Scene Statistics
Part 2

Odelia Schwartz

2016

Summary

- We've considered bottom-up scene statistics, efficient coding, and relation of linear transforms to visual filters

- This class: going beyond learning V1 like linear filters

Beyond linear

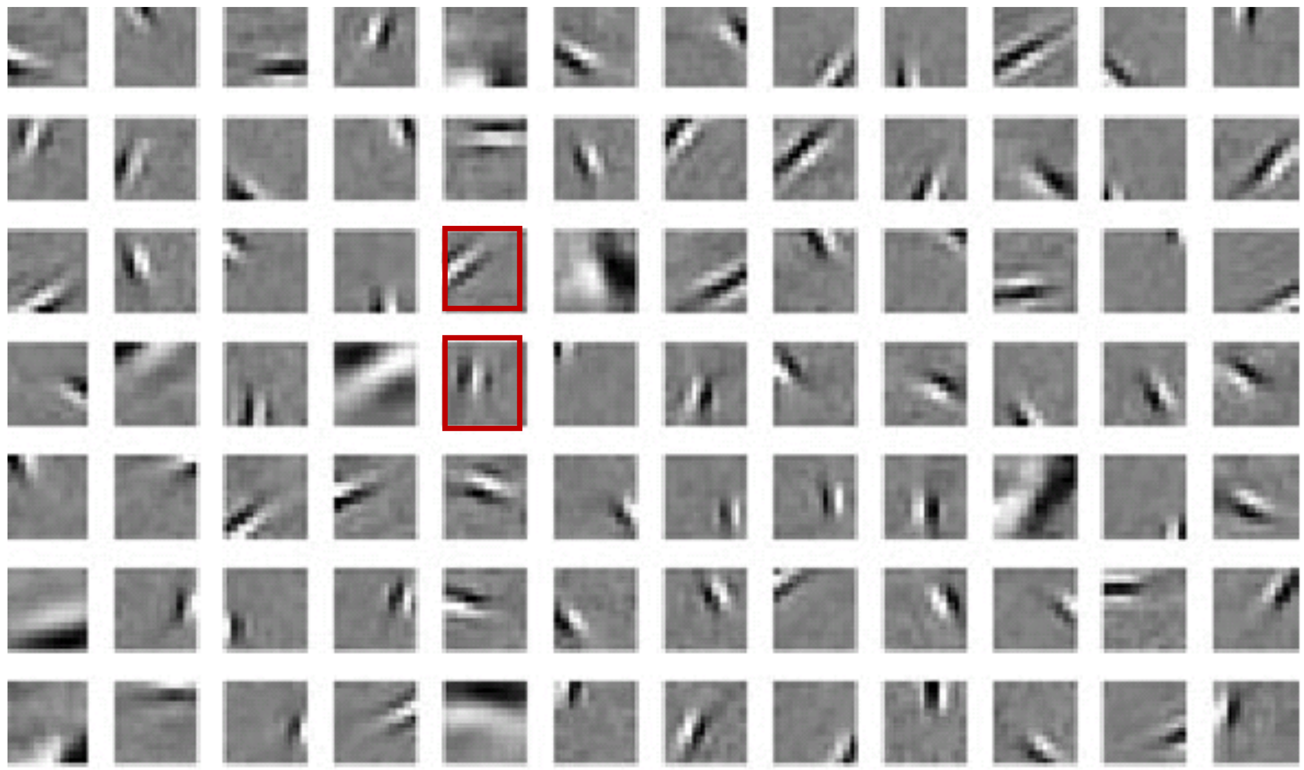

- Filter responses as independent as possible assuming a linear transform
- But are they independent?

Bottom-up Joint Statistics

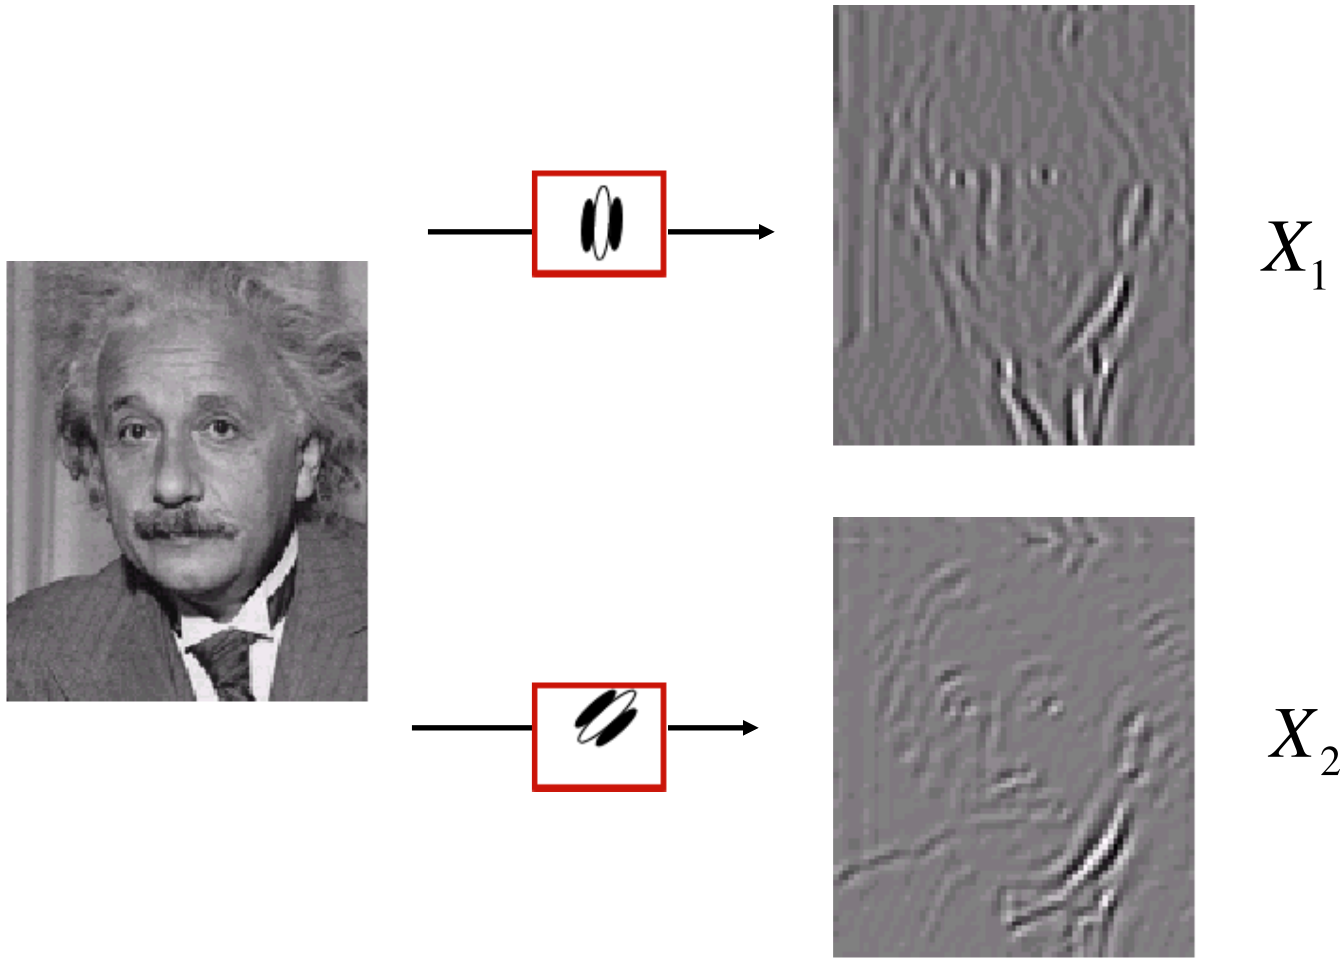

Bottom-up Joint Statistics

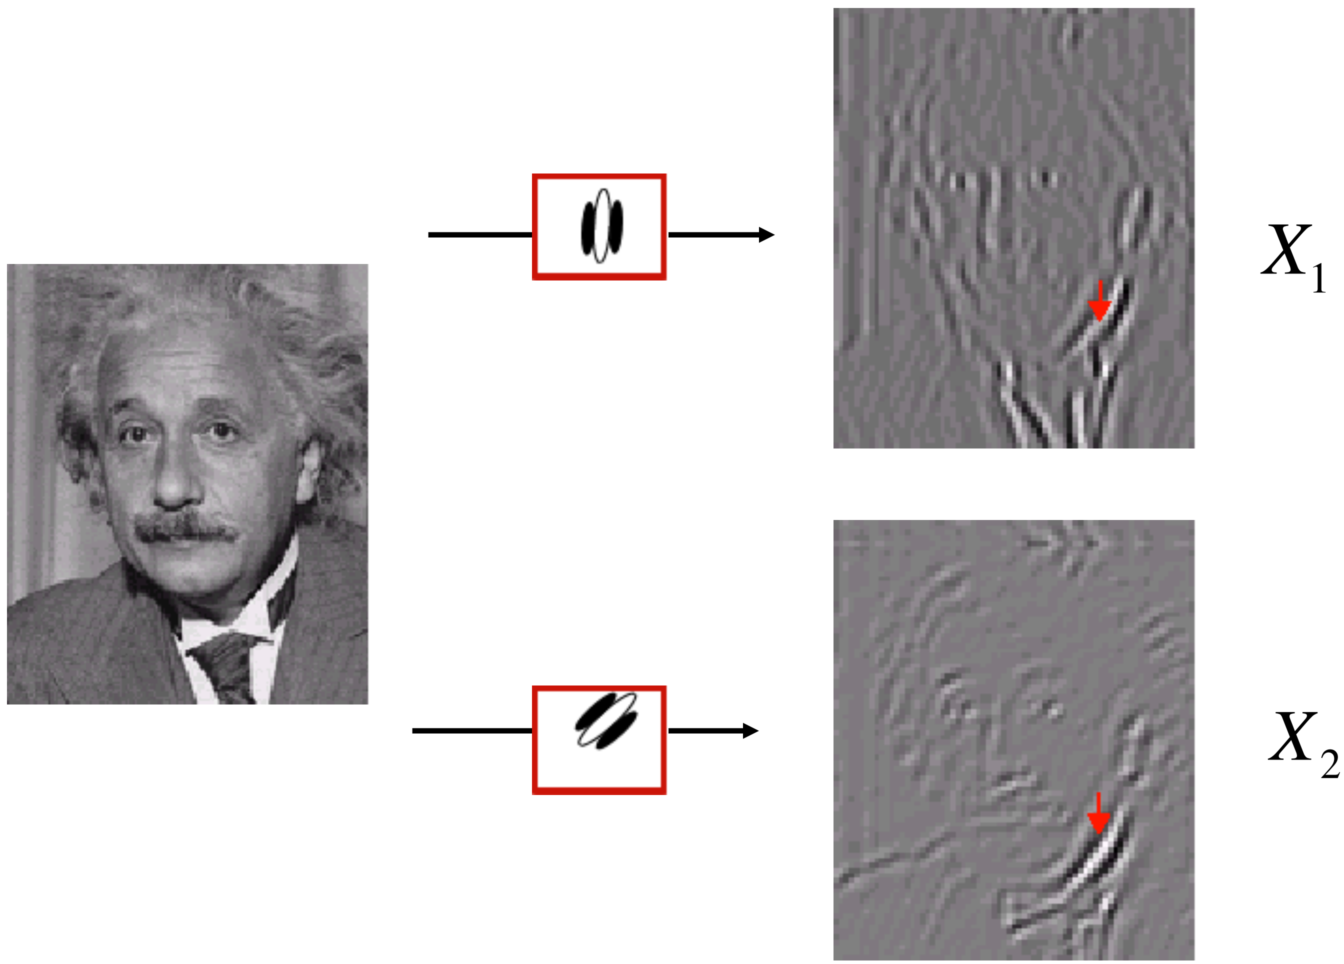

Bottom-up Joint Statistics

$histo(X_1 \mid X_2 \approx 0.1)$

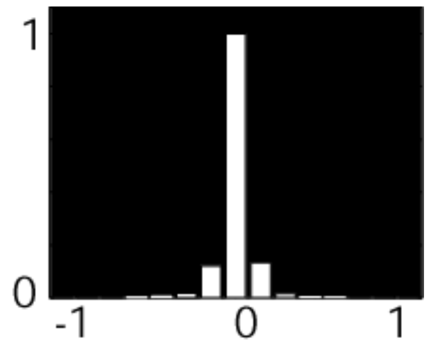

$histo(X_1 \mid X_2 \approx 0.8)$

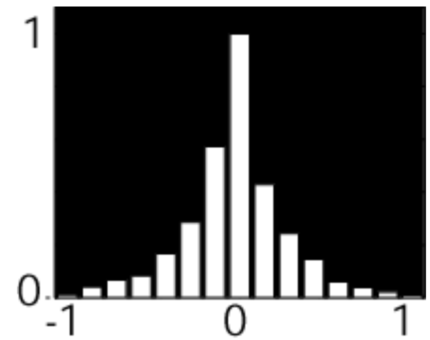

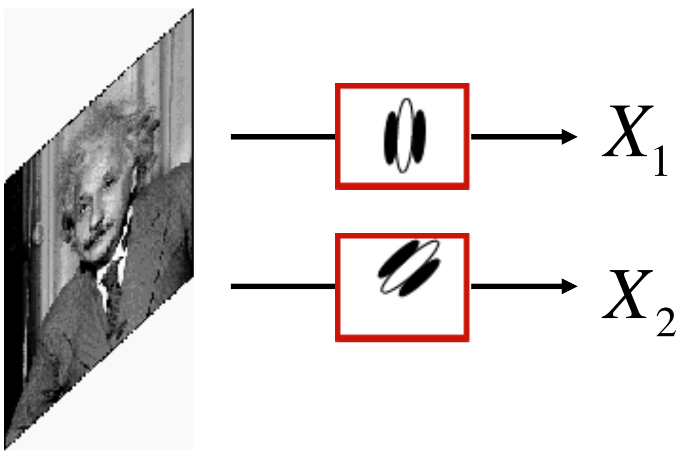

Are X_1 and X_2 statistically independent?

Bottom-up Joint Statistics

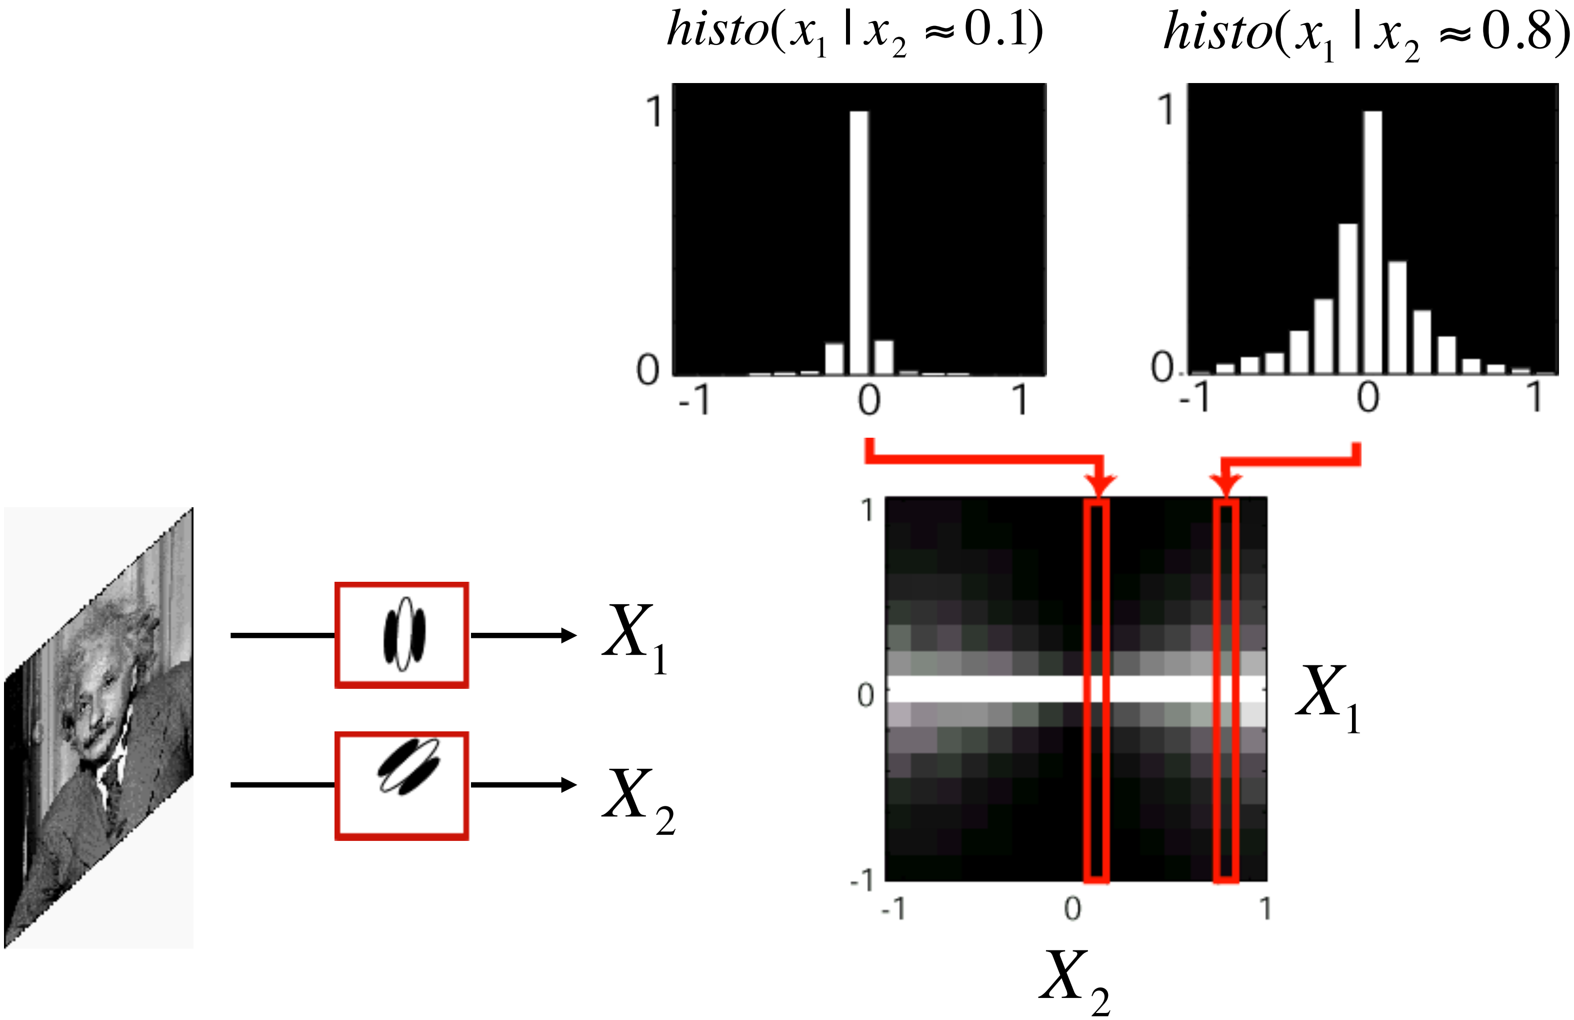

X_1 and X_2 are not statistically independent

Schwartz and Simoncelli, 2001

Bottom-up Joint Statistics

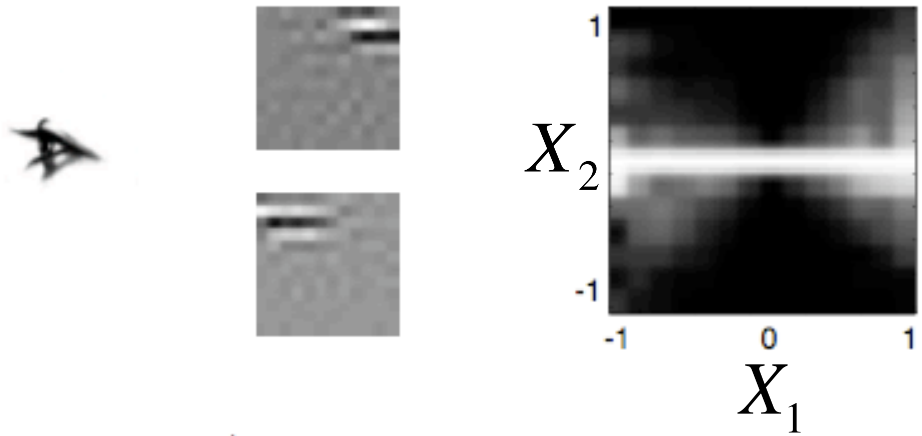

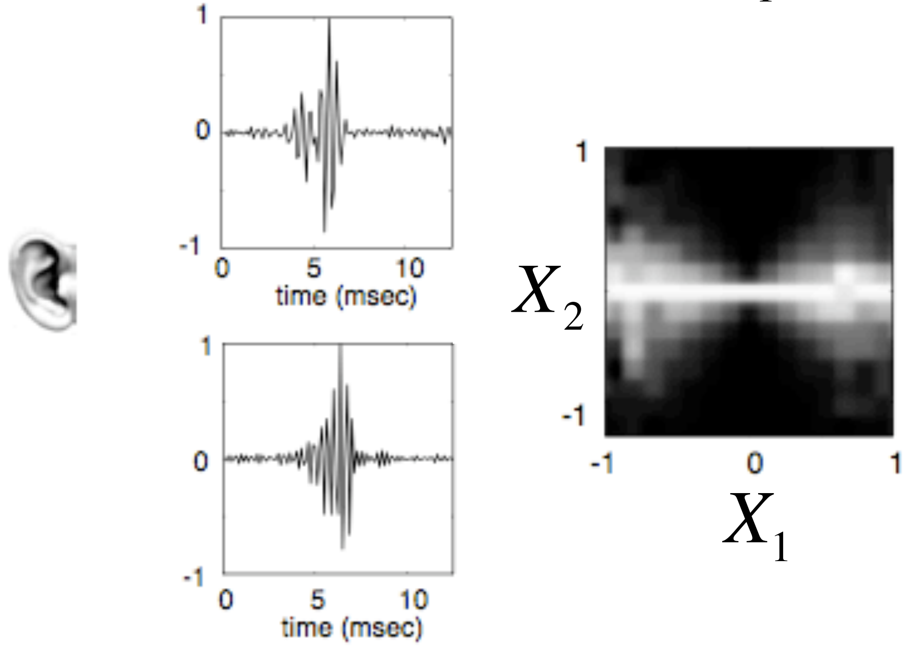

Bottom-up Statistics

Filter pair and different image patches…

$\longrightarrow X_1$

$\longrightarrow X_2$

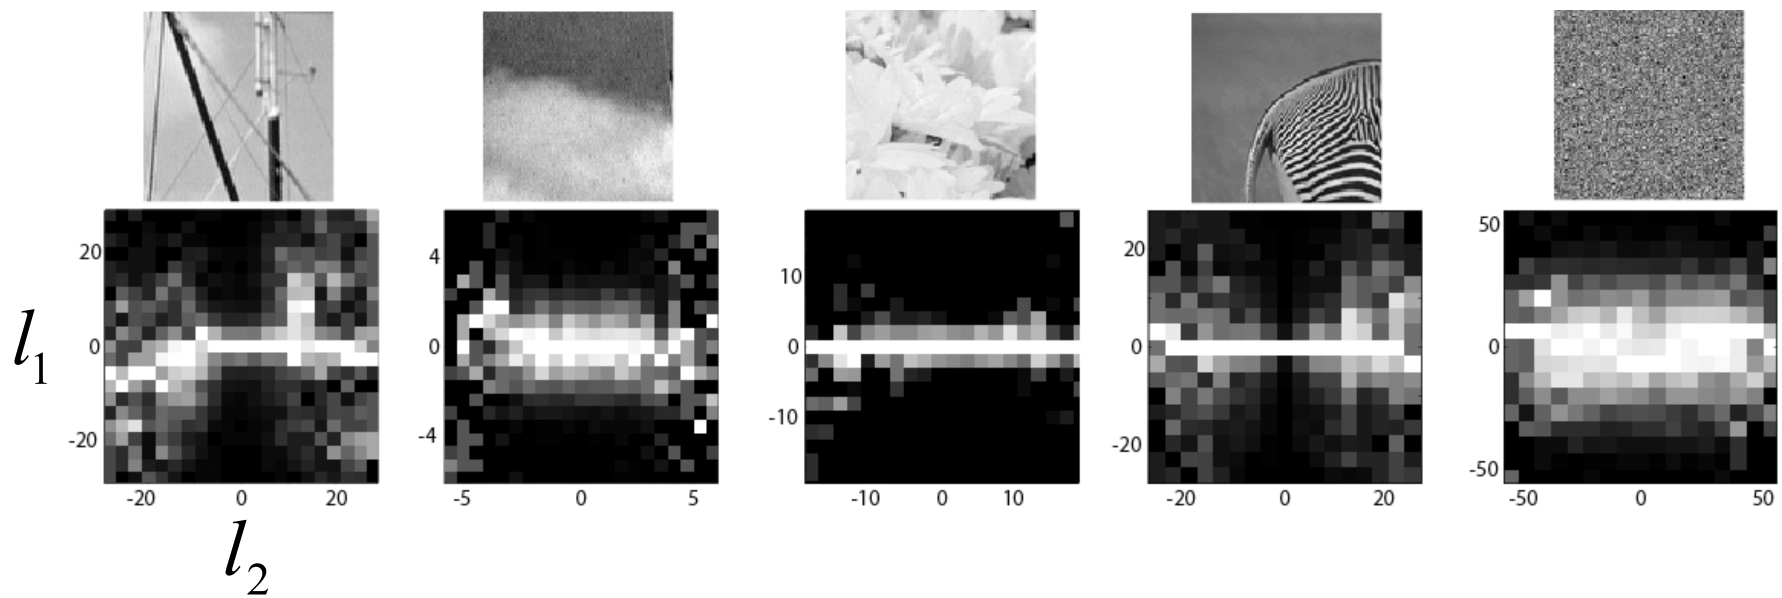

Bottom-up Statistics

Image patch and different filter pairs…

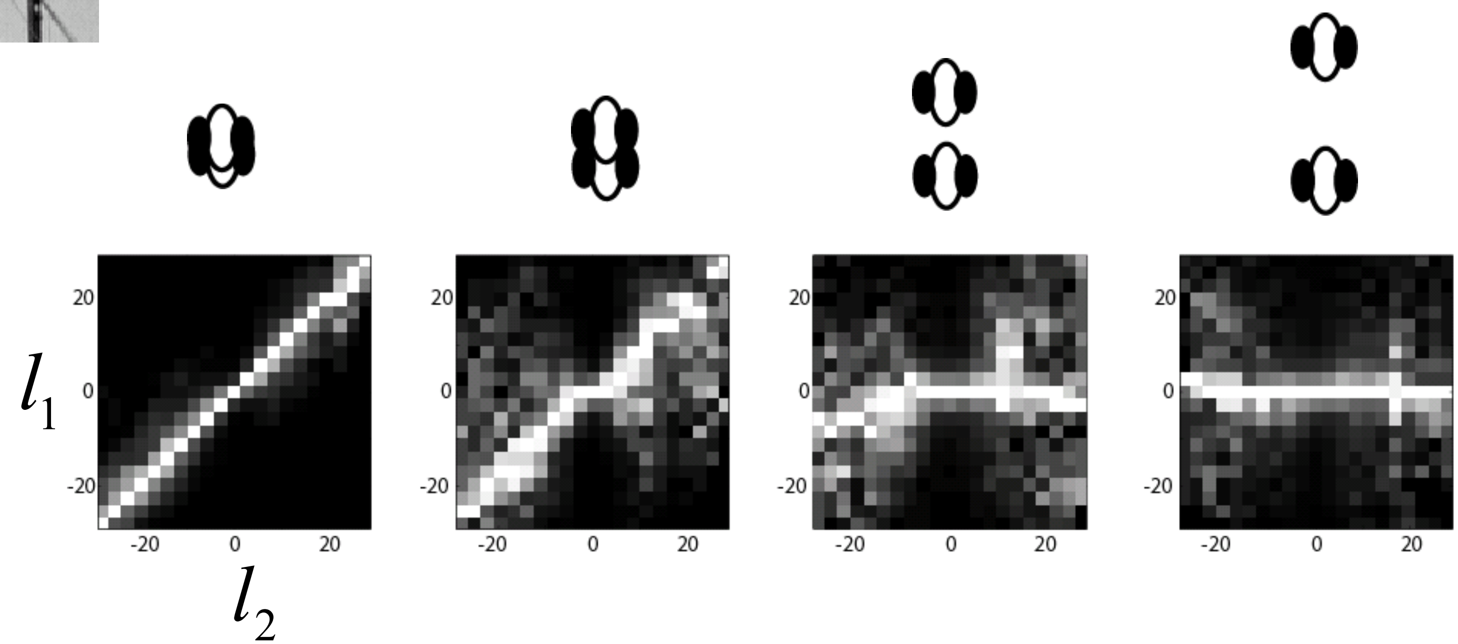

Modeling filter coordination

Modeling filter coordination in images

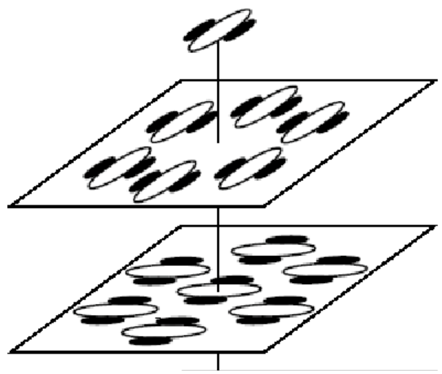

- Learning how more complex representations build up from the structure of dependencies in images

- Reducing dependencies further via nonlinear: divisive normalization

Modeling filter coordination

Modeling filter coordination in images

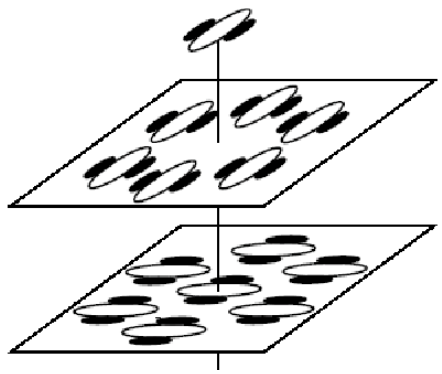

What kind of complex representations?

Modeling filter coordination

Modeling filter coordination in images

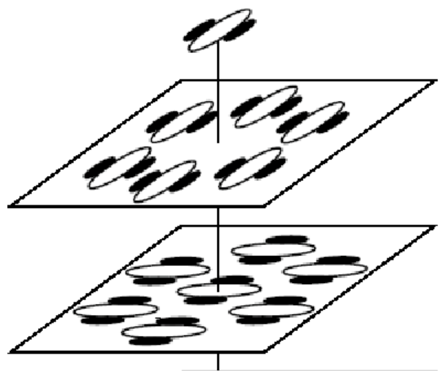

What kind of complex representations?

1. In V1, eg complex cells
2. Higher visual areas

Modeling filter coordination

Modeling filter coordination in images

First what we know; then learning from dependencies in images

More complex representations

In primary visual cortex (capturing an invariance)

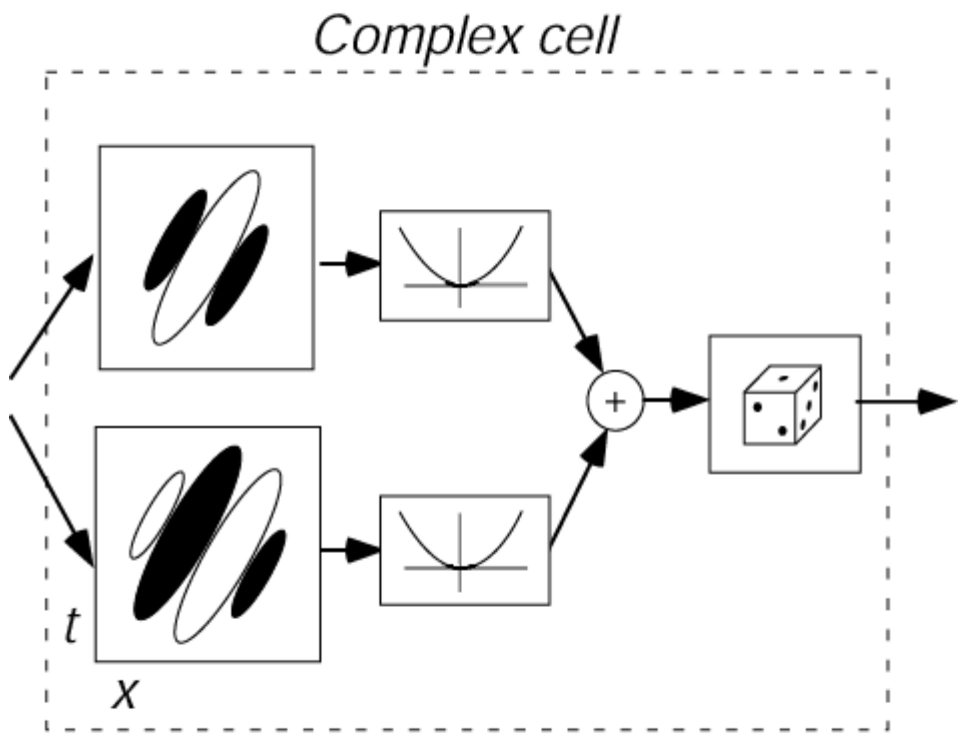

Adelson & Bergen (1985)

Beyond Primary Visual Cortex

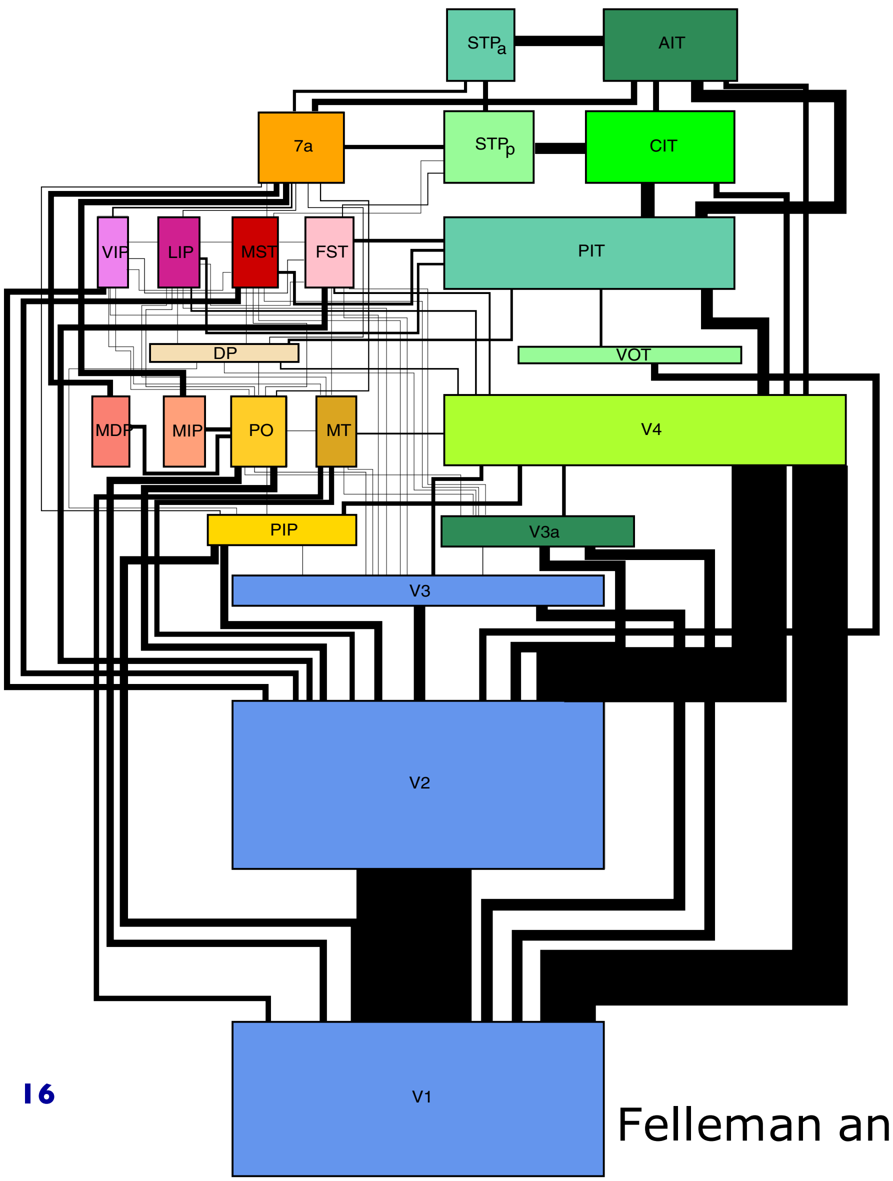

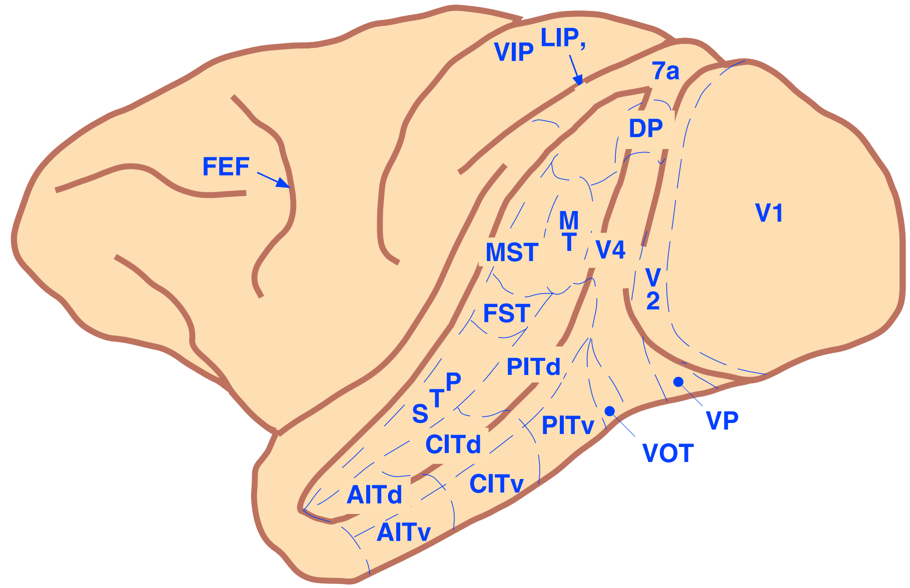

Felleman and Van Essen, 1991

RF size increases at higher levels

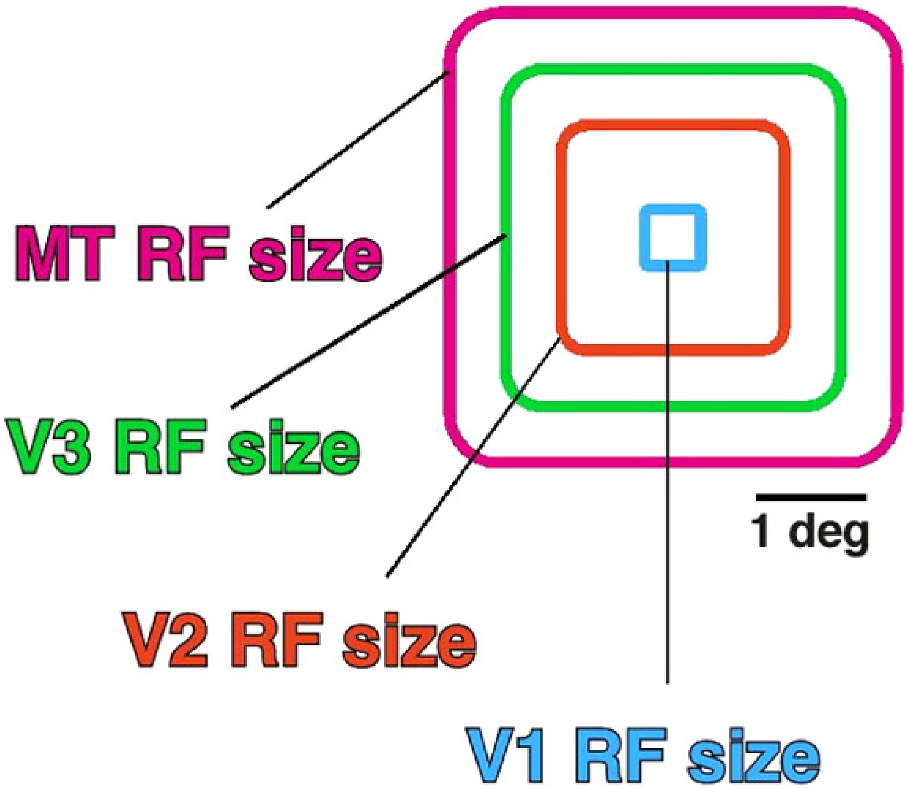

Beyond Primary Visual Cortex

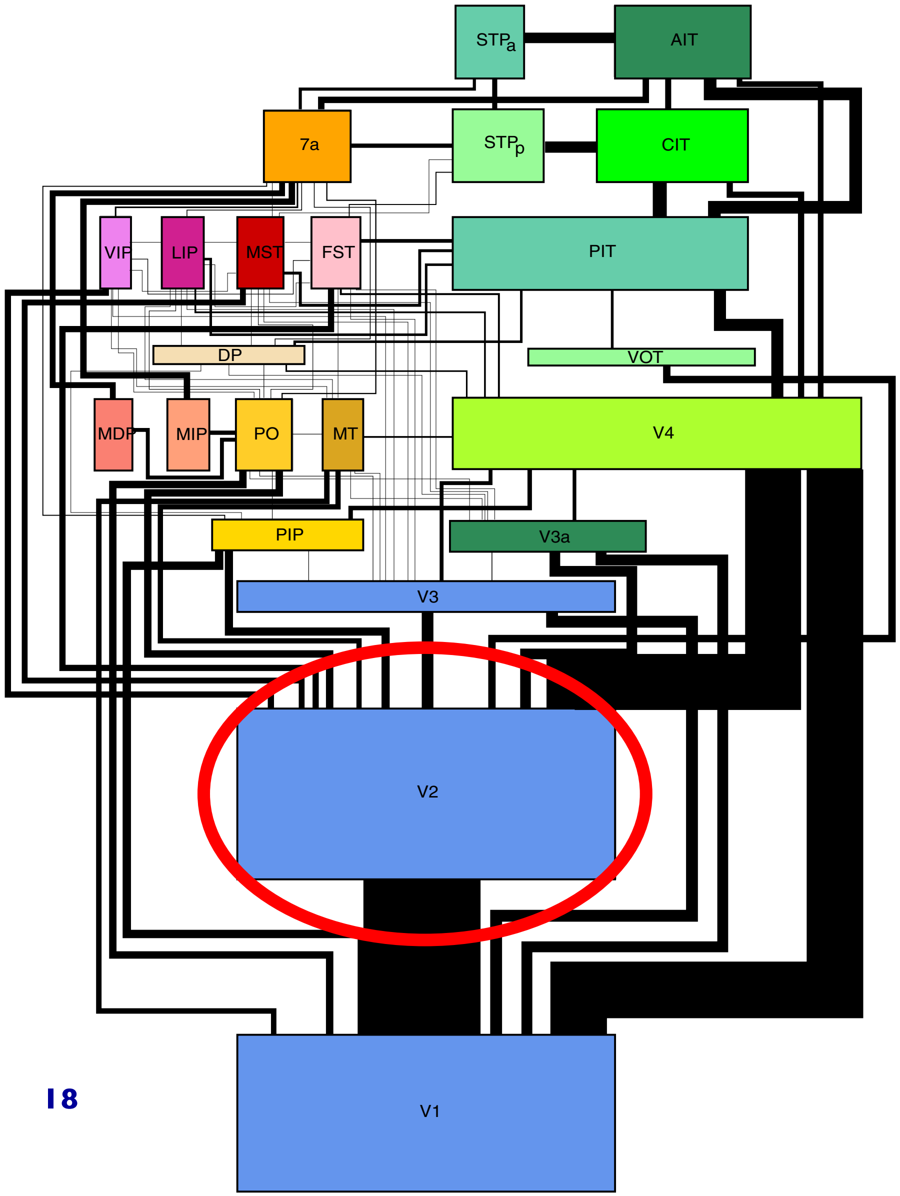

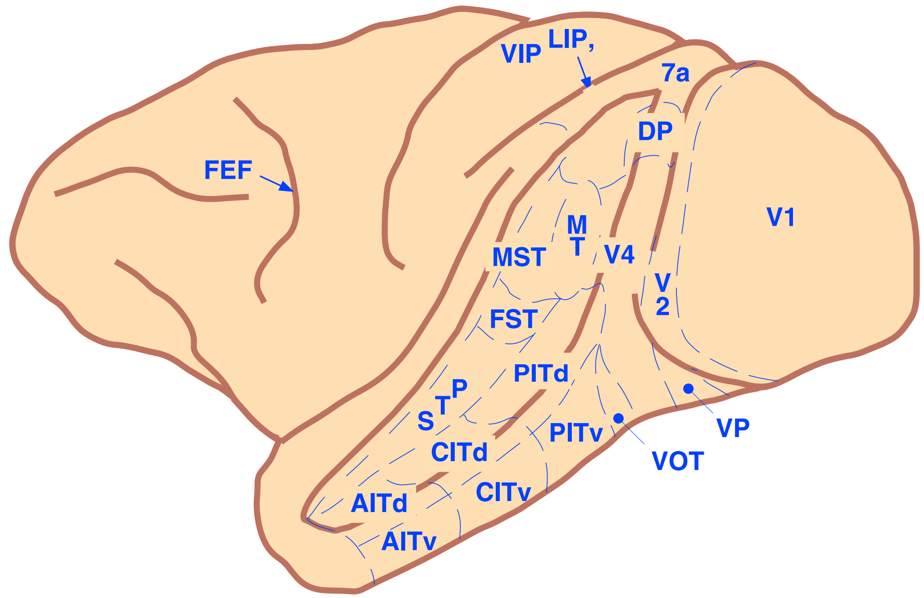

More complex representations

Example of V2 neurophysiology

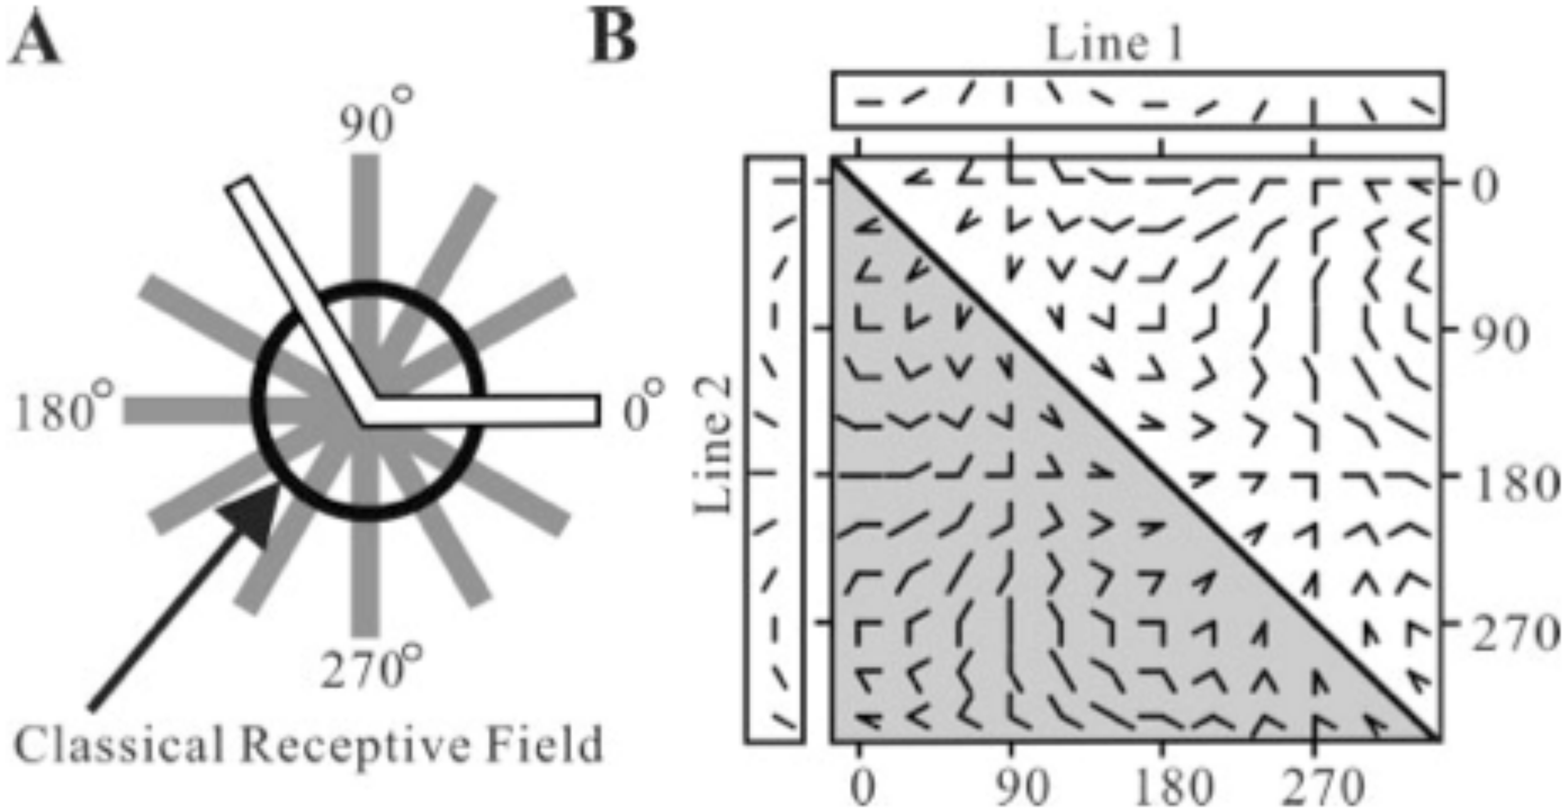

Ito and Komatsu, 2005

More complex representations

Example of V2 neurophysiology

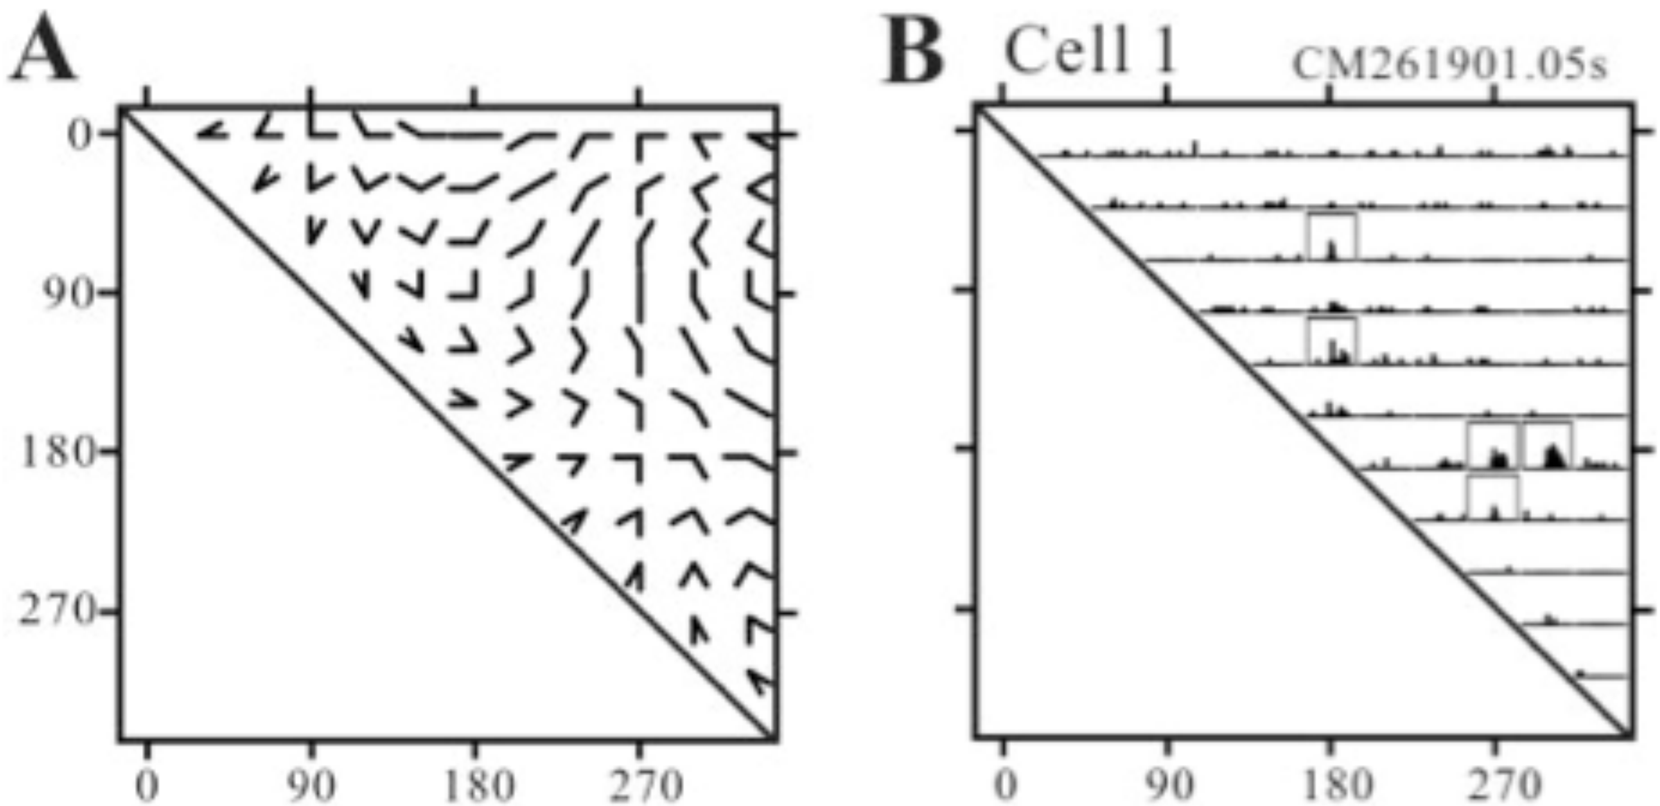

Ito and Komatsu, 2005

Beyond Primary Visual Cortex

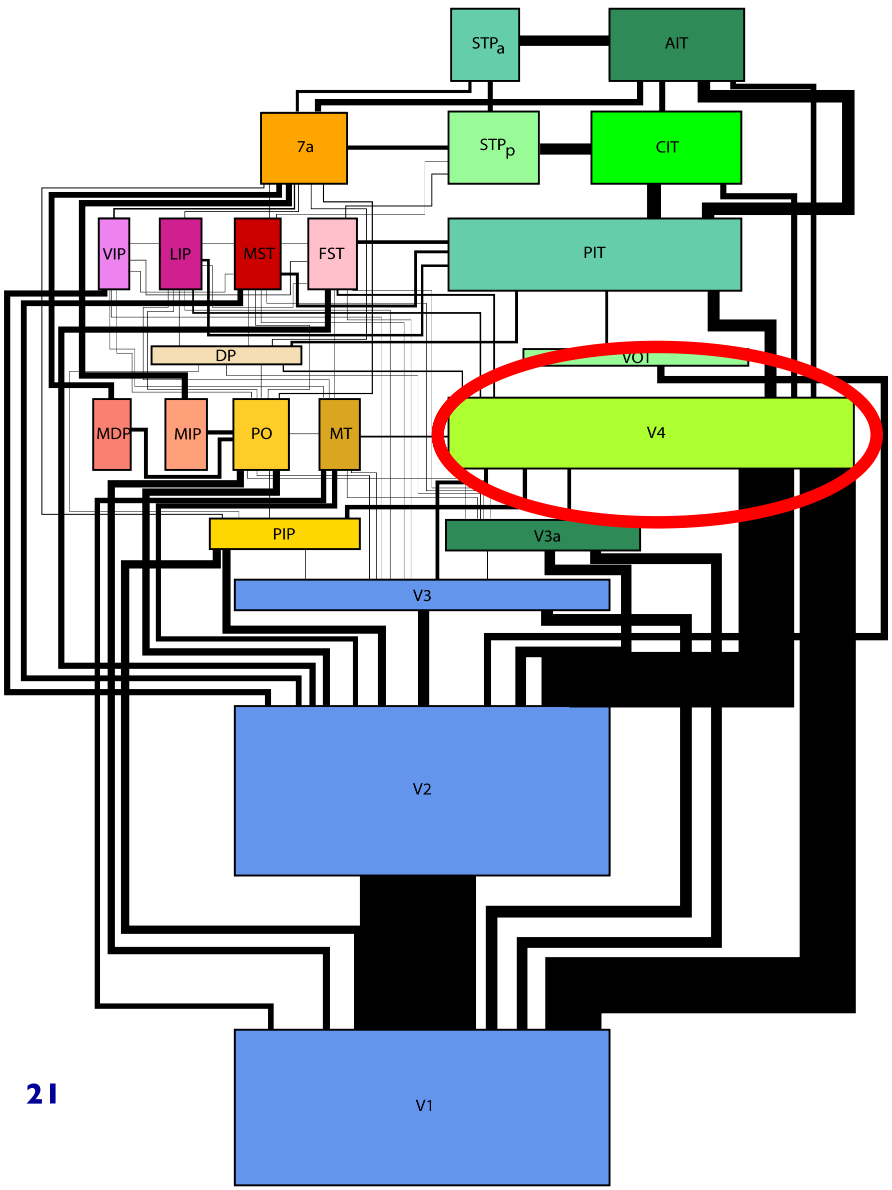

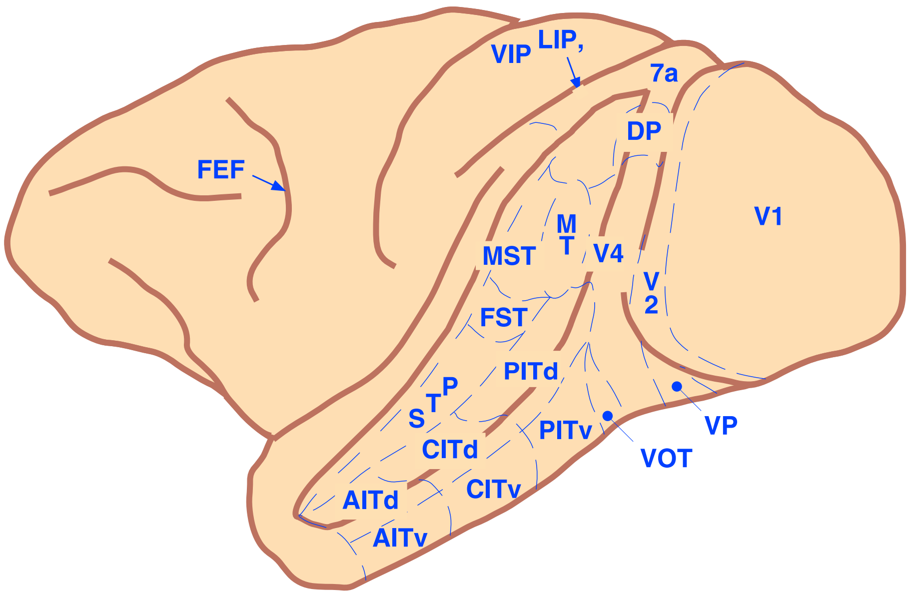

From Adam Kohn

More complex representations

Example of V4 neurophysiology

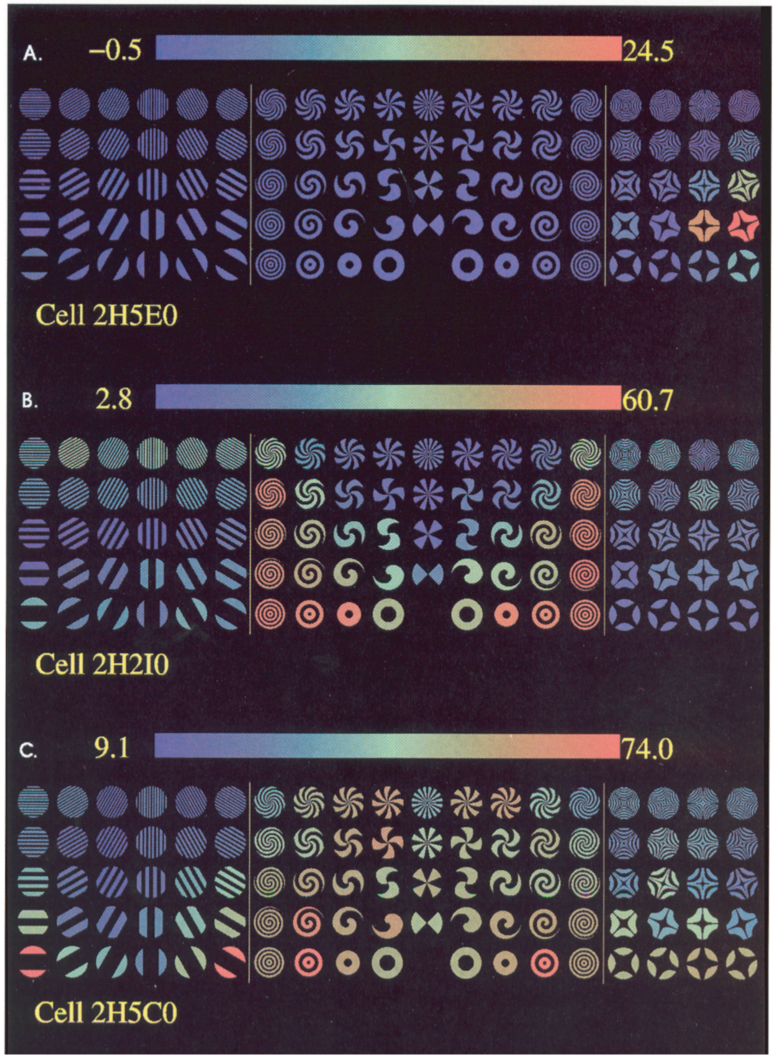

Beyond Primary Visual Cortex

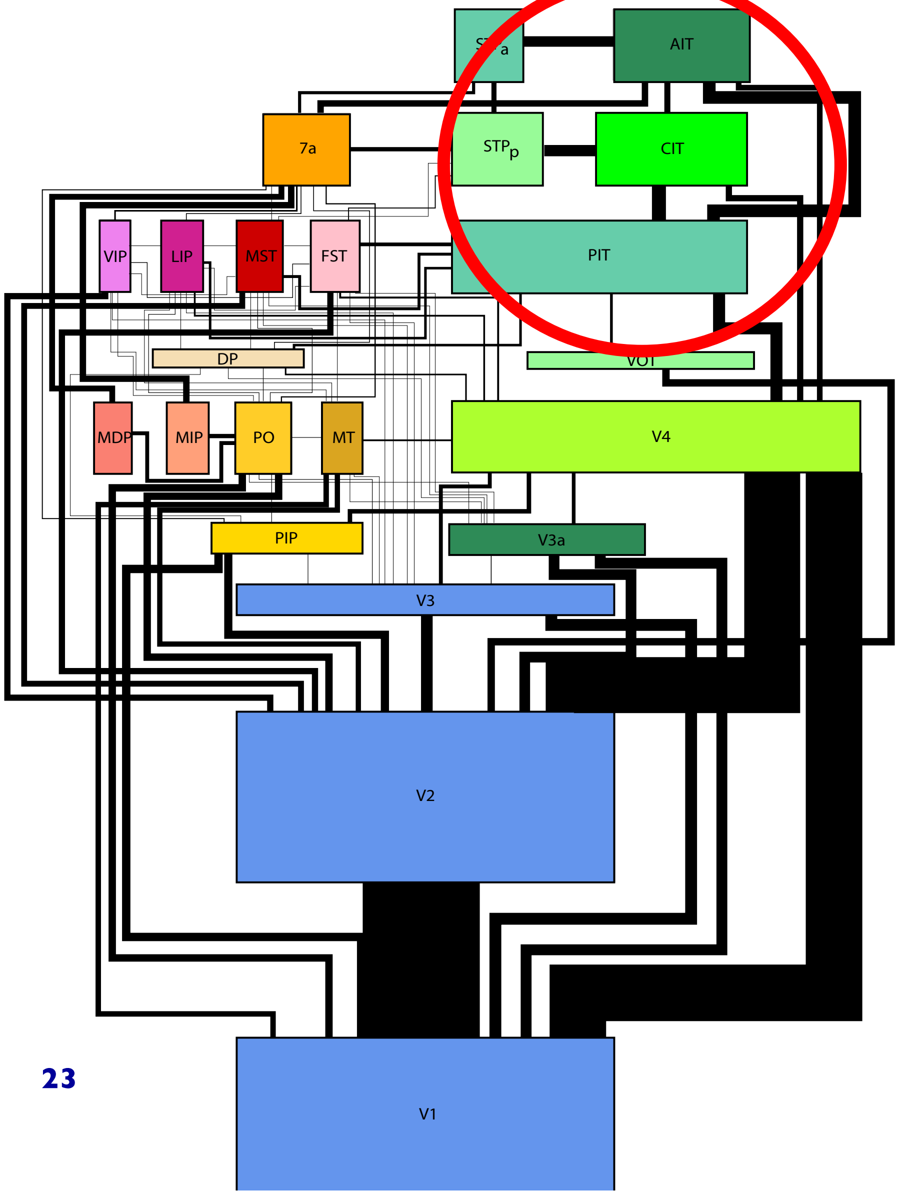

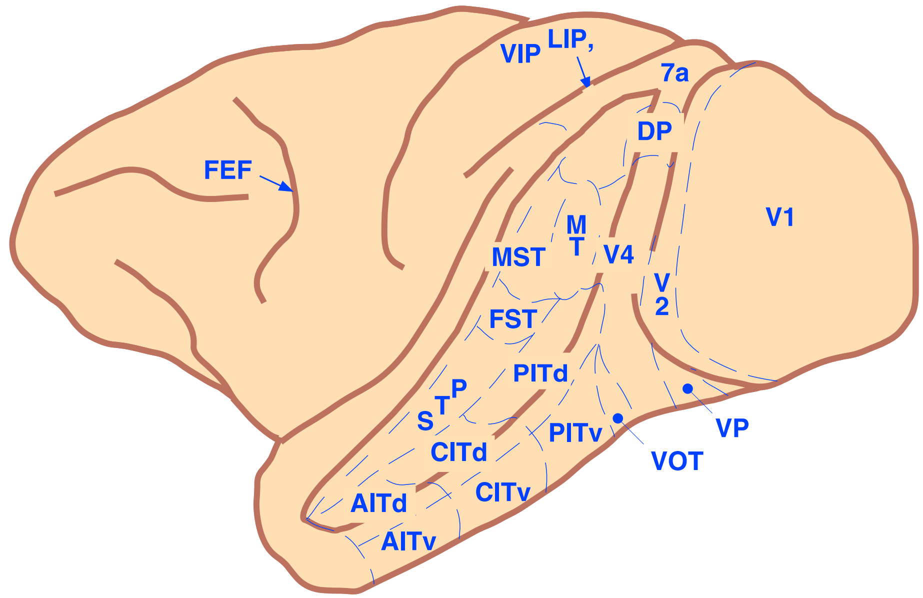

From Adam Kohn

More complex representations

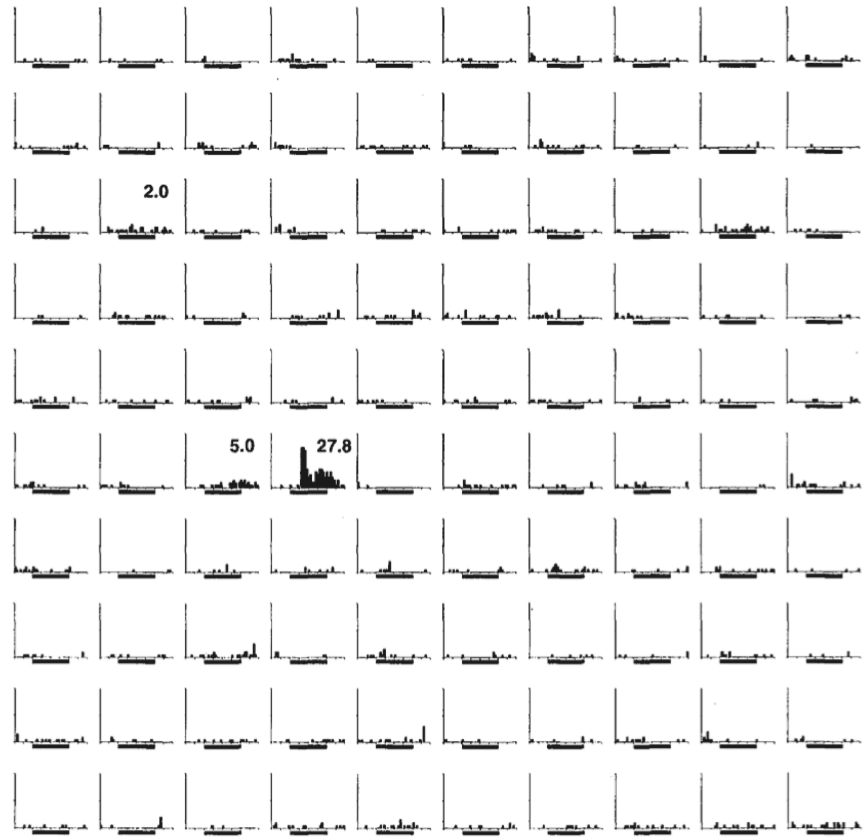

More complex representations

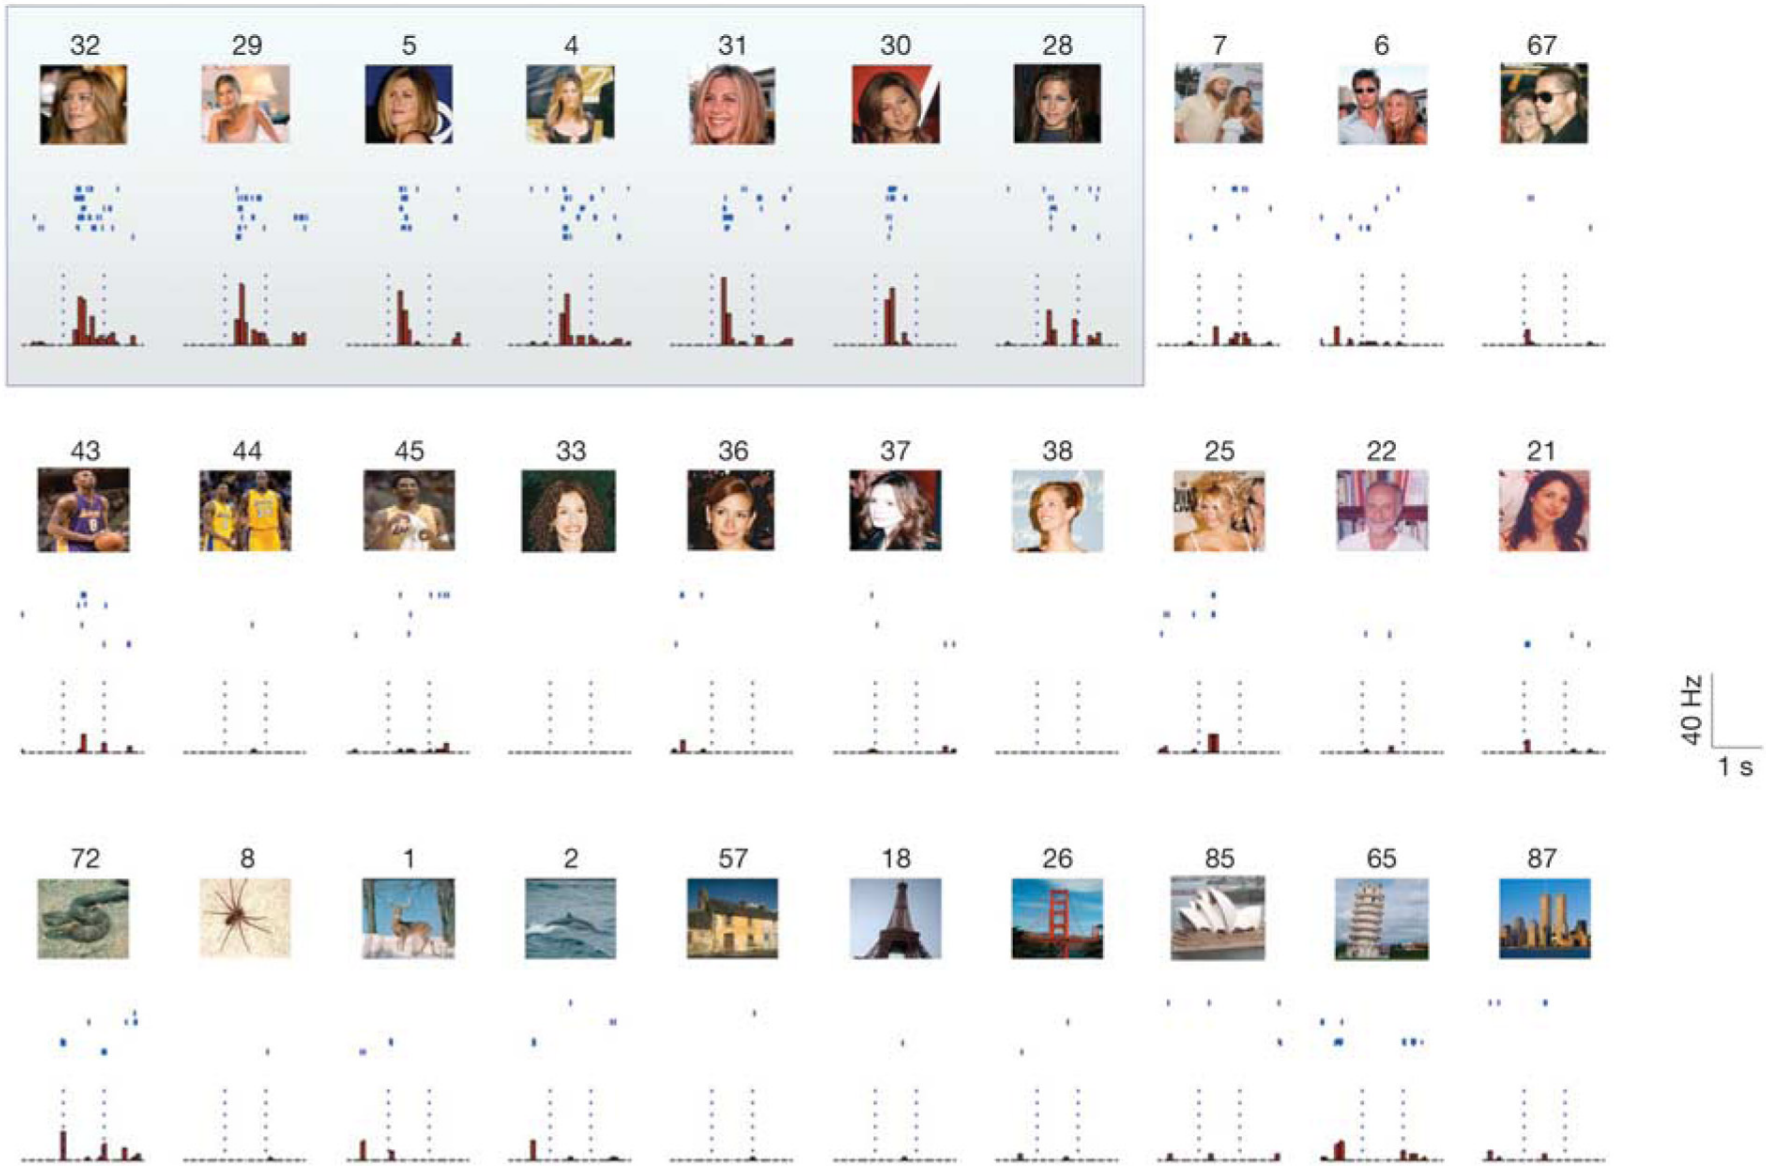

Building selective and tolerant representations

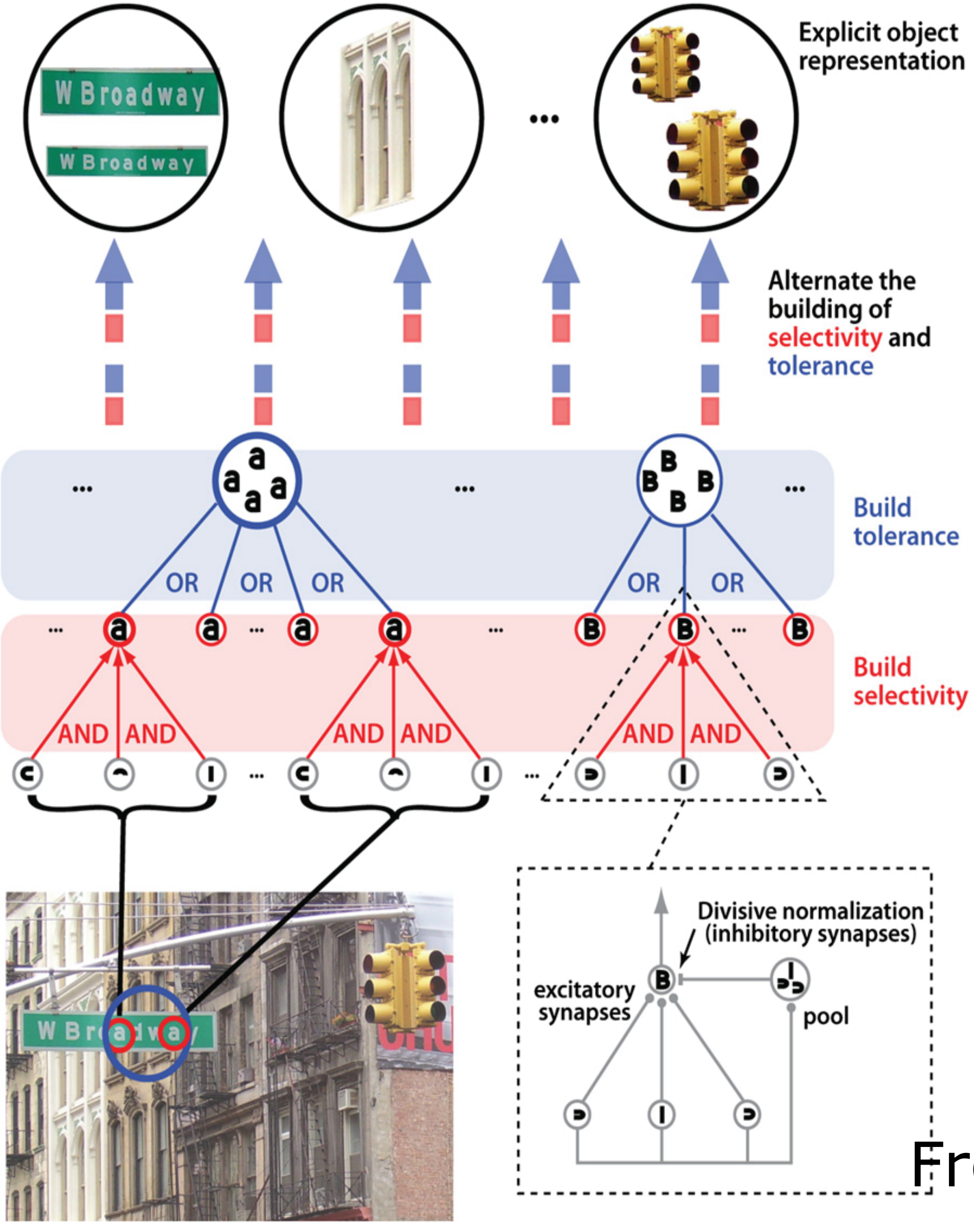

From Adam Kohn

Selectivity and tolerance increase at higher levels

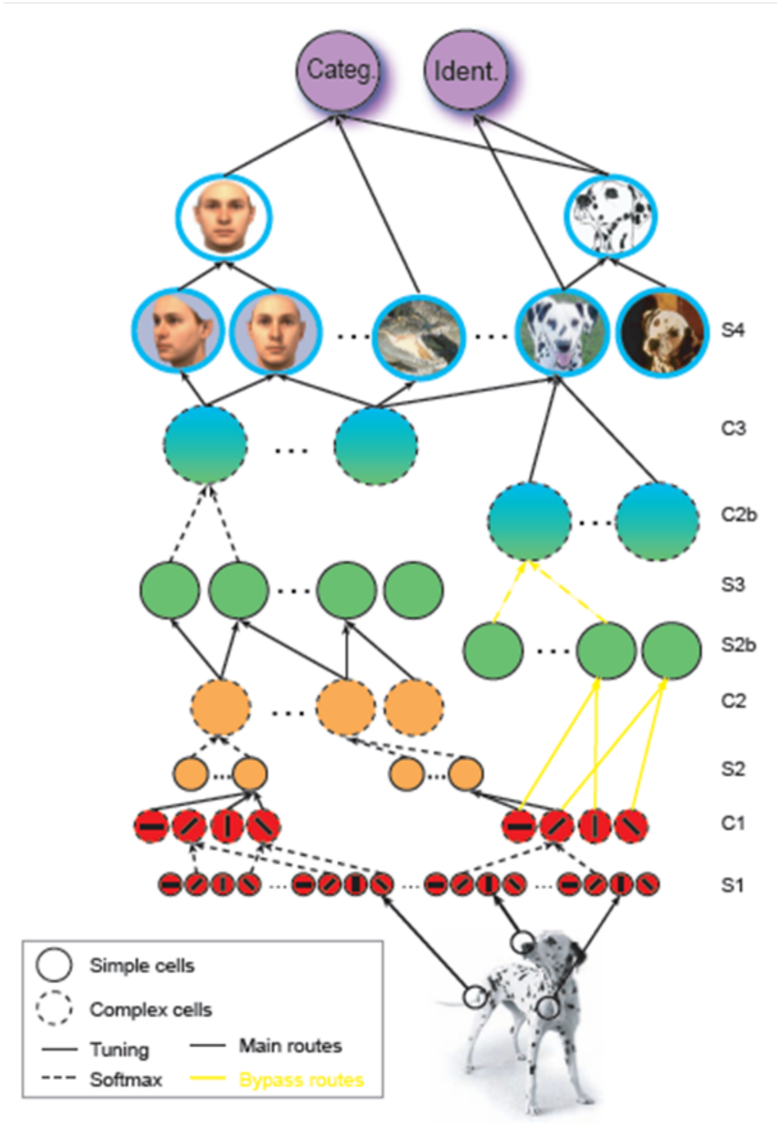

Reisenhuber and Poggio

More complex representations

What about learning from natural images beyond V1 like filters ?

Extensions to ICA

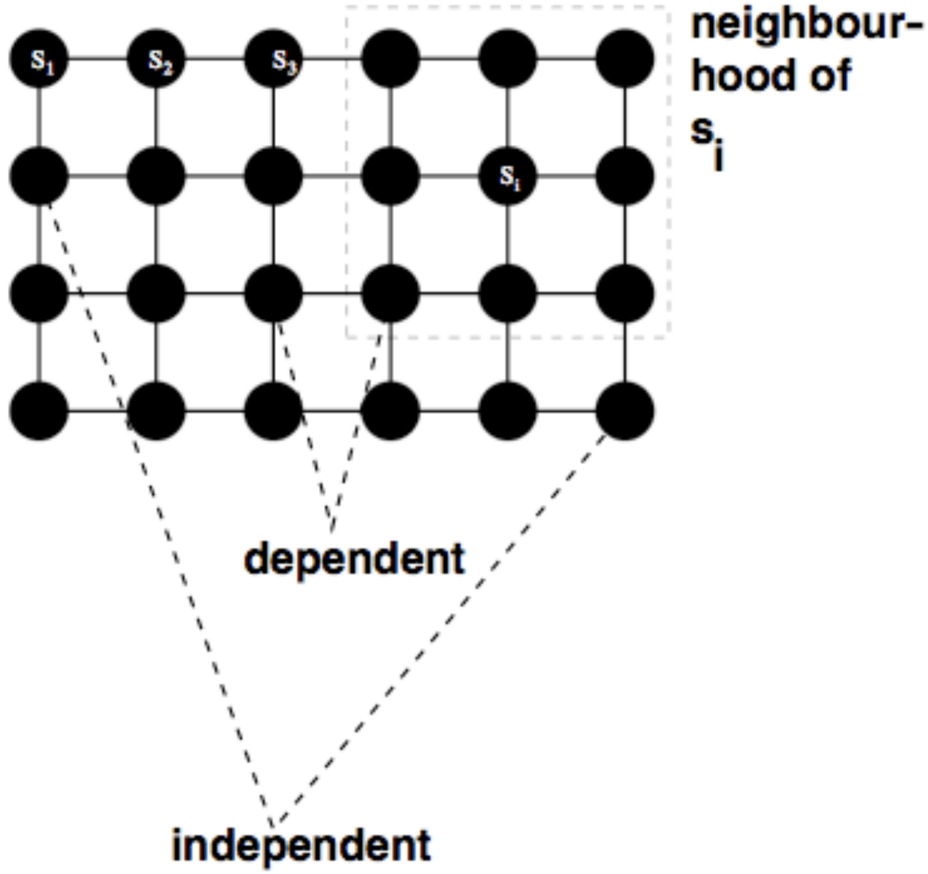

- from Hyvarinen and Hoyer; relax independence assumption; nearby units no longer independent; but different neighborhoods independent of one another...

Extensions to ICA

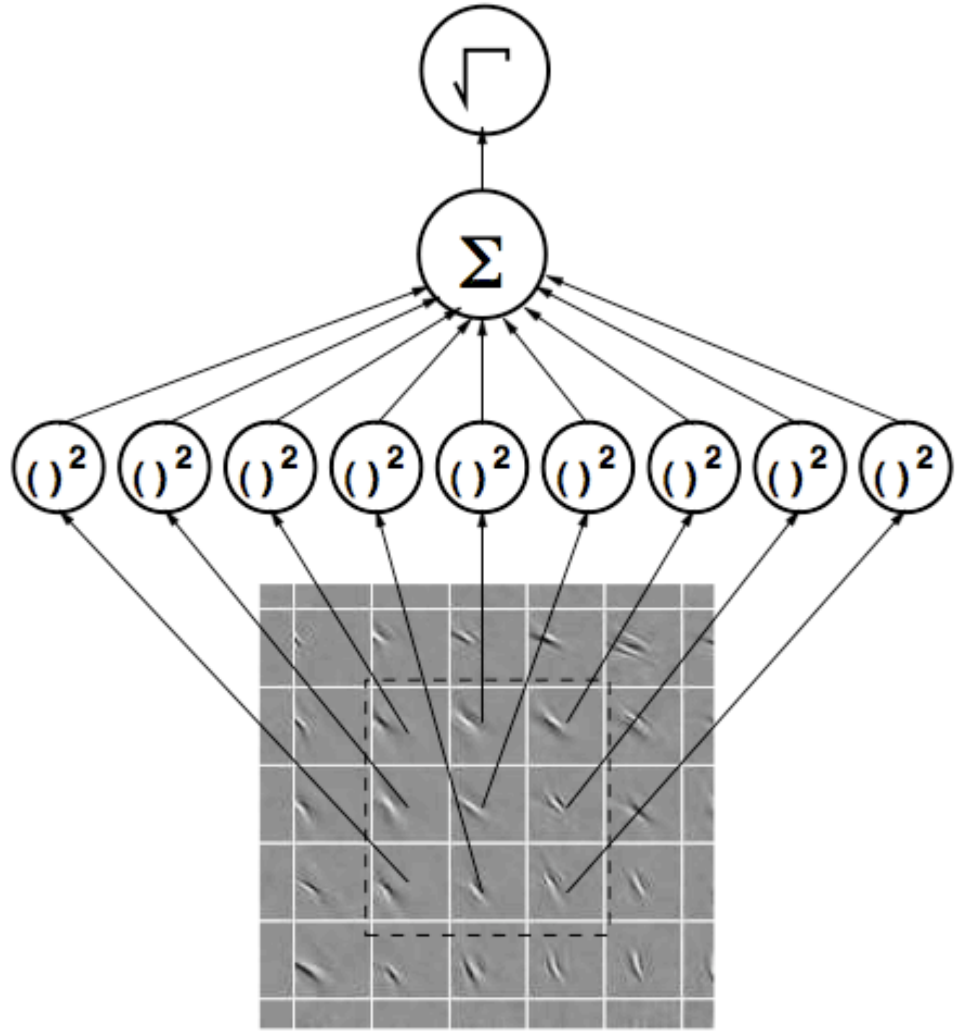

- Hyvarinen and Hoyer

Extensions to ICA

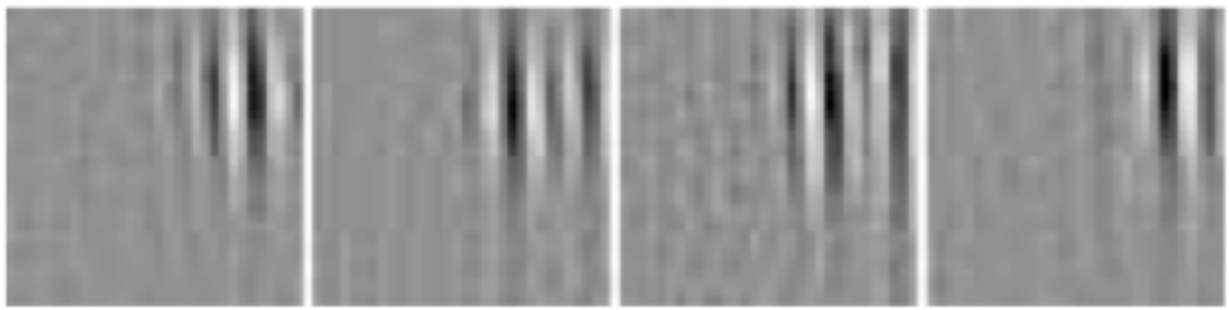

- Hyvarinen book: shown smaller group of dependent filters

Complex cell

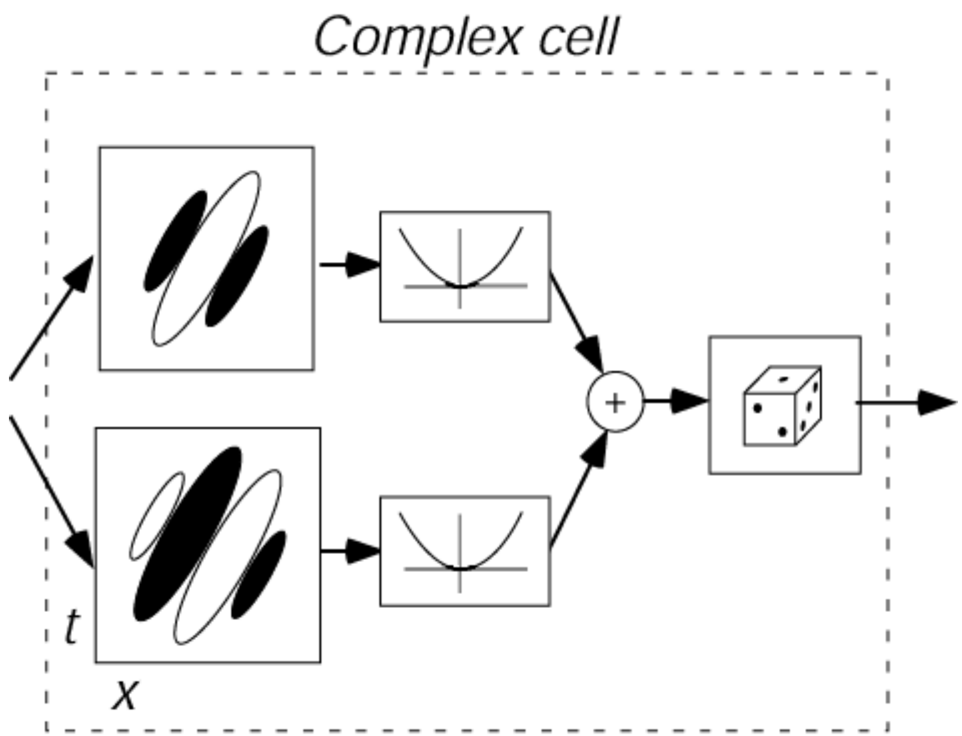

Adelson & Bergen (1985)

Relates to complex cells and invariances…

Hierarchical ICA

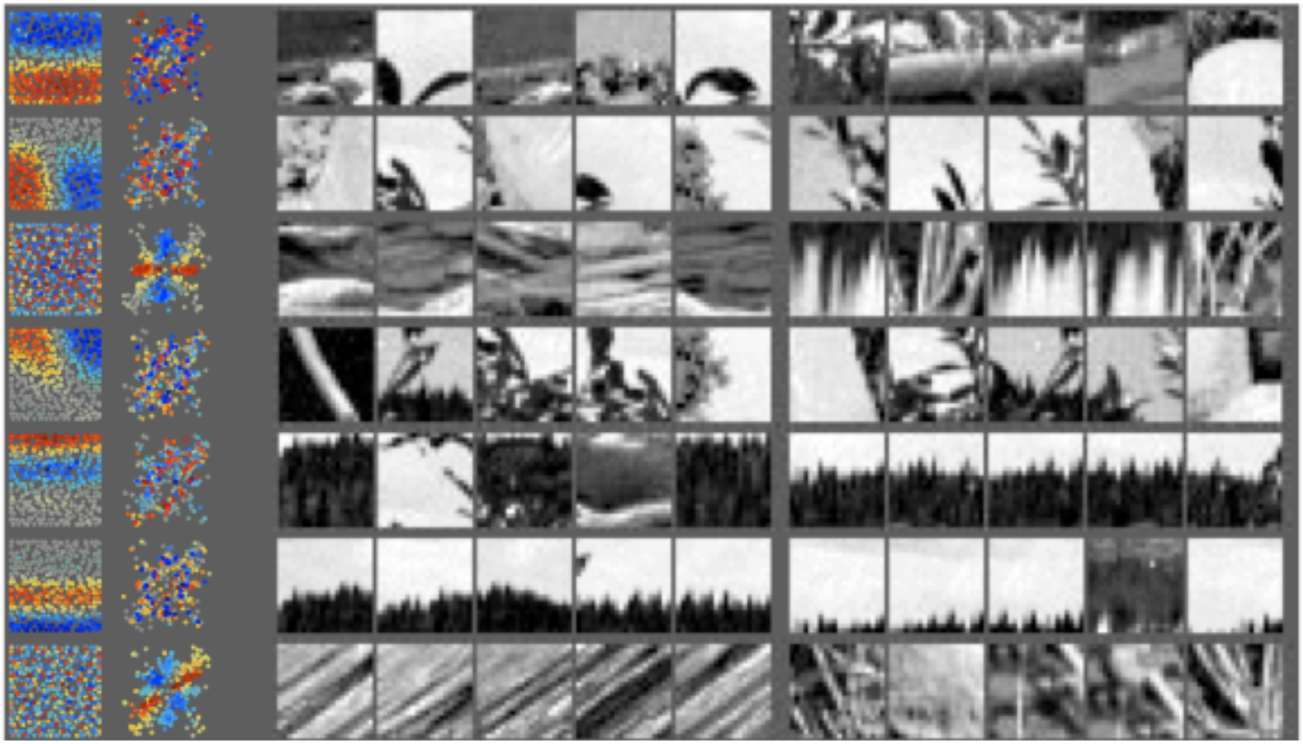

Karklin & Lewicki, 2003; 2005: higher order units are a linear combination of lower order units; learning patterns of dependencies

Hierarchical ICA

- Everything we have seen thus far: Unsupervised Learning
- There is no supervision about what object is in the image (eg, car versus tree)

Deep networks in machine learning

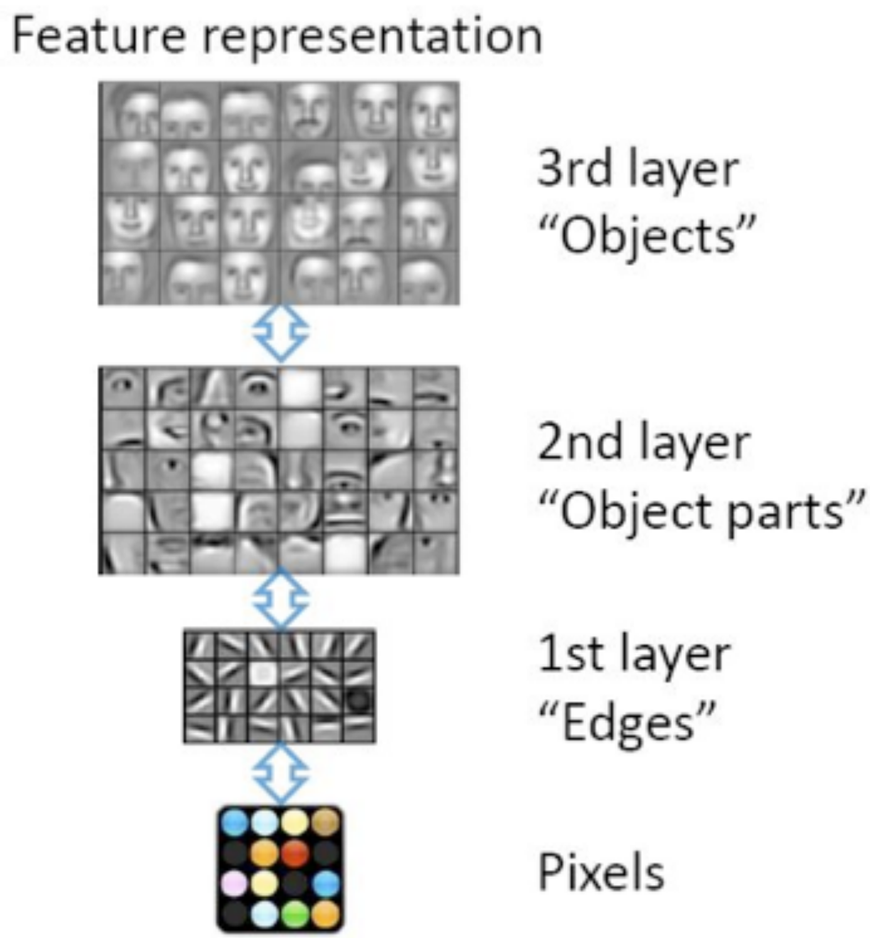

Figure from Honglak Lee, NIPS 2010 workshop

Deep networks in machine learning

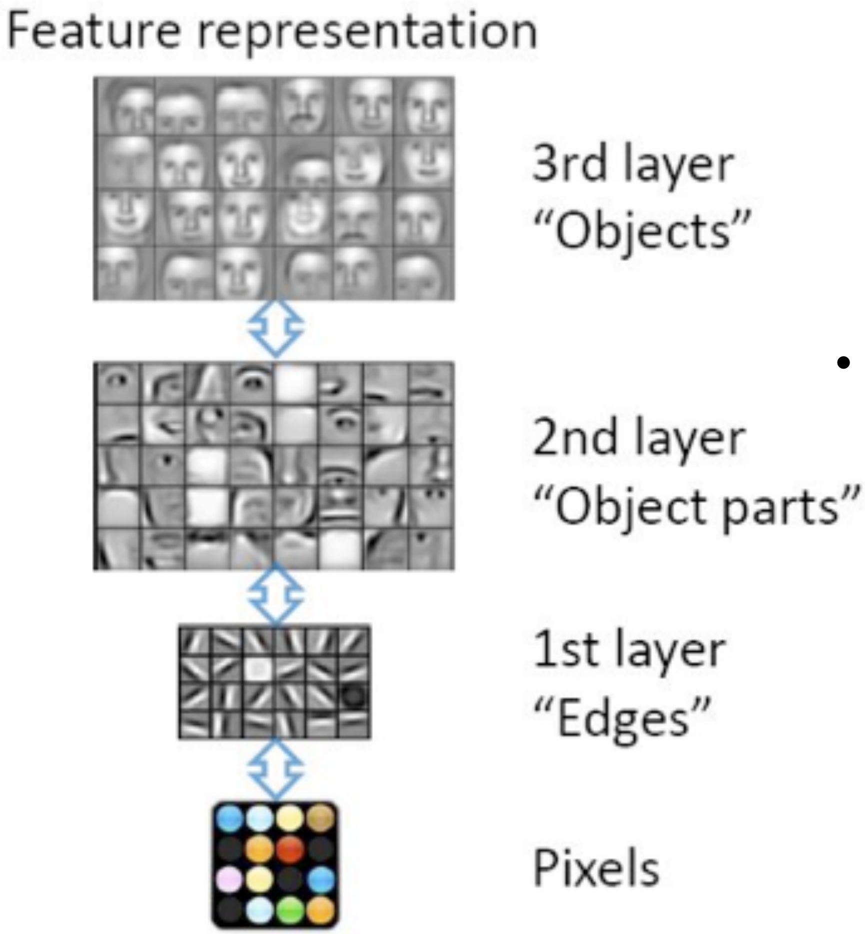

- Desirable for representing/learning multiple levels of visual processing

Figure from Honglak Lee, NIPS 2010 workshop

Deep networks in machine learning

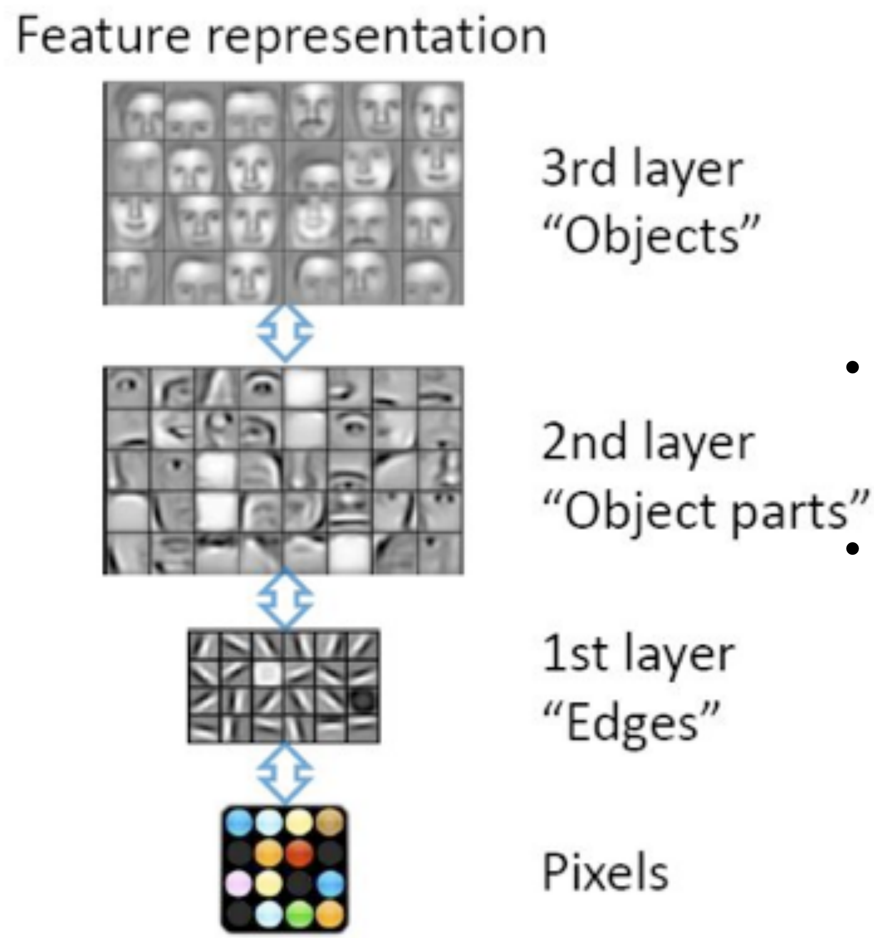

- Desirable for representing/learning multiple levels of visual processing
- Desirable for machine learning applications such as visual recognition

Figure from Honglak Lee, NIPS 2010 workshop

Deep networks in machine learning

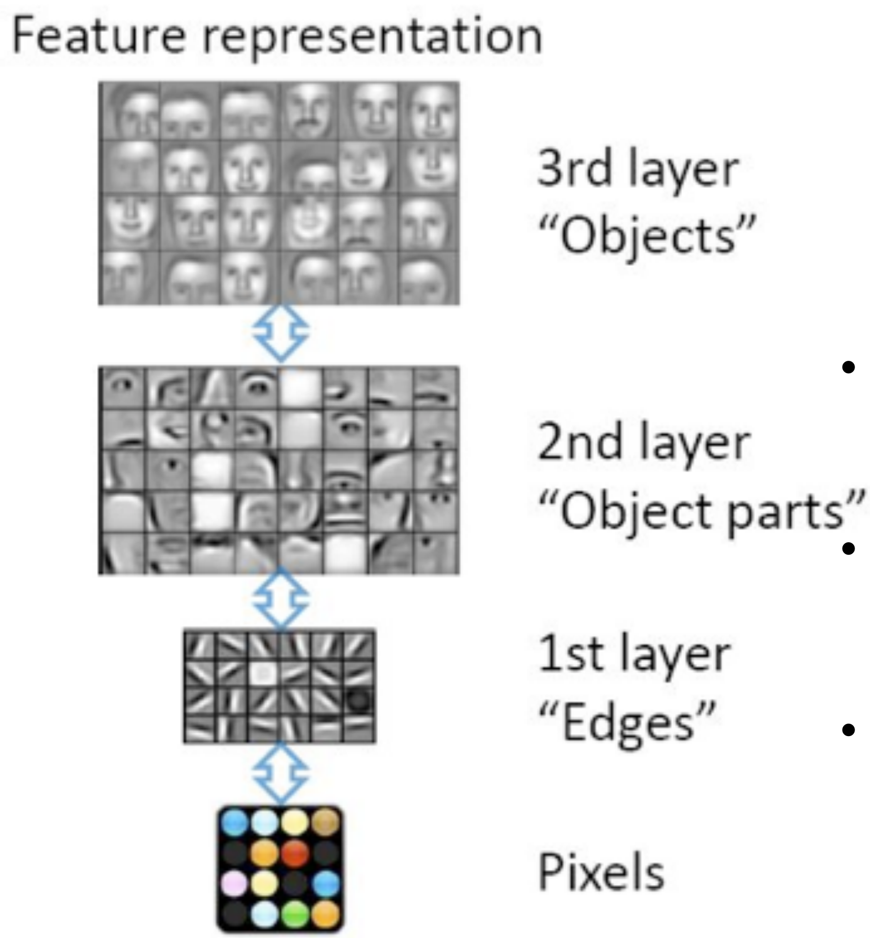

- Desirable for representing/learning multiple levels of visual processing
- Desirable for machine learning applications such as visual recognition
- More recently: interplay between deep learning in machine learning, and neuroscience

Figure from Honglak Lee, NIPS 2010 workshop

Deep networks

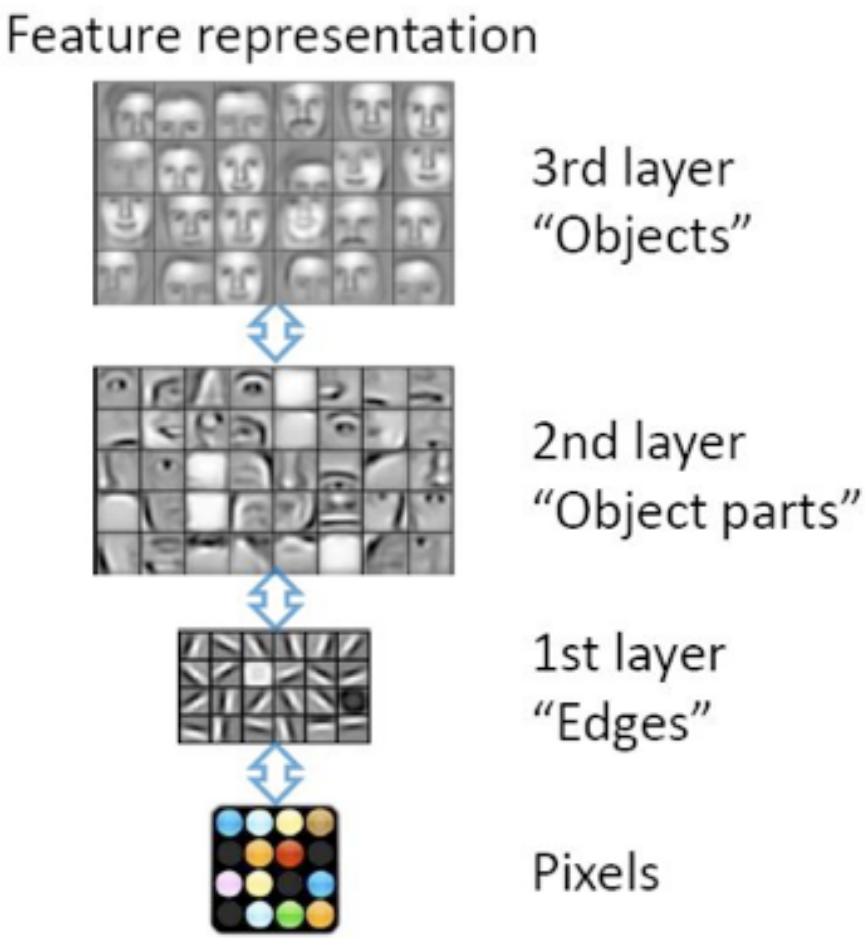

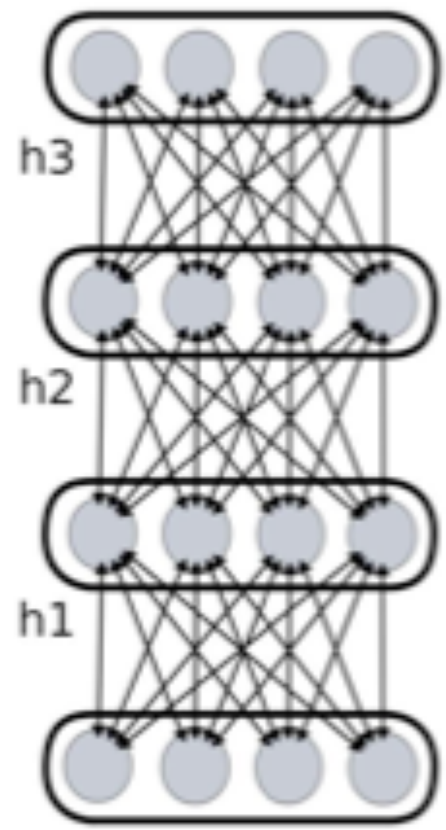

However: Training a deep network is difficult

Figures from Honglak Lee, NIPS 2010; Arnold et al. 2011

Deep networks

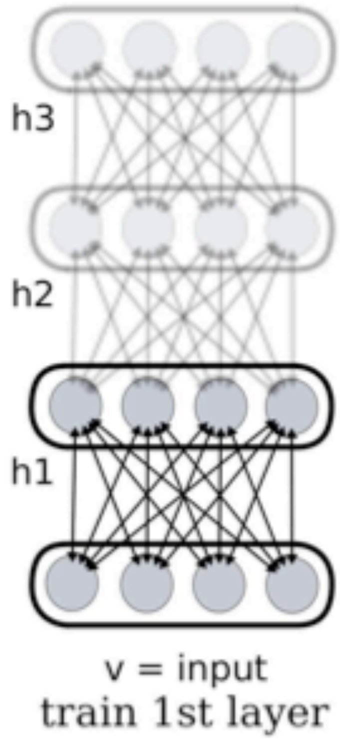

Solution: It helps to do unsupervised learning one layer at a time, and then to fine tune with supervision

Figure from Arnold et al., 2011; concept in Hinton et al. 2006

Deep networks

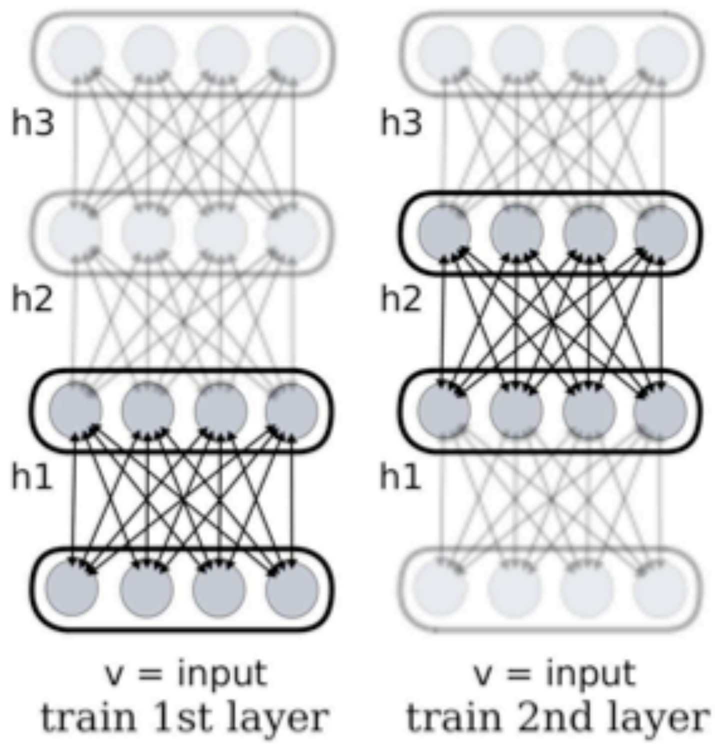

Solution: It helps to do unsupervised learning one layer at a time, and then to fine tune with supervision

Figure from Arnold et al., 2011; concept in Hinton et al. 2006

Deep networks

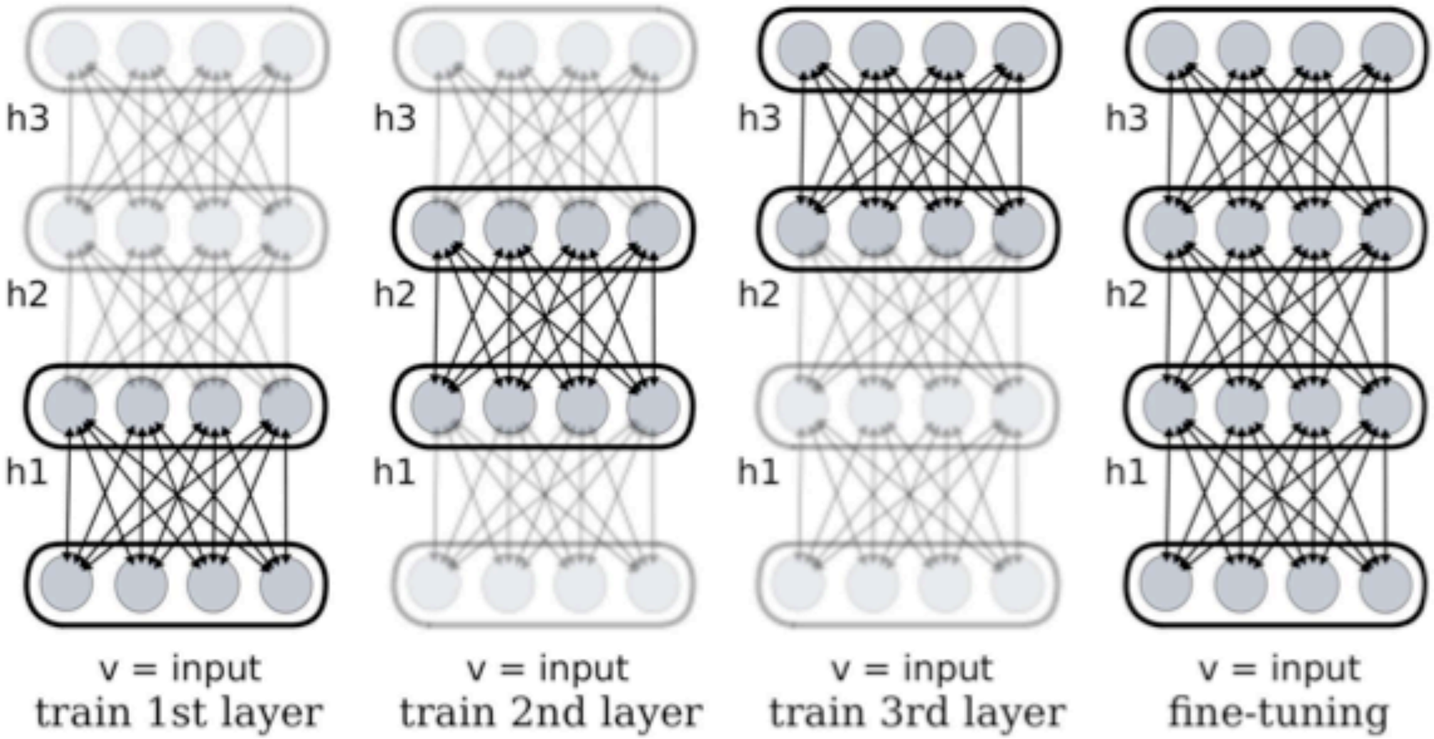

Solution: It helps to do unsupervised learning one layer at a time, and then to fine tune with supervision

Figure from Arnold et al., 2011; concept in Hinton et al. 2006

Deep networks: example unsupervised

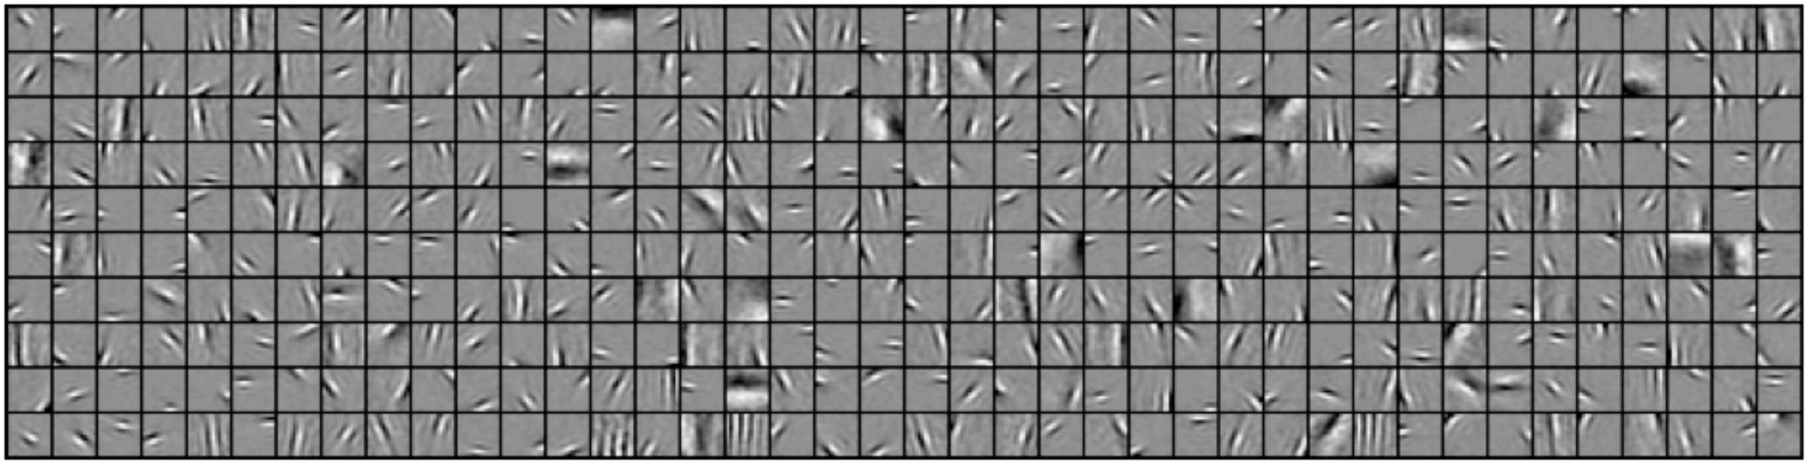

Lee, Ekanadham, NG, 2007:

- 2-layer unsupervised network with Sparsity constraint; First layer (what happens without sparsity?)

Deep networks: example unsupervised

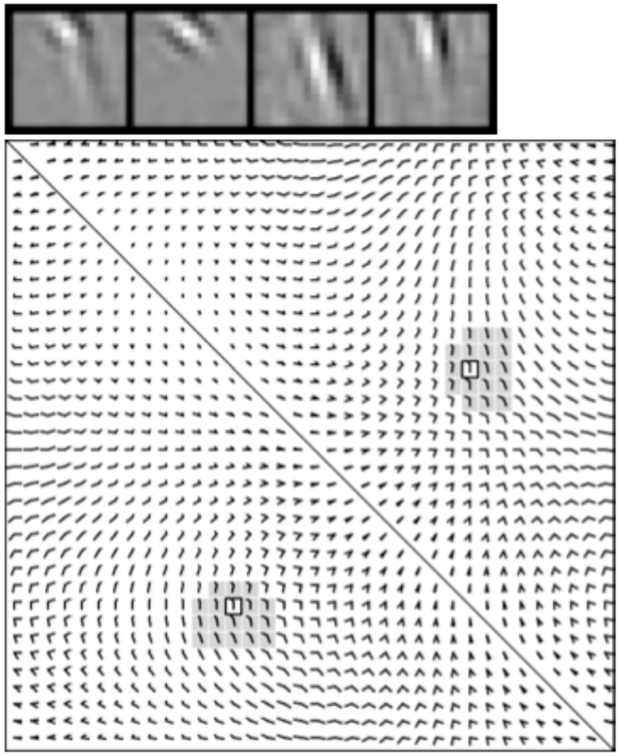

Lee, Ekanadham, NG, 2007:

• 2-layer unsupervised network with Sparsity constraint; V2-like structure

Deep networks: supervised more layers

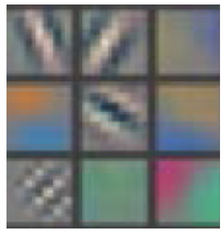

Layer 1

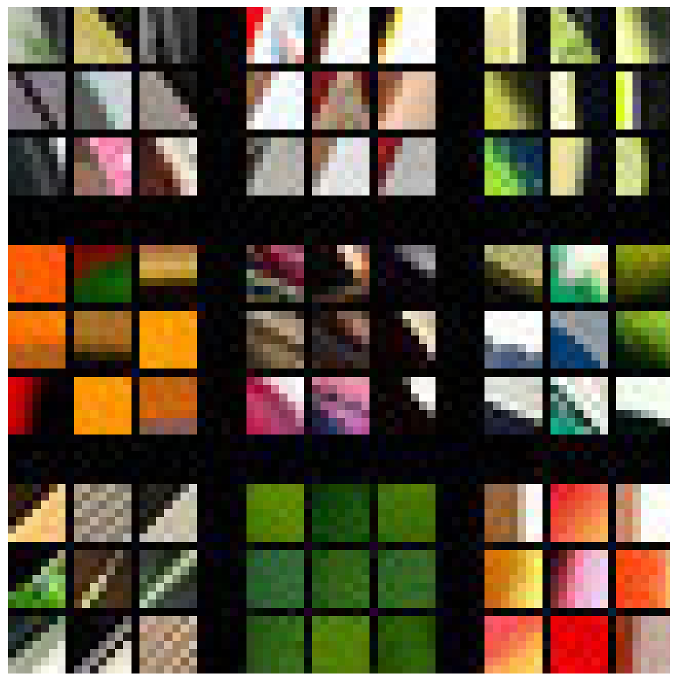

Zeiler, Fergus 2014

Deep networks: supervised more layers

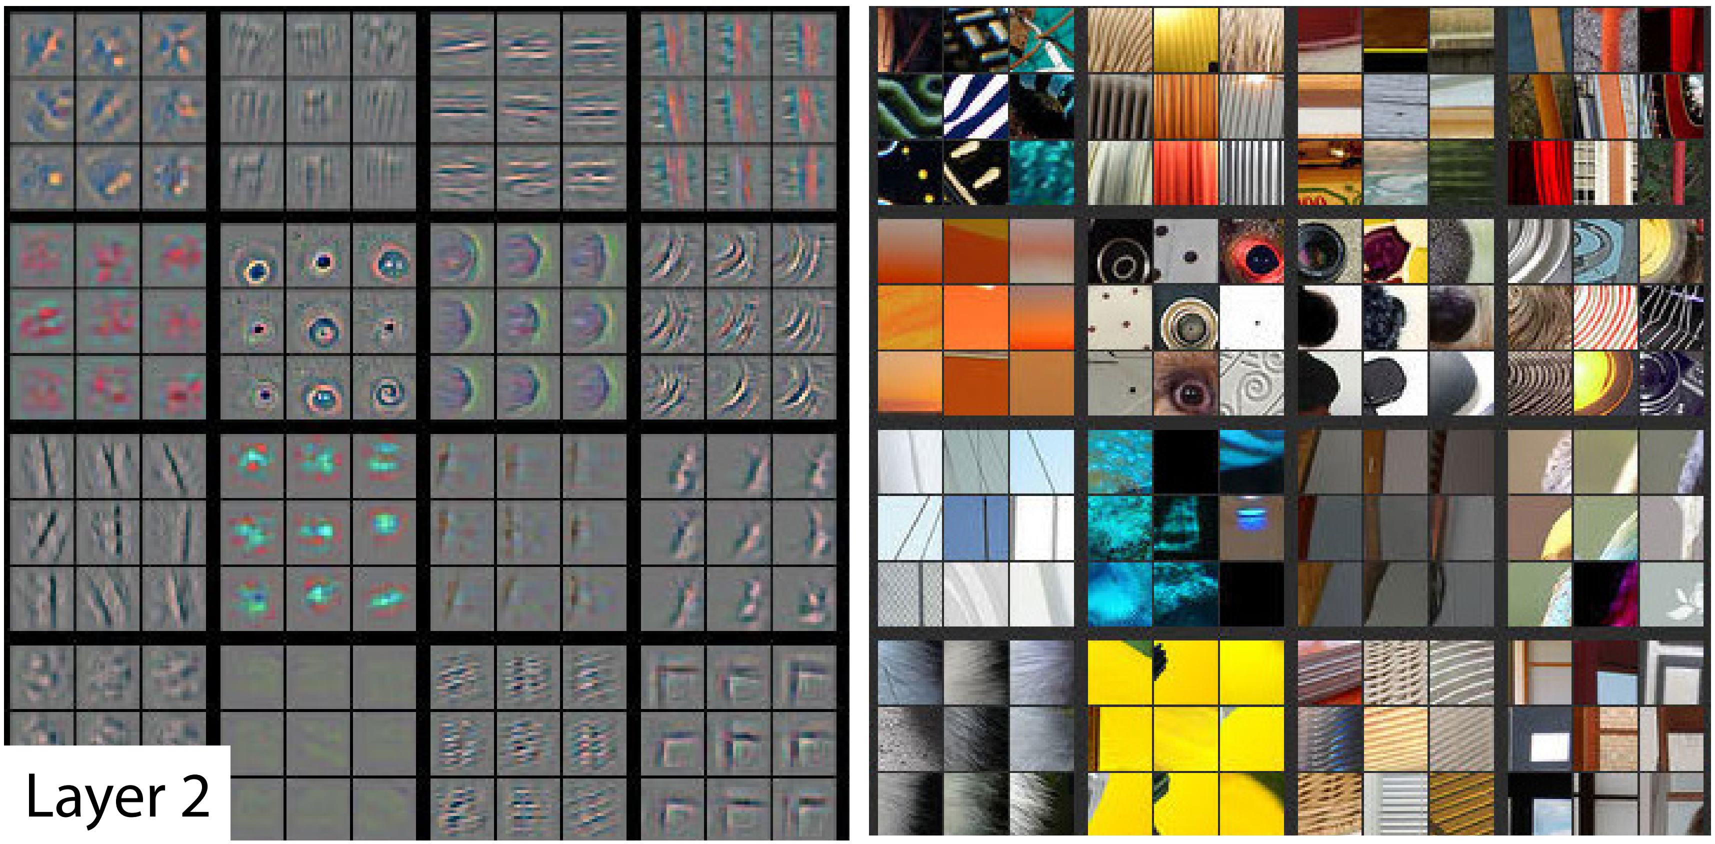

Zeiler, Fergus 2014

Deep networks: supervised more layers

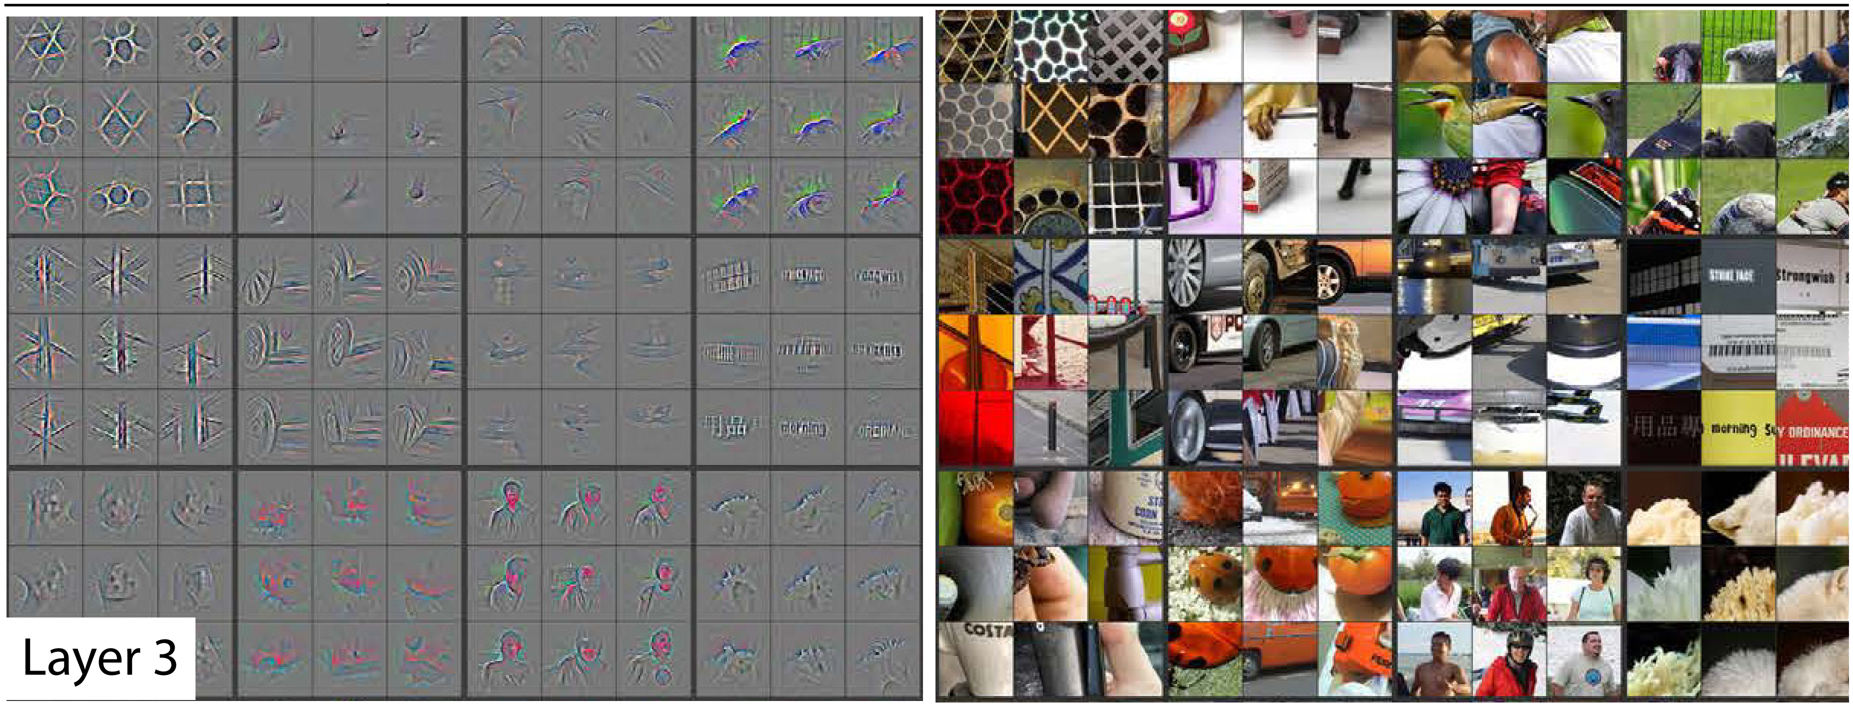

47

Zeiler, Fergus 2014

Deep networks: supervised more layers

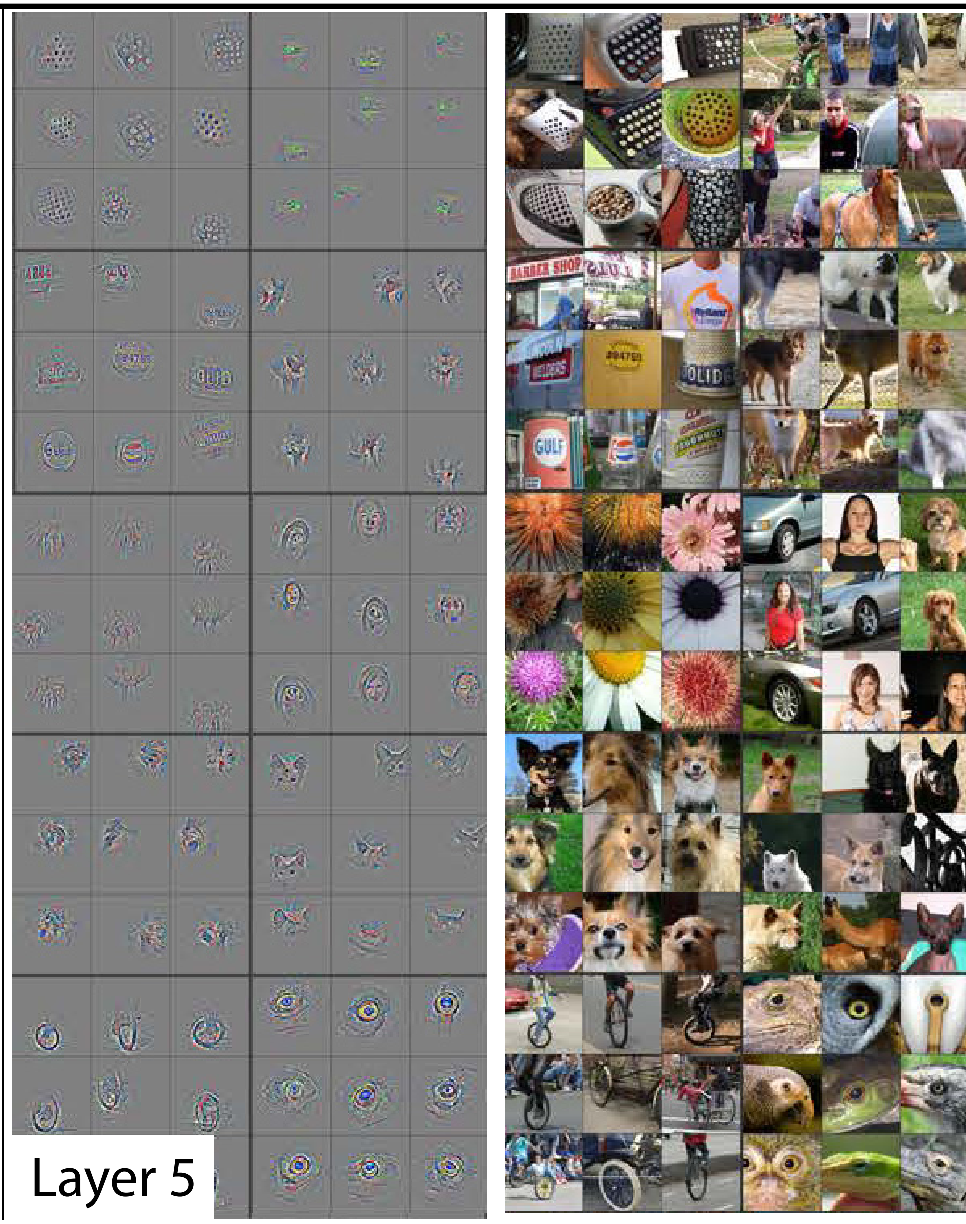

Zeiler, Fergus 2014

Deep networks: nonlinearities

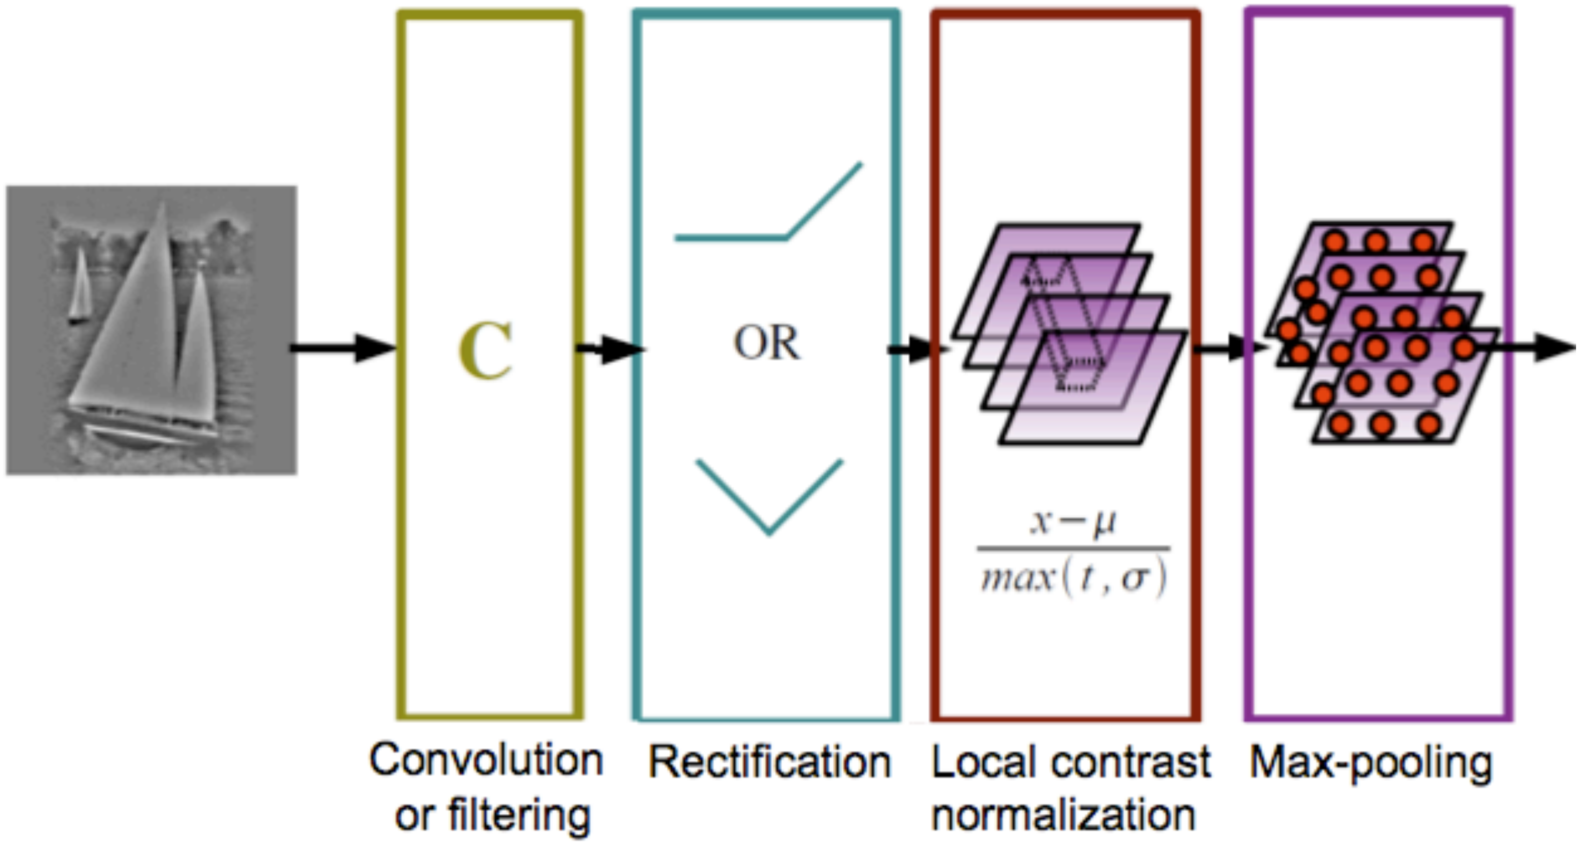

The importance of nonlinearities (From Lee NIPS 2010 workshop; Jarrett, LeCun et al. 2009)

Deep networks: nonlinearities

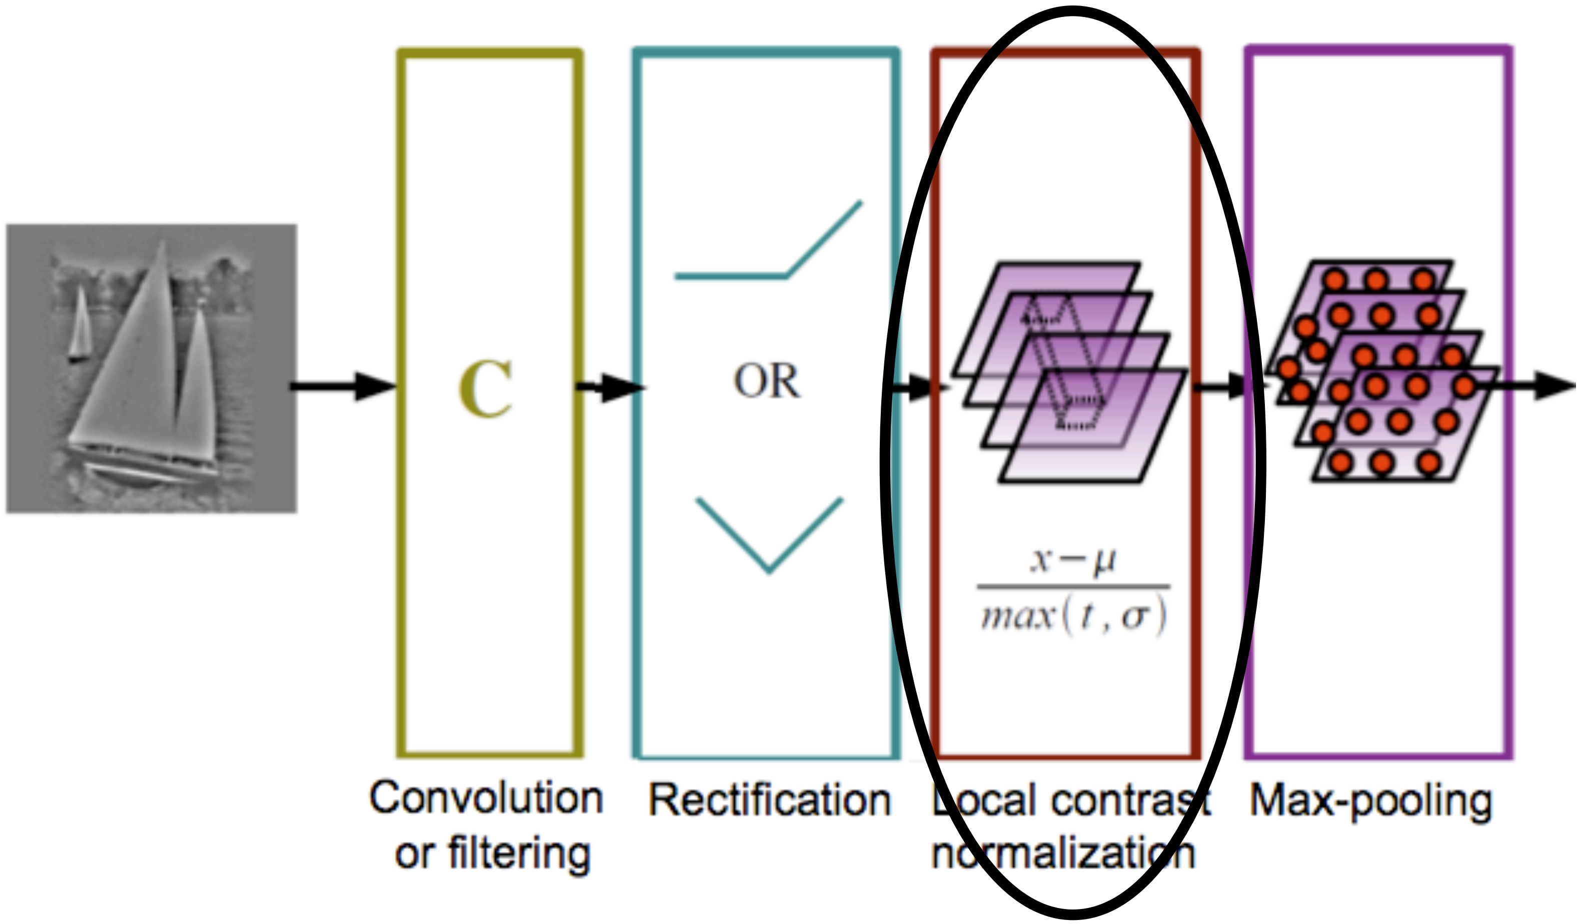

The importance of nonlinearities (From Lee NIPS 2010 workshop; Jarrett, LeCun et al. 2009)

Deep networks: nonlinearities

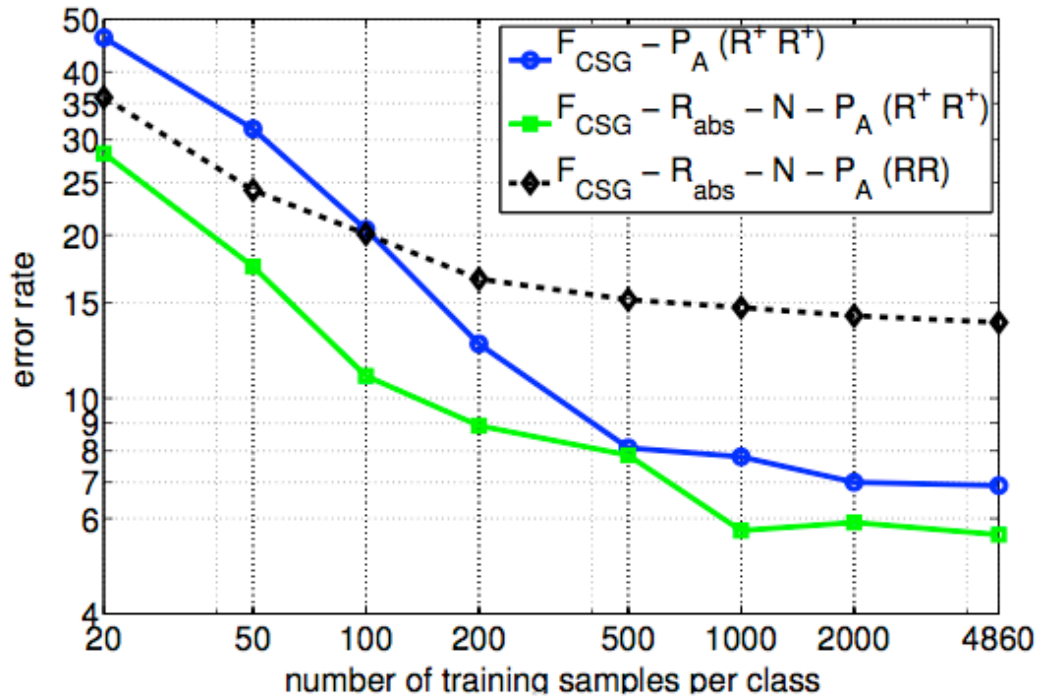

The importance of nonlinearities (Jarrett, LeCun et al. 2009)

Generative Model (nonlinear)

Modeling filter coordination in images

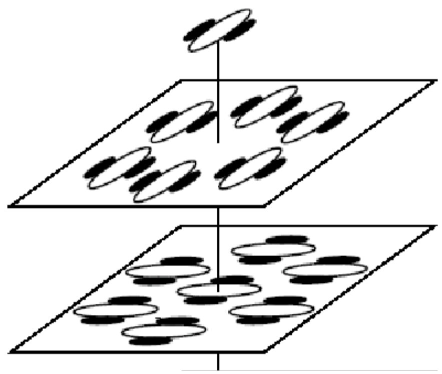

- Learning how more complex representations build up from the structure of images
- Next: Reducing dependencies further via divisive normalization