Human Factors in the Security of Online and Mobile Systems

Gang Wang
Assistant Professor
Department of Computer Science
Virginia Tech

A Bit of Background: Gang Wang

- **Assistant Professor of Computer Science**
  - Ph.D. from UC Santa Barbara (2016)
  - B.E. from Tsinghua University (2010)
- **Research interests**
  - Security and Privacy
  - Data Mining
  - Human Computer Interactions
- Outstanding Dissertation Award (UCSB)
- Best Practical Paper Award (SIGMETRICS'13)
- Research at Microsoft Research and LinkedIn (2011, 2012, 2014)
- Press coverage: *MIT Technology Review*, *Fusion*, *Boston Globe, etc.*

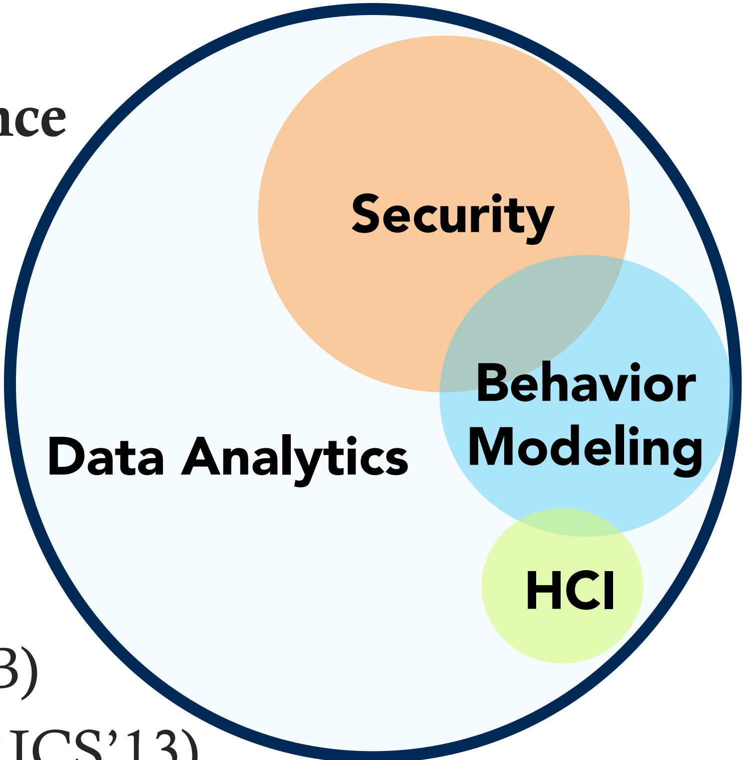

Looking for bright PhD/MS students to work with me!

Humans: The Weakest Link

- Data breaches caused by **human factors**
  - Anthem: largest breach in 2015
  - 80 Million records leaked (SSN, name, birthday)

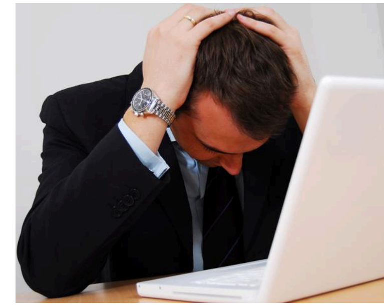

Employee revealed password to attacker

Anthem BlueCross BlueShield

UC University Health Insurance

Victim

- A growing concern
  - More recently: MySpace leaked 400 Million passwords (May 2016)
  - 1564 breaches, 1.5 Billion records leaked (2014 - 2015)
  - 95% security incidents involved human factors [1]

myspace JP Morgan COSTCO WHOLESALE CVS pharmacy STAPLES TARGET Hilton HOTELS & RESORTS

[1] IBM Security Services 2014 Cyber Security Intelligence Index

Attacks Targeting Users Now Common

- Malicious content target human users daily
  - Massive email/social spam, scam
  - Targeted spear phishing, like this one:

LinkedIn

Hi Gang,

I am a recruiter here with Amazon Data Science in Ireland. I am hoping to talk to you about a Systems Engineering role which I am hiring for at the moment.

This position is based on our data science team here in Dublin, Ireland and offers a competitive compensation plan, as well as a fantastic opportunity for continuous career growth and professional development in a challenging work enviror be a good match :)

Shortened URL to a phishing site
http://amazen.xxxx.com

Please find at the link below some informati role itself and please let me know if you would considering applying. http://tinyurl.com/qxadbqf

Reply   Not Interested

Understanding Human Factors

- Key questions
  - What are human's roles in online attacks?
  - How to understand user behaviors in online systems?
  - How to leverage this understanding to improve security?

- Traditional user study has limitations
  - Interview/survey: trade breadth with depth
  - High costs, does not scale

Need a scalable approach to study human factors in security

Potential solution: leverage detailed data on user behavior!

User Interaction

User generated content

Web clicks

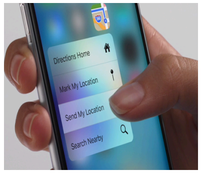

Mobility

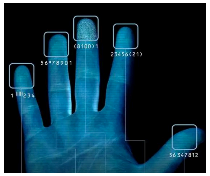

Biometrics

Data-driven Approach to Improving Online Security Through Users

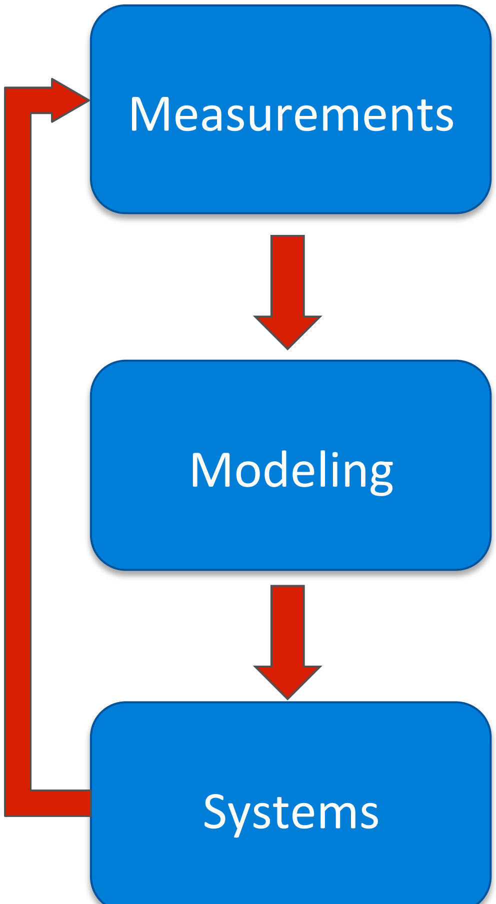

Understanding online users and attacks

- Large-scale online/mobile communities
- Spam, malware, malicious crowdsourcing

[WWW'12] [DSN'13] [IMC'13] [IMC'14] [SIGMETRICS'13] [ICWSM'16] [MobiSys'16] [IMC'16]

User behavior modeling

- Graph-based behavior models
- Clickstream based behavior models

[MobiCom'11] [WWW'13] [CSCW'15] [CHI'16]

This Talk

Data-driven security systems

- Behavior based security defenses
- Adversarial machine learning

[NDSS'13] [USENIX Security'13] [USENIX Security'14]

Talk Outline

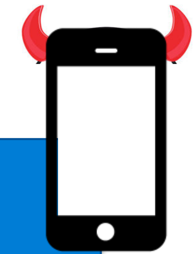

1. Emerging Threat of Sybil Devices

- Simulated mobile devices pretending to be real users
- Manipulate online services at a large-scale
- Example attacks: location tracking on Waze

[MobiSys'16]

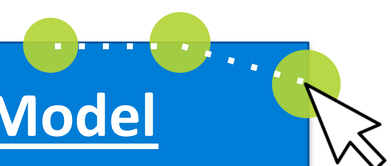

2. Clickstream based User Behavior Model

- Build hierarchy of behavior clusters
- Automatically extract key distinguishing features
- Detect fake accounts, track dynamic behavior changes

[CHI'16]

Mobile Phone = Your Identity?

- Mobile phones for content, payment, authentication

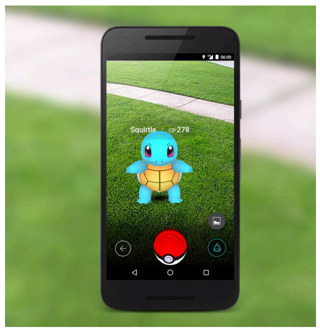


- Mobile devices are virtual representations of ourselves.

=

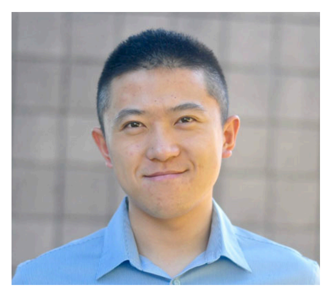

But Is This a Safe Assumption?

- An app user = 1 real phone + 1 real person

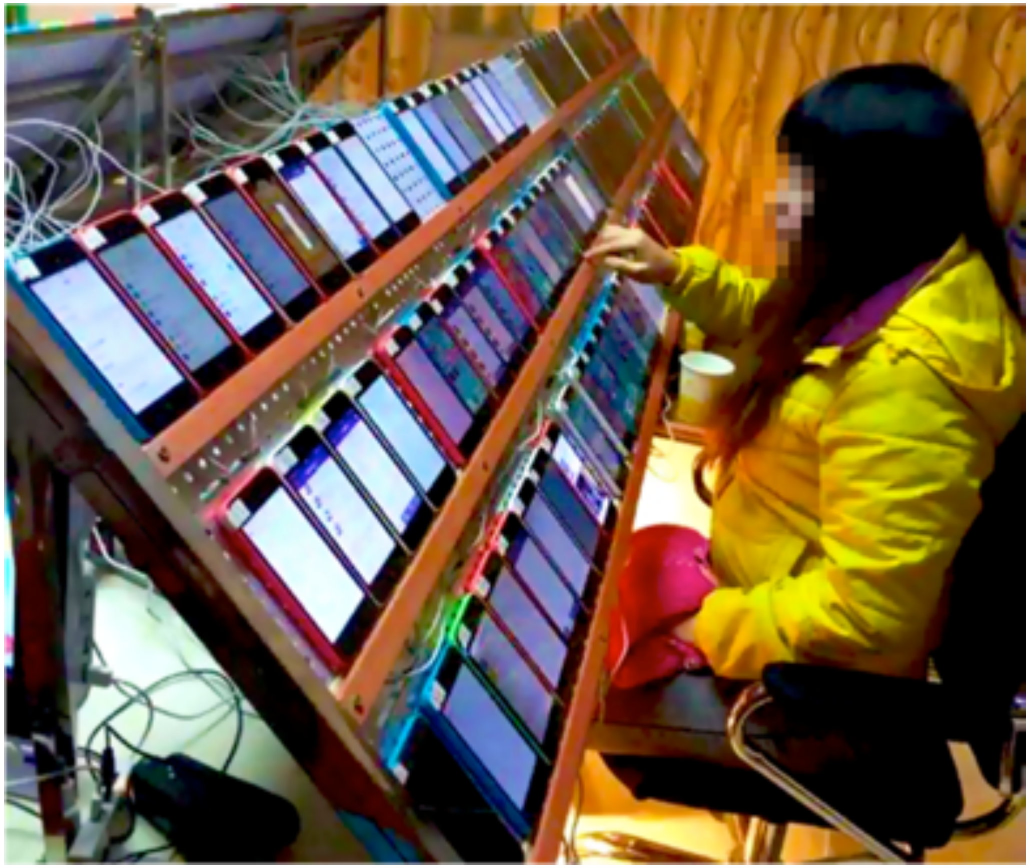

Can We "Authenticate" Devices?

- Register via email account
- Require CAPTCHAs
- 2FA via phone number
- Validate IMEI number

- Create fake email account
- Out-source to third party
- Temporary SMS services
- Spoofed IMEI

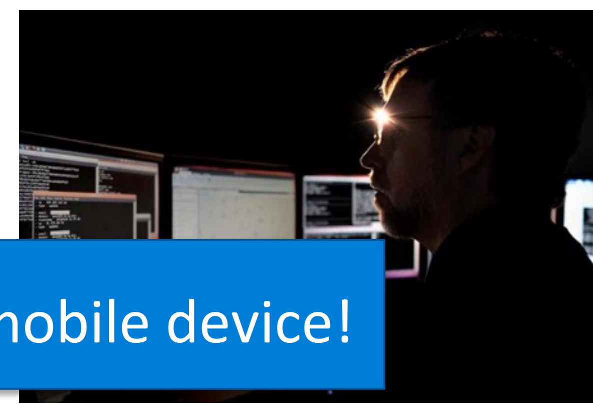

Highly challenging to authenticate a mobile device!

Threat of Sybil Devices

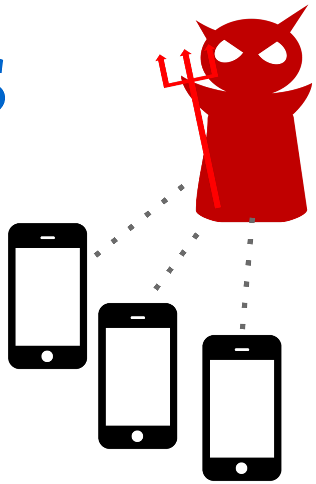

- Sybil devices
  - Software scripts emulating as real devices
  - Allowing a single user to control many devices

- In the context of Waze (popular navigation app)
  - Creating a large number of Sybil devices with low costs
  - Attacks: injecting fake events, user location tracking
  - Generalizable to other mobile communities

waze

waze Key Features

- 50M active users
- Real-time traffic update using millions of users' locations

User reported events
- Accidents, police trap, etc.
- Alert users of nearby events

Social features
- See nearby users on the map
- Say "hi"/msg nearby users

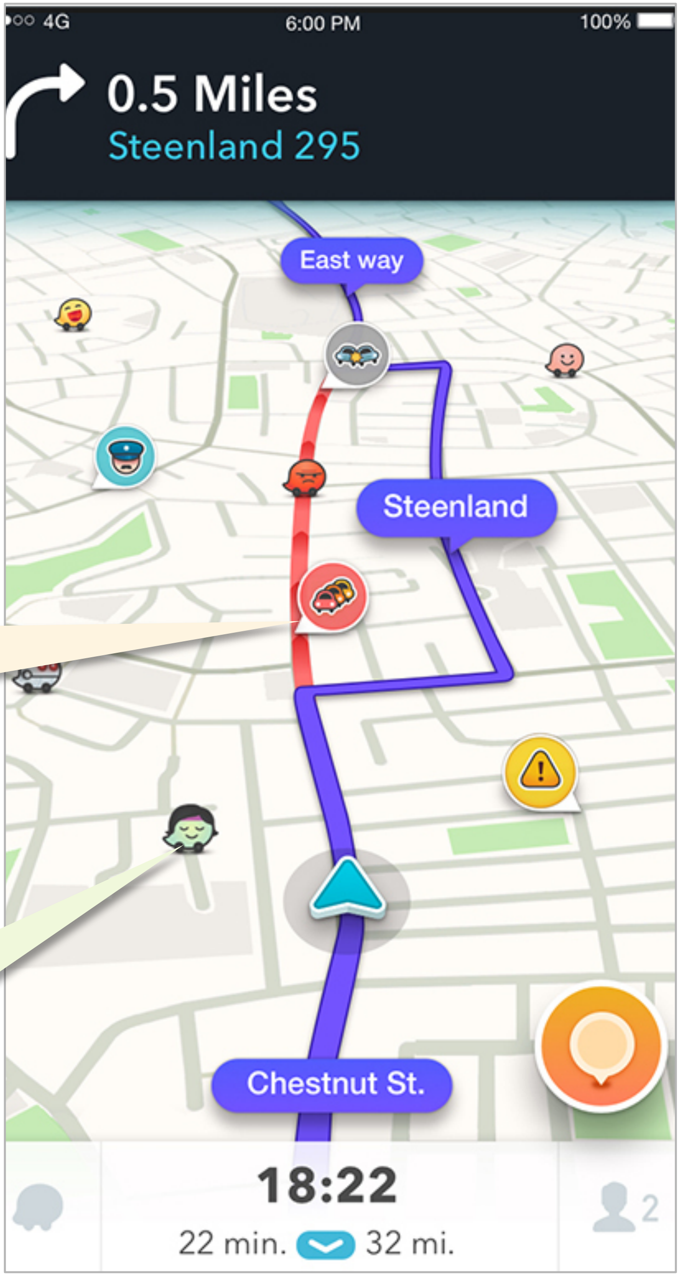

Creating Sybil Devices

- Naïve approach: mobile emulators
  - Not scalable: ~10 emulators per PC
- Our way: emulate a mobile client using scripts
  - Server communicates with client via limited APIs
  - Mimic API calls to replace full client

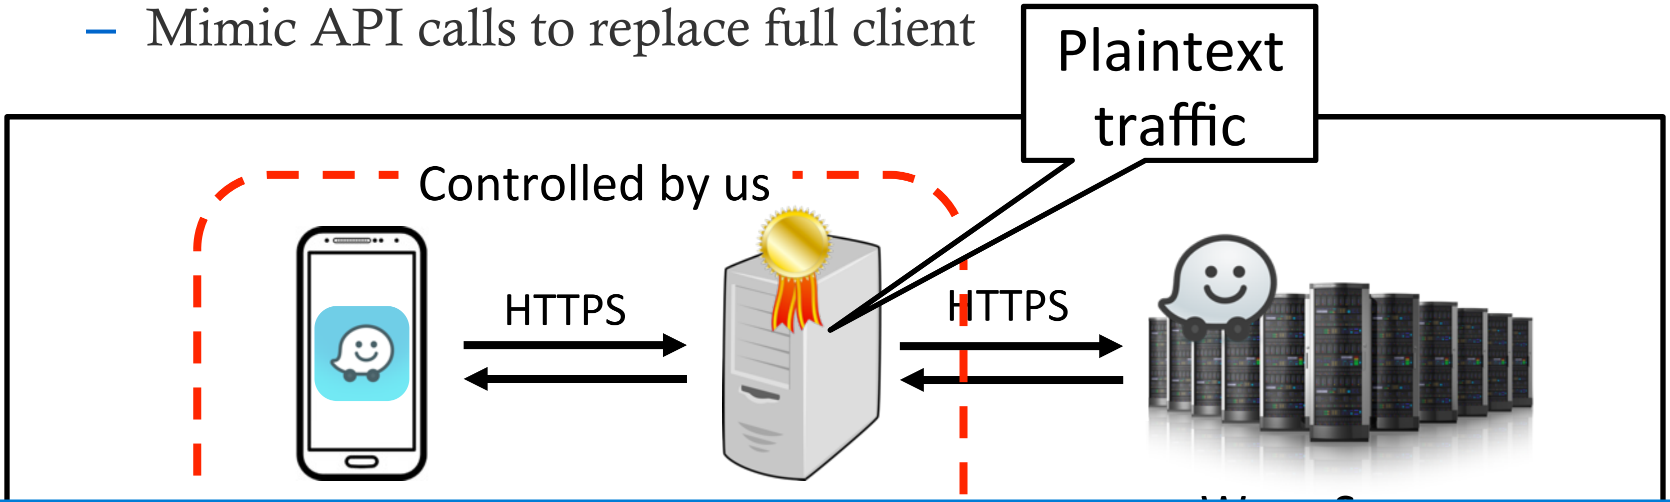

We can create 10,000 Sybil devices on a single PC

Attack #1: Polluting Waze Database

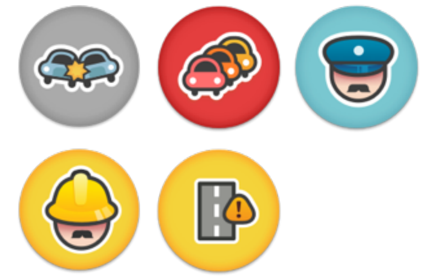

- Fake road-side events.
  - Any type of event at any location
  - Potentially affect 1+billion Google Maps users

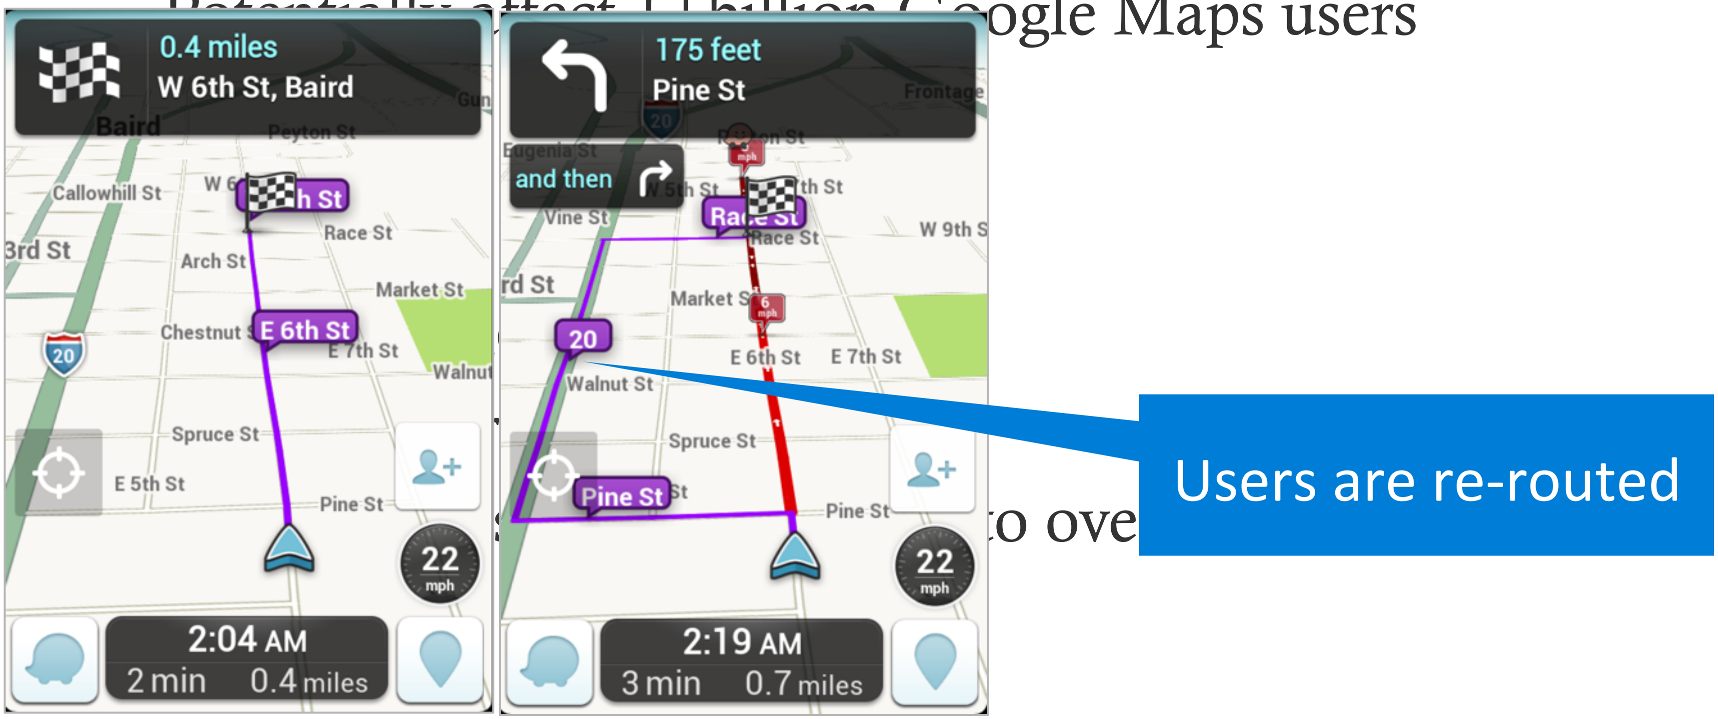

Before After

o ove

Attack #2: User Location Tracking

- Follow (stalk) any Waze user in real-time
  - Waze marks nearby users on the map

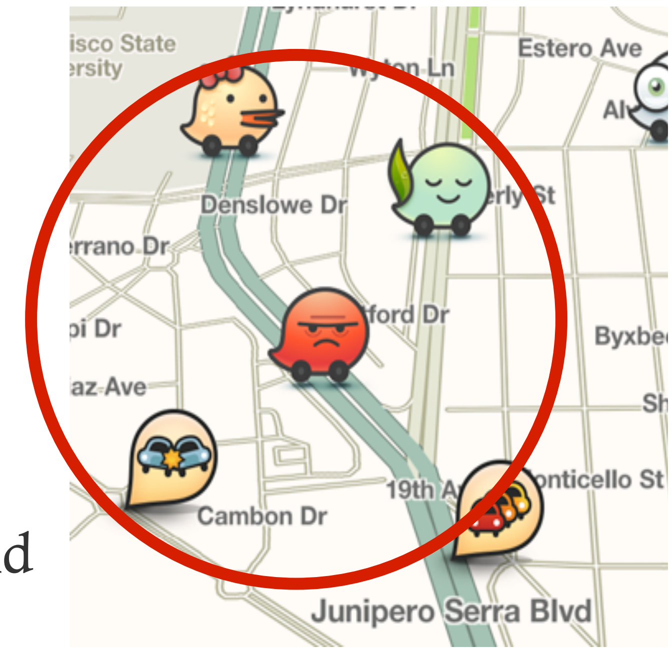

- Pinpoint to exact GPS location
  - Specific hotels, gas stations, etc.
- Remain invisible
  - Move in and out quickly
- Track users in the background
  - Waze uploads GPS in the background
- Track users across days
  - Use creation time as GUID

A Tracking Example

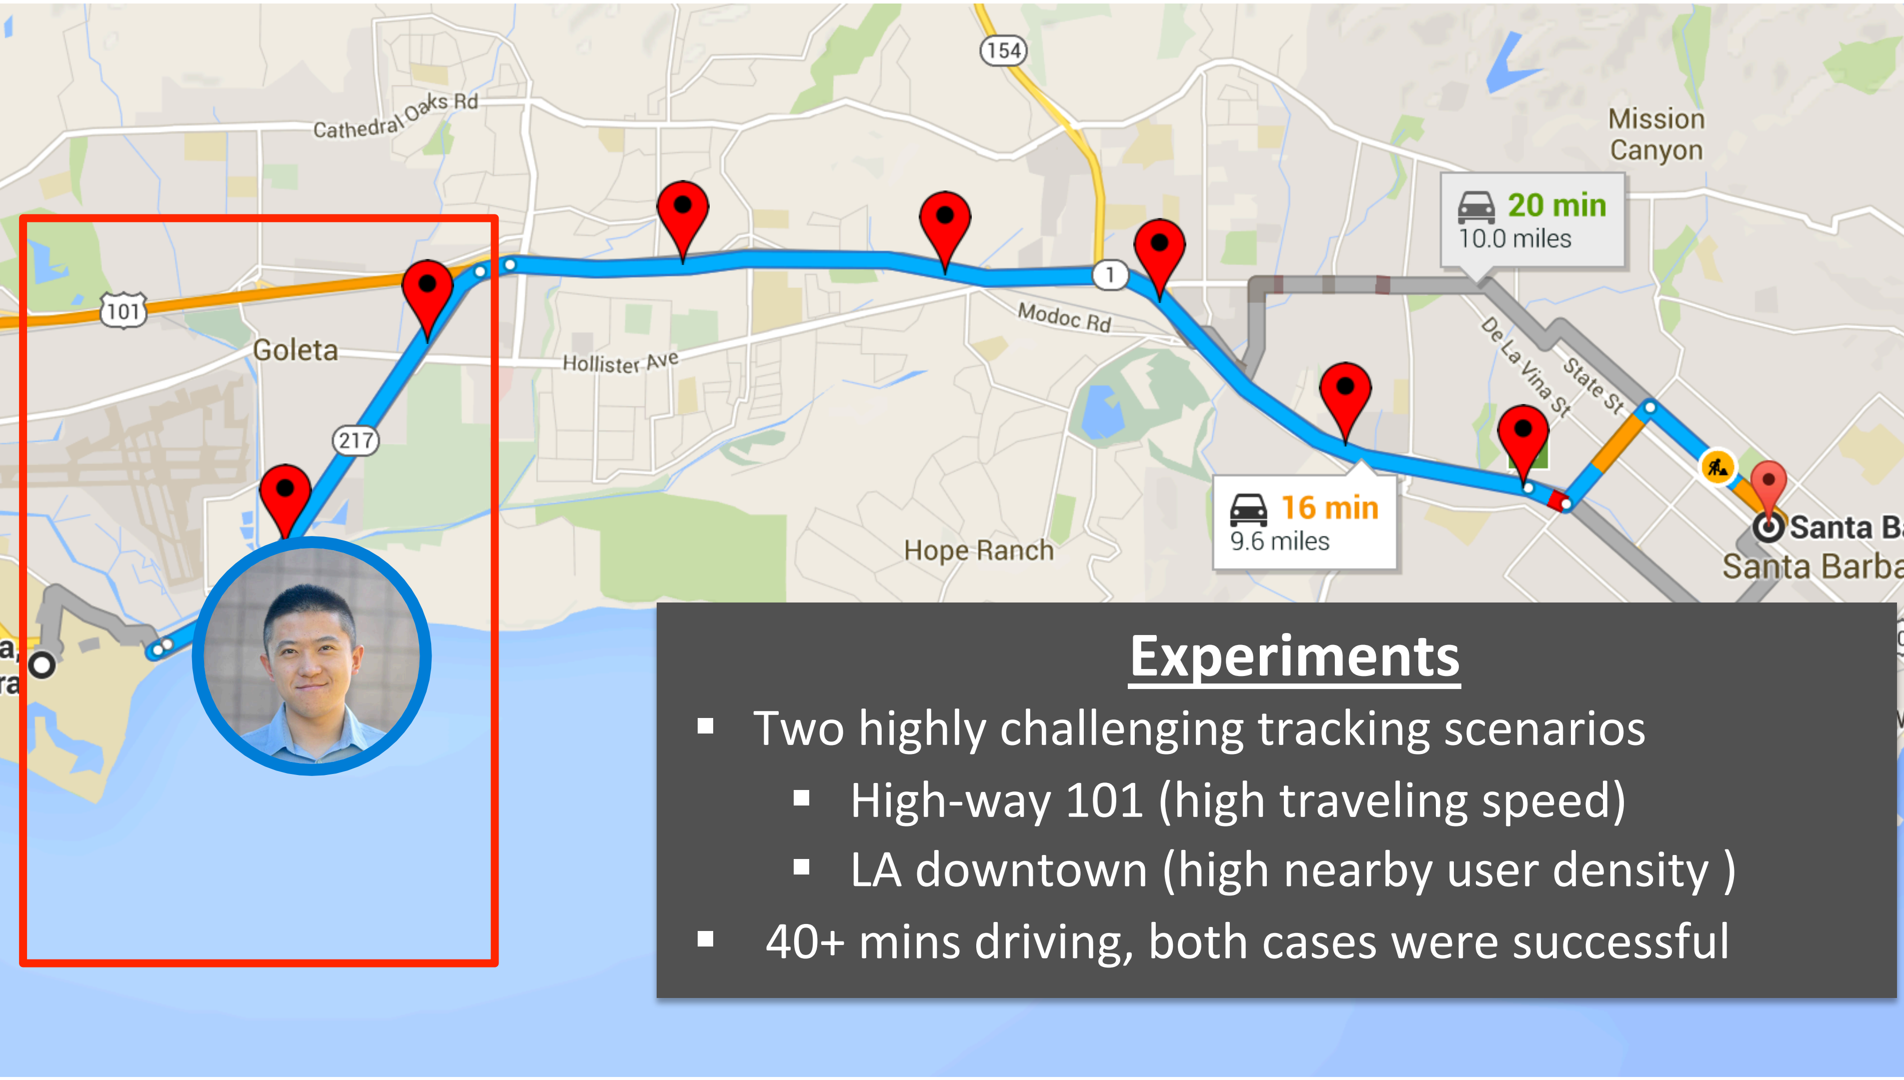

Experiments

- Two highly challenging tracking scenarios
  - High-way 101 (high traveling speed)
  - LA downtown (high nearby user density)
- 40+ mins driving, both cases were successful

Conversation With Waze

Notify Waze and Google

Nov. 14 2014

Pitch work to Fusion

Oct. 18 2015

1st code change: remove background GPS upload

Apr. 16 2016

- Fusion report on tracking
- Media attention

Apr. 26 2016

+21 more

More news coverage

Apr. 27 2016

Public PR release
2nd code change: disable social function for v3.5

+16 more

Working with Waze

May. 11 2016

Time

Waze's Security Measures

- No background GPS
- Hide GPS if not moving
- Hide start/end location

- Disable social feature (v5.3-)
- Special encoding for APIs

- Remove username
- Scramble creation time
- Require SMS verification

Oct. 18 2015 — Apr. 27 2016 — Apr. 29 2016 — May 11 2016 — May 17 2016 — May 23 2016 → Time

- Track active users

- Start collaboration

- Yes, we can still track Waze users
- Much less location information being shared

Broad Implications on Other Apps

- Sybil de
  - Fours
  - Rever
- Tinder/
  - Locat
- Uber/L
  - Track

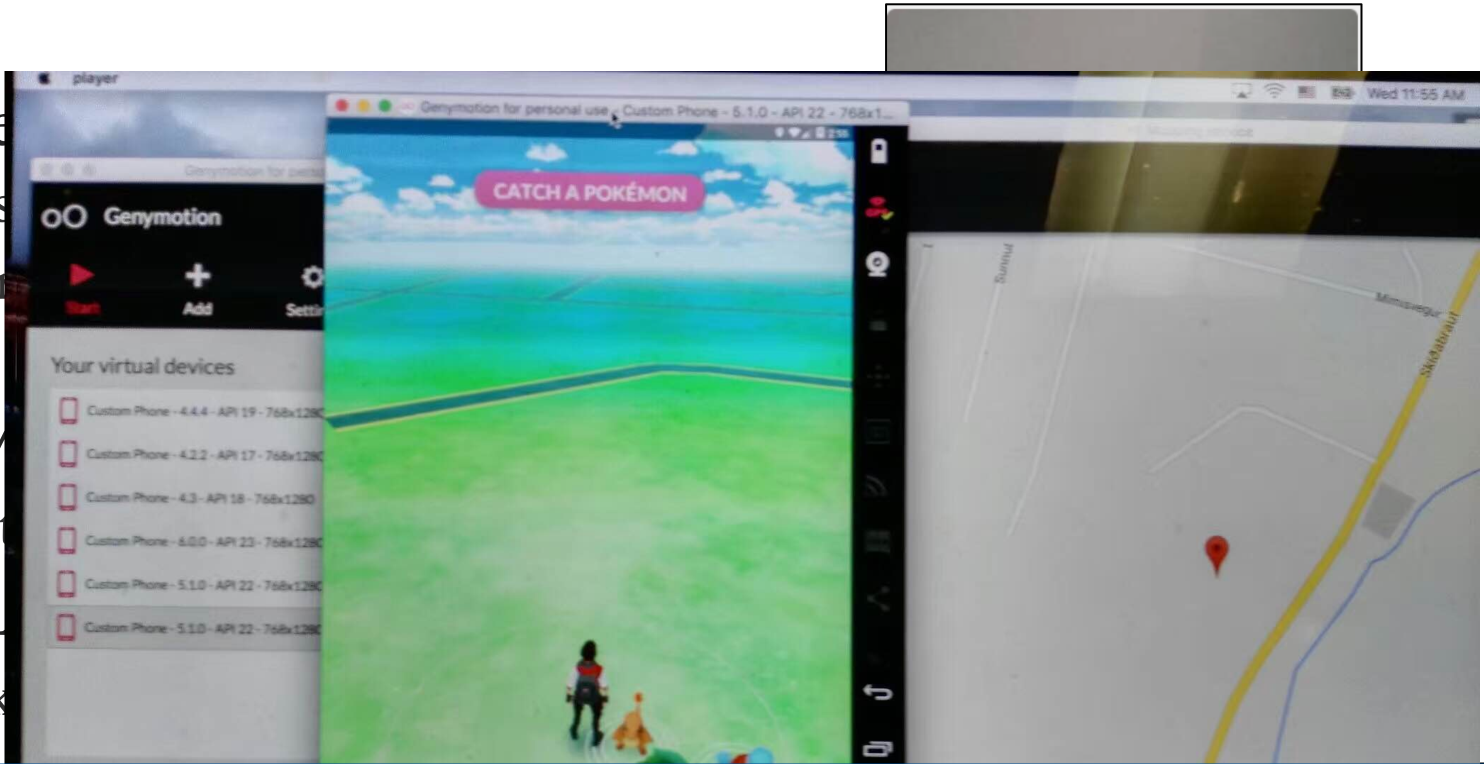

Key Takeaway

- Apps that support "human-to-human" interactions → leak user data
- Sybil devices make this a bigger concern

Talk Outline

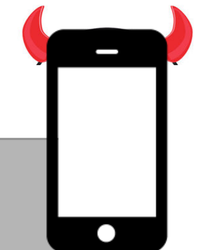

1. Emerging Threat of Sybil Devices

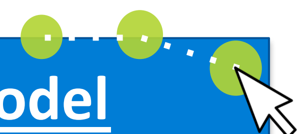

2. Clickstream based User Behavior Model

- Build hierarchy of behavior clusters
- Automatically extract key distinguishing features
- Detect fake accounts, track dynamic behavior changes

Understanding Online Users

- An increasing need to understand user behavior
  - What are the prevalent types of user behaviors?
  - How to identify and understand these behaviors?
  - Do user behaviors evolve/change over time?

LinkedIn

Job seekers

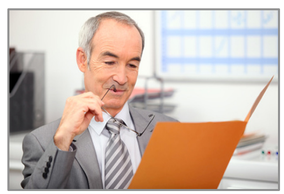

Happily Employed

Job hoppers

Recruiters

……

Are there undesired behaviors (job scams)?

Is the company doing well?

Can we predict key trends in professional/stock market?

Clickstream: You are How You Click

- Clickstream analysis for behavior modeling
  - Clickstream: a sequence of click events (and time gaps)
  - Suitable for identifying fine-grained user behaviors

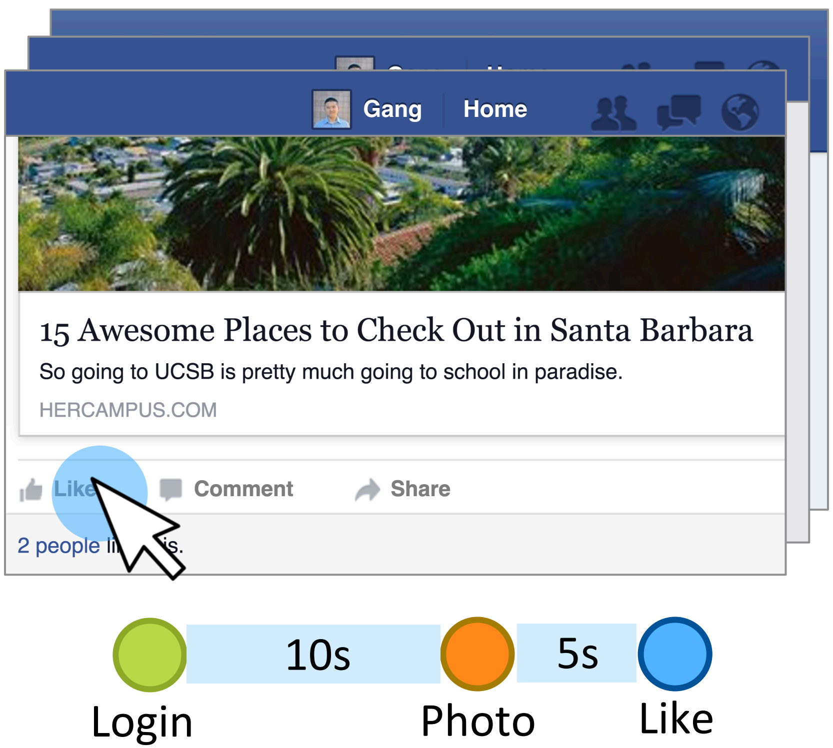

Our Goals

1. Identify natural clusters of user behavior based on clickstreams
2. Extract semantic meanings for captured behaviors
3. Scalable for large online services

User Behavior Model

- Key intuitions
  - Users naturally form clusters
  - More fine-grained user clusters are hidden within big clusters

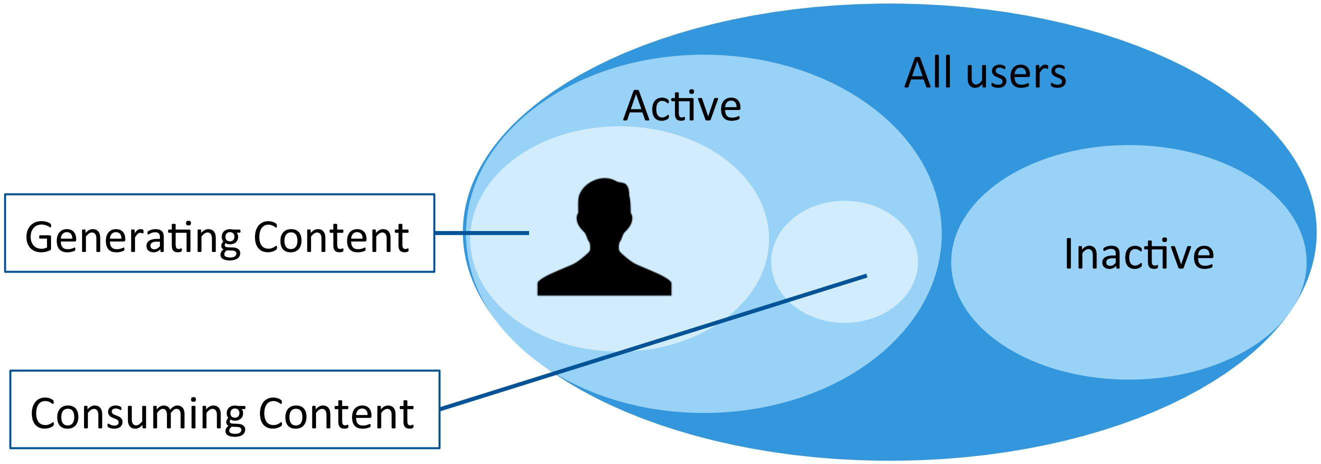

Automatically capture hierarchical structure of behavior clusters

Clickstream Similarity Graph

Identify user clusters that share similar behaviors

1. Map user's clickstreams to a similarity graph
   - Clickstreams are nodes
   - Edge weighted by the similarity of clickstreams

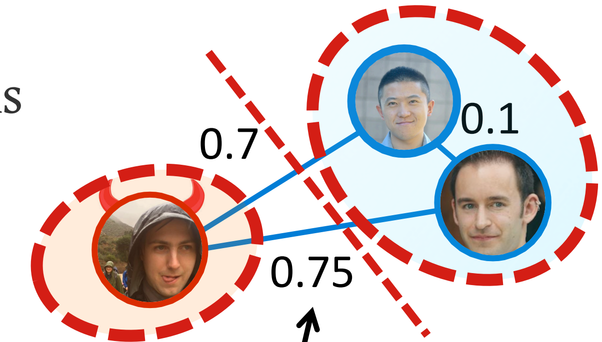

Similarity: common subsequence (count)

Cosine Distance

S_1= AAB
S_2= BBC

$\rightarrow$

$ngram_1$= {A(2), B(1), AA(1), AB(1), AAB(1)}
$ngram_2$= {B(2), C(1), BB(1), BC(1), BBC(1)}

$\rightarrow$

V_1=(2,1,0,1,1,0,0,1,0)
V_2=(0,2,1,0,0,1,1,0,1)

Hierarchical Clustering
with "Iterative Feature Pruning"

- Partition a clickstream similarity graph
  - Identify fine-grained clusters within big clusters
  - Select features to interpret each cluster

- No pre-defined features / constraints
- Cons ickstream (ngram)

Feature selection based on Chi-square statistics

Based on clustering quality convergence (modularity)

1. Start f
2. Partition the gra
3. Select
   new clu
4. Prune top features, re-compute similarity graph, detect sub-clusters
5. Iteratively repeat 2-4 for new graphs, terminate if no clear cluster structures

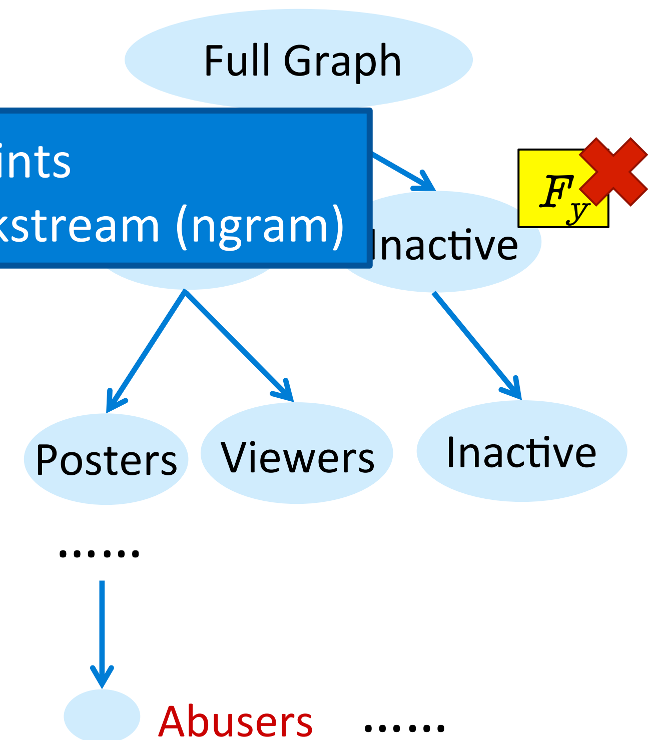

Application #1: Behavior Analysis

Based on 100K Whisper users, 142M clicks

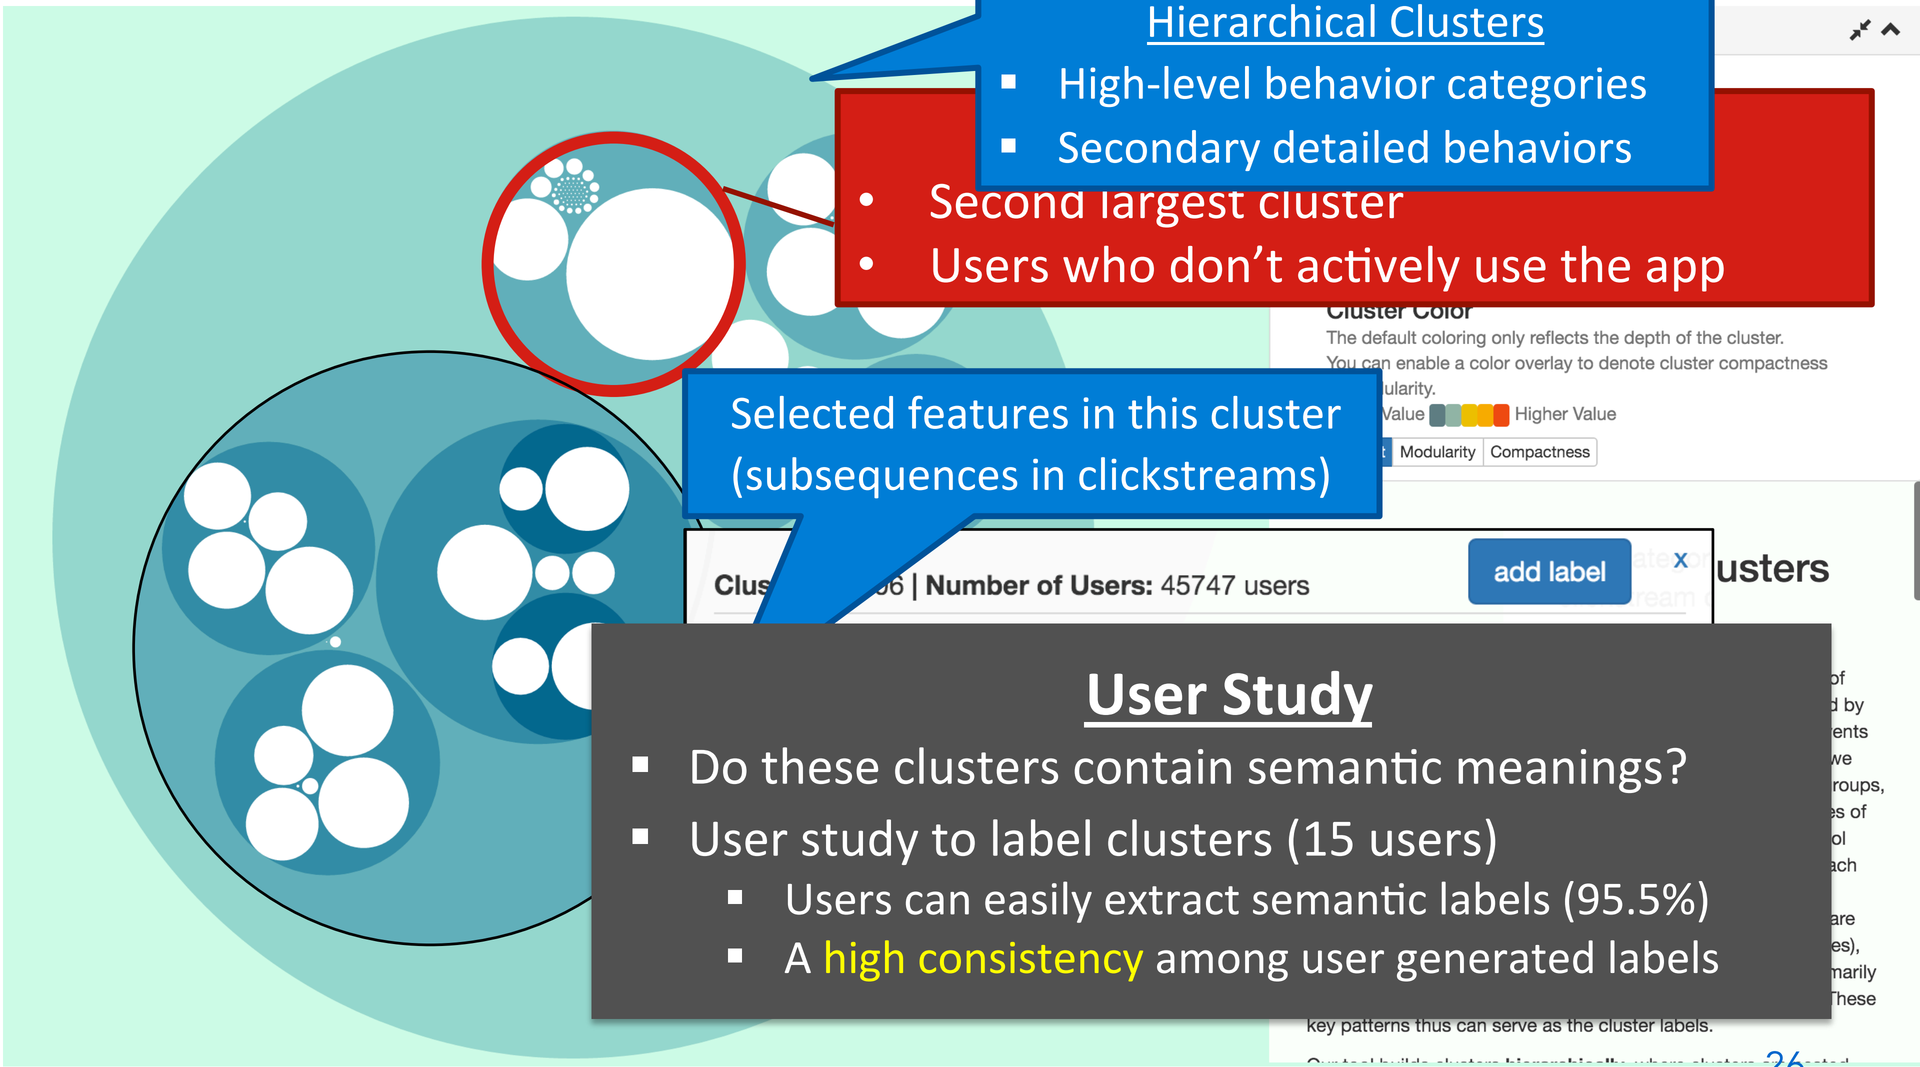

Hierarchical Clusters

- High-level behavior categories
- Secondary detailed behaviors

- Second largest cluster
- Users who don't actively use the app

Selected features in this cluster (subsequences in clickstreams)

User Study

- Do these clusters contain semantic meanings?
- User study to label clusters (15 users)
  - Users can easily extract semantic labels (95.5%)
  - A high consistency among user generated labels

Tracking Behavior Changes

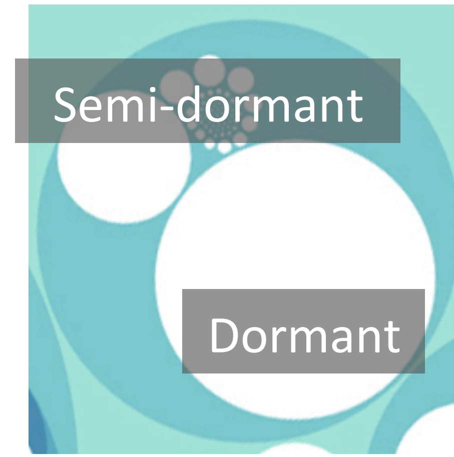

- Users within the inactive cluster
  - Dormant: zero active actions
  - Semi-dormant: only login occasionally
- **Hypothesis:** users in inactive cluster will migrate to "dormant" cluster over time
- Analyzing user migration
  - Split clickstream data into three snapshots, 2-week each
  - Compare user behavior clusters across snapshots

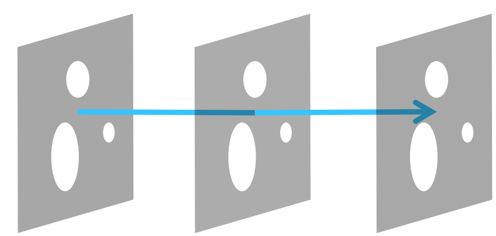

Predicting User Dormancy

- Users turning dormant within adjacent snapshots
  - Dormant users are likely to remain dormant (94%)
  - Semi-dormant users are more likely to turn dormant (17% vs. 1%)

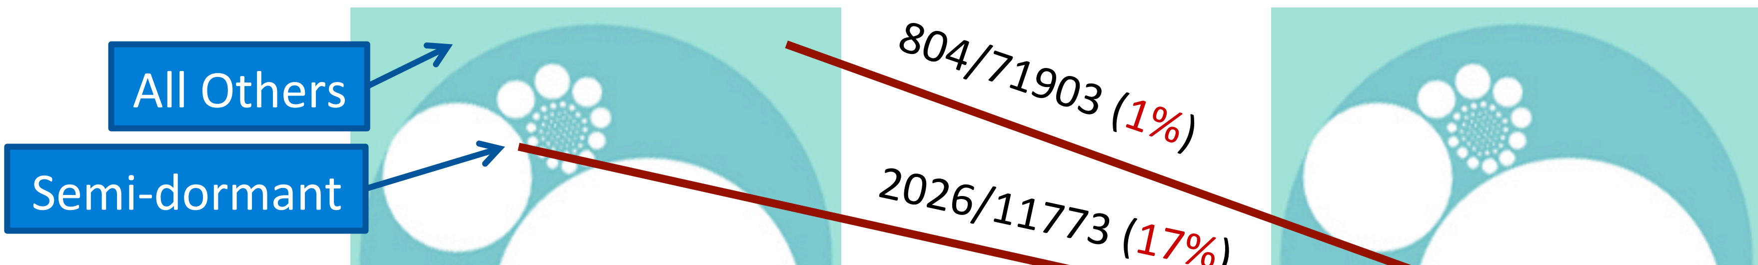

- Predict user dormancy by monitoring the inactive cluster
- Implement necessary interventions to retain users

Ongoing: identify “paths” of behavior changes

“What makes a user turn into a bully/troll?”

Application #2: Sybil Detection

- Detecting fake accounts in social networks [USENIX SEC'13]
  - Real users and fake users behave differently → different clusters

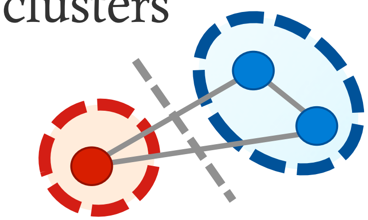

- Ground-truth evaluation
  - Clickstream data from Renren (10K Sybil + 6K normal)
  - Highly accurate: 0.7% false positive rate, 4% false negative rate
- Shipped our prototype code to renren LinkedIn
  - LinkedIn: detected 200 new Sybils in a set of 36K "good" users
  - Renren: detected new type of spam attack (image spammers)

"Image" Spammers in Renren

- Embed spam content in images
- Easy to evade text/URL based detectors

Talk Outline

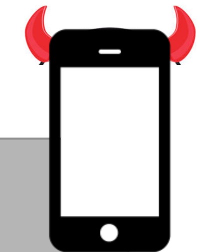

1. Emerging Threat of Sybil Devices

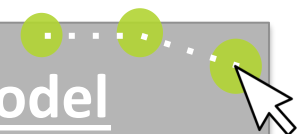

2. Clickstream based User Behavior Model

Conclusion

- Highlights of My Work
- Ongoing and Future Projects

Research Summary

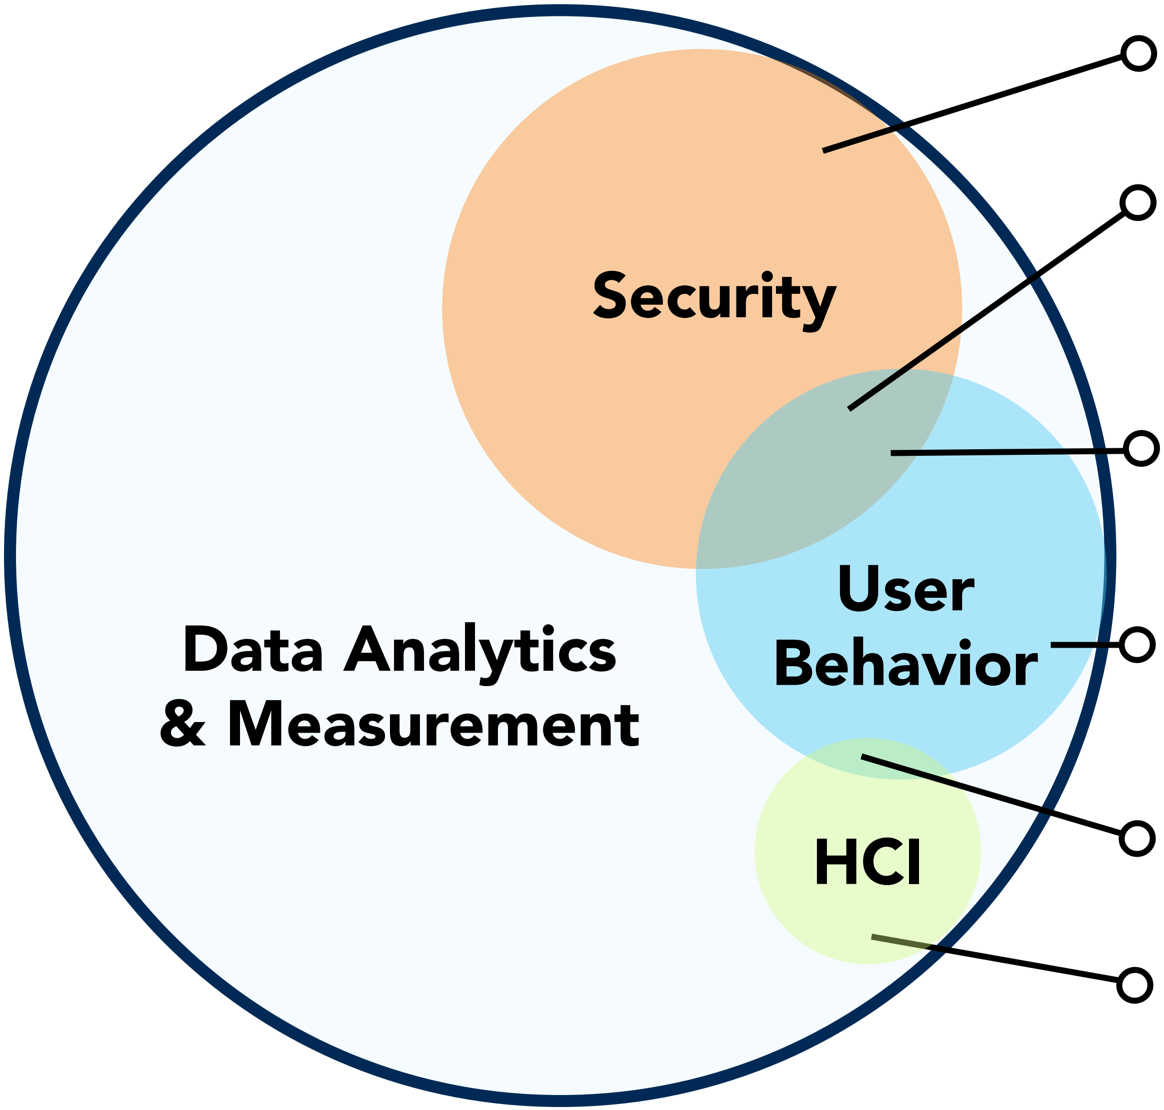

Malware Distribution
[DSN'13]

Spam, Crowdturfing
[WWW'12] [NDSS'13]
[Usenix Security'13] [IMC'13]
[Usenix Security'14]

Mobile Networks
[MobiCom'11] [HotNets'13]
[SIGMETRICS'13] [MobiSys'16]

Social & Crowdsourcing
[WWW'13] [HotMobile'13]
[IMC'14] [ICWSM'16] [IMC'16]

User Behavior Modeling
[CHI'16]

Crowds in Fin. Markets
[CSCW'15]

Impact of Research

- **Academic Impact**
  - Broad publications in Security, Measurement, Mobile, HCI
  - Frequent media coverage

- **Industry Impact**
  - Deployed: malware/Sybil detection, location anonymity scheme
  - Actively protecting millions of users in production systems

400M 220M 20M

Short Term: Sybil Devices Defense

- How to defend against Sybil devices?
- Apps: protecting APIs against reverse-engineering
  – Waze: special encoding on data fields of API calls
  – Yik Yak: use HMAC for message integrity
  – Periscope: SSL Pinning
- Lack of empirical understanding at a large scale
  – What apps are vulnerable to API reverse-engineering?
  – What security approaches are used to protect APIs?
  – How effective are these security approaches?

Top 100,000 Apps

GET IT ON Google Play

Simulator + Traffic analysis

- Security approaches used
- Is APIs visible?
- Can APIs be simulated?

Short Term: $ in Mobile Systems

- Mobile digital wallet
  - Wide adoption
  - Many integrate with social features
  - How do users use the system? Are there malicious activities there?

Venmo Data: 90 million public transactions from 7 million users
- Infer who you are based on how you make transactions (Gambling bookies, merchants, drug dealers)

- Mobile payment based social Q&A (FenDa)
  - Ask experts questions directly on your phone
    - Pay $50 to ask a doctor a question
    - Get paid $1 from anyone who listens to the answer
  - Is money a good incentive to obtain/archive knowledge?

FenDa Data: 65K users/experts/celebrities and their answers

Future Directions: Long Term

- Explosive growth of Internet devices
  - Smartphones, wearable/medical devices, smart vehicles, smart city

Future trends

- Massive data from both cyber and physical world
- Opens up new attacking surface

User-centric security

- Identify real security threats by understanding user behaviors
- Statistical user behavior analysis that can scale

Thank You!

http://people.cs.vt.edu/~gangwang
gangwang@vt.edu