

# Artificial Intelligence: Current Methods and Prospects

Sargur N. Srihari  
srihari@cedar.buffalo.edu

# Topics

- Artificial Intelligence
  - History, Challenges
- Representation Learning
  - Autoencoder
- Factors of Variation
  - VAE, GAN
- Natural Language Processing
  - Word-to-vec
- Comparison of approaches to AI
  - Future Prospects

# Artificial Intelligence Origins

- Thinking machines traceable to Greece?



Mythical inventors Pygmalion, Daedalus, Hephaestus can be interpreted as inventors

## Pygmalion of Cyprus

Sculpts ivory.

Makes form of Galatea that came to life

(and he fell in love with it)



## Pygmalion of GBS:

Higgins teaches Eliza to speak Queen's English

(and falls in love as well)

# Early History of AI

- Advent of *programmable* computers led people to wonder whether AI was possible
  - Lovelace 1842
    - 100 years before first computer
  - Turing 1940s
    - Turing's test for AI



# Closed world: Machine > Computer

- GO has  $10^{10^{48}}$  games, Chess has  $10^{45}$  moves,
  - Computer cannot solve problem but only has to win
  - Gary Kasparov on loss to Deep Blue:
    - Human uses 1% calculation, 99% *understanding*
      - based on patterns, drawing information from experience
    - Machine opposite: 99% calculation 1% understanding
      - though this understanding is growing
  - Google Deep Mind's Alphago:
- Closed-world (Chess- 62 squares, 32 pieces):
  1. Needs only to reason about next move
  2. Near-infinite training data available




# Challenge of AI

- Early successes of AI:
  - Solved problems intellectually difficult for humans
    - problems described by small set of rules
      - Sterile formal environment
      - Little knowledge about the world
- True challenge of AI
  - Solve tasks easy for people but hard to describe formally
    - solved intuitively, that feel automatic
      - e,g recognize spoken words, or faces in images
- Today's AI:
  - It is about solving these more intuitive problems

# Solution to Intuitive Problems

1. Allow computers to learn from experience
  - By gathering knowledge from experience
    - Avoids need for human operators to specify knowledge that the computer needs
2. Understand the world as hierarchy of concepts
  - Thereby learn complicated concepts by building them out of simpler ones
    - A graph of how these concepts are built on top of each other is deep, with many layers



# Everyday life needs knowledge



- A person's everyday life requires immense amount of knowledge of the world
  - Much of this knowledge is intuitive and subjective
    - Difficult to articulate in a formal way
- Computers need to capture same knowledge to behave intelligently
- Key challenge of AI is how to get this informal knowledge into a computer

# The Knowledge-based Approach

- Hard-code knowledge in a formal language
  - Computer can reason about statements in these languages using inference rules
  - Ex: Cyc is an inference engine and database of statements in CycL
- Unwieldy process
  - Staff of human supervisors
  - People struggle formalize rules with enough complexity to describe the world

# The ML approach


- Difficulties of hard-coded approach suggests:
  - AI systems need ability to acquire knowledge
    - By extracting patterns from raw data
- This approach is known as machine learning
  - It allowed tackling problems requiring knowledge of real world and make decisions that seem subjective
  - Simple ML algorithms are:
    - Logistic Regression
      - Caesarian delivery using uterine scar feature
    - Naïve Bayes
      - Decide whether email is spam



# Simple ML depends on representation

- Simple ML algorithms depend heavily on *representation* of given data
- Ex: Logistic regression to recommend delivery does not examine patient directly
  - Instead doctor provides information such presence of uterine scar (called a feature)
  - If Logistic regression was provided with MRI scan, it could not make useful predictions

Individual pixels have negligible correlation with complications in delivery




# Dependence on Representation

- This dependence appears in computer science and daily life
  - In CS: search is exponentially faster if data is data is structured and indexed intelligently 
  - People: arithmetic on Arabic numerals 1,2,...9,10, easier than using Roman numerals I, II,...IX, X,
  - Choice of representation has enormous effect on ML algorithm performance

- Ex: straight line separation

- Impossible in Cartesian
- Simple straight line in polar



# Designing right set of features

- Useful feature for speaker ID: size of vocal chord
  - Provides strong clue for: man, woman or child
- Difficult to know what set of features are good for detecting a car in photographs
  - Presence of wheel as a feature
    - Has a simple geometric shape but difficult to describe in terms of pixel values
      - Shadows, glare from metal parts, fender or object in front obscures




# 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

# Example of Representation Learning

- Autoencoder:
  - Quintessential example of representation learning
  - Encoder:
    - Converts input into a representation with nice properties
  - Decoder:
    - Converts the representation back to input




New designs from representation




# Illustration of Deep Learning

- Function mapping from pixels to object identity is complicated
- Series of hidden layers extract increasingly abstract features



# Factors of Variation (FoVs)

- When designing features or algorithms for learning features,
  - we seek invariance to FoVs in the observed data
- Factors refer to separate sources of influence
  - Often not directly observed
- Examples of FoVs
  - Face:
    - pose, expression, morphology
  - Speech:
    - age, sex, accent, words spoken
  - Car image:
    - position, color, angle of sun



# Disentangling/Preserving FoVs


- Many FoVs influence every data we observe and they interact
  - FoVs for a car:
    - Night: pixels of red car are close to black at night
    - Viewing angle: decides shape of car silhouette
- Discriminative models: *disentangle* FoVs and discard those we do not care about
- Generative models: *preserve* informative FoVs

# Disentangling FoVs

- During training, only supervision is class labels
- Specified FoVs
  - Images captured from different viewpoints
  - Strong supervision: pairs of images with two different objects at same viewing angle
- Unspecified FoVs
  - Labels unavailable
- A disentanglement method
  - Combine *variational autoencoder* with *adversarial training*

# Variational Autoencoder (VAE)

- Combines two types of models: *discriminative* and *generative* models into a single framework



The models are trained jointly using the variational EM framework

# The VAE model

- Method for modeling a data distribution using a collection of independent latent variables
  - **Generative model:**  $p(x, z) = p(x|z)p(z)$ 
    - $x$  is a r.v. representing observed data
    - $z$  is a collection of latent variables
  - $p(x|z)$  is parameterized by a deep neural network (decoder)
  - Components of  $z$  are independent Bernoulli or Gaussian
  - Learned approx inference trained using gradient descent
    - $q(z|x) = N(\mu, \sigma^2 I)$  whose parameters are given by another deep network (encoder)
  - Thus we have  $z \sim \text{Enc}(x) = q(z|x)$  and  $y \sim \text{Dec}(z) = p(x|z)$


# Generative Adversarial Net (GAN)

- Simultaneously train two models
  - A generative model  $\mathcal{G}$  that captures the data distribution
  - A discriminative model  $\mathcal{D}$  that estimates a sample came from the training data rather than  $\mathcal{G}$
  - Training procedure for  $\mathcal{G}$  is to maximize probability that  $\mathcal{G}$  makes a mistake
- Framework corresponds to a minimax two-player game
  - For arbitrary  $\mathcal{G}$  and  $\mathcal{D}$ , a unique solution exists
  - If  $\mathcal{G}$  and  $\mathcal{D}$  are MLPs can be trained using backprop

# GAN Image Synthesis

- Architecture/hyperparams are carefully selected
  - Latent variables capture important FoVs

Images of bedrooms generated by Deep Convolutional GAN






Images of churches generated by Laplacian Pyramid of GANs



# Natural Language Processing

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



# Word-to-vec:

Represent noun by co-occurrences with 25 verbs\*

Semantic feature values:

**“celery”**

0.8368, eat  
0.3461, taste  
0.3153, fill  
0.2430, see  
0.1145, clean  
0.0600, open  
0.0586, smell  
0.0286, touch  
...  
...  
0.0000, drive  
0.0000, wear  
0.0000, lift  
0.0000, break  
0.0000, ride

Semantic feature values:


**“airplane”**

0.8673, ride  
0.2891, see  
0.2851, say  
0.1689, near  
0.1228, open  
0.0883, hear  
0.0771, run  
0.0749, lift  
...  
...  
0.0049, smell  
0.0010, wear  
0.0000, taste  
0.0000, rub  
0.0000, manipulate

\* in a trillion word text collection


# Neural model of language

[Mitchell et al., *Science*, 2008]



# Summary of different AI Models

- Relationship of different parts of system in different AI disciplines
- Shaded boxes involve ML



# History of Technology

- 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, data and people connected in new era



# Future prospects of AI

- Survey of AI researchers

- AI will outperform humans in:

- Translating languages 2024
    - Writing high-school essays 2026
    - Driving a truck 2027
    - Working in retail 2031
    - Writing a best-selling book 2049
    - Working as a surgeon 2053
    - Outperform humans in all tasks: 50% chance in 45 years
    - Automating all human jobs 120 years

- Survey population: 2015 NIPS/ICML authors

- Questions on AI capabilities (e.g. folding laundry, language translation), superiority at specific occupations (e.g. truck driver, surgeon), superiority over humans at all tasks.