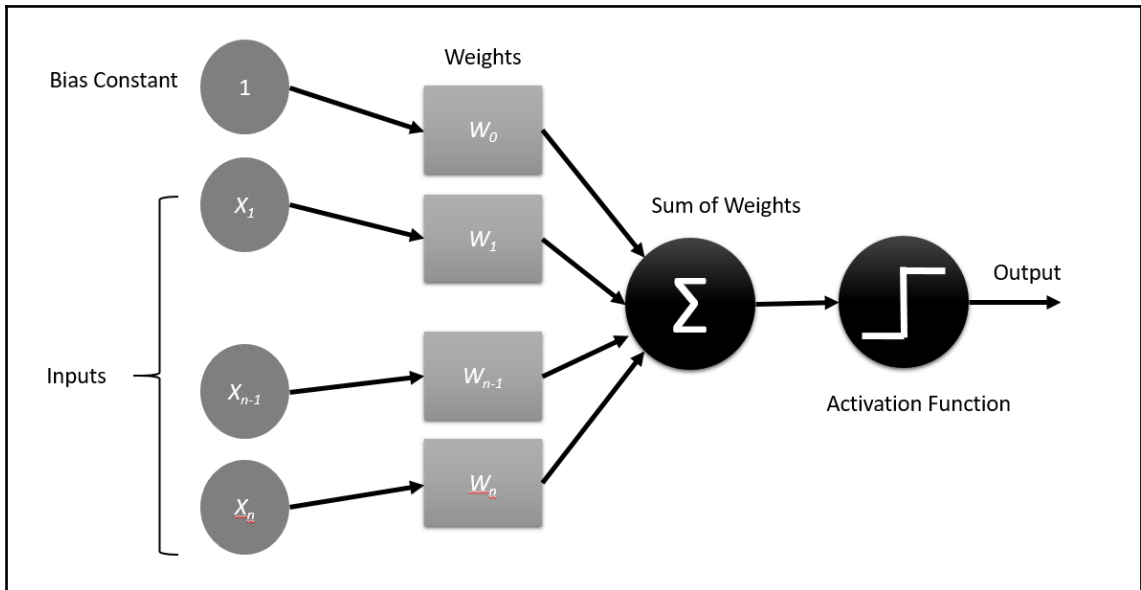
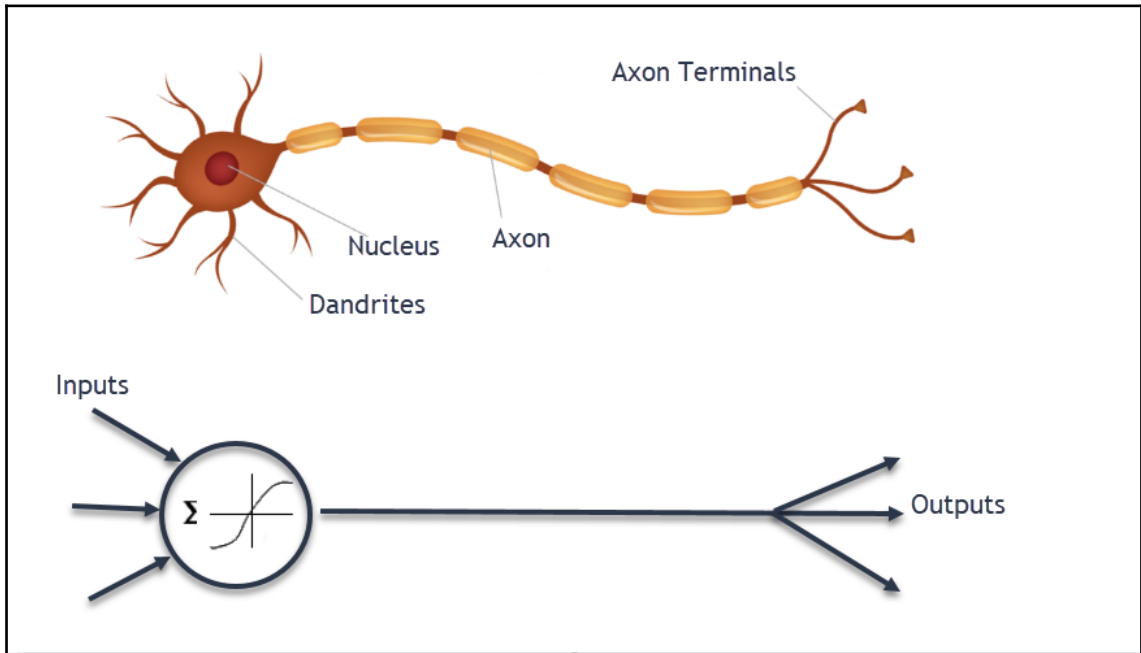




# Chapter 1: Deep Learning for Games






1D TENSOR /  
VECTOR

8
7
45
12
-6
3
22
1
8
3
-9


2D TENSOR /  
MATRIX

-9	4	2	5	7
3	0	12	8	61
1	23	-6	45	2
22	3	-1	72	6


3D TENSOR /  
CUBE



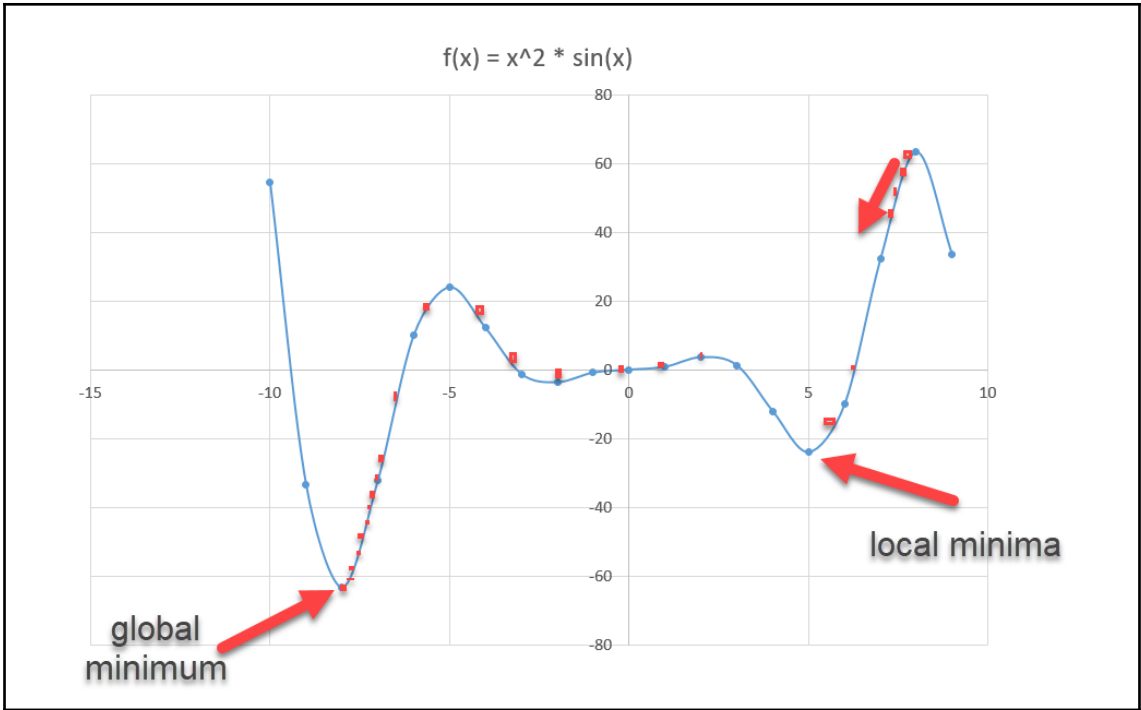
-9	4	2	5	7
3	0	12	8	61
1	23	-6	45	2
22	3	-1	72	6





4D TENSOR  
VECTOR OF CUBES



5D TENSOR  
MATRIX OF CUBES




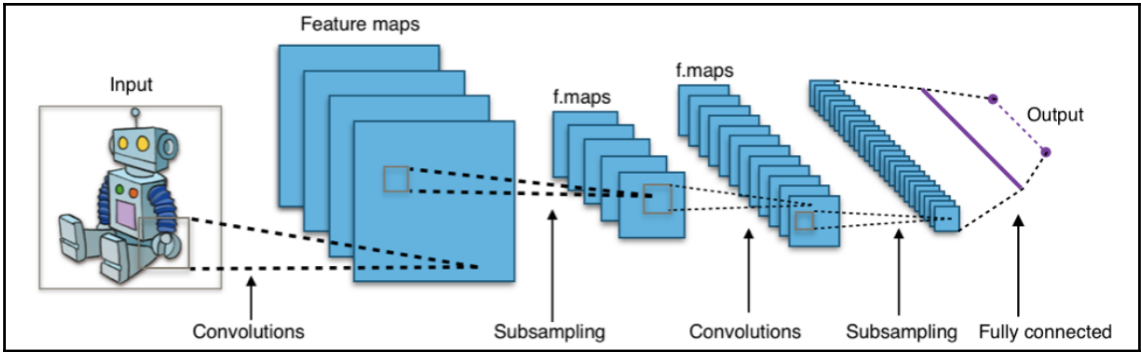





---

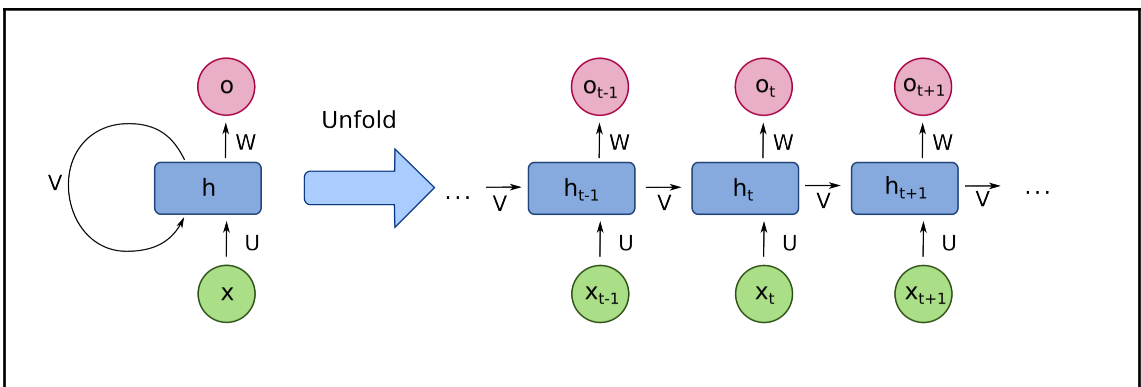

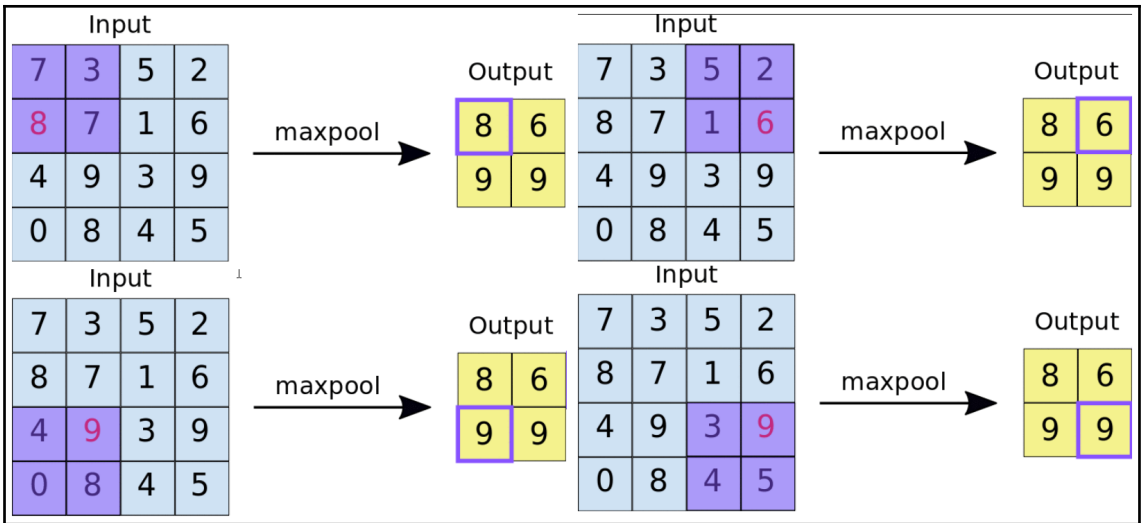
# Chapter 2: Convolutional and Recurrent Networks

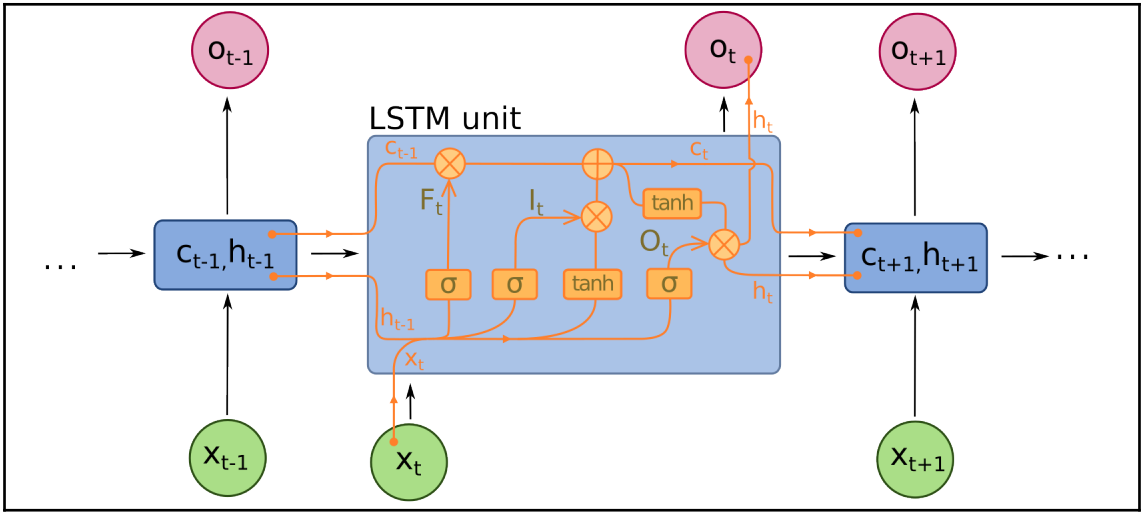






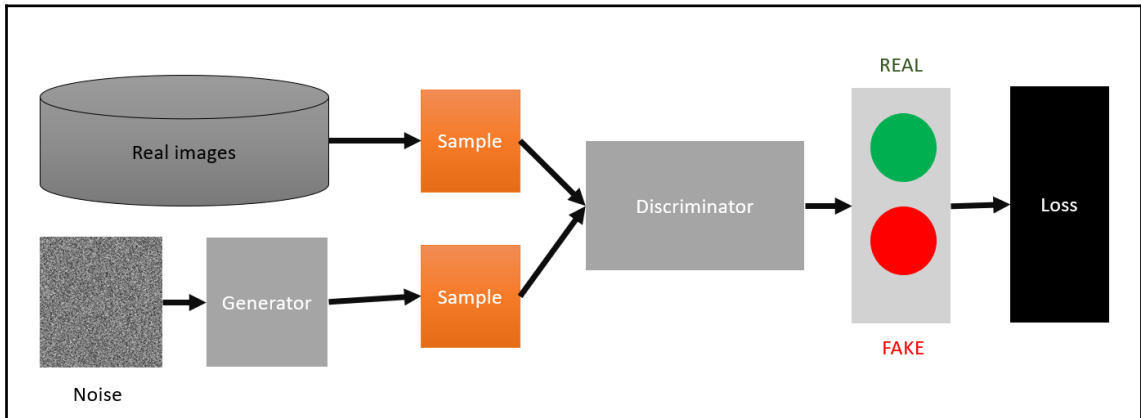









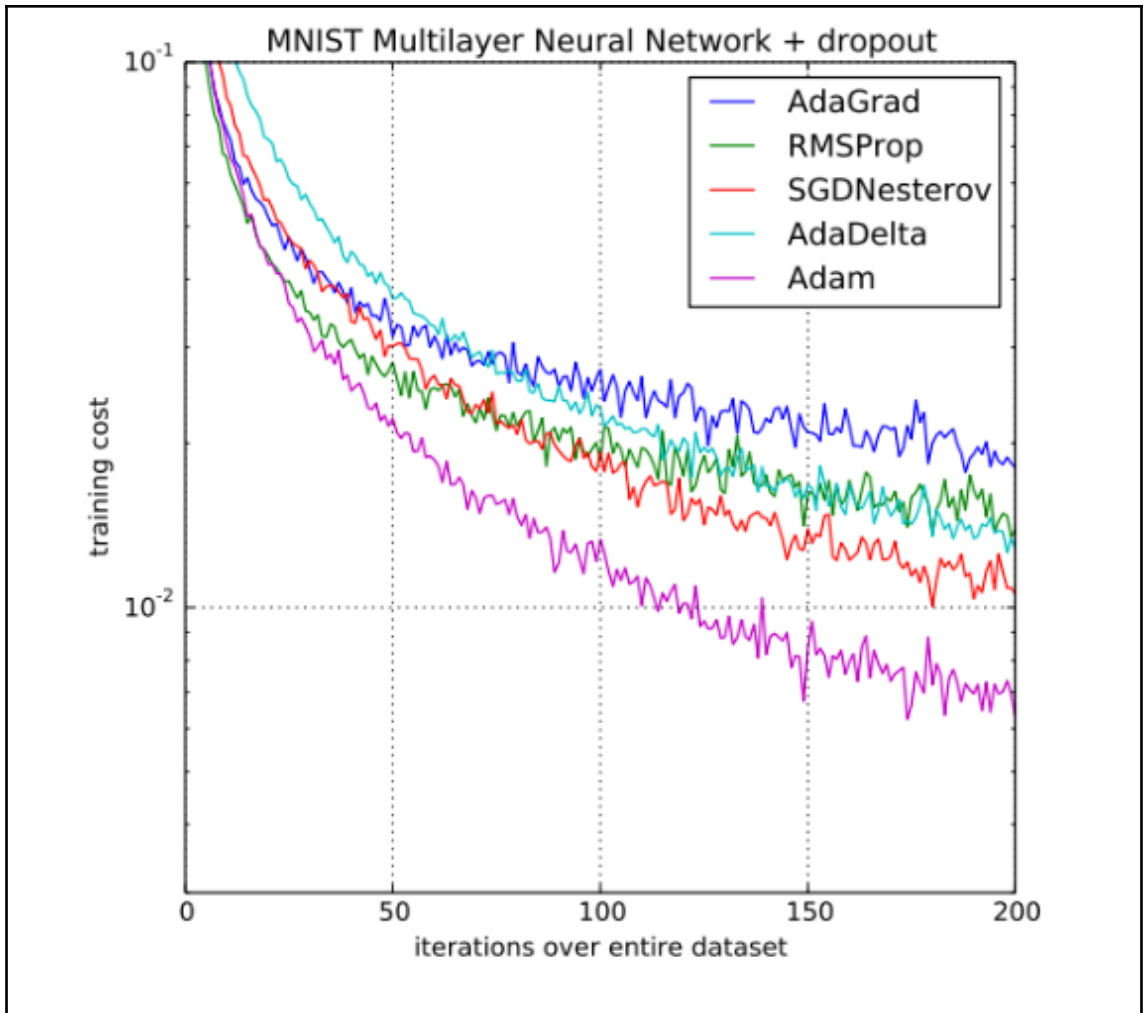
---

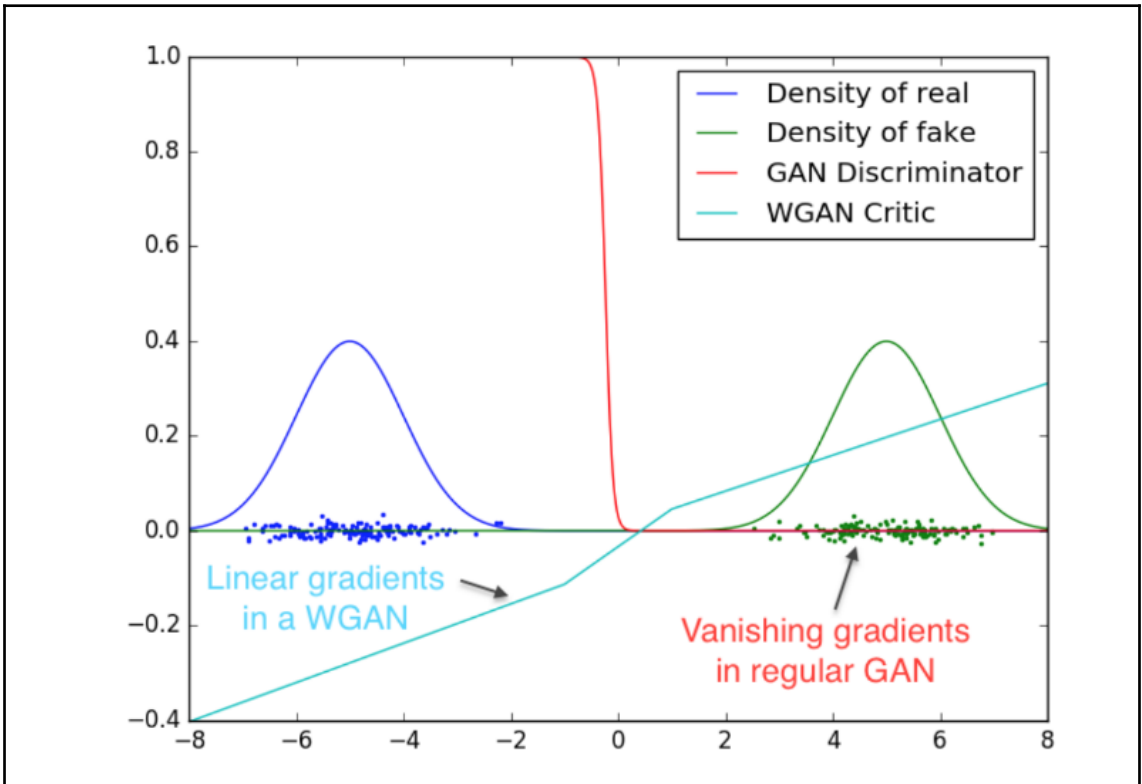

## Chapter 3: GAN for Games







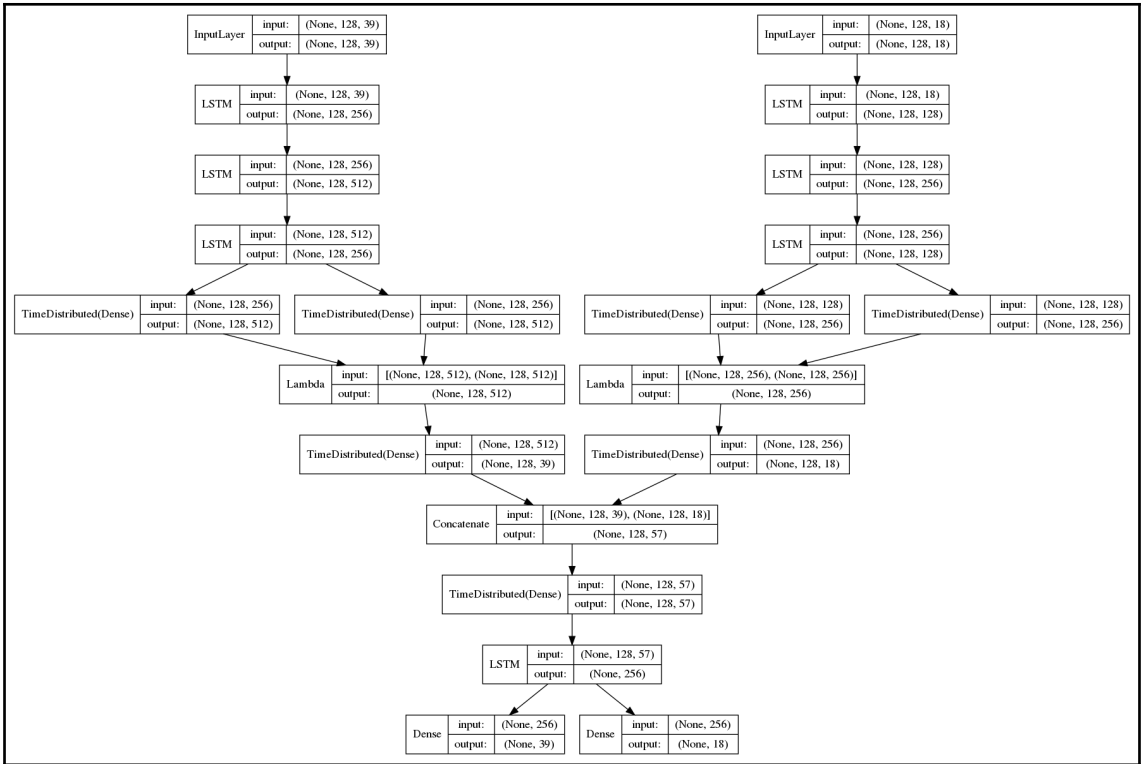




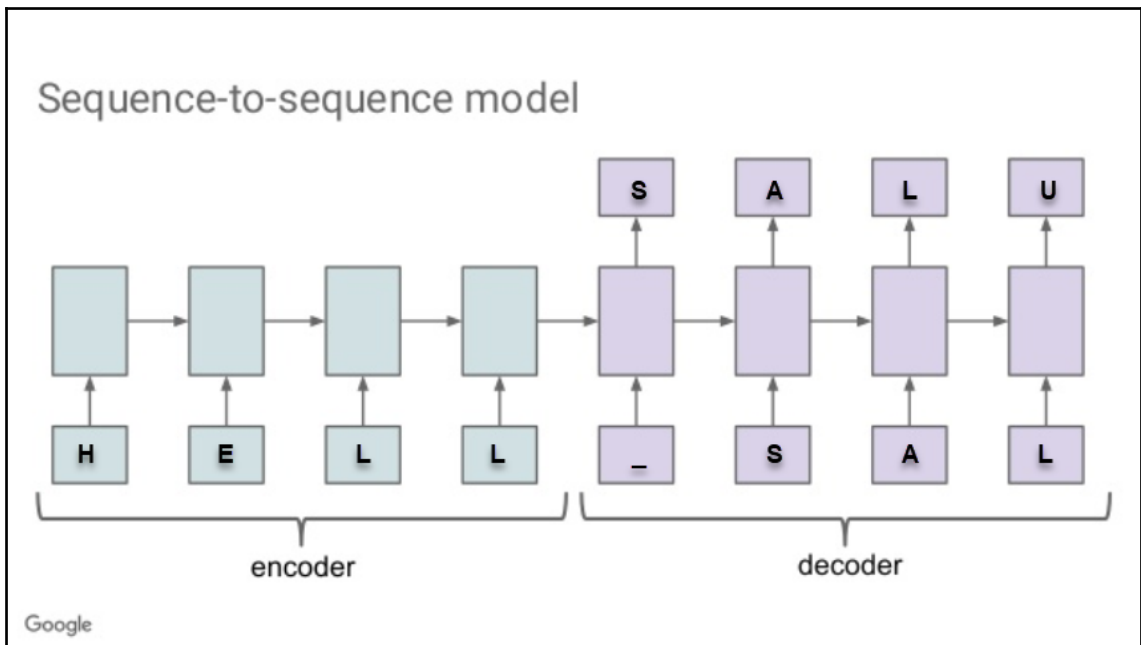



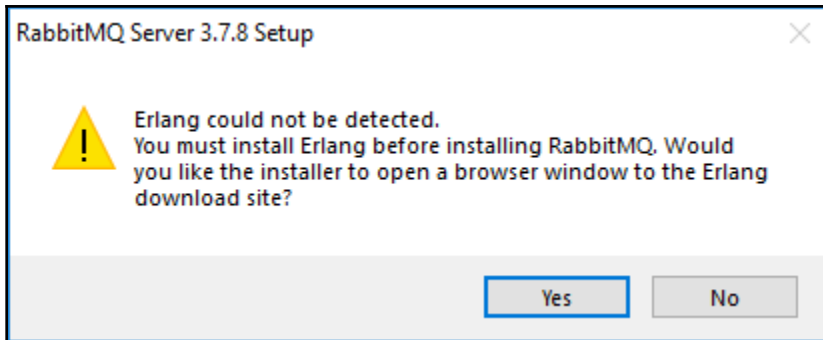
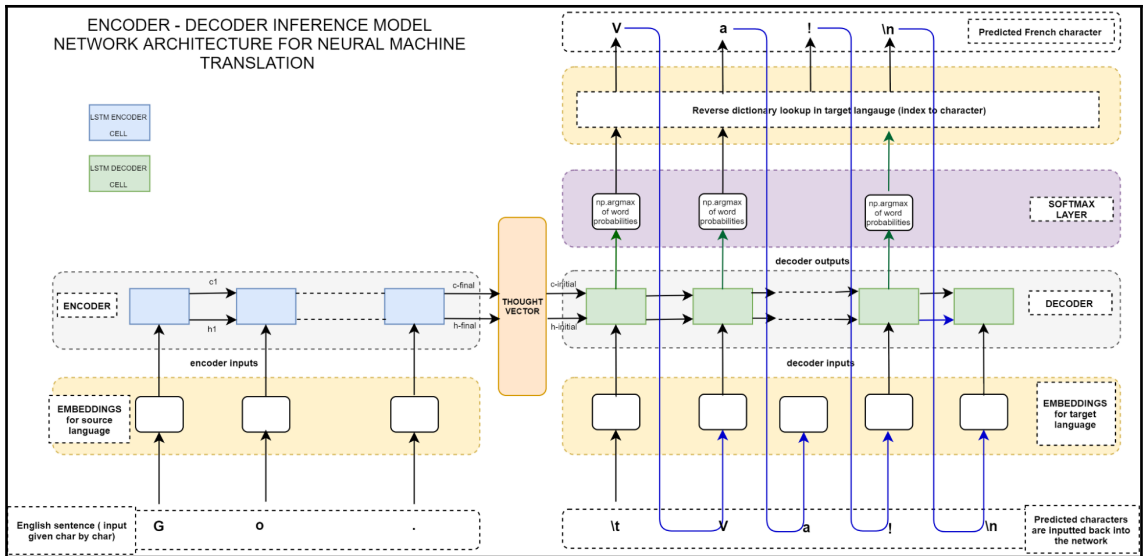




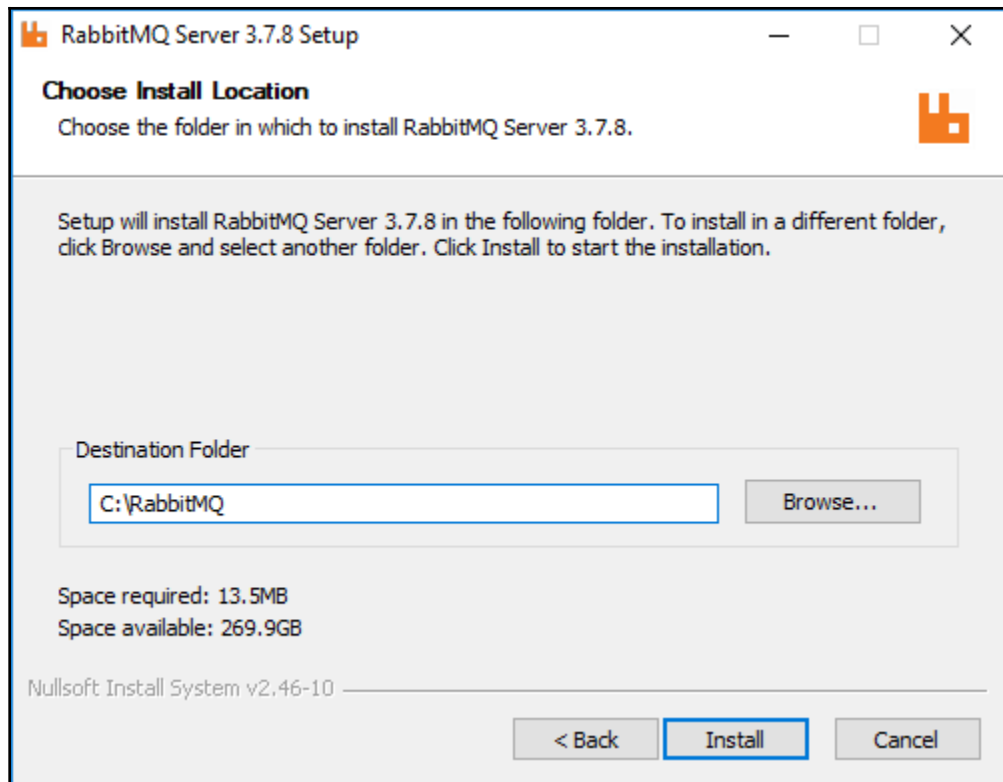


# Chapter 4: Building a Deep Learning Gaming Chatbot









```
C:\Windows\System32\cmd.exe
Microsoft Windows [Version 10.0.17134.286]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\RabbitMQ\rabbitmq_server-3.7.8\sbin>rabbitmq-plugins enable rabbitmq_management
Enabling plugins on node rabbit@DESKTOP-V2J9HRG:
rabbitmq_management
The following plugins have been configured:
  rabbitmq_management
  rabbitmq_management_agent
  rabbitmq_web_dispatch
Applying plugin configuration to rabbit@DESKTOP-V2J9HRG...
The following plugins have been enabled:
  rabbitmq_management
  rabbitmq_management_agent
  rabbitmq_web_dispatch

started 3 plugins.
```

RabbitMQ Management localhost:15672/#/

Refreshed 2018-10-21 14:04:24 Refresh every 5 seconds

3.7.8 Erlang 21.1

Virtual host All

Cluster rabbit@DESKTOP-V2J9HRG

User guest Log out

Overview Connections Channels Exchanges Queues Admin

## Overview

▼ Totals

Queued messages last minute ?

Currently idle

Message rates last minute ?

Currently idle

Global counts ?

Connections: 0 Channels: 0 Exchanges: 7 Queues: 0 Consumers: 0

▼ Nodes

Name	File descriptors ?	Socket descriptors ?	Erlang processes	Memory ?	Disk space	Uptime	Info	Reset stats	+/-
rabbit@DESKTOP-V2J9HRG	0 8192 available	0 7280 available	401 1048576 available	82MB 6.4GB high watermark	270GB 48MB low watermark	39m 20s	basic disc 1 rss	This node All nodes	

► Ports and contexts

► Export definitions

► Import definitions

HTTP API Server Docs Tutorials Community Support Community Slack Commercial Support Plugins GitHub Changelog

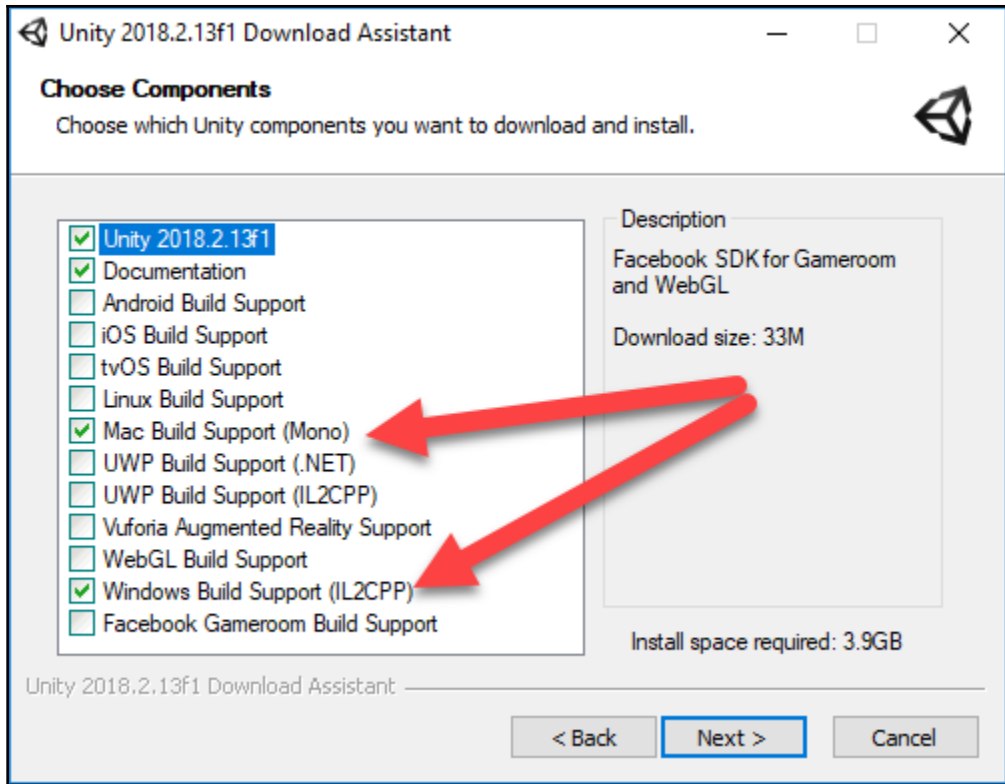
Ready: 1, Unacked: 0, Total: 1

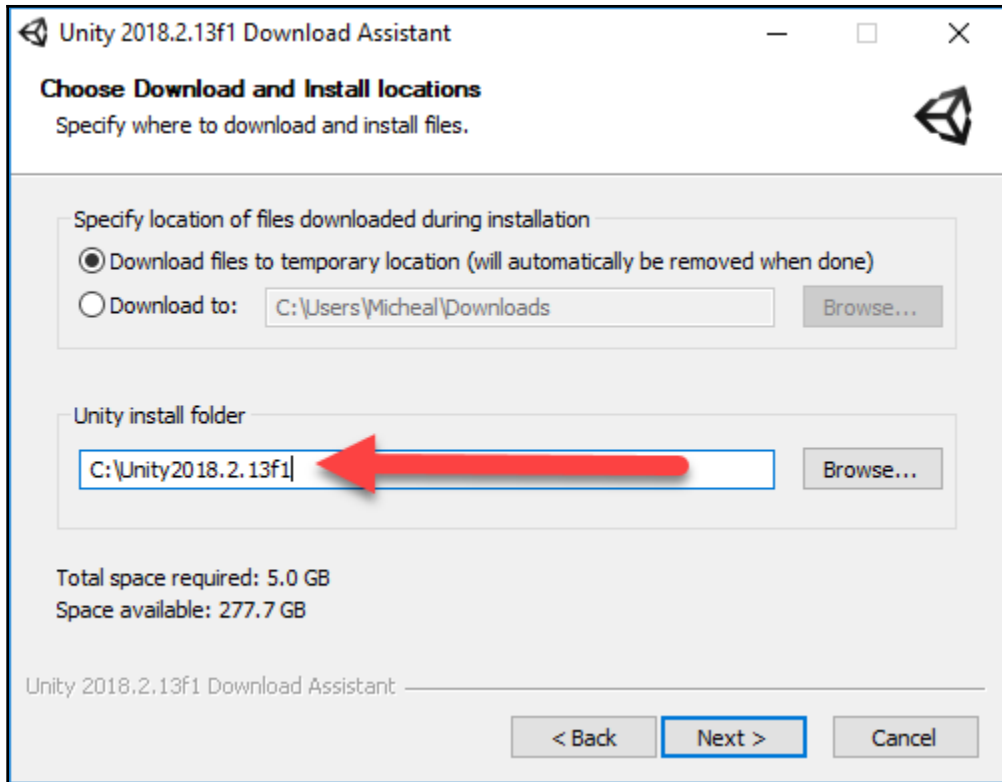
Message rates last minute ?

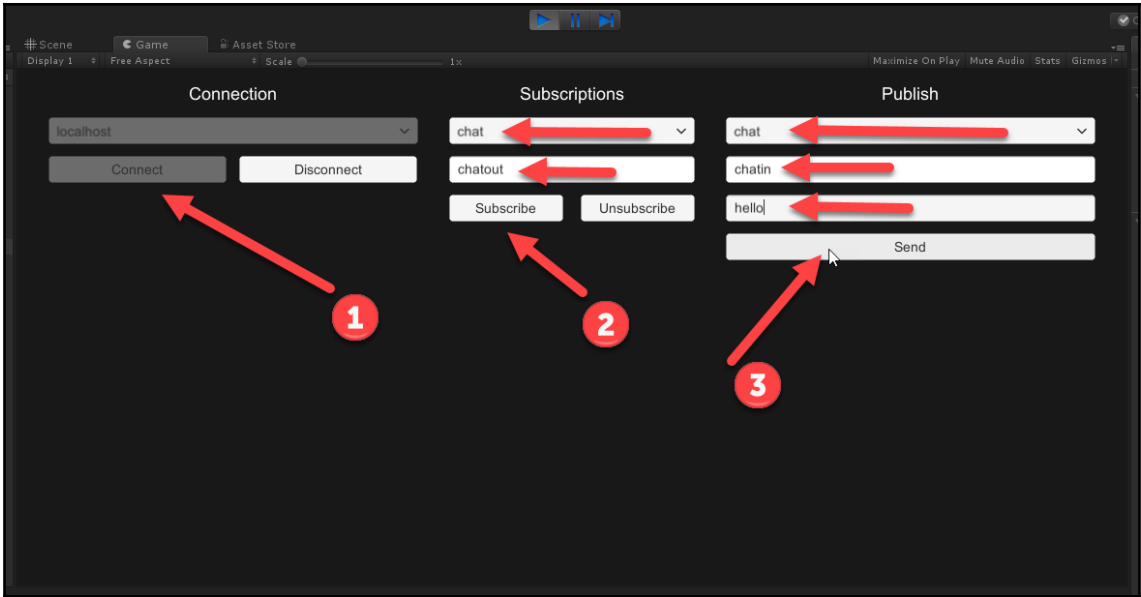
Publish: 0.00/s, Publisher confirm: 0.00/s, Return: 0.00/s

Disk read: 0.00/s, Disk write: 0.00/s

Global counts ?

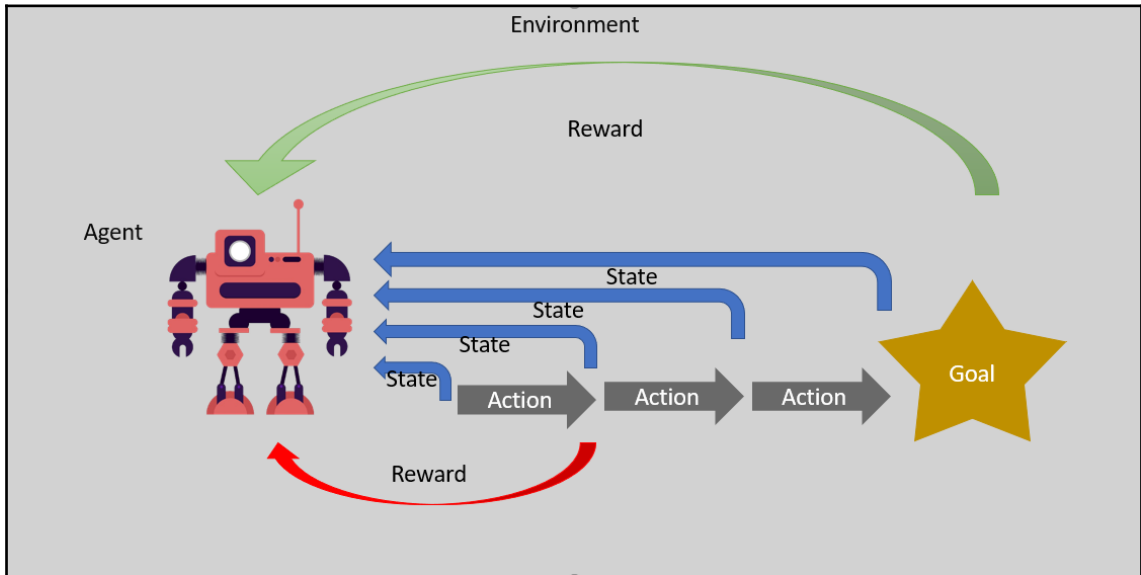


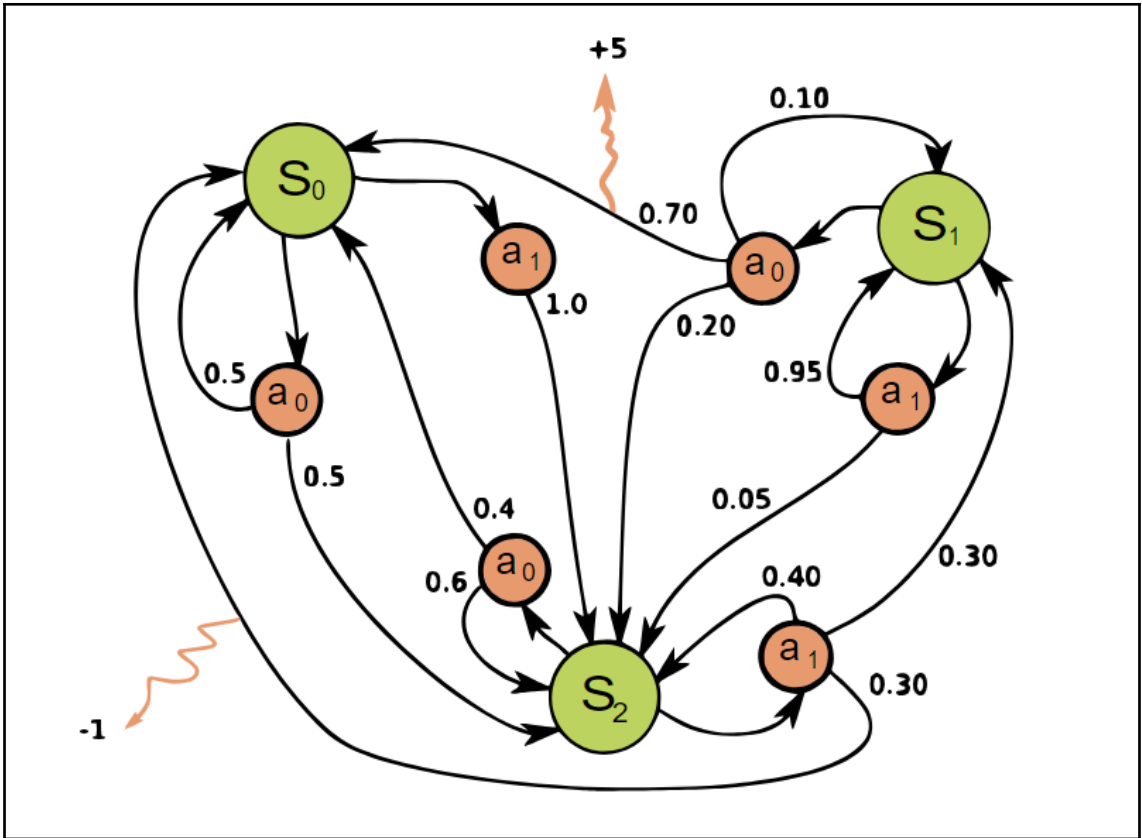




---

## Chapter 5: Introducing DRL






(Left)


```

SFFF
FHFH ←
FFFH
HFFG
-


```

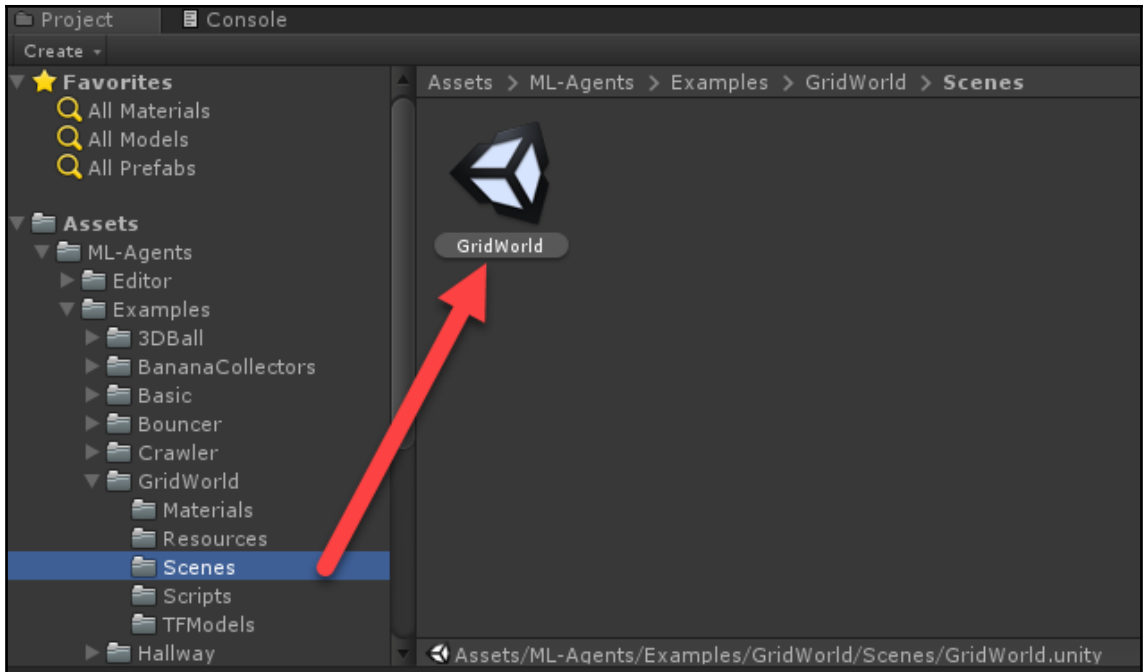


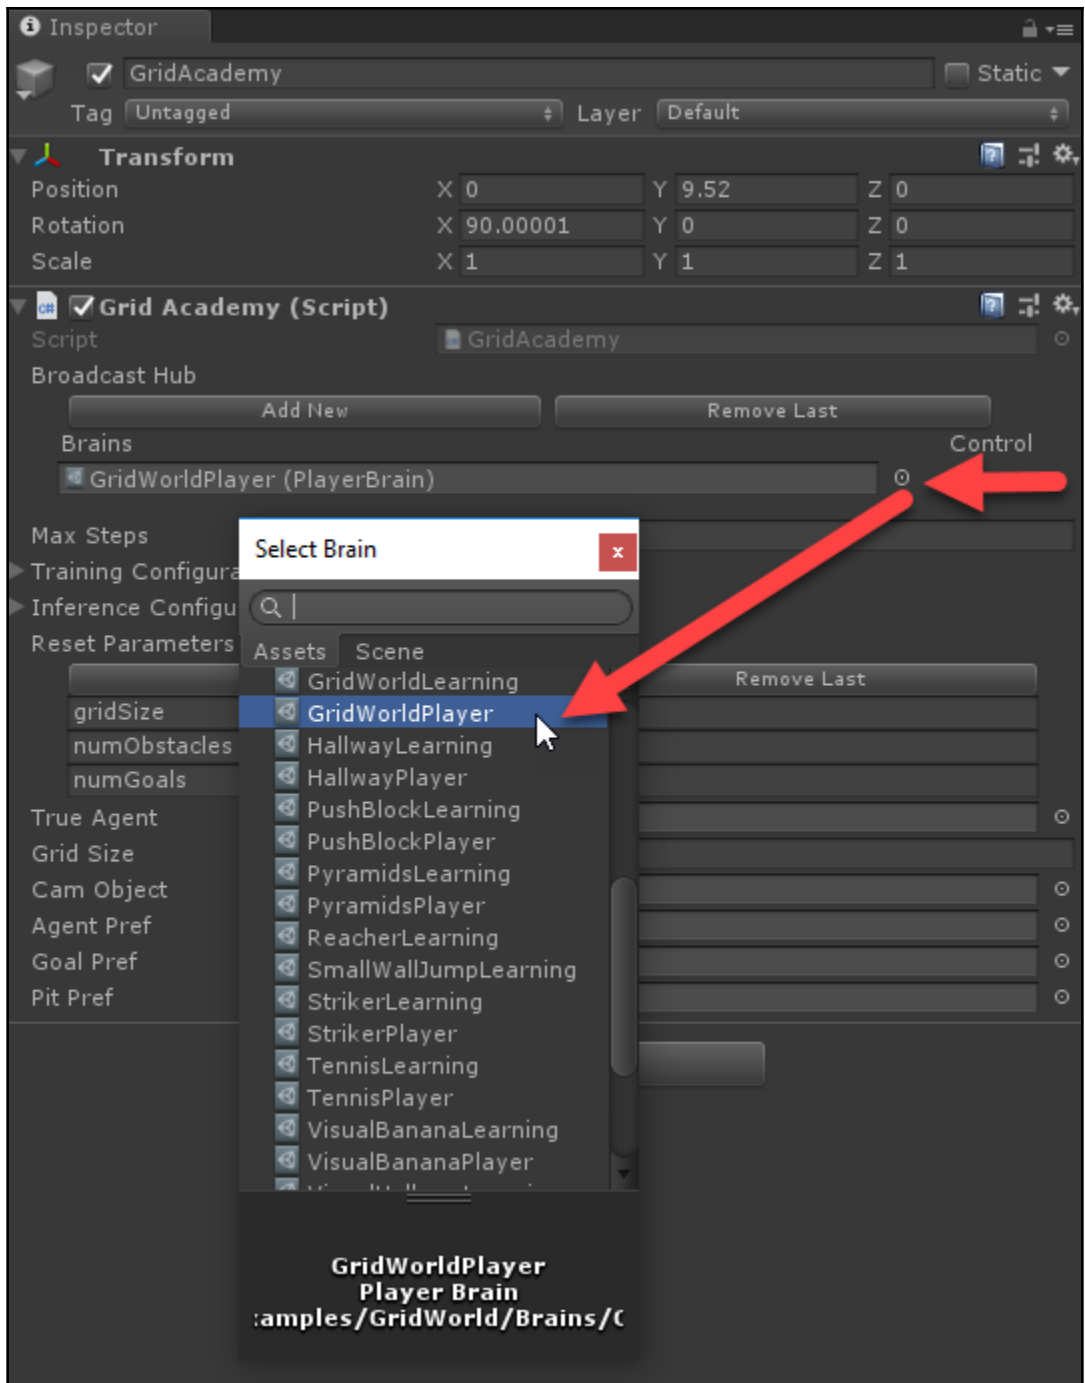


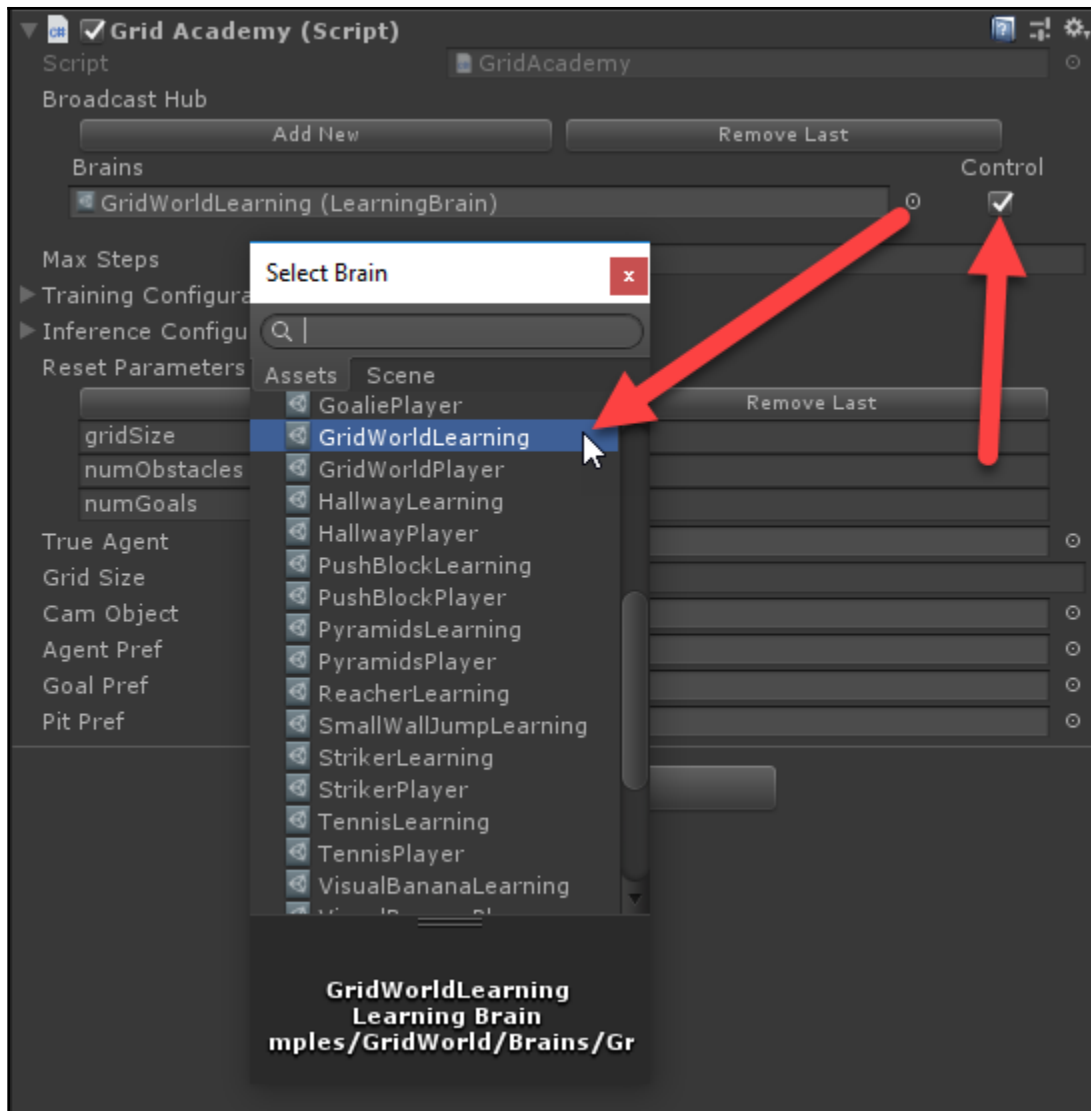


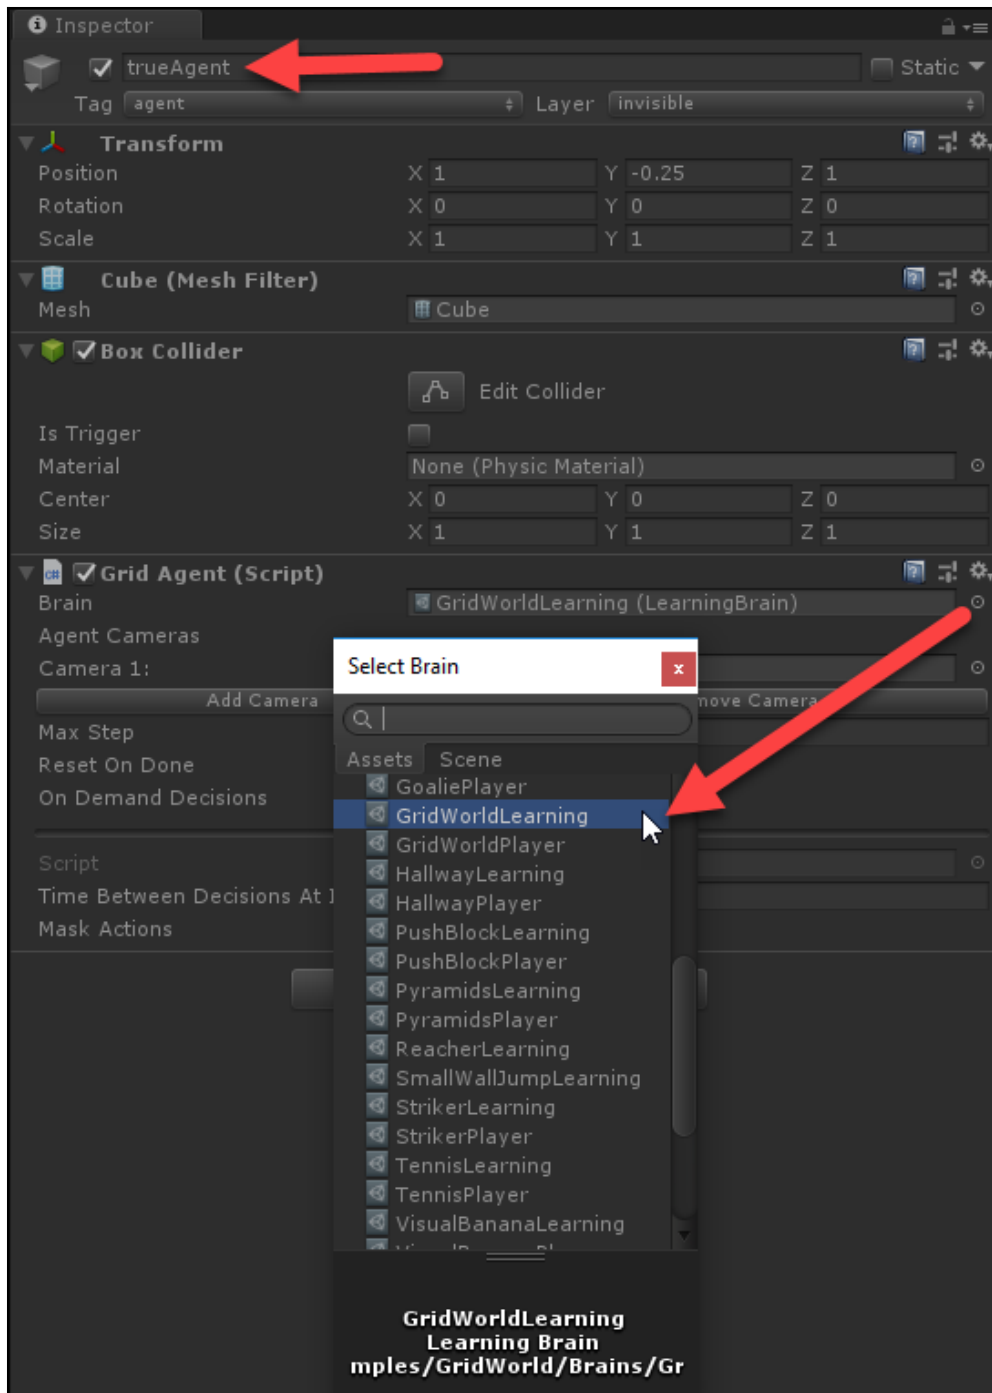
# Chapter 6: Unity ML-Agents




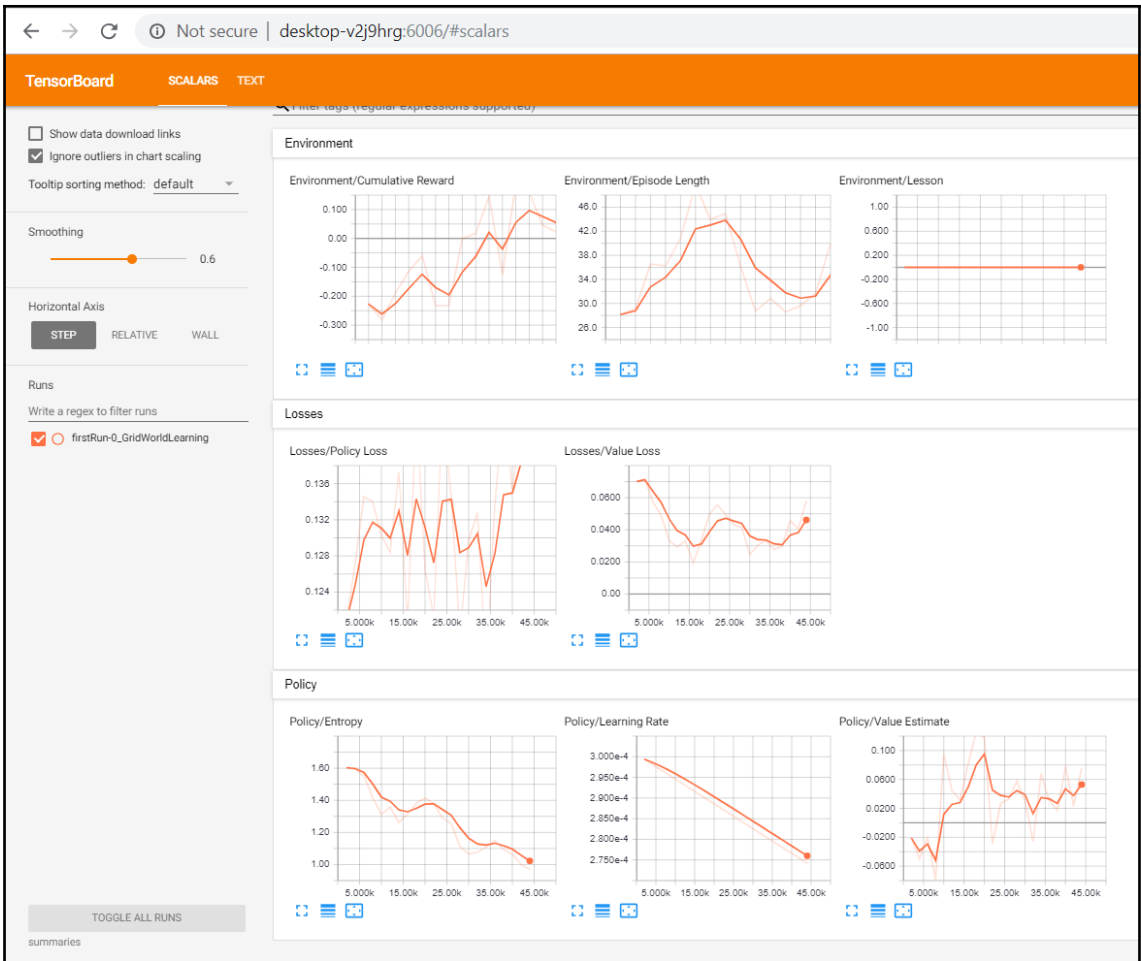









```
Anaconda Prompt - mlagents-learn config/trainer_config.yaml --run-id=firstRun --train
--worker-id: '0',
<trainer-config-path>: 'config/trainer_config.yaml')
INFO:mlagents.envs:Start training by pressing the Play button in the Unity Editor.
INFO:mlagents.envs:
'GridAcademy' started successfully!
Unity Academy name: GridAcademy
  Number of Brains: 1
  Number of Training Brains : 1
  Reset Parameters :
    numGoals -> 1.0
    numObstacles -> 1.0
    gridSize -> 5.0
Unity brain name: GridWorldLearning
  Number of Visual Observations (per agent): 1
  Vector Observation space size (per agent): 0
  Number of stacked Vector Observation: 1
  Vector Action space type: discrete
  Vector Action space size (per agent): [5]
  Vector Action descriptions:
2019-01-07 15:48:18.862448: I T:\src\github\tensorflow\tensorflow\core\platform\cpu_feature_guard.cc:140] Your CPU supports instructions that this TensorFlow binary was not compiled to use: AVX2
INFO:mlagents.envs:Hyperparameters for the PPO Trainer of brain GridWorldLearning:
  batch_size: 32
  beta: 0.005
  buffer_size: 256
  epsilon: 0.2
  gamma: 0.9
  hidden_units: 256
  lambda: 0.95
  learning_rate: 0.0003
  max_steps: 5.0e5
  normalize: False
  num_epoch: 3
  num_layers: 1
  time_horizon: 5
  sequence_length: 64
  summary_freq: 2000
  use_recurrent: False
  summary_path: /summaries/firstRun-0_GridWorldLearning
  memory_size: 256
  use_curiosity: False
  curiosity_strength: 0.01
  curiosity_enc_size: 128
  model_path: /models/firstRun-0/GridWorldLearning
INFO:mlagents.trainers: firstRun-0: GridWorldLearning: Step: 2000. Mean Reward: -0.227. Std of Reward: 1.029. Training.
INFO:mlagents.trainers: firstRun-0: GridWorldLearning: Step: 4000. Mean Reward: -0.282. Std of Reward: 1.002. Training.
INFO:mlagents.trainers: firstRun-0: GridWorldLearning: Step: 6000. Mean Reward: -0.190. Std of Reward: 1.016. Training.
```






## Losses

### Losses/Policy Loss




### Losses/Value Loss




## Policy


### Policy/Entropy

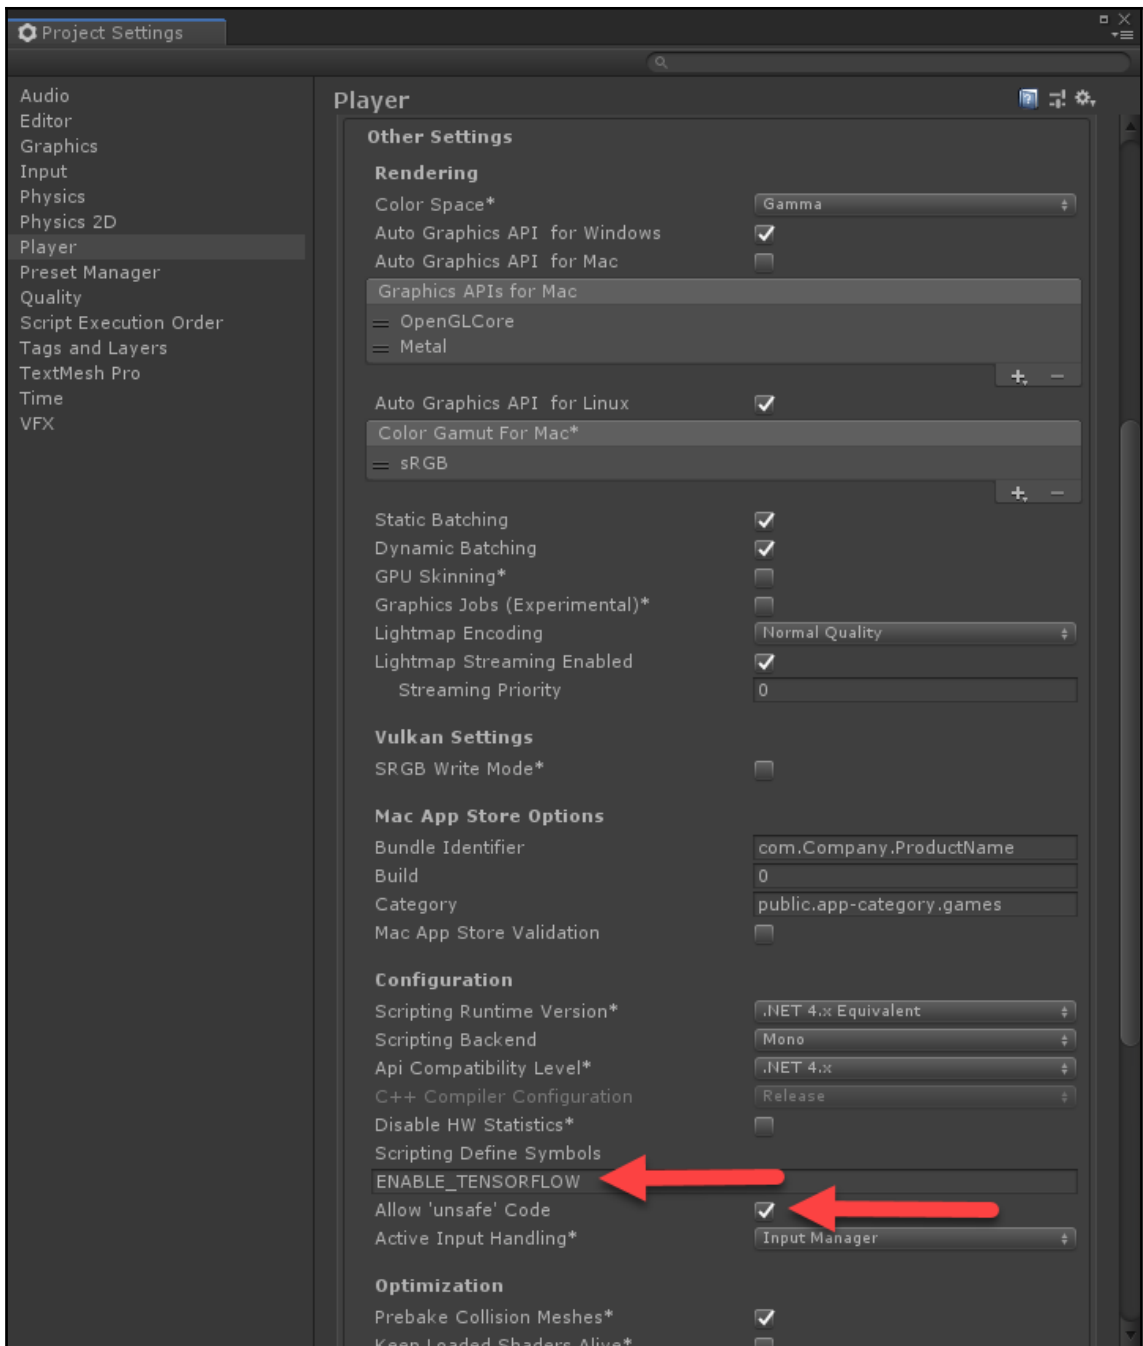


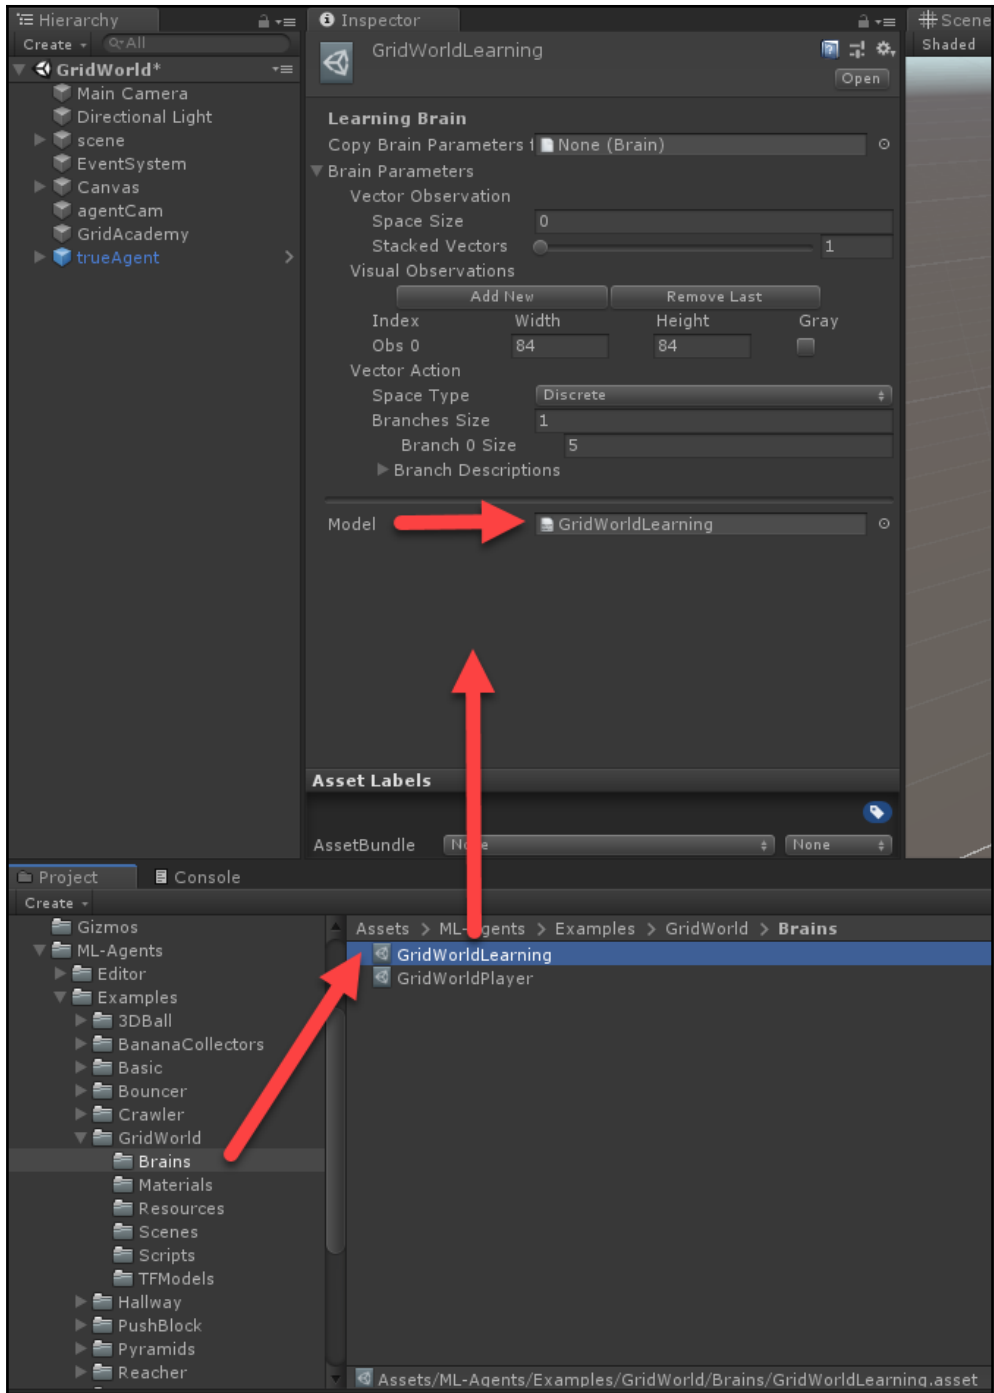
### Policy/Learning Rate

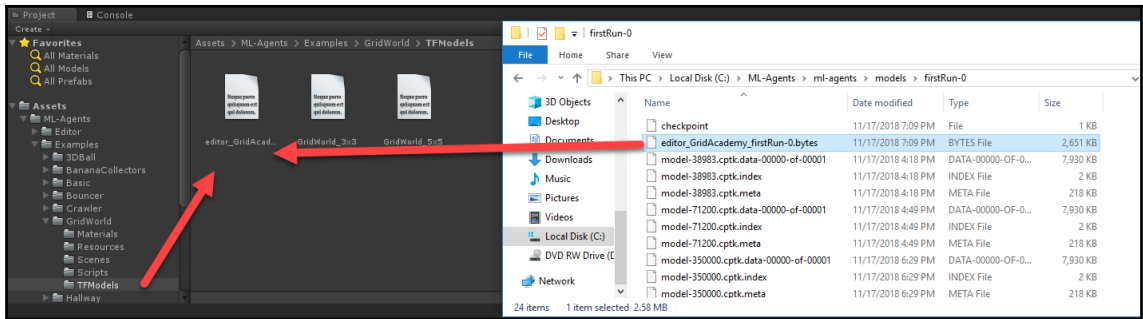


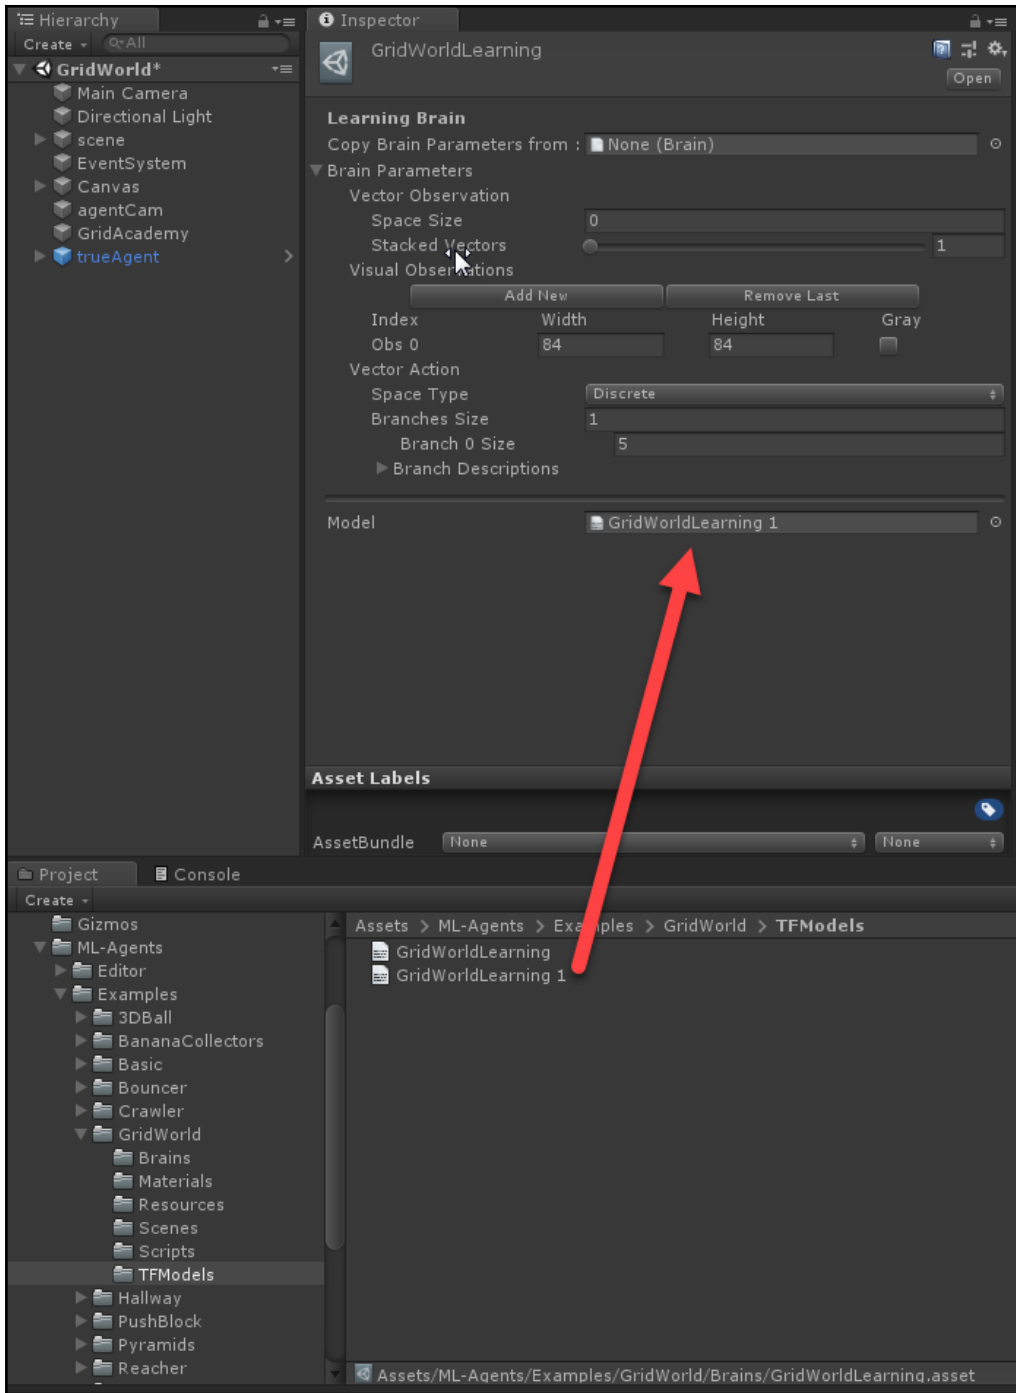
### Policy/Value Estimate





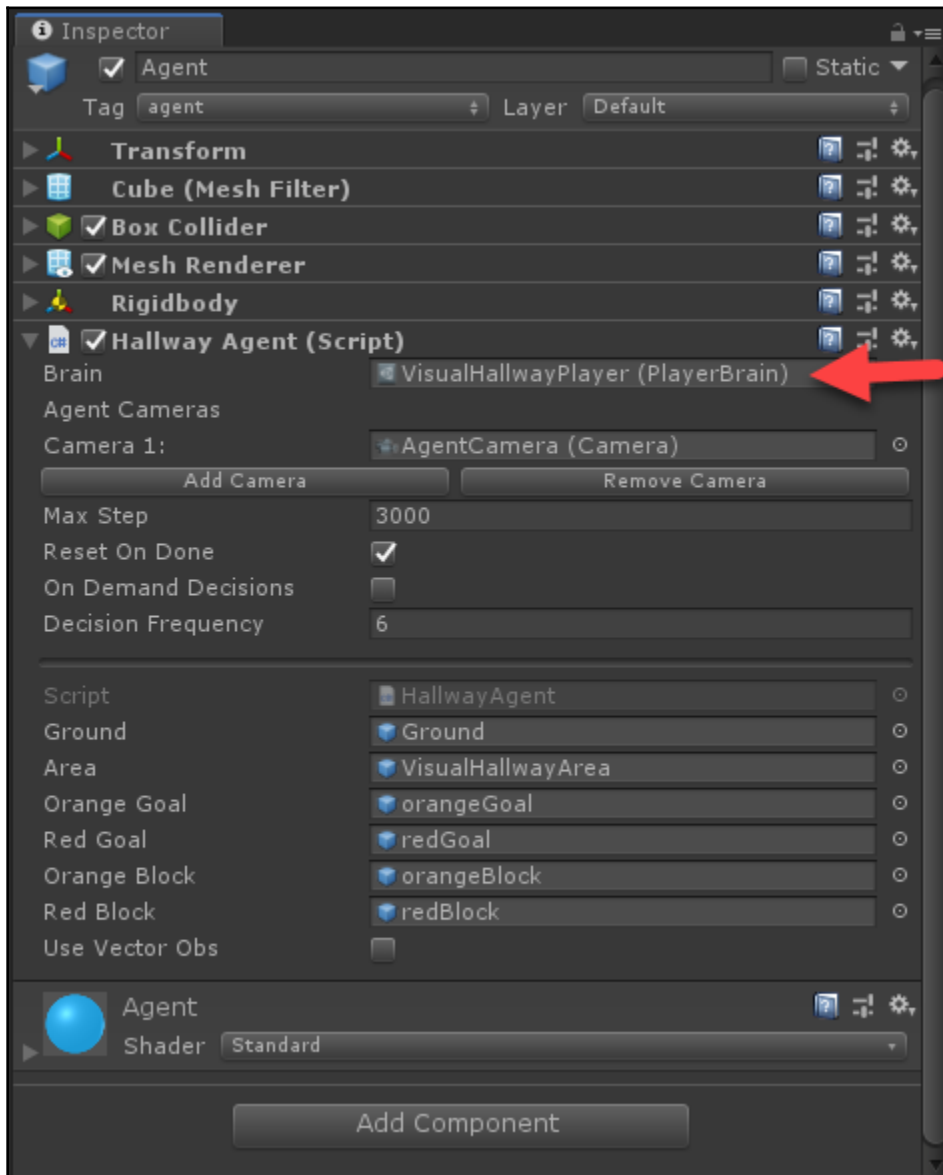


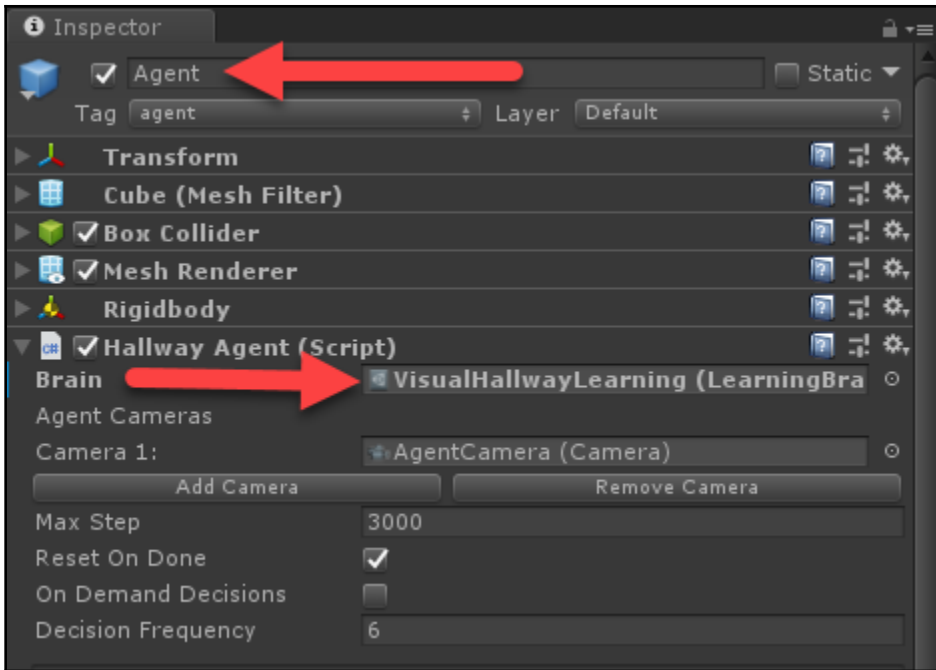


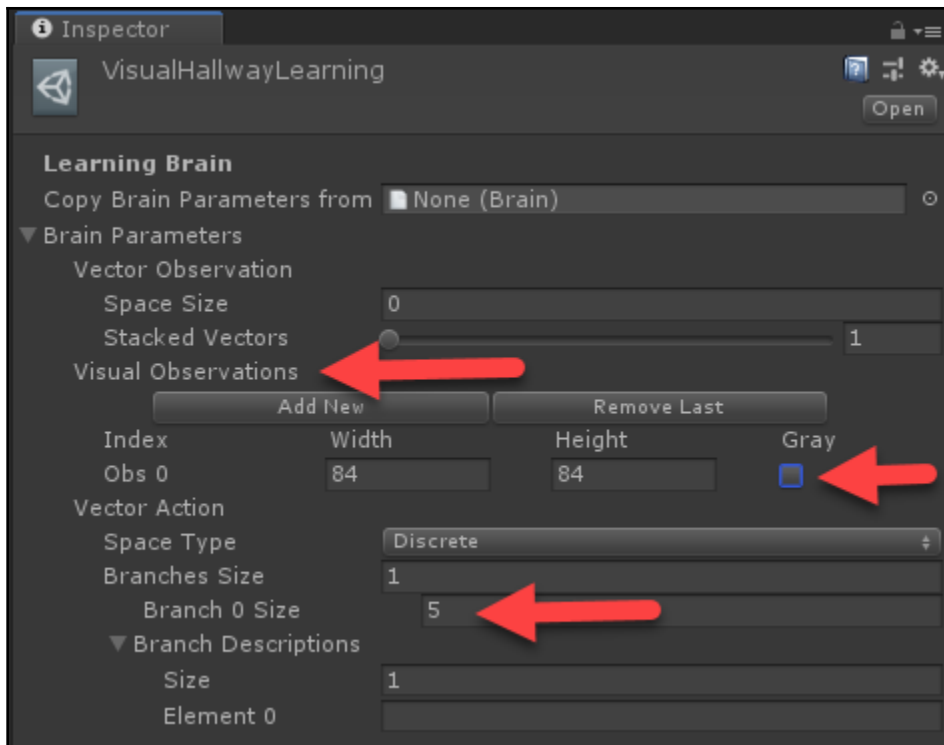


---

## Chapter 7: Agent and the Environment




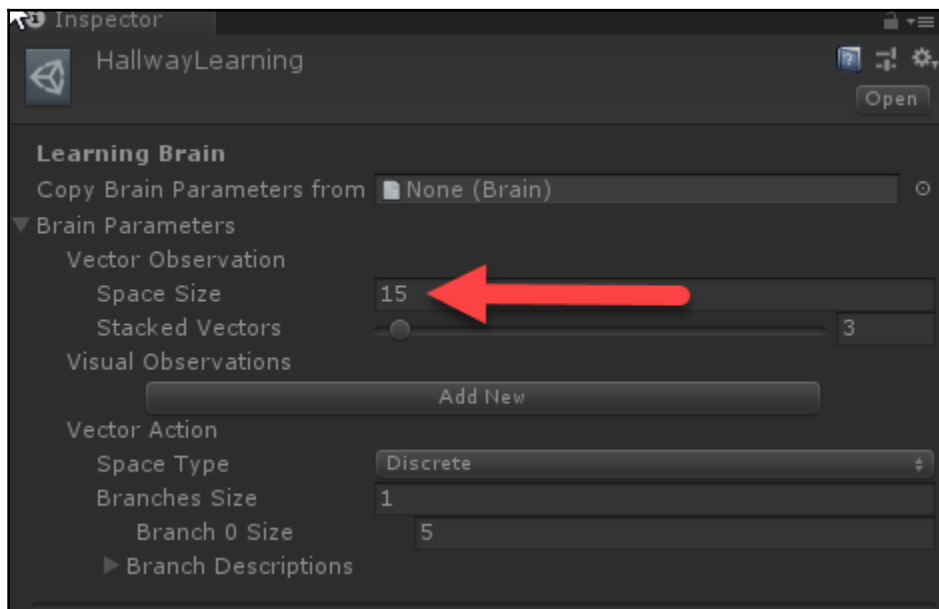
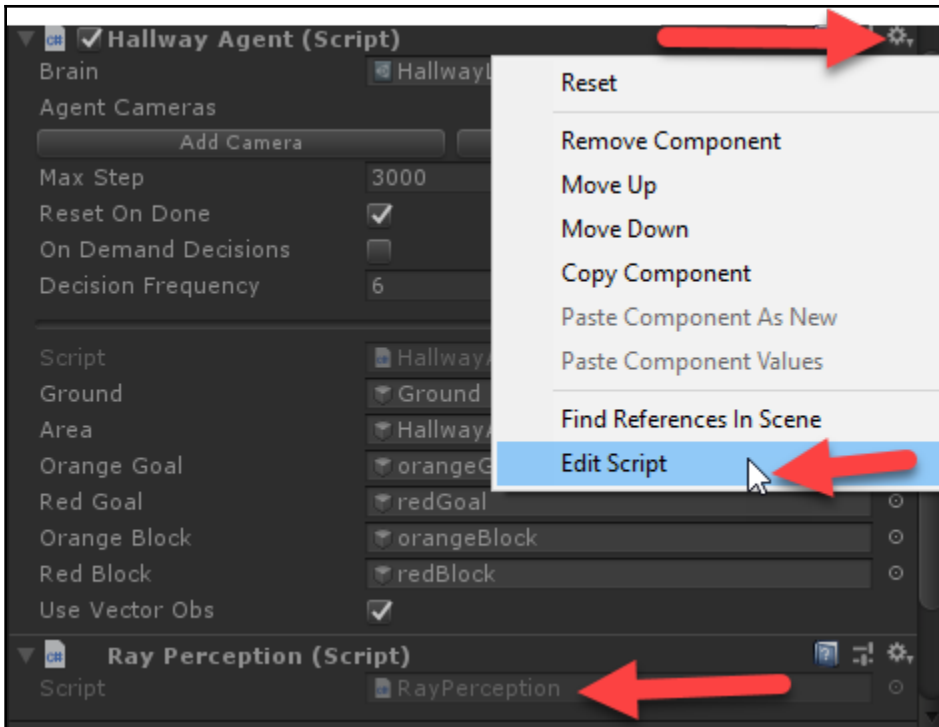


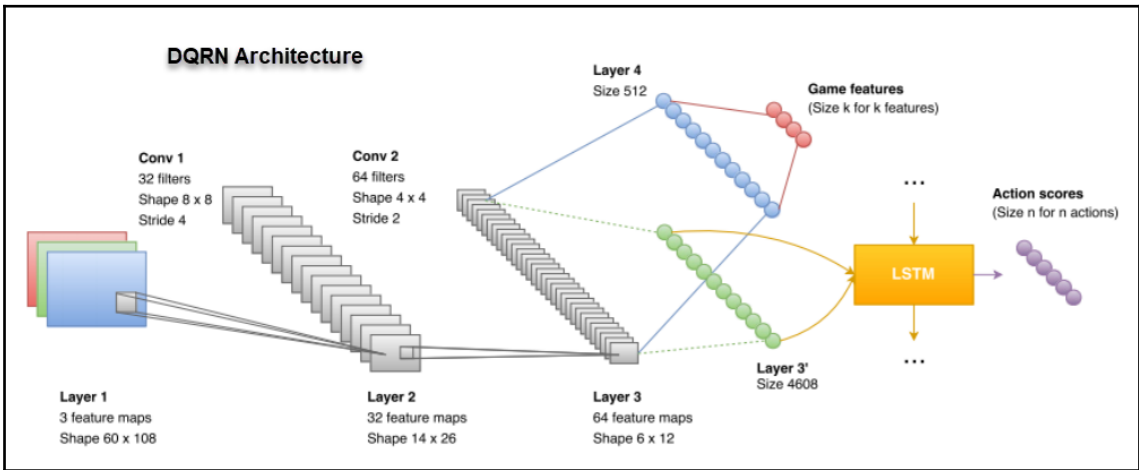
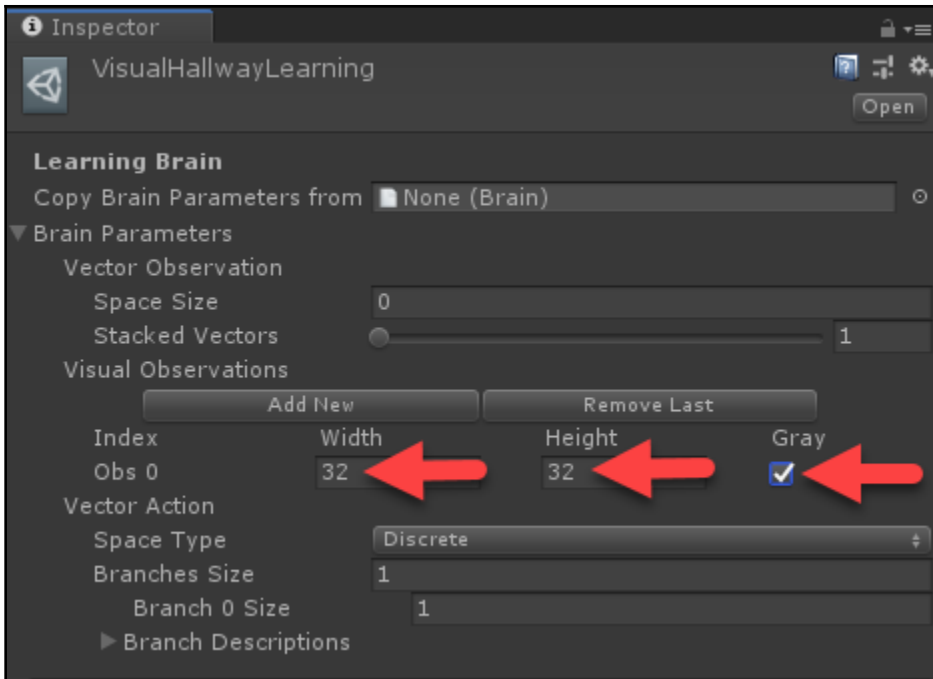




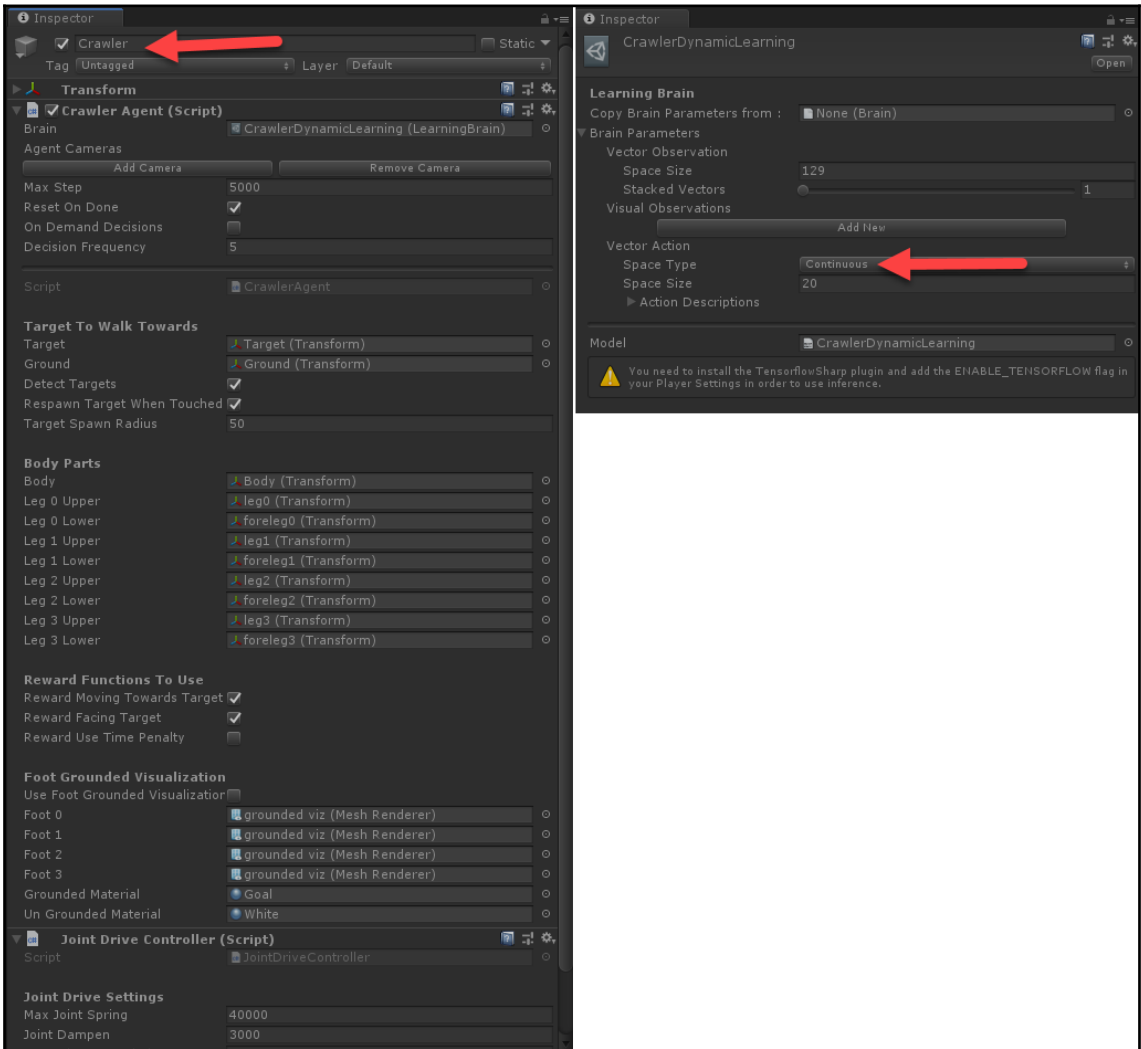
```
Anaconda Prompt
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 471000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 472000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 473000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 474000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 475000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 476000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 477000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 478000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 479000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 480000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 481000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 482000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 483000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 484000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 485000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 486000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 487000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 488000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 489000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 490000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 491000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 492000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 493000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 494000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 495000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 496000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 497000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 498000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 499000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.envs: Saved Model
NFO:magents.trainers: visualhallway2-0: VisualHallwayLearning: Step: 500000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
NFO:magents.envs: Saved Model
NFO:magents.trainers: List of nodes to export for brain :VisualHallwayLearning
NFO:magents.trainers: is_continuous_control
NFO:magents.trainers: version_number
NFO:magents.trainers: memory_size
NFO:magents.trainers: action_output_shape
NFO:magents.trainers: recurrent_out
NFO:magents.trainers: action_probs
NFO:magents.trainers: action
NFO:magents.trainers: value_estimate
NFO:tensorflow: Restoring parameters from ./models/visualhallway2-0/VisualHallwayLearning\model-500001.cptk
NFO:tensorflow: Froze 15 variables.
onverted 15 variables to const ops.
```





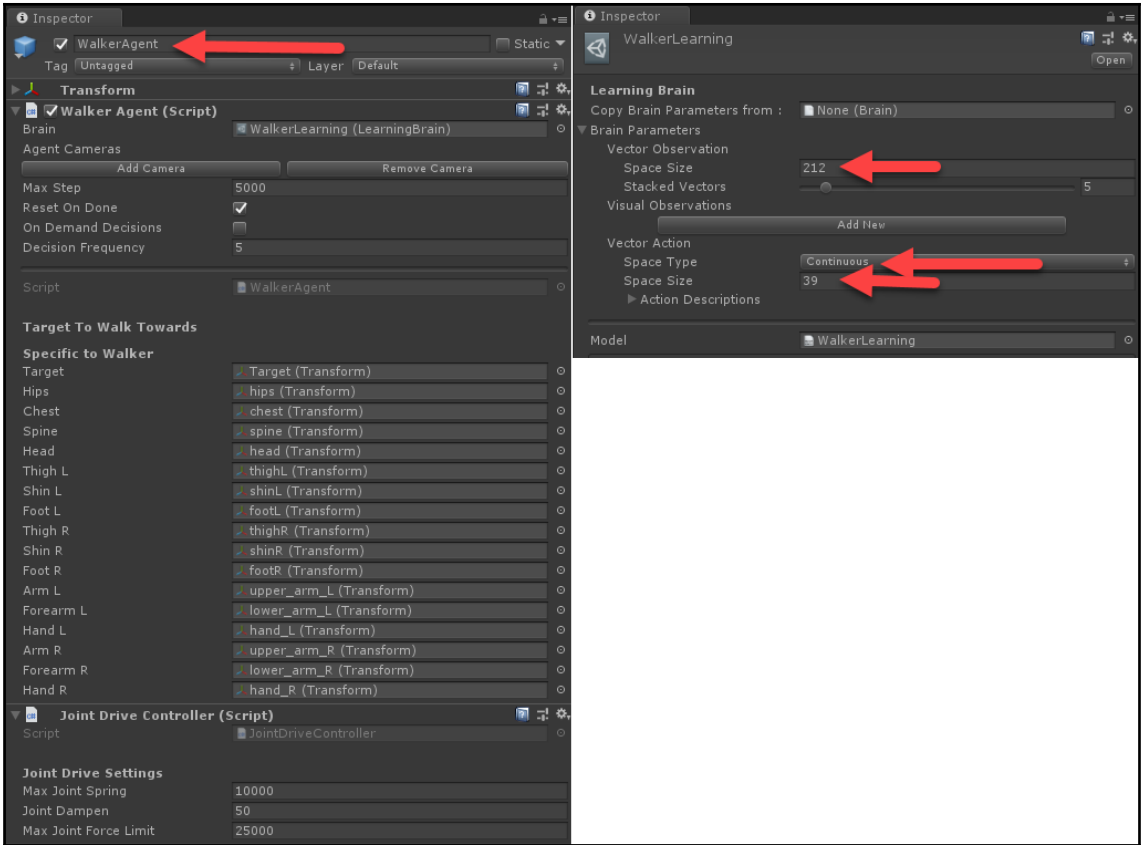




# Chapter 8: Understanding PPO

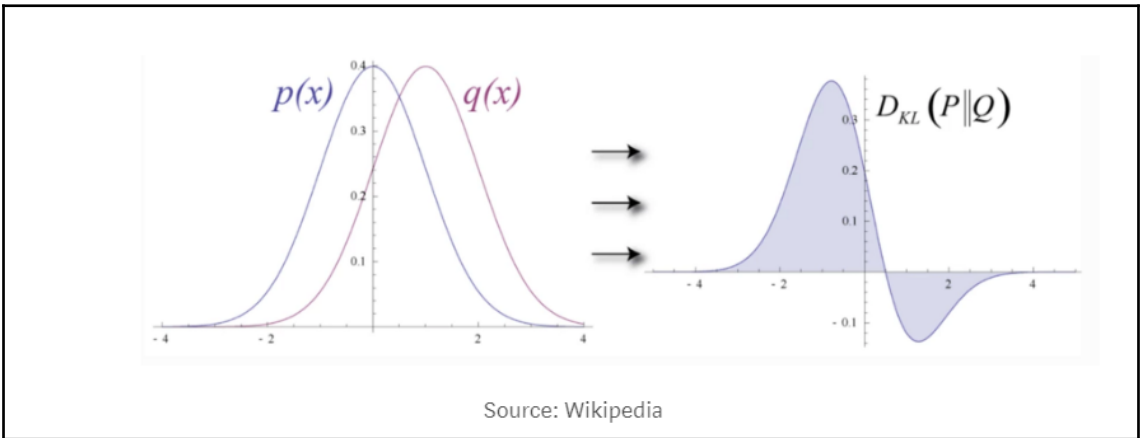







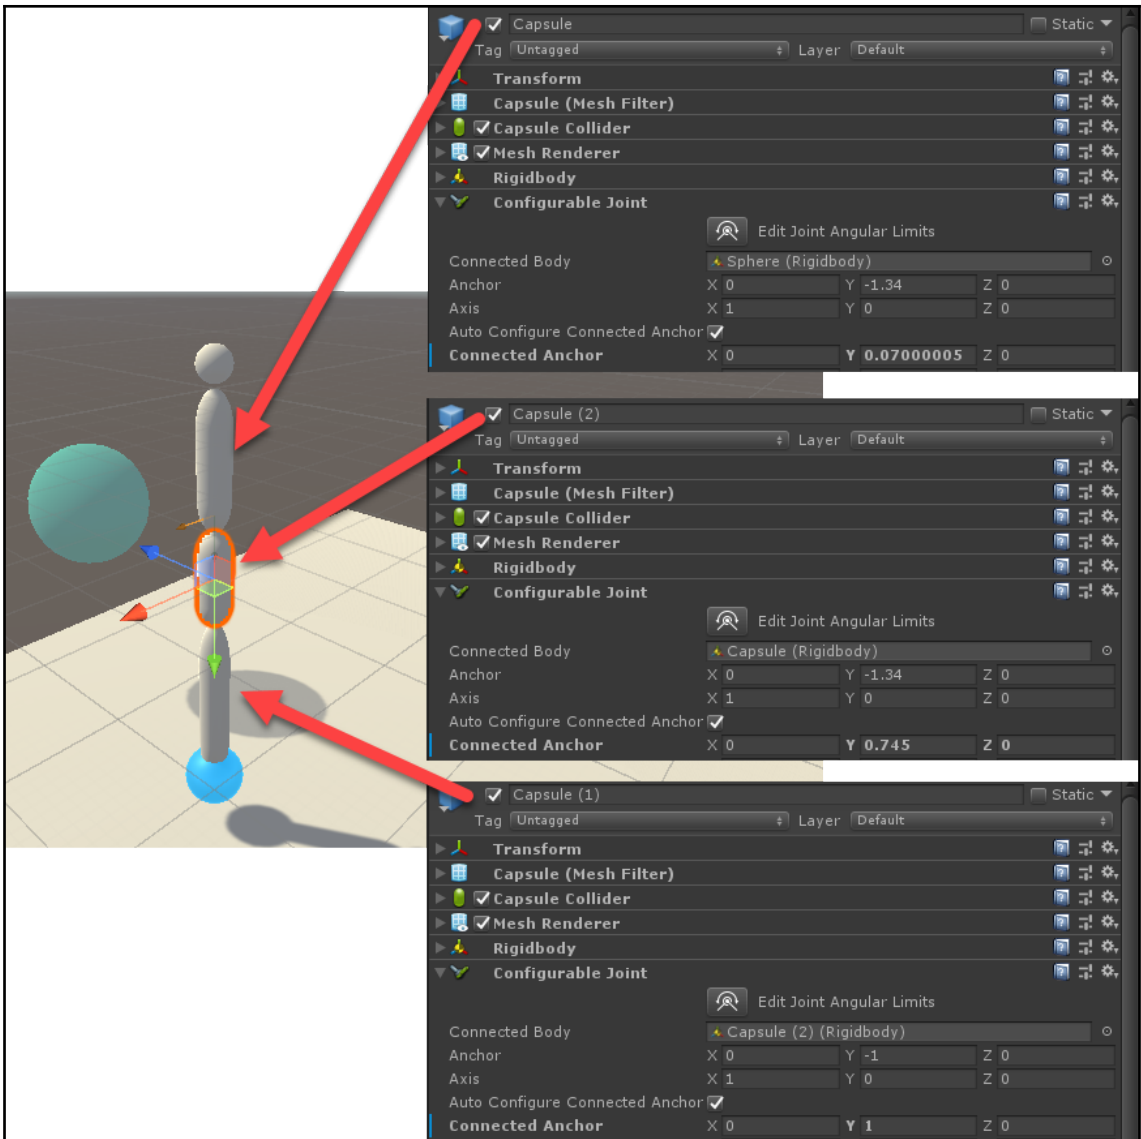


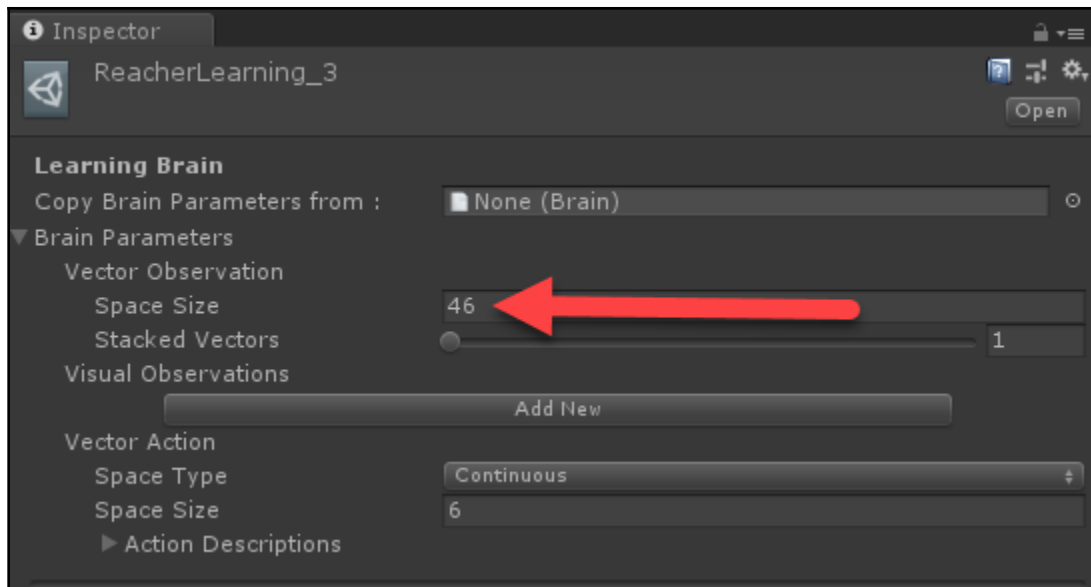






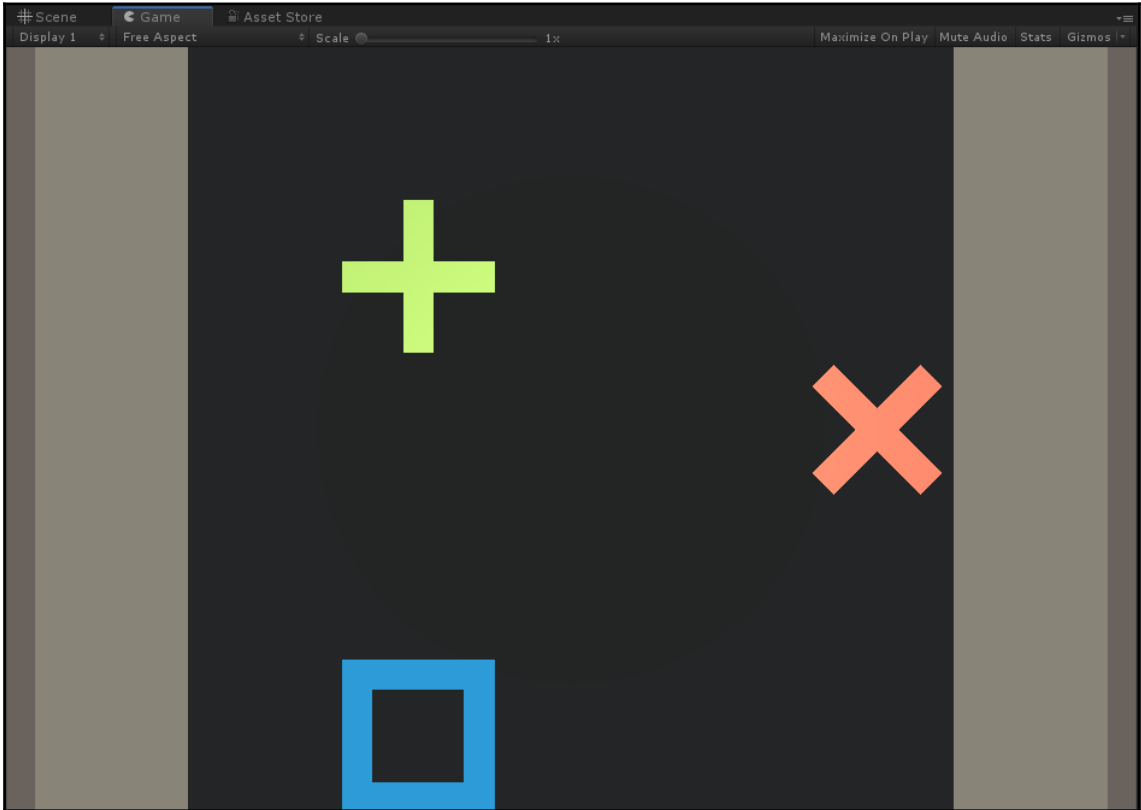


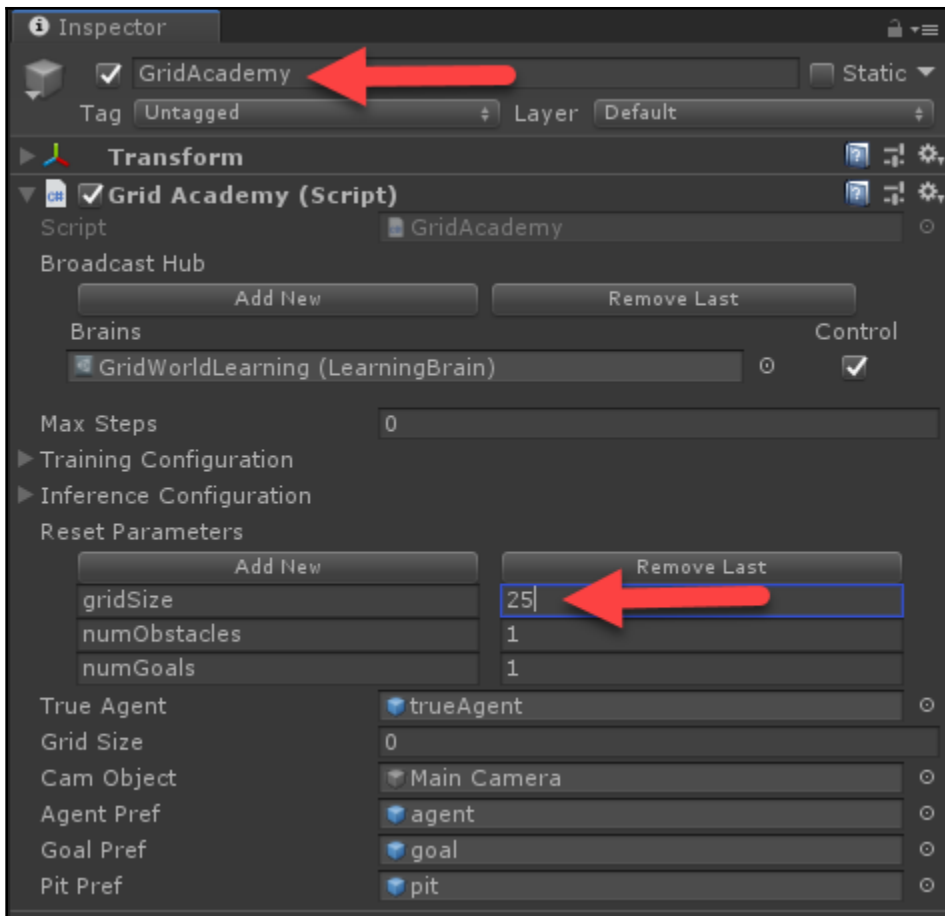





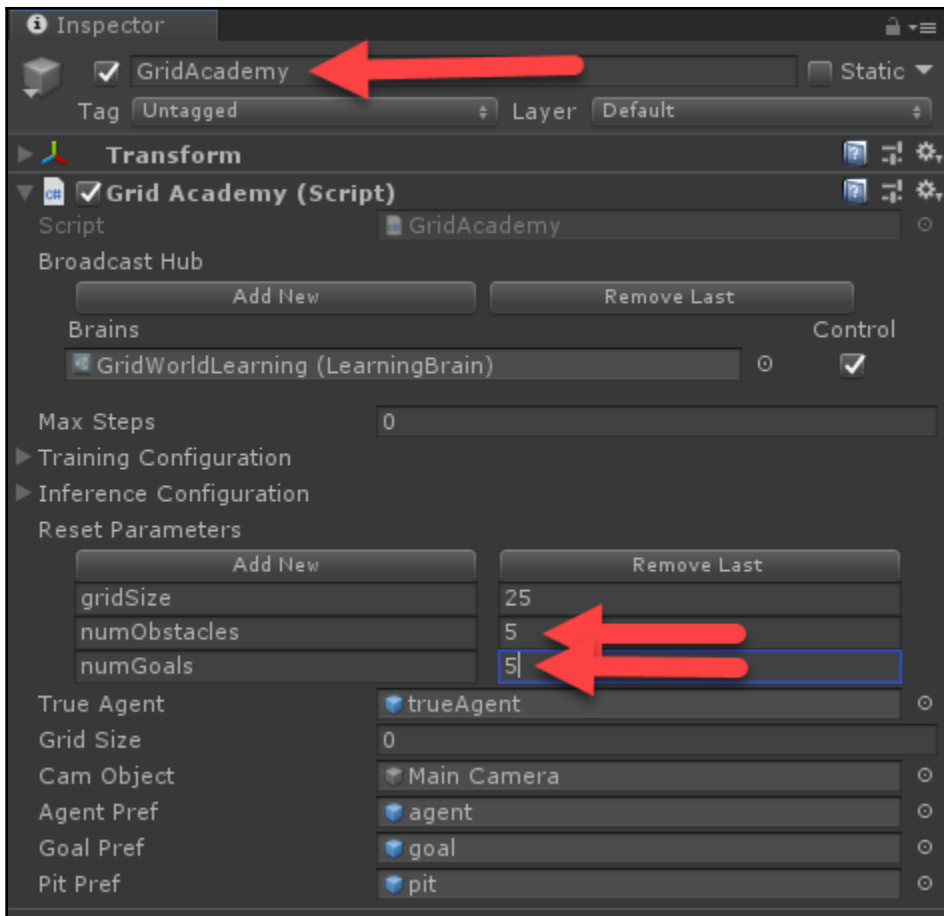
---

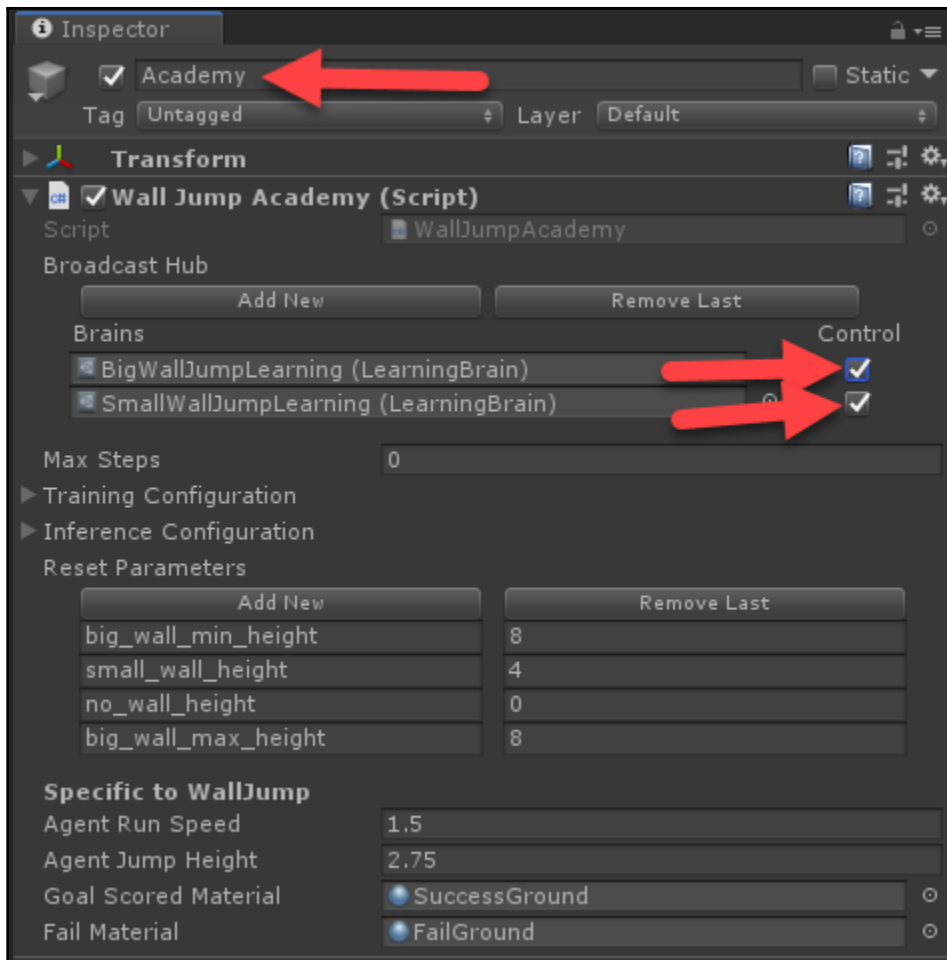
# Chapter 9: Rewards and Reinforcement Learning



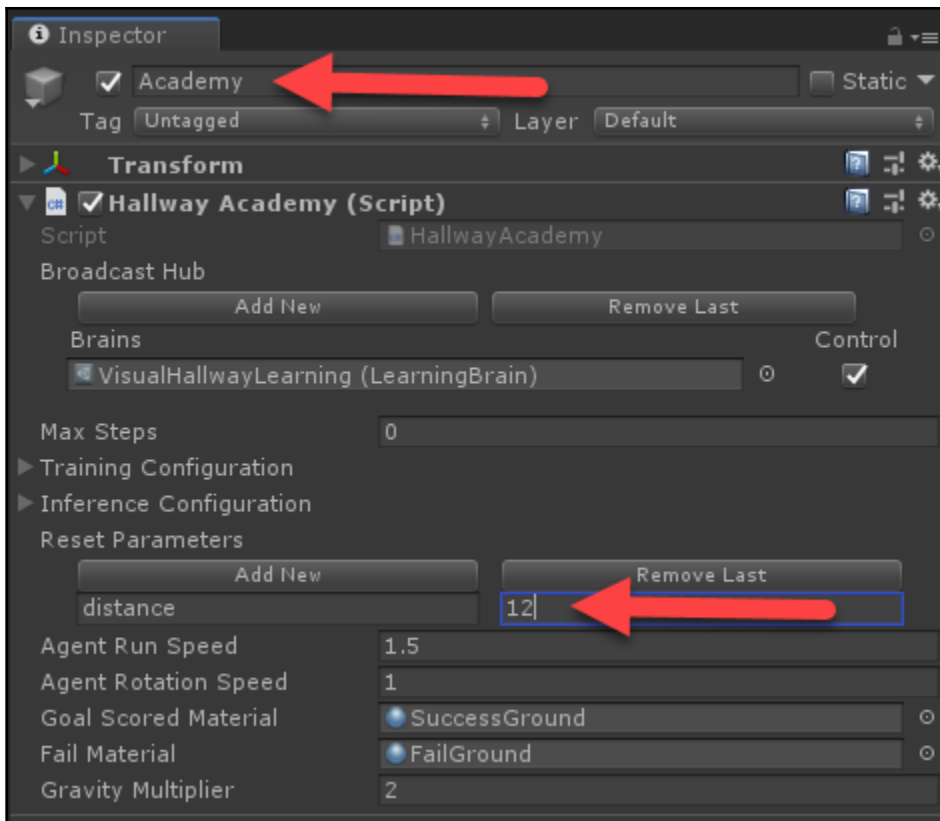




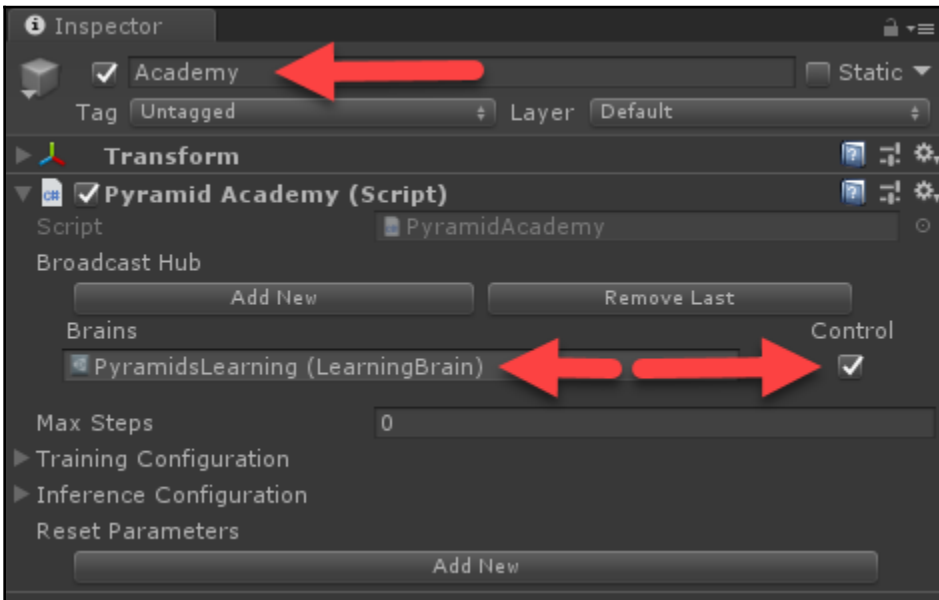
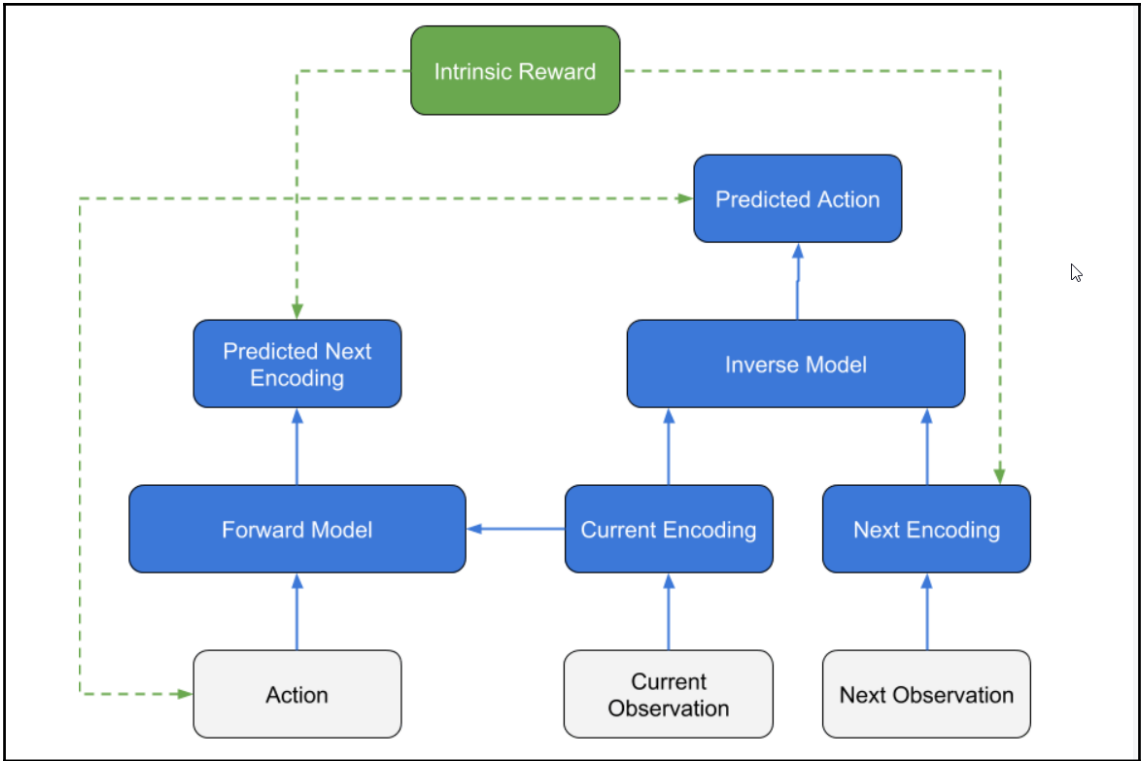






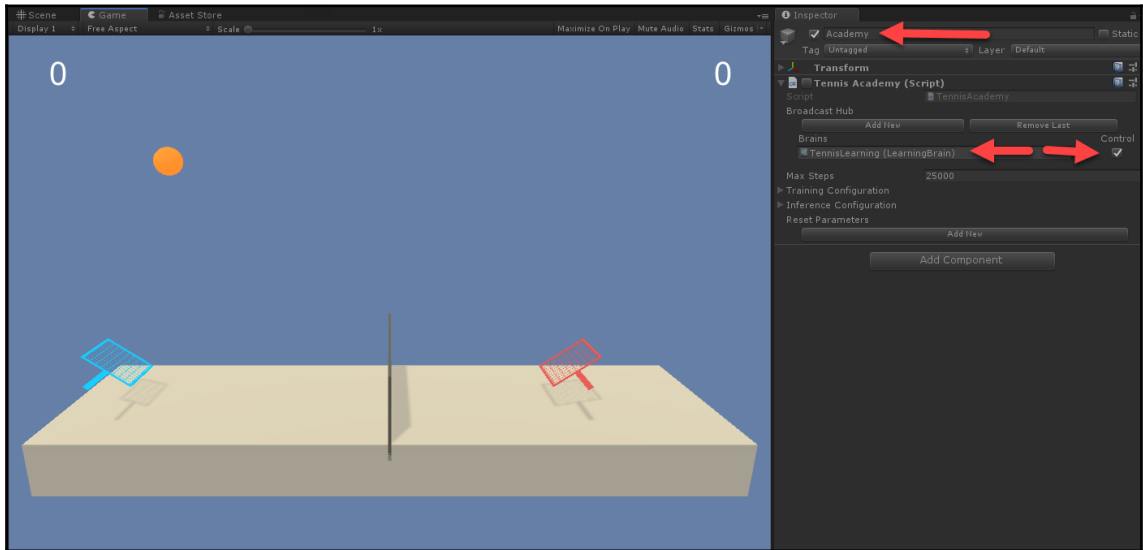



```
Anaconda Prompt - mlagents-learn config/trainer_config.yaml --curriculum=config/curricula/hallway/ --run-id=hallway-curriculum --train
Number of stacked Vector Observation: 1
Vector Action space type: discrete
Vector Action space size (per agent): [1]
Vector Action descriptions:
2019-02-03 21:38:15.126765: I T:\src\github\tensorflow\tensorflow\core\platform\cpu_feature_guard.cc:140] Your CPU supports instructions that th
is TensorFlow binary was not compiled to use: AVX2
INFO:mlagents.envs:Hyperparameters for the PPO Trainer of brain VisualHallwayLearning:
  batch_size: 64
  beta: 0.01
  buffer_size: 1024
  epsilon: 0.2
  gamma: 0.99
  hidden_units: 128
  lambda: 0.95
  learning_rate: 0.0003
  max_steps: 5.0e5
  normalize: False
  num_epoch: 3
  num_layers: 1
  time_horizon: 64
  sequence_length: 64
  summary_freq: 1000
  use_recurrent: True
  summary_path: ./summaries/hallway-curriculum-0_VisualHallwayLearning
  memory_size: 256
  use_curiosity: False
  curiosity_strength: 0.01
  curiosity_enc_size: 128
  model_path: ./models/hallway-curriculum-0/VisualHallwayLearning
INFO:mlagents.envs:Academy reset with parameters: distance -> 12
INFO:mlagents.trainers: hallway-curriculum-0: VisualHallwayLearning: Step: 1000. Mean Reward: -1.002. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: hallway-curriculum-0: VisualHallwayLearning: Step: 2000. Mean Reward: -0.336. Std of Reward: 0.939. Training.
INFO:mlagents.trainers: hallway-curriculum-0: VisualHallwayLearning: Step: 3000. Mean Reward: -0.336. Std of Reward: 0.939. Training.
INFO:mlagents.trainers: hallway-curriculum-0: VisualHallwayLearning: Step: 4000. Mean Reward: -0.703. Std of Reward: 0.420. Training.
INFO:mlagents.trainers: hallway-curriculum-0: VisualHallwayLearning: Step: 5000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
```

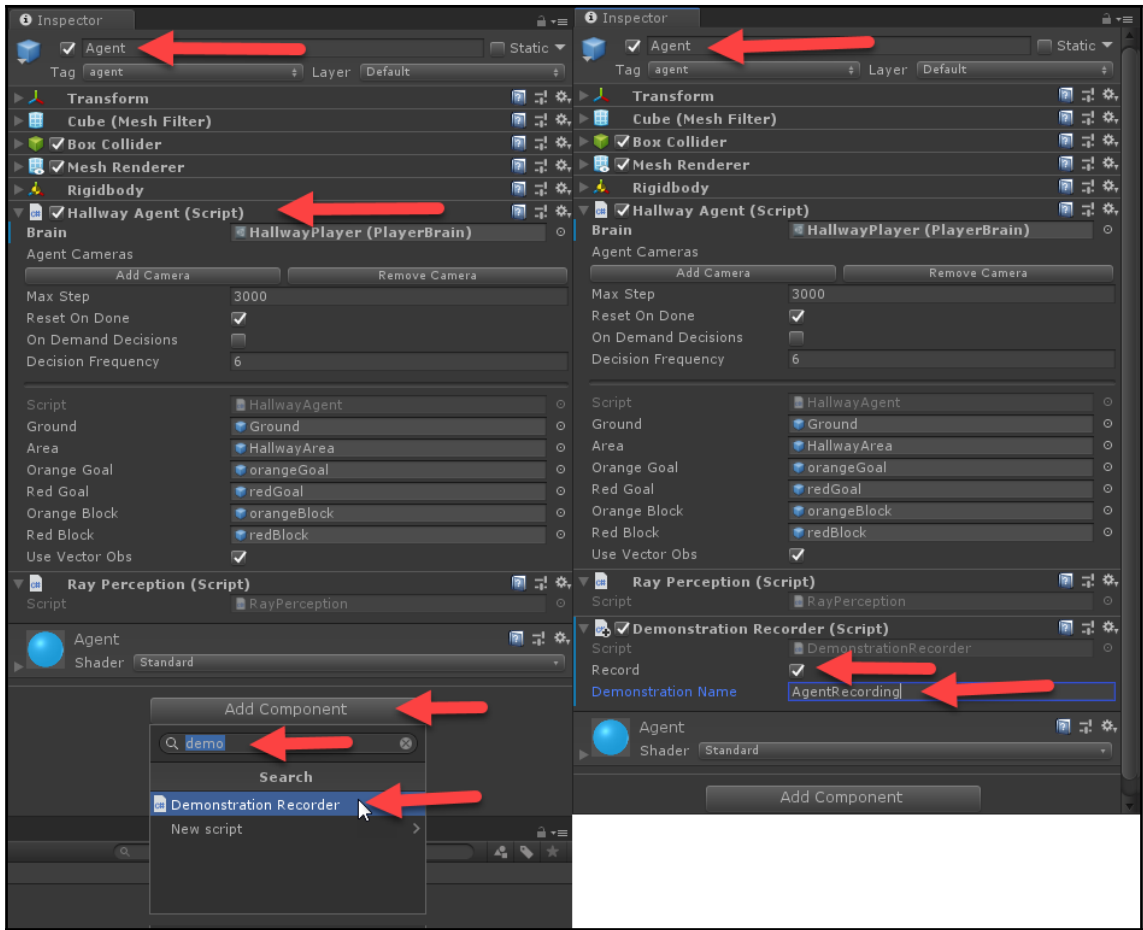


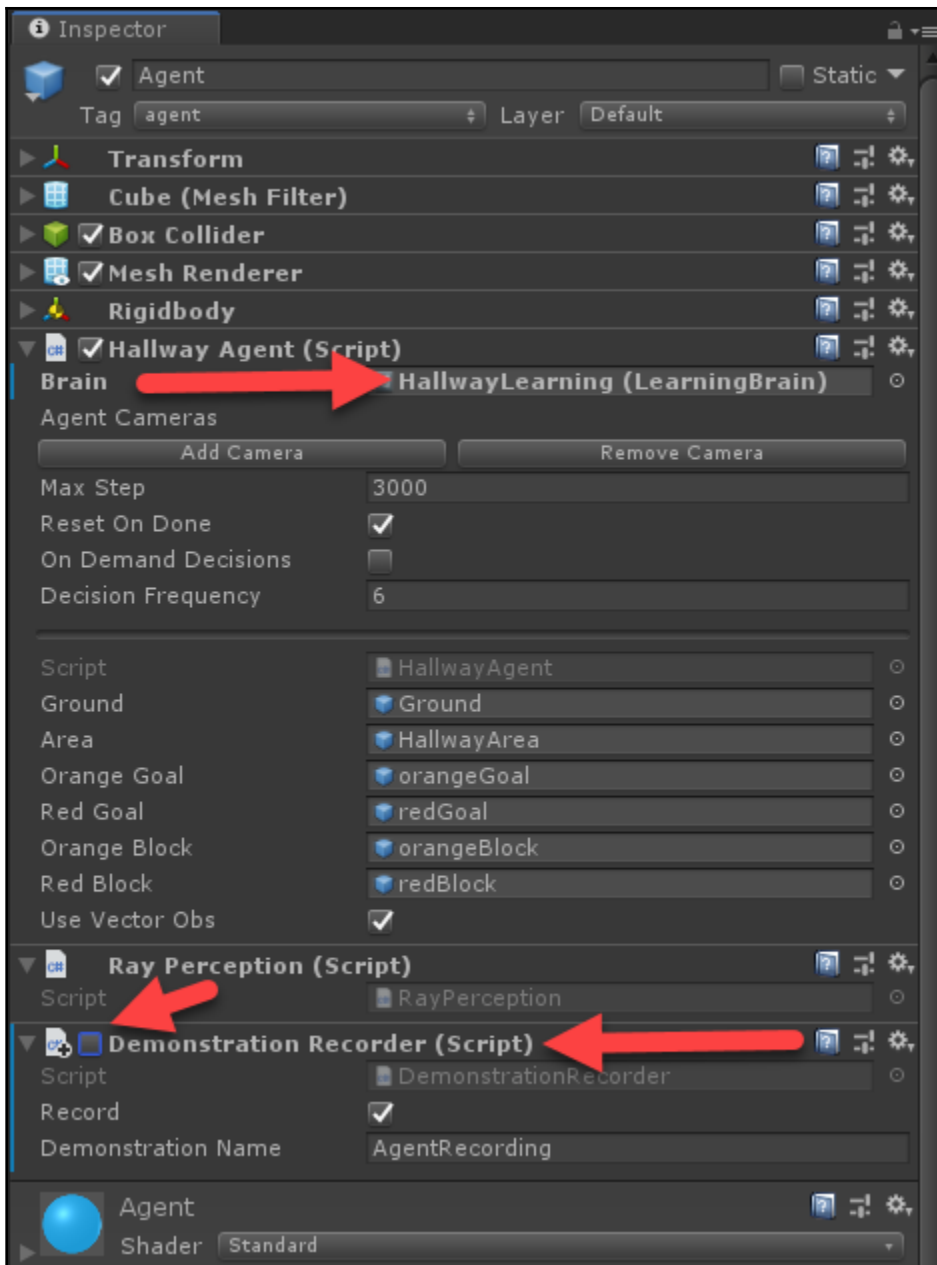
---

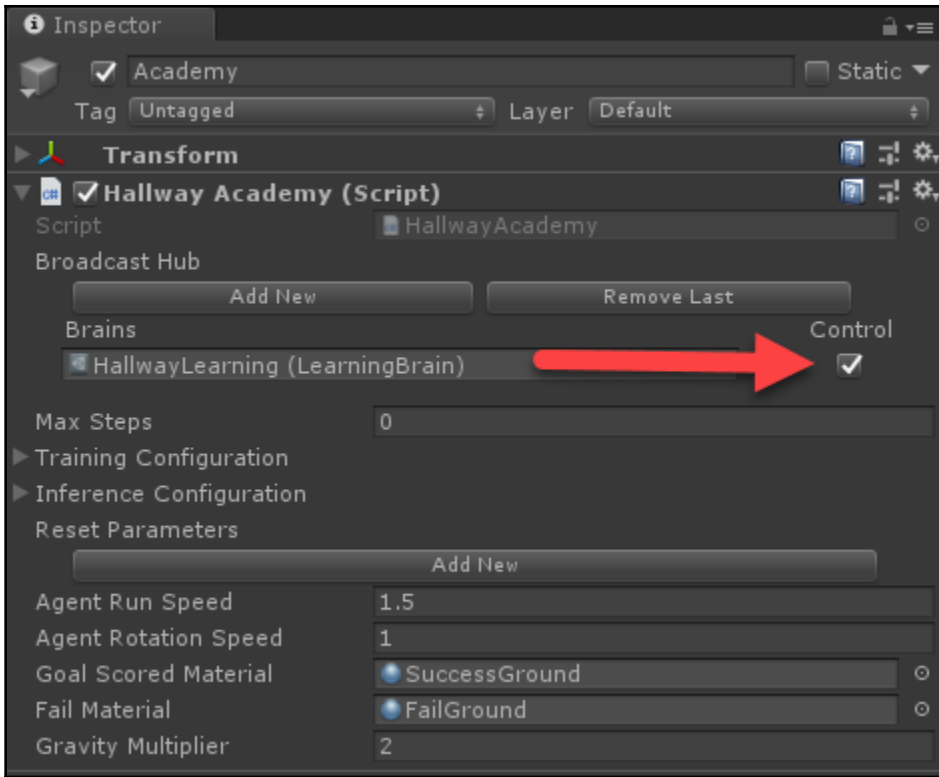
# Chapter 10: Imitation and Transfer Learning



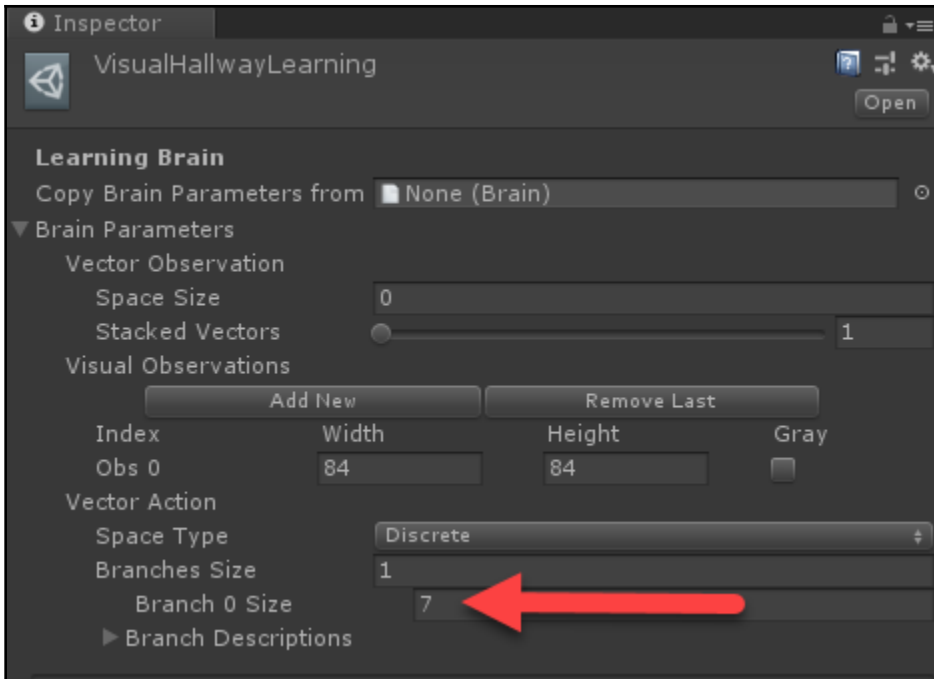












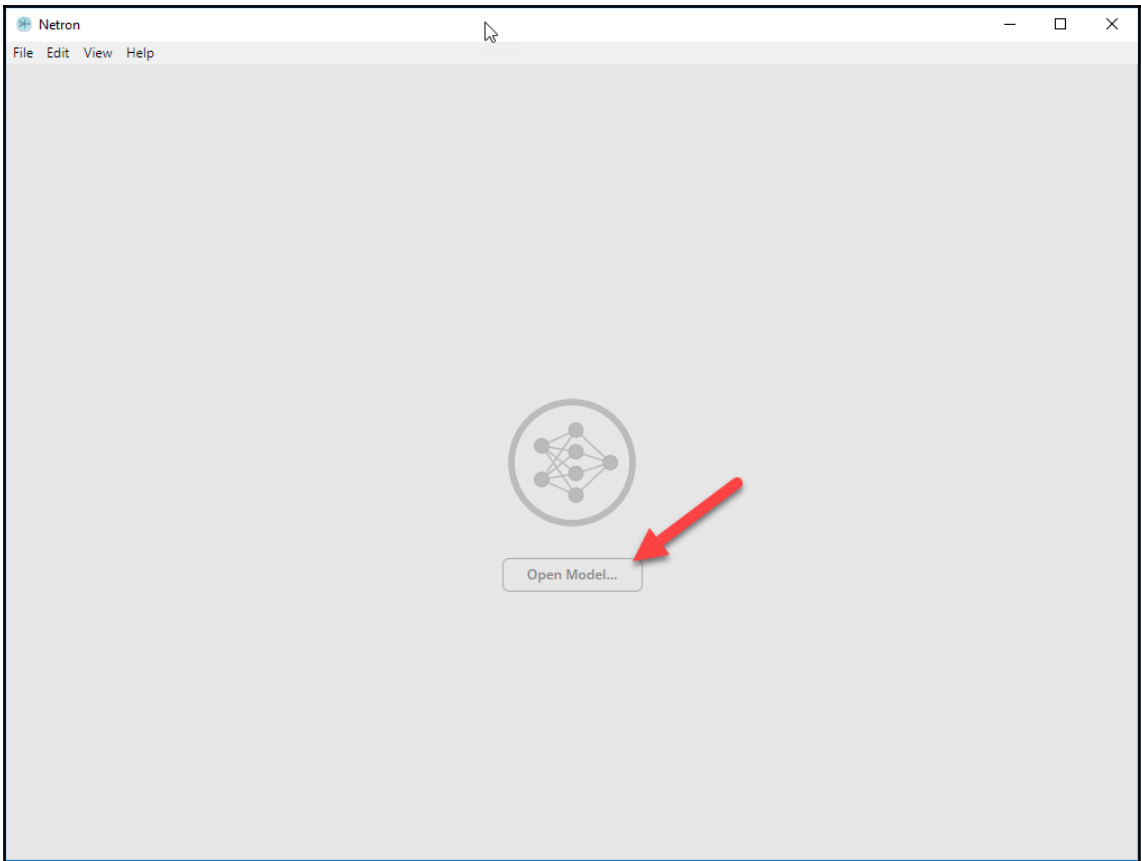
```

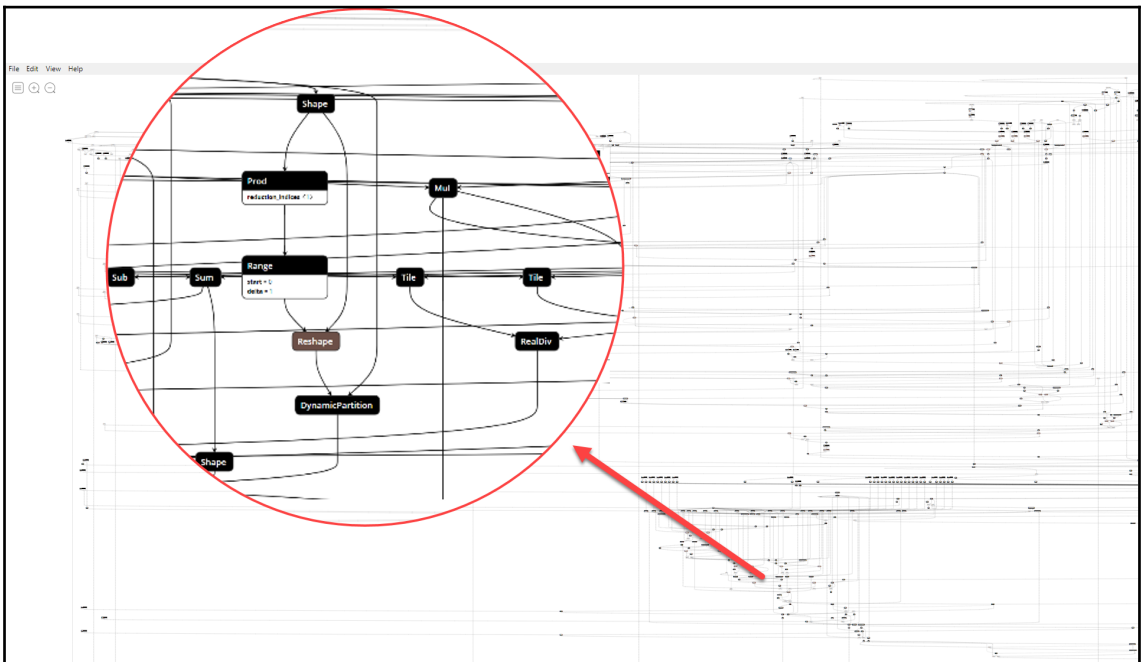
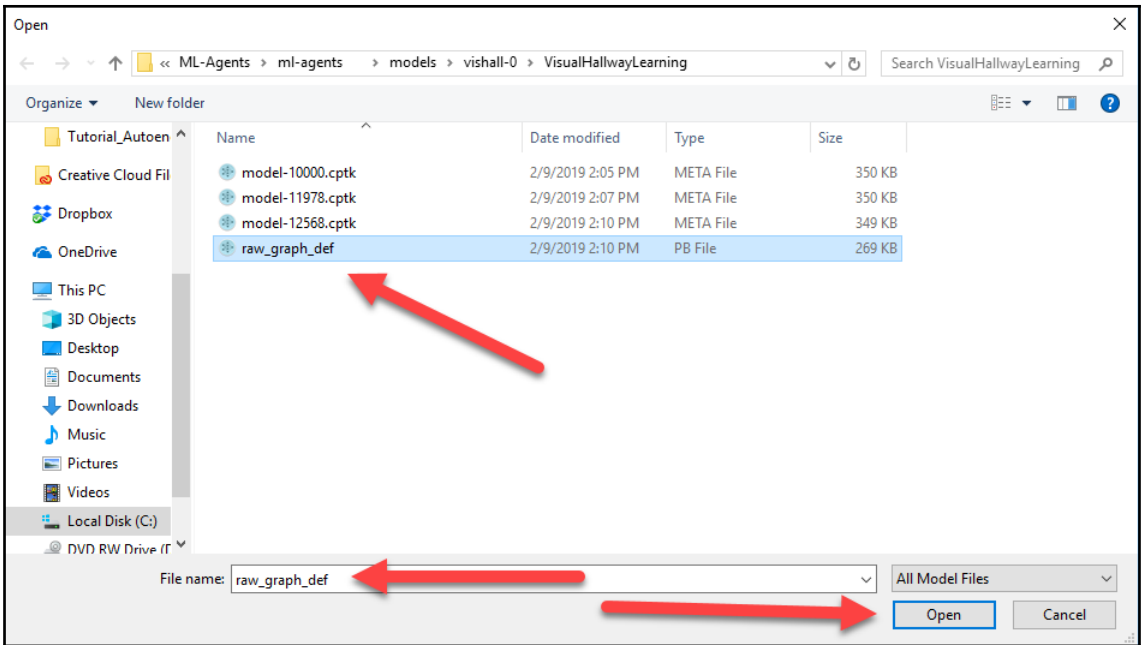
Anaconda Prompt - mlagents-learn config/trainer_config.yaml --run-id=visualhallwaylearning --train --save-freq=10000
is TensorFlow binary was not compiled to use: AVX2
INFO:mlagents.envs:Hyperparameters for the PPO Trainer of brain VisualHallwayLearning:
  batch_size: 64
  beta: 0.01
  buffer_size: 1024
  epsilon: 0.2
  gamma: 0.99
  hidden_units: 128
  lambda: 0.95
  learning_rate: 0.0003
  max_steps: 5.0e5
  normalize: False
  num_epoch: 3
  num_layers: 1
  time_horizon: 64
  sequence_length: 64
  summary_freq: 1000
  use_recurrent: True
  summary_path: ./summaries/visualhallwaylearning-0_VisualHallwayLearning
  memory_size: 256
  use_curiosity: False
  curiosity_strength: 0.01
  curiosity_enc_size: 128
  model_path: ./models/visualhallwaylearning-0/VisualHallwayLearning
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 1000. Mean Reward: -1.002. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 2000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 3000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 4000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 5000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 6000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 7000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 8000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 9000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.envs:Saved Model
INFO:mlagents.trainers: visualhallwaylearning-0: VisualHallwayLearning: Step: 10000. Mean Reward: -1.000. Std of Reward: 0.000. Training.

```

---

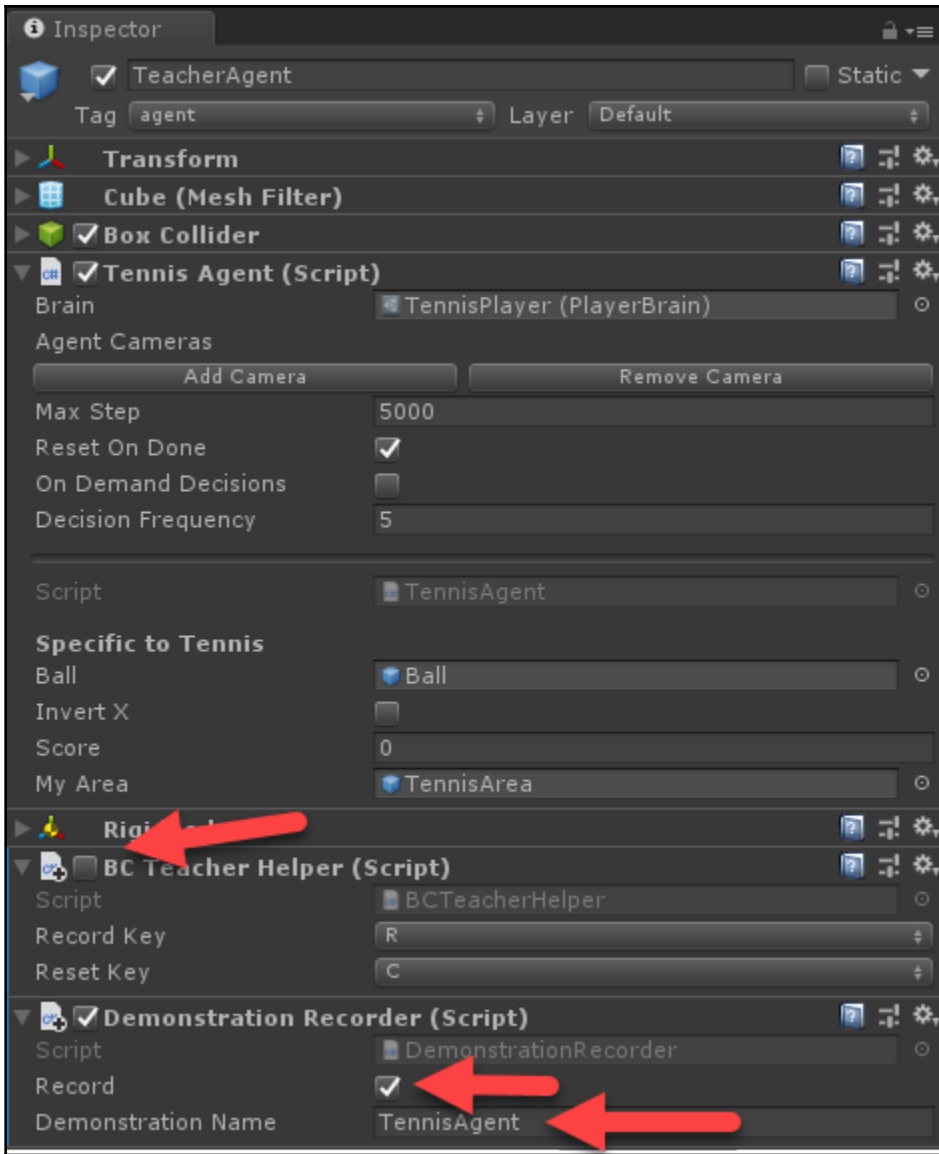
Name	Date modified	Type	Size
VisualHallwayLearning			
VisualHallwayLearning.bytes			
VisualPushBlockLearning	2/9/2019 2:10 PM	File folder	
VisualPushBlockLearning.bytes	2/9/2019 2:10 PM	BYTES File	1,897 KB






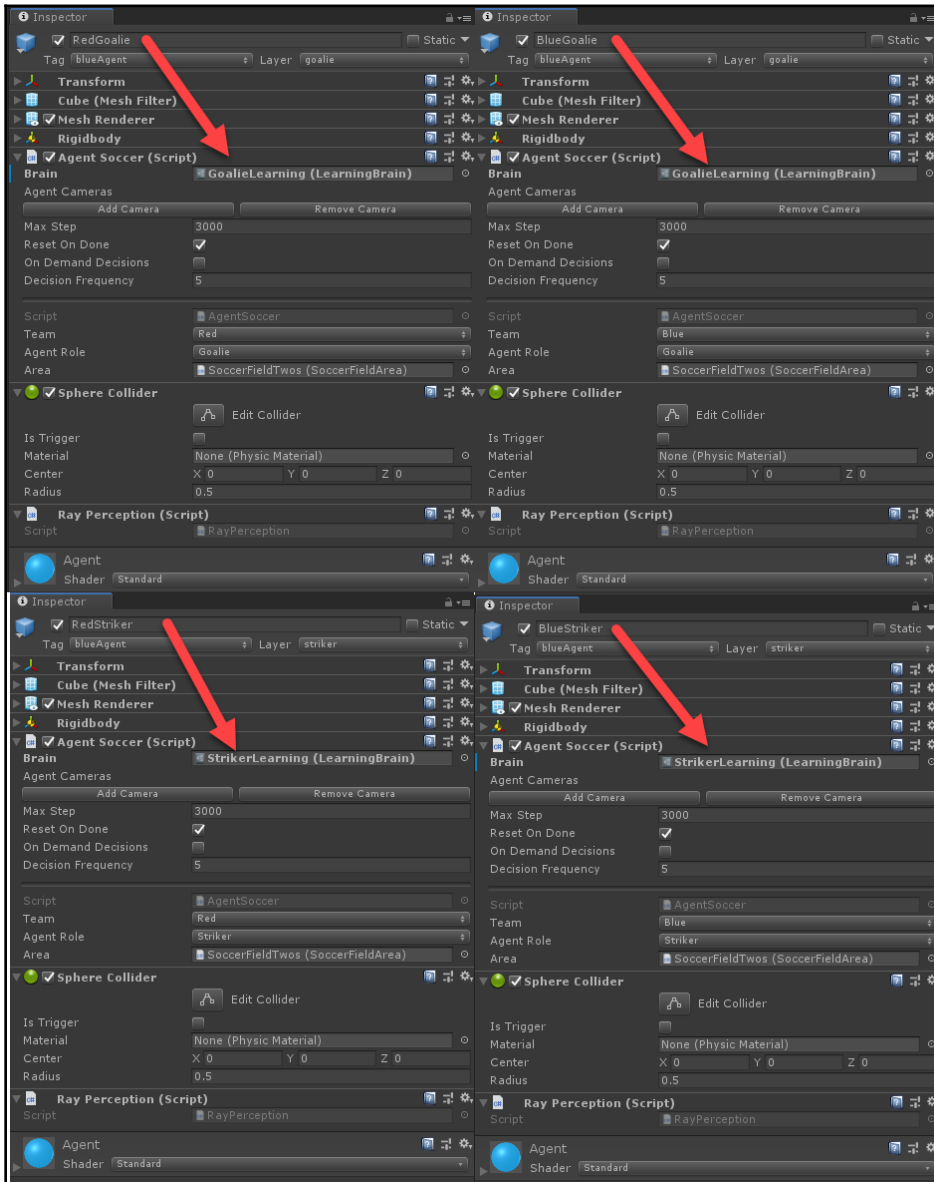
The image shows a neural network architecture diagram on the left and a 'MODEL PROPERTIES' panel on the right. The diagram includes nodes for 'advantages', 'SquaredDifference', 'DynamicPartition', and 'ZeroLike', with arrows indicating data flow. The 'MODEL PROPERTIES' panel is titled 'MODEL PROPERTIES' and contains the following information:

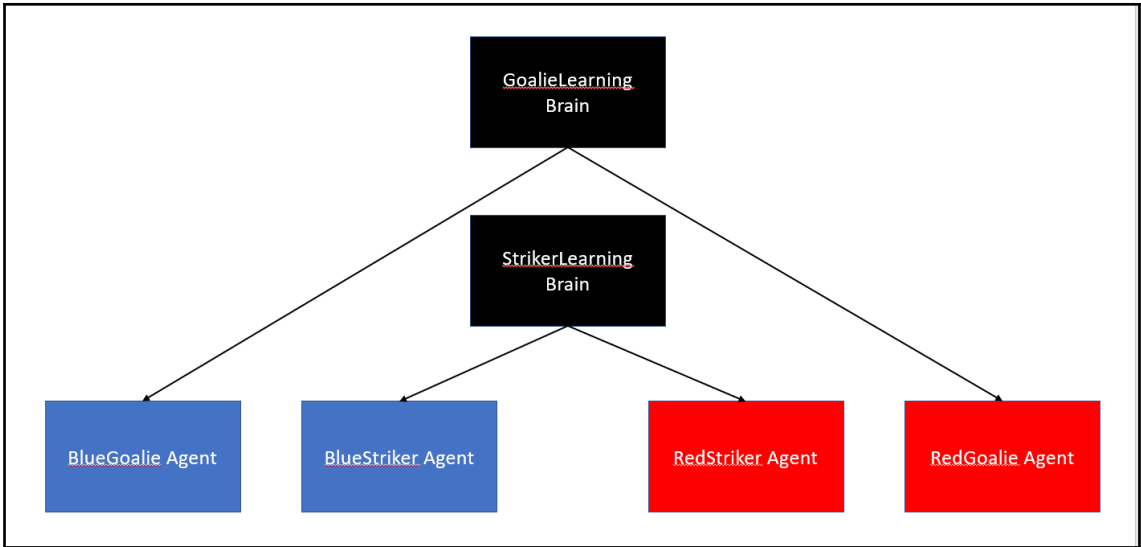
- format:** TensorFlow Graph
- Graph**
- version:** v26
- operators:** Total: 1643
- Inputs:**
  - sequence\_length: Id: sequence\_length, type: int32[-]
  - masks: Id: masks, type: float32[?]
  - new\_reward: Id: new\_reward, type: float32
  - visual\_observati...: Id: visual\_observation\_0, type: float32[?, 84, 84, 3]
  - prev\_action: Id: prev\_action, type: int32[?, 1]
  - recurrent\_in: Id: recurrent\_in, type: float32[?, 256]
  - action\_masks: Id: action\_masks, type: float32[?, 7]
  - action\_holder: Id: action\_holder, type: int32[?, 1]
  - old\_probabilities: Id: old\_probabilities, type: float32[?, 7]
  - discounted\_rewards: Id: discounted\_rewards, type: float32[?]
  - advantages: Id: advantages, type: float32[?, 1]
  - old\_value\_estimates: Id: old\_value\_estimates, type: float32[?]







# Chapter 11: Building Multi-Agent Environments






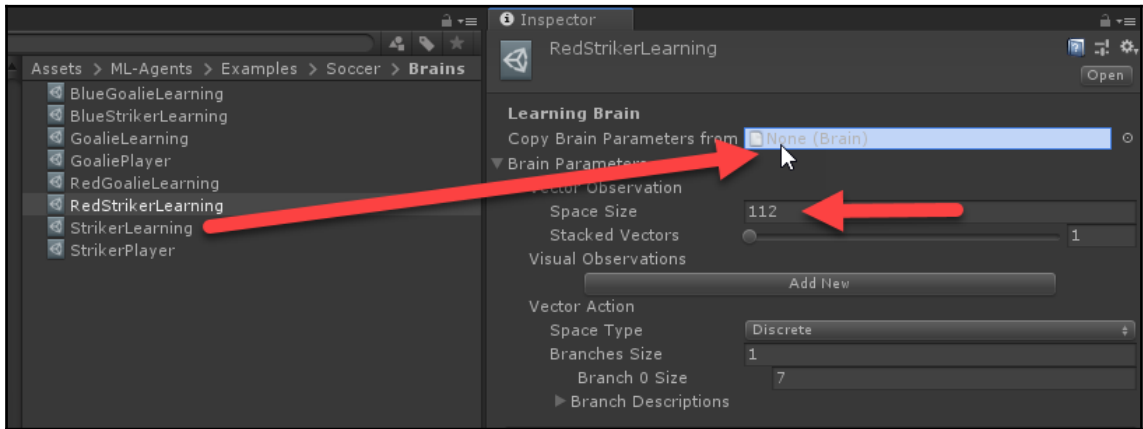
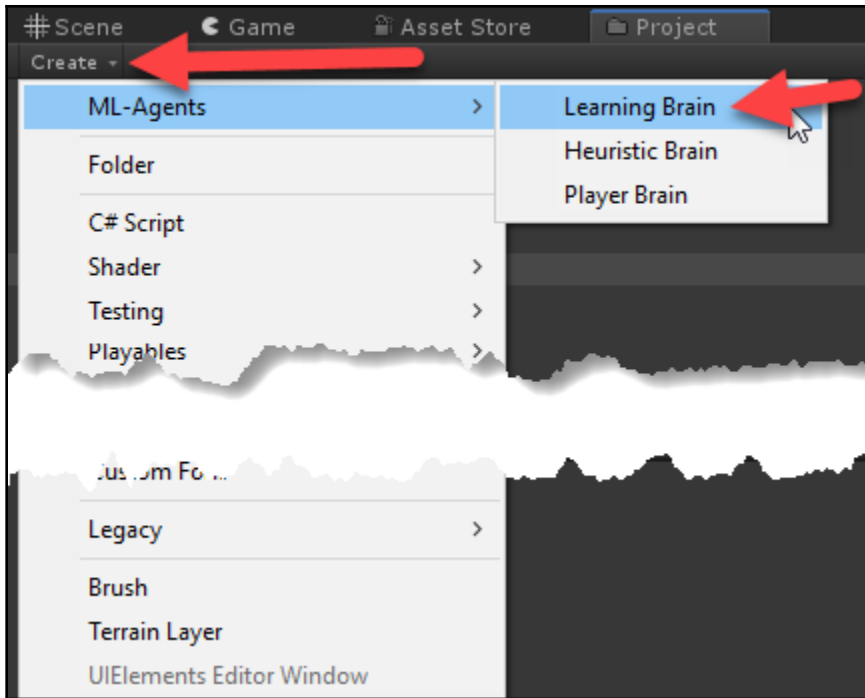


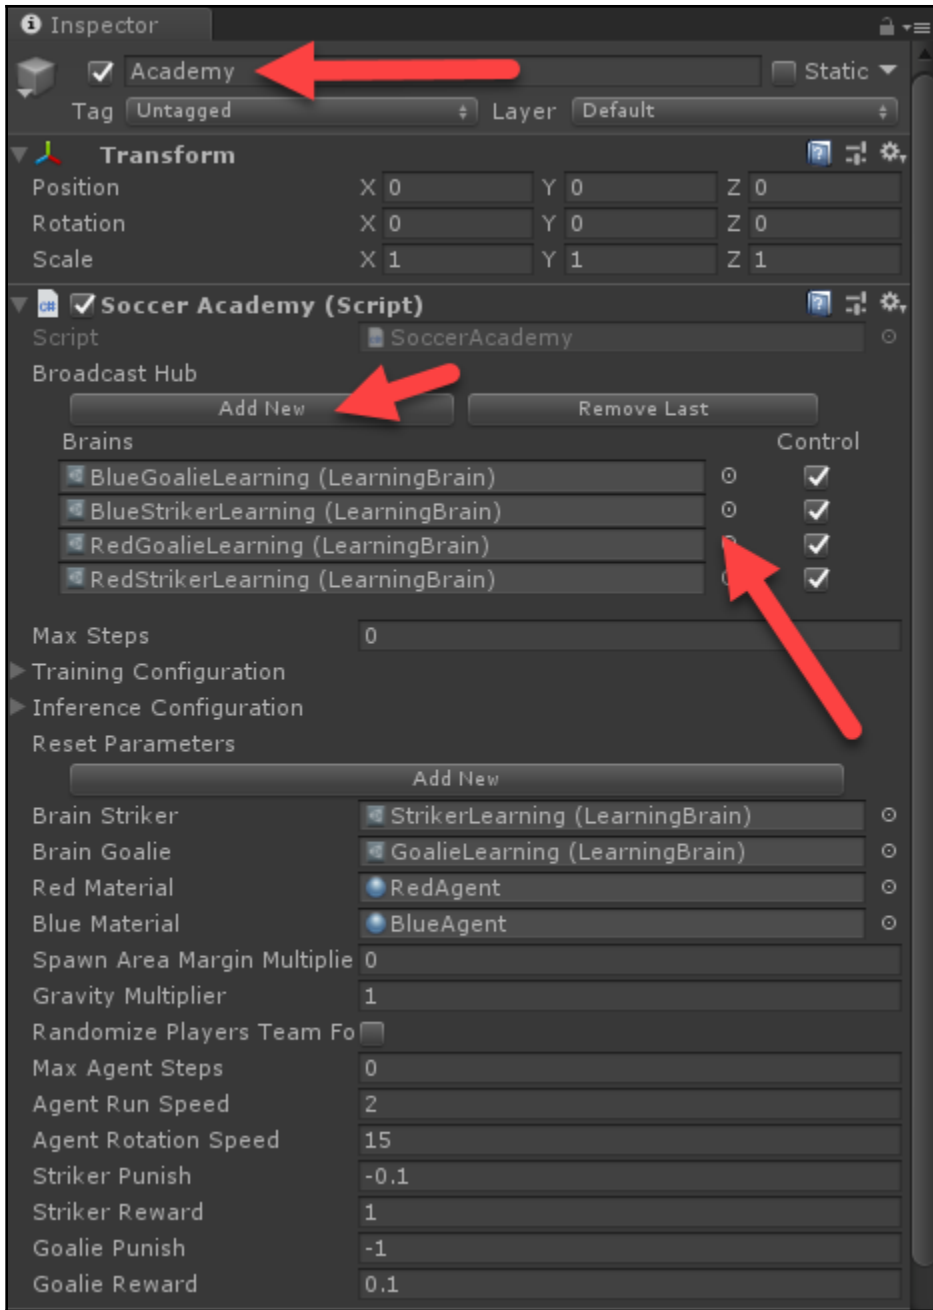




```
Anaconda Prompt - mlagents-learn config/trainer_config.yaml --run-id=soccer--train
epsilon: 0.2
gamma: 0.99
hidden_units: 256
lambda: 0.95
learning_rate: 1e-3
max_steps: 5.0e5
normalize: False
num_epoch: 3
num_layers: 2
time_horizon: 128
sequence_length: 64
summary_freq: 2000
use_recurrent: False
summary_path: ./summaries/soccer-0_GoalieLearning
memory_size: 256
use_curiosity: False
curiosity_strength: 0.01
curiosity_enc_size: 128
model_path: ./models/soccer-0/GoalieLearning
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 2000. Mean Reward: -0.240. Std of Reward: 0.708. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 2000. Mean Reward: 0.240. Std of Reward: 0.708. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 4000. Mean Reward: -0.755. Std of Reward: 0.505. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 4000. Mean Reward: 0.755. Std of Reward: 0.505. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 6000. Mean Reward: -0.482. Std of Reward: 0.724. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 6000. Mean Reward: 0.482. Std of Reward: 0.724. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 8000. Mean Reward: -0.822. Std of Reward: 0.406. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 8000. Mean Reward: 0.822. Std of Reward: 0.406. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 10000. Mean Reward: -1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 10000. Mean Reward: 1.000. Std of Reward: 0.000. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 12000. Mean Reward: -0.545. Std of Reward: 0.718. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 12000. Mean Reward: 0.545. Std of Reward: 0.718. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 14000. Mean Reward: -0.198. Std of Reward: 0.749. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 14000. Mean Reward: 0.198. Std of Reward: 0.749. Training.
INFO:mlagents.trainers: soccer-0: StrikerLearning: Step: 16000. Mean Reward: -0.788. Std of Reward: 0.437. Training.
INFO:mlagents.trainers: soccer-0: GoalieLearning: Step: 16000. Mean Reward: 0.788. Std of Reward: 0.437. Training.
```








All Assets ▾ Type here to search assets 🔍 ⋮

! You downloaded this item on Feb 10, 2019. Write a Review



POLYGON BLACKSMITH

# Toony Tiny People Demo

FREE

★★★★☆ 3 user reviews

Popular Tags

Add a new tag right now?

Add tags

Demo version of the Toony Tiny characters series


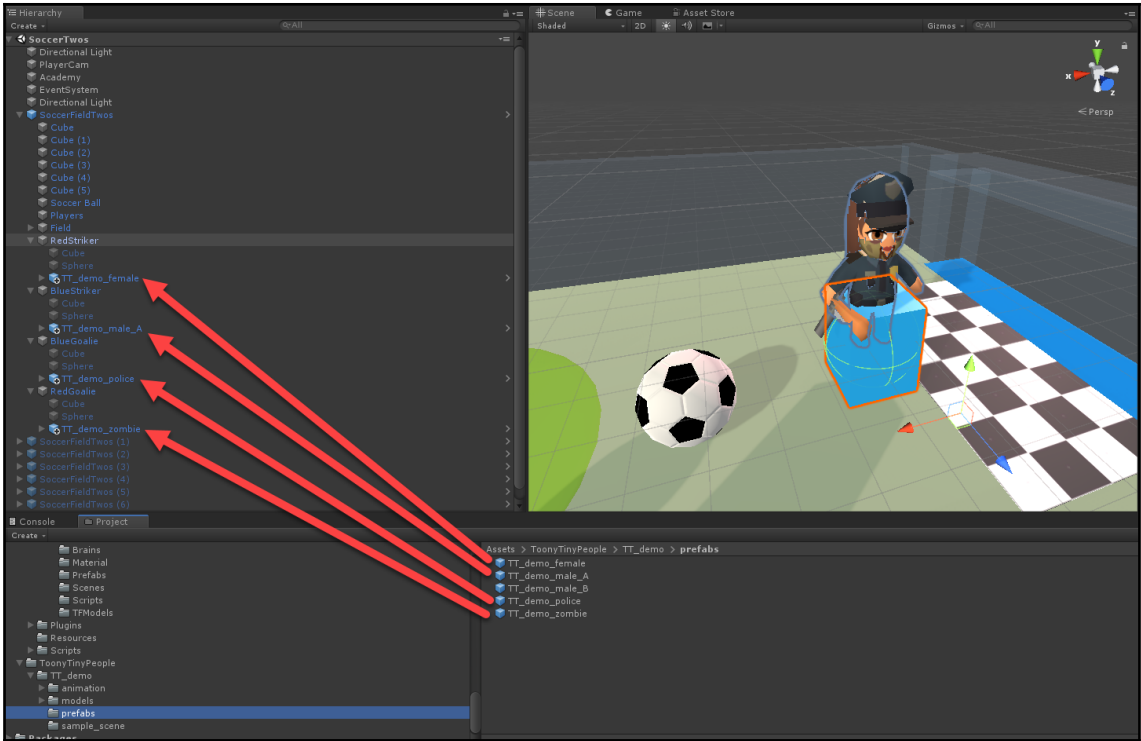
Pack includes:

- Male citizen character x2
- Female citizen character
- Policeman character
- Zombie character

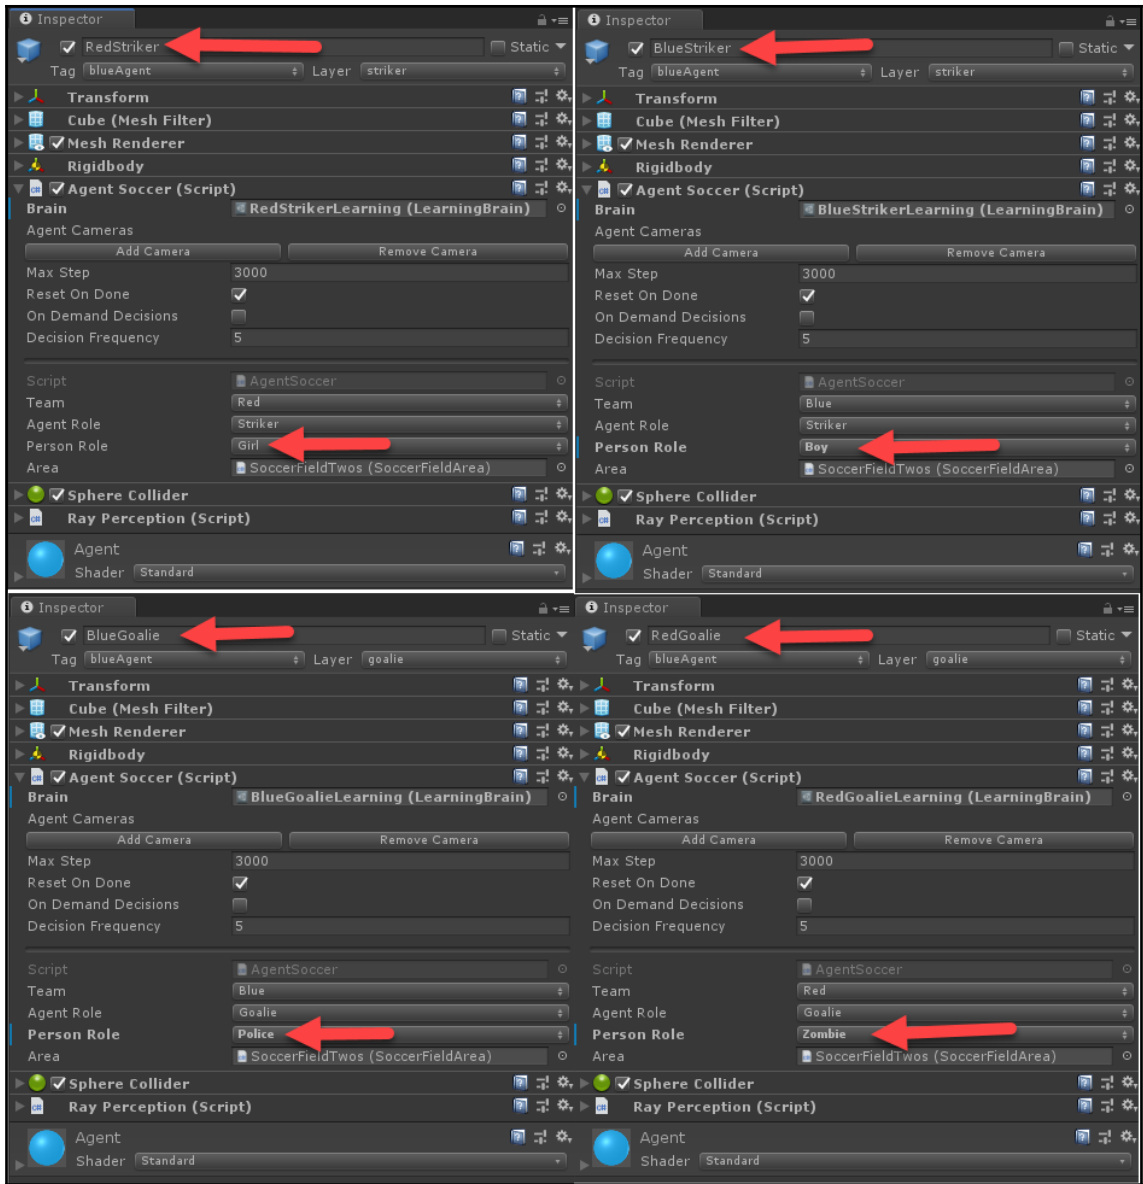
30 animation files.

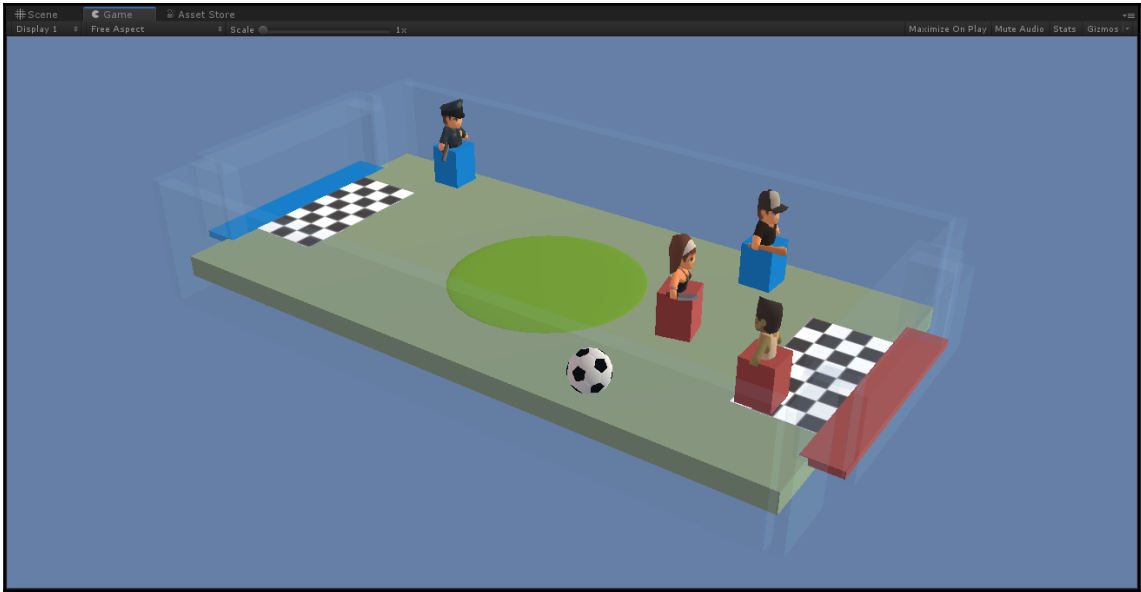
Feedback

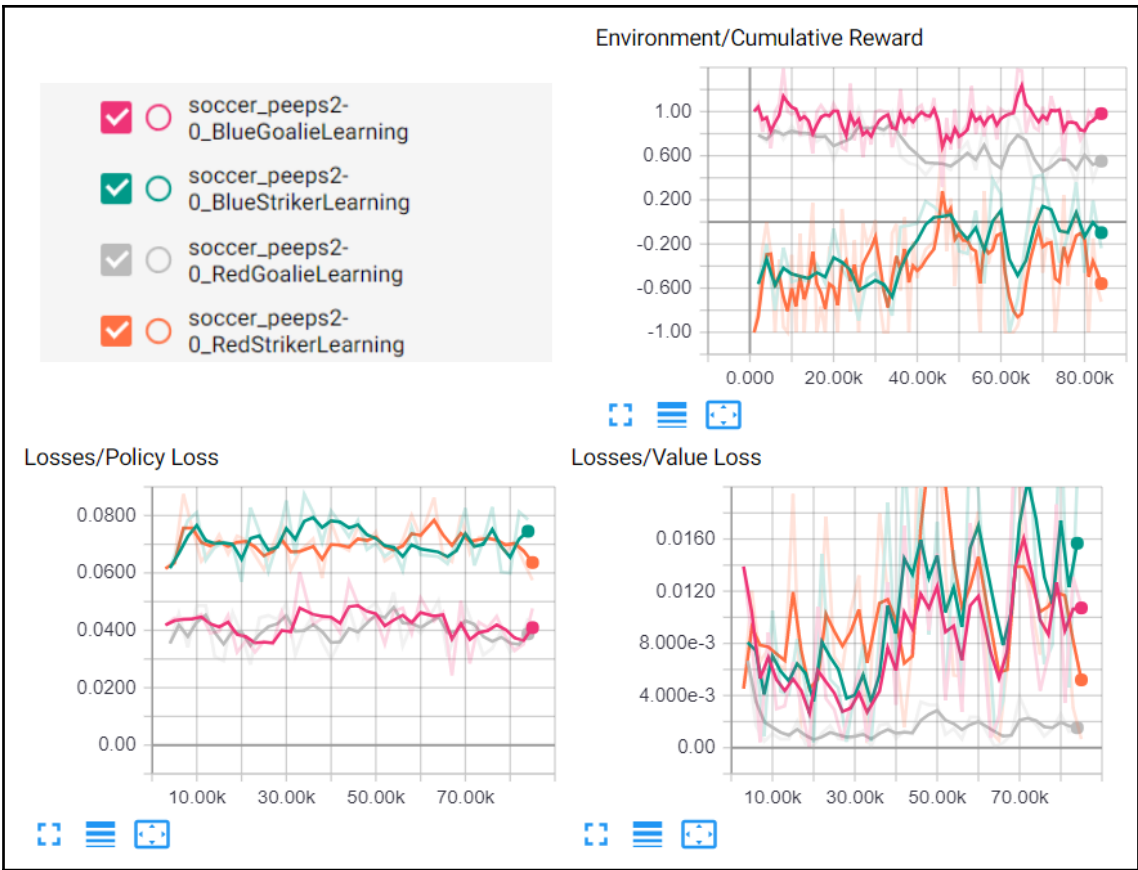
▾ Package contents 1.9 MB




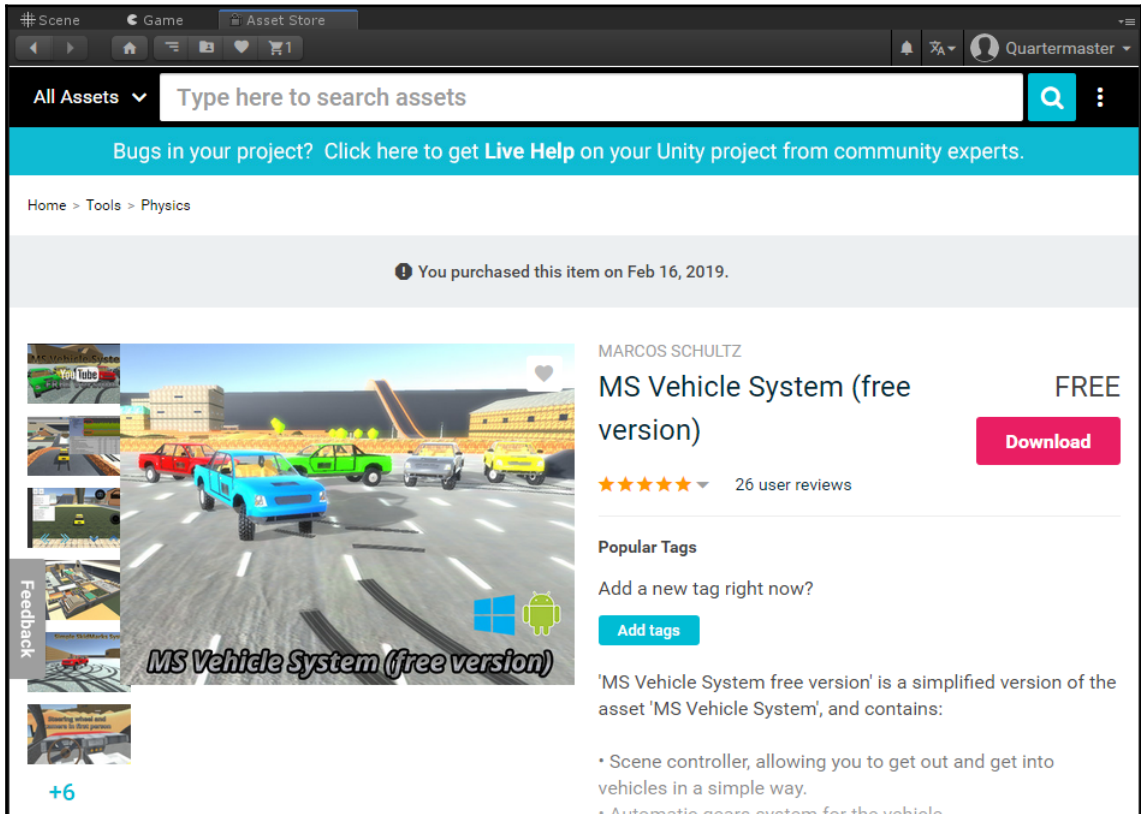





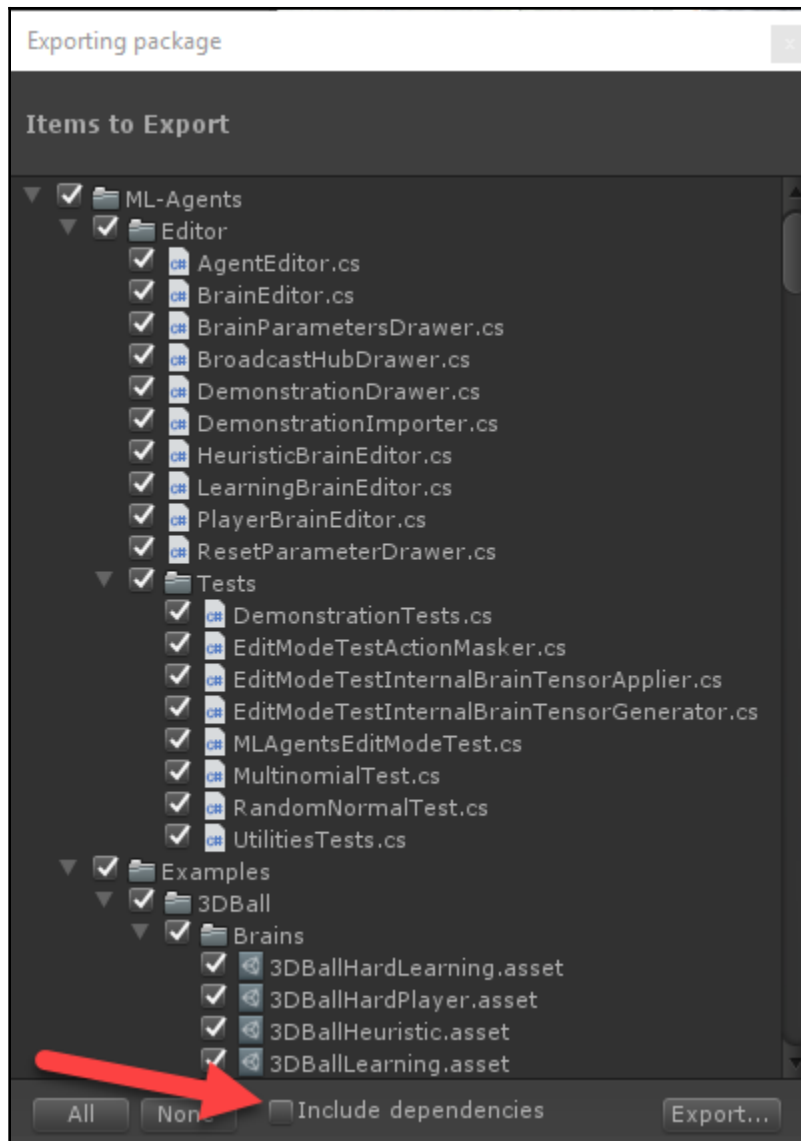


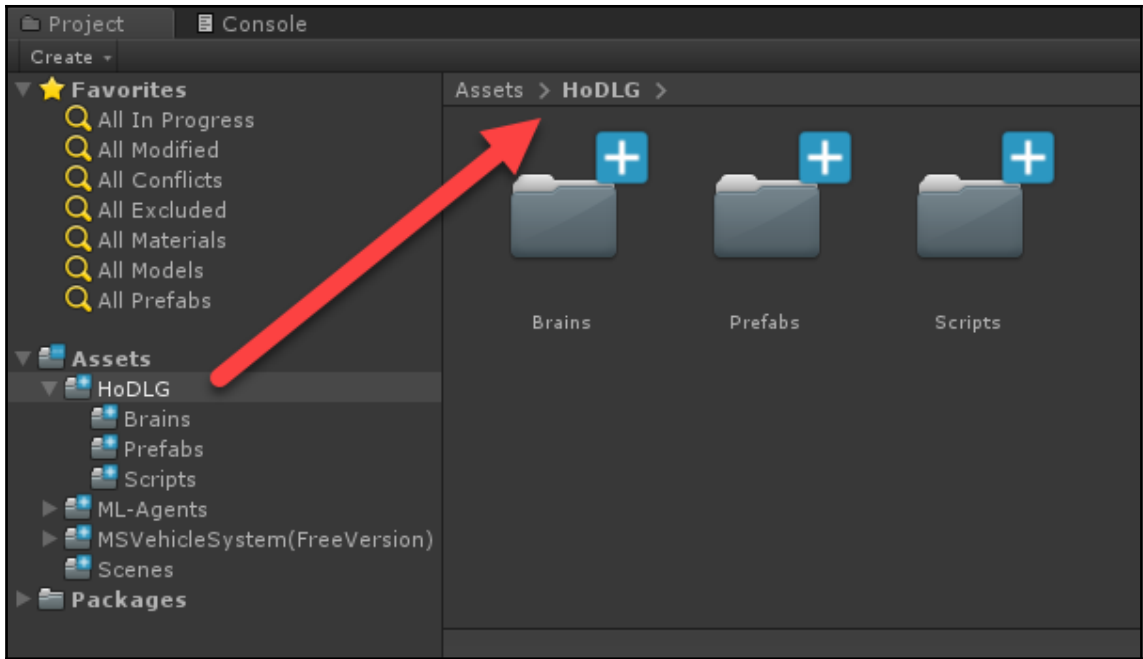


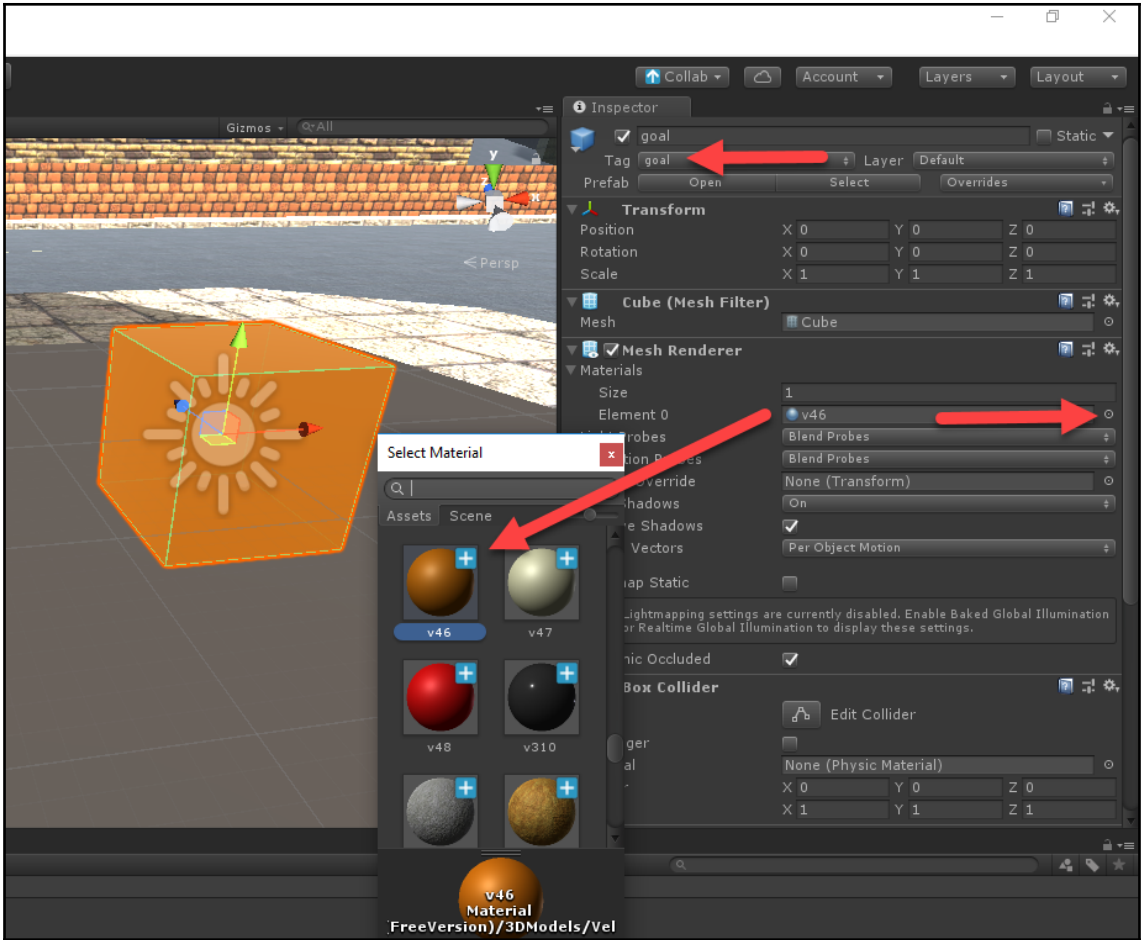
# Chapter 12: Debugging/Testing a Game with DRL



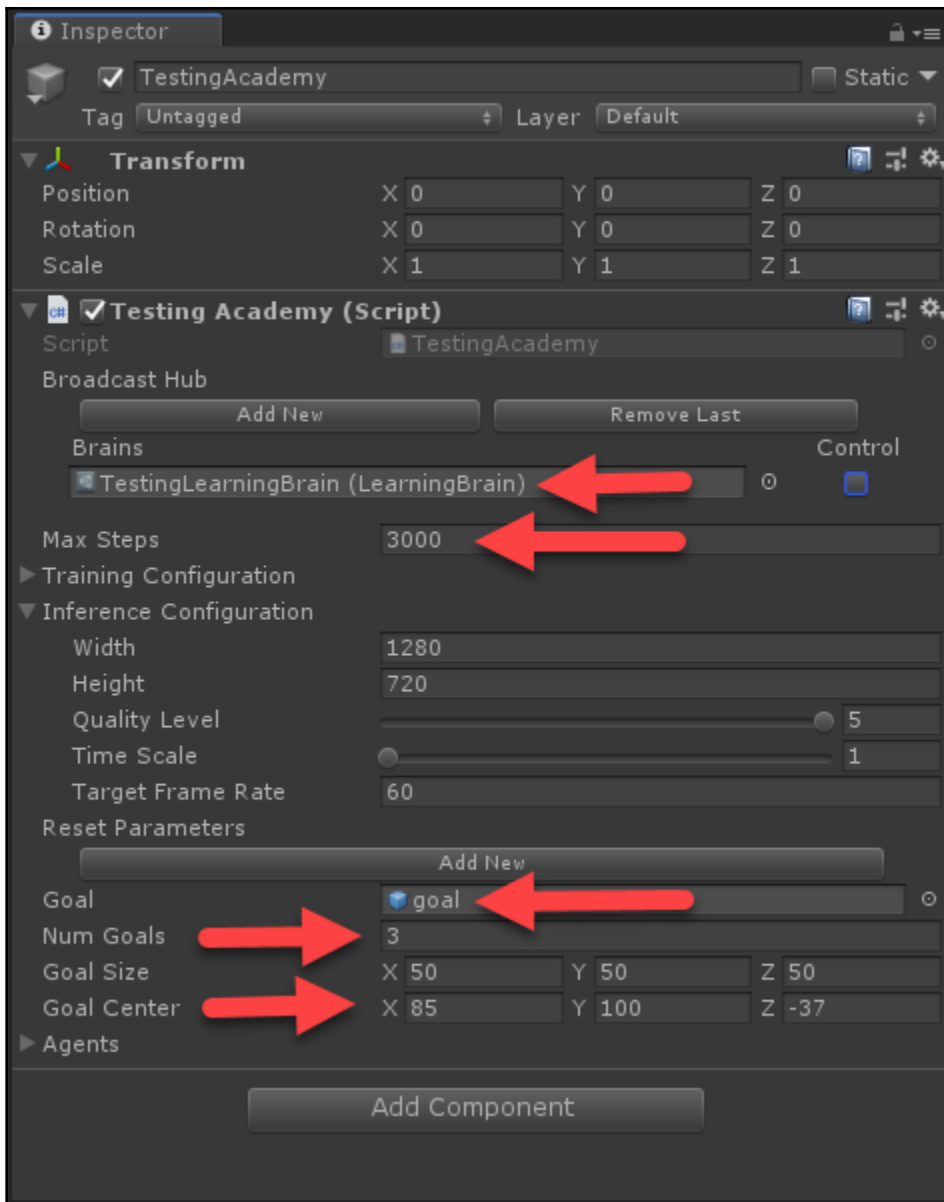





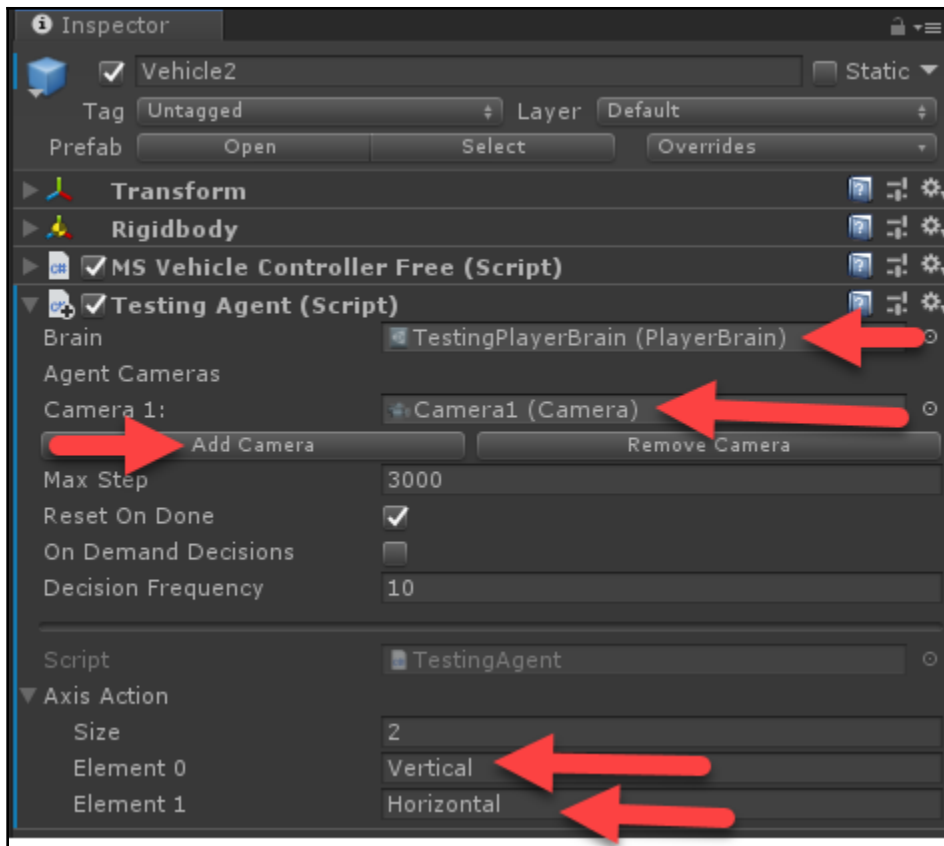


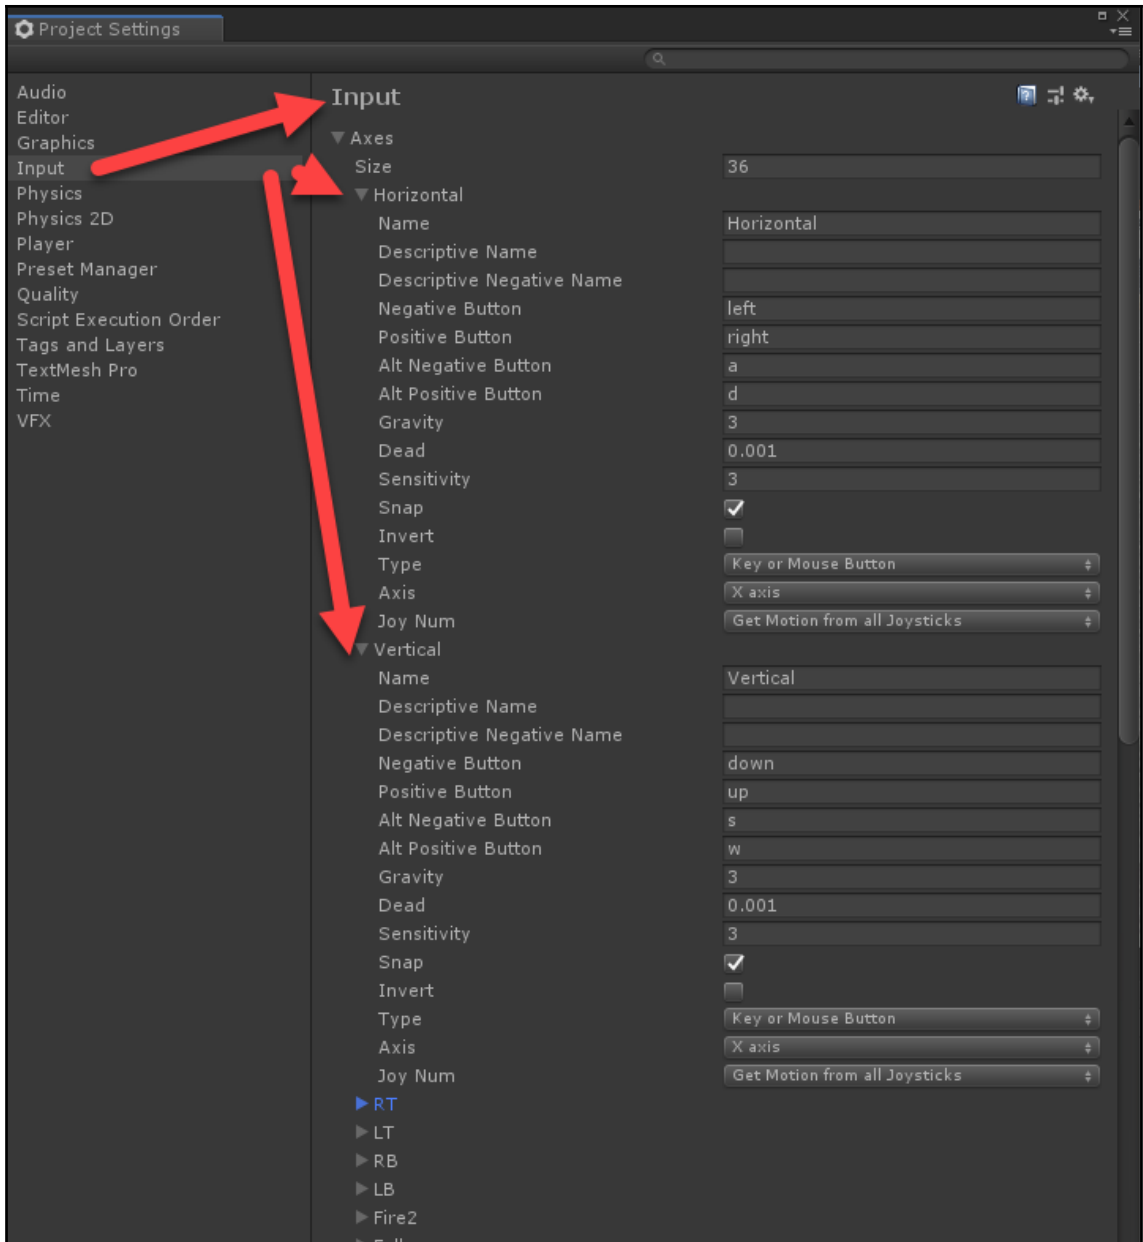





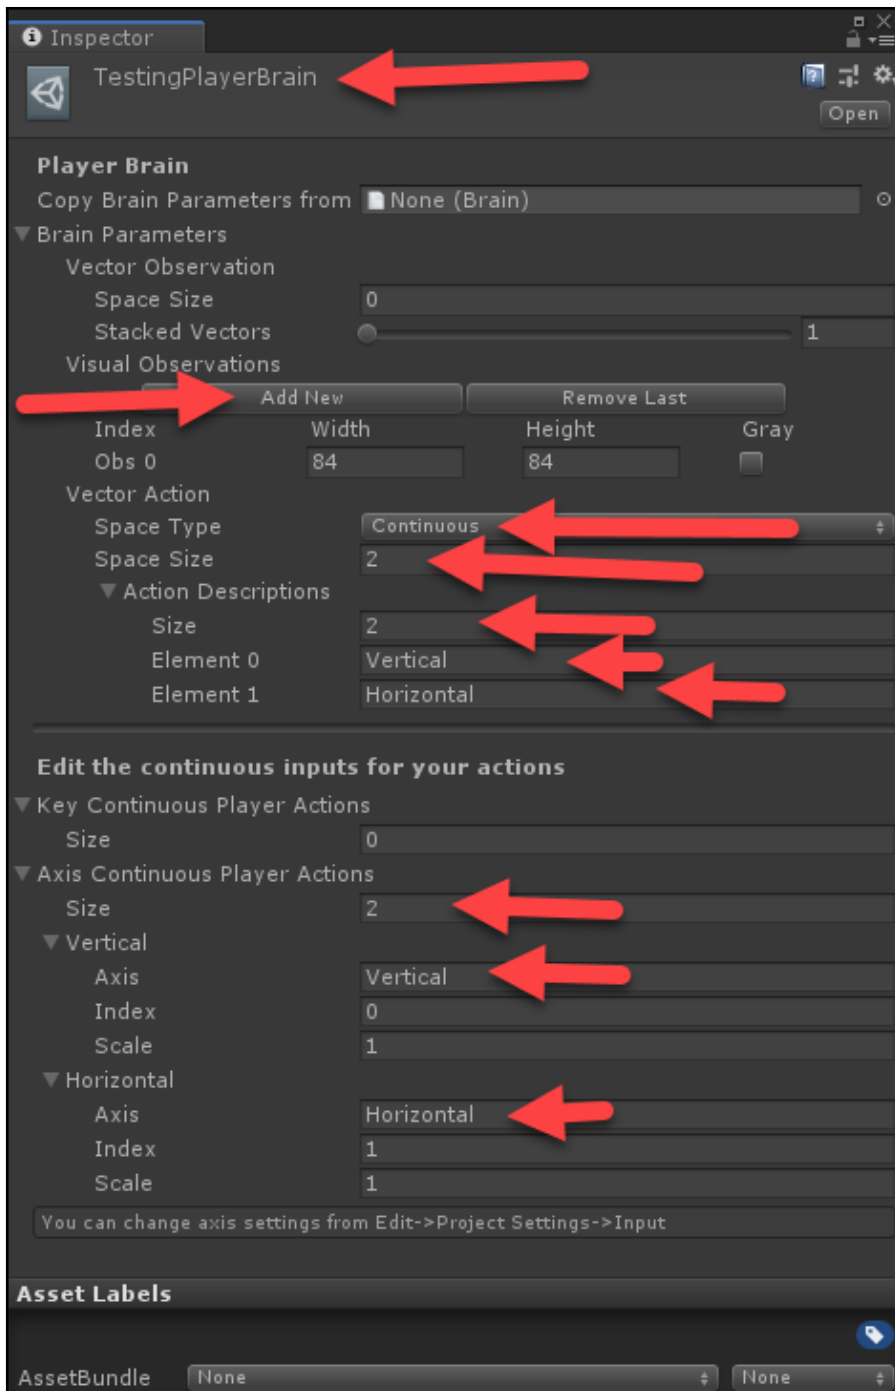


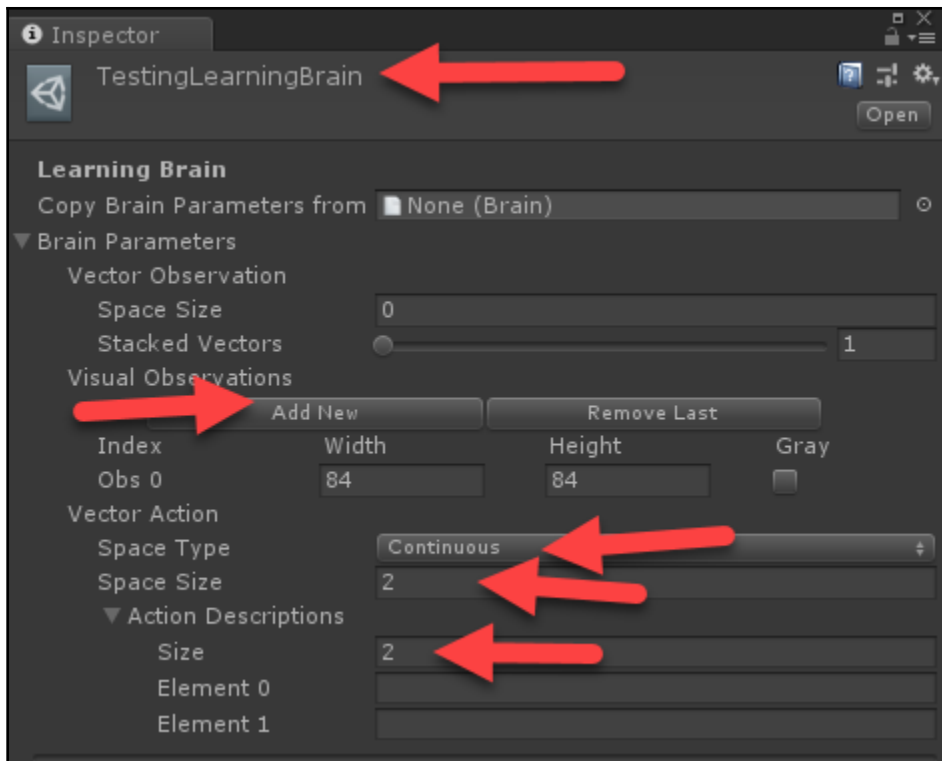






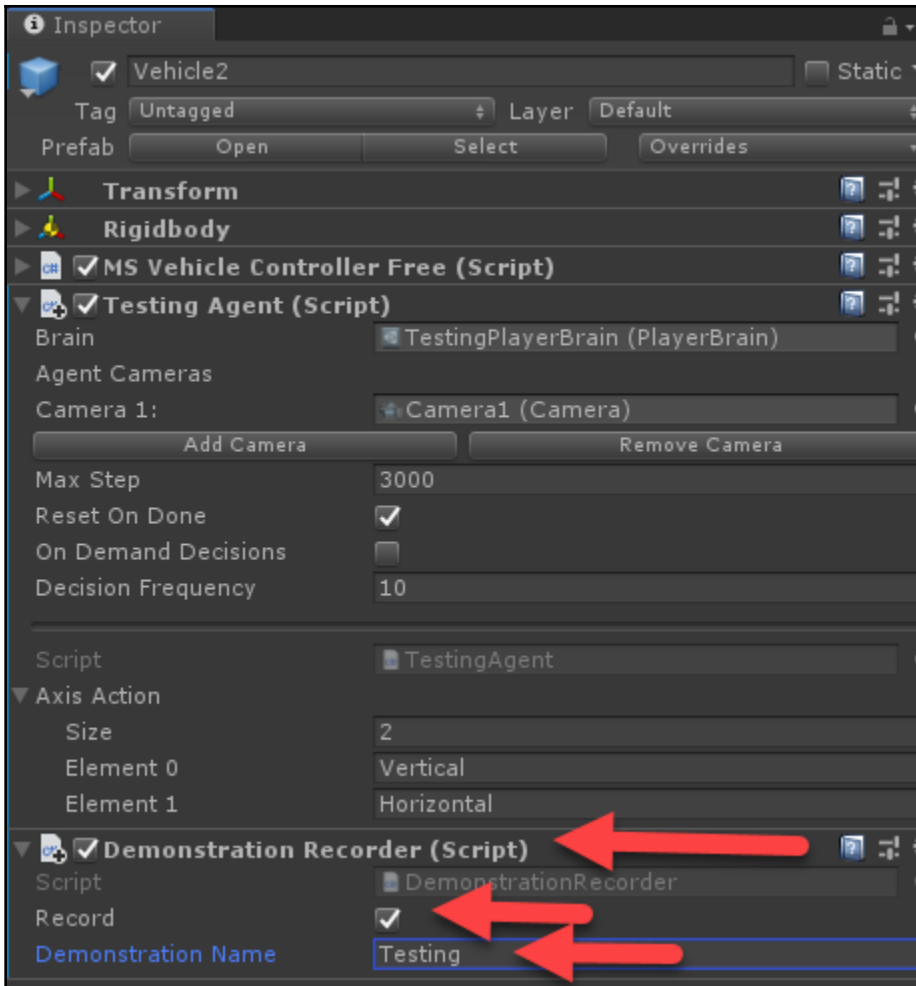


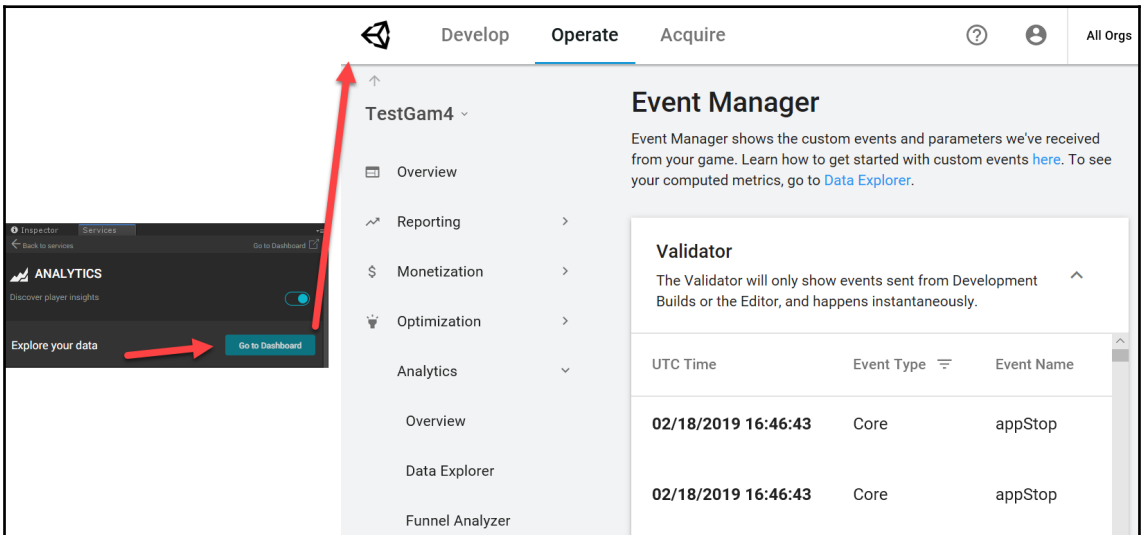
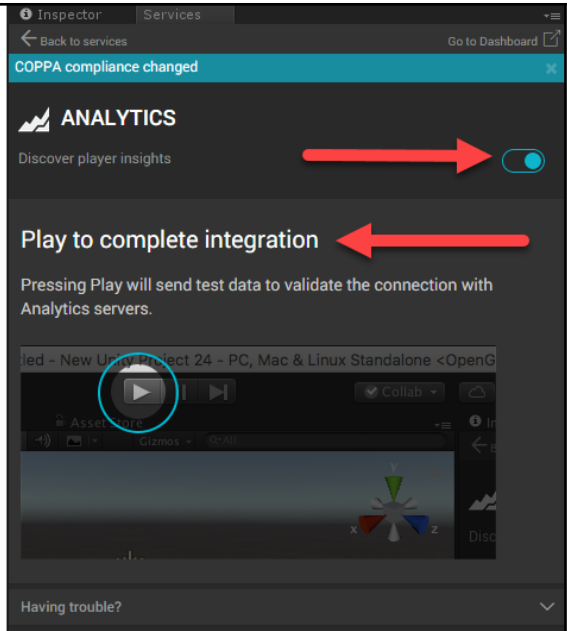
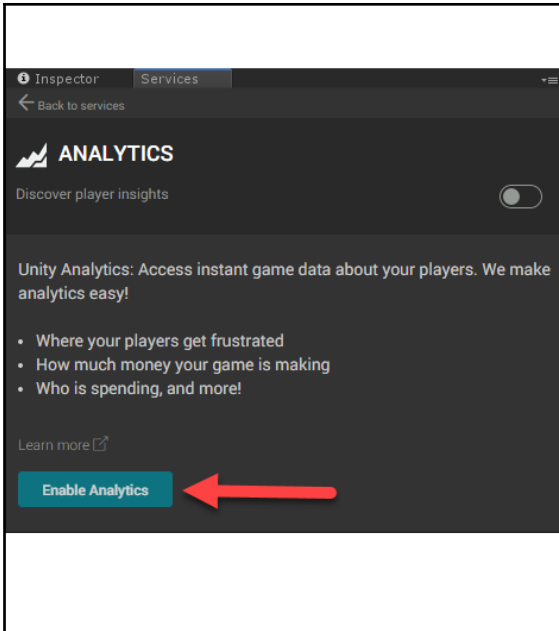


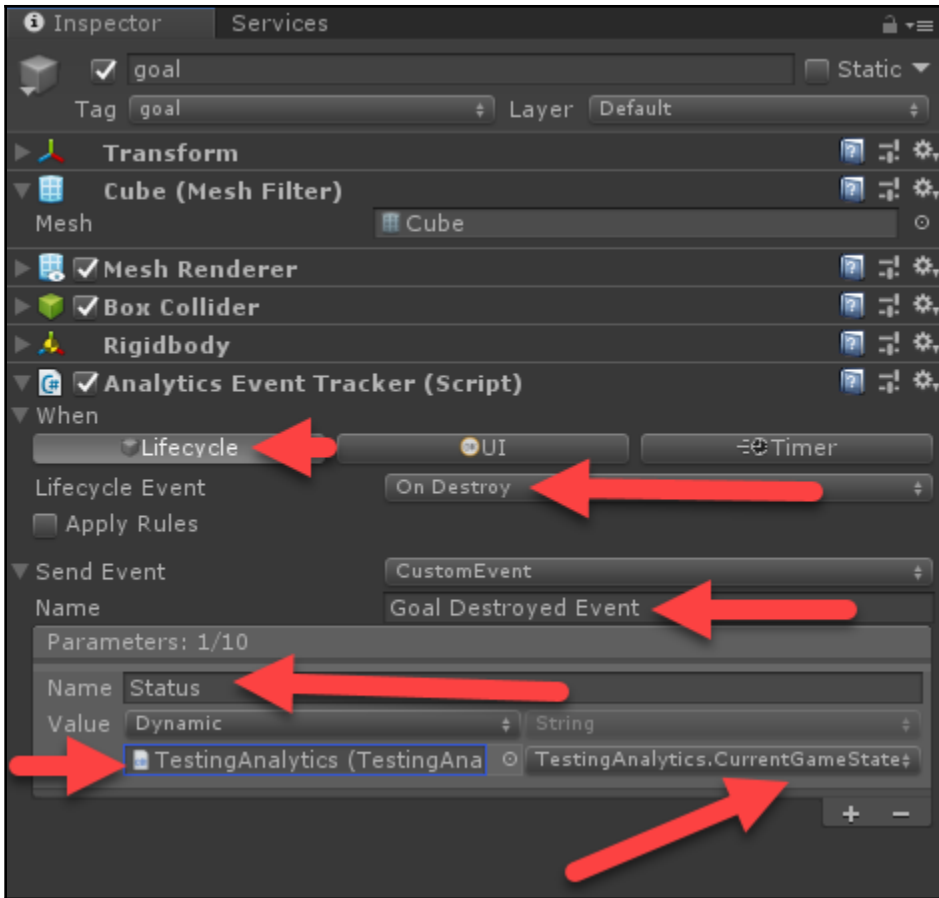












---

# Chapter 13: Obstacle Tower Challenge and Beyond

