

CS342 - Section 4

- Homework 3
 - Due 9/20, can still turn it in before 9/23!
- Homework 4
 - Posted!
 - Classification on SuperTux data with ConvNets (and get higher accuracy)
- Homework 2
 - Grades posted

Agenda

- Recap of ConvNets from lecture
- (Cool) Applications that used ConvNets
- PyTorch code walk through

Recap of ConvNets


- Components
 - Convolutional layers
 - Pooling layers
 - Striding/Padding
- Refer the doc
 - <http://www.philkr.net/cs342/doc/conv2d/>
 - <http://www.philkr.net/cs342/doc/maxpool/>
 - <http://www.philkr.net/cs342/doc/avgpool/>
 - http://www.philkr.net/cs342/doc/stride_padding/

“Convolutional-like” operations

- Moves a sliding window over input tensors and apply the same operations
 - Conv2D: linear transformation over input regions
 - Max Pooling: channel-wise maximum
 - Average Pooling: channel-wise average
- Striding
 - How many pixels do we skip when applying the sliding windows
- Padding
 - How many pixels do we add to the input tensors


Conv2D

- Example of Conv2D with kernel size 3 (no padding)
 - Credit: Michael Plotke




Max Pooling

- Example of Max Pooling with kernel size 2 and stride 2 (no padding)
 - Credit: cs231n.github.io




(Cool) Applications that use ConvNets

- Image Classification
 - HW3!



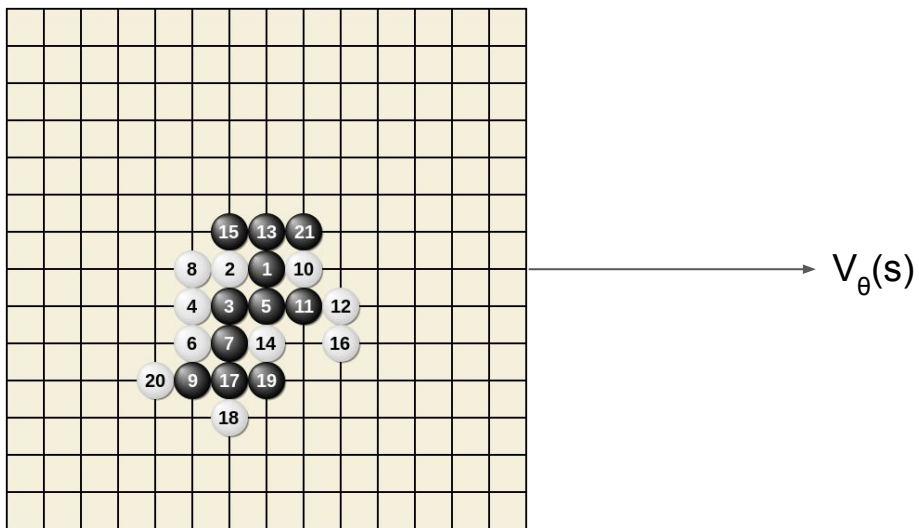
(Cool) Applications that use ConvNets

- Image Segmentation




(Cool) Applications that use ConvNets

- Alpha Go



(Cool) Applications that use ConvNets

- ConvNets has a great inductive bias for visual data
- Filters of different layers can learn different representation
 - Credit: Jason Yosinsky



(Cool) Applications that use ConvNets

- Computer Vision
 - Image Classification
 - Image Segmentation
 - Object Detection
 - Image Generation
 - ...
- Reinforcement Learning
 - Playing Atari Games
 - Alpha Go
- Will be introduced in lecture!

PyTorch Code Walk Through