
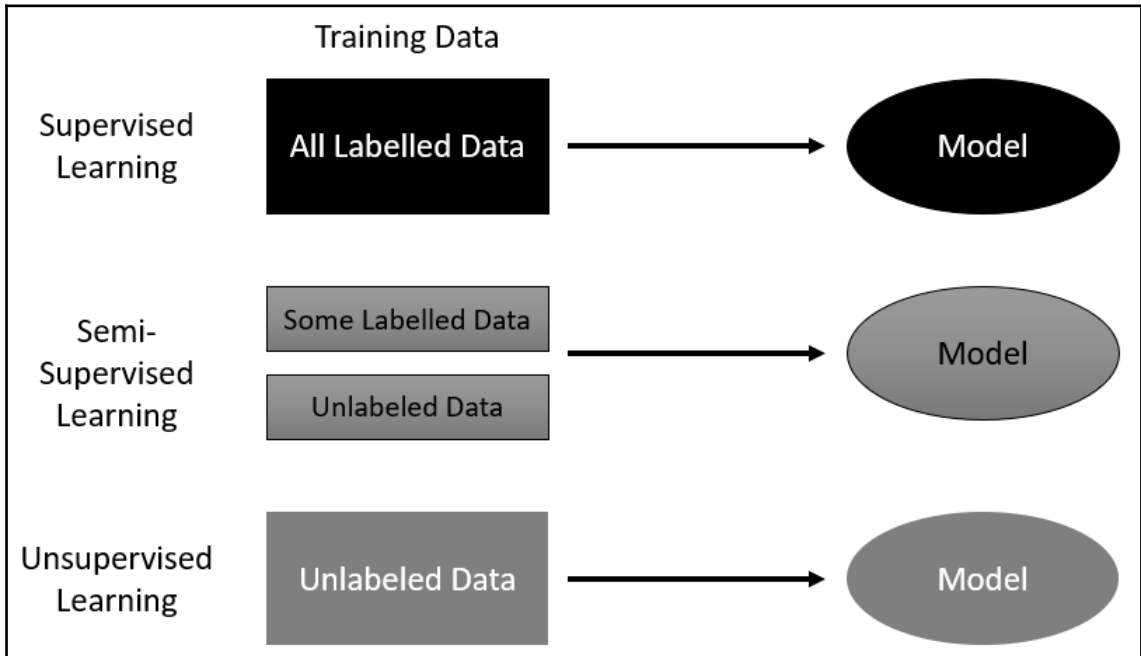
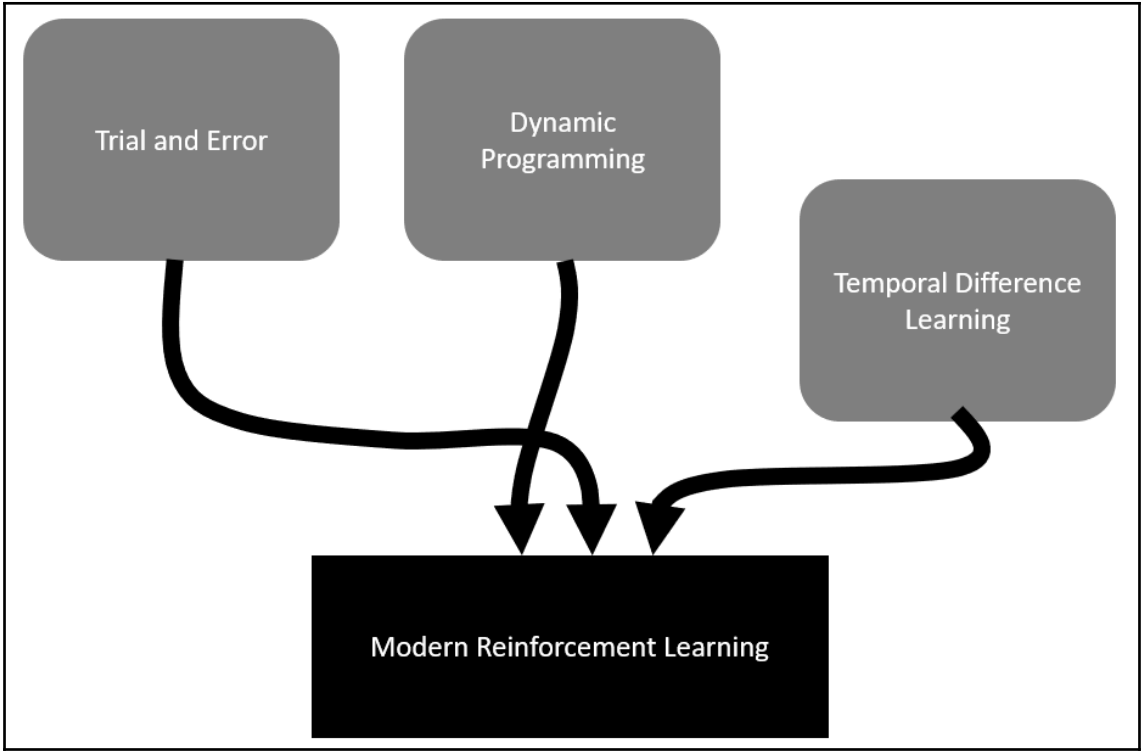

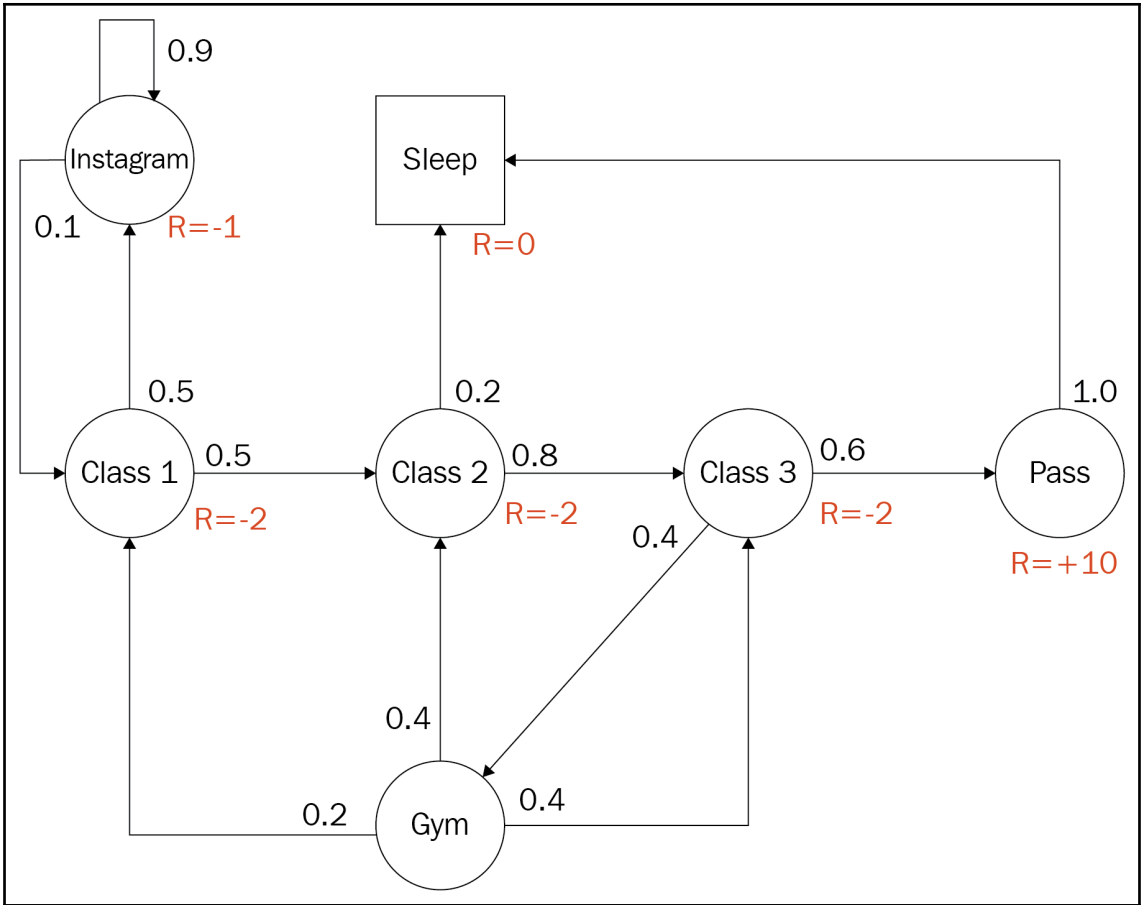


# Chapter 1: Understanding Rewards-Based Learning





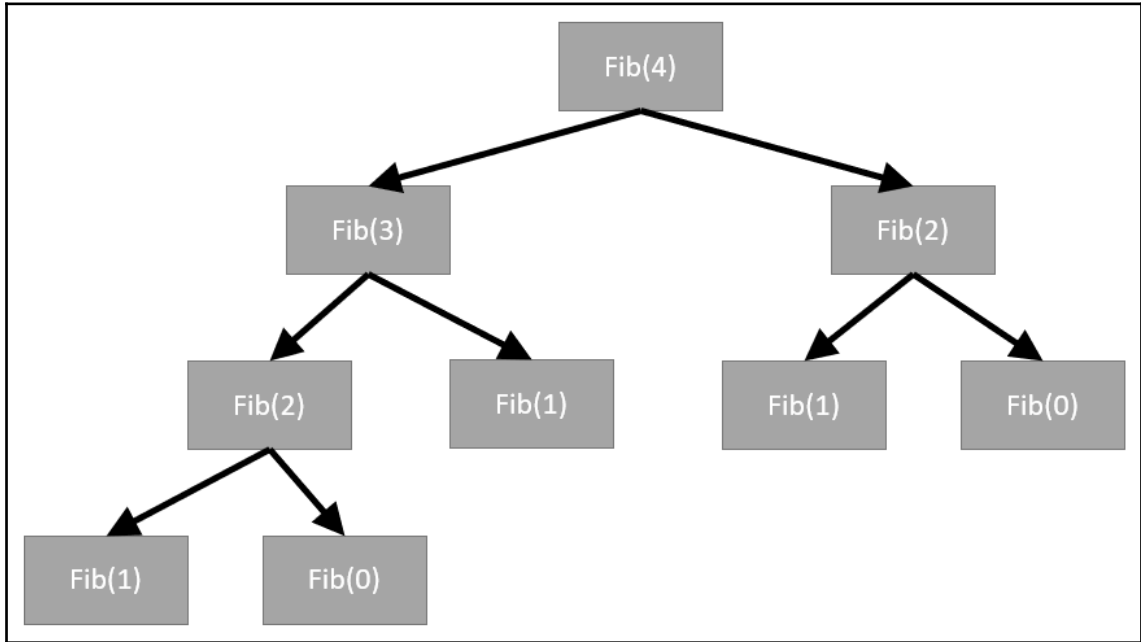


```
C:\Users\Micheal\Anaconda3\python.exe
[0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0]
[0.7941088679053511, 0.37290670858355, 0.17298296564654017, 0.3062897555000004,
0.4888187886688896, 0.0814697981114816, -0.3587852317595]
Press any key to continue . . .
```

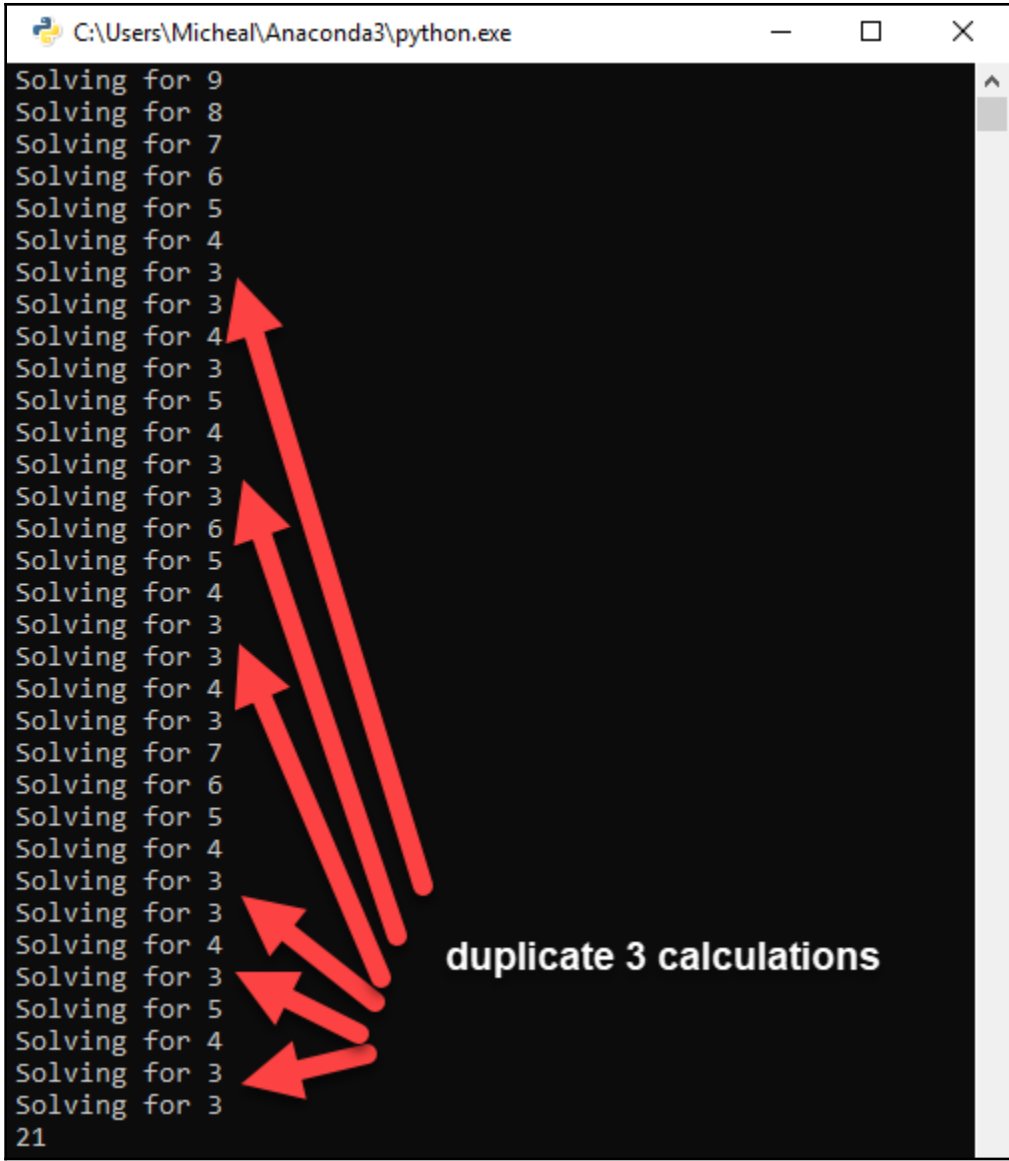
```
C:\Users\Micheal\Anaconda3\python.exe
[5.0, 5.0, 5.0, 5.0, 5.0, 5.0, 5.0]
[1.0000000000000009, 0.9923854510918054, 0.9204542494255955, 0.9923854510918054, 0.9899692772647101, 0.9171819032836618, 0.8980262055809499]
Press any key to continue . . .
```



## Chapter 2: Dynamic Programming and the Bellman Equation



```
C:\Users\Micheal\Anaconda3\python.exe
Solving for 9
Solving for 8
Solving for 7
Solving for 6
Solving for 5
Solving for 4
Solving for 3
Solving for 3
Solving for 4
Solving for 3
Solving for 5
Solving for 4
Solving for 3
Solving for 3
Solving for 6
Solving for 5
Solving for 4
Solving for 3
Solving for 3
Solving for 4
Solving for 3
Solving for 7
Solving for 6
Solving for 5
Solving for 4
Solving for 3
Solving for 3
Solving for 4
Solving for 3
Solving for 5
Solving for 4
Solving for 4
Solving for 3
Solving for 3
21
```


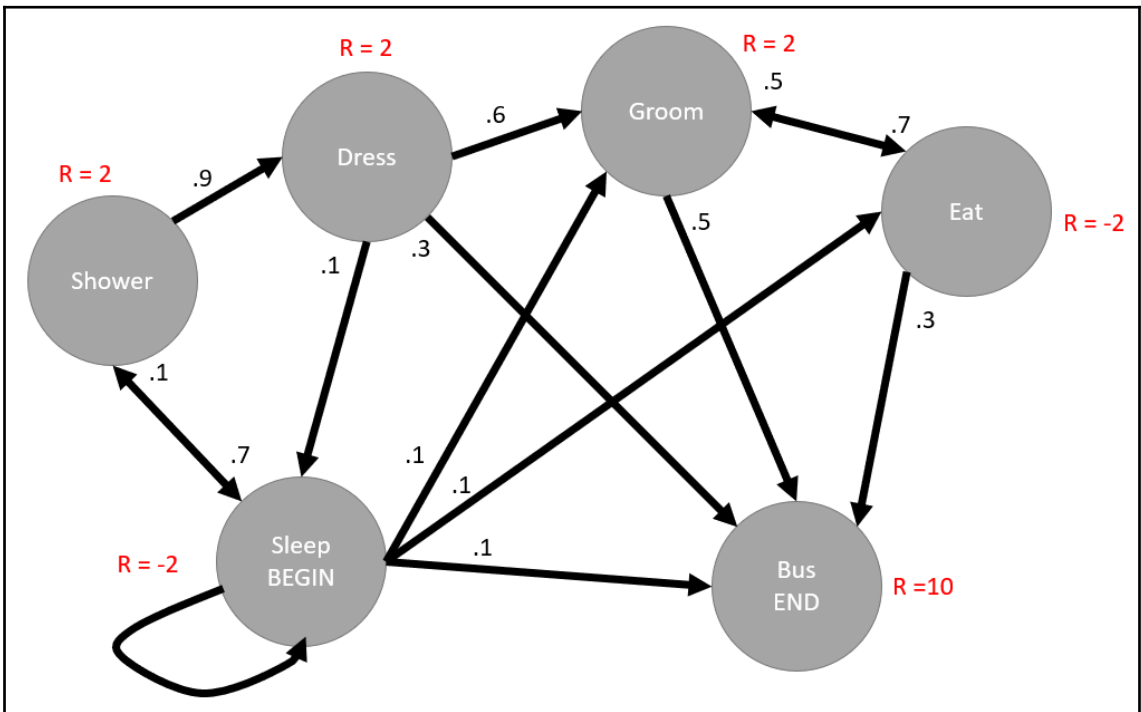


**duplicate 3 calculations**



```

C:\Users\Micheal\Anaconda3\python.exe
Solving for 9
Solving for 8
Solving for 7
Solving for 6
Solving for 5
Solving for 4
Solving for 3
21
Press any key to continue . . .

```



## Backup diagrams:



C:\Users\Micheal\Anaconda3\python.exe

(Up)

SFFF

FHFH

FFFH

HFFG

```
C:\Users\Micheal\Anaconda3\python.exe
[0. 0. 0. 0. 0. 0.
 0. 0. 0. 0. 0.
 0. 0. 0.25694444 0. ]

C:\Users\Micheal\Anaconda3\python.exe
[0.00227497 0.00368463 0.01073147 0.00400858 0.00465974 0.
 0.02684317 0. 0.01505108 0.05126114 0.09704873 0.
 0. 0.09365075 0.31028739 0. ]

C:\Users\Micheal\Anaconda3\python.exe
[0.00353151 0.00472702 0.01166752 0.00441536 0.00531099 0.
 0.02723946 0. 0.0154987 0.05165132 0.09732289 0.
 0. 0.09383682 0.31041555 0. ]
[[0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]] [0.00353151 0.00472702 0.01166752 0.00441536 0.00531099 0.
 0.02723946 0. 0.0154987 0.05165132 0.09732289 0.
 0. 0.09383682 0.31041555 0. ]
Press any key to continue . . .
```

```
C:\Users\Micheal\Anaconda3\python.exe
[0.01070127 0.01147851 0.0178007 0.00593357 0.02062529 0.
0.03804234 0. 0.06308208 0.12130938 0.13436865 0.
0. 0.16647742 0.37812288 0. ]
[[0. 1. 0. 0. ]
[0. 0. 0. 1. ]
[0.25 0.25 0.25 0.25]
[0. 0. 0. 1. ]
[0.25 0.25 0.25 0.25]
[0.25 0.25 0.25 0.25]
[0.25 0.25 0.25 0.25]
[0.25 0.25 0.25 0.25]
[0. 0. 0. 1. ]
[0. 1. 0. 0. ]
[0.25 0.25 0.25 0.25]
[0.25 0.25 0.25 0.25]
[0.25 0.25 0.25 0.25]
[0. 0. 1. 0. ]
[0. 0. 1. 0. ]
[0.25 0.25 0.25 0.25]] [0.01070127 0.01147851 0.0178007 0.00593357 0.02062529 0.
0.03804234 0. 0.06308208 0.12130938 0.13436865 0.
0. 0.16647742 0.37812288 0. ]
Press any key to continue . . .
```



```
C:\Users\Micheal\Anaconda3\python.exe
[[1. 0. 0. 0.]
[0. 0. 0. 1.]
[0. 0. 0. 1.]
[0. 0. 0. 1.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[0. 0. 0. 1.]
[0. 1. 0. 0.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[1. 0. 0. 0.]
[0. 0. 1. 0.]
[0. 1. 0. 0.]
[1. 0. 0. 0.]] [0.82352939 0.82352939 0.82352939 0.82352938 0.8235294 0.
0.52941175 0. 0.8235294 0.8235294 0.76470587 0.
0. 0.88235293 0.94117647 0. ]
Press any key to continue . . .
```


```
C:\Users\Micheal\Anaconda3\python.exe
[0.01402847 0.01350705 0.01987033 0.00662344 0.02338524 0.
0.04180654 0.          0.06957365 0.13315616 0.14735612 0.
0.          0.18253931 0.41446192 0.          ]
[[1.  0.  0.  0.  ]
 [0.  0.  0.  1.  ]
 [0.25 0.25 0.25 0.25]
 [0.  0.  0.  1.  ]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.  0.  0.  1.  ]
 [0.  1.  0.  0.  ]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.25 0.25 0.25 0.25]
 [0.  0.  1.  0.  ]
 [0.  1.  0.  0.  ]
 [0.25 0.25 0.25 0.25]] [0.01402847 0.01350705 0.01987033 0.00662344 0.02338524 0.
0.04180654 0.          0.06957365 0.13315616 0.14735612 0.
0.          0.18253931 0.41446192 0.          ]
767 ← policy iteration
[[1.  0.  0.  0.]
 [0.  0.  0.  1.]
 [0.  0.  0.  1.]
 [0.  0.  0.  1.]
 [1.  0.  0.  0.]
 [1.  0.  0.  0.]
 [1.  0.  0.  0.]
 [1.  0.  0.  0.]
 [0.  0.  0.  1.]
 [0.  1.  0.  0.]
 [1.  0.  0.  0.]
 [1.  0.  0.  0.]
 [1.  0.  0.  0.]
 [0.  0.  1.  0.]
 [0.  1.  0.  0.]
 [1.  0.  0.  0.]] [0.82352939 0.82352939 0.82352939 0.82352938 0.8235294 0.
0.52941175 0.          0.8235294 0.8235294 0.76470587 0.
0.          0.88235293 0.94117647 0.          ]
851 ← value iteration
Press any key to continue . . .
```

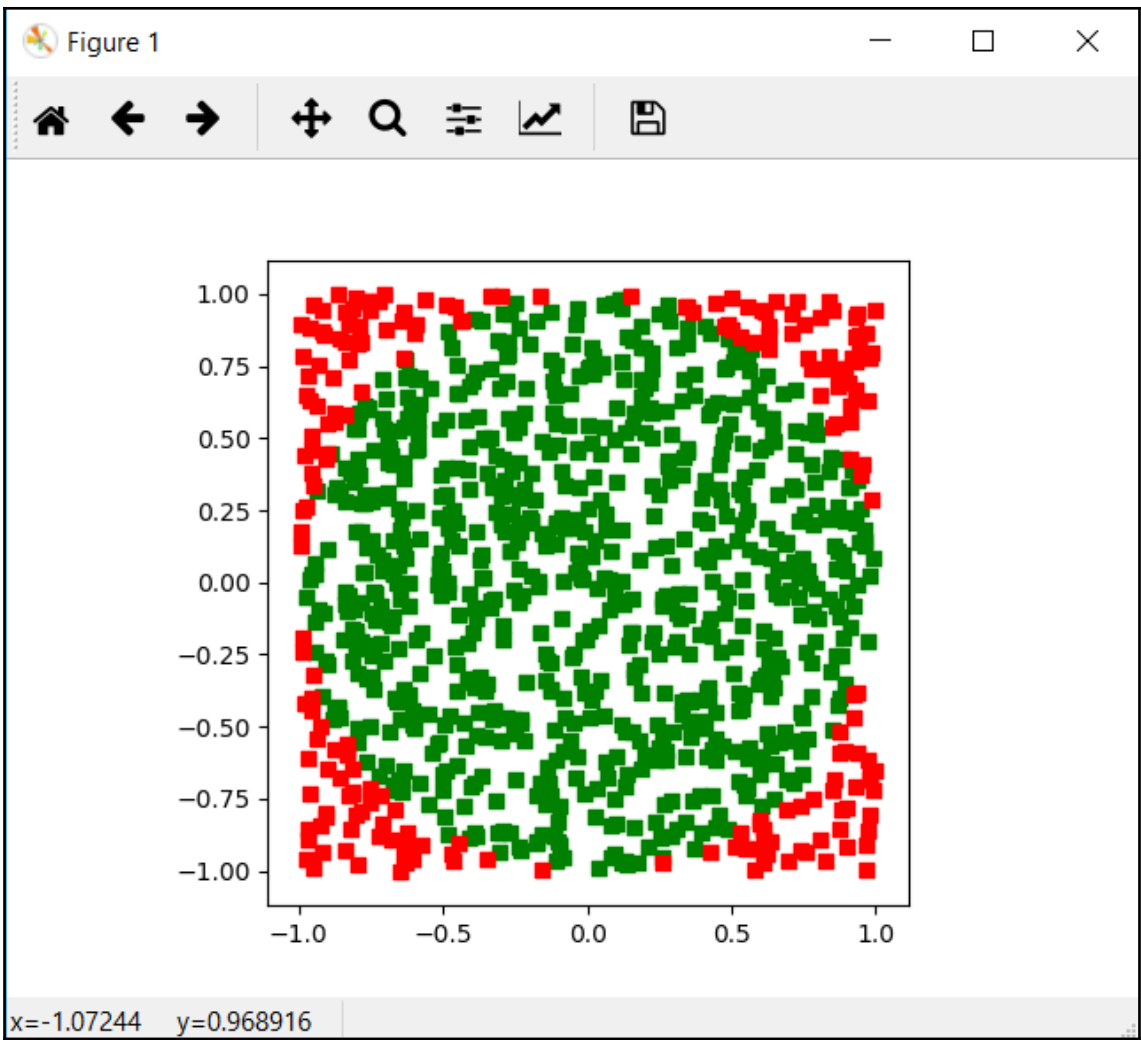
# Chapter 3: Monte Carlo Methods

S	F	F	F
F	H	F	H
F	F	F	H
H	F	F	G



$2 \times 2$  units = 4 square units





```
C:\Users\Micheal\Anaconda3\python.exe
SFFFFFFF
FFFFFFF
FFFHFFF
FFFFHFF
FFFHFFF
FHHFFFH
FHHFFFH
FFFHFFG
  (Up)
SFFFFFFF
FFFFFFF
FFFHFFF
FFFFHFF
FFFHFFF
FHHFFFH
FHHFFFH
FFFHFFG
  (Down)
SFFFFFFF
FFFFFFFF
FFFHFFF
FFFFHFF
FFFHFFF
FHHFFFH
FHHFFFH
FFFHFFG
```

```
C:\Users\Micheal\Anaconda3\python.exe
Test policy for episode 25000 wins % = 0.32
Test policy for episode 26000 wins % = 0.37
Test policy for episode 27000 wins % = 0.35
Test policy for episode 28000 wins % = 0.35
Test policy for episode 29000 wins % = 0.29
Test policy for episode 30000 wins % = 0.32
Test policy for episode 31000 wins % = 0.29
Test policy for episode 32000 wins % = 0.38
Test policy for episode 33000 wins % = 0.3
Test policy for episode 34000 wins % = 0.4
Test policy for episode 35000 wins % = 0.36
Test policy for episode 36000 wins % = 0.31
Test policy for episode 37000 wins % = 0.31
Test policy for episode 38000 wins % = 0.31
Test policy for episode 39000 wins % = 0.38
Test policy for episode 40000 wins % = 0.3
Test policy for episode 41000 wins % = 0.36
Test policy for episode 42000 wins % = 0.39
Test policy for episode 43000 wins % = 0.34
Test policy for episode 44000 wins % = 0.26
Test policy for episode 45000 wins % = 0.38
Test policy for episode 46000 wins % = 0.33
Test policy for episode 47000 wins % = 0.41
Test policy for episode 48000 wins % = 0.39
Test policy for episode 49000 wins % = 0.4
0.34
Press any key to continue . . .
```

Initialize, for all  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}(s)$ :

$Q(s, a) \leftarrow$  arbitrary

$Returns(s, a) \leftarrow$  empty list

$\pi \leftarrow$  an arbitrary  $\varepsilon$ -soft policy

Repeat forever:

(a) Generate an episode using  $\pi$

(b) For each pair  $s, a$  appearing in the episode:

$R \leftarrow$  return following the first occurrence of  $s, a$

Append  $R$  to  $Returns(s, a)$

$Q(s, a) \leftarrow$  average( $Returns(s, a)$ )


(c) For each  $s$  in the episode:

$a^* \leftarrow \arg \max_a Q(s, a)$

For all  $a \in \mathcal{A}(s)$ :

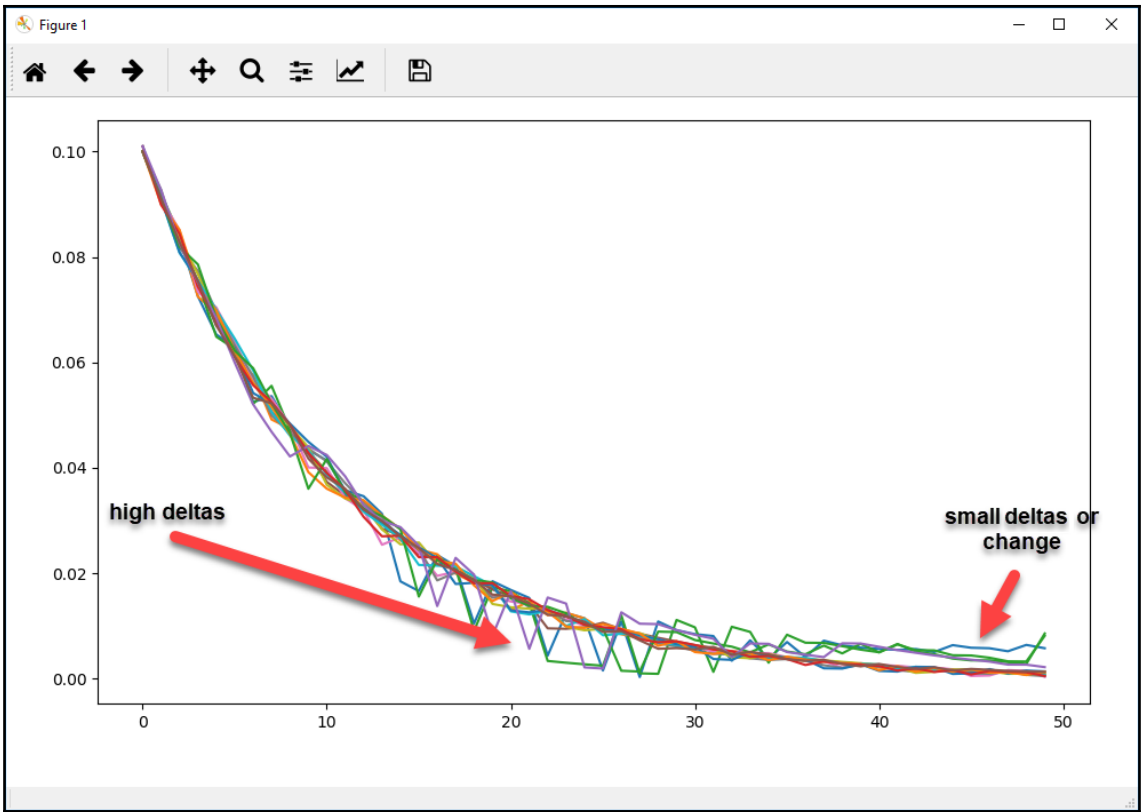
$$\pi(s, a) \leftarrow \begin{cases} 1 - \varepsilon + \varepsilon/|\mathcal{A}(s)| & \text{if } a = a^* \\ \varepsilon/|\mathcal{A}(s)| & \text{if } a \neq a^* \end{cases}$$

# Chapter 4: Temporal Difference Learning

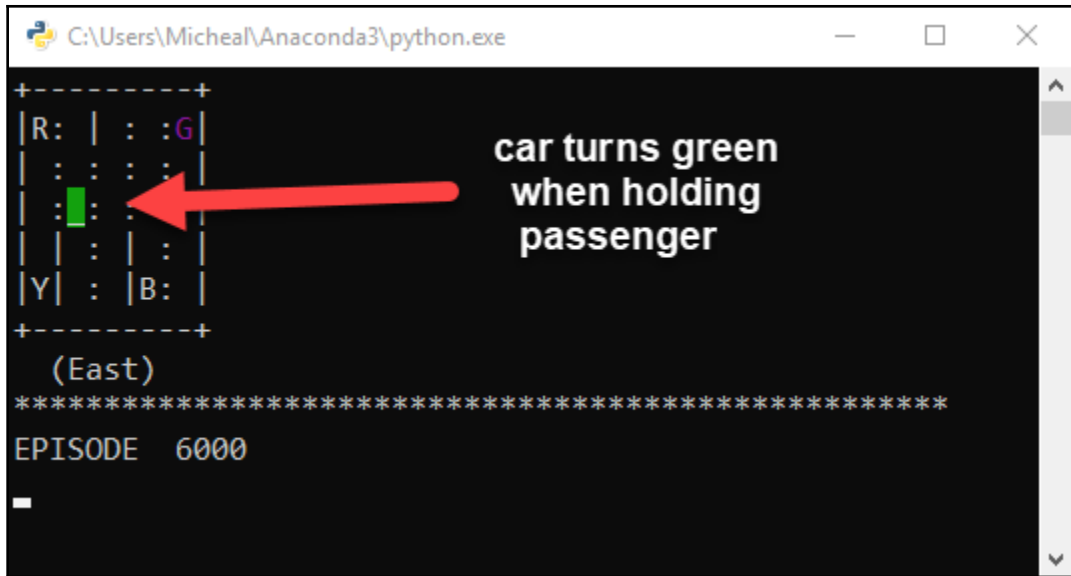


```

C:\Users\Micheal\Anaconda3\python.exe
100% ##### | 10000/10000 [00:21:00:00, 458.54it/s]
[[ 0.          -1.49423831 -1.9958849  -1.99776133]
 [-1.2447176  -1.98811069 -1.98747494 -1.98064978]
 [-1.86806477 -1.98964328 -1.88924666 -1.49072661]
 [-1.98202495 -1.90378281 -1.55590434  0.        ]]
Press any key to continue . . .
  
```






```
C:\Users\Micheal\Anaconda3\python.exe  
(Left)  
SFFF  
FHFH  
FFFH  
HFFG  
*****  
EPISODE 10000  
|
```

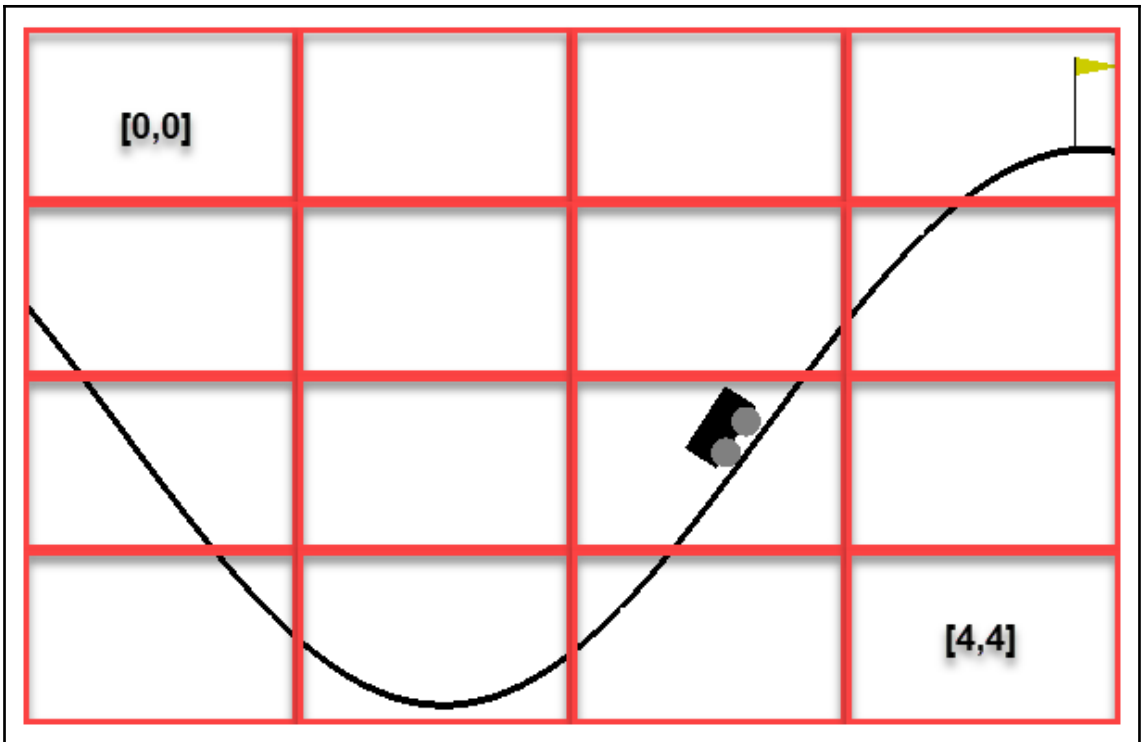


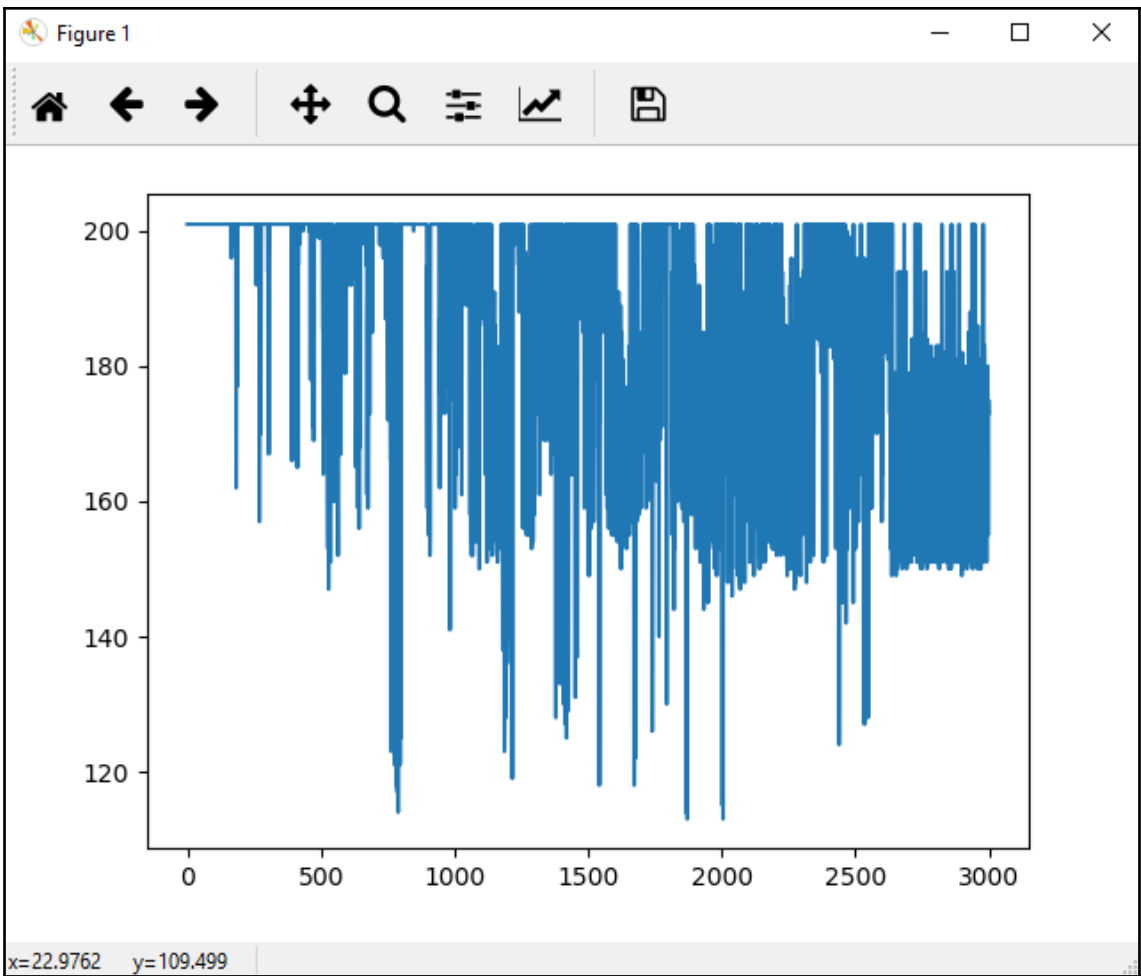


## Chapter 5: Exploring SARSA









**Mountain Car****Observation**

Type: Box(2)

Num	Observation	Min	Max
0	position	-1.2	0.6
1	velocity	-0.07	0.07

**Actions**

Type: Discrete(3)

Num	Action
0	push left
1	no push
2	push right

**Cart Pole****Observation**


Type: Box(4)

Num	Observation	Min	Max
0	Cart Position	-2.4	2.4
1	Cart Velocity	-Inf	Inf
2	Pole Angle	~ -41.8°	~ 41.8°
3	Pole Velocity At Tip	-Inf	Inf


**Actions**

Type: Discrete(2)

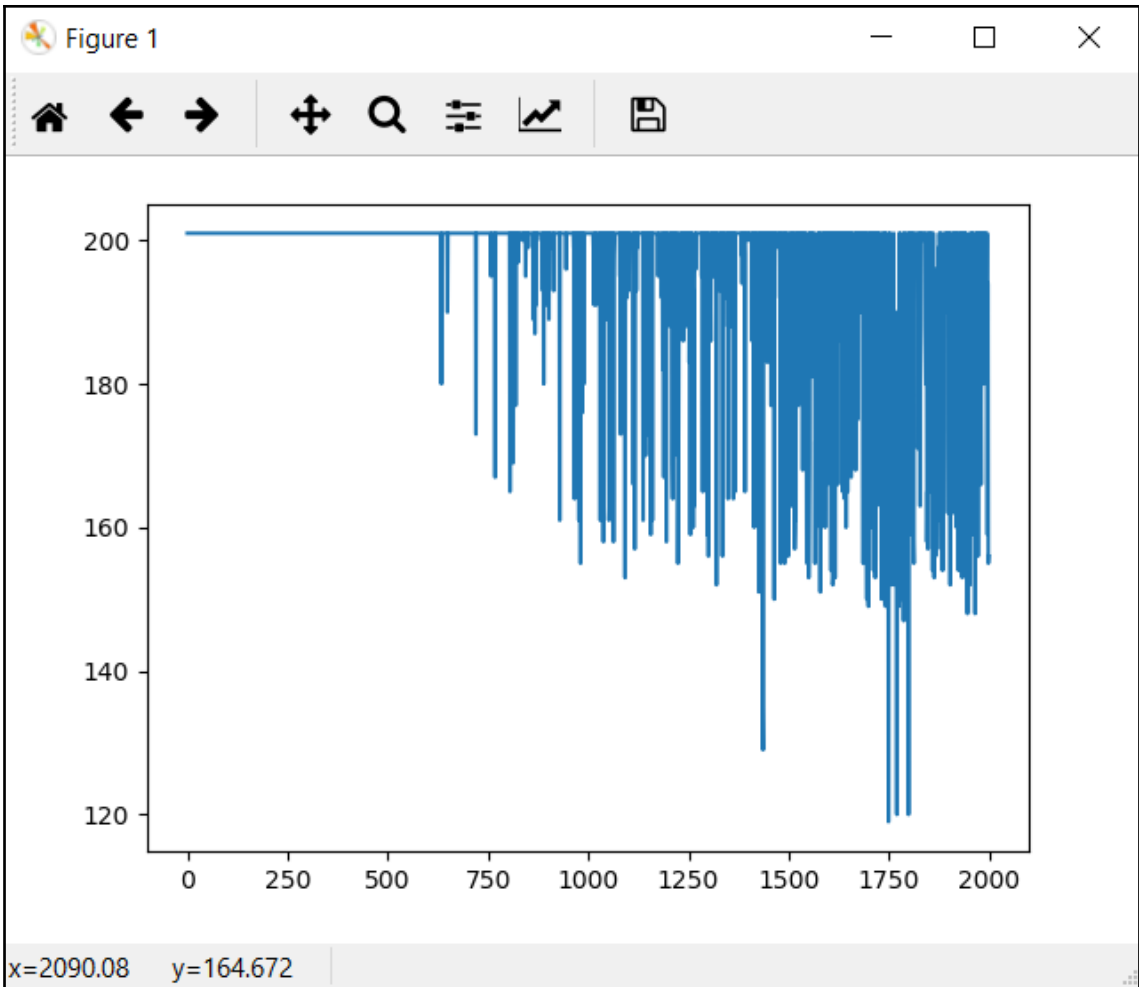
Num	Action
0	Push cart to the left
1	Push cart to the right



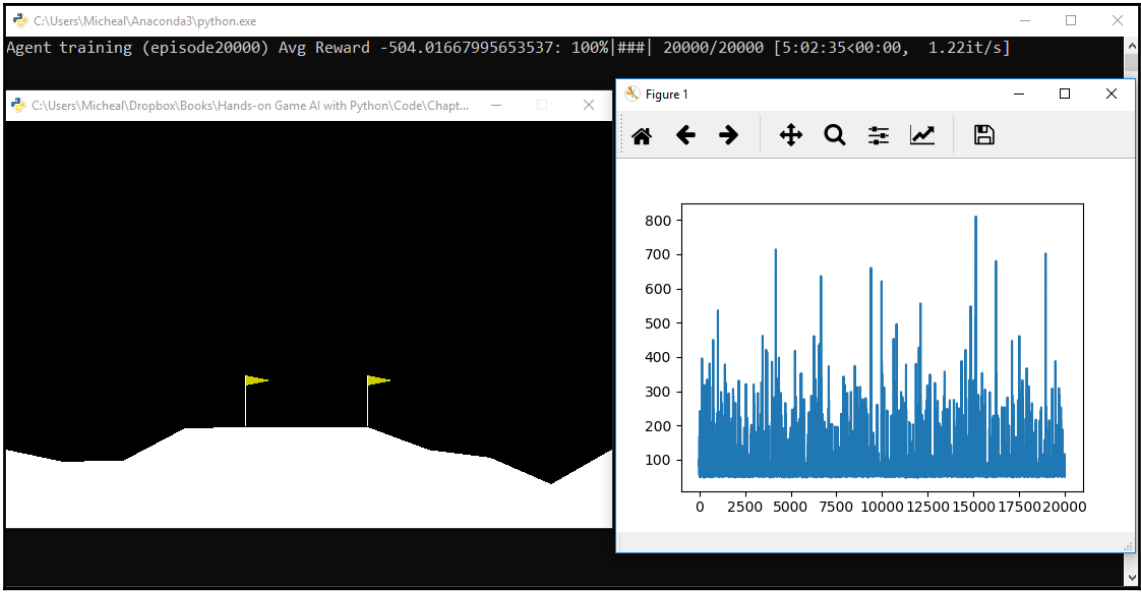
TD( $\lambda$ )



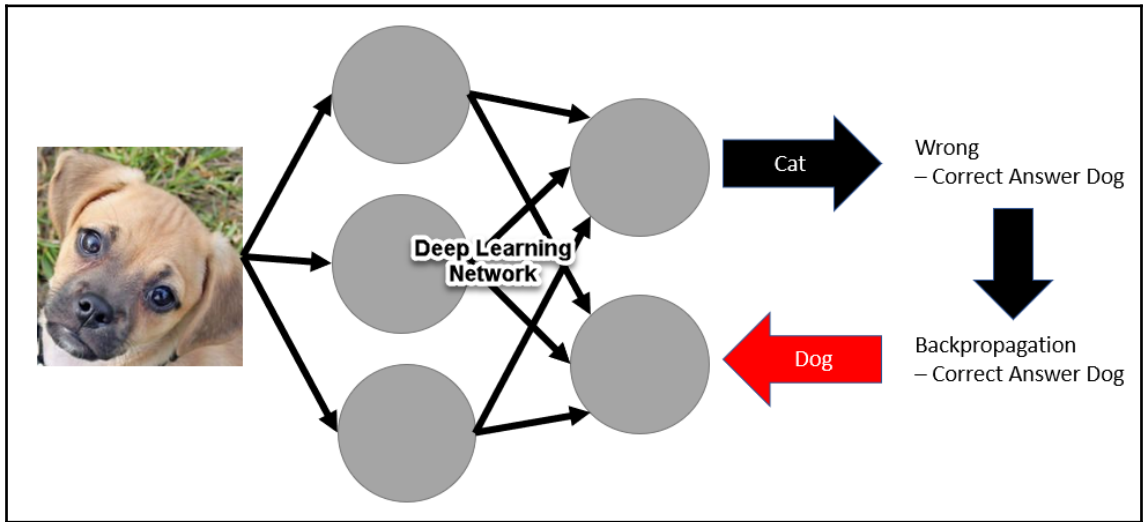
Tabular TD( $\lambda$ )	SARSA( $\lambda$ )	Q( $\lambda$ )
<p>Initialize <math>V(s)</math> arbitrarily and <math>e(s) = 0</math>, for all <math>s \in S</math></p> <p>Repeat (for each episode):</p> <p>Initialize <math>s</math></p> <p>Repeat (for each step of episode) :</p> <p><math>a \leftarrow</math> action given by <math>\pi</math> for <math>s</math></p> <p>Take action <math>a</math>, observe reward, <math>r</math>, and next state <math>s'</math></p> <p><math>\delta \leftarrow r + \gamma V(s') - V(s)</math></p> <p><math>e(s) \leftarrow e(s) + 1</math></p> <p>For all <math>s</math> :</p> <p><math>V(s) \leftarrow V(s) + \alpha \delta e(s)</math></p> <p><math>e(s) \leftarrow \gamma \lambda e(s)</math></p> <p><math>s \leftarrow s'</math></p> <p>Until <math>s</math> is terminal</p>	<p>Initialize <math>Q(s,a)</math> arbitrarily and <math>e(s,a) = 0</math>, for all <math>s,a</math></p> <p>Repeat (for each episode):</p> <p>Initialize <math>s,a</math></p> <p>Repeat (for each step of episode):</p> <p>Take action <math>a</math>, observe <math>r,s'</math></p> <p>Choose <math>a'</math> from <math>s'</math> using policy derived from <math>Q</math> (e.g. <math>\epsilon</math>-greedy)</p> <p><math>\delta \leftarrow r + \gamma Q(s',a') - Q(s,a)</math></p> <p><math>e(s,a) \leftarrow e(s,a) + 1</math></p> <p>For all <math>s,a</math> :</p> <p><math>Q(s,a) \leftarrow Q(s,a) + \alpha \delta e(s,a)</math></p> <p><math>e(s,a) \leftarrow \gamma \lambda e(s,a)</math></p> <p><math>s \leftarrow s'; a \leftarrow a'</math></p> <p>Until <math>s</math> is terminal</p>	<p>Initialize <math>Q(s,a)</math> arbitrarily and <math>e(s,a) = 0</math>, for all <math>s,a</math></p> <p>Repeat (for each episode):</p> <p>Initialize <math>s,a</math></p> <p>Repeat (for each step of episode):</p> <p>Take action <math>a</math>, observe <math>r,s'</math></p> <p>Choose <math>a'</math> from <math>s'</math> using policy derived from <math>Q</math> (e.g. <math>\epsilon</math>-greedy)</p> <p><math>a' \leftarrow \arg \max_b Q(s',b)</math> (if <math>a</math> ties for the max, then <math>a' \leftarrow a</math>)</p> <p><math>\delta \leftarrow r + \gamma Q(s',a') - Q(s,a)</math></p> <p><math>e(s,a) \leftarrow e(s,a) + 1</math></p> <p>For all <math>s,a</math> :</p> <p><math>Q(s,a) \leftarrow Q(s,a) + \alpha \delta e(s,a)</math></p> <p>If <math>a' = a</math>, then <math>e(s,a) \leftarrow \gamma \lambda e(s,a)</math></p> <p>else <math>e(s,a) \leftarrow 0</math></p> <p><math>s \leftarrow s'; a \leftarrow a'</math></p> <p>Until <math>s</math> is terminal</p>







# Chapter 6: Going Deep with DQN



PyTorch Build	<b>Stable (1.2)</b>	Preview (Nightly)			
Your OS	Linux	Mac	<b>Windows</b>		
Package	<b>Conda</b>	Pip	LibTorch	Source	
Language	Python 2.7	Python 3.5	<b>Python 3.6</b>	Python 3.7	C++
CUDA	9.2	10.0	<b>None</b>		
Run this Command	<code>conda install pytorch torchvision cpuonly -c pytorch</code>				

```
layer1 = torch.randn(inputs, hidden, device=device, dtype=dtype)
layer2 = torch.randn(hidden, outputs, device=device, dtype=dtype)

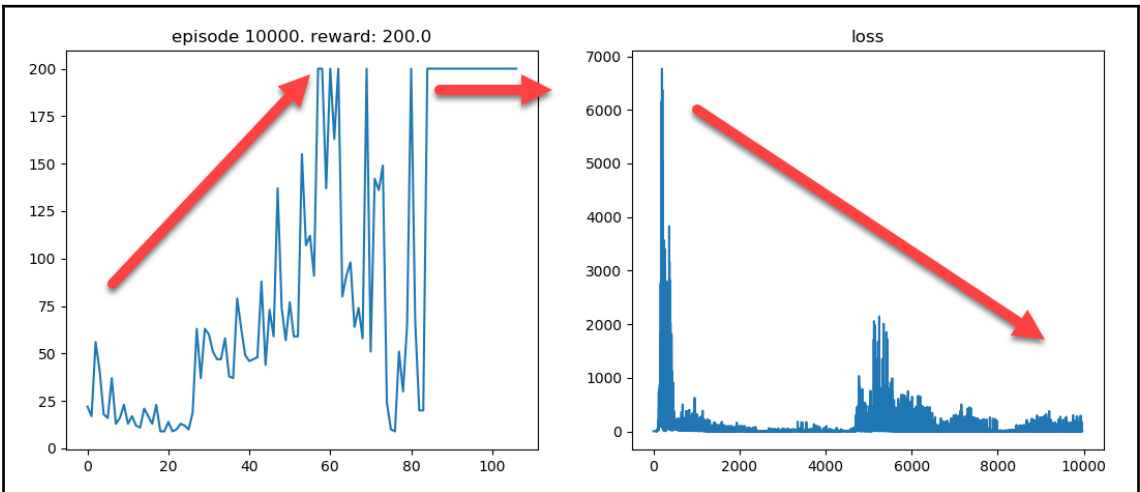

tensor([[ 1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
         1.0212e-01, -8.1076e-01,  7.1715e-01, ...,
         1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
         1.0212e-01, -8.1076e-01,  7.1715e-01, ...,
         1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
         1.0212e-01, -8.1076e-01,  7.1715e-01, ...])
> T: tensor([[ 1.1902e+00,  1.0212e-01, -8.1076e-01, ...,
              1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
              1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
              1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
              1.1902e+00, -1.8583e-01, -1.8020e+00, ...])
> data: tensor([[ 1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
                 1.0212e-01, -8.1076e-01,  7.1715e-01, ...,
                 1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
                 1.0212e-01, -8.1076e-01,  7.1715e-01, ...,
                 1.1902e+00, -1.8583e-01, -1.8020e+00, ...,
                 1.0212e-01, -8.1076e-01,  7.1715e-01, ...])
> device: device(type='cpu')
> dtype: torch.float32
grad: None
grad_fn: None
is_cuda: False
is_leaf: True
is_mkl_dnn: False
is_quantized: False
is_sparse: False
> layout: torch.strided
name: None
ndim: 2
output_nr: 0
requires_grad: False
> shape: torch.Size([1000, 100])
_backward_hooks: None



layer one tensor

tensor([[ 3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
         3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
         3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
         3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
         3.1835e-02, -8.1076e-01,  7.1715e-01, ...])
> T: tensor([[ 3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
              3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
              3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
              3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
              3.1835e-02, -8.1076e-01,  7.1715e-01, ...])
> data: tensor([[ 3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
                 3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
                 3.1835e-02, -8.1076e-01,  7.1715e-01, ...,
                 3.1835e-02, -7.6702e-01,  7.1715e-01, ...,
                 3.1835e-02, -8.1076e-01,  7.1715e-01, ...])
> device: device(type='cpu')
> dtype: torch.float32
grad: None
grad_fn: None
is_cuda: False
is_leaf: True
is_mkl_dnn: False
is_quantized: False
is_sparse: False
> layout: torch.strided
name: None
ndim: 2
output_nr: 0
requires_grad: False
> shape: torch.Size([100, 10])
_backward_hooks: None


layer two tensor
```

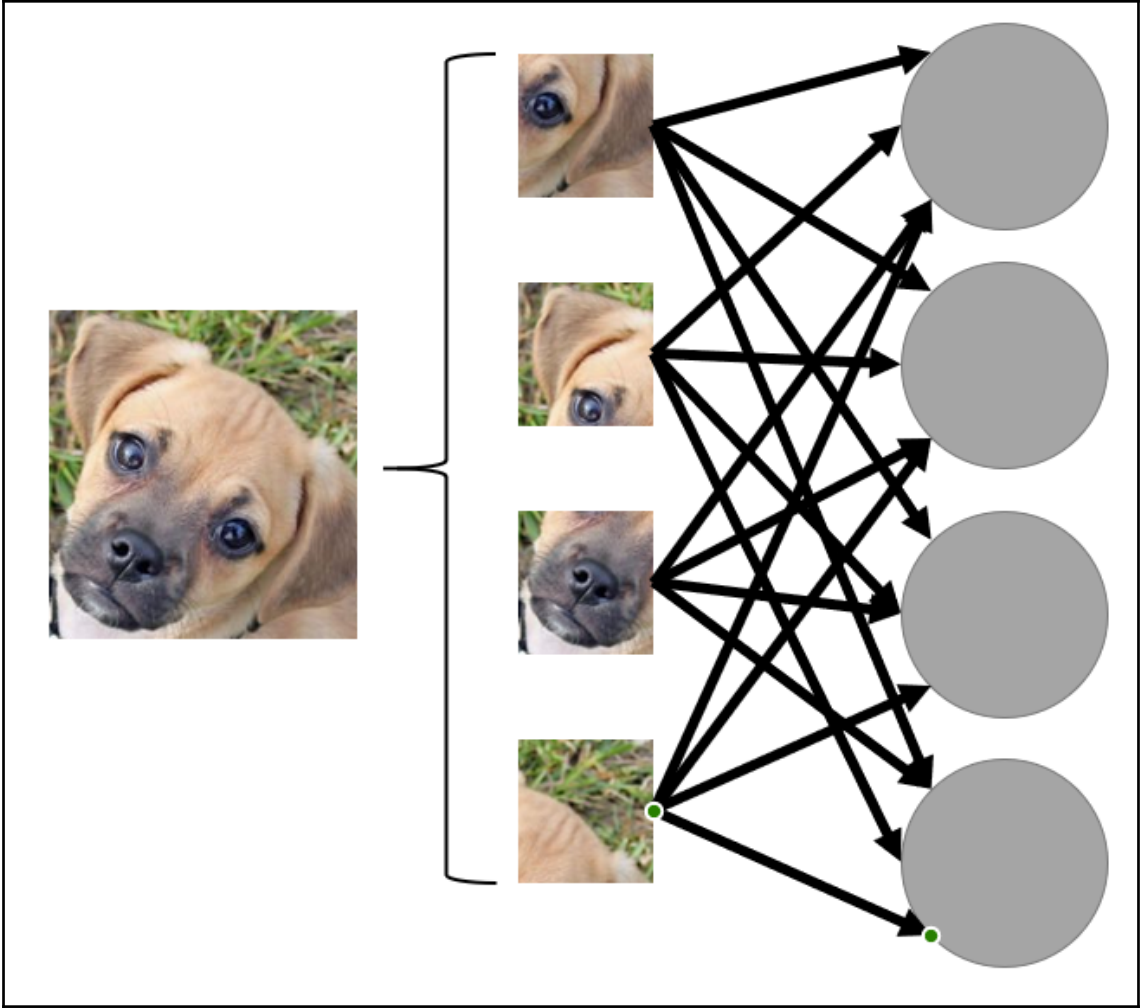
```
99 635.5177001953125
199 3.109295606613159
299 0.026640085503458977
399 0.0005209375522099435
499 6.454912363551557e-05
```







# Chapter 7: Going Deeper with DDQN





← → ↻ 🔒 tensorspace.org/html/playground/... ☆ U A 🗑️ | 🌐 🟢

 **TensorSpace.js** EN | 中 

## TensorSpace Playground

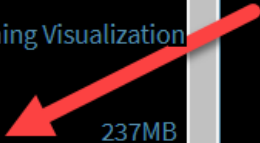
TensorSpace Playground is the place we designed for presenting different pre-built models. In the "Playground", we can experience different pre-trained deep learning models, including object classification, object detections, image generations etc.

All models in the "Playground" are interactive. We can move the mouse to see the relation lines among layers; we can click the layer aggregation to check feature maps; we can move the camera to view from any direction... We highly recommend to try playground models with a better network condition due to the model sizes (e.g. VGG-16 has > 500MB, AlexNet has > 250MB...). It may take some time to load some large models. To have a better experience in this amusement park, we highly recommend to use medium or large device (device width > 750px).


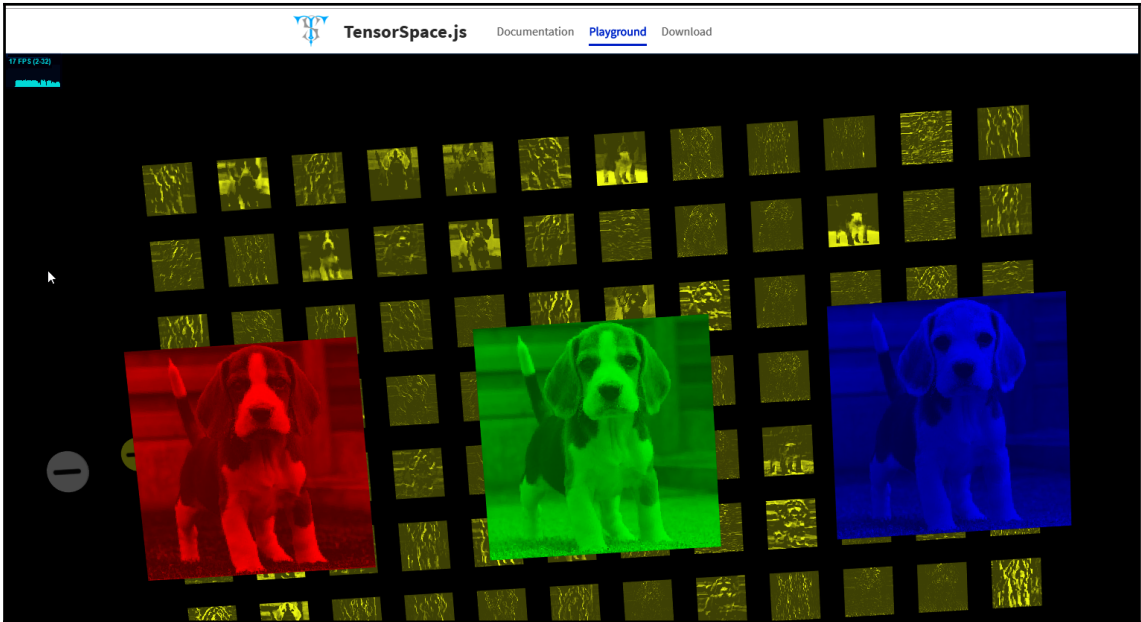
While we are still developing and expanding our model collections, enjoy the models and have fun~


### Interaction Guide:


Introduction	
LeNet	0.2MB
LeNet Training Visualization	
AlexNet	237MB
InceptionV3	96MB
ResNet-50	98MB
MobileNetv1	16MB
YOLOv2-tiny	60MB
ACGAN	3MB

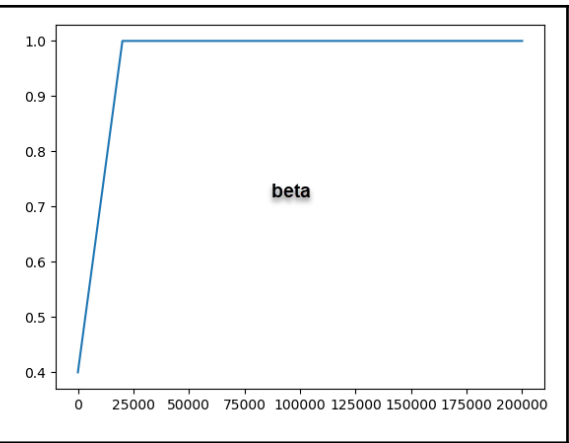

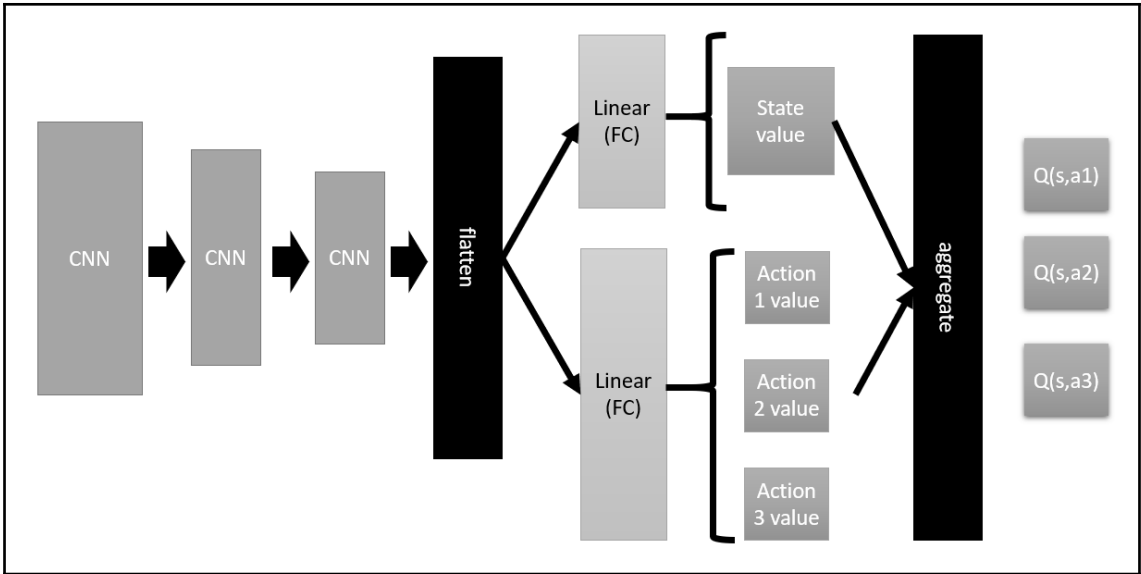












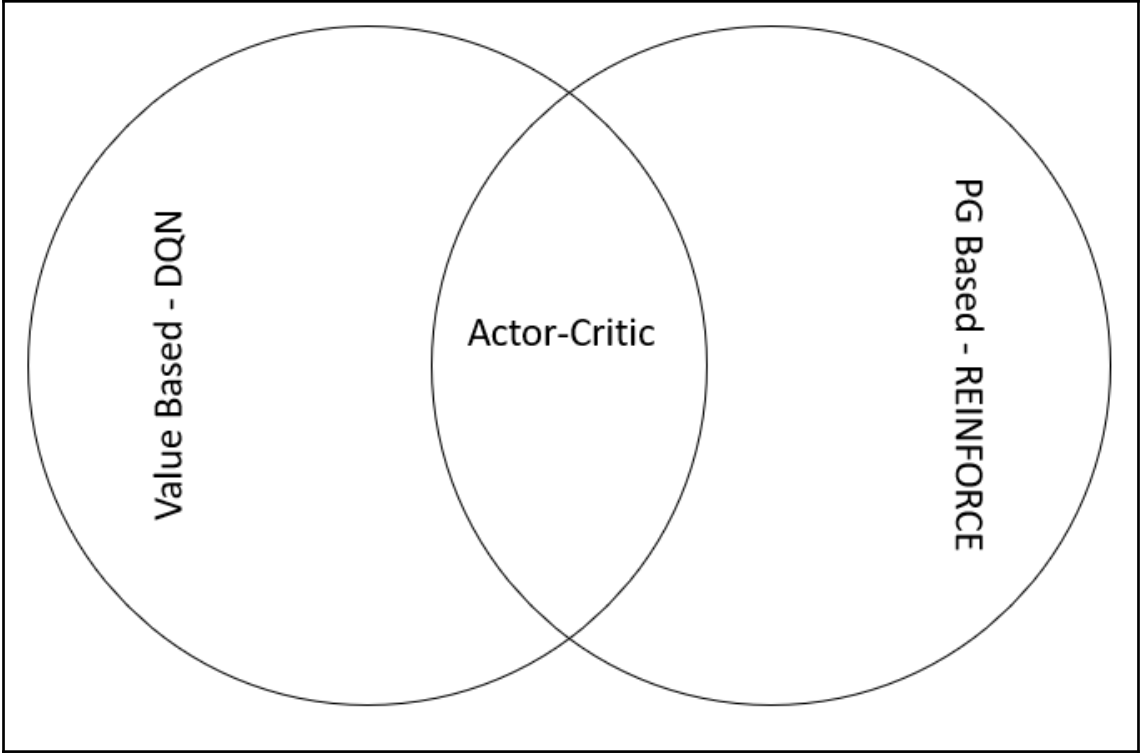


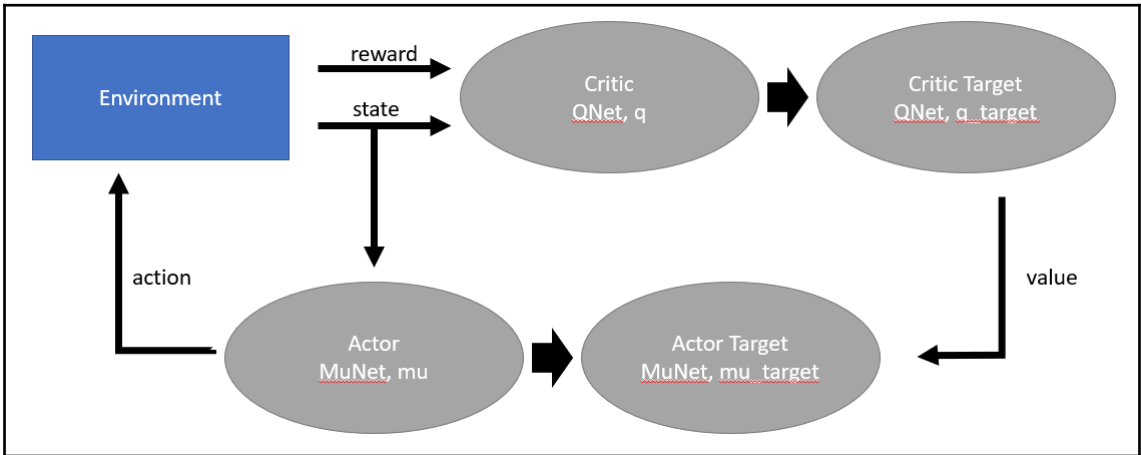

**the eagle has landed!!!**



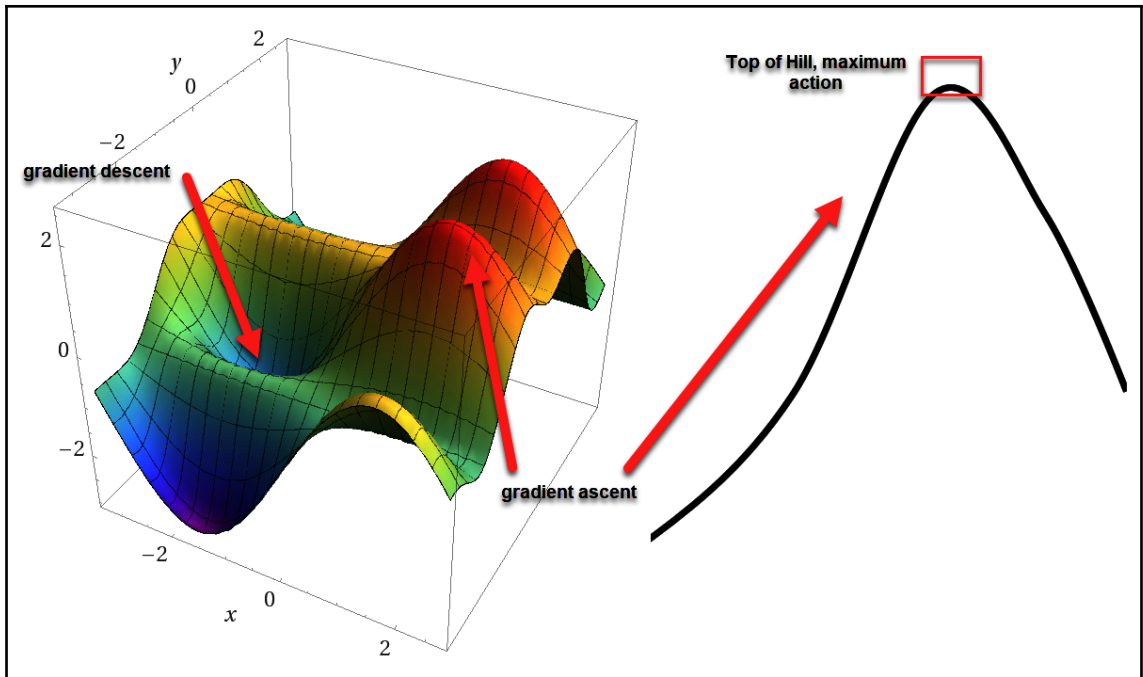
# Chapter 8: Policy Gradient Methods











```
C:\ProgramData\Anaconda3\envs\game\python.exe
fval before -2.056834235095027e-17
a/e/r 0.004196146307487529 0.004245106934584498 0.988466573904631
fval after -0.00419614630748755
Episode 91      Last reward: -912.6011060549151 Average reward -973.64
('lagrange multiplier:', tensor(0.2289), 'grad_norm:', tensor(0.0912))
fval before 7.105427357601002e-17
a/e/r 0.004674936250927307 0.0045777398718830084 1.021232394536284
fval after -0.004674936250927236
Episode 92      Last reward: -1011.7306568770845      Average reward -998.55
('lagrange multiplier:', tensor(0.2264), 'grad_norm:', tensor(0.1863))
fval before -9.349246523159213e-18
a/e/r 0.004193475940531191 0.0045323531795798495 0.925231502131069
fval after -0.0041934759405312
Episode 93      Last reward: -998.7409040861277 Average reward -1002.09
('lagrange multiplier:', tensor(0.1980), 'grad_norm:', tensor(0.1286))
fval before 1.6828643741686585e-17
a/e/r 0.0036979960207453337 0.003958261570809577 0.9342475110832528
fval after -0.0036979960207453168
Episode 94      Last reward: -903.6812992964333 Average reward -1006.95
('lagrange multiplier:', tensor(0.2623), 'grad_norm:', tensor(0.2498))
fval before -4.113668470190054e-17
a/e/r 0.005604256367298496 0.005227386032703227 1.0720953708483585
fval after -0.005604256367298536
Episode 95      Last reward: -1004.4161781600075      Average reward -1010.09
('lagrange multiplier:', tensor(0.2762), 'grad_norm:', tensor(0.3254))
fval before 1.3088945132422898e-17
a/e/r 0.005118301075806763 0.005511908796890441 0.9285895802002869
fval after -0.00511830107580675
Episode 96      Last reward: -1017.8063539769834      Average reward -1003.34
```



Line search  
(like gradient ascent)



Trust region

# Chapter 9: Optimizing for Continuous Control




Full name

Email address

Computer id  [Win32](#) [Win64](#) [Linux](#) [OSX](#)


Acceptance  I agree to the [terms and conditions](#) of the Trial License.

```
Run speed = 1.000 x real time [S]lower, [F]aster
Render every frame           On
Switch camera (#cams = 5)    [Tab] (camera ID = -1)
[O]nly contact forces        On
Referenc[e] frames           On
T[ra]nsparent                On
Display [M]ocap bodies       On
Stop                          [Space]
Advance simulation by one step [right arrow]
[H]ide Menu
Record [V]ideo (OHT)
Capture frame
Start [I]pdb
Toggle geomgroup visibility  0-4
```




```
C:\ProgramData\Anaconda3\envs\game\python.exe
# of episode :740, avg score : 31.4
# of episode :760, avg score : -1.4
# of episode :780, avg score : -38.4
# of episode :800, avg score : -61.1
# of episode :820, avg score : -95.9
# of episode :840, avg score : -120.6
# of episode :860, avg score : -6.4
# of episode :880, avg score : 12.4
# of episode :900, avg score : 7.5
# of episode :920, avg score : -93.8
# of episode :940, avg score : -152.1
# of episode :960, avg score : -122.7
# of episode :980, avg score : -77.0
# of episode :1000, avg score : -69.9
# of episode :1020, avg score : -4.5
# of episode :1040, avg score : -55.6
# of episode :1060, avg score : -114.5
# of episode :1080, avg score : -19.3
# of episode :1100, avg score : 63.3
# of episode :1120, avg score : -114.9
# of episode :1140, avg score : -118.9
# of episode :1160, avg score : -122.5
# of episode :1180, avg score : -156.1
# of episode :1200, avg score : -133.7
# of episode :1220, avg score : -144.2
# of episode :1240, avg score : -106.1
# of episode :1260, avg score : 10.5
# of episode :1280, avg score : -46.7
# of episode :1300, avg score : -133.6
```

```
C:\ProgramData\Anaconda3\envs\game\python.exe
# of episode :20, avg score : -187.0
# of episode :40, avg score : -224.8
# of episode :60, avg score : -217.2
# of episode :80, avg score : -179.0
# of episode :100, avg score : -150.2
# of episode :120, avg score : -172.5
# of episode :140, avg score : -201.0
# of episode :160, avg score : -124.4
# of episode :180, avg score : -67.1
# of episode :200, avg score : -136.7
# of episode :220, avg score : 4.4
# of episode :240, avg score : -21.2
# of episode :260, avg score : 21.2
# of episode :280, avg score : 9.3
# of episode :300, avg score : -10.8
# of episode :320, avg score : -45.0
# of episode :340, avg score : 42.8
# of episode :360, avg score : -33.2
# of episode :380, avg score : -40.6
# of episode :400, avg score : -88.6
# of episode :420, avg score : -26.5
# of episode :440, avg score : -106.6
# of episode :460, avg score : 39.5
# of episode :480, avg score : 68.1
```





```
C:\ProgramData\Anaconda3\envs\game\python.exe
Step # :46000, avg score : -2.9
Step # :46500, avg score : 11.6
Step # :47000, avg score : -37.8
Step # :47500, avg score : 36.3
Step # :48000, avg score : 4.6
Step # :48500, avg score : -38.6
Step # :49000, avg score : -11.2
Step # :49500, avg score : 2.6
Step # :50000, avg score : 9.6
Step # :50500, avg score : 28.0
Step # :51000, avg score : -22.9
Step # :51500, avg score : -3.7
Step # :52000, avg score : 31.0
Step # :52500, avg score : -4.9
Step # :53000, avg score : 3.2
Step # :53500, avg score : 64.0
Step # :54000, avg score : -0.9
Step # :54500, avg score : 23.7
Step # :55000, avg score : 19.9
Step # :55500, avg score : -18.8
Step # :56000, avg score : -76.0
Step # :56500, avg score : -12.5
Step # :57000, avg score : -37.2
Step # :57500, avg score : -65.4
Step # :58000, avg score : -20.4
Step # :58500, avg score : 42.4
Step # :59000, avg score : -50.8
Step # :59500, avg score : -15.9
Step # :60000, avg score : 10.0
Press any key to continue . . .
```



```
C:\ProgramData\Anaconda3\envs\game\python.exe C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapt...  
# of episode :20, avg score : -202.9  
# of episode :40, avg score : -133.7  
# of episode :60, avg score : -31.4  
# of episode :80, avg score : 46.0  
# of episode :100, avg score : -7.1  
# of episode :120, avg score : -34.6  
# of episode :140, avg score : -92.3  
# of episode :160, avg score : -101.3  
# of episode :180, avg score : -123.4  
# of episode :200, avg score : -110.6
```


The screenshot shows a game environment with a black background. A blue robot-like character is positioned in the upper center. There are two yellow flags on the ground, one on the left and one on the right. Several red stars are scattered in the upper half of the screen, possibly representing targets or rewards. The interface includes a terminal window on the left displaying episode scores and a window title bar at the top.




C:\ProgramData\Anaconda3\envs\game\python.exe

```
# of episode :80, avg score : -205.5, buffer size : 840
# of episode :100, avg score : -204.8, buffer size : 1063
# of episode :120, avg score : -214.8, buffer size : 1294
# of episode :140, avg score : -240.3, buffer size : 1532
# of episode :160, avg score : -179.0, buffer size : 1791
# of episode :180, avg score : -169.1, buffer size : 2054
# of episode :200, avg score : -124.0, buffer size : 2427
# of episode :220, avg score : -209.3, buffer size : 2806
# of episode :240, avg score : -176.4, buffer size : 3161
# of episode :260, avg score : -141.7, buffer size : 3469
# of episode :280, avg score : -123.3, buffer size : 3834
# of episode :300, avg score : -75.7, buffer size : 4509
# of episode :320, avg score : -58.7, buffer size : 5365
# of episode :340, avg score : -2.9, buffer size : 6000
# of episode :360, avg score : -32.5, buffer size : 6000
# of episode :380, avg score : 45.0, buffer size : 6000
# of episode :400, avg score : 75.0, buffer size : 6000
# of episode :420, avg score : 30.6, buffer size : 6000
# of episode :440, avg score : 87.6, buffer size : 6000
# of episode :460, avg score : 75.5, buffer size : 6000
# of episode :480, avg score : 63.8, buffer size : 6000
# of episode :500, avg score : 29.6, buffer size : 6000
# of episode :520, avg score : 68.5, buffer size : 6000
# of episode :540, avg score : 72.6, buffer size : 6000
# of episode :560, avg score : 73.6, buffer size : 6000
# of episode :580, avg score : 55.7, buffer size : 6000
# of episode :600, avg score : 54.9, buffer size : 6000
# of episode :620, avg score : 29.8, buffer size : 6000
# of episode :640, avg score : 45.1, buffer size : 6000
```

# Chapter 10: All about Rainbow DQN

```
2019-11-02 11:40:20.756778: I T:\src\github\tensorflow\tensorflow\core\platform\cpu_feature_guard.cc:140] Your CPU
instructions that this TensorFlow binary was not compiled to use: AVX2
TensorBoard 1.7.0 at http://DESKTOP-V2J9HRG:6006
```



**TensorBoard** INACT...   

**No dashboards are active for the current data set.**


Probable causes:

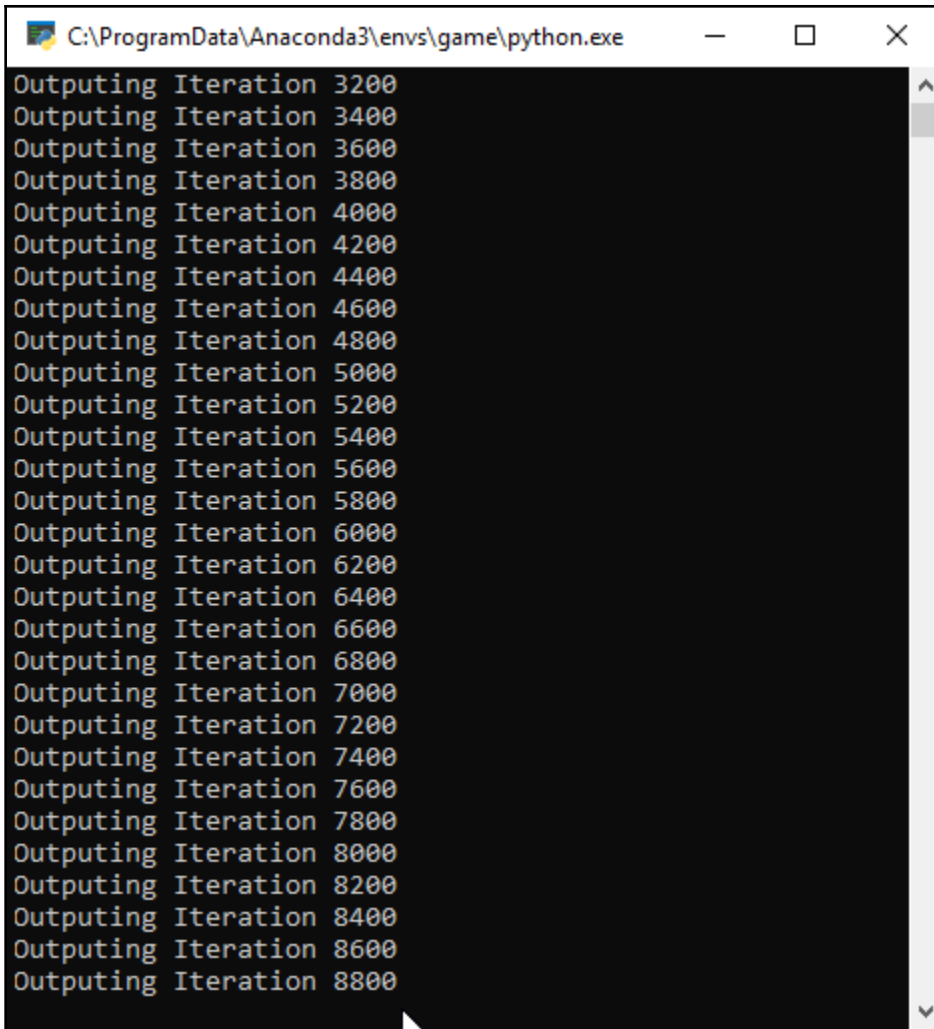
- You haven't written any data to your event files.
- TensorBoard can't find your event files.

If you're new to using TensorBoard, and want to find out how to add data and set up your event files, check out the [README](#) and perhaps the [TensorBoard tutorial](#).

If you think TensorBoard is configured properly, please see [the section of the README devoted to missing data problems](#) and consider filing an issue on GitHub.

*Last reload: Sat Nov 02 2019 11:47:26 GMT-0600 (Mountain Daylight Time)*

*Data location: runs* 




A screenshot of a terminal window titled "C:\ProgramData\Anaconda3\envs\game\python.exe". The terminal displays a list of 24 lines of text, each starting with "Outputing Iteration" followed by a number. The numbers range from 3200 to 8800 in increments of 200. The terminal has a black background and a white scrollbar on the right side.

```
Outputing Iteration 3200
Outputing Iteration 3400
Outputing Iteration 3600
Outputing Iteration 3800
Outputing Iteration 4000
Outputing Iteration 4200
Outputing Iteration 4400
Outputing Iteration 4600
Outputing Iteration 4800
Outputing Iteration 5000
Outputing Iteration 5200
Outputing Iteration 5400
Outputing Iteration 5600
Outputing Iteration 5800
Outputing Iteration 6000
Outputing Iteration 6200
Outputing Iteration 6400
Outputing Iteration 6600
Outputing Iteration 6800
Outputing Iteration 7000
Outputing Iteration 7200
Outputing Iteration 7400
Outputing Iteration 7600
Outputing Iteration 7800
Outputing Iteration 8000
Outputing Iteration 8200
Outputing Iteration 8400
Outputing Iteration 8600
Outputing Iteration 8800
```

- Show data download links
- Ignore outliers in chart scaling
- Tooltip sorting method: default

Smoothing



Horizontal Axis

- STEP**
- RELATIVE
- WALL

Runs

Write a regex to filter runs

- Oct28\_20-25-11\_D  
ESKTOP-V2J9HRG


TOGGLE ALL RUNS

runs


Train

4


Episode  
tag: Train/Episode




Exploration  
tag: Train/Exploration

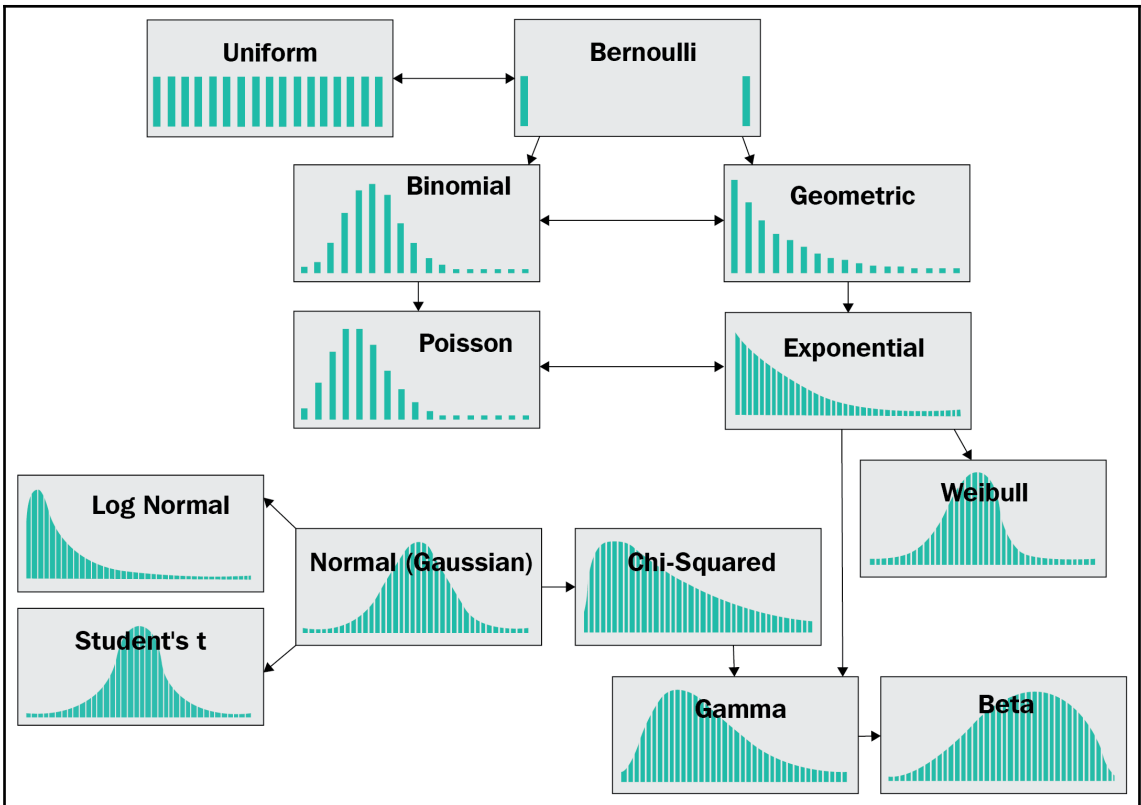


Losses  
tag: Train/Losses




Rewards  
tag: Train/Rewards







**Beta**  
tag: Train/Beta




**Episode**  
tag: Train/Episode



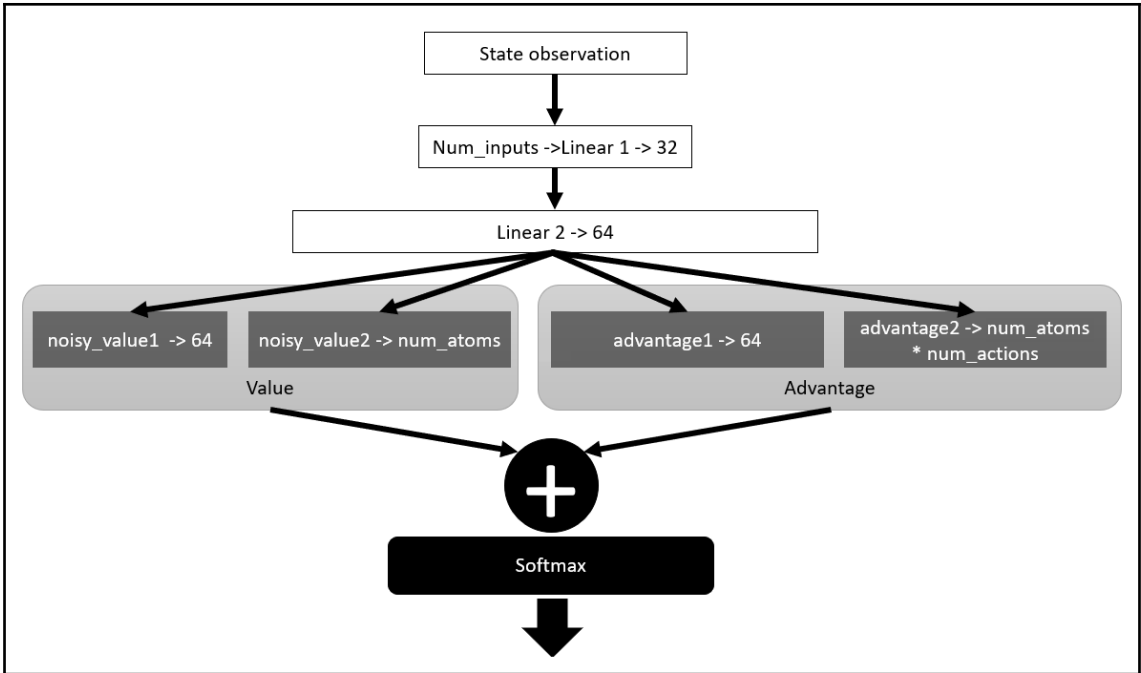
**Losses**  
tag: Train/Losses



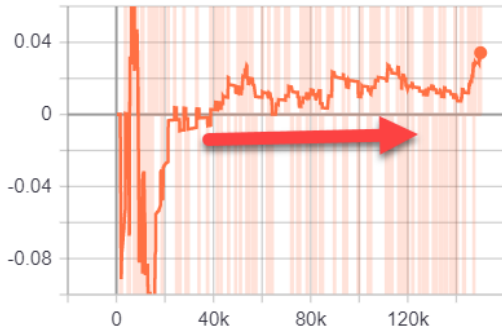
**Rewards**  
tag: Train/Rewards









Advantage  
tag: Train/Advantage



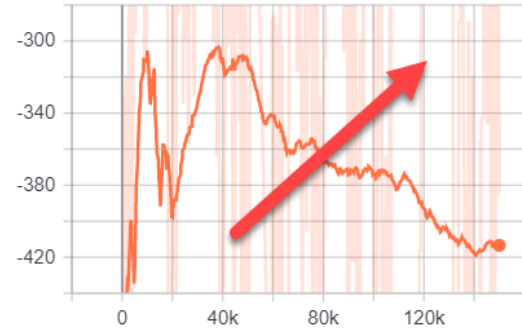
Episode  
tag: Train/Episode



Losses  
tag: Train/Losses




Rewards  
tag: Train/Rewards



# Train


Beta  
tag: Train/Beta




Episode  
tag: Train/Episode




Losses  
tag: Train/Losses




Rewards  
tag: Train/Rewards




Advantage  
tag: Train/Advantage




Episode  
tag: Train/Episode



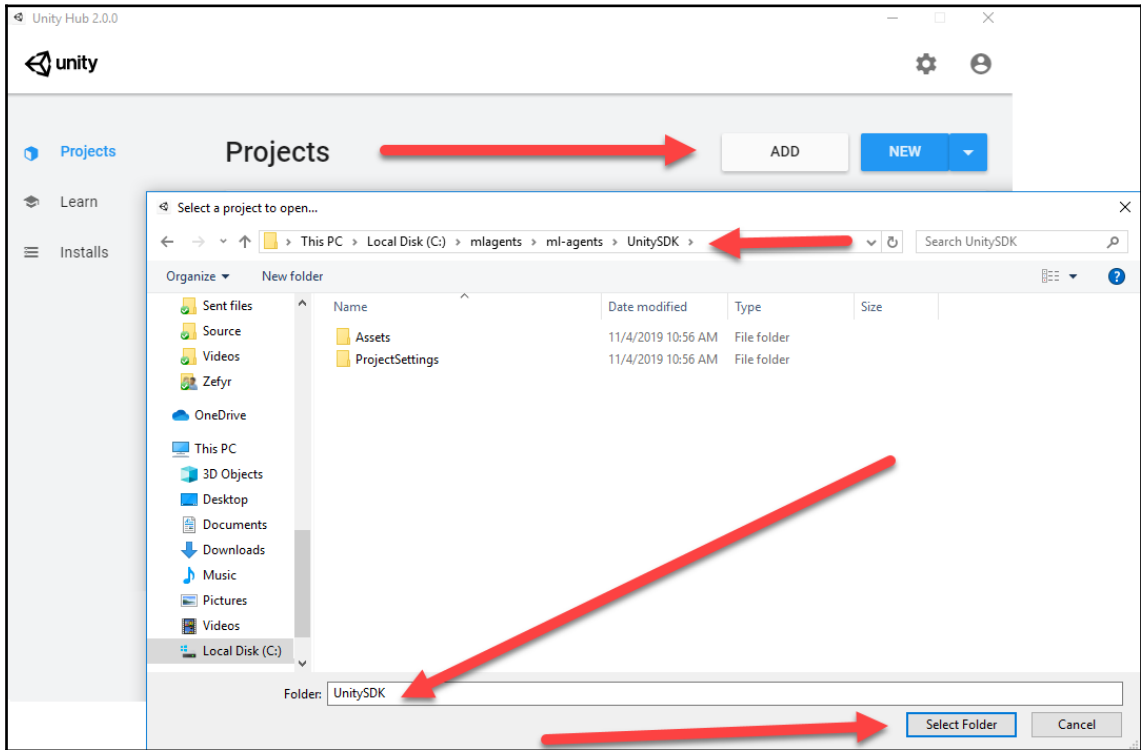
Losses  
tag: Train/Losses




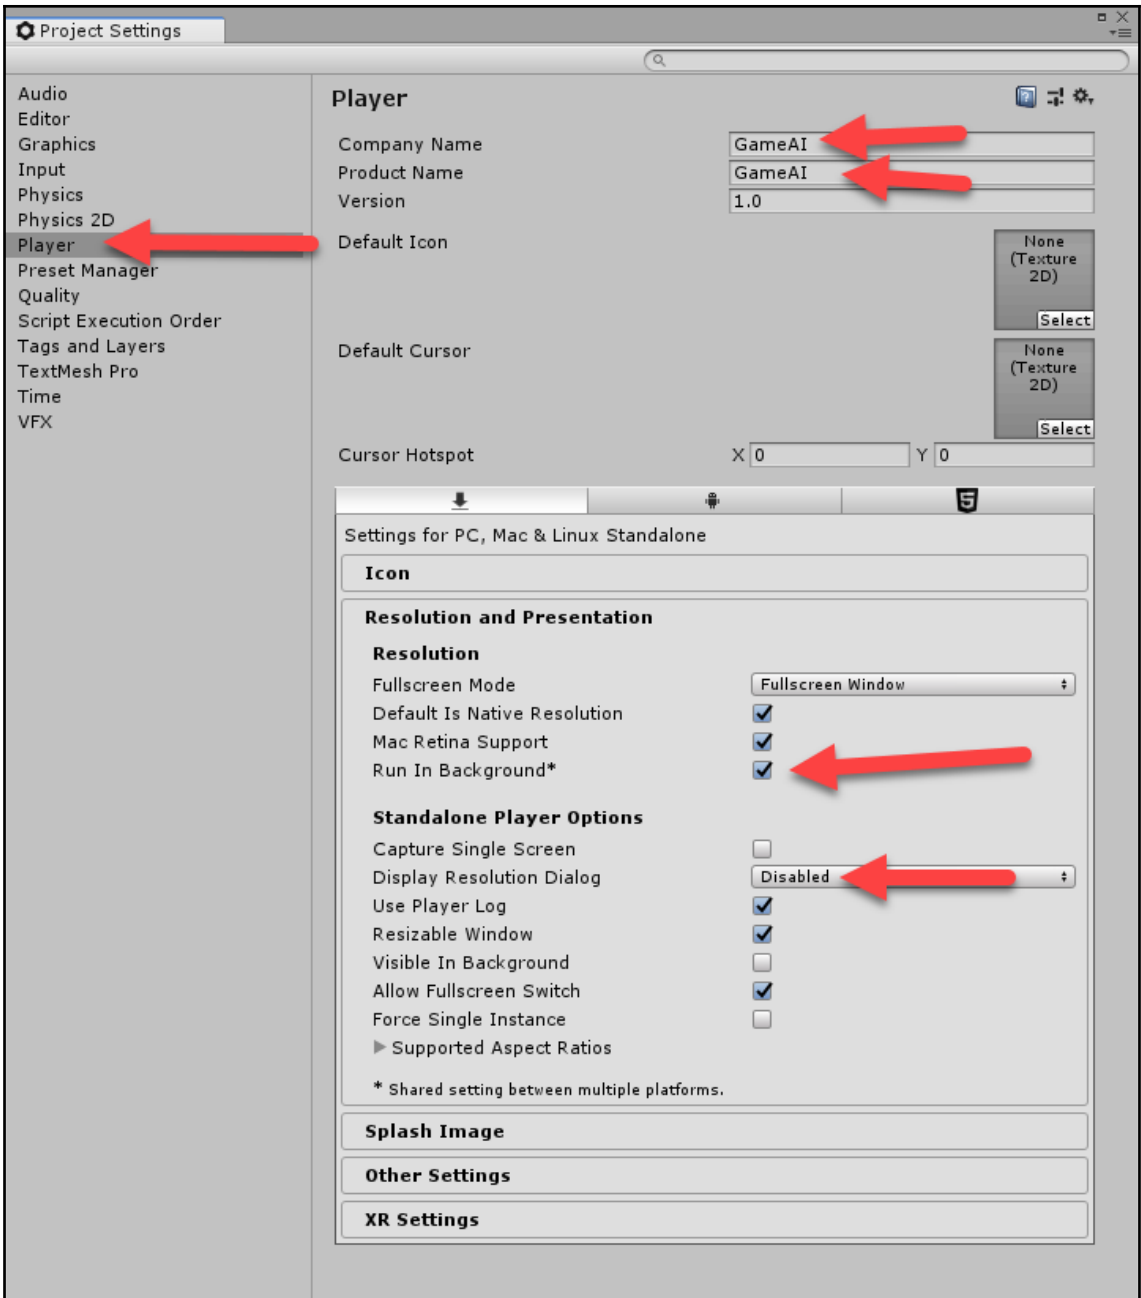
Rewards  
tag: Train/Rewards

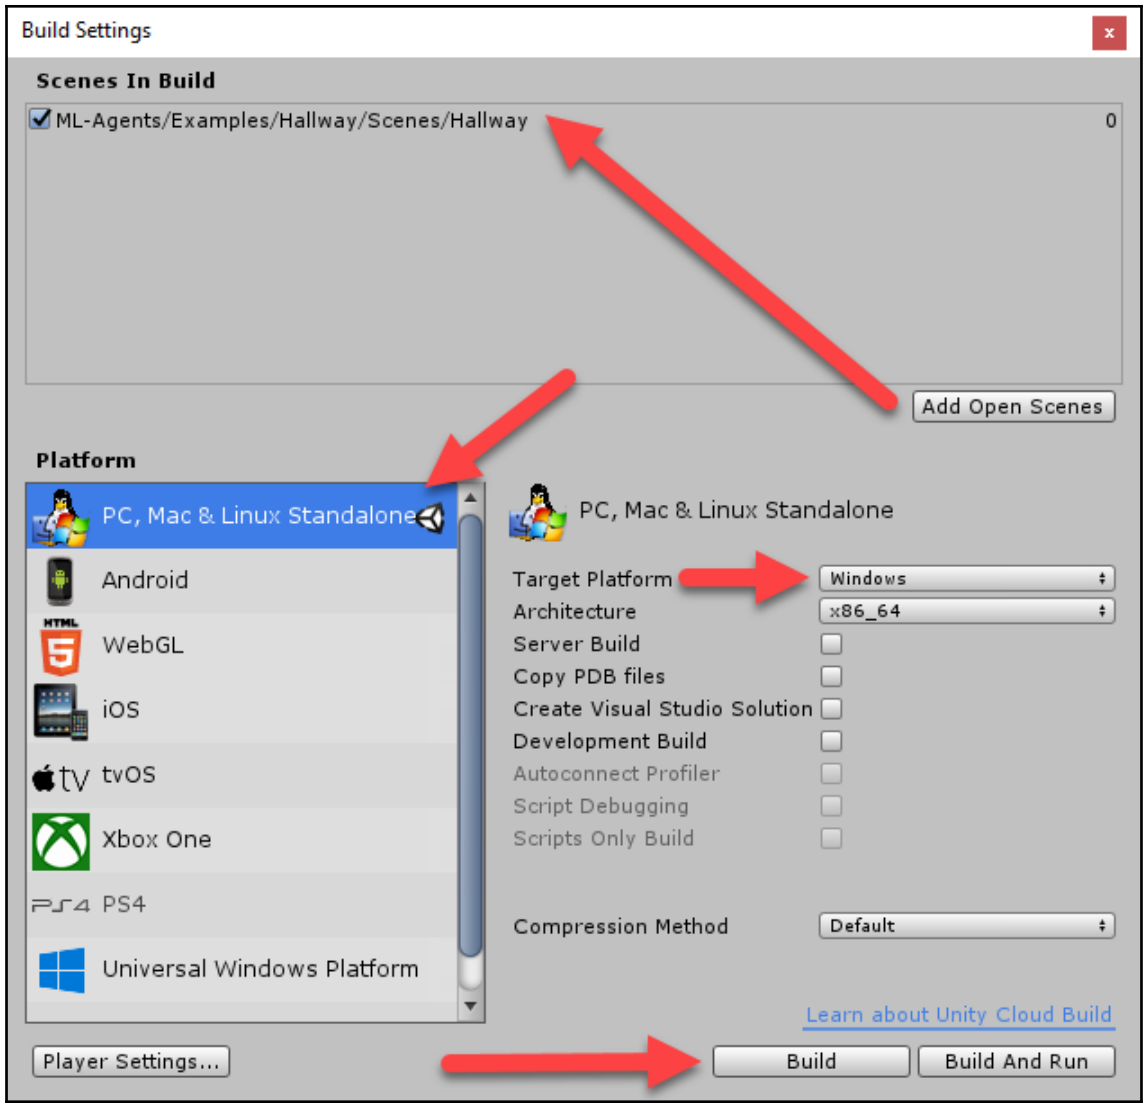


# Chapter 11: Exploiting ML-Agents

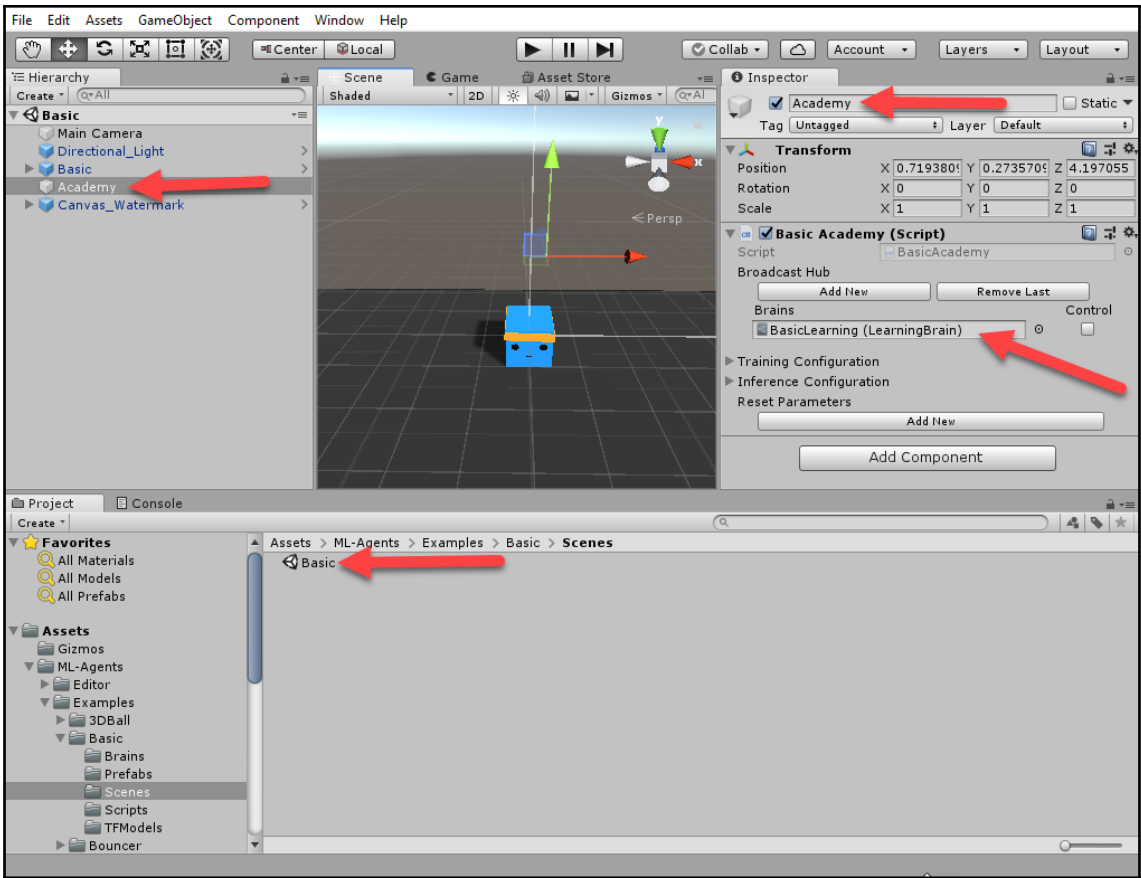


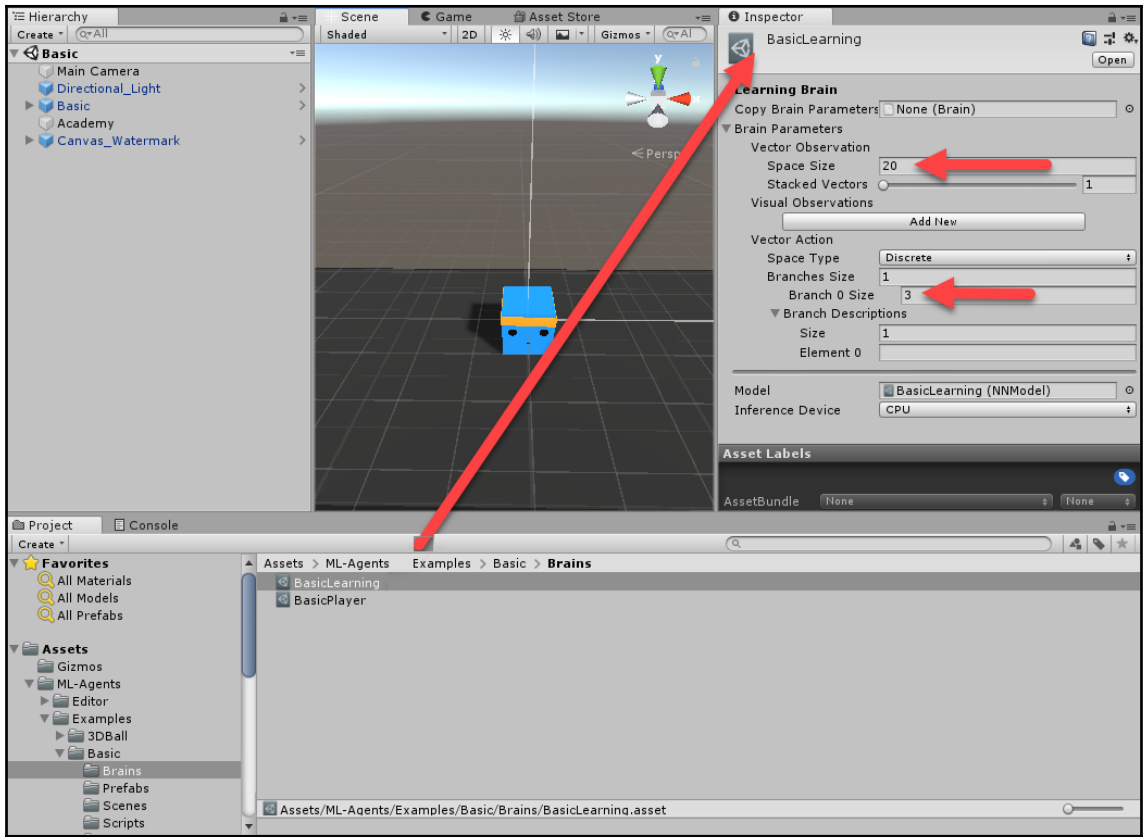


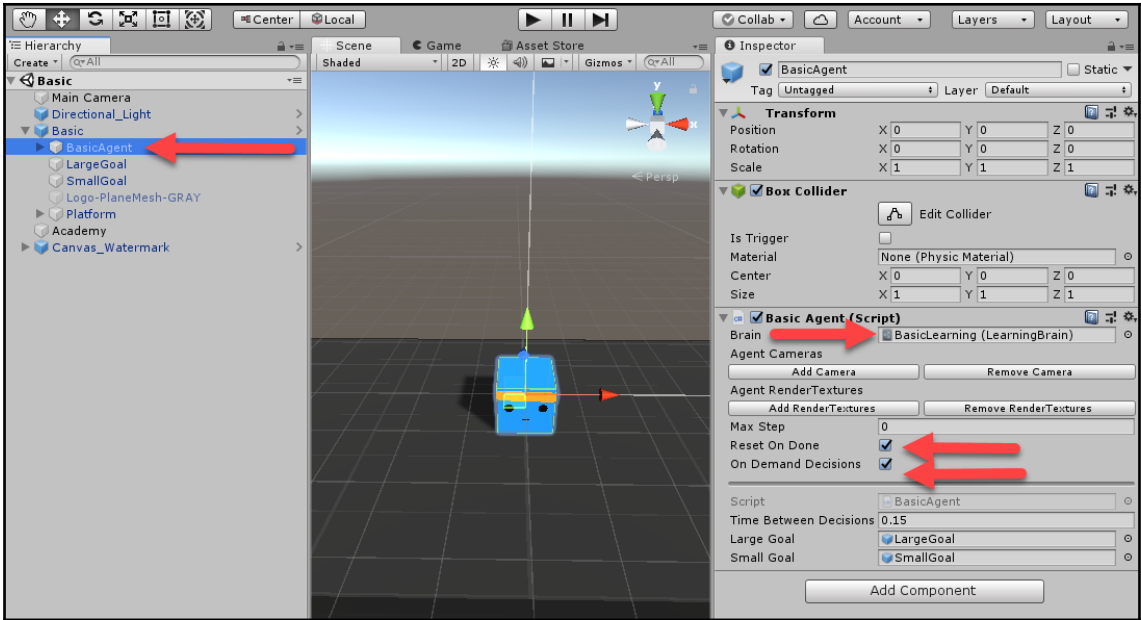








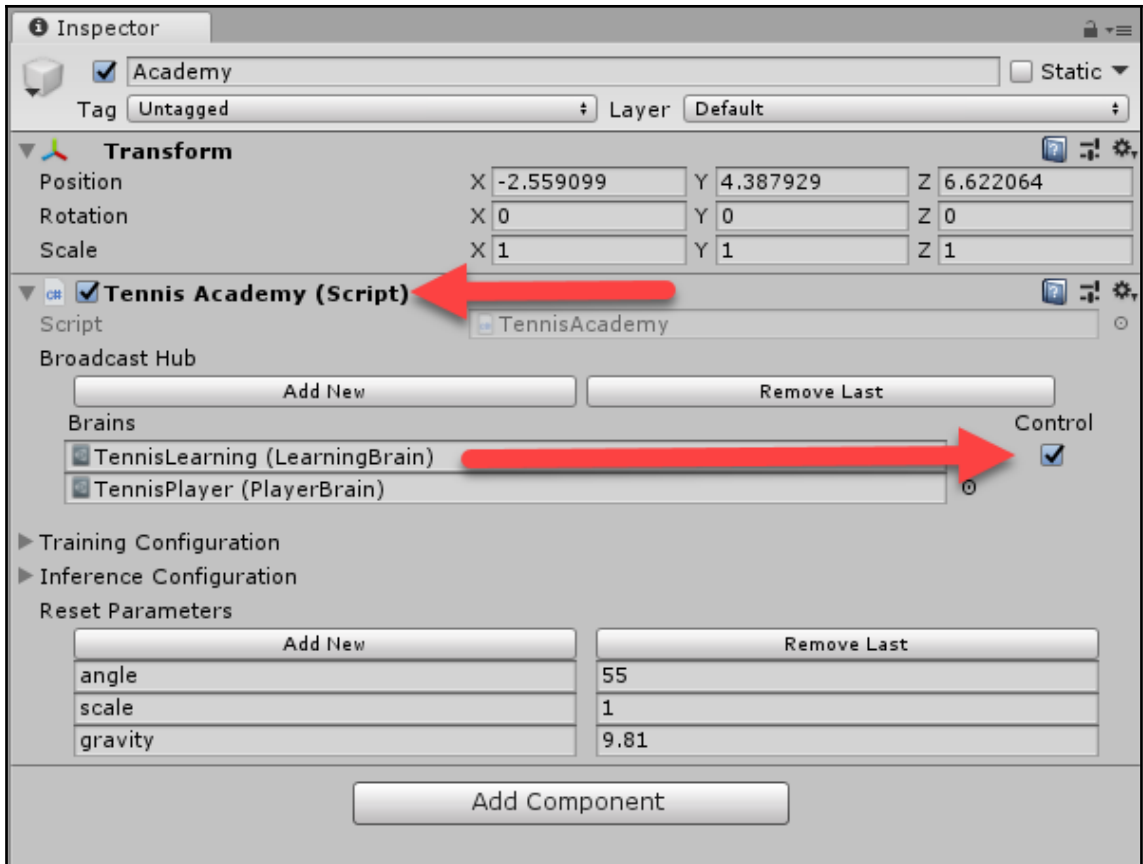


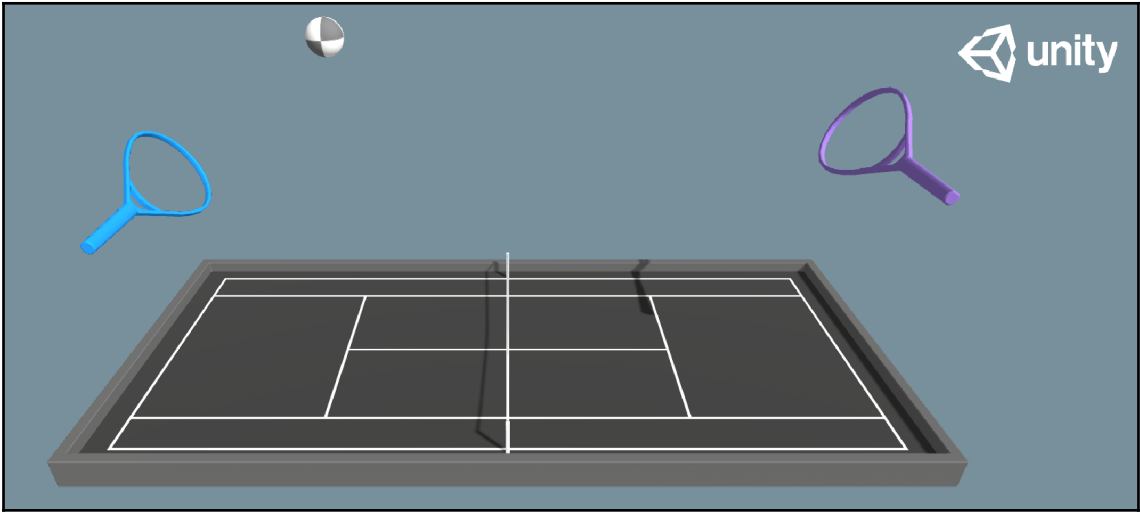









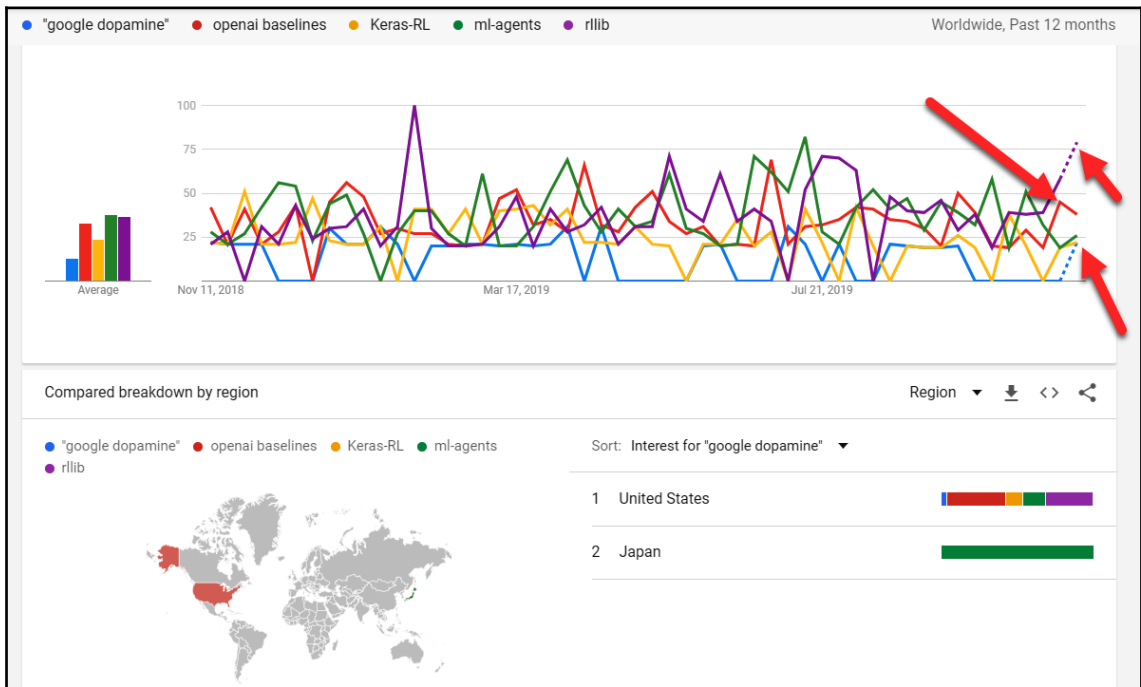








# Chapter 12: DRL Frameworks




## Notebook settings

Runtime type  
Python 3

Hardware accelerator  
GPU

Omit code cell output when saving this notebook

CANCEL SAVE

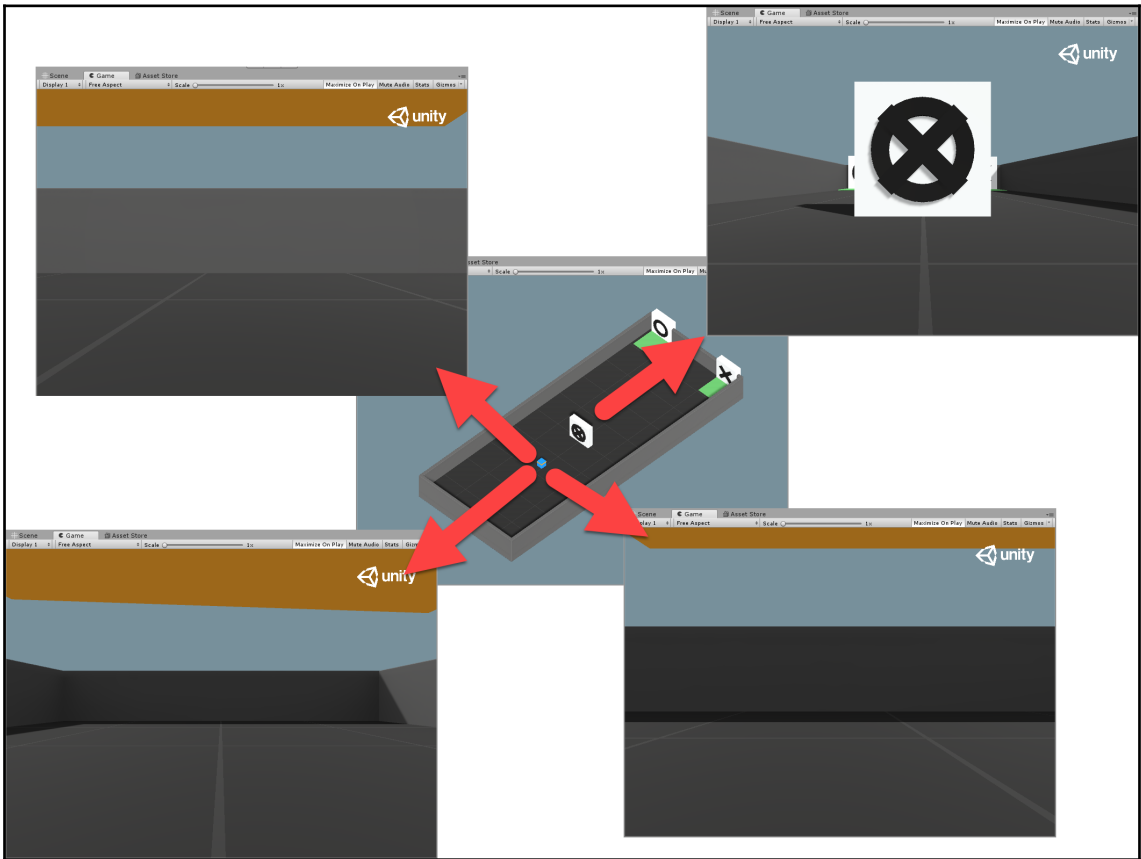



```
RLlib.ipynb ☆
File Edit View Insert Runtime Tools Help All changes saved

+ Code + Text

(pid-491) 'type': 'MultiAgentBatch'
(pid-491)
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/fc_1/kernel:0' shape=(4, 256) dtype=float32>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/fc_1/bias:0' shape=(256,) dtype=float32>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/fc_out/kernel:0' shape=(256, 256) dtype=float32>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/fc_out/bias:0' shape=(256,) dtype=float32>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/action_value/hidden_0/kernel:0' shape=(256, 256) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/action_value/hidden_0/bias:0' shape=(256,) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/action_value/dense/kernel:0' shape=(256, 2) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/action_value/dense/bias:0' shape=(2,) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/state_value/dense/kernel:0' shape=(256, 256) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/state_value/dense/bias:0' shape=(256,) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/state_value/dense_1/kernel:0' shape=(256, 1) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_policy.py:358 -- Optimizing variable <tf.Variable 'default_policy/q_func/state_value/dense_1/bias:0' shape=(1,) dtype=float32_ref>
(pid-491) 2019-11-11 00:36:55,819 INFO tf_run_builder.py:92 -- Executing TF run without tracing. To dump TF timeline traces to disk, set the TF_TIMELINE_DIR environment variable.
Result for DQN_CartPole-v0_1_lr=0.001:
custom_metrics: {}
date: 2019-11-11_00-37-05
done: false
episode_len_mean: 18.3
episode_reward_max: 40.0
episode_reward_mean: 18.3
episode_reward_min: 8.0
episodes_this_iter: 59
episodes_total: 107
experiment_id: c939f9681a9c4a4646f4dd24c6b7b40
hostname: 74d68c92a77
Info:
grad_time_ms: 9.305
learner:
  default_policy:
    cur_lr: 0.0010000000474974513
    max_q: 2.99139484296875
    mean_q: 2.078482151031494
    mean_td_error: 0.18657320737838745
    min_q: 0.9282804131507874
    model: {}
  max_exploration: 0.902
  min_exploration: 0.902
  num_steps_sampled: 2000
  num_steps_trained: 8000
  num_target_updates: 3
  opt_peak_throughput: 3438.853
  opt_samples: 32.0
  replay_time_ms: 2.895
  sample_time_ms: 14.544
  update_time_ms: 11.078
  iterations_since_restore: 2
```

# Chapter 13: 3D Worlds





Inspector  
VisualHallwayLearning

**Learning Brain**

Copy Brain Parameters from :

▼ Brain Parameters

Vector Observation

Space Size

Stacked Vectors

Visual Observations

Index	Width	Height	Gray
Obs 0	<input type="text" value="84"/>	<input type="text" value="84"/>	<input type="checkbox"/>

Vector Action

Space Type

Branches Size

Branch 0 Size

▼ Branch Descriptions


Size

Element 0

---

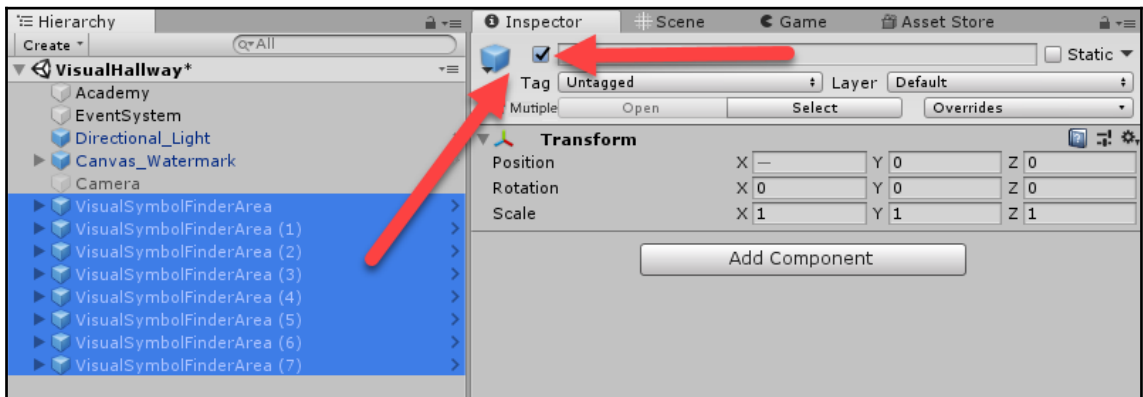
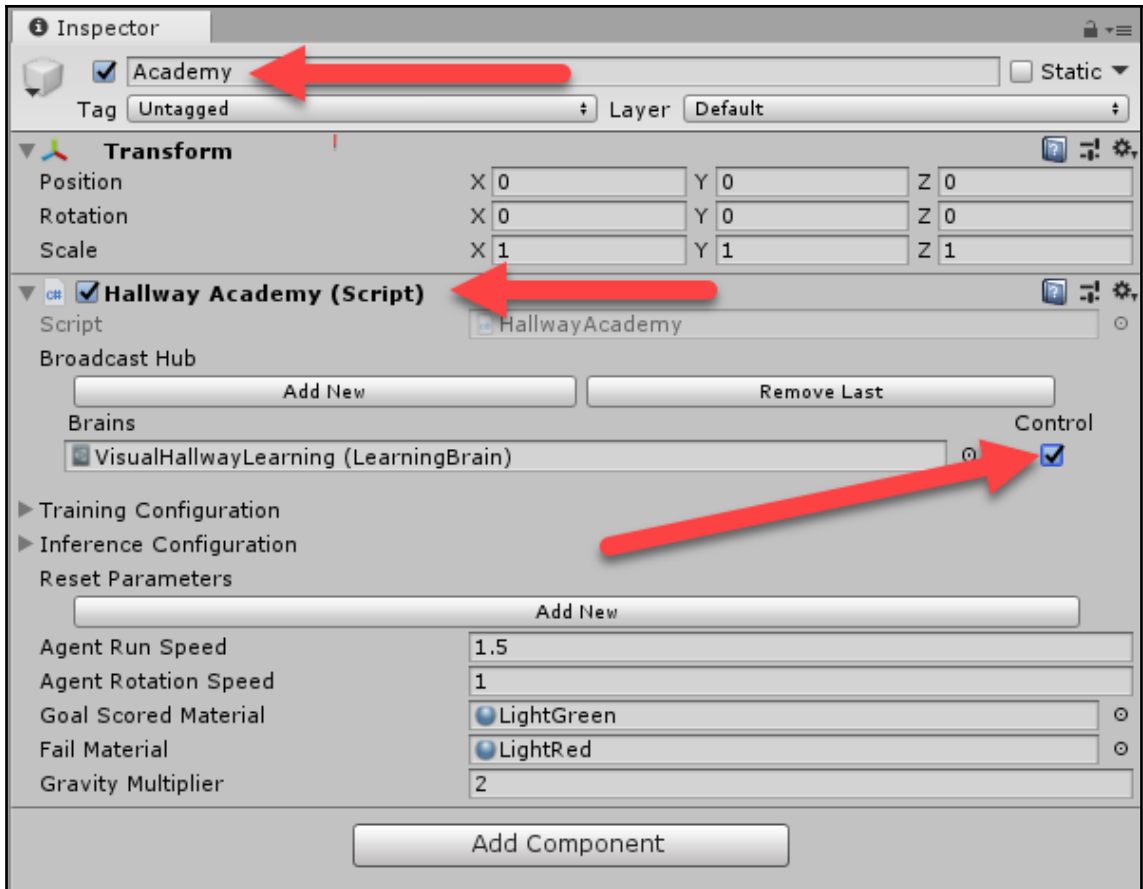
Model

Inference Device

 There is no model for this Brain, cannot run inference. (But can still train)

**Asset Labels**

AssetBundle



```
2019-11-16 15:25:19.214278: I tensorflow/stream_executor/cuda/cudart_stub.cc:29] Ignore above cudart dlerror if you do not have a GPU set up on your machine.
```

```
WARNING:tensorflow:
```

```
The TensorFlow contrib module will not be included in TensorFlow 2.0.
```

```
For more information, please see:
```

```
* https://github.com/tensorflow/community/blob/master/rfcs/20180907-contrib-sunset.md
```

```
* https://github.com/tensorflow/addons
```

```
* https://github.com/tensorflow/io (for I/O related ops)
```

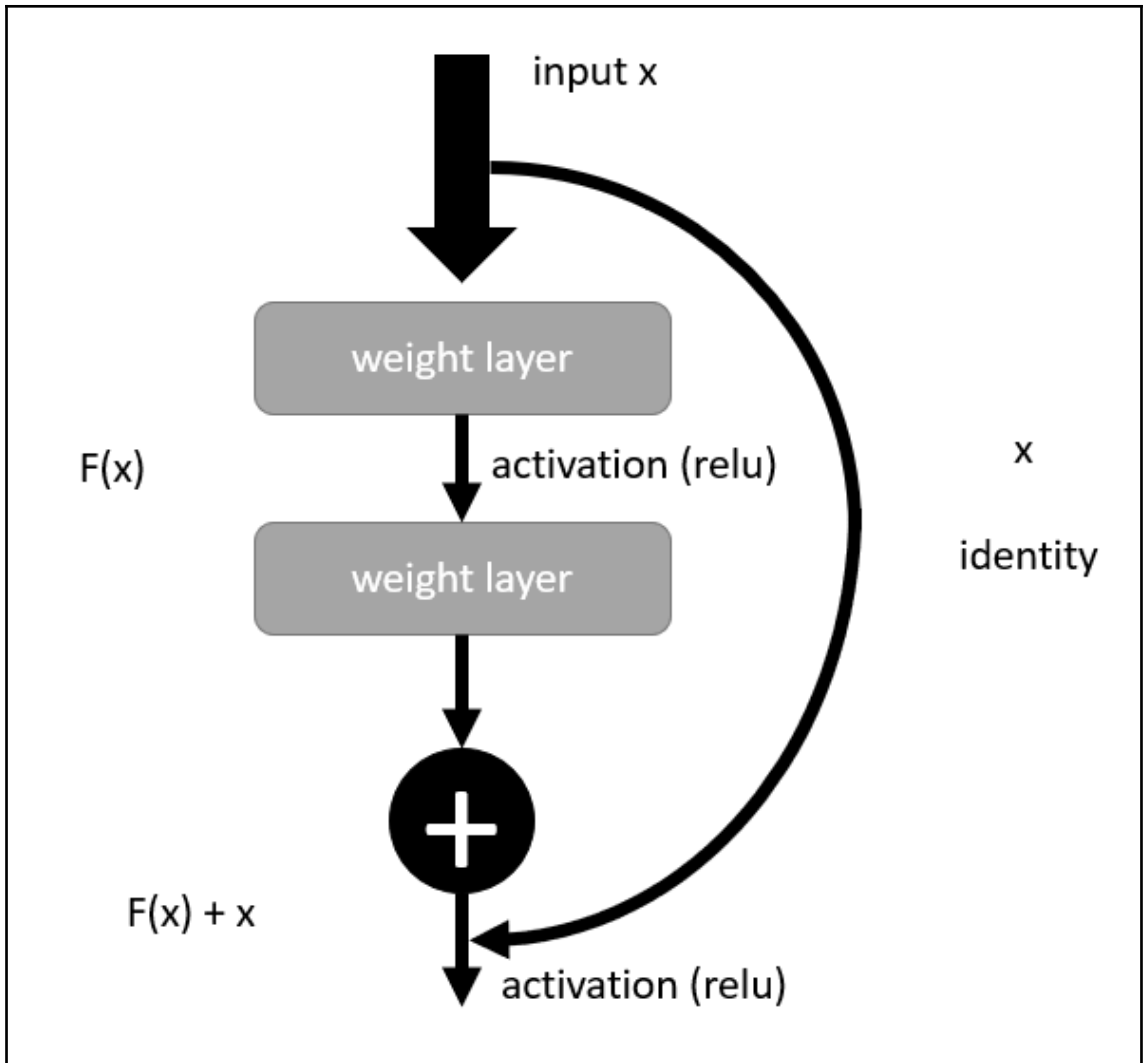
```
If you depend on functionality not listed there, please file an issue.
```



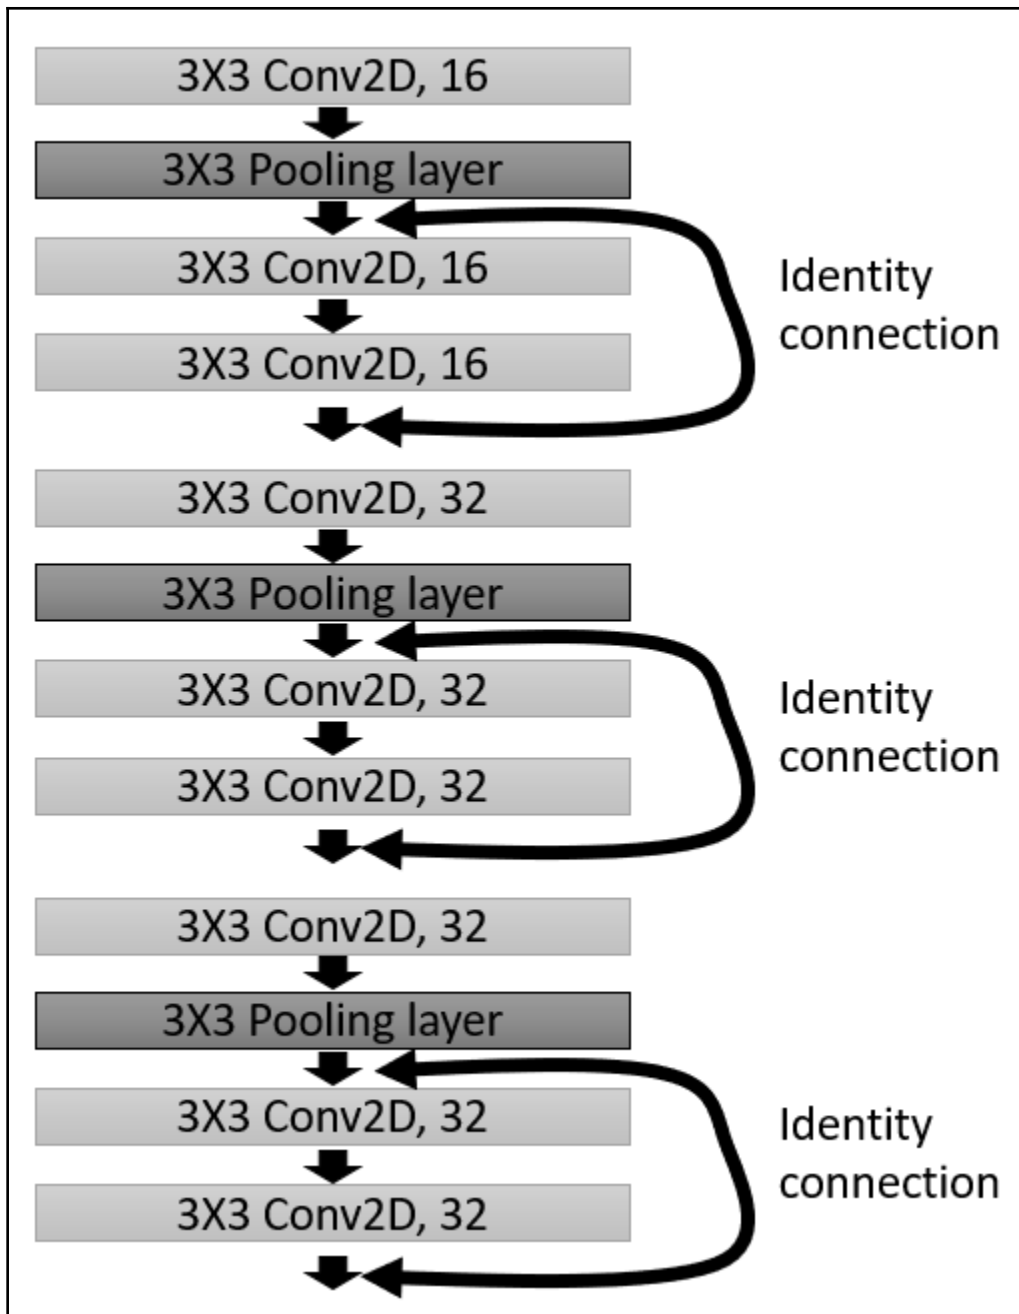
```
INFO:magents.trainers.CommandLineOptions(debug=False, num_runs=1, seed=-1, env_path=None, run_id='cob_1', load_model=False, train_del=True, save_freq=50000, keep_checkpoints=5, base_port=5005, num_envs=1, curriculum_folder=None, lesson=0, slow=False, no_graphics=False, multi_gpu=False, trainer_config_path='config/trainer_config.yaml', sampler_file_path=None, docker_target_name=None, env_args=one, cpu=False)
```


```
INFO:magents.envs:Start training by pressing the Play button in the Unity Editor.
```

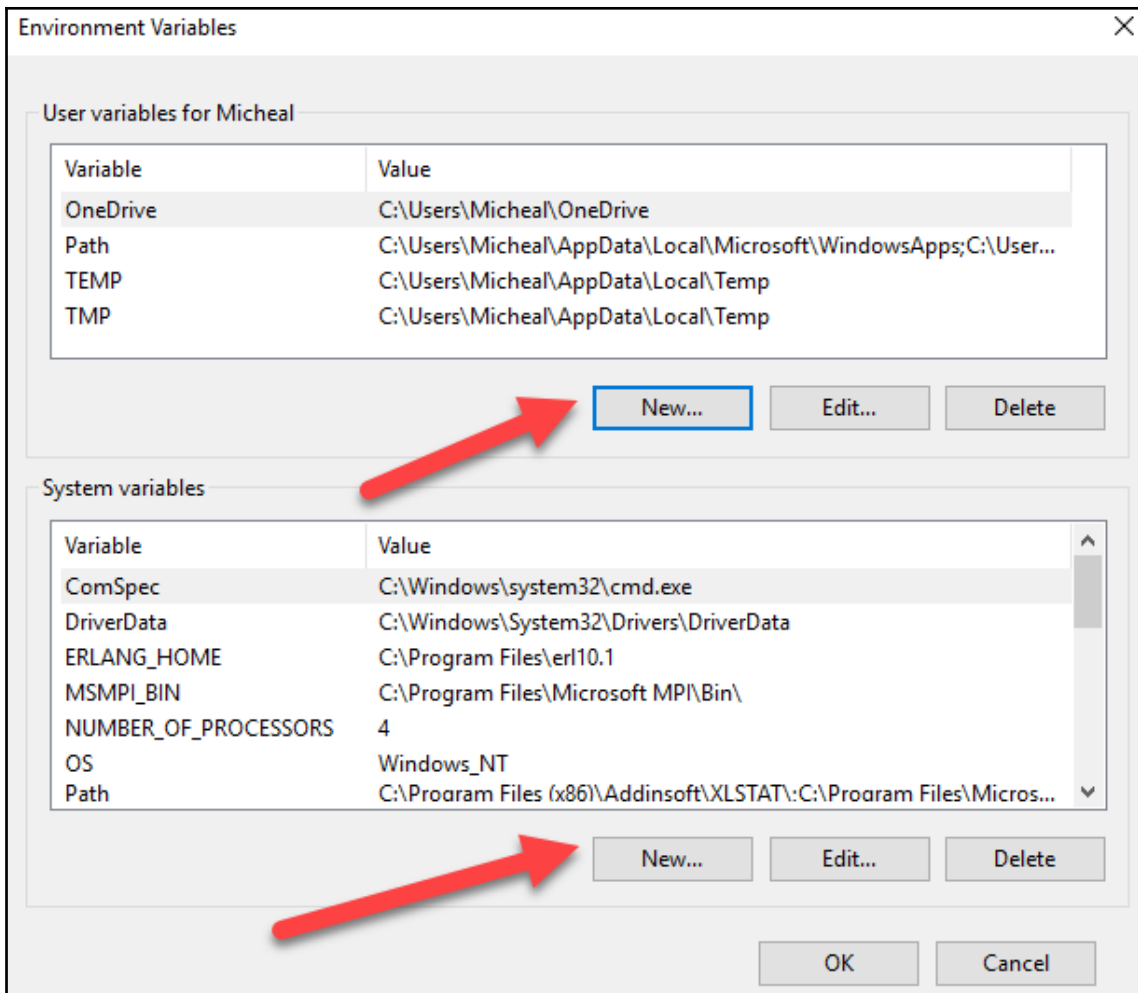














Locked Door

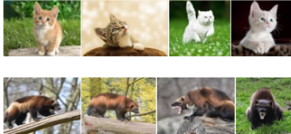




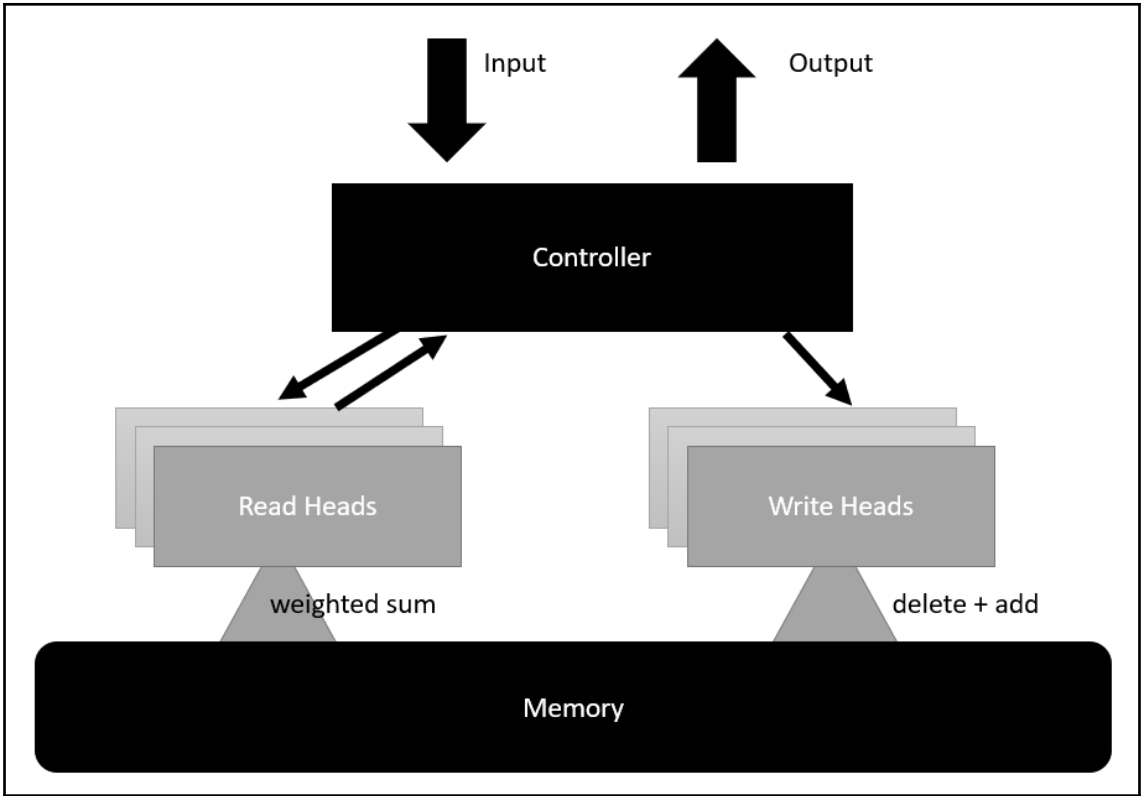
- Closed door
- Locked door
- Boxed door
- Open door
- Key
- Box
- Hurtle
- Orb
- Goal
- Box target
- Box undo

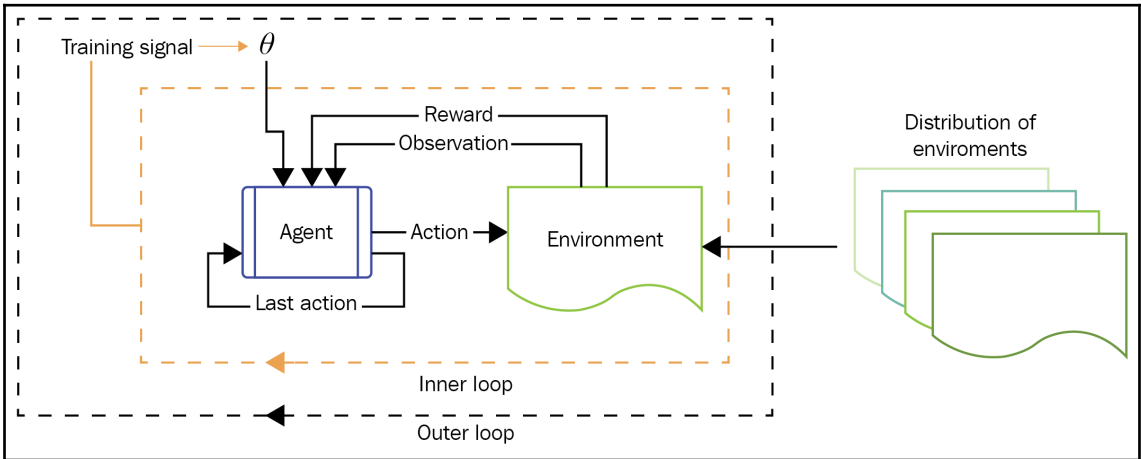
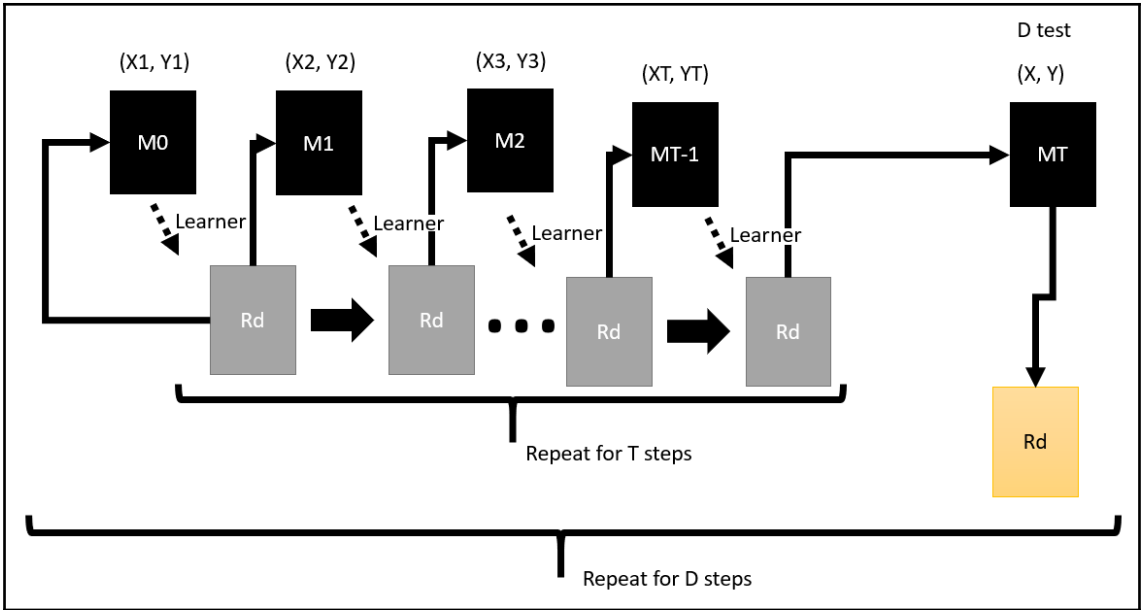
Save



# Chapter 14: From DRL to AGI

Training Data	Test Data
<p data-bbox="357 388 499 414">Cat or Wolverine</p> <div data-bbox="194 381 674 578"></div> <p data-bbox="357 617 499 643">Mouse or Moose</p> <div data-bbox="194 610 674 807"></div>	<p data-bbox="942 493 1085 520">Tiger or Bear</p> <div data-bbox="759 486 1239 684"></div>












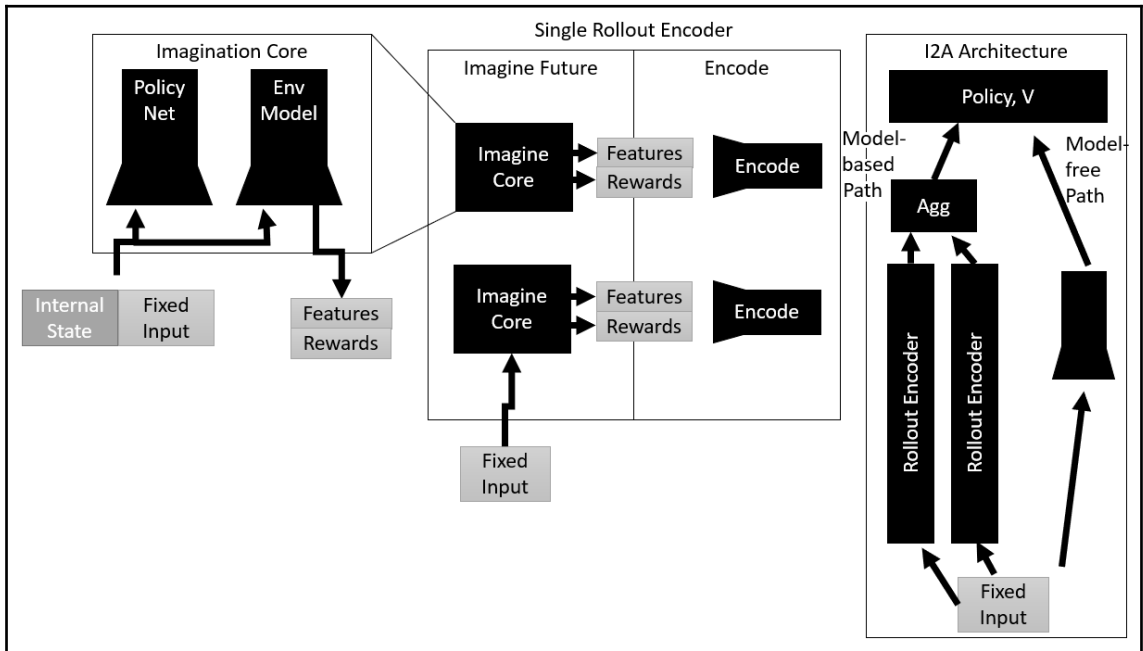
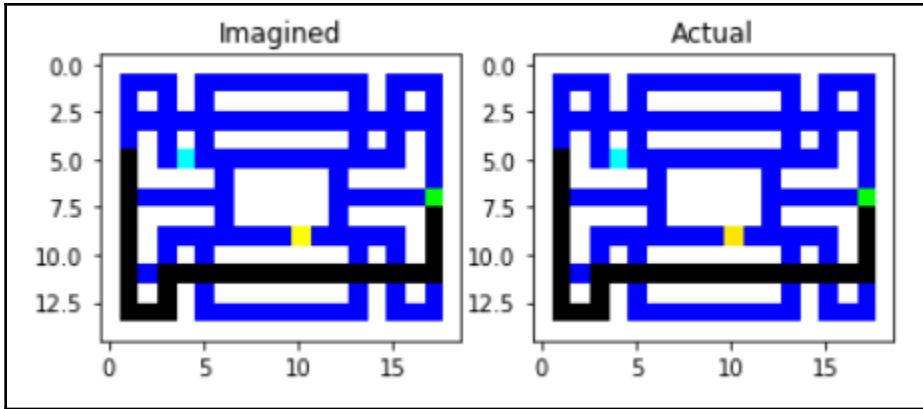
Mean Reward: -6.608695652173913. frame: 190000



loss







```
C:\ProgramData\Anaconda3\envs\i2a\python.exe
g: Implicit dimension choice for softmax has been deprecated. Change the call to include dim=X as an argument.
  imagined_state = F.softmax(imagined_state).max(1)[1].data.cpu()
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:203: UserWarn
g: Implicit dimension choice for softmax has been deprecated. Change the call to include dim=X as an argument.
  imagined_reward = F.softmax(imagined_reward).max(1)[1].data.cpu()
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:215: UserWarn
g: volatile was removed and now has no effect. Use `with torch.no_grad():` instead.
  action = self.distil_policy.act(autograd.Variable(state, volatile=True))
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\actor_critic.py:17: UserWarning:
Implicit dimension choice for softmax has been deprecated. Change the call to include dim=X as an argument.
  probs = F.softmax(logits)
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:295: UserWarn
g: volatile was removed and now has no effect. Use `with torch.no_grad():` instead.
  next_value = actor_critic(autograd.Variable(rollout.states[-1], volatile=True))
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\actor_critic.py:29: UserWarning:
Implicit dimension choice for softmax has been deprecated. Change the call to include dim=X as an argument.
  probs = F.softmax(logits)
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\actor_critic.py:30: UserWarning:
Implicit dimension choice for log_softmax has been deprecated. Change the call to include dim=X as an argument.
  log_probs = F.log_softmax(logits)
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:310: UserWarn
g: Implicit dimension choice for softmax has been deprecated. Change the call to include dim=X as an argument.
  distil_loss = 0.01 * (F.softmax(logits).detach() * F.log_softmax(distil_logits)).sum(1).mean()
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:310: UserWarn
g: Implicit dimension choice for log_softmax has been deprecated. Change the call to include dim=X as an argument.
  distil_loss = 0.01 * (F.softmax(logits).detach() * F.log_softmax(distil_logits)).sum(1).mean()
C:\Users\Micheal\Dropbox\Books\Hands-on Game AI with Python\Code\Chapter_14\Chapter_14\Chapter_14_I2A.py:310: UserWarn
g: torch.nn.utils.clip_grad_norm is now deprecated in favor of torch.nn.utils.clip_grad_norm_
  nn.utils.clip_grad_norm(actor_critic.parameters(), max_grad_norm)
0% | 2/1000000 [01:11<10920:15:46, 39.31s/it]
```