

Neural nets unit:

- ▶ General introduction (week-before Wednesday)
- ▶ Trying out neural nets (week-before Friday, in lab)
- ▶ How to train your perceptron (last week Monday)
- ▶ The feed-forward and back-propagation algorithms (last week Wednesday)
- ▶ Deep learning: CNNs (last week Friday and **Today**)
- ▶ Deep learning in practice (Wednesday, in lab)

Today:

- ▶ Review of general concepts
- ▶ Tensors and other implementation details
- ▶ Pooling
- ▶ Architecture overview




Figure 14-6. Convolutional layers with multiple feature maps, and images with three color channels

Equation 14-1. Computing the output of a neuron in a convolutional layer

$$z_{i,j,k} = b_k + \sum_{u=0}^{f_h-1} \sum_{v=0}^{f_w-1} \sum_{k'=0}^{f_{n'}-1} x_{i',j',k'} \times w_{u,v,k',k} \quad \text{with} \quad \begin{cases} i' = i \times s_h + u \\ j' = j \times s_w + v \end{cases}$$

In this equation:

- $z_{i,j,k}$ is the output of the neuron located in row i , column j in feature map k of the convolutional layer (layer l).
- As explained earlier, s_h and s_w are the vertical and horizontal strides, f_h and f_w are the height and width of the receptive field, and $f_{n'}$ is the number of feature maps in the previous layer (layer $l-1$).
- $x_{i',j',k'}$ is the output of the neuron located in layer $l-1$, row i' , column j' , feature map k' (or channel k' if the previous layer is the input layer).
- b_k is the bias term for feature map k (in layer l). You can think of it as a knob that tweaks the overall brightness of the feature map k .
- $w_{u,v,k',k}$ is the connection weight between any neuron in feature map k of the layer l and its input located at row u , column v (relative to the neuron's receptive field), and feature map k' .




Figure 14-8. Max pooling layer (2×2 pooling kernel, stride 2, no padding)

Geron, pg 457




Figure 14-9. Invariance to small translations




Figure 14-11. Typical CNN architecture

Geron, pg 461

Coming up:

Due Fri, Apr 11:

*Read excerpt from Geron introducing convolutional neural nets
(See Canvas)*

Due Wed, Apr 16:

Implement perceptron training, feed-forward, and back-propagation

Sometime between Mar 31 and Apr 17:

*Make an office-hours appointment for project check-in
(Originally the deadline was Apr 11)
(So far I have met with only two groups)
(There will be one more ethics reading)*

Apr 28 and 30: Final versions of the project

(That's all.)