

# Neuro-Symbolic Artificial Intelligence


Sargur N. Srihari  
srihari@buffalo.edu

# Human and Biological Intelligence

<https://www.youtube.com/watch?v=PlevD5Mjeak>




# Neurons

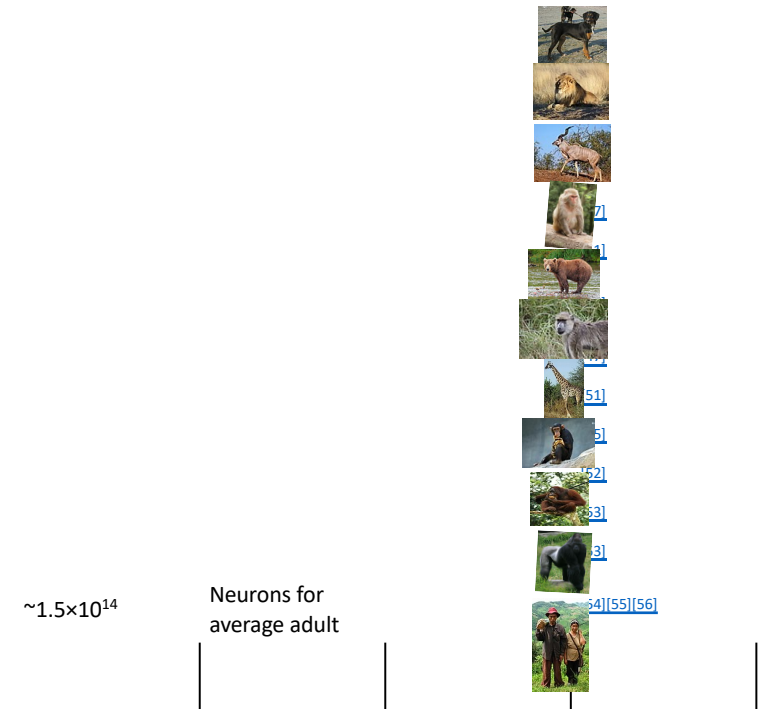


# Number of Neurons


## Whole nervous system

<a href="#">Tardigrade</a>	about 200	Brain only	
<a href="#">Ciona intestinalis</a> larva (see <a href="#">squirt</a> )	231	8,617 (central nervous system only)	
<a href="#">Caenorhabditis elegans</a> (roundworm)	302	~7,500	<a href="#">[13]</a>
<a href="#">Jellyfish</a>	5,600	<i>Hydra vulgaris</i> ( <i>H. attenuate</i> )	<a href="#">[14]</a>
<a href="#">Megaphragma mymaripenne</a>	7,400		<a href="#">[15]</a>
<a href="#">Box jellyfish</a>	8,700–17,500	adult <a href="#">Tripedalia cystophora</a> (8 mm diameter) – does not include 1000 neurons in each of the four <a href="#">rhopalia</a>	<a href="#">[16]</a>
<a href="#">Medicinal leech</a>	10,000		<a href="#">[17]</a>
<a href="#">Pond snail</a>	11,000		<a href="#">[18]</a>
<a href="#">Sea slug</a>	18,000		<a href="#">[19]</a>
<a href="#">Amphioxus</a>	20,000		<a href="#">[20]</a> <a href="#">[21]</a>
<a href="#">Larval zebrafish</a>	100,000		<a href="#">[22]</a>
<a href="#">Lobster</a>	100,000	central nervous system only	<a href="#">[23]</a>
<a href="#">Fruit fly</a>			

<a href="#">Dog (Komondor)</a>	$3.99 \times 10^9$
<a href="#">Dog (Transylvanian Hound)</a>	$4.39 \times 10^9$
<a href="#">Lion</a>	$4.667 \times 10^9$
<a href="#">Greater kudu</a>	$4.91 \times 10^9$
<a href="#">Rhesus macaque</a>	$6.376 \times 10^9$
<a href="#">Brown bear</a>	$9.586 \times 10^9$
<a href="#">Giraffe</a>	$1.075 \times 10^{10}$
<a href="#">Yellow baboon</a>	$1.095 \times 10^{10}$
<a href="#">Chimpanzee</a>	$2.8 \times 10^{10}$
<a href="#">Orangutan</a>	$3.26 \times 10^{10}$
<a href="#">Gorilla</a>	$3.34 \times 10^{10}$
<a href="#">Human</a>	$8.6 \times 10^{10}$
<a href="#">African elephant</a>	$2.57 \times 10^{11}$



# No of neurons



# List of animal species by forebrain (cerebrum or pallium) neuron number

Asian elephant

6,775,000,000

Human

16,340,000,000  
21,000,000,000\*

Isotropic  
fractionator  
Optical  
fractionator


Pallium (cortex)

*Homo sapiens*


# Stories and Abstraction (YN Harari)

- Success of Homo-Sapiens over other primates
  - Ability to abstract high-level *stories*
    - Allow collaboration with others at scale
    - Thus have a broader impact/influence
- Abstracting concepts is necessary for stories
  - Conscious decision beyond involuntary reaction
    - To abstract, generalize, and reason about scenarios and concepts that do not physically exist
    - To focus attention to certain, limited set of features and process them in depth, while disregarding others


# Foundations of Neurosymbolic AI



775 pages



1,232 pages



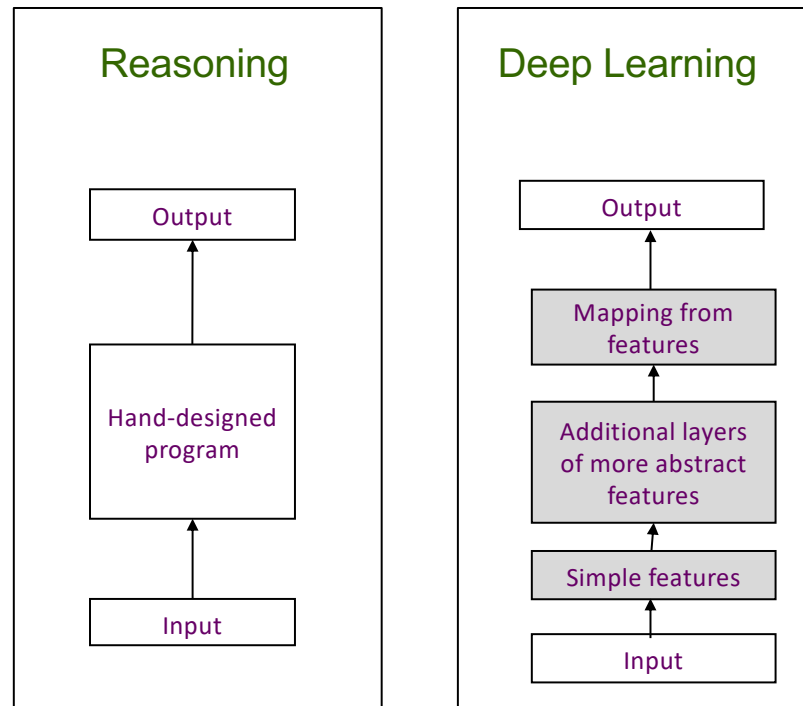
499 pages



# Topics

- Paradigms of Artificial Intelligence (AI)
- Causal Reasoning
- Deep Learning

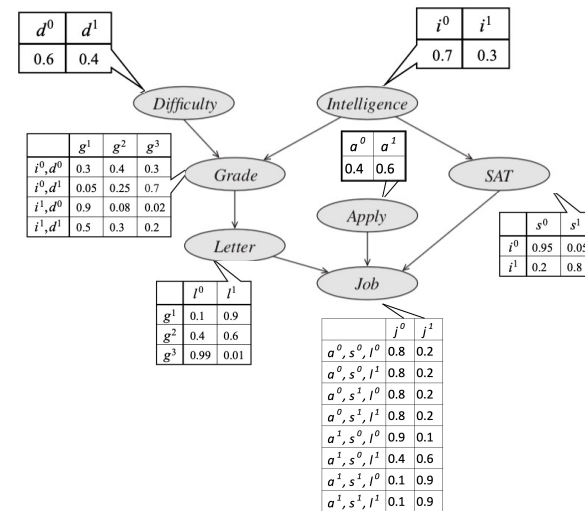
# Two Paradigms of AI



■ Shaded boxes indicate components that can learn from data

# Reasoning

- It is intuitive and subjective
- Knowledge-based Approach
  - Hard-code knowledge
  - Need to have deep knowledge of the problem domain



$$P(I = 1 | J = 1) = 0.0629 / (0.0629 + 0.0387) = 0.619$$

# Deep Learning


1. Neural networks:
  - Inspired by neural networks of the brain
2. Learning:
  - Machines to discover rich/useful representations
3. Depth:
  - As a composition of learned features and functions

# Deep Learning


Understand the world as a hierarchy of concepts

- How these concepts are built on top of each other is deep, with many layers
- Weights  $\mathbf{w}$  learnt by gradient descent


$$\mathbf{w}^{(t+1)} = \mathbf{w}^t - \varepsilon \nabla_{\mathbf{w}} L(\mathbf{w}^t)$$



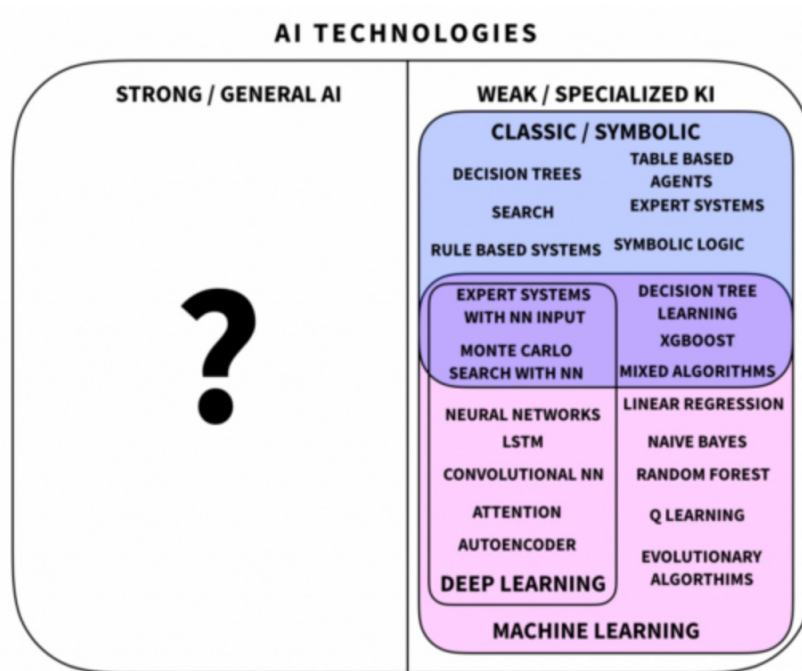
Discriminative Model  
CNN



Generative Model  
VAE



# AI Technologies



# Current status of AI

- Human capabilities still lacking in AI
  - Adaptability
  - Generalizability
  - Robustness
  - Explainability
  - Abstraction
  - Common sense
  - Causal reasoning
- Current research addresses these limitations

# Research Goal

- Can we embed the causal components?
- Complex and seamless integration of learning and reasoning supported by both implicit and explicit knowledge
- Debate: whether end-to-end DL can achieve this goal or whether we need to integrate ML with symbolic and logic-based AI techniques
- Hypothesis: Integration route is most promising




# Cognitive Approach

- Comprehension of how humans have, and have evolved to obtain, these advanced capabilities can inspire innovative ways to imbue AI systems with these competencies
- Study and exploit cognitive theories of human reasoning and decision making as a source of inspiration for the causal source of these capabilities,

# Kahneman's Theory

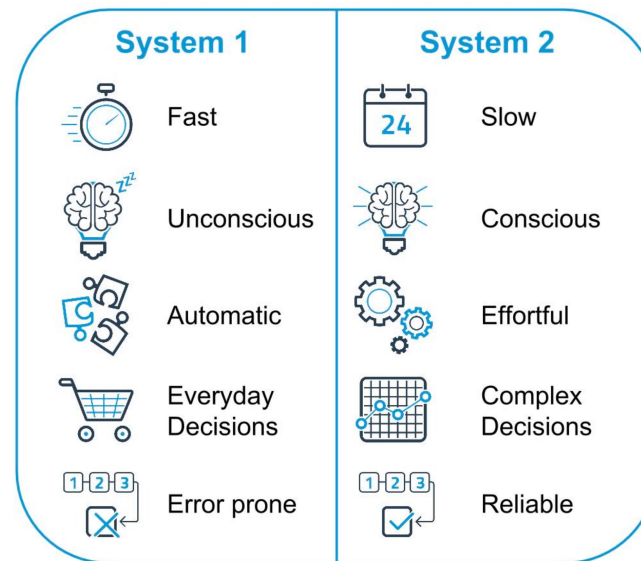

- Human decisions are supported and guided by the cooperation of two main kinds of capabilities
  - System 1: Thinking Fast
    - Provides tools for intuitive, imprecise, fast, and often unconscious decisions
  - System 2: Thinking Slow
    - Handles more complex situations where logical and rational thinking is needed to reach a complex decision



# Kahnemann on Two Systems

[Kahnemann](#)


# Summary of two systems



Source:  
<https://medium.com/@ryansheffer/founders-need-to-think-slow-move-fast-6b683e94c110>

# Further Summary

<b>System 1</b>	<b>System 2</b>
drive a car on highways	drive a car in cities
come up with a good chess move (if you're a chess master)	point your attention towards the clowns at the circus
understands simple sentences	understands law clauses
correlation	causation
hard to explain	easy to explain



# System 1 Characteristics

- Guided mainly by intuition not deliberation
  - Fast answers to very simple questions.
  - Answers are sometimes wrong, due to unconscious bias or reliance on heuristics
- Usually do not have an explanation
- Able to build models of the world
  - Although inaccurate/imprecise, can fill knowledge gaps through causal inference
  - Allow us to respond well to the many stimuli of our everyday life

# Examples of System 1 Tasks

1. Finding the answer to a very simple arithmetic calculation
2. Reaching out to grab something that is going to fall
  - We use our system 1 about 95% of the time when we need to make a decision

# Role of System 2

- When the problem is too complex for System 1
- System 2 kicks in
- Solves it with access to:
  - Additional computational resources
  - Full attention, and
  - Sophisticated logical reasoning.



## Example of System 2 Task

- Solving a complex arithmetic calculation, or a multi-criteria optimization problem
- To do this, humans need to be able to recognize that a problem goes beyond a threshold of cognitive ease and therefore the need to activate a more global and accurate reasoning machinery.
- Hence, introspection is essential in this process

# When is System 2 invoked

- A new/difficult problem is handled by system 2
- Certain problems over time pass on to system 1
  - Procedures used by system 2 are used to accumulate examples that system 1 can later use readily with little effort
    - Problems, initially solvable only by system 2 reasoning tools, can become manageable by system 1
    - E.g., Reading text in our own native language
- However, this does not happen with all tasks.
  - Finding correct solution to complex arithmetic

# System 2 is supported by System 1

- System 2 does not usually work by itself
  - Supported by system 1 in its elaborate calculations
  - When searching for a solution in a very large solution space, system 2 does not usually explore the whole search space but may employ heuristics that are provided by system 1 and that help in focusing the attention only on the most promising parts of the space.
- This allows system 2 to work with manageable time and space

# Data-level causality by System 1

- System 1 also capable of basic causal reasoning
  - Allows it to build a (possibly imprecise and biased) model of the world even if it has incomplete knowledge, and to use it to tackle simple tasks
- System 1 data-level causality skills support more complex/accurate reasoning of system 2 on more complex problems

# Multi-tasking by System 1

- System 1 reasoning:
  - At various levels of abstraction
  - Adapt to new environments
  - Generalize from experience to other problems
  - Easily multi-task
- System 2 is sequential
  - Requires full attention to devise/execute appropriate procedure for complex or new problem
  - Only one complex task at a time handled

# Local vs Global

- System 1 works at a local level
  - Activates only a specific part of the brain, e.g., recognize a familiar face, or when we speak
- System 2 involves more regions of the brain
  - Combines contributions-- works at a global level
- Multiple specialized agents/components/skills supports the functioning of both systems
  - Not as multi-agent architecture terminology but as metaphors for two broad classes of information processing

# Experimental Justification

- Two forms of consciousness (Graziano)
  1. I(Information)-consciousness
    - Related to System 2
  2. M(Mystery)-consciousness
    - Related to System 1

# I(Information)-consciousness

- Ability to solve (possibly complex) problems
  - By recognizing necessary processing steps in specific (even new) context, to tackle a desired problem
- Related to System 2's high-fidelity processing mode
  - Considering a problem and harnessing the relevant faculties of our cognition to devise a plan to solve it



# M(Mystery)-consciousness

- Ability to build a simplified, approximate, and subjective model of peoples', both ourselves and others, mind, beliefs and intentions.
- Such low-fidelity model building can be linked to system 1, as system 1 is able to form a rapid but usually inexact model of the world

# Humans excel at I-consciousness

- While many animals and primates possess various levels of M-consciousness, and also limited forms of I-consciousness, humans excel at I-consciousness
- M-consciousness can even sustain in presence of limited I-consciousness, yet, a sophisticated I-consciousness needs to rely on M-consciousness:
  - without a model of the mind (of self and other agents), it is difficult to devise how to adapt to new environments and solve complex problems.

# Summary

- Neural networks are what appear to make humans intelligent
- Two paradigms of AI : Causal Reasoning and Deep learning
- Can we embed causal reasoning into neural networks?
- I-consciousness and M-consciousness are two levels
- I-consciousness is what humans excel at