



# Project Ideas

---

CS 175: Projects in AI (in Minecraft)

PROF: SAMEER SINGH

# Reinforcement Learning

Agent learns to do things by trying things, and succeeding/failing

Navigation


- Explore the map without dying
- Solve mazes
- Learn the best way home from anywhere
- Get to the highest hill in the map

Learn Recipes


- Figure out best way to make items
- Without any knowledge of the recipes

Combat




- Learn to hide/find shelter
- Learn to fight, [example paper](#)



(a) I-Maze



(b) Pattern Matching



# Reinforcement Learning

Agent learns to do things by trying things, and succeeding/failing



# Reinforcement Learning

---

Next few lectures will go into details (and more ideas)

For now, let's look at non-RL ideas

# Describe the Scene

---



Houses and a pig on a grassy field during the day.

Pig staring at me in a village.

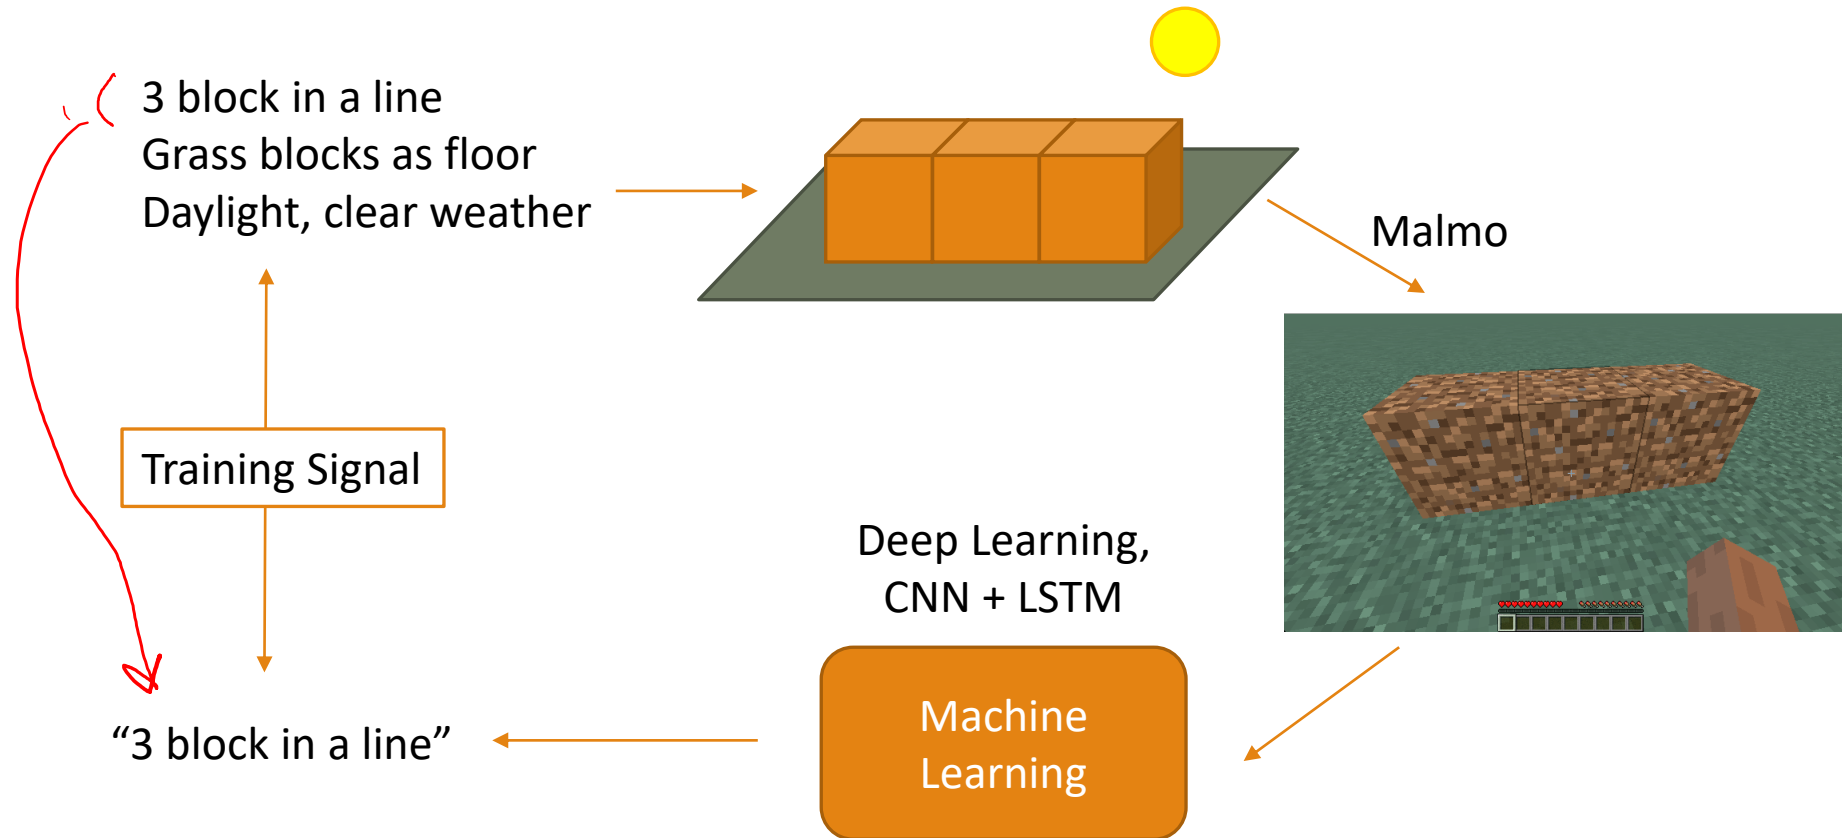
# Live Commentator

---

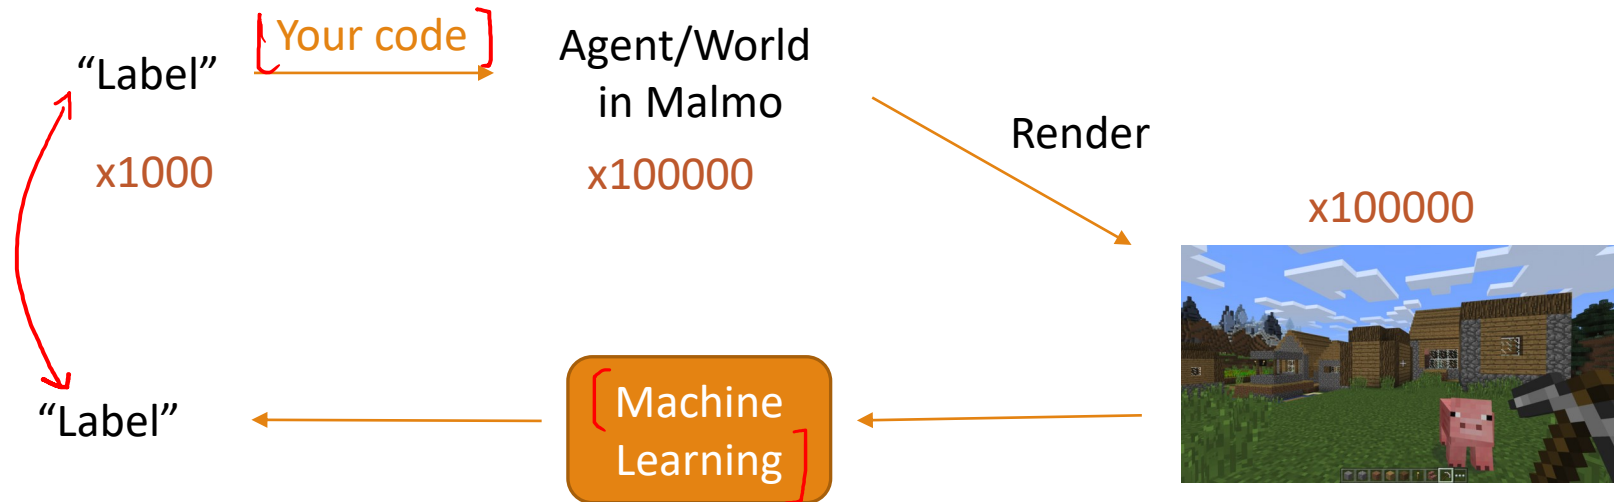


“Hit a rabbit”

# How is this even possible?



# Many Variations of These



*Label*

object

objects

action

depth of pixel

*Task*

object detection

~caption generation

~action detection, “commentary”

~stereoscopy, depth/distance prediction

# Captions to Speech

---



Pig staring at me in a village.



# Natural Language Navigation



Quite Difficult!



- > Go forward till you hit a wall
- > Go to the pig
- > Go to the house on the right
- > Go behind the house

~~trivial~~ simple



hardest

# Natural Language Interface



Quite Difficult!



- > Choose steel pickaxe and dig
- > Go and destroy that window
- > Put the blue block on the closest wall
- > Find a tree and chop it

trivial




hardest

# SHRDLU (from 1970!)

---

Person: WHAT DOES THE BOX CONTAIN?  
Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.  
Person: WHAT IS THE PYRAMID SUPPORTED BY?  
Computer: THE BOX.  
Person: HOW MANY BLOCKS ARE NOT IN THE BOX?  
Computer: FOUR OF THEM.  
Person: IS AT LEAST ONE OF THEM NARROWER THAN THE ONE WHICH I TOLD YOU TO PICK UP?  
Computer: YES, THE RED CUBE.  
Person: IS IT SUPPORTED?  
Computer: YES, BY THE TABLE.  
Person: CAN THE TABLE PICK UP BLOCKS?  
Computer: NO.  
Person: CAN A PYRAMID BE SUPPORTED BY A BLOCK?  
Computer: YES.  
Person: CAN A PYRAMID SUPPORT A PYRAMID?  
Computer: I DON'T KNOW.  
Person: STACK UP TWO PYRAMIDS.  
Computer: I CAN'T.



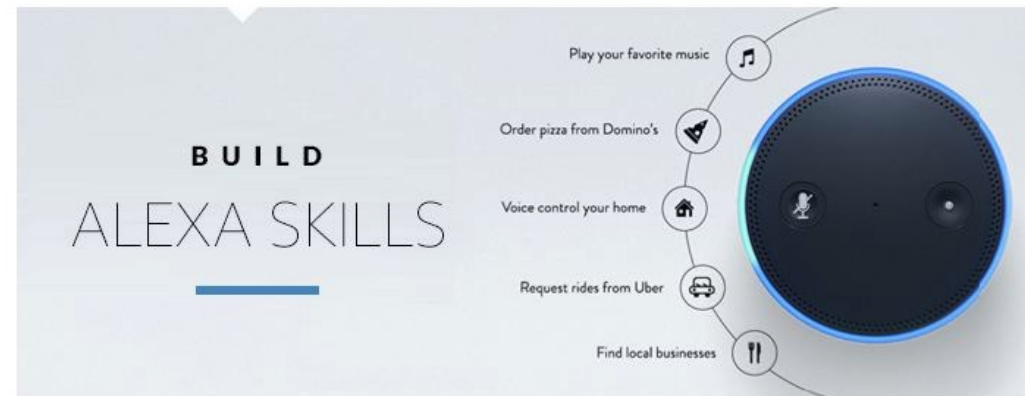
# Natural Speech to Commands

---

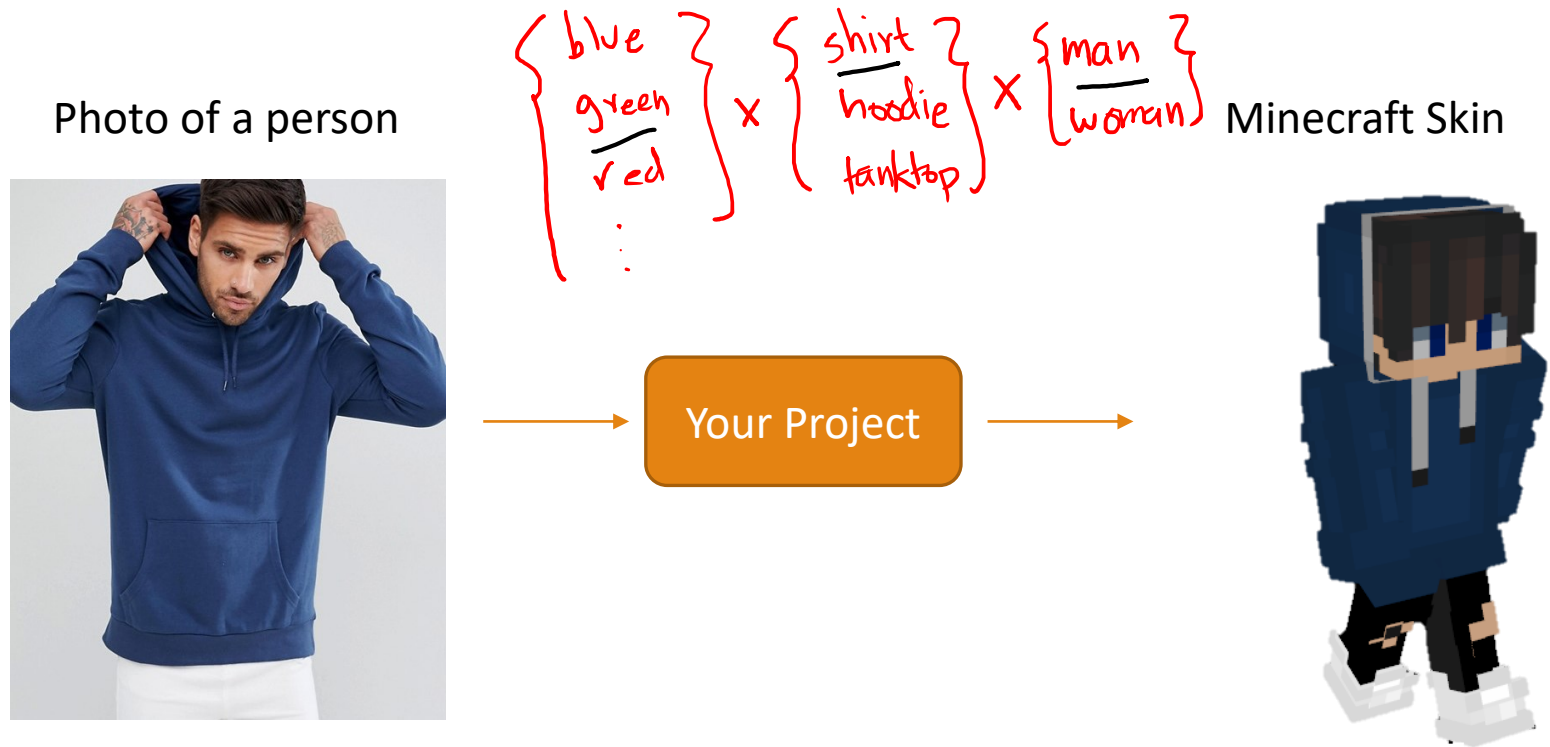


Off the shelf  
Speech to Text systems

Online Speech to Text APIs



# Photo to Minecraft Character <sup>Skin</sup>



Need to label data?  
Can you use existing classifiers, like Visual QA?

# Recipe Planners

Inventory




“Need”(s)



Steps



~~> Get 2 wood planks~~ 

~~> Make a stick~~ 


~~> Get 2 diamonds~~ 

> Make diamond sword



# Lots of other possibilities

---



Many other games in Minecraft

Create AI for those?

One AI that works for all of those?


# Try not to do...

---

Combat



Maze/Navigation



I won't disallow it, but you will have to do a better job convincing me that it is unique from previous projects, interesting to do, and likely to work



# Project Ideas

---

CS 175: Projects in AI (in Minecraft)

PROF: SAMEER SINGH