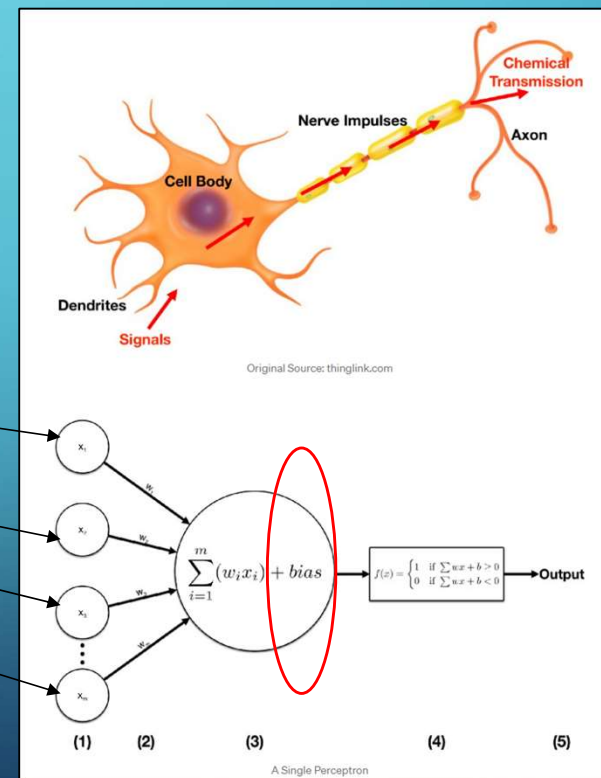
## DISCUSSION QUESTION 5

What might this teach us about how to (and how NOT to) train AI?

# NOW WE HAVE (RE)LEARNED WHAT A CAT LOOKS LIKE... WHAT DOES THIS HAVE TO DO WITH AI?



<https://towardsdatascience.com/what-is-image-classification-8e78c8235acb>



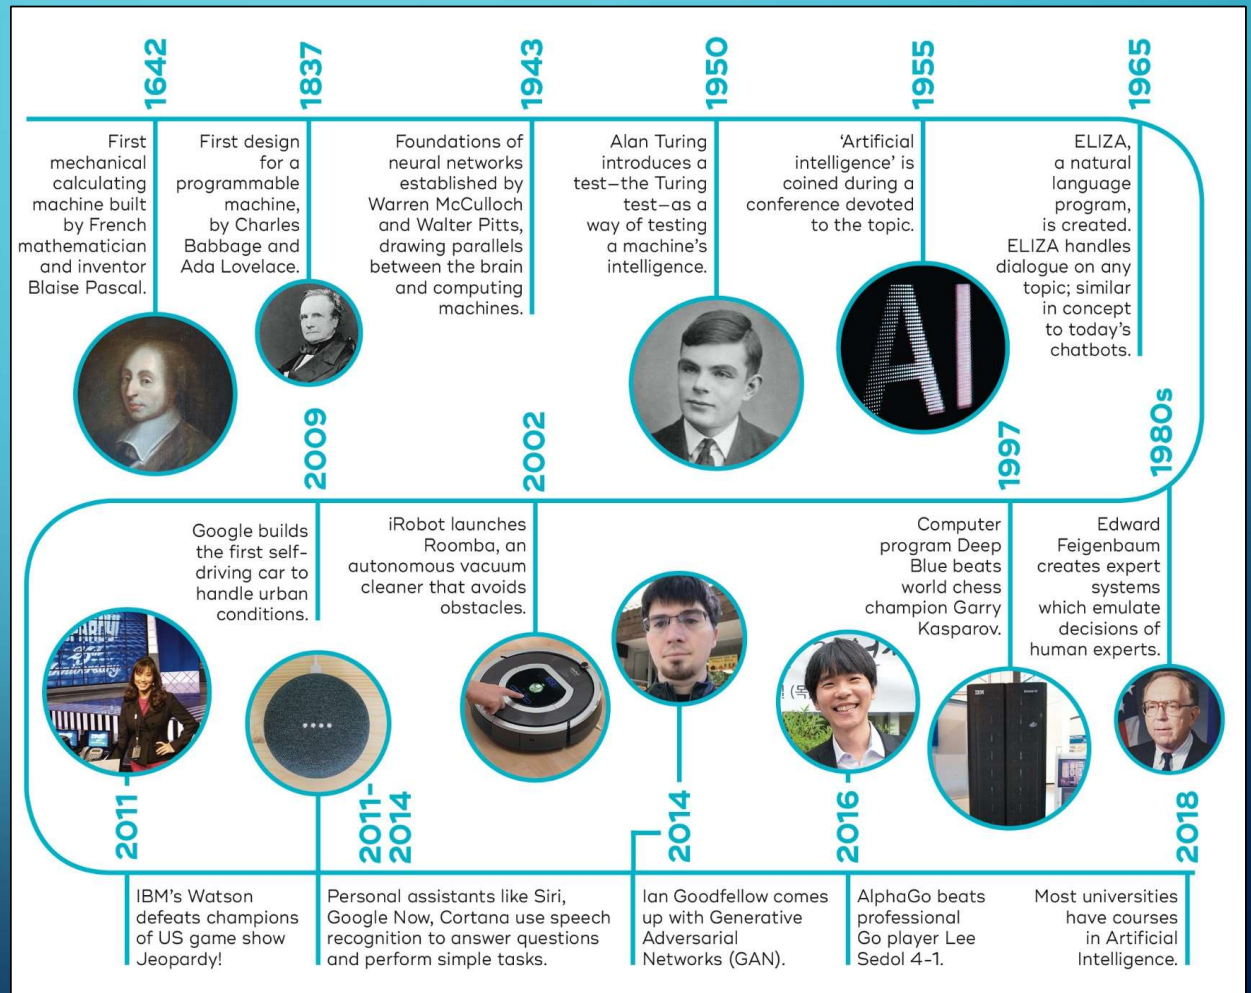
<https://towardsdatascience.com/introducing-deep-learning-and-neural-networks-deep-learning-for-robotics-1-bd68f9df5883>

CAT = YES  
(PROBABLY)

# AI HISTORY AND OVERVIEW

The background is a teal-to-blue gradient. In the corners, there are white line-art illustrations of circuit traces and nodes, resembling a printed circuit board (PCB) layout. A solid dark blue horizontal bar is positioned below the title text.

# AI TIMELINE




## DISCUSSION QUESTION 6

If I told you that AI is designed to solve TWO problems, what do you think those problems are?

# PROBLEM TO SOLVE

- Businesses thrive when they can do TWO key things better than their competitors
- 1. Have omniscient power over the current status
- 2. Have oracle power when seeing what might happen next



Clarke's Third Law: Any sufficiently advanced technology is indistinguishable from magic.

(Arthur C. Clarke)

izquotes.com

I don't think I ordered this... but I needed it.





# LET'S DIVE INTO A.I., BUT FIRST...

## WHAT EXACTLY IS A.I.?

### A DARPA Perspective on Artificial Intelligence

John Launchbury  
Director I2O, DARPA

**perceive**  
rich, complex and subtle information  
**learn**  
within an environment  
**abstract**  
to create new meanings  
**reason**  
to plan and to decide

*Artificial intelligence is a programmed ability to process information*

### Three waves of AI



Handcrafted Knowledge  
Statistical Learning  
Contextual Adaptation

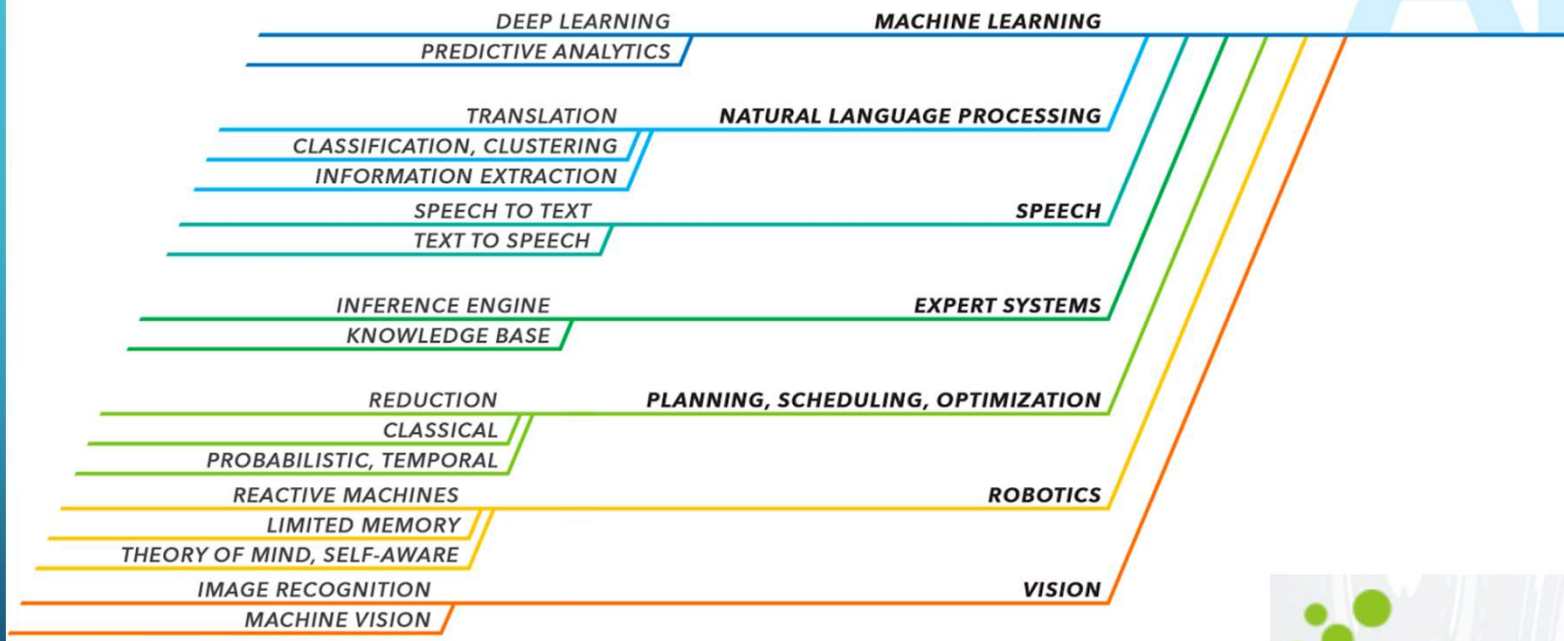


*We live  
here today*

"The science and engineering of making intelligent machines."

"The ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings."

# TYPES OF ARTIFICIAL INTELLIGENCE



DESIGN: CLOUD-NQB.COM

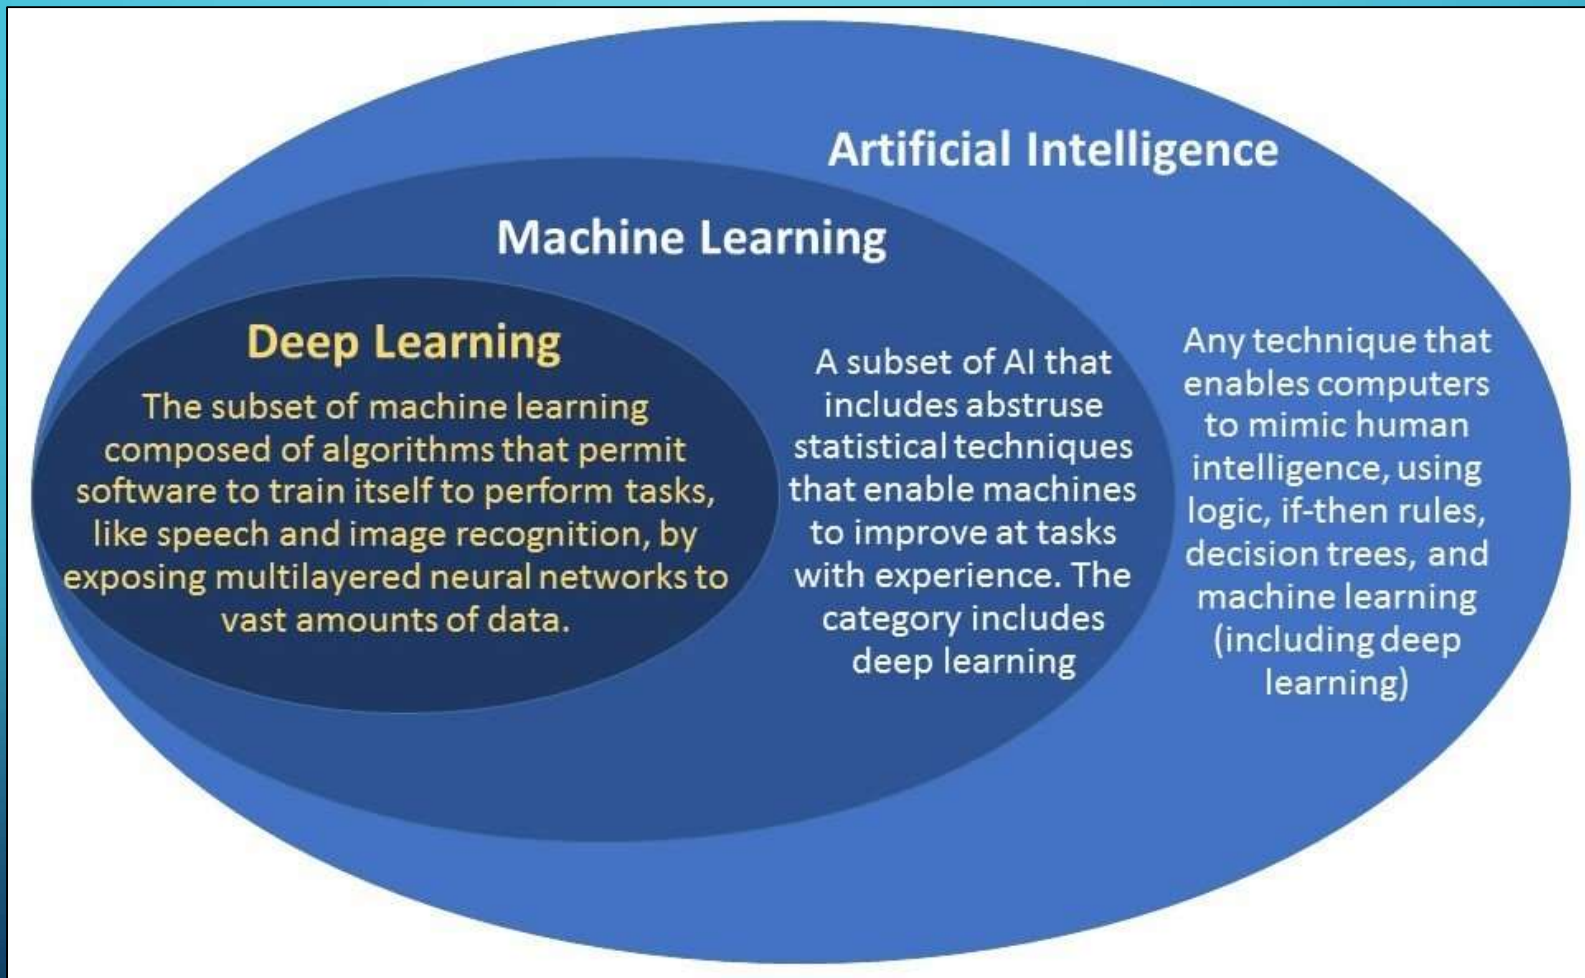


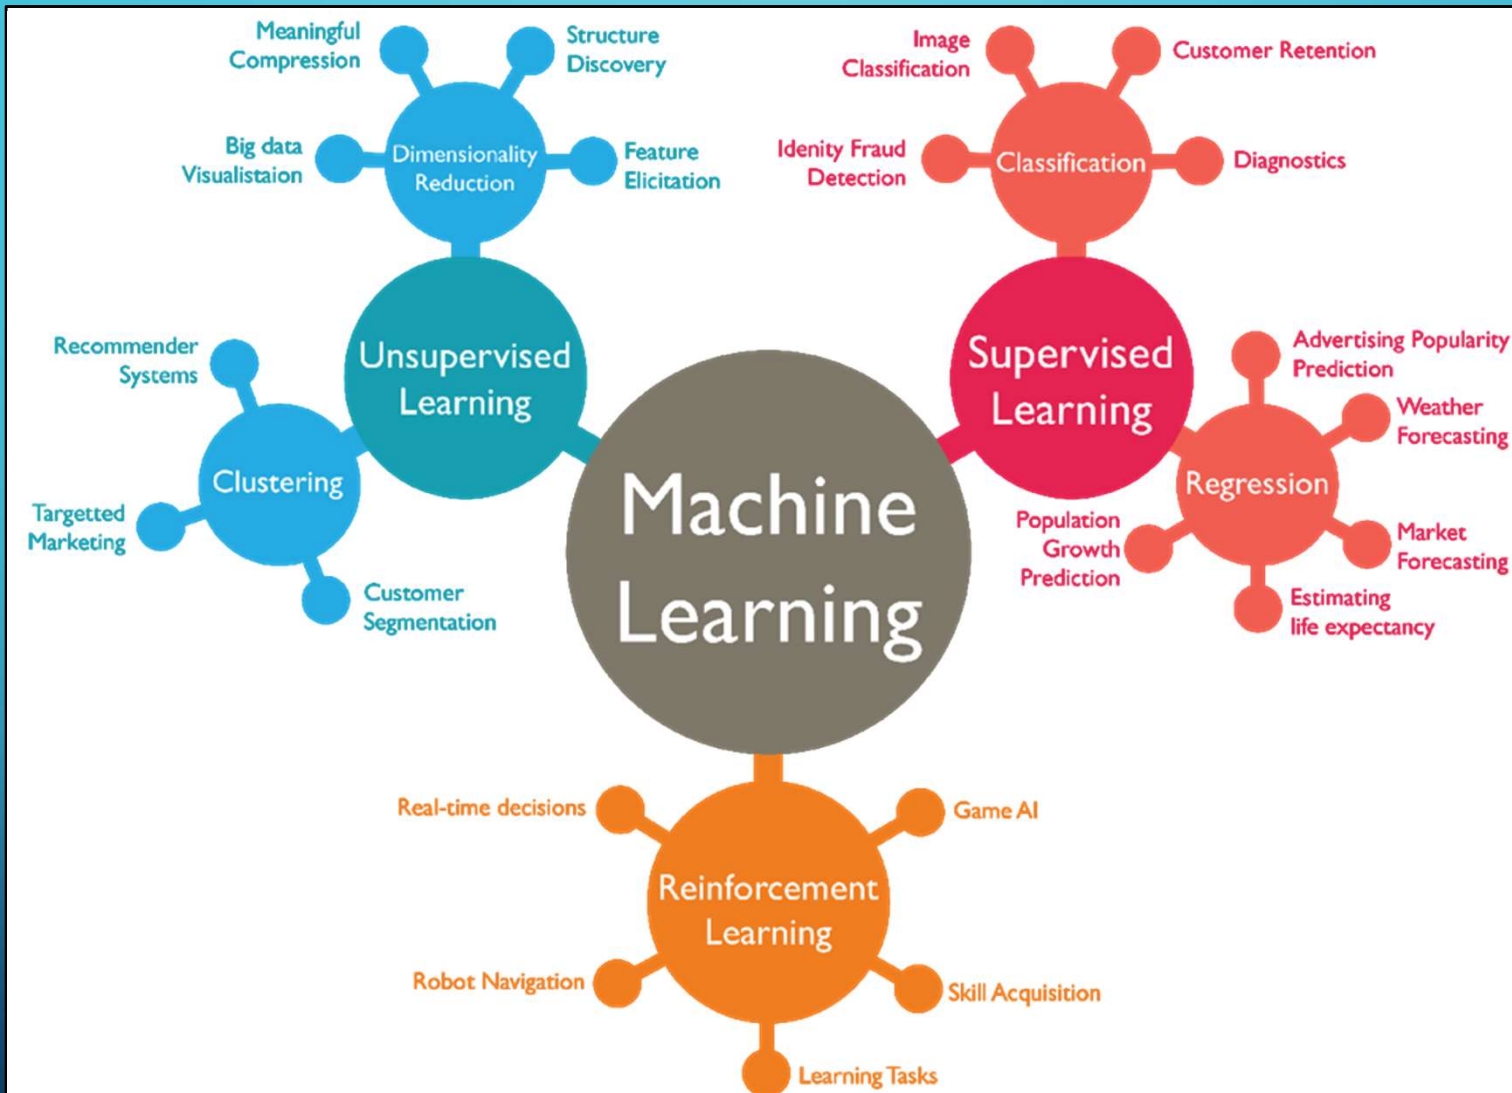
## DISCUSSION QUESTION 7

There are three key terms in defining AI:

- Artificial Intelligence
- Machine Learning
- Deep Learning

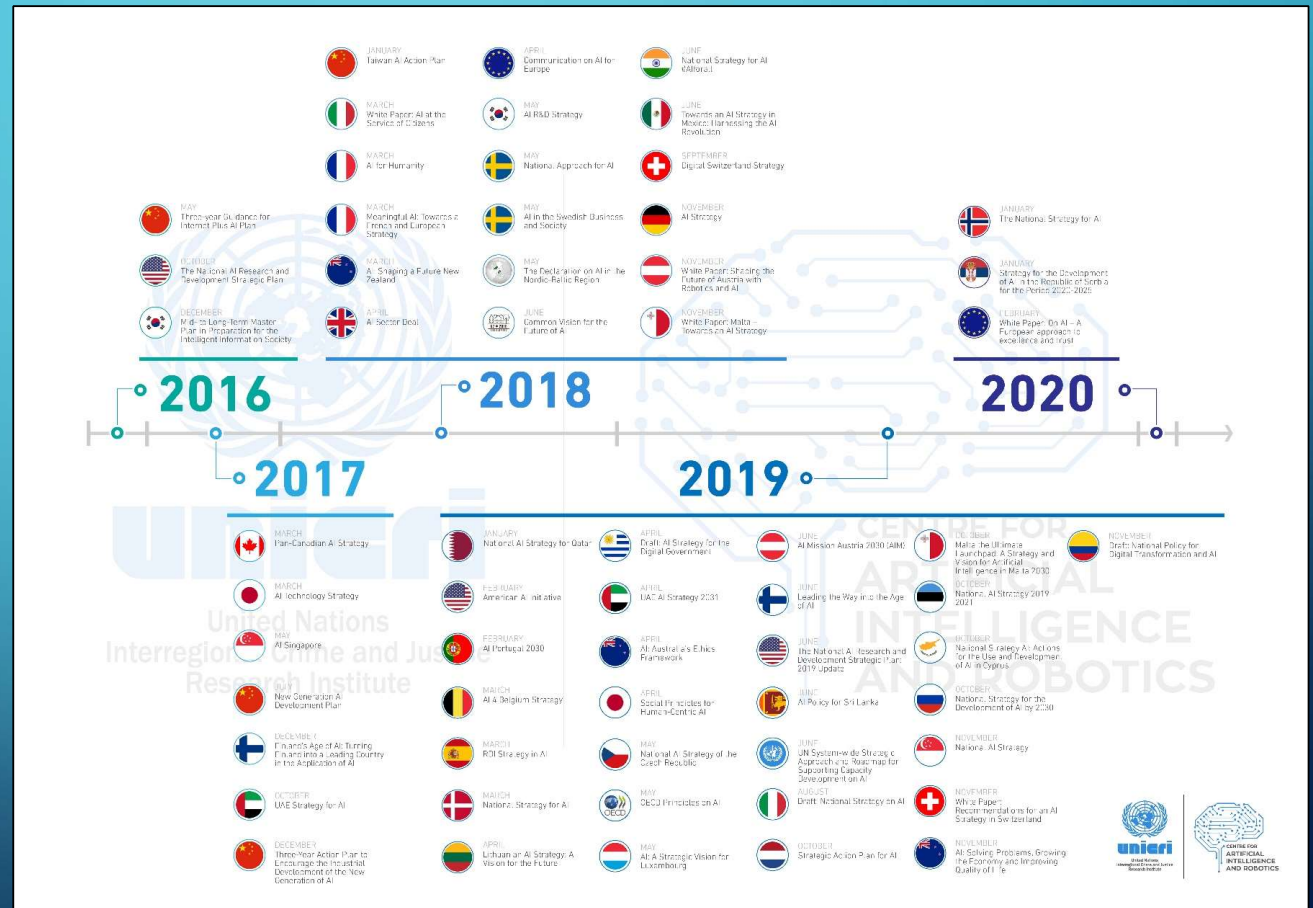
How do you think these three items relate to each other?







# GOVERNMENTS ARE CATCHING UP



# AI IS PROGRESSING FAST



Technique: Image-GPT  
June 2020



Technique: VQ-VAE  
December 2020

This was the lower half  
of the real image

<https://www.youtube.com/watch?v=c7daGcLDf0A&list=PLH5dlbFSXR1yHgdPzX--47eru-Zr2lKk&index=15>



Source: [Esser and Rombach, Ommer 2020]

The slide features a teal-to-blue gradient background with white circuit-like line art in the corners. A dark blue horizontal bar is positioned below the title text.

# AI STRENGTHS AND WEAKNESSES



## TO WORK, AN AI MODEL NEEDS:

- Data (more data usually allows more advanced AI)
- Ability to process data into information
- Ability to make sense of as many variables as possible, including their relationships
- Ability to improve decision making over time
- Ability to reason why it made a decision, how good that decision was, why it might be wrong, and the consequences if it is wrong

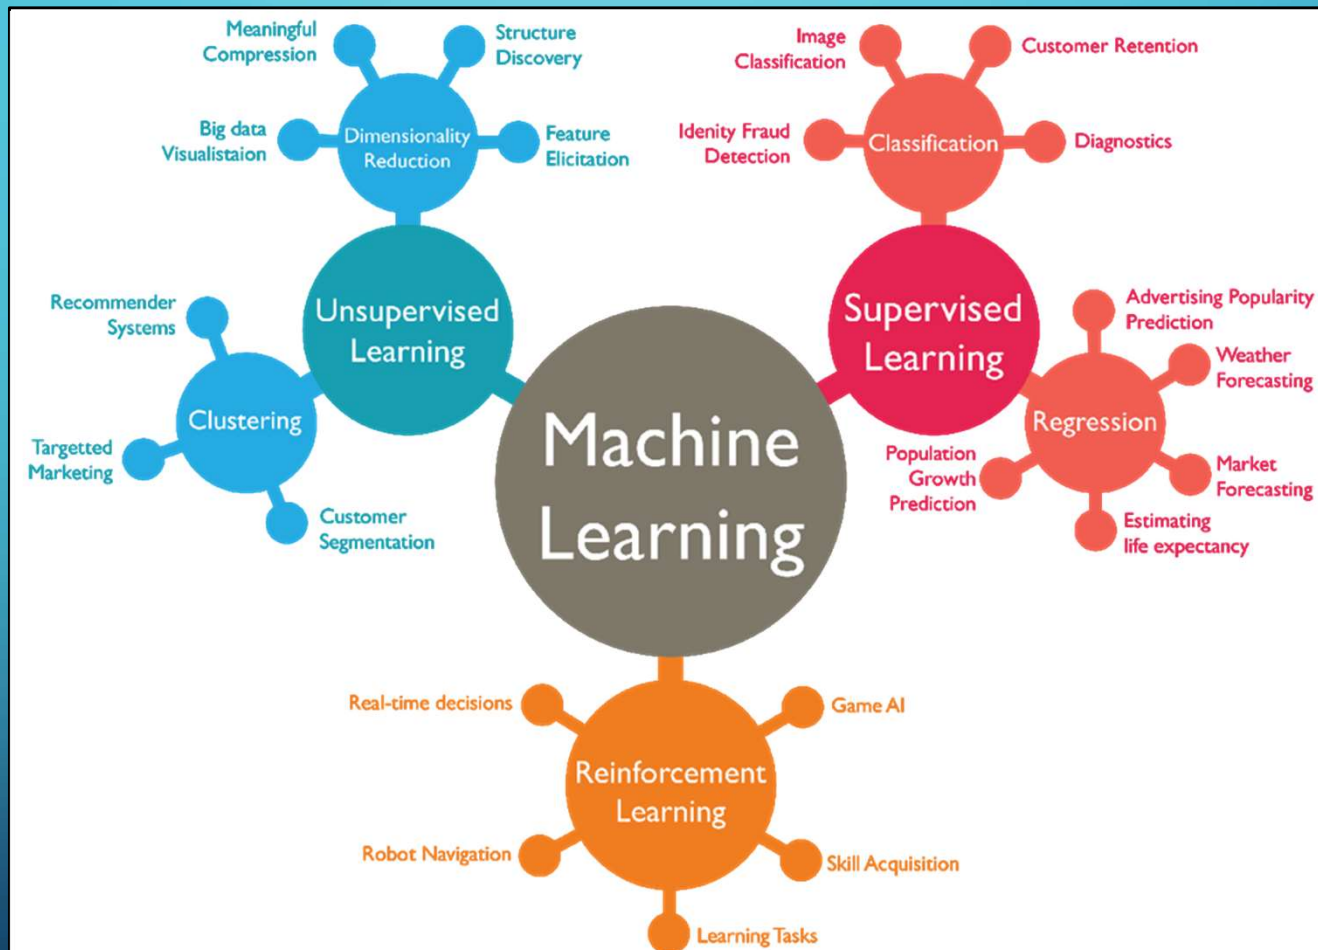
## DISCUSSION QUESTION 8

From the last slide, which item is often impossible (for now)?

### TO WORK, AN AI MODEL NEEDS:

- Data (more data usually allows more advanced AI)
- Ability to process data into information
- Ability to make sense of as many variables as possible, including their relationships
- Ability to improve decision making over time
- Ability to reason why it made a decision, how good that decision was, why it might be wrong, and the consequences if it is wrong

# SO—AI IS REALLY GOOD AT A FEW KEY THINGS...



<https://blogs.oracle.com/datascience/types-of-machine-learning-and-top-10-algorithms-everyone-should-know-v2>


# BUT STILL HAS A WAY TO GO FOR TWO BIG REASONS

The biggest risk of using A.I. is the higher cost of a failure—measured by an irrational human bias--consider the driverless car version of the Murder Game

THE PRICE OF INACTION IS FAR GREATER THAN THE COST OF MAKING A MISTAKE.

...but not with AI

QUOTEHD.COM Meister Eckhart German Philosopher



a young boy is holding a baseball bat

Statistically impressive, but individually unreliable

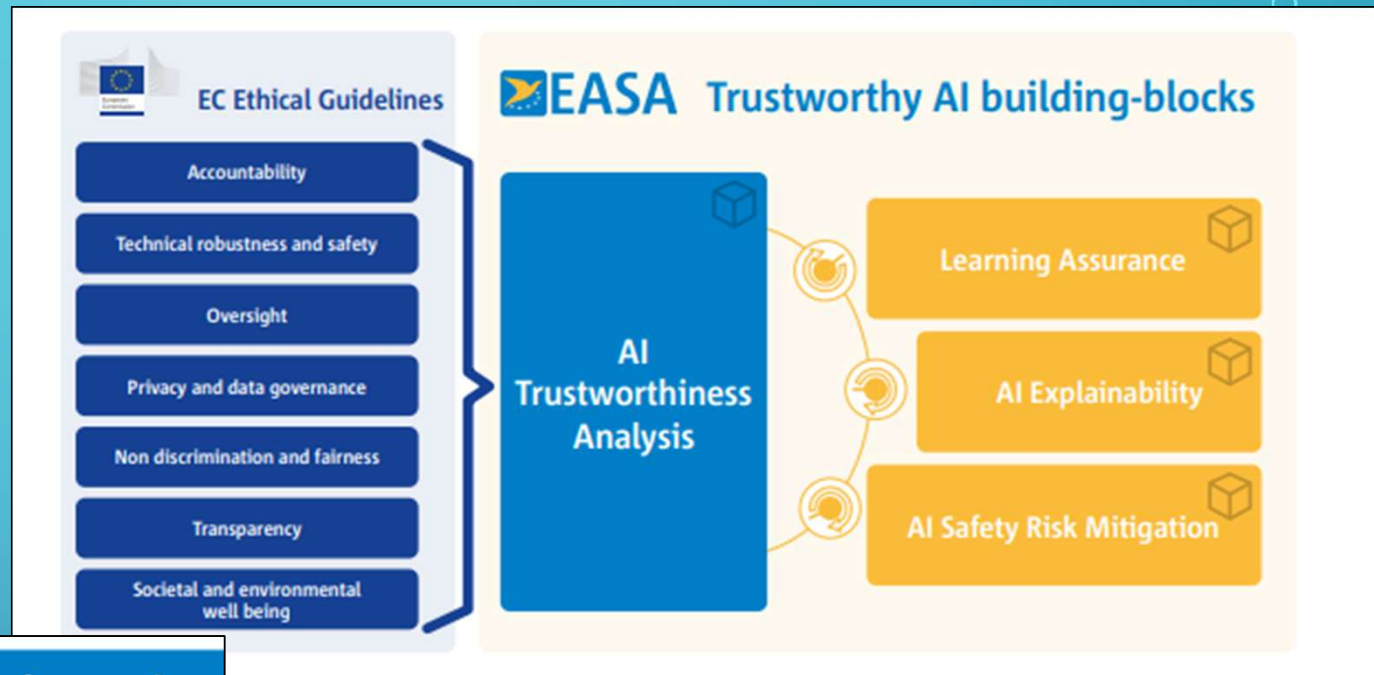


What if an AI robot surgeon makes a similar mistake?

## DISCUSSION QUESTION 9

What is an example where AI could be a significant benefit to society, but a major disaster if it failed even a few times?

# AI TRUSTWORTHINESS IS KEY



## Top 5 EASA AI Roadmap Objectives

- 1 Develop a human-centric AI Trustworthiness framework
- 2 Make EASA a leading certification authority for AI
- 3 Support European Aviation leadership in AI
- 4 Contribute to an efficient European AI research agenda
- 5 Contribute actively to EU AI strategy and initiatives



# APPLICATION: VISUAL INSPECTION FOR QUALITY CONTROL

---



# VISUAL INSPECTION FOR QUALITY CONTROL

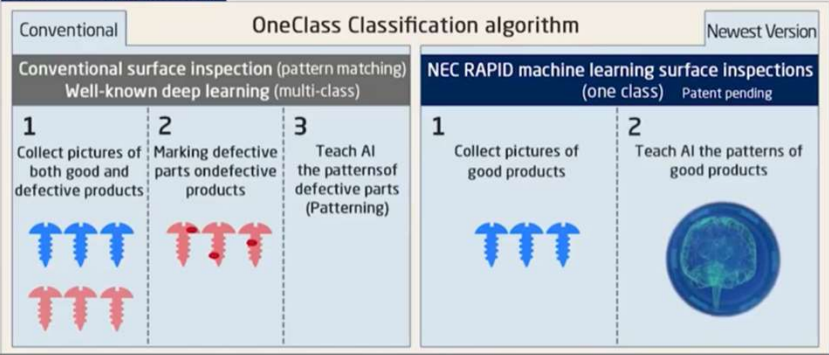
- Problem: For precision manufacturing with few defects, how do we ensure quality, AI-driven quality inspection?
- For high quantity, this is straightforward
- For low quantity, this can be difficult
- What “tricks” can be used to increase defect detection accuracy?



What RAPID machine learning can do

**Manufacturing industry**

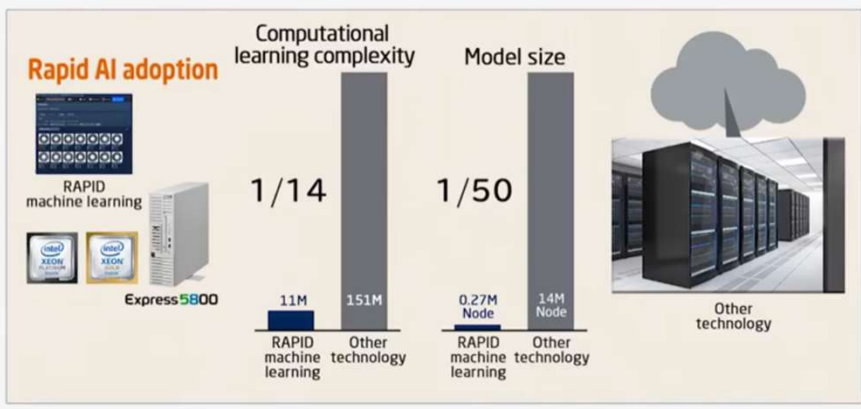
**1 Image analysis**  
for high-accuracy classification support



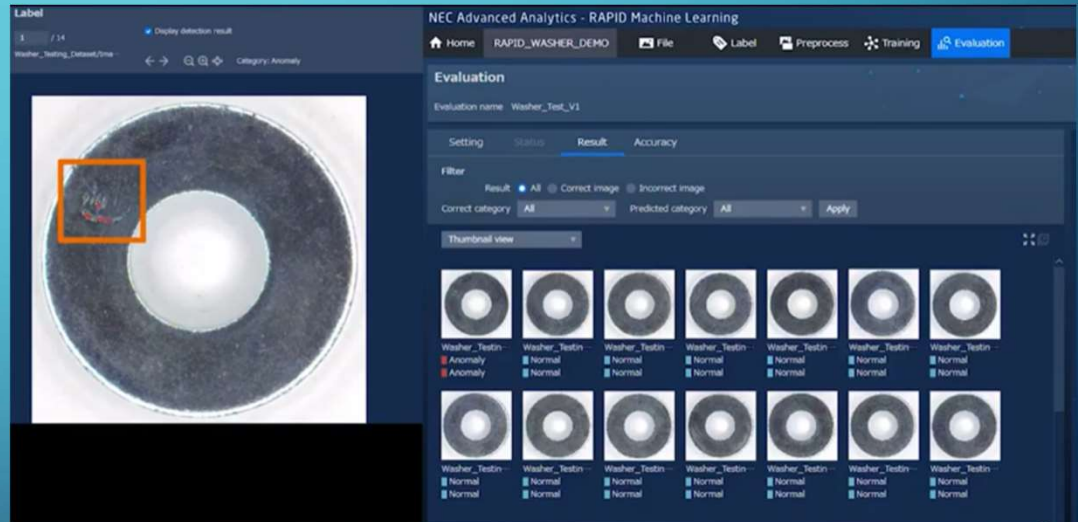
NEC's Proprietary Technology enables us to classify defective products by just collecting pictures of good products

Technology supporting RAPID machine learning

Software optimized for Intel® processors  
Fast, lightweight, high-accuracy

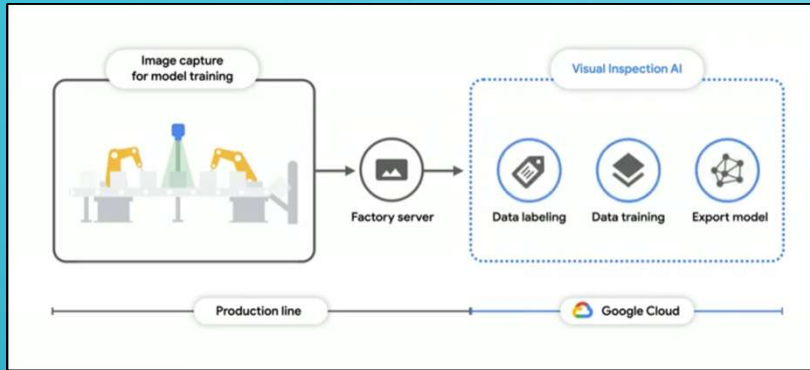


# NEC "RAPID" TECHNIQUE

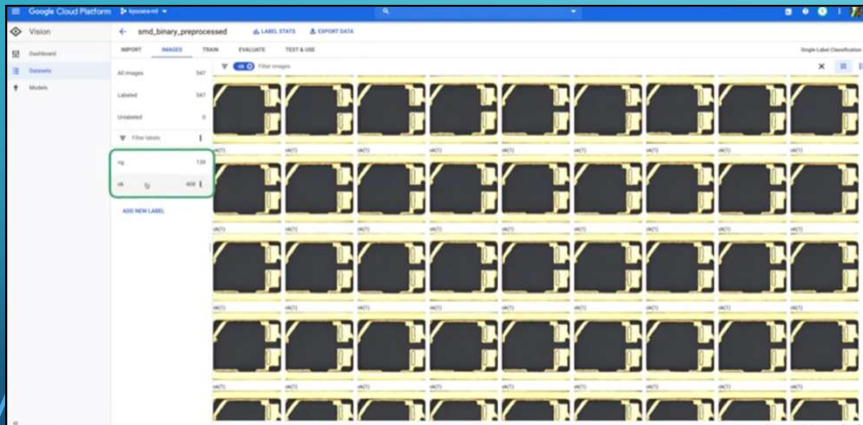


**Key—Part of training is identifying the actual defect vs. simply labelling the part defective**

# GOOGLE CLOUD TECHNIQUE



## Manufacturing use cases



### Confusion Matrix

	Predicted NG	Predicted OK
True NG	95.10%	4.90%
True OK	1.47%	98.53%


### Performance

Average inference latency	8.98 ms
---------------------------	---------



**Key—Small details like alignment and edge detection can increase accuracy drastically**

# VISUAL INSPECTION IN ACTION



Multiple camera setup for high quantity





# VISUAL INSPECTION IN ACTION

SAM|XL

TU Delft



High def camera (possible stereoscopic use)

Robotic arm for multiple angle, precise photos




## DISCUSSION QUESTION 10

Confusion Matrix		
	Predicted NG	Predicted OK
True NG	95.10%	4.90%
True OK	1.47%	98.53%

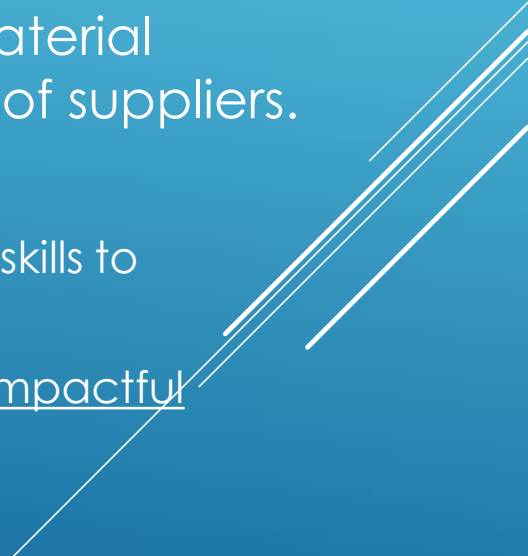
  

Performance	
Average inference latency	8.98 ms



- Is 95-98% acceptable for spacecraft?
- What could be done to improve these numbers?
- If this is best, what process changes ensure this is still useful?

# A DAY AT THE MOONSHOT FACTORY : SUPPLY CHAIN OPTIMIZATION/MANAGEMENT


- ▶ Problem: Need to coordinate hugely complex Bills of Material (BOMs) with long lead times, low yield, and many layers of suppliers.
  - ▶ Need to combine omniscience of the current situation, and oracle skills to predict the most likely outcome.
  - ▶ The key goals are to minimize surprises, and always know the most impactful action you can take in any given moment.
- 

Visualize: If you were an Omniscient Oracle, how would you run your business?

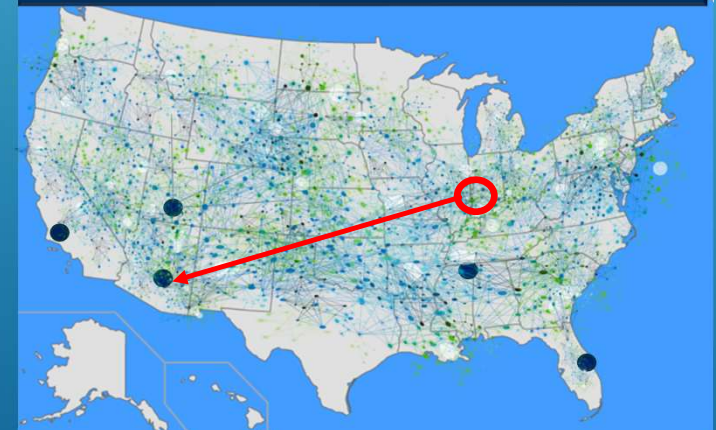
## Application: Shortest Time to Best Action

1. Forecasted schedule at risk

Supervised  
Regression



LIVE SUPPLY CHAIN STATUS






Visualize: If you were an Omniscient Oracle, how would you run your business?

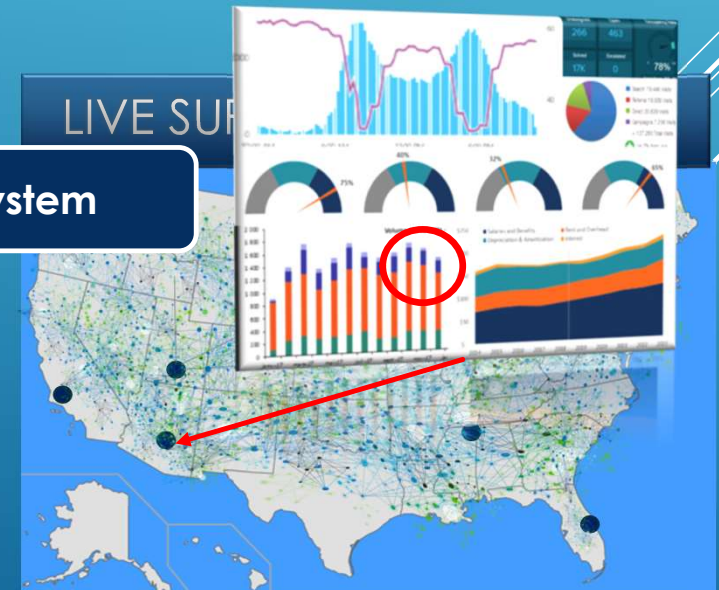
## Application: Shortest Time to Best Action

1. Forecasted schedule at risk
2. Highlight concern with data

Supervised  
Regression



Expert System




Visualize: If you were an Omniscient Oracle, how would you run your business?

## Application: Shortest Time to Best Action

1. Forecasted schedule at risk
2. Highlight concern with data
3. Pull context data based on prediction

Supervised  
Regression



Expert System




Visualize: If you were an Omniscient Oracle, how would you run your business?

## Application: Shortest Time to Best Action

1. Forecasted schedule at risk
2. Highlight concern with data
3. Pull context data based on prediction
4. View live feed, contact process owner

Supervised  
Regression



Unsupervised  
Classification

Expert System



Visualize: If you were an Omniscient Oracle, how would you run your business?

## Application: Shortest Time to Best Action

1. Forecasted schedule at risk
2. Highlight concern with data
3. Pull context data based on prediction
4. View live feed, contact process owner



**System in Action**



QUESTIONS?

## 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=PLnr1O8OWc6boN4WHeuisJWmeQHH9D\\_Vg](https://www.youtube.com/playlist?list=PLnr1O8OWc6boN4WHeuisJWmeQHH9D_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!**



