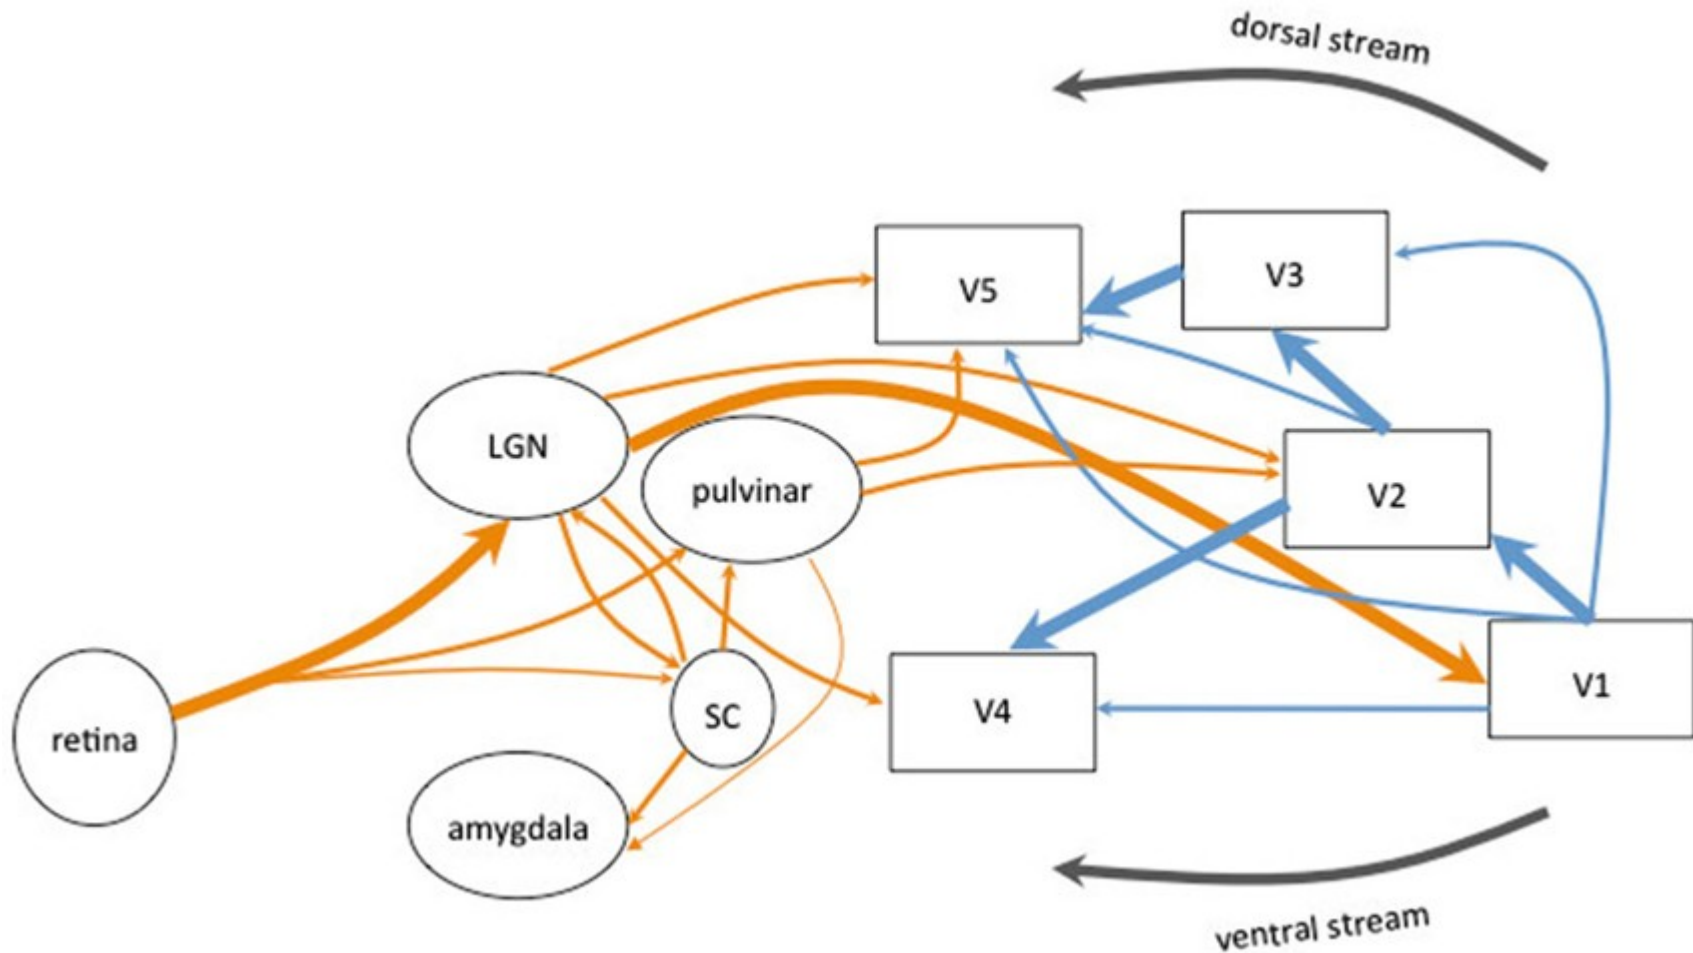


Convolutional Neural Networks

Nathan Sprague

Human Visual System




Urbanski, Marika, Olivier A. Coubar, and Clémence Bourlon. "Visualizing the blind brain: brain imaging of visual field defects from early recovery to rehabilitation techniques." *Neurovision: Neural bases of binocular vision and coordination and their implications in visual training programs* (2014).

Convolutional Neural Networks

- Convolutional neural networks use the same trick of learning layers of localized features...
- CNN's were actually being used by Yann Lecun at Bell Labs around 1990


Convolutions

Grayscale Image
1 convolutional filter




Convolutions

Grayscale Image
1 convolutional filter




Color Image
5 convolutional filters




Pooling Layers

- Pooling layers down-sample the filter outputs to
 - Reduce dimensionality and computational requirements
 - Increase the spatial extent of subsequent filters



Complete Network


- A “traditional” CNN is composed of convolutional layers, each followed by non-linearities, followed by pooling layers, with a dense (non-convolutional) layer at the end:



Chen, Xianjie, and Alan L. Yuille. "Articulated pose estimation by a graphical model with image dependent pairwise relations." Advances in Neural Information Processing Systems. 2014.

Residual Networks

- How deep can we make these networks? Simply stacking more convolutional layers eventually degrades performance.
- One solution is to introduce “skip connections”:




- “Residual learning”

He, Kaiming, et al. "Deep residual learning for image recognition." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.


> 70,000 citations!

Residual Networks

- ResNet-34:



- Get ResNet-50 by introducing “bottleneck” blocks:



- The 1x1 convolutions can be used to increase or decrease the number of channels