Sponsored by the Center for Consciousness Studies

Consciousness: The Webcourse

Week 3. Bernard J. Baars

Lecture 3. The theater metaphor:

Limited *conscious* capacity goes along with gigantic *unconscious* capacity.

Summary:

- The limited capacity of conscious contents is a puzzle, because our brains have such vast processing and storage capacities.
- Our giant mind-brains have numerous specialized functions, which are unconscious most of the time. Useful numbers to keep in mind: tens of billions of neurons in cortex, and trillions of connections.
- A theater metaphor combines limited capacity (in the bright spot on stage) and vast capacity (in the audience and back-stage). Several modern models are based on this way of thinking.
- The pros and cons of a theater metaphor.

Readings: Finish reading Chapter 2 in the textbook. Please pay special attention to the notion of conscious limited capacity and the developing theater metaphor.

Baars, B.J. (2002) The conscious access hypothesis: History and recent evidence. Trends in Cognitive Sciences.

Baars, B.J. (1997). In the Theater of Consciousness. Journal of Consciousness Studies.

Parts of Baars & Gage (Eds., in press), Cognition, Brain & Consciousness: An Introduction to Cognitive Neuroscience. Academic Press.

The latest version of a theater & spotlight figure for global workspace theory: Only the bright tip of the spotlight is conscious. A number of other models today are based on this view, including brain models. Notice how it combines *both* limited capacity and vast access to the audience. Conscious experience also has limited capacity and enormous access to unconscious memories, automatic habit components, contextual information, and sub-cortical mechanisms (including emotions and motor control). The spotlight is controlled to select objects of interest, a traditional function of selective attention.
(This image is modified from Carl Carpenter's theater figure).

This week I would like to talk about the "limited capacity paradox." I imagine I'm not the only one who tries to do a lot of multi-tasking --- listening to music while trying to read my email, for example. Unfortunately for us, multitasking has pretty narrow limits. For activities that require real attention --- that is, those that need to be brought to consciousness on a sustained basis --- we seem to be limited to just one content at a time. However, if we have a lot of practice at a task that is very predictable, we can walk and chew gum at the same time. What I notice when I'm multi-tasking is that when I have to do something difficult or demanding on one task, I don't hear the music anymore, or my mind wanders from reading my email.

Limited capacity is closely associated with consciousness (though there are some nonconscious aspects of it). The famous Necker cube is a good example. Here is one:

The Necker Reversible Cube.
Copyright, 2005, Olivia Carter.

As you know, you can flip it back and forth, so that the blue surface either seems to stick out at you, or so that it looks like the back of the cube, on the same plane as the computer screen. Now I'd like you to try to see *both* views simultaneously. I do not think that is possible. If you think you can do it – step into my lab! (Just kidding, but try).

Why is this kind of limit a paradox? Because we have a humongous brain --- 20-40 billion neurons, by a common estimate. And all those neurons are firing about 10 times per second. So in one second, we have 10x100 billion events taking place, or one trillion. The interconnectivity is such that there are trillions of connections between neurons. That's a lot of stuff going on.

We know that the cortex can be seen in terms of many specialized regions. The ones I want to show you here are called Brodmann Areas, first defined by careful microscopic examinations of the cells of cortex by Korbinian Brodmann in 1908. This version of the Brodmann map --- which is sort of the zip code of the brain --- is from Professor Mark Dubin of the University of Colorado (with permission, copyright M. Dubin, 2005).

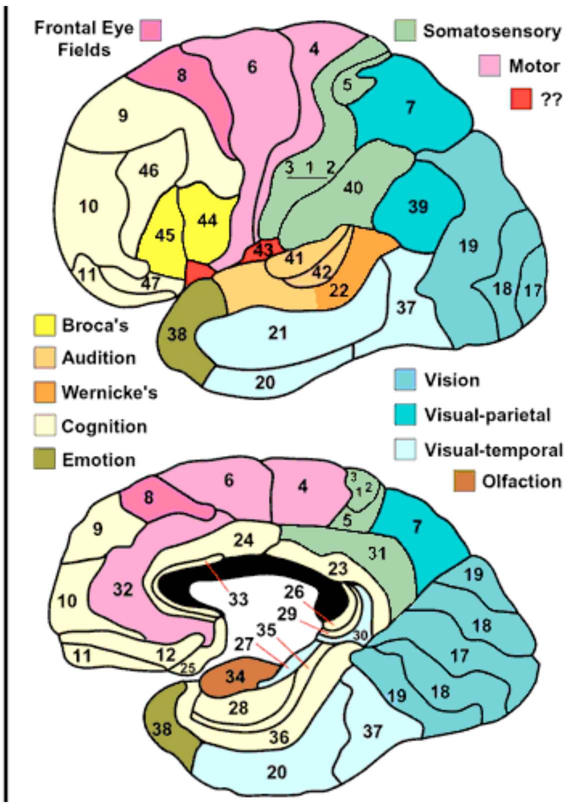

The top figure is the left hemisphere seen from the left side, "looking" to the left. The bottom one is the right hemisphere, also seen from the left side, and also "looking" to the left. You might think of it as a sliced apple, with the outside visible in the top brain, and the inside in the bottom brain. Dubin's picture is color coded, so you can tell the visual areas, the auditory ones, and so on. As well as the language areas (Broca for speech output, and Wernicke for speech input), emotion regions, cognitive regions. All of them involve a lot of complex and sophisticated processing. But most of this processing is unconscious most of the time!

So ---- the question becomes: how can such a massively complex brain be so limited when it comes to the Necker Cube? Or any other ambiguous figure? Why can't we do many things at the same time consciously? Our answer, as you know from your textbook, is the theater metaphor, and the specific models that have been developed by various people (Baars, 1988; 1997; 2002; Franklin and coworkers http://ccrg.cs.memphis.edu/; Dehaene and coworkers, http://cogimage.dsi.cnrs.fr/publications/2005/SBD05/, Murray Shanahan, http://www.doc.ic.ac.uk/~mpsha/Publications.html, and others).

.

Here are some more details.

The central puzzle: Conscious limits vs. unconscious vastness.

Limited capacity phenomena include immediate memory, the selectivity of attention, coherent binding of perceptual features, the limits of voluntary control, and the fact that we cannot do two demanding actions at the same time. If we look only at such evidence, the brain seems to be slow, serial and error-prone. But when the brain is observed directly, it seems dramatically different: It is a massive collection of neural assemblies, cells, layers and connections, each specialized in some specific task like analyzing visual shape, maintaining body temperature or mapping body space. The great bulk of these functions happen at the same time, in parallel, as one great "society of mind." Together, their processing capacity is enormous, though unconscious. The great puzzle is, why is the conscious aspect of the brain so limited when the unconscious part is so vast?

Here are some aspects of the limited capacity puzzle.

1. Working memory limits.

We can only keep about four to nine separate items in working memory. This is the famous upper bound on the number of items in immediate memory, the reason why telephone numbers are generally limited to seven digits, and why psychological rating scales never exceed nine or ten steps. In a brain of tens of billion of neurons, this capacity limit is fantastically small. A cheap calculator can store more numbers than human working memory. Why has the brain not evolved a bigger capacity for immediate memory? The obvious answer is that the brain simply did not evolve to remember phone numbers, any more than human legs grew to run along railroad tracks. But what is it that the brain has optimized in the trade-off between conscious and unconscious capacity?

2. Selectivity: consciousness is limited to only one dense stream of information.

If we try to take in two dense streams of information, like two spoken stories, we can understand only one at a time. That is not limited to language: when two different ball games are shown on a screen we can also follow only one at any time. Nor can an automobile driver talk with a passenger while coping with heavy traffic. When we are preoccupied with a train of inner thoughts we cannot be fully conscious of the flow of external events (Singer & Antrobus, 19xx). Conscious involvement with one dense flow of information always blocks others. This is why the reader's experience of this very sentence can be interrupted by competing thoughts or feelings.

Brain evidence for vast capacity of unconscious processes.

Looking directly at the brain we see great orderly forests of neurons. A structure like the cortex has an estimated 55 billion cells. The interconnectedness of neurons is remarkable; in less than seven steps we can reach any single neuron in the brain from any other. Cortical neurons branch out into vast elegant fibre bundles running between thalamus and cortex; 600 million fibres connect the two hemispheres, and comparable numbers hang in great loops to connect distant points within each hemisphere. It is quite a beautiful and regular arrangement. The brain is massively parallel: many things are happening at the same time. Most of them by far are unconscious.

The brain shows a distributed style of functioning, in which the detailed work is done by millions of specialized neuronal groupings without instructions from some command centre. By analogy, the human body works cell by cell; unlike an automobile, it has no central engine that does all the work. Each cell is specialized for a specific function according to its DNA, its developmental history, and chemical influences from other tissues. In its own way the human brain shows the same distributed style of organization as rest of the body.

It is a remarkable fact that we can create access to almost any part of the brain using consciousness. To gain control over alpha waves in cortex we merely sound a tone when alpha is detected in the EEG, and shortly we will be able to increase alpha at will. To control a single spinal motor neuron we can play back its electrical spikes over headphones; in a half-hour subjects have been able to play drumrolls on their single motor neurons. Of course we are not conscious of the details of control; rather, conscious feedback seems to mobilize unconscious systems that handle the details. Conscious feedback control over single neurons and even large populations of neurons is well established (Buchwald, 1974).

Psychological evidence for vast unconscious capacity.

We come to similar conclusions when we look at the mental lexicon, at semantics or grammar. In a fraction of a second after you glance at a word in this essay, your visual input is converted into a semantic code able to interpret its meaning. Going from words to meanings is believed to require a large, unconscious mental lexicon. The lexicon of educated speakers of English contains about 100,000 words. While we do not use all of them in everyday speech, we can understand each one instantly, as soon as it is presented in a sentence. But words are complicated things. The Oxford English Dictionary, for example, devotes 75,000 words to the numerous meanings of the word "set."

Consider your experience of these examples.

sun: set
temperature: set
tennis: set
tool: set
down: set
stage: set
jet: set
theory: set

ready: set

Notice how rapidly we change our minds about the meaning of “set” when it is primed by a conscious word like “tool.” This ability to combine two separate words has not been found for unconscious words (e.g. Greenwald, 1992; MacKay, 1973). It seems that consciousness is required to combine of individual words in ways that make sense. Of course everyday language always combines words: that is the role of grammar and semantics, both great unconscious bodies of knowledge. It seems that understanding language demands the gateway of consciousness; the same may be true of any other domain that integrates many different chunks of content, such as scene analysis, chess playing, and path finding through a city. This may be another example of the principle that consciousness enables widespread access to unconscious sources of knowledge.

Autobiographical memory.

The size of long-term memory is unknown, but we know that simply by paying attention to as many as 10,000 distinct pictures over several days, we can learn to recognize each of them without any attempt at memorizing. Stephen Kosslyn writes that “The capacity of our visual memories is truly staggering; it is so large that it has yet to be estimated. . . . Perhaps the most staggering results are reported by Standing (1973) who showed some of his subjects 10,000 arbitrarily selected pictures for 5 seconds each. . . . His findings showed that there is no apparent upper bound on human memory for pictures. Moreover, with immediate recall, Standing estimated that if one million vivid pictures were shown, 986,300 would be recognized if one were tested immediately afterward; even after a delay, he estimates that 731,400 would be recalled.” (Kosslyn, 1980, p. 129).

Remarkable results like this are common when we just ask people to choose between known and new pictures. Such recognition tests work so well because they re-present the original conscious experience in its entirety. Here the brain does a marvellous job of memory search, with little conscious effort. We can get an everyday sense of this remarkable performance from recognizing a film seen only once, many years ago, with as sudden sense of familiarity. Often we can even predict the following scene. It seems that we create memories of the stream of experience merely by paying attention to something; but human beings are always paying attention to things, suggesting that autobiographical memory must be very large indeed. Once again we have a vast unconscious domain, and we gain access to it using consciousness. Mere consciousness of an event helps to store its memory, and when we experience the same event again that experience also helps us to retrieve it from among millions of other memories.

In sum, it seems that the very limited stream of consciousness gives us access to billions of neurons in the brain and body, to the mental lexicon, and to an inestimably large source of autobiographical memories. Mere unaided consciousness may be sufficient to create rapid learning and accurate recognition. Consciousness is also needed to trigger a great number of automatic routines that make up specific actions. All these effects of consciousness are unconscious. Consciousness may be considered as the gateway to the brain's unconscious sources of knowledge and control.

Global Workspace Theory can be thought of in terms of a "theater metaphor." A theater combines a narrow focus --- the actor in the spotlight --- with a broad convergence of incoming influences ---- the playwright, the director, the set designer --- and a broad divergence of outgoing information to the audience. If you're keeping up with the textbook reading you already know something about the theater metaphor. But before we go on, let me talk about the uses (and abuses) of metaphors in science.

A lot of the time we can learn from metaphors things that are hard to express directly. "Love is a journey" is one metaphor. "Love is a explosion" is another one. They compare love to different kinds of events, thereby bringing out different aspects.

In the history of science, metaphors have many uses. We still speak about a "flow" of electricity (which is not at all the same as a "flow" of electrons). We speak about "current," and as on, all of which reflects a river metaphor. That made a lot of sense to Benjamin Franklin 230 years ago, but it doesn't tell us much about electron shells, hole flow, and other fancy physics. And yet, the "water flow" metaphor is so useful that even experts use it when it is convenient. And it helps in communicating basic concepts. Good metaphors are easy to visualize, and that helps us in understanding. But scientific metaphors are always partially wrong!!!

Below is a picture of an orrery, a mechanical clockwork that helps us to understand the solar system. Orreries were first built around Newton's time. They show a "sun" at the center, surrounded by spheres of different sides to represent the planets. All very nice. But wrong!!! The real solar system has no metal arms holding the planets in their orbits. The distances are way out of proportion. Most obviously, the dimensions of the planets and the sun are all wrong. The point: To use scientific metaphors properly, we have to understand the part that's metaphor, and the part that is denotative theory. For Newton, the expression F=ma is denotative theory. You can directly measure m (mass) and a (acceleration). So that is not metaphorical.

The theater metaphor is a useful way to think about consciousness. It can help us come up with denotative or literal hypotheses. But we have to use it sensibly.

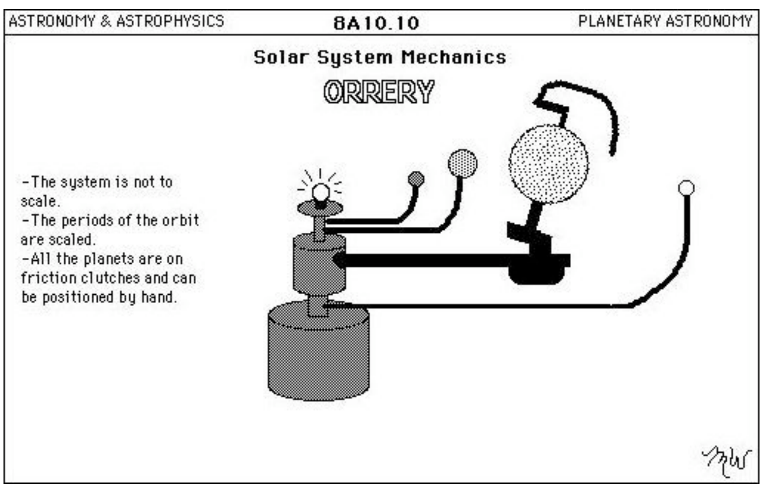

(Copyright belongs to owner.) Orreries were invented in the 17th century to represent the Copernican theory of the solar system. They are machine metaphors --- they're not the real thing, because the solar system doesn't have metal arms to keep the planets in orbit, and it doesn't have a light bulb for the sun, and it doesn't have a little motor to keep the planets moving. Orreries nicely capture the spatial relationships in the solar system. We have to be clear about what it does and doesn't do.

So here's the metaphor: Global Workspace Theory (GWT) can be compared to a theater of mind, in which conscious contents resemble a bright spot on the stage of immediate memory, selected by a spotlight of attention under executive guidance. Only the bright spot is conscious; the rest of the theater is dark and unconscious.

And here is some theory: GWT has been implemented in a number of explicit and testable global workspace models (GWM's). These specific GW models suggest that conscious experiences recruit widely distributed brain functions that are mostly unconscious (unreportable). A large body of new findings support that view. For example, brain experiments show that while unconscious visual stimuli evoke high activity in visual cortex, identical conscious stimuli reveal an additional spread of high brain activity to frontal and parietal lobes (Dehaene, 2001). Similar results have been found for hearing, touch, pain, and sensorimotor skills (Baars, 2002). The conscious waking state supports such fast, flexible, and widespread brain interactions, while unconscious states do not (Baars et al, 2004). These findings illustrate the ability of the GW framework to suggest novel and falsifiable hypotheses.