Institute for Pure & Applied Mathematics
IPAM
University of California, Los Angeles

UCLA

Modeling human action understanding as inverse planning

Chris Baker, Josh Tenenbaum, Rebecca Saxe
MIT

Heider and Simmel demo

Intuitive psychology

- How do we infer hidden mental states of other agents that cause their observed behavior?
  - Beliefs, desires, plans, intentions, emotions.
- How do we use mental-state inferences to learn about the world?
  - Pulling out into traffic, jumping off a summit…
- What is the structure of intuitive theories of psychology that support these inferences, and how are those theories acquired?

Why did the man cross the street?

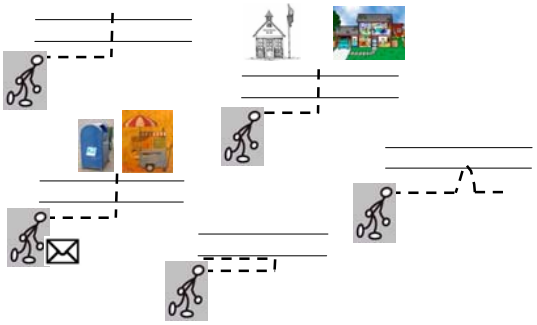

Principle of rationality

- Intuitively, we assume other agents will tend to take sequences of actions that most effectively achieve their goals given their beliefs.
- More formally: *inverse planning* in a goal-based Markov Decision Process (MDP).

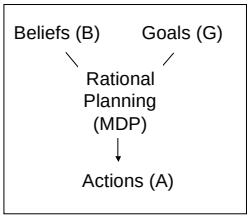

$$p(B,G \mid A) \propto p(A \mid B,G)\; p(B,G)$$

Caution!

- We are not proposing this as a computational model of how people plan! It is a computational model of people's mental model of how people plan.
- Whether planning is "rational" or well-described as solving an MDP is an interesting but distinct question.

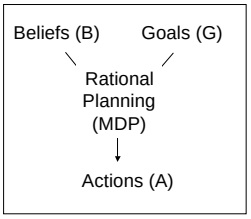

Rational action understanding in infants
(Gergely & Csibra)

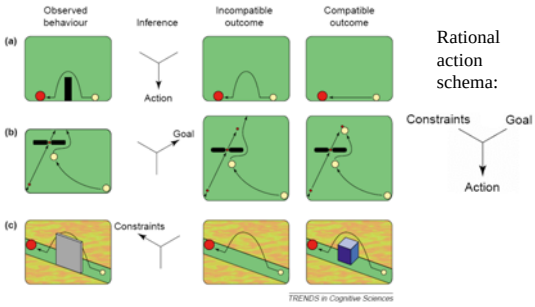

The present research

- Aims
  - To test how accurately and precisely the inverse planning framework can explain people's intuitive psychological judgments.
  - To use the inverse planning framework as a tool to test alternative accounts of people's intuitive theories of psychology.
- Experiments 1 & 2: goal inference
- Experiments 3 & 4: action prediction

Experiment 1: goal inference

- Method
  - Subjects (*N*=16) view animated trajectories in simple maze-like environments.
  - Subjects observe partial action sequences with several candidate goals and are asked to rate relative probabilities of goals at different points along each trajectory.

Experiment 1: goal inference

- Method
  - Subjects (*N*=16) view animated trajectories in simple maze-like environments.
  - Subjects observe partial action sequences with several candidate goals and are asked to rate relative probabilities of goals at different points along each trajectory.
- Set up
  - Cover story: intelligent aliens moving in their natural environment.
  - Assume fully observable world: agent's beliefs = true states and transition functions of the environment.
  - 100 total judgments, with 3-6 judgment points along each of 36 different trajectories (= 4 goal positions x 3 kinds of obstacles x 3 goals).

Specific inverse planning models

- Model M1(β): fixed goal
  - The agent acts to achieve a particular state of the environment, which is fixed for a given action sequence.
  - Small negative cost for each step that does not reach the goal.

G: goal state
S_t: environmental state at time t
A_t: action at time t
β: parameter controlling "noisiness" of actions

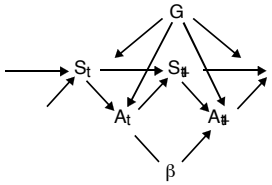

Specific inverse planning models

- Model M2(β,γ): switching goals
  - Just like M1, but the agent's goal can change at any time step with probability γ.
  - Agent plans greedily, not anticipating its own potential goal changes.

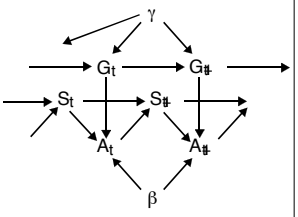

Sample behavioral data

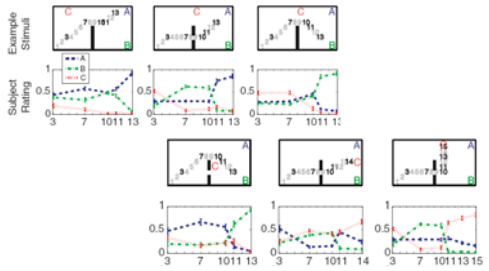

Modeling results

M1(β)
Fixed goal

M2(β,γ)
Goal switching

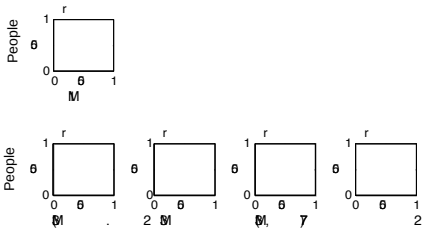

Modeling results

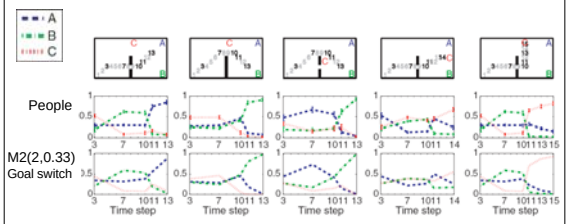

An alternative heuristic account?

- We can rule out most *very* simple heuristics
  - Infer the goal that is closest
  - Infer goal based on direction
- What about more complex heuristics?
  - Last-step: infer goal based on *only* the last movement (instead of entire path)
  - Delta-D: infer goal based on the derivative of the distance from the goal with respect to time
  - Last-step and Delta-D heuristics yield similar, sometimes identical predictions

An alternative heuristic account?

- Last-step heuristic: infer goal based on only the last movement (instead of the entire path)
  - a special case of M2, equivalent to M2(β,.67).
- This model correlates highly with people's judgments in Experiment 1.
- However, there are qualitative differences between this model's predictions and people's judgments that suggest that people are using a more sophisticated form of temporal integration.

An alternative heuristic account?

A thought experiment…

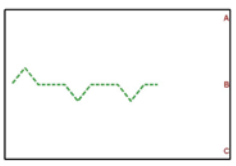

M2(2,0.1)
Goal switch

M2(2,0.67)
One-step
heuristic

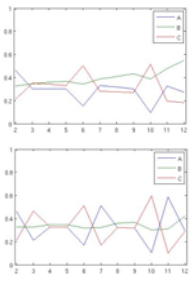

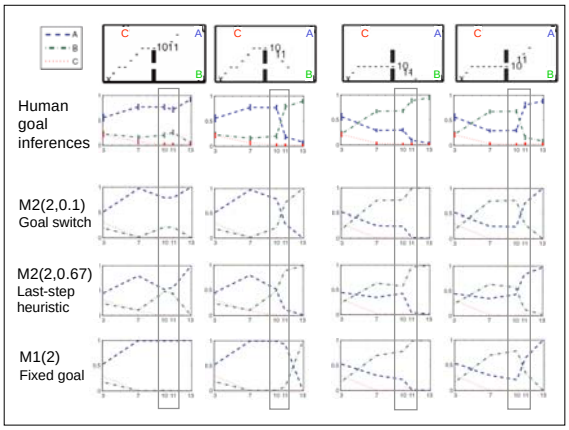

Analysis with full data set

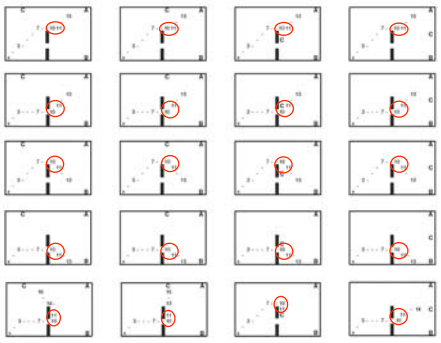

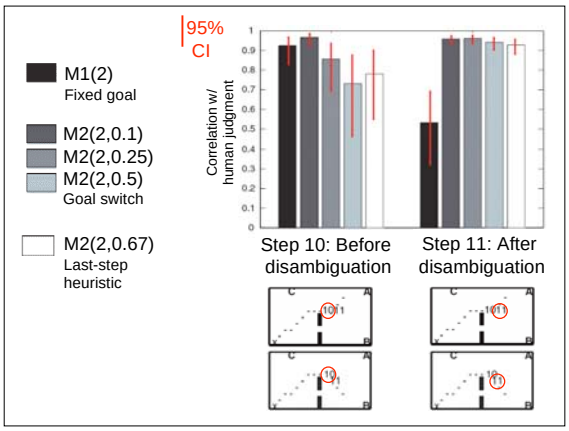

Summary

- Inverse planning is a framework for inferring mental states from behavior, assuming a rational agent.
- Inverse planning can be used to predict people's goal attributions with high accuracy (at least in simple environments).
- Goal attributions are better explained by inverse planning with a dynamic space of goals than a simpler model with fixed goals or various "one-step" heuristics.
- Intuitive psychology appears to be based on a precise predictive model, much like intuitive physics.
- Intuitive psychology may be "more rational" than actual psychology….

Open directions

- More complex environments.
- Hierarchical goal structures, plans.
- Richer mental-state representations, e.g. recursive belief: "I'm guessing that you think Mary is wrong, but trust me, she isn't."
  - Competitive interactions (e.g., Jun Zhang)
  - Language understanding
- The acquisition of intuitive psychology.
- The relation between psychology (how people actually think and plan) and intuitive psychology.