Information Seeking, Curiosity, and Attention: Computational and Neural Mechanisms

Jacqueline Gottlieb, Pierre-Yves Oudeyer, Manuel Lopes, and Adrien Baranes

Shrey Sahgal

Curiosity Drives Humanity

(Curiosity Rover)

Exploration

- Actions with the goal of obtaining information
- Alter the observer's epistemic state
- **Imperative: reduce uncertainty**
- Extrinsic vs. intrinsic motivations

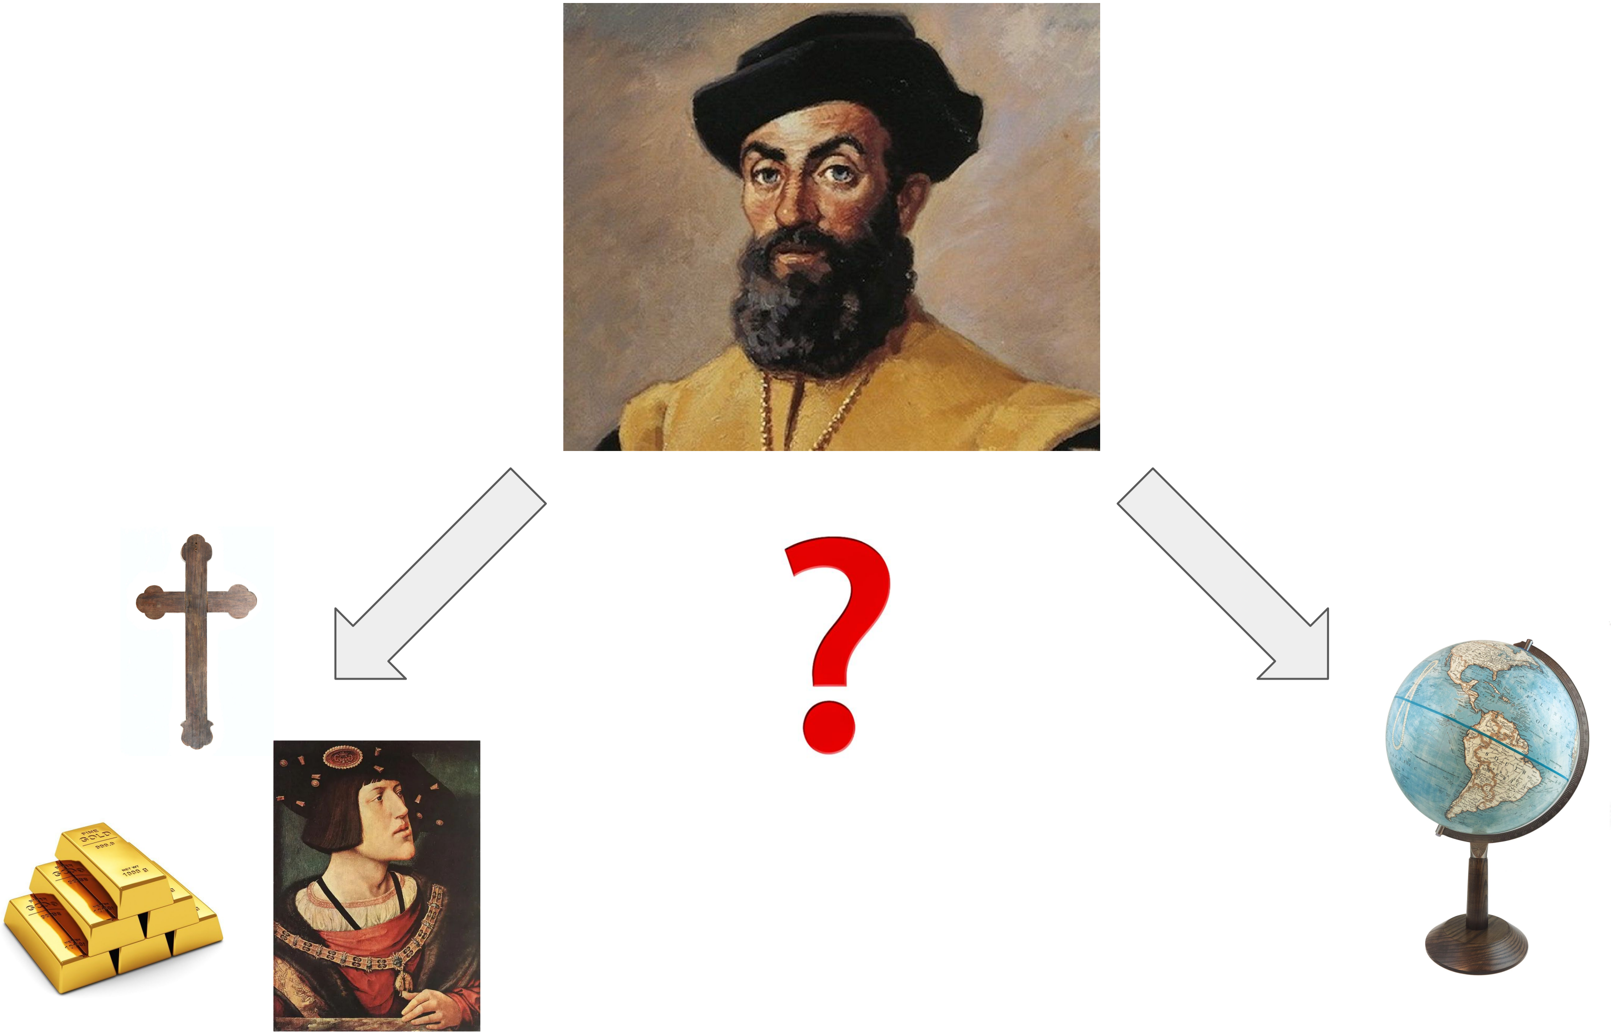

Extrinsic vs. Intrinsic Motivation

Extrinsic

- Information gathering as a means to an end
- Maximize an agents' progress towards a separate goal (e.g. wealth, proselytization, or political capital)

Intrinsic

- Search for information is itself a goal
- "Curiosity" or "interest"
- Intrinsic rewards that assign value to information

Task-Directed Information Search

- Select actions that maximize immediate or future rewards
- Partially-observable Markov decision processes
- Information-gathering as an intermediate step
  - The significance of the information must be known
  - Use information to plan future actions
- Depends on the observer's momentary uncertainty

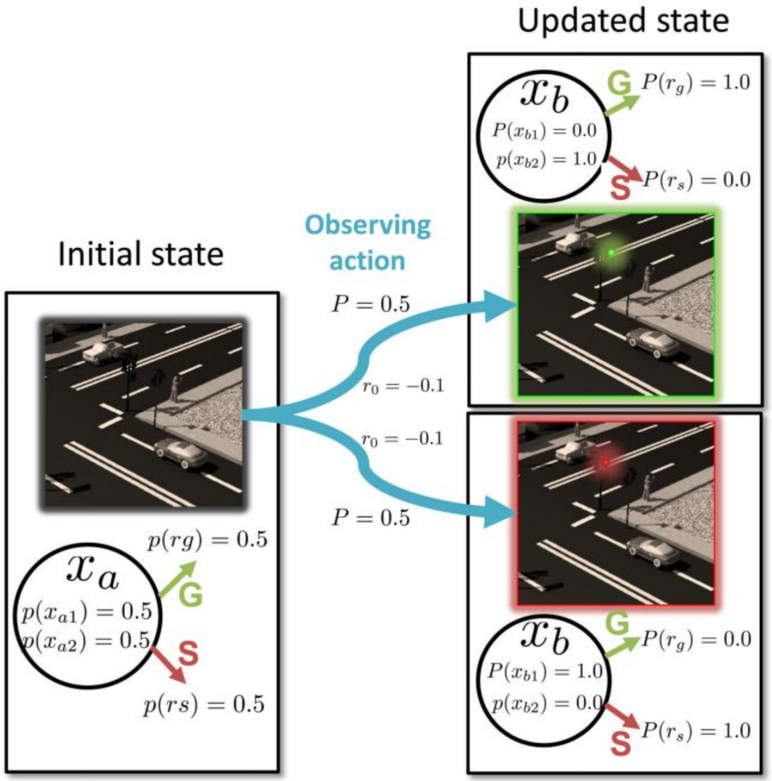

Searching for a Task Strategy

- Machine learning: minimize a cost function
  - Dozens of examples: reinforcement learning, stochastic optimization, evolutionary techniques, etc.
- Requires a lot of exploration
  - But still aim to accomplish a particular extrinsic goal

DeepMind Atari
Deep Q Learner

Active Information Search: Eye Movements

- Visual information is sampled with saccadic eye movements
    - Information gathering is an active process
- Sample information to reduce uncertainty
- Visuomotor coordination example: eyes move ahead of the hand
- Pedestrian tracking example: subject preferentially looked at oncoming pedestrians with uncertain trajectories

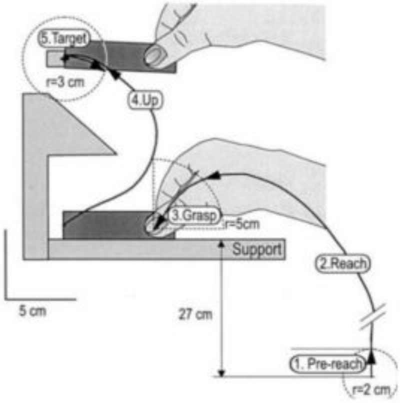

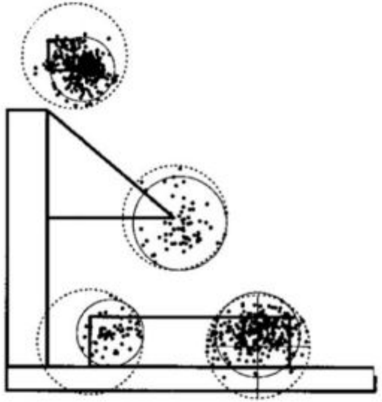

Curiosity and Autonomous Exploration

- Assumption: brain creates intrinsic reward for exploration
- Dopaminergic system is:
    - Sensitive to intrinsic rewards
    - Responds to anticipated information about rewards (in monkeys)
    - Activated by paradigms that induce curiosity
- Computational: *why* does the brain generate these rewards?
- Algorithmic/physical: *how* are these rewards calculated?

Evolutionary Advantages to Exploratory Behavior

- Maximize long-term evolutionary fitness in rapidly changing environment
- Computer-simulated models
  - Limited cognitive capacity of the agent
  - Information that is not immediately useful could be re-used in the future
- Free-energy principle: having a large array of skills can be useful to avoid future surprises
- Gestalt psychology: humans' "need for cognition"

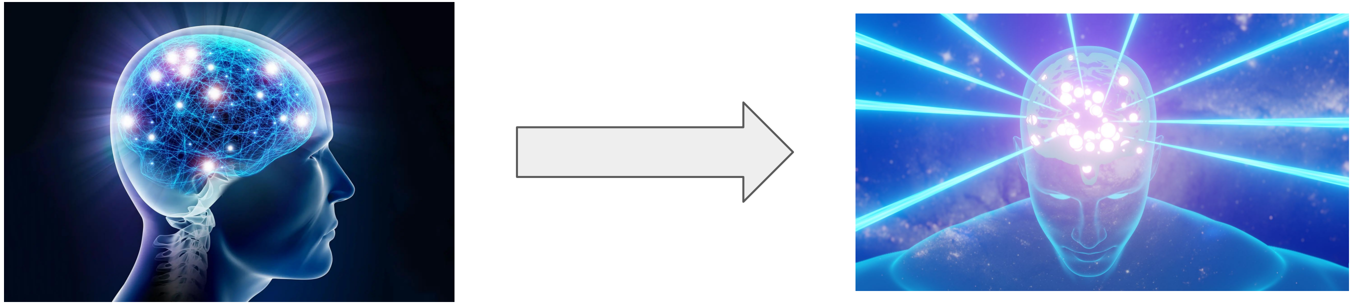

Challenges of Learning

- Incredibly vast learning space for infants
  - Sensorimotor domain: learning a huge repertoire of complex actions, potentially infinite ways of interacting with objects
- *Unlearnable* situations: the agent can't detect regularities or improve
  - Eg. trying to learn to run before learning to crawl

Randomness, Novelty, Surprise, and Uncertainty

- Random action selection
- Biases towards novel, surprising, or uncertain events
  - Aligns with the theory there intrinsic rewards are given for novel actions and for exploring states that produce high empirical prediction error
- Contextual novelty (surprise!)
  - Bayesian inference: difference between prior and posterior probabilities
  - High prediction error for high-confidence states
- Bias towards actions and states with high entropy/variance

Heuristic Methods: No Guarantees on Learning

- Just because an event is novel, surprising, or uncertain doesn't mean it will result in detectable, generalizable, or useful regularities
- Still efficient in small, closed spaces
- But in large open-ended spaces, agents can be caught in unlearnable situations

Information Gap Hypothesis

- Curiosity is the result of a discrepancy between what an observer *knows* and what they would *like to know*
- Aims to explain “specific epistemic curiosity”
- Curiosity as a deprivation phenomenon: analogous to a need to fulfill physical needs
- Agent needs prior knowledge to set the starting and reference points
  - You can’t be curious about something you don’t know!

Learning Progress Exploration: Developmental Robotics

- Choose action that maximizes learning progress
  - This ensures that something is always learned; unlearnable situations are avoided
- Allows the robot to learn a large repertoire of skills in high dimensions
- System gradually shifts interest from simple to more complex tasks
- Hypothesis: infants' sensorimotor progression isn't pre-programmed but rather emerges from intrinsic motivation to interact with the environment
- No empirical evidence of quantitative response to learning progress

Discussion

Task-Specific vs. General Curiosity?

Haoyuan Ma 2 days ago
The paper made clear that we should implement intrinsic motivation for exploration that is not task based. Because of limited cognitive ability of the agent and that collected information might not be immediate rewarding, the intrinsic curiosity performs better than external curiosity. Does it mean ultimately instead of focusing on AI that solves problem task specific, we should focusing on building this intrinsic curiosity which ultimately solves all the problems?
helpful! | 0

Lance Bassett (lanbas) 4 hours ago
Yes I think it's still interesting to think about a higher level of AI that has general knowledge and learns through intrinsic curiosity. But as other's have said it would be difficult and likely perform worse than task specific solutions for a long time. So maybe there would be a lack of motivation for this problem because it's not that we want robots to replace humans, but we want robots to replace humans in specific situations. Still, if a flexible general knowledge robot could replace a task-specific robot that would be best case.
helpful! | 0

Zixuan 3 hours ago
A recent paper "LEARNING INVARIANT REPRESENTATIONS FOR REINFORCEMENT LEARNING WITHOUT RECONSTRUCTION" presents a different view. It shows that instead of emphsizing reconstruction/prediction accuracy, which is the focus of curiosity-driven approach, we should learn the model in a task-driven manner.

As an illustrative example, when we are driving, we don't really care how the cloud in the sky move and deform. Instead, we always focus on the critical region, such as the other vehicles on the road. A pure curiosity-driven approach will waste the limited computing power on the things that are not important for the task.
helpful! | 0

IAC vs. Novelty/Surprise/Uncertainty Bias

- There is no empirical evidence that there is a quantitative response to learning progress
- Does Intelligent Adaptive Curiosity (IAC) capture the intrinsic rewards for novelty/surprise/uncertainty that has been empirically seen?

Visual Information Gathering

- Cost/benefits of oculomotor systems for active information gathering?
- From Piazza: It’s interesting that eye movements are a natural indicator of the brain’s active information search. How does this relate to previous readings about eye movements/saccading? If eye movements are so important in our visuospatial perception, how can we creating computing systems that perform similar ‘eye movement?’

Brain Cells on a Chip Learn to Play Pong

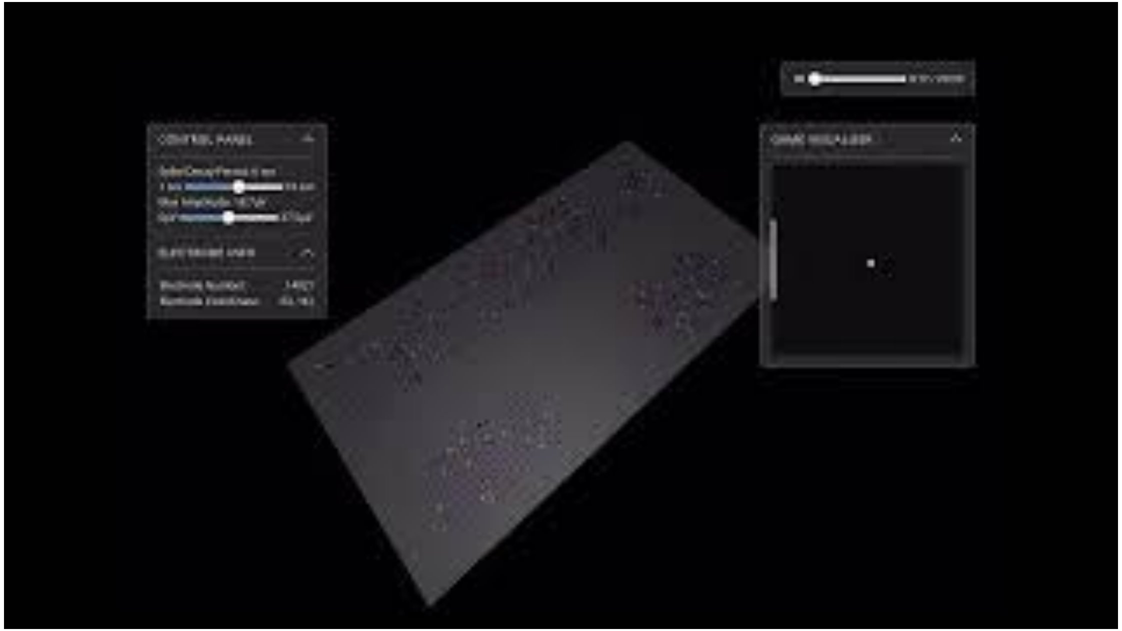

Active Perception

- From this reading and from many others in class, there seems to be quite a bit of evidence that shows that our perception in nature is much more active than it is static. How can we utilize this active quality of perception in nature to inform how we create computer vision models? Is there still benefit in using traditional, static methods for computer vision?