Social Authentication: Harder than it Looks

This appears to be:

- Hyoungshick Kim
- John Tang
- Ross Anderson

How personal is this knowledge?

Social Authentication on Facebook

- Facebook began using additional measures to authenticate users in novel locations
- If you usually log in from London, but the system sees someone trying to log in to your account from Cape Town, it will show you a few pictures of your friends and ask you to name a selected person in each photo
- Facebook called this feature “social authentication”

An Example

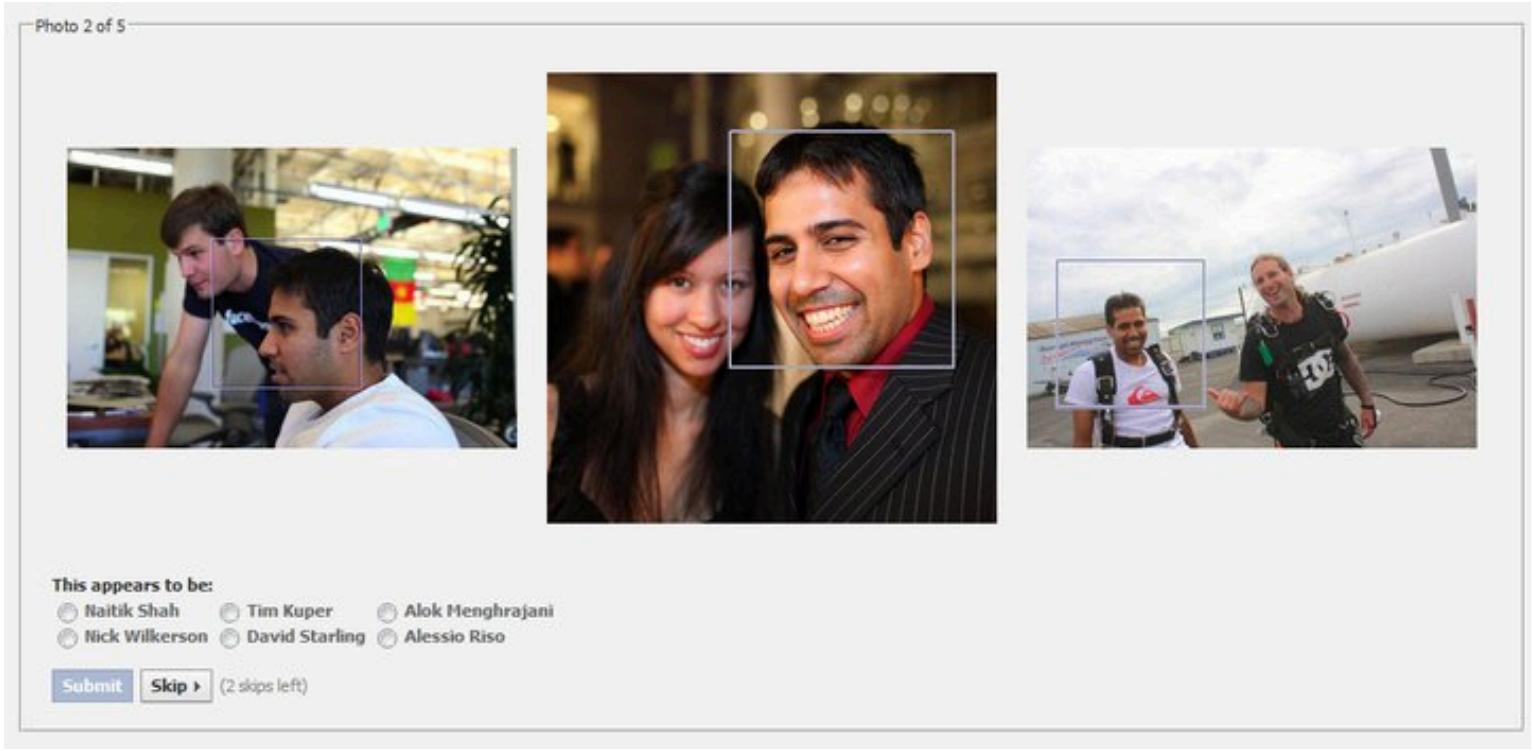
Photo 2 of 5
This appears to be:
Naitik Shah
Tim Kuper
Alok Menghrajani
Nick Wilkerson
David Starling
Alessio Riso
Submit
Skip
(2 skips left)

Main Observations (1)

- We set out to formally quantify the guessing probability through quantitative analysis of real social network structures
- We found that being able to recognise friends is not in general enough for authentication if the threat model includes other friends
- Community-based challenge selection can significantly reduce the insider threat; when a user's friends are divided into well-separated communities, we can select one or more recognition subjects from each.

I Know Him!

But so do many other people.

Friends or frenemies?

- If you're doing something embarrassing, then from whom do you need privacy?
- If you're a celeb, everyone – but the rest of us only have to worry about a few hundred friends
- So: if someone who can recognise a random subset of *k* of my friends can attack me, to whom am I vulnerable?
- We calculate the attack possibility from such users (your friends, or friends of friends)

Attack Advantage of Impersonation

Given *k* challenge images of friends chosen at random, the impersonation attack probability for user u can be calculated as:

$$\mathbf{Adv}_{\mathcal{R}}(u,k,\rho) \geq \max_{a \in A_u} \left\{ \prod_{i=1}^{\min\{k,|f_u|\}} \frac{|f_{ua}| - (i-1)}{|f_u| - (i-1)} \cdot \rho \right\}$$

where f_{ua} is the intersection of f_u and $\{f_a \cup a\}$

A_u is the set of users who share mutual friends with u.

Real Datasets

Table 1. Summary of datasets used. $\langle d \rangle$ and n_{cc} represent the "average number of friends" and the "number of connected components", respectively. The sub-networks of universities are highly connected compared to those of regions.

Network	Type	$\lvert U \rvert$	$\lvert E \rvert$	$\langle d \rangle$	n_{cc}
Columbia	University	15,441	620,075	80.32	16
Harvard	University	18,273	1,061,722	116.21	22
Stanford	University	15,043	944,846	125.62	18
Yale	University	10,456	634,529	121.37	4
Monterey Bay	Region	26,701	251,249	18.82	1
Russia	Region	116,987	429,589	7.34	3
Santa Barbara (SB)	Region	43,539	632,158	29.04	1

We display histograms of the vulnerability of users in each sub-network.

Histogram of Attack Advantage

When the number of challenge images is 1,

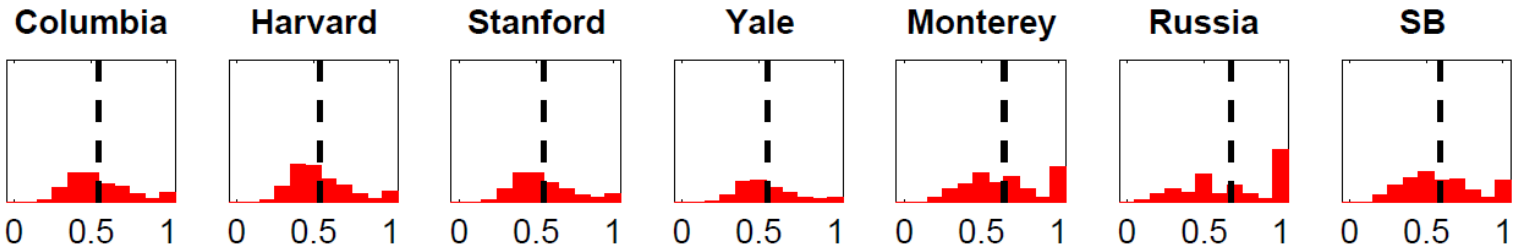

many people are vulnerable to impersonation.

Even for 5 challenge images,

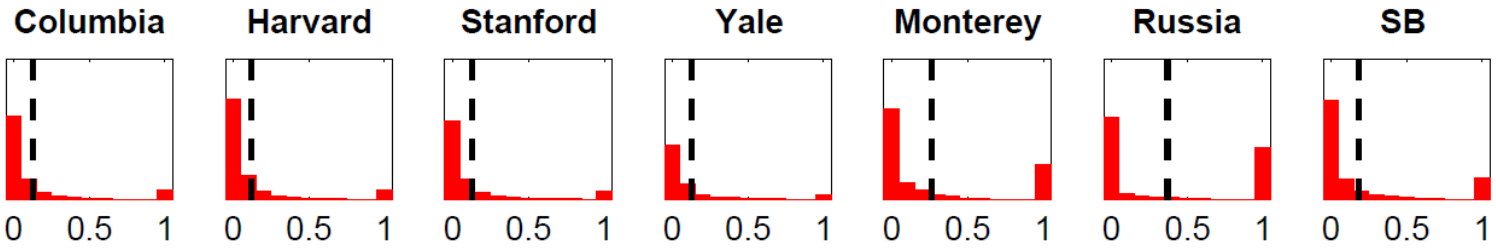

some people can be impersonated with probability 100%.

Who is the most vulnerable?

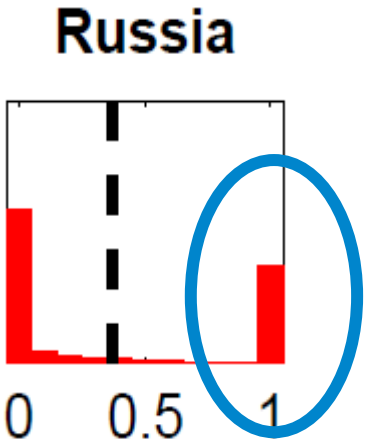

Some people can still be impersonated with probability 100%. Who?

Social authentication is not effective for users with only a few friends

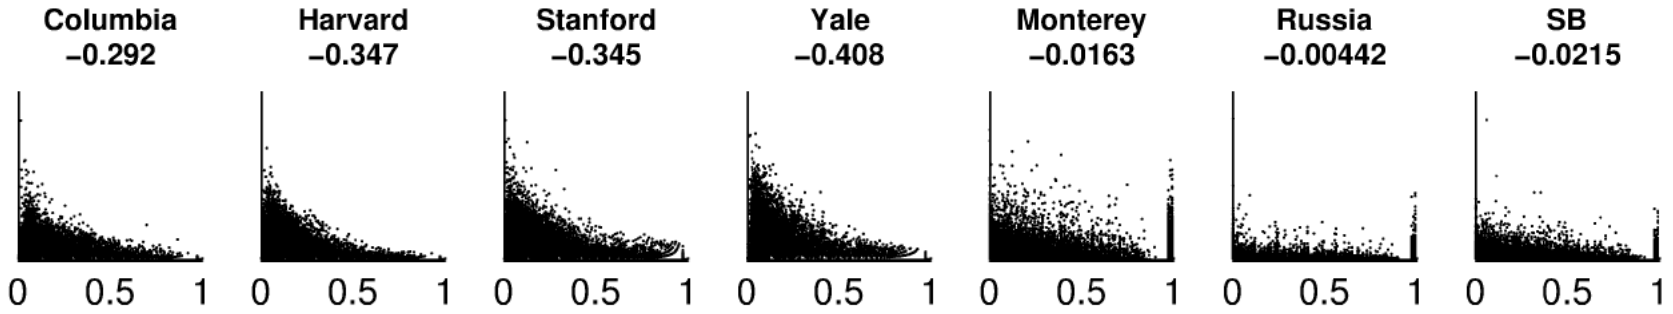

Correlation between number of friends and attack advantage

Social authentication is not effective for users with a high clustering coefficient

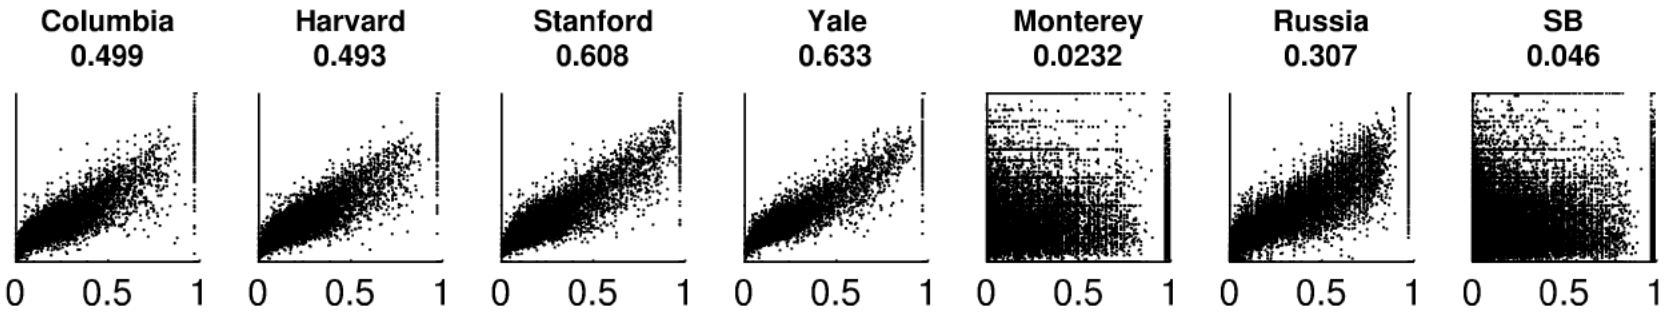

Clustering coefficients vs attack advantage

The clustering coefficient of node *u* measures the probability that its neighbours are each others' neighbours too

Community-based selection is better

If user *u*'s friends split into two communities, we can cut the risk by selecting friends' photos from different groups.

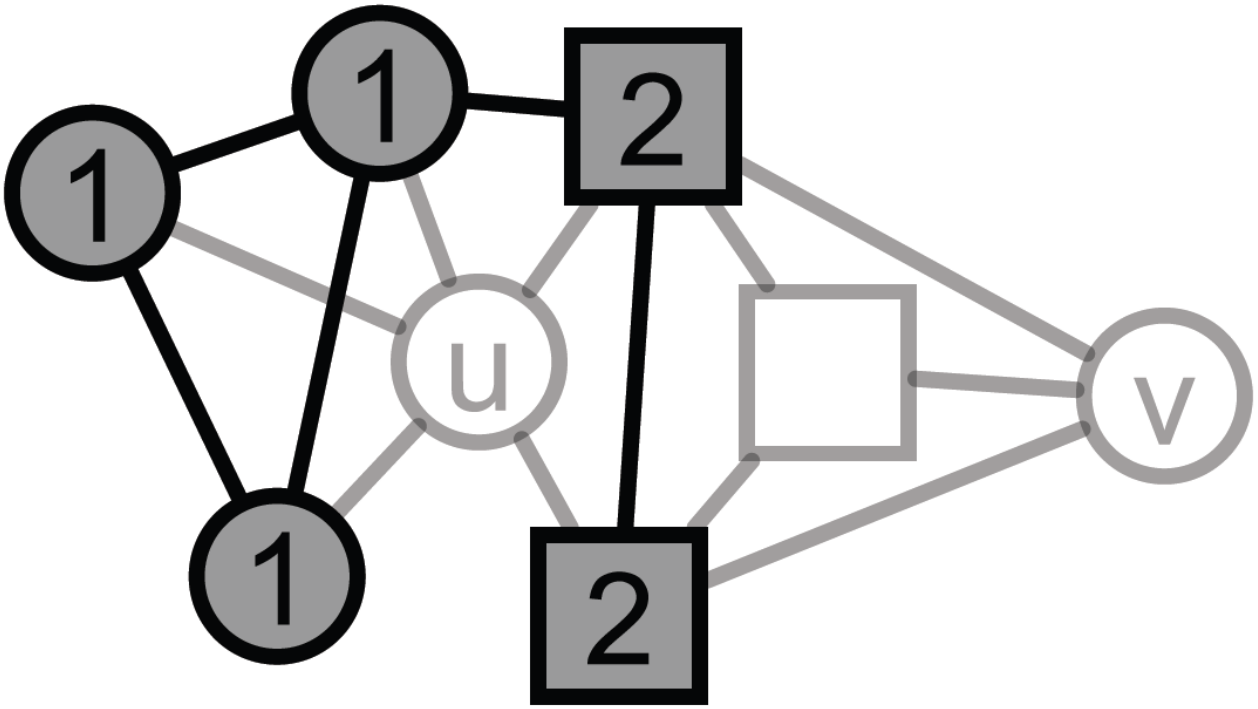

With 3 challenge images

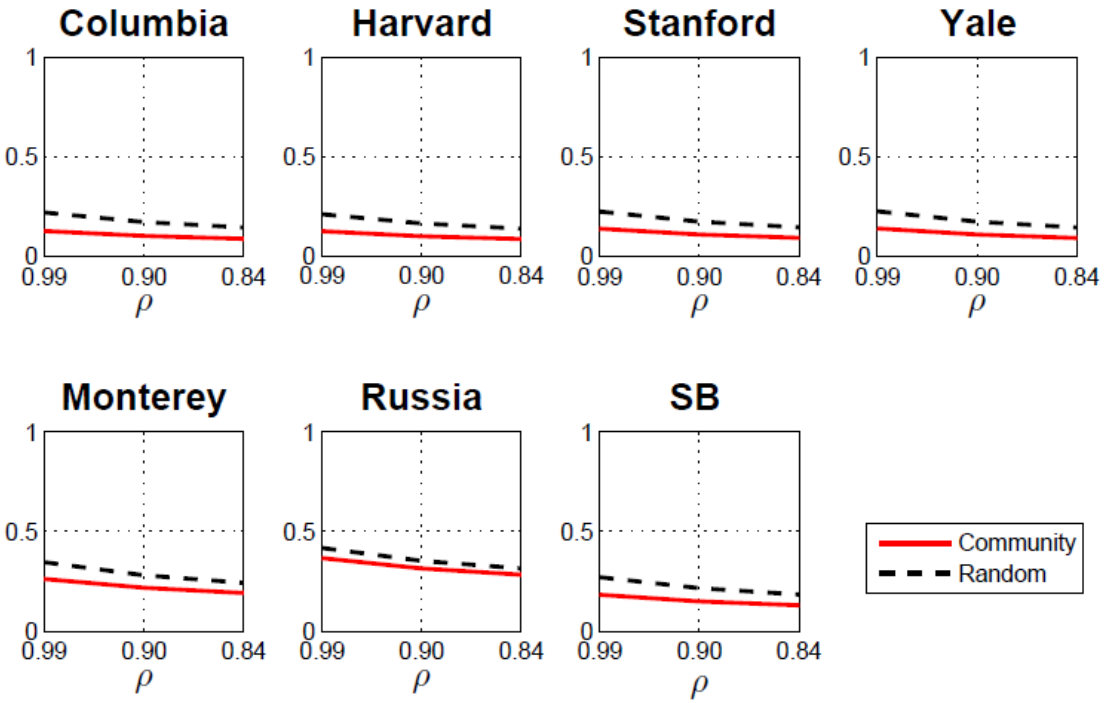

Table 2. The average number of communities for each user's friends.

Columbia	Harvard	Stanford	Yale	Monterey	Russia	Santa
3.779	3.371	3.227	2.812	3.690	3.099	4.980

Main Observations (2)

- Facebook's social authentication is an extension of the idea of CAPTCHAs. So it shares their problems
- Many users display tagged photos, and Facebook provides APIs to get images with Facebook ID
- The best performing face-recognition algorithms achieve about 65% accuracy using 60,000 facial images of 500 users
- Acquisti et al. did an attack using a larger database of images taken from Facebook profiles only, across the CMU campus (accuracy was about one third)

Current selection criteria

- Facebook used to use any pictures on your friends' albums
- Recently they have started screening photos with face detection software to improve usability
- For the same reason, Facebook selects friends who communicate frequently with the user they wish to authenticate

Remaining usability issues…

Bad Example (1)

Bad Example (2)

Discussion with Facebook

- After this paper was accepted, Facebook's security team got a copy
- Claimed: they knew it was weak against your jilted former lover; and you can log in easily from friends' machines as a matter of policy
- Argued: local police and courts are the proper remedy for the 'insider' threat
- Also: sure, anyone can use it for targeted attacks (not seen much – Indonesian attacks on casinos)
- What this system did was to kill industrial scale phishing, which used to be a bother. Spammers now use malware instead

Conclusion

- Facebook implemented a new security system based on social CAPTCHAs for people who log in from remote machines
- This may have provided some reassurance of privacy to ordinary users like us…
- But it's not doing security for me – it's doing security for them
- As service firms get ever larger, is this the way of the future?