

VOYAGER: An Open-Ended Embodied Agent with Large Language Models

Presenter: Yuqi Xie
10/31/2023

What is necessary for an embodied agent?

World Knowledge



Continuous Exploration

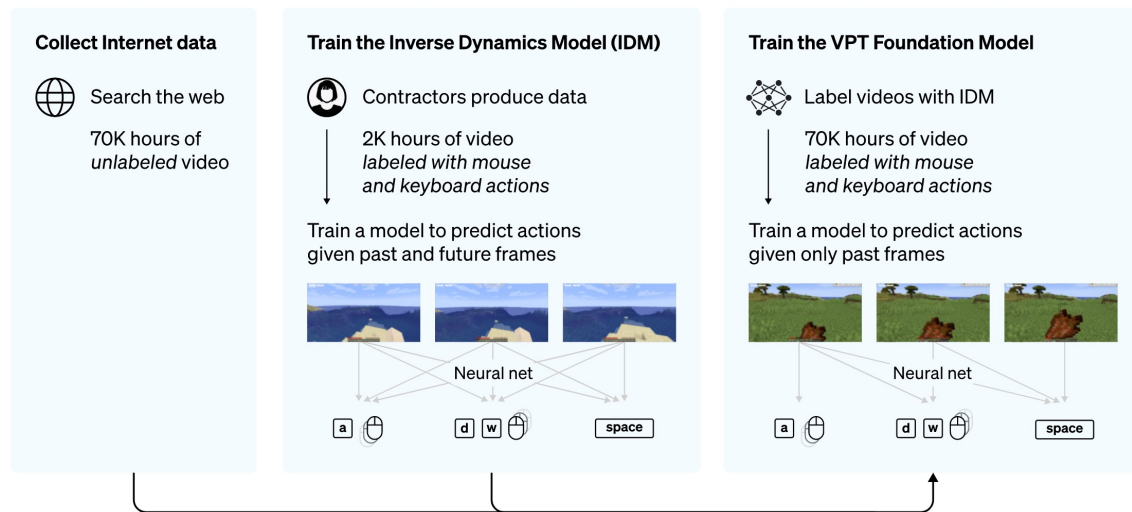


Open-Ended Objectives



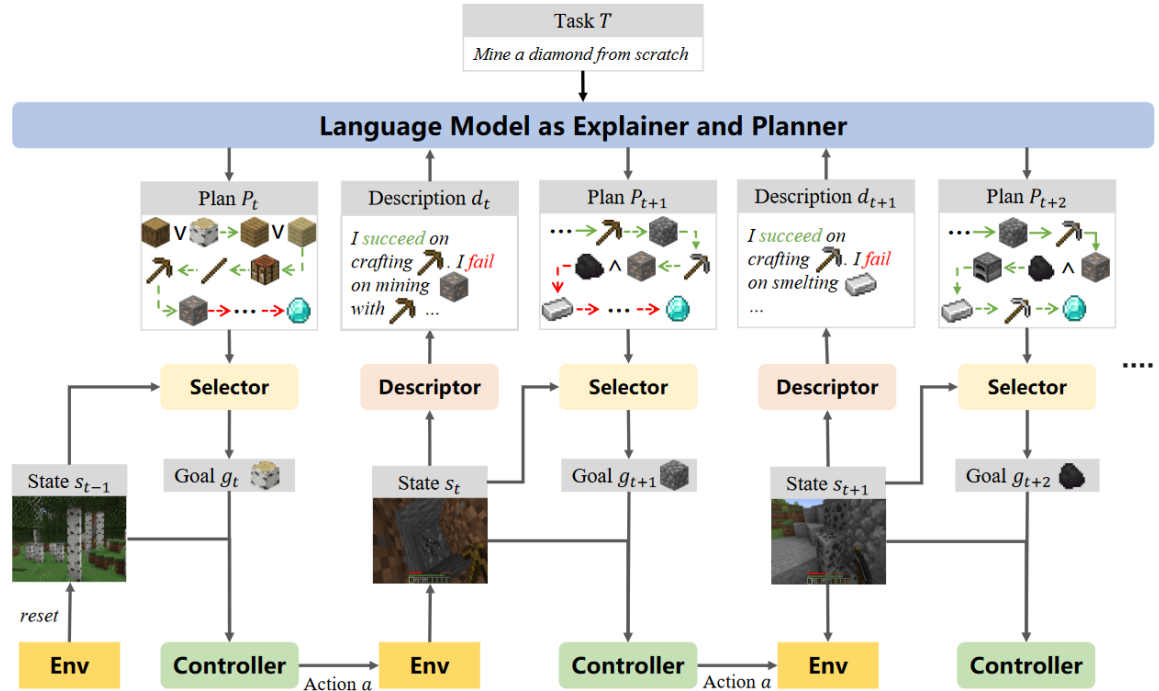
Video-Pretraining (VPT)

- OpenAI VPT: Extract policies from human-play videos
 - Use keyboard and mouse control
 - Internet-scale data
 - Continuous exploration
 - Do not have any objectives

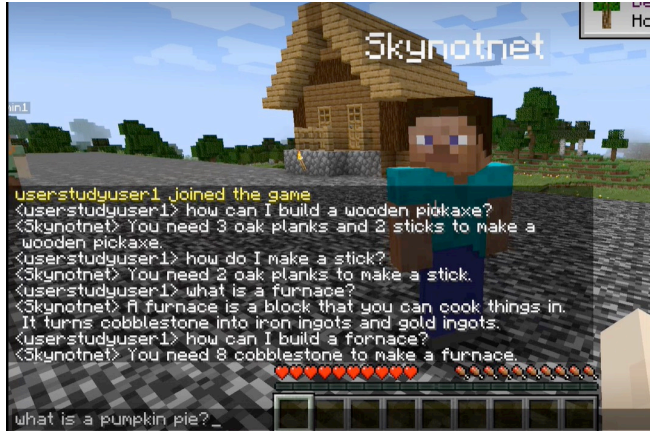


Describe, Explain, Plan and Select

- DEPS: Use LLM as explainer and planner
 - Task driven
 - Top-down decomposition
 - Not continuous



Minecraft AI NPC



In game Q&A



Task Completion using code

Finish task with supervision



Why not let the agent propose the task itself?

Automatic Curriculum

Curriculum
(GPT-4)



Wood
Log



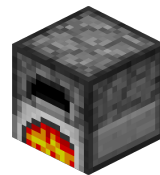
Crafting
Table



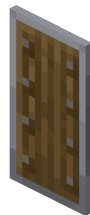
Wooden
Pickaxe



Stone
Sword



Make
Furnace



Make
Shield



Hunt
Cow



Cook
Steak



Combat
Zombie



Make
Clock



Diamond
Sword



Beat
Enderman



Automatic Curriculum



Inventory (5/36): {'oak_planks': 3, 'stick': 4, 'crafting_table': 1, 'stone': 3, 'wooden_pickaxe': 1}

GPT-4

Reasoning: Since you have a wooden pickaxe and some stones, it would be beneficial to upgrade your pickaxe to a stone pickaxe for better efficiency.
Task: Craft 1 stone pickaxe.



Biome: river
Inventory (4/36): {'sugar_cane': 3, 'paper': 3, 'fishing_rod': 1, 'wooden_pickaxe': 1}

GPT-4

Reasoning: Since you have a fishing rod in your inventory and you are near a river biome, it would be a good opportunity to catch some fish for food and experience.
Task: Catch 1 fish.



Nearby entities: pig, cat, villager
Health: 12/20
Hunger: 0/20

GPT-4

Reasoning: Your hunger is at 0, which means you need to find food to replenish your hunger. Since there are pigs nearby, you can kill one to obtain raw porkchops.
Task: Kill 1 pig.



Inventory (6/36): {'furnace': 1, 'stone_pickaxe': 1, 'oak_planks': 7, 'crafting_table': 1, 'raw_iron': 4, 'coal': 1}

GPT-4

Reasoning: You have raw iron and coal, and you have a furnace. It's time to smelt the iron to make iron ingots, which can be used to craft better tools and armor.
Task: Smelt 4 raw iron.



Time: night
Nearby entities: zombie
Equipment: 'stone_sword', 'shield'

GPT-4

Reasoning: Since it's night and there's a zombie nearby, it's a good opportunity to try killing the zombie now that you have a stone sword and a shield equipped.
Task: Kill 1 zombie.

Action-Critic Loop Overview



Code as Action Space



Environment Feedback
Execution Errors

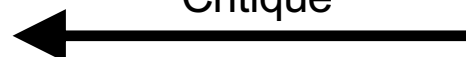


Agent's
State



Critic (GPT-4)

Critique



Self-Critic Success Check

Inventory (10/36): {'raw_copper': 9, 'copper_ingot': 3, 'acacia_planks': 1, 'raw_iron': 1, 'stick': 1, 'iron_sword': 1, 'iron_pickaxe': 1, 'iron_ingot': 3, 'crafting_table': 1, 'furnace': 1}
Task: Craft a spyglass

GPT-4

Reasoning: To craft a spyglass, you need 2 copper ingots and 1 amethyst shard. You have 3 copper ingots, but you don't have any amethyst shards.
Success: False
Critique: Find and mine an amethyst shard underground.

Inventory (7/36): {'oak_log': 5, 'oak_planks': 1, 'wooden_pickaxe': 1, 'wooden_sword': 1, 'porkchop': 2, 'white_wool': 2, 'mutton': 6}
Task: Kill 3 sheep

GPT-4

Reasoning: You have 2 white_wool and 6 mutton in your inventory, which indicates that you killed 2 sheep. You needed to kill 3 sheep.
Success: False
Critique: Find and kill one more sheep to complete the task.

Inventory (9/36): {'string': 4, 'coal': 1, 'rotten_flesh': 1, 'iron_sword': 1, 'furnace': 1, 'dirt': 6, 'stone_shovel': 1, 'wooden_pickaxe': 1, 'granite': 5}
Task: Kill 1 zombie

GPT-4

Reasoning: You have 1 rotten_flesh in your inventory, which means you have killed at least 1 zombie.
Success: True

Environment and Error Feedback

Environment Feedback

I cannot make stick because I need: 2 more planks
I cannot make stone_shovel because I need: 2 more stick

Actor
(GPT-4)

```
async function craftStoneShovelWithTable(bot) {  
  // If not enough cobblestone, mine cobblestone  
  ...  
+ // If not enough sticks, check if there are  
+ // enough planks in the inventory  
  if (sticksCount < 2) {  
+   const planksCount = bot.inventory.count(  
+     mcData.itemsByName.acacia_planks.id);  
+   if (planksCount < 2) {  
+     // Collect or craft planks  
+     await mineBlock(bot, "acacia_log", 1);  
+     await craftItem(bot, "acacia_planks", 1);  
+   }  
+   // Craft sticks using planks  
    await craftItem(bot, "stick", 1);  
    await craftItem(bot, "stone_shovel", 1);  
  }  
}
```

Execution Error

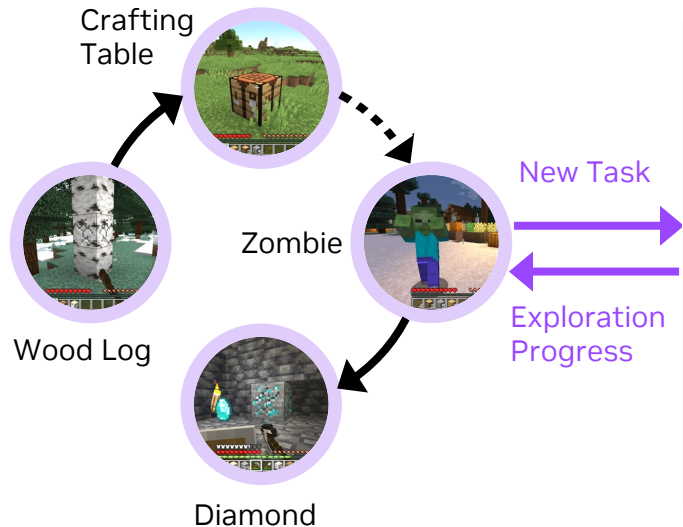
throw new Error(`No item named \${name}`);
No item named acacia_axe
at line 18:await craftItem(bot, "acacia_axe", 1);

Actor
(GPT-4)

```
-async function craftAcaciaAxe(bot) {  
+async function craftWoodenAxe(bot) {  
  // Craft anacia planks and sticks  
  ...  
  
  // Place the crafting table near the bot  
  ...  
  
- // Craft an acacia axe using 3 acacia planks  
- // and 2 sticks  
- await craftItem(bot, "acacia_axe", 1);  
- bot.chat("Acacia axe crafted.");  
+ // Craft a wooden axe using 3 acacia planks  
+ // and 2 sticks  
+ await craftItem(bot, "wooden_axe", 1);  
+ bot.chat("Wooden axe crafted.");  
}
```

Structure Overview

Automatic Curriculum



Iterative Prompting Mechanism

```
async function combatZombie(bot) {  
  // Equip a weapon  
  if (sword) {  
    await bot.equip(sword, "hand");  
  } else {  
    await craftStoneSword(bot);  
    ...  
  }  
  // Craft and equip a shield  
  await craftShield(bot);  
  ...  
  // Look for and combat a zombie  
  ...  
}
```

Retrieve

Record

Skill Library

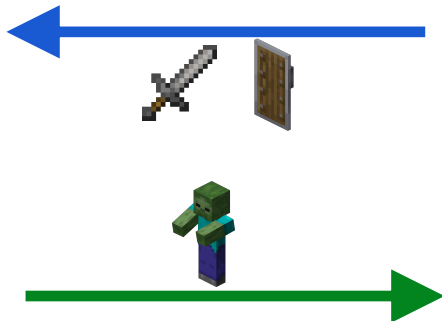
-  Mine Wood Log
-  Make Crafting Table
-  Craft Stone Sword
-  Craft Shield
-  Cook Steak
- ...
-  Combat Zombie

Skill Library

Actor (GPT-4)

```
async function combatZombie(bot) {  
  // Equip a weapon  
  const sword = bot.inventory.findInventoryItem(  
    mcData.itemsByName["stone_sword"].id  
  );  
  if (sword) {  
    await bot.equip(sword, "hand");  
  } else {  
    await craftStoneSword(bot);  
    ...  
  }  
  // Craft and equip a shield  
  await craftShield(bot);  
  ...  
  // Look for and combat a zombie  
  ...  
}
```

Retrieve old skill



Record new skill

Skill Library



Mine Wood Log



Make Crafting Table



Craft Stone Sword



Craft Shield



Cook Steak

...



Combat Zombie

Adding new skills

Program Generated by GPT-4

```
async function combatZombie(bot) {  
  // Equip a weapon  
  const sword = bot.inventory.findInventoryItem(  
    mcData.itemsByName["stone_sword"].id  
  );  
  if (sword) {  
    await bot.equip(sword, "hand");  
  } else {  
    await craftStoneSword(bot);  
    ...  
  }  
  // Craft and equip a shield  
  await craftShield(bot);  
  ...  
  // Look for and combat a zombie  
  ...  
}
```

GPT-3.5

Program Description

```
async function combatZombie(bot) {  
  // The function is about  
  equipping a stone sword to combat  
  a zombie. If a stone sword is not  
  found, it will craft one.  
  Additionally, it crafts and equips  
  a shield for added protection.  
  Afterwards, it actively searches  
  for a zombie and engages in combat  
  with it.  
}
```

Embedding

Key

Value

Skill Library



Mine Wood Log



Make Crafting Table



Craft Stone Sword



Craft Shield



Cook Steak

...



Combat Zombie

Retrieving skills

New Task:
Craft Iron Pickaxe

How to craft
an iron pickaxe
in Minecraft?

GPT-3.5

Query Content

To craft an iron pickaxe, you need to 3 iron ingots and 2 sticks. Once you have gathered the materials, ...

**Environment
Feedback**

Querying Skill Library

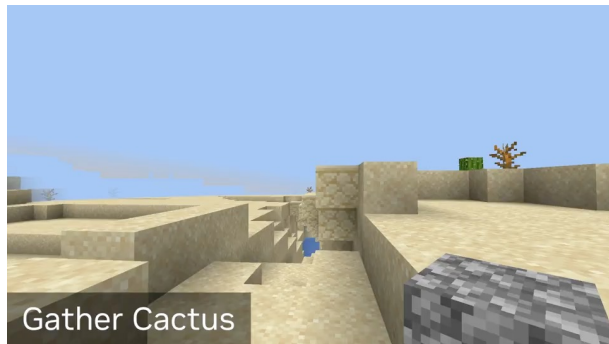
Embedding

Top-5 Relevant Skills

-  Smelt Iron Ingot
-  Craft Stick
-  Make Crafting Table
-  Make Furnace
-  Craft Wooden Pickaxe



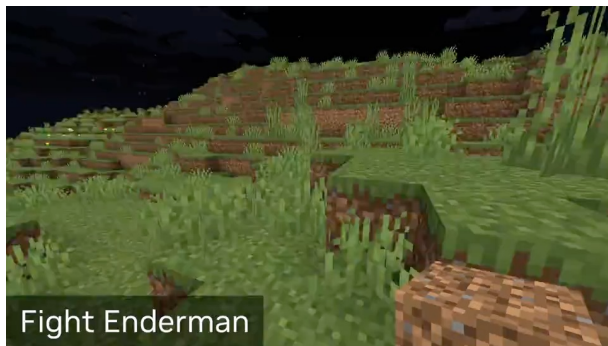
Mine Amethyst



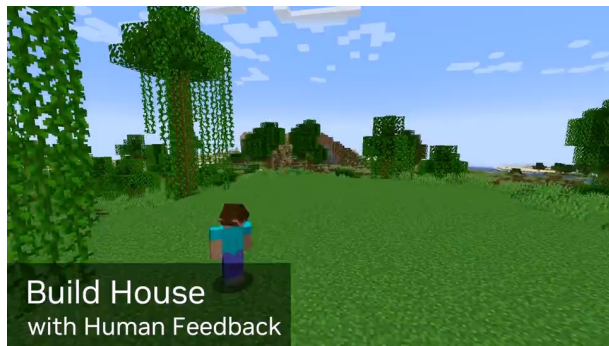
Gather Cactus



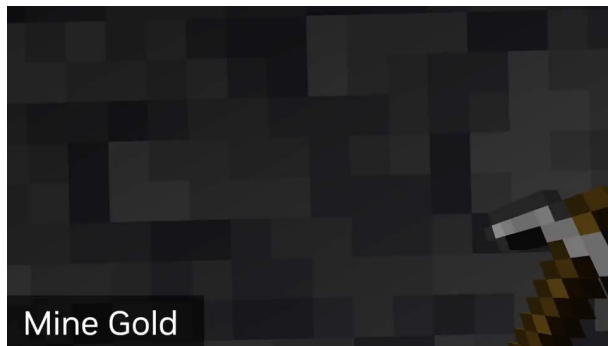
Build Base



Fight Enderman



Build House
with Human Feedback



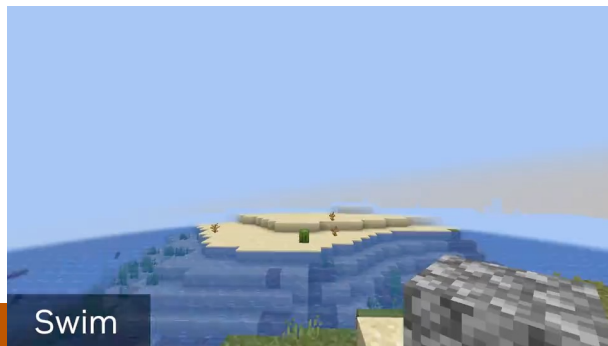
Mine Gold



Build Nether Portal
with Human Feedback



Hunt Pig



Swim

Baseline: ReAct

- Reason before Act
 - Just like “Let’s think step by step”
- Experiment setting:
 - No curriculum
 - No error feedback
 - No critic
 - No skill library

(2b) ReAct (Reason + Act)

Act 1: **Think**[First I need to find a pepper shaker...more likely to appear in cabinets (1-6), countertops (1-3), ...]

Act 2: Go to cabinet 1

Obs 2: On the cabinet 1, you see a vase 2.

(Here the agent go to cabinet 1, then cabinet 2, 3, then countertop 1 and 2)

Act 6: Go to countertop 3

Obs 6: On the countertop 3, you see a apple 1, a bread 1, a pepper shaker 1, and a vase 3.

Act 7: Take pepper shaker 1 from countertop 3

Obs 7: You pick up the pepper shaker 1 from the countertop 3.

Act 8: **Think**[Now I find a pepper shaker 1. Next, I need to put it in/on drawer 1.]

Act 9: Go to drawer 1

Obs 9: Drawer 1 is closed.

Act 10: Open drawer 1

Obs 10: You open Drawer 1 ...

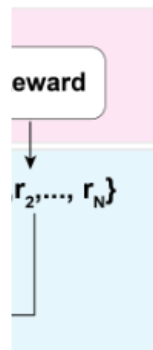
Act 11: Put pepper shaker 1 in/on drawer 1

Obs 11: You put pepper shaker 1 in/on the drawer 1.



Baseline: Reflexion

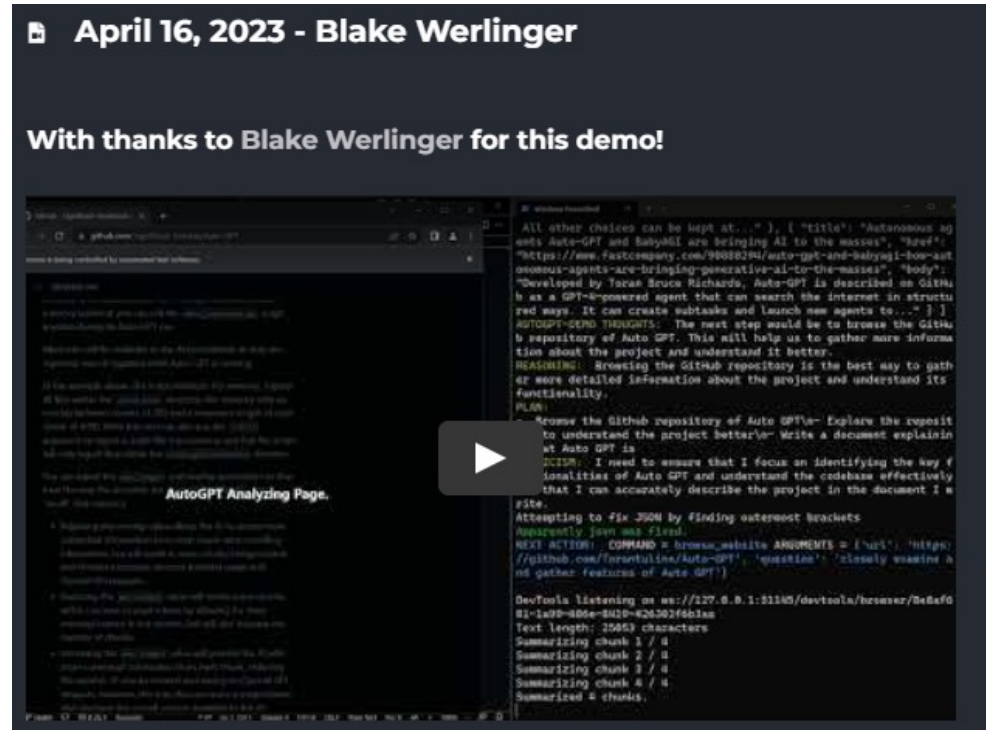
- A > go to desk 1
Nothing happens.
> go to desk 2
On the desk 2, you see a alarmclock 1, a bowl 1, a mug 3, a pencil 3, and a pencil 2.
> think: Now I find a desklamp (1). Next, I need to use it.
OK.
> use desklamp 1
Nothing happens.
> use desklamp 1
- E Nothing happens.
Status: Fail



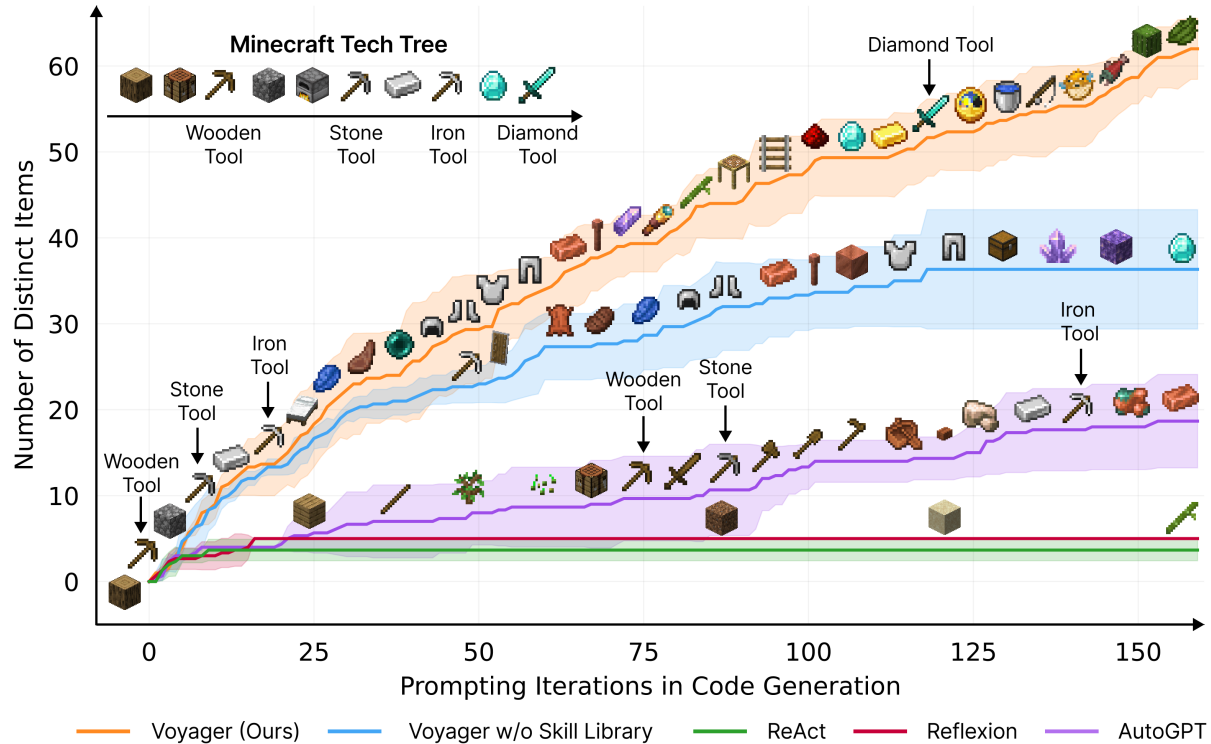
Reflection: In this environment, my plan was to find a mug then find and use a desklamp. However, the task says to examine the mug with the desklamp. I should have looked for the desklamp first, then looked for the mug. I noticed that the desklamp was found on desk 1. In the next trial, I will go to desk 1, find the lamp, then look for the mug and examine it with the desklamp.

Baseline: AutoGPT (April 16th)

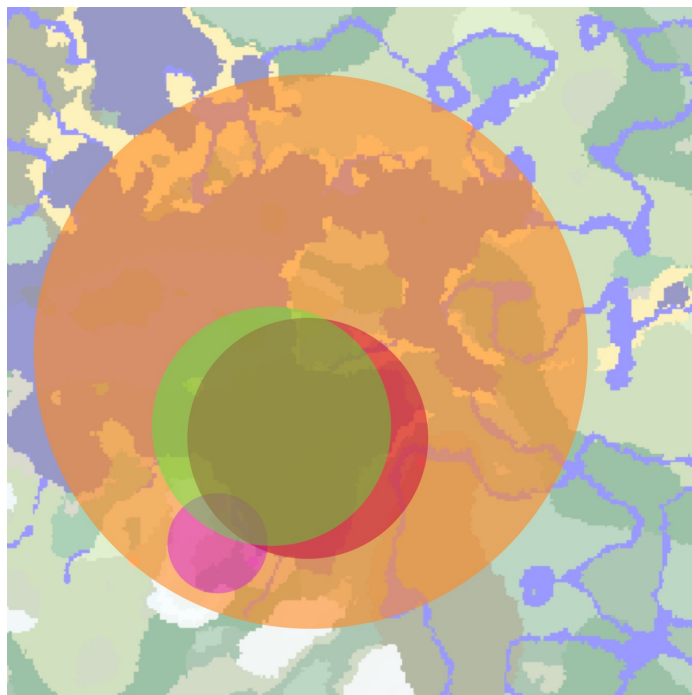
- AI agent for complex tasks
- Experiment setup:
 - Task decomposition
 - Replanning over iterations
 - No curriculum
 - No critic



Open-Ended Exploration

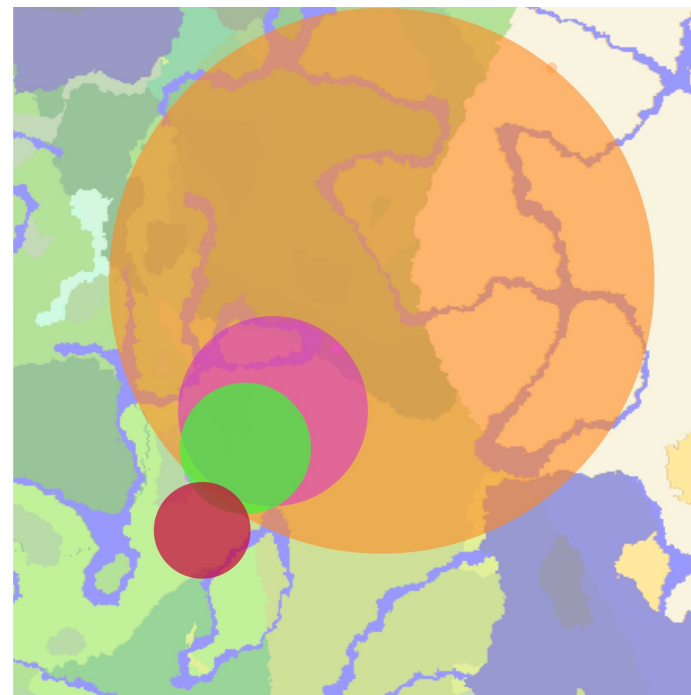


Map Coverage



—●— Voyager (Ours)

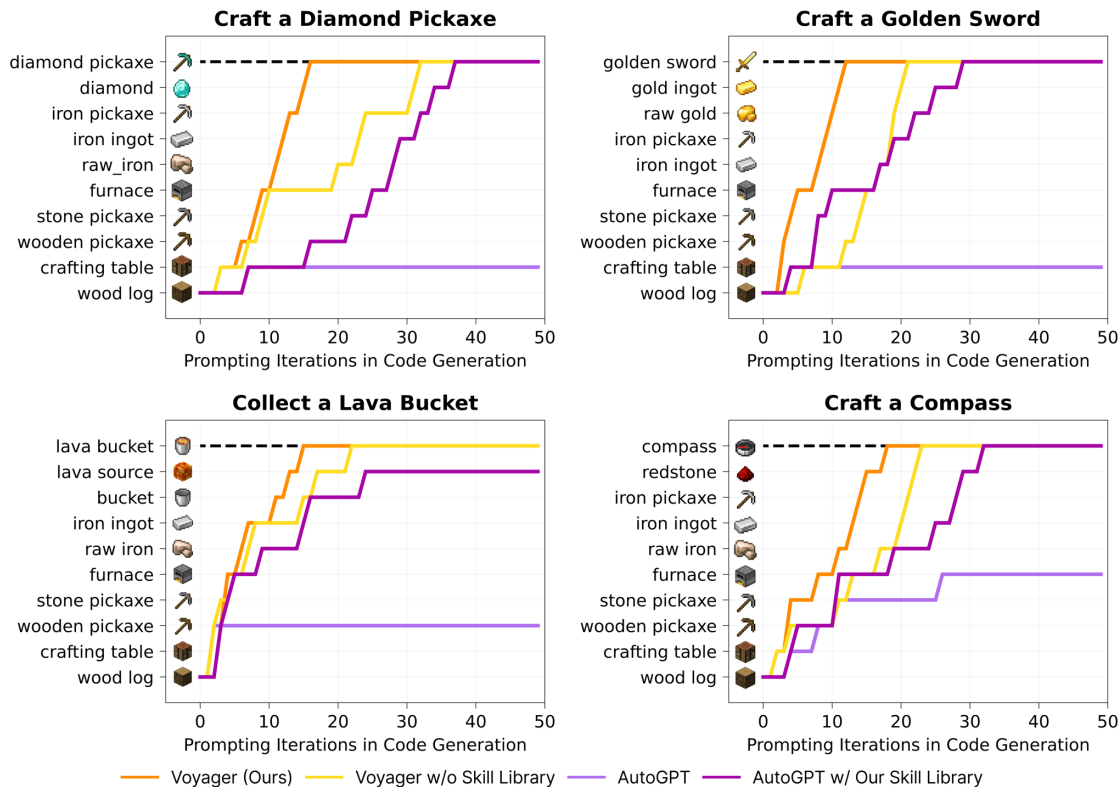
—●— ReAct



—●— Reflexion

—●— AutoGPT

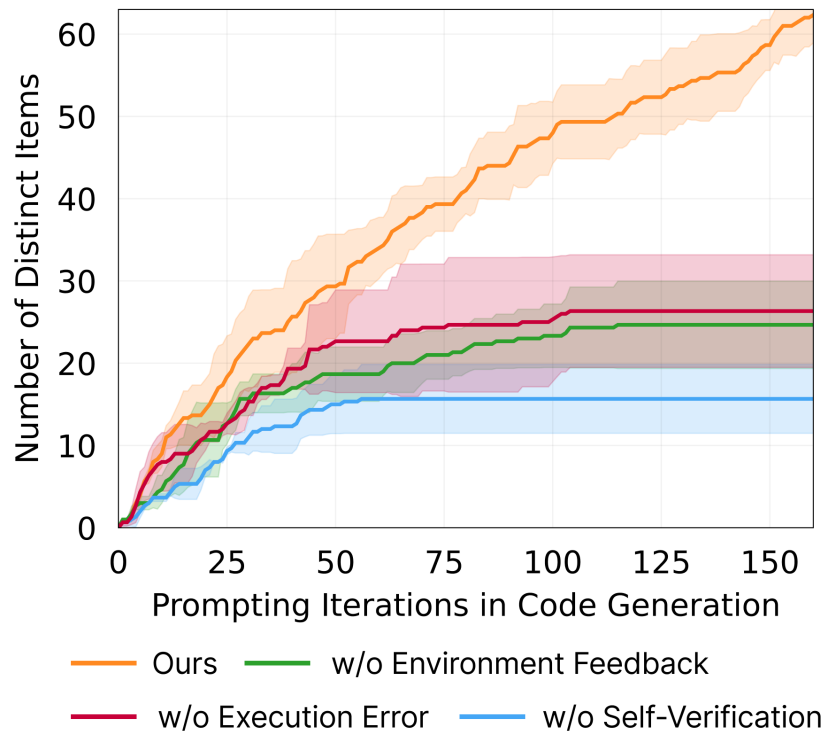
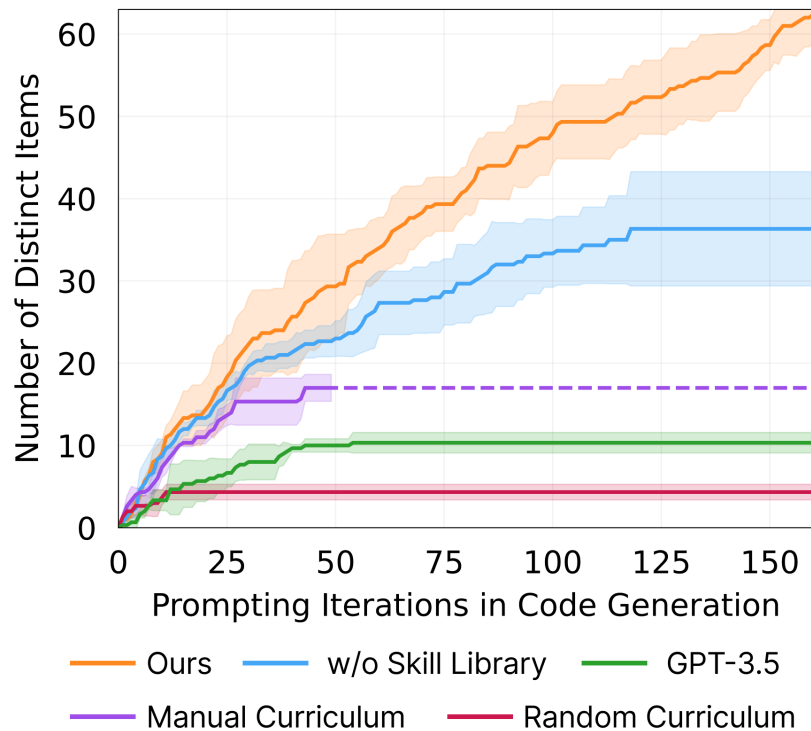
Zero-shot Generalization to Unseen Tasks



Human as Multimodal Feedback



Ablation Studies



Q&A