Gridworld Mapper Model

Matt Taylor
5/6/2022

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Overview

UNIVERSITY OF ILLINOIS
URBANA-CHAMPAIGN

Motivation

- Long term goal for MARL algorithms
- Capture the Flag (CTF):
  - Agents have view of grid radius 4 or 5
- Test algorithms with real robots
- **Goal**: Find a way to create a top-down gridworld from a robot's ground view

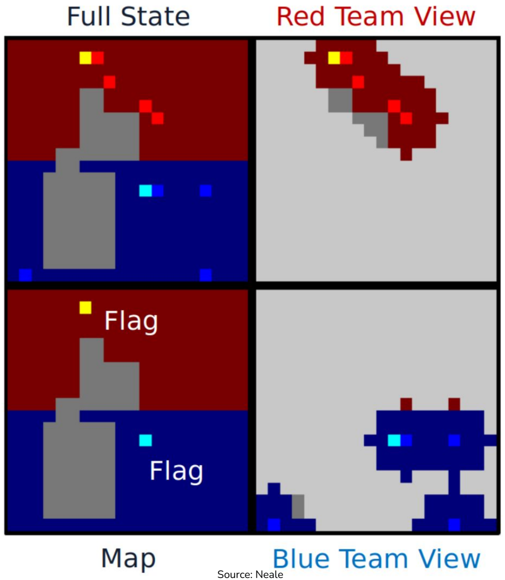

Source: Neale

Source: Neale

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

JetBot

- AI robot kit powered by the NVIDIA Jetson Nano
- Wide angle camera with 136° field of view
- Jupyter notebook interface to control the robot

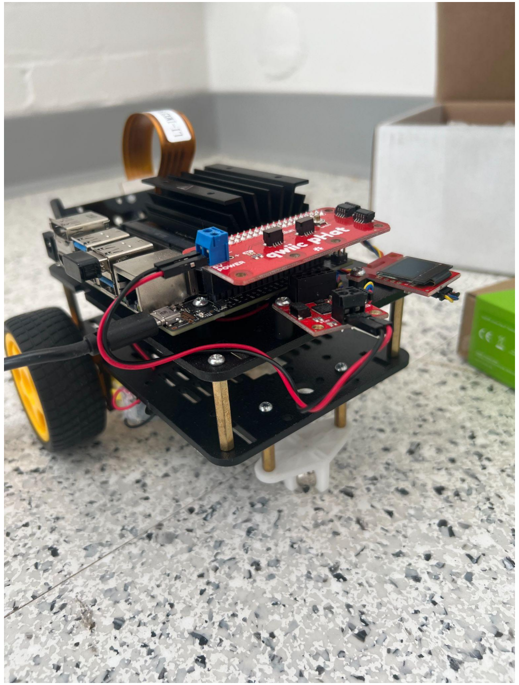

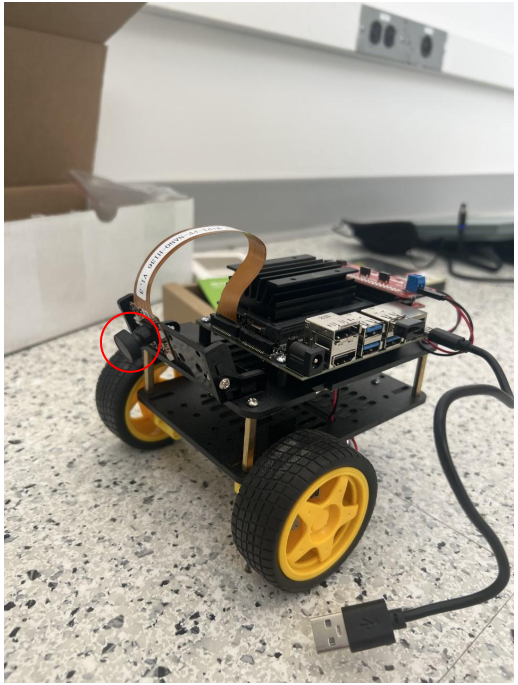

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Project Goals

1. Collect data in the lab using the JetBot.
2. Create a pipeline for the gridworld mapper model.
3. Produce visualizations for model training and evaluation.
4. Tweak the mapper model for best results with the given data.

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Data Collection

- Move JetBot to a square on the grid
    1. Take snapshot.
    2. Label image with top down view.
    3. Back out and move into the same square.
    4. Repeat steps 1-3 until 5 images are taken.
- ~ 180 hand labeled images for 3x3 grids

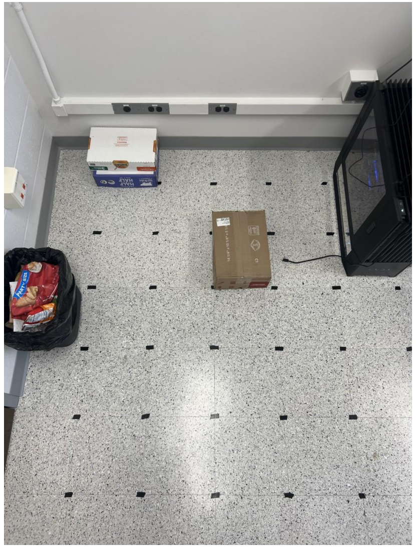

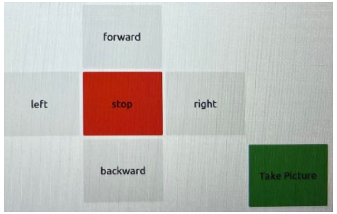

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Labeling Convention

[0, 0, 1, 0, 0, 0, 1, 0, 0]

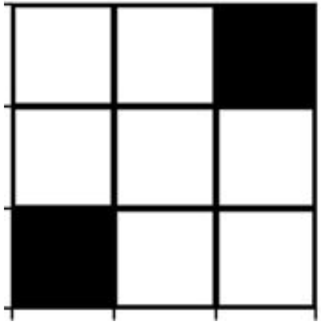

[[0, 0, 1],
[0, 0, 0],
[1, 0, 0]]

[0, 0, 0, 0, 0, 1, 1, 0, 1]

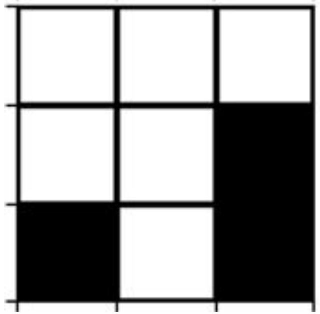

[[0, 0, 0],
[0, 0, 1],
[1, 0, 1]]

General Pipeline

- Preprocessing
  - JetBot dataset class
    - Parse labels from image filenames
    - Convert images to PyTorch float tensors and normalize
  - Split into train and validation sets (15% validation)
- CNN
  - **Input shape**: (batch_size, 3, 256, 256)
  - 6 convolutional layers with ReLU activations
  - Flatten
  - First fully connected linear layer with ReLU activation
  - Second fully connected layer
  - **Output shape**: (batch_size, 9)
- Training
  - Take batch_size and epochs as command line inputs
  - Use Adam optimizer and MSE for the loss

Results

Initial Results

- Bad . . .
- Loss plot seemed to show convergence after 30-40 epochs
- Mapper model averaged locations of all objects in the dataset

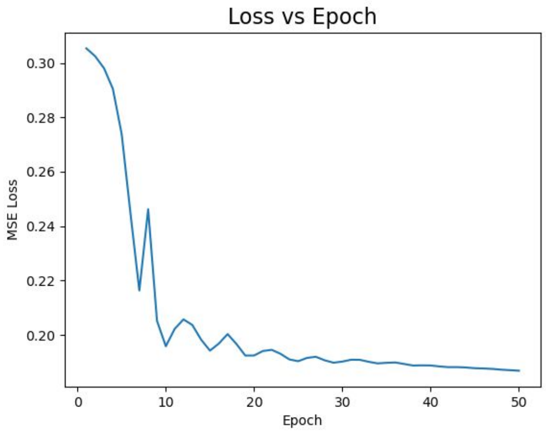

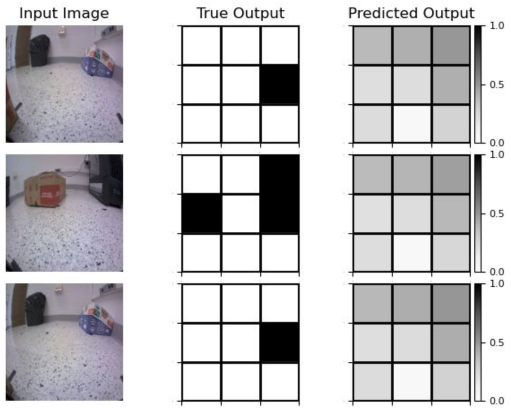

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Final Results

Training:

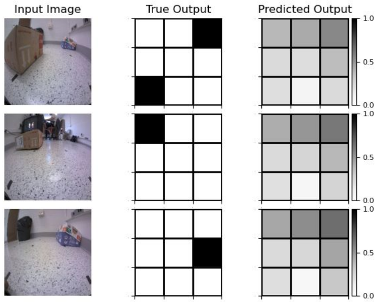

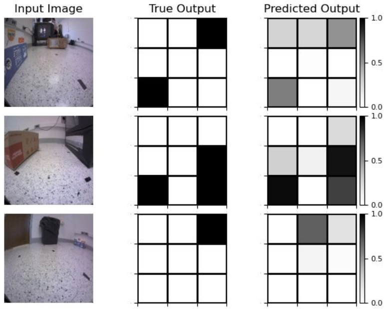

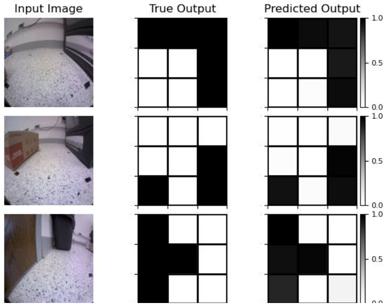

Validation:

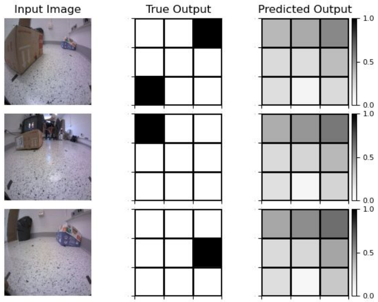

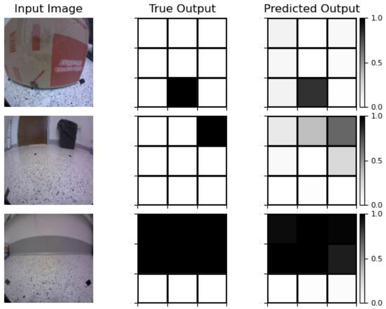

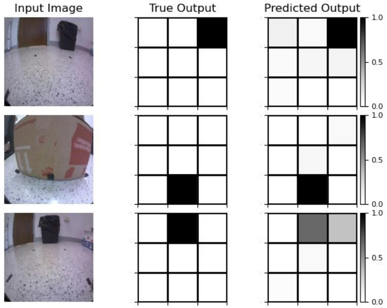

Epoch 10 Epoch 50 Epoch 80

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Batch Size: 64

Accuracy:
Train: 98.4 %
Validation: 88.5 %

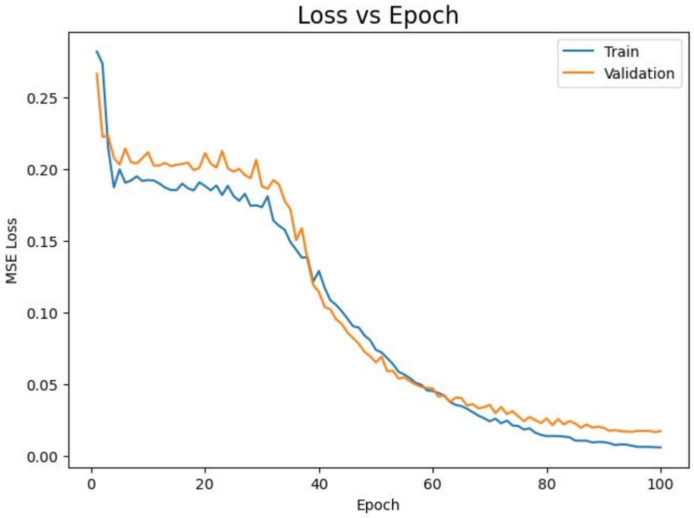

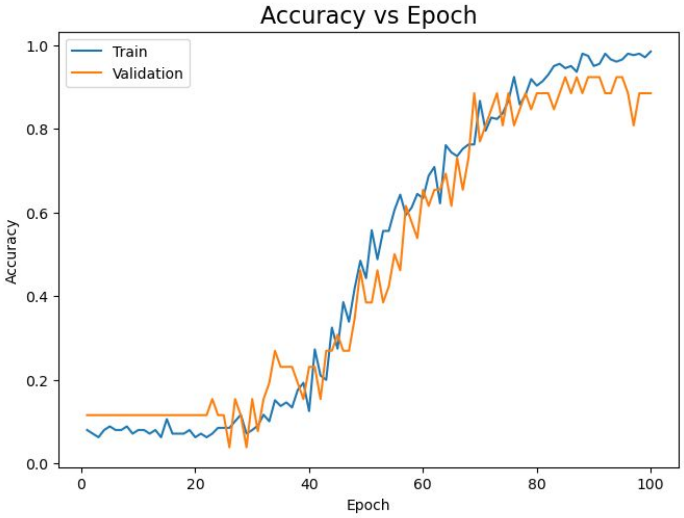

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Samples from Training Set

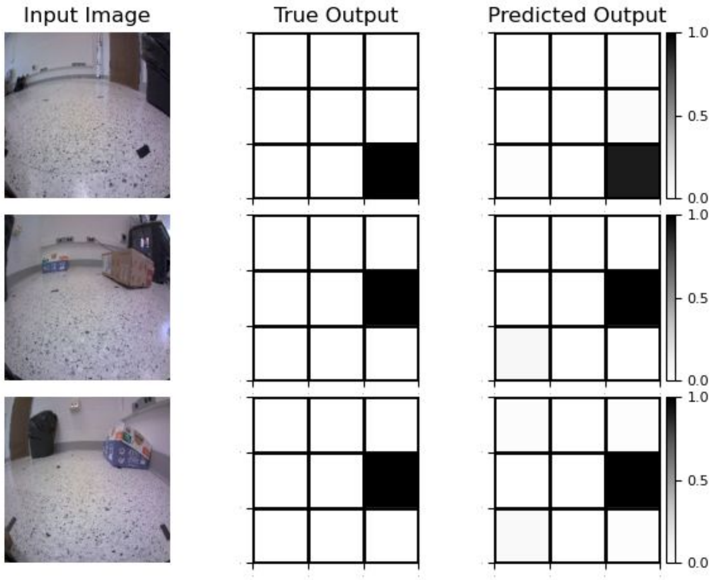

Samples from Validation Set

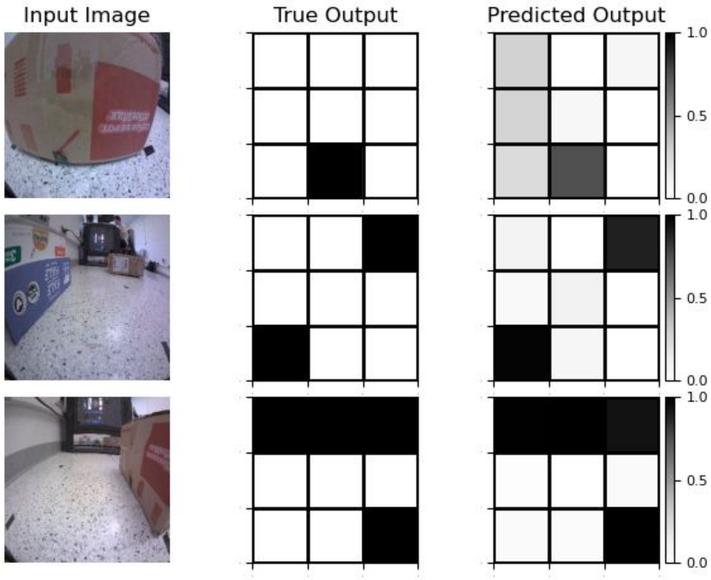

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Batch Size: 32

Accuracy:
Train: 99.38 %
Validation: 88.5 %

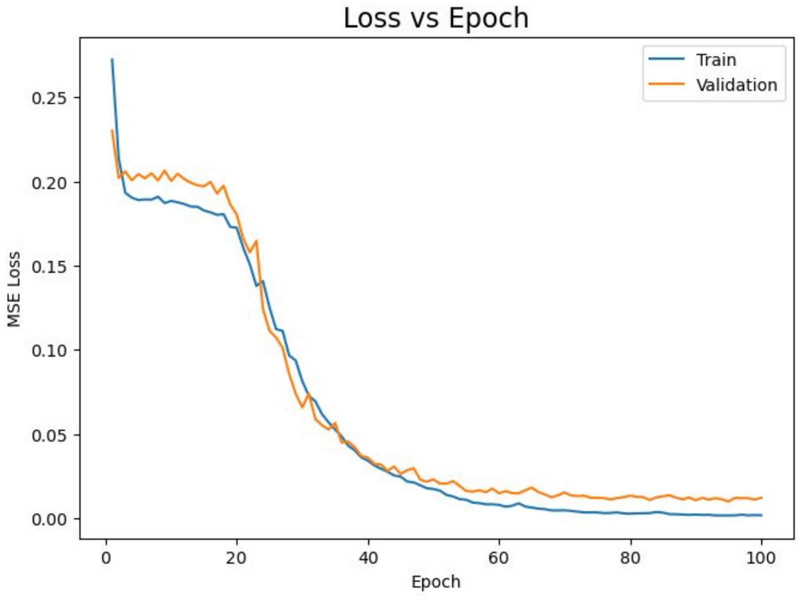

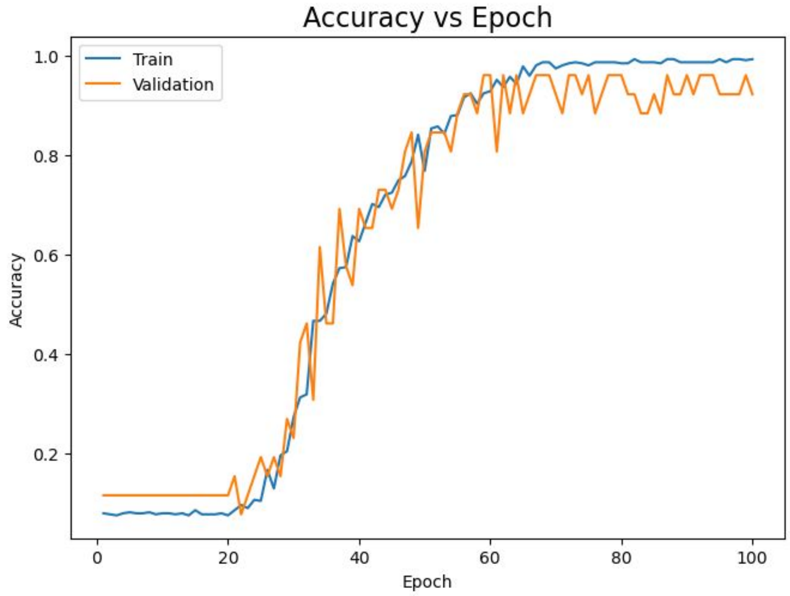

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Samples from Training Set

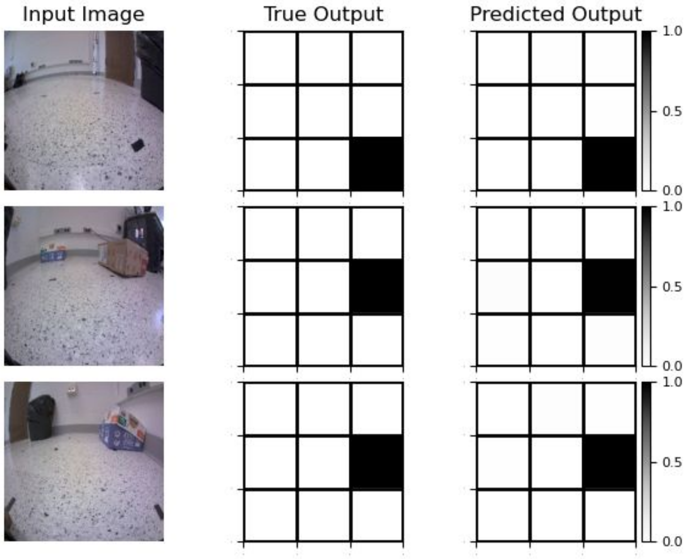

Samples from Validation Set

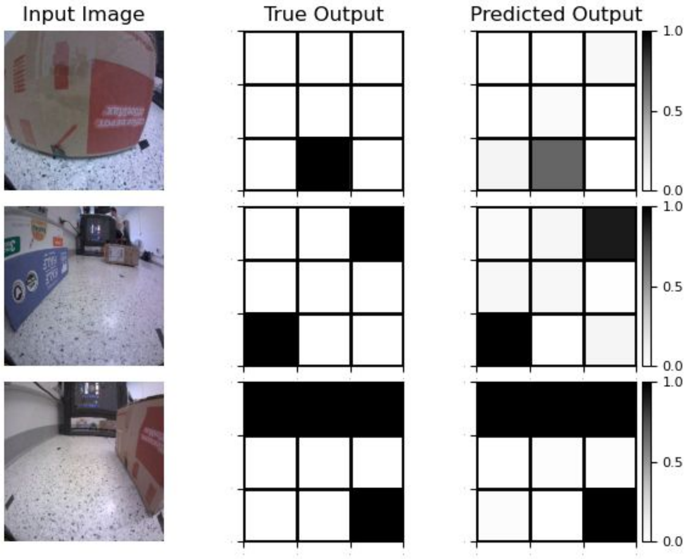

Conclusions and Future Work

Conclusions and Future Work

- Gridworld mapper model had ~88% accuracy on unseen data.
- Shows that an accurate top town view can be produced with little data.
- One step closer to using real agents for MARL algorithms.

<u>Future Work</u>

- Continue to try and reduce overfitting with other regularization techniques.
- Collect more data with more variation in obstacles and obstacle placement.
- Robustness testing
    - Change model to have certainty levels for hidden locations.
- Increase the grid size to 4 or 5 to better fit CTF or other MARL applications.

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN