A Computational Approach To Understanding The Response Properties Of Cells In The Visual System

David C. Noelle
Assistant Professor of Computer Science and Psychology
Vanderbilt University
david.noelle@vanderbilt.edu

March 3, 2004

Who Is This Guy?

My research involves the construction and analysis of *computational models* of human behavior and brain function.

I am particularly interested in the phenomena surrounding learning and memory.

- learning from both direct instruction and practice
- the role of rule-like knowledge in concept learning
- implicit/explicit learning
- working memory
- prefrontal cortex, the basal ganglia, and rapid explicit learning
- the shaping of representations in prefrontal cortex

What can a machine learning researcher tell you about the visual system?

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Issues Of Scale In Computational Modeling

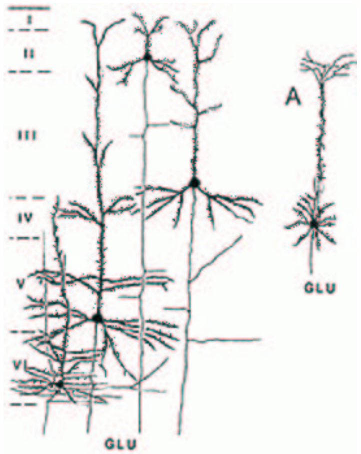

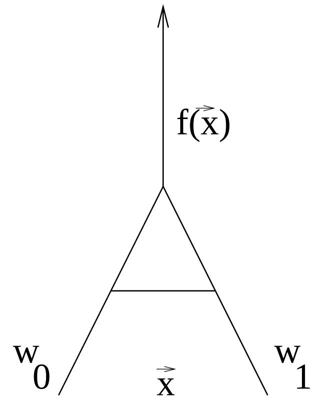

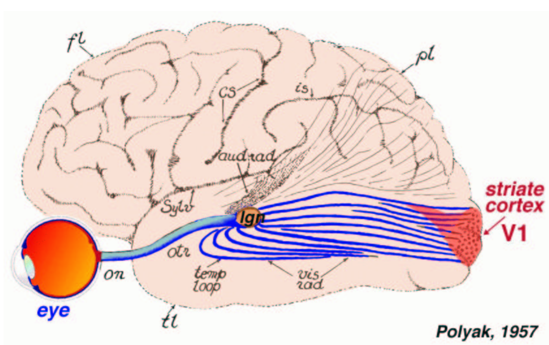

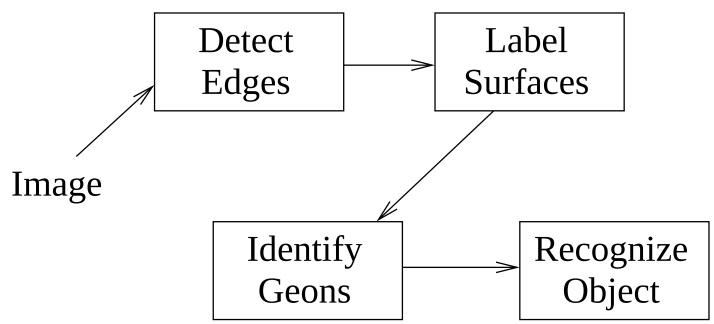

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Levels Of Scale

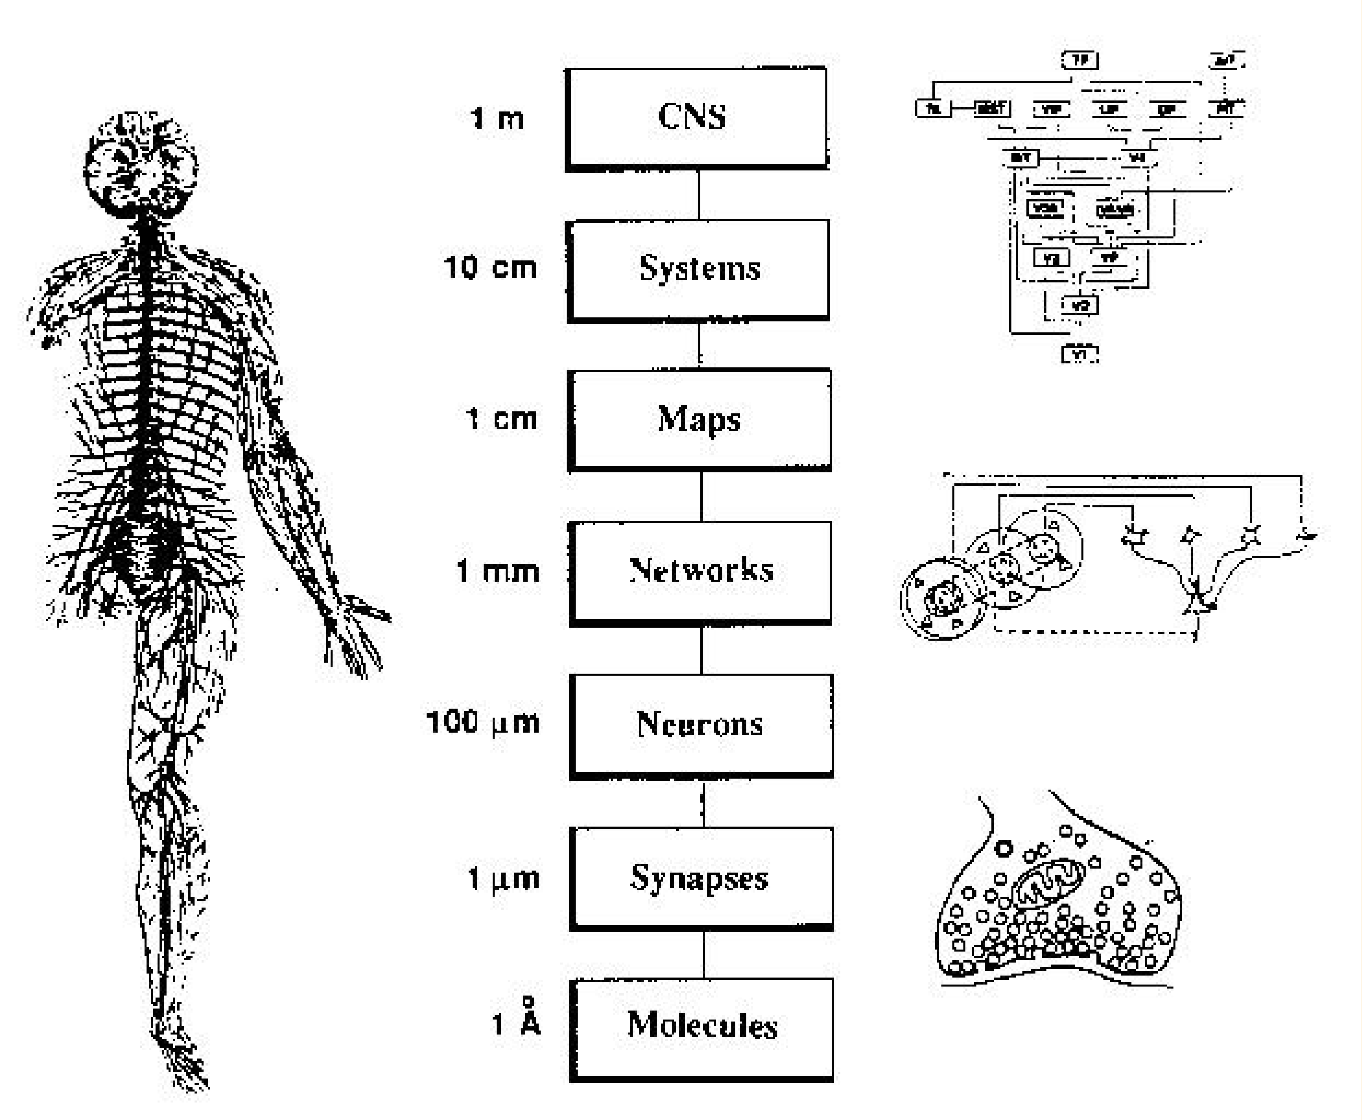

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Levels Of Task Abstraction

- *Computational Theory* — “What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?”

- *Representation and Algorithm* — “How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation?”

- *Hardware Implementation* — “How can the representation and algorithm be realized physically?”

(Marr, 1982)

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Rational Analysis

1. Identify the task faced by the system.
2. Precisely characterize the objective of the task.
3. Enumerate the essential properties of solutions which ideally meet this objective.
4. Look for these properties in the system being studied.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

The Questions Of The Day

What does a rational analysis of the task of representing natural scenes tell us about the properties of an ideal vision system?

To what degree are those properties present in the human visual system?

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Forget Studying Vision! Let's Party!

How can you separate one source from the background noise?

How can you separate all sources from each other?

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Two General Strategies For Separating Sources

- *Model The Sources*

  Use knowledge about the different sources in order to associate different components of the received signal with those sources.

- *Model-Free*

  Carefully examine the received signal for statistical patterns which suggest contributions from independent sources.

  This is called *blind source separation*.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

How Is Blind Source Separation Possible?

Imagine two source variables (e.g., sound amplitudes) from two independent sources. Imagine two sensors (e.g., microphones), the first of which picks up 80% of the first source and 20% of the other, and the second of which picks up an 80/20 split favoring the other source.

Imagine that the sources are completely unstructured, acting essentially as random variables, sampled uniformly from some fixed range.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Independent Components

If we plot readings from the two sensors, we can make out the *independent components* of the distribution of sensor readings — the set of axes that have the property that knowing a data point's coordinate along one axis tells you nothing about its coordinate along the others.

These independent components correspond to the two sources.

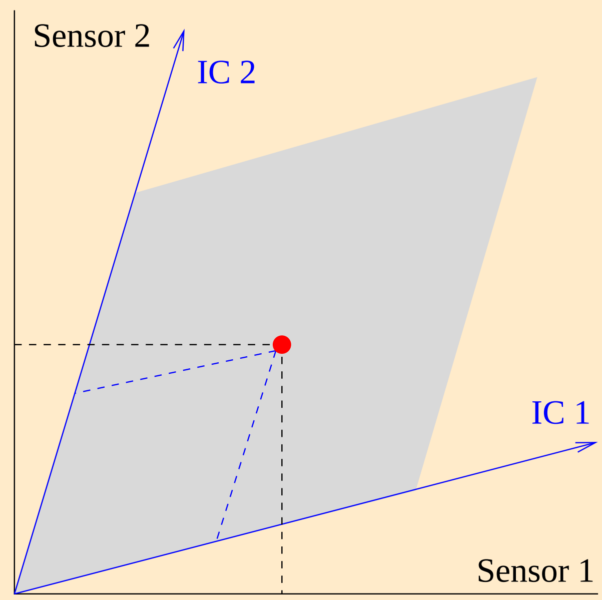

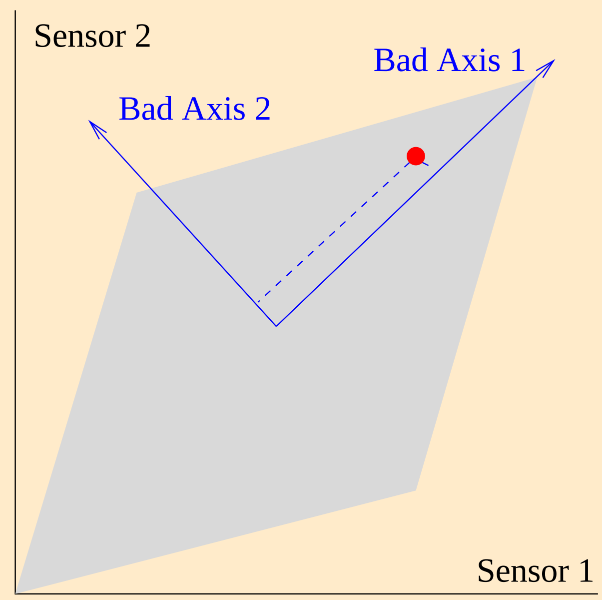

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Formalizing The Problem

Assume each sensor receives a linear combination of the sources, but we don't know the weightings of these linear mixtures. Find the *unmixing* weights — the weights which determine how to linearly combine the sensor readings in order to recreate the sources.

(If we knew the matrix of mixing weights, we could find the matrix of unmixing weights by computing the matrix inverse!)

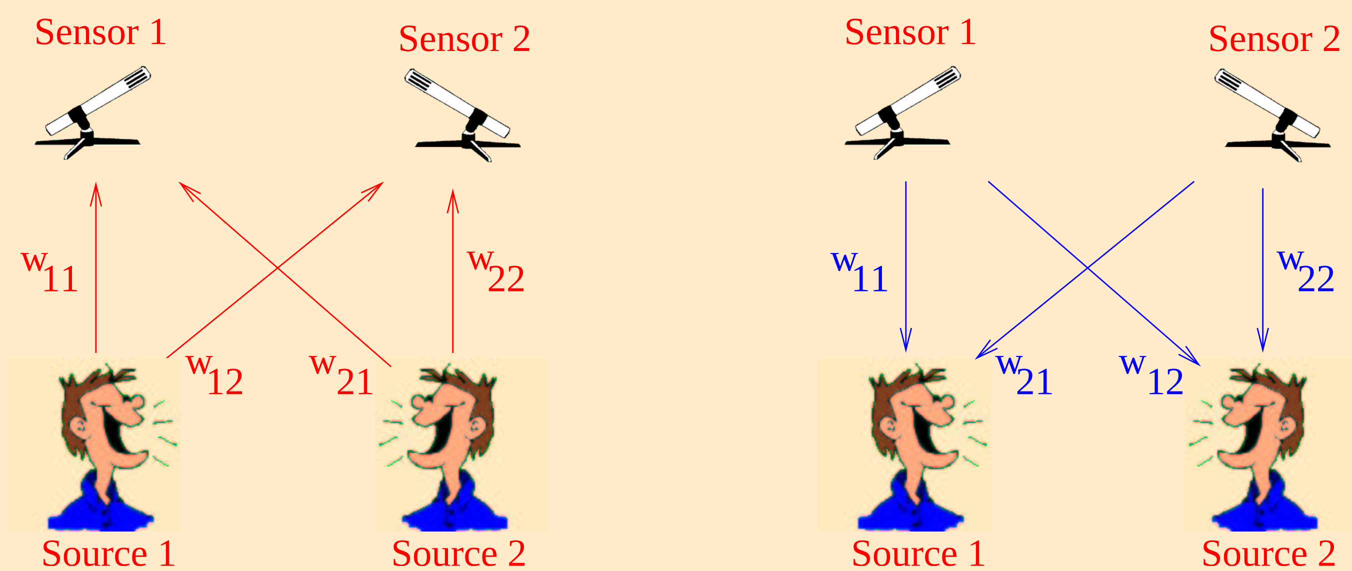

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

An Information Theoretic Objective

Since we don't know the original sources, how can we find the unmixing weights that recreates them?

We search for the weights that, given our sensor readings (**x**), produce maximally independent reconstructed signals (**y**).

This objective of independent reconstructed signals can be formalized as finding the set of weights which maximizes the *entropy* of the signals, **y**.

$$H(Y) = -\sum p(\mathbf{y}) \log p(\mathbf{y})$$

$$I(y1, y2) = \sum \sum p(y1, y2) \log \left(p(y1, y2) \;/\; p(y1)p(y2) \right)$$

$$H(y1, y2) = H(y1) + H(y2) - I(y1, y2)$$

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Learning Weights

Our goal, then, is to find weights mapping from our mixed signals to the reconstructed sources which maximize the entropy of the resulting reconstructed signals. How do we find such weights?

We can use a numerical optimization technique called *gradient ascent*, which is a common machine learning technique.

1. start with small random weights
2. sample the sensors and compute the entropy, H(Y), of the resulting reconstructed signals
3. calculate the partial derivative of the entropy with respect to each weight (e.g., (cof w / det W) + x (1 - 2y))
4. adjust each weight in proportion to its derivative, and go to step 2

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

The Results

This strategy is called Independent Components Analysis (ICA), and it does a pretty good job at the task of blind source separation.

(Bell & Sejnowski, 1995)

What Does This Have To Do With Vision?

The task faced by early parts of the visual system might be seen as one of representing the current scene in an efficient manner. One sign of lack of efficiency is how redundant the internal representation is. Thus, a central task of early visual processing might be the reduction or removal of redundant information.

Independent components analysis (ICA) excels at reducing redundancy. It seeks out underlying source signals which are statistically independent — knowing about one source tells you nothing about the others. Thus, it explicitly attempts to minimize redundancy across the sources.

Consider that a portion of an image falling on your retina might be viewed as being composed of a weighted sum of independently appearing visual features — knowing if one feature is present tells you nothing about the presence of the others. By applying ICA to natural images, we can find out what these independent features may be like.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Formalizing The Task

Each feature detector responds with an activity level which is a monotonic function of the weighted sum of all input pixel intensity values. What weight values cause the feature detectors to detect statistically independent features, thereby encoding the image across the feature detectors in a manner which minimizes redundancy?

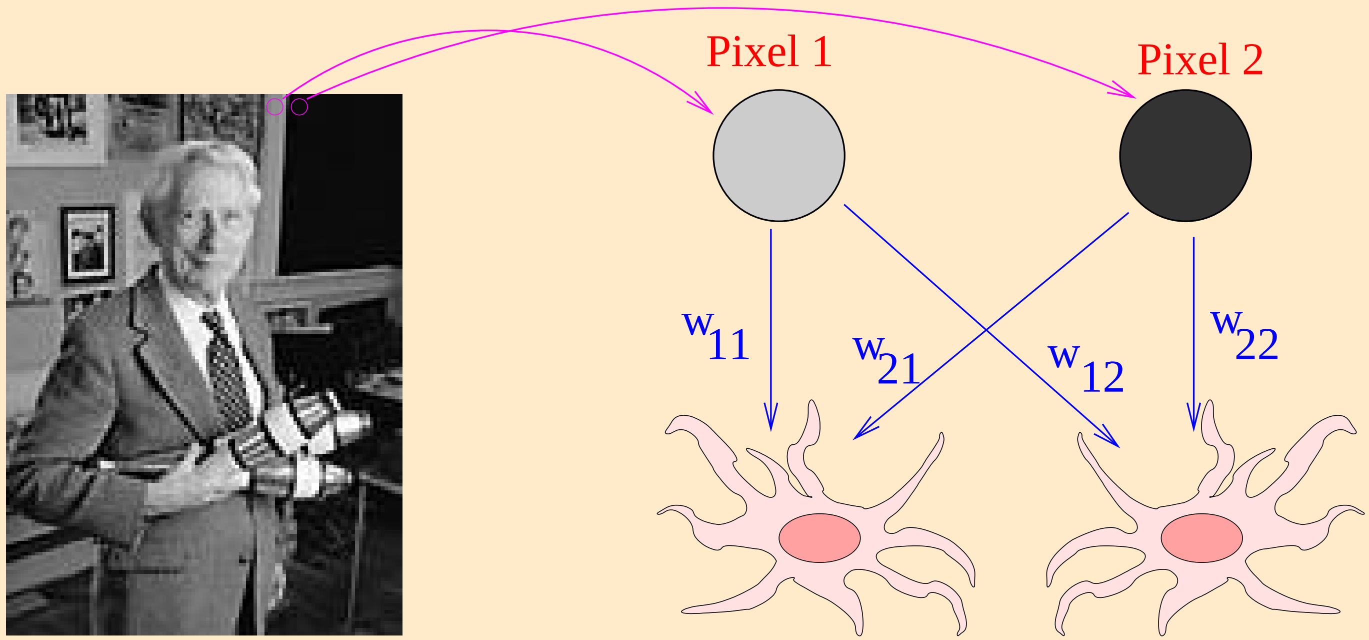

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Edges Are The Independent Components Of Natural Images

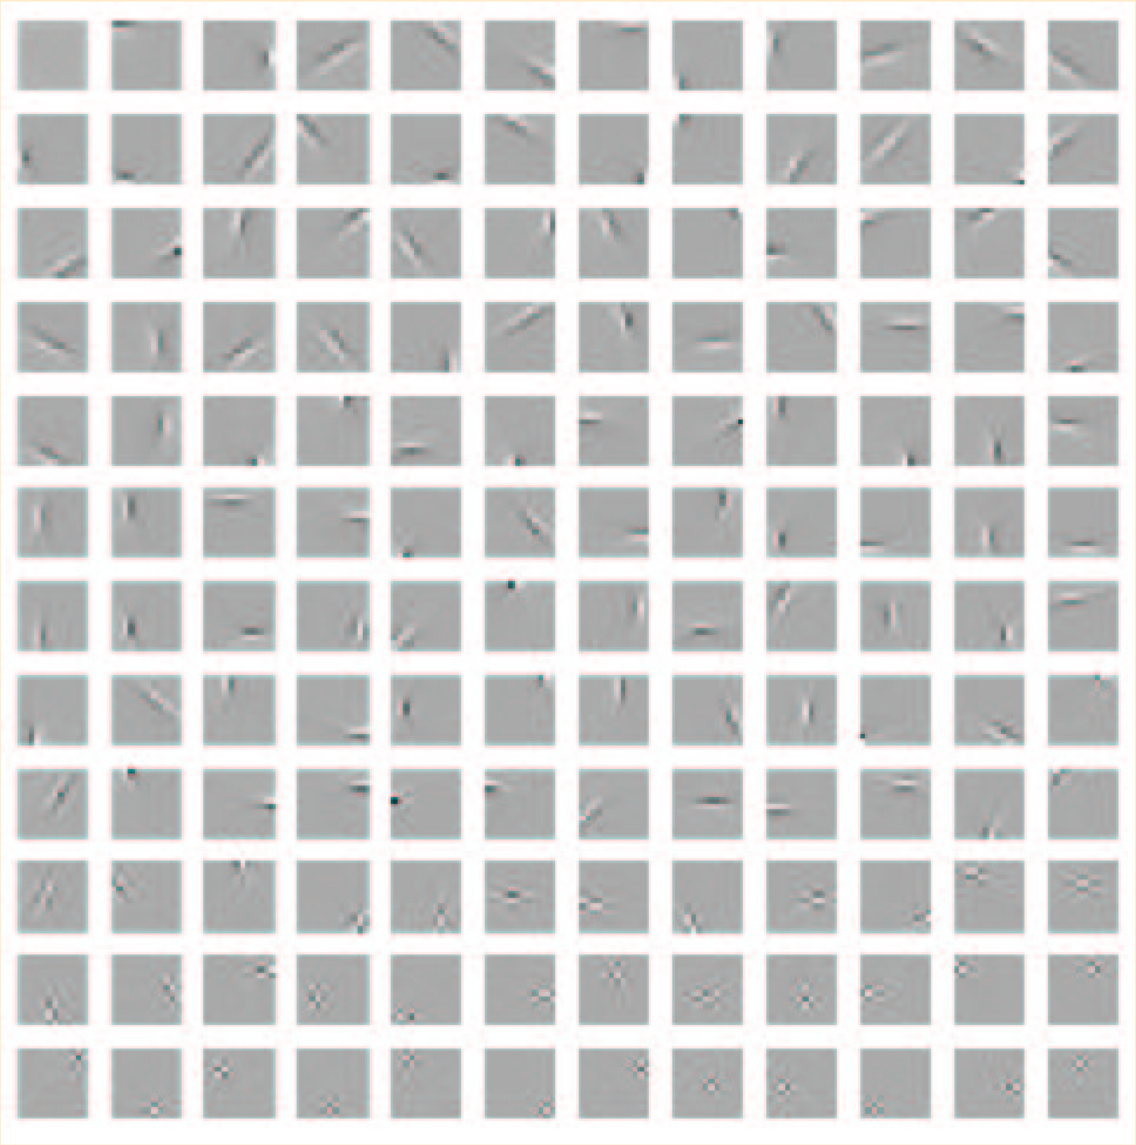

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

ICA Features Match Response Properties of V1 Simple Cells

Gabor Filters

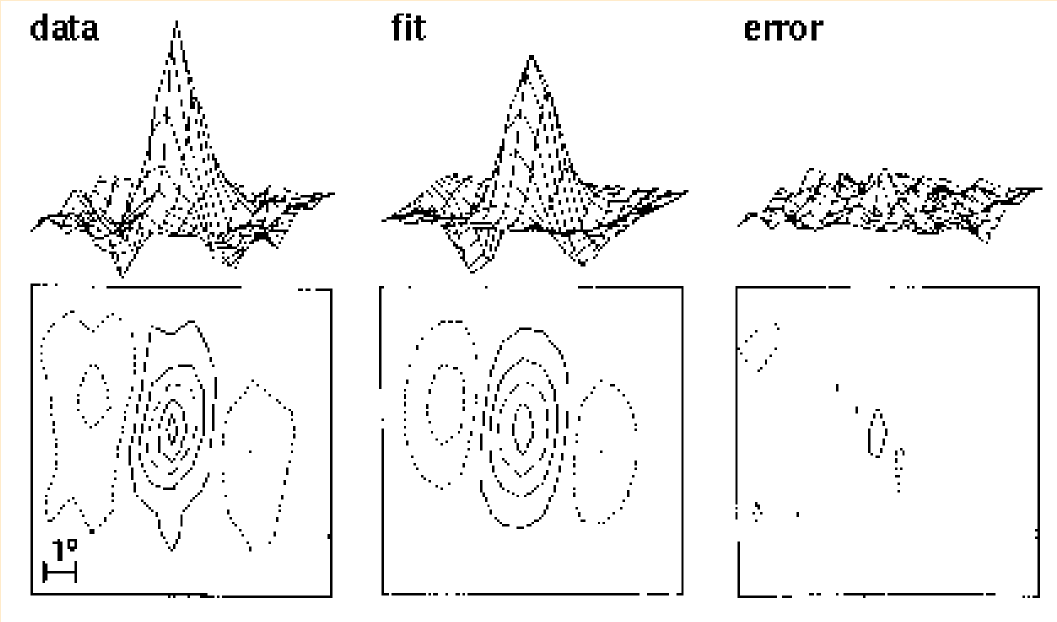

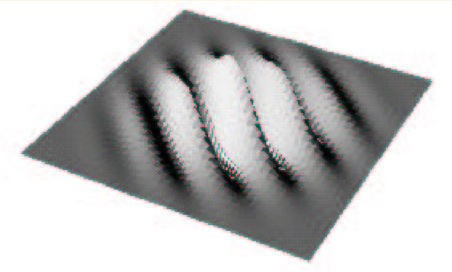

(Bell & Sejnowski, 1996)

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

What Does This Mean?

We've performed a *rational analysis* of the task of efficiently encoding visual information — of producing a visual representation with minimal redundancy. The result of this analysis is a class of feature detectors which are strikingly similar to the feature detectors embodied by neurons in the earliest stages of cortical processing.

This suggests that this part of the brain may have adapted very well (either over evolutionary time or over the course of individual development) to the task of producing such an efficient visual code.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle

Alternative Theories

There are other possible explanations of these response properties . . .

An encoding of visual information based on independent components tends to be *sparse* — few feature detectors are strongly active for any specific scene.

In fact, other mechanisms for producing *sparse codes* also tend to produce Gabor-filter-like response properties. For example, Randy O'Reilly's Leabra framework for constructing computational models of brain function includes a powerful mechanism for fast *lateral inhibition* between cells. This inhibition mechanism tends to produce sparse codes. When a Hebbian learning mechanism is used to adapt synaptic strengths during the presentation of natural scenes, Leabra produces internal representations with Gabor-like properties.

Summary

Using the mathematical tools of machine learning, it is possible to analyze a signal containing a mixture of independent sources and recover those sources as the *independent components* of the mixed sensory signal. This can be done even if nothing is known about the properties of the sources, other than the fact that they are independent. This is called *blind source separation*, and it can be accomplished using *independent components analysis (ICA)*.

ICA can be used as part of a *rational analysis* of the task faced by the initial stages of the visual system. Such an analysis reveals that the goal of *reducing redundancy* in an internal visual representation leads to Gabor-filter feature detectors. The fact that simple cells in V1 act as such feature detectors suggests that this part of the brain may be adapted to produce low redundancy codes.

David C. Noelle, Ph.D.
david.noelle@vanderbilt.edu
people.vanderbilt.edu/˜david.noelle