Machine Learning

How Smart Can It Really Be?

March 11, 2019

Speaker

Shayne Champion

CISO

Conversant Group

Overview: Today's Agenda

1) Definitions
2) Machine Learning
3) Neural Networks
4) Pitfalls
5) Cyber Applications

Definitions

What is Learning?

learn·ing /ˈlərniNG**/** *noun*

The ability to acquire and apply knowledge and skills.

Source: google.com “define learning”; Machine Learning: For Beginners

What is Intelligence?

in·tel·li·gence /inˈteləjəns**/** *noun*

The acquisition of knowledge or skills through experience, study, or by being taught.

Change influenced by previous experience

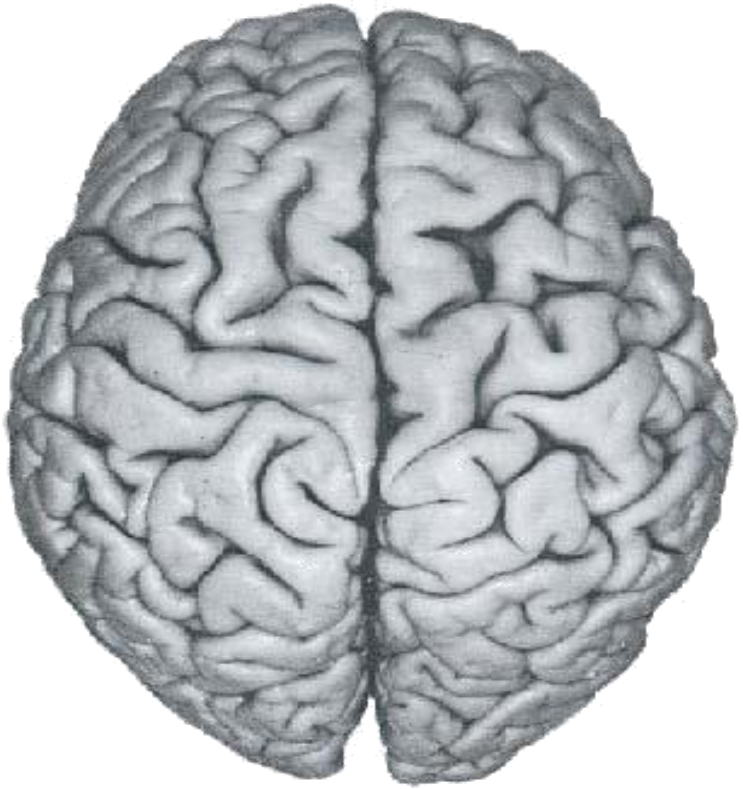

Source: google.com “define intelligence”

Artificial Intelligence (AI)

Machines that can perform tasks that are characteristic of **human intelligence**
(*e.g., planning, understanding language, recognizing objects and sounds, learning, and problem solving*)

Two types of AI

a) **General AI**:
Has *all* of the characteristics of human intelligence

b) **Narrow (Specific) AI**:
Exhibits *some* facet(s) of human intelligence, and can do that facet extremely well, but is lacking in other areas

Machine Learning

The Turing Test

The exhibition of natural intelligence in a machine a machine would be indistinguishable from a human being in natural language conversation.

Source: Machine Learning: For Beginners

Recap: Artificial Intelligence

Intelligence:
The ***ability*** to acquire and apply knowledge and skills

Learning:
The ***acquisition*** of knowledge or skills

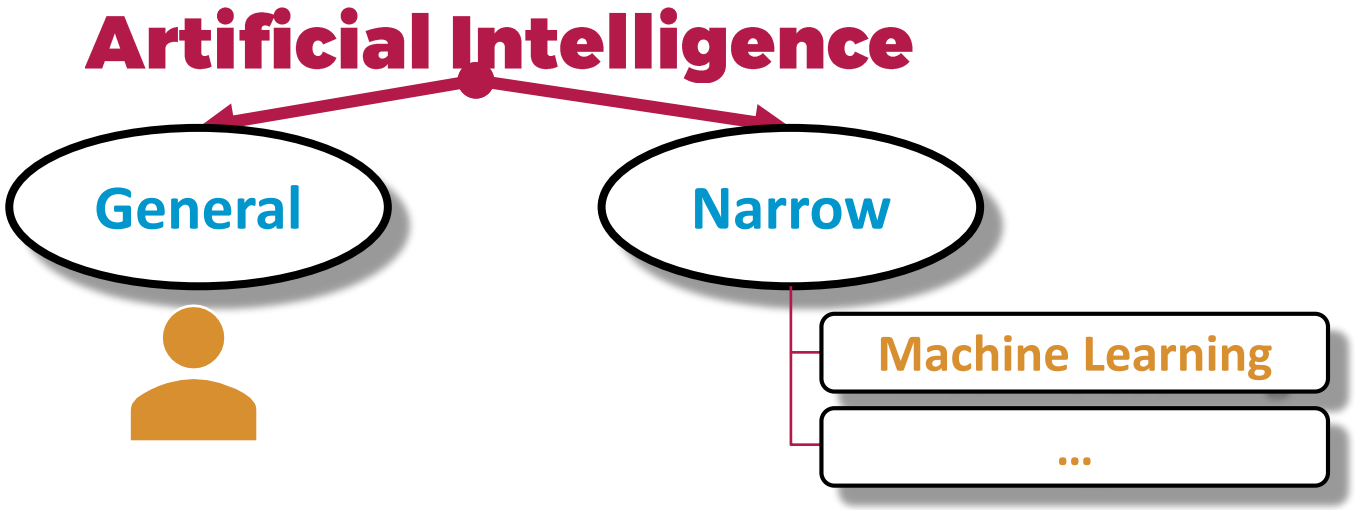

Machine Learning

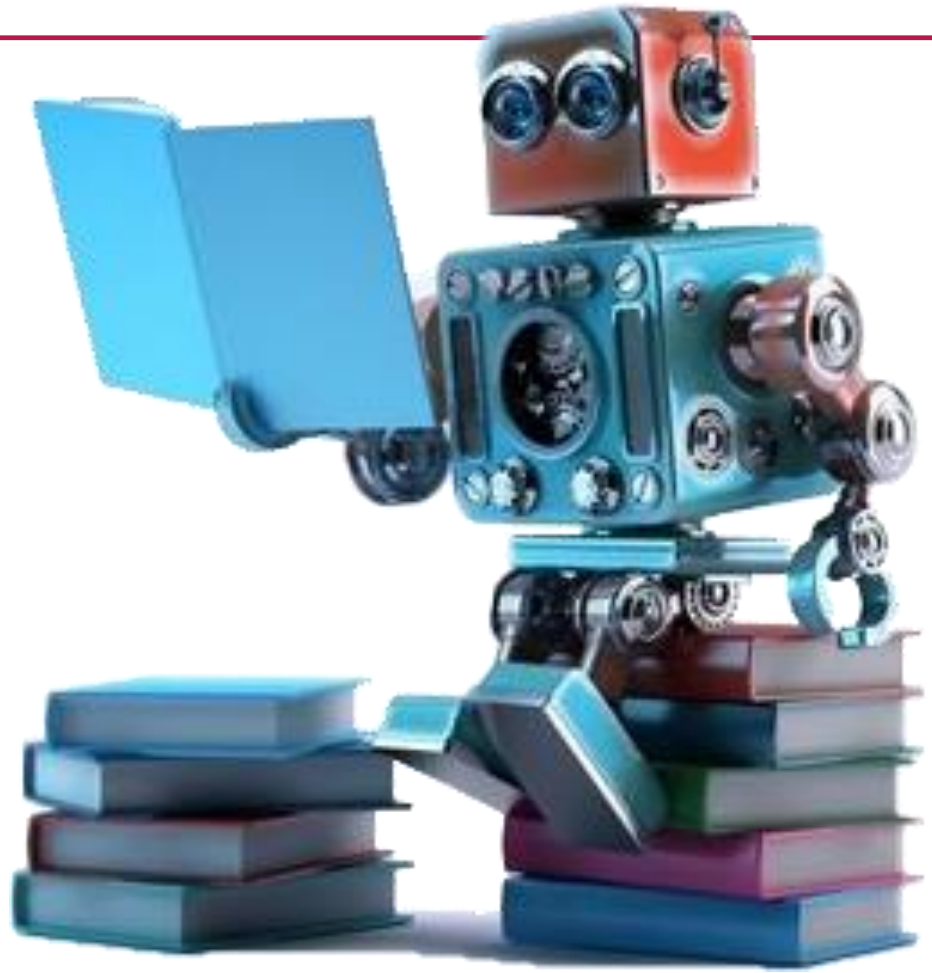

Machine Learning (ML)

Algorithms that receive data and apply statistical analysis to predict the output data within an acceptable range.

Goals of ML

a) Adapt and change from previous experience based on **pattern recognition** & **iteratively** adjust response without human intervention (the algorithm outputs become new inputs)

b) Standardize' the development of AI… 'without programming'

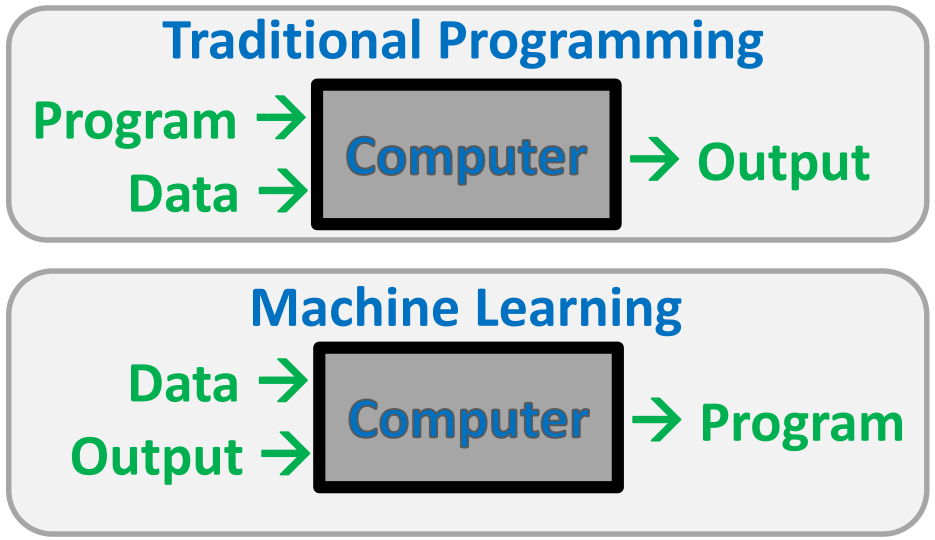

Source: Machine Learning: For Beginners, How will AI and machine learning impact CSPs? https://inform.tmforum.org/data-analytics-and-ai/2017/08/will-ai-machine-learning-impact-csps/

ML Types

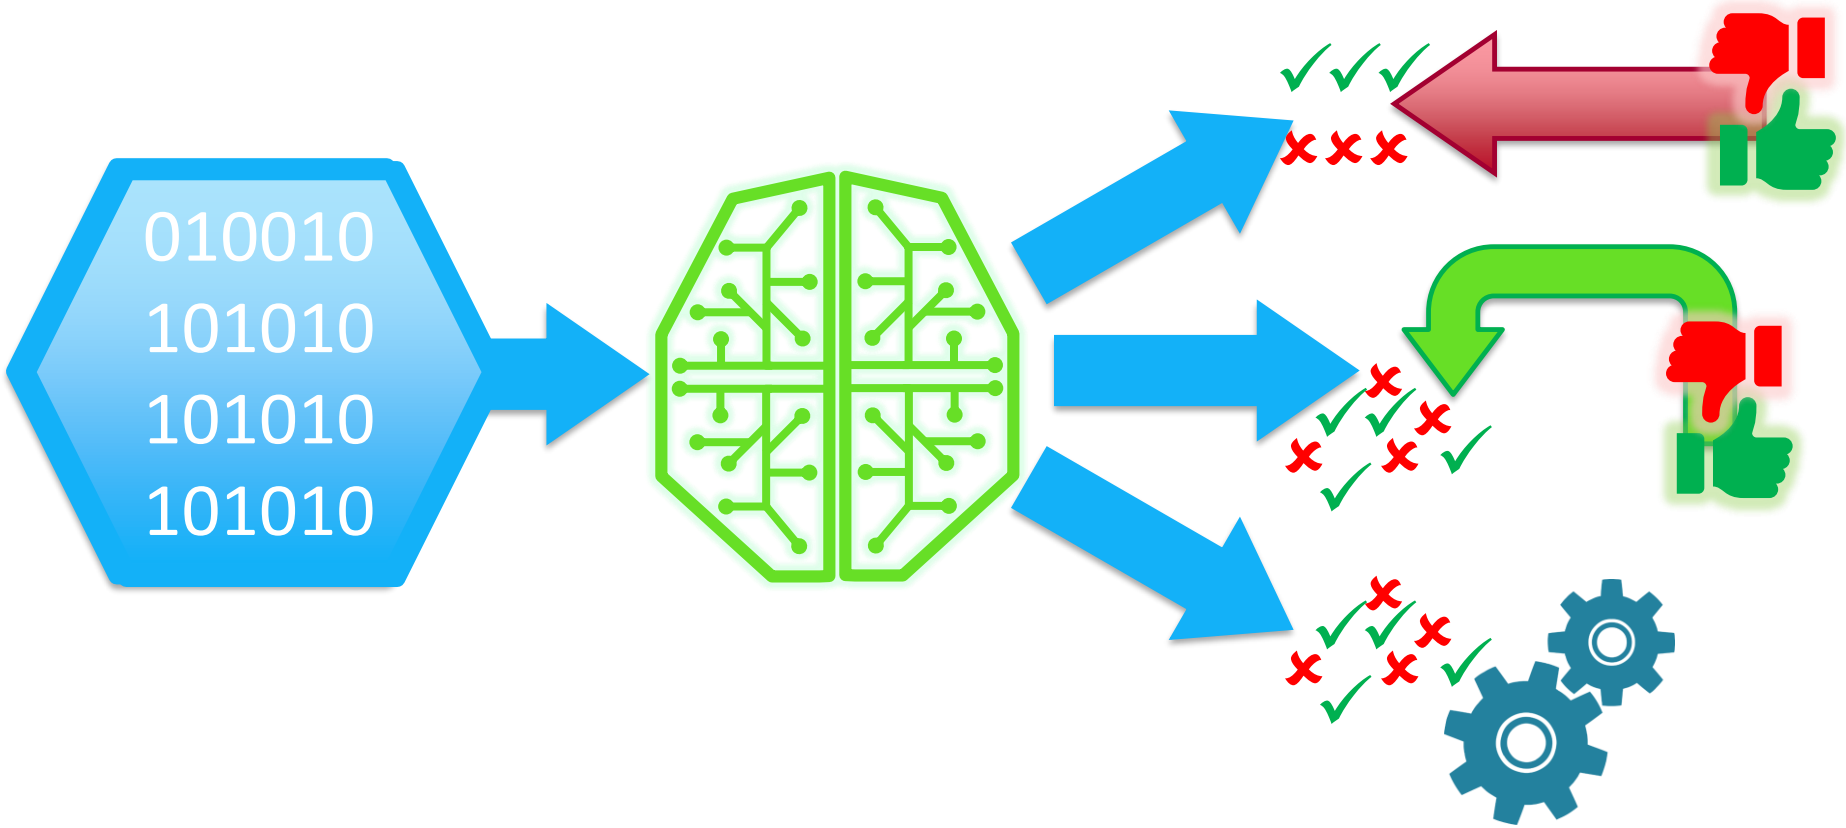

Machine Learning Types

1) Supervised

- Regression (numeric)
- Classification (class || tag)

EXAMPLE:

Source: Machine Learning: For Beginners, http://cdn2.hubspot.net/hubfs/305377/Supervised_vs_Unsupervised_ML.png

Machine Learning Types

1) Supervised

2) Re-enforced Learning

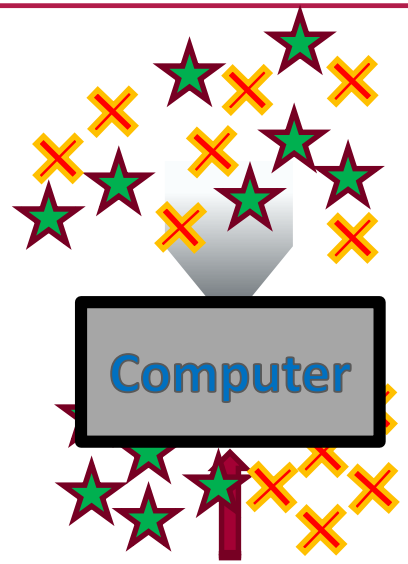

EXAMPLE:

Source: Machine Learning: For Beginners, http://cdn2.hubspot.net/hubfs/305377/Supervised_vs_Unsupervised_ML.png

Machine Learning Types

1) Supervised

2) Re-enforced Learning

3) Unsupervised

Deep Learning

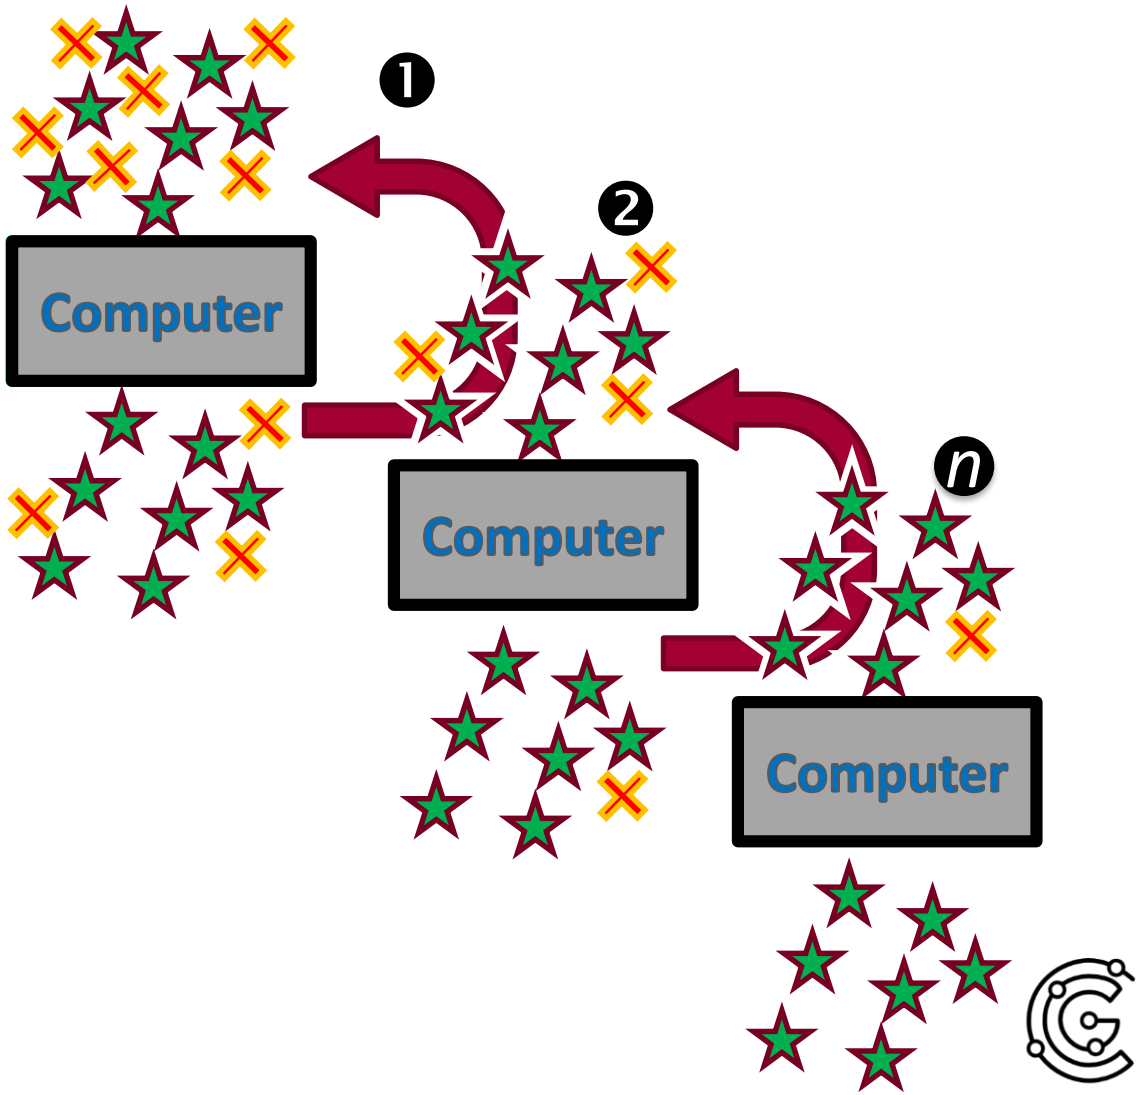

Source: Machine Learning: For Beginners, http://cdn2.hubspot.net/hubfs/305377/Supervised_vs_Unsupervised_ML.png

Machine Learning Types

1) Supervised

2) Re-enforced Learning

3) Unsupervised

EXAMPLE:

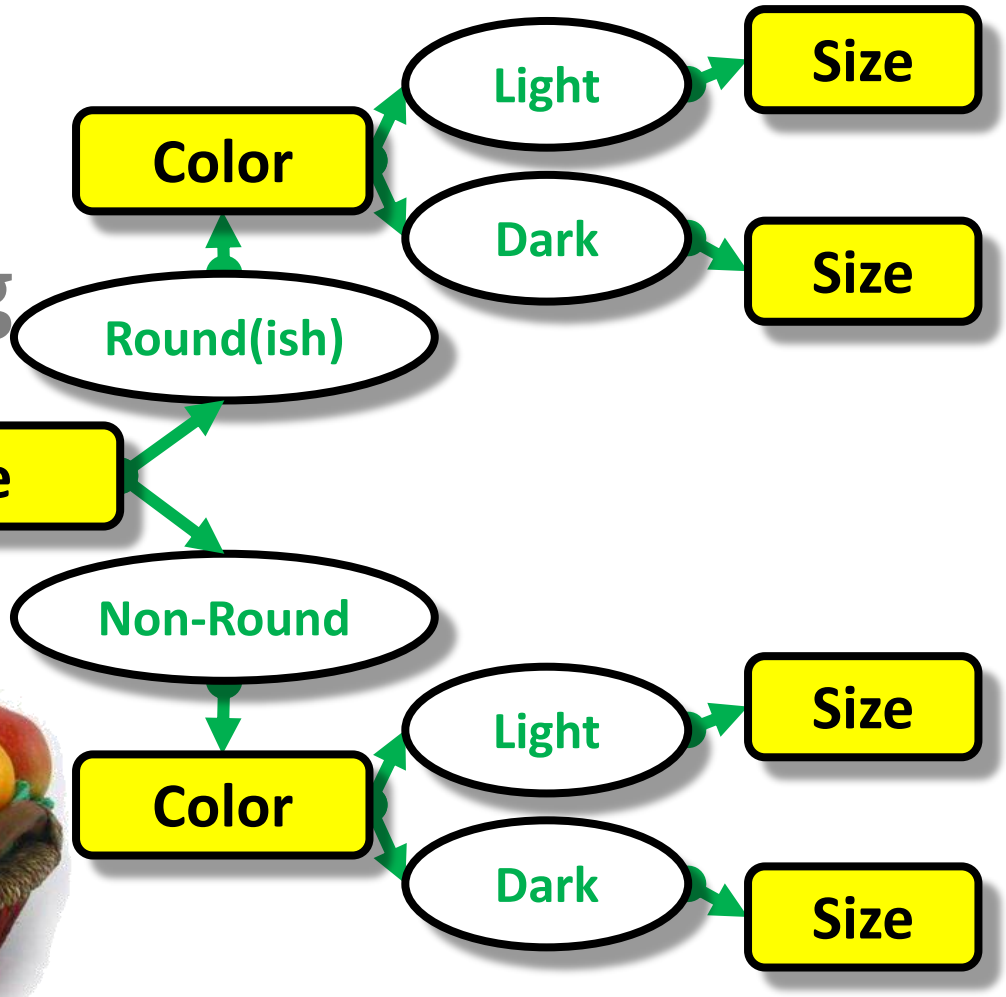

Source: Machine Learning: For Beginners, http://cdn2.hubspot.net/hubfs/305377/Supervised_vs_Unsupervised_ML.png

Machine Learning Types

1) Supervised

2) Re-enforced Learning

3) Unsupervised

These are <CLASS>
(human defined)

These are Similar
(no value judgements)

Source: Machine Learning: For Beginners, http://cdn2.hubspot.net/hubfs/305377/Supervised_vs_Unsupervised_ML.png

Deep Learning

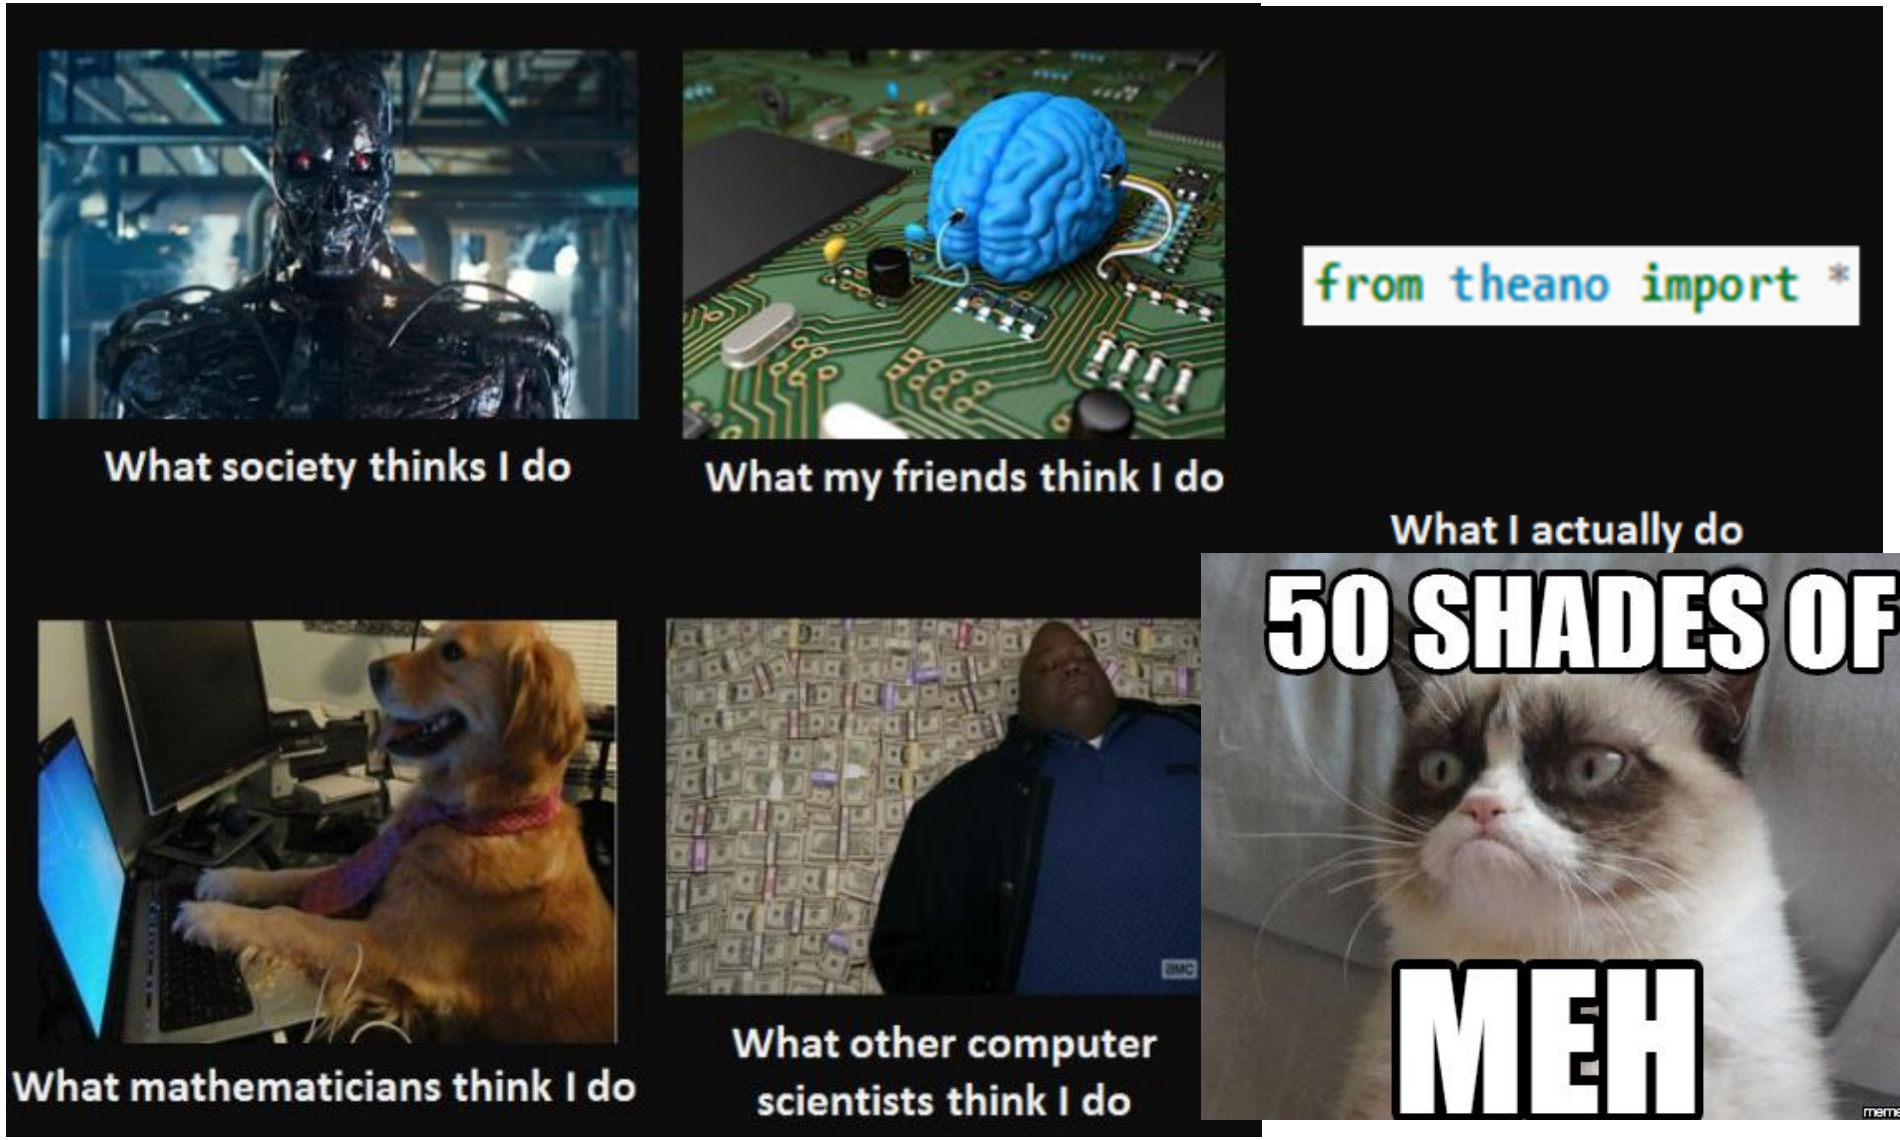

Recap: ML Types

1) Supervised
2) Re-enforced Learning
3) Unsupervised

INPUT	OUTPUT	FEEDBACK
Human Sorted	Human Review	Human
Unsorted	**Human Review**	**Human**
Unsorted	**Algorithm**	**Algorithm**

- Unsupervised is more extensible
- Unsupervised cannot make value judgements

ML Methods

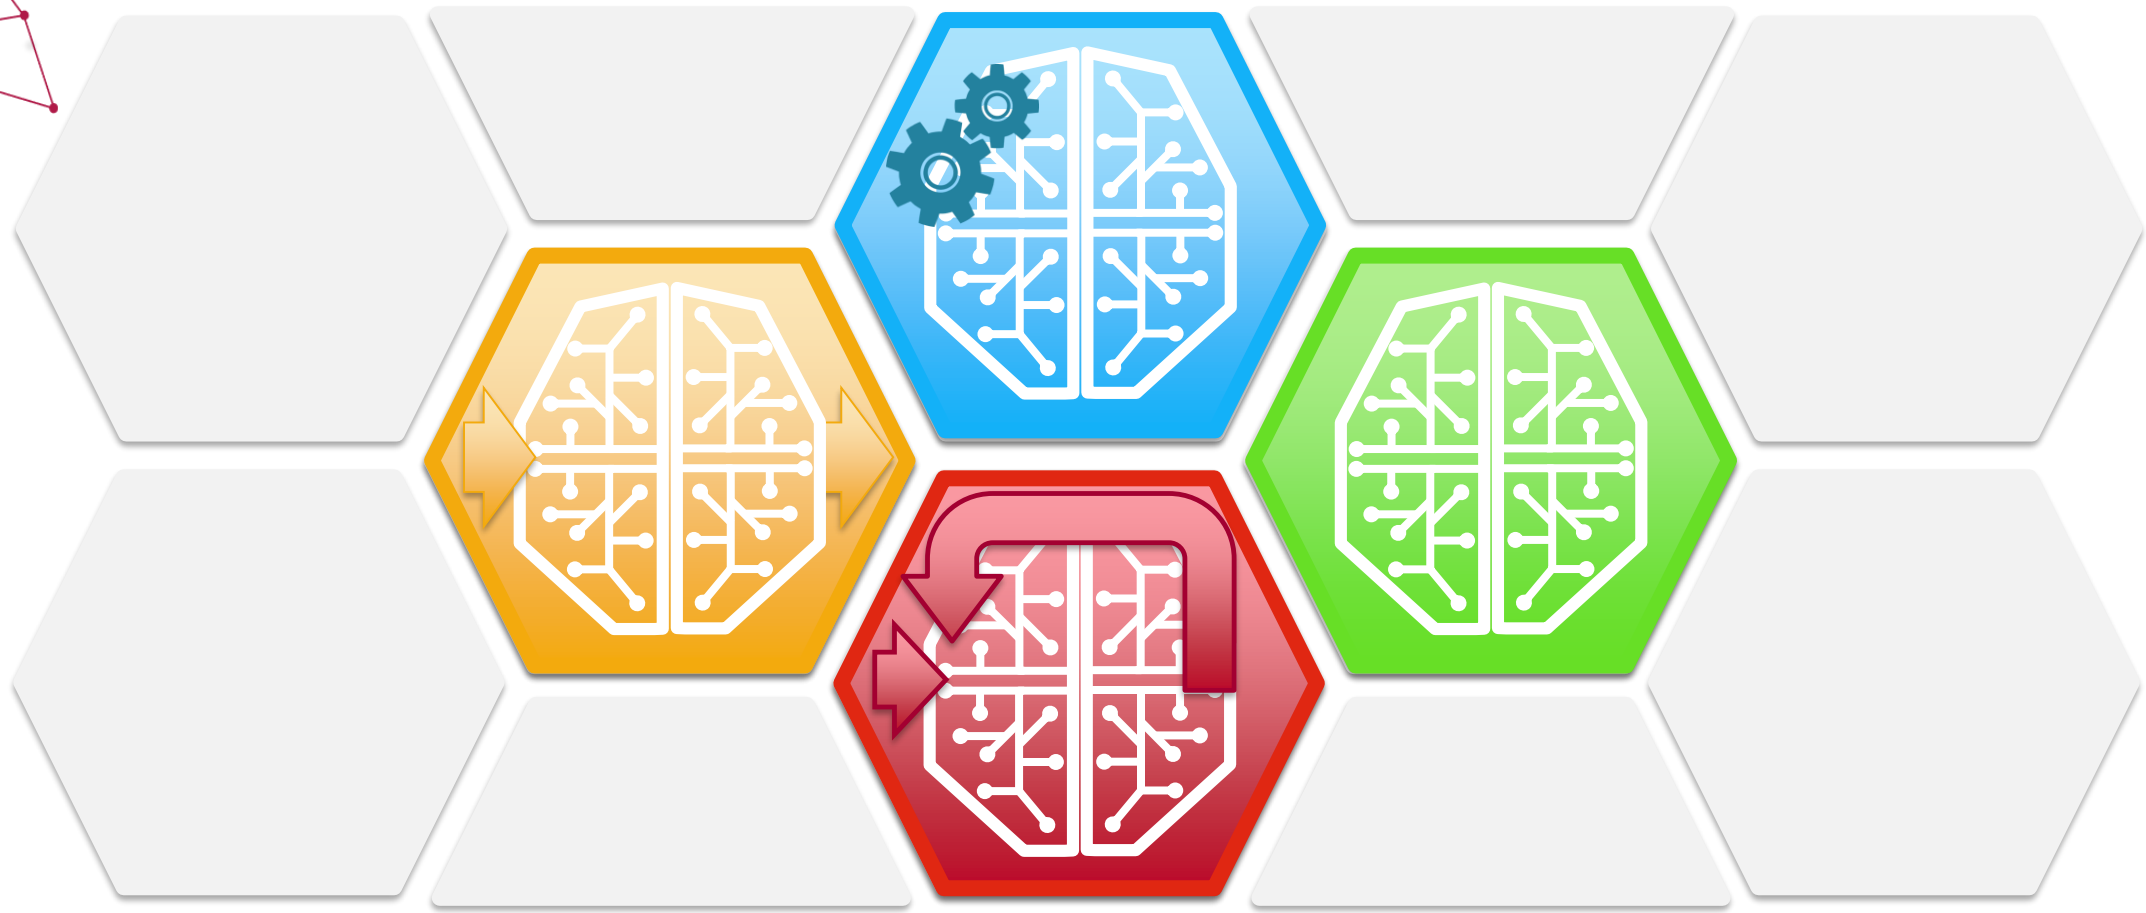

ML Methods: Linear Regression

A model of the relationship between a scalar **dependent variable** *Y* and **one or more explanatory variables** (or independent variables) denoted *X*.

$$y = mx + b$$

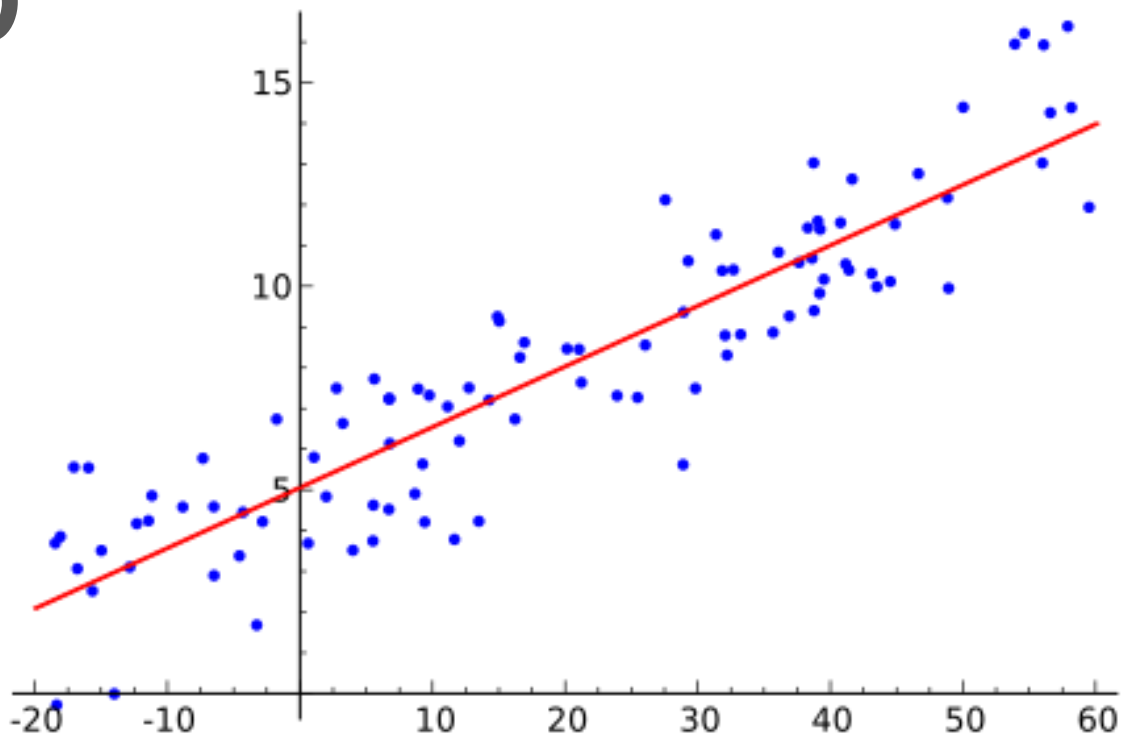

Source: https://en.wikipedia.org/wiki/Linear_regressio

ML Methods: Decision Trees

Data is continuously split according to certain parameters based on human input. (Supervised)

Source: Types of Machine Learning Algorithms, http://www.intechopen.com/books/new-advances-in-machinelearning

ML Methods: Bayesian Networks

A **probabilistic** graphical model representing **variables** and the **relationships** between them

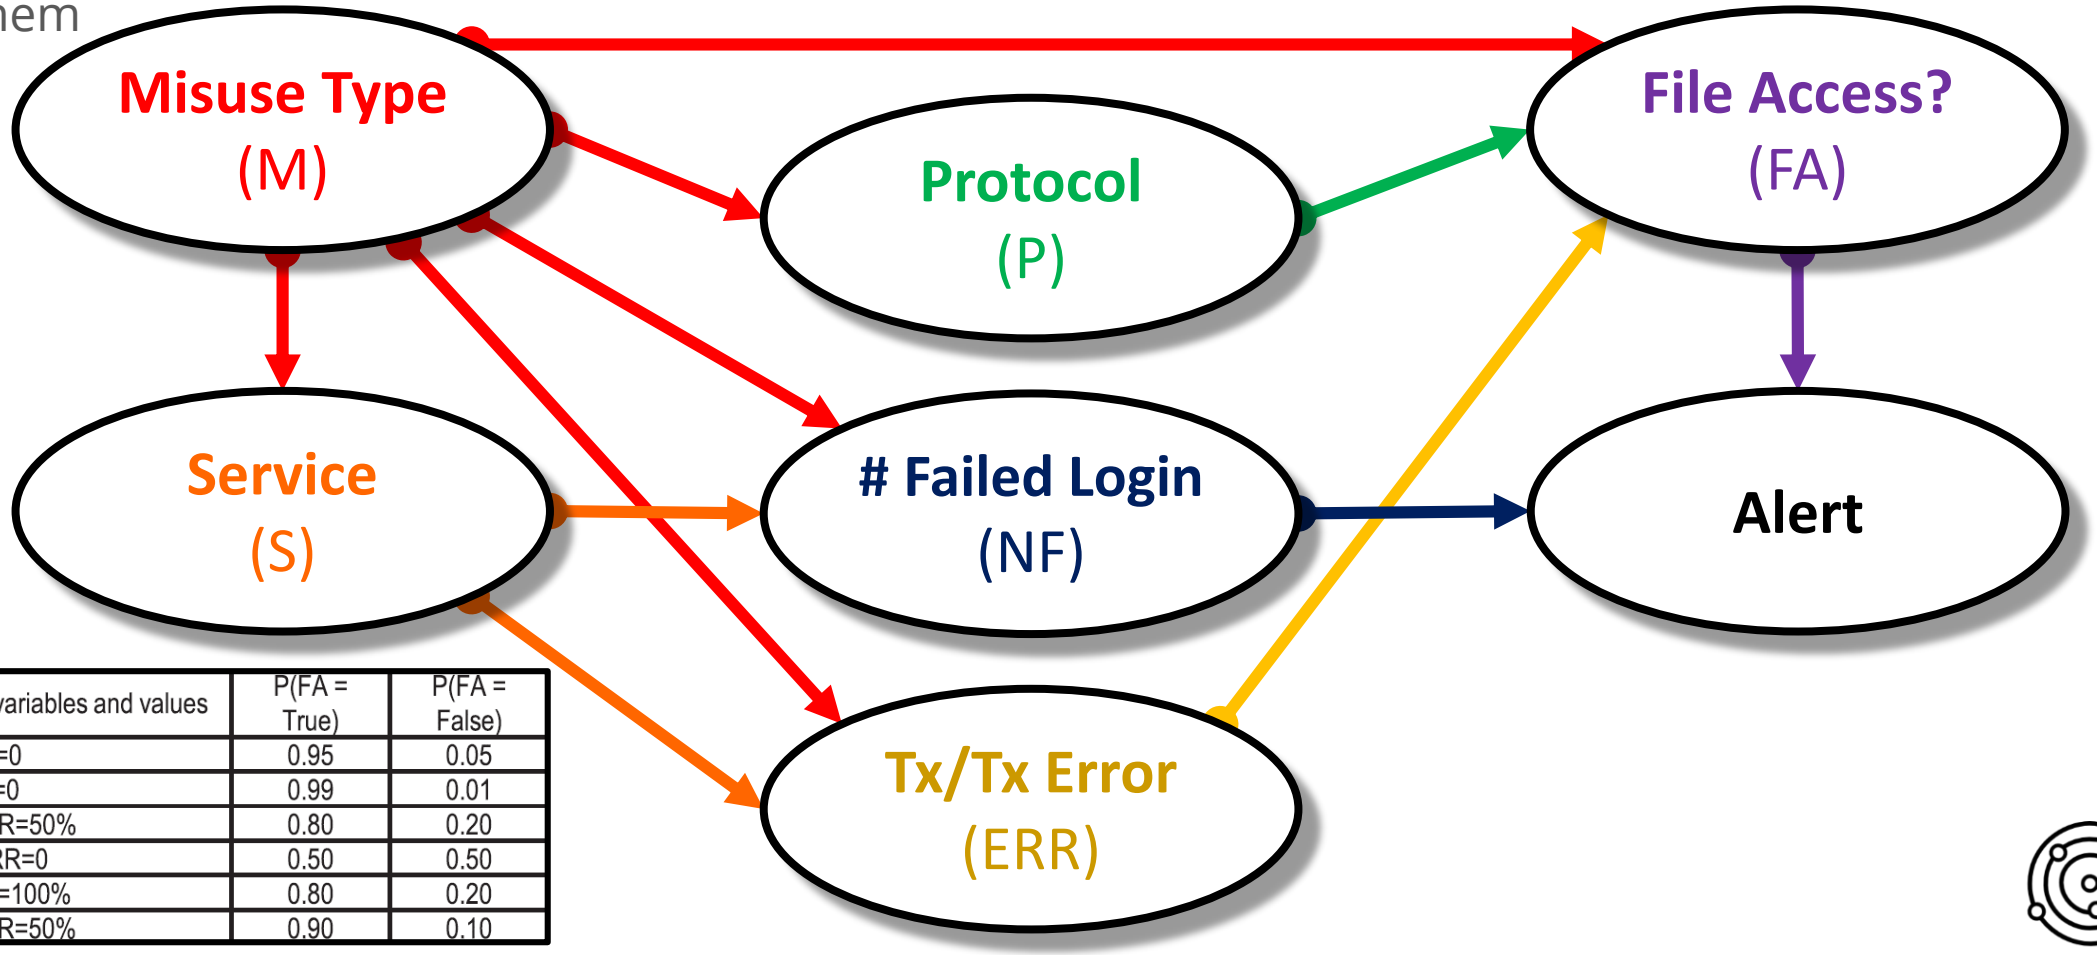

File Access state input variables and values	P(FA = True)	P(FA = False)
M=R2H, PT=NSF, ERR=0	0.95	0.05
M=R2H, PT=FTP, ERR=0	0.99	0.01
M=Probe, PT=none, ERR=50%	0.80	0.20
M=Probe, PT=PING, ERR=0	0.50	0.50
M=DoS, PT=POP, ERR=100%	0.80	0.20
M= DoS, PT=HTTP, ERR=50%	0.90	0.10

Source: : "A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection."

ML Methods: Clustering

Grouping a set of objects in such a way that objects in the same group (called a 'cluster') are **more similar** (in some sense) to each other **than to those in other groups**.

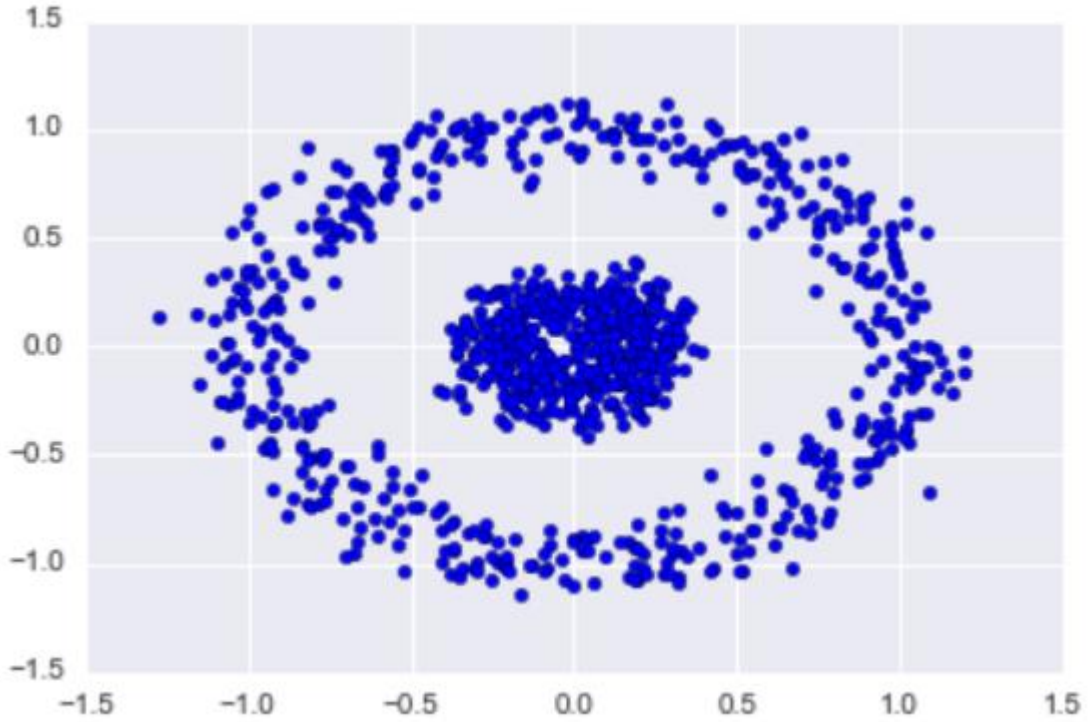

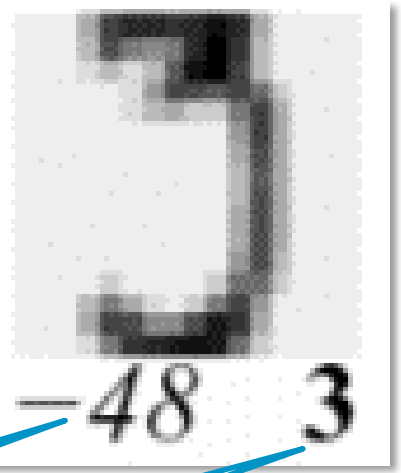

Novelty

Class

Source: https://en.wikipedia.org/wiki/Cluster_analysis, Introduction to Machine Learning, https://github.com/sinanuozdemir/sfdat28/blob/master/notebooks/06_logistic_regression.ipynb

Just Covering the Basics

Source: https://jixta.wordpress.com/2015/07/17/machine-learning-algorithms-mindmap/

And Your Point...?

Neural Networks

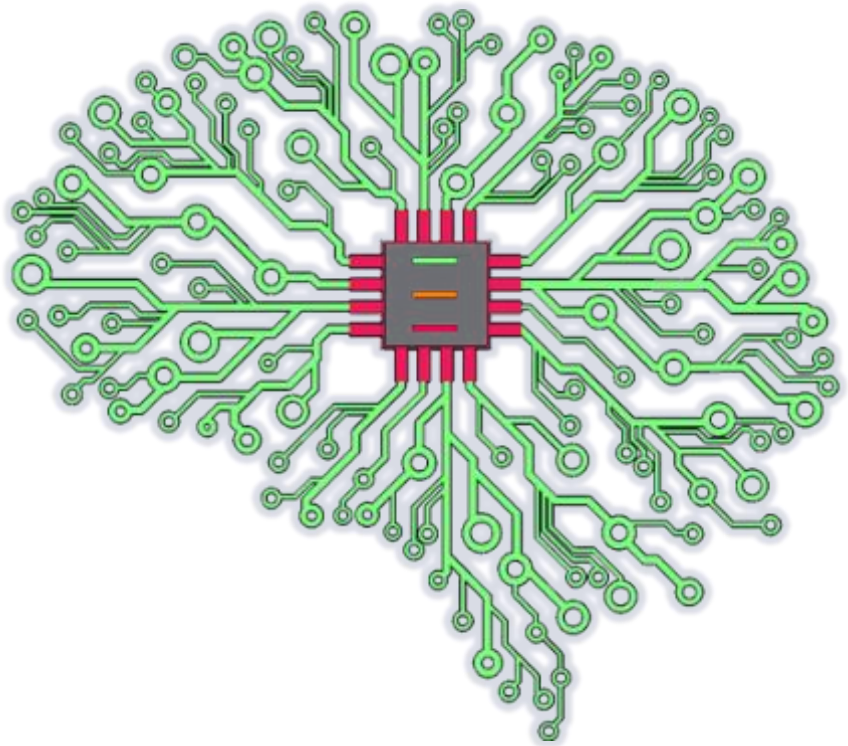

Automated Neural Networks (ANN)

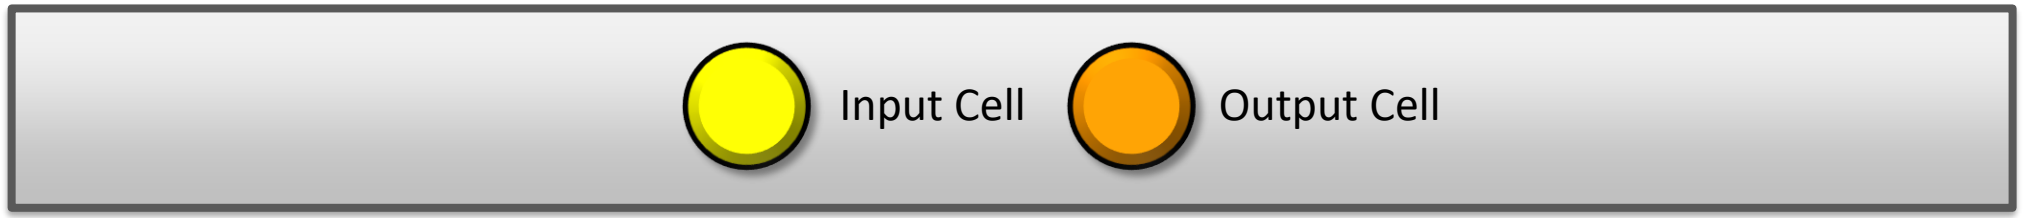

What House to buy?

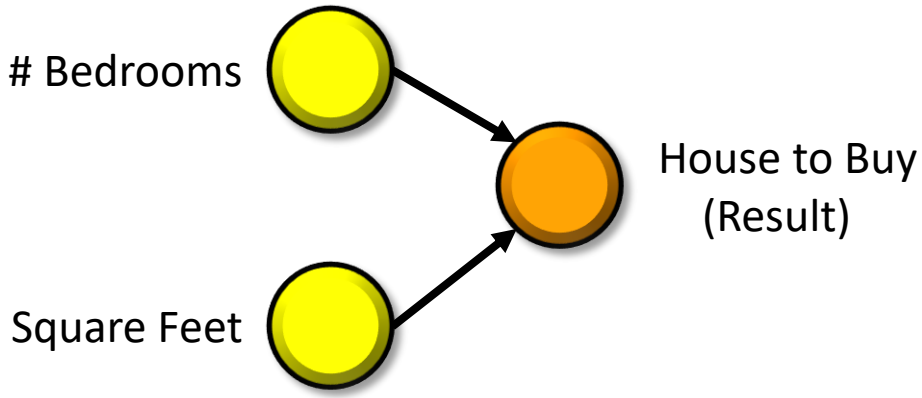

Source: https://towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464

Neural Network Example

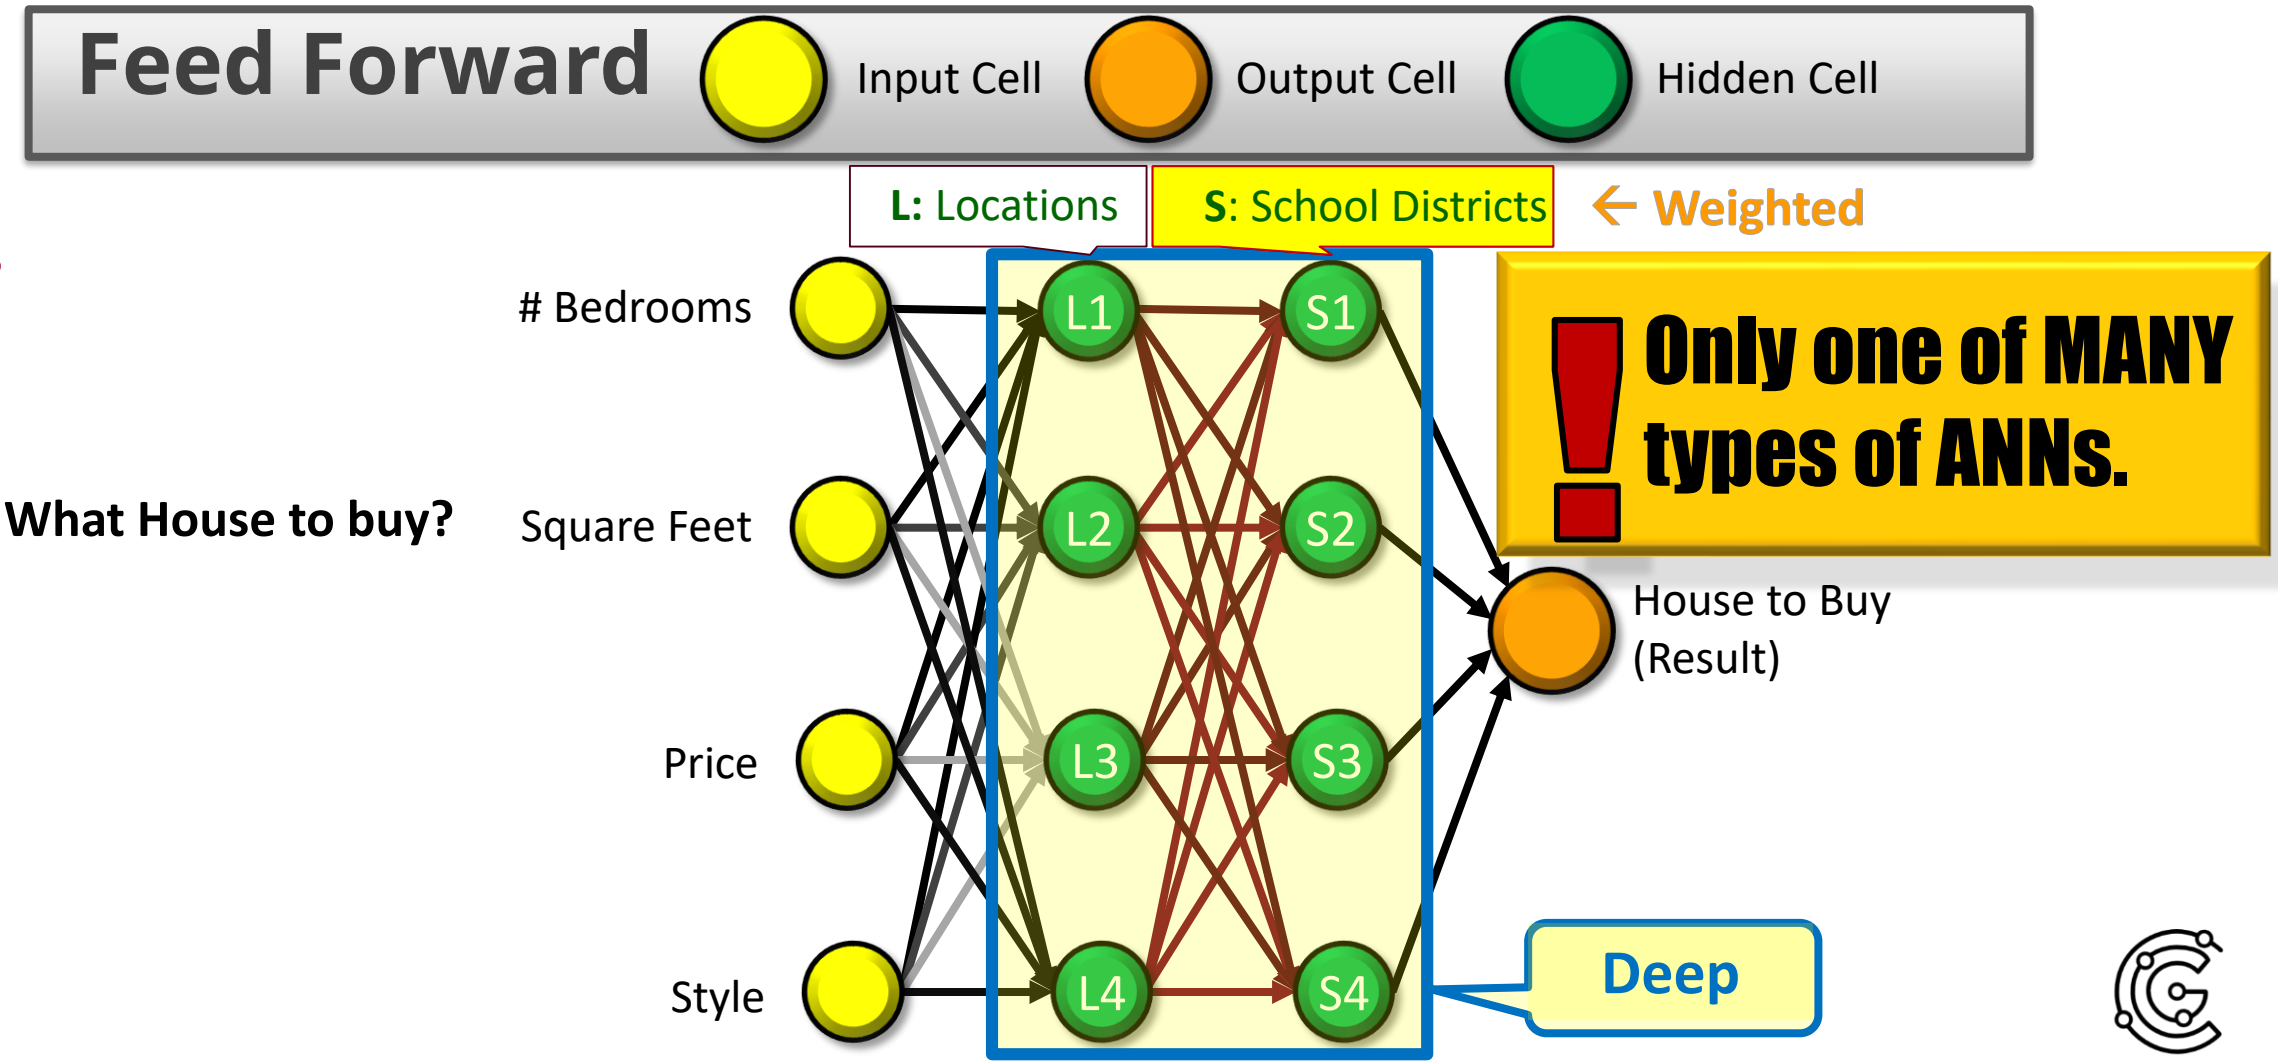

Source: https://towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464

Recap: AL & ML Terms

- **Artificial Intelligence:**
  When Machines act like a real person; “General” AI
  (In all aspects)
- **Machine Learning:**
  Uses algorithms to predict patterns; all current AI is “Specific”
  (Type of AI – limited application [like a savant])
- **Deep Learning**
  Algorithm that uses Neural Networking
  (Type of ML, multiple levels of variables)

Pitfalls

Training Errors

Getting good training data can be freaking HARD.

Underfitting

Model performs poorly

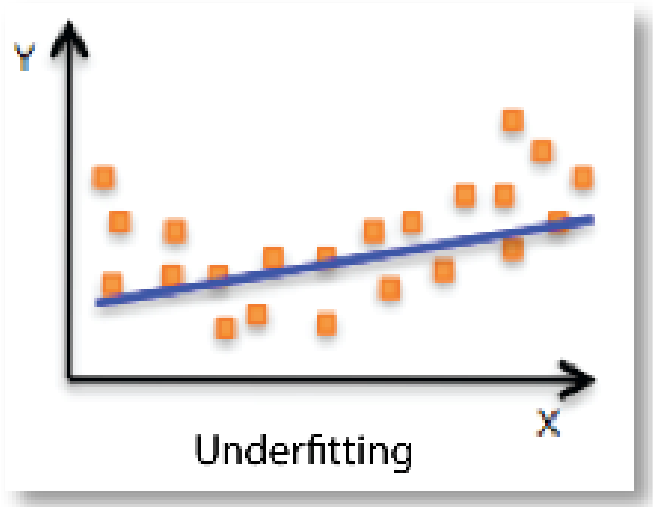

High bias, low variance

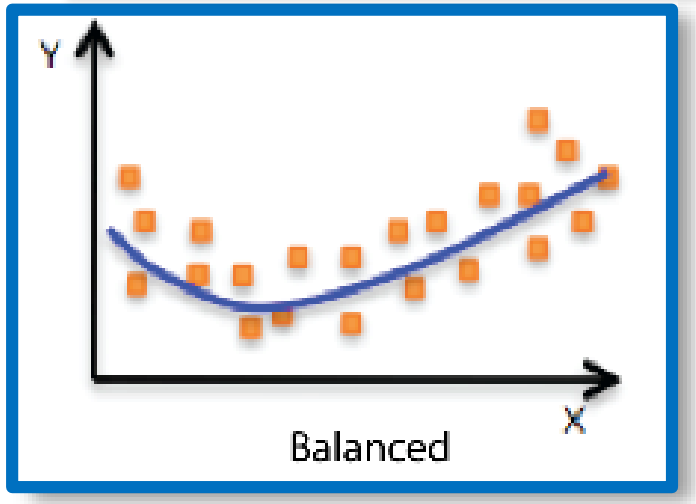

Still some false positives

Overfitting

Model performs too well

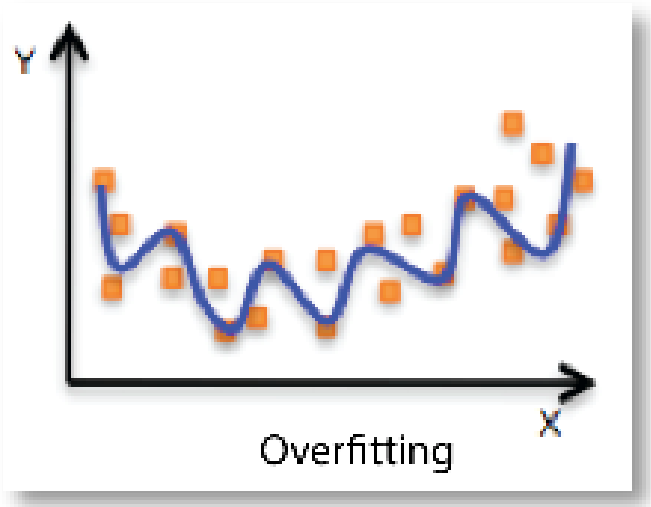

Low bias, high variance

- *More features*
- *Decrease data regularization*

- *Reduce feature count*
- *Trim/normalize data*

Source: https://docs.aws.amazon.com/machine-learning/latest/dg/model-fit-underfitting-vs-overfitting.html

Manipulation Biases

X Mean: 54.2659224 **X SD:** 16.7649829 **Corr.:** -0.0642526
Y Mean: 47.8313999 **Y SD:** 26.9342120

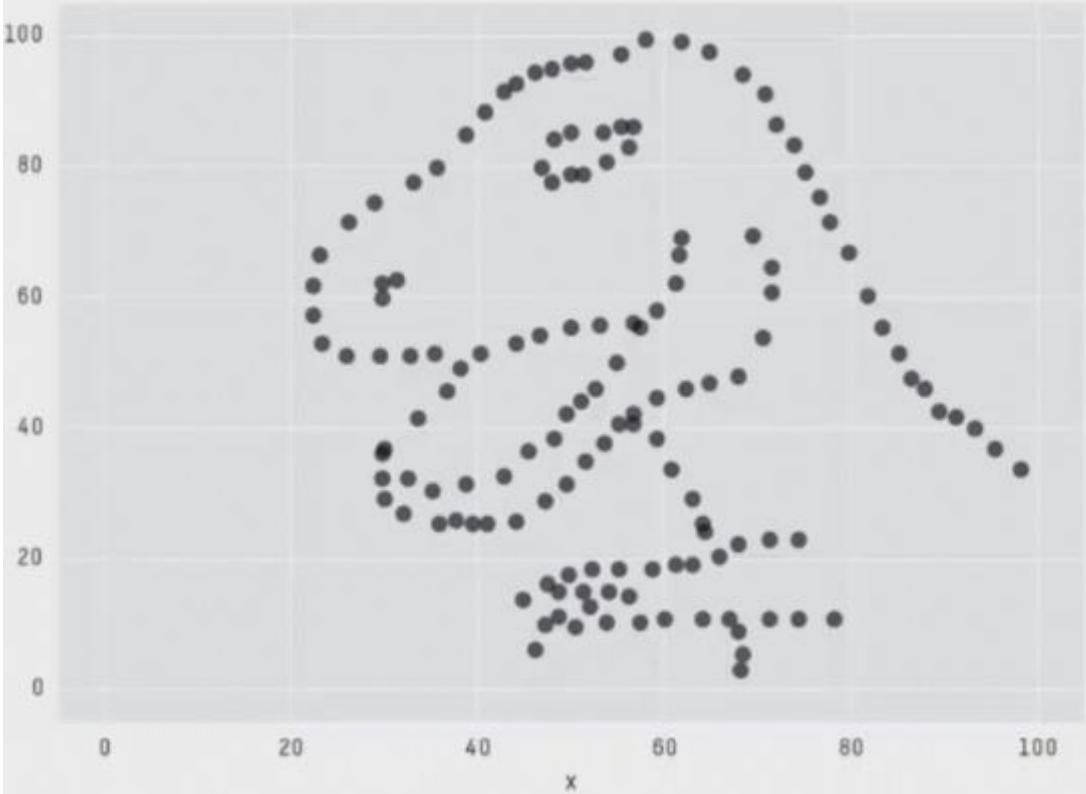

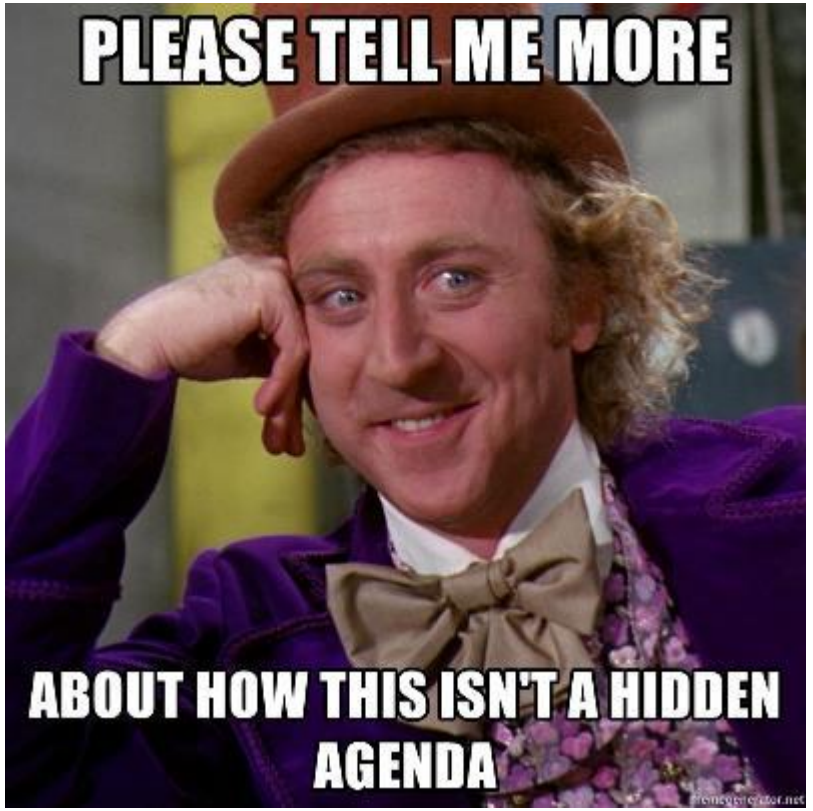

Source: https://github.com/stephlocke/lazyCDN/blob/master/DinoSequential.gif?raw=true

Manipulation Biases

Clustering Example

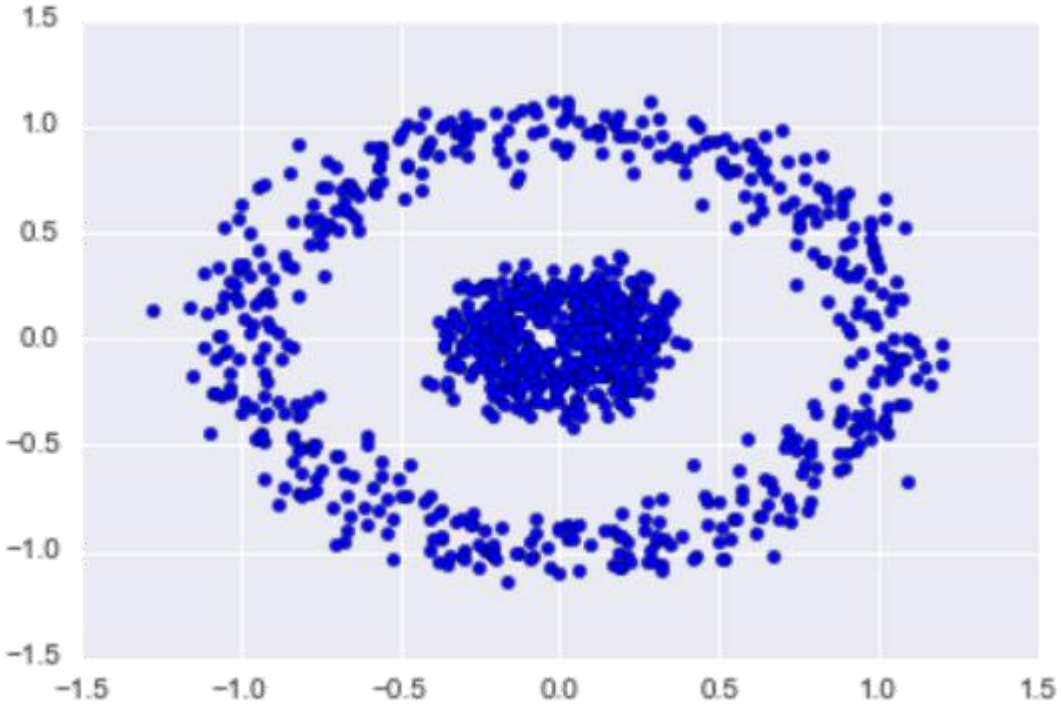

Source: https://github.com/sinanuozdemir/sfdat28/blob/master/notebooks/06_logistic_regression.ipynb,

Not all Algorithms are the Same

Parametric

Circle classification Logistic Regression

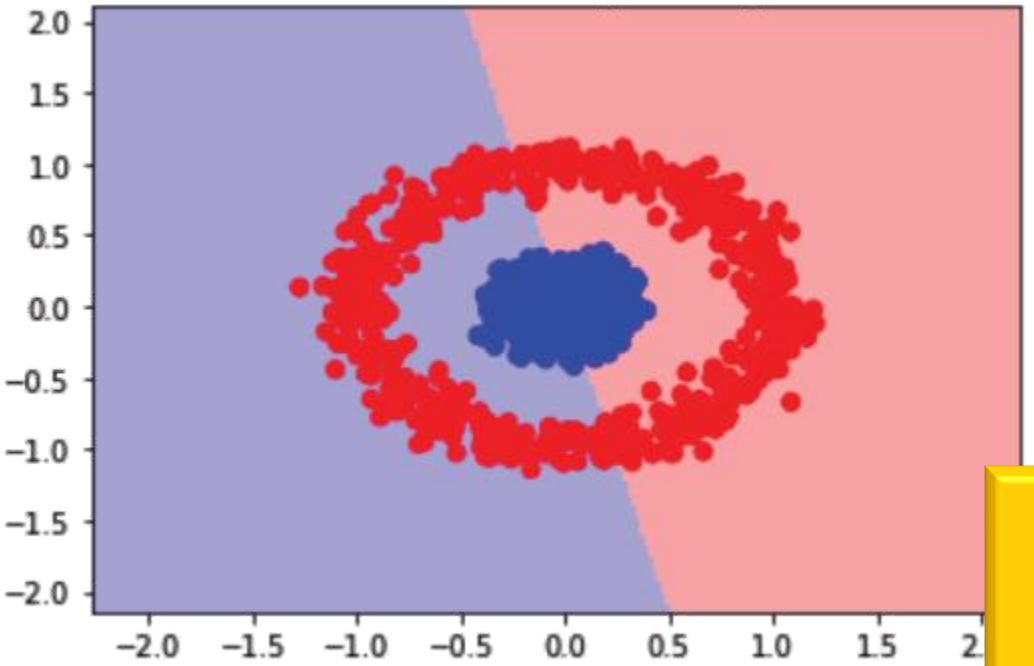

Assumes normally distributed data

Non-Parametric

Circle classification (k=9)

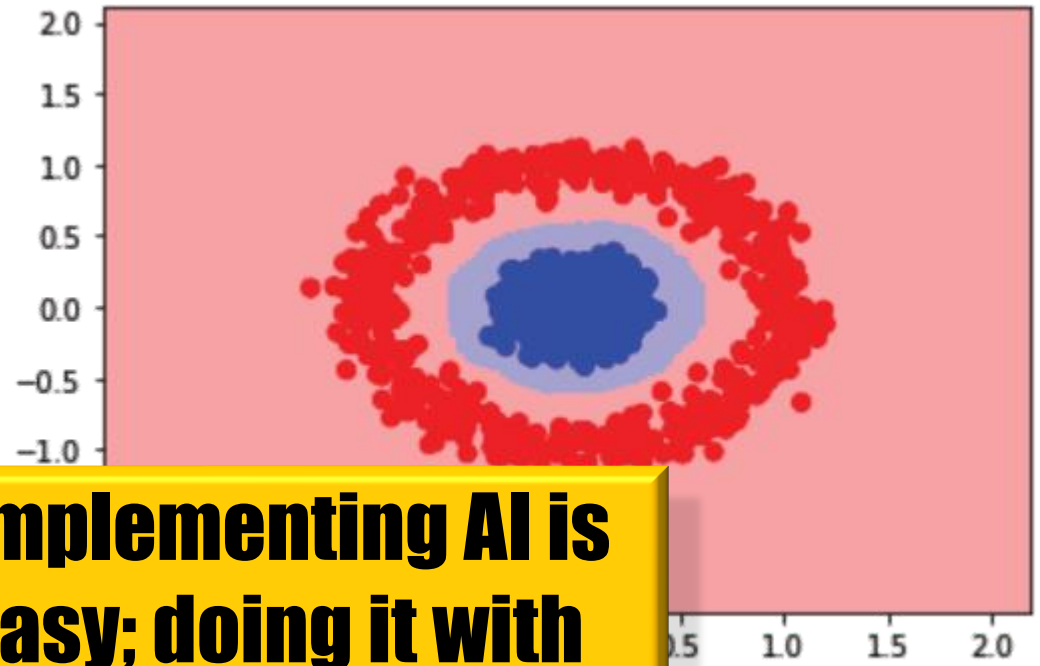

G1

G2

Implementing AI is easy; doing it with intelligence is not.

Source: https://github.com/sinanuozdemir/sfdat28/blob/master/notebooks/06_logistic_regression.ipynb,

Neural Network Assumptionns

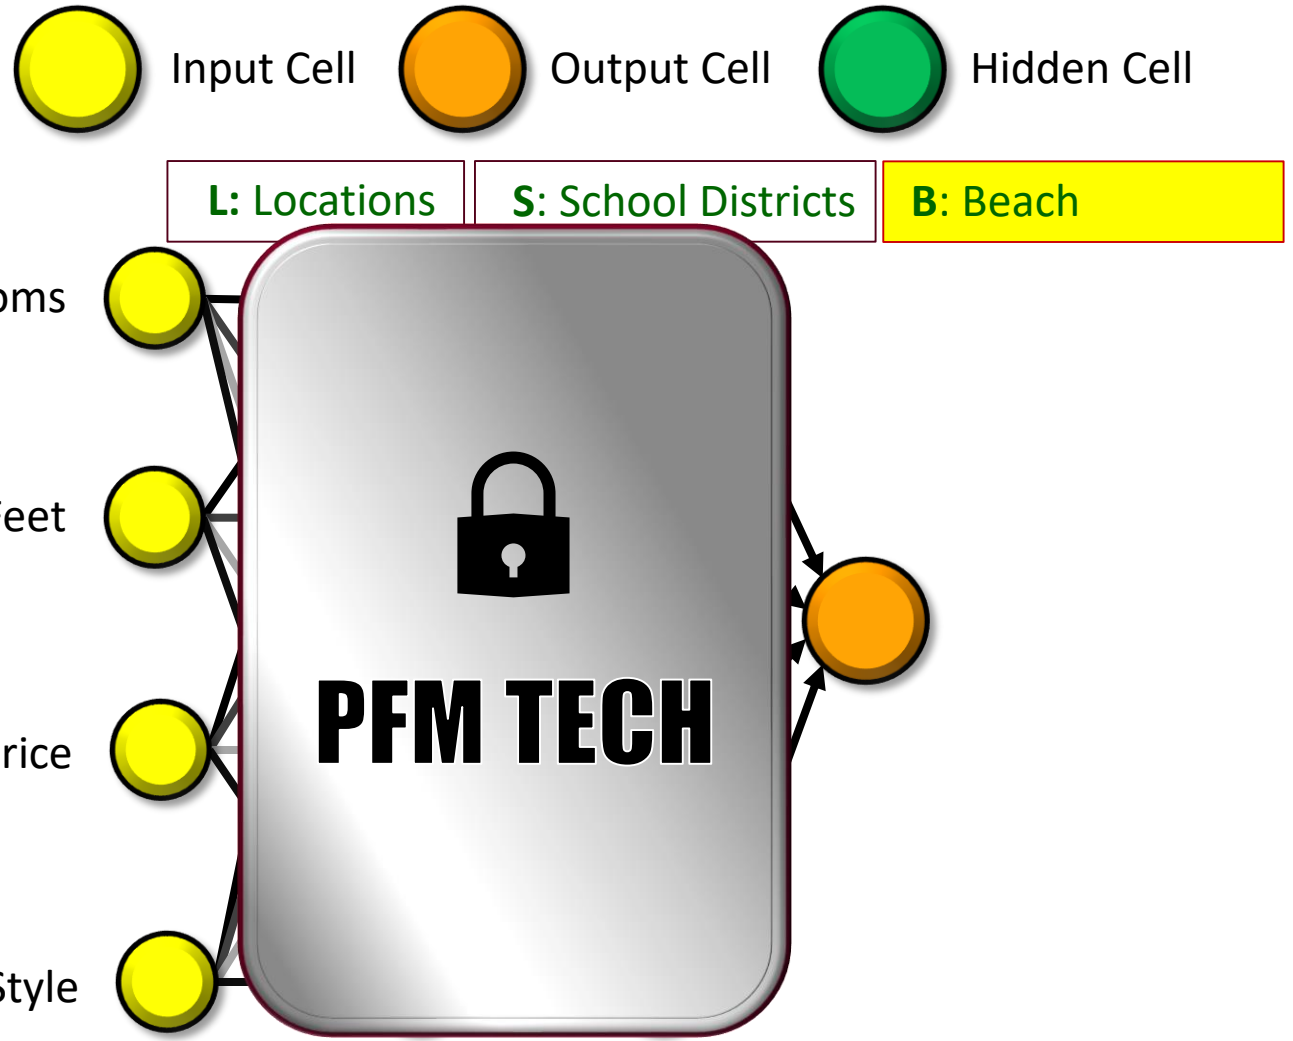

What House to buy?

”…many times, organizations have a **lack of control** over the AI output and outcome.”

- *Matt Sanchez*
  (CTO and co-founder of CognitiveScale)

Source: https://towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464

More “A” than “I”...

Malicious Code

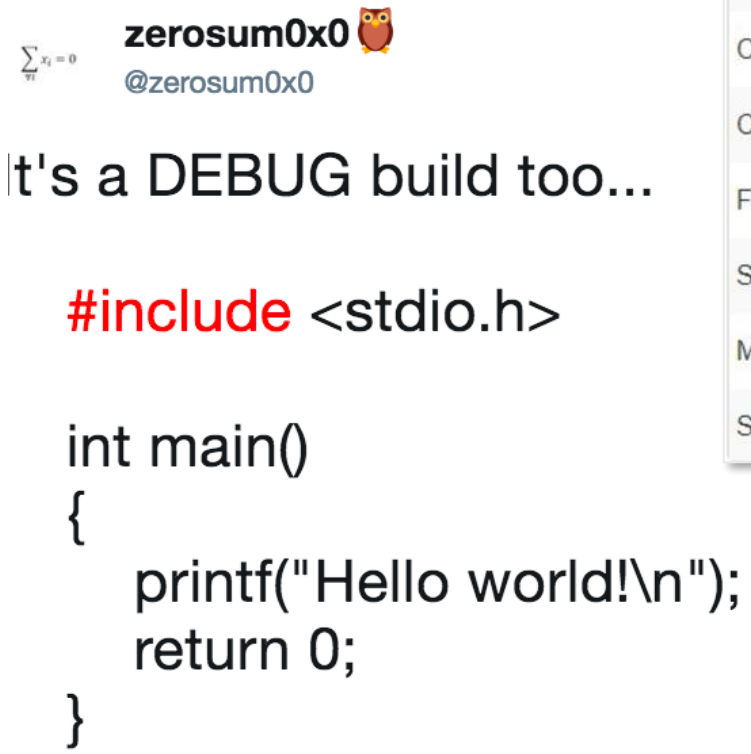

Antivirus	Result
CrowdStrike Falcon (ML)	malicious_confidence_80% (D)
Cylance	Unsafe
Cyren	W32/S-d2b5872a!Eldorado
F-Prot	W32/S-d2b5872a!Eldorado
Sophos ML	heuristic
McAfee-GW-Edition	BehavesLike.Win32.Trojan.nt
SentinelOne (Static ML)	static engine - malicious

Source: https://arxiv.org/abs/1412.6572ificial-intelligence, “Deep Learning with Python”

AI-n't Perfect

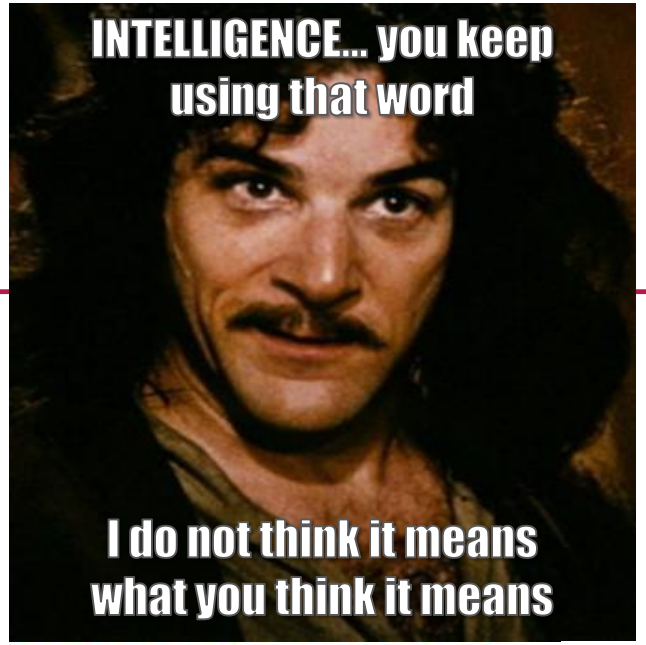

"panda"
57.7% confidence

"gibbon"
99.3 % confidence

Source: https://www.theverge.com/2017/4/12/15271874/ai-adversarial-images-fooling-attacks-artificial-intelligence

AI-n't Perfect: Chihuahua or chocolate chip muffin?

http://imgur.com/a/K4RWn

AI-n't Perfect: Labradoodle or fried chicken?

http://imgur.com/a/K4RWn

AI-n't Perfect: Sheepdog or Mop?

http://imgur.com/a/K4RWn

AI is easily confused

≠ Intuition

≠ Instinct

≠ ‘6th Sense’

≠ Morality

Abstraction Modeling

The Problem with Abstraction

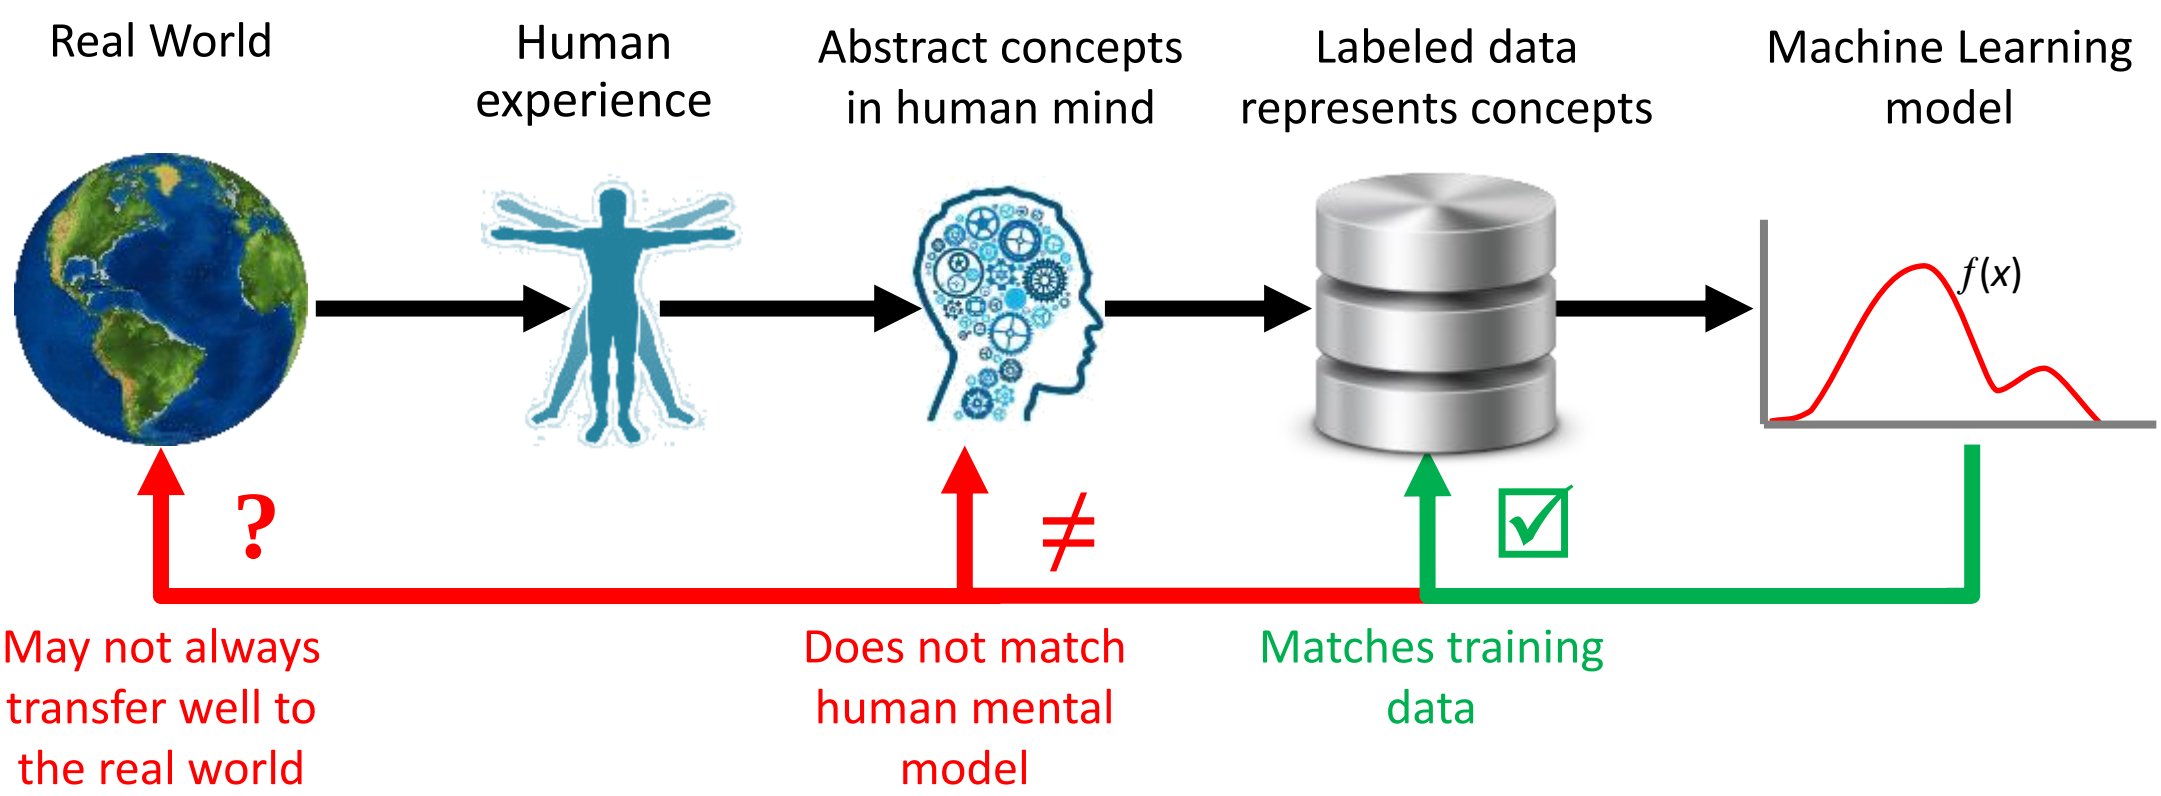

Source: “Deep Learning with Python”

AI Failures: Real-World

First death due to self-driving car

March 18, 2018

Source: https://www.nytimes.com/2018/03/23/technology/uber-self-driving-cars-arizona.html

Pitfalls: Summary

- Training
  - Test data
  - Over & under fitting
- Abstraction
- Bias & hidden weighting
- Data assumptions
- Errors
- Perturbation

”…**we have to bias our algorithms** so that you never trust any one individual or any one team. It is a careful(ly) controlled dance to build these types of systems to produce general purpose, general results that applies to all organizations.”

-Greg Martin, JASK (jask.ai)

If ‘cloud computing’ is just someone else’s data center, most Machine Learning is just someone else’s assumptions.

Cyber Applications

White Hat AI

> "…while AI systems can exceed human performance in many ways, they **can also fail in ways that a human never would.**"
>
> *-The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*

- Network Management
- Data: Visualization, Log patterns, UBA
- First-Level SOC analysis?
- Augment (not replace) the Human
- Reverse Engineering (GHIDRA)
- IoT = ANN 'sense' organs

GHIDRA

It **IS NOT** a silver bullet!

Source: The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation
https://www.wired.com/story/nsa-ghidra-open-source-tool/

White Hat AI: 5 Questions to Ask

1) Technical Components
2) Flexibility
3) Applications
4) AI/ML Updates
5) Your Security Team's Skillset

Source: https://www.darkreading.com/vulnerabilities---threats/cutting-through-the-jargon-of-ai-and-ml-5-key-issues/a/d-id/1333595

Black Hat AI

> "AI systems and the knowledge of how to design them can be put toward both … beneficial and harmful ends … **artificial intelligence is dual-use in the same sense that human intelligence is**."
>
> *-The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*

- Data Poisoning
- Scales the Attack (#, speed, & targets)
- Discover New Attack Vectors – *FAST*
- Exploit AI Vulnerabilities
- Increase anonymity & psychological distance

Including physical, voice, images…

Unique Data

Source: The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation

AI/ML Musts

- AI must be able to explain '*why*' (for Audit and Compliance)
- Plan for & design continuous learning feedback loops
- Governance processes = assurance/guardrails for AI insights & recommendations

Source: https://dzone.com/articles/artificial-intelligence-will-automate-business-pro

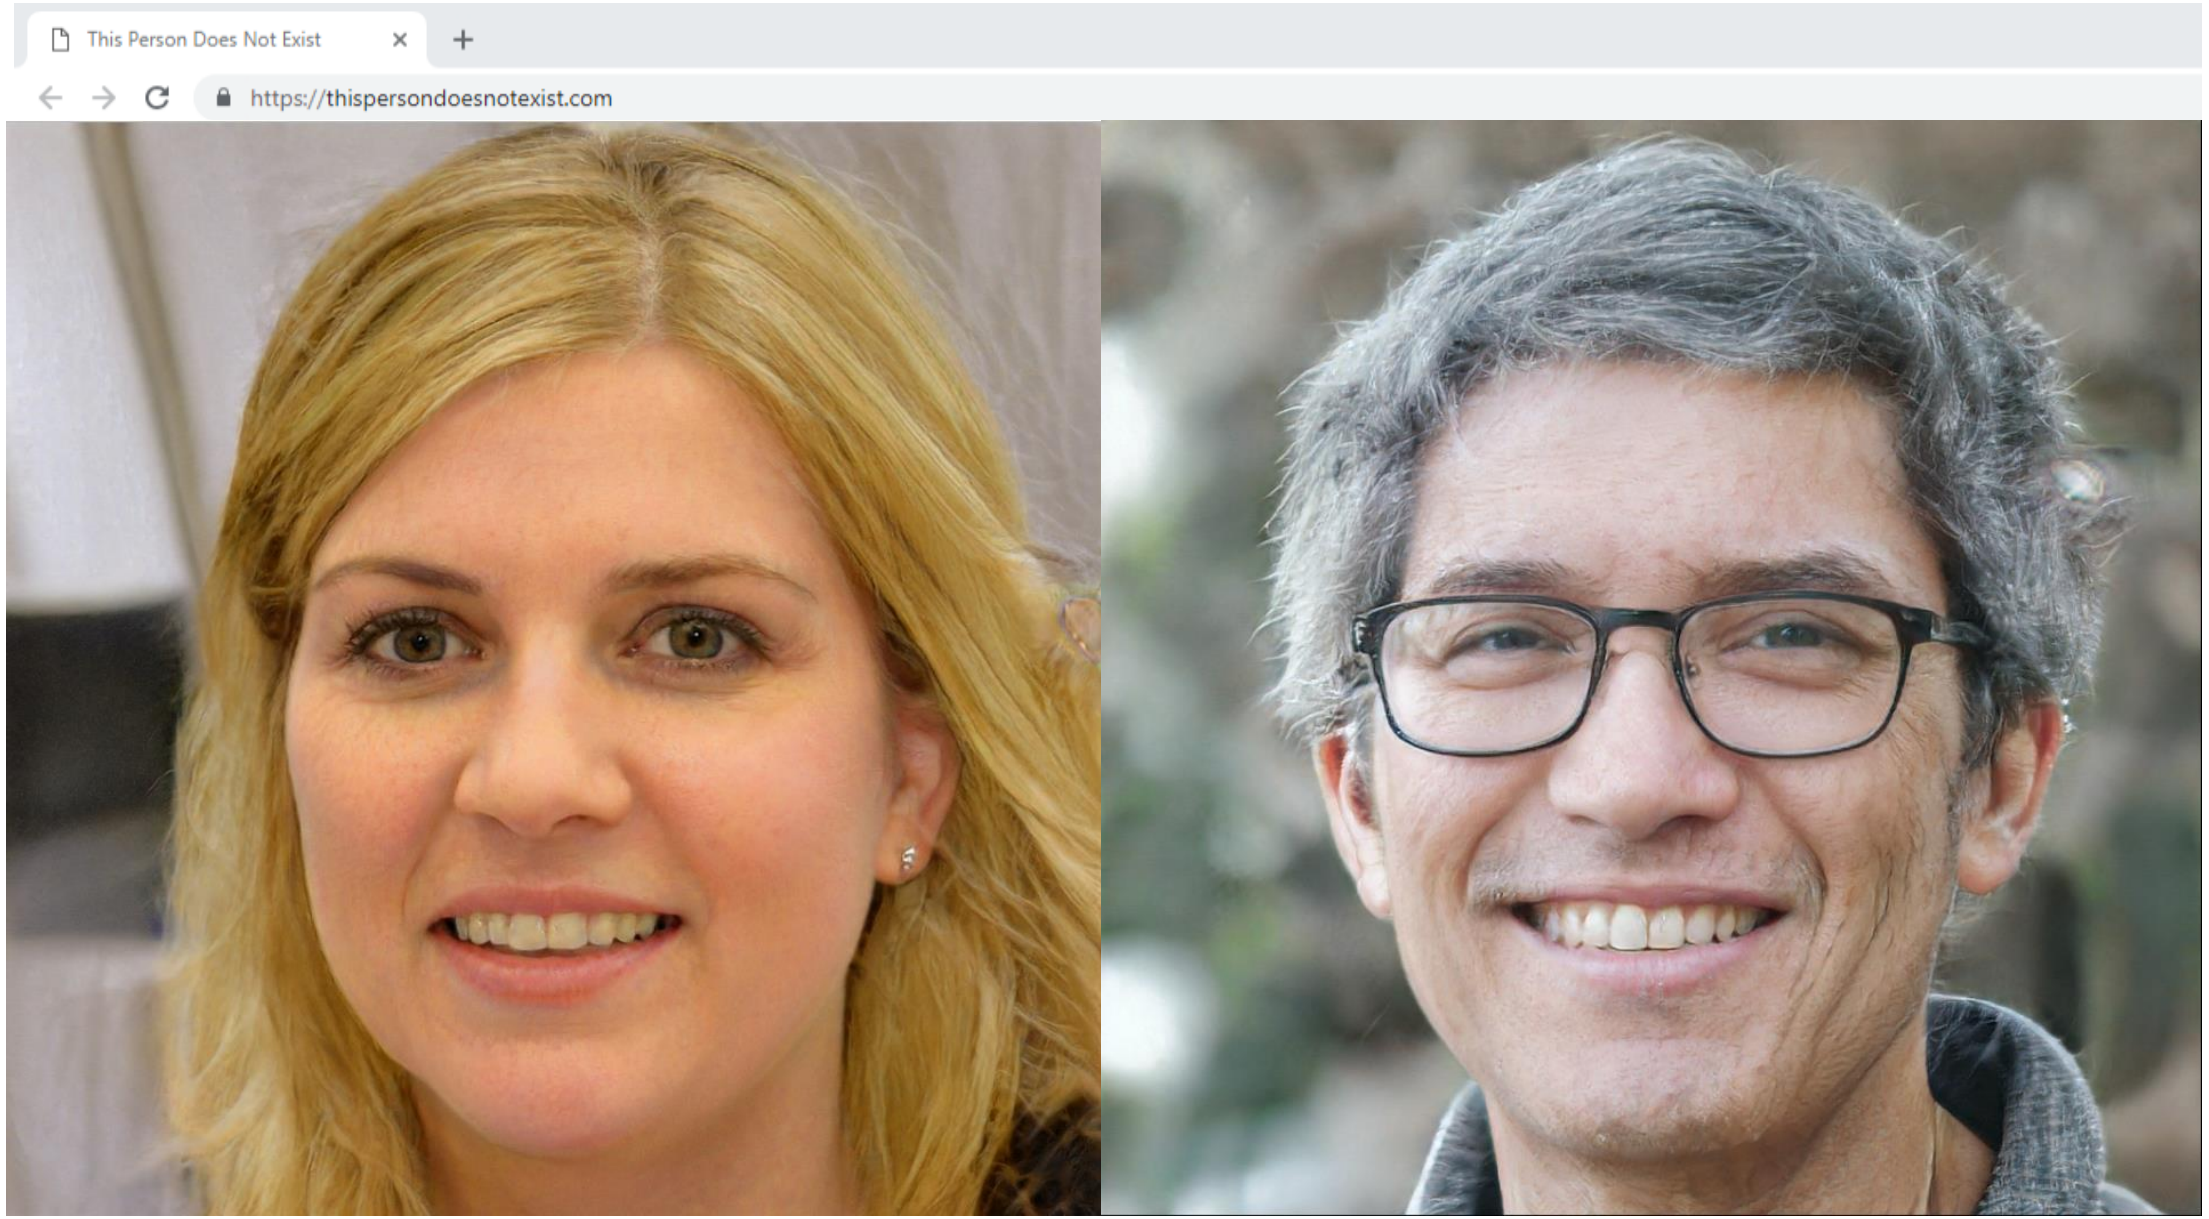

Source: https://thispersondoesnotexist.com

Cyber Applications: Summary

- Good at similar, predictable data
- Good at sifting data w/ patterns
- Assist – *not* replace – the human
- Speeds everything up – including bad guys & how fast we break crap
- New & faster exploits
- Increase of psychological distance

In Review...

Review: What We Covered Today

1) Definitions
2) Machine Learning
   - Types
   - Methods
3) Neural Networks
4) Pitfalls
   - Errors
   - Assumptions
   - Biases
5) Cyber Applications

References

Ayodele, Taiwo Oladipupo. “Types of Machine Learning Algorithms.” University of Portsmouth, http://www.intechopen.com/books/new-advances-in-machinelearning.

Benjamin, Paul. “US7784099B2 - System for Intrusion Detection and Vulnerability Assessment in a Computer Network Using Simulation and Machine Learning.” *Google Patents*, Google, 18 Feb. 2005, patents.google.com/patent/US7784099B2/en.

Brundage, M. , Avin, S. , Clark, J. , Toner, H. , Eckersley, P. , Garfnkel, B. , Anderson, H. , Flynn, C. , Farquhar, S. , Page, M. , Dafoe, A. , Roff, H. , Ó hÉigeartaigh, S. , Lyle, C. , Bryson, J. , Scharre, P. , Allen, G. , Beard, S. , Yampolskiy, R. , Zeitzoff, T. , Steinhardt, J. , Belfeld, H. , Evans, O. , Amodei, D. , Filar, B. . “The Malicious Use of Artificial Intelligence: Forecasting Prevention and Mitigation”, https://maliciousaireport.com, Feb. 2018.

Buczak, Anna L., and Erhan Guven. “A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection.” *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, 26 Oct. 2015, pp. 1153–1176., http://ieeexplore.ieee.org/document/7307098/.

Chollet, Francois. “Deep Learning with Python.” Manning Publications, Chapter 9 section 2, Nov 2017, https://blog.keras.io/the-limitations-of-deep-learning.html

Dzone. “The DZone guide to Artificial Intelligence: Machine Learning and Artificial Intelligence volume 1”. https://dzone.com/articles/10-enterprise-machine-learning-predictions-for-201

Levy, Brian. “How will AI and machine learning impact CSPs?”. August 31, 2017, https://inform.tmforum.org/data-analytics-and-ai/2017/08/will-ai-machine-learning-impact-csps/

References

Nicholson, Chris V., Gibson, Adam, Skymind team. “Introduction to Deep Neural Networks.” Deeplearning4j: Open-Source, Distributed Deep Learning for the JVM”, deeplearning4j.org/neuralnet-overview.

Ozdemir, Sinan. Logistic Regression; https://www.linkedin.com/in/sinan-ozdemir.

Omerisk, Joh. “Cutting Through the Jargon of AI & ML: 5 Key Issues.” https://www.darkreading.com/vulnerabilities---threats/cutting-through-the-jargon-of-ai-and-ml-5-key-issues/a/d-id/1333595

Richards, Ken. Machine Learning: For Beginners - Your Starter Guide For Data Management, Model Training, Neural Networks, Machine Learning Algorithms: Volume 1.

Smith, Tom. ‘Artificial Intelligence will Automate Business Processes’. Interview with Matt Sanchez, November 9, 2017. https://dzone.com/articles/artificial-intelligence-will-automate-business-pro

Smola, Alex and Vishwanathan, S.V.N.. Introduction to Machine Learning. Cambridge University Press, 2008. http://alex.smola.org/drafts/thebook.pdf

Tchircoff, Andrew. “The Mostly Complete Chart of Neural Networks, Explained.” Towards Data Science, Towards Data Science, 4 Aug. 2017, towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464.

Zeolla, Jon. “Cutting Through the Buzz: Machine Learning and Artificial Intelligence”. http://www.threeriversinfosec.com/wp-content/uploads/2017/07/2017-10-Cutting-Through-The-Buzz-Machine-Learning-and-AI.pdf; October 20, 2017. Video at https://www.youtube.com/watch?v=61qJnY9njgs

Have questions?


shayne.champion@conversantgroup.com


423.602.7789


https://www.linkedin.com/in/shaynechampion


@TNfoSec

Machine Learning

How Smart Can It Really Be?

This presentation will evaluate Machine Learning (ML), Deep Learning (DL), and Artificial Intelligence (AI) as used within cyber security. During the session we will explore the difference between ML, DL, and AI, and show how these technologies work - as well as their shortcomings. Finally, we will discuss how these tools could work to help reduce risk and how to apply them in your security environment.

- Definitions
- Machine Learning
  - Types
  - Methods
- Neural Networks
- Errors
  - Assumptions
  - Biases
  - Pitfalls
- Cyber Applications