Bayesian vision: The early years

Dan Kersten
David Knill Memorial Symposium
VSS 2015

The “state of affairs” ~ 1985

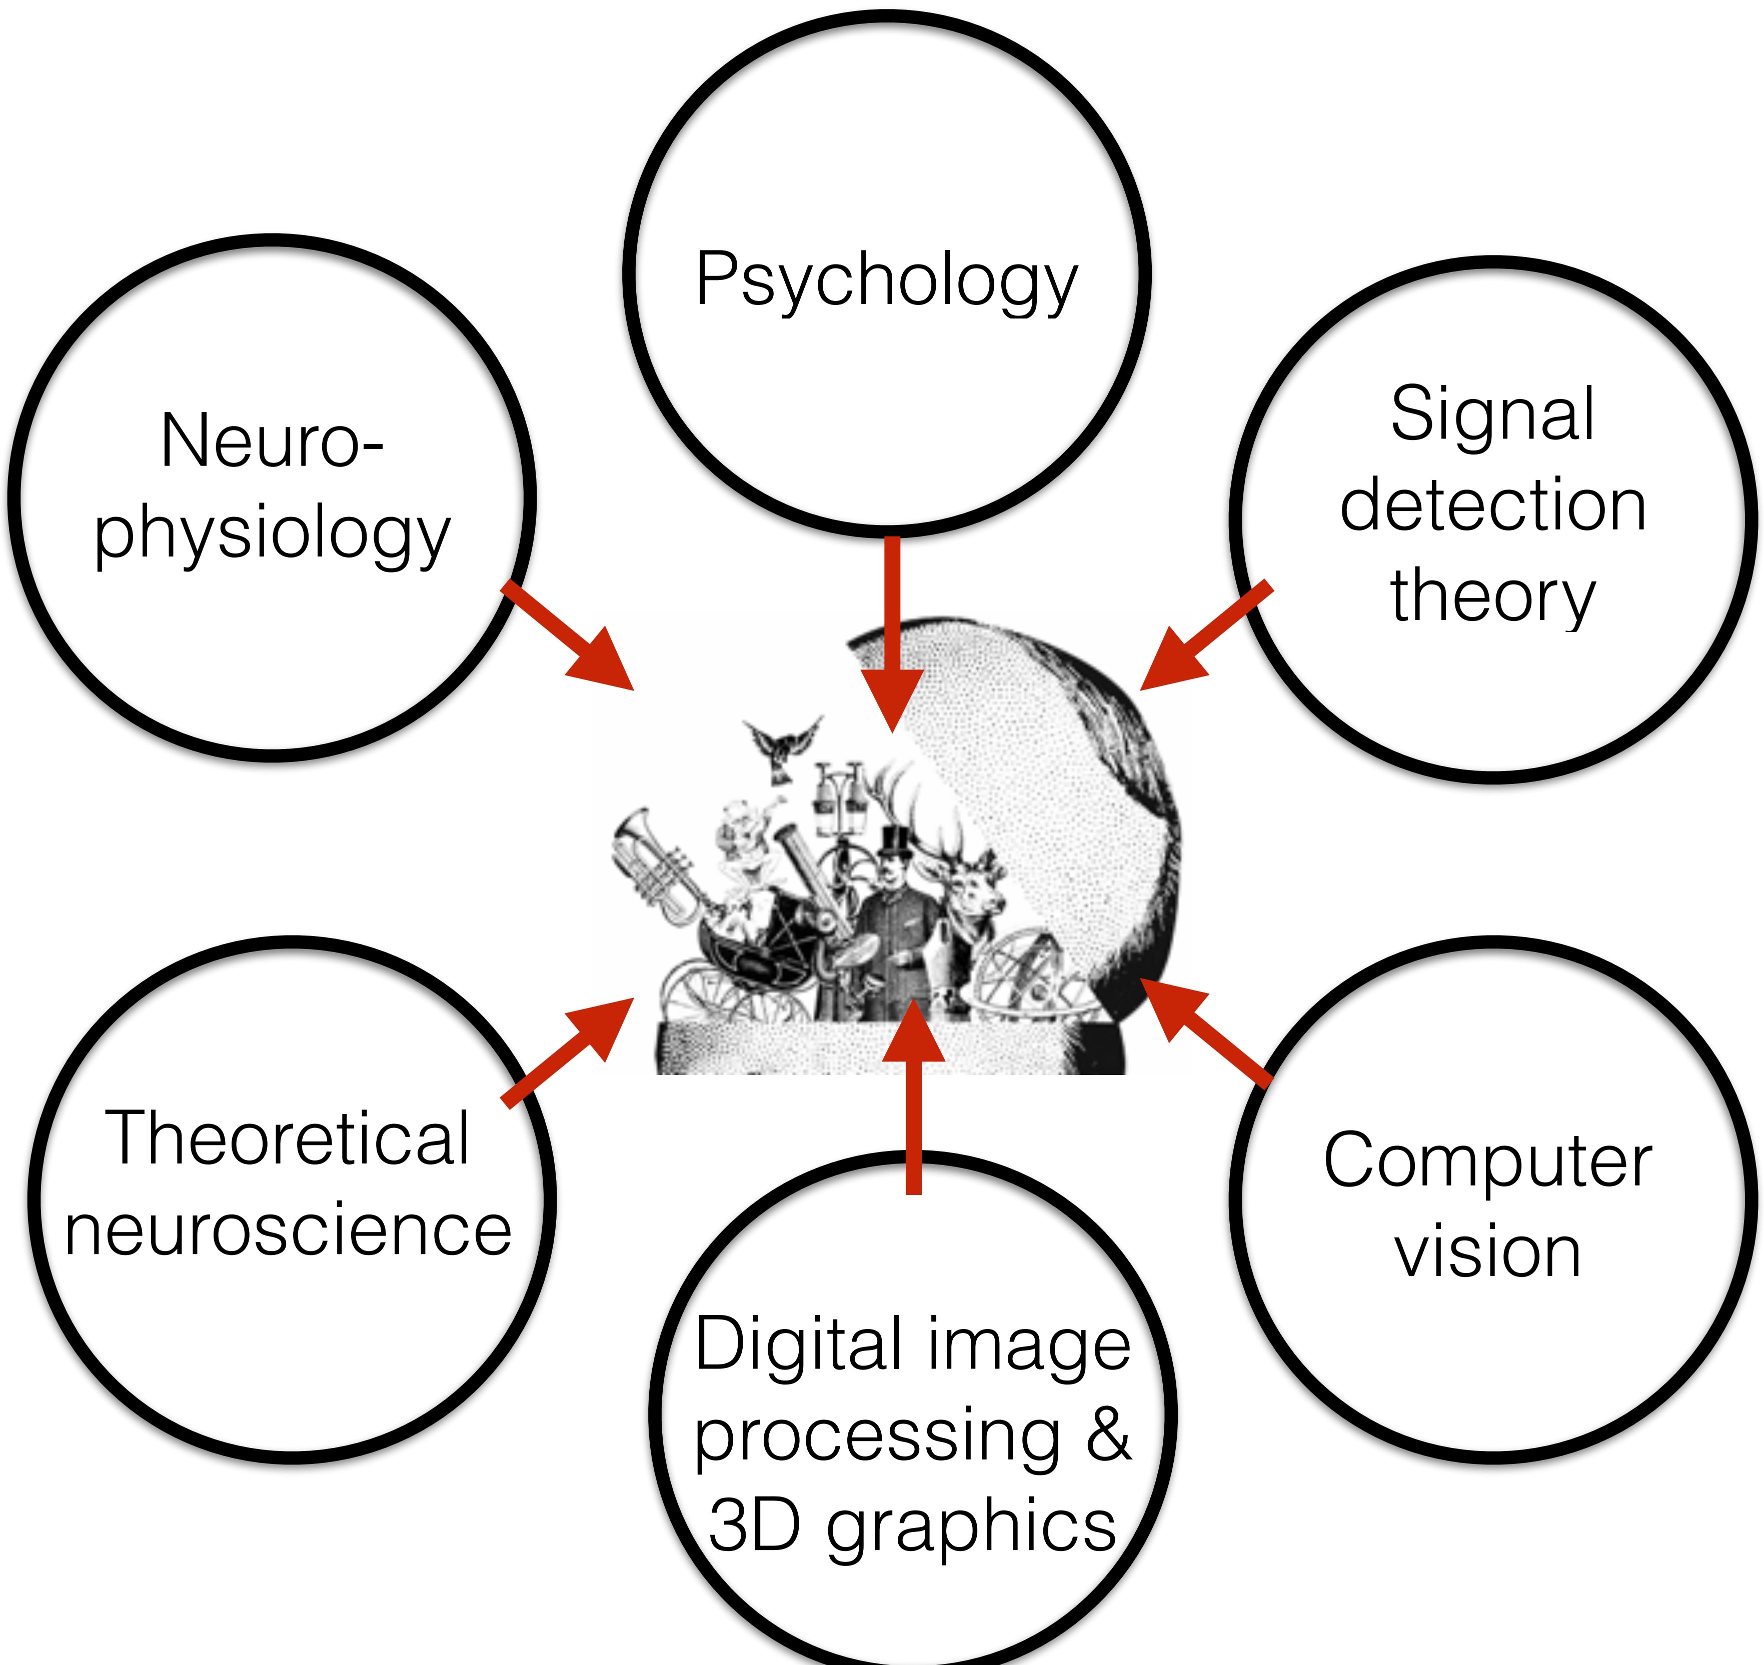

The “state of affairs” ~ 1985

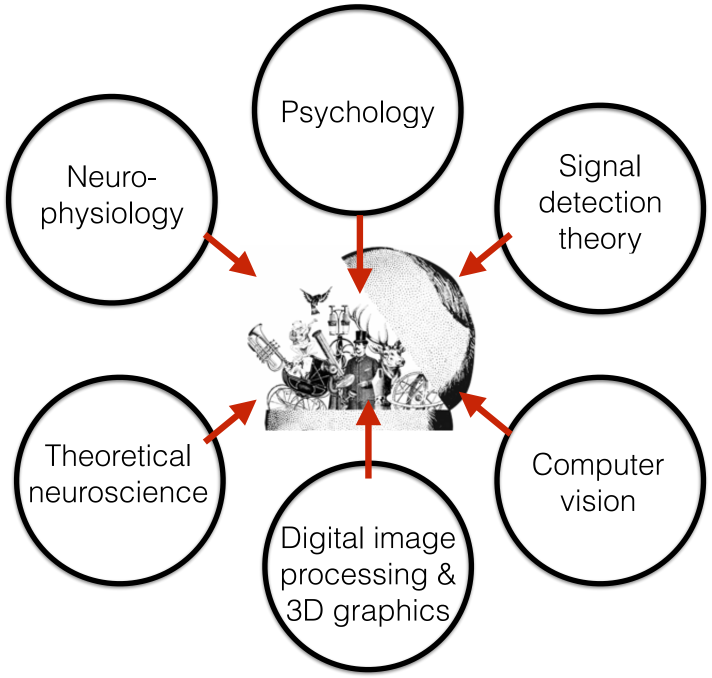

The “state of affairs” ~ 1985

ideal observer analysis

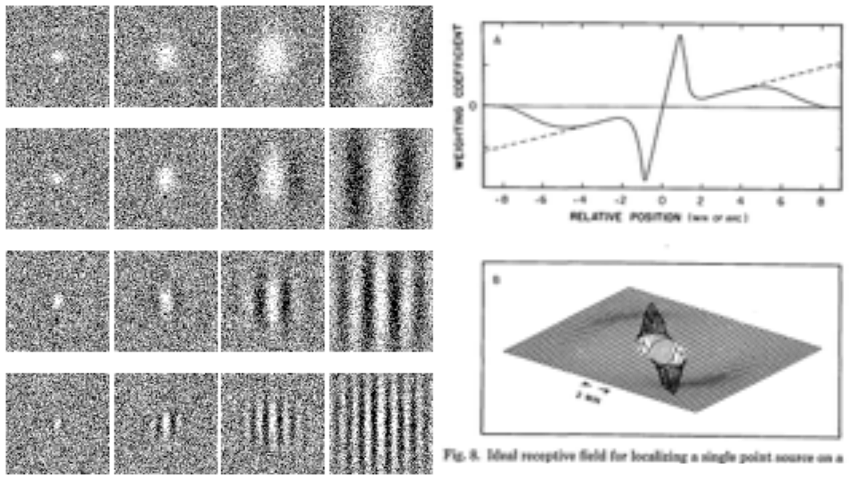

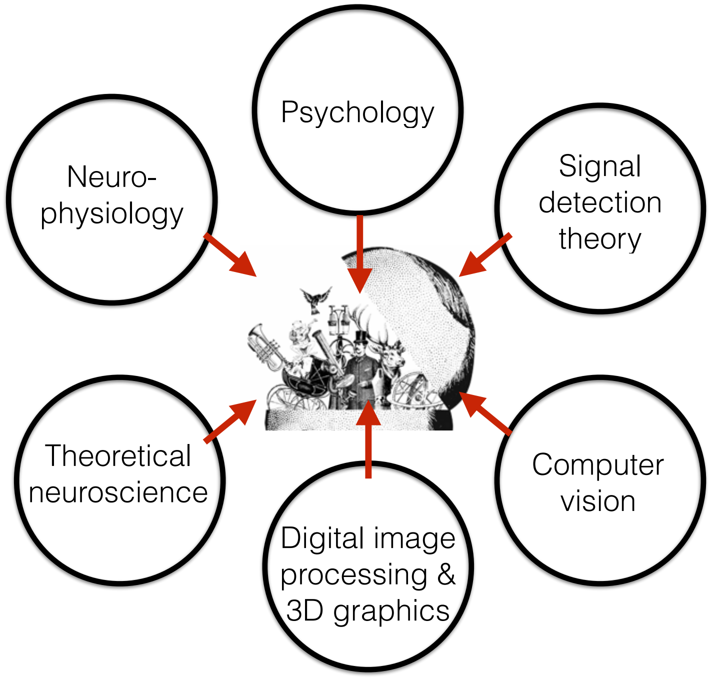

The “state of affairs” ~ 1985

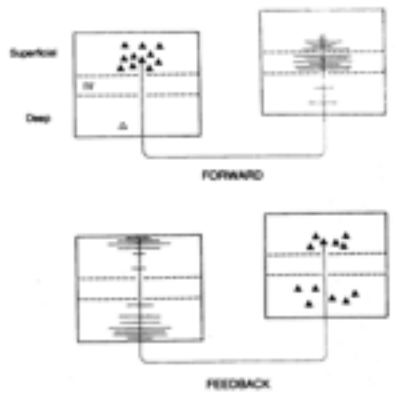

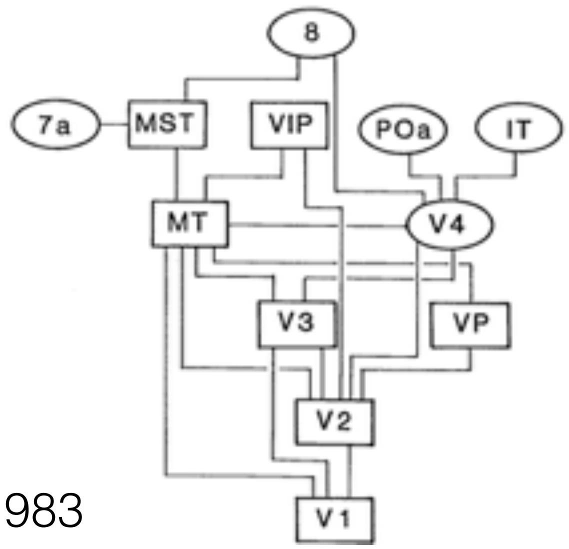

1983

ideal observer analysis

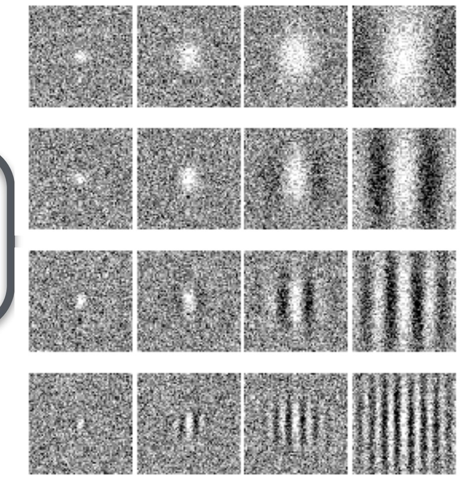

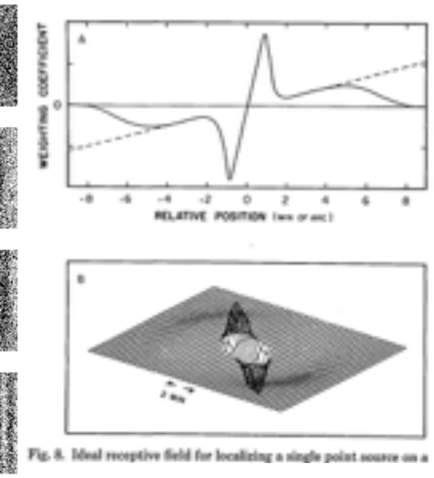

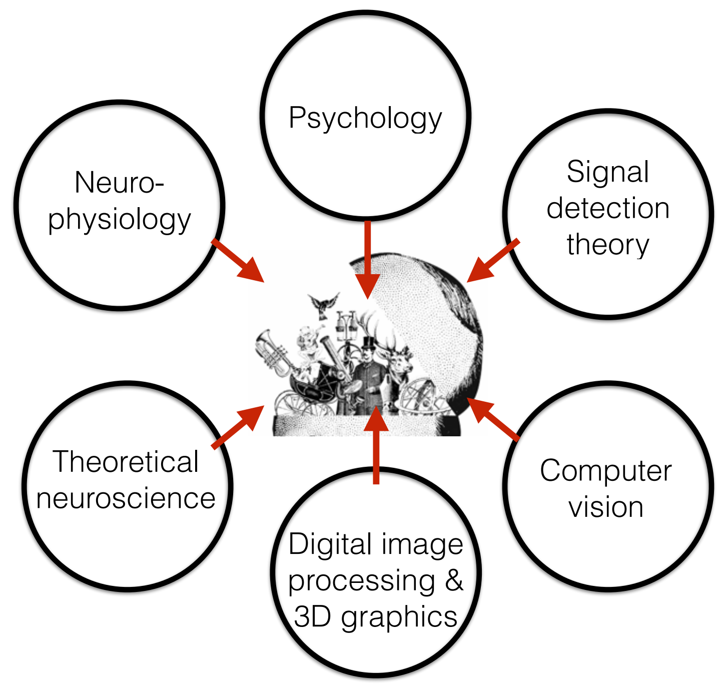

The “state of affairs” ~ 1985

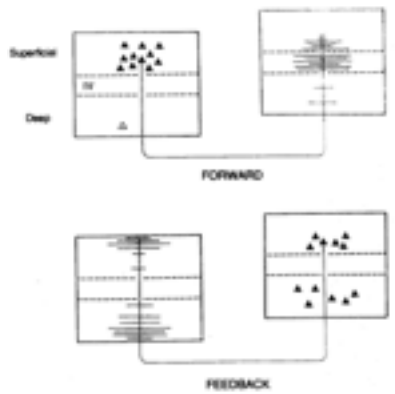

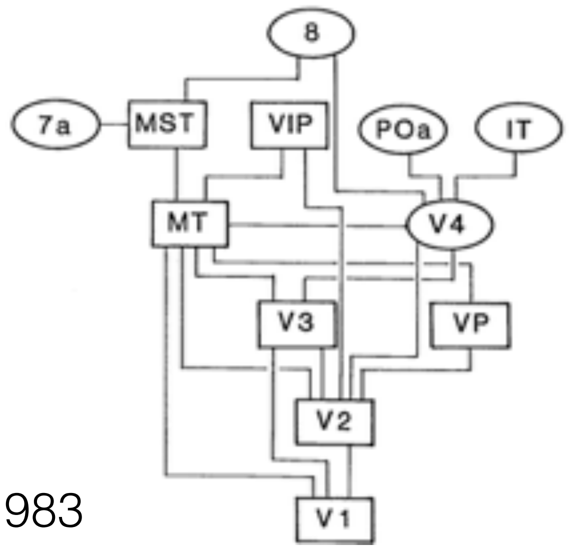

1983

ideal observer analysis

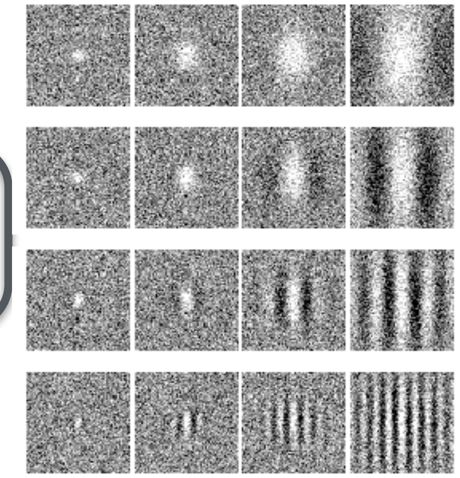

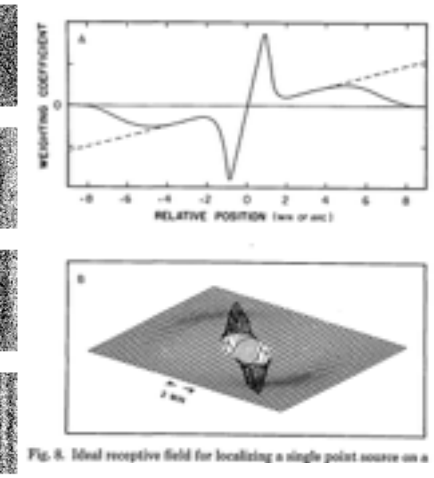

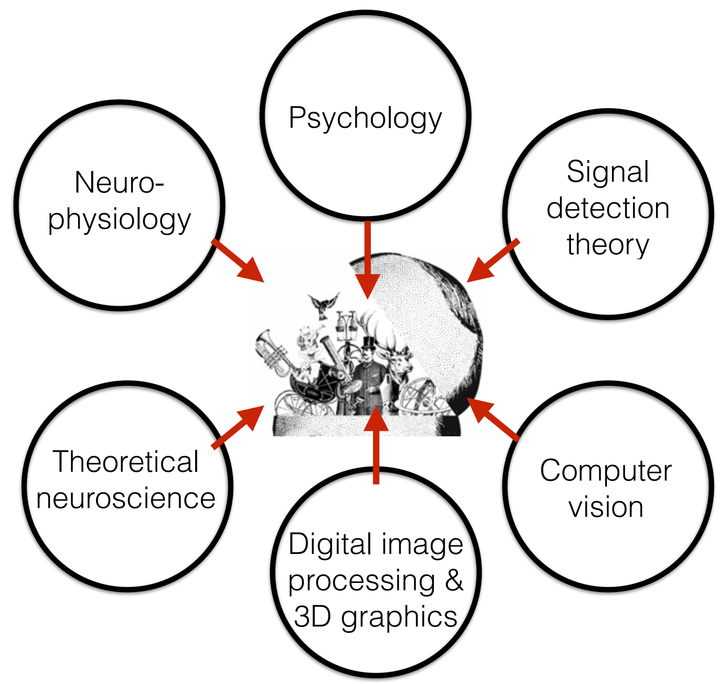

A Learning Algorithm for Boltzmann Machines*

1985

The “state of affairs” ~ 1985

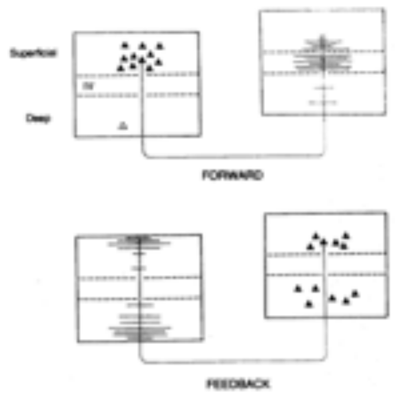

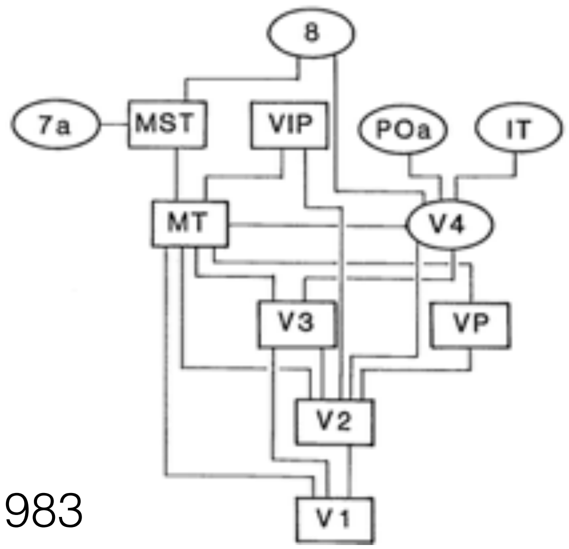

1983

ideal observer analysis

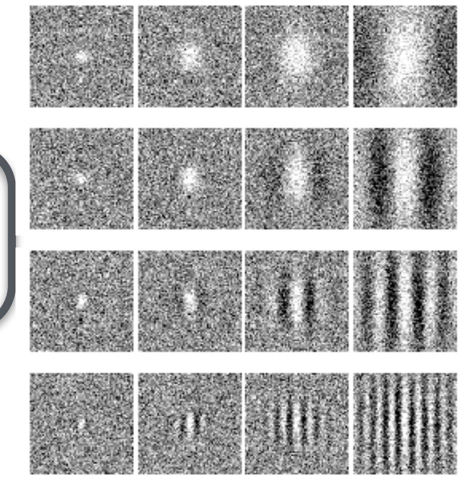

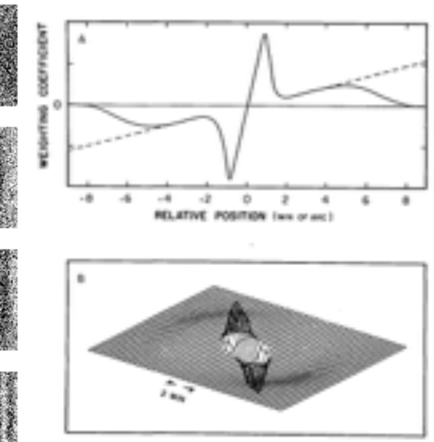

Fig. 8. Ideal receptive field for localizing a single point source on a

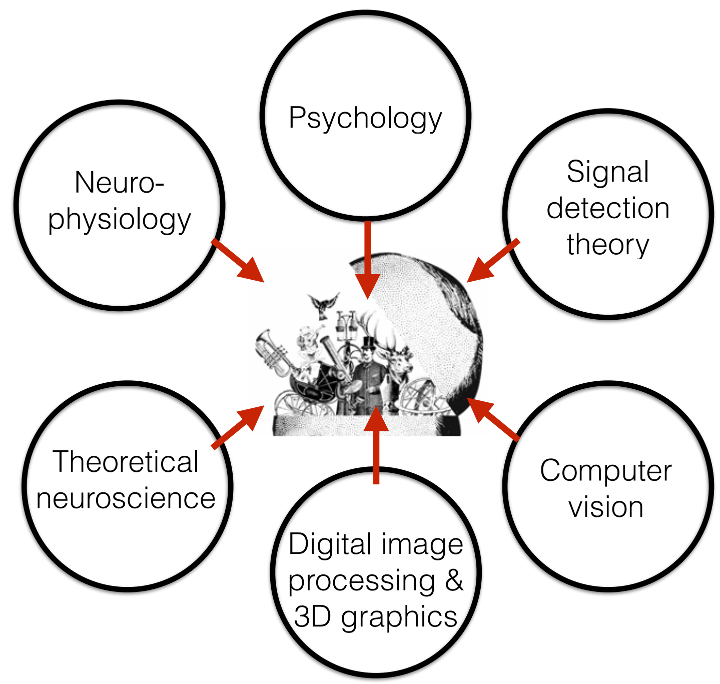

A Learning Algorithm for Boltzmann Machines*

1985

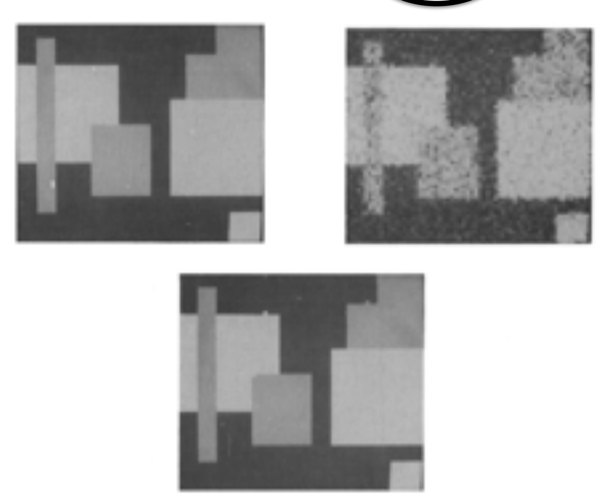

1984

The "state of affairs" ~ 1985

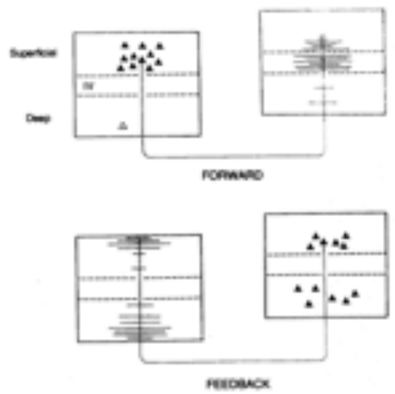

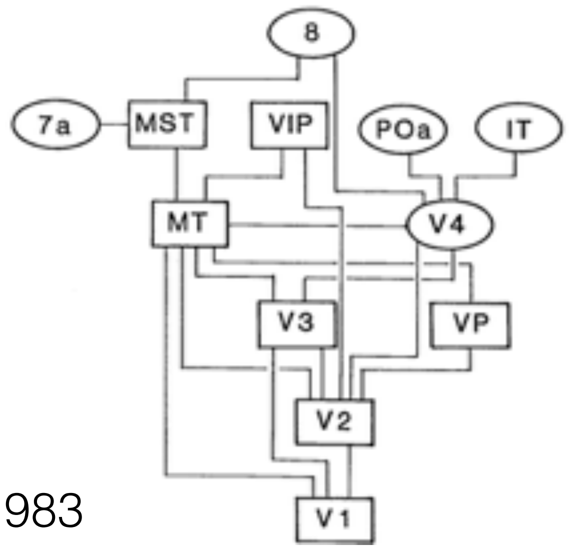

1983

ideal observer analysis

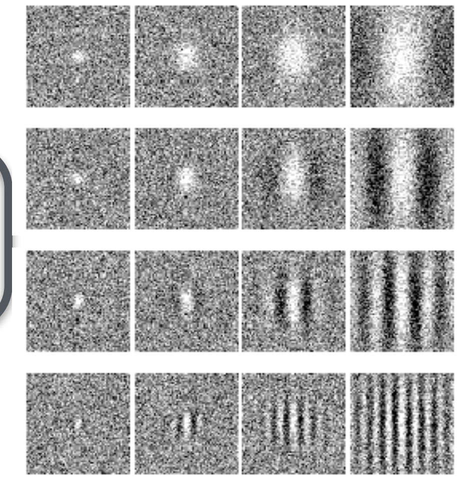

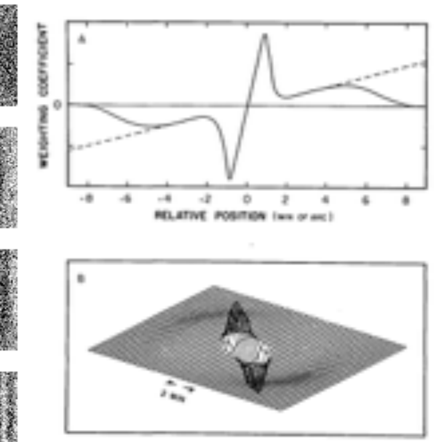

Fig. 8. Ideal receptive field for localizing a single point source on a

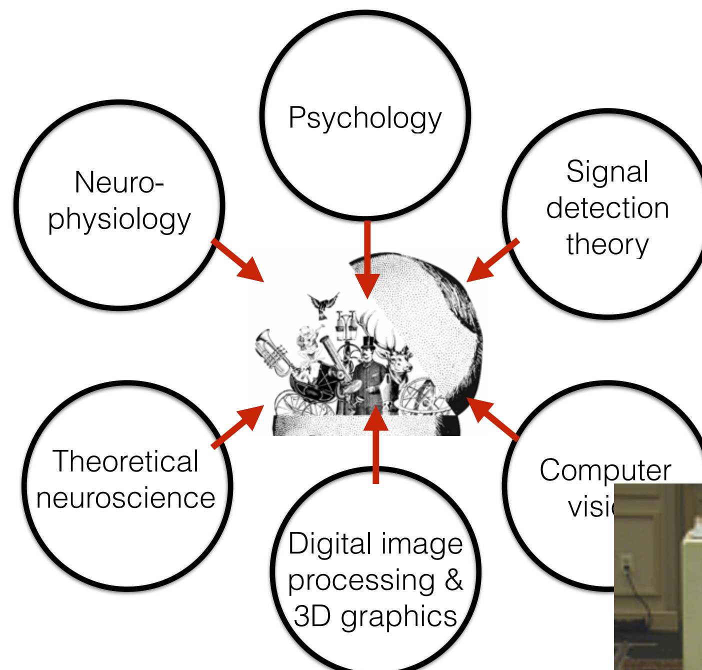

A Learning Algorithm for Boltzmann Machines*

1985

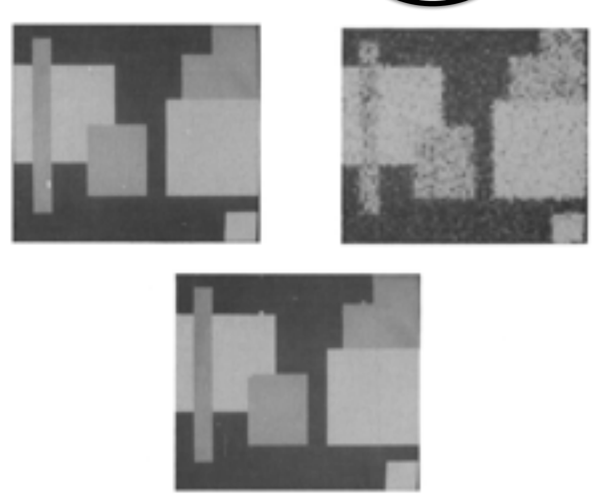

1984

1983

The “state of affairs” ~ 1985

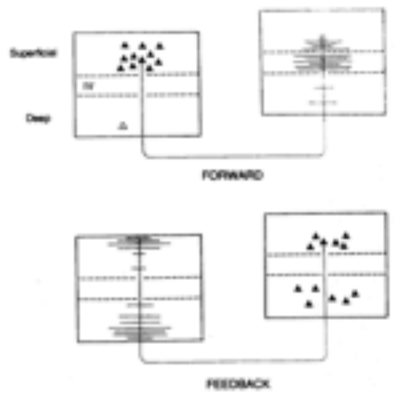

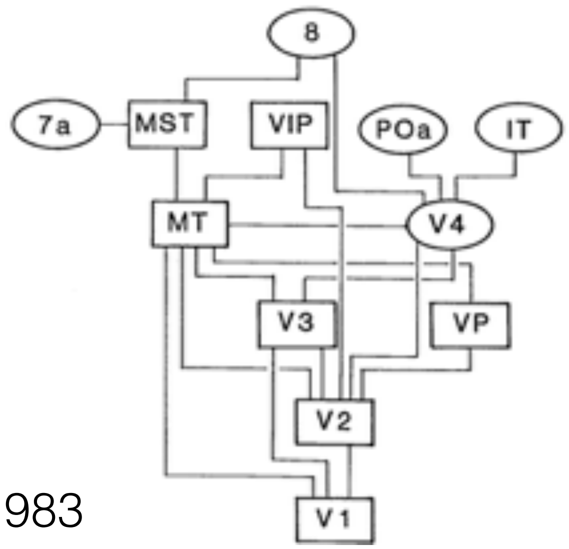

1983

ideal observer analysis

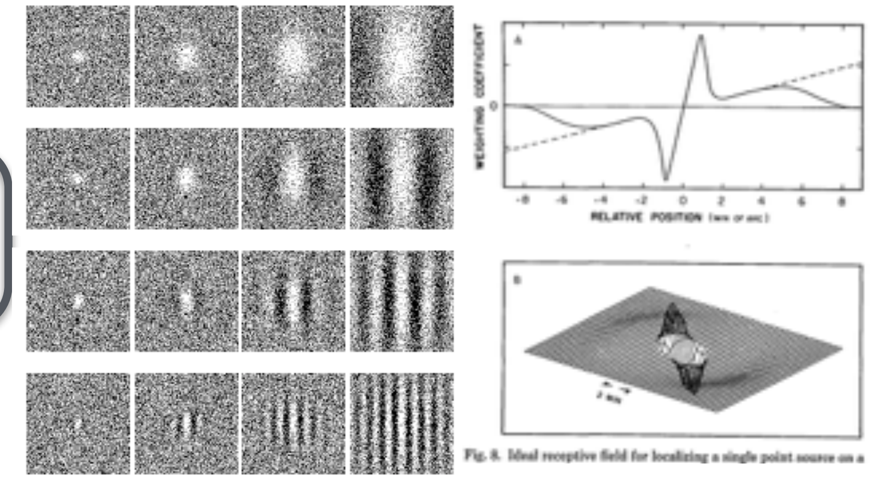

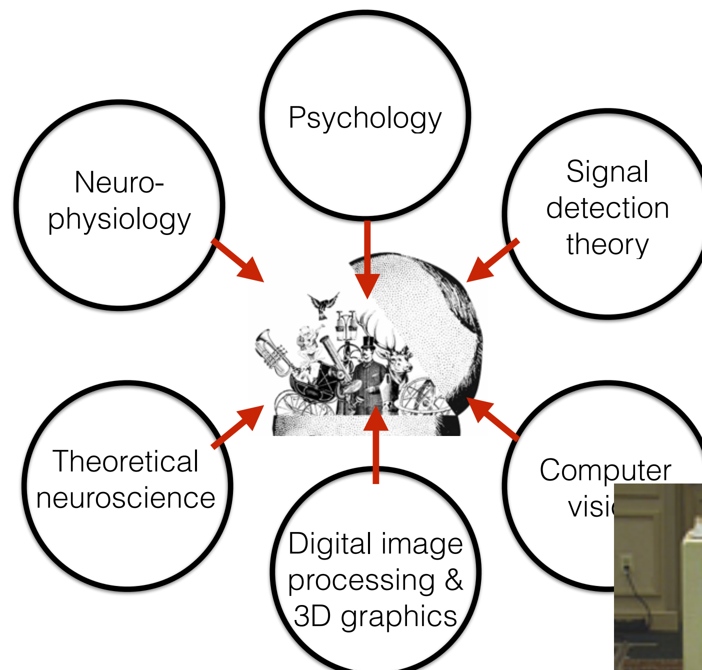

A Learning Algorithm for Boltzmann Machines*

1985

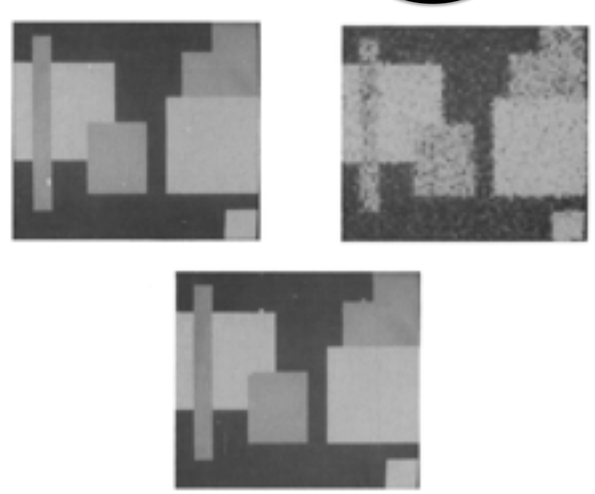

1984

1983

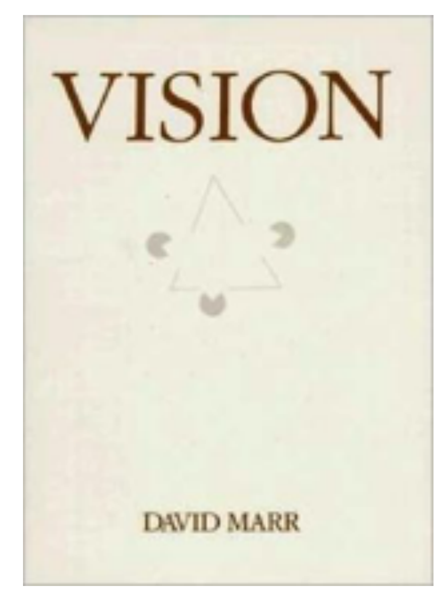

1983

The “state of affairs” ~ 1985

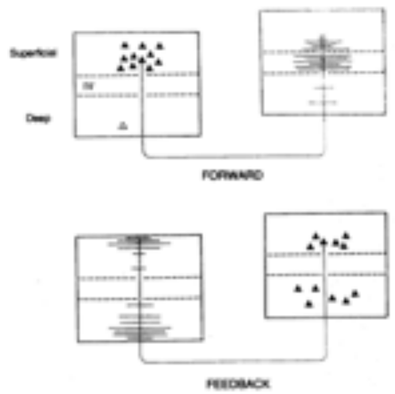

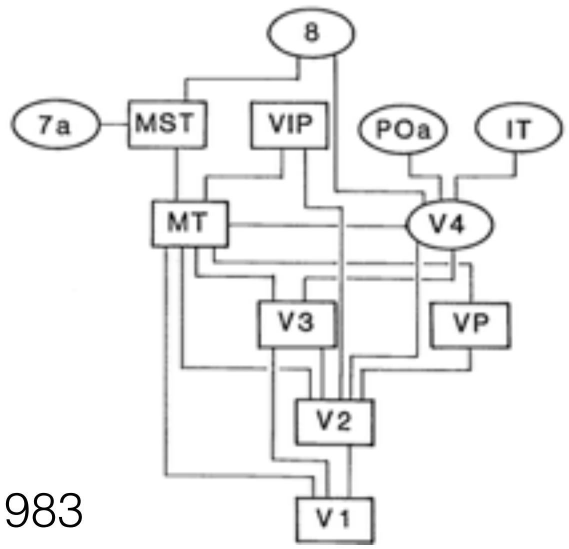

1983

ideal observer analysis

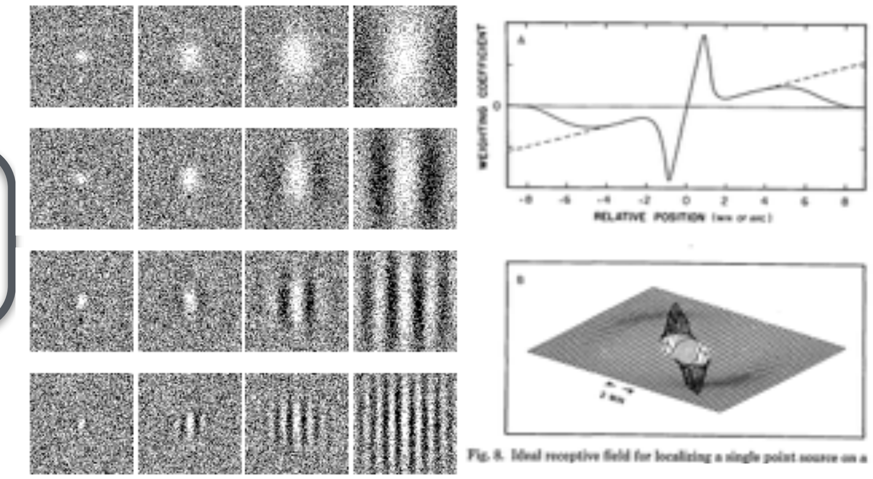

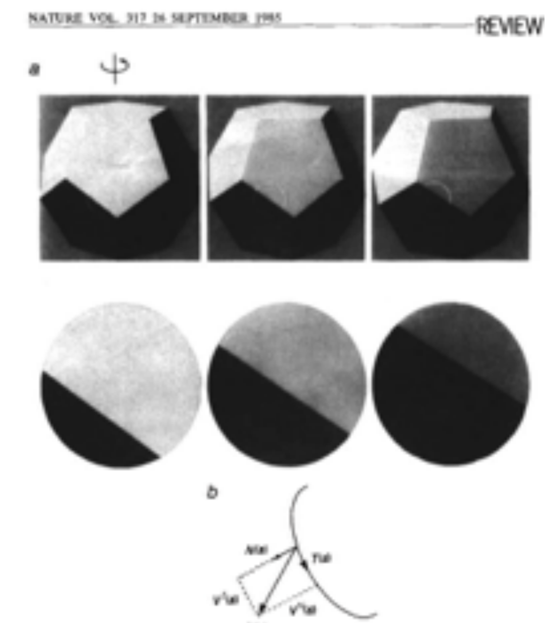

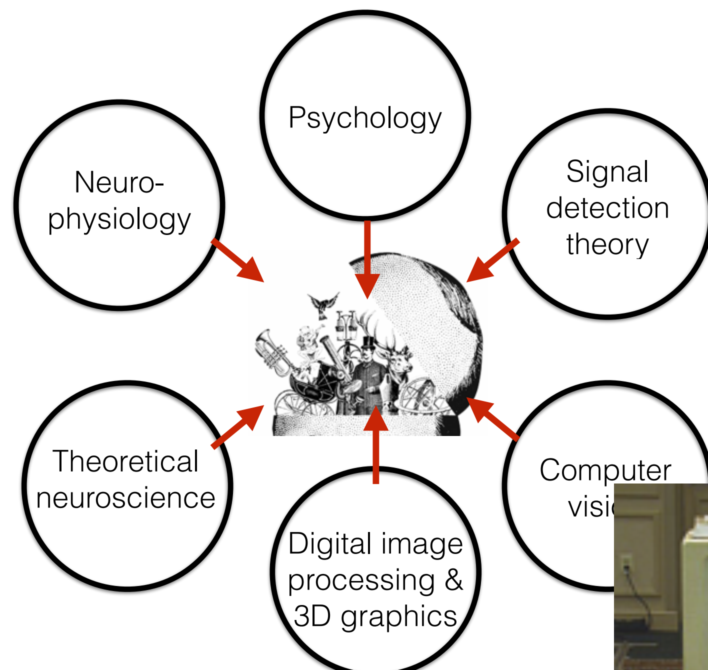

A Learning Algorithm for Boltzmann Machines*

1985

$$\| Az - y \|^2 + \lambda \| Pz \|^2$$

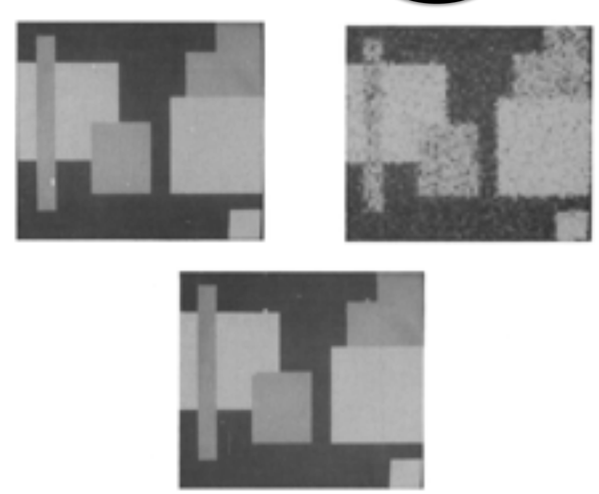

1984

1983

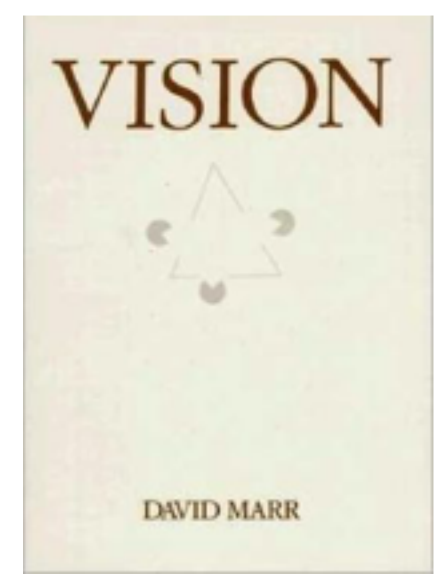

1983

What was “Bayesian vision” back then?

- mathematics of inference given uncertainty
- common language to integrate disciplines
- tools to model image and scene regularities

What was “Bayesian vision” back then?

- mathematics of inference given uncertainty
- common language to integrate disciplines
- tools to model image and scene regularities

…beginning to hint as a set of conceptual and analytical tools to understand how humans infer causes (scenes, objects) from data (images)

In those “early days”, it was strongly motivated by the idea of perception as inverse optics

forward optics

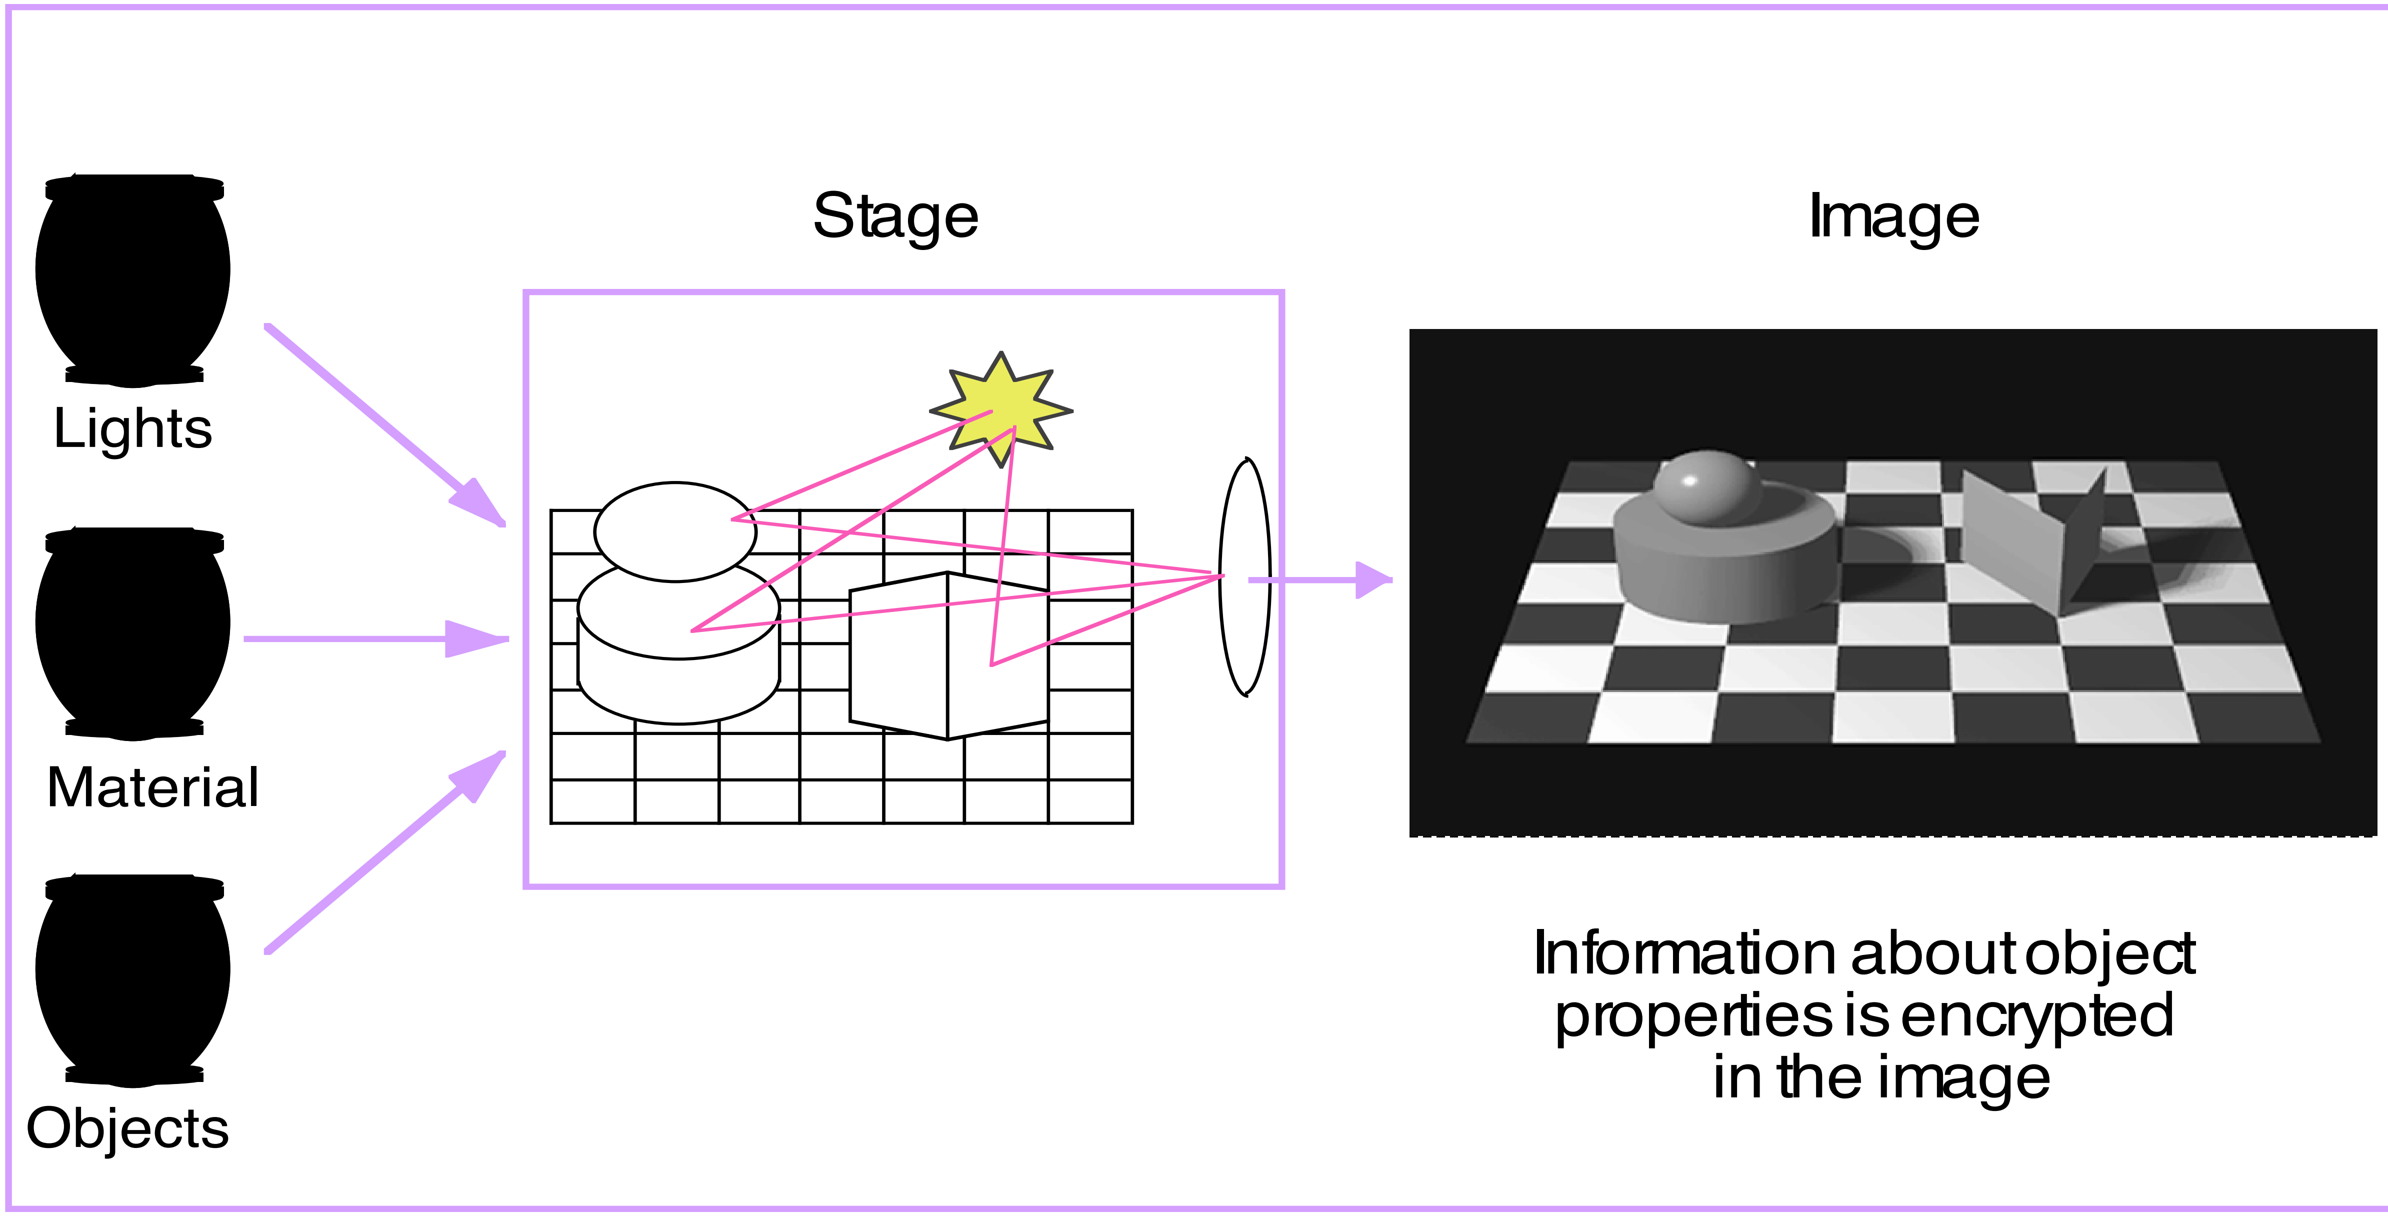

$p(S)$ prior

$p(I \mid S)$ likelihood

$p(I)$

the *generative* components

uncertainties

Given a small intensity patch, what caused it in the scene?

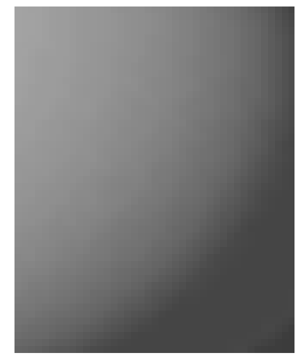

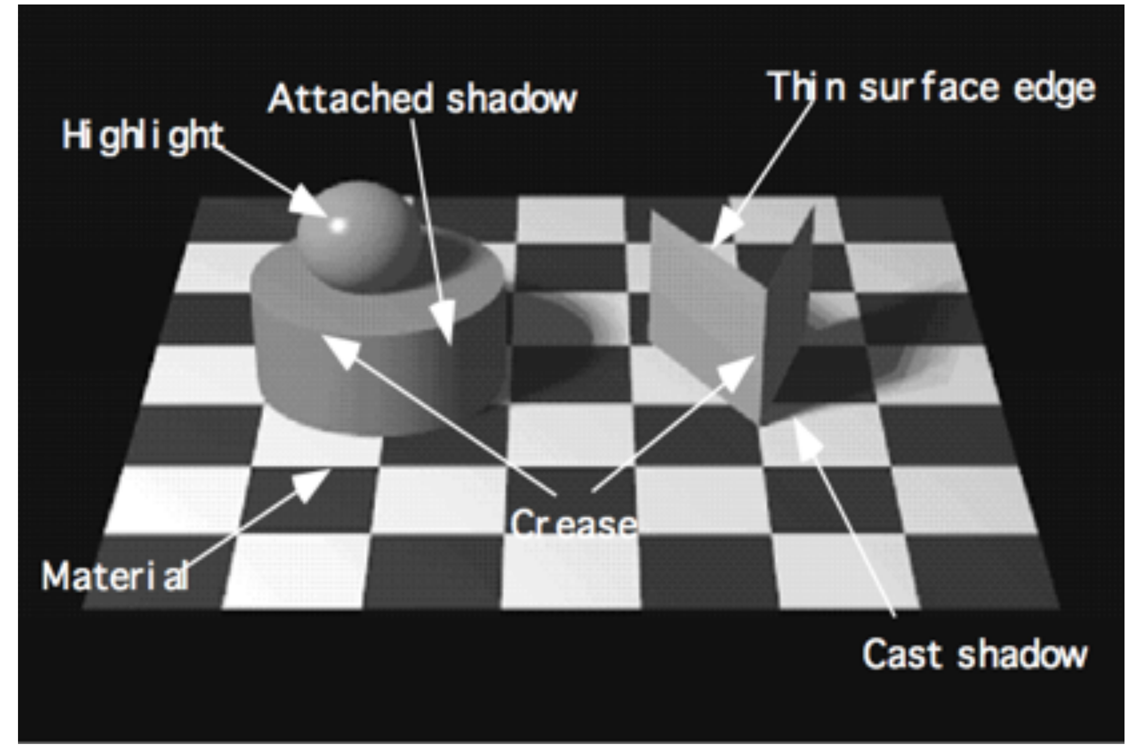

Given a 2D image, which 3D shape?

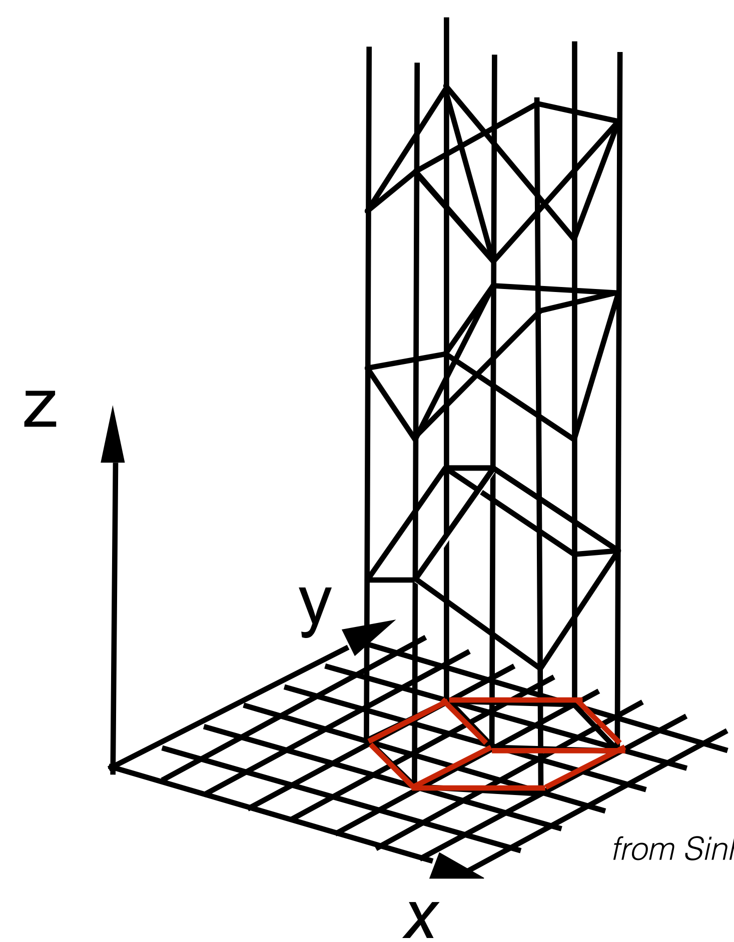

from Sinha & Adelson

vision as inverse optics

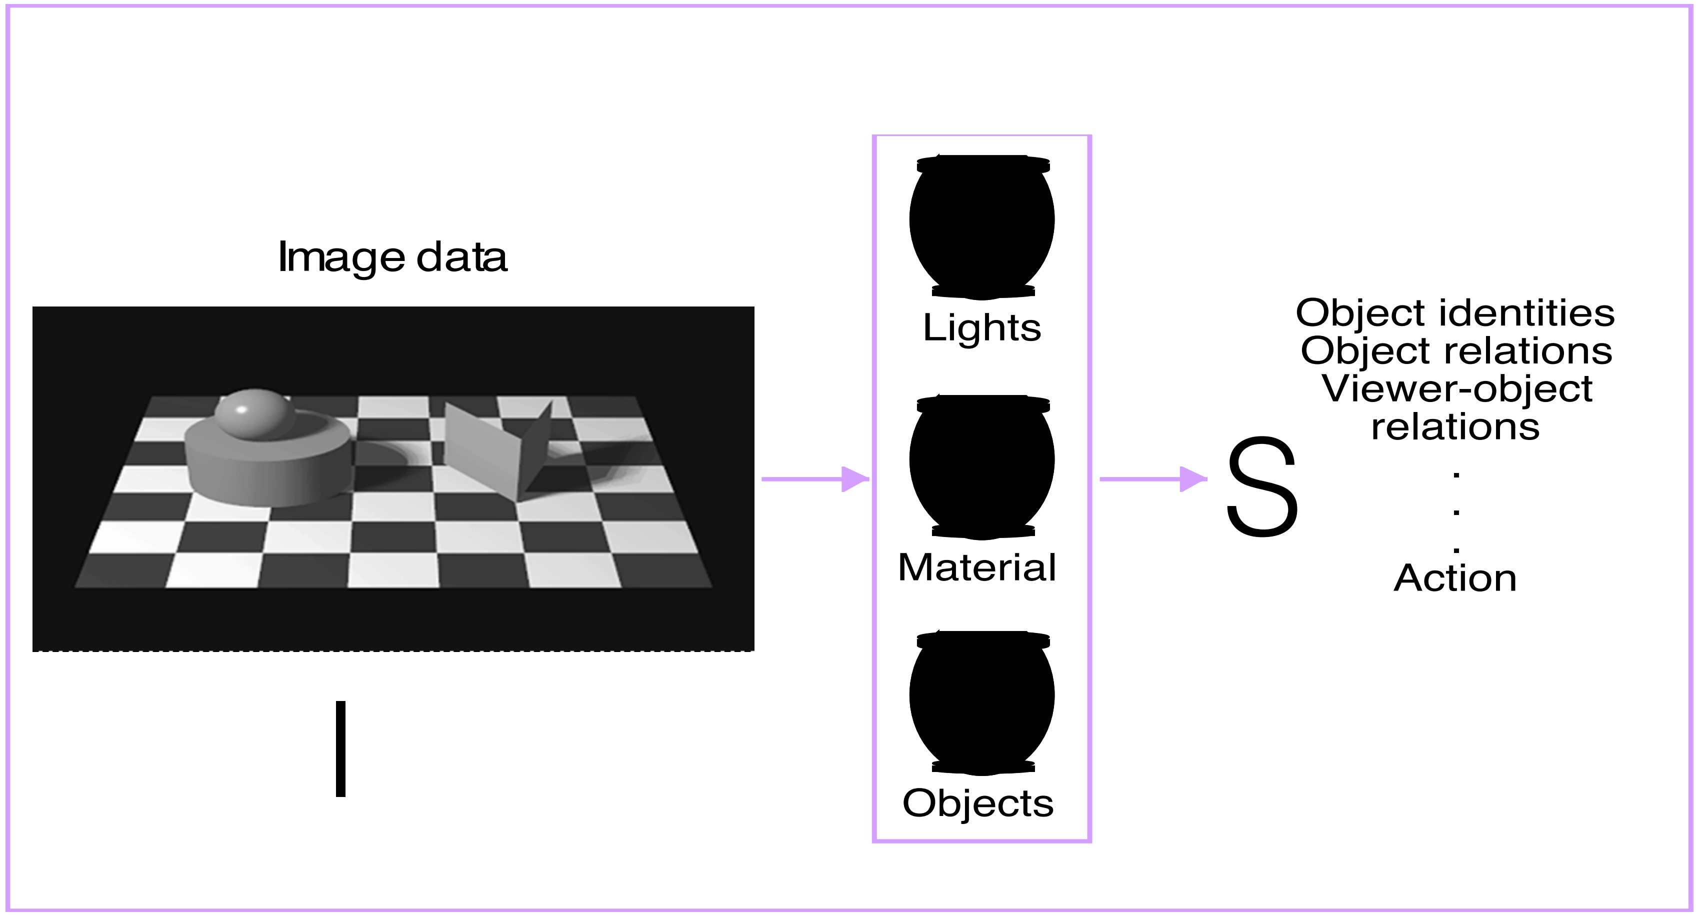

Given image pattern I, what combinations
of lights, material, object properties (S) caused it?

vision as inverse optics

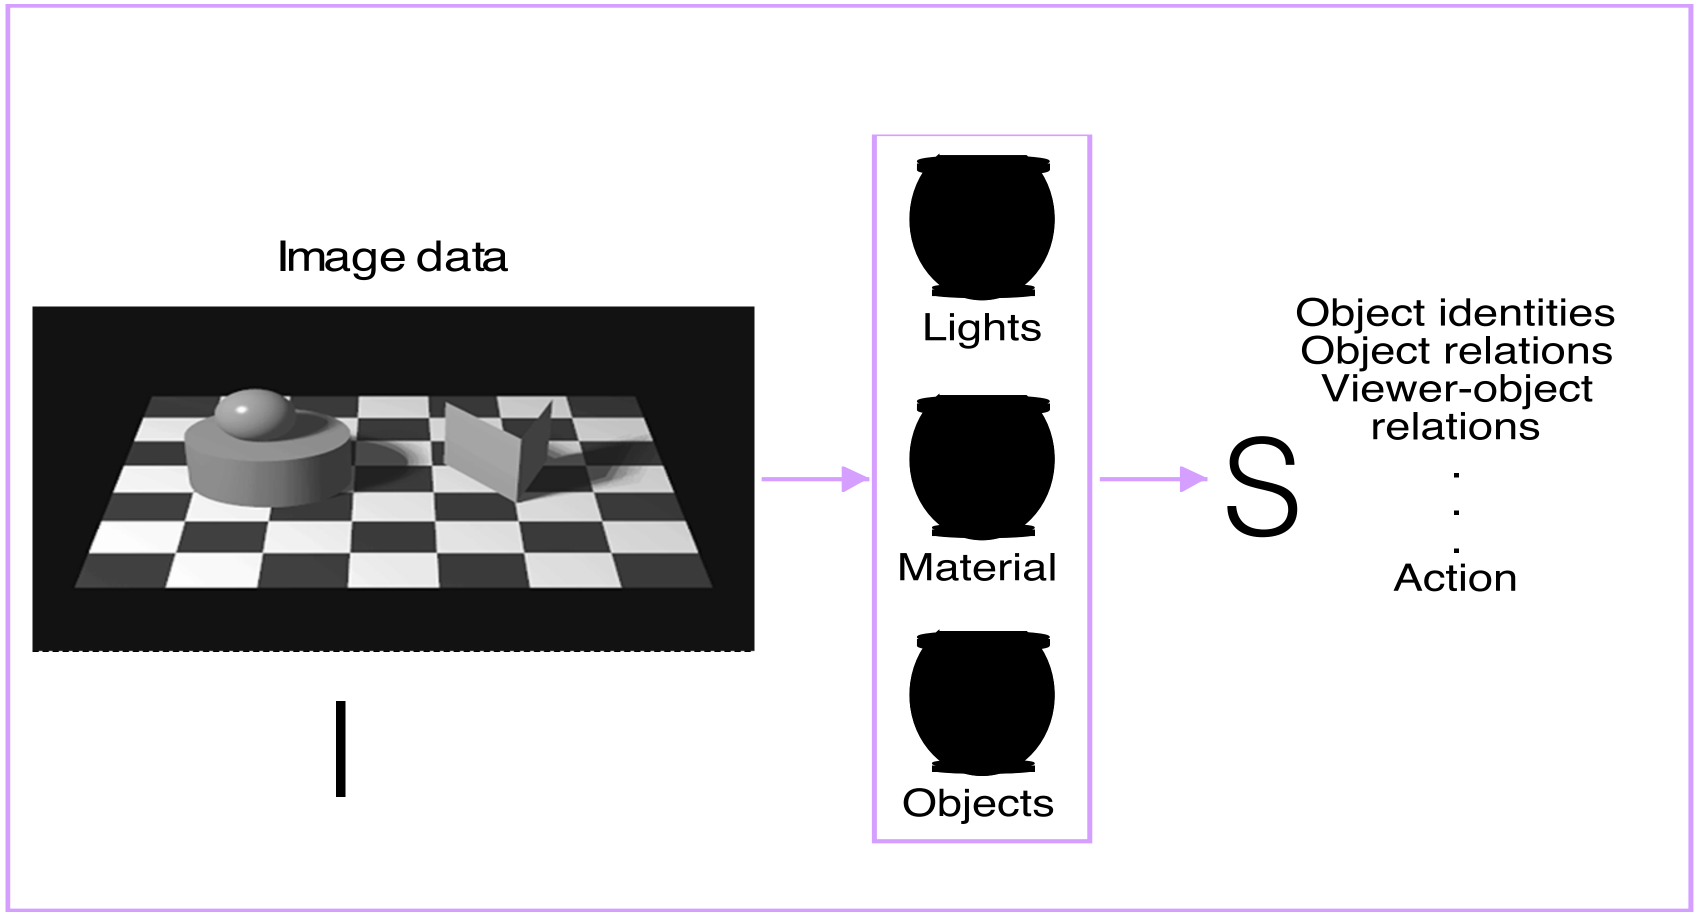

Given image pattern I, what combinations
of lights, material, object properties (S) caused it?

Bayes theorem $$p(S \mid I) \propto p(I \mid S)\ p(S)$$

likelihood modeled using forward optics

prior assumed or measured properties of scenes

Brown University 1985-1990











shape from shading

1986-87: Dave, student "glue"

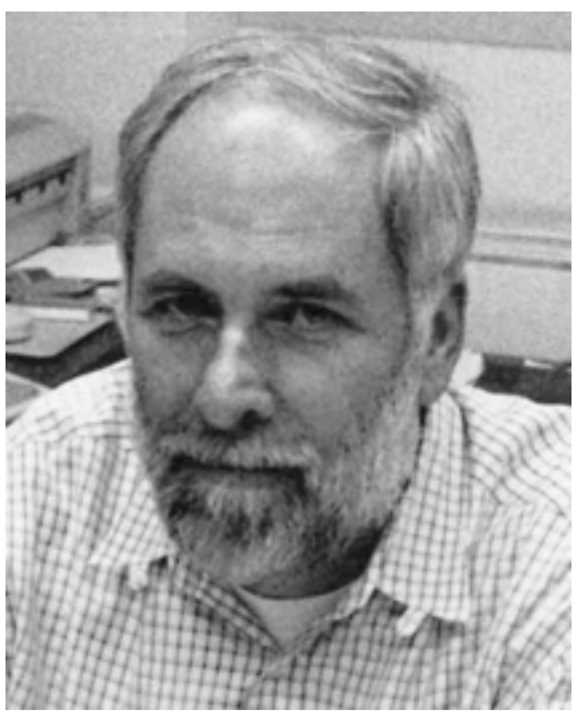

Jim Anderson

Dave

Dan Kersten

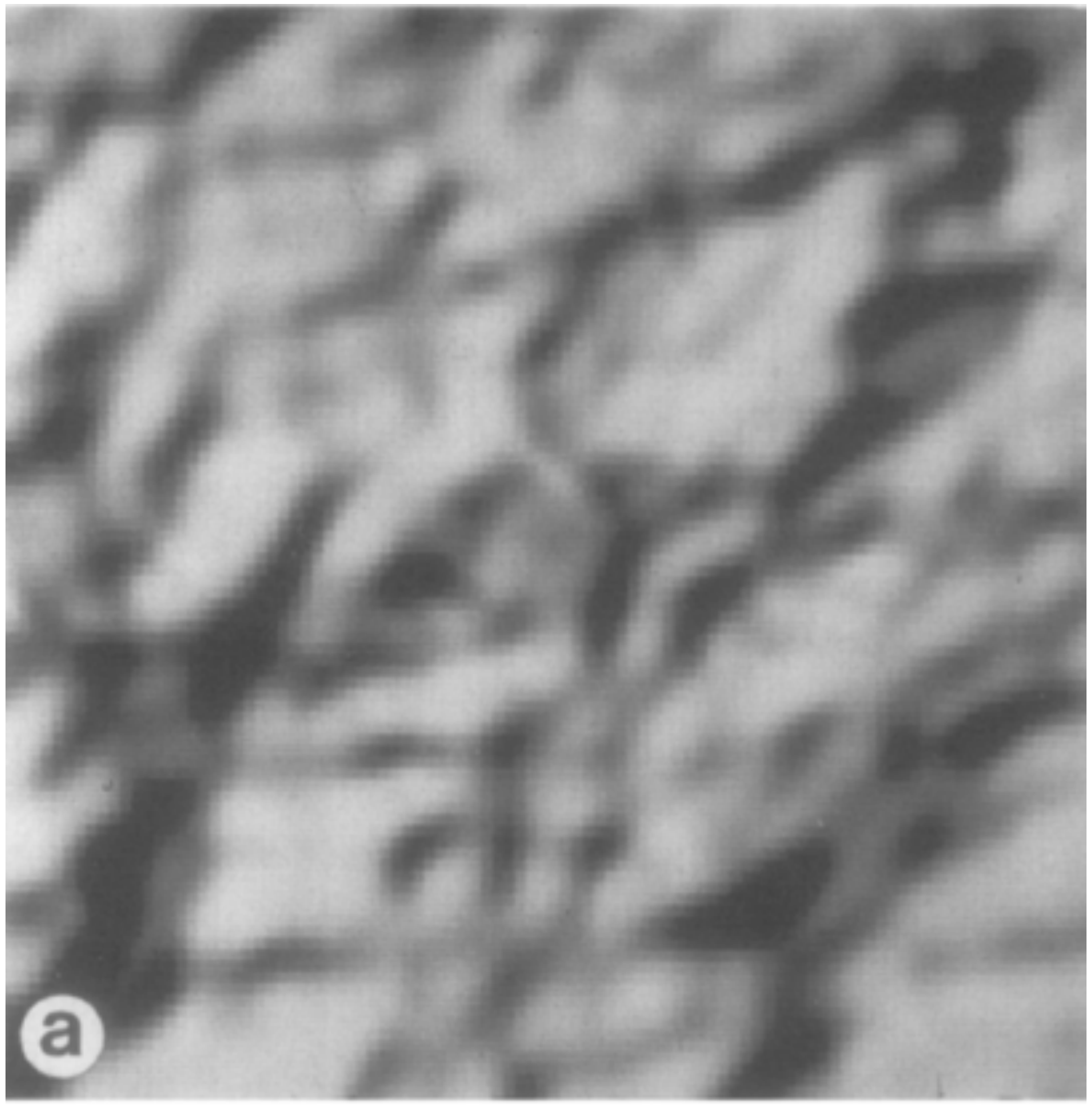

image intensities

the “shape from shading” problem

perceived geometry?

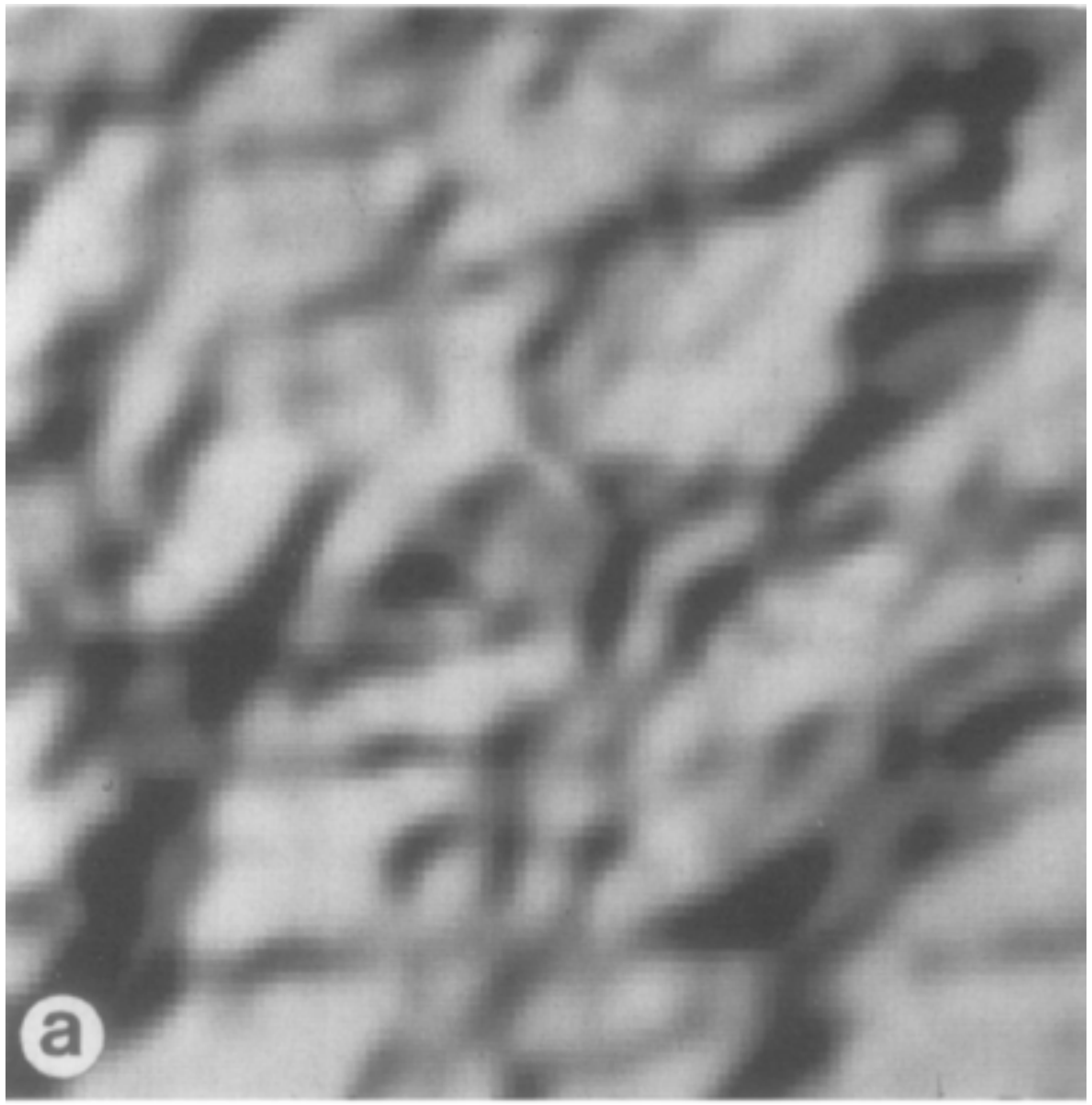

image intensities

the "shape from shading" problem

→ perceived geometry?

$$p(S \mid I) \propto p(I \mid S)\, p(S)$$

lambertian,
single light source,
fixed view

fractal

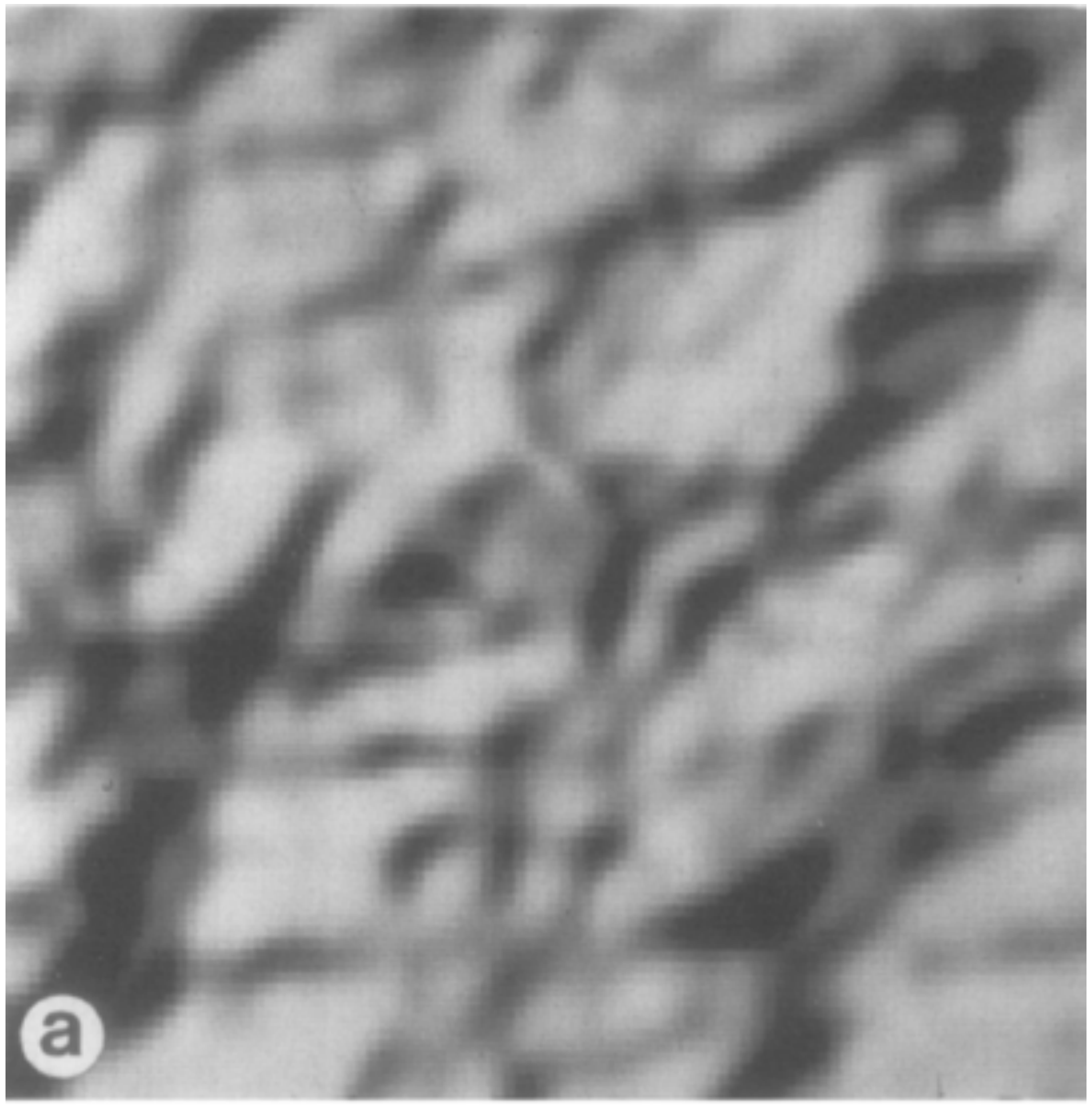

the “shape from shading” problem

perceived geometry?

image intensities

exploit the generative aspect of Bayes

$$p(S \mid I) \propto \boxed{p(I \mid S)\, p(S)}$$

lambertian,
single light source,
fixed view

fractal

generate lots of surfaces in 3D

render the surfaces to make lots of images

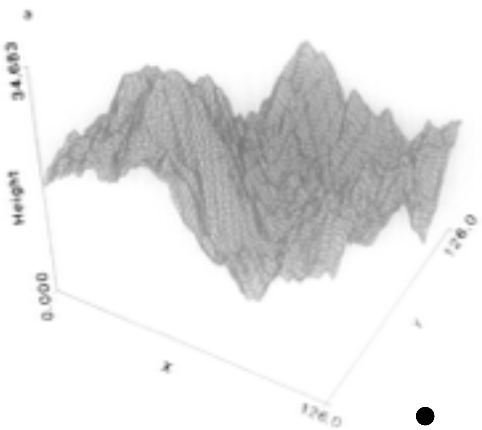

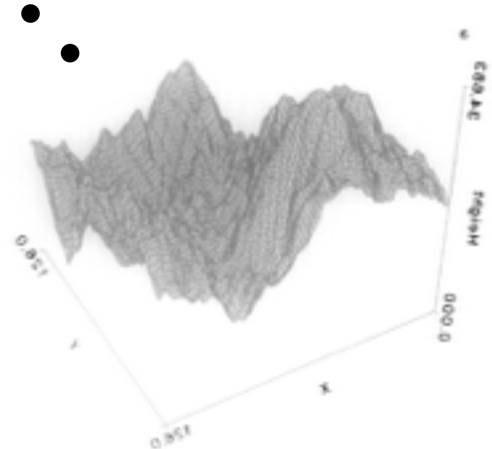

fractal prior

3D to 2D projection

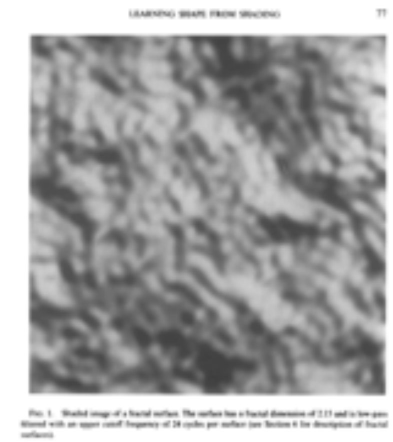

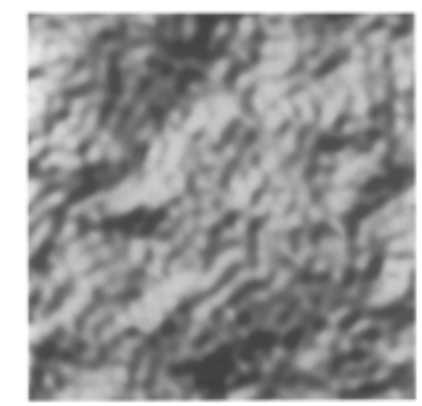

Use supervised learning to construct an estimator for 3D surface shapes

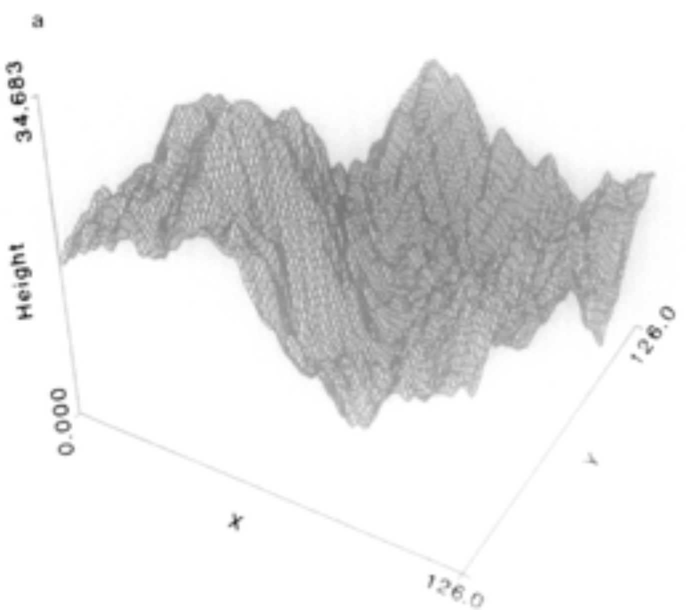

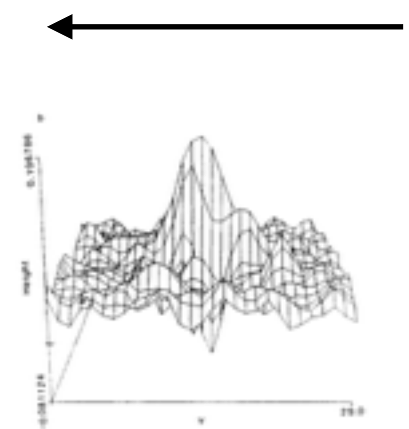

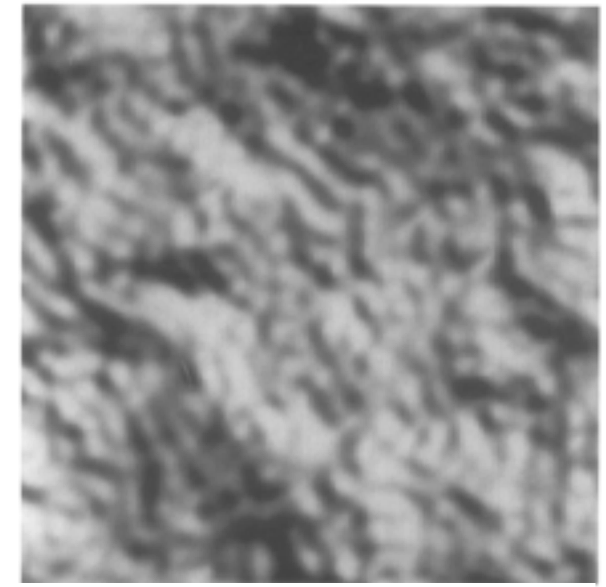

Knill, D. C., & Kersten, D. (1990). Learning a near-optimal estimator for surface shape from shading. *Computer Vision, Graphics, and Image Processing*, *50*(1), 75–100.

Use supervised learning to construct an estimator for 3D surface shapes

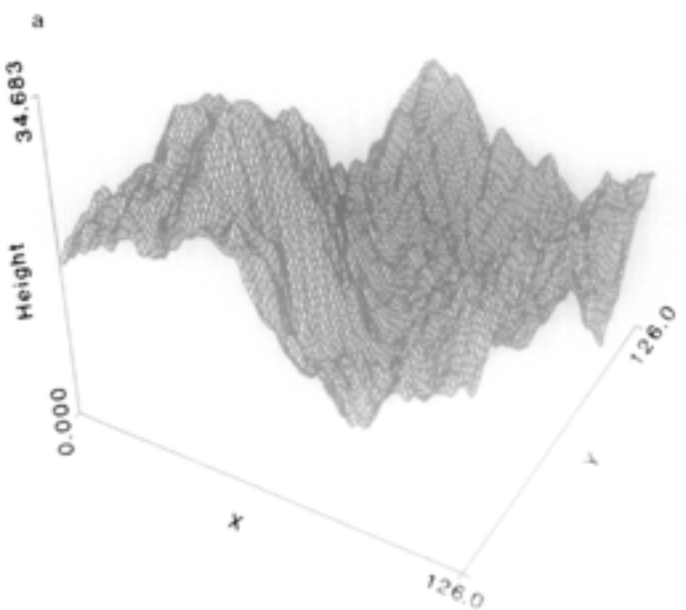

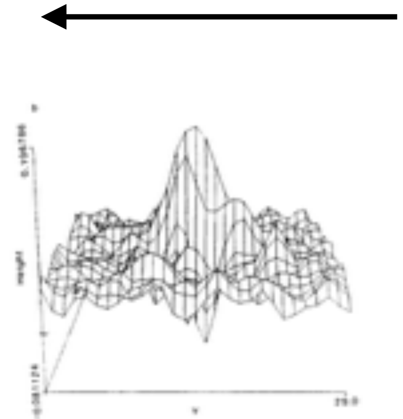

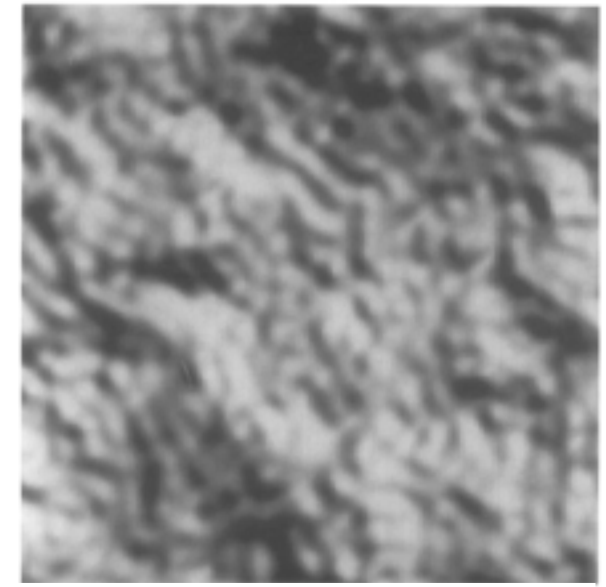

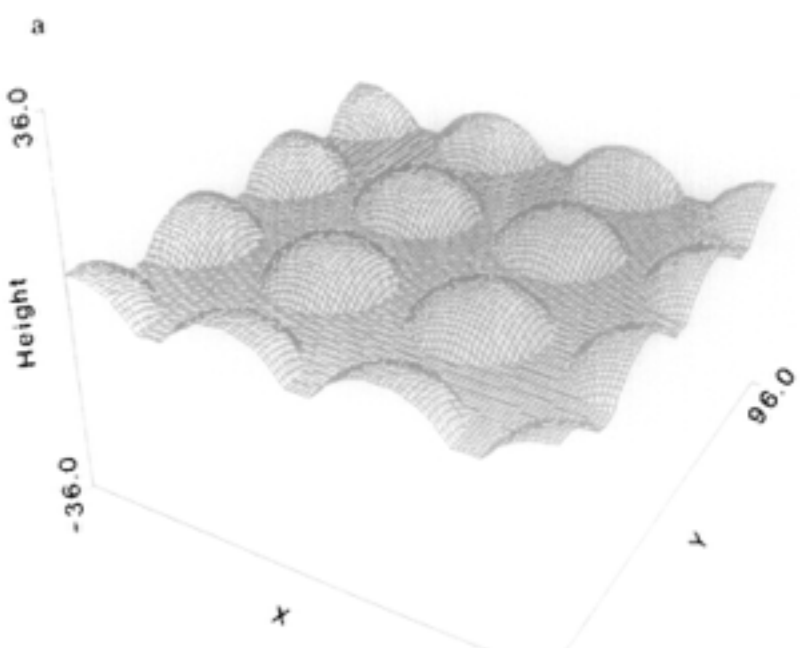

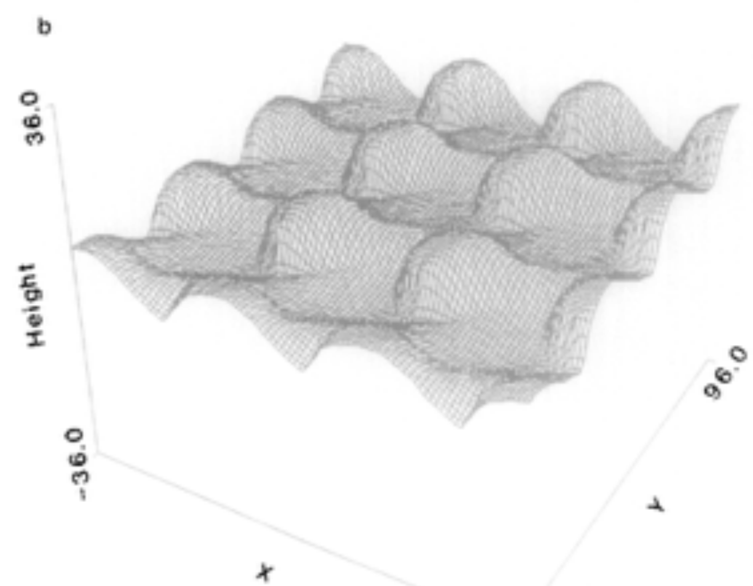

True shape

Estimated shape

Knill, D. C., & Kersten, D. (1990). Learning a near-optimal estimator for surface shape from shading. *Computer Vision, Graphics, and Image Processing*, *50*(1), 75–100.

an interesting method…but how to take this
forward to explain human perception of shape,
object properties?

perhaps there were bigger issues…

the causes of image patterns were more complicated
e.g. discontinuities are important, but causes of discontinuities
are not all the same

lightness as reflectance estimation

1988-91: Dave, precocious experimentalist and “closer”

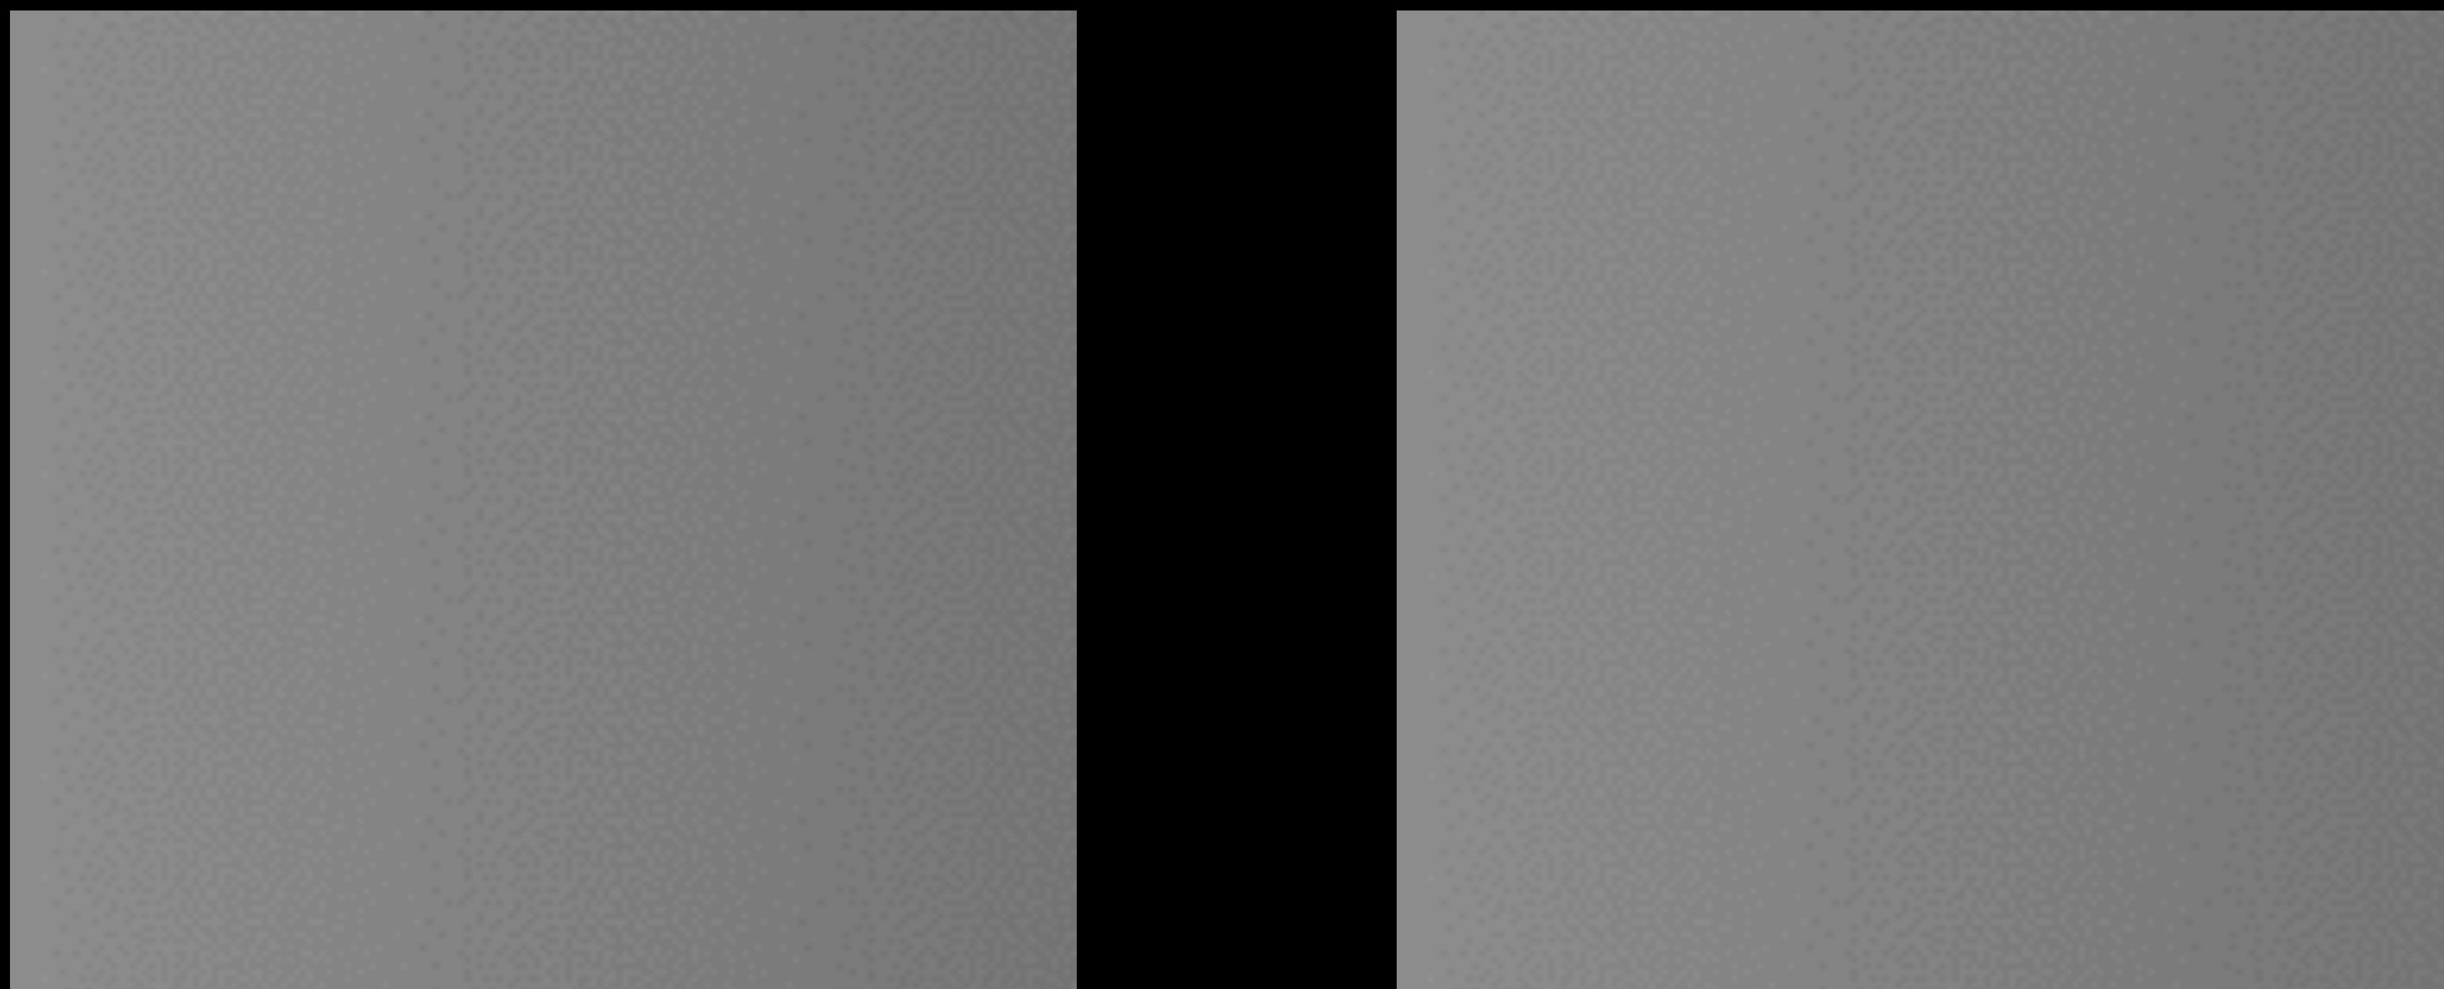

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*(1), 1–11.

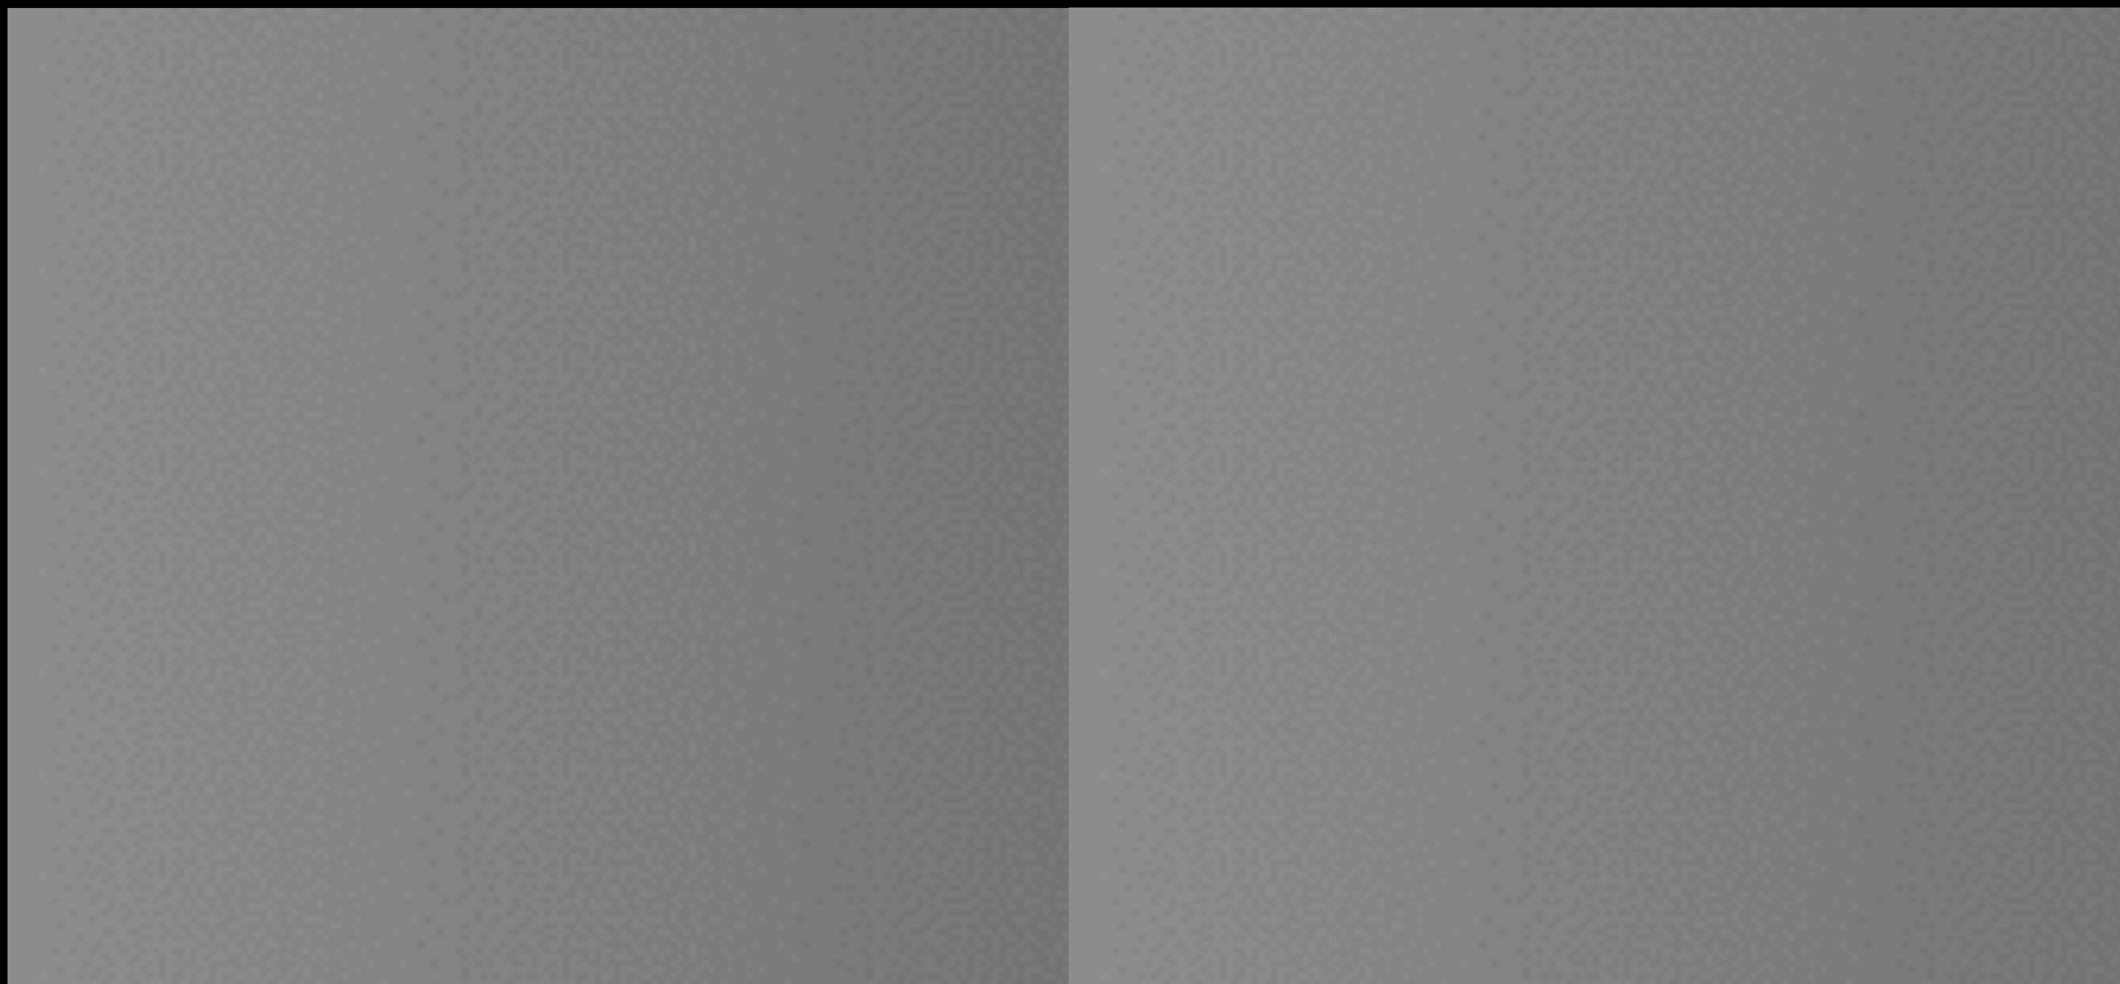

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*(1), 1–11.

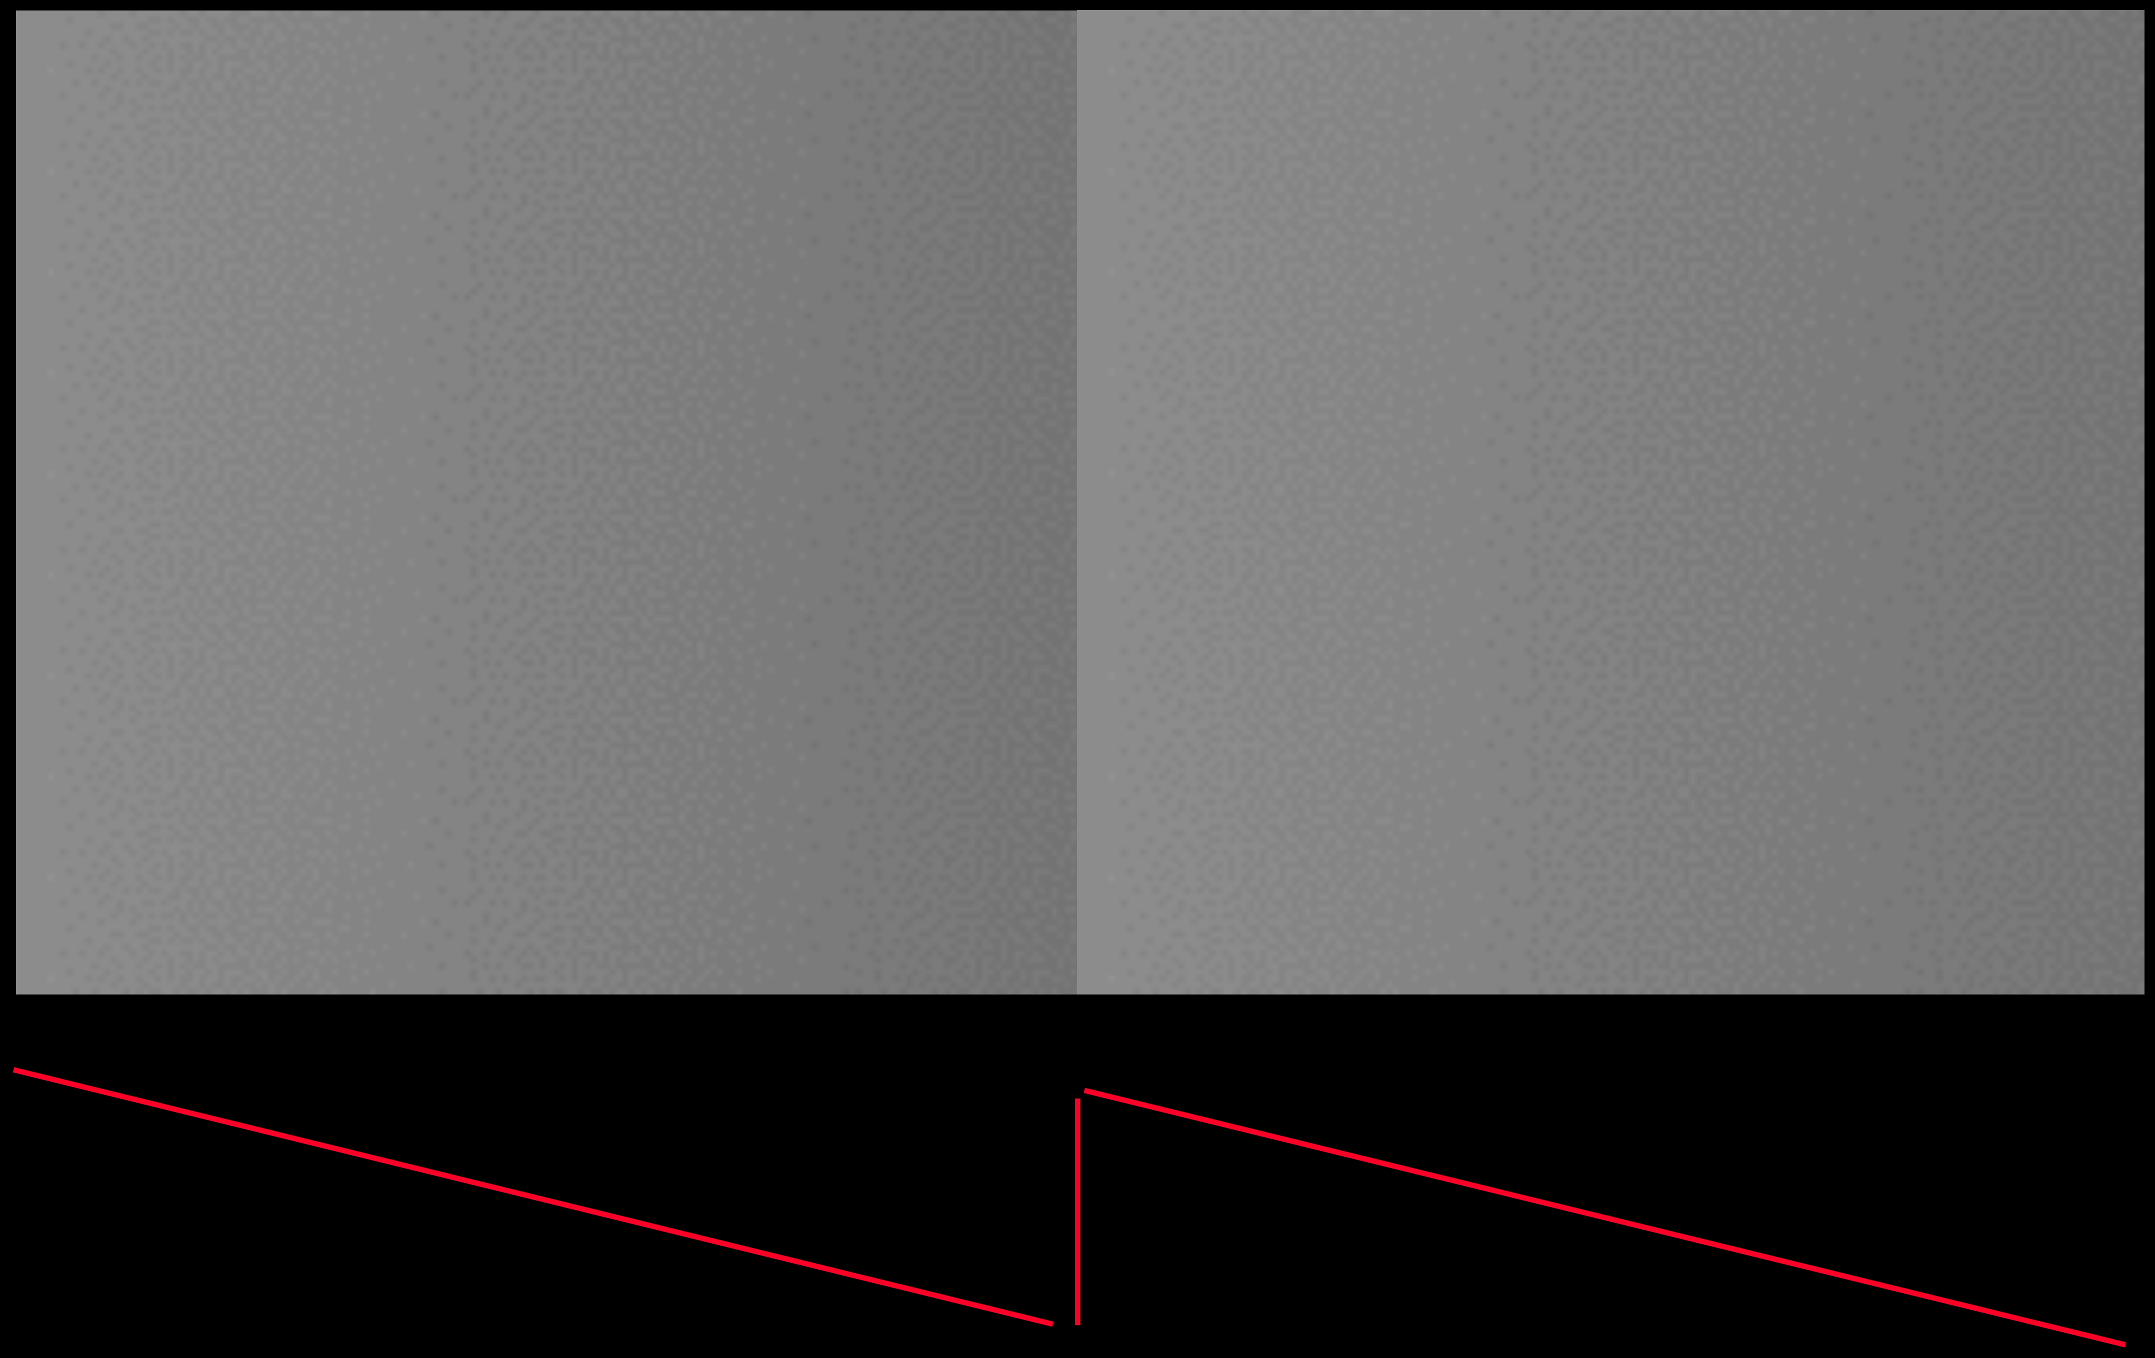

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*(1), 1–11.

mechanistic view

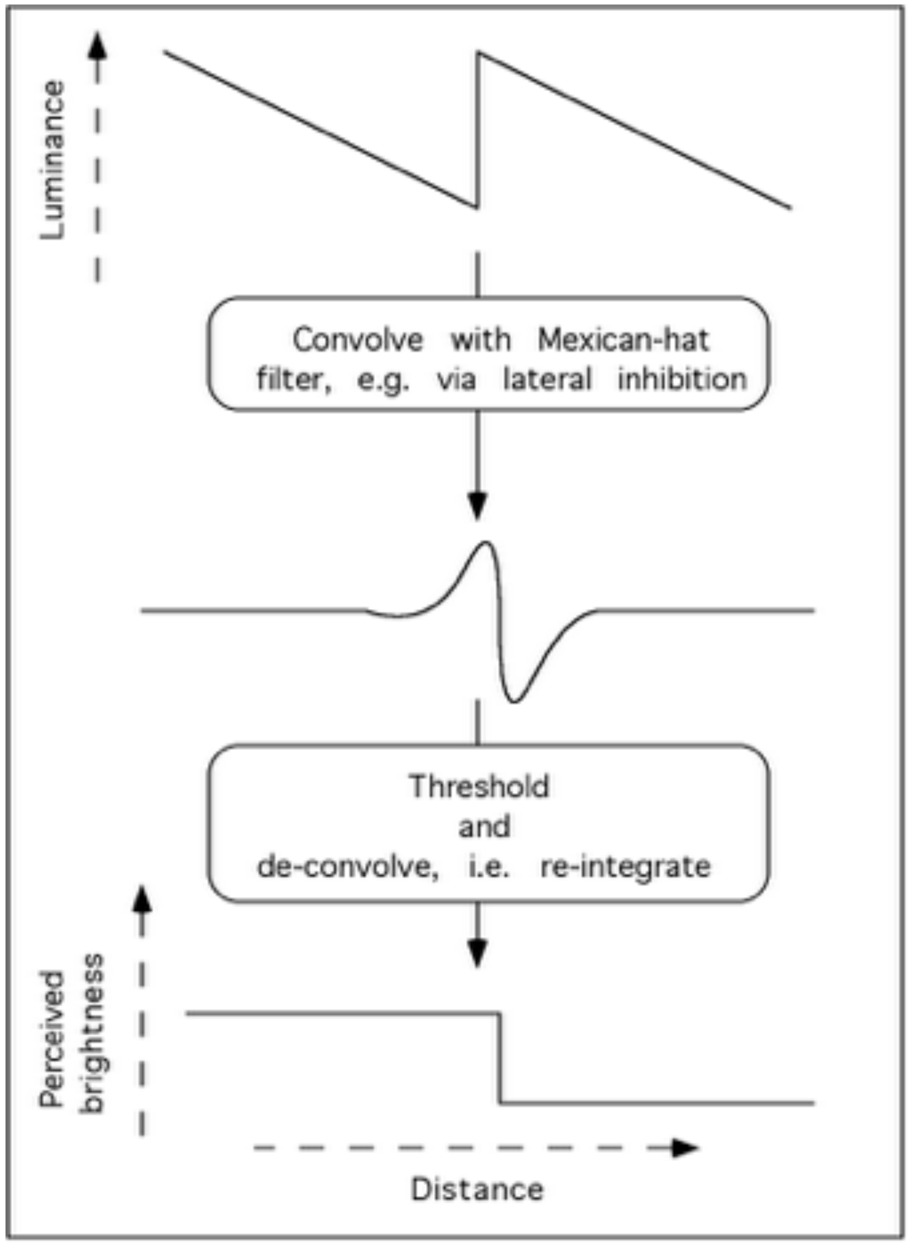

mechanistic view

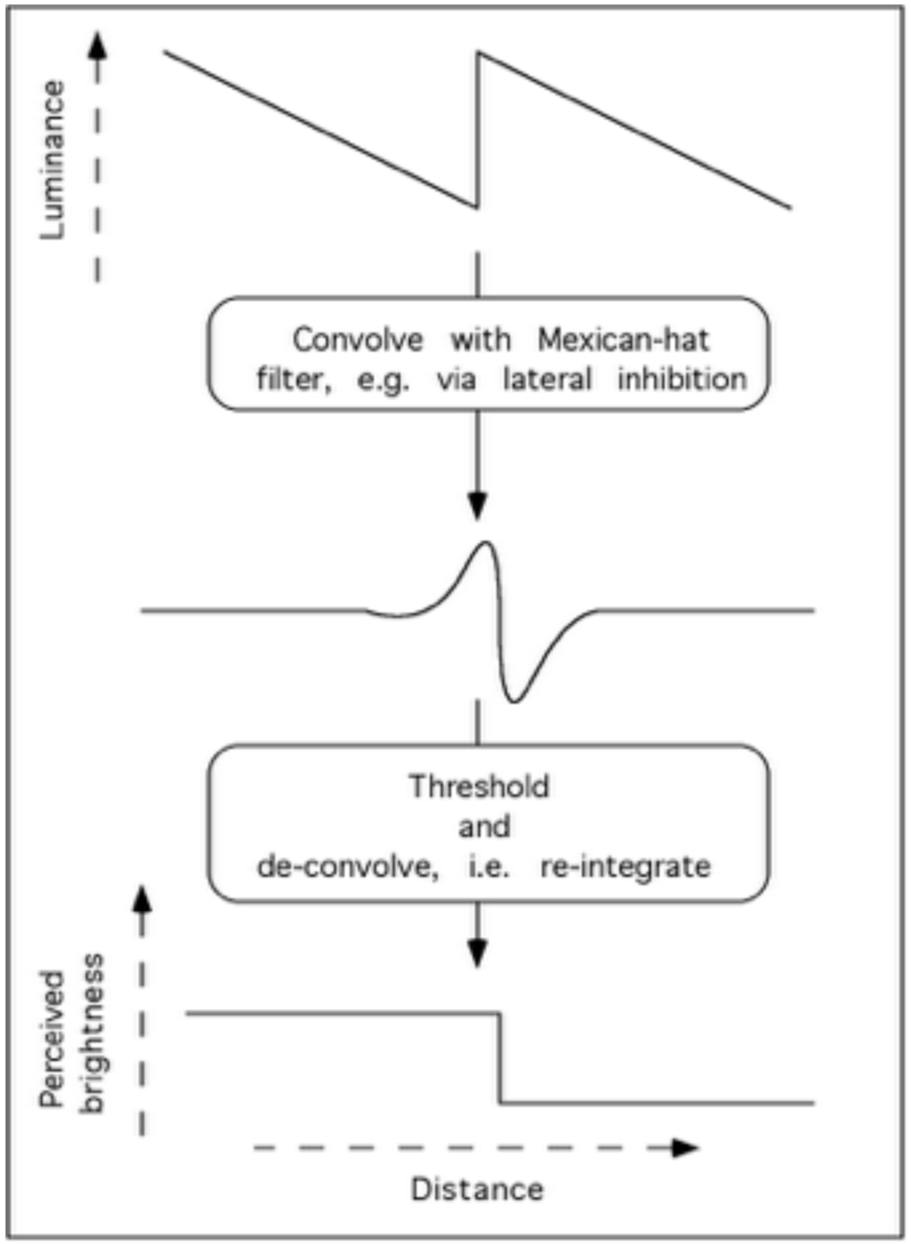

causal view

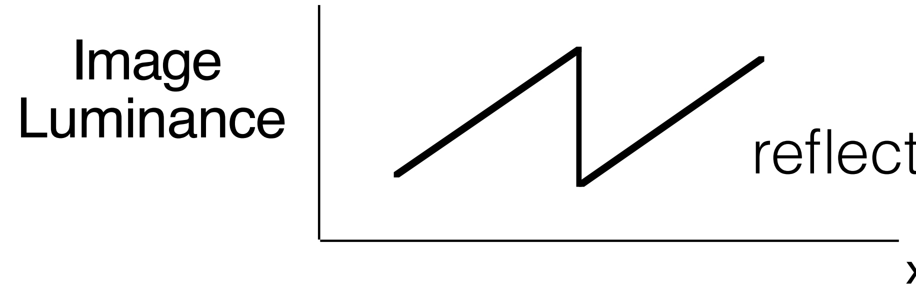

image = reflectance x illumination

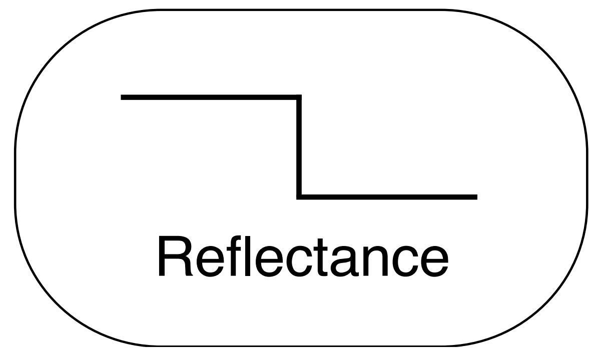

piece-wise constant prior

X

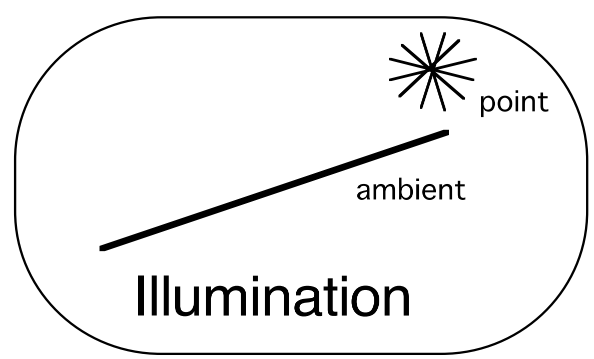

spatially smooth prior

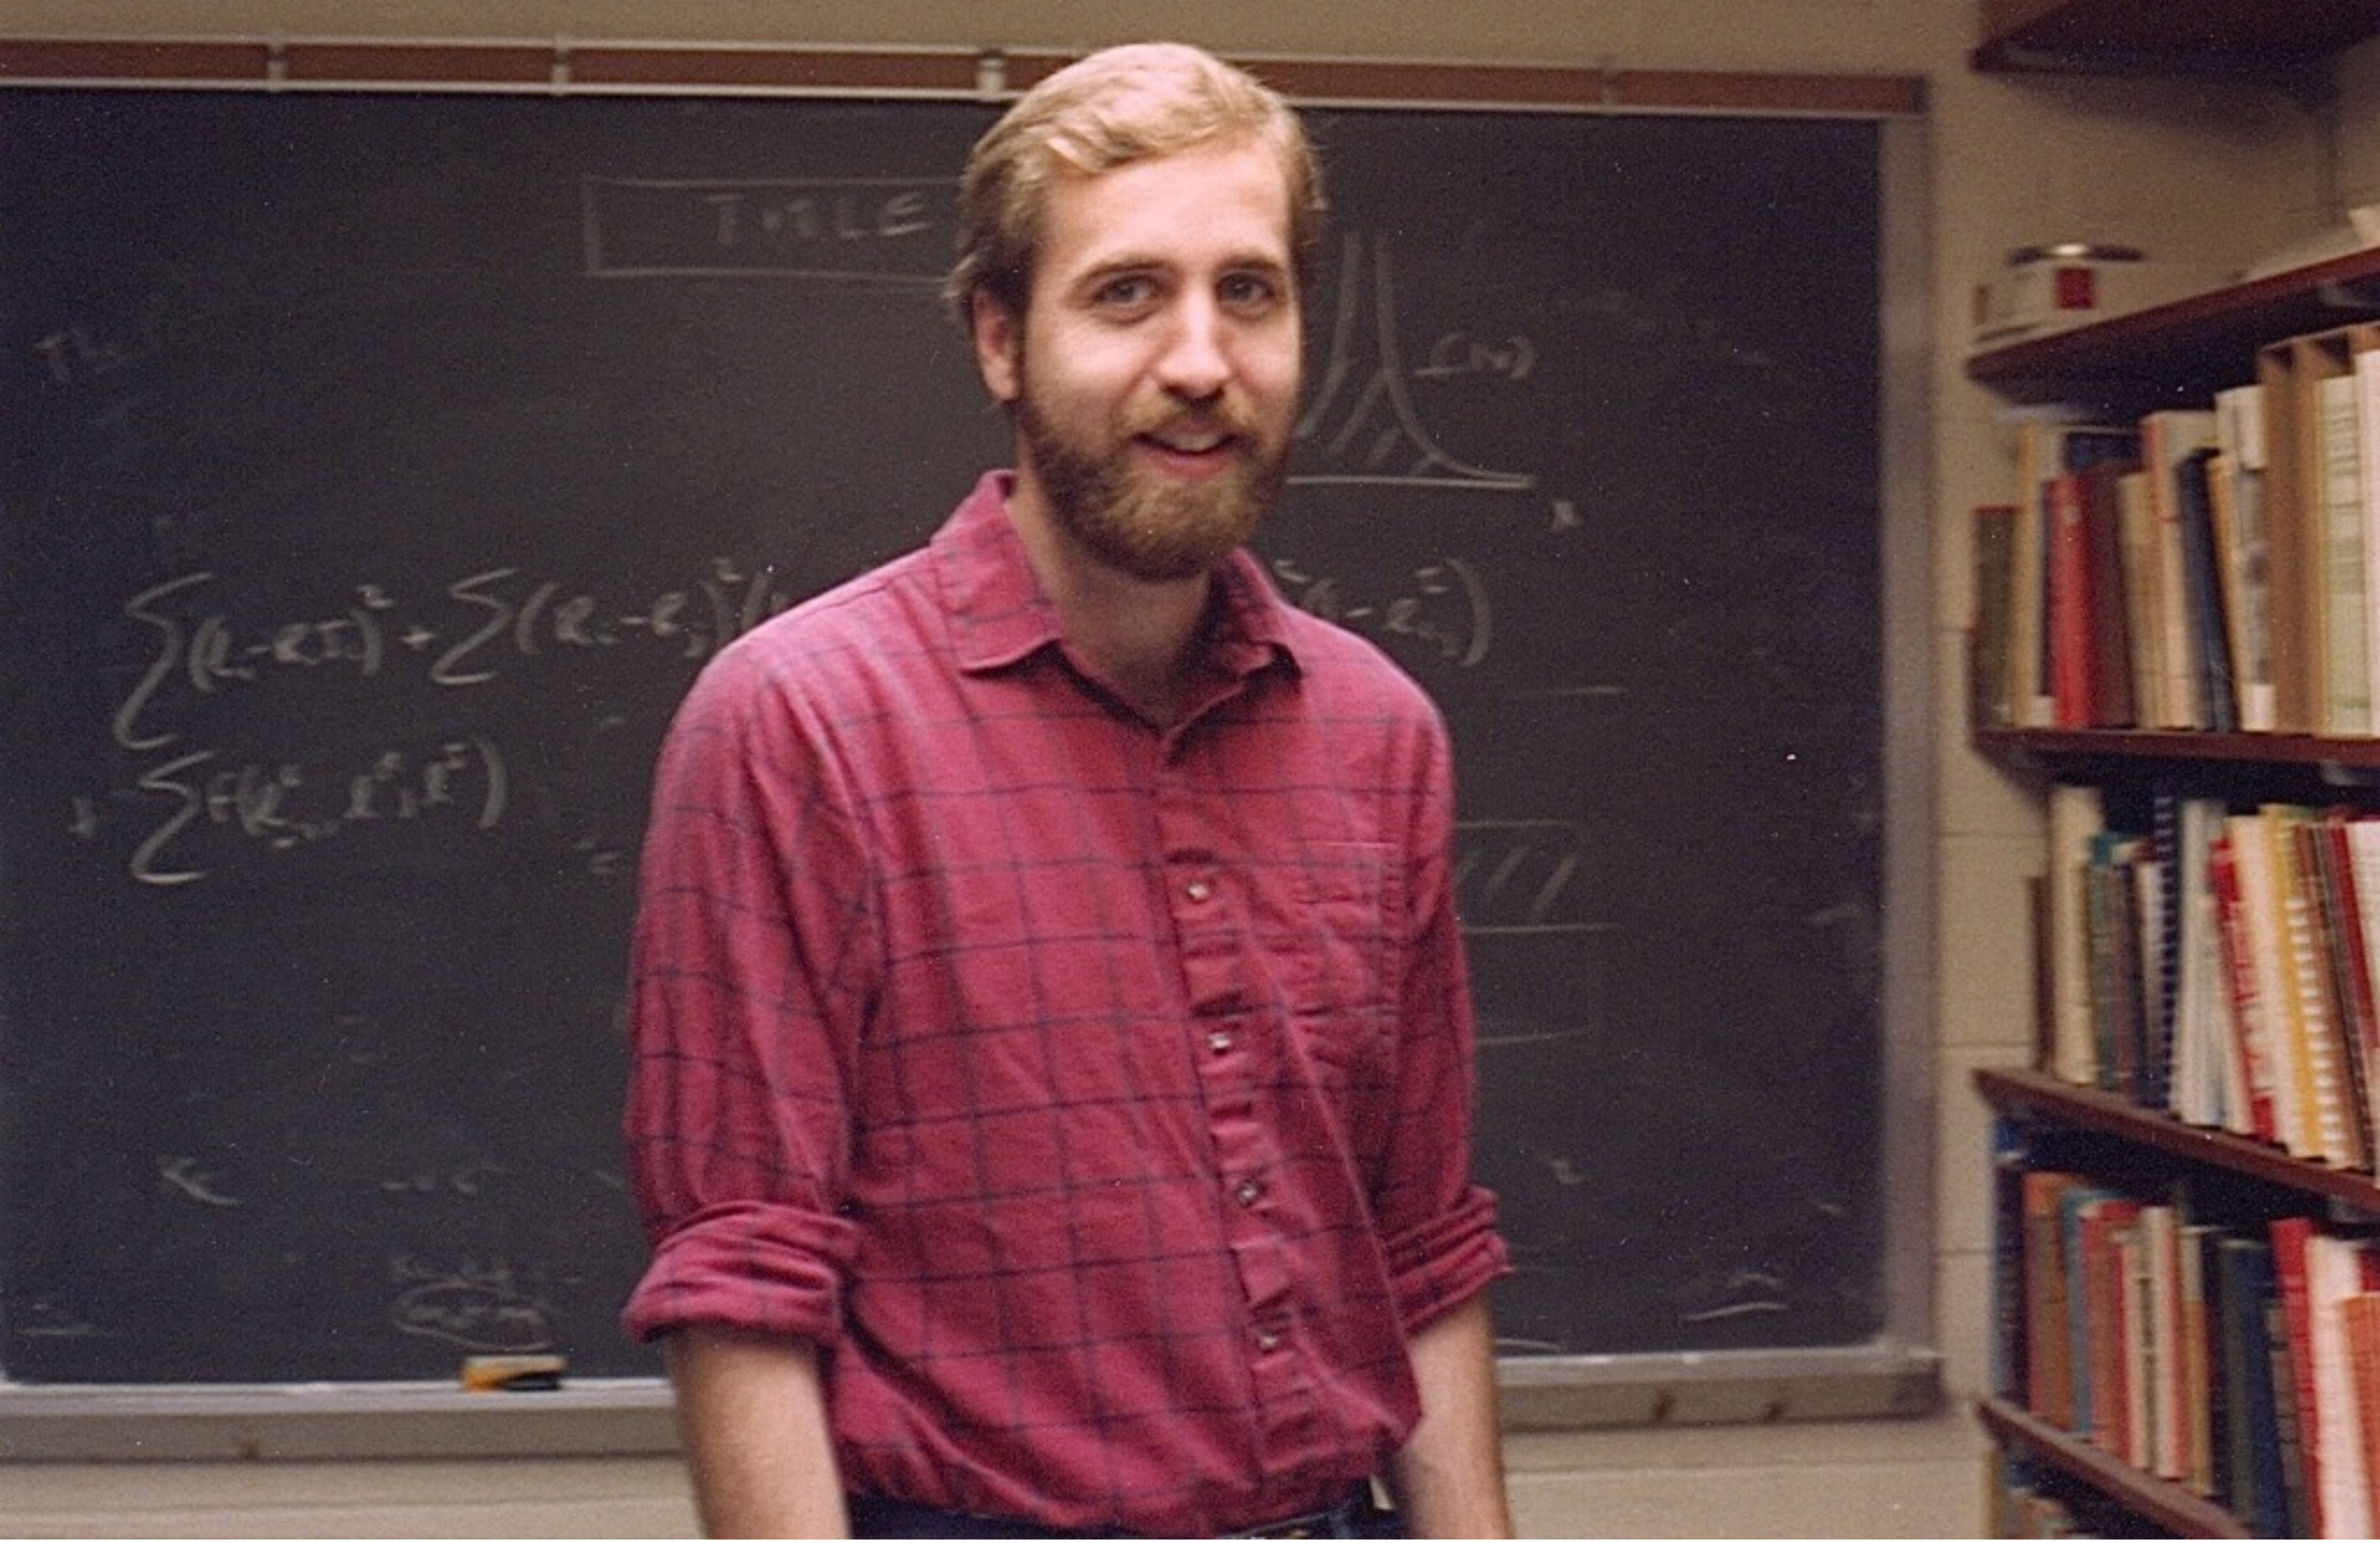

Dave Knill, Brown University, ca. 1986

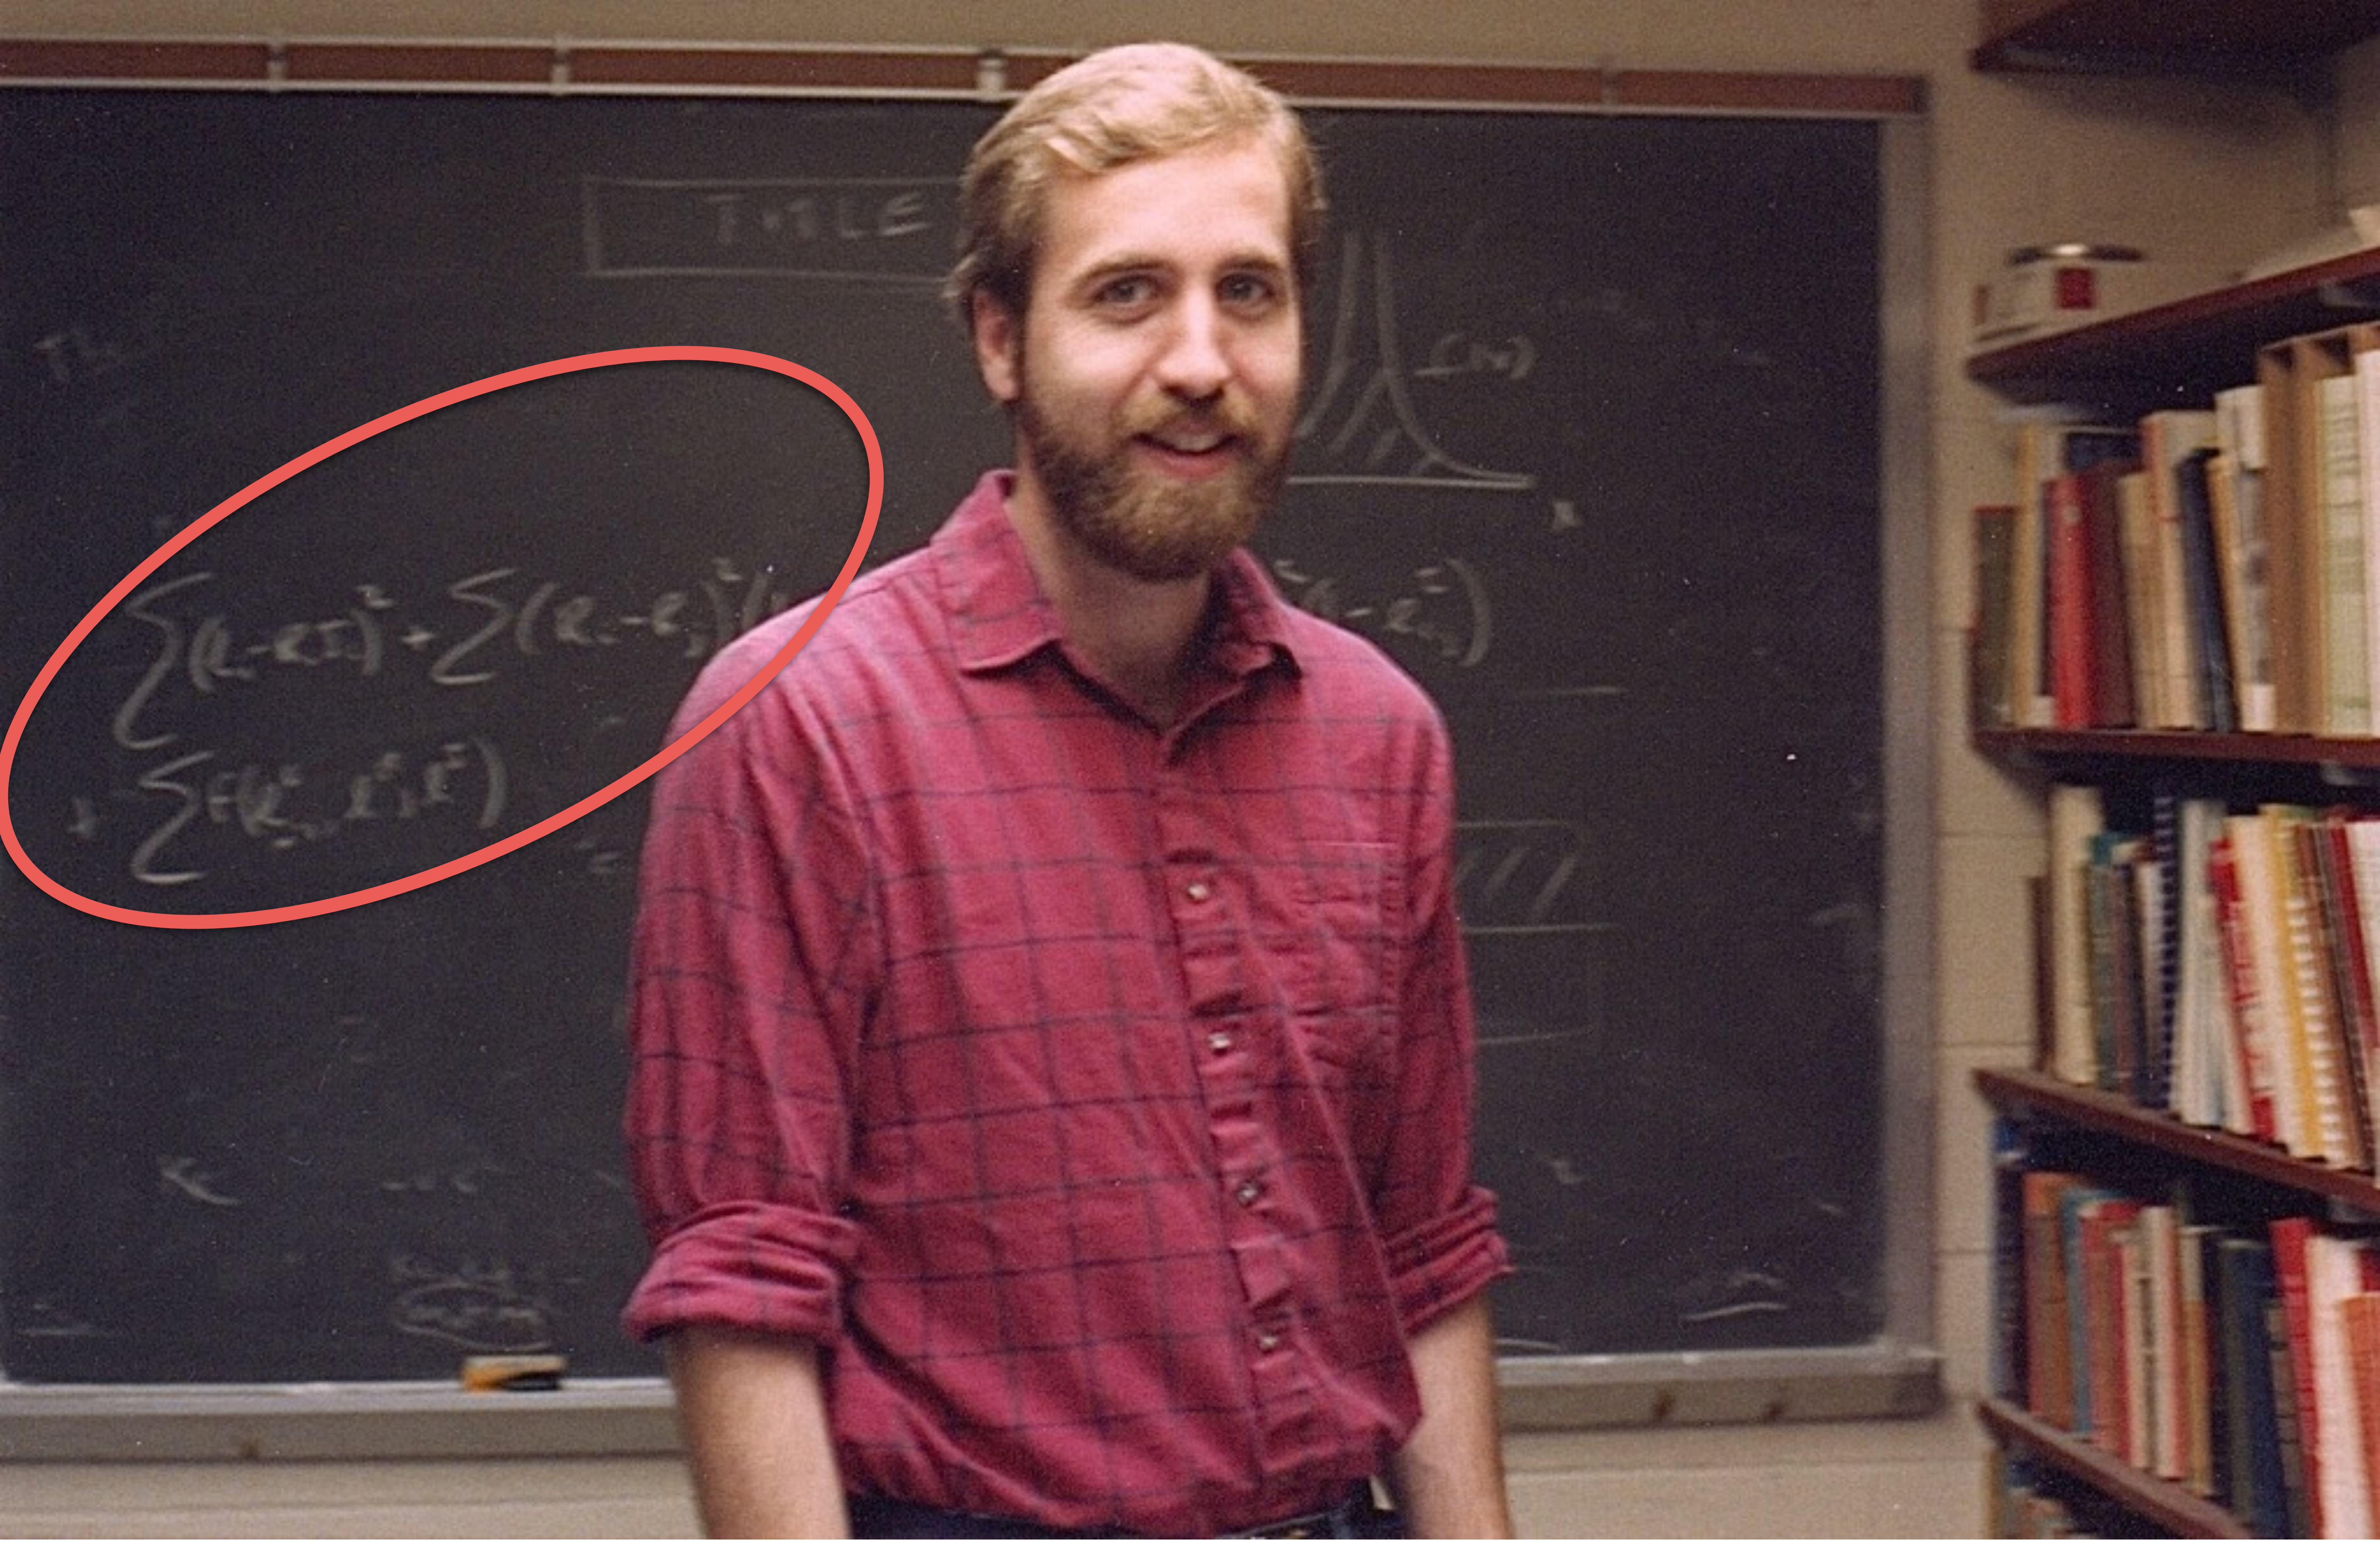

Dave Knill, Brown University, ca. 1986

inferring causes

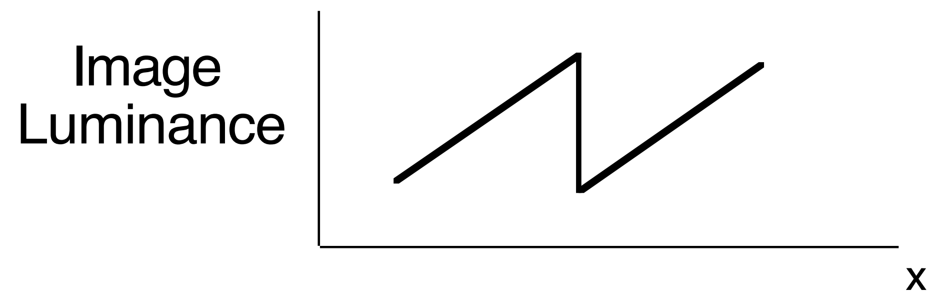

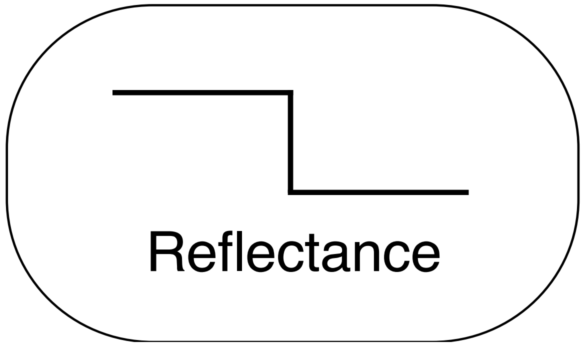

X

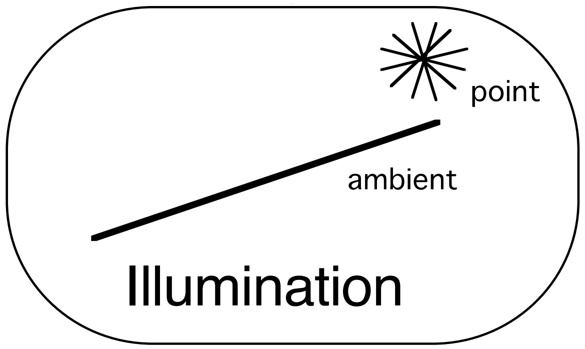

causal view

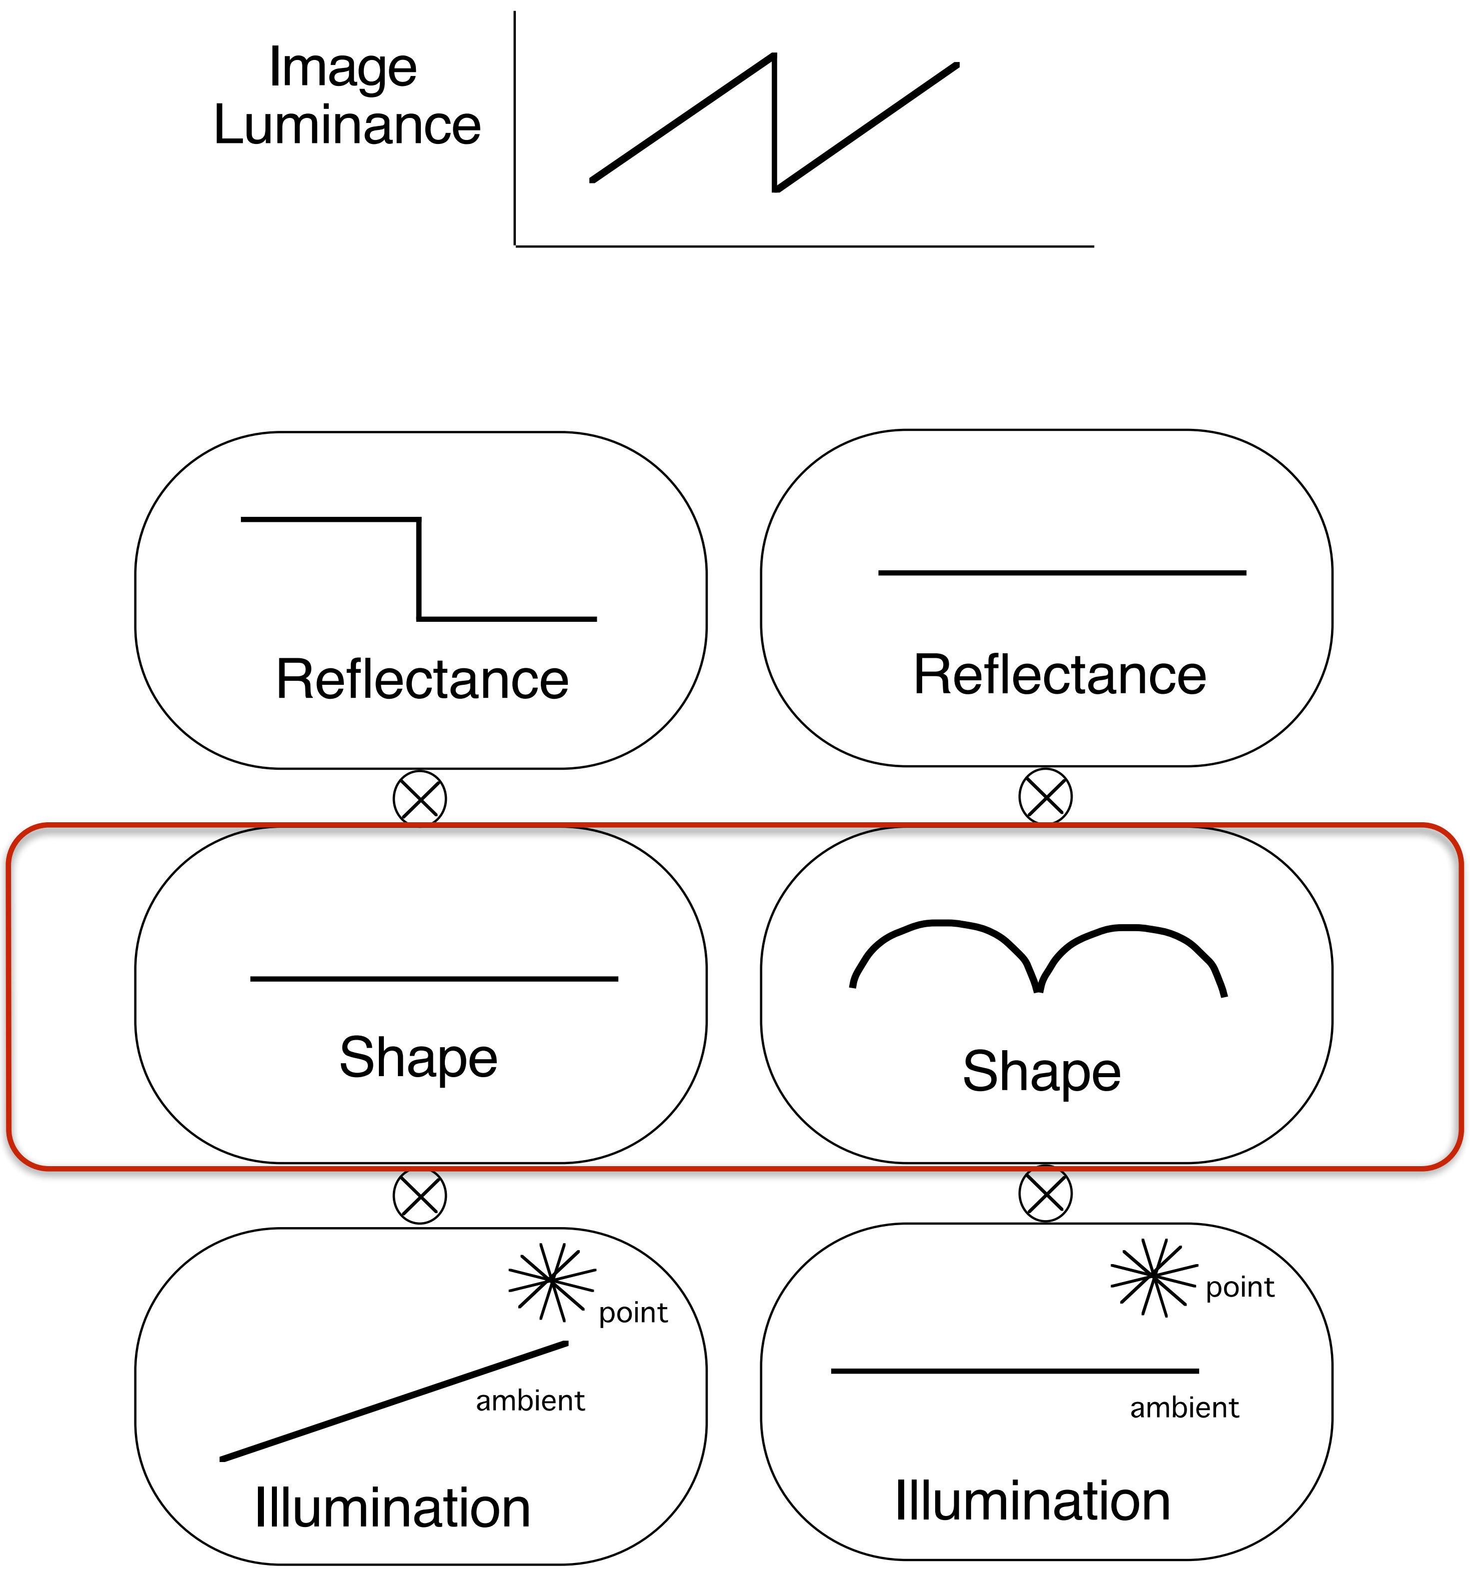
Image
Luminance
Reflectance
Reflectance
Shape
Shape
point
ambient
Illumination
point
ambient
Illumination

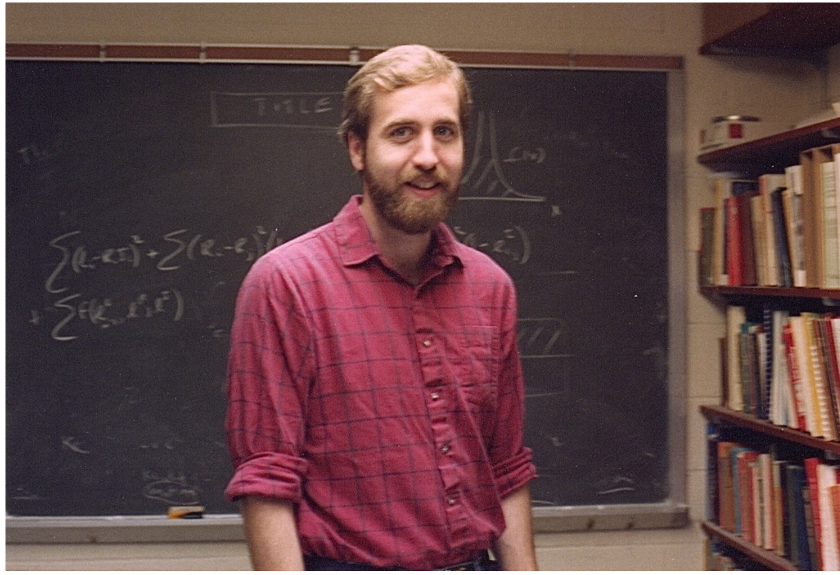

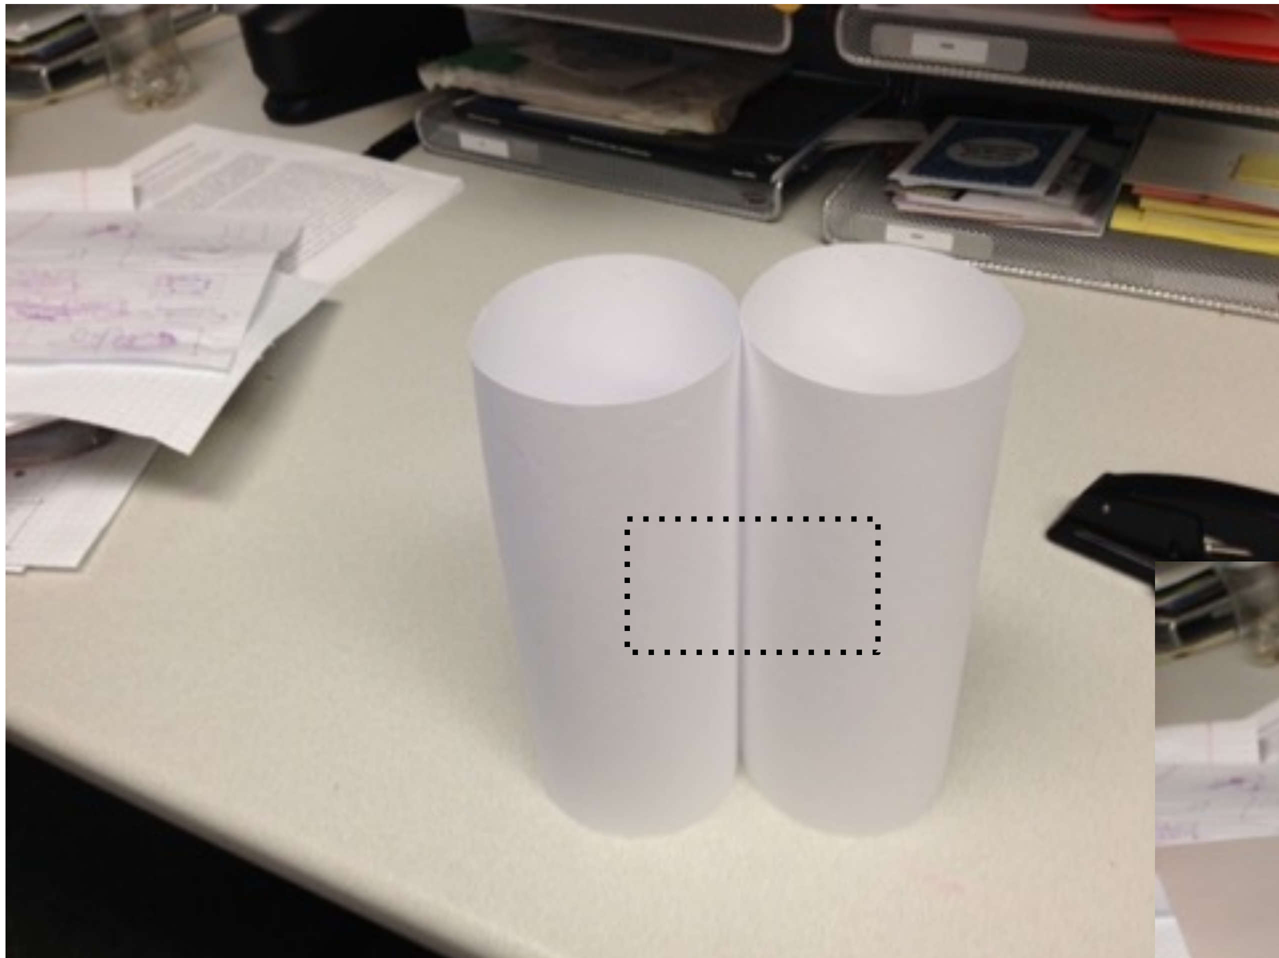

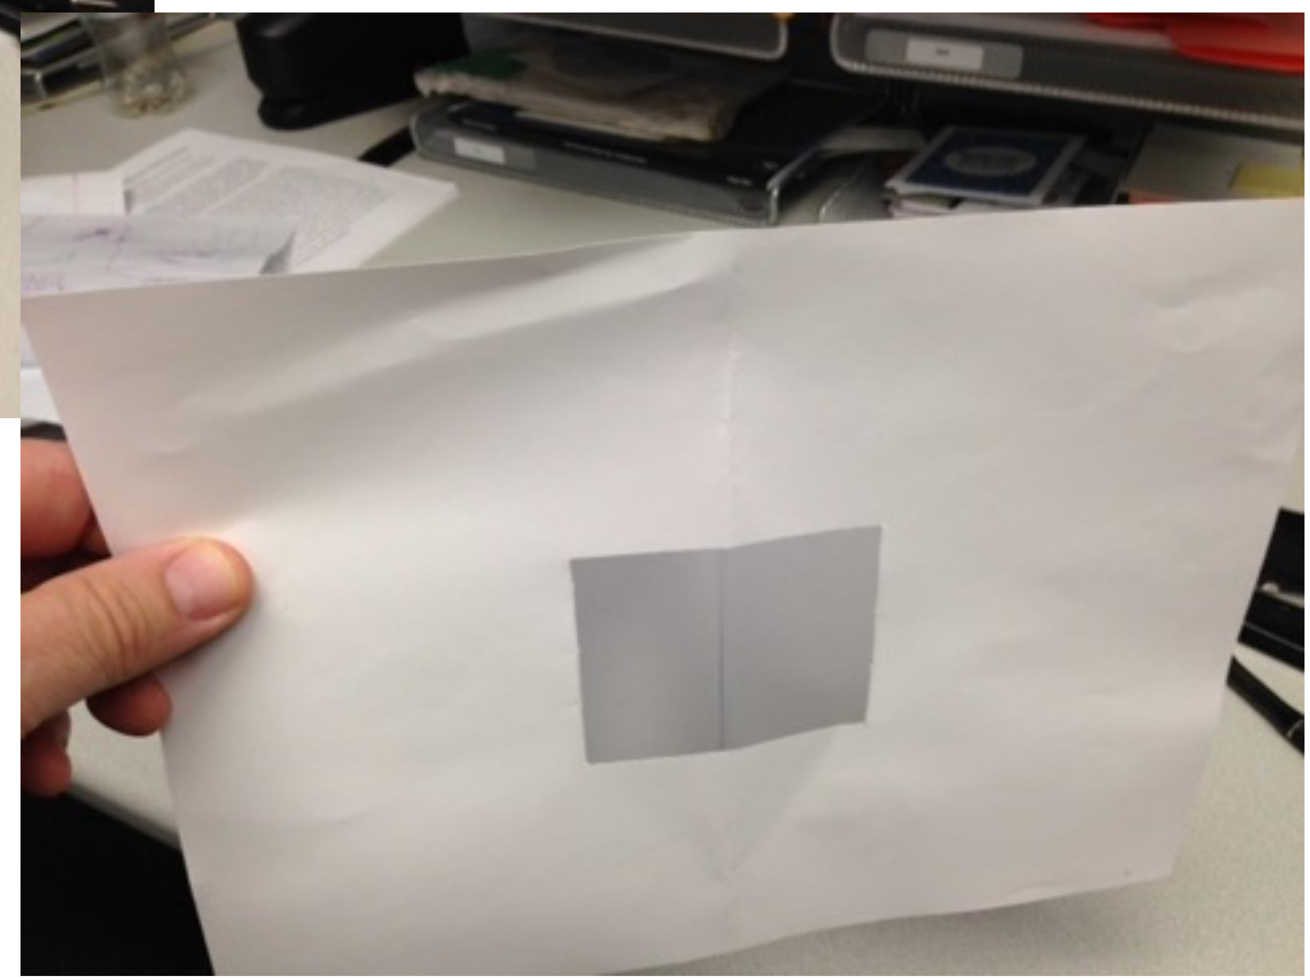

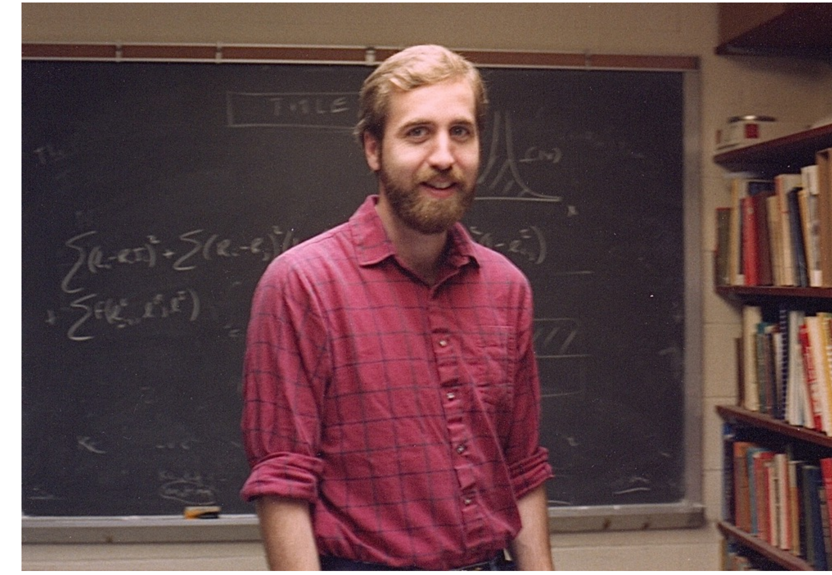

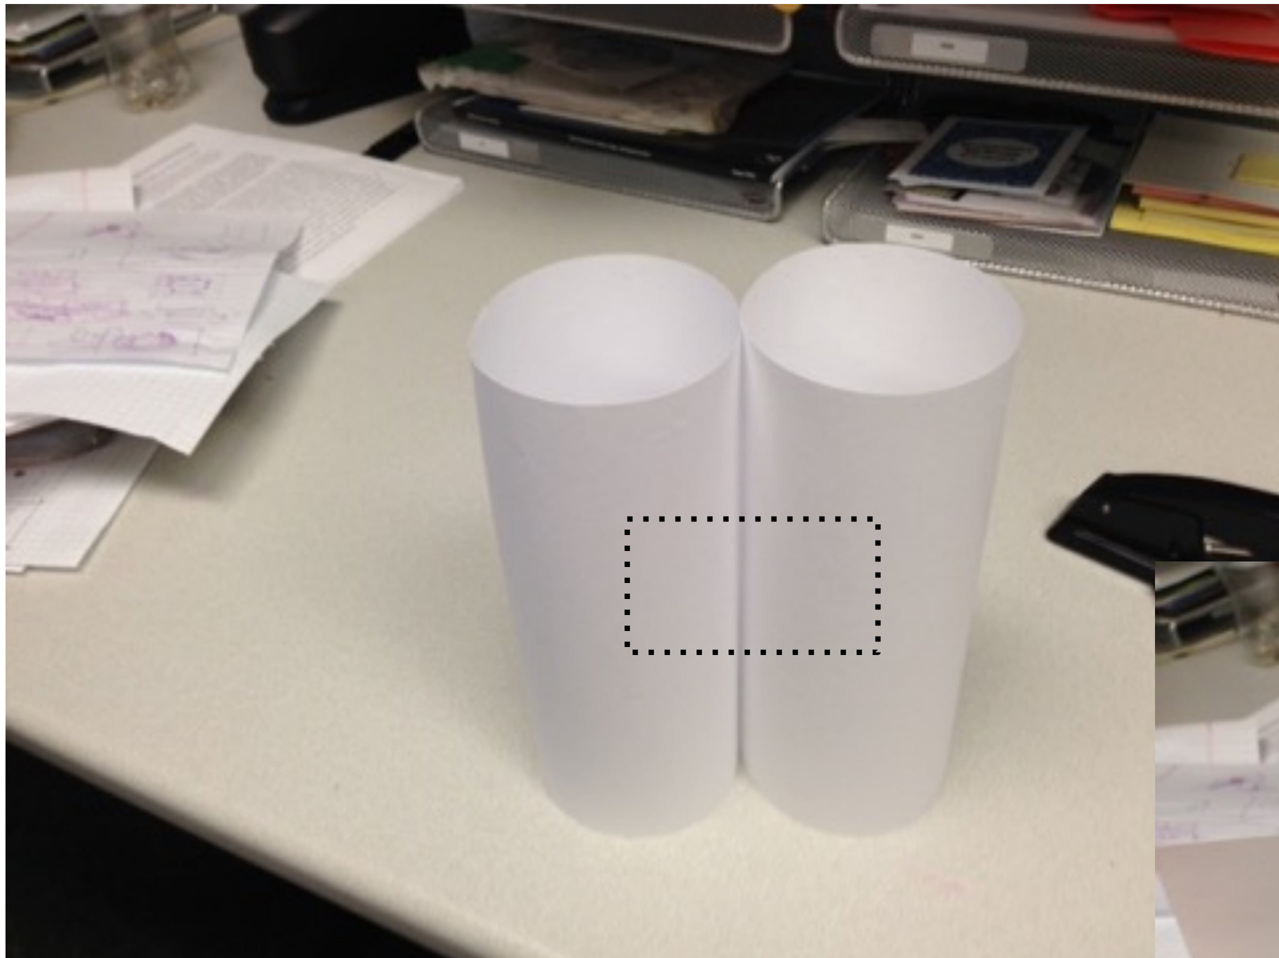

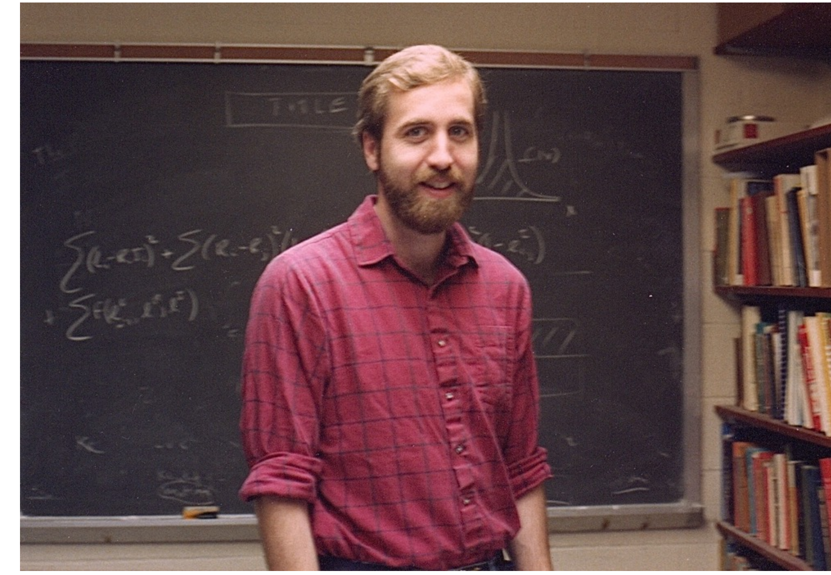

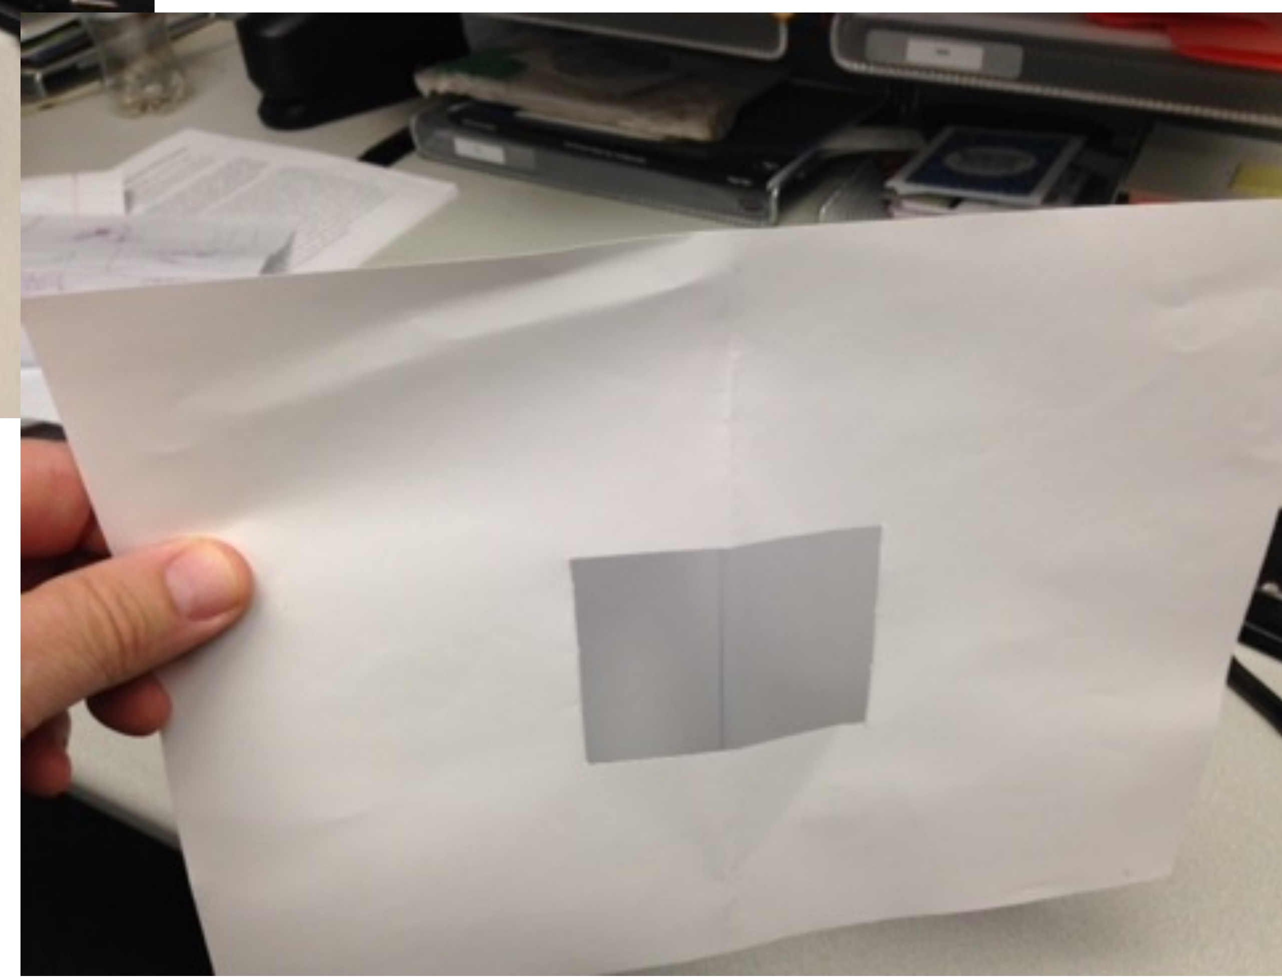

"harumph!"

(anonymous senior professor)

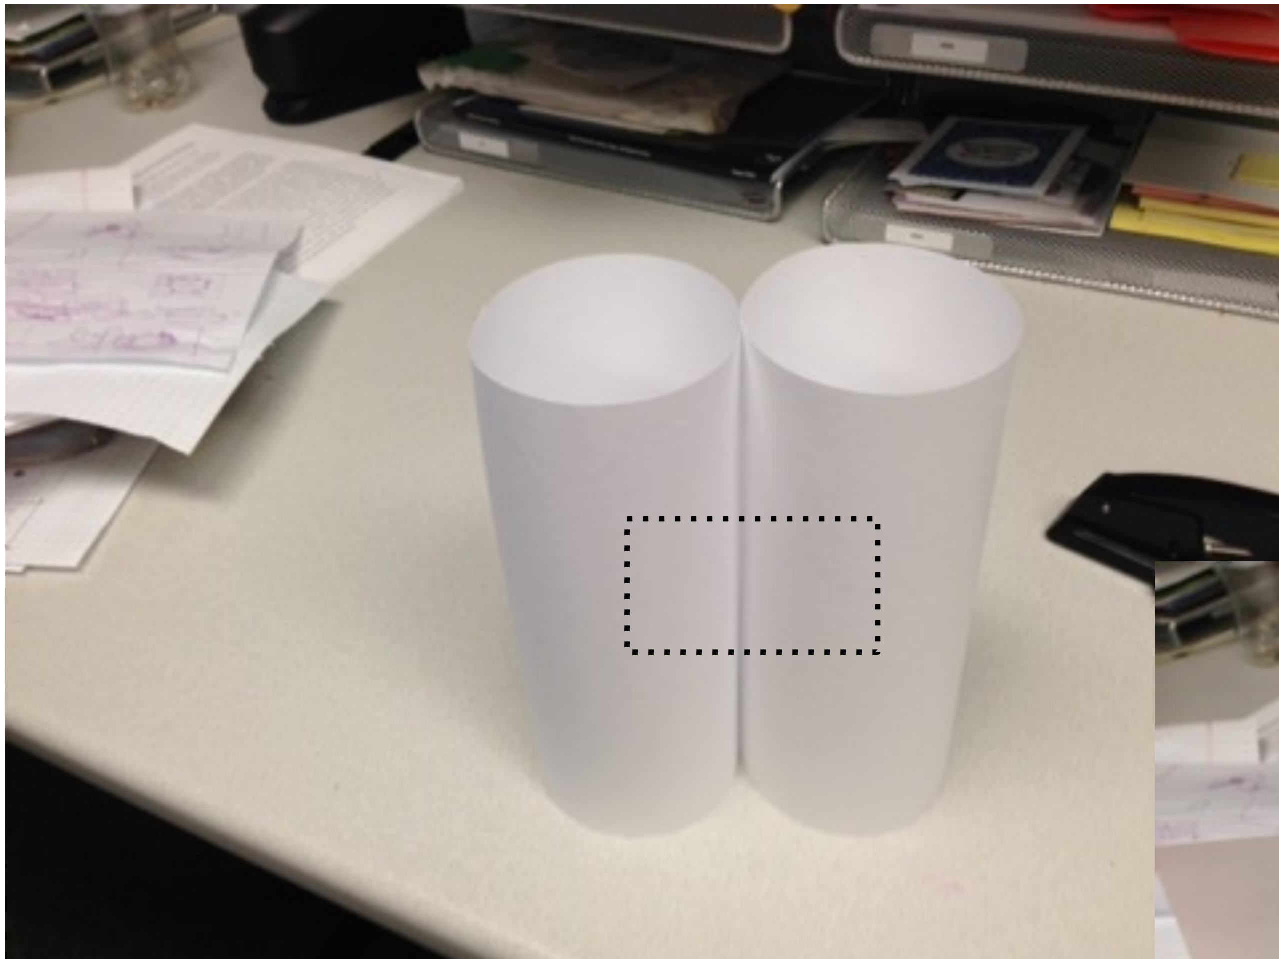

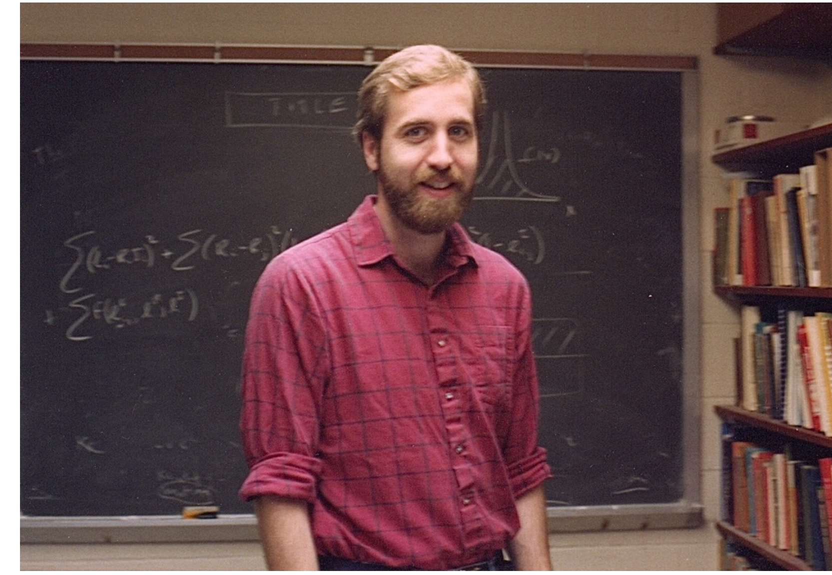

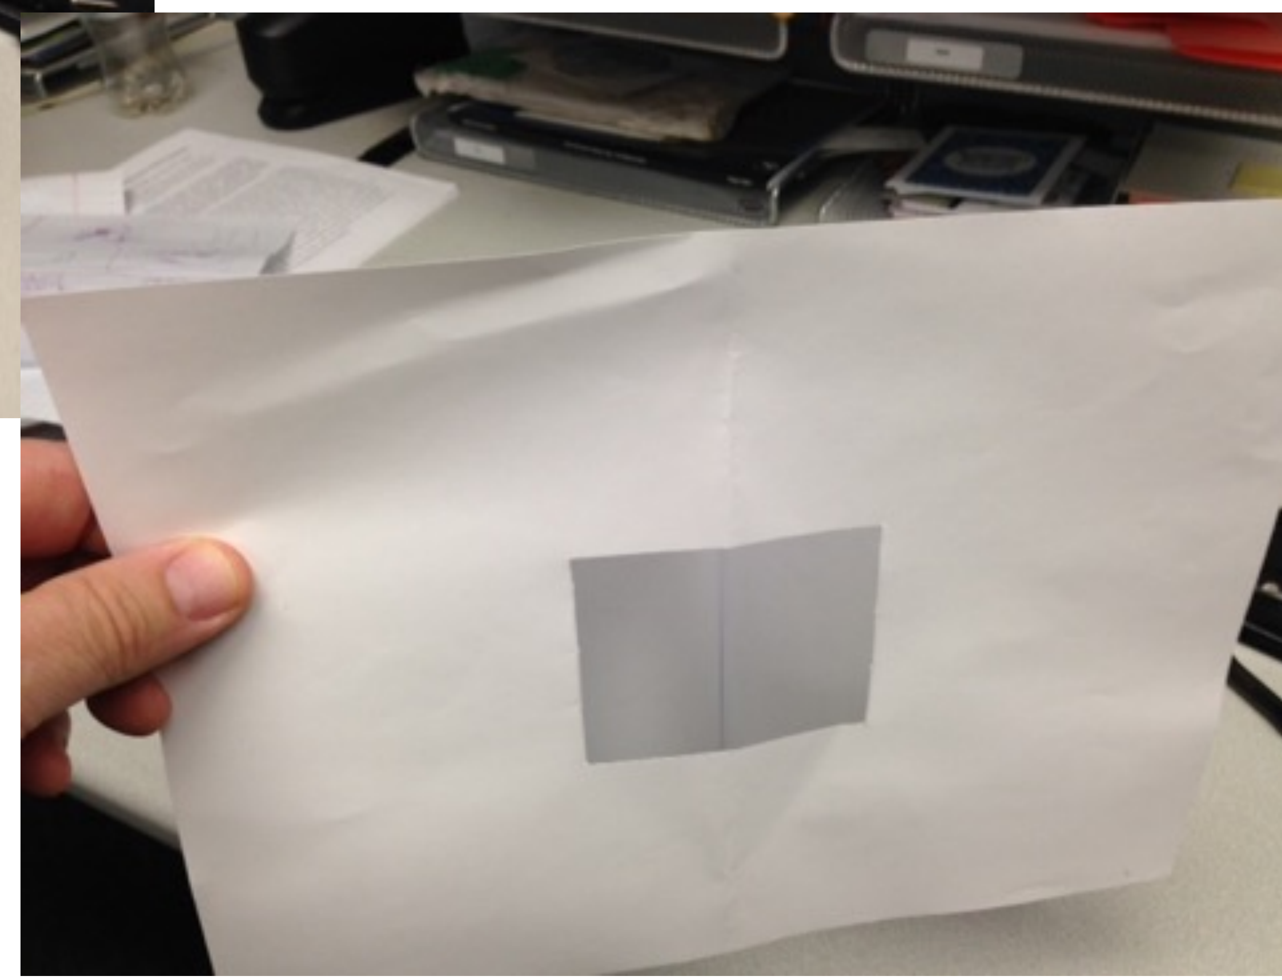

"harumph!"

(anonymous senior professor)

"An expression of disdain, disbelief, protest, refusal or dismissal" - en.wiktionary.org

Dave’s response

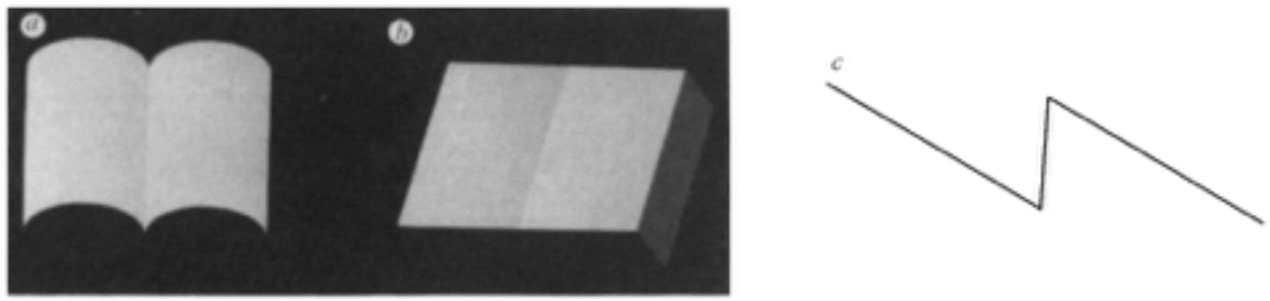

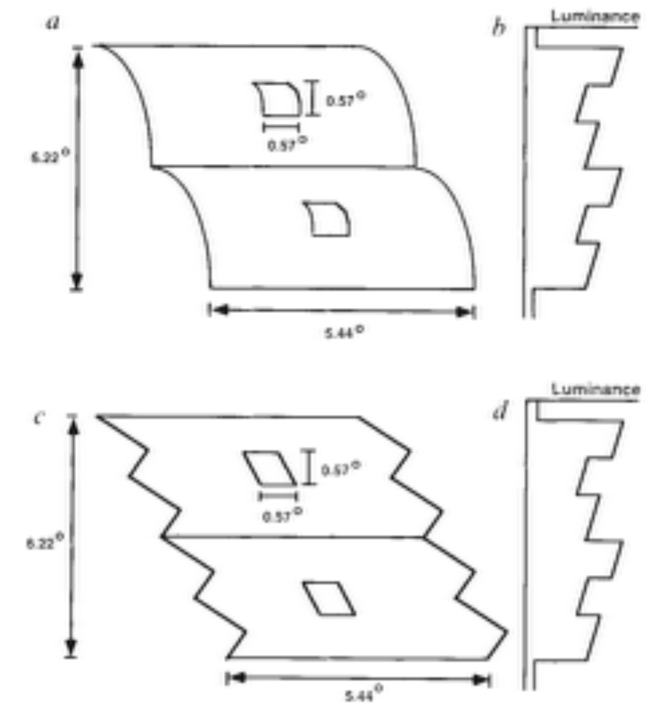

Use 3D graphics to make the stimuli

Do psychophysics with an indirect matching task of lightness

Knill, D. C. & Kersten, D. (1991) Apparent surface curvature affects lightness perception. Nature, 351, 228-230

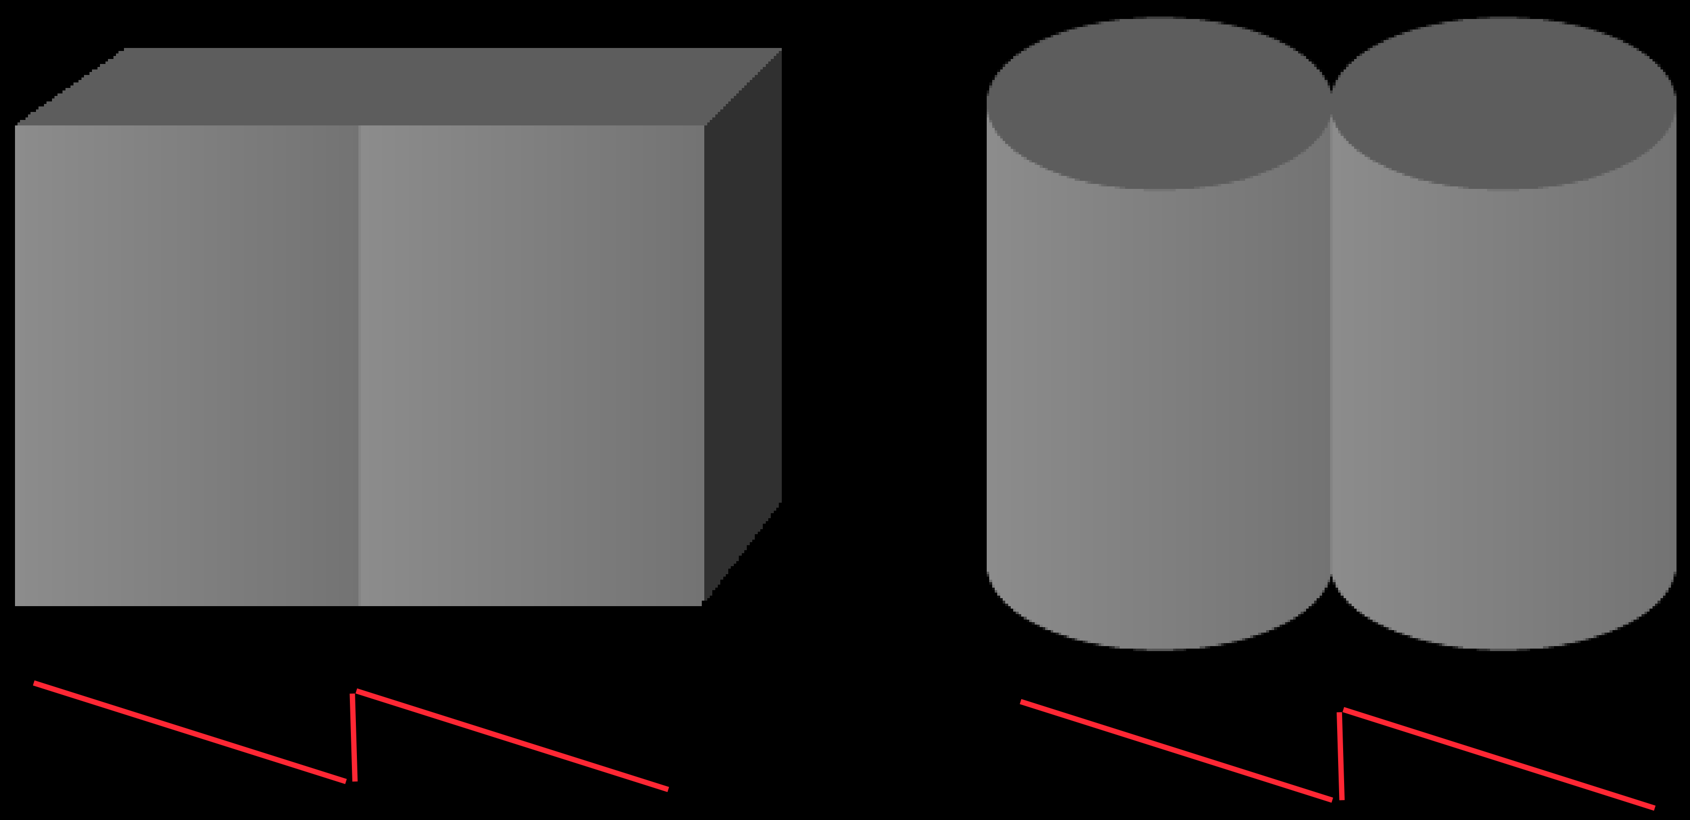

   

contours, shape (and material)

1988: Dave goes solo

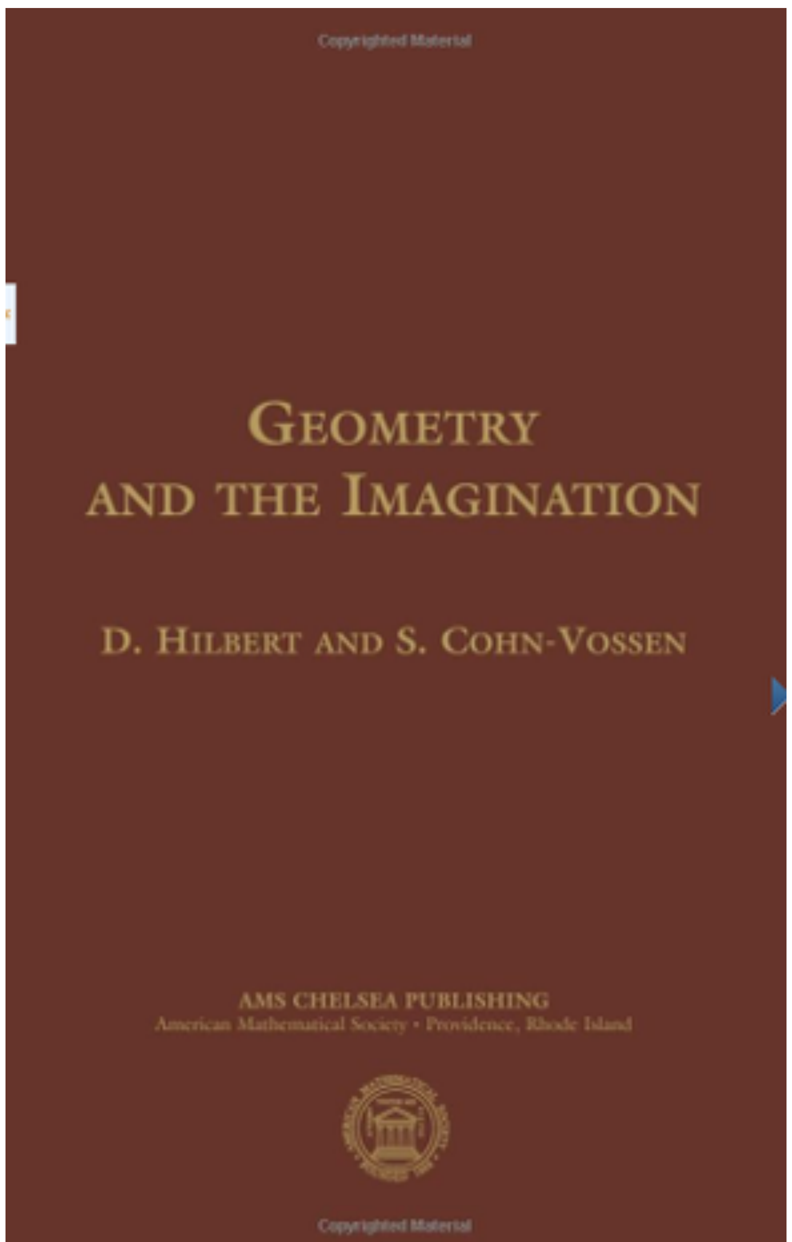

contours, shape (and material)

1988: Dave goes solo

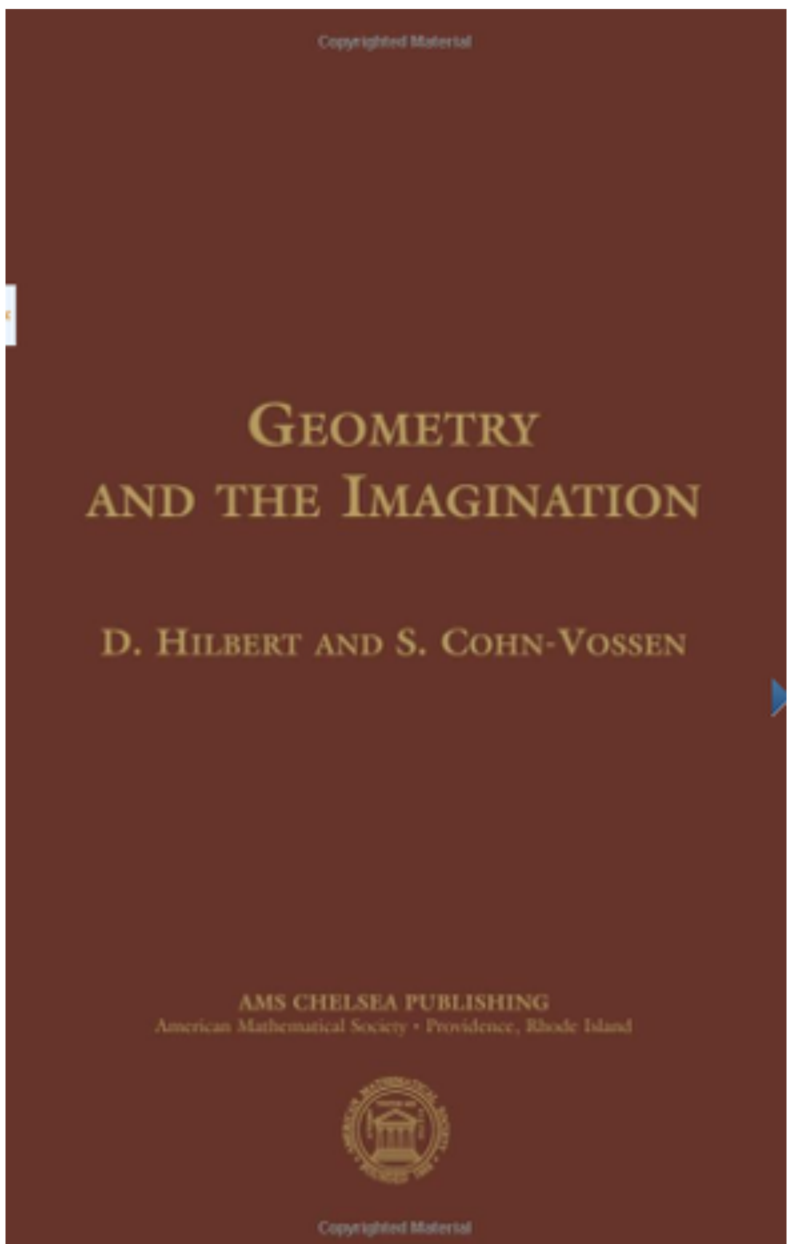

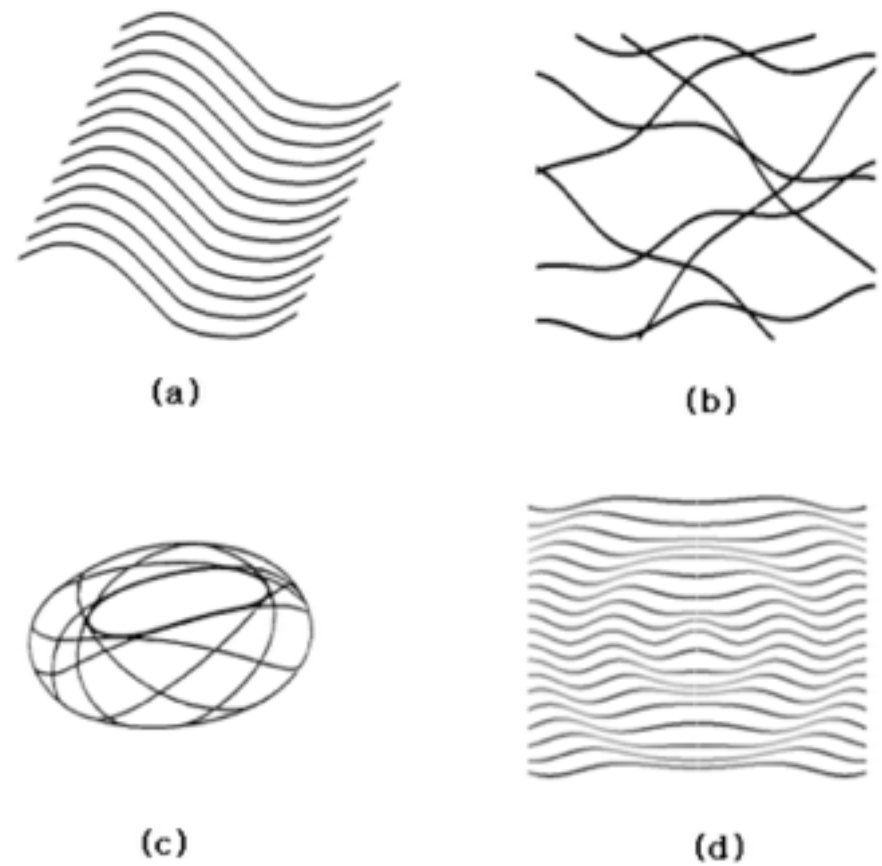

the "geodesic" constraint

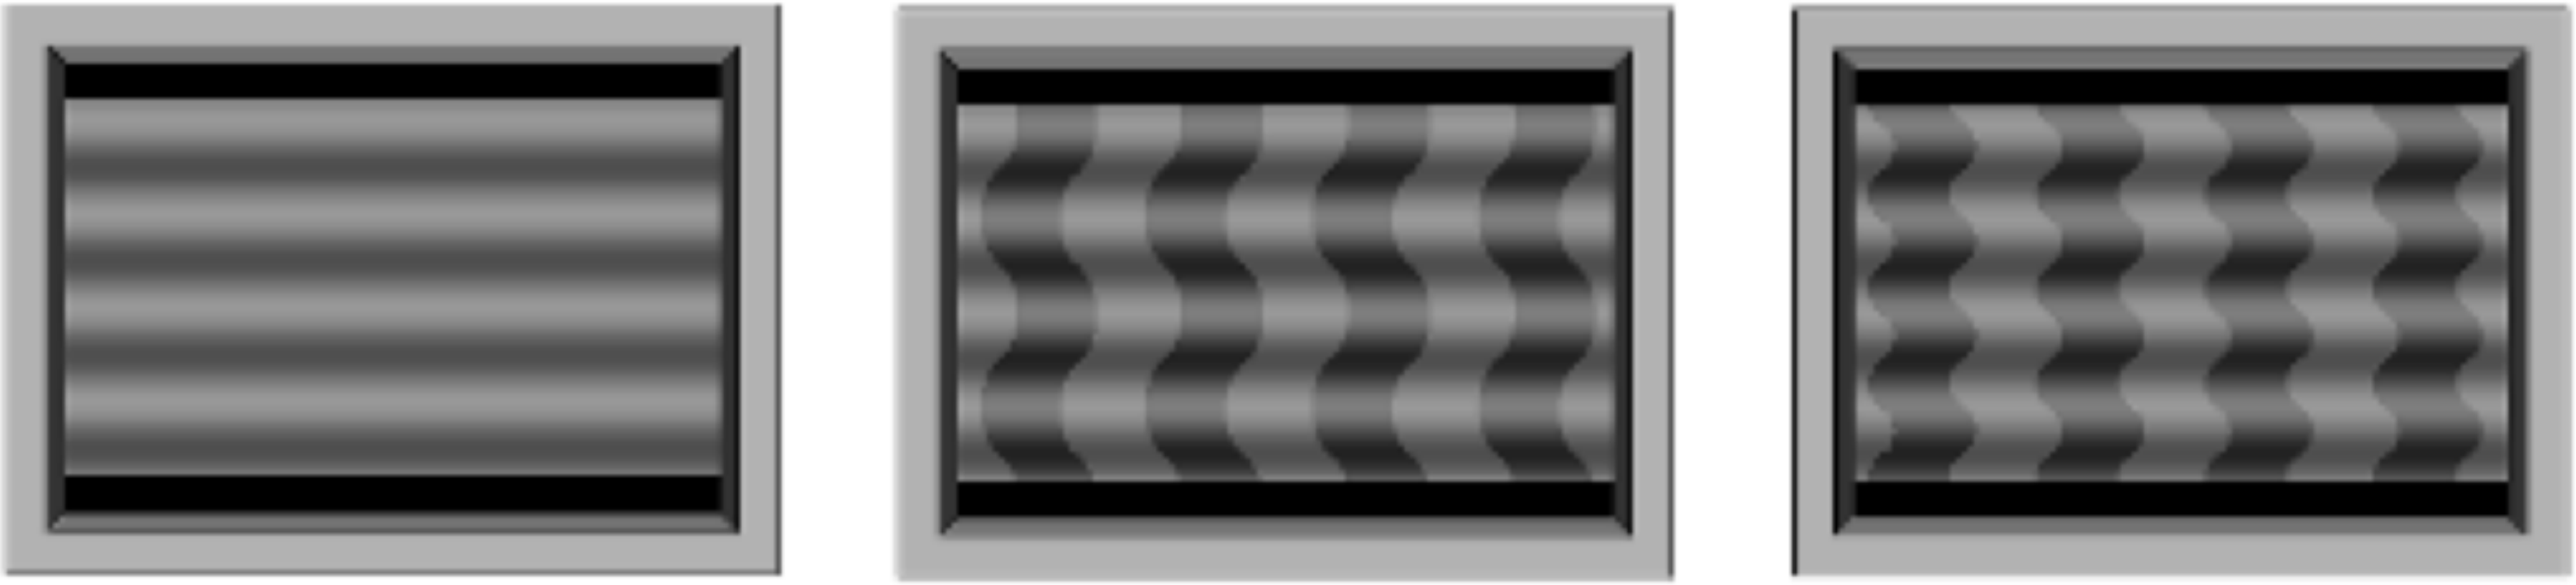

Knill, D. C. (1992) The perception of surface contours and surface shape: from computation to psychophysics. Journal of the Optical Society of America A., 9 (9), 1449- 1464.

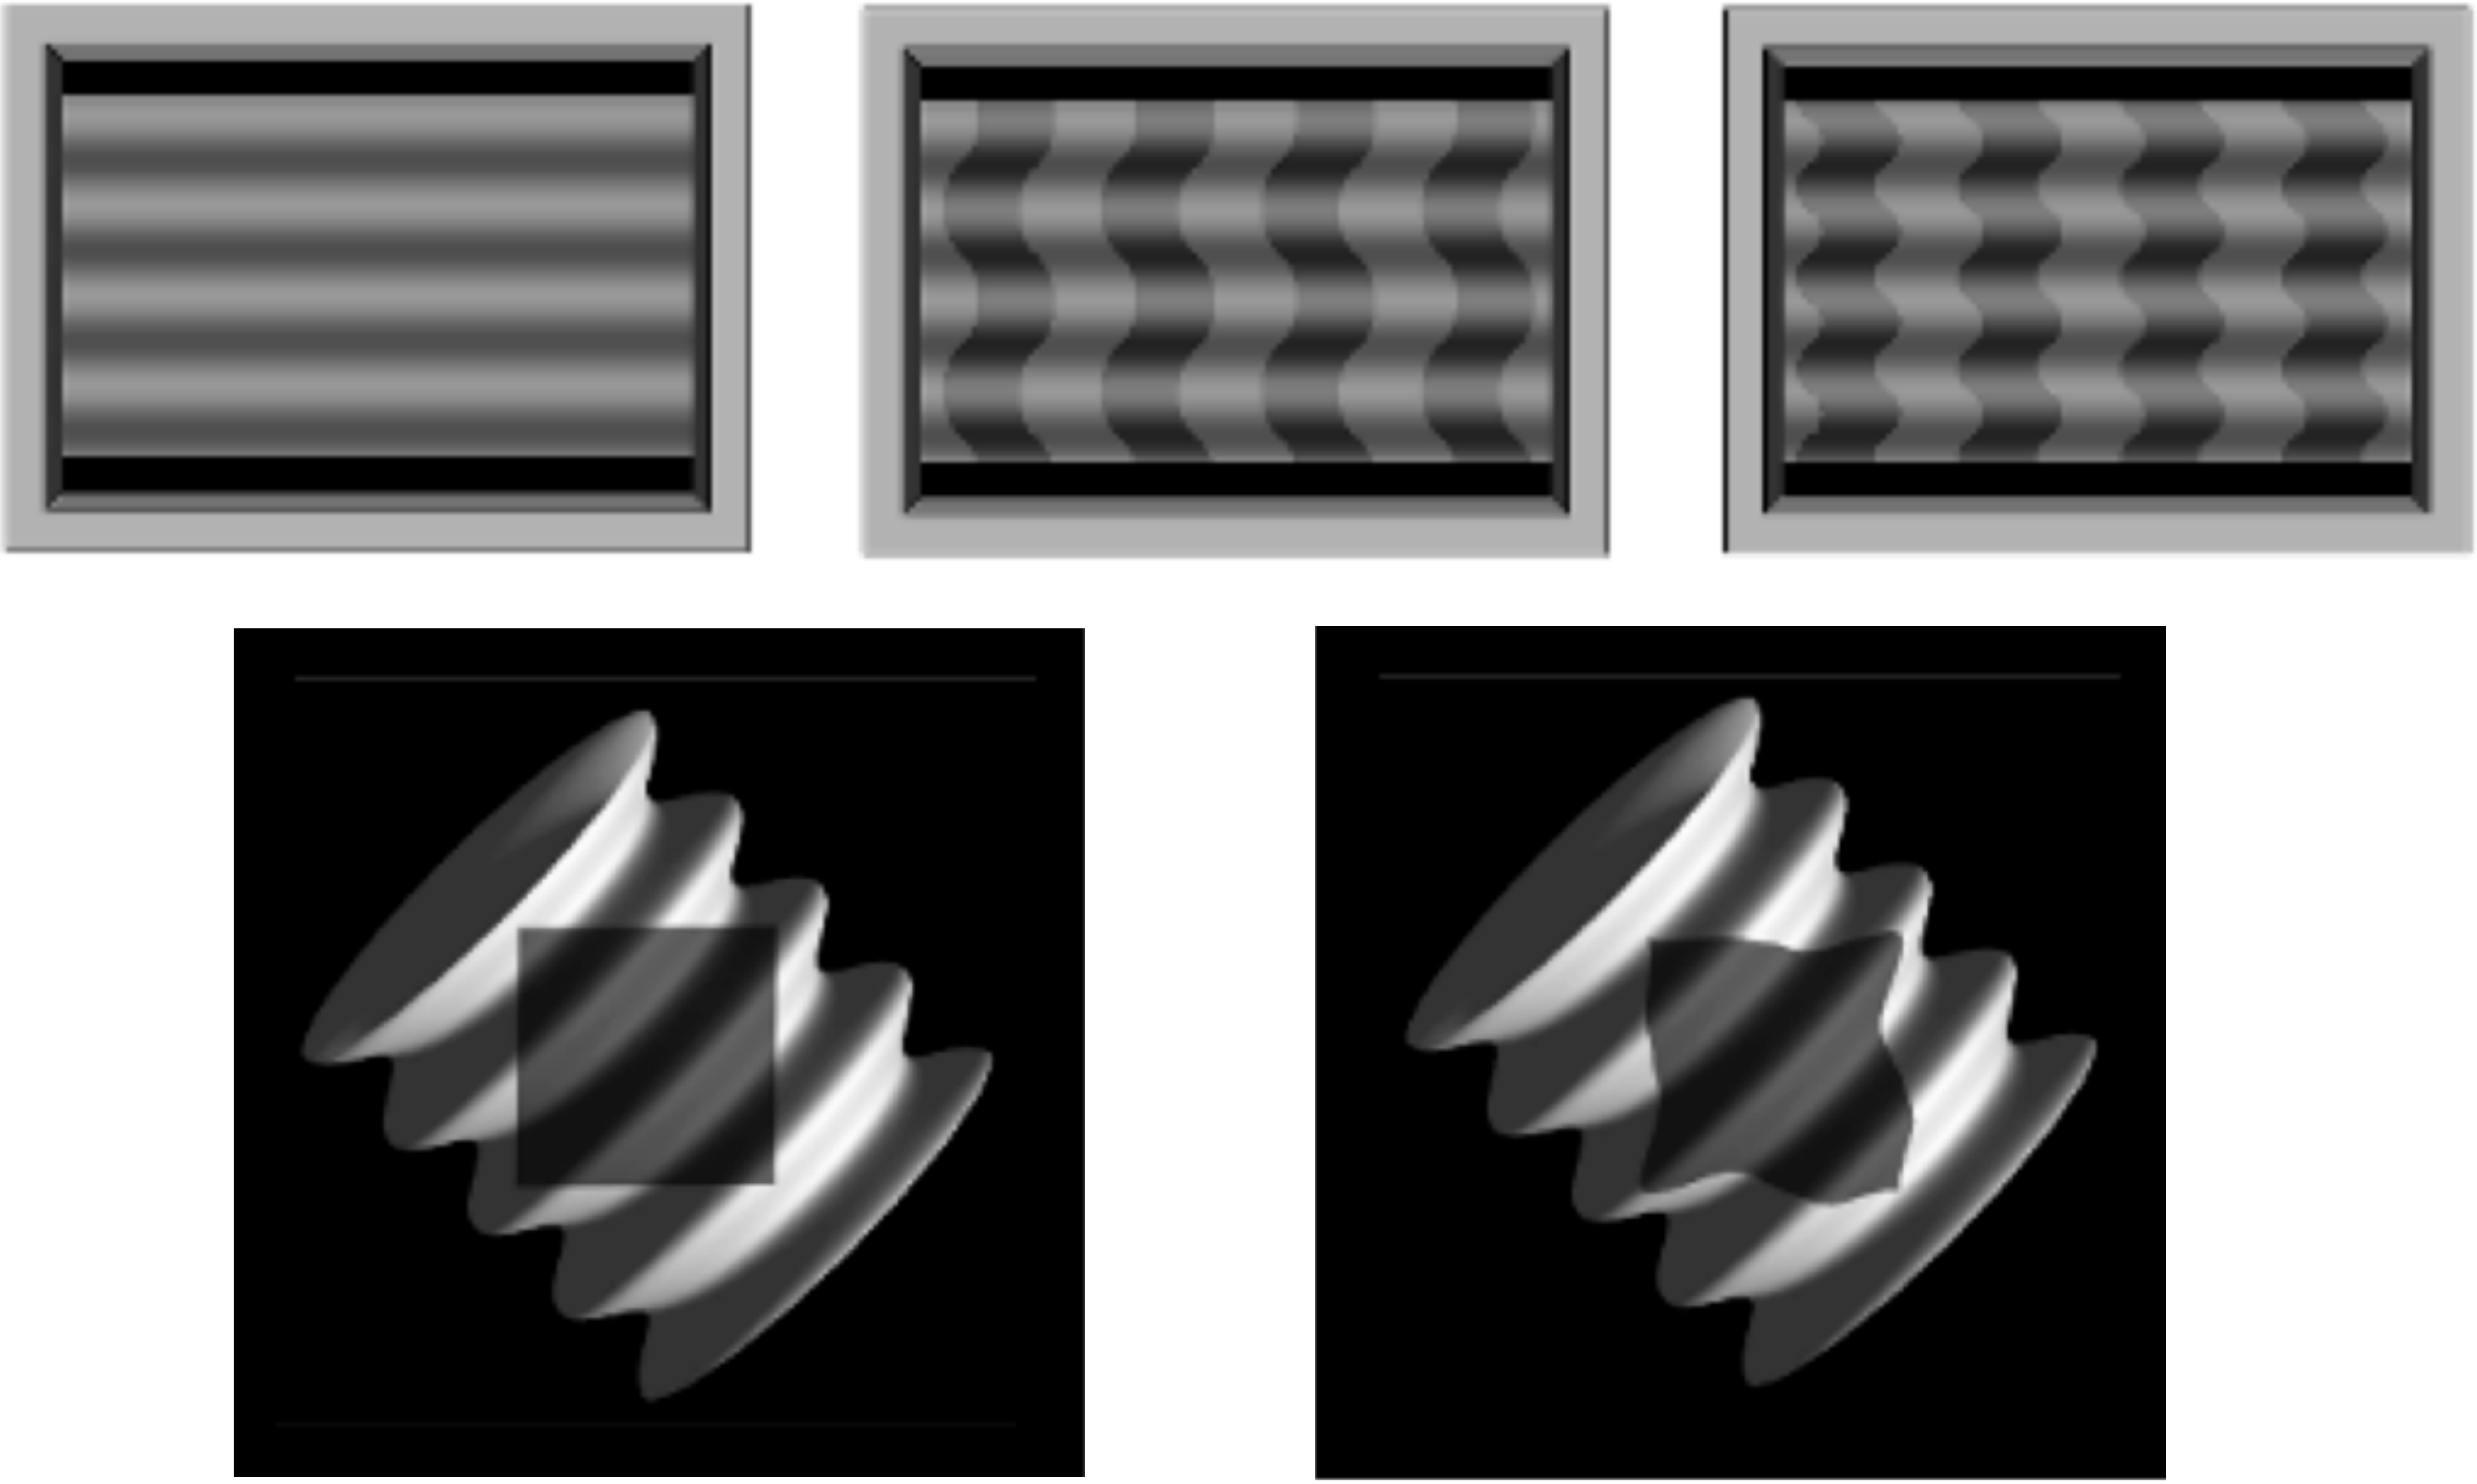

Knill, D. C. (1992) The perception of surface contours and surface shape: from computation to psychophysics. Journal of the Optical Society of America A., 9 (9), 1449- 1464.

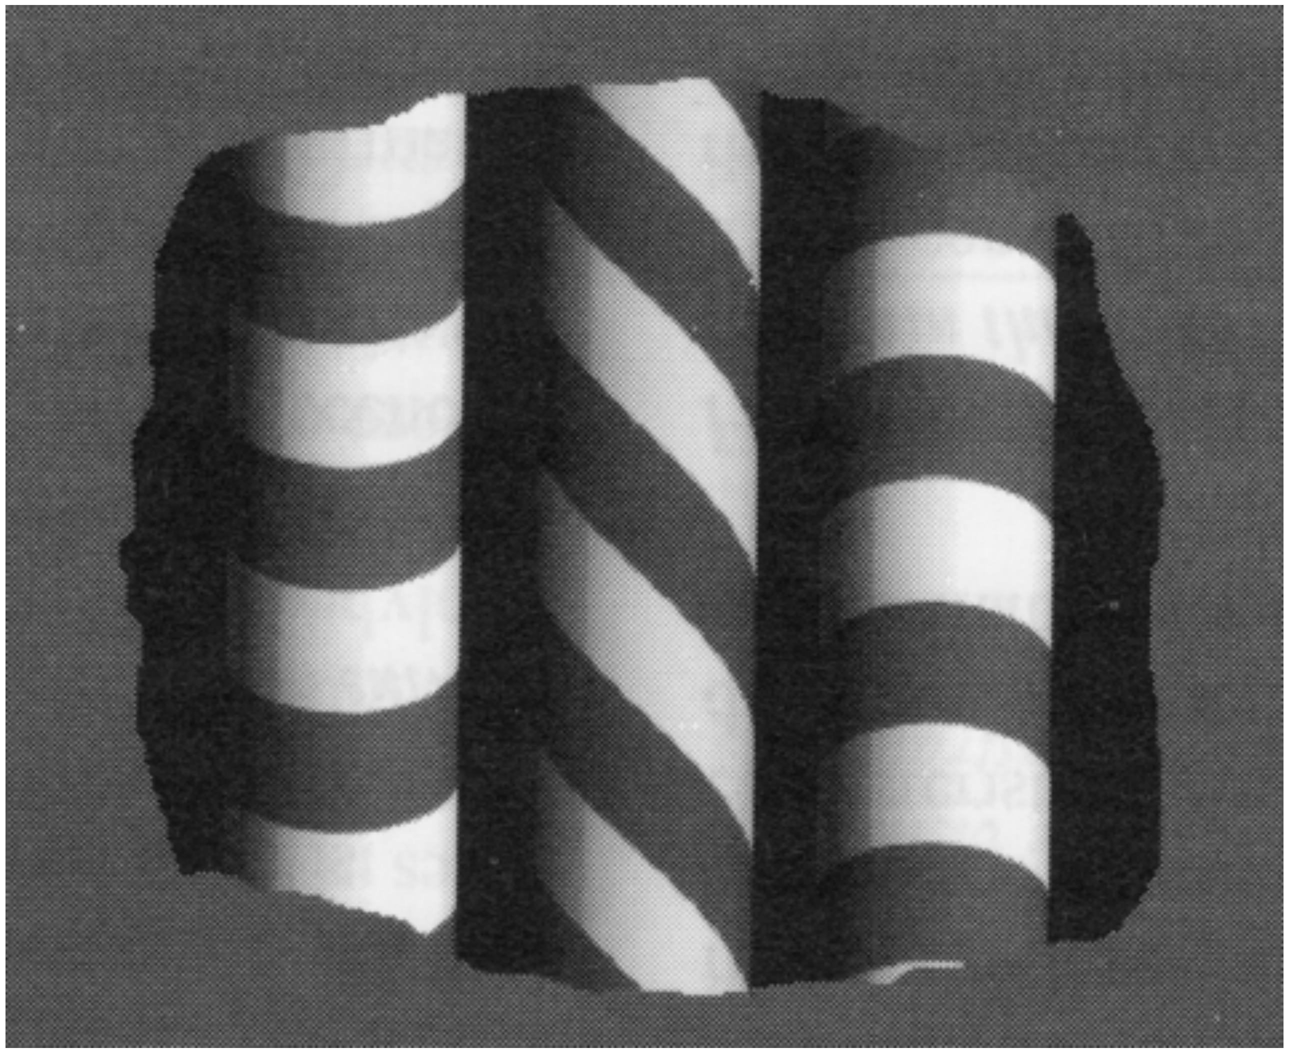

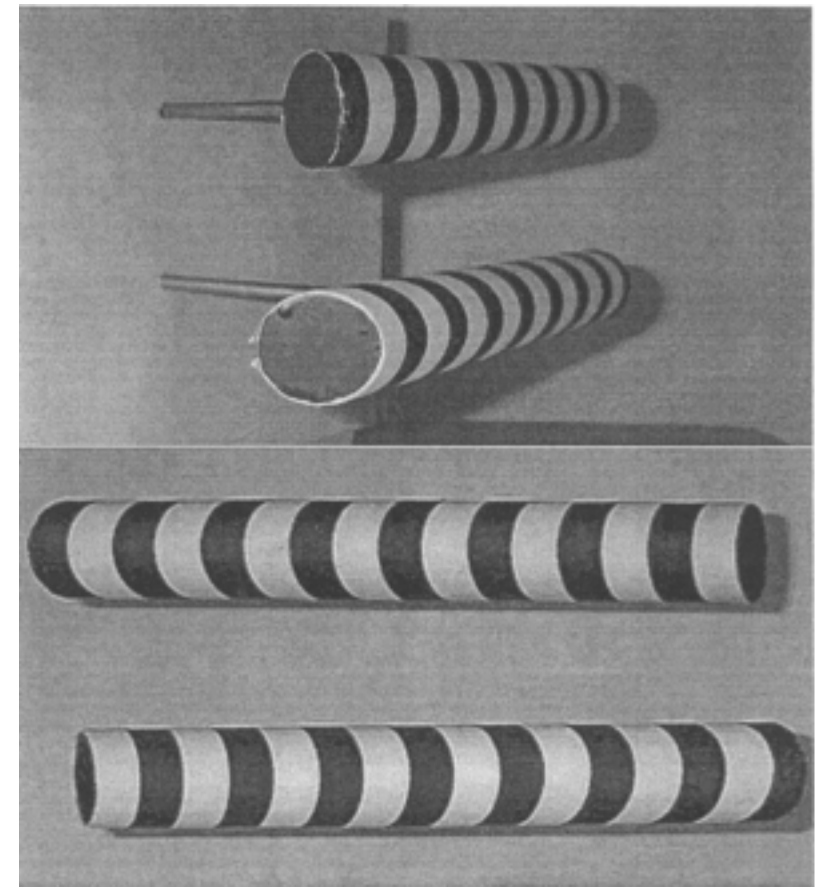

Sen, M. G., Yonas, A., & Knill, D. C. (2001). Development of infants' sensitivity to surface contour information for spatial layout. Perception, 30(2), 167–176

same interior shading pattern
different contour shapes

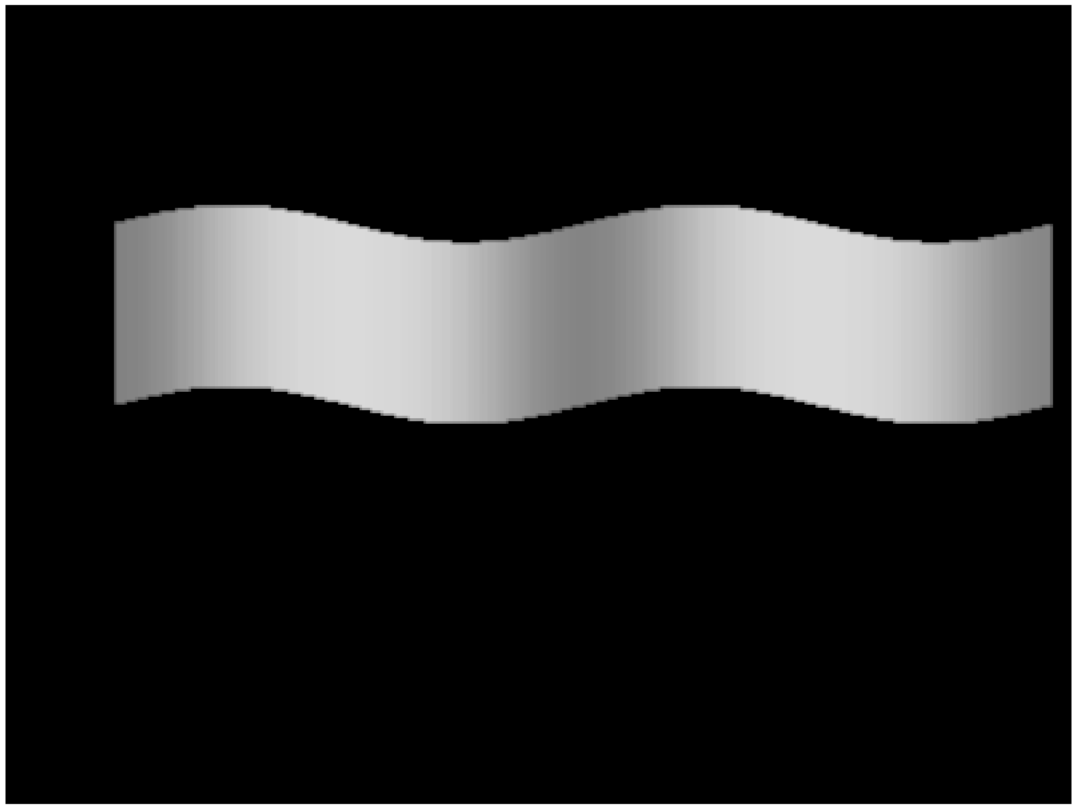

matte

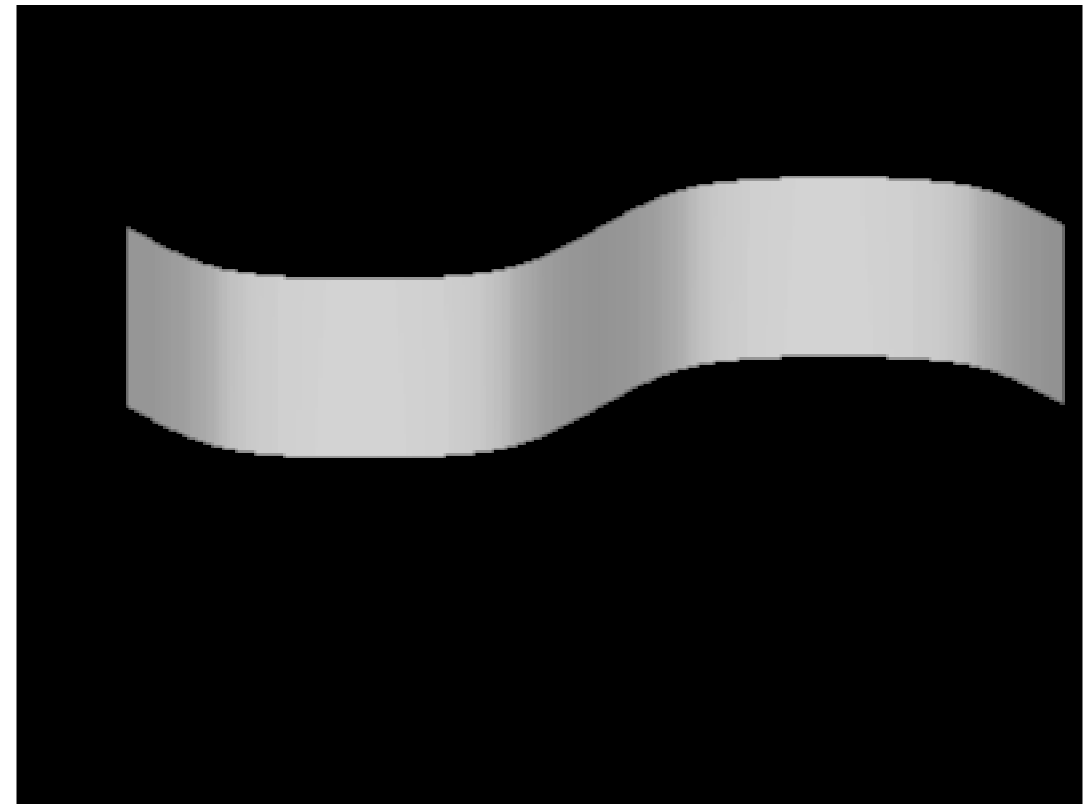

appears more specular

same interior shading pattern
different contour shapes

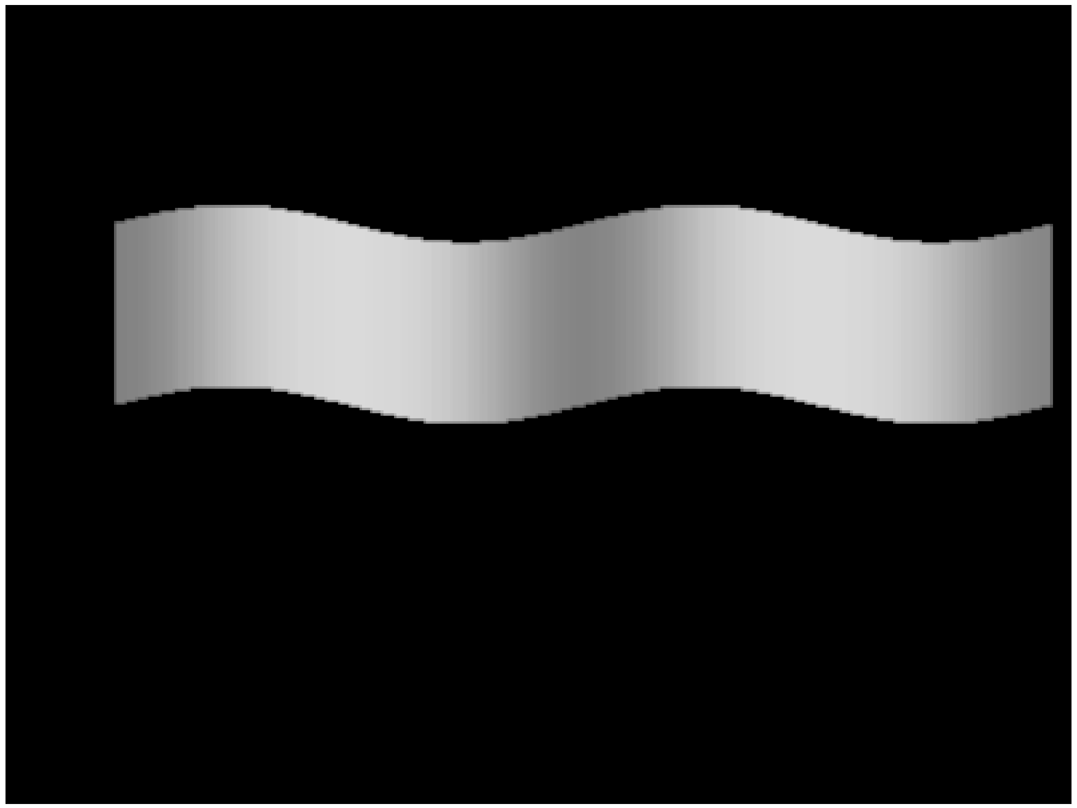

matte

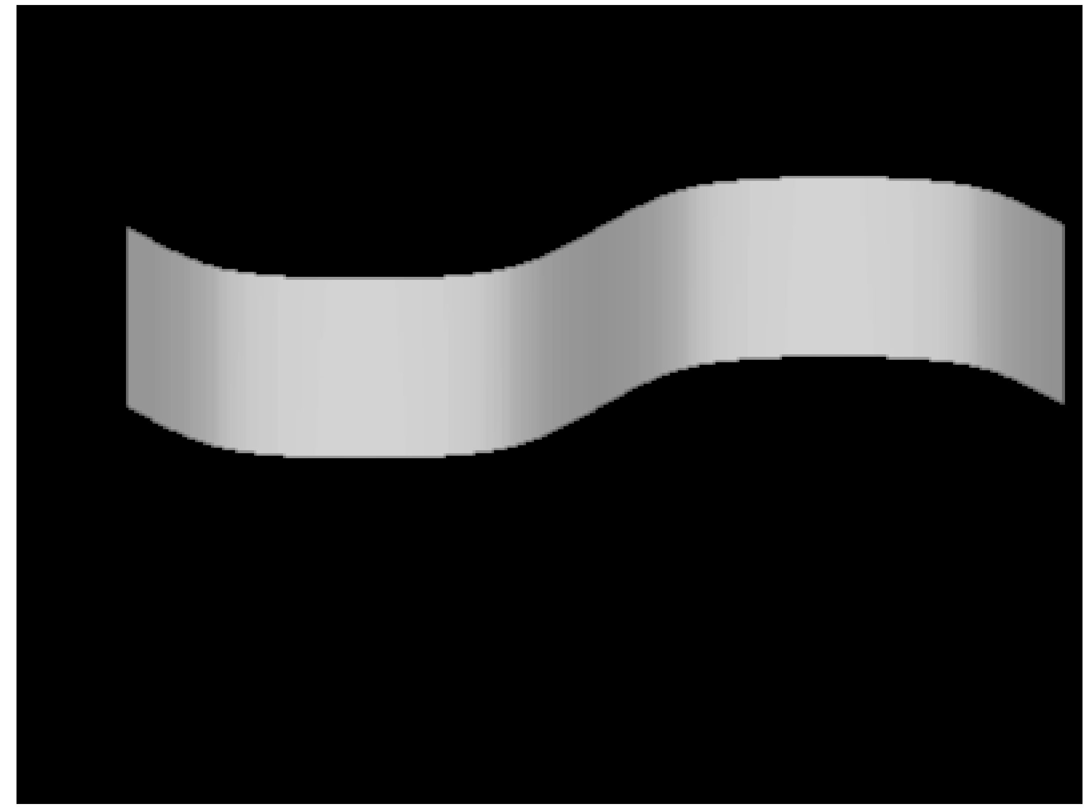

appears more specular

Braje, W. L. and Knill, D. C. (1994) Apparent surface shape influences perceived specular reflectance of curved surfaces. Poster presented at the annual meeting of the Association for Vision and Ophthalmology; Sarasota, FL.

same interior shading pattern
different contour shapes

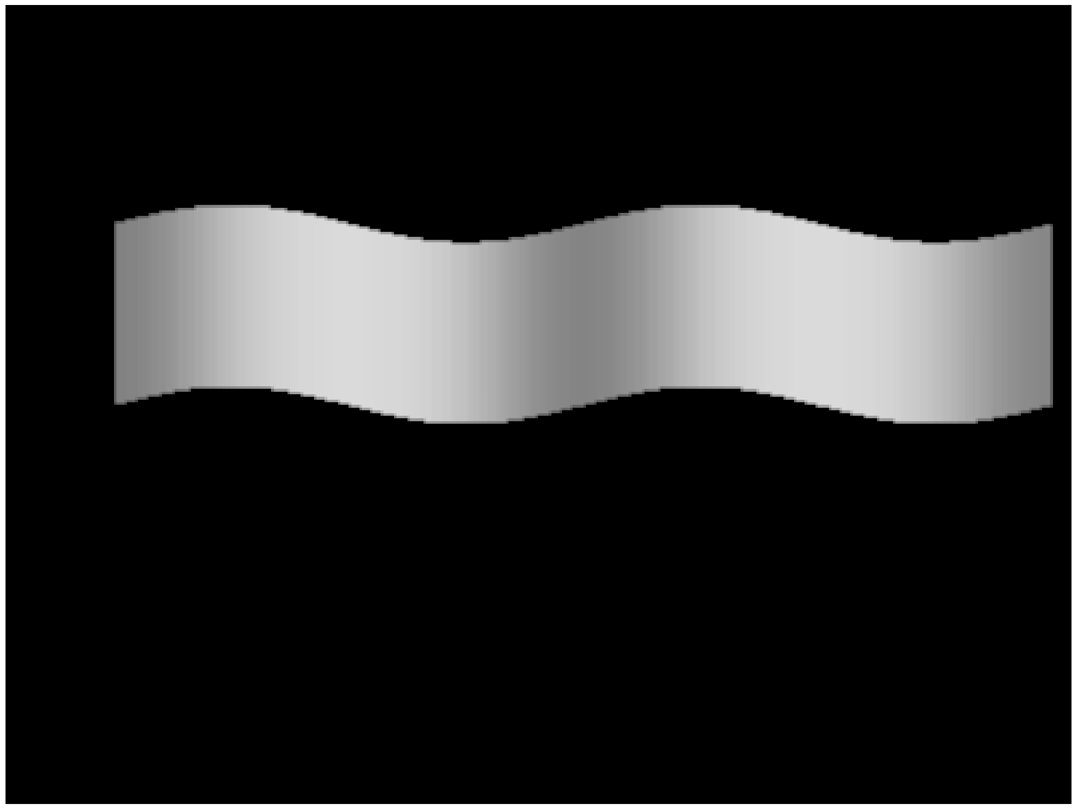

matte

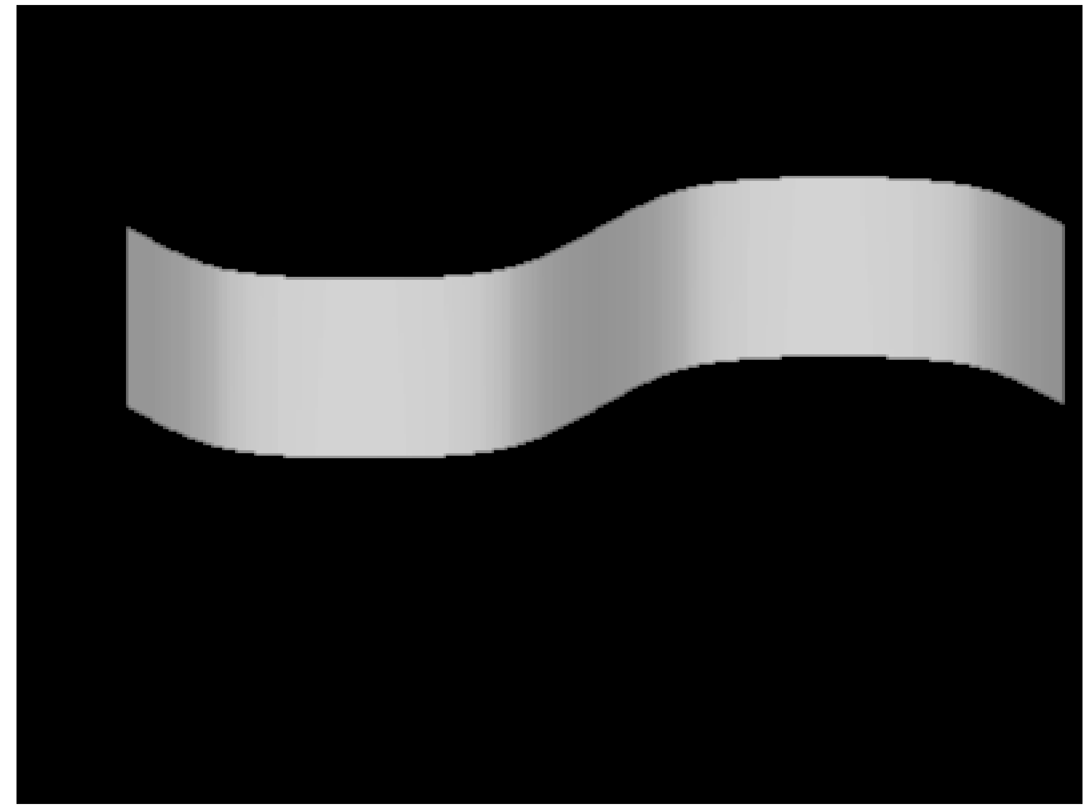

appears more specular

Braje, W. L. and Knill, D. C. (1994) Apparent surface shape influences perceived specular reflectance of curved surfaces. Poster presented at the annual meeting of the Association for Vision and Ophthalmology; Sarasota, FL.

Marlow, P. J., Todorović, D., & Anderson, B. L. (2015). Coupled computations of three-dimensional shape and material. Current Biology, 25(6), R221–R222.

Minnesota: 1990-1994

Zili Liu, Pascal Mamassian, Wendy Braje, Suthep Madarasmi, Bosco Tjan..

Al Yonas, Irv Biederman, Gordon Legge

Visiting professors: Heinrich Bülthoff, Alan Yuille, Mel Goodale

1991: Dave the seer

"Maybe the brain represents probability distributions, not just estimates" — Dave Knill, ca. 1991

1991: Dave the seer

“Maybe the brain represents probability distributions, not just estimates” — Dave Knill, ca. 1991

1991: Dave the seer

"Maybe the brain represents probability distributions, not just estimates" — Dave Knill, ca. 1991

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. Neural Computation, 10(2), 403–430

.

.

.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. TINs, 27(12), 712–719.

The Chatham meeting and book

Ted Adelson, Horace Barlow, Peter Belhumeur, Bennett, Andrew Blake, Heinrich Bülthoff, Jacob Feldman, Bill Freeman, Stu Geman, Don Hoffman, Alan Jepson, Dan Kersten, Dave Knill, Pascal Mamassian, David Mumford, Ken Nakayama, Alex Pentland, Chetan Prakash, Whitman Richards, Scott Richman, Ron Rensink, Dave Sheinberg, Shin Shimojo, Alan Yuille

John Tangney, AFOSR

1993: Dave the organizer, integrator, and conversant

The Chatham meeting and book

Ted Adelson, Horace Barlow, Peter Belhumeur, Bennett, Andrew Blake, Heinrich Bülthoff, Jacob Feldman, Bill Freeman, Stu Geman, Don Hoffman, Alan Jepson, Dan Kersten, Dave Knill, Pascal Mamassian, David Mumford, Ken Nakayama, Alex Pentland, Chetan Prakash, Whitman Richards, Scott Richman, Ron Rensink, Dave Sheinberg, Shin Shimojo, Alan Yuille

John Tangney, AFOSR

1993: Dave the organizer, integrator, and conversant

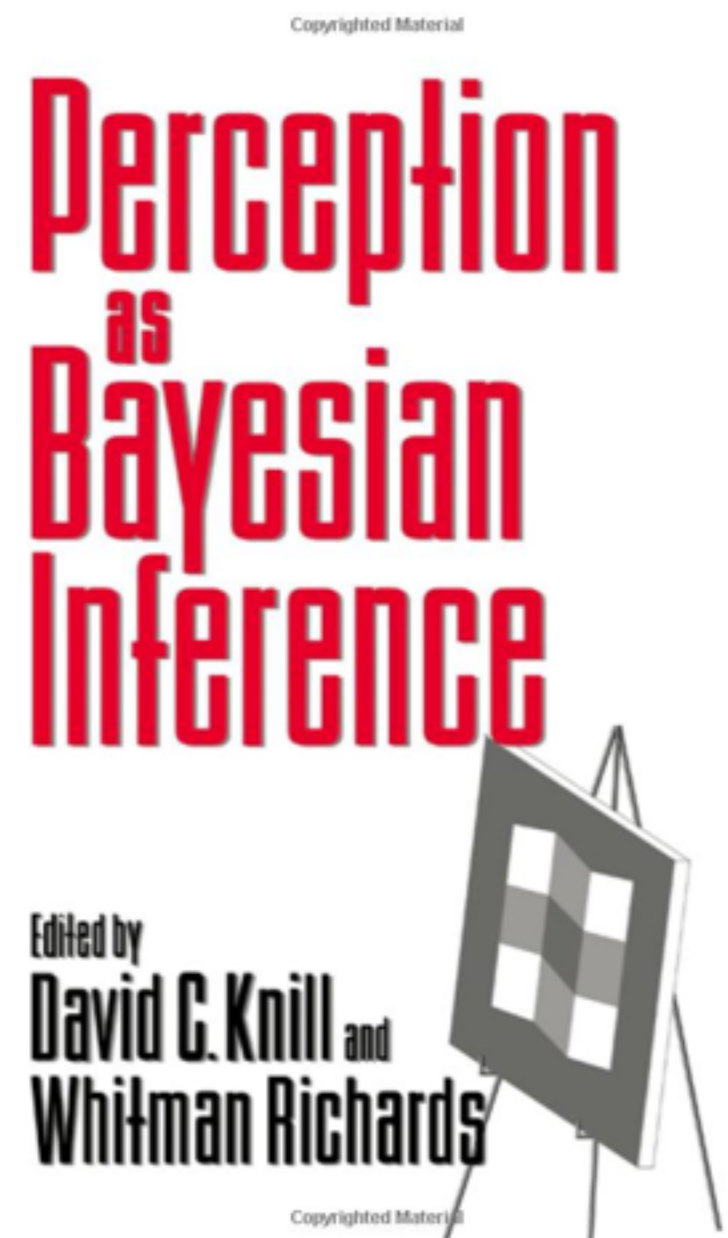

1996

1991-94: Dave the mentor, co-advisor, collaborator

Pascal Mamassian

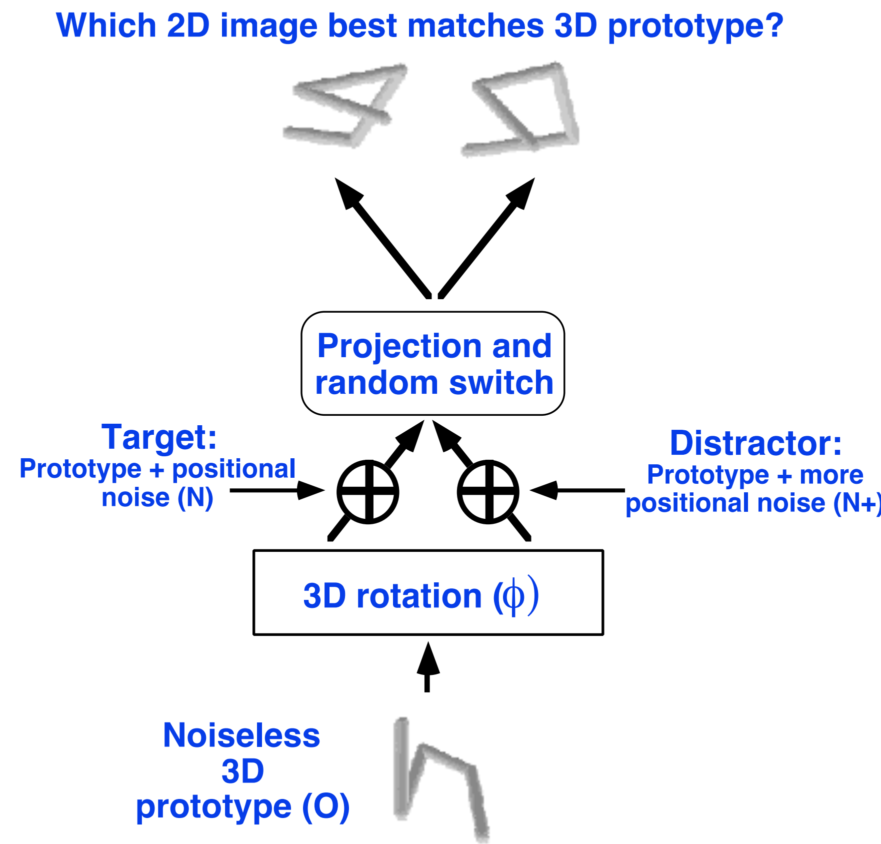

Zili Liu

Knill, D. C., Mamassian, P., & Kersten, D. (1997). Geometry of shadows. JOSA A, 14(12), 3216–3232.

Kersten, D., Knill, D. C., Mamassian, P., & Bülthoff, I. (1996). Illusory motion from shadows. Nature, 379(6560), 31.

Liu, Z., Knill, D. C., & Kersten, D. (1995). Object classification for human and ideal observers. Vision Research, 35(4), 549–568.

..in closing

Dave the problem solver, not an ideologue

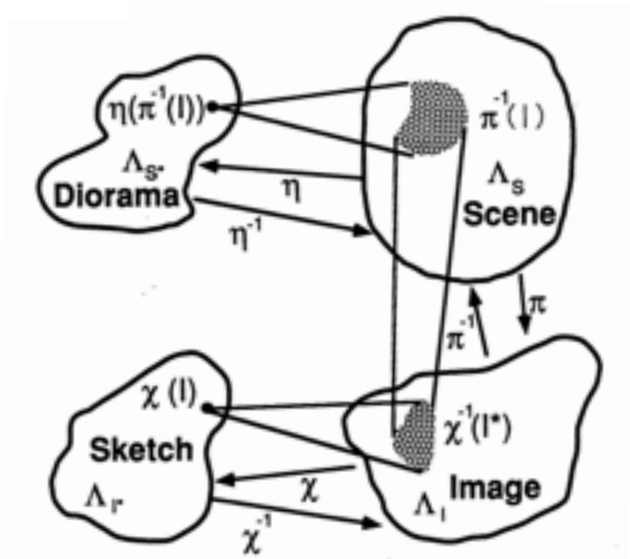

Fig. 7.3 *Schematic including the image-to-sketch map, χ: $\Lambda_I \rightarrow \Lambda_{I^*}$. χ is often many-to-one, and its inverse, one-to-many, as shown here.*

..by 1995, studies leading to some 17 articles, 10 as first or sole author

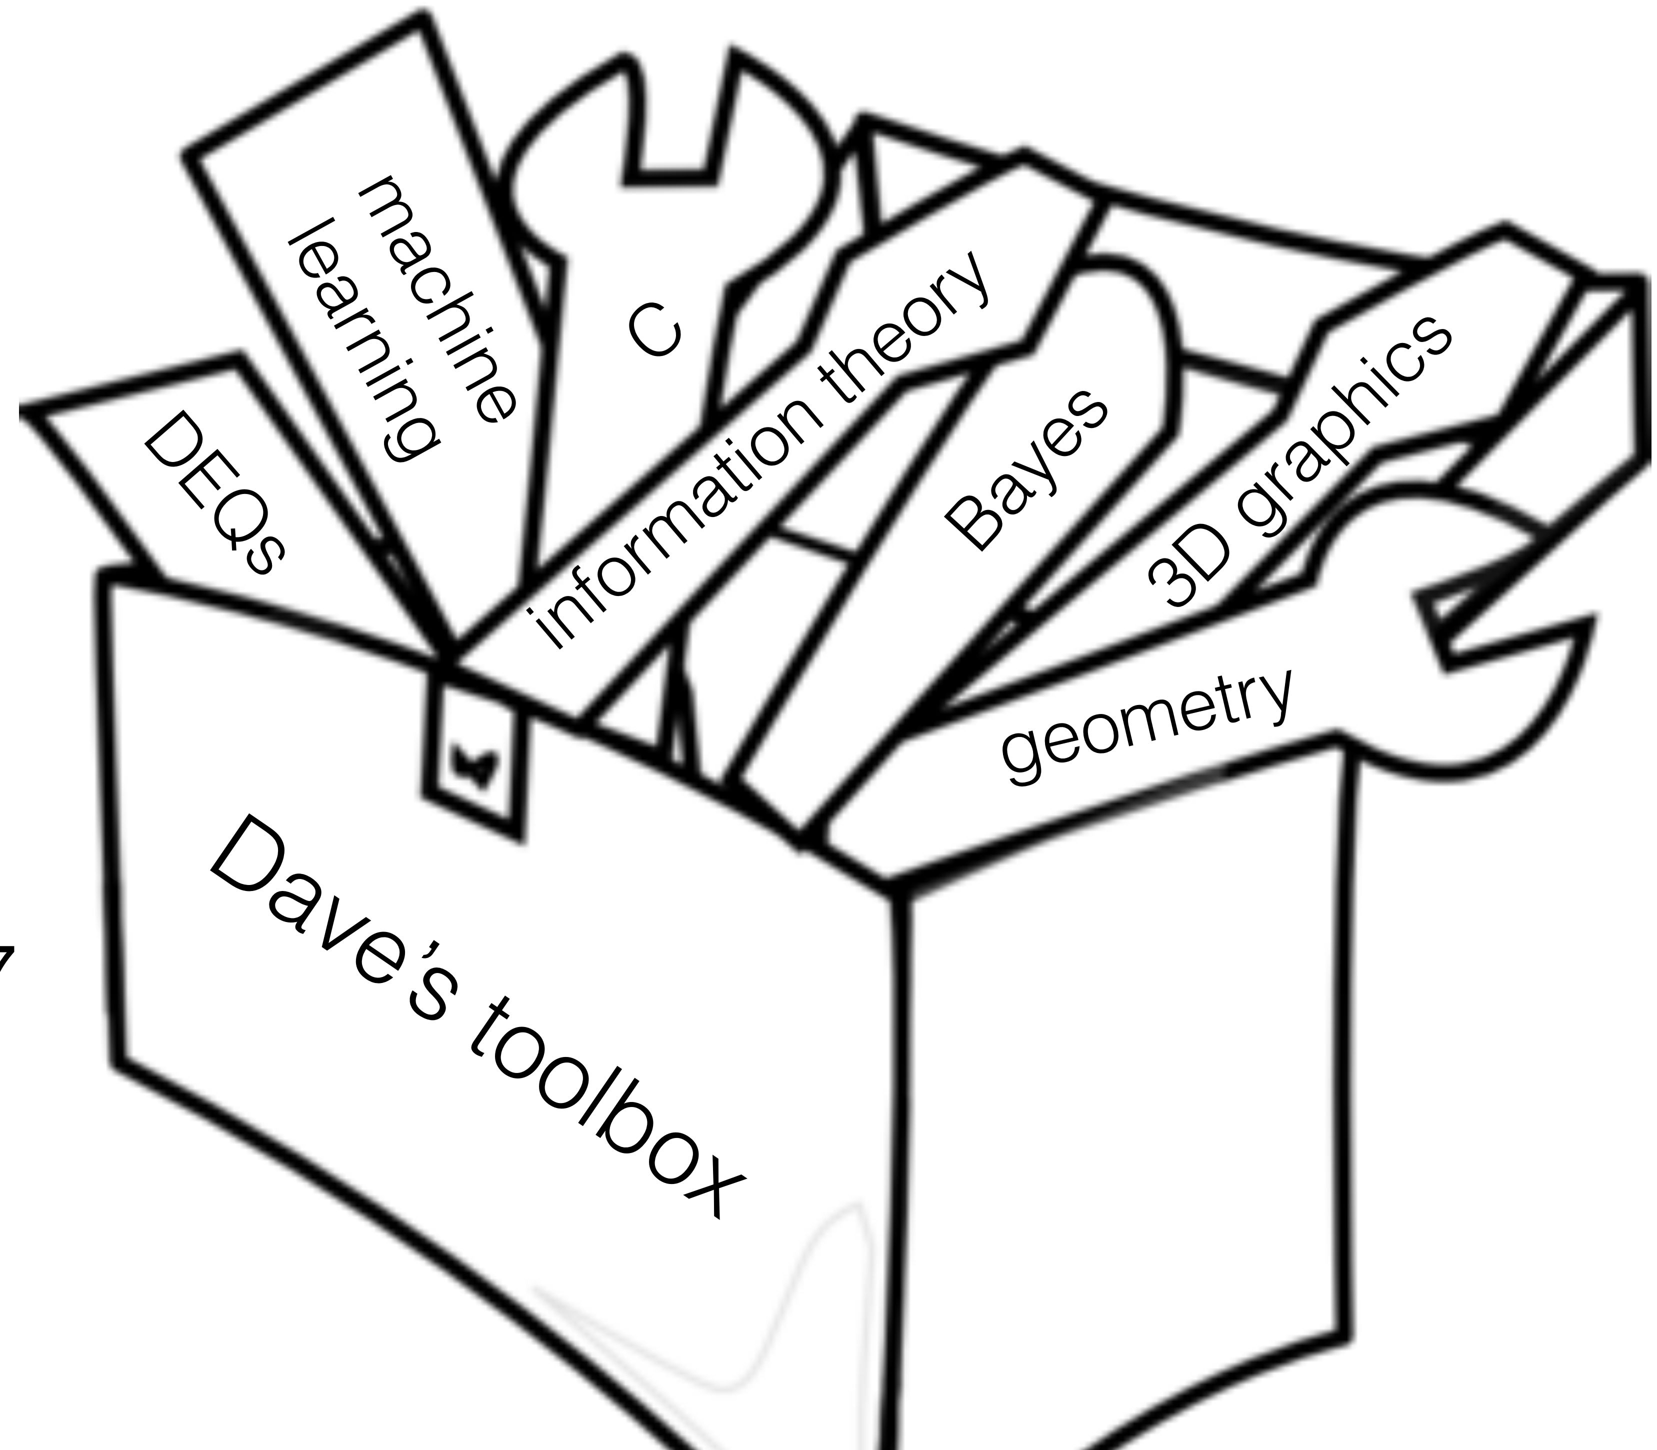

What is Bayesian vision today? …by 1998

No longer as simple as "inverse optics"

Narragansett Bay, Rhode Island — late 1980s

David Knill and the Rational Analysis of mid-level vision

Paul Schrater

Graduate student with Dave, 1994-1998

20 Years ago today

1995 University of Pennsylvania Psych Building

As a joke only we though was funny, all lab members wore that hat for lab profile pics.

Knill became almost synonym for the Bayesian Brain

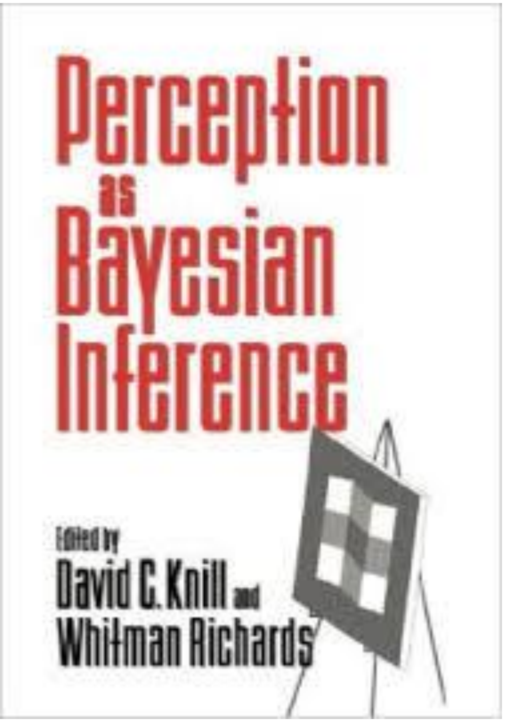

- Ecological perception complex and ***ambiguous***
- Ambiguity generates ***uncertainty*** *which must be handled well to guide actions*

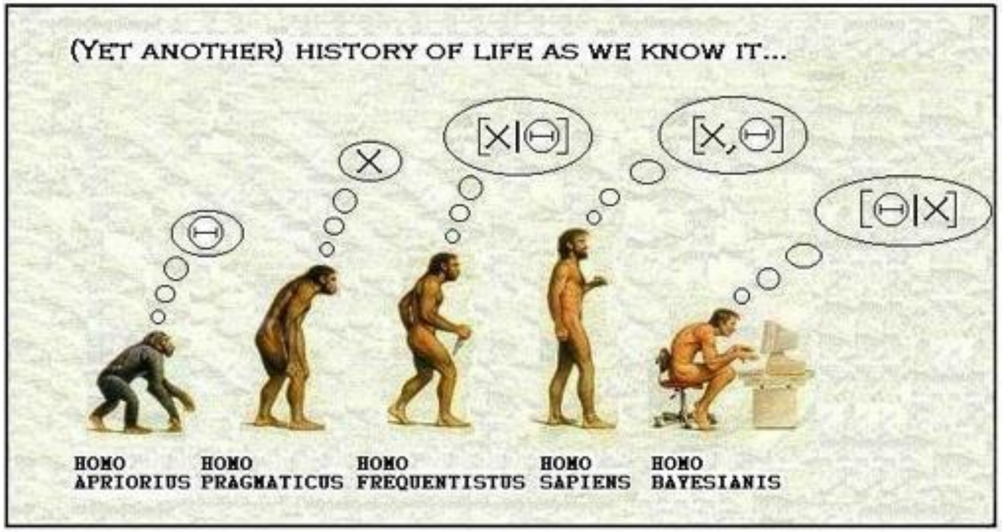

Traditional Levels of Analysis

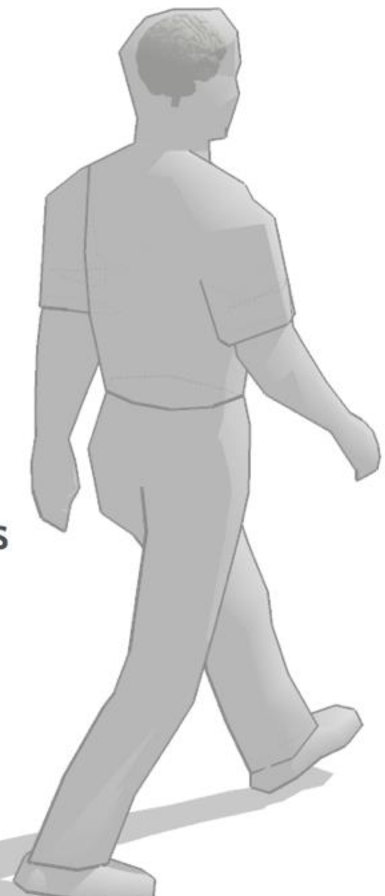

Computational

Why do things work the way they do?
What is the goal of the computation?
What are the unifying principles?

maximize:

$$R_t = r_{t+1} + r_{t+2} + \cdots + r_T$$

Algorthmic

What representations can implement such computations?
How does the choice of representations determine the algorithm?

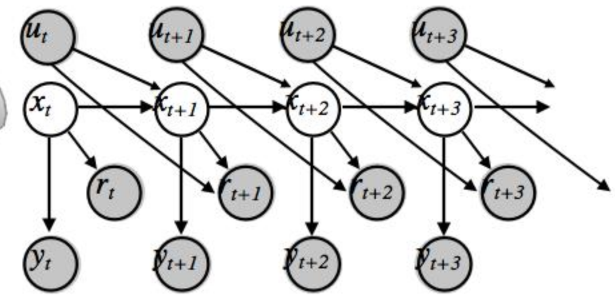

Implementational

How can such a system be built in hardware?
How can neurons carry out the computations?

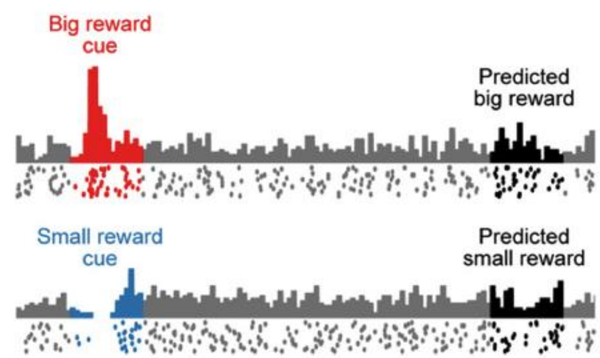

Reductionist Explanations

Craik O'Brien Cornsweet illusion

Illusion results from **byproduct** of early sensory processing

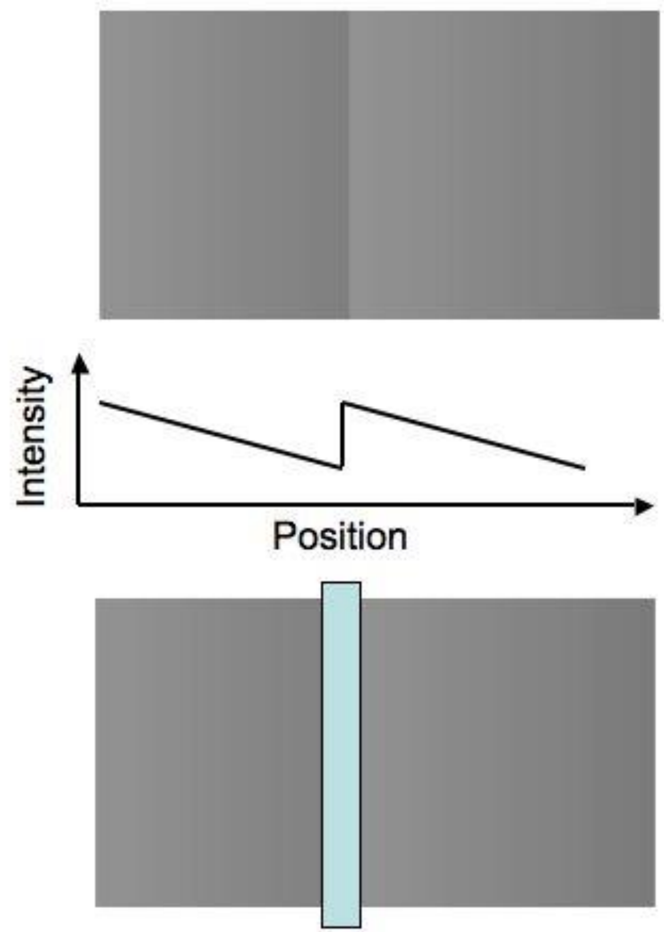

Computational Epistomology

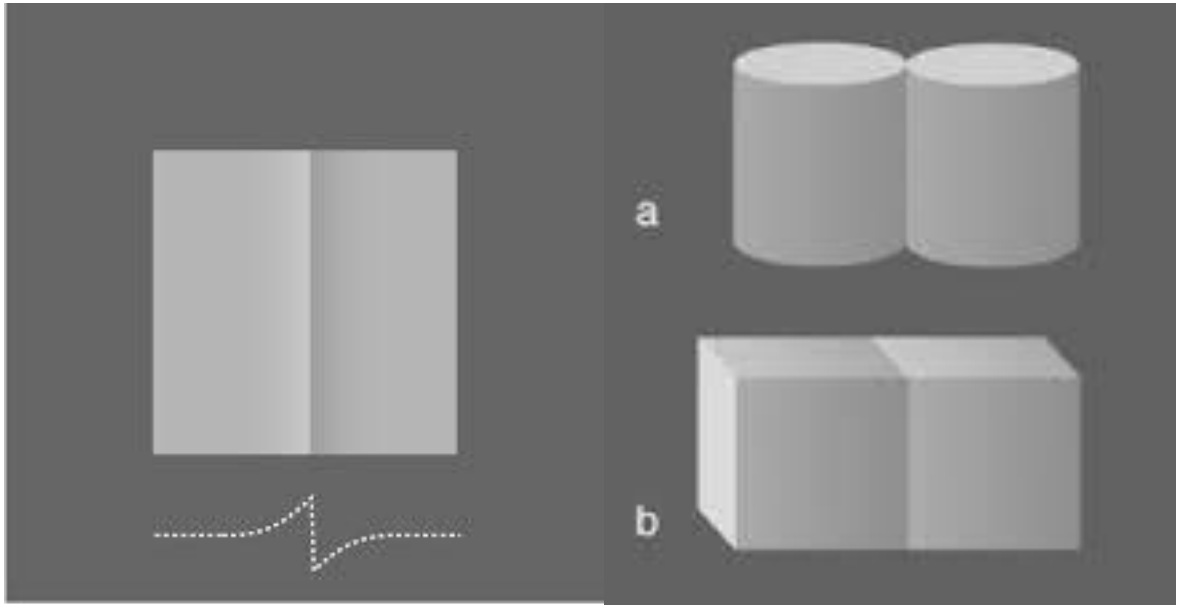

Illusion results from **rational analysis** of the scene

Dave was not Anti-Reductionist

- But some complain that a Bayesian approach is "vague" and of "questionable merit"

'While Marr's original attack on reductionism was justified it is no longer tenable'

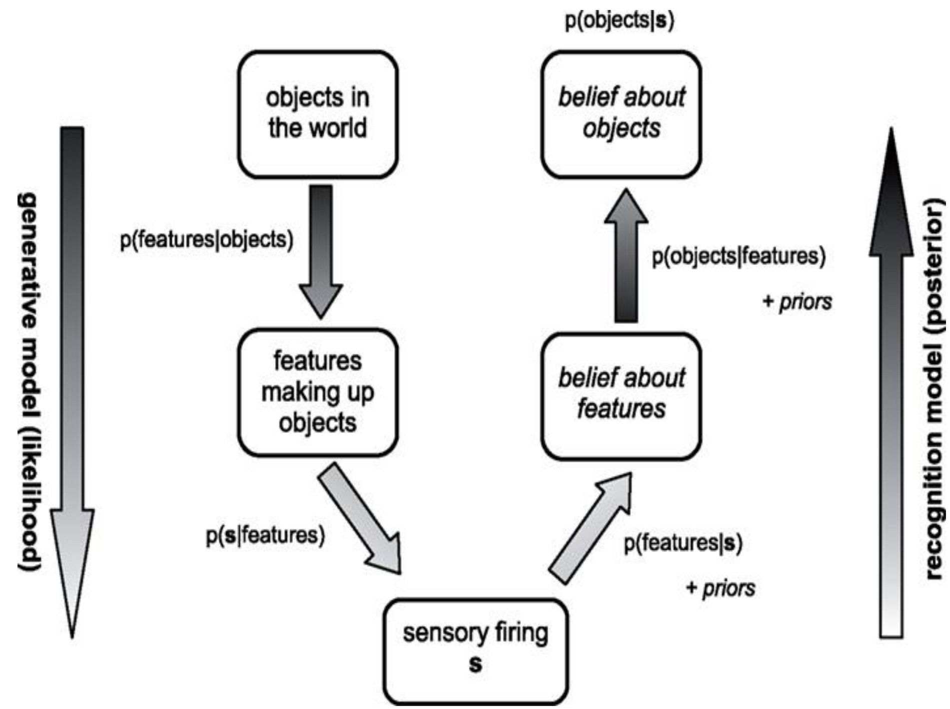

Some properties of a useful psychophysical framework

- Support building predictive models of perceptual performance.
- Support bridging statements between models and descriptions of behavior.
- Explain "why" perception / sensorimotor control works the way it does.
- Help guide psychophysical research
  - Suggests new and interesting theoretical questions.
  - Supports scaling *down* perceptual / sensorimotor problems to bring them into the lab.
  - Scales *up* naturally

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

World model

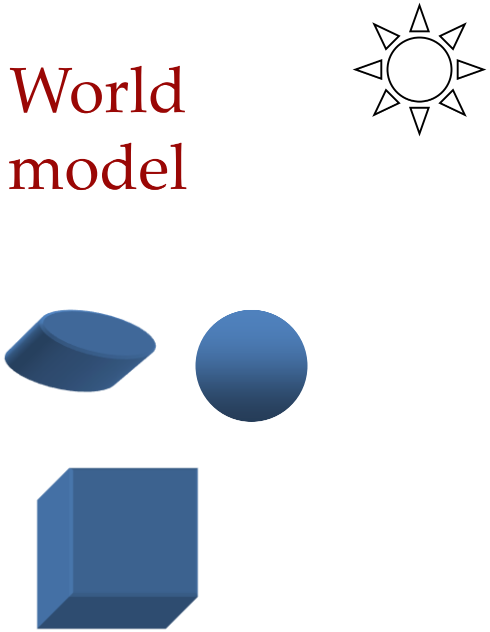

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

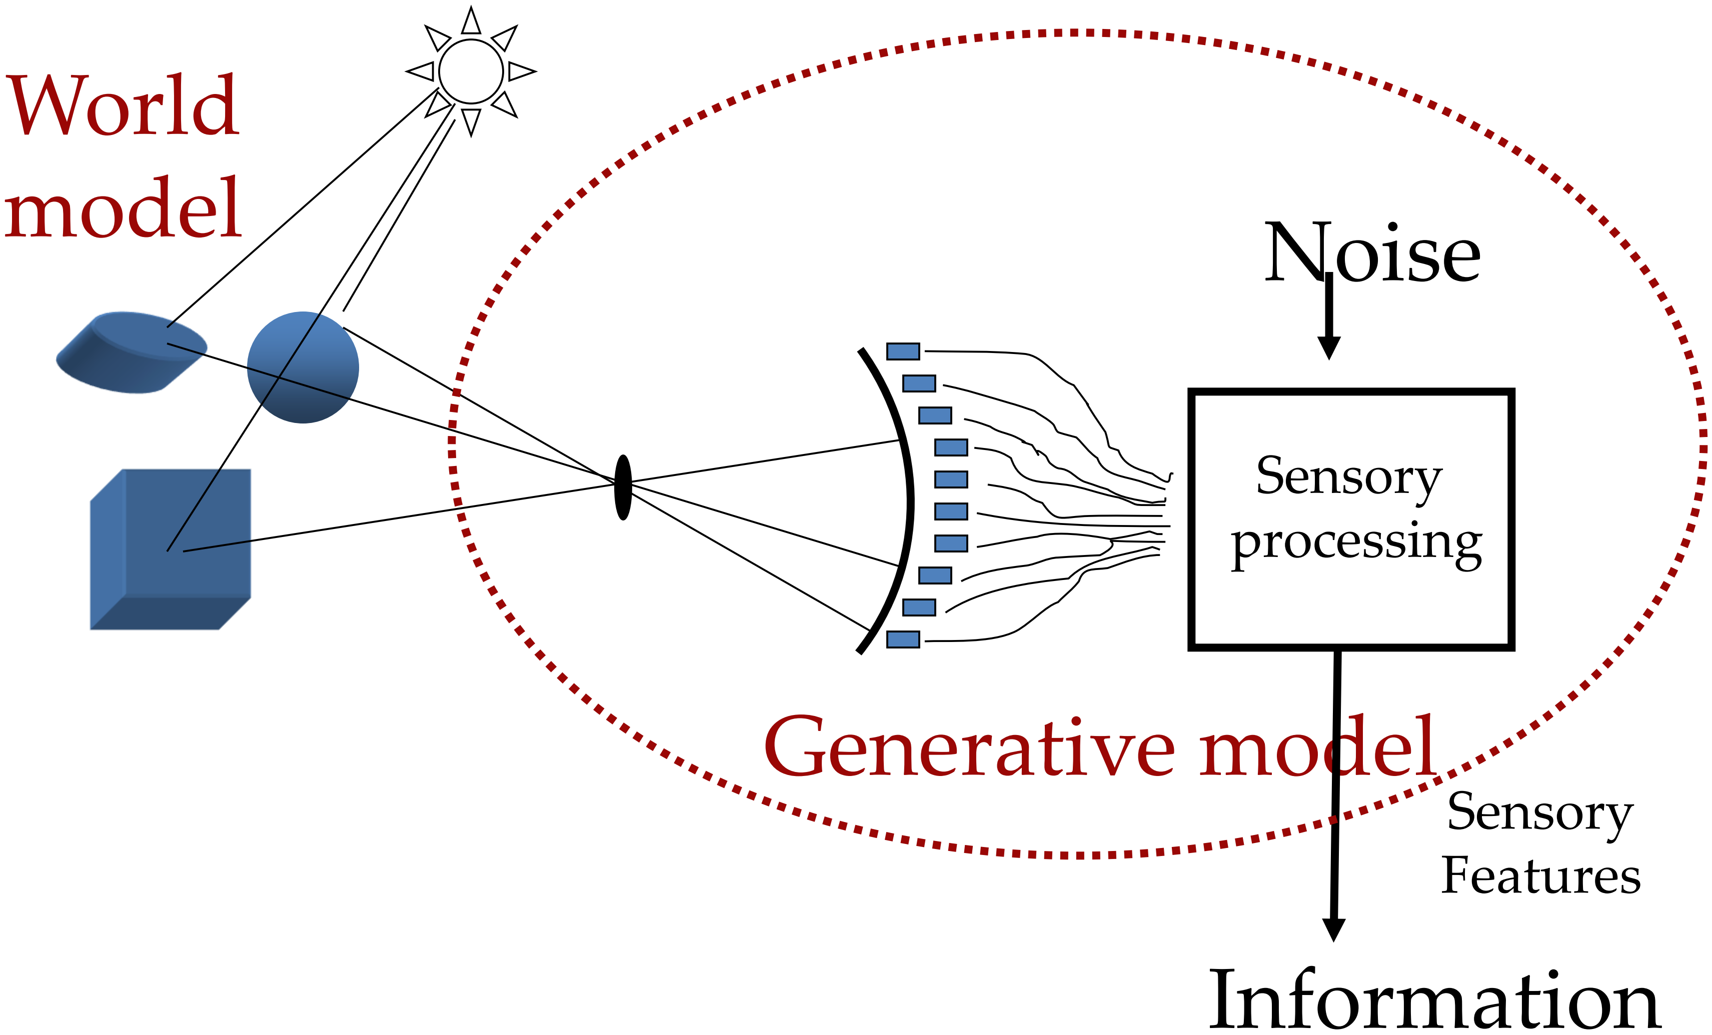

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

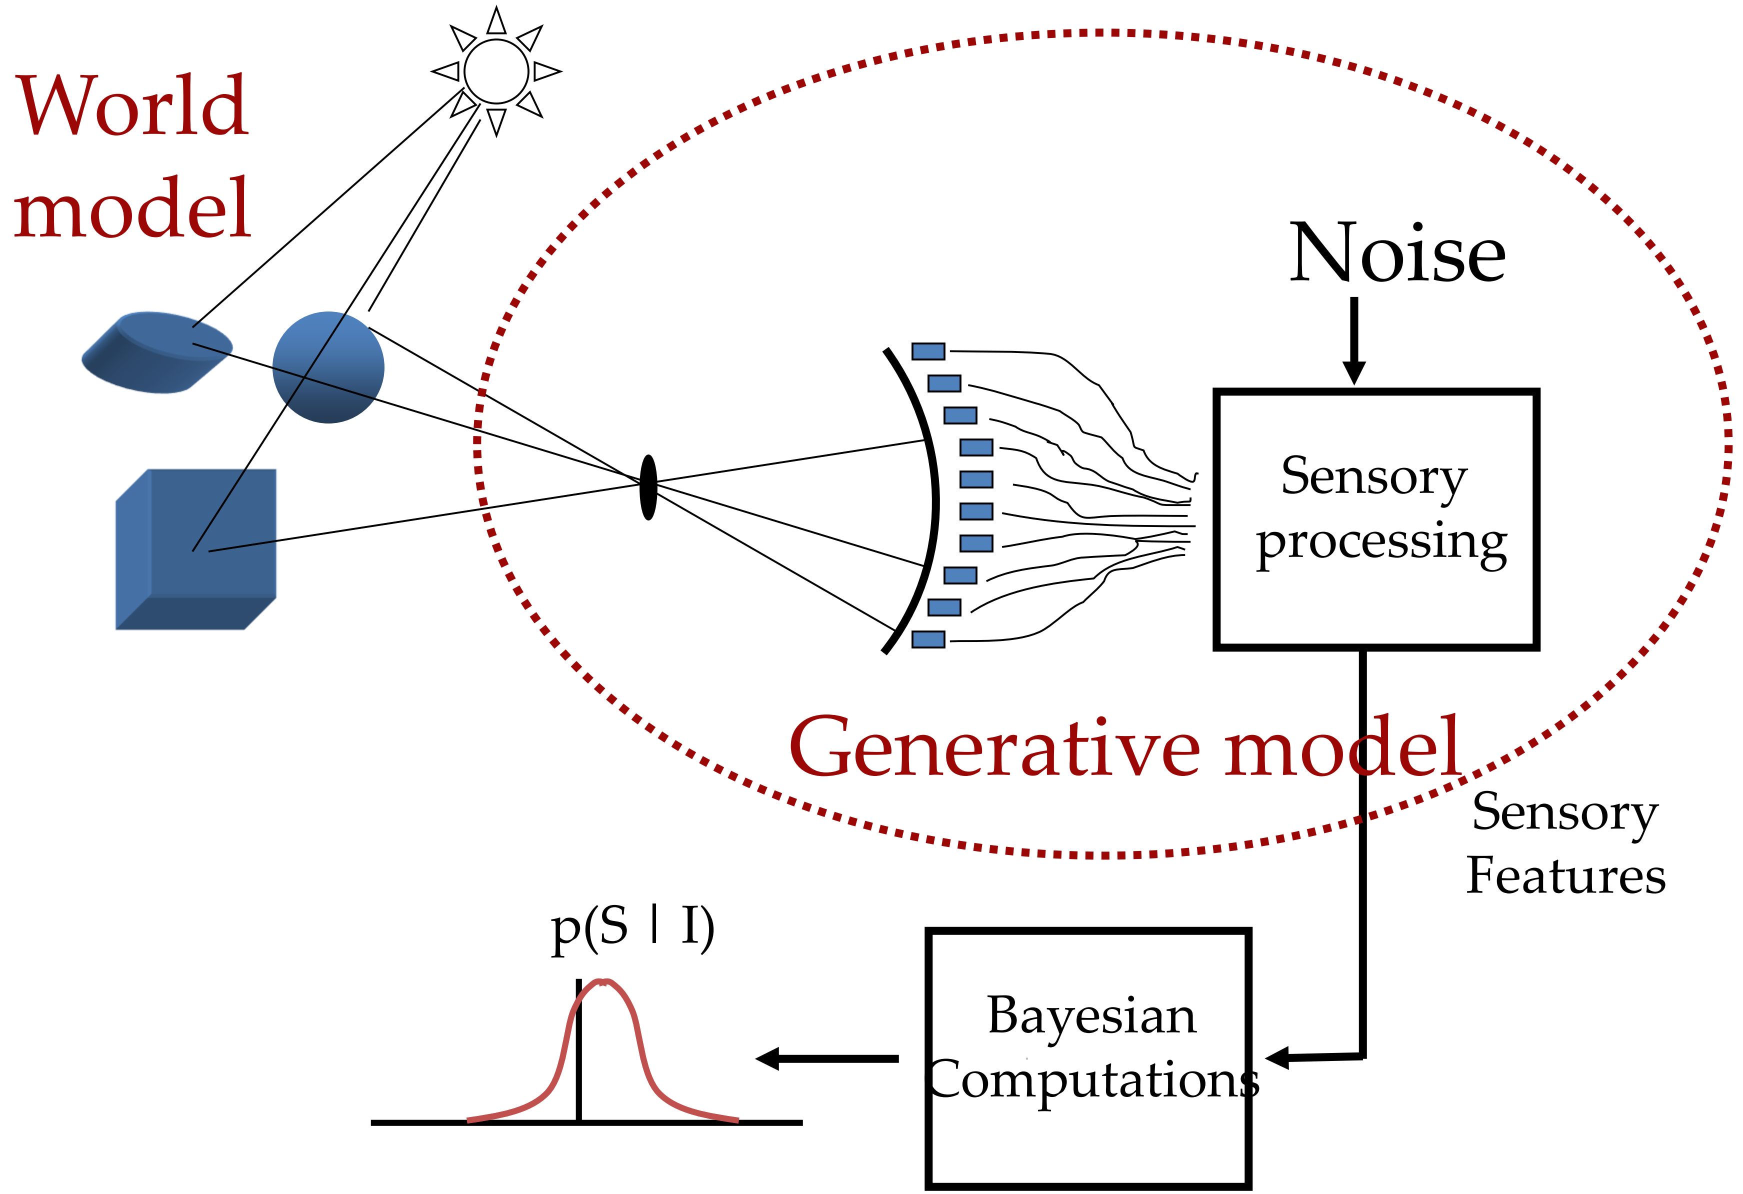

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

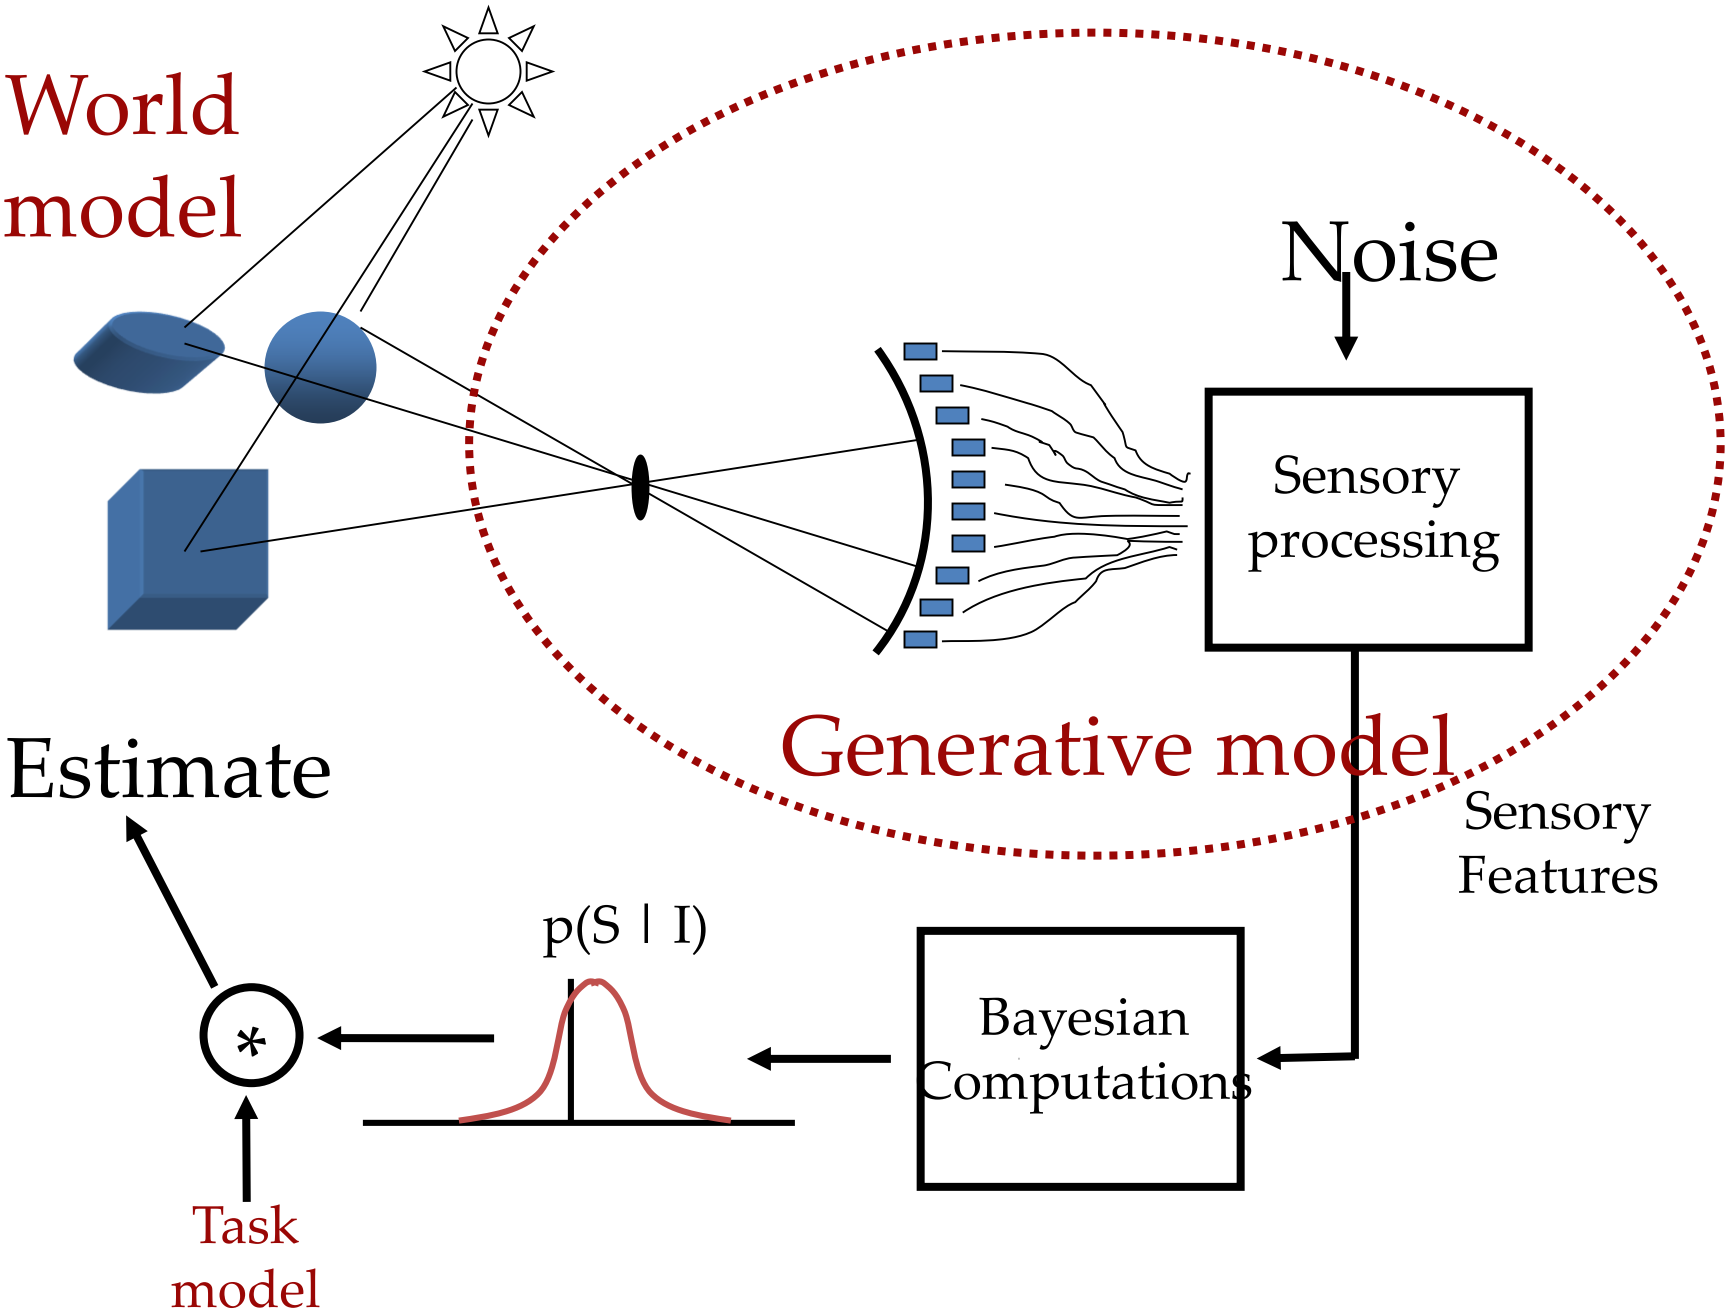

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

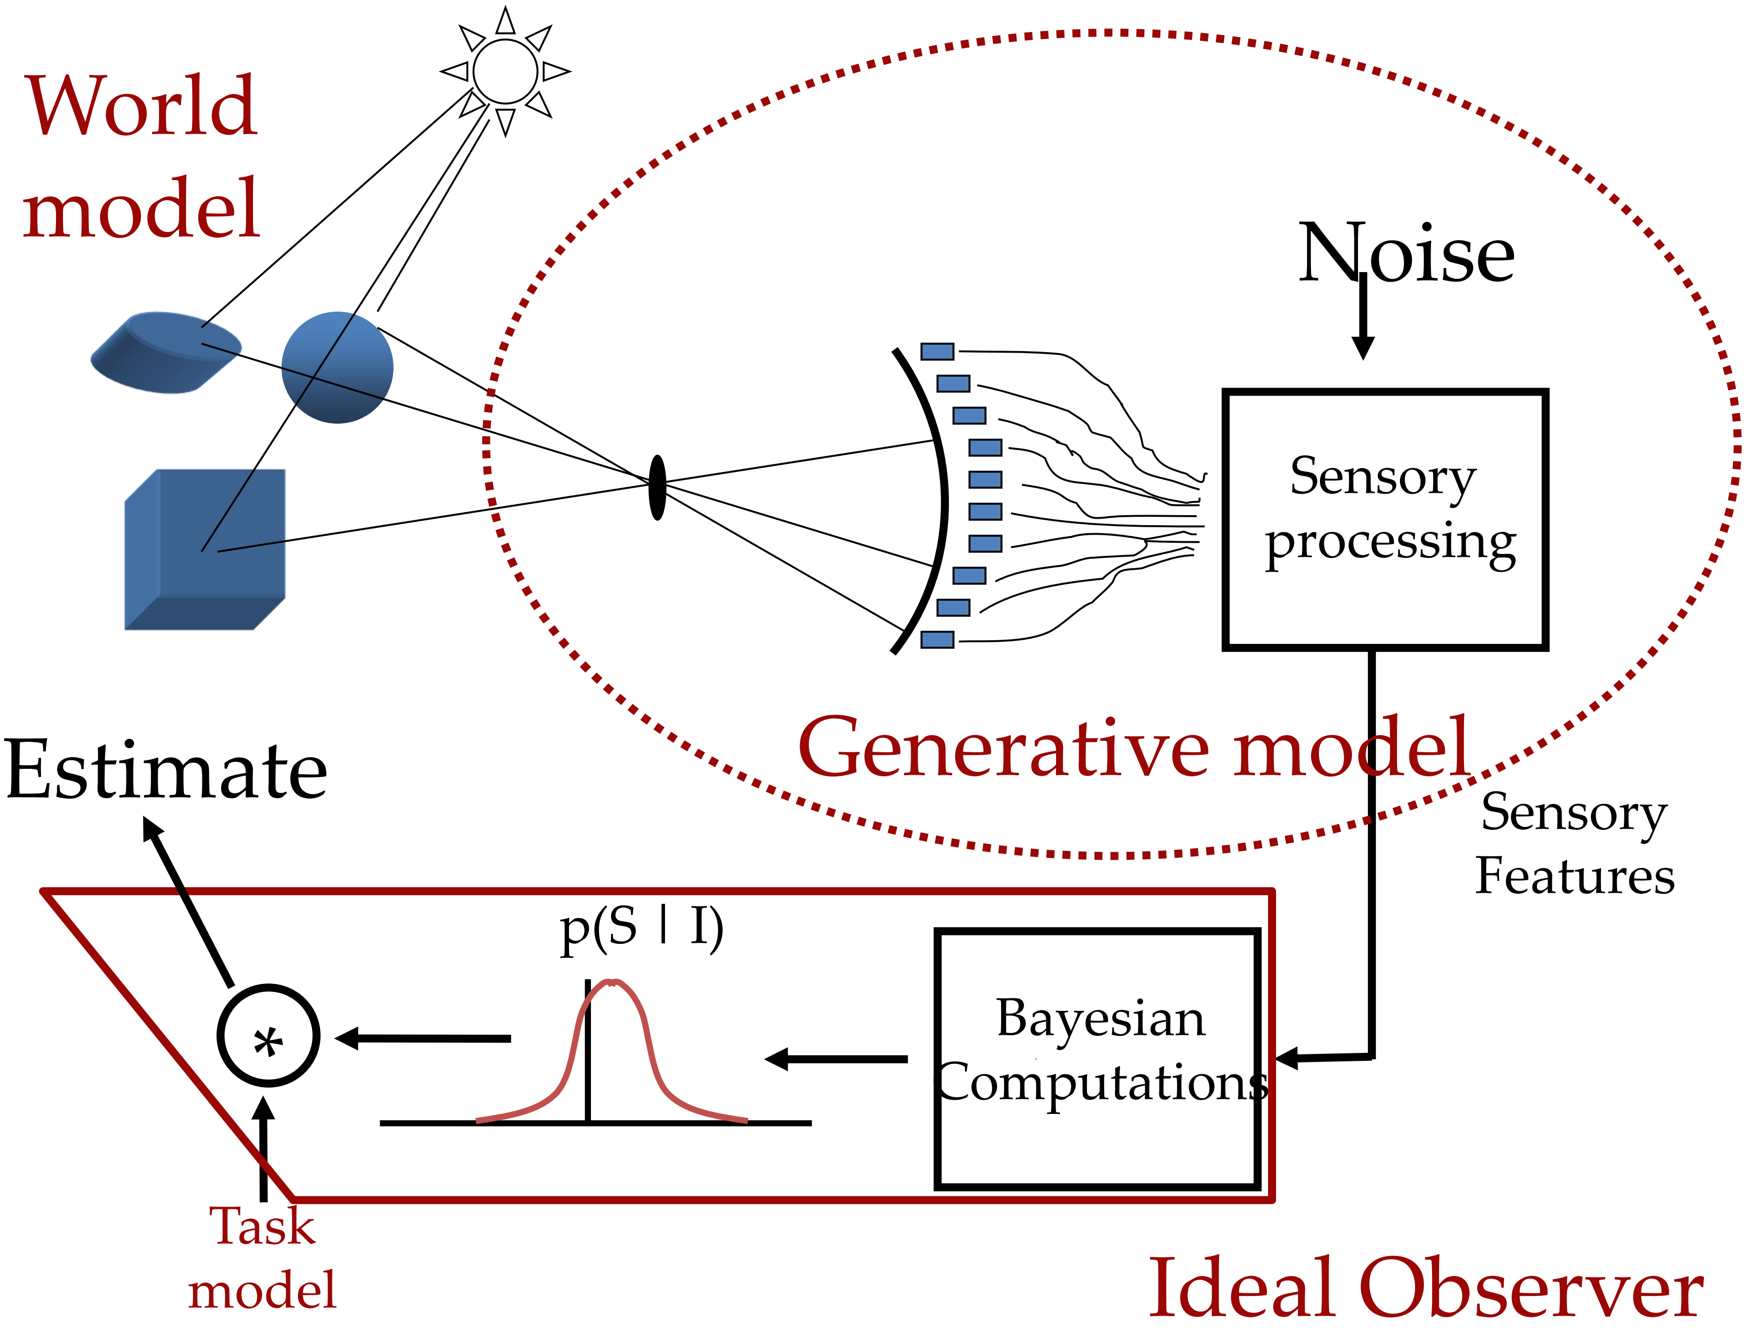

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

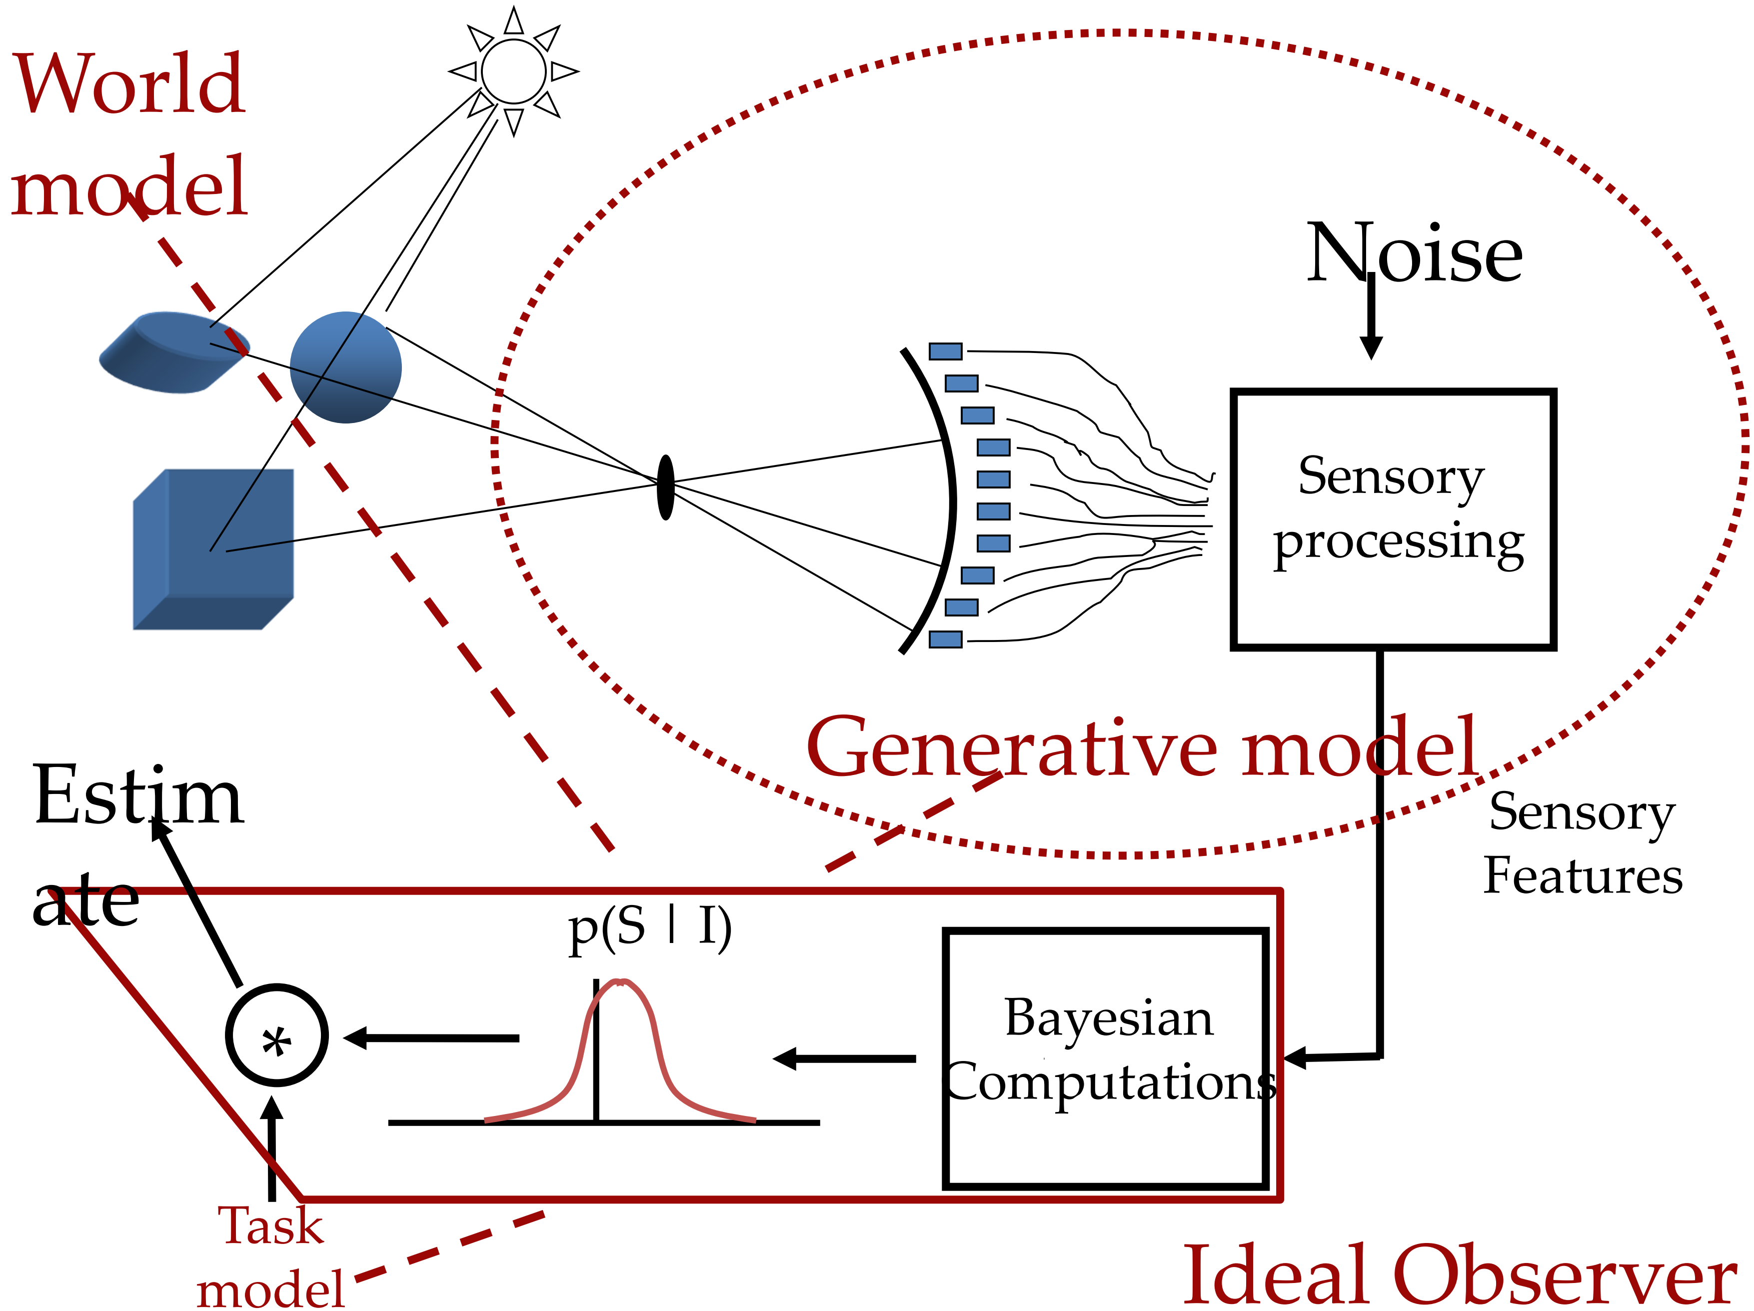

Slide from Dave's talk at Battaglia and Schrater's 1998 VSS symposium

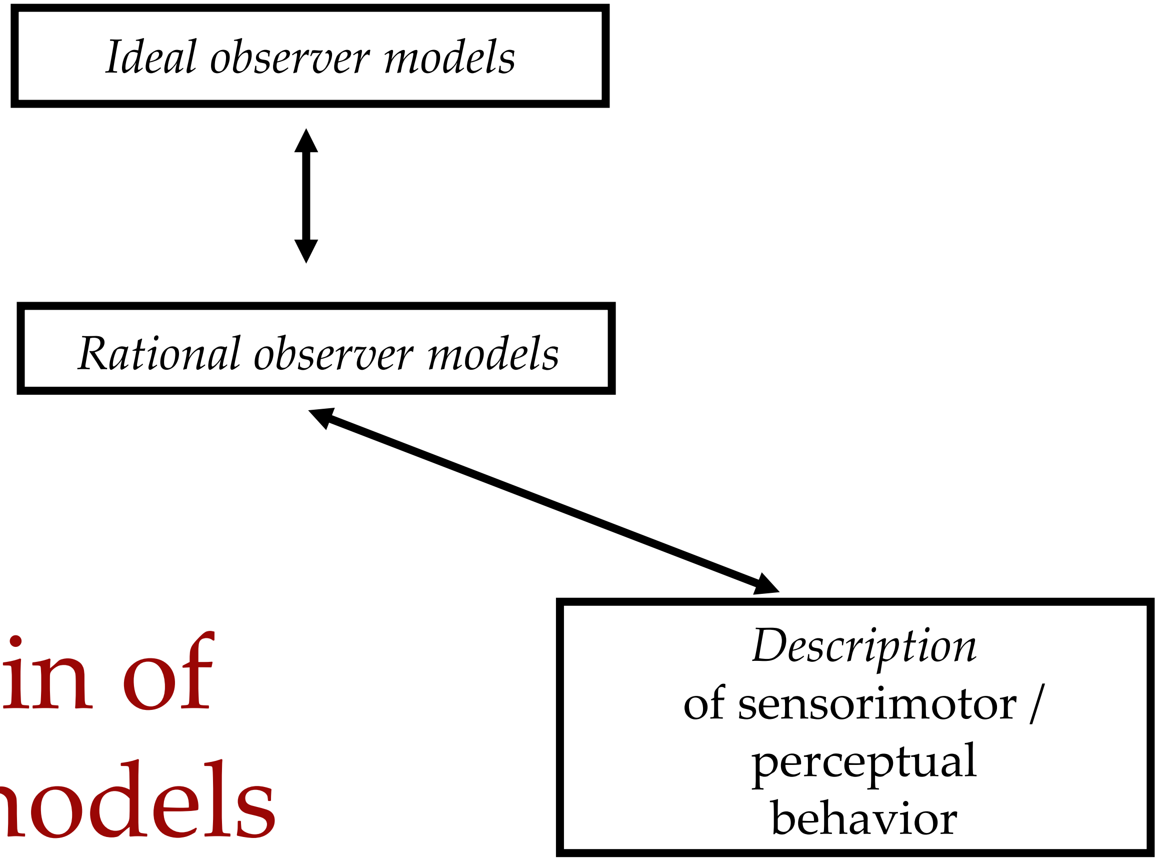

The domain of Bayesian models

Threefold Knowledge

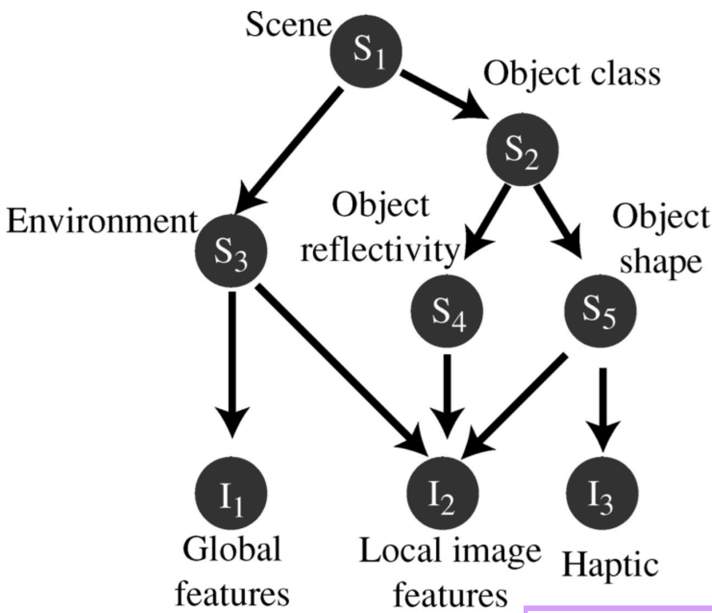

1) Image physics

2) Environmental regularities

3) Human task requirements

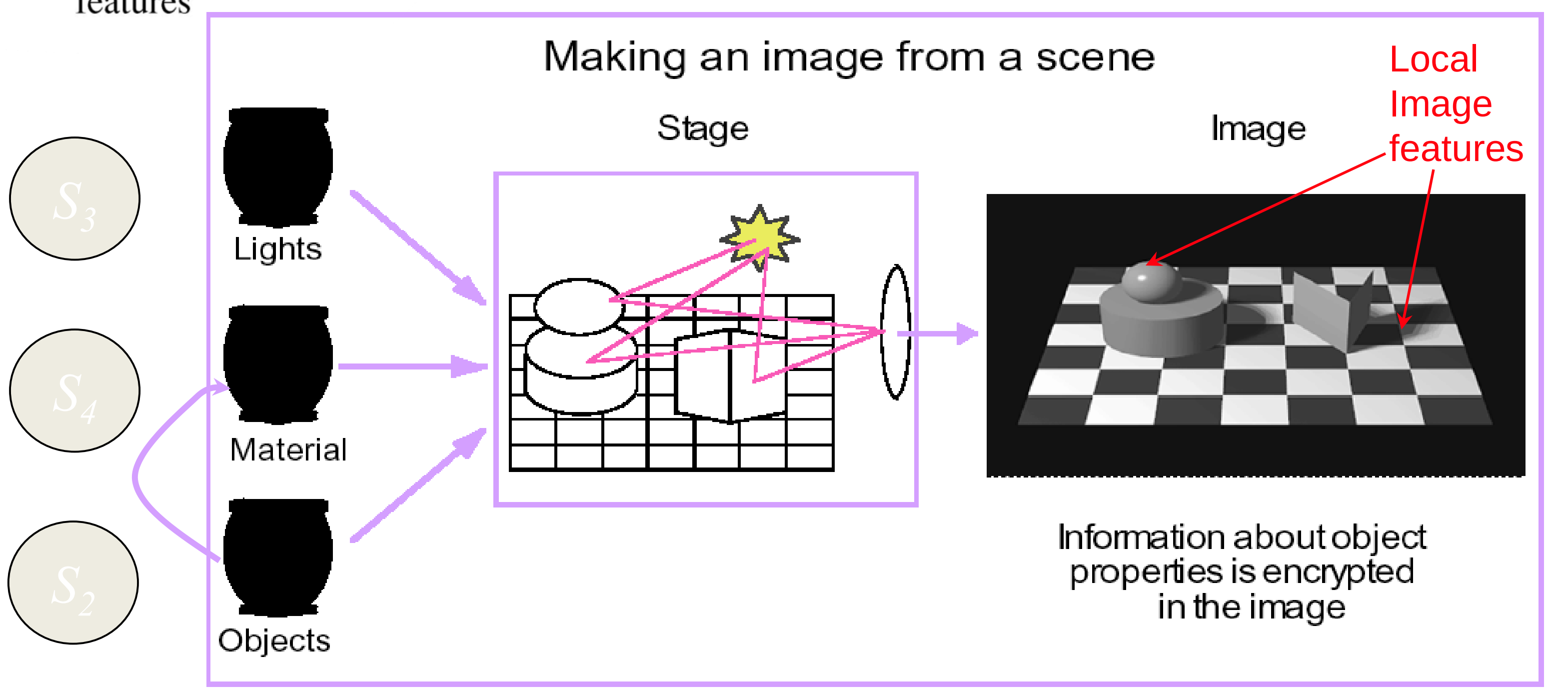

Generative model

- Sample a scene type
- Sample N object classes
- Sample Objects from each class (locations and attributes for each object)
- Sample rendering variables (lights, viewpoint)
- Sample image features from rendered scene

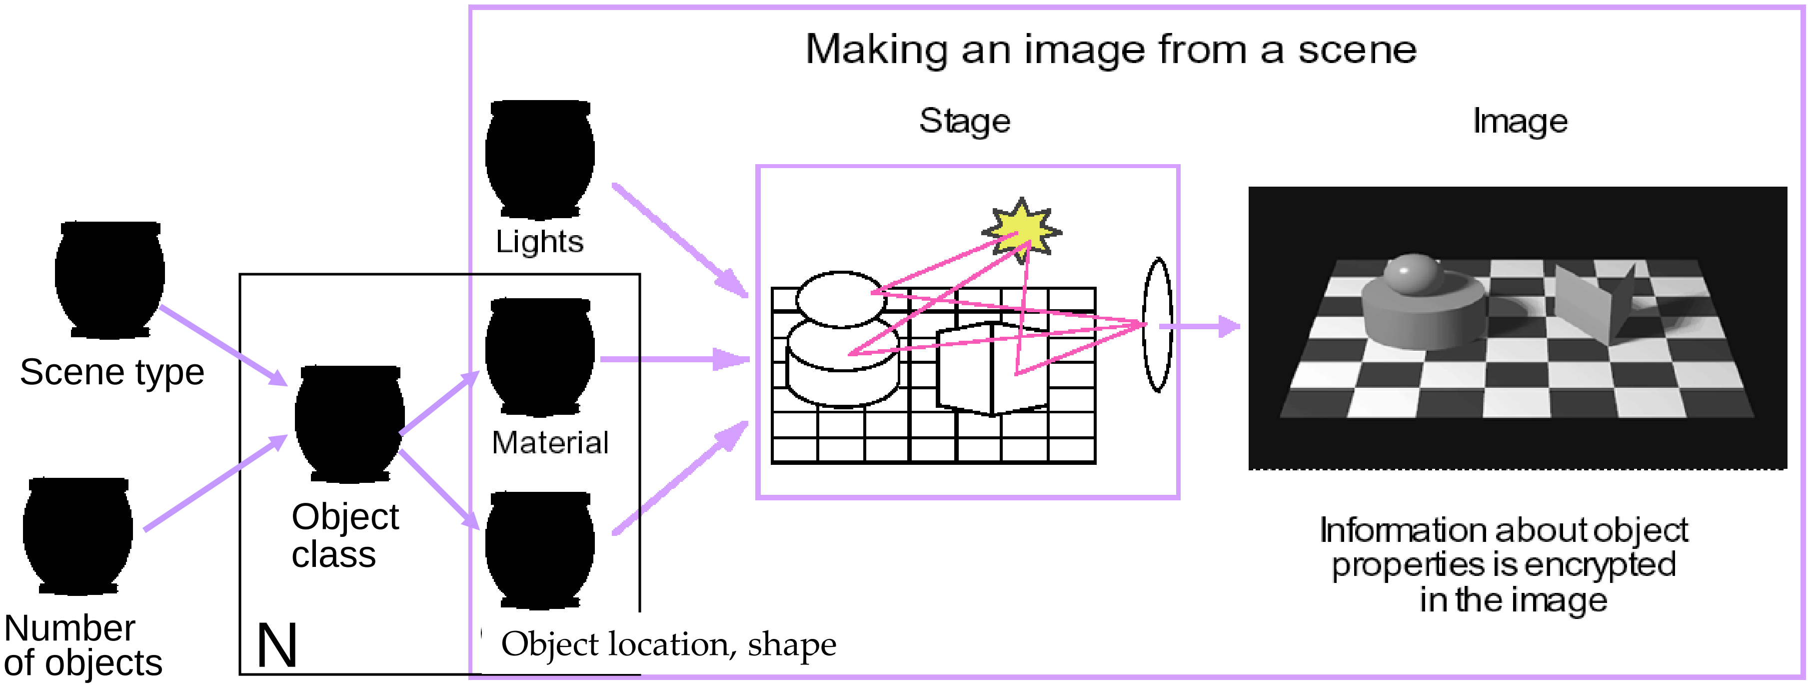

Rational Analysis for Mid-level vision

- What are the evolutionary pressures and environmental features that shape perception?
- These lead to a ***family of computational problems***
  - Natural visual tasks and behavior
    - Getting reliable estimates of object geometry and material

  - Statistical structure of the environment
    - What regularities can be exploited?

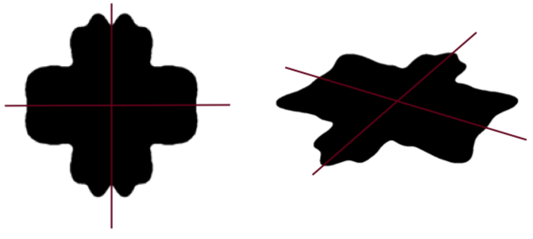

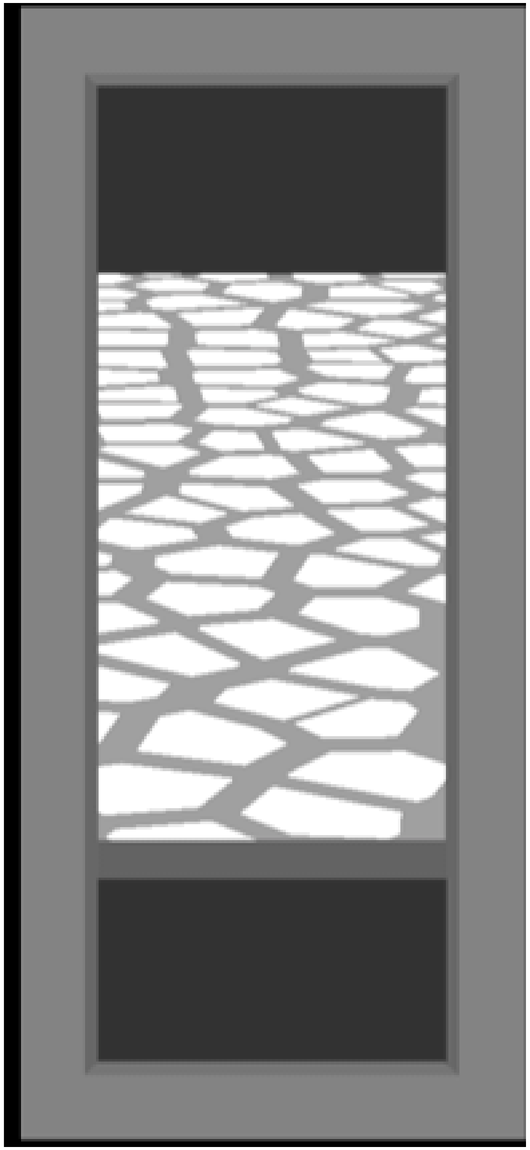
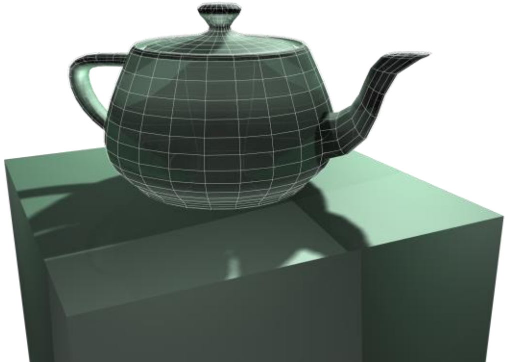

Ideal observer analysis

How well can ***any*** observer compute surface orientation from texture?

"Surface orientation from texture: ideal observers, generic observers and the information content of texture cues" Vision Research, 1998

Knowledge needed

- Geometric

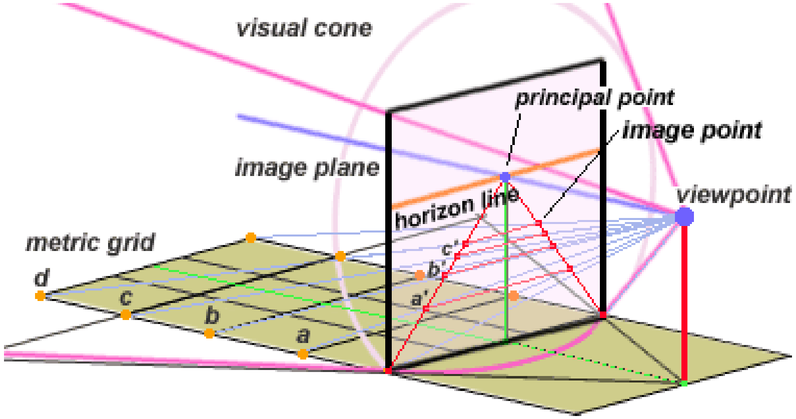

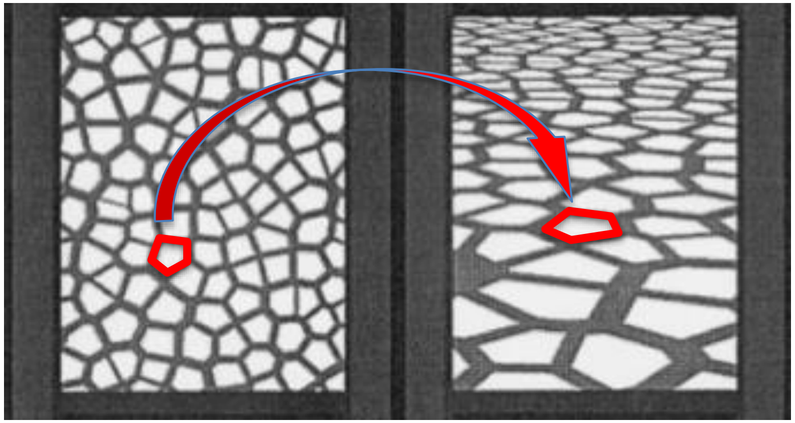

Knowledge needed

- Geometric

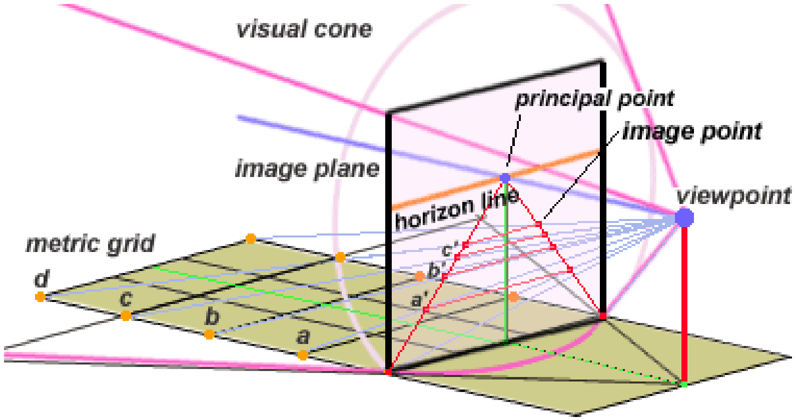

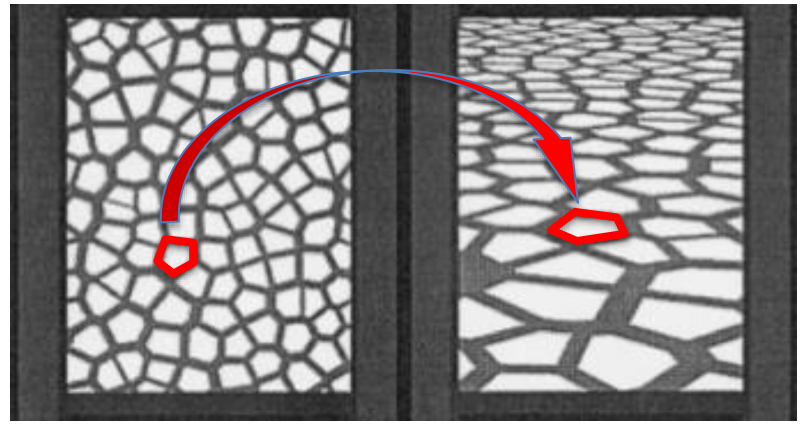

- Statistical

Homogeneity/isotropy

Inhomogeneity

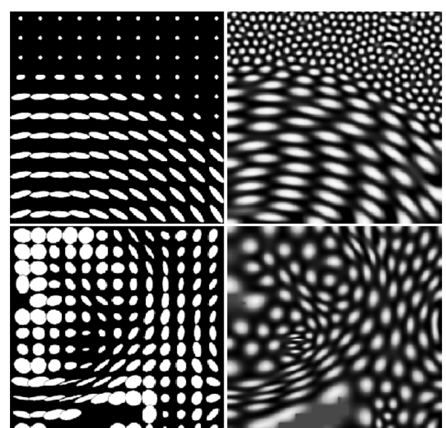

(a) Ellipse Array (b) Reaction-Diffusion

anisotropy

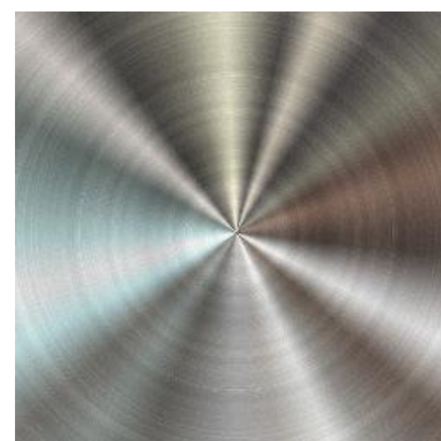

Defining "Cues"

Elliptical Approximation

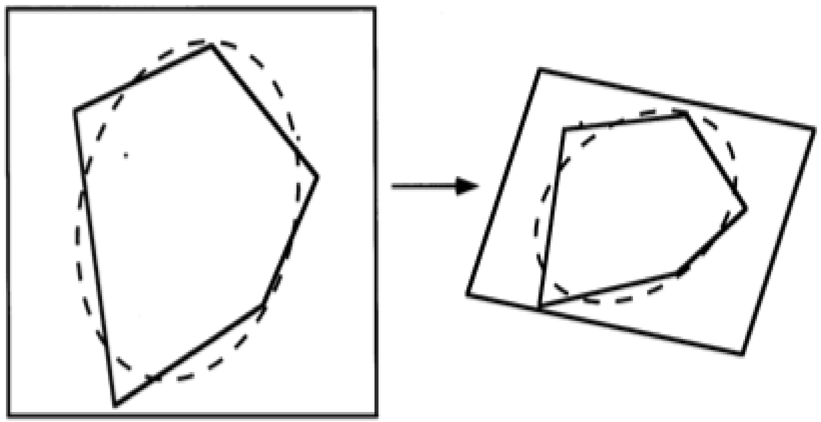

Defining "Cues"

Elliptical Approximation

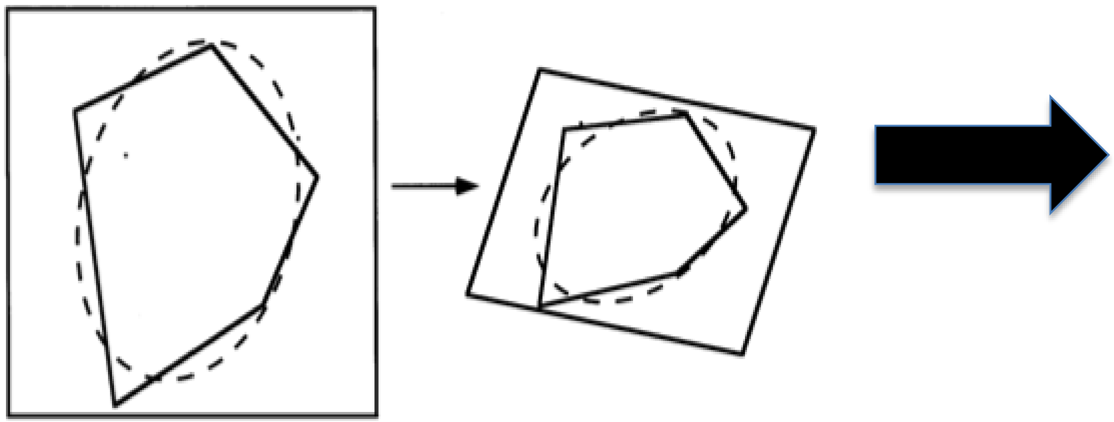

3 parameters

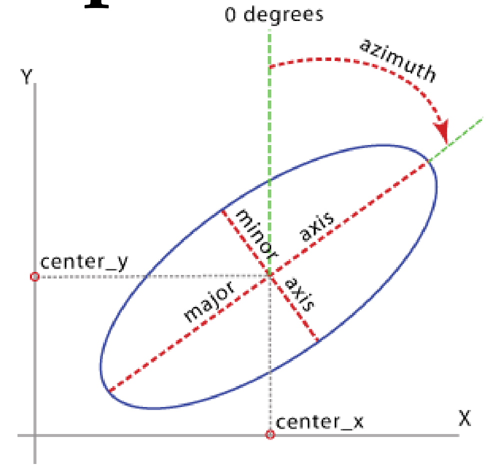

Defining "Cues"

Elliptical Approximation

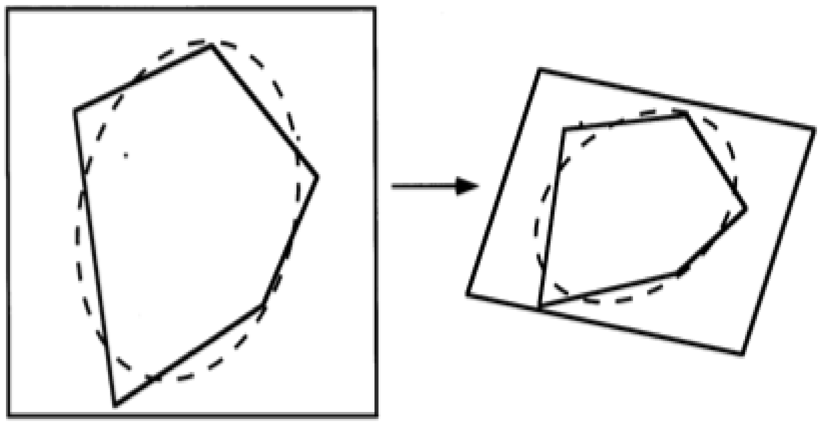

3 parameters

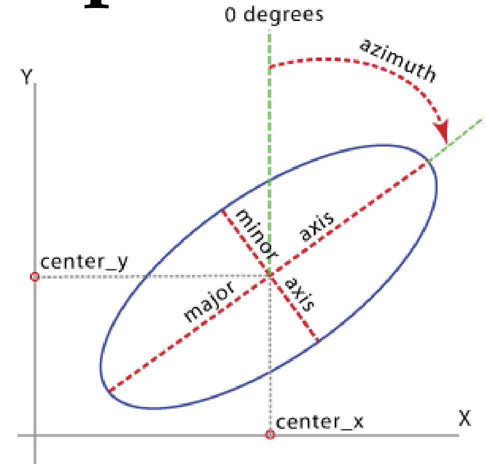

... Per element

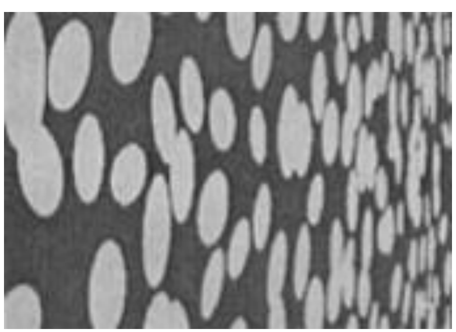

Defining "Cues"

Elliptical Approximation

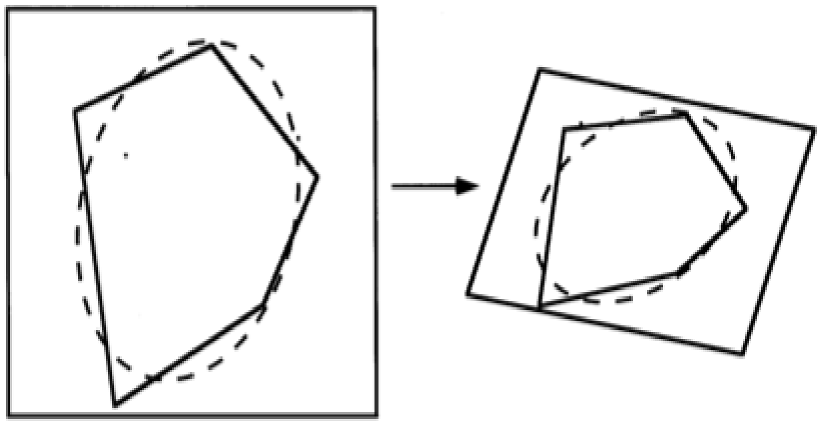

3 parameters

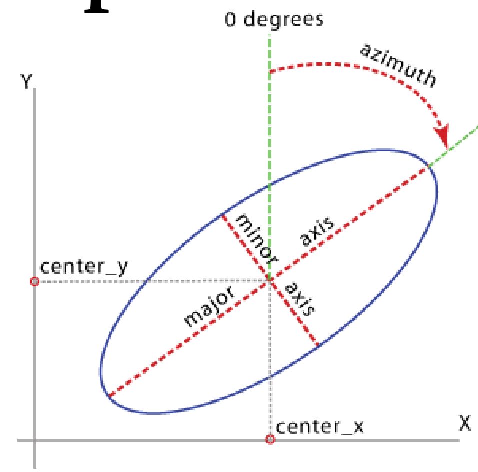

... Per element

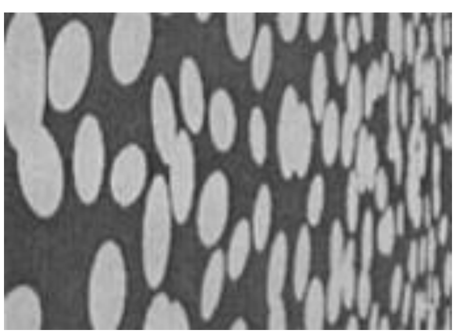

Add spatial gradient

Voila! Cues!

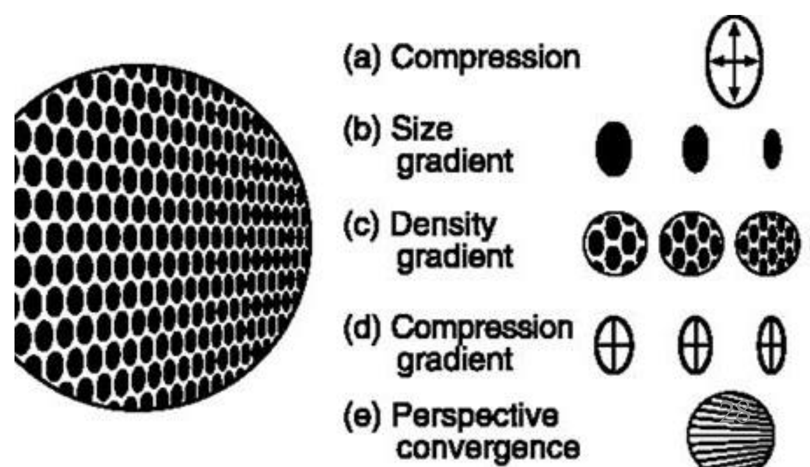

Ideal observer vs. Human

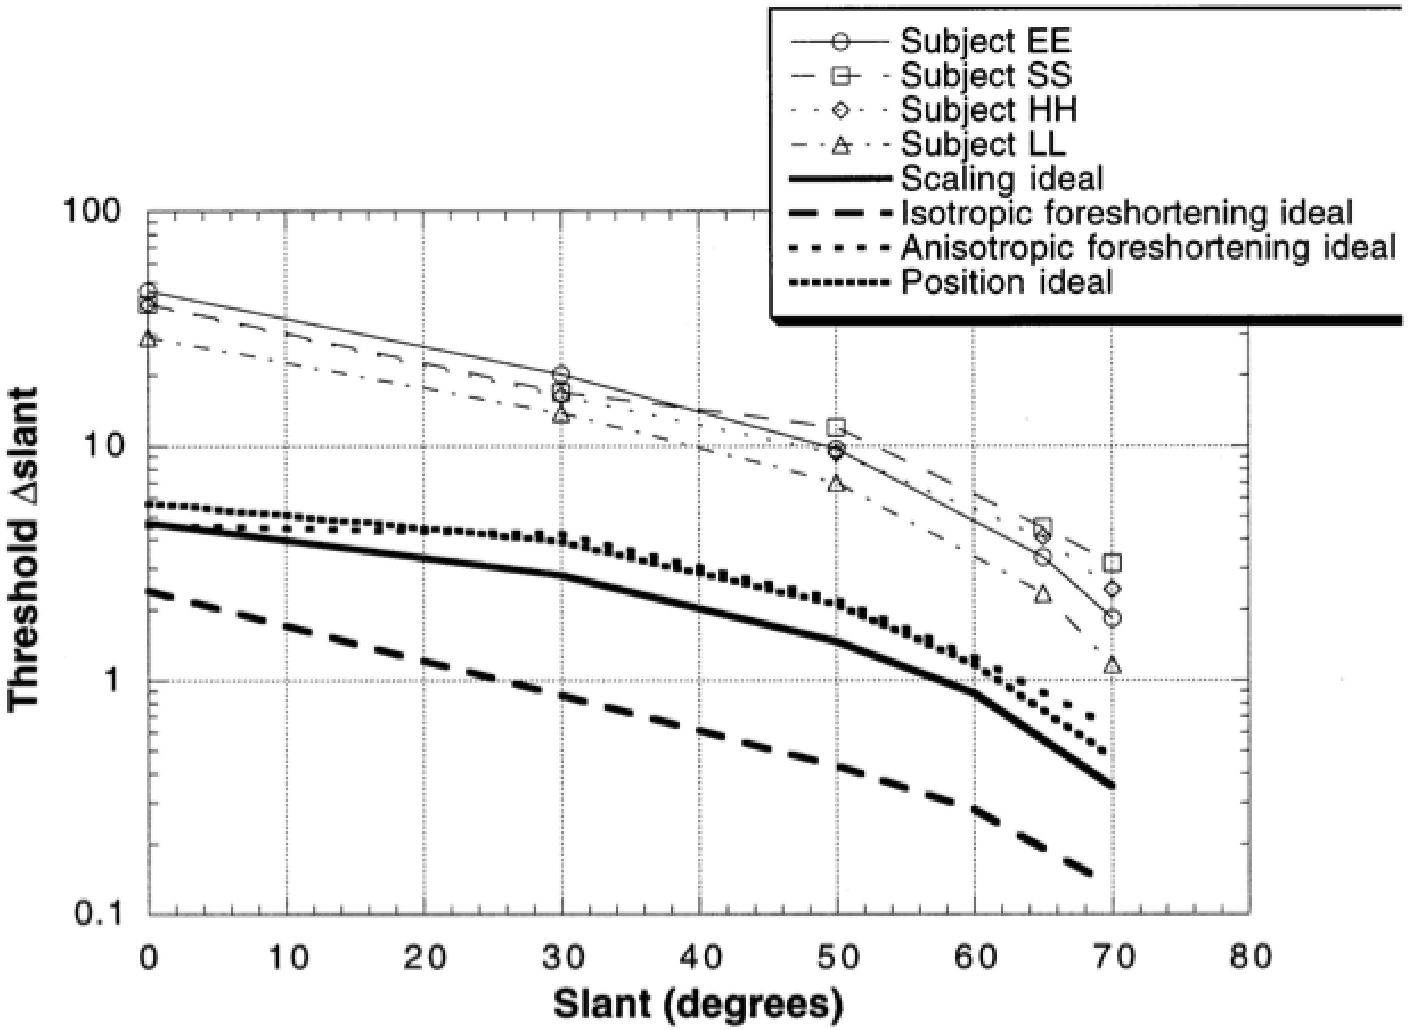

People kind of suck, but that's expected!

Why so BAD?

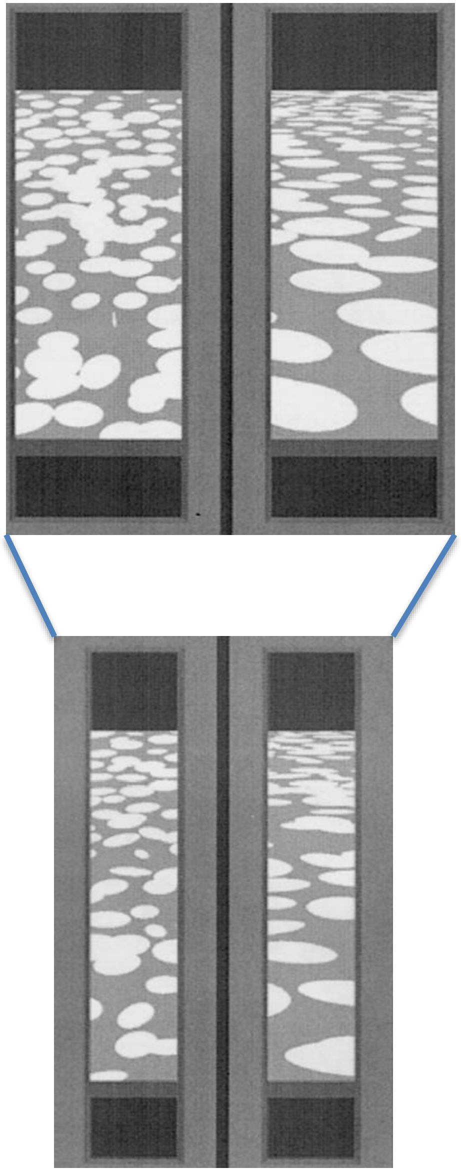

Vary Field of View

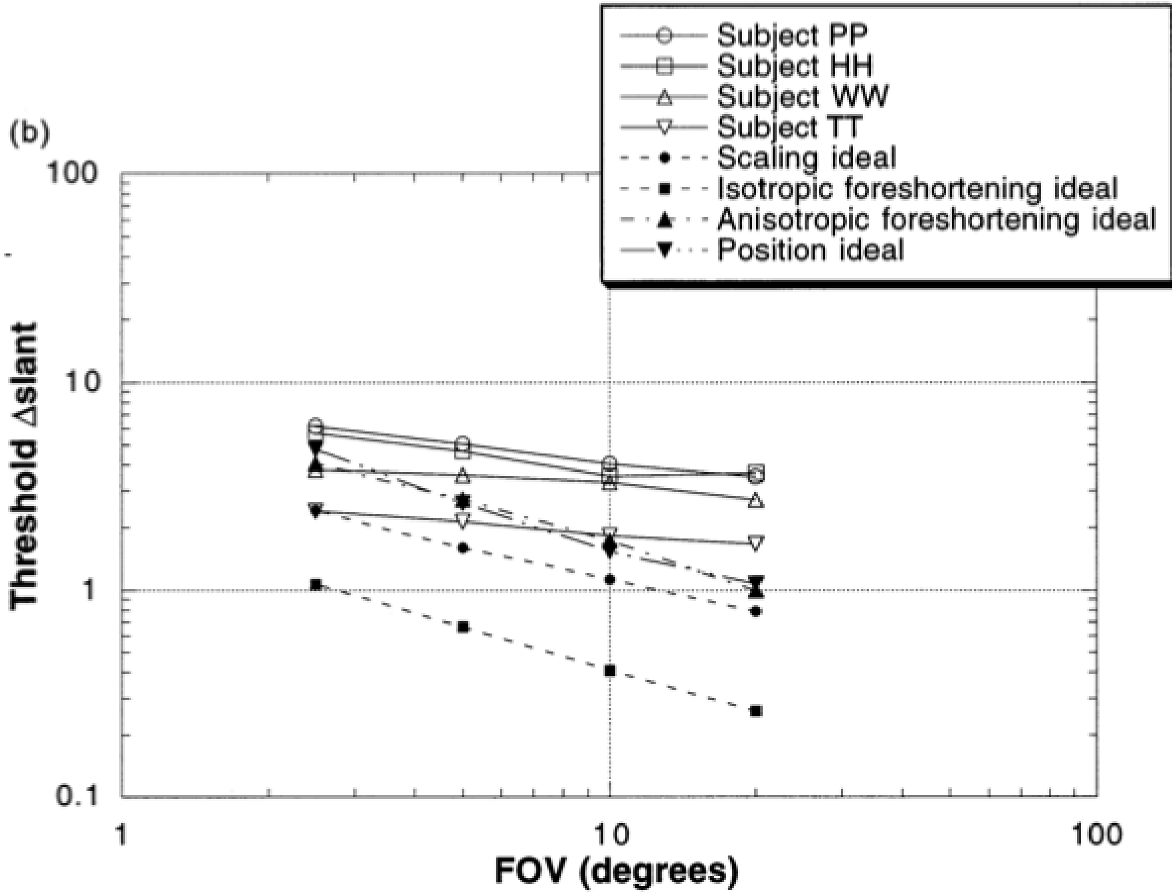

People *barely improve* with FOV
=> Not much of image is used

Natural "reverse correlation"

Natural Cue fluctuations

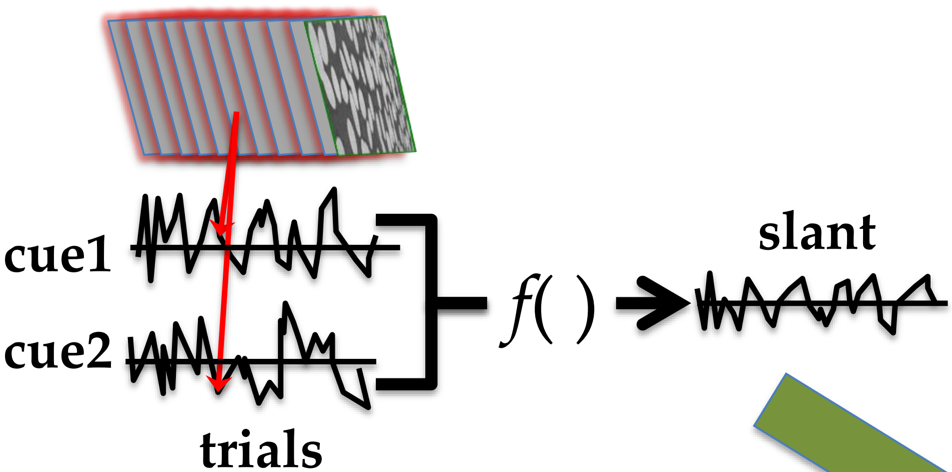

Bayesian
Estimation
with
Family of observer
models

Cue weights

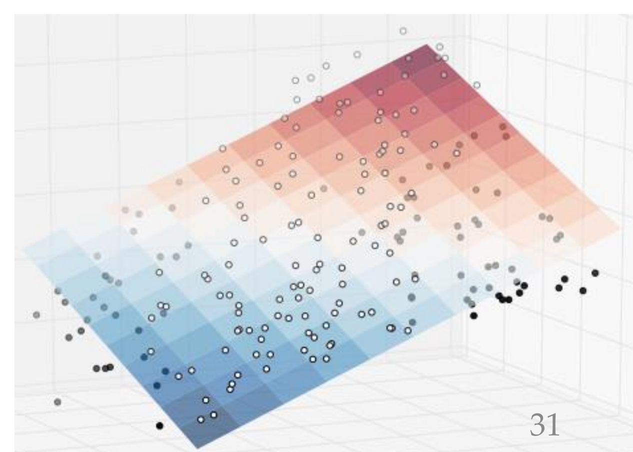

Cue weights

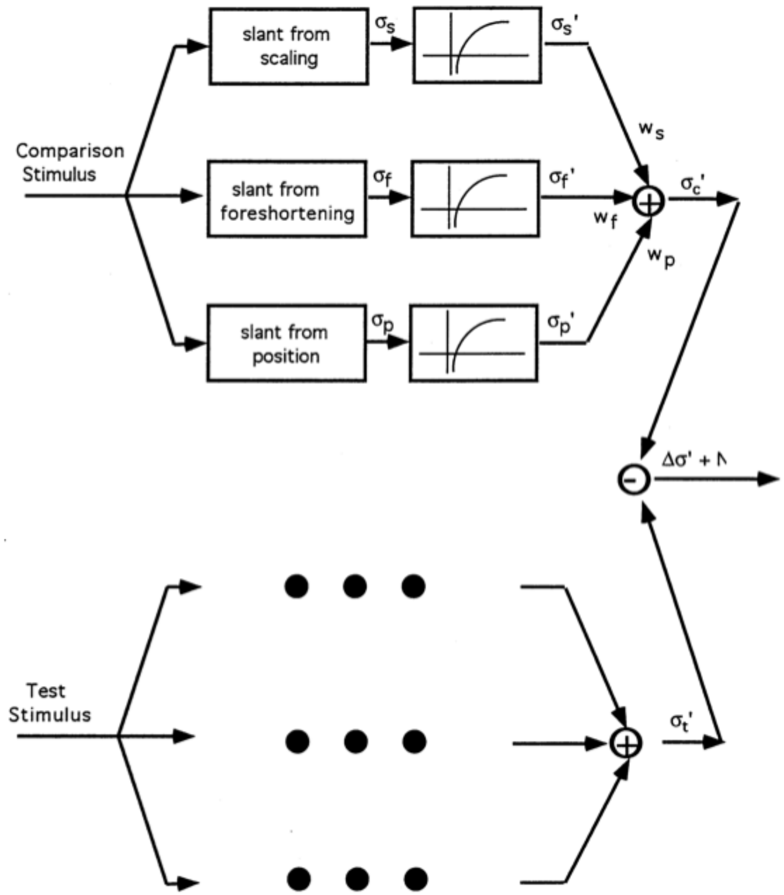

Fig. 3. The signal detection model we used for our analysis (see text for description).

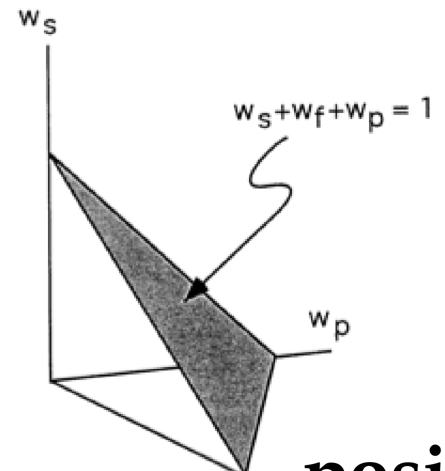

position

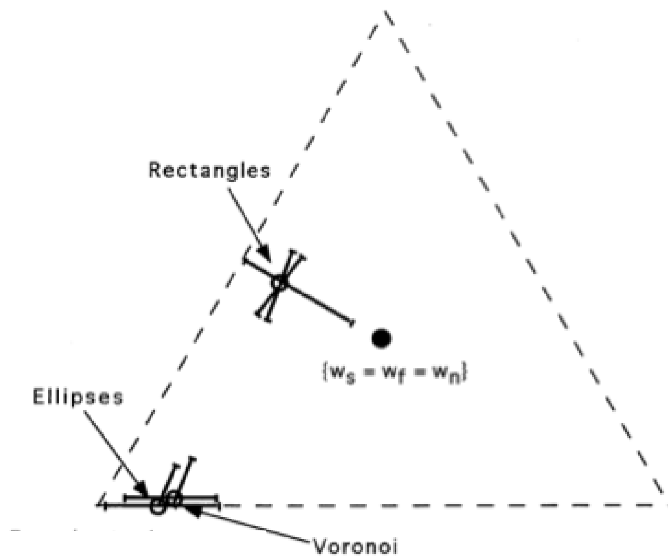

foreshortening

scaling

Change strategy with environmental regularities

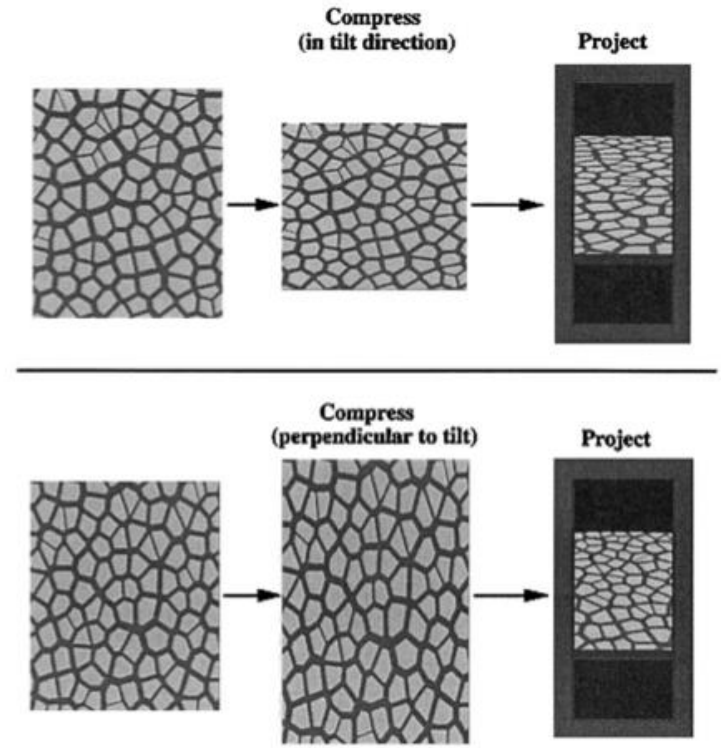

Increasing texture compression

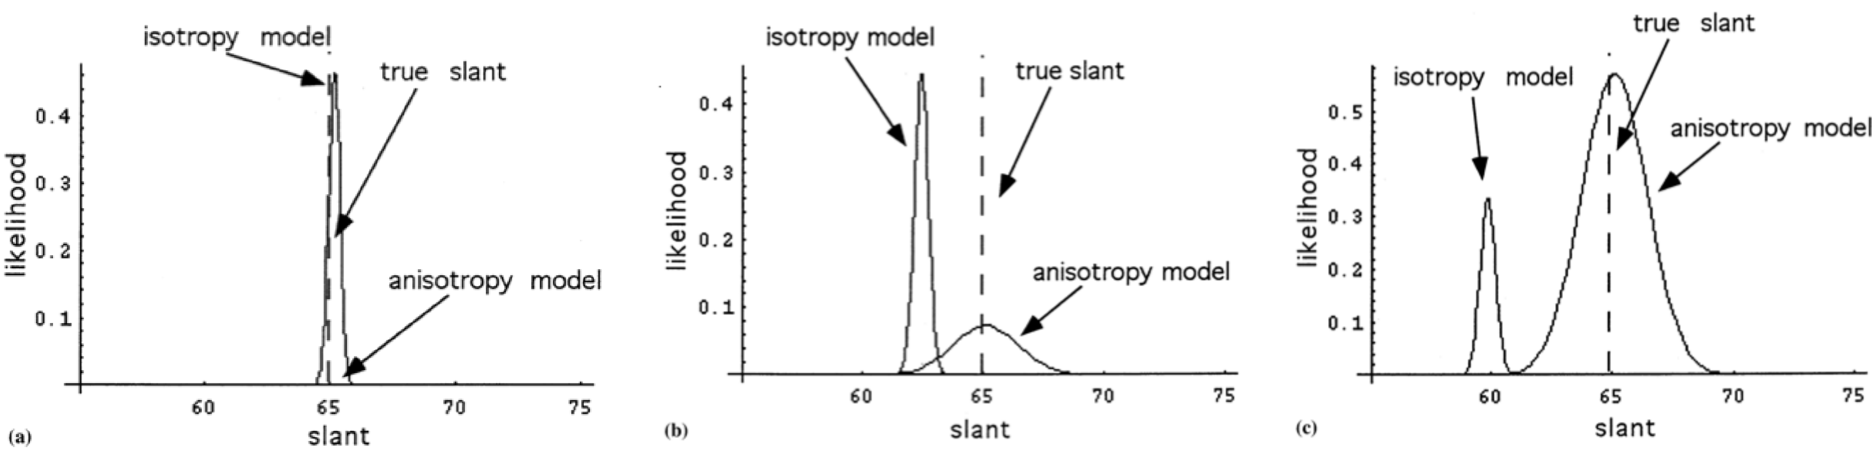

What are cue weights?

What are cue weights?

- Summary descriptions of perceptual performance.

What are cue weights?

- Summary descriptions of perceptual performance.
- Summary descriptions of the information available for a task.

What are cue weights?

- Summary descriptions of perceptual performance.
- Summary descriptions of the information available for a task.
- Support logical links between behavior and rational / normative models of performance.

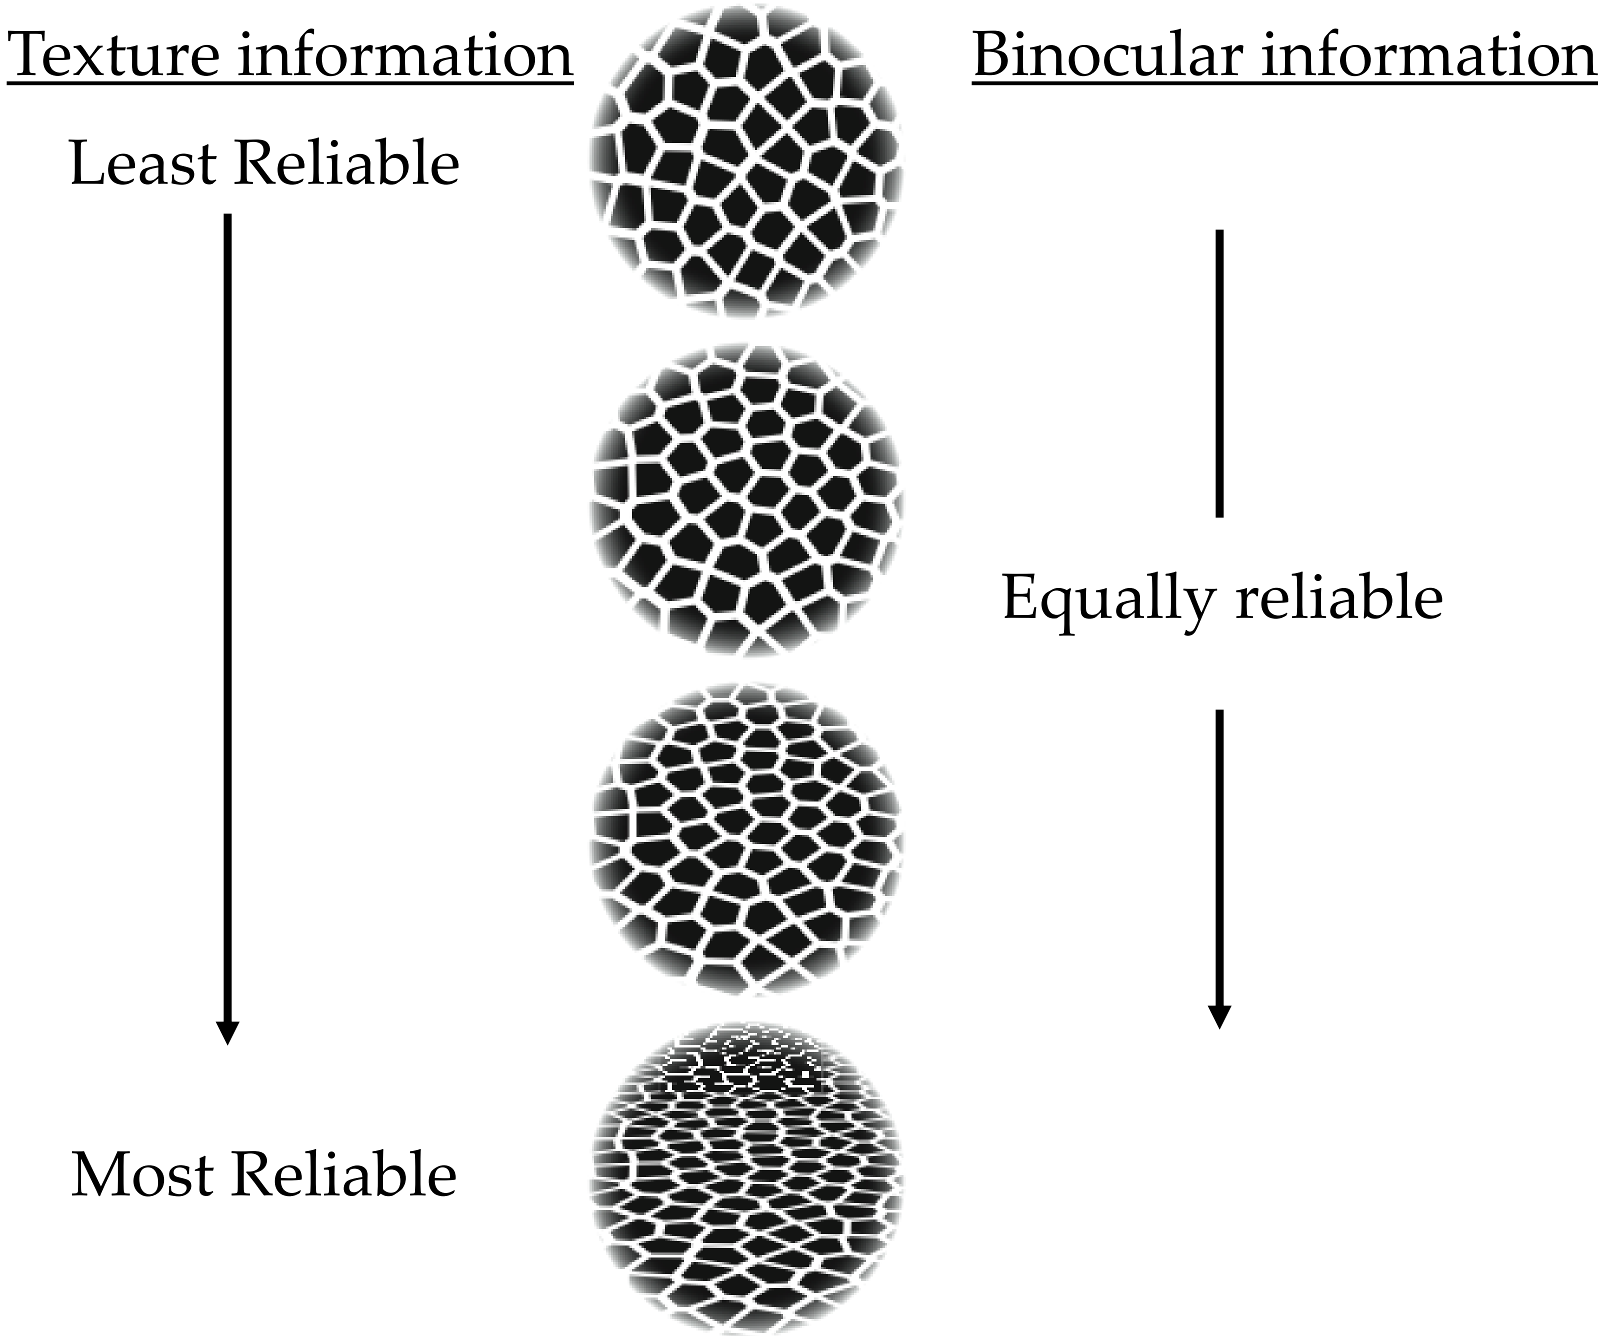
Texture information
Binocular information
Least Reliable
Most Reliable
Equally reliable

Humans weight sensory cues "optimally"

- Discrimination thresholds in single cue conditions predict weights measured in multi-cue experiments.
  - Ernst and Banks, 2002; Knill and Saunders, 2003; Alais and Burr (2004); etc., etc., etc.

Humans weight sensory cues "optimally"

- Discrimination thresholds in single cue conditions predict weights measured in multi-cue experiments.
  - Ernst and Banks, 2002; Knill and Saunders, 2003; Alais and Burr (2004); etc., etc., etc.

Linear process model

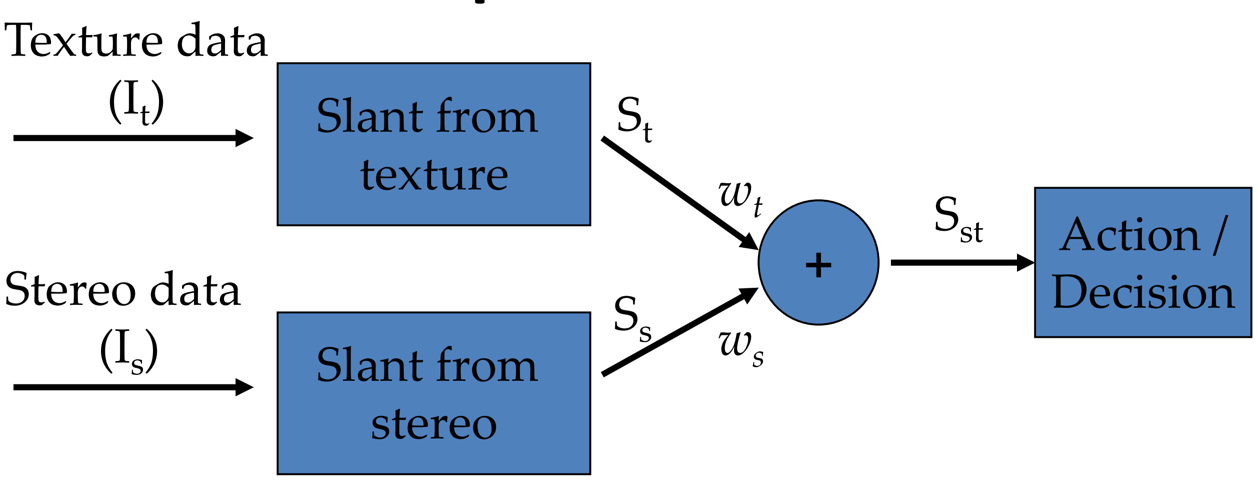

Humans weight sensory cues "optimally"

- Discrimination thresholds in single cue conditions predict weights measured in multi-cue experiments.
  - Ernst and Banks, 2002; Knill and Saunders, 2003; Alais and Burr (2004); etc., etc., etc.

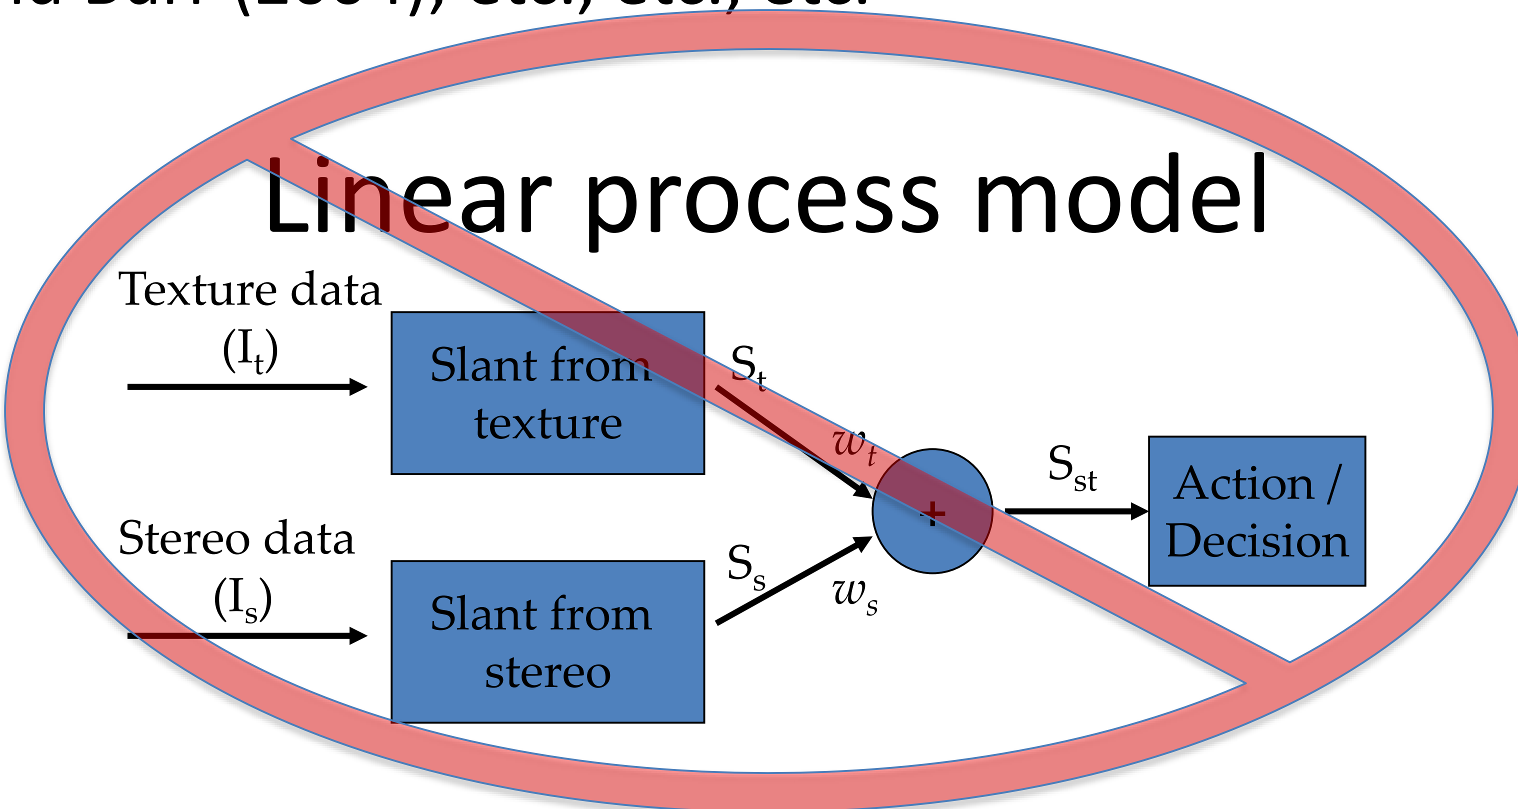

Why depth?

Depth cues: Vision vs. Motor Control

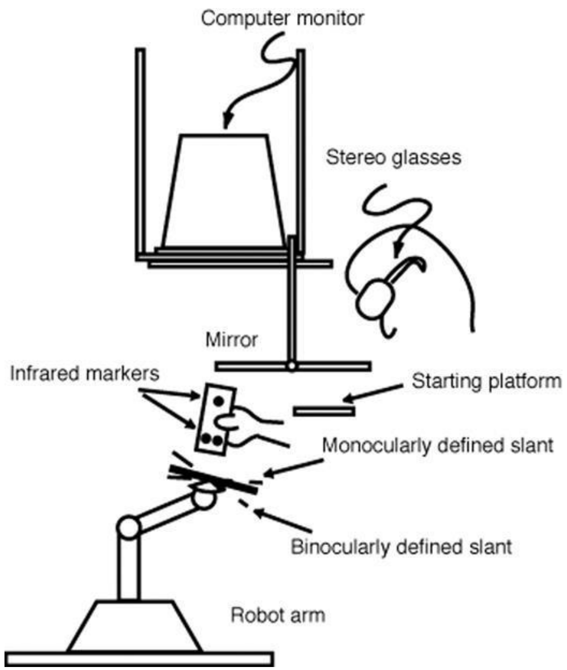

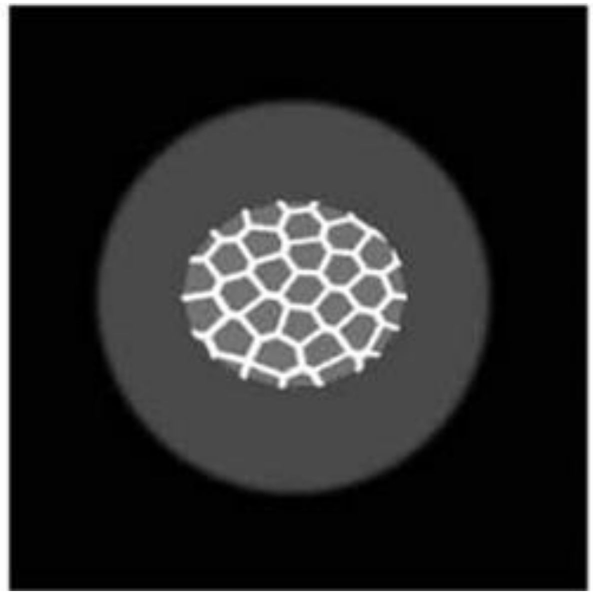

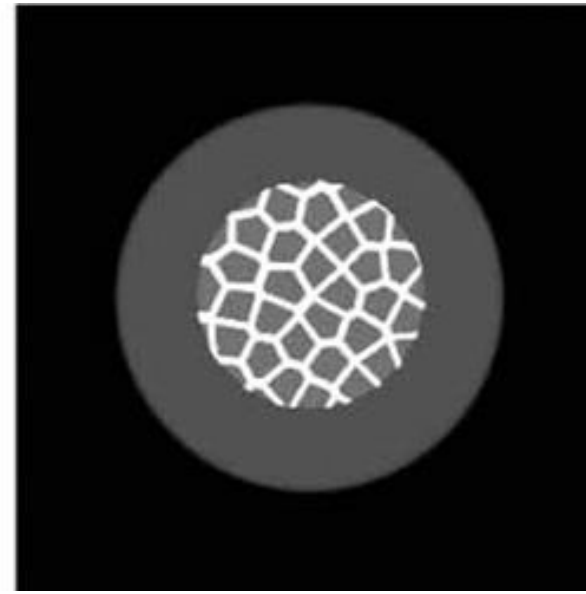

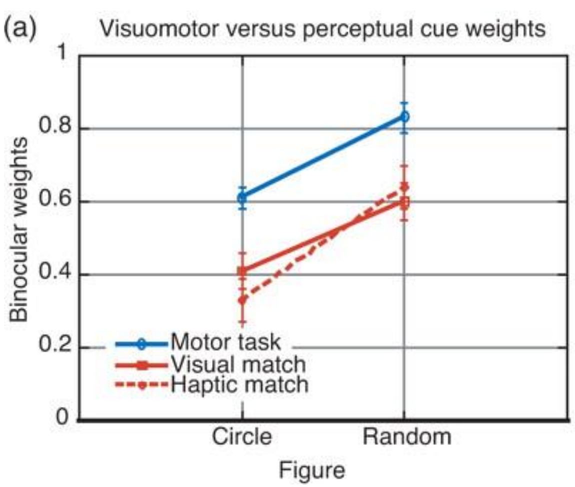

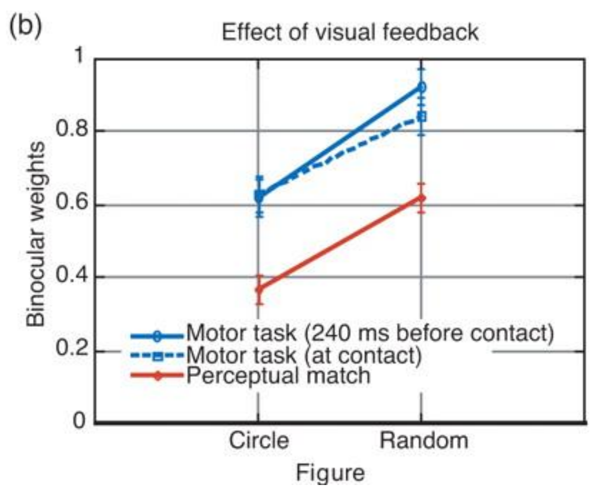

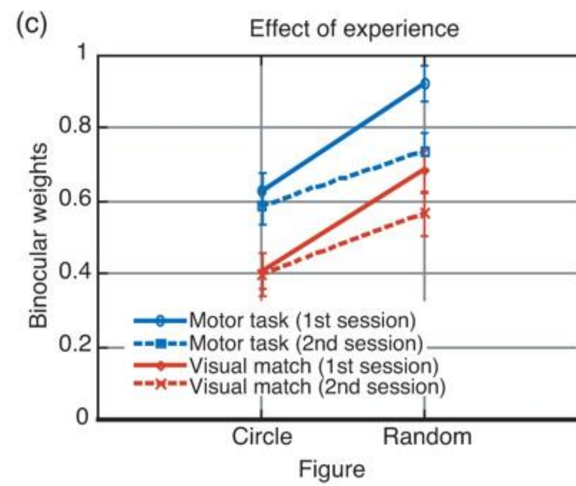

My Work: Complex inference in reaching to depth

This model represents the decomposition:

$$P(X_1, X_2, X_3, X_4) = P(X_4 \mid X_2)\, P(X_3 \mid X_1, X_2)\, P(X_1) P(X_2)$$

Nodes: random variables

$X_1, \ldots , X_4$

Each node has a conditional probability distribution

Links: direct dependencies

Data: observations of X_3 and X_4

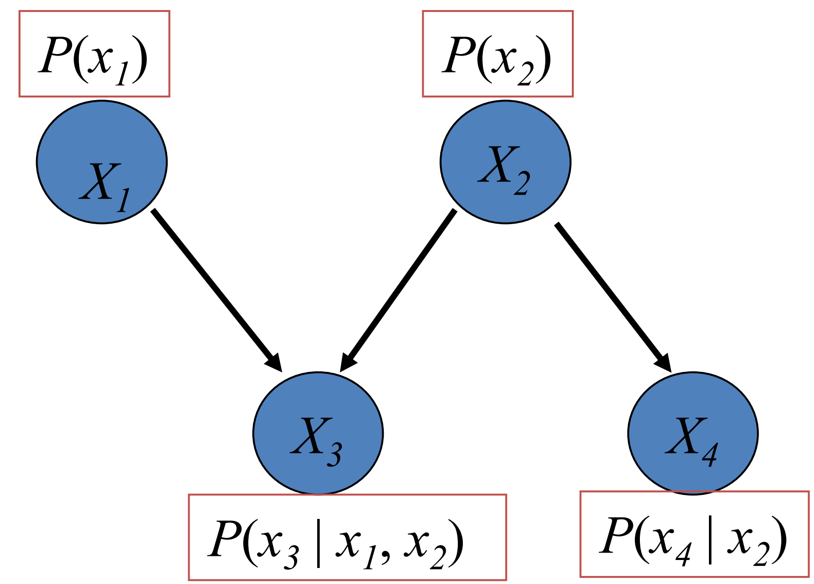

EXAMPLE

X_1 object size

X_2 object distance

X_3 image size

X_4 "felt" distance

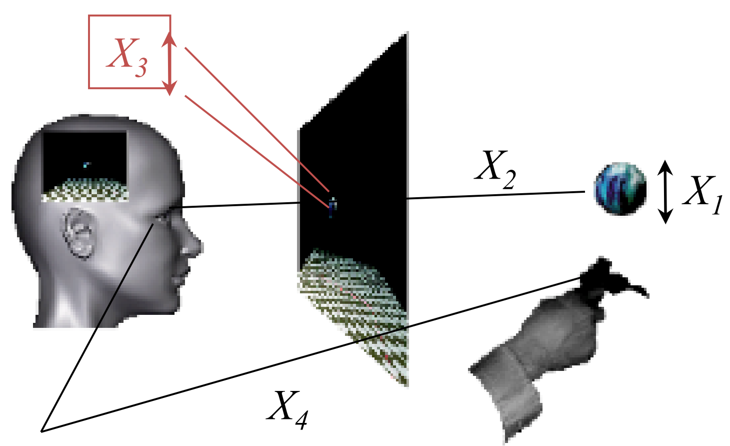

Main Lesson: ***Theory Matters***

It's not a theory of vision unless it can handle real stimuli and tasks

- *Functional analysis*
  - *Why and what come before How*
  - Develop whatever theory you need

- Design airtight psychophysical experiments
  - **BUT** *Embed experiments in near-ecological contexts*

Dave’s Impact

- Professional
  - Key champion of computational level modeling
    - Why do we have vision at all?
    - What’s the brain for?
    - Only given it’s purpose can you make sense of details
  - Key champion of Bayesian analysis
  - Combined rigor, depth and hard problems like very few in the field can. Tough act to follow.
- Personal
  - Taught me how to balance high standards with the joy of discovery
  - How to concoct a story on the spot
  - Never prepare a talk before the night before!

With enough details, all are credulous

- Chinese influence on the origins of Appalachian folk music

IRCS Progress Report 1995

Knill has recently completed a series of studies that examines the role of texture in the perception of orientation. Texture cues indicate surface orientation vis-a-vis the change in shape associated with more distant vs. nearer elements. For example, the grade of a cobblestone road is cued in part by changes in the size and shape of the road elements.

Knill's work indicates that changes in both size and shape of the texture elements contribute to the sense of surface tilt, and that the contributions are approximately equal.

In collaboration with Tjeerd Dijkstra (IRCS postdoctoral fellow), an evaluation of the contribution of highly oriented textures, or texture flow, for the perception of orientation has begun. In the past year, Simoncelli and Knill have begun collaboration on experiments that evaluate the role of temporal deformations in the perception of the shape of texture patterns.

Theoretical Approaches to Multisensory Perception

Robert Jacobs
Department of Brain & Cognitive Sciences
Center for Visual Science
University of Rochester

Multisensory Perception

- Collaborating with Dave was productive and fun:
  - Experiment: Cue reliability and cue recalibration

- Dave loved talking about science. Not only his own science, but **your** science too:
  - Experiment: Generalization from perception to motor production
  - → Implications for perceptual learning

Cue Reliability and Cue Recalibration

- Collaborators
  - Joseph Atkins (Colby College)
  - David Knill

- Atkins, J. E., Jacobs, R. A., & Knill, D. C. (2003). Experience-dependent visual cue recalibration based on discrepancies between visual and haptic percepts. *Vision Research*, *43*, 2603-2613.

“Touch Educates Vision”

- Bishop George Berkeley
  - *An Essay Towards a New Theory of Vision (1709)*

- Perception of visual space results from associations between visual sensations and sensations of touch and motor movement
  - “Touch educates vision”

Research Question

Question:

Can observers adapt their interpretations of a stereo cue on the basis of consistencies (and inconsistencies) between depth-from-stereo and depth-from-haptics percepts?

Visual Stimuli

- Scenes consisted of two fronto-parallel surfaces
- Narrow surface was closer to subject, and it occluded the middle portion of the wide surface
- Subjects viewed scene head-on (orthogonal view)
- Stereo only reliable visual cue to depth between two surfaces

Visual Stimuli

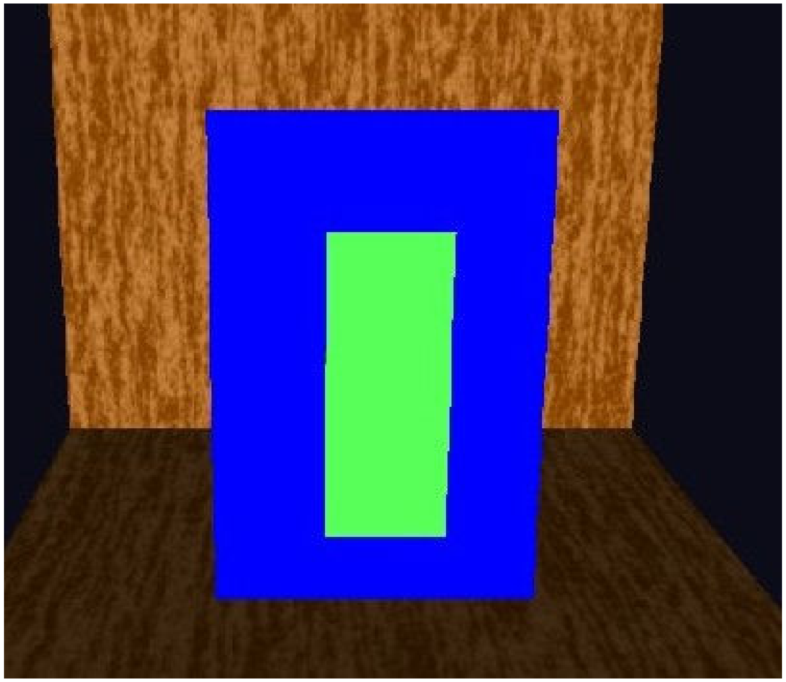

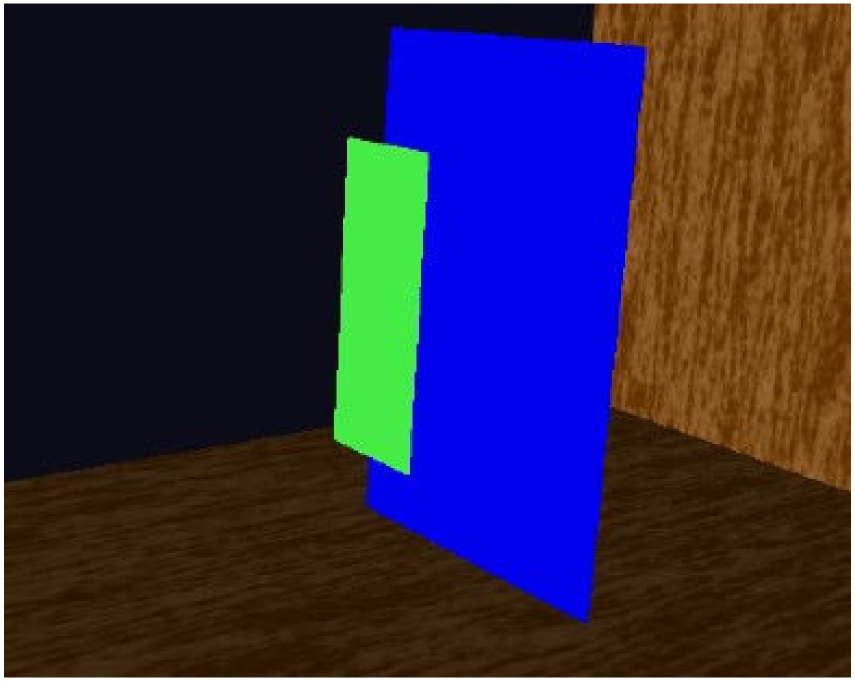

Virtual Reality Environment

Procedure

- Judgment: Is width of front surface greater or less than the depth between the two surfaces?
  - Based on visual cues
  - Based on visual and haptic cues
  - No corrective feedback

- Four stages:
  - Consistent-cue training trials
  - Pre-test trials (visual information only)
  - Inconsistent-cue training trials
  - Post-test trials (visual information only)

Cue Conflict

- Independent control of:
  - Depth indicated by visual stereo
  - Depth indicated by haptics

- Trials with inconsistent cues:
  - Reaching distance greater than viewing distance by 60mm
  - Binocular disparities consistent with both reaching and viewing distances

    → Scaled depth between front and rear surfaces so that depth indicated by haptics > depth indicated by stereo

Prediction

- Based on inconsistent-cue training trials, subjects will adapt their depth-from-stereo estimates so that these estimates become more similar to their depth-from-haptics estimates

Experimental Results

- Subject TL:

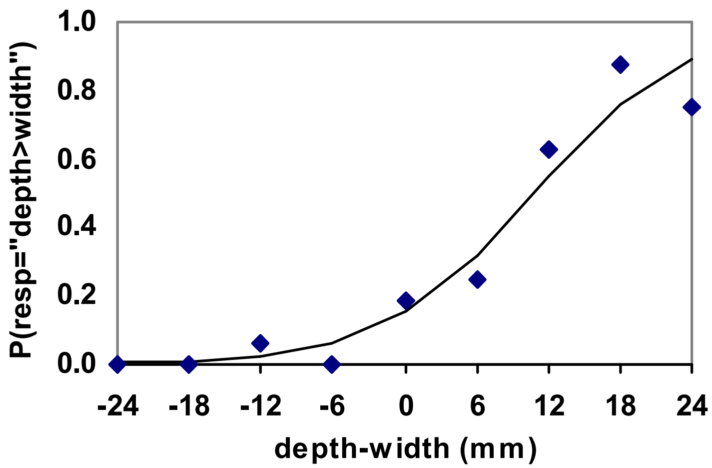

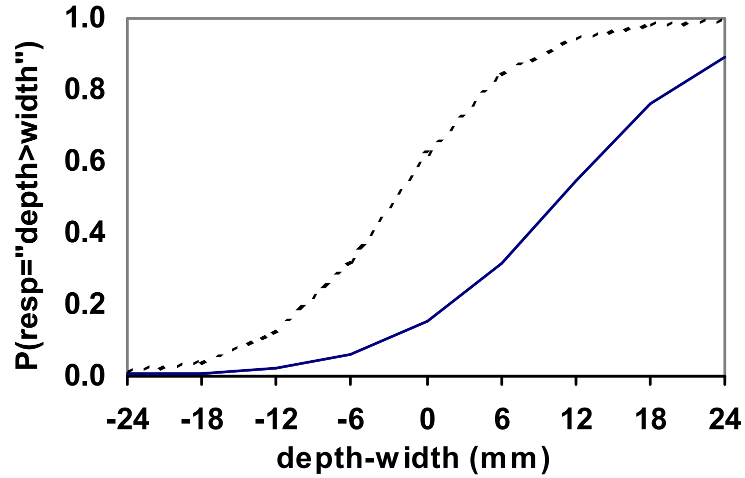

Experimental Results

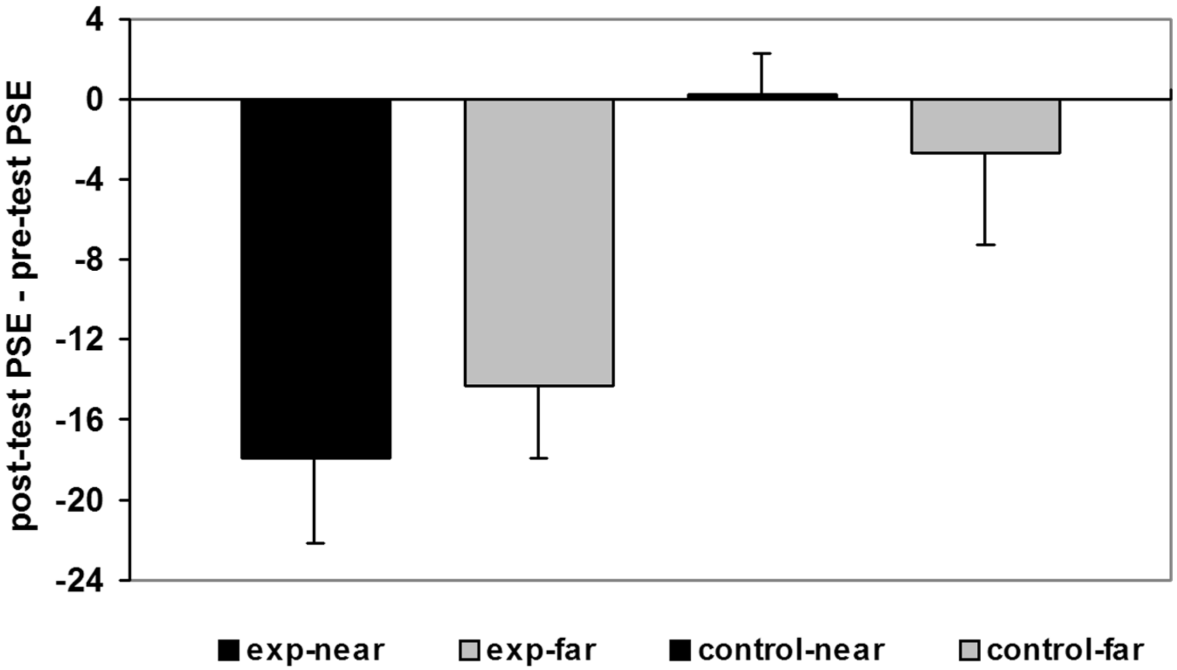

- **Bishop Berkeley was right!**

Multisensory Perception

- Collaborating with Dave was productive and fun:
  - Experiment: Cue reliability and cue recalibration
- Dave loved talking about science. Not only his own science, but **your** science too:
  - Experiment: Generalization from perception to motor production
  - → Implications for perceptual learning

Generalization from Perception to Motor Production

- Collaborators:
  - Daniel Meegan (University of Guelph)
  - Richard Aslin (University of Rochester)

- Meegan, D. V., Aslin, R. N., & Jacobs, R. A. (2000). Motor timing learned without motor training. *Nature Neuroscience*, **3**, 860-862.

Generalization From Perception to Motor Production

- Experiment:
  - Motor production tests (Tasks 1 and 2)
  - Perceptual training
  - Motor production tests (Tasks 1 and 2)

Generalization From Perception to Motor Production

- Motor production tasks:
  - Produce two finger taps separated by a target temporal interval
    - Task 1: target interval = 300 ms
    - Task 2: target interval = 500 ms
  - Feedback: actual temporal interval

Generalization From Perception to Motor Production

- Perceptual training:
  - Auditory temporal interval duration discrimination task
    - Temporal intervals indicated by 2 auditory tones
  - On each trial, subjects heard two intervals (standard and comparison) and judged which one was longer
    - Group 1: standard interval = 300 ms
    - Group 2: standard interval = 500 ms

Generalization From Perception to Motor Production

- Prediction:
  – Subjects will show more motor improvement when the temporal requirements of the perceptual and motor tasks are identical
    - Subjects trained to perceptually discriminate 300 ms (500 ms) intervals from other intervals will show the most improvement on producing 300 ms (500 ms) intervals

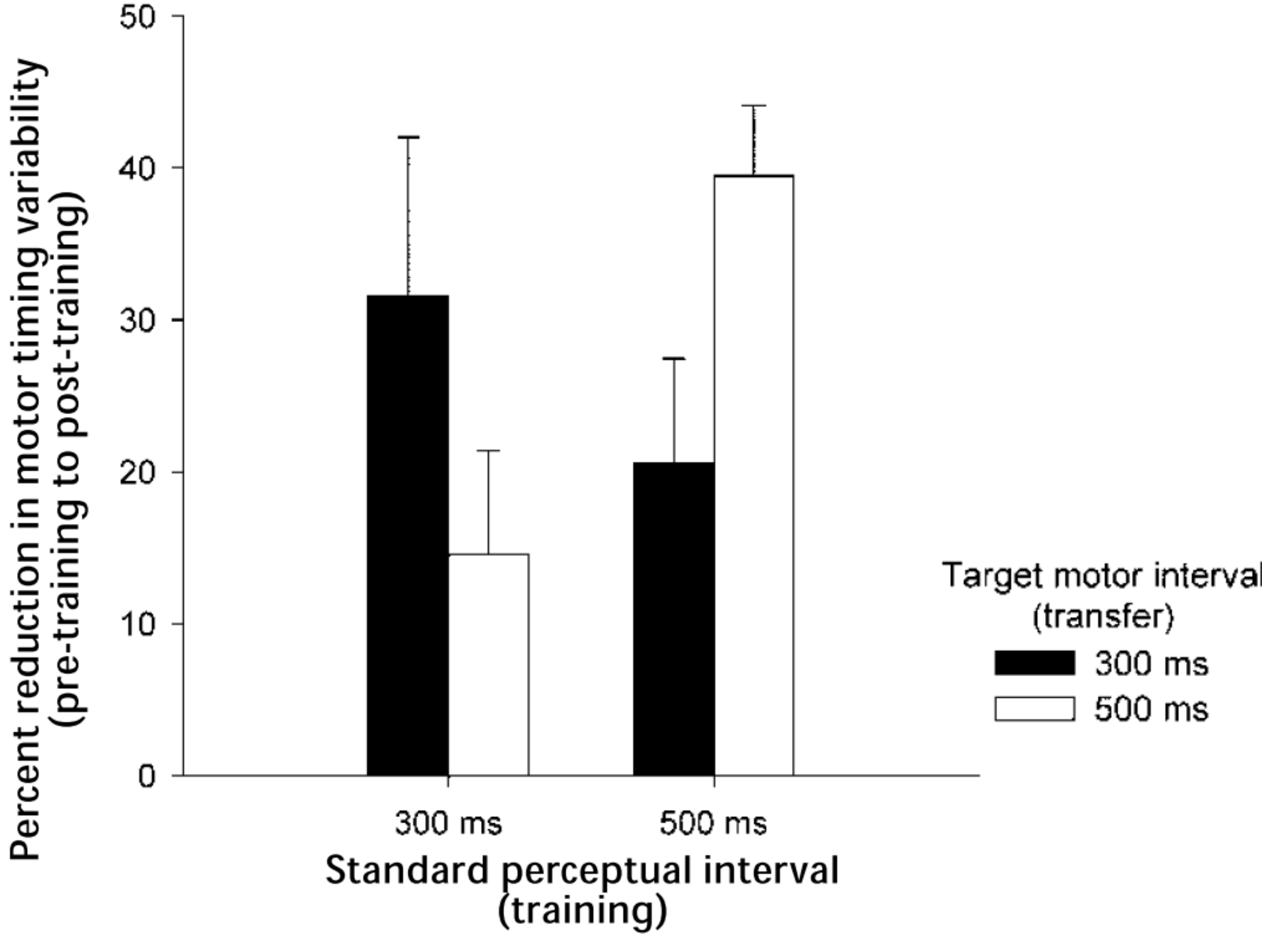

→ Subjects showed more motor improvement when the temporal requirements of the perceptual and motor tasks were identical

Implications for Perceptual Learning

- Cross-modal transfer
  - Acquire knowledge about the environment through one sensory modality
  - Apply acquired knowledge when the environment is sensed through a different sensory modality

- Example: If you learn to visually categorize a novel set of objects, you can also often categorize the same (and similar) objects when they are grasped but not seen (Yildirim & Jacobs, 2013)

Implications for Perceptual Learning

- To us, cross-modal transfer and transfer from perception to motor production are closely related phenomenon
  - Both suggest the existence of amodal representations
  - If so, then experiment on transfer from perception to motor production has implications for perceptual learning

Narrow vs. Broad Generalization

- Perceptual learning
  - Many studies report that perceptual learning is often stimulus-specific (narrow generalization)
    - Fiorentini & Berardi (1980, 1981), Shiu & Pashler (1992), Fahle, Edelman, & Poggio (1995), Liu & Vaina (1998)
  - However…cross-modal transfer of knowledge is, by definition, **not** stimulus-specific

- Q: When is generalization narrow and when is it broad?

Are People Biased Toward Cross-Modal Transfer?

- In our experiment, subjects simultaneously generalized both narrowly and broadly
  - Narrow
    - Transfer of learning better for trained temporal interval
  - Broad
    - Transfer of learning from perception to motor production

- Hypothesis: Cross-modal transfer has a privileged status
  - People are biased toward generalizing cross-modally even under circumstances in which they simultaneously fail to generalize (or generalize narrowly) along other dimensions

Thank you!!!

Looking back and moving forward:

Dave Knill's contributions to visual memory and motor control

Chris R. Sims

DREXEL UNIVERSITY
Applied Cognitive &
Brain Sciences

Knill Lab, ~2009–2013

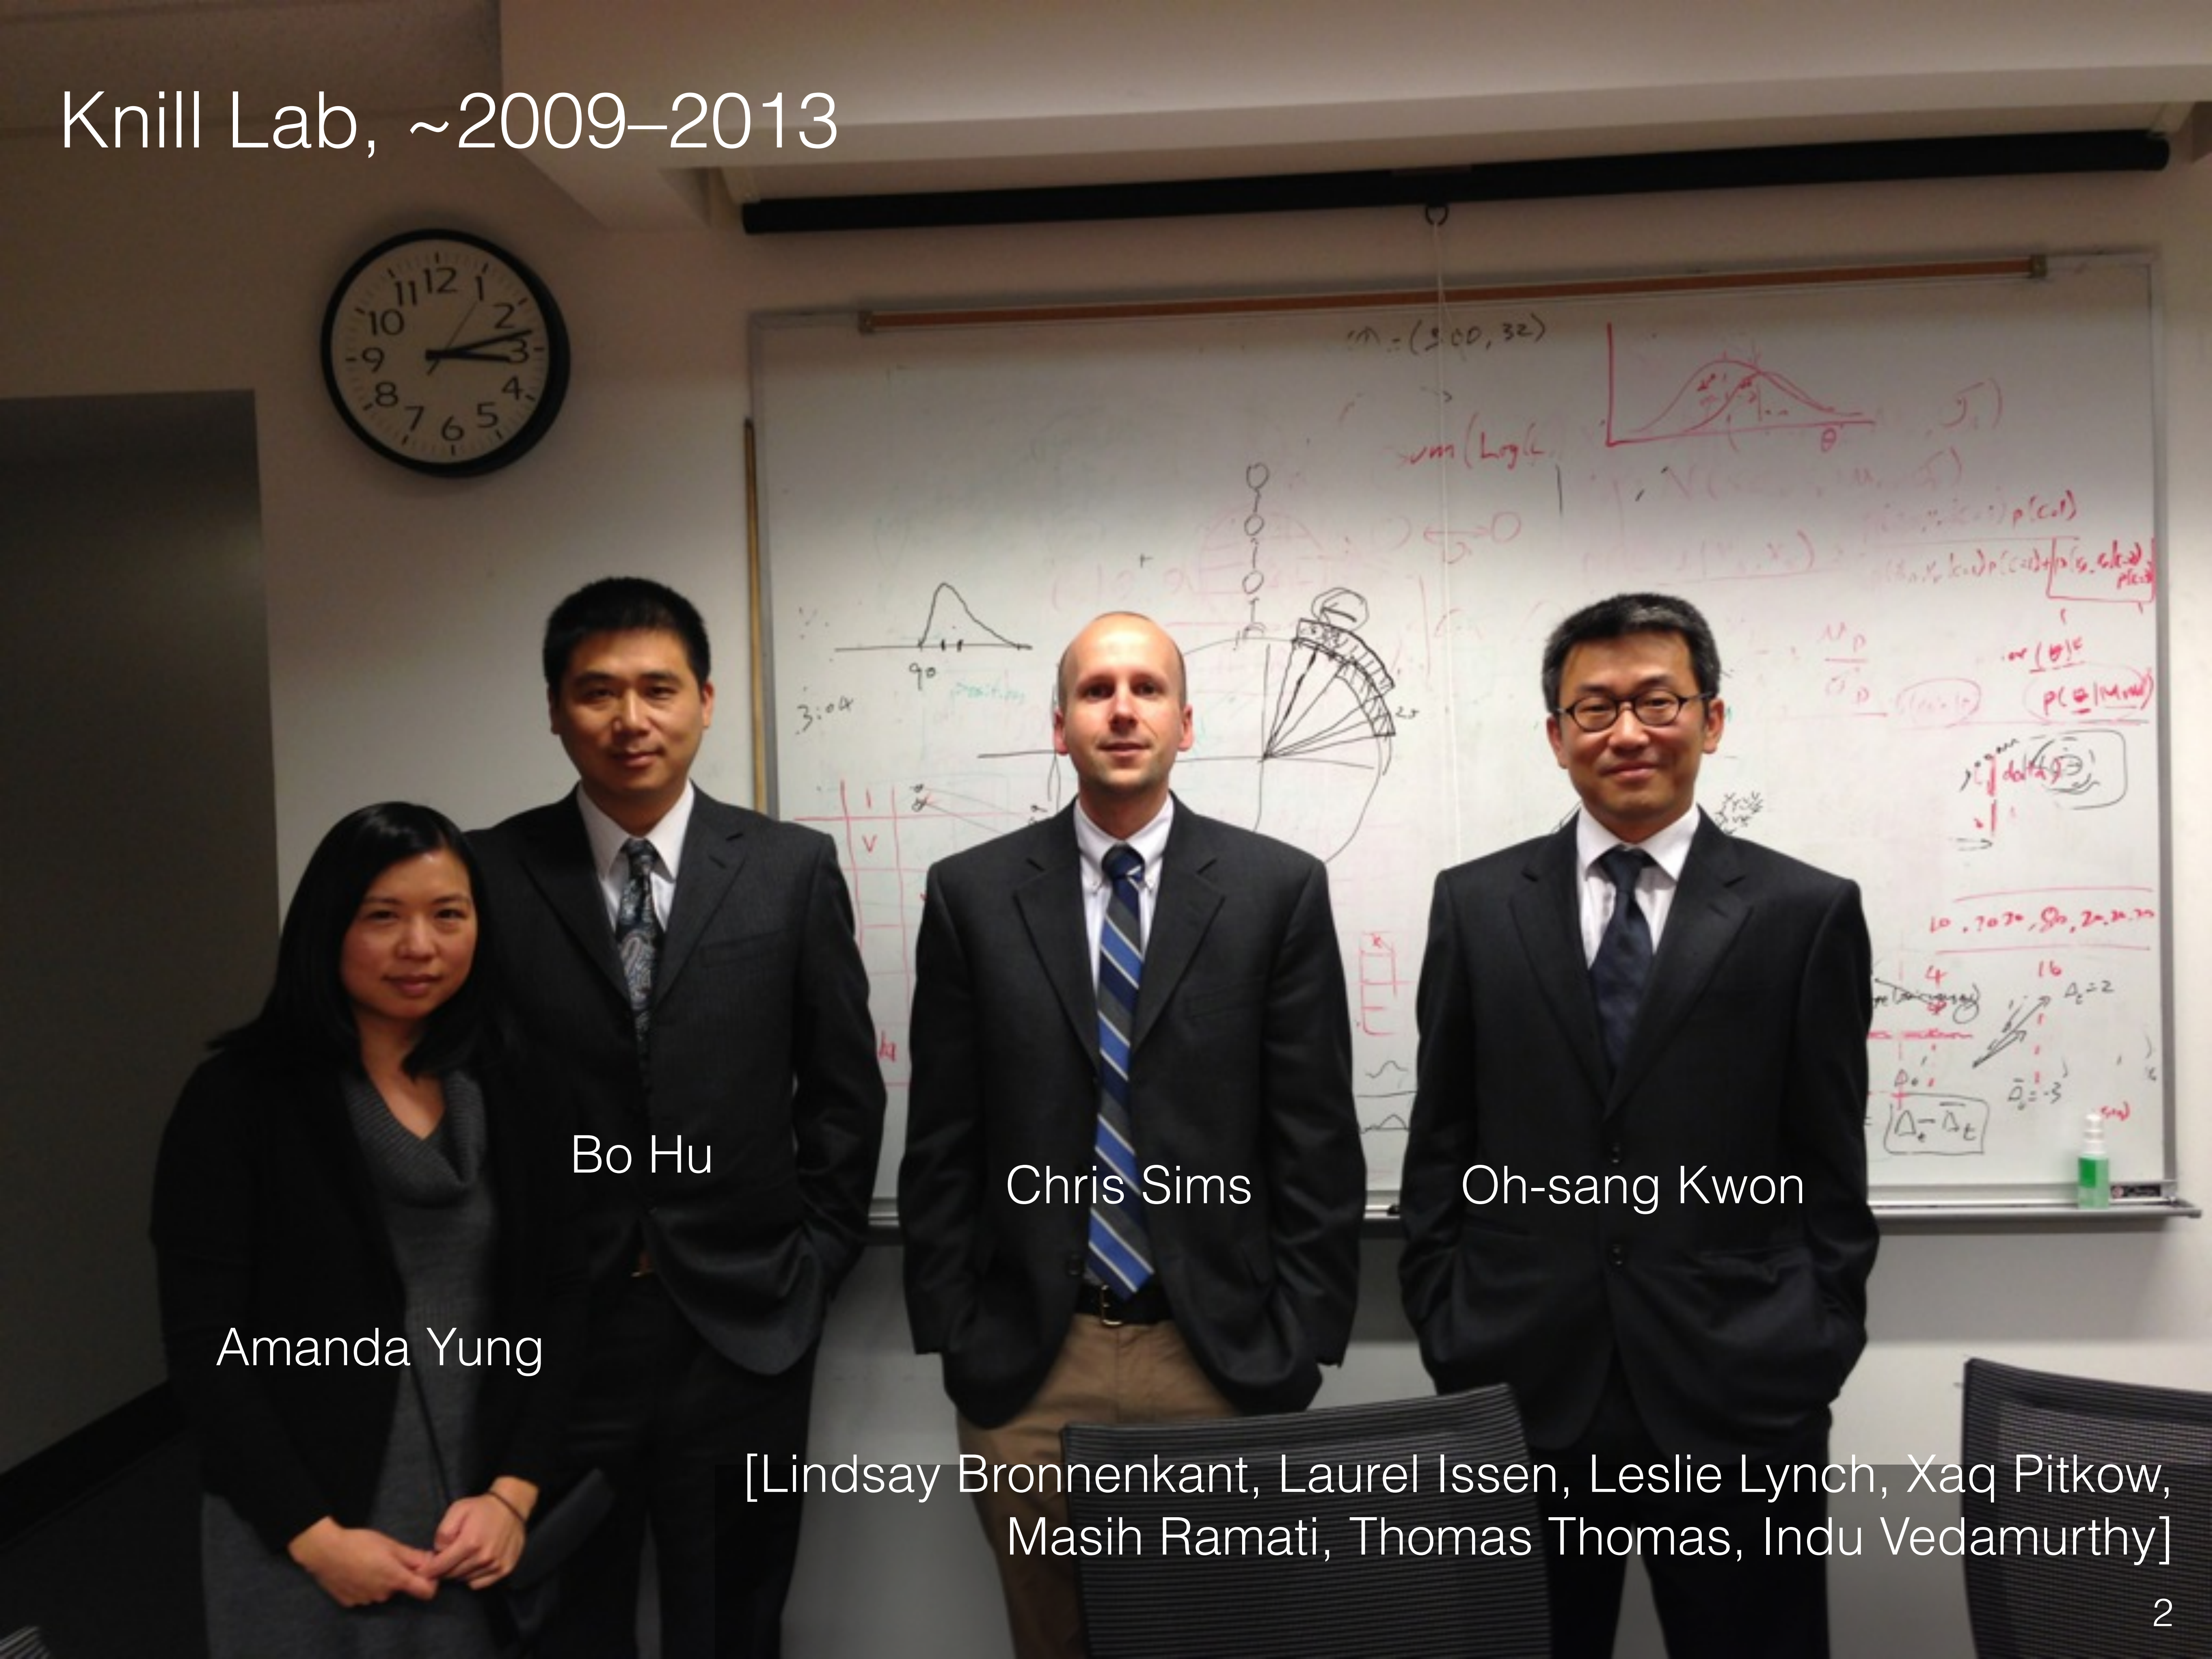

[Lindsay Bronnenkant, Laurel Issen, Leslie Lynch, Xaq Pitkow, Masih Ramati, Thomas Thomas, Indu Vedamurthy]

What I learned from Dave Knill

1. Think harder
2. Don't be satisfied with inelegant solutions
3. Enjoy the journey

Outline (aka, this is an impossible task)

- Sensorimotor control and coordination

Saunders, J. A., & Knill, D. C. (2003). Humans use continuous visual feedback from the hand to control fast reaching movements. Experimental Brain Research, 152(3), 341-352.

Saunders, J. A., & Knill, D. C. (2004). Visual feedback control of hand movements. The Journal of neuroscience, 24(13), 3223-3234.

Saunders, J. A., & Knill, D. C. (2005). Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. Experimental Brain Research, 162(4), 458-473.

Greenwald, H. S., Knill, D. C., & Saunders, J. A. (2005). Integrating visual cues for motor control: A matter of time. Vision research, 45(15), 1975-1989.

Knill, D. C., Bondada, A., & Chhabra, M. (2011). Flexible, task-dependent use of sensory feedback to control hand movements. The Journal of Neuroscience, 31(4), 1219-1237.

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2011). Adaptive allocation of vision under competing task demands. The Journal of Neuroscience, 31(3), 928-943.

- Visual memory

Brouwer, A. M., & Knill, D. C. (2007). The role of memory in visually guided reaching. Journal of Vision, 7(5), 6.

Brouwer, A. M., & Knill, D. C. (2009). Humans use visual and remembered information about object location to plan pointing movements. Journal of vision, 9(1), 24.

Issen, L. A., & Knill, D. C. (2012). Decoupling eye and hand movement control: visual short-term memory influences reach planning more than saccade planning. Journal of vision, 12(1), 3.

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2012). An ideal observer analysis of visual working memory. Psychological review, 119(4), 807.

Orhan, A. E., Sims, C. R., Jacobs, R. A., & Knill, D. C. (2014). The adaptive nature of visual working memory. Current Directions in Psychological Science, 23(3), 164-170.

Sims, C. R. (2015). The cost of misremembering: Inferring the loss function in visual working memory. Journal of vision, 15(3), 2.

Sensorimotor control

$\hat{x} = f(x_{obs})$	State estimation
$u = g(\hat{x})$	Feedback control law
$\mathcal{L} = h\,(x, u)$	Cost function

Goal: Minimize $\mathcal{L}$ w.r.t. f, g

Stochastic optimal feedback control

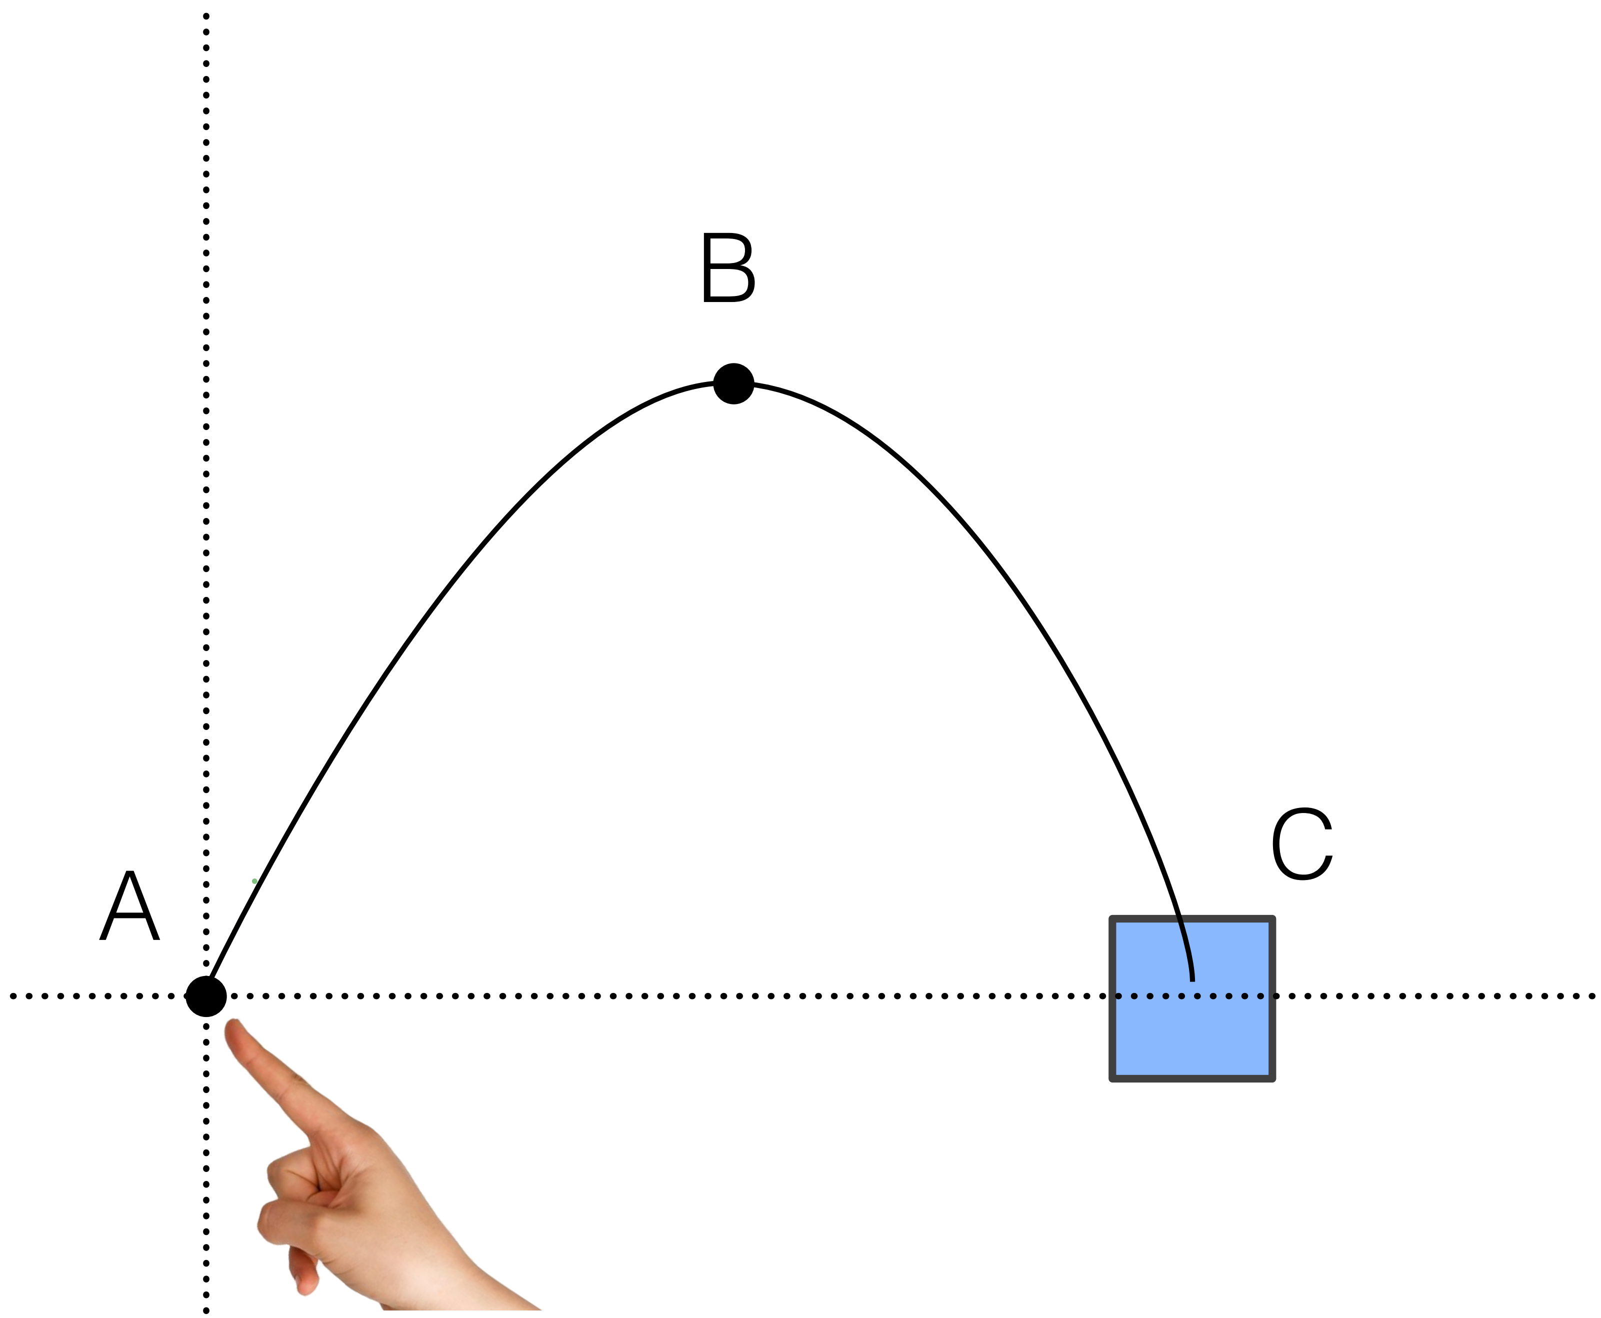

Stochastic optimal feedback control

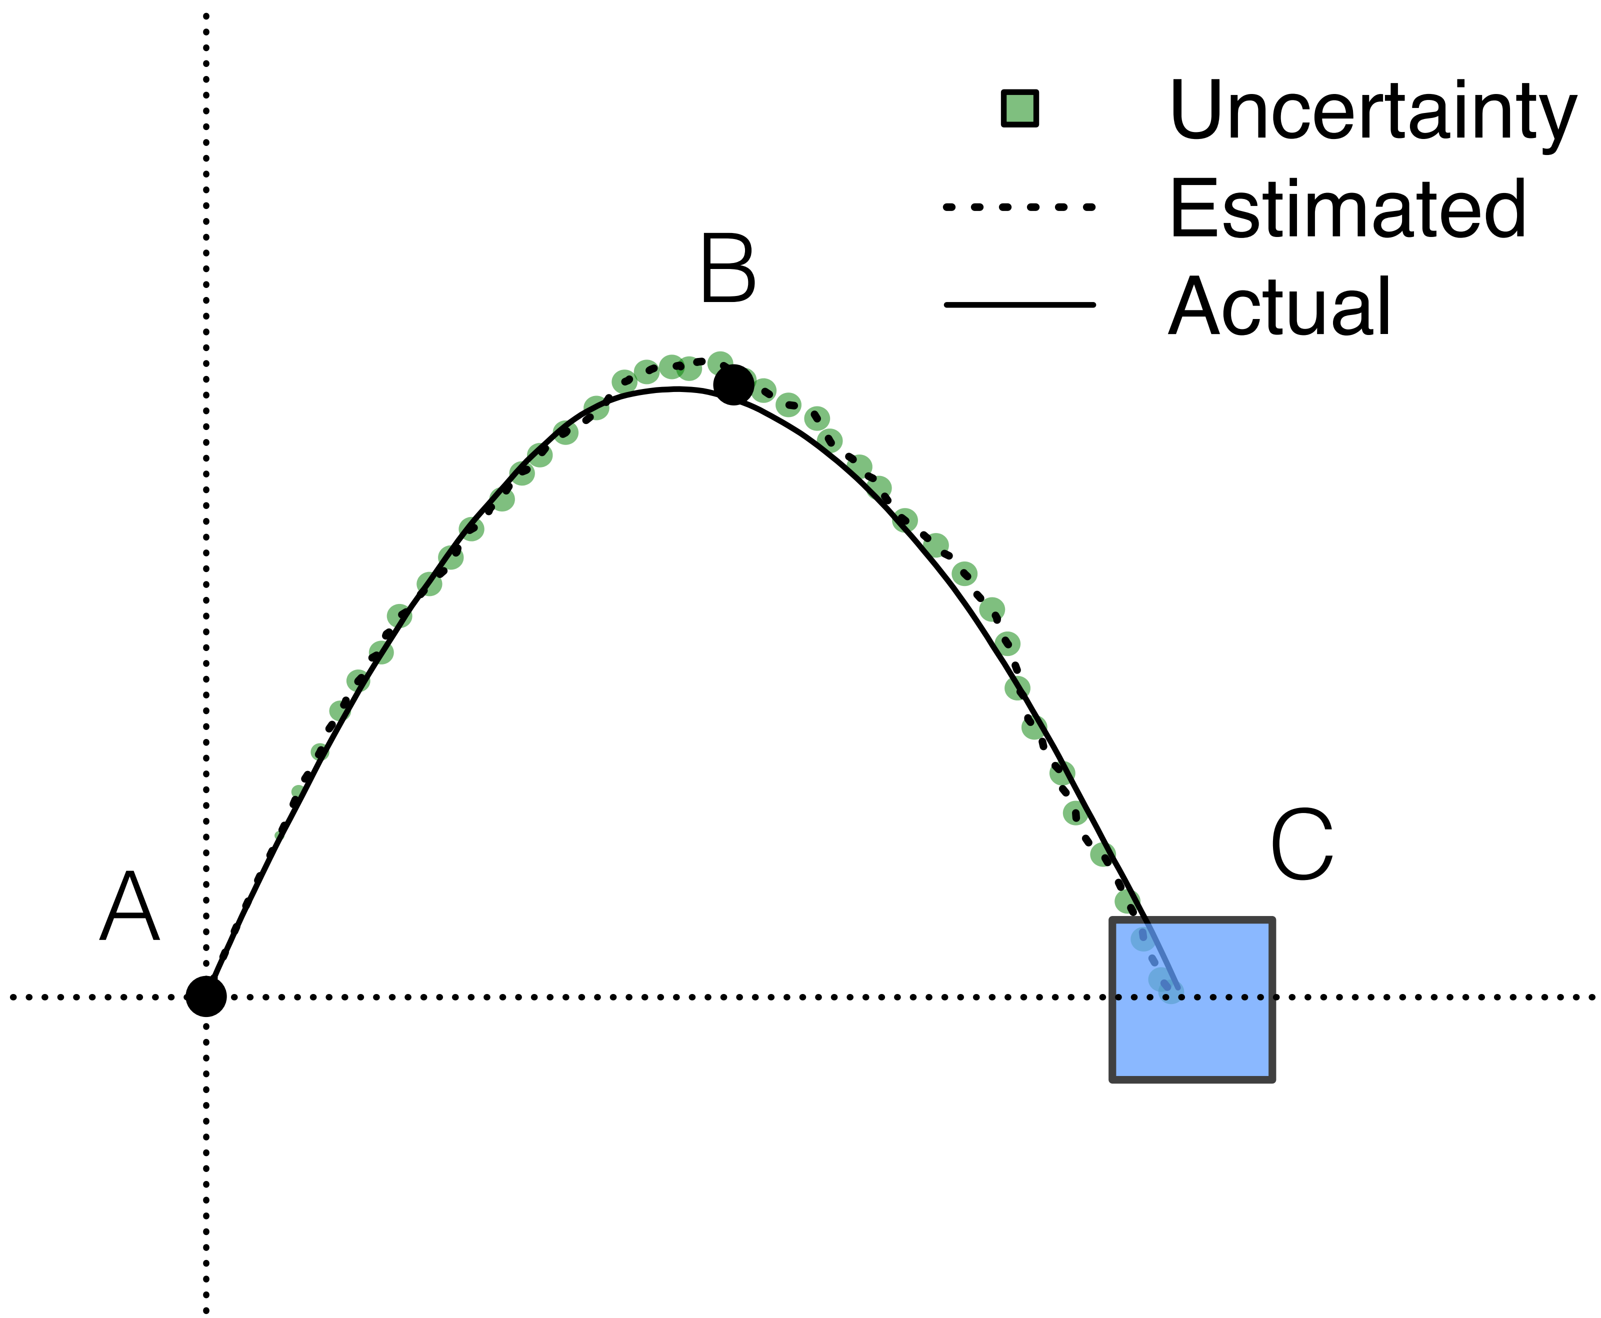

Stochastic optimal feedback control

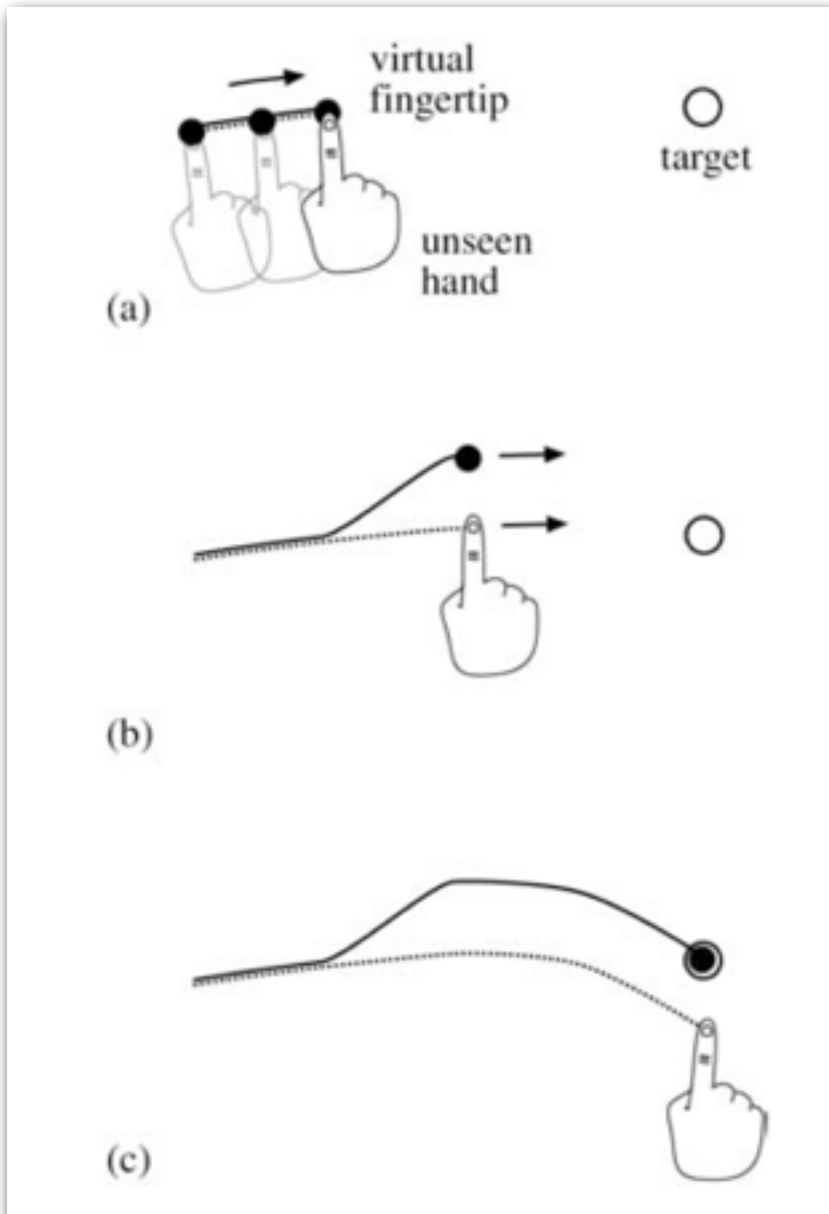

(Saunders & Knill, 2003)*

"The results of the current experiment provide the first direct evidence for continuous, on-line visual control of the moving hand that extends throughout the course of reaching movements. We hope that these results will help to settle the long-running debate concerning the role of visual feedback in the control of reaching movements. The technique of perturbing a virtual hand during reaching movements provides a promising tool for further exploring the nature of the visual feedback that the brain uses to control reaching movements."

* Research also presented at first VSS meeting in 2001

Sensorimotor coordination

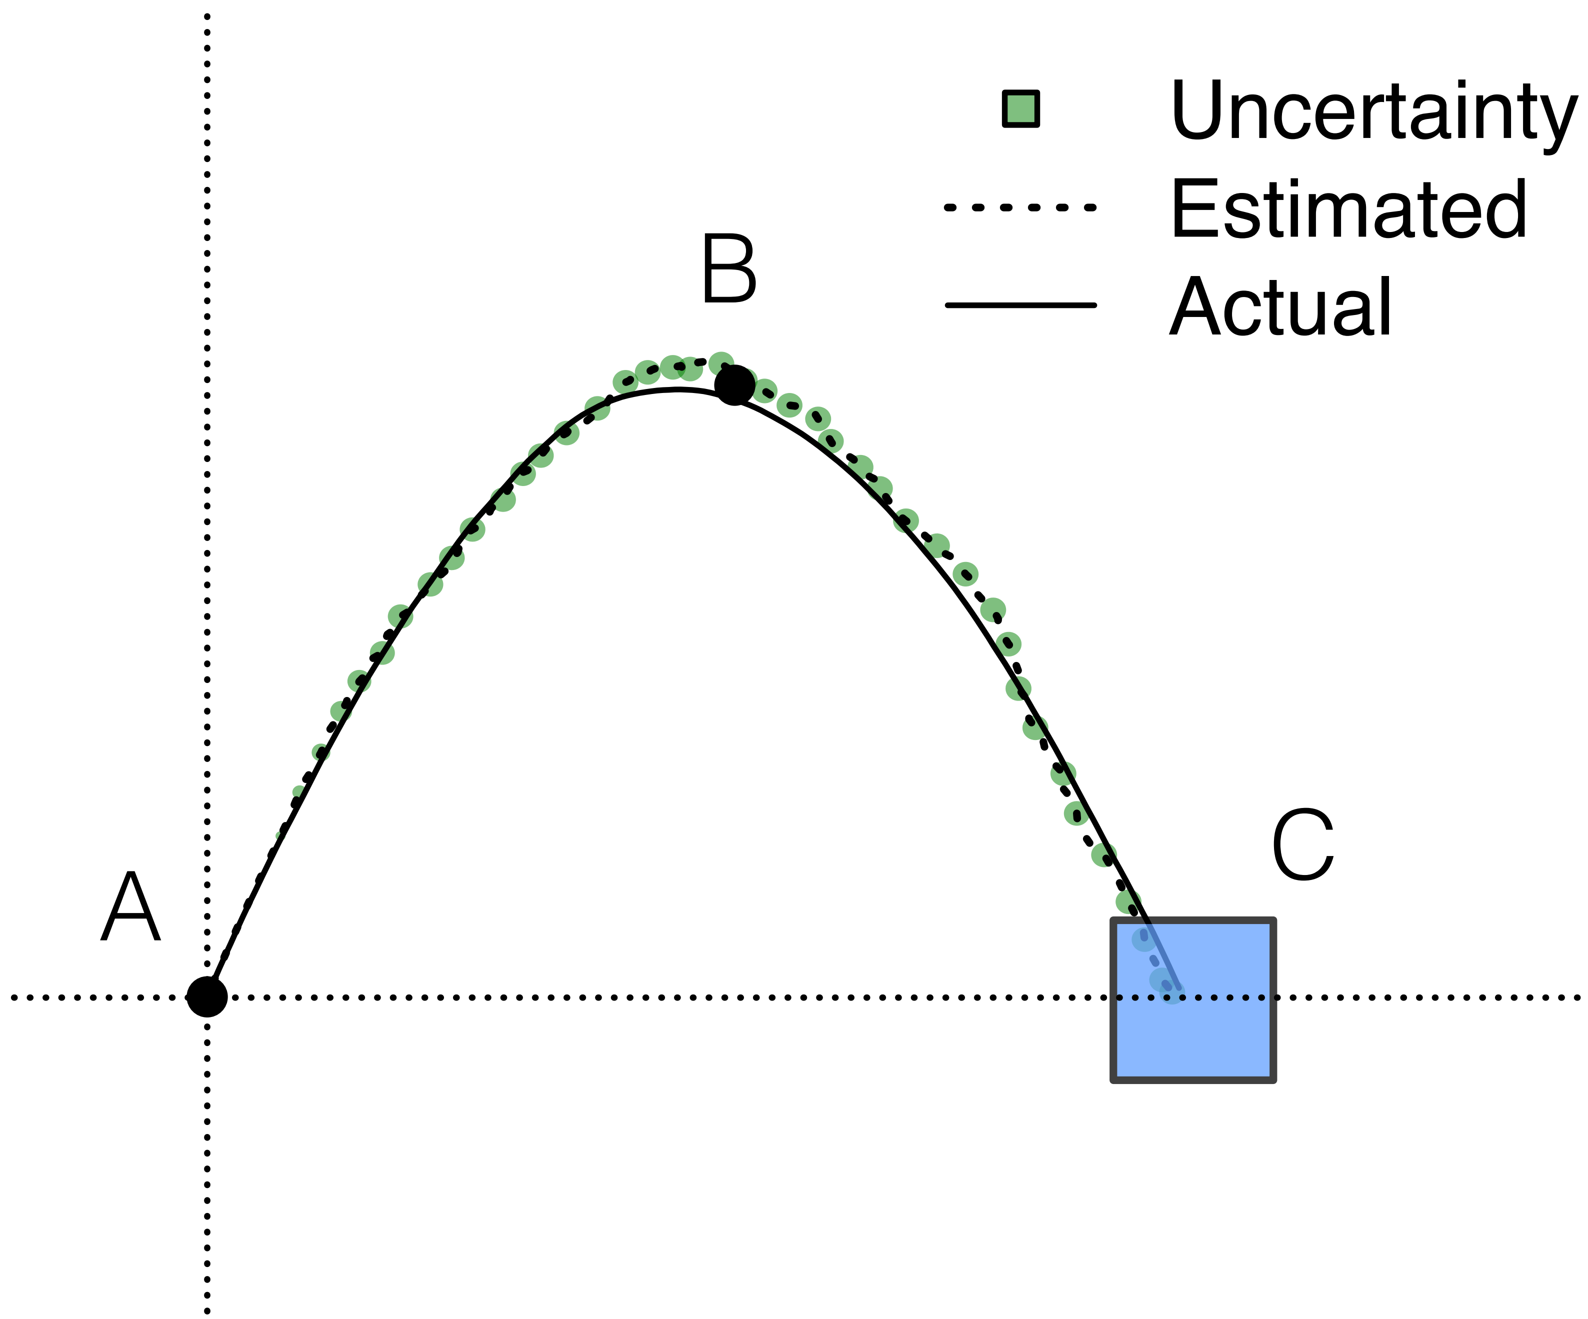

Sensorimotor coordination

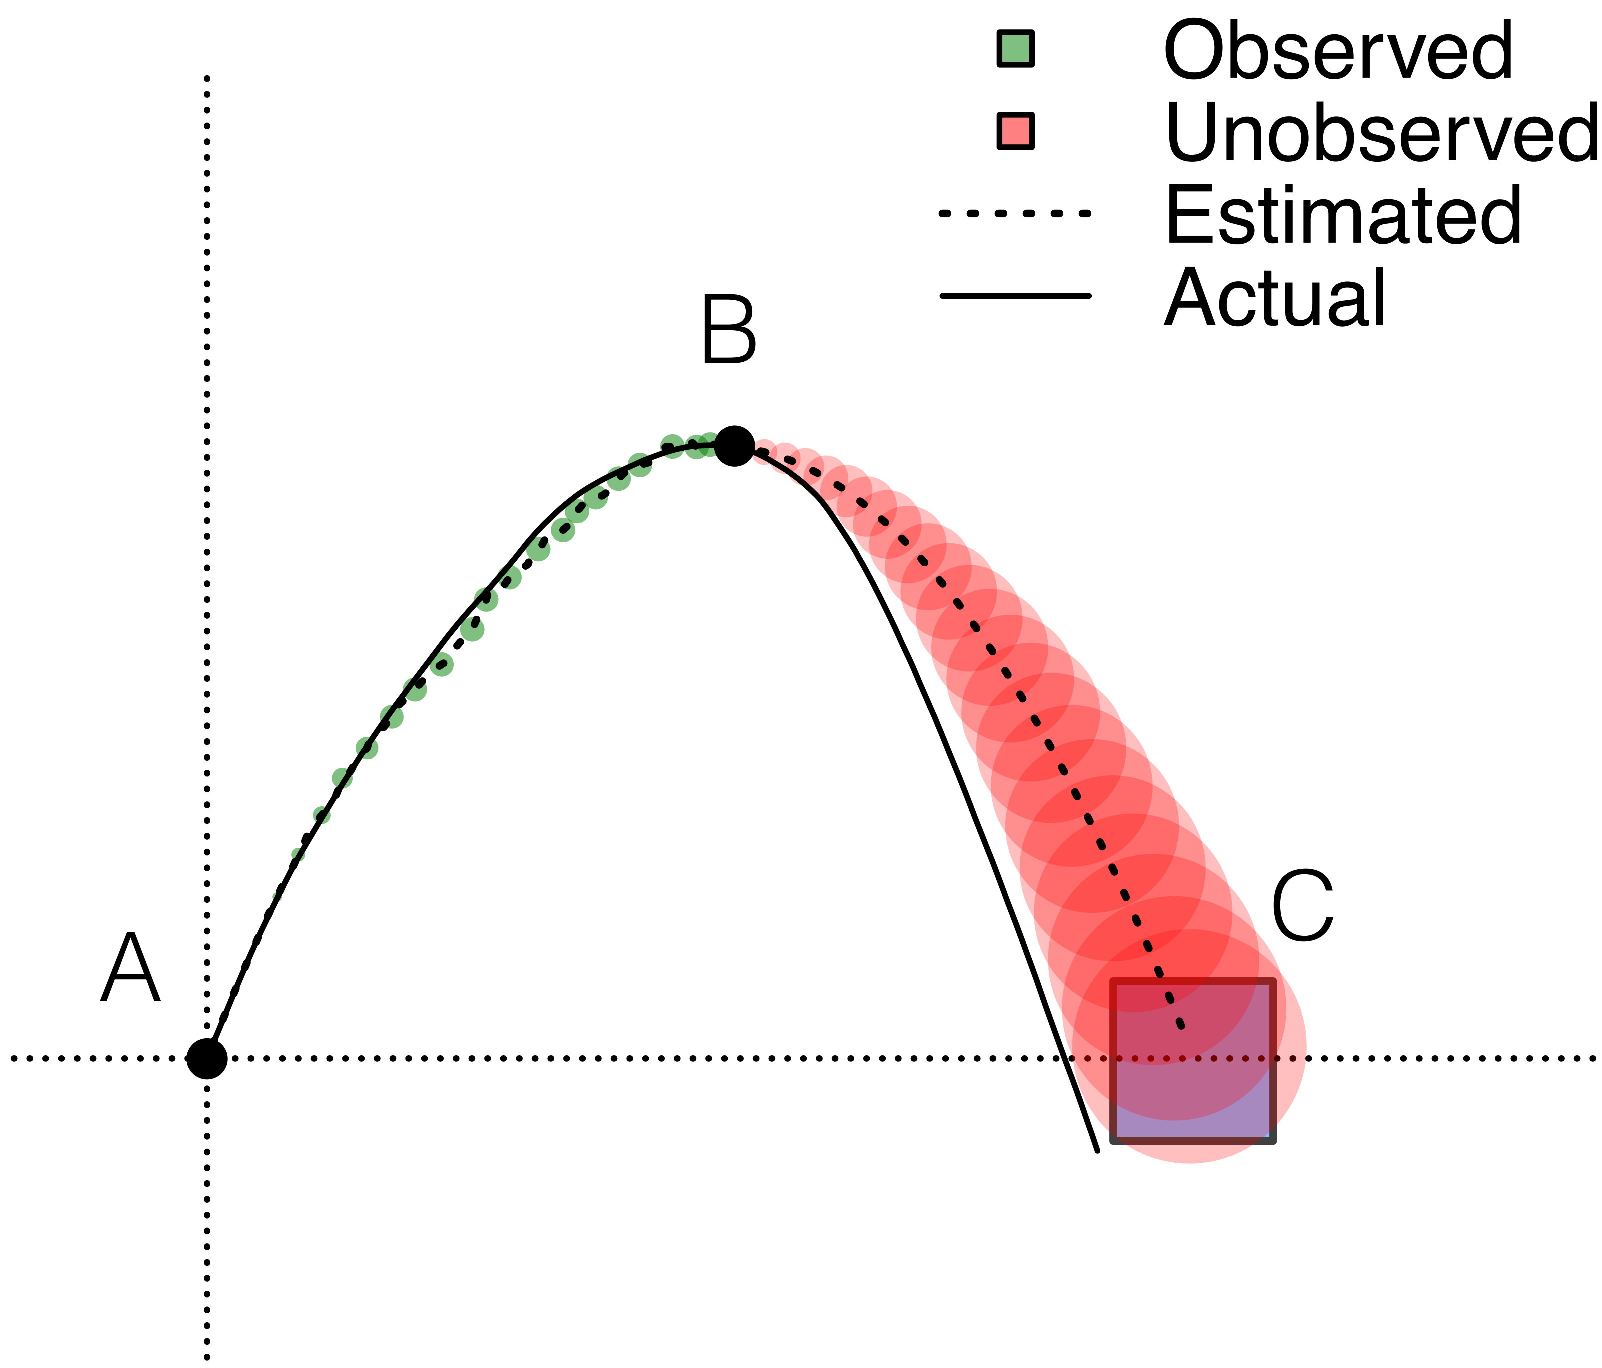

Sensorimotor coordination

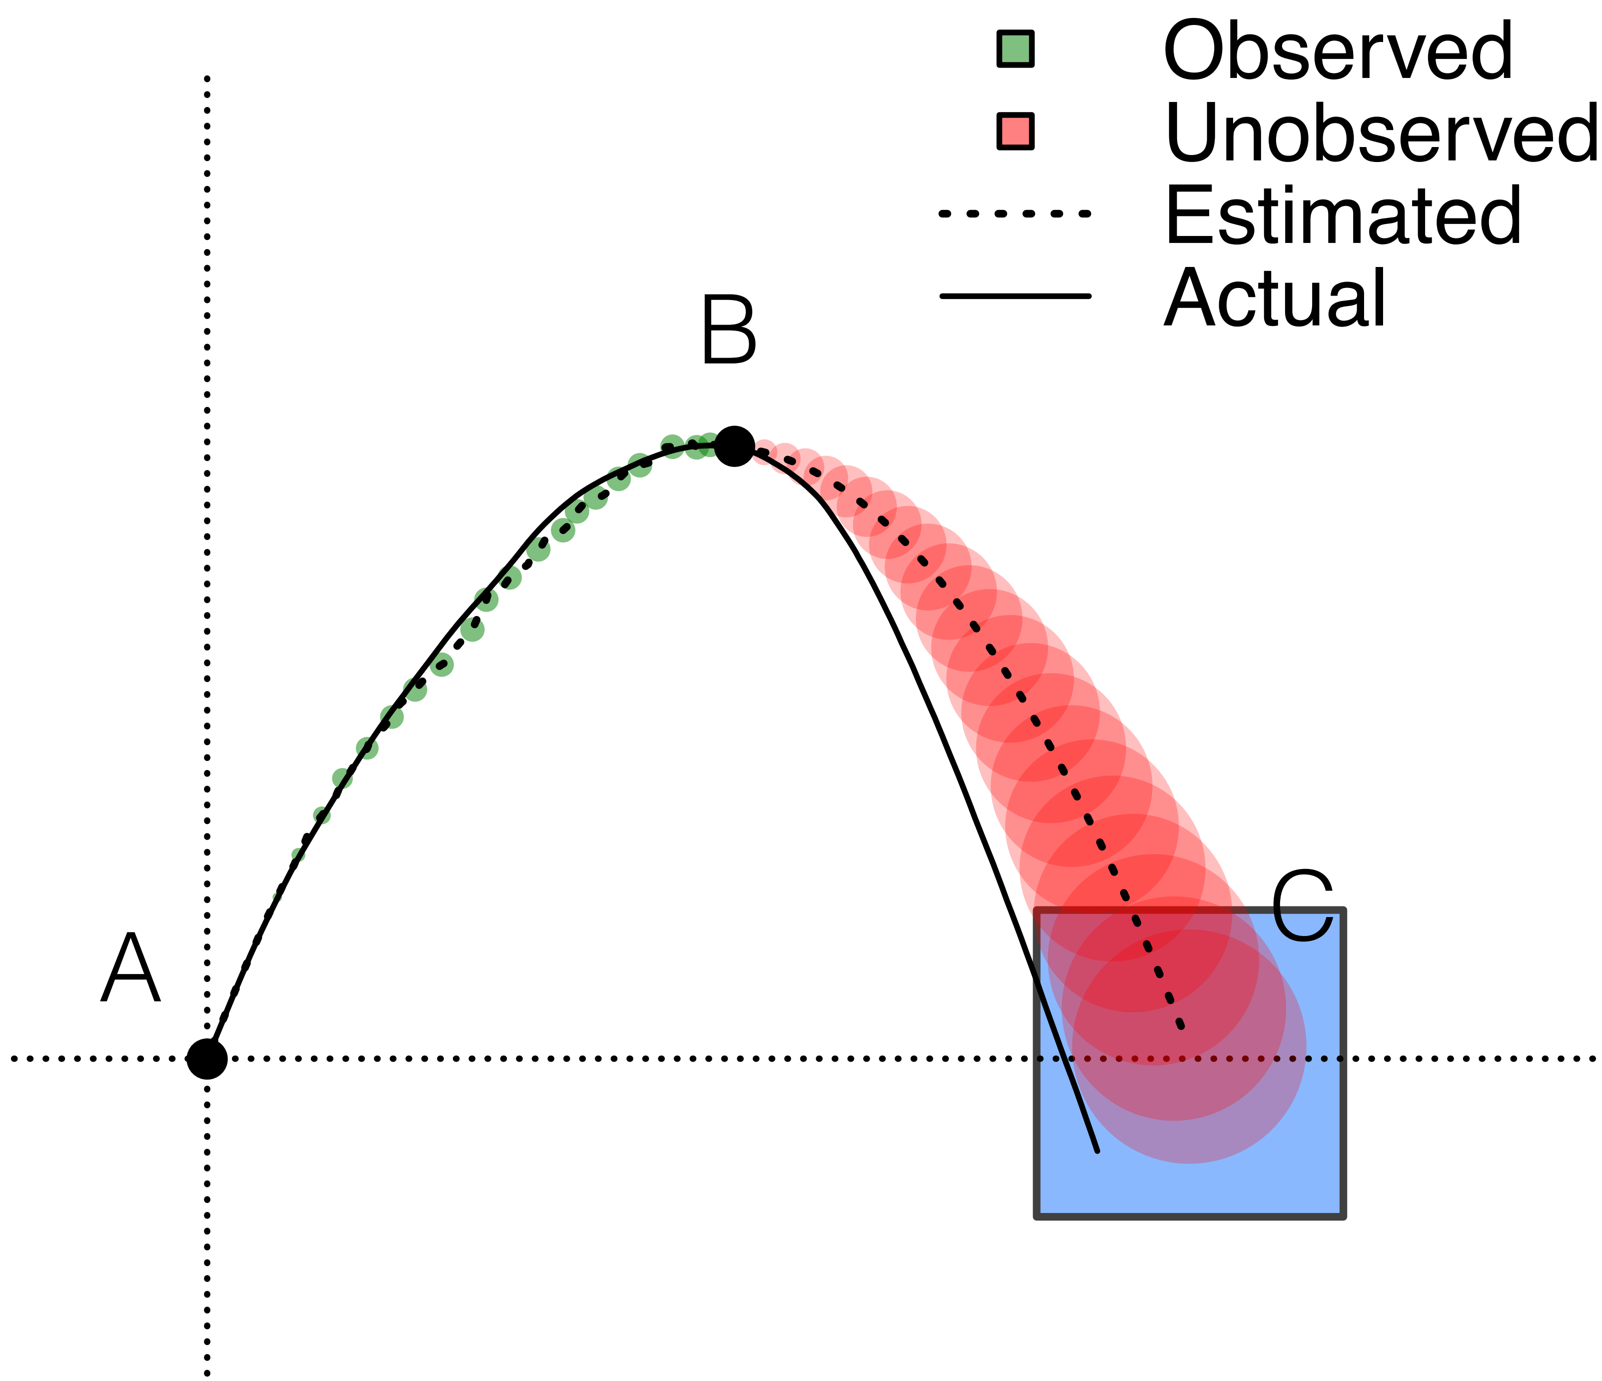

Experiment

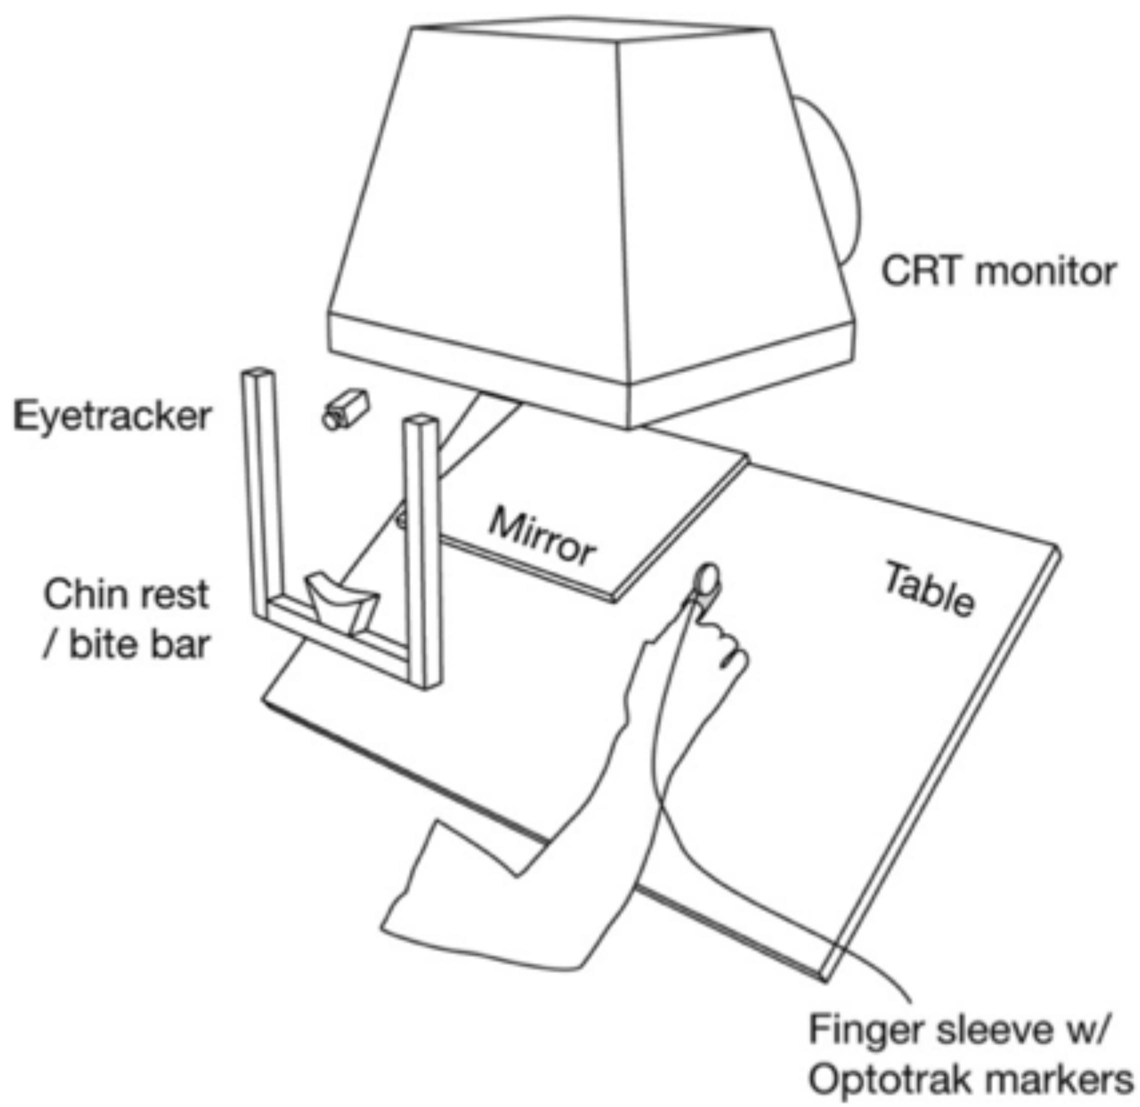

(Sims, Jacobs, & Knill, 2011)

- Task: Sort a bunch of objects into two piles
- Demands on vision:
  — Motor guidance
  — Information acquisition/planning
- Manipulate:
  — Difficulty of motor task
  — Difficulty of perceptual discrimination
- Examine adaptive timing of eye movements

Saccade timing as utility maximization

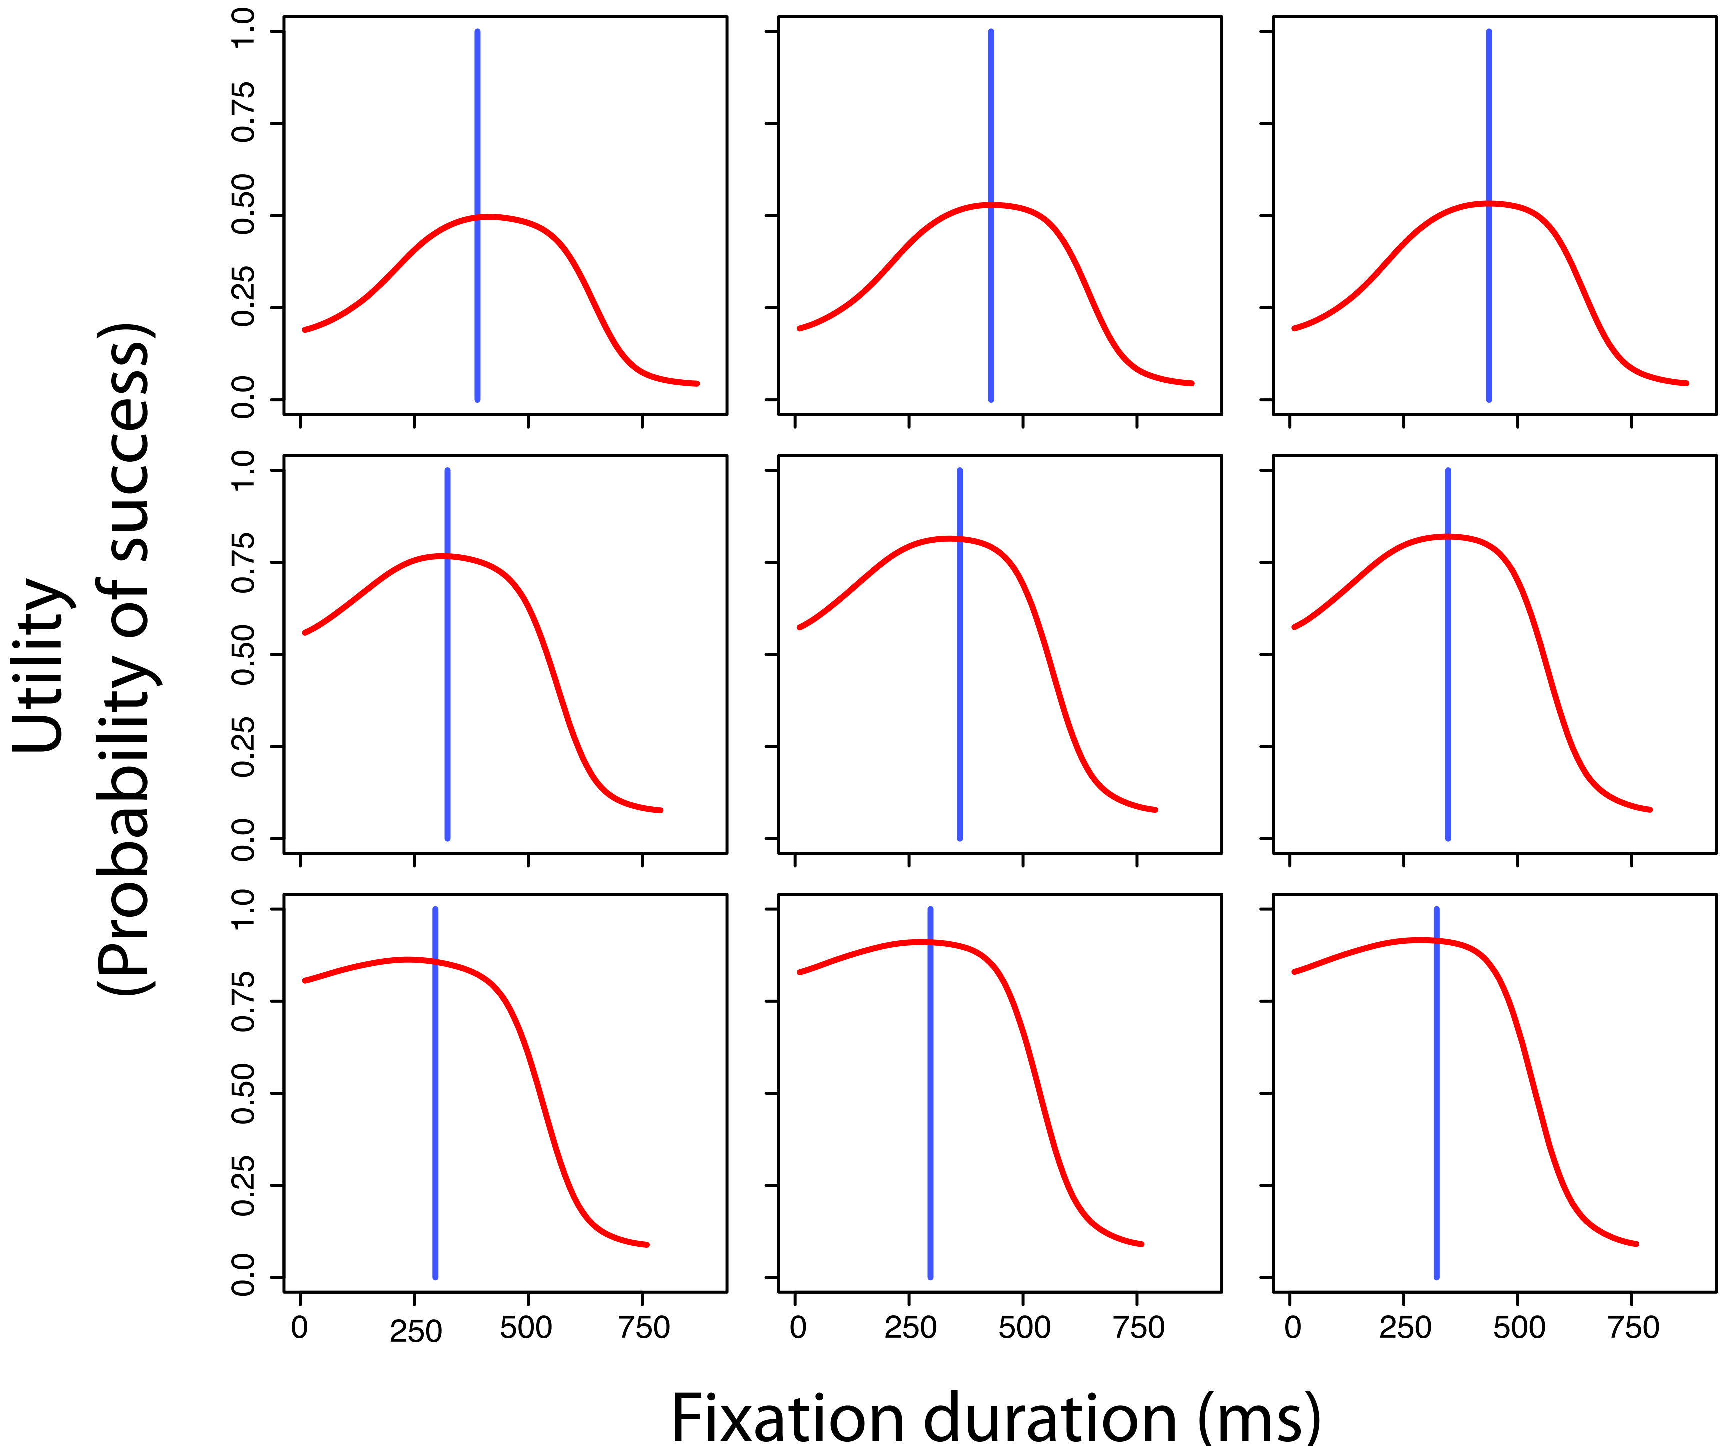

(Sims, Jacobs, & Knill, 2011)

How to understand sensorimotor behavior

Motor control is decision-making

Daniel M Wolpert[1] and Michael S Landy[2]

Motor behavior may be viewed as a problem of maximizing the utility of movement outcome in the face of sensory, motor and task uncertainty. Viewed in this way, and allowing for the availability of prior knowledge in the form of a probability distribution over possible states of the world, the choice of a movement plan and strategy for motor control becomes an application of statistical decision theory. This point of view has proven successful in recent years in accounting for movement under risk, inferring the loss function used in motor tasks, and explaining motor behavior in a wide variety of circumstances.

(Wolpert & Landy, 2012)

II. Perceptual memory

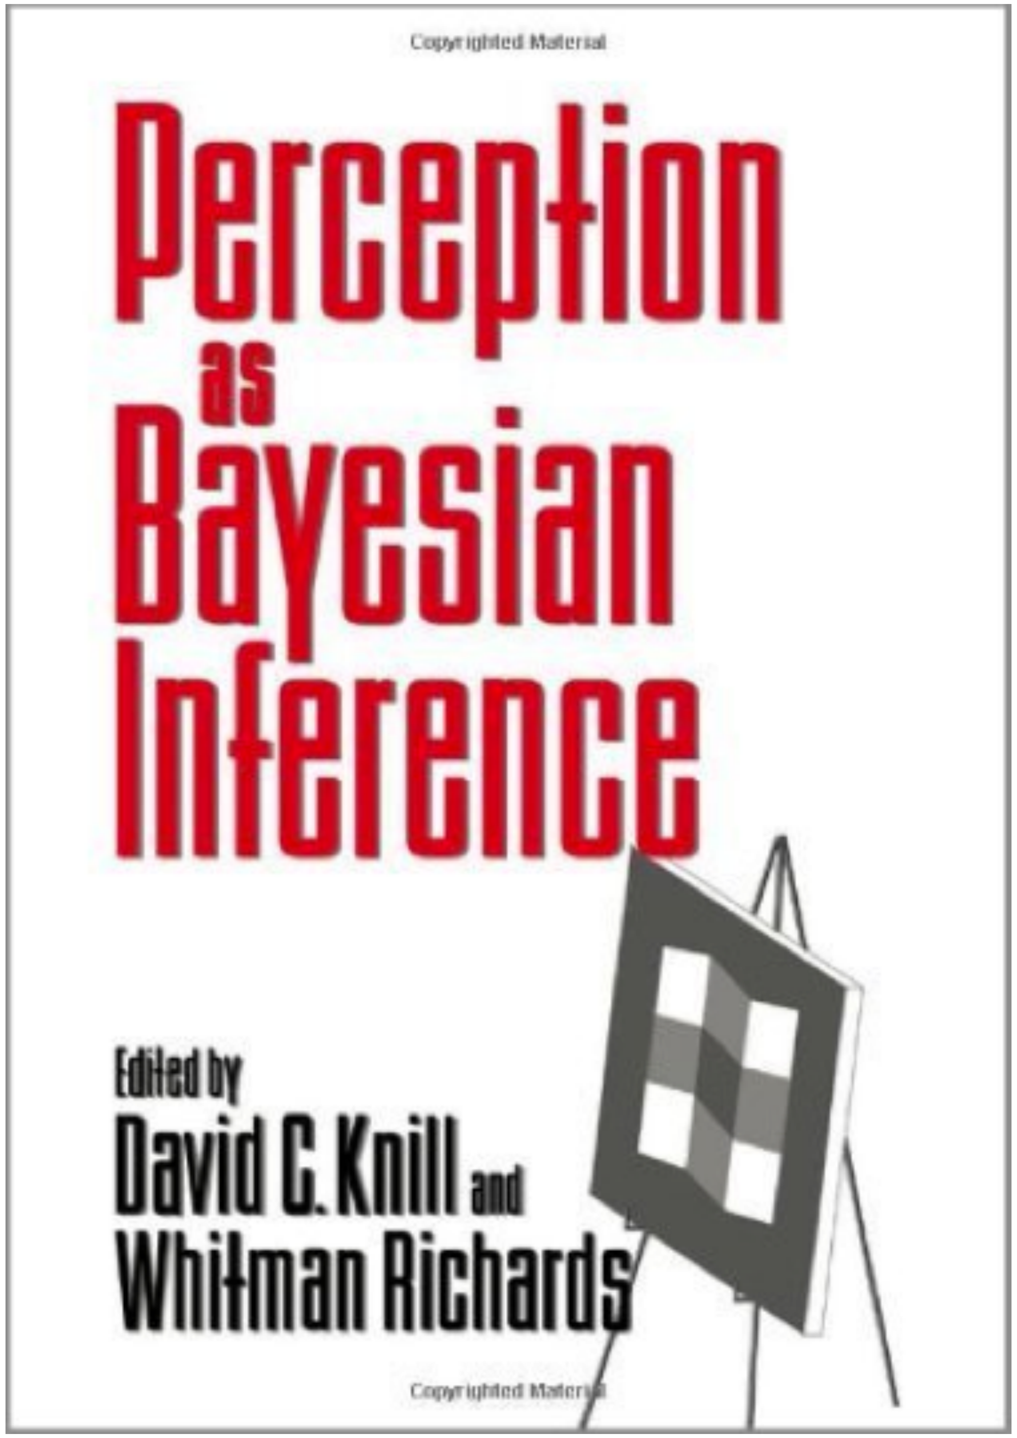

$\Rightarrow$ **Memory as Bayesian inference?**

Memory as Bayesian inference?

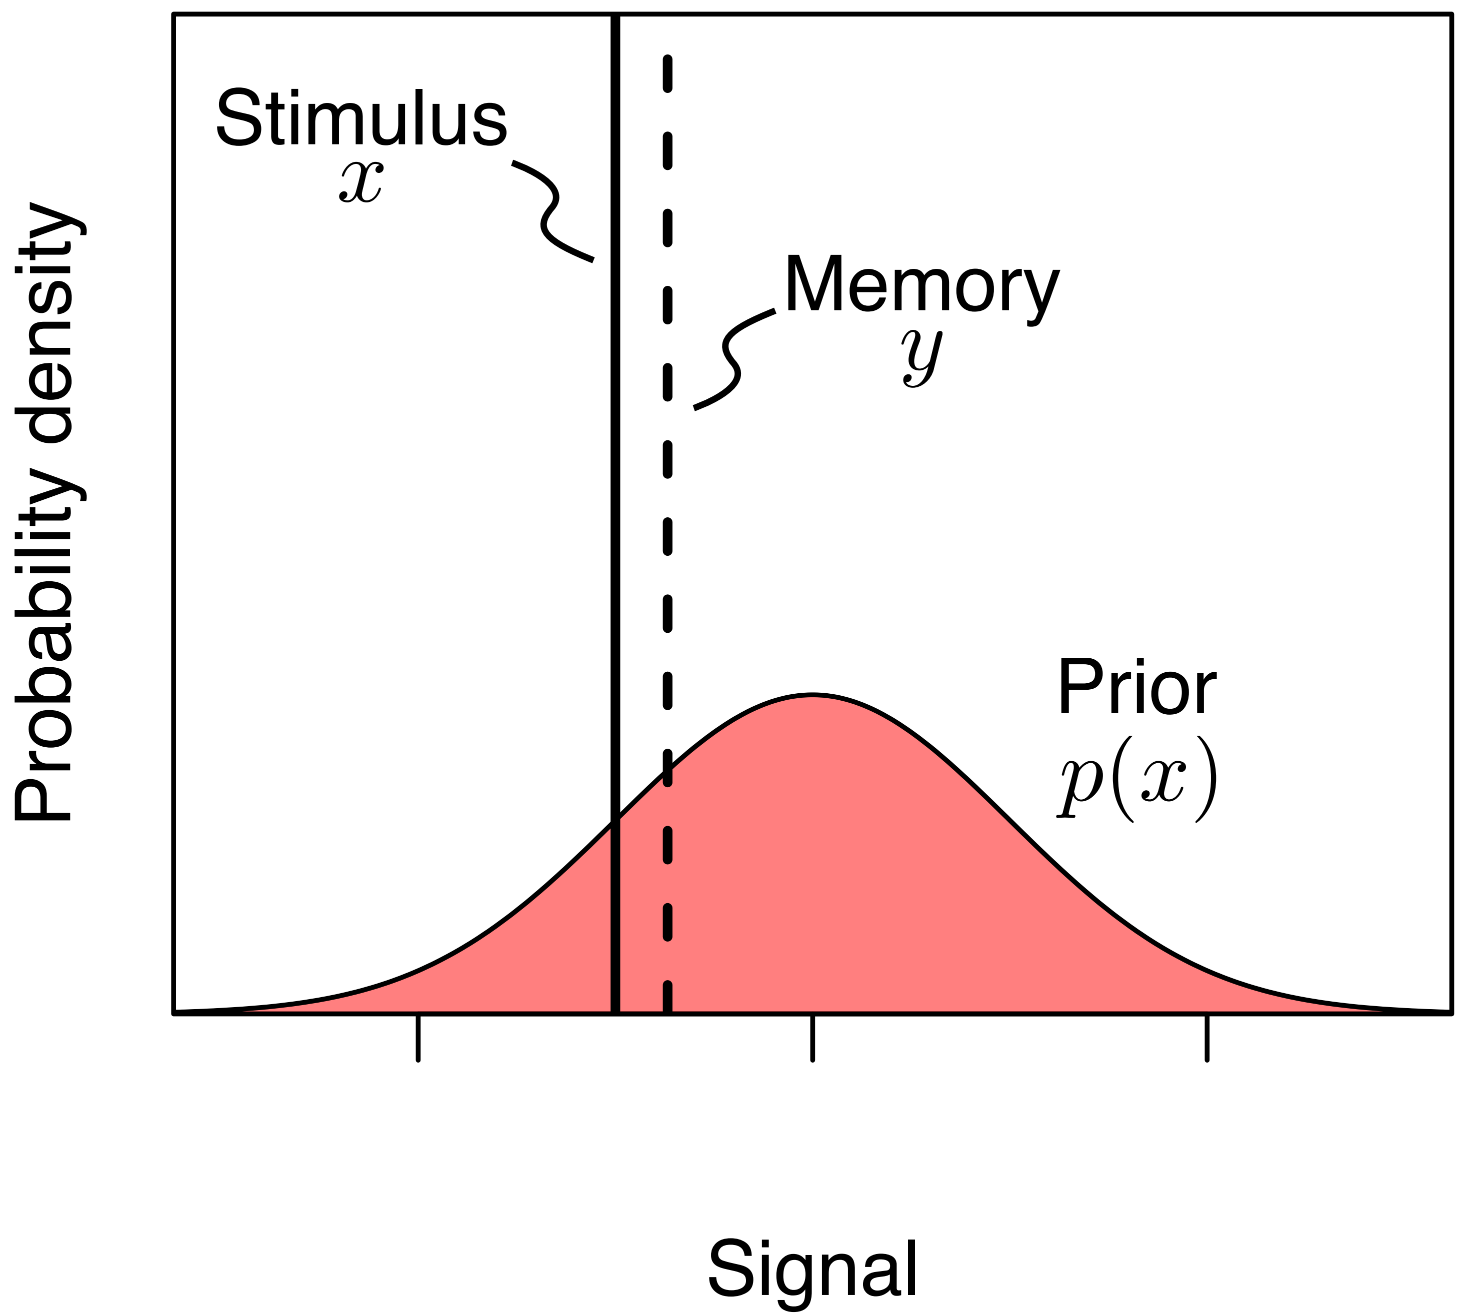

Memory as Bayesian inference?

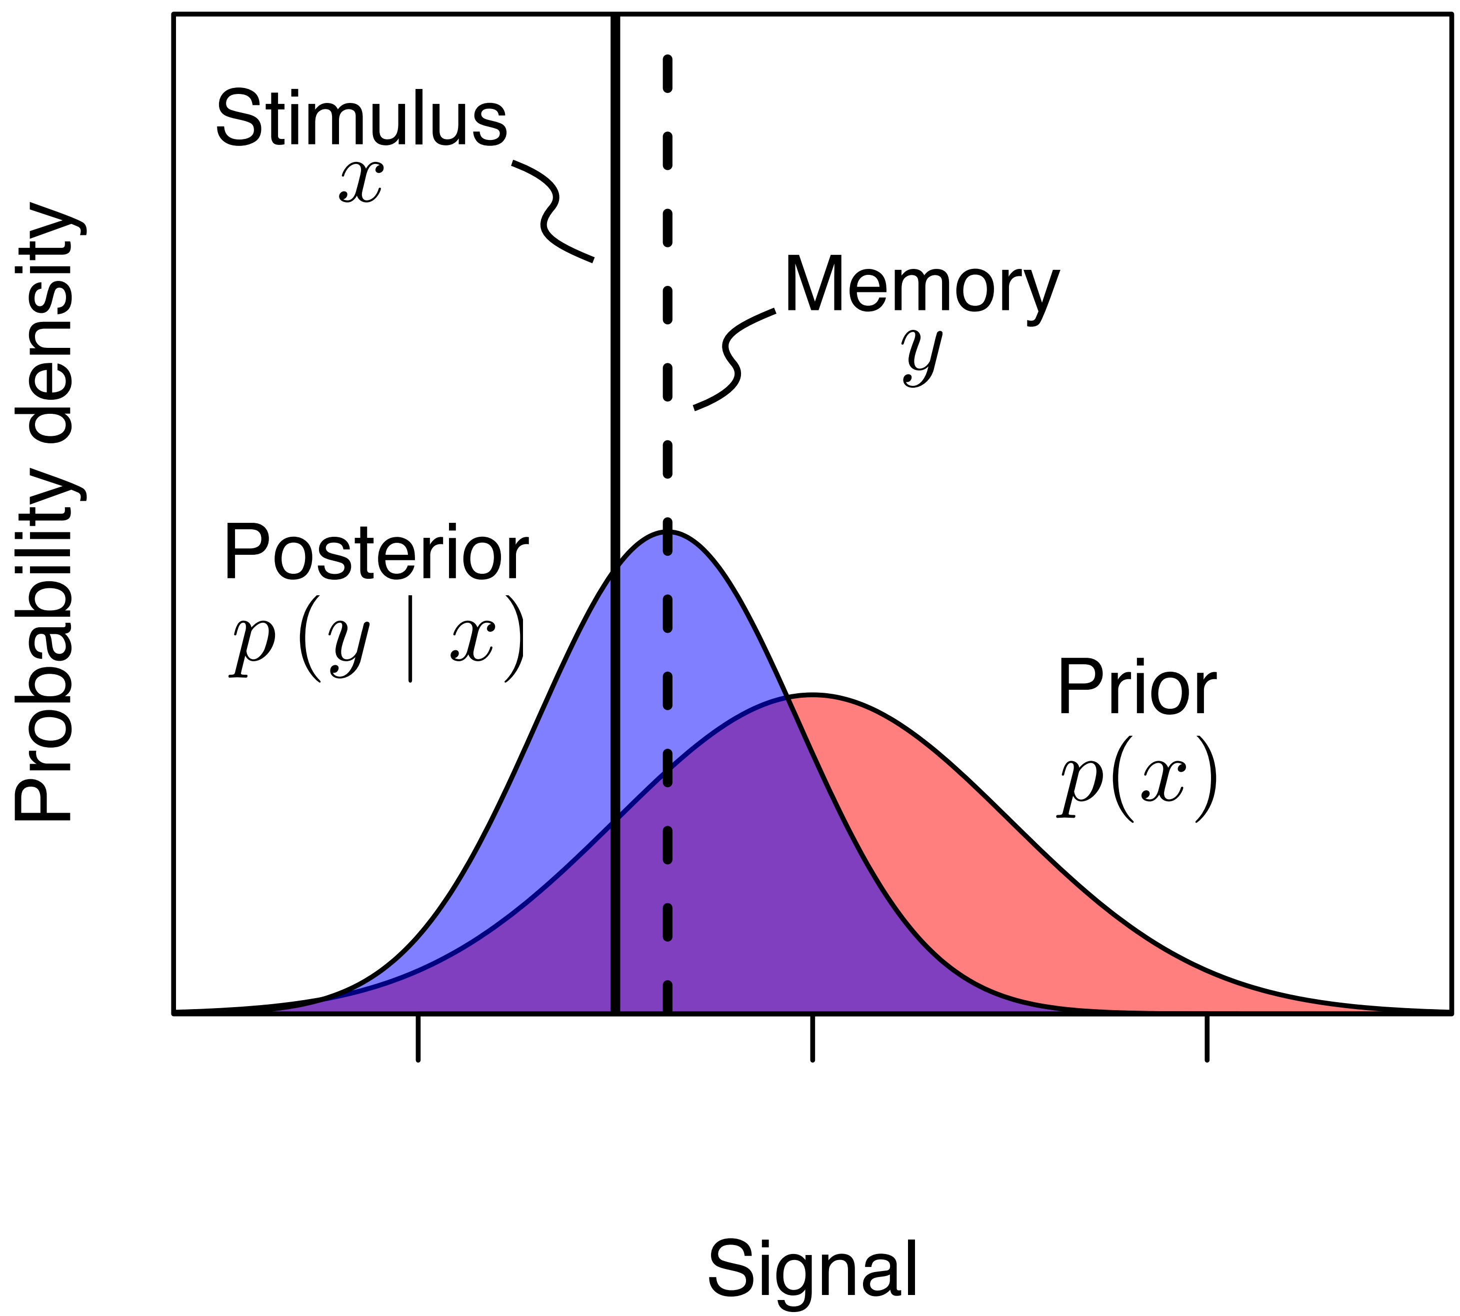

Memory as Bayesian inference?

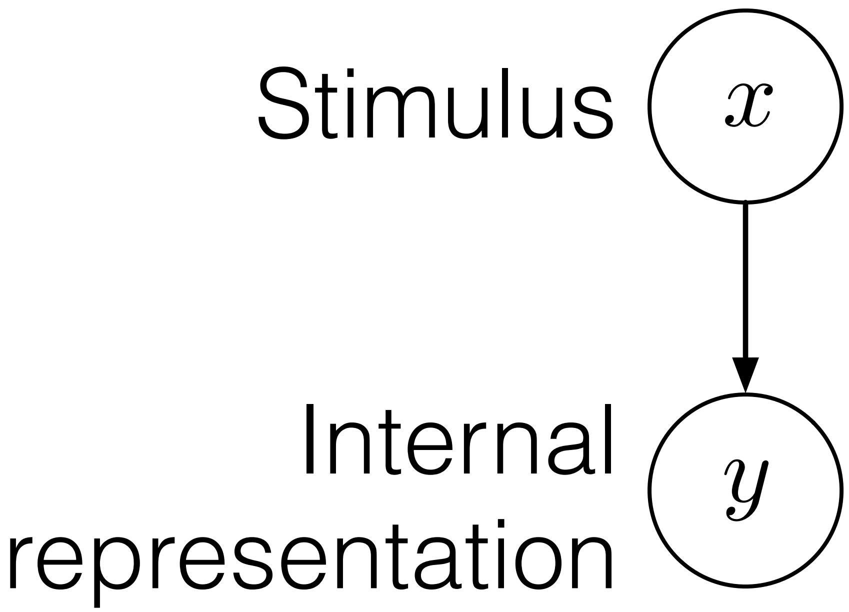

$$x \sim \mathrm{Normal}\left(\mu, \sigma_x^2\right)$$

$$y \mid x \sim \mathrm{Normal}\left(x, \sigma_y^2\right)$$

Memory = $p(x \mid y)$?

Memory as Bayesian inference?

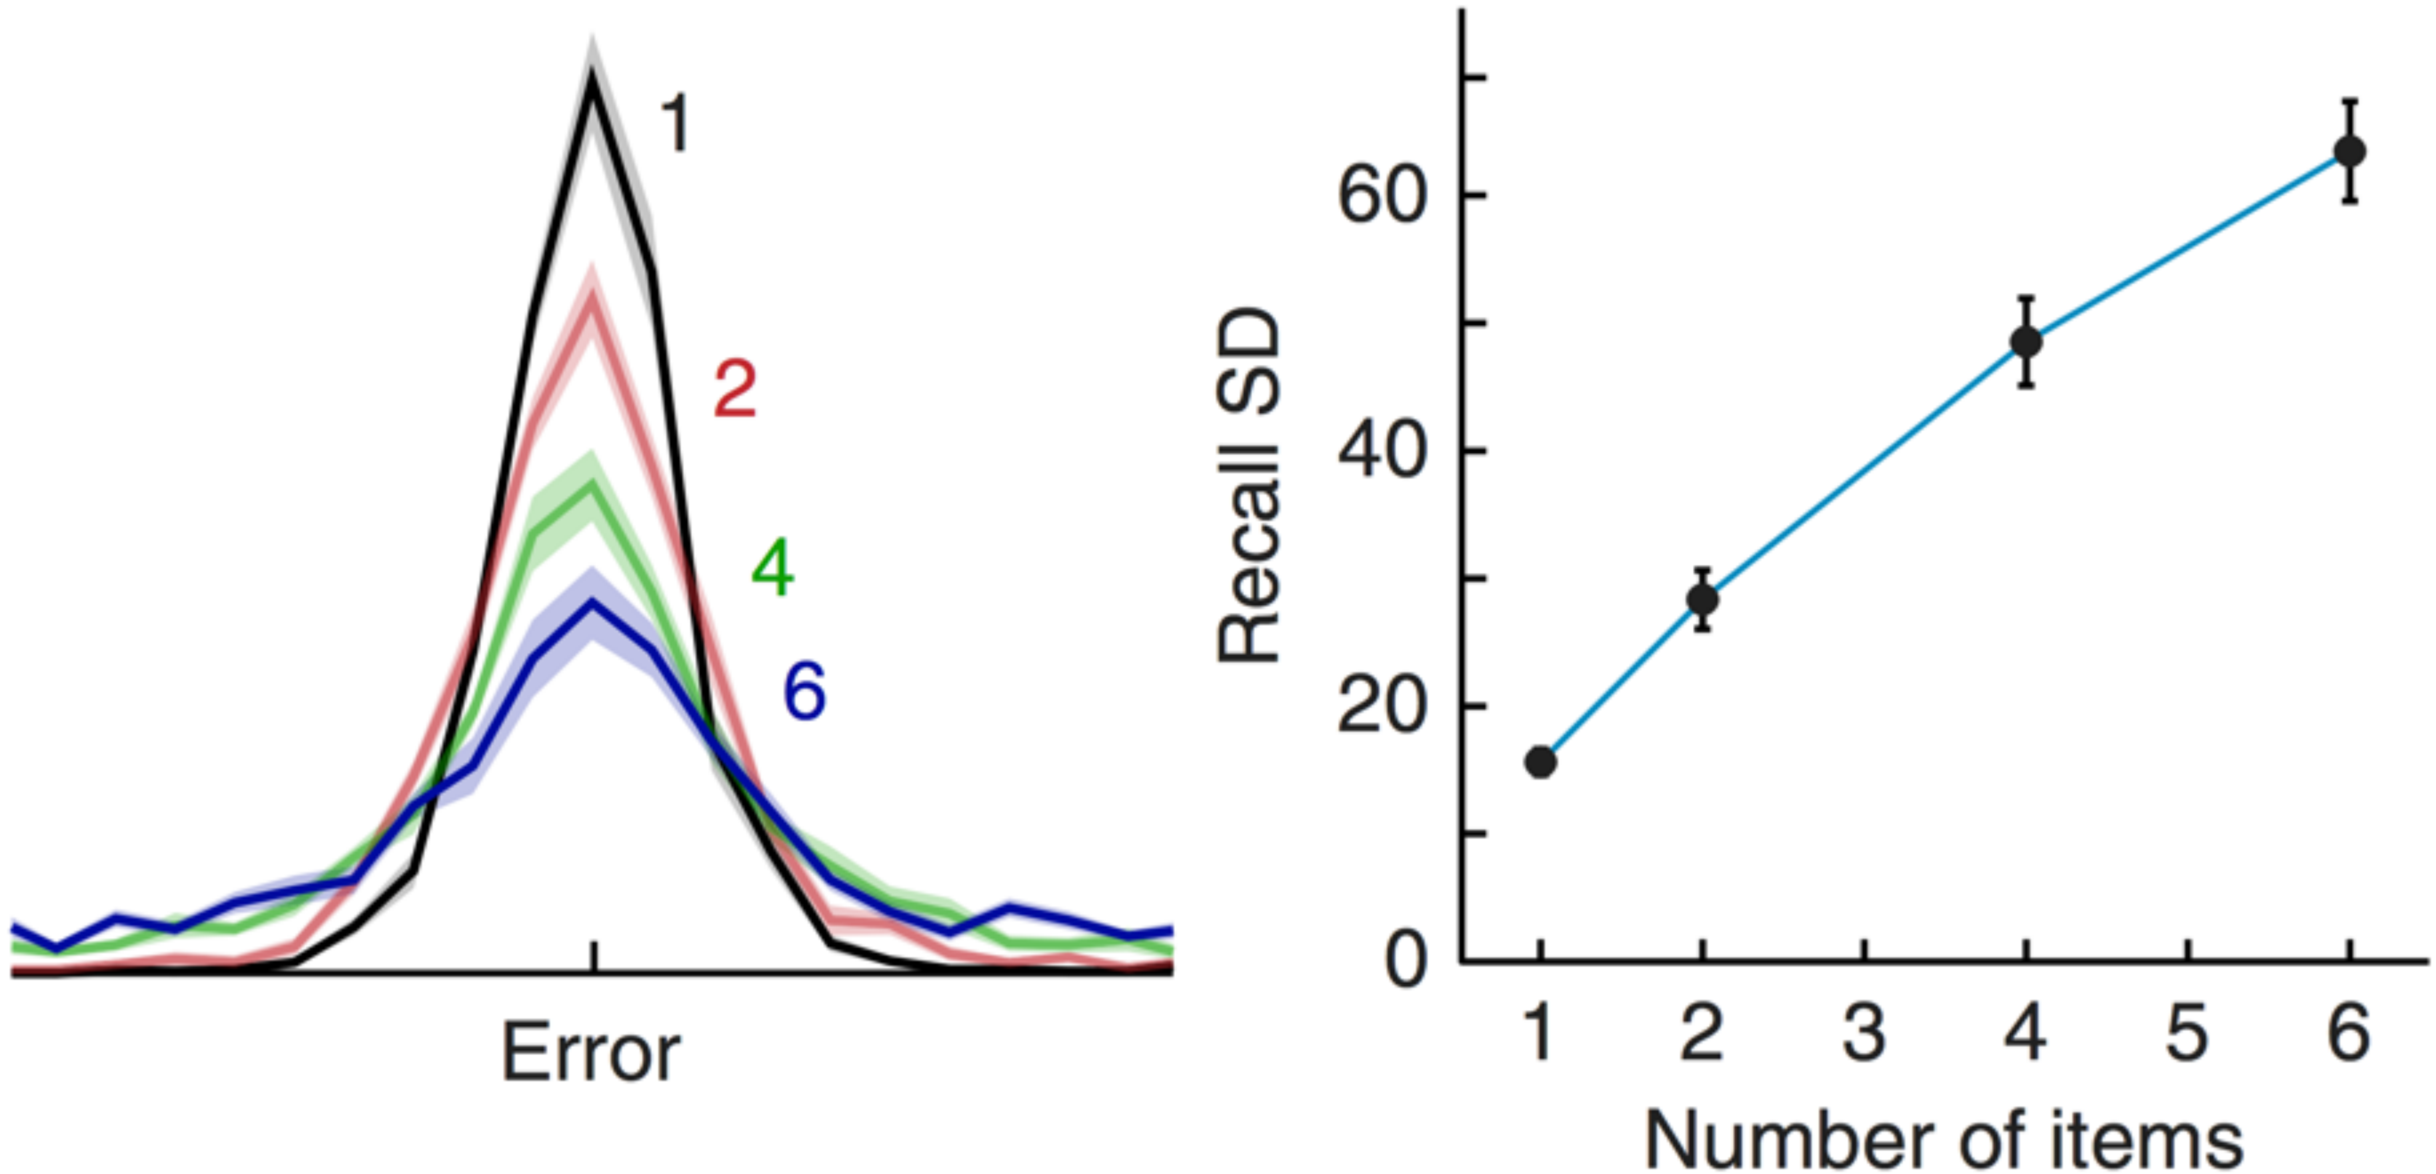

(Ma, Husain, & Bays, 2014)

Memory as ~~Bayesian inference~~ *efficient communication*

$$x \longrightarrow \boxed{} \longrightarrow y$$

$p(x)$: Visual statistics

$\mathcal{L}(x, y)$: Cost function

C : Channel capacity

Goal: Minimize $\mathcal{L}(x, y)$ w.r.t. $p(y \mid x)$
subject to $I(x, y) \leq C$

Memory as efficient communication

$$p(x) = \mathrm{Normal}(\mu, \sigma)$$

$$\mathcal{L} = (y - x)^2$$

$$C = \frac{C_{tot}}{n}$$

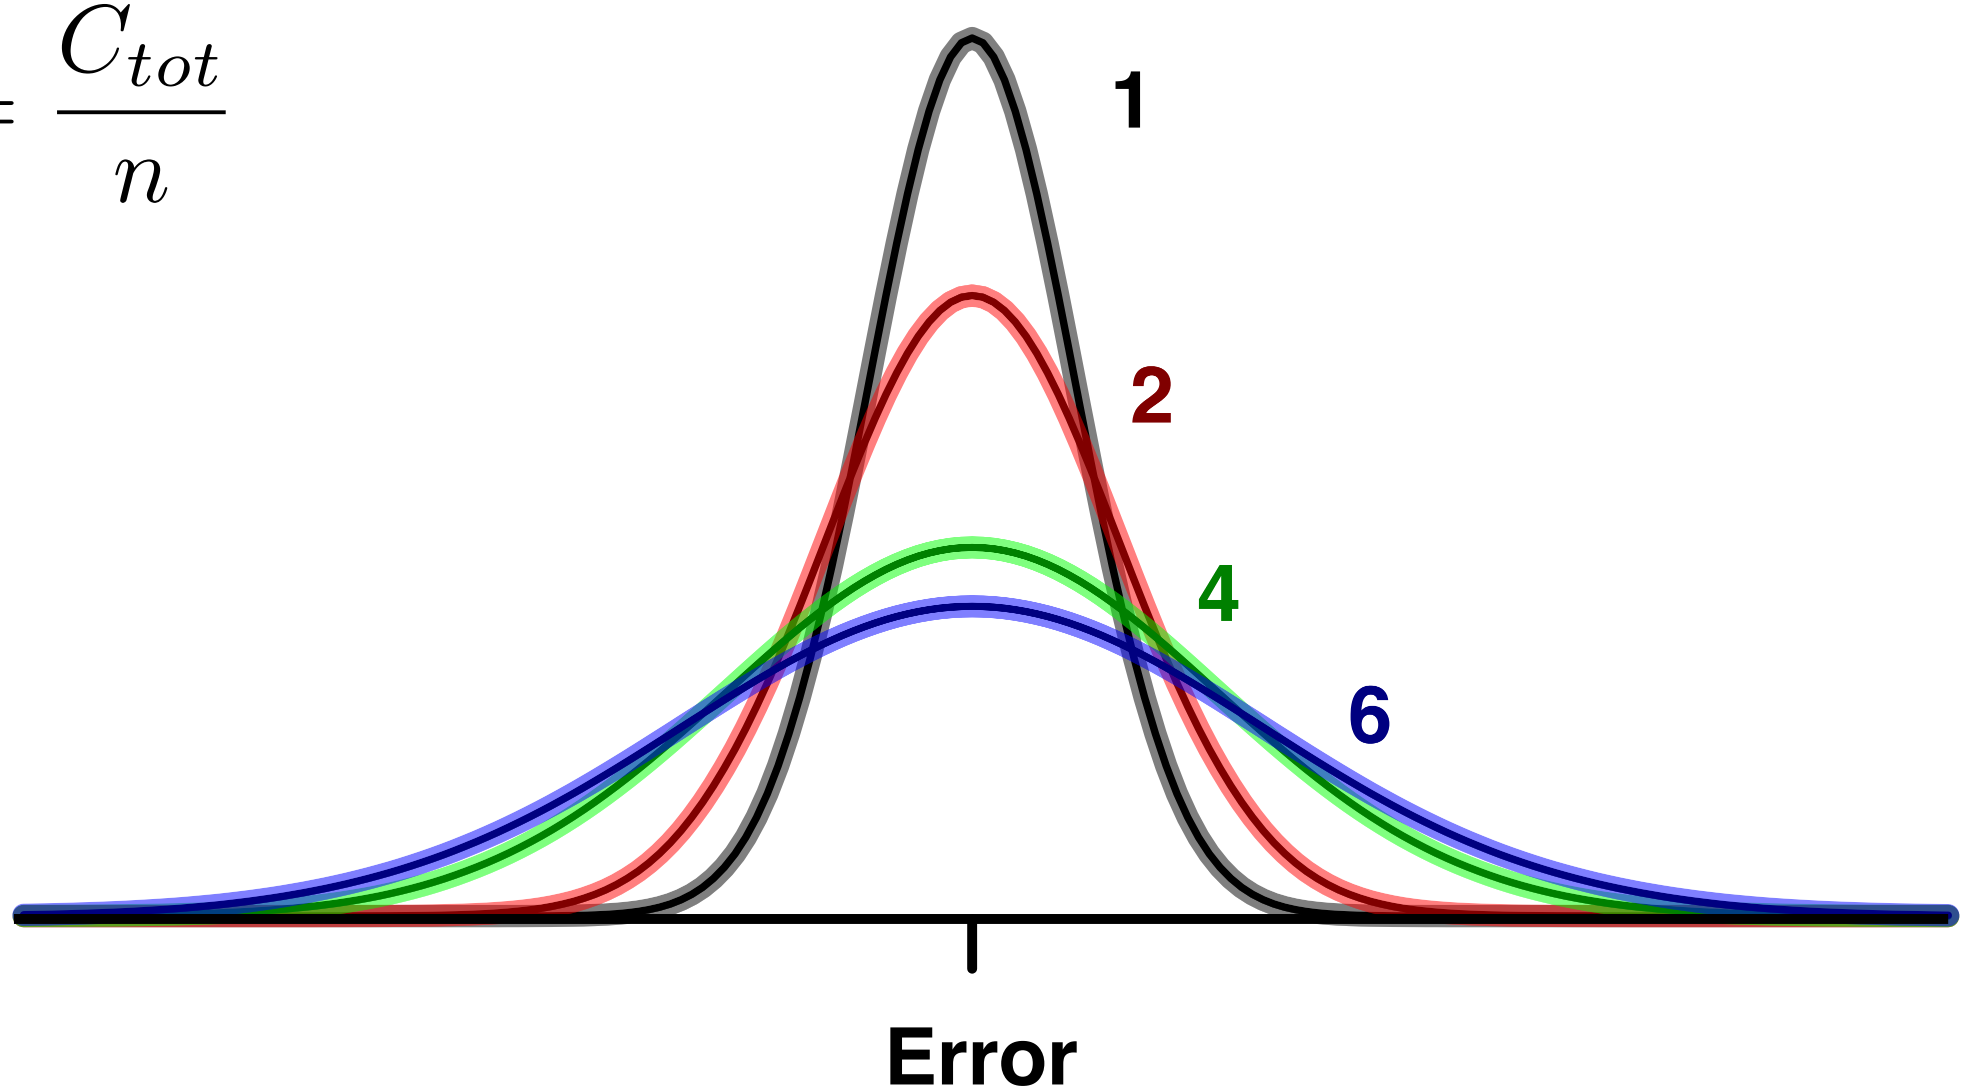

What is the cost of misremembering?

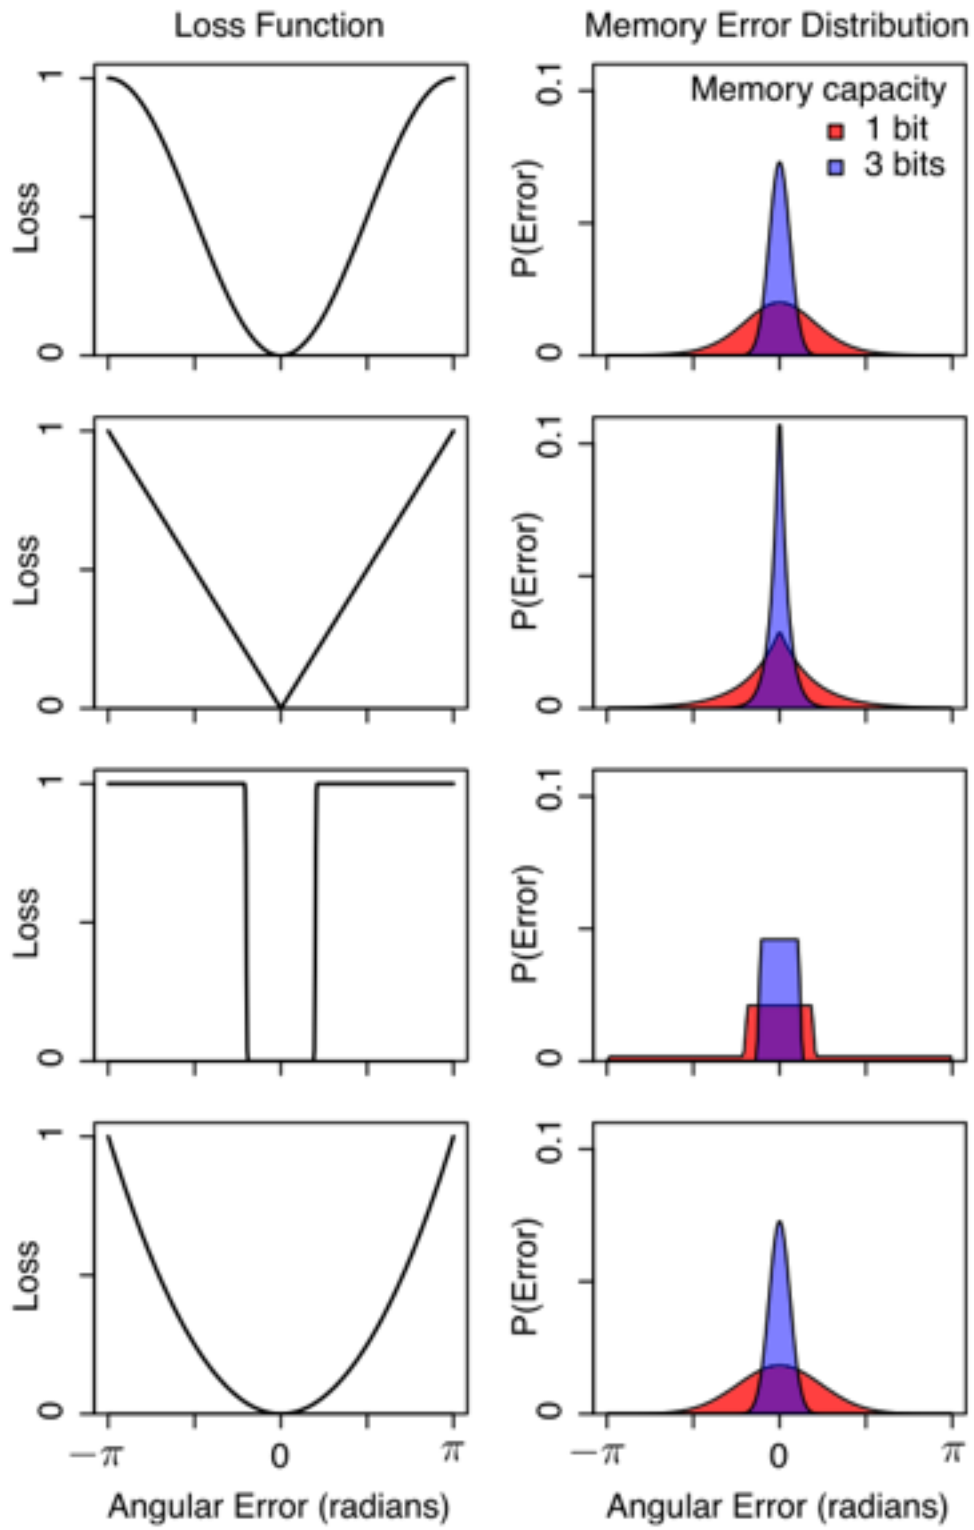

- Different cost functions imply different optimal distributions of memory error, given the same channel capacity

What is the cost of misremembering?

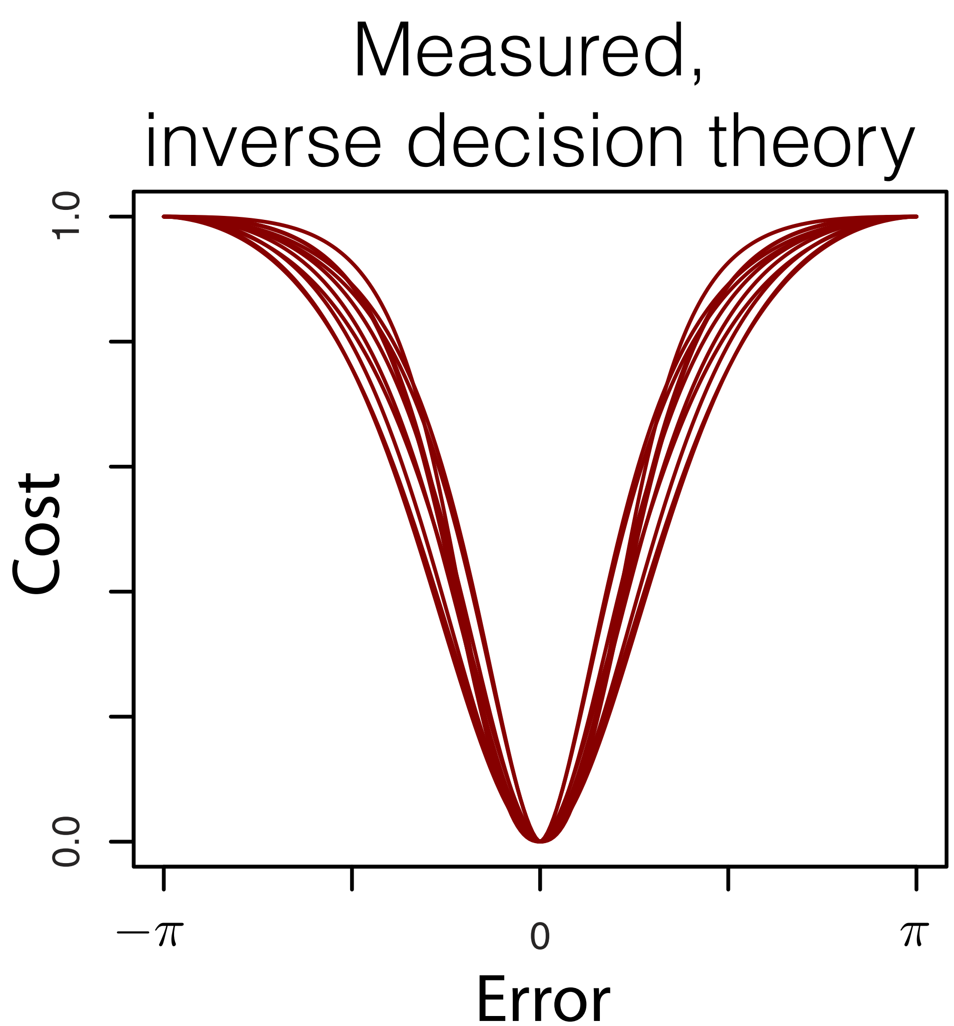

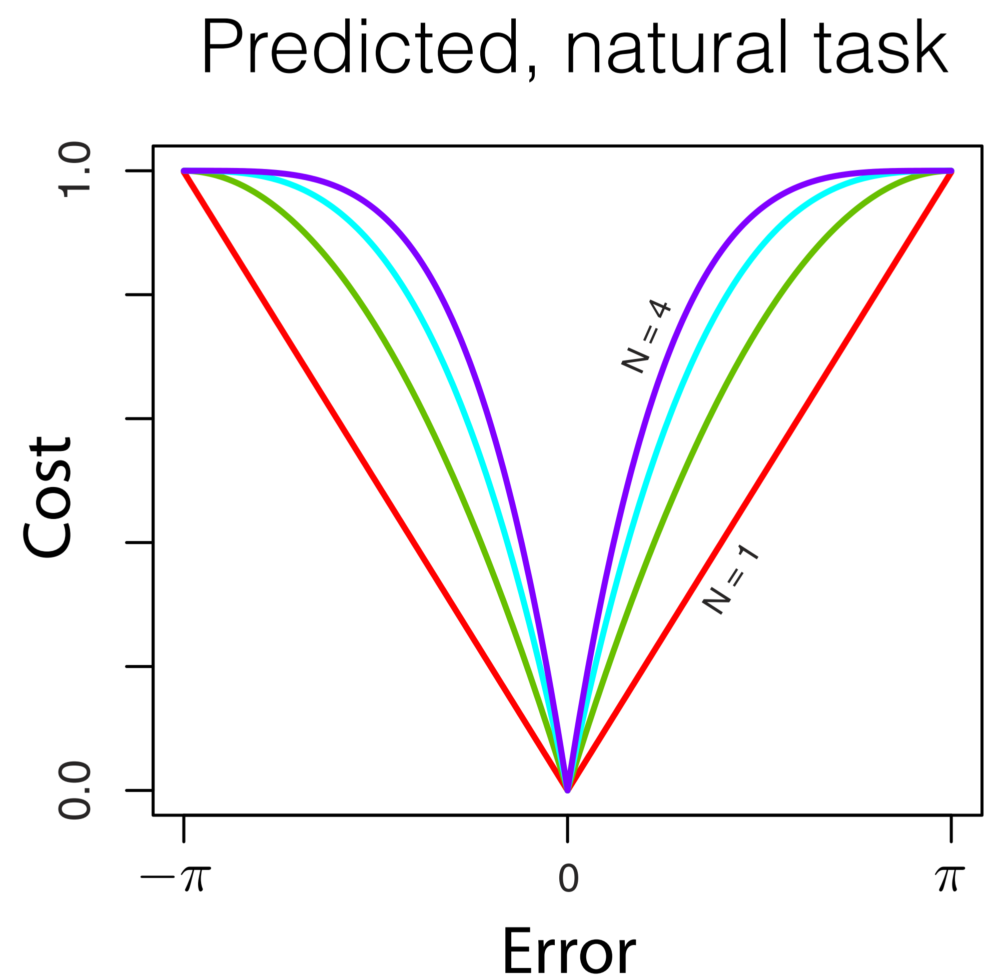

(Sims, 2015; JOV)

Looking back and moving forward

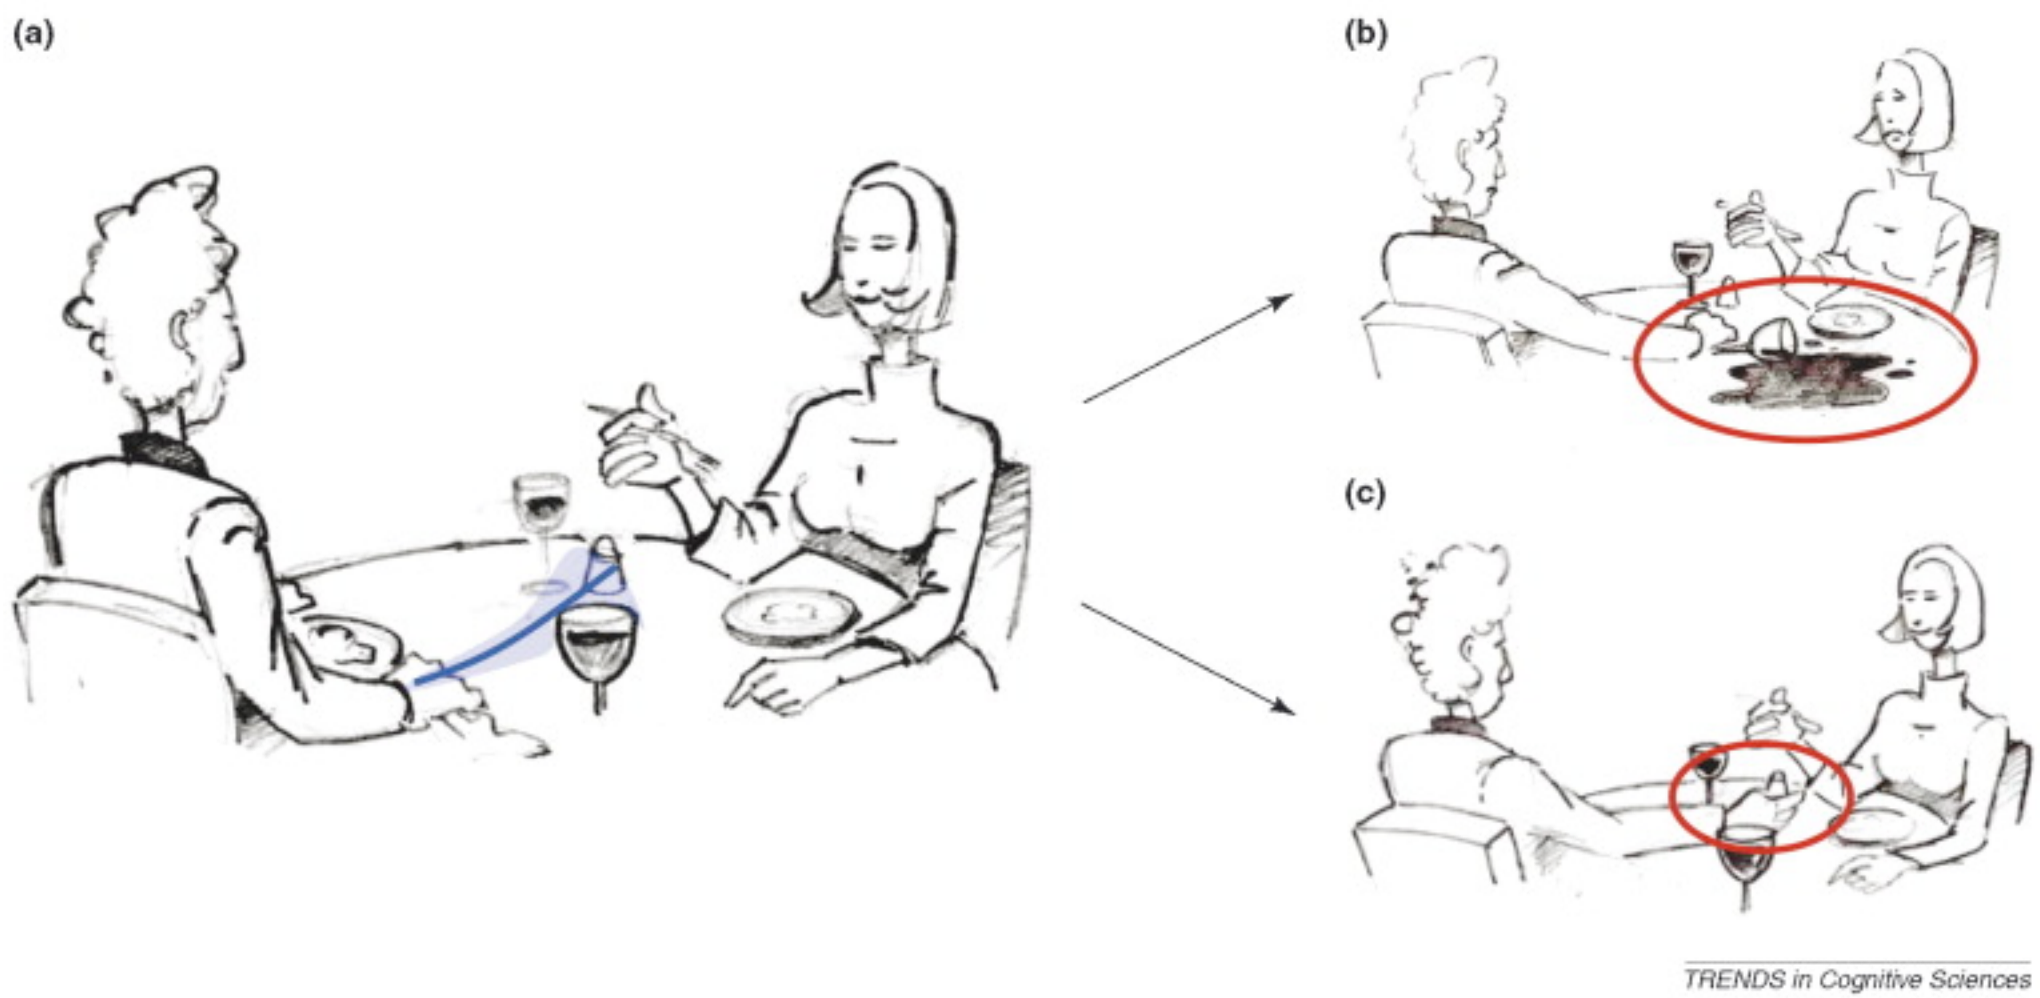

(Trommershaüser, Maloney, & Landy, 2008)

Vision lost and regained

Impact on direction of heading estimates from optic flow

Laurel Issen, Krystel Huxlin and David Knill

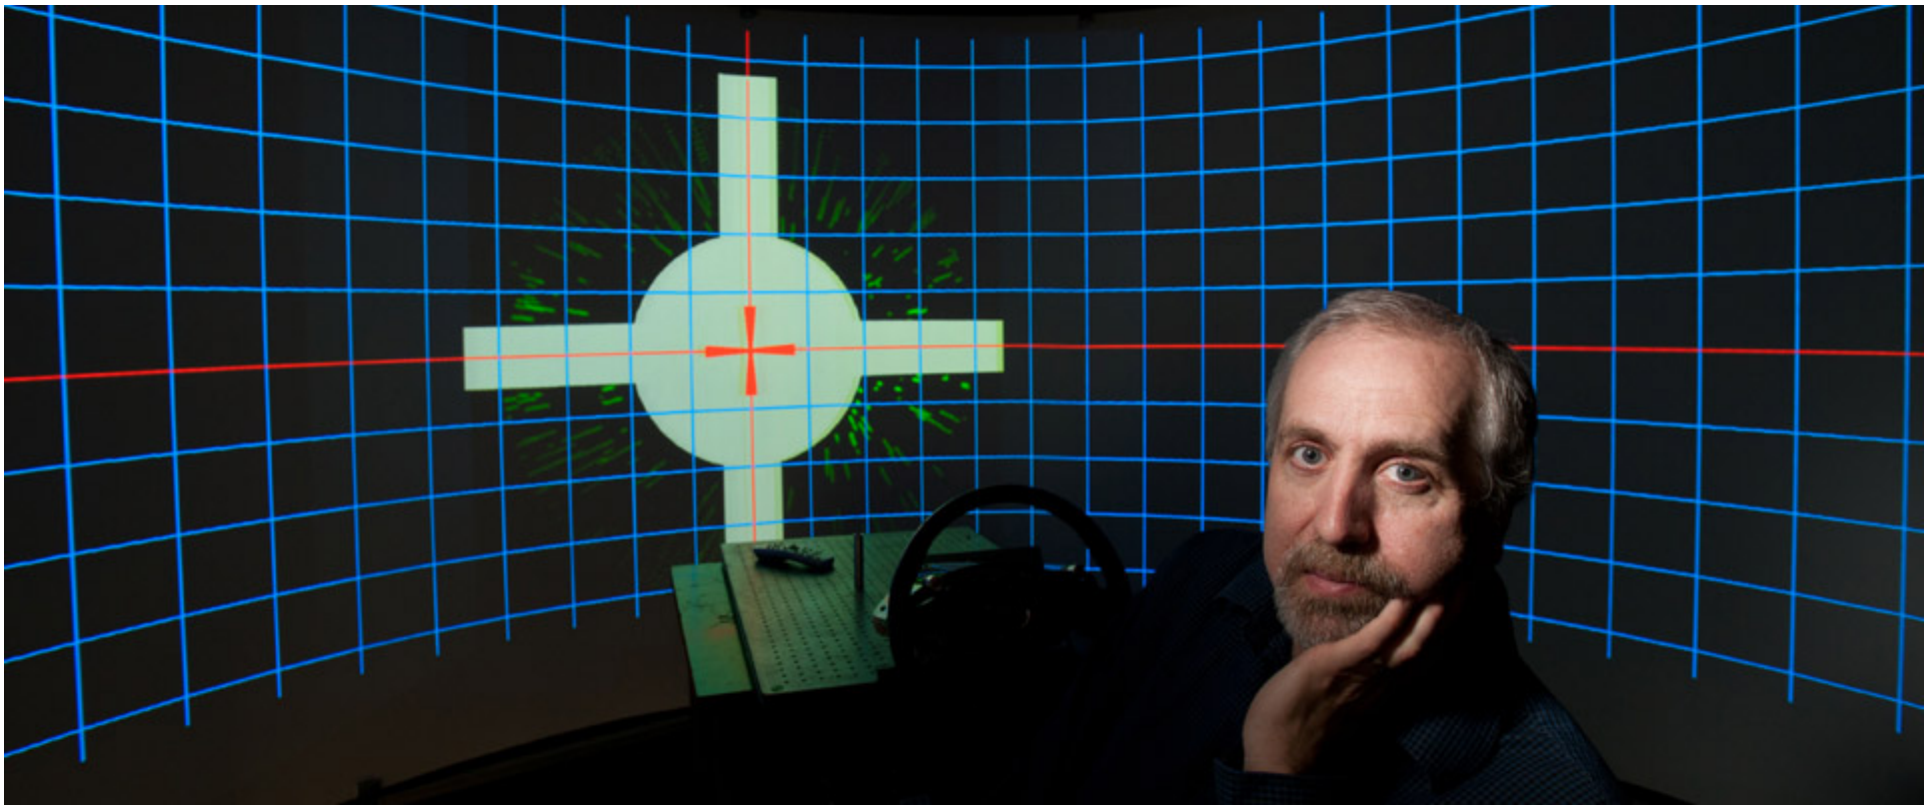

Vision lost after V1 damage

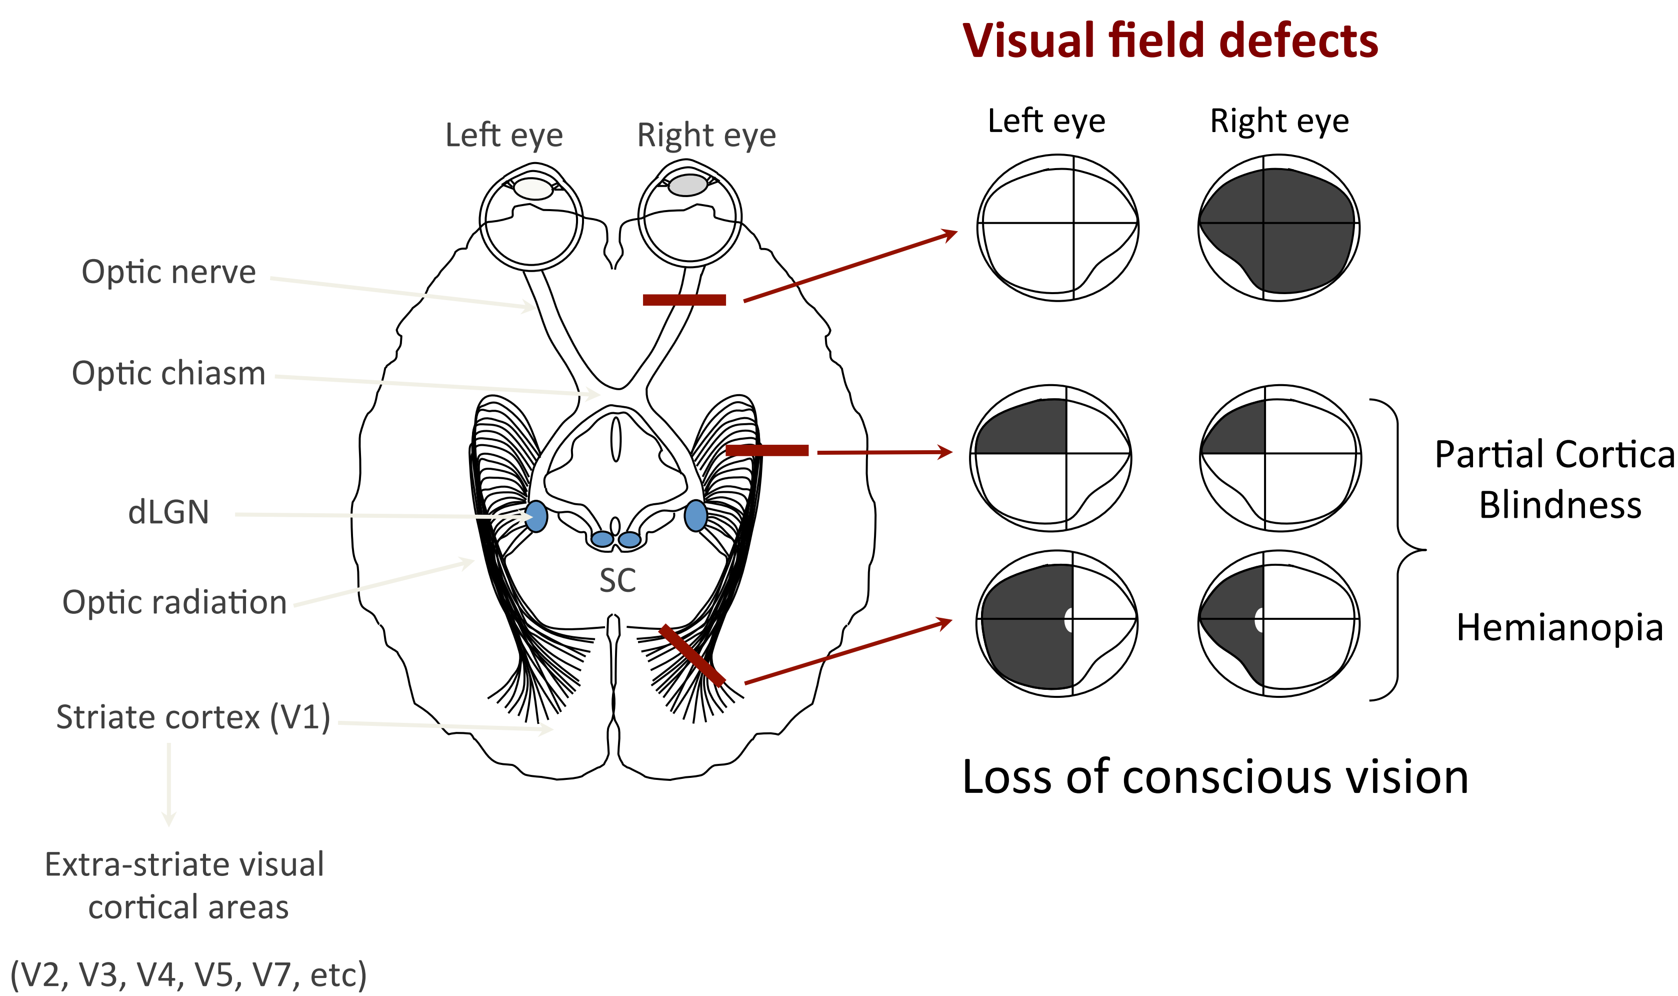

Loss of conscious vision

Damaging V1 – hemianopia

Causes:

Stroke – PCA, MCA

Tumors

Trauma, incl. TBI

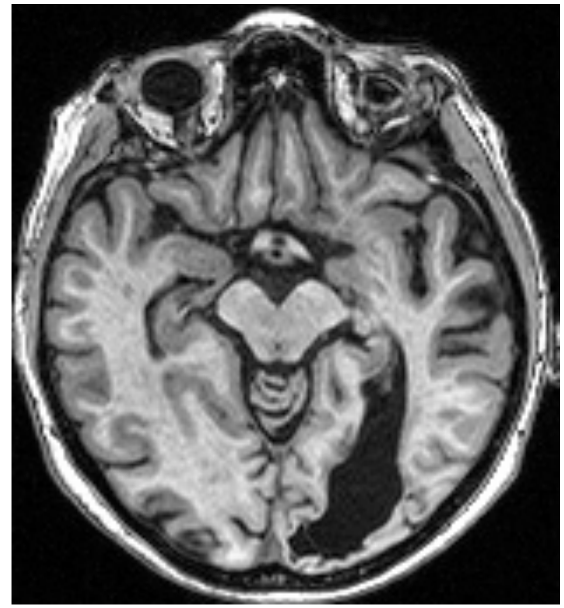

Incidence:

0.8% population > 49 yrs old

(Blue Mountains Eye Study, Australia)

Up to 50% of stroke victims

Prognosis:

Spontaneous improvements in first 2-3 months

Deficit stable and permanent after that

Dogma: blindness cannot be recovered

Why damaging V1 causes blindness?

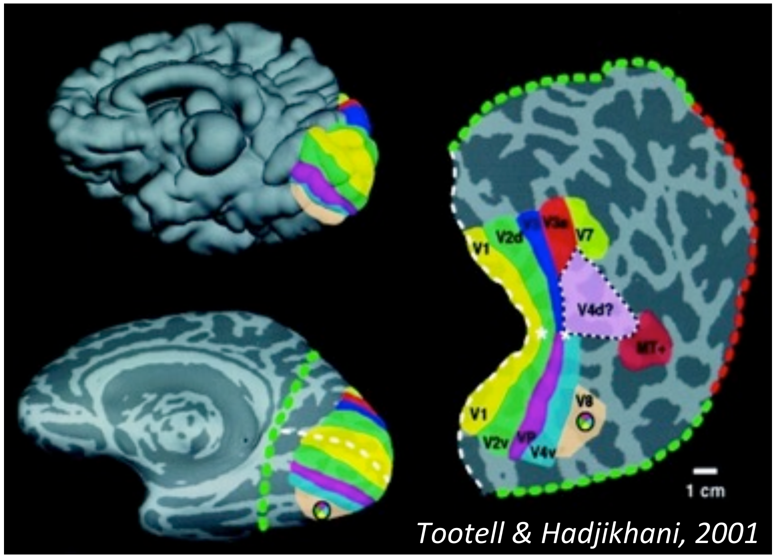

Tootell & Hadjikhani, 2001

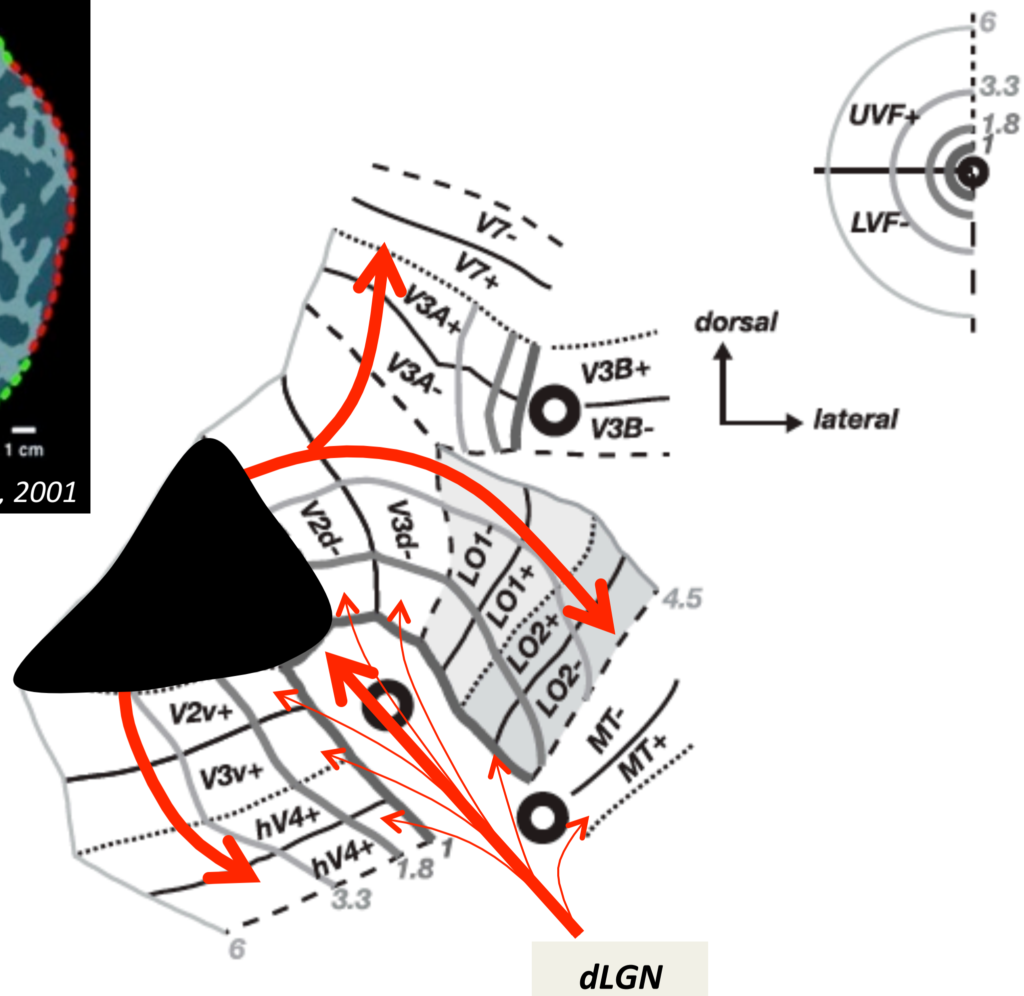

Modified from Larsson and Heeger, J. Neurosci. (2006)

Residual visual processing after V1 damage

Blindsight (Weiskrantz et al., 1974; Weiskrantz, 1986)

Unconscious ability to detect, match, discriminate orientation, wavelength, speed (Morland et al., 1999)

But "blindsight" is not seeing

In spite of blindsight, even *unilateral* V1 damage dramatically alters visually-guided functions in daily life:

–Difficulties reading

–Inability to drive

–Bumping into objects

–Difficulties navigating

WHY?

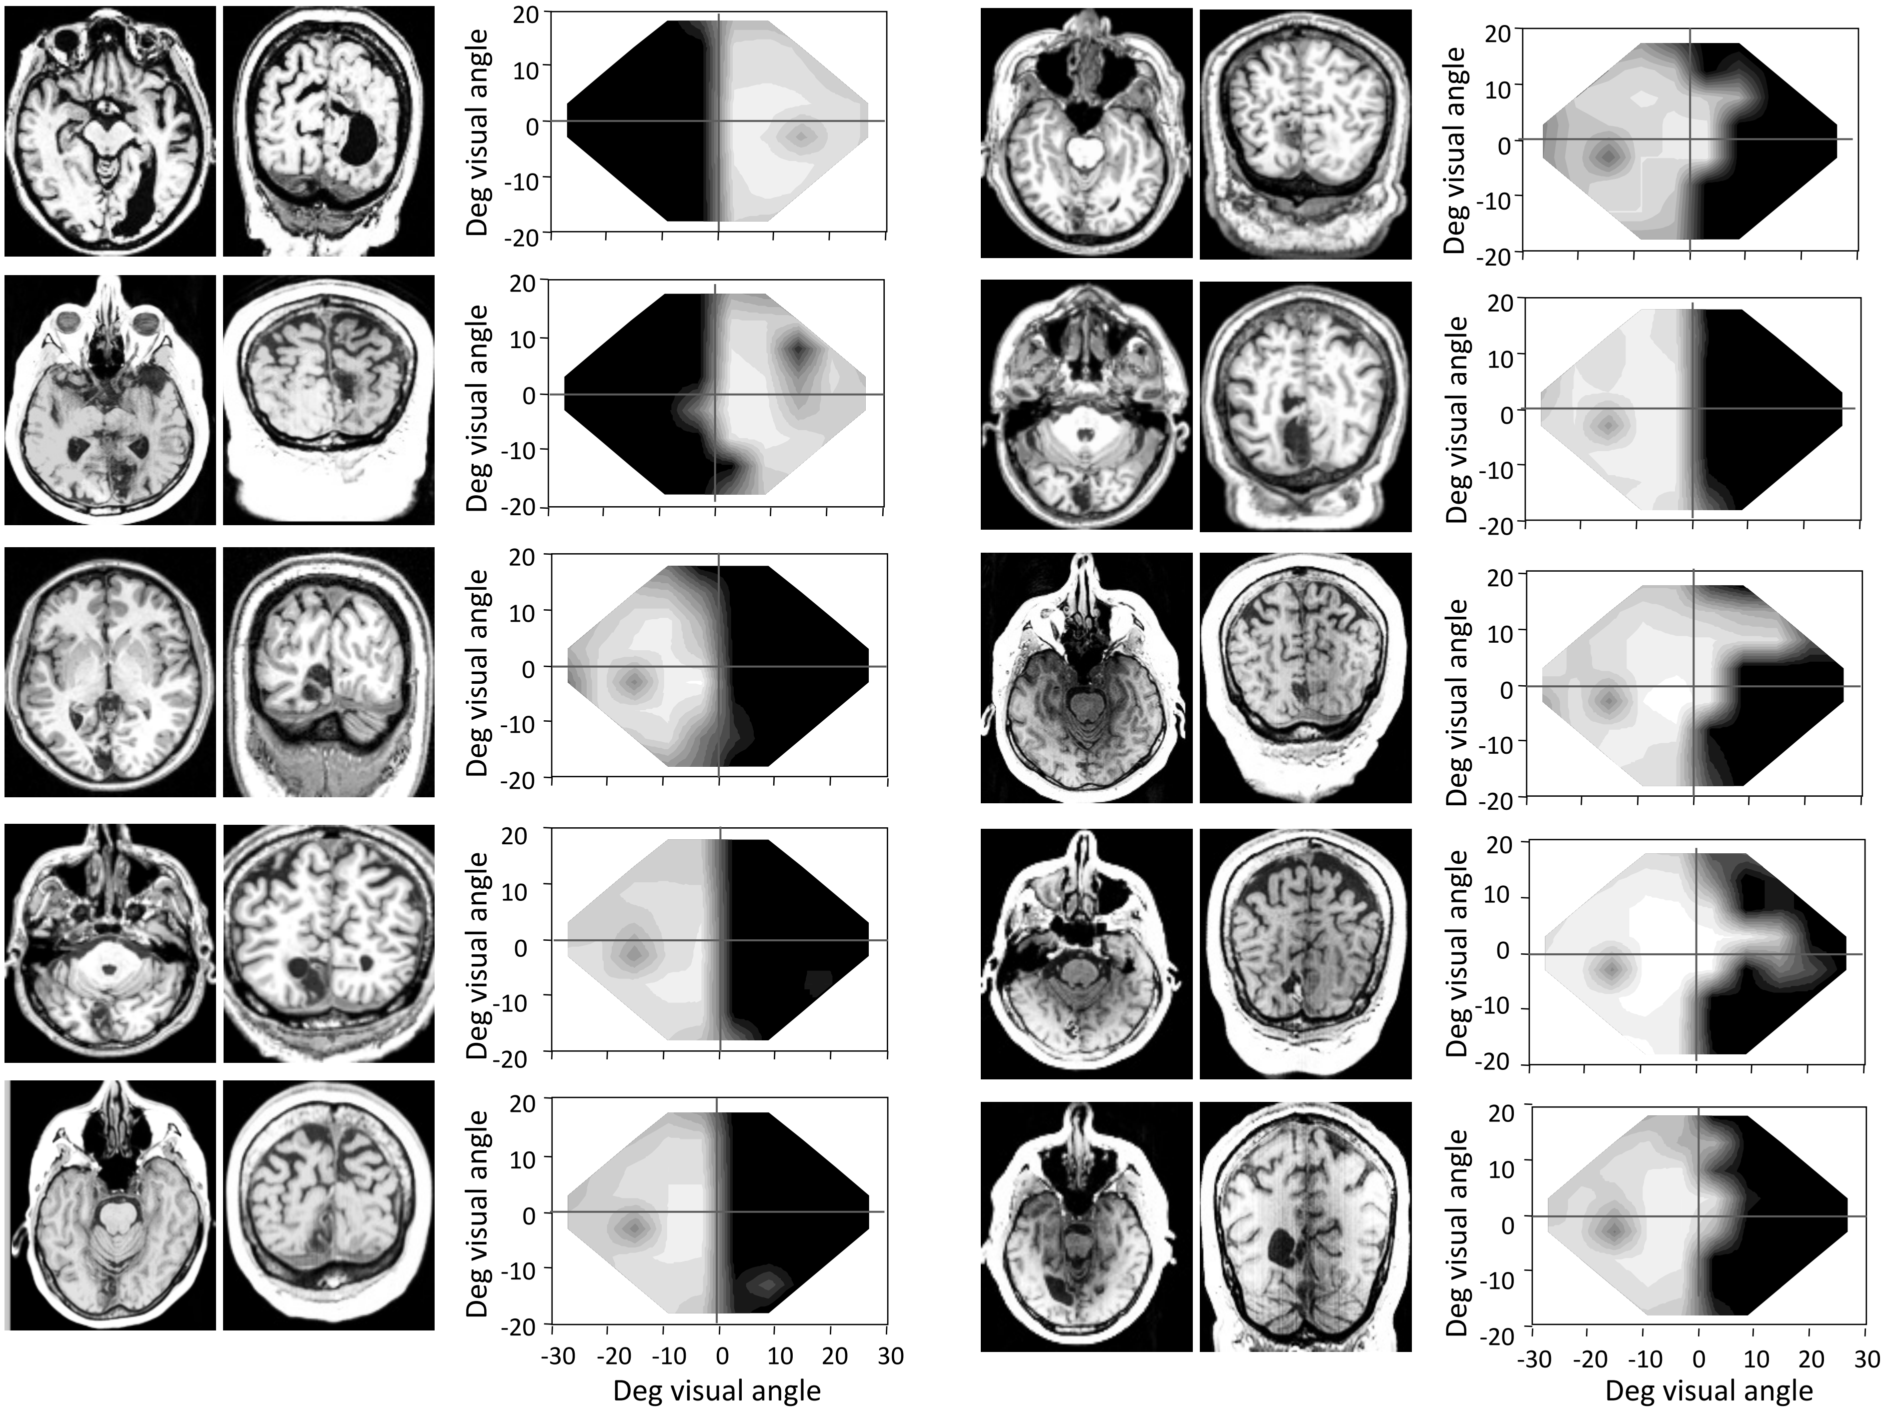
Deg visual angle
20
10
0
-10
-20
-30
-20
-10
0
10
20
30
Deg visual angle

Optic flow

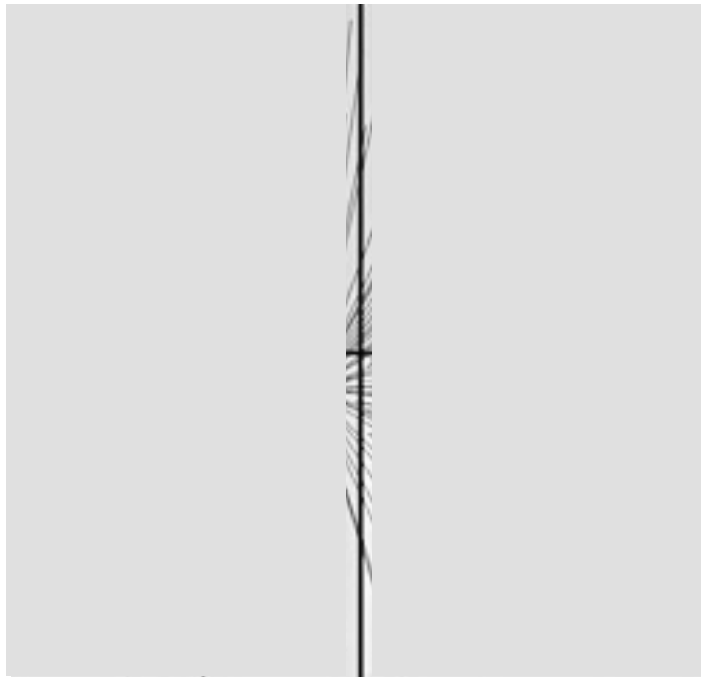

Warren & Kurtz, 1992
Crowell & Banks, 1996

Contributes to walking
Warren et al., Nat Neuro 2001

May not contribute to walking
Rushton et al., Curr Bio 1998
Harris and Bonas, Vis Res 2002

Contributes to walking depending on fidelity of info
Li & Niehorster, J Neurophys 2014
Li et al., JOV 2014

How impaired are hemianopes at estimating direction of heading (DOH) from optic flow?

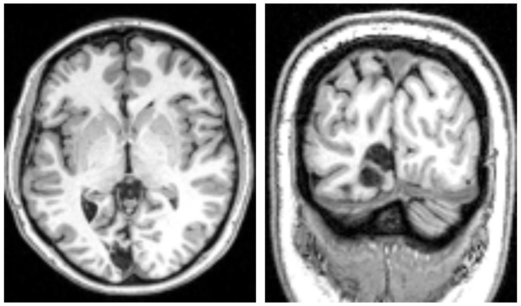

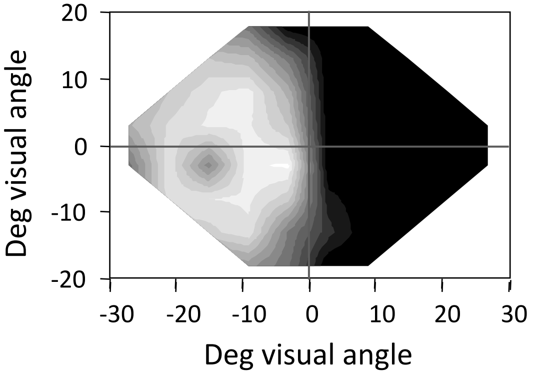

- How well do normal observers estimate DOH from optic flow when focus of expansion is obscured?
- How important are 4 quadrants/2 hemifields of vision in DOH estimation?

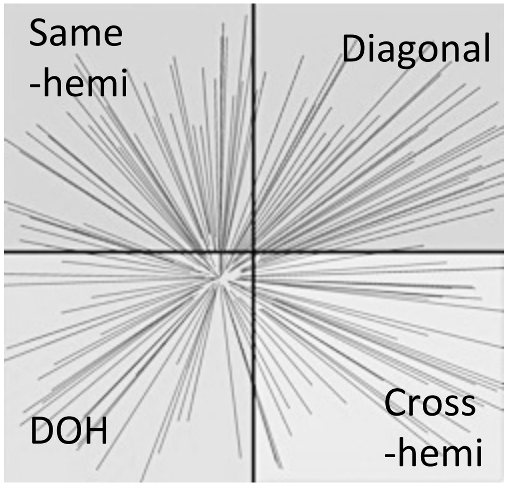

Experimental paradigm for direction of heading estimation task

Issen, Huxlin & Knill, JOV 2015

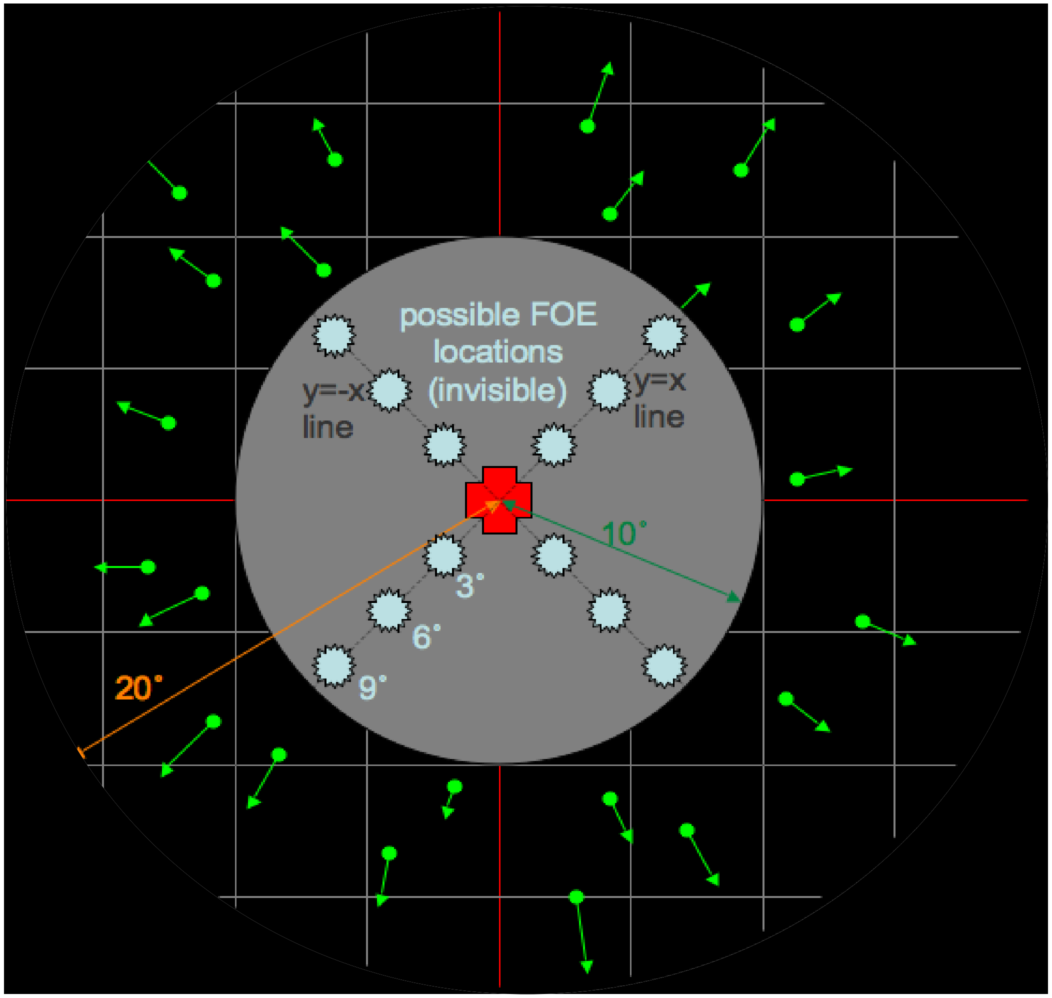

Experimental paradigm for direction of heading task

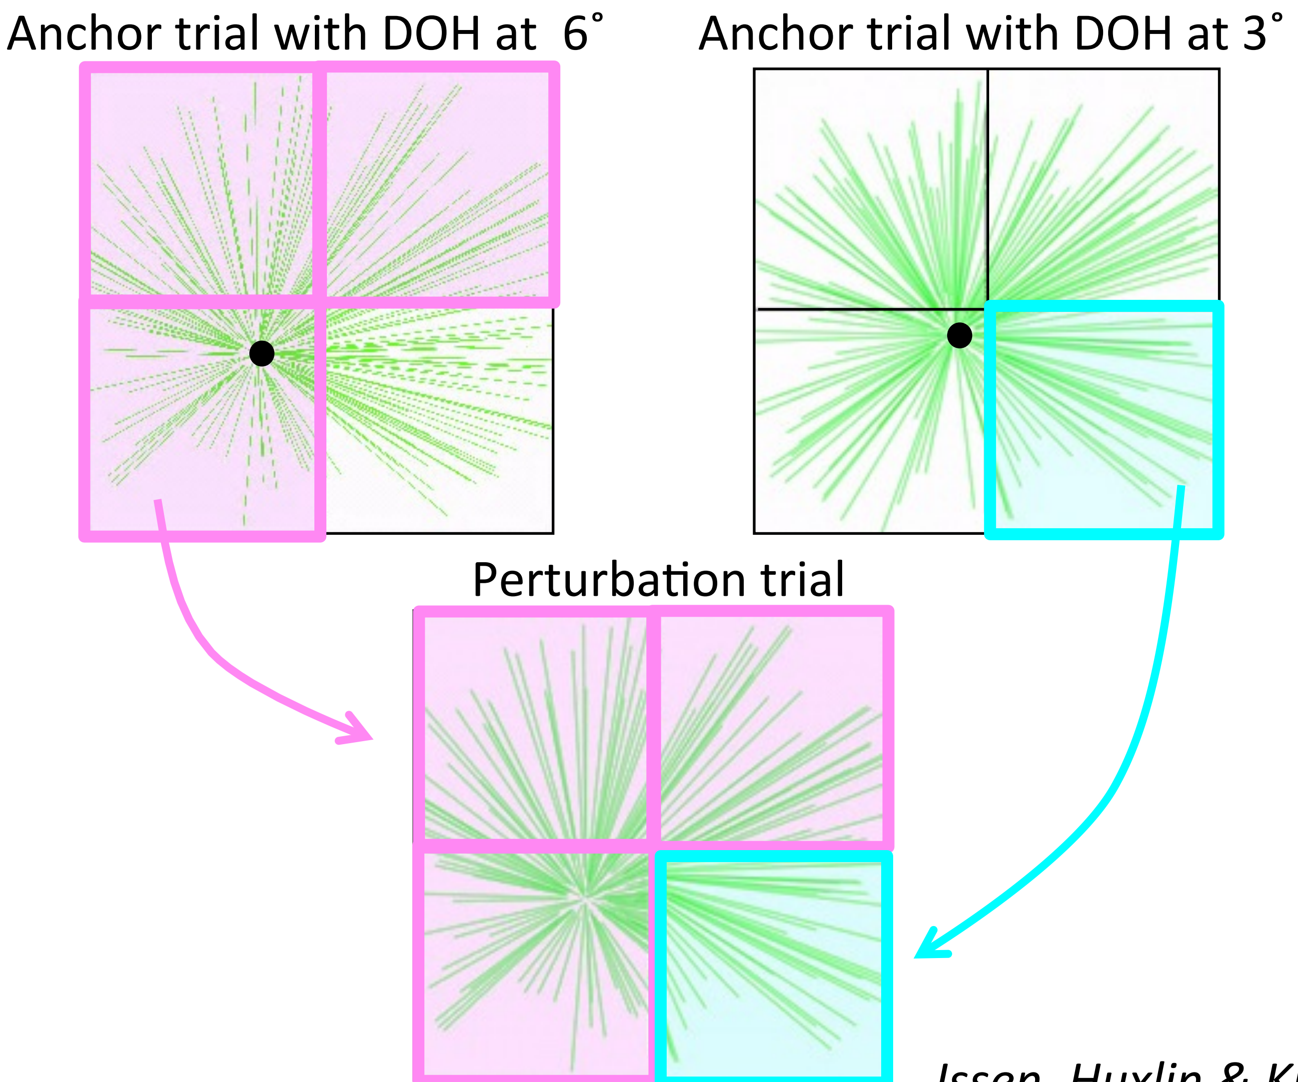

Issen, Huxlin & Knill, JOV 2015

Intact humans are *almost* ideal observers

They give weight to different visual field quadrants according to relevance of information content for DOH task

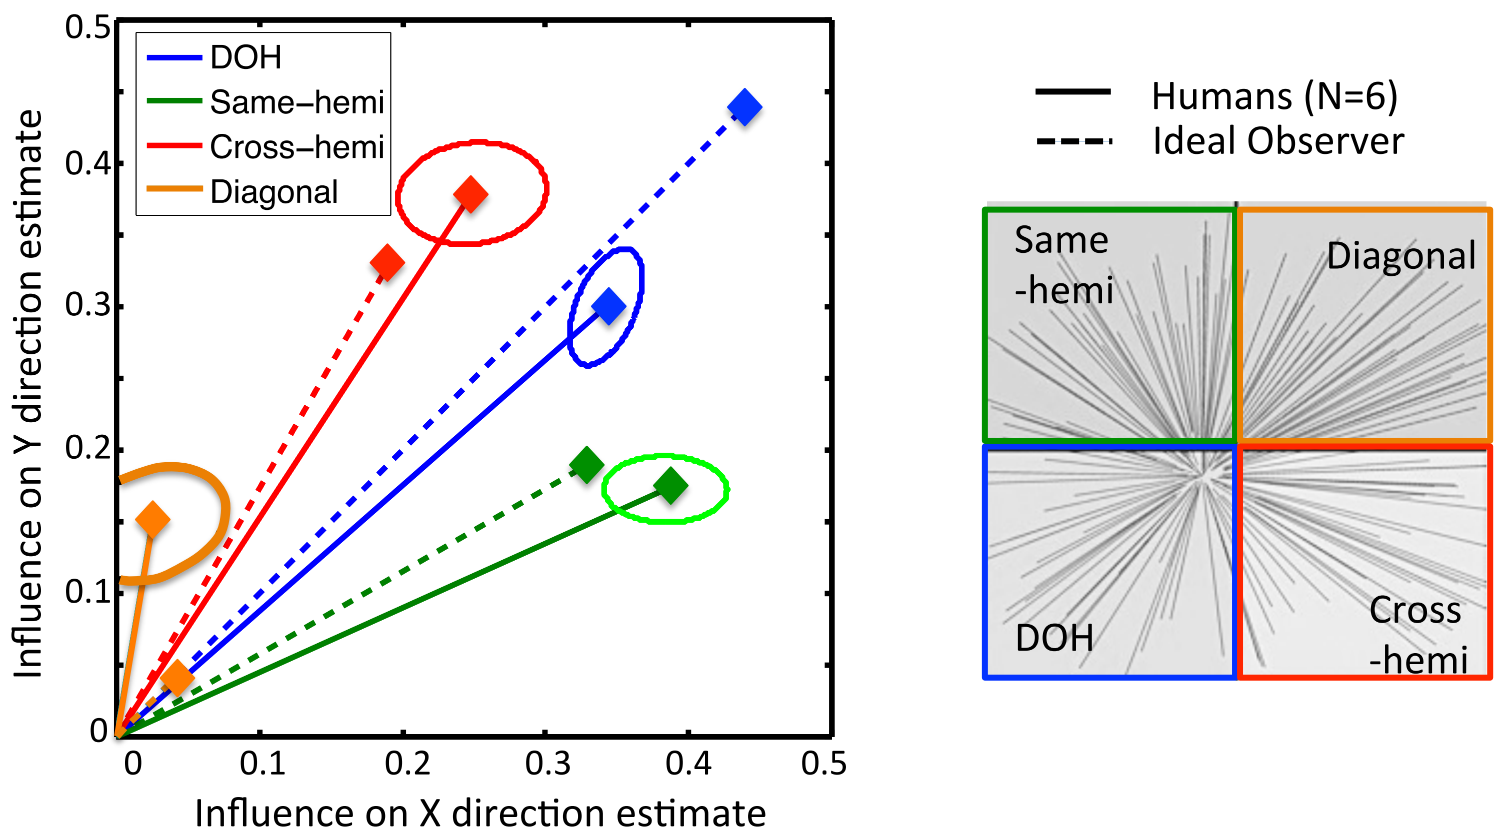

Issen, Huxlin & Knill, JOV 2015

How impaired are hemianopes at estimating DOH from optic flow?

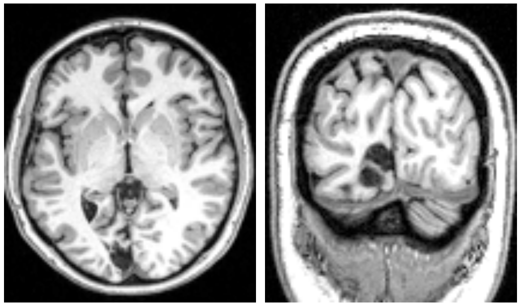

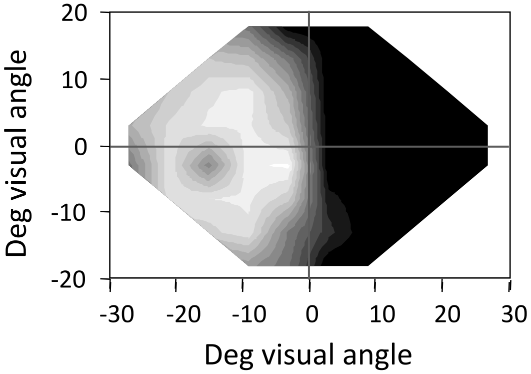

Characterize and model behavior

- *Can the effects of hemianopia be modeled by simulating field loss?*
- *Are heading estimates in the intact visual field affected by hemianopia?*
- *We know hemianopes can "sense" some motion in their blind field – is it used for DOH tasks?*

Estimating impact of hemifield loss

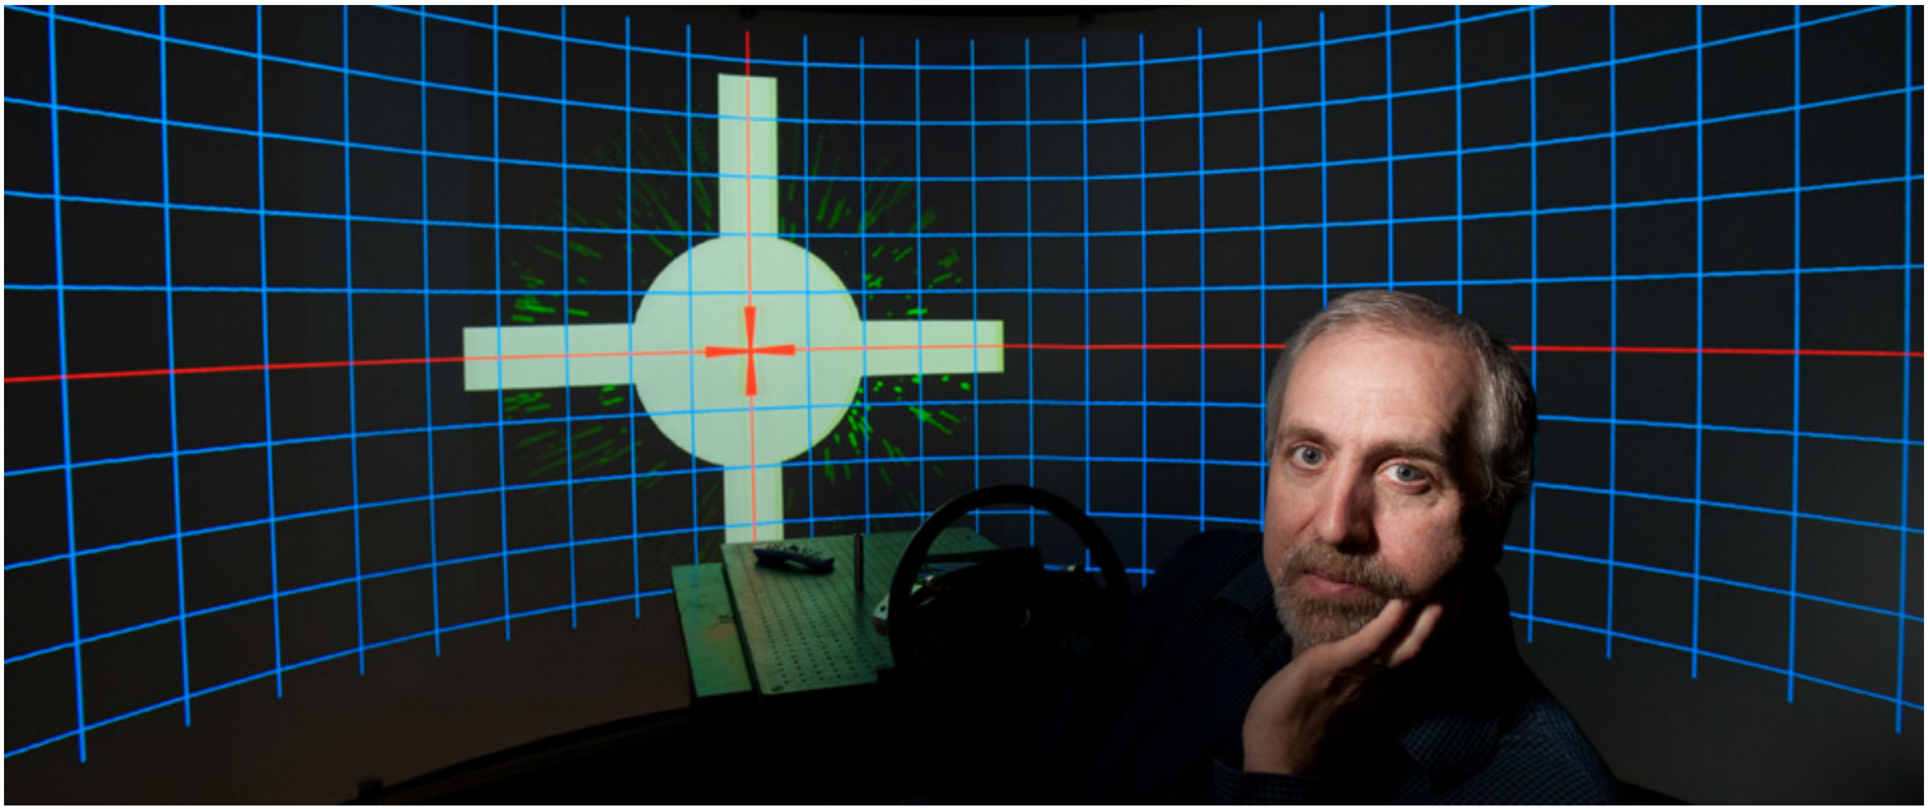

Estimating impact of hemifield loss

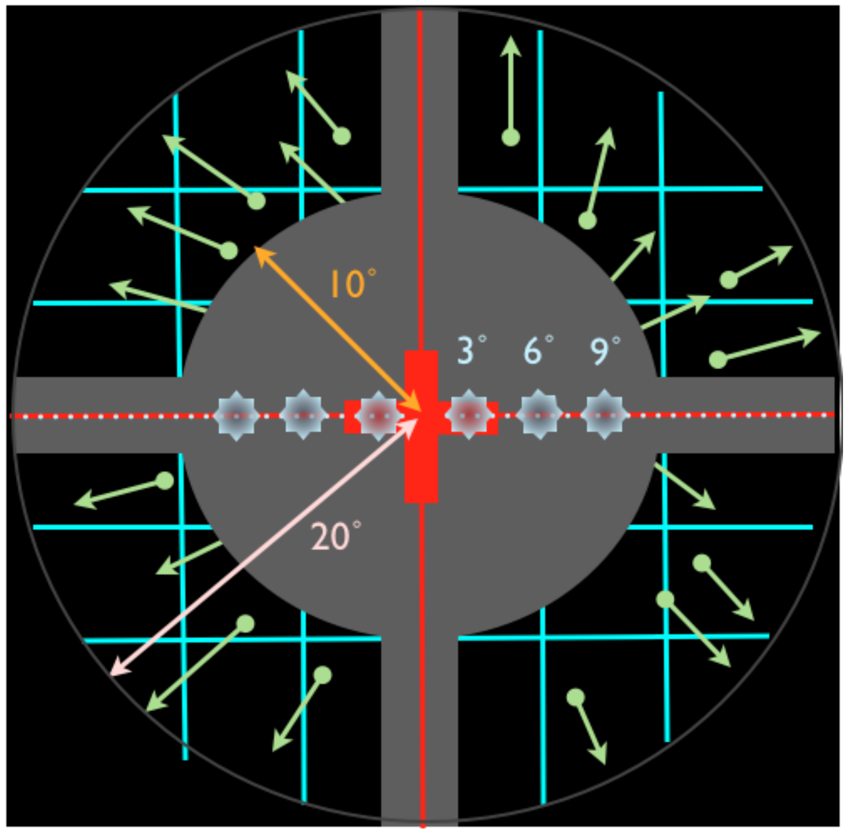

Subjects: visually-intact controls (8 young, 8 older) and 7 hemianopes (older)

800 trials: perturbation (in left or right hemifields), anchor and feedback trials

Conditions: full field or simulated hemianopia (for older controls)

Older adults' DOH estimates are more compressed towards fixation

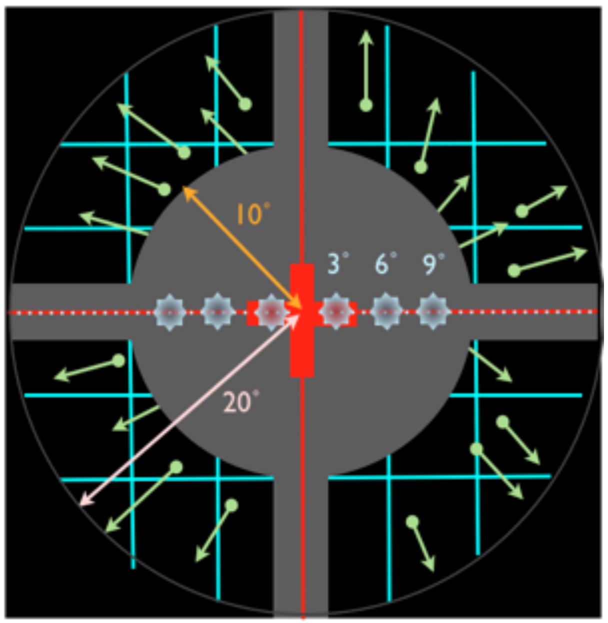

8 young adults: 18-21 yrs, mean 19 yrs

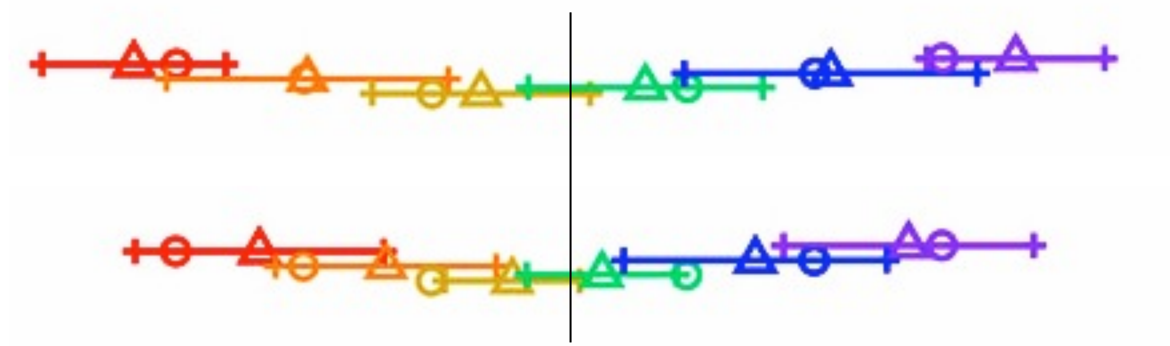

8 older adults: 54-75 yrs, mean 68 yrs

○ Real target position

△ Perceived target position

Analyzing DOH estimates

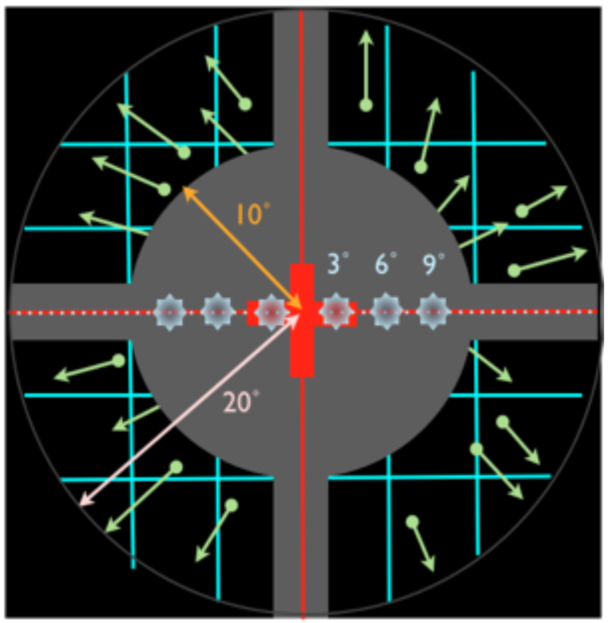

$$x_{estimate\mathtt{left}} = \alpha_{\mathtt{left}}[w^*_{x\mathtt{left}} x_{\mathtt{left}} + w^*_{x\mathtt{right}} x_{\mathtt{right}}] + \beta_{\mathtt{left}} + noise$$

$$x_{estimate\mathtt{right}} = \alpha_{\mathtt{right}}[w^*_{x\mathtt{left}} x_{\mathtt{left}} + w^*_{x\mathtt{right}} x_{\mathtt{right}}] + \beta_{\mathtt{right}} + noise$$

x	X-coordinate of response
α	Compression/expansion factor
ω^*	Weight given to hemifield, normalized
β	Additive response bias

All adults give more weight to hemifield containing DOH

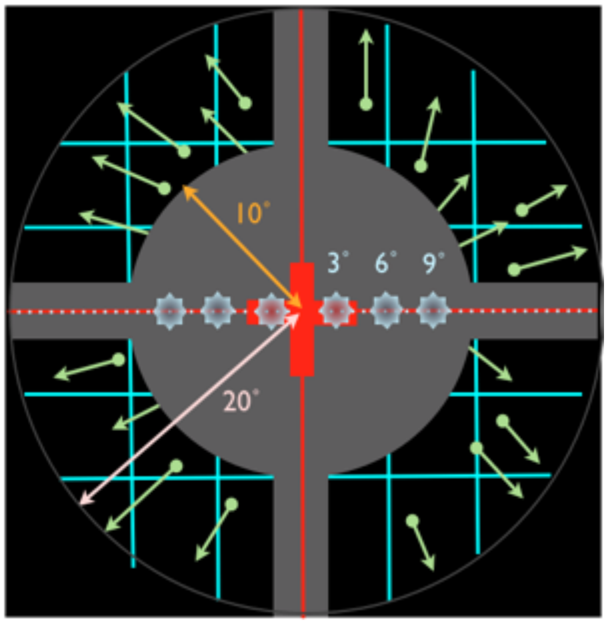

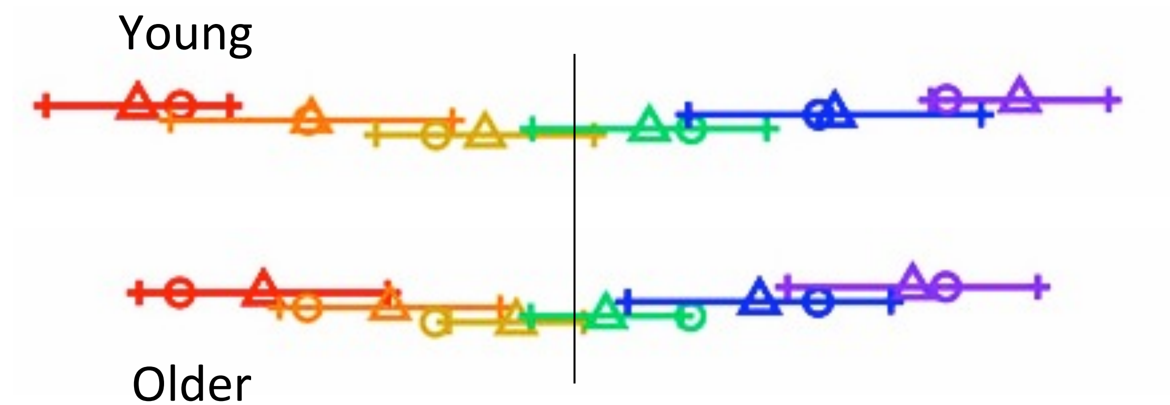

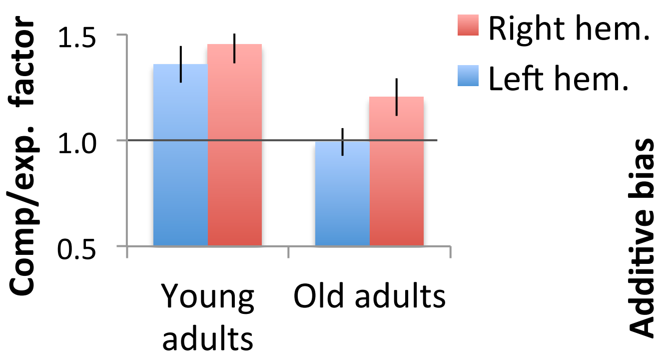

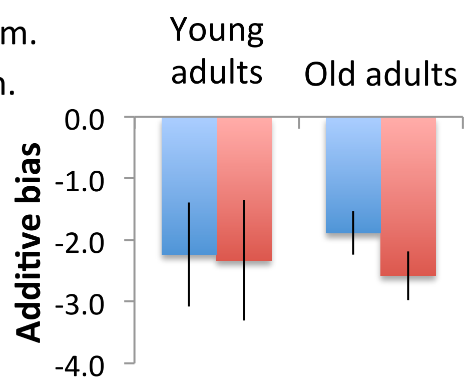

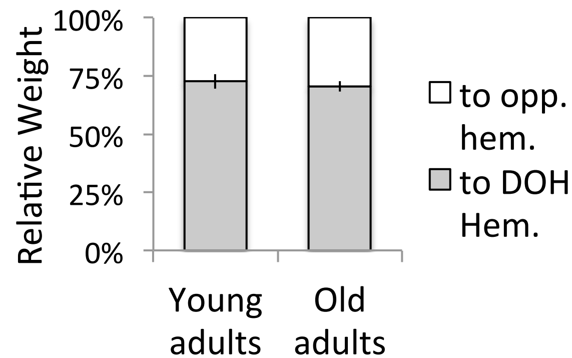

Simulating hemianopia changes bias and compression

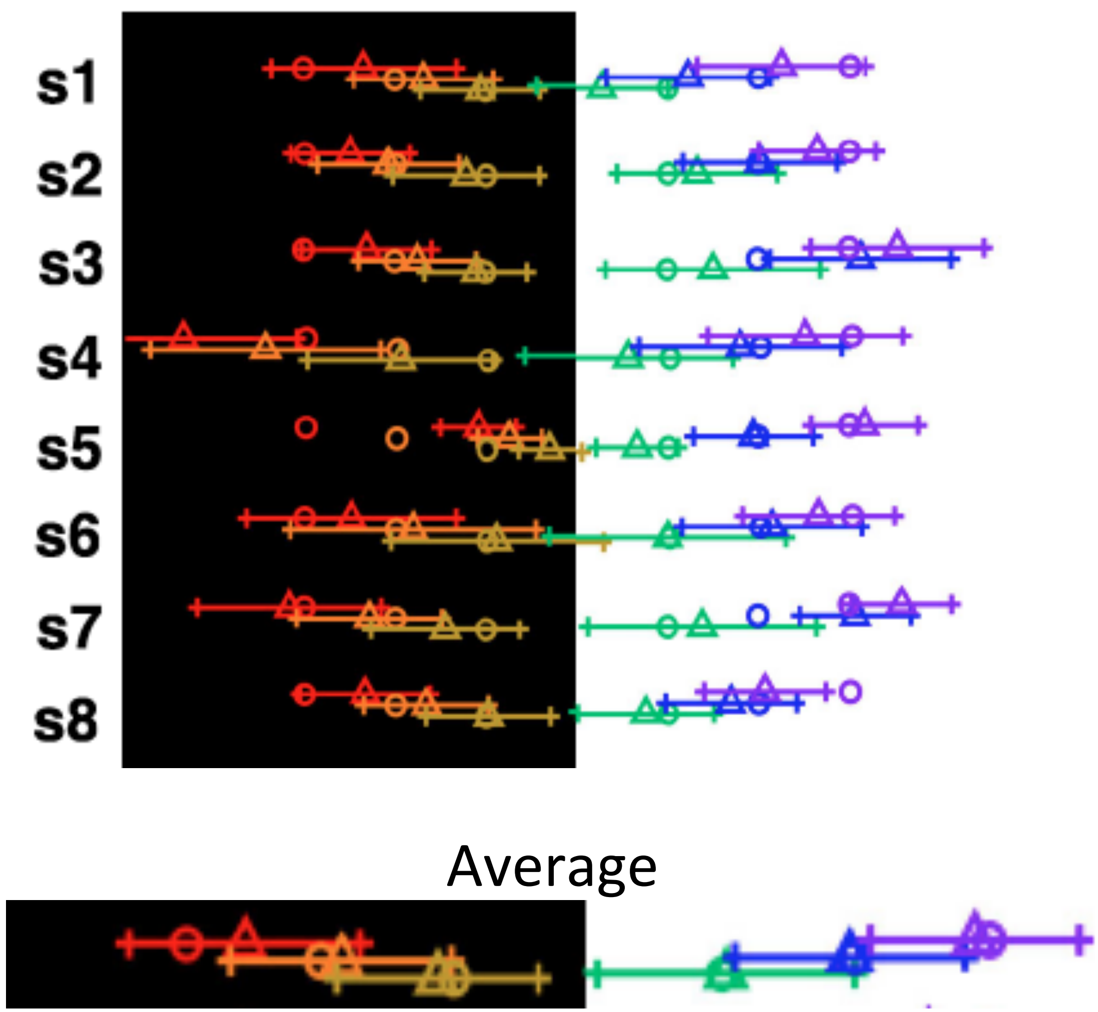

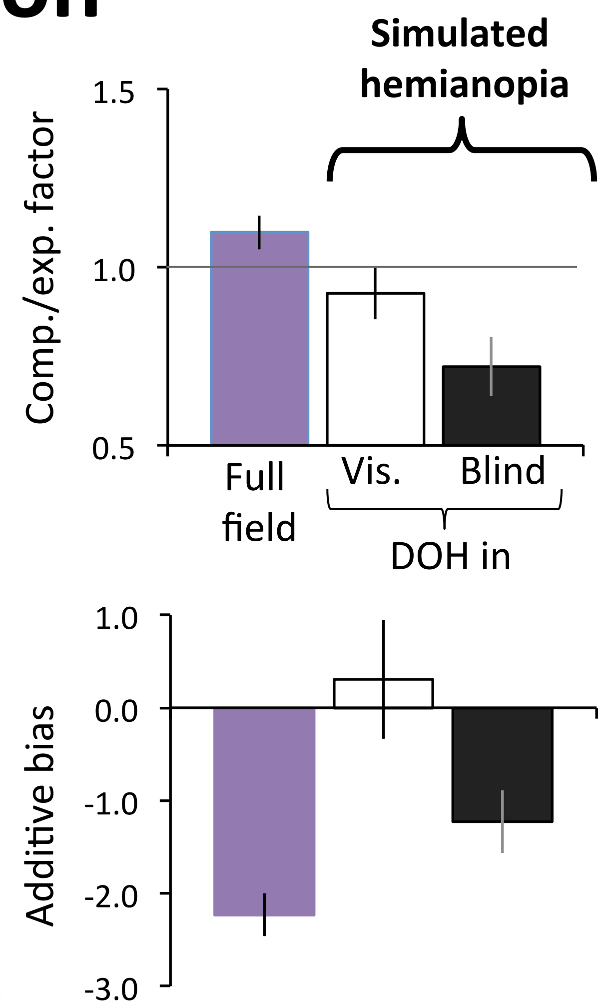

Real hemianopia alters compression and bias in BOTH hemifields

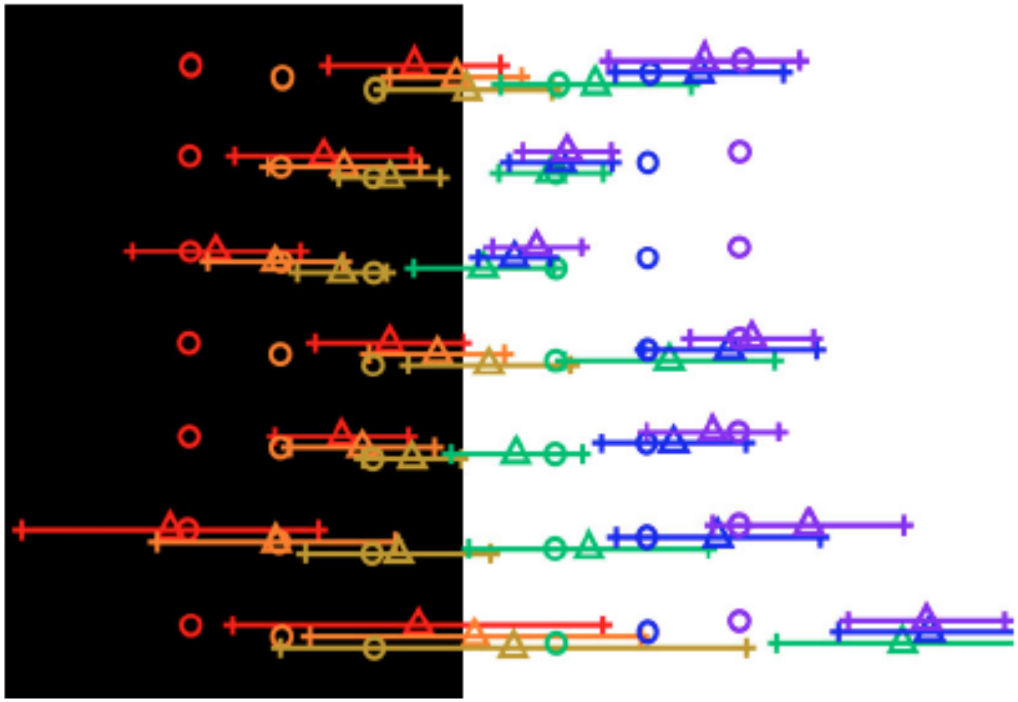

Average

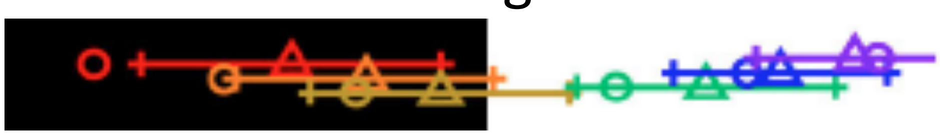

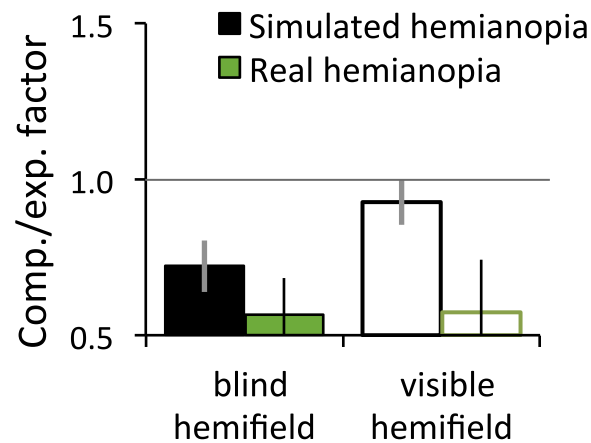

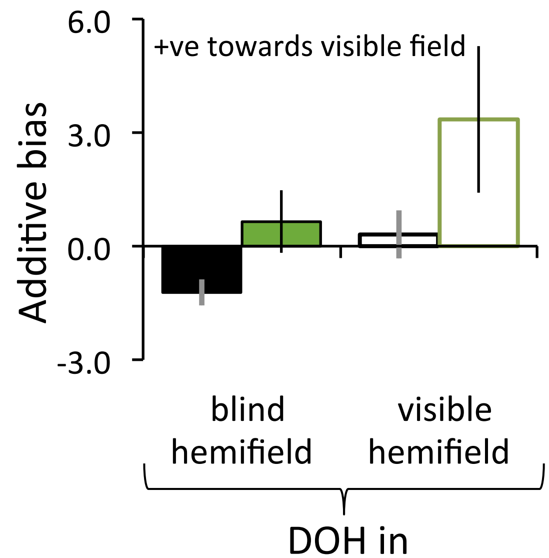

Some hemianopes give weight to blind field information when it contains DOH

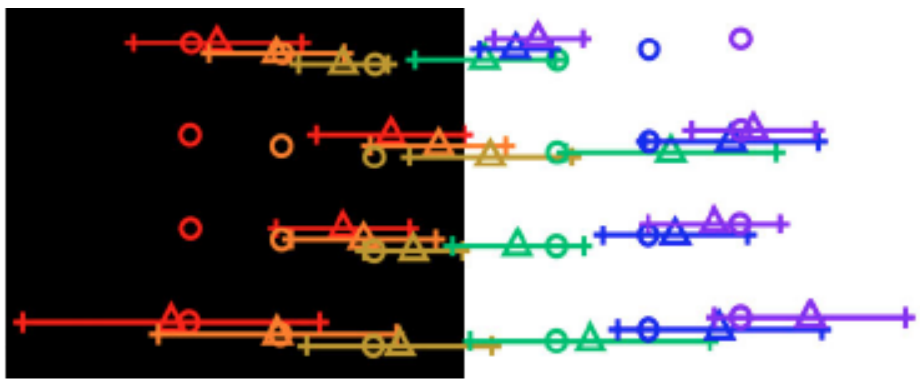

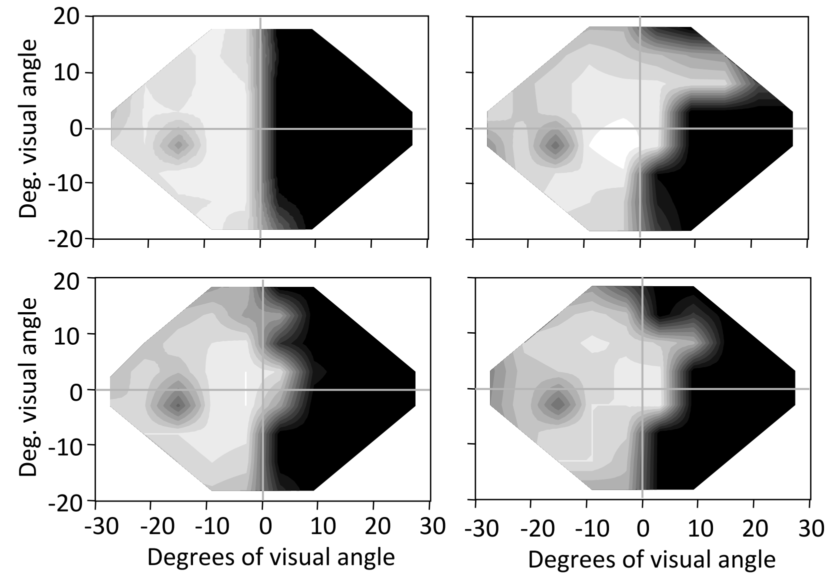

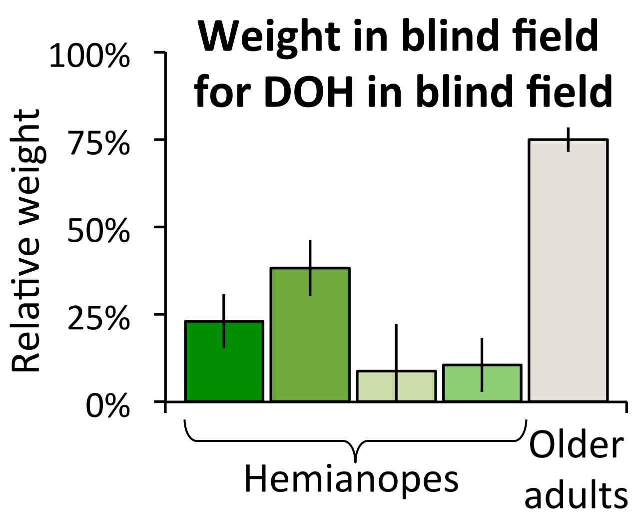

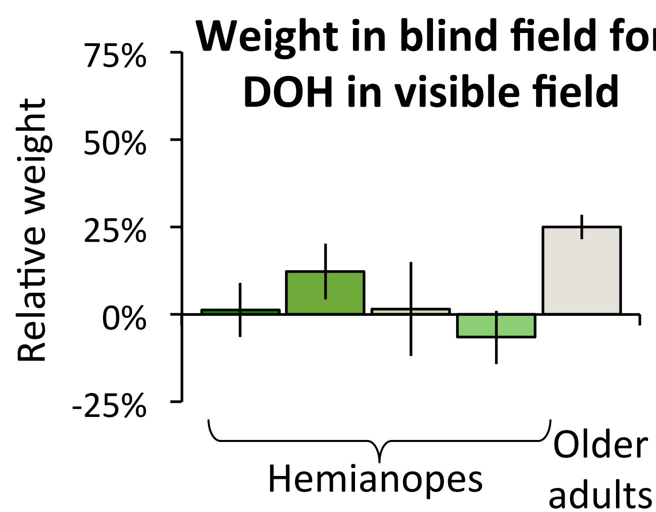

Summary

Hemianopes are impaired at estimating direction of heading from optic flow

- **Simulated** hemianopia underestimates real behavior
- Deficit affects **intact** hemifield performance
- Weight is given to **blind** hemifield information

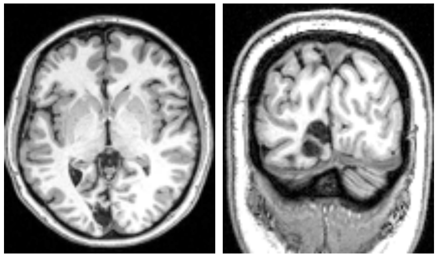

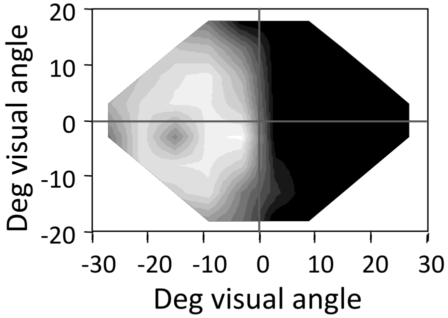

Implications?

Real hemianopia is more exaggerated than simulated deficit

Adaptation over time since stroke, additional factors?

DOH judgments impaired across whole hemifield

May explain persistent problems navigating

Some hemianopes give significant weight to blind field information - automatically

Improving motion processing in the blind field could help DOH estimation

Acknowledgments

Collaborators

David Knill, Ph.D.

Duje Tadin, PhD

David Heeger, PhD

Marisa Carrasco, PhD

Tatiana Pasternak, PhD

Mary Hayhoe, PhD

Post-docs

Antoine Barbot, PhD

Elisha Merriam, PhD

Tim Martin, PhD

Students

Laurel Issen, PhD

Matt Cavanaugh

Anasuya Das, PhD

Michael Melnick

Aaron Levi

Danielle Shaked

Noah Elkins

Our study participants

Center for Visual Science

Rochester Center for Brain Imaging

Technical Staff

Keith Parkins

Adin Reisner

Pat Weber

Margaret DeMagistris

Sponsors

Schmitt Foundation, RPB, NEI

UR MEDICINE | FLAUM EYE INSTITUTE

Thank you!

Dave's work on mixture priors and causal inference

Wei Ji Ma
New York University

Mixture models and the probabilistic structure of depth cues

David C. Knill

Center for Visual Sciences, University of Rochester, 274 Meliora Hall, Rochester, NY 14627, USA

Received 1 March 2002; received in revised form 19 September 2002

Vision Research, 2003

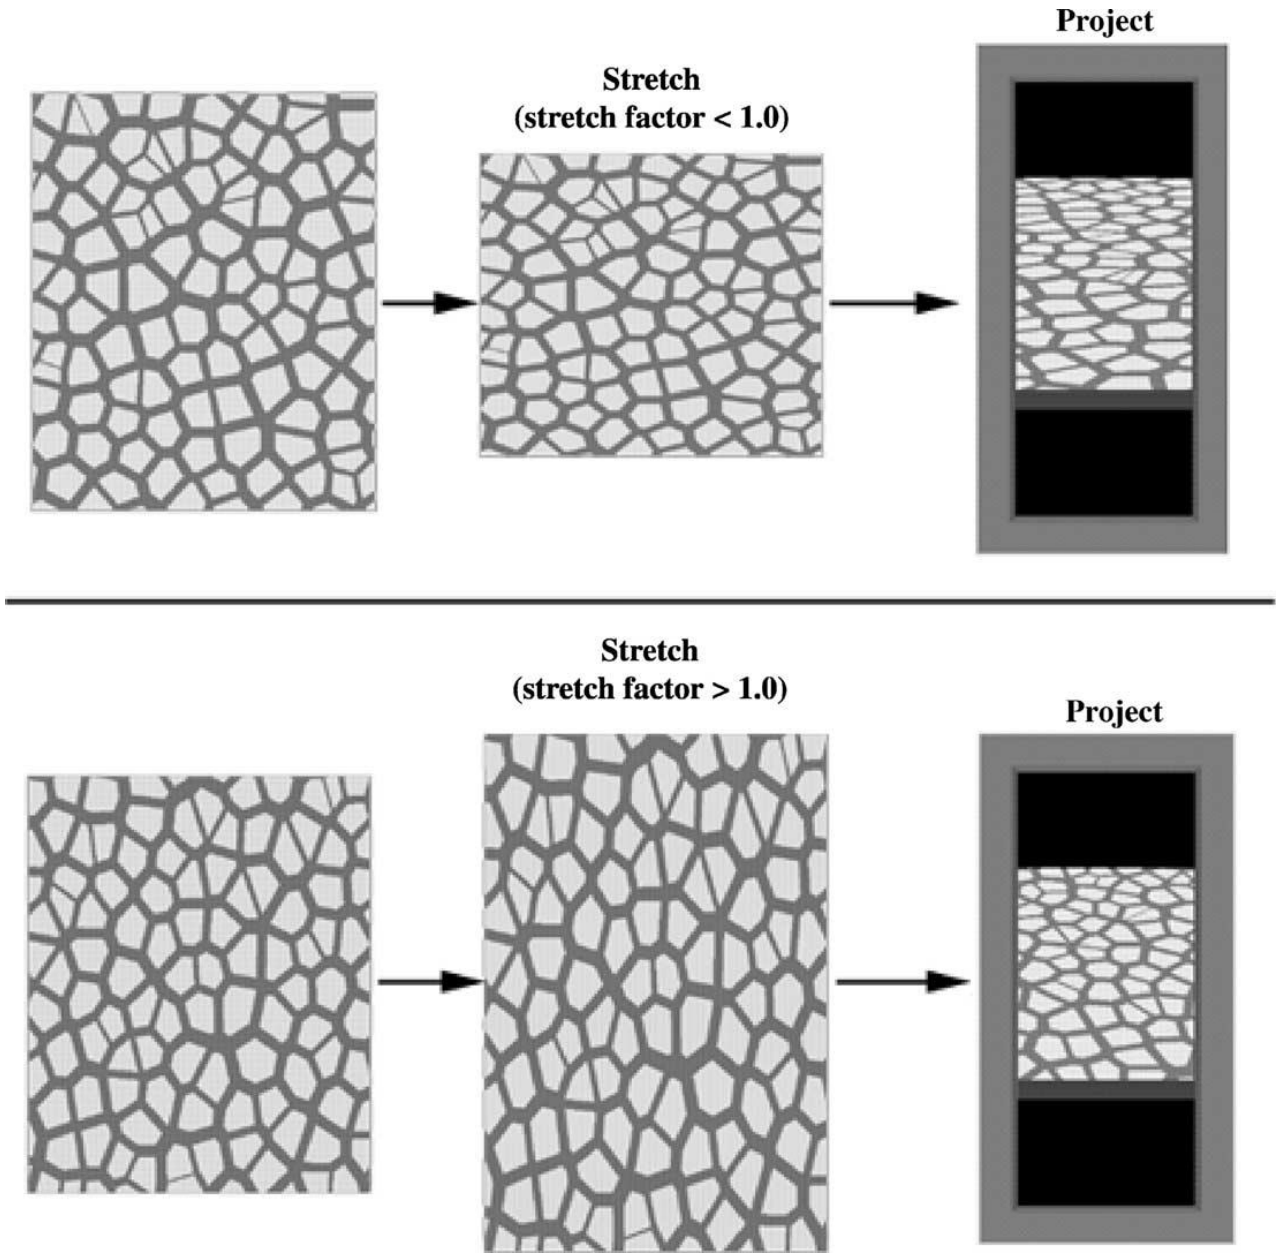

Fig. 10. Stimuli for the experiments were created in three stages. First, a random, isotropic texture pattern was generated. This was then stretched by some amount in the vertical direction. The resulting texture was projected into the image at a slant of 65° and a vertical tilt. A subject that assumes surface textures are isotropic would overestimate the slant of the top stimulus and underestimate the slant of the bottom one.

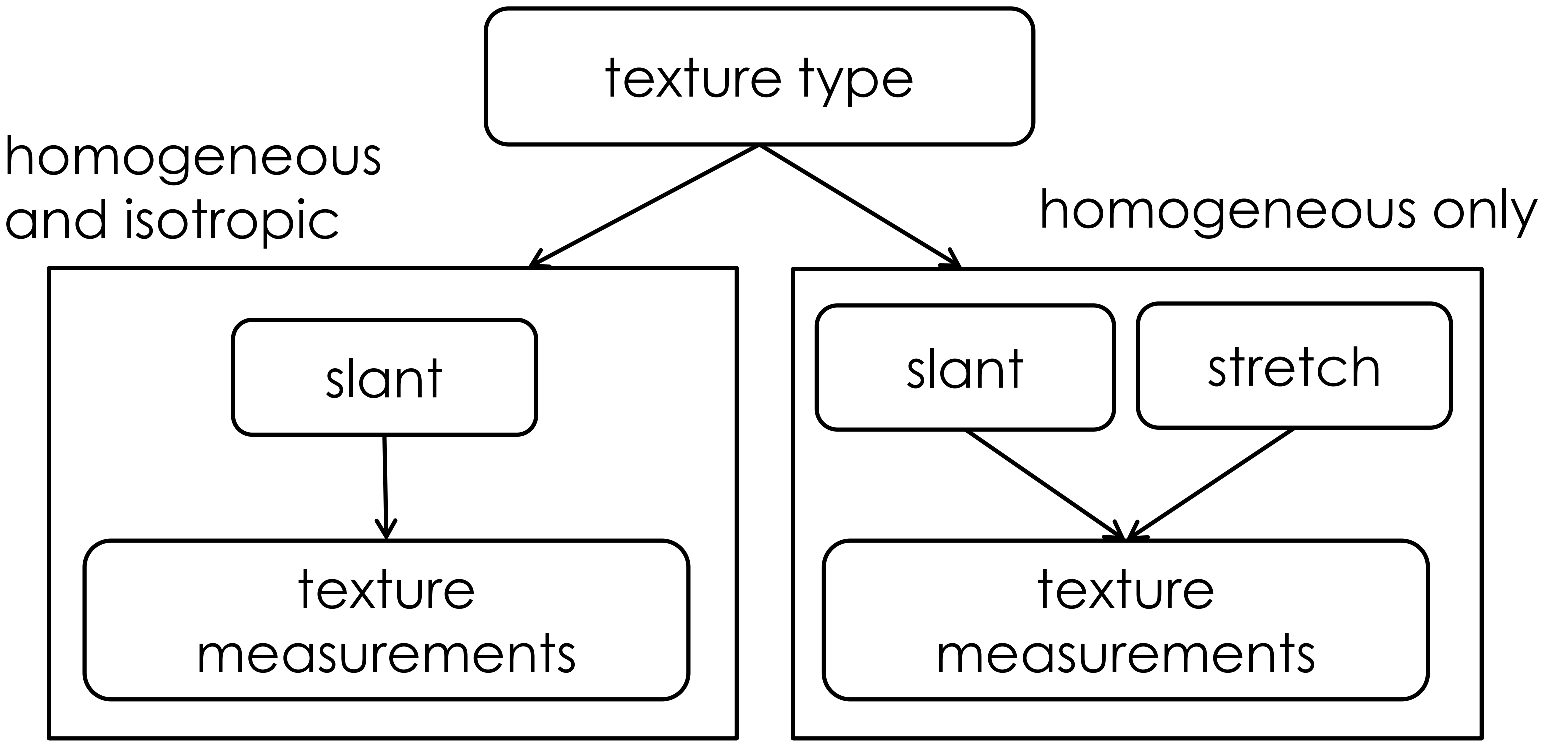

Two (or more) categorically different scenarios that could have given rise to the sensory observations.

The role of memory in visually guided reaching

Anne-Marie Brouwer
Center for Visual Science, University of Rochester, Rochester, NY, USA

David C. Knill
Center for Visual Science, University of Rochester, Rochester, NY, USA

Journal of Vision, 2007

Anne-Marie Brouwer

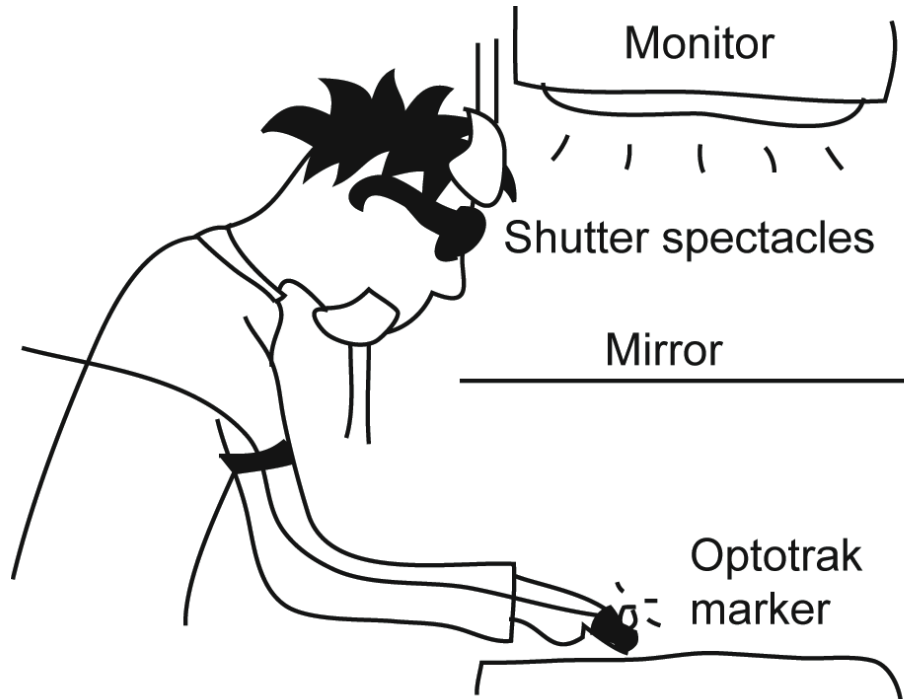
Monitor
Shutter spectacles
Mirror
Optotrak
marker

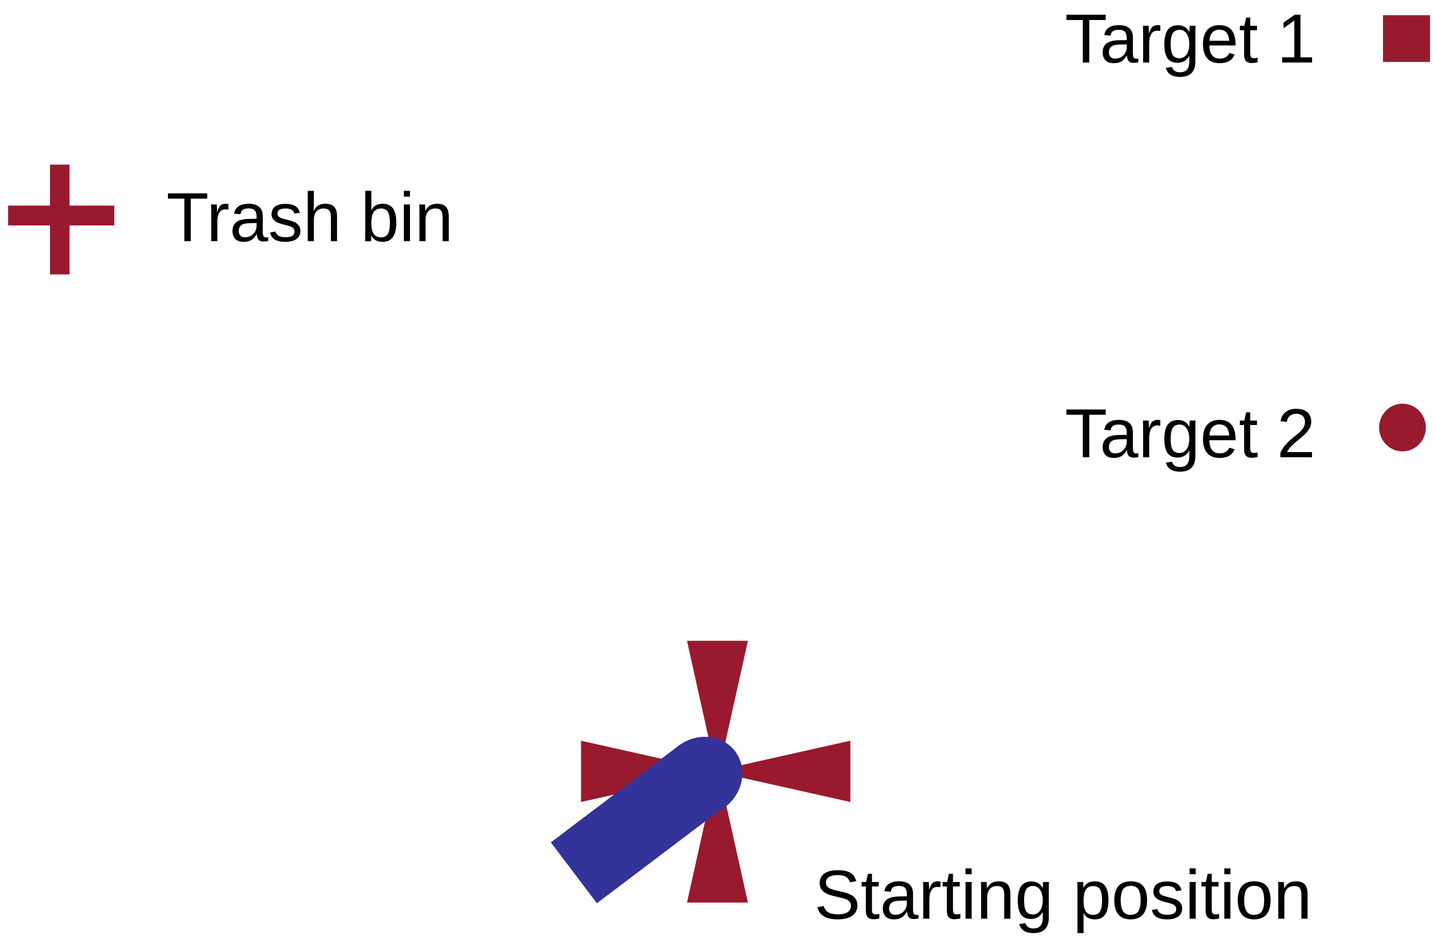
Target 1
Trash bin
Target 2
Starting position

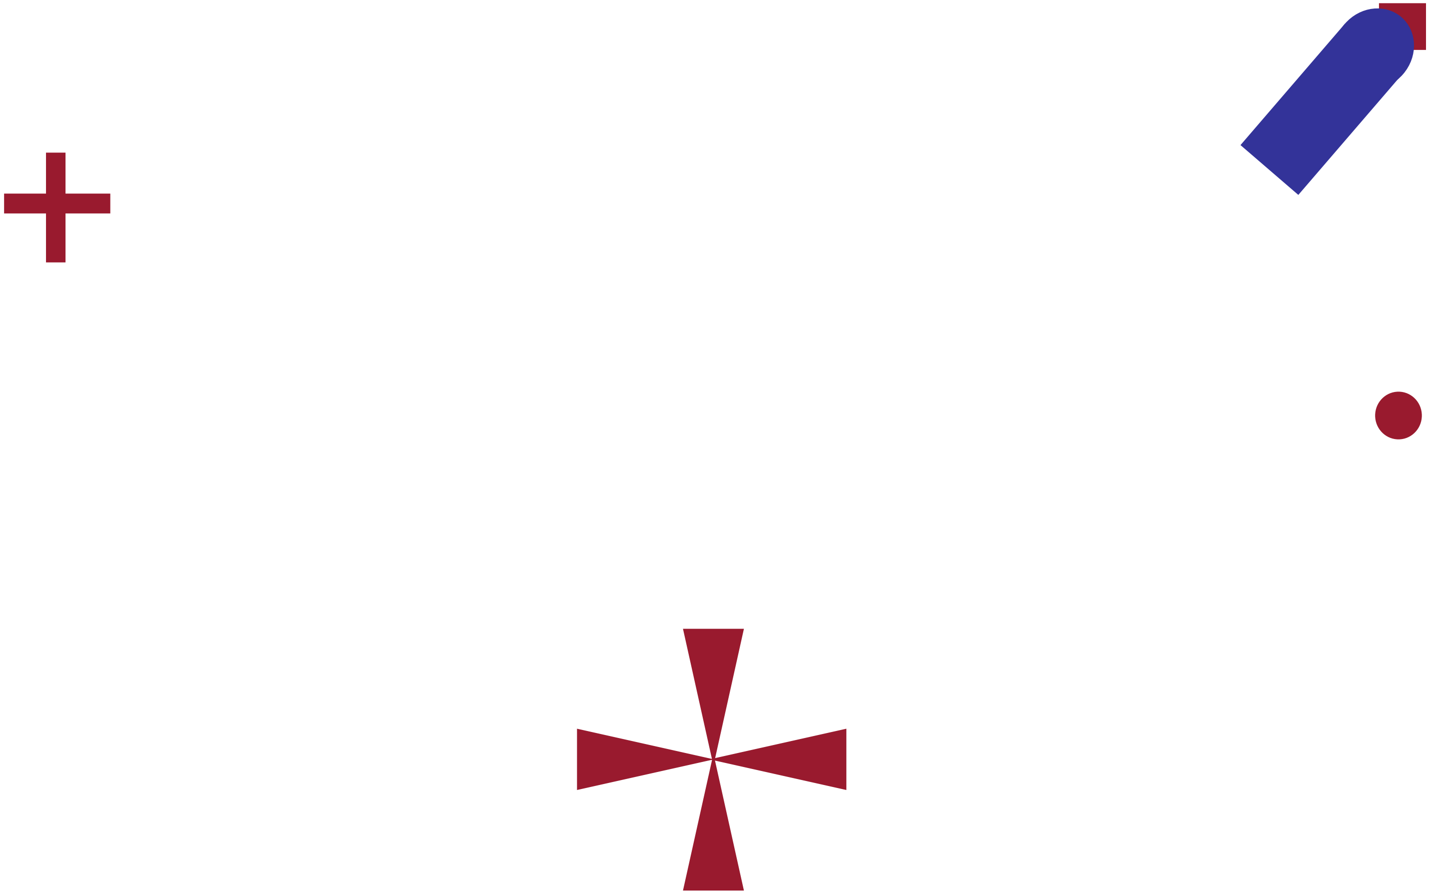

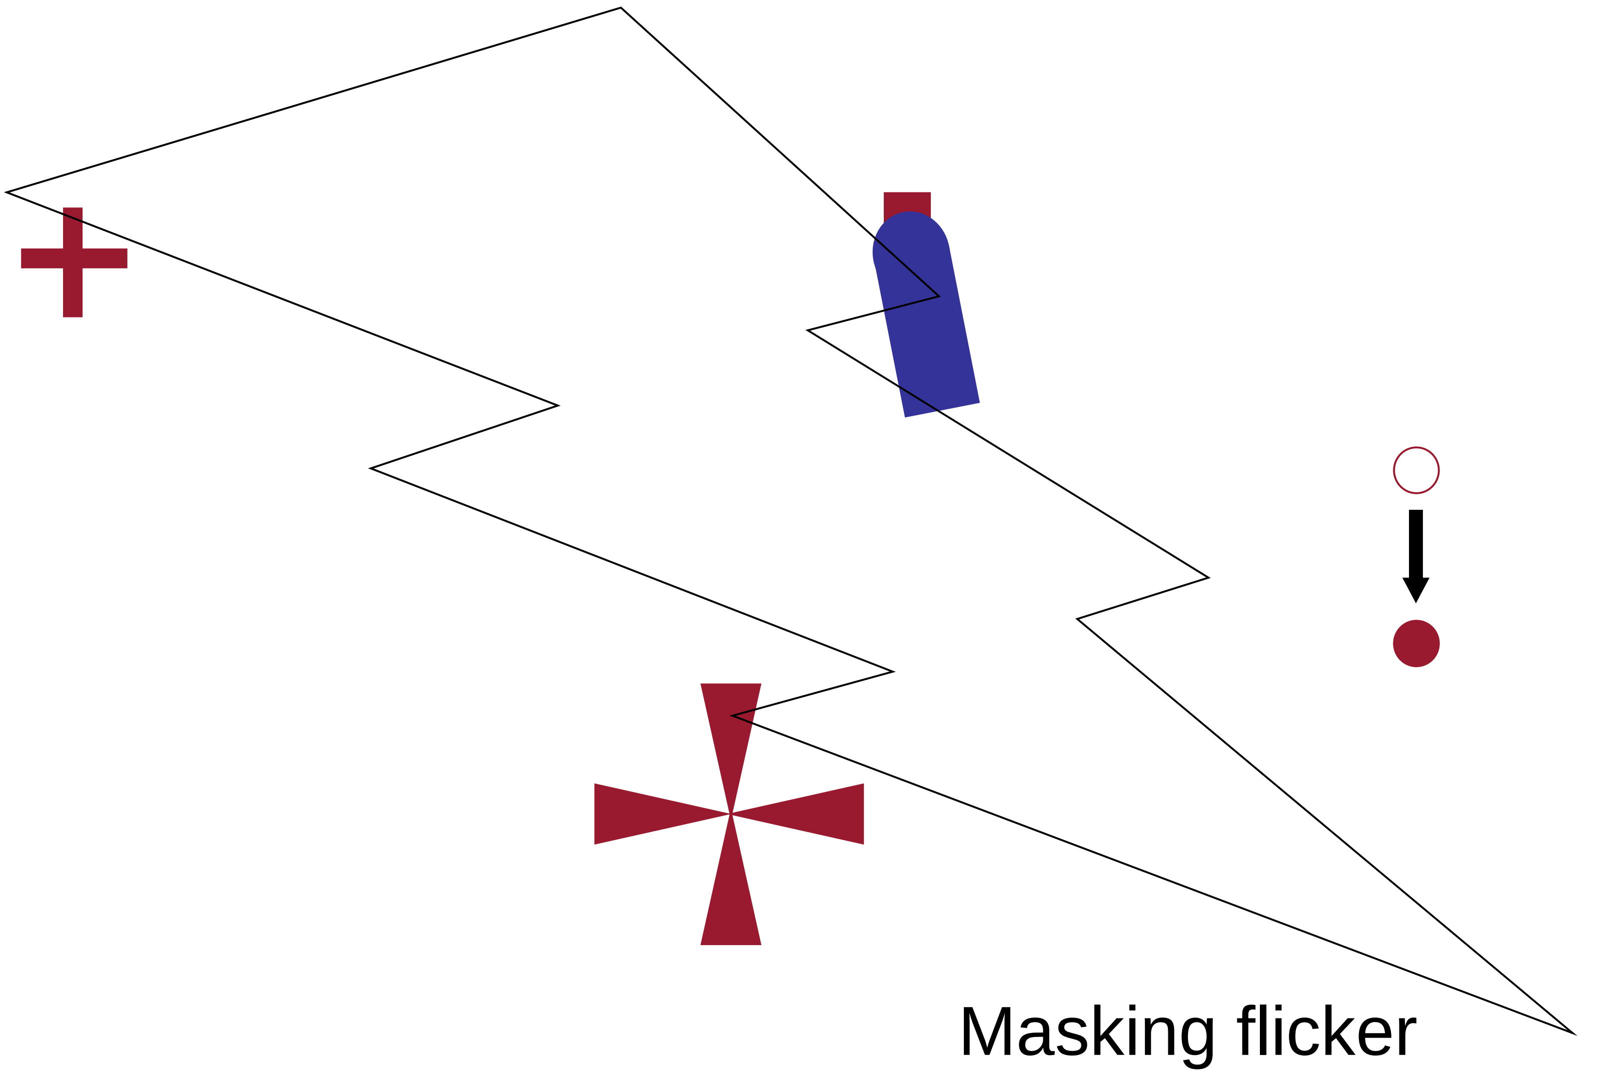
Masking flicker

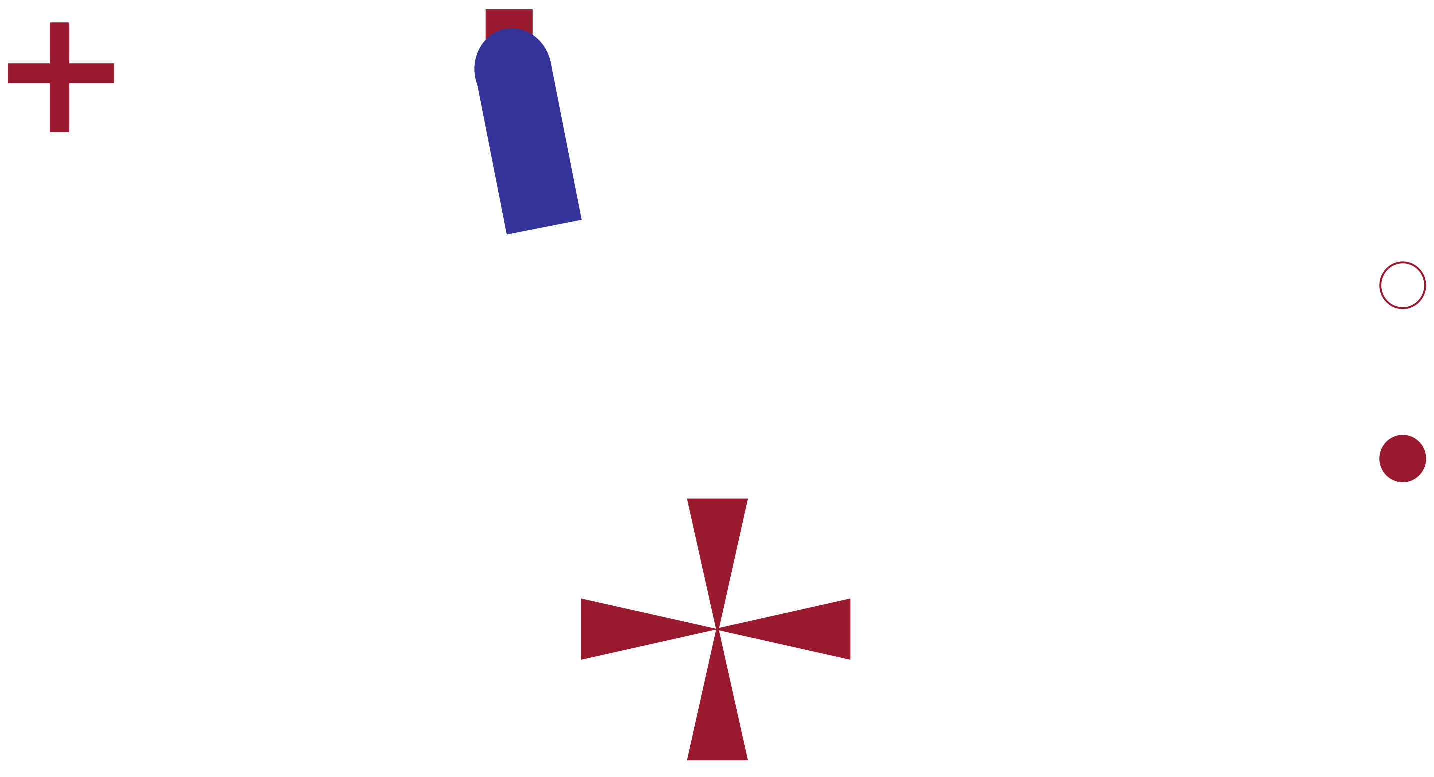

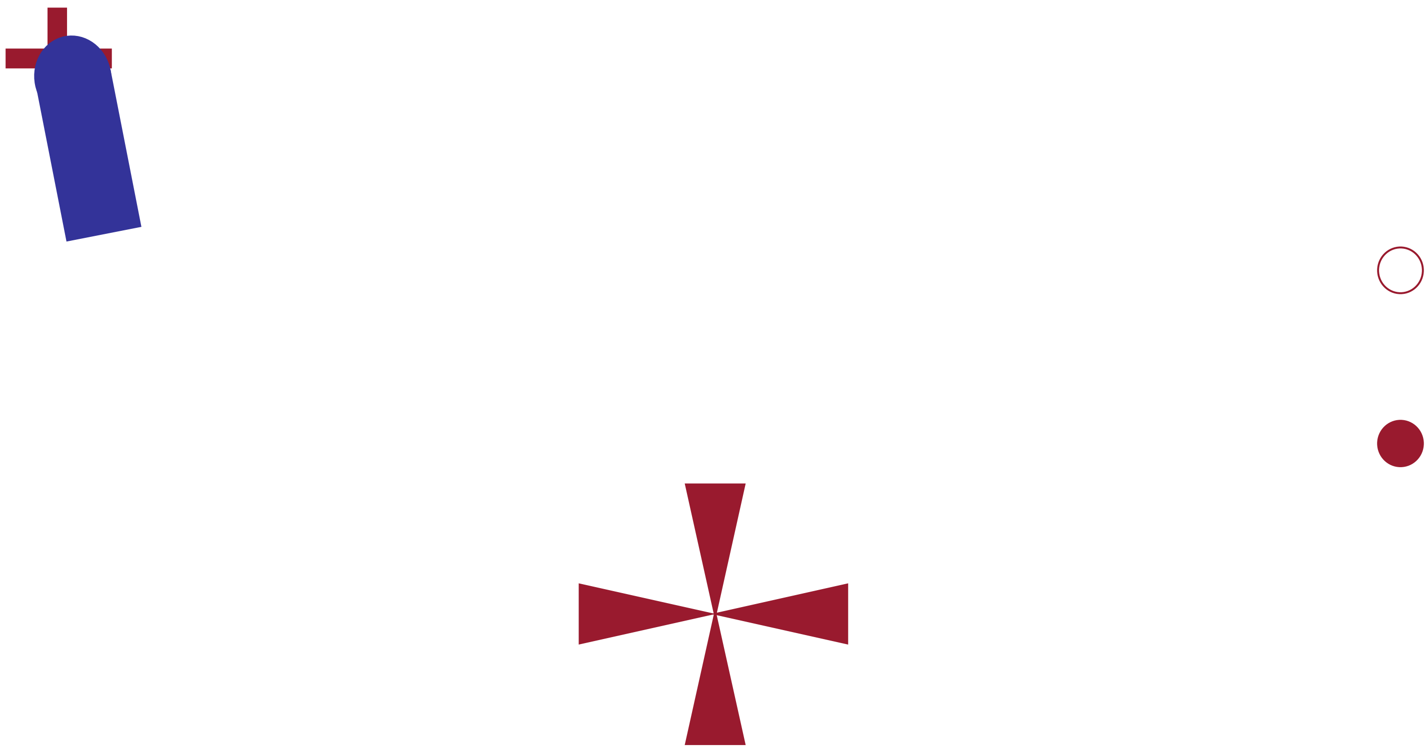

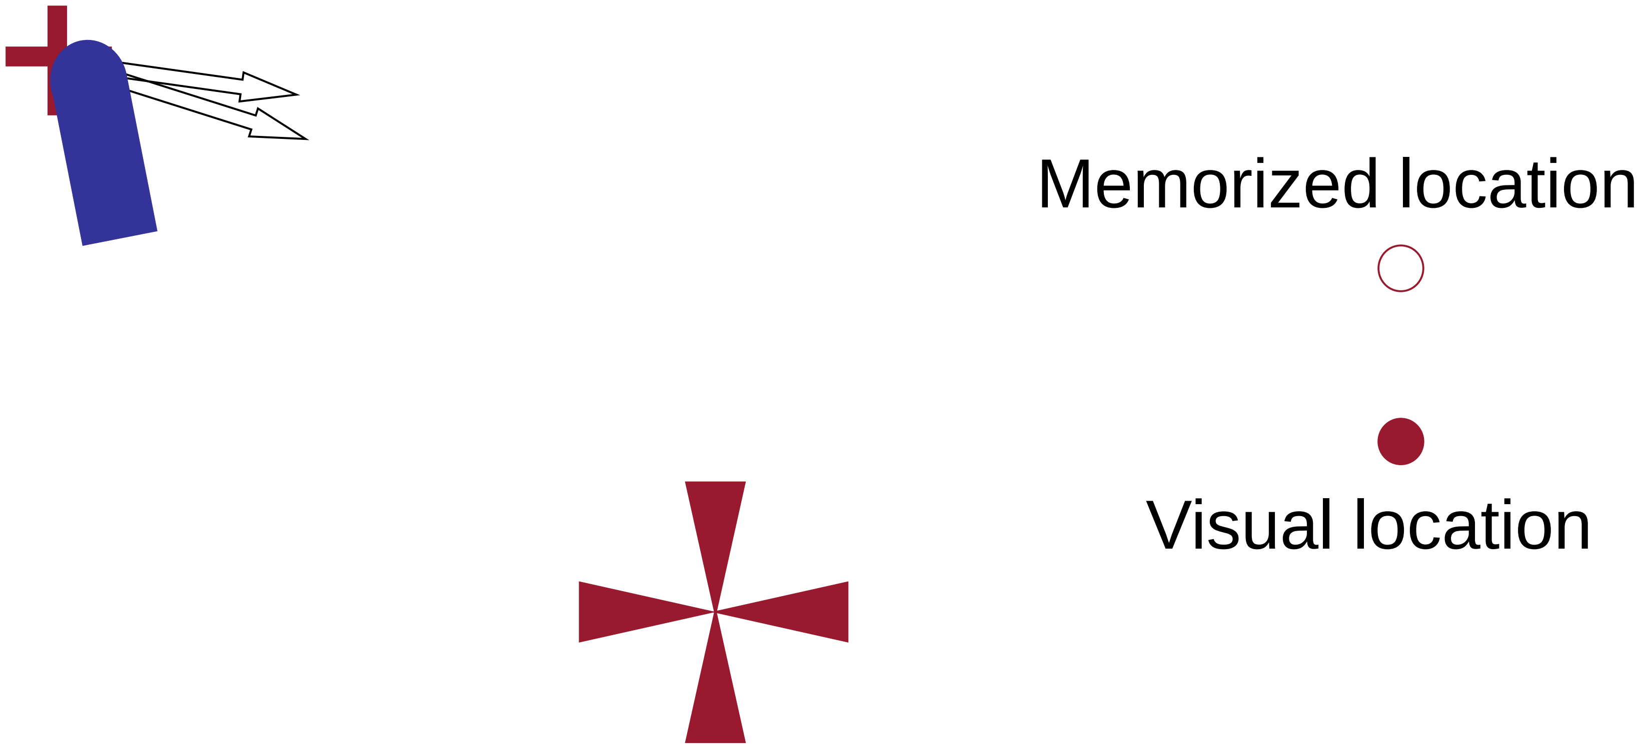
Memorized location
Visual location

Flicker to extend the time that reflects information used for planning

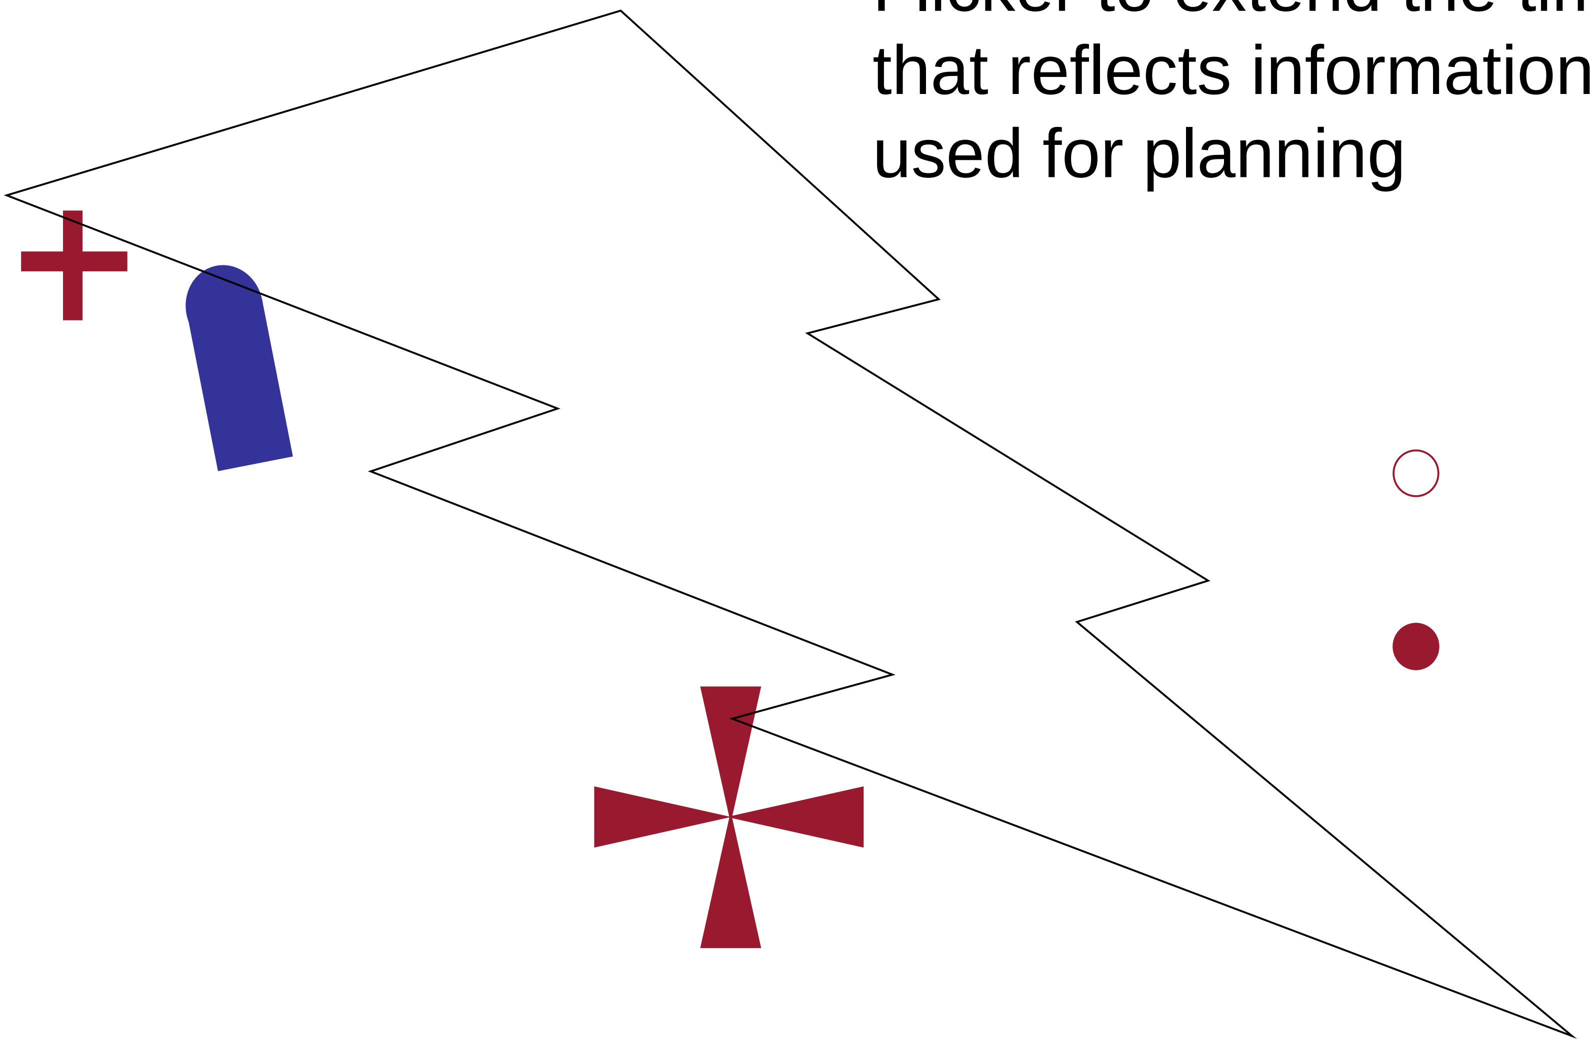

"I was a postdoc of Dave's. I had my job interview at VSS 2005, in the sun. Dave immediately struck me as a very friendly, thoughtful person. Even though in Rochester the environmental circumstances were quite different (no windows in the whole department!), Dave indeed turned out to be the friendly, thoughtful person consistent with my first impression, and so much more. To me he represents the true scientist, who wholeheartedly wants to get at the bottom of it, rather than being distracted by status and petty politics. I very much value the time spent in his lab, and will not forget our discussions and the occasional TGIF, having a beer with Ross, Hal, Brian, Bo and Dave amidst the experimental setups (-: "

Anne-Marie Brouwer
(Dave's postdoc from Sep 2005 - May 2007)

2007-8: Extending to large jumps

Feb 2009

Hi Dave,

I hope all is well. I would like to discuss with you if you would prefer me to step down from the project we have been doing together. I am well aware that between setting up my lab, writing grants, and preparing a course from scratch, I have neglected working on our project. Although I greatly enjoy it and would still be interested in carrying it through to a conclusion, unfortunately I do not anticipate having more than a few days a month to devote to it, as has been the case in the past half year. I can imagine that you need the results faster, or that you have students or postdocs who would be interested in this project. If that is the case, I do not want to be an obstacle. I could easily transfer my analysis files to someone else. Let me know!

Best,
Weiji

Hi Wei Ji,

February 24, 2009

Thanks for your thoughtful note. I'm afraid that I've been as bad as you about putting time into this - maybe we should both step down :). At the moment I do not have someone to step right into the breach, so it's ok with me if you want to stay with it. Right now, we seem to be in a place where we might have to re-design the experiments and collect more data. I understand, though, if you want to step down. Sometimes these side projects end up being more of a psychological burden than anything else. If you are feeling that way, I completely understand and I'm ok with you stepping away. You don't need to feel badly about it. If you are interested, though, in continuing, even if it is at a slow pace for now, I'm also ok with that. I have a new post-doc starting sometime this summer who is interested in some of these issues. If you decide to stay on, it's possible that you could shift your role at that point to a more conceptual one and he would take on more of the detailed data analysis and modeling. I'm not sure what he's going to end up working on, so I guess what I'm really saying is that if I find someone to take on more of the project, you could still stay involved (pretty much in the way I am currently).

Let me know what you decide.

BTW, how are things going there? How do you like life as a faculty member?

Best,
Dave

A unifying account of visual motion and position perception

Oh-Sang Kwon[1], Duje Tadin[1] and David C. Knill[1,#]

Kwon, Tadin, Knill, PNAS, in press

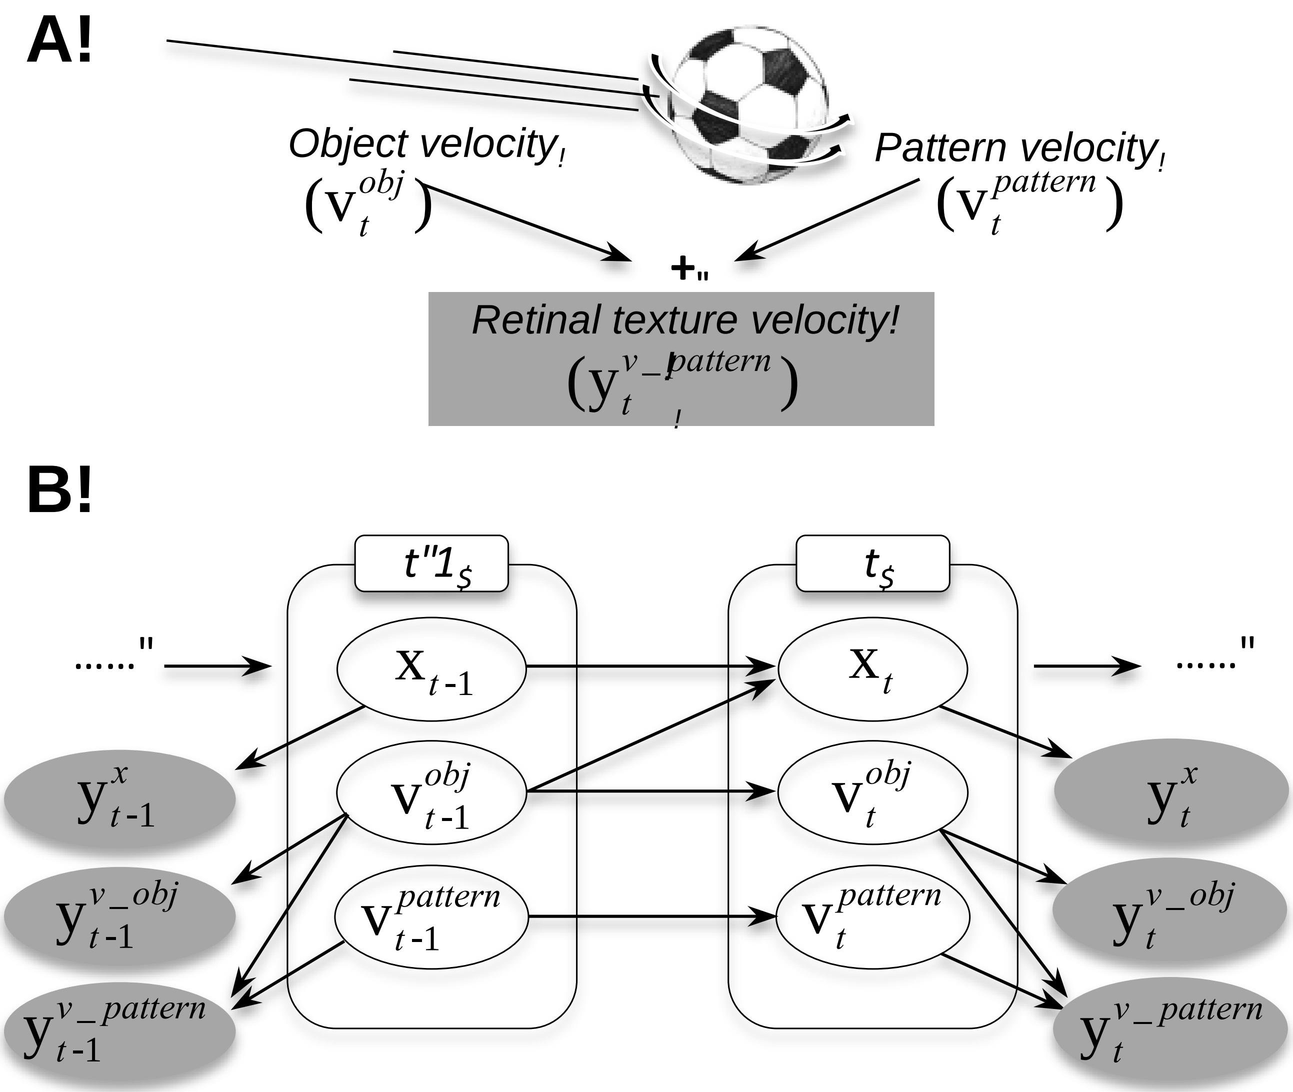

Kwon, Tadin, Knill, PNAS, in press

- motion-induced shifts in perceived position
- peripheral slowing
- curveball illusion

"Dave was my postdoctoral mentor. He was brilliant yet strict in research, but generous in life. It was always enlightening to have a meeting with him. He could see through to the core of my vague ideas, which I might have been thinking about for weeks. Usually before my full description finished, he would come up with several better ones in mathematically organized form. I adored his ability. He encouraged me to explore fundamental principles governing human behaviors rather than to search for eye-catching effects, while urging every bit of research to be crystal clear. I am greatly indebted to him."

Oh-Sang Kwon
(postdoc 2009-2014)

Talk 21.14, tomorrow morning

Rochester, Feb 2011

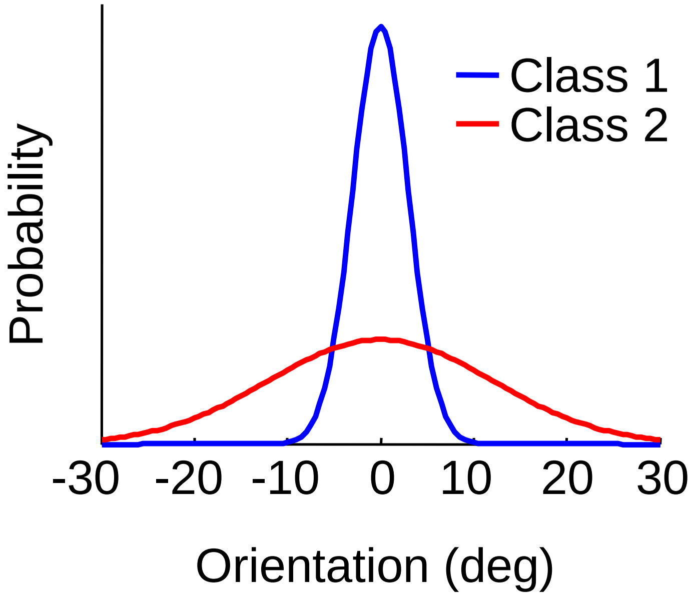

Liu, Knill, Kersten 1995

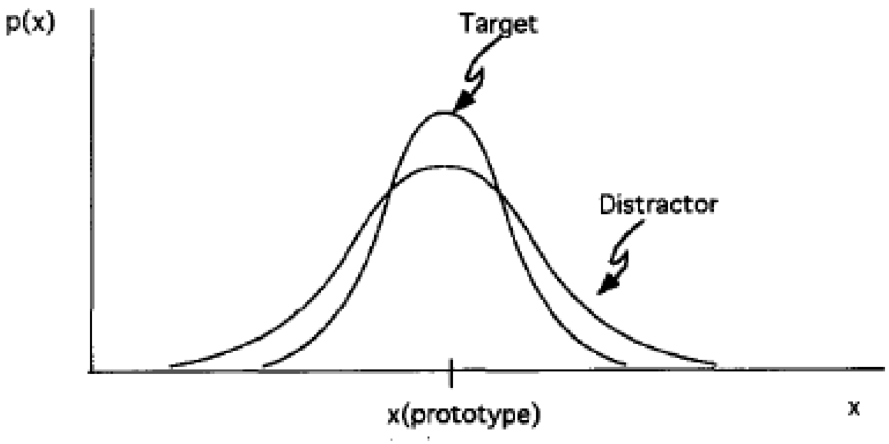

Portland, 2012

VSS 2014