### Artificial Intelligence and Machine Learning Overview

### Sargur N. Srihari University at Buffalo, State University of New York USA

### Outline

- 1. What is AI?
- 2. Three paradigms of AI
- 3. Disruption due to Al
- 4. Ethics in Al
- 5. Explainable AI

## History of Al

### Greek Mythology



Pygmalion of Cyprus

Sculpts marble Galatea that came to life (falls in love) GBS' Pygmalion: Higgins teaches Eliza to speak Queen's English



Eliza is first AI program (Weizenbaum, MIT)

### Indian Mythology



Ganesa (Buddhipriya) Parvati forms sandalwood Ganesa, whose head is transplanted from an elephant

## Today AI is ubiquitous

### Automate routine labor



Google						
1						



Understand speech
 – SIRI, Alexa





Autonomous Vehicles





# What tasks require intelligence?

- Reasoning
  - Puzzles, Judgments
- Planning
  - Action sequences
- Learning
  - Improve with data
- Natural language
- Integrating skills
- Abilities to sense, act



### Al Paradox

- Hard problems for people are easy for AI
- Easy problems are hard for AI
  - Narrow Intelligence General Intelligence

People easy tasks:



## Everyday life needs knowledge

- Knowledge is intuitive and subjective
  - Key challenge of AI is how to get this informal knowledge into a computer
- Knowledge-based Approach
  - Hard-code knowledge in a formal language
  - Computer can reason about statements in these languages using inference rules
  - Examples:
    - Expert systems for diagnosis (MYCIN, CADUCEUS)
    - Design (VAX)

Machine Learning

Sriha

### **Knowledge-Based Al**

Rule-based System



Disadvantage: Unwieldy process

Time of human experts

People struggle to formalize rules with enough complexity to describe the world

### The Machine Learning approach

- Difficulties of hard-coded approach suggests:
   Allow computers to learn from experience
- First determine what features to use
- Learn to map the features to outputs

Linear classifier



Quadratic classifier



### The ML Approach





Testing)

### Inference

Find responses to queries

10

### Learning Problem Definition

- Improving some measure of performance *P* when executing some task *T* through some type of training experience *E*
- Example: Learning to detect credit card fraud
- Task T
  - Assign label of fraud or not fraud to credit card transaction



- Performance measure P
  - Accuracy of fraud classifier
     With higher penalty when fraud is labeled as not fraud
- Training experience E
  - Historical credit card transactions labeled as fraud or not

## **ML Problem Types**

- 1. Based on Type of Data
  - 1. Supervised, Unsupervised, Semi-supervised
  - 2. Reinforcement Learning
- 2. Based on Type of Output
  - Regression, Classification
- 3. Based on Type of Model
  - Generative, Discriminative

### Periodic Table of ML algorithms



#### We will look at examples first proceeding along the horizontal axis and then along the vertical

### **Supervised Learning**

- Most widely used methods of ML, e.g.,
  - Spam classification of email
  - Face recognizers over images
  - Medical diagnosis systems
- Inputs *x* are vectors or more complex objects
   documents, DNA sequences or graphs
- Outputs are binary, multiclass(*K*),
  - Multi-label (more than one class), ranking,
  - Structured:
    - y is a graph satisfying constraints, e.g., POS tagging
  - Real-valued or mixture of discrete and real-valued 14



## Ability to Generalize

- ML algorithms need to perform well not just on training data but on new inputs as well
  - 1. Parameter Norm Penalties ( $L^2$  and  $L^1$  regularization)
  - 2. Data Set Augmentation



and	and	and	th	キア	th
and	and	and	th	th	th
and	and	and	th	th	th
and	and	and	th	チア	th
and	and	and	th	42	th
and	and	and	th	ヤア	th
and	and	and	th	チア	th
and	and	and	th	ith	th
and	and	and	th	ith	th

1. Early Stopping

4. Dropout



### Two paradigms in Al



Shaded boxes indicate components that can learn from data

## Designing right set of features

- Simple Machine Learning depends heavily on representation of given data
- For detecting a car in photographs
  - Tire shape difficult in terms of pixel values

- Shadows, glare, occlusion







#### 19

### Srihari Feature Learning for Classification

- Function to map pixels to object identity is complicated
- Series of hidden layers extract increasingly abstract features
- Final decision made by a simple classifier



## **Deep Learning**

- Understand the world as hierarchy of concepts
  - How these concepts are built on top of each other is deep, with many layers
  - Weights learnt by gradient descent



### Summary of AI Models



Shaded boxes indicate components that can learn from data

## Al's Fruit-fly

1846 1996 One Handred Files Years	JNIVERSITY AT BUFFA Wate University of New York EDAR meter of Excellence for focument Analysis and Recognition Decement Analysis	LO RAN FL PM 2 OCT. 1996	star roth at 32
	$\left( \right)$	Saugur N. Srihan 276 Meadourien L	ana
	1422	Williamsnille, NY 1	4221 ulludillulull



- Many handcrafted rules and exceptions
- Better learn from training set
- Handwriting rec cannot be done without ML!



Handwriting is the fruit fly of AI



#### NYU, Toronto, Montreal, Baidu 22

## **Unsupervised Learning**

- Unlabeled data assuming underlying structure
  - 1. Clustering to find partition of data
  - 2. Identify a low-dimensional manifold
    - PCA, Autoencoder
  - 3. Topic modeling
    - Topics are distributions over words
    - Document: a distribution across topics
      - Methods: SVD, Collaborative Filtering
  - 4. Recommendation Systems
  - Data links between users and items
  - Suggest other items to user
  - Solution: SVD, Collaborative Filtering







### **Regression: Learning to Rank**



#### **Target Variable**

- **Point-wise (**0,1,2,3**)**
- Regression returns continuous value
   Allows fine-grained ranking of URLs

### ML as a Software Area

- 1. ML is programming computers for AI:
  - Perform tasks that humans perform well but difficult to specify algorithmically
- 2. ML is principled way to build IT systems
  - Probabilistic responses to queries—IR
  - Adaptive user interfaces, personalized assistants (information systems)
  - Scientific/engineering applications

## Sribari Disruption in Software Development

- Machine Learning is not just another tool

   It is a fundamental shift in how software is written
- Software 1.0 (Classical "stack")
  - It is code we write
    - e.g., LAMP(Linux, Apache, MySQL, Python/Perl)
- Software 2.0 (Code written by Optimizer)
  - It is in a user unfriendly language
    - There are millions of weights
  - No human involved in coding

![](_page_25_Figure_9.jpeg)

2Machine Learning

### <sup>7Srihari</sup> Software 1.0 and 2.0: Fizzbuzz

### • Print i = 1 to 100, except:

- if divisible by 3 print fizz,
- if divisible by 5 print buzz,
- if divisible by both 3 and 5 print fizzbuzz
- Two approaches:
  - Software 1.0: C++

### Software 2.0: Python/Tensorflow

![](_page_26_Picture_9.jpeg)

 $def model(X, w_h, w_o):$ h = tf.nn.relu(tf.matmul(X, w\_h)) return tf.matmul(h, w\_o)

![](_page_26_Figure_11.jpeg)

### Program Space: Software 1.0 vs 2.0

Software 1.0

By writing each line of code, programmer identifies a point in program space with some desirable behavior

Software 2.0

Restrict search space to continuous subset of program space Search is made efficient by using stochastic gradient descent

![](_page_27_Figure_6.jpeg)

\*A. Karpathy, Tesla

### Benefits of Software 2.0

- Computationally homogeneous
  - Sandwich of two operations: matrix multiply, RELU
- Simple to bake into silicon
  - Small instruction set
- Constant run time
  - Every iteration of forward pass has same FLOPS
  - Constant memory use
- Highly portable: sequence of matrix multiplies is easier
- Very agile
  - C++ is hard to speed-up, instead remove half of channels
- Can meld into optimal whole
  - Software often has modules, can jointly optimize
- It is better than you

# Machine Learning imitations of Software 2.0 stack

- After training we have large networks that work very well, but hard to tell how
  - -90% accurate model we understand
  - -99% accurate model we don't
- Can fail unintuitively (Adversarial examples)

![](_page_29_Figure_5.jpeg)

## How ML disrupts CS

- Software Developer
  - Present: Write and maintain layers of tangled code
  - Future: A teacher- curate training data & analyze results
- Mathematics
  - Present: Logic, Discrete mathematics
  - Future: Probability Theory, Calculus, Linear Algebra
- Programming Environments
  - Present: C++, Java
  - Future: Tensorflow/Pytorch/Gluon/...
    - In ten years most software jobs won't involve programming

### Hardware

- Present: CPUs
- Future: GPUs

### AI and Ethics

- Bias Anecdotes
  - Common dataset is 74% male and 84% white
    - Google program identified African Americans as Gorillas
    - Amazon's screening rejected applicants mentioning women's groups or colleges
    - Facial recog. misidentified 35% dark-skinned women as men, while error rate for white men was 0.8%
    - Creditworthiness of those who drove long distances
- Solutions
  - Diverse group of developers
  - Up-front logic (Software 1.0) to handle situations

### **Explainable AI**

![](_page_32_Figure_2.jpeg)

![](_page_32_Picture_3.jpeg)

![](_page_32_Figure_4.jpeg)

Why did you do that? Why not something else? When do you succeed? When do you fail? When can I trust you? How do I correct an error?

### Anecdote: Medical Al Decisions can be worse with Al e.g., Patient discharge to a nursing home

![](_page_32_Figure_7.jpeg)

### Summary

- Machine Learning as an Al approach

   Overcomes limitations of knowledge-based systems
- Types of AI tasks
  - Data: Supervised, Unsupervised, Semi-supervised, Reinforcement
  - Output: Classification, Regression, Ranking
  - Model: Generative, Discriminative
- Disruption of computer science
  - Principled approach to develop IT systems
  - Drivers are mobile systems (big data), personalization