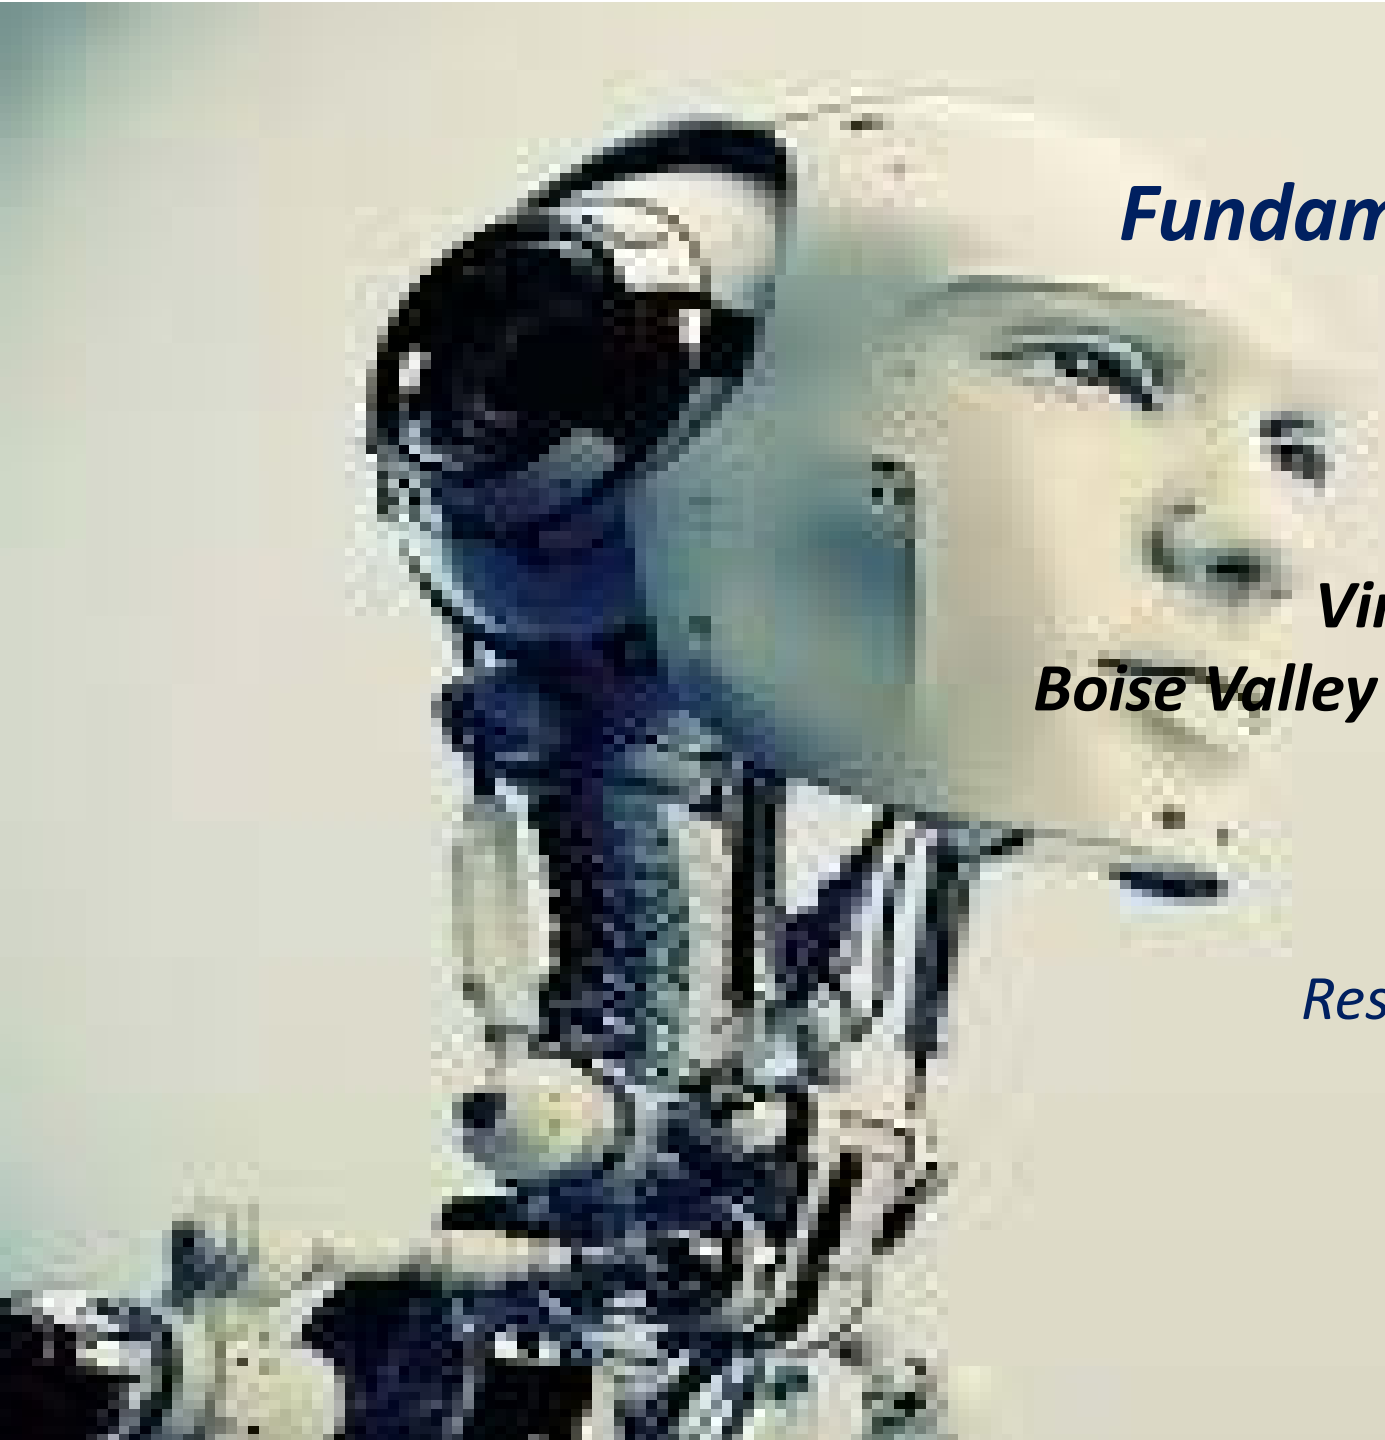

Fundamentals of Artificial Intelligence for Non-Computer Scientists

Virtual Coffee with Zasio Consulting and Boise Valley Idaho Chapter of ARMA International

March 10, 2022

Maura R. Grossman, J.D., Ph.D.
Research Professor @ University of Waterloo
Principal @ Maura Grossman Law

University of Waterloo

Maura Grossman Law

TODAY'S AGENDA

- What AI Is and Why Now
- How AI Works
- What AI Can Do
- Bias and Other Issues Implicated by AI

What AI Is and Why Now

WHAT IS "ARTIFICIAL INTELLIGENCE"?

- Umbrella term first used at Dartmouth in 1956
- Computers doing intelligent things (*i.e.*, performing cognitive tasks) once thought to be the sole province of humans
- Not a single technology or function
- Whatever computers can't do … until they can
- Called "software" after we get used to it
- Different from automation and robotics
- Generally involves algorithms, machine learning, and/or natural language processing ("NLP")

TYPES OF AI

Narrow (Weak) AI vs. General (Strong) AI

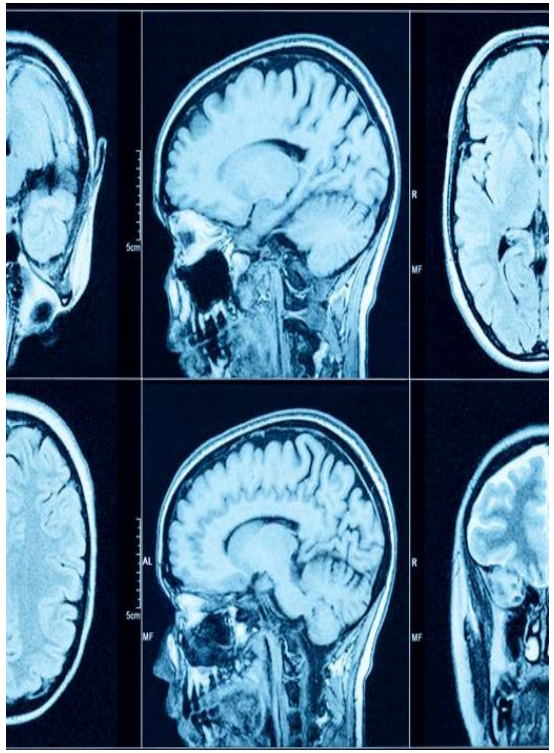

BUT ISN'T MACHINE LEARNING AN EMERGING TECHNOLOGY?

- Logistic Regression – 1944
- Stochastic Gradient Descent ("SGD") – 1950s
- Artificial Neural Network ("ANN") – 1958
- Hidden Markov Model – late 1950s/1960
- Decision Tree – 1975
- Expectation Maximization – 1977
- Q-Learning – 1989
- Support Vector Machine ("SVM") – 1990s
- K-Nearest Neighbor ("K-NN") – 1992
- Random Forest – 2001

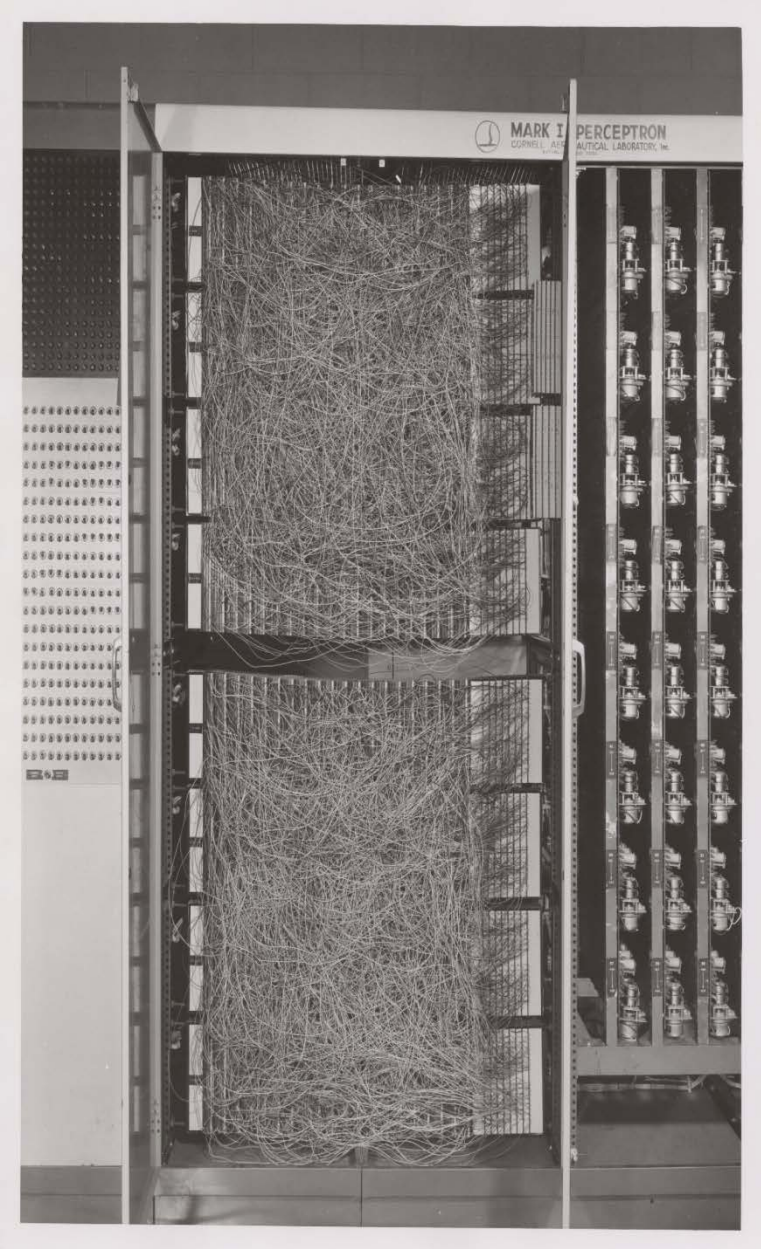

MORE COMPUTING POWER

Rosenblatt’s Mark I Perceptron
Cornell University, 1957

MORE DATA

KRYDER'S LAW

- **Purchase price for storage of one GB of data over 20 years:**
  - **1981 – $300,000**
  - **1987 – $50,000**
  - **1990 – $10,000**
  - **1994 – $1,000**
  - **1997 – $100**
  - **2000 – $10**
  - **2004 – $1**
  - **2011 – $0.10**
  - **2021 – Practically free**

MOORE'S LAW

- **The number of transistors on a microchip (and hence the speed of processing) doubles every 18 to 24 months**

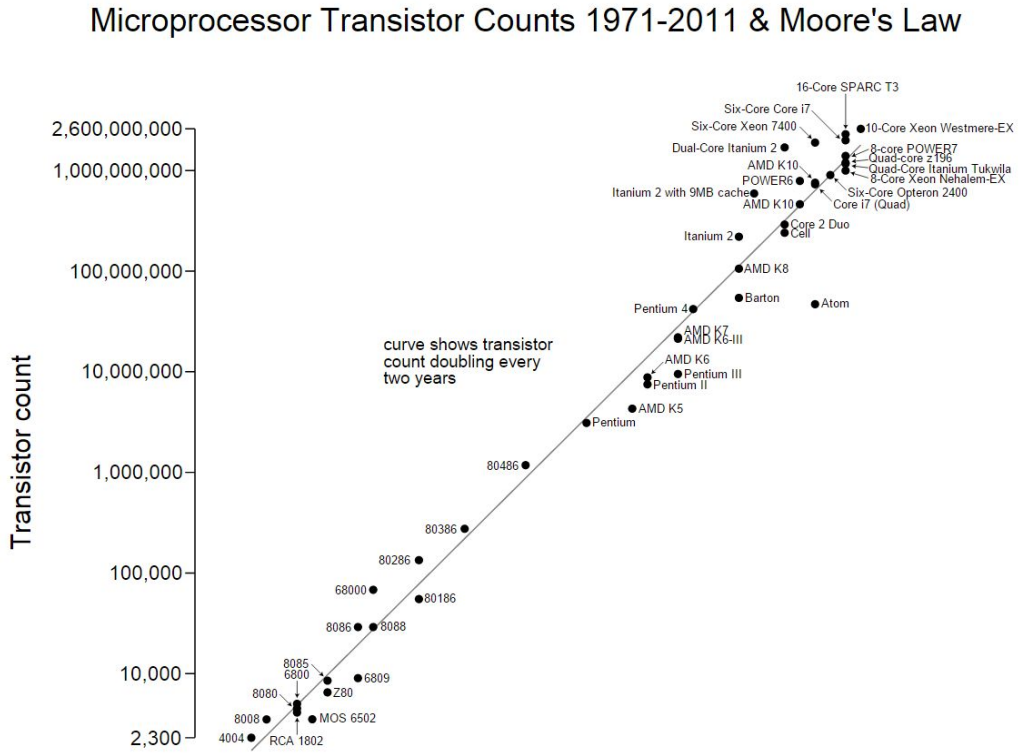

NEXT GENERATION HARDWARE

OPEN SOURCE COMMUNITIES LOWER BARRIERS TO ENTRY

Google open-sources mobile-first computer vision models for TensorFlow

BLAIR HANLEY FRANK @BELRIL JUNE 14, 2017 11:09 AM

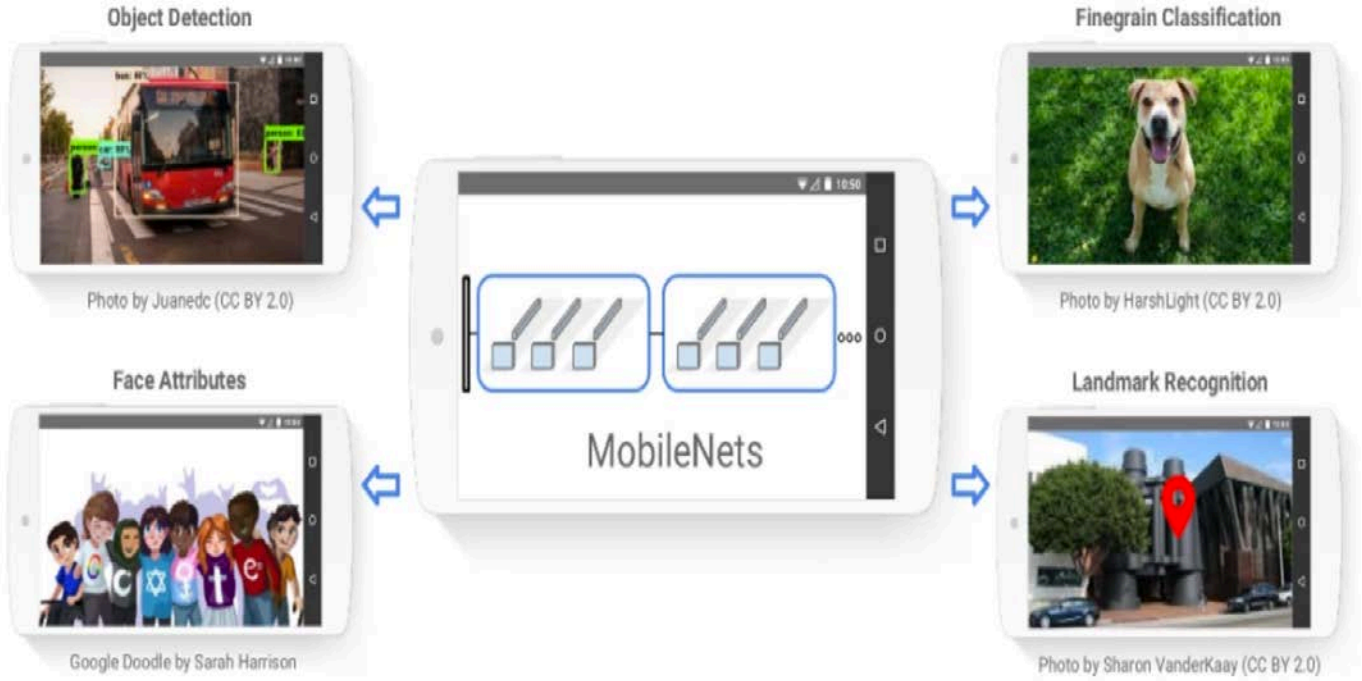

How AI Works

"Over the past decades computers have broadly automated tasks that programmers could describe with clear rules and algorithms. Modern machine learning techniques now allow us to do the same for tasks where describing the precise rules is much harder."

Jeff Bezos' 2017 Letter to Shareholders

UNSUPERVISED MACHINE LEARNING

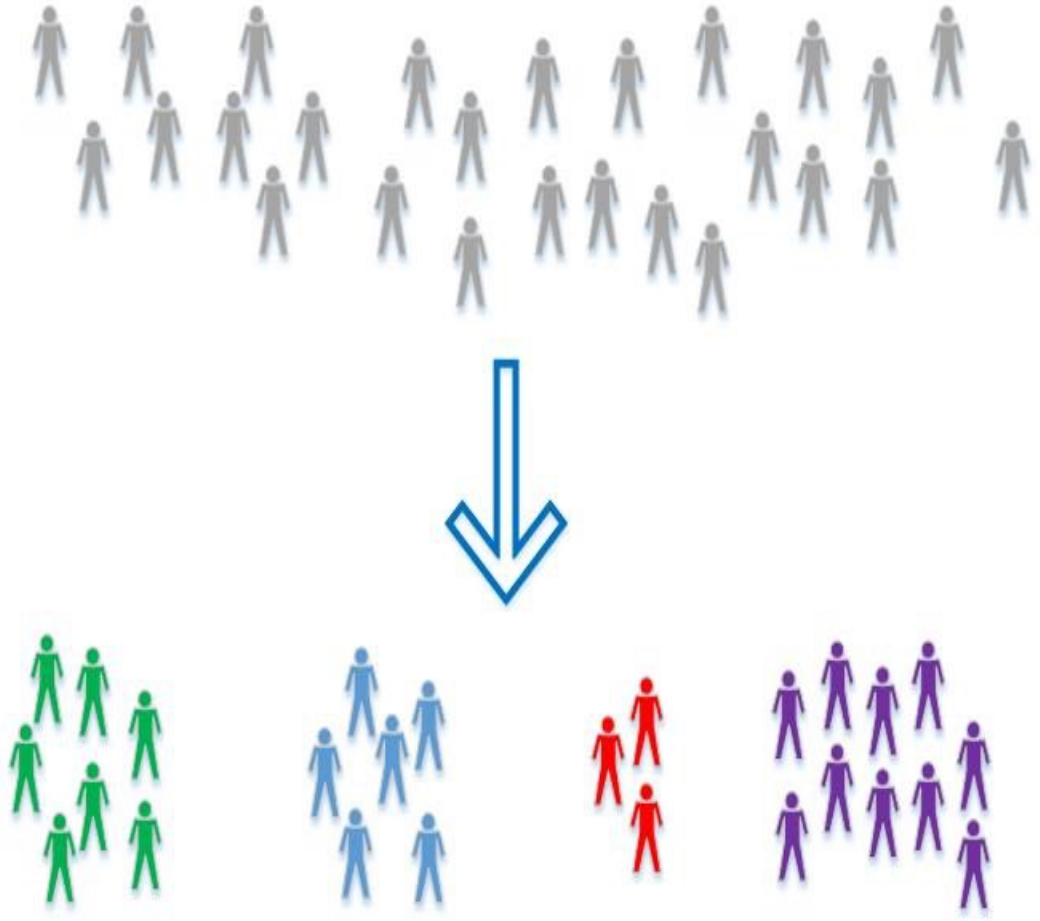

System looks for naturally occurring patterns, clusters, groupings, or anomalies

SUPERVISED MACHINE LEARNING

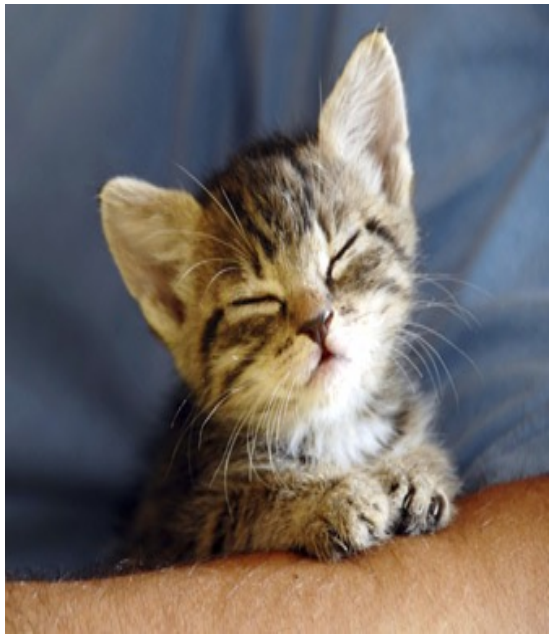

Training a system to distinguish between two or more categories by providing it with labeled examples from which it learns the rules

HOW SUPERVISED MACHINE LEARNING WORKS

Supervised machine learning systems infer mathematical functions from old data to help make educated guesses about new data.

JUST MEMORIZE THIS!

$$
\begin{cases}
\frac{\partial L_1}{\partial w} = 0 \quad \rightarrow \quad w = \sum_{i=1}^{N} \alpha_i y_i \phi(x_i), \\
\frac{\partial L_1}{\partial b} = 0 \quad \rightarrow \quad \sum_{i=1}^{N} \alpha_i y_i = 0, \\
\frac{\partial L_1}{\partial \xi_i} = 0 \quad \rightarrow \quad 0 \leq \alpha_i \leq c, \ i = 1, \ldots, N.
\end{cases}
$$

HOW SUPPORT VECTOR MACHINES WORK

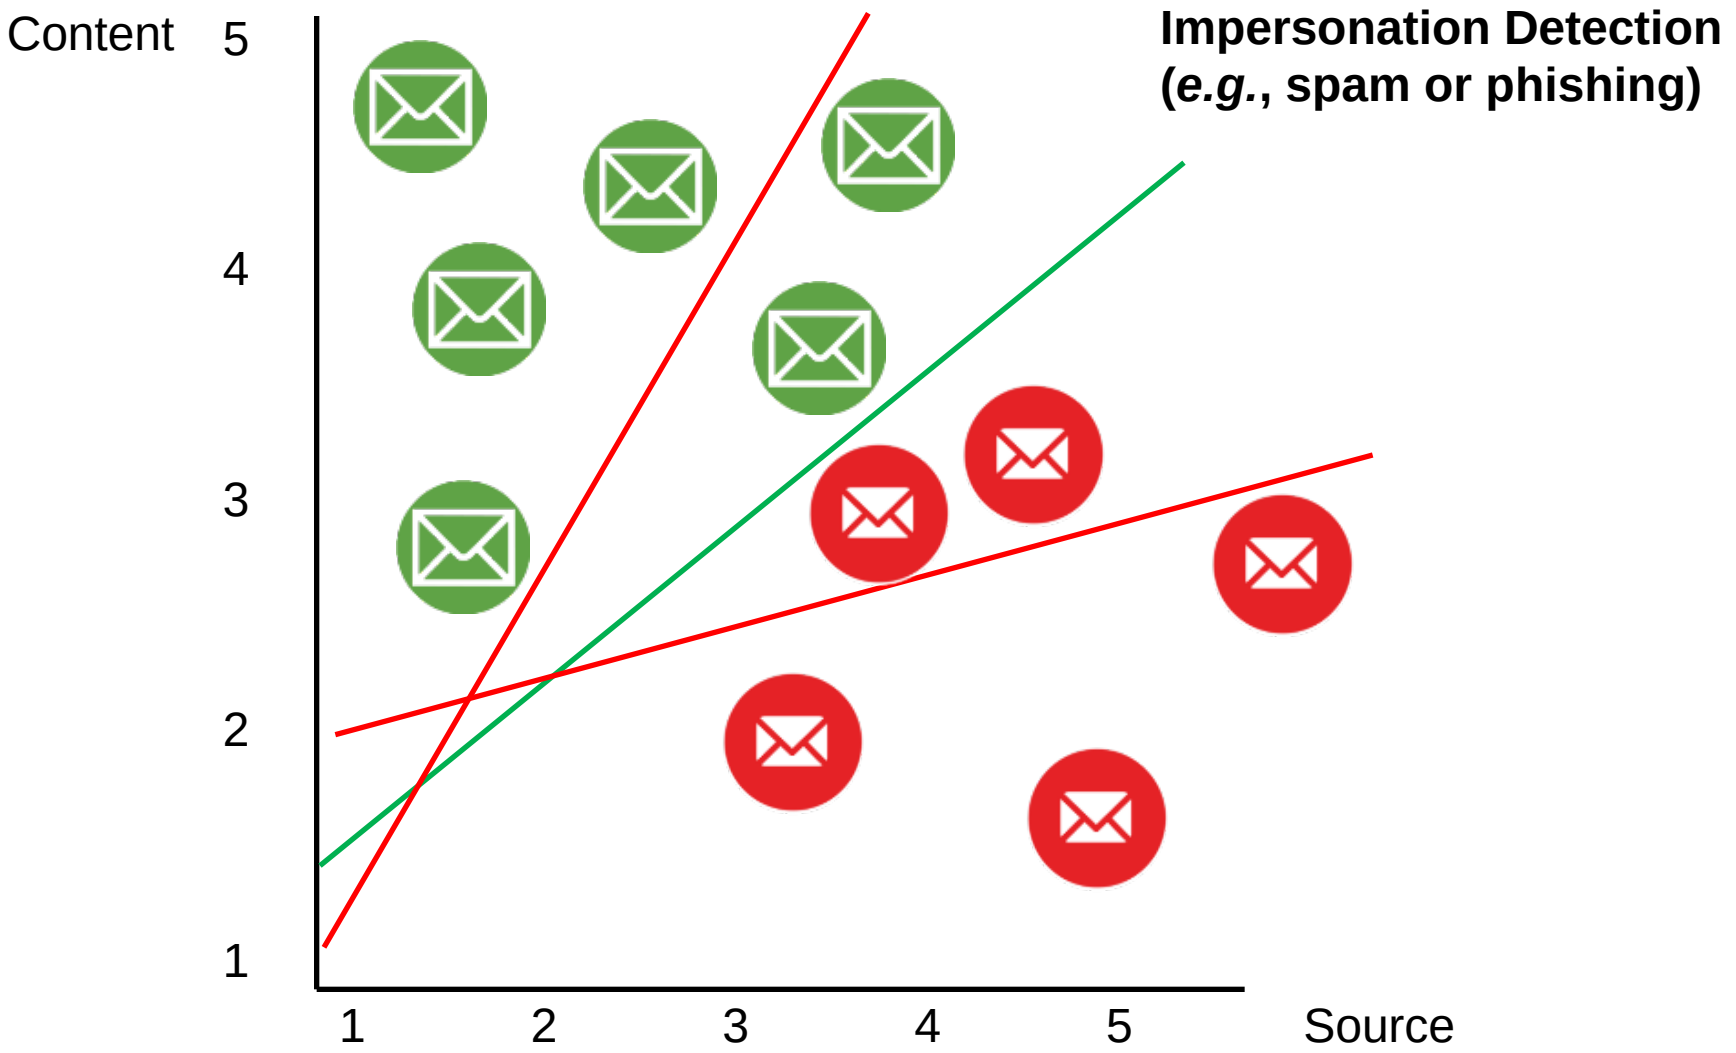

NOW IN MULTIPLE DIMENSIONS

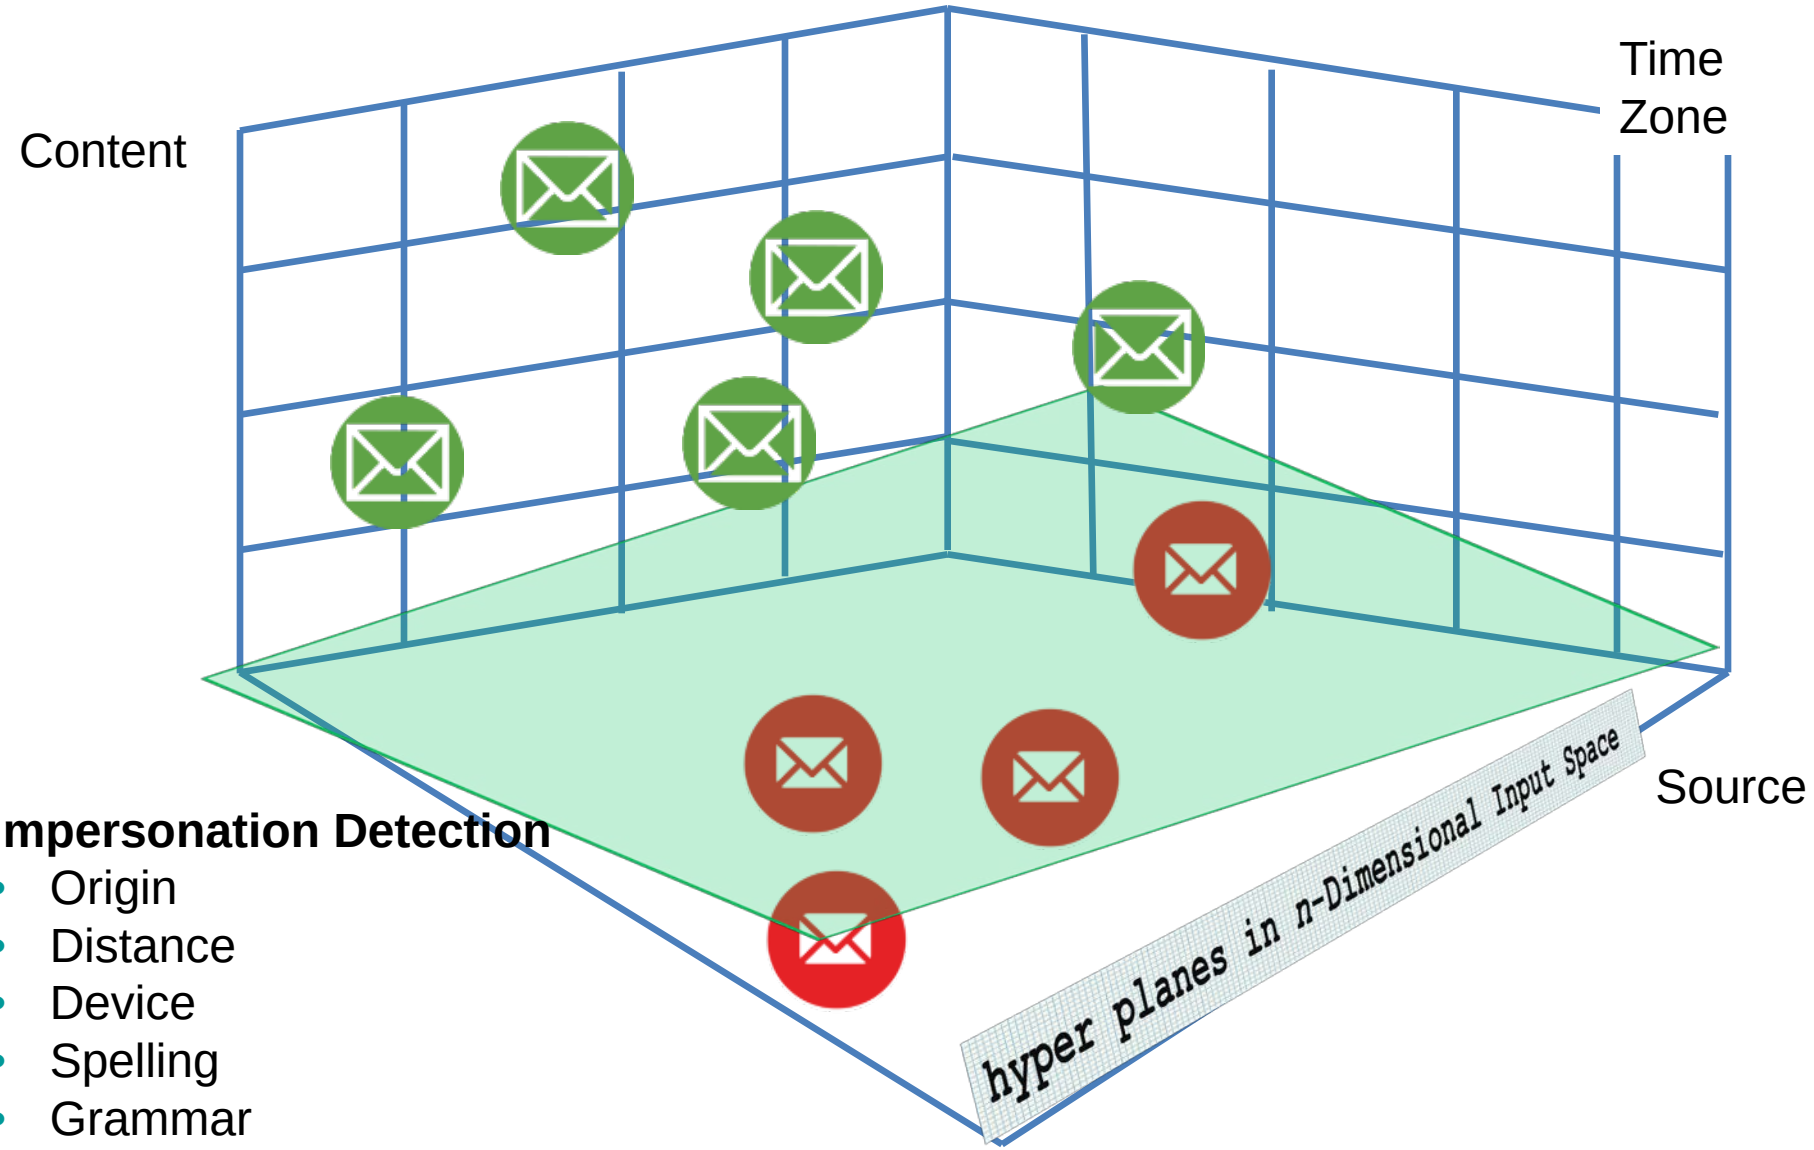

Impersonation Detection

- Origin
- Distance
- Device
- Spelling
- Grammar
- Punctuation
- …

REINFORCEMENT LEARNING

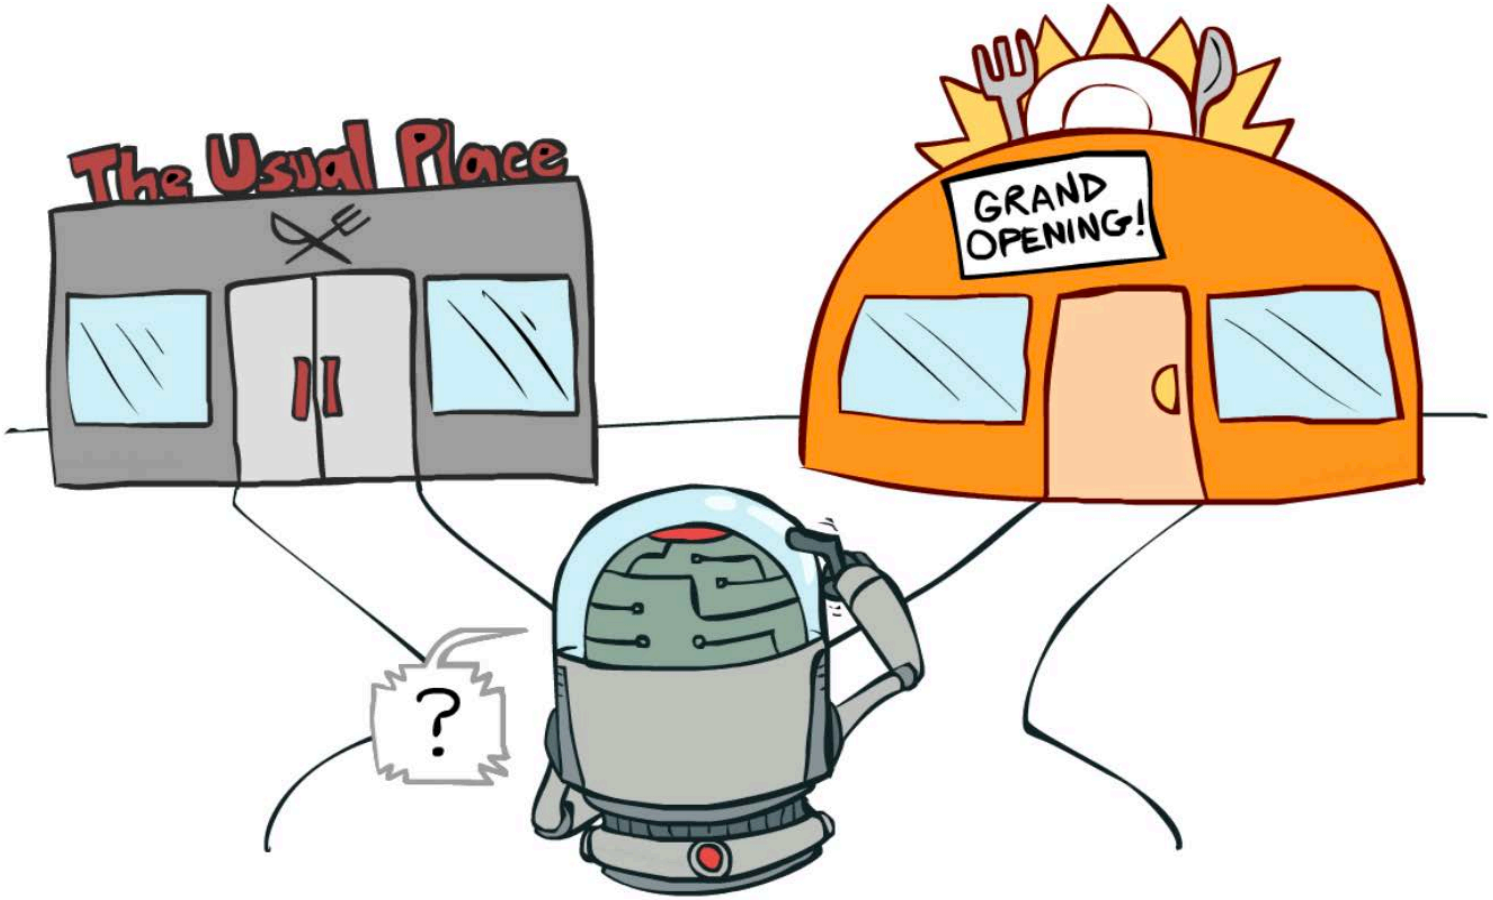

System begins at random but learns goal from reinforcement provided by training

Exploration versus Exploitation

DEEP LEARNING

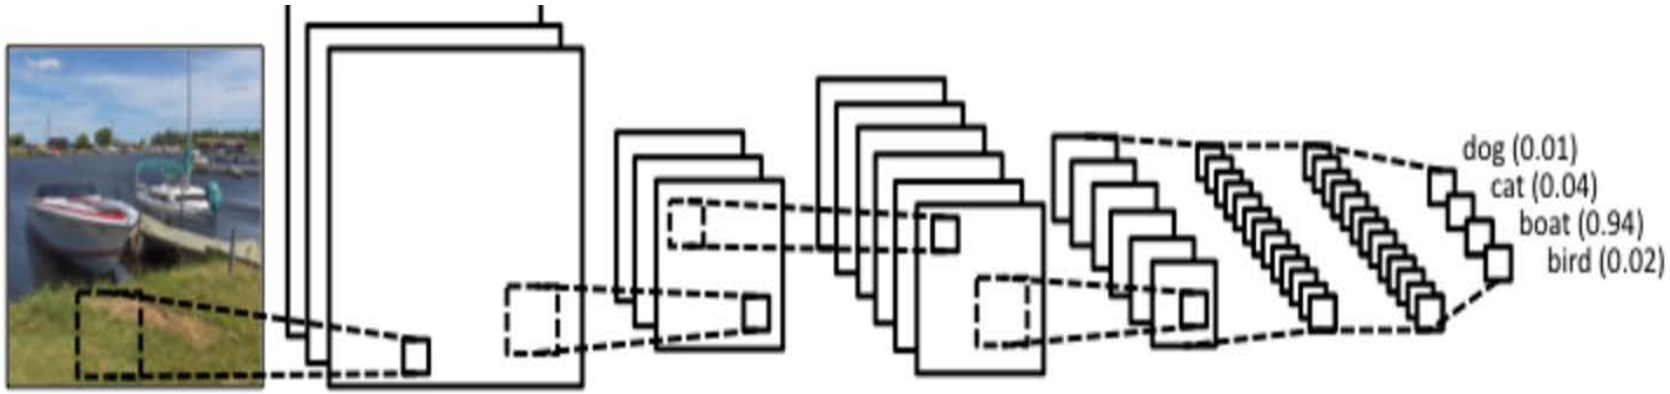

- Uses multiple layers to transform complex input into mathematical representations
- Information from each layer is combined at the next layer
- Requires massive amounts of labeled training data to work

NATURAL LANGUAGE PROCESSING

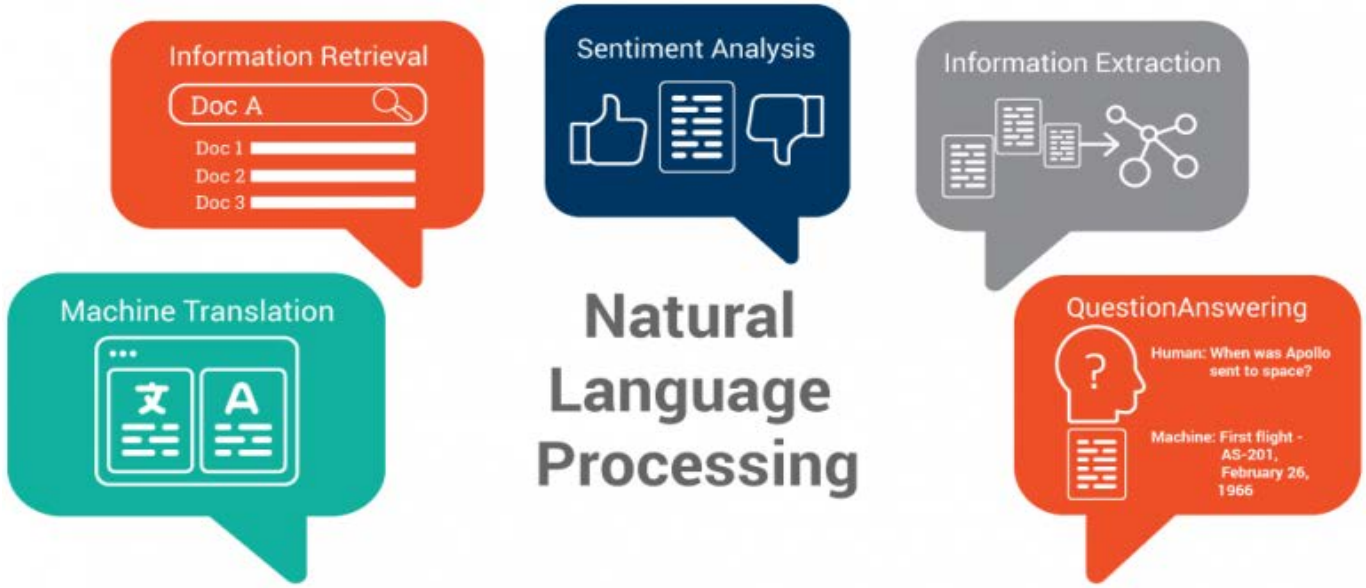

Uses a computer to "understand " human language as it is written or spoken, or creates a computer representation of language (including both syntax and semantics)

- **Tokenization** → Splits longer strings into smaller pieces; determines word boundaries
- **Stemming** → Eliminates prefixes and suffixes from words
- **Bag of words** → Looks for co-occurrences of words in a document
- **Stop words** → Removes words that are noise and don't add meaning
- **Tf-idf** → Determines how important a word is to a document by its frequency
- **Disambiguation of content** → Polisemy (*i.e.*, lead vs. lead)
- **Topic modeling** → Statistical models to discover abstract concepts

What AI Can Do

WHAT KINDS OF PROBLEMS CAN AI SOLVE?

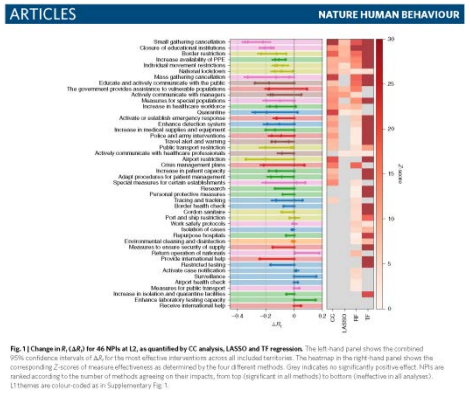

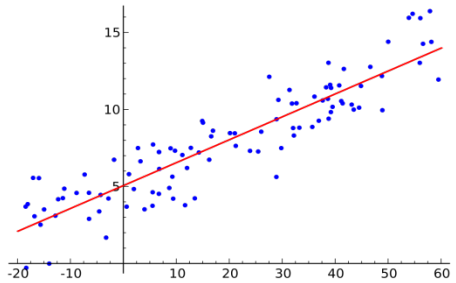

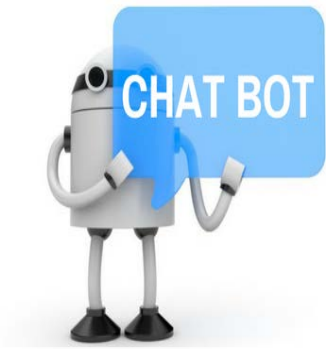

Classification or Prediction

Relevant / Not Relevant

Ranking

Most to least likely to be relevant

Regression

Numerical estimate of probability of relevance

Conversing or Translating

“It’s relevant!”

AI TODAY: THE SKY'S THE LIMIT

- Health Care
  - Detecting cancer in MRIs
  - Diagnosing illness or prescribing treatment
- Education
  - Evaluating teacher performance
- Employment
  - Evaluating job applicants for interviews
- Transportation
  - Autonomous vehicles
  - Drones
  - Logistics
- Finance
  - Credit forecasting
  - Mortgage underwriting
- Law Enforcement
  - Predictive policing and facial recognition
  - Bail evaluation and sentencing
- Government
  - Benefits determination
  - Fraud detection
  - Autonomous weapons

AI IN THE LAW

- Technology-assisted review (“TAR”) and other analytics in eDiscovery
- M&A due diligence / contract analysis and review
- Public disclosure analytics
- Legal research / summarization / drafting of memos and initial pleadings
- Analysis of briefs for missing citations
- Predictions about opposing counsel and courts
- Litigation outcome forecasting for financing purposes
- Jury pool evaluation
- Analysis of claims and automated forms completion
- Billing
- …

Bias and Other Issues Implicated by AI

AI RAISES A SLEW OF LEGAL AND OTHER ISSUES

- Bias
  - Data
  - Algorithm
  - Human
    - ✓Algorithm Aversion
    - ✓Automation Bias
    - ✓Confirmation Bias
- Inadequate Testing
  - Validity → Does in measure what it is supposed to measure?
  - Reliability → Does it measure it consistently?
- "Function Creep"

OTHER LEGAL AND ETHICAL ISSUES

- Claims of trade secret leading to lack of due process
- Lack of transparency and explainability
- Lack of resilience to intentional and unintentional efforts to cause the system to perform other than as designed
- Lack of competence
- Lack of data privacy / security
- Lack of accountability and regulation
- . . .

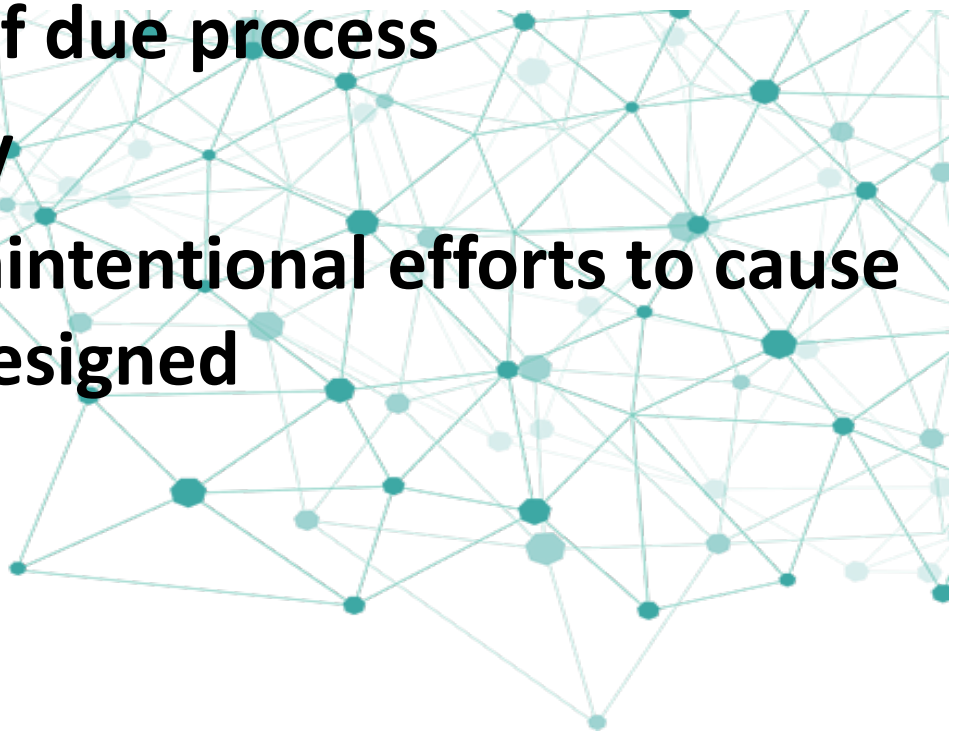

OPPORTUNITY OR DISPLACEMENT?

"Using technology to solve specific real problems using the advantages that computers have with the special abilities that we as human beings share."

Maura R. Grossman, J.D., Ph.D.

School of Computer Science
University of Waterloo
- and -
Maura Grossman Law

maura.grossman@uwaterloo.ca
maura@mauragrossman.com

212.861.8097

Tha e!

UNIVERSITY OF WATERLOO

Maura Grossman Law