

# DSC190: Machine Learning with Few Labels

## Overview

**Zhiting Hu**

Lecture 1, September 27, 2024

**UC San Diego**

**HALICIOĞLU DATA SCIENCE INSTITUTE**

# Logistics

- Class webpage: <http://zhiting.ucsd.edu/teaching/dsc190fall2024>

DSC190-Fall2024

Logistics Lectures Homework Project



## Machine Learning with Few Labels

DSC 190 • Fall 2024 • UC San Diego

Machine learning is about computational methods that enable machines to learn concepts from experience. Many of the successful results of machine learning rely on learning with massive amounts of data labels. However, in many real problems we do not have enough labeled data, but instead have access to other forms of experience, such as structured knowledge, constraints, feedback signals from the environment, auxiliary models from related tasks, etc.

This course focuses on those learning settings with few labels. This course is designed to give students a holistic understanding of related problems and methodologies (such as **large language/multi-modal models, world models, self/weakly/un-supervised learning, transfer learning, meta-learning, reinforcement learning, adversarial learning, knowledge constrained learning, panoramic learning**), different possible perspectives of formulating the same problems, the underlying connections between the diversity of algorithms, and open questions in the field.

Students will read, present, and discuss papers, and complete course projects.

# Logistics



Instructor: **Zhiting Hu**

Email: [zhh019@ucsd.edu](mailto:zhh019@ucsd.edu)

Office hours: Mon 3:30-4:30pm

Location: HDSI 442



TA: **Yuheng Zha**

Email: [yzha@ucsd.edu](mailto:yzha@ucsd.edu)

Office hours: Thu 3:30-4:30pm

Location: HDSI 437

- Discussion forum: Piazza
- Homework & writeup submission: Gradescope

# Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)

# Logistics: grading

- 2 Homework assignments (30% of grade)
  - Theory exercises, implementation exercises
  - 3 total late days without penalty
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)

# Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
  - Each **individual** student will give an oral presentation on a research paper
    - 10 mins = 8 mins presentation + 2 mins QA (*tentative*)
      - Timing -- hard time constraint: if you run over the expected time limit (8min), there will be no QA session for your presentation, and thus no credits for the QA component
    - **Critical thinking:** discuss both strengths and limitations of the paper
    - Sign up in a google sheet (TBA)
    - Design quiz questions for audience
  - **Peer grading:** other students will rate and give feedback (5% of grade)
  - Starting later part of the quarter, after the class size is stabilized
- Course project (46%)
- Participation (4%)

# Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
  - 3 or 4-member **team** to be formed and sign up in a google sheet (TBA)
  - Designed to be as similar as possible to researching and writing a **conference-style paper**:
    - Due to tight timeline, fine to use synthetic/toy data for proof-of-concept experiments + explanation of theory/intuition of why your approach is likely to work
  - **Proposal** : 2 pages excluding references (10%) -- **due in 2 or 3 weeks (TBA)**
    - Overview of project idea, literature review, potential datasets and evaluation, milestones
  - **Midway Report** : 4-5 pages excluding references (20%)
  - **Presentation** : oral presentation, 15-20mins (20%)
    - Peer grading (5%)
  - **Final Report** : 6-8 pages excluding references (50%)

# Logistics: grading

- 2 Homework assignments (30% of grade)
- Paper presentation (20%)
- Course project (46%)
- Participation (4%)
  - Submission of quiz answers and feedback on paper/project presentations
  - Contribution to discussion on Piazza
  - Completion of mid-quarter and final course evaluation
  - Any constructive suggestions



# Have Large Language Models (LLMs) Solved ML/AI?



What are a couple of meals I can make with this?

Based on the items visible in the fridge, here are a couple of meal ideas:

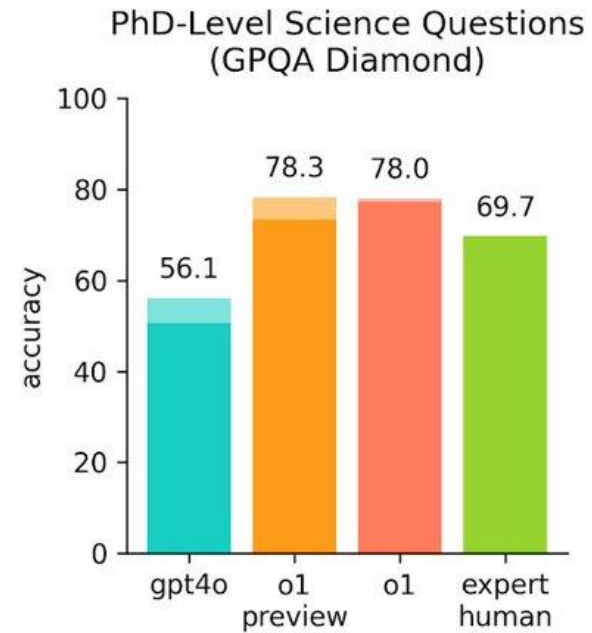
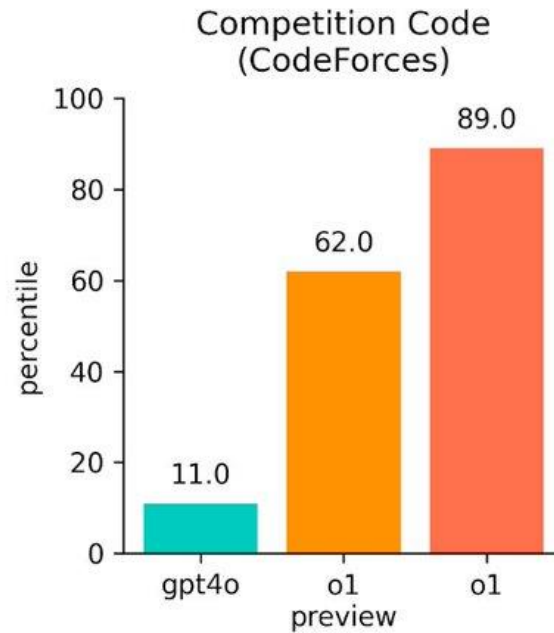
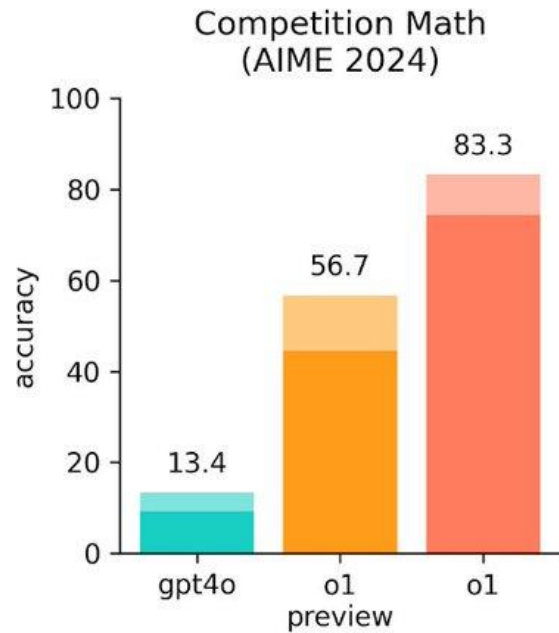
- 1) Yogurt parfait: Layer yogurt, strawberries, and blueberries in a bowl or jar. Top with honey or nuts for added texture and flavor.
- 2) Carrot and hummus wrap: Spread hummus on a tortilla or wrap. Top with shredded carrots and any other desired veggies. Roll up and enjoy!

**GPT-4**



**Sora**

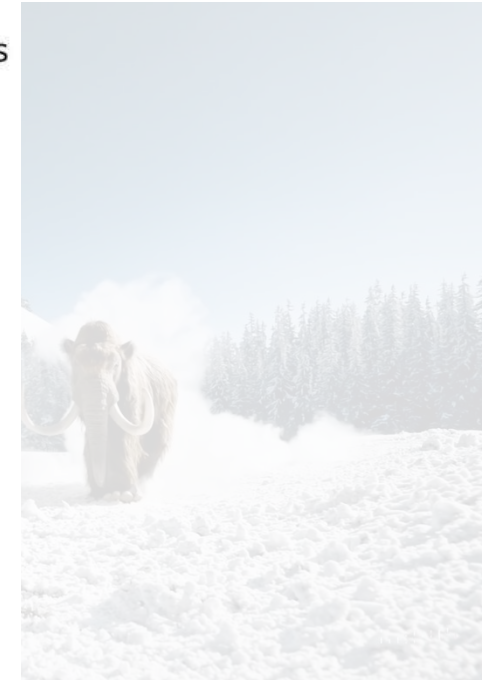
# Have Large Language Models (LLMs) Solved ML/AI?



**OpenAI o1**

**Sora**

**GPT-4**



# Have Large Language Models (LLMs) Solved ML/AI?

A year ago

Blog

## Introducing ChatGPT

We've trained a model called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests.

[Try ChatGPT ↗](#)

[Read about ChatGPT Plus](#)

AI doomsday imagined by  
ChatGPT + DALL-E3



# LLMs are great... but fundamentally limited

Language  
Reasoning

Embodied  
Reasoning

Social  
Reasoning

# Limitations in LLMs

- Language reasoning (e.g., math, logics, common sense, ...)

Explain the process of  $3452 \times 253$

 GPT-4

```
3452
x 253
-----
10356 (3452 x 3)
6904 (3452 x 5, shifted one position to the left)
+3452 (3452 x 2, shifted two positions to the left)
-----
873256
So,  $3452 \times 253 = 873256$ .
```

*(correct answer:  
873356)*

# Limitations in LLMs

- Language reasoning (e.g., math, logics, common sense, ...)

Problem: Track the same group of students over time

Here is a plan for surveying the mental health of undergraduate students on our campus. This year, we will ask freshmen; next year, sophomores; the year after, juniors; and the year after that, seniors.

**Are there any problems with this plan?**

 GPT-4

...

**1. Longitudinal Data and Cohort Effects:** This approach **doesn't track the same group** of students over time...

2. ...

Language Reasoning

Embodied Reasoning

Social Reasoning

# Limitations in LLMs

- Language reasoning (e.g., math, logics, common sense, ...)

Human: Internal **world model**

- Physical common sense
- Object permanence
- Spatial relation (e.g., support)

Emily found a desk and place **the cell phone** on top of it. *[Irrelevant Actions]* ... putting the **lime** down next to the cell phone. *[Irrelevant Actions]* ... She final finally put an **apple** on the desk. **How many items are there on the desk?**

 GPT-4

There are **two** items

*(correct answer: three)*

Language  
Reasoning

Embodied  
Reasoning

Social  
Reasoning

# Limitations in LLMs

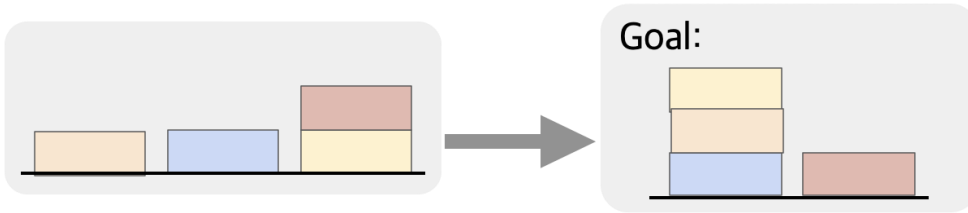
- Building embodied agents requires **embodied reasoning**





# Limitations in LLMs

How to move the blocks to the goal state?



## LLMs: Autoregressive plan generation



1. Pick up the orange block.
2. Stack it on the blue block.
3. Pick up the yellow block. ❌
4. Stack it on the orange block.
5. Pick up the red block.
6. Put it on the table.

### Invalid Action!

The yellow block is still under the red one.

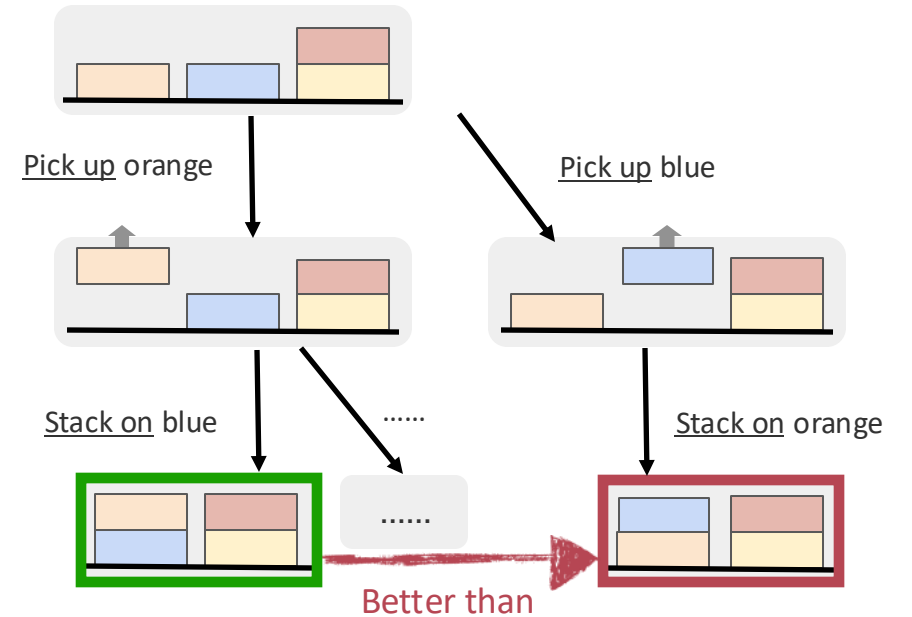
Language Reasoning

Embodied Reasoning

Social Reasoning

## Human: strategic planning

- Internal **world model** to predict states
- **Simulation** of alternative plans
- Assess outcomes to refine/pick the best



Language  
Reasoning

Embodied  
Reasoning

Social  
Reasoning

# Limitations in LLMs

- Building **socially** intelligent systems requires **social reasoning**

AI Drivers



AI Assistants



AI Coworkers



AI Companions



AI Teachers



# Limitations in LLMs

- Social reasoning: understanding humans

Explain why this is funny



GPT-4V

... The final panel reveals the punchline: the robot has merely produced a pile of crumpled paper, just like the human did, suggesting that **the robot also suffers from writer's block** ... highlighting a situation where the human and the AI are **equally challenged**



Language  
Reasoning

Embodied  
Reasoning

Social  
Reasoning

## Limitations in LLMs

- Social reasoning: understanding humans



Does this person need help?

# Limitations in LLMs

- Social reasoning: understanding humans

Social reasoning also requires a **world model**: Humans' actions are constrained by physics



Does this person need help?



GPT-4V

It looks like the person in the image is carrying a heavy load of boxes, which might suggest they are in the process of moving. While **I can't determine the actual need for help without context or the ability to interact with the situation,**

...

Language  
Reasoning

Embodied  
Reasoning

Social  
Reasoning

## Limitations in LLMs

- Social reasoning: understanding humans



Is the person on the right helping?

# Limitations in LLMs

- Social reasoning: understanding humans

Understanding humans  $\neq$   
understanding actions



GPT-4V

The person on the right **appears to be helping** by holding the chair, probably to steady it or to move it into place. The way they are reaching out and positioning themselves suggests they are assisting the person who is reaching across the table.



Is the person on the right helping?

Correct action: **Holding the chair**  
Wrong intent: **Helping**

Language Reasoning

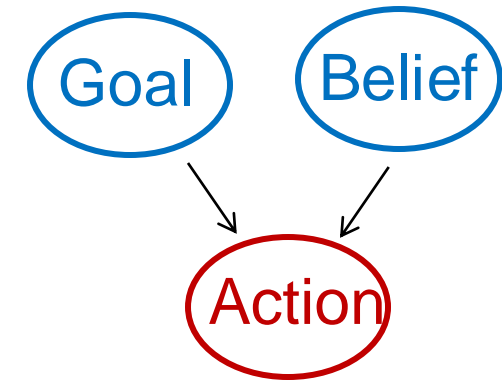
Embodied Reasoning

Social Reasoning

# Limitations in LLMs

- Social reasoning: Theory of Mind

Reasoning about **hidden mental** variables that drive **observable actions**



**Goal:** Sit down

**Belief:** The chair was not moved

**False belief**

caused by the person on the right



Language Reasoning

Embodied Reasoning

Social Reasoning

# Limitations in LLMs

- Social reasoning: Theory of Mind

## Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs

Maarten Sap<sup>◆</sup> Ronan Le Bras<sup>♣</sup> Daniel Fried<sup>◇</sup> Yejin Choi<sup>♣♥</sup>

<sup>♣</sup>Allen Institute for AI, Seattle, WA, USA

<sup>◇</sup>Language Technologies Institute, Carnegie Mellon University, Pittsburgh, USA

<sup>♥</sup>Paul G. Allen School of Computer Science, University of Washington, Seattle, WA, USA

## Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks

Tomer D. Ullman  
Department of Psychology  
Harvard University  
Cambridge, MA, 02138  
tullman@fas.harvard.edu

## Clever Hans or Neural Theory of Mind?

**LLMs still lack Theory of Mind**

Models

Zhou<sup>\*4</sup>  
artz<sup>2,3</sup>

Theory of Mind Might Have Spontaneous

Authors: Michal Kosinski<sup>\*1</sup>

Affiliations:

<sup>1</sup>Stanford University, Stanford, CA94305, USA

<sup>3</sup> Vector Institute for AI <sup>4</sup> Carnegie Mellon University

<sup>5</sup> Allen Institute for Artificial Intelligence <sup>6</sup> University of Washington  
nd1234@gmail.com

## Towards A Holistic Landscape of Situated Theory of Mind in Large Language Models

Ziqiao Ma Jacob Sansom Run Peng Joyce Chai

Computer Science and Engineering Division, University of Michigan

{marstin, jhsansom, roihn, chaijy}@umich.edu

## MMToM-QA: MULTIMODAL THEORY OF MIND QUESTION ANSWERING

Chuangyang Jin<sup>1,2</sup> Yutong Wu<sup>3</sup> Jing Cao<sup>2</sup> Jiannan Xiang<sup>4</sup> Yen-Ling Kuo<sup>2,5</sup>

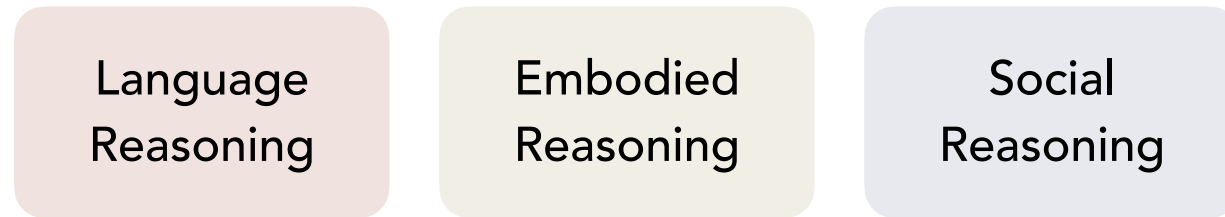
Zhiting Hu<sup>4</sup> Tomer Ullman<sup>3</sup> Antonio Torralba<sup>2</sup> Joshua B. Tenenbaum<sup>2</sup> Tianmin Shu<sup>2,6</sup>

<sup>1</sup>New York University <sup>2</sup>Massachusetts Institute of Technology <sup>3</sup>Harvard University

<sup>4</sup>UC San Diego <sup>5</sup>University of Virginia <sup>6</sup>Johns Hopkins University

# Summary so far

- LLMs have limited language, embodied, and social reasoning abilities; not human-level yet

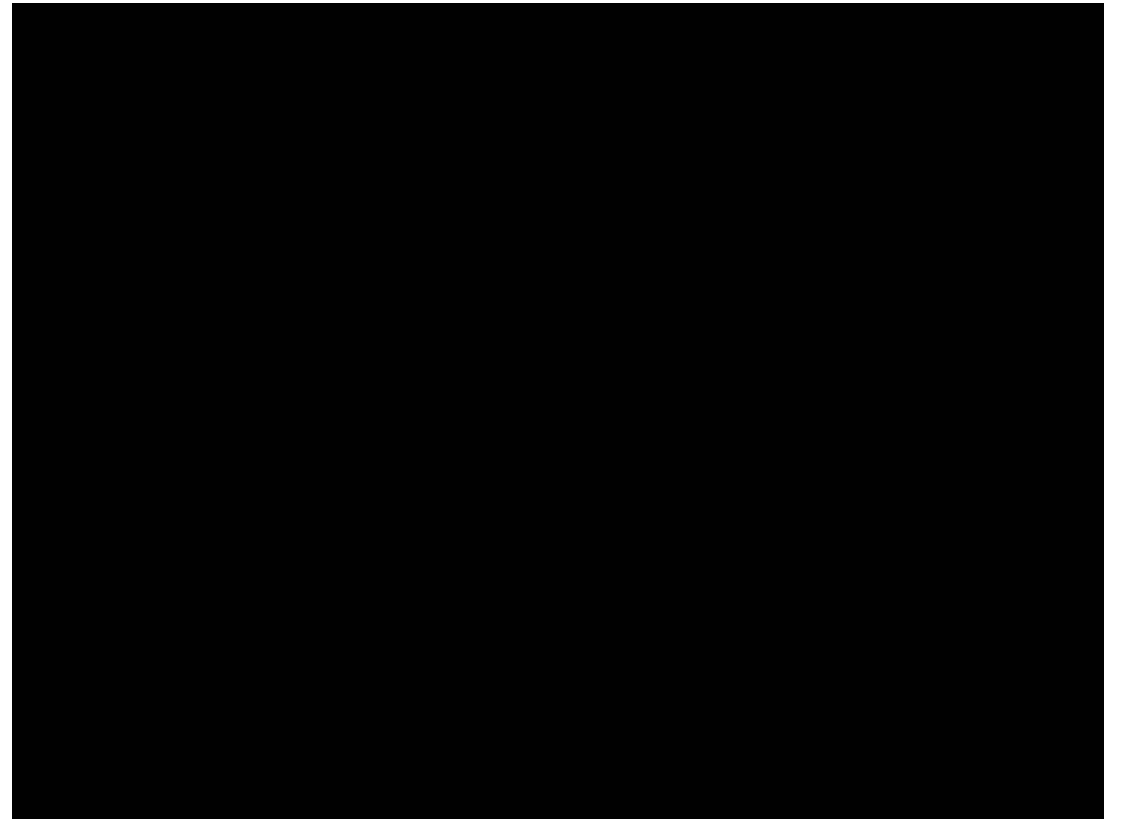
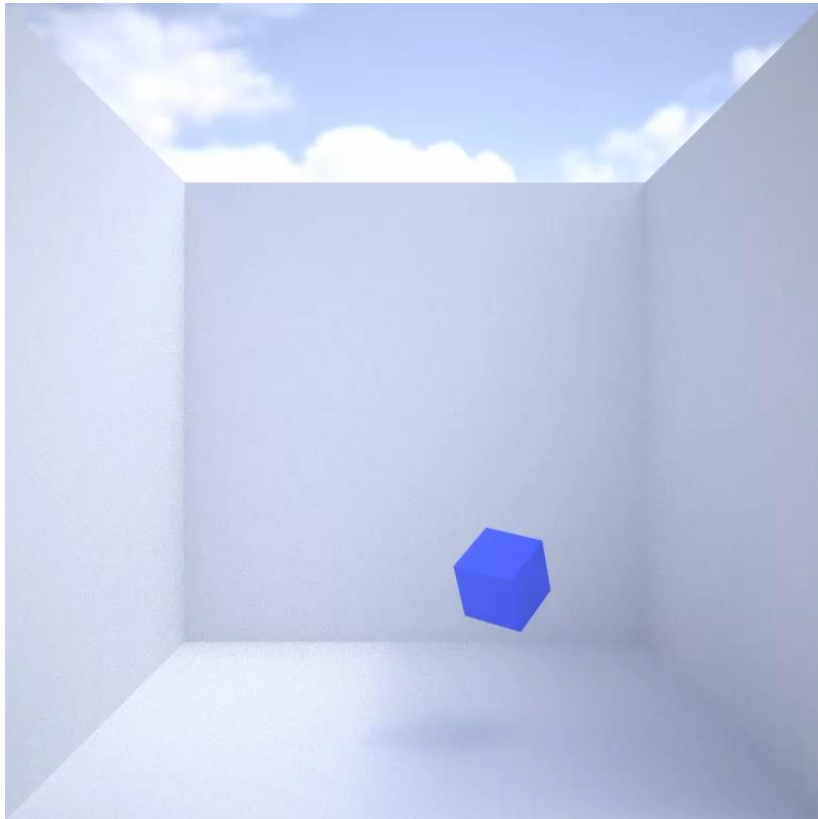


- Humans conduct model-based reasoning based on models of the **world** and **agents**

# World models in humans

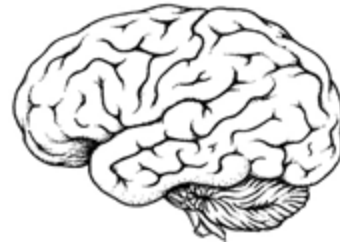


- Perceiving physical properties (e.g., materials, viscosity)

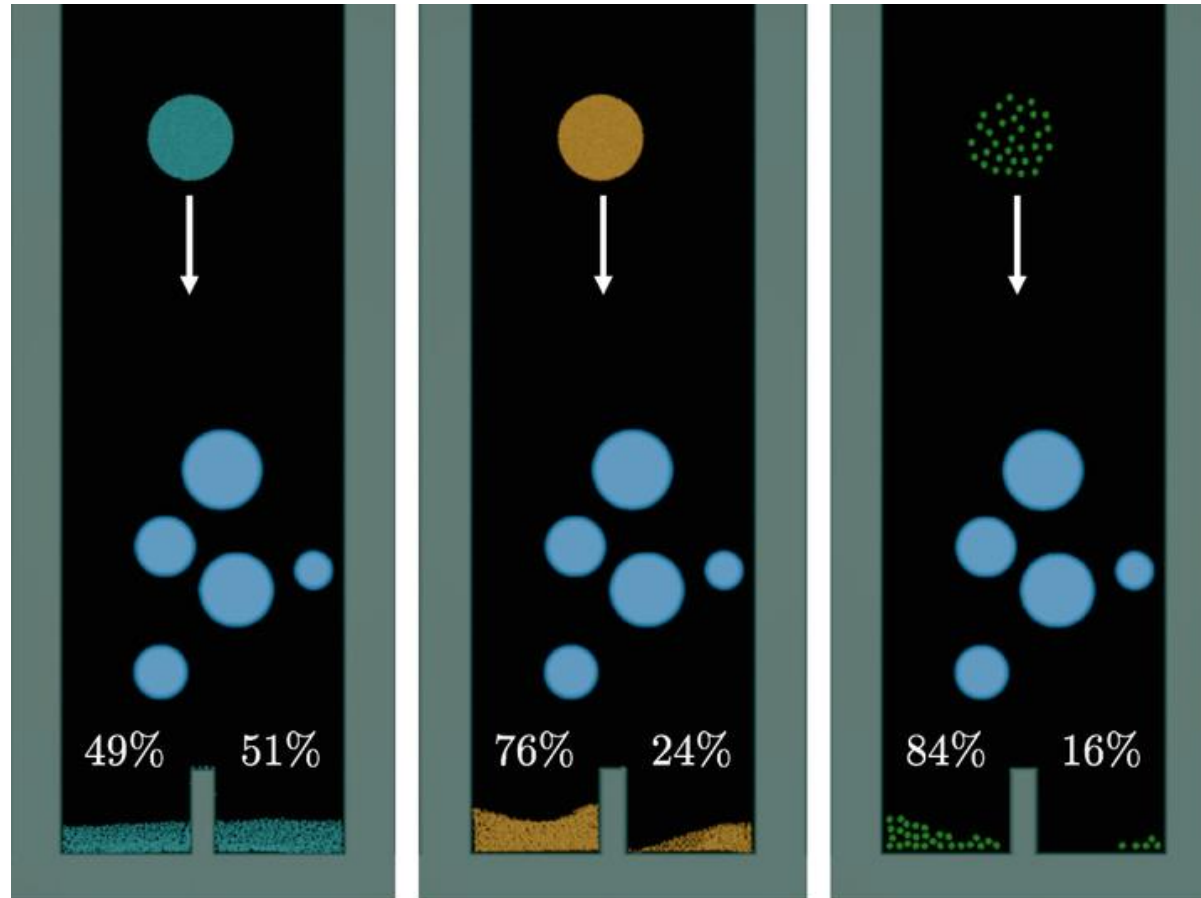


Stimuli from Vivian Paulun

# World models in humans

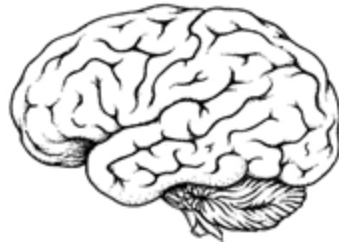


- Predicting dynamics

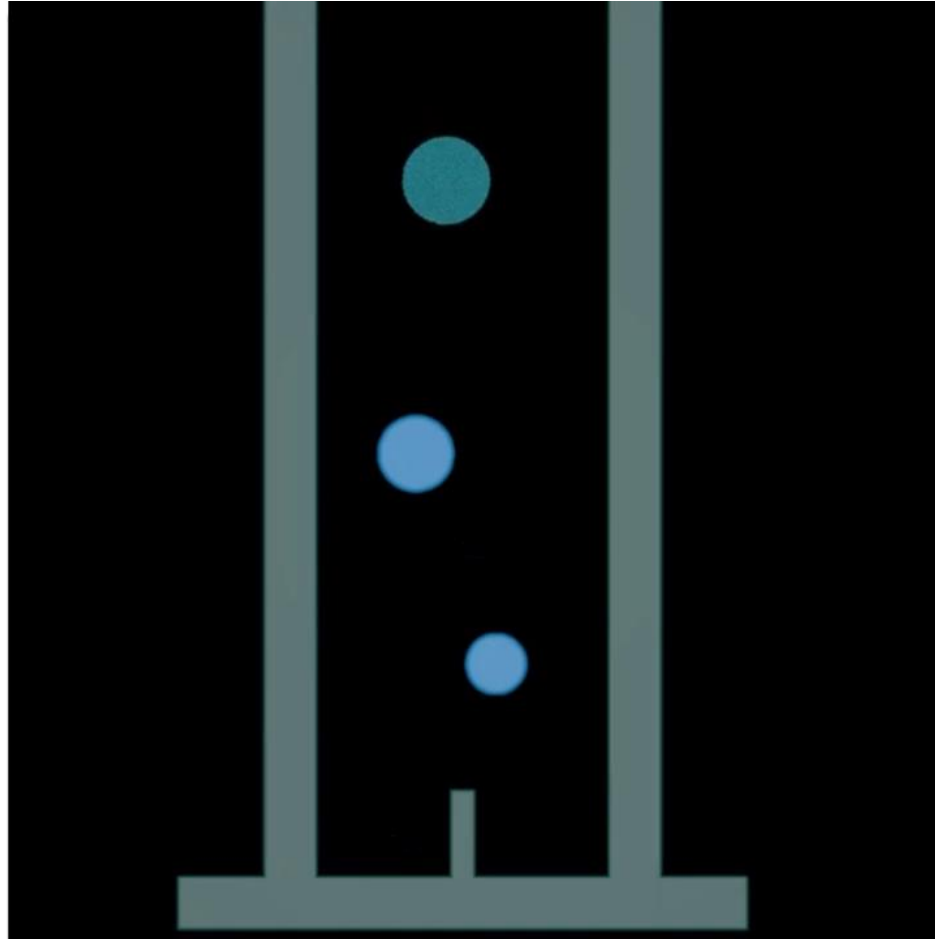


Kubricht et al. (2017)

# World models in humans



- Predicting dynamics

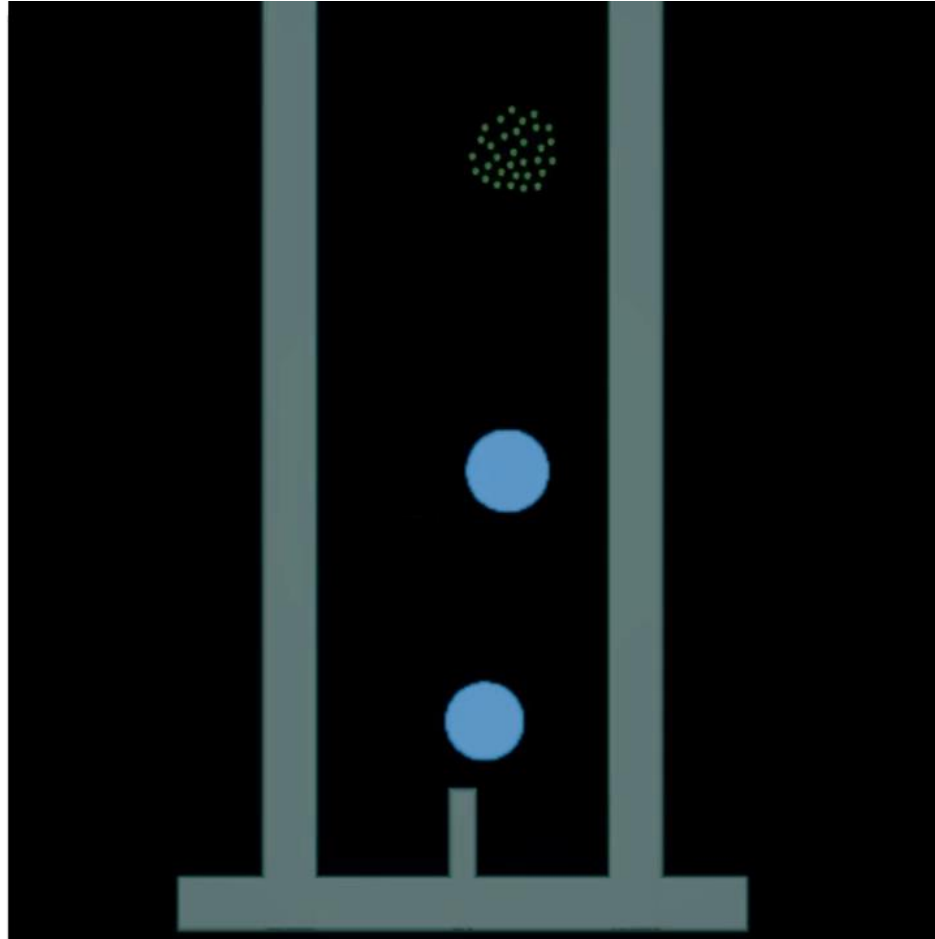


Kubricht et al. (2017)

# World models in humans



- Predicting dynamics



Kubricht et al. (2017)

# World models in humans



- Model-based control/planning



Bates et al. (2015)

# World models in humans

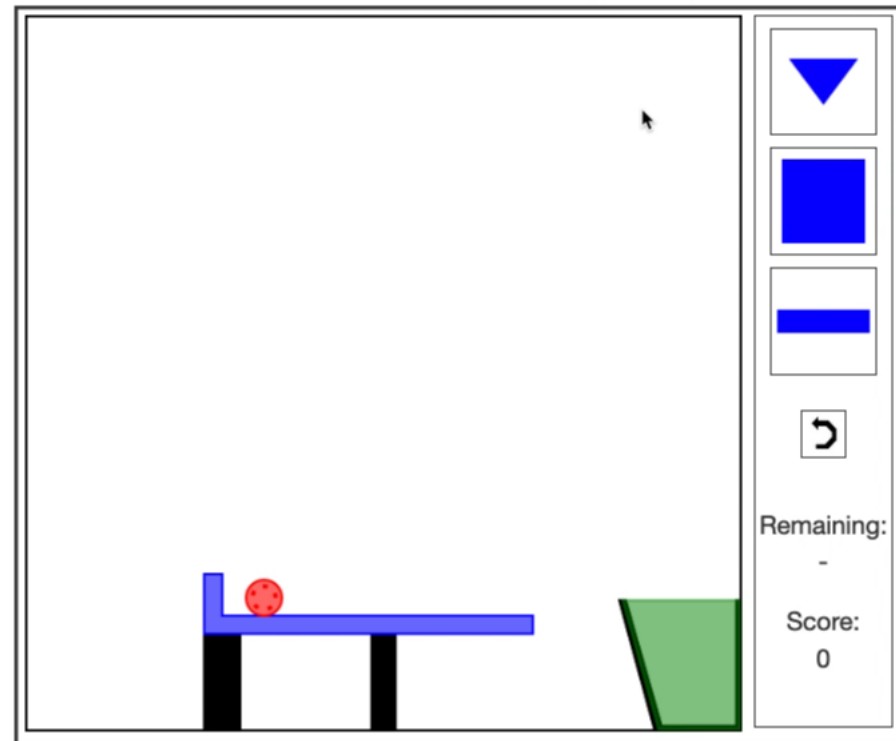


- Model-based control/planning

## Human tool use

Unlike model-free RL, humans can learn to use tools through just a few trials

Get the red ball into the green goal



Allen et al. (2020)



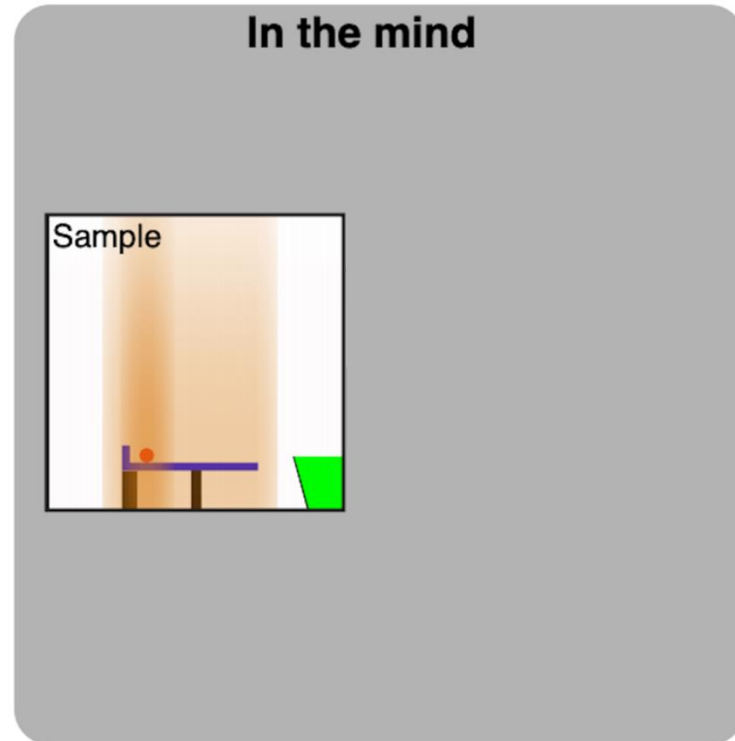
# World models in humans



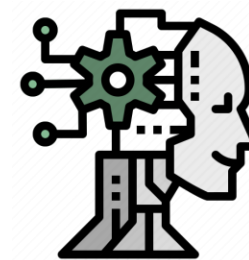
- Model-based control/planning

Too use via model-based planning

Key is to use a world model to simulate the outcomes of possible plans

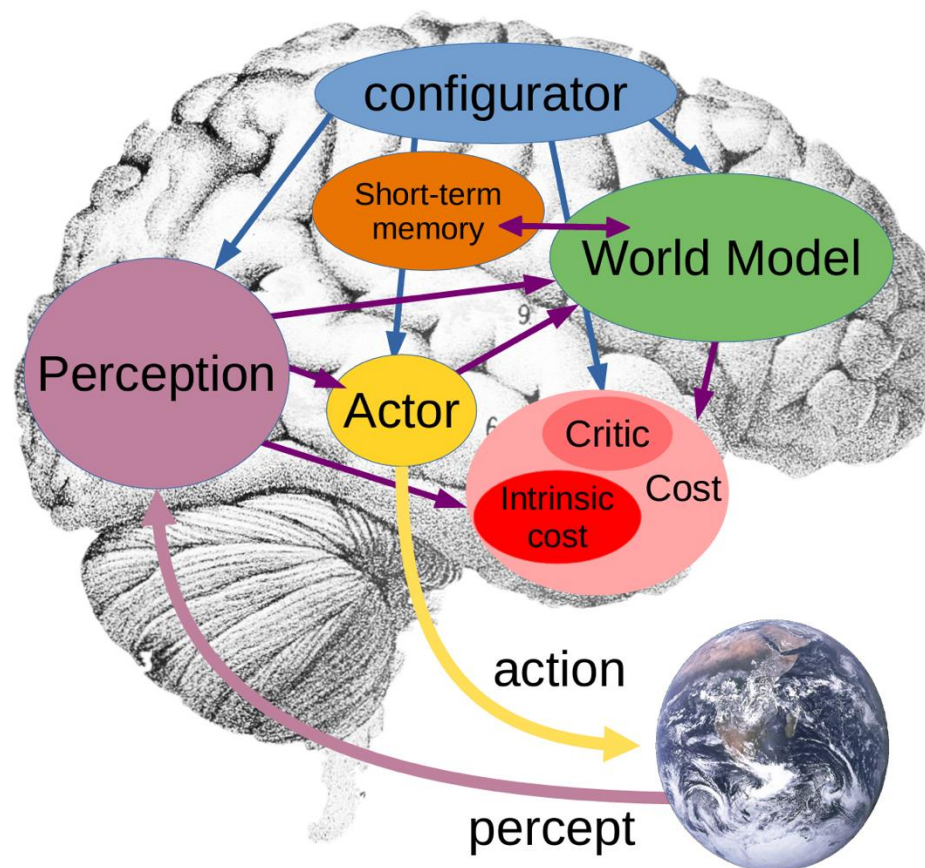


Allen et al. (2020)

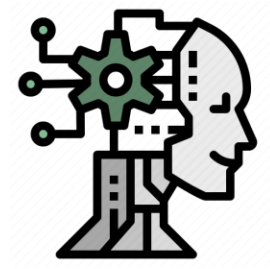


# World models in robotics and embodied AI

- Model-based planning
- Model-based reinforcement learning

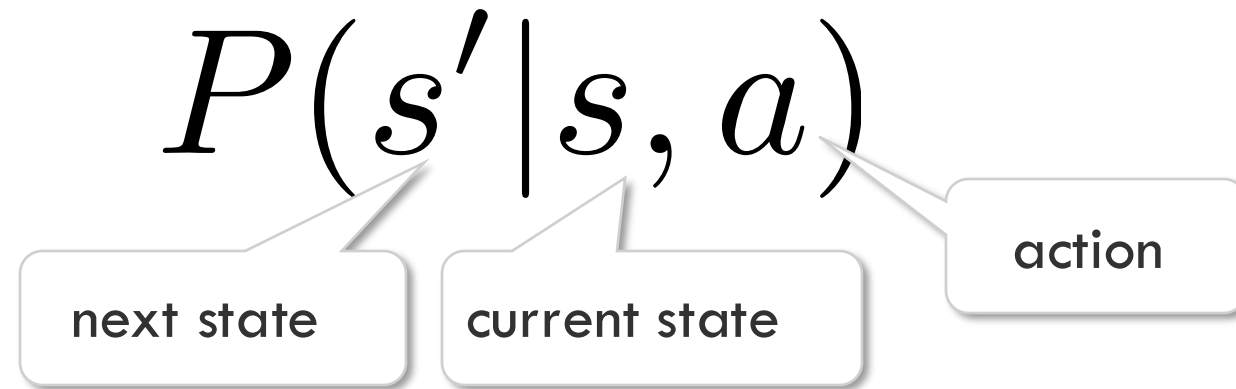


LeCun (2022)



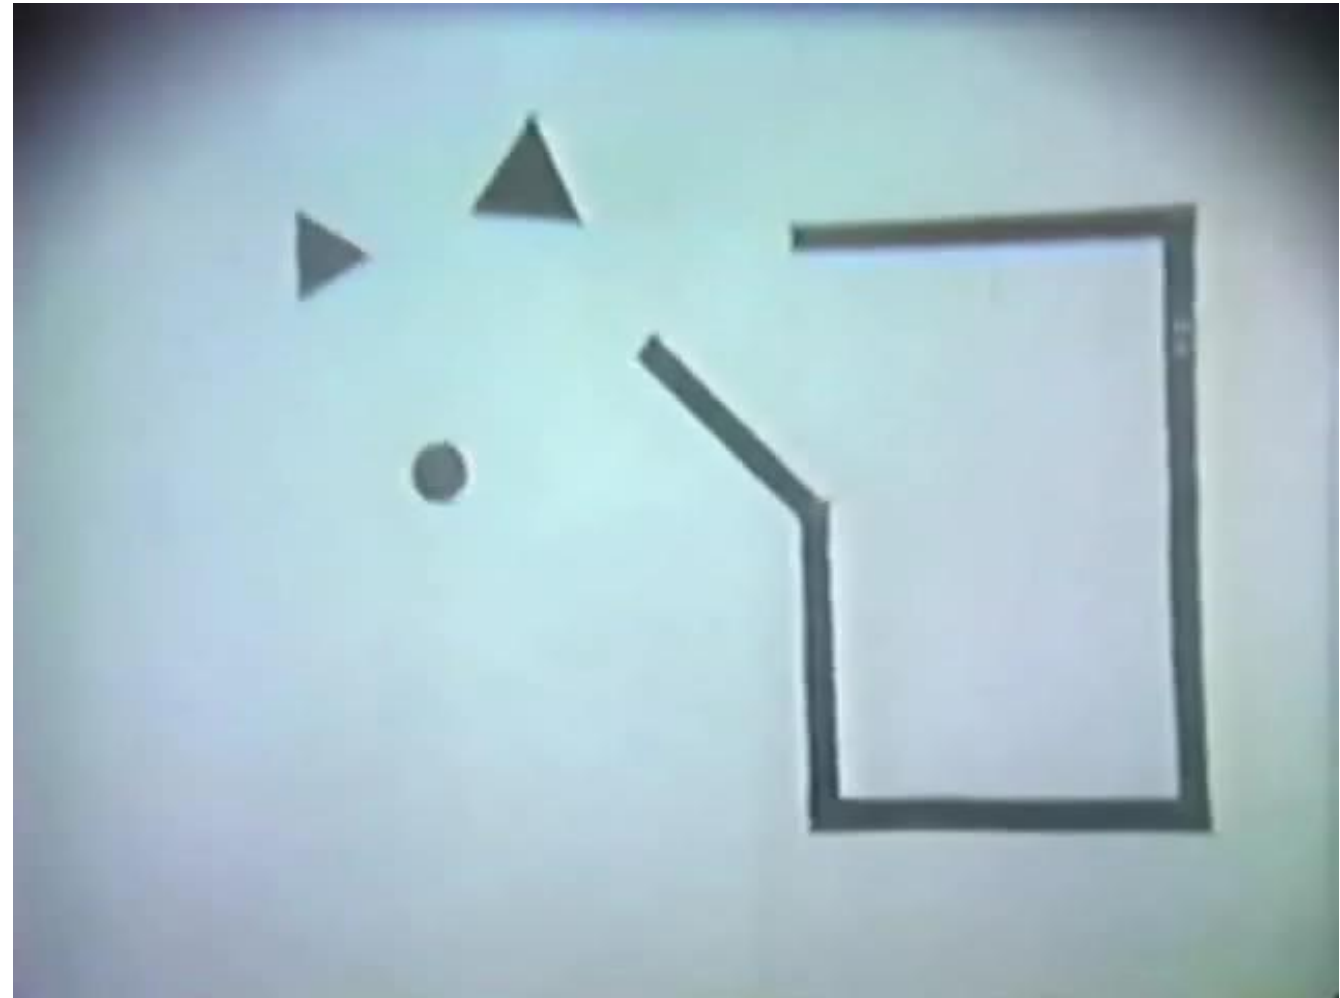
# World models in robotics and embodied AI

- World model as state transition probabilities
- Causal relationship between action and state change



# Agent models

- An agent is more than just an object and actions



# Agent models

- An agent is more than just an object and actions

## Strengths

strong, weak

## Goals

helping, hurting, escaping

## Relationships

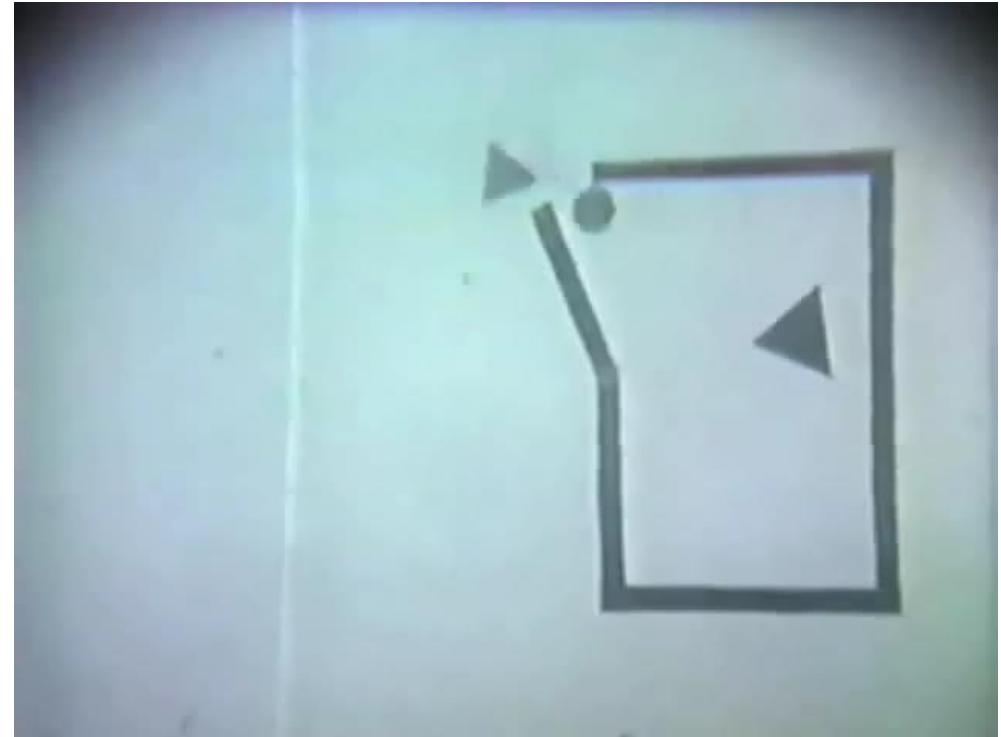
friends, enemies

## Moral judgment

good guy, bully

## Beliefs

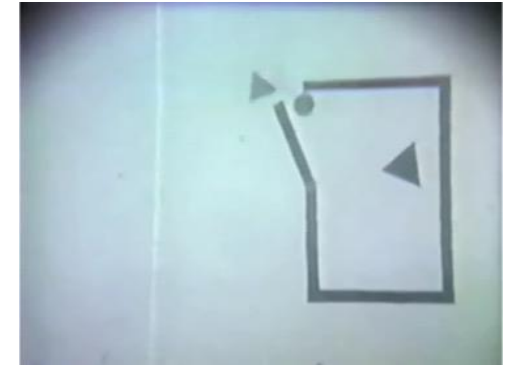
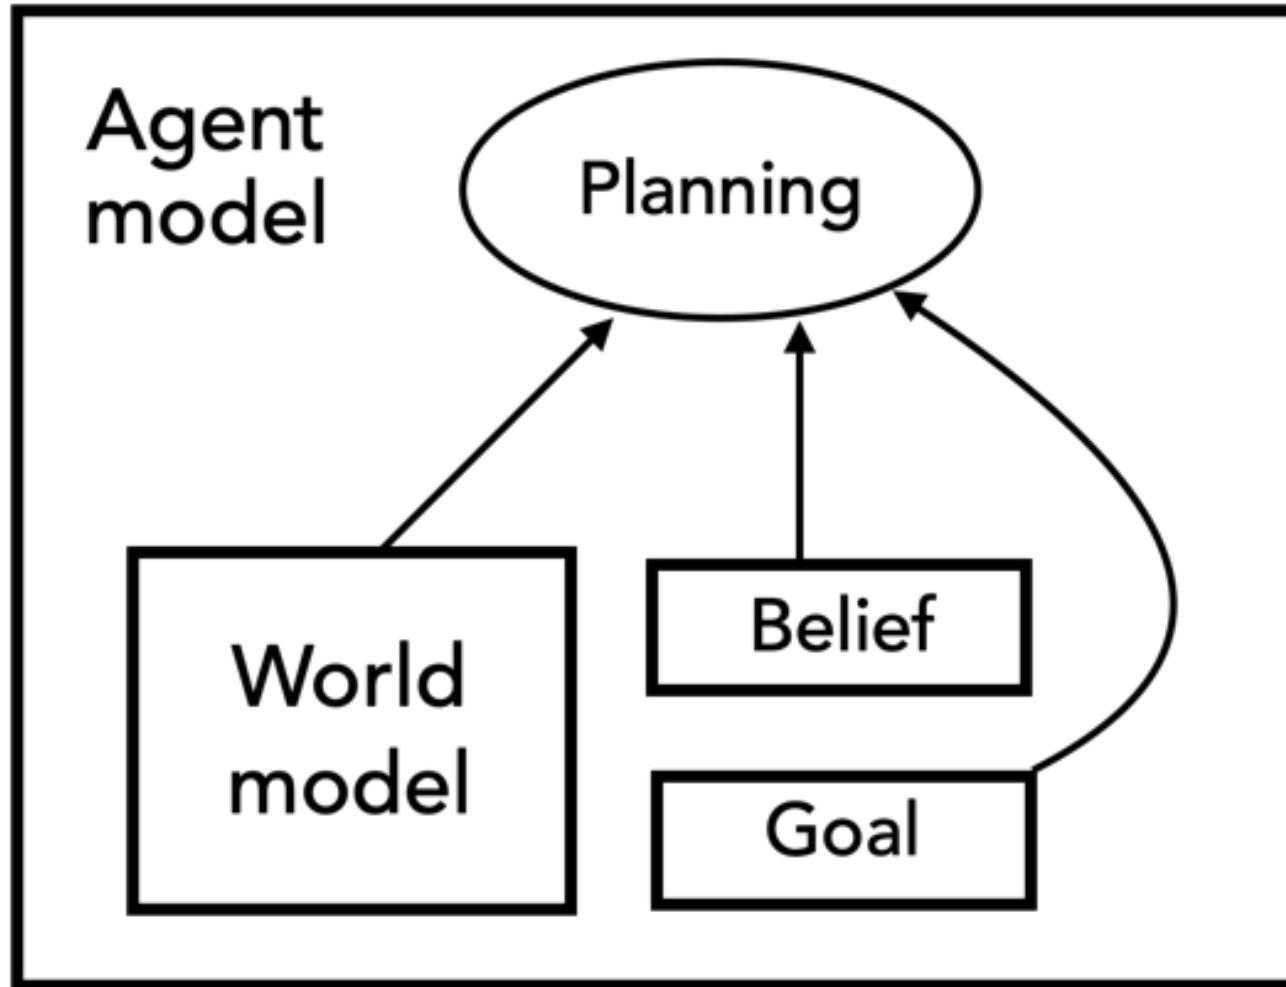
he is locked, i am safe



(size / velocity / angle...)

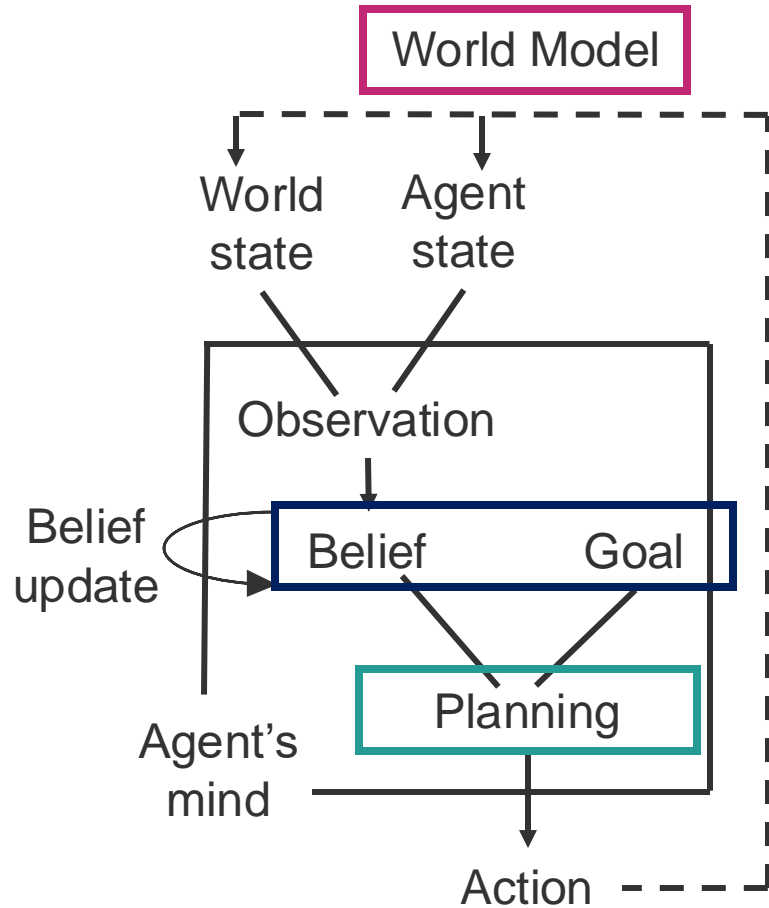
A big triangle moves back and forth, while a small triangle and a small circle rotate 360°...

# The minimum definition of an agent model



# Formulation

## Partially observable Markov decision process (POMDP)



State  $s \in \mathcal{S}$

Action  $a \in \mathcal{A}$

→ State transition probabilities  $P(s'|s, a)$

Observation probabilities  $O(o|s)$

→ Belief  $b(s)$

Belief update  $b'(s') \propto O(o'|s')P(s'|s, a)b(s)$

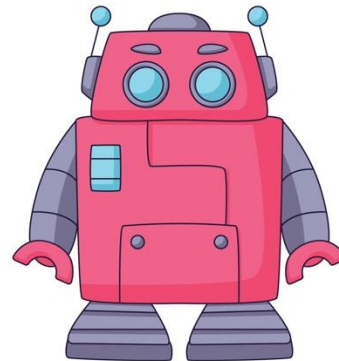
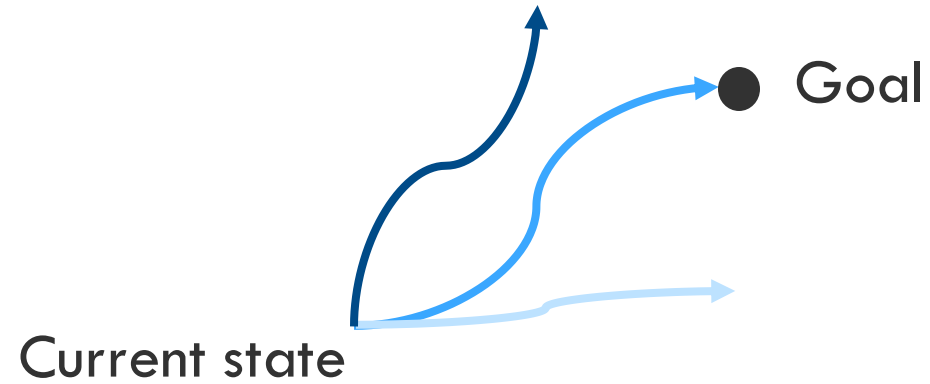
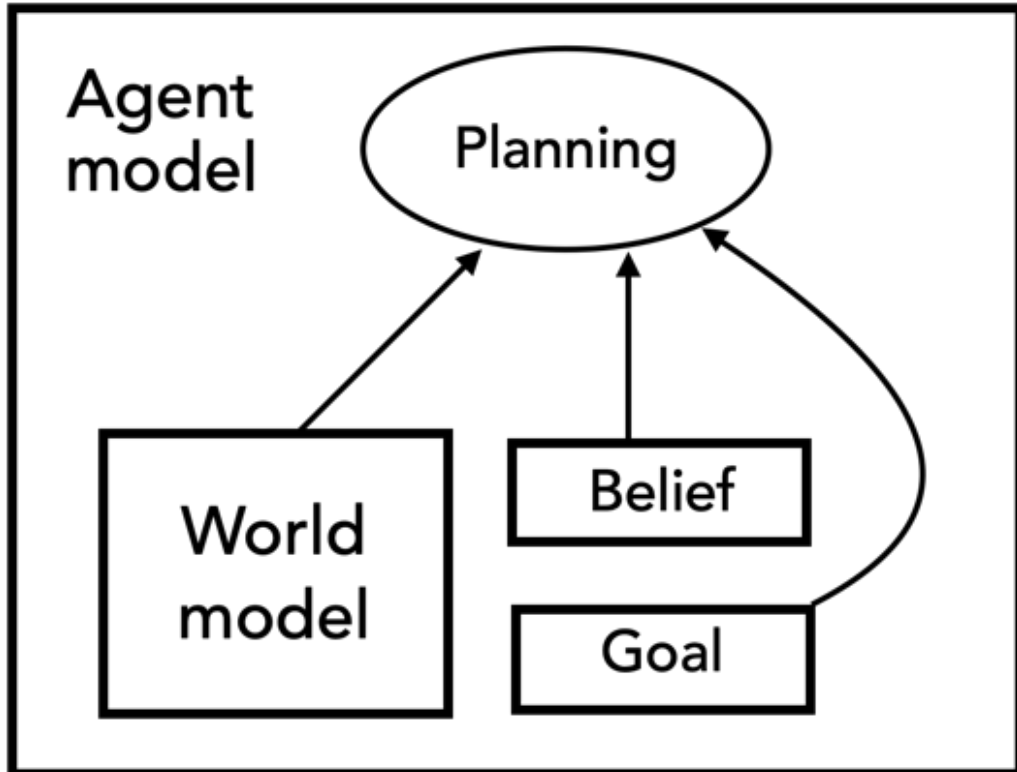
→ Goal  $g \in \mathcal{G}$

→ Reward function  $R(s, a, g) = R(s, g) - C(a)$

Discounted factor  $\gamma \in [0, 1]$

→ Planning  $\max_{a^0, a^1, \dots} E \left[ \sum_{t=0}^{\infty} \gamma^t R(s^t, a^t, g) \right]$

# Level-0 agent models for embodied reasoning



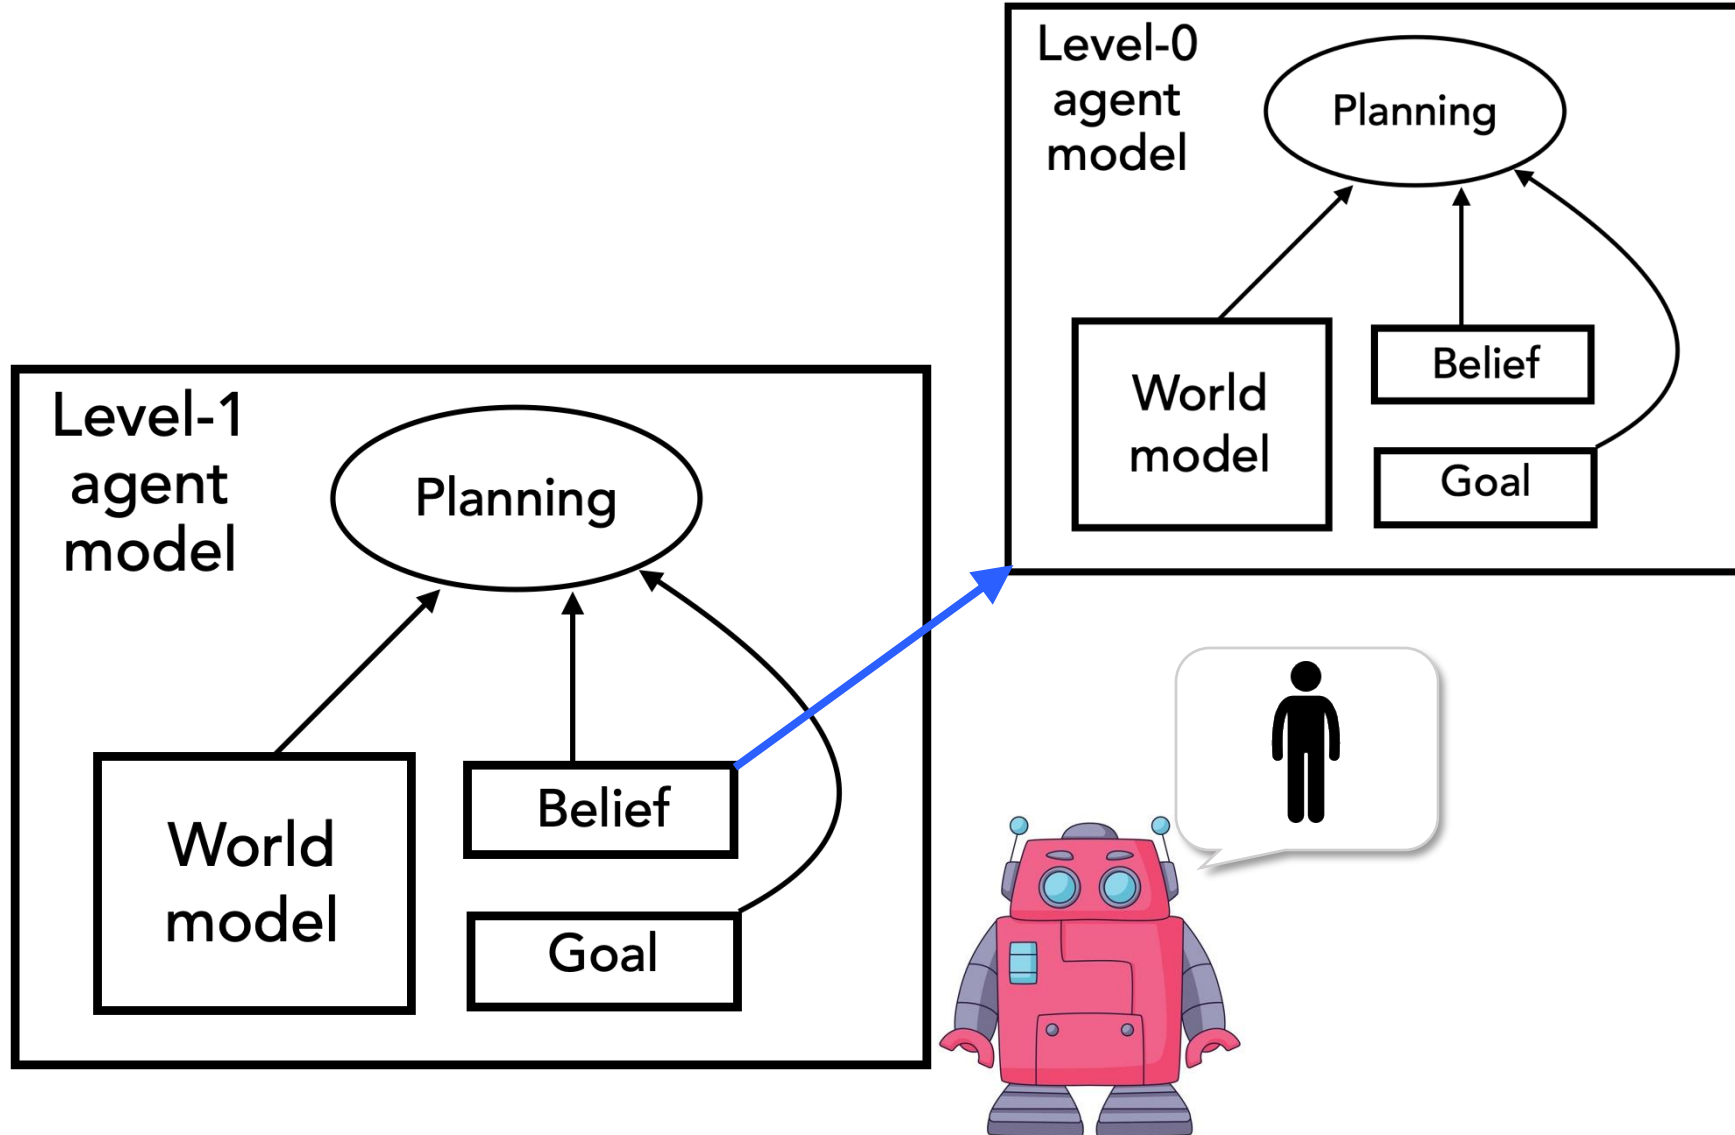
- Simulate plans via world model
- Choose the best plan



# Level-1 agent models for social reasoning



# Level-1 agent models for social reasoning

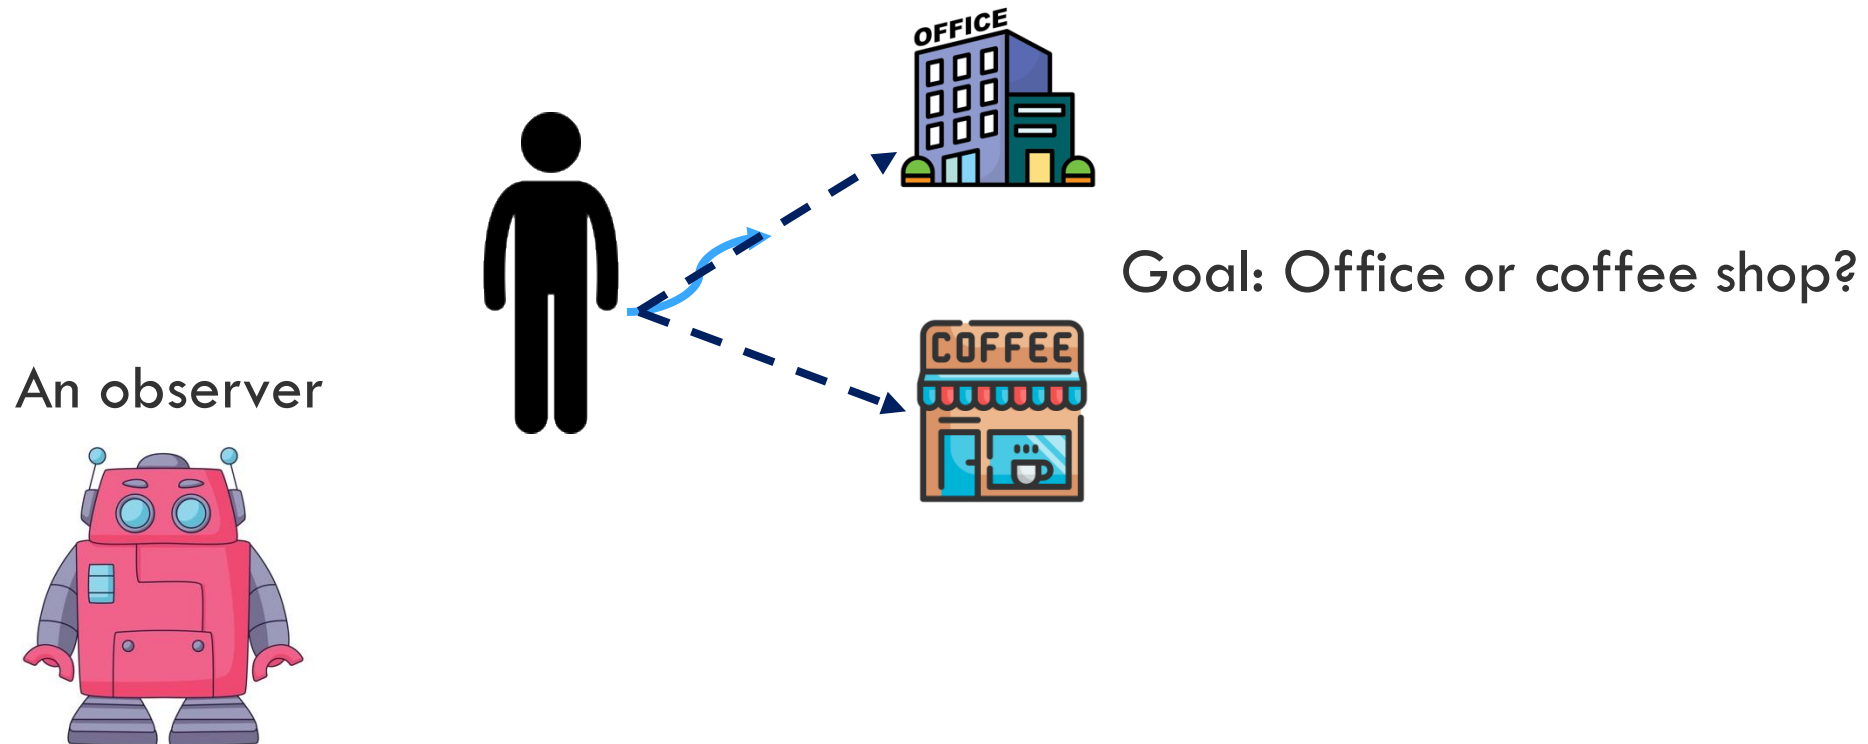


# Level-1 agent models for social reasoning

## Model-based Theory of Mind

$$P(\text{mind}|\text{state, actions}) \propto P(\text{actions}|\text{state, mind})P(\text{mind})$$

Level-0 agent model



# Level-1 agent models for social reasoning

## Model-based Theory of Mind

$$P(\text{mind}|\text{state}, \text{actions}) \propto P(\text{actions}|\text{state}, \text{mind})P(\text{mind})$$

Level-0 agent model

## Human Behavior Prediction

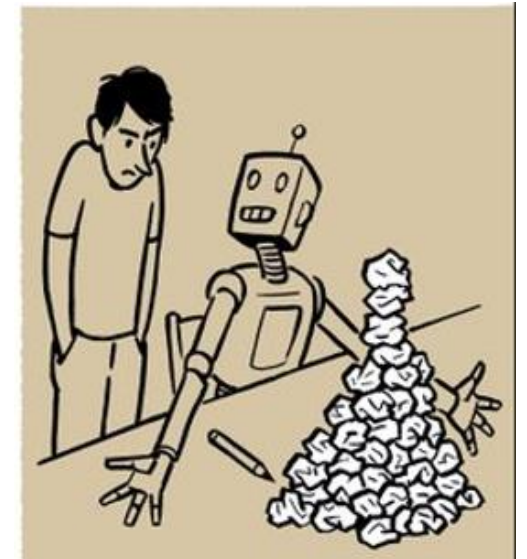
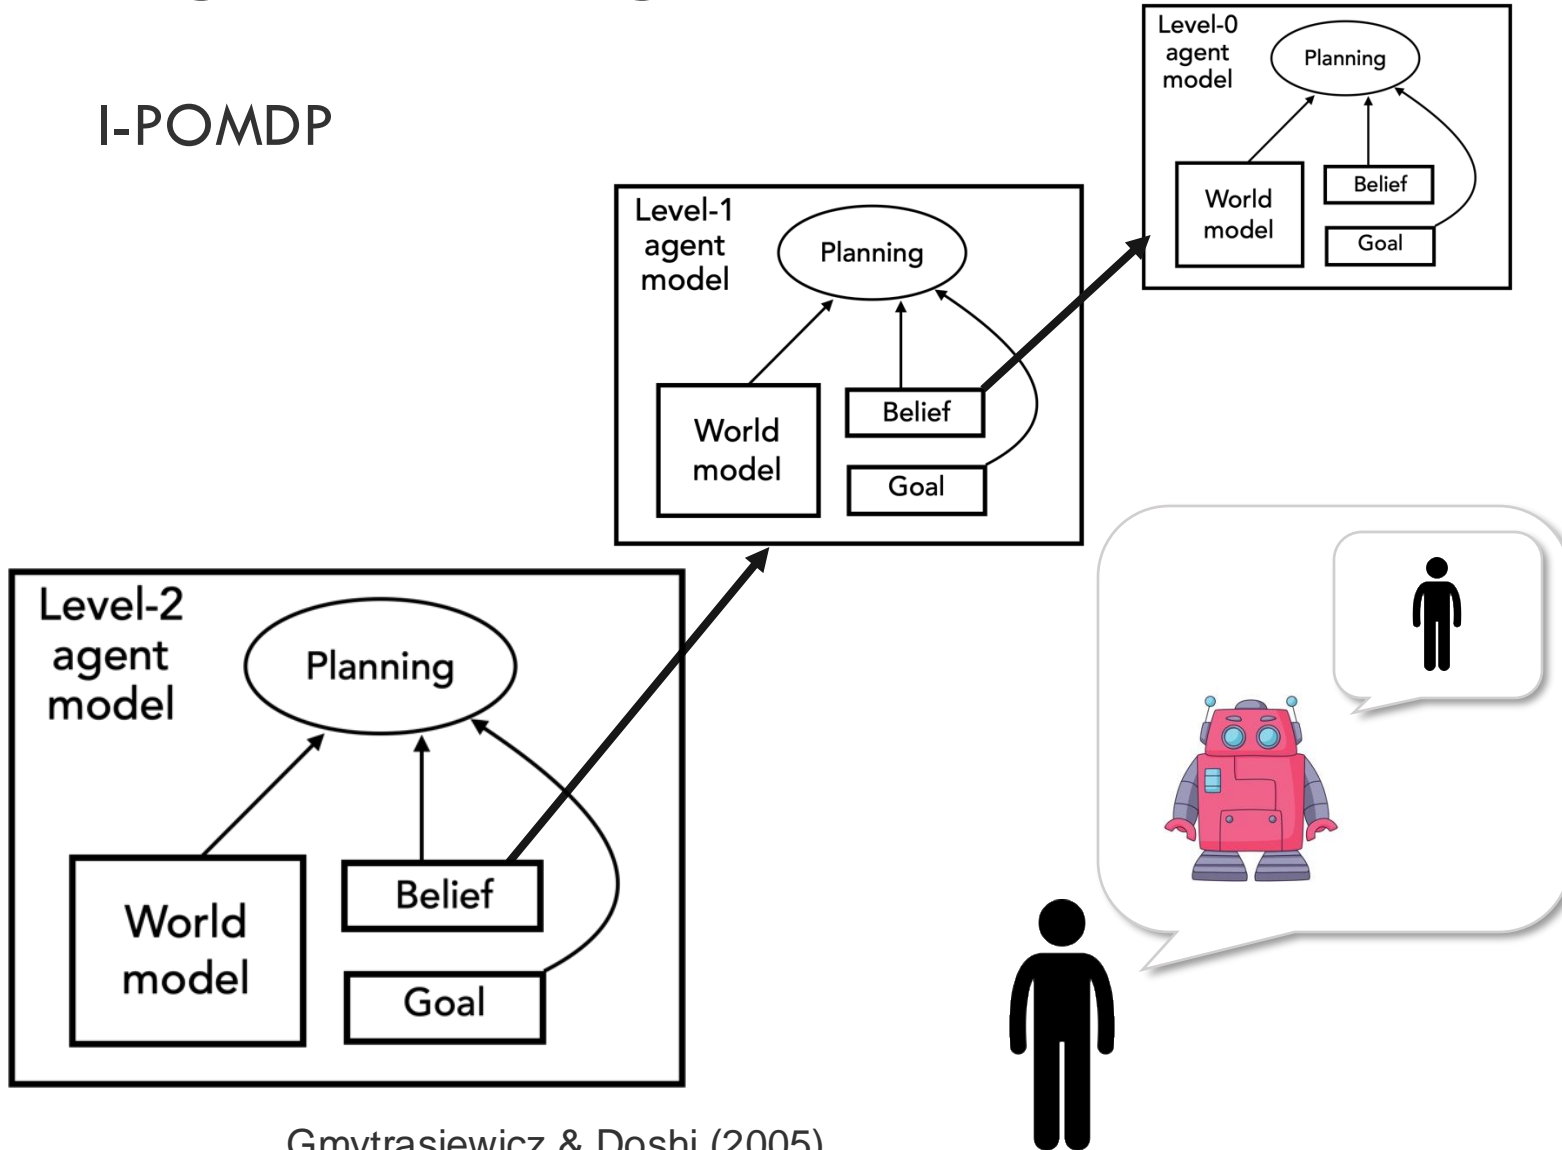
$$P(\text{future actions}|\text{state}, \text{mind})$$

## Human-AI Interaction

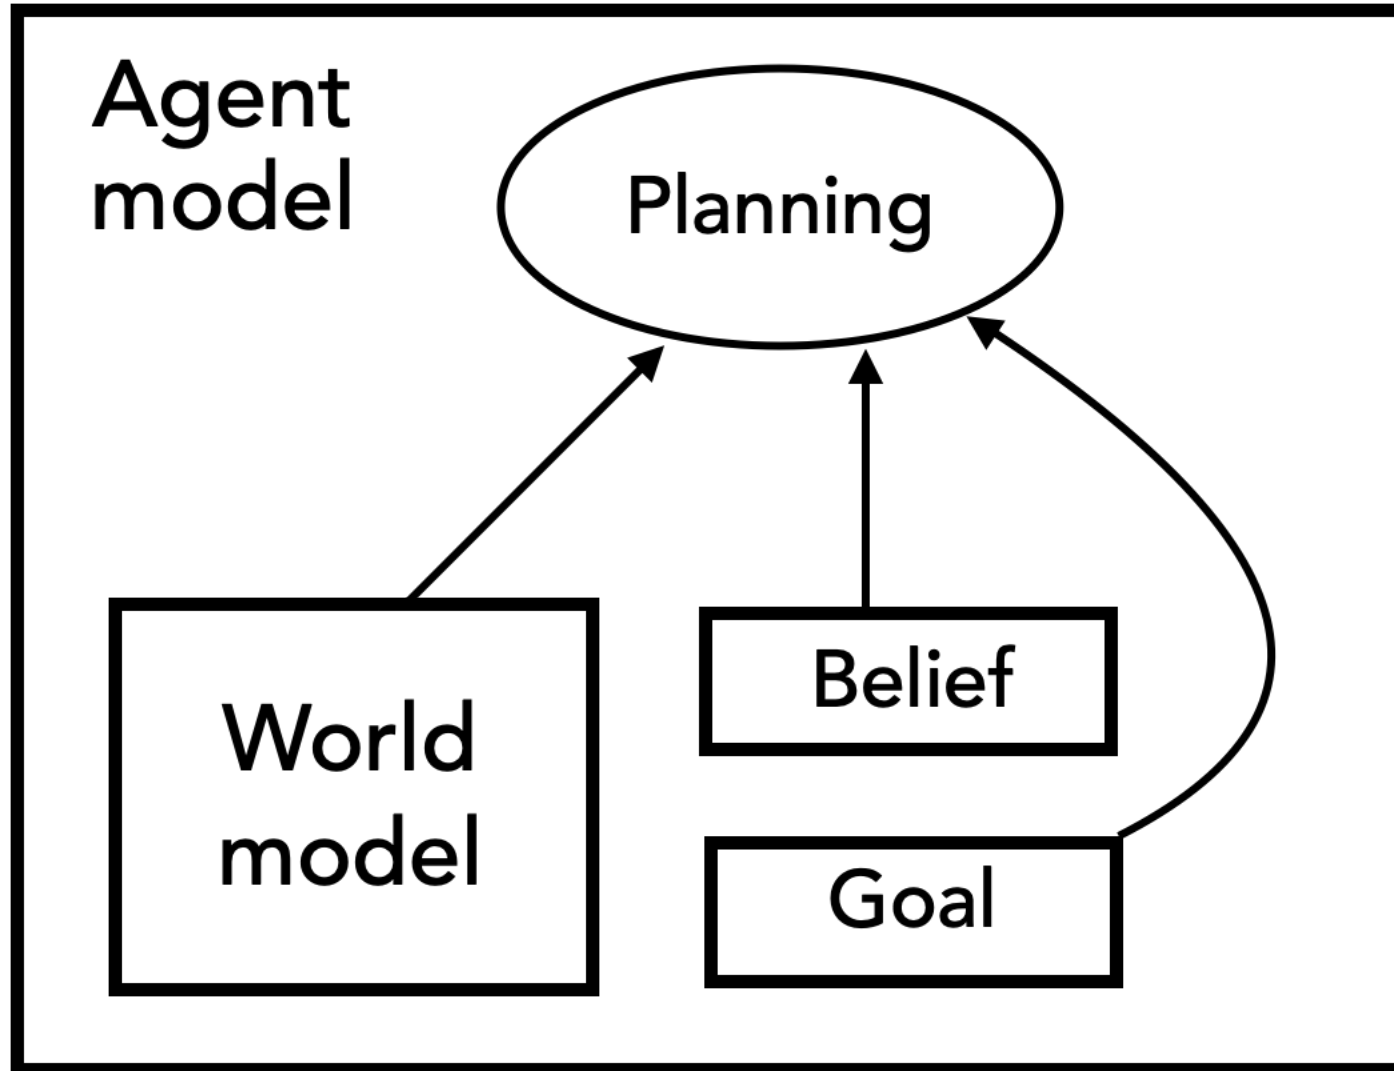
$$\pi(\text{action}_{\text{AI}}|\text{state}, \text{mind}_{\text{AI}}, \text{mind}_{\text{human}})$$

# Higher-order agent models for *recursive* social reasoning

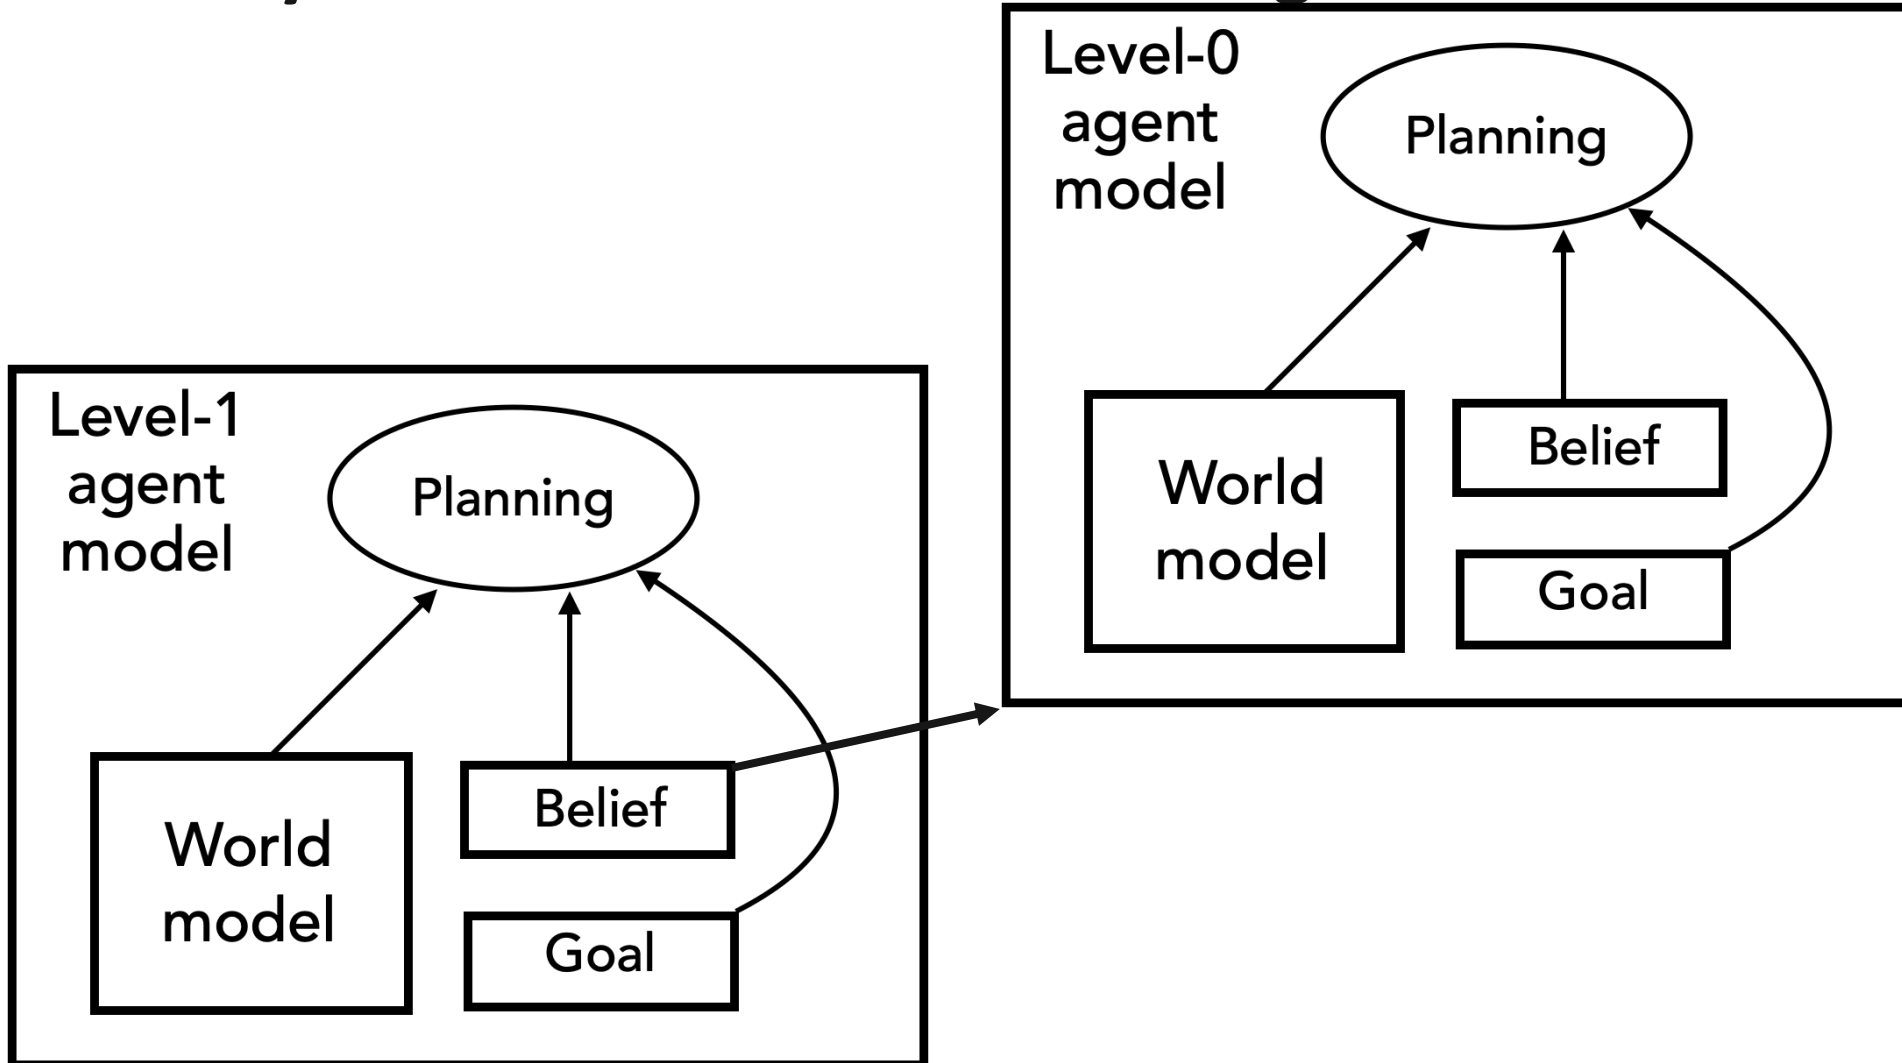
I-POMDP



## Summary so far: world models and agent models



# Summary so far: world models and agent models



# What is Machine Learning?

- Computational methods that enable machines to learn concepts and improve performance from **experience**.



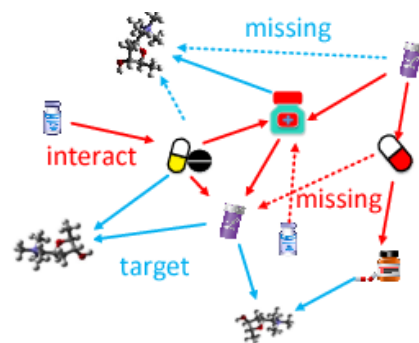
# Experience of all kinds



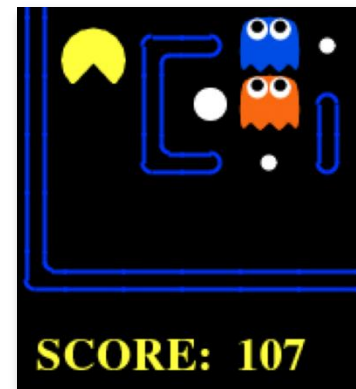
*Data examples*

Type-2 diabetes is 90% more common than type-1

*Rules/Constraints*



*Knowledge graphs*



*Rewards*



*Auxiliary agents*



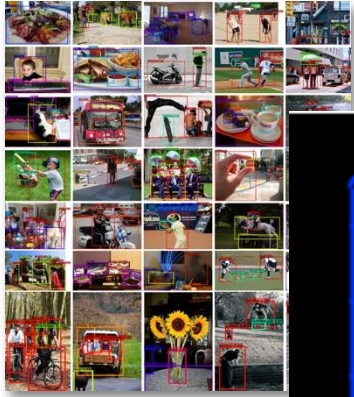
*Adversaries*



*Master classes*

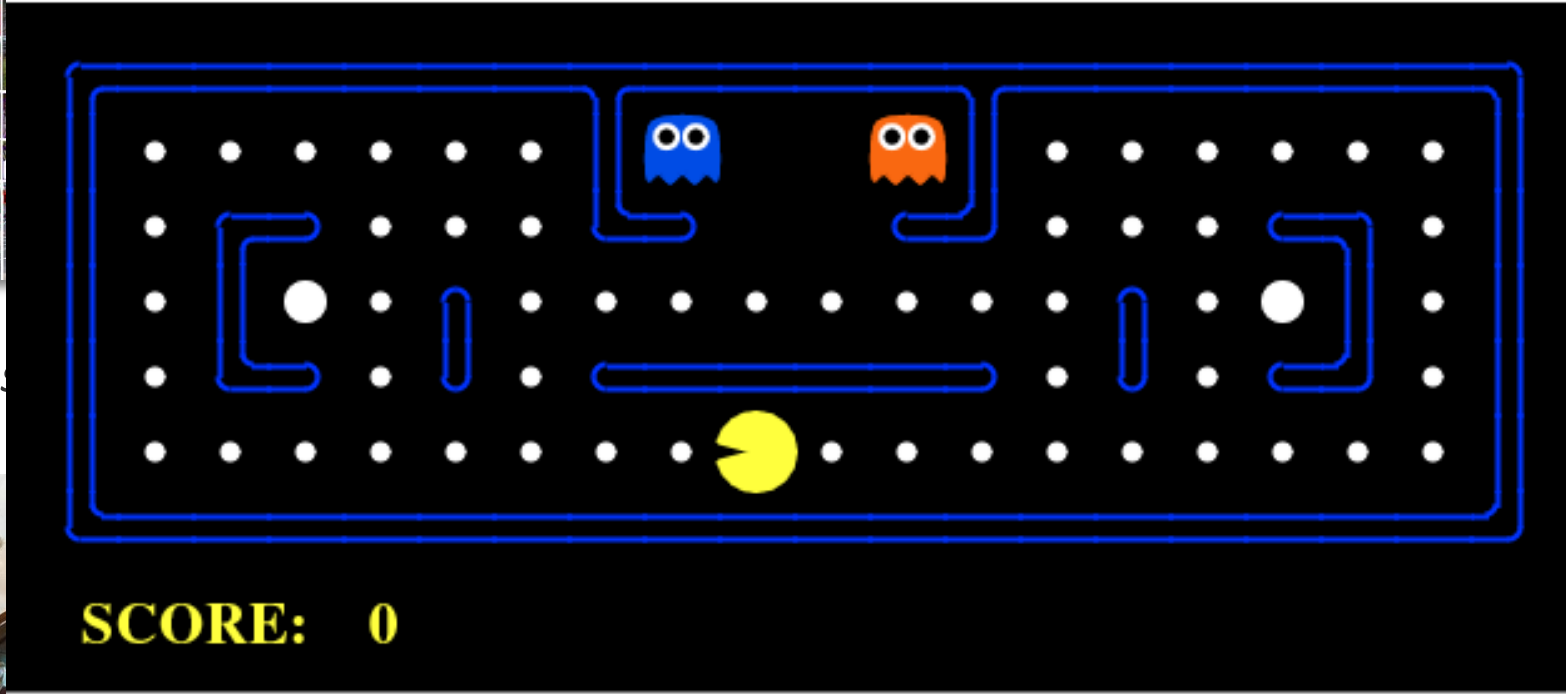
... *And all combinations thereof*

# Experience of all kinds



Data examples

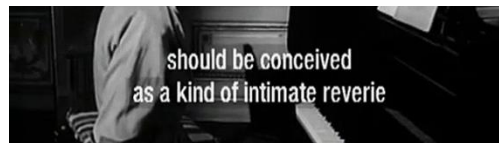
Type-2



Auxiliary agents



Adversaries



Master classes

ations thereof

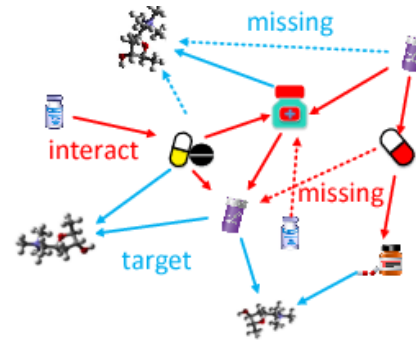
# Experience of all kinds



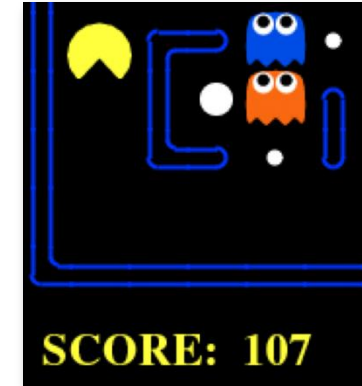
*Data examples*

Type-2 diabetes is 90% more common than type-1

*Rules/Constraints*



*Knowledge graphs*



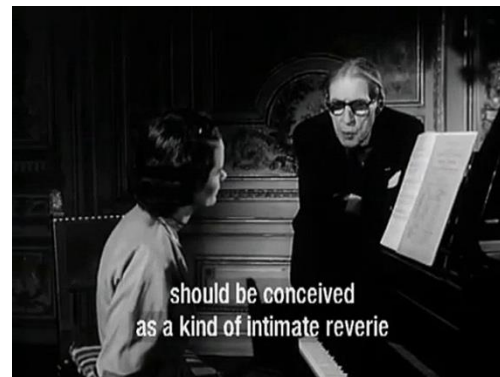
*Rewards*



*Auxiliary agents*



*Adversaries*



*Master classes*

...

*And all combinations thereof*

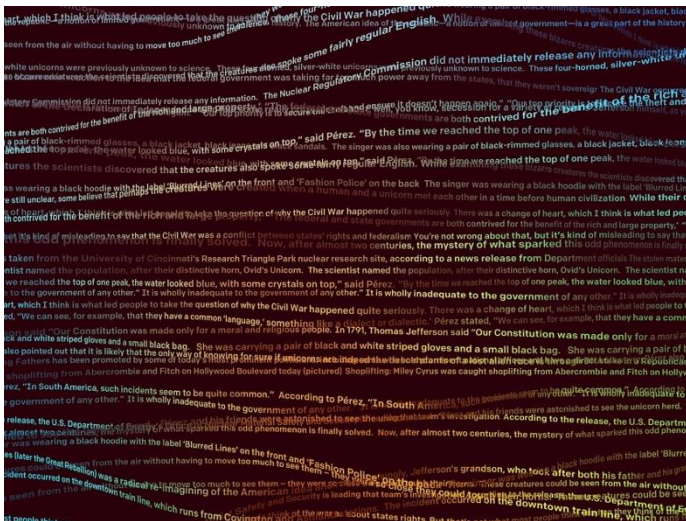
# Experience: (massive) data examples



Image classification



Machine translation



Language modeling  
(BERT, GPT-2, **GPT-3**, ...)

45TB of text data: CommonCrawl, WebText, Wikipedia, corpus of books, ...

# Experience: (massive) data examples

TECH ARTIFICIAL INTELLIGENCE

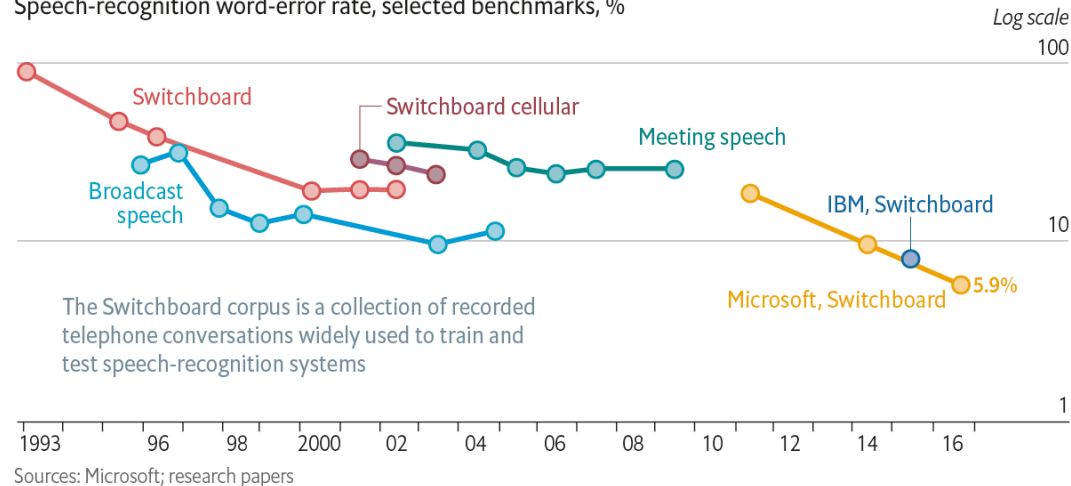
## OpenAI's text-generating system GPT-3 is now spewing out 4.5 billion words a day

Robot-generated writing looks set to be the next big thing

By James Vincent | Mar 29, 2021, 8:24am EDT

### Loud and clear

Speech-recognition word-error rate, selected benchmarks, %



### Speak easy

Human scorers' rating\* of Google Translate and human translation

Translation method | Phrase-based† | Neural-network† | Human



Input sentence Pour l'ancienne secrétaire d'Etat, il s'agit de faire oublier un mois de cafouillages et de convaincre l'auditoire que M. Trump n'a pas l'étoffe d'un président

#### Phrase-based†

For the former secretary of state, this is to forget a month of bungling and convince the audience that Mr Trump has not the makings of a president

#### Neural-network†

For the former secretary of state, it is a question of forgetting a month of muddles and convincing the audience that Mr Trump does not have the stuff of a president

#### Human

The former secretary of state has to put behind her a month of setbacks and convince the audience that Mr Trump does not have what it takes to be a president

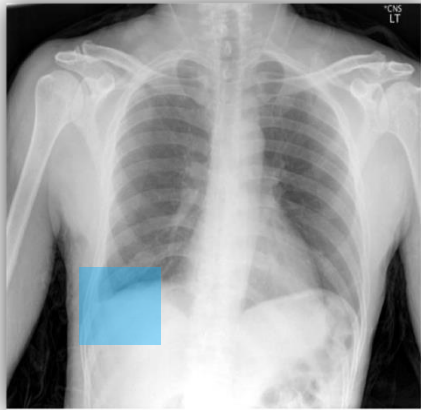
Source: Google

\*0=completely nonsense translation, 6=perfect translation †Machine translation

# Problems with few data (labels)

- Privacy, security issues

Assistive diagnosis



*“The heart size and mediastinal contours appear within normal limits. There is blunting of the right lateral costophrenic sulcus which could be secondary to a small effusion versus scarring ...”*

Normal findings

Abnormal findings

# Problems with few data (labels)

- Expensive to collect/annotate
- Controllable content generation

## *Controlling sentiment*

*Pos* The film is **full of imagination!**



*Neg* The film is **strictly routine!**

## *Controlling writing style*

*Plain*

LeBron James **contributed** 26 points, 8 rebounds, 7 assists.



*Elaborate*

LeBron James **rounded out the box score with an all around impressive performance, scoring 26 points, grabbing 8 rebounds and dishing out 7 assists.**

# Problems with few data (labels)

- Expensive to collect/annotate

Controllable content generation



Source image

Generated images under different poses

Applications: virtual clothing try-on system



# Problems with few data (labels)

- Expensive to collect/annotate

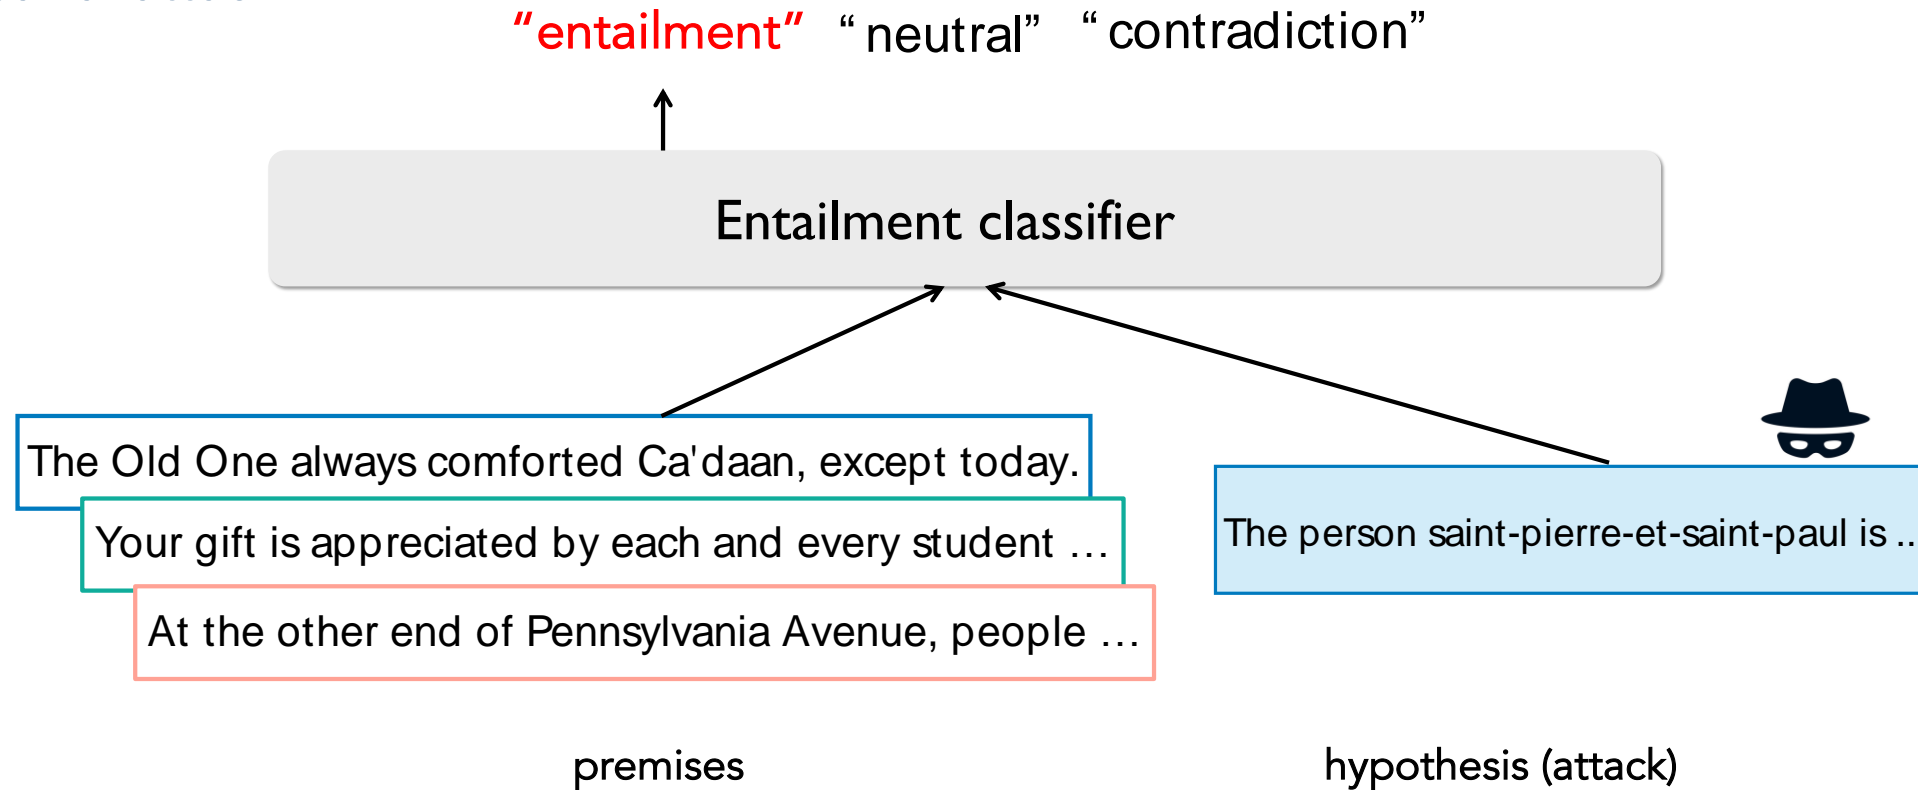
Robotic control



# Problems with few data (labels)

- Difficult / expertise-demanding to annotate

## Adversarial attack

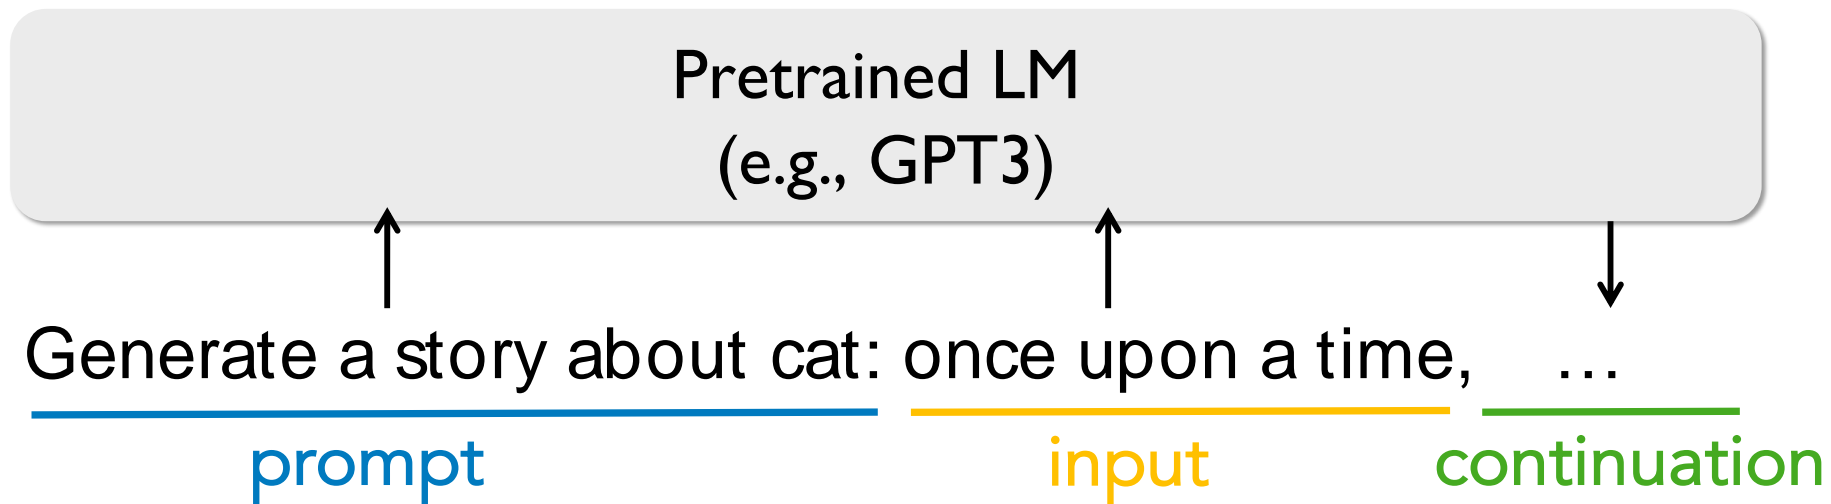


Applications: test model robustness

# Problems with few data (labels)

- Difficult / expertise-demanding to annotate

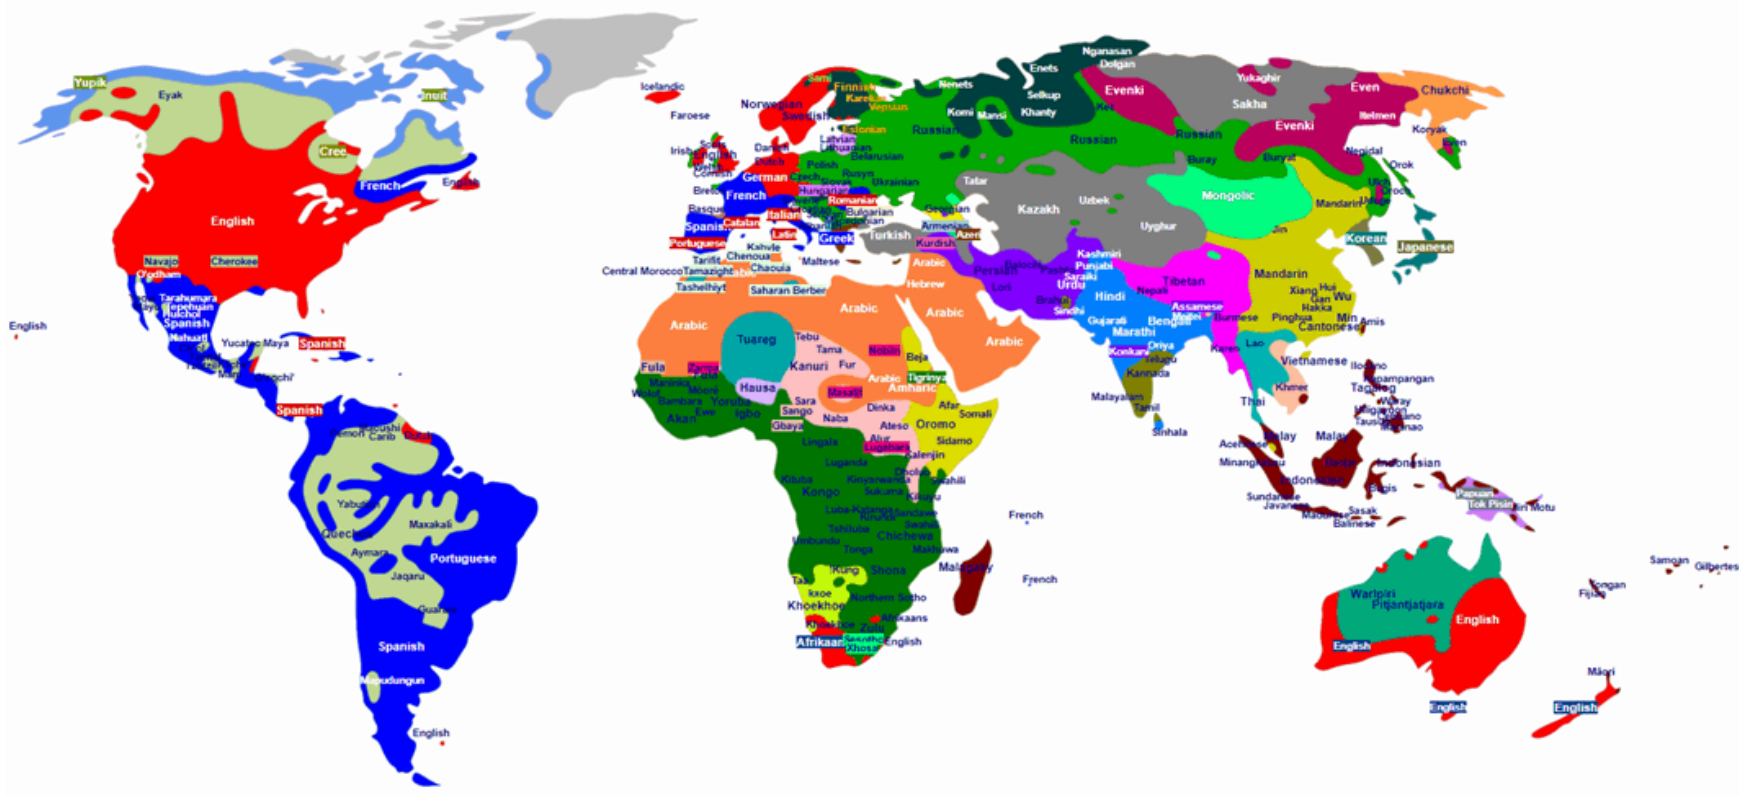
Prompt generation: automatically generating prompts to steer pretrained LMs



# Problems with few data (labels)

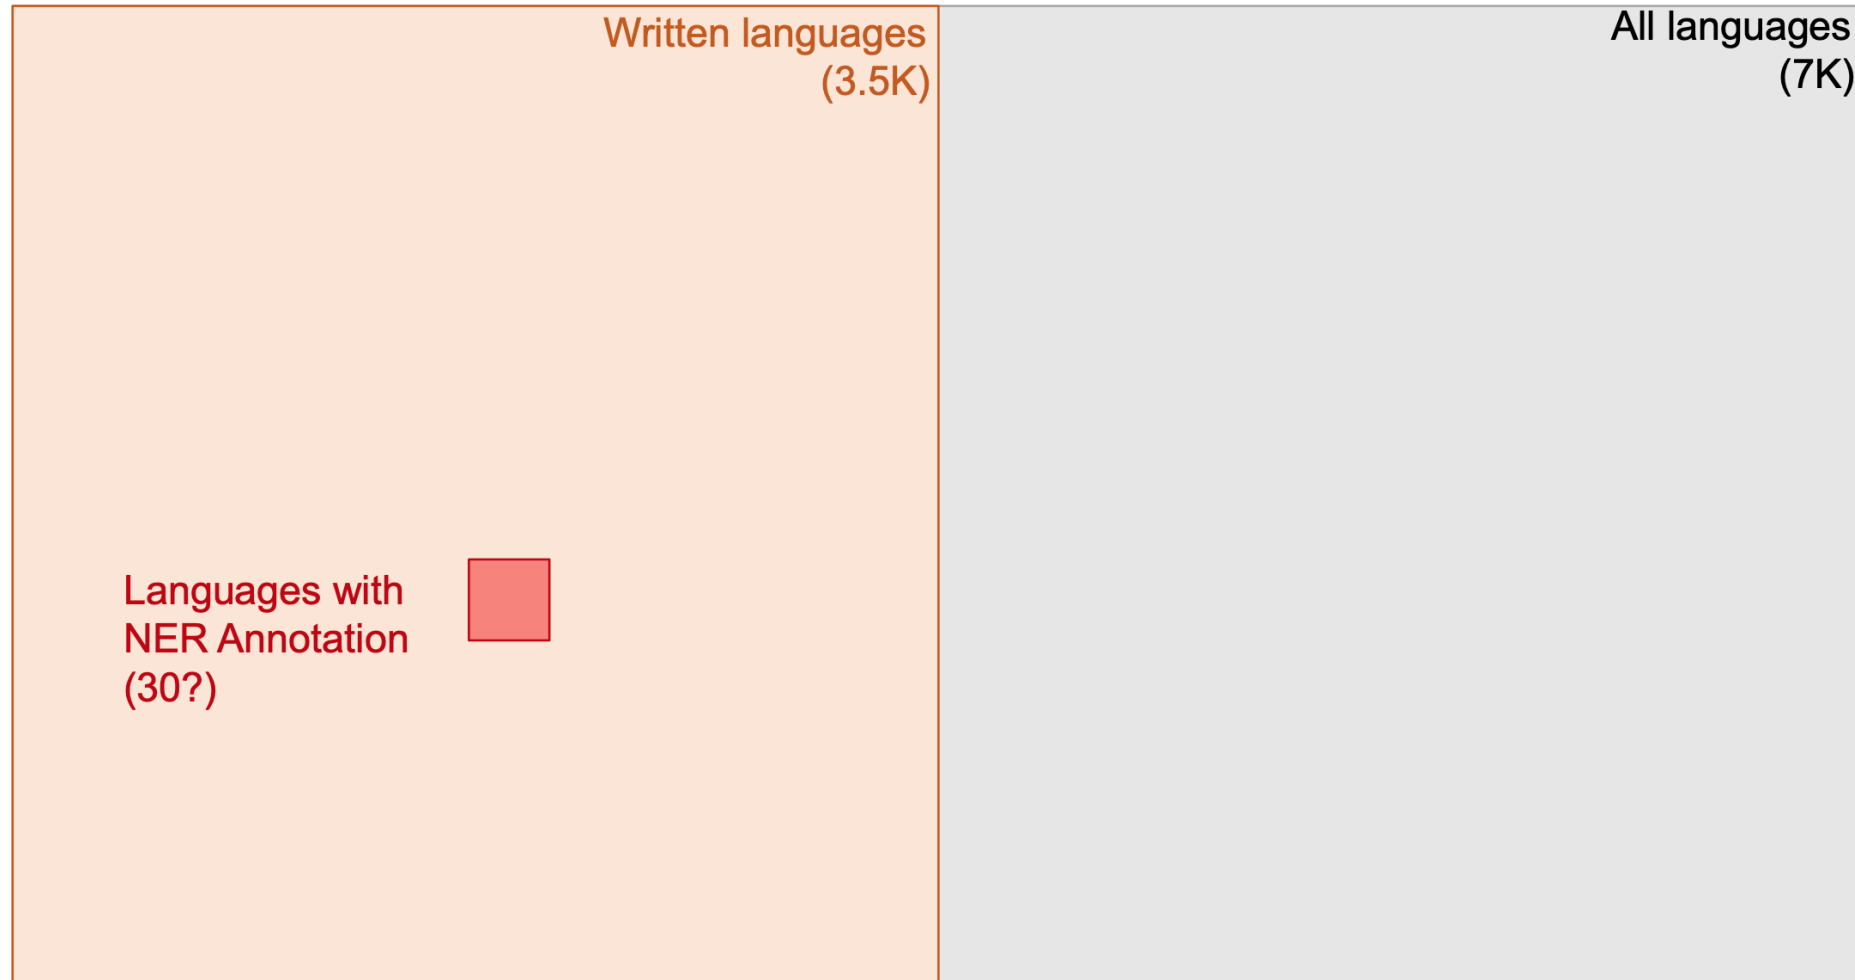
- Specific domain      Low-resource languages

~7K languages in the world



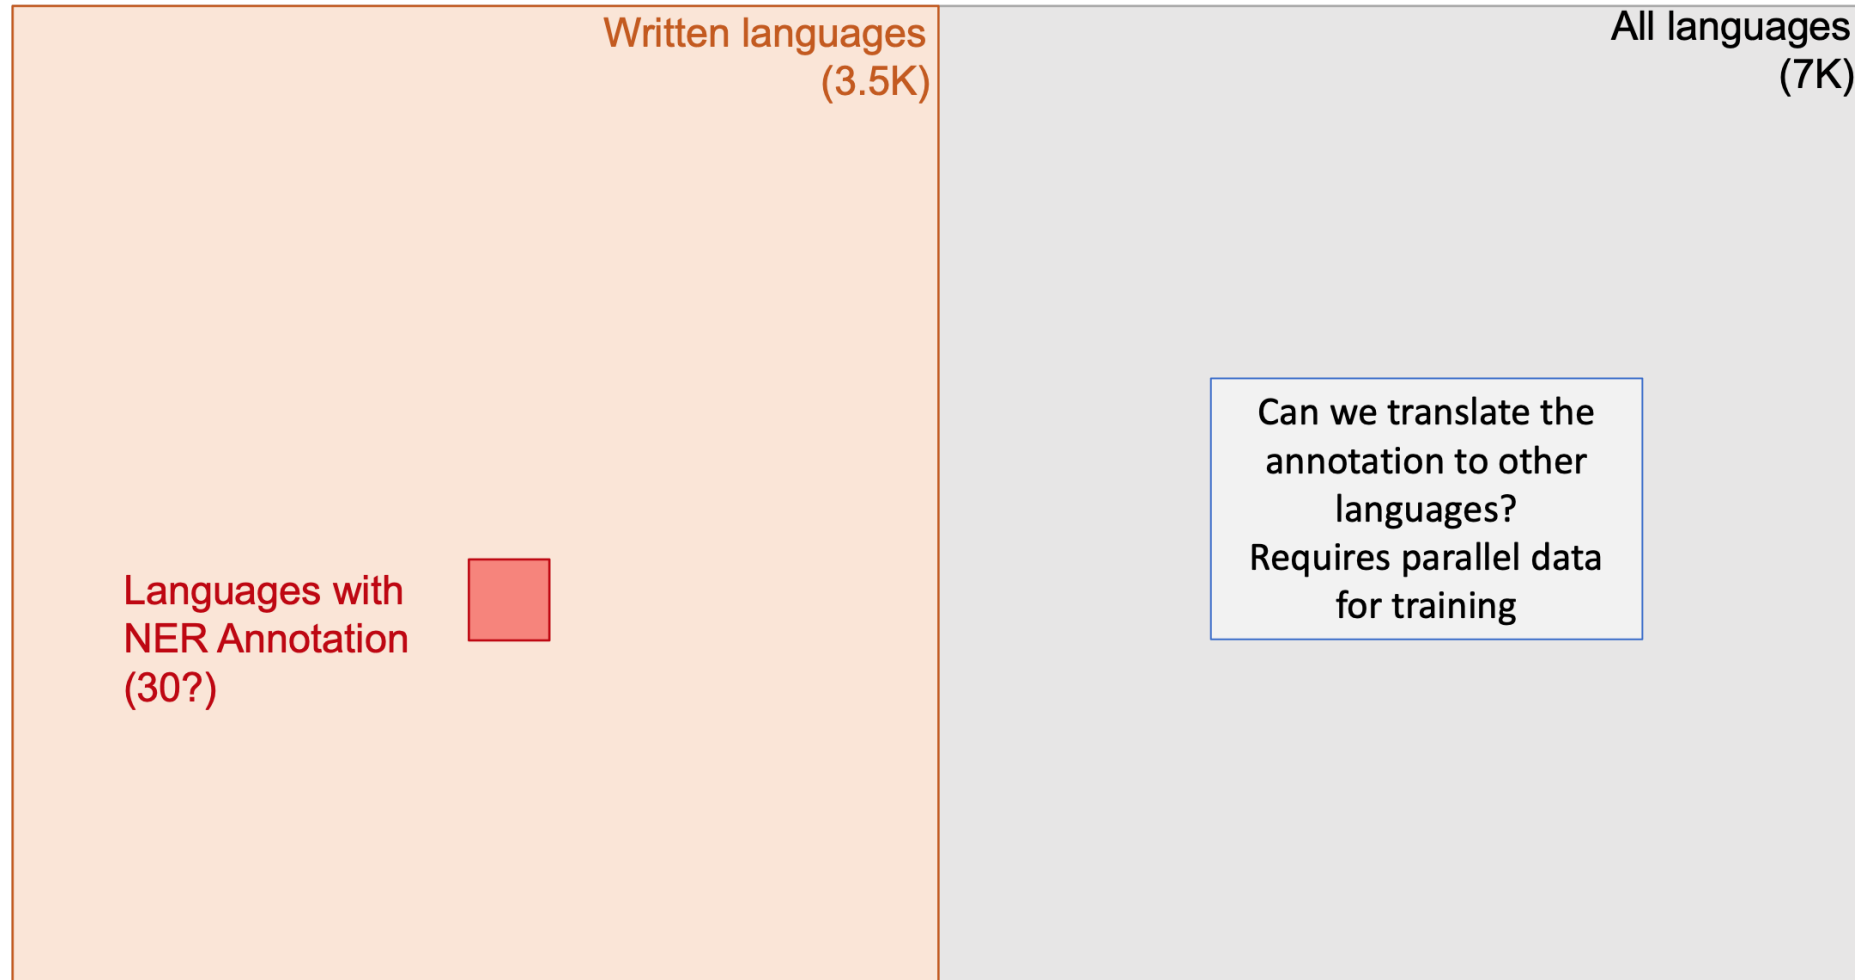
# Problems with few data (labels)

- Specific domain      Low-resource languages



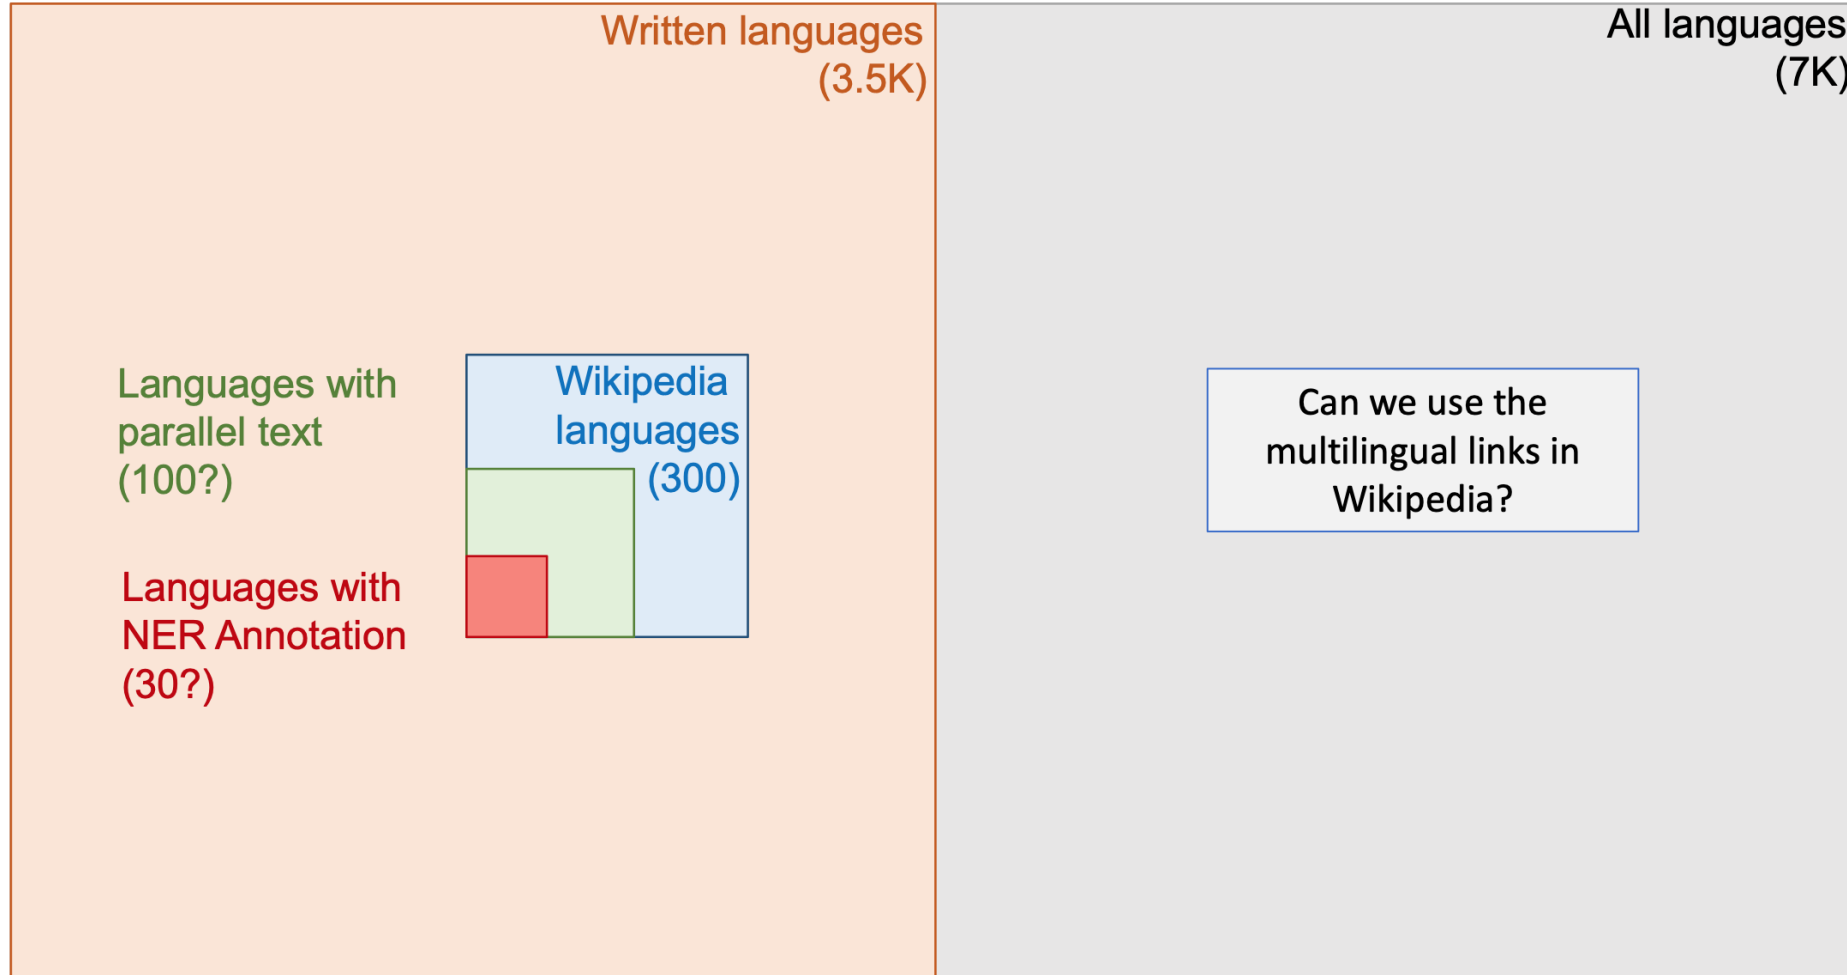
# Problems with few data (labels)

- Specific domain      Low-resource languages



# Problems with few data (labels)

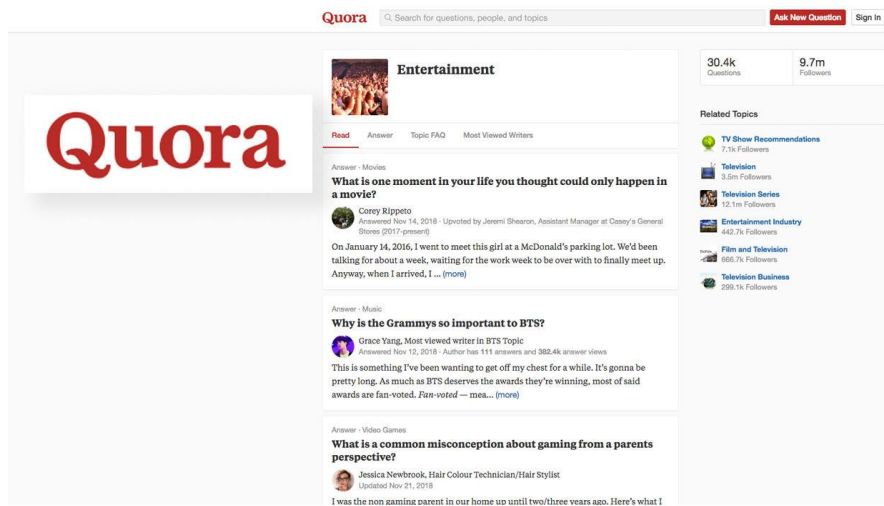
- Specific domain
- Low-resource languages



# Problems with few data (labels)

- Specific domain

Question answering



QA based on car manual?





# Problems with few data (labels)

- Privacy, security issues
- Expensive to collect/annotate
- Difficult / expertise-demanding to annotate
- Specific domain

# Machine learning solutions given few data (labels)

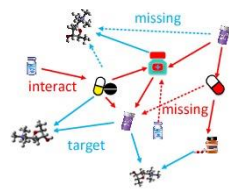
- How can we make more efficient use of **data**?
  - Clean but small-size
  - Noisy
  - Out-of-domain
- Can we incorporate **other types of experience** in learning?



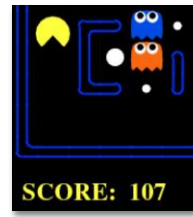
*Data examples*

Type-2 diabetes is 90% more common than type-1

*Rules/Constraints*



*Knowledge graphs*



*Rewards*



*Auxiliary agents*



*Adversaries*

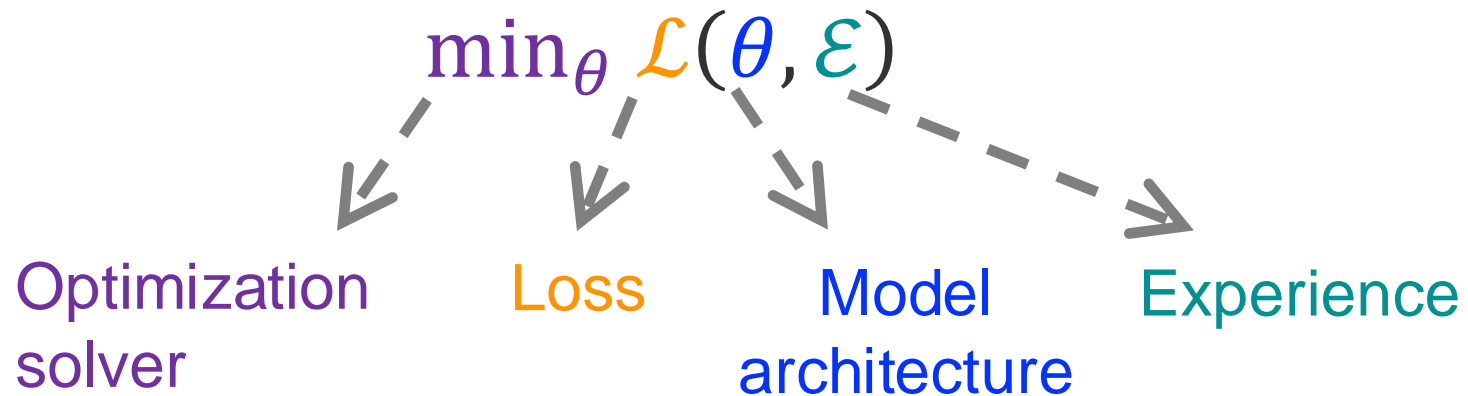


*Master classes*

... *And all combinations thereof*

# Components of a ML solution (roughly)

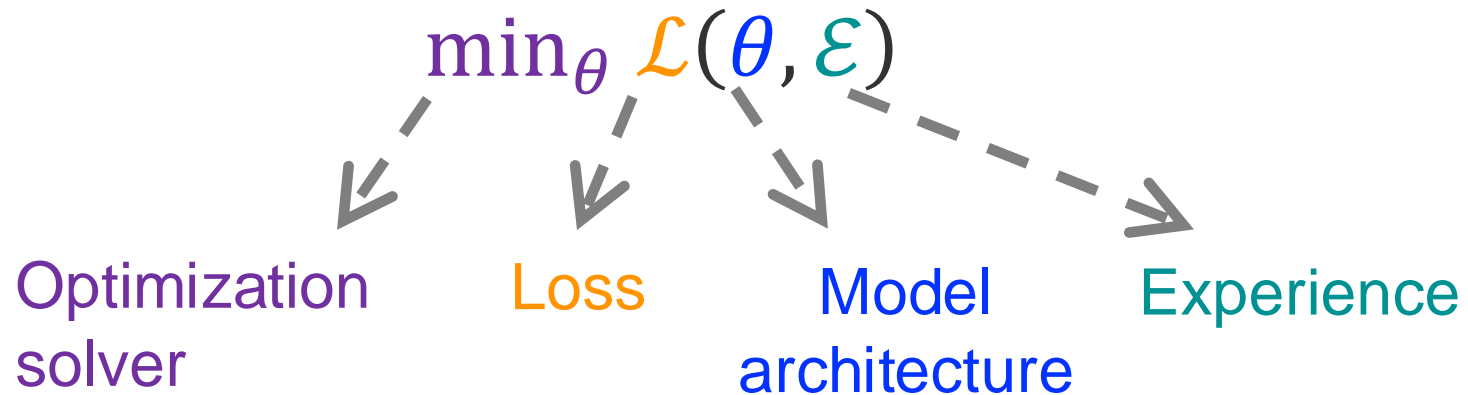
- Loss
- Experience
- Optimization solver
- Model architecture



# Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course does **not** discuss model architecture



# Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned,  $p_{\theta}(x, y)$  or  $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures

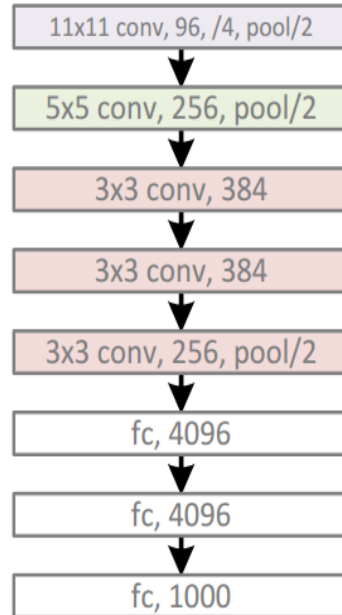
# Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

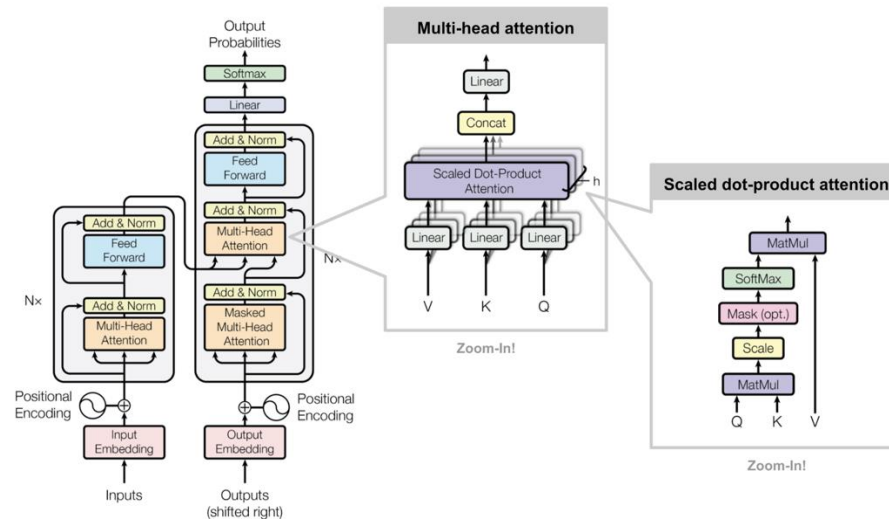
This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned,  $p_{\theta}(x, y)$  or  $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures



Convolutional networks



Transformers

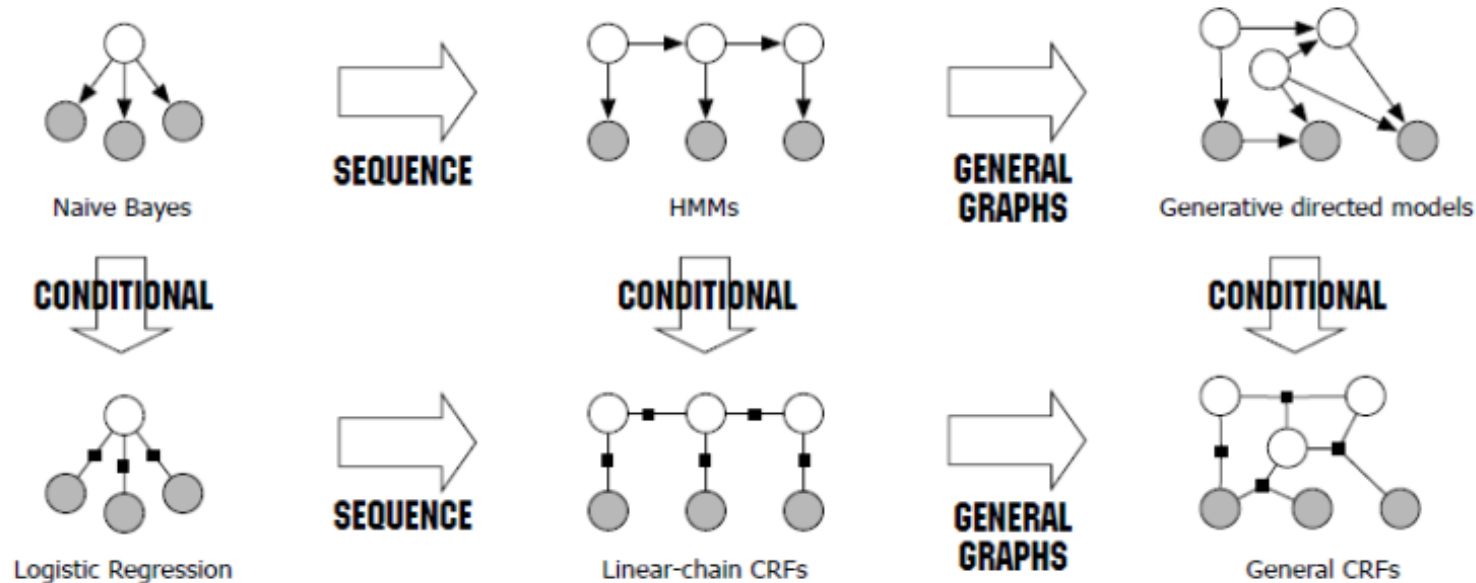
# Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned,  $p_{\theta}(x, y)$  or  $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures



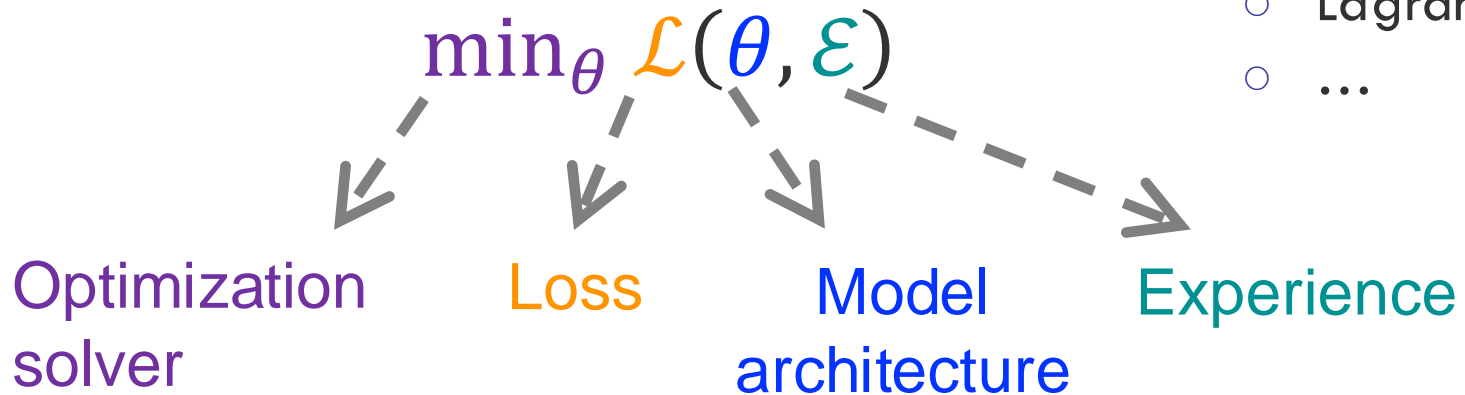
# Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course discusses *a little* about optimization

Assuming you know basic procedures:

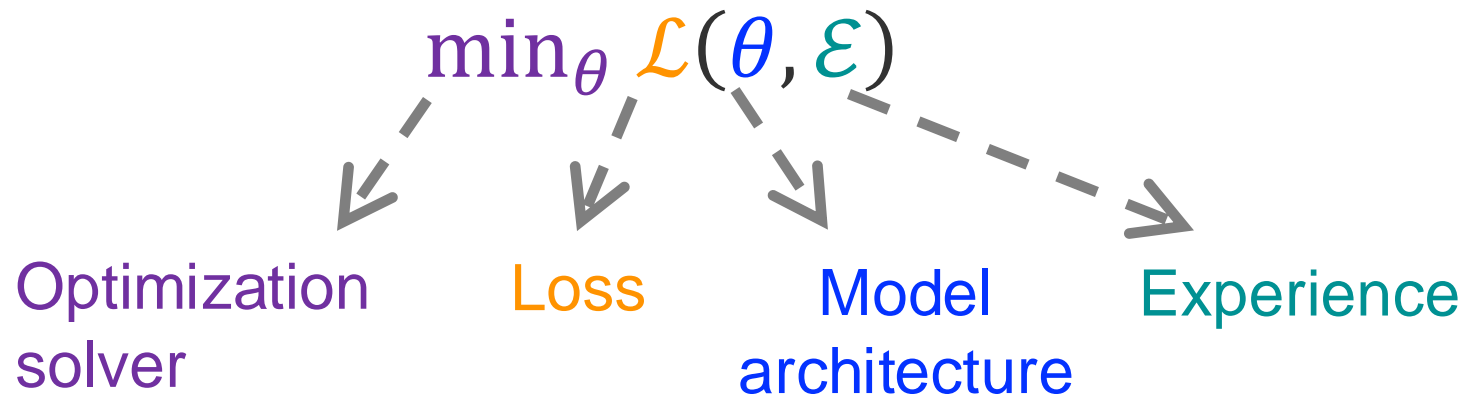
- (Stochastic) gradient descent
- Backpropagation
- Lagrange multiplier
- ...





# Components of a ML solution (roughly)

- **Loss** This course discusses a lot of loss & experience
- **Experience** Core of most learning algorithms
- Optimization solver
- Model architecture



# Machine learning solutions given few data (labels)

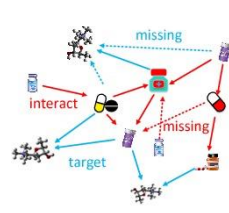
- (1) How can we make more efficient use of **data**?
  - Clean but small-size, Noisy, Out-of-domain
- (2) Can we incorporate **other types of experience** in learning?



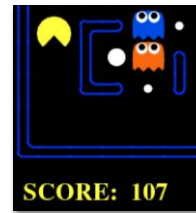
*Data examples*

Type-2 diabetes is 90% more common than type-1

*Rules/Constraints*



*Knowledge graphs*



*Rewards*



*Auxiliary agents*



*Adversaries*



*Master classes*

...

*And all combinations thereof*

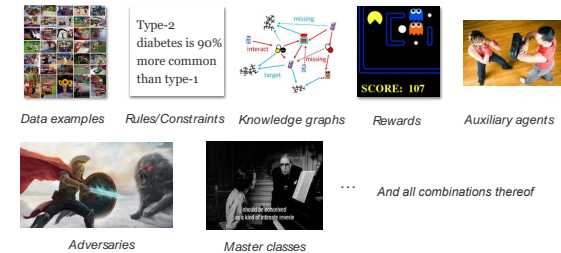
# Machine learning solutions given few data (labels)

- (1) How can we make more efficient use of **data**?
  - Clean but small-size, Noisy, Out-of-domain, ...
- Algorithms
  - **Supervised learning**: MLE, maximum entropy principle
  - **Unsupervised learning**: EM, variational inference, VAEs
  - **Self-supervised learning**: successful instances, e.g., BERT, GPT-3, contrastive learning, applications to downstream tasks
  - **Distant/weakly supervised learning**: successful instances
  - **Data manipulation**: augmentation, re-weighting, curriculum learning, ...
  - **Meta-learning**

Mostly first half of the course

# Machine learning solutions given few data (labels)

- (2) Can we incorporate **other types of experience** in learning?
  - Learning from auxiliary models, e.g., adversarial models:
    - Generative adversarial learning (GANs and variants), co-training, ...
  - Learning from structured knowledge
    - Posterior regularization, constraint-driven learning, ...
  - Learning from rewards
    - Reinforcement learning: model-free vs model-based, policy-based vs value-based, on-policy vs off-policy, extrinsic reward vs intrinsic reward, ...
  - Learning in dynamic environment (*not covered*)
    - Online learning, lifelong/continual learning, ...



# Algorithm marketplace

Designs driven by: experience, task, loss function, training procedure ...

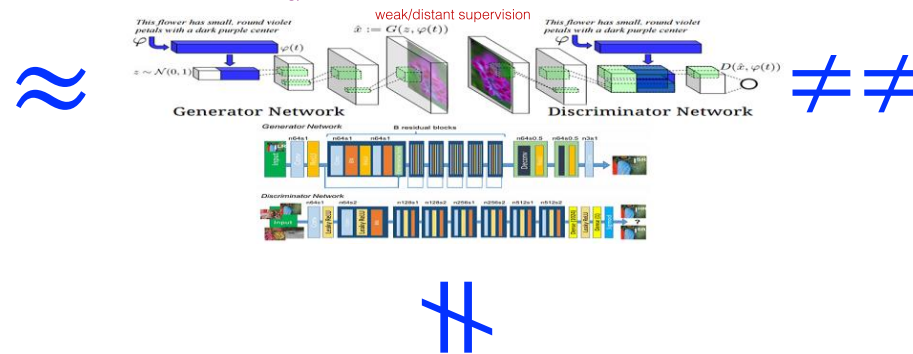


maximum likelihood estimation    reinforcement learning as inference  
data re-weighting    inverse RL    active learning  
policy optimization  
data augmentation    reward-augmented maximum likelihood  
label smoothing    imitation learning    softmax policy gradient  
actor-critic    adversarial domain adaptation  
GANs    posterior regularization  
knowledge distillation    intrinsic reward    constraint-driven learning  
prediction minimization    generalized expectation  
regularized Bayes  
energy-based GANs    learning from measurements  
weak/distant supervision

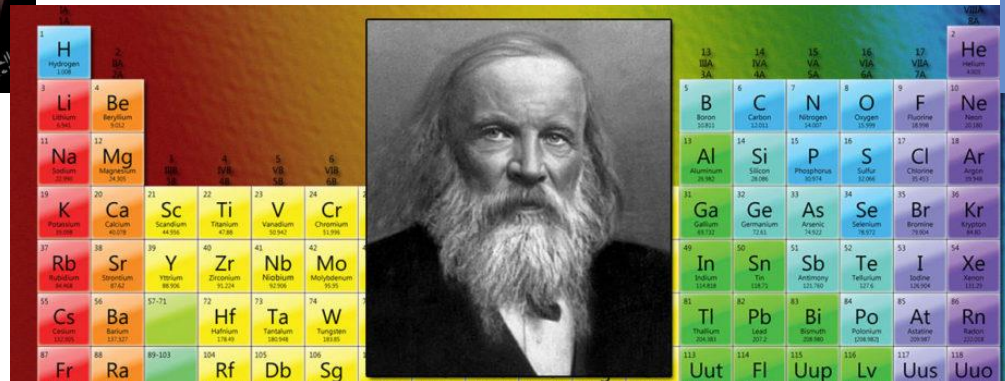
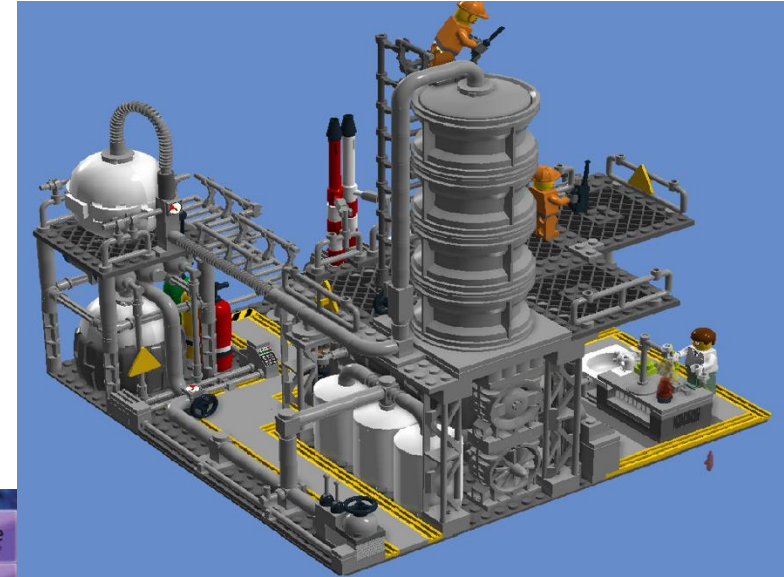
# Where we are now? Where we want to be?

- Alchemy vs chemistry

maximum likelihood estimation    reinforcement learning as inference  
 data re-weighting    inverse RL    active learning  
 data augmentation    policy optimization    reward-augmented maximum likelihood  
 label smoothing    imitation learning    softmax policy gradient  
 actor-critic    adversarial domain adaptation  
 GANs    posterior regularization  
 knowledge distillation    intrinsic reward    constraint-driven learning  
 prediction minimization    generalized expectation  
 regularized Bayes    learning from measurements  
 energy-based GANs



© Mind Juice Media Inc. All rights reserved



# Quest for more standardized, unified ML principles

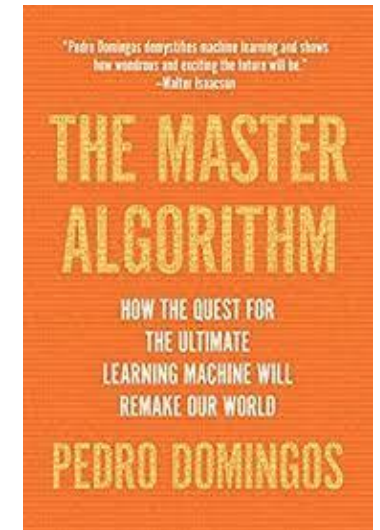
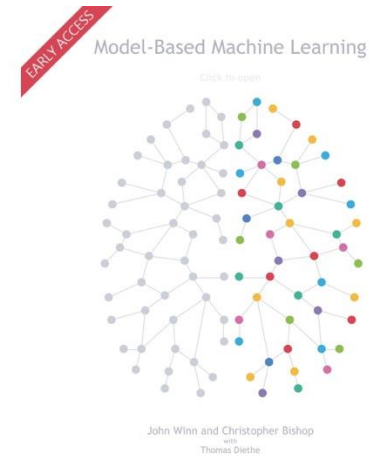
Machine Learning 3: 253–259, 1989

© 1989 Kluwer Academic Publishers – Manufactured in The Netherlands

EDITORIAL

Toward a Unified Science of Machine Learning

[P. Langley, 1989]



REVIEW 

---

 Communicated by Steven Nowlan

**A Unifying Review of Linear Gaussian Models**

**Sam Roweis\***

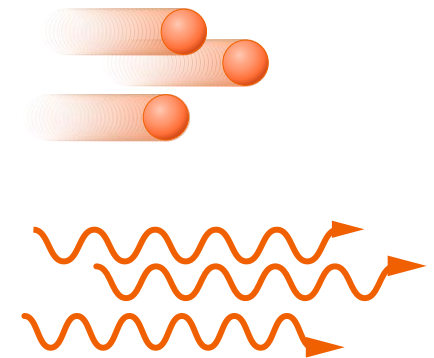
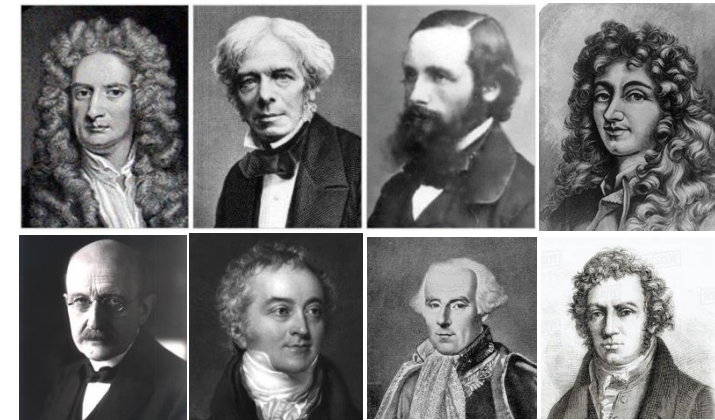
*Computation and Neural Systems, California Institute of Technology, Pasadena, CA 91125, U.S.A.*

**Zoubin Ghahramani\***

*Department of Computer Science, University of Toronto, Toronto, Canada*

# Physics in the 1800's

- Electricity & magnetism:
  - Coulomb's law, Ampère, Faraday, ...
- Theory of light beams:
  - Particle theory: Isaac Newton, Laplace, Plank
  - Wave theory: Grimaldi, Chris Huygens, Thomas Young, Maxwell
- Law of gravity
  - Aristotle, Galileo, Newton, ...





# “Standard equations” in Physics

Maxwell's Eqns:  
original form

$e + \frac{df}{dx} + \frac{dg}{dy} + \frac{dh}{dz} = 0$	(1) Gauss' Law
$\mu\alpha = \frac{dH}{dy} - \frac{dG}{dz}$ $\mu\beta = \frac{dF}{dz} - \frac{dH}{dx}$ $\mu\gamma = \frac{dG}{dx} - \frac{dF}{dy}$	(2) Equivalent to Gauss' Law for magnetism
$P = \mu \left( \gamma \frac{dy}{dt} - \beta \frac{dz}{dt} \right) - \frac{dF}{dt} - \frac{d\Psi}{dz}$ $Q = \mu \left( \alpha \frac{dz}{dt} - \gamma \frac{dx}{dt} \right) - \frac{dG}{dt} - \frac{d\Psi}{dy}$ $R = \mu \left( \beta \frac{dx}{dt} - \alpha \frac{dy}{dt} \right) - \frac{dH}{dt} - \frac{d\Psi}{dx}$	(3) Faraday's Law (with the Lorentz Force and Poisson's Law)
$\frac{d\gamma}{dy} - \frac{d\beta}{dz} = 4\pi p'$ $\frac{d\alpha}{dz} - \frac{d\gamma}{dx} = 4\pi q'$ $\frac{d\beta}{dx} - \frac{d\alpha}{dy} = 4\pi r'$	(4) Ampère-Maxwell Law
$P = -\xi p \quad Q = -\xi q \quad R = -\xi r$	Ohm's Law
$P = kf \quad Q = kg \quad R = kh$	The electric elasticity equation ( $\mathbf{E} = \mathbf{D}/\epsilon$ )
$\frac{de}{dt} + \frac{dp}{dx} + \frac{dq}{dy} + \frac{dr}{dz} = 0$	Continuity of charge

Maxwell's Eqns simplified w/ rotational symmetry

$$\nabla \cdot \mathbf{D} = \rho_V$$

$$\nabla \cdot \mathbf{B} = 0$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

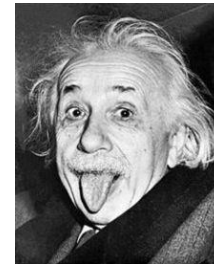
$$\nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}$$



Maxwell's Eqns further simplified w/ symmetry of special relativity

$$\epsilon^{uvk\lambda} \partial_v F_{k\lambda} = 0$$

$$\partial_v F^{uV} = \frac{4\pi}{c} j^u$$



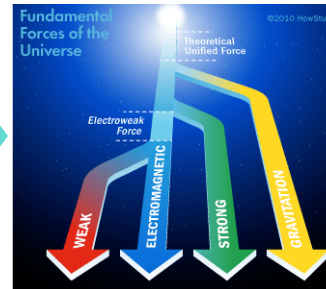
Standard Model w/ Yang-Mills theory and US(3) symmetry

$$\mathcal{L}_{gf} = -\frac{1}{2} \text{Tr}(F^2)$$

$$= -\frac{1}{4} F^{\alpha\mu\nu} F_{\mu\nu}^{\alpha}$$



Unification of fundamental forces?



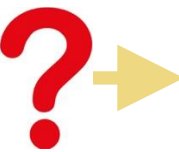
Diverse electro-magnetic theories



1861

1910s

1970s



# A “standardized formalism” of ML



Data examples

Type-2 diabetes  
is 90% more  
common than  
type-1

Constraints



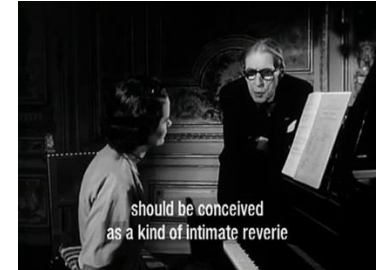
Rewards



Auxiliary agents



Adversaries



Imitation

$$\min_{q, \theta} - \mathbb{H} + \mathbb{D} - \mathbb{E}$$

Uncertainty      Divergence      Experience

- Panoramically learn from all types of experience
- Subsumes many existing algorithms as special cases

Will discuss in later in the class

**Questions?**