



1

Copyright © 2012 Chris Sumner and the Online Privacy Foundation. This work may be used without restriction for the purposes of academic research, teaching, and journalism, on condition that all copies and reproductions retain this notice intact. Permission should be obtained for the use and reuse of this work for any other purpose. Conditions and a fee may apply. Please contact chris@onlineprivacyfoundation.org

General Flow of talk

Before I continue, and by means of an introduction, I'm Chris Sumner. Together with a small number of likeminded individuals, I co-founded the volunteer organisation, The Online Privacy Foundation (Twitter @The_OPF). It is in this capacity that I am speaking today. I'm joined today by Randall Wald, a PhD student at the Florida Atlantic University. I ran into Randall at the Q&A room after our Defcon 19 talk on Facebook personality (https://media.blackhat.com/bh-ad-11/Sumner/bh-ad-11-Sumner-Concerns_w_Facebook_WP.pdf). Randall asked if we'd considered machine learning techniques, we explained that we didn't have those skills and so a relationship was born.

Note: All images via Shutterstock unless otherwise stated.



2

General Flow of talk

- So can you or can't you determine psychopathy from Twitter? If so, to what extent and what are the limitations. That's what we attempted to find out.
- The question itself can be interpreted in a number of ways. e.g.
 - Can we expose psychopaths (defined, for arguments sake, as a PCL-R score over 25) - "No"
 - Can we expose traits associated with psychopathy? - "To some extent"
 - Using Twitter, can we predict who scores high in psychopathy and who doesn't? - "We'll let you decide, based on the results we share today"
 - Is it useful? - Maybe
- It is the second question we focused on and NOT the first
- This might not be the talk for you if you're expecting a black and white answer, as most of you can probably guess, the reality is a little bit more complicated than that, but is interesting never the less.



3

General Flow of talk

Our default position approaching social media based personality research is one of caution and healthy skepticism, we don't have to look far to see where people got things wrong in the past, i.e. Phrenology and



4

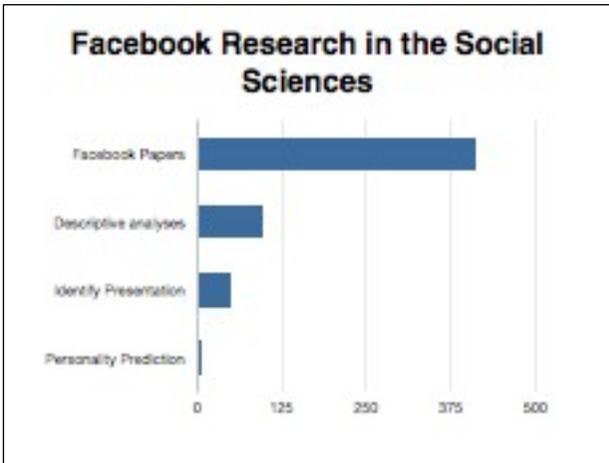
General Flow of talk (slide didn't make it into DEFCON presentation)

Graphology <http://www.nytimes.com/1993/08/03/news/03iht-grapho.html>

This sort of research, social media research, is in its infancy and as such we (all) have an opportunity to get involved and really examine/question the results and methods.

For the curious, here's a list of Facebook research in the social sciences: <http://psych.wustl.edu/robertwilson/> Danah Boyd has a comprehensive list of Twitter based studies here <http://www.danah.org/researchBibs/twitter.php>

We, The Online Privacy Foundation, are consolidating pointers to relevant research, with an initial focus on Social Media based Personality Prediction papers. <https://www.onlineprivacyfoundation.org/research/>



5

General Flow of talk (slide didn't make it into DEFCON presentation)

Illustrative breakdown of Facebook research in the social sciences: <http://psych.wustl.edu/robertwilson/>

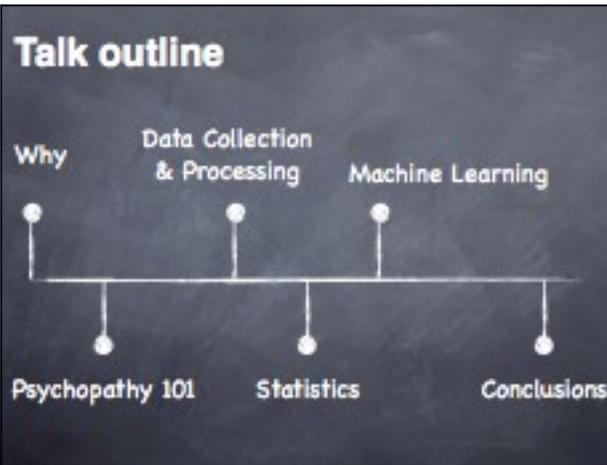
i.e. this is a nascent area of study.



6

General Flow of talk

Our approach has been to have a good look at social media behaviour residue research and come to you with our findings. If you're expecting a live demo where we unmask the highest scoring psychopath in the room, you're at the wrong talk, but if you want to see the early results from our study, our methodology and some psychopathy 101, then you're very much in the right talk.



7

General Flow of talk
So in this talk we'll cover

- Why? what drove us to study this field in the first place
- Psychopathy 101
- Data Collection & processing
- Statistics
- Machine Learning (this is going to move from covering some basic concepts, to something more in-depth, if you want more of the in-depth, please find us in the Q&A room or drop us an email)
- Conclusions



8

General Flow of talk
Starting with "why?"



9

General Flow of talk
Social media research is something we've been looking at for a while, last year we spoke about a similar experiment on Facebook. I spoke about this experiment a week ago (July 14th 2012) at the European Conference on Personality Psychology and Randall has a paper to be shared at IEEE Information Reuse in a few days time here in Las Vegas. We felt the results were very interesting, but not enough to base a critical decision on (such as a hiring decision), on it's own.

Wald, R., Khoshgoftaar, T. M., and Sumner, C. (2012). Machine Prediction of Personality from Facebook Profiles, 13th IEEE International Conference on Information Reuse and Integration, 2012, pp. 109-115.

Sumner, C., Byers, A., & Shearing, M. (2011). Determining personality traits & privacy concerns from Facebook activity. Black Hat Briefings '11. Abu Dhabi, United Arab Emirates.

Image is property of The Online Privacy Foundation.

Online searches of candidates



Jobvite Social Recruiting Survey 2012

10

General Flow of talk

One of our concerns related to the accuracy of automated feature extraction when used on it's own. It's worth pointing out that there are two key approaches to social media personality prediction.

1) **Observer rated** - where a human rates one or more personality traits of a subject.

2) **Automated Feature Extraction** - where certain features are extracted programatically, such as the number of Twitter or Facebook friends someone has.

We are concerned because according to a recently released Jobvite survey, 48% (This figure was 45% last year, by the way.) of hiring managers are routinely checking candidates social media profiles. There are important considerations regarding discrimination and unfair invasion of privacy, but that aside, what's the false positive, false negative rates? Is there enough empirical evidence to support the vetting of potential employees based on their social media use? Indeed, is it even ethical? Observer based studies tend to show greater accuracy overall, but there's concern that companies will spring up offering automated pre-screening.

For the 48% of hiring managers who do look at social media activity, their choices may, in part be supported by headlines such as those in the next two slides

"Facebook profiles reveal personality traits to researchers"

11

General Flow of talk (slide didn't make it into DEFCON presentation)

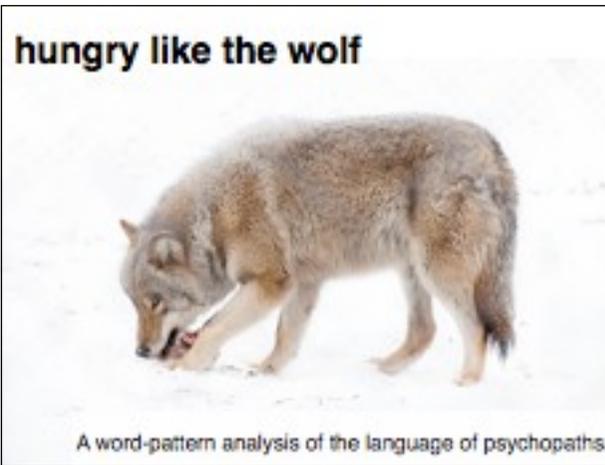
On May 9th 2011 we saw this article.

"Facebook can serve as a personality test"

12

General Flow of talk (slide didn't make it into DEFCON presentation)

Here's another report...same topic much more sure...same thing....but this seems to suggest that Facebook can be used as a reliable replacement for a personality test. Can it?

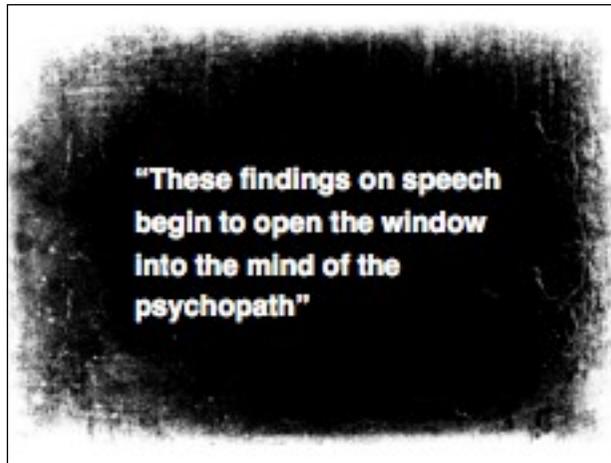


13

General Flow of talk

In September 2011, Prof Jeff Hancock and colleagues released a paper titled “Hungry like the wolf: A word pattern analysis of the language of psychopaths”. Although a reasonably small sample, he did have access to clinical psychopaths and his was the first study to examine language in psychopaths through statistical analysis of a corpus. IMO, it’s a very nice paper.

[http://www.cbc.ca/fifth/37/episodes/murderhewrote/images/Hancock%20Woodworth%20%20Porter%20\(2011\)Hungry%20Like%20The%20Wolf%20-%20The%20Language%20of%20the%20.pdf](http://www.cbc.ca/fifth/37/episodes/murderhewrote/images/Hancock%20Woodworth%20%20Porter%20(2011)Hungry%20Like%20The%20Wolf%20-%20The%20Language%20of%20the%20.pdf)



14

General Flow of talk

The Hancock paper stated “these findings on speech begin to open the window into the mind of the psychopath, allowing us to infer that the psychopath’s world view is fundamentally different from the rest of the human species”.

However, that paper led to this...



15

General Flow of talk

...and the title of our talk “Can Twitter help expose psychopath killers’ traits?”

Naturally, this led to a raised eye-brow or two, although the headline is more speculative than anything. Other papers have been released on Twitter personality prediction, most notable by Golbeck et al and Quercia et al, so studying psychopathy on Twitter isn’t a stretch of the imagination,

<http://www.cs.umd.edu/~golbeck/pubs/Golbeck%20et%20al.%20-%202011%20-%20Predicting%20Personality%20from%20Twitter.pdf>

<https://www.cl.cam.ac.uk/~dq209/publications/quercia11twitter.pdf>

As mentioned earlier...The question itself can be interpreted in a number of ways. e.g.

- Can we expose psychopaths (define as a PCL-R score over 25 - for arguments sake)
- Can we expose traits associated with psychopathy?



"I worry about people being judged by these criteria. Psychopaths make up about 1% of the population, so even a small false-positive rate can be a significant problem."

16

General Flow of talk

Bruce Schneier commented on the press surrounding the paper (not the paper itself) on his blog highlighting the issue of false positives.

http://www.schneier.com/blog/archives/2011/11/detecting_psych.html

Link to original image/store <http://geekz.co.uk/shop/store/show/schneier-tshirt.html>

The Telegraph

HOME NEWS WORLD SPORT FINANCE COMMENT BLOG CULTURE TRAVEL LIFE FAMILY
Technology News Technology Companies Technology Reviews Video Games Technology Video

FBI to use Twitter to predict crimes

The FBI is trying to develop software that will scan Twitter and Facebook to look for "emerging threats".



17

General Flow of talk

But now it gets interesting because the headlines are transcending employee pre-screening and recommender systems and turning to crime prediction (not necessarily through personality, in-fact I'm guessing not through personality).

<http://www.telegraph.co.uk/technology/news/9043690/FBI-to-use-Twitter-to-predict-crimes.html#>

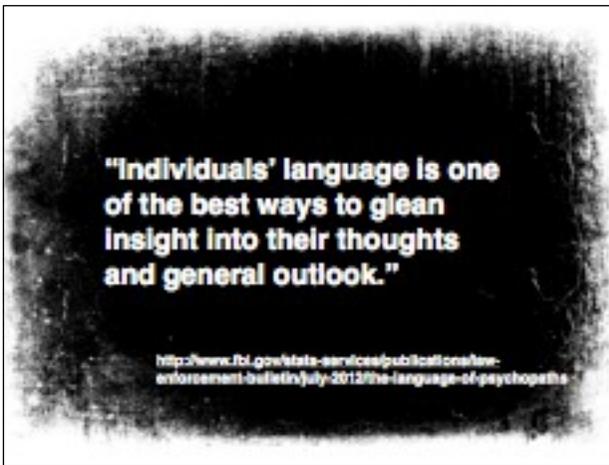
Automatic Crime Prediction Using Events Extracted from Twitter Posts

Xiaofeng Wang, Matthew S Gerber and Donald E. Brown, 2012

18

General Flow of talk

In fact there's already a paper out on Twitter Crime Prediction by Wang and colleagues (although this isn't based on personality, but it is based on linguistic analysis)



19

General Flow of talk

On Thursday (July 26th 2012) we saw this on the FBI.gov website talking specifically about Psychopathy, relating back to the Hancock paper.

“Individuals’ language is one of the best ways to glean insight into their thoughts and general outlook. Recent advances in technology make it possible to examine more closely the language of various clinical populations through automatic linguistic analysis programs. These applications can differentiate between a variety of individual and personality factors.⁴ The tools range from simple to sophisticated, but they all essentially identify linguistic patterns and count their frequency relative to a control language”

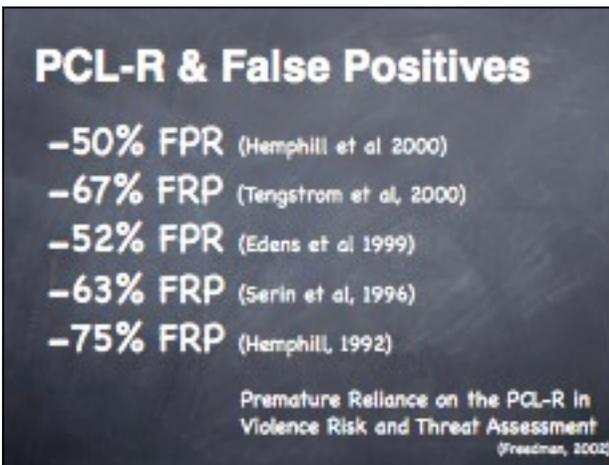


20

General Flow of talk

Highlighting people who may need help (Intervention) isn't new either

<http://blog.thoughtpick.com/2010/04/crime-prediction-software-implications-dangers.html>



21

General Flow of talk

To address an obvious and valid concern about psychopathy and crime, for some balance, some papers suggest there's a high false positive rate when looking at people who score highly in Psychopathy versus whether they go on to commit crime, although others have argued that the remaining number just aren't getting caught.

Just because someone scores high in psychopathy doesn't make them a criminal.

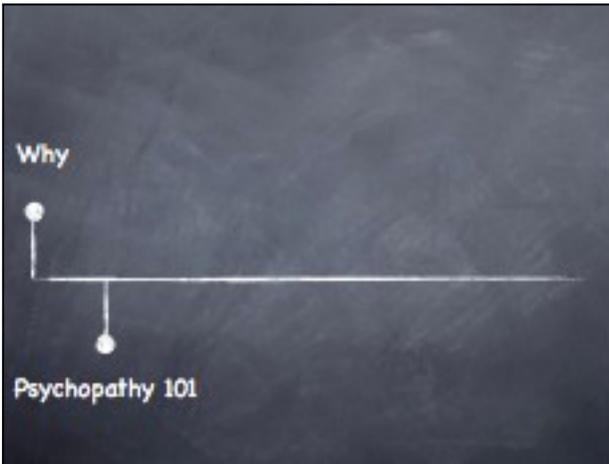


22

General Flow of talk

It's also important to note that we examined Psychopathy as a scale. There's an important difference between investigating Psychopathy and detecting people who meet the criteria of a "Hare" Psychopath (typically 25-30+ on the PCL-R scale). That aside, can psychopathic traits be observed in Twitter usage?

We will provide conclusions and you can tell us whether you think the headlines are correct and indeed whether the practice of social media vetting really requires much more scrutiny.



23

General Flow of talk

So why the interest in Psychopathy?

We examined 8 personality traits in total, but psychopathy captures peoples imagination in ways other personality traits do not.



24

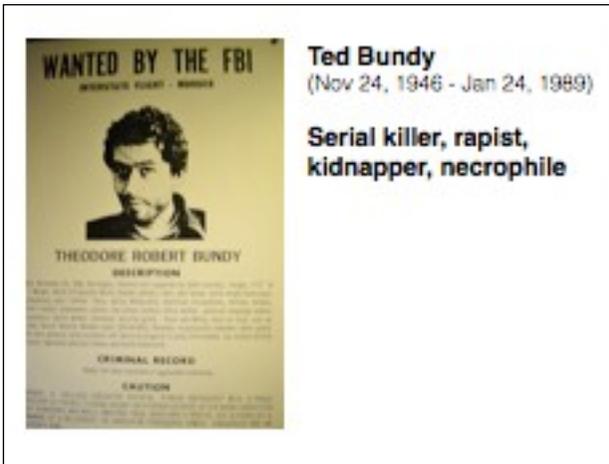
General Flow of talk

Timothy McVeigh



General Flow of talk
Richard Ramirez

25



General Flow of talk
Ted Bundy

26



General Flow of talk
And there fascination continues in movies...The movie Psycho, although Norman Bates appears to be Psychotic rather than a psychopath as such. Still, the term Psycho is generally used as a short form for Psychopath.

27



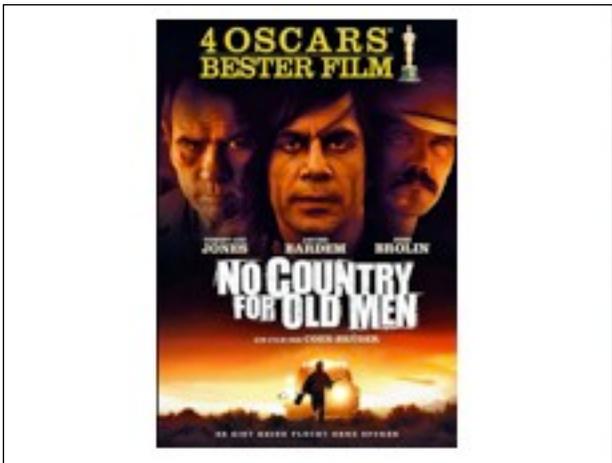
General Flow of talk
The Shining

28



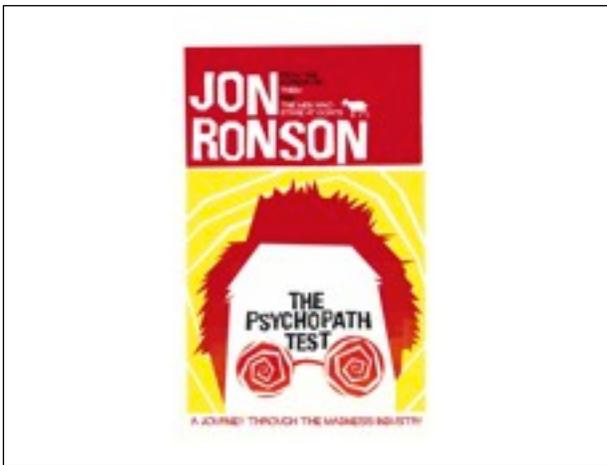
General Flow of talk
Silence of the Lambs, perhaps the most famous movie Psychopath

29



General Flow of talk
And one of the best films depicting a possible psychopath in my opinion, No Country for Old Men

30

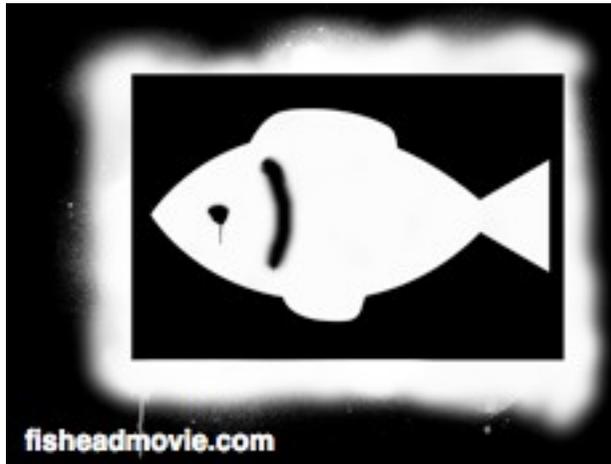


31

General Flow of talk (slide didn't make it into DEFCON presentation)

And of course, Jon Ronson's book brought psychopathy onto the metaphorical coffee table (it's worth a read if you haven't already read it).

Read the Bob Hare comments regarding this book here: http://www.psychopathysociety.org/index.php?option=com_content&view=article&id=115&Itemid=132&lang=us



32

General Flow of talk (slide didn't make it into DEFCON presentation)

and last year (2011), this film was released. fisheadmovie.com - available at <http://www.youtube.com/watch?v=KUbjaI3X5Qk>
Also available here (Password = fhmovie) <http://www.fisheadmovie.com/watch-the-movie>

Other videos worth hunting out

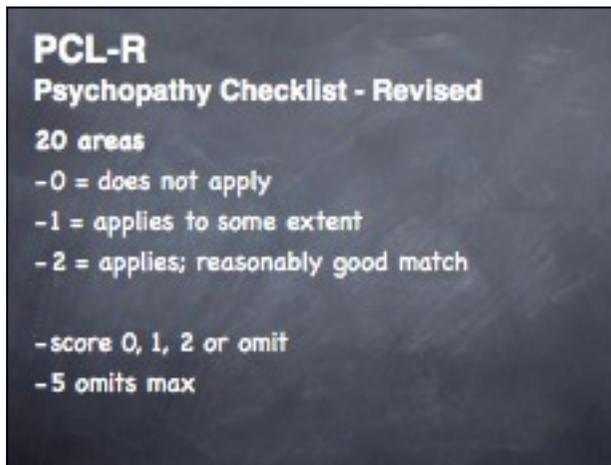
Look also for BBC Horizon - Are you good or Evil?

<http://www.bbc.co.uk/programmes/b014kj65>

http://www.youtube.com/watch?v=u88lYs4FMTY&feature=list_other&playnext=1&list=SP4B8EC020E34E227F

and Jim Fallon TED talk - Jim Fallon: Exploring the mind of a killer

http://www.ted.com/talks/jim_fallon_exploring_the_mind_of_a_killer.html



33

General Flow of talk

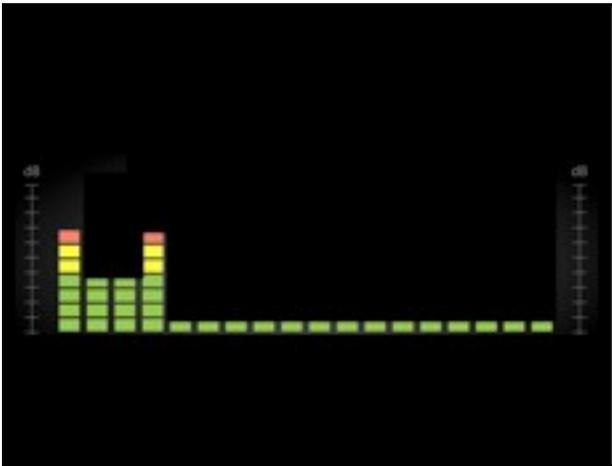
So, let's take a look at the most commonly used Psychopathy test, the Psychopathy Checklist -Revised (or PCL-R) by Robert D Hare. the test (or instrument) that give rise to the title of Ronson's book.

It consists of 20 areas where you can either be scored a 0, 1, 2 or an "omit".



General Flow of talk
 Many people might score nothing at all

34



General Flow of talk
 Others might score high in some areas, but not others, and certainly not enough to be labeled a "Psychopath"

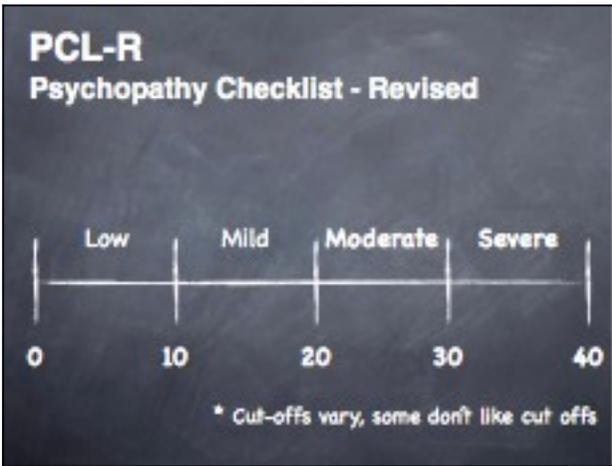
35



General Flow of talk
 At the extreme end, a minority of people might score high in many categories.

It's worth noting that the equaliser graphic might suggest these dimensions are dynamic or that they fluctuate with the winds, this is NOT the case.

36

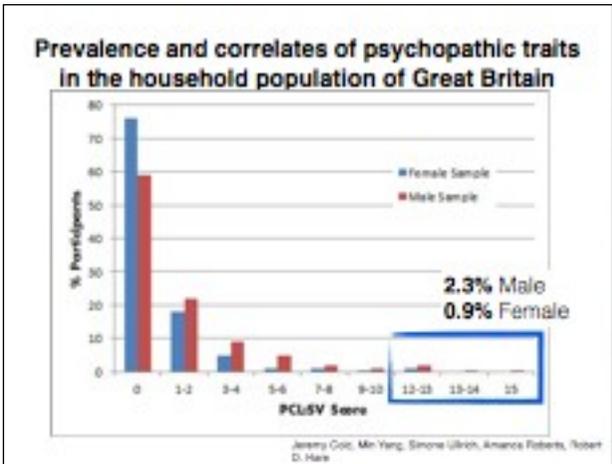


37

General Flow of talk

This is how Dr. Reid Meloy, author of the Psychopathic Mind, generally segments people who score on the PCL-R continuum

It should also be stated that PCL-R should only be conducted by trained psychiatrists and also that some psychologists are not at all fond of cut-off points or indeed PCL-R. The discussion on that continues.



38

General Flow of Talk:

So how many of these people are out there?

This study of the population of Great Britain shows 2.3% Male and 0.9% Male scoring between 11 and 15 (pretty high on the PCL:SV). Note '12-13 should say 11-12.

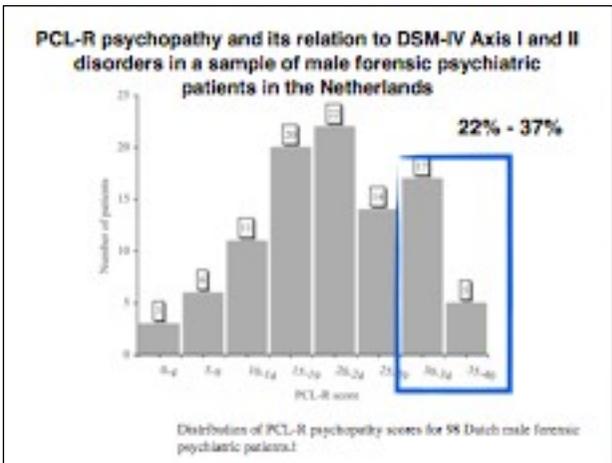
PCL:SV is the Screening Version of the PCL-R. Read more here. <http://www.hare.org/scales/pclsv.html>

For the curious, these figures for high scoring PCL-R individuals came from the PCL-R course.

- 1.49% - 5.66% (US) V 0.1 - 0.22 Taiwan

For the GB Study....

The weighted prevalence of "possible" psychopathy, using a suggested cut score of 11 or more points in this population was 2.3% (95% CI: 1.2– 3.8); 3.7% (95% CI: 1.8–6.6) in men and 0.9% (95% CI: 0.2–2.8) in women. The prevalence of possible psychopathy using a cut score of 13 or more was 0.6%; 1.3% in men (95% CI: 0.3–3.4) and 0% in the females.



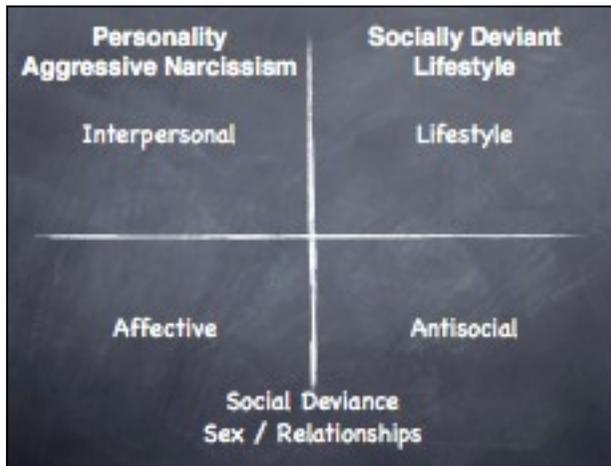
39

General Flow of Talk:

However, the distribution of psychopathy is notably different in prison and psychiatric populations. This study is of Dutch male patients. Other studies put the figure between 10 and 25% of the prison population, with a further ~50% having APSD and the remaining ~25% or so trying to keep off the radar.

Study: PCL-R psychopathy and its relation to DSM-IV Axis I and II disorders in a sample of male forensic psychiatric patients in the Netherlands

Martin Hildebranda, *, Corine de Ruiterc



40

General Flow of Talk:

Many (not all) psychologists concur that Psychopathy consists of 2 factors (Personality Aggressive Narcissism and Socially Deviant Lifestyle) with 4 facets. The list below shows the 4 facets and the numbers which map to the PCL-R categories/questions (this also explains why the numbers on the next few slides will sometimes be out of sequence):

- Interpersonal - 1, 2, 4, 5
- Affective - 6, 7, 8, 16
- Lifestyle - 3,9, 13, 14, 15
- Antisocial - 10, 12, 18, 19, 20

2 categories are not loaded to any factor. Social Deviance - 11, 17



41

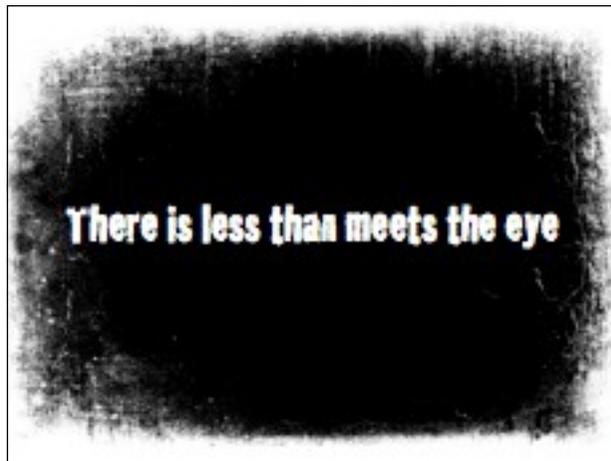
General Flow of talk

Interpersonal Facet:

The first category, Glibness & Superficial charm... They talk the walk, but they don't necessarily walk it. Many of the contestants on the TV show the Apprentice demonstrate glibness remarkably well. People fall for the patter very easily indeed....

...it's important to note that this doesn't make them psychopaths or even "Almost Psychopaths" (see the book of the same title) <http://www.amazon.co.uk/Almost-Psychopath-Someone-Problem-Manipulation/dp/1616491027>

This superficial charm led to....



42

General Flow of talk

this oft used line... "There's nothing behind the mask." This concept gave rise to the title of Cleckley's seminal work, "The Mask of Sanity", available free for download. http://www.cassiopaea.org/cass/sanity_1.PdF

2. Grandiose Sense of Self Worth

- Cocky
- Opinionated
- Self-assured
- Exaggerated regard for his/her own abilities



43

General Flow of talk
Interpersonal Facet:

Another apparently common trait of many (but not all) of the contestants in the TV show the Apprentice, although we're not still suggesting they're psychopaths, a grandiose sense of self worth.

This is strongly related to narcissism. They may claim success for projects that they added little real value to, a narcissistic takeover. (although people sometimes do that for other reasons)...

They attend training classes thinking "I'm not going to learn anything here", even if the class leader is the worlds leading expert on the topic.

They maintain their self worth through devaluation of other people. Their objective is domination.

4. Pathological Lying

"His readiness to lie can be quite remarkable"



44

General Flow of talk
Interpersonal Facet:

People typically lie to manage anxiety, but not the psychopath, their motivation is manipulation and domination
They are possibly outed by non-verbal cues (i.e. gestures don't match the words), although research is thin here. (symptom of hollowness, they're play acting, might not nail the part completely.)

5. Conning / Manipulation



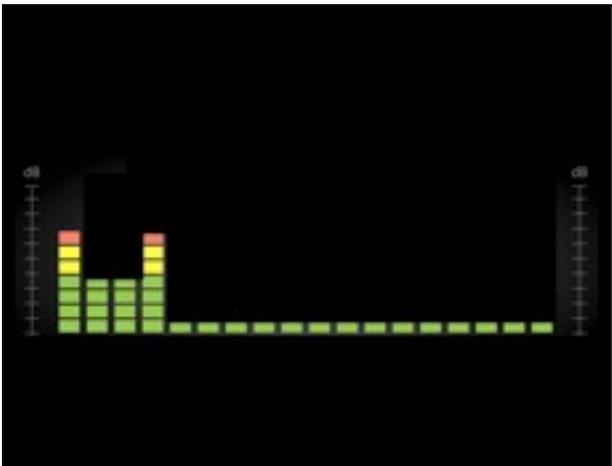
45

General Flow of talk - Interpersonal Facet:

Which comes in very hand here. Using all the tools in their toolbox, Charm offensive, self worth, the plausible stories, which turn out to be lies. Hardly a week goes by without a retired person falling victim to a scam that robs them of their life savings. Here's one from this week

http://www.thenorthernecho.co.uk/news/9839402.Darlington_pensioner_conned_out_of_400_000_in_savings/

It's impossible for many of us to comprehend how someone could does this...more of that shortly. t's also important to note, that someone who may have scored higher in the previous points, may actually show some restraint here, which would indicate at least some level of empathy/guilt (more on that later too). This could equip them to do well in the TV show the Apprentice but still be pretty decent people in the grand scheme of things



46

General Flow of talk

So if you score HIGH on those previous four categories, would this make you a psychopath? No. not exactly a saint, but certainly not a psychopath.

Also important to remember that these behaviours/traits need to be a constant feature of someone's life, not something someone did once. It's possible that a pressured sales person could go in for the con and feel extreme guilt once his/her transgressions are pointed out to him/her.



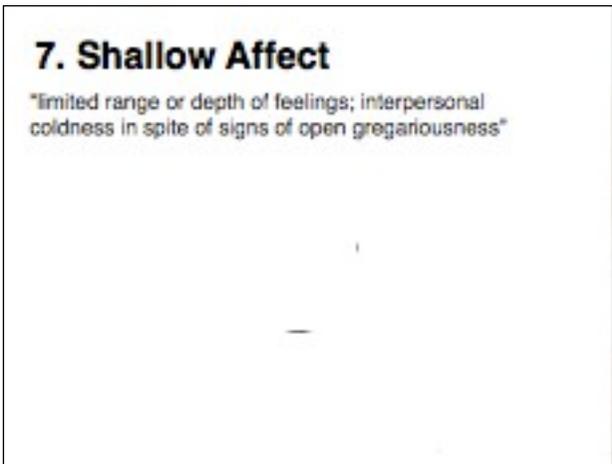
47

General Flow of talk

Affective Facet:

So what makes it possible for someone to con old people and not lose any sleep? These people are in effect, equipped to do it and feel NO guilt and have a long term lack of remorse (i.e. not realise how bad there actions were).

There's a subtle difference between Guilt and Empathy by the way. We'll cover empathy shortly. In a nutshell, it's possible to not feel guilt for running a red traffic signal, but you can still have empathy...



48

General Flow of talk

Affective Facet:

They have limited depth of feelings in spite of that superficial charm, Stoney emotionless. They know the words, but not the music

7. Shallow Affect

"limited range or depth of feelings; interpersonal coldness in spite of signs of open gregariousness"



49

General Flow of talk (Slide builds to this)

Affective Facet:

They have limited depth of feelings in spite of that superficial charm, Stoney emotionless.
They know the words, but not the music

Image licensed through : <http://www.cartoonstock.com/directory/p/psychopath.asp>

7. Shallow Affect

"I'm the most cold-blooded son of a bitch that you'll ever meet."

Ted Bundy

50

General Flow of talk

Affective Facet:

This quote was from Ted Bundy...may also load on item 2.

8. Callous / Lack of Empathy

A general lack of concern for the impact of their actions.

General Flow of talk

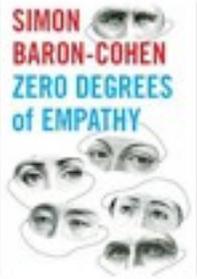
Affective Facet:

We touched on Empathy...there's also a general lack of concern for the impact of their behaviour.
Someone with empathy may still perform the action of leaving the dog, rationalising that
- someone will come by soon, it's better off in a good home - defence mechanism.
- Since they Understand how dog in feel and have no empathy, there's no guilt response triggered.

Someone scoring high on this item, won't think twice about it.

51

Zero Degrees of Empathy



52

General Flow of talk
Affective Facet:

To read more about Empathy, you might be interested in this book. The name might be familiar....it's Sasha Baron-Cohen's brother (Sasha Baron-Cohen being Borat, Ali G etc).

16. Failure to Accept Responsibility



**"I was pissed off.
He stepped into my space.
I did what I had to do"**

53

General Flow of talk
Affective Facet:

It's always someone else's fault, there's always a reason.... An example in relation to violent crime.. "It was their own fault, anyone could have seen I was in a rotten mood that night".

Some people wont accept responsibility through a defensive stance, e.g. "Someone shot my mother, so I shot him back." (affective rather than predatory). A psychopaths response can be both affective or predatory.

16. Failure to Accept Responsibility



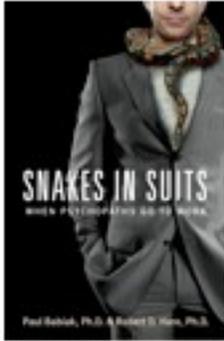
"an unfortunate series of events"

54

General Flow of talk
Affective Facet:

Some are now pointing to psychopathy in the corporate world, these are some recent quotes regarding interest rate fixing. This doesn't indicate a Psychopathy, but it's a good example of something a white-collar psychopath might say.

Snakes in Suits



55

General Flow of talk
Affective Facet:

if you're looking for a great book on Psychopathy in the corporate world, I really can't recommend this enough, it's by Babiak and Hare, two of the leading experts on Psychopathy.

3. Need for stimulation



56

General Flow of talk
Lifestyle Facet:

Psychopathy is also characterised by a need for stimulation / proneness to boredom, however risk taking doesn't necessarily transfer. For example, there's some evidence that skydivers make reasonably sound financial choices

3. Need for stimulation



57

General Flow of talk
Lifestyle Facet:

...but in psychopathy you'll see repeated risk taking, often including drug taking. When questioned, it's not unusual for criminals convicted of shooting to explain that "I was a bored and wanted to see what would happen". e.g. <http://www.kare11.com/news/article/938787/396/Man-charged-with-attempted-murder-for-shooting-at-Mpls-cops>

"After tracking Gomez's whereabouts investigators learned he was in the Hennepin County Jail on other charges. They say Gomez admitted to shooting 9 shots at the officers because he was bored and "wanted to fire the gun off."

9. Parasitic



58

General Flow of talk

Lifestyle Facet:

A parasitic lifestyle, finding people to con/manipulate into providing money, shelter, sex in return for nothing. When they've got what they wanted, they'll be off to the next victim. In a relationship context, partners are often left bewildered as to what went wrong.

9. Parasitic @ work



TEAMWORK
NOT FOR EVERYONE

59

General Flow of talk

Lifestyle Facet:

The parasitic lifestyle also applies in work. It's not uncommon for people scoring in this category to quickly identify people who will do some, if not all of their work for them. They are adept at hiding within the quick changing environments of most large corporates. Coupled with this is the ability for a narcissistic takeover (see Item 2), where they will claim success for projects they had little to no involvement in.

13. Lack of realistic long term goals



60

General Flow of talk

Lifestyle Facet:

The astronaut is a metaphor for other implausible goals, although it's worth noting that adept corporate psychopaths can often rise quite high, although their performance is typically poor (once found out). It's arguably more common to see unrealistic goals being cited by convicted Psychopaths, who, due to a criminal record would have limited chances to attain the positions they talk (convincingly) about.

14. Impulsivity

"I wasn't thinkin',
I wasn't plannin',
I was just doin'....",
Gary Gilmore

61

General Flow of talk Lifestyle Facet:

Somewhat related to the need for stimulation is impulsivity. Current thinking suggests that psychopaths are not great planners This quote is from Gary Gilmore -

"I did it because I felt like it", "it was a spur of the moment thing"

Regarding Texas murderer Gary Gilmore (Gained attention for legally pursuing his own execution. http://en.wikipedia.org/wiki/Gary_Gilmore

"If you hadn't been caught that night, do you think there would have been a third or fourth murder?"

"Until I got caught or shot to death by the police or something like that...I wasn't thinkin', I wasn't plannin', I was just doin'. It was a damned shame for those two guys...I'm just saying that murder vents rage. Rage is not reason. The murders were without reason. Don't try to understand murder using reason" excerpts from "Without a conscience", Robert D. Hare Guilford Press.

15. Irresponsibility

Repeated....

- Stoned at work
- Not honouring debts
- Fired
- DUI

62

General Flow of talk

Lifestyle Facet:

There's sustained and ongoing examples of irresponsibility.

At 44Con.com I suggested that I didn't know of any doctors like this. Someone in the audience suggested Dr. Harold Shipman. Another book "The Wisdom of Psychopaths" also indicates that some degree of psychopathic traits may be a distinct advantage in some professions, such as surgery.

10. Poor behaviour controls



RAGE

It's impossible to satisfy "Rage Boy" and his ilk. It's stupid to try.
- Christopher Hitchens

63

General Flow of talk

Lifestyle Facet:

This will be someone often described as short tempered or hot headed. "Rage Boy", although graphically depicting "rage", may not be a good analogy, as "rage boy" appears to be generally level headed see - <http://www.dailymail.co.uk/news/article-492864/The-surprising-truth-Rage-Boy-Americas-hated-poster-boy-Islamic-radicalism.html>

To meet the requirement of this category, hotheadedness would need to be a regular feature in someone's personality and not a photograph from a demonstration which can be taken out of context.

http://en.wikipedia.org/wiki/Islamic_Rage_Boy

12. Early Behaviour Problems

(ages 12 & under)



History of

- Truancy
- Robbery
- Vandalism
- Class disruption
- Bullying
- Cruelty to animals

64

General Flow of talk
Anti-social facet

Important to stress that the psychopath's history will be more extensive and serious when compared to peers in similar setting.

A high scorer told a story from when he was 10 or 11 when he shot "an irritating mutt" with a pellet gun. "I shot him in the ass and he cried and crawled around a while and died"

Another put a noose around the neck of a cat, tied the other end of the string to the top of a pole, and bat the cat around the pole with a tennis racket"

Both excerpts from "Without a conscience", Robert D. Hare Guilford Press.

18. Juvenile Delinquency cont..

(ages 13-17)

- Violence
- Fire-starting
- Murder or
- Attempted Murder
- Rape



65

General Flow of talk
Anti-social:

At this age bracket you'd expect to see formal contacts with the criminal justice system

19. Revocation of Conditional Release



66

General Flow of talk
Anti-social Facet:

You'd see a trend of skipping bail, escape from detention.

11. Promiscuous Sexual Behaviour

**GIRLS
DIRECT
TO YOUR ROOM
IN 20 MINUTES
1-877-8-I-BE-OVA**

67

General Flow of talk

Not loaded to any factor - They will typically have many encounters with prostitutes and one night stands.

17. Many short term relationships



"Playa: The skill of a playa is measured by the extent of his or her "game". The more "game" a playa has, the more respect they command in their community."

Source: Urban Dictionary

68

General Flow of talk

Not loaded to any factor. There will likely be a string of previous sexual relationships, often not lasting terribly long. They may well have a lot of children. There's some discussion about this being an evolutionary advantage.

20. Criminal Versatility

- Theft
- Robbery
- Drug offences
- Assault
- Murder / Manslaughter
- Possession of weapons
- Sexual offences
- Criminal negligence
- Fraud / Forgery
- Escape / Jump ball
- Kidnapping
- Arson
- Obstruction of justice
- Crimes against state
- Misc (Vandalism etc)

8 or more types = score 2 | 4 or 5 types = score 1

69

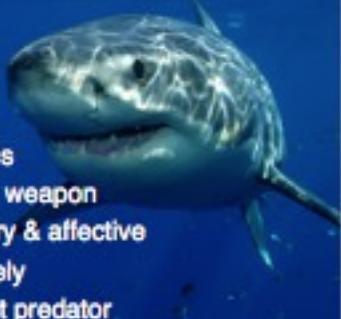
General Flow of talk

Anti-social

Criminal Versatility... Note, it's not about doing a lot of one crime, this is about have a broader portfolio of crimes. (4 or 5 types gets you a score of 1, while 6 or more gets you a score of 2).

Violence & Psychopathy

- More violent
- Commit more crimes
- More likely to use a weapon
- Violence is predatory & affective
- Stranger victims likely
- No guilt...the perfect predator



70

General Flow of Talk:

High scoring psychopaths are generally

- more violent
- commit more crimes
- more likely to use a weapon
- their violence is both predatory and affective
- Stranger victims are likely
- and they have no guilt.

...but as we saw earlier, not all psychopaths murder....

Treatment ?

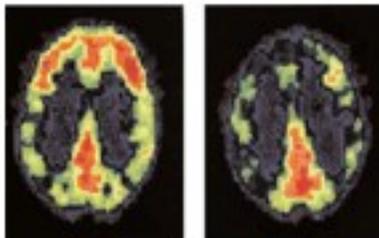


71

General Flow of talk

So, what can be done to treat these individuals? The generally held belief is that nothing can be done, that there is no treatment. However, there's some alternative positions emerging, which we'll quickly explore...

How "free" is our will?



Control

Murderer

72

General Flow of Talk:

A common statement regarding psychopaths is that they are born that way. Research of the brain suggested this is so, with much less activity in the pre-frontal cortex.

<http://www.dana.org/news/cerebrum/detail.aspx?id=3066>

Look also for BBC Horizon - Are you good or Evil?

<http://www.bbc.co.uk/programmes/b014kj65>

http://www.youtube.com/watch?v=u88lYs4FMTY&feature=list_other&playnext=1&list=SP4B8EC020E34E227F

and Jim Fallon TED talk - Jim Fallon: Exploring the mind of a killer

http://www.ted.com/talks/jim_fallon_exploring_the_mind_of_a_killer.html



73

General Flow of talk

But t's emerging that upbringing is what can determine what happens to someone with those brain related "birth gifts".

This quote was found on the Guardian website. The guardian is a British paper. It reads

"You might think, "This guy's being a total tosser," says Dave Owens, an officer for 19 years. 'But when he describes the abuse he had as a child from his father, the jigsaw fits.

What's incredibly interesting though is news from this week relating to a piece of research published in 2010 and then updated in 2011.

Source: <http://www.guardian.co.uk/society/2005/nov/20/prisonsandprobation.ukcrime>



74

General Flow of talk

"Without exception, these people have been injured in the company of their caregivers," she says. "And many of the descriptions made it clear that their later ruthlessness was an attempt to address this damage, but in an inappropriate or bad way."

"we are one step closer to treating the disorder. But this task is not an easy one. As health providers, it is deeply rooted in us that when some- one expresses pain, we must help."

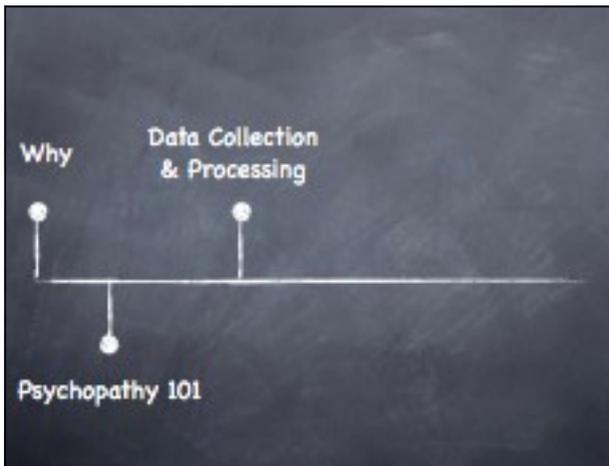
source: <http://www.sciencedaily.com/releases/2012/07/120713122925.htm>



75

General Flow of talk

Indeed, Simon Baron-Cohen indicated that during the formative years of a child's development, its parents can provide it with a pot of gold, by providing a stable, positive environment. Conversely, I imagine that this is where the fate of many "born psychopaths" is sealed.



76

General Flow of talk

Broadly speaking, there are two key camps in behavioural personality studies.

- Target v Rater (or observer), where a human examines, in this context, a Twitter feed, or
- Automated feature extraction, which pulls specific data points, often anonymising those points when/where the data is collected. i.e. rather than pulling a post, it tallies certain features, e.g. frequency of a certain word.

Based on current research on social network sites, Target - Rater generally performs better, possibly because humans are better at spotting low validity clues.



77

General Flow of talk

We created an application, first as a beta (not collecting data), then a live app (collecting data), then back to not collecting data because some people wanted to take the test, but we didn't need more participants.

The application collected a user id and answer to self assessment questions. This data was encrypted and temporarily stored outside the "doc root" folder. Periodically this data would be downloaded and removed from the server.



78

General Flow of talk

The app measured 8 personality traits in total, 3 of which we're the Dark Triad of narcissism, Machiavellianism and psychopathy. All of us have these traits to some degree.

The Dark Triad scale we used was something called the Short Dark Triad (SD3) by Prof Del Paulhus at the University of British Columbia. For psychopathy, the cross correlations between the SD3 and the SRP-III (Paulhus, Neumann & Hare) at .81**. SRP-III has a cross correlation of .626** with PCL-R. It's important to note that SD3 is therefore NOT a direct mapping to PCL-R.

Despite these limitations, still very useful to explore for research like ours, because it's short. In an online context, the longer the questionnaire, the higher the drop out rate.

We discuss these limitations in greater detail in our paper.



79

General Flow of talk

We then wrote software to query the twitter api, pulling features.

All done with open source software and on a HP Microserver. Cracking little box for the job

Over 3 million tweets, from 2927 people in 80 countries, although it's not really possible to tell if they really live in 80 countries

We looked at basic twitter attributes such as friends, followers, length of time on twitter, list membership etc

We also examined Linguistic Inquiry and Word Count (Or LIWC) of original tweet content, replies, retweets. Take a look at Professor James Pennebaker's work on LIWC for more information.

Note: It near impossible to determine whether twitter content is truly original or a cut&paste. For our study, we classified anything that wasn't a retweet or reply as "original content".



80

General Flow of talk

We used HP Proliant Microserver. Cheap, quiet (relatively), 4 drive bays

Ubuntu

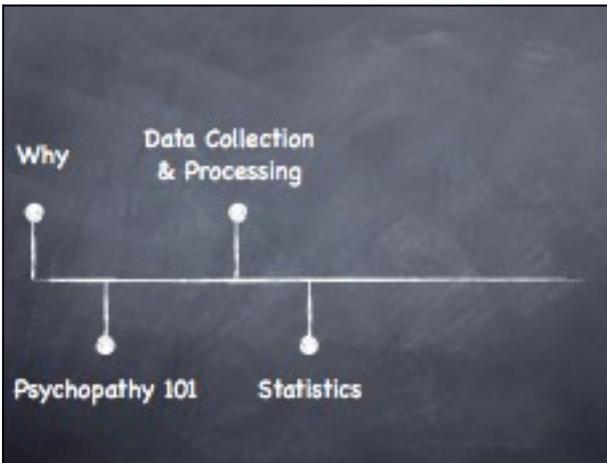
Python

CouchDB

MySQL

LIWC (liwc.net)

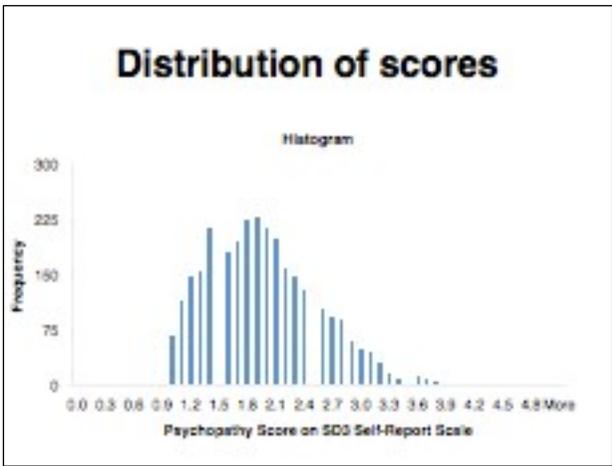
SofaStat, SPSS, XLStat, Excel, R



81

General Flow of talk

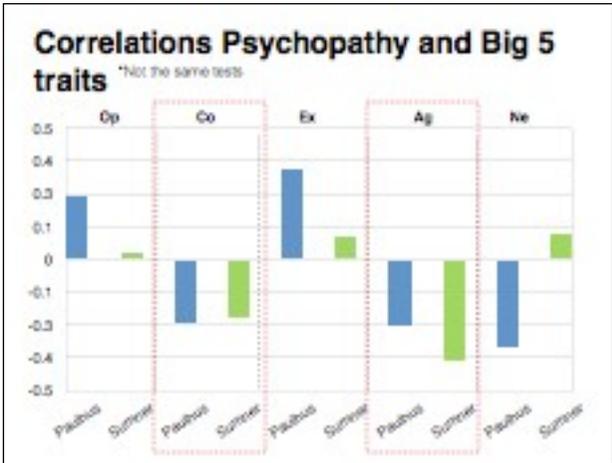
In the next section we look at some of the more interesting statistics.



82

General Flow of talk

First, the distribution and descriptive statistics were very similar to those noted in the creation of the original SD3 measure. Note the heavy skew to the lower end. (1 is actually the lowest score, not 0)
 Although not a perfect match, we see a similar curve to that in the GB Population study we spoke of earlier. Very few people score high. (note. different measures were used in the two studies mentioned)



83

General Flow of talk

Looking at the relationship between the psychopathy and the big 5 (note, the measures in the two papers were different, but it was one of the few comparisons between Big 5 and Psychopathy) we see some interesting similarities with Conscientiousness and Agreeableness. Variations may be explained by the different measures of both Psychopathy and Big 5 traits.

Note: These differences were also reflected in twitter activity. i.e. Where there was a significant positive correlation between some twitter activity and Psychopathy, we saw an almost inverse relationship to Agreeableness. The relationship between Psychopathy and agreeableness in this study still needs some investigation.

Attributes

| Attribute | Significance |
|----------------------|--------------|
| Number of Tweets* | 0.057* |
| Num Original Tweets | 0.050* |
| Followers per Friend | 0.044* |

N = 2,927
 ** p < 0.0001
 * p < 0.01

84

General Flow of talk

Full details are in the paper

klout

| Attribute | Significance |
|-------------|--------------|
| Klout Score | 0.051* |

N = 2,583
* p < 0.05

General Flow of talk
Full details are in the paper

85

linguistics (+ve correlations)

| Attribute | Examples | Significance |
|-------------|----------------------|--------------|
| Swearing | damn, piss, fuck | 0.187** |
| Anger | hate, kill, annoyed | 0.151** |
| Death | bury, coffin, kill | 0.094** |
| Neg Emotion | hurt, ugly, nasty | 0.083* |
| Filler | blah, imean, youknow | 0.073** |

N = 2614
** p < 0.0001
* p < 0.01

General Flow of talk
Full details are available in the paper

86

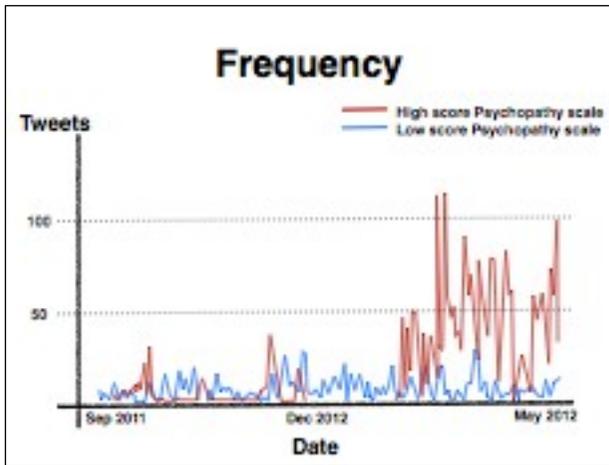
linguistics (-ve correlations)

| Attribute | Examples | Significance |
|-------------|-------------------------|--------------|
| Pos Emotion | love, nice, sweet | -.124** |
| Preps | to, with, above | -.086** |
| Family | daughter, husband, aunt | -.076** |
| Work | job, majors, xerox | -.077** |
| We | we, us, our | -.071** |

N = 2614
** p < 0.0001
* p < 0.01

General Flow of talk
Full details are in the paper

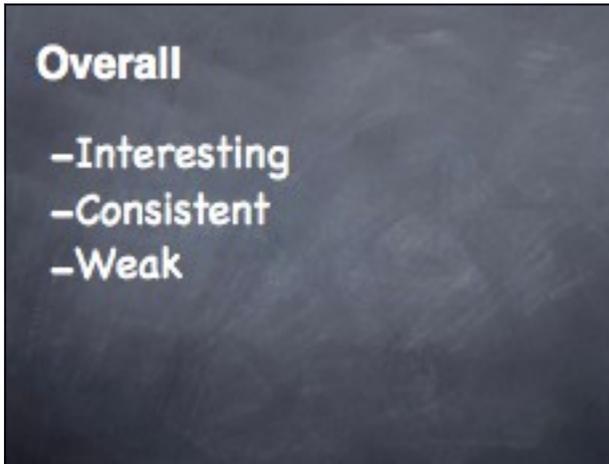
87



88

General Flow of talk

We also took an exploratory look at Frequency of tweeting. The blue line is representative of the lowest 10% scorers (avg), while the red line (the more jagged line) is representative of the top 10% scorers (avg). This could be associated with proneness to boredom and/or temper control, or nothing in particular. Nevertheless, the difference warrants more study.

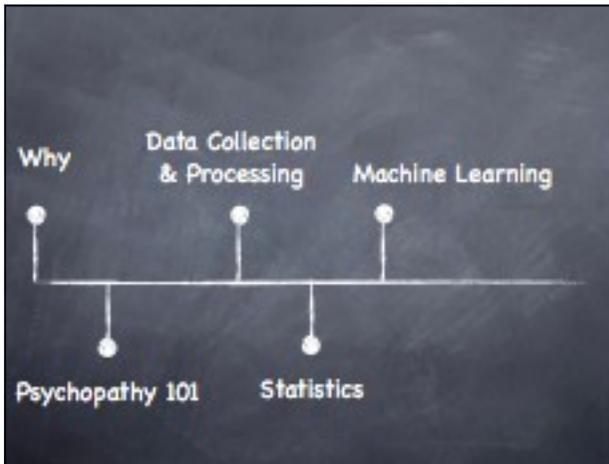


89

General Flow of talk

Overall, the results are interesting and largely comparable to the study's from Hancock et al and Boochever (<http://ecommons.library.cornell.edu/bitstream/1813/29536/2/Boochever,%20Rachel%20-%20Research%20Honors%20Thesis.pdf>)

However, in a practical sense, the relationships between psychopathy and twitter usages are weak.



90

General Flow of talk

After our previous study on Facebook and personality, Randall Wald from Florida Atlantic University asked whether we'd considered using machine learning against our data set.

Research over the past 12-18 months has highlighted that machine learning can be employed to predict personality.

For this reason, we teams up with Randall and FAU and also with Kaggle.com (a data science focused company, more of that later) to objectively examine to what extent if any, it was possible to predict a users SD3 score based on their twitter usage.

- Too much data to understand quickly
- Too many variables for easy statistical analysis
- Stakeholders want results that can improve decisions

Machine Learning

- Collect data into datasets

- Instances: samples (e.g., individual)
- Independent attributes (features)
- Dependent attributes (classes)

- Build models

- Models take features, predict classes
- Build using training data, validate with test data



91

General Flow of talk

Data mining / Machine learning.

A great book on Data Mining is this <http://www.amazon.com/Data-Mining-Practical-Techniques-Management/dp/0123748569>

What question are you trying to answer?

- Spotting the top n of psychopathic scorers?
- Splitting a group into high versus low?
- Predicting a psychopathy score?

- We used this experiment to try and answer all 3

92

General Flow of talk

Models are generally created to address a particular question or set of questions.

What question are you trying to answer?

- Spotting the top n of psychopathic scorers?
- Splitting a group into high versus low?
- Predicting a psychopathy score?

- We used this experiment to try and answer all 3

93

General Flow of talk

For this presentation, we focused on looking at the top n scores on the SD3 scale. In our paper, we also examine high/low median splits.

The question you're asking is very important, but so is the...



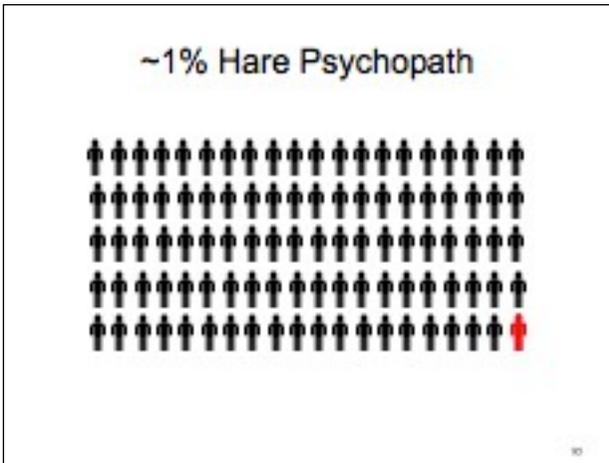
94

General Flow of talk
...way you evaluate the results.

As an example, one paper (Golbeck et al) posits that Facebook can predict personality within, just over 10% based on MAE figures.
You can have good performance (based on MAE) overall but still get the outliers pretty wrong. Generally there are inconsistent approaches to machine learning evaluation metrics in behavioural papers and results are therefore not easily digestible.

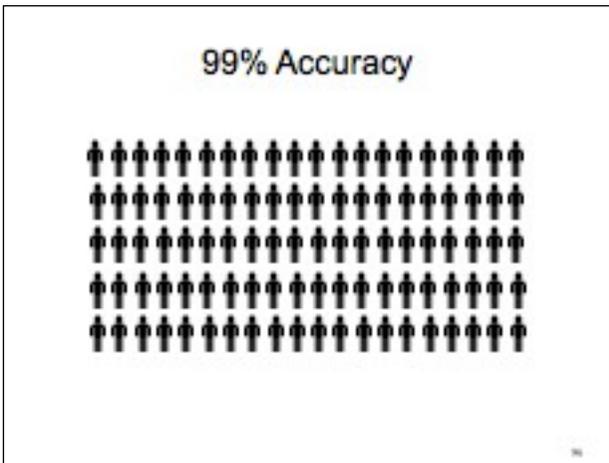
This is an important consideration in evaluating personality prediction. To demonstrate this problem....

Golbeck paper: <http://www.cs.umd.edu/~golbeck/pubs/Golbeck%20et%20al.%20-%202011%20-%20Predicting%20Personality%20from%20Twitter.pdf>



95

General Flow of talk
Lets look at trying to predict top 1%, for arguments sake, the highest scores of the psychopathy scale we used. The red person depicts our Psychopath



96

General Flow of talk
Simply predicting every instance (person) as a low-scorer will get you 99% accuracy. It's an impressive result for a gambler, but not much use if trying to identify high scoring psychopaths. This is why evaluation criteria play a very important part in articulating the practical performance of predictive models.

Confusion Matrix

| | | Predicted as | |
|--------|------|--------------|-----|
| | | High | Low |
| Actual | High | TP | FN |
| | Low | FP | TN |

97

General Flow of talk

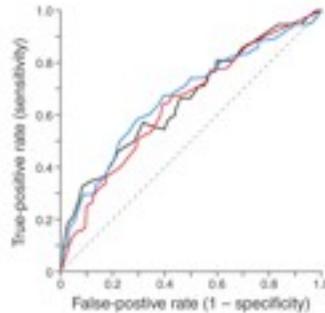
So one option is to look beyond accuracy and use the confusion matrix.

Psychopathy = "Positive"

- True Positives: positive instances labeled correctly
- False Positives: negative instances labeled incorrectly
- True Negatives, False Negatives: analogous

ROC Curve

- Curve plotting TPR vs. FPR
- AUC: Area under the ROC Curve



98

General Flow of talk

We can also use the ROC curve, which plots TPR against FPR.

The AUC is the area under the curve.

The grey, dotted line represents a random guess.

As the performance increase, the line will get nearer to the top left hand corner. It's a nice visual representation of the performance of your predictive models

Performance Metrics

- TPR: True Positives/All Positives
- TNR: True Negatives/All Negatives
- FPR: False Positives/All Negatives
- Arithmetic Mean of TPR and TNR
- Geometric Mean of TPR and TNR

99

General Flow of talk

To provide the most critical evaluation, we used a number of evaluation criteria to demonstrate the issues highlighted in the previous slide and to communicate how well our models work (from a practical sense, e.g. finding people in the top n%).



100

General Flow of talk

Got a data science problem? there's an app for that. Well, not an app exactly, Kaggle.com

What is Kaggle.com?

Kaggle host data science competitions, either public or private. Typically the competition organiser offers a price.

You can therefore crowd source your data science problem, **which is one of the things we did.**

We had over 100 teams create over 1000 models. We felt this would also provide the most critical evaluation of the problem. It should be noted that some competitors later actively engaged in helping us write up the results, for no reward other than pure interest.

If you don't have the skills "in house", I highly recommend checking out Kaggle.com

Results

| Evaluation | Kaggle Winner |
|------------|---------------|
| G-TPR | 0.640 |
| G-TNR | 0.664 |
| AUC | 0.678 |

101

General Flow of talk (Not shown in presentation)

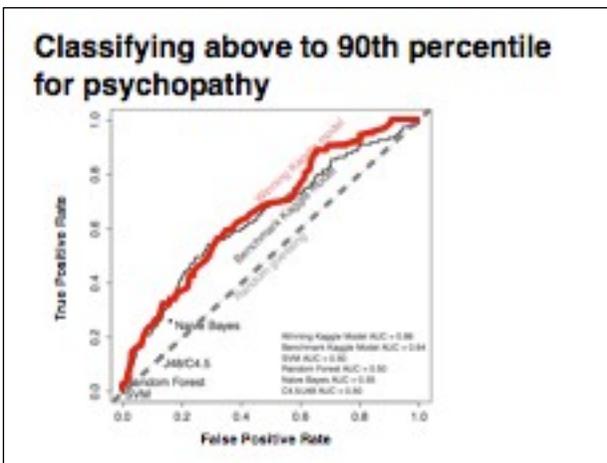
Notes on calculating G-Mean and A-Mean

Values of TP TN FP and FN are all a function of the threshold used to classify cases as positives or negatives. First, we chose the threshold that maximizes the A-mean, apply that threshold, and that generates TP, TN, FP, and FN. From that I can calculate the metrics. Then we chose the threshold that maximizes the G-mean and repeat the procedure.

"The Kaggle model scores each case with a value from 1.000 - 5.000 (roughly). I have a small function that will apply a threshold of 1.000, classify all cases below the threshold as negative and all equal to or above as positive. Then it's possible to count the TP, FP, TN, and FN and all the related metrics. Then it moves the threshold a tiny bit, like 1.001, and recalculates everything. Repeat this until the threshold has been stepped all the way to 5.000. Then you can look back and see which threshold values resulted in the maximum G-mean and A-mean - reapply those and calculate any other additional metrics needed."

$$A\text{-mean} = ((TP / (TP + FN)) + (TN / (TN + FP))) / 2$$

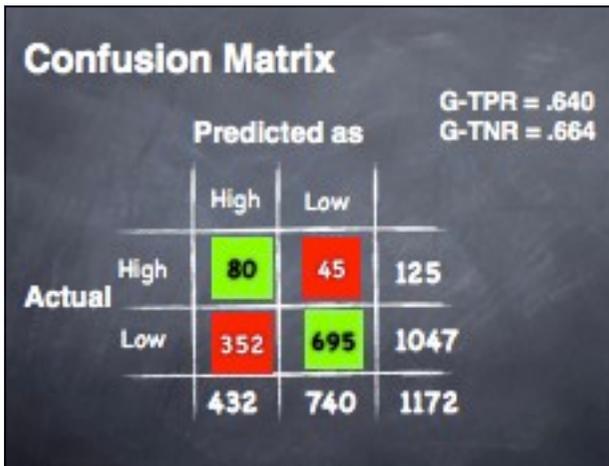
$$G\text{-mean} = \sqrt{((TP / (TP + FN)) * (TN / (TN + FP)))}$$



102

General Flow of talk

We can judge models by comparing them to random guessing. This plot (the ROC plot) compares several different models to random guessing. Models farther away from the diagonal random guessing line perform better than random guessing, and the two Kaggle models are always the best.



103

General Flow of talk

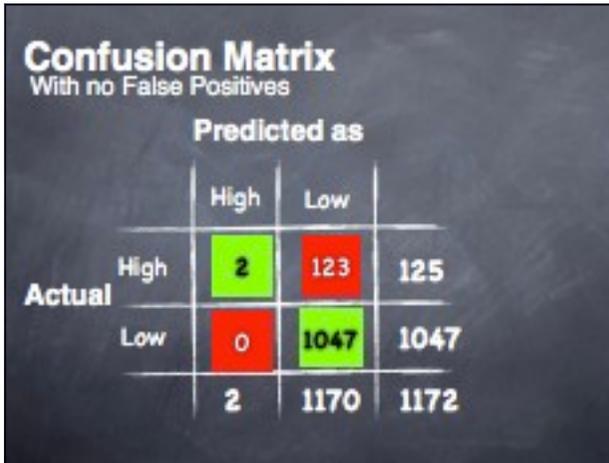
Good or bad? it depends on your goal. You can widen the net, but you'll increase false positives.

Notes on calculating G-Mean and A-Mean

Values of TP TN FP and FN are all a function of the threshold used to classify cases as positives or negatives. First, we chose the threshold that maximizes the A-mean, apply that threshold, and that generates TP, TN, FP, and FN. From that I can calculate the metrics. Then we chose the threshold that maximizes the G-mean and repeat the procedure.

"The Kaggle model scores each case with a value from 1.000 - 5.000 (roughly). I have a small function that will apply a threshold of 1.000, classify all cases below the threshold as negative and all equal to or above as positive. Then it's possible to count the TP, FP, TN, and FN and all the related metrics. Then it moves the threshold a tiny bit, like 1.001, and recalculates everything. Repeat this until the threshold has been stepped all the way to 5.000. Then you can look back and see which threshold values resulted in the maximum G-mean and A-mean - reapply those and calculate any other additional metrics needed."

$$A\text{-mean} = ((TP / (TP + FN)) + (TN / (TN + FP))) / 2 \quad G\text{-mean} = \sqrt{((TP / (TP + FN)) * (TN / (TN + FP)))}$$



104

General Flow of talk

This is what our confusion matrix looks like if we don't want to get False Positives



105

General Flow of talk

But what about looking at the highest scorers? The top 1.4%? (we choose 1.4% as it was ~3 SD's from the mean).

This is where Randall Wald and Florida Atlantic University came in.

Sampling

- Modify dataset to fix balance
- Random sampling
 - Discard/duplicate instances
- Random Under-sampling
 - Simple yet effective
 - Reduces dataset



106

General Flow of talk
No notes in this section

Ensembles

- Combine multiple weaker models
- Help when individual models don't perform well
- How to introduce diversity of models?



107

General Flow of talk

RUSBoost

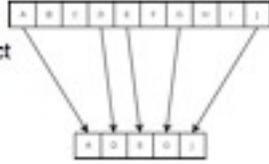
- Combines Random Undersampling and Boosting (a form of ensemble)
- Each round will undersample data, then build model
- Instances which are misclassified are given more weight when building future models
- All models combined at the end (weighted based on their accuracy)

108

General Flow of talk

Feature Selection

- Choose best subset of features to use for building model
- Rank features and select top N
- Reduce dataset size and improve performance

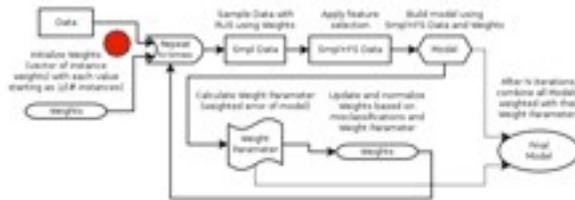


General Flow of talk

109

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

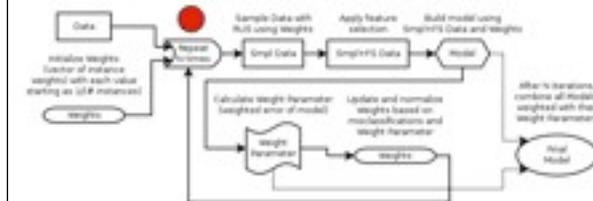


General Flow of talk

110

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm



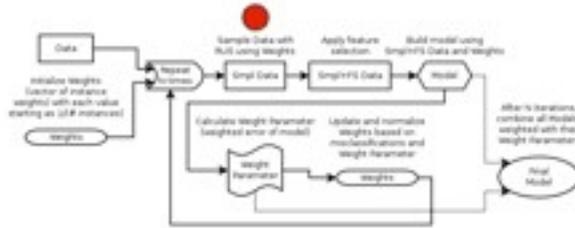
General Flow of talk

111

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

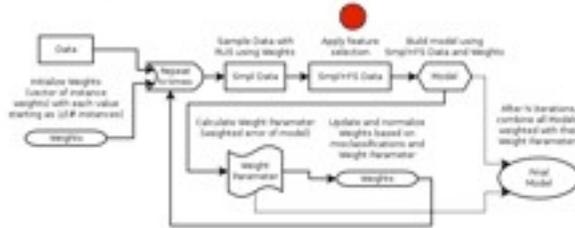


112

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

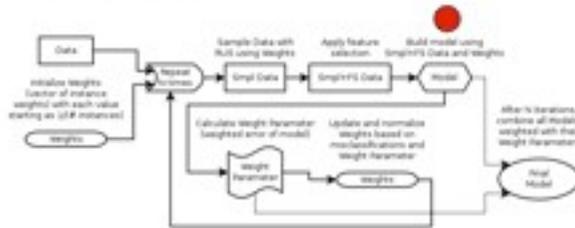


113

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

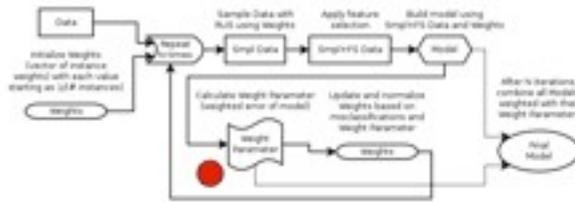


114

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

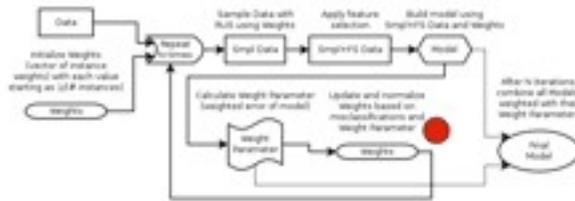


115

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm

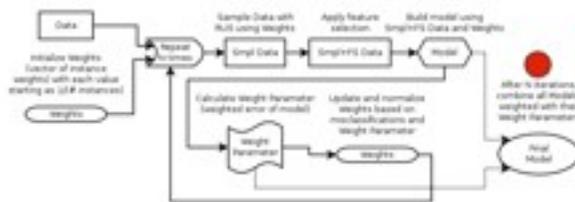


116

General Flow of talk

SelectRUSBoost

- Same as RUSBoost, but adds feature selection into the algorithm



117

Results

- With six of seven learning algorithms, RUSBoost was improving over no boosting
- SelectRUSBoost improved performance for all learners
- Statistical analysis proves that either boosting approach is better than no boosting
- Discussion of best learners/rankers omitted
- Best results: TPR of .707, TNR of 0.719, AUC 0.746

118

Results

| Evaluation | Kaggle Winner | SelectRUSBoost |
|------------|---------------|----------------|
| G-TPR | 0.640 | 0.707 |
| G-TNR | 0.664 | 0.719 |
| AUC | 0.678 | 0.746 |

119

119

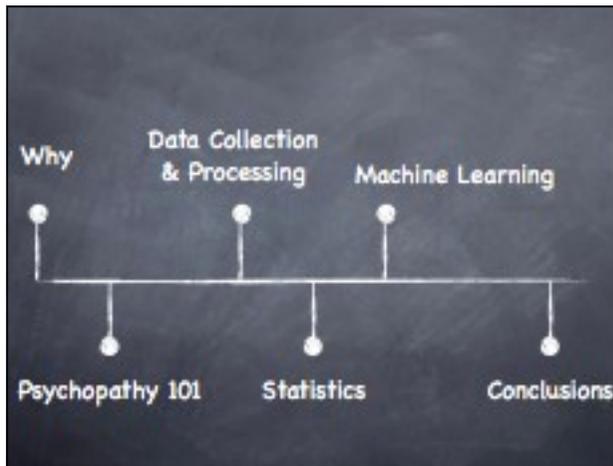
Confusion Matrix (top 1.4%)

| | | Predicted as | | |
|--------|------|--------------|------|------|
| | | High | Low | |
| Actual | High | 29 | 12 | 41 |
| | Low | 808 | 2067 | 2875 |
| | | 837 | 2079 | 2916 |

G-TPR = .707
G-TNR = .719

Good or bad? Clearly a better TPR and TNR, but look at the False Positives. Do you care about the false positive?

120



General Flow of talk
So to some conclusions

121

Conclusions

- Beware evaluation criteria
- Poor when applied to individuals
- Personality prediction improves when more than one measure is used
- ...but could combine with closed source data...

General Flow of talk

For spotting people who are potentially higher on the psychopathy scale, there is a performance gain over a random guess, but as we saw, you have a large problem in the shape of False Positives. Further, again Psychopathy does not necessarily equal criminal.

That said, models such as these, could be combined with closed source data (such as criminal records etc) to further home in on people (we're not suggesting that's a good thing or bad thing).

It's also been shown that combining personality measures can increase the predictability of personality traits, so it's not inconceivable that predictive improvements are possible when other factors are taken into account.

So, should an individual/organisation wish to highlight (or eliminate from a hiring process) people scoring higher on a psychopathy scale, models such as these **may** end up playing a part. (There are very important questions on ethics, privacy etc which we discuss shortly).

Used alone, these models are, in our opinion far to error prone to be useful for identifying individuals, although it's not say future models performance will not improve.

122

Conclusions

- Maybe useful for observing changing levels of anti-social personality traits
- Depending on your view point, that's good or bad.

Below the text is a row of five gummy bears: two red, one yellow, and two more red.

General Flow of talk

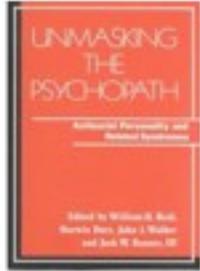
A more practical use may be in observing anti-social traits in society through social media. i.e. Are anti-social traits increasing/ decreasing and is that a good or a bad thing? Are some groups/geographies more anti-social than others?

Depending on your view point, this may be a good idea or a terrible idea.

Note: Far more research would be required even for use in this context and again, there are important issues regarding ethics and privacy. Just because it MAY be possible, doesn't mean we should.

123

Unmasking the psychopath



Chapter 4

Social Crises and Psychopathy: Toward a Sociology of the psychopath

José Sánchez

General Flow of talk

Societal trends for Anti Social personality are interesting and not new... In chapter 4 of this book, the author talks about the impact of industrialisation on the rise in psychopathy. We also see theories (published) emerging which associate the global financial crisis with rising levels of psychopathic traits in finance.

Clive Boddy - The Corporate Psychopaths Theory of the Global Financial Crisis

124



General Flow of talk

Could insights from research in this field be useful to understand more about internet trolls/cyberbullies ? Maybe. It's certainly a problem that will impact the next generation, many of which will live much of their lives in public.

125



General Flow of talk

As with Hancock et al, understanding more about linguistic differences may help us understand a little more about people in relation to psychopathy and indeed a variety of traits. Far more research is required in this area, but currently research into the language and psychopathy appears to be pointing in the same direction.

Can models such as these really spot a psychopath (scoring over say 25 on PCL-R) ? From what we can tell using the SD3 scale, the answer as yet is **no**

126

Limitations

- Self-Report
- Not a widely used assessment
- Selection bias?
- Linguistic Analysis against Tweets

General Flow of talk

There are a number of limitations:

- It's a self-assessment, so there's always a chance that x% lied.
- The self assessments have not been widely used in previous research.
- The self assessment is not designed to detect "clinical psychopathy"
- Most participants were recruited after a tweet from Stephen Fry, possibly creating a selection bias.
- The use of LIWC against tweeting is relatively unexplored, although it has been used before.

We focus more on these in the accompanying paper

127



General Flow of talk

Ultimately we're here (the tip of the iceberg), but research is all pointing in the same direction in terms of showing linguistic correlations.

The point is really that social media based personality research is in its infancy. Almost no one has looked at the more anti-social aspects on our personality and social media, but a growing number have looked at other traits.

A lot more research is required before any of these claims can really be supported fully, even with the Facebook studies (in our opinion). If people are concerned of false positives with this particular study (as they should be), they should also be critically examining the validity and reliability of reviewing a potential employees social media profile or NetRep.

128



General Flow of talk

We therefore need to be asking, should we and in what (if any circumstances) attempt to predict personality and behaviour? We're not aware of anyone having initiated that debate or study yet, but its clear that such debate is necessary.

129



General Flow of talk

www.onlineprivacyfoundation.org

Comments/feedback to - [chris at onlineprivacyfoundation.org](mailto:chris@onlineprivacyfoundation.org)

NOTE: A further paper focusing entirely on machine prediction is planned for release in December 2012 at the 11th International Conference on Machine Learning and Applications. <http://www.icmla-conference.org/icmla12/>

130