AGENT

A BENCHMARK FOR CORE PSYCHOLOGICAL REASONING

Enrico Meloni

SAILab, University of Siena

Introduction

SOCIALLY AWARE AGENTS

- Growing interest in **socially aware** agents
- **Human-like** interaction with humans
- Need for understanding **motivations and actions**
- It is an ability that comes **naturally to people**
- The so-called **intuitive psychology**

INTUITIVE PSYCHOLOGY

- The ability to **reason** about **other people mental states**
- Intuition from **observed actions**
  - Differentiate **agents** from **objects**
  - Expect agents to follow **physical constraints**
  - Expect agents to achieve goals in an **efficient way**
- A skill already developed in **pre-verbal infants**
- Even in the case of **partially observed** actions

EVALUATION OF AN AGENT CORE PSYCHOLOGY

- Need for a **rigorous evaluation process** of such psychology
- Assess how artificial agents learn about **core psychological reasoning**
- Assess how learned representations generalize to **new agents and environments**
- The authors propose AGENT, a benchmark inspired by **cognitive development experiments**
- Probe the agent understanding of intuitive psychology **as if it was a child**.

AGENT benchmark

DATASET

- It consists of a **large-scale dataset of 3D animated scenes**.
- An agent moves under **physical constraints** achieving **given goals**
- Organized in four categories of trials:
  - **Goal Preferences**
  - **Action Efficiency**
  - **Unobserved Constraints**
  - **Cost-Reward Trade-Offs**
- Cover the concept of agents as entities that **value some states of the world** over others
- And try to maximize their own **rewards** minimizing the **costs**
- The dataset is validated by external **human evaluators**

TRIALS OVERVIEW

- Every trial has two phases:
  - **Familiarization**: shows the typical behavior of an agent
  - **Test**: shows a video of the same agent in a different situation
- Each test video is assigned a category:
  - **Expected**: The agent behaves consistently to the familiarization phase
  - **Surprising**: The agent behaves inconsistently (e.g. goal inconsistency or physics violation)
- The evaluated model needs to correctly evaluate test videos as **expected** or **suprising**

SCENARIOS

- Four macro-types of trial, called Scenarios.
- The reasoning model needs to understand that the agent:
  - **Goal Preferences**: pursues a preferred goal
  - **Action Efficiency**: tends to take the most efficient actions to reach the goal
  - **Unobserved Constraints**: infers unobserved obstacles by assuming action efficiency
  - **Cost-Reward Trade-Offs**: understands the level of cost an agent is willing to pay for the preferred goal

SCENARIOS

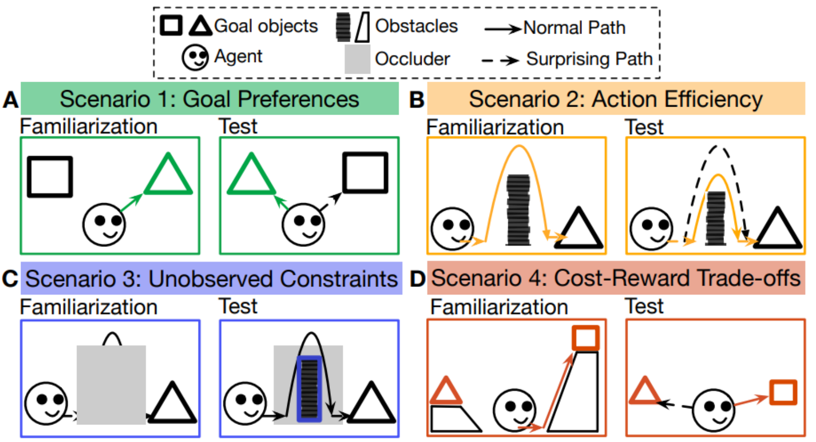
Goal objects
Obstacles
Normal Path
Agent
Occluder
Surprising Path
A Scenario 1: Goal Preferences
Familiarization
Test
B Scenario 2: Action Efficiency
Familiarization
Test
C Scenario 3: Unobserved Constraints
Familiarization
Test
D Scenario 4: Cost-Reward Trade-offs
Familiarization
Test

DATASET GENERATION

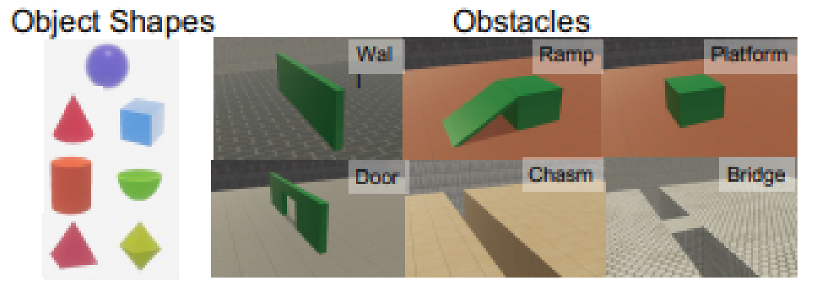

- The dataset is **procedurally generated** in **TDW**
- Obstacles, environment, agent preferences are **randomly picked**
- Motions are **hand-crafted heuristics**

DATASET CONTENT

- 8400 video, 5s to 25s, 35fps
- A total of 3360 trials
  - Training: 1920
  - Validation: 480
  - Test: 960
- Training and validations are pairs of familiarization and expected test
- Test set is composed of 480 pairs of expected/surprising videos that share the same familiarization
- The data contains: **RGB-D video**, **instance segmentation**, **camera parameters**, **3D bounding boxes**

EXAMPLES: SCENARIO 1

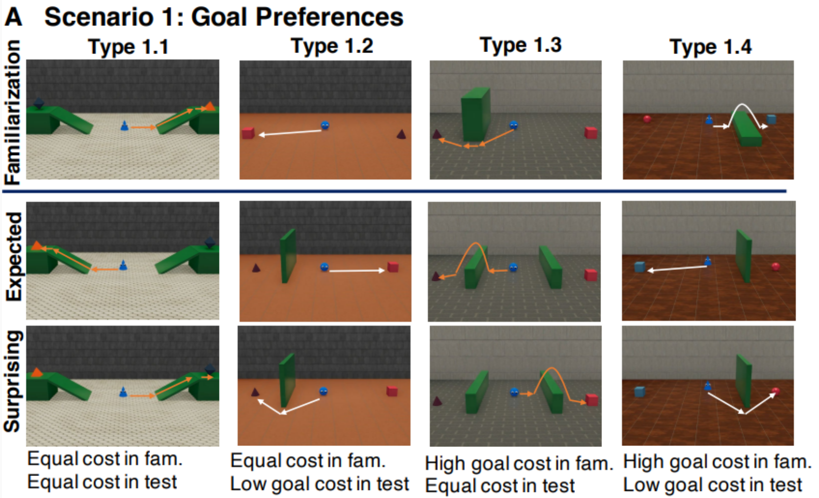

EXAMPLES: SCENARIO 2

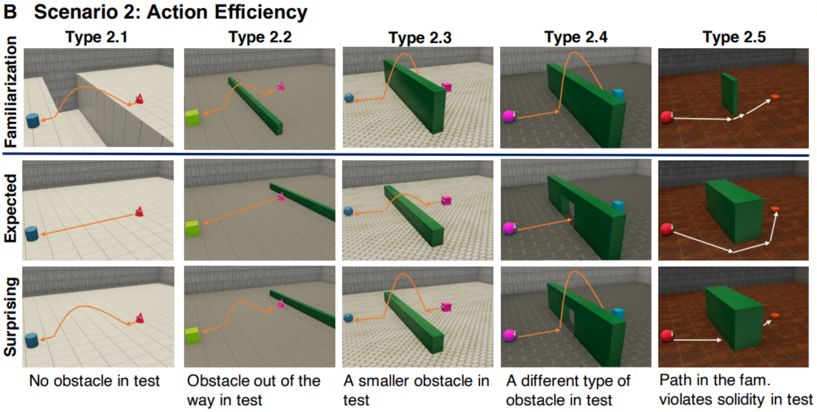

EXAMPLES: SCENARIO 3

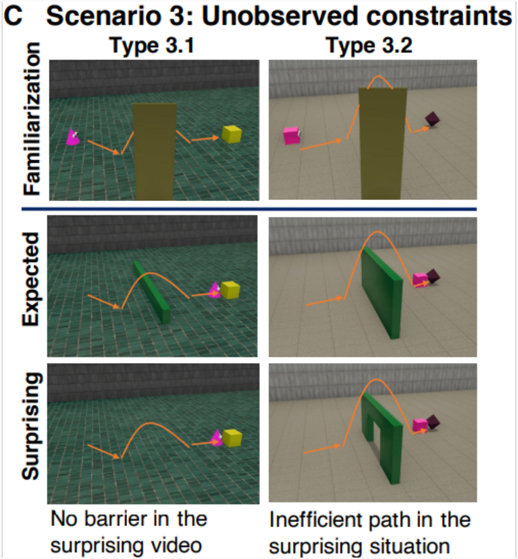

EXAMPLES: SCENARIO 4

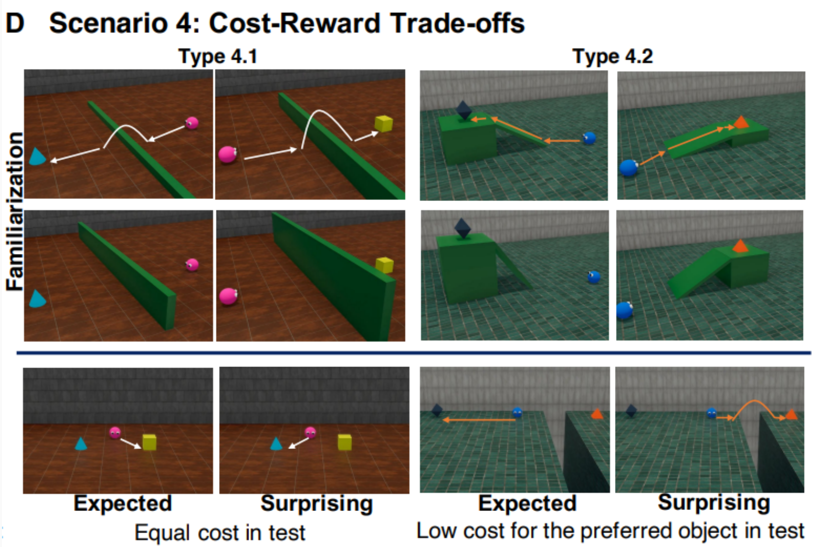

Experimental Results

BASELINE MODELS

- The dataset is evaluated with two baseline models
  - Bayesian Inverse Planning and Core Knowledge (BIPaCK)
  - Theory of Mind Neural Network
- The two models are based on Theory of Mind reasoning
- The paper sketches some high-level details of the models

BIPACK

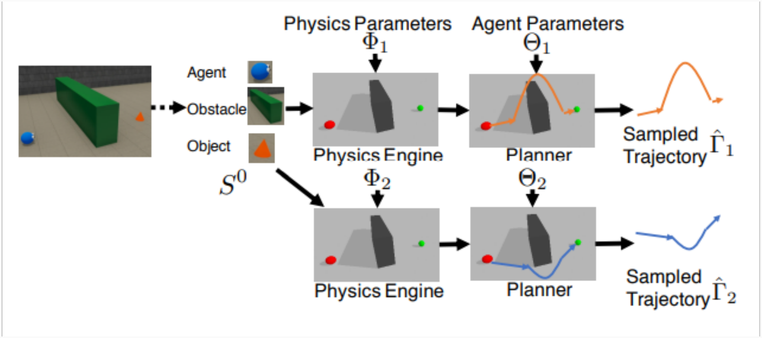

- Core idea: infer **hidden mental states** through a **generative model** of the agent's plans, during familiarization
- Combines **core knowledge of physics** and **physical simulation**

BIPACK

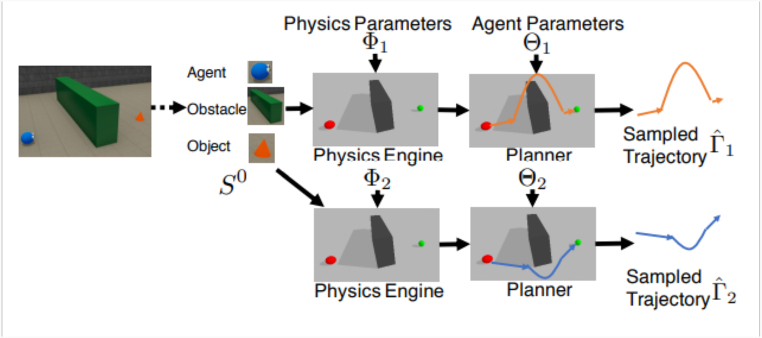

- The model estimates physical parameters and the agent's parameters (i.e. rewards and costs)
- Then it **indirectly** estimates the agent's trajectory using a built-in physics engine

TOMNET

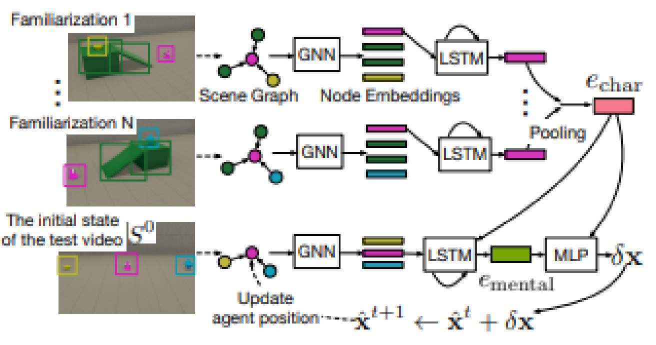

- Core idea: summarize **hidden mental states** into a **character embedding** during familiarization
- Combine it with the **embedding of the state** of the test video to infer a trajectory

LEAVE-ONE-OUT EXPERIMENTS

Condition	Method	Goal Preferences 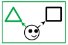					Action Efficiency 						Unobs. 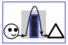			Cost-Reward 			All
		1.1	1.2	1.3	1.4	All	2.1	2.2	2.3	2.4	2.5	All	3.1	3.2	All	4.1	4.2	All	
	Human	.95	.95	.92	.97	.95	.87	.93	.86	.95	.94	.91	.88	.94	.92	.82	.91	.87	.91
All	ToMnet-G	.57	1.0	.67	1.0	.84	.95	1.0	.95	1.0	1.0	.98	.93	.87	.89	.82	.97	.89	.90
All	BIPaCK	.97	1.0	1.0	1.0	.99	1.0	1.0	.85	1.0	1.0	.97	.93	.88	.90	.90	1.0	.95	.96
G1	ToMnet-G	.50	.90	.63	.88	.75	.90	.75	.45	.90	.05	.66	.58	.77	.69	.48	.48	.48	.65
G1	BIPaCK	.93	1.0	1.0	1.0	.98	1.0	1.0	.80	1.0	1.0	.97	.93	.82	.86	.88	1.0	.94	.94
G2	ToMnet-G	.37	.95	.63	.88	.71	.35	.60	.75	.68	.85	.65	.63	.80	.73	.55	.95	.75	.71
G2	BIPaCK	.93	1.0	1.0	1.0	.98	1.0	1.0	.75	1.0	.95	.95	.88	.85	.87	.83	1.0	.92	.94

- When using **All** familiarization videos from every scenario, ToMNet and BIPaCK perform very well
- **G1**: for every scenario, train on **every type but one** and test on the **left out type**. ToMNet has some issues generalizing, but BIPaCK performs well
- **G2**: train on **every scenario but one**, and evaluate on the **left out scenario**. Results similar to G1.

SINGLE-TYPE EXPERIMENTS

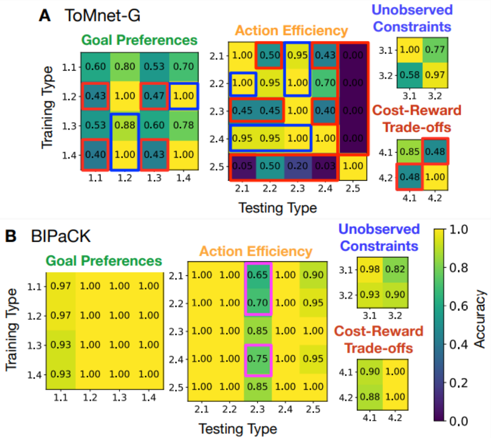

- **G3**: for every scenario, train on **a single type** and test on **all other types**.

SINGLE-SCENARIO EXPERIMENTS

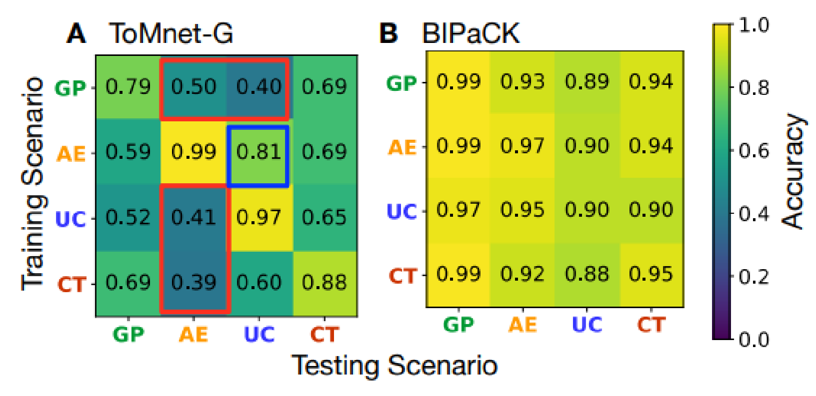

- **G4**: train on **a single scenario** and test on **all other scenarios**.

Conclusions

CONCLUSIONS

- AGENT, benchmark for core psychology reasoning
- Large-scale dataset of cognitively inspired tasks
- Probe artificial agents understanding of intuitive psychology
- Showcase the benchmark on two baseline models
- Show that the benchmark can help distinguish the performance of the two models on different generalization capabilities
- The benchmark is a well-structured diagnostic tool for developing better models of intuitive psychology

Thank you for listening!