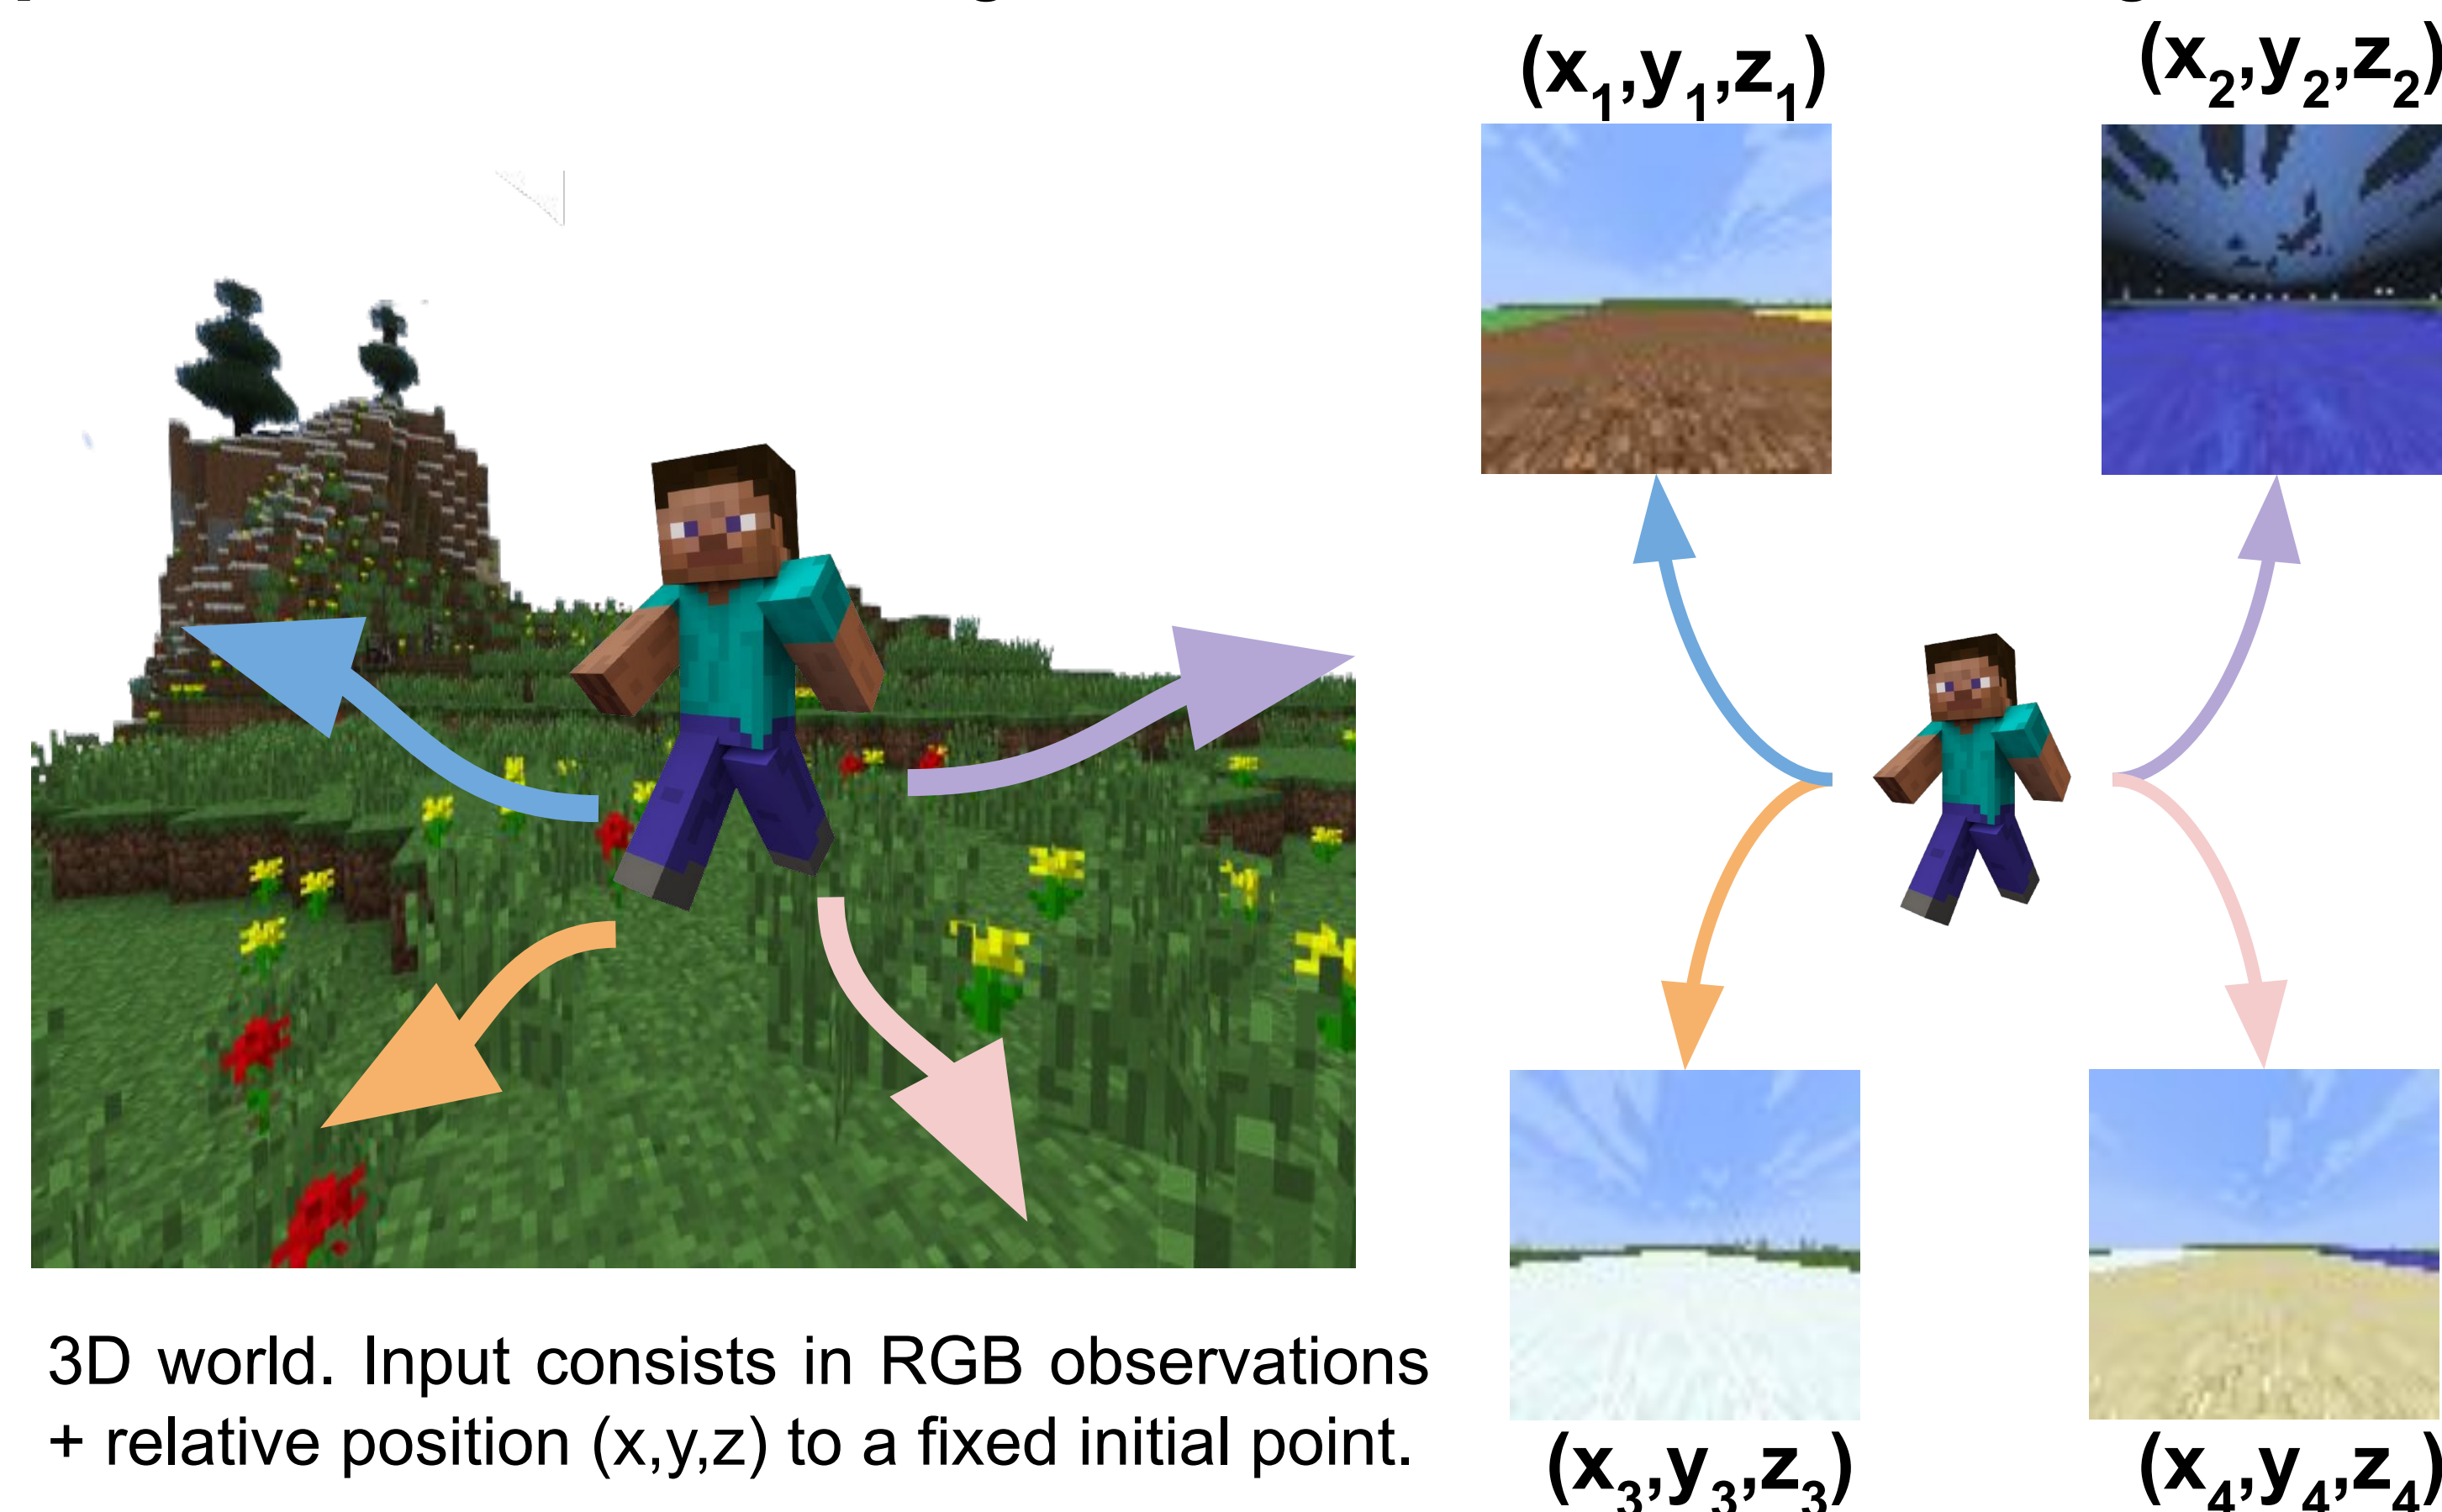


PiCoEDL: Discovery and Learning of Minecraft Navigation Goals from Pixels and Coordinates

Task description

Discover navigation goals in Minecraft by using an intrinsic motivation (**empowerment**). Then, **learn** some **conditioned policies** that drive the agent towards the discovered goals.



Motivation

EMPOWERMENT

Maximize influence over the environment.

SKILL DISCOVERY

Abstract a set of temporal actions into skills.



INFORMATION THEORETIC OBJECTIVE

Maximize skill discovery → Maximize mutual information between states and some latents

$$I(S; Z) = H(S) - H(S|Z) = H(Z) - H(Z|S)$$

Challenges

1. **Intrinsic reward** derived from the **forward form** of the Mutual Information **does not match a distance** in the environment.
2. Discover disjoint navigation goals in simple maps by taking into account **high-dimensional inputs** like pixel observations.
3. Discovering skills from pixels does not scale to realistic maps.



Navigation goals are too spread out over the map and **some latent skills dominate** over the others.

Our approach

- Based on **EDL [1]** (Explore, Discover and Learn).
- Use the **reverse** form of the Mutual Information.
 - Add agent's **relative position** respect to a initial point.

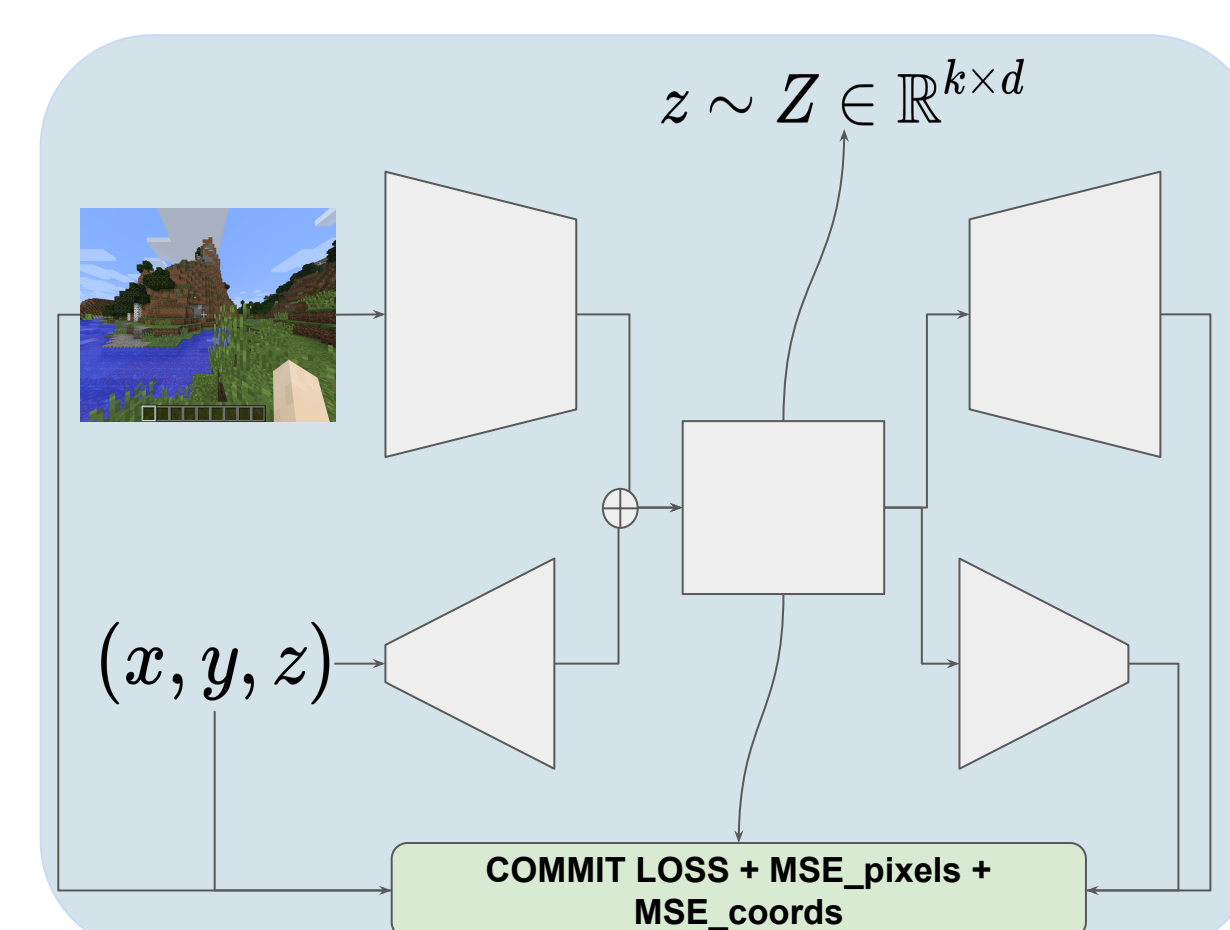
EXPLORE

GENERATE RANDOM TRAJECTORIES




DISCOVER

VARIATIONAL



LEARN



SKILL DISCOVERY

SKILL LEARNING
(one skill example)

Results

