

Re: Robot Evolution

Source: <http://sci.tech--archive.net/Archive/sci.bio.evolution/2006-12/msg00327.html>

- *From:* "Kent Paul Dolan" <xanthian@xxxxxxxx>
 - *Date:* Wed, 27 Dec 2006 02:03:29 -0500 (EST)
-

Phil Roberts, Jr. wrote:

Tim Tyler wrote:

John Lucas's 'Godel' argument has been much-criticized – and Penrose's views in this area are essentially a variation on it.

I concede that there is a clear majority who disagree with the Lucas/Penrose position. On the other side of the equation, however, we have:

a. Hofstadter, Dennett, Penrose, Clarke and Chaitin, in various ways acknowledging that Godel at least SUGGESTS a disconnect between formalism and mathematical reasoning.

Well, no. Goedel proved a very limited thing about limitations on the power of proof productions generated from systems of axioms "at least as powerful as the Peano postulates". That proof in and of itself had nothing to convey about how humans accomplish mathematical reasoning.

That humans would like to appear somehow superior to such productions lacks a first demonstration to satisfy that liking.

So far, no human has ever proved any mathematical proposition that can be proved unprovable using Goedel's mechanism.

Re: Robot Evolution

Quite on the contrary, again and again humans employ the work done by Goedel to identify _other_ unprovable propositions.

In each case, as with the Hilbert program or, after the fact, Frege's work, such a "proof of unprovability" serves as an excellent "stop wasting time here" signpost [not that a universe of kooks won't keep trying to square the circle or trisect the angle, not all are literate enough to read "give up, what you seek cannot be done" and take it as solid advice].

b. Little unanimity as to what exactly is wrong with the Godel argument, with dozens and dozens of different sorts of objections, many based on impenetrable confabulations.

There's nothing "wrong with" the Goedel argument, unless you and I are talking at cross purposes and you mean something other than Goedel's proof.

That proof is simple enough to follow that anyone with a minimally competent education in mathematics can work through it in an evening, and once worked through, it becomes "self-evident"; too simple to challenge.

If on the contrary you are using "the Goedel argument" as the false hypothesis that humans are doing something "super-computational" in their mathematical reasoning, I'll just refer you back to my prior posting: no, we aren't.

c. Papers still being published criticizing the Godel argument against mechanism almost 80 years after Godel first published his theorem.

So? Twenty years ago or more, "special relativity scientists" supposedly universally agreed to stop using "mass" to describe the combination of rest mass and energy of motion represented as mass [because the term used that way is dependent on the frame of reference of the observer], yet just last year there was a paper published arguing that "scientists" should do just what, for the most part, they have done, change their terminology.

"Publish or perish" is a terrible blight, since it results in so much re-publication of what is already known.

d. The universal abandonment of Hilbert's program of formalizing mathematical reasoning by mathematicians all over the world subsequent to Godel's proof.

Right -- Goedel proved that the goal was unattainable, everyone who could read his proof with understanding immediately realized that he was correct, and diverted their attention elsewhere.

This in NO WAY argues for the superiority of human mathematical reasoning to computational mathematical reasoning.

To the contrary, each mathematician who changed courses was thereby agreeing that indeed his/her work was not going to somehow "beat Goedel", and that trying to do so is a fool's game.

Each such diversion was a vote by an intelligent and well informed participant that human reasoning was limited to the same limits as "computational reasoning".

e. Intersubjectively reproducible empirical evidence (feelings of worthlessness) suggesting that not even Mother Nature herself seems to be able to constrain rationality within a formalism (the program for "trying to stay alive").

That's not, and never has been, how nature "programs" species. Evolution doesn't work that way. In fact, it cannot, since it works at the level of gene allele frequencies of occurrence in an entire population.

Protection and propagation of genes shared in common with ones own genome can lead to some wonderfully counterintuitive behaviors. Study any good writeup on the genetic basis of altruism for more details.

Re: Robot Evolution

f. Evidence (e.g., Parfit, 'Reasons and Persons', p. 12) that any theory that attempts to constrain rationality within a formal structure (e.g., a fixed objective) can be shown to sanction rational irrationality (i.e., can be shown to be self-defeating).

But that is still agreement with Goedel's proof, recast as "there are some good behaviors that cannot be proven (computationally) to be good behaviors".

Again, no argument that humans are "better than computational", more an argument for some findings of game theory, that sometimes the only way to win is to randomize your choices in some careful way.

If you always jink left in fleeing from the tiger, the tiger will learn the shortcut that makes you lunch. If you randomize your choices between left and right options, the tiger cannot improve over following your trail as you followed it, and your chance of making the tiger into a robe improves.

Brief version of what's wrong:

"A mathematician often makes judgments about what mathematical statements are true. If he or she is not more powerful than a computer, then in principle one could write a (very complex) computer program that exactly duplicated his or her behavior.

Assumes what is being questioned.

No, it doesn't. That's the subjective portion of a longer argument.

But any program that infers

Programs don't infer,

Before correcting the language of someone speaking in his own field of expertise, you'd be well advised to check your knowledge of the meanings a word can take. Computer programs doing mathematical reasoning, which for the most part use formal deductive logic to draw conclusions, very much _do_ "infer":

<quote>

infer

v 1: reason by deduction; establish by deduction [syn: deduce, deduct, derive]

2: draw from specific cases for more general cases [syn: generalize, generalise, extrapolate]

3: conclude by reasoning; in logic [syn: deduce]

4: guess correctly; solve by guessing; "He guessed the right number of beans in the jar and won the prize" [syn: guess]

5: believe to be the case; "I understand you have no previous experience?" [syn: understand, gather]

</quote>

<http://dict.die.net/infer/>

they model logical relations that have been found to underly human inferences on most occasions.

Nonsense.

As to whether these relations are actually being followed or simply EMBEDDED IN our inferences remains to be seen.

Bafflegab.

mathematical statements can infer no more than can be proved within an equivalent formal system of mathematical axioms and rules of inference,

True,

Then why did you intrude the above line noise?

but Lucas/Penrose assumes we can go beyond this, that the intuiting of mathematical truth is not simply a matter of logical proof:

So? They're wrong, or so all those who voted with their feet when they abandoned the Hilbert program concluded.

Continuing to cite their arguments, without balancing them with the known rebuttals, doesn't seem particularly useful or integrous when they are fairly universally considered to be incorrect.

The immediate consequence is that truth cannot be defined in terms of provability.

Nonsense.

The issue isn't "we must take as truth what we really only have on faith". The issue is (and it is good for humankind to be so humbled) that the list of things we can know to be "true" [identical for intellectually honest persons to the list of things we can prove to be "true"] is severely circumscribed, and there's no wriggle room to go around that circumscription.

We have, for example, no way to prove that the lights in the sky will not spell out "Drink Coca Cola" starting tomorrow, but we can, if we are sane, live our lives in perfect peace without such a proof.

The supposition that pretty much everything should be subject to being known is an especially pernicious form of hubris in a universe where Heisenberg's uncertainty principle holds sway.

In any serious intellectual endeavor we shall be able to formulate simple mathematical arguments, and thus shall be subject to Godel's incompleteness theorem. However far we go in formalizing our canons of proof, we shall be able to devise propositions which are not, according to those canons, provable, but are none the less, true. So it is one thing

Re: Robot Evolution

to be provable, and a different thing to be true. Truth outruns provability. (J.R. Lucas).

Lucas makes the mistake of assuming that truth exists separate from provability. This is of course the error of theism. What Goedel proved is that there are propositions whose truth or falsehood cannot be determined.

The assumption that some of those propositions must therefore be "true ones" misunderstands what "true" should mean.

Lucas wants it to have a separate meaning from "provable", but I cannot see how allowing such a separate meaning can be anything but incredibly dangerous.

Allowing that, allows frauds and mountebanks to assign the token "true" to any proposition which they can cast as a Goedel-style unprovable proposition, merely on their self-interested say-so.

This argument won't fly if the set of axioms to which the human mathematician is formally equivalent is too complex for the human to understand.

What is the basis for the assumption that the intuiting of mathematical truth is based on a set of axioms, let alone that they must be too complex to understand?

That was the assumption to which you objected above as:

> Assumes what is being questioned.

on which this ongoing argument is being still being argued by the prior poster.

"If human reasoning about mathematics is computational, then the argument that the human can somehow understand and then by computational means exceed his own axiom set is incorrect if that axiom

set is too complex for the human to understand" —
to summarize badly.

[It's worth noticing that any computational device, including the human mind, is probably storage-limited from entirely understanding its own operation, in any case. It would probably need to know, for example, the bit by bit storage reliability (probability of failure) of its entire storage mechanism to accomplish such a task, more than a bit of data per bit of available storage.]

These are amazing claims, which Penrose hardly bothers to defend. Reviewers knowledgeable about Godel's work, however, have simply pointed out that an axiom system can infer that if its axioms are self-consistent,

An axiom system can infer?

That's shorthand. The longhand is something like "any correct production of sentential calculus using the rules of logic and axioms of a logical system can infer". Live with the usual form, please, rather than arguing vocabulary to distract from the weakness of your arguments.

then its Godel sentence is true.

That's a very strong claim of equivalence.

An axiom system just can't determine its own self-consistency.

That's just another way of stating what Goedel proved (as Tim already knows, I'm just saying that for the rest of the audience).

But then neither can human mathematicians know whether the axioms they explicitly favor (much

Re: Robot Evolution

less the axioms they are formally equivalent to) are self-consistent.

That's a bit overstated; if you can derive a contradiction from them, they definitely are inconsistent. Once proved inconsistent in that way, one knows them to be inconsistent.

The problem is that no guaranteed plan for deriving such a contradiction exists; to claim that one does violates Goedel's proof.

To the best of my understanding, the situation with which one is confronted is that mathematics becomes just another branch of science whose theorems are subject to issues of falsifiability.

If what you proved leads to a contradiction, and if your proof is formally correct when cast in terms of your base axioms, then your system of axioms is inconsistent and what is derived from it is all cast into suspicion.

The nationfulls of mansions of mathematics that will fall if the Riemann hypothesis proves incorrect would be hard to count by today, I'd guess, but mathematicians have, slowly, gotten used to dealing with such uncertainty.

Cantor and Frege's proposed axioms of set theory turned out to be inconsistent, and this sort of thing will undoubtedly happen again."

Agreed.

Why then do you go on to say:

But we can nonetheless "know" them to be true in the sense that we all agree we have good reason to believe.

You are undergoing a massive failure to understand what it meant for Frege's system of axioms to be

"inconsistent".

That's not just something we can, or he could, or his colleagues would, let slide by with a "but we know it's true anyway, nudge, nudge, wink, wink".

Starting from an inconsistent set of axioms, anything at all can be proved.

That his axioms proved to be inconsistent, meant that his life's work crumbled like a house of cards, and he died self-perceived to be a failure [despite that today is he greatly respected for having established much of the logical and philosophical foundation of the concept of "number", up until then a concept with no particular rigor attached to it].

That partial success in achieving rigor didn't leave the set of his axioms that proved inconsistent the least bit acceptable to mathematicians, then or now.

As to what this has to do with evolution – if humans can do things no machine can do – or will ever be able to do – that may impact the hypothesis that machine-based organisms may replace humans as the dominant life form on earth over the next century or so.

And on the contrary, if humans are "computational" in a very rich sense, then a steady program of replicating those computational capabilities into integrated mechanisms is a reasonable prospect to accomplish, if not our replacement, then our augmentation by peer much faster intellects, especially as such a program can feed on its own successes as it goes along, by using its outputs as intellect augmenters to create the next generation of outputs.

More importantly, it would mean that there is reason to suspect that E. O. Wilson may have gotten it wrong in asserting genetic determinism:

Can the cultural evolution of higher ethical values gain a direction and momentum its own

Re: Robot Evolution

and completely replace genetic evolution? I think not. The genes hold culture on a leash. The leash is very long, but inevitably values will be constrained in accordance with their effects on the human gene pool (E. O. Wilson).

That is a wholly separate issue, a discussion whose demesne is meme-space, not cyberspace.

Cultural evolution is not about robots [except perhaps by coincidence of mechanisms likely to be employed by cultural evolution], but about evolution at the level of ideas despite slowing (by medical advances and pacification of the environment) of genetic evolution to something running short on forces of natural selection.

I find that Wilson is perhaps correct that we cannot become more moral than our genome will allow, but long term confinement of the most massively immoral of our population is an extremely effective contraceptive.

So, that particular kind of genetic evolution may perhaps be among the most strongly "unnaturally" selected ones which humankind is presently experiencing [though I don't know quite how well "chop off the distribution tail" works in natural selection as opposed to "bias the whole population distribution" mechanisms].

and that Dawkins may have actually gotten it right in asserting the converse:

We, alone on earth, can rebel against the tyranny of the selfish replicators" (Dawkins, 1976, p. 215).

The most immediate question would be: as a beneficiary of the dice tosses of the selfish replicators, how wise or secure should we feel in engaging in such a rebellion?

I know that there are many, many misfeatures in selfish replication. Forced sex will

easily stand for the lot as it is fairly easy to comprehend.

But we need to be very careful in what we are willing to give away in forbidding the best rapists their reproductive spoils: bigger, stronger, smarter, healthier, sneakier, more dexterous — how much of that can humanity afford to forego promoting in its genome?

That's flame bait, and I have a daughter and a granddaughter that make that an uncomfortable issue to raise, but still, it exists whether we discuss it or not.

The other side of that issue of course is that the target who successfully resists forced sex gets (usually "her") choice of mates, and may select one on other, better criteria for "best genes to propagate in a cooperative society", since we humans are the biggest part of our own environment these days.

We, as a society, may (and usually do) choose to remove forcers of sex from the gene pool, but we should be doing that advisedly, with evolutionary as well as social consequences fully exposed to discussion.

It would be ironic if humankind in rebelling against its "selfish replicators" hastened its own extinction by losing competition to species choosing and undergoing no such rebellion, species like fire ants or killer bees, for example.

However, this particular argument for the qualitative superiority of humans is simply wrong – and (IMO) rather obviously so for anyone who knows anything about Godel's work.

Why then is one of the papers you referenced written in 2004? Shouldn't this have all been over and done with decades ago for a flaw that is so "obvious"?

See above in my response to your item "c". This situation is rife throughout science. Second raters rewrite what first raters first publish in more arcane terms, and sometimes decades later. Sometimes, also, papers merely summarize the "state of an art" and are a more convenient citation than are the original sources.

[quote from Penrose]

The many arguments that computationalists and other people have presented for wriggling around Godel's original argument have become known to me only comparatively recently; perhaps we act and perceive according to an unknowable algorithm, perhaps our mathematical understanding is intrinsically unsound, perhaps we could know the algorithms according to which we understand mathematics, but are incapable of knowing the actual roles that these algorithms play. All right, these are logical possibilities. But are they really plausible explanations?

For those who are wedded to computationalism, explanations of this nature may indeed seem plausible. But why should we be wedded to computationalism?

Primarily because the alternative seems to be mysticism, and descents into mysticism are the historical stopping point of scientific progress.

I do not know why so many people seem to be.

Yet, some apparently hold to such a view with almost religious fervour.

Typical theistic BS, trying to call "atheism" a religion, trying to call maintaining that "humans compute like machines do" a religious claim.

Each is of course the exact opposite, but tarring them with the historical record of abject idiocy of theism reduces them in the minds of naive observers of the discussion to "mere alternatives" rather than the "deliberate diametric opposites" they are.

(Indeed, they may often resort to unreasonable rudeness when they feel this position to be threatened!)

Like Penrose accusing anyone who opposes his view of exercising "religious fervor"? The irony is intense here.

Perhaps computationalism can indeed explain the facts of human mentality -- but perhaps it cannot. It is a matter for dispassionate discussion, and certainly not for abuse!

Yet when only mysticism is proposed as an alternative, why would Occam's Razor not come into immediate employment?

I find it curious, also, that even those who argue passionately may take for granted that computationalism in some form -- at least for the objective physical universe -- HAS to be correct. Accordingly, any argument which seems to show otherwise MUST have a "flaw" in it.

Certainly any argument whose only basis is mysticism does have a flaw in it, by that very fact.

Why would Penrose contend otherwise?

Even Chalmers, in his carefully reasoned commentary, seeks out "the deepest flaw in the Godelian arguments". There seems to be the presumption that whatever form of the argument is presented, it just HAS to be flawed. Very few people seem to take seriously the slightest possibility that the argument might perhaps, at root, be correct! This I certainly find

puzzling.

Okay, for sure here he is using "Goedelian arguments" to be shorthand standing for "because Goedel is correct, therefore human (mathematical) reasoning must be other (and better) than computational".

That's incredibly sloppy usage, seemingly attributing to Goedel something he had no part in formulating, but merely inspired by work elsewhere.

As to his finding opposition to descent into mysticism among scientists (mathematicians, here) "puzzling", he need merely re-read a history of the end of the life of Galileo to cure his puzzlement.

Mysticism is the avowed enemy of science, and not the least bit shy when allowed to acquire power about imposing the death penalty for use of science in preference to mysticism.

Nevertheless, I know of many who (like myself) do find the simple "bare" form of the Godelian argument to be very persuasive. To such people, the long and sometimes tortuous arguments that I have provided in 'Shadows of the Mind' may not add much to the case -- in fact, some have told me that they think that they effectively weaken it! It might seem that if I need to go to lengths such as that, the case must surely be a flimsy one. (Even Feferman, from his own particular non-computational standpoint, argues that my diligent efforts may be "largely wasted!") Yet, I would claim that some progress has been made.

Why? Mysticism has been promoted, benefiting exactly whom?

I am struck by the fact that none of the present commentators has chosen to dispute my conclusion G (in 'Shadows', p. 76) that "Human mathematicians are not using a knowably sound algorithm in order to ascertain mathematical truth". (Roger Penrose, 'Psyche' Vol 2)

Nothing humans do is "knowably sound" [among other reasons, because there is no absolute metric for "right behavior"]. Why should mathematics be an exception?

The success of evolutionary algorithms in optima search shows beyond refutation that "looking for an answer by wandering around lost" is a perfectly functional mechanism if one applies appropriate biases to the process.

Peer review, collegial cooperation, and "backtracking search on perceived failure" are several splendid such biasing mechanisms that turn a random search methodology into a frequently converging one.

That's how meme evolution happens, among many other ways.

FWIW

xanthian.

.