


Discovering and Modeling Mechanisms



Review-Feedback Relations

- Although sometimes operations are organized sequentially, in many mechanisms operations later in a sequence feed back and alter those earlier in the process
 - Negative feedback
 - Positive feedback
- These relations often serve to integrate components into a mechanism in which parts work together to produce a phenomenon

Review: Small Worlds

- In a network in which all connections are between neighbors, the average number of operations needed to get from one node to another is very large
- Yet in the real world, there are often very few connections to get from any one node to another
- A few long distance connections reduce the distance dramatically
 - Often into the range of 6 plus or minus 3 connections

Clicker Question

- What makes the human population a small world
 - A. that all all individuals are located very near each other
 - B. there some individuals know someone elsewhere in the world
 - C. that the network of humans is random—each individual has equal likelihood of knowing every other individual
 - D. that each of us has a low Bacon number

Review - Scale-free Networks and Hubs

- In many real world networks, the number of connections originating from a node is not distributed normally but according to a power law
 - Highly connected nodes serve as hubs, linking other nodes into collectives
- Scale-free networks are robust to random distribution
 - But can collapse if hubs are destroyed

Hierarchy of Mechanisms


- In many real-world systems, nodes are collected into local systems that works together to achieve some end—constitute mechanisms
 - But these mechanisms are also interconnected through a more limited set of interactions between nodes—themselves constituting mechanisms
 - And these interconnected mechanisms are again interconnects
- In such a hierarchy, researchers need to both
 - look down to study what the components do (reductionism)
 - look up to study how they are connected into larger systems that affect how the components behave

Discovering How Mechanisms Work

- What are organized in mechanisms are parts (entities) performing operations (activities or causal processes)
- When it is possible, the most effective way to understand parts and their operations is to experiment
- Experimentation involves intervention and manipulation of a possible cause to determine what effect it has—and CONFOUNDS remain the worry!
- The only differences from our earlier discussion are
 - This is being done within the context of a organized system
 - The goal is to understand how the parts contribute to the working of the mechanism

First step: Getting the Phenomenon Right

Pick a card and think real hard about it. Don't forget it.



Whoops. I thought you were cheating and tried to shoot you. But I got your card instead.



Check it out—did I take out your card



Correctly characterizing perception

One view: Perception is like photography—creating a picture for a homunculus to view

This fits our phenomenal awareness, but that may be a false lead—

- As you just saw
- Remember change blindness

Alternative: Perception involves extraction of information in format useable by down-line systems

There may be no place where all the information comes together

We may not in fact even perceive much of the scene we are looking at

The Mechanism

Your task—figure out how this mechanism produces the phenomenon of seeing (only one of the phenomena for which it is responsible)



Intervening to discover how a mechanism works

The operations of the components of a mechanism result in causal interactions between components
Just as with simple causal interactions between a independent and dependent variable

Correlated activity suggests causal linkages


But the best evidence for causation comes from manipulating and so controlling changes

- Manipulating the input to the mechanism and determining the effects
- Manipulating components of the mechanism and determining their effects

Figuring out how it works

Three basic strategies for figuring out what the components of a mechanism do

- Recording from an individual component while the mechanism is operating
and inferring from the conditions in which it is active what operation it (or components prior to it in the pathway) might be involved in



Figuring out how it works


- Inhibition, lesion, or ablation studies:
Lesioning or ablating a component

and inferring from the deficit in the behavior of the whole what operation the

component likely contributed to

- Excitation or stimulation studies: stimulating a component

and inferring from its effect on the whole system what operation it likely contributed to



Clicker Question

If you removed a part from your MP3 player and it no longer produces sound, you can infer

- A. The part you removed was itself causally responsible for generating sound
- B. The part you removed probably had no direct role in generating sound
- C. The part you removed probably figured in the process of generating sound
- D. The part you removed probably performed a lot of other operations as well

Start with Lesions


Until the 1940s, there was no way to record the activity of individual neurons in the brain

Crude stimulation (with the electrodes of the 19th century) activated very wide areas, and so not sufficiently specific

That left ablation or lesion as initially the tool of choice

Occipital Lobe

OCCIPITAL LOBE



Bartolomeo Panizza, 1855

Proposed occipital lobe as locus of vision based on patients who experienced blindness after strokes (experiments of nature)

Hermann Munk, 1870s

Unilateral removal of the occipital lobe in dogs resulted in partial blindness
Bilateral removal resulted in total blindness

Getting More Detailed

Tatsuji Inouye, a young Japanese ophthalmologist, was assigned to assess visual loss in Japanese soldiers who had suffered brain injury so as to determine how large their pensions would be, but he decided to make the job more interesting and map deficits in particular parts of the visual field onto the area damaged


The Russians had developed a new high-velocity rifle, (Mosin–Nagant Model 91), that fired a 7.62-mm hard-jacketed bullet. This bullet pierced the skull without shattering, leaving tidy entrance and exit wounds

This made it possible to trace the trajectory and determine just what parts of the occipital lobe was damaged



Retinotopic map of the visual field


Inouye correlated the parts of the visual field in which his patients were blind with areas of brain damage, and mapped the visual field onto the visual cortex




Gordon Holmes (1918) constructed a similar map based on studies of soldiers injured during World War I.

So Visual Cortex is Needed—but What does it do?

- Lesions can show that a part of the mechanism seems to be necessary for it to perform a specific phenomenon
 - But are not able to show more specifically what the damaged part does in normal situations
- A powerful *complement* is to record from the component while it is in operation to see what elicits its activity
- Made possible by the development of electrodes wired to amplifiers and speakers




Single-cell recording



Single-cell recording in Retina and LGN

Stephen Kuffler applied the technique of single-cell recording to the retina and lateral geniculate nucleus (LGN) and found cells that fired either to

- Light dots with dark backgrounds (on-center, off-surround)
- Dark dots with light backgrounds (off-center, on-surround)





Turning to Cortex



When a technique works once, it makes sense to try it again
David Hubel and Thorsten Wiesel tried to replicate Kuffler's achievements in occipital lobe

And failed, and failed, and FAILED

BUT, one day while they were inserting a glass slide into their projecting ophthalmoscope, it stuck, creating a bar of light on the screen


Hubel reports that "over the audiomonitor the cell went off like a machine gun"

- Bars of light (edges), not dots, activate occipital cortex

Hubel and Wiesel's Simple Cortical Cells

Many of the cells Hubel and Wiesel tested in occipital lobe responded to bars of light


If they were properly oriented



Hubel and Wiesel's complex cells


Some cells Hubel and Wiesel tested responded to bars of light anywhere in the receptive field of the cell

or
If they were moving in a preferred direction across the field




How do Simple and Complex Cells do it?

Hubel and Wiesel proposed simple model wiring diagrams to show how simple and complex cells could perform their different tasks



Simple cells: Fire if enough LGN cells with centers on the bar are active



Complex cells: Fire if one or another simple cell detecting a bar is active (or if several become active in sequence)

Clicker Question

What can lesion studies show that recording studies cannot?


- A. That the part in question is actually involved in performing the activity
- B. That the part in question might be sufficient for performing the activity
- C. That if the part in question were stimulated, it would enhance the activity
- D. That the part in question actually performs a wide range of operations

Beyond edge detection

The cells Hubel and Wiesel found are all located in one part of the occipital lobe known as the striate cortex, Brodmann's area 17, or V1 (visual area 1).


Detecting edges is important to seeing, but it isn't the whole story, as Hubel and Wiesel recognized:

"Specialized as the cells of 17 are, compared with rods and cones, they must, nevertheless, still represent a very elementary stage in the handling of complex forms, occupied as they are with a relatively simple region-by-region analysis of retinal contours. How this information is used at later stages in the visual path is far from clear, and represents one of the most tantalizing problems for the future" (Hubel and Wiesel, 1968, p. 242).



Additional areas in extrastriate cortex

V2—adjoining V1: cells respond to illusory contours



V4—further forward from V1 and V2—cells responded to color: Zeki “in every case the units have been colour coded, responding vigorously to one wavelength and grudgingly, or not at all, to other wavelengths or to white light at different intensities”

The Woman Who Couldn't See Motion




- When Gisela Leibold tried to pour coffee, she could see the cup's color, shape, and position, and could tip the pot
 - But what she saw was like a frozen waterfall
 - She couldn't see anything moving
 - Just a sequence of still life paintings
- Result of a stroke she had suffered that affected area known as MT

Cells that process motion

The phenomenon of perceived motion:

Two stimuli, moving in different directions, will sometimes be seen as one stimulus moving in a combined direction




V1 cells respond only to actual motion, not perceived motion

So they do not compute perceived motion

But cells in area V5 (MT) do respond to perceived motion

Adding microstimulation



William Newsome trained monkeys to indicate the direction of motion they perceived correlated motion in ambiguous displays

Recording from MT cells showed that the responses of those cells predicted the animal's behavior

Microstimulation of those cells biases the behavioral response

This combination of recording and stimulation studies (when combined with the lesion results) offers powerful evidence about what these components are doing

Clicker Question

Making a mechanism produce the phenomenon of interest by stimulating a part of it serves to establish

- A. That the part is sufficient for the mechanism to perform the phenomenon of interest
- B. That the part is necessary for the mechanism to perform the phenomenon of interest
- C. That the part can initiate a causal process resulting in the phenomenon of interest
- D. That if one lesioned the part, the phenomenon would be destroyed


Understanding Motion Perception

- Lesion deficit pointed to MT as a likely motion area
- Single cell recording provided further evidence that the area was involved in detecting motion—perceived, not just actual motion
- The ability to stimulate the area and enhance the effect further confirmed the result

- All three methods brought to bear to figure out this mechanism

Object Recognition in Inferotemporal Cortex


Recoding from single cells in Inferotemporal Cortex, Charles Gross found cells that responded to specific shapes: e.g., hands.



New tools for recording activity

Introduction of PET and fMRI in the 1980s and 1990s provided a way to record (indirectly via blood flow) activity in brain areas of humans while engaging in actual tasks

Nancy Kanwisher identified an area (in the fusiform gyrus) that responds particularly to faces



Prosopagnosia: Face Blindness

A self report:

- When I look at a face, I see the same thing that I suspect you do. My vision works fine (other than some autistic difficulties that aren't relevant to this discussion). My brain sees a face much like any other object. The problem I have is in associating that face with a particular person I know.
- "I recognize people by three primary methods - general body size/shape, hair, and the sound of their voice. These three methods are not nearly as effective as the normal way of recognizing people - by recognizing a face. Thus, I often mistake someone I don't know for someone that I do know or I fail to recognize someone I know. For instance, I have been unable to recognize my father on multiple occasions, since his body size and shape are not very distinctive, nor does he have long or distinctive hair."

Clicker Question

Given the evidence that the fusiform face area is active when people view faces and that face blindness results when it is damaged, why might other researchers still resist treating the area as a *face* area

- A. The area might have nothing to do with recognizing faces
- B. The operation performed in the area may not be limited to faces
- C. The area might not be a sufficient cause of recognizing faces
- D. Researchers have not yet done a stimulation study to determine if stimulation results in reports of faces

But is it a “face” area?

There is little doubt that the area Kanwisher identified responds particularly well to faces


But like any recording study, we don't yet know what else it might respond to

- Some evidence that it responds to objects for which detecting individual identity is important



38

Putting it back together: The visual system as a complex,






Ungerleider and Mishkin (1982): two pathways of visual processing

- Pathway into temporal cortex: what
- Pathway into parietal cortex: where

Different Deficits when Pathways Damaged

When deficit in parietal (where) pathway, patients are unable to put hand through slot in correct orientation





When deficit in temporal (what) pathway, copying is slow and slavish and patients cannot name object

Mechanism for Visual Processing

Van Essen: Schema of overall organization of visual processing

Represents the combined efforts of recording, lesion, stimulation

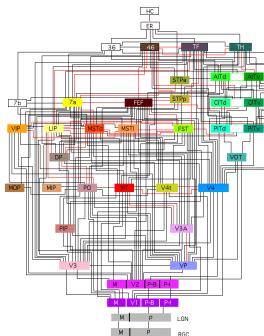


The visual mechanism only partly understood

We know of more components in the visual system than we know what they do

The mechanism is only partly understood

As we learn about other component, revisions in the current account will be required



Each technique is limited

Each approach—recording, lesioning, and stimulating—requires inference *and inference is fallible*

- Just showing that a component is active given a specific stimulus does not tell you
 - Specifically what about the stimulus it is responding to
 - What it is doing in response to that feature
- Just showing that lesioning a component interrupts an activity does not tell you
 - That it alone was responsible for the activity
 - What it contributed to the activity
- Just showing that stimulating a component increases the performance of the activity
 - Does not tell you how it figured in generating the activity

Need to orchestrate multiple techniques

- There is no foolproof strategy for figuring out how a mechanism works
- The best results stem from combining different strategies to determine what the components of a system are and what they do

Where we have been in this class

- Logic: structure of arguments for confirmation and falsification
- Observation: variables and their measurement
- Correlation: predictions based on correlations and statistically significant differences within samples
- Causation: experimental and non-experimental evidence, and strategies for controlling confounds
- Mechanisms: discovering how component parts, operations, and their organization yield a system that exhibits the phenomenon of interest

Final Thought—Or a First Thought Repeated

Reasoning and making decisions, whether about
Perception
Correlation
Causation
Mechanism
is fallible

We can (and should) strive to come closer to the truth and
rely on the best information available now

But we must also recognize that tomorrow
something might be discovered that makes us
revise our best conclusions of today
