

Chapter 1: Deep Learning Walkthrough and PyTorch Introduction

PyTorch

Get Started Features Ecosystem Blog Tutorials Docs Resources GitHub

QUICK START LOCALLY

Select your preferences and run the install command. Stable represents the most currently tested and supported version of PyTorch 1.0. This should be suitable for many users. Preview is available if you want the latest, not fully tested and supported, 1.0 builds that are generated nightly. Please ensure that you have met the prerequisites below (e.g., `numpy`), depending on your package manager. Anaconda is our recommended package manager since it installs all dependencies. You can also [install previous versions of PyTorch](#). Note that LibTorch is only available for C++.

PyTorch Build	Stable (1.0)	Preview (Nightly)			
Your OS	Linux	Mac	Windows		
Package	Conda	Pip	LibTorch	Source	
Language	Python 2.7	Python 3.5	Python 3.6	Python 3.7	C++
CUDA	8.0	9.0	10.0	None	

Run this Command:


```
pip3 install torch torchvision
```

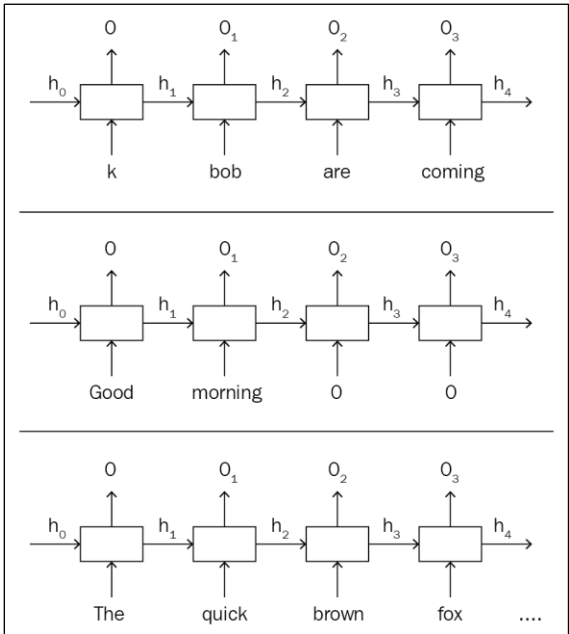


[Previous versions of PyTorch](#)


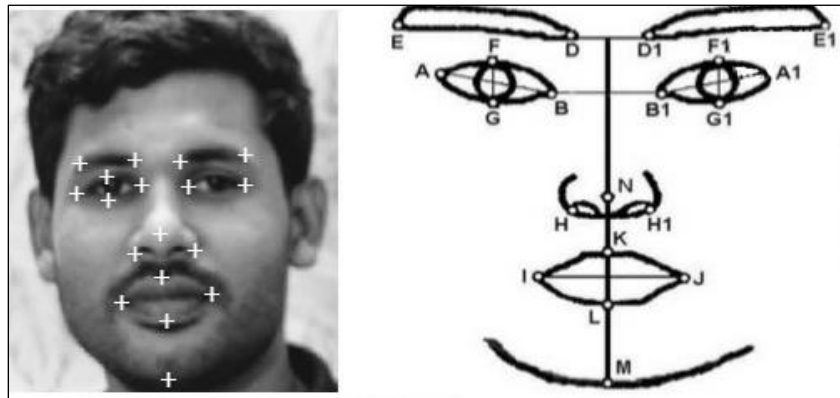

QUICK START WITH CLOUD PARTNERS

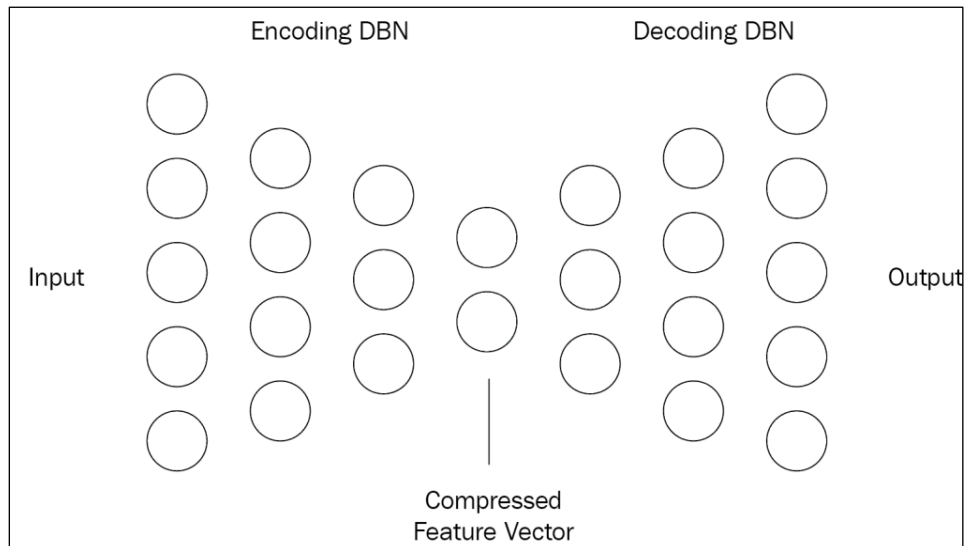
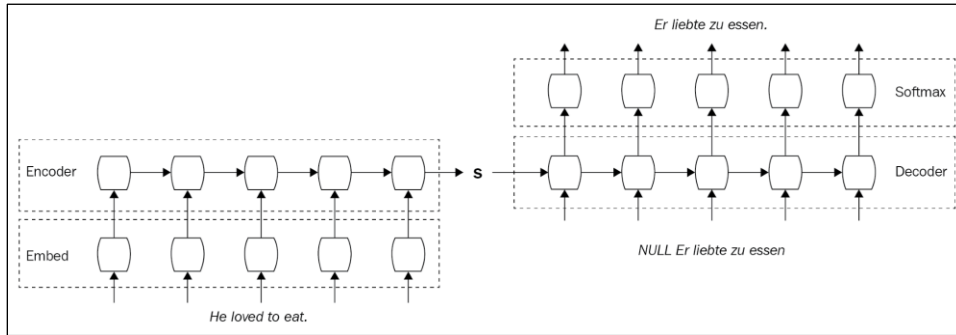

Get up and running with PyTorch quickly through popular cloud platforms and machine learning services.

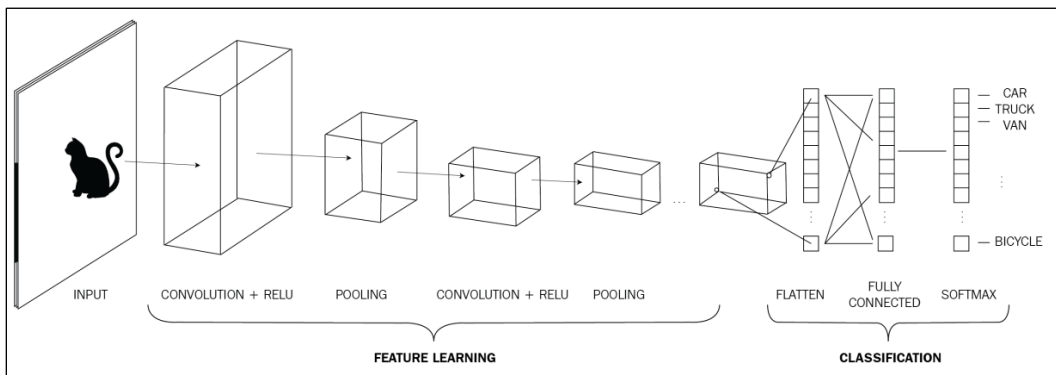
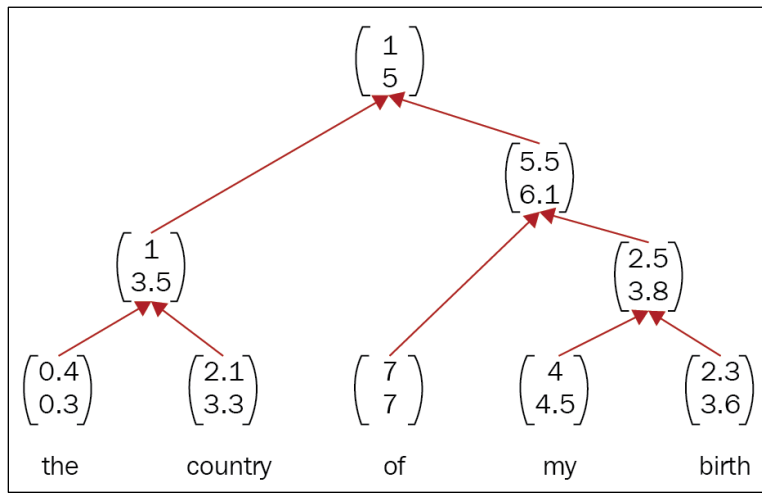
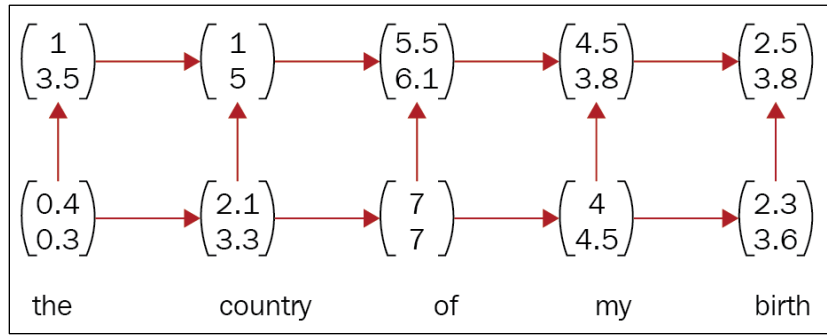
- AWS
- Google Cloud Platform
- Microsoft Azure

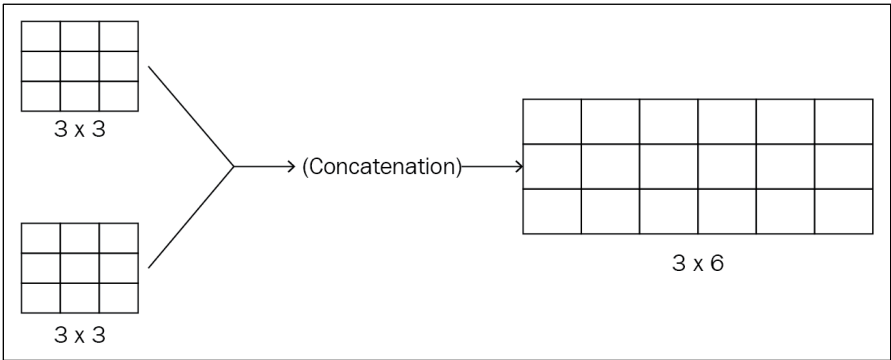
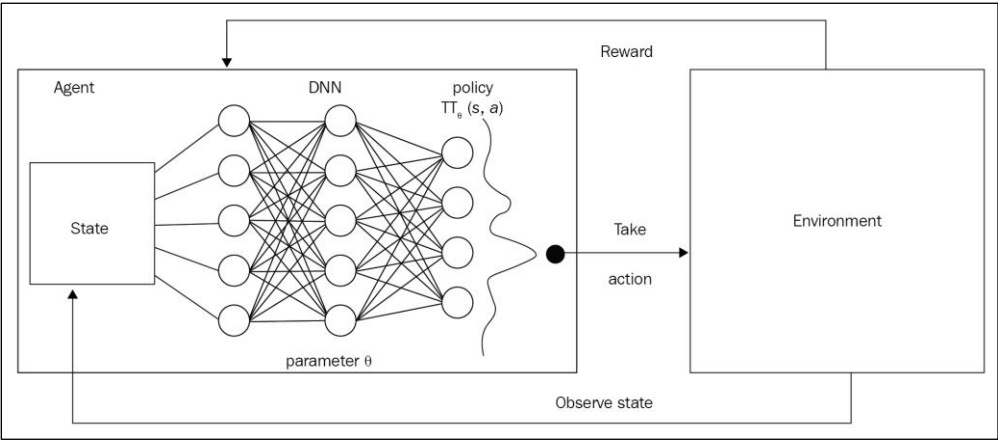
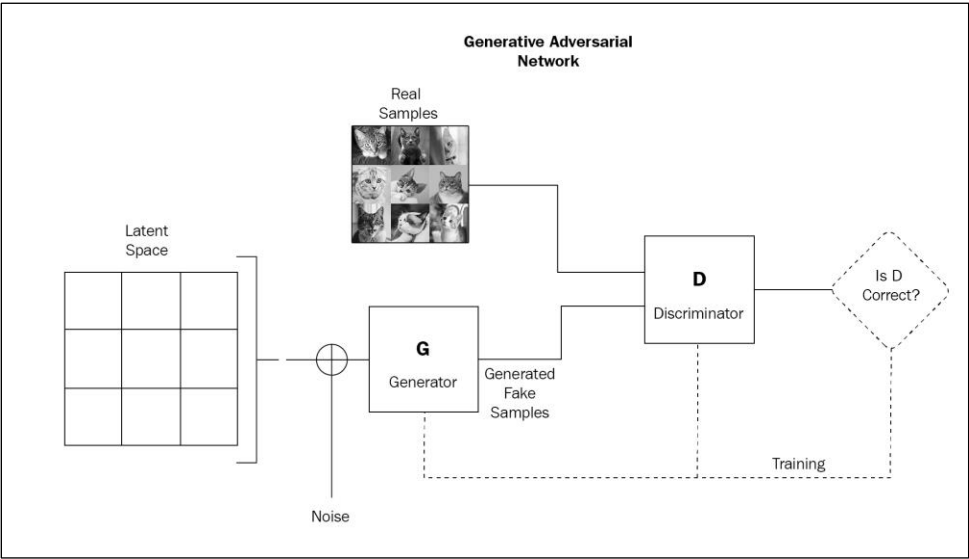


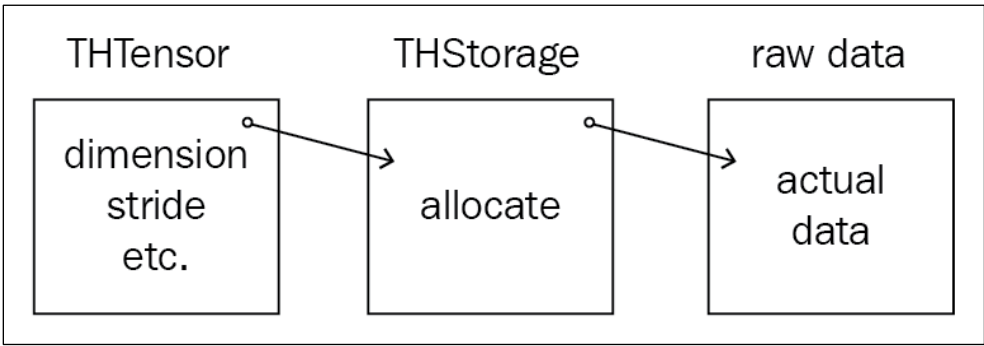






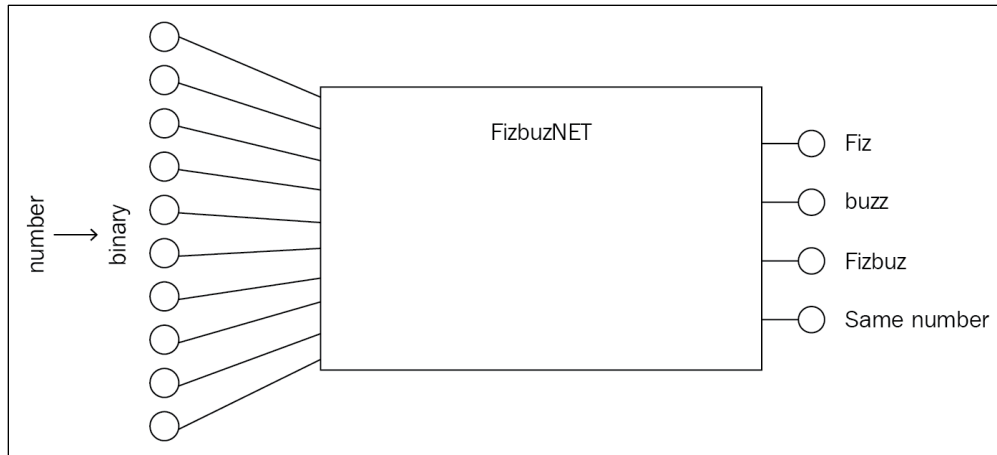









Chapter 2: A Simple Neural Network





$$\theta_1 := \theta_1 - \alpha \frac{\partial}{\partial \theta_1} J(\theta_1)$$

if α is too small, gradient decent can be slow.

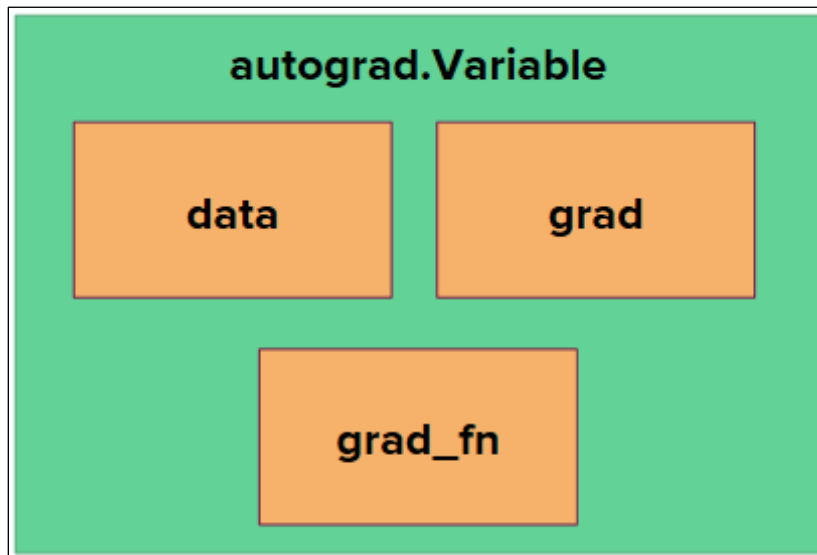



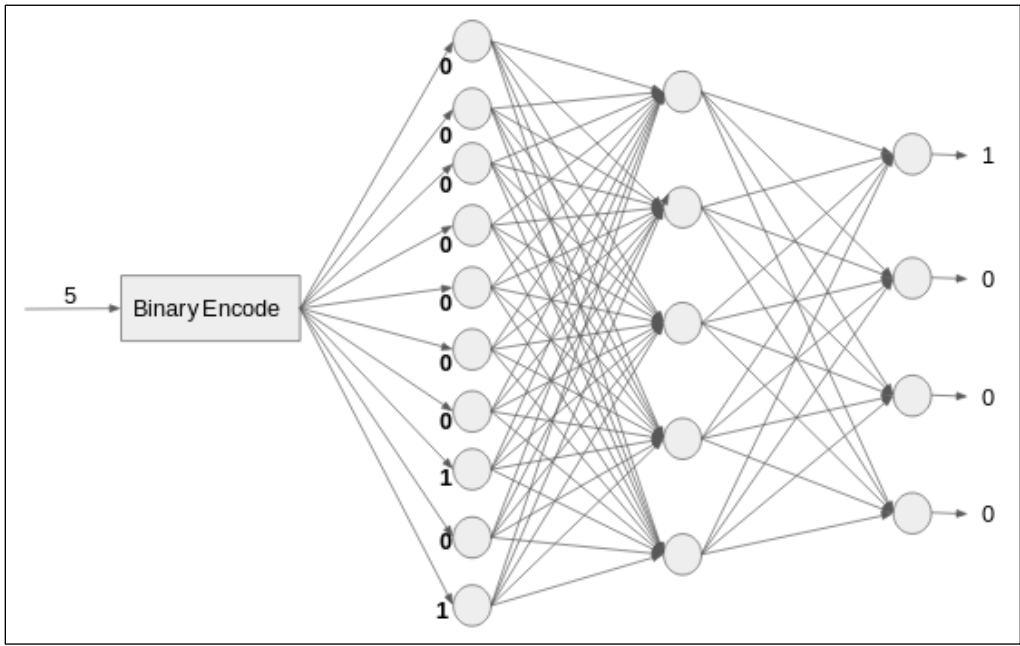
if α is too large, gradient decent can overshoot the minimum. It may fail to converge, or even diverge.

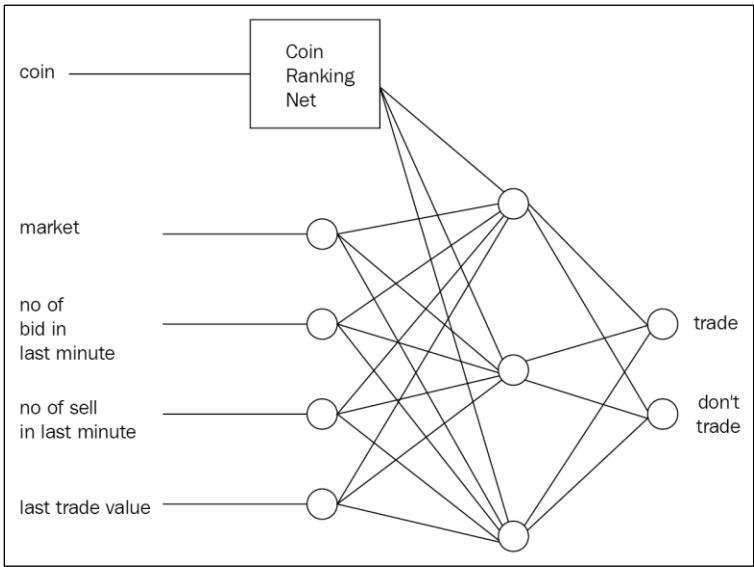
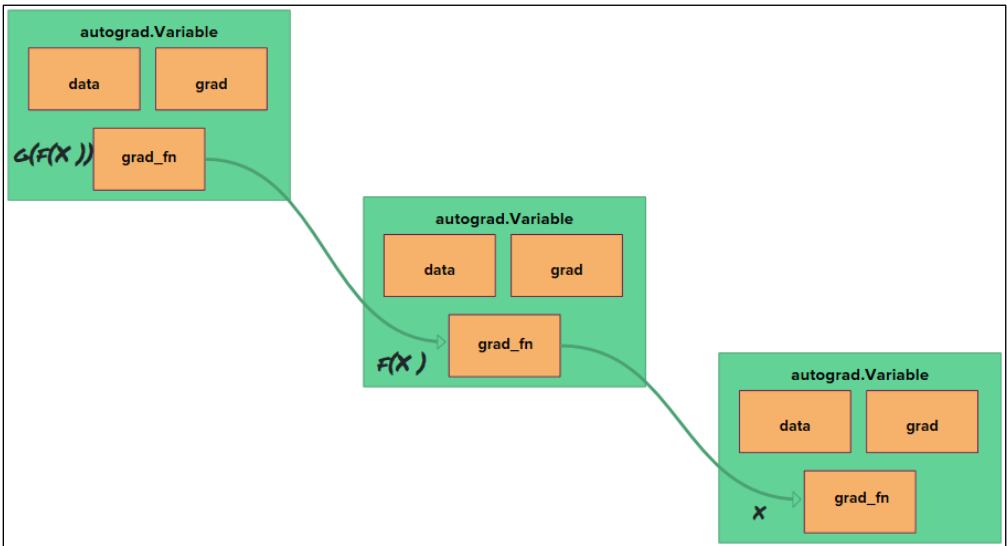




$$\text{weight gradient} = \frac{\partial t}{\partial \text{weight}} = \frac{\partial t}{\partial \text{weight}} * \frac{\partial \text{out}}{\partial t} \times \frac{\partial \text{out}}{\partial \text{out}}$$






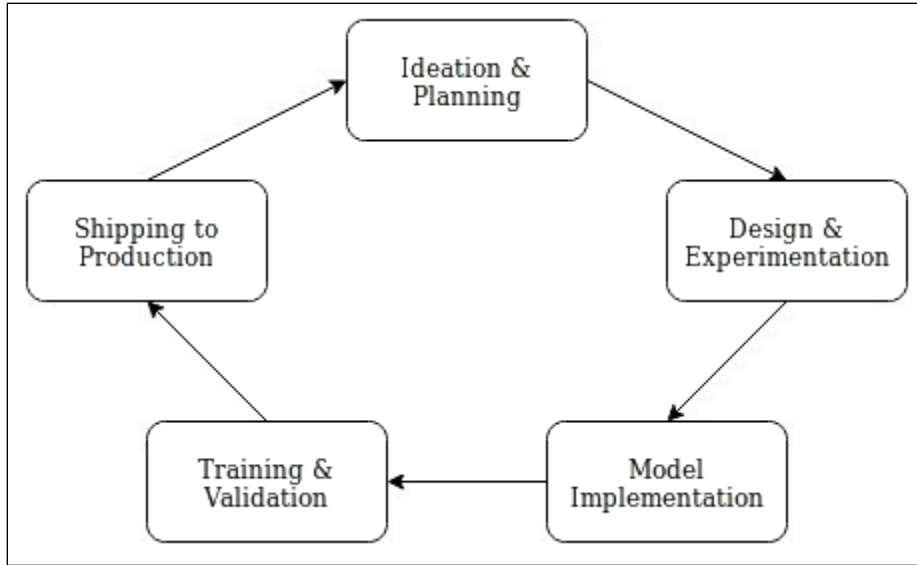


Sigmoid

$$\sigma(z) = \frac{1}{1+e^{-z}}$$

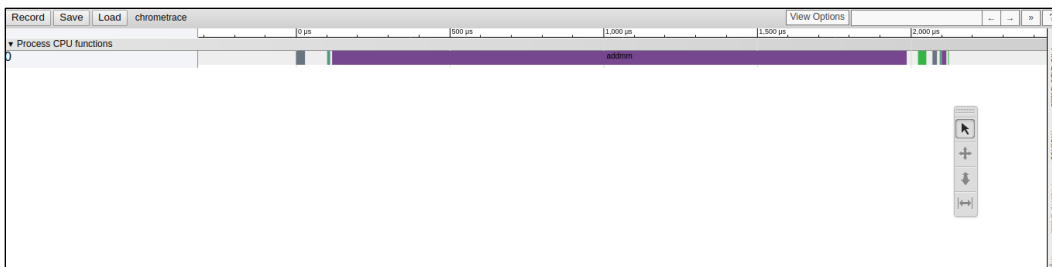


Chapter 3: Deep Learning Workflow



Name	CPU time	CUDA time	Calls	CPU total	CUDA total
t	15.758us	0.000us	2	31.516us	0.000us
expand	6.203us	0.000us	2	12.405us	0.000us
addmm	898.371us	0.000us	2	1796.742us	0.000us
sigmoid	14.462us	0.000us	2	28.923us	0.000us

Name	CPU time	CUDA time	Calls	CPU total	CUDA total
sigmoid	2.524us	0.000us	1	2.524us	0.000us
expand	4.325us	0.000us	1	4.325us	0.000us
t	5.140us	0.000us	1	5.140us	0.000us
expand	9.168us	0.000us	1	9.168us	0.000us
addmm	16.646us	0.000us	1	16.646us	0.000us
sigmoid	29.335us	0.000us	1	29.335us	0.000us
t	30.314us	0.000us	1	30.314us	0.000us
addmm	1858.968us	0.000us	1	1858.968us	0.000us



```

-----
Environment Summary
-----
PyTorch 2018.05.07 compiled w/ CUDA 8.0.61
Running with Python 3.6 and

`pip3 list` truncated output:
msgpack-numpy (0.4.1)
numpy (1.14.2)
torch (2018.5.7, /home/sherin/miniconda3/lib/python3.6/site-packages)
torchaudio (0.1, /home/sherin/mypro/audio)
torchtext (0.2.3)
torchvision (0.2.1)

```

```

-----
autograd profiler output (CPU mode)
-----
top 15 events sorted by cpu_time_total
-----
Name                CPU time      CUDA time      Calls      CPU total     CUDA total
-----
AddmmBackward       233.993us     0.000us        1          233.993us     0.000us
t                   100.119us     0.000us        1          100.119us     0.000us
addmm               79.638us     0.000us        1          79.638us     0.000us
_cast_float         70.043us     0.000us        1          70.043us     0.000us
mm                 61.899us     0.000us        1          61.899us     0.000us
_mm                53.449us     0.000us        1          53.449us     0.000us
AddmmBackward       47.908us     0.000us        1          47.908us     0.000us
zeros_like          46.894us     0.000us        1          46.894us     0.000us
ExpandBackward      46.219us     0.000us        1          46.219us     0.000us
uniform_            44.824us     0.000us        1          44.824us     0.000us
addmm               42.784us     0.000us        1          42.784us     0.000us
sqrt                40.520us     0.000us        1          40.520us     0.000us
sqrt                39.984us     0.000us        1          39.984us     0.000us
sum                 39.054us     0.000us        1          39.054us     0.000us
MseLossBackward    39.017us     0.000us        1          39.017us     0.000us

```

```


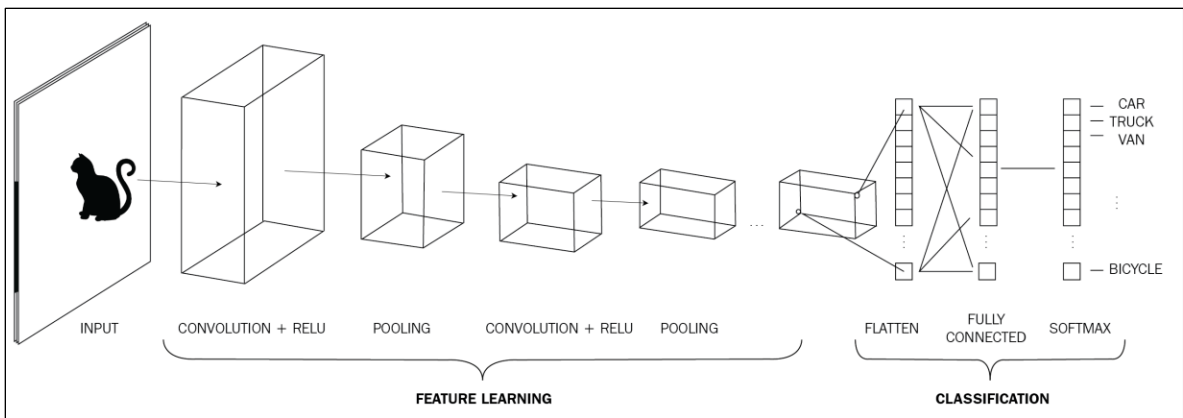
-----
cProfile output
-----
6532 function calls (6524 primitive calls) in 0.149 seconds

Ordered by: internal time
List reduced from 81 to 15 due to restriction <15>

ncalls  tottime  percall  cumtime  percall  filename:lineno(function)
1      0.071    0.071    0.071    0.071    {method 'run_backward' of 'torch._C_EngineBase' objects}
2      0.015    0.008    0.015    0.008    {method 'sigmoid' of 'torch._C_TensorBase' objects}
8      0.015    0.002    0.015    0.002    {method 'add_' of 'torch._C_TensorBase' objects}
1      0.011    0.011    0.011    0.011    {built-in method ones_like}
1000   0.008    0.000    0.008    0.000    bottleneck_support.py:16(<listcomp>)
2      0.005    0.003    0.005    0.003    {built-in method addmm}
2      0.004    0.002    0.004    0.002    {built-in method numpy.core.multiarray.array}
1000   0.004    0.000    0.013    0.000    bottleneck_support.py:15(wrapper)
1      0.003    0.003    0.003    0.003    {built-in method torch._C_nn.mse_loss}
1      0.003    0.003    0.021    0.021    bottleneck_support.py:39(get_data)
4      0.002    0.001    0.002    0.001    {method 'sqrt' of 'torch._C_TensorBase' objects}
1004   0.001    0.000    0.001    0.000    {method 'format' of 'str' objects}
2      0.001    0.000    0.001    0.000    {method 'type' of 'torch._C_TensorBase' objects}
1      0.001    0.001    0.149    0.149    bottleneck_support.py:1(<module>)
1      0.001    0.001    0.001    0.001    {method 'permutation' of 'mtrand.RandomState' objects}

```

Chapter 4: Computer Vision

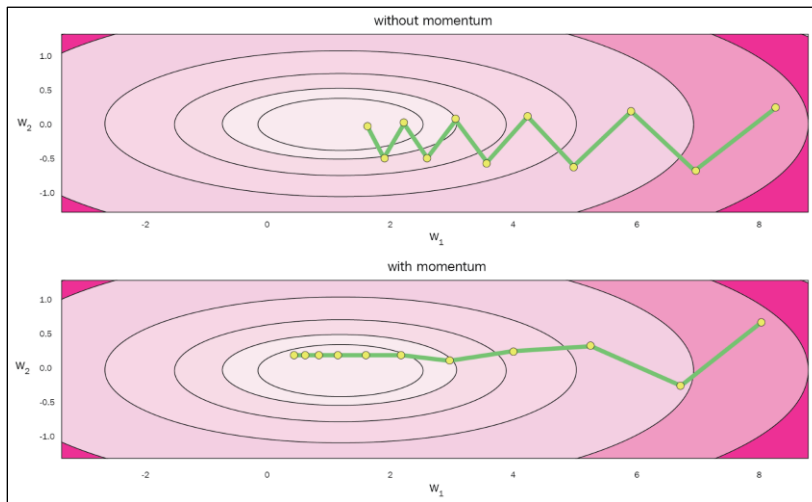




	1	2	2	3	4	4	2
	3	3	1	4	2	5	8
	1	9	5	9	6	6	7
	1	7	4	1	8	4	5
	7	9	7	5	3	2	8
	6	1	3	6	5	4	3
	3	6	2	8	7	5	2

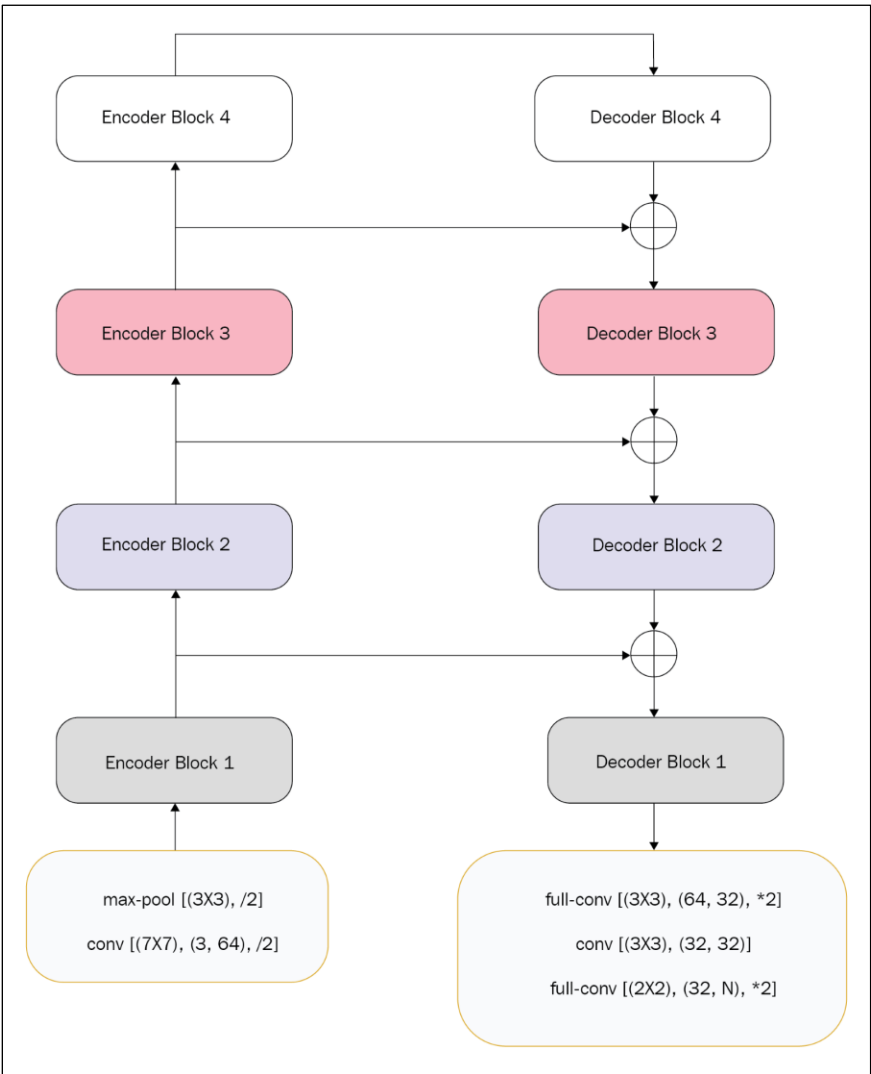
$$\begin{aligned}
 &1 \times 0 + \\
 &3 \times 0 + \\
 &4 \times 0 + \\
 &7 \times 0 + \\
 &6 \times 1 + = 61 \rightarrow \\
 &2 \times 2 + \\
 &5 \times 0 + \\
 &8 \times 3 + \\
 &9 \times 3 \\
 &\downarrow
 \end{aligned}$$

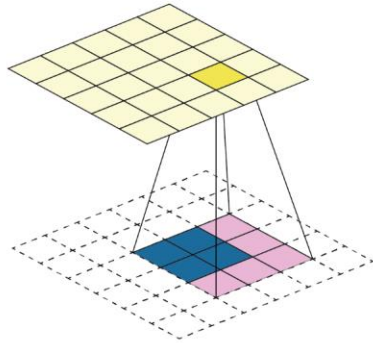
7 x 7 Input Volume

5 x 5 Output Volume












No padding, no strides, transposed




Arbitrary padding, no strides, transposed




Half padding, no strides, transposed



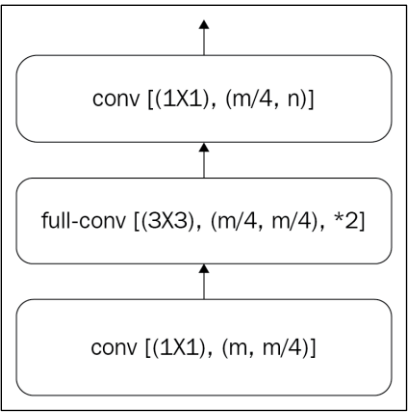
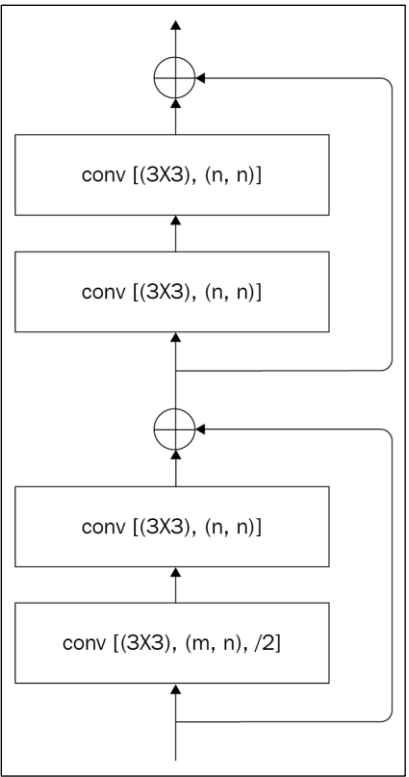
Full padding, no strides, transposed



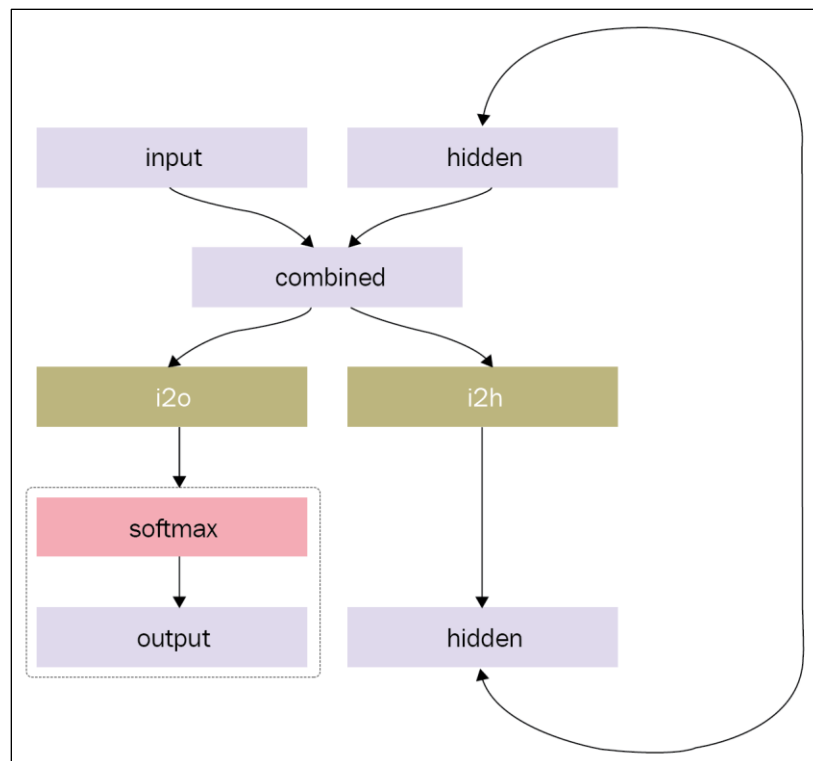
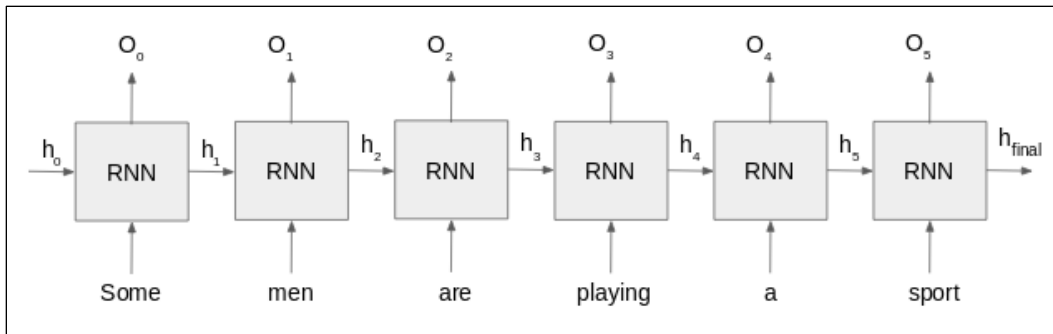

No padding, strides, transposed

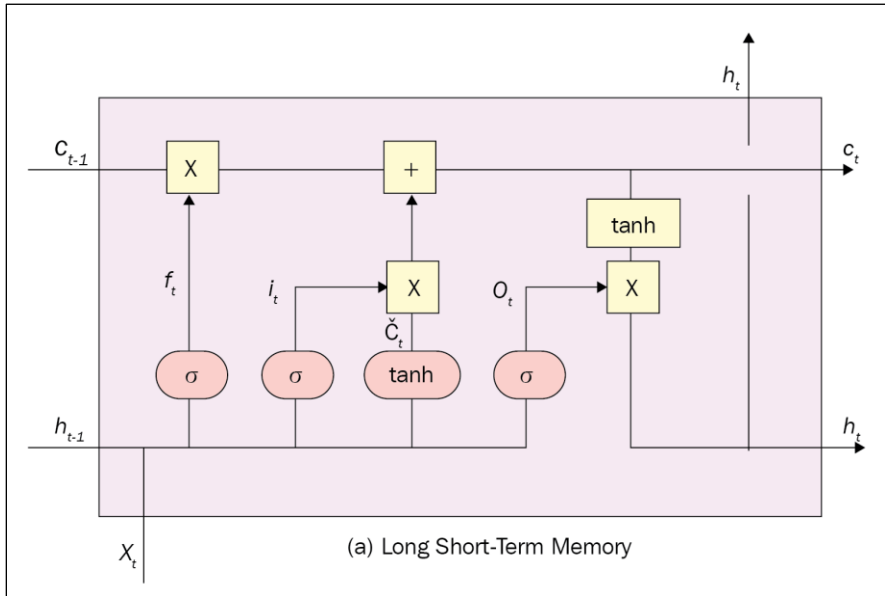

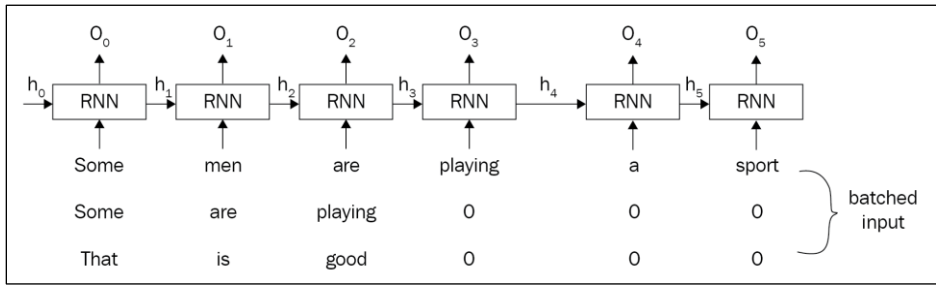


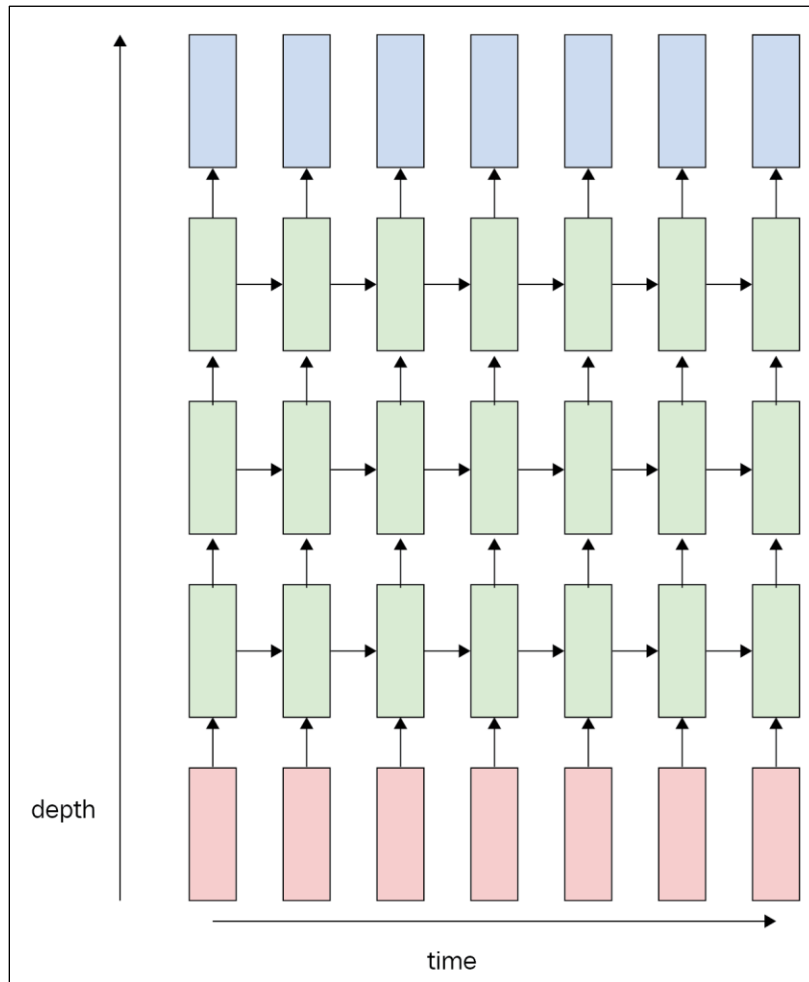
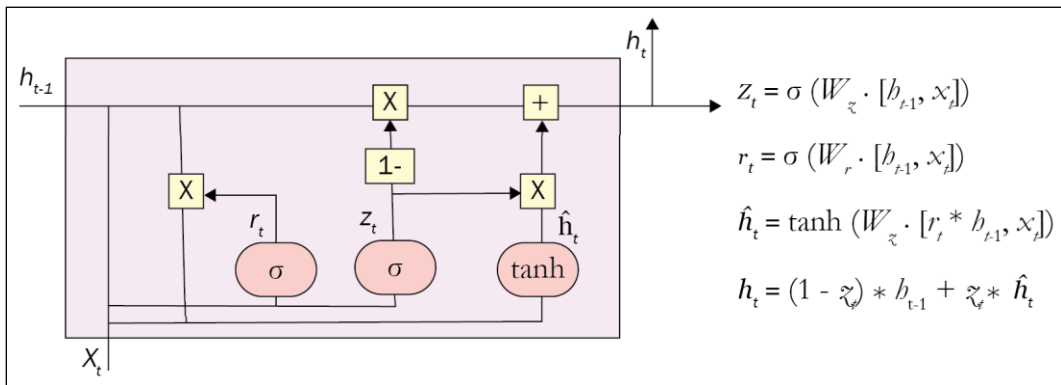
Padding, strides, transposed

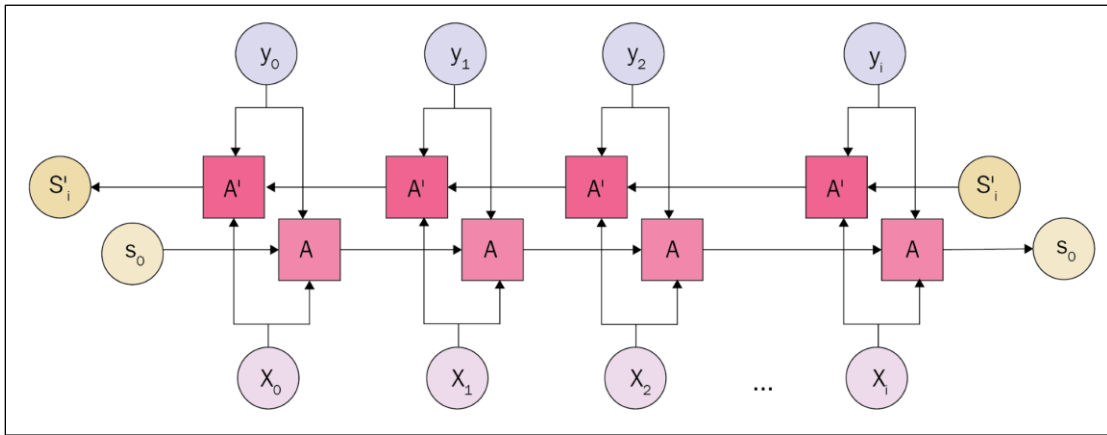


Chapter 5: Sequential Data Processing







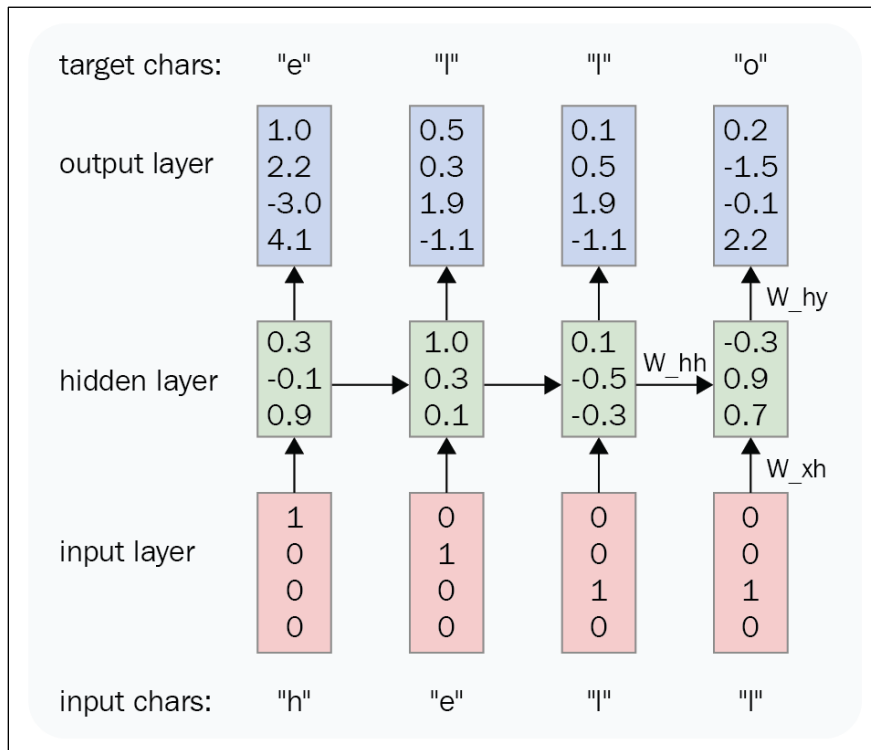
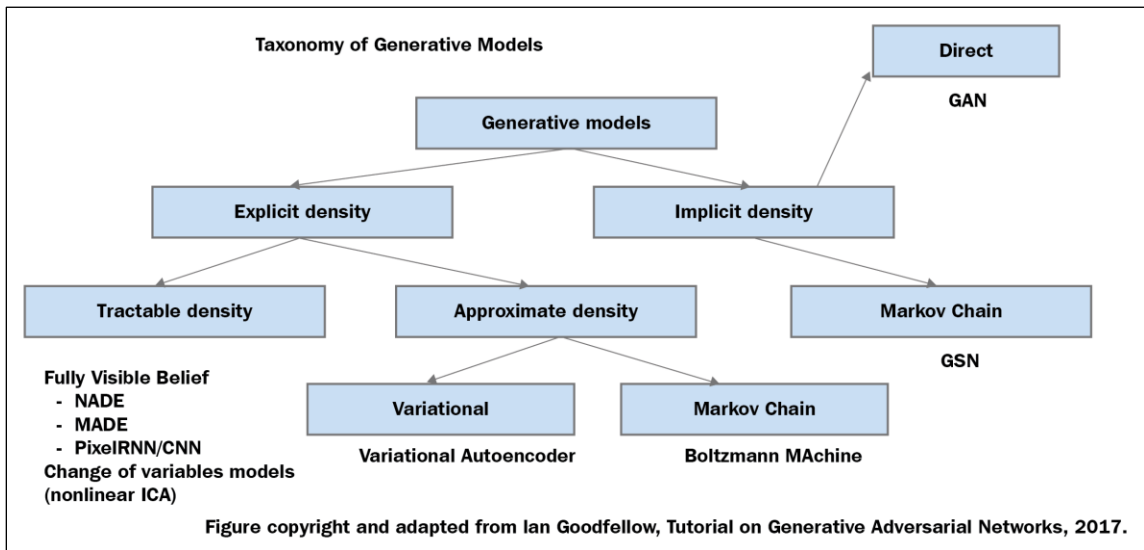




Reduce for each blank space
Shift for each character

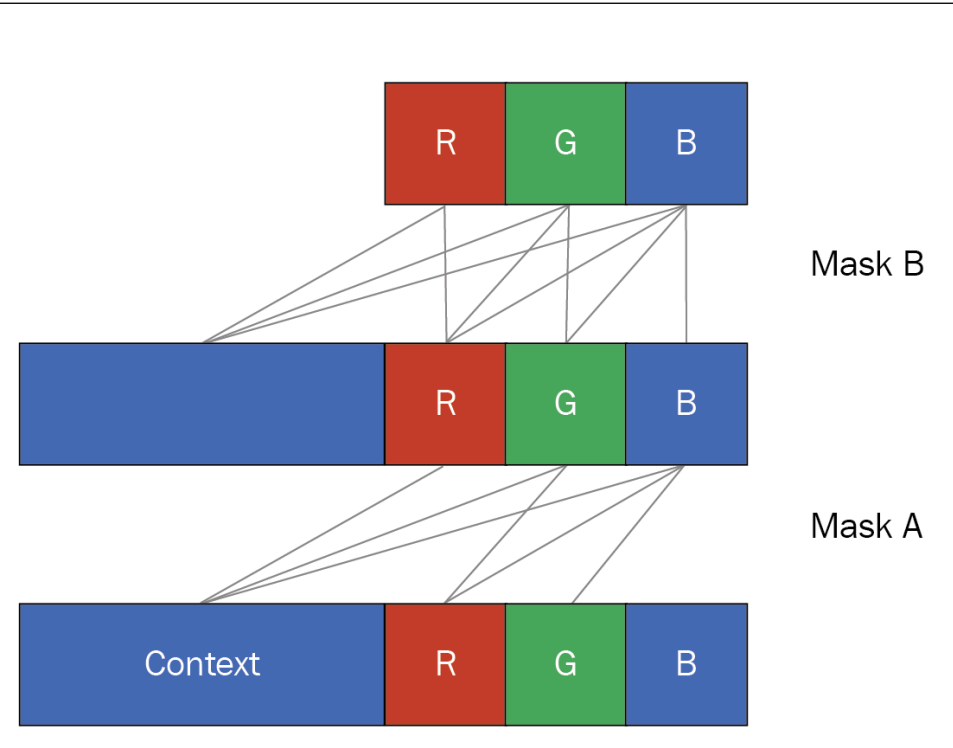

<u>input</u>	<u>stack</u>
HE_LL_O_ _ _ _	
E_LL_O_ _ _ _	H
_LL_O_ _ _ _	E H
LL_O_ _ _ _	H
L_O_ _ _ _	L H
O _ _ _	L L H
O_ _ _ _	L H
_ _ _ _	O L H
_ _ _	L H
_ _	L H
_	H


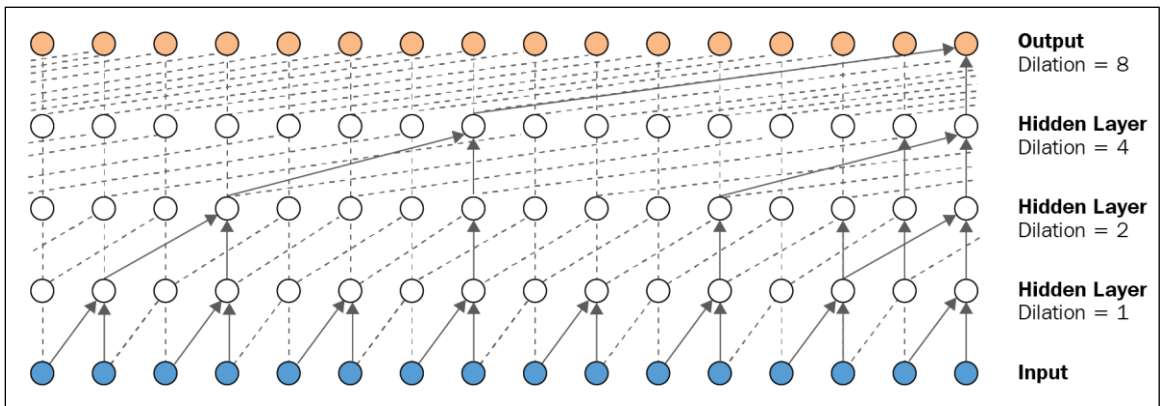

Chapter 6: Generative Networks



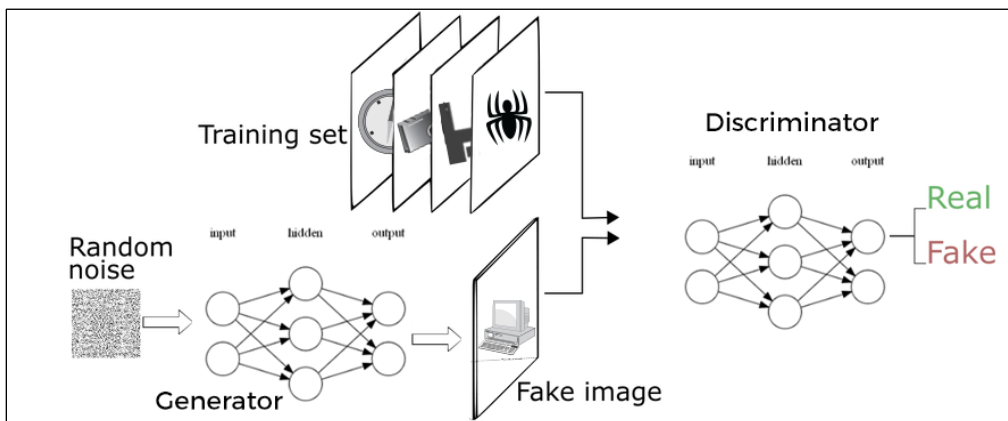
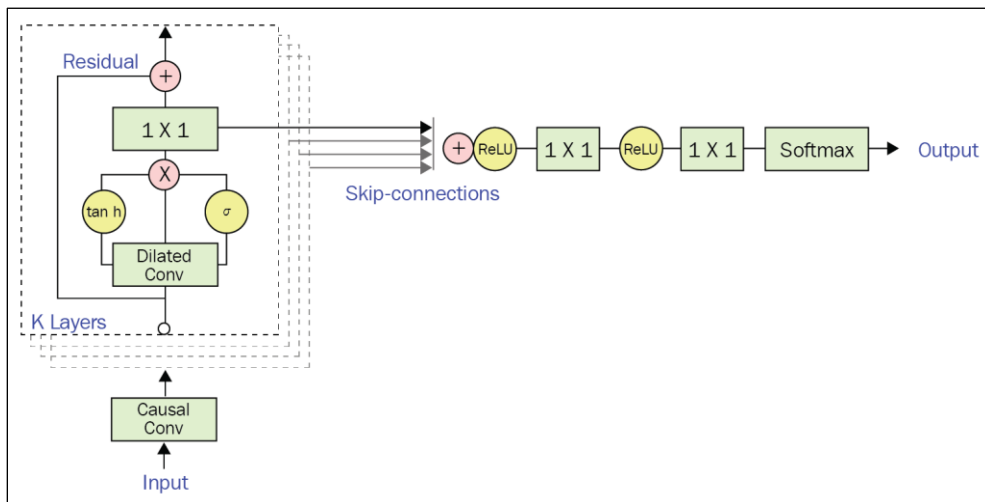







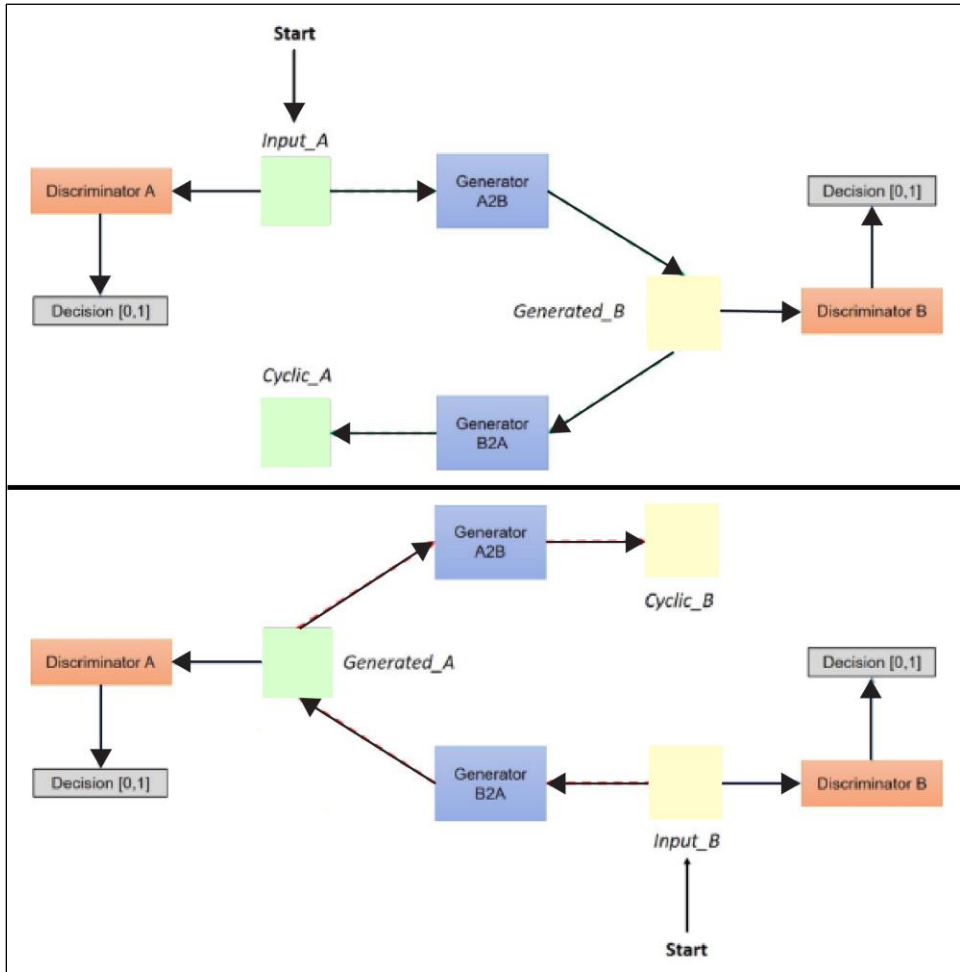
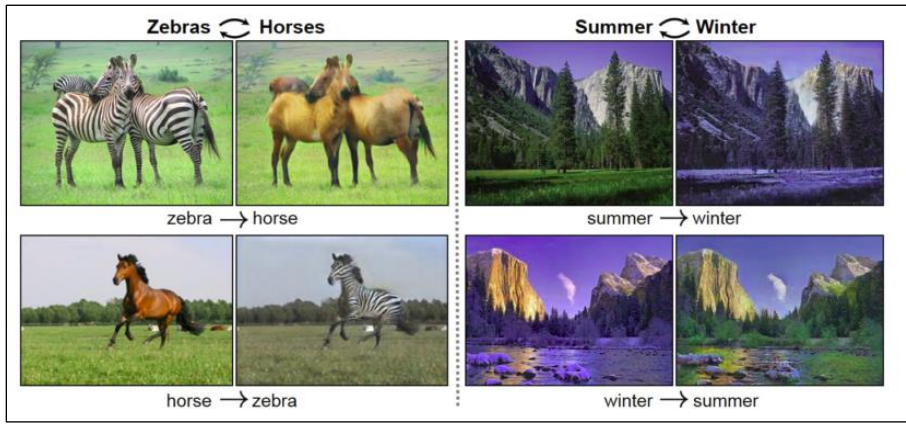





1	0	1	0	1
0	0	0	0	0
1	0	1	0	1
0	0	0	0	0
1	0	1	0	1








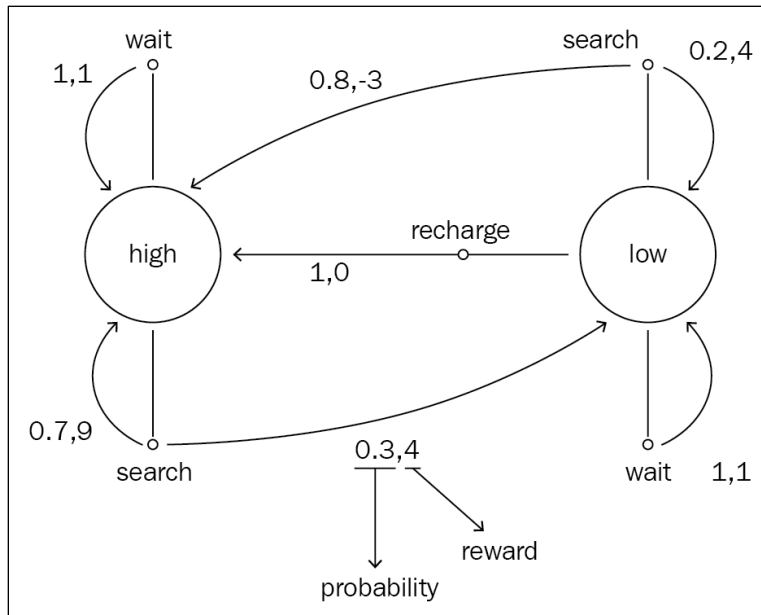

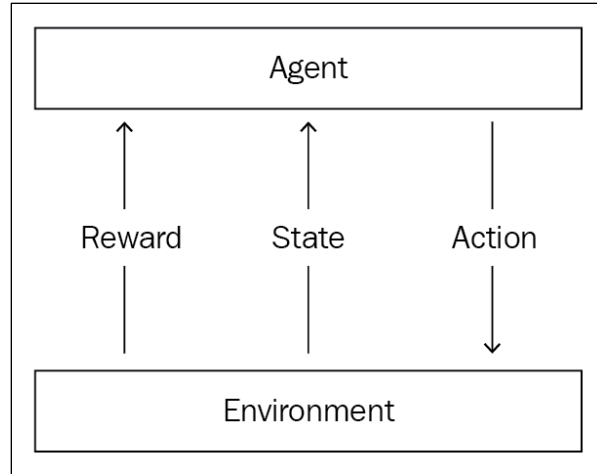
Batch Norm





Instance Norm



Chapter 7: Reinforcement Learning





```

class CartPoleEnv(gym.Env):
    """
    Description:
        A pole is attached by an un-actuated joint to a cart, which moves along a frictionless track. The pendulum starts
        Source:
        This environment corresponds to the version of the cart-pole problem described by Barto, Sutton, and Anderson

    Observation:
        Type: Box(4)
        Num Observation      Min      Max
        0   Cart Position    -4.8     4.8
        1   Cart Velocity    -Inf     Inf
        2   Pole Angle       -24°     24°
        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
    """

```

```

tensor([[2.0429, 1.4886],
        [1.2952, 1.2798],
        [1.1960, 1.1665],
        [1.3114, 1.1780],
        [1.2970, 1.2814],
        [1.4016, 1.4096],
        [1.5460, 1.2322],
        [2.1189, 1.5717],
        [1.4563, 1.1823],
        [1.2912, 1.2759],
        [2.0797, 1.6504],
        [1.2814, 1.2050],
        [1.3184, 1.3216],
        [1.3782, 1.3824],
        [1.4194, 1.4275],
        [1.4445, 1.1700]])

```

```

tensor([[0.9818],
        [0.8832],
        [1.2682],
        [0.9230],
        [0.9572],
        [0.8275],
        [1.0659],
        [1.1392],
        [1.2381],
        [1.1048],
        [0.9397],
        [0.8558],
        [1.0015],
        [1.0669],
        [1.0863],
        [1.1538],
        [1.0786],
        [0.9248],
        [0.9540],
        [0.9916]])

```

```

tensor([1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
        1, 1, 1, 1, 1, 1, 1, 1, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
        1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1], dtype=torch.uint8)


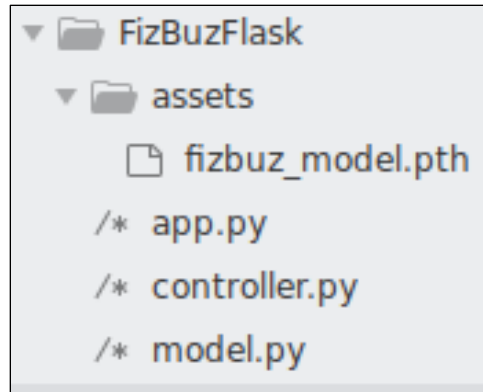
```

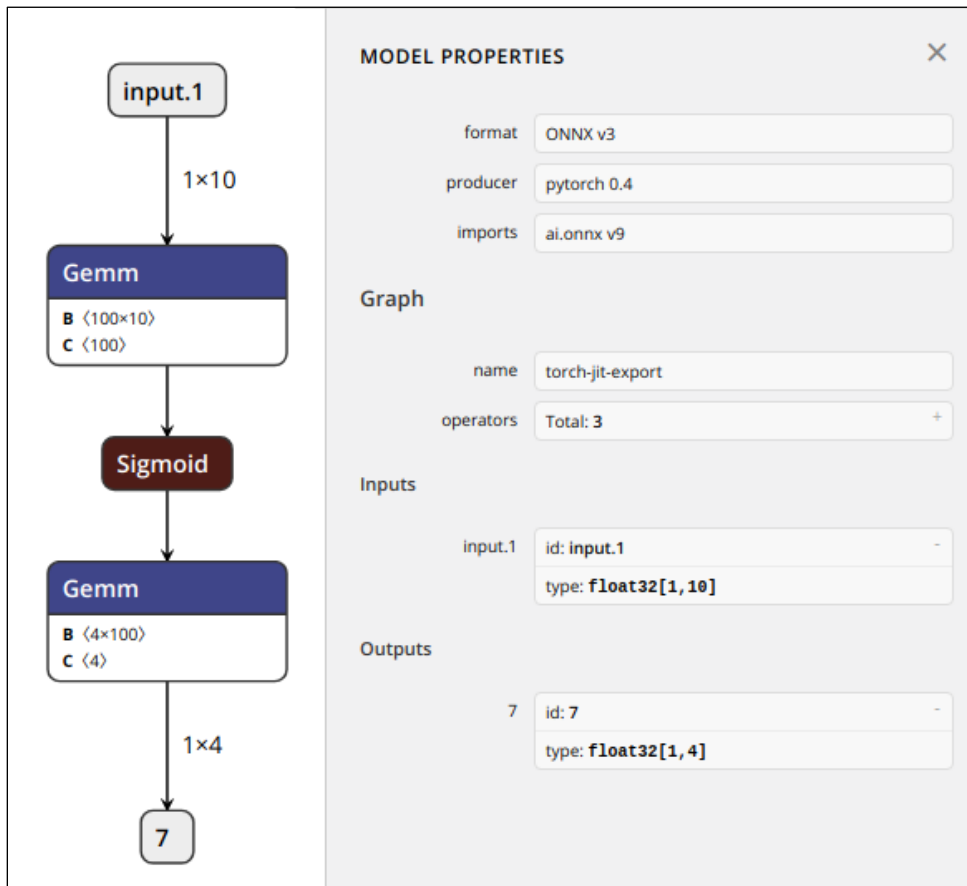
```

tensor([-0.0286, -0.0289, -0.0286, -0.0287, -0.0287, -0.0287, -0.0281, -0.0285,
        -0.0285, -0.0285, -0.0285, -0.0284, -0.0285, -0.0281, -0.0288, -0.0280,
        -0.0281, -0.0286, -0.0283, -0.0285, -0.0281, -0.0289, -0.0282, -0.0285,
        -0.0286, -0.0281, -0.0288, -0.0284, -0.0284, -0.0281, -0.0288, -0.0280,
        -0.0282, -0.0291, -0.0285, -0.0282, -0.0287, -0.0288, -0.0287, -0.0286,
        0.0000, -0.0284, -0.0285, -0.0283, -0.0289, -0.0282, -0.0286, -0.0285,
        -0.0283, -0.0286, -0.0285, -0.0284, -0.0288, -0.0287, -0.0283, -0.0280])

```

Chapter 8: PyTorch to Production





- ▼ **fizbuz_package**
- 📄 fizbuz.onnx
- /* fizbuz_service.py
- /* signature.json

The image shows a browser window with the address bar containing `localhost:8089`. The page features the Locust logo on the left and the text "HOST http://localhost:8080" and "STATUS READY" on the right. The main content area has a dark green background and contains the following elements:

- Section header: **Start new Locust swarm**
- Label: "Number of users to simulate" above a white input field.
- Label: "Hatch rate (users spawned/second)" above another white input field.
- A green button labeled "Start swarming".
- A small link labeled "About" in the bottom right corner.