

# Introduction to Deep Learning

Sargur N. Srihari  
srihari@cedar.buffalo.edu

This is part of lecture slides on [Deep Learning](http://www.cedar.buffalo.edu/~srihari/CSE676):  
<http://www.cedar.buffalo.edu/~srihari/CSE676>


# Topics

- Artificial Intelligence Paradigm shifts
- Definition of Deep Learning
- Representation Learning
- Software 1.0 vs Software 2.0

# AI is ubiquitous

- Automate routine labor

- Search




- Understand speech

- SIRI, Alexa



- Autonomous Vehicles



# AI is based on Machine Learning

- ML powers many aspects of modern society
  - Web searches
  - Content filtering on social networks
  - Recommendations on e-commerce websites
  - Consumer products:
    - cameras, smartphones
- ML used to:
  - Identify objects in images
  - Transcribe speech to text
  - Match news items, posts or products with user interests
  - Select relevant results of search
- Increasingly these use deep learning techniques

# Challenge of AI

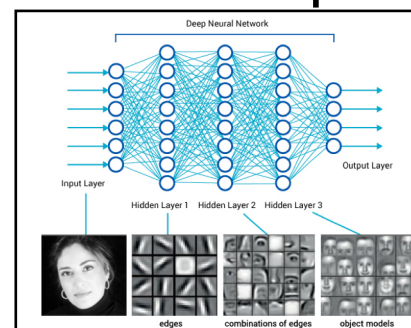
- Early successes of AI:
  - Solved problems intellectually difficult for humans
    - Problems described by small set of rules
      - Sterile formal environment
      - Little knowledge about the world
- True challenge of AI
  - Solve tasks easy for people but hard to describe formally
    - solved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

# 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

# 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





# 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




# Limitations of Conventional ML

- Limited in ability to process natural data in raw form
- PR and ML systems require
  - careful engineering and domain expertise to transform raw data, e.g., pixel values, into a feature vector for a classifier

# 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

- 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



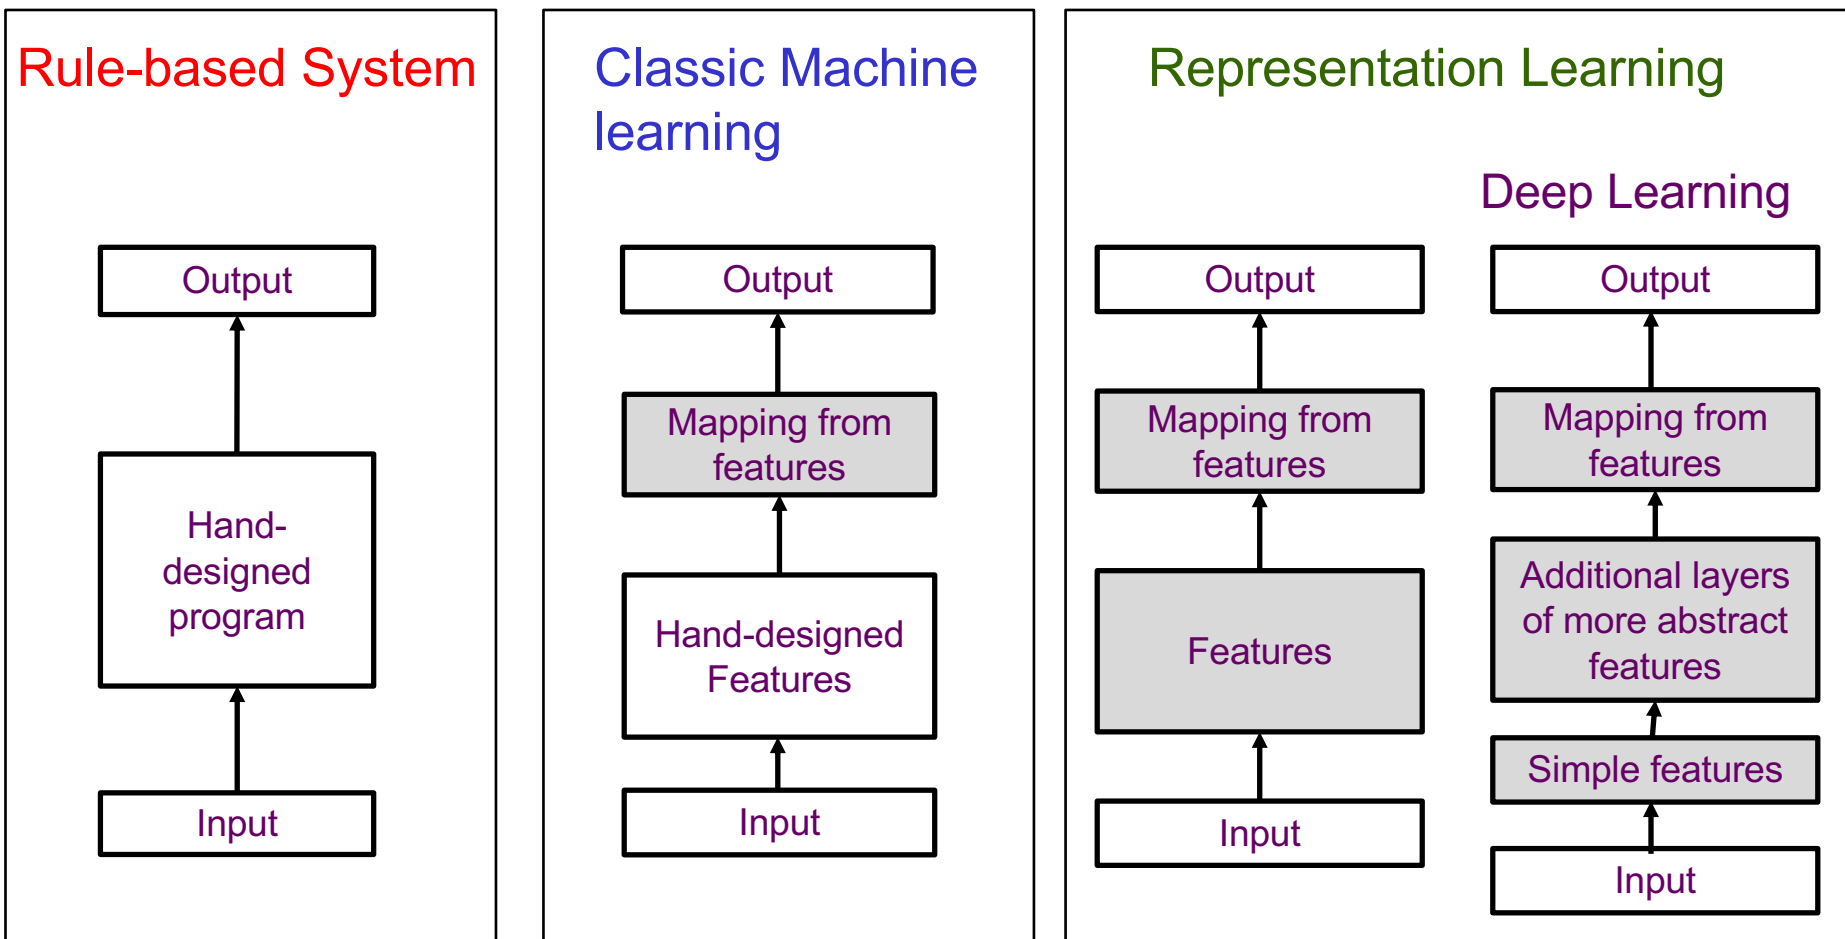
# DL is Representation Learning

- 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
- Deep Learning methods are Representation Learning Methods
  - Methods allow a machine to be fed with raw data to automatically discover representations needed for detection or classification

# Advantage of multiple levels

- Compose simple but non-linear modules that transform representation at one level (starting with raw input) into a representation at a higher slightly more abstract level
- Complex functions can be learned
  - Higher layers of representation amplify aspects of input important for discrimination and suppress irrelevant variations

# Three paradigms of AI



■ Shaded boxes indicate components that can learn from data



# Physics and AI Paradigm Shifts

- Physics paradigm shift
  - Newtonian Physics
    - Cannot explain black-body radiation
  - Quantum Mechanics
- AI paradigm shift
  - Knowledge-based systems
    - Cannot perform simple recognition tasks
  - Simple machine learning methods
    - Cannot perform complex recognition tasks
  - Deep Learning methods
- Coexisting paradigms



# Definition of Deep Learning

- Deep learning is inspired by neural networks of the brain to build learning machines which discover rich and useful internal representations, computed as a composition of learned features and functions
- This definition is a goal and does not say much about HOW we achieve that
  - E.g., by adding priors to learn better high-level representations
  - The term deep learning is indeed aspirational, like 'AI' or 'machine learning'

# Characteristics of Deep Learning

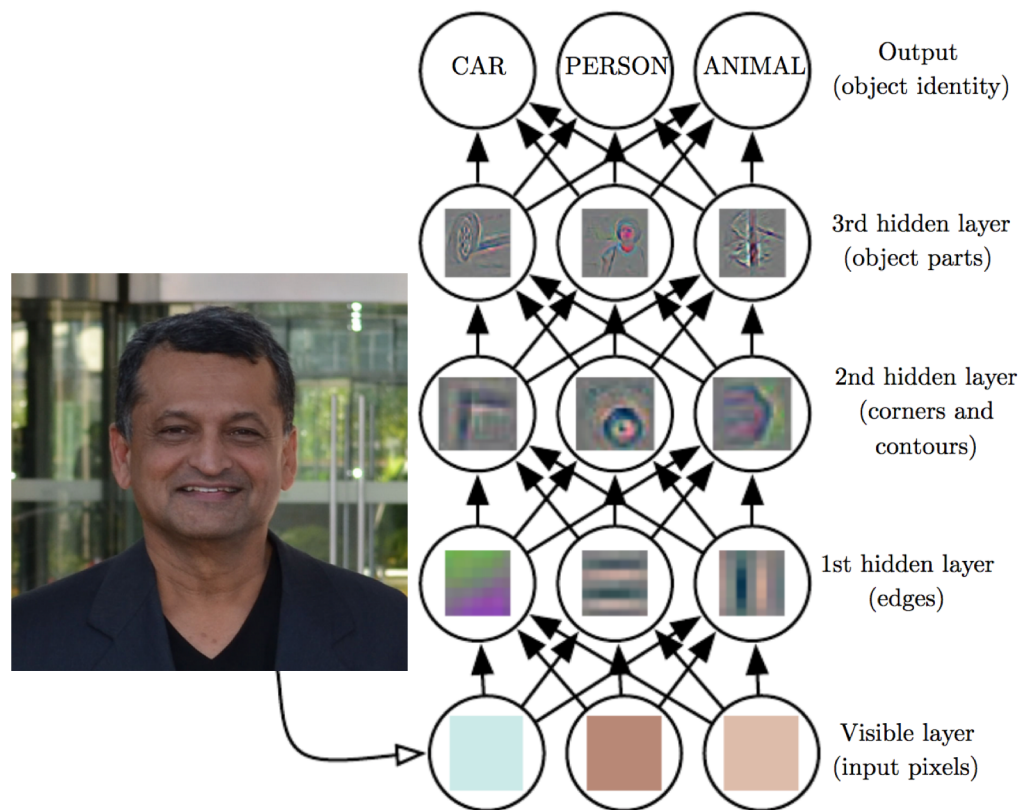
1. A type of Machine Learning that improves with experience and data
2. Only viable approach to building AI systems in real-world environments
3. Obtains its power by a nested hierarchy of concepts
  - each concept defined by relationship to simpler concepts
    - More abstract representations computed in terms of less abstract ones

# Image Example

- Input is an array of pixel values
  - **First Layer:** presence or absence of *edges* at particular locations and orientations of image
  - **Second layer:** detect *motifs* by spotting arrangements of edges, regardless of small variations in edge positions
  - **Third layer:** assemble motifs into larger combinations that corresponds to *parts* of familiar objects
  - **Subsequent layers:** detect *objects* as combinations of these parts
- Key aspect of deep learning:
  - **Layers not designed by human engineers**
    - Learned from data using a general purpose learning procedure


# Example of Image Deep Learning

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



# Unsupervised Representation Learning

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

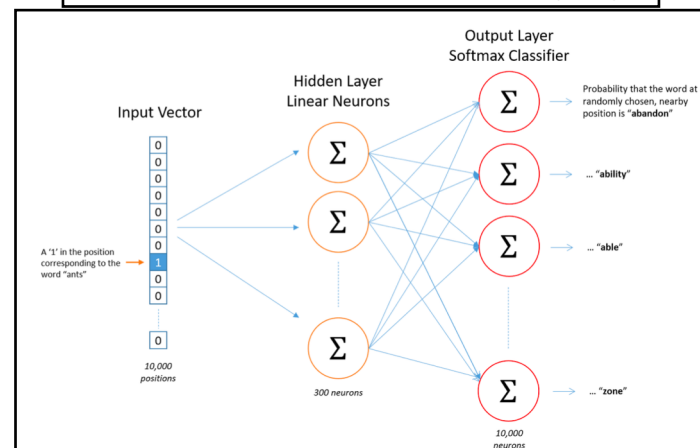
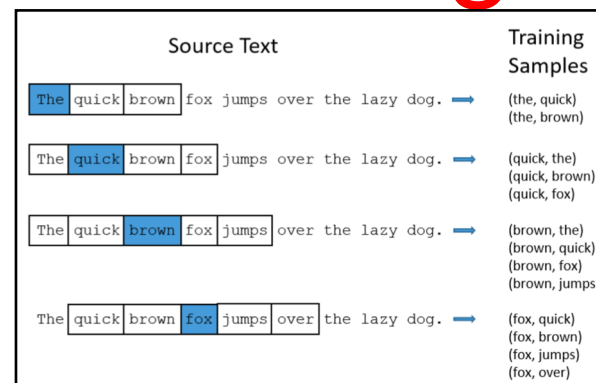


New designs from representation



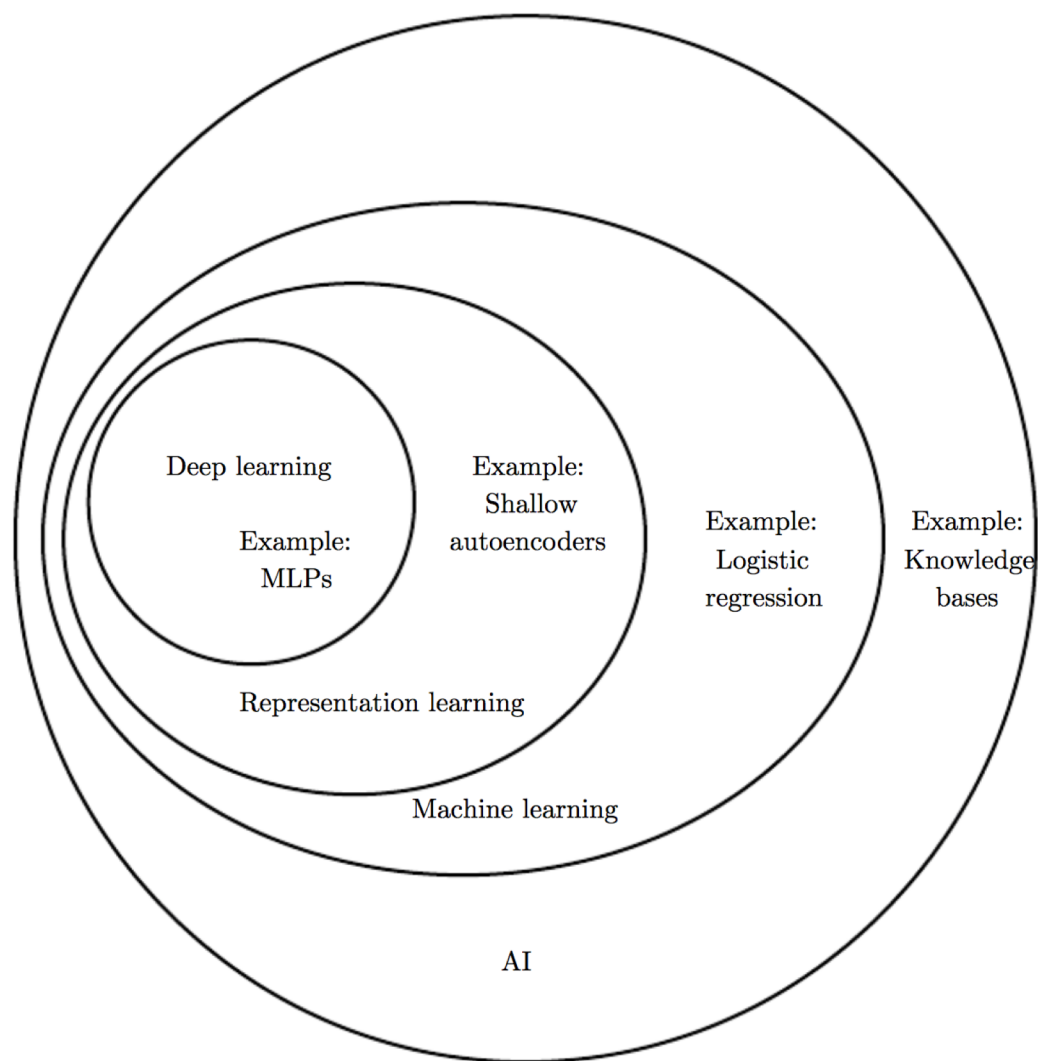
# Natural Language Processing

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



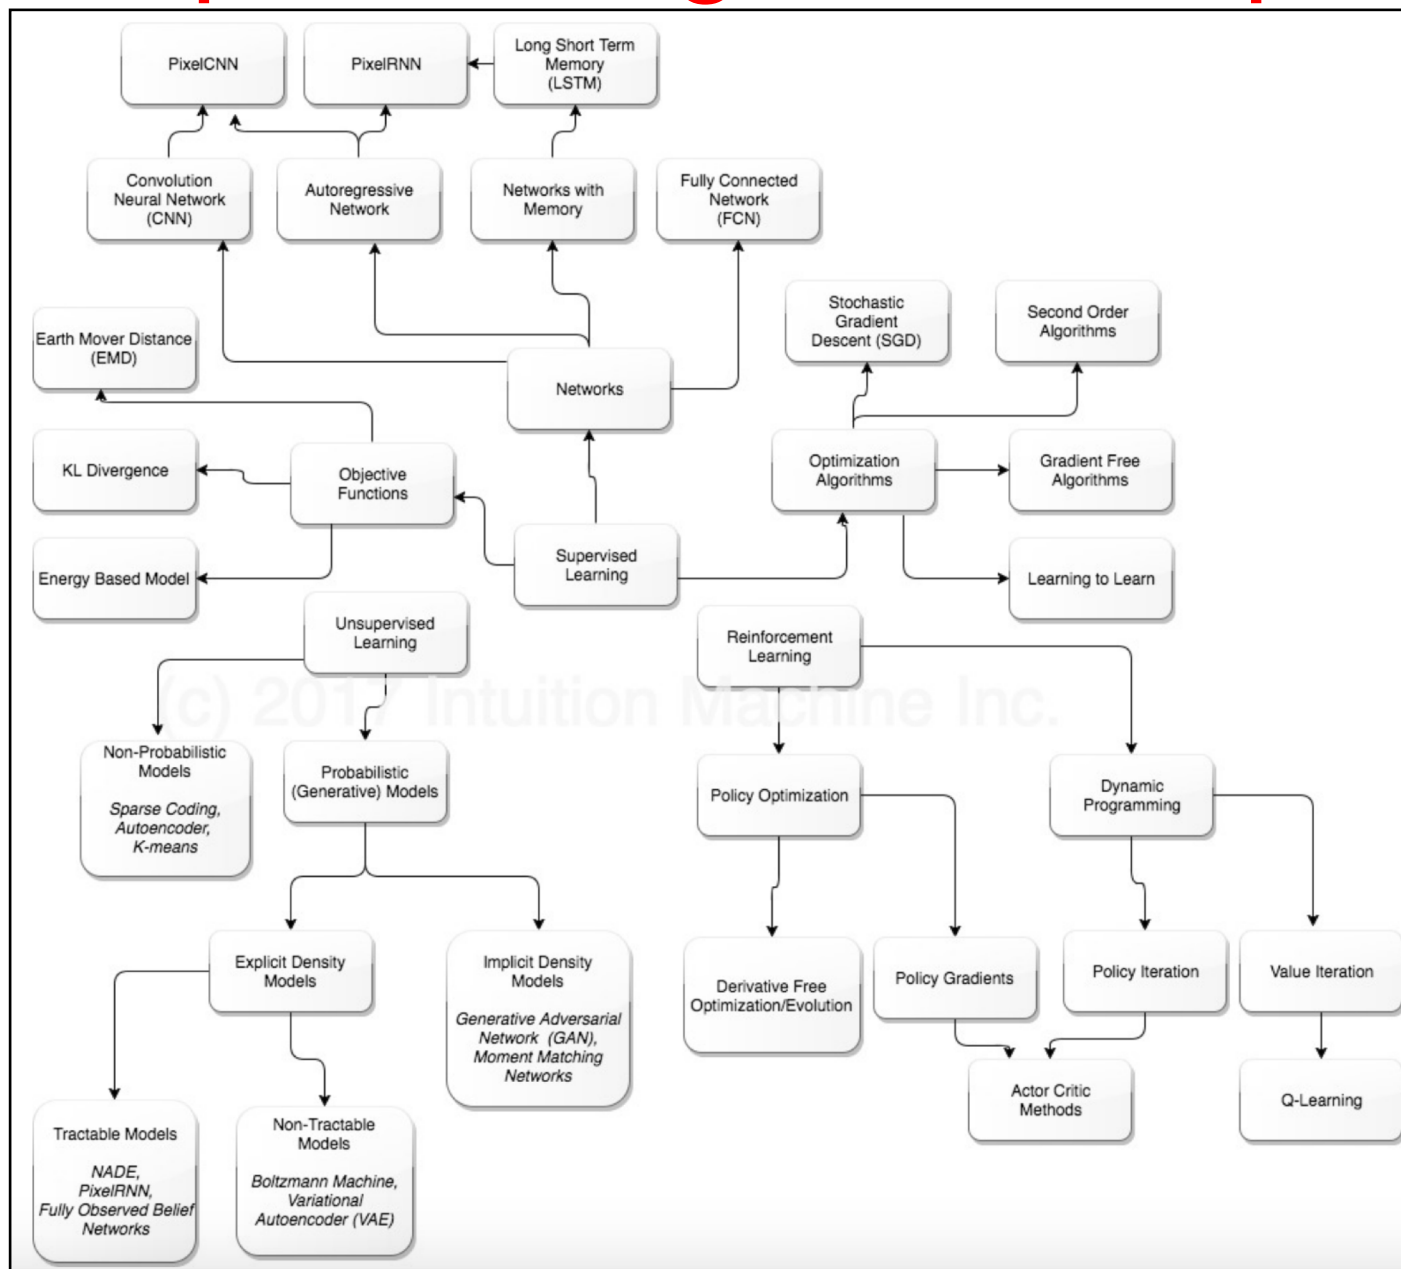
# Approaches to AI: Venn Diagram

- AI methods
- Deep Learning is a type of representation learning
- An example AI technology is included in each
- Also see





# Deep Learning Road Map



# Study of Deep Learning

