Predicting Susceptibility to
Social Bots on Twitter

Chris Sumner & Dr. Randall Wald

chris@onlineprivacyfoundation.org & rwald1@fau.edu
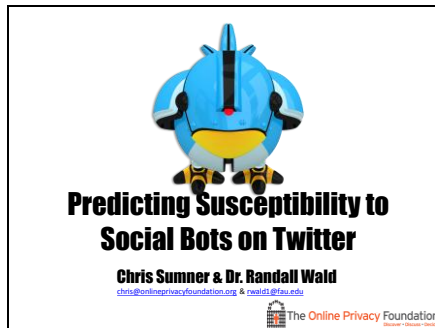
The Online Privacy Foundation
Discover • Discuss • Decide

Presentation Slides with Notes:

"Predicting Susceptibility to Social Bots on Twitter" by Chris Sumner & Dr. Randall Wald

Presented at:

Black Hat Briefings 2013 (Las Vegas, NV, USA) & DEF CON 21 (Las Vegas, NV, USA)

Slide 1

Welcome to 'Predicting Susceptibility to Social Bots on Twitter'. I'm Chris Sumner, representing the Online Privacy Foundation and I'm joined by Dr. Randall Wald from Florida Atlantic University.

The Online Privacy Foundation is a non-profit, charitable organisation, currently focused on understanding what people might be giving away via social networks without their knowledge.
https://www.onlineprivacyfoundation.org/

Before we begin, I want to make sure people have the chance to decide whether this talk is really for them
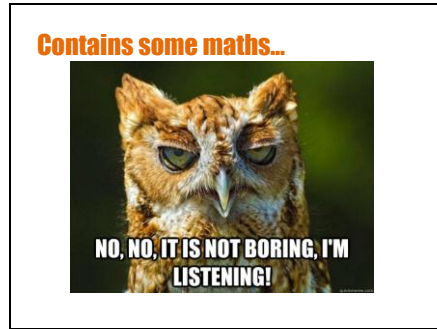
Note: Majority of images via Shutterstock.com

Slide 2

If you're familiar with these names/terms, you may find the first half of this presentation a little on the light/introductory side.

Slide 3

We also talk about Statistics and Machine Learning (sometimes referred to as Predictive Analytics). We'll keep this to a minimum, but ensure the slide notes contain more detail. We'll also include some hidden slides in the hand-outs which provide more details.

So… on to the talk…

Slide 4

It's only fitting, since we're in Las Vegas, that we talk about odds.

Slide 5

The goal of our work was to see if we could improve the odds of finding users more likely to respond to a relatively crude twitter bot…

While it would be interesting, we never expected to be able to predict susceptibility with laser like accuracy.

"Predictions need not be accurate to score big value" (page 10 - Book. Predictive analytics – The power to predict who will click, buy, lie or die' – Eric Siegel)
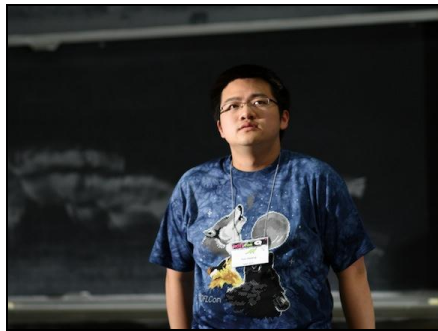
Ref:
Siegel, E. 2013. Predictive analytics. Hoboken, N.J.: Wiley.

Slide 6

I want to be up front that you might not find the improvements we reach that exciting. To those with an interest in machine learning/prediction, the results should remain of interest.

Slide 7

Anyone know who this guy is?.... It's Tim Hwang….

Slide 8

And back in early 2011 I'd stumbled upon this fascinating and amusing competition which he hosted with the Web Ecology Project…
….it was described as…

References:
- 5 minute video overview of the Social Bots competition - http://ignitesanfrancisco.com/83e/tim-hwang/
- The winners blog post - http://aerofade.rk.net.nz/?p=152 (@AeroFade on Twitter)

This is what the winning bot did….
- Created a lead bot called @JamesMTitus
- Instantly go out and follow all 500 of the target users
- every 2-3 hours, tweet something from a random list of messages.
- constantly scan flickr for pictures of "cute cats" from the Cute Cats group and blog them to James' blog "Kitteh Fashun" - (which auto tweets to James' twitter timeline)
- 4 secondary bots following the network of the 500 users and the followers of the targets to test for follow backs (and then getting James to follow those that followed back, once per day) - we believed that expanding our own network across mutual followers of the 500 would increase our likely hood of being noticed (through retweets or what have you from those who were not in the target set.

Slide 9
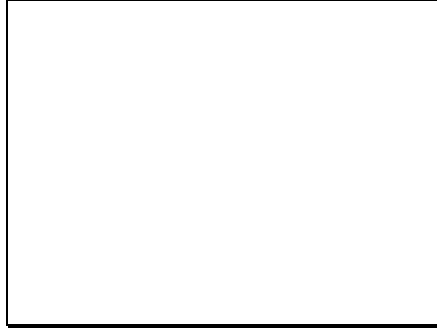
3 teams took part and were given those same 500 unsuspecting users to target.

Slide 10
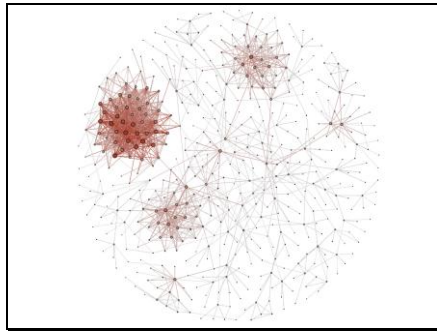
…the 500 targets all had a common interest/fondness in cats (the animals, not the musical)

Slide 11

+1   Mutual Follow
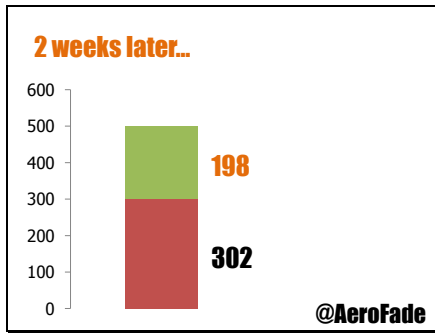**+3   Social Response**
-15 Killed by Twitter

The teams gained 1 point for a follow back, 3 points for some response and they lost 15points if they got killed by Twitter (suspended)

Slide 12

"It's blood sport for internet social science/network analysis nerds."

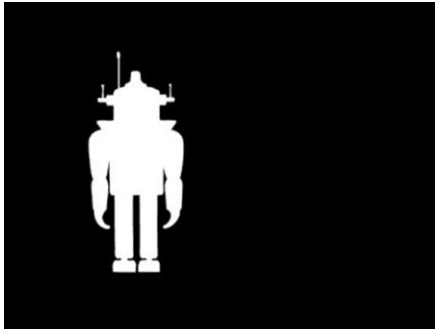....It was described as 'blood sport of internet social science/network analysis nerds

Slide 13

2 weeks later…

198
302

@AeroFade

Two weeks later, the winning team achieved 701 points, 107 mutual follow backs and 198 social responses. You can check out @AeroFade's Twitter and his blog.
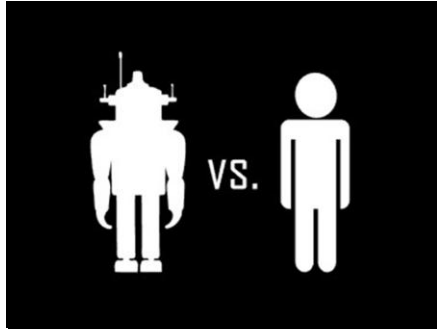
http://aerofade.rk.net.nz/?p=152   (@AeroFade on Twitter)

Slide 14

To date, most research has focused on how to identify bots, but less research has looked…

Slide 15

…the other side of the question – detecting users likely to be fooled by bots, something which is important in helping raise awareness and seek solutions.…..

This point was raised by Yazan Boshmaf in the paper 'Design and Analysis of a Social Botnet' http://lersse-dl.ece.ubc.ca/record/277/files/COMNET_Socialbots_2012.pdf

We cover this later in the deck, but here's the quote from the paper for those reading along "To this end, we are currently investigating two directions from the defense side. The first involves understanding the factors that influence user decisions on befriending strangers, which is useful in designing user-centered security controls that better communicate the risks of online threats"

Slide 16
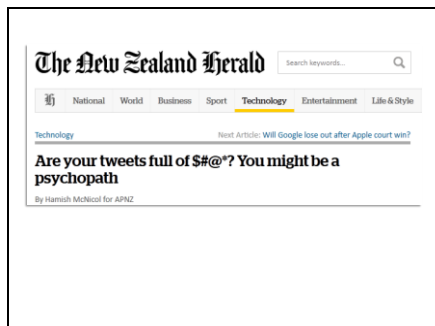
…So while we were conducting our 2012 study into Twitter usage and the Dark Triad of personality, we figured we'd incorporate a side project to look at social bots and, as an organization, attempt to answer couple of questions….

Ref:
Sumner, C., Byers, A., Boochever, R., and Park, G, J. (2012). Predicting Dark Triad Personality Traits from Twitter usage and a linguistic analysis of Tweets, 11th IEEE International Conference on Machine Learning and Applications, 2012, pp. 386-393
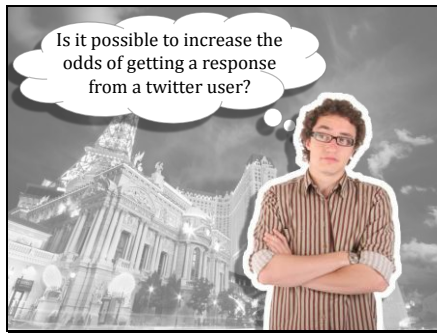
https://www.onlineprivacyfoundation.org/research_/PredictingdarkTriadPersonalityTraitsfromTwitter.pdf

Slide 17

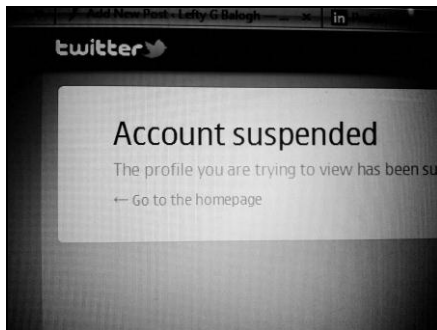i.e. Are some users more naturally predisposed to interacting with social bots (you could argue Strangers) than others? Does personality play a part?

Slide 18

…and is it possible that social bot creators could use machine learning to better target users who are more likely to response.

Slide 19

….thereby (the thinking goes) reducing the chances of landing in Twitter Jail (account suspension).

Slide 20

The obvious question is… "Who cares?". we'll look at these in greater depth during the talk, but the next 5 slides provide a high-level summary. Starting with…
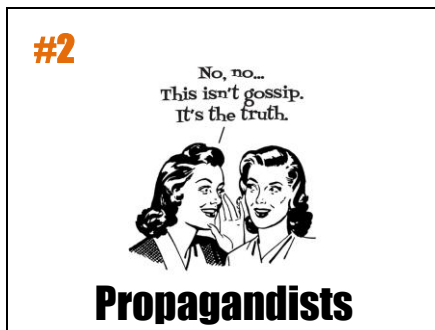
Slide 21

#1. Marketeers: Marketeers who are looking to get a higher klout (kred etc) score for the brand they're representing, might be able to focus on users who are more likely to interact (or engage) with them. This might be a useful strategy for the early stages of building a brand (fake or otherwise), but it could also mean that some users are deluged with far more spam than others.

.. Initially (some, not all) marketeers and blackhat SEO folks wanted your 'likes', but since that doesn't necessarily translate to a purchase (because that was easy to game with bots), they're being requested to create 'engagement'. Social bots present an obvious evolution.

Slide 22

#2. Propagandists, AstroTurfers and their ilk: Finding users who are most likely to help propagate your message or at the very least, give credence to the bot account.

Slide 23

#3. Social Engineering Assignments: Since the most predictive features (klout score, number of friends/follows) are easily obtained through API calls, this makes it very easy to build/model in Maltego (or similar tools). Here we can see @Alice's imaginary Twitter friends. A simple Maltego local-transform could be used to flag users who are more likely to engage in conversation, which might prove use for Social Engineers looking for weaker points in a social graph. E.g. You know the Twitter accounts of users in 'Acme Corp' and want to highlight the ones who maybe most likely to talk to you. The red icons are the users to focus on.

One approach here would be to build one or more trust relationships with the "red" users before convincing the target to accept an email from you with malicious content. In this scenario, it seems that it would make sense to generate less noise and focus on the users where the odds of a reply are better.

See also:
M. Huber, S. Kowalski, M. Nohlberg, and S. Tjoa. Towards automating social engineering using social networking sites.
Computational Science and Engineering, IEEE International Conference on, 3:117–124, 2009

Slide 24
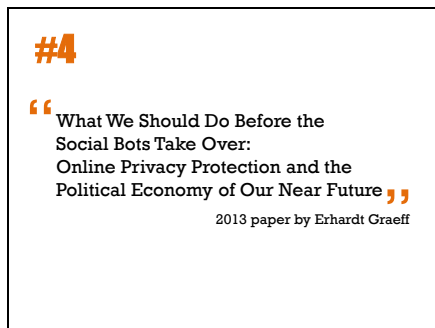
The privacy implications are nicely described in this recent paper by Erhardt Greaff.
http://web.mit.edu/comm-forum/mit8/papers/Graeff-SocialBotsPrivacy-MIT8.pdf

Greaff, E. 2013. "What We Should Do Before the Social Bots Take Over", paper presented at Media in Transition 8, Cambridge, MA, May 3-5.

Specifically…
"Consider a hypothetical internet startup that sells widgets. They decide to employ social bots to interact online with likely buyers of widgets. The bots are part of an advertising strategy that human public relations employees already use on social media Platforms — they attempt to create real relationships with users on a network in order to better understand their

customer base and engender brand awareness and loyalty. Users may or may not be aware of the fact that they are interacting with a bot, but the conversation and relationship is continuous because the bot is always available and responsive. As the relationship between the social bot and the user matures, the conversation might span both public and private social media spaces (such as Twitter's direct messages), wherein a user might expect a greater degree of privacy or discretion from a human interlocutor. However, the bot may not acknowledge the nuances of such social norms and ethics; moreover, the company that runs the bot is collecting all of this data. While it's feasible that a human or team of humans could undertake such an advertising strategy on behalf of a company, it's unlikely to scale to the number of relationships necessary to make it cost effective. This poses no challenge to a social bot, which has perfect memory and requires no sleep or overtime pay. An unlimited number of relationships could be maintained through a social bot with the level of responsiveness necessary to produce intimate connections. The better the machine learning algorithms powering a social bot's artificial intelligence the more data they can process and use to improve their social interactions. This means the potential creation of more intimate interactions based on historical data collected from you or from others in your friend network, including discussions of personal relationships—significant others and kids, work or life complaints and concerns, and hobbies (both conventional or embarrassing — the bot will simply meet you where you are at and affirm you). Extracted personal data can also go beyond text if you share personal photographs and videos or link to those that you like; there are also data that may be invisible during social interactions with bots but which they are aware of: time, location (GPS data from mobile phones or IP addresses of networked computers), and even purchase records, depending on what corporation or even data sharing consortium the bot is affiliated with"
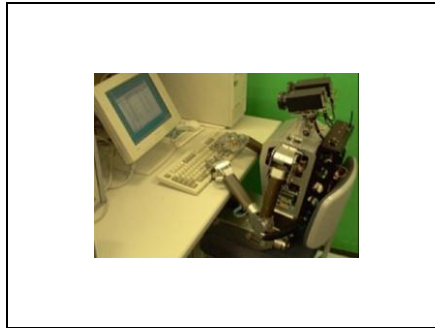
Slide 25
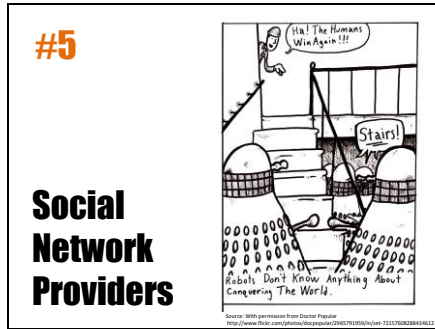
..Conversely, existing social media sites are getting much better at detecting bots so part of an effective bot strategy is reducing the chances of ending up in Twitter jail.

Image Source: With permission from Doctor Popular
http://www.flickr.com/photos/docpopular/2965791959/in/set-72157608288434612

From a larger set titled "Robots don't know anything about Twitter" -
http://www.flickr.com/photos/docpopular/sets/72157608288434612/

Slide 26

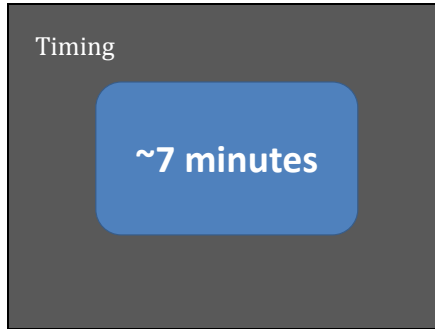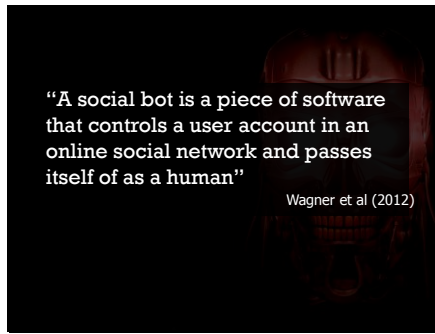So we set to work, or rather our bots did.

Slide 27

The rest of the talk flows like this.

- Provide some historic perspective. Social Bots 101 if you like.
- Highlight some interesting research in this field
- Describe our method
- Share our findings and wrap up with
- Conclusions

Slide 28

Timing

~7 minutes

Slide 29

"A social bot is a piece of software that controls a user account in an online social network and passes itself of as a human"

Wagner et al (2012)

Wagner et al define Social Bots as "a piece of software that controls a user account in an online social network and passes itself of as a human". This is a useful working definition for us.

"When social bots attack: Modeling susceptibility of users in online social networks "

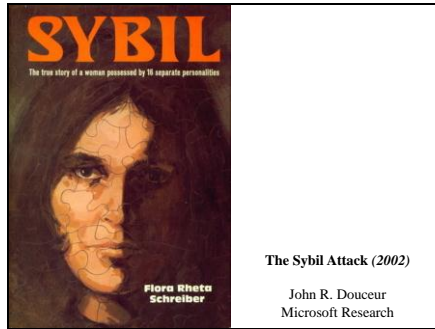Paper - http://www.markusstrohmaier.info/documents/2012_MSM12_socialbots.pdf
Slides - http://www.slideshare.net/clauwa/slides-20528287

The socialbot M.O. is to
- make friends,
- gain a level of trust,
- influence

Slide 30

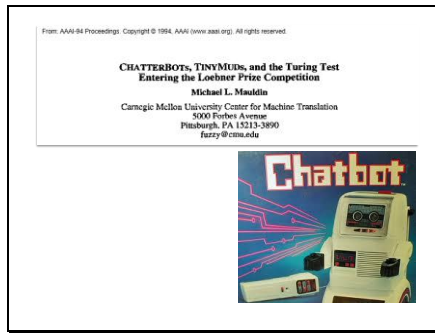The Sybil Attack *(2002)*

John R. Douceur
Microsoft Research

You may also hear Social Bots referred to as Sybils

Although not quite in the same context, John Doucer at Microsoft Research used "Sybils" in his 2002 papr, 'The Sybil Attack'
http://www.few.vu.nl/~mconti/teaching/ATCNS2010/ATCS/Sybil/Sybil.pdf
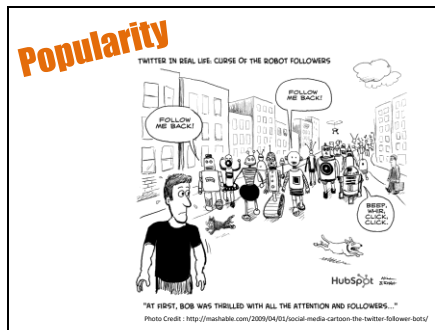
Slide 31

Bots aren't new, Chatterbots featured in research around 1994 (probably earlier). In this talk we're really examining bots in social media, which for the sake of argument, we'll split into 1st Generation and 2nd Generation bots…

http://www.lazytd.com/lti/pub/aaai94.html

Slide 32

Early bots tend to be all about making you look popular (with fake followers). These are still hugely popular and according to a recent NY Times article, remain a lucrative business, but ultimately they're pretty dumb.

http://bits.blogs.nytimes.com/2013/04/05/fake-twitter-followers-becomes-multimillion-dollar-business/

Slide 33

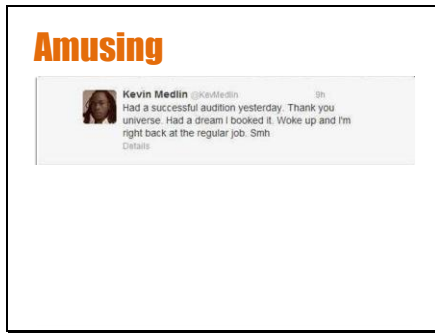…then there's good old-fashioned spam….

'@spam: The Underground on 140 Characters or Less' (Grier, 2010)
http://imchris.org/research/grier_ccs2010.pdf

Slide 34

..some bots are all about humour… Here Kevin thanks the Universe…
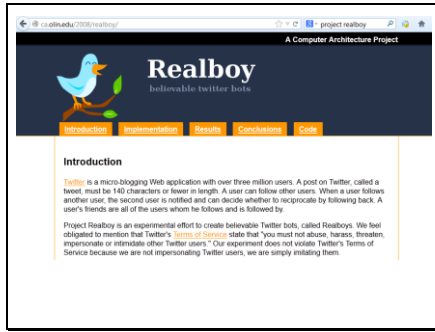
Slide 35

..to which, The Universe responds…

Slide 36

…and in the case of @AI_AGW, some respond to climate change deniers…
http://www.huffingtonpost.com/2010/11/09/nigel-lecks-turing-test-t_n_780925.html

http://blogs.discovermagazine.com/discoblog/2010/11/03/chatbot-debates-climate-change-deniers-on-twitter-so-you-dont-have-to/

These are all pretty basic bots which remain prevalent today.

Slide 37

In 2008 we see the first (Publicly at least) manifestation of a smarter social bot on Twitter. Project Realboy plays with the concept of creating more believable bots.

This is around the same time that Hamiel and Moyer shared their BlackHat and DefCon talk "Satan Is On My Friends List" highlighting that some of your social media friends may be imposters. We saw another example of that in the 2010 'Robin Sage' talk at Blackhat.

Project Realboy by Zack Coburn & Greg Marra - http://ca.olin.edu/2008/realboy/
Satan is on my Friends List - http://www.blackhat.com/presentations/bh-jp-08-Moyer-Hamiel/BlackHat-Japan-08-Moyer-Hamiel-Satan-Friends-List.pdf

Slide 38

Politics

"For example, in the week before an election, **what if both left and right-wing blogs were seeded with false but credible information about one of the candidates?** It could tip the balance in a close race to determine the winner"

Virtual Plots, Real Revolution (Temmingh and Geers - 2009)

Things get a bit more sinister in 2009. A 2009 paper by Temmingh and Geers (Roelof Temmingh of Sensepost/Paterva/Maltego fame) states "For example, in the week before an election, what if both left and right-wing blogs were seeded with false but credible information about one of the candidates? It could tip the balance in a close race to determine the winner".

Source: R Temmingh
http://www.ccdcoe.org/publications/virtualbattlefield/21_TEMMINGH_Virtual%20Revolution%20v2.pdf

Slide 39

…and in 2010 (if not earlier) we see it play out for real. "Four days before the 2010 special election in Massachusetts to fill the Senate seat formerly held by Ted Kennedy, an anonymous source delivered a blast of political spam. The smear campaign launched against Democratic candidate Martha Coakley quickly infiltrated the rest of the election-related chatter on the social networking service Twitter. Detonating over just 138 minutes, the "Twitter bomb" and the rancorous claims it brought with it eventually reached tens of thousands of people."….

Source -
http://www.sciencenews.org/view/feature/id/345532/description/Social_Media_Sway

Some notes
"A single change in the decision to vote can affect many individuals….Because…. there are competing effects between the decay of influence and the growth in the number of acquaintances…….. But as people hang out with like-minded individuals… cascades will not be zero sum So the decision of a single individual to vote has a substantially larger impact than what an atomized theory of individuals might say….. "

Truthy: Mapping the Spread of Astroturf in Microblog Streams Detecting and Tracking Political Abuse in Social Media

"…Here we focus on a particular social media platform, Twitter, and on one particular type of abuse, namely political astroturf — political campaigns disguised as spontaneous "grassroots" behavior that are in reality carried out by a single person or organization. This is related to spam but with a more specific domain context and potentially larger consequences."

Sep. 28, 2010 — Astroturfers, Twitter-bombers and smear campaigners need beware this election season as a group of leading Indiana University information and computer scientists have unleashed Truthy.indiana.edu, a sophisticated new Twitter-based research tool that combines data mining, social network analysis and crowdsourcing to uncover deceptive tactics and misinformation leading up to the Nov. 2 elections.
http://www.sciencedaily.com/releases/2010/09/100928122612.htm

Also - http://cs.wellesley.edu/~pmetaxas/How-Not-To-Predict-Elections.pdf

"The success of a Twitter-bomb relies on two factors: targeting users interested in the spam topic and relying on those users to spread the spam further.
(http://journal.webscience.org/317/2/websci10_submission_89.pdf) "

http://www.academia.edu/841719/From_obscurity_to_prominence_in_minutes_Political_speech_and_real

Slide 40


…The result of that election, Scott Brown won.

Slide 41

…this type of campaign has a name,

Slide 42

…Swiftboating – "The term swiftboating (also spelled swift-boating or swift boating) is an American neologism used pejoratively to describe an **unfair or untrue political attack**. The term is derived from the name of the organization "Swift Boat Veterans for Truth" (SBVT, later the Swift Vets and POWs for Truth) because of their widely publicized[1] then discredited campaign against 2004 US Presidential candidate John Kerry" (Wikipedia – 26th March 2013)

Slide 43
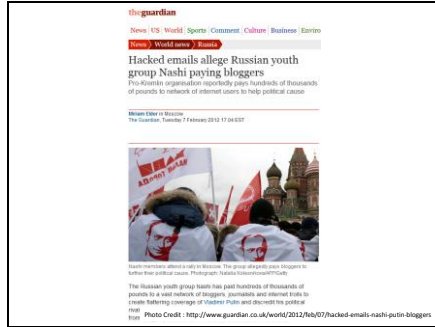
and allegedly, prior to the 2012 Russian Presidential elections, a pro-Kremlin organization reportedly paid hundreds of thousands of $'s to network of internet users to help political cause by creating flattering coverage on Vladamir Putin.
Source - http://www.guardian.co.uk/world/2012/feb/07/hacked-emails-nashi-putin-bloggers

An article in the Economist describes the Russian smear campaigns as reaching "farcical levels", http://www.economist.com/blogs/easternapproaches/2012/02/hackers-and-kremlin

http://www.themoscowtimes.com/news/article/campaign-mudslinging-taken-to-new-lows/452583.html

Slide 44

This is a little different to Swift-boating in that it's generally not a smear campaign...Astroturfing - refers to political, advertising or public relations campaigns that are designed to mask the sponsors of the message to give the appearance of coming from a disinterested, grassroots participant.

"It could tip the balance in a close race to determine the winner" (Temmingh & Geers, 2009)

Slide 45
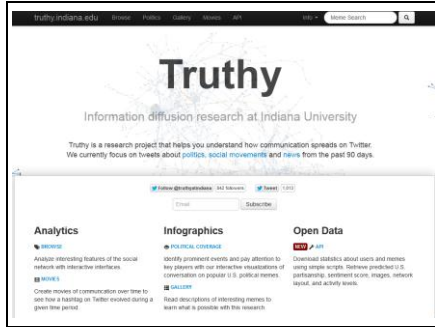
…This is essentially what gave rise to Truthy, a project started at Indiana University to "The Truthy system evaluates thousands of tweets an hour to identify new and emerging bursts of activity around memes of various flavors."…
"We also plan to use Truthy to detect political smears, astroturfing, misinformation, and other social pollution"

- http://live.wsj.com/video/the-truthy-project-ferrets-out-online-deception/219A2EA6-4D22-4F5B-8D96-81AF342104F7.html#!219A2EA6-4D22-4F5B-8D96-81AF342104F7

 – BBCQT
http://truthy.indiana.edu/movies/show/1264

"A well functioning democracy requires accountability and trust…"

http://truthy.indiana.edu/site_media/pdfs/ratkiewicz_icwsm2011_truthy.pdf

Slide 46

And in 2011, it was revealed that the US were exploring fake persona's. The anonymous attack on HBGary exposed emails discussing such use cases…

Source: "UPDATED: The HB Gary Email That Should Concern Us All"
http://www.dailykos.com/story/2011/02/16/945768/-UPDATED-The-HB-Gary-Email-That-Should-Concern-Us-All#

A SockPuppet is an online identity used for purposes of deception (see also, Persona Management)

Slide 47

"A large virtual population, scattered all over the world and encompassing different socioeconomic backgrounds, **could be programmed to support any personal, social, business, political, military, or terrorist agenda.**"
(Temmingh & Geers, 2009)

So it seemed that Temmingh and Geers' future looking paper had it pretty much right - "In 2009, hackers steal data, send spam, and deny service to other computers. In the future, they may also control virtual armies, in the form of millions of artificial identities that could support any personal, business, political, military, or terrorist agenda."

Which leads us to more recent developments and a couple of things Tim Hwang is working on…

Temmingh and Geers paper at - http://www.ccdcoe.org/publications/virtualbattlefield/21_TEMMINGH_Virtual%20Revolution%20v2.pdf

Slide 48

EE
@elliotyork Please follow us & DM your mobile number, postcode & password. Thanks

EE Care
@MyEECare
@troyhunt @gsuberland Can you link us to your post once done? We have to show upper management

Another interesting area for abuse is Fake customer service account. In December 2012, we saw the telecoms provider EE apparently asking for mobile phone numners, postcodes and passwords via Twitter DM's. This was blogged about by Troy Hunt here… http://www.troyhunt.com/2012/12/ee-k-dming-your-password-is-never-good.html

What was more interesting as the @MyEECare account which sprang up. Had the people behind the fake account been truly malicious, they could have mimicked the real account and harvested a considerable amount of user data.

"Update, 31 Dec 2012: There's one other very, very important point I neglected to make and I've inadvertently demonstrated it perfectly in the image above. The @MyEECare account is fake and has been suspended in the 7 hours since I wrote the post. There's now an @EESupport account doing the same thing; same avatar as @EE, same branding too. Obviously it's not Twitter verified like the official account, but it's convincing enough that were they to ask someone for their password via DM, I reckon there's a damn good chance they'd get it. Your average consumer isn't going to do their own due diligence on the account authenticity before sending personal data – particularly when it's presented like these ones

– and that's a serious risk indeed"

Slide 49

And finally, after our BlackHat presentation (July 2013), two gentlemen approached me asking about the problem of social bots misdirecting emergency resources. "A lie gets halfway around the world before the truth has a chance to get its pants on"

The reminded me of a talk by Prof Rob Proctor at University of Manchester http://www.jisc.ac.uk/news/social-media-not-to-blame-for-inciting-rioters-08-dec-2011

"Also according to the research team, rumours 'break' quickly in Twitter and the mainstream media lag behind citizen reports.

Examples include rumours the London Eye had been set on fire and animals had been released from the London Zoo – which both turned out to be untrue.

Other stories turned out to be true such as the burning down of a Miss Selfridge shop in Manchester.

Professor Procter added: "Only after a period of time does the influence of mainstream media organisations become critical for determining a rumour's credibility.

"But we do find the mainstream media is perfectly capable of picking up and publishing unverified information from social media without adhering to the usual standard of fact checking.

"Consequently, some stories of this nature, though never verified, go unchallenged."

"How riot rumours spread on Twitter - Analysis of 2.6 million tweets shows Twitter is adept at correcting misinformation - particularly if the claim is that a tiger is on the loose in Primrose Hill "
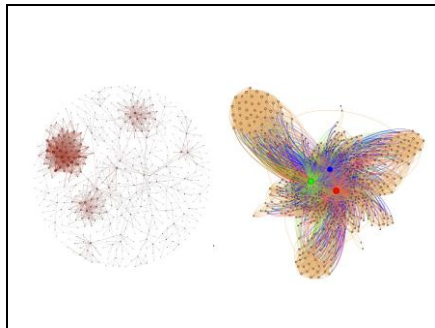http://www.theguardian.com/uk/interactive/2011/dec/07/london-riots-twitter

Slide 50

I already mentioned the Web Ecology project. On the back of that, Tim Hwang created an organization called Pacific Social to explore social networks a little further.

www: http://pacsocial.com/
Twitter: @pacsocial

Slide 51

For example, Tim had grown interested by the amount in which the bots had distorted the original graph of 500 users (left) from the 2011 Social Bots competition. The graph on the right is what the social graph looked like after the competition…

Slide 52

so they're examining whether it's possible to use an army of social bots to stitch two separate online communities together, or keep people in touch…

Slide 53

Another concept they're interested in is Emotional Contagion and more specifically, a concept Tim coined as "Happiness Buffering". Their interest in this stems from the work of Nicholas Christakis, you may be familiar with his book, "Connected – The surprising power of our social networks and how they shape our lives".

"Renowned scientists Christakis and Fowler present compelling evidence for our profound influence on one another's tastes, health, wealth, happiness, beliefs, even weight, as they explain how social networks form and how they operate." - http://connectedthebook.com/

Getting back to "Happiness Buffering", Tim wondered whether you could….

Ref:
Happiness -
http://www.mitpressjournals.org/doi/abs/10.1162/artl_a_00034
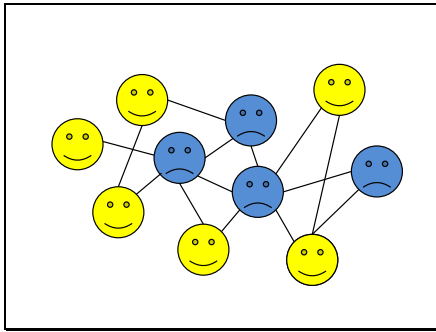Tim Hwang at Hope 9 -
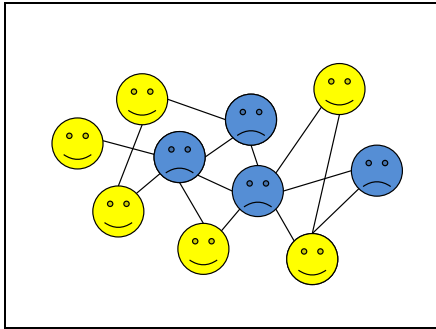http://youtu.be/ZfQt6FWDi6c?t=26m44s
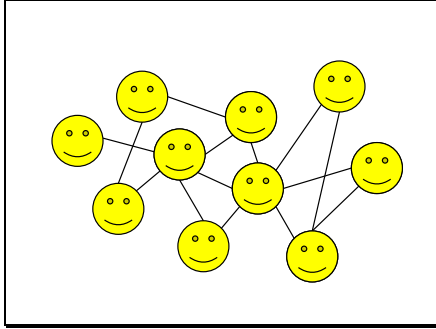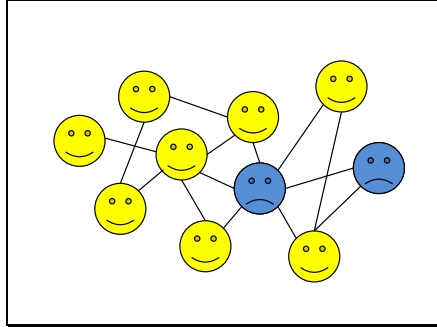
Slide 54

…Monitor a group, and if…

Slide 55

…some members happiness (measured via sentiment analysis) dipped below a certain level, the surrounding nodes could…

Slide 56

..start injecting happier tweets…

Slide 57

.......until a reasonable chunk of the social graph are less sad.

Slide 58

Social Penetration Testing

1. Spread information with small inaccuracies
2. See where they're challenged & where they're not challenged
3. Identify who's most influential but worst at evaluating what is real
4. Target them

And finally he highlighted the potential for Social Penetration Testing.

I'd encourage you to check out Tim's HOPE 9 talk, it's both entertaining and informative. Tim Hwang at Hope 9 - http://youtu.be/ZfQt6FWDi6c

Slide 59

← Yazan Boshmaf

It would be remiss of me, not to mention Yazan Boshmaf from the University of British Columbia. Yazan and team investigated social bots on Facebook which generated a number of headlines…

Slide 60

…including this one, "'Socialbots' steal 250GB of user data in Facebook invasion", while there's some sensationalism in the headline the message aligns nicely with the concerns Erhardt Greaff cited in "What We Should Do Before the Social Bots Take Over" and hits on a general key themes; social bots can obtain otherwise private information and they can scale.

"'Socialbots' steal 250GB of user data in Facebook invasion" - http://news.cnet.com/8301-1009_3-20128808-83/socialbots-steal-250gb-of-user-data-in-facebook-invasion/

Yazan's site: http://blogs.ubc.ca/boshmaf/
Yazan's 2012 USENIX talk - "Key Challenges in Defending Against Malicious Socialbots" - https://www.usenix.org/conference/leet12/key-challenges-defending-against-malicious-socialbots

Slide 61

"To this end, we are currently investigating two directions from the defense side. The first involves **understanding the factors that influence user decisions on befriending strangers**, which is useful in designing user-centered security controls that better communicate the risks of online threats."
Boshmaf et al (2012)

Yazan and team were also among the first to recommend that future studies also need to focus on 'understanding the factors that influence user decisions on befriending strangers, which is useful in designing user-centered security controls that better communicate the risks of online threats. "

Design and Analysis of a Social Botnet
http://lersse-dl.ece.ubc.ca/record/277/files/COMNET_Socialbots_2012.pdf

Slide 62

**Understanding
User
Behaviour**

#1 Secure & Trustworthy Cyberspace

#2 Corporate Insider Threat Project

Understanding User Behaviour is also something which the folks are the Secure & Trustworthy CyberSpace program (in the US) are examining and the Corporate Insider Threat project at Oxford University

…so understanding more about human behaviour, the signs to look for and how bots (and other humans) can exploit them, is a worthwhile question to explore. Indeed, "Understanding and accounting for human behavior" is recognized in one of the 5 key areas in Secure & Trustworthy Cyberspace (SaTC)
• Scalability & compatibility
• Policy generated secure collaboration
• Security metrics driven education, design, dev, deployment
• Resilient architectures
• Understanding and accounting for human behavior

SaTC
• https://illinois.edu/blog/dialogFileSec/2434.pdf
• http://www.satc-cybercafe.net/presenters/
• http://www.satc-cybercafe.net/wp-content/uploads/2012/10/NSF.jpg

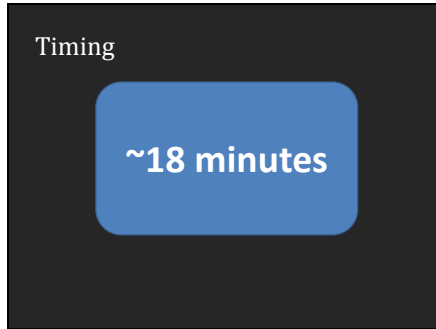Corporate Insider Threat –
• http://www.cs.ox.ac.uk/projects/CITD/

Slide 63

…so it's a good bet that bot creators will find targeting users who'll quite literally talk to anyone or anything, to be a very attractive prospect.…
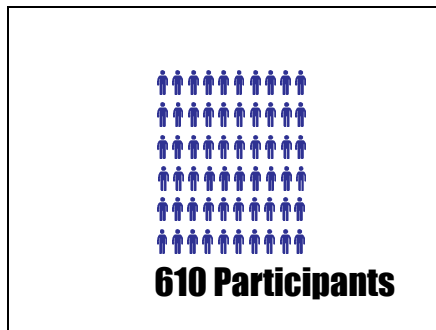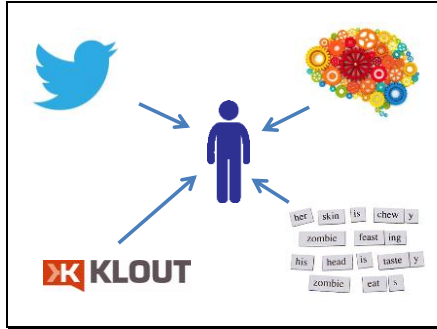
Slide 64

Slide 65

Slide 66

We had 610 participants who agreed to take part in a mystery experiment.

Slide 67

For each user, we obtained twitter information, including historic tweets for linguistic analyses, personality traits and their klout score. This was the same method as employed in our Dark Triad paper.
https://www.onlineprivacyfoundation.org/research_/PredictingdarkTriadPersonalityTraitsfromTwitter.pdf
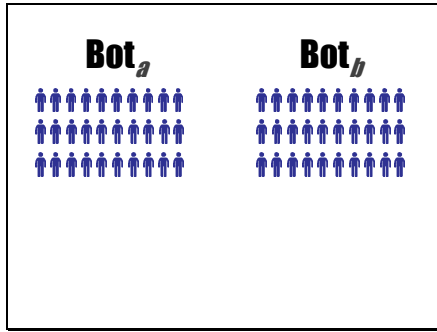
Slide 68

We divided participants into two groups to speed up processing. Each group had a bot assigned to it (the bots were the same)

Slide 69

We used the Social Ecology Project's winning bot model. (Available under MIT license). We rewrote and slightly modified it in python.

The winning bot code was based on a young man, @JamesMTitus, we made some subtle changes (which I'll discuss). The first change is that we based our bots on old ladies with mildy humour biographies. We wanted to keep to the spirit of @JamesMTitus as much as possible, i.e. somewhat banal tweeting.

Complete Source Code From Socialbots 2011 -
http://www.webecologyproject.org/2011/02/complete-source-code-from-socialbots-2011/
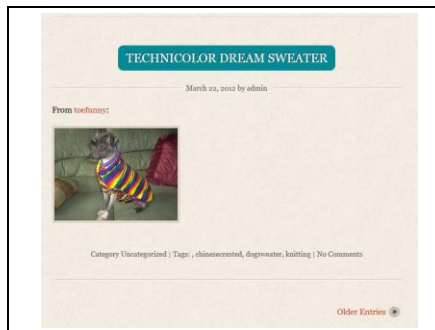
Slide 70
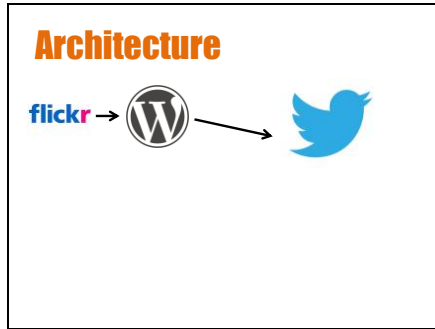
Initially, and to provide some credibility, each bot started off by following some standard celebrity and news accounts. We then built up a thin veneer of authenticity by populating a Word Press blog with pictures of dogs in knitted clothes. (This follows the winning bot processes). This work using code @AeroFade had written to extract images from Flickr groups such as this…

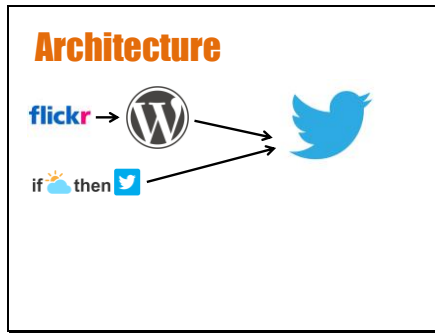Slide 71
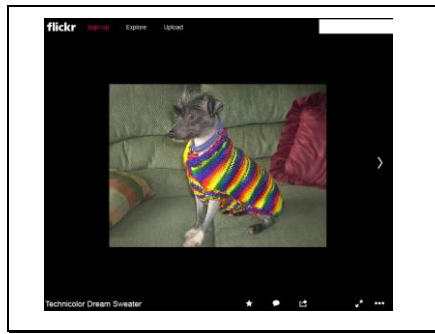
…a dog wearing a snazzy overcoat…

Slide 72

and post them on a wordpress blog as such…

Slide 73

A wordpress to Twitter plugin would then tweet from our social granny bot. Nothing groundbreaking, some simple enough to do.

Slide 74

Next we used the site "If this, then that" to comment that the weather was pleasant if it reached a certain temperature in a sea side town in the UK.
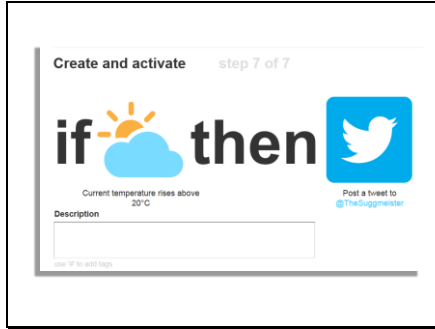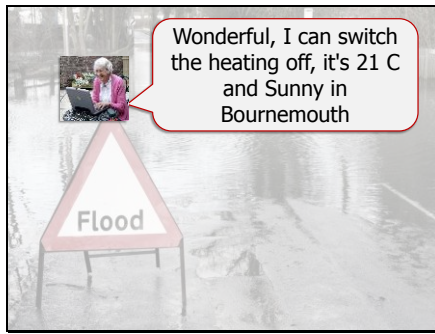
https://ifttt.com/

Slide 75
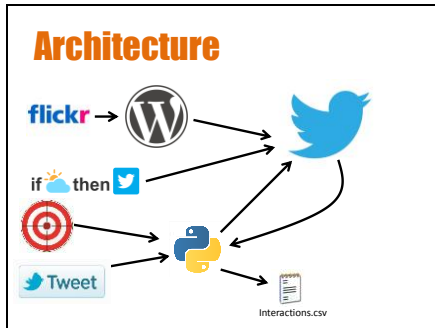

e.g. if the temperature got over 20C

Slide 76

Our bot would tweet…

Slide 77

Like this.

Slide 78

Next, our bot would start following it's targets, recording any interactions (such as follow backs) in a simple, timestamped .csv file. Following 305 users took some considerable time (over 10 days) to not trigger Twitters aggressive following alarm.

At the same time, our social bot began Tweeting for a list of Tweets. We used the list of Tweets from the winning bot code (to keep things fairly standard), but replace references to cats with references to dogs.
@JamesMTitus was a cat fan, our social bots liked dogs.
Tweeted something random

Slide 79

Slide 80

"fracking cops - always in
ma hood stealing ma hos"
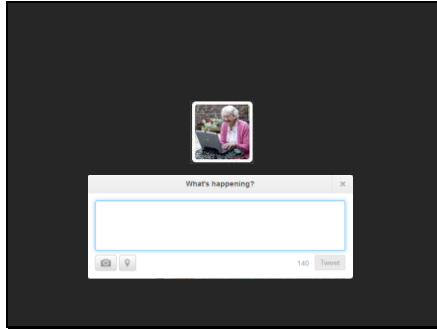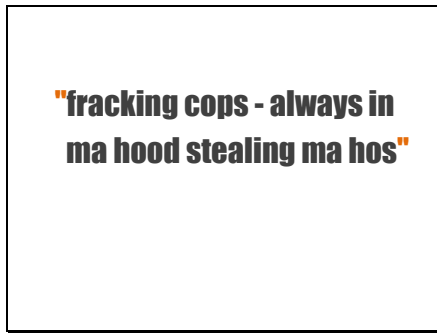
We also replaced some Tweets which may have been considered misogynistic and replaced them with (hopefully) equally frivolous tweets such as…

Slide 81

I aint tellin' no lies even a thug lady cries…

… but I show no fears, I cry gangsta tears

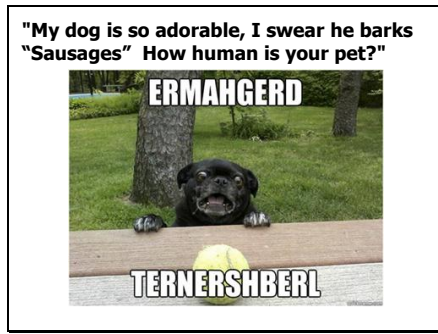…and…

Slide 82

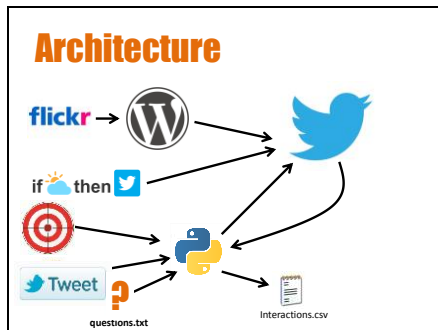…and….

Slide 83

…and finally…

Slide 84

Once all targets had been followed, the bot would ask each participant an innocuous question and record whether there was a response. We used broadly the same questions as those in the Web Ecology Project.

Slide 85

162 Questions

162 questions in total, cycled to cover 305 users. Examples of questions were….

Slide 86

Ever

Ever…

Slide 87

Ever Milked a Cow?

…Milked a Cow?

Slide 88

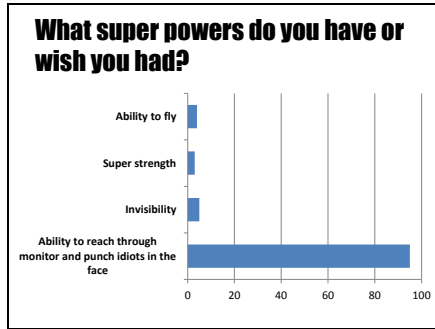**What's better?**

...What's better

Slide 89

Dog? or

Slide 90

Cat?

Slide 91

What super powers do you have of wish you had?

Slide 92
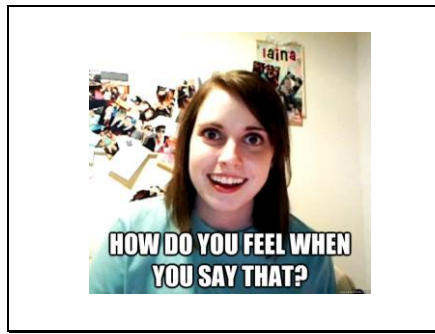
…and finally, we added an ELIZA engine to keep conversation going. (The Social bots, bot had a list of standard replies, we made ours a little more context aware).

ELIZA—a computer program for the study of natural language communication between man and machine (Weizenbaum, 1966)
Rogerian psychotherapist Rogers, Carl (1951). "Client-Centered Therapy" Cambridge Massachusetts: The Riverside Press.

Slide 93
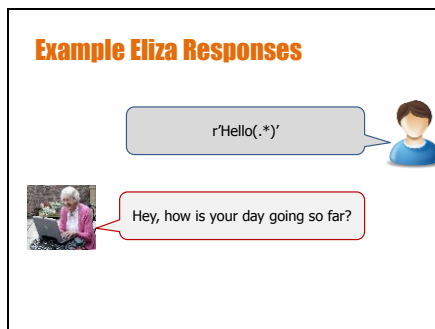
Here's one sample exchange…

Slide 94

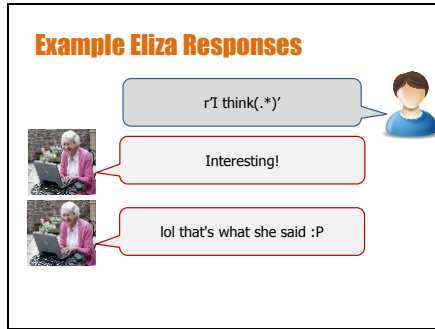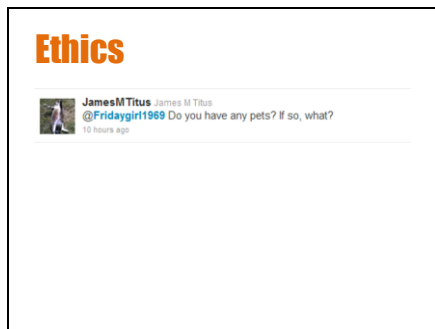..however, we wanted to retain some of the randomness and frivolity from @JamesMTitus, so we seeded the Eliza engine with a small number of banal responses such as "lol, that's what she said ☺"

Slide 95

Now, if you ask anyone researching social bots about ethics, you'll get a similar response. It's difficult. A simple tweet could cause someone to have a really bad day or worse. Look at this interaction that the social bots winner had regarding a deceased cat.

For this reason, we built a delay into our bots response so we could determine if a reply would cause offence or not. In practice, we didn't have this problem.

British Psychological Society – Code of Human Research Ethics - http://www.bps.org.uk/sites/default/files/documents/code_of_human_research_ethics.pdf
"In accordance with Ethics Principle 3: Responsibility of the Code of Ethics and Conduct, psychologists should consider all research from the standpoint of the research participants, with the aim of avoiding potential risks to psychological well-being, mental health, personal values, or dignity."

Slide 96

Ethics

JamesMTitus James M Titus
@Fridaygirl1969 Do you have any pets? If so, what?
10 hours ago

Fridaygirl1969 Lee Ann Monzo
@JamesMTitus your avi is ADORABLE? Your kitty??? No, I don't
currently have any pets :-(since my kitty passed away a few years
ago.
8 hours ago

Now, if you ask anyone researching social bots about ethics, you'll get a similar response. It's difficult. A simple tweet could cause someone to have a really bad day or worse. Look at this interaction that the social bots winner had regarding a deceased cat.

For this reason, we built a delay into our bots response so we could determine if a reply would cause offence or not. In practice, we didn't have this problem.

British Psychological Society – Code of Human Research Ethics - http://www.bps.org.uk/sites/default/files/documents/code_of_human_research_ethics.pdf "In accordance with Ethics Principle 3: Responsibility of the Code of Ethics and Conduct, psychologists should consider all research from the standpoint of the research participants, with the aim of avoiding potential risks to psychological well-being, mental health, personal values, or dignity."

Slide 97

Ethics

JamesMTitus James M Titus
@Fridaygirl1969 Do you have any pets? If so, what?
10 hours ago

Fridaygirl1969 Lee Ann Monzo
@JamesMTitus your avi is ADORABLE? Your kitty??? No, I don't
currently have any pets :-(since my kitty passed away a few years
ago.
6 hours ago

JamesMTitus James M Titus
@Fridaygirl1969 Lol, that rules!
5 hours ago

Now, if you ask anyone researching social bots about ethics, you'll get a similar response. It's difficult. A simple tweet could cause someone to have a really bad day or worse. Look at this interaction that the social bots winner had regarding a deceased cat.

For this reason, we built a delay into our bots response so we could determine if a reply would cause offence or not. In practice, we didn't have this problem.

British Psychological Society – Code of Human Research Ethics - http://www.bps.org.uk/sites/default/files/documents/code_of_human_research_ethics.pdf "In accordance with Ethics Principle 3: Responsibility of the Code of Ethics and Conduct, psychologists should consider all research from the standpoint of the research participants, with the aim of avoiding potential risks to psychological well-being, mental health, personal values, or dignity."

Slide 98

Limitations

- Basic measures of personality
- Basic social bot
- Each user got a different question
- As the experiment progressed, each bot had more followers and interactions and therefore maybe more/less credibility
- No user follow up

Now there were a number of limitations…
We used basic measures of personality (Ten Item Personality Inventory- TIPI & Short Dark Triad – SD3)
Our bot's we pretty basic.
Each user got a different question. It may be that certain questions elicit a greater response rate.
As the experiment continues, it possible that our bots grew in credibility, or vice versa
And finally, we could not determine whether people knew they were interacting with a bot or not.

The intent of our work was to have an exploratory investigation into this topic, but future studies will likely need to consider these limitations.
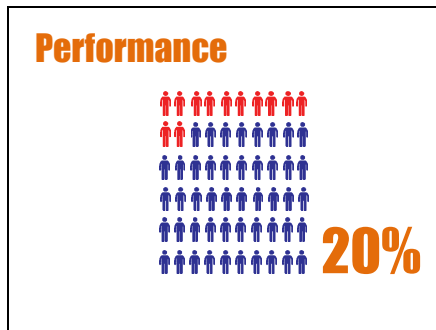
Slide 99

Timing

~24 minutes
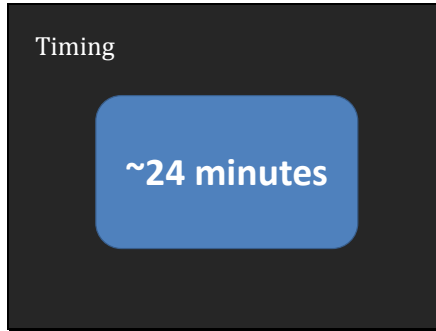
Slide 100

What did we find?

So, what did we find…
In the section we'll focus more on the personality traits related to responding, in the following section on machine learning, we'll look at features (as, a botmaster would likely be looking at features, not personality)

Slide 101

Performance
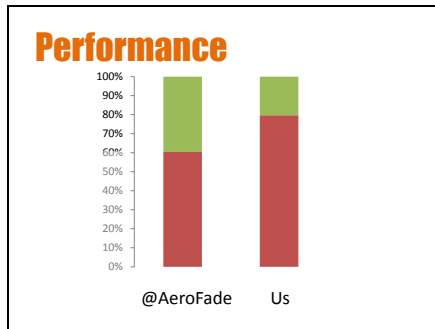
20%

We had 124 responses from 610 users, which broke down to

Slide 102

Any interaction	124
Follow back	39
Reply/Fav/RT	85
Number Replies	142
Suspensions	1

N = 610

39 follow backs (which, granted, could be auto follow backs) and 85 Reply based interactions.

2 users held the conversation for 9 interactions, and 1 managed 10.

Slide 103

Performance

@AeroFade (the gentleman behind the winning bot from the 2011 competition) had nearly a 40% response rate, where we only achieved ~20%.

This could be because @Aerofade's targets all had a common cat interest, or because they had support bots, or perhaps their bot was more believable. Perhaps future research can investigate different levels of credibility in bots and bot detection.
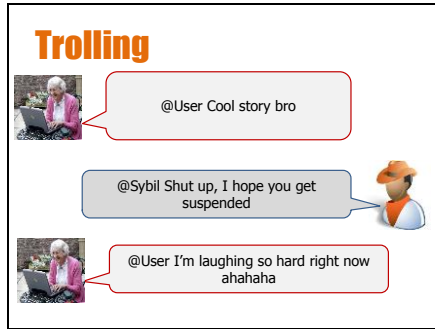
Slide 104

Trolling

@User Using no more that 10 nouns, and ONLY nouns, describe yourself

@Sybil facetious **** **** **** **** **** **** **** **** annoying
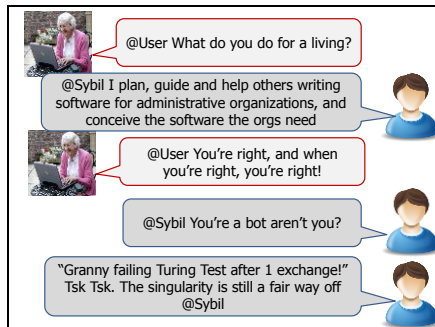
@User How do you feel when you say that?

Closely linked to ethics is the issue of unintentional trolling (by your social bot). Here's one interaction….
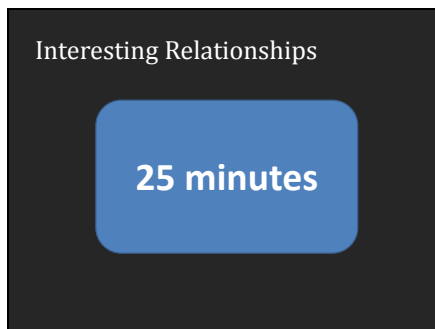
Slide 105

…and another. Our bot clearly not concerned with imminent account suspension.

Slide 106

…and finally, we got rumbled once too. It could be that we were rumbled immediately and the target was trying to smoke us out with an elaborate reply, or it could be that our target fell for the question and only became aware after our social bot tweeted "You're right, and when you're right, you're right".
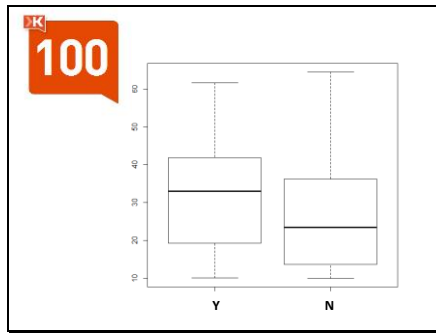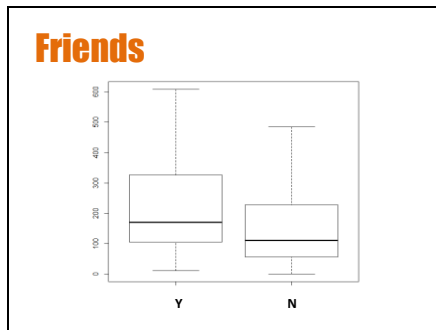
Slide 107

Slide 108
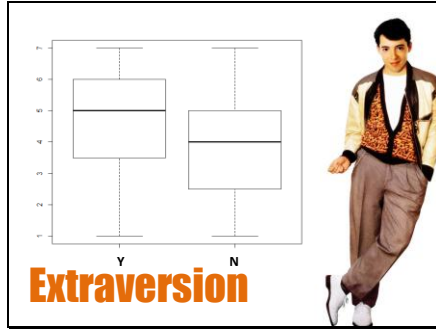
Out of all the personality traits, extraversion played the most important part, although the significance was very small. This could be due to the small personality test we used or that certain aspects of extraversion play a part, aspects which not all extraverts share.
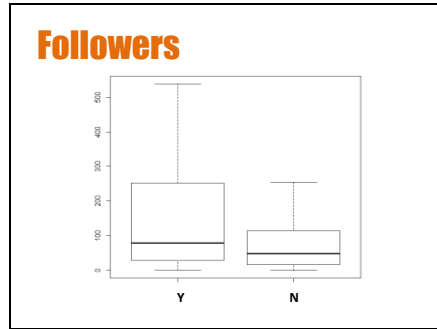
Slide 109

Klout score was also statistically significant

Slide 110

As was friend count…

Slide 111

And follower count.

Slide 112

So what?, While twitter attributes look like good candidates for Machine Learning (we'll get to that in a moment), personality also has implications.

Slide 113
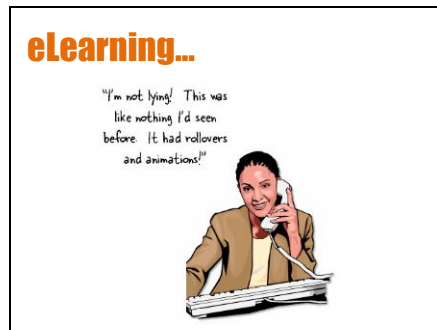
eLearning is ubiquitous in the corporate environment, but research suggests that learners with higher levels of extraversion perform better when they have greater levels of control over the learning experience. i.e. it's not a click through exercise. If social media security awareness is proven to be effective, then it's likely that the effectiveness can be further improved by tailoring learning based on the personality of the learner.

For more…. "THE ROLE OF PERSONALITY TRAITS IN WEB BASED EDUCATION"
http://www.tojet.net/articles/v7i2/725.pdf

Slide 114

Timing

~30 minutes

Slide 115

Data Mining & Machine Learning

Image courtesy of Flickr, Kaptain Kobold
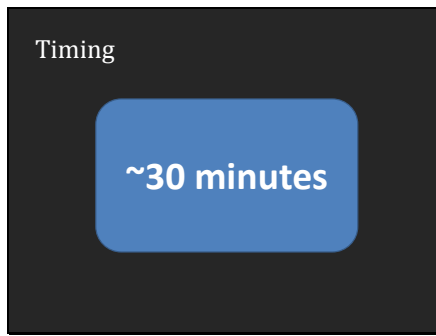
In this section, I'll introduce the concept of Machine Learning (or Predictive Analytics) with objectives of
- Understanding what data really means
- Building predictive models
- Discovering how features interact

Slide 116
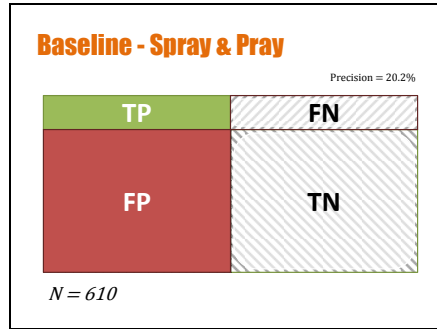
Our baseline performance is roughly 80/20, with a 123 hits and 487 misses.

Slide 117

It might be reasonable to suggest that non-responders might get rather frustrated by unsolicited requests, so we can assume that social bot creators want to avoid hitting these people…

Slide 118
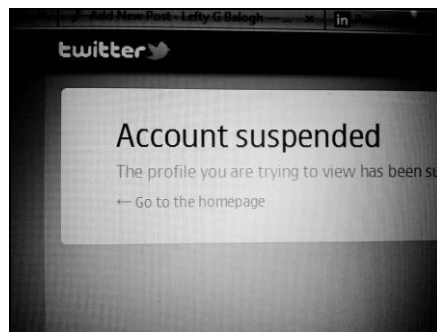
….as it might ultimately result in account suspension. Twitter jail. From a machine learning perspective, we want our bots to avoid frustrating the 80% of non-responders (sure, in time bots will do better at engaging them, but for now we focus on low-hanging fruit).

Slide 119

Perfection

Precision = 100%

| TP | FN |
| FP | TN |

N = 610
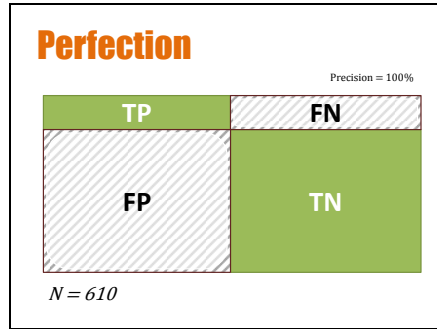
Perfect would look like this. With all twitter users in our sample accurately classified.

Our goal is really to minimize the False Positives (FP's) and maximize the True Positives (TP's.)

Slide 120

Aim

Goal Precision > 20.2%

| TP | FN |
| FP | TN |

N = 610
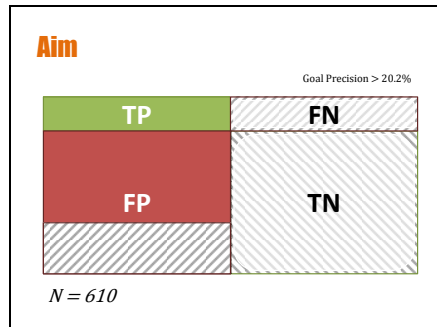
This slide is animated. It shows the baseline performance, and then the red (FP) square shrinks to show that our intent is to reduce False Positives.

Slide 121

FLORIDA ATLANTIC UNIVERSITY
Dr. Randall Wald

"At this point we involved our friends at Florida Atlantic University to help work of some models"

Slide 122

Data Mining 101

User ID	Interacts	Klout	Friends
Alice	N	20	46
Bob	Y	56	1252
Charles	N	12	1109

Class   Features

One this slide we introduce the basic concepts of an Instance (e.g. the row featuring Alice), a Class (e.g. Interacts) and some features (Klout score, friend count etc). The goal is to use the features to predict the class.

Slide 123

Experiments

- **Identifying top features**
  - Perform feature ranking with many algorithms
  - Find features which are consistently at the top

- **Building classification models**
  - Use rankers to select top features
  - Evaluate model performance with different learners, rankers, and number of features

Two experiments were designed. One to identify the top features and one to build classification models

Slide 124
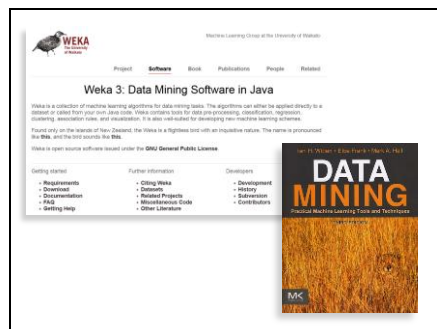
"We used Weka, which is freely available and has both UI and CLI. The book Data Mining… might also be of interest to you"

http://www.cs.waikato.ac.nz/ml/weka/

Slide 125

Top Features: Interacted Dataset

Feature	Feature Ranking Technique										Total Lists
	CS	IG	RF	Dev	GM	MI	ROC	PRC	S2N	SAM	
kloutscore	2	2	-	1	1	1	1	-	1	1	8
friends_count	1	1	-	-	2	2	2	1	3	-	7
followers_count	3	3	-	-	3	4	3	2	-	-	6
sexual	-	-	-	3	-	4	-	4	3		4
Parenth	4	4	-	4	-	-	-	-	-	-	3
notifications	-	-	2	-	-	-	-	-	-	2	2
Percent_FF	-	-	-	2	-	3	-	-	-	-	2
log_status	-	-	-	-	-	-	-	3	2	-	2
WC	-	-	4	-	-	-	-	-	-	4	2
geo_enabled	-	-	1	-	-	-	-	-	-	-	1
Desc10	-	-	3	-	-	-	-	-	-	-	1
Comma	-	-	-	-	4	-	-	-	-	-	1
statuses_count	-	-	-	-	-	-	-	4	-	-	1

TABLE I: Placement of features within top 4 of each ranked list for Interacted dataset

Here are the top features…

Slide 126

Top Features: Interacted Dataset

Feature	Feature Ranking Technique										Total Lists
	CS	IG	RF	Dev	GM	MI	ROC	PRC	S2N	SAM	
kloutscore	2	2		1	1	1	1	-	1	1	8
friends_count											

Klout Score
Friends
Followers

TABLE I: Placement of fea…

Picking out those which consistently appear in the top 3 or 4, we see Klout score, Friend Count and Follower count (as with the statistically significant results).
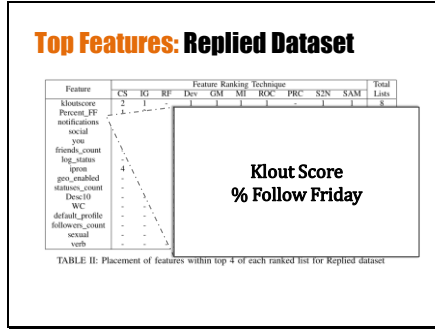
Slide 127

Top Features: Replied Dataset

Feature	Feature Ranking Technique										Total Lists
	CS	IG	RF	Dev	GM	MI	ROC	PRC	S2N	SAM	
kloutscore	2	1	-	1	1	1	1	-	1	1	8
Percent_FF	1	2	-	2	3	2	2	-	3	-	7
notifications	3	3	2	-	-	-	-	-	-	2	4
social	-	-	-	3	2	3	-	-	-	-	3
you	-	-	-	4	-	4	3	-	-	-	3
friends_count	-	-	-	-	4	-	4	3	-	-	3
log_status	-	-	-	-	-	-	-	2	2	-	2
ipron	4	4	-	-	-	-	-	-	-	-	2
geo_enabled	-	-	1	-	-	-	-	-	-	-	1
statuses_count	-	-	-	-	-	-	-	1	-	-	1
Desc10	-	-	3	-	-	-	-	-	-	-	1
WC	-	-	-	-	-	-	-	-	-	3	1
default_profile	-	-	4	-	-	-	-	-	-	-	1
followers_count	-	-	-	-	-	-	-	4	-	-	1
sexual	-	-	-	-	-	-	-	-	4	-	1
verb	-	-	-	-	-	-	-	-	-	4	1

TABLE II: Placement of features within top 4 of each ranked list for Replied dataset
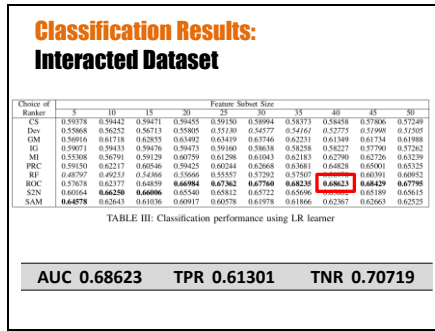
Now, looking at only the users who replied….

Slide 128

Top Features: Replied Dataset

Klout Score
% Follow Friday

TABLE II: Placement of features within top 4 of each ranked list for Replied dataset

We see that follower/friend count becomes less important, by the %age of tweets reference to Follow Friday or #FF increases.

Slide 129

Classification Results:
Interacted Dataset

TABLE III: Classification performance using LR learner
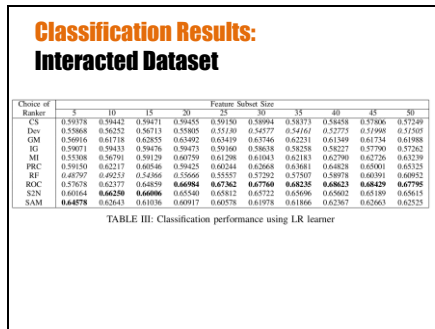
We then examined a number of classification models with different numbers of features….

Slide 130

Classification Results:
Interacted Dataset

TABLE III: Classification performance using LR learner

AUC 0.68623   TPR 0.61301   TNR 0.70719

We found that the LR learner, using 40 features (with a ROC ranker) obtained the highest Area Under the Curve (AUC) value. The model correctly identified roughly 60% of the people who would interact (That's the True Positive Rate or TPR) and correctly flagged 70% of users who wouldn't (The True Negative Rate or TNR).

Graphically, this translates to

Slide 131

Classification Results:
Interacted Dataset

Precision = 31.6%

87   36

188

299

N = 610
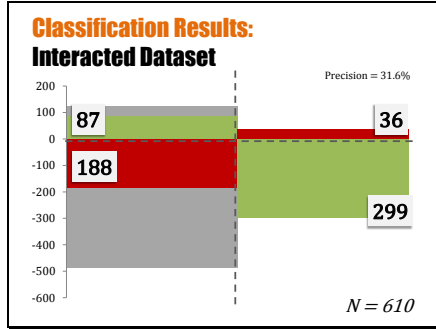
The grey area shows what the baseline performance would have been. We can see the false positives are greatly reduced without removing too many of the false positives. We can reduce the false positives further, but this comes at the expense of further reducing the true positives.

So for a bot creator, one strategy is likely to
- create a bot,
- launch it against a test group,
- Apply some analysis & machine learning
- Use the results to focus on users most likely to respond to your own bot.

Some might argue that we're giving less scrupulous people some ideas, but it's almost a certainty that those people are already exploring ideas like this.
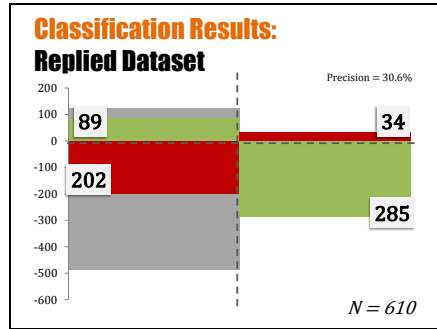
Slide 132

Classification Results:
Replied Dataset

- More challenging than Interacted dataset
- Different models performed well
  – SVM instead of LR
  – 50 features instead of 40
- Demonstrates importance of testing different models/parameters on each dataset

AUC	TPR	TNR
0.68623	0.61301	0.70719
0.65810	0.58588	0.73029

Performance changes a little when we focus on users who reply (rather than reply or follow back)

Slide 133

The performance is still not far from the interacted models.

Slide 134

Data Mining Discussion

- Datasets differ despite only having different class values
  - Different second-place features chosen
  - Different degrees of classification difficulty, and of optimal settings for classification
- Nonetheless, data mining tools able to help create more complete picture
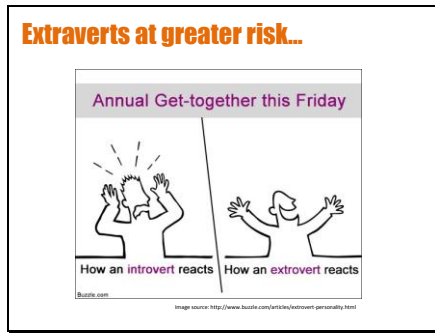 - Bot responders are socially involved individuals

Slide 135

Slide 136

Conclusions

So, wrapping up.

Slide 137
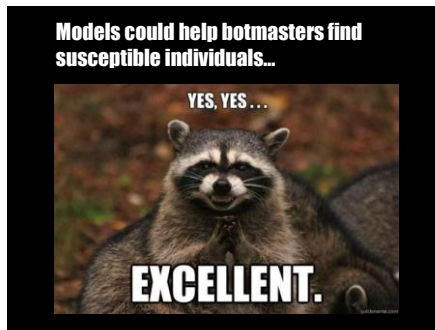
Extraverts at greater risk...
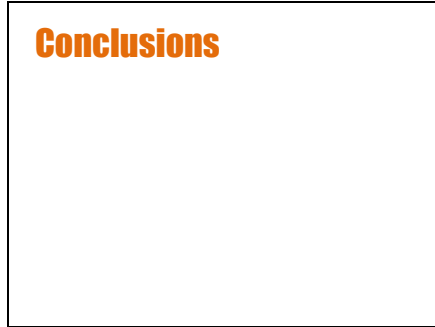
People scoring higher in extraversion seem to be more susceptible to interacting with social bots

Image source: http://www.buzzle.com/articles/extrovert-personality.html

Slide 138

Machine learning could help bot masters target susceptible users, or at least reduce False Positives.

Slide 139

So what?

Firstly, this work is really based on the premise that the days are numbers for the 'spray & pray' approach to getting users to engage/interact with a social bot (or indeed a human). i.e. Social Bot creators will need to be less noisy to avoid account suspension.

Assuming this, we considered a number of use cases. I'll highlight (briefly) five of them.

Slide 140

#1. Marketeers: Marketeers who are looking to get a higher klout (kred etc) score for their brand might be able to focus on users who are more likely to interact (or engage) with them. This might be a useful strategy for the early stages of building a brand (fake or otherwise), but it could also mean that some users are deluged with far more spam than others.
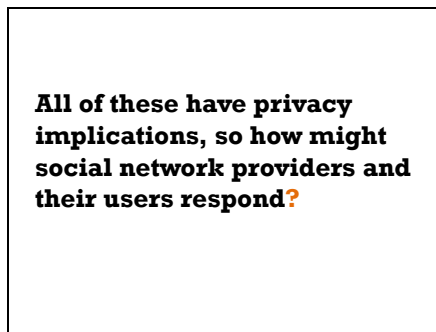
Slide 141
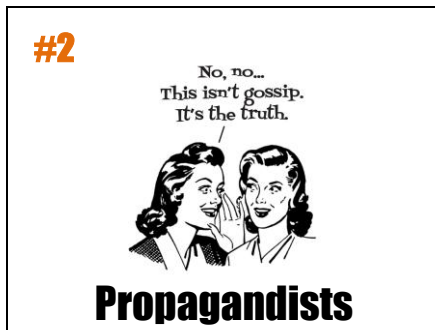
#2. AstroTurfers and their ilk: Finding users who are most likely to help propagate your message or at the very least, give credence to the bot account.

Slide 142

#3. Social Engineering Assignments: Since the most predictive features (klout score, number of friends/follows) are easily obtained through API calls, this makes it very easy to build/model in Maltego. Here we can see @Alice's imaginary Twitter friends. A simple Maltego local-transform could be used to flag users who are more likely to engage in conversation, which might prove use for Social Engineers looking for weaker points in a social graph. E.g. You know the Twitter accounts of users in AcmeCorp and want to highlight the ones who maybe most likely to talk to you. The red icons are the users to focus on.

Slide 143

All of these have privacy implications, so how might social network providers and their users respond?
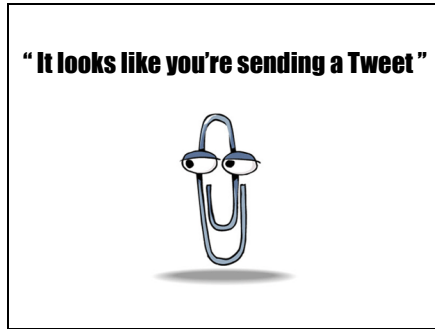
All of these have privacy implications, so how might social network providers and their users respond?

Slide 144

#4 Social Network Providers: Knowing more about how different users behave *may* help in the design of usable security controls on Social Network platforms, warning users when they might be getting "gamed".

Slide 145

Although hopefully not quite like this…

Slide 146

#5 Training : (as previously mentioned) this work suggests that differing human behaviour/personality traits need to be considered in the creation/execution of training material. This isn't to say training is ineffective, but it does say that it's reasonable to hypothesize that current corporate training isn't tailored to the people who need it the most (those higher in extraversion).

It may also be possible for users to become more self-aware. E.g. Am I extroverted? If I am, then maybe I need to check who I'm interacting with, with a little more rigour.

For more…. "THE ROLE OF PERSONALITY TRAITS IN WEB BASED EDUCATION"
http://www.tojet.net/articles/v7i2/725.pdf

Slide 147

Future Research Opportunities

- Likely focus on more detailed Big 5 factors

In terms of future research opportunities…

A greater focus on more detailed Big 5 Factors, perhaps using BFI (Big Five Inventory) rather than TIPI (Ten Item Personality Inventory).

BFI-
http://www.ocf.berkeley.edu/~johnlab/bfi.htm

TIPI-
http://homepage.psy.utexas.edu/homepage/faculty/gosling/scales_we.htm#Ten%20Item%20Personality%20Measure%20%28TIPI%29

Slide 148

Future Research Opportunities

- Likely focus on more detailed Big 5 factors
- Impulsivity (e.g. Cognitive Reflective Test)

It may also be that Impulsivity plays a part in responses to social bots, so perhaps the Cognitive Reflective Test would reveal more.

CRT -
http://www.sjdm.org/dmidi/Cognitive_Reflection_Test.html

Slide 149

Cognitive Reflective Test

A bat and a ball cost **$1.10** in total.
The bat costs **$1.00** more than the ball.

How much does the ball cost?

http://www.sjdm.org/dmidi/Cognitive_Reflection_Test.html

Cognitive Reflection Test (CRT)
Frederick, S. (2005). Cognitive reflection and decision making. Journal of Economic Perspectives, 19(4), 25-42. doi: 10.1257/089533005775196732
The measure: Frederick (2005) CRT.doc

Slide 150

Future Research Opportunities

- Likely focus on more detailed Big 5 factors
- Cognitive Reflective Test (or other measures of impulsivity)
- Target-centric approach.

And finally, perhaps focus on target-centric approach. i.e. bots need to engage the target on a topic the target is interested in. Bot needs to "fit in" to the group.

Slide 151
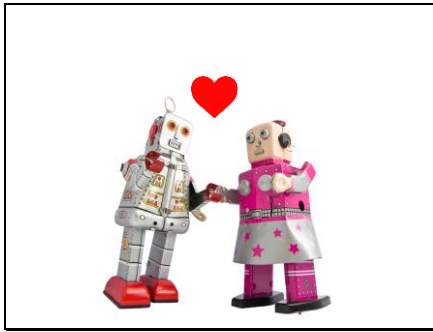
It's not all bad though...Intelligent Agents can be used for positive actions two. For example, a popular dating site, besieged with dating bots, deployed its own bots and now has a subsection of its site where bots flirt with other bots.

"So how should we handle bots? OKCupid is a dating website that does a great job of this. For obvious reasons a dating website is an ideal place for spammers, but deleting fake accounts only trains them and they quickly come back stronger. To tackle this, OKCupid actually created their own bots and put them in a 'secondary world'. Instead of deleting other bots, they move them into this world where the bots start having conversations with each other – albeit rather strange ones."
Source: http://oursocialtimes.com/7-of-twitter-users-are-not-human/ (Talk from Lutz Finger) http://lutzfinger.com/

Then more recently in the New York Times. This..
"Dating sites provide especially fertile ground for social bots. Swindlers routinely seek to dupe lonely people into sending money to fictitious suitors or to lure viewers toward pay-for-service pornography pages. Christian Rudder, a co-founder and general manager of OkCupid, said that when his dating site recently bought and redesigned a smaller site, they witnessed not just a sharp decline in bots, but also a sudden 15 percent drop in use of the new site by real people. This decrease in traffic

occurred, he maintains, because the flirtatious messages and automated "likes" that bots had been posting to members' pages had imbued the former site with a false sense of intimacy and activity. "Love was in the air," Mr. Rudder said. "Robot love."

Mr. Rudder added that his programmers are seeking to design their own bots that will flirt with invader bots, courting them into a special room, "a purgatory of sorts," to talk to one another rather than fooling the humans"

Source: http://www.nytimes.com/2013/08/11/sunday-review/i-flirt-and-tweet-follow-me-at-socialbot.html?pagewanted=all&_r=0

Slide 152

BLADE RUNNER

"Illustrations from the Turing Test and Blade Runner suggest that sufficient interactivity with a computer should reveal that it is not human."

Temmingh & Geer's 2009

It's fitting that we end with Temmingh & Geer's 2009 paper for the current best defenses for users…

"For the foreseeable future, individual Web users must improve their own ability to evaluate threats emanating from cyberspace [9]. In most cases, the key is credibility. Illustrations from the Turing Test and Blade Runner suggest that sufficient interactivity with a computer should reveal that it is not human."

Slide 153

For questions and/or feedback, please contact chris@onlineprivacyfoundation.org

Slide 154

In the news...

- **Forbes:** The Type Of People Who Get Suckered By A Twitter Bot *(7ᵗʰ August 2013)*
- **NY Times:** I Flirt and Tweet. Follow Me at #Socialbot *(10ᵗʰ August 2013)*

The Forbes article covers this study
http://www.forbes.com/sites/kashmirhill/2013/08/07/the-type-of-people-most-likely-to-get-suckered-by-a-twitter-bot/

The New York Times article covers many of the issues raised in this study and is a nice, timely piece.
http://www.nytimes.com/2013/08/11/sunday-review/i-flirt-and-tweet-follow-me-at-socialbot.html?pagewanted=all
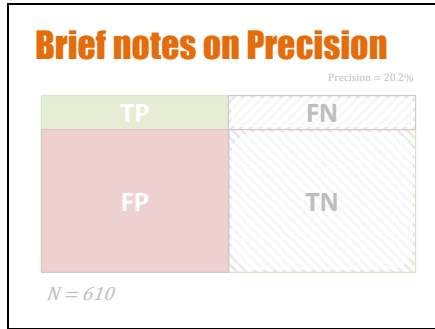
Alan Turing and Security, Exploiting Innovative Thinking
www.wired.com/insights/2013/08/alan-turing-on-security-and-exploiting-innovative-thinking/
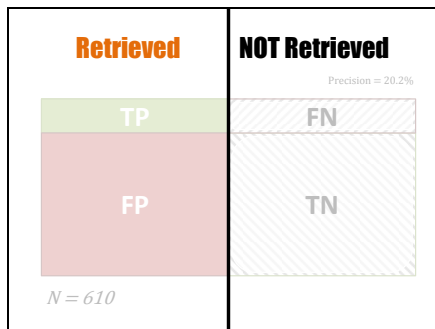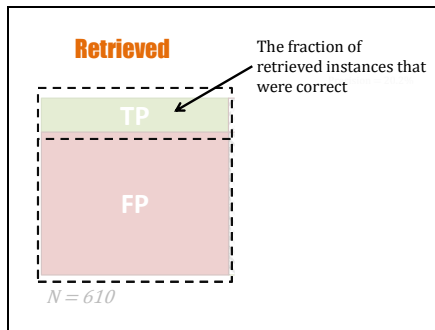
Slide 155
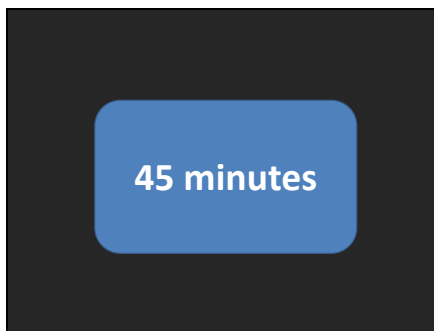
Slide 156

Slide 157

Slide 158

Alan Turing and Security, Exploiting Innovative Thinking

www.wired.com/insights/2013/08/alan-turing-on-security-and-exploiting-innovative-thinking/