Inverted Color Spectrums

[This is an email I sent out to friends on Feb. 23, 2010, together with an article by a "machine language" researcher which touches on the issue of inverted color spectrums. –S.H.]

Hi everybody,

Appended below is a short interview with the artificial intelligence researcher (er, ... excuse me, he prefers "machine intelligence" researcher), Tom Mitchell. He discusses a number of interesting things, including the very alarming and dangerous new ability to "read people's minds" to some degree, the growing problem of privacy in this modern technological (and bourgeois!) age, and so forth.

But in these introductory comments I'm going to focus on an issue Mitchell obliquely raises that is far less politically important, but much more interesting philosophically. Namely, the old puzzle about the supposed possibility of "inverted color spectrums": While those of us who are not colorblind in some way all agree that certain things are red and other things are blue, how do we know that the internal experience *that I have* when I see red things is not the same as the internal experience *that you have* when you see blue things, and vice versa?

Yah, I know, this is pretty much useless, idle philosophical speculation of the sort that gives philosophy its bad name! But the puzzle has intrigued many people, and really bothered some of them.

Actually, I think the solution to this puzzle is not really that difficult, and I already have posted my solution (that such inverted color spectrums do NOT exist!) at http://www.massline.org/Dictionary/I.htm#inverted_spectrum and in more generalized form at http://www.massline.org/Dictionary/Q.htm#qualia. I claim the hypothesized "internal subjective experiences" of seeing red or blue (as something beyond simply *perceiving* something as red or blue) also do not exist. I further claim that those who consider such inverted spectrums to be possible have a philosophical idealist point of view, rather than a materialist point of view.

However, the interesting thing in this regard in the Mitchell interview is that he says that neurological science is on the verge of being able to provide an empirical answer to this long-standing puzzle! He says that it is already possible to show that different people have essentially the same concept of "hammer", and that by scanning the brain it is now possible to determine if someone is thinking about a hammer. He suggests that this is not yet quite possible in the case of more complicated concepts such as color adjectives like red, but apparently it won't be long!

I don't really think an empirical demonstration is actually necessary in this case; I think the correct solution can be determined without it. But empirical evidence is always welcome in support of theory (or in some cases to force the alteration of theory).

OK, now comes the fun part, the wager! I'm willing to bet that it soon *will* be possible to

determine whether someone is thinking about (or looking at) a red thing. What this will mean is that there is some one particular brain state that corresponds to the perception (or thought) of red, a brain state that is distinct and different from the one that corresponds to the perception (or thought) of blue. This will show empirically that inverted color spectrums (of the hypothesized sort) do not exist.

Anybody want to take the bet? Say for a bag of nickels?

Scott

From:
http://www.computerworld.com/s/article/346917/The_Grill_Tom_Mitchell?taxonomyId=11&pageNumber=1

The Grill: Tom Mitchell

This Carnegie Mellon researcher predicts a revolution in psychology and neuroscience. What's on his mind is learning what's on yours.

By Robert L. Mitchell
February 22, 2010 06:00 AM ET

Computerworld - *Tom Mitchell is head of the Machine Learning Department at Carnegie Mellon University, where he oversees 45 doctoral students and 20 faculty members. He has spent more than 30 years performing research in the machine learning field. He recently discussed how machine-learning algorithms that analyze MRI data can know what you're thinking, and why enacting strict privacy regulations to bring order to the "Wild West" atmosphere that pervades the data collection business could shut down real-time data mining, for better or worse.*

As a researcher, you're known for your work in machine learning. What's important about your work? The question that defines computer science is, how can we get machines to perform different algorithms, and what algorithms can we write? Machine learning is like that, but with a twist. Instead of hand-coding what the computer does, we train it. We show it examples. Machine learning has to do with how we build computer programs that improve with experience or find trends in historical data that make good predictions in the future. Face recognition, speech recognition and many other kinds of perceptual sensing problems end up being in the sweet spot for machine-learning algorithms.

What breakthroughs have you achieved? We applied machine learning to problems in neuroscience, looking at brain image data. We're starting to understand how the brain uses neural activity to represent the meaning of different words. We've trained a program that can look at

functional MRI images of someone's brain activity and tell whether they're thinking about a house or hammer, for example.

So you can identify, based on a previous analysis of other people's brain image patterns, what object I might be thinking of? That is correct.

What are the practical applications of this ability to read people's minds? We're at the beginning of a revolution in psychology and neuroscience. Suddenly you can look inside the brain and turn what used to be fun philosophical questions into empirical science. We can look inside your brain when you see the color red, and we can look inside my brain when I see the color red, and we can ask, "Is it or is it not the same pattern of neural activity?"

So is my red the same as your red? I don't have the answer for red. I have it for hammer. When you think "hammer," it's the same as when I think "hammer."

Could people be networked to share this information, so that one person knows what the other is thinking? That's not too far out. There are certain medical patients who are "locked in," who don't have the ability to speak and can't move. It's very tedious to try to communicate. A number of people are working on brain-computer interfaces, devices that can allow a person to have their thoughts decoded.

In a recent opinion piece in the journal *Science*, you argued that we need better privacy regulations due to the explosion of personal data. Why is that important to you? Think about the iPhone. You've got a camera, microphone, all of my email goes through there, a GPS locator, an accelerometer that can tell if I'm sitting still or walking. There's a tremendous amount of data about me that can be collected by those sensors.

If you add to that the stuff that's available about you in the cyberworld, your electronic footprint as you go from site to site, there's a tremendous increase in the amount of information being collected.

It's pretty much the Wild West. If a company collects data about you, they get to decide what they do with it.

That was not worth agonizing about 30 years ago, when there wasn't that much data. But [today], if you can get your hands on all of the data about me, it wouldn't be hard to reconstruct my day.

What are the good things that come from mining all that real-time data? Already it's starting to be used for things like traffic detection and reporting. If you go to Google Maps, some roads blink red because they're congested. Google doesn't exactly tell us how that data is collected, but there are some small companies that have deals with Google, providing them information of this type from cell phones.

And the downside to sharing and aggregating all that personal data? Suppose I go into the emergency room tomorrow and I'm diagnosed with H1N1. My cell phone has been with me for

the last week, so the phone company knows where I've been. It can tell where my cell phone location intersected with the cell phone location of other people.

There could be a service that would call and say, "You might be interested to know that Tom, whom you were with in Starbucks yesterday, was just diagnosed with an infectious disease." That service could be offered with data that's already being collected. But it's not being offered, and the privacy issues are apparent in that case.

What are the consequences of inaction? People will react at some point to an overwhelming invasion of privacy and shut it down. Then there will be good things that won't even get on the table to be discussed. We're on a course that may lead to that.

Why should technologists care about data privacy issues? A lot of the debates I've heard are underinformed about the technology available for assisting in protecting privacy. So it's really important that technologists insert themselves into this discussion -- to make sure that when we're weighing all of the trade-offs, we're informed about what the options really are.