The Psychology of Computer Insecurity

Peter Gutmann

University of Auckland

Why can’t users get security right?

Users are idiots

Re:Good. (Score:2)

by jrockway (229604) * <jon-nospam@jrock.us> on Wednesday July 12, @01:13AM (#15698213)

> The MiM is the hardest security problem by far there are no easy answers.

Umm, SSL was designed to solve this problem. When you visit your online bank, make sure the cert is valid and that the URL matches the one on your printed bankbook or credit card.

Pretty simple.

(People being too dumb/lazy to check, though, is the hard problem. Fortunately this is evolution at work.)

- Developers build security applications
- Users apply them incorrectly
- Users are idiots
- QED

Why can't users get security right?

OK, so users are irrational

Definition: "Rational"

- How geeks wish that users would behave

Definition : "Irrational"

- ¬ (How geeks wish that users would behave)

Users are "irrational" simply because they don't behave in the manner arbitrarily tagged "rational" that's defined as "How users should be using my software, dammit!"

- This type of "rational" behaviour does sometimes exist… in people with psychiatric disorders
- (Later slides will go into this in more detail)

Why can't users get security right? (ctd)

The field of psychology provides a great deal of insight into how people deal with security, but this resource is rarely used

> The heavenly laws of logic and probability rule the realm of sound reasoning: psychology is assumed to be irrelevant. Only if mistakes are made are psychologists called in to explain how wrong-wired human minds deviate from these laws […] Many textbooks present first the laws of logic and probability as the standard by which to measure human thinking, then data about how people actually think. The discrepancy between the two makes people appear to be irrational
>
> — Gerd Gigerenzer,
>
> "Adaptive Thinking: Rationality in the Real World"

How Users Make Decisions

Economic decision-making model (Bayesian decision-making-model) is based on standard economic thinking

- Goes back to (at least) John von Neumann's work on game theory in the 1940s

Assumes that people always know what they want and will choose the optimal course for getting it

> [This model] took its marching orders from standard American economics, which assumes that people always know what they want and choose the optimal course of action for getting it
>
> — Baruch Fischhoff,
>
> "Decision making in complex systems"

How Users Make Decisions (ctd)

The formalisation of this model, Subjective Expected Utility (SEU) theory, makes the following assumptions about the decision-making process

1. The decision-maker has a utility function that allows them to rank their preferences based on future outcomes
2. The decision-maker has a full and detailed overview of all possible alternative strategies
3. The decision-maker can estimate the probability of occurrence of outcomes for each alternative strategy
4. The decision-maker will choose between alternatives based on their subjective expected utility

How Users Make Decisions (ctd)

To apply the SEU model, execute the following algorithm

```
for each possible decision alternative
   x = all possible consequences of making a
     decision (which includes recursive
     evaluation of any carry-on effects);
   p(x) = quantitative probability for x;
   U(x) = subjective utility of each
     consequence;
   p(x) × U(x) = probability multiplied by
     subjective utility;
   SEU total = Σ_{i=0}^{n} p(xi) × U(xi);
```

SEU Example

Certificate dialog designed for SEU-based decision making

SEU Example

Case study: Evaluate the possibility of a server misconfiguration

SEU Example (ctd)

Can evaluate this based on an evaluation in turn of

- The competence of the remote system's administrators
- The chances that they've made an error
- The chances of a software bug
- …

Assign probabilities and utilities to each of these

- Competence of admins = 0.6
- Subjective utility = 0.85
- …

SEU Example (ctd)

Assign weights to other factors

- Risk of credit card info being phished/misused
- Risk of identity theft
- … *other negative outcomes* …
- Mitigating factors like credit card consumer protection measures
- Intangible factors like the satisfaction of making a purchase
  - (Emotional trauma of not making a purchase if it's someone's birthday)

(Rather lengthy and tedious, particularly since it's an arbitrarily recursive process)

SEU Example (ctd)

Evaluate the sum total to get the selective expected utility for this option

- Then repeat for all other possible options

Finally, pick the option with the highest subjective expected utility value

SEU Example (ctd)

There was a tiny flaw in the plan

What was that, sir?

It was bollocks

Fixing the SEU Model

This method of decision-making requires total omniscience

- This is a quality that's generally lacking in humans

OK, so we'll patch the model by introducing the concept of stopping rules

- Bail out when it's obvious that there's no (cost-effective) benefit to going any further

Fixing the SEU Model (ctd)

How do we decide when to stop?

- Use the SEU model to tell us

Oops

- The stopping-rule patch attempts to model limited search by assuming total omniscience

If stopping rules were practical, you wouldn't be reading this but would be in Las Vegas applying the stopping rule "Stop playing just before you start losing"

How Users *Really* Make Decisions

Two approaches to determining how things really work

1. Empirical evaluation
   - Examine what users do in practice
2. Conceptual modelling
   - Take a set of conceptual models and see which one best approximates reality

How Users *Really* Make Decisions (ctd)

In the 1980s the US DoD sponsored research into improving battlefield decision-making

- Found that people under pressure don't use anything remotely like the economic decision-making model

People under pressure don't weigh up the relative merits of a set of options and choose the most optimal one

- They don't even make a choice from a cut-down subset

They generate one option at a time and take the first one that works

How Users *Really* Make Decisions (ctd)

In evolutionary terms, if a lion turns up in front of a monkey, it runs up the first tree it sees rather than stopping to think about it at length

- It's better to be wrong than to be eaten

Model is called the singular evaluation model

- And various other things: Recognition-primed decision making, heuristic decision making, the take-the-best heuristic, …
- (Shows independent reproducibility, if not consistent nomenclature)

How Users *Really* Make Decisions (ctd)

Singular evaluation is used when

- The decision-maker is under pressure
  - Computer users being prevented from getting a job done are automatically in the "under pressure" category
- The conditions are dynamic
  - The situation may change by the time you've performed a long detailed analysis
- The goals are ill-defined
  - Most users have little grasp of the implications of security-related choices

This is almost a mirror image of the SEU theoretical model!

How Users *Really* Make Decisions (ctd)

Other researchers examined the problem from a conceptual modeling angle

- Which conceptual model best matches how humans make decisions under pressure?

The best conceptual model to actual human behaviour was singular evaluation (under a heuristic name)

- Both empirical and conceptual-modelling approaches reached the same conclusion

How Users *Really* Make Decisions (ctd)

Simple heuristics are popular because it's very difficult to learn from feedback from complex decision-making processes

- Diffusive reinforcement provides insufficient information to single out any one strategy as being particularly effective
- False correlations and biased attributions of success lead to superstition-based decision support

Popularity of "systems" for gambling and stock trading is because participants like to think that they're doing something that's better than just guessing

How Users *Really* Make Decisions (ctd)

Gambling-based decision-making relies on obvious feedback and provides instant gratification

- Coin toss/die roll
- Immediate feedback

Security decision-making doesn't work this way

- No immediate feedback
- No obvious feedback
- Silent failures

How Users *Really* Make Decisions (ctd)

Result: Any action that fails to trigger immediate negative feedback appears to be a win

- This is one reason why users dismiss warning dialogs

How Users *Really* Make Decisions (ctd)

When there's no immediately obvious choice, people's decision-making abilities go downhill rapidly

- Look for some arbitrary distinction, no matter how useless, and go by that
  - "All of these DVD players are near-identical. I'll get this one because it has a karaoke function and the others don't"
- Procrastinate
- Decide based on irrational emotions

This is an appalling way to perform security-related decision making!

How Users *Really* Make Decisions (ctd)

Any form of strong emotion (not just job stress) causes inflexible thinking/decision-making

- External stimuli reduce our ability to gather information and use working memory to sort out the information that we have

Example: Soldiers were trained on how to safely exit a plane

- Overheard a (rehearsed) conversation among the pilots discussing how the plane was about to crash
- Had great difficulty in recalling their instructions
  - (… and needed a change of underwear afterwards)
- Soldiers who weren't exposed to the conversation fared much better

How Users *Really* Make Decisions (ctd)

These limits on reasoning ability are exploited by stage magicians, who anticipate how observers will reason and then choose ways of doing things that fall outside our ability to think of possibilities

- Distraction and misdirection also play a role

Example: How to make an item disappear

- Ask a bunch of people to list all the explanations that they'd have for how the disappearance worked
- Come up with a way of doing it that doesn't involve any of these expectations
- (If the item that disappears is someone else's money or valuables then you didn't learn this strategy here)

It’s not a Bug, it’s a Feature!

The ability to sort out relevant details from the noise is what makes it possible for humans to function

The entire human sensory/information-processing system acts as a series of filters to reduce the vast flow of incoming information to the small amount that’s actually needed

- Did you notice the sensation of the clothes on your skin before you read this bit?

Selective attention processes allow things like the cocktail party phenomenon/source separation problem

It’s not a Bug, it’s a Feature! (ctd)

Imagine if humans couldn’t take shortcuts in reasoning

- They’d never get anything done

AI researchers have already run into this problem

- Programs had to mechanistically grind through vast numbers of implications to come to a conclusion

Known in AI as the frame problem: How do you frame a problem so that it’s practically solvable?

- In problem-solving literature it’s called analysis paralysis

May be a cause of OCD in humans

- People become trapped in a labyrinth of implications
- OCD is a means of dealing with the anxiety that results

It’s not a Bug, it’s a Feature! (ctd)

Researchers have run into the problem of analysis paralysis when evaluating browser security indicators

- Users were asked to switch off heuristic decision-making and carefully verify security indicators to check the validity of a site

This is the standard “best-practice” advice given to users

It’s not a Bug, it’s a Feature! (ctd)

Researchers had to abort the experiment

- Users spent “absurd amounts of time” trying to verify the site’s legitimacy

Switching off singular evaluation lead to a false-positive rate of 63%

- Even with singular evaluation switched off, users still failed to detect 36% of false sites

It’s not a Bug, it’s a Feature! (ctd)

There is one small class of people who use the SEU model for decision-making

- People who have sustained damage to the frontal lobes of the brain

The medical term for this when it’s done deliberately as part of a medical procedure is “lobotomy”

It’s not a Bug, it’s a Feature! (ctd)

Neurology professor Antonio Damasio’s account of SEU decision making in a patient with ventromedial prefrontal lobe damage:

> For the better part of a half-hour, the patient enumerated reasons for and against [the two possible dates for his next appointment ...] he was now walking us through a tiresome cost-benefit analysis, an endless outlining and fruitless comparison of options and possible consequences. It took enormous discipline to listen to all of this without pounding on the table and telling him to stop

It’s not a Bug, it’s a Feature! (ctd)

In extreme cases overanalysis can cause a complete failure to function

- People suffering from somatising catatonic conversion are paralysed by the overhead of having to analyse in infinite detail every decision that they make

Evaluating Heuristic Reasoning

Researchers have run detailed evaluations of the relative performances of different conceptual models

Tests involved applying various strategies to deciding which of two objects scored higher for given criteria

- City populations
- High school dropout rates
- Homelessness rates
- House prices
- Professor’s salaries
- Obesity at age 18
- Fish fertility (!!)
- …

Evaluating Heuristic Reasoning (ctd)

Information available to guide the decision included (for the example of house prices)

- Current property taxes
- Number of bathrooms
- Number of bedrooms
- Property size
- Total living area
- Garage space
- Age of the house
- *… various other factors, up to 18 in total …*
- (Different strategies used different numbers of factors)

Evaluating Heuristic Reasoning (ctd)

Example of heuristic reasoning applied to the city population problem

- For non-US citizens: Does San Diego have a larger population than San Jose?
- For US citizens: Does Munich have a larger population than Dortmund?

People tend to choose San Diego/Munich because they've heard of them

- Better-known → bigger
  - – Hosting a major beer festival doesn't hurt either
- (People do this without even thinking about it)

Evaluating Heuristic Reasoning (ctd)

When researchers compared this with full-blown multiple regression analysis using all 18 factors, M-R was only slightly better than the simple heuristic!

This can't be right…

- Researchers hired independent programming teams in the US and Germany to reproduce the results

Published all of their data so that others could replicate it

- Others got the same result

It's not a Bug, it's a Feature! (ctd)

Why did simple heuristics perform as well as multiple linear regression?

- Linear regression makes use of large numbers of free parameters and assumes that each is relevant
- Problem is known as overfitting
- Simple heuristics reduce overfitting by filtering out noise

It’s not a Bug, it’s a Feature! (ctd)

Overfitting problem was confirmed by investigating how the models handled new data after being fed the training data set (generalisation)

- Performance of linear regression dropped by 12%
  – c.f. an IDS trained with Lincoln Labs test data

Simple heuristics was left as the overall winner

It’s not a Bug, it’s a Feature! (ctd)

Another experiment compared a Bayesian network to simple heuristics

- The ultimate expression of the economic decision-making model

Full-blown Bayesian network performed only marginally better than the simple heuristics, but at a massively higher cost

It’s not a Bug, it’s a Feature! (ctd)

OK, sometimes this can be a bit of a bug…

- Marketers exploit it through techniques like brand recognition
  - Active penetration attack on the human decision-making process
- Consumers use heuristic decision making to choose recognised brands over unrecognised ones
  - See later slides on geeks vs. normal humans for more on this

(Fraudsters and marketers figured this out empirically long before psychologists had explored it)

Conseq.of the Decision-making Process

Psychologists distinguish between two types of actions taken in response to a situation

Controlled processes

- Slow and costly in terms of mental effort
- Provide a great deal of flexibility in handling unexpected situations

Automatic processes

- Quick, little mental overhead
- Acting on autopilot, little control or flexibility

Conseq.of the Decision-making Process (ctd)

Example: Novice vs. experienced drivers

- Novice driver has to manually and consciously check mirrors, change gears, …
- Experienced driver performs these as an automatic process
- Novice drivers deal with this by load-shedding
  - Sacrifice driving speed for steering control

This effect is particularly nasty when it occurs with complex control systems

- Aircraft cockpits (situational awareness problem)
- Nuclear reactors
- …

Conseq.of the Decision-making Process (ctd)

You can experience load-shedding during a controlled process yourself by writing the weekday repeatedly on a piece of paper

- At some point start counting backwards from 100
- Look at what happens to your handwriting quality or speed
  - This is your brain load-shedding

Now try it again, but this time sign your name (automatic process)

Conseq.of the Decision-making Process (ctd)

Automatic processes are people acting on autopilot

- Once the correct stimulus is presented, it's very hard to stop

People click away warning dialogs without thinking

- This is an automatic process, performed without conscious awareness

The action is not only automatic, but people aren't even aware afterwards that they've done it

- "Did I lock the door/leave the iron on/…?"
- Can you recall the driving-related actions you performed while driving to work?

Conseq.of the Decision-making Process (ctd)

Microsoft has encountered this in its automatic update system

- Users swatted away update dialogs without even knowing that they'd done it
- Many Windows systems are so riddled with adware and popups that this would be a natural action for users

Windows XP SP2 changed the update process to nagware to get around this

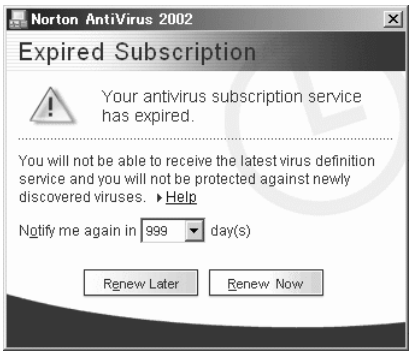

Conseq.of the Decision-making Process (ctd)

This occurs outside the computer world as well…

- British trains have a safety feature called the Automatic Warning System (AWS) in which the engine driver has to press a button within 3s of passing a danger signal
- If they fail to do so, the brakes are automatically applied

The system had design flaws that habituated drivers into cancelling unnecessary warnings

- Technical term is Signal Passed At Danger or SPAD
- In 1989 a driver went through two successive signals in this manner and killed five people
- If your security system has a defect significant enough to have its own acronym then it's probably a sign that you need to fix it

Confirmation Bias

Humans are bad at generating testable hypotheses

- Phenomenon is called confirmation bias
- Try and prove, rather than disprove, a theory

Humans will look for (or cook) the facts in order to support the conclusions that they want to reach

- Dissonance-motivated selectivity, look for material that avoids cognitive dissonance (challenging your opinions)

Confirmation Bias (ctd)

How do you check whether a web site is valid?

- Enter your name and password
- If the site accepts the password, it's valid

(If the security geeks had actually designed the mechanisms properly, this would be a valid site test)

Confirmation Bias (ctd)

US Navy addressed the problem of confirmation bias in tactical decision making after the Vincennes shootdown of an Iranian airliner in 1988

Introduced the STEP cycle for decision-making

- Create a Story (hypothesis)
- Test the hypothesis
- Evaluate the results

Makes creating a testable hypothesis an explicit part of the decision-making cycle

- Unfortunately a person in front of a security dialog hasn't had US Navy training and constant drilling to assist them

Other Biases

Disconfirmation bias

- People are more likely to accept an invalid but plausible conclusion than a valid but implausible one
- "This site looks and acts like my bank site, even if it's in eastern Europe. The browser must have got the URL wrong or something"

Blind-spot bias

- We can't see our own cognitive biases

You just can't win!

Other Biases (ctd)

CIA has published a special manual on dealing with biases

- Agency was particularly concerned with projection bias, a.k.a. "everyone thinks like us" bias
- "Psychology of Intelligence Analysis", available online from `www.cia.gov`

Well worth reading, since some of the techniques are also useful for performing security analyses

- The other side has authenticated themselves, from now on we can trust anything that they send us

Other Biases (ctd)

Has hit numerous SSH implementations (client and server)

- Only check data validity before the user-auth phase
- The peer would never dream of authenticating itself and only *then* sending a malformed packet

Other Biases (ctd)

Widespread in other implementations as well

- Unix access control: Only check security on the first access
- Signed ActiveX controls: It's signed, it's gotta be OK
  - Signed anything: All it means is that someone paid a CA for a magic token to turn off the warning dialogs
- Confused deputy problem
  - Solaris automountd, …
- Internet kiosks
  - "Are you sure you want to install ikat.xpi?"
- Firewalls and the firewall mentality
- …

Rationality

Once users adopt a particular belief, they're remarkably reluctant to change it

- "Einstellung", from Gestalt psychology

Rationality (ctd)

Example: Analysing suicide notes as a "problem-solving exercise"

- Users were evaluated based on their performance in distinguishing fake and genuine suicide notes
- Were given feedback that they'd done well or poorly
- Finally, they were told that the ratings that they'd been given were completely random (with supporting paperwork)

Users continued to rate themselves based on this completely fictitious, randomly-chosen information

Rationality (ctd)

Receiving a particular type of feedback creates a search for further confirmatory evidence to support it

The Barnum effect, “we’ve got something for everyone”

- More formally the subjective validation effect

Rationality (ctd)

Also known as the Forer effect

- Bertram Forer gave students personality analyses assembled from horoscopes and asked them to rate their applicability on a 5-point Likert scale
- Average rating given by the students was 4.26
  - Very high rating when you consider central tendency bias
- They had all been given the same generic “personality analysis”

This (combined with cold reading) is the bread-and-butter of generations of psychics, tarot readers, and crystal-ball gazers

Rationality (ctd)

Example: Professional palm-reader was given a high accuracy rating by his customers

- Researchers asked him to tell them the opposite of what his readings showed for a one-week trial
- Customers rated him equally well for accuracy

Example: Subjects were given a political statement to read

- Those told that it was by Thomas Jefferson thought it advocated political debate
- Those told that it was by Lenin thought it advocated violent revolution

Rationality (ctd)

There are many, many results like this

- Experimental psychologists like doing this with people :-)

Real-world example: Try reading a description of some new medicine or therapy

- How many of the symptoms of whatever it's meant to cure do you suffer from?

Security and Rationality

Our brains evolved for survival and reproduction, not to automatically seek the truth

- Quick and dirty techniques serve evolution better than purely rational ones
- It's better to be wrong than …

We can rationalise away almost anything

Security and Rationality (ctd)

Example: Subjects were given a canned biography on someone along with a random snippet of information like "He joined the Navy" or "He committed suicide"

- In every case they could explain the snippet via some item in the short bio
  - Sometimes the same item was used to explain away diametrically opposite facts
- When subjects were told that the information was fictitious, they still maintained their beliefs
  - c.f. earlier suicide-note analysis experiment

Security and Rationality (ctd)

Example: Researchers created “inexplicable” situations by giving subjects two sentences covering totally unrelated events

> Kenneth made his way to a shop that sold TV sets. Celia had recently had her ears pierced

- Subjects had ten seconds to come up with an explanation
- 71% of them could

Sentences were changed to contain a common referent

> Celia made her way to a shop that sold TV sets. She had recently had her ears pierced

- 86% of subjects were able to come up with an explanation

Security and Rationality (ctd)

People will concoct plausible explanations for something and continue to believe it even if they’re shown that the evidence for their conclusion is wrong

- This plays straight into the hands of con artists and phishers

Security and Rationality (ctd)

Example: Humans going to a phishing site (part of a phishing study)

- www.ssl-yahoo.com must be a “subdirectory” of Yahoo!
- sign.travelocity.com.zaga-zaga.us is probably an outsourcing site for travelocity.com
- The company running the site probably had to register a different name from its brand because the name was already in use by someone else
- Other sites use IP addresses instead of domain names so this IP-address-only site must be OK
- Sites use redirection to a different site so this one must be OK
- …

Reasoning this way is normal human behaviour!

Security and Rationality (ctd)

Extreme example: Patients whose brain hemispheres have been separated in order to treat severe epileptic attacks

- Split-brain/corpus callosotomy
- Left brain was able to rationalise away what the right brain was doing even though it literally had no idea why it was doing it

An example of a phenomenon called illusory correlation

- People see connections where there aren’t any

Security and Rationality (ctd)

Example: Subjects were shown drawings of humans supposedly done by people who had been matched to random psychiatric disorders

- Reported various signs in the drawings that were indicative of the disorders
- (Like the Rorschach test, the Draw-a-Person diagnosis method used to be common in psychiatry until experimental psychologists showed that the "results" obtained were meaningless)
  - They show something about the person who set the test, but not the subject

(Experimental psychologists *really* like messing with people :-)

Security and Rationality (ctd)

Example: Human vision

- Humans have a blind spot at the back of the retina where the optic nerve attaches to the eyeball
  - No visual nerves present to register an image
- The mind fills in the blanks based on data from the surrounding area to patch over this bug

Octopi have a properly-designed eyeball

- Photoreceptors are located in the inner portion of the eye, optic nerves are located in the outer portion of the retina
- This is great for annoying intelligent design advocates: The designer of humans made a mistake
  - Junk DNA is actually a bunch of FIXME / TODO / BUG comments

Security and Rationality (ctd)

Example: Confabulation across saccades

- Our eyes are constantly making small jerky movements called saccades
- The mind smoothes out our vision during these movements
  - We're actually blind during saccades, so changes in a scene don't register because they're smoothed over by the mind
  - Frozen second-hand phenomenon on a watch (chronostasis)

Example: Filling in obliterated words in a conversation

- Use a cough to mask words in a recording
- Different listeners "hear" different words that are filled in by the brain (phonemic restoration)

Security and Rationality (ctd)

Self-deception isn't a bug but a psychological defence mechanism

Depressed people have a *better* grasp of reality than non-depressed people, not the other way around

- Phenomenon is called depressive realism

Depressives suffer from a deficit in self-deception

- Depressed people's decision-making is closer to data-driven SEU
- Non-depressed people follow a more flexible heuristic approach
- Serotonin deficiency gives a narrow focus of attention and discourages potentially risky heuristics

Security and Rationality (ctd)

What about removing emotions from decision-making?

- People occasionally incur damage to the amygdala, a part of the primitive lymbic system involved in the processing of emotions
- This should lead to completely rational, logical decision-making

In practice people with this type of brain damage are very poor decision-makers

- They don't know what they care about any more

Security and Rationality (ctd)

High levels of self-deception are strongly correlated with conventional notions of good mental health

- If the self-deception is removed, various mental disorders may emerge

No matter which "fixes" (accidental or deliberate) you try and apply to improve things, they invariably make things worse rather than better

The “Simon Says” Problem

Users are required to change their behaviour in the *absence* of a stimulus

Problem is well-known to social psychologists

- Experts in some cases will notice the absence of a particular cue
- Novices don’t know what’s supposed to happen and so won’t notice when it doesn’t happen

The “Simon Says” Problem (ctd)

Example: Subjects are shown sets of trigrams with a special feature

- After (on average) 34 sets of trigrams, they figured out that the special feature was the presence of the letter ‘T’
- No-one was able to detect the absence of the letter ‘T’, no matter how many trigrams they saw
  - Users were totally unable to detect the absence of a stimulus

This is exactly what browser UI designers expect us to be able to do!

- We have to detect the *absence* of a stimulus like a padlock

The "Simon Says" Problem (ctd)

People find negative information far more difficult to process than positive information

- Educational psychologists advise educators to present information as positively-worded truths, not negatively-worded non-facts

The "Simon Says" Problem (ctd)

Example: Propositional calculus problems used by psychologists

If today is not Wednesday, then it is not a public holiday.

Today is not a public holiday.

- Is today not Wednesday?
  - People find these problems far harder to evaluate than positive-information ones

Example: Browser security indicators

If the padlock is not showing then the security is not present.

- You couldn't make this any worse if you deliberately designed it this way!

Inattentional Blindness

People don't register objects unless they're consciously paying attention to them

- Humans have a deficit of something called "attention"
- Whatever this is, we don't have enough of it to go round

Attention is tied to change and the motion signals that accompany it

- Think of a predator creeping up on its prey
- Lack of motion signals → inability to spot change

fMRI has shown that we really are totally blind to the stimulus

Inattentional Blindness (ctd)

Best-known example is "Gorillas in our Midst"

- Subjects were asked to watch a basketball game with players dressed in black and white
- Told to count the number of times that each team bounced the ball
- In the middle of the game, a person in a gorilla suite pranced across the court
- Only 54% of users noticed
  - This amazing demonstration is often shown in pop-psychology programs on TV

Inattentional Blindness (ctd)

Commonly encountered on the road

- Drivers are looking for cars (and in some cases pedestrians), but not non-cars
  - Technical term from aircraft control is “tunneling of attention”
- Bike riders are non-cars and therefore practically invisible to motorists

It’s possible to change your bike’s profile from “not-a-car” to “car” by mounting two driving lights far apart on a frame

- (Then your bike looks really ugly)

Inattentional Blindness and Security

The padlock and other security indicators are a perfect match for inattentional blindness

- Researchers have found up to 100% failure rates for these indicators

IE6 SP2 added a security bar to warn users of security issues

- One usability test found that not one user had noticed its presence

Even the more obvious security indicators like browser security toolbars fall victim to this

- One study found that 39% of users were fooled by phishing sites across the entire range of toolbars

Inattentional Blindness and Security (ctd)

Windows Vista added UAC dialogs to warn users of (potential) security issues

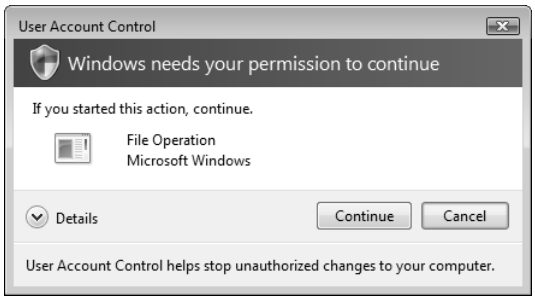

- Informal tests revealed that no-one had noticed that it had different colours in different situations
- Now try and find out what the colours actually signify…

Geeks vs. Humans

The geeks who build computer software don't think like normal humans

Example: Brand recognition

- Consumer-electronics store has two DVD players, a Philips and a Kamakuza
- Normal humans look at the Philips brand and buy it
  - Simple heuristics (RPD)
- Geeks will see that the Kamakuza player supports DivX, XviD, WMA, and Ogg Vorbis, has an external USB input and SD card slot for playing alternative media, and buy it
  - Economic/Bayesian model

Geeks vs. Humans (ctd)

For both sides this is a perfectly natural, sensible way to make a decision

Source: XKCD

- Both have come to completely opposite conclusions

Geeks vs. Humans (ctd)

Geek vs. human MBTI traits

- MBTI is a widely-used psychometric for personality traits
  - (Classifies personalities by Jungian personality types, not necessarily a personality test)

Geeks tend to be *TJ types

- Computer security people have a preponderance of INTJ's
  - (Geek → *TJ doesn't mean *TJ → Geek)

Why does this make geeks weird?

- Only 7% of the population has the *TJ profile

93% of users that geeks build software for think entirely differently from them

Geeks vs. Humans (ctd)

Final example of the difference between geeks and normal humans

All of Anne's children are blond

Does it follow that some of Anne's children are blond?

- (For logicians, assume that Anne has a nonzero number of children)

Geeks vs. Humans (ctd)

Most geeks would agree that the inference (a subalternation in Aristotlean logic) from "all A are B" to "some A are B" is valid

70% of normal humans consider this invalid

- This result is consistent across different cultures and re-phrasings of the problem (in the jargon, it is robust)

The people creating the security software just don't think like the majority of the people using it

Conclusion

Humans' minds work very differently from geeks' minds

- Many applications are written by geeks for geeks
- (Even supposedly user-friendly ones)

The mind works in very counterintuitive ways

- There are good reasons for the behaviour, but they're not at all obvious

Geeks are weird

- (No, really)