

Consciousness

PSY 200
Greg Francis
Lecture 32

Do you see red like I see green?


What is consciousness?

- Awareness of events, stimuli, thoughts, self
- A sequence of meaningful items
- Stream of thoughts
- Distinct from unconscious processing (e.g., hearing a sentence, retrieving information from memory,...)

History



- Descartes' dualism (Cartesian dualism)
 - pineal gland link between body and spirit
 - how they could connect was a real problem




- Mind-body problem
- Materialism (the brain *is* the mind)
 - or the mind derives from the brain

Materialism

- Nearly all scientists are materialists, but old ideas die hard
- A lot of work (PET scans, fMRI,...) looks for the *site* of consciousness
 - a special physical transformation
 - thalamus
 - reticular formation
 - quantum mechanics
 - distributed awareness

A turning point

- A common view is that there is a moment/ place which/where before something was *not* conscious and which after it *is* conscious
- But this is not true in the brain



Distributed processing

- Information processing is spatially and temporally *distributed* in the brain
- Processing changes with new stimuli
- There really is no "moment of consciousness"
 - different brain areas know different things at different times

An analogy

- When did the British empire learn of the end of the War of 1812?
 - treaty signed in London months before the Battle of New Orleans
 - word was not received by British troops in America until two weeks after the Battle of New Orleans (January 8, 1815)



British Empire in 1815

Purdue University

An analogy

- For complicated systems like the British empire (and human brains)
 - different parts know different things at different times
 - there is no official moment of knowledge
 - no official moment of consciousness!
- Demonstration
 - when does the class know/understand?



Purdue University


How / Why?

- There is no “moment” because information is *distributed* in the brain
 - Both in space and time
- Can distributed processing really produce consciousness, or must there be something else to “put it all together”?
 - can consciousness arise from non-conscious processors? (artificial intelligence?)

Purdue University

Turing test

- How do you know a person is conscious/intelligent?
 - They behave in a way that we interpret as consistent with a conscious being
- Turing test: apply the same logic to a computer
 - if a conversation with a computer is indistinguishable from a conversation with a human
 - Then conclude the computer is intelligent



Purdue University


Artificial intelligence

- No computer has passed anything but a weak form of the Turing test
 - lack sufficient schemas, creativity, general knowledge
- It is worth noting that other things would also not pass a Turing test
 - children
 - mentally impaired people
 - mute people
 - people who speak a language we do not understand
- Passing a Turing test is not *necessary* for consciousness

Purdue University

Captcha

- The basic ideas are implemented in several methods for computer security
- Completely Automated Public Turing test to tell Computers and Humans Apart



Purdue University

Turing test


- The Turing test is only one way to demonstrate intelligence
 - and a rather strict one at that
 - not passing the Turing test does not mean that a computer is not intelligent
 - of course, it doesn't mean the computer *is* intelligent either
- Variations on Turing test
 - discriminate conversation between a child and a computer
 - look at a conversation and decide which was the computer
 - ...

Doubters

- Many people have suggested that computers cannot, in principle, become intelligent
 - they argue that purely symbolic computations cannot lead to consciousness
 - and humans use emotion, insight, intuition, intentionality instead of simple computation
- Let's look at two arguments against "strong AI"

1. The Chinese room (Searle)

- Imagine you are in a room with two slots and a book
 - Slot 1: someone sends you notes with Chinese characters on them
 - Book (written in English): in the book you can look up the Chinese characters and write down corresponding Chinese characters on another piece of paper
 - Slot 2: you can send your piece of paper out this slot




1. The Chinese room (Searle)

- If the book provides rules on how to answer questions in Chinese
 - then you can answer written questions in Chinese
 - even though *you* do not know Chinese!
- Consciousness (in general, understanding) is not a function of the thing (or person) who *implements* the rules
- But consider it from the point of view of a person outside the room
 - Who is sending messages in

1. The Chinese room (Searle)

- You are having a conversation with someone
 - You have to decide if the person understands what you are saying (it's the Turing test)
 - You ask them to describe the wall of their room
 - They report it is green
 - ...and so on...




1. The Chinese room (Searle)

- Searle's point is that
 - We know the person in the room does not understand Chinese
 - We might be fooled into thinking they do based on their responses to the questions
 - Thus, the Turing test is a bad test
- Because the Turing test is essentially the same structure
 - The computer plays the role of the person in the room


However,...


- Searle has set up a deceptively simple scenario
 - the Chinese room may be an impossibility
- You can *imagine* a situation where one has a book with rules to answer questions in Chinese
 - but only if you do not think too hard
 - in reality, there may be no such book!
 - if the questions can be on almost any topic, then understanding is *required* for that type of complex processing
- And understanding is generally restricted to consciousness
 - Or maybe one needs to conclude that such an advanced book has *potential* consciousness

Purdue University 

And moreover...


- At a smaller level of computation, it is hard to see how consciousness could *not* be (theoretically) possible in computers
- Each cell in your head is data in - data out
 - suppose cells were gradually replaced by tiny computers that kept *all* processing the same
 - » Neuromorphic chips
 - would you claim that at some point you are no longer conscious?
- There is nothing fundamental about *organic* consciousness



Purdue University 


2. Qualia


- Some researchers object to the very idea that computers could become conscious
 - They argue that some things in consciousness are *not* just computation
 - e.g., consider the color **red**
 - There seems to be a particularly subjective experience of seeing something red


Purdue University 

2. Qualia

- Consider two people who see the world in color opposites


Qualia for person 1  "A red apple with a green leaf"

Qualia for person 2  "A red apple with a green leaf"

Purdue University 


2. Qualia


- Clearly, there's a big difference in the perceptual experience of these people, but their behavior is essentially the same
 - And there seems no way to distinguish one experience from the other
 - It's the unmeasurable experience that is a qualia

Purdue University 

2. Qualia

- Qualia proponents argue, for example,
 - you can learn all there is to know about light waves, photoreceptors, neural transduction and coding of color,...
 - But suppose you never see any red objects
 - Your knowledge will not tell you what you will experience when you first see the red of an apple
 - Indeed, you could be tricked into believing a green apple was red (if you had never seen green either)



Purdue University 

2. Qualia

- But this is a defeatist argument, or a pointless one
 - *if* I knew *everything* about light, photoreceptors, and neural representation of colors, then I *would* be able to know what I will experience when I see red
 - it is difficult (maybe impossible for any single human) to know (or even imagine knowing) all that information in an academic sense
 - but that doesn't mean that such information does not exist
- It's partly an empirical question
 - But no one can do the experiment

Purdue University



Conclusions

- Consciousness
- distributed processing in the brain
 - no *site* of consciousness
 - no *time* of consciousness
- Chinese room
- Qualia
- Artificial Intelligence
- Daniel Dennet *Consciousness Explained* (1991)

Purdue University



Next time

- Review for exam 4
- After exam 4
 - Decision making
 - Framing effects
 - Risks
 - Alternatives
- *What every consumer should know before they buy.*

Purdue University

