Perception, cognition and human performance

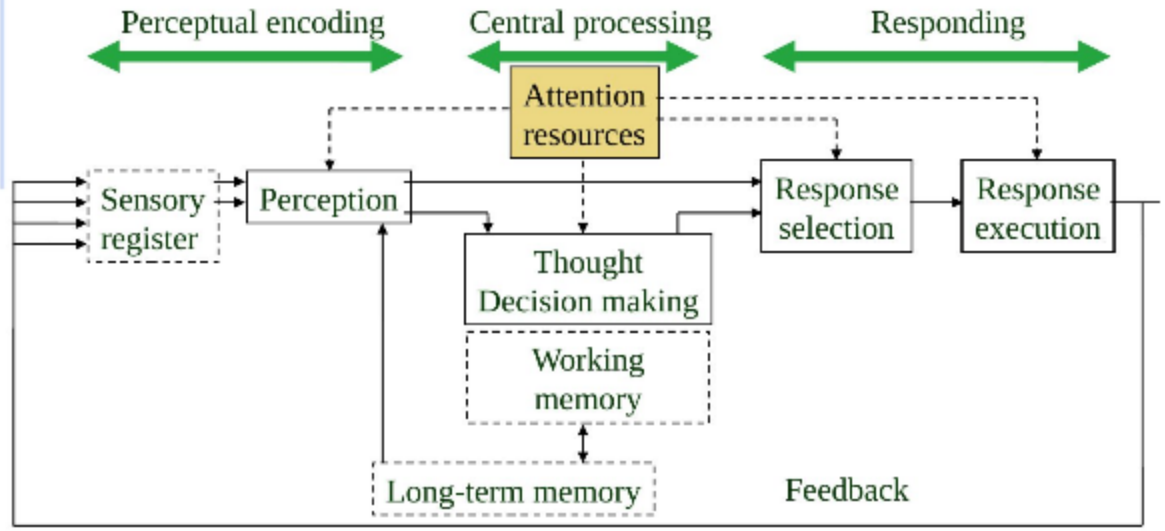
Lecture 2 · Week 2

CSE566 · Mehmet Göktürk

From interface difficulty to an explanation

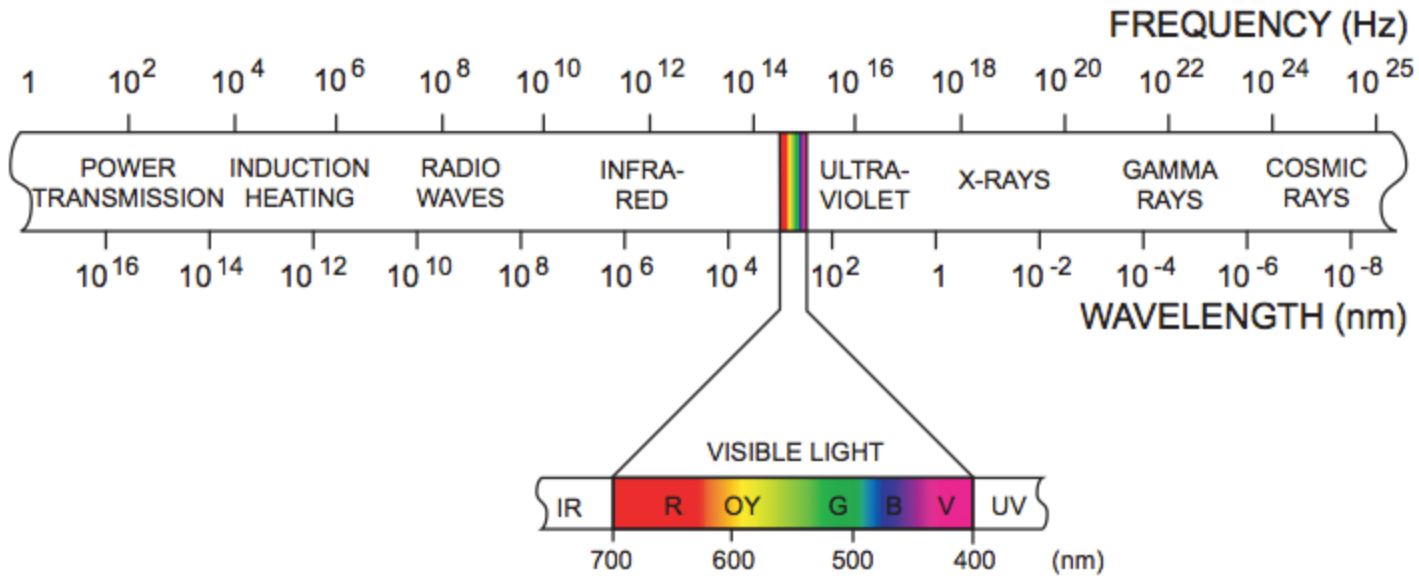
An interface asks people to perceive, attend, remember, interpret and act. These demands interact through feedback.

Our goal is to explain a difficulty, predict a useful change, and identify what the prediction leaves out.



Wickens et al., 1998

Visible light: a small part of the spectrum



Vision responds to a narrow band of electromagnetic wavelengths, approximately 400–700 nm.

Wavelength relates to spectral color; intensity and the mixture of wavelengths also affect appearance.

The illuminant changes what reaches the eye

Daylight, incandescent lamps and fluorescent lamps have different spectral distributions.

A surface does not have the same appearance under every lamp: the illuminant changes the light available for reflection.

Check color-coded physical controls in the lighting where they will be used.

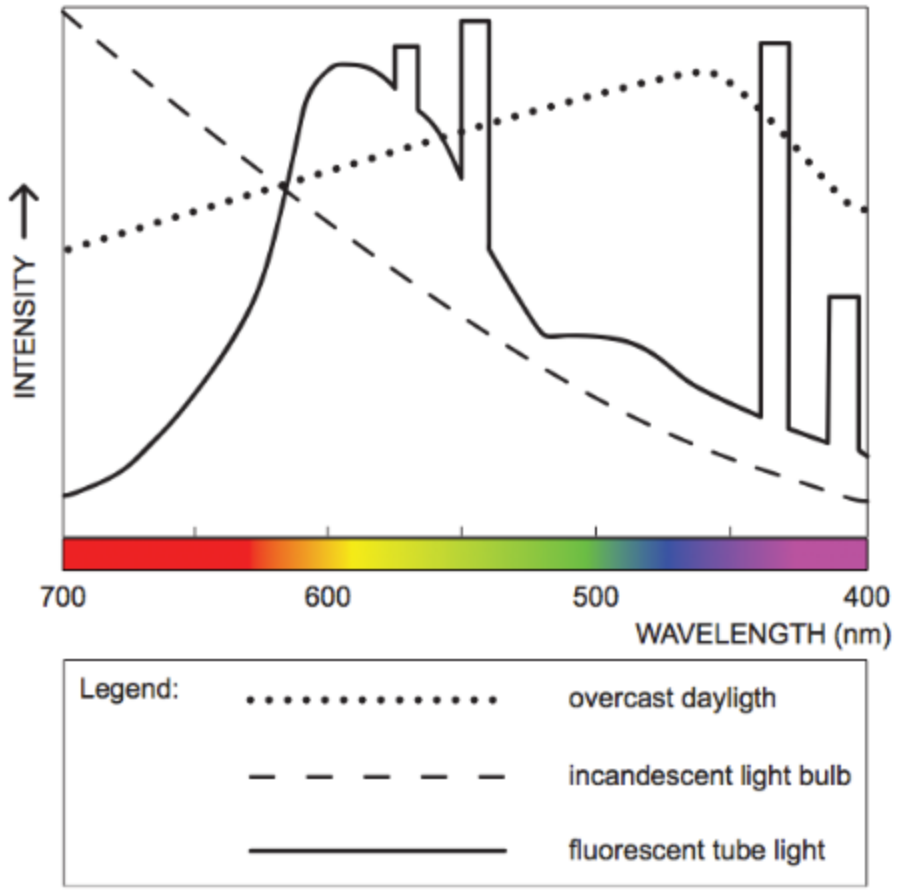


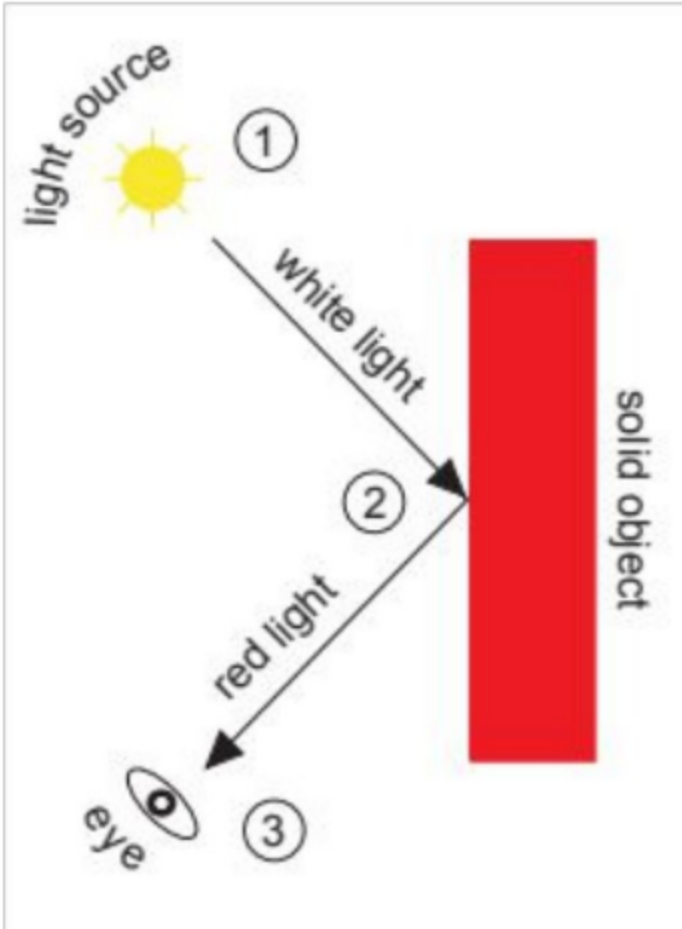
Figure 2: Spectrum of daylight, a light bulb and a tube lamp

Seeing a surface: illumination × reflectance

Light from a source reaches an object; some wavelengths are absorbed and others are reflected toward the eye.

A red surface typically reflects more long-wavelength light. It does not return only one exact “red” frequency.

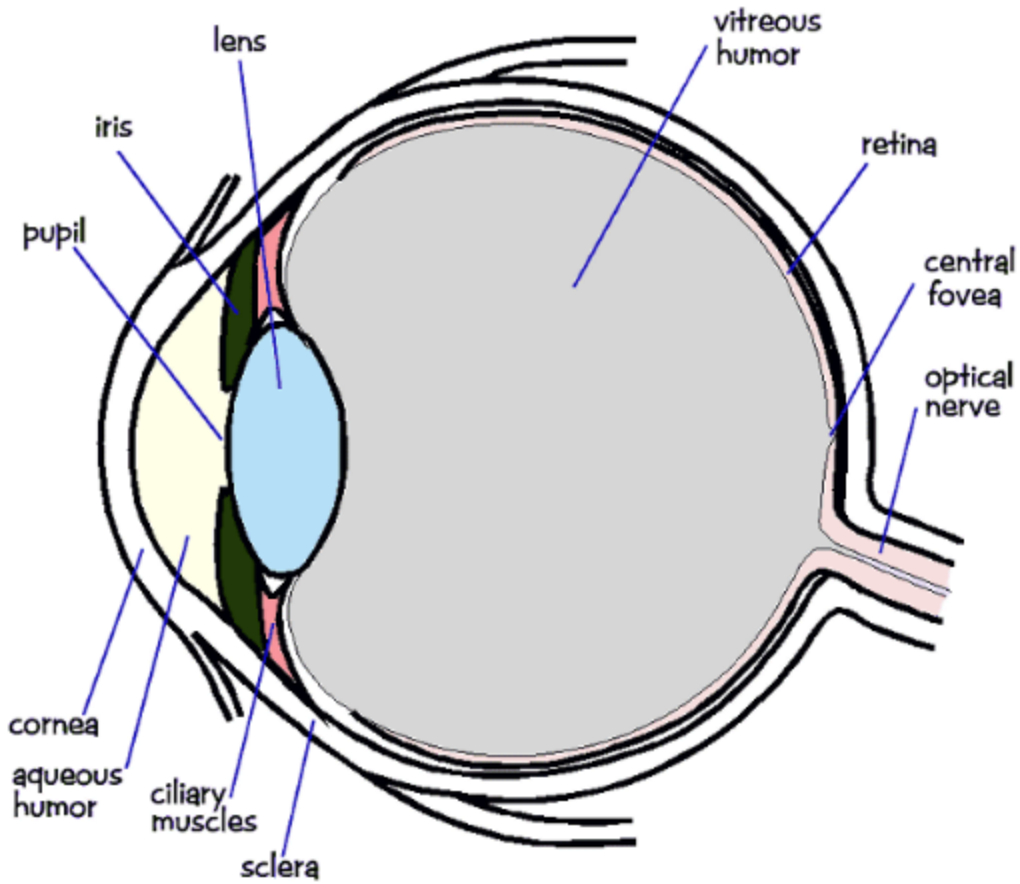
The perceived color depends on the illuminant, surface, surroundings and visual adaptation.



Seeing a stimulus and interpreting it

The eye receives light; perception interprets structure and meaning. A visible mark is not yet an understood instruction.

The same input can be interpreted differently as experience, context and the task change.

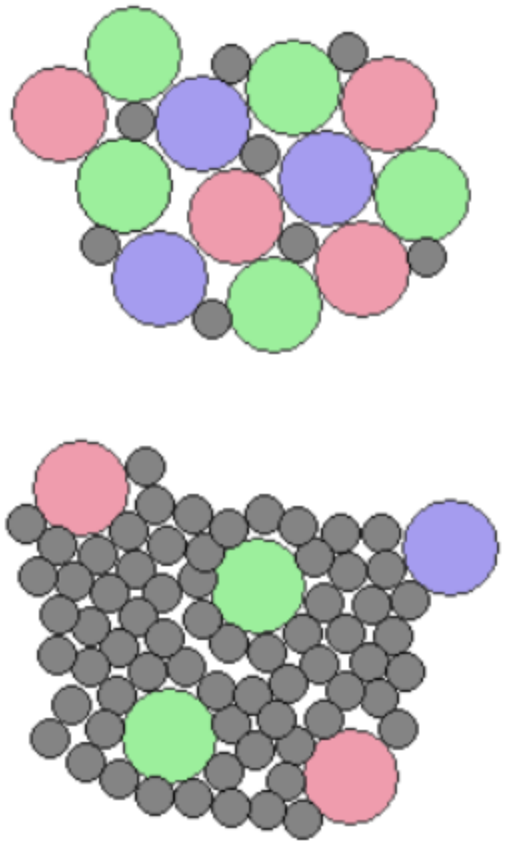


Rods and cones support different kinds of vision

Cones support fine spatial detail and color in brighter conditions. They are especially densely packed in central vision.

Rods are highly sensitive in dim light. Pooling their signals improves sensitivity while sacrificing spatial detail.

The receptor mosaic changes with position on the retina.

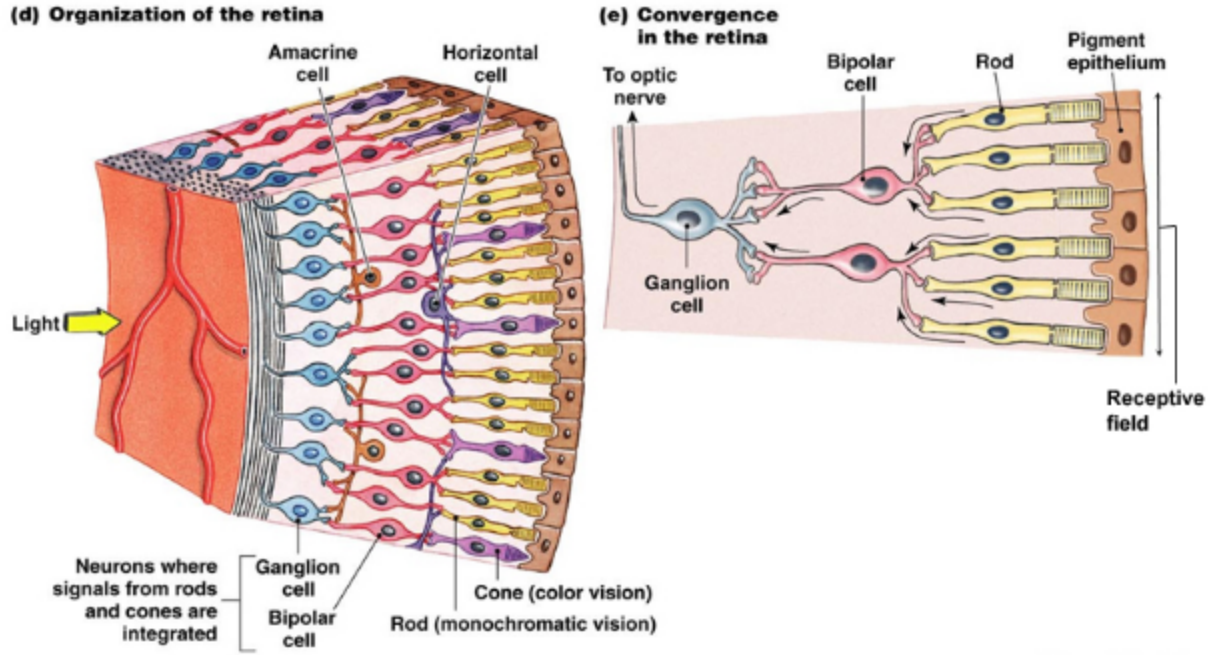


The retina begins processing the image

The optics form an inverted image on the retina. Photoreceptors convert absorbed light into changes in electrical activity.

Signals pass through retinal circuits before ganglion-cell axons carry the output into the optic nerve.

The eye is not a camera that simply sends a complete pixel array to the brain.



Copyright © 2007 Pearson Education, Inc., publishing as Benjamin Cummings.

Fig. 10-35

Retinal cells: from light to neural output

Several interacting cell types transform the incoming signal.

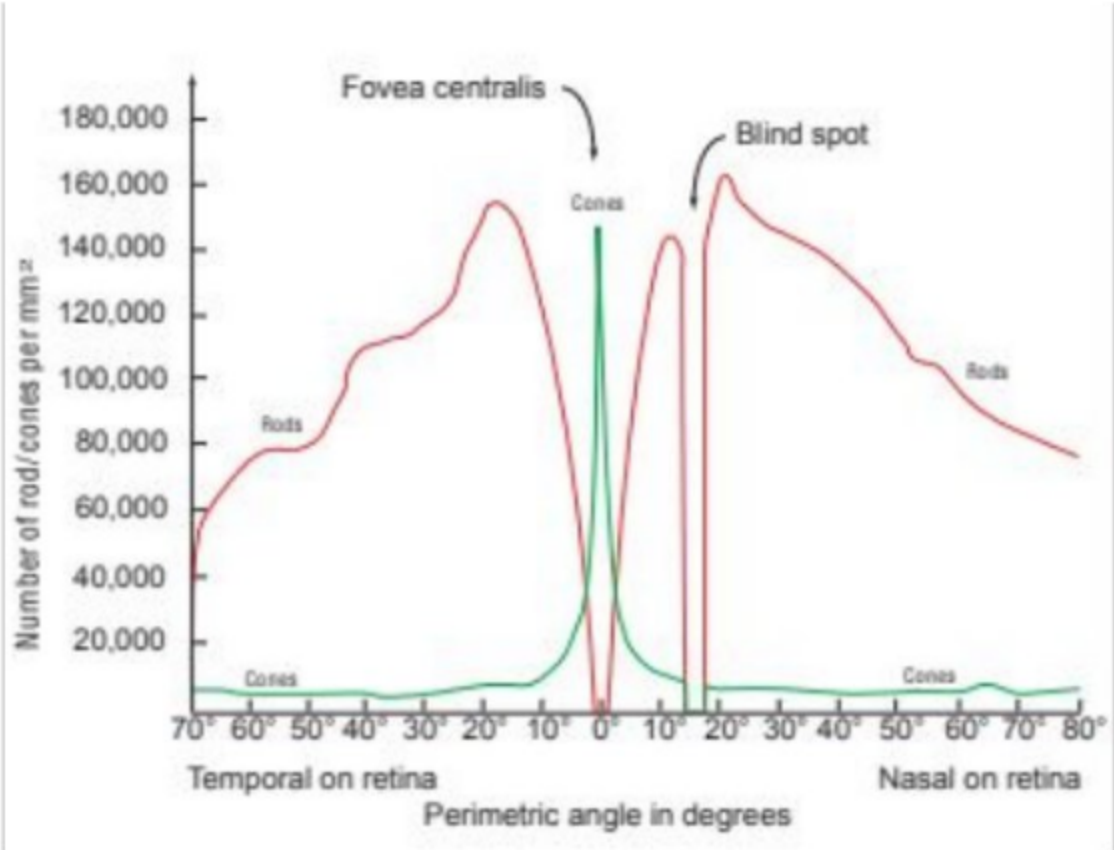
Cell or structure	Main role
Rods and cones	Phototransduction: respond to absorbed light
Bipolar cells	Relay receptor signals toward ganglion cells
Horizontal cells	Lateral interactions in the outer retina
Amacrine cells	Temporal and other interactions in the inner retina
Ganglion cells / optic nerve	Retinal output; axons form the optic nerve

Visual sensitivity is not uniform across the retina

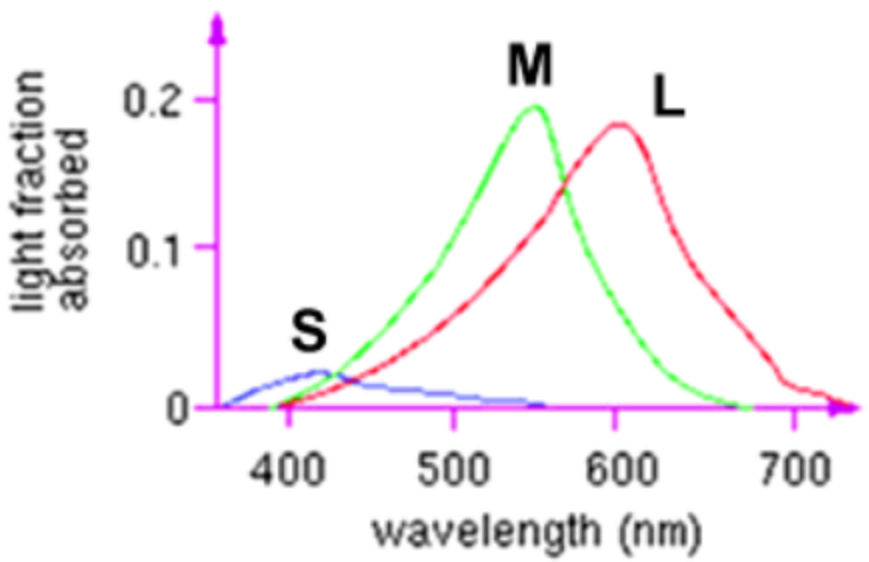
The foveal center is cone-rich and rod-free. Rod density rises away from this center.

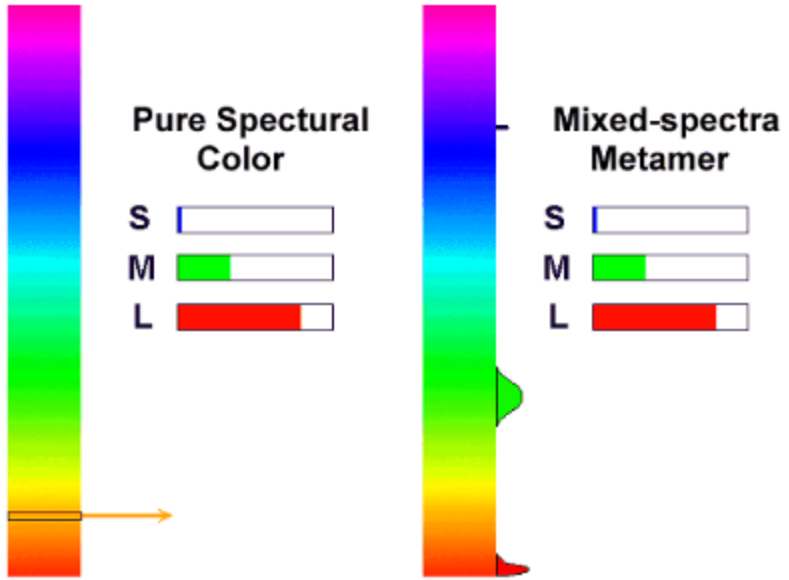
The optic-nerve exit has no photoreceptors: this is the anatomical basis of the blind spot.

S-cones are sparse and absent in the very center of the foveola; they are not confined to the far periphery.



Why three display primaries can reproduce many colors





S-, M- and L-cones have broad, overlapping wavelength sensitivities. Their relative responses support trichromatic matching.

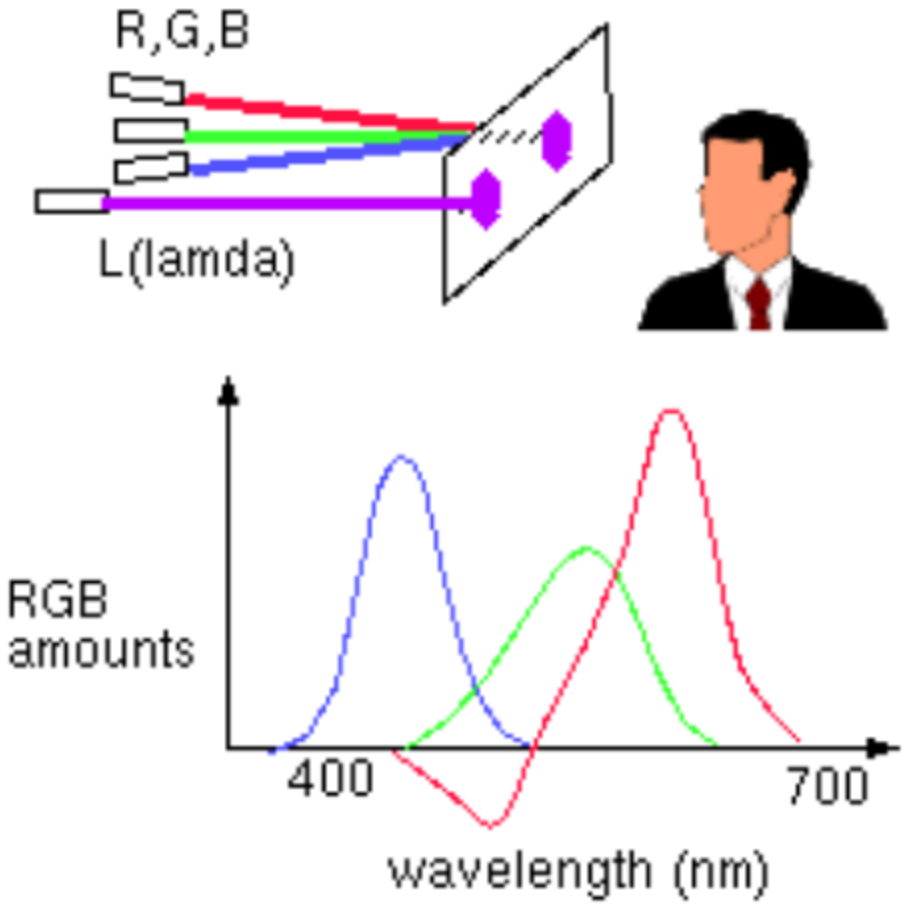
Different spectra can give similar cone responses: these matches are called metamers. RGB displays exploit this, within their gamut.

Discriminating a color is different from naming it

We can distinguish many more color appearances than we can reliably name. A tiny color difference is a poor categorical label.

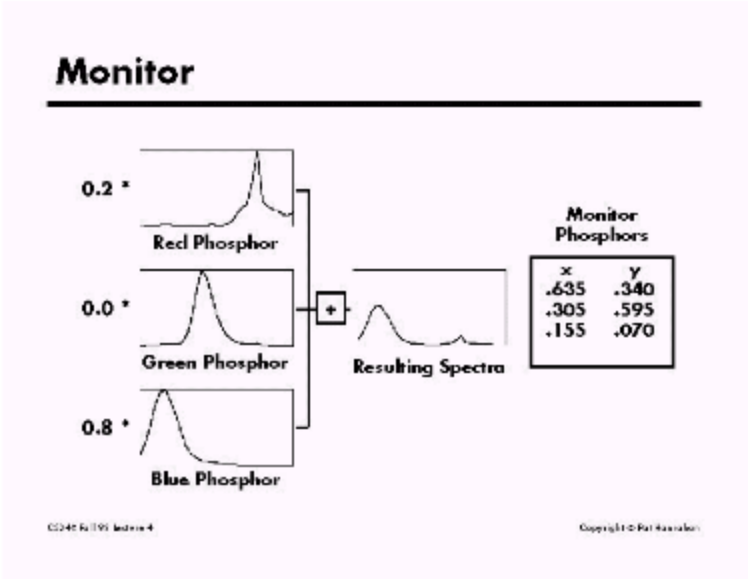
Historical estimates: about 2 nm discrimination near green versus 10 nm near spectral ends; about 7 million distinguishable colors but only 11 common names.

These are illustrative estimates, not fixed capacities. Discrimination varies with wavelength, luminance, adaptation and task.



CRT primaries: physical spectra and perceived color

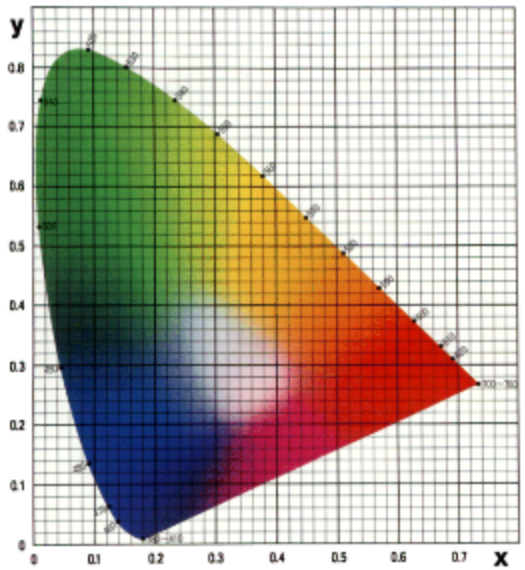


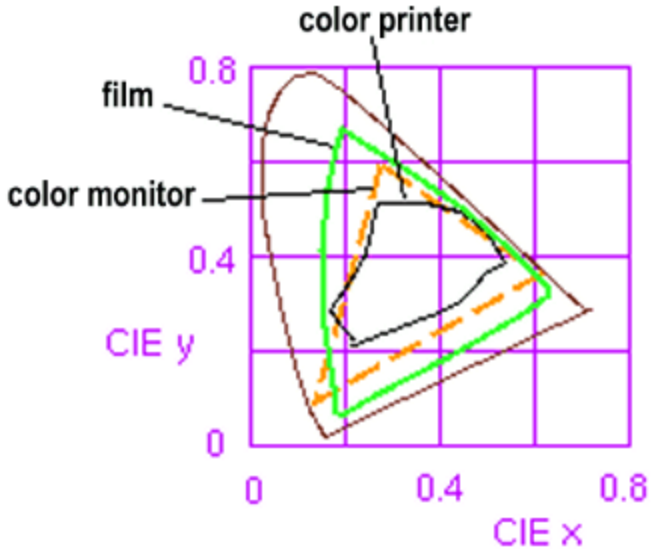


A CRT produces light with red, green and blue phosphors. Adjusting their relative output changes the mixture.

The three primary spectra are not identical to the cone sensitivity curves. Cone responses determine the resulting color match.

Gamut: a display cannot reproduce every color





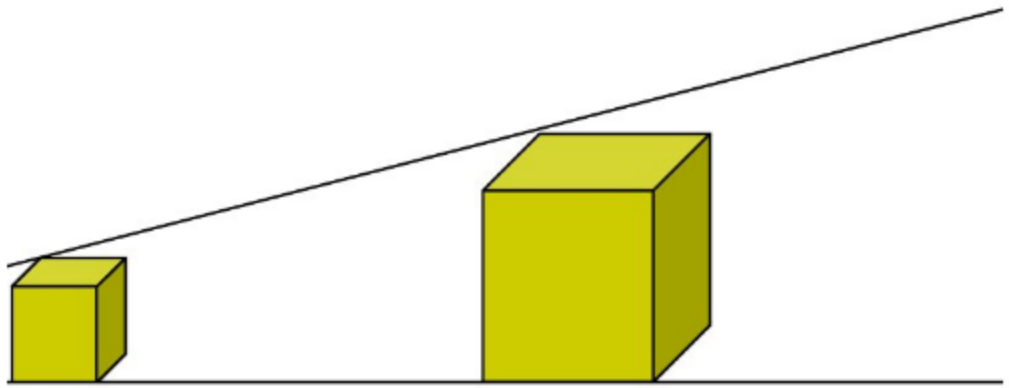
A device's primaries constrain its reproducible gamut. Monitors, printers and film can cover different regions.

A color specified on one device may need mapping on another. The chromaticity plots show color relationships, not a perfectly reproduced atlas of visible colors.

Detail depends on where we look

We move our eyes to inspect detail.
Information away from the point of fixation may signal that something happened without revealing exactly what it means.

Text size must be considered with viewing distance. A label that works at a desk may fail across a room.

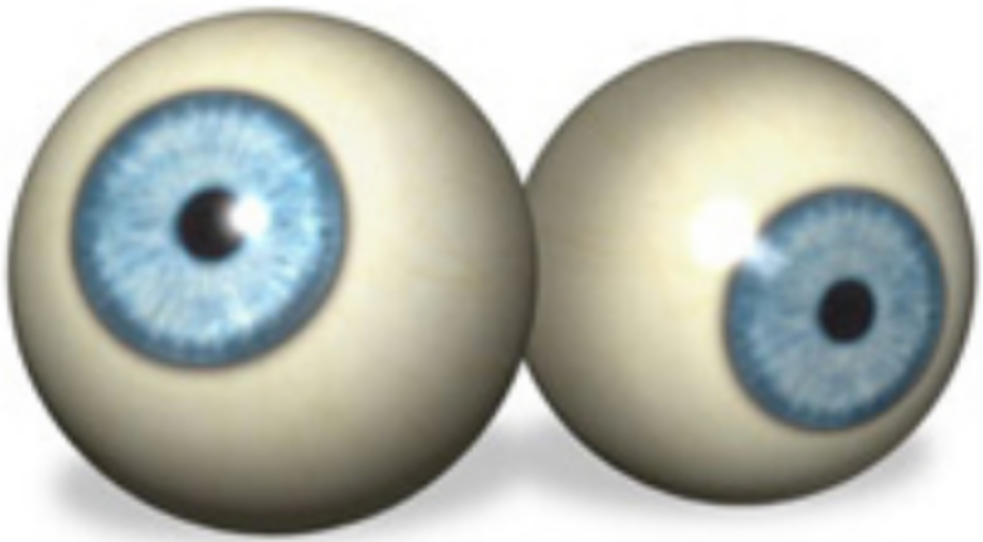


Central detail and peripheral awareness

High-acuity vision covers only a small central region. We move our eyes to inspect details rather than seeing the whole display sharply at once.

Peripheral vision can signal motion or change—for example, a nearby car moving—but it is less suitable for reading small labels.

Place essential detail near the current task; do not expect a peripheral status change to be fully interpreted.



Visual angle: connect physical size to viewing distance

$$\theta = 2 \arctan(h / 2d)$$

For h = 10 mm at d = 600 mm: θ ≈ 0.955°. The same object at 1,200 mm subtends about 0.477°.

A region spanning 2° at 600 mm is about 20.9 mm wide. This illustrates central viewing geometry; it is not a rule that a whole word must fit within 2°.

Point size is not retinal image size

12 pt = 4.23 mm for the nominal em square. At 600 mm, that em subtends approximately 0.404°.

Actual letter height is smaller and depends on the typeface. A point-size label alone cannot determine the visual angle of the visible glyph.

The same slide viewed on a laptop and projected across a room has different physical geometry.

6 pt: Aa

8 pt: Aa

12 pt: Aa

16 pt: Aa

20 pt: Aa

30 pt: Aa

40 pt: Aa

48 pt: Aa

72 pt: Aa



Retinal dimensions: useful orders of magnitude

Approximate anatomical reference values; measurements and definitions vary.

Quantity	Earlier lecture estimate
Retinal surface / thickness	About 5 × 5 cm unfolded / about 0.4 mm
Retinal distance per degree	About 0.3 mm for 1°
Fovea diameter	1.5 mm ≈ 5.2°
Rod-free center / foveola	0.5 mm ≈ 1.7° / 0.3 mm ≈ 1°
Optic disc (blind spot)	1.5 × 2.1 mm ≈ 5° × 7°; about 15° nasal

Photoreceptors: counts and dimensions

The scale helps explain why central detail and dim-light sensitivity differ.

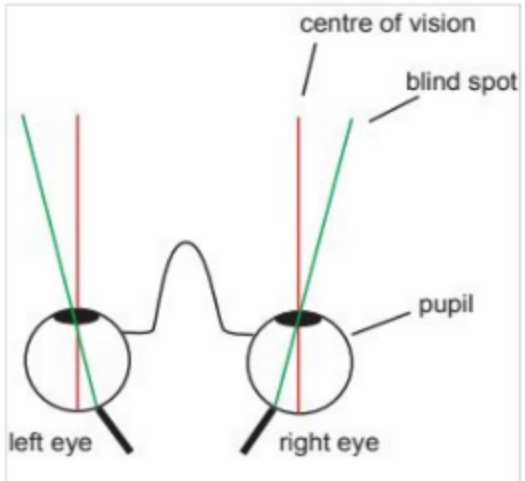
Quantity	Approximate reference value
Cone and rod populations	About 5–6 million cones; 100–120 million rods
Peak cone density	About 160,000 per mm²
Foveal cone: diameter × length	1–4 µm × 50–80 µm
Extrafoveal cone	4–10 µm × about 40 µm
Rod dimensions	About 1 µm × 60 µm

Retinal sampling and neural convergence

These historical estimates describe different regions and pathways, not one uniform sampling grid.

Quantity	Earlier lecture estimate
Central L/M-cone spacing	About 0.5 arcmin
S-cone spacing	About 10 arcmin; spatially variable
(L + M) to S-cone ratio	About 14:1; varies by region and observer
Optic-nerve fibers	About 1–1.5 million
Photoreceptor : ganglion ratio	Foveal circuitry ≈ 1:3; whole-retina mean ≈ 125:1

The blind spot demonstration



+

Cover your left eye. With your right eye, fixate the + while changing your distance from the slide until the dot disappears. Keep looking at the +, not the dot.

The optic disc has no photoreceptors. The missing input is usually not experienced as a black hole: perception supports a continuous view.

Adaptation changes what we can see

A display is experienced in an environment, not in isolation. Moving between bright and dark surroundings changes visual sensitivity.

For design, consider illumination, glare and adaptation together. Increasing brightness alone may not restore readable detail.

Context	Visual demand
Bright surroundings	Detail competes with glare and ambient light
Dark surroundings	Excessive brightness can be uncomfortable
Change of lighting	Sensitivity adjusts over time

Adaptation changes the operating range of vision

Photopigment bleaching and regeneration contribute to adaptation, alongside neural adjustments and pupil changes.

After moving from a bright room into darkness, cone sensitivity improves first; rod-mediated sensitivity recovers more slowly.

Allow time for adaptation. A suddenly bright display can interrupt dark-adapted viewing even when the surrounding task remains dim.

The retina extracts differences, not just brightness

Early visual processing emphasizes edges, local contrast, temporal changes and motion-related information.

Parallel pathways carry different aspects of the scene. Color-related and luminance-related signals overlap rather than forming completely separate pictures.

Data reduction begins before the optic nerve: the number of output fibers is far smaller than the number of receptors.

Relative intensity: a Weber-fraction example

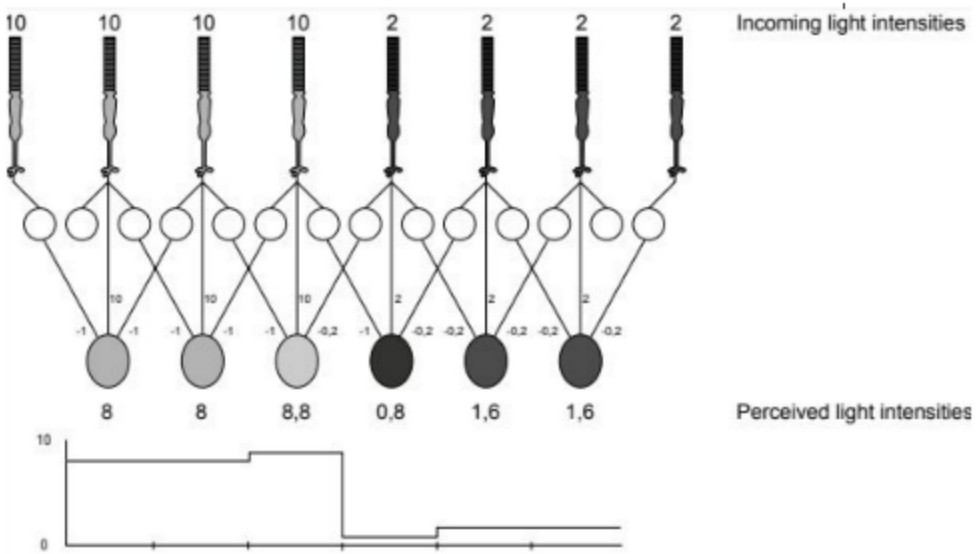
$$\Delta I / I = k$$

Increasing intensity from 0.10 to 0.11 gives $\Delta I/I = 0.01/0.10 = 0.10$. Increasing 0.50 to 0.55 gives the same fraction.

Under conditions where Weber's law is a useful approximation, proportional changes predict discriminability better than equal absolute increments.

Display gamma describes a device transfer function; it is not the same thing as Weber's law.

Lateral inhibition emphasizes a boundary



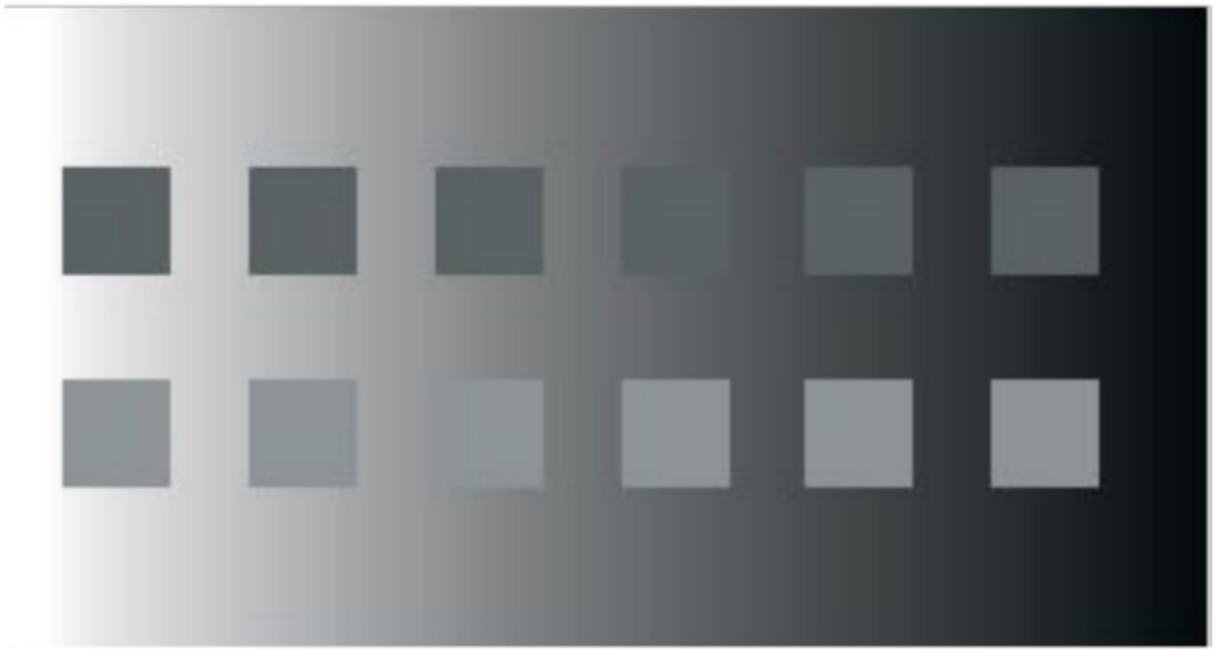
Simplified model: subtract 10% of each neighboring input.

Bright interior: 10 − 1 − 1 = 8. Bright edge: 10 − 1 − 0.2 = 8.8.

Dark edge: 2 − 1 − 0.2 = 0.8. Dark interior: 2 − 0.2 − 0.2 = 1.6.

Equal values can look different

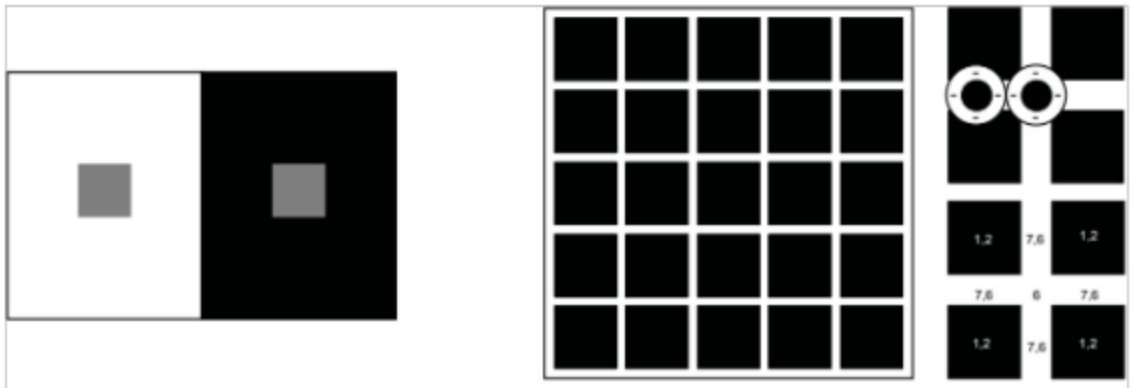
Compare the small gray squares against their different backgrounds. Surrounding brightness changes apparent contrast.



A color value does not determine its appearance by itself. Evaluate text, icons and status cues in their actual interface context.

Local context and visual illusions

Look at the grid centrally, then away from an intersection. What seems to appear or disappear?



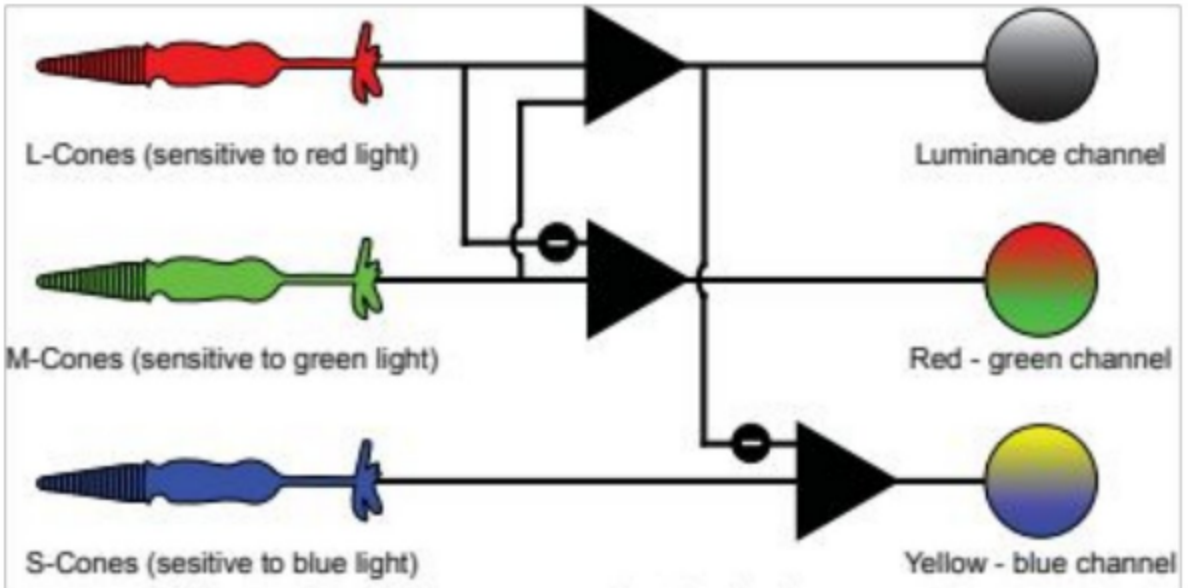
Vision processes spatial context. Apparent structure can change as fixation changes, even though the pixels on the display remain unchanged.

Opponent processing separates useful comparisons

A simplified model combines cone signals into an achromatic channel and opponent color channels.

Luminance-like: L + M. Red–green: L − M. Blue–yellow: S − (L + M), with pathway-specific weighting.

These are conceptual comparisons, not literal RGB pixel channels. They help explain why color and brightness contrasts can behave differently.



Color difference is not enough for legibility

Color difference and luminance contrast are different. Two strongly different hues can still produce difficult-to-read text.

Use sufficient light–dark contrast for fine detail. Check the actual foreground/background combination rather than the color names.



Text contrast: compare both polarities





White on black and black on white can both produce strong luminance contrast. Lower-contrast gray or warm-colored lettering is harder to read in these examples.

Judge the foreground–background combination, viewing conditions and text size; a foreground color alone does not guarantee readability.

A hue difference may still give weak text contrast

Green and brown are distinguishable color categories, yet this text remains difficult to read.

The letters need sufficient luminance contrast against their background. Thin strokes and small type make weak contrast more damaging.

Do not make readers recover a label from hue differences alone.



Figure 55: Text and background only with chromatic difference.

Shape can disappear without luminance contrast

The boundary of a shape must remain distinguishable. Hue alone can carry less structural information than we expect.

Add a second cue when a boundary or state is important: luminance, outline, label, shape or position.

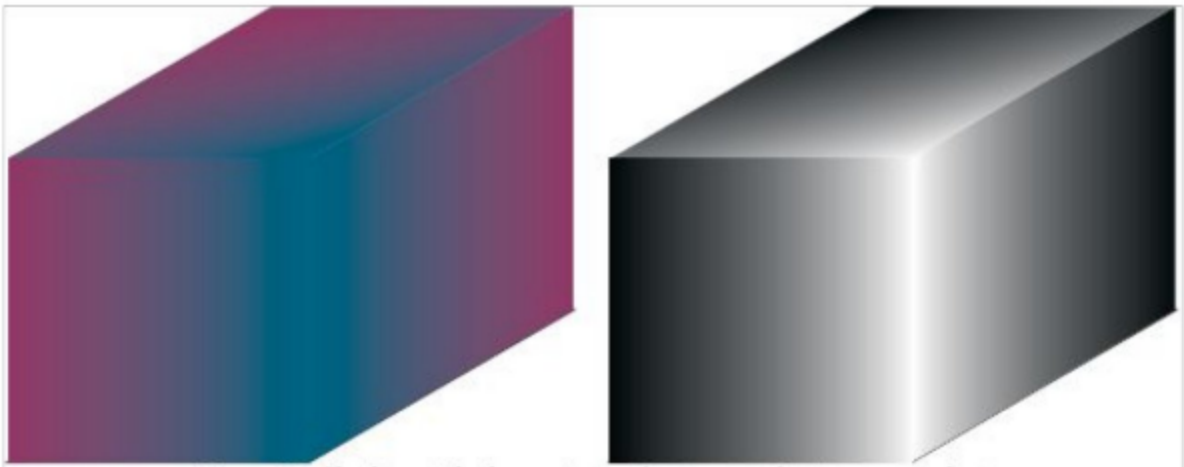


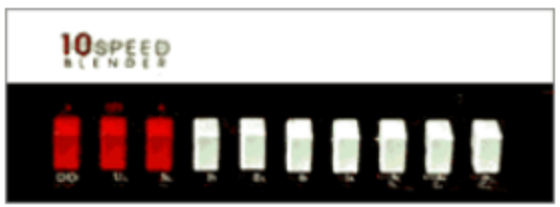
Figure 56: Shading with chromatic gradient versus luminance gradient

What does this red control mean?

Before reading the explanation: what would you expect the red controls to do?



Consider stop, danger, high power and low speed. Which expectation comes from the object, and which comes from previous experience?





The control violates the expected association

In this example, the red indication denotes slow speeds. That differs from the “high power” interpretation some users may bring.



Color should support an understandable label and control mapping. Its meaning is learned in context, not supplied universally by the color itself.



Color conventions carry learned meaning

Historical workplace palette: compare meaning across applications.

Color	Typical use in the earlier lecture
Red	Danger, stop, fire-protection equipment
Fluorescent orange / orange-red	Biological hazards
Yellow	Caution: tripping, falling, striking; some container labels
Orange	Warning around dangerous machinery

Color conventions need labels and context

A convention helps only when the audience recognizes it.

Color or combination	Typical use in the earlier lecture
Green	First aid and safety equipment
Blue	Information; railway equipment under repair
Black / white / yellow combinations	Boundaries, stairways and direction markings
Magenta or purple on yellow	Ionizing-radiation hazards

Redundant cues survive color differences

If the task depends only on distinguishing two colors, it can fail for some users and viewing conditions.



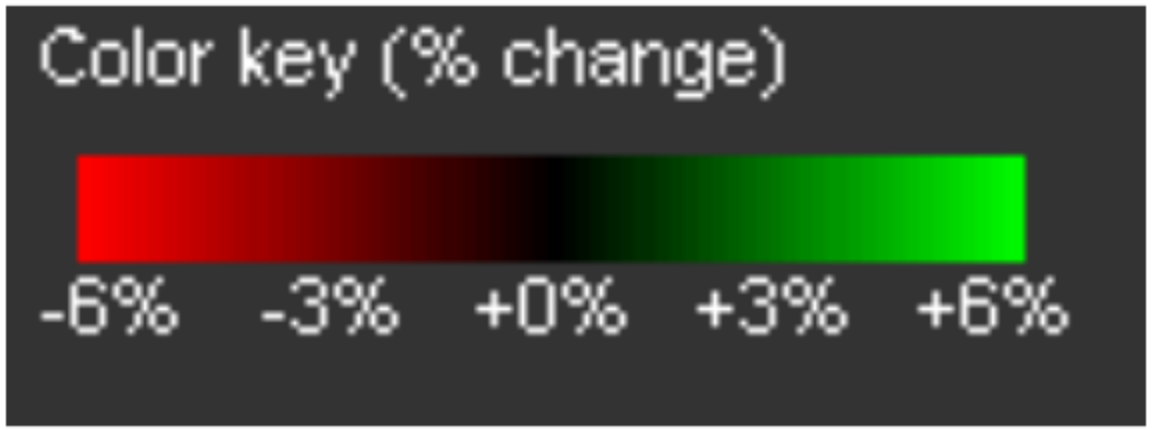


Combine color with a label, distinct shape or stable position. A red STOP control can remain understandable when its word and outline are visible without its color.





Color coding should have a purpose





A palette can encode a measured value, identify a control, or evoke a familiar meaning. Each use imposes a different interpretation task.

The color key explains a continuous value. Controls also need meaningful grouping and labels; too many unrelated colors compete for attention.

Color-vision differences affect interpretation

Color-vision deficiencies alter which color differences are easy to distinguish. Red–green differences are especially common.

Often-quoted rates around 8% of men and about 1% of women describe particular populations and definitions, not a global constant.

Pair color with text, shape or position. These plates illustrate color-dependent discrimination; classroom projection is not a diagnostic test.

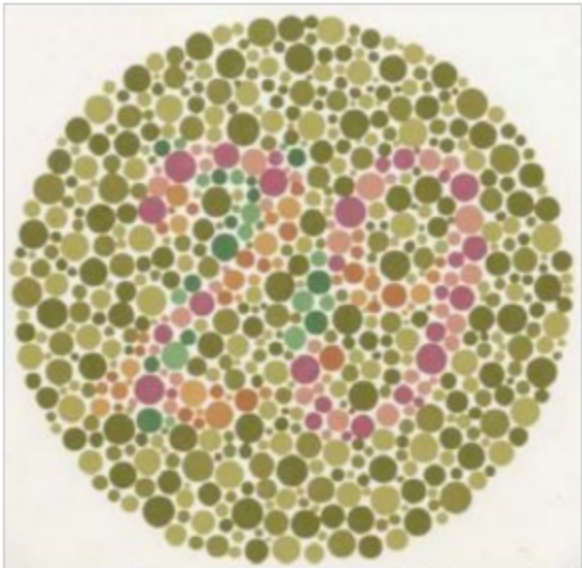


Figure 53: Test plate for colour blindness [7].

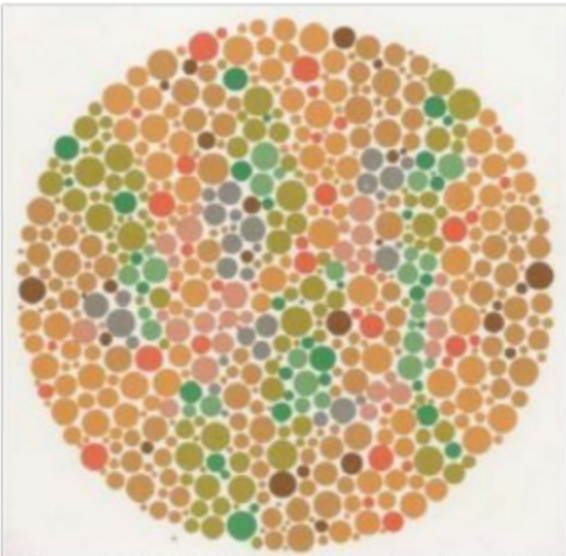


Figure 54: Test plate for colour blindness [7].

Why use color? Connect each choice to a task

Color is useful when its role is clear and alternatives remain available.

Role	Interface example
Label or distinguish	Identify related controls or categories
Show a variable	A scale with a legend and readable values
Use an analogy	Warm colors for heat; verify the convention
Support satisfaction	A coherent appearance without hiding task content

Color roles in familiar controls











Which device uses color to attract attention, distinguish controls, or communicate a value?

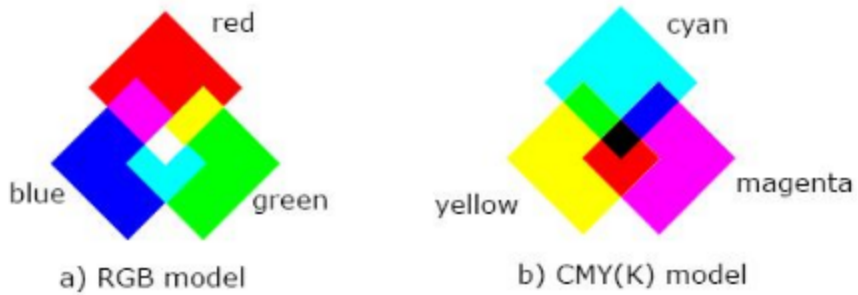
Worked interpretation: the emergency button emphasizes action; the meter communicates a range; the remote uses color to distinguish groups. Meaning still depends on labels and context.

Additive and subtractive color

RGB combines emitted light: display primaries add to form brighter mixtures.

CMY/CMYK controls reflected light through colorants. Black ink supports dark tones and practical reproduction.

A display preview cannot exactly predict every printed color.



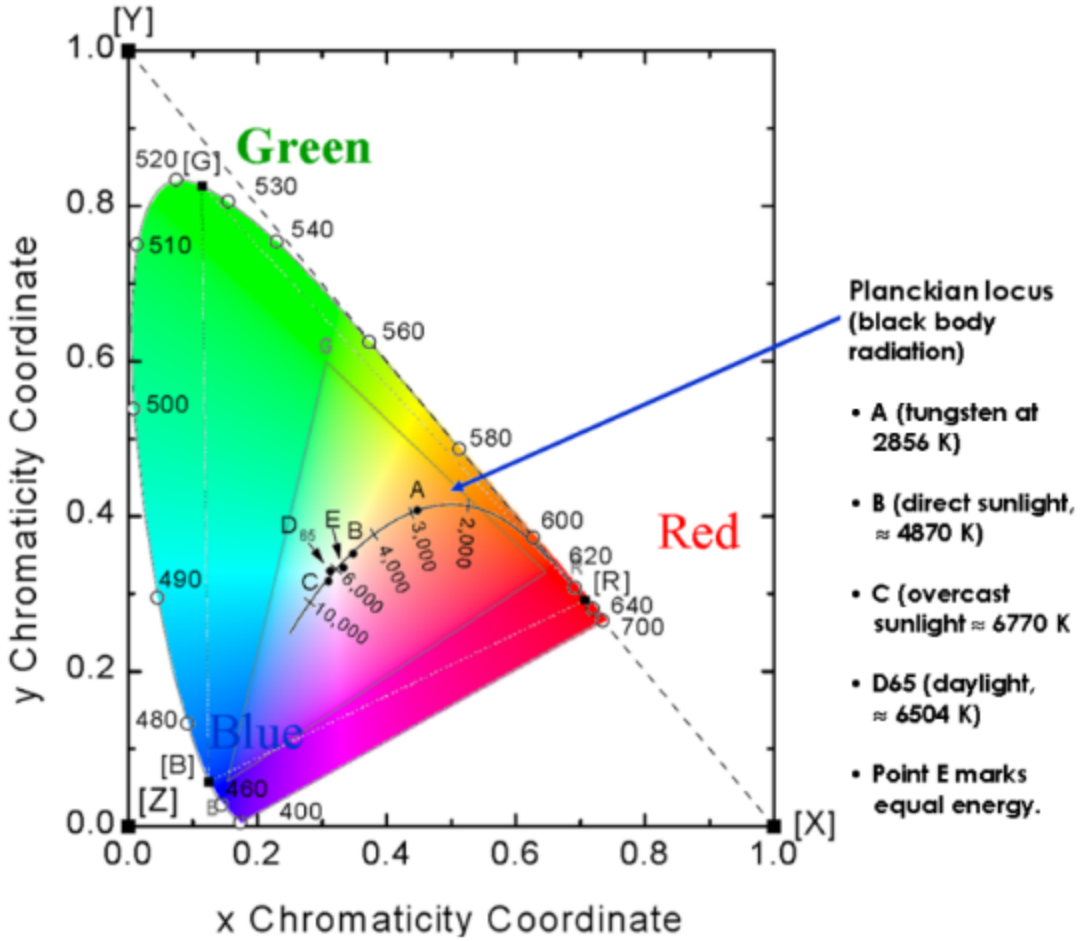
a) RGB model

b) CMY(K) model

Different color models answer different questions

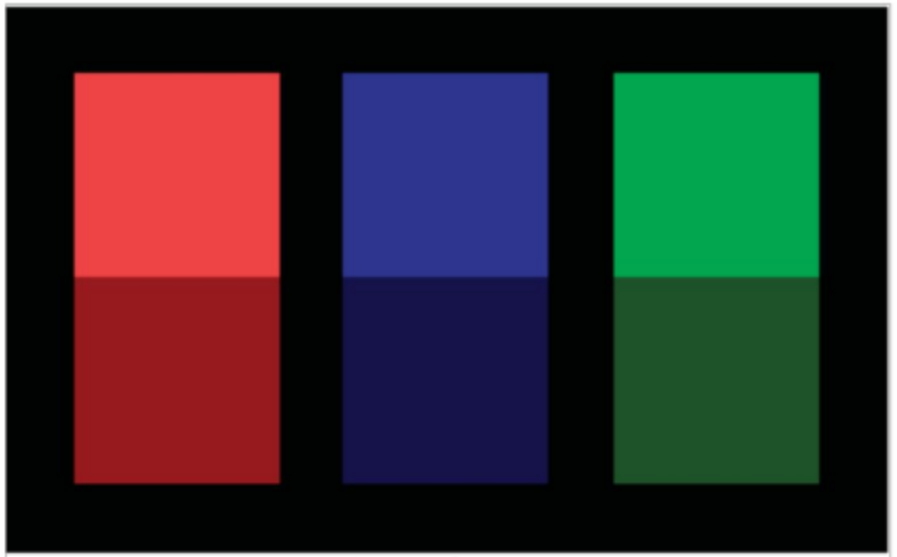
CIE coordinates describe colorimetry; RGB describes primary mixtures for a specified system. HSV supports hue-oriented selection; YUV-type models separate luma from chroma.

They are different coordinate systems and engineering conventions. Numeric values require a defined model and color space to be meaningful.



Hue and lightness are different contrast dimensions





Left: strong hue differences between red, yellow and blue. This traditional art palette is not the same as additive RGB display primaries.

Right: lighter and darker versions of a hue vary in lightness as well as appearance. Use the dimension that fits the information being encoded.

Warm–cool and complementary relationships





Figure 49: Complementary colour wheel

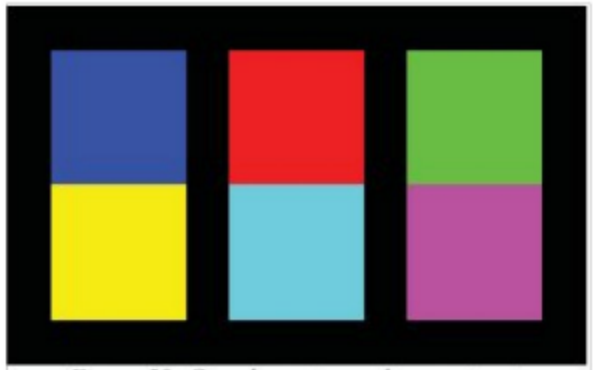


Figure 50: Complementary colour contrasts

Warm–cool comparisons use learned and perceptual associations across a palette. They can establish hierarchy or mood, but do not replace readable contrast.

Complementary relationships depend on the chosen color model. A color wheel organizes palette choices; it is not a universal perceptual equation.

Saturation contrast: vivid versus muted



The original “contrast of quality” example varies how vivid or muted a color appears.

A saturated accent among muted elements can draw attention. If everything is saturated, the accent loses its priority.

Check both saturation and luminance: a vivid color can still be hard to read.

Area changes the balance of a palette



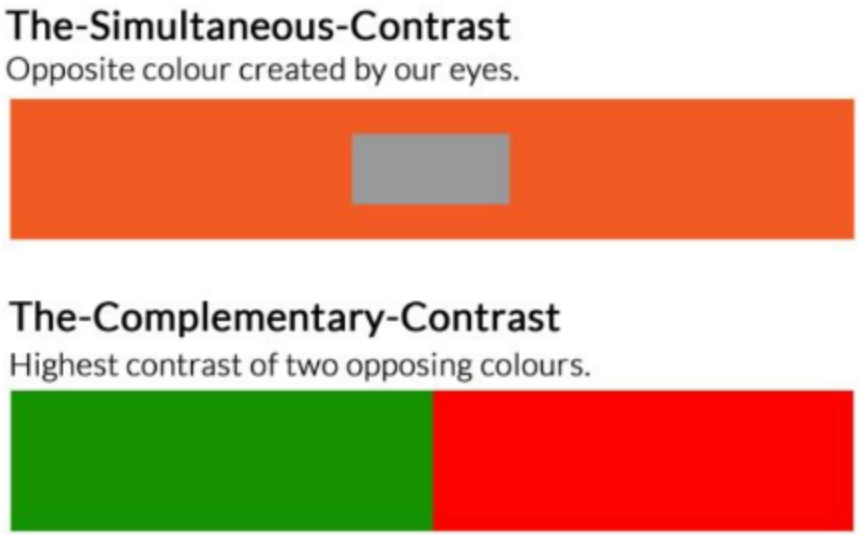
A small patch of one color can stand out against a large field of another. Color balance depends on area as well as the colors themselves.

Historical Goethe-style scores: yellow 9, orange 8, red 6, violet 3, blue 4, green 6. These are artistic heuristics, not measured visual laws.

If score × area is balanced, 6Ared = 3Aviolet: violet needs twice the area of red.

Several contrasts can act at the same time





Hue, warm–cool and light–dark differences can reinforce or compete with one another.

Simultaneous contrast: a patch's appearance is influenced by its surround. Complementary pairings can appear especially strong; the actual text/background contrast still needs checking.

Split-complementary palettes

A split-complementary scheme combines a base hue with hues on either side of its complement.

This can create variety while retaining a relationship between accents. The wheel is a palette-construction aid.

For interface categories, test whether users can discriminate and name the resulting colors under actual viewing conditions.



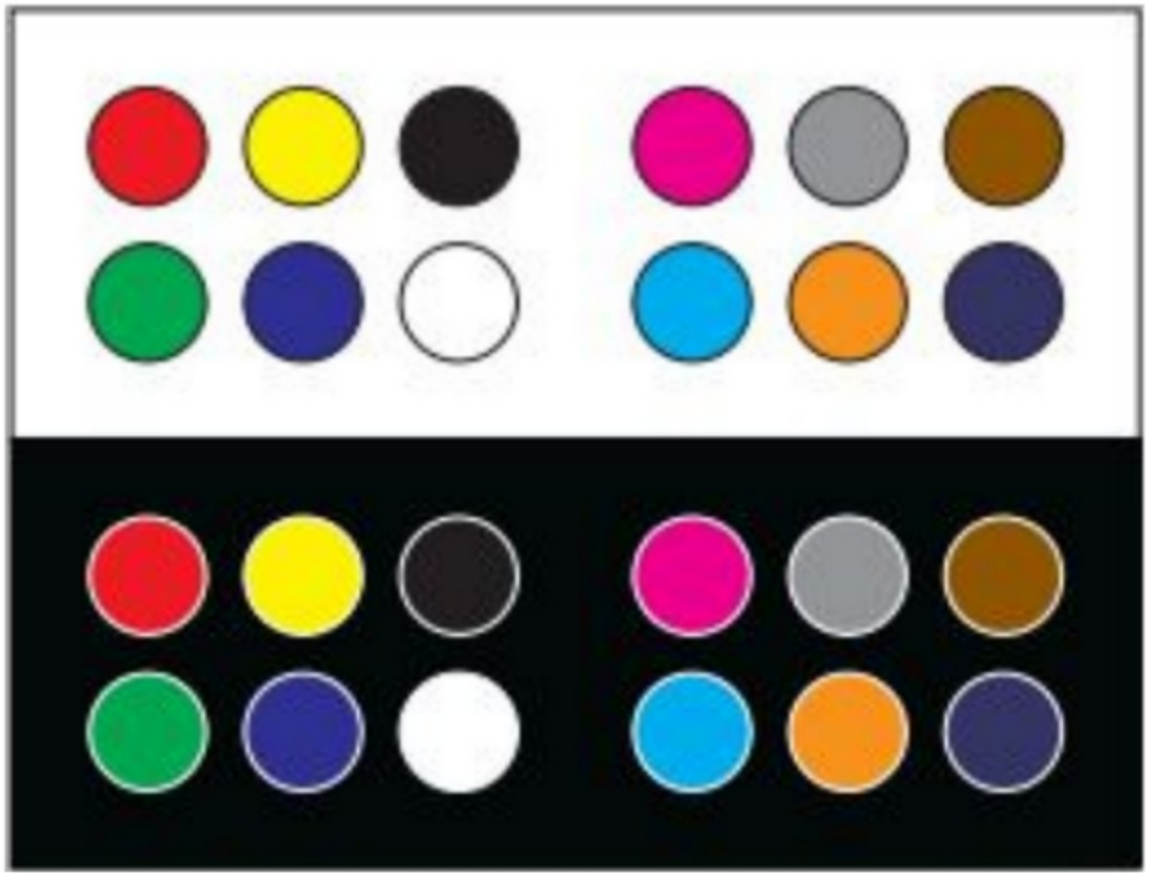
One hue plus two hues adjacent to its complement.



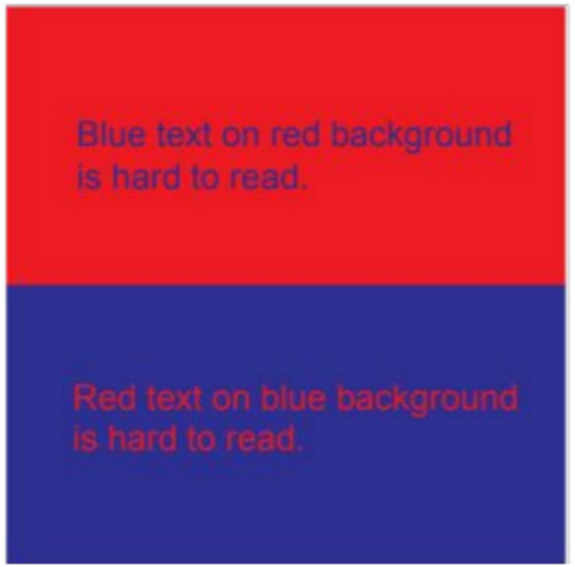
Color labels depend on language and background

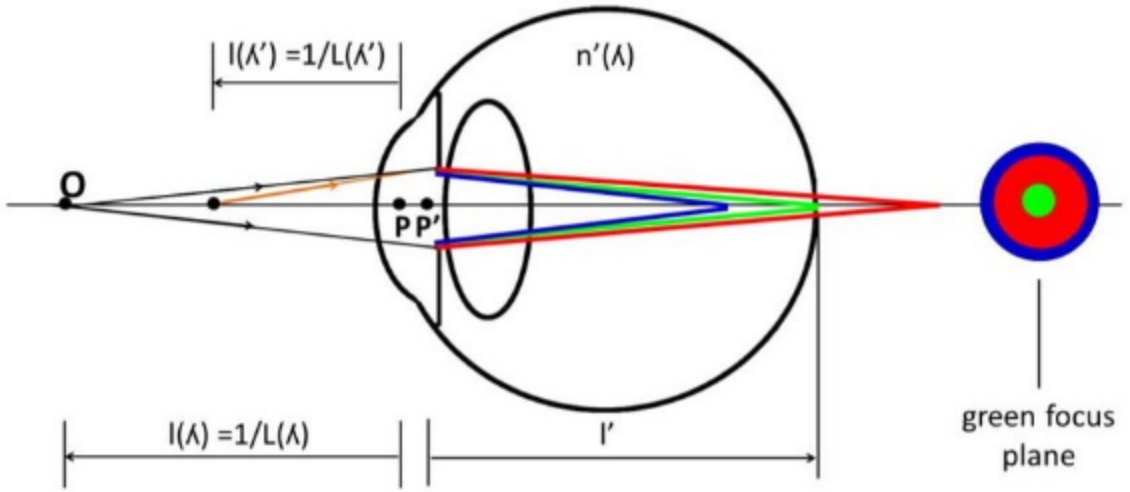
Common labels include red, green, yellow, blue, black, white, pink, cyan, gray, orange, brown and purple. Boundaries between named categories are not identical for every observer.

The same swatch changes in apparent contrast against white and black. Supply a legend when the color carries data, and avoid requiring a verbal name as the only identifier.



Red–blue text and chromatic aberration





The eye's optics focus different wavelengths at different positions. Strong red–blue pairings can produce difficult or unstable-looking edges.

The examples also have weak luminance contrast. Avoid using hue opposition as a substitute for text readability.

Red illumination and dark adaptation

Rods are less sensitive to long-wavelength red light. Dim red illumination can therefore help preserve dark-adapted sensitivity.

Brightness still matters: intense red illumination can disrupt dark adaptation. Red is not automatically safe at every level.

A dashboard must support legibility without flooding a dark viewing environment with light.



Reading depends on more than font size



Reading involves fixations and eye movements; layout guides where readers look next.

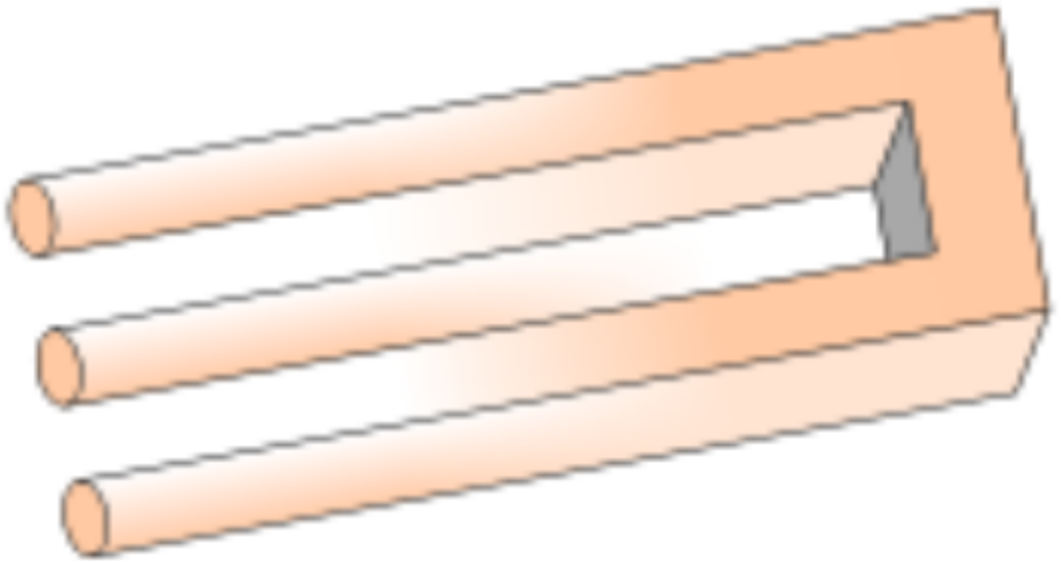
Legibility depends on visual angle, contrast, spacing and context. Use a clear line structure and distinguishable characters so that attention can follow the meaning.

Reading depends on more than letter recognition

Reading uses fixations and saccades; perception and interpretation are active processes.

Legibility depends on size, typeface, spacing, line length, contrast and layout, alongside the reader's experience and task.

The impossible trident shows how locally plausible contours can imply an inconsistent whole. Inspecting the parts does not guarantee a coherent interpretation.



Historical reading findings: keep the conditions

Earlier studies provide useful hypotheses to test on the actual medium.

Earlier finding	Design interpretation
Suggested line width: 60–130 mm	Physical width must be related to type size and distance
All capitals: about 14% slower	Favor familiar word forms for continuous prose
Uppercase for isolated emphasis	Short labels differ from paragraphs
Screen reading: 15–30% slower	A historical device comparison, not a timeless screen penalty

Gestalt laws: how we organize what we see

Gestalt means form or configuration. We tend to perceive organized wholes and relationships, rather than an unrelated collection of visual elements.

Why it matters in HCI: grouping tells users what belongs together, which label describes which control, and where one object ends and another begins.

1. **Proximity**
2. **Similarity**
3. **Continuity**
4. **Symmetry**
5. **Closure**
6. **Relative size**
7. **Figure–ground**

Proximity creates groups

Describe the groups you see before counting the dots. The elements are similar, but their spacing changes the organization.

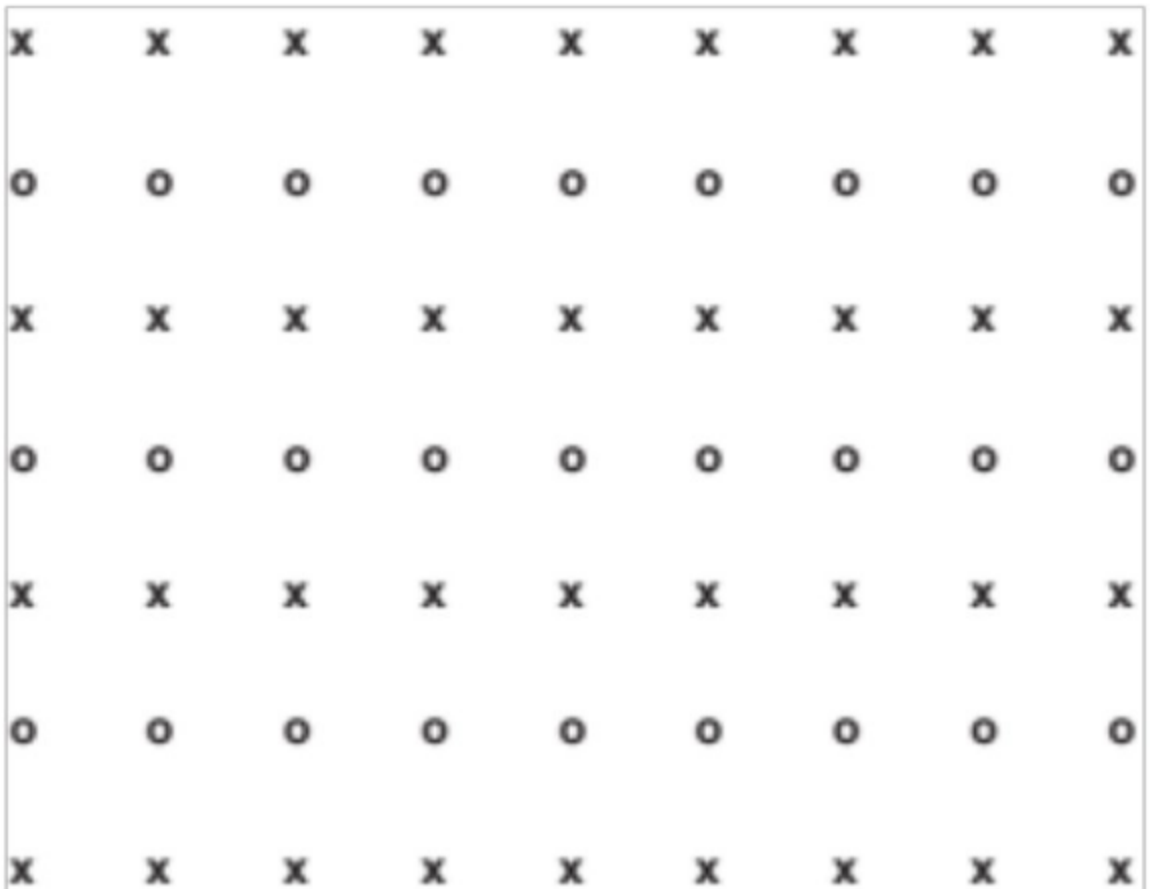


In a form, a label placed closer to the wrong input can imply the wrong relationship even when both labels are readable.

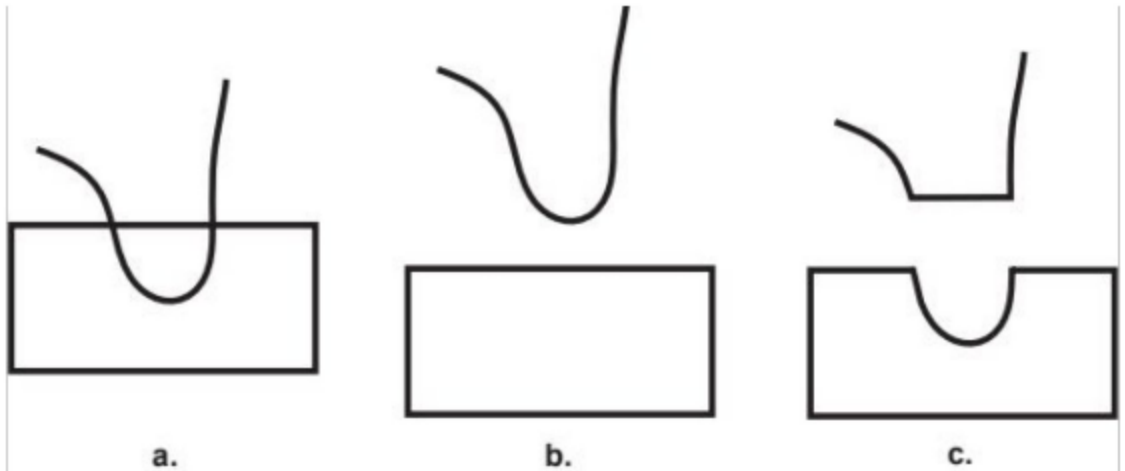
Similarity and common region can compete

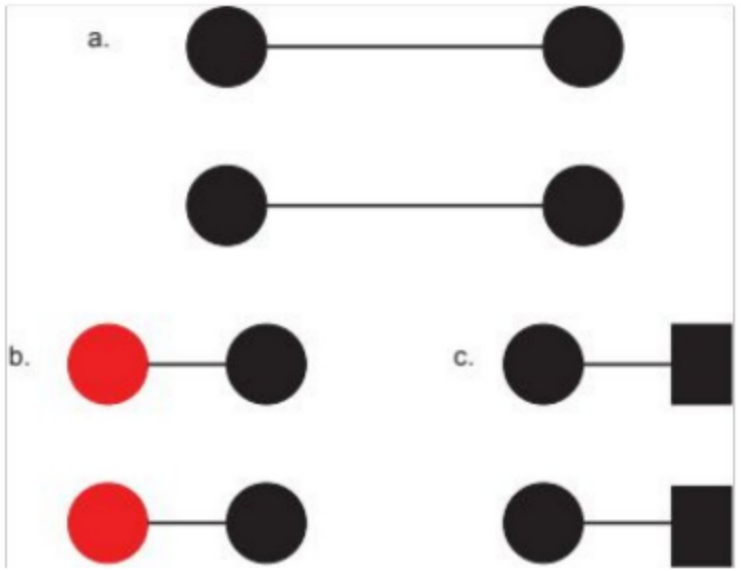
Similar appearance suggests membership. Enclosure suggests a common region. Those cues can reinforce each other—or conflict.

For repeated records, make the visual grouping agree with the logical grouping. Do not let alternating bands connect the wrong label and value.



Continuity guides tracing

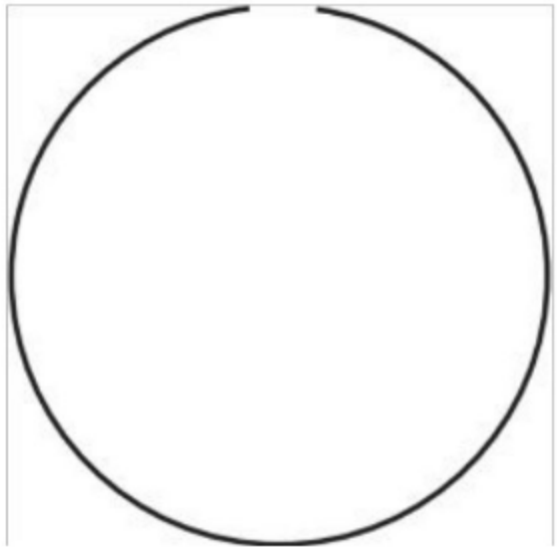


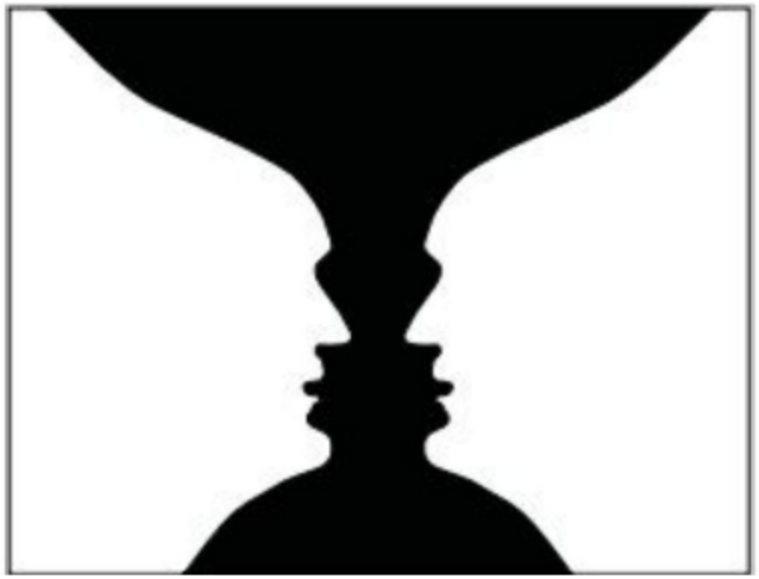


We tend to follow a smooth continuation through a crossing. A line is interpreted as a path, not just a set of pixels.

In diagrams, preserve an intelligible route. A visually smooth path that connects the wrong objects can create a convincing error.

Closure and figure–ground organize incomplete input

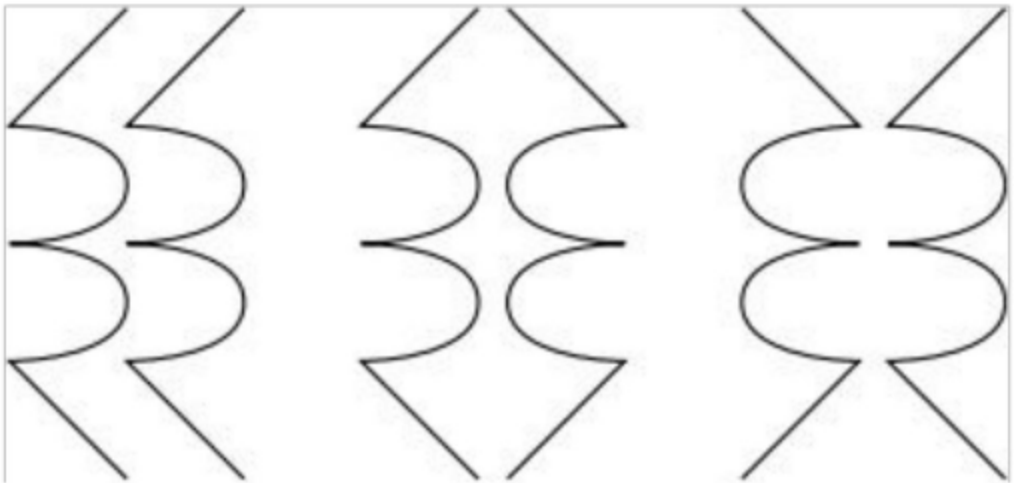




A contour can be perceived as a complete figure despite missing segments. A shared boundary can also support two competing foreground interpretations.

These effects explain why outlines, backgrounds and negative space can change what an interface seems to contain.

Symmetry and relative size guide grouping





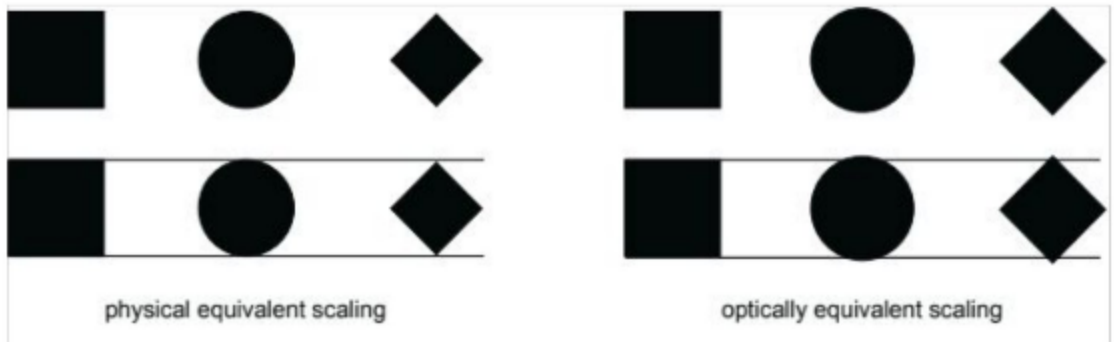
Symmetry can make contours belong to a coherent object. In the left example, the preferred grouping changes how we interpret the lines.

Smaller enclosed regions often appear as figures against a larger ground. On the right, relative area and contour organization influence which sectors seem to be objects.

Optical alignment differs from measured alignment



Figure 42: To align curvilinear elements properly with the baseline they have to project farther beyond the margin than linear elements (because of the human vision).



Equal coordinates do not always produce equal perceived alignment. Shape and visual weight change the apparent edge or center.

Compare the letter forms and symbols. Optical adjustment may be appropriate when geometrical alignment looks uneven.

A grouping error in a form

Which entries form one record? Here, the shaded band joins an amount with the next traveller, conflicting with the order of the fields.

Discuss the grouping before proposing a repair. Adding more color will not help if it reinforces the wrong grouping.

Field	Value
Traveller	Aylin
Amount	€60
Traveller	Bora
Amount	€90

The grouping repair and its tradeoff

Make each record a stable unit, with the label and value aligned within the same row.

The repair clarifies the relationship without requiring the user to remember a color code. Keep spacing consistent; avoid competing borders and unrelated emphasis.

Traveller	Amount due
Aylin	€60
Bora	€90

What perception explains—and what it does not

The item is visible, readable and grouped correctly. Is that enough?

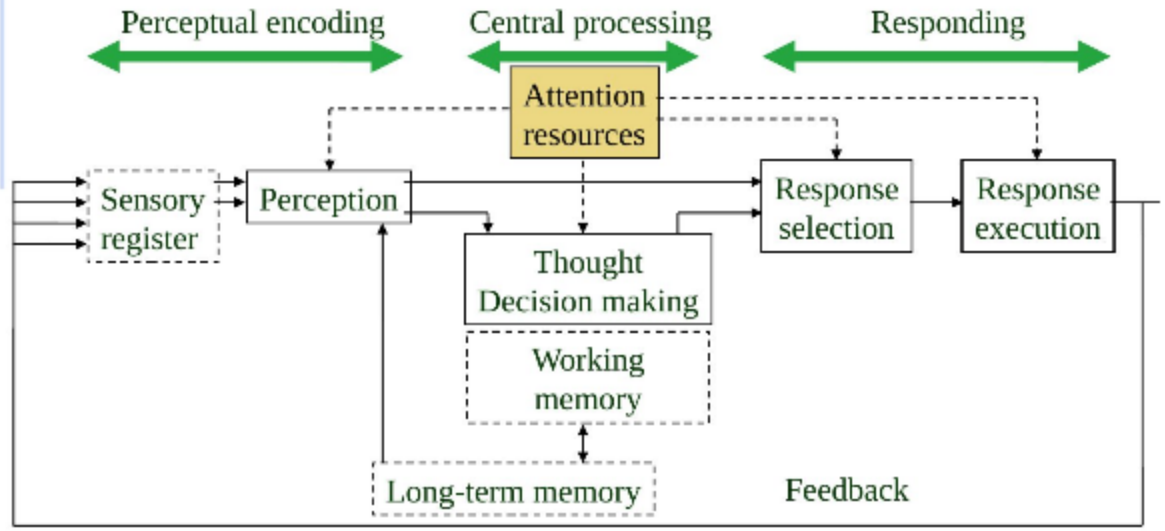
A person occupied with another task may still miss it. Perception describes what can be resolved; attention concerns what receives processing priority.

Question	Possible difficulty
Can I resolve it?	Small or low-contrast label
Can I organize it?	Ambiguous grouping
Do I attend to it?	Competing task or signal

Visibility does not guarantee attention

Attention is selective. The task, expectations and competing signals influence which information receives processing.

A prominent item may attract attention, but prominence is not a guarantee. A quiet cue can be effective when it appears where the task directs the user.



Wickens et al., 1998

A search task changes what stands out

Find the only Q. Then find every letter A. Notice how changing the goal changes your search.

An interface's visual hierarchy should be assessed against the user's task. This demonstration is not a timed performance test.

A O R P A O R P
P R O A P R O A
O A P R O Q P R
R P A O R P A O

Visual hierarchy should reflect the task

The user must verify the payment amount. Making the cardholder's name the dominant element competes with that task.

The repair gives the amount and payment action clear priority. Equal-looking "Cancel" actions should also behave consistently.

Element	Distracting emphasis	Task-relevant emphasis
Cardholder	Largest text	Secondary detail
Amount	Small, buried	Primary information
Actions	Similar prominence	Clear primary and cancel

Illustrative payment comparison.

A change can go unnoticed

Inspect view A, then cover it and inspect view B. What changed? Use a hand or paper if needed.

A stable side-by-side comparison supplies an external reference. A brief interruption forces more of the comparison into memory.

View A	Booking
Flight	GT 402
Seat	14A
Gate	16

View B	Booking
Flight	GT 402
Seat	14A
Gate	18

Worked answer: the gate changes from 16 to 18. The other fields remain the same.

Stroop interference: a color can compete with a word

YELLOW RED BLUE GREEN BLUE RED

BROWN GREEN ORANGE SARI GREEN BROWN

RED YELLOW ORANGE BLUE RED BLUE YELLOW RED



Name the ink colors rather than reading the words. When the two conflict, the automatic reading response competes with the requested task.

The game example makes the same conflict explicit: an easy-looking instruction can require inhibition and careful attention.

Stroop: a practiced response interferes

Name the ink color rather than reading the word. A practiced reading response can interfere with the instructed task.

Automatic responses can compete with deliberate goals. When an interface requires users to override a familiar meaning, it adds interference to an otherwise simple choice.

RED	BLUE	GREEN
PURPLE	BLUE	RED

Interruption leaves a resumption problem

You are entering an order when a message interrupts you. On return, remembering the goal is different from finding the last completed step.

Persist progress and entered values. A clear resumption cue can reduce reconstruction and repeated work.

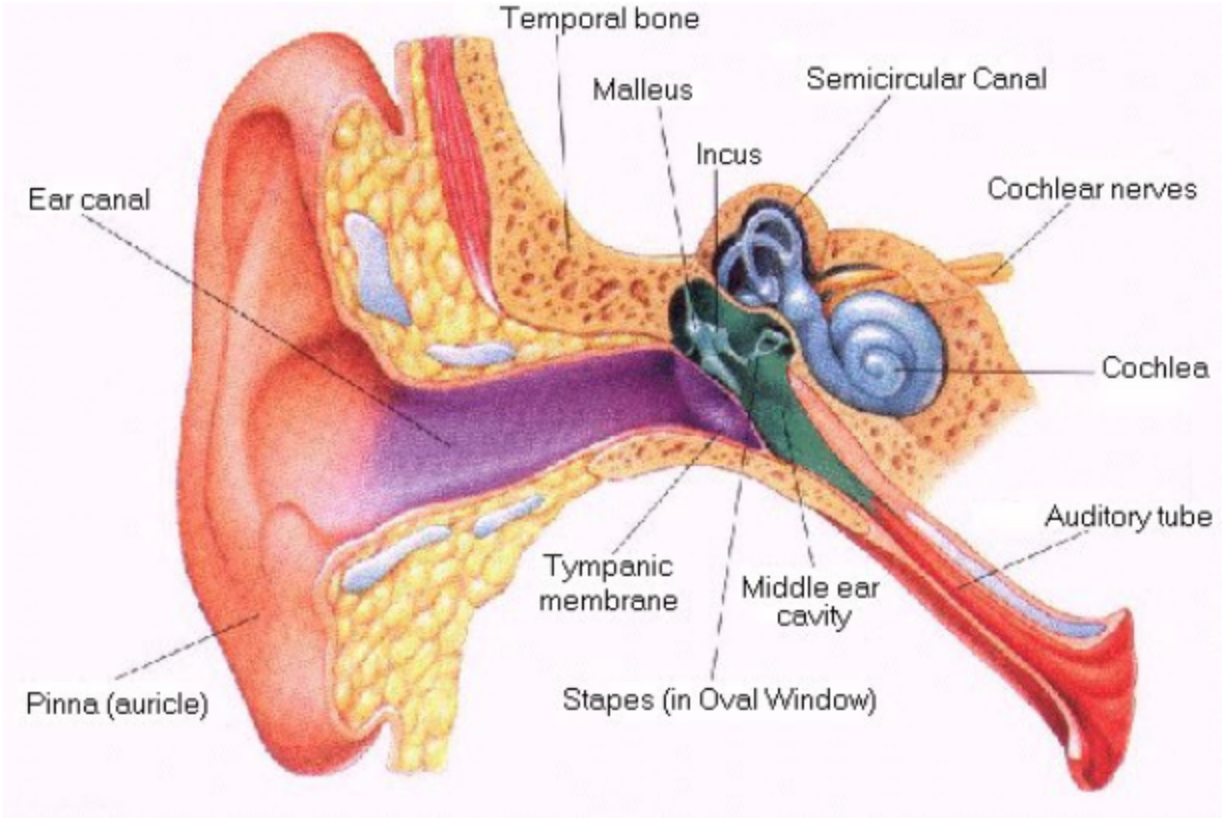
Before interruption	After interruption
Address entered	Still visible
Delivery option chosen	Selection preserved
Payment not confirmed	Next action: review and pay

The ear converts vibration into neural signals

The outer ear and canal guide sound to the eardrum. The middle-ear bones transmit vibration toward the cochlea.

Movement within the cochlea stimulates sensory hair cells; auditory-nerve signals carry information onward.

Pitch, timing and intensity provide different cues. Sound can convey an event even when the user is looking elsewhere.



Ear properties and auditory interface design

Hearing varies between people and with age, exposure and environment.

Property	Implication for interaction
Range: roughly 20 Hz–20 kHz in young listeners	Do not rely on high pitches being audible to everyone
Frequency-dependent sensitivity	Equal sound pressure does not mean equal loudness
Selective attention / “cocktail party” effect	People can attend to a stream, but masking still occurs
Sound pressure is expressed logarithmically	Level changes are not linear loudness steps
Beeping, errors and status tones	Use meaningful, distinguishable cues; avoid annoyance

Use sound where it helps the task

An auditory cue can reduce visual demand or reinforce a visible status change. The user should still have another way to obtain essential information.

Design for the environment: background noise can mask an alert, while a quiet room can make repeated beeps intrusive.

Choose pitch, duration and rhythm deliberately. A louder or higher-pitched alarm is not automatically more informative.

Sound and touch can support attention









Sound can announce an event when eyes are occupied. Touch can distinguish controls without a fresh visual search.

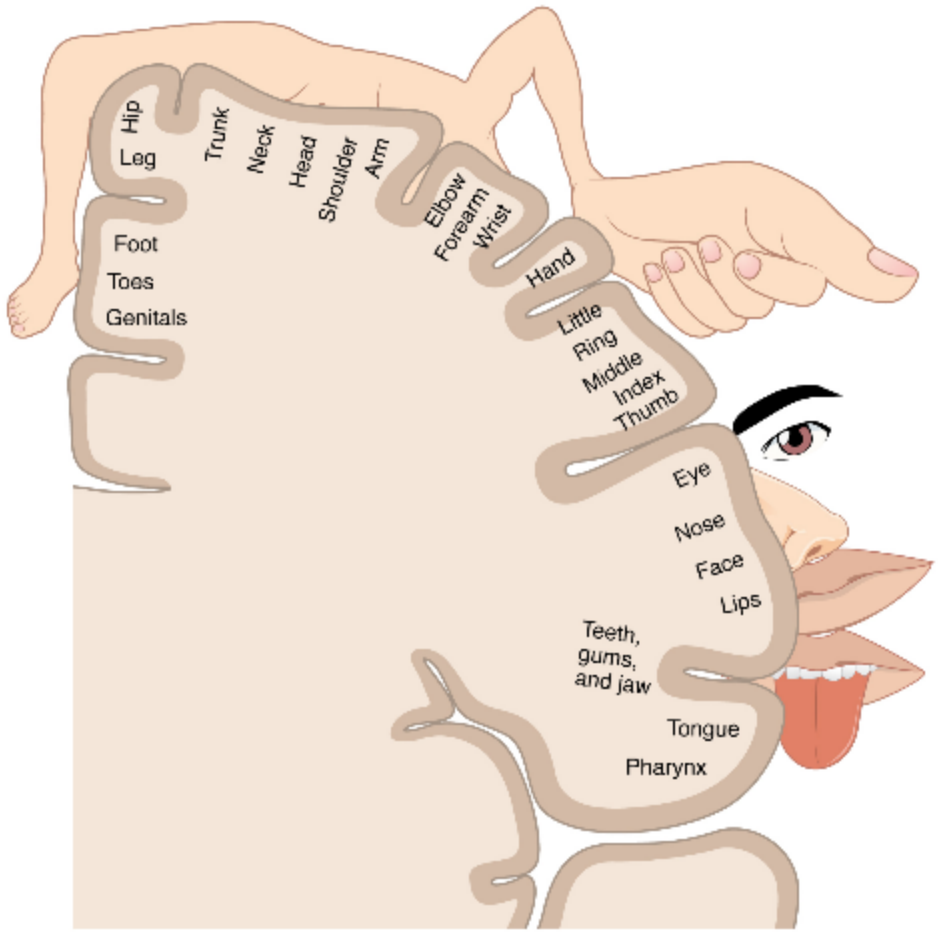
The crossing device and microwave examples communicate function. Repeated beeps can impose costs on everyone nearby; adding a channel is not automatically an improvement.

Somatosensory representation: the sensory homunculus

Somatosensation includes touch, pressure, vibration, body position, temperature and pain. The picture represents the body map in primary somatosensory cortex.

Hands and lips occupy disproportionately large cortical regions. Map size reflects sensory representation, not physical body size.

HCI implication: fingertip feedback can support precise interaction, but a vibration is not equally discriminable at every body location.



OpenStax College · CC BY 3.0

Tactile switches and haptic feedback

Haptic feedback engages our sense of touch, including skin sensations and forces or movement. Tactile feedback is the skin-sensation component.

A tactile switch produces a noticeable change in resistance, a bump or snap, as it is pressed. Many dome or membrane keypads provide this feedback.

A felt click confirms the press, not that the TV received a command. The system still needs feedback showing that the requested action actually happened.









The switch and keypad examples

An attention-aware redesign

For an interrupted order, a persistent “Payment not confirmed” message preserves the unresolved state. A transient beep alone cannot do that.

Use sound for a time-sensitive event, a visible cue for later recovery, and meaningful tactile differences where physical interaction permits them.

Signal	Helps with	Possible cost
Brief sound	Noticing an event	Distraction; missed meaning
Persistent status	Returning to the task	Screen space
Tactile difference	Identifying a control	Must be learnable

Memory supports ongoing action

Sensory persistence briefly retains input. Working memory maintains and manipulates information needed for the current task.

Long-term knowledge supports interpretation and skilled action. Attention, rehearsal and retrieval connect these functions.

Sensory persistence	Working memory	Long-term knowledge
Brief trace of input	Maintain and manipulate	Knowledge and skills

Attention → ↔ Retrieval and encoding

Rehearsal and external reminders can help sustain the information needed for action.

Holding information versus working with it

Hold the code 4729. Now compare it with 4792 and decide which digits changed position.

The second task requires both maintaining information and operating on it. Interfaces often ask users to do these two jobs at once.

Task	Demand
Repeat 4729	Maintain a short sequence
Compare 4729 and 4792	Maintain and manipulate
Show both codes together	Move part of the comparison into the display

Why one capacity number is misleading

“7 ± 2” is not a rule that menus should contain seven items. Capacity estimates depend on task, chunking, rehearsal and what counts as one item.

Do not replace that shortcut with a universal four-item rule. Ask what must be maintained during this particular task—and what can stay visible.

Chunking depends on knowledge

Consider 1 9 2 3 2 0 2 6 as isolated digits, then as 1923 / 2026. Does the grouping mean something to you?

Chunking draws on existing knowledge. A meaningful group for one user may remain several unfamiliar items for another.

Sequence	Possible grouping
1 9 2 3 2 0 2 6	Eight separate digits
1923 / 2026	Two meaningful dates—for some readers
Unfamiliar domain	The same groups may not be meaningful

Interfaces create temporary memory work

A server carries item names, quantities and modifications from a customer's request to the order record. Interruptions can disrupt that temporary information.



A paper note or handheld order device keeps the order visible. Consider how entering and checking it also competes with attention to the customer.





External state reduces the carrying task

Keep the source and destination together when comparison matters. Preserve entered values and show the current choice.

External state does not eliminate thinking. It changes the task from carrying a value in memory to finding and interpreting it on the display.

Recall-heavy task	Visible support
Remember order details	Visible items, quantities and changes
Remember prior entries	Persistent summary
Remember current step	Clear progress state

Long-term knowledge changes interpretation

Prior knowledge supplies categories, conventions and procedures. What does this mark suggest to someone who reads Chinese, and to someone who does not?

The same shape can evoke different interpretations. Learning helps, but familiarity is not shared equally by all users. Test the intended audience.



Recognition and recall impose different demands

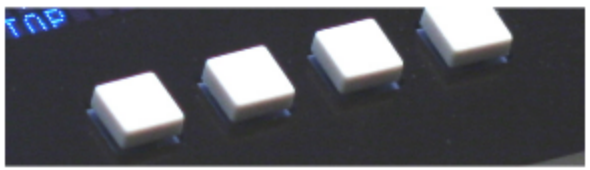
Recall requires producing information, such as a command name. Recognition provides candidates, such as items on a menu.





Visible controls help only if their choices are recognizable. An expert's remembered shortcut may be faster than searching a large menu.





Practice can make a sequence fluent

Practice can make a sequence fluent. The same habit can cause an error when a familiar control changes meaning or position.

Consistency helps transfer, but users still need feedback when the state changes. A well-practiced action is not necessarily the right action in a new context.



Interference and prospective memory

Remembering information and remembering to act later are different problems. A person can know what to do and still miss the appropriate moment.

A saved value supports retention. A persistent unresolved-state cue supports returning to the intended action.

What must be remembered?	Appropriate support
The order code	Keep it visible
Where I stopped	Show completed steps
That I must submit	Show unresolved state at the relevant moment

Memory demands in a multi-step form

A booking form hides the fare while the user enters passenger details. It then asks for the total, a baggage choice and a final confirmation.

Identify what must be held, manipulated, retrieved from knowledge and remembered for later. Propose a change for each demand.

Task detail	Demand to diagnose
Fare on the previous screen	Hold a value
Fare × number of passengers	Manipulate values
Baggage abbreviation	Retrieve prior knowledge
Final confirmation after interruption	Remember an intention

A worked memory-demand reduction

Keep the fare and calculated total visible, explain baggage terms, preserve entries and show whether submission is complete.

Each change answers a specific demand. The tradeoff is extra screen space and possible search effort, so grouping and hierarchy still matter.

Demand	Worked redesign
Hold fare	Persistent fare summary
Calculate total	Show calculation and total
Interpret abbreviation	Plain label with explanation
Remember confirmation	Visible unsubmitted state

A mental model explains and predicts

A mental model is a learned explanation that lets a user predict what a system will do. It can be incomplete, inconsistent or wrong.

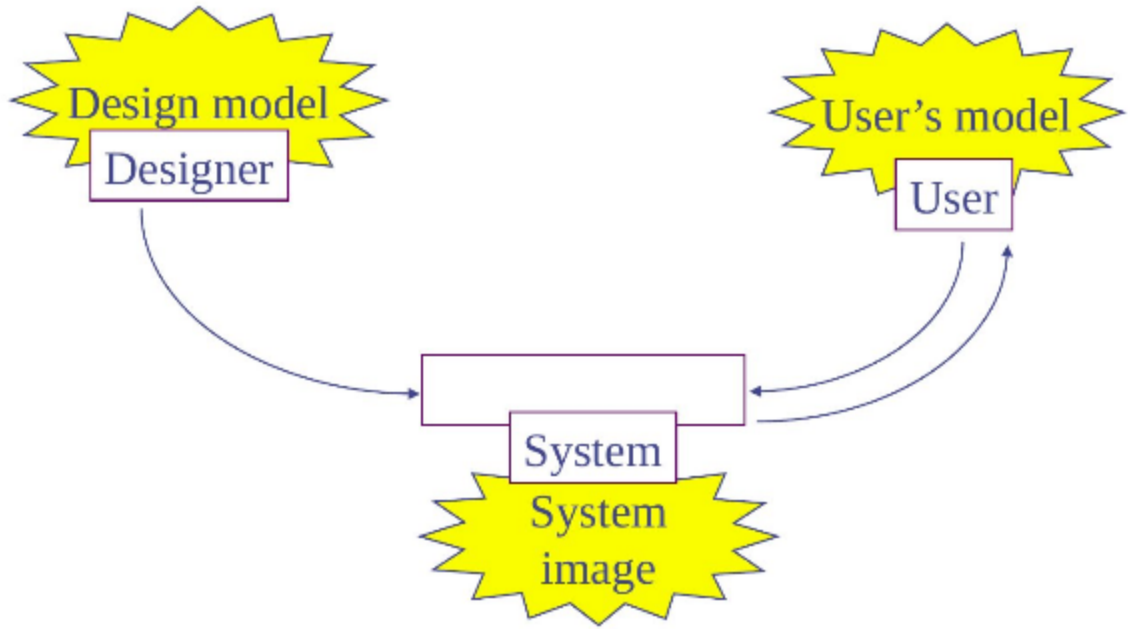
Models draw on experience and long-term knowledge; they are not a separate memory store. Ask what the user expects to happen next.



Designer model, system image, user model

The designer has an intended explanation. The user constructs an explanation from the system image: controls, behavior, feedback and documentation.

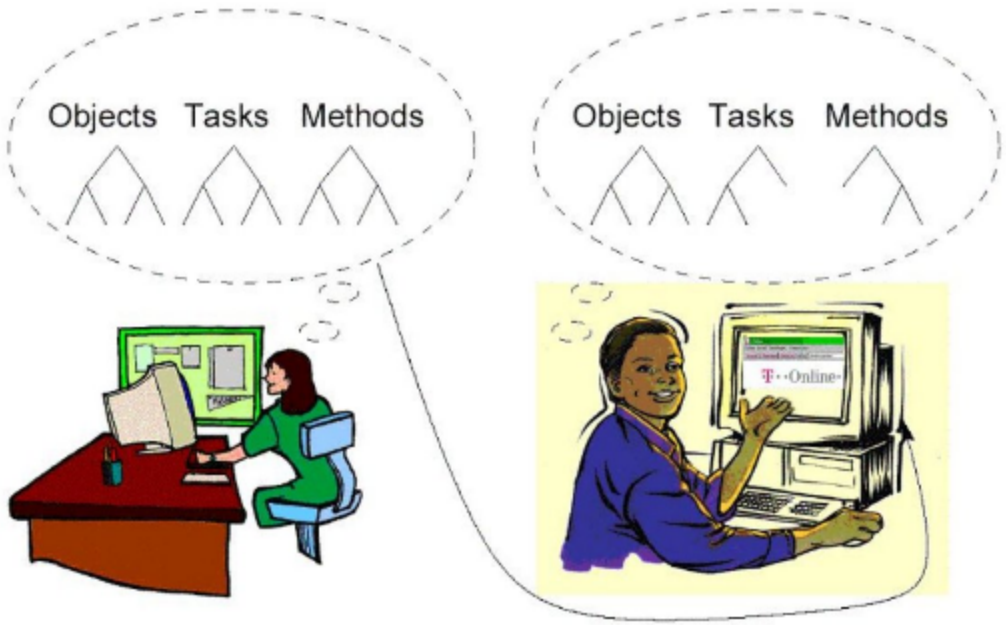
The user cannot inspect the designer's intentions directly. A convincing but misleading system image can teach the wrong model.



The system image teaches a model

Two separate-looking controls may suggest two independent effects. That suggestion can survive until feedback reveals a coupling.

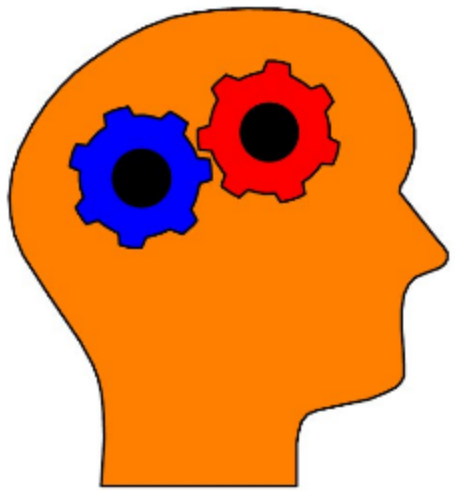
A design should make consequential relationships understandable. The next case asks us to distinguish what the interface suggests from how the mechanism could work.



The refrigerator goal

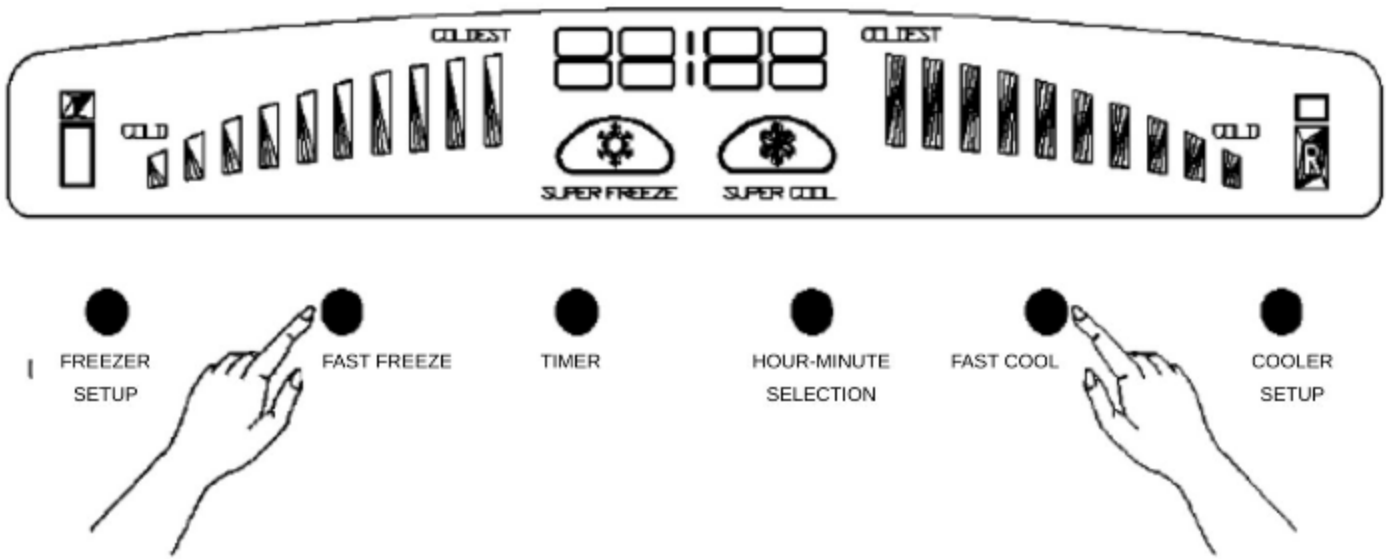
The freezer is too cold, but the fresh-food compartment is satisfactory. Which control would you change first?

State your predicted result before discussing the mechanism. Different causal models can produce different, reasonable-looking actions.





Vestel: what does the panel suggest?



Which controls appear to belong to each compartment? What do the displays and the “fast” functions imply?

The panel suggests relationships, but its appearance alone does not establish the cooling hardware.

Beko: two representations of control



BEKO 8290 Refrigerator

(from operators manual)



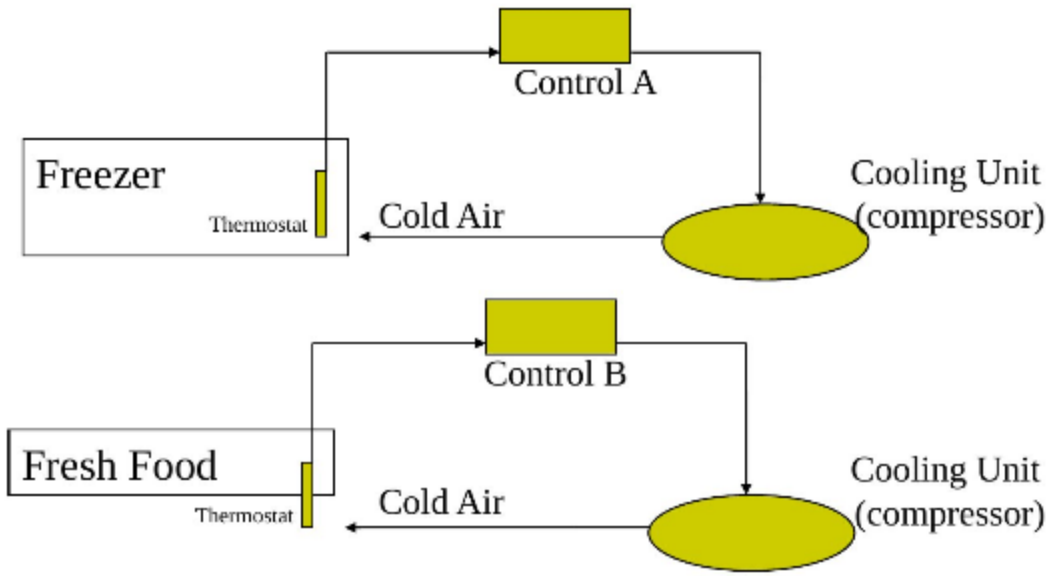
The fresh-food panel and freezer control use different representations. Compare the displayed temperature with the ordinal freezer settings.

What does increasing a setting mean, and what would you expect to happen in the other compartment?

Model A: independent cooling controls

Under this candidate model, each compartment has its own cooling loop. Changing control A need not change the setting of control B.

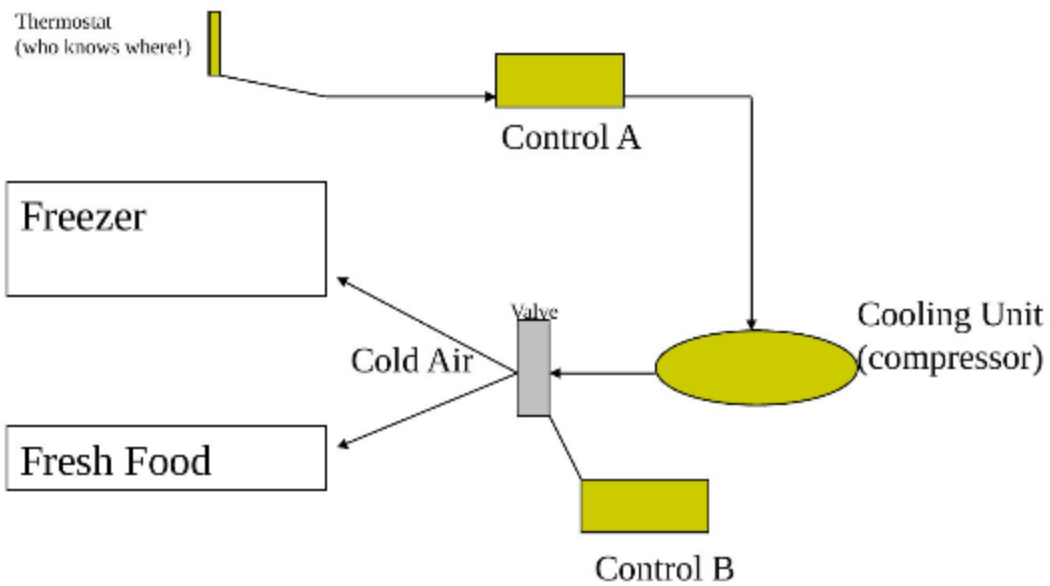
This diagram explains an independence hypothesis. It is not a verified specification of either named refrigerator.



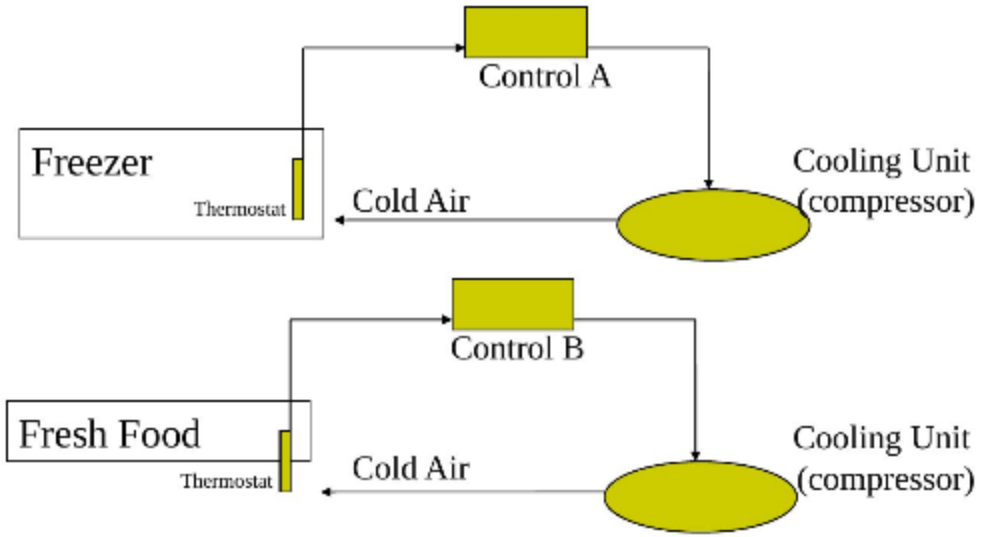
Model B: shared cooling and air distribution

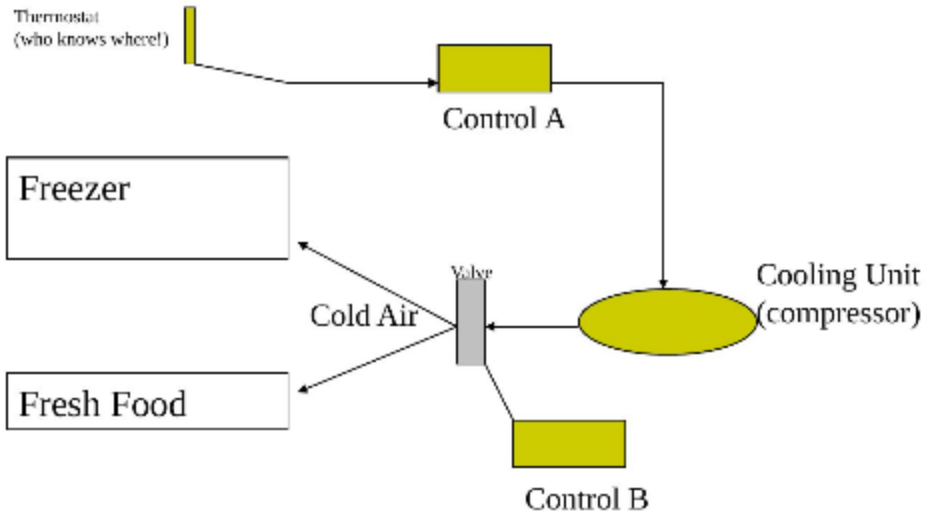
Under this alternative, one cooling unit serves both compartments. One control can affect cooling operation while another changes air distribution.

A change aimed at one compartment can therefore affect the other. Predictions depend on the actual coupling and feedback.



Two models make different predictions





For independent loops, “leave the fresh-food side unchanged” is a plausible prediction. For a shared system, the same action may alter both temperatures.

The important question is which causal relationship the controls communicate—and whether observed behavior supports it.

Klemmer’s refrigerator reveals hidden coupling

In Klemmer’s separate example, one dial changes overall cooling while another divides the cold air between compartments.

Two dials can invite a mistaken “one dial per temperature” model. Feedback and an explanation of the coupling are needed to predict the result.

Control function	Consequence in this example
Overall cooling	Affects available cooling
Air distribution	Changes how cooling is shared
Two independent temperature dials	An interpretation the appearance may invite

Prediction questions for the original case

Can the fresh-food section be switched off while the freezer remains cold? Can the freezer be used as fresh-food storage?

Independent loops and shared cooling support different predictions. Settings alone are insufficient evidence; operating documentation and measured behavior would decide.

Question	What must be established?
Turn one compartment off?	Independent operation or shared cooling?
Use freezer as fresh-food space?	Supported temperature range and control behavior?

A clearer control explanation

A clearer explanation names the effect of each control and makes any coupled consequence visible. It should support the user's goal, not demand a guess about hardware.

For this fictional shared-cooling design, the interface explicitly warns that changing distribution affects both compartments.

Control label	Explanation shown with it
Overall cooling	Changes cooling supplied to both spaces
Air distribution	Balances cold air between compartments
Before applying	Both temperatures may change

Illustrative redesign; not operating instructions for a named appliance.

Slip or mistaken understanding?

A slip: the user intends to press the freezer control but accidentally presses the adjacent control.

A mistaken understanding: the user deliberately adjusts a dial believing it affects only one compartment, although the system is coupled. Prevention must address the actual cause.

Failure	Possible response
Unintended press	Spacing, differentiation, undo
Wrong causal prediction	Clearer mapping, feedback and explanation

Representation exercise: make fifteen

Two players alternately choose an unused number from 1–9. The first player holding any three numbers that sum to 15 wins.

Try a short round. You must track your choices, the opponent's choices and available winning triples.

1 2 3 4 5 6 7 8 9

Example winning set: 4 + 5 + 6 = 15

The same game in a different representation

The same numbers can be arranged as a magic square. Each winning triple is a row, column or diagonal.

8 + 1 + 6 = 15. Also, 8 + 5 + 2 = 15. The board exposes threats that were harder to see in a list.

8	1	6
3	5	7
4	9	2

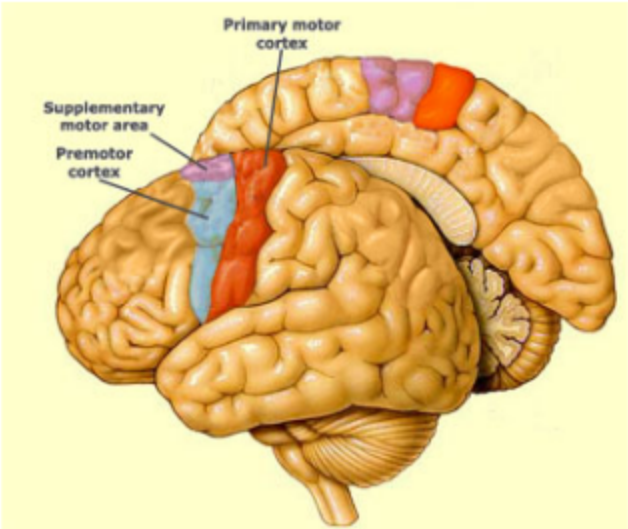
What representation changes

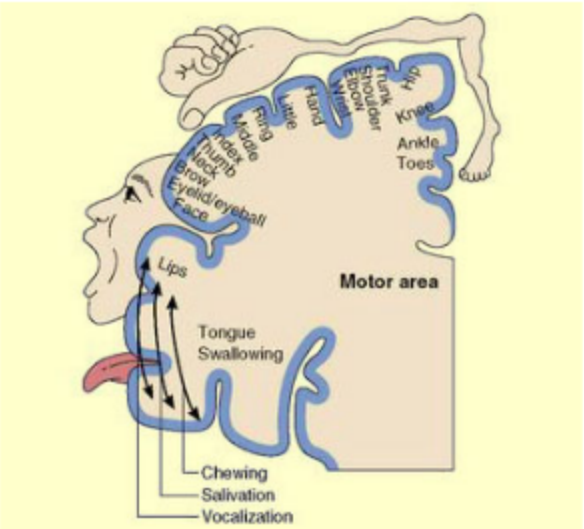
The rules did not change. The representation made relationships visible and reduced the bookkeeping needed to find a move.

The same principle applies to controls and forms: external structure can support reasoning, but users still need to understand what the relationships mean.

Representation	What becomes easier to see
Number list	Which numbers remain
Magic-square board	Winning lines and threats
Explained control relationship	Consequences of an adjustment

Sensory and motor maps have different roles



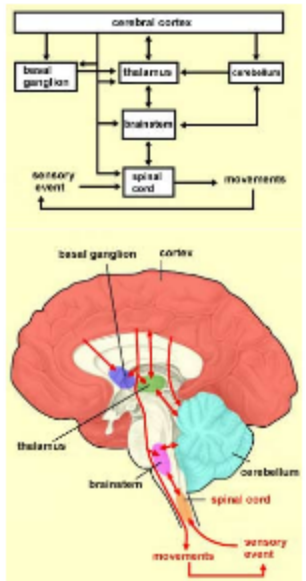


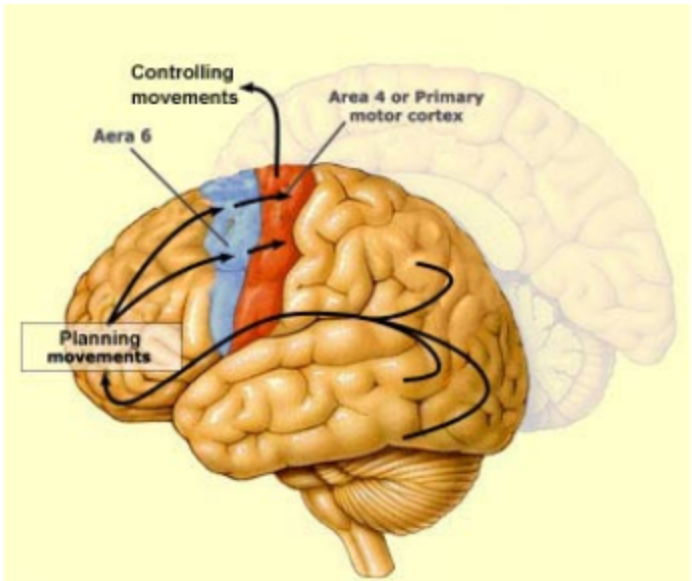
The motor homunculus represents cortical organization of movement; the sensory homunculus represents incoming somatic sensation.

Both devote substantial representation to the hands and face. The pictures should not be interchanged merely because their distorted body proportions look similar.

Interaction forms a loop: movement changes the world, and sensory feedback helps guide the next action.

Movement combines planning, control and feedback





Motor control involves interacting cortical and subcortical systems. Planning an action and correcting its execution are related but distinct functions.

Feedback about position and outcome supports corrections. This motivates the next question: how do target distance and size constrain pointing time?

Finding a target and reaching it are different

Before predicting a click, distinguish finding the target, deciding to use it, moving to it and waiting for system response.

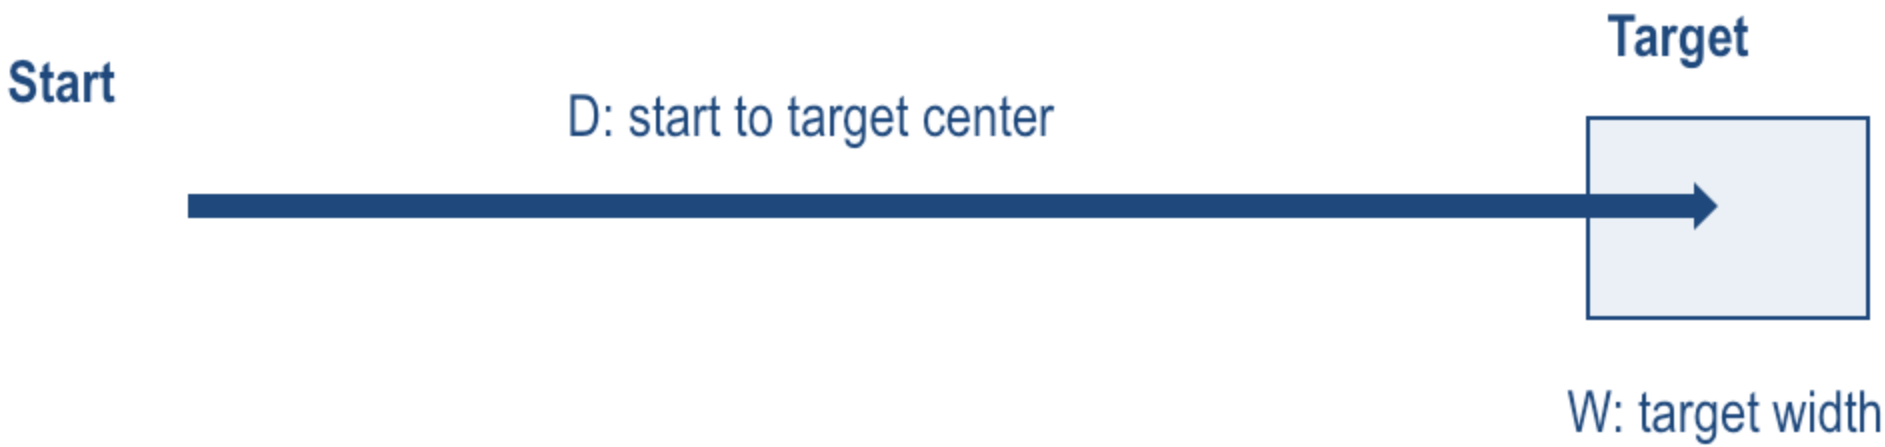
Fitts's law addresses movement and target acquisition under specified conditions. It does not by itself predict the duration of the whole task.

Stage	Example demand
Find	Locate the button
Decide	Interpret its action
Move/select	Acquire the target
System response	Wait for completion

Distance and target width

D is the movement amplitude. W is the target width measured along the approach direction. Use the same distance units for both.

The visible icon and the interactive hit area need not be identical. State which boundary defines success.



The target boundary defines which endpoints count as a successful selection.

Fitts's law: a predictive model

$$MT = a + b \log_2(1 + D/W)$$

ID = $\log_2(1 + D/W)$ is the index of difficulty. D/W is dimensionless; ID is expressed in bits.

a has time units and b has time-per-bit units. They are fitted coefficients for a task and device, not universal human constants.

Worked example: narrow target

Illustrative coefficients: a = 50 ms; b = 100 ms/bit. Target distance D = 400 px; width W = 20 px.

$ID = \log_2(1 + 400/20) = \log_2(21) \approx 4.392$ bits

$MT = 50 + 100 \times 4.392 \approx 489$ ms

Teaching values, not measured participant results.

A wider or nearer target

Doubling target width reduces the model's predicted time. Halving distance gives the same ratio in this example.

Compared with the baseline, either change predicts about 93 ms less movement time, using the same coefficients.

D (px)	W (px)	ID (bits)	MT (ms)
400	20	4.392	489
400	40	3.459	396
200	20	3.459	396

Illustrative model: a = 50 ms; b = 100 ms/bit.

A button-layout prediction

A third layout has D = 300 px and W = 30 px. Predict its movement time before revealing the calculation below.

$ID = \log_2(1 + 300/30) = \log_2(11) \approx 3.459$ bits. $MT \approx 396$ ms.

The equal D/W ratio gives the same model prediction. It does not prove equal visual-search, decision or total task time.

Screen edges and touch are different conditions

A pointer constrained by a screen boundary can benefit from reduced overshoot in that direction. This depends on the actual boundary and interaction.

A physical touch target has different constraints. Do not transfer a desktop edge argument to every touch interface or treat every edge as an infinite target.

The TiVo remote combines several mechanisms

Large, differentiated controls can help targeting and interpretation. Related controls can also suggest a useful mapping.

The remote's symmetry creates a tradeoff: it may be harder to tell its orientation by touch. Size alone does not explain the whole design.



Historical example: MIT 6.831 (2004)

Limits and a small measurement plan

Record both movement time and errors. Repeat trials, vary target geometry and balance order so practice is not mistaken for a layout effect.

Device, approach direction, motion, familiarity and error handling affect interpretation. A model prediction is a testable claim, not proof of usability.

Keep track of	Why it matters
Time + errors	Speed can be bought with inaccuracy
Order + practice	Later trials may be faster
Device + task	Coefficients may change

From prediction to observation

Compare two target layouts and calculate a prediction using clearly stated coefficients.

Explain the assumptions, propose how you would measure performance, and discuss a result that could contradict your prediction.

Proposed Assignment 1 direction: connect the calculation to an observable interaction task.

One interface, several explanations

An interrupted booking form can fail for several reasons. A useful diagnosis identifies the mechanism before proposing a remedy.
Some changes help one demand while increasing another. Keep the task and the user in the explanation.

Difficulty	Mechanism	Candidate response
Unreadable label	Perception	Contrast and size
Missed status	Attention	Persistent task-relevant cue
Forgotten fare	Memory	Visible summary
Wrong control effect	Mental model	Mapping and feedback
Missed small button	Motor demand	Larger hit area

Closing questions and further study

Explain one interface difficulty using a mechanism from this lecture. Make a prediction and state one limitation.

Further study: Scott Klemmer's Mental Models and Representation Matters lectures; the original examples and cited readings in the slide notes.

Next: understanding users, tasks and context—finding out whose knowledge and goals the design must support.