Windfallen

Darby Brown
Erik Sambrailo
River Schieberl
Tim Tung

Index

Presentation

Our Team

Darby Brown
Project Manager

Erik Sambrailo
Machine Learning

River Schieberl
Simulation

Tim Tung
Infrastructure

We're on a mission to reduce food waste

- **92 billion pounds** of food is wasted in the U.S. each year
- Field losses of edible produce are **~34% of marketed yield**
- "Windfall" fruit is a key contributor to waste on modern commercial farms

Existing Technology

- Manually operated equipment that specializes in windfalls (Zim)
- Robotic platforms that harvest fruit from the plant
    - Each platform tends to work on only a small number of specific fruits
    - Platforms often are designed to work in specialized environments
        - Indoor or greenhouse farms with rails between rows (advanced.farm)
        - Vertical wall trellis systems for orchards (Tevel)

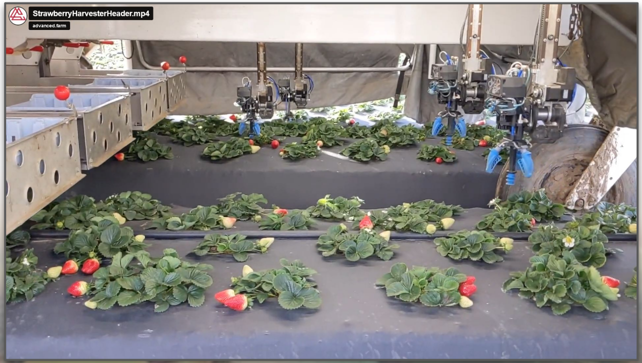

Target Customer (Primary)

Commercial Farms

- apple growers w/ 10+ acres
- orchard layout conducive to automation
- acceptable apple variety for processing
- desire to harvest windfalls
- open to implementing robotic technology

Business Model: subscription or rental model for technology deployment

Target Customer (Secondary)

Apple Processing Companies

- Juice and Cider Manufacturer’s
- Applesauce and Apple Butter Producers
- Dried Fruit Manufacturers
- Animal Feed Producers
- Distilleries
- Pectin Manufacturers

Business Model: acquiring the rights to harvest windfalls and selling direct to processors

Merging technology with agriculture to revive windfall

- Windfall loss is a very challenging problem that is often solved with manual labor or not at all.
- Extending cutting edge Multi-Agent Reinforcement Learning (MARL) research to food sustainability
- Improving yield could drive higher revenues by ~$220M in the U.S. market
- Paving the way for additional research through extensible and generalizable models and environments

Our secret sauce is APPLESAUCE

Farmers are open, and the industry enthusiastic

We interviewed:

Farmers

- **Labor too costly** to harvest windfall.
- **Timeliness is key** (3-4 day window).
- **Open to robotics** as a potential solution.

Robotics Experts

- **Multi-agent orchestration** is a burgeoning field.
- **Software-only orchestration** is a better business than E2E development.

RL Developers

- **Key research areas** include novel environments and novel training.
- **Research libraries** under active development.

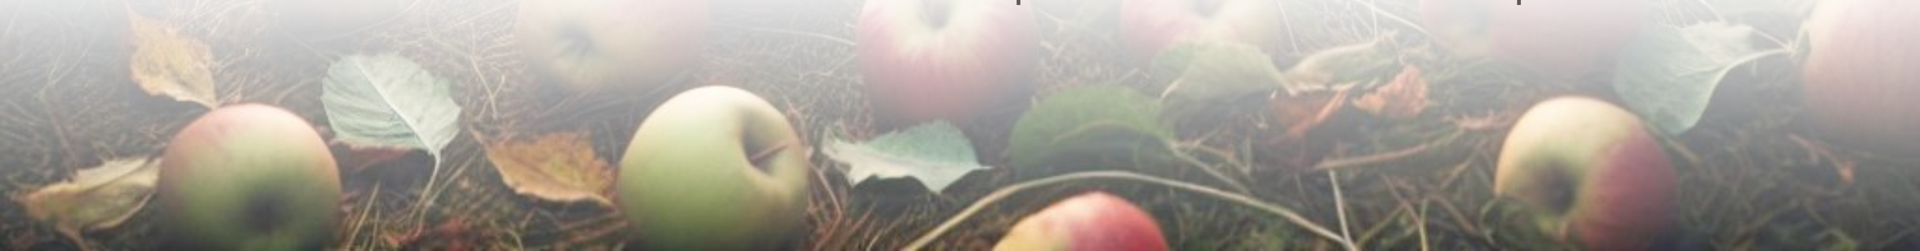

Research leads to an MVP to Test Efficacy in Agriculture Settings

Research Question:

- Can multi-agent reinforcement learning **increase the efficiency** of robotically recovering dropped apples?

Key Features of MARL Research:

- Multiple agents with different jobs
- Performant on unseen scenarios
- **Environment setup based on real-world harvesting scenarios**

MVP:

- **Trained neural network for MARL implementation in an orchard setting**
- Trained model leveraged by downstream robotic fleet owners & operators

Existing multi-agent reinforcement learning environments do not represent real world agricultural scenarios

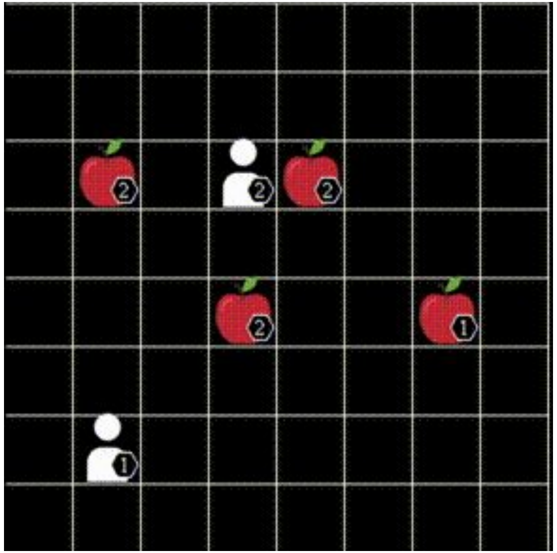

Existing Simulated Environments

Aerial View of Operating Orchards

We took field measurements from ten local California orchards to inform our environment layout

Custom environment developed for use case

- **Incorporates real-world dimensions, randomness and inconsistencies**

  Ranges for key parameters generate a variety of possible orchard layouts.

  tree_diameter = [9, 44]
  distance_btw_rows = [365, 762]
  distance_btw_trees = [150, 762]
  canopy diameter = [150, 762]
  apple_diameter= [5, 9]
  vacancy_rate = [0, 0.3] (percent)
  big_bin_dimension = 122 x 122 x 72
  small_bin_dimension = 41 x 61 x 30

- **Task-specific bots designed for minimally disruptive deployment**

  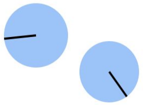

  "Picker bots" pick up one apple at a time and deposit into a bin.

  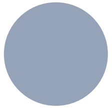

  Movable bins auto-locate to reduce distance traveled by pickers
  (on roadmap)

- **Continuous environment space**
  (not grid world)

APPLESAUCE Overview

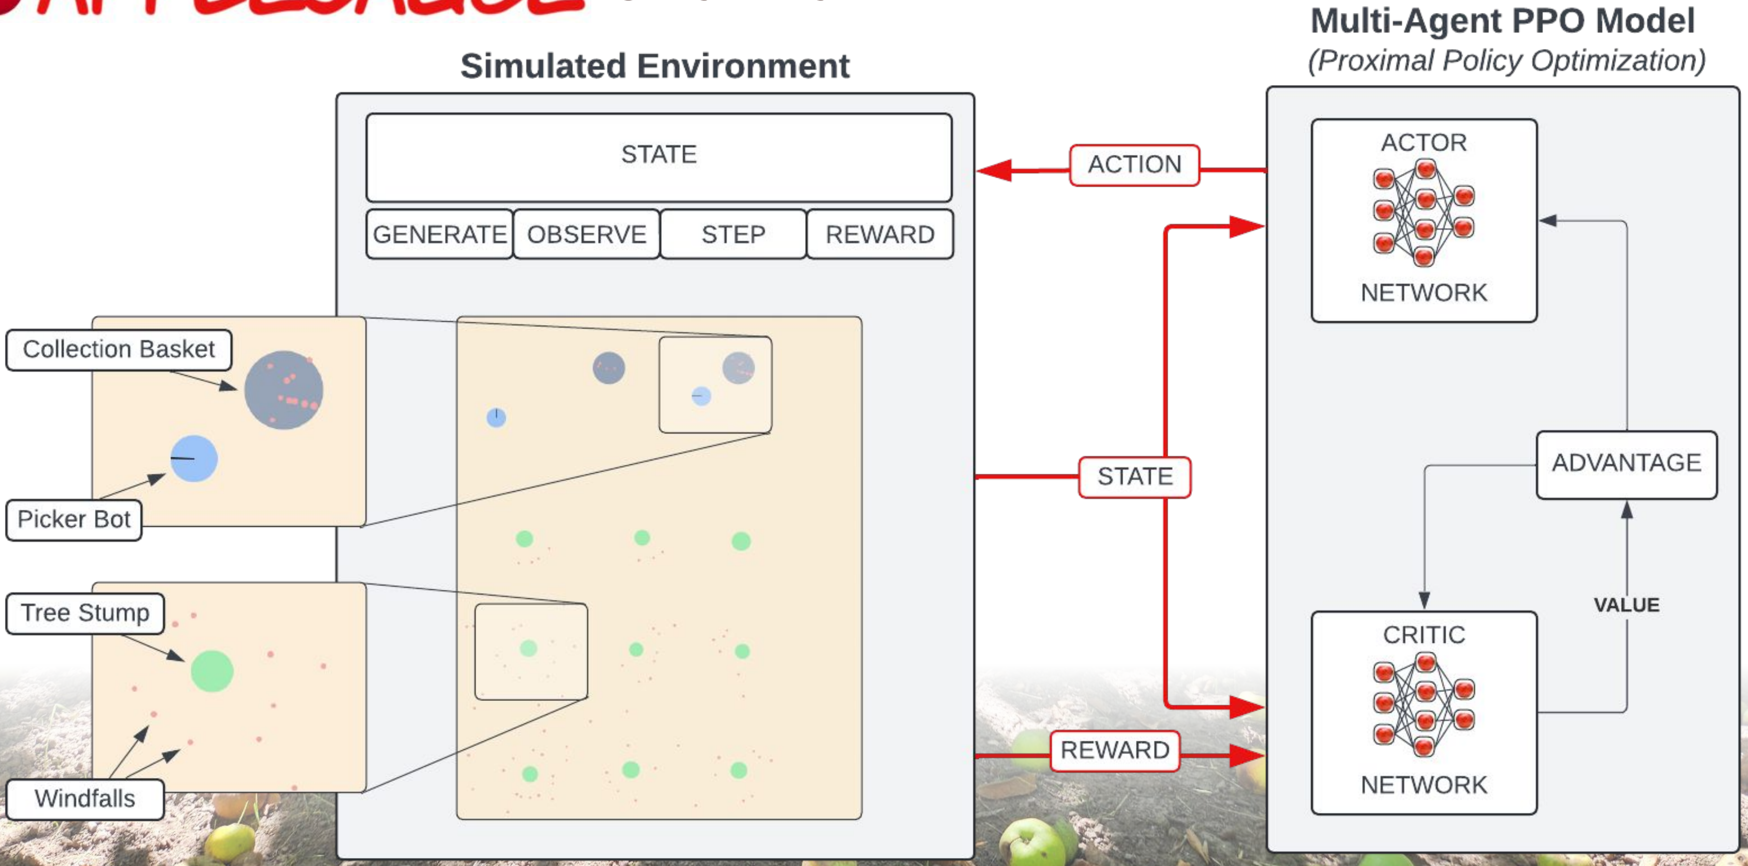

Technical Approach

- **reinforcement learning requires parallelization**
- deep study of:
    - Jumanji/Mava
    - Jax/Flax
- APPLESAUCE: 4,000 lines of code built upon InstaDeep libraries

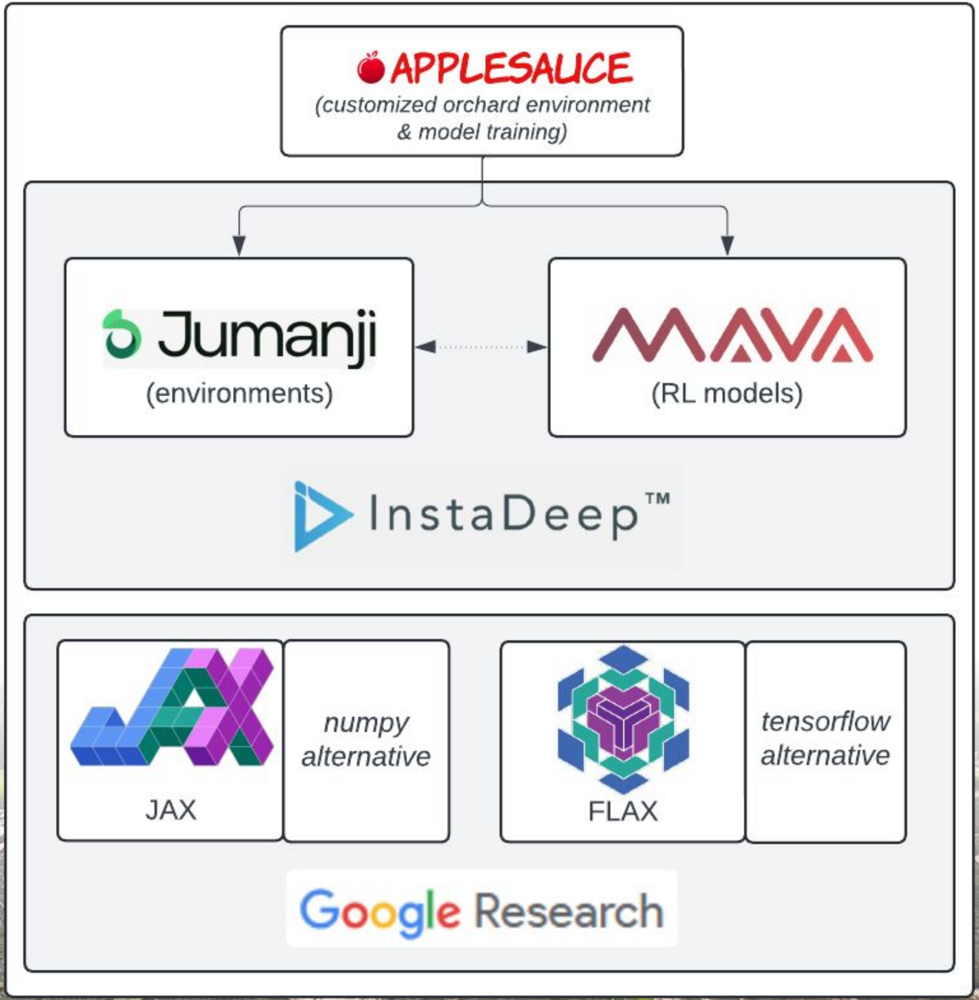

Model Demonstration

- Hundreds of simulations simultaneously = quick learning

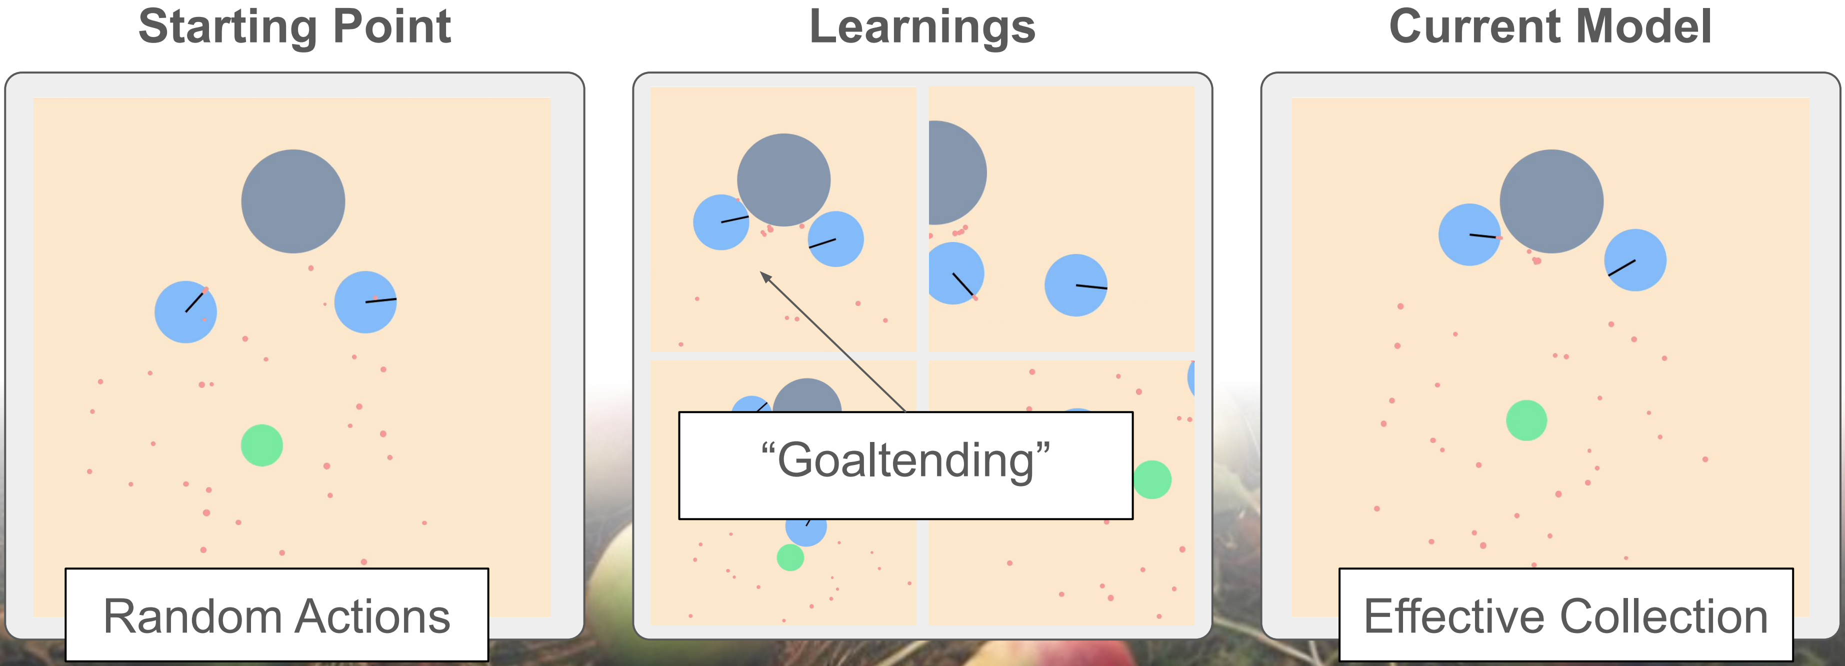

Model Demonstration

- Hundreds of simulations simultaneously = quick learning

Starting Point

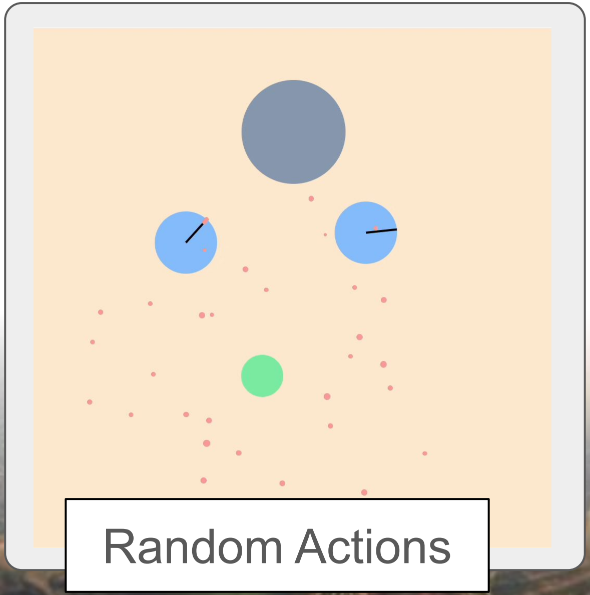

Learnings

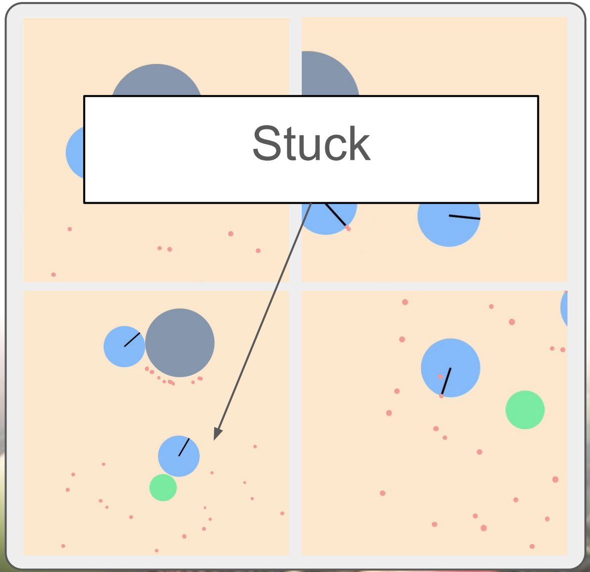

Current Model

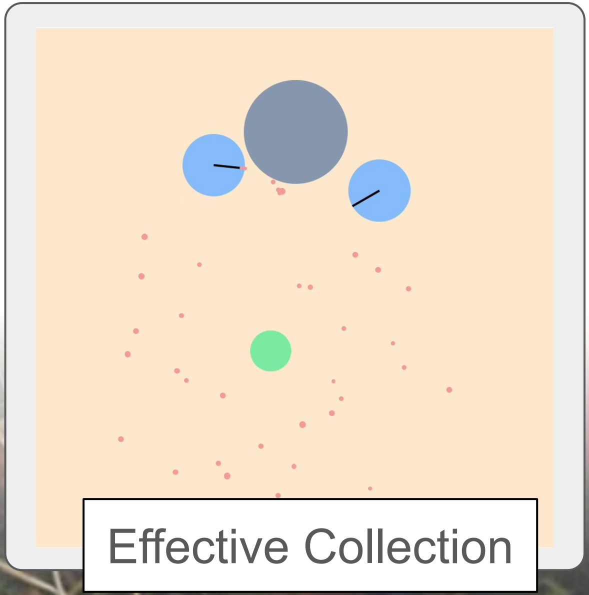

Model Demonstration

- Hundreds of simulations simultaneously = quick learning

Starting Point

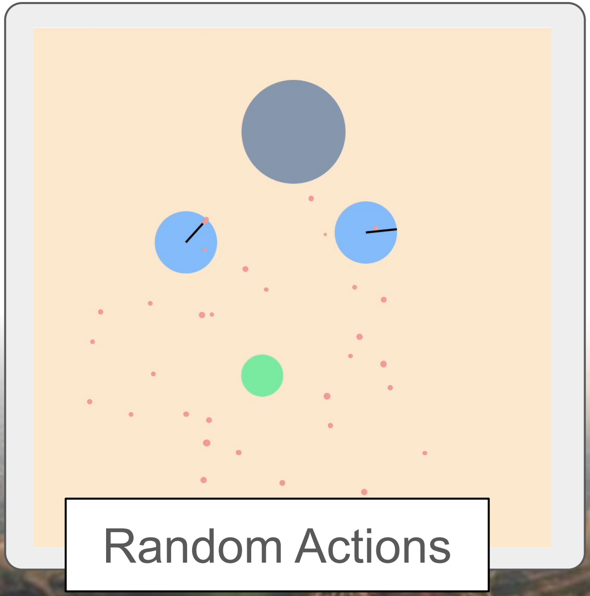

Learnings

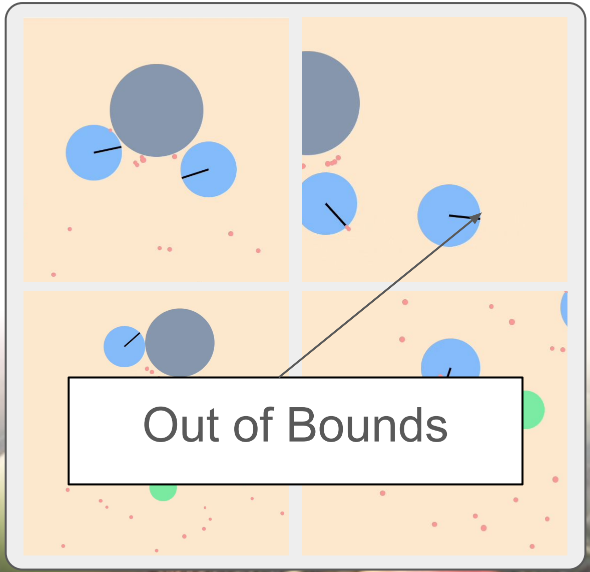

Current Model

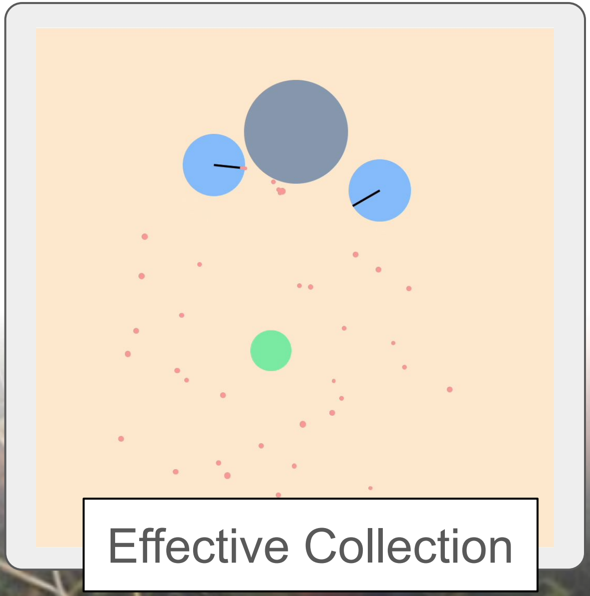

Model Demonstration

- Hundreds of simulations simultaneously = quick learning

Starting Point

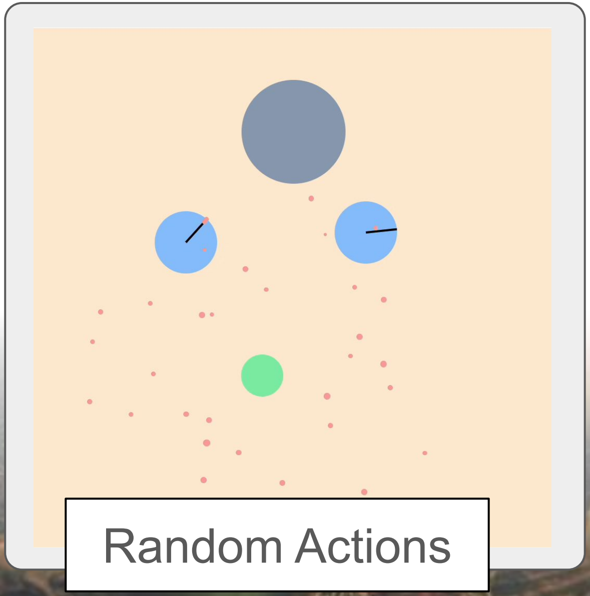

Learnings

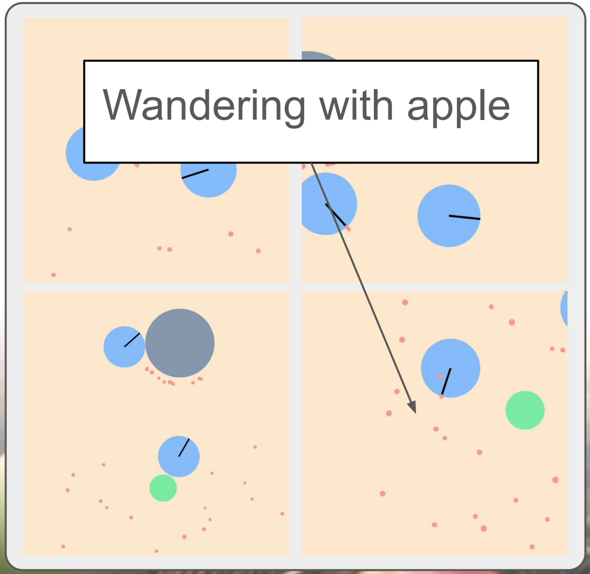

Current Model

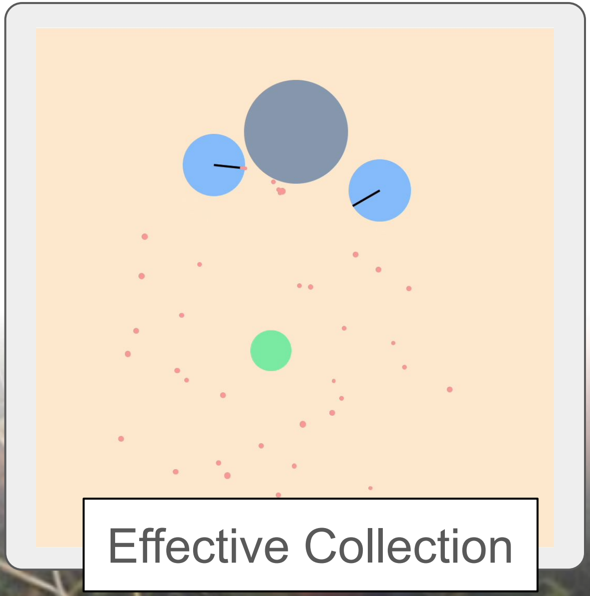

Model Demonstration

A* Baseline
(18 of 35 Collected in 3mins)

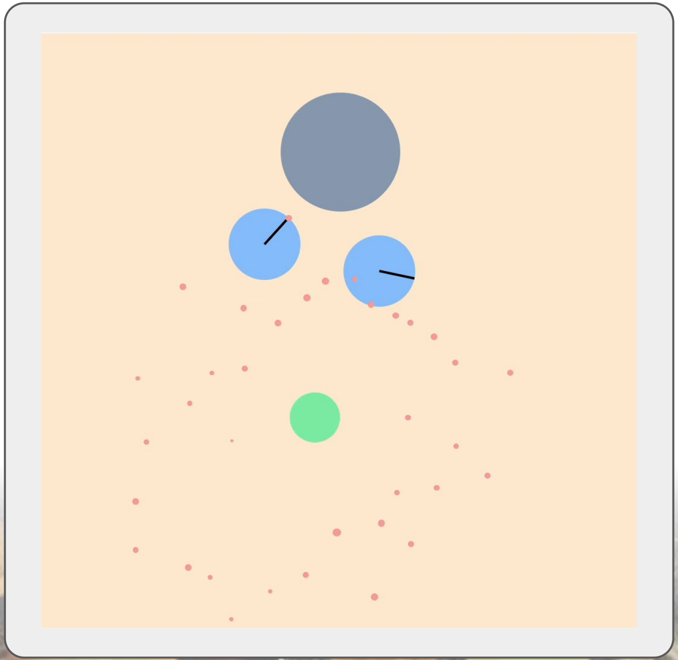

MARL Model
(25 of 35 Collected in 3mins)

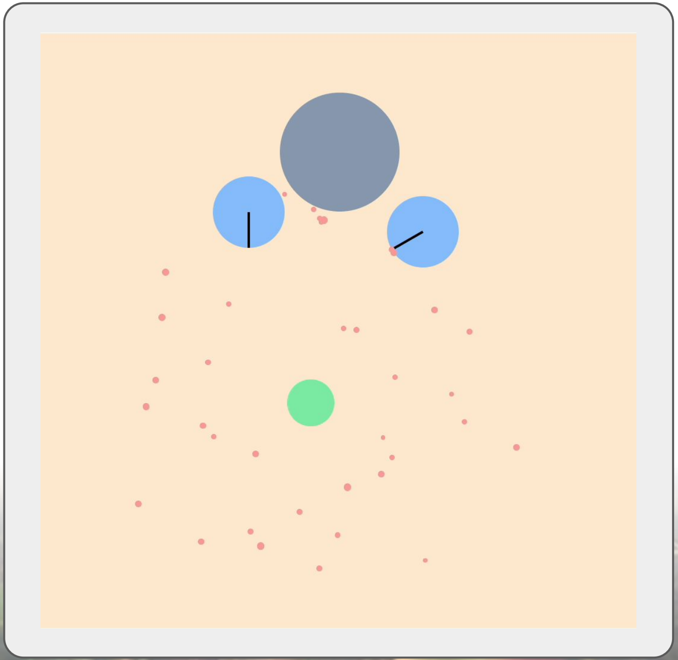

Hyperparameter Optimization - Grid Search

Training Duration →

Gradient Update Frequency ↓

	256 Steps	512 Steps	1024 Steps	2048 Steps	4096 Steps
400 / Epoch	264 🍎/hr	266 🍎/hr	84 🍎/hr	288 🍎/hr	264 🍎/hr
800 / Epoch	252 🍎/hr	204 🍎/hr	252 🍎/hr	288 🍎/hr	**300 🍎/hr**
1600 / Epoch	264 🍎/hr	276 🍎/hr	288 🍎/hr	276 🍎/hr	300 🍎/hr

Reward = 20

Reward	**5**	**10**	**20**	**100**	**200**
Result	144 🍎/hr	192 🍎/hr	**300 🍎/hr**	216 🍎/hr	216 🍎/hr

Duration = 4096 / Updates = 800

Training Statistics

- 102 models trained
- 20 virtual machines
- 85 compute-hours
- 24,000,000,000 timesteps
- 52,000 unique orchards visited

Appears to Plateau

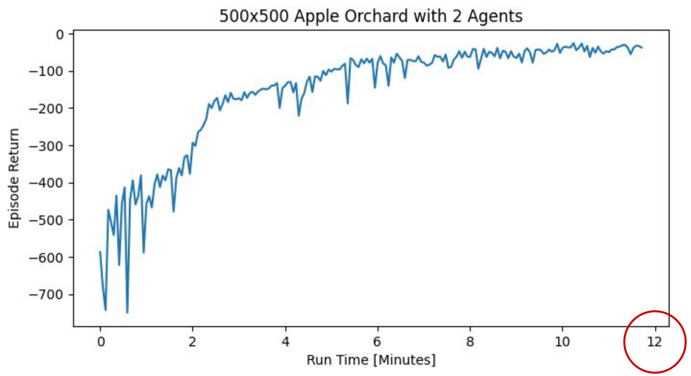

BREAKTHROUGH!

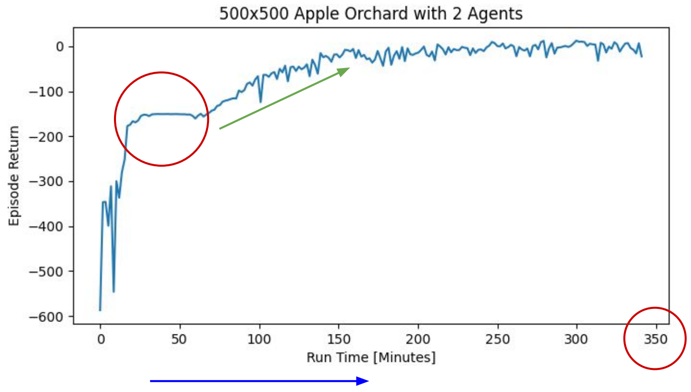

Model Evaluation

Human Baseline

A* Baseline

MARL Model

957 apples/hr**
$25/hr

201 apples/hr*
$2/hr

300 apples/hr*
$2/hr

= 10 apples
**Assumes bots only hold one apple at a time.*
***Assumes human is able to hold 5 apples at a time, but has imperfect efficiency.*

Model Evaluation

Human Baseline

MARL Model

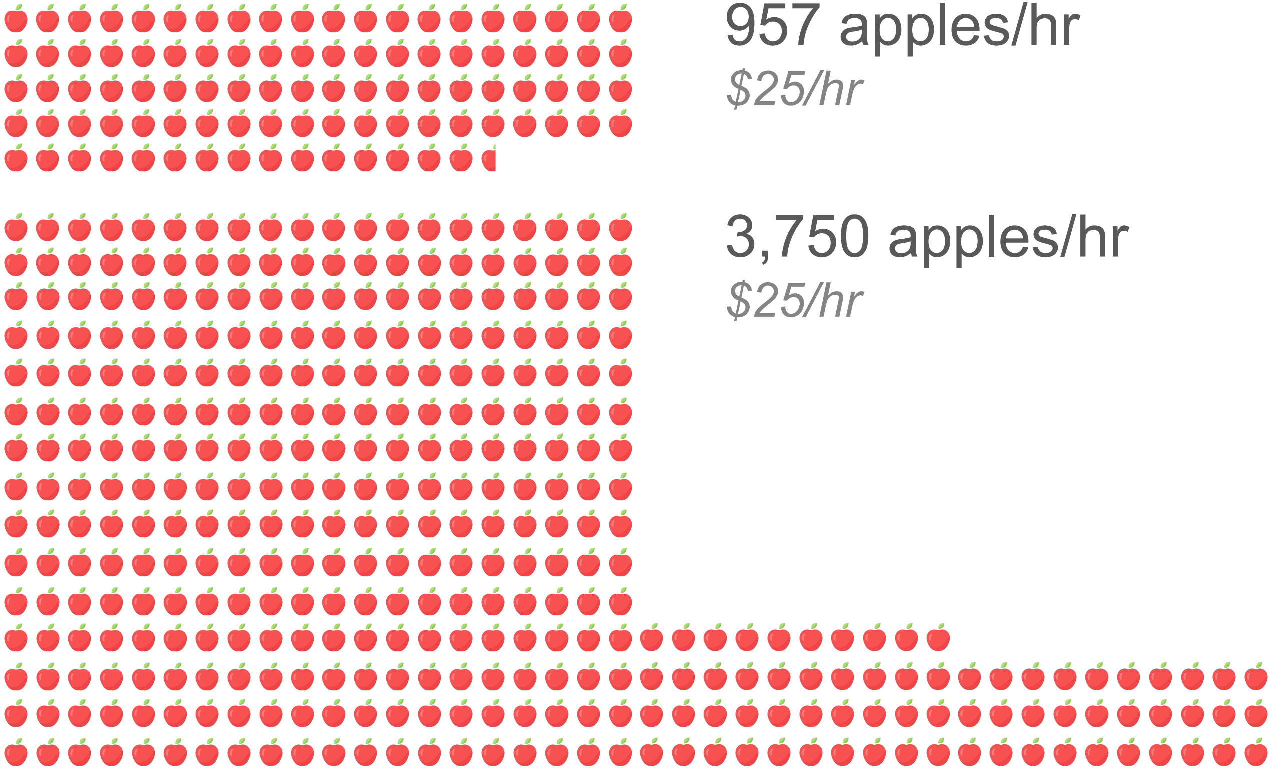

957 apples/hr
$25/hr

3,750 apples/hr
$25/hr

Key takeaways: Technical Innovation

Research Contribution:

- MARL research has seen significant progress over the past decade
- Yet, minimal research available for MARL applied to agriculture

- **Our research** explores the application of **MARL in the agricultural space** and provides a **new environment** for further development of this use case.

Key Technical Challenges:

1. **Environment complexity** and transferability to real world
2. **High compute** requirements to complete billions of training steps
3. Nascent research area with **changing libraries**

Roadmap

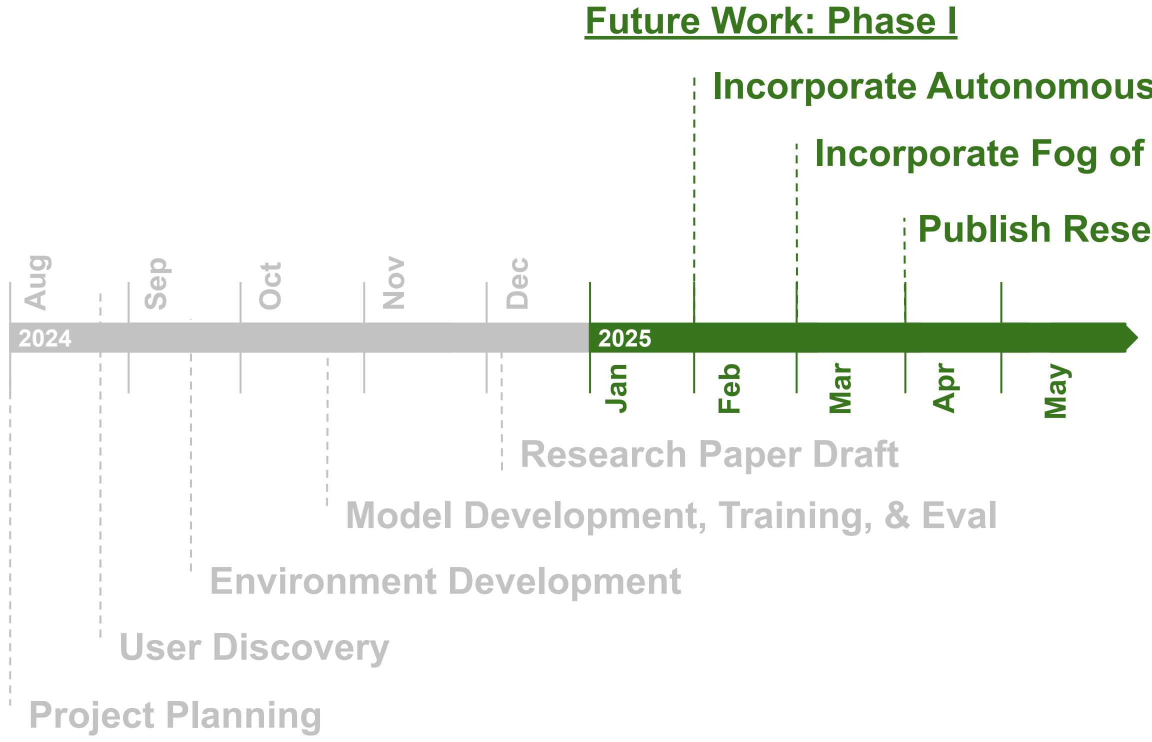

Future Work: Phase II

- Time-series component for short-term memory
- 3D environment
- Tune communication methods between bots

We're on a mission to reduce food waste by extending cutting edge multi-agent reinforcement learning research to food sustainability.