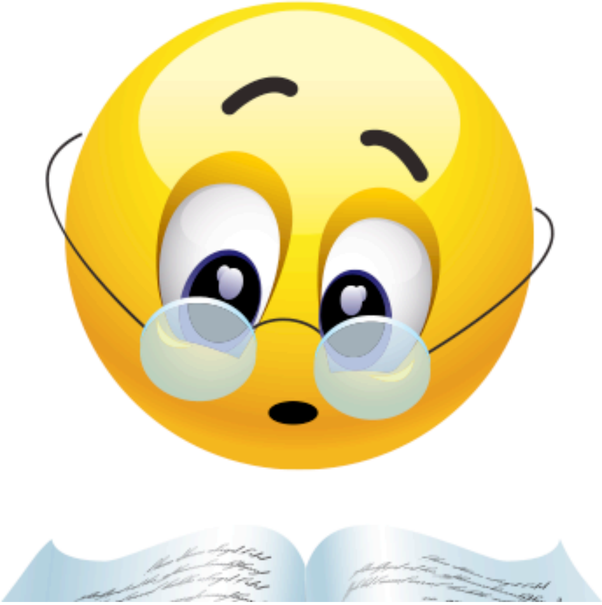

Learning in Order to Reason: The Approach

Dan Roth*
Dept. of Appl. Math. & CS,
Weizmann Institute of Science
Rehovot 76100
Israel
danr@wisdom.weizmann.ac.il
http://www.wisdom.weizmann.ac.il/~danr

A presentation by **Aditya M. Kashyap**

Introduction

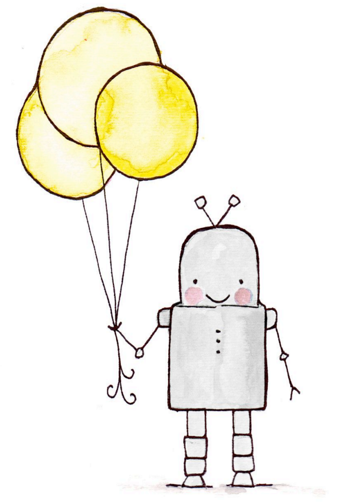

- Consider a baby robot
- If it was human → provided safe Environment

  → safe to learn about the environment
- Expected to have “full functionality” after this “grace period”
- Performance depends on the environment
  - On the amount of time it interacted with it
- Central role of learning in cognition is widely acknowledged
- Early theories of intelligent systems → Learning $\perp\!\!\!\perp$ Reasoning

Traditional Learning Theory

Traditional Reasoning Theory

Introduction

- This paper presents a new framework for the study of reasoning
- FRAMEWORK (Learning to Reason)
  - Given
    - Agent given access to its favorite learning interface.
    - Given a grace period
  - Expected
    - Knowledge Base (KB) of the world W
    - Reasoning evaluation after “grace period”
  - Could also learn in an online fashion
- KB not given to the agent → It must be learned

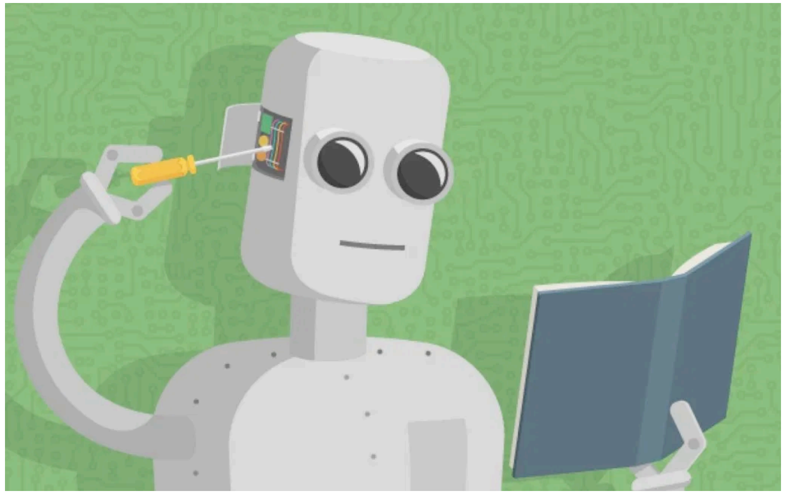

Introduction

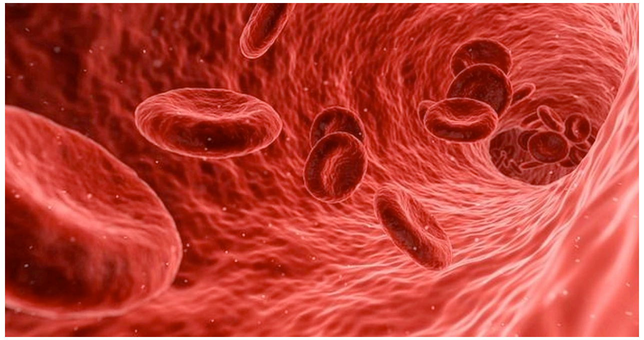

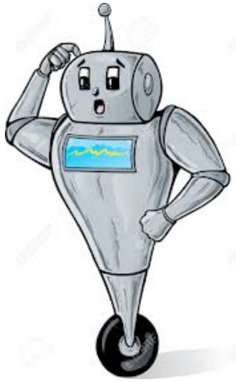

- Assumption:
  - Reasoner need not answer all possible queries
- Performance measured relative to the environment
- Quality of KB depends on whether:
  - It is efficiently learnable
  - It supports efficient reasoning performance

Introduction

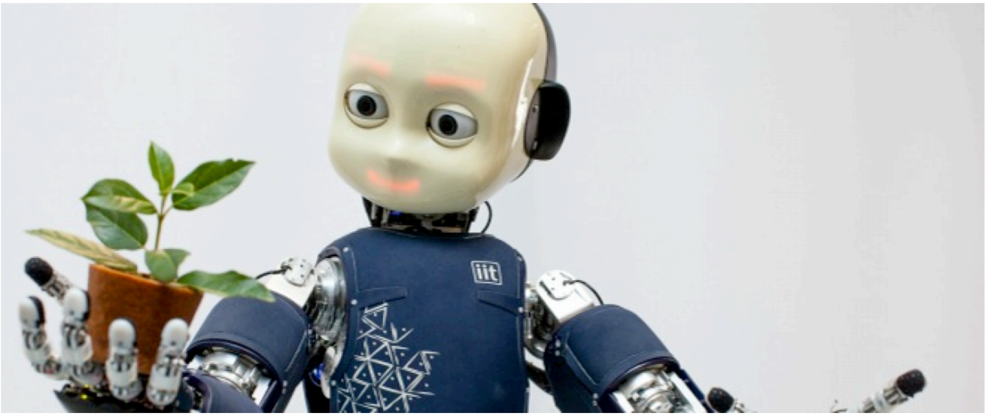

- Claim:
  - Interaction with the world → Agent gains additional reasoning power

- This claim is proven for cases where:
  - Reasoning from a given KB don't have efficient solutions (traditional reasoning)
    - Deduction breaks down into satisfiability which is NP hard to solve
  - Learning representations of the world don't have efficient solutions (traditional learning)
    - There are no polynomial algorithms for learning general Boolean functions

Motivation

- General framework to study reasoning in intelligent systems → KB
- KBs can occur in different forms. 2 examples:
  - A set of logical rules
  - A probabilistic network
- Normally, we do not consider:
  - the method in which the knowledge is acquired, or
  - whether this should influence how the performance of the reasoning system is measured

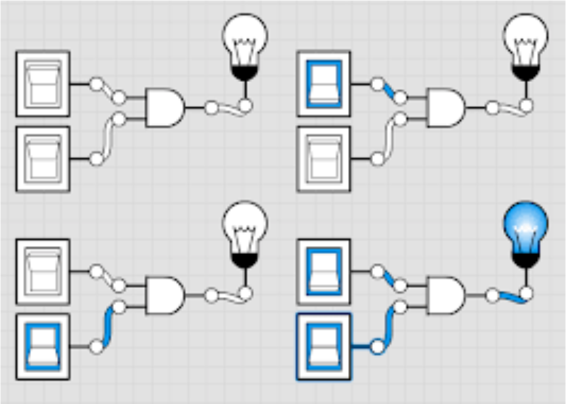

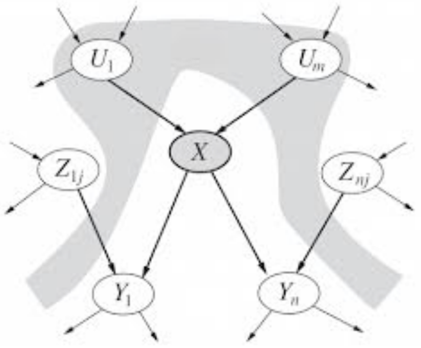

Motivation

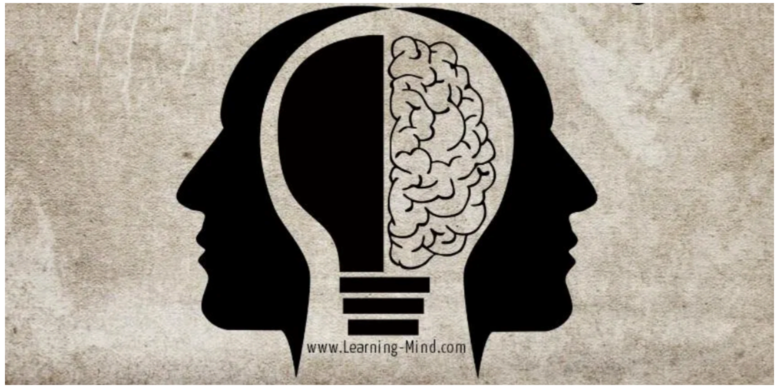

- This is because of the following observation:

  Observation: *If there is a learning procedure that can learn an exact description of the world in representation R, and there is a procedure that can reason exactly using R, then there is a complete system that can learn to produce "intelligent behavior" using R.*

- Separate study of learning and the rest in cognition is partly motivated by the assumption of the converse being true

  i.e If there is a system that can Learn to Reason, then there is a learning procedure that can learn a representation of the world, and reason with it.

- Computational considerations, however, render the traditional self-contained reasoning approach inadequate.
- Recent works in reasoning aim at identifying classes of limited expressiveness, with which one can perform some sort of reasoning efficiently.

Quick Diversion: Propositional Logic

- **Literals**: An atomic formula (Binary Variables). Ex: Rain, Win
- **Clause**: Disjunction (OR) of literals. Ex: $l_1 \vee \cdots \vee l_n$
- **Conjunctive Normal Form (CNF)**: *(AND of ORs)*
  - Conjunction of 1 or more clauses

    Examples: $(A \vee C) \wedge (B \vee C)$

    $A \wedge (B \vee D) \wedge (B \vee E)$
- **Disjunctive Normal Form (DNF)**: *(OR of ANDs)*
  - Disjunction of 1 or more conjunctions

    Examples: $(A \wedge \neg B \wedge \neg C) \vee (\neg D \wedge E \wedge F)$

    $(A \wedge B) \vee C$
- Why is this important? CNFs and DNFs are often how the KBs and the queries are represented.

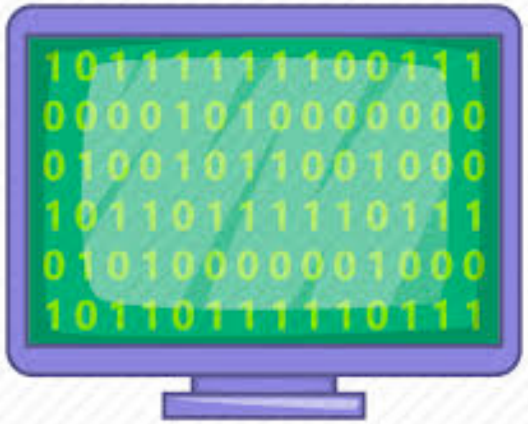

Motivation

- Perhaps the most open questions in learning theory (1996):
  - Learnability of CNF and DNF formulas
- Even if there were solutions for the learning task, it would not be relevant for the reasoning task, because:
  - CNF expression output of learning algorithm → computationally hard
  - Learning DNF → not interesting → as it does not relate to rule-based representation
  - DNF → “list” of examples and not a set of rules

    (conjunctions; which are partial variables since not all variables are specified)

Learning to Reason Principles

1. Intelligent systems are not omniscient
   - Agent must function in a complex environment
   - environment maybe hard to exactly model
   - We care about how well it performs on a fairly wide, restricted set of tasks
   - Therefore, requirements from the reasoning stage may be relaxed

2. The goal of the learning stage depends on the required functionality
   - Learning stage is not evaluated on how well it models the world, but rather how well it supports the required functionality

Learning to Reason Principles

3. The interaction with the environment is a key issue
   - Unlike previous frameworks like the *traditional learning framework* and the *traditional reasoning framework,* the agent interacts with the world.

4. The knowledge representation used may depend on the functionality
   - KB is evaluated based on:
     - Its learnability
     - How well it supports inference
   - KB is not evaluated on its comprehensibility

Learning to Reason Principles

5. The performance of the agent is measured with respect to the world it functions in, and not in any absolute terms
   - The world in which the agent performs its task is the same world which supplies the agent the information when learning
   - The performance of the agent is evaluated on a collection of tasks that are “common” or “relevant” in the environment

6. Rigor and efficiency
   - The aim was to define a framework that was rigorous and amenable to analysis
   - Learning to Reason is usually required to be done in time that is polynomial in the natural complexity parameters

General Framework

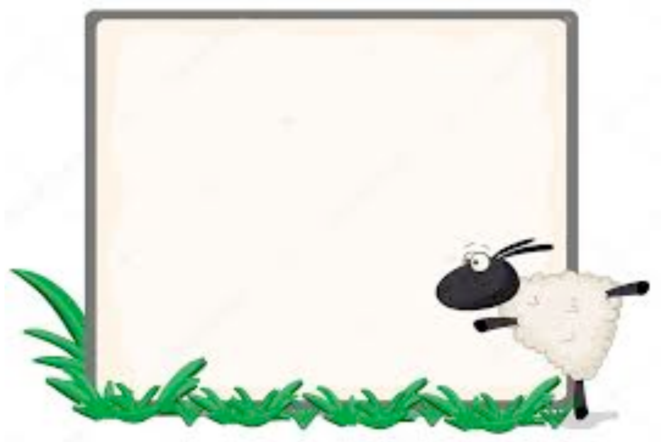

- Two previous frameworks:
  - Traditional Reasoning framework → Deductive task
  - Traditional Learning framework → PAC Learning
- A question concerns the possibility of using the past two frameworks to create the “Learning to Reason” framework.
- Some problems are:
  - Output of learning algorithm doesn’t support effective reasoning
  - Even if you can combine them, it exhibits limitations as discussed previously

Deductive Reasoning

- Reasoning from one of more statements to reach a logically certain conclusion
- Example:
    1. *All men are mortal* ***(major premise)***
    2. *Pythagoras is a man* ***(minor premise)***
    3. *Pythagoras is mortal* ***(conclusion)***

- Most striking evidence of usefulness of “Learning to Reason” framework
- Separate learning and reasoning tasks are not tractable
- “Learning to Reason” algorithms can be used instead

Deductive Reasoning

- Learning to Reason without Reasoning:
  - Consider the reasoning problem $\mathrm{W} \models \alpha$.
    - **W**: some CNF formula
    - **α**: log(n)CNF formula (a CNF formula with at most log(n) literals in each clause

  - When W has a polynomial size DNF, there is an exact and efficient Learning to Reason algorithm for this problem

  - The traditional reasoning problem (with a CNF representation as input) is NP hard

Deductive Reasoning

- Learning to Reason without Learning to Classify:
  - Consider the reasoning problem $\mathrm{W} \models \alpha$.
    - **W**: some Boolean formula with a polynomial size DNF
    - **α**: log(n)CNF formula (a CNF formula with at most log(n) literals in each clause
  - There is an exact and efficient Learning to Reason algorithm for this problem
  - The class of Boolean formulas with a polynomial size DNF is not known to be learnable in the traditional (Learning to Classify) sense

Deductive Reasoning

- Learning to Reason algorithms that use formulas as their knowledge representation are also considered
  - Consider the reasoning problem $\mathrm{W} \models \alpha$.
    - **W**: any Boolean formula that has Horn approximation with a polynomial size
    - **α**: Horn Expression
  - Horn clause: Disjunction of literals with at most 1 positive literal

$$(A \vee \neg B) \wedge (\neg A \vee \neg C \vee D)$$

  - There is an exact and efficient Learning to Reason algorithm
  - Problem of Learning W exactly is not known
  - Problem of Reasoning from a representation of W is not tractable

Deductive Reasoning

- The new positive results are made possible by the combination of several features based on the principles mentioned previously:
  - The inference problem is relaxed by restricting classes of queries considered
  - W does not have be modeled exactly.
    - Instead, a least upper-bound approximation of W can be used.
    - This approximation is a function closest to W in the class of queries

Least Upper-Bound Approximation

W does not have be modeled exactly.

- Instead, a least upper-bound approximation of W can be used.

- This approximation is a function closest to W in the class of queries

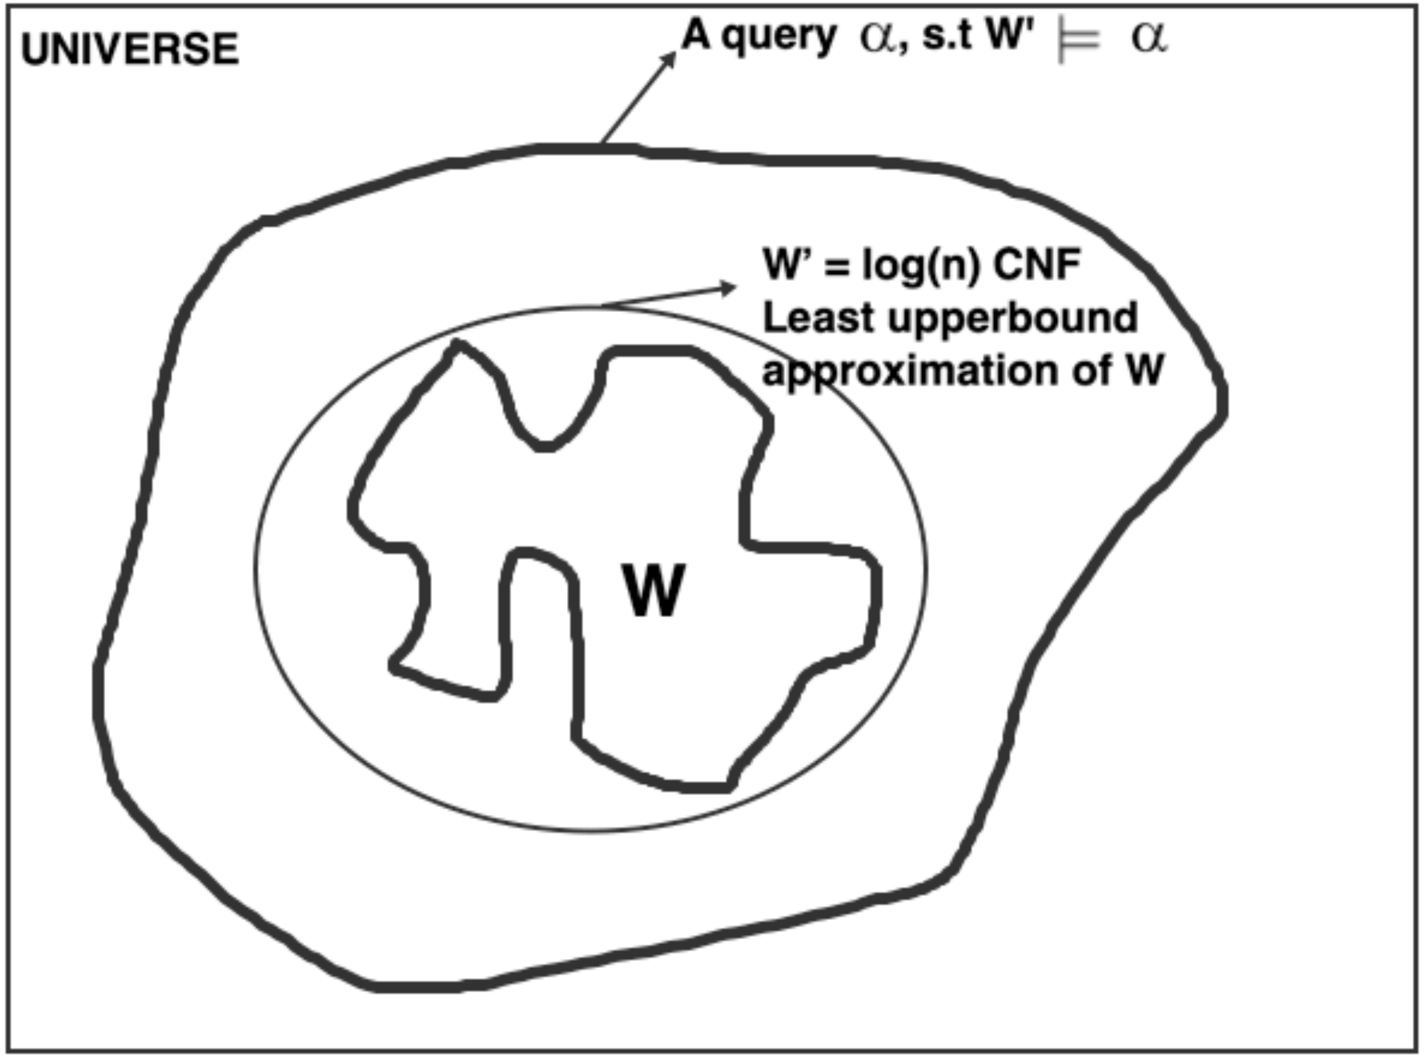

Other Learning to Reason Results: *Abductive Reasoning*

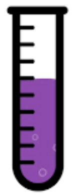

- Logical inference task that starts with an observation/ set of observations and then seeks to find the simplest and most likely explanation
- Unlike deductive reasoning, it does not positively verify the explanation

- Example

The scientist observes the test tube and sees the chemical turn purple. She abduces that either there is potassium in the sample, or her colleague is playing yet another prank on her.

- Using model-based representations for abductive reasoning shown by previous work, *Learning to Reason* algorithms can be used

Other Learning to Reason Results: *Default Reasoning*

- A form of non-monotonic reasoning where plausible conclusions are inferred based on general (default) rules which may have exceptions

- Example:
    1. *Tweety is a bird*
    2. *Birds fly* ***(This has exceptions since not all birds fly)***
    3. *Therefore, Tweety flies*

- Similar to Abductive reasoning, using model-based representations for default reasoning shown by previous work, Learning to Reason algorithms can be used

Other Learning to Reason Results: *Reasoning with Partial Assignments*

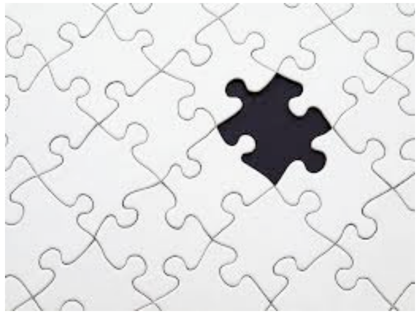

- Partial assignments may occur because the world cannot be learned efficiently
- More natural way of thinking about learning.

- Complete Observation → Variables take on values in {0,1}
- Partial Observation → Variables take on values in {0,1,*}
- Example v = {1 * 0} → v_1 is True, v_2 is unknown and v_3 is False
- Total Assignment → case where the value of every variable is known (x = {0,1,1})

Other Learning to Reason Results: *Reasoning with Partial Assignments*

- **Universal Interpretation**: For all possible extensions of v (partial) to v' (total), W(v') = 1
  v = {1 * 0} → {1 1 0} and {1 0 0} hold in the world W

- **Existential Interpretation:** There exists an extension of v to v', such that W(v') = 1
  v = {1 * 0} → {1 1 0} or {1 0 0} hold in the world W

- **Abbreviated Interpretation:** All variables assigned * in v are assumed to have value 0 in v'
  v = {1 * 0} → {1 0 0} hold in the world W

- The deductive reasoning approach has been extended to handle such cases

- **An important conclusion:** When dealing with partial information in the interface, classification problems become deductive problems

Other Learning to Reason Results: *Non-Monotonic Reasoning*

- Example:
    1. *Typically birds fly*
    2. *Penguins do not fly*
    3. *Tweety is a bird*

It is plausible to conclude that “Tweety flies”. However, if the following sentence is added:

4. *Tweety is a penguin*

The previous conclusion must be retracted and, instead, the new conclusion “Tweety does not fly” will hold.

Other Learning to Reason Results: *Non-Monotonic Reasoning*

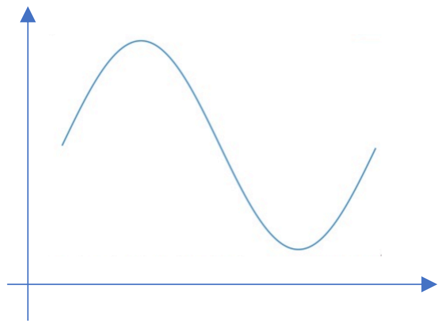

- Problem of reasoning from incomplete information can be presented as a problem of learning attribute functions over a generalized domain

- These representations can be learned efficiently, yielding Learning to Reason algorithms

Other Learning to Reason Results: *Learning to take actions*

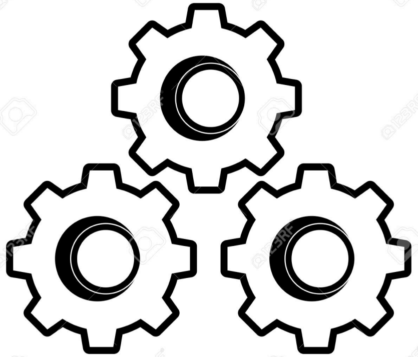

- Supervised learning problem
- The learning problem is in a dynamic, stochastic environment
- Agent receives observations and learns from it an acting strategy
- Implements Learning to Reason principles

- Most important difference: The agent acts in the world, thereby changing it

Conclusion

- Presented the Learning to Reason framework, which is intended to overcome some of the fundamental problems in earlier approaches of reasoning
- In some cases like Deductive Reasoning, the new framework allows for successful Learning to Reason algorithms
- Provided the main principles on which this framework is built
- Displayed a few areas in which the *Learning to Reason Framework* is applied

Learning to Reason with Neural Module Networks

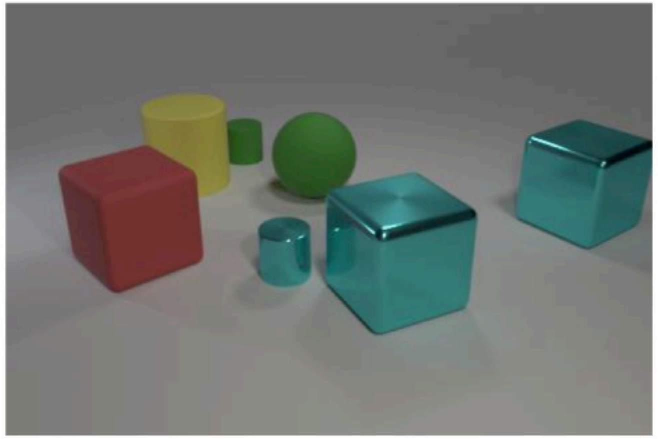

What color is the thing with the same size as the blue cylinder?

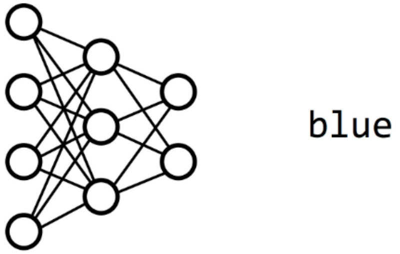

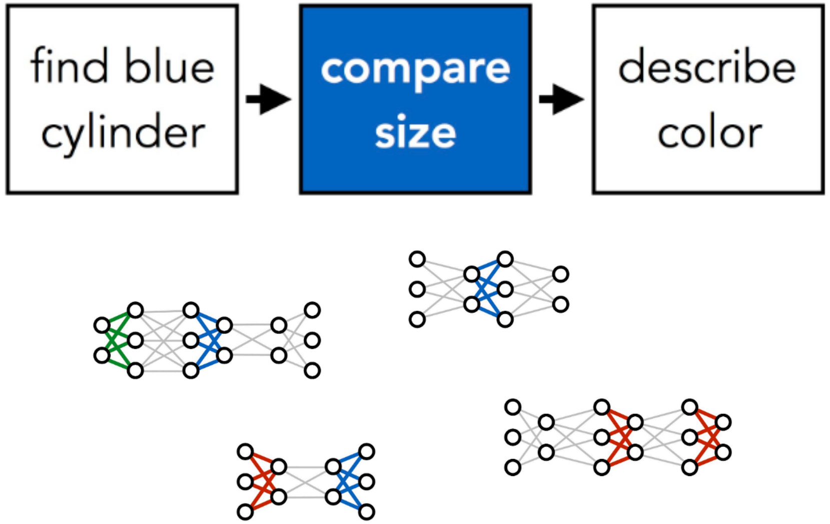


That's all Folks!