

Information Processing in Neural Populations

selective tutorial introduction

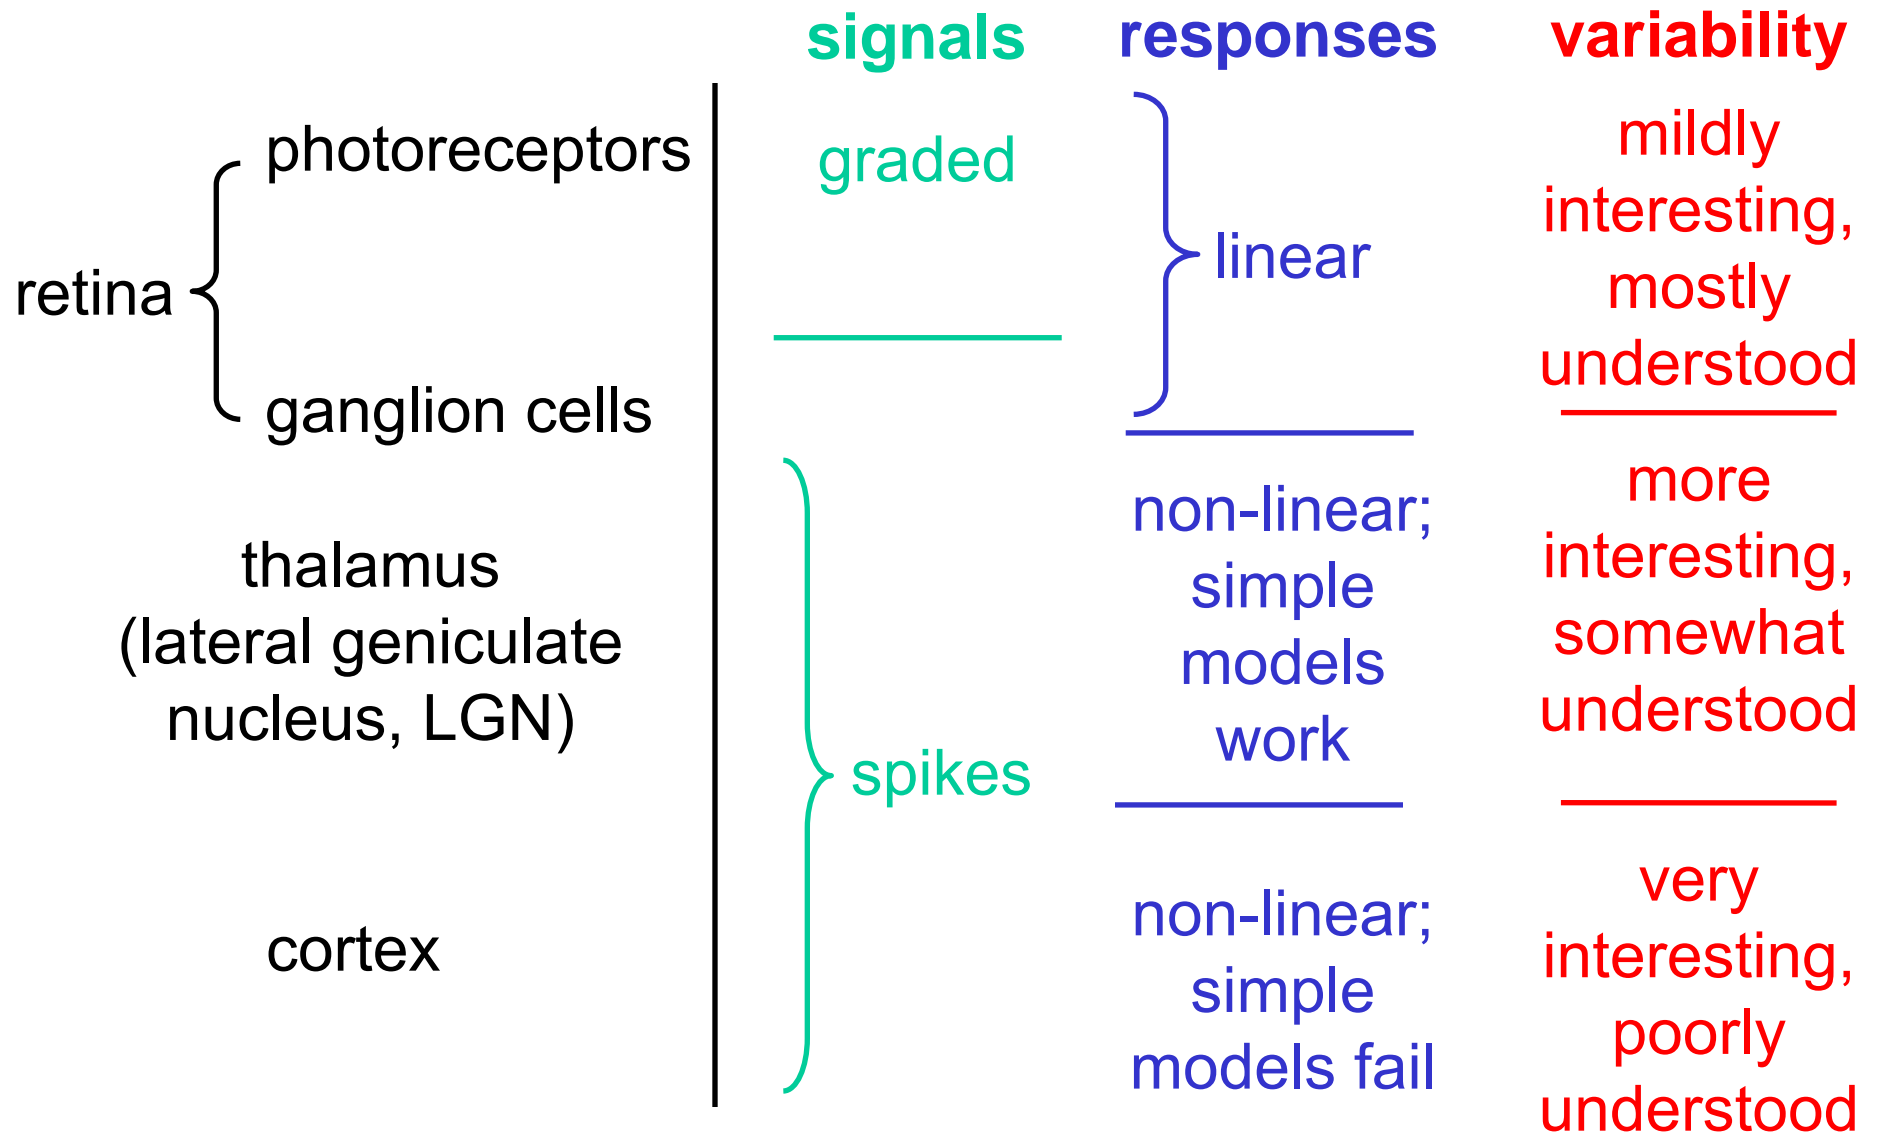
KITP, UCSB
July 2011

Jonathan Victor
Neurology and Neuroscience
Weill Cornell Medical College



Disclaimers

- The visual system is taken as a model (but it's a good one!)
- I will skip many details to try to paint a big picture
- I will emphasize other details because I don't often get a chance to do so

The Big Picture






Overview: retina, thalamus, cortex




Schmolesky: <http://webvision.med.utah.edu/imageswv/>

Retinal Anatomy 101



Kolb, Fernandez, Nelson:
<http://webvision.med.utah.edu/imageswv/>

Visual Cortex



V1
 =
 primary visual cortex
 =
 striate cortex
 =
 area 17

Van Essen (1992); Schmolesky: <http://webvision.med.utah.edu/imageswv/>

Visual Processing is NOT Serial

Anatomic evidence

- Retinal synapses are (nearly) always reciprocal
- Thalamic “relay” neurons receive the retinal output BUT 90% of their synapses are not from the retina
 - feedback from cortex
 - non-visual inputs from brainstem
- Cortical areas have a definable hierarchy based on laminar pattern of inputs and outputs BUT
 - ascending and descending projections are equally prominent
 - ascending and descending projections are always reciprocal

Visual Processing is NOT Serial

Physiologic evidence

- Lateral interactions within V1 rely on thin unmyelinated fibers (slow)
 - $(3 \text{ mm}) / (0.1 \text{ m/sec}) = 30 \text{ ms}$
- Lateral interactions between V1 and V2 rely on myelinated fibers (fast)
 - $(20 \text{ mm}) / (6 \text{ m/sec}) = 3 \text{ ms}$
- Signal spread between cortical areas is faster than within areas
- Perhaps, the processing "unit" spans multiple cortical areas

After Bullier

Acknowledgments

WCMC Neurology and Neuroscience

Keith Purpura
Ferenc Mechler
Anita Schmid
Ifije Ohiorhenuan
Qin Hu
Mary Conte

Former students

Dmitriy Aronov
Danny Reich
Michael Repucci
Bob DeBellis
Bu Qing Mao

Collaborators

Bruce Knight (Rockefeller)
Larry Sirovich (Mt. Sinai, NYC)
Ehud Kaplan (Mt. Sinai, NYC)
Tatyana Sharpee (Salk)

Support

NEI
NINDS
WMC Neurology and
Neuroscience

Framing the Problem

$$R(t) = F[S(x, t)]$$

- Stimulus
 - light intensity or contrast ($L=L_0+L_1 S(x, t)$)
 - current injection
- Response
 - intracellular voltage
 - transmembrane current
 - firing rate

Really want to study

$$R(x, t) = F[S(x, t, \lambda, \textit{behavior})]$$

Framing the Problem, II

- Full probabilistic formulation: study $p(R|S(x,t))$
- Typical simplification
 - For non-spiking neurons: $R(t)=F[S(x,t)]+noise$
 - For spiking neurons: assume an underlying "rate", and then a model for generating spikes
 - Inhomogeneous Poisson (possibly with refractory period)
 - Inhomogeneous renewal
- Ideally, study $p(R|S(x,t,\lambda,behavior))$, and do this for many neurons at the same time

Simplifying the deterministic part: some generic models

$$R(t) = F[S(x, t)]$$

- Linear

$$R(t) = \int K(x, \tau) S(x, t - \tau) dx d\tau$$

- Nonlinear but analytic (Volterra)

$$R(t) = K_0 + \int K_1(x_1, \tau_1) S(x_1, t - \tau_1) dx_1 d\tau_1 + \\ \iint K_2(x_1, \tau_1, x_2, \tau_2) S(x_1, t - \tau_1) S(x_2, t - \tau_2) dx_1 dx_2 d\tau_1 d\tau_2 + \dots$$

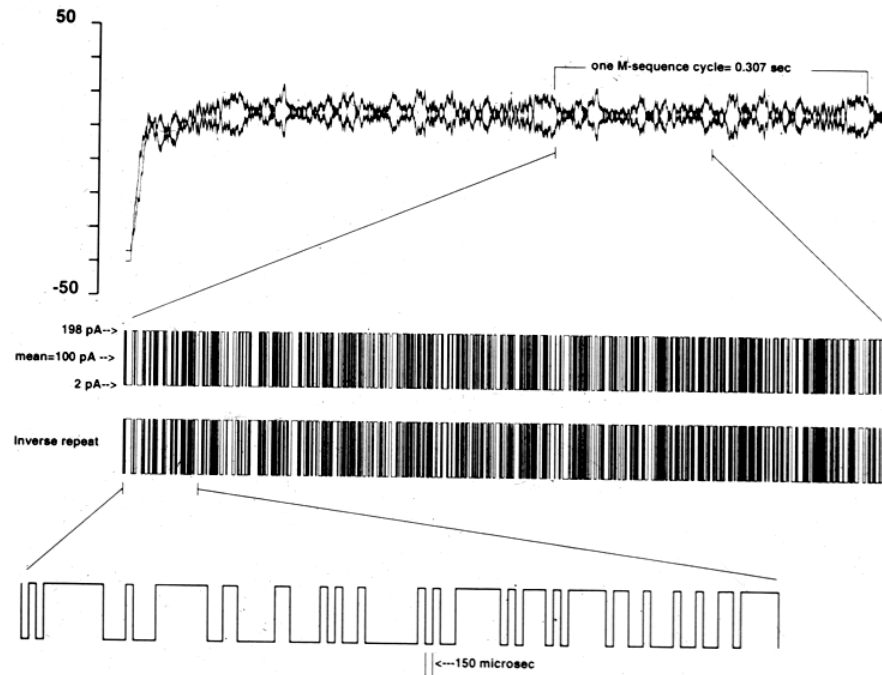
- Orthogonal expansion (Wiener)
- Other forms?

Kernel Measurement: Linear System

Assume: single input $S(t)$, and $R(t) = \int K(\tau)S(t - \tau)d\tau$

$R(t)$: intracellular voltage

$S(t)$: injected current
(e.g., pseudorandom binary
“m-sequence”)



Since the spectrum of $S(t)$ is white, cross-correlation of S and R yields an estimate of impulse response of best-fitting linear system


$$K(\tau) \approx \frac{1}{N} \langle R(t)S(t - \tau) \rangle$$

Photoreceptors

- Graded responses to light
 - depolarized in dark
 - *hyperpolarize* to light (vertebrates)
- Approximately linear for moderate depths of modulation (<30%)
- Linear kernel is separable

$$K(x, \tau) = X(x)T(\tau)$$

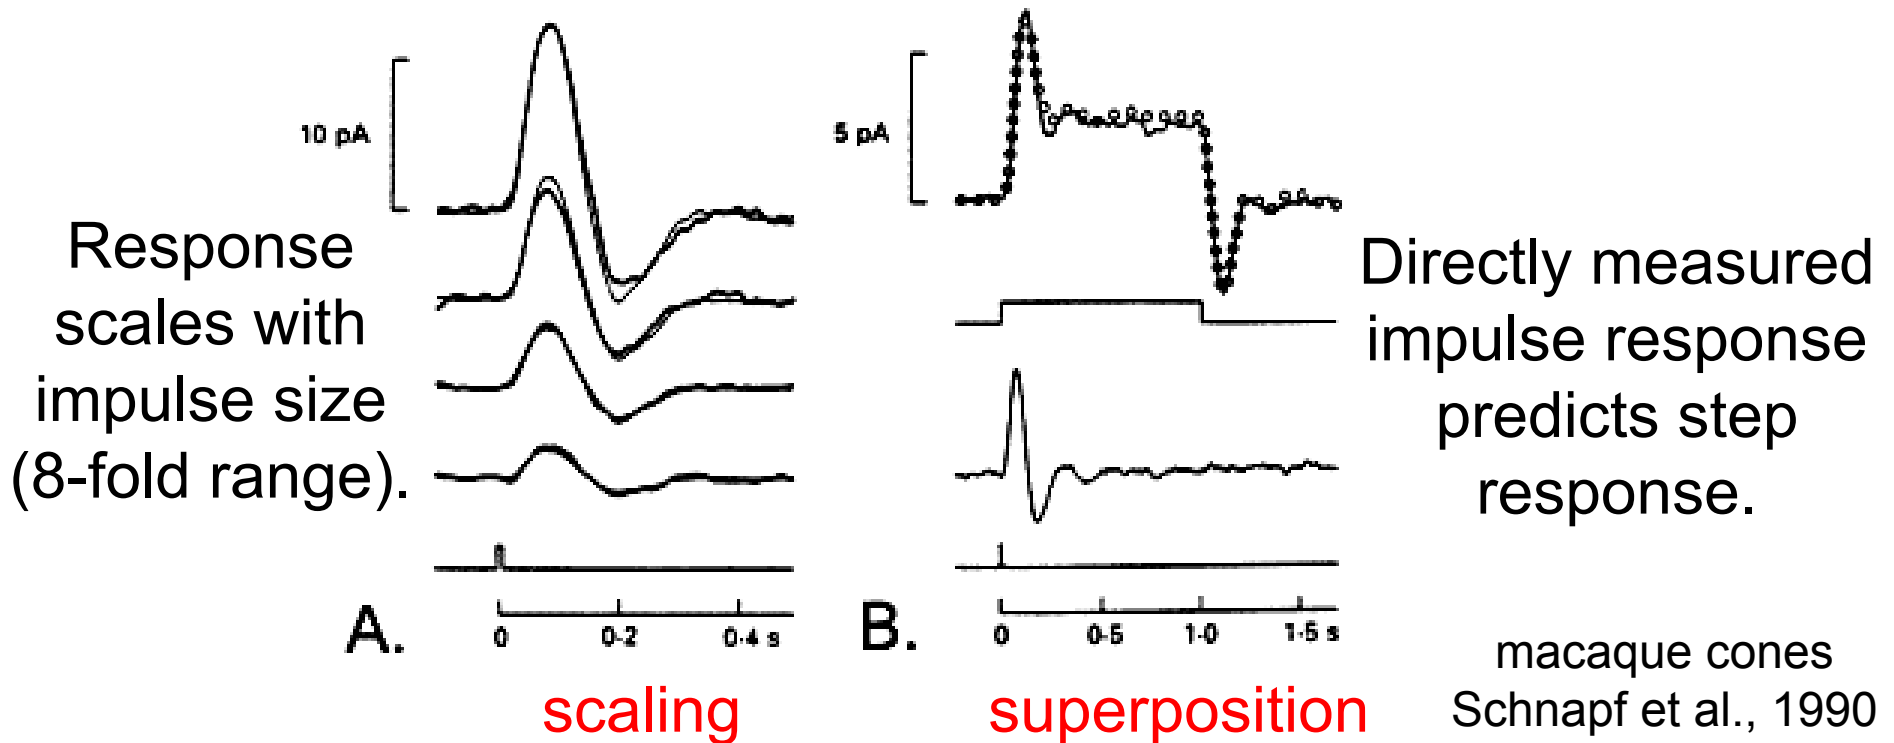
- Spatial profile $X(x)$ determined by optics, waveguide properties, electrical coupling (gap junctions)



Photoreceptor Dynamics I

Assume linearity: $R(t) = \int K(\tau)S(t - \tau)d\tau$


The impulse response $K(t)$ is the response $R(t)$ to $S(t) = \delta(t)$.



Linear behavior for moderate input range

Photoreceptor Dynamics II

Nonlinear behavior (change in dynamics)
over wide input range



70000-fold change in
background intensity,
1000-fold change in
flash intensity


turtle cones

Daly and Normann, 1985

With increasing light level, sensitivity decreases
and response speed increases.


Horizontal Cells

- Basic dynamical features similar to photoreceptors
 - Non-spiking
 - Approximately linear
- A functional syncytium
- Spatiotemporal kernel only approximately separable, due to “cable” (i.e., disk) properties




Bipolar Cells

- Non-spiking
- On-off dichotomy
 - Sublaminar organization
 - Consequences of Dale's Law
- Center-surround organization
 - Non-separable for light input:
$$K(x, \tau) = X_C(x)T_C(\tau) + X_S(x)T_S(\tau)$$
 - Consequence: spatiotemporal coupling
- Approximately linear but only for very small inputs



Bipolar Cell Dynamics


Impulse response varies with mean input current



salamander
Mao et al., 2002

Amacrine Cells


- Many (~40) morphological types
- Recurrent connectivity
- Complex dynamics
 - Some are highly nonlinear
 - On/Off behavior
 - Directional selectivity
 - Some are spiking



Amacrine Cells: On-Off Responses

Intracellular voltage records of responses to abrupt increases and decreases of illumination

Note depolarizing responses at both **ON** and **OFF** transients



Toyoda et al., 1973

Kernel Measurement: Nonlinear System

$$R(t) = K_0 + \int K_1(\tau_1)S(t - \tau_1)d\tau_1 + \iint K_2(\tau_1, \tau_2)S(t - \tau_1)S(t - \tau_2)d\tau_1d\tau_2 + \dots$$

Estimation of K_0, K_1, K_2, \dots is also a linear regression.

White noise (Wiener-Lee-Schetzen): If $S(t)$ is drawn from Gaussian white noise, design matrix is (eventually) simple, and:

$$K_0 \approx \langle R(t) \rangle$$

$$K_1(\tau_1) \approx \frac{1}{N_1} \langle R(t)S(t - \tau_1) \rangle$$


$$K_2(\tau_1, \tau_2) \approx \frac{1}{N_2} \langle R(t)S(t - \tau_1)S(t - \tau_2) \rangle \quad (\tau_1 \neq \tau_2)$$

K 's are not universal; they depend on power.

Lots of parameters.


Not clear what kind of input signal will yield the best design matrix. e.g., m-sequences have $S(t - \tau_1)S(t - \tau_2) = S(t - \lambda(\tau_1, \tau_2))$.

Amacrine Cells: Second-order Kernel




second-order kernel of catfish
amacrine cell intracellular voltage
response to light

first- and second-order
kernels provide a good
approximation to response



Retinal Ganglion Cells

- This is the output of the retina to the (rest of the) brain
- Linear center-surround is a caricature
 - All ganglion cells show changes in gain and dynamics as contrast varies ("contrast gain control")
 - Some ganglion cells are highly nonlinear even for small inputs



Analysis in the Frequency Domain: Linear System

Assume: single input $S(t)$, and $R(t) = \int K(\tau)S(t - \tau)d\tau$

Consider the Fourier transforms $\tilde{K}(\omega) = \int e^{-i\omega\tau} K(\tau)d\tau$, etc.


The convolution theorem yields: $\tilde{R}(\omega) = \tilde{S}(\omega)\tilde{K}(\omega)$

So, $\tilde{K}(\omega)$ may be measured by sinusoids, white noise, sums of discrete sinusoids, ...

Why work in the frequency domain?


Boxes turn into algebra

parallel




$$\tilde{K}(\omega) = \tilde{G}(\omega) + \tilde{H}(\omega)$$

serial



$$\tilde{K}(\omega) = \tilde{G}(\omega)\tilde{H}(\omega)$$

feedback

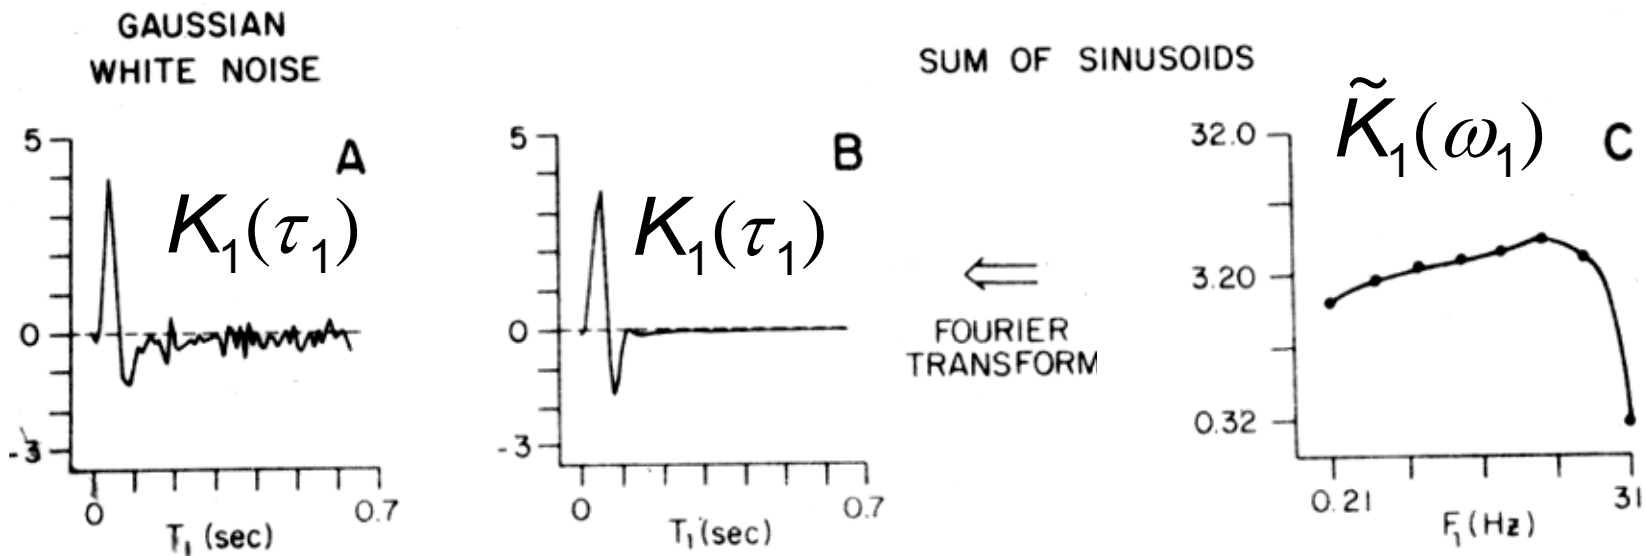


$$\tilde{K}(\omega) = \frac{\tilde{G}(\omega)}{1 - \tilde{G}(\omega)\tilde{H}(\omega)}$$

Comparison: time vs. frequency domain

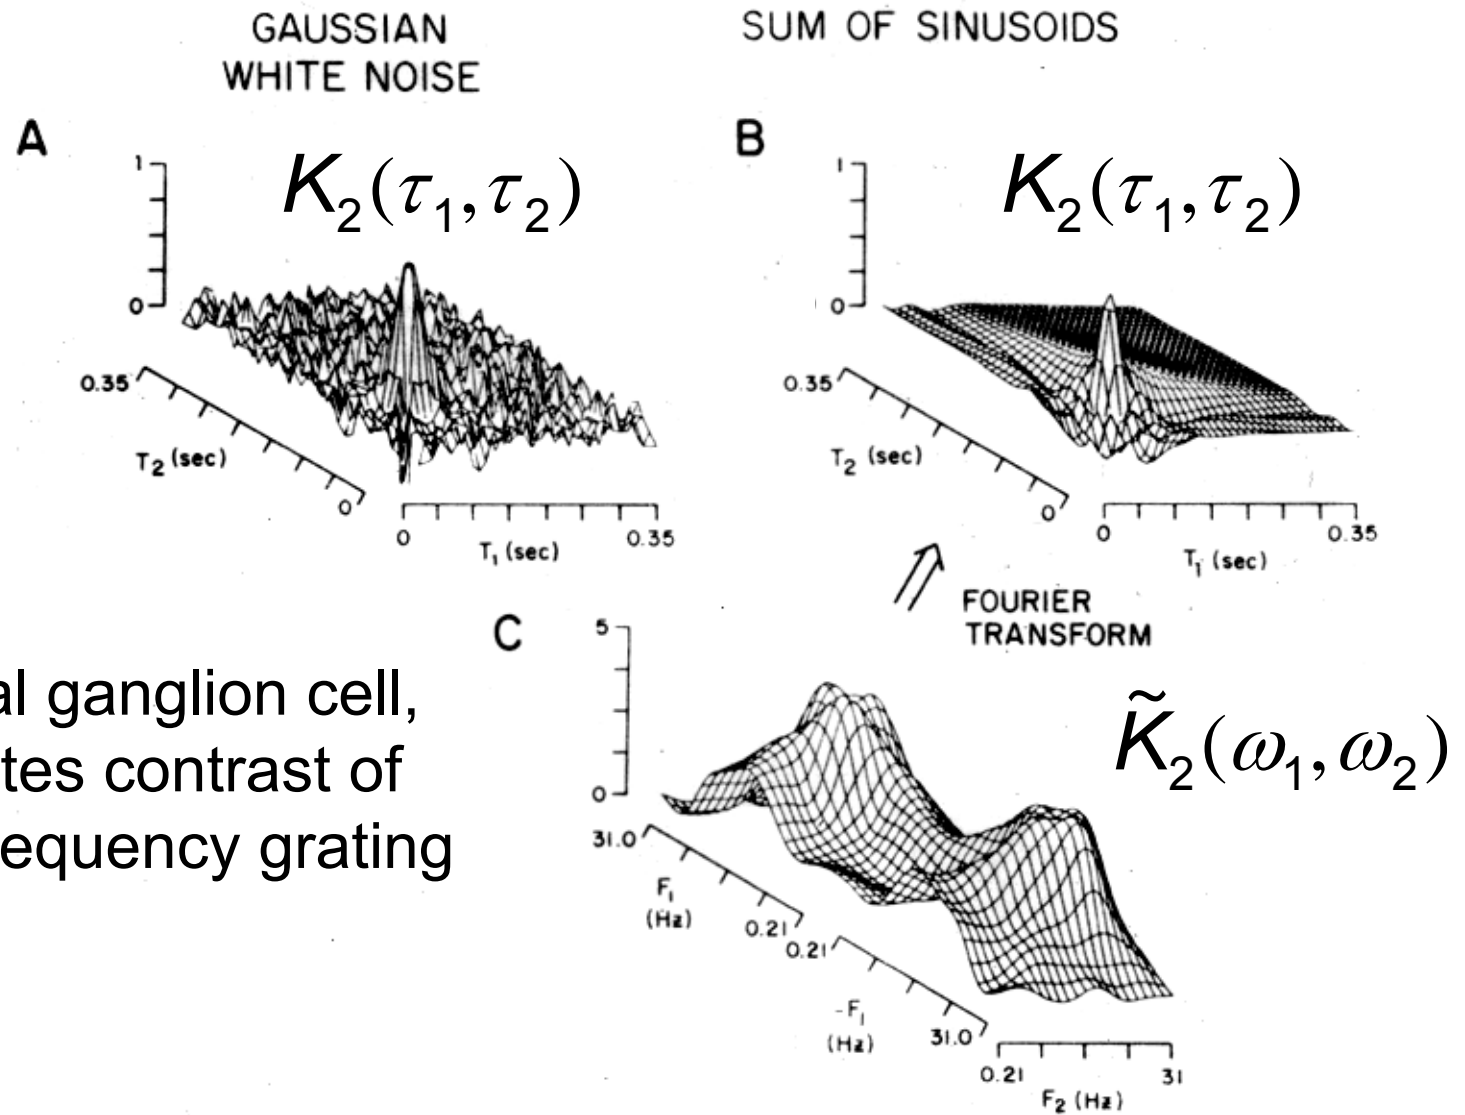
X-type retinal ganglion cell

high spatial frequency grating, $S(t)$ modulates its contrast




Note smoothness of kernel in frequency domain.

Comparison: Time vs. Frequency Domain



Y-type retinal ganglion cell,
 $S(t)$ modulates contrast of
high spatial frequency grating


Frequency domain kernels can have revealing functional forms



$$\tilde{K}_2(\omega_1, \omega_2) \propto \tilde{F}(\omega_1)\tilde{F}(\omega_2)\tilde{G}(\omega_1 + \omega_2)$$

but


$$K_2(\tau_1, \tau_2) \propto \int F(\tau_1 - \tau)F(\tau_2 - \tau)G(\tau)d\tau$$



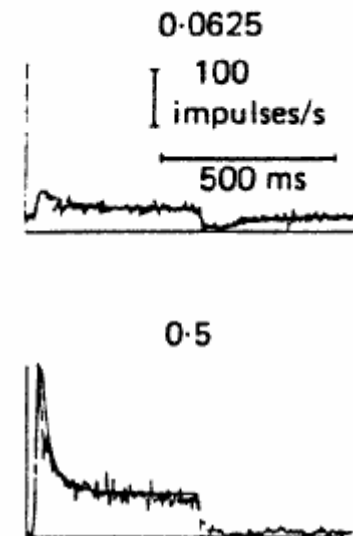
Retinal ganglion cells: Dynamic adaptation to contrast

X cell responses to contrast modulation of a grating

sum-of-
sinusoids
responses



step
responses



Contrast also modulates second-order response,
primarily via the first filter of $F \rightarrow N \rightarrow G$.

Shapley and Victor
1980, 1981


Spikes




Retinal ganglion cells: detailed firing pattern

"linear" ganglion cell responses to sinusoidal modulation

contrast=0.32



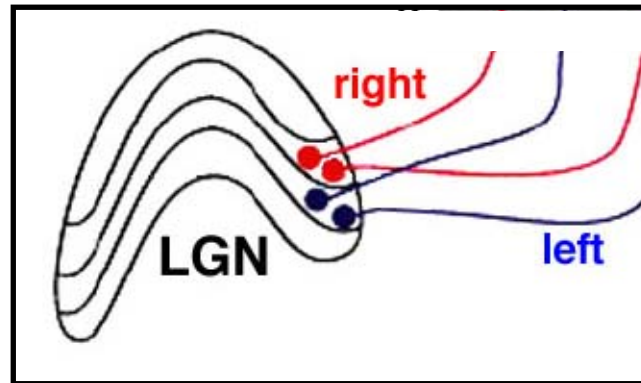
contrast=1.0



At high contrasts, spikes lock to stimulus phase.
This behavior is consistent with a "noisy leaky
integrate-and-fire model"

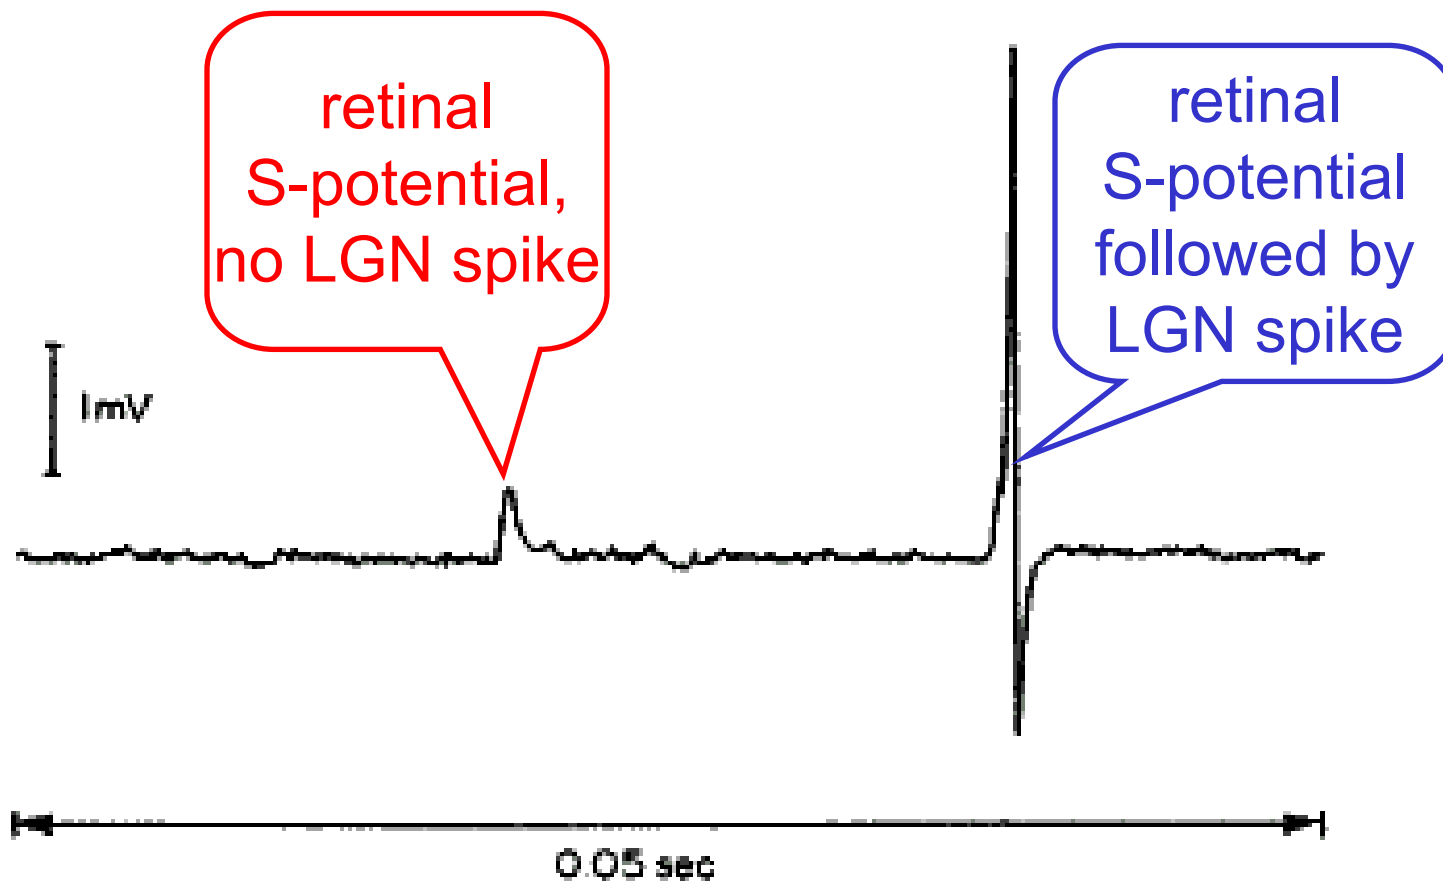
cat
Reich et al. 1997

Relay neurons of the thalamus (LGN)



- Each thalamic neuron has a retinal output neuron as its primary input
- The retinal neuron's spike is necessary but not sufficient for an output event


Spike Editing by Thalamic Relay Neurons



Thalamic editing is not just deleting

- Most "relay" neuron inputs are not retinal
 - From visual cortex
 - From brainstem (? arousal)
- Depending on the recent past history, a relay neuron's output event can be
 - a spike
 - no spike
 - a burst
- This is typical of thalamic relay neurons, not only visual

Firing Modes in LGN Relay Neurons




Bursts rely on a voltage-dependent Ca^{++} channel. Brainstem inputs modulate membrane potential on 100-ms timescale, appropriate to the activation/inactivation dynamics of the channel.

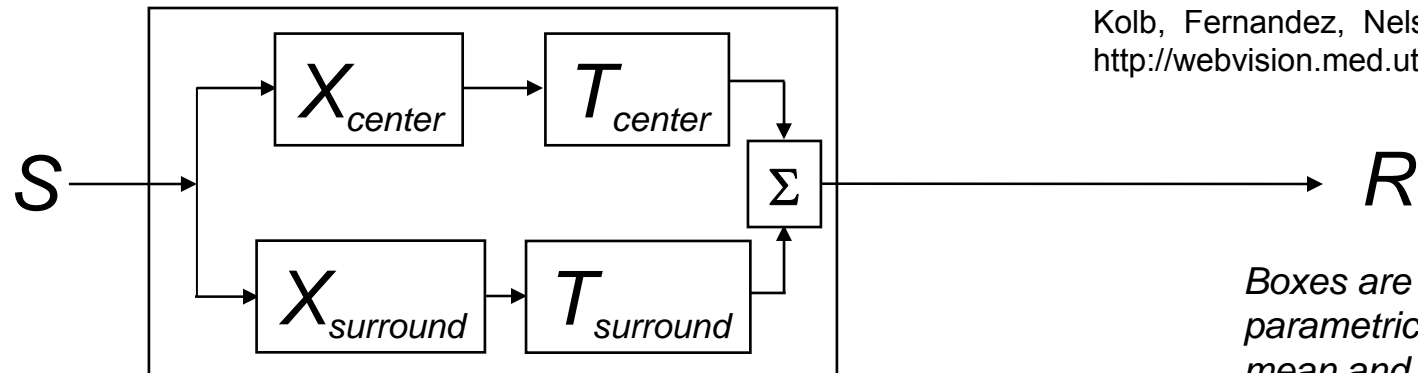
Stretch!



A reasonably satisfying picture




Kolb, Fernandez, Nelson:
<http://webvision.med.utah.edu/imageswv/>



*Boxes are gently
parametric in input
mean and variance*

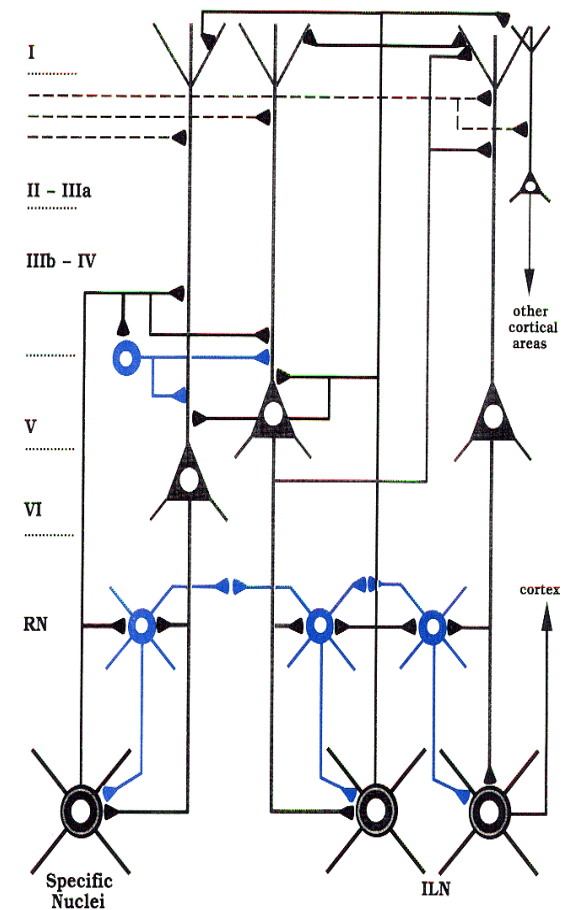
OR,



- Processing steps correspond to anatomy
- Goal of processing is clear: redundancy reduction for efficient coding to get through the bottleneck of the optic nerve

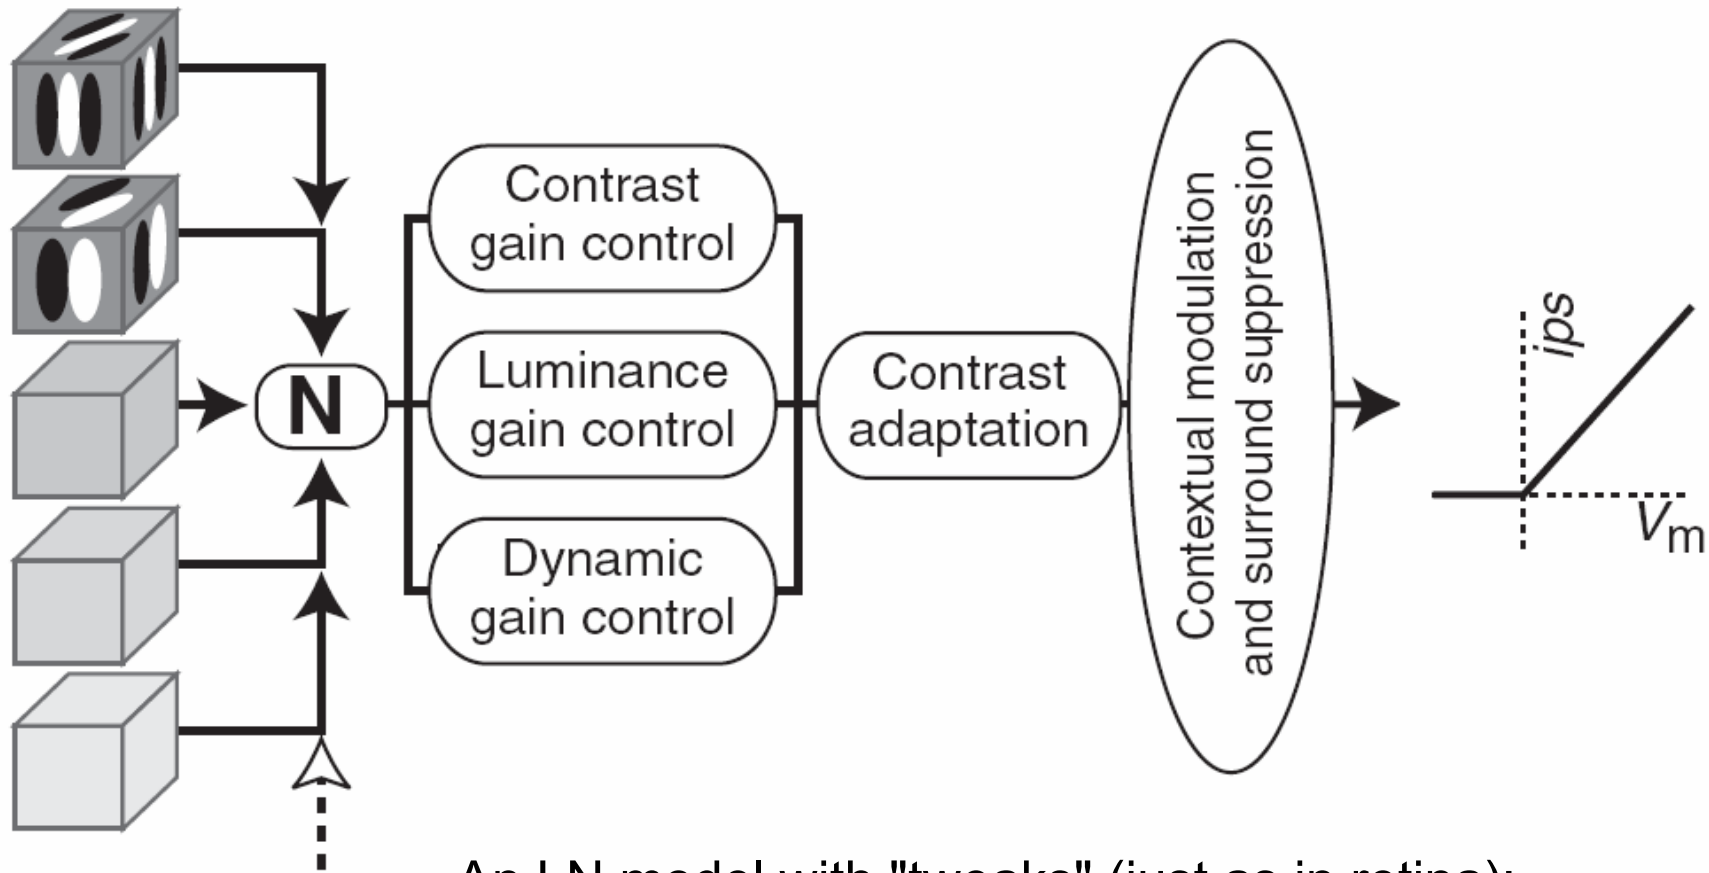
Onward to primary visual cortex

- Major differences between retinal and cortical anatomy and physiology
 - Multilaminated structure, even more cell types
 - All neurons spike
 - Even "input" layer synapses are mostly intrinsic
 - Anatomical and physiological substrate for top-down influences
- Unlikely that cortical processing has the same goals
 - No "bottleneck"
 - Much redundancy has already been removed
- But current computational models have the same computational structure -- LN cascades, with tweaks



Adapted from Llinas et al. 1994, by Purpura and Schiff 1997

The "New Standard Model" for V1 neurons




An LN model with "tweaks" (just as in retina);
selectivity is governed by the initial linear stage

adapted from Rust and Movshon, 2005


Is something qualitative missing?

- Models are built from neural responses to simple stimuli but have only fair predictive accuracy for natural scenes
 - Because of “top-down” factors: attention?
 - Because of low-level factors: high-order correlations
 - They distinguish local features (lines, edges) from noise
 - They distinguish natural scenes from traditional analytical stimuli
 - Unclear whether explaining V1’s computations requires a departure from the “new standard” architecture
 - We don't have a concise predictive model for the gain controls
 - We can't collect enough data to characterize them
- Strategy: Use *designed* stimuli that *neutralize* the gain controls

Two-Dimensional Hermite Functions



Neutralizing the gain controls




- All elements have the same mean-squared contrast.
- Within each rank, the two sets have the same spatial extent, frequency spectrum, and linear span.


Strategy: build LN-type models from neural responses to each set. The inferred filters should match. If not, then we cannot blame some special property of natural scenes, or the effect of gain controls.

Responses of a typical V1 neuron


Cartesian



polar




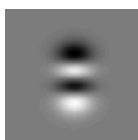
layer 3, non-directional, broad orientation tuning

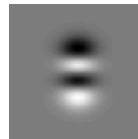


Testing the model

Fitting an LN-type model to the responses



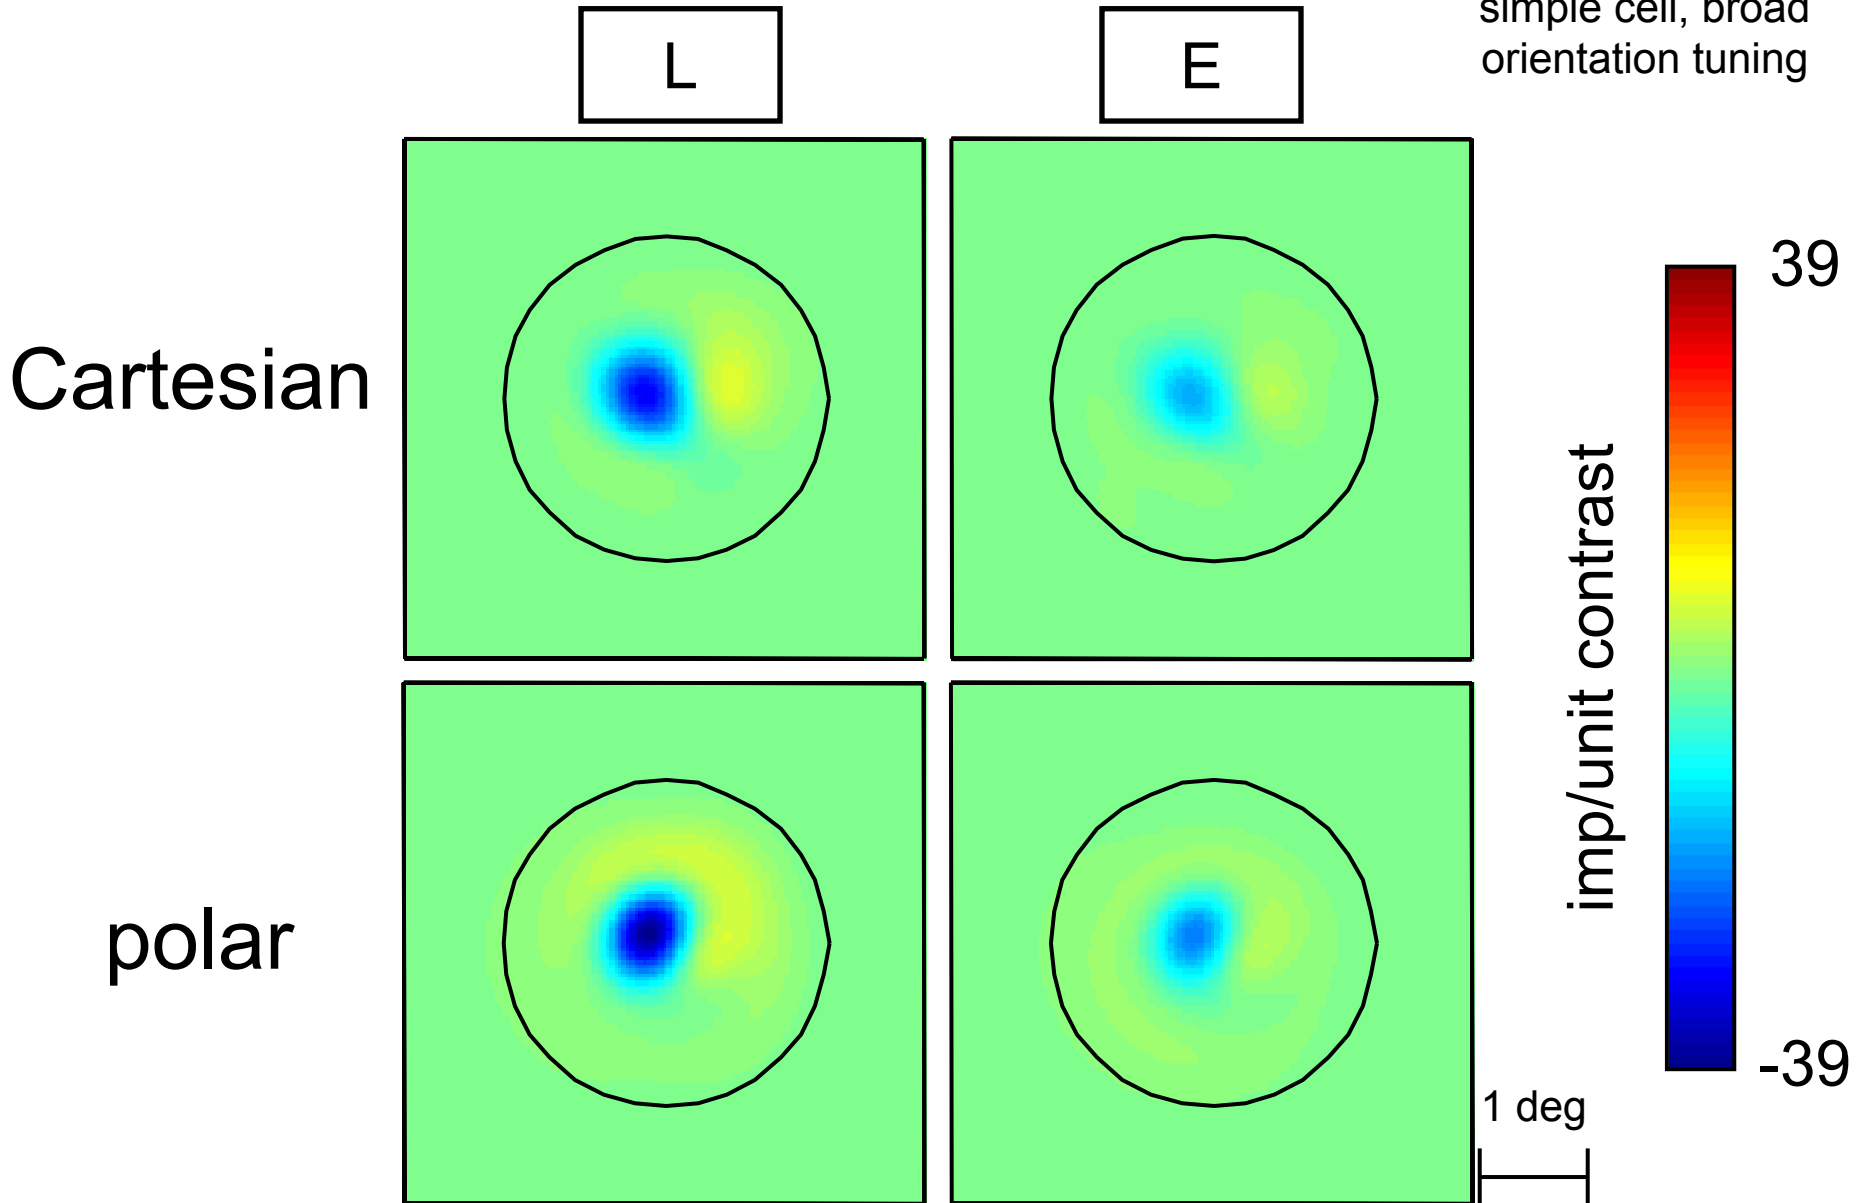
sensitivity of L to  = $\frac{\text{resp}[\text{Gabor}] - \text{resp}[\text{Gabor}]}{2}$

sensitivity of E to  = $\frac{\text{resp}[\text{Gabor}] + \text{resp}[\text{Gabor}]}{2}$

*"New Standard Model prediction:
Cartesian and polar stimuli yield the same filters*

First cell

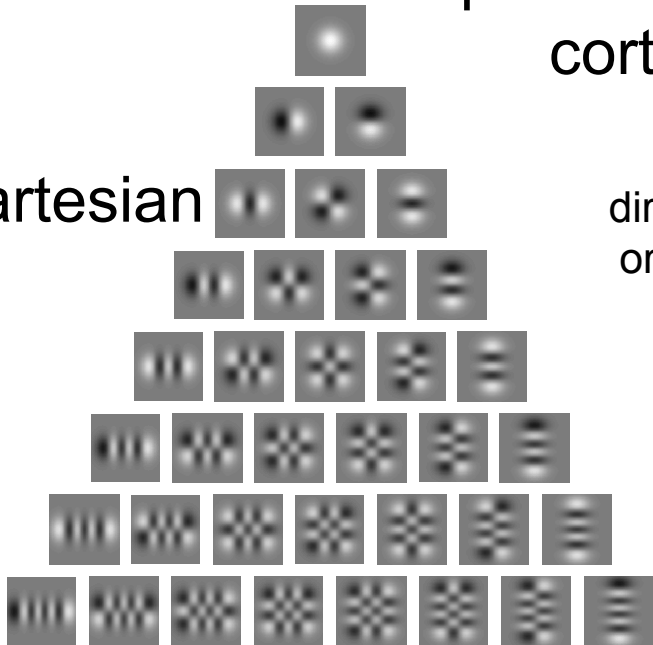
layer 3, non-directional
simple cell, broad
orientation tuning




Prediction holds: filters determined from C and P are similar

Responses of another typical cortical neuron


Cartesian



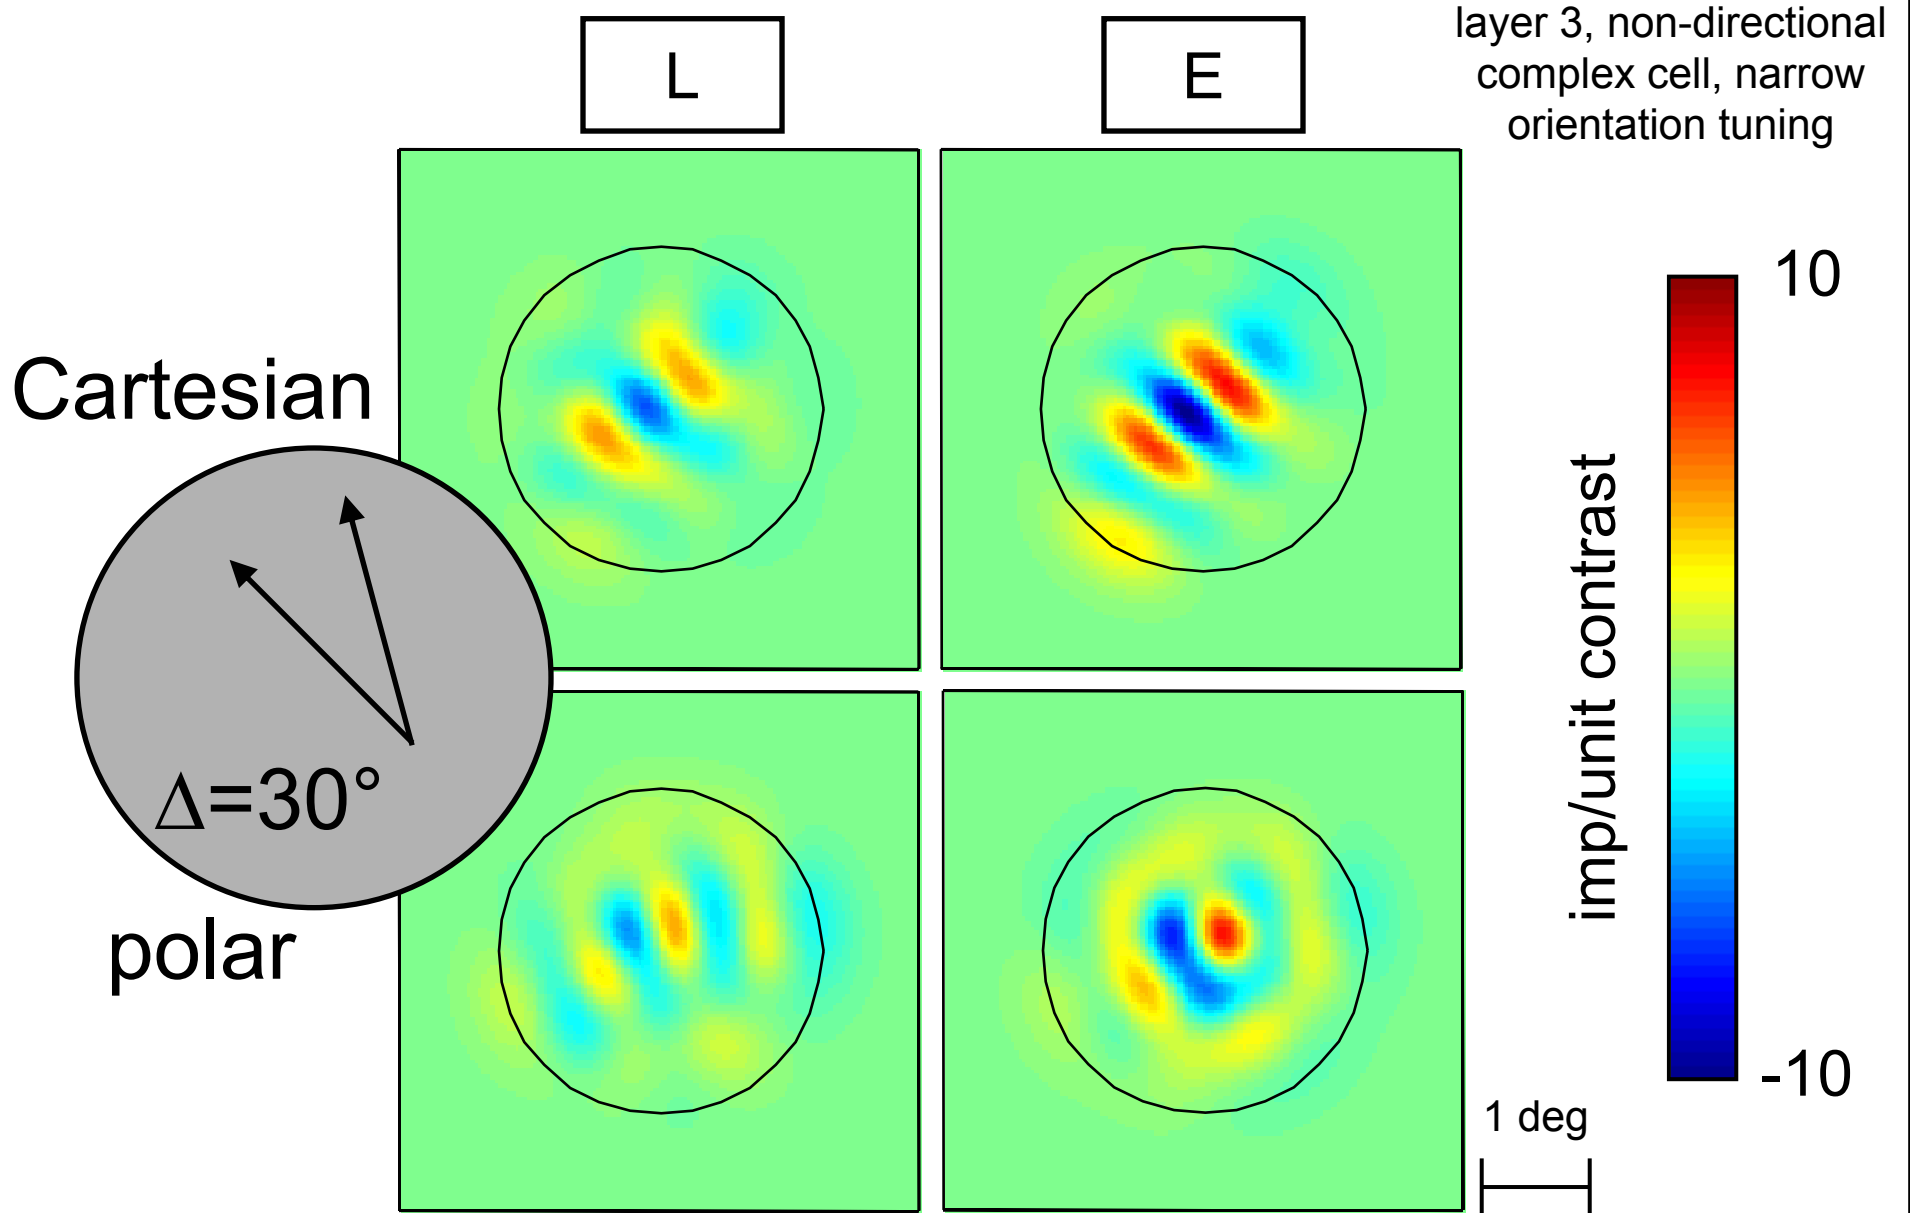
polar



layer 3, non-directional, narrow orientation tuning




Second cell



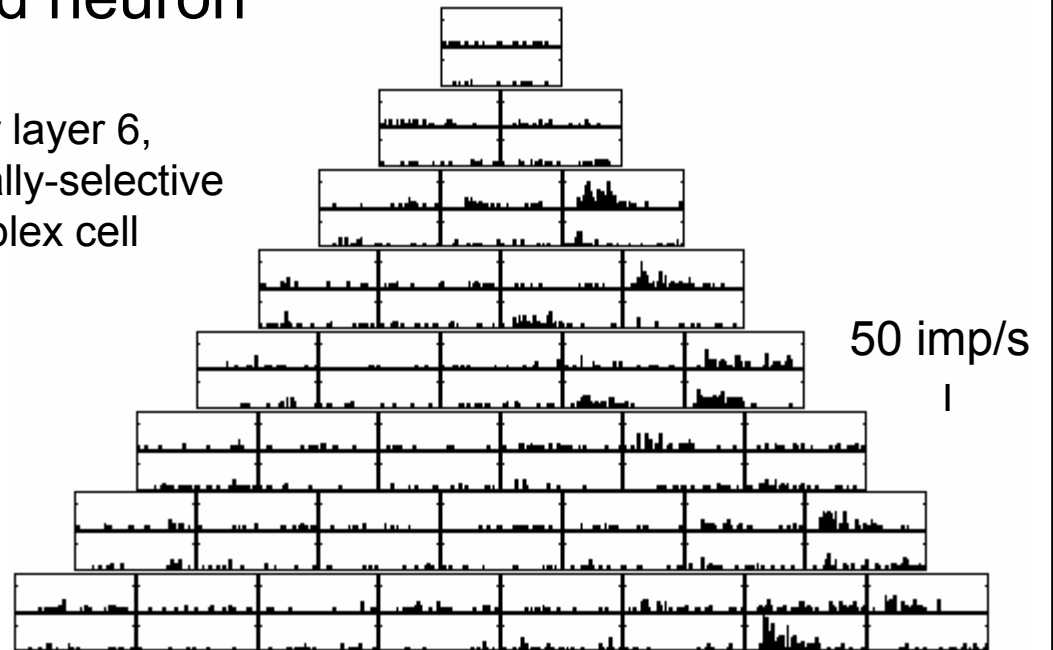
Prediction fails: filters determined from C and P differ in shape

A third neuron



Cartesian



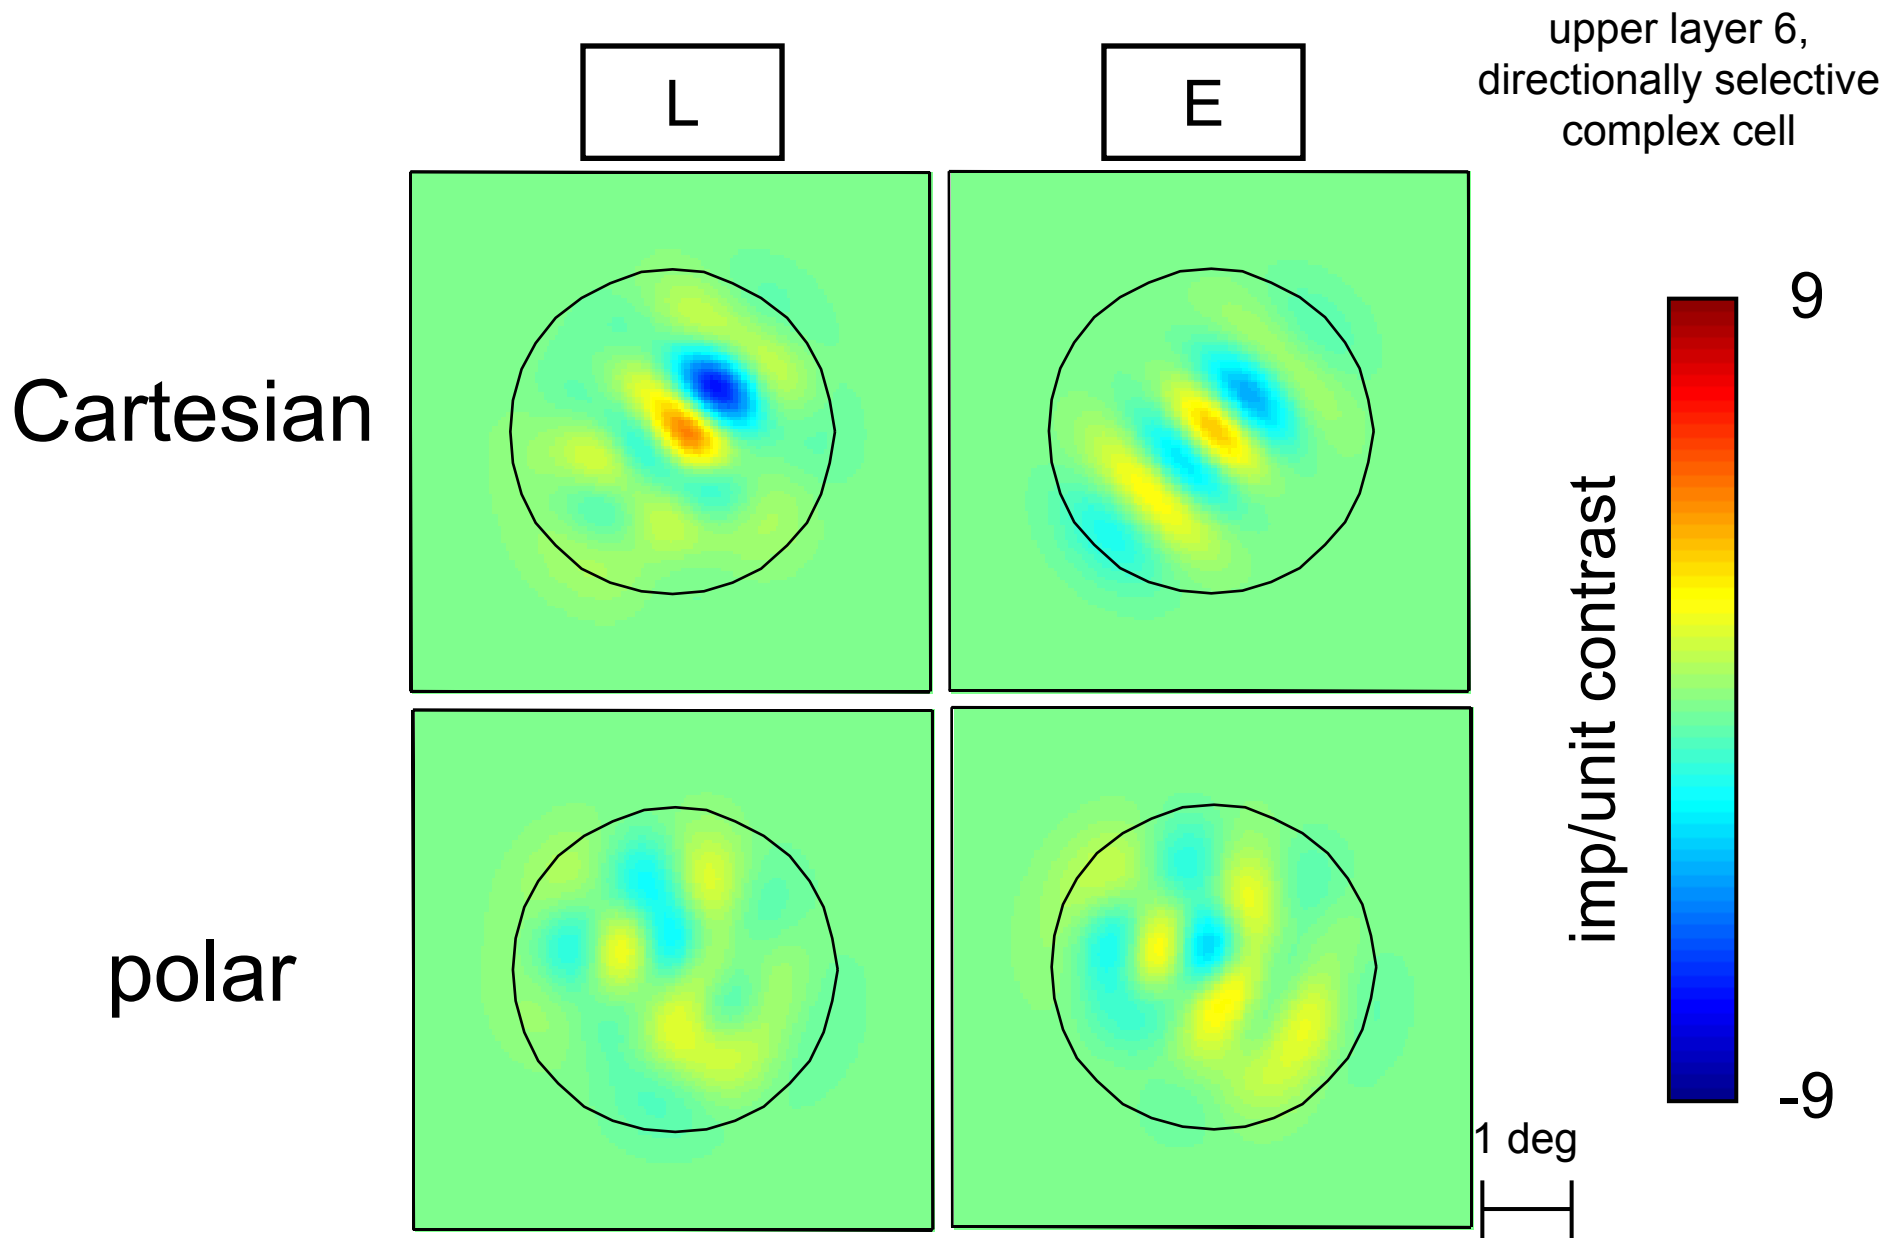
upper layer 6,
directionally-selective
complex cell



polar




Third cell



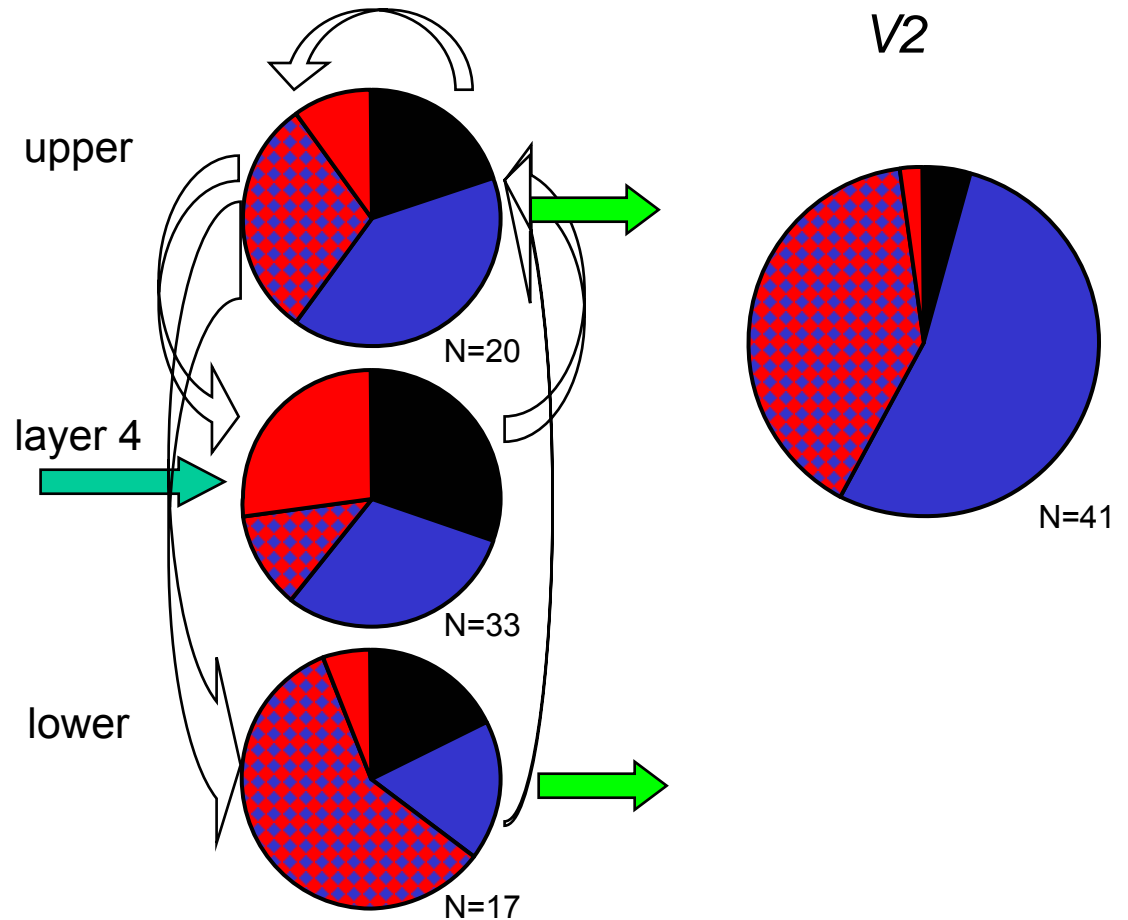
Prediction fails: filters determined from C and P differ in size

Population summary, and relationship to laminar organization



- 41/70 neurons: prediction fails, filters differ in shape
- 32/70 neurons: prediction fails, filters differ in size
- 17/70 neurons: prediction holds (neither difference)

Laminar analysis, V1




Prediction failure is typical in all layers, even in layer 4 (input)


Contour orientation is key

The apparent change in the filters is a signature of departure from LN behavior. What stimulus characteristics are driving it?


Cartesian, **aligned** to preferred orientation



polar



Cartesian, **oblique** to preferred orientation



$r=0.76$

$r=0.76$

$r=0.49$


Changing contour orientation has a larger effect than removing them.

Summary so far


- Most V1 neurons show qualitative departures from the predictions of cascade models
 - for simple non-natural stimuli
 - that neutralize the gain controls
 - even in the input layers
- The presence of oriented contours drives this departure
- Since orientation selectivity first appears in V1, this suggests recurrent nonlinear processing
- But is this finding specific to these peculiar matched basis sets? Can we more directly test the idea that high-order correlations matter?

Variation on a familiar theme

mapping with random binary stimuli



Cross-correlate the spike trains with the stimulus, pixel by pixel, to generate a receptive field sensitivity profile.




This works if the stimulus pixels are uncorrelated *in pairs*.


We can create stimuli in which pixels are uncorrelated in pairs, but correlated at higher orders -- and use them to study whether V1 neurons care about these correlations.

Implementing the idea

- No pairwise correlations
- Strong fourth-order correlation: every "glider" has an even number of white checks

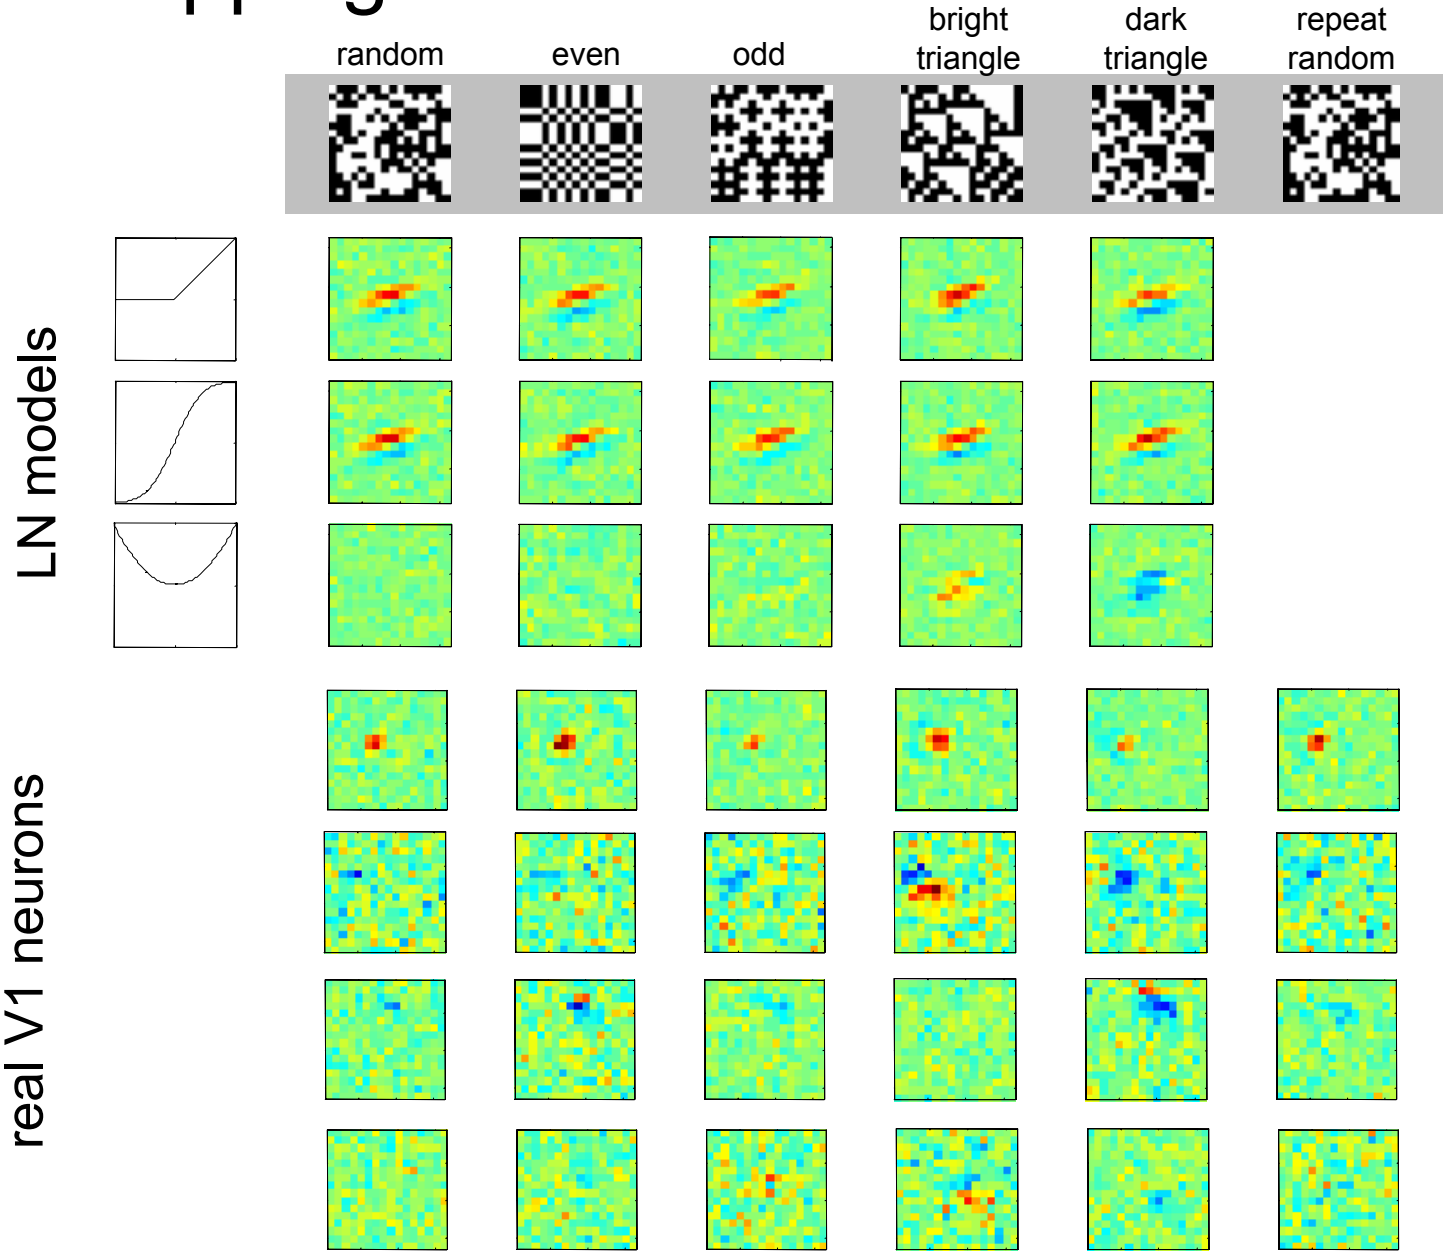


other glider shapes and parities



Each texture is a probe for neural sensitivity to a specific kind of high-order correlation.

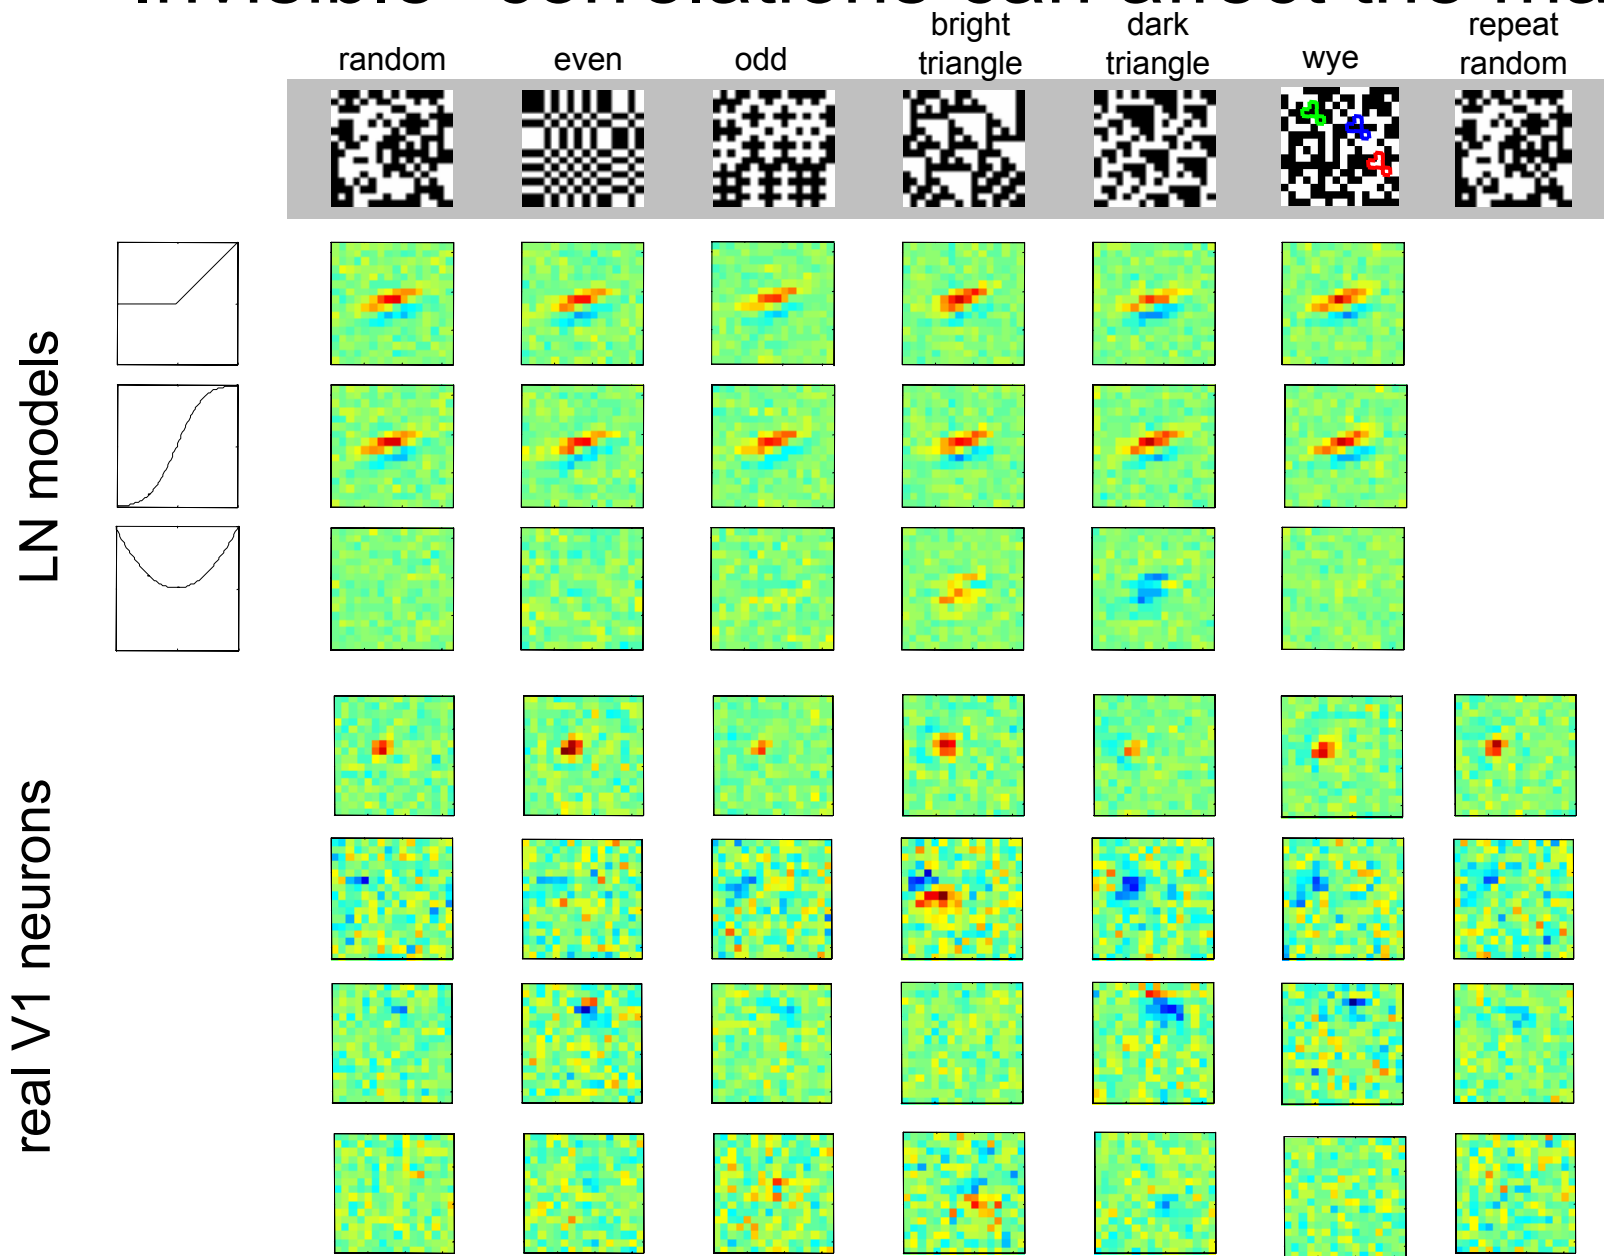
Mapping V1 neurons with correlated stimuli



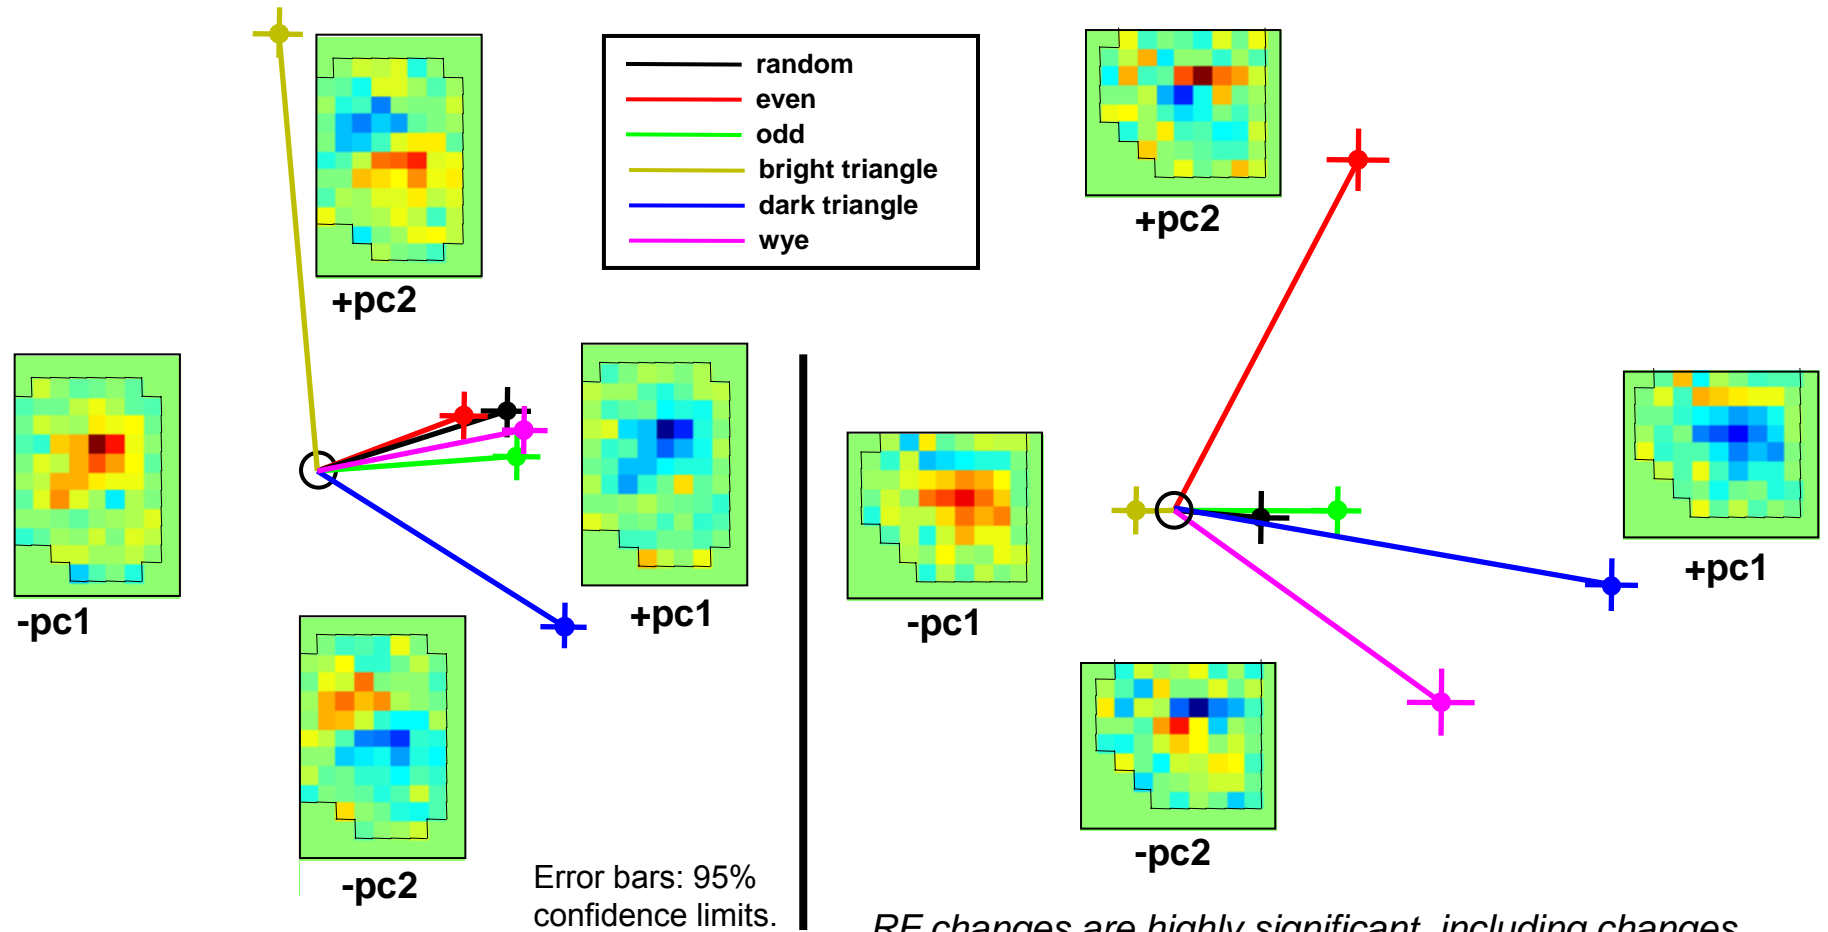

But maybe we could build dedicated, parallel combinations of LN models to recover this behavior.

Large changes in sensitivity profiles; some neurons only "mappable" with correlated stimuli

"Invisible" correlations can affect the map



Principal components analysis of RF maps




Error bars: 95% confidence limits.

*RF changes are highly significant, including changes induced by "invisible" correlations (the **wye** stimulus).*


Recurrence makes it sensible

- With feedforward architecture
 - We would have to build a parallel set of LN modules, with dedicated circuitry for each kind of correlation
 - We would have to include circuitry for correlations that aren't perceptible (and don't seem to correspond to edges and regions)
- But what if the recurrence dominates?



- Each path through the network traverses a different combination of nonlinearities
- This generates lots of useful combinations (e.g., local edge detection followed by interactions along extended contours)
- But it also generates some crosstalk -- accounting for sensitivity to "invisible" correlations
- And it meshes well with anatomy

What is a “natural scene”?



“I know it when I see it”.

Potter Stewart, *Jacobellis v. Ohio* (1964)

What is the goal of the computations
in primary visual cortex?

Lab meeting, April 2011