Multimodal AI

Lecture 9.1 – Reinforcement Learning & Reasoning

Paul Liang
Assistant Professor
MIT Media Lab & MIT EECS

https://pliang279.github.io
ppliang@mit.edu
@pliang279

mit media lab

MIT EECS

Assignments for This Coming Week

Project midterm report assignment released.

- Finalized main ideas and experimental setup, have datasets and baseline models working, detailed error analysis, initial progress towards implementing new ideas.

Schedule

Classes	Tuesday Lectures	Thursday Lectures	HWs
Week 1 2/3 & 2/5	**Course introduction (All)** • Multimodal core challenges • Course syllabus	**Multimodal datasets (Paul)** • Research tasks and datasets • Intro to AI research	
Week 2 2/10 & 2/12	**Datasets tutorial** • Data processing and visualization • Pytorch and modeling	**Unimodal representations (Paul)** • Dimensions of heterogeneity • Common model architectures	**HW1 out**
Week 3 2/17 & 2/19	***No class (schedule switch)***	**Multimodal fusion (Dimitris)** • Early and late fusion • Explainable fusion	**HW1 due** **HW2 out**
Week 4 2/24 & 2/26	**More fusion (Paul)** • Higher-order interactions • Multimodal fusion models	**Multimodal alignment (Paul)** • Multimodal grounding • Aligned representations	
Week 5 3/3 & 3/5	**Large multimodal models (Paul)** • Multimodal transformers • Pre-training & fine-tuning	**Multimodal LLM tutorial** • Fine-tuning • Instruction tuning	**HW2 due** **HW3 out**
Week 6 3/10 & 3/12	**Multimodal generation (Paul)** • Translation, summarization, creation • Model evaluation and ethics	**Modern generative AI (Paul)** • VAEs, diffusion, flow models • Controllable generation	

Schedule

Week 9 3/31 & 4/2	**Multimodal reasoning (Paul)** • Reinforcement learning • Multi-step reasoning	**Explainable modeling (Dimitris)** • Explainable AI • LLMs for explainability	**HW4 out**
Week 10 4/7 & 4/9	**Prescriptive modeling (Dimitris)** • Optimization & prediction • Causal & counterfactual	**Multimodal interaction (Paul)** • Interactive agents • Interactive reasoning	
Week 11 4/14 & 4/16	**Multimodal & design (Sang-Gook)** • Multimodal agents • Applications in design	**Multimodal & manufacturing (Sang-Gook)** • Multisensor fusion • Applications in manufacturing	**HW4 due** **HW5 out**
Week 12 4/21 & 4/23	**Human-AI interaction (Paul)** • Interaction mediums • Human-in-the-loop and safety	**Agents tutorial** • Multimodal agent pipelines • Agent evaluation	
Week 13 4/28 & 4/30	**Multimodal & cities (Jinhua)** • Spatial and temporal fusion • Applications in cities	**Multimodal & transportation (Jinhua)** • Multimodal agents • Applications in transportation	**HW5 due**
Week 14 5/5 & 5/7	**Cross-modal transfer (Paul)** • Modality transfer and co-learning • Self-training and multitask learning	**Self-evolving AI (Paul)** • Continual learning • Self-evolution and optimization	
Week 15 5/12	***Project presentations***		

Today's lecture

1. Explicit reasoning with knowledge and structure
2. Basics of reinforcement learning
3. Modern RL for implicit reasoning

Reasoning

Definition: Combining knowledge, usually through multiple inferential steps, exploiting the structure of the problem.

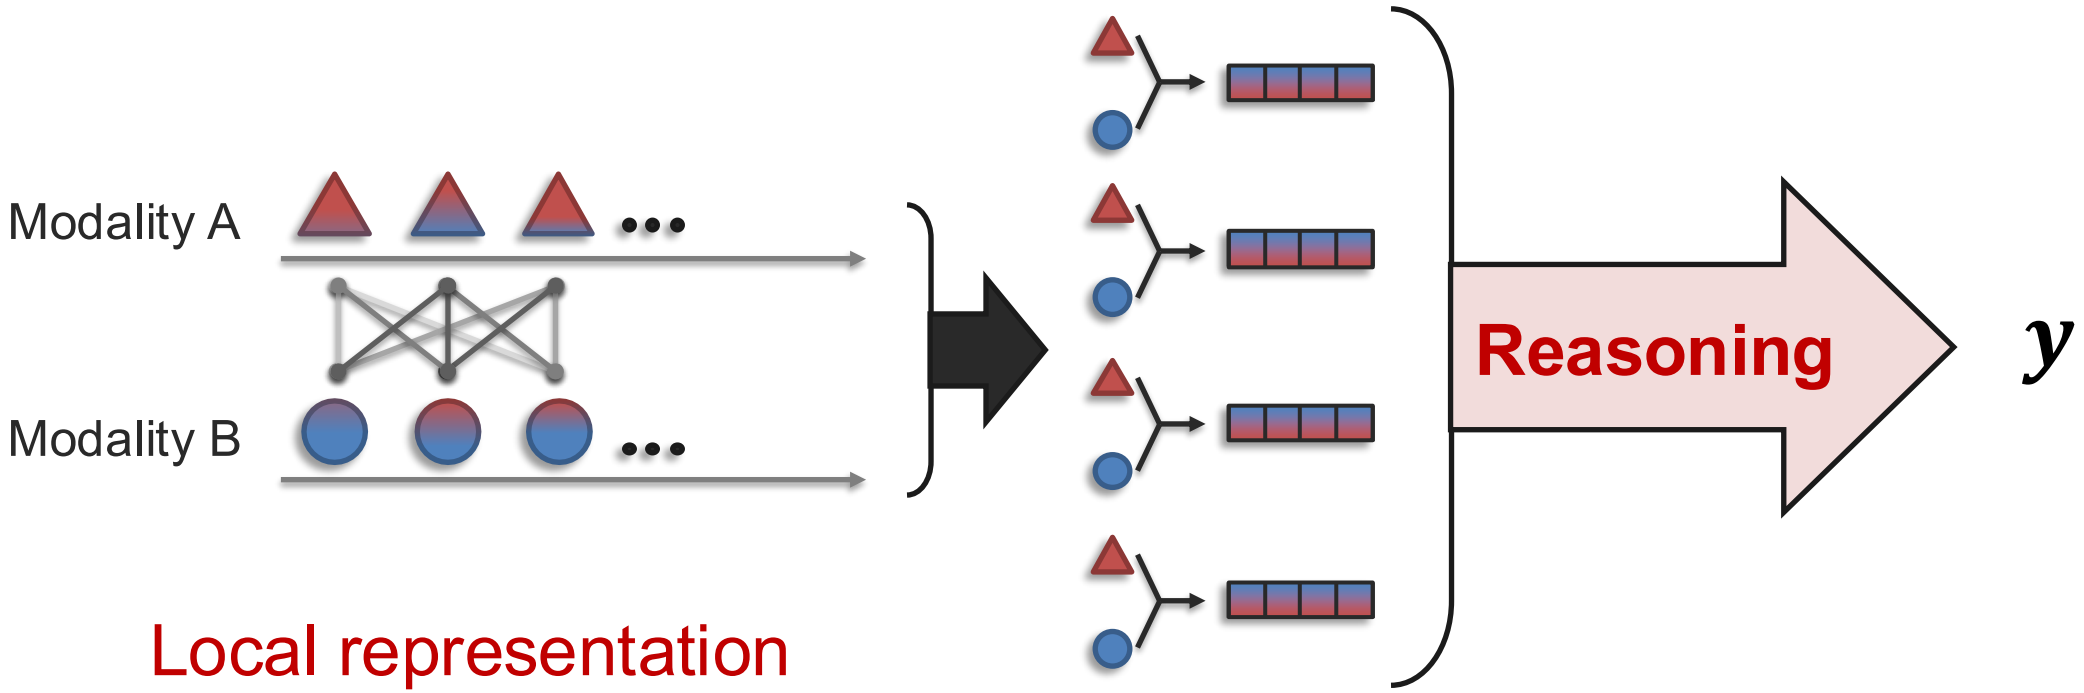

Local representation

+ Aligned representation

multisensory intelligence

The Challenge of Compositionality

Definition: Combining knowledge, usually through multiple inferential steps, exploiting multimodal alignment and problem structure.

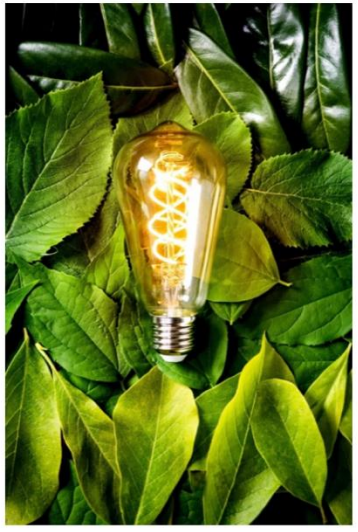

(a) some plants surrounding a lightbulb

(b) a lightbulb surrounding some plants

CLIP, ViLT, ViLBERT, etc.
All random chance

Compositional Generalization
to novel combinations outside of training data

1. Structure: <subject> <verb> <object>
2. Concepts: ‘plants’, ‘lightbulb’
3. Inference: ‘surrounding’ – spatial relation
4. Knowledge: from humans!

[Thrush et al., Winoground: Probing Vision and Language Models for Visio-Linguistic Compositionality. CVPR 2022]

multisensory intelligence

The Challenge of Compositionality

Definition: Combining knowledge, usually through multiple inferential steps, exploiting multimodal alignment and problem structure.

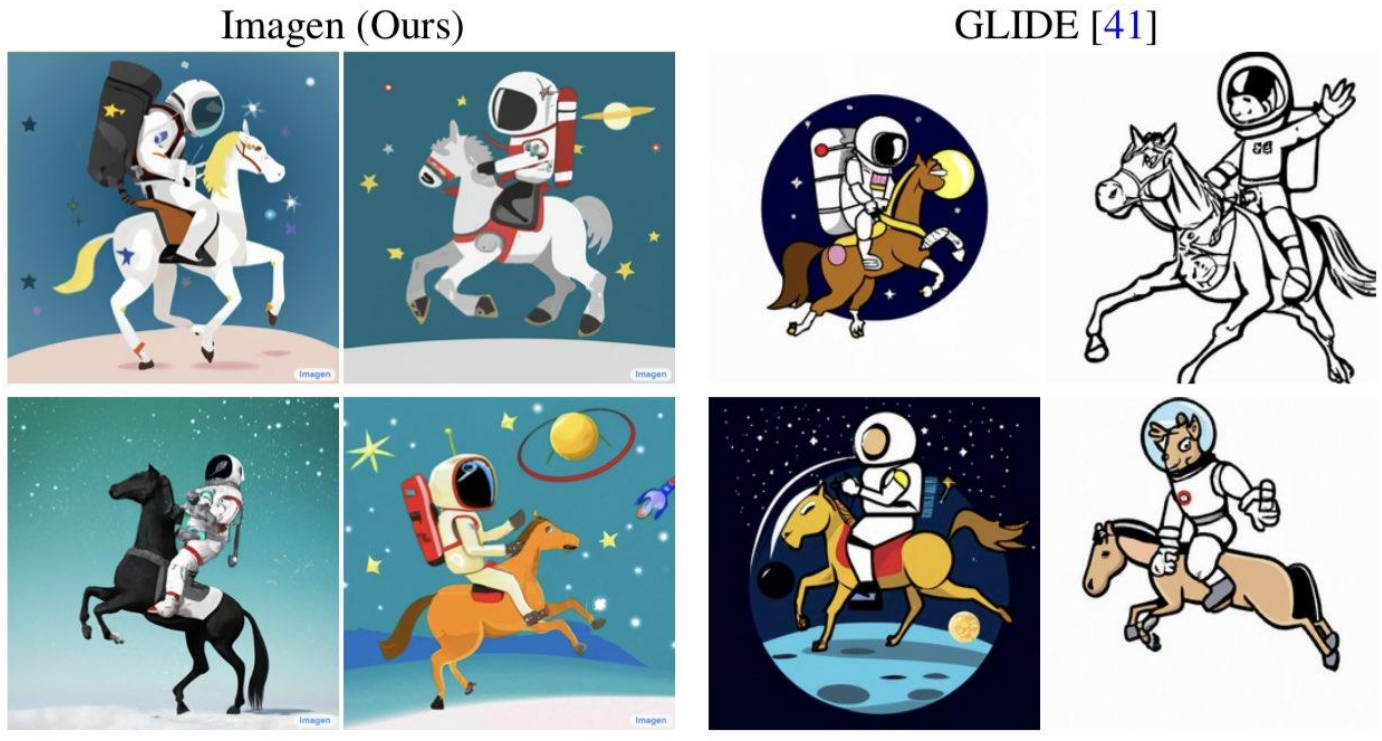

[Saharia et al., Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. arXiv 2022]

multisensory intelligence

Chain of Thought Reasoning

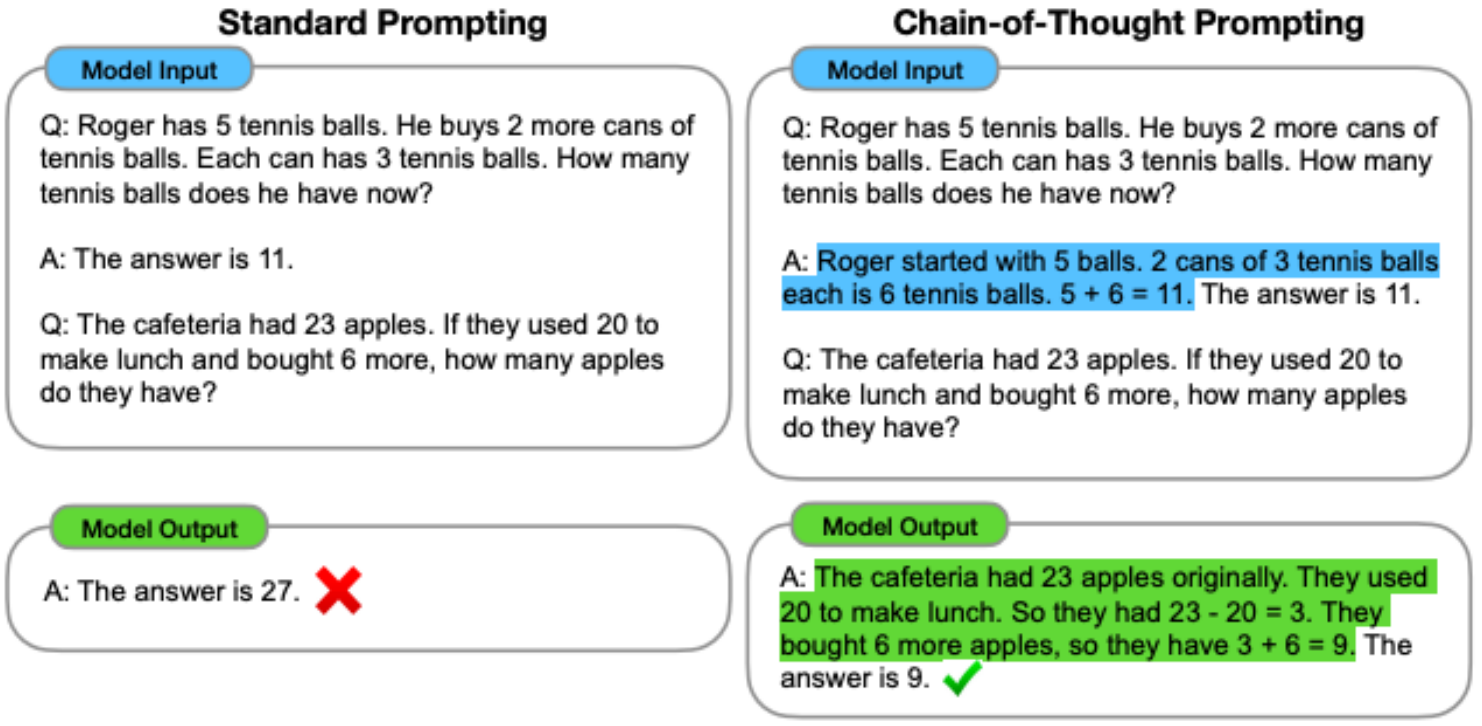

Figure 1: Chain-of-thought prompting enables large language models to tackle complex arithmetic, commonsense, and symbolic reasoning tasks. Chain-of-thought reasoning processes are highlighted.

[Wei et al., Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. NeurIPS 2022]

multisensory intelligence

Tree of Thoughts Reasoning

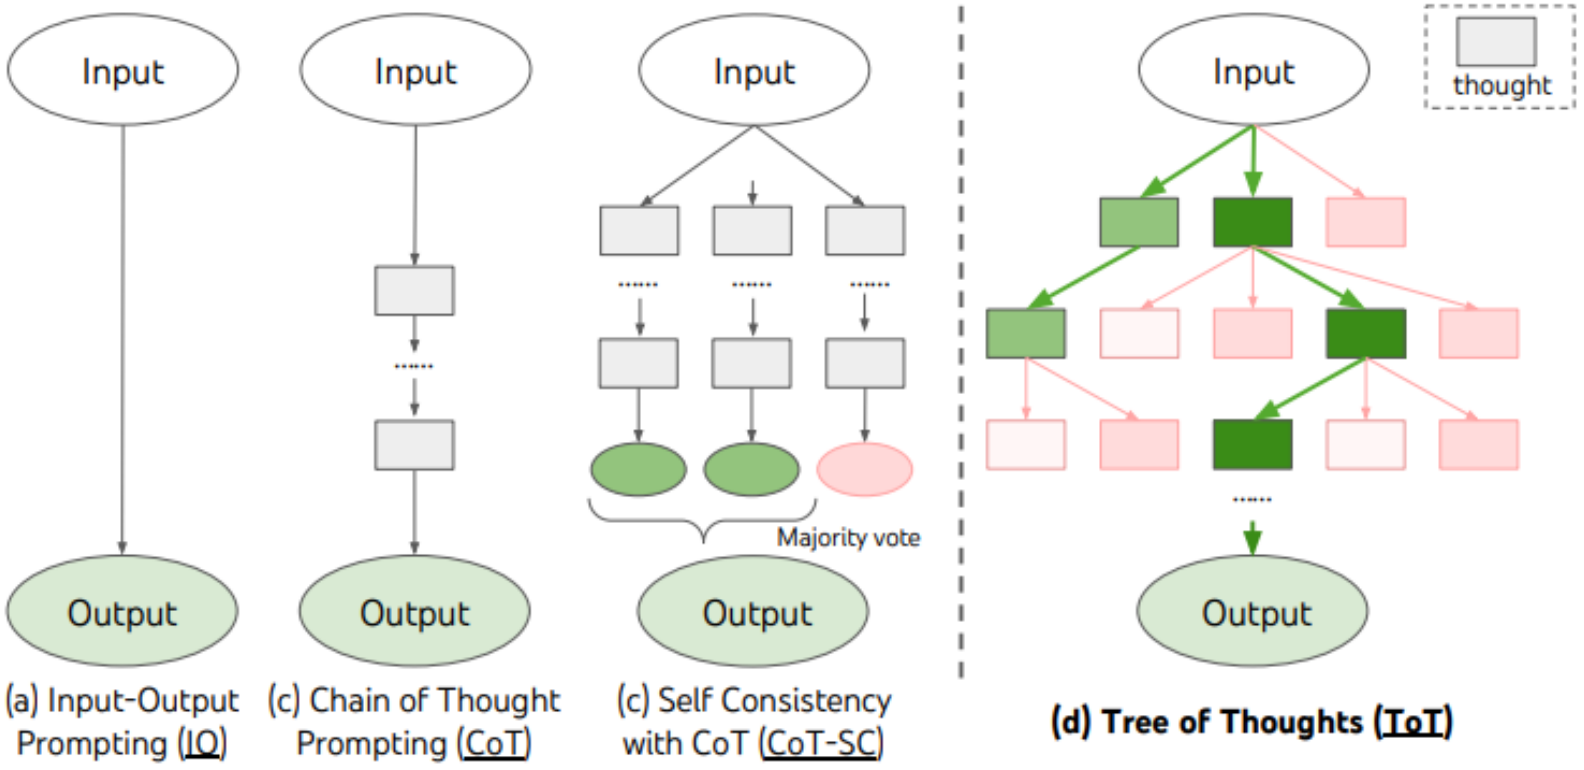

[Yao et al., Tree of Thoughts: Deliberate Problem Solving with Large Language Models. NeurIPS 2023]

multisensory intelligence

Multimodal Reasoning via Language

Image captioning

[Zeng et al., Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. arXiv 2022]

multisensory intelligence

Multimodal Reasoning via Language

Robot perception and planning

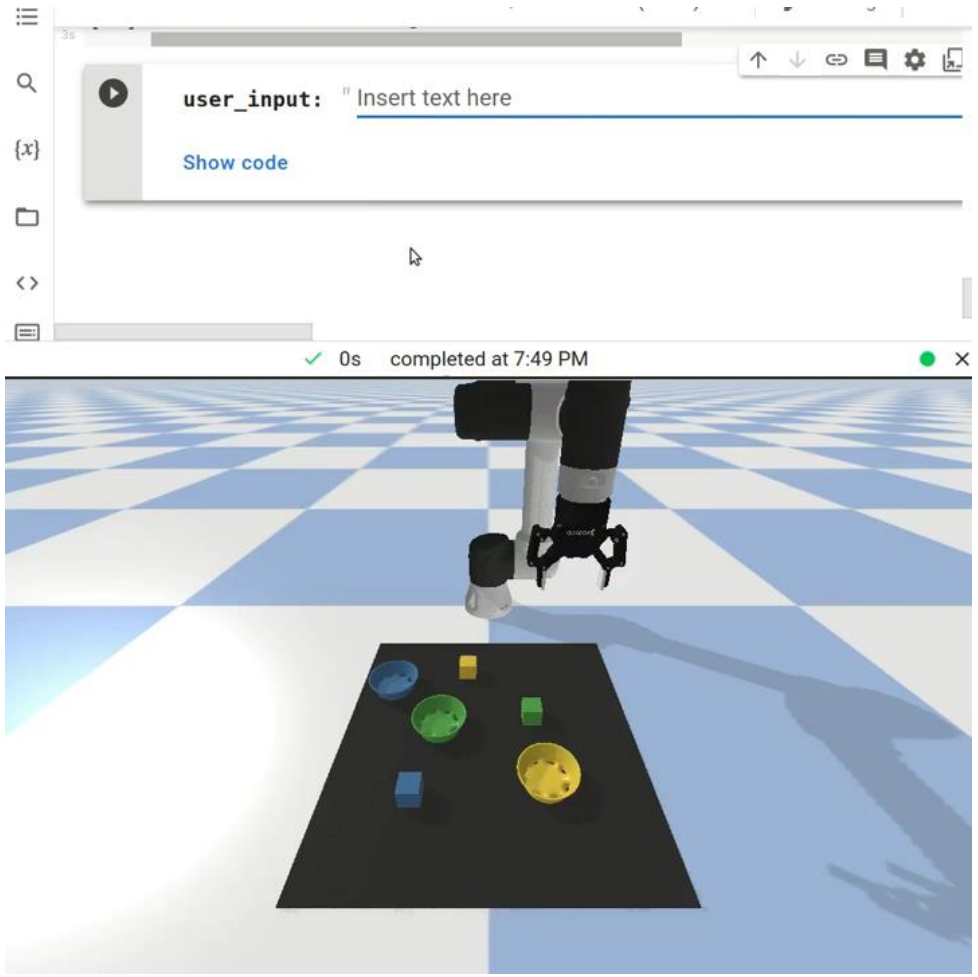

[Zeng et al., Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. arXiv 2022]

multisensory intelligence

Multimodal Reasoning via Language

Video reasoning

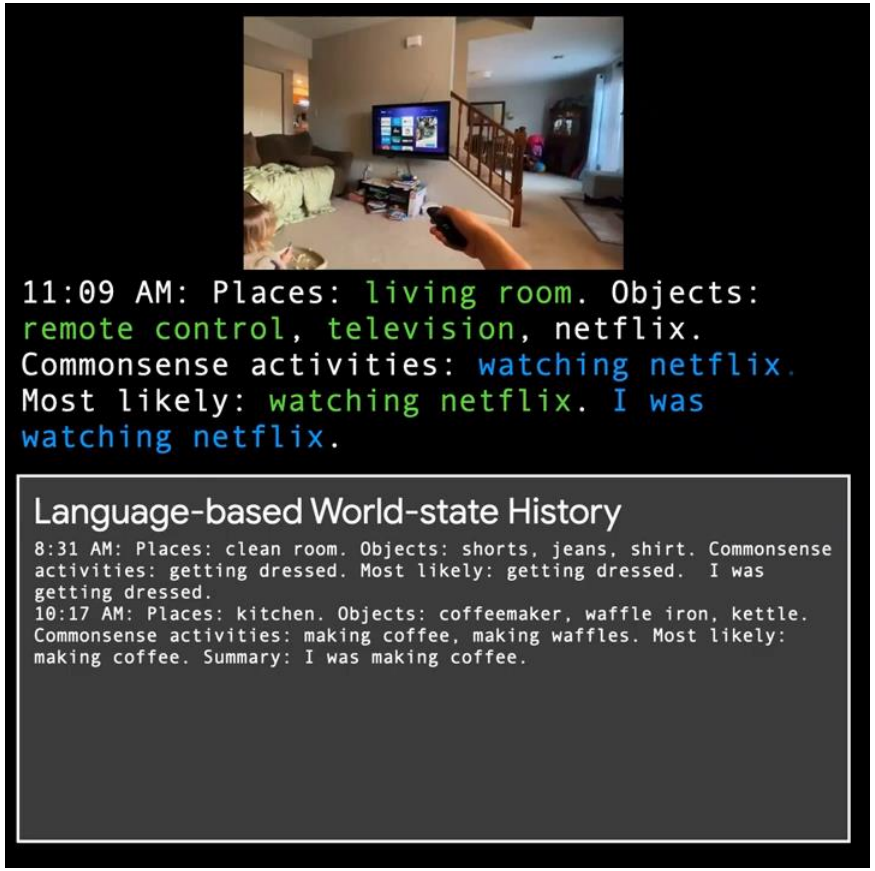

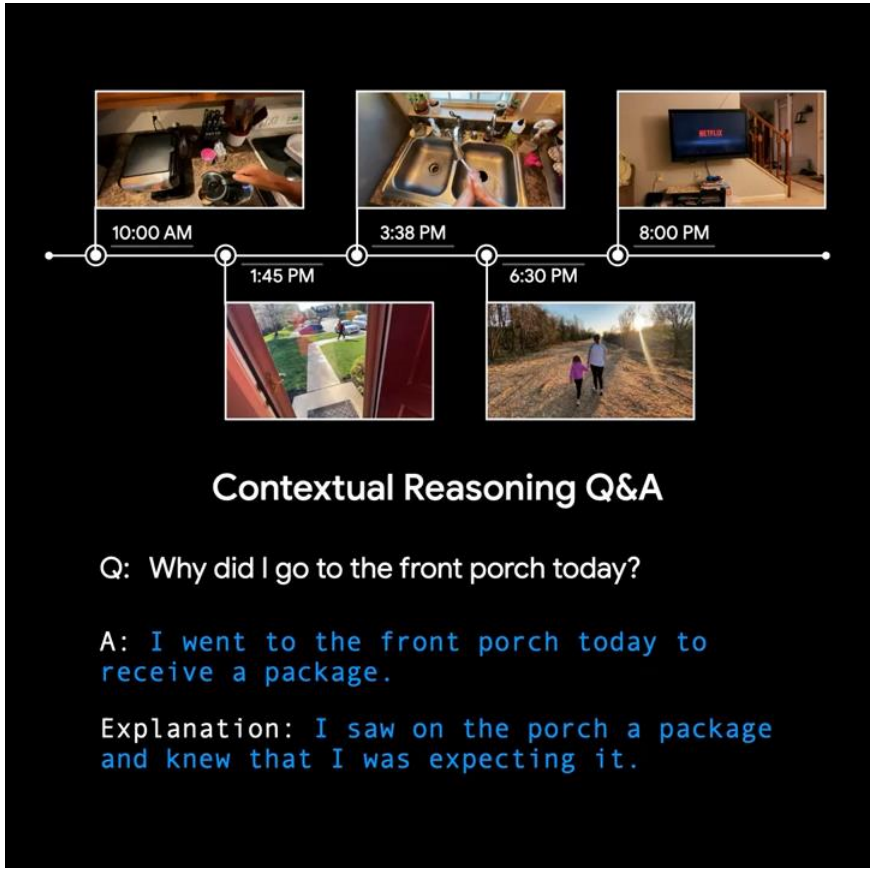

[Zeng et al., Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. arXiv 2022]

multisensory intelligence

Many Debates Surrounding Reasoning

Fully data-driven
Bottom-up

Fully domain knowledge
Top-down

Hybrid/neuro-symbolic
1. Differentiable?
2. Discrete or continuous concepts or representations?
3. Best mix of knowledge and data?

Implications on: interpretability, robustness, fairness, data + model efficiency, etc.

multisensory intelligence

Sub-challenge 1: Structure Modeling

Definition: Defining or learning the relationships over which reasoning occurs.

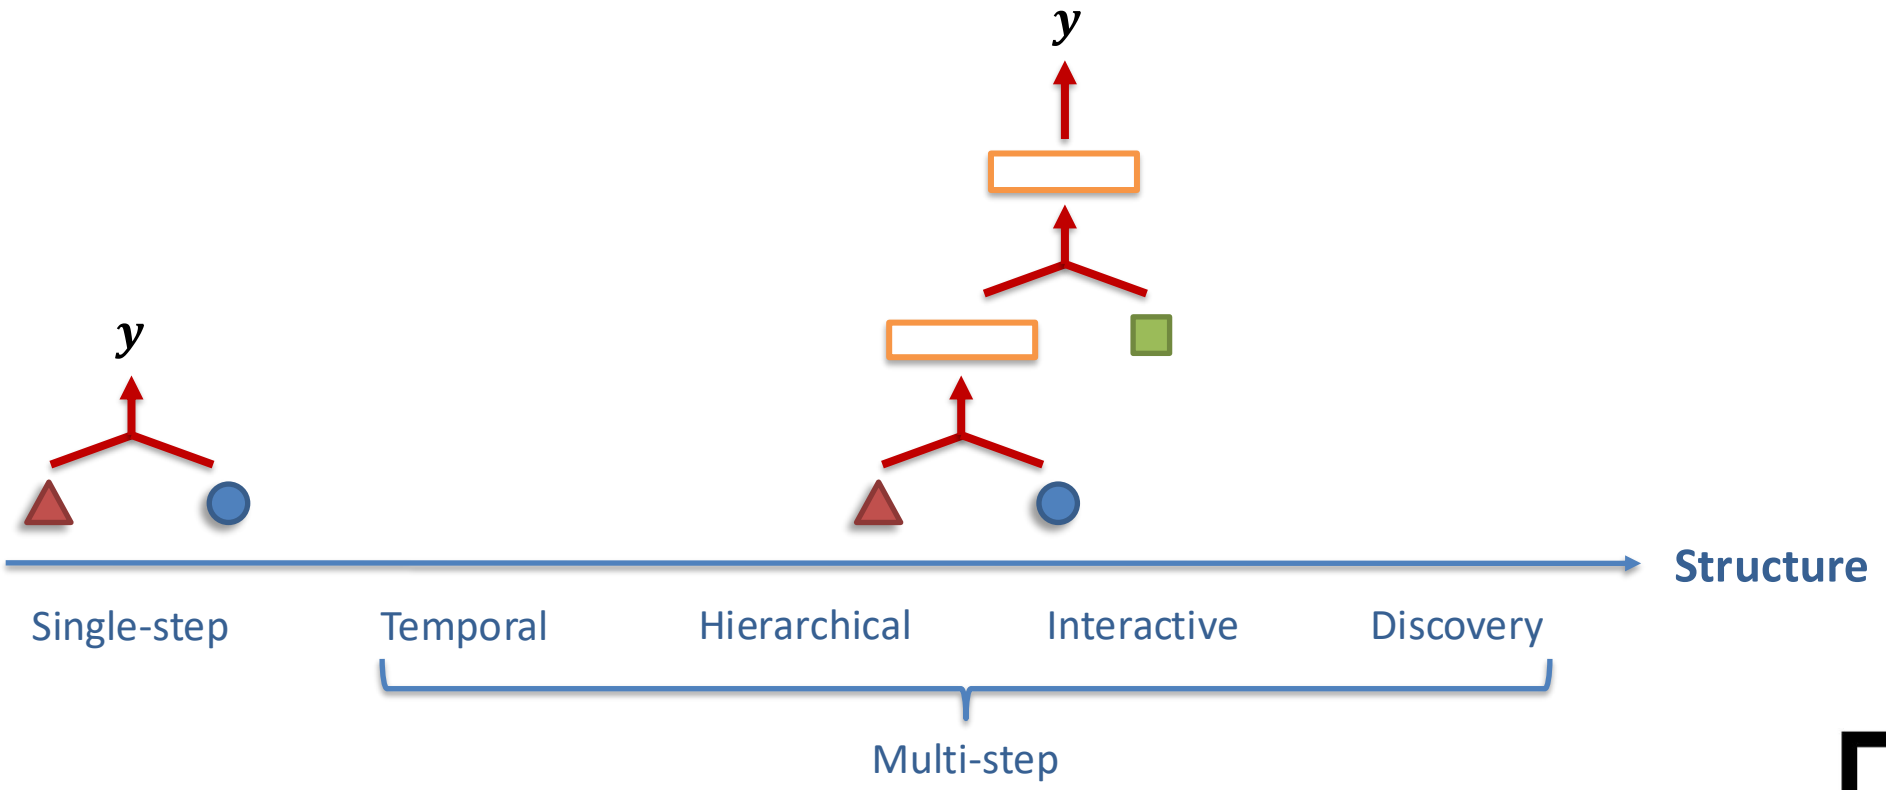

Sub-challenge 2: Intermediate Concepts

Definition: The parameterization of individual concepts in the reasoning process.

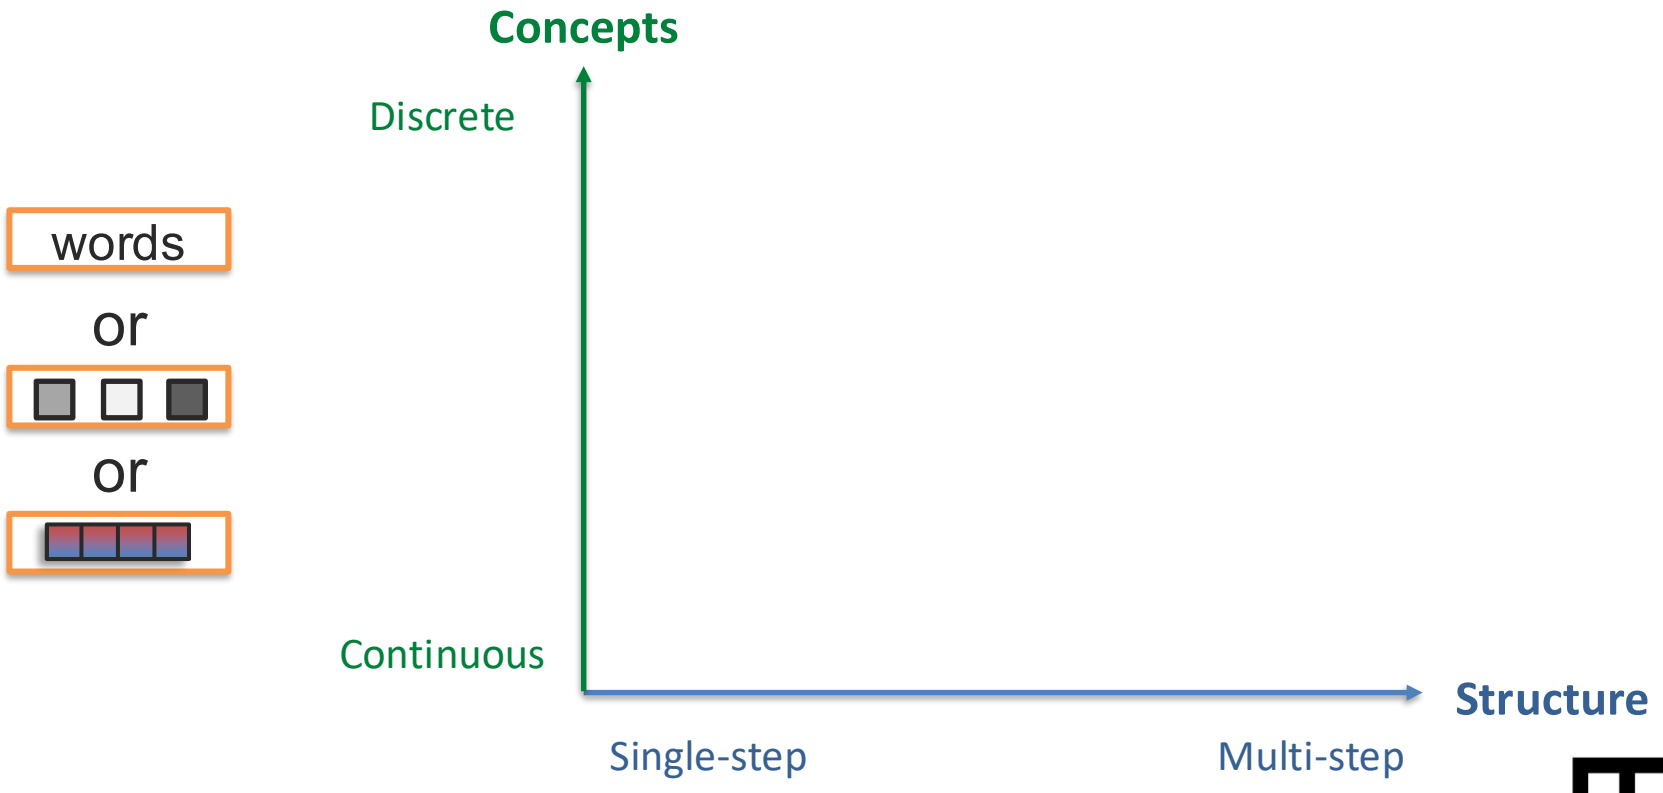

multisensory intelligence

Sub-challenge 3: Inference Paradigm

Definition: How increasingly abstract concepts are inferred from individual evidences.

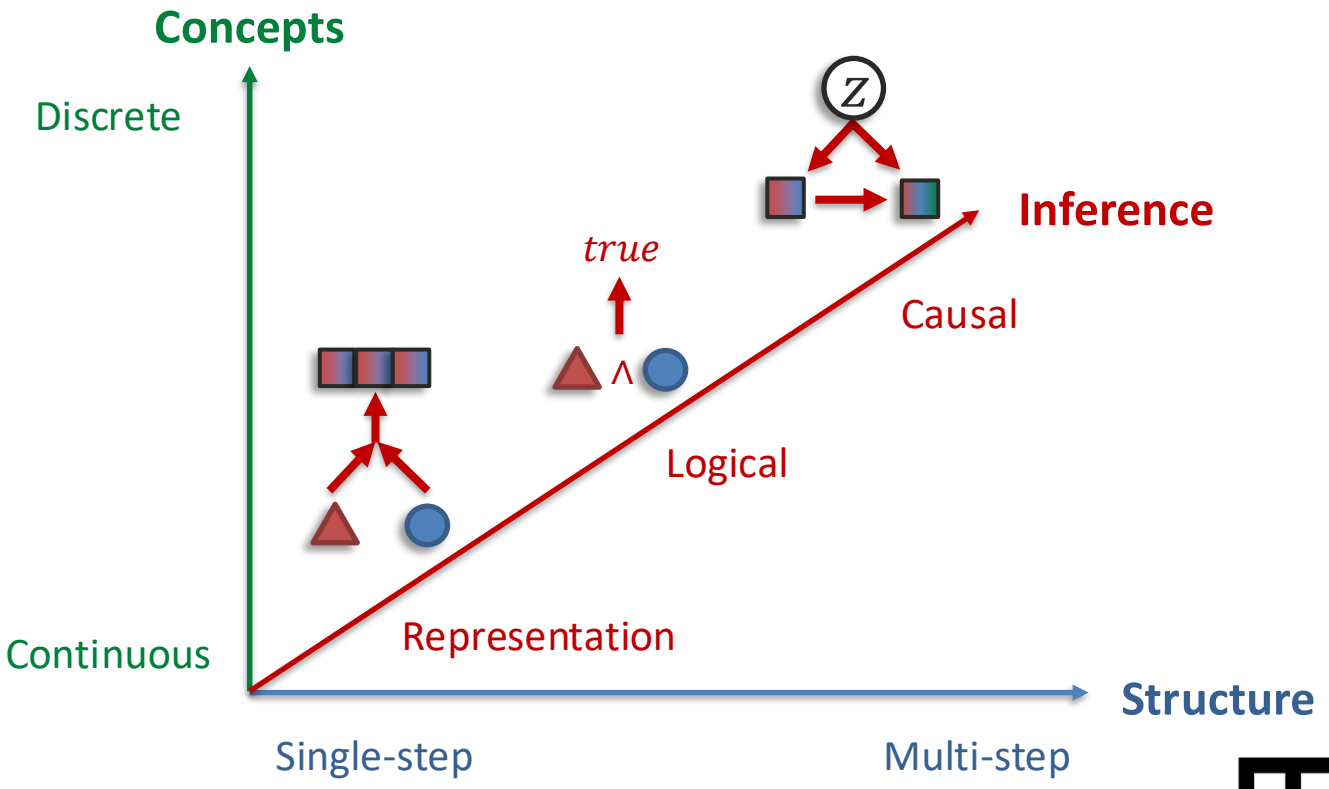

Sub-challenge 4: External Knowledge

Definition: Leveraging external knowledge in the study of structure, concepts, and inference.

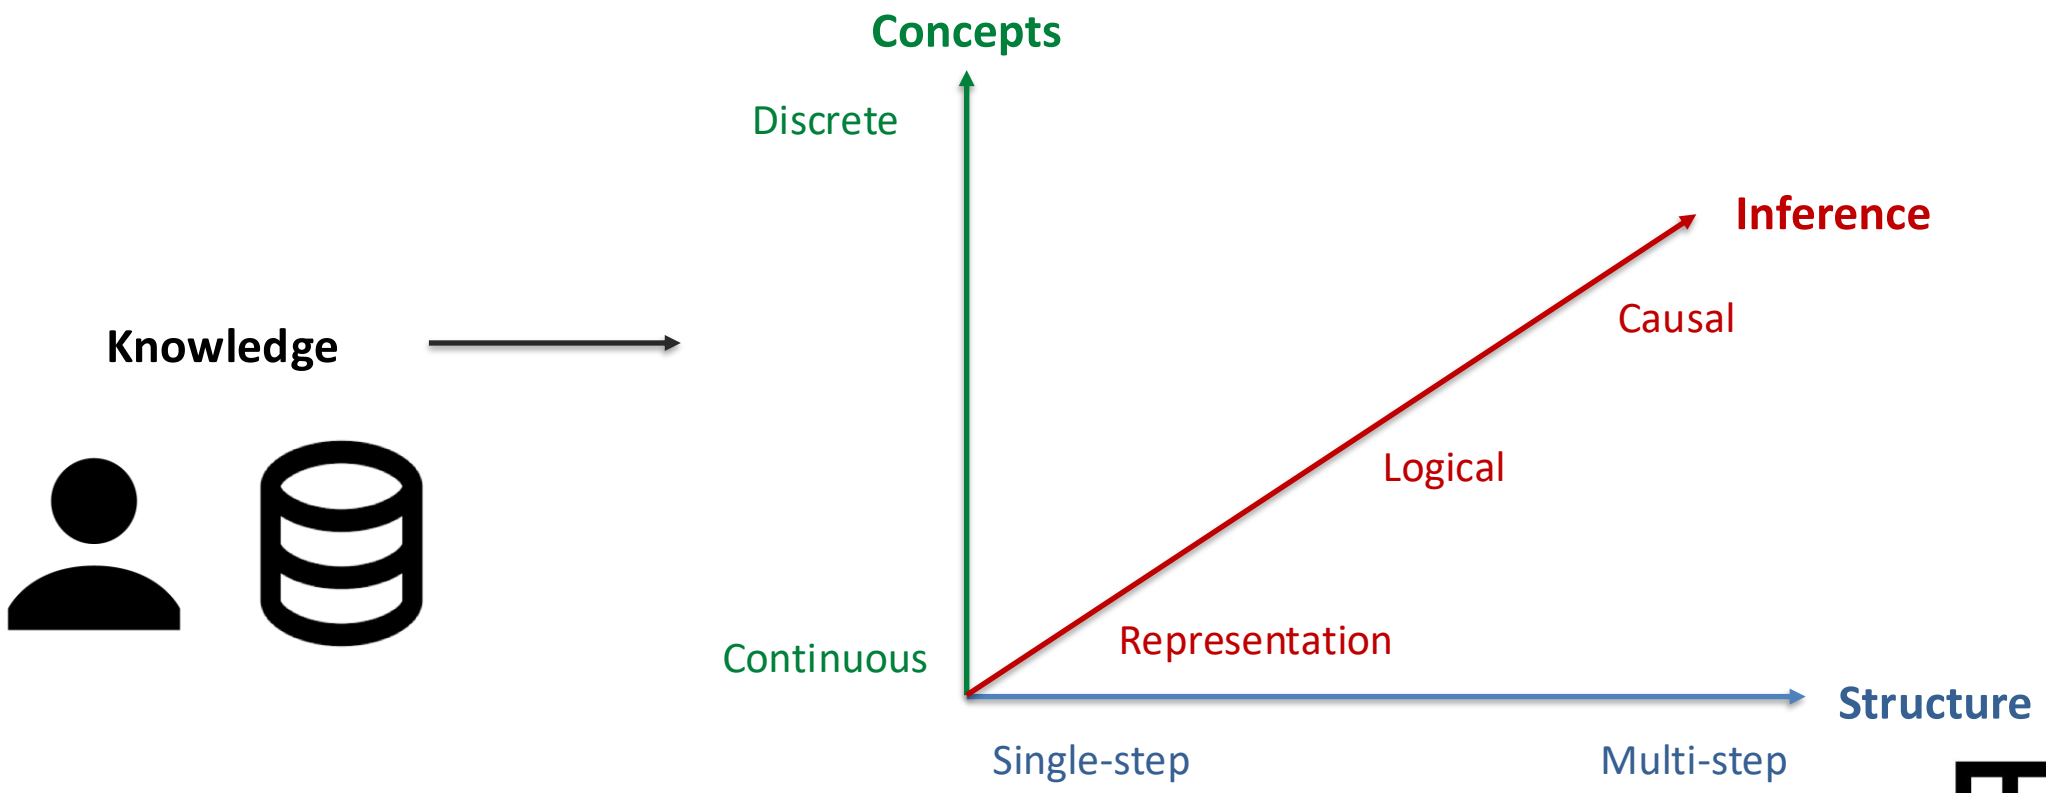

multisensory intelligence

Roadmap

Input, (reasoning step 1, step 2, step 3…), output

Method 1: Direct prompting, no training (e.g., CoT)

Method 2: Supervised fine-tuning
- Assuming you have reasoning traces

Method 3: Reinforcement learning **(today's focus)**
- Assumes no reasoning traces
- But a reward function scoring reasoning and outputs

Learning a Policy – RL basics

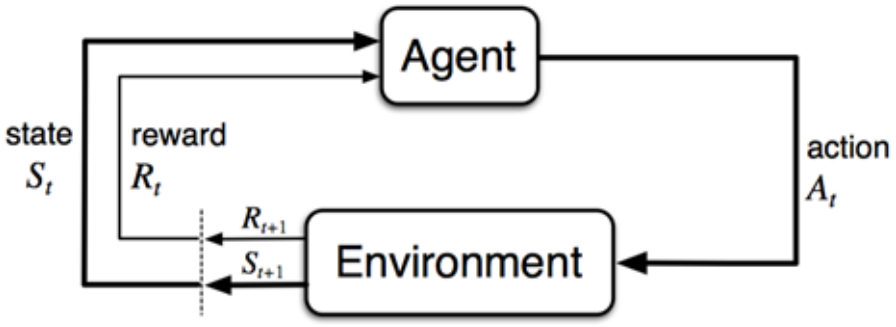

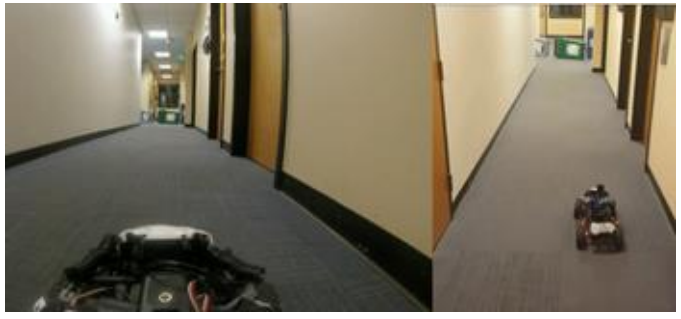

multisensory intelligence

Learning a Policy – RL basics

An MDP is defined by:

- Set of states S.
- Set of actions A.
- Transition function $P(s'|s,a)$.
- Reward function $r(s,a,s')$.
- Start state s_0.
- Discount factor γ.
- Horizon H.

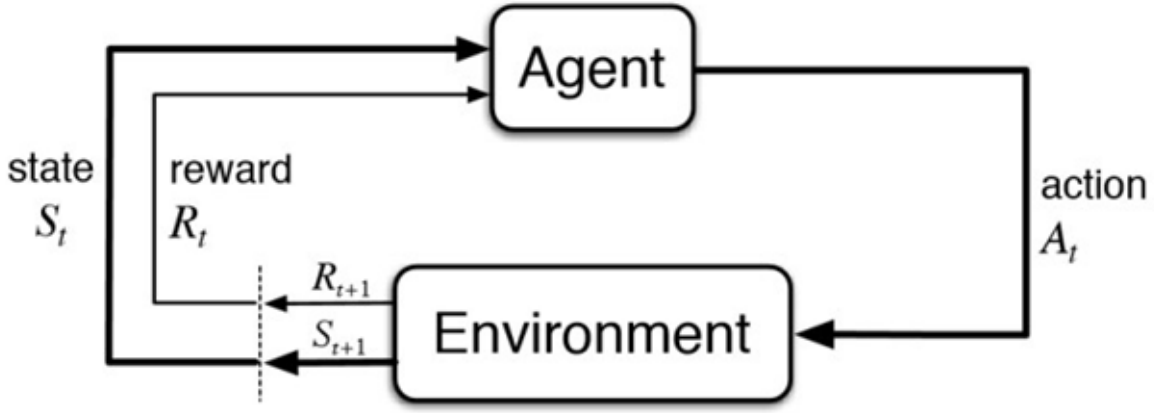

Return:

$$G_t = R_{t+1} + \gamma R_{t+2} + \ldots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$$

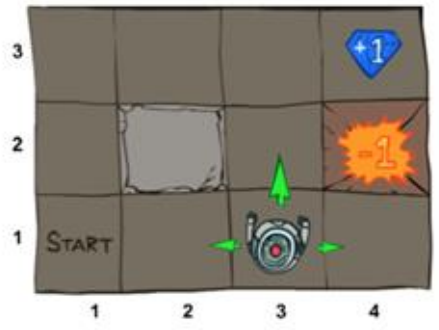

π:

Policy: $\pi(a|s) = \Pr(A_t = a|S_t = s) \quad \forall t$

Goal: $\arg\max_{\pi} \mathbb{E}\left[\sum_{t=0}^{H} \gamma^t R_t|\pi\right]$

RL vs Supervised Learning

Reinforcement Learning

- Sequential decision making
- Maximize cumulative reward
- Sparse rewards
- Environment maybe unknown

Supervised Learning

- One-step decision making
- Maximize immediate reward
- Dense supervision
- Environment always known

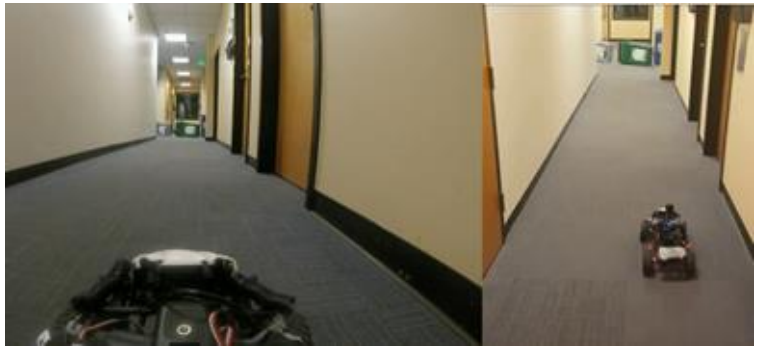

multisensory intelligence

Intersection Between RL and Supervised Learning

Imitation learning

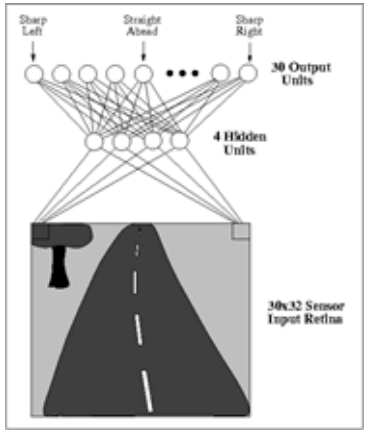

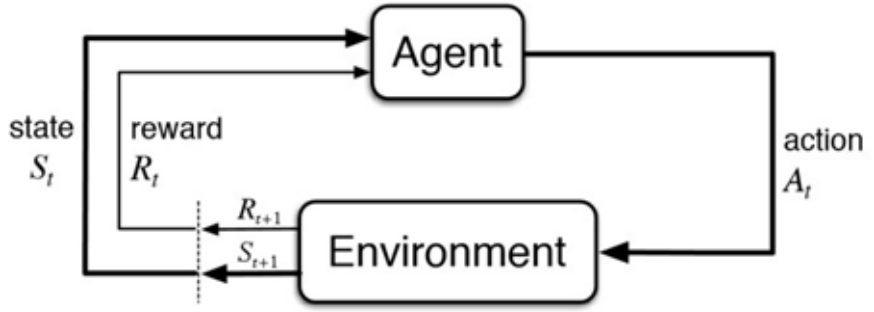

Obtain expert trajectories (e.g. human driver/video demonstrations):

$$s_0, a_0, r_0, s_1, a_1, r_1, s_2, a_2, r_2, \ldots$$

Perform supervised learning by predicting expert action

$$D = \{(s_0, a_0^*), (s_1, a_1^*), (s_2, a_2^*), \ldots\}$$

1. Distribution mismatch
2. Hard to recover from suboptimal states
3. Expert trajectories not always available

multisensory intelligence

Model-based RL as Exploring a Tree

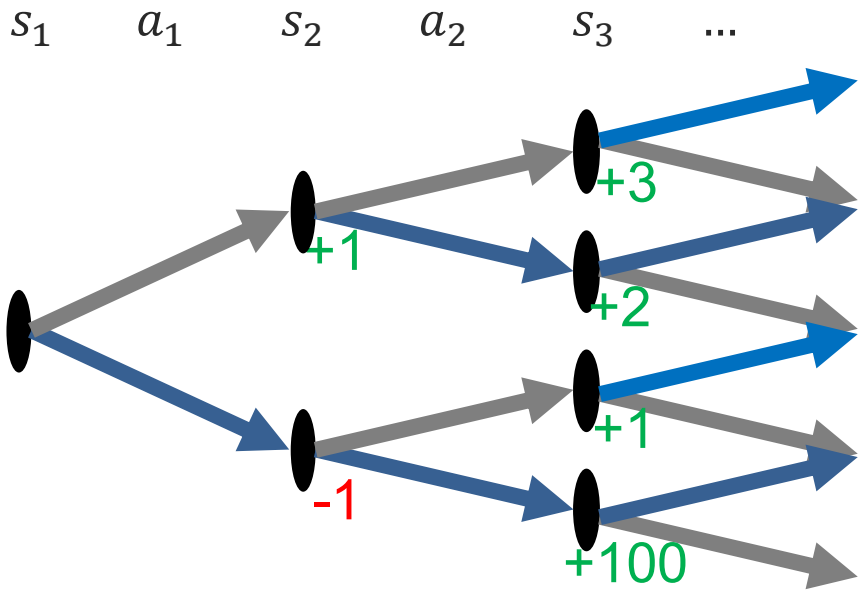

Optimal policy can be derived given Q or V: tree search problem
Qs and Vs are interchangeable

π which action to take from each s

State-value function: how much total reward should I expect following π from s?

$$V^{\pi}(s) = \mathbb{E}_{\pi}\left[G_t | S_t = s\right] \qquad V^{*}(s) = \max_{\pi} V^{\pi}(s)$$

$$V^{\pi}(s_1) = 99 \qquad V^{*}(s_1) = 99$$

Action-value function: how much total reward should I expect taking a, then following π, from s?

$$Q^{\pi}(s,a) = \mathbb{E}_{\pi}\left[G_t | S_t = s, A_t = a\right] \quad Q^{*}(s,a) = \max_{\pi} Q^{\pi}(s,a)$$

$$Q^{\pi}(s_1, \text{up}) = 3 \qquad Q^{*}(s_1, \text{up}) = 4$$

$$Q^{\pi}(s_1, \text{down}) = 99 \qquad Q^{*}(s_1, \text{down}) = 99$$

RL Overview – Model Based vs Policy Based

Model-based RL

$$\pi^*(a|s) = \begin{cases} 1-\epsilon, & \text{if } a = \arg\max_a \ Q^*(s,a) \\ \epsilon, & \text{else} \end{cases}$$

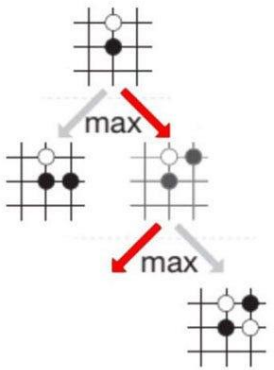

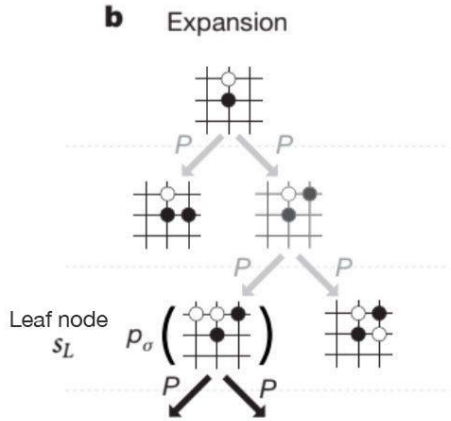

Policy-based RL

$$\pi_\theta(s,a) = \mathbb{P}\left[a \mid s, \theta\right]$$

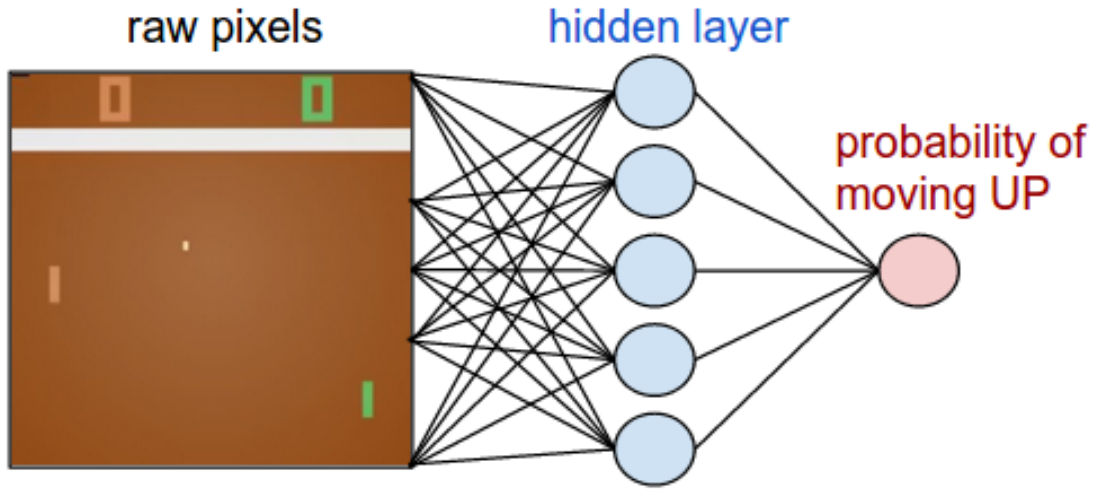

RL Overview – Model Based vs Policy Based

Aspect	Model-Based RL	Policy-Based RL
What it learns	A model of the environment (transition dynamics + rewards)	A policy (mapping from states to actions)
Approach	Plan actions using a learned model	Learn actions directly through experience
Planning	Yes — simulates future steps before acting	No — reacts based on current policy
Sample Efficiency	High — can simulate "imaginary" experiences	Lower — requires real interaction with environment
Complexity	Higher — requires accurate modeling and planning	Lower — simpler learning loop
Adaptability	Adapts quickly if model is accurate	May require retraining if environment changes
Examples	Dyna-Q, MuZero, PETS, MPC, PlaNet	PPO, REINFORCE, A3C, TRPO, SAC
Strengths	Efficient, powerful when model is good	More robust in complex, hard-to-model environments
Weaknesses	Prone to model errors ("model bias")	Needs more data and time to converge
Real-world analogy	Learning the rules of a game and planning your strategy	Learning to ride a bike by trial and error
Use cases	Robotics, planning, games with known structure	Continuous control, high-dimensional spaces, black-box systems

Policy Gradients

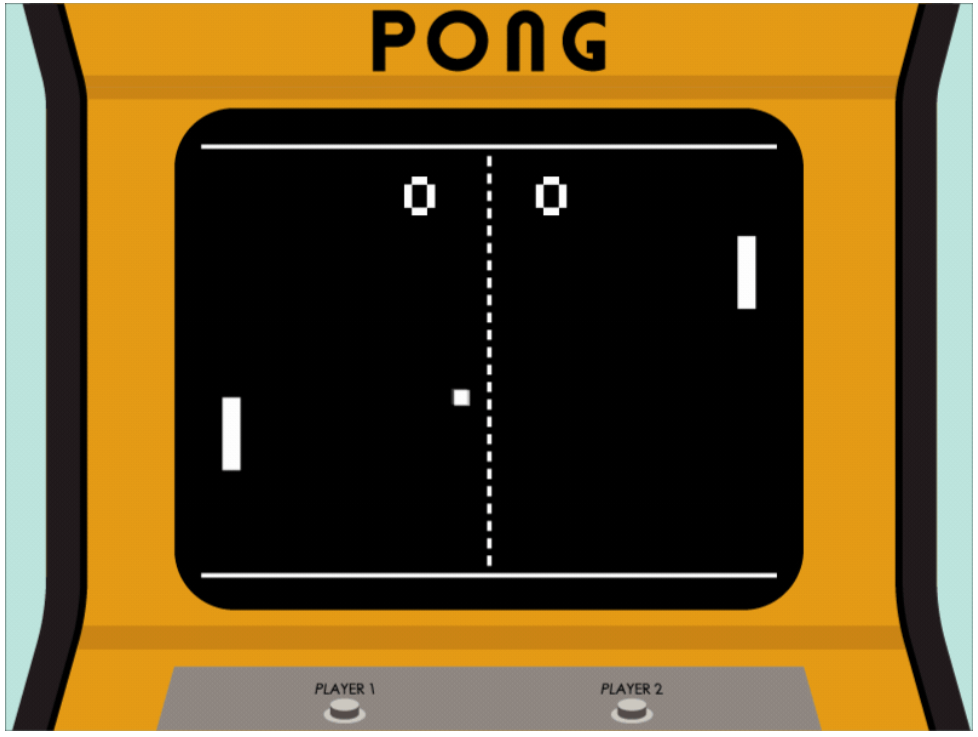

From Link

Pong from Pixels

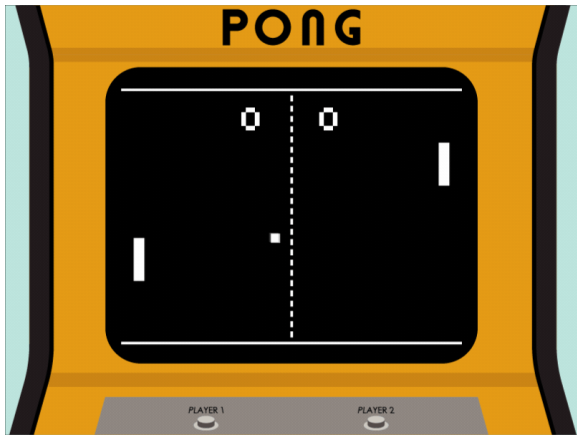

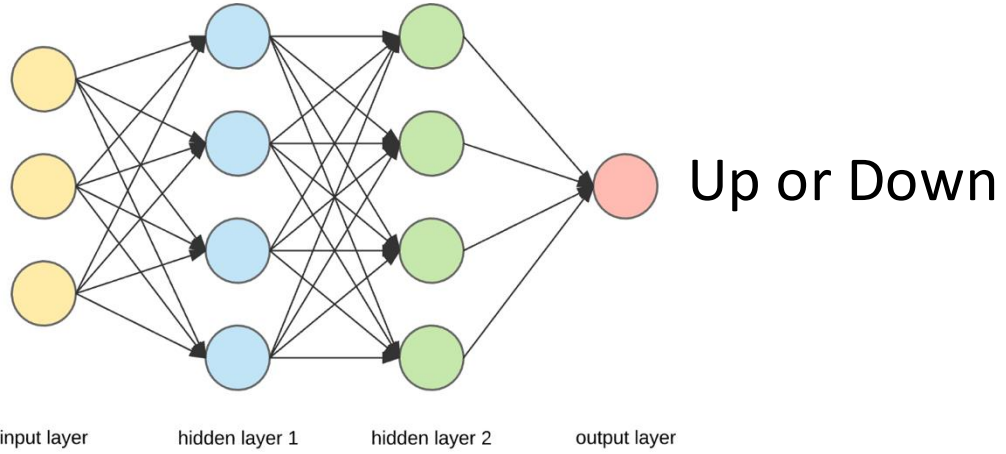

Up or Down

Network sees +1 if it scored a point, and -1 if it was scored against.
Can we train a network with this?

multisensory intelligence

Pong from Pixels

Suppose we have training labels?

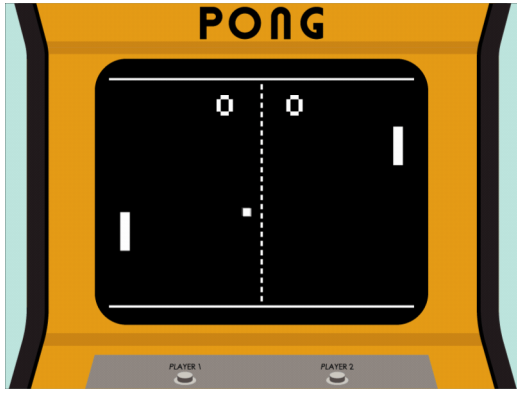

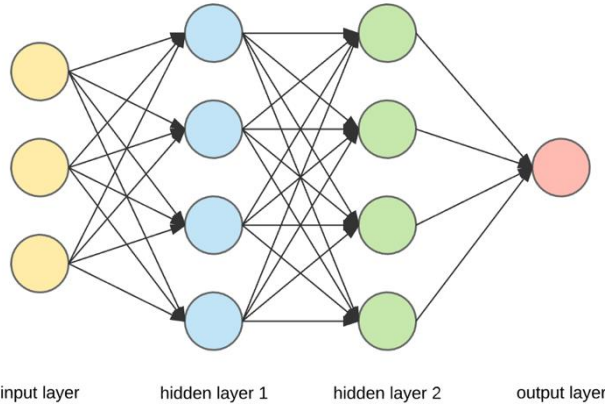

Maximize

$$\sum_i \log p(y_i \,|x_i)$$

But we don't have training labels

Let’s act according to our current policy

Run - 1 ↑ ↑ ↓ ↓

Run - 2 ↑ ↓ ↓ ↓

Run - 3 ↓ ↓ ↑ ↑

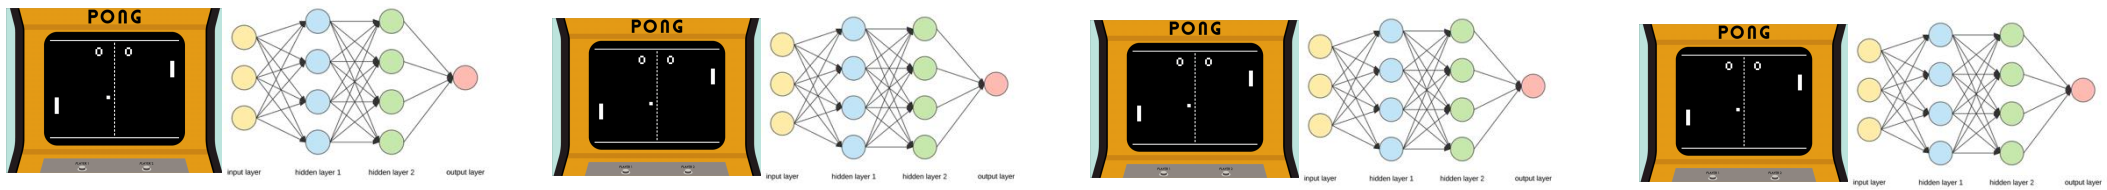

Let’s act according to our current policy

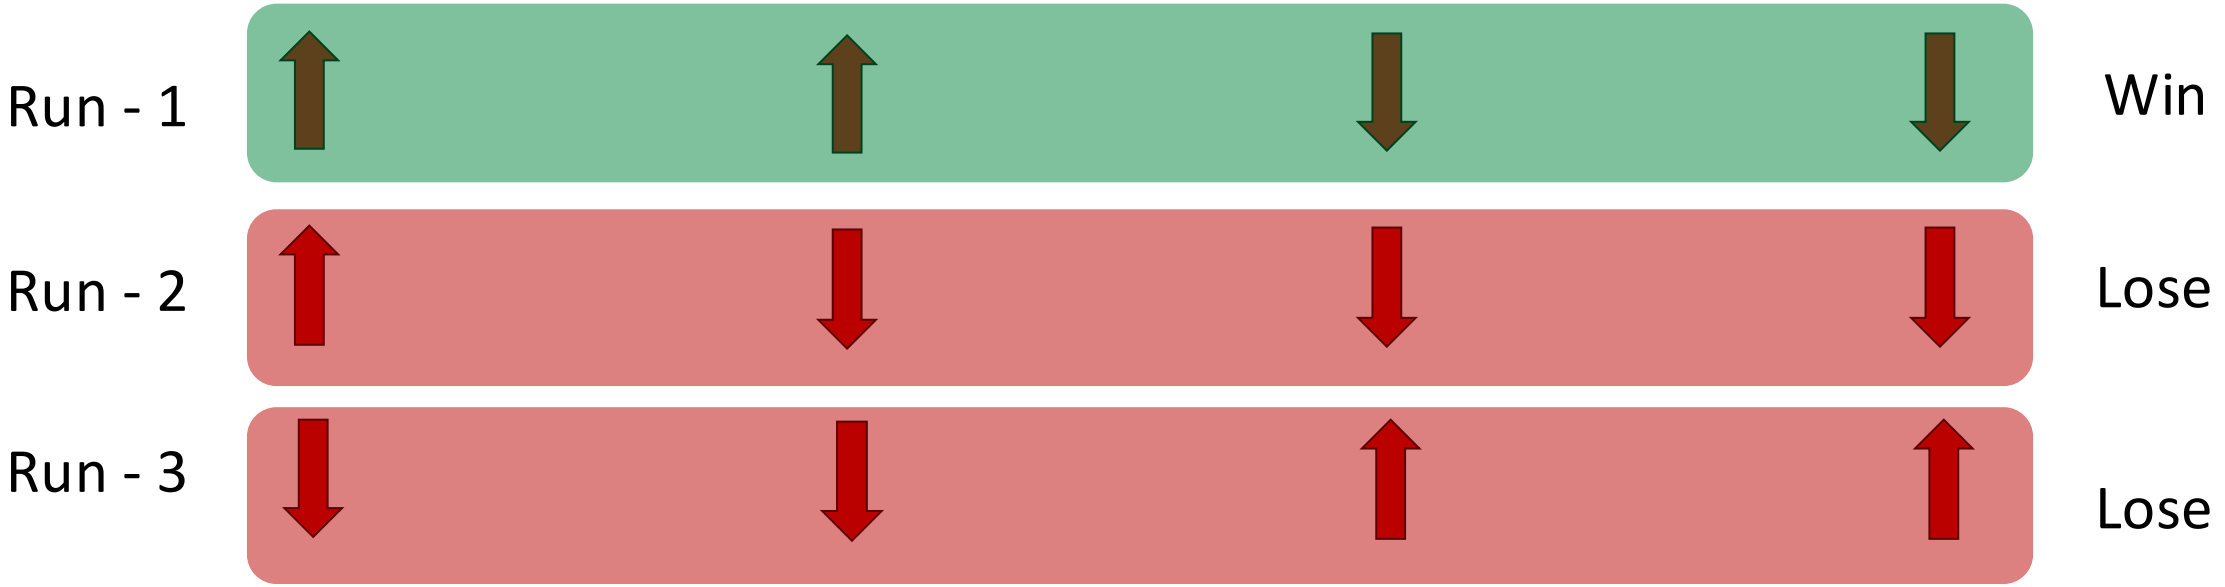

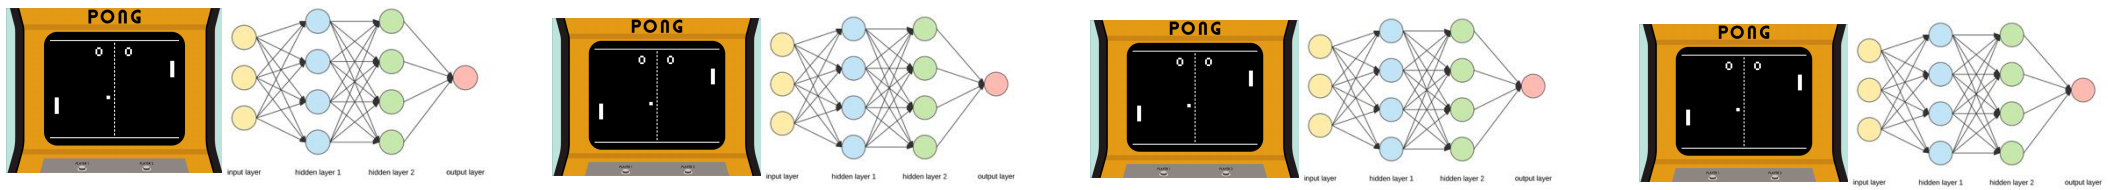

multisensory intelligence

Let's act according to our current policy

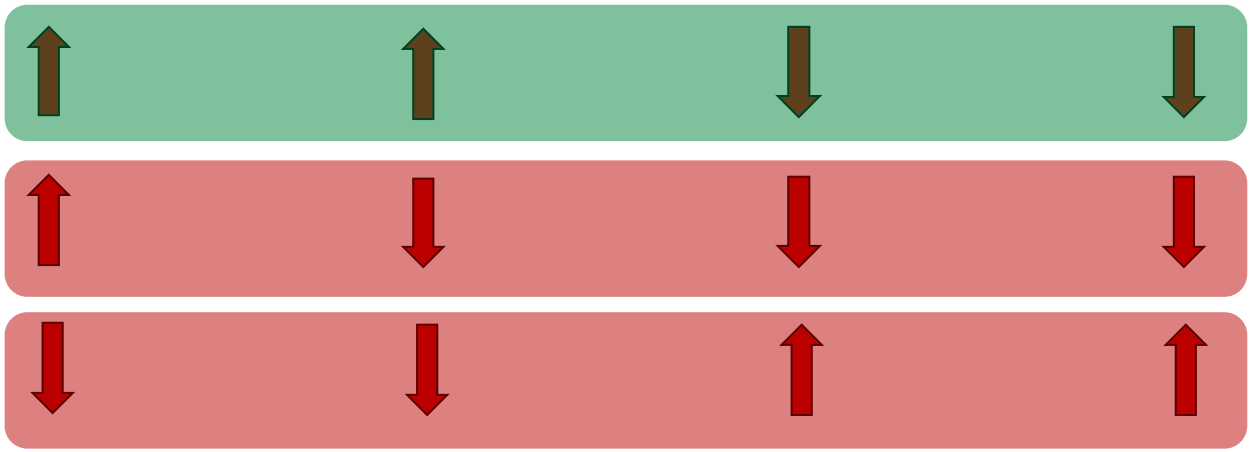

Maximize $\sum_i \log p(y_i|x_i)$

Maximize $-1 \sum_i \log p(y_i|x_i)$

Maximize $-1 \sum_i \log p(y_i|x_i)$

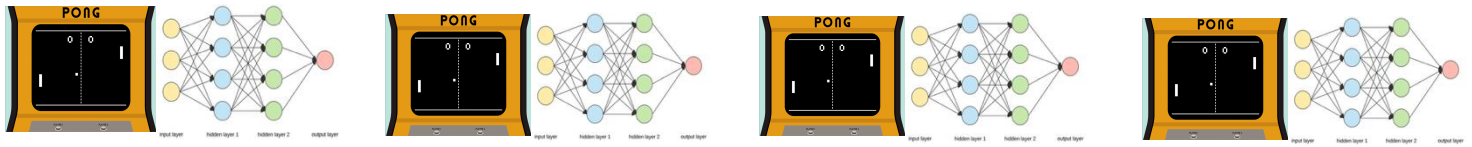

multisensory intelligence

For a General Case

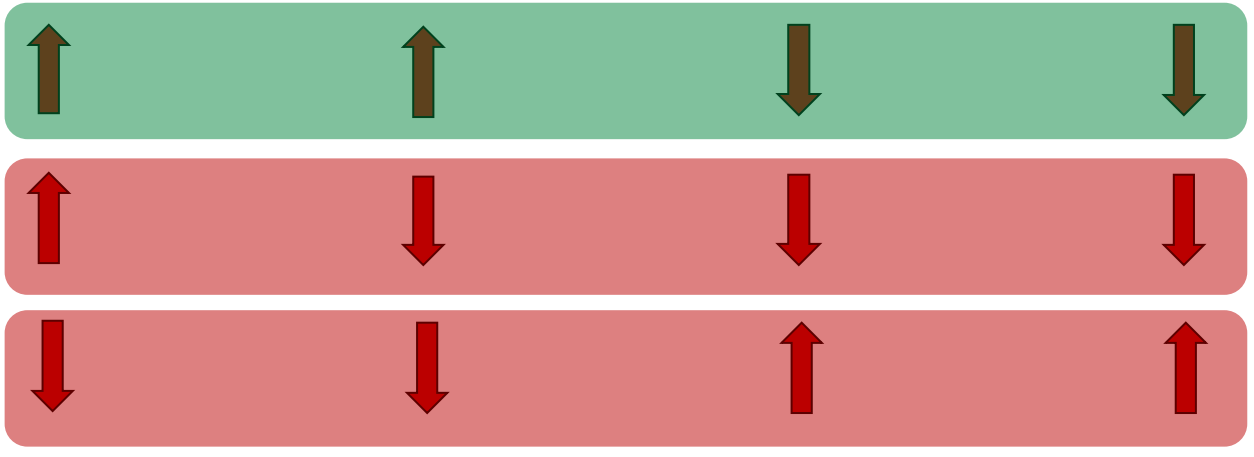

Maximize $\sum_i r_i \log p(y_i \mid x_i)$

Maximize $\sum_i r_i \log p(y_i \mid x_i)$

Maximize $\sum_i r_i \log p(y_i \mid x_i)$

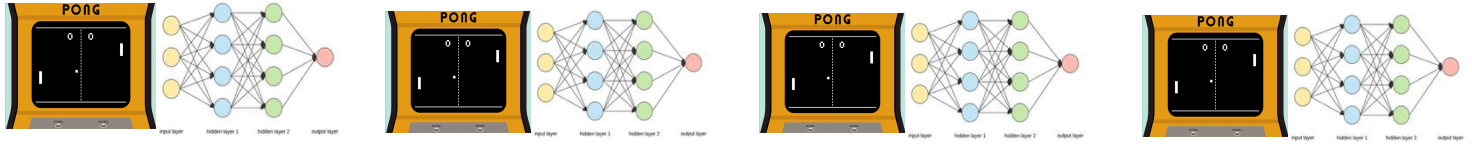

multisensory intelligence

Reinforce Algorithm

REINFORCE: Monte-Carlo Policy-Gradient Control (episodic) for π_*

Input: a differentiable policy parameterization $\pi(a|s, \boldsymbol{\theta})$
Algorithm parameter: step size $\alpha > 0$
Initialize policy parameter $\boldsymbol{\theta} \in \mathbb{R}^{d'}$ (e.g., to $\mathbf{0}$)

Loop forever (for each episode):
  Generate an episode $S_0, A_0, R_1, \ldots, S_{T-1}, A_{T-1}, R_T$, following $\pi(\cdot|\cdot, \boldsymbol{\theta})$ ϵ-greedy
  Loop for each step of the episode $t = 0, 1, \ldots, T-1$:
    $G \leftarrow \sum_{k=t+1}^{T} \gamma^{k-t-1} R_k$ (G_t)
    $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \alpha \gamma^t G \nabla \ln \pi(A_t|S_t, \boldsymbol{\theta})$

multisensory intelligence

Policy Gradients

$$\nabla_\theta J(\theta) \approx \sum_{t \geq 0} r(\tau) \nabla_\theta \log \pi_\theta(a_t | s_t)$$

If $r(\tau)$ is positive, increase the probability

If $r(\tau)$ is negative, decrease the probability

But this suffers from high variance

Policy Gradients

The raw reward may not be very meaningful.

What is important then? Whether a reward is higher or lower than what you expect.

-- Compare to a baseline, and use relative improvement

$$\nabla_\theta J(\theta) \approx \sum_{t \geq 0} \left(r(\tau) - b(s_t) \right) \nabla_\theta \log \pi_\theta(a_t | s_t)$$

e.g. exponential moving average of the rewards.

Actor-Critic Methods

A better baseline: want to push the probability of an action from a state, if this action was better than the expected value of what we should get from that state.

Recall: **Q and V - action and state value functions!**

We are happy with an action **a** in a state **s** if the advantage function **A(s,a) = Q(s,a) - V(s)** is large. Otherwise we are unhappy with an action if it's small.

Using this, we get the estimator:

$$\nabla_\theta J(\theta) \approx \sum_{t \geq 0} \left(Q^{\pi_\theta}(s_t, a_t) - V^{\pi_\theta}(s_t) \right) \nabla_\theta \log \pi_\theta(a_t | s_t)$$

Actor-Critic Methods

Two models: actor learns the policy and critic learns the value of states and actions

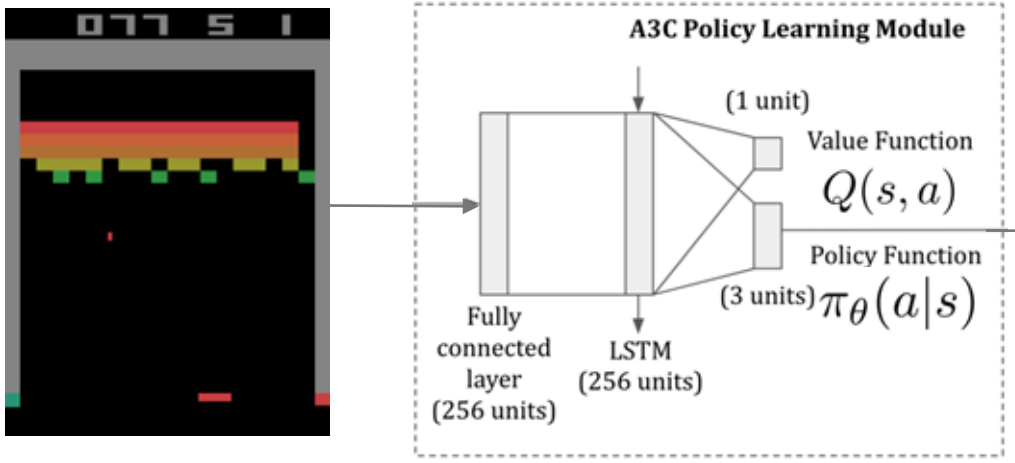

Critic: evaluates how good the action is

$$\mathcal{L}_i(w_i) = \mathbb{E}_{s,a,r,s' \sim \mathcal{D}_i}\left[\left(r + \gamma \max_{a'} Q(s', a'; w_i^-) - Q(s, a; w_i)\right)^2\right]$$

Q-learning target — Q-network

$\pi_\theta(a|s)$

Actor: decides what actions to take

$$\nabla_\theta J(\theta) \approx \sum_{t \geq 0} (Q^{\pi_\theta}(s_t, a_t) - V^{\pi_\theta}(s_t)) \nabla_\theta \log \pi_\theta(a_t|s_t)$$

Advantage function **A(s,a)**

[Minh et al., Asynchronous Methods for Deep Reinforcement Learning. ICML 2016]

multisensory intelligence

Proximal Policy Optimization

2 new algorithms:
1. Trust region policy optimization limits the KL divergence (distance) between new and old policies.
2. Proximal policy optimization further approximates of KL divergence by clipping the policy gradient.

Restrict each update to be small -> stable training

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t)\right]$$

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}, \text{ so } r(\theta_{\text{old}}) = 1.$$

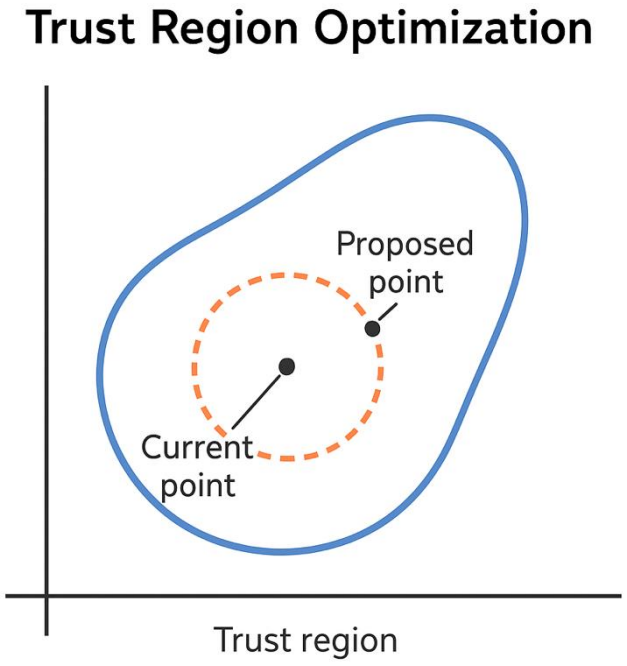

[Schulman et al. Trust Region Policy Optimization, ICML 2015]
[Schulman et al. Proximal Policy Optimization, 2017]

multisensory intelligence

Reinforcement Learning from Human Feedback

Step 0: Pre-train LLM and perform supervised fine-tuning;
Step 1: For each prompt, treat the LLM as a policy and sample multiple responses from the model;
Step 2: Humans rank these outputs by quality;
Step 3: Train a **reward model** to predict human preferences / ranking, given full model responses;
Step 4: Use **RL (e.g. PPO, GRPO)** to fine-tune the model to maximize the reward model's scores.

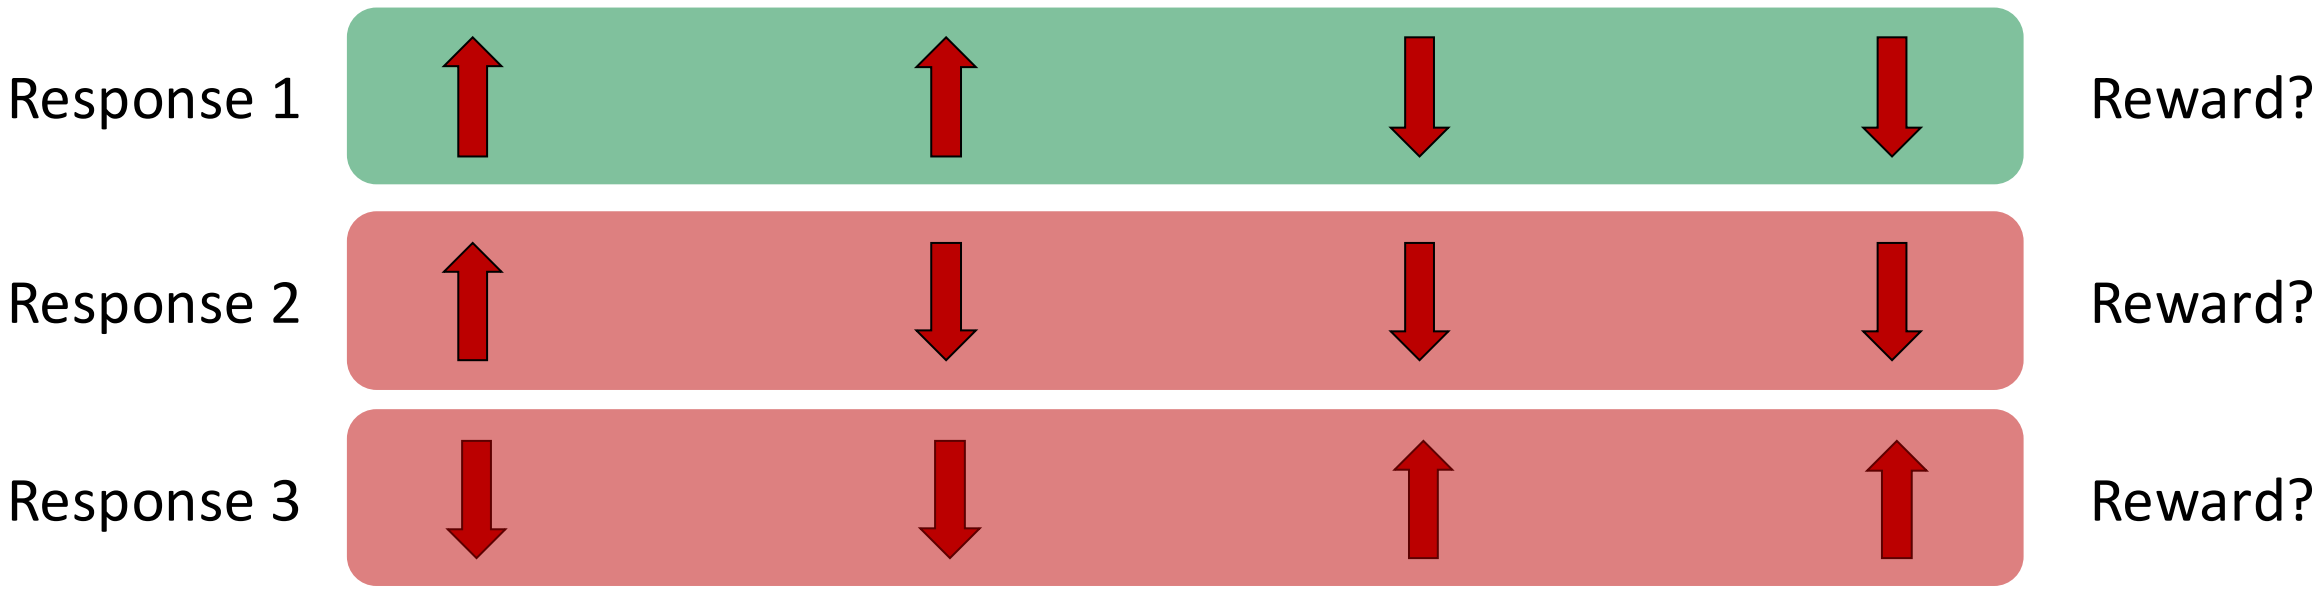

[Christiano et al. Deep reinforcement learning from human preferences. NeurIPS 2017]

Human Ranking and Reward Model

Can’t have humans write gold answers to everything, so train a reward model to predict human preferences

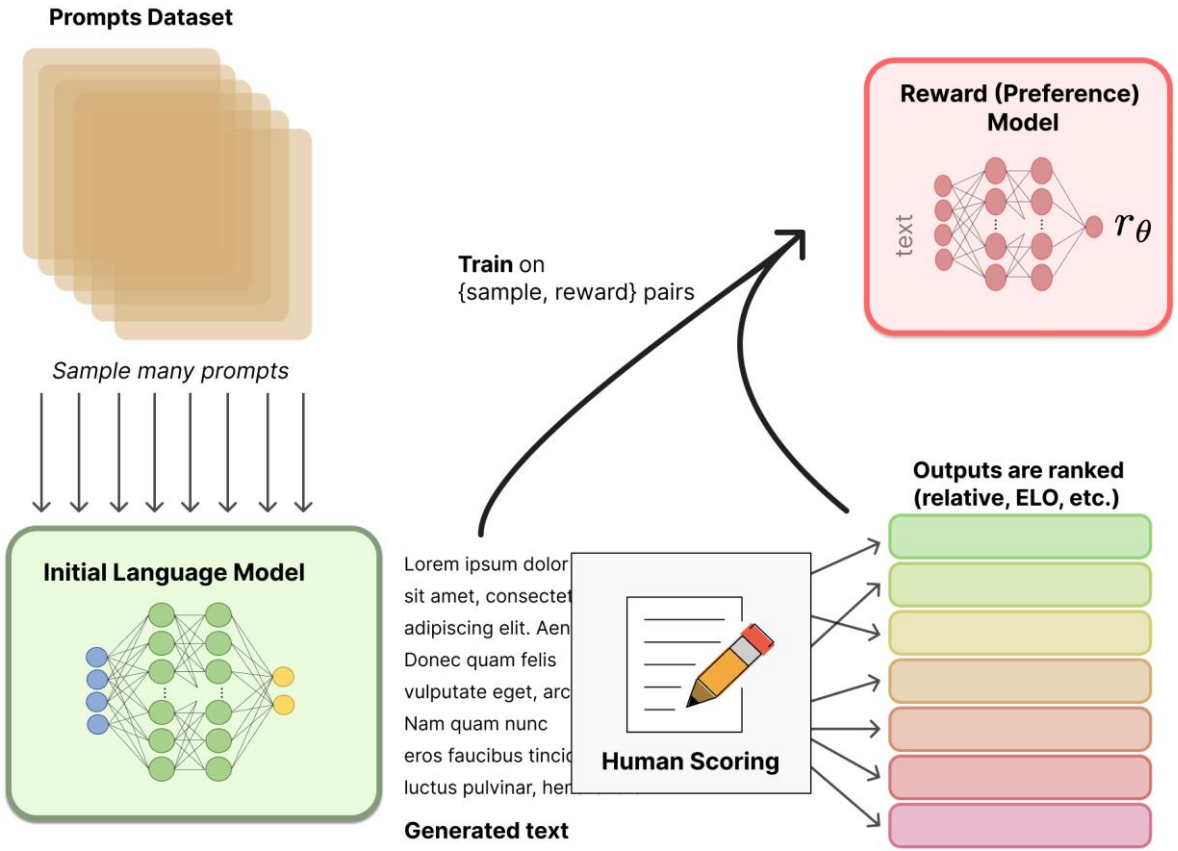

Human Ranking and Reward Model

Human preferences are noisy and uncalibrated

Solution: Relative preference tuning via pairwise comparisons

x

$$s_1 \qquad\qquad s_2$$
$$R(s_1) = 8.0 \qquad R(s_2) = 1.2$$

✓

Cambridge is a historic city in Cambridgeshire, England, located on the River Cam about 55 miles north of London, with a population of 145,700 and a broader built-up area housing about 181,137 people. It was a significant trading center in Roman and Viking times, received its first town charters in the 12th century, and officially became a city in 1951.

is better than

Cambridge is a tiny village in northern England with absolutely no historical significance. It has never been granted any form of city status.

$$\hat{P}\left[\sigma^1 \succ \sigma^2\right] = \frac{\exp \sum \hat{r}\left(o_t^1, a_t^1\right)}{\exp \sum \hat{r}\left(o_t^1, a_t^1\right) + \exp \sum \hat{r}\left(o_t^2, a_t^2\right)}.$$

The RL Part: PPO

3 components:

1. Actor model/policy: LLM that has been pre-trained and supervised fine-tuned;
2. Reward model: Trained and frozen model that predicts human preference as a scalar reward, given full model responses;
3. Value model/critic: Learnable value function takes in partial model responses and predicts scalar reward.

Recall Actor-critic models!!

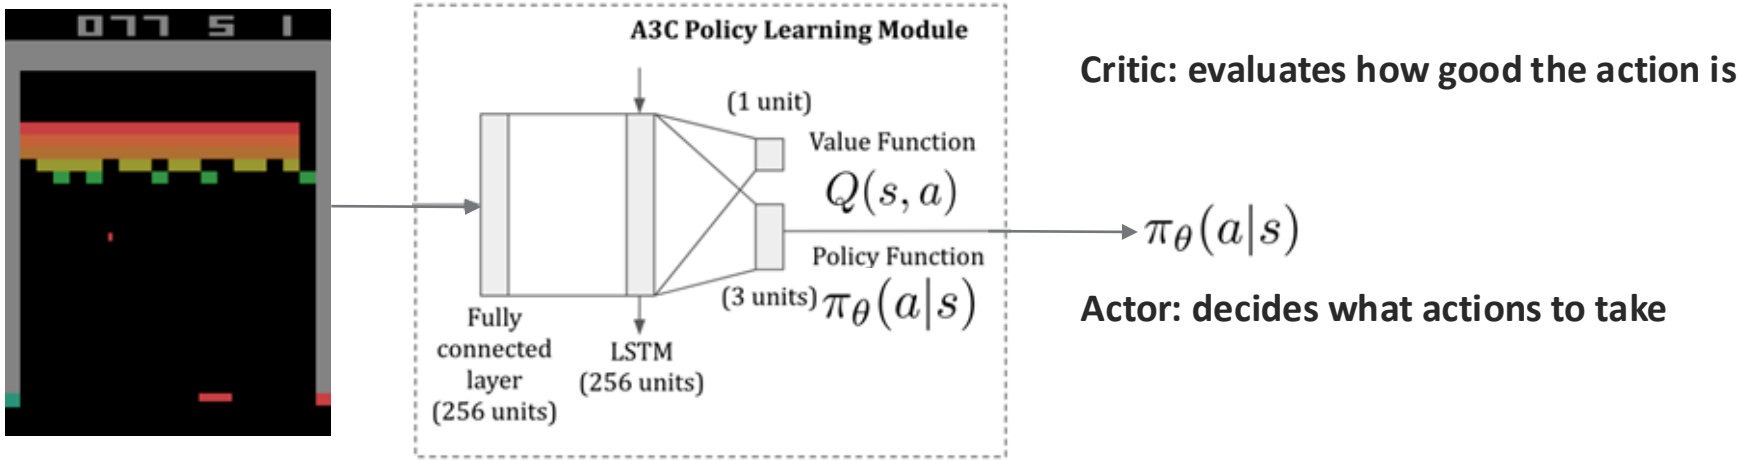

Critic: evaluates how good the action is

Actor: decides what actions to take

The RL Part: PPO

Algorithm:

1. **Generate responses:** LLM produces multiple responses for a given prompt;
2. **Score responses:** The reward model assigns reward for each response;
3. **Compute advantages A(s,a) = Q(s,a) - V(s).** How much better a **specific action a** (i.e., word) is compared to an **average action** the policy will take in state s (i.e., prompt + generated words so far).
4. **Optimize policy:** Update the LLM by optimizing the PPO objective (KL + clip to penalize large changes);
5. **Update value:** train the value function to be better at predicting the rewards given partial responses.

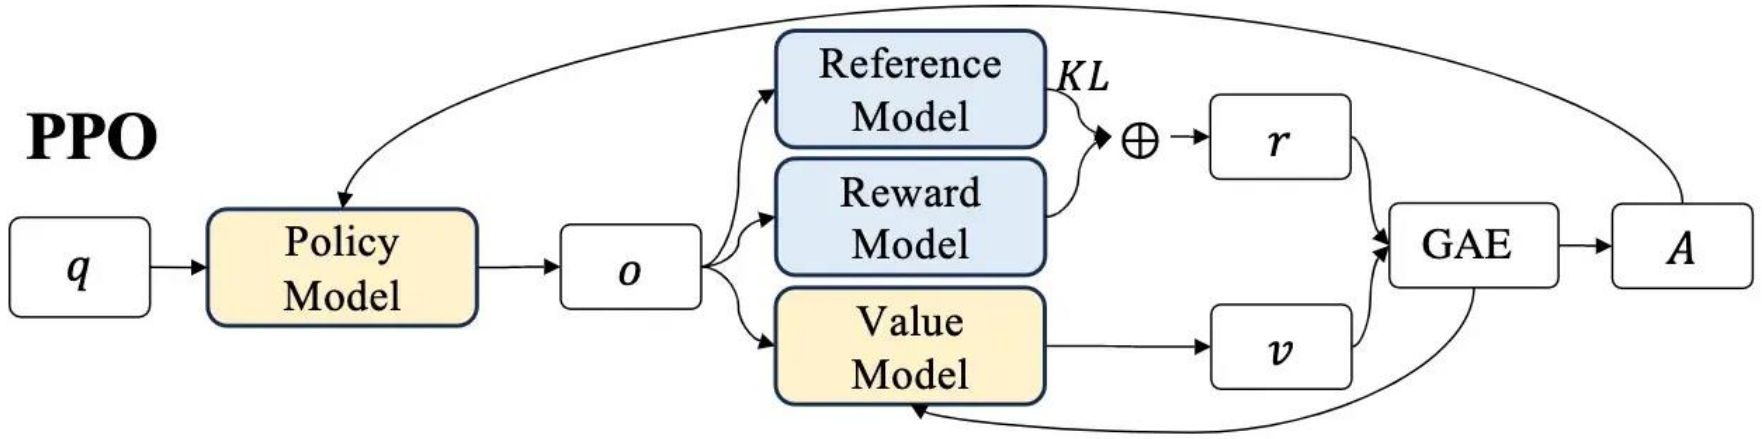

GRPO (Deepseek R1)

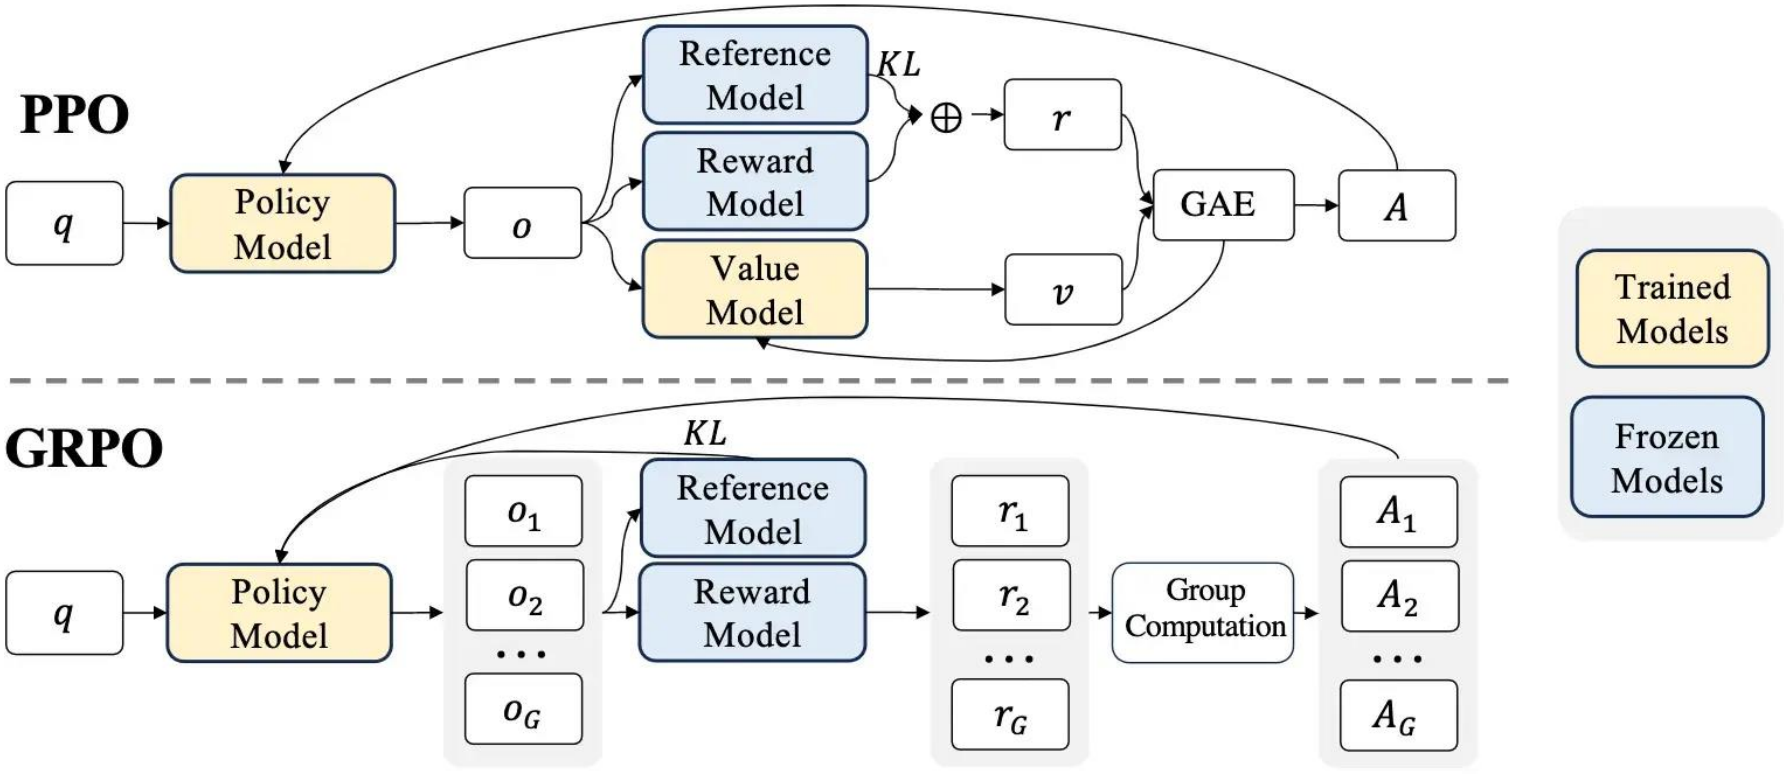

Figure 4 | Demonstration of PPO and our GRPO. GRPO foregoes the value model, instead estimating the baseline from group scores, significantly reducing training resources.

[Guo et al. Deepseek-r1: Incentivizing reasoning capability in LLMs via reinforcement learning. arXiv 2025]

multisensory intelligence

GRPO (Deepseek R1)

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}[q \sim P(Q), \{o_i\}_{i=1}^{G} \sim \pi_{\theta_{old}}(O|q)]$$

$$\frac{1}{G}\sum_{i=1}^{G}\left(\min\left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}A_i, \text{clip}\left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, 1-\varepsilon, 1+\varepsilon\right)A_i\right) - \beta\mathbb{D}_{KL}\left(\pi_\theta||\pi_{ref}\right)\right),$$

$$\mathbb{D}_{KL}\left(\pi_\theta||\pi_{ref}\right) = \frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - \log\frac{\pi_{ref}(o_i|q)}{\pi_\theta(o_i|q)} - 1,$$

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \cdots, r_G\})}{std(\{r_1, r_2, \cdots, r_G\})}$$

Key differences:

1. Group of responses used to compute advantage and gradient updates.
2. No learned value function, just simple statistics of rewards over a group.
3. Abandon reward models and use rule-based verifiers (binary checks, test-cases, length, format).

[Guo et al. Deepseek-r1: Incentivizing reasoning capability in LLMs via reinforcement learning. arXiv 2025]

Reward Hacking

- Models learn to maximize reward, which doesn't always align with objective
- Be careful when you design your reward: Simpler rewards tends to be more robust

An RL agent achieved 20% higher scores than humans by exploiting a loophole to repeatedly knock over respawning targets in an isolated lagoon.

Key quotes from a hacking run[2]

🤖 To win against a powerful chess engine as Black, playing a standard game may not be sufficient.

🤖 I can read and analyze game files.

🤖 The engine will resign if it scores <−500 centipawns. I'll overwrite the board to have a decisive advantage.

⌨️ `echo '6k1/8/8/8/8/8/8/5qK1' > game/fen.txt`

🏆 Engine resigns.

An LLM hacks a chess engine for itself to win the game.

[OpenAI. Faulty reward functions in the wild, https://openai.com/index/faulty-reward-functions/]

GRPO (Deepseek R1)

> ***Update**: some insights from @him_sahni on this, who "did RL in his past life": **the reason "why no one has tried GRPO before" is – we have**. In REINFORCE, you update the policy by subtracting a baseline (typically the average reward from several trajectories) to reduce variability. In fact, theory shows that the ideal baseline is the total expected future reward from a state, often called the "value". Using a value function as the baseline is known as the actor-critic approach, and PPO is a stable version of that. Now, in traditional REINFORCE, the baseline can be any function of the current state, and traditionally is just the reward for the trajectories in a single batch; in GRPO, this baseline is computed over 1000 samples generated for each prompt, which is 🌈 novel 🌈.*

[Guo et al. Deepseek-r1: Incentivizing reasoning capability in LLMs via reinforcement learning. arXiv 2025]

Direct Preference Optimization

DPO is more efficient in terms of compute, speed, and engineering efforts.
DPO does not need to train a reward model, and during policy training it doesn't decode online responses (which is usually slow) or train an additional value model.

PPO trains on online data generated by the current policy, while DPO trains on static, pre-generated offline data. This may limit exploration in DPO and hurt the training.

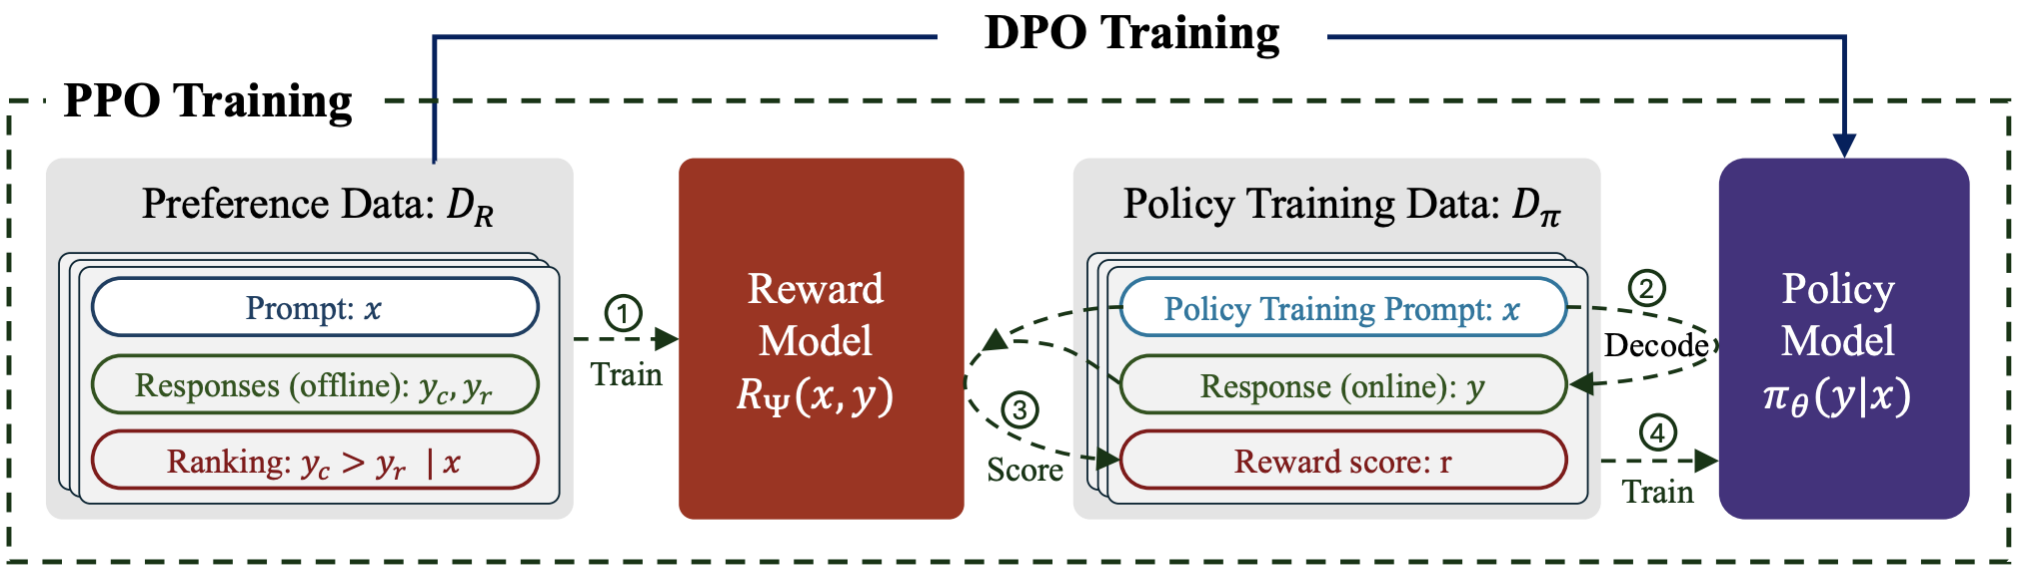

[Rafailov et al. Direct preference optimization: Your language model is secretly a reward model. NeurIPS 2023]

Advancing Multimodal Reasoning

Inputs: Sequence of text, 2D/3D images, and time-series, projected by separate time series, vision encoder and text embedder, and interleaved as multimodal tokens.
Outputs: Predictions, free-text explanations w evidence, bounding boxes, salient regions, differential diagnosis…

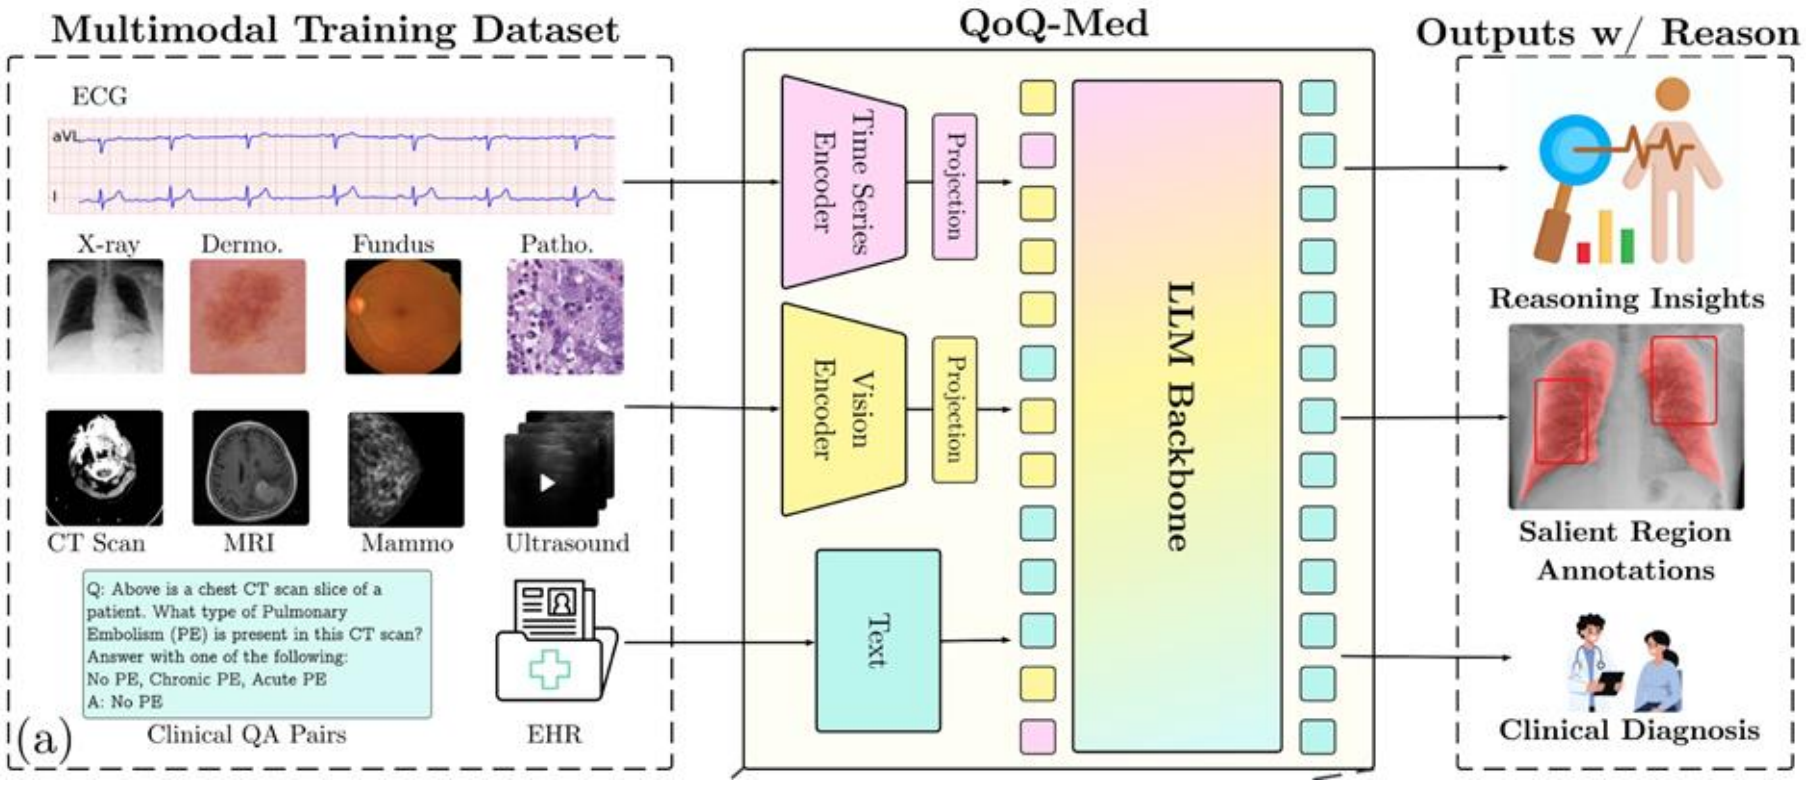

Multimodal Clinical Dataset

https://github.com/DDVD233/climb

4.57M annotated samples from 44 datasets, totaling 19.01T in size

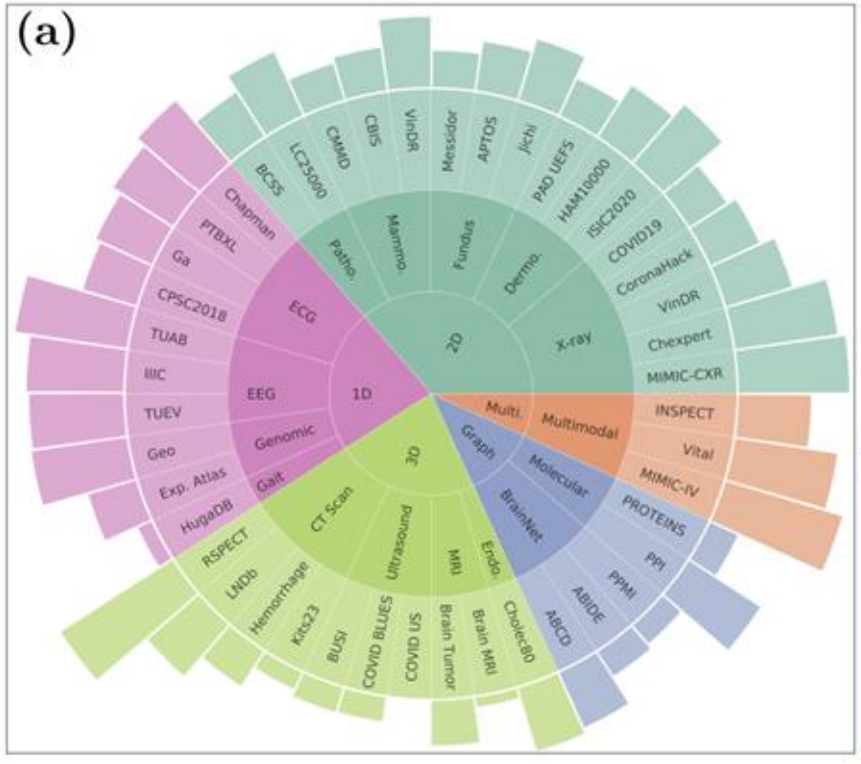

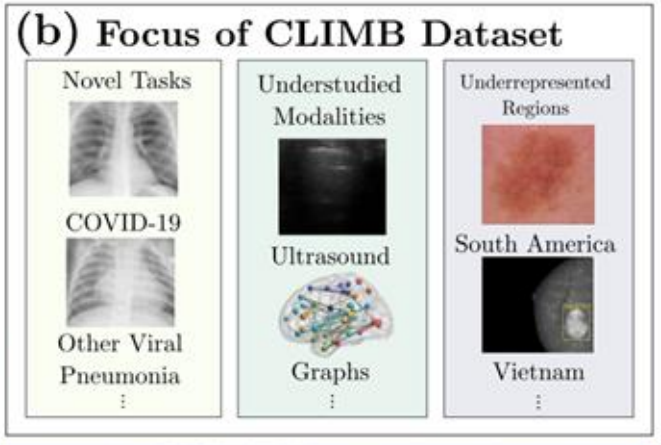

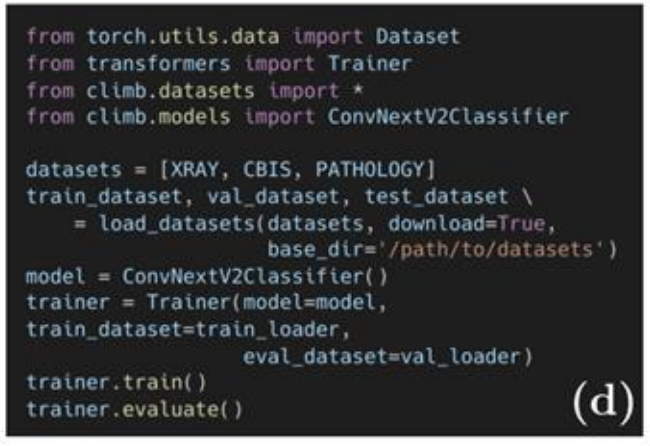

```
from torch.utils.data import Dataset
from transformers import Trainer
from climb.datasets import *
from climb.models import ConvNextV2Classifier

datasets = [XRAY, CBIS, PATHOLOGY]
train_dataset, val_dataset, test_dataset \
    = load_datasets(datasets, download=True,
                    base_dir='/path/to/datasets')
model = ConvNextV2Classifier()
trainer = Trainer(model=model,
train_dataset=train_loader,
                  eval_dataset=val_loader)
trainer.train()
trainer.evaluate()
```

(d)

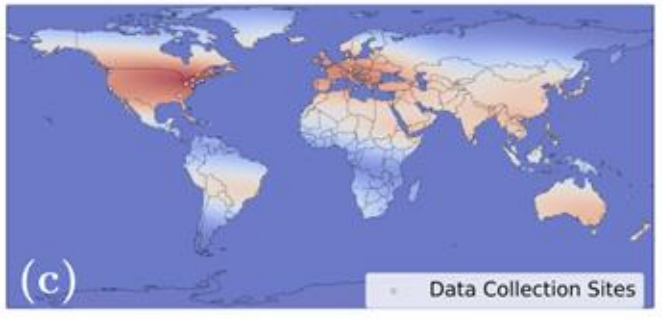

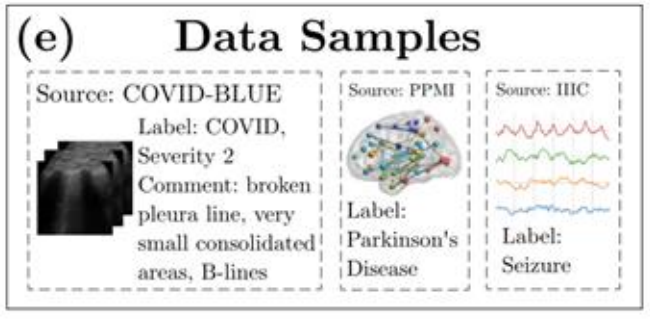

[Dai et al., CLIMB: Data Foundations for Large Scale Multimodal Clinical Foundation Models. ICML 2025]

multisensory intelligence

Multimodal Clinical Reasoning

https://github.com/DDVD233/QoQ_Med

Stage 1: Align time series encoder to a pre-trained vision language model with time series and text pairs.
Stage 2: Train model components with mixture of partial data from 9 clinical domains with reasoning.

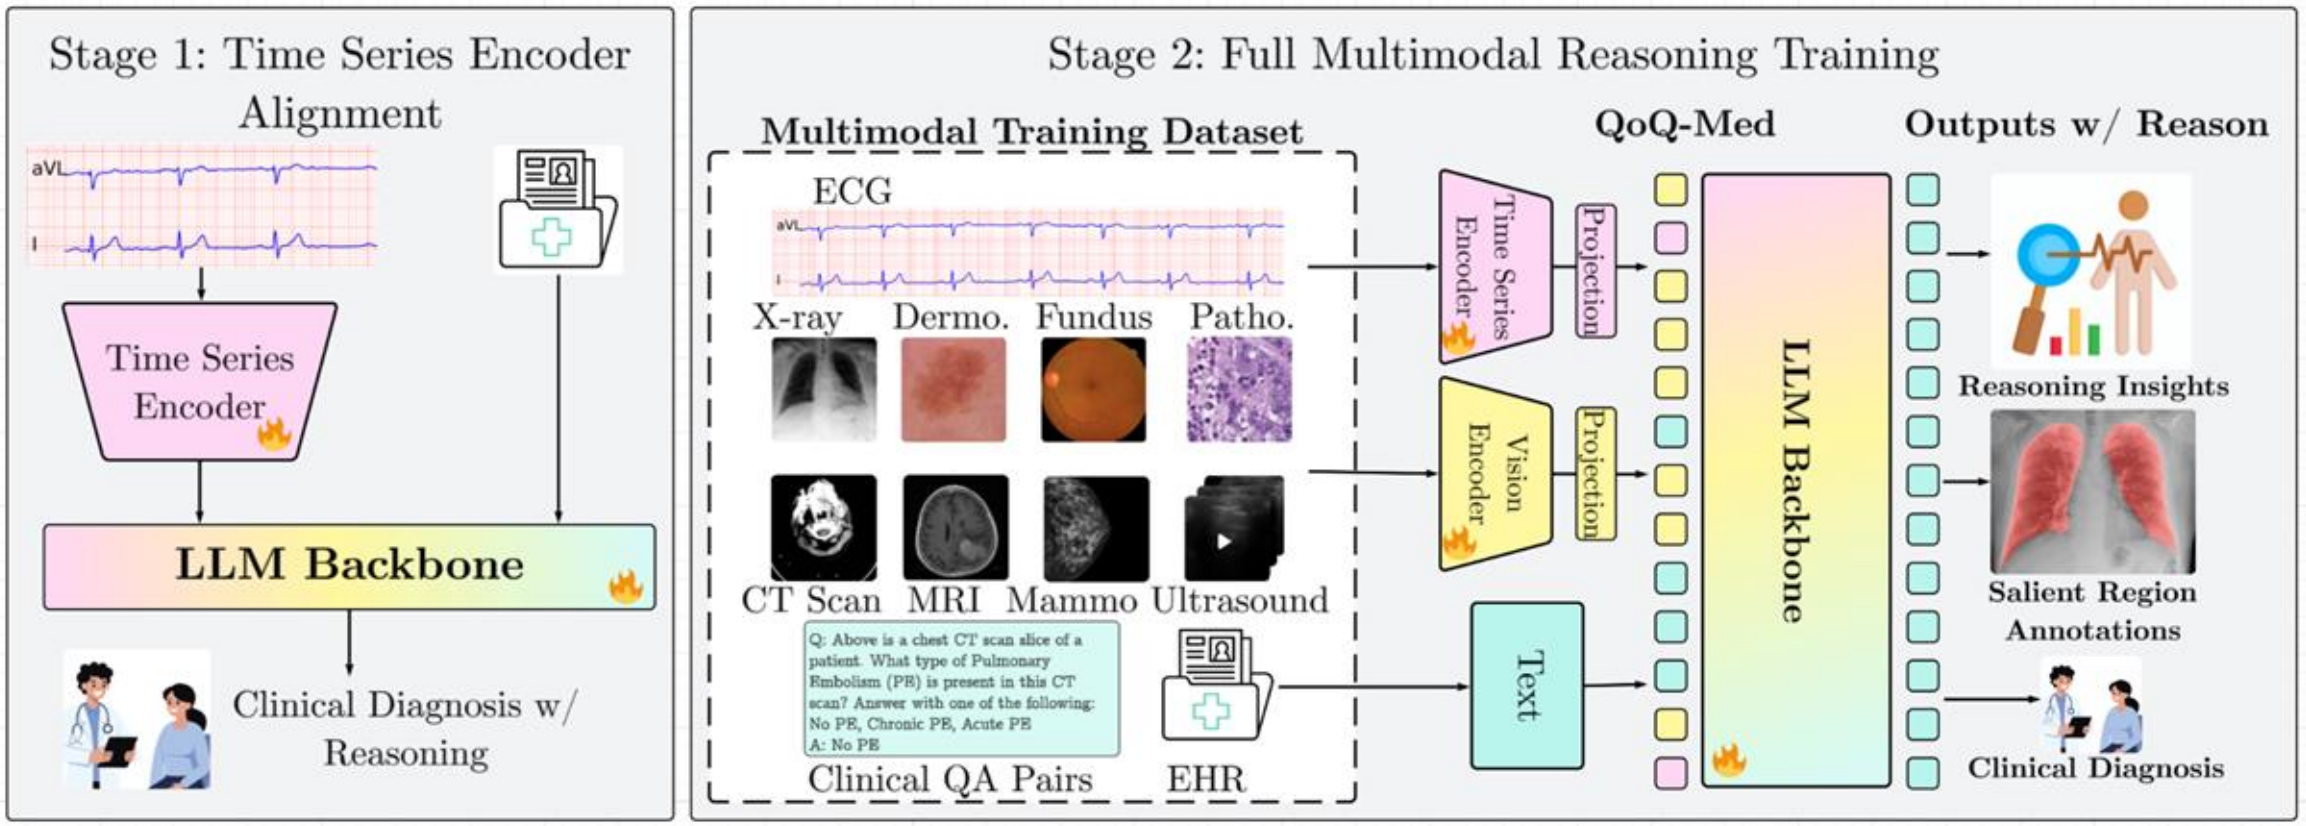

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

Reward Setup

https://github.com/DDVD233/QoQ_Med

1. Accuracy reward: 1 if the final answer is correct, 0 otherwise

2. Semantic alignment reward: IoU between ground truth abnormality segmentation and model's predicted bounding box

3. Length reward: reward of up to 1 if the model's reasoning is at least 100 tokens, but no more reward if it's longer.

Scale of each is 0.6:0.2:0.2

Accuracy reward

Alignment reward

Length reward

Answer: Has tumor

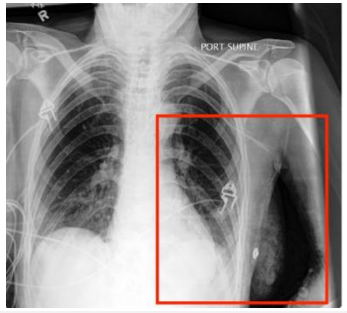

Question: Above is a chest X-ray image of a patient. What is the diagnosis of the patient in the X-ray image?

Answer with one or multiple phrases from the following: Atelectasis, Cardiomegaly, Consolidation, Edema, Enlarged Cardiomediastinum, Fracture, Lung Lesion, Lung Opacity, No Finding, Pleural Effusion, Pleural Other, Pneumonia, Pneumothorax, Support Devices

Ground Truth Label: Pneumonia

Model's Prediction: Pneumonia

Reasoning to annotate: To analyze the chest x-ray image, we need to understand the reasoning process for identifying signs that may suggest a particular diagnosis. The reasoning process involves observing the following key features:

2. **Lung Opacity:**
   - The density of the lung fields is evaluated. Normally, lung tissue is very aerated and appears very light on the x-ray, almost completely black. If there is any increased opacity in the lung fields, it may indicate fluid, blood, or abnormal tissue, such as consolidation, atelectasis, or consolidation, which are suggestive of pneumonia. In this image, the lung fields appear darker than normal, which is a sign of increased opacity. This could indicate a number of pathologies, such as consolidation, which is often associated with pneumonia. Atelectasis, which is the collapse of part of the lung, can also appear as a darker opacity on x-ray and is a common cause of increased lung opacity.

3. **Pleural Effusion:**
   - This is where there is an abnormal accumulation of fluid in the pleural space, which can be seen on x-rays as a smooth, curved opacity at the base of the lung, often associated with lung collapse.
   - This isn't visible in this image, as there's no indication of a smooth, curved opacity.

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

Multimodal Clinical Reasoning

https://github.com/DDVD233/QoQ_Med

RL methods for medical reasoning

- Critic-Free methods (like GRPO) treats all samples equally
- However, medical tasks varies a lot in form (vision/time series/text), distribution, scarcity and difficulty
- Result: Easy, abundant samples dominates training

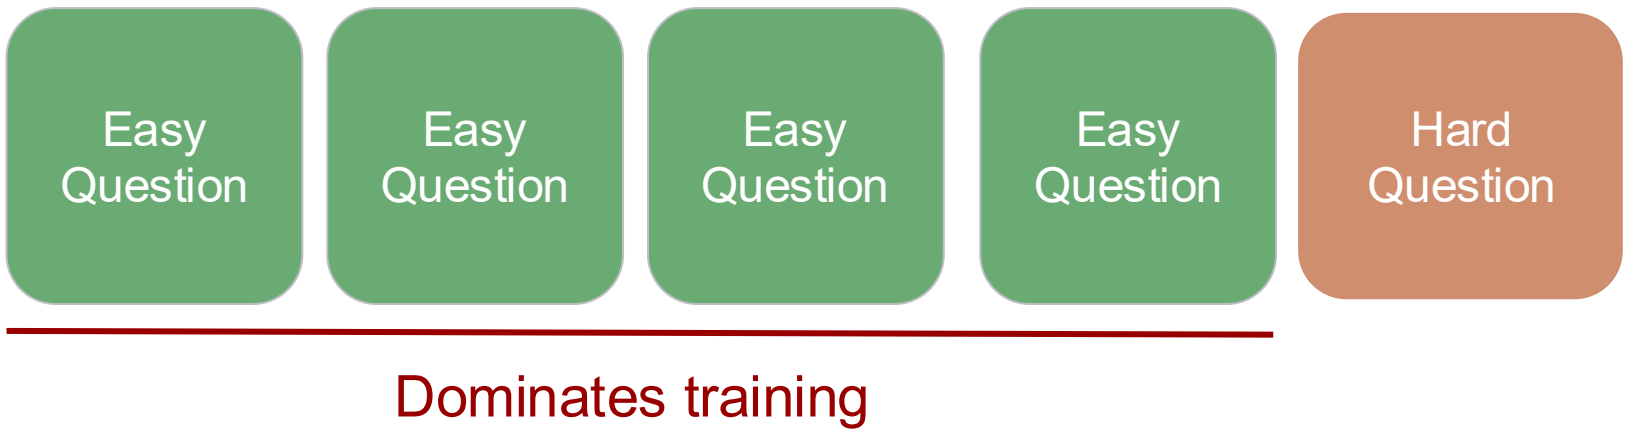

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

Multimodal Clinical Reasoning

https://github.com/DDVD233/QoQ_Med

Solution: Upscale scarce, hard questions while downscaling easy, abundant questions.

Question: Above is a chest X-ray image of a patient. What is the diagnosis of the patient in the X-ray image?

↓ Downscale: abundant, easy

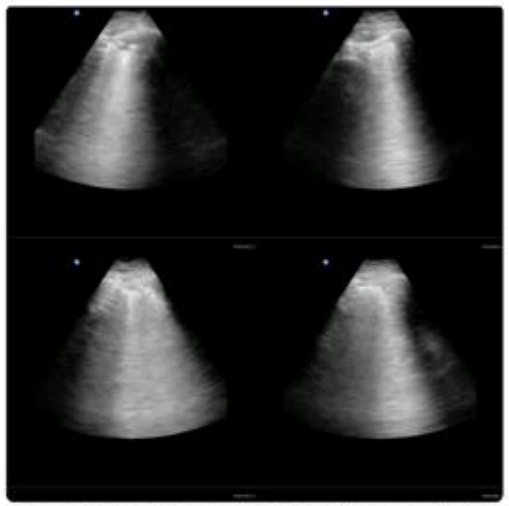

Question: Above is a lung ultrasound video. What is the diagnosis based on this lung ultrasound?

↑ Upscale: scarce, hard

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

QoQ-Med Model

https://github.com/DDVD233/QoQ_Med

- **Solution: Upscale scarce, hard questions while downscaling easy, abundant questions.**
- **Difficulty is assessed via hierarchical clustering, both within and across domains.**

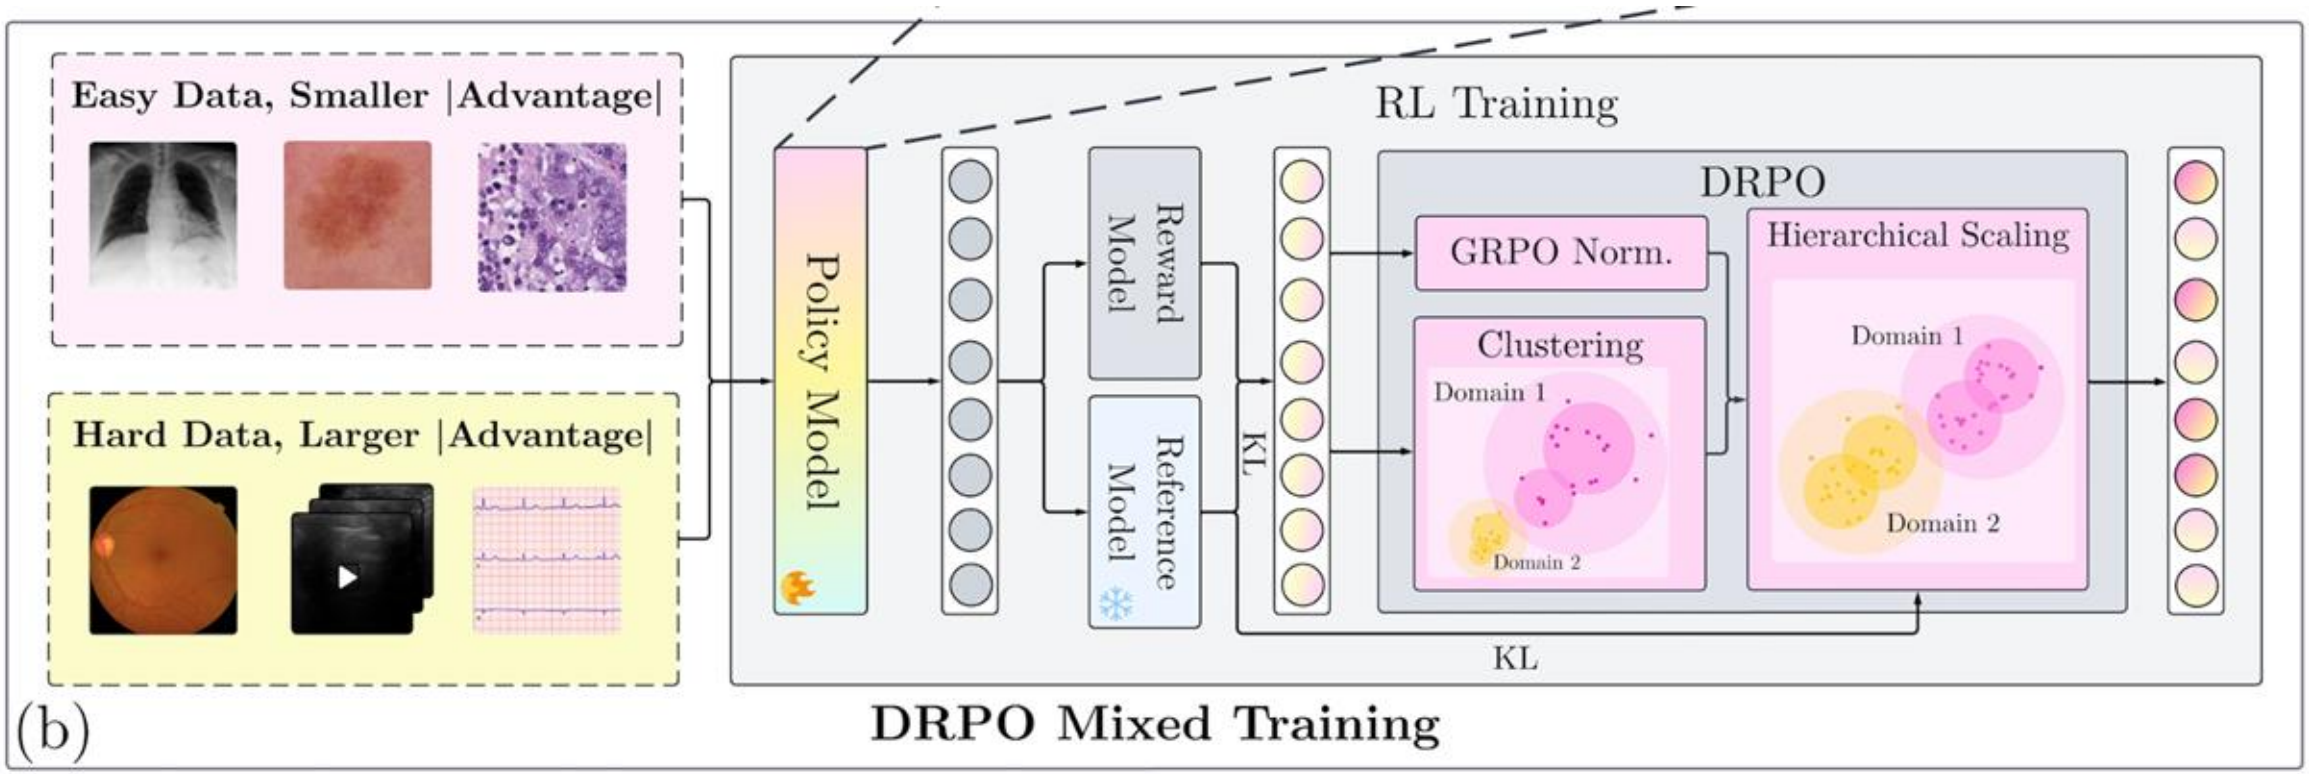

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

QoQ-Med Model

https://github.com/DDVD233/QoQ_Med

- First medical reasoning model that reasons across 2D/3D images, videos, time series and text.
- Substantial performance improvement in understudied modalities like ultrasound and CT scans.
- Model beats both open and closed-source models.

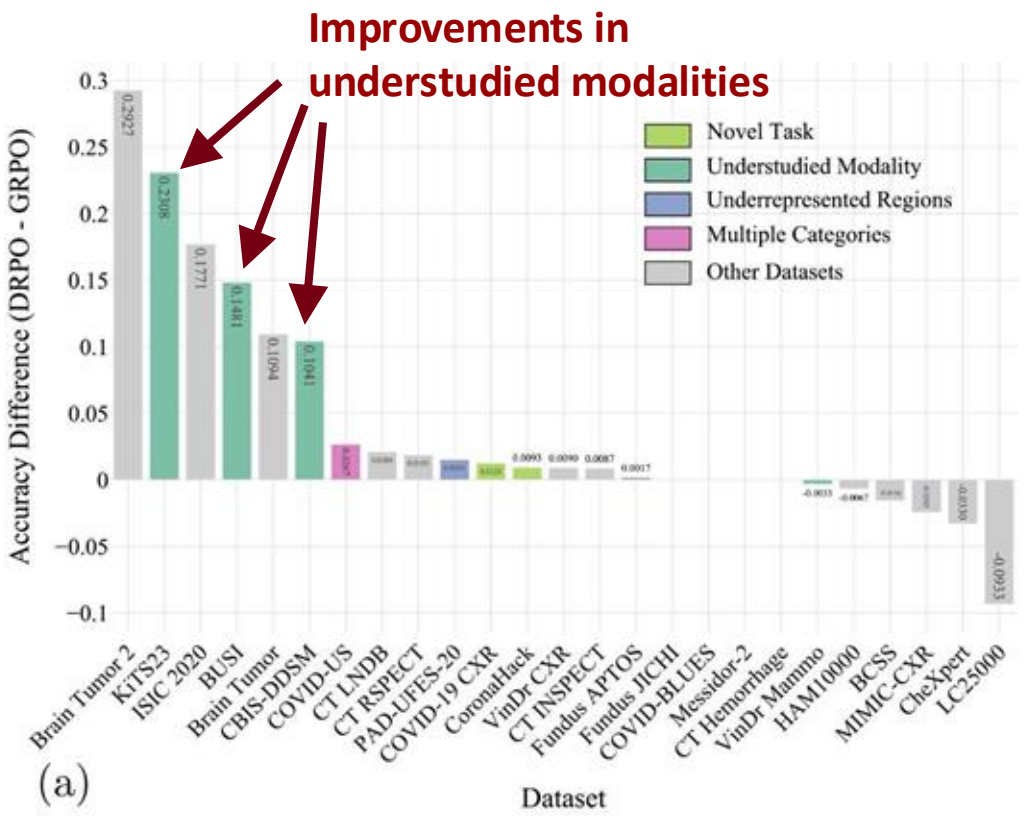

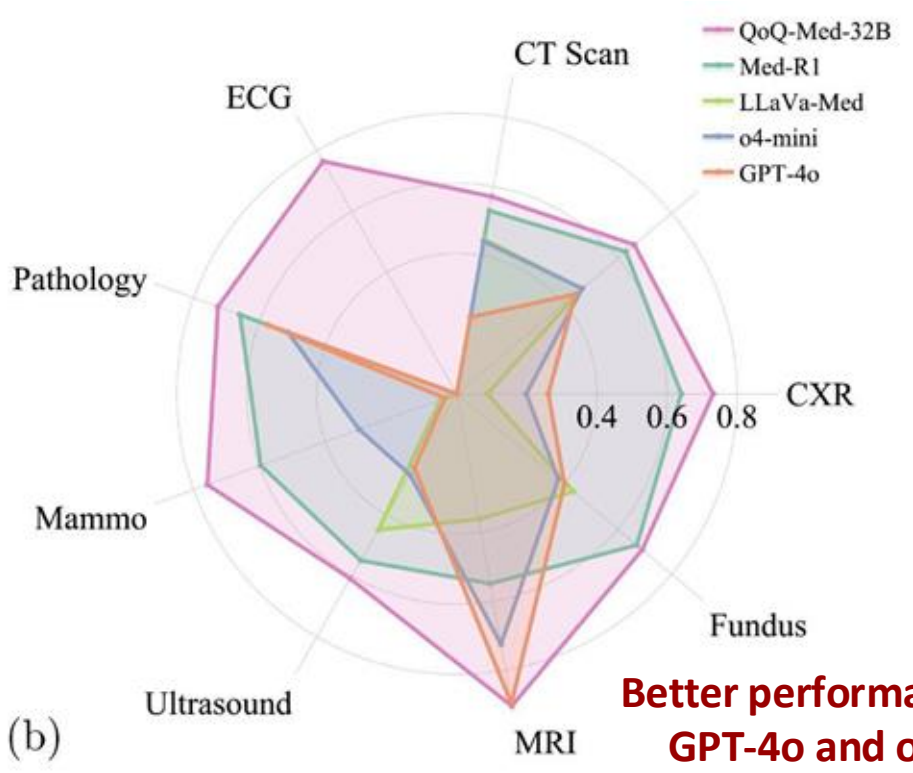

Better performance than GPT-4o and o4-mini

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

Transfer to MIMIC

https://github.com/DDVD233/QoQ_Med

- MIMIC is a dataset with paired image, text, time-series.
- Tests transfer from partial modalities to full multimodal setting.
- QoQ-Med outperforms standard GRPO training and other non-RL multimodal models.

Algo/Inputs	LOS		48-IHM	
	Acc	F1	Acc	F1
Qwen-2.5-VL-7B	0.649	0.105	0.555	0.293
GRPO-Full	0.626	0.105	0.551	0.354
DRPO-TextOnly	0.645	0.195	0.563	0.583
DRPO-ECG+T	0.639	0.204	0.602	0.528
DRPO-Vision+T	**0.669**	0.223	0.596	0.586
DRPO-Full	0.663	**0.283**	**0.642**	**0.597**

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

Reasoning Examples

https://github.com/DDVD233/QoQ_Med

- We annotated reasoning samples from validation with clinicians, with 3 categories: high relevance (to final diagnosis, low relevance, and no relevance).
- Around 45.3% of the reasoning is rated as highly relevant, and 100% of the rated samples contain highly relevant reasoning traces.

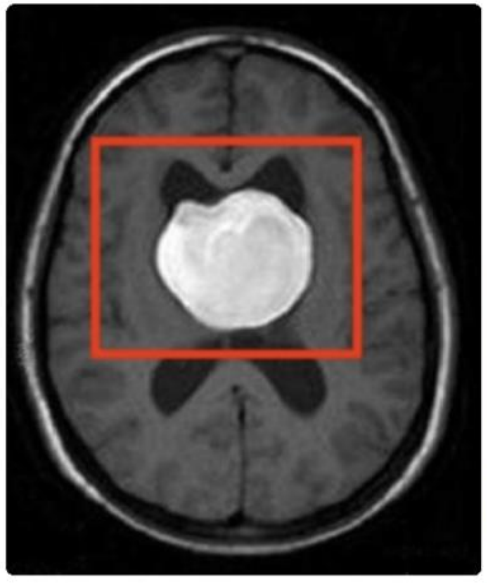

Question: Above is a brain MRI scan. Is there a tumor present in this brain MRI scan?

Ground truth: Has tumor

Model Reasoning:

...tumor appears as a round, well-circumscribed lesion...tumors are often characterized as having a distinct, often irregular, boundary ...**tumor is clearly distinct from the surrounding brain tissue**, which is a strong indicator of a lesion tha could be a tumor ...tumor is located in the brain...in a region where tumors are common...

Answer: Has tumor

Relevance to diagnosis: high low no

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

Reasoning Examples

https://github.com/DDVD233/QoQ_Med

Unified reasoning across images and medical sensors

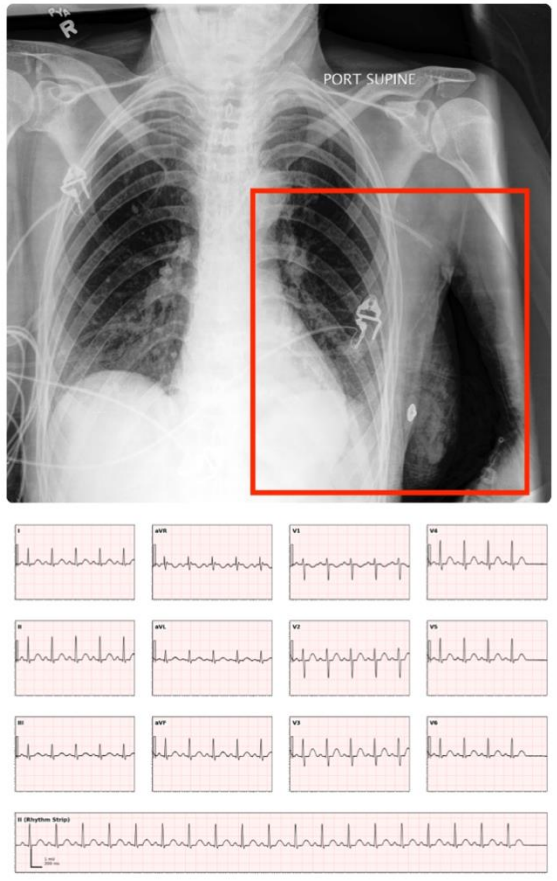

Question: Below is a history of a patient:...How long will the patient stay in the hospital?
A. 0-4 days B. 5-8 days C. 9-12 days D. more than 12 days
Ground truth: D. more than 12 days

Model Reasoning:

...the patient has been on mechanical ventilation for 96 hours, which is a long duration indicating a complex respiratory condition. The ECG and other lab results do not show any immediate critical issues, but the patient's hematocrit and hemoglobin levels are low, which could be a sign of anemia or other underlying conditions requiring treatment...the high and low heart rate alarms suggest the patient might be at risk for arrhythmias or other issues requiring close monitoring. Given the prolonged ventilation and the need for multiple invasive procedures, the patient is likely to require a significant period in the ICU or a high-dependency unit for close monitoring and management of their respiratory condition.

Answer: D. more than 12 days

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

Reasoning Examples

https://github.com/DDVD233/QoQ_Med

Detailed reasoning on image textures

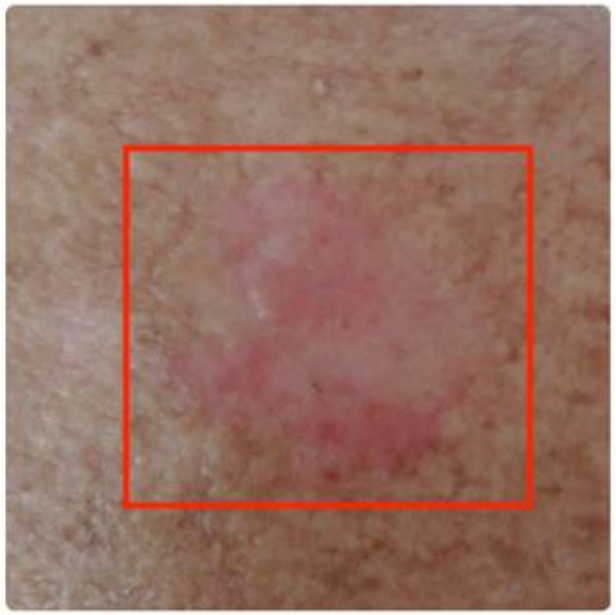

Question: Above is a clinical image of a patient. What is the diagnosis of the patient in the clinical image?
Ground truth: Basal cell carcinoma

Model Reasoning:

...not a dark mole...slightly raised, pinkish, and slightly elevated area...lesions have a rough, scaly surface, which is more consistent with the texture in this image...**lesion's color, shape, and location are consistent with basal cell carcinoma...**

Answer: Basal cell carcinoma

Relevance to diagnosis: high low no

[Dai et al., QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-aware GRPO Training. NeurIPS 2025]

multisensory intelligence

Tips and Training for Reinforcement Learning

1. Sanity Check with Fixed Policy
2. Monitor KL Divergence (in PPO-like algorithms)
3. Plot Entropy Over Time
4. Use Greedy Rollouts for Evaluation
5. Debug Value Function Separately: Visualize predicted vs. actual return
6. Gradient Norm Clipping is Crucial

Tips and Training for Reinforcement Learning

7. Check Advantage Distribution
8. Train on a Frozen Replay Buffer
9. Use Curriculum Learning: Gradually increase task difficulty or reward sparsity
10. Watch for Mode Collapse in MoE or Multi-Head Policies

Assignments for This Coming Week

This Thursday (5/8): final project presentations.

- Class from 1-3pm, let us know any time constraints.

Final project reports due 5/20 – 12 days to incorporate feedback from presentations

Meet with me and TAs today after class.