


Conscious Entities

Laughing computers

November 24, 2007



We've discussed here previously the enigma of what grief is for; but almost equally puzzling is the function of laughter. Apparently laughter is not unique to human beings, although in chimps and other animals the physical symptoms of hilarity do not necessarily resemble the human ones very closely. Without going overboard on evolutionary explanation, it does seem that such a noteworthy piece of behaviour must have some survival value, but it's not easy to see what a series of involuntary and convulsive vocalisations, possibly accompanied by watering eyes and general incapacitation, is doing for us. Shared laughter undoubtedly helps build social solidarity and good feeling, but surely a bit of a hug would be fine for that purpose - what's with the *cachinnation*?

Igor M. Suslov has a [paper](#) out, building on earlier thoughts, which presents an attempt to explain humour and its function. He thinks it would be feasible for a computer to appreciate jokes in the same way as human beings; but the implication of his theory seems to be that a sophisticated computer - certainly one designed to do the kind of thinking humans do - would actually *have to* laugh.

Suslov's theory draws on the idea (not a new one) that humour arises from the sudden perception of incongruity and the resulting rapid shift of interpretation. When cognitive processes attain a certain level of sophistication, the brain is faced with many circumstances where there are competing interpretations of its sensory input. Is that a bear over there, or just a bush? The brain has to plump for one reading - it can't delay presenting a view to consciousness until further observations have resolved the ambiguity for obvious practical reasons - and it constructs its expectations about the future flow of events on that basis: but it has the capacity to retain one or two competing interpretations in the background just in case. In fact, according to Suslov, it holds a number of branching future possibilities in mind at any one time.

The brain's choice of scenario can only be based on an assessment of probability, so it is inevitably wrong on occasion - hey, it's not a bear, after all! In principle, the brain could wait for the currently assumed scenario to drain away naturally when it reached its current end: but the disadvantages of realising one's error slowly are obvious. Theoretically another alternative would be to delete all recollection of the original mistake: but the best approach seems to be to tolerate the fact that our beliefs about the bush conflict with what we remember believing. The sudden deletion of the original interpretation is the source of the humorous effect.

Suslov has drawn on the views of Spencer, which had it that actual physical laughter was caused by the discharge of nervous energy from mental process into the muscles. This theory, once popular, suffered the defect that there really is no such thing as 'nervous energy' which behaves in this

pseudo-hydraulic style; but Suslov thinks it can be at least partially resurrected if we think of the process as excess energy arising from the clearance of large sections of a neural network (when a scenario is deleted). He recognises that this is still not really an accurate biological description of the way neurons work, but he evidently still thinks there's an underlying truth in it.

One further point is necessary to the plausibility of the theory, namely that humour can be driven out by other factors. We may laugh when we realise the 'bear' is really a bush, but not when we make the reverse discovery. This is because the 'nervous energy', if we can continue to use that term, is directed into other emotions, and hence goes on to power shaking with fear rather than laughter. Suslov goes on to explain a number of other features of humour in terms of his theory with a fair degree of success.

An interesting consequence if all this were true, it seems to me, is that a network-based simulation of human consciousness would also necessarily be subject to sudden discharges. It seems to me this could go two ways. Either the successful engineers are going to notice this curious and possibly damaging property of their networks, or at some stage they are going to encounter problems (the frame problem?) which can in the end only be solved by building in a special rapid-delete facility with a special provision for the tolerance of inconsistency. Use of this facility would amount to the machine laughing.

Would it, though? There would be no need, from an engineering point of view, to build in any sound effects or 'heave with laughter' motors. Would the machine enjoy laughing, and seek out reasons to laugh? There seems no reason to think so, and it is a definite weakness of the theory that it doesn't really explain why humour is anything other than a neutral-to-unpleasant kind of involuntary shudder. Suslov more or less dismisses the pleasurable element in humour: it's more or less a matter of chance, he suggests, just as sneezing happens to be pleasant without that being the point of it. It's true that humans are good at taking pleasure in things that don't seem fun at first sight; making the capsaicin which is designed to deter animals from eating peppers into the very thing that makes them taste good, for example. But it's hard to accept that funny things are only pleasant by chance; it seems an essential feature of humour is being left on one side.

It's also possible to doubt whether *all* humour is a matter of conflicting interpretations. It's true that jokes typically work by suddenly presenting a reinterpretation of what has gone before. Suslov claims that tickling works in a similar way - our expectations about where the sensation is coming from next are constantly falsified. Are we also prepared to say that the sight of someone slipping on a banana skin is funny because it upsets our expectations? That might be part of it: but if conflicting interpretations are the essence of humour, optically ambiguous figures like the Necker cube should be amusing and binocular rivalry ought to be hilarious.

There are of course plenty of technical issues too, apart from the inherent doubtfulness of whether the metaphor of 'nervous energy' can really be given a definite neurological meaning.

One aspect of Suslov's ideas ought to be testable. It's a requirement of the theory that the discarded interpretation is *deleted*, otherwise there is no surplus 'nervous energy'. But why shouldn't it simply recede to the status of alternative hypothesis? That seems a more natural outcome. If that were what happened, we should be ready to change our minds back as quickly as we changed them the first time: if Suslov is right and the discarded reading is actually deleted, we should find it difficult to switch back to the 'bear' hypothesis once we've displaced it with the 'bush' reading. That ought to show up in a greater amount of time needed for the second change of mind. I doubt whether experiments would find that this extra delay actually occurs.

10 Comments »

1. I'm just thinking of all the times when the brain deals with situations which, according to the theory you describe, ought to be funny. The 'duck/rabbit' for example. It would be interesting

to see which bits of the brain kick off when, for example, 'getting' a pun. I'm sure this has been done but I've not seen it.

Anyway, just found this website - I have not read much but the spread seems really impressive.

Cheers, Matt

Comment by Matt — November 24, 2007 @ 8:29 pm

2. Thanks. Yes, ambiguity alone doesn't seem to be it, somehow. It seems there's some vital element missing - I can't quite put my finger on what, though.

Comment by Peter — November 26, 2007 @ 7:36 am

3. Returning to the issue, it does rather seem to be an 'agent level' phenomenon that ties into human interaction. I cannot think of anything funny that doesn't track back to interaction. I wonder if it a learned thing... *drags out copy of The Investigations*

Comment by Matt — November 26, 2007 @ 5:54 pm

4. As a concept, 'deleting a scenario' might be nearly the right concept, though 'flipping between scenarios' could be closer - the perceived probabilities being suddenly drastically revised or reversed.

'Discharge of nervous energy' as a concept, meanwhile, seems only marginally helpful. My take would be this: laughter at such moments 'covers up' a potentially damaging sense of error. Pleasure is induced in a hurry to dismiss or replace what could otherwise be an incorrectly negative learning experience.

Comment by Rollo — November 26, 2007 @ 11:31 pm

5. Hmm, yes, that's an interesting idea. Humour does seem to go with embarrassment and misfortune, and that would tend to explain why mere ambiguity isn't funny.

I suppose you could then say that that the actual display - the cachinnation - is a mechanism for recruiting others to share in the dismissal or replacement.

Comment by Peter — November 28, 2007 @ 10:00 pm

6. Hang on a moment...

isn't there a danger that we are invoking some rather bright homunculi here? Surely, 'recruiting others' and avoiding 'incorrect learning experiences' are the sort of thing that only self conscious agents get up to? If these were the mechanisms of humour then surely we would know they were for the simple reason that we would be aware, as self conscious agents, that that was why we laughed - unless we invoke some other hidden agent to do this bit for us. I have to admit that I'm no great fan of homunculi!

Comment by Matt — November 29, 2007 @ 11:22 pm

7. why we have laughter might still be a mystery, but the causes/effects are not. there was a book (or article, forgot which) that goes into detail about this. they have studied stuff like lesions to

the brain that can cause someone to laugh uncontrollably (and die from laughing) to the effects of nitrous oxide (laughing gas).

Comment by Devy — November 30, 2007 @ 5:18 am

8. I don't think we mean to get homuncular. When I talk about recruiting others, I have in mind the sort of thing that happens with yawning: one person does it and then everyone in the group starts to do the same. I imagine this has a purpose in some sense - perhaps it's valuable in getting everyone to sleep at the same time or something - but not a conscious one. You don't (normally) yawn with the deliberate intention of making other people sleepy.

It's true we do know a bit about the neurology of laughter. Here are a couple of [links](#) that might be worth a look. It's one of the weaknesses of Suslov's theory that it doesn't really make much sense in neurological terms.

If he were backed into a corner, I expect he might say that the bits of the brain that evoke laughing are like a drain, or like a special neural clearance mechanism. But I don't really want to defend that idea because it doesn't seem very plausible to me!

Comment by Peter — November 30, 2007 @ 7:31 am

9. I agree that you don't normally yawn with the intention of influencing other people. Personally I would not consider a natural yawn or a natural laugh to be any sort of intentional action. (The fact that we can intentionally simulate both yawns and laughter is a red herring.)

I was more concerned with the statement: "laughter at such moments 'covers up' a potentially damaging sense of error. Pleasure is induced in a hurry to dismiss or replace what could otherwise be an incorrectly negative learning experience." Given that laughter is not intentional I cannot see who is doing such intentional things as 'covering up' and 'inducing pleasure in a hurry'. I cannot see how such purposive activity could be non consciously achieved without recourse to (secret)homunculi.

Moving on, the brain lesions stuff reminds me of Dennett's radio repairman with a radio makes a squealing noise when he removes a part - leading him to label the part 'the squeal inhibitor'.

It seems to me that laughter is one of those issues in which one has to tread particularly carefully as it can only be an artifact of neural architecture, and thus just something that this particular sort of dynamic system does in the process of falling into the lowest energy state. This also includes learned stuff as there is no principled distinction between content and architecture in the parts of the brain that are relevant here. This level is in no way intentional and no intentional descriptions are relevant.

My reason for making this claim (while accepting that everything is instantiated (under a different description) in the brain or local environment (I'm just that sort of monist!)) is that intentional states are conscious states. Thus, if it was intentional we would know we were doing it. I may be wrong but I just don't think that there is anything in the middle.

I know that for a Dennett fan (if you are) that this might seem odd. However, I assume that you would allow that heterophenomenological intentional states are always conscious states. We would still know we were doing it. So whether you are homophenomenological or heterophenomenological (joke) the result is the same.

Frankly, unless the state of the art has moved on a great deal since I was last in a lab we still

cannot explain why twenty node stochastic neural nets do half the things that they do. If there isn't an agent level explanation then we are years away from any premises that will help this discussion.

Comment by Matt — November 30, 2007 @ 11:26 am

10. There need not always be a conflict of interpretation - for example consider President Bush jokes. The theme behind such a joke is that Bush is too dumb in interpreting something said to him in a subtle fashion. And it is considered a joke only if this theme is maintained. If there is a contrary interpretation where Bush is made out to be smarter than he is expected to be, it would not evoke any laughter, but only raised eyebrows.

Comment by Shankar — December 8, 2007 @ 4:45 pm

[RSS feed for comments on this post.](#) [TrackBack URI](#)

Leave a comment

Name

Mail (will not be published)

Website

- Search

- About

- [About 'Conscious Entities'](#)
- ['Disgressed' - my blog.](#)

- People



[David Chalmers](#)



[Daniel Dennett](#)