

Dennett (and Searle) Discussion debate on the Philosophy Forum

12/2004

All posts by John Donovan unless noted otherwise

Quote:

Originally Posted by **Minty**

If materialism is correct, is it possible that the psychological explanations for our behaviour - eg the intention to get a drink because one feels thirsty - are not reducible to motions of elementary particles or atoms inside our brains? In other words if materialism is true, is it possible for the whole brain to influence the behaviour of the parts comprising it, rather than exclusively the other way round?

Minty,

What then "influences" the brain as a whole? If we are trying to explain brain processes we need to start with something less than the brain. Consciousness must be explained by parts of the brain that are themselves less conscious than the whole.

However, I'll grant that cognitive science recognizes both specialist and (spatially limited) global brain processes, so in a banal sense you are correct, but I think a more important point you are trying to make is Dennett's rejection of "greedy reductionism". Yes, we generally shouldn't (even though in principle we could) try to explain consciousness as a chemical reaction, just as we generally shouldn't (even though in principle we could) try to explain bio-chemistry using quantum physics.

Just as one should always program computers in the highest level language that gets the job done, we should explain brain processes in the most high level (broad) and powerful theories **that gets the job done**.

The real problem is that most philosophers are going about the job backwards. They are assuming that their folk intuitions about consciousness are "obviously" accurate accounts of what is going on in their brains and then try to find justification for those introspective notions by invoking ideas along the lines of intrinsic properties, non-material properties, soul-stuff, and other ghost in the machine ideas.

Dennett and some others instead start from what we already know from cognitive science (which is considerable) and then using ideas from evolutionary psychology and AI, try to come up with a theory of the mind that actually explains those observations, but also explains these intuitively "obvious" aspects that we introspectively observe. That we have these notions is not in dispute. That they are accurate depictions of the actual processes inside our brains is.

Here is a philosophy student term paper that although is not perfect, is well written and distills down much of Dennett's ideas in a very readable way. It's short and deals head-on with many of the complaints by Dennett's critics. I think that everyone in this forum interested in the mind should read this so that we can move forward in ways we have been unable to do so far.

<http://www.stanford.edu/group/duali...pdfs/newman.pdf>

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

<http://forums.philosophyforums.com/...ead.php?t=12459>

Let us know what you think.

The Chinese Room, Qualia and the Zombie are all very clever appeals to our folk intuitions that purport to

demonstrate that consciousness will never be explained by merely mechanical processes. When I first read about the Chinese Room Argument I had to admit it was clever, but it wasn't until later when I starting reading more that I realized in how many different ways it was fallacious. But before I start on the specifics let us merely note that all such attempts to place a naturally observable phenomena outside the realm of scientific methodology have historically been utter failures. The fact that all humans have a preprogrammed storehouse of sometimes but not always reliable folk intuitions about nature, needs itself to be explained in any cognitive theory of the mind. But these all too fallible intuitions should not be counted as accurate evidence in scientific investigations. The unexplained is not the inexplicable.

I would also make the point that like brain in the vat type of thought experiments, all such philosophical ideas deal with such an abstracted presentation of facts in which simple realistic parameters like evolutionary history and the ability to learn are automatically excluded from consideration. In the Chinese Room not only is the situation presented without any history, but it is static, unable to move forward in ways that would naturally occur. This will become apparent once one contemplates the likely origins of communication and language. The earliest imaginable instances of affirmative, negative and interrogative grunts and other sounds could only have achieved meaning and usefulness in environments where social interaction occurred and correlation could be established.

The Chinese Room is not a new idea of course. It probably first originated with Leibniz as seen in the quote below:

"Moreover, it must be confessed that perception and that which depends upon it are inexplicable on mechanical grounds, that is to say, by means of figures and motions. And supposing there were a machine, so constructed as to think, feel, and have perception, it might be conceived as increased in size, while keeping the same proportions, so that one might go into it as into a mill. That being so, we should, on examining its interior, find only parts which work one upon another, and never anything by which to explain a perception. Thus it is in a simple substance, and not in a compound or in a machine, that perception must be sought for." (Leibniz, *Monadology*, 1714arag. 17[Latta translation])

Thus the call for mind stuff, soul, non-material properties, quantum fluctuations, whatever. Although there are many detailed philosophical ways in which the Chinese Room Intuition fails on close examination, I think that the bird's eye view or system problem most clearly breaks the intuitively hypnotic spell. The cleverness of the intuition is essentially the use of a "person" doing the translating without knowing what they are doing outside the room (skull). But the problem is that the "parts" of the mind can't be minds themselves- the mind has to be explained by means of successively less intelligent and often more specialized parts that in and of themselves are not intelligent minds. Yes, our visual cortex is pretty fancy, but no it is not a little man with a video camera watching the world from inside our heads. In fact the visual subsystems along with many other perceptual systems (along with a few behavioral processes) have been quite well explained in merely mechanistic processes.

The point is, the parts of my brain don't have to know what or why they do what they do, any more than the cells in my heart or lungs have to know what or why they do what they do. The Chinese Room, by ignoring the evolutionary history of the brain and perceptual learning, fails to demonstrate that mere mechanical processes cannot bring about the behavior of minds. Dennett puts it like this:

"Might it be that somehow the organization of all the parts which work one upon another yields consciousness as an emergent product? And if so, why couldn't we hope to understand it, once we had developed the right concepts? This is the avenue that has been enthusiastically and fruitfully explored during the last quarter century under the twin banners of cognitive science and functionalism—the extrapolation of mechanistic naturalism from the body to the mind. After all, we have now achieved excellent mechanistic explanations of metabolism, growth, self-repair, and reproduction, which not so long ago also looked too marvelous for words. Consciousness, on this optimistic view, is indeed a wonderful thing, but not that wonderful—not too wonderful to be explained using the same concepts and perspectives that have worked elsewhere in biology.

Consciousness, from this perspective, is a relatively recent fruit of the evolutionary algorithms that have given the planet such phenomena as immune systems, flight, and sight."

I might add, that what Leibniz didn't know about, but we do today is: computers. Everyone will agree that there is nothing soulful or non-material about computers, yet as simple as they are compared to the human mind, clearly they can exhibit mindlike behavior and sometimes behavior far more complex in certain specialized examples than our own minds. Will we ever build an inorganic human brain? I think not. We could, in principle, build a perfectly functional human kidney atom for atom but why bother? It's easier to grow one.
By Probeman (John Donovan)

I'm not a hundred percent certain, but I think Monroe here is the first person to actually address something in the book itself, even if very indirectly! Perhaps I am only dreaming, but that seems to be the case. At any rate, Monroe's post was interesting and thoughtful, so that and the fact that this is my favorite thread are my justification this time around for being long-winded.

Quote:

Originally Posted by **Monroe**

It takes reports of mental events and then goes to see if there's anything we can find from in the third person scientific perspective that matches up to the things as described in the reports. Then it concludes that these reports were really about these things.

Well, as I will repeat later on, perceptual reports can also be interpreted as being about things in the world, or at least the things which caused our brain events. But in the highly unusual context of experiments in cognitive science, we sometimes switch the meaning of "about" to refer to brain events. This especially is true in the case of illusions. More on this later, but it gets slippery.

Quote:

Originally Posted by **Monroe**

This automatically discounts the view that mental events are private, and that descriptions of them are about things that are only privately viewable.

Viewable? From inside the Cartesian Theater, no doubt, where a show goes on for the Self to watch and enjoy, then describe for the benefit of scientists.

So, does the brain provide the stage? Manage the set? Dress the actors? Provide lighting? And the self viewing this private show: if you aren't proposing that it is a soul interacting with the brain via the pineal gland (or whatever), then the brain must be managing it, too. So the brain has to perceive the show, process what it means, then process the reactions the Self will have to it. . .and give the show, too. That's a lot to do.

This is deliberately described to make the whole affair look silly, of course. Feel free to point how your own view is different.

Essentially, explaining experiences as events that are "privately viewable" creates your standard infinite regress: a man views a red apple, and this is explained by a something (?) which "views" a private experience. Of course, this is no explanation at all, for now we must ask the essentially the same question of this thing inside viewing a private experience: what is this process, and how does it get explained scientifically?

Good explanations mean explaining why people give the reports they do and why they talk about mental states in the way they do. Some of these tasks are handed off to neurology, some can be explained by looking at how our intellectual culture has trained itself to discuss mental states, and there are some vocabularies which owe a little to both.

I used to be convinced that Cartesian metaphors such as the idea of things “privately viewable” were completely an artifact of Western culture and the pernicious influence of Christian dualism handed down through Descartes himself. But I’ve heard that recent work from neurology suggests that these metaphorical habits are actually hard wired into us and begin to manifest in early childhood, long before anyone is exposed to Descartes. I haven’t actually read the evidence, only heard about it third hand.

Regardless of whether this way of talking is owed to neurology or tradition, the way you stated the issue merely begs the question. You are assuming without argument that speaking of mental states as things “privately viewable”—with all the metaphysical baggage that this vocabulary entails—is a FACT to be explained rather than an alternative conjecture in need of its own support. It is certainly a fact that this is how we all talk about mental states. But the implications of this way of talking are another matter entirely.

Maybe we’re too early in the chapter summaries for this to be clear, but to me the whole discussion of Shakey should have been enough to show how and why the heterophenomenological method is neutral with regard to such claims. Until we’ve done a full scientific investigation, we don’t know which version of Shakey is closest to our own case. We could be completely confabulating everything, or we could be talking about something real, or a mix of the two. Evidence and not armchair philosophizing is what we need to decide.

Quote:

Originally Posted by **Monroe**

The third person scientific search, by definition, can never find these things, and Dennett's method then leaves us to conclude that they don't exist at all.

Again, you beg the question by assuming these things must exist and possess all the untouchable-by-science properties your intuitions tell you they have. If you want to insist on deep mental properties which are in principle unrecognizable by any form of science, you are free to have this faith. But faith is all it is. CE confines itself to views on consciousness which are evidence-based, and in this sense your views are being excluded by his approach, and by science generally.

This is destined to continue forever until people who think as you do can produce some sort of experiment that would show how the world would be different if you were right. If your position cannot be confirmed or falsified by any conceivable evidence—it you can’t show how it matters--then science doesn’t need to pay it any attention. It can never be part of a scientific approach to the mind, nor can it be used as a criticism.

Quote:

Originally Posted by **Monroe**

There is a certain important subset of honest reports about how things seem that cannot be mistaken: Reports about how they seem. (i.e. If someone says, "This is the way it seems to me: Blahblah...") Reports about phenomenal consciousness are mostly about this. How would Dennett's approach handle this?

Heterophenomenology grants everything you want to reports about how things seem. It naturally assumes full honesty and integrity in subjects. Honest reports about how things seem are, one could say, not falsifiable. (Richard Rorty once called this the “hallmark of the mental.”) But you possess this sort of “Papal infallibility” only about how things seem, not about how things are.

For instance, you might report that it certainly seems to you as if your entire visual field is filled with colors and has no blank or invisible areas. That is how things seem to you, no doubt about it.

But that isn’t the way things really are. You in fact cannot see most colors at the edge of your visual field, even

though it seems as though colors exist all the way to the periphery. You have a blind spot front and center where no information from the world gets through, even though it seems as if there is no such blind spot. Both of these facts about how things really are can be confirmed by simple experiments done at the office or home, as we touched upon earlier in the thread. And when these experiments are done, subjects come to change their minds about how things seem.

This is all because your brain is designed by evolution to make instant and unconscious judgments about the environment based on meager evidence. When you report how things seem, you are reporting the content of those judgments. Whether that content is true or not is another matter. It takes very unusual experimental contexts to expose flaws in these leaps of conjecture by your brain.

Quote:

Originally Posted by **Monroe**

He would find some brain process and say, "This is the way seeing red seems to you." Then the subject would say, "No, it's not. That's something entirely different. It doesn't resemble red experiences at all."

There is nothing in Dennett's work to suggest he would endorse the absurdity of pointing to a brain processes and saying, "This is the way seeing red seems to you." All anyone could point to would be something in the heterophenomenological record where a subject reported how something seemed to her.

Quote:

Originally Posted by **Monroe**

Even if the brain process can be shown to be the cause of red experiences, the subject is not referring to the brain process as such, but the experience to which it gives rise.

Ah, but Dennett is most certainly NOT attempting to argue that brain processes cause red experiences.

This is what your quote tacitly assumes: first, an event in the world happens, say a red apple coming in view of the subject. This causes an event in the brain. This event causes another event entirely—an "experience". That is, I interpret you to be saying that the experience is not a brain event.

Now normally, I would say that if a subject reports, "I see a red apple now," she is referring to an event in the world occurring roughly in front of her, not a brain event. This is what you could have meant by "the subject is not referring to the brain process as such". But apparently you didn't, because you go on to say that the subject is instead referring to the experience. Oh. She isn't talking about the apple? What happened to the apple?

It gets worse. The subject reports seeing the apple by making some sort of physical motion—hitting a button, speaking a sentence, whatever. This can't happen without the correct signals firing from the motor centers of the brain. And the motor centers aren't going to just do this spontaneously—they are going to be caused to do their thing by other physical events, events elsewhere in the brain. In fact, short of the discovery of utterly magical events uncaused by anything (sort of like a miniature poltergeist tickling neurons), there is going to be a complete causal chain leading from photons bouncing off the apple into the subject's eyes, to the subject's report.

As far as I can see it, you have three ways to go.

You can continue to posit non-physical experiences but grant them no physical effects in the world. This is epiphenomenalism. In this case, you couldn't say that subjects are EVER referring to experiences when they make reports, because all the causes of their reports are the physical events in the brain, and the experiences drop out as unimportant—you'd make the same reports whether brain events produced experiences or not. (Think of Wittgenstein's beetle box.)

You can continue to posit non-physical experiences but grant them physical consequences which are necessary for subjects to make reports, such that if experiences did not have these causal powers, no one would ever report them. But there is no evidence of non-physical causes in the brain, which would literally be a form of magic. At least this approach is open to scientific confirmation, though.

Or you can treat mentalistic vocabularies, which produce a metaphorical space “inside” where experiences happen, as just that: metaphors, habits of talk. And here is where I think it’s important to pay attention to the story of Shakey in the chapter summaries—when asked, different versions of Shakey can report with different levels of competence about what happens “inside” when they process visual stimuli and monitor what they do. The point of heterophenomenology is to collect the reports and other third hand evidence, then decide which version of Shakey we are most like.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Monroe**

It's something empirically confirmed. We know that we have our own inner lives, by direct experience. We assume others have similar things (by inference to the best explanation I suppose). And we also observe that we do not have access to other people's conscious minds.

We have inner lives to be sure, but my larger point was that you are making question begging assumptions about what saying that means, then using those undefended assumptions to critique a theory which doesn’t even recognize them as true. I don’t think you realize how virtually every sentence in the paragraph of yours that I quoted has interpretations that carry metaphysical baggage—there are a lot of theoretical implications smuggled in as so-called “facts”.

It’s one thing to call mental states private, and quite another to mean by this that mental states cannot be addressed by third person science. The particular form of privacy you are endorsing is not something that could be empirically verified, but is rather a theoretical position which has been under contention from Wittgenstein on. So citing it as a “fact” that CE’s heterophenomenology can’t handle is question-begging to the extreme.

When we get to the parts of the book that begin to discuss qualia in more detail, this will be clearer.

Quote:

Originally Posted by **Monroe**

In light of Hume's criticism of causation, what scientific evidence shows what kind of causes of brain events there are?

I’m really not sure what you are asking here. We’ve been watching brain events via various kinds of scans for decades, and mapped countless functional zones within it. We know how neurons cause and respond to biochemical events. There is nothing in Hume that would have the slightest bearing on this subject.

Quote:

Originally Posted by **Monroe**

Would a philosophical argument that mental events are not reducible to physical ones, plus whatever kind of indirect evidence justifies beliefs about causal structures showing that the mind and brain are causally tied, be evidence for nonphysical causes in the brain?

Again, I can’t really follow your question. Dennett’s theory of consciousness is largely one that is anti-reductionistic to begin with. And I am aware of no evidence whatsoever in favor of there being “nonphysical” causes anywhere, let alone in the brain. A nonphysical cause would literally be magic.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Banno**

The answer given above (#59) to the systems approach is sufficient to show that syntax is not sufficient to explain meaning.

But the evolution of language shows us that syntax in a behavioral context is enough.

Searle is a realist up to a point. That point is that the human brain is (somehow) more than a "mere" mechanism and therefore no intelligent consciousness machine is possible. His Chinese Room "thought experiment" is an attempt to "prove" his intuitions on this point. But the systems problem does refute his argument because Searle's "thought experiment" actually proves too much. Searle in response to the system problem says it is ridiculous to say "that while [the] person [in the room] doesn't understand Chinese, somehow the conjunction of that person and bits of paper might." This is like saying that since each cell in the brain doesn't understand Chinese, that a Chinese brain can't understand Chinese either.

Here is a summary of this argument (from <http://www-users.york.ac.uk/~twcs1/C&C/Lecture%203.pdf>):

"The man is just part of the system, he is playing the role of CPU in the Turing Machine. The sufficiency claim does not say that there is some TM such that every part of it understands Mandarin, only that there is some TM such that the whole of it understands Mandarin." Searle responds to this by saying that the man could internalize all the rules etc and then: 'If he doesn't understand, then there is no way the system could understand because the system is just part of him'.

But: (a) this is even less of a genuine possibility than the original example; (b) Searle's official reason for rejecting it is incoherent, since the thought-experiment requires him to be part of the system, so it cannot also be part of him; (c) the man would not know in advance that he could understand Mandarin, but he might come to believe it of himself; (d) if this were possible, it would give us reason to doubt the unity of consciousness

As I said in the beginning, basing philosophy on our heartfelt and sincere intuitions of "how it SEEMS to us" (folk psychologies) is the wrong way to explain the mind. We need theories based on the evidence and then we need to see if our intuition notions can be explained (away).

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

Computers don't have souls, so how could they become conscious?

Souls? Well, so far as we can tell, neither do humans. Is your argument simply based on intuition?

Quote:

Originally Posted by **Minty**

Searle claims he is a materialist (although I must admit he sounds like a dualist from 2nd hand reports I've read). Now, if he is a materialist, then *necessarily* he must believe that an intelligent consciousness "machine" (as in something which is made or produced) is possible; even if this amounts to creating an exact duplicate of a human being. It's just that he is maintaining that the execution of algorithms doesn't somehow equate to, or produce,

consciousness.

The point is that even if Searle is correct in this unsupported assertion, his "Chinese Room" example does not "demonstrate" it.

As Dennett has said- never underestimate the power of algorithms. Especially heuristic (self-modifying) algorithms. Here is a quote from a review (<http://www.scientificexploration.org...-2/dennett.html>) of "Darwin's Dangerous Idea"; a book I highly suggest you read:

"Much of the hostility both toward evolution and an engineering approach to the mind rests on the fear that such reasoning will subvert our sense of self, drain life of meaning and purpose, and explain away our very minds. This hidden agenda of fear, Dennett argues, misdirects scientific debate about evolution. Behind the hot-tempered controversy, the announced revolution that changes little or nothing, and "the tremendous -- and largely misguided -- animosity" to Darwinian accounts of language and the human mind, Dennett detects a failure of nerve. It is not that the "Modern Synthesis" is in dispute, it is rather that its consequences are too hard to bear.

Dennett wants to cut through the smoke screens of avoidance, confront and disarm the animosity, and work out answers to responsible objections. In this regard, he singles out a number of distinguished thinkers for special criticism: paleontologist, Stephen Jay Gould; linguist, Noam Chomsky; philosopher, John Searle; and mathematical physicist, Roger Penrose. Gould's anti-adaptationism and insistence on "radical contingency" and "punctuated equilibrium," Chomsky's suggestion that evolutionary theory has as yet little to say about language, Searle's argument that only human minds have "original intentionality," and Penrose's conviction that our ability to "see" and "understand" mathematical truth is non-algorithmic -- all these positions, Dennett suspects, represent attempts to refute the idea that evolution is an algorithmic process and to shield the mysteries of free will, language, and the mind from Darwinian mechanisms.

Each of these thinkers, Dennett claims, betrays a yearning for "skyhooks," when they should be looking only for "cranes." Skyhooks are, in Dennett's inventive terminology, impossible, imaginary devices that spring the frame of mechanical, algorithmic explanation. They are "mind first" forces or processes, moments of special creation, exempt from, and discontinuous with the mindless mechanics of design. Cranes, on the other hand, are the real lifters in the evolutionary process. Cranes are complex intermediary mechanisms that arise from the process of evolution itself, and in turn, speed the process along by promoting the development of still more complex structures. In Dennett's view, God is a skyhook; sex is a crane. "

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

So all you have to do to show he is wrong is to produce an example of a syntactical system that supplies its own

meaning...

No, the burden is on him to show that his intuitions on this subject have any validity at all. Since so far the science shows no evidence for anything except algorithmic processes. Cranes, not skyhooks as Dennett would say.

Quote:

Originally Posted by **Banno**

I find your contempt for what you call "intuition" rather odd. Would it not be the case that any theory of mind that was seriously at odds with the way in which we perceive ourselves to think would have to be considered wanting, perhaps even falsified? How else would we judge the accuracy of a theory of mind apart from its ability to explain what it is like to have a mind? What else could count as evidence?

Just as the heliocentric theory did not utilize the intuition that the Sun SEEMS to revolve around the earth (in fact it required much rational effort to overcome), any scientific explanation of the human mind will eventually require us to give up our heartfelt intuitions on this topic as well. My contempt for intuition is not universal- for the topics our intuitions were actually evolved for (mate selection, personal safety, etc) intuitions should be given much attention- but for scientific problems, we have already seen that they are simply not reliable sources of knowledge.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Probeman, one ought not throw away one's intuition without good reason – as the history of science shows. Searle clearly points to an aspect of the mind that is not explained.

You have not given an account of how intentionality arises from a heuristic algorithm. You seem to take it as an article of faith that somehow it just will arise. This is neither rational nor scientific. Furthermore, you seem unaware of the lack of support for your faith. I can only assume that you have reason for your belief that you have not yet presented.

Each of us has direct experience of intent. To describe this experience as an "intuition" is to try to be rid of it by calling it names. Present an explanation of how this arises in an algorithmic system, and you will have made your point. If you cannot present such an argument, explain why your intuition is more valid than Searle's.

Your language is so skewed to your intuitions that it doesn't even make sense to me. To say that we "experience" intent is simply an appeal to intuition. You want me to explain your intuitions of intentionality and how they "arise" from algorithms, but this presupposes that intentionality is a physical property when it is merely an arbitrary (though often useful) description of certain types of outcomes. I am not going to provide a reductionist account of intentionality "arising" from algorithms because I think the whole notion is question begging. Besides that, it would be inappropriately reductionist. Like explaining how water feels wet using quantum mechanics, the explanation would be difficult and tedious beyond practicality. You are essentially claiming that because water is wet, that atoms must have intrinsic properties of "wetness"? All Dennett is saying is that intentionality is an emergent observable behavior in certain systems. The subsystems themselves are not intentional in the same way, but are intentional in their own way. (e.g., the heart "tries to keep up" with the oxygen level needs.) And so it proceeds down to the level of the cell and thence to atoms.

The fact (and I agree it is a fact) that intentionality SEEMS to be an intrinsic property of living things does not demonstrate that it really is an intrinsic property of living things. Dennett argues that intentionality **as a behavioral description** is appropriately applicable to many systems, including organisms (from the amoeba to humans) and that it is also useful in describing the "behavior" of even some non-living entities like chess playing computers, but it is not an intrinsic property of certain kinds of objects.

So far as science can tell, living things are composed of the same atoms that non-living things are composed of. Your thoughts today are implemented on last week's potatoes. The only differences appear to be in the way the atoms are arranged.

Here is link to a paper called "Evolution, Error and Intentionality". You should read the whole paper to understand the problem. If you want to see in detail how algorithms can explain complex behavior and how our narrative sense of self can be produced, I suggest Consciousness Explained by Dennett. If you want to really understand this issue as opposed to simply reinforcing your intuitions, it's going to take some effort on your end.

<http://ase.tufts.edu/cogstud/papers/evolerr.htm>

By Probeman (John Donovan)

Quote:

Originally Posted by probeman

As Dennett has said- never underestimate the power of algorithms.

Quote:

Originally Posted by **Minty**

Hee hee, I'll try not to. Amazing they could be capable of such deep magick ie produce consciousness

Magic is not involved- these scientific ideas are exactly opposed to our natural tendency for intuitively magical and supernatural stone-age thinking. But I agree- it is amazing- as are the many unintuitive aspects of the natural world- if only you took the time to learn about them instead of relying on your emotional introspections.

By Probeman (John Donovan)

Quote:

Originally Posted by **Nonblack Raven**

Now it seems to me that Searle's story about the Chinese Room is interesting because it suggests a way in human minds are quite different than syntactic algorithms. To prove to me that Searle's story is a bad intuition, we would need both an AI program as good as the Chinese Room is imagined to be (which we have not yet got, and may never have) **and** a demonstration the neurophysiology of the human brain is functionally just like the Chinese Room, which we also have not got.

Science doesn't have prove anybody's intuitions wrong (regarding souls, gods, intrinsic intentionality, or whatever)- all it has to do it come up with a natural explanation that fits the available evidence.

As part of this effort I have come across a paper by William Calvin that describes from a neurobiology perspective how Darwinian processes within the brain itself could create consciousness. Very interesting:

<http://www.williamcalvin.com/1990s/...onscStudies.htm>

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Yes, sir, and I'll have a report in on Monday, Mr Probeman, sir! Probeman, your posts are, increasingly tedious. Again rather than present Dennett's arguments you have referred us to a long article – the same approach you took in another thread. It would be much more entertaining if you at least presented a summary of the argument, or paraphrased it.

Oh, you want entertainment! I thought you wanted knowledge. In that case- stick with your intuitions- they seem to entertain you well enough.

Quote:

Originally Posted by **Banno**

But I suppose if you feel the need to rely so heavily on authority, you should be allowed. It does make me wonder if, that you are not able to present the arguments in your own words indicates that you really do not understand them.

I'll not respond to your silly taunts except to say, that like all scientists, I rely on evidence. Citing prior work is not "authority" (it can and often is cited in opposition as well) but is simply part of the scientific method of which you seem to be rather unaware of.

I've paraphrased enough for you already, but since you apparently only want to see your intuitions reinforced and you also refuse to attempt learn anything new that might challenge your beliefs, I'd just as soon not waste too much more time on you. At least until you are seriously interested in understanding why your intuitions are simply never going to explain anything useful to understanding.

Quote:

Originally Posted by **Banno**

Are you wishing to claim that you do not have an experience of intention? Then why are you writing these posts?

Because like all organisms and many artifacts I exhibit behavior than is well described by the "intentional stance". But more to your point- because I get pleasure from teaching- it's my job in fact.

Quote:

Originally Posted by **Banno**

Apparently you do not agree, as non-black raven pointed out so succinctly in post #80, that your position is simply not demonstrated. That it is as dependent on intuition as is Searle's. This seems to me to be the pivotal point of the argument – not that you are wrong, but that your case is not demonstrated.

The analogy (again, you have argued indirectly, but as if your argument were definitive) with QM and water is flawed. A better one would be that I am asking you to demonstrate how a spectrum is explained using QM.

Wetness is incidental to the nature of water; intention is *the* defining characteristic of the mind.

Wetness incidental to the nature of water? It's an analogy, that's all. But you obviously missed the point which is that the human mind- as arguably the most complicated object in the universe, is probably not going to ever be explained in a completely detailed reductionist manner- nor does it need to be. Just as we can confidently assume that all the properties of water are ultimately simple sub-atomic processes (without being able to "demonstrate" it), so the mind is ultimately a number of simple chemical processes, even though it can't be "demonstrated" as you would like.

As for intention, intention is a defining characteristic of all self-organizing and/or algorithmic processes- mind is not necessary for basic intention. And if you had actually read the Dennett link on evolution and intentionality carefully you might start to get a small inkling of what I'm talking about.

Quote:

Originally Posted by **Banno**

As for the chapter you cite, I've had a quick look, and when I have some time I might give it a read. But if the

basic argument is, as it appeared, that human intentionality is itself derived from our genetics, and so derivative in the same way a computer's intentionality is supposed to be, it seems to be a flawed argument. Evolution is not a teleological process, and so genes simply do not have some inherent purpose. That they are described as *wanting* to survive is anthropomorphism.

Good! I'd take a much closer look (re-read it several times and think about it carefully) if you are really interested in understanding this issue. Yes, evolution is not teleological, science does not even recognize teleology. From the scientific perspective, teleology is merely the all too human attempt to provide meaning and purpose for things that have no intrinsic meaning or purpose. That idea scares many folks and that is one possible reason why our intuitions may have some evolutionary selective advantage. After all, if one is wandering in the desert for 40 years, it might help one to believe that there is a purpose for it!

But genes do indeed have a non-teleological "purpose": replication. They may not be aware of that purpose (as you said it's anthropomorphic to say that) but that is all they "do". In any case our much evolved "purposes" and our genes "purposes" don't have to exactly coincide (see the section in the link on cryo-preservation of a person in a robot vehicle). The two generally do for obvious reasons, but you can (for just one example) decide to skip reproduction by using birth control (it may not always be easy in some situations especially for the younger of us!).

"Our" intentions have evolved far beyond the original "intentions" of our genes through the "cranes" of language and culture, though we are still very closely tied to them, e.g., the debate over gay marriage and abortion.

This is an amazing subject, but to understand it you might have to do some work. I've been reading about it for 30 years, but I still have so much to learn. I will say that the appreciation I've gained for how science tackles difficult (and unintuitive) questions has been very worthwhile.

If you re-read the link:

<http://ase.tufts.edu/cogstud/papers/evolerr.htm>

and return with some specific questions, I'll try to answer them as best as I can. Though they might be better posted in the Dennett discussion thread. In fact in case you missed it, here is another very related post (to follow) by Faustus for another poster that might help. It explains another common intuition that appears misplaced as well (experience).

I'll just post a link to Faustus' very well written comments:

<http://forums.philosophyforums.com/...3&postcount=604>

By Probeman (John Donovan)

The roots of Intentionality

Here is a concluding quote from the Dennett article that gets to the root problem with ascribing intentionality as a fundamental essence as opposed to an emergent property. The whole paper is worth reading for the thought experiments and intuition pumps that Dennett provides.

"Certainly we can describe all processes of natural selection without appeal to such intentional [stance] language, but at enormous cost of cumbersomeness, lack of generality, and unwanted detail. **We would miss the pattern that was there, the pattern that permits prediction and supports counterfactuals.** The "why" questions we can ask about the engineering of our robot, which have answers that allude to the conscious, deliberate, explicit reasonings of the engineers (in most cases) have their parallels when the topic is organisms and their "engineering". If we work out the

rationales of these bits of organic genius, we will be left having to attribute--but not in any mysterious way--an emergent appreciation or recognition of those rationales to natural selection itself.

How can natural selection do this without intelligence? It does not consciously seek out these rationales, but when it stumbles on them, the brute requirements of replication ensure that it "recognizes" their value. The illusion of intelligence is created because of our limited perspective on the process; evolution may well have tried all the "stupid moves" in addition to the "smart moves", but the stupid moves, being failures, disappeared from view. All we see is the unbroken string of triumphs. When we set ourselves the task of explaining why those were the triumphs, we uncover the reasons for things--the reasons already "acknowledged" by the relative success of organisms endowed with those things.

The original reasons, and the original responses that "tracked" them, were not ours, or our mammalian ancestors', but Nature's. Nature appreciated these reasons without representing them. And the design process itself is the source of our own intentionality. We, the reason-representers, the self-representers, are a late and specialized product. What this representation of our reasons gives us is foresight: the real-time anticipatory power that Mother Nature wholly lacks. As a late and specialized product, a triumph of Mother Nature's high tech, our intentionality is highly derived, and in just the same way that the intentionality of our robots (and even our books and maps) is derived. A shopping list in the head has no more intrinsic intentionality than a shopping list on a piece of paper. What the items on the list mean (if anything) is fixed by the role they play in the larger scheme of purposes. We may call our own intentionality real, but we must recognize that it is derived from the intentionality of natural selection, which is just as real--but just less easily discerned because of the vast difference in time scale and size.

So if there is to be any original intentionality--original just in the sense of being derived from no other, ulterior source--the intentionality of natural selection deserves the honor. What is particularly satisfying about this is that we end the threatened regress of derivation with something of the right metaphysical sort: a blind and unrepresenting source of our own sightful and insightful powers of representation. As Millikan (forthcoming, ms. p.8) says, "The root purposing here must be unexpressed purposing."

This solves the regress problem only by raising what will still seem to be a problem to anyone who still believes in intrinsic, determinate intentionality. Since in the beginning was not the Word, there is no text which one might consult to resolve unsettled questions about function, and hence about meaning. But remember: the idea that a word--even a Word--could so wear its meaning on its sleeve that it could settle such a question is itself a dead end....

We cannot begin to make sense of functional attributions until we abandon the idea that there has to be one, determinate, right answer to the question: What is it for? And if there is no deeper fact that could settle that question, there can be no deeper fact to settle its twin: What does it mean?"

By Probeman (John Donovan)

Quote:

Originally Posted by **NoSoul**

How would we get people to accept highly intelligent AI's with something like the modicum of humanity & dignity we now try to extend to all humans & many animal species?

I don't know. But if we can convince people that we should judge all such entities (including ourselves) by both their actual capacities and their demonstrated behaviors, as opposed to their external appearances and/or imagined metaphysically "intrinsic" or "essential" properties, we will have taken a large step in right direction.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

This is a philosophy forum, not a science forum. I'd much rather you presented your own summation, rather than the references to articles or the long quotes that you rely on. And yes, I do come here for entertainment. I'd hoped that you might be able to provide some insight into the Chinese room, and hence my invitation for you to join.

So you only will accept philosophical explanations for your intuitions? Well then, since it is obviously intuitive that the Sun circles the Earth, then should I assume that you are a geocentrist? Because that's what Aristotle thought. Look at it pragmatically- do you want to understand consciousness or simply confirm your intuitive beliefs? If it's the former then you might have to turn to science to get your explanations. It's your choice however.

I already explained the "systems" argument in some detail and cited additional references for you to read which explain how the Chinese Room is merely a metaphor that ends up describing a system that actually "understands" Chinese as much as any Chinese person. Do the atoms in a Chinese person's brain understand Chinese? No. Do the parts of the Chinese person's brain understand Chinese? No. Searle's claim is like saying that because atoms don't photosynthesize, then plants can't photosynthesize.

Consider this for some historical perspective. Before the 20th century many philosophers were convinced that living organic tissues must have some vital "essence" (not unlike Searle's intrinsic "understanding") for them to actually be "alive". See Vitalism:

<http://www.skeptdic.com/vitalism.html>

However, today, even most philosophers accept that living organisms (and their organs) are more or less complex arrangements of atoms. So is your brain. Searle's dismissal of the systems argument as not obviously intuitive is simply besides the point. As has been endlessly shown, the natural world (and that includes human nature) is often quite unintuitive. That's why science takes some effort to learn.

I agree it may not be intuitive that the atoms that comprise our brains do not have intrinsic "meaning", "purpose" and "intentionality", but that's the way it seems to be. Now if you want to hold out for non-material properties or "intentionality particles" when there is no evidence for such things, you can certainly do that. But it's not going to explain anything for you.

Finally, your refusal to learn new material that might help you overcome your intuitive beliefs is evidence to me that you are not serious about learning. Rather it would seem that you only want to confirm your heartfelt intuitions. That

is not the path to knowledge and understanding and I would rather not waste my time with someone who refuses to make any effort to challenge themselves. When you are ready to read and discuss the specifics, let me know and ask a question. I'll try to answer it.
By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**
both Searle and I entirely agree with Dennett and yourself, that the human mind is the result of physical processes. What is at question is the nature of those processes. Both sides of the debate also agree that the human mind is a product of evolution.

Maybe. Searle is very reluctant to invoke evolutionary explanations. He commonly refers to the "vulgarity" of Darwinism. But let's continue- I'm thrilled to see some actual discussion.

Quote:

Originally Posted by **Banno**
Both sides agree that any formal language remains a set of symbols until it is provided with an interpretation. This implies that syntax alone is not capable of providing semantics.

I already agreed that many additional conditions are required for language understanding. Evolutionary and cultural context for just two examples. Consider the few cases where a child has raised themselves entirely alone. They have no language.

Quote:

Originally Posted by **Banno**
Both sides agree that the human minds include the capacity to provide a semantic interpretation – to give the system a purpose or intent.

But NOT, a human mind by itself!

Quote:

Originally Posted by **Banno**
Dennett, and I suppose Probeman himself, think that a computer can provide a model of the human mind. Note that such a model does not as yet exist; what they propose is only that it is possible.

Not quite. Models describing mechanistic explanations of various types of mind like behavior and activity do exist, though not in the atomic detail you seem to require. As far as actual artificial human minds are concerned, this is something only possible in principle. Like creating a kidney from atoms, it is possible in principle but will never be demonstrated due to practical considerations. Rather what Dennett and I would say is that semantic interpretation, intentionality, and intelligence are possible on any number of "substrates", both organic and inorganic.

Quote:

Originally Posted by **Banno**
Searle points out a distinct difference between computers and minds – that computers, being algorithmic, prima facie cannot provide an interpretation for their calculations. This is the guts of the Chinese room argument – that the syntactic system of rules does not provide anything – room, inhabitant or total system – with an understanding of Chinese.

"Prima facie"? This is simply the argument from intuition again!

Quote:

Originally Posted by **Banno**

Searle does not say that the brain did not evolve; nor that the mind is not a product of the physical processes in the head; nor that there is some transcendent aspect of the mind. He is saying that the mind cannot be modelled using only algorithms.

Yeah, I ready know that. So?

On Searle's side is the argument from intuition that algorithms can never provide understanding or intentionality. On Dennett's side is 30 years of science that, although is just scratching the surface of consciousness understanding, has already shown that algorithmic processes explain much of our behavior and perception and social interaction.

Quote:

Originally Posted by **Banno**

As for the excursion into evolution in the article Probeman cites, I think it a bit of a furphy. In order make his point, Dennett must maintain that genes have intent. Doing so is not only anthropomorphic and teleological, but begs the question.

I have no idea what a "furphy" is, but genes have basic "intent" in the sense that they only exist to replicate (or replicate only to exist- it amounts to the same thing really). In that sense only, they have a most primitive form of intent. The "intent" that is to replicate. Of course our much more evolved "intentions" are much more complicated, but just like our "intention" to keep parasites outside our bodies, our intentions are evolved from the very distant and basic pseudo-intentions that helped the earliest replicators distinguish themselves (their own boundaries) from the rest of the universe.

You clearly only read the first couple of pages of the article. Try again.

By the way the following two posts might help with the "teleology" problem you're having) (it's from Faustus' and my chapter summaries of Dennett's Consciousness Explained).

By Probeman (John Donovan)

Chapter 7, The Evolution of Consciousness, 1. Inside The Black Box of Consciousness

Taking a new tack, Dennett suggests that we pause in our external (heterophenomenological) scrutiny of the "black box" of consciousness for a moment, and instead consider how consciousness might have arisen evolutionarily. Since human consciousness is obviously a relatively recent phenomenon (evolutionarily speaking), it must have evolved from prior processes that themselves weren't actually conscious. The reason an evolutionary line of thought might be profitable for us, is that it is easier to imagine the behavior of a "device" that one "builds" or synthesizes from the inside out, than it is to try and analyze a "black box" and try to figure out what is going on inside.

Up till now we have been taking the behavior or phenomenology of the brain as a given and wondering what hidden mechanisms inside could explain what we observe. Now let's think about the evolution of brains or nervous systems for doing "this or that" and see if by this we can explain some of the puzzling "behaviors" of our consciousness. Dennett proposes to tell a story, one that is not necessarily complete or scholarly, but in the interests of keeping it short and interesting, more like a hundred word summary of War and Peace. In our particular case- this document is therefore a summary of a summary, so please read Dennett's book to get even "the hundred word summary of War and Peace."

The story of the origins of consciousness will be analogous to other stories from the evolution of biology, for example the origins of sex. Originally all was asexual reproduction and then slowly by some imaginable series of steps, some of these organisms must have evolved into organisms with gender and eventually into us. How, and even more importantly, why did this happen?

The parallels between the evolution of sex and consciousness are intriguing: there is almost nothing “sexy” (for humans at least) about the sex life of flowers, oysters and other simple forms of life, but we recognize in these apparently “joyless routines of reproduction the foundations and principles of our much more exciting world of sex.” In the same way, there is nothing especially “selfy” (as Dennett coins the term) about the primitive precursors of human consciousness, but they lay the foundations for our “particularly human innovations and complications.” Dennett suggests that our conscious minds are the result of three successive evolutionary processes, piled on top of each other, each one successively much more powerful and complicated than its predecessor.
By Probeman (John Donovan)

2. Early Days (Genesis)

“In the beginning, there were no reasons: there were only causes. Nothing had a purpose, nothing had so much as a function; there was no teleology in the world at all.”

This is because there was nothing that actually had “interests.” But after a while there emerged simple replicators. Though they had no inkling of their interests and it would be proper to say they indeed had no interests, we, looking back from our “god-like” perspective can assign them certain interests by defining them an “interest” in self-replication. Of course their replication didn’t really matter to anyone and really made no difference whether they replicated or not (though perhaps we might be grateful they did replicate), but we can say that if these simple replicators are to survive and replicate in the face of increasing entropy (disorder), their immediate environment must be conducive to replication at least some of the time.

To put it anthropomorphically: if these simple replicators want to replicate they should “hope and strive” to avoid “bad” things and seek “good” things. The “good” for such an entity (by our non-teleological definition) is to, however primitively, avoid its dissolution and decomposition. This is the simple replicator’s “point of view” if you will. In this “point of view” there are three kinds of world events: the favorable, the unfavorable and the neutral. Any behavior, even simple chemical causes, of these simple replicators that improves its replication, is a reason or interest in our limited sense, however “unself” recognized that behavior might be to the organism itself.

Now as soon as something is in the business of (preserving) self-replication, boundaries start to become important. Simply because if you are preserving your replicating self, you don’t want to waste your energy on preserving the rest of the universe. So you need to draw a line. The replicator becomes, in a word, “selfish”.

Obviously this primordial “selfishness” does not have most of the variety and breadth of human selfishness, but this “selfishness” is distinctly different from non-life. A piece of granite can in no sense imaginable, be said have an interest in where its boundaries are. Nothing “works” to protect a fracture boundary, no mechanism pushes the boundary back to preserve itself. All things biological have the imperative- “me against the world”. Not just ingestion and excretion, respiration and transpiration but also other processes. Consider the immune system, with its millions of different antibodies arrayed in defense of millions of alien intruders. The fundamental problem that the immune system must deal with is: “recognition” of what are “friendly” forces (belonging to the organism) and what are “unfriendly” forces (those seeking its dissolution). It is worth pointing out that this “army” of the immune system is an army without generals, without headquarters, or even a description of the “enemy.” The antibodies represent their enemies only in the way a million locks represent the keys that open them.

Several other points are worth noting here:

1. Although evolution depends on history, Mother Nature doesn’t “care” how the organism acquires its prowess, just as long as the job gets done. This has important implications for later in the explanation of consciousness.

2. Because evolution is “blind”, as in “undesignated”, there is no way to “foresee” evolutionary or especially phenotypic side-effects. Most genetic evolution is so haphazard from an engineering point of view, that side-effects are plentiful and sometimes these side-effects have serendipitous effects, especially in combination with other

functional systems, that are then effective in dealing with a problem in a new way. As Dennett puts it:

“Multiple functions are not unknown in human engineered artifacts, but they are relatively rare; in nature they are everywhere, as we shall see, one of the reasons theorists have had such a hard time finding plausible designs for consciousness in the brain is that they have tended to think of brain elements as serving just one function each.”

To sum up our primordial facts so far:

1. There are reasons to recognize. (replication itself is a primitive sort of reason)
2. Where there are reasons, there are points of view from which to recognize or evaluate them. (survival of replication, "intentional" or not, implicitly defines a "point of view")
3. Any agent must distinguish “here inside” from “the external world.”
4. All recognition [of boundaries] must ultimately be accomplished by myriad “blind, mechanical” processes.
5. Inside the defended boundary, there need not always be a Higher Executive or General Headquarters.
6. In nature, handsome is as handsome does; origins don't matter (where a mechanistic process is co-opted from, does not matter so long as it provides a survival advantage)
7. In nature, elements often play multiple functions within the economy of a single organism.

Next Dennett looks at the evolutionary survival value for a new "trick". That is “producing” or more simply, predicting or anticipating (no matter how primitively), the future. There are many ways to survive- an organism can armor itself like a tree or a clam and “hope for the best” or it can develop methods for getting out of harms way. If you perform this latter strategy, you are an animal, and the question on your mind is always: Now what do I do?

To do this, you need a nervous system, to control your activities through time and space. For navigating through the sea for a suitable home the sea squirt has a rudimentary nervous system. But once rooted, it eats its brain since it is not needed anymore. Brains are anticipation machines. Even the armored clam cannot always stay closed- it snaps shut as a crude but effective harm-anticipator/avoider.

For more primitive organisms, simple withdrawal and approach responses are tied to bad things (recoil) or good things (engulf). How these two classes of events are discriminated, is of course the job of the perceptual system, which at the lowest level may be a simple chemical reaction. Organisms are “wired” for these responses at the some primitive level and these pre-wired circuits are present in subtle and not so subtle ways even in our advanced human brains.

These early nervous systems depended on avoiding noxious contacts and seeking out nutritious bits (and mating opportunities once sex had appeared, of course), but this could still be improved upon by short range anticipation processes.

The ducking from looming object response is hard wired in us and other animals. It can be observed in newborn infants and is a gift from all our (human and non-human) ancestors that learned to duck “instinctively” and survive to reproduce. We are also hard wired to be sensitive to vertical axes of symmetry. Why? Because this type of visual perceptual pattern usually means another animal is looking at us. Maybe to eat us, so it's better to have an alarm go off that you are being looked at by another animal, as opposed to waiting until you feel its teeth digging into you. Of course, sex adds new dimensions to this type of alarm response too.

An important point Dennett makes, that is crucial for understanding human brain functions, is that there is always a tradeoff between “truth and accuracy” and “speed and economy” in perceptual or anticipatory brain processes. This vertical symmetry detection is one example of an “orienting” response which has been interpreted as a sort of “all hands on deck” alarm for the brain. That is, “we just got an alarm from a crude and specialized harm avoidance circuit, so take us out of auto-pilot and let's do a sensor scan and see what's out there.”

Now these brief episodes of brain process interruption and heightened activity are not themselves episodes of human-style “conscious awareness” (as people redundantly say). But they probably are precursors in the evolution of our human/primate conscious states.

Once this “all hands on deck” alarm circuit existed, it cost little or nothing to keep it turned on all or most of the

time (though I myself wonder if this extra duty cycle time is involved in the increased need for sleep of mammals). So regular vigilance could evolve to regular scanning or exploration, and hence a new strategy evolved: gathering information for information's own sake. But this new brain process of gathering information was itself cobbled out of existing systems and this evolutionary history has left its emotional and affective overtones on our consciousness (as cognitive science has discovered in the mammal brain). That is, the innate links of informing states to withdrawal and engulfment, avoidance and reinforcement were not thrown away, but only attenuated and re-directed.

The dorsal/ventral division of these brain process labors were developed further in the primates into the celebrated right-hemisphere/left hemisphere specializations: the global right brain and the analytic left brain.

If you want to read the rest of this very interesting post it's posted here:

<http://forums.philosophyforums.com/...31&page=1&pp=20>

starting at post #35

It should be pointed out before we get off-track again that we are really talking about heuristic algorithms because it has been useful to see human behavior (and that of a few animals) as self-modifying algorithms due to Baldwinian effects.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

My post #59 answered the system argument – but just to repeat it, there is a distinction to be made between following a set of rules by rote and understanding what they do. Searle's reply appears to me to stand; notwithstanding your counterargument in post #71, which I critiqued in post #76, to which I don't think you replied.

Your counter argument was just a repetition of the argument from intuition, that is "I just don't see how...". But I already know you don't see how- nor do you seem to want to see how.

Tell me, are the cells in your body "just following a set of rules by rote" or do they each actually know and understand the greater purposes to which they labor?

It's no different for the cells in your brain.

Quote:

Originally Posted by **Banno**

Really? I'd like to see a citation for that – not that I doubt you, just that I'd like to know the context.

It's cited in the Dennett article on evolution and intentionality which demonstrates to me that you didn't even read the first few pages. You really don't want to challenge your intuitions do you?

"I'm not sure that this teleology/optimalty story is false, but I do find it thoroughly unsatisfying. . . .I think we can get a theory of error without relying on notions of optimality and teleology. All else being equal, the less such Pop-Darwinian apparatus the better, surely. (ms. p.6)"

Quote:

Originally Posted by **Banno**

I assume you mean "But not *only* human minds". The point would then be moot, since human minds are as yet the only ones we have to examine.

No I don't mean that at all. I meant that human minds require a social and cultural context in order to function the way we observe them to do so today. I also mean that algorithms are indifferent to the substrate they run on. One can do long division with calculators, paper and pencil or a stick in the sand. They will all "work".

Quote:

Originally Posted by **Banno**

I don't follow this explanation at all. Are these "substrates" in some way not algorithmic? How else could they "transcend" the basic problem of interpreting an algorithmic system? So, how is it that a program *provides* an interpretation that is not derivative?

But I'm saying our interpretations are derivative. Ultimately from our genes, but that's a large step to take so that's why I posted the evolution posts from the Consciousness Explained thread. I have a forlorn hope you might actually try and read it- as opposed to skimming it as quickly as you can to avoid comprehension.

Quote:

Originally Posted by **Banno**

Not really. On Searle's side is a few hundred years of formal logic, which indicates that formal systems need to be *provided* with an interpretation. As I said before, saying that this is an intuition is just name-calling. Do you have an argument that shows how a formal system can provide its own interpretation?

Actually Searle's intuitions go all the way back to Aristotle and his "essences" (and as "folk" psychology and religion probably further than that). No, I have something better- an explanation that is consistent with all the scientific evidence and evolution by natural selection. If you don't like me pointing out your argument by intuition (Searle's "unsatisfied" intuitions included), make an argument based on the data as opposed to "intrinsic" essences.

And besides, who said anything about "formal" systems- you just snuck that in. Heuristic algorithmic systems are not formal systems. They are not bound by a fixed set of rules. There is nothing in human or animal behavior that requires formal system proof. In fact that would be very "unevolutionary" since organisms acting in real-time must sacrifice accuracy for the ability to react usefully and in time. Philosophical intuitions just aren't going to "cut the mustard."

Quote:

Originally Posted by **Banno**

Genes simply do not have any purpose beyond what *we* give them. Genes just are, just like rocks. Describing them as having intent is giving an interpretation to a purely syntactic system – the point at issue. Since it is Dennett who provides this interpretation, the intent is, as he says, derivative.

Genes are just like rocks? I think you ought to reconsider that statement. Seriously.

Dennett is only saying that regardless of whether they actually know their purposes, they have purposes. Genes replicate and your heart cells pump your blood. Do the genes know their purposes and do the heart cells know their purposes? No. But you might want to thank them for their efforts anyway.

Can you describe any part of your body without providing an intentional or purposeful "stance". Not really. That is Dennett's whole point of the "intentional stance". Our much more evolved and self-aware, language described "intentions" and "purposes" are a far cry from the blind and unaware "purposes" of genes and heart cells, but we ultimately derive our intentions, like the parasite avoidance example I gave, from this primitive source. This is the standard model accepted by science since Darwin's time. It may not "feel right" or "satisfy" our emotional intuitions, but it does explain what has stumped philosophers since Aristotle.

Quote:

Originally Posted by **Banno**

You will no doubt think I am just being obstructive, but this post supposes that sex and consciousness both evolved – not something with which I would take issues – and that in some way "self" is analogous to sexual pleasure. A bit weird, but not really very helpful as far as I can see.

As you said- they both evolved. The sex of oysters isn't very sexy is it? Yet it is the ultimate origin of our mammalian caring, loving, infatuated, head over heels, nuturing and self-sacrificing emotions of love.

Just as the amoeba has no self-aware sense of self. It's not very "selfy" is it? But it does indeed "know" where it ends and the rest of the universe begins. As I paraphrased Dennett:

"Now as soon as something is in the business of (preserving) self-replication, boundaries start to become important. Simply because if you are preserving your replicating self, you don't want to waste your energy on preserving the rest of the universe. So you need to draw a line. The replicator becomes, in a word, "selfish"."

Weird is good- that means you are starting to think a little. Keep re-reading it. It will start to make sense if you try.
By Probeman (John Donovan)

Quote: Originally Posted by probeman

I don't know. But if we can convince people that we should judge all such entities (including ourselves) by both their actual capacities and their demonstrated behaviors, as opposed to their external appearances and/or imagined metaphysically "intrinsic" or "essential" properties, we will have taken a large step in right direction.

Quote:

Originally Posted by **Nonblack Raven**

I may be misreading you (the term "actual capacities" is vague), but it does seem to me that there is something between demonstrated behaviors and metaphysical properties. That something is the process that produced the behavior.

Oh, I agree- it does SEEM that way. Look at it this way, metaphysical properties are simply metaphors for trying to describe the universe based on our intuitions. From that perspective they serve minimal purpose in creating scientific theories, except as a possible side interest, much as scientific explanations for optical illusions may shed some light on how our visual perceptions operate.

By capacities I mean actual functional capacity. For example, an infant for example doesn't have the capacity for moral behavior (yet).

Quote:

Originally Posted by **Nonblack Raven**

The process that produces behavior is under-determined by the behavior. I think this is true in two senses...

I agree and there are many examples from evolution such as the bird wings and insect wings.

Quote:

Originally Posted by **Nonblack Raven**

First, by assumption the Chinese Room produces verbal behavior sufficiently similar to that of a human who understands Chinese that human beings could not tell that the computer does not understand Chinese. However, this does not prove that the Chinese Room understands Chinese. The process that the computer uses to produce the verbal behavior could be quite different than processes that understands Chinese. The similarity of the verbal behavior proves nothing in this respect. This is part of Searle's point.

I agree that the Chinese room is an overly simple metaphor for a system with "understanding". But that is why it doesn't prove anything. It certainly is not a scientific model of how the brain works. That task would require much more detail, some of which isn't completely filled in yet.

Quote:

Originally Posted by **Nonblack Raven**

Second, Searle's Chinese Room does not produce verbal behavior by some unknown process, but through a quite specific, carefully described process. The Chinese Room is not just a black box that produces verbal behavior, but a process that produces verbal behavior in a very particular way. By reading this process, I conclude, in spite of similarities in behavior, the Chinese Room does not understand Chinese. For example, if I understand a language, I do not just produce verbal behavior, I am capable of picturing, hearing, feeling some of the kinds of things the symbols stand for. From the description of the process provided, the Chinese Room simply cannot do these kinds of things. This is not metaphysics, not "folk psychology", it is a capacity I have and associate with understanding that the Chinese Room, as described, simply does not have.

It is both "folk psychology and metaphysics based on intuitions. The argument that "I just can't see how..." merely tells me that you haven't really tried to understand how algorithmic processes can produce incredibly complex behavior.

Consider IBM's Deep Blue- a hundred years ago it would have been unthinkable that a machine could beat a human chess player. Yet today many computers can beat all grandmasters and Deep Blue can regularly beat Kasparov. Now I think a case can be made that Deep Blue, using algorithmic processes, actually "understands" how to play chess. You may disagree, but only based on metaphysical intuitions. In any case, the point is that Deep Blue is hundreds of orders of magnitude less complex than the human brain.

Quote:

Originally Posted by **Nonblack Raven**

The Chinese Room is not a refutation of physicalism. There may be other kinds of devices and other kinds of sets of programs that do understand Chinese. The Chinese Room only offers a problem for first generation AI--the kind of algorithms that produced expert systems, chess playing programs, etc-- it shows that this kind of program, though capable of producing very complex mind-like behavior, is not actually performing the way human minds do.

Oh, I agree completely that the Chinese Room is not actually performing the way human minds do. But both are algorithmic.

Quote:

Originally Posted by **Nonblack Raven**

Let me add, to take a point of Dennett's, the Chinese Room does not show that no possible human-designed device cannot be similar to human mind (I think Searle and I agree on this point), nor does it show that some vastly more complex program, perhaps incorporating neural networks, evolutionary programming, and sensors that show the device what is happening in the world, could not understand Chinese.

Conclusions: AI researchers provide a useful service in finding programs that can produce mind-like behavior. However, the production of mind-like behavior does not mean the programs are human minds or should be treated as such. To determine this would require a demonstration that the program and the human mind produce the behavior in the same kind of way. This is not a matter of metaphysics, it is a matter of understanding the processes used in programs and in the human mind.

Hmmm... so you are saying that if an alien showed up tomorrow that appeared as intelligent and moral as humans, we shouldn't treat it as if it were as intelligent and moral as humans merely because it's "insides" might be operating in a similar way but not exactly the same as us humans?

Again, I think we should judge entities by their capacities and demonstrated behaviors not by some imagined metaphysical essences, appearances or even different "substrates". You do see where your position is heading?
By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

Anyway, you surely jest in thinking it is conscious! LOL Machines are just machines, they do not have souls

And based on all the available scientific evidence, neither do you.

By Probeman (John Donovan)

Quote:

Originally Posted by **Nonblack Raven**

An example to consider. I and any other programmer could write in five minutes a program that, given certain inputs, would yield a message that said "That Hurts." Do you think the program feels pain? Why or why not?

This is most disingenuous. I can "write in five minutes a program that, given certain inputs, would yield a message" that said "I understand how to play chess." Would that program know how to play chess?

By Probeman (John Donovan)

Originally Posted by **Minty**

Suppose the soul is non-physical. What scientific evidence is there for its non-existence?

There is no scientific evidence for the non-existence of the soul. There is also no scientific evidence for the non-existence of Fairies and Santa Claus, nor evidence for the non-existence of anything else. This is Philosophy 101.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Ah, Probeman. In post #95 you quoted *Searle* as talking about "vulgar" Darwinism. I asked you for a citation.

You refer to the article you cited earlier, in which the only such mention is *Fodor*, not *Searle*. The mention is not referenced.

Yes, you are correct. I got confused because in the next paragraph Dennett lumps the two together:

"I appreciate the candor with which Fodor expresses his discomfort with appeals to evolutionary hypotheses. (Elsewhere he finds he must help himself to a bit of "vulgar Darwinism" to buttress an account he needs of the functions of transducers.) Why, though, should he be so unwilling to follow down the path? Because he sees (I gather) that the most one can ever get from any such story, however well buttressed by scrupulously gathered facts from the fossil record, etc., is a story with all the potential for indeterminacy that we found in the tale of the transported two-bitser. And Fodor wants

real, original, intrinsic meaning--not for the states of artifacts, heaven knows, for Searle is right about them!--but about our own mental representations. "

The point is that both Fodor and Searle want to appeal to an "intrinsic" intentionality that transcends mere mechanism. A sort of metaphysical property. The possibility that human intentionality, like sex and other more prosaic behaviors could have their origins in our evolutionary history is intuitively repugnant to them.

In any case, you didn't respond to my discussion at all. Do you find it odd that your heart cells aren't self-aware of their purpose? As I stated above:

Dennett is only saying that regardless of whether they actually know their purposes, they have purposes. Genes replicate and your heart cells pump your blood. Do the genes know their purposes and do the heart cells know their purposes? No. But you might want to thank them for their efforts anyway.

Can you describe any part of your body without providing an intentional or purposeful "stance". Not really. That is Dennett's whole point of the "intentional stance". Our much more evolved and self-aware, language described "intentions" and "purposes" are a far cry from the blind and unaware "purposes" of genes and heart cells, but we ultimately derive our intentions, like the parasite avoidance example I gave, from this primitive source. This is the standard model accepted by science since Darwin's time. It may not "feel right" or "satisfy" our emotional intuitions, but it does explain what has stumped philosophers since Aristotle.

By Probeman (John Donovan)

Quote:

Originally Posted by **Nonblack Raven**

Your response is also most "disingenuous". I cannot help but note that you fail to answer the question, "Would this program feel pain?" I think it would not. What do you think?

Not in the way that we feel pain clearly. You seem to want to give our phenomenal perceptions a special status. It appears that you are headed for the zombie argument which is another metaphysical dead end and I will refer you to Death Monkey's thread with Monroe for that.

Quote:

Originally Posted by **Nonblack Raven**

However, a program that wrote "I understand how to play chess" and could not play chess would not know how to play chess.

Obviously. But what if Deep Blue said "I understand how to play chess". Would you disagree with it? In what sense is it incorrect? In what sense does Deep Blue not understand how to play chess?

Now, you might answer: "Well it's not self-aware that it understands how to play chess". And I would agree. Not like we are self aware that we understand how to play chess. But the issue of self-awareness is another discussion which Dennett also responds to. But if you just want to know if a mere mechanism can "understand", then I suspect that Deep Blue does "understand" how to play chess- better than any human in fact.

By Probeman (John Donovan)

Quote:

Originally Posted by **Nonblack Raven**

I am sorry, but I do not see where you have actually addressed the question of whether we can talk about

processes without being metaphysical. My contention is that we can.

If using the "intentional stance" is metaphysical then every biologist is guilty:

". . . A much more demanding *task* for these enzymes is to *discriminate* between similar amino acids. . . .However, the observed *error* frequency in vivo is only 1 in 3000, indicating that there must be subsequent *editing* steps to enhance fidelity. In fact the synthetase *corrects* its own *errors*. . . . How does the synthetase *avoid* hydrolyzing isoleucine-AMP, the *desired* intermediate? (pp.664-5; Rosenberg's emphases) "

Quote:

Originally Posted by **Nonblack Raven**

I agree, the details are not yet filled in. As I said, I do not view the Chinese Room as a proof that brains/minds cannot use understandable programs, only that they do not *seem* to use this kind of program.

I already agreed with that statement. Actually I agree with a stronger statement. Brains do not use any self-aware, language understanding homoculi (in "rooms") that have levels of self-awareness and language understanding similar to the whole brain.

Quote:

Originally Posted by **Nonblack Raven**

I can agree that Deep Blue can play chess, and that this a major achievement for AI. Let me note in passing that I have written a number of expert systems that my clients were very happy with. I am delighted to have access to such programs. But I do not think that such programs have a clue as to what they are talking about.

Because they weren't "designed" for that ability. But our self-narrative sense of awareness (for that is all it appears to be) WAS evolved for as a useful selective advantage in social interaction.

Quote:

Originally Posted by **Nonblack Raven**

I am uncertain that Deep Blue "understands how to play chess". I suspect not, because Deep Blue does not know it is playing chess. To know this would seem to me to require kinds of programs that Deep Blue does not possess.

Granted. I responded to this in a previous post that self-awareness is a separate capability. Ask yourself this, does a bird "understand" how to build a nest? Does an infant "understand" how to learn language? Just because something cannot discuss their understanding does not mean it doesn't have understanding. As you said with your five minute program example, visa versa, just because something can discuss understanding doesn't mean it does understand. I agree that understanding and self-awareness of understanding are two separate issues.

Quote:

Originally Posted by **Nonblack Raven**

However, I think that there is more than adequate evidence to suggest that Deep Blue does not play chess in the way a human being plays chess.

Oh, I agree with that. But like Dennett, I'm substrate neutral when it comes to the ability for understanding, and self-awareness for that matter.

Quote:

Originally Posted by **Nonblack Raven**

This is where I think we part ways. Depending on the meaning of the term "algorithmic" I think you have done nothing to prove that human minds are "algorithmic". You may be able to prove this in the future, but have not yet done so. By all means continue with your promising research program, but do not imagine that it is now the only possible research program.

Oh, we definitely part ways here. I'm only a chemist but I've read enough of the literature in cognitive science to

know that every behavior that has been examined in detail, from the amoeba on up is algorithmic. No magic essences or vital properties so far as can be determined.

Quote:

Originally Posted by **Nonblack Raven**

Actually, I would have no idea if some imagined alien that appeared as intelligent and moral as human beings is as intelligent and moral as human beings. For moral reasons, I think we should treat it as such until proven otherwise.

You mean for pragmatic reasons I think. Do unto others...

Quote:

Originally Posted by **Nonblack Raven**

However, the question of whether the creature is intelligent and moral, and should be treated as a moral entity ultimately depends on questions concerning the process that produce such behavior, and what constitute sufficiently similar processes that I do not think we yet have answers to.

Ooops. You just stepped in it!

So if, hypothetically, another human race were shown to have non-identical processes that produce behavior identical to say, the white race, you would feel that process difference should be included as a criteria for considering whether they should be treated morally?

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

I also am convinced, by Searle's argument, that they do not arise from syntactic processes *alone*, since syntax does not provide semantics. The question of how intentional states arise will not be answered by adopting Probeman's strategy of calling them names.

I also agree intention does not arise from syntax alone for the simple reason that intentionality seems to be a behavior exhibited by all organisms and many devices, most of which do not have language abilities.

The question of how intentionality did evolve from primitive intentions into the intentions that we can talk about, has been explained by vulgar, bottom-up evolutionary considerations.

I will add regarding "name calling" that if my pointing out someone's appeals to intuitions is "name calling", then all of science is "name calling". Because it has been endlessly shown by science that our heartfelt intuitions are simply unreliable sources of knowledge.

It is intuition that will never explain these questions as Aristotle would tell you, were he still alive today.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

As for you argument, I apologise for not responding – I had not recognised it for what it was. Are you now claiming that it is impossible to describe the function of the various cells in a body without reference to intentional states? That would be a curious turnaround.

No it is not impossible, just impractical. Just like trying to describe wetness with QM. Intentionality is a useful concept that merely and usefully describes behavioral outcomes of organisms and devices. That is Dennett's "intentional stance". There are no intentional atoms or intentional particles, just the behavior of mechanistic algorithms that are well described by that "stance".

Quote:

Originally Posted by **Banno**

I accept that cells can be *described* as having a purpose. I have put the point – several times – that such descriptions are inherently anthropomorphic. The cells of course do not perceive themselves as having a purpose, nor do they act *because of* that purpose. This is not true of the human mind.

The cells in your heart do indeed act "because of the purpose" of pumping blood. Unguided, blind, evolution "designed" the cells in your heart for that "purpose". The cells in your brain are designed the same way, individually they are unaware, not understanding. But their purpose as a whole still exists as observable intentionality.

Even if we couldn't TALK about it.
By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

You said there's loads of scientific evidence that the soul doesn't exist, didn't you?

Nope. Just "loads" of scientific evidence for an evolutionary (algorithmic) explanation of organisms and their behaviors. Science has nothing to say about non-material or supernatural claims except that it can offer a better and more reasonable explanation.

Minty, I don't want to blow you off, but if you are truly interested in these questions you should read a little of the background material first. Once you do that I'd be happy to try and answer any questions you have. I suggest that you start with Darwin's Dangerous Idea by Dennett.
By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Well, I suppose I'll have to bight the bullet and address the article more directly. I had thought the best critique was Probeman's own, in post #87: "*But genes do indeed have a non-teleological "purpose": replication.*" This oxymoronic phrase sums up the problem admirably: Dennett wants to have his cake and eat it, too, by constructing purposes that are without purpose.

Yes, I'd say you are close.

Consider the concept of "design". Before Darwin the word design was two things, one, by God as in creation, and two by humans as in artifacts. Darwin gave a new and extremely useful meaning to the word "design", as in evolutionary design. Now this new concept of

design had nothing to do with intelligence or purpose as it was traditionally conceived. But today, we can understand how intricate and complex "design" is produced by blind, purposeless algorithms. Life is actually engineering by natural selection.

Dennett is doing the same thing to the word intention. He is excising the metaphysical baggage long associated with the word to show that simple (and complex) mechanisms have intentionality by their observable behavior. There's nothing "intrinsic" about this intention, just like there's nothing "intrinsic" about "design" in evolution. It's simply an observable process due to the outcome of certain algorithms.

Quote:

Originally Posted by **Banno**

That genes have reproduction as their *purpose* is a particular way of describing them, from the point of view of a creature that is capable of setting out and acting on its own intentions. Genes are not capable of setting out and acting on their own intentions. Although it is fun to talk of them as if they did, they do not *desire, wish, long, crave, yearn or need* to reproduce.

Of course genes have reproduction as their "purpose". What else are genes for? What else do they do? And they are perfectly capable of setting out and acting on their "own" (blind, unintelligent, unaware) "intentions." How else do you suppose us animals got here?

I realize that Dennett is expanding the meaning of the word "intention" here, just like Darwin expanded the meaning of the word "design." It's "design" without "Design", just like it's "intention" without "Intention."

Quote:

Originally Posted by **Banno**

Probeman, the genes do not really *want* to reproduce; they just want you to think that, while they make other plans...

You have a grain of truth in your little joke here. Indeed, our gene's purposes are no longer perfectly co-incident with our much more evolved self aware purposes. This is primarily due to Baldwinian selection (learning to learn better). Although if one observes a singles bar scene, genes in action is just about all one will see. On the other hand, we can decide to use contraception, so we can thwart our gene's purposes to some (often limited extent).

Here's a quote from Dennett that might help: <http://ase.tufts.edu/cogstud/papers/evoltele.htm>

"The apparent differences between adaptationist theorizing in biology and intentionalist theorizing in psychology are due, in my view, to the huge differences in time scale, and--more evident in the discussions of both Ringen and Bennett--a downplaying of the importance of the implications of the ubiquitous idealizing assumptions in both enterprises. When we grasp the nettle and confront the ineliminable "practical difficulties" (Ringen) that beset the evolutionary theorist intent on distinguishing actual cases of selection for, and the parallel practical inability of the intentionalist psychologist to cash out the idealizing assumptions that permit talk (in Bennett's example) about a "class of environments . . . unified with help from the concept of food-getting" we see that both enterprises continue to avail themselves--quite appropriately and defensibly--of what Quine called the "dramatic idiom"; the sense-making interpretation-talk of the intentional stance. I claim that since there is just one sort of explanation going on in both quarters, the choice Ringen offers me must be rejected: teleology is neither as illusory as his neo-Darwinians claim, nor as real and irreducible as his Aristotelian Bennett claims. "

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

Read [this](#).

The article seems to argue against Chalmers and for Searle and Nagel which I find amusing because all three are mysterians positing non mechanistic/algorithmic "forces" to explain what evolutionary explanations already provide.

Hasker is also an "intelligent design" proponent, which is a 200 year old religious belief pretending to be science and I have no interest in religious explanations. From Paley's eye, or Behe's blood clotting, the claim that some biological features couldn't possibly be derived from evolutionary mechanisms is simply the argument from ignorance. The argument that whatever we haven't yet scientifically explained is somehow explained by God is not only not science, it's also a science "stopper". After all, why try to explain anything if the explanation is God? In any case, all examples of "irreducible complexity" from Paley's eye to Behe's blood clotting and rotating flagellum have been explained in evolutionary detail since they were proposed.

But you should have noticed that I have been posting many warnings about "greedy" reductionism if you've been reading the thread carefully. The fact that all biological design and intentionality is ultimately derivable from mechanistic algorithms does not mean that we should always resort to lowest level explanations (e.g., QM or chemistry) when a higher level explanation (adaptive design or algorithmic intentionality) will suffice.

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

So what's wrong with his arguments?

What's wrong with his argument is that it's metaphysically (supernaturally) based and therefore not only untestable but also not useful (except possibly for trying to convince one's self that God loves us).

If one's goal is to try and demonstrate the existence, love or purposes of gods or fairies, there are many intuitively appealing metaphysical routes one can take. I have no interest in metaphysics. My interest is epistemology, especially scientific epistemology, since that appears to be the best (only?) method humans have for gaining reliable knowledge about the natural world (and that includes human nature so far as we can tell).

By Probeman (John Donovan)

In a nut shell

This following quote from Dennett pretty much sums it up:

"How can natural selection do this without intelligence? It does not consciously seek out these rationales, but when it stumbles on them, the brute requirements of replication ensure that it "recognizes" their value. The illusion of intelligence is created because of our limited perspective on the process; evolution may well have tried all the "stupid moves" in addition to the "smart moves", but the stupid moves, being failures,

disappeared from view. **All we see is the unbroken string of triumphs.** When we set ourselves the task of explaining why those were the triumphs, we uncover the reasons for things--the reasons already "acknowledged" by the relative success of organisms endowed with those things.

The original reasons, and the original responses that "tracked" them, were not ours, or our mammalian ancestors', but Nature's. Nature appreciated these reasons without representing them. And the design process itself is the source of our own intentionality. We, the reason-representers, the self-representers, are a late and specialized product. What this representation of our reasons gives us is foresight: the real-time anticipatory power that Mother Nature wholly lacks. As a late and specialized product, a triumph of Mother Nature's high tech, our intentionality is highly derived, and in just the same way that the intentionality of our robots (and even our books and maps) is derived. A shopping list in the head has no more intrinsic intentionality than a shopping list on a piece of paper. What the items on the list mean (if anything) is fixed by the role they play in the larger scheme of purposes. We may call our own intentionality real, but we must recognize that it is derived from the intentionality of natural selection, which is just as real--but just less easily discerned because of the vast difference in time scale and size. "

I suspect that 150 years from now, the word "intentionality" will have the same implication of being an observable outcome of algorithmic process, just as the word "design" has today. Of course there will always be those who for psychological reasons cannot give up the comfort of "intelligent design" or "intrinsic intentionality."
By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Firstly, Probeman makes the claim that a new meaning has been attached to *design* since the advent of *the Origin of Species*. Now this is curious. And also rather difficult. (I will ignore the strange ontological power apparently attributed to algorithms. Of course Probeman does not really want to claim that it is the algorithm that drives evolution).

Banno, I'm afraid once again your intuitions mislead you. Let's examine your strawmen characterizations of my points and see they lead us.

First of all I argue against any ontological powers or properties. Evolution is indeed simply algorithmic processes. If you want to say evolution is "driven" by algorithms that's fine too, although one could also say that evolution is driven by genes, or bio-chemistry. It just depends on what level of reductionism is appropriate for optimal explanatory power. What do you think "drives" evolution? God? Aliens?

Quote:

Originally Posted by **Banno**

Probeman appears to be claiming that we can, since Darwin, use the word *design* when talking about things that are, quite specifically, not designed, but the product of evolution. Now, I could not find a reference to such in any dictionary I checked; and, although I am an avid reader of science journals and books, I cannot recall seeing such a usage before.

You miss my point entirely. What Dennett, Gould, Dawkins and other scientists have pointed out is that it only makes sense to understand evolutionary mechanisms and adaptation as the "appearance of design". As I took great

pains to point out and you apparently took equal pains to ignore, Nature's natural selective processes can produce "design" **without** a designer. Engineering **without** an engineer. In fact using computer models there is even a enormous field with applications from aeronautics, to genetics to yacht racing that use evolutionary engineering techniques to let computers "discover" improved designs by evolutionary selective algorithms.

Are these computer program "aware" that they are "designing"? Of course not. Yet, they mindlessly do actual "design" work, just like "Mother Nature."

For someone who claims to be well read in science I find it fascinating that you have missed this entire field of scientific applications. But that is once again merely your old argument from ignorance.

Here are a few evolutionary design links for you to browse:

<http://www.soton.ac.uk/~ajk/opt/welcome.html>

<http://portal.acm.org/citation.cfm?id=1016037>

<http://www.cs.ucl.ac.uk/staff/P.Bentley/evdes.html>

<http://www.msci.memphis.edu/~dasgup...y/wc3paper.html>

Quote:

Originally Posted by **Banno**

Next, Probeman claims that Dennett is doing to *intention* what it was that Darwin did (or, perhaps, didn't do) to *design*. Now I have no qualms about re-defining a term, or adding a new meaning, if it seems worthwhile. But it is worthy of note that this process does not remove the old meaning. Probeman cannot re-define intention so as to extinguish its meaning as, to quote from the *Oxford*, "The action of straining or directing the mind or attention to something". Saying that you did something *intentionally* is just too useful for us to do without.

Your strawman is that I'm "extinguishing" the meaning of the word intention, when in fact I took great pains to say that Dennett is **expanding** the word intention. In exactly the same way that the word "design" has been expanded to include "mindless", algorithmic design.

It is Searle's usage of "Intention" (with a capital "I" for intrinsic) that is entirely metaphysical and without any scientific basis. There are no intentional properties of atoms, there are no intentionality particles. Intention is simply a useful and efficient way of describing behavior by organisms and some devices. That's not to say that we can't still say "I intend to have a cup of tea". We can also say "I designed this clock". But the origins of the "design" of our bodies and the "intentions" of our brains are rooted in algorithmic processes that are themselves mindless and unaware of design or intention.

I tired to explain this by analogy to heart cells or brain cells which perform useful and apparently purposeful behaviors without themselves understanding those purposes. Where do you think this "intrinsic intention" comes from? You want to elevate "Intention" to some magical metaphysical status, just like the creationists want to elevate "Design" to a supernatural status.

Quote:

Originally Posted by **Banno**

Probeman, if Darwin and Gould talk of the Panda's thumb as having been *designed* (and I am not convinced that they ever actually do), they are speaking *metaphorically*. Gould makes it abundantly clear that he does not think that evolution has a purpose. I agree with him. If Gould or Darwin speak of one body part being a better *design* than another, one does not conclude that they think there is a designer. Similarly, one ought not conclude that they think that nature has a plan or purpose in mind for that body part. And in exactly the same way, one ought not infer that they think that nature has an *intent* for that body part.

I agree evolution has no "purpose" or "intent". The point that Darwin and Dennett have made is that evolution has "mindlessly" (without a designer), produced hearts that are for the (apparent) purpose of pumping blood, just like evolution has produced organisms with (apparent) intentional behavior to seek good things and avoid bad things.

The only thing that "guided" these processes was natural selection and in the case of human behavior some

additional Baldwinian processes. But they are still all Darwinian algorithmic processes.

Quote:

Originally Posted by **Banno**

Probeman might well have us live in a world in which genes intend to reproduce, rocks intend to weather, water plans to run down hill, and balls fall because of their innate love of the earth. But this is poetry, not science.

Well since you claimed before that there is no difference between genes and rocks I can see why you might be so confused. However, genes are algorithmic and rocks are not.

Quote:

Originally Posted by **Banno**

But if physics is more than a set of algorithms, why on earth would one suppose that the mind is nothing but a set of algorithms? Yet that appears to be what Probeman would have us believe.

Wrong again. Of course the mind is much more than a set of algorithms. Just as the physics books contain examples, so does the human mind. We call these examples "memories." But these processes are still mechanistic and algorithmic at root.

This misunderstanding of yours reminds me of your earlier claim that syntax alone can't produce "meaning." I agree—"meaning" also requires context and interaction and correlation with an environment. Your attempts (and Searle's Chinese Room attempts) to create black and white metaphysical distinctions will not provide you with "meaning" either.

You're not, by any chance, one of those praying for QM indeterminism to "save the mind"?

By Probeman (John Donovan)

Banno,

I don't mean to shoot any more water under your wheel, but a friend of mine saw your previous post and had these comments (some of which I have already alluded to in the previous post):

1. It's pretty standard practice in biology for scientists to wonder what a specific behavior or organ is FOR. This implies what function it has. What the heck else could it mean? Dawkins makes this explicit in a recent article on Beliefnet (reprinted from Free Inquiry, I think). He discusses the behavior in some birds called "anting," where a bird sits atop a swarming nest of ants and lets them run rampant through its feathers. What is this for?
2. Function implies design. Does the heart function—was it designed—to pump blood, or is this just an amazing coincidence? Do wings have the function of flight, or are they decorative racks for feathers (actually, even that implies a function). Does the fusiform face area of the brain function to help primates recognize faces, or does it just "light up" when faces are seen completely by accident? When a parent bird begins to fake a broken wing and leads a predator away—behavior which isn't taught and must be programmed by the genes—what is the function? A happy coincidence again?
3. "Design" does not always imply a conscious designer, or in other words that the goals and plans of the design were ever represented to anyone (hence Dennett's "free floating rationales"). Before natural selection, there would have been no reason to imagine design with no consciousness. Darwin shows us how we can. That's what evolution is.
4. Artificial life: Examples from this area are legion—or at least they were in the 1990's, when I did a lot of reading on the subject. In countless experiments, artificial life organisms (often no more than a few lines of code) have, WITHOUT HAVING BEEN INSTRUCTED TO BY THEIR PROGRAMMERS, evolved parasitic behaviors, then specific adaptations designed to thwart the specific parasitic strategies of the new parasites, and so on—a typical

arms race as found in nature. It's quite clear in these cases what the newly evolved lines of code were "for"—that is, clear they had a function designed by virtually natural selection.

5. Genetic algorithms. Sometimes a conscious designer isn't as smart as an unconscious process that produces design—like natural selection. With this thought in mind, genetic algorithms (another form of artificial life) have been grown, and have reproduced, and have produced designs for various bits of equipment. In these cases, programmers create a fitness landscape consisting of sought for capabilities and restraints that the final widget is supposed to have. This landscape functions as the environment of selection pressures, just as any a real animal may find itself in. The end results have been brilliantly efficient widgets of various sorts with features unlikely to have been thought of by even the most clever of engineers.

6. Given all the above, once you forbid yourself from using the terminology of design simply for ideological reasons, you rob yourself of the ability to recognize and quickly understand vast patterns in nature. And you are also forced into the "it's all a convenient coincidence" refrain when confronted with obvious examples of function and design.

By Faustus (Brain Peterson)

Quote:

Originally Posted by **Banno**

In trying to make sense of your account it seems necessary to make the obvious point that algorithms do not have causal power. Algorithms are just algorithms. So it appears incoherent to say that "evolution is indeed simply algorithmic processes". Evolution is a description of the way in which living things change over time. This description includes some processes that might be described algorithmically, but it is silly to say that evolution is nothing but a series of algorithms.

Causal power? I'm not discussing Aristotle's "causes". I'm presenting a few aspects of the scientific explanation of evolution and the mind. Scientists don't care about "ultimate" causes- that's metaphysical speculation. What scientists care about is description and explanation.

At a certain level of description evolution is algorithms. At another level of description it is genes. At another level of description it is bio-chemical reactions. At another level of description it is atomic interactions, at another level, etc, etc.

All these levels of description are useful depending on one's explanatory goals and the level of detail required. This is what scientists do everyday in their research. I should know.

Quote:

Originally Posted by **Banno**

But if you do indeed think that evolution is just a series of algorithms, it might explain why you find it difficult to see why the mind is not also just a series of algorithms. It comes back to that pivotal issue, of what links the algorithm to the real world. Algorithms by themselves explain very little. It is not until they, like all formal languages, are interpreted that they quite literally achieve significance.

What links algorithms to the real world is sensory perception. This is well understood by cognitive science. Of course algorithms "by themselves" don't explain everything. But at a certain level of explanation they explain exactly what is needed without calling on "souls" or other magical non-material forces. And this has nothing to do with "formal languages" by the way.

Quote:

Originally Posted by **Banno**

I am quite happy to agree that a bird's wing could have a design (sb.) – that is, as it were, that one might lay out a drawing of the wing, setting out the various parts. But you appear to take this to imply that it was designed (v.) to

a purpose. This leads to a bit of gymnastics as for instance when you say “function implies design (v)”, and then feel obliged to say that “design (sb)” does not imply a designer, presumably to avoid being trapped into creationism; or again when you say: “I agree evolution has no "purpose" or "intent". The point... is that evolution has "mindlessly" (without a designer), produced hearts that are for the (apparent) purpose of pumping blood”. I can only assume that this is a rhetorical device on your part. You wish again to have your cake and to eat it. Either evolution has a purpose, in which case one can attribute intent to it, or it does not, in which case you may not coherently attribute intent to it unless you are doing so metaphorically.

Rhetorical? Metaphorical? One could argue that science is rhetorical and metaphorical I guess. Scientific theories and models could be described as metaphors if you like. But the point is that they are useful and predictive metaphors. They actually work and provide useful results- unlike metaphysical metaphors. So I think we should distinguish between the two types of metaphor.

The point I'm trying to make is twofold. One, we all use the words "design" and "intention" in scientifically "metaphorical" ways to describe organisms that weren't designed intelligently. Dennett suggests, that just as the appearance of "Design" in nature is an illusion, we can also see that the appearance of "Intention" in nature is also an illusion. There are mechanisms that are usefully described as "designed", just as there are behavioral outcomes that are usefully described as "intentional". But no magical force is directing the design, just as there is no magical force directing the intention. Just algorithms interacting with the environment.

Two, levels of explanation are what science is all about. If it helps scientists to describe the "function" or "purpose" of the heart to understand what it does and how it works, that is a useful "metaphor." Biologists still understand that the heart was designed by mindless algorithmic processes. If you disagree with that, then it's you that is calling on Gods or Aliens to the rescue.

Dennett suggests that it is helpful to understand the mind by showing that the mind can be understood by using levels of explanation that are increasingly less mindlike, less self-aware and less intentional. This is what science in general has been doing very successfully for the last several hundred years and it's no different in cognitive science. Searle is a dead end because he holds out for metaphysical properties that are simply unnecessary for science. His "intrinsic intentionality" will go away without providing any benefit just like the "vital essences" of the 19th century did.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

I was very happy to see you agree with both myself and Searle when you said (in #144) :Perhaps you are beginning to see the point. Intentionality comes about as a result of the causal relationship between the environment and a mind. Ergo, the mind is more than an algorithm.

Right. The mind is a set of algorithms interacting with an environment (and itself). No metaphysically intrinsic properties or essences are required. I've always agreed with that point (that was in fact one of my original objections to Searle's Chinese Room- no environmental context is provided).

And that is why Searle is wrong that the Chinese room demonstrates what he claims. Because heuristic algorithms (on many different substrates) interacting with the environment can produce behavior that is **well described** as having understanding, meaning and intentionality. There are no understanding “particles”, meaning “rays” or intentionality “essences”, and none are required.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Not quite. The apparatus also has to be able to act on the world, as well as to sense it. Brains do this very well.

Look, are you asking for a list of everything the brain requires to function? It requires a blood supply too! The point is that algorithmic mechanisms and brain processes of various sorts explain quite well what we observe behaviorally. But you want more than that- you want magic.

Quote:

Originally Posted by **Banno**

(my bolding) **Do you man that they are badly designed?**

Yes, I mean exactly that. For example, you have a blind spot right near the center of your field of vision because the retina is penetrated by the optic nerve there. Does that sound intelligently designed? But, the point again is, given enough iterations, "blind" natural selection can accomplish some impressive engineering.

Quote:

Originally Posted by **Banno**

The trouble is that I have direct knowledge of my intentions; if I stand up in order to get my toast, then I know that it was my intention to get the toast. So I conclude that any argument that claims that my intentions are somehow an illusion must be wrong. Call them intuitions if you like, but I've got 'em, and they are real.

There is no more abused term than "direct knowledge" in philosophy. You have no idea of how your brain actually functions. You aren't even aware (normally) that you have a blind spot near the center of your vision. You have "direct knowledge" of your field of vision too, but it fails you there! Your beliefs are real, but as cognitive science has showed repeatedly, your intuitions about those processes simply aren't reliable. **Science accepts that you have these intuitions- it does not accept that they are reliable sources of knowledge.**

Quote:

Originally Posted by **Banno**

Here you could be quoting Searle...Perhaps you should state exactly what it is that you think Searle wishes to claim, since you have already agreed with what I think he wishes to claim. Certainly I don't think he would disagree with your last sentence...

I have no doubt that we agree on some issues. The point is I disagree with him on what his Chinese Room demonstrates. He thinks it demonstrates that one can't have "understanding" without "intrinsic" intentionality and Dennett and I and think it's simply an unreliable intuition. His position is like the creationists that claim it's intuitively obvious that organisms must have been "designed" by an intelligence.

Intuitions, "direct knowledge" and other metaphysical properties like intrinsic intentionality have no explanatory value in cognitive science.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

Now you are just being silly. And why do you feel the need to yell?

Bolding is not yelling, it is just emphasis. Yelling is like this: **THINKING THAT ONE'S INTUITIONS ARE RELIABLE SOURCES OF KNOWLEDGE IS SILLY** (except for certain evolutionary situations like mate selection, personal safety, etc).

By the way, I read the Searle reference you gave and I agree with some of what he said, but here's an example of his

problem. He says 'See, the 'systems reply' says "you really do understand Chinese, you just don't know that you do.'

It's not that we don't know that we understand Chinese, we know that because we know we've established communication with someone using the Chinese language. What we don't know are the brain processes by which we understand Chinese. That's where his intuition fails.

By Probeman (John Donovan)

Quote:

Originally Posted by **Banno**

I do not know how intentional states arise in the brain. But I do believe that they do. I also am convinced, by Searle's argument, that they do not arise from syntactic processes *alone*, since syntax does not provide semantics. The question of how intentional states arise will not be answered by adopting Probeman's strategy of calling them names.

I am a bit confused over this whole debate, perhaps you can clear it up for me.

Searle's argument has always seemed peculiar to me. Where exactly is the line drawn between this syntax and semantics? If I understand Searle's claim correctly, he is saying that a robot could speak and interact with people -as well as people do- but that robot really doesn't understand anything? Is this correct?

Would Searle also say that any given phrase out of the robot's mouth, would be meaningless, and lacking in semantics? I believe that is what he would say.

But what about the following scenario.

What if one of your neurons -just one- became irreparably damaged. A scientist offers to fix it for you, and you agree. He inserts a small (super tiny) chip into your brain, that functions in the same on/off fashion as the damaged neuron. Subsequently, over the next little while, more and more of your neurons are replaced, until eventually your entire brain is composed of these chips.

Since you are now a robot, Searle is forced to claim that the words you are saying are now meaningless, even though every indication we get from your outward behaviour says you are the exact same person. Isn't this a peculiar notion to you?

The intuition for me, is that if I underwent such a process, I would continue to be a thinking being. Searle might not have the same intuition. But Searle has a bigger problem with this, I think, and it is the following: At what point would such a person lose "intentionality"? When would their words simply stop meaning anything, and become just symbols spat out by a robot (even though they appear to be meaningful)

The onus is on Searle to tell us when this happens. But, as you no doubt are aware, Searle cannot answer such a question. there is no logical reason that indicates a particular point, yet Searle's position requires it. This indicates there is something wrong with his position.

We can avoid this problem by taking an alternative stance: the robots words are meaningful.

By Andrew Saunders

Quote:

Originally Posted by **Minty**

But Searle's thought experiment shows this is absurd. Therefore I think we have to hold that the self is non-physical (even if ontologically dependent on processes within the brain).

Gee that's actually an extremely funny and revealing conclusion.

First off, as a matter of sociological fact, there are very few people in the sciences and artificial intelligence who have EVER taken the Chinese Room thought experiment seriously. As far as I'm concerned, it has already been refuted numerous times, starting with the edition of *Behavioral and Brain Sciences* where it first appeared.

By and large, the argument only has seemed to impress non-scientists terrified of materialism. Due to its popularity in that crowd, cognitive scientists writing books on their expertise for the general public have felt compelled to mention it and offer their own dismissals. This happens all the time. I'm constantly reading stuff by such people, and I can't remember a single scientist of consciousness who ever believed the argument had merit (though there must be some out there I've missed). So, when I see philosophers try to resurrect this dinosaur, it always strikes me as an act of desperation.

What's particularly revealing about the reasons people bring it up is that they rarely pay attention to what Searle himself actually says in the infamous paper. They follow Searle as long as he's questioning A.I. and have an amazing ability to blank out every line of the paper that goes against their mystical dualism.

Take Minty's quote—no one who has actually paid attention to the paper would ever suggest that it concludes the self is non-physical. That is certainly not Searle's own conclusion. Several times, he asserts such things as "I believe we are machines" and that consciousness is "due to the causal powers of brains". He even thinks we can build machines that are conscious—from the ground up--so long as it had these causal powers.

Oops.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Minty**

Chaos theory is interesting. Just imagine going back in time one million years, **and just staying there for a split second**. On getting back probably the world would be unrecognisable!

There's an Asimov or Clarke SF story with the same idea- but just a metal sphere goes back in time for a second and the infinitesimal effects cascade amusingly.

Quote:

Originally Posted by **Minty**

The question I'm wondering is why some physical laws should produce consciousness. What is it about those physical laws that could do this? And surely, whatever function the processes in the brain perform, they could **in principle** be captured by some algorithm, even if you add in a bit of innate randomness. Therefore a computer could in principle be conscious.

Yes. That's the conclusion that cognitive science has drawn. As for why physical laws can produce consciousness, why not? In fact the same question can be asked of life too. Why should "mere" physical laws be capable of producing life? In the 19th century it was believed (by intuition and metaphysical argument) that life just had to be more than arrangements of matter, but organic chemistry and Darwin provided a natural explanation that completely eliminated the need for a "vital life force".

Quote:

Originally Posted by **Minty**

But Searle's thought experiment shows this is absurd. Therefore I think we have to hold that the self is non-physical (even if ontologically dependent on processes within the brain).

I disagree. You can claim if you like that Searle's problematic idea is evidence for a non-material mind stuff (though he wouldn't agree). But since we have no evidence for such magical properties in the brain, I would be more inclined to suspect that there is just no such "ghost in the machine".

In fact I was recently told that in a personal communication Searle agreed that a person whose brain had been replaced by chips would still be conscious. But then what is the difference between a robot's brain made of chips and a human's brain made of chips, that produces consciousness? Intel versus Motorola chips?

Quote:

Originally Posted by **Banno**

The point is, of course, that you have not established communication, because you do not know what you said. You just shuffled symbols around, not had a conversation. Extraordinary how, after all this time, you miss this simple point...

This simple point was easily refuted in a recent debate on this forum so I won't repeat it since it's was well put by Socrastein...

<http://forums.philosophyforums.com/...55&postcount=10>

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

I said up above that if one is a materialist, I think they are compelled to be functionalists. I know that Searle calls himself a materialist. But it doesn't seem consistent to me.

I agree Searle's position is inconsistent. He claims to be a materialist yet he feels compelled to hold out for "intrinsic" or "original" Meaning, Understanding and Intentionality. However, these metaphysical problems can be made explained if one is willing to understand that "function" provides the missing ingredient. When that is included in the mix, along with context and interaction, one can see how meaning, understanding and intentionality are merely observable outcomes of mechanistic heuristic algorithms. Searle's position is much like the intelligent design folks that hold out for intrinsic or original "Design", simply because they are terrified of the ultimate implications of unintelligent design by natural selection. Searle is terrified of the ultimate implications of noncausal intention by algorithmic function.

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

I said up above that if one is a materialist, I think they are compelled to be functionalists. I know that Searle calls himself a materialist. But it doesn't seem consistent to me.

I agree Searle's position is inconsistent. He claims to be a materialist yet he feels compelled to hold out for "intrinsic" or "original" Meaning, Understanding and Intentionality. However, these metaphysical problems can be made explained if one is willing to understand that "function" provides the missing ingredient. When that is included in the mix, along with context and interaction, one can see how meaning, understanding and intentionality are merely observable outcomes of mechanistic heuristic algorithms. Searle's position is much like the intelligent design

folks that hold out for intrinsic or original "Design", simply because they are terrified of the ultimate implications of unintelligent design by natural selection. Searle is terrified of the ultimate implications of noncausal intention by algorithmic function.

Note: If the moderator that keeps deleting my response to Mariner would like to edit this response, please have the courtesy to remove his personal attacks against me from his above post as well.

Quote:

Originally Posted by **Mariner**

It is a well-known fact that probeman only ever reads posts that agree with his viewpoint... Of course, probeman did not understand the debate itself. He doesn't want to. He's one of those materialists who are so terrified of losing their faith that they prefer to cling to it in the face of logic. "It is more logical to ditch logic", things like that...

Oh, omniscient one... of course you know all that I read and understand.

I know this much however: that we've all seen your "god of the gaps" tortured logic and your conclusion that everything you don't understand comes from god, but we've moved on. We aren't interested in debating the supernatural powers of ghosts, gods, fairies or santa claus to explain the natural world (and human nature) in this thread. Though I believe you will find plenty of interested parties in the Philosophy of Religion forum.

By Probeman (John Donovan)

Quote:

Originally Posted by **Minty**

I just wish I could see some of these wonderful refutations!

Then do some homework and your wish will be granted.

I got my copy of the original paper and the responses by going to the closest university library. Took me maybe fifteen minutes. I also watch for books written by cognitive scientists and buy a few every year. If you pay attention to what people in the field say about these subjects, the status of the Chinese Room argument as rather disregarded among scientists becomes quite obvious. They have moved on. Too bad so many others haven't.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Reformed Nihilist**

This is a philosophy forum, not a science forum. It shouldn't surprise that cognitive scientists disregard Searle. Cognitive science is based on eliminativism, and the assumption that the mind is material. Many philosophers have refuted eliminativism (Thomas Nagel for example). You may disagree with their refutations, but don't take the lazy man's approach. Tell us why.

Just because we're in a philosophy forum doesn't mean we can't bring up science. Besides, faustus only brought up cognitive scientists as a side note -the first thing he mentioned was getting the original paper and the responses.

As for simply disagreeing with their refutations, responses to the Chinese Room have been offered several times

already on this forum, yet for some reason they are met with ignorance. In the face of logic, proponents of Searle seem to feel that shouting "Intentionality", "Semantics", or "Chinese Room" is a valid defense.

Faustus is absolutely right -the Chinese Room is dead. It has been proved wrong many times. In fact, Searle himself has admitted it to be flawed, and recanted his objection to the system response.

I will sum up the line of argument for you to see why clearly he is wrong:

1) Searle attempts to answer the question "Could a system be said to have thought solely in virtue of running the right program?"

2) Searle proposes the Chinese Room argument. The man inside the room can be said to be running the right program, yet he has no understanding. So, Searle concludes, running the right program cannot guarantee thought.

3) System Response: Searle is cheating -the man doesn't understand, but he is only a part of the system. The system, which is composed of the man, the rulebook, the paper and pencils, might in fact understand Chinese.

4) Searles original response: Have the man internalize the system. Have him memorize the rules, do the calculations in his head. Now the man comprises the whole system, yet still doesn't understand. Therefore the system response is invalid- the whole system doesn't understand Chinese, as opponents have suggested.

5) Why Searle is wrong: there are at least two reasons, but I will present the most convincing. Consider again the question, does the man (once he memorizes the rules and does the calculations in his head) understand Chinese? What are we asking? We can translate this question as "Could a system be said to have thought solely in virtue of running the right program?"

This is the question Searle set out to answer. To simply claim yes, as Searle does, would be simply just stating your conclusion as your premise -we must leave it possible that the man, once memorizing the rules, is endowed with an understanding of Chinese.

To put Searles argument in logical form:

hypothesis one: A system could NOT be said to have thought solely in virtue of being a computer with the right sort of program.

Therefore:

A system could NOT be said to have thought solely in virtue of being a computer with the right sort of program.

That's a bad argument.

By Andrew Saunders

Quote:

Originally Posted by **Reformed Nihilist**

I personally do not support Minty's position. I just dislike the implication that one poster is being lazy, by another poster who is being lazy. Faustus responded to a legitimate request for refutation with an Ad Hominem implication that Minty wasn't well enough educated in the area, rather than giving a reference (it takes 2 seconds to type the name of someone who has refuted it, if you're that well read) or even referencing post numbers in this thread. It's not Faustus' position I was criticizing, it was his discourse.

I agree with you.

The post made by Faustus was uninformative and unhelpful. I was quick to take his side, only because I am fed up with people like Minty, whose response to a long and detailed objection is to say things like "No you're wrong, because of the Chinese Room"

Here's a classic example from Minty:

"But you cannot build an electronic brain from ground up (because of Searle's Chinese room argument)."

By Andrew Saunders

Quote:

Originally Posted by **Reformed Nihilist**

This is a philosophy forum, not a science forum. It shouldn't surprise that cognitive scientists disregard Searle. Cognitive science is based on eliminativism, and the assumption that the mind is material. Many philosophers have refuted eliminativism (Thomas Nagel for example). You may disagree with their refutations, but don't take the lazy man's approach. Tell us why.

Cognitive science (and science in general) doesn't eliminate anything. It completely accepts that we have beliefs about some things. But what science often shows is that some of these intuitive beliefs are mistaken when the evidence is examined.

Do you want evidence that your intuitions are mistaken? Try this experiment (assuming you are not color blind):

Now take a regular deck of playing cards and shuffle them and pick one card without looking at it. Look straight ahead with a fixed stare. Does it seem to you that your whole field of vision (from edge to center to edge) is in color? That's the way it seems to everyone alright. But your eyes can only detect color in a very narrow range of vision and only in the direction you are looking. In fact it's a small cone of angle about 5 to 20 degrees for most people depending somewhat on age.

Take the playing card without looking at it and hold it in your hand to one side as far back as you can with the face of the card towards your head. Now (and this is critical) find a spot on a wall in front of you, say 5-10 feet away and stare fixedly at this point (the center of a wall clock works well but anything will do). Do not avert your gaze from this point- if you do move your eyes at all from the spot, pick another card and start again.

Now slowly, very slowly swing your arm forward to bring the card into your peripheral vision. At some point, probably about 90 degrees from your "stare point" you will probably be able to tell that your hand is vaguely visible and maybe that it's holding something but that's about all. Keep your eye fixed on the point in front of you and keep bringing the card slowly forward. Eventually at some point around 30 degrees from front you will eventually be able to guess whether the card is a face card or not, but not what color it is and certainly not what number it is. Keep swinging your arm slowly forward and when the card is somewhere around 10 degrees from your straight ahead stare point, you might be able to guess what the color is. And it's not until the card is almost directly in front of you that you will be able to determine what suit and number it is.

Don't take my word for this, try the experiment yourself. If you can keep your stare fixed straight ahead you will be able to demonstrate that your entire field of vision being in color is only a belief, unsubstantiated by the evidence. Yet it is a first person direct experiential qualia- the way it seems to you!

There are many intuitive aspects to consciousness that cognitive science has shown to be simply mistaken. Not eliminated- your belief itself is real, yet what your belief attempts to establish is wrong.

By Probeman (John Donovan)

I'm fully capable of refuting the Chinese Room, having done so in two different papers while I was getting my philosophy degree (Saunders adequately describes one avenue). But that was years ago, and since then I've realized

just how little regarded this exercise in arm chair thinking is by the people best able to understand what real challenges to A.I. look like. Plus, it's been dealt with already in this forum. So my attitude is now "It's done with, get over it. Let's talk about something interesting."

I have a similar attitude towards, for instance, creationists. Why take the time to refute something that really doesn't carry any weight, especially when there are plenty of sources out there for people to look at? It's not like I could do better than the folks who replied to Searle in the original BBS article, or those who came after.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Reformed Nihilist**

I personally do not support Minty's position. I just dislike the implication that one poster is being lazy, by another poster who is being lazy. Faustus responded to a legitimate request for refutation with an Ad Hominem implication that Minty wasn't well enough educated in the area, rather than giving a reference (it takes 2 seconds to type the name of someone who has refuted it, if you're that well read) or even referencing post numbers in this thread. It's not Faustus' position I was criticizing, it was his discourse.

Actually several opposing arguments were offered including the "one at a time neuron replacement" intuition pump, but they were dismissed or ignored. I myself have suggested that Minty read a few papers and even supplied the links, but he laughs them off. He's clearly needs to educate himself. If he's only going to offer heartfelt arguments from intuition he won't get far in his understanding. If he were in my class I'd assign him several papers and a written summary. But in this forum it's really up to him to make that effort.

In any case, science is an epistemological method and therefore it is entirely reasonable to have this position presented. Especially since that system of knowledge seems to be the only one that actually has produced useful and convincing results.

By Probeman (John Donovan)

Quote:

Originally Posted by **Reformed Nihilist**

Cognitive science is based on eliminativism . . .

Baloney. Very few cognitive scientists I've ever read endorse eliminativism (I can only name two: Paul and Patricia Churchland). More often than not, if the E-word is invoked, it's merely because the person being accused of endorsing the doctrine refuses to go along with the sorts of conceits Nagel and other mysterians insist upon.

Quote:

Originally Posted by **Reformed Nihilist**

. . .and the assumption that the mind is material.

Do let me know when evidence contradicting this ever appears. And it's evidence, not armchair thought experiments, that interests me.

By Faustus (Brian Peterson)

Quote:

Originally Posted by **Banno**

[*]Computer programs are entirely defined by their formal, syntactic structure

Computer programs are defined by their formal syntactic structure, but computers are not. We are not just computer programs, we are computers that interact with the environment and each other, therefore we can algorithmically gain "understanding" and "significance" of our behavior.

I'd be interested in seeing your response to Saunders "one at a time neuron replacement by chips" thought experiment. When does the person undergoing this procedure stop understanding what they are saying?
By Probeman (John Donovan)

Quote:

Originally Posted by **Reformed Nihilist**

'Mind' is an ill defined term, but my take is that 'mind' refers to a gestalt of subjective things, so I would find it hard to think that there could be any empirical evidence for something that is by definition subjective. I think that was Kant's view as well.

Well, if we're defining a biological phenomenon in such a way as to make it immune to empirical considerations, perhaps we need to abandon such definitions and any approaches that use them.
By Faustus (Brian Peterson)

Quote:

Originally Posted by **Reformed Nihilist**

I see it as neither. It is simply a question of point of view.

Actually this is close to Dennett's "center of narrative gravity" idea. Here is a link you might find interesting:

<http://ase.tufts.edu/cogstud/papers/selfctr.htm>

By Probeman (John Donovan)

Quote:

Originally Posted by **Reformed Nihilist**

I didn't think it was likely that it was a truly original idea, although I find that his use of the word 'fiction' is questionable. It has negative connotations. People would get quite upset (maybe rightly so) if I say that 'infinity', or 'Pi', and most certainly 'existence' are fiction, yet they are concepts that cannot be empirically measured.

I think some of these things can be measured (we can see differences in brain behavior between a conscious and self-aware person and a comatose or not self-aware person for example), but then let's just say these ideas are very, very useful fictions. Or if you prefer, "idealizations", like the "center of gravity" which can be measured based on some set of rules, but is still just a very, very, useful idealization.

By Probeman (John Donovan)

