Lecture 13.4 Normative Models and Conclusion

Reading Assignments

From the Textbook

Section 13.6–13.9

Suggestions for Further Reading

Normative models: Olshausen & Field (1996), Hyvärinen & Hoyer (2001)

Q4: What is the information-processing goal of neurons in maps?

- Previous slides focused on how initial topography is formed, where feature map patterns come from, mechanistic models of these processes in animals, and detailed tests and predictions of these models.
- The models suggest some aspects of map function as well:
    - Maps provide good coverage of a multidimensional input space.
    - Lateral correlation via feature-specific connectivity helps to decorrelate neural responses.
    - Contextual modulation reveals the effects of this specific connectivity.
    - Aftereffects reveal short-term adaptation of this specific connectivity.
- Crucially, these models start from the mechanisms, so function can only be inferred indirectly.
- For a different class of *normative* models considered in the following section, functional criteria form part of the model itself—the models are derived from a stated objective, with a given neural mechanism as a specific example.

Generative models with latent variables

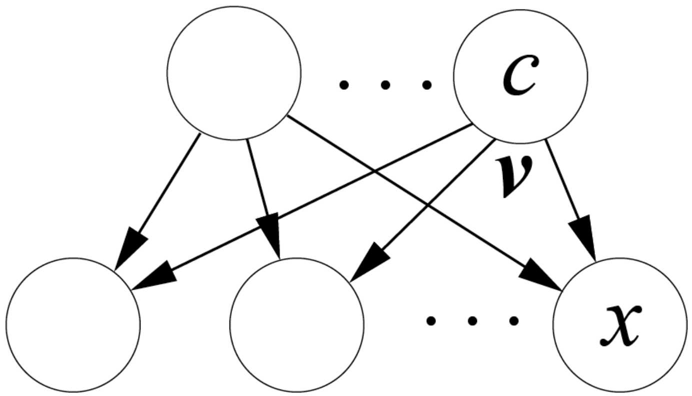

- Why would we want a map-like collection of neurons, with their specific set of visual response properties?
- One explanation: visual cortex maps form a generative model for images.
- Given an observed image $\boldsymbol{x}$, each map neuron responds with activity c such that the c-weighted sum of each weight (basis) vector $\boldsymbol{v}$ can approximate $\boldsymbol{x}$:

$$\boldsymbol{x} \approx \sum_i c_i \boldsymbol{v}_i$$

- The set of c_i coefficients models the latent variables (in the world) giving rise to that image, as estimated by the brain of the observer.
- The real latent variables are objects and light sources, but inferring all of that structure is neither tractable nor necessary for most visual tasks.
- Inferring causes for small patches of images is more tractable but still ill-posed, so it requires additional assumed constraints.

Sparse generative models for images

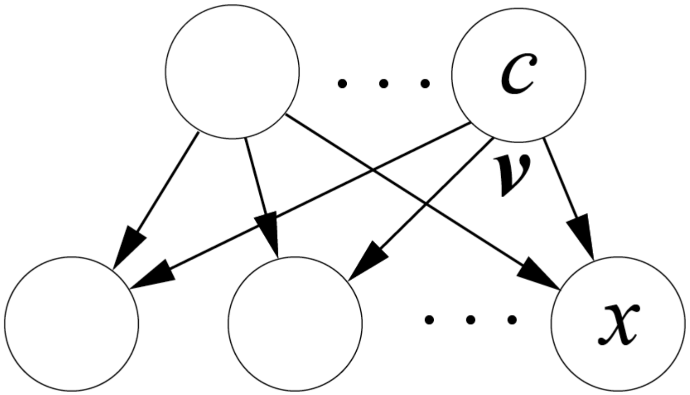

- Generative model parameters are typically set by optimizing an objective function, e.g., requiring low reconstruction error and high sparsity on ***c*** (such that most c_i are zero).
- If ***c*** is sparse, the activity of a few V1 neural responses will suffice to represent the image.
- Sparse representations reduce metabolic requirements and could be useful for further processing.
- Extreme example: Require $c_i = 0$ for all but one index *i* (winner-take-all).
- Effectively clusters the input patches, which does form V1-like Rfs.
- Reconstruction error would be high from such an approach, though.
- More biologically meaningful approach allows multiple neurons to respond, forming a faithful componential representation of the input.

Reconstruction of an image patch from basis vectors

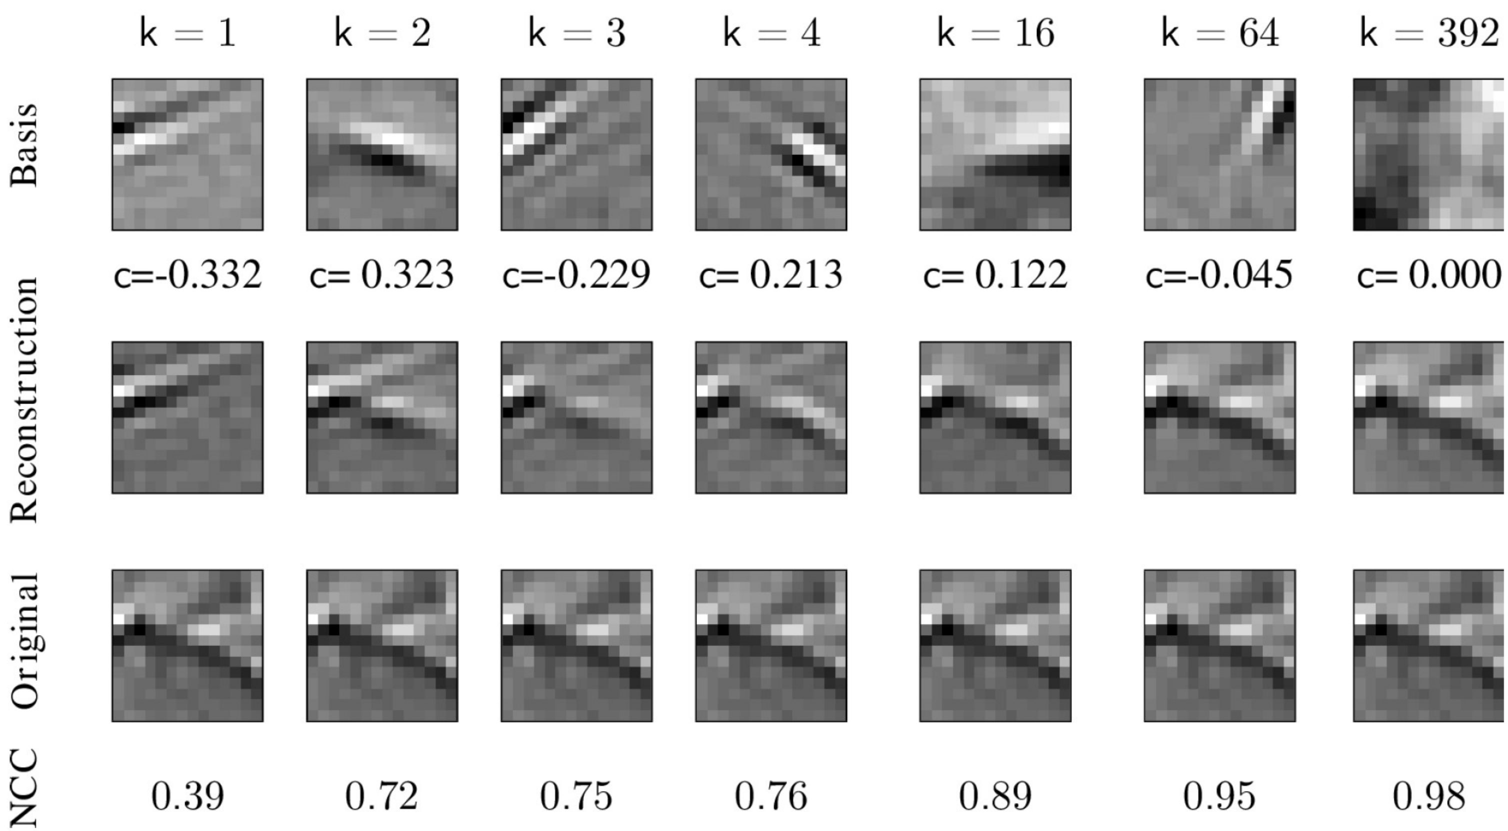

- Basis vectors approximating V1 neurons can be activated in combination to represent an image patch as a weighted sum.
- With this basis set developed to optimize sparsity, even just 4 or 16 active units gives a faithful representation of the 196 pixels in this image patch.

Developing basis vectors for image patches

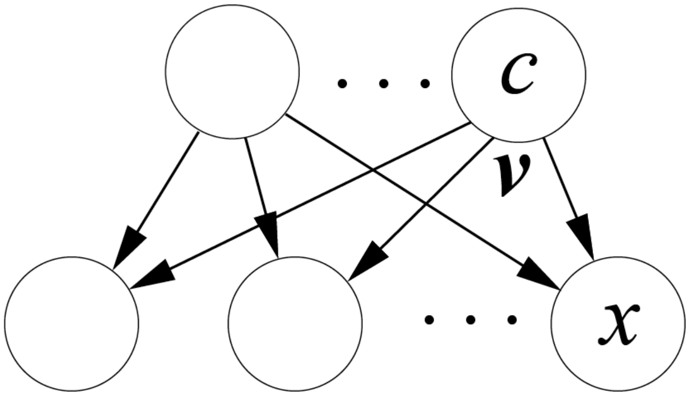

- How are these basis vectors constructed?
- Olshausen & Field (1996) model objective function:

$$|\boldsymbol{x} - \sum_i c_i \boldsymbol{v}_i|^2 + \lambda \sum_i g(c_i)$$

- The model iteratively minimizes the reconstruction error $|\boldsymbol{x} - \sum_i c_i \boldsymbol{v}_i|^2$ and a penalty function $g(c_i)$, balanced with λ.
- This non-neural process involves estimating the coefficients $\{\hat{c}_i\}$ that best explain a given training image, then updating the basis vectors $\boldsymbol{v}_i$ to minimize the reconstruction error.
- Alternative implementations are more easily related to neural mechanisms, but this formulation explicitly identifies what the network is trying to achieve: good reconstruction from a small number of active units.

Macaque and generative model V1 RFs

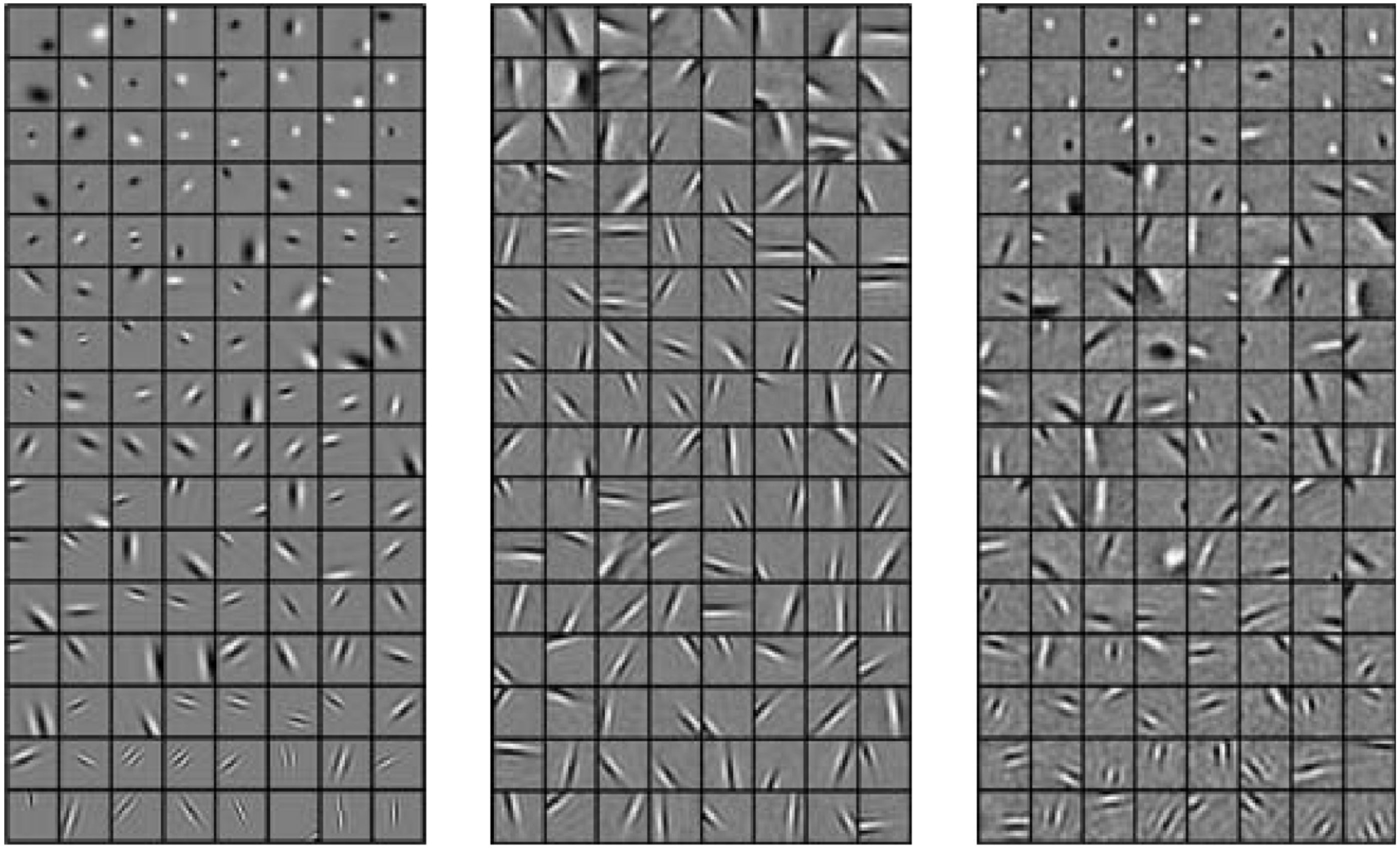

(a) Macaque (b) (Olshausen & Field, 1996) (c) SSC

(a) Macaque RFs (here fitted with Gabors) span a wide range of shapes.

(b) Original Olshausen model projective fields do not, but

(c) They can with a different sparseness criterion.

Discussion

- A wide variety of map models have been proposed, addressing different aspects of map function and development:
  - Chemoaffinity models of topographic mapping
  - Geometrical models of multidimensional feature spaces
  - Mechanistic models of feature map development
  - Mechanistic models of visual phenomena
  - Normative models expressed in information-processing goals
- Complete explanations would address all these aspects consistently, but important gaps and some minor conflicts remain among all of these approaches.
- Explanations of why map patterns should be smooth and what functional purpose is served by the specific map patterns remains speculative.
- Whether the map patterns themselves serve a function or not, the patterns do provide important constraints for models of map development, as they reflect the underlying function-related circuitry.

Conclusions

- Neural maps are ubiquitous but diverse.
- Many neural maps are topographic, mapping the sensory surface.
- Topography is established under genetic control, with activity-based refinement.
- Maps for other features appear to develop as a result of neural activity.
- Geometrical models suggest that feature-map patterns result from folding a multidimensional input space onto a 2D cortical space.
- Mechanistic implementation of such dimensionality reduction can explain lateral connectivity, contextual modulation, and aftereffects.
- Normative models suggest that neurons uncover latent causes in sensory inputs, forming a sparse representation suitable for further computation and action.
- Future work can focus on the relationships between geometrical, mechanistic, and normative models of maps.