Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 19: Motion Illusions & Bayesian models

Initialize

■ **Spell check off**

In[360]:=
```
Off[General::spell1];
<< VectorFieldPlots`
```

In[362]:=
```
SetOptions[ArrayPlot, ColorFunction → "GrayTones", DataReversed → True,
  Frame → False, AspectRatio → Automatic, Mesh → False,
  PixelConstrained → True, ImageSize → Small];
SetOptions[ListPlot, ImageSize → Small];
SetOptions[Plot, ImageSize → Small];
SetOptions[DensityPlot, ImageSize → Small, ColorFunction → GrayLevel];
nbinfo = NotebookInformation[EvaluationNotebook[]];
dir =
 ("FileName" /. nbinfo /. FrontEnd`FileName[d_List, nam_, ___] :→
    ToFileName[d]);
```

Outline

Last time

• Early motion measurement--types of models

•Functional goals of motion measurements

• Optic flow

Cost function (or energy) descent model

A posteriori and a priori constraints

Gradient descent algorithms

Computer vs. human vision and optic flow

-- area vs. contour

Today

■ Motion phenomena & illusions

Neither the area-based nor the contour-based algorithms we've seen can account for the range of human motion phenomena or psychophysical data that we now have.

Look at human motion perception

■ Local measurements & neural systems

Representing motion, Orientation in space-time

Fourier representation and sampling

Optic flow, the gradient constraint, aperture problem

Neural systems solutions to the problem of motion measurement.

Space-time oriented receptive fields

■ Global integration

Sketch a Bayesian formulation--the integrating uncertain local measurements with the right priors can be used to model a variety of human motion results.

Human motion perception

Demo: area-based vs. contour-based models

Last time we asked: Are the representation, constraints, and algorithm a good model of human motion perception?

The answer seems to be "no". The representation of the input is probably wrong. Human observers often give more weight to contour movement than to intensity flow. Human perception of the sequence illustrated below differs from "area-based" models of optic flow such as the above Horn and Schunck algorithm. The two curves below would give a maximum correlation at zero--hence zero predicted velocity. Human observers see the contour move from left to right--because the contours are stronger features than the gray-levels. However we will see in Adelson's missing fundamental illusion that the story is not as simple as a mere "tracking of edges" --and we will return to spatial frequency channels to account for the human visual system's motion measurements. At the end of this lecture, we'll review a Bayesian model that integrates local motion information according to reliability, providing a theory that may explain a diverse set of motion illusions.

In[367]:=
```
size = 120;
Clear[y];
low = 0.2; hi = .75;
y[x_] := hi /; x<1
y[x_] := .5 Exp[-(x-1)^2]+.1 /; x >= 1
```

In[372]:=
```
ylist = Table[y[i],{i,0,3,3/255.}];
width = Dimensions[ylist][[1]];
```

In[374]:=
```
picture1 = Table[ylist,{i,1,width/2}];
picture2 = .9 - Transpose[Reverse[Transpose[picture1]]];
```

In[376]:=
```
g1 = ListPlot[picture1[[size/2]],PlotStyle->{Hue[.3]}];
g2 = ListPlot[picture2[[size/2]],PlotStyle->{Hue[.6]}];
Show[g1,g2]
```

Out[378]=

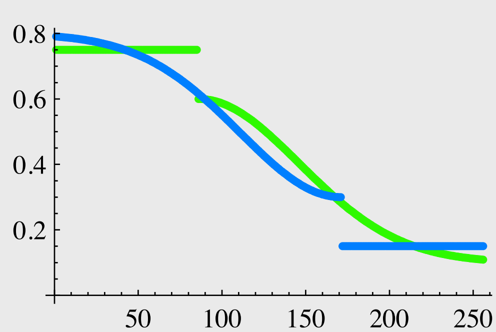

In[379]:=
```
ga1=ArrayPlot[picture1,Frame->False,Mesh->False,
                PlotRange->{0,1}, AspectRatio->Automatic]
ga2=ArrayPlot[picture2,Frame->False,Mesh->False,
                PlotRange->{0,1}, AspectRatio->Automatic]
```

Out[379]=

Out[380]=

There is a clear sense of motion of the edge, even though the signal inferred from an intensity, region-based integration of optic flow would produce little or no optic flow in that direction.

Aperture effects

In[381]:=
```
niter = 8; width = 32;
theta1 = Pi/4.; contrast1 = 0.5;
freq1 = 4.; period1 = 1/freq1;
stepx1 = Cos[theta1]*(period1/niter); stepy1 = Sin[theta1]*(period1/niter);

grating[x_,y_,freq_,theta_] := Cos[(2. Pi freq)*(Cos[theta]*x + Sin[theta]*
```

■ Circular aperture

In[386]:=

```
Animate[DensityPlot[If[(x-0.5)^2+(y-0.5)^2<0.3^2,grating[x+i*stepx1,y+i*ste
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width],{i,1,n
]
```

Out[386]=

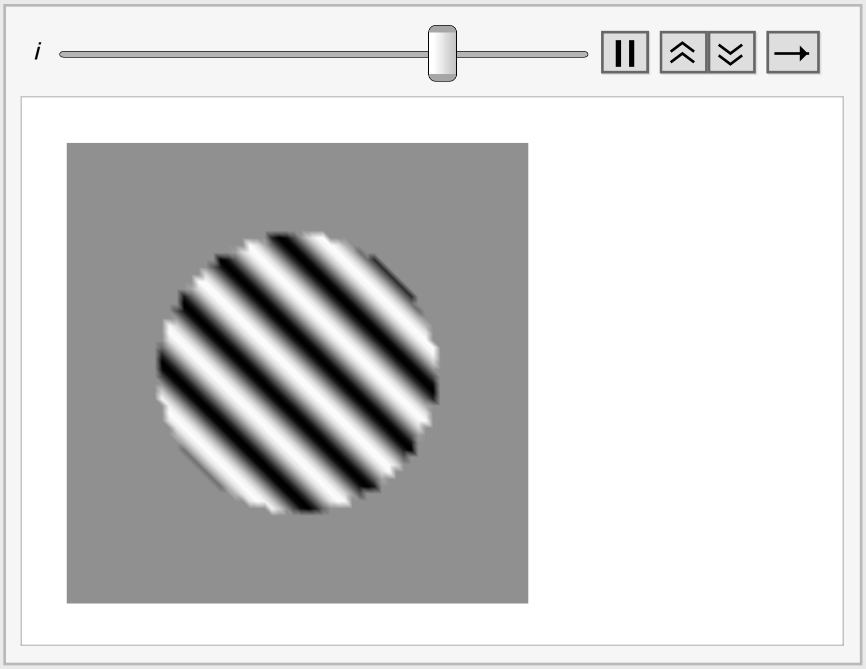

■ Square aperture

In[387]:=
```
Animate[DensityPlot[grating[x+i*stepx1,y+i*stepy1,freq1,theta1],{x,0,1},{y,
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width],{i,1,n
```

Out[387]=

What do you see at the vertical boundaries? The horizontal boundaries?

■ Rectangular horizontal aperture

In[388]:=
```
Animate[DensityPlot[grating[x+i*stepx1,y+i*stepy1,freq1,theta1],{x,0,1},{y,
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width, Aspect
```

Out[388]=

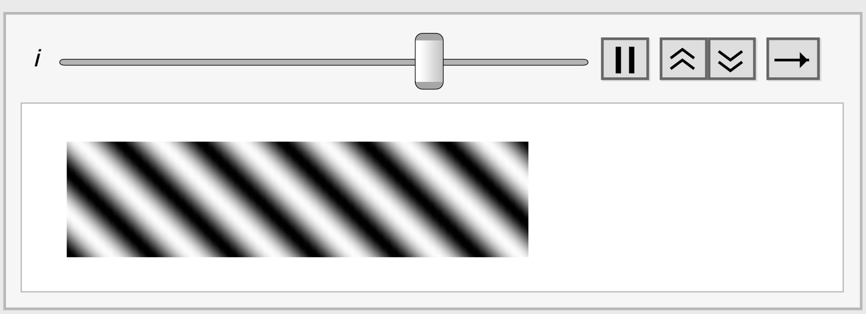

■ Rectangular vertical aperture

In[389]:=
```
Animate[
    DensityPlot[grating[x+i*stepx1,y+i*stepy1,freq1,theta1],{x,0,.25},{y,0,
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width, Aspect
```

Out[389]=

Project idea: Try the above with stereo-defined apertures

Adelson's missing fundamental motion illusion

We first make a square-wave grating.

In[390]:=
```
realsquare[x_,y_,phase_] := Sign[Sin[x + phase]];
```

And make a four-frame movie in which the grating gets progressively shifted LEFT in steps of Pi/2. That is we shift the grating left in 90 degree steps.

In[391]:=
```
Table[DensityPlot[realsquare[x,y,i Pi/2],
        {x,0,14},{y,0,1}, Frame->False,
                Mesh->False,PlotPoints -> 60, Axes->None, PlotRange->{-4,4},ImageSize->Tiny],{i,1,4,1}
]
```

Out[391]=

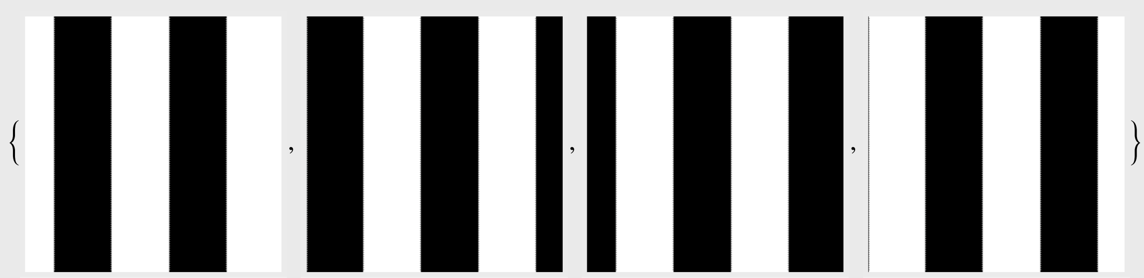

In[392]:=
```
Plot[realsquare[x, .5, Pi / 2], {x, 0, 14}]
```

Out[392]=

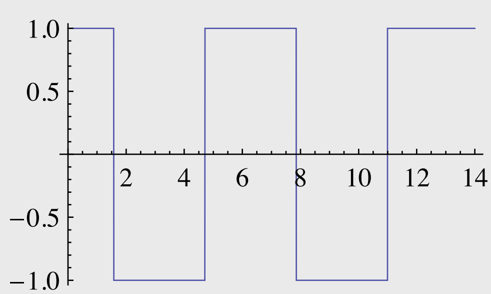

In[393]:=
```
gsq = Table[
  ArrayPlot[Table[realsquare[x, y, i Pi / 2], {y, 0, 14, .1}, {x, 0, 14, .1}],
   PlotRange → {-8, 8}, Frame -> False, ColorFunction -> "GrayTones",
        Mesh -> False, Axes -> None], {i, 1, 4, 1}
];
ListAnimate[gsq]
```

Out[393]=

A square wave can be decomposed into its Fourier components as:

realsquare(x) = (4/π)*{sin(x) + 1/3 sin(3x) + 1/5 sin(5x) + 1/7 sin(7x) + ...}

Now subtract out the fundamental frequency from the square wave

...leaving (4/π)*{1/3 sin(3x) + 1/5 sin(5x) + 1/7 sin(7x) + ...}

In[394]:=
```
realmissingfundamental[x_,y_,phase_] := realsquare[x,y,phase] - (4.0 / Pi) Sin[x + phase];
```

And make another four-frame movie in which the missing fundamental grating gets progressively shifted LEFT in steps of Pi/2. That is we shift the grating left in 90 degree steps.

It is well-known that a low contrast square wave with a missing fundamental appears similar to the square wave (with the fundamental). (There is a pitch analogy in audition.) One reason is that we are more sensitive to sharp than gradual changes in intensity. If you look at the luminance profile with the missing fundamental, you would probably guess that the perceived motion for this sequence would appear to move to the left, as before. But it doesn't. Surprisingly, the missing fundamental wave appears to move to the right!

In[395]:=
```
gsqr = Table[
  ArrayPlot[Table[realmissingfundamental[x, y, i Pi / 2], {y, 0, 14, .1},
    {x, 0, 14, .1}], PlotRange → {-8, 8}, Frame -> False,
   ColorFunction -> "GrayTones",
        Mesh -> False, Axes -> None], {i, 1, 4, 1}
];
ListAnimate[gsqr]
```

Out[395]=

Play the above movie. It typically appears to be moving to the right. You can generate movies with different contrasts by adjusting the PlotRange parameters.

In fact the missing fundamental frequency moves towards the left as you can see by playing the movie below.

In[396]:=
```
Animate[Plot[Sin[x + i Pi/2],{x,0,14},
        PlotPoints -> 60, Axes->None],{i,1,4,1}]
```

Out[396]=

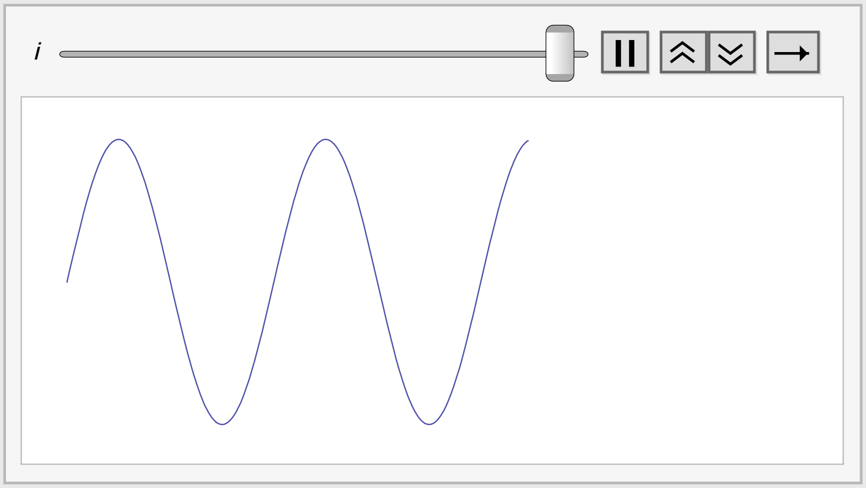

What in the stimulus does move to the right?

Why might this be? Probably the best explanation comes from looking at the dominant frequency component in the pattern, which is the 3rd harmonic. It turns out that the third harmonic is jumping in 1/4 cycle steps to the right, even though the pattern as a whole is jumping in 1/4 cycle steps (relative to the missing fundamental) to the left, as shown in the figure below:

Make a movie with Plot[ ] that shows the third harmonic. Which way does it move?

In[397]:=

```
Animate[
        Plot[Sin[3 (x + i Pi/2)],{x,0,14},PlotPoints -> 60],{i,1,4,1}
]
```

Out[397]=

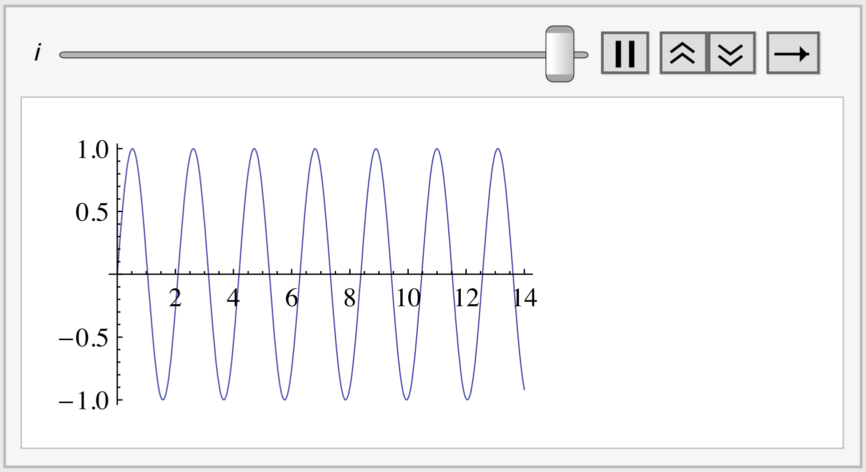

And here is the movie with just the third harmonic. Which way does it move?

In[398]:=
```
Animate[
        DensityPlot[Sin[3 (x + i Pi/2)],
                {x,0,14},{y,0,1}, Frame->False,
                Mesh->False,PlotPoints -> 60, Axes->None, PlotRange->{-4,4}],{i,1,4,1}
]
```

Out[398]=

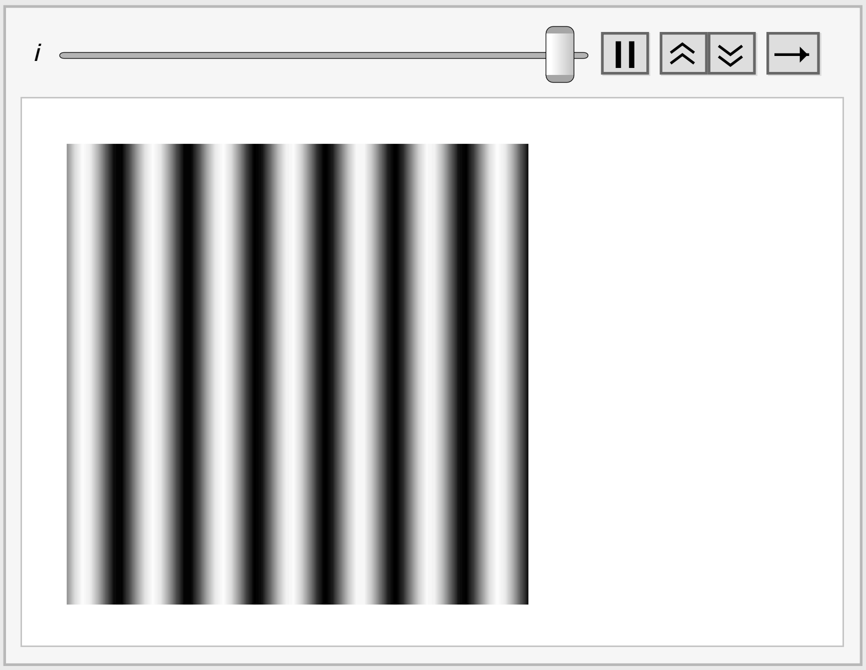

The main conclusion drawn from this demonstration is that human motion measurement mechanisms are tuned to spatial frequency.

How can the inferred biological mechanisms be pieced together to compute optic flow? We can construct the following rough outline. (For an algorithm for optic flow based on biologically plausible spatiotemporal filters see Heeger, 1987). Assume we have, at each spatial location, a collection of filters tuned to various orientations (q) and speeds (s) over a local region. (Already we run into problems with this simple interpretation, because many V1 cells are known to be tuned to spatial and temporal frequency in such a way that the spatio-temporal filter is the product of the space and time filters. This means that there is a favored temporal frequency that is the same across spatial frequencies, so the filter will be tuned to different speeds depending on the spatial frequency).

In this scheme, the optic flow measurements are distributed across the units, so if we wanted to read off the velocity from the pattern of activity, we would need some additional processing. For example, the optic flow components could be represented by the "centers of mass" across the distributed activity. Because these measurements are local, we still have the aperture problem. We will look at possible biological solutions to this problem later.

Project idea: Try the above with contours of low amplitude, rather than contrast gratings

---

Moving rhombus illusions

http://www.cs.huji.ac.il/~yweiss/Rhombus/rhombus.html

Motion Plaids

Two overlapping (additive transparent) sinusoids at different orientations and moving in different directions are, under certain conditions seen as a single pattern moving with a velocity consistent with an intersection of constraints. Under other conditions, the two individual component motions are seen.

■ Adding two gratings, single frame

Grating 1

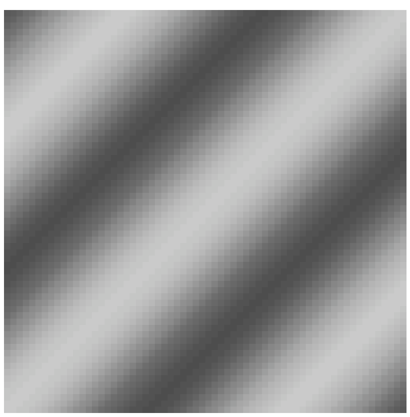

Grating 2

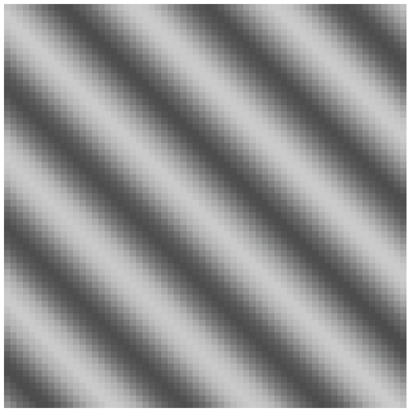

Plaid grating: Grating 1 + Grating 2

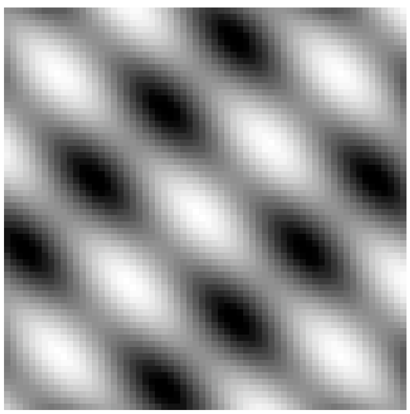

■ Initialize parameters

In[399]:=

```
niter = 16; width = 64;
theta1 = Pi/4.; contrast1 = 0.5; theta2 = -Pi/4.; contrast2 = 0.25;
freq1 = 8.; period1 = 1/freq1; freq2 = 2.; period2 = 1/freq2;
stepx1 = Cos[theta1]*(period1/niter); stepy1 = Sin[theta1]*(period1/niter);

stepx2 = Cos[theta2]*(period2/niter); stepy2 = Sin[theta2]*(period2/niter);
(*stepx = Min[stepx1,stepx2]; stepy = Min[stepy1,stepy2];*)

grating[x_,y_,freq_,theta_,contrast_] := contrast*Cos[(2. Pi freq)*(Cos[the
```

■ **Display plaid grating**

In[406]:=
```
For[i=1,i<niter + 1,i++,
    DensityPlot[grating[x+i*stepx1,y+i*stepy1,freq1,theta1,contrast1]+ grat
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width];
];
```

■ **Display grating 1 only**

In[407]:=
```
For[i=1,i<niter + 1,i++,
    DensityPlot[grating[x+i*stepx1,y+i*stepy1,freq1,theta1,contrast1],{x,0,
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width];
];
```

■ **Display grating 2 only**

In[408]:=
```
For[i=1,i<niter + 1,i++,
    DensityPlot[grating[x+i*stepx2,y+i*stepy2,freq2,theta2,contrast2],{x,0,
        Mesh->False,Frame->None,PlotRange->{-2,2},PlotPoints->width];
];
```

Now try the above motion plaid with equal spatial frequencies and contrasts

---

Orientation in space-time

In this section, we'll see how viewing motion measurement as detecting orientation in space-time is related to neurophysiological theories of neural motion selectivity.

Representation of motion

■ *Mathematica* demo

In[409]:=
```
size = 32; x0 = 4; y0 = 4; pw = 12; xoffset = 1;
A1 = Table[Random[], {size}, {size}]; (*A2 = A1;*)
A2 = Table[Random[], {size}, {size}];
A2[[Range[y0, y0 + pw], Range[x0, x0 + pw]]] =
  A1[[Range[y0, y0 + pw], Range[x0 - xoffset, x0 + pw - xoffset]]];
```

In[413]:=
```
ListAnimate[{ArrayPlot[A1, Mesh → False],
  ArrayPlot[A2, Mesh → False]}, 1 / 2]
```

Out[413]=

In[414]:=
```
nframes = 8;
xt = {};
For[i=0,i<nframes,i++,
    A2[[Range[y0,y0+pw],Range[x0,x0+pw]]] =
    A1[[Range[y0,y0+pw],Range[x0+i,x0+pw+i]]];
    xt = Join[xt,{A2[[8]]}]
];
```

In[417]:=
```
ListDensityPlot[Transpose[xt],Mesh->False,Axes->True, AxesLabel->{t,x},Imag
```

Out[417]=

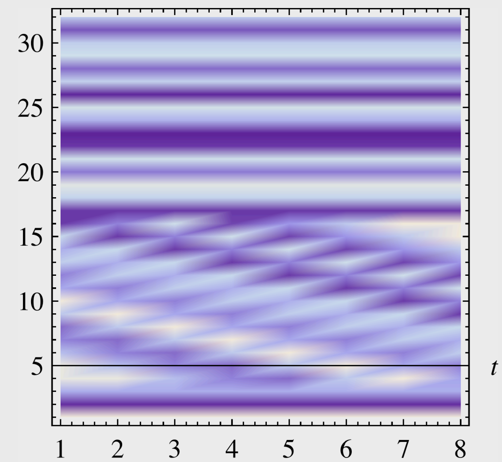

■ x-y-t space

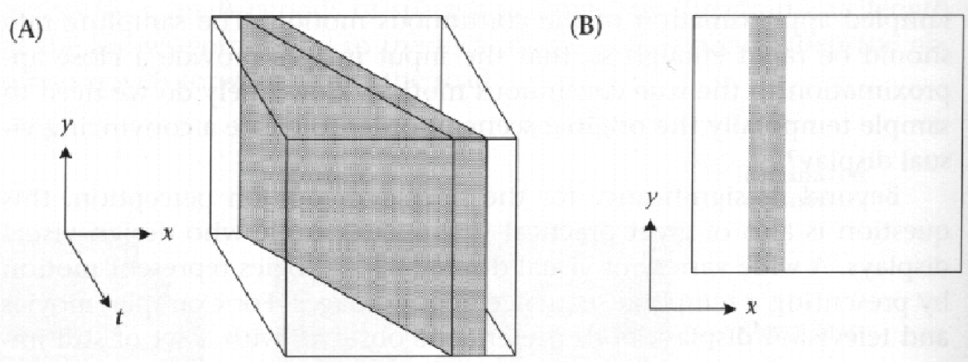

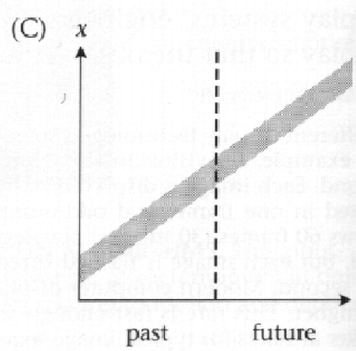

10.3 A MOTION SEQUENCE is a series of images measured over time. (A) The motion sequence of images can be grouped into a three-dimensional volume of data. (B) Cross sections of the volume show the spatial pattern at a moment in time. (C) Time (t) may be plotted against one dimension (x) of space. When the spatial pattern is one-dimensional, the (t, x) cross-section provides a complete representation of the stimulus sequence.

Neurophysiological filters

■ **Space-time filters for detecting orientation in space-time**

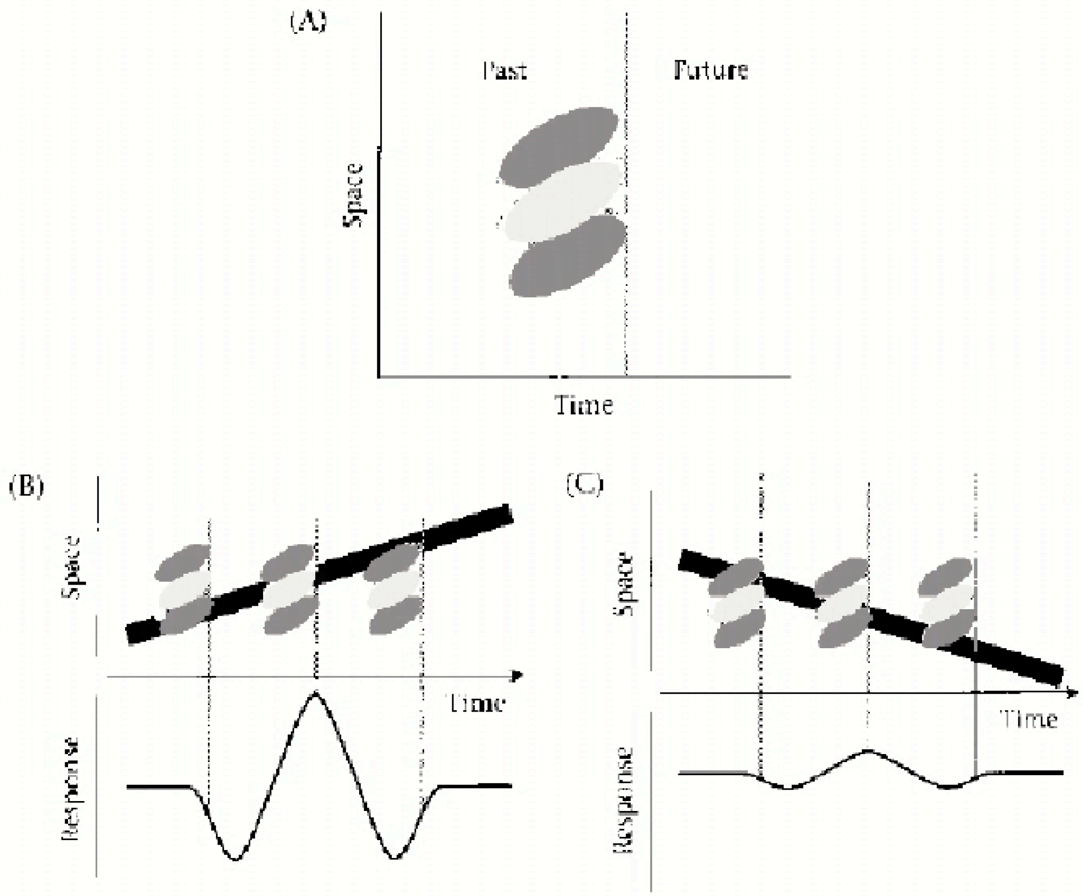

10.6 SPACE-TIME-ORIENTED RECEPTIVE FIELD. (A) The space–time receptive field of a neuron is represented on a (t, x) plot. The neuron always responds to events in the recent past, so the receptive field moves along the time axis with the present. The dark areas show an inhibitory region, and the light area shows an excitatory region. (B) The upper portion of the graph shows a (t, x) plot of a moving line and the space–time receptive field of a linear neuron. The neuron's receptive field is shown at several different moments in time, indicated by the vertical dashed lines. The common orientation of the space–time receptive field and the stimulus motion produce a large amplitude response, shown in the bottom half of the graph. (C) When the same neuron is stimulated by a line moving in a different direction, the stimulus motion aligns poorly with the space–time receptive field. Consequently, the response amplitude is much smaller.

From **Wandell, "Foundations of Vision", 1995**

■ A possible mechansim for building space-time filters from two spatial filters with a temporal delay

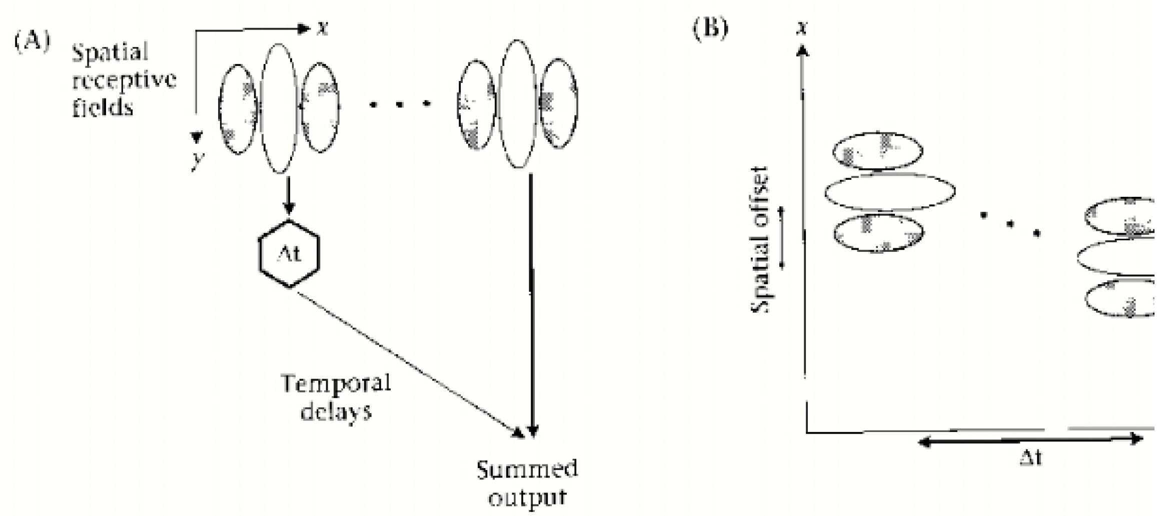

10.7 A METHOD FOR CREATING A SPACE–TIME-ORIENTED RECEPTIVE FIELD. (A) A pair of spatial receptive fields, displaced in the x direction, is shown at the top. The response of the neuron on the left is delayed and then added to the response of the neuron on the right. (B) The (t, x) receptive field of the output neuron in panel (A). The temporal response of the neuron on the left is delayed compared to the temporal response of the neuron on the right. The combination of spatial displacement and temporal delay yields an output neuron whose receptive field is oriented in space–time.

Wandell, "Foundations of Vision", 1995

■ Relationship of the gradient constraint to oriented space-time filters

$$v_x \frac{\partial L}{\partial x} + v_y \frac{\partial L}{\partial y} + \frac{\partial L}{\partial t} = 0$$

v_x and v_y correspond to u and v used in the previous lecture.

$$v_x \frac{\partial L}{\partial x} + \frac{\partial L}{\partial t} = 0$$

Image L(x,y,t) -> blurred in space and smeared in time, g(x,y,t). Consider just one spatial dimension, (t,x) space.

$$v_x \frac{\partial g}{\partial x} + \frac{\partial g}{\partial t} = 0$$

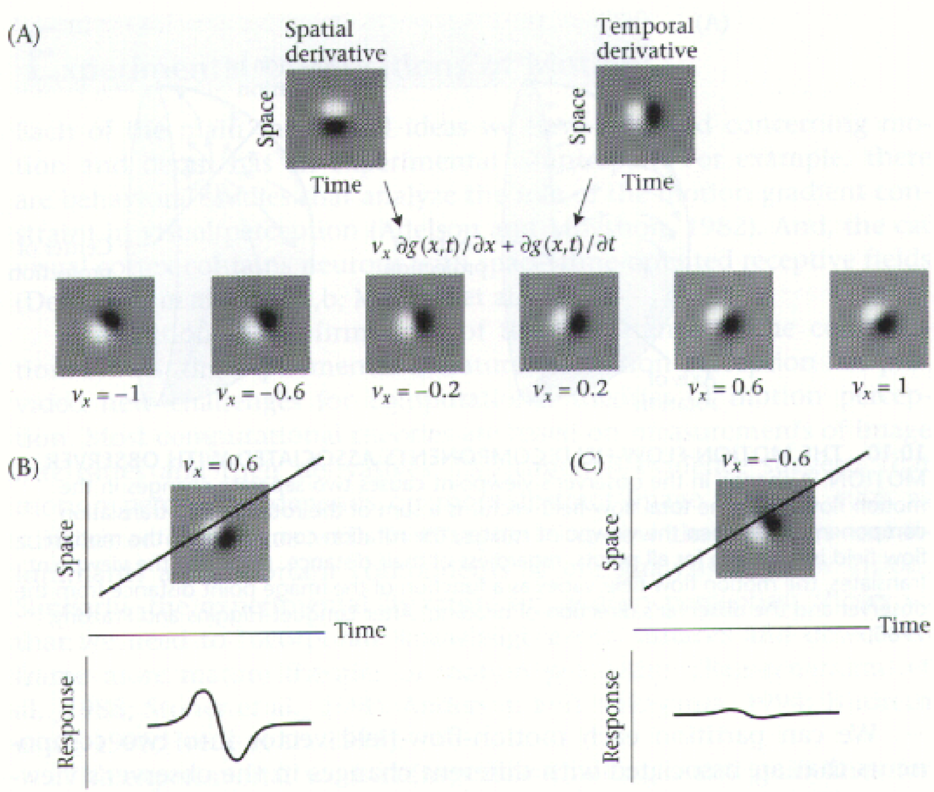

10.9 THE MOTION GRADIENT CONSTRAINT REPRESENTED IN TERMS OF SPACE–TIME RECEPTIVE FIELDS. (A) The spatial and temporal derivatives can be computed using neurons whose (t, x) receptive fields are shown at the top. We can form weighted sums of these neural responses to create new receptive fields that are oriented in space–time. The response amplitudes of these neurons can be used to identify the motion of a stimulus. The receptive field of the neuron represented in (B) responds strongly to the stimulus motion while the receptive field of the neuron in (C) responds weakly. By comparing the response amplitudes of the array of neurons, one can infer the stimulus motion.

http://www.amazon.com/Foundations-Vision-Brian-Wandell/dp/0878938532/ref=sr_1_1?ie=-UTF8&s=books&qid=1226337031&sr=1-1

In[418]:=
```
grating[x_, t_, fx_, ft_, φ_, σ_] :=
  Exp[-((x^2 + t^2) / σ)^2] * Sin[2 π (fx x + ft t) + φ];
fx = 1; ft = 0; φ1 = 0; σ = .20;

dgdx = Table[grating[x, t, fx, ft, φ1, σ], {x, -2, 2, .05}, {t, -2, 2, .05}];
fx = 0; ft = 1; φ1 = 0; σ = .20;
dgdt = Table[grating[x, t, fx, ft, φ1, σ], {x, -2, 2, .05},
   {t, -2, 2, .05}];
Manipulate[ArrayPlot[vx * dgdx + dgdt], {vx, 0, Pi}]
```

Out[423]=

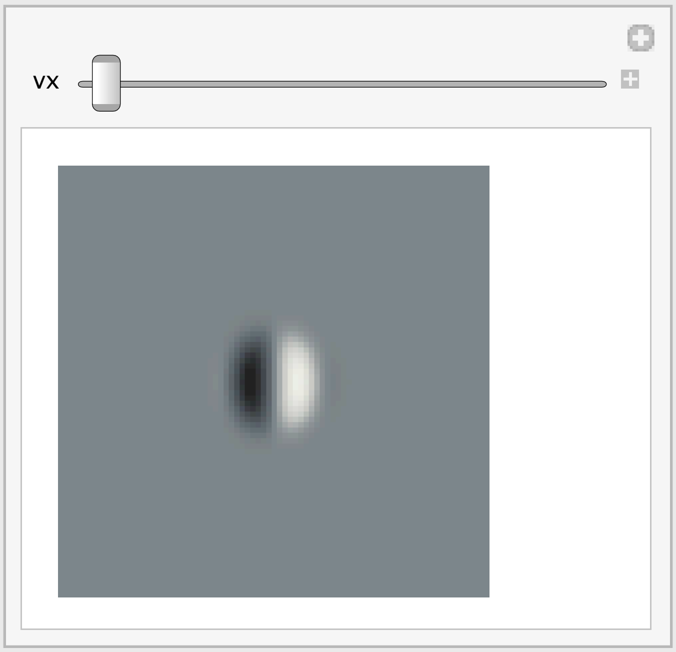

Bayesian model for integrating local motion measurements

Global integration.

Yuille, A., & Grzywacz, N. (1988);

Lorenceau & Shiffrar's demo (http://gandalf.psych.umn.edu/users/kersten/kersten-lab/courses/Psy5036W2008/Lectures/18.MotionOpticFlow/aperturedemomovie.mov)

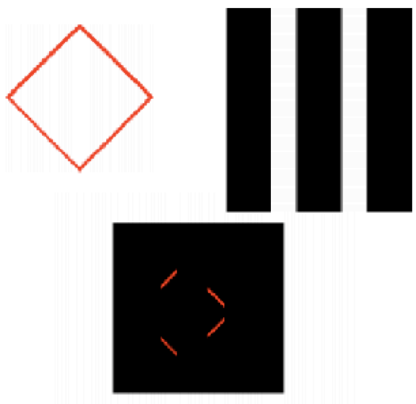

General problem

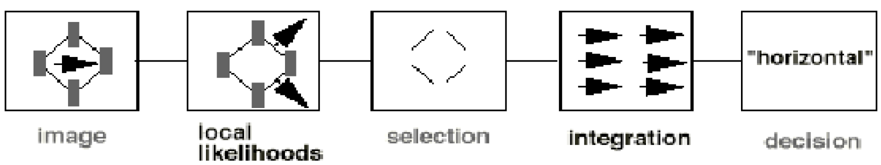

Intersection of constraints revisited

Grating plaids sometime seen as coherent, other times as two overlapping transparent gratings moving separately.

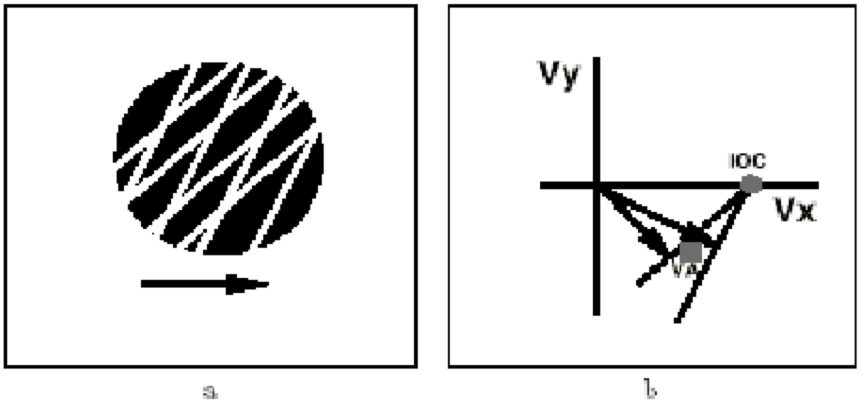

Bayes model for integration

Yuille, A., & Grzywacz, N. (1988)

Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. Nat Neurosci 5:598-604.

■ Probabilistic interpretation of intersection of constraints

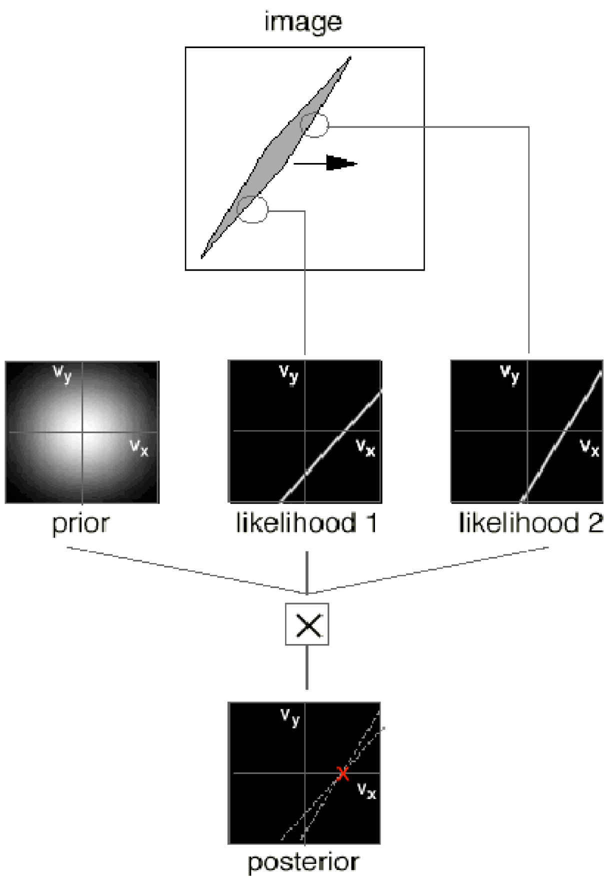

The plot illustrates the calculation of the posterior:

p(v_x,v_y| perpendicular component1, perpendicular component 2) ∝ p(perpendicular component1 | v_x)p(perpendicular component 2 | v_x) p(v_x,v_y)

■ Probabilistic interpretation with noisy measurements

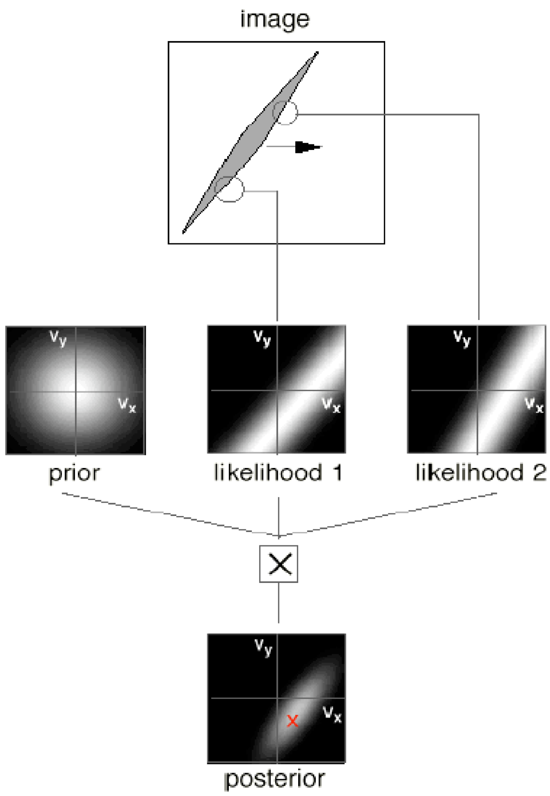

■ Key ideas

A. Information for motion direction and speed comes from two sources: 1) the data, which involves many measurements of local velocity that produce a likelihood of (v_x v_y) for each measurement. (Two are shown above). These likelihoods have various degrees of uncertainty (i.e. variance in the possible values of (v_x v_y)) that depend on image signal-to-noise ratio, e.g. contrast. 2) Prior assumptions that assume the speeds are slow, i.e. a probability distribution of (v_x v_y) with a mean of zero.

B. A Bayes optimal solution multiplies the prior and likelihoods to obtain the posterior. The important qualitative idea is that estimates based on this posterior effectively weight information from the data and prior according to reliability. So if there is more certainty in the measurements (e.g. high contrast), this will bias the estimates of (v_x v_y) towards the intersection of constraints and away from the prior. If it is hard to see the motion, then the estimates should be biased towards the prior, i.e. slower.

Exercise: Assume Gaussian distributions and prove that the maximum a posteriori estimate of a parameter given two measurements is given by:

---

v = $v_1 r_1/(r_1 + r_2) + v_2 r_2/(r_1 + r_2)$

Where v_1 and v_2 are the best estimates based on measurements 1 and 2 separately, and $r_i = 1 / \sigma_i^2$, $i.e.$ the reciprocal of the variance of each. The math for this is identical to that for cue integration ($e.g.$, see Lecture 6, and Ernst, $M.O.$, & Banks, $M.S.$(2002).Humans integrate visual and haptic information in a statistically optimal fashion.Nature, 415 (6870), 429 − 433.)

■ Generalize to other types of motion stimuli

Requirements for generalizatoin:

Base likelihoods on actual image data

spatiotemporal measurements

Include “2D” features

E.g. corners

Rigid rotations, non-rigid deformations

Stage 1:local likelihoods

Stage 2: Bayesian combination

- Prior

slowness -- wagon wheel example, quartet example

smoothness - e.g. translating rigid circle

■ Overview of Weiss & Adelson theory

http://www-bcs.mit.edu/people/yweiss/intro/intro.html

Dense to sparse: $\mathbf{v}(\mathrm{r}) = \Phi(\mathrm{r})\theta$

$$v_x(x, y) = \sum_{i=1}^{N/2} \theta_i G(x - x_i, y - y_i)$$

$$v_y(x, y) = \sum_{i=1+N/2}^{N} \theta_i G(x - x_i, y - y_i)$$

Likelihood: $$L(v) \propto e^{-\sum_r w(r)(I_x v_x + I_y v_y + I_t)^2 / 2\sigma^2}$$

$$L_r(v) \to p(I \mid \theta) \propto \prod_r L_r(\theta)$$

Prior: $$P(V) \propto e^{-\sum_r (Dv)^t (r)(Dv)(r) / 2}$$

$$P(V) \to P(\theta)$$

Posterior: $$P(\theta \mid I) \propto P(I \mid \theta) P(\theta)$$

Log posterior is quadratic in θ, –> linear estimator for θ

Weiss & Adelson, 1998

Tests of theory

■ Rhombus experiment

The above figure shows how as the rhombus gets skinnier, the peak of the posterior moves towards the lower right quadrant of velocity space, consistent with psychophysics.

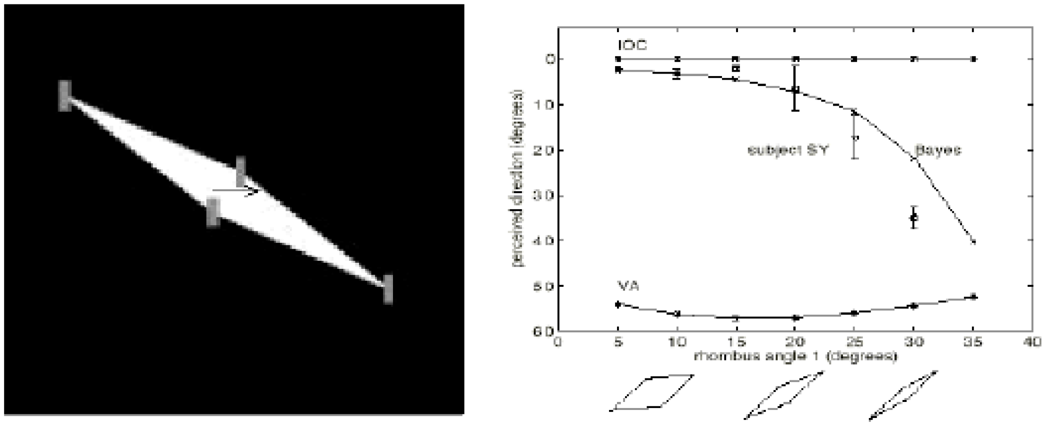

a b

■ Aperture effects

Imagine a corrugated surface moving up, but viewed through apertures. For a circular aperture, information around the boundary is symmetric, so the bias for a certain direction balances out, leaving the interior velocity measurements to dominate.

For a rectangular aperture, corner information can provide a strong bias.

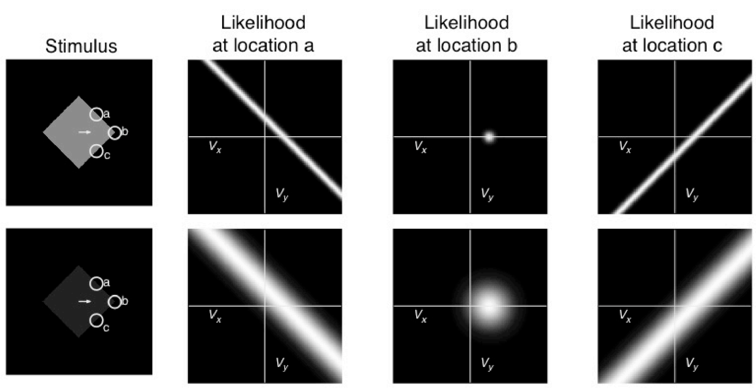

Fig. 3. Likelihood functions for three local patches of a horizontally translating diamond stimulus, computed using equation (4). Intensity corresponds to probability. Top, high-contrast sequence. Bottom, low-contrast sequence, with the same parameter σ. At edges, the local likelihood is a 'fuzzy' constraint line; at corners, the local likelihood peaks around the veridical velocity. The sharpness of the likelihood decreases with decreasing contrast.

So for the rectangular aperture below, there is more and overall stronger evidence for leftward motion.

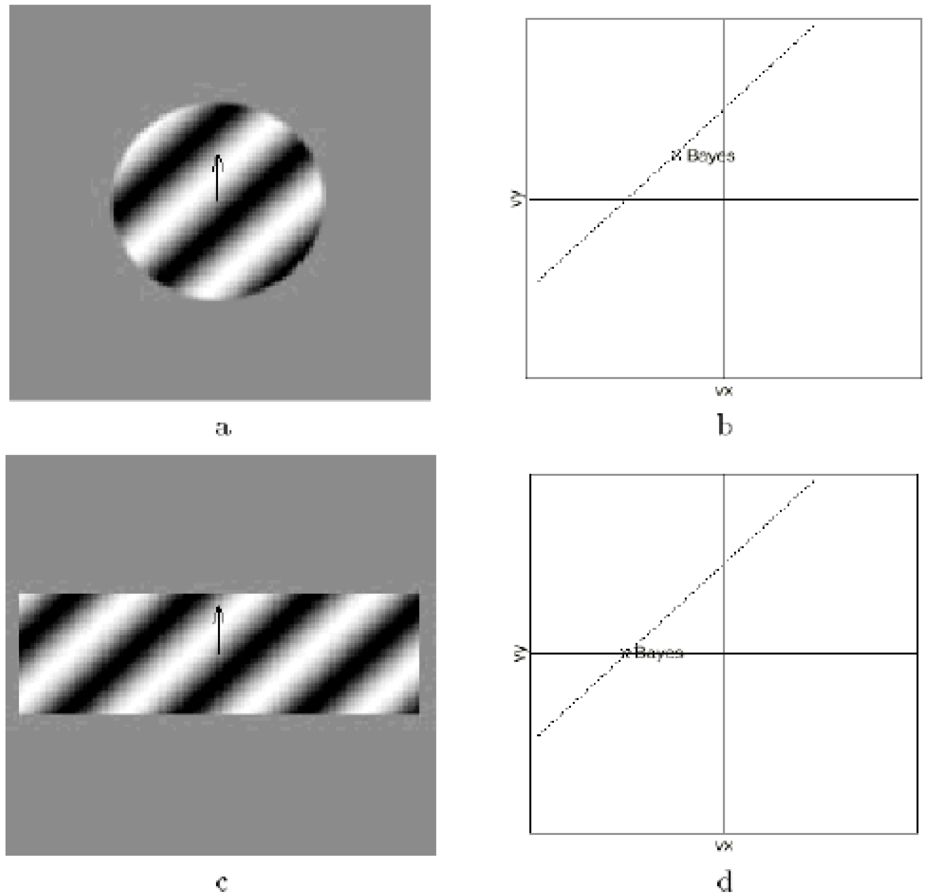
Bayes
vy
vx
Bayes
vy
vx
a
b
c
d

■ Plaids

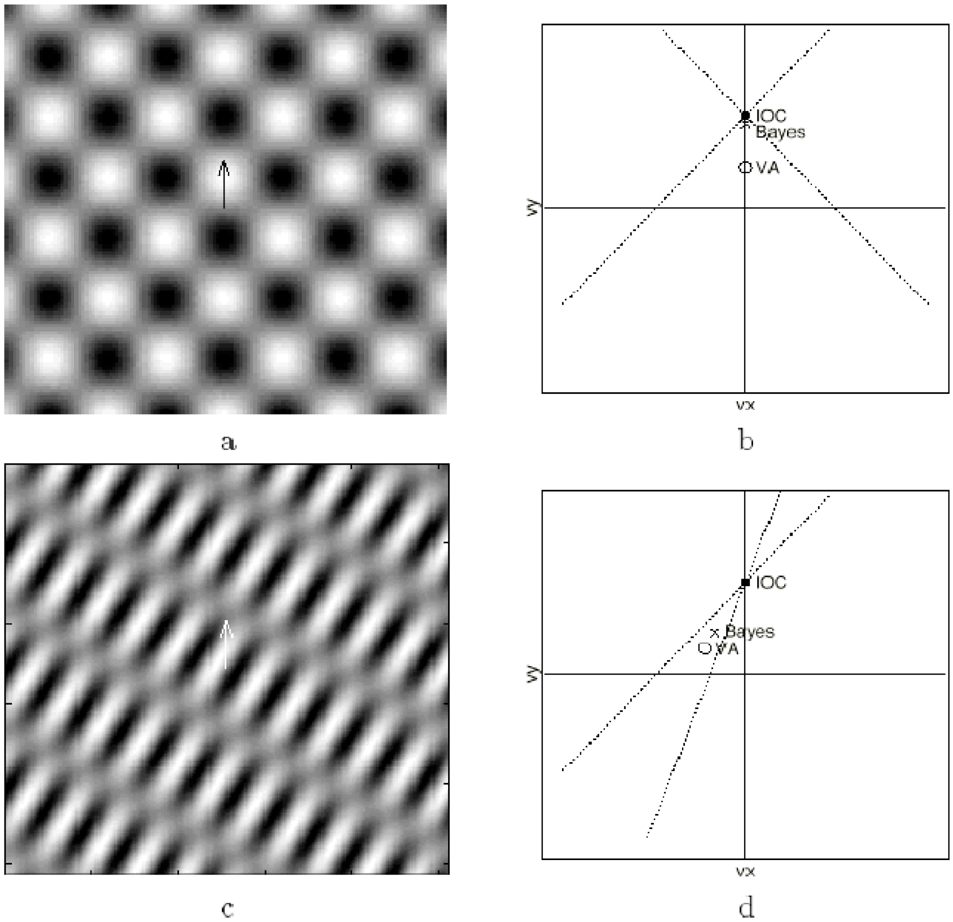

a b c d

From Weiss and Adelson, 1998. Type I and I plaids. (Yo and Wilson, 1992)

References

Adelson EH, Movshon JA (1982) Phenomenal Coherence of Moving Visual Patterns. Nature 300:523-525.

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal Energy Models for the Perception of Motion. Journal of the Optical Society of America, 2(2), 284-299.

Barlow, H. B., & Levick, R. W. (1965). The Mechanism of Directional Selectivity in the Rabbit's Retina. Journal of Physiology, 173, 477-504.

Barlow, H. (1996). Intraneuronal information processing, directional selectivity and memory for spatio-temporal sequences. Network: Computation in NeuralSystems, 7, 251-259.

Braddick, O. J. (1974). A short-range process in apparent motion. Vision Research, 14, 519-527.

Hassenstein, B., & Reichardt, W. (1956). Functional Structure of a Mechanism of Perception of Optical Movement. Proceedings First International Congress on Cybernetics, 797-801.

He, S., & Masland, R. H. (1997). Retinal direction selectivity after targeted laser ablation of starburst amacrine cells. Nature, 389(6649), 378-82.

Hildreth, E. C. (1984). The Measurement of Visual Motion. Cambridge, MA: MIT Press.

Koch, C., Torre, V., & Poggio, T. (1986). Computations in the Vertebrate Retina: Motion Discrimination Gain Enhancement and Differentiation. Trends in Neuroscience, 9(5), 204-211.

Movshon JA, Adelson EH, Gizzi MS, Newsome WT (1986) The Analysis of Moving Visual Patterns. In: Pattern Recognition Mechanisms (Chagas C, Gattas R, Gross CG, eds), pp 117-151. Rome, Italy: Vatican Press.

Schrater PR, Simoncelli EP (1998) Local velocity representation: evidence from motion adaptation. Vision Res 38:3899-3912.

Schrater PR, Knill DC, Simoncelli EP (2000) Mechanisms of visual motion detection. Nature Neuroscience 1:64 - 68.

Van Santen, J. P. H., & Sperling, G. (1985). Elaborated Reichardt Detectors. Journal of the Optical Society of America, 2(7), 300-321.

Wandell, B. A. (1995). Foundations of Vision. Sunderland, Massachusetts: Sinauer.

Weiss, Y., & Adelson, E. H. (1998). Slow and smooth: a Bayesian theory for the combination of local motion signals in human vision (A.I. Memo No. 1624). M.I.T.

http://www-bcs.mit.edu/people/yweiss/intro/intro.html

Weiss Y, Simoncelli EP, Adelson EH (2002) Motion illusions as optimal percepts. Nat Neurosci 5:598-604.

Yuille, A., & Grzywacz, N. (1988). A computational theory for the perception of coherent visual motion. Nature, 333, 71-74.