

urth-urth.net Digest, Vol 5, Issue 41

The Urth list — 13 messages, 6 voices, 27 Jan 2005

<https://urth.darkrealm.vip/thread/urth/4345>

Exported from an unofficial mirror of a public mailing list. Copyright in each message remains with the person who wrote it. Email addresses are obscured as `user at host`, the form the original archive published. This file was rebuilt from parsed fields rather than retained headers, so it is faithful in content but not byte-exact.

turin — 27 Jan 2005, 18:16

turing called it the imitation game.

it is not enough for a machine to produce the same outputs as a person, it must arrive at its outputs in the same fashion, that is whatever the functional apparatus, it must have the same cognition and intentional and phenomenal states as a person. the machine would obviously have to have a certain degree of self recursion or sentience, and something which resembles linear consciousness, it is debatable whether there is any other kind. the imitation game taken literally is the same kind of reductionism as behaviorism. and turing was really just talking about how we identify other people, especially under certain circumstances. the test itself was just the "chinese room", but i wont explain that here. om short, turing is just the jumping off point for cognitive science. for a mahine to be a person it has to have, and in the 40s and 50s this was a taboo subject in psychology, a mind. the rigidity of your criteria for cognition and whether phenomenal states or emotions are even meccesary are something else entirely. "creativity" is just a way to describe solving complex problems when there is incomplete information, indefinite goals, or both. most "real world" problems are like this. what we really mean is that the machine will have to be flexible and proficient in a wide variety of tasks. it also means that the machine will have to be self correcting, possess learning algorithms, and that it will be able to reset its own goals, though its programming would not have to be completely transparent or malleable to itself. ours arent. i'd love ot be able to reset some of my goals. EOT;

Dan'l Danehy-Oakes — 28 Jan 2005, 13:13 · in reply to turin

On Thu, 27 Jan 2005 18:16:28 -0800 (PST), turin <turin at hell.com> wrote:
it is not enough for a machine to produce the same outputs as a person,
it must arrive at its outputs in the same fashion, that is whatever the
functional apparatus, it must have the same cognition and intentional
and phenomenal states as a person.

How would you be able to demonstrate that it had these "states"? How can you define them in humans? To the best of my knowledge neither of these questions has any empirically testable answer to date.

the imitation game taken literally is the same kind of reductionism
as behaviorism.

The imitation game stands as the only empirical test thus far proposed for identifying a machine intelligence, if one should arise. It has several flaws: I find most significant among these its failure to account for the probability that the experiential life of a machine intelligence would differ in important ways from that of a human; thus, it would tend to eliminate a (hypothetical as of today) class of genuinely sentient meachines which were not good at imagining and empathizing with the experiential life of humans.

the test itself was just the "chinese room",

You betray your prejudices by even mentioning Searle's quibble.

for a machine to be a person it has to have, and in the 40s and 50s this was a taboo subject in psychology, a mind.

This grotesquely misrepresents the psychological institutions of the period: this was the heyday of behaviorism, but behaviorism was never the sole dominant school of psychology.

--Dan'I

"We're going to sit on Scorsese's head"

-- The Goodfeathers

James Wynn — 28 Jan 2005, 14:13 · in reply to Dan'I Danehy-Oakes

the imitation game taken literally is the same kind of reductionism as behaviorism.

The imitation game stands as the only empirical test thus far proposed for identifying a machine intelligence, if one should arise.

But the very term "the imitation game" suggests that in tFHoC maybe...just maybe...long before the Long Sun and the Short Sun...Wolfe was considering the Turing's postulates in terms of the validity of apparent consciousness vs real consciousness.

~ Crush

Chris — 28 Jan 2005, 14:30 · in reply to Dan'I Danehy-Oakes

Dan'I said:

Turin said:

the imitation game taken literally is the same kind of reductionism as behaviorism.

The imitation game stands as the only empirical test thus far proposed for identifying a machine intelligence, if one should arise. It has several flaws: I find most significant among these its failure to account for the probability that the experiential life of a machine intelligence would differ in important ways from that of a human; thus, it would tend to eliminate a (hypothetical as of today) class of genuinely sentient machines which were not good at imagining and empathizing with the experiential life of humans.

I think it's possible to accept the test without engaging in full-blown reductionism, although yes, it does smack of behaviorism regardless. Basically by endorsing the test you are making an ought-claim (possibly of the ethical variety, although the ethical isn't the only type of "ought"): that, regardless of the truth behind the currently impenetrable curtain, we *should* make our judgments concerning intelligence based on evidence of behavior. It is, granted, a short step from there to saying that the evidence of behavior *is* the only truth of the matter there is, that there's nothing to see behind the curtain - but this additional step isn't really necessary to accept the test for what it is.

I have often thought about the particular flaw that Dan'I identifies. It seems to me that the class of sentient machines that you're saying would be excluded would really be comparable, all in all, to madmen. A madman might have trouble passing as "sentient" by rigid standards, and yet they are - at least to some extent - somehow identifiable as intelligent, but insane. Labeling a machine as "insane" might be quite unfair, but pragmatically would have to be done. A machine intelligence with perceptions and motivations sufficiently different from our admittedly arbitrary norm would be far more dangerous than any homicidal madman.

the test itself was just the "chinese room",

You betray your prejudices by even mentioning Searle's quibble.

I really don't follow what you mean by this. He may have been mentioning "Searle's quibble" because I brought it up earlier, and I don't think that I have a bias on the matter, nor do I think simply mentioning him in this context is irrelevant or biased.

-- Civet

"You must not be soft, but cover yourself up; for you must discover an intellection abstractional and fraudulent."

Dan'l Danehy-Oakes — 28 Jan 2005, 15:19 · in reply to Chris

My apologies to both Turin and Civet: I had missed the latter's earlier mention of Searle's "Chinese Room."

I do think that the "Chinese Room" is a quibble, and a bad one at that: it involves a particular kind of essentialism that I don't have any use for, and I perceive it as a desperate attempt to defend the epistemologically unique status of the human "mind," in advance of any serious assault made on that status by (for example) a machine capable of winning the Imitation Game with any regularity.

Basically by endorsing the test you are making an ought-claim (possibly of the ethical variety, although the ethical isn't the only type of "ought"): that, regardless of the truth behind the currently impenetrable curtain, we *should* make our judgments concerning intelligence based on evidence of behavior.

I don't think it's an ethical but an epistemological "ought," but I agree with your interpretation here. But then, what evidence - other than some kind of behavior - is available for or against the existence of any intelligence, machine, animal, human, or alien?

I don't think I'm "endorsing" the test, though: I'm simply challenging those who think it's invalid to come up with something better, ideally something that doesn't rely upon intuition. (The Turing test relies on a statistical summation of multiple people's intuitions over many runs; the Searle-Penrose axis rely totally on their intuition that intelligence is "something more" than massively parallel information processing.)

It is, granted, a short step from there to saying that the evidence of behavior *is* the only truth of the matter there is,

It is in fact a huge leap: one, I gather, that some behaviorists have made, but very few. I certainly have never encountered a behaviorist who would deny his or her own consciousness.

I have often thought about the particular flaw that Dan'l identifies. It seems to me that the class of sentient machines that you're saying would be excluded would really be comparable, all in all, to madmen.

H'mmm. No. I'm suggesting that its experiential world might be far more different from mine and mine than that of, say, a sociopath, or even a catatonic: at least as different as the experiential world of one of the "higher" animals. As far as I can tell, a sentient machine would operate from, at least some, entirely different kinds of sensory input from humans, for example; nor would it likely be bashed about by hormonally-induced "moods."

Labeling a machine as "insane" might be quite unfair, but pragmatically would have to be done. A machine intelligence with perceptions and motivations sufficiently different from our admittedly arbitrary norm would

be far more dangerous than any homicidal madman.

Only if it also had the ability to act upon its motivations. I see no reason why such a machine should be granted, for example, arms or any other manipulating appendage.

"We're going to sit on Scorsese's head"
-- The Goodfeathers

maru — 28 Jan 2005, 15:27 · in reply to Dan'l Danehy-Oakes

That's not nearly certain: Some think that emotions have a dual use: first, to provide some a priori biases in the decision-making process, to speed things up (this view derives from certain psychological disorders where emotions are largely stripped out. The sufferers experienced great indecision, and often made very sub-optimal choices.) and as a 'super-rational' evolutionarily developed strategy. Ie, if everyone 'knows' that you do not make rational choices, and are entirely willing to escalate, they will treat you with deference. Sorta like the Russian or N. Korean nuclear strategies :)

~Maru

Dan'l Danehy-Oakes wrote:

H'mmm. No. I'm suggesting that its experiential world might be far more different from thine and mine than that of, say, a sociopath, or even a catatonic: at least as different as the experiential world of one of the "higher" animals. As far as I can tell, a sentient machine would operate from, at least some, entirely different _kinds_ of sensory input from humans, for example; nor would it likely be bashed about by hormonally-induced "moods."

Labeling a machine as "insane" might be quite unfair, but pragmatically would have to be done. A machine intelligence with perceptions and motivations sufficiently different from our admittedly arbitrary norm would be far more dangerous than any homicidal madman.

Only if it also had the ability to act upon its motivations. I see no reason why such a machine should be granted, for example, arms or any other manipulating appendage.

--Blattid

Chris — 28 Jan 2005, 16:30 · in reply to Dan'l Danehy-Oakes

Dan'l said:

I do think that the "Chinese Room" is a quibble, and a bad one at that: it involves a particular kind of essentialism that I don't have any use for, and I perceive it as a desperate attempt to defend the epistemologically unique status of the human "mind," in advance of any serious assault made on that status by (for example) a machine capable of winning the Imitation Game with any regularity.

I don't entirely disagree, although I think, to be fair to Searle, he didn't intend it to be quite the desperation-position it seems to be. In answer to a rhetorical question you ask later on, "what other evidence is available?" Well, there's empirical evidence about the processes that we presume produce our consciousness, which interestingly *would* catch some subset of possible

"weird" intelligences that would never be picked up by a Turing test, so you could say there is a return to Searle's angle there. The very obvious problem with this is that it is difficult/impossible to quantify which

aspects of these processes are the ones that make them essential to intelligence. And nobody's seriously going to fall back on Searle's apparent presumption that, for lack of knowing any better, basically *all* of them are essential.

Basically by endorsing the test you are making an ought-claim (possibly of the ethical variety, although the ethical isn't the only type of "ought"): that, regardless of the truth behind the currently impenetrable curtain, we *should* make our judgments concerning intelligence based on evidence of behavior.

I don't think it's an ethical but an epistemological "ought," but I agree with your interpretation here. But then, what evidence - other than _some_ kind of behavior - is available for or against the existence of _any_ intelligence, machine, animal, human, or alien?

I'm not really certain it's either of the above, which is why I'm pretty wishy-washy here. I think the "ought" can be framed any number of ways, which pretty much determines how you'll then choose to characterize it. I just have a vague suspicion that if you tried to get down to one essential formulation that sums up all the rest, you'd end up with either an aesthetic or ethical claim.

It is, granted, a short step from there to saying that the evidence of behavior *is* the only truth of the matter there is,

It is in fact a huge leap: one, I gather, that some behaviorists have made, but very few. I certainly have never encountered a behaviorist who would deny his or her own consciousness.

Point taken, although I still think it is a relatively easy step to take (not that I find it desirable). The way has been prepared ahead of time by other epistemological arguments.

H'mmm. No. I'm suggesting that its experiential world might be far more different from thine and mine than that of, say, a sociopath, or even a catatonic: at least as different as the experiential world of one of the "higher" animals. As far as I can tell, a sentient machine would operate from, at least some, entirely different _kinds_ of sensory input from humans, for example; nor would it likely be bashed about by hormonally-induced "moods."

I was using, or was trying to use, "mad" in a sort of abstract sense. "Madness" as that which lies outside the realm of intelligibility; incapable of being reduced or expressed logically in human terms. So your garden-variety crazy person is only a little mad in this sense. A machine could do it much better.

Labeling a machine as "insane" might be quite unfair, but pragmatically would have to be done. A machine intelligence with perceptions and motivations sufficiently different from our admittedly arbitrary norm would be far more dangerous than any homicidal madman.

Only if it also had the ability to act upon its motivations. I see no reason why such a machine should be granted, for example, arms or any other manipulating appendage.

The problem is in knowing just how many limitations would be enough to render such a machine harmless. It is, potentially, much much much smarter than any human meatbag.

Dan'l Danehy-Oakes — 30 Jan 2005, 13:54 · in reply to Chris

Turin wrote...

I don't entirely disagree, although I think, to be fair to Searle, he didn't intend it to be quite the desperation-position it seems to be. In answer to a rhetorical question you ask later on, "what other evidence is available?" Well, there's empirical evidence about the processes that we presume produce our consciousness, which interestingly *would* catch some subset of possible "weird" intelligences that would never be picked up by a Turing test, so you could say there is a return to Searle's angle there. The very obvious problem with this is that it is difficult/impossible to quantify which *aspects* of these processes are the ones that make them essential to intelligence.

First of all, to say this involves a set of a priori assumptions, not least of which is that there are some "aspects of the processes" that make them "essential (note this key word) to intelligence." This in turn involves an interesting (to me, at any rate) set of as-of-not unverified (and to my thinking unlikely ever to be verified) assumptions about the *nature* of "the processes that we presume produce our consciousness": specifically, that there exists some empirically-unique set of processes that etc. But as of now (30 Jan 05 @1.30PM PST) I know of no evidence that supports this hypothesis and, in fact, conclude from evidence available-to-me at this time, the (empirically observable) "processes that produce our consciousness" occur in many contexts that do *not* produce consciousness, at least that we can tell. Rather, consciousness appears to occur as an emergent phenomenon, and there is no a priori reason to assume that other physical processes cannot produce similar emergent phenomena. The only "reason" I can see to make such an assumption is the desire to protect the ontologically privileged status of the human nervous system or "soul."

(Incidentally, I brainfarted in my previous post and repeatedly used "epistemological" where I meant "ontological." Sorry for any confusion this may have caused.)

This (in brief) is why I regard the "Chinese room" thoughtexperiment as a desperation move taken by essentialists.

And nobody's seriously going to fall back on Searle's apparent presumption that, for lack of knowing any better, basically *all* of them are essential.

Searle does. And, for all I can tell, so does Penrose. Both seem to me to be engaged in a holding battle for the ontologically unique essence of "human mind." (In fact, I am inclined to agree with their position, but on purely "emotional" grounds: I see not one bit of empirical evidence to justify the position, and so believe it has no proper place in any scientific debate.)

I don't think it's an ethical but an epistemological "ought," but I

Again, this should have been "ontological."

agree with your interpretation here. But then, what evidence - other than *some* kind of behavior - is available for or against the existence of *any* intelligence, machine, animal, human, or alien?

I'm not really certain it's either of the above, which is why I'm pretty wishy-washy here. I think the "ought" can be framed any number of ways, which pretty much determines how you'll then choose to characterize it. I just have a vague suspicion that if you tried to get down to one essential formulation that sums up all the rest, you'd end up with either an aesthetic or ethical claim.

H'mmmm. An interesting point. I definitely do agree that the ontological "ought" here results in ethical issues - i.e., if we are compelled to conclude based on empirical evidence that this artificial system is sentient-sapient-intelligent-conscious (pick your favorite magick word), we have to consider our ethical stance toward it differently than if it is "just" an artifact: we have to consider what ethically might allow or oblige us to regard something other than a human as a "person" in resolving ethical questions. (Which in turn is one of the great values of SF for us today: it doesn't resolve the questions, but it at least gives us a sandbox in which to play with them until a real situation arises, and so be better prepared to deal with them.)

I was using, or was trying to use, "mad" in a sort of abstract sense. "Madness" as that which lies outside the realm of intelligibility; incapable of being reduced or expressed logically in human terms. So your garden-variety crazyperson is only a little mad in this sense. A machine could do it much better.

Point taken and accepted.

Only if it also had the ability to act upon its motivations. I see no reason why such a machine should be granted, for example, arms or any other manipulating appendage.

The problem is in knowing just how many limitations would be enough to render such a machine harmless. It is, potentially, much much much smarter than any human meatbag.

That seems to me another presupposition. I don't assume that it couldn't be, but, as long as we don't know how consciousness and intelligence actually arise in physical systems like brains, we don't know whether any c-like limit to what a consciousness can do exists. There is a reasonable (but not empirically testable at this time) argument to the effect that a single consciousness cannot handle more than some (undetermined) amount of "stuff" at a time without splitting into multiple consciousnesses.

--Dan'l

"We're going to sit on Scorsese's head"
-- The Goodfeathers

maru — 30 Jan 2005, 14:08 · in reply to Dan'l Danehy-Oakes

I'm afraid I'll have to disagree here: there are good reasons to think that our current consciousness is fairly low on the intelligence possibility-totem pole. (I'm not talking about any 'c' -ish limit on intelligence here: as long as current physics hold, intelligences are limited by the Berkenstein bound and the Hubble volume. O' course, the caveat is assuming that physics as we know it has no loopholes. But then we're getting into deep anarchism, and I'd rather not go there. Too many transfinities.) Consider: human brains are massively parallel. They have to be, since neurons are limited to 200 hz. You could very easily bring that up several megs or gigahertz. Without any splitting of consciousness or anything- subjectively the mind simply has more time. You can vastly increase the amount of "stuff" (amount/time is what's relevant) and it makes no difference. Caffeine speeds up the brain slightly, and that shows on IQ tests. Having been 'wired' a few times, I can say my consciousness felt as unitary as ever (Unless that's what each split section *wants* me to think!). The only way I see greater intelligence causing a split is if either the communications lag is so great that each segment thinks enough on its own to be considered another consciousness; or if there is so much computing horsepower that it becomes feasible enough to literally devote a mind to each problem/task.

Dan'l Danehy-Oakes wrote:

That seems to me another presupposition. I don't assume that it _couldn't_ be, but, as long as we don't know how consciousness and intelligence actually arise in physical systems like brains, we don't know whether any _c_-like limit to what a consciousness can do exists. There is a reasonable (but not empirically testable at this time) argument to the effect that a single consciousness cannot handle more than some (undetermined) amount of "stuff" at a time without splitting into multiple consciousnesses.

--Dan'l

maru — 30 Jan 2005, 14:48 · in reply to turin

Why must it be in the same way? If over the range of all realistic outputs, it returns the same output, (which is feasible: Weiner spent considerable time in his 'Cybernetics' and 'God & Golem' working out how functions can be generalized from a set of input. Curiously, this leads to the corollary that machines can reproduce- but that is a tangent.) then it doesn't matter. If the algorithm followed was even more powerful than the human equivalent, would you disqualify it because it was deviating from the God-given human example?

~Maru

turin wrote:
turing called it the imitation game.

it is not enough for a machine to produce the same outputs as a person, it must arrive at its outputs in the same fashion, that is whatever the functional apparatus, it must have the same cognition and intentional and phenomenal states as a person. the machine would obviously have to have a certain degree of self recursion or sentience, and something which resembles linear consciousness, it is debatable whether there is any other kind. the imitation game taken literally is the same kind of reductionism as behaviorism. and turing was really just talking about how we identify other people, especially under certain circumstances. the test itself was just the "chinese room", but i won't explain that here. om short, turing is just the jumping off point for cognitive science. for a machine to be a person it has to have, and in the 40s and 50s this was a taboo subject in psychology, a mind. the rigidity of your criteria for cognition and whether phenomenal states or emotions are even necessary are something else entirely. "creativity" is just a way to describe solving complex problems when there is incomplete information, indefinite goals, or both. most "real world" problems are like this. what we really mean is that the machine will have to be flexible and proficient in a wide variety of tasks. it also means that the machine will have to be self correcting, possess learning algorithms, and that it will be able to reset its own goals, though its programming would not have to be completely transparent or malleable to itself. ours aren't. i'd love it to be able to reset some of my goals. EOT;

maru — 30 Jan 2005, 14:51 · in reply to James Wynn

And there's the crux of the issue: we can never get into the subjective experience of the consciousness, and asking does not suffice. Any decent fake would of course protest that it felt conscious. We can only see the seeming of the thing, never the 'ding an sich', the thing in itself.

~Maru

James Wynn wrote:
But the very term "the imitation game" suggests that in tFHoC maybe...just

maybe...long before the Long Sun and the Short Sun...Wolfe was considering the Turing's postulates in terms of the validity of apparent consciousness vs real consciousness.

~ Crush

Chris — 30 Jan 2005, 22:08 · in reply to Dan'l Danehy-Oakes

Whoa whoa hold up. I am definitely not Turin. Civet or Chris will do.

Turin wrote...

I don't entirely disagree, although I think, to be fair to Searle, he didn't intend it to be quite the desperation-position it seems to be. In answer to a rhetorical question you ask later on, "what other evidence is available?" Well, there's empirical evidence about the processes that we presume produce our consciousness, which interestingly *would* catch some subset of possible "weird" intelligences that would never be picked up by a Turing test, so you could say there is a return to Searle's angle there. The very obvious problem with this is that it is difficult/impossible to quantify which *aspects* of these processes are the ones that make them essential to intelligence.

First of all, to say this involves a set of a priori assumptions, not least of which is that there are some "aspects of the processes" that make them "essential (note this key word) to intelligence."

I'll grant that there are a priori assumptions, although these need not be as extensive as you suggest. In the end not much is required other than the same general assumptions necessary for any causal inquiry. For example, we know that certain creatures (humans) seem to behave as if conscious, but if certain processes centered around the brain stop functioning they cease to behave as if conscious. This doesn't mean that there's no other way to be conscious, but it does suggest that the functions of those processes contribute toward human consciousness, and allows the suggestion that "relevantly similar" processes might also fulfill the same role.

This line of reasoning can lead to hasty generalization depending on the depth of your knowledge of what's really going on with consciousness. Which in our case is pitifully small. But my point was that this avenue could conceivably open up, and there's every reason to believe it will, at least to some extent.

This in turn involves an interesting (to me, at any rate) set of as-of-not unverified (and to my thinking unlikely ever to be verified) assumptions about the *nature* of "the processes that we presume produce our consciousness": specifically, that there exists some empirically-unique set of processes that etc. But as of now (30 Jan 05 @1.30PM PST) I know of no evidence that supports this hypothesis

There is actually a great deal of empirical evidence surrounding you every day: how often do you see something that you believe is conscious without a brain on the average day? Of course,

it takes little imagination to reason hypothetically that the weight of this empirical evidence is systematically skewed by whatever you want to call it - the evolutionary path of this planet, etc. The possibility, or even likelihood, of this logical possibility is what would lead you to disbelieve the hypothesis as you are thinking of it.

But really the a priori assumption here isn't so grandiose. The "aspects of processes" and such are functional descriptions rather than physical, and fixed more or less by definition (although, again, we have a very poor definition of consciousness, but we may get a better one later). As a loose example, to describe something as a "car" you need not specify the exact workings of its engine, but you do need something that **functions** as an "engine" that makes it go.

The only "reason" I can see to make such an assumption is the desire to protect the ontologically privileged status of the human nervous system or "soul."

I guess the best way to sum up here is that I wasn't talking about this, I was simply suggesting that advances in cognitive science could make it possible to identify other possible physical processes capable of producing consciousness, and that this might help identify consciousnesses that wouldn't necessarily be picked up by a test of intelligent behaviour. Nothing logically precludes it, not that any test based on such a method would ever be any more certain of its results than the Turing test would.

(Incidentally, I brainfarted in my previous post and repeatedly used "epistemological" where I meant "ontological." Sorry for any confusion this may have caused.)

It sort of hung together as "epistemological", it didn't seem unreasonable. I don't know about "ontological", though. An ontological claim is a claim about what "is", and "is" claims are distinguished from "ought" claims essentially by definition. I don't think that there can actually **be** such a thing as an ontological "ought".

Searle does. And, for all I can tell, so does Penrose. Both seem to me to be engaged in a holding battle for the ontologically unique essence of "human mind." (In fact, I am inclined to agree with their position, but on purely "emotional" grounds: I see not one bit of empirical evidence to justify the position, and so believe it has no proper place in any scientific debate.)

I am not familiar with Penrose. Searle is not quite so bald about it. There's this whole aspect of the Chinese Room thought experiment that we generally strip out (and thank god we do, because thinking about it makes me want to scream) where supposedly the Chinese Room is a purely syntactical calculator **with no semantics**, and if you accept this supposition then to make a long story short it gives Searle some ground to stand on. But there are 100 pretty good reasons not to accept this, and if you reject it then Searle will appear to be standing in the rain shouting "Intelligence is made by brains!" for no particular reason.

But yeah, he was probably motivated by the same emotional reasons you suggest. He's just not **quite** as sloppy a thinker as all that, or at least not sloppy in quite the same way.

The problem is in knowing just how many limitations would be enough to render such a machine harmless. It is, potentially, much much much smarter than any human meatbag.

That seems to me another presupposition. I don't assume that it *_couldn't_* be, but, as long as we don't know how consciousness and intelligence actually arise in physical systems like brains, we don't know whether any *_c_*-like limit to what a consciousness can do exists. There is a reasonable (but not empirically testable at this time) argument to the effect that a single

consciousness cannot handle more than some (undetermined) amount of "stuff" at a time without splitting into multiple consciousnesses.

--Dan'l

I'm willing to accept this as a hypothesis, but I wouldn't be willing to wager on it.

In any event here we are nicely back in the realm of sci-fi. I am thinking of cautionary tales along the lines of the beginning of Vinge's "Fire Upon the Deep". There's probably a better example, and I know I've read multiple stories that picked up on the idea, but that's the only one I can think of.

Chris

Iorwerth Thomas — 31 Jan 2005, 04:50 · in reply to Chris

Civet wrote...

I am not familiar with Penrose. Searle is not quite so bald about it. There's this whole aspect of the Chinese Room thought experiment that we generally strip out (and thank god we do, because thinking about it makes me want to scream) where supposedly the Chinese Room is a purely syntactical calculator *with no semantics*, and if you accept this supposition then to make a long story short it gives Searle some ground to stand on. But there are 100 pretty good reasons not to accept this, and if you reject it then Searle will appear to be standing in the rain shouting "Intelligence is made by brains!" for no particular reason.

But yeah, he was probably motivated by the same emotional reasons you suggest. He's just not *quite* as sloppy a thinker as all that, or at least not sloppy in quite the same way.

Ok, forgive the ignorance of a physicist... I was under the impression that Searle's argument was intended to attack strong AI, in particular the idea that algorithms alone are sufficient to simulate consciousness without actually enquiring into the nature of the brain, and would therefore only be able to attack a Turing test based on the assumptions of strong AI. I was also under the impression that AI research had moved on from that position towards things like neural networks, which sound more likely to work (largely because they imitate the only conscious things we know of), IMO.

I'm probably wrong, though.

Iorwerth