CS342 - Section 8

- Homework 7
  - Due yesterday, can still turn it in until Sunday
  - What is your IoU?
- Homework 8
  - Posted on webpage
  - On image super resolution
- Homework 6
  - Grades posted

Questions?

Agenda

- Upconvolution vs Linear upsample
- Image super-resolution: L1/L2 loss vs perceptual loss
- PyTorch code walkthrough

Upconv vs Linear Upsample

- Recap
  - Linear upsample
    - Pro: Accurate, no artifacts
    - Con: Slow
  - Upconv
    - Faster, learnable upsampling
    - Con: can have severe checkerboard artifacts

Checkerboard artifacts

- When kernel size can be divided by striding

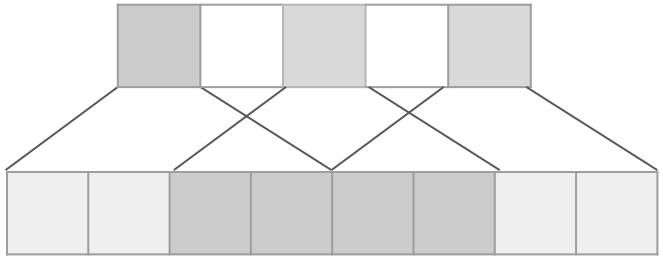

Kernel: 4

Stride: 2

Checkerboard artifacts

- When kernel size cannot be divided by striding

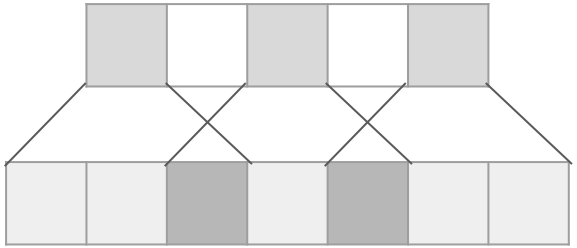

Kernel: 3

Stride: 2

- Severe artifacts on output image

Credit: https://distill.pub/2016/deconv-checkerboard/

Checkerboard artifacts

- If we use bilinear upsampling + convolution w/o stride to upsample

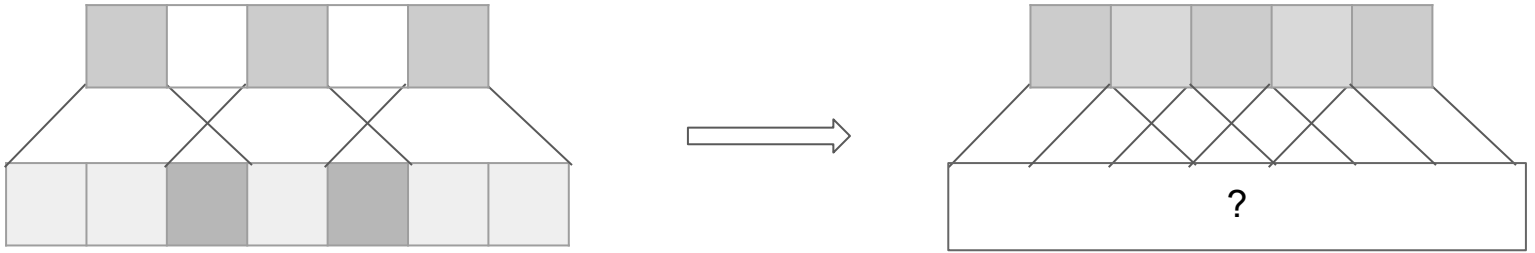

Checkerboard artifacts

- If we use bilinear upsampling + convolution w/o stride to upsample

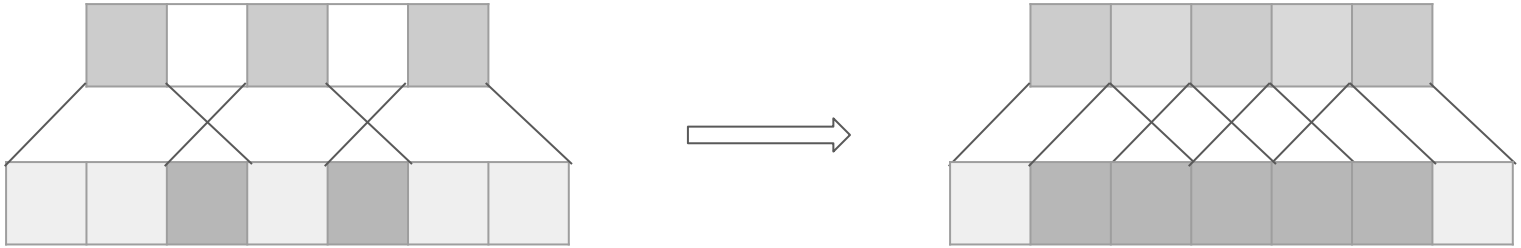

Checkerboard artifacts

- If we use bilinear upsampling + convolution w/o stride to upsample

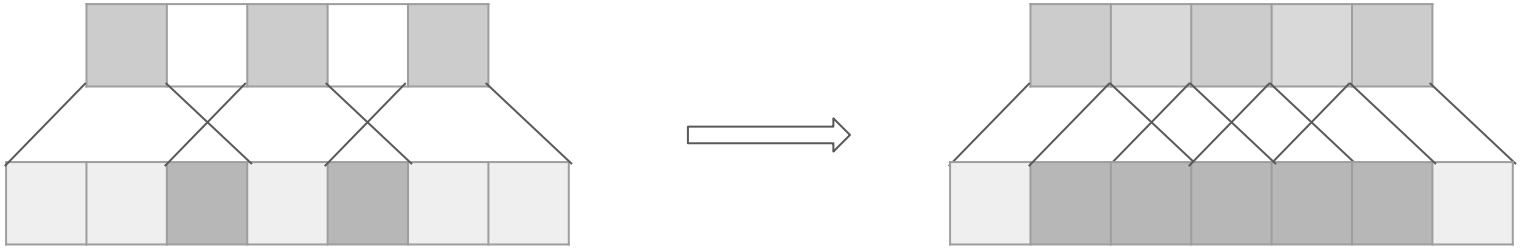

- Checkerboard effects gone

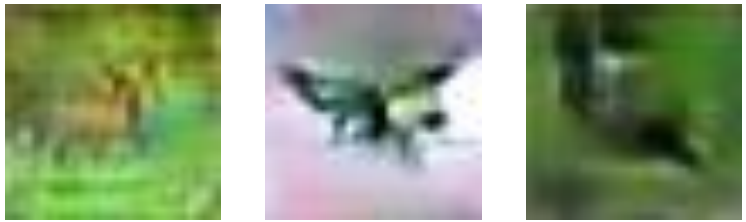

Credit: https://distill.pub/2016/deconv-checkerboard/

Checkerboard artifacts

- Online interactive demo
  - https://distill.pub/2016/deconv-checkerboard/
- Note
  - No need to worry about checkerboard effect in homework 8

Image super-resolution: L1/L2 vs Perceptual loss

- Supervision/Loss
  - L1/L2 per-pixel loss: directly optimize for output pixel
  - Problem: L1/L2 encourage blurry images

Image super-resolution: L1/L2 vs Perceptual loss

- If network is able to generate sharp images, but off by 1 pixel

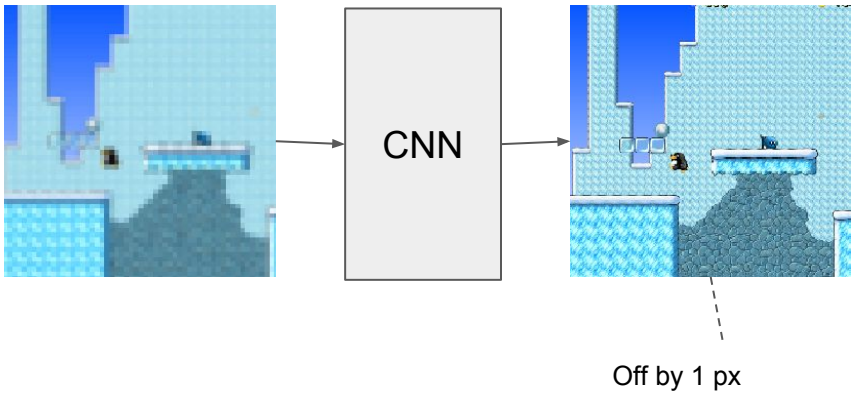

Image super-resolution: L1/L2 vs Perceptual loss

- Then L1/L2 loss will be high, because per-pixel error is large due to the offset

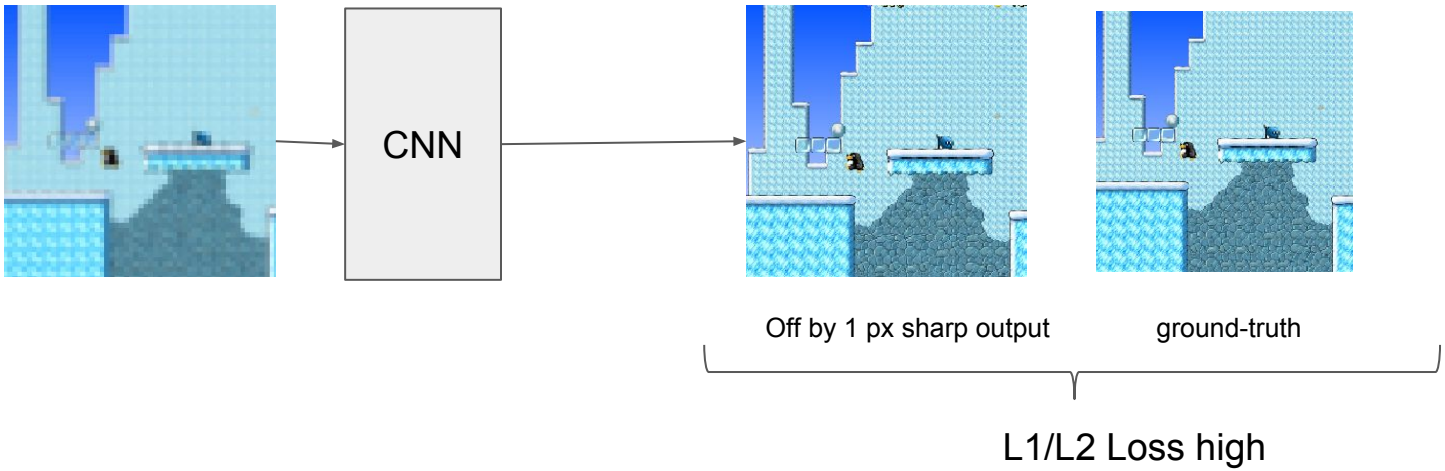

Image super-resolution: L1/L2 vs Perceptual loss

- In such scenario, it is ‘safer’ for network to output blurry image to avoid high L1/L2 penalty rather than outputting sharp images with offsets

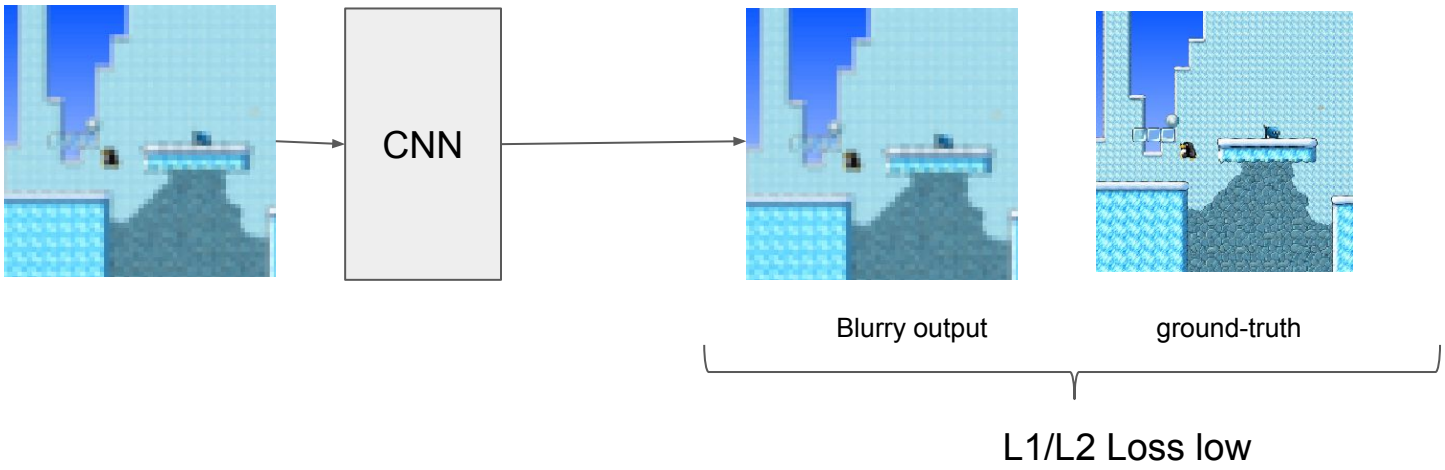

Image super-resolution: L1/L2 vs Perceptual loss

- Therefore, we need a loss that does not penalize sharp images even if there are some spatial shift/offset

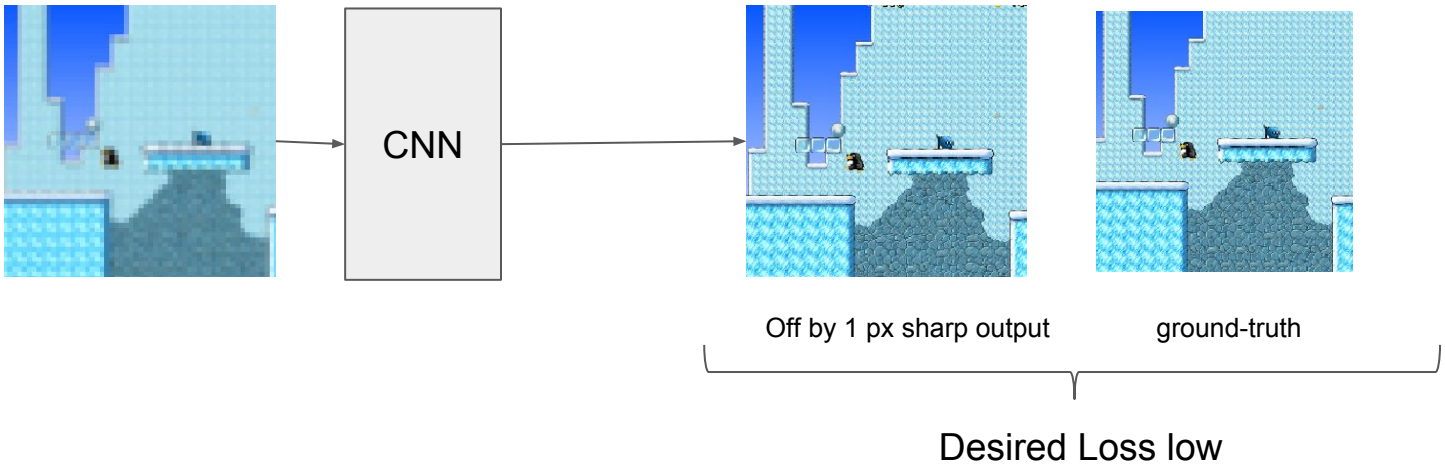

Image super-resolution: L1/L2 vs Perceptual loss

- Such loss could be implemented as differences between visual features

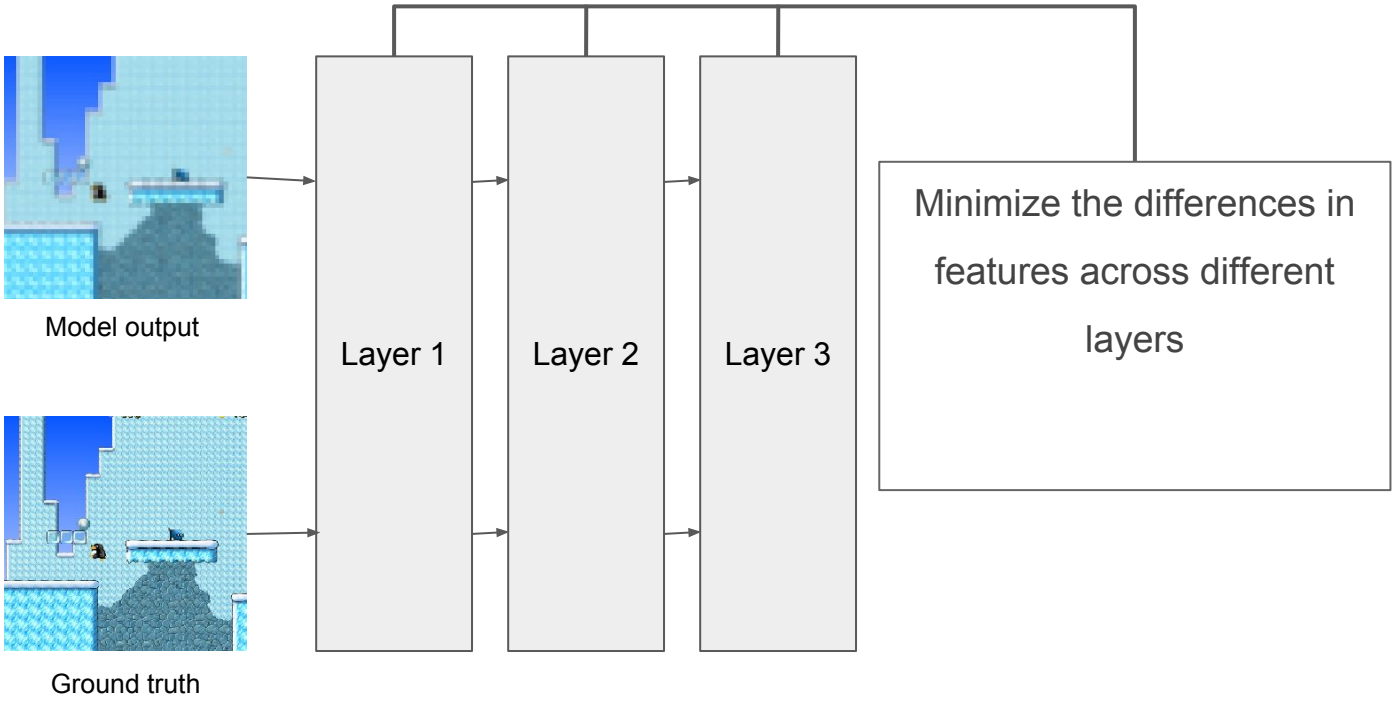

Image super-resolution: L1/L2 vs Perceptual loss

- The loss network below are pre-trained and produce visual features

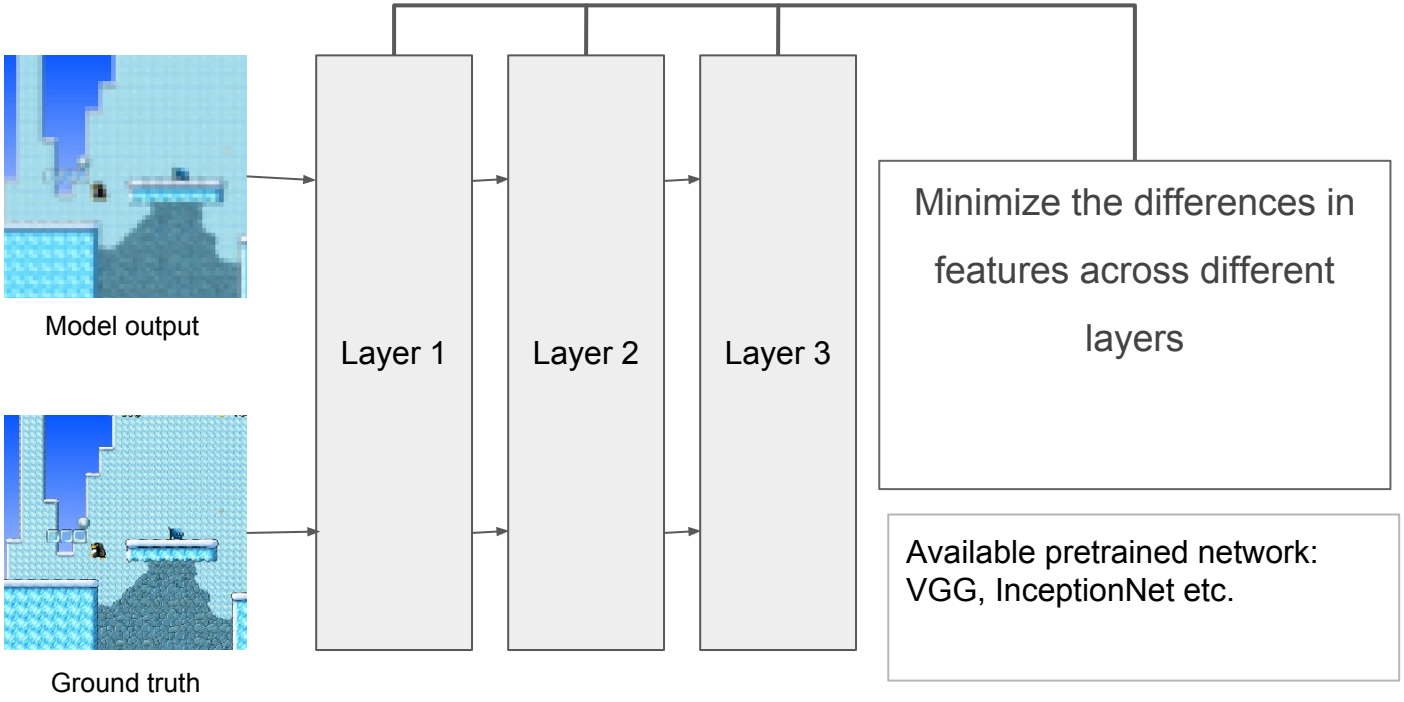

Image super-resolution: L1/L2 vs Perceptual loss

- Note
  - Such loss is called “Perceptual loss”
  - This is not required for Homework 8

Image super-resolution: L1/L2 vs Perceptual loss

- Supervision/Loss
  - Solution/Alternative to L1/L2 loss
    - Idea: Blurry images are not real. Force the images to be real!
    - GAN/adversarial loss
      - Supervision comes from the discriminator which predicts whether images are real or fake/generated. Network try to fool the discriminator by outputting real/sharp images, whereas discriminator is learning not to be fooled.
    - Perceptual loss
      - Supervision comes from semantic features. Generated images forced to have visual features that are close real images.