Deep Questions

Richard Sutton

Reinforcement Learning and Artificial Intelligence Laboratory
Department of Computing Science
University of Alberta

Questions and answers

- Every learning system involves both a *question* and an *answer*, but usually we focus on the answers
  - how good are the answers?
  - how can we find the best answers?
- Questions are overlooked. They are too obvious
- We need to pay more attention to questions
  - Questions are a major part of the problem of representation learning, maybe half of it

Examples of learning questions

and their formalizations

- Is there a cat in this photo? Where is the cat?
- What will the next frame of this video be?
- What is the probability that we will win this chess game?
- What will the final score be from this Atari game state?

- What will the label be? (supervised learning)
- What is the expected discounted cumulative future reward from this state? (value-function learning)
- What will the next input be (unsupervised distribution learning)

Need a compositional way of building up complex questions

- We know how to build up answers compositionally
  - in the layers of deep networks
  - features of features of features
  - to form arbitrarily complex answers

- We need an analogous compositional complexity on the question side
  - questions about questions about questions

- Then we will be beginning to address the second half of the problem of representation learning

Examples of predictive questions and answers

- Q: How much rain will fall in the next 24 hours? A: 0.5 centimeters
- Q: Will i win this chess game? A: with Probability 0.9
- Q: What will the dow jones index be at the end of the year? A: 18,000

- Q:What will be the discounted sum of rewards from here forward? A: 5.7 (or whatever the value of the state is)

- The question describes the procedure for calculating the target
- The answer is the expected value of the target (say)
- The answer process is familiar; it might be a deep/neural network
- The question process is less familiar, possibly more important

A policy-conditional predictive question

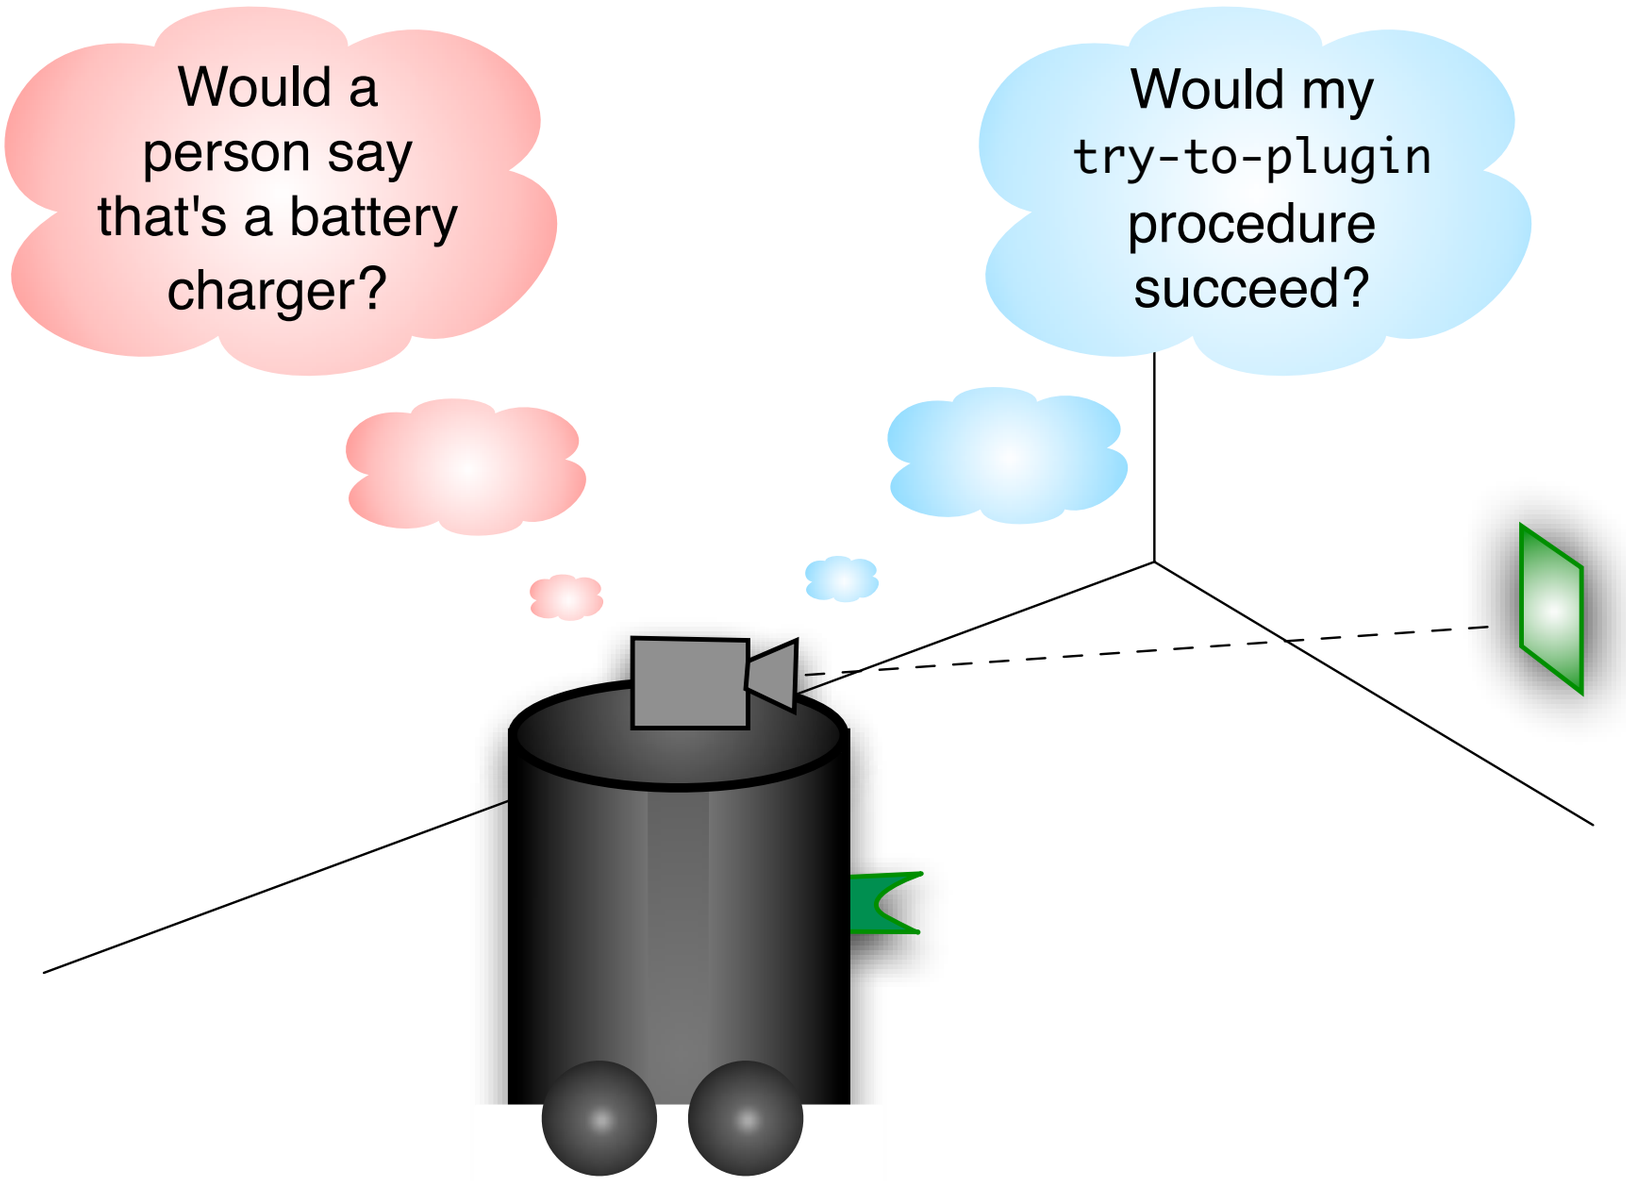

Where we are now — the state of machine learning

- With deep learning, we have gotten good at answering the formal question: What is the label of this example?
  - But the original question (is this a cat?) is known only to the people who were asked to label the training set
- We also explore “unsupervised learning” where the labels are constructed automatically
  - these are the cases in which the question is in the machine—and thus learning can proceed without external teachers—but still the question is hardcoded and cannot be changed or manipulated by the machine
  - Reinforcement learning (value-function learning) is also like this

The next step is for the machine to choose its questions

- Questions being in the machine is good, gives us self-labeling
- But the real real reason for questions to be in the machine is so that the machine can change them
  - try out different questions
  - search through them, evaluate them
  - determine which are most useful
- There are many senses of useful
  - their answers may be good features
  - they may be good subtasks, directing feature construction
  - they may form a high-level model for planning, etc

For a machine to choose questions, it needs a formal language of questions

- So that there is a space to explore
- General Value Functions (GVFs) are one such language, good for multi-step predictive questions answerable by RL/TD methods
- Temporal-difference networks are another
- I am sure others are possible, maybe even non-predictive questions

Answers are relatively easy to represent; it's questions that are hard

- e.g., flipping a coin
  - Question: what is the probability of heads
  - Answer: 0.5
- How to represent *flipping*, *coin*, and *heads*?

- What is *heads*?
- It's *not* a sensation
- It's another *prediction*
- We need to be able to ask questions about *predicting predictions*
- We need *compositionality*
  - predictions that can be built out of other predictions
- We need *abstraction*
  - predictions that capture similarities

Need a compositional way of building up complex questions

- We know how to build up answers compositionally
  - in the layers of deep networks
  - features of features of features
  - to form arbitrarily complex answers

- We need an analogous compositional complexity on the question side
  - questions about questions about questions

- Then we will be beginning to address the second half of the problem of representation learning