Artificial Intelligence and Deep Learning

Sargur N. Srihari srihari@buffalo.edu

Topics

- Artificial Intelligence History
- Al paradigm shifts
 - Deep Learning
- Disruption caused by AI
 - Software development
 - Societal
- India and 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





Handwriting as "Fruit Fly" of Al







NYU, Toronto, Montreal, Baidu

Technology Ages (50-year spans)

- 1771 Industrial Revolution
 Arkwright's mill in Cramford
- 1826 Age of Steam and Railways
 - "Rocket steam" engine for Manchester railway
- 1875 Age of Steel, Electricity & Heavy Eng.
 Carnegie Bessemer steel plant in Pittsburgh
- 1908 Age of Oil, Automobile, Mass Production
 First model T comes out in Detroit
- 1971 Age of Information & Telecommunications
 Intel microprocessor announced in Santa Clara
- 2017 Age of Artificial Intelligence
 - Machines and people connected
 - Data is new oil, gold







Disruptions due to Al



International (Rich and mid-income): 14% of jobs in 32 countries have 70% risk For 32% of jobs, that probability is 60% Thus 210m jobs at risk in those countries alone. Rich-countries less at risk than mid-income states

AI will outperform humans in*:

Translating languages 2024 Driving a truck 2027 Working in retail 2031 Working as a surgeon 2053 Outperform humans in all: 50% chance in 45 years Automating all human jobs 120 years *NIPS/ICML Survey

National AI Initiatives



"Artificial Intelligence is poised to disrupt the world" NITI Aayog

Al development (much to be done)

- Paradox
 - Hard problems for people are easy
 - Easy problems for people are hard
- Early AI (Artificial Narrow Intelligence)
 Chess, GO
- Today's AI (Artificial General Intelligence)
 - Solve easy tasks for people:
 - Recognize faces in images

What tasks require intelligence?

- Reasoning
 - Solve puzzles, make judgments under uncertainty
- Planning
 - Strategies or action sequences, e.g., for robots
- Learning
 - Improve performance on a task with data
- Communicating in natural language
- Integrating all these skills towards goals
- Abilities to sense, act

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

Knowledge-Based Al



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
- Determine what feature representations to use
- Map the features to outputs

Decide whether email is spam





Two paradigms in AI



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







Representation Learning

 Solution: use ML to not only learn mapping from representation to output but representation itself

Better results than hand-coded representations

- Allows AI systems to rapidly adapt to new tasks
 - Designing features can take great human effort
 - Can take decades for a community of researchers
- Does not need programmer to have deep knowledge of the problem domain

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



Unsupervised Representation Learning

- Autoencoder
 - Encoder:
 - Converts input into a representation
 - Decoder:
 - Converts representation
 back to input



New designs from representation



Language Processed as Numbers

- Training Data
- Word-to-vec
 - One-hot vector mapped to vector of 300
- Word embedding
 - Similar words are close together



		him them		Iraq Israel Pr Russia Britain France _G ermany Japan London China	ime Clintor
Noven Sæp tember July August	7.3. 7 .8.30				Reuter
	Saturday		South	New United Nationa Int	l ternational
	Friday Wednesday Thursday Monday			British Australi an nerio	an
late	Tuesday earlier		U.S.	Russian Europea Japanese Germarench inte	n ernational
		past	ce	ntral domestic local foreign	1
ne	lasf ^{ext}	same	top	national private senior	public

Three Paradigms of Al



Shaded boxes indicate components that can learn from data

Disruption in Software Development

- Deep 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



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++

2 2

Software 2.0: Python/Tensorflow



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





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



*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

Limitations 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)



How Deep Learning 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

The transition

- Conventional software can be replaced with a deep learning solution with improvement, e.g.*,
 - Upgrading search ranking
 - Data center energy usage
 - Language translation
 - Solving GO
- Repositories
 - Github is home for Software 1.0 code
 - Software 2.0 repositories are datasets
 - Commits are additions and edits of the labels

*Projects at Alphabet:

Al Development: US, China lead

Deep Learning Frameworks (US) IJCAI 2018 Papers (China)



India's future role in Al

- Strong fundamentals in mathematics
- Demographic strength
- 50% participation by women
- Large data sets/curation
- NITI Aayog #Alforall



Conclusion

- Al is disrupting CS
- Al will disrupt the world
 - Deep learning is better than you
- India's role:
 - Long-ramp up as knowledge diffuses through developer community
 - India is positioned well for
 - Development of AI and
 - Al applications