

# DeepRacer Car



# Motivation

The AWS Deepracer is a complex environment for reinforcement learning. I can learn how the agent training with complex action space and reward function. It is also a good practice for learning autonomous vehicle. The most important is that the model we get from training can apply to the physical agent. This closely resembles a real-world use case.

For this project, we also need to process this system on local machine. This is a good way for us to get familiar with Docker, Cuda, Gazebo, etc, if we did not use these things before. We can find different repositories for the locally training. If you have time, try different repositories is also a chance to know how they organize the applications and what they modified for the Deepracer.



Agent: RC car with camera  
Action: Go forward, turn left,  
turn right

# Action Space

Maximum steering angle

degrees

Max values are between 1 and 30.

Steering angle granularity

▼


Maximum speed

m/s

Select values between 0.1 and 4.

Speed granularity

▼



## Action list

| Action number | Steering angle | Speed   |
|---------------|----------------|---------|
| 0             | -30 degrees    | 0.5 m/s |
| 1             | -30 degrees    | 1 m/s   |
| 2             | 0 degrees      | 0.5 m/s |
| 3             | 0 degrees      | 1 m/s   |
| 4             | 30 degrees     | 0.5 m/s |
| 5             | 30 degrees     | 1 m/s   |

## 课程 7: The League

继续学习课程 7

AWS DeepRacer Course

100% 已学完

Welcome!

✓ 已完成

Get Started with AWS DeepRacer

✓ 已完成

Test Drive DeepRacer

✓ 已完成

Reinforcement Learning

✓ 已完成

Tuning Your Model

✓ 已完成

The course from Udacity provides good instructions for the deepracer car. It gives basic knowledges about RL.


It also described the main contents of the training steps and details about the hyperparameters and reward function.

1846



# Envrionment





In reinforcement learning for AWS DeepRacer, an **agent** (vehicle) learns from an **environment** (a track) by interacting with it and receiving rewards for performing specific actions.



☒ **re:Invent 2018**

The official 2019 DeepRacer League Summit Circuit track.


Length: 17.6 m (57.97')  
Width: 76 cm (30")



☐ **Baadal Track**

Baadal is the Hindi word for cloud. The Baadal track combines long arching straightaways perfect for passing opportunities coupled with tight windings corners.

Length: 39 m (128')  
Width: 107 cm (42")



☐ **SOLA Speedway**

The first track of the 2020 AWS DeepRacer League Virtual Circuit season is dedicated to our reigning champion: SOLA-DNP. A dominating force of the AWS DeepRacer League in 2019, SOLA holds world record times on multiple tracks as she paved her way to taking home the AWS DeepRacer Championship Cup.

Length: 38 m (124')  
Width: 106 cm (42")



We can choose different tracks as the environment.

The state is the position that the car on the track, and the image from There is a center lane on the track to help locate the car.

There are two conditions can be considered as finish. One is complete the track and the other one is drive out of the track.

# First Try before Checkpoint



Code editor

Reward function examples

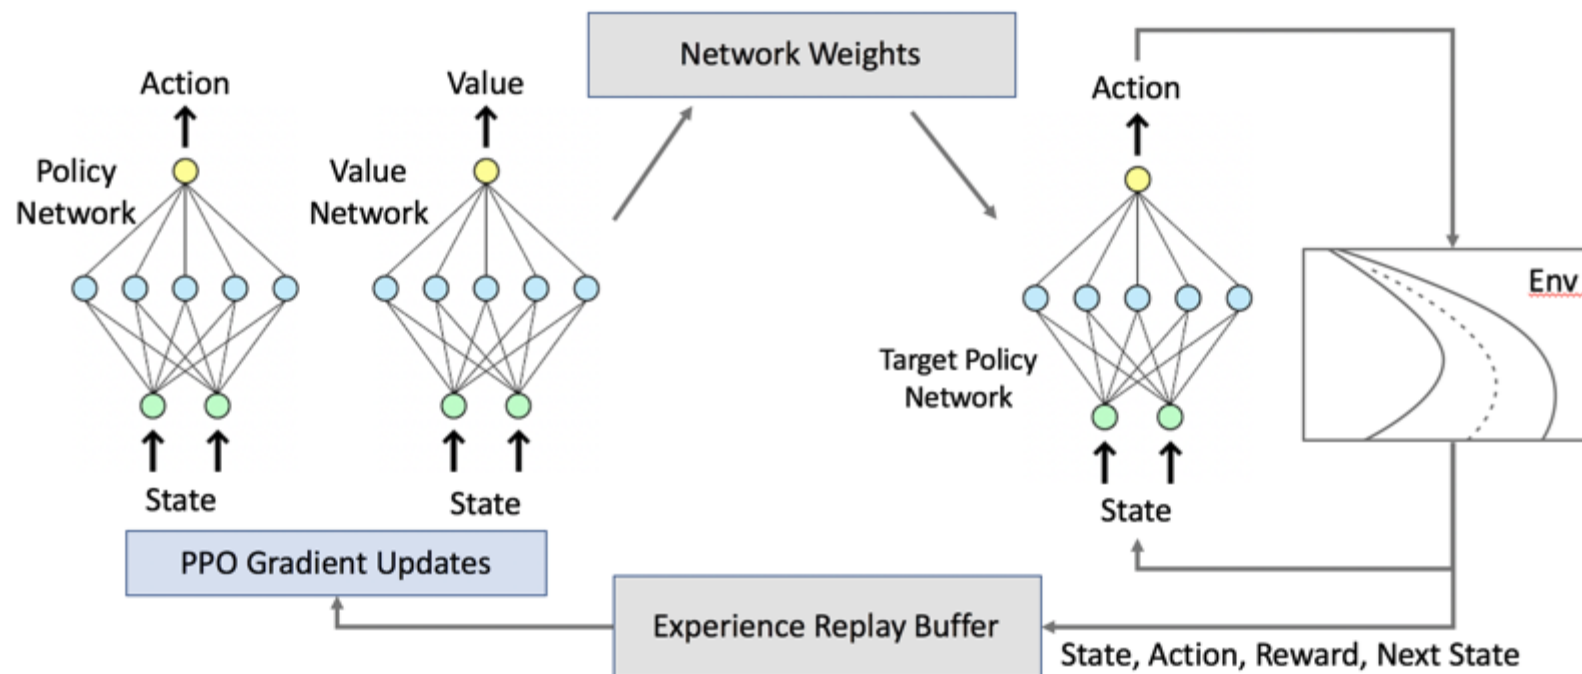
Reset

Validate

```
1 def reward_function(params):
2     '''
3     Example of rewarding the agent to follow center line
4     '''
5
6     # Read input parameters
7     track_width = params['track_width']
8     distance_from_center = params['distance_from_center']
9
10    # Calculate 3 markers that are at varying distances away from the center line
11    marker_1 = 0.1 * track_width
12    marker_2 = 0.25 * track_width
13    marker_3 = 0.5 * track_width
14
15    # Give higher reward if the car is closer to center line and vice versa
16    if distance_from_center <= marker_1:
17        reward = 1.0
18    elif distance_from_center <= marker_2:
19        reward = 0.5
20    elif distance_from_center <= marker_3:
21        reward = 0.1
22    else:
23        reward = 1e-3 # likely crashed/ close to off track
24
25    return float(reward)
```

This is the default reward function that helps agent to learn. It use the center line as the reference. The nearer to the center line the more rewards it will get.


We can modify it to gain better performance.



The default algorithm in DeepRacer is Proximal Policy Optimization algorithm. PPO uses two neural networks during training: a policy network and a value network. The policy network (also called actor network) decides which action to take given an image as input. The value network (also called critic network) estimates the cumulative reward we are likely to get given the image as input. Only the policy network interacts with the simulator and gets deployed to the real agent, namely an AWS DeepRacer vehicle.

# Hyperparameters of PPO

- **Gradient descent batch size:** The number recent vehicle experiences sampled at random from an experience buffer and used for updating the underlying deep-learning neural network weights.
- **Number of epochs:** The number of passes through the training data to update the neural network weights during gradient descent.
- **Learning rate:** During each update, a portion of the new weight can be from the gradient-descent (or ascent) contribution and the rest from the existing weight value.
- **Entropy:** A degree of uncertainty used to determine when to add randomness to the policy distribution.
- **Discount factor:** A factor specifies how much of the future rewards contribute to the expected reward.
- **Loss type:** Type of the objective function used to update the network weights.
- **Number of experience episodes between each policy-updating iteration:** The size of the experience buffer used to draw training data from for learning policy network weights.



| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:29.562 | 100%                              | Lap complete |
| 2     | 00:00:15.609 | 50%                               | Off track    |
| 3     | 00:00:27.667 | 100%                              | Lap complete |

Here is the result from the training using default rewards function and PPO algorithm. The result is pretty good and I even did not finish the training with some issues on the AWS. I setted training time to 1 hour and the training just ran for around half an hour, then it stopped and return error. The guide said I can still evaluate my result. So I got the above evaluation. Two run completed and one failed, not bad.



# Install on Local





## Repository From ARCC

- The recommended OS is Ubuntu 18.04, but I use 16.04 and it works
- Nvidia GPU needed
- CUDA and CUDNN
- Docker Docker compose and Nvidia-Docker
- AWS-cli(used for interaction with AWS)
- vncviewer(simulation visualization)

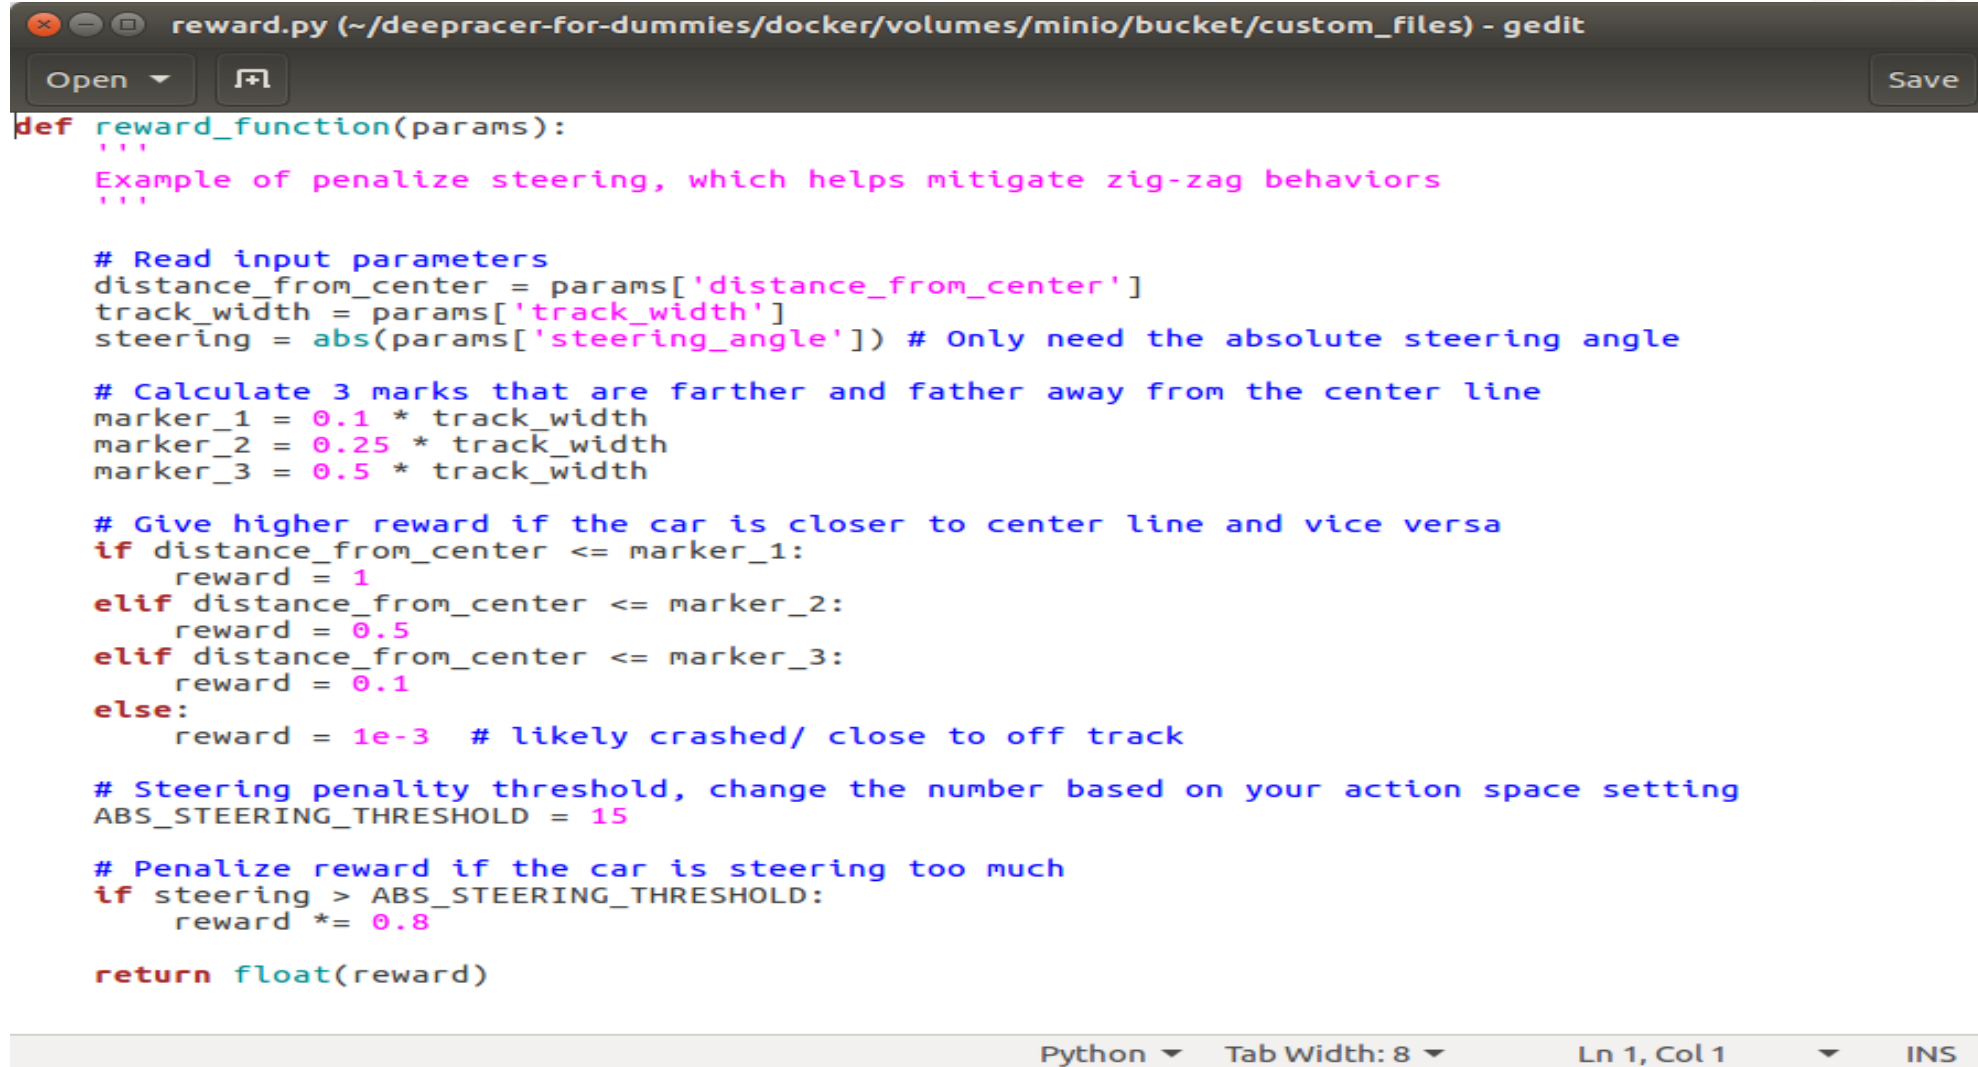
<https://github.com/ARCC-RACE/deepracer-for-dummies.git>

Git clone the repository from above link.

Then go into the folder and run the script init.sh for initialization.

Setup aws-cli(<https://www.youtube.com/watch?v=FOK5BPy30HQ>)

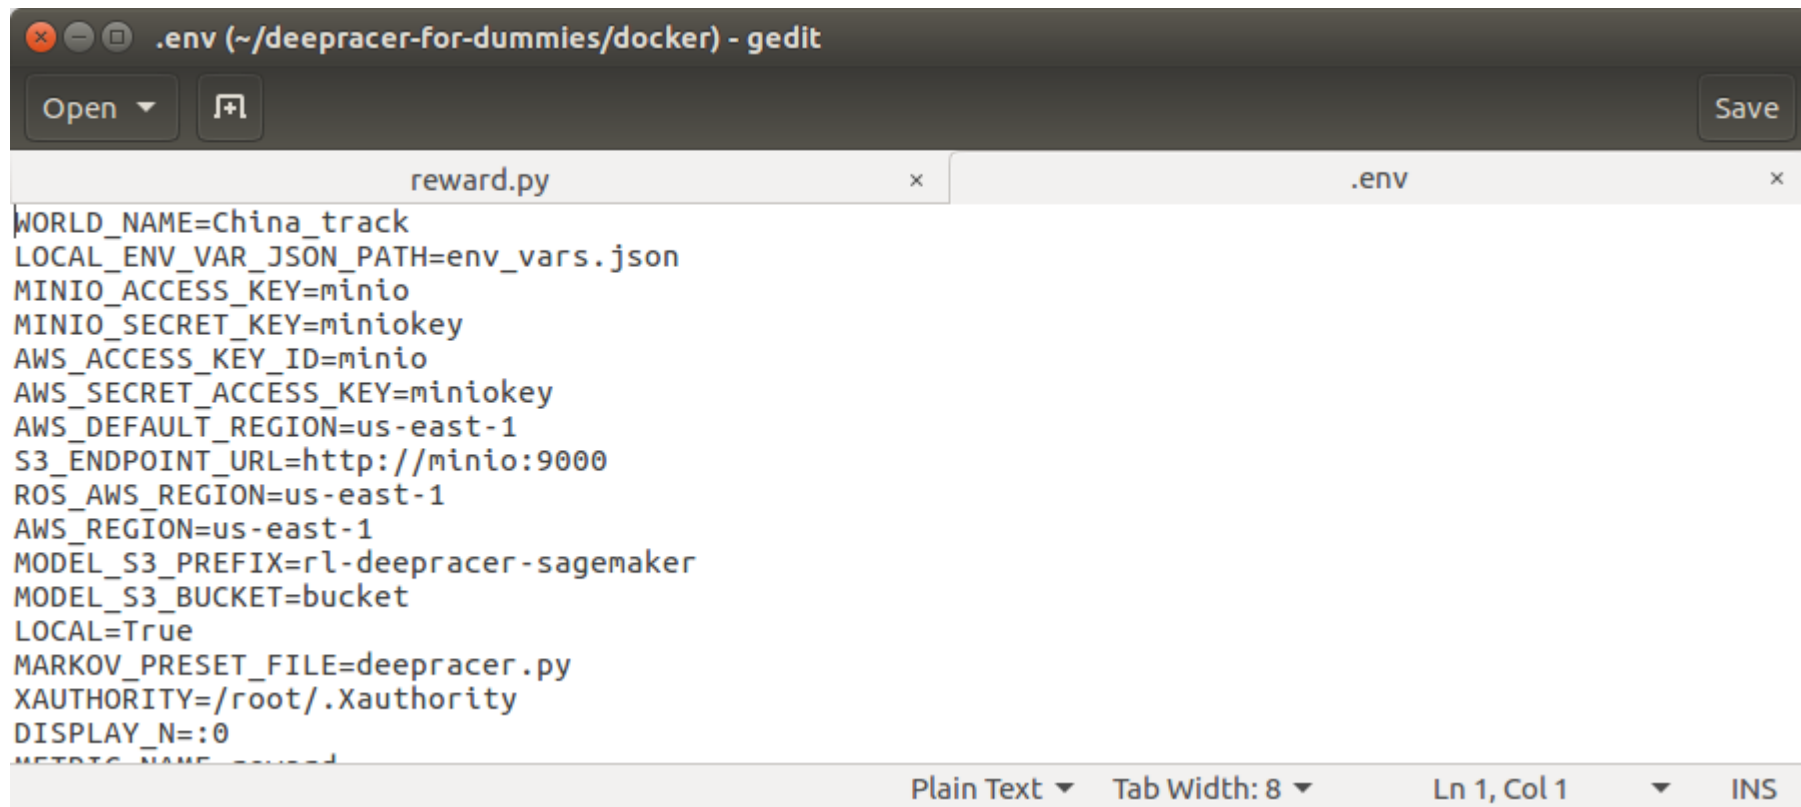
Edit the reward function in the `deepracer-for-dummies/docker/volumes/minio/bucket/custom_files/reward.py` file.



```
def reward_function(params):  
    """  
    Example of penalize steering, which helps mitigate zig-zag behaviors  
    """  
  
    # Read input parameters  
    distance_from_center = params['distance_from_center']  
    track_width = params['track_width']  
    steering = abs(params['steering_angle']) # Only need the absolute steering angle  
  
    # Calculate 3 marks that are farther and father away from the center line  
    marker_1 = 0.1 * track_width  
    marker_2 = 0.25 * track_width  
    marker_3 = 0.5 * track_width  
  
    # Give higher reward if the car is closer to center line and vice versa  
    if distance_from_center <= marker_1:  
        reward = 1  
    elif distance_from_center <= marker_2:  
        reward = 0.5  
    elif distance_from_center <= marker_3:  
        reward = 0.1  
    else:  
        reward = 1e-3 # likely crashed/ close to off track  
  
    # Steering penalty threshold, change the number based on your action space setting  
    ABS_STEERING_THRESHOLD = 15  
  
    # Penalize reward if the car is steering too much  
    if steering > ABS_STEERING_THRESHOLD:  
        reward *= 0.8  
  
    return float(reward)
```

Python ▾ Tab Width: 8 ▾ Ln 1, Col 1 ▾ INS

Change the track in `deepracer-for-dummies/docker/.env`.



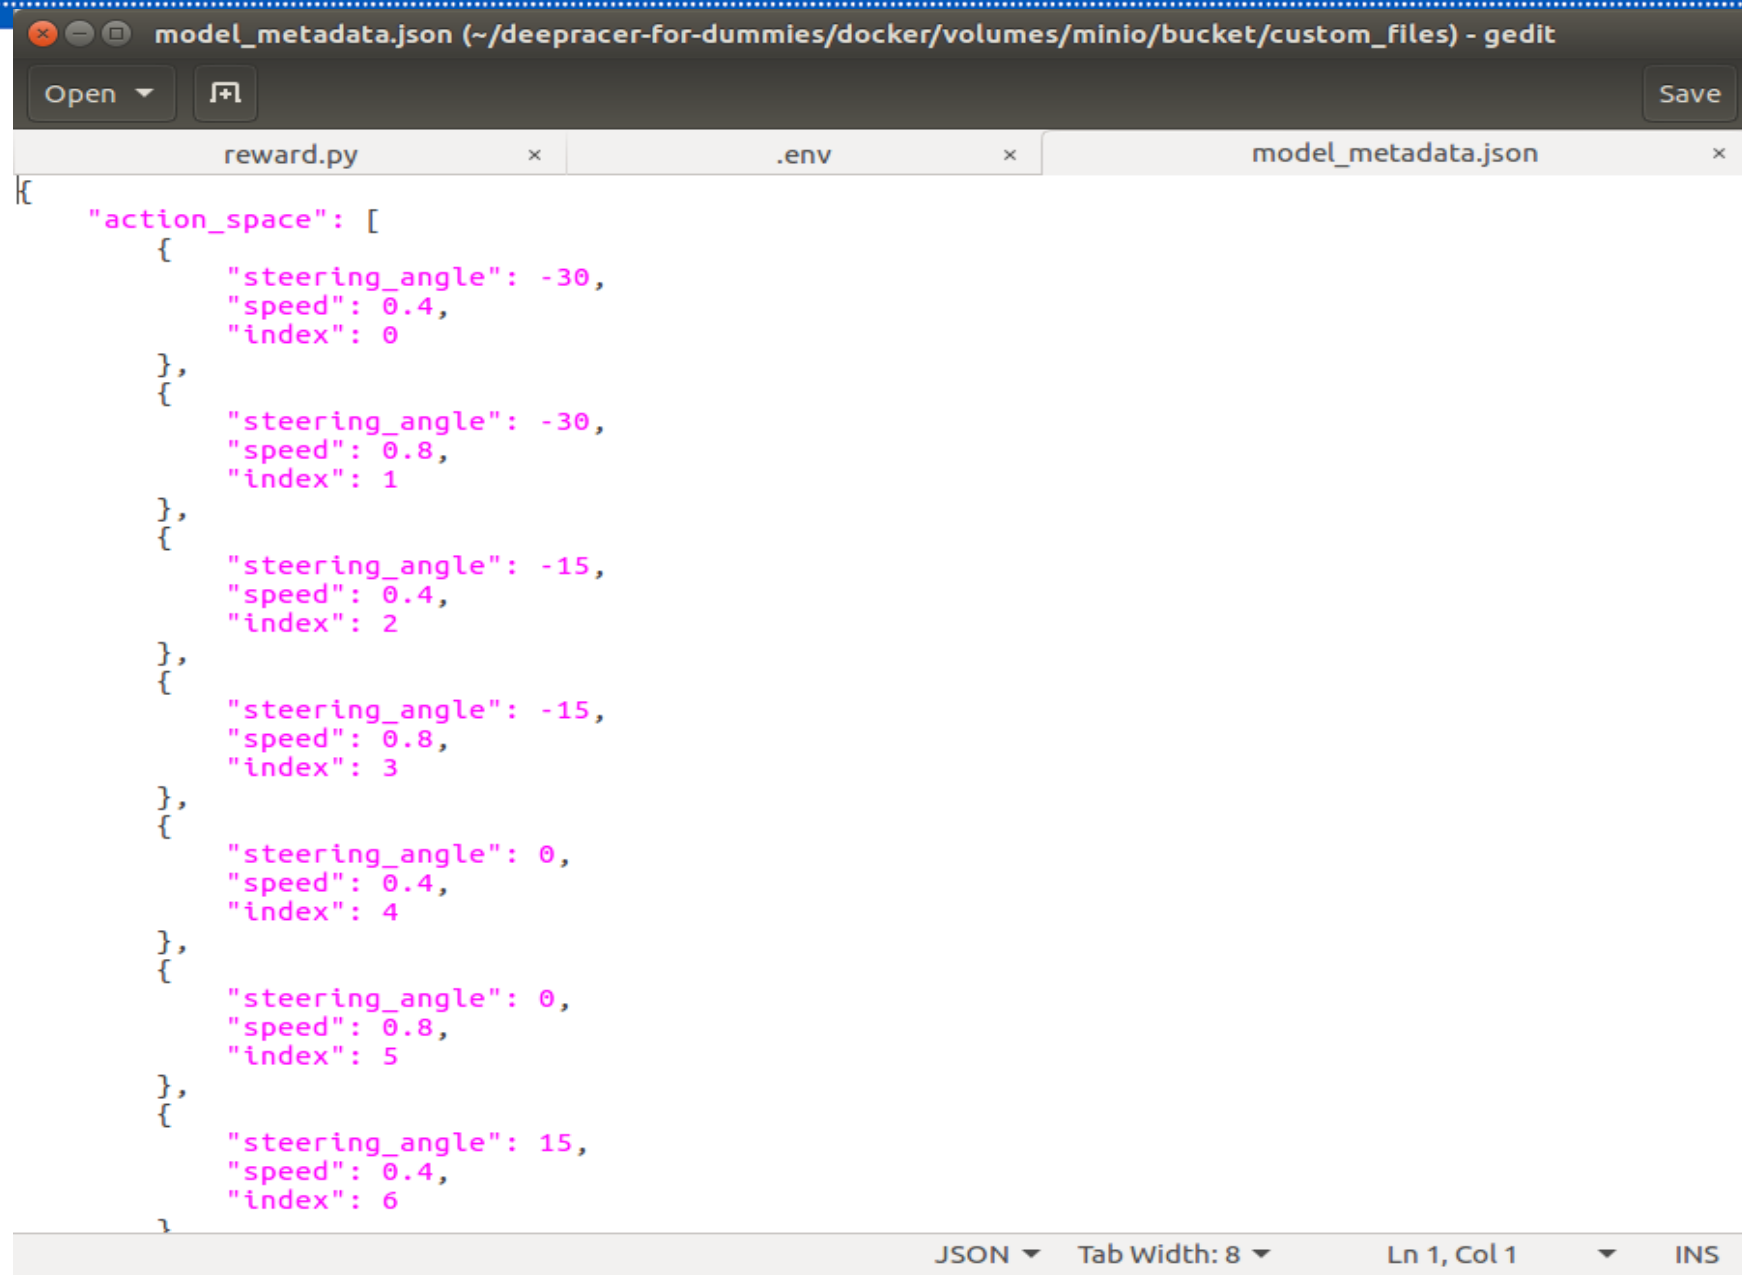
The screenshot shows a gedit editor window titled ".env (~/deepracer-for-dummies/docker) - gedit". The window has a dark theme and contains two tabs: "reward.py" and ".env". The ".env" tab is active, displaying the following environment variables:

```
WORLD_NAME=China_track
LOCAL_ENV_VAR_JSON_PATH=env_vars.json
MINIO_ACCESS_KEY=minio
MINIO_SECRET_KEY=miniokey
AWS_ACCESS_KEY_ID=minio
AWS_SECRET_ACCESS_KEY=miniokey
AWS_DEFAULT_REGION=us-east-1
S3_ENDPOINT_URL=http://minio:9000
ROS_AWS_REGION=us-east-1
AWS_REGION=us-east-1
MODEL_S3_PREFIX=rl-deepracer-sagemaker
MODEL_S3_BUCKET=bucket
LOCAL=True
MARKOV_PRESET_FILE=deepracer.py
XAUTHORITY=/root/.Xauthority
DISPLAY_N=:0
METRIC_NAME=reward
```

The status bar at the bottom indicates "Plain Text", "Tab Width: 8", "Ln 1, Col 1", and "INS".

Adjust the hyperparameters list in the deepracer-for-dummies/docker/volumes/minio/bucket/custom\_files/model\_metadata.json file.

Make sure to update the action indices as you add more actions.

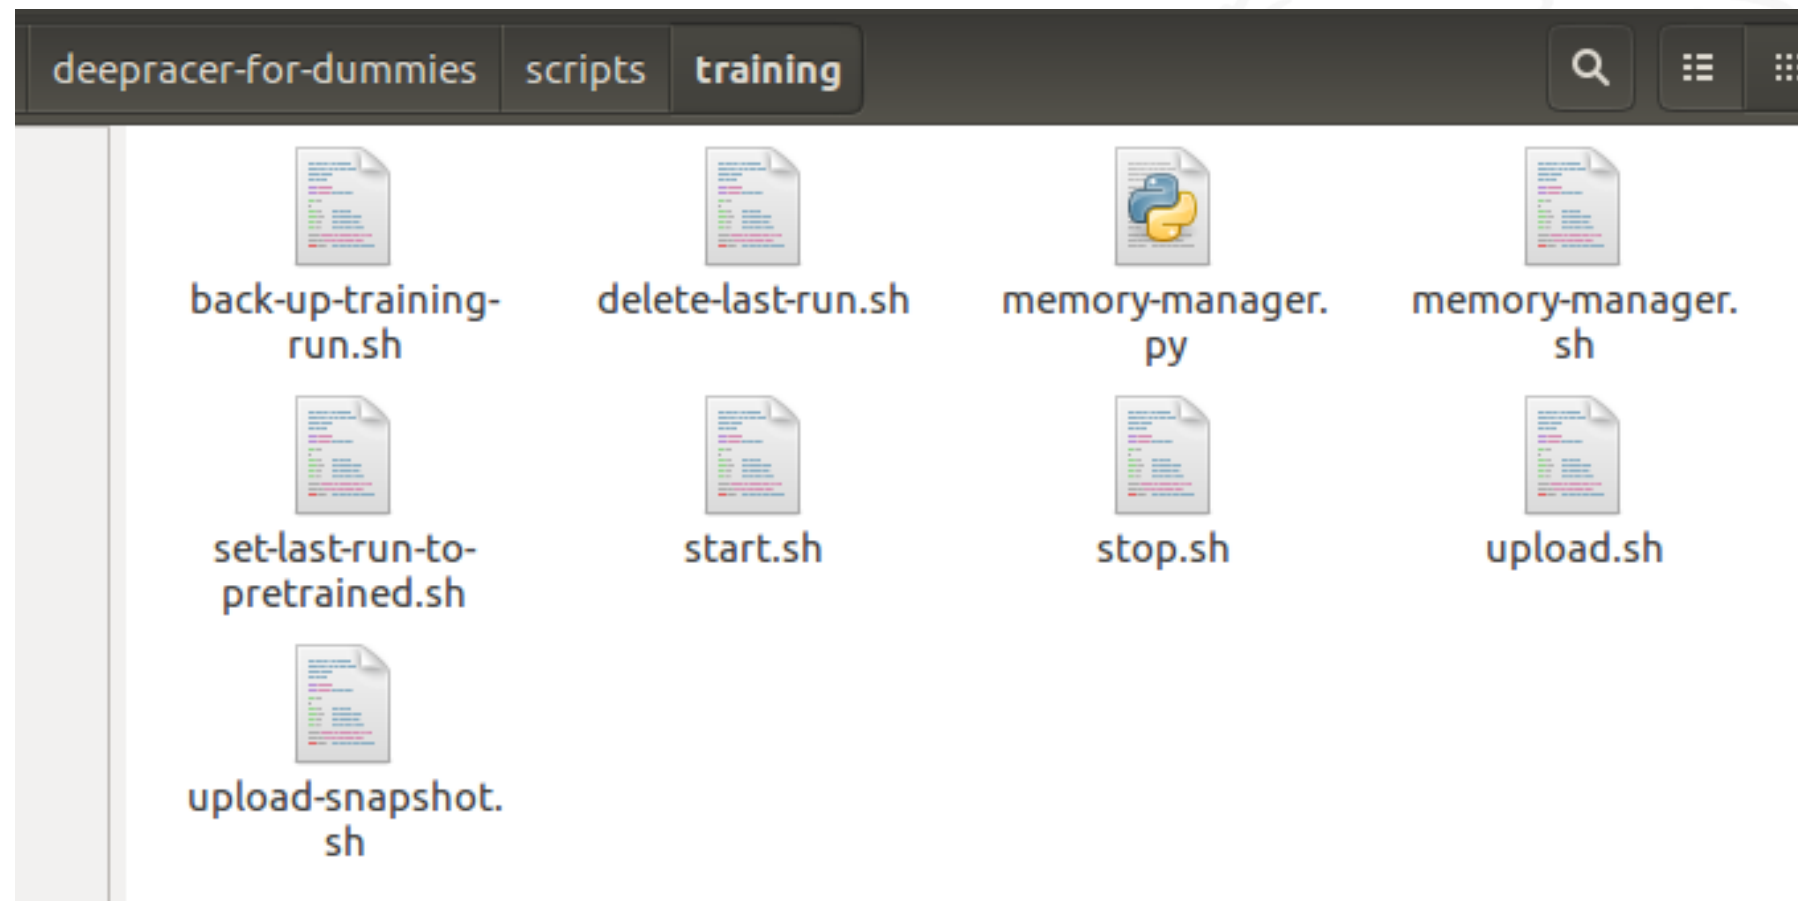


```
{
  "action_space": [
    {
      "steering_angle": -30,
      "speed": 0.4,
      "index": 0
    },
    {
      "steering_angle": -30,
      "speed": 0.8,
      "index": 1
    },
    {
      "steering_angle": -15,
      "speed": 0.4,
      "index": 2
    },
    {
      "steering_angle": -15,
      "speed": 0.8,
      "index": 3
    },
    {
      "steering_angle": 0,
      "speed": 0.4,
      "index": 4
    },
    {
      "steering_angle": 0,
      "speed": 0.8,
      "index": 5
    },
    {
      "steering_angle": 15,
      "speed": 0.4,
      "index": 6
    }
  ]
}
```

Adjust the hyperparameters in the  
 deepracer-for-  
 dummies/rl\_deepracer\_coach\_robo  
 maker.py file.

```
rl_deepracer_coach_robomaker.py (~/deepracer-for-dummies) - gedit
Open Save
reward.py x .env x model_metadata.json x rl_deepracer_coach_robomaker.py x
framework=RLFramework.TENSORFLOW,
sagemaker_session=sage_session,
# bypass sagemaker SDK validation of the role
role="aaa/",
train_instance_type=instance_type,
train_instance_count=1,
output_path=s3_output_path,
base_job_name=job_name_prefix,
image_name=image_name,
train_max_run=job_duration_in_seconds, # Maximum runtime in seconds
hyperparameters={"s3_bucket": s3_bucket,
                  "s3_prefix": s3_prefix,
                  "aws_region": aws_region,
                  "model_metadata_s3_key": "s3://{}/custom_files/
model_metadata.json".format(
    s3_bucket),
                  "RLCOACH_PRESET": RLCOACH_PRESET,
                  "batch_size": 32,
                  "num_epochs": 10,
                  "stack_size": 1,
                  "lr": 0.00035,
                  "exploration_type": "categorical",
                  "e_greedy_value": 0.05,
                  "epsilon_steps": 10000,
                  "beta_entropy": 0.01,
                  "discount_factor": 0.999,
                  "loss_type": "mean squared error",
                  "num_episodes_between_training": 20,
                  "term_cond_max_episodes": 100000,
                  "term_cond_avg_score": 100000 # place comma here if
using pretrained!
                  # "pretrained_s3_bucket": "{}".format(s3_bucket),
                  # "pretrained_s3_prefix": "rl-deepracer-pretrained"
                },
metric_definitions=metric_definitions,
s3_client=s3Client
# subnets=default subnets # Required for VPC mode
Python Tab Width: 8 Ln 1, Col 1 INS
```


After edit the previously files, we can start training by run the script start.sh. Run stop.sh to stop the training and delete-last-run.sh to clean the space before next run.





## Local training using vncviewer for visualization

The track in the image is  
China\_track map






# Change reward function



## Reward graph [Info](#)



☒ ☐ ☐ Average Reward

☒ ☐ ☐ Average percentage completion (Training)


☒ ☐ ☐ Average percentage completion (Evaluating)

## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:33.428 | 100%                              | Lap complete |
| 2     | 00:00:31.085 | 100%                              | Lap complete |
| 3     | 00:00:31.228 | 100%                              | Lap complete |

Default reward function which is focus on follow the center line to get more rewards.

## Reward graph [Info](#)



☒ -●- Average Reward

☒ -◆- Average percentage completion (Training)


☒ -■- Average percentage completion (Evaluating)

## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:12.101 | 46%                               | Off track    |
| 2     | 00:00:25.681 | 100%                              | Lap complete |
| 3     | 00:00:11.271 | 48%                               | Off track    |

I only use the `all_wheels_on_track` parameter to define the rewards. The rule is giving reward if there is no wheels go off the track.

## Reward graph [Info](#)



☒ ☐ ☐ Average percentage completion (Training)


☒ ☐ ☐ Average percentage completion (Evaluating)

## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:24.187 | 100%                              | Lap complete |
| 2     | 00:00:17.232 | 71%                               | Off track    |
| 3     | 00:00:22.838 | 100%                              | Lap complete |

This is based on the follow the center line function, and it add formula to penalize reward if the agent is steering too much.

Reward graph [Info](#)



 Average Reward

 Average percentage completion (Training)

 Average percentage completion (Evaluating)


## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:24.267 | 100%                              | Lap complete |
| 2     | 00:00:23.426 | 100%                              | Lap complete |
| 3     | 00:00:22.516 | 100%                              | Lap complete |


This is based on the follow the center line function, and it add formula to give more reward if the agent using high speed to drive.

I use the parameter speed to increase the reward for every step.

## Reward graph [Info](#)



 Average Reward

 Average percentage completion (Training)

 Average percentage completion (Evaluating)


## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:09.370 | 35%                               | Off track    |
| 2     | 00:00:28.909 | 100%                              | Lap complete |
| 3     | 00:00:21.626 | 74%                               | Off track    |


This reward function is different from the previously reward functions. I use the waypoints as the reference to give rewards.

I firstly draw a shortest path on the track map, then I define where the agent locate to the center line, and use closest waypoints as reference.

reInvent Track Waypoints - Inner, Center, and Outer








| Variable Name        | Type    |
|----------------------|---------|
| all_wheels_on_track  | Boolean |
| distance_from_center | Float   |
| is_left_of_center    | Boolean |
| speed                | Float   |
| steering_angle       | Float   |
| track_width          | Float   |
| closest_waypoints    | Integer |
| progress             | Float   |




## Reward graph [Info](#)



## Evaluation results

| Trial | Time         | Trial results (% track completed) | Status       |
|-------|--------------|-----------------------------------|--------------|
| 1     | 00:00:26.232 | 100%                              | Lap complete |
| 2     | 00:00:21.251 | 74%                               | Off track    |
| 3     | 00:00:09.964 | 35%                               | Off track    |

This is based on the waypoints following function.  
 I add formula to give higher reward if the speed is  
 fast and using less steps to finish the lap.

-  Average Reward
-  Average percentage completion (Training)
-  Average percentage completion (Evaluating)

# Thank you

