Chapter 1:
Introducing Machine Learning and ML-Agents

shot	velocity	distance
1	5	2.28
2	10	9.12
3	20	36.49
4	25	57.02
5	35	111.75
6	40	145.96
7	50	228.06
8	60	328.41
9	65	385.42
10	75	513.14

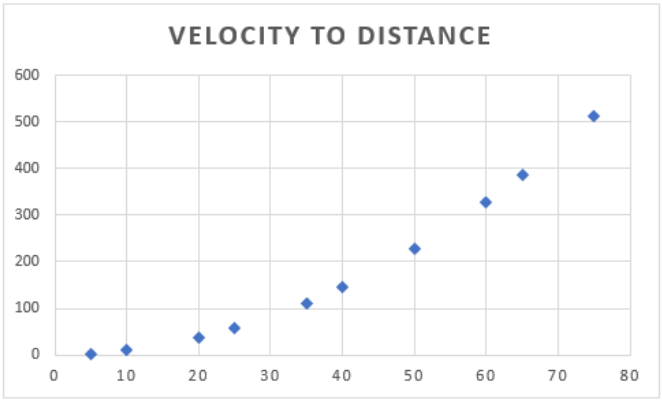

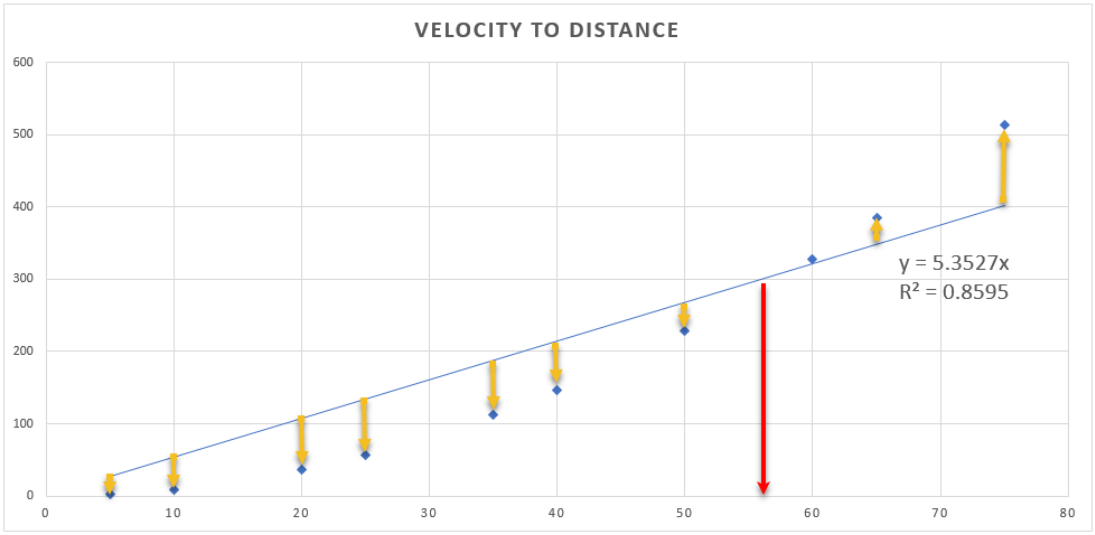

$$\boxed{d = 5.3527v}$$

$$\boxed{v = d/5.3527}$$

$$\boxed{v = 300/5.3527 = 56.05}$$

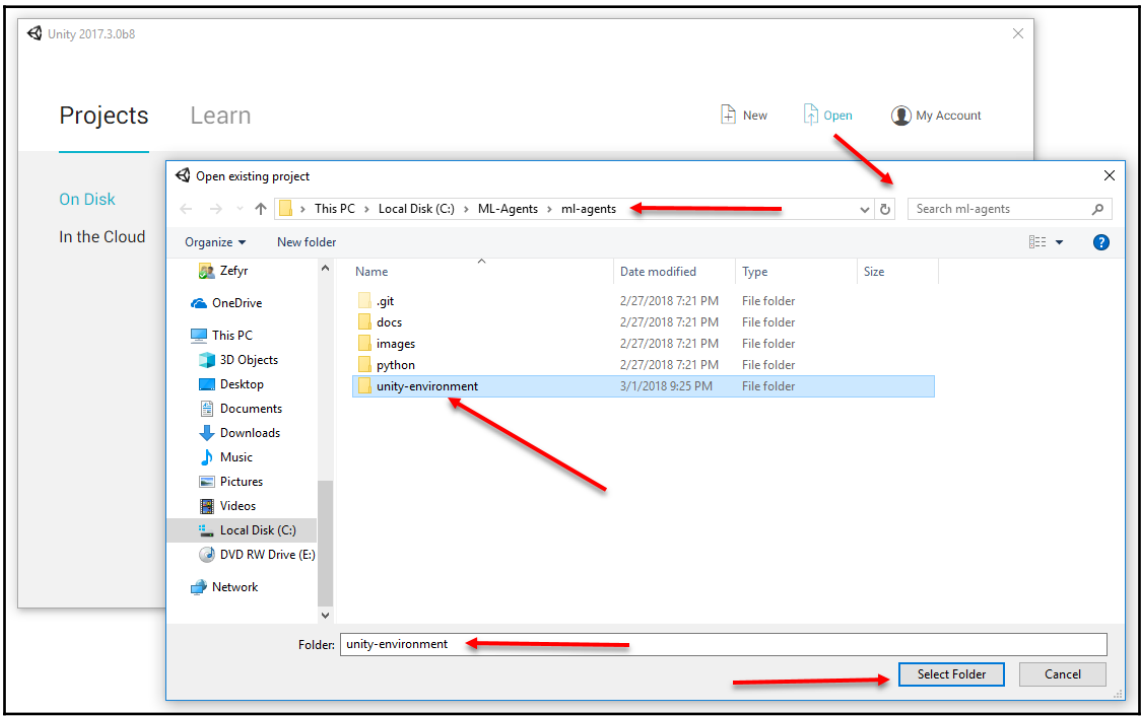

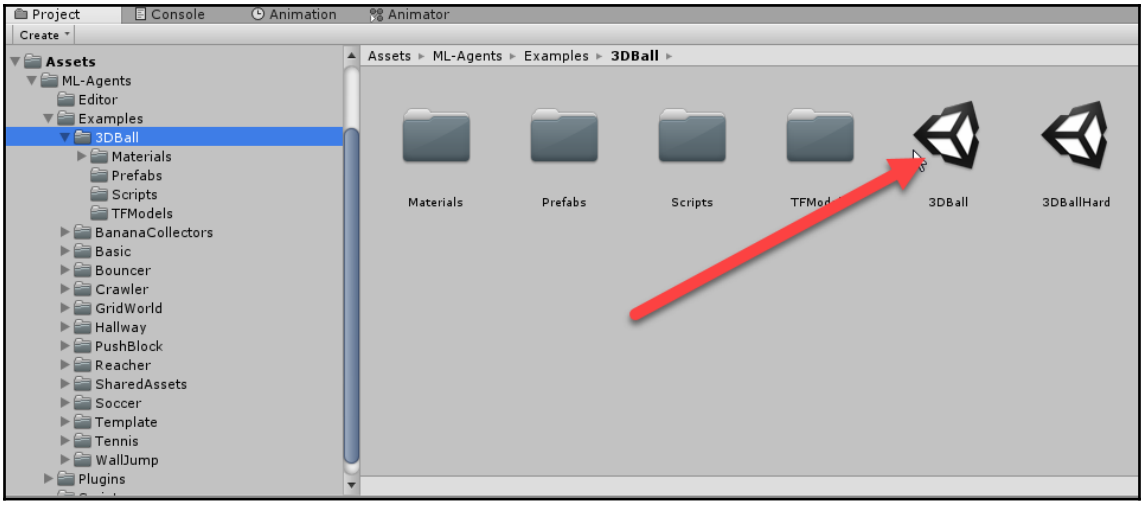
Project
Console
Animation
Animator
Create
Assets
ML-Agents
Editor
Examples
3DBall
Materials
Prefabs
Scripts
TFModels
BananaCollectors
Basic
Bouncer
Crawler
GridWorld
Hallway
PushBlock
Reacher
SharedAssets
Soccer
Template
Tennis
WallJump
Plugins
Assets ▸ ML-Agents ▸ Examples ▸ 3DBall ▸
Materials
Prefabs
Scripts
3DBall
3DBallHard

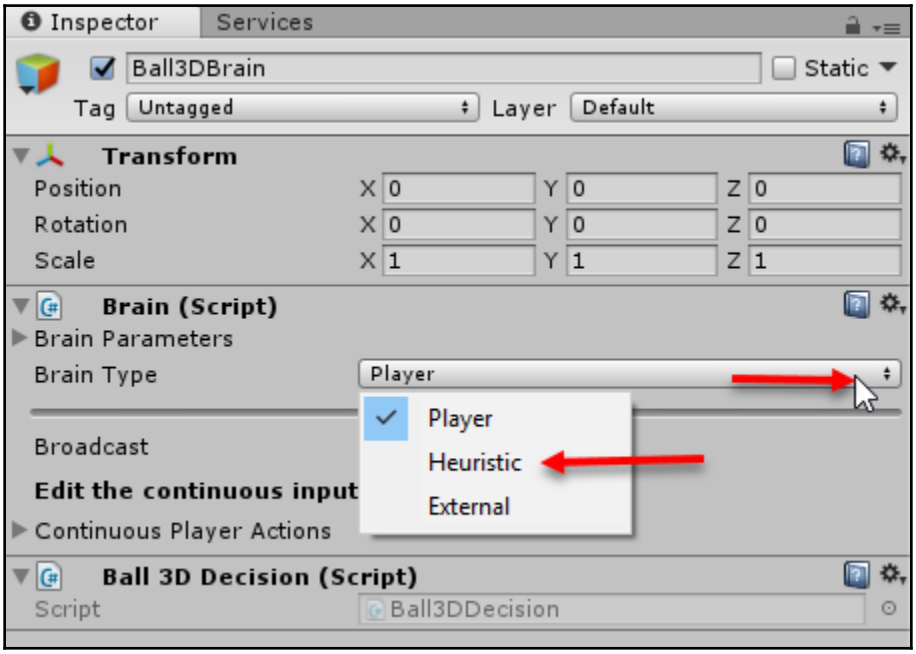
Inspector
Services
Ball3DBrain
Static
Tag
Untagged
Layer
Default
Transform
Position
Rotation
Scale
X 0
Y 0
Z 0
X 0
Y 0
Z 0
X 1
Y 1
Z 1
Brain (Script)
Brain Parameters
Brain Type
Player
Broadcast
Edit the continuous input
Continuous Player Actions
Player
Heuristic
External
Ball 3D Decision (Script)
Script
Ball3DDecision

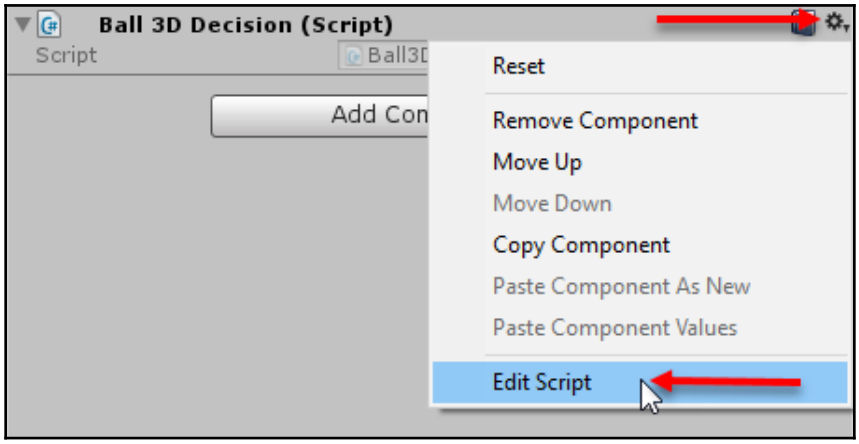
Ball 3D Decision (Script)
Script
Ball3D
Add Con
Reset
Remove Component
Move Up
Move Down
Copy Component
Paste Component As New
Paste Component Values
Edit Script

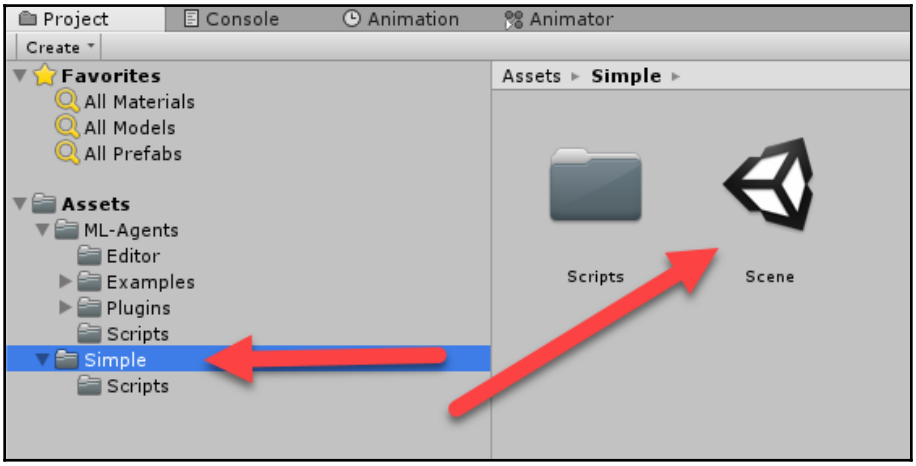
Project
Console
Animation
Animator
Create
Favorites
All Materials
All Models
All Prefabs
Assets
ML-Agents
Editor
Examples
Plugins
Scripts
Simple
Scripts
Assets ▸ Simple ▸
Scripts
Scene

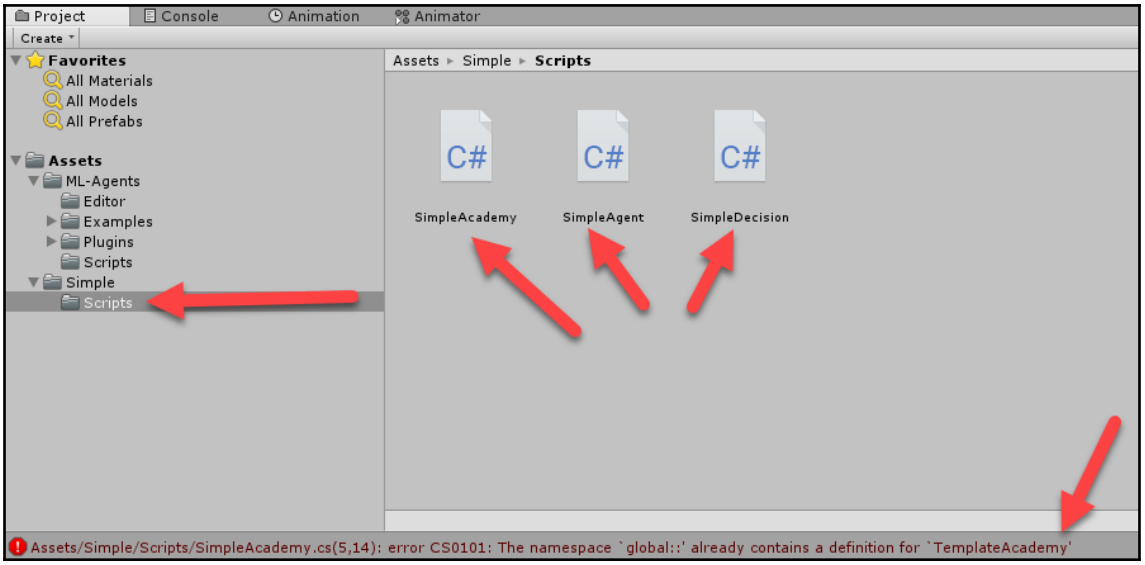
Project
Console
Animation
Animator
Create
Favorites
All Materials
All Models
All Prefabs
Assets
ML-Agents
Editor
Examples
Plugins
Scripts
Simple
Scripts
Assets ▸ Simple ▸ Scripts
C#
SimpleAcademy
C#
SimpleAgent
C#
SimpleDecision
Assets/Simple/Scripts/SimpleAcademy.cs(5,14): error CS0101: The namespace `global::' already contains a definition for `TemplateAcademy'

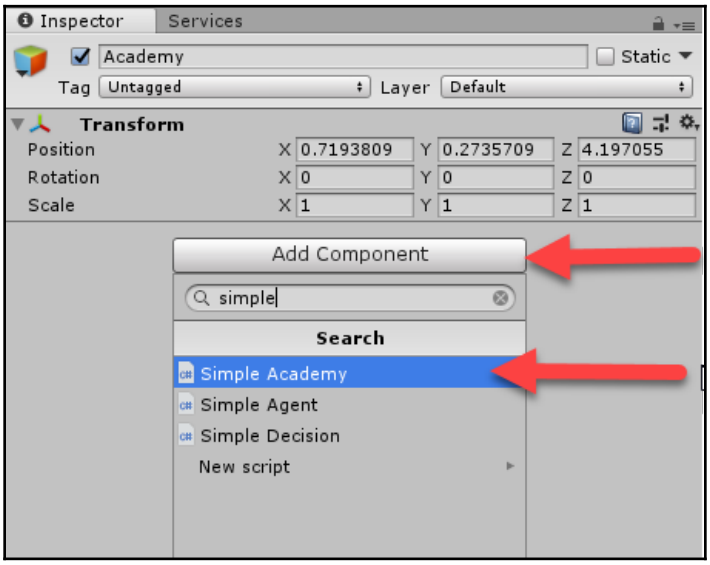
Inspector
Services
Academy
Static
Tag
Untagged
Layer
Default
Transform
Position
X 0.7193809
Y 0.2735709
Z 4.197055
Rotation
X 0
Y 0
Z 0
Scale
X 1
Y 1
Z 1
Add Component
simple
Search
Simple Academy
Simple Agent
Simple Decision
New script

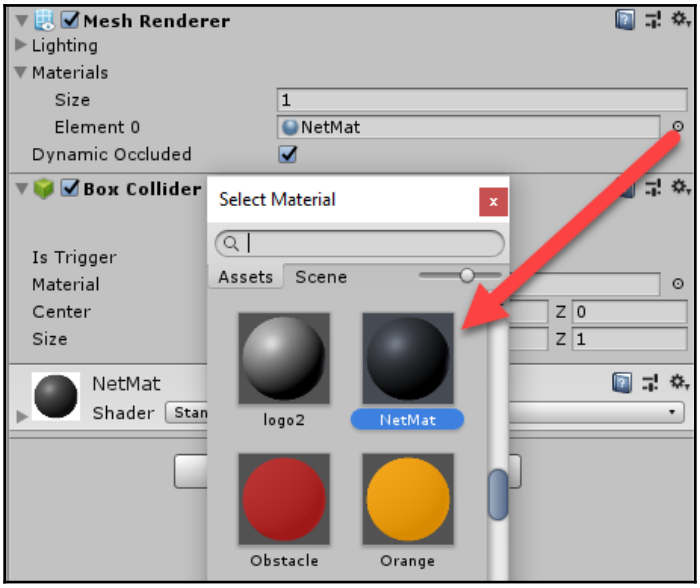
Mesh Renderer
Lighting
Materials
Size
1
Element 0
NetMat
Dynamic Occluded
Box Collider
Is Trigger
Material
Center
Size
Z 0
Z 1
Select Material
Assets
Scene
logo2
NetMat
Obstacle
Orange
NetMat
Shader

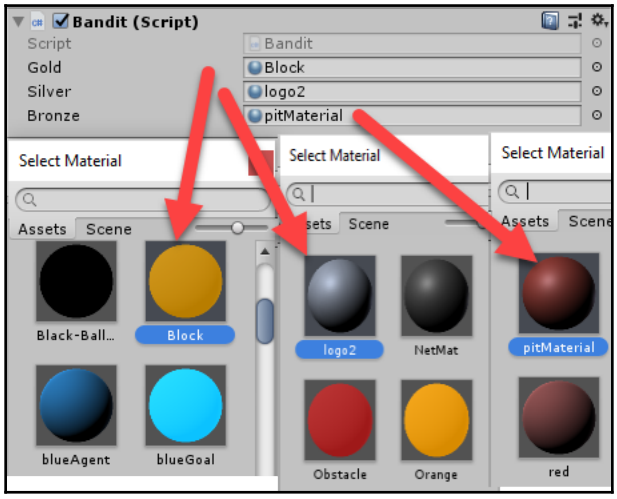
Bandit (Script)
Script
Bandit
Gold
Block
Silver
logo2
Bronze
pitMaterial
Select Material
Assets
Scene
Black-Ball...
Block
blueAgent
blueGoal
logo2
NetMat
Obstacle
Orange
pitMaterial
red

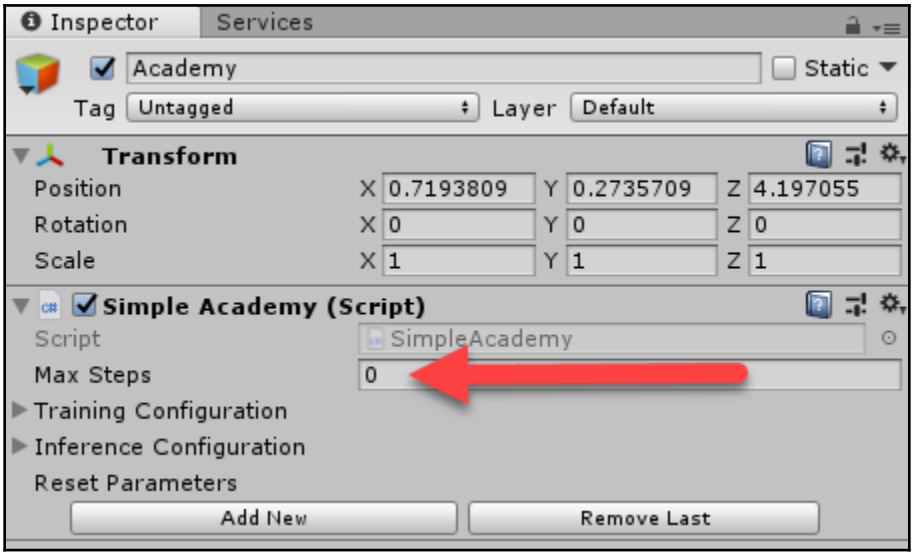
Inspector
Services
Academy
Static
Tag
Untagged
Layer
Default
Transform
Position
X 0.7193809
Y 0.2735709
Z 4.197055
Rotation
X 0
Y 0
Z 0
Scale
X 1
Y 1
Z 1
Simple Academy (Script)
Script
SimpleAcademy
Max Steps
0
Training Configuration
Inference Configuration
Reset Parameters
Add New
Remove Last

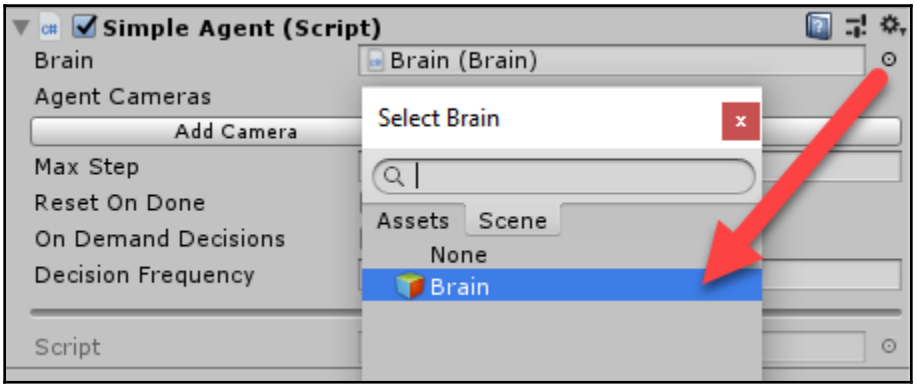
Simple Agent (Script)
Brain
Brain (Brain)
Agent Cameras
Add Camera
Max Step
Reset On Done
On Demand Decisions
Decision Frequency
Script
Select Brain
Assets
Scene
None
Brain

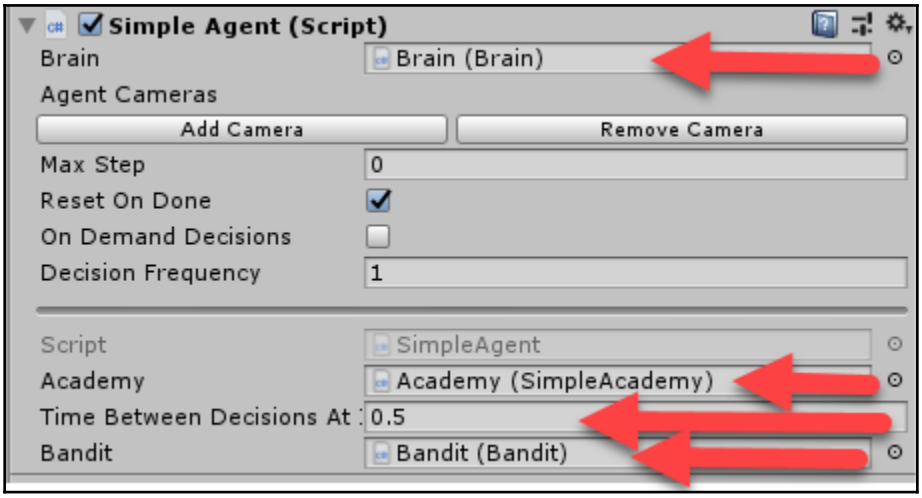
Simple Agent (Script)
Brain
Brain (Brain)
Agent Cameras
Add Camera
Remove Camera
Max Step
0
Reset On Done
On Demand Decisions
Decision Frequency
1
Script
SimpleAgent
Academy
Academy (SimpleAcademy)
Time Between Decisions At
0.5
Bandit
Bandit (Bandit)

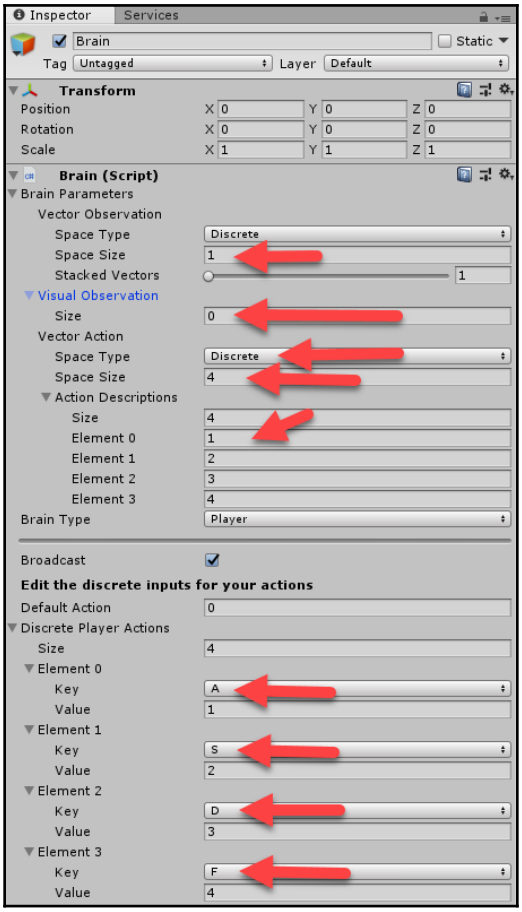
Inspector
Services
Brain
Static
Tag
Untagged
Layer
Default
Transform
Position
Rotation
Scale
Brain (Script)
Brain Parameters
Vector Observation
Space Type
Discrete
Space Size
1
Stacked Vectors
Visual Observation
Size
0
Vector Action
Space Type
Discrete
Space Size
4
Action Descriptions
Size
4
Element 0
1
Element 1
2
Element 2
3
Element 3
4
Brain Type
Player
Broadcast
Edit the discrete inputs for your actions
Default Action
Discrete Player Actions
Size
Element 0
Key
A
Value
Element 1
Key
S
Value
Element 2
Key
D
Value
Element 3
Key
F
Value

Chapter 2:
The Bandit and Reinforcement Learning

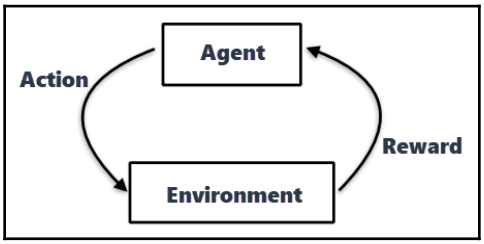

$$V(a) = V(a) + \alpha \times (r - V(a))$$

$$V(a) =$$

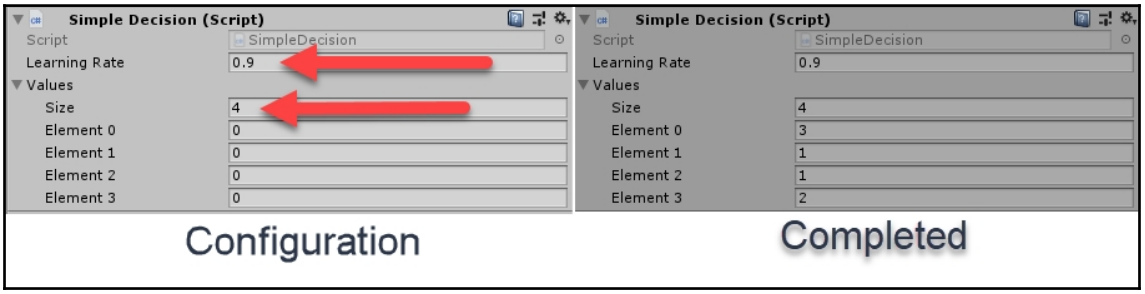

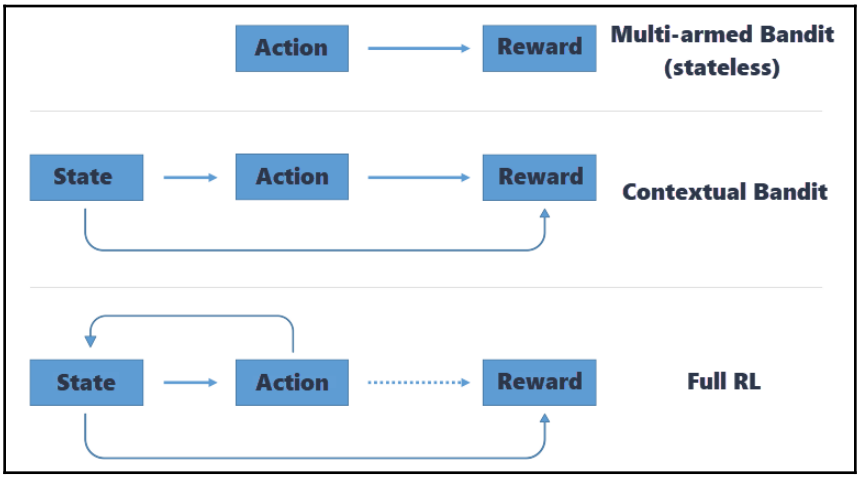

$$\boxed{Q[s,a] = Q[s,a] + \alpha \times (r - Q[s,a])}$$

$$\boxed{Q[s,a] =}$$

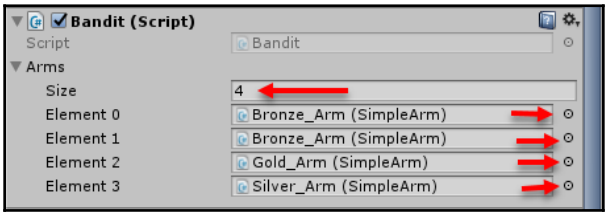

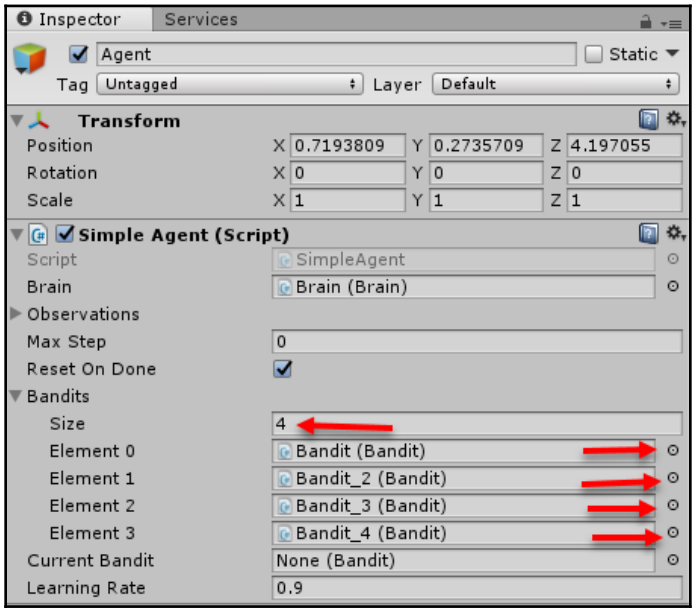

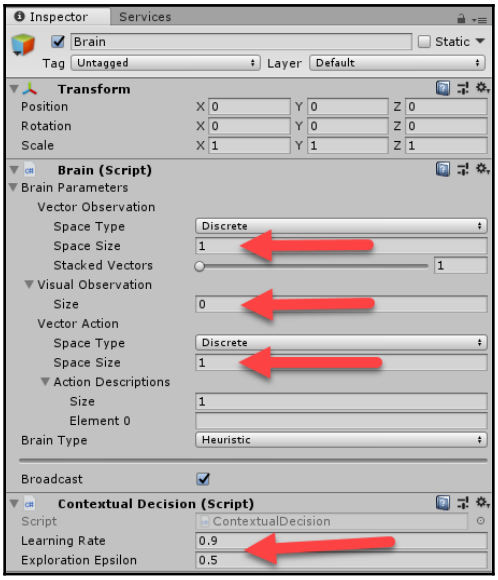

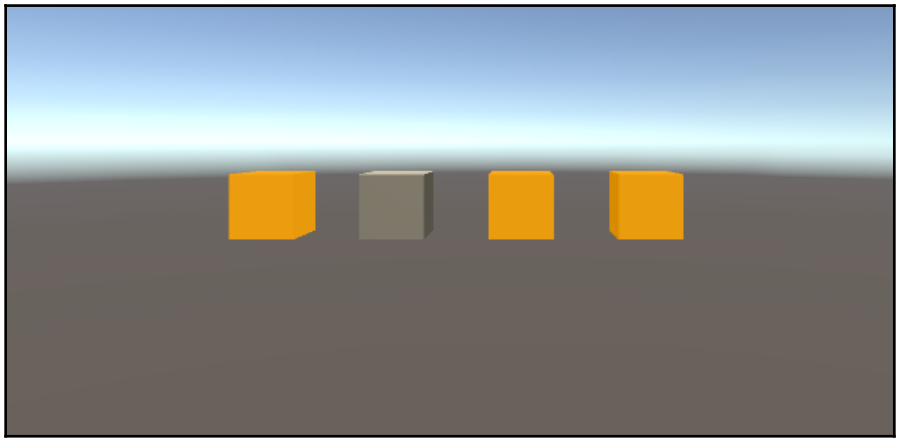

$$\boxed{Q(s,a) = r + \gamma \max_{a'} Q(s',a')}$$

$$\boxed{r = reward}$$

$$\boxed{\gamma = gamma\ (reward\ discount\ factor\ 0 - 1.0)}$$

$$\boxed{\max_{a'} = maximum\ of\ all\ actions\ for\ state}$$

$$\boxed{Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha(r_{t+1} + \gamma \max_a Q_t(s_{t+1}, a) - Q_t(s_{t+1}, a_t))}$$

$$\boxed{\alpha = learning\ rate}$$

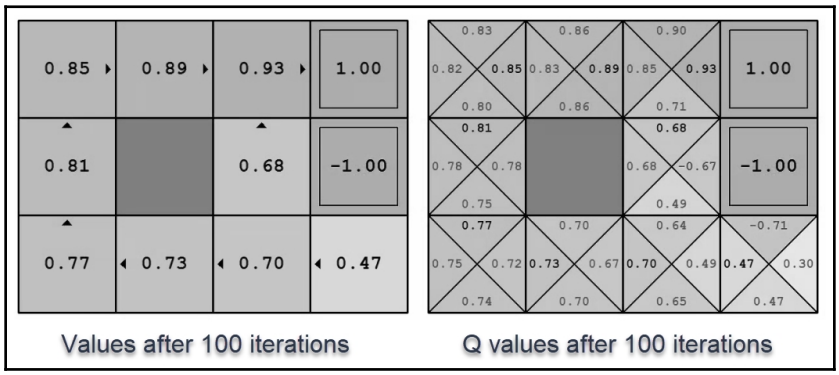

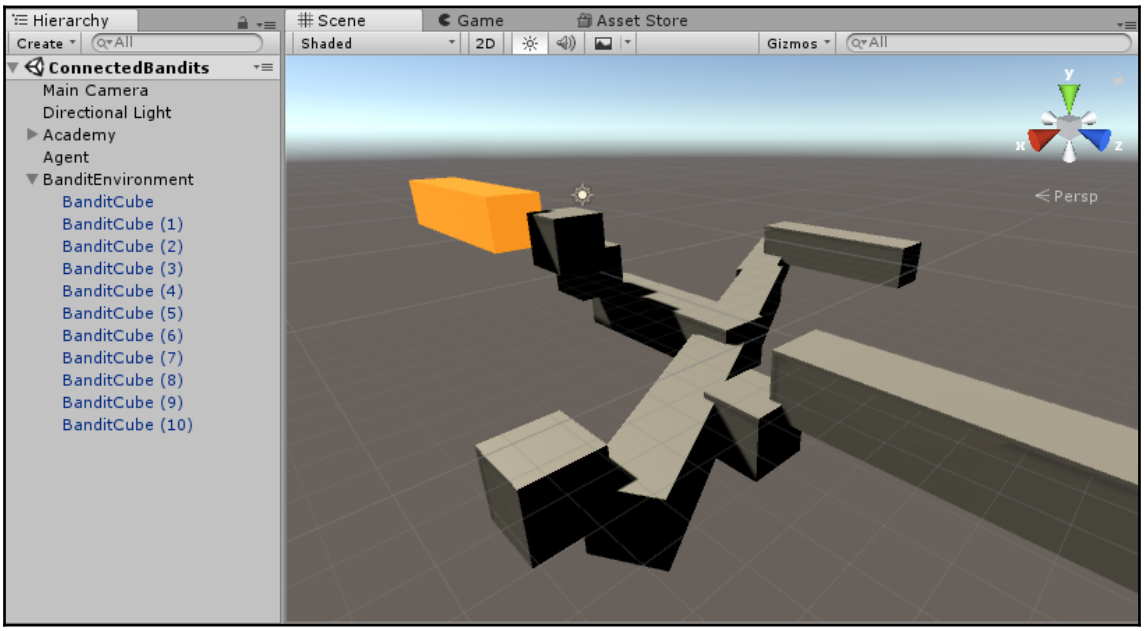

Chapter 3:
Deep Reinforcement Learning with Python

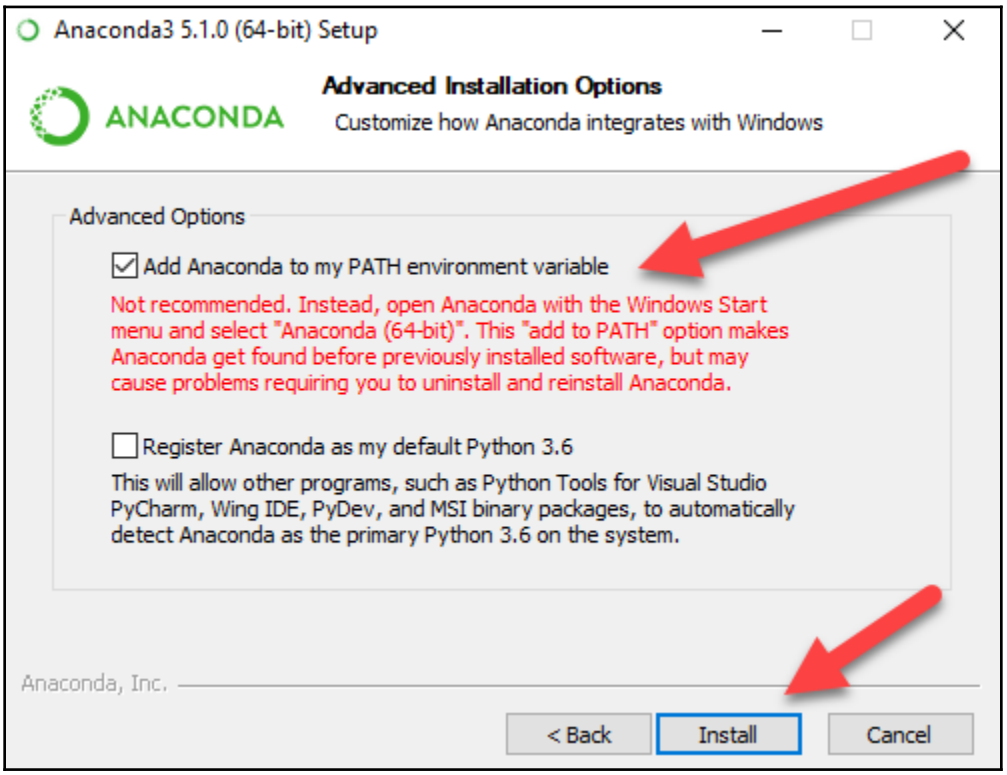

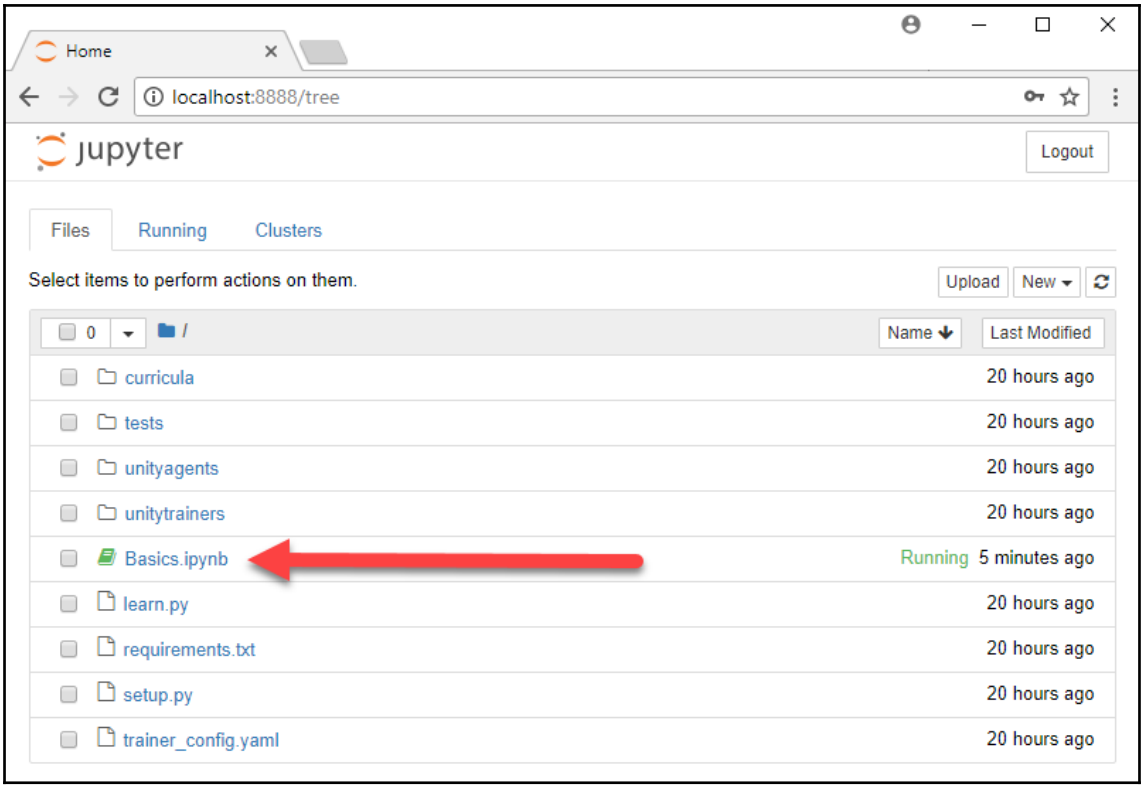
Home
localhost:8888/tree
jupyter
Logout
Files
Running
Clusters
Select items to perform actions on them.
Upload
New
0
/
Name
Last Modified
curricula
20 hours ago
tests
20 hours ago
unityagents
20 hours ago
unitytrainers
20 hours ago
Basics.ipynb
Running 5 minutes ago
learn.py
20 hours ago
requirements.txt
20 hours ago
setup.py
20 hours ago
trainer_config.yaml
20 hours ago

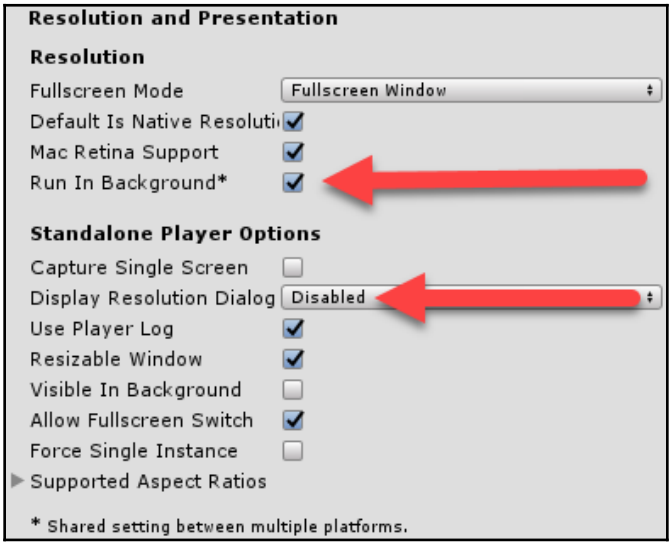
Resolution and Presentation
Resolution
Fullscreen Mode
Fullscreen Window
Default Is Native Resoluti
Mac Retina Support
Run In Background*
Standalone Player Options
Capture Single Screen
Display Resolution Dialog
Disabled
Use Player Log
Resizable Window
Visible In Background
Allow Fullscreen Switch
Force Single Instance
Supported Aspect Ratios
* Shared setting between multiple platforms.

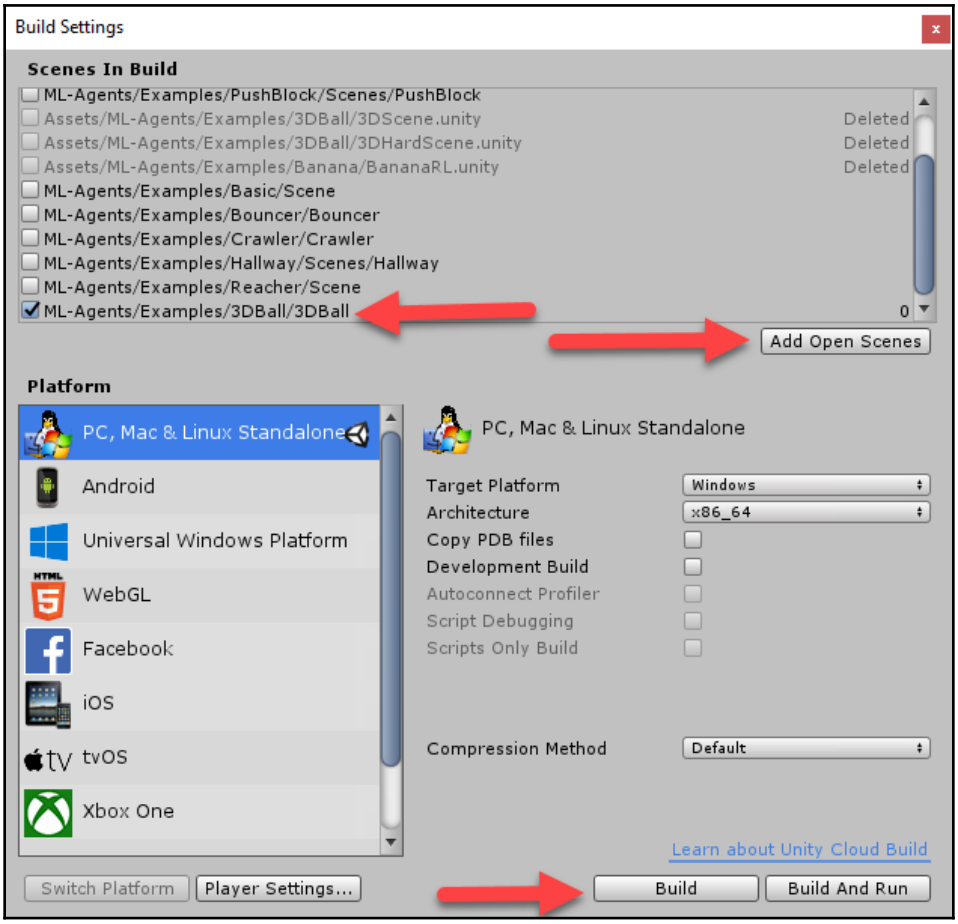
Build Settings
Scenes In Build
ML-Agents/Examples/PushBlock/Scenes/PushBlock
Assets/ML-Agents/Examples/3DBall/3DScene.unity
Deleted
Assets/ML-Agents/Examples/3DBall/3DHardScene.unity
Deleted
Assets/ML-Agents/Examples/Banana/BananaRL.unity
Deleted
ML-Agents/Examples/Basic/Scene
ML-Agents/Examples/Bouncer/Bouncer
ML-Agents/Examples/Crawler/Crawler
ML-Agents/Examples/Hallway/Scenes/Hallway
ML-Agents/Examples/Reacher/Scene
ML-Agents/Examples/3DBall/3DBall
0
Add Open Scenes
Platform
PC, Mac & Linux Standalone
Android
Universal Windows Platform
WebGL
Facebook
iOS
tvOS
Xbox One
PC, Mac & Linux Standalone
Target Platform
Windows
Architecture
x86_64
Copy PDB files
Development Build
Autoconnect Profiler
Script Debugging
Scripts Only Build
Compression Method
Default
Learn about Unity Cloud Build
Switch Platform
Player Settings...
Build
Build And Run

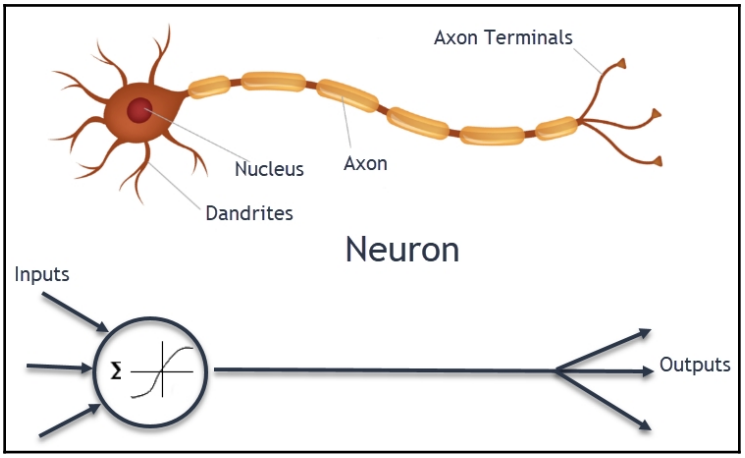
Axon Terminals
Nucleus
Axon
Dandrites
Neuron
Inputs
Σ
Outputs

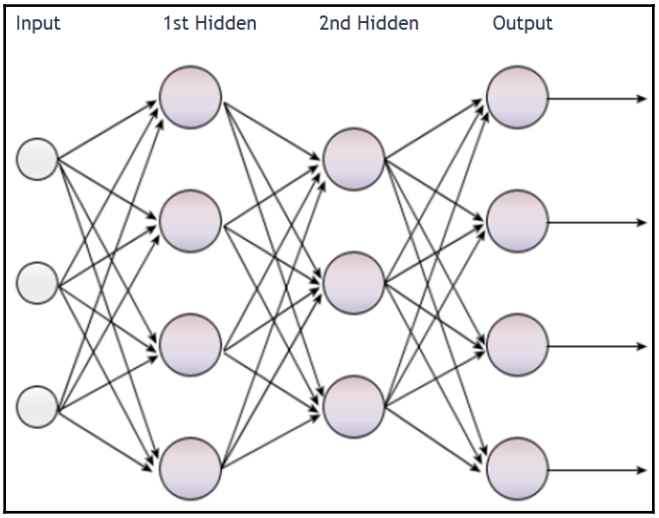
Input
1st Hidden
2nd Hidden
Output

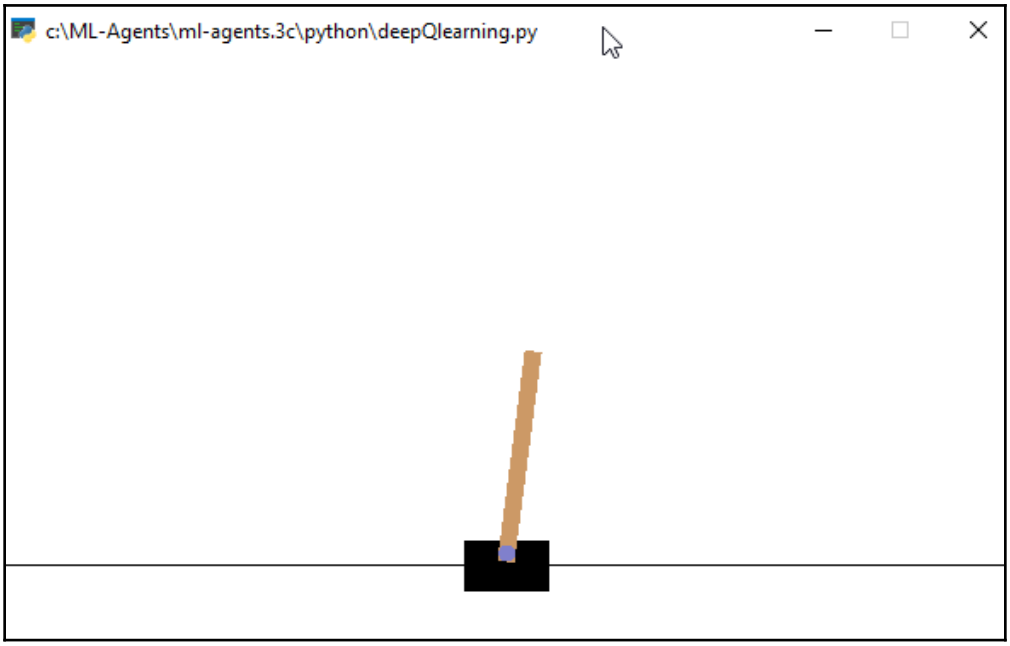


```
20    model = Sequential()
21    model.add(Flatten(input_shape=(1,) + env.observa
22    model.add(Dense(16))
23    model.add(Activation('relu'))
24    model.add(Dense(nb_actions))
25    model.add(Activation('linear'))
26    print(model.summary())
27
28    policy = EpsGreedyQPolicy()
29    memory = SequentialMemory(limit=50000, window_le
30    dqn = DQNAgent(model=model, nb_actions=nb_action
31    target_model_update=1e-2, policy=policy)
32    dqn.compile(Adam(lr=1e-3), metrics=['mae'])
```

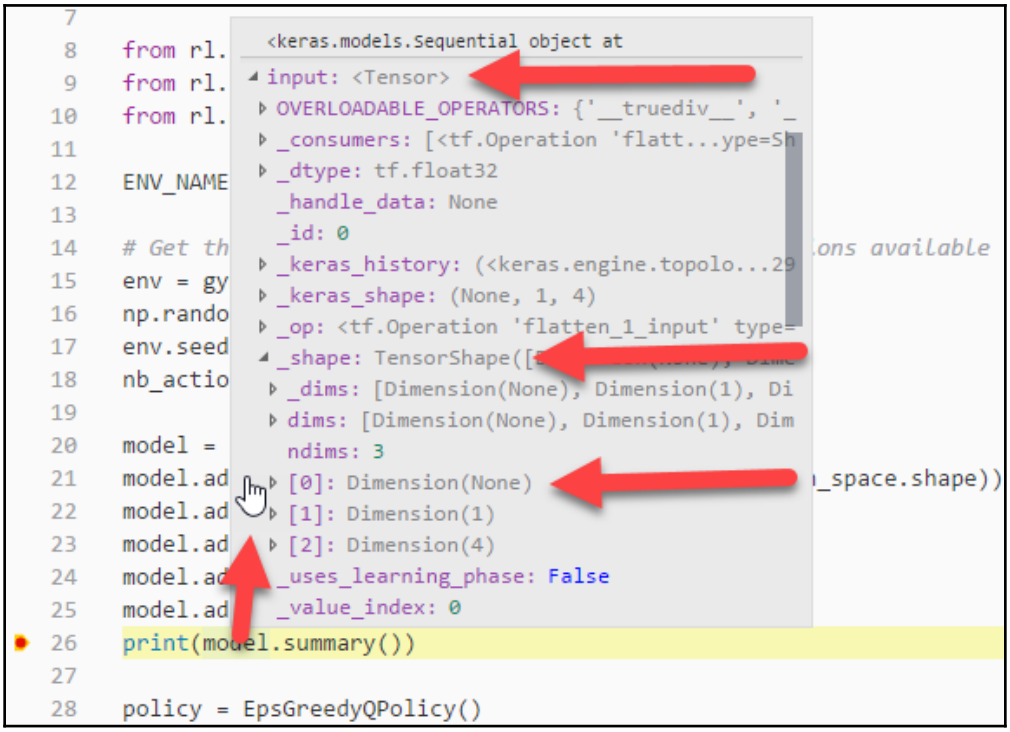
from rl.
from rl.
from rl.
ENV_NAME
Get th
ons available
env = gy
np.rando
env.seed
nb_actio
model =
model.ad
_space.shape))
model.ad
model.ad
model.ad
model.ad
print(model.summary())
policy = EpsGreedyQPolicy()
<keras.models.Sequential object at
input: <Tensor>
OVERLOADABLE_OPERATORS: {'__truediv__', '_
_consumers: [<tf.Operation 'flatt...ype=Sh
_dtype: tf.float32
_handle_data: None
_id: 0
_keras_history: (<keras.engine.topolo...29
_keras_shape: (None, 1, 4)
_op: <tf.Operation 'flatten_1_input' type=
_shape: TensorShape([
_dims: [Dimension(None), Dimension(1), Di
dims: [Dimension(None), Dimension(1), Dim
ndims: 3
[0]: Dimension(None)
[1]: Dimension(1)
[2]: Dimension(4)
_uses_learning_phase: False
_value_index: 0

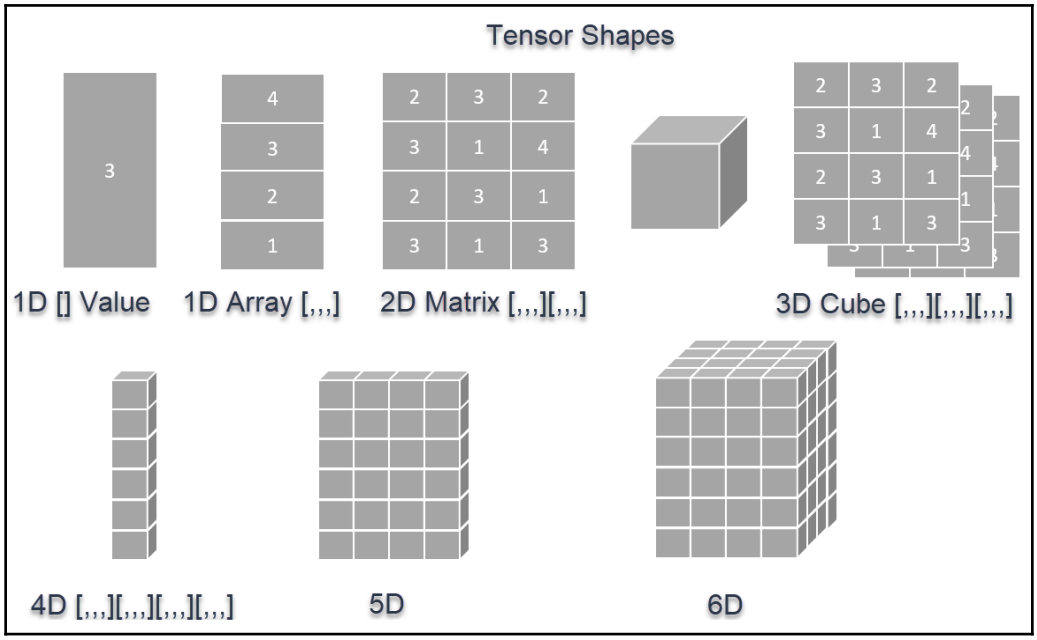
Tensor Shapes
3
4
3
2
1
2 3 2
3 1 4
2 3 1
3 1 3
2 3 2
3 1 4
2 3 1
3 1 3
1D [] Value
1D Array [,,,]
2D Matrix [,,,][,,,]
3D Cube [,,,][,,,][,,,]
4D [,,,][,,,][,,,][,,,]
5D
6D

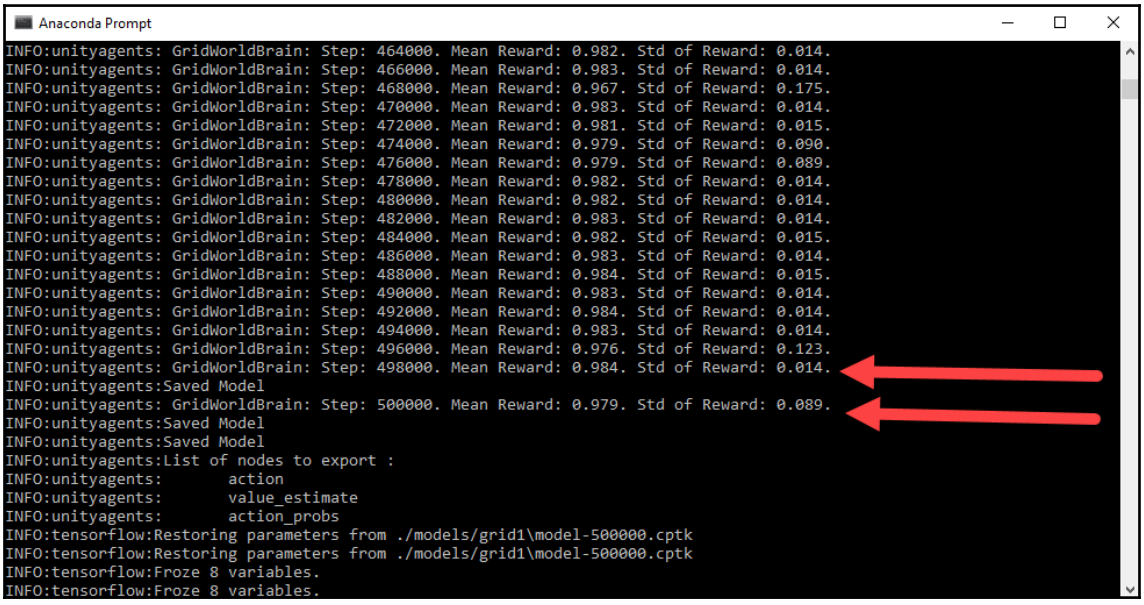
Anaconda Prompt
INFO:unityagents: GridWorldBrain: Step: 464000. Mean Reward: 0.982. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 466000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 468000. Mean Reward: 0.967. Std of Reward: 0.175.
INFO:unityagents: GridWorldBrain: Step: 470000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 472000. Mean Reward: 0.981. Std of Reward: 0.015.
INFO:unityagents: GridWorldBrain: Step: 474000. Mean Reward: 0.979. Std of Reward: 0.090.
INFO:unityagents: GridWorldBrain: Step: 476000. Mean Reward: 0.979. Std of Reward: 0.089.
INFO:unityagents: GridWorldBrain: Step: 478000. Mean Reward: 0.982. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 480000. Mean Reward: 0.982. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 482000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 484000. Mean Reward: 0.982. Std of Reward: 0.015.
INFO:unityagents: GridWorldBrain: Step: 486000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 488000. Mean Reward: 0.984. Std of Reward: 0.015.
INFO:unityagents: GridWorldBrain: Step: 490000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 492000. Mean Reward: 0.984. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 494000. Mean Reward: 0.983. Std of Reward: 0.014.
INFO:unityagents: GridWorldBrain: Step: 496000. Mean Reward: 0.976. Std of Reward: 0.123.
INFO:unityagents: GridWorldBrain: Step: 498000. Mean Reward: 0.984. Std of Reward: 0.014.
INFO:unityagents:Saved Model
INFO:unityagents: GridWorldBrain: Step: 500000. Mean Reward: 0.979. Std of Reward: 0.089.
INFO:unityagents:Saved Model
INFO:unityagents:Saved Model
INFO:unityagents:List of nodes to export :
INFO:unityagents: action
INFO:unityagents: value_estimate
INFO:unityagents: action_probs
INFO:tensorflow:Restoring parameters from ./models/grid1\model-500000.cptk
INFO:tensorflow:Restoring parameters from ./models/grid1\model-500000.cptk
INFO:tensorflow:Froze 8 variables.
INFO:tensorflow:Froze 8 variables.

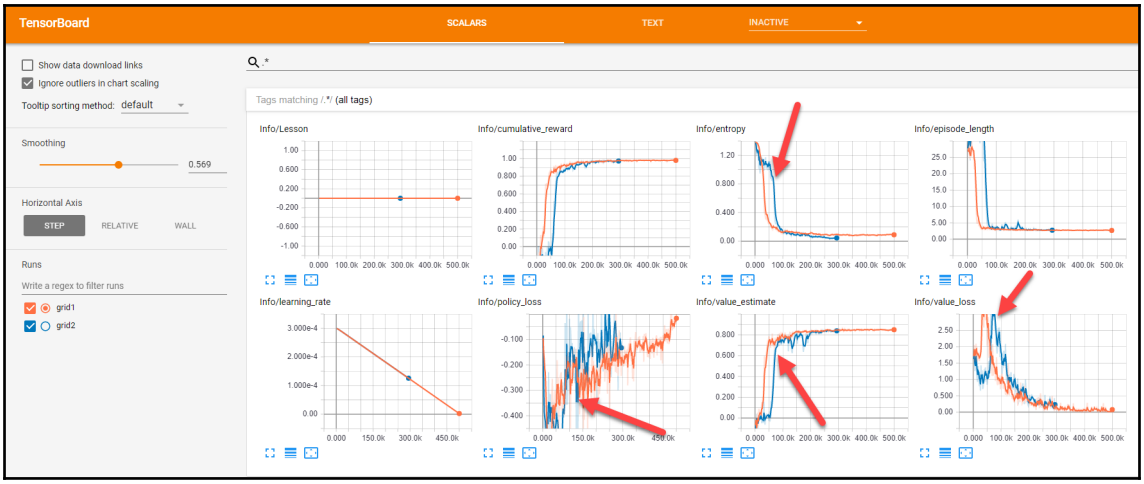
TensorBoard
SCALARS
TEXT
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method: default
Smoothing
0.569
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
grid1
grid2
Tags matching /.*/ (all tags)
Info/Lesson
Info/cumulative_reward
Info/entropy
Info/episode_length
Info/learning_rate
Info/policy_loss
Info/value_estimate
Info/value_loss

Chapter 4: Going Deeper with Deep Learning

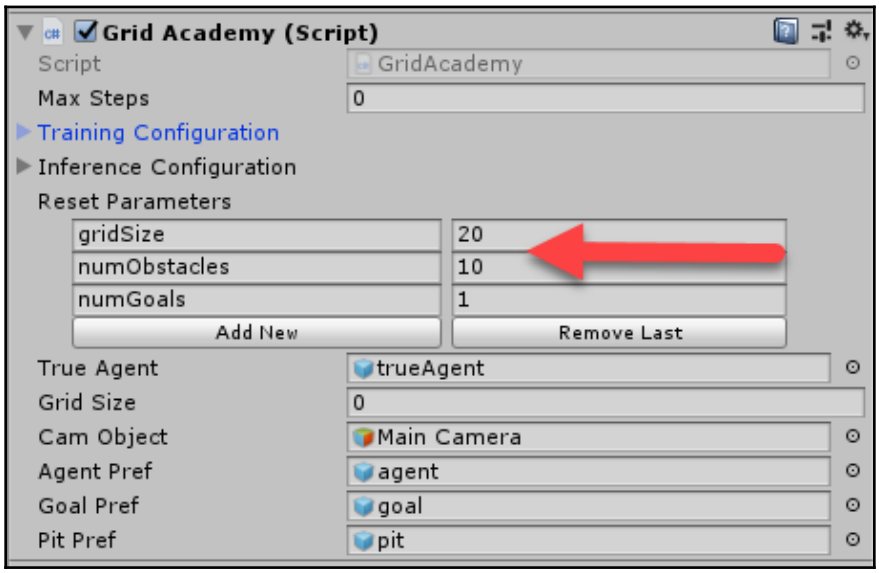

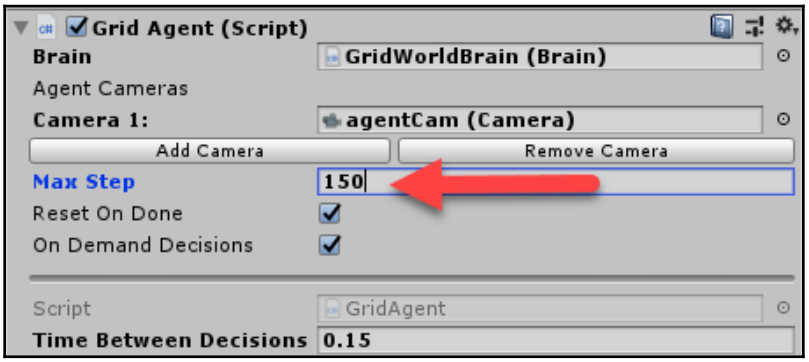


```
INFO:unityagents: GridWorldBrain: Step: 2000. Mean Reward: -1.394. Std of Reward: 0.468.
INFO:unityagents: GridWorldBrain: Step: 4000. Mean Reward: -1.390. Std of Reward: 0.733.
INFO:unityagents: GridWorldBrain: Step: 6000. Mean Reward: -1.319. Std of Reward: 0.490.
INFO:unityagents: GridWorldBrain: Step: 8000. Mean Reward: -1.273. Std of Reward: 0.683.
INFO:unityagents: GridWorldBrain: Step: 10000. Mean Reward: -1.158. Std of Reward: 0.647.
INFO:unityagents: GridWorldBrain: Step: 12000. Mean Reward: -1.413. Std of Reward: 0.274.
INFO:unityagents: GridWorldBrain: Step: 14000. Mean Reward: -1.109. Std of Reward: 0.883.
INFO:unityagents: GridWorldBrain: Step: 16000. Mean Reward: -1.251. Std of Reward: 0.472.
INFO:unityagents: GridWorldBrain: Step: 18000. Mean Reward: -1.203. Std of Reward: 0.593.
INFO:unityagents: GridWorldBrain: Step: 20000. Mean Reward: -1.199. Std of Reward: 0.672.
```

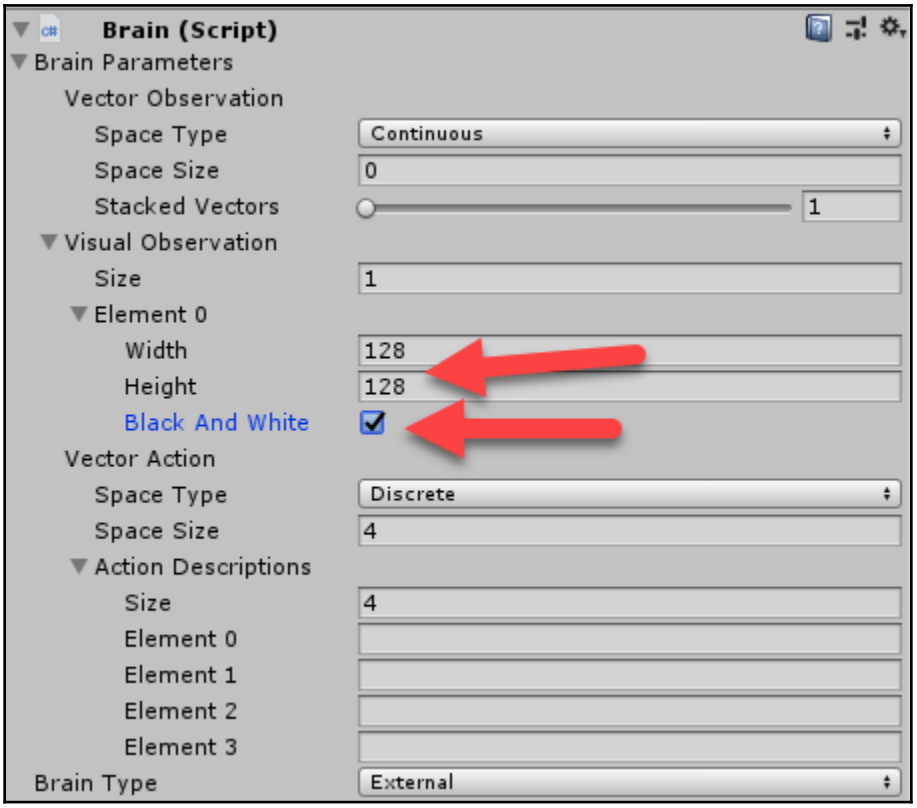
Brain (Script)
Brain Parameters
Vector Observation
Space Type
Continuous
Space Size
0
Stacked Vectors
1
Visual Observation
Size
1
Element 0
Width
128
Height
128
Black And White
Vector Action
Space Type
Discrete
Space Size
4
Action Descriptions
Size
4
Element 0
Element 1
Element 2
Element 3
Brain Type
External

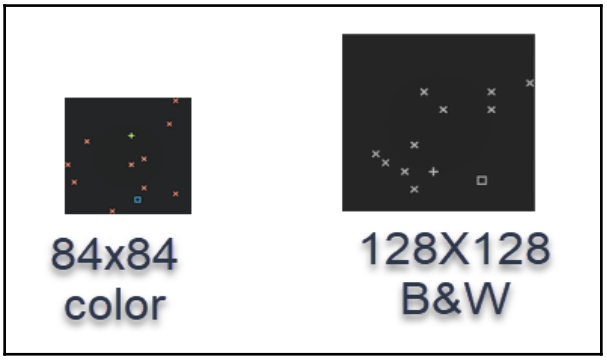
84x84
color
128X128
B&W

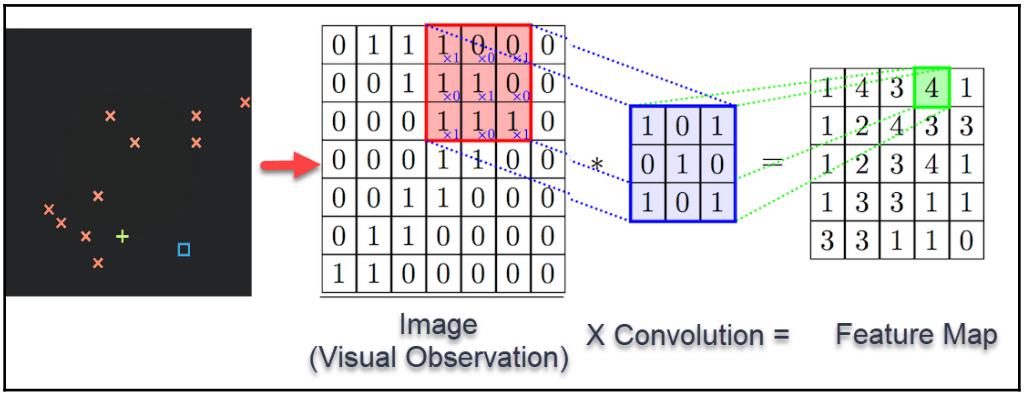
Image
(Visual Observation)
X Convolution =
Feature Map

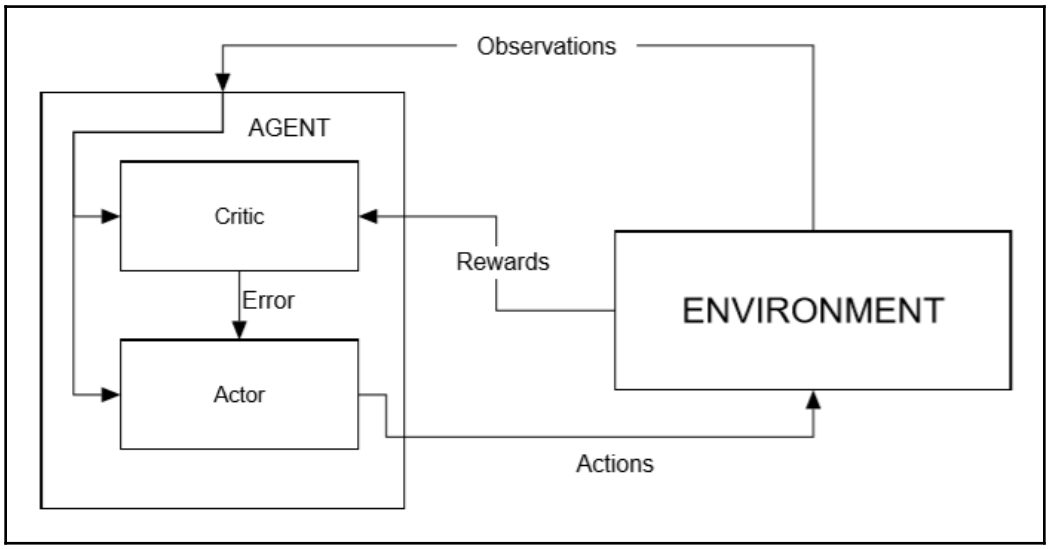
Observations
AGENT
Critic
Error
Actor
Rewards
ENVIRONMENT
Actions

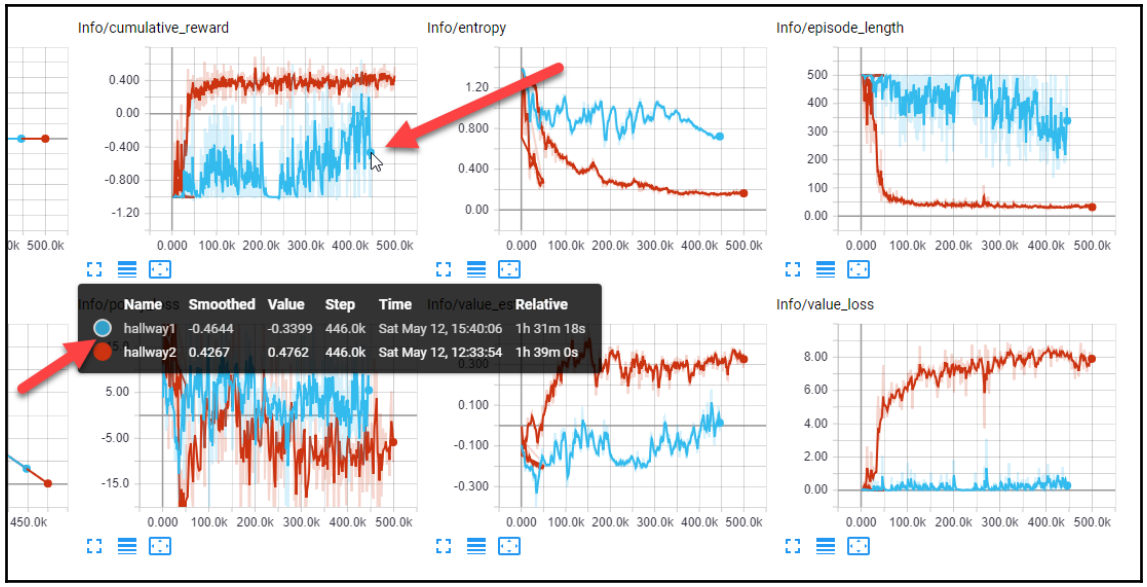
Info/cumulative_reward
Info/entropy
Info/episode_length
Info/value_loss
Name Smoothed Value Step Time Relative
hallway1 -0.4644 -0.3399 446.0k Sat May 12, 15:40:06 1h 31m 18s
hallway2 0.4267 0.4762 446.0k Sat May 12, 12:33:54 1h 39m 0s

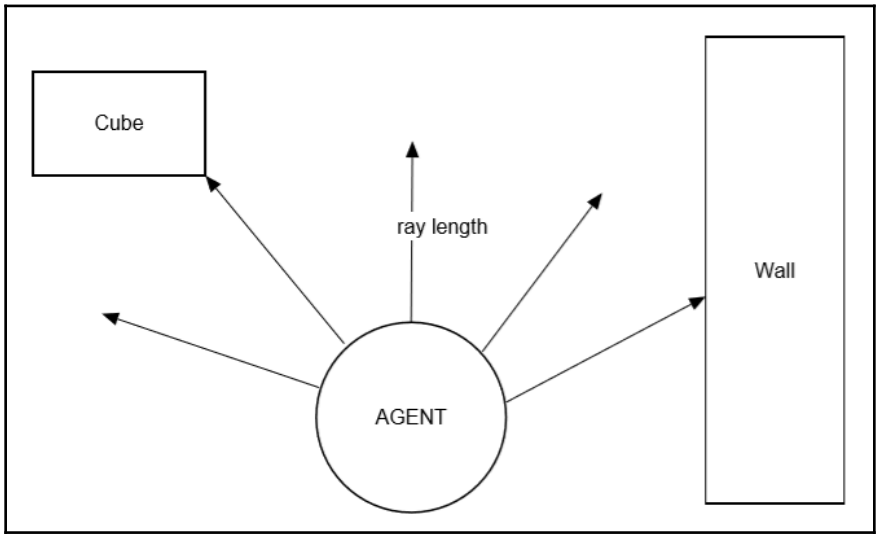
Cube
ray length
Wall
AGENT

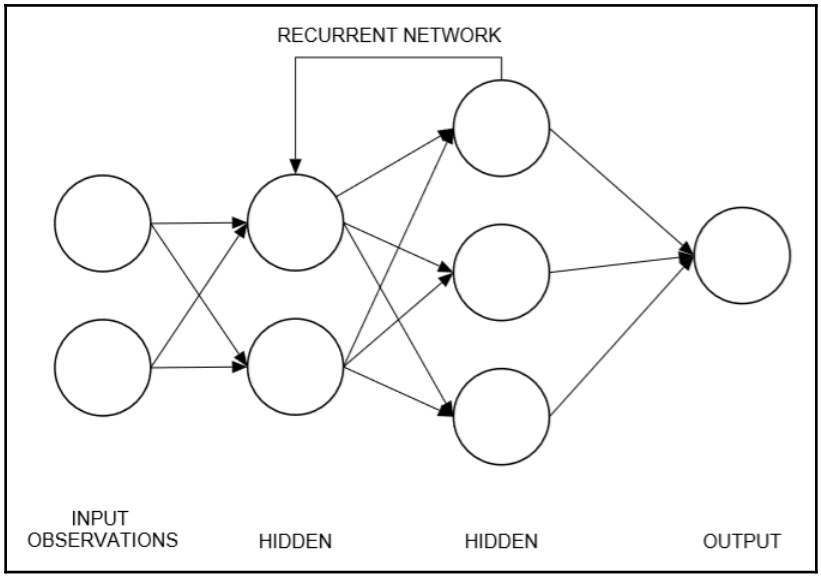
RECURRENT NETWORK
INPUT
OBSERVATIONS
HIDDEN
HIDDEN
OUTPUT

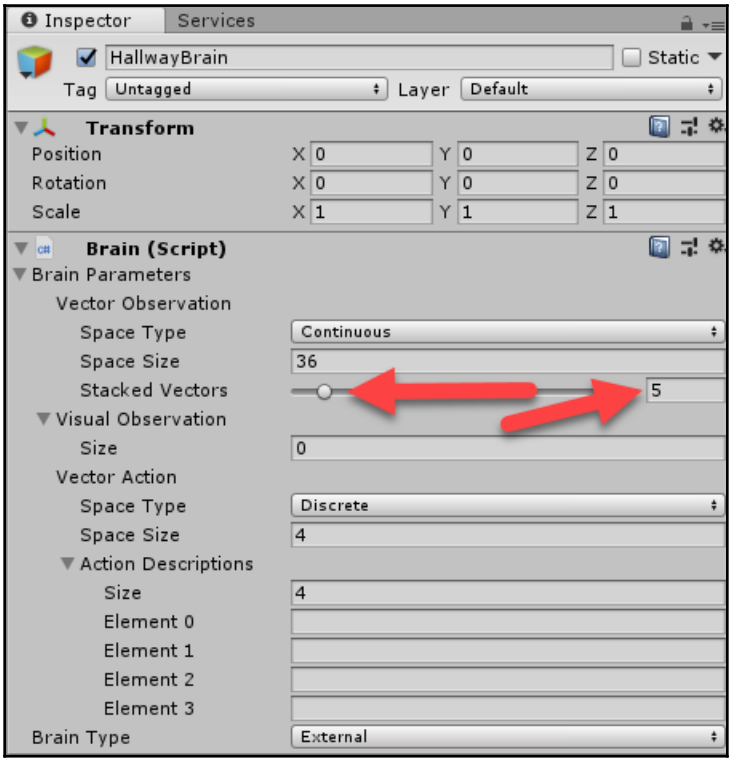
Inspector
Services
HallwayBrain
Static
Tag
Untagged
Layer
Default
Transform
Position
Rotation
Scale
Brain (Script)
Brain Parameters
Vector Observation
Space Type
Continuous
Space Size
36
Stacked Vectors
5
Visual Observation
Size
0
Vector Action
Space Type
Discrete
Space Size
4
Action Descriptions
Size
4
Element 0
Element 1
Element 2
Element 3
Brain Type
External

$$\boxed{Advantage: A = Q(s,a) - V(s)}$$

$$\boxed{Advantage: A = R - V(s)}$$

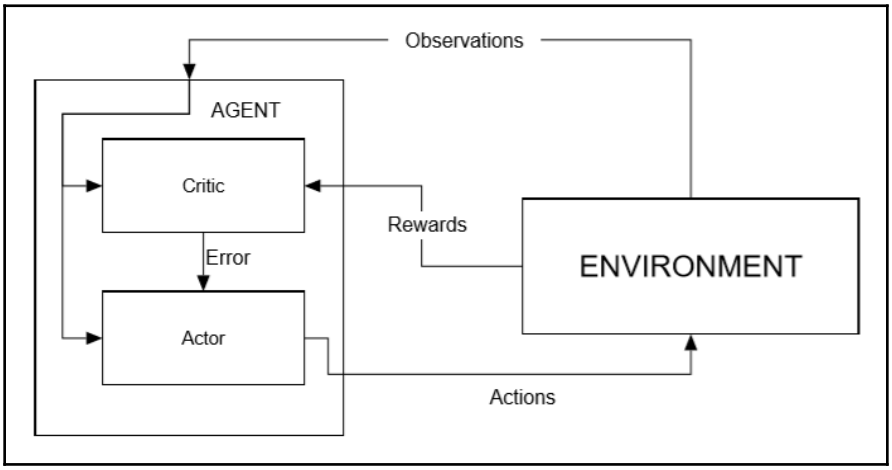

$$\boxed{ValueLoss: L = \Sigma(R - V(s))^2 (SumSquaredError)}$$

$$\boxed{PolicyLoss: L = -log(\pi(a|s)) * A(s)}$$

$$\boxed{H(\pi) = -\Sigma(P(x)log(P(x))}$$

$$\boxed{PolicyLoss: L = -log(\pi(a|s)) * A(s) - \beta * H(\pi)}$$

$$\boxed{L = 0.5 * \Sigma(R - V(s))^2 - log(\pi(a|s)) * A(s) - \beta * H(\pi)}$$

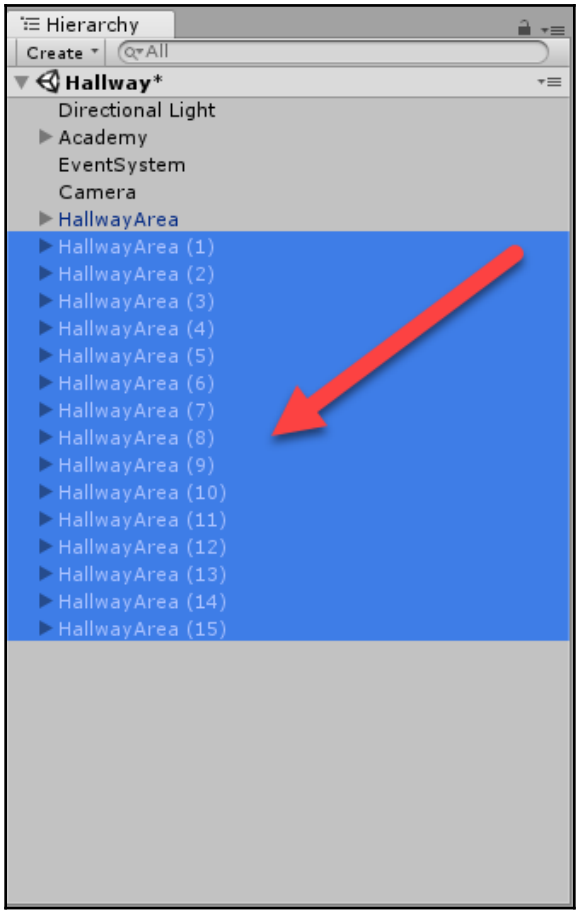
Hierarchy
Create
All
Hallway*
Directional Light
Academy
EventSystem
Camera
HallwayArea
HallwayArea (1)
HallwayArea (2)
HallwayArea (3)
HallwayArea (4)
HallwayArea (5)
HallwayArea (6)
HallwayArea (7)
HallwayArea (8)
HallwayArea (9)
HallwayArea (10)
HallwayArea (11)
HallwayArea (12)
HallwayArea (13)
HallwayArea (14)
HallwayArea (15)

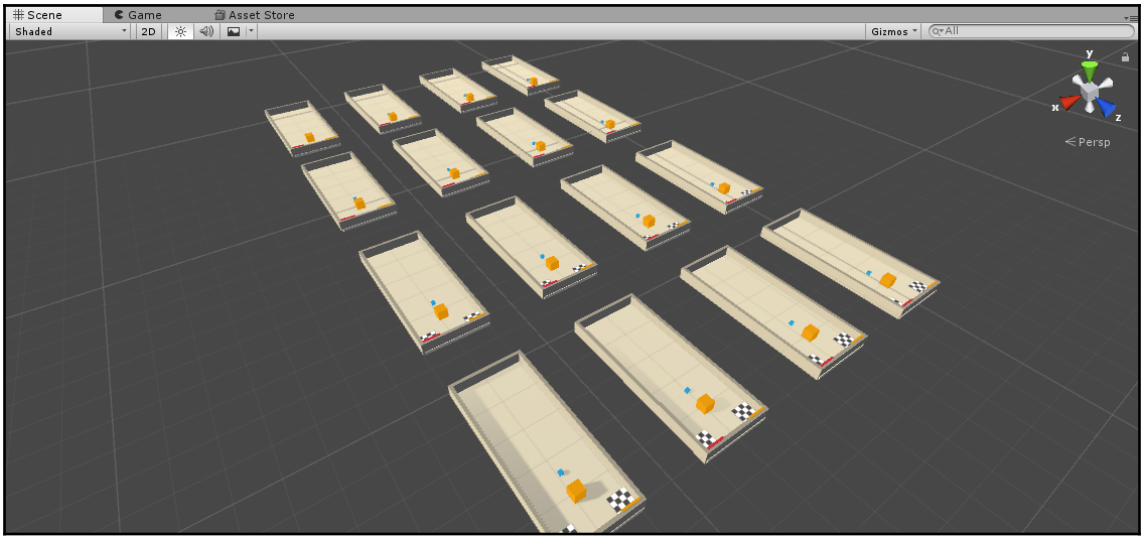
Scene
Game
Asset Store
Shaded
2D
Gizmos
Persp

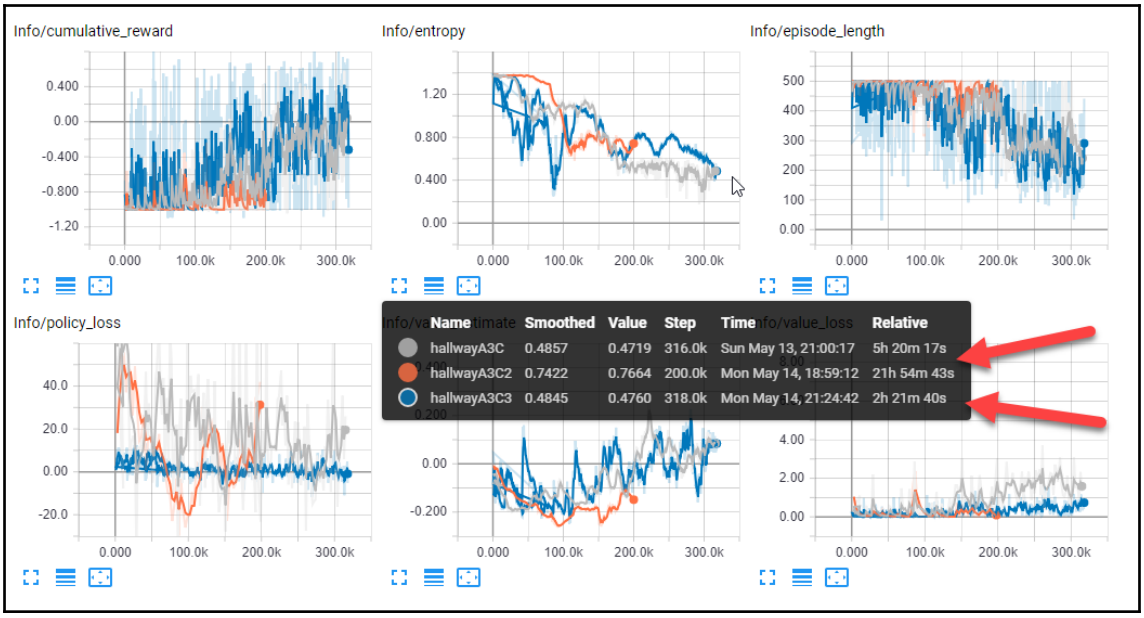
Info/cumulative_reward
Info/entropy
Info/episode_length
Info/policy_loss
Name Smoothed Value Step Time Relative
hallwayA3C 0.4857 0.4719 316.0k Sun May 13, 21:00:17 5h 20m 17s
hallwayA3C2 0.7422 0.7664 200.0k Mon May 14, 18:59:12 21h 54m 43s
hallwayA3C3 0.4845 0.4760 318.0k Mon May 14, 21:24:42 2h 21m 40s

Chapter 5: Playing the Game

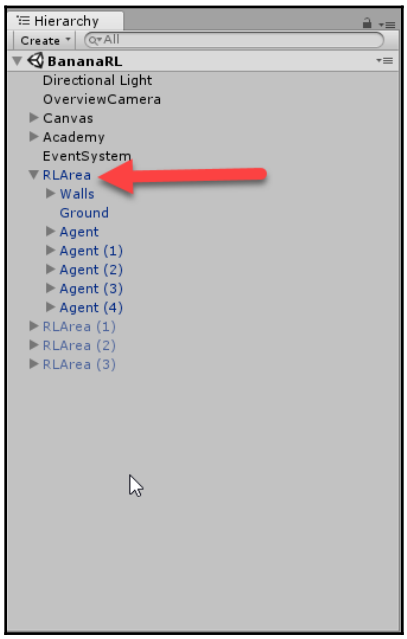

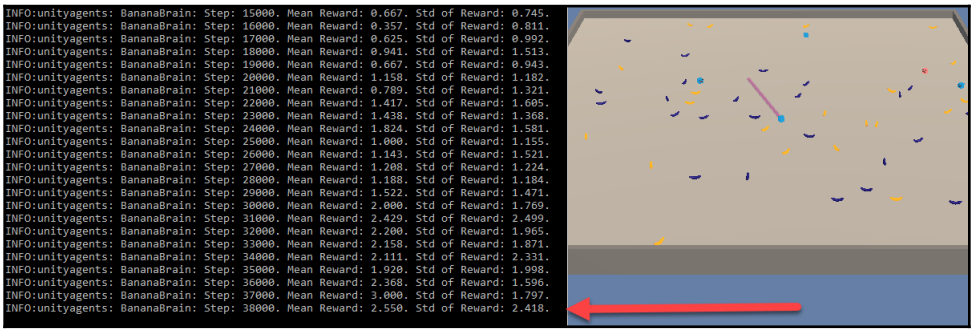


Anaconda Prompt - python python/learn.py python/python.exe --run-id=banana1 --train
use_recurrent: False
graph_scope: GoalieBrain
summary_path: ./summaries/banana1_GoalieBrain
memory_size: 256
INFO:unityagents:Hypermarameters for the PPO Trainer of brain StrikerBrain:
batch_size: 128
beta: 0.005
buffer_size: 2048
epsilon: 0.2
gamma: 0.99
hidden_units: 256
lambd: 0.95
learning_rate: 0.0003
max_steps: 1.0e5
normalize: False
num_epoch: 3
num_layers: 2
time_horizon: 128
sequence_length: 64
summary_freq: 2000
use_recurrent: False
graph_scope: StrikerBrain
summary_path: ./summaries/banana1_StrikerBrain
memory_size: 256
INFO:unityagents: GoalieBrain: Step: 2000. Mean Reward: 0.198. Std of Reward: 0.782.
INFO:unityagents: StrikerBrain: Step: 2000. Mean Reward: -0.198. Std of Reward: 0.782.
INFO:unityagents: GoalieBrain: Step: 4000. Mean Reward: 0.736. Std of Reward: 0.583.
INFO:unityagents: StrikerBrain: Step: 4000. Mean Reward: -0.736. Std of Reward: 0.583.
INFO:unityagents: GoalieBrain: Step: 6000. Mean Reward: 0.313.
INFO:unityagents: StrikerBrain: Step: 6000. Mean Reward: -0.313.
3DBall

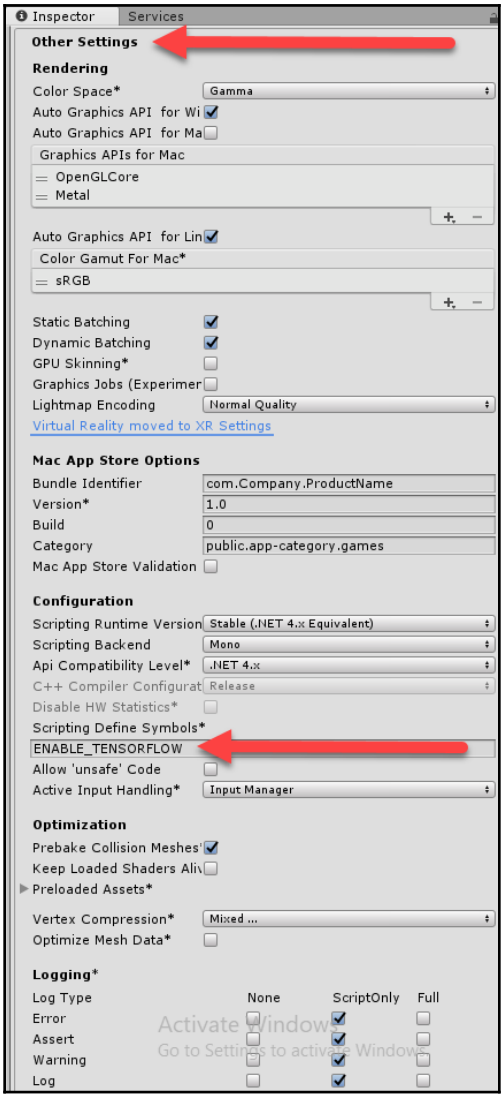
Inspector
Services
Other Settings
Rendering
Color Space*
Gamma
Auto Graphics API for Wi
Auto Graphics API for Ma
Graphics APIs for Mac
OpenGLCore
Metal
Auto Graphics API for Lin
Color Gamut For Mac*
sRGB
Static Batching
Dynamic Batching
GPU Skinning*
Graphics Jobs (Experimer
Lightmap Encoding
Normal Quality
Virtual Reality moved to XR Settings
Mac App Store Options
Bundle Identifier
com.Company.ProductName
Version*
1.0
Build
0
Category
public.app-category.games
Mac App Store Validation
Configuration
Scripting Runtime Version
Stable (.NET 4.x Equivalent)
Scripting Backend
Mono
Api Compatibility Level*
.NET 4.x
C++ Compiler Configurat
Release
Disable HW Statistics*
Scripting Define Symbols*
ENABLE_TENSORFLOW
Allow 'unsafe' Code
Active Input Handling*
Input Manager
Optimization
Prebake Collision Meshes
Keep Loaded Shaders Aliv
Preloaded Assets*
Vertex Compression*
Mixed ...
Optimize Mesh Data*
Logging*
Log Type
None
ScriptOnly
Full
Error
Assert
Warning
Log
Activate Windows
Go to Settings to activate Windows.

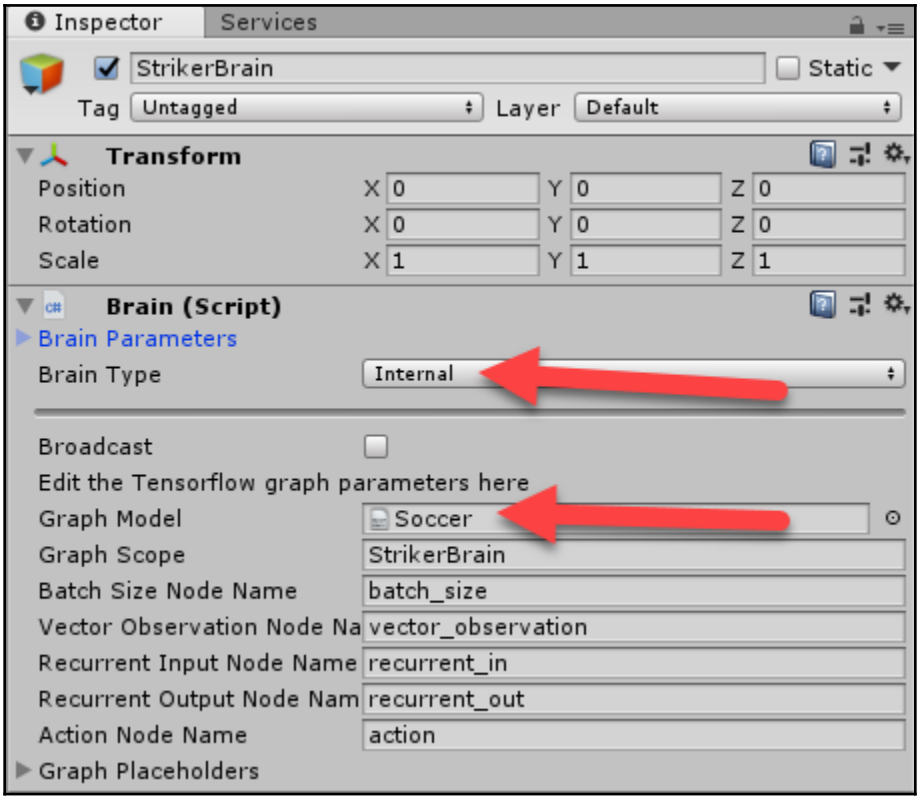
Inspector
Services
StrikerBrain
Static
Tag
Untagged
Layer
Default
Transform
Position
X 0
Y 0
Z 0
Rotation
X 0
Y 0
Z 0
Scale
X 1
Y 1
Z 1
Brain (Script)
Brain Parameters
Brain Type
Internal
Broadcast
Edit the Tensorflow graph parameters here
Graph Model
Soccer
Graph Scope
StrikerBrain
Batch Size Node Name
batch_size
Vector Observation Node Na
vector_observation
Recurrent Input Node Name
recurrent_in
Recurrent Output Node Nam
recurrent_out
Action Node Name
action
Graph Placeholders

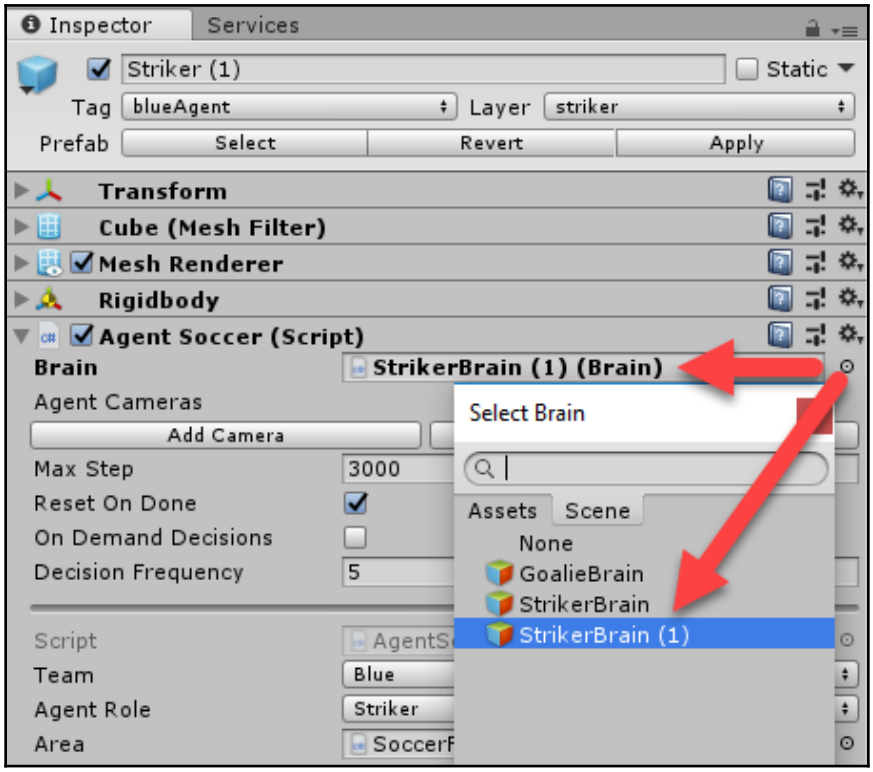
Inspector
Services
Striker (1)
Static
Tag
blueAgent
Layer
striker
Prefab
Select
Revert
Apply
Transform
Cube (Mesh Filter)
Mesh Renderer
Rigidbody
Agent Soccer (Script)
Brain
StrikerBrain (1) (Brain)
Agent Cameras
Add Camera
Max Step
3000
Reset On Done
On Demand Decisions
Decision Frequency
5
Script
Team
Blue
Agent Role
Striker
Area
Select Brain
Assets
Scene
None
GoalieBrain
StrikerBrain
StrikerBrain (1)

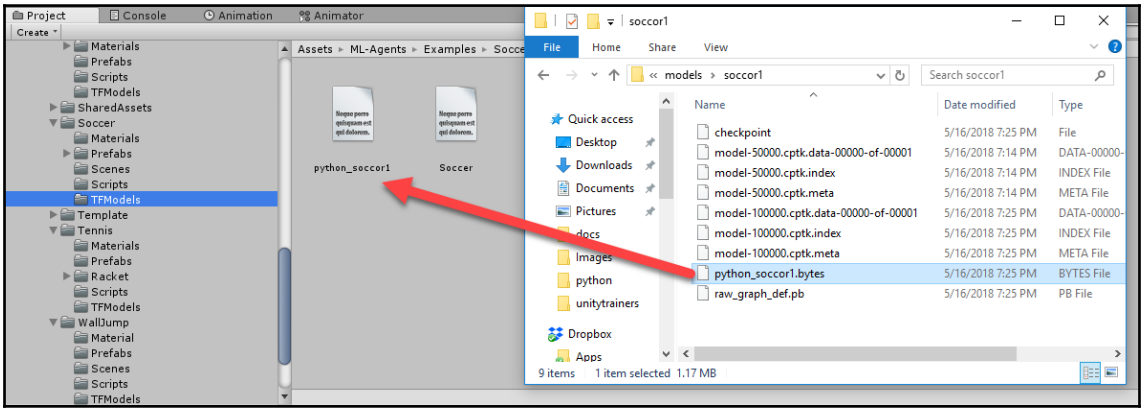
Project
Console
Animation
Animator
Create
Materials
Prefabs
Scripts
TFModels
SharedAssets
Soccer
Scenes
Template
Tennis
Racket
WallJump
Material
Assets ▸ ML-Agents ▸ Examples
python_soccer1
Soccer
soccer1
File
Home
Share
View
models ▸ soccer1
Search soccer1
Quick access
Desktop
Downloads
Documents
Pictures
docs
Images
python
unitytrainers
Dropbox
Apps
Name
Date modified
Type
checkpoint
model-50000.cptk.data-00000-of-00001
model-50000.cptk.index
model-50000.cptk.meta
model-100000.cptk.data-00000-of-00001
model-100000.cptk.index
model-100000.cptk.meta
python_soccer1.bytes
raw_graph_def.pb
9 items
1 item selected 1.17 MB

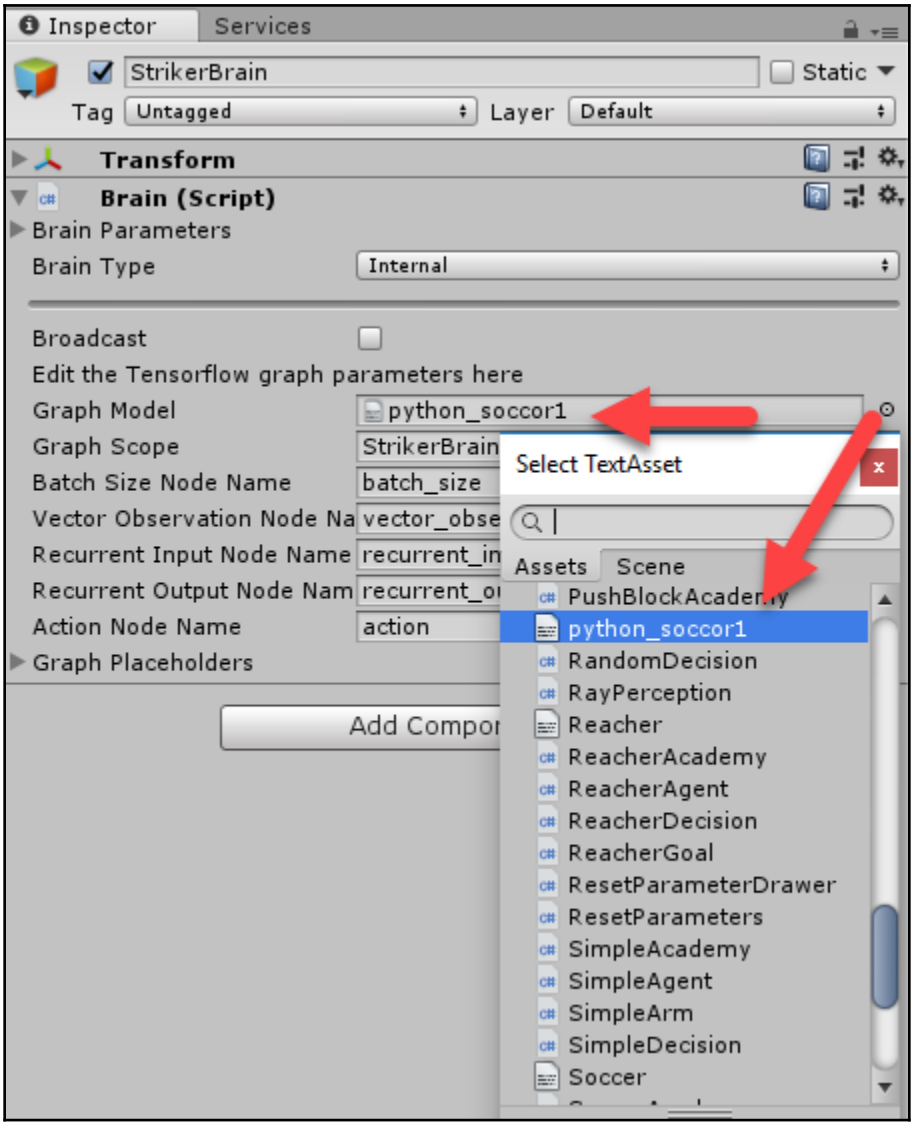
Inspector
Services
StrikerBrain
Static
Tag
Untagged
Layer
Default
Transform
Brain (Script)
Brain Parameters
Brain Type
Internal
Broadcast
Edit the Tensorflow graph parameters here
Graph Model
python_soccor1
Graph Scope
StrikerBrain
Batch Size Node Name
batch_size
Vector Observation Node Na
vector_obse
Recurrent Input Node Name
recurrent_in
Recurrent Output Node Nam
recurrent_ou
Action Node Name
action
Graph Placeholders
Add Compor
Select TextAsset
Assets
Scene
PushBlockAcademy
python_soccor1
RandomDecision
RayPerception
Reacher
ReacherAcademy
ReacherAgent
ReacherDecision
ReacherGoal
ResetParameterDrawer
ResetParameters
SimpleAcademy
SimpleAgent
SimpleArm
SimpleDecision
Soccer

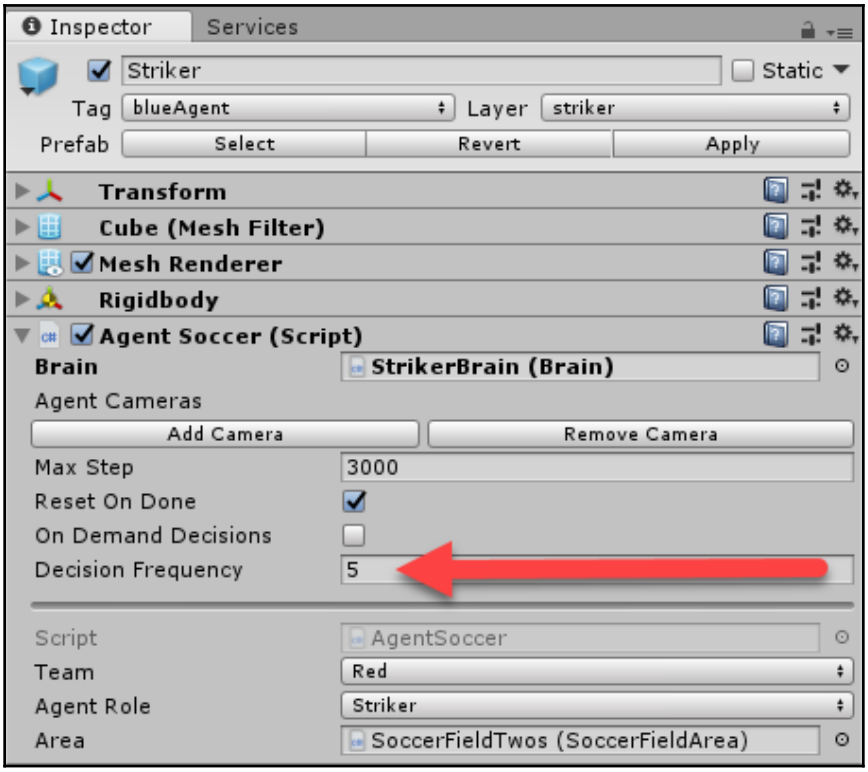
Inspector
Services
Striker
Static
Tag
blueAgent
Layer
striker
Prefab
Select
Revert
Apply
Transform
Cube (Mesh Filter)
Mesh Renderer
Rigidbody
Agent Soccer (Script)
Brain
StrikerBrain (Brain)
Agent Cameras
Add Camera
Remove Camera
Max Step
3000
Reset On Done
On Demand Decisions
Decision Frequency
5
Script
AgentSoccer
Team
Red
Agent Role
Striker
Area
SoccerFieldTwos (SoccerFieldArea)

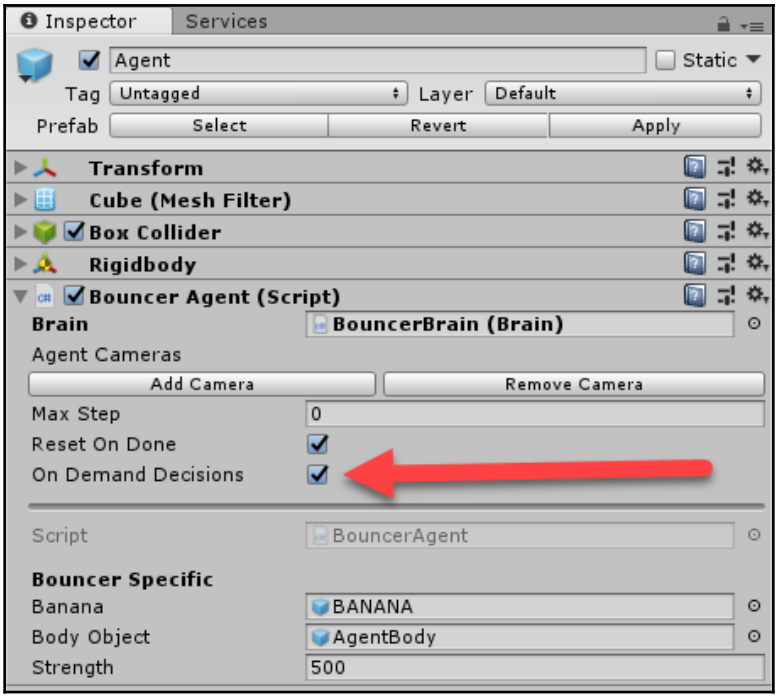
Inspector
Services
Agent
Static
Tag
Untagged
Layer
Default
Prefab
Select
Revert
Apply
Transform
Cube (Mesh Filter)
Box Collider
Rigidbody
Bouncer Agent (Script)
Brain
BouncerBrain (Brain)
Agent Cameras
Add Camera
Remove Camera
Max Step
0
Reset On Done
On Demand Decisions
Script
BouncerAgent
Bouncer Specific
Banana
BANANA
Body Object
AgentBody
Strength
500

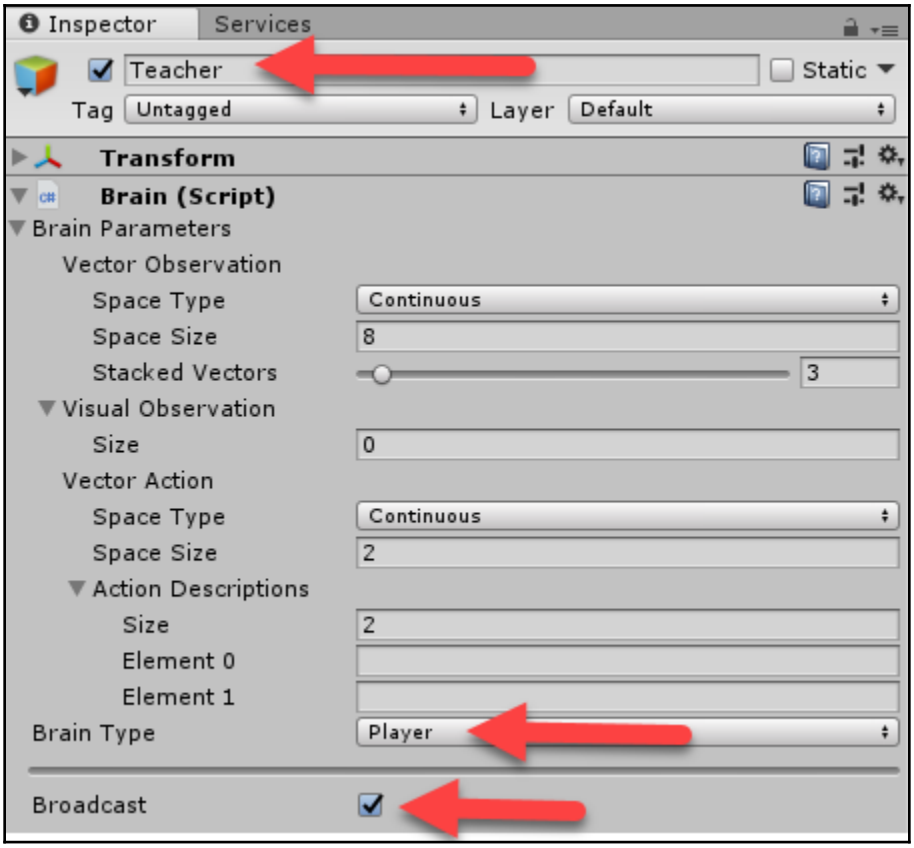
Inspector
Services
Teacher
Static
Tag
Untagged
Layer
Default
Transform
Brain (Script)
Brain Parameters
Vector Observation
Space Type
Continuous
Space Size
8
Stacked Vectors
3
Visual Observation
Size
0
Vector Action
Space Type
Continuous
Space Size
2
Action Descriptions
Size
2
Element 0
Element 1
Brain Type
Player
Broadcast

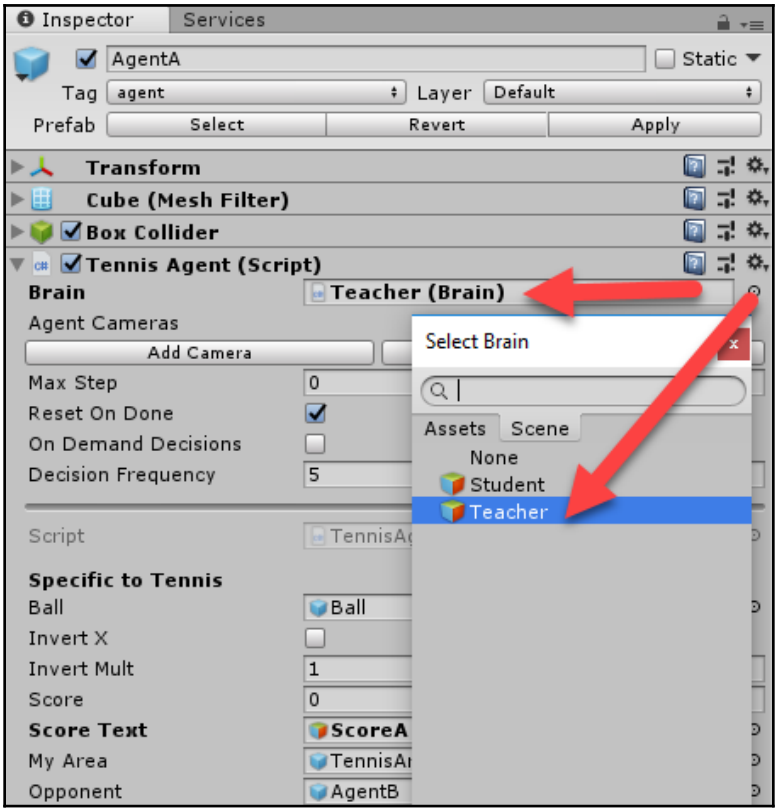
Inspector
Services
AgentA
Static
Tag
agent
Layer
Default
Prefab
Select
Revert
Apply
Transform
Cube (Mesh Filter)
Box Collider
Tennis Agent (Script)
Brain
Teacher (Brain)
Agent Cameras
Add Camera
Max Step
0
Reset On Done
On Demand Decisions
Decision Frequency
5
Script
Specific to Tennis
Ball
Ball
Invert X
Invert Mult
1
Score
0
Score Text
ScoreA
My Area
Opponent
AgentB
Select Brain
Assets
Scene
None
Student
Teacher

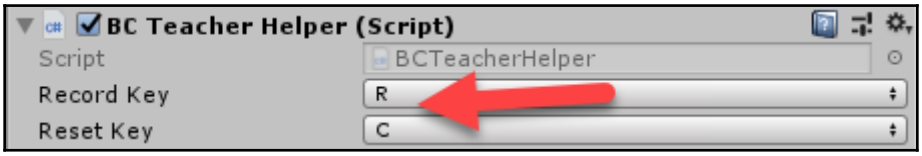
BC Teacher Helper (Script)
Script
BCTeacherHelper
Record Key
R
Reset Key
C

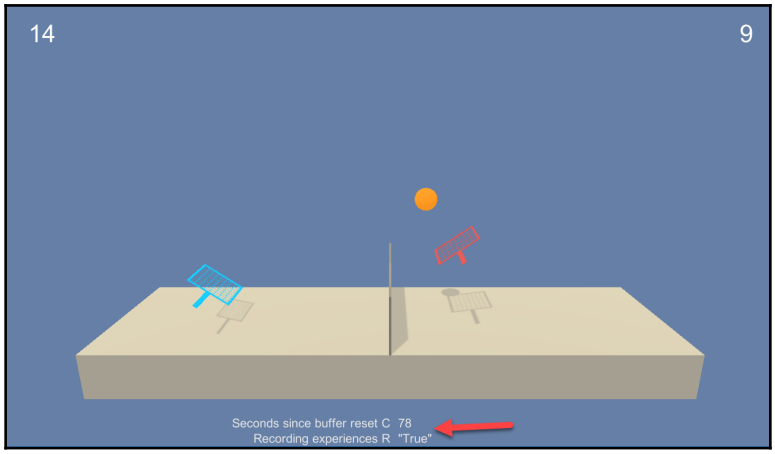
14
9
Seconds since buffer reset C 78
Recording experiences R "True"

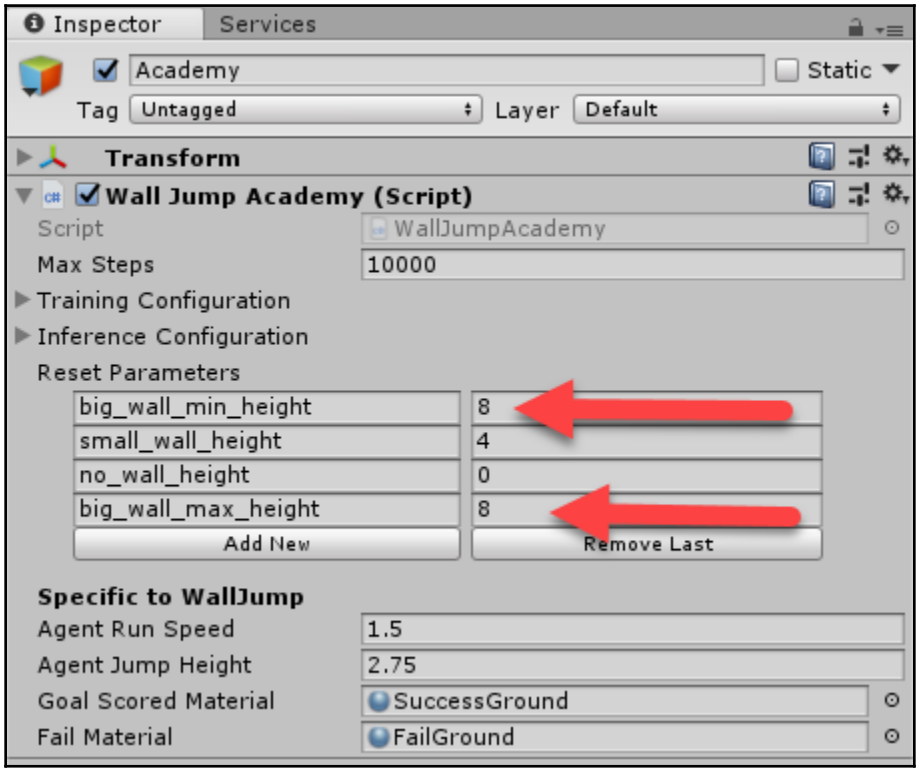
Inspector
Services
Academy
Static
Tag
Untagged
Layer
Default
Transform
Wall Jump Academy (Script)
Script
WallJumpAcademy
Max Steps
10000
Training Configuration
Inference Configuration
Reset Parameters
big_wall_min_height
8
small_wall_height
4
no_wall_height
0
big_wall_max_height
8
Add New
Remove Last
Specific to WallJump
Agent Run Speed
1.5
Agent Jump Height
2.75
Goal Scored Material
SuccessGround
Fail Material
FailGround

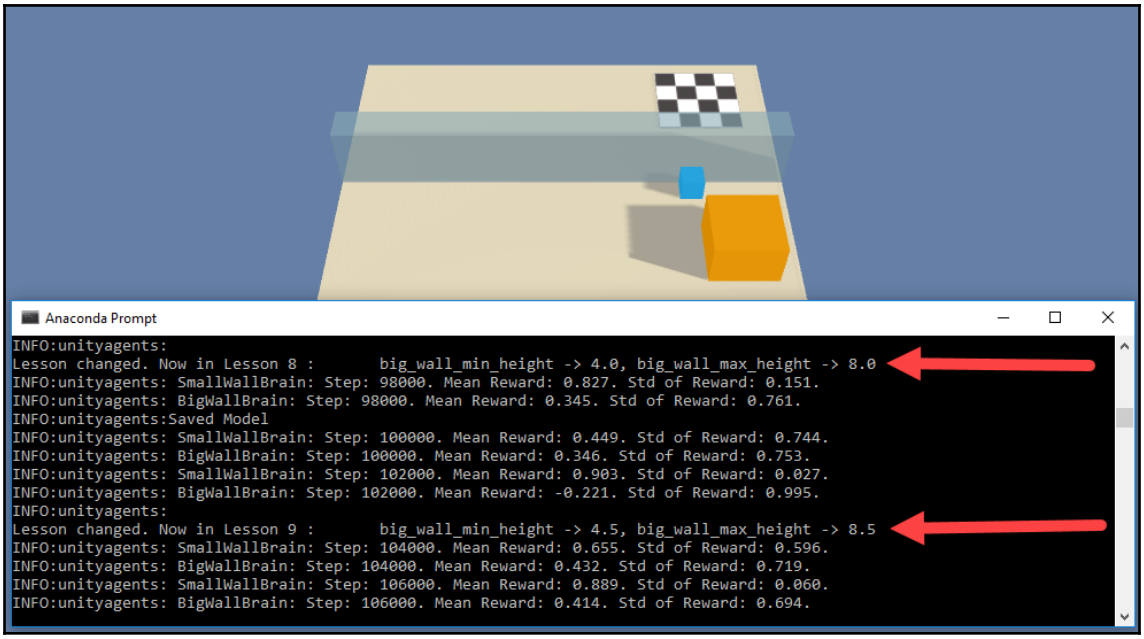
Anaconda Prompt
INFO:unityagents:
Lesson changed. Now in Lesson 8 : big_wall_min_height -> 4.0, big_wall_max_height -> 8.0
INFO:unityagents: SmallWallBrain: Step: 98000. Mean Reward: 0.827. Std of Reward: 0.151.
INFO:unityagents: BigWallBrain: Step: 98000. Mean Reward: 0.345. Std of Reward: 0.761.
INFO:unityagents:Saved Model
INFO:unityagents: SmallWallBrain: Step: 100000. Mean Reward: 0.449. Std of Reward: 0.744.
INFO:unityagents: BigWallBrain: Step: 100000. Mean Reward: 0.346. Std of Reward: 0.753.
INFO:unityagents: SmallWallBrain: Step: 102000. Mean Reward: 0.903. Std of Reward: 0.027.
INFO:unityagents: BigWallBrain: Step: 102000. Mean Reward: -0.221. Std of Reward: 0.995.
INFO:unityagents:
Lesson changed. Now in Lesson 9 : big_wall_min_height -> 4.5, big_wall_max_height -> 8.5
INFO:unityagents: SmallWallBrain: Step: 104000. Mean Reward: 0.655. Std of Reward: 0.596.
INFO:unityagents: BigWallBrain: Step: 104000. Mean Reward: 0.432. Std of Reward: 0.719.
INFO:unityagents: SmallWallBrain: Step: 106000. Mean Reward: 0.889. Std of Reward: 0.060.
INFO:unityagents: BigWallBrain: Step: 106000. Mean Reward: 0.414. Std of Reward: 0.694.

Chapter 6: Terrarium Revisited – A Multi-Agent Ecosystem

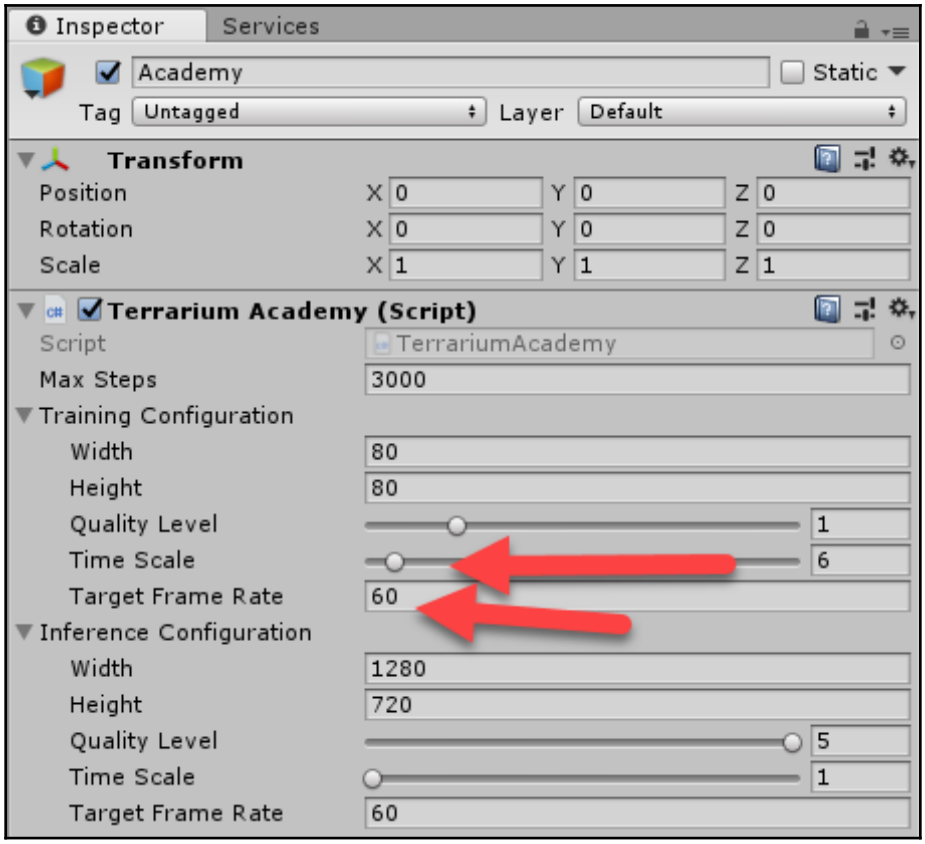
Inspector
Services
Academy
Static
Tag
Untagged
Layer
Default
Transform
Position
Rotation
Scale
Terrarium Academy (Script)
Script
TerrariumAcademy
Max Steps
3000
Training Configuration
Width
80
Height
80
Quality Level
1
Time Scale
6
Target Frame Rate
60
Inference Configuration
Width
1280
Height
720
Quality Level
5
Time Scale
1
Target Frame Rate
60


Filters
Free Low Polygon_Animal
3D Models/Characters/Animals
GLOOMY STUDIO
★★★★★ (15)
FREE
Import
3 type cute little poly mobile friendly animals and birds!
+ Mobile friendly!
+ Low Poly (100~250 tris) !
+ 512 x 512 textures !
Animal Pack
GLOOMY STUDIO
Home
3D Models
Characters
Animals
Birds
Fish
Insects
Mammals
Reptiles
Other
Creatures
Humanoids
Robots
Toons
UMA Avatars
Other
Environments
Props
Vegetation
Vehicles
Other
Animation
Audio
Complete Projects

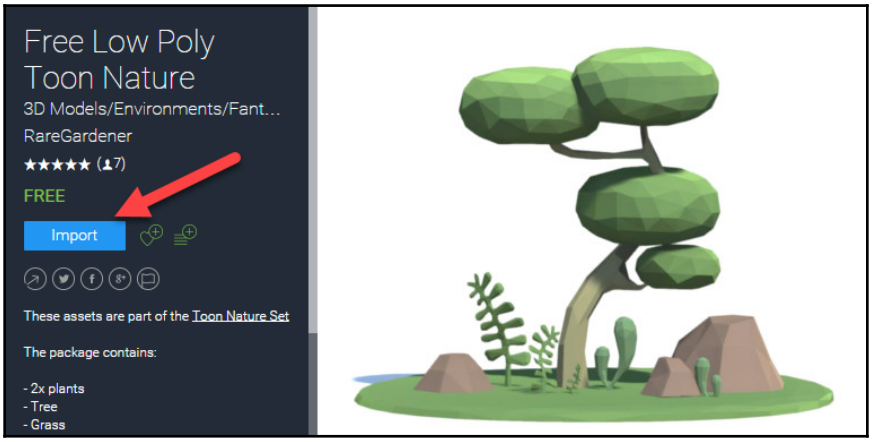
Free Low Poly Toon Nature
3D Models/Environments/Fant...
RareGardener
★★★★★ (17)
FREE
Import
These assets are part of the Toon Nature Set
The package contains:
- 2x plants
- Tree
- Grass

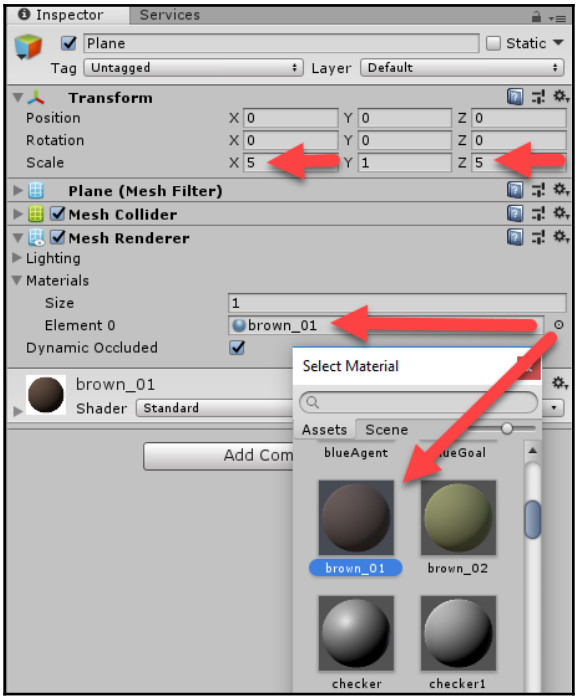
Inspector
Services
Plane
Static
Tag
Untagged
Layer
Default
Transform
Position
Rotation
Scale
Plane (Mesh Filter)
Mesh Collider
Mesh Renderer
Lighting
Materials
Size
Element 0
brown_01
Dynamic Occluded
brown_01
Shader
Standard
Add Com
Select Material
Assets
Scene
blueAgent
brown_01
brown_02
checker
checker1

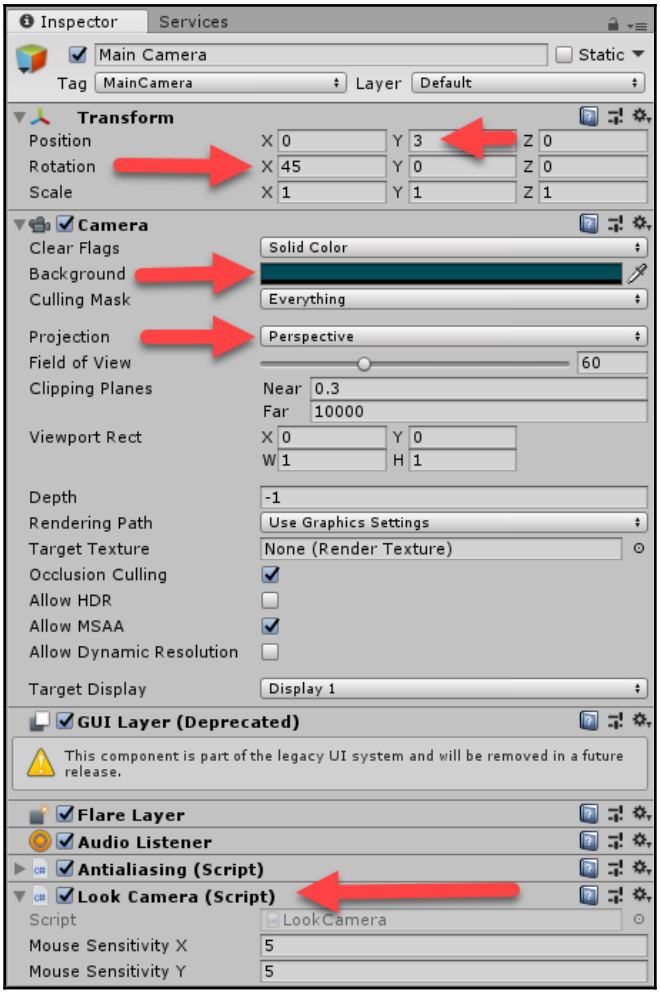
Inspector
Services
Main Camera
Static
Tag
MainCamera
Layer
Default
Transform
Position
X 0
Y 3
Z 0
Rotation
X 45
Y 0
Z 0
Scale
X 1
Y 1
Z 1
Camera
Clear Flags
Solid Color
Background
Culling Mask
Everything
Projection
Perspective
Field of View
60
Clipping Planes
Near 0.3
Far 10000
Viewport Rect
X 0
Y 0
W 1
H 1
Depth
-1
Rendering Path
Use Graphics Settings
Target Texture
None (Render Texture)
Occlusion Culling
Allow HDR
Allow MSAA
Allow Dynamic Resolution
Target Display
Display 1
GUI Layer (Deprecated)
This component is part of the legacy UI system and will be removed in a future release.
Flare Layer
Audio Listener
Antialiasing (Script)
Look Camera (Script)
Script
LookCamera
Mouse Sensitivity X
5
Mouse Sensitivity Y
5

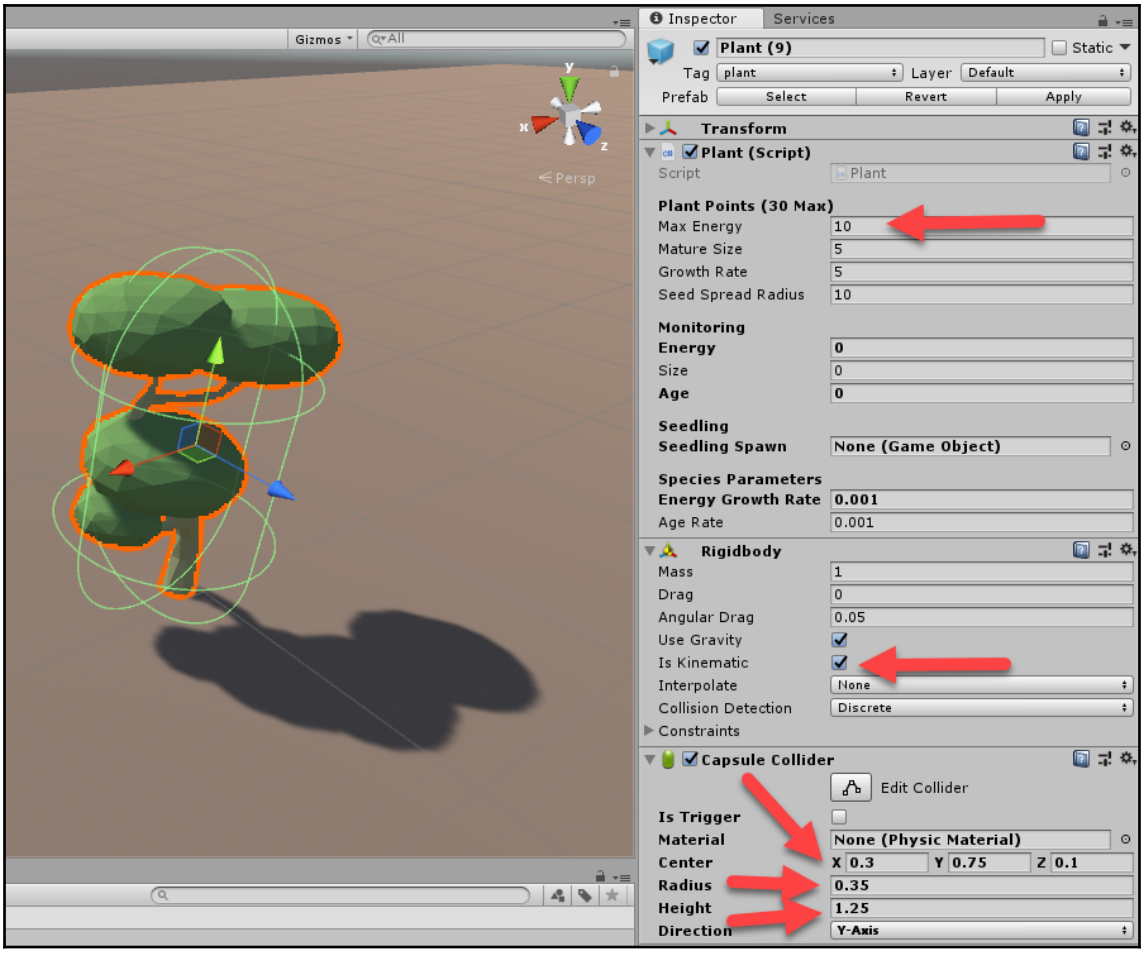
Gizmos
All
Persp
Inspector
Services
Plant (9)
Static
Tag
plant
Layer
Default
Prefab
Select
Revert
Apply
Transform
Plant (Script)
Script
Plant
Plant Points (30 Max)
Max Energy
10
Mature Size
5
Growth Rate
5
Seed Spread Radius
10
Monitoring
Energy
0
Size
0
Age
0
Seedling
Seedling Spawn
None (Game Object)
Species Parameters
Energy Growth Rate
0.001
Age Rate
0.001
Rigidbody
Mass
1
Drag
0
Angular Drag
0.05
Use Gravity
Is Kinematic
Interpolate
None
Collision Detection
Discrete
Constraints
Capsule Collider
Edit Collider
Is Trigger
Material
None (Physic Material)
Center
X 0.3
Y 0.75
Z 0.1
Radius
0.35
Height
1.25
Direction
Y-Axis

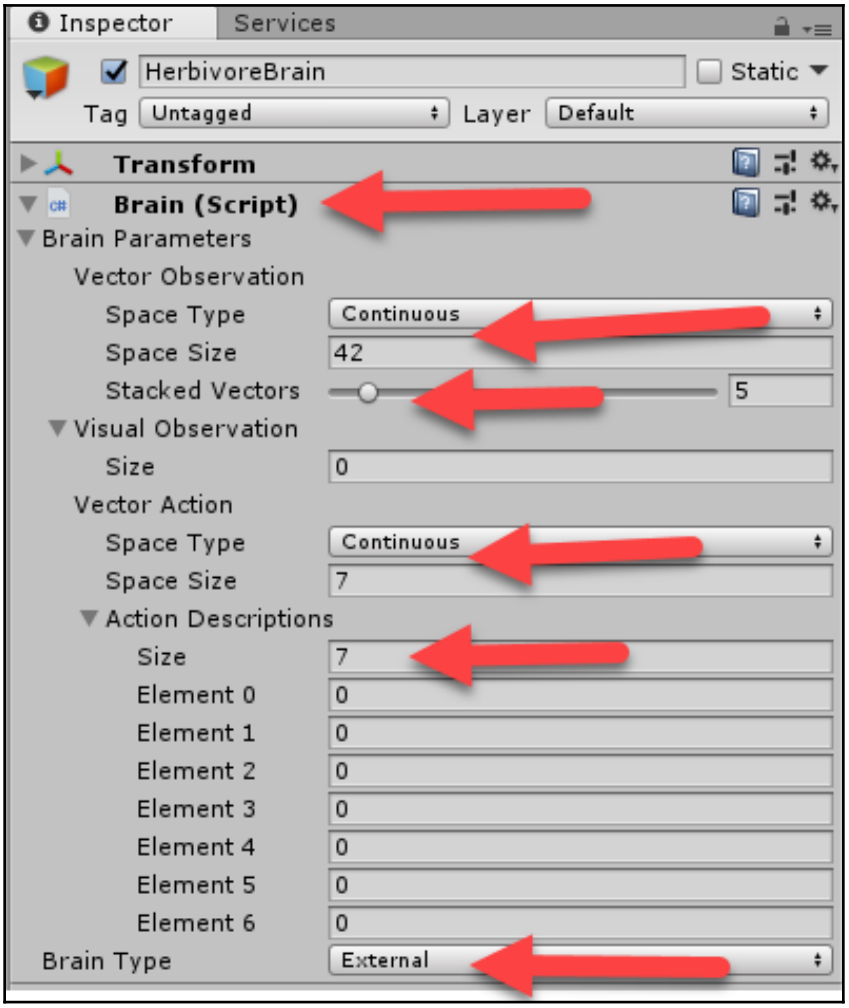
Inspector
Services
HerbivoreBrain
Static
Tag
Untagged
Layer
Default
Transform
Brain (Script)
Brain Parameters
Vector Observation
Space Type
Continuous
Space Size
42
Stacked Vectors
5
Visual Observation
Size
0
Vector Action
Space Type
Continuous
Space Size
7
Action Descriptions
Size
7
Element 0
0
Element 1
0
Element 2
0
Element 3
0
Element 4
0
Element 5
0
Element 6
0
Brain Type
External

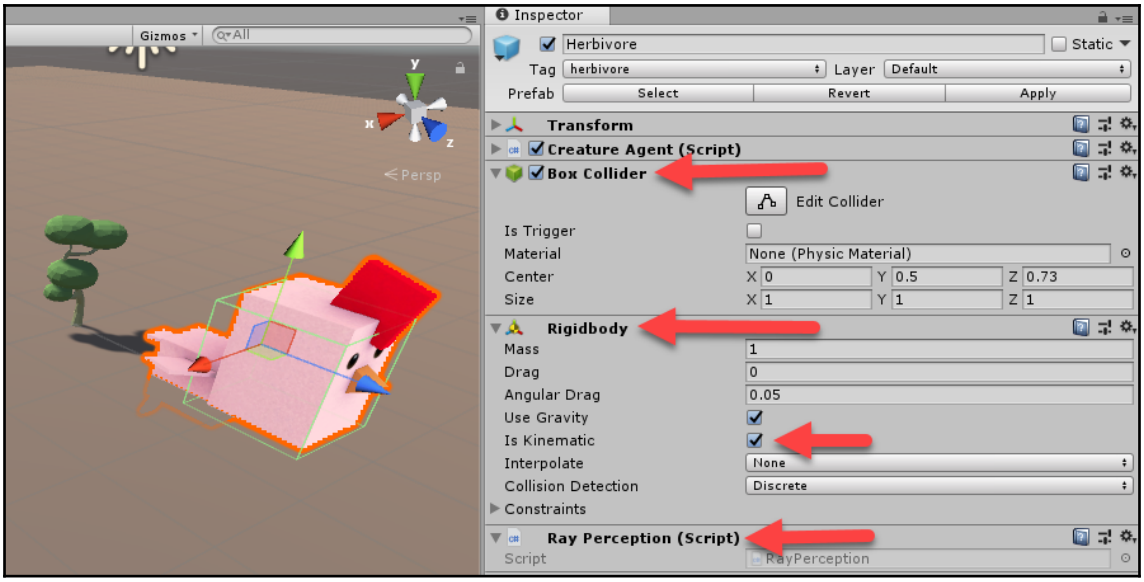
Gizmos
All
Persp
Inspector
Herbivore
Static
Tag
herbivore
Layer
Default
Prefab
Select
Revert
Apply
Transform
Creature Agent (Script)
Box Collider
Edit Collider
Is Trigger
Material
None (Physic Material)
Center
X 0
Y 0.5
Z 0.73
Size
X 1
Y 1
Z 1
Rigidbody
Mass
1
Drag
0
Angular Drag
0.05
Use Gravity
Is Kinematic
Interpolate
None
Collision Detection
Discrete
Constraints
Ray Perception (Script)
Script
RayPerception

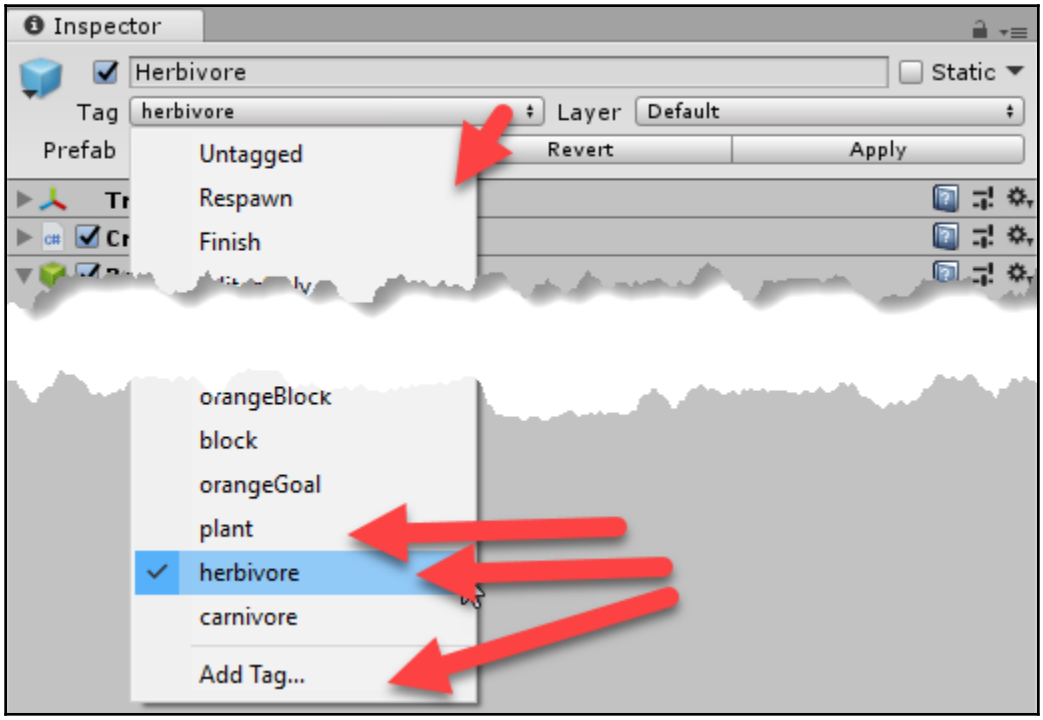
Inspector
Herbivore
Static
Tag
herbivore
Layer
Default
Prefab
Revert
Apply
Untagged
Respawn
Finish
orangeBlock
block
orangeGoal
plant
herbivore
carnivore
Add Tag...

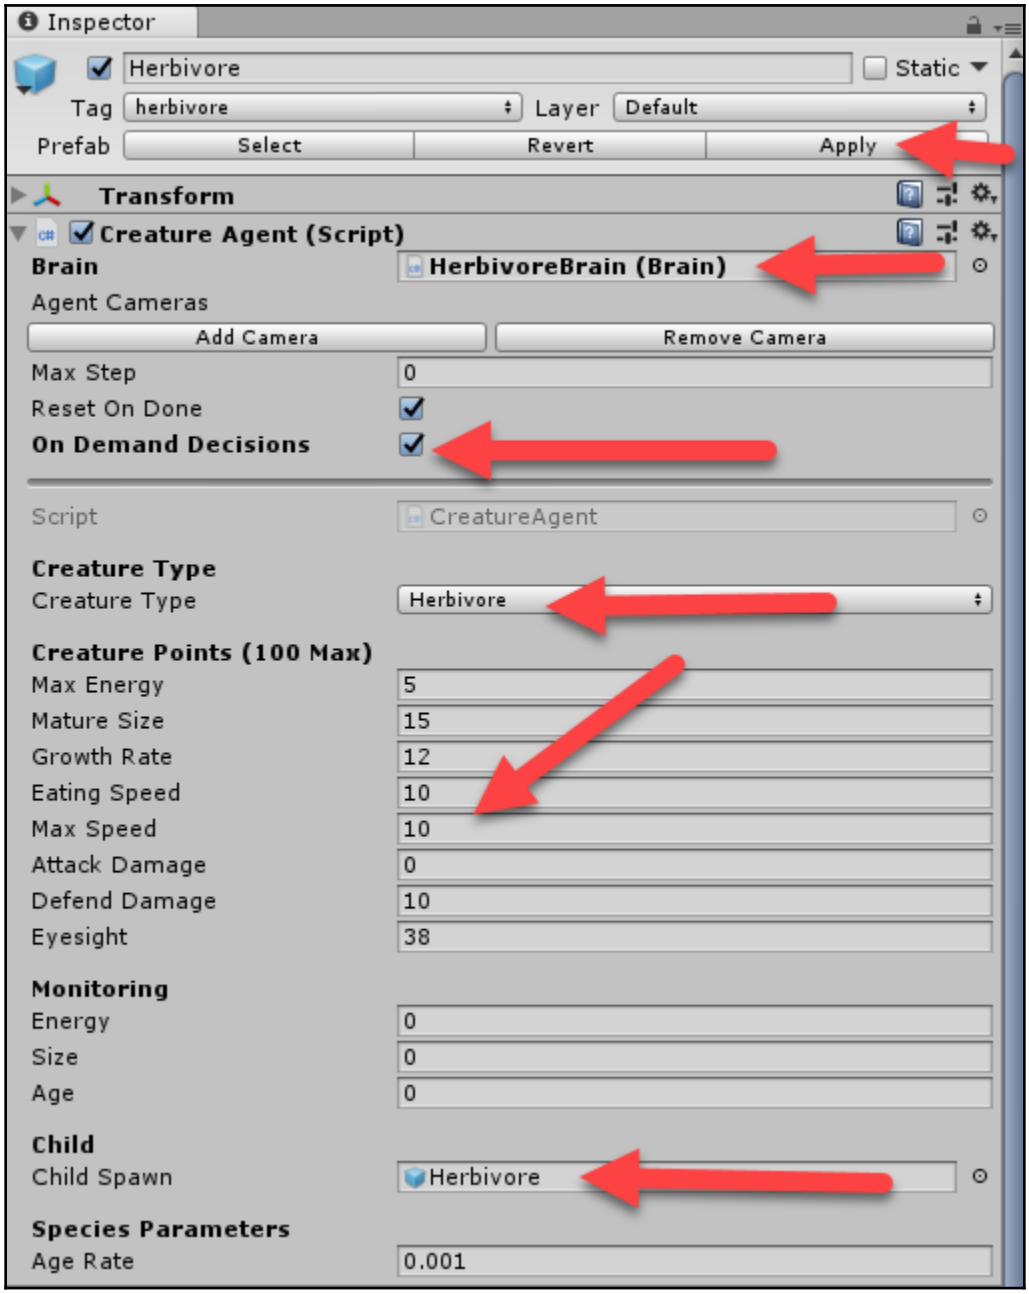
Inspector
Herbivore
Static
Tag
herbivore
Layer
Default
Prefab
Select
Revert
Apply
Transform
Creature Agent (Script)
Brain
HerbivoreBrain (Brain)
Agent Cameras
Add Camera
Remove Camera
Max Step
0
Reset On Done
On Demand Decisions
Script
CreatureAgent
Creature Type
Creature Type
Herbivore
Creature Points (100 Max)
Max Energy
5
Mature Size
15
Growth Rate
12
Eating Speed
10
Max Speed
10
Attack Damage
0
Defend Damage
10
Eyesight
38
Monitoring
Energy
0
Size
0
Age
0
Child
Child Spawn
Herbivore
Species Parameters
Age Rate
0.001


Age
Energy
Size
Action "Idle"
Age
Energy
Size
Action "Idle"
Age
Energy
Size
Action "Idle"
Age
Energy
Size
Action "Idle"
Age
Energy
Size
Action "Idle"

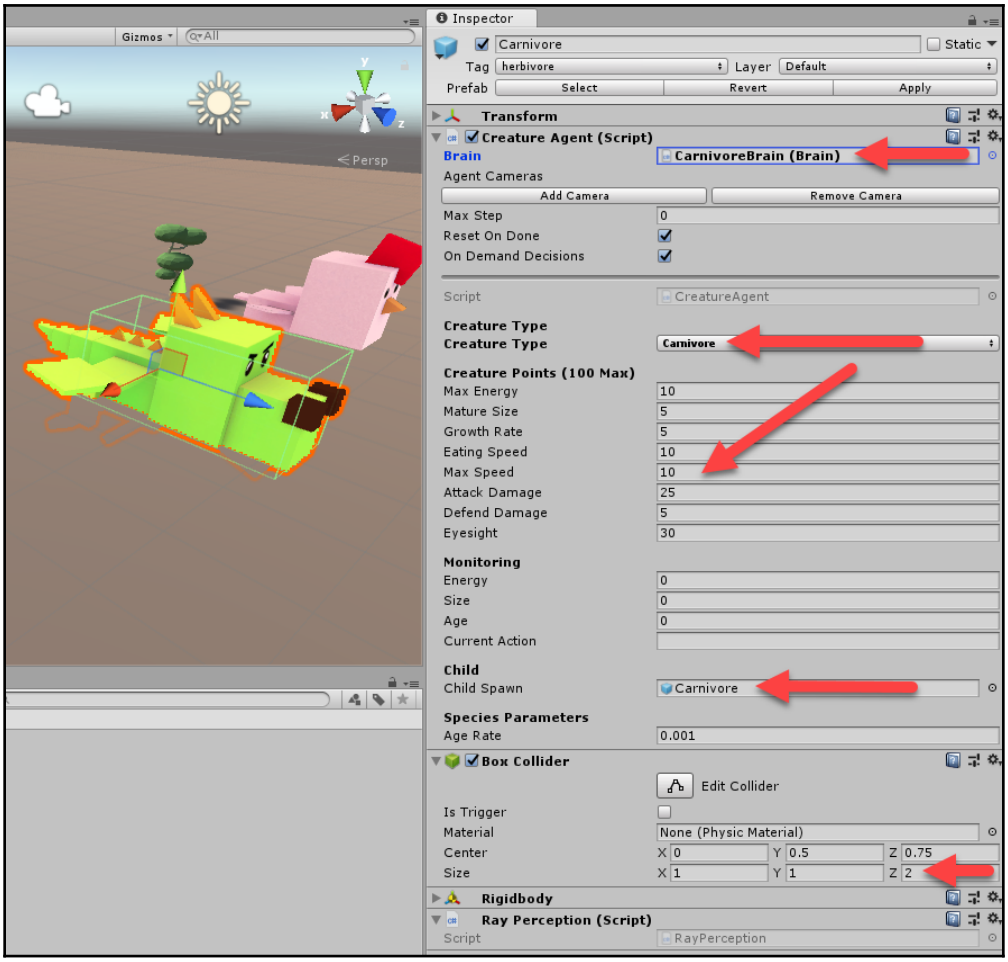
Gizmos
All
Persp
Inspector
Carnivore
Static
Tag herbivore
Layer Default
Prefab
Select
Revert
Apply
Transform
Creature Agent (Script)
Brain
CarnivoreBrain (Brain)
Agent Cameras
Add Camera
Remove Camera
Max Step 0
Reset On Done
On Demand Decisions
Script CreatureAgent
Creature Type
Creature Type Carnivore
Creature Points (100 Max)
Max Energy 10
Mature Size 5
Growth Rate 5
Eating Speed 10
Max Speed 10
Attack Damage 25
Defend Damage 5
Eyesight 30
Monitoring
Energy 0
Size 0
Age 0
Current Action
Child
Child Spawn Carnivore
Species Parameters
Age Rate 0.001
Box Collider
Edit Collider
Is Trigger
Material None (Physic Material)
Center X 0 Y 0.5 Z 0.75
Size X 1 Y 1 Z 2
Rigidbody
Ray Perception (Script)
Script RayPerception

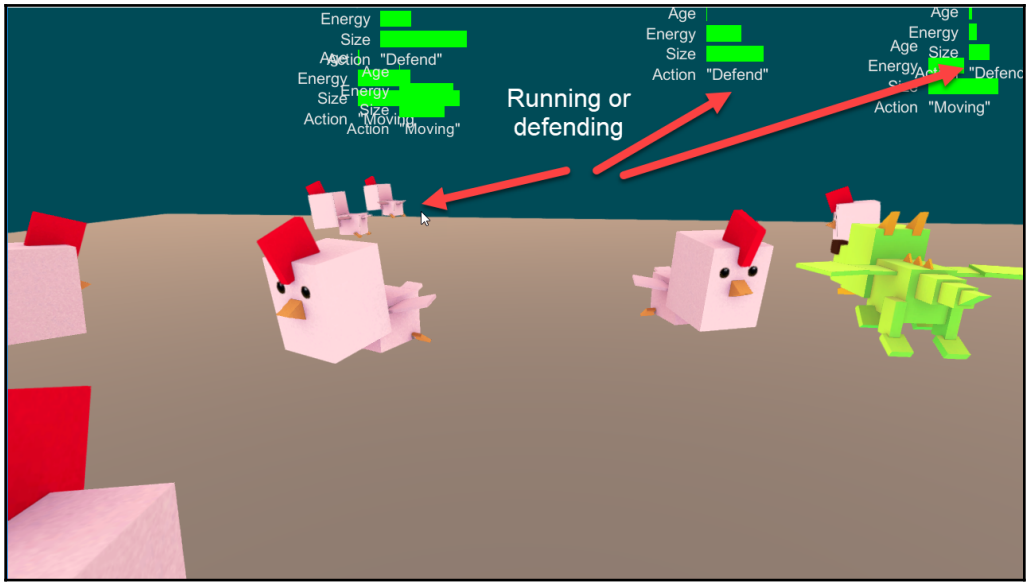
Energy
Size
Action "Defend"
Energy
Size
Action "Moving"
Running or
defending
Age
Energy
Size
Action "Defend"
Age
Energy
Size
Action "Moving"