

Neural networks

Computer vision - parameter sharing

COMPUTER VISION


Topics: computer vision

- We can design neural networks that are specifically adapted for such problems
 - ▶ must deal with very high-dimensional inputs
 - 150×150 pixels = 22500 inputs, or 3×22500 if RGB pixels
 - ▶ can exploit the 2D topology of pixels (or 3D for video data)
 - ▶ can build in invariance to certain variations we can expect
 - translations, illumination, etc.
- Convolutional networks leverage these ideas
 - ▶ local connectivity
 - ▶ **parameter sharing**
 - ▶ pooling / subsampling hidden units

COMPUTER VISION

Topics: parameter sharing

- Second idea: share matrix of parameters across certain units
 - units organized into the same “feature map” share parameters
 - hidden units within a feature map cover different positions in the image



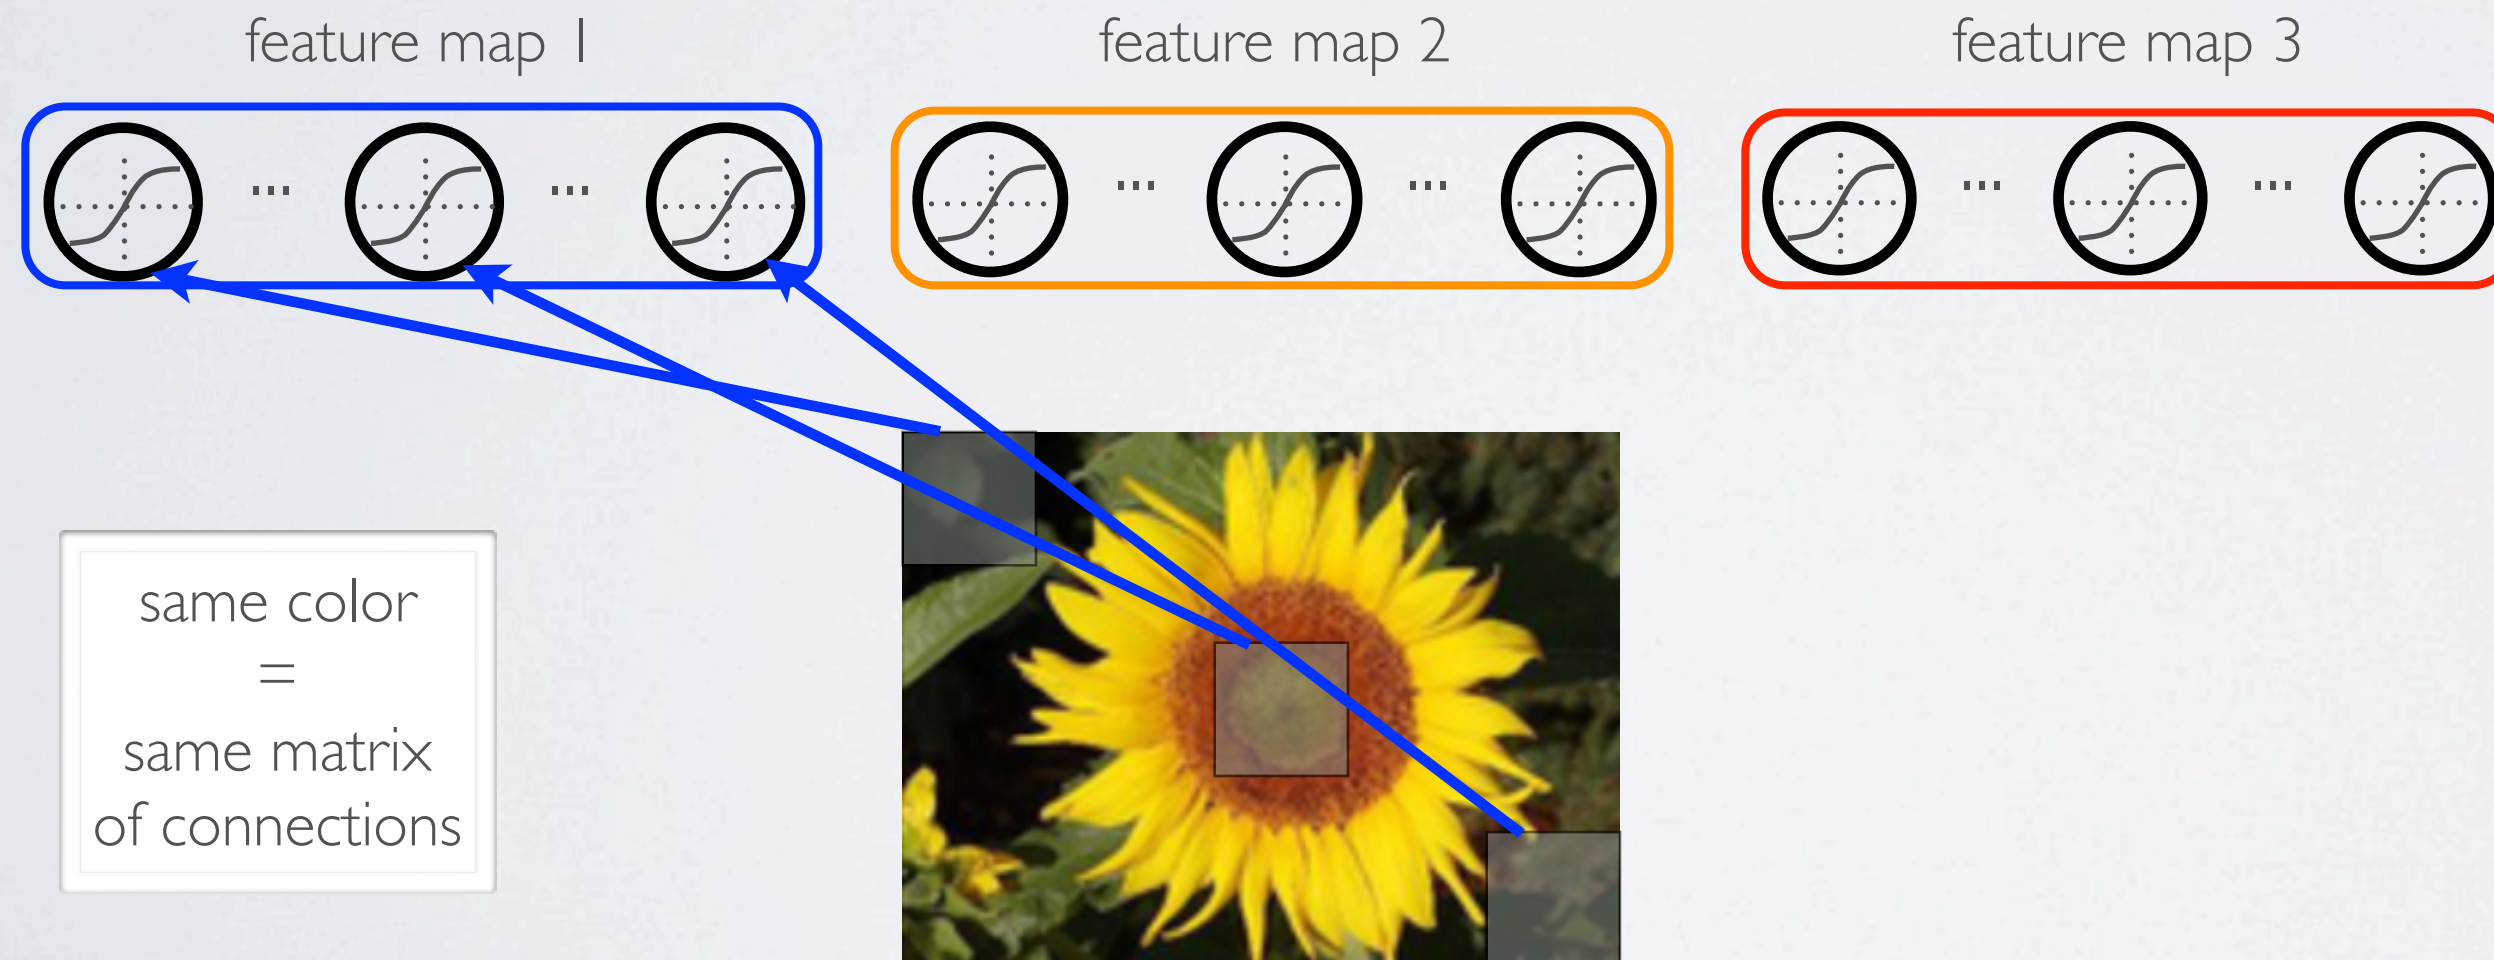
same color
=
same matrix
of connections



COMPUTER VISION

Topics: parameter sharing

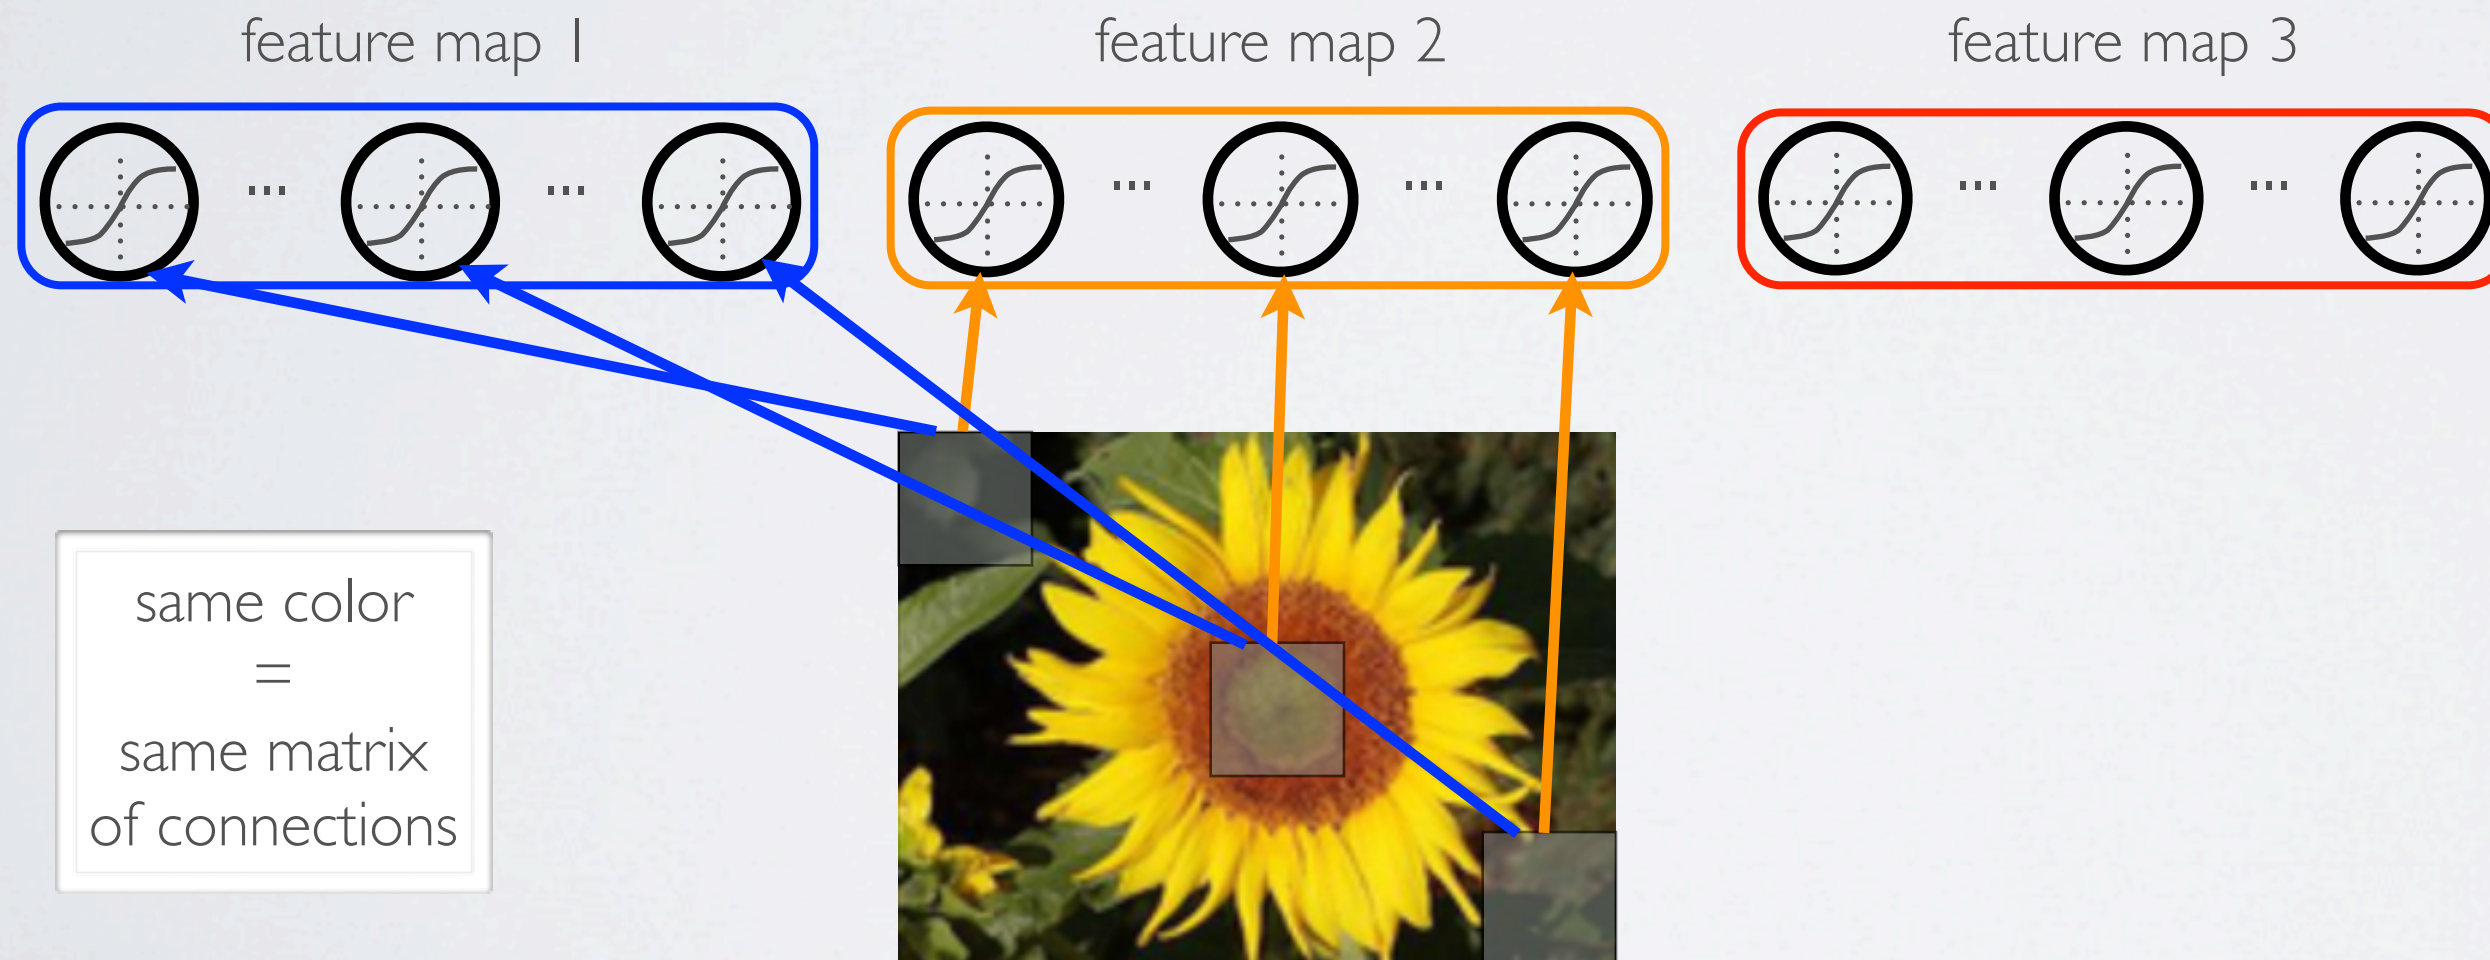
- Second idea: share matrix of parameters across certain units
 - units organized into the same “feature map” share parameters
 - hidden units within a feature map cover different positions in the image



COMPUTER VISION

Topics: parameter sharing

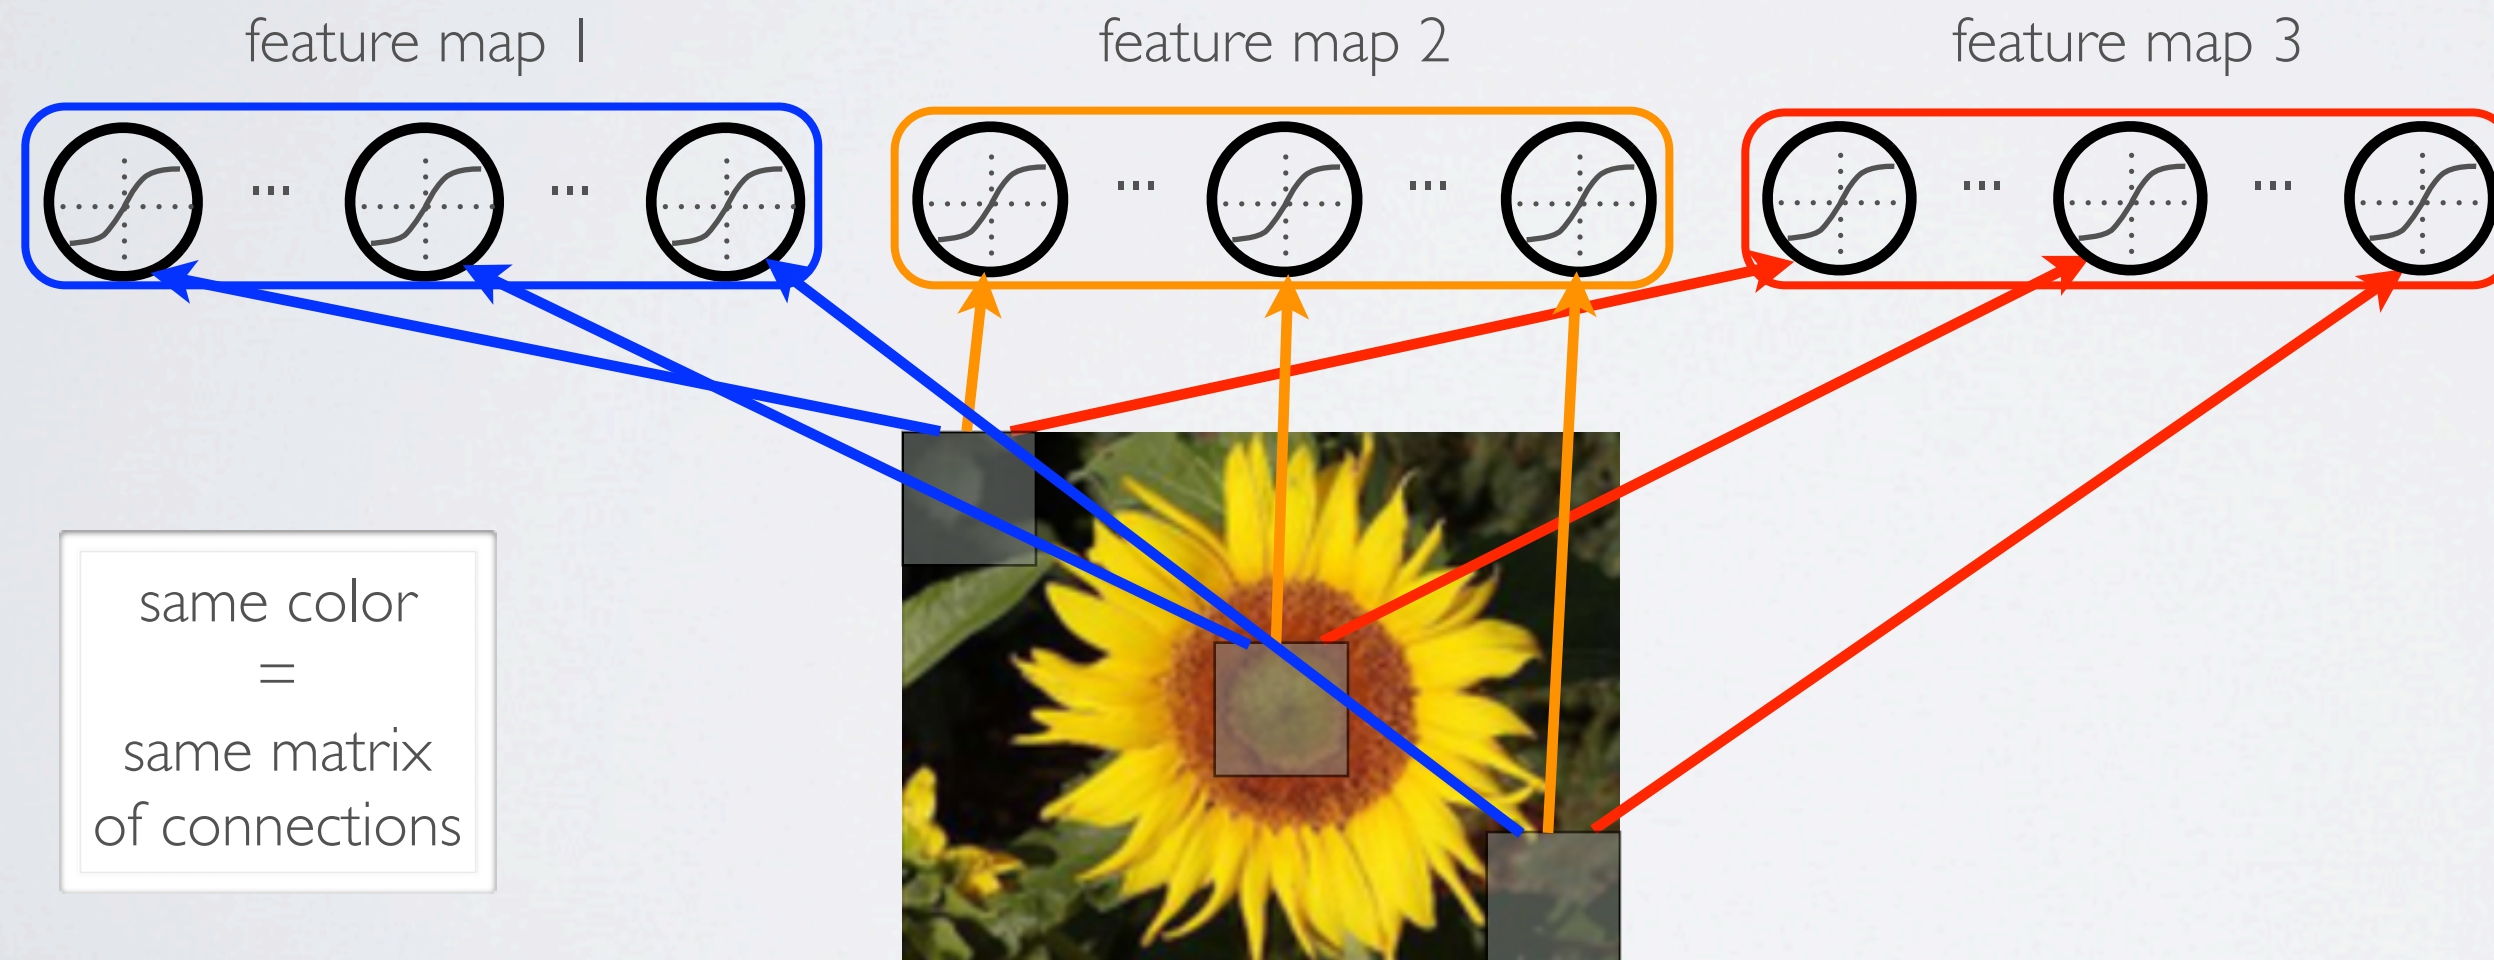
- Second idea: share matrix of parameters across certain units
 - units organized into the same “feature map” share parameters
 - hidden units within a feature map cover different positions in the image



COMPUTER VISION

Topics: parameter sharing

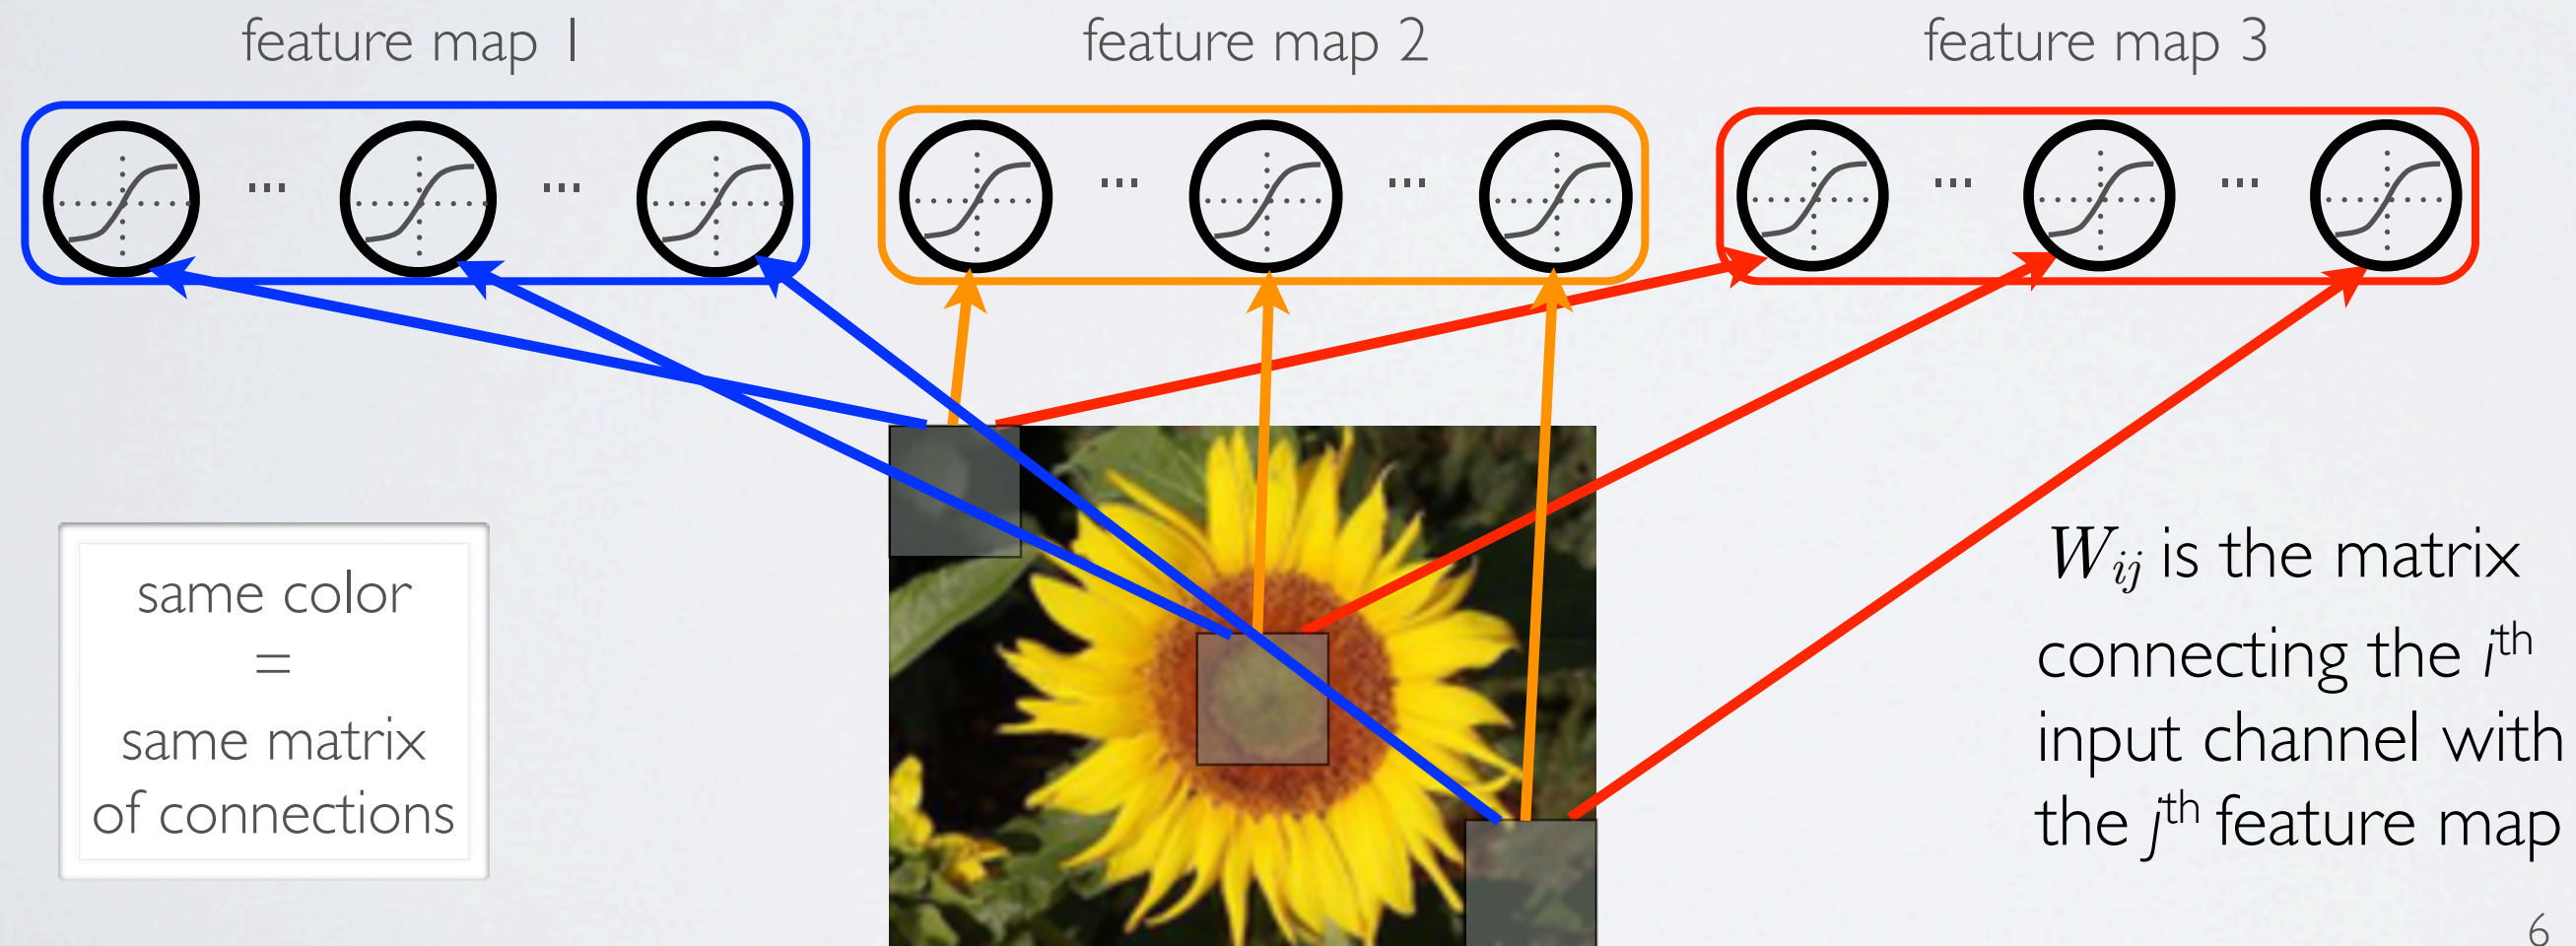
- Second idea: share matrix of parameters across certain units
 - units organized into the same “feature map” share parameters
 - hidden units within a feature map cover different positions in the image



COMPUTER VISION

Topics: parameter sharing

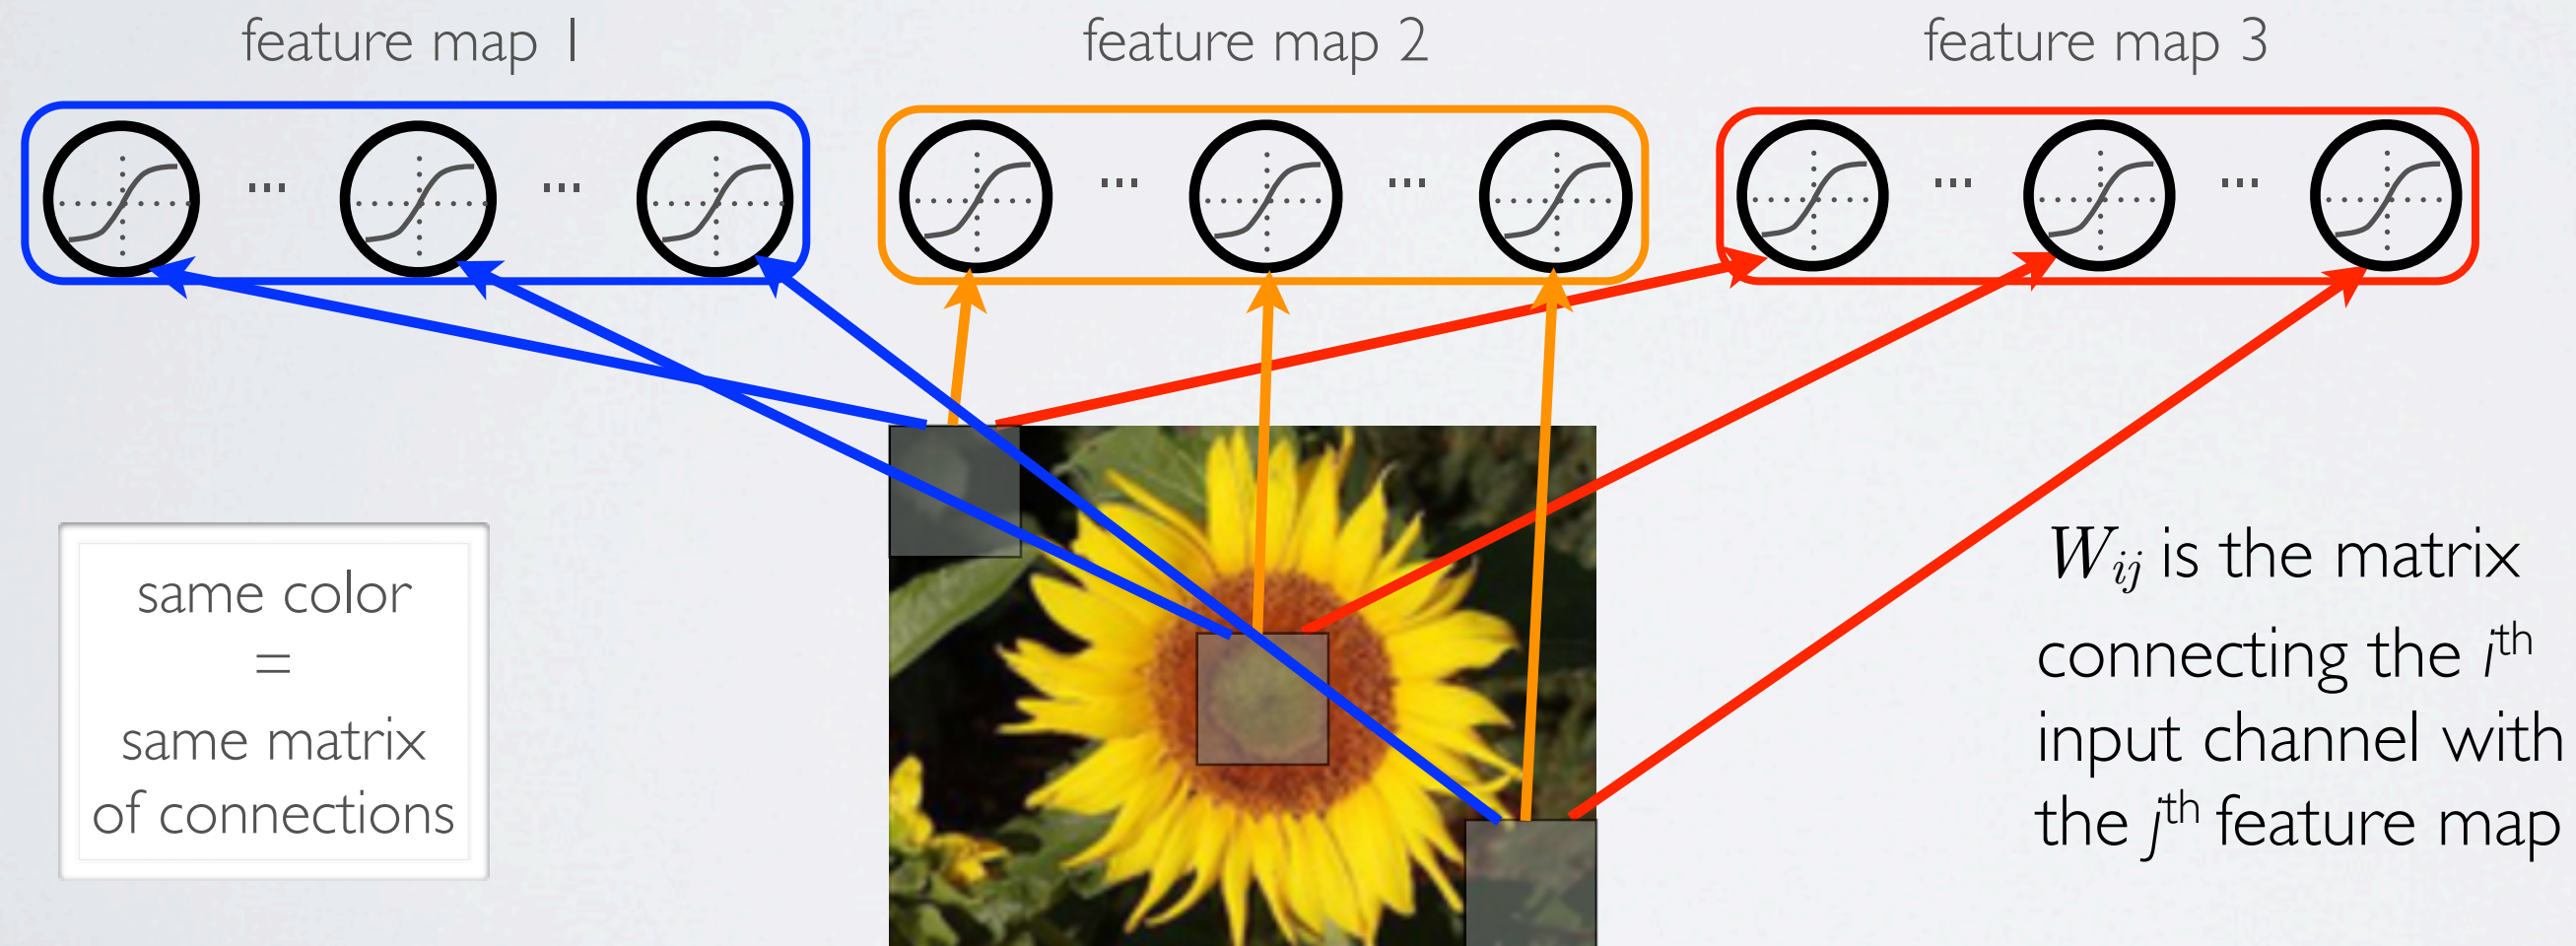
- Second idea: share matrix of parameters across certain units
 - units organized into the same “feature map” share parameters
 - hidden units within a feature map cover different positions in the image



COMPUTER VISION

Topics: parameter sharing

- Solves the following problems:
 - reduces even more the number of parameters
 - will extract the same features at every position (features are “equivariant”)




COMPUTER VISION

Topics: parameter sharing

Jarret et al. 2009

- Each feature map forms a 2D grid of features
 - can be computed with a discrete convolution ($*$) of a kernel matrix k_{ij} which is the hidden weights matrix W_{ij} with its rows and columns flipped



- x_i is the i^{th} channel of input
- k_{ij} is the convolution kernel
- g_j is a learned scaling factor
- y_j is the hidden layer

(could have added a bias)

$$y_j = g_j \tanh\left(\sum_i k_{ij} * x_i\right)$$