CS6301.001 Robotics

# Trajectory-Aware Human Feedback for Efficient Hierarchical Reinforcement Learning in Robotics

Naveen Gururmurthy
Nikunj Gohil

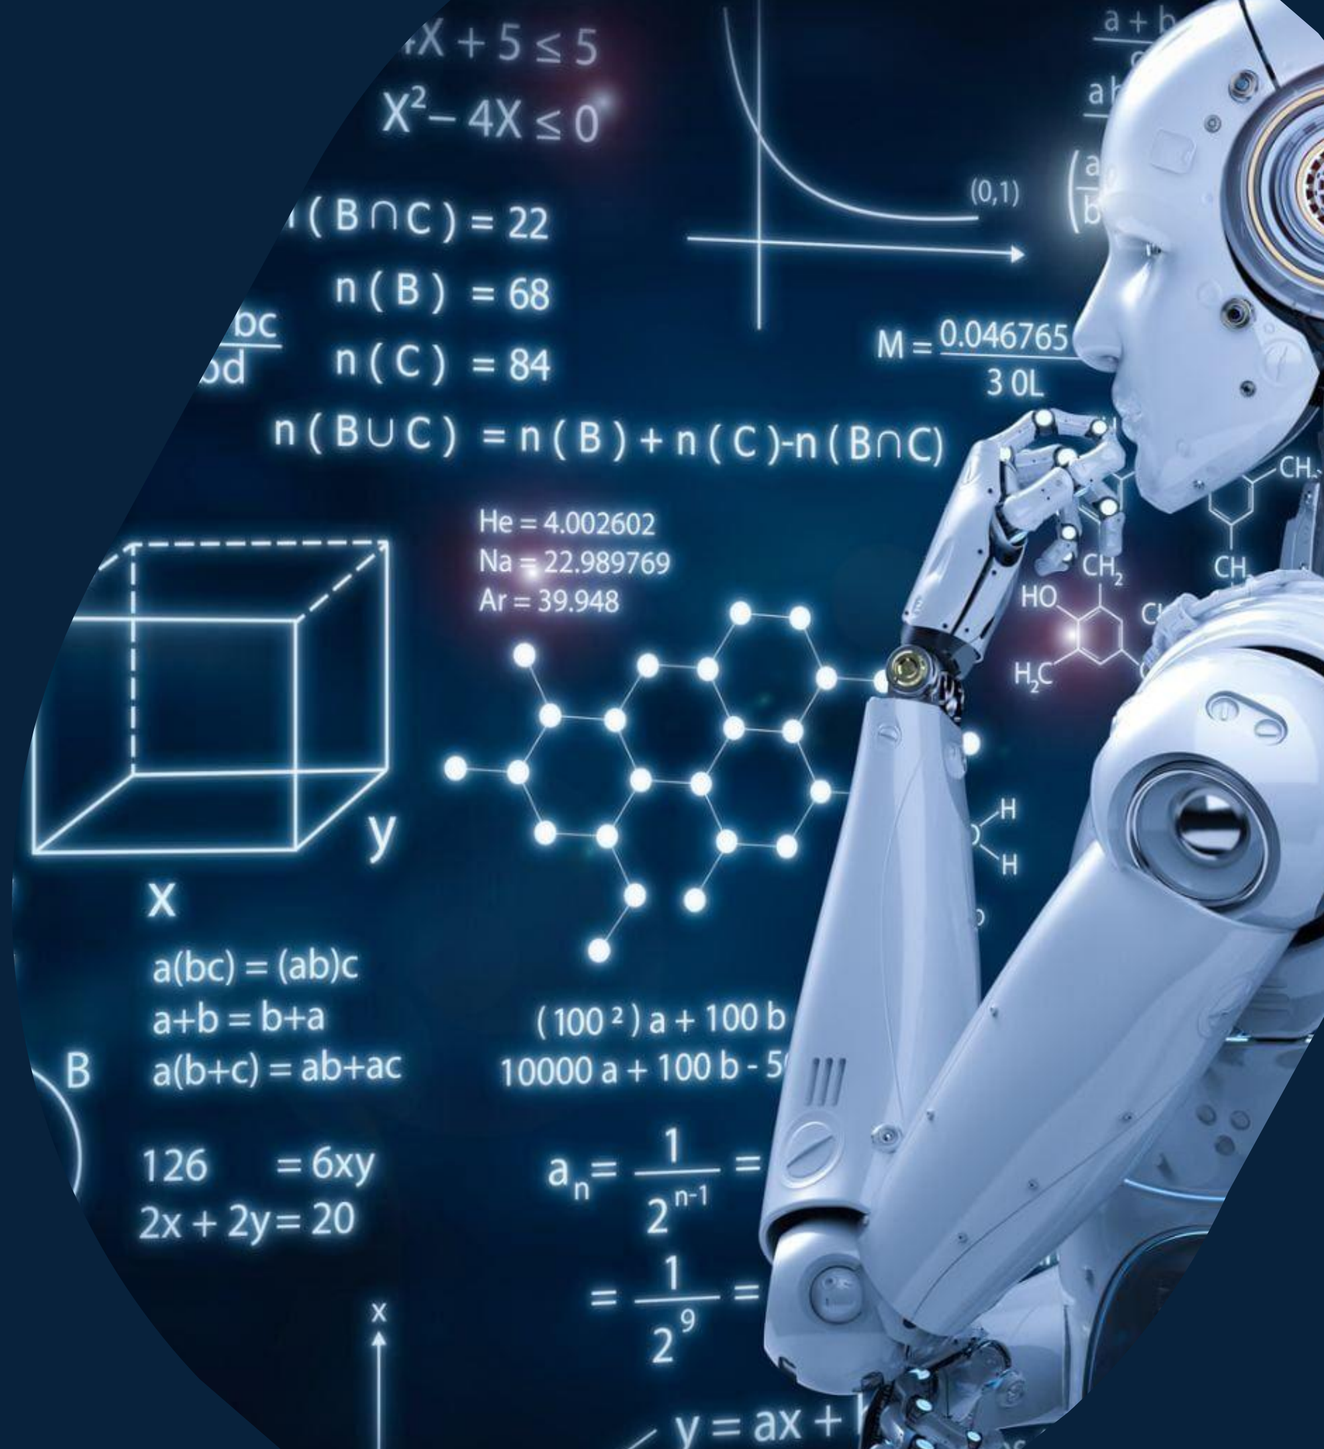

# Agenda

Introduction

Literature

Method

Demonstration

Evaluation

# What is RL

A machine learning technique that uses rewards or punishments to train a model to make decisions.

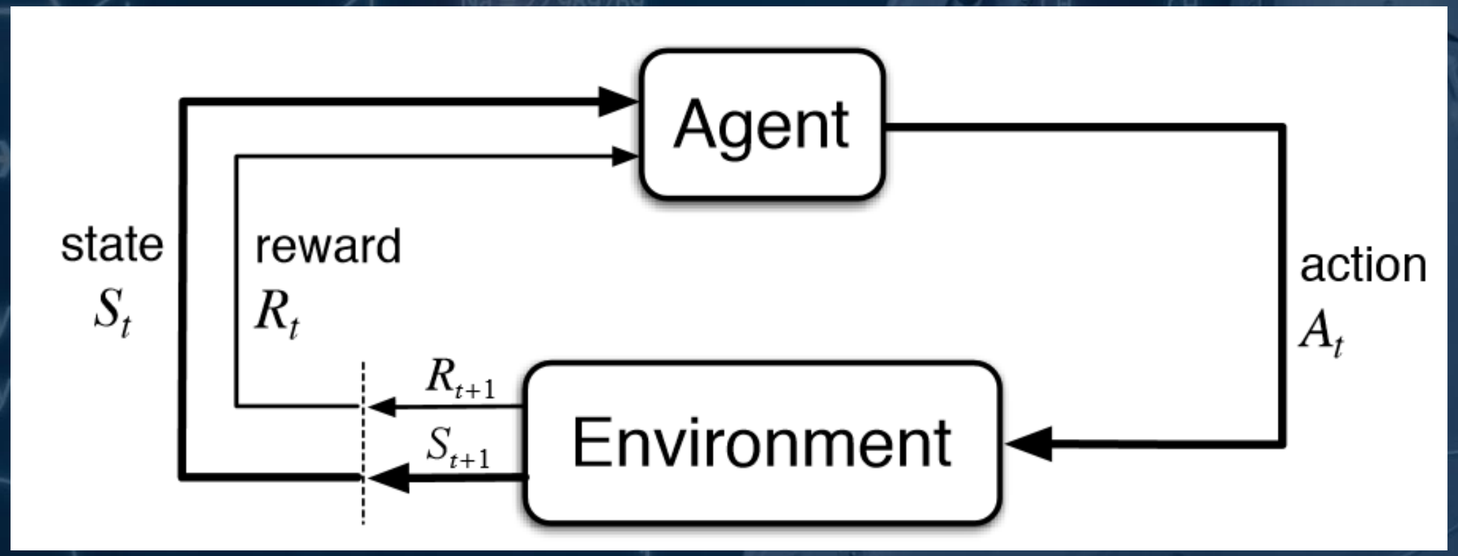

# What is HRL

HRL, instead of using a single policy to learn a task, the task is divided into subtasks. A meta-policy chooses the subtask first and then a low level policy learns the specific subtask, which helps in solving complex tasks.

So, the essence lies in effective subgoal generation and their efficient completion.

In the next slides, we propose an extension to HRL

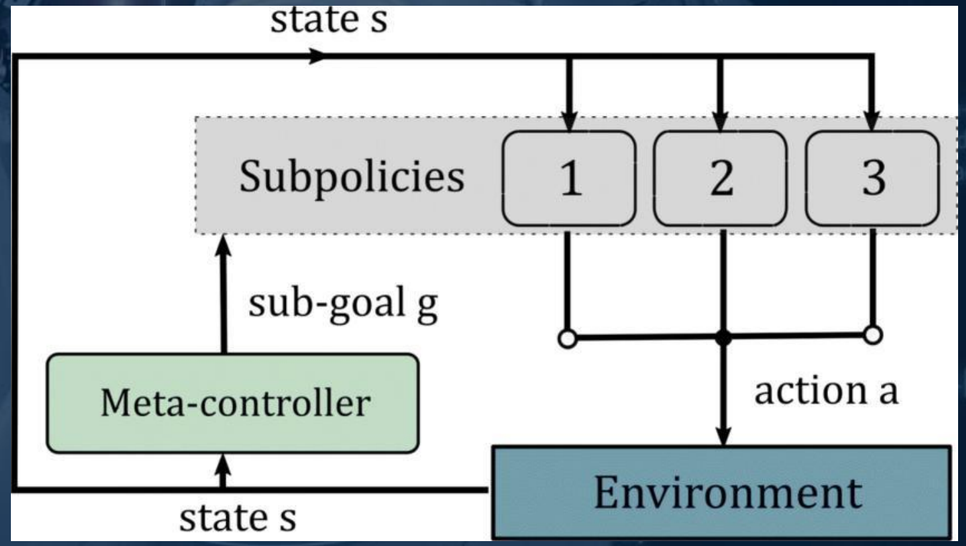

# Literature

Roadblocks in HRL are effective subgoal generation and sparse rewards.
Humans can generate subgoals for the robot but it can become costly and challenging.
One way to overcome sparse rewards obstacle is using Goal Conditioned Reinforcement Learning. Wherein, a failed task is viewed as a way to do that specific task.

# Method

Using Human Guided HRL, it is ensured that the subgoals are both instructive and achievable.
The feedback can be in the form of approving and ranking subgoals, providing specific instructions.
A reward model can be trained according to human reference relaying information for efficient task completion.
The difficulty of subgoals can be dynamically adjusted, so that the agent is challenged and does not get stuck on overly difficult subgoals.

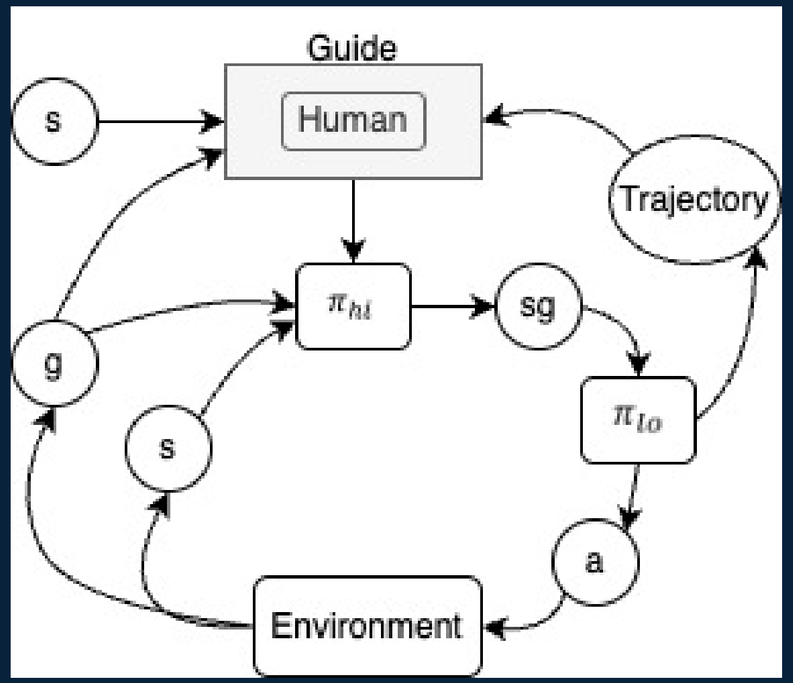


s - current state
g - desired goal
sg - subgoal
a - action

# Progress

Created a reward model based on human feedback which will in turn train the robot in carrying out tasks.

Set up a PyBullet simulation environment. Wherein the robot is trained to pick up an object whose location varies in each sample.

# Demonstration

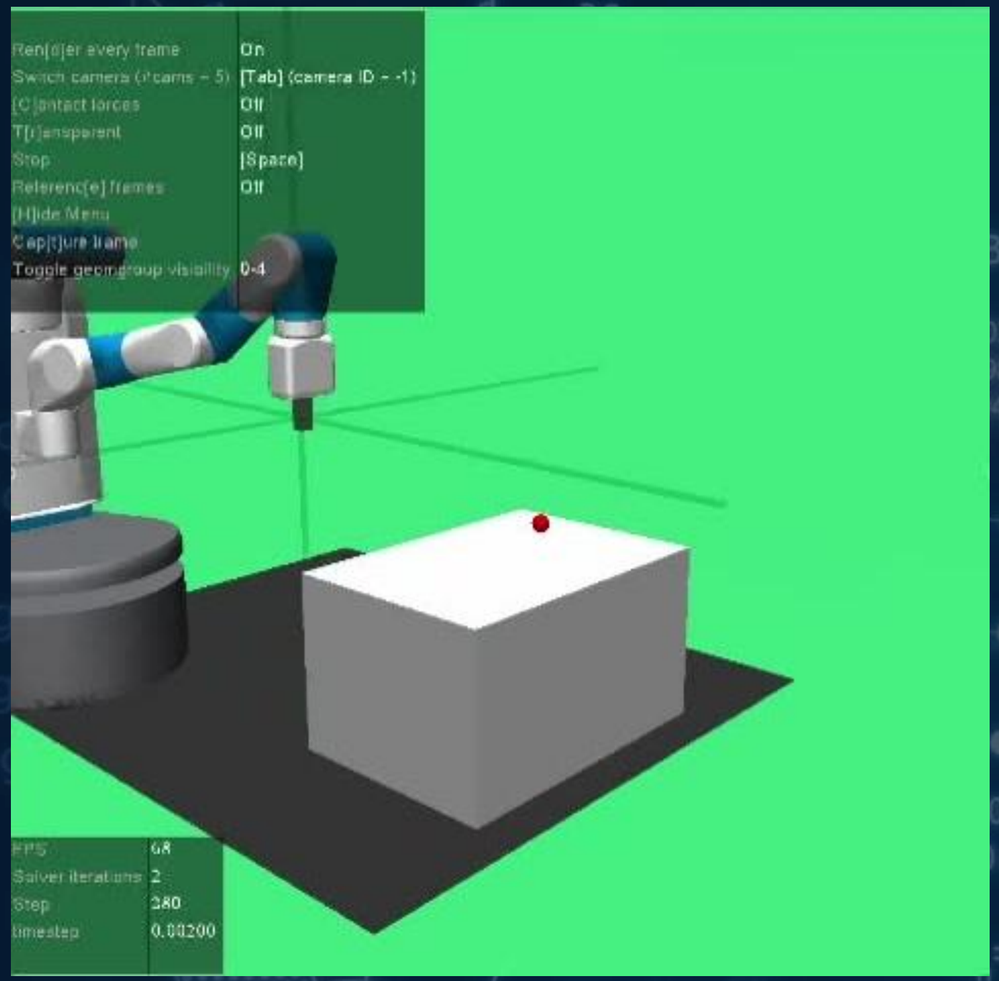


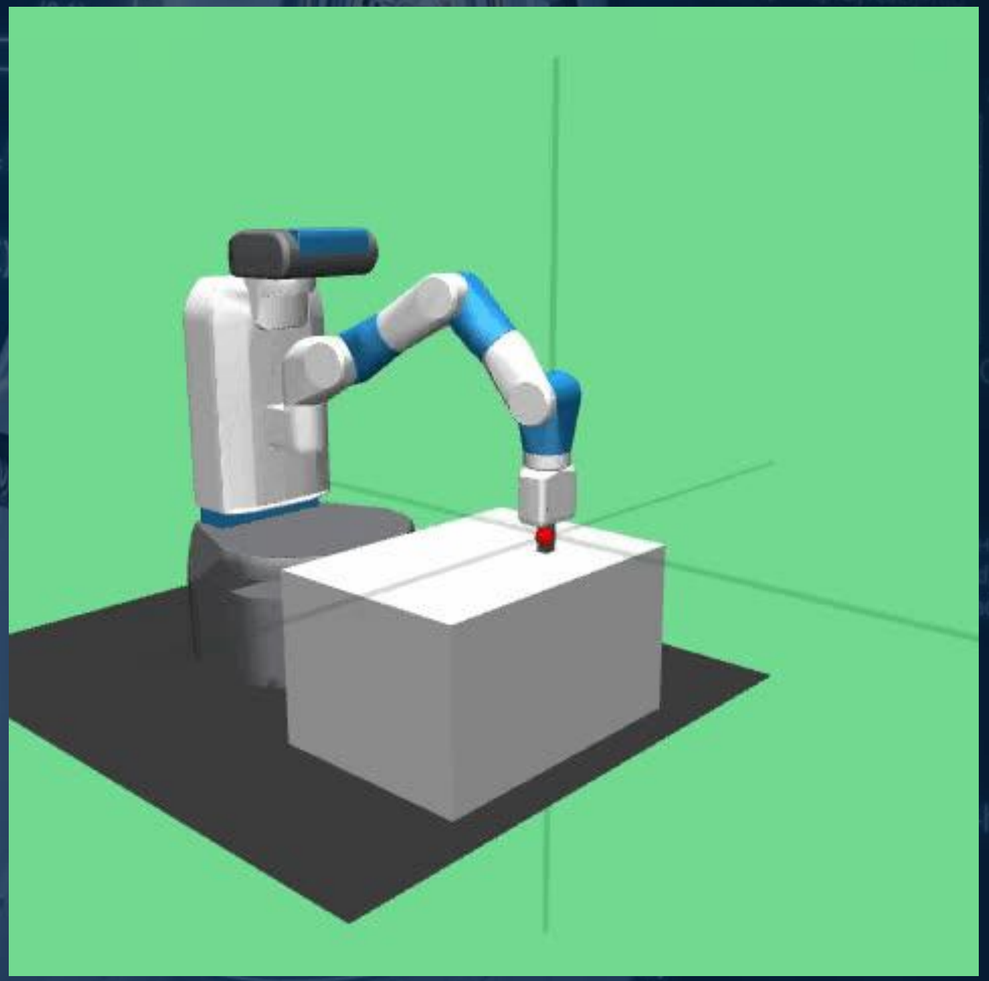

# Evaluation

Task Success Rate - Measure how many tasks were successfully completed

Learning Efficiency - How well is the model performing and the results when the model is applied to the robot.

Additionally, Subgoal Quality ( how well are subgoals generated ) and Human Feedback Quality ( effectiveness of particular human inputs )

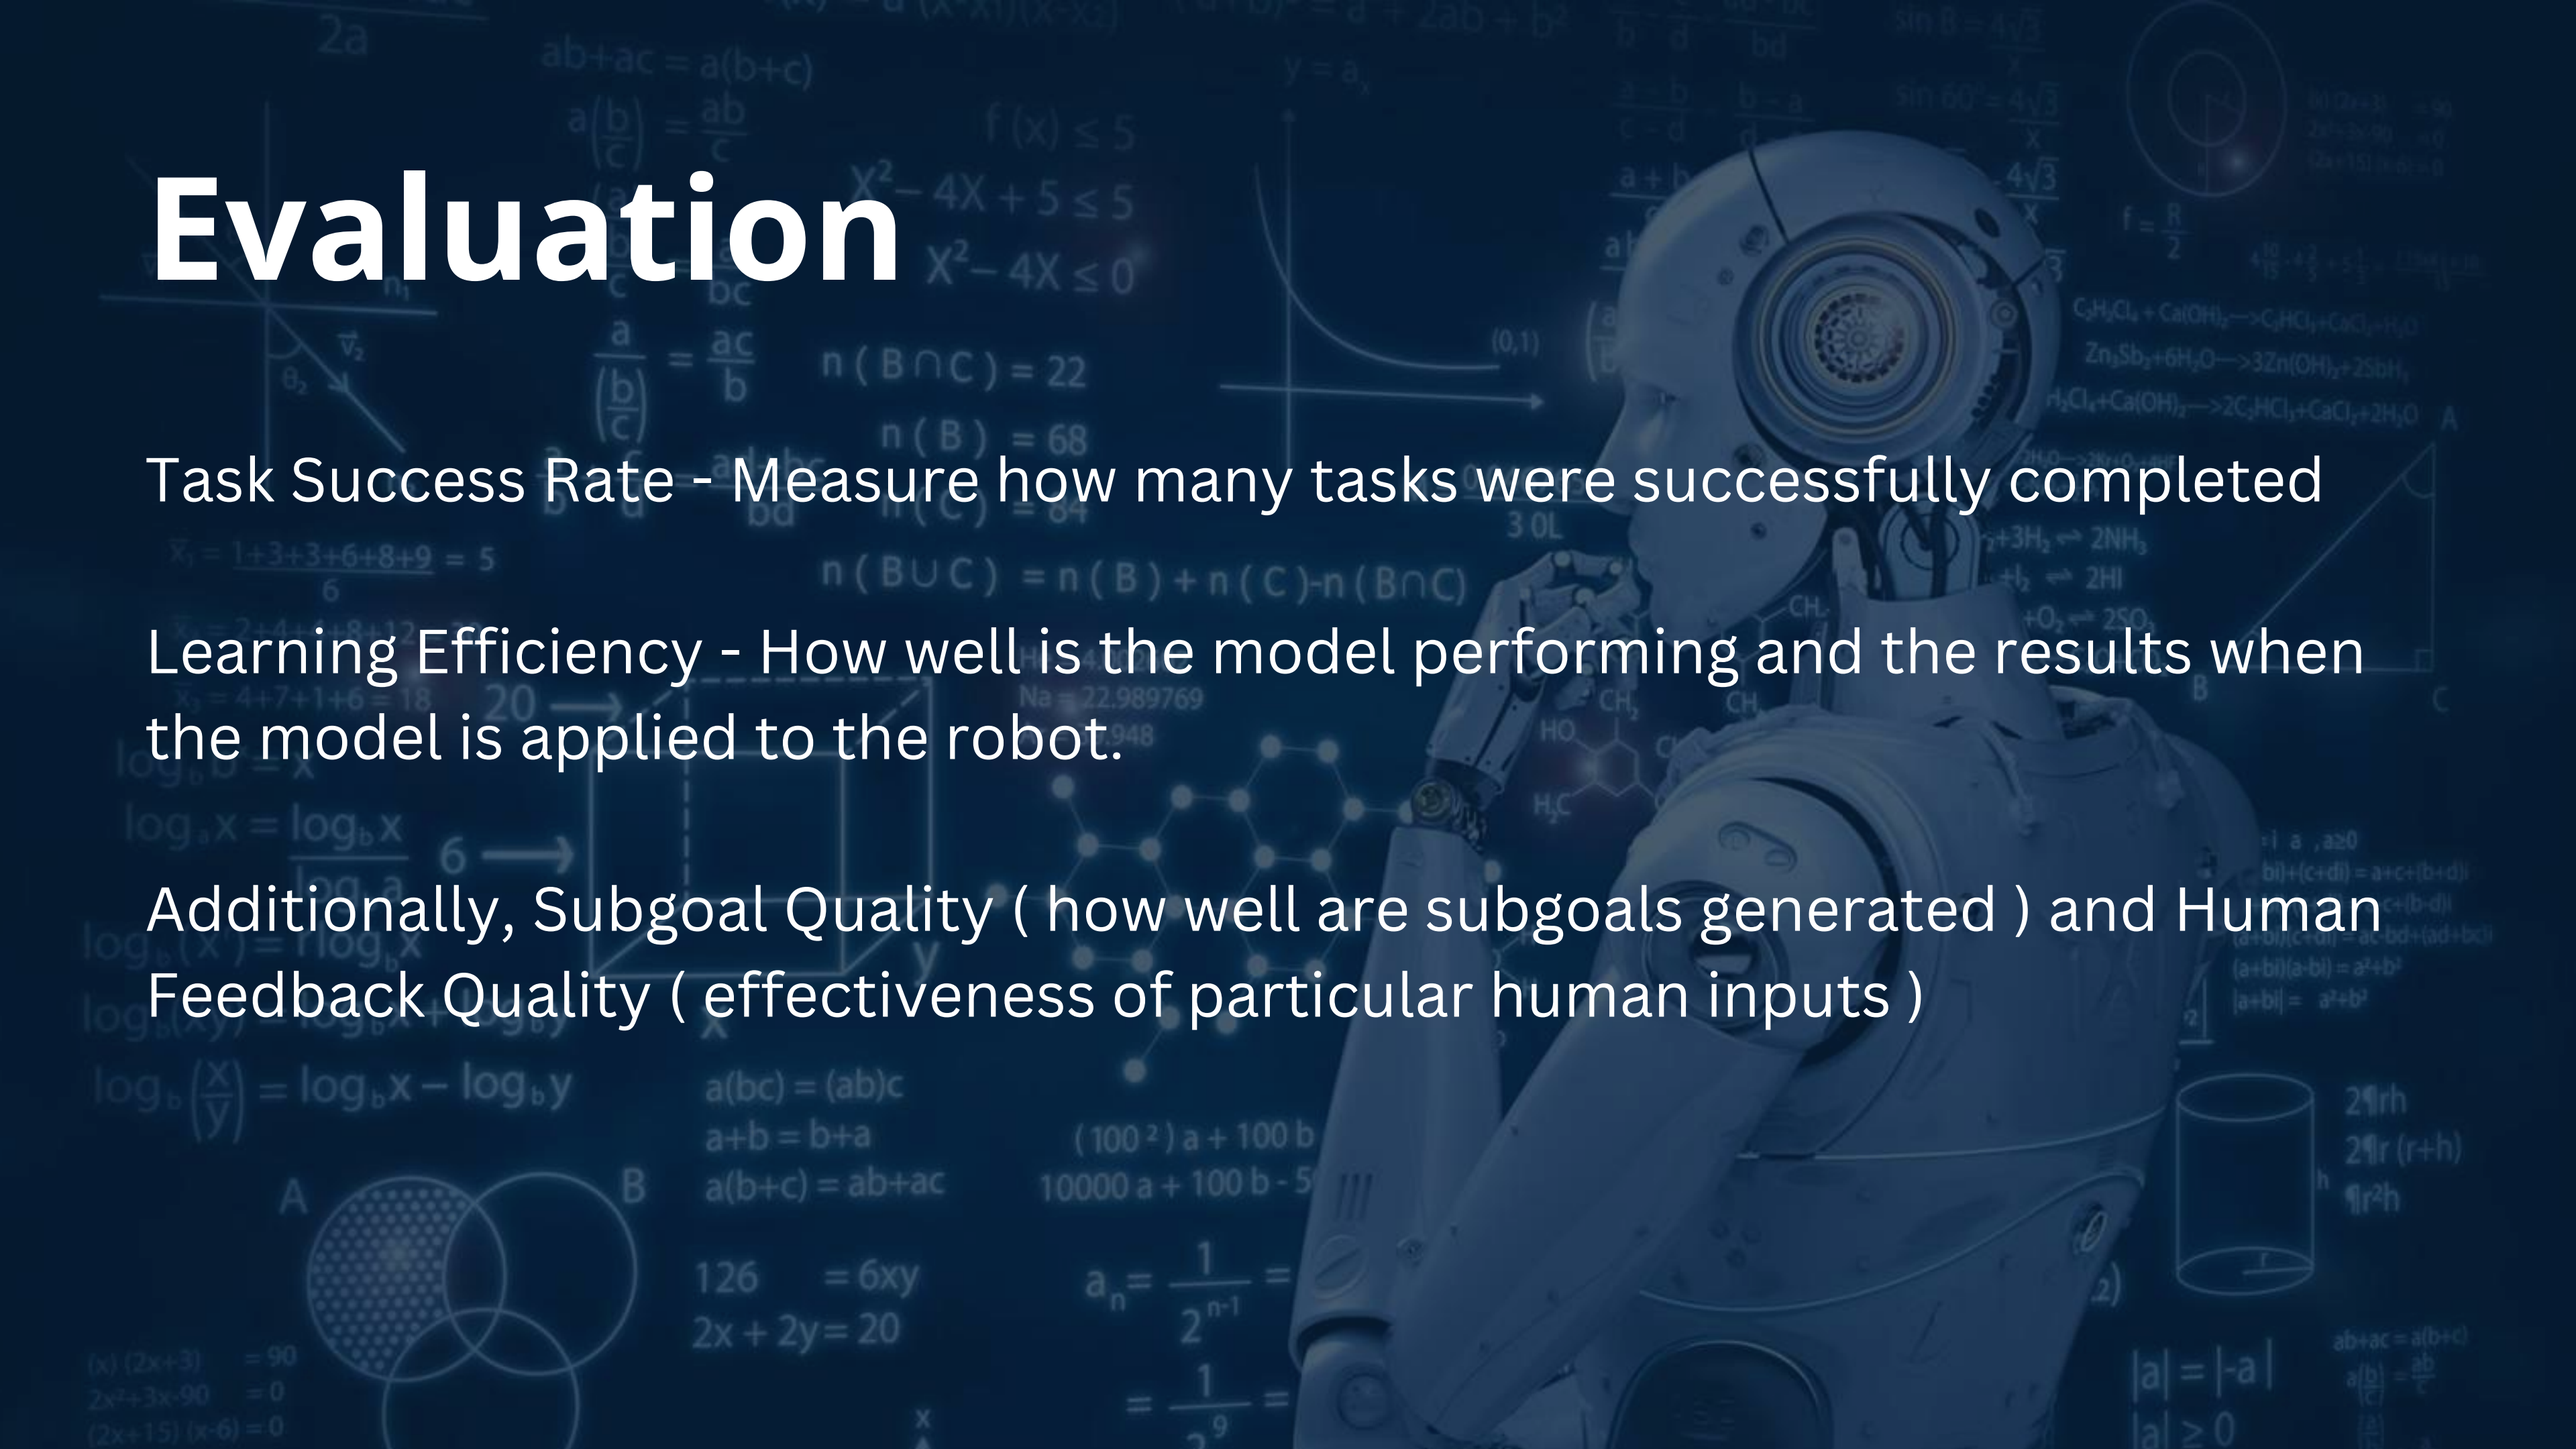

# Thank You !