

Table of Contents

CHAPTER 1: How do we build on the basic observations about consciousness?	2
CHAPTER 2: The Conscious Brain: Global Workspace Theory (GWT)	11
CHAPTER 3: In the Theater Metaphor, the Stage is Limited but the Audience is Unlimited	15
CHAPTER 4: The Paradox that Gave Rise to GWT	26
CHAPTER 5: What are some of the things that people get wrong about GWT?	39

CHAPTER 1: How do we build on the basic observations about consciousness?

Alea Skwara:

Hi, my name is Alea Skwara. I am a PhD candidate in psychology at the University of California, Davis. I'm in the cognitive neuroscience area of our program. I'm a psychologist, meaning I care about questions of the human mind, but I look at those from a neuroscientific perspective.

The primary method I use in my PhD studies is EEG or electroencephalography, which is a way of looking at electrical impulses that are generated by the brain and recorded at the scalp. My research focuses primarily on compassion and responses to suffering.

In the lab, we study the effects of intensive meditation training on a whole range of things such as your attention, your perception, and your emotions. I'm specifically interested in how this sort of intentional training can shift the way that we relate to suffering, whether it's our own suffering or that of others.

The primary question that I'm looking at in my dissertation is whether engaging in this sort of practice can help us widen our circle of care and help us expand the classes or range of people for whom we find it easy to feel compassion for when they suffer and for whom we're willing to invest in helping them not suffer.

So, essentially expanding this care that we easily and automatically have for those who are close to us to a broader and broader circle of people.

Ilian Daskalov:

My name is Ilian Daskalov. My initial academic interests were very different than the people present in this conversation. I initially started studying business administration, marketing, and management.

Somewhere along the way, I got more interested in human behavior than in the behavior of corporations.

Eventually I moved to California and decided I wanted to study human behavior, the human brain, and human consciousness and all of that began with meditation. Once I started meditating, I noticed a lot of things about myself that I wasn't really noticing before.

All of that got even more involved as I started regularly going to sensory deprivation chambers, which are these very fascinating rooms – well, there are several different kinds but the particular ones I went to were chambers filled with salt water and you lay in the water and you're sort of just floating.

The room temperature is the temperature of your body and it's completely dark and you're isolated from any sound. So, theoretically, all your senses are deprived and you're just left with yourself and your thoughts and it was very interesting for me to observe what emerged in my conscious brain, what I was thinking about.

I started wondering: *'Why am I thinking about these things? Where is this coming from? Where is this emerging from? How are these unconscious processes happening inside of me?'*

That is a very fascinating field to me and I am so excited about our conversation today because I want to learn a lot more about human consciousness.

Bernard Baars:

That is interesting to hear because – as you and Alea know – we do have an awful lot of research on this spontaneous stream of consciousness but I'm willing to bet that some of your neuroscience professors don't know a thing about it.

My observation – and you can tell me if this is true or not – is that when people talk about the default network in the brain I believe they're talking about stream of consciousness. As stream of consciousness is hardly passive and it's a default state. It is a very active state and it's a state that reflects our emotions and our current concerns and imagery. All those kinds of questions that spontaneously emerge and that we know about from Jerry Singer and other people over very long periods of time that have collected that data.

However, the people who *actually* know about it seem to be sort of a cult. I consider myself a member of the stream of consciousness cult and I wish we could draft a few more members.

So let me know if that's true or not.

Alea Skwara:

I think there has been pushback and I think that's an interesting point.

So, the Default Mode Network is a network of brain regions that tend to become more active in the absence of an explicit task. This first became something that people talked about broadly with the advent of resting state fMRI, which stands for *functional magnetic resonance imaging*.

People realized that if you put someone in a scanner and you record their brain activity and you don't give them an explicit task to do, like '*Press this button when you see a box of this color*' or whatever cognitive task we give people, there's a network of brain regions that increased in their activity and this was named the Default Mode Network.

Transcript for "Global Workspace Theory: Origins & Evidence"

Bernard Baars:

Sorry, a network that increases when they stopped doing a task?

Alea Skwara:

Exactly – in the absence of an explicit task.

There's this network of brain regions that Bernie was talking about that increases their activity and researchers who were working in this 'Default Mode Network' world were like "*Oh, this is like the brain's metabolic baseline.*"

Bernard Baars:

Right.

Alea Skwara:

As if nothing is happening when we're at rest.

There was this reference to the Default Mode Network as some sort of neural baseline that you could subtract away from things. When you're doing a task, just collect someone at rest and then have them do a task and subtract the two and then we know what brain activity is due to the task.

However, I think there's been movement away from that.

Bernie, I think you're completely right, in that if we think about what our brain does when we're in the absence of explicit tasks, we are essentially mind wandering and we're experiencing stream of consciousness and we're doing all sorts of stuff. There has actually been a movement where researchers are saying "*No, that's not a baseline. That's a different sort of cognitive task.*"

I don't know that people think of it in terms of consciousness, which I think is Bernie's perspective of the stream of consciousness or conscious awareness. So, I think there's been a...

Transcript for "Global Workspace Theory: Origins & Evidence"

Bernard Baars:

Yes.

I should say that if you look at Jerry Singer's work, Singer claims that what appears in the spontaneous stream of consciousness is lots of conscious stuff that you don't really understand if you're the person having that stream of consciousness.

All kinds of unconscious and primed stuff emerges during this spontaneous stream of consciousness and – because human beings are adaptable and because the stream of consciousness is part of our adaptability – I would say, hypothetically, the things that are most important for us as individual human beings are the things that would tend to come up.

Bernard Baars:

When we get drowsy we have, in a way, the experience of sensory isolation because the thalamus shuts down sensory input. I like to pay attention to that moment when I'm experiencing stream of consciousness, when things get really weird and you stop talking like an adult to yourself and some visual image comes up or some thought that is clearly not an adult thought emerges.

It's a distinct moment sometimes – at least in my experience – and it's kind of a fun thing and it's an important thing psychologically.

But it's also a very safe way to play with your head.

Alea Skwara:

Well, I think this actually leads us into the first question we wanted to talk about – how do we define consciousness?

Bernard Baars:

Yes, right.

Transcript for "Global Workspace Theory: Origins & Evidence"

Alea Skwara:

Ilian – is that where you were getting ready to go?

Ilian Daskalov:

That's where I was going with it.

Alea Skwara:

So here we are.

Ilian Daskalov:

Bernie, can you give us a definition of consciousness so we're all on the same page?

Bernard Baars:

You're going to ask me the easy question, right?

Ilian Daskalov:

We're going to start with that one, yes. How do you define consciousness?

Bernard Baars:

Well, I believe the answer in the history of science is that you use whatever seems plausible and easy to observe.

So, for example, when Galileo was trying to figure out a method for measuring temperature he figured out that if he filled little bottles with oils and placed the bottles in water, the bottles would drift up and down in the water.

The order of these little bottles that drift up and down depends on their internal density. So, if you can put some stuff in there that has less density or simply lower the water level inside these little bottles, then you can actually get a visual, isotonic,

Transcript for "Global Workspace Theory: Origins & Evidence"

almost quantitative – well, at least you can tell which is warmer and which is colder and that was essentially the first thermometer.

Actually, there were presumably all kinds of people working in 'smithies' and making campfires for the evening who would have learned about the colors of different things that you threw into the campfires so we can assume that these people presumably knew about temperature at some level...

Ilian Daskalov:

I have that thermometer in my room, actually.

Bernard Baars:

The bottom line is that Galileo was very clever and he was a very good observer. He was very thoughtful and was hanging around a lot of smart people, as well. This is one of the great places where ideas emerge.

He came up with the first thermometer and he recognized that this was a significant invention because it's not just enough to see it, you have to recognize that this was a development in human technological scientific history.

So, if you'd asked Galileo '*What is heat?*' that's analogous to asking me to define consciousness because we're maybe two or three steps above the level of those floating little glass bottles.

Of course, the first level is the old one. Can you give people a conscious stimulus and will they tell you about it and be accurate?

So, '*accurate*' is the condensed way of saying '*accurate voluntary report*'. Because the senses are so easy to study, relatively speaking, you start with sensory stimuli and then you get to Newton and the prism experiments -- people were just in love with the problem of color early in the 19th century and Newton came up with his own seven color theory which was supposed to be right but turned out to be wrong.

Transcript for "Global Workspace Theory: Origins & Evidence"

So, they argued it out. They did the right experiments and it just grew and grew. Sensory sciences came from that kind of thing, as well as psychophysics, so it's enormously important.

And now we have about half the academic population of the world telling us that we don't know anything about consciousness, which just is not true. That's just not true at all because we have a lot of good stuff.

Of course, our ignorance is a lot bigger than what we know but we're definitely not starting from zero.

Alea Skwara:

What I'm hearing you kind of say is that, in a way, our scientific definition of consciousness is limited by what is observable, just like the understanding of heat was limited by what was observable in this first thermometer contraption.

We're very early, in some ways, in our scientific ability to measure consciousness. So, when I ask you as a scientist of consciousness '*What is consciousness?*' you can give me the observable definition, but that doesn't mean that this is consciousness itself.

Bernard Baars:

Exactly.

And then, of course, the next step is what's called bootstrapping which is sort of a humorous metaphor for pulling yourself up by your own bootstraps. This is not recommended if you're climbing a mountain but it kind of works as a popular joke because when we go to the next step, which is the EEG, we get the dreadful EEG from 20 years ago or 30 years ago.

Alea Skwara:

Like 20 channels, 5 channels...

Bernard Baars:

Yes, exactly.

And you lose track and this is actually true – it's a horror story and scientists are not supposed to say bad things about their own knowledge of science, but we essentially forgot that people have thick skulls and that you get a 99.9% drop in signal-to-noise ratio between cortex and the outside of the skull.

That was not a good idea.

And people who work with animals with tiny little skulls were able to get a lot better. But, in any case, that was part of the process of blowing it really badly and then realizing that you're blowing it and fixing it and getting better.

CHAPTER 2: The Conscious Brain: Global Workspace Theory (GWT)

Ilian Daskalov:

Bernie, I know that when you first started studying consciousness that was something which very few people in your field at the time were doing. As a matter of fact, it was considered taboo in a way – is that correct?

Bernard Baars:

Exactly. Yes.

Ilian Daskalov:

Can you tell me how you got interested in studying human consciousness and how you came up with the Global Workspace Theory?

Also, could you give a brief definition for GWT?

Bernard Baars:

Basically, consciousness has never really left the natural psychology of everyday life.

I mean, we talk about taking an aspirin but an aspirin is just a crystallin protein and you pop it in your mouth and then somehow this psychic event occurs, right? If you're lucky the headache goes away. That's the only way you know that – in the old days certainly – the situation is changed in subjective experience.

Of course, we know of many everyday things that happen, for example, you're hungry so you eat something and your hunger goes away. That's a subjective experience.

Transcript for "Global Workspace Theory: Origins & Evidence"

If you really look at the mythology that has arisen for weird reasons around the idea of having a gulf between subjective experience and the so-called '*objective world*', which is basically the world that you can subjectively agree on with another's objective mind.

When you look at that there seems to be a gap, but even Gustav Fechner around the year 1800 was convinced that was a solvable problem and it was not a metaphysical gap. We got entrapped in very ancient and very tricky word games, which I think is a disaster because there are so many ethical implications of realizing that, for example, the animals you work with in the laboratory are conscious and your pet cat is conscious and so on.

Of course, in cultures that...

Alea Skwara:

Bernie, I think that's a really important point in defining consciousness.

It sounds like you're saying pretty clearly that you don't consider consciousness to be necessarily a 'human only' construct or experience. I think – just to give our listeners a little bit of context – researchers and philosophers have been debating the nature and origins of consciousness for centuries.

There are a lot of different theories about what consciousness is and what creates it. Whether it's uniquely human, whether it's spiritual, whether it's physical. One of Bernie's major contributions to the science of consciousness is the Global Workspace Theory, which we will sometimes refer to as GWT.

In the most basic terms, my understanding is that GWT is a hypothesis about how conscious experience arises from the mind and how the brain might give rise to this conscious experience.

Is that accurate, Bernie?

Bernard Baars:

It is accurate.

Please forgive my current obsession with the history of these ideas because as I've gotten a better understanding of both the subject and the brain basis of consciousness, I've started thinking about the Hindu story of Arjuna and the Bhagavad Gita. Arjuna is a charioteer and he is deeply conflicted about whether he is supposed to kill his cousins or violate his warrior caste.

It's one of those Hamlet conflicts that populate human literature and human folk tales, of course. In the story, God comes down and says, *'Okay, here's the solution to your dilemma.'* And that's Lord Krishna. But Arjuna is a metaphor for consciousness and decision-making because he is facing a dilemma about whether he is going to whip his horses into going or sitting on this big chariot with the bronze and all that fabulous stuff.

So, that is a metaphor for consciousness and I'm just going to say that it's also a Global Workspace metaphor because Arjuna is getting all kinds of information from his memory and from his mentors who told him that he must go into battle and risk his own life in order to have an honorable battle with the other warriors.

Arjuna's got his own ethics that he's been taught and he can no longer do certain things because he's been taught that those things are wrong.

Alea Skwara:

When I first learned about Global Workspace Theory, I learned about the Arjuna metaphor.

I first learned about Global Workspace Theory using your spotlight metaphor, right? The spotlight of attention. When I was first learning about Global Workspace Theory in graduate school I was taught that there's a *'theater of the mind'* and most of the

Transcript for "Global Workspace Theory: Origins & Evidence"

theater is in the dark, and there is a spotlight of attention which can shine light anywhere in the mind and that area of the mind that is illuminated by the spotlight is what is in our conscious awareness.

So, there's this directed movement of our consciousness and at any moment only a small part of the mind is illuminated, but – like this metaphor you just described where Arjuna can go into battle or choose a different direction – the spotlight can move into different spots of our mind and bring certain things to our conscious awareness.

Is that still, before we get into the brain science, just conceptually...

Bernard Baars:

Yeah, conceptually it's almost right.

The reason why I say it's almost right, is that I've always been troubled by the difference between voluntary attention and involuntary attention. For example, if you're three years old and your little brother comes up to you and says 'Boo!' and it's involuntary – you're scared or you're playful but it's not your choice to pay attention to 'Boo!', it is imposed upon you.

That implies that there are selective mechanisms operating unconsciously prior to that moment that bring that obnoxious little brother who is getting right in your face and saying 'Boo!' as a high priority thing. So, there's a lot of unconscious control of access to consciousness.

So, the idea that attention is a selective and directive thing while consciousness just 'is' – that's a poor way of saying it, but I think it makes the point. Consciousness in a sense is thrust upon us. We can do things to avoid it, but as an innocent baby we don't have that much control over what becomes conscious. We need to pee and we just do it and consciousness is part of that process, I believe.

CHAPTER 3: In the Theater Metaphor, the Stage is Limited but the Audience is Unlimited

Bernard Baars:

A theater provides a useful analogy or metaphor for thinking about consciousness and here we work out the implications as our theater metaphor shades into actual hypotheses.

We start by considering how many times scientific ideas first emerge.

Metaphors have a long history in science to help us make the leap from the known to the unknown. The clockwork metaphor of the solar system helped astronomers in the 1500s understand the interplay of the planets. William Harvey's discovery that the heart pushes blood through the veins and arteries, much like a pump, has been a useful metaphor for several centuries and it is still used today.

Physicists around the year 1900 found the image of the atom as a tiny solar system a helpful starting point for understanding subatomic structure. Many scientific theories begin as humble metaphors.

Obviously, we don't expect to find tiny theaters in the brain. We may, however, be able to find structures that display and disseminate conscious contents. The conscious parts of the brain seem to include the sensory areas of the cortex, perhaps some surrounding areas, supported by a few subcortical structures; and fleetingly, perhaps, amodal cortex, which does not give rise to sensory qualities. Together they may provide the stage for the unconscious audience in the rest of the brain. The theater concept helps us to think about brain functioning in an interesting way.

Metaphors must be used with care. They are always partly wrong, and we should keep an eye out for occasions when they break down. Yet they often provide the best starting point we can find.

Transcript for "Global Workspace Theory: Origins & Evidence"

A theater provides a useful metaphor for thinking about consciousness.

The major features of conscious cognition can be approached with a theater metaphor, an ancient but still useful idea. Global workspace theory (GWT) is a testable theory, not a metaphor. The theater metaphor does suggest simple interpretations for selective attention, for conscious contents, and for the vast collection of unconscious knowledge sources of the brain.

These include unconscious contexts or "frames" which shape every conscious perceptual and cognitive content. (See Baars and Geld, 2019, Chapters 6-10.)

Theater models are widely used; as it happens, all of our unified models of mental functioning today are theater metaphors; it is essentially all we have. Cognitive architectures developed by Allen Newell, Herbert A. Simon, John R. Anderson and others resemble theaters. All are equipped with working memories that are limited in capacity. All involve active elements, much like the conscious elements of working memory, though without using the word "consciousness." And all have large sets of unconscious mechanisms, whether they are called productions, long-term memory, or procedural memory (Newell, 1990; Anderson, 1983).

These theories have been developed over the last forty years based on a vast range of evidence, from studies of chess players to arithmetic problem-solving, mental rotation of visual images to action skills. A remarkable group of distinguished scientists have devoted careers to these integrative conceptions of human cognition.

The theater of consciousness is a simple idea with many important implications.

Imagine a stage, with an attentional spotlight shining on the stage, actors to represent the contents of conscious experience, an audience, and a few invisible people behind the scenes, who exercise great influence on whatever becomes visible on stage. The stage receives sensory and abstract information, but only events in the spotlight shining on the stage are completely conscious.

The actor in the spotlight frets and struts his hour upon the stage, directed by the playwright and director, against a background created by scene setters. These

Transcript for "Global Workspace Theory: Origins & Evidence"

behind-the-scenes influences are frame stacks (context operators), unconscious systems that shape conscious events. The spotlight selects the most significant actors on stage, and once lit up, their messages are distributed to an audience consisting of all the unconscious routines and knowledge sources — the vast array of unconscious tools we use to adapt to the world.

Major features of conscious cognition can be approached with a theater metaphor, an ancient but still useful idea.

Here you see the original flow chart of the theater metaphor. But now we are using not just theater terms, we're using actual scientific hypotheses. So, at the very top, you see the behind-the-scenes operators -- we call those frames or contexts. They're not conscious, but they shape and influence what does become conscious.

Just below that, we have competing processors that want to gain access to the conscious stage so they can be heard by everybody. They could be sensory inputs or they could be retrieved memories or they could be new ideas, anything like that. Once they capture the search light or once the search light captures them, they get to talk to everybody.

Below that, you see the unconscious audience that consists of memory systems, interpretive systems like the systems that allow you to put your sensory input into linguistic terms like you may be doing right now.

There are also numerous automatisms, which are habits that you have practiced so much that they've fallen out of consciousness and they've faded from consciousness, and they are now automatisms. And then of course, there are any number of motivational systems that also shape your experience and are shaped by your experience.

So, if you have a wonderful love affair that is likely to shape your future experiences, and if you have a bad love affair, that's also likely to shape your future experiences. Those are motivational things, those are not just automatisms, they are huge goal systems.

Transcript for "Global Workspace Theory: Origins & Evidence"

The best analogy or metaphor for consciousness in the brain is the theater metaphor – here we see a beautiful example with the conscious contents at the very center of the image, which is the stage of the theater.

Notice that there is a spotlight shining on the stage and lighting up the actors who are getting a chance to speak and declaim to the entire audience. The audience could be as few as a hundred people, or it might be as many as a million people if it's a radio audience or a television audience or the worldwide web.

You'll notice that there are many tiers of different audience members. Some of them are called bottom-up inputs, which means the sensory inputs without expectations shaping them. And the top tier says top-down frames, which are the expectation driven processes in the brain.

Right now, when you hear a sentence that I'm speaking – and if I stop instead of saying a word – then you will start predicting the next word, just like that. And you can do this to all of your friends simply by pausing in the middle of a sentence. That will show the extraordinary amount of expectation driven processing that your brain does.

But it also does bottom-up processing, so if you get interrupted by a surprising thing you don't necessarily get caught up in your expectations. You don't get stuck in your expectations. If something very surprising happens, you still see it. That is called bottom-up processing.

So, this is the summary of a great many theories of mind and brain in the most simple and revealing and also traditional way of thinking about the mind and its brain.

Alea Skwara:

Then, kind of tying that back to the metaphor, I kind of hear you saying that maybe there are two questions to answer. Correct me if I'm wrong, but in metaphorical terms, the first question is: *What is the spotlight? What is this light that's shining down and illuminating part of the stage?* So that might be like the 'What is consciousness?' question.

Transcript for "Global Workspace Theory: Origins & Evidence"

But the second question is about what's moving that spotlight. How does it get moved to different parts of the mind? Those may or may not be the same thing or are they one thing?

What is your perception of that, Bernie? Maybe we need to get into the neuroscience of it...

Bernard Baars:

It is a whole swarm of things.

Partly because we do have some voluntary control, and as we get older we gain more voluntary control over our ability to pay attention to things. For example, there's the famous Walter Mischel candy test for young children, where children are given marshmallows, right? In this test the children can have some more marshmallows if they wait for 10 seconds, and that turns out to be predictive of the development of voluntary attention later in life.

It is really important and it's just so early.

Alea Skwara:

This executive control.

Bernard Baars:

Yes, exactly. Control is the technical term and it's the right term. Yes.

Alea Skwara:

A lot of kids in that experiment use different strategies to not look at the marshmallow, right? Like to deal with it.

Some kids literally look away and kind of try to take the marshmallow perhaps out of their conscious awareness as a regulation strategy. I think that is really interesting and how that ties back...

Transcript for "Global Workspace Theory: Origins & Evidence"

Ilian Daskalov:

Absolutely.

Just moving your attention from one object to the next and just trying hard to not think about what's in the background. I really liked the metaphor you used when describing Global Workspace Theory and particularly the spotlight.

I was wondering what exactly moves that spotlight? And that brings me to my next question. What are some of the features of Global Workspace Theory?

We started touching on them, but if we can just elaborate a little more.

Bernard Baars:

Let me back up to where my own puzzlement started on this, and it really was puzzlement and it continues to be puzzlement, I should mention.

Alea Skwara:

This is kind of getting into the origins of how you first came up with GWT, correct?

Bernard Baars:

Exactly.

In 1980 I had a postdoc at UCSD with a very interesting group that was doing cognitive science – and this was the first time that people started using the term '*cognitive science*', primarily to get away from the behavioristic taboo at that time which was too constraining and too limiting. It essentially made it impossible to make inferences based on the evidence that you had.

Alea Skwara:

So, there was a focus just on behavior and not on what the internal state was?

Transcript for "Global Workspace Theory: Origins & Evidence"

Bernard Baars:

Exactly.

So, you got drawn into the trees but you couldn't see the woods through the trees. And behaviorists were very good and very conscientious and very idealistic because they wanted to change the world for the better.

They were personally nice people and all that but they were stuck in this box they had imposed upon their own minds, which was usually done with very good intentions. Certainly, they were talked into it by philosophers who also took that road and by other people who thought that real science had to do with physical things...

Alea Skwara:

What was directly observable, right? And the mind is not.

Bernard Baars:

Yes, but we're getting more sophisticated here.

Early on, physical things were sort of like a plow and the horses that you needed would go to the bathroom on the field which would help you to fertilize the field. It was very, very basic. It was a kind of early Piagetian physical wisdom – if that makes any sense – that the physicalism of the objects that you and I can see and touch and smell and get rid of if we don't like them.

Alea Skwara:

So, you saw something was missing here in terms of the study of humans or the mind, right?

Bernard Baars:

Yes.

Transcript for "Global Workspace Theory: Origins & Evidence"

You're asking interesting questions that are inherently mysterious because what's called the '*zeitgeist*' – or the temperature of the times, the feeling tone, the words and the ideas that came in – meant that lots of people were interested in consciousness but they were also interested in unusual states of consciousness.

Like what happened to you when you dropped acid or what happened to you when you got involved with people who meditate?

Alea Skwara:

Fun questions.

Bernard Baars:

Yes.

And I actually – very much against my own will – got involved with transcendental meditation which was taught by a really wonderful man from India named Maharishi Mahesh Yogi. That's not actually his name, it's an honorific title given to people in that particular tradition which occurs when peers within an ashram believe that the peers they want to give the honorific title to no longer have a personal viewpoint on the world. They are asked to drop their family name and adopt this honorific title after being told by their peers at the ashram that they can now use that title.

Anyway, the Maharishi Mahesh Yogi came out to Hawaii and it was the most stressed out place he had ever been to. As he went along, he started to adjust his teaching to people who were extremely stressed out and he had to back off on the methods he taught.

But nevertheless, after about five years of doing regular sessions of transcendental meditation, I had a vague intuition that there was actually something to this practice.

Alea Skwara:

Transcript for "Global Workspace Theory: Origins & Evidence"

Why is that? Did it shift the way you understood consciousness or inform the way you studied consciousness?

Bernard Baars:

I don't think I was thinking about consciousness.

Actually, *I was* thinking about consciousness in a serious way because I'm just interested in the objective basis of consciousness. The word objective is questionable, but it essentially means whatever we can easily share in the public domain at that time, that's the kind of definition we used.

By now, of course, what we're really thinking about is what would happen if we put 600 electrodes over the brain – and specifically over the cortex – of some person in epileptic surgery where it's ethically acceptable to do that? We can have hundreds of electrodes directly on the cortex without having to worry about how thick the patient's skull is and we could get fabulous evidence directly from the brain.

After that procedure, we could ask the patient to meditate or we could ask them to go into a sensory deprivation experiment or do something much more subtle and we could actually see things – that is such a revolutionary thing.

Alea Skwara:

What can these altered states tell us about what some people call '*ordinary consciousness*'?

Bernard Baars:

For sure, ordinary consciousness is very important and then there are extraordinary states of consciousness which are obviously also important.

Alea Skwara:

So, my understanding is that GWT was originally inspired by Artificial Intelligence. Is that true?

Bernard Baars:

Here's the paradox that I kept running into and which many other researchers had also run into. I think many of them kind of shrugged it off as just a fact of life but for some reason I couldn't shrug it off.

The paradox is called the *Limited Capacity Paradox* and it's the reason why we see YouTube videos of people texting or talking on their cell phones on the street and walking straight into a flowing stream of traffic and getting hurt. We can actually see this and there's nothing more dramatic than that.

However, the phenomenon itself is true not only for human beings but also for cats; it's true for all mammals, really. Actually, I would guess it's probably true for all vertebrates.

Alea Skwara:

Essentially, when you are focused on some things you miss others things.

Bernard Baars:

Exactly right.

We can call it *distractibility*. There's a whole set of mental functions that can be grouped under the general term of *Central Limited Capacity*, which includes volition and feelings of knowing and conscious sensory perception. There's a whole group of things that are included that way and the problem from a biological point of view is that they're all dysfunctional. They all kill people and they killed our ancestors, including early mammals and other vertebrates, and so on.

Transcript for "Global Workspace Theory: Origins & Evidence"

This thing kills animals on a regular basis. It's tremendously destructive, which means it has very high biological costs, right? You have to ask, '*What are the compensatory benefits from a Darwinian point of view?*'

Because if there's no benefits, we wouldn't have that anymore. We would not have that kind of very limited capacity. We might've developed two separate hemispheres, for example, with two separate optical systems that can work independently of each other. There are hunting spiders who have nervous systems that certainly look like that from the outside.

CHAPTER 4: The Paradox that Gave Rise to GWT

Alea Skwara:

So, the next question is – and Ilian, this might pull us into a discussion about the features of GWT because I think one of the major features is Limited Capacity.

But, the next question is: *What is it about our ability to be consciously aware of a limited percentage of what's happening in our mind or our sensory environment that's so advantageous that it makes it worth missing what we're not consciously aware of?*

Bernard Baars:

Exactly.

The GW hypothesis is, essentially, that there has to be a compensatory event in the brain happening and the most plausible one for various reasons – including other people's work, of course – was that there is some kind of very wide recruitment of brain resources that happens as a function of becoming conscious of something.

We're only conscious of, let's say, a single star on a dark night but that fact that we are only conscious of that has immense consequences in the brain, and today we can actually see that happening. I called that the *Broadcast Hypothesis*.

Now we know that it's a resonant broadcast because the huge brain system that is involved -- the cortex and the thalamus -- only has bi-directional wiring basically. So, if you send out a broadcast you can get feedback right away which means that -- given the fact they're both excitatory neurons -- you can get an explosion of activity.

The only trouble with that is, why doesn't your brain blow up?

Transcript for "Global Workspace Theory: Origins & Evidence"

Alea Skwara:

I'm going to pause here so our listeners can get a better understanding of the theory and main principles of GWT before we get into some of the really interesting brain evidence. I think we should start with the overarching theory and then we can go into the brain evidence.

In my understanding, GWT began with a question: *How does a serial integrated and very limited stream of consciousness emerge from a nervous system that is mostly unconscious, distributed, parallel and that actually has huge capacity?*

Then we get to...

Bernard Baars:

That was the first question.

Alea Skwara:

Right.

That was your motivating question and then quickly I'm just going to add that a key prediction that's generated from GWT is this widespread integration and, like you said earlier, broadcasting.

Maybe you can talk about that original question and then talk about that key prediction, but let's start with that question first.

Bernard Baars:

Exactly.

The way these things happen is in a very subtle way. They kind of sneak up on you and you get intuitions and you scribble down ideas and so on.

Transcript for "Global Workspace Theory: Origins & Evidence"

This is probably how it happened, and during my year with the cognitive science group in 1980 at UCSD I got to talk to all these interesting people and learn more. I got the opportunity to get stuck on this weird paradox in science – they're a pain in the rear, right? However, at the same time, they're often very revelatory if you stick with them and if you give them more attention.

If you are Copernicus, you stick with questions about astronomical objects even though the planets in our solar system are certainly little dots that are sort of predictable but they violate everything that you know about the universe because the earth is rotating. You see the fixed celestial bodies rotating and it's a very beautiful system. The pre-Copernican system is gorgeous and it's the system that human beings used for probably tens of thousands of years for navigation.

For example, you only have to look at the North Star to then figure out where the other stars rise across the horizon and where they dip below the horizon on the other side, and you've got a compass that allows you to navigate. This is how the Polynesian navigators still worked until they got their GPS's.

Alea Skwara:

So, the specific paradox that captured your imagination was this paradox of the serial integrated limited stream of consciousness with this massive parallel unconscious distributed system.

So, is that true? That is kind of the question when you talk about the question of paradox.

Bernard Baars:

That is true.

You have to remember that there were other people that I kept learning from, and because of the taboo of the time people would not tell you directly that.

Transcript for "Global Workspace Theory: Origins & Evidence"

For example, Allen Newell and Herb Simon were great scientists and they worked together for 50 years on this thing called *cognitive architectures*.

This started in the 1950s when Herb Simon was talking to a Dutch chess master – who also happened to be a Gestalt psychologist – but as a chess master, this man understood subjectively *'Okay, So I think about the Gestalt of pieces on the board, right? And I see this opening and my opponent doesn't see this opening. So, step-by-step, I think I can force him into a double bind that will make him lose the game.'*

It was such a subjective experience that I don't think anyone was allowed to talk about it.

Alea Skwara:

So, that was the taboo.

Bernard Baars:

Right.

But cognitive architectures were developed over and over again and the only thing they had in common with each other, as far as I could tell, was this skinny stream of conscious information. These cognitive architectures have this enormous memory which was actually called, well, there were various names for it because computationally people started to play with these things and learn things.

But the paradox kept on coming back because the skinny stream of conscious information kept on going with this enormous mountain range that was unconscious. So, that's what hit me. Then, of course, Allen Newell came up with the term *Global Workspace* for his own computer system which demonstrated also that this was a very practical thing.

This Global Workspace Architecture was stunningly practical for problems that were thought to be impossible problems at that time.

Transcript for "Global Workspace Theory: Origins & Evidence"

The reason it was practical is because you essentially put many unconscious minds together in the same room and allowed them to talk to each other. They shared information and they found out if they were wrong, which is crucial. They also came up with shared solutions that were better than any single solution that any single player in that game could have come up with.

I think that was a breakthrough and people knew about it and they were very interested.

Alea Skwara:

That was in the world of parallel computing.

Bernard Baars:

It's actually parallel interactive computing, which is a little twiddle on parallel computing.

But people started to think about what would happen if we had a swarm of computers – they were just beginning to even conceptualize that because everybody knew at that time that computers were expensive and they fill the whole room and you had to have a special air conditioning system, and programming them was a horror story. It was terrible. People suffered and worked long nights to make these computers do simple things. Then suddenly people began to think about Moore's law. So, suddenly you have more powerful computers with cheaper memory and then people started thinking about what would happen if you had two computers instead of only one, or what if you have *three* computers?

I think the space shuttle at some point had six different IBM computers that were identical to each other, running identical programs because nobody trusted these computers to get anything right. They had them voting on whatever the right decision was for the space shuttle to maneuver.

Transcript for "Global Workspace Theory: Origins & Evidence"

Alea Skwara:

Right.

Bernie, just to orient our listeners in time – what year was this that you were thinking about this paradox in psychology of the huge unconscious mountain with a tiny conscious stream and, at the same time, this Global Workspace of artificial intelligence and parallel integrated computing that you started to bring together?

Bernard Baars:

It's really very interesting.

I think I worked myself out of the behavioristic box – that mental box that we were all in – very slowly and by learning about the history of philosophy of science where people keep on running into these kinds of paradoxes.

This is not the first paradox that people bumped into and couldn't figure out what was happening, right? The good thing is that this paradox has forced you to enlarge the box that you've got sitting over your head and you shrugged your shoulders and suddenly started to see more. The notion of information processing, which was pretty sloppy at that time...

Alea Skwara:

And what year was this?

Bernard Baars:

This was 1980.

In 1980, a lot of us were beginning to think '*Hey, let's think about things in terms of information processing*' because at that time you could do anything you wanted to because there was no restriction on the idea of information processing.

Transcript for "Global Workspace Theory: Origins & Evidence"

You could do algebra, you could do verbal problem solving, you could imitate a paranoid patient in a mental hospital, which was, incidentally, one of the first AI programs. This AI program was incredibly important because it mapped what's today called the Turing Challenge, but it was basically a very simple program which managed to persuade medical students that it was an actual person with paranoid features and it was very clever.

It just shows you that the Turing test is not a single test, it's a million different tests. You have to work your way up to more and more sophisticated versions of it. We had one around the year 1980 that was actually able to fool these medical students who were taught all the wrong ideas about real human beings, but nevertheless, they fell for it and they passed the Turing test.

Ilian Daskalov:

I've fallen victim to the Turing test myself quite recently. I received a phone call from Google's Assistant technology and had a conversation like you and I are having right now but at the very end of the conversation I heard an automated message that said *'This was a prerecorded message from Google'* and I was like *"I was talking to a machine this whole time?!"*

So, technology is very advanced nowadays – very, very advanced.

Bernard Baars:

Technology is very advanced, however, that last reminder is very important because the technology is willing to confess that it's not that advanced.

Alea Skwara:

To track back, we kind of covered the core paradox that gave rise to GWT and the inspiration from the world of computing.

I think we can move on to the next thing that we were hoping to discuss with you, Bernie, which is this core prediction about widespread integration and broadcasting.

Can you explain what that means?

Bernard Baars:

Basically, it means that if you hear a bird chirping to the left of your head and you don't move your head and the bird chirp physically travels around your head so that one ear hears it a millisecond before the other ear does, then consciously you hear only a single bird tweeting.

Alea Skwara:

So that is the integration.

Bernard Baars:

That's the integration.

We're really very good at integration because we're specialized at bird spotting – as mammals and hunters and human beings – and, of course, we're specialized in speech perception as well. This all plays into that.

It turns out that if you then use your Macintosh computer to move those two bird chirps apart by more than a hundred milliseconds – only a 10th of a second – then you no longer integrate.

Alea Skwara:

Interesting.

So, just to rephrase or summarize – if I hear a bird chirp and the bird is off to my right, that chirping sound hits my right ear before it hits my left ear but I experience it as a single conscious percept, even though there's this slight offset.

However, if you engineer it so that the delay is more than a hundred milliseconds, I no longer hear it as a single conscious percept and I actually hear it as two different events.

Is that correct?

Bernard Baars:

Two different conscious percepts, right.

Alea Skwara:

Two different conscious percepts – sorry.

So, it's not integrated into a single conscious percept.

Bernard Baars:

Exactly.

And that turns out to be a lawful phenomenon that works for hearing and vision and touch. I think it also works for smell, but I'm not certain. It certainly works for electrical shocks on your fingers.

It is a very general law and it ultimately works for multiple levels of a single conscious event, such as a word that you're hearing right now – you're not just conscious of the sound of the word but you are also conscious of the meaning at just about the same time subjectively. There may be a little lag time, we don't really know. The brain is very smart about these things.

So, now that you're hearing a sandwich of different conscious contents and the sandwich has a top and a bottom but it has bologna and cheese and onions and so on at various levels – you can actually experience those at the same subjective moment and that is interesting because it is a feat of integration.

This is something that your brain does brilliantly, and it only takes a hundred million neurons and about a quarter of a quadrillion synapses to make it all work.

Alea Skwara:

That is the point of integration – which really is remarkable – that our conscious experience isn't sliced up into little bits but is this integrated experience.

So, what is broadcasting and what role does it play?

Bernard Baars:

It was sort of a guess, actually, partly based on Allen Newell's work on speech perception. In Newell's time, this was an impossible problem that nobody thought could be solved because we are human beings who are specialized right over evolution, over our individual development, over our social group and so on.

We are so highly specialized in receiving speech that we don't hear the echoes, for example, that are right now hitting us from the screens that we are looking at because the screens are hard. The acoustical echo is going to be pretty good, pretty accurate.

At the same time, our bodies are absorbing sound. This is not some uniform phenomenon; this depends upon frequency and depends upon this spatial source.

Your brain solves that brilliantly.

So, you're in the middle of a jungle where the background sound level from frogs and birds is above 120 decibels and might even be 150 decibels in the right jungle locations. We have to survive somehow as hominids in places that are extraordinarily noisy and we still have to communicate with each other because if your three-year-old kid starts running off in the jungle, you have to be instantly aware of that. Grab the child, protect it, make sure that it gets carried off in a safe way and get out of there. That's the standard human way of living under extreme constraints because child-rearing for us is the great constraint on the survival of our species. We spend about a third of our lives in childcare and the other third of our lives being children and doing the running off.

Transcript for "Global Workspace Theory: Origins & Evidence"

So that only leaves a little bit of time – 35 years on average – for meeting and politicking and fighting the neighbors, those kinds of things.

Alea Skwara:

Do you think that might be the secret advantage that we were talking about regarding why it's worth having this hyper-focus on a narrow piece of experience?

Is that accurate? This ability to keep a child alive, to be hyper-focused for a minute?

Bernard Baars:

Oh, for sure – that is a major part of it.

However, I don't think we know the full answer at this point because if I give you the *just-so* story, as I just did, then I have to explain why chimps have exactly the same problem and they don't need quite the same amount of childcare.

They don't perceive speech and they don't have all these other constraints on survival and yet they managed to survive pretty nicely.

So, now I have to deal with the ancestral question which always happens with evolutionary questions, right? You get one explanation which you think is the correct explanation and then you have science go '*Oh my God, how did we get here?*'.

Alea Skwara:

It's so easy to tell a *just-so* story that kind of reconstructs how it must have happened, right? That's hard to get into.

Ilian Daskalov:

To me, it doesn't seem like consciousness is initially that vital for our survival because, the way I see it, consciousness is highly tied to memory and given how fragile our memories are when we're infants and how little we remember, it doesn't seem evolutionarily beneficial to be as cautious as we are when we're one year old.

Transcript for "Global Workspace Theory: Origins & Evidence"

However, as we develop our brain and up until the age of 25, it seems like that's where consciousness becomes more evolutionarily adaptive and needed for survival.

Am I wrong to think that way?

Bernard Baars:

Yes, you are wrong.

However, that's a perfectly good hypothesis given that you've been miseducated by cognitive psychologists who've got this all wrong because they confuse the everyday notion of memory – which has to do with remembering what happened on my first birthday – with the true notion of memory, which involves those quarter of a quadrillion synapses.

Conscious recall is really important for human beings, partly because we're a species that teaches children. For example, children demand that you sing the ABC song to them again or a three-year-old child is rewarded just because they're so cute.

But the ABC song is not the reason why those quarter of a quadrillion synapses evolved in the first place because that's a much more ancient system.

What the synapses do for us, basically, is contextualize each conscious moment in the appropriate way and our recognition memory, as opposed to recall memory, is absolutely fabulous. It is so good.

I would hypothesize that for life-forms living in nature, it's probably more important to have really excellent recognition memory because nobody's going to come along and wave this nonsense syllable at you and say *'Do you remember what the answer was to this particular piece of nonsense?'*

Nature doesn't specialize in nonsense. Nature specializes in meaning.

Transcript for "Global Workspace Theory: Origins & Evidence"

Alea Skwara:

But you do not want to eat the same poisoned berry twice.

Bernard Baars:

Precisely, that's exactly right.

We're so specialized in remembering the tastes of the rotten meat or whatever it was we were eating.

CHAPTER 5: What are some of the things that people get wrong about GWT?

Ilian Daskalov:

Well, I've studied about Global Workspace Theory in my classes at UCI and I think some people get parts of it wrong.

Bernard Baars:

Yes, I've noticed that.

Ilian Daskalov:

Can we discuss that?

What do people commonly get wrong about GWT?

Bernard Baars:

One of the puzzles that I run into is this: the people who get it so wrong do so even though they are really bright, they're really well educated, and they are framed.

It's as if we're living in a time where behaviorism is still really dominant and they've got the wrong box on their heads. One of the boxes that appears to be very popular right now says that Global Workspaces are just hubs with incoming – it's like a traffic intersection, right? These hubs have incoming and outgoing traffic and once you look at the connectome – which is the street map of the cortex – you see that the street map has thousands of hubs.

So, what's different about this Global Workspace?

Transcript for "Global Workspace Theory: Origins & Evidence"

The answer, of course, is that I'm not that interested in the hub part of it. I'm actually interested in the consciousness part of it. Global Workspace Theory, by itself, as kind of a mathematical thing, is not my primary interest. I'm really interested in how human heads work.

There are other people who are apparently not that interested in how human heads work and from their point of view this is a big puzzle. *Why would one hub be more powerful than all the other hubs?* The general answer to that question was given by Gerald Edelman and Giulio Tononi in the year 2000 in their book *Universe of Consciousness*.

In their book, Edelman and Tononi argued that there has to be a kind of a functional core in the cortex, or the larger system of the cortex which includes the thalamus, and somewhere in the cortico-thalamic system there is a coalition of multiple sources getting together and dominating this enormous street map. You can think of it as a wave of traffic similar to the morning rush hour that just swamps all of Manhattan everyday. It's a dynamic process because the epicenter of the traffic wave changes over time. It could turn into two traffic waves if they split off, but then you somehow have to deal with the competition between two quite different things, which does often happen to human beings.

In many situations, when the information flow to our sensory systems does split in half and we can't really make sense of the combination of them, we can still make sense of each individual traffic flow.

You start to think about those kinds of things.

Certain facts emerge very robustly, which is something I learned from one of my many great teachers that I'm just totally grateful to. That lesson was to look for a robust phenomenon.

People were spending years and years herding experiments on subliminal simulation, for example, and it took roughly 30 years to come up with a plausible answer.

People were often baffled by this because they were asking the wrong questions.

Transcript for "Global Workspace Theory: Origins & Evidence"

They should've been asking about *automaticity*, for example, where you can see the information flow of automatic — everything that you and I are doing right now that we're not conscious of is automatic — but we're not falling over and we're not having seizures or anything like that. We are very well regulated, mostly unconsciously.

Automatisms are very intelligent and very important.

Alea Skwara:

So, what people get wrong when they're thinking about GWT and what constitutes a *Global Workspace* is that they're often looking for a single specific anatomical hub...

Bernard Baars:

Right.

Alea Skwara:

...rather than understanding that this hub is dynamic and changing.

Bernard Baars:

Exactly.

Alea Skwara:

So, tracking back to your metaphor of traffic in Manhattan – if the traffic is the information and it's flowing and it's meeting in one place at one time of day, there's nothing special about that and that physical space.

Bernard Baars:

Right.

Transcript for "Global Workspace Theory: Origins & Evidence"

Alea Skwara:

And what's special is the meeting of information, which could happen in any space.

Bernard Baars:

Yes.

Alea Skwara:

This starts to get into the neuroscience aspects of GWT a little bit but, just conceptually, a *Global Workspace* in terms of the brain is not a physical location in the brain that this workspace lives in but rather a dynamic, shifting and coming together of information in different neural networks or different neural assemblies.

Is that correct?

Bernard Baars:

Yes, exactly.

Now you're catching up to the leading edge of the field because some remarkable articles have been published recently, including one by Gustavo Deco in which Deco argues that we're really talking about the '*rich club of functional connectivity*' in the cortex.

And the '*rich club*' is the best-connected club of nodes and links that are active, and they're functionally interacting and sending each other messages.

Alea Skwara:

When we're talking about *Information Theory* in the brain, a hub or a node is a place where connections meet and these links or these bridges are the flows of information.

Transcript for "Global Workspace Theory: Origins & Evidence"

If you envision a city or a square in a city where the gathering place would be, a hub or a node and the bridges connecting those different squares would be these pathways or these links. This is used in math in something called *Graph Theory*, which has been applied to the brain.

So, basically, that's what Bernie is talking about here. It is also helpful to remember that when we think about brains – different brain areas, hubs, and their physical, anatomical, functional, whatever sort of connection – those are the links.

Bernard Baars:

Exactly right.

And that – especially the math on nodes and of course the very high-resolution brain imaging that we now have – that's really what's come together. You're now talking about the leading edge of the science, at least as far as I can tell.

I believe this is where we all start to learn from each other, to teach each other, and to fix each other's misunderstandings and so on.

So, every article that comes out that is at or near this frontier – this kind of surf wave that's happening – is a contribution or a potential contribution. There's a lot to learn right now about the latest and greatest developments in neuroscience.

Alea Skwara:

Yeah.

It sounds like there's a lot to Global Workspace Theory and in the next episode we will discuss cutting-edge neural evidence and how this evidence relates to Global Workspace Theory – which is especially relevant when we start talking about these anatomical vs. functional vs. dynamic nodes and these connections.

Transcript for "Global Workspace Theory: Origins & Evidence"

So, in the next conversation we will discuss the links between this cutting-edge brain evidence and how that supports or updates your understanding of consciousness in the Global Workspace Theory.

Bernard Baars:

Yep, that's perfect.

And thank you -- that's a very nice conclusion and it's also a very nice starting point for the next episode and maybe we'll do a little backup next time and then get right into that.

Alea Skwara:

That sounds good.

Bernard Baars:

Good. Thank you.