

strangeness

sometimes it makes sense
other times you just have to go with it

[Who writes this stuff anyway?](#)

[The Equation of the Devil](#)

[Electronic Hippies and the Police State](#)

[The 99 Gazillion Laws of Robotics](#)

[The Awful Truth About Teaching Math](#)

Who writes this stuff anyway?

By: *irv* on June 4, 2010

The title of this post is something I once heard a newsroom editor yell (in slightly less family-friendly form) while editing the news. Being a sciency type myself, I am most likely to have that feeling when looking over the science news. The headlines reproduced below are from the last few days and I just couldn't resist commenting on them.

[New gene therapy proves effective in treating severe heart failure](#)

You mean, there's such a thing as mild heart failure? For the record, I don't want that either.

[Link identified between lower IQ scores and attempted suicide in men](#)

The key word is "attempted." The smart ones succeed.

[Eyes of cattle may become new windows to detect mad cow disease](#)

Yes. Especially when they're red and glow. Stay away from those cows. (believe it or not, the article actually discusses looking for glowing bits in the retina, under a microscope though. Much less funny when you put it that way).

[First Paper 'Dipstick' Test for Determining Blood Type](#)

Wait - is this a repeat of the one about suicide? (Dipstick test. get it?)

[Apologies may fuel settlement of legal disputes, study says](#)

Well, more so than bullets, I suppose. I want to know what government agency actually spent money to find this out.

[Visual system interprets sign languages](#)

Ummm, isn't that why it's **sign** language and not , y'know, spoken? Also repeat comment above. (Okay, okay. The visual system discussed in the article isn't the human vision system. It's an artificial system. Couldn't the headline have said, "device" or something?)

[Study finds poker players using drugs to enhance performance](#)

The other guy's performance, actually. At least that's what my father taught me. He said

when in a hot game, try to drink less beer than the other folks at the table so you won't play as drunk as they do. Also, repeat comment repeated above.

[What are the most effective strategies for secondary suicide prevention?](#)

Maybe it's just me but I tend to think that the primary suicide would prevent any secondary. Not to mention pretty much everything else. Also, repeat ...

(Note: I wanted to work in a dipstick or mad cow joke on this one but couldn't think of one and still make it to dinner on time. Let me know if you think of a good one.)

Believe it or not, I don't read the science headlines just to make fun of them. Sometimes, though, it's almost mandatory. Usually the articles are not nearly as silly. Usually.

The Equation of the Devil

By: *irv* on September 6, 2009

I've discovered a new science. I'd like to say I founded it or invented it but there are already brilliant people doing interesting work in the field. They just don't know they share a common field.

To begin, consider this story from Wired about a bizarre scientific paper on the development of a [zombie plague](#). The paper itself (link) is a little dry, though it's interesting if you can wade through the math. If not, read the Wired story. The basic idea is this: Some mathematicians (with quite a bit of time on their hands, apparently) developed the math to model the spread of a zombie infection. They concluded that, unless humans respond quickly with extremely large amounts of violence, the zombies win, civilization collapses and the human race is ultimately annihilated.

The paper assumes slow zombies, not fast or smart ones. It seems reasonable that both of those situations would make things harder for humanity, most likely. It also assumes that normal human replacement (birth and death) does not take place, since newborns eventually die and the newly dead are a perpetual source of zombies, which means the zombies win. The paper models multiple scenarios, including medical treatment for zombieism and the effect of quarantine procedures on the spread. Factors considered in developing their solutions include rates of transmission, the outcome of encounters (fights) and the effect on the spread of destroying zombies so that they can no longer spread the infection. In other words, despite the seemingly whimsical nature of the subject, this is real science.

The authors point out that, while a zombie apocalypse may not be likely, the same general principles apply in real world situations, such as a disease that has a period of dormancy. Presumably, this dormancy would mask the infection so that, like with zombies, it would be impossible to know who to worry about much of the time. Their model is more complicated than usual disease models, yet it has some real world applications. That's interesting.

From another point of view, the important thing about this model is that it shows that the human race is not well prepared to withstand the onslaught of an evil force hell bent (so to speak) on causing its destruction. The model does not show what happens when human casualties become so great that the level of organization of the response is degraded. It's an interesting question to consider whether steps could be taken to make civilization more

resilient against chaotic forces. Maybe we'll consider that question some other time.

Another piece of seemingly whimsical or at least offbeat research that fits in with the theme of a new science, was described in an October 2008 Scientific American article about [modeling an evil human](#). This article describes a project at Rensselaer Polytechnic Institute to develop a computer program that could simulate an evil human being. That is, they have been working on an artificial intelligence that embodies all the traits we generally don't want an AI to have (See my previous post [The 99 Gazillion Laws of Robotics](#)). An AI that intentionally causes harm (albeit, within the confines of a computer generated environment).

The idea of modeling an evil intelligence rather than just reading a book about psychopathy (For example the brilliant work of Robert Hare, such as [Without Conscience](#)) is fascinating all by itself. It involves problems of defining evil as well as building human-like responses into software, plus the need to develop some kind of environment for the AI to respond to. If this effort is successful, it will allow experiments in behavior that can not, ethically be performed in the real world. Think about it. If you want to know what stressor's will cause Able to kill Cain (just for a change), you have the problem that if your experiment succeeds, someone dies.

No such problems exist in a purely computer generated world. True, you then run into the problem of checking your results against reality (would Able really kill Cain in an argument about the best thing to feed goats?) but at least you have data to try to check. Without that, you don't have much beyond speculation and hindsight.

This is why I consider these two separate efforts to be a part of a single scientific study of evil. Maybe I'll call it evilology (to distinguish it from evil science, which could mean any science misused). This is a science that has gone beyond merely observing the things that do us harm but uses the tools of mathematics and computer science to actively study those forces. This is a science that is capable of developing real, data based predictions about both behavior and potential responses (kill the zombies or we all die!).

During the height of the Cold War, the CIA had a huge Russia desk, populated by people who spoke Russian and had spent years studying Soviet leaders' interactions, policies and anything else they could get their hands on. The general idea was Sun Tzu's dictum that it is better to know the enemy, and yourself, than to know only one or, worse, neither.

The study of evil, in all its forms, helps us know the enemy, even when we are our own enemies. It also gives us a way to look beyond ourselves, at artificial intelligences (which **are** developing), unusual plagues and maybe even aliens. I'm giving serious thought to how to model an alien invasion (with or without genetically engineered zombies).

The advancement of this study out of the realms of history and psychology into scientific modeling is very much to be encouraged.

Electronic Hippies and the Police State

By: *irv* on May 18, 2009

How can you not love a domain named "cryptohippie.com?"

Okay, so it's a business that sells unusual and interesting services that broadly fall under the heading of "security." I say broadly because this is not the usual anti-virus or hacker proofing

kind of stuff. Check out the website if you like. For now let's just say that CryptoHippie lives up to its name.

What I really want to discuss is CryptoHippie's report on the Electronic Police State, 2008. (Available [here](#)). The title caught my eye immediately, partly because I recently finished a class that included in the reading list a couple books that were chock full of scare stories about that same topic, more or less [See [No Place to Hide](#) by Robert O'Harrow, Jr. and [Darknet: Hollywood's War against the Digital Generation](#) by J.D. Lasica]. The class wasn't quite about that, though. It was about the law as it relates to computer and internet security and privacy (It was also brutal but it looks like I got the A).

Of course, some of what we covered included the hoops the government has to jump through to gather and the way that was changed by the USA PATRIOT Act. Privacy policies and the laws that govern or even require them were also a large part of the class. And other interesting things. Never did the phrase "Electronic Police State" come up. That would be worth another class by itself and I hope to take it one of these days.

The first topic should be **What does "Electronic police state" mean?**

First, what is a "regular" police state? According to Wikipedia, the term "describes a state in which the government exercises rigid and repressive controls over the social, economic and political life of the population" ([Police state](#)). This is a nice start but doesn't tell the half of it. A police state is one where citizens have few, if any, rights. It's a place where they can be arrested at any time with, or without a reason. In the old Soviet Union the crime of committing "anti-soviet activities" (or was it un-Soviet?) was a catchall that could be used to collect dissidents or prostitutes with equal ease (the story goes that it was used against prostitutes because there were no laws against prostitution, since that was said to exist only in decadent western countries like the U.S.A. But that law could be used to nab almost anybody for almost anything, so it worked just fine).

In a police state, you might need to have the equivalent of a passport to travel a couple hundredmiles to visit your sick mother, which you would rather do than call her because, like East Germany before the fall of Communism there might be more people employed to tap citizen phones than to service the lines and keep the system running (note: I don't know the actual numbers but Stasi, the East German secret police, was really big). The upshot is, a police state works hard at keeping people in their place.

According to the Cryptohippie report, the modern electronic police state is a slightly different animal. It is characterized by "State use of electronic technologies to record, organize, search and distribute forensic evidence against its citizens." This is distinguished from the effort to compile electronic dossiers of people's purchases, habits, movements and behavior that is mostly done for marketing purposes (though the government has access to an amazing amount of this kind of data. See the book [No Place to Hide](#), mentioned above). It is also different from having your employer read your emails or monitor your web browsing habits while you work. These things come more under the heading of a surveillance state, which is worth studying, just not what the Electronic Police State is about. So far.

Getting to the meat of the report, the fine and thoughtful people of Cryptohippie developed a list of 17 aspects of an EPS (I'm getting tired of typing it all out) and gathered data (also available at the link given) to score a number of different countries on those factors. From that they computed an average and ranked the countries. The higher the average score, the more oppressive the EPS in that country. The factors include the requirement for personal identity documents, mandatory retention of phone and ISP records, and the ability of the

government to electronically gather more information, with or without warrants, financial tracking, gag orders (such as are associated with the National Security letters that the FBI uses to gather information. The people served with those letters are not allowed to tell the subjects of the investigation, or anyone else, that there even **is** an investigation) and more.

The list looks to me to be heavily U.S.-centric and heavily weighted toward the negative. Where, for example, does the strong protection of the European Union Privacy Directive fit? Possibly under "constitutional protection" though that doesn't sound like a perfect match for it. There is no "good laws about privacy" category and not enough information provided to be sure that such laws are even considered.

I also had some questions about the "Loose Warrants" category, which applies to a lot more than electronic measures. In some countries, warrants may be loose because little evidence is required to justify them. In the U.S., it is more likely that warrants are loose because the judges reviewing them are simply lazy. Does that count?

That is probably just nit-picking, though, I'm sure in the coming years they will refine their list and their method (if they continue to update the report, anyway). One point that is not nit-picking is that there is virtually no information provided about how the numbers for each country in each category were arrived at. There is no way to argue or agree with the numbers because there is no background. That makes this seem more like a propaganda tool than an attempt at actual debate. **This is the truth from on high. Trust us.** It's not that I don't trust them, I just like to see for myself.

The ranking of the countries is interesting but contains few surprises. With a score of 3.588 (out of a possible 5), China is the number one worst EPS, according to this report. Number 2 is a bit of a surprise: North Korea. It's such a poor country, I didn't think they would have enough computers to even show up on the chart! But the number of electronic transactions is not one of the factors considered. It would be too hard to estimate, I imagine. And once estimated, it would have to be divided up by population, or possibly by number of computers.

Hmm. That brings up an interesting point: What proportion of EPS activity is directed against companies rather than individuals? In the U.S., the answer might be quite a bit. The SEC and the FBI both enforce laws against corporate misbehavior and following the paper trail (or email trail) is a standard part of that. RICO investigations (Racketeering Influenced Corrupt Organization, historically used against the mob) also use a lot of electronic material. This may be almost entirely the province of very advanced countries like the U.S.

I don't get the impression the Cryptohippie people considered investigations of organizations vs. investigations of individuals in their calculations. It would be very hard to study but might shed light on how electronic police tactics are used and how they really affect freedom and privacy. I wonder if I could get a grant? ...

Anyway, the United States came in at number 7 with a score of 3.118, just behind The United Kingdom and Russia, which were tied at 3.176. Again, the question of infrastructure is apparent. The United States may score higher than, say, Ecuador, simply because it can. This is the dark side of modernization and economic development. Countries with little expertise in the digital world, will not (yet) score high on EPS measures. It may also be a trifle mis-leading. Just because a country does not have an EPS doesn't mean it's more free than others. It might still have an old-fashioned police state, with rubber hoses and secret dungeons. That's not a flaw in the report. It's just a caution not to read more into it than it says.

The EPS is an important subject that deserves a lot of attention, study and debate. This report, short and vague though it is, is a fascinating start. I hope that next year, Cryptohippie will update the report and provide more insight into the research behind those 17 scores for each country. Until then, it will be left as an exercise for students to each pick a country and submit their own scoring for the class to discuss.

And yes, there will be a test afterwards.

The 99 Gazillion Laws of Robotics

By: *irv* on August 4, 2009

Robots are in the future. They are in the present, of course, but most people today don't consider some preprogrammed floating arm on an assembly line to be a true "robot." We learned what a robot is from science fiction and that's what we're all waiting for, often with dread (Don't think so? Try googling "robot apocalypse." Wait, let me try it first. 139,000 results. Hey, cool! T-shirts!)

Anyway, in anticipation of the day when robots are the smart, helpful servants/terminators of science fiction fame, lots of people have tried to come up with rules that robots could be programmed to follow to make everything better. Obviously the trend began with Isaac Asimov's infamous [3 laws of robotics](#) (Follow the link. I'm not going to repeat them here).

Asimov's laws were pretty good, though his own stories involving them pointed out some flaws at least in potential implementations. Speaking as a programmer, believe me that implementation is an important point with any software. Give 2 programmers the same 3 rules to implement in a very complex system and you will find the two systems do not act quite the same. One programmer checks for compliance at the beginning of a decision, the other checks afterwards. Maybe they have different ways of checking, besides. The outcomes are often the same but there may be huge differences in some situations.

That different people approach the same problem in different ways is just a fact of life that may result in great differences between robot behavior, too. Anyway, because of these and other considerations there have been numerous attempts to update Asimov's laws. For example a hilarious one I found a few years ago (and can't seem to find the link for anymore) expanded the 3 laws to 10 (I think) and claimed to have patented them - thus ensuring no one would ever have the slightest interest in using them, even if they turned out to be perfect.

No set of robotics laws could possibly be **perfect** (see above) and personally I question whether such laws, themselves are even possible. But it's an important exercise to try to figure out how to make robots safe and controllable, you know, to avoid the robot apocalypse. An interesting attempt to update Asimov's laws came out of Ohio State University recently, where some researchers reformulated the laws to make less sense and have even more loopholes than in the original version.

I tried to buy the original paper online but the system was down or something (ironic that the IEEE **Computer** Society not only charges for electronic documents but then makes it impossible to get them) but an article about the paper ([here](#)) reproduces these updated laws and that's enough for now. According to this article, the first law as advanced in the paper reads:

A human may not deploy a robot without the human-robot work system meeting the highest

legal and professional standards of safety and ethics.

This formulation of the law recognizes that the ethical and legal complexities of behavior, especially when robots interact with humans, probably can't be summed up in a single law. However, it fails miserably as a law that robots can use, which is what Asimov's laws of robotics were about. It also throws robot behavior into the realm of lawyers, which is absolutely not a good thing. Using this rule I can easily foresee long complicated EULAs requiring robot owners to hold the manufacturer harmless for any damage due to lapses in ethics or ethical judgment. We have gone from "Thou (robots) shalt not kill" to "It ain't my fault if your robot accidentally offs you. Didn't you read the license?"

There is no known [Turing Test](#) for ethical behavior. How do you certify that robots have ethics without getting someone killed? Maybe we need something like an FDA for robots. Yeah! That's it! We'll let bureaucrats decide!

On second thought, I don't think I want a robot unless I've programmed it myself. At least then I'll know where to find the kill switch.

I won't go through the other two proposed laws because they are just as bad as the first. They sound like nice, reasonable statements of how robots should be made, if excessively vague, but they completely fail to provide any guidance **for the robots themselves**.

In the long run, the argument that, since no one can possibly anticipate every situation a robot will encounter, no rule or set of rules will ever be good enough at forcing them to behave the way we want them to, is irrefutable. Developing more vague rules and hoping someone implements them well is not much of a solution though.

That doesn't mean there is no solution. As the implementation of robotic judgment proceeds, we may need to accept that, rather than giving them laws to obey, we will merely be able to influence them strongly. I have some ideas about how to influence robot behavior. I'll probably write about them in a future post. For now let's just say there's a lot more research to be done.

The Awful Truth About Teaching Math

By: *irv* on June 6, 2009

Every time the power goes out, I have to re-install the driver for my wife's printer. Every time, including (once again) today. I live in the country, beyond the suburbs into cow country, where it seems sometimes that the power goes out every time there's a high wind. It doesn't stay out for very long. Usually no more than ten minutes or so. That's still enough to make me reinstall the driver. Oh! And VMWare player, which I use on my own computer. I figure that's a bug that was probably fixed in a newer version but the last time I tried to upgrade, it completely hosed my network connections. After about 4 hours of fighting with it, I downgraded again. It works.

There's actually a point to this other than just complaining about computers. I make my living (such as it is) with the things. Complaining about them is just part of the job. The bigger point is that, believe it or not, the computer age is still very young and there's a lot we don't fully understand about how to make software operate to our satisfaction. Things that should be easy aren't always and benefits we think we should see sometimes don't materialize.

Which brings me to the subject of a very interesting recent report (actually a thesis)

summarizing studies of how students use software intended to help them learn to do arithmetic word problems. For a short article about the paper, see [here](#). For the paper itself, go [here](#). Three studies are considered. The purpose was to learn about how students interact with educational software when there has been a breakdown situation. That is, when they get the wrong answer, what do they do?

Anyone who has worked user support or who has even been around people who work with computers can probably answer that without the need for an academic study (or three). What do students do when they can't figure out the answer to a problem given them by a computer? They do what everybody does: They blame the computer.

Specifically, when students entered the wrong answer to a problem and the computer rejected it, the first thing they tried was entering the answer **in a different way** - such as using a comma or fraction rather than a decimal point - to see if the computer had simply failed to recognize the syntax they used. I wouldn't be too surprised if, when that didn't work either, they tried rebooting the computer and starting over from the beginning. It's what I'd do.

In another study, the researcher noted that students didn't just read a problem and try to solve it, they tried to understand what the point of the exercise was in the context of school. That is, they didn't think "What math concepts are needed to understand this problem?" They thought, "What does the teacher (in this case an unseen teacher who wrote software) want me to do?"

Not surprisingly, my interpretation of these results is a bit different from that of the researcher, who talked about framing concepts and social understandings. None of that really has anything to do with math. What was going on in these studies was that the students didn't have good enough math skills to see right off why their answers were wrong (if they did, they might have gotten the right answers in the first place), so they tried to game the system instead.

This could show a weakness in the software. It wasn't giving them good enough feedback to understand that their answer really was wrong. From the sound of it, the only feedback the students got was a yes/no type response. Not even a hint as to how to do the problem correctly. Without feedback and with rudimentary ability to do the work, the students can be forgiven for feeling lost (And I mean **the students** can be forgiven. Many of the irate users I used to field calls from should have known better, but that's an entirely different rant). The software used sounds more like it was intended to drill the students in (theoretically) already existing skills. It was not true teaching software. At the current state of development of the market, there is probably little of the latter. It seems likely that to really teach, software would need a quality of artificial intelligence that does not yet exist.

Another possibility is that there is a problem with math education in general, or more accurately in the things that need to be taught for most people to learn math well. Several times when trying to help people with their math homework I've seen that they had the wrong idea about how math problems are done. They thought that, since they could figure out $2+2$ intuitively (meaning, with little or no conscious thought), they should be able to do all math the same way. They conjured answers out of thin air and hoped for the best. They relied on what worked in second grade, apparently feeling that anything harder than that was more trouble than it was worth. This may sound like an attitude that would be found mostly in children but I've found it in adults, too. Probably no one ever showed them the beauty of math.

Both of these explanations would tend to show that there is still no substitute for a good teacher. There's another article about the role of teachers in using software as an aid to teaching math [here](#).

The paper concludes with the common sense suggestion that the claims of companies that sell educational software should be taken with a grain of salt and the use of that software in the classroom thought through very carefully. This I agree with. Computers are wonderful things but they are not magic bullets. Even with computers, math is still challenging to learn.

Someday computers will be better teachers. Hopefully by then they will be able to reinstall their own #^*&\$! drivers.