


How would a non-expert assess the limits and capabilities of an AI system?

## INTRODUCTION

Objective: Learn an interpretable model of a black-box agent by interrogating it.



Key technical challenge:

- Which sequence of queries to ask?

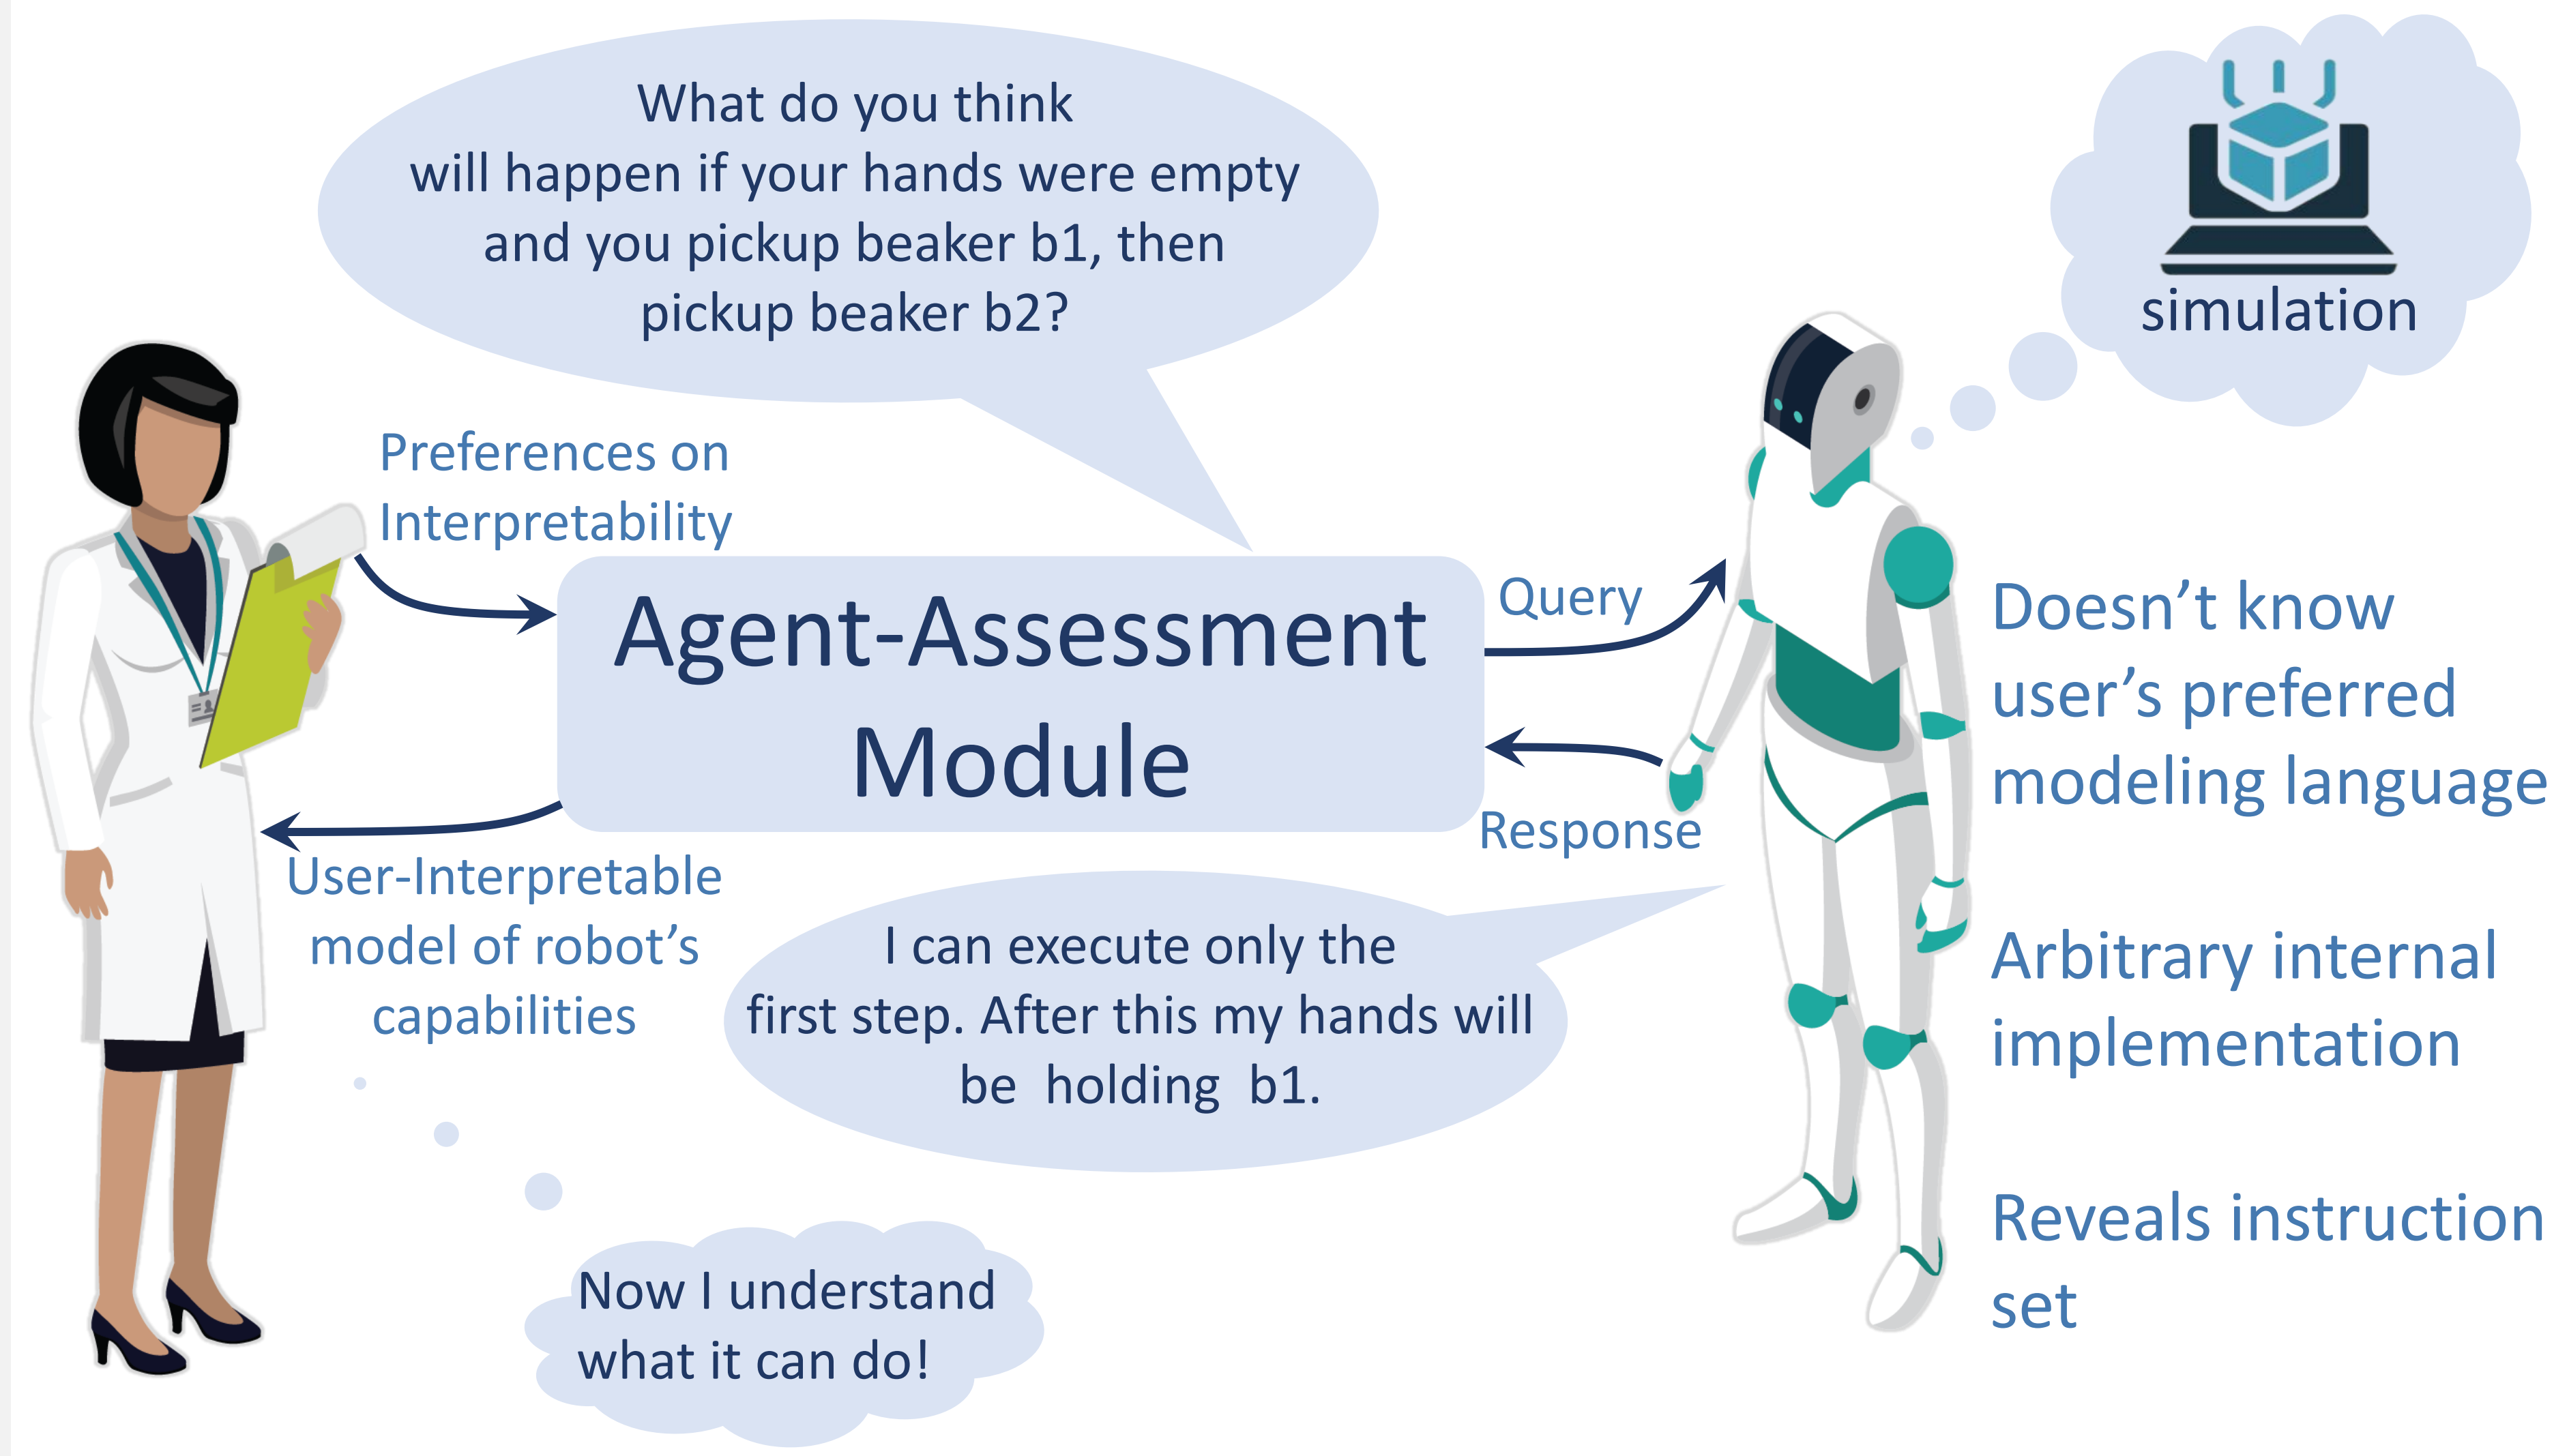
## ALGORITHM (AIA)

- 1 Start with the most abstracted node in lattice.
- 2 Pick abstraction candidates in some order.
- 3 For each candidate, generate three models and for each pair of models:
- 4 • Generate a distinguishing query  $Q$  and pose it to the agent.
- 5 • Get the response  $R$  from the agent.
- 6 • Prune out the incorrect variants of candidate models.
- 7 • Repeat steps 3-6 till the model is fully estimated.
- 8 Return the final set of model(s).

## SALIENT FEATURES

- Efficiently learns internal model of an autonomous agent in a STRIPS-like form.
- Needs no prior knowledge of the agent model.
- Only requires an agent to have rudimentary query answering capabilities.
- Queries can be answered using a simulator.

## EXAMPLE OF AGENT INTERROGATION



## ABSTRACTION IN SPACE OF MODELS

```
(:action load_truck
:parameters (?package ?truck ?location)
:precondition (and (at ?truck ?location)
(+/-/∅) (at ?package ?location))
:effect (and (not (at ?package ?location))
(in ?package ?truck)))
```

Abstracted model

Abstracted model

This predicate can appear in three forms:

- positive
- negative
- absent

```
(:action load_truck
:parameters (?package ?truck ?location)
:precondition (and (at ?truck ?location)
(at ?package ?location))
:effect (and (not (at ?package ?location))
(in ?package ?truck)))
```

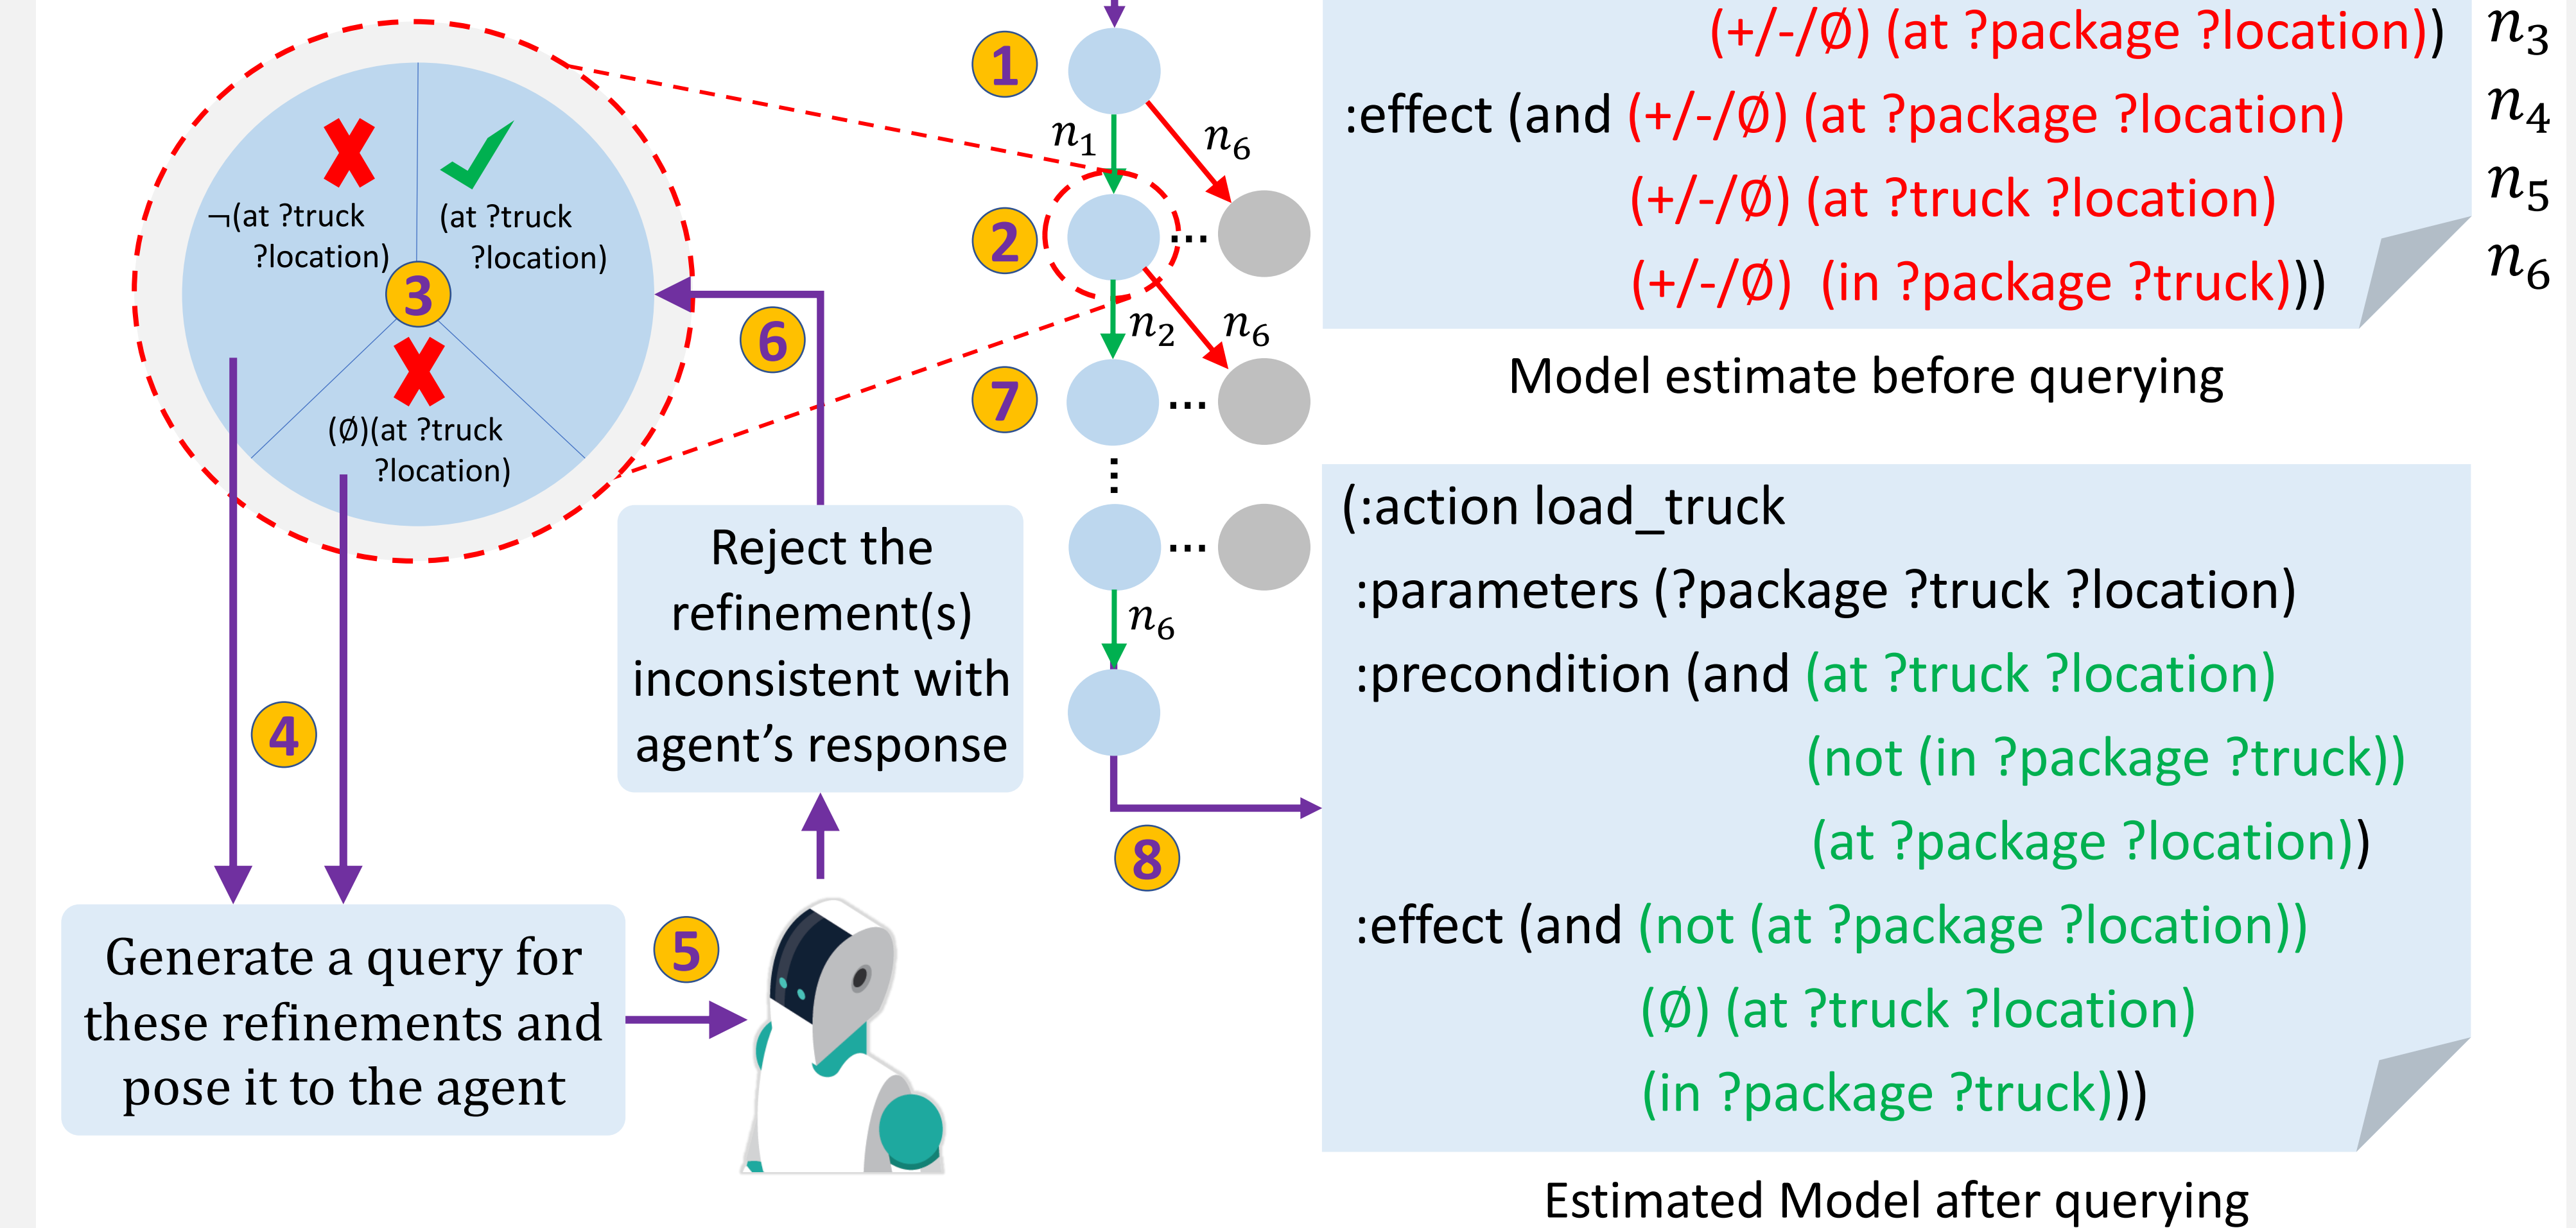
Concrete model

Concrete model



abstraction ↑

## KEY ALGORITHMIC PRINCIPLE

**Key feature of the algorithm**  
Each time we prune an abstracted model, we prune a very large number of models at the most concrete node.



## RESULTS



- AIA efficiently derives interpretable agent models for a range of agents.
- AIA is much faster than state of the art methods for deriving models based on passive observations.
- AIA offers better convergence guarantees.

Refer to the paper for detailed results

[bit.ly/3p4cVRu](https://bit.ly/3p4cVRu)

