K-means++ vs Behavioral biometrics: One Loop to Rule Them All

Parimarjan Negi, Prafull Sharma, Vivek Jain, Bahman Bahmani
Stanford University

What is behavioral biometrics?

Historically:

- Handwriting recognition
- Telegraph Operators in WWII

Behavioral Biometrics: Modern Version

- Typing (Keystroke Dynamics)
- Mouse movements
- Typing, or swiping on a smartphone
- Through other smartphone sensors, e.g., gait analysis

Secondary Authentication

- Most secondary authentication methods involve the user actively doing something, e.g. two factor authentication.
- Behavioral Biometric methods function in the background

Popular

AI-based typing biometrics might be authentication's next big thing

SECURITY

The Future of Biometrics Could Be in What You Type

Behavioral Biometrics "stole the show"* at Google I/O

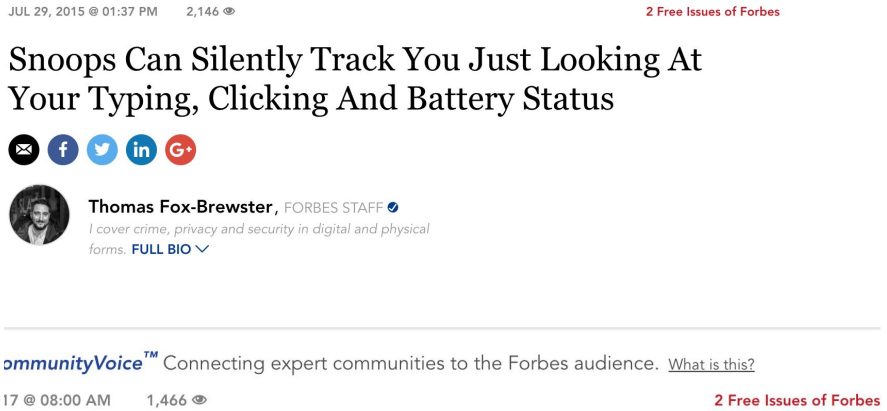

JUL 29, 2015 @ 01:37 PM 2,146

2 Free Issues of Forbes

Snoops Can Silently Track You Just Looking At Your Typing, Clicking And Battery Status

Thomas Fox-Brewster, FORBES STAFF

I cover crime, privacy and security in digital and physical forms. FULL BIO

ommunityVoice™ Connecting expert communities to the Forbes audience. What is this?

17 @ 08:00 AM 1,466

2 Free Issues of Forbes

Biometrics: A Stepping-Stone To Eliminating The Password Forever

Quantifying Errors

False Rejection Rate: How many genuine samples get rejected?

False Acceptance Rate: How many impostor samples get accepted?

Equal Error Rate: Threshold where FAR = FRR

General Scenario

- Attacker **knows** the target user's password
- Target user's account protected using keystroke dynamics system
- Attacker does not have access to typing data from user

Attacker Aim

- Produce timings (key-press time, duration between keys) for a given password

How many tries does it take an attacker to “fool” such systems?

Targeted Attack Scenario

- Idealized scenario for the adversary
- has unlimited to attack single target
- Can generate a lot of timing samples for the target’s password from MTurk

Indiscriminate Attack Scenario

- Leaked database of passwords - attacker wants to quickly try these passwords for all accounts
- Too expensive to collect samples for each password
- Has access to precomputed datasets of typing data from the general population

Example Password: “Mustang”

- **mu**tter, **mu**mble
- b**us**, f**us**s
- try**st**, li**st**
- da**ta**, io**ta**
- th**an**, cr**an**e
- ba**ng**, ra**ng**

Is everyone’ behaviour unique?

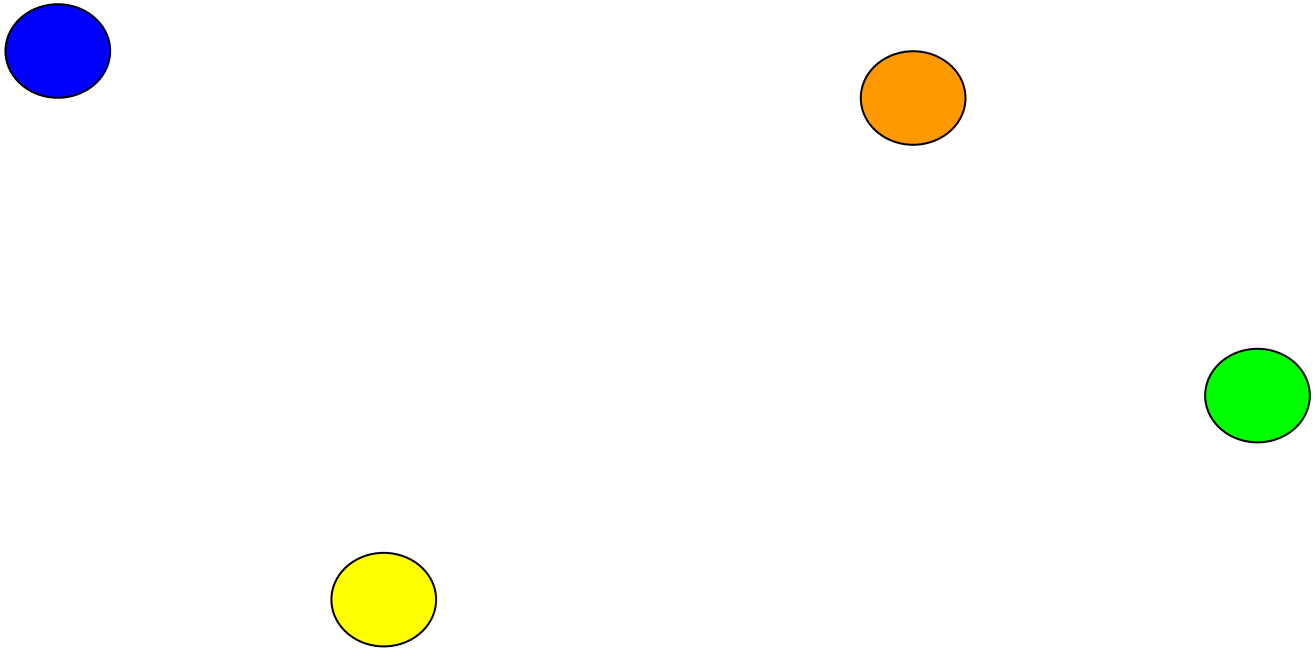

Hypothesis: belongs to a bigger family of similar patterns

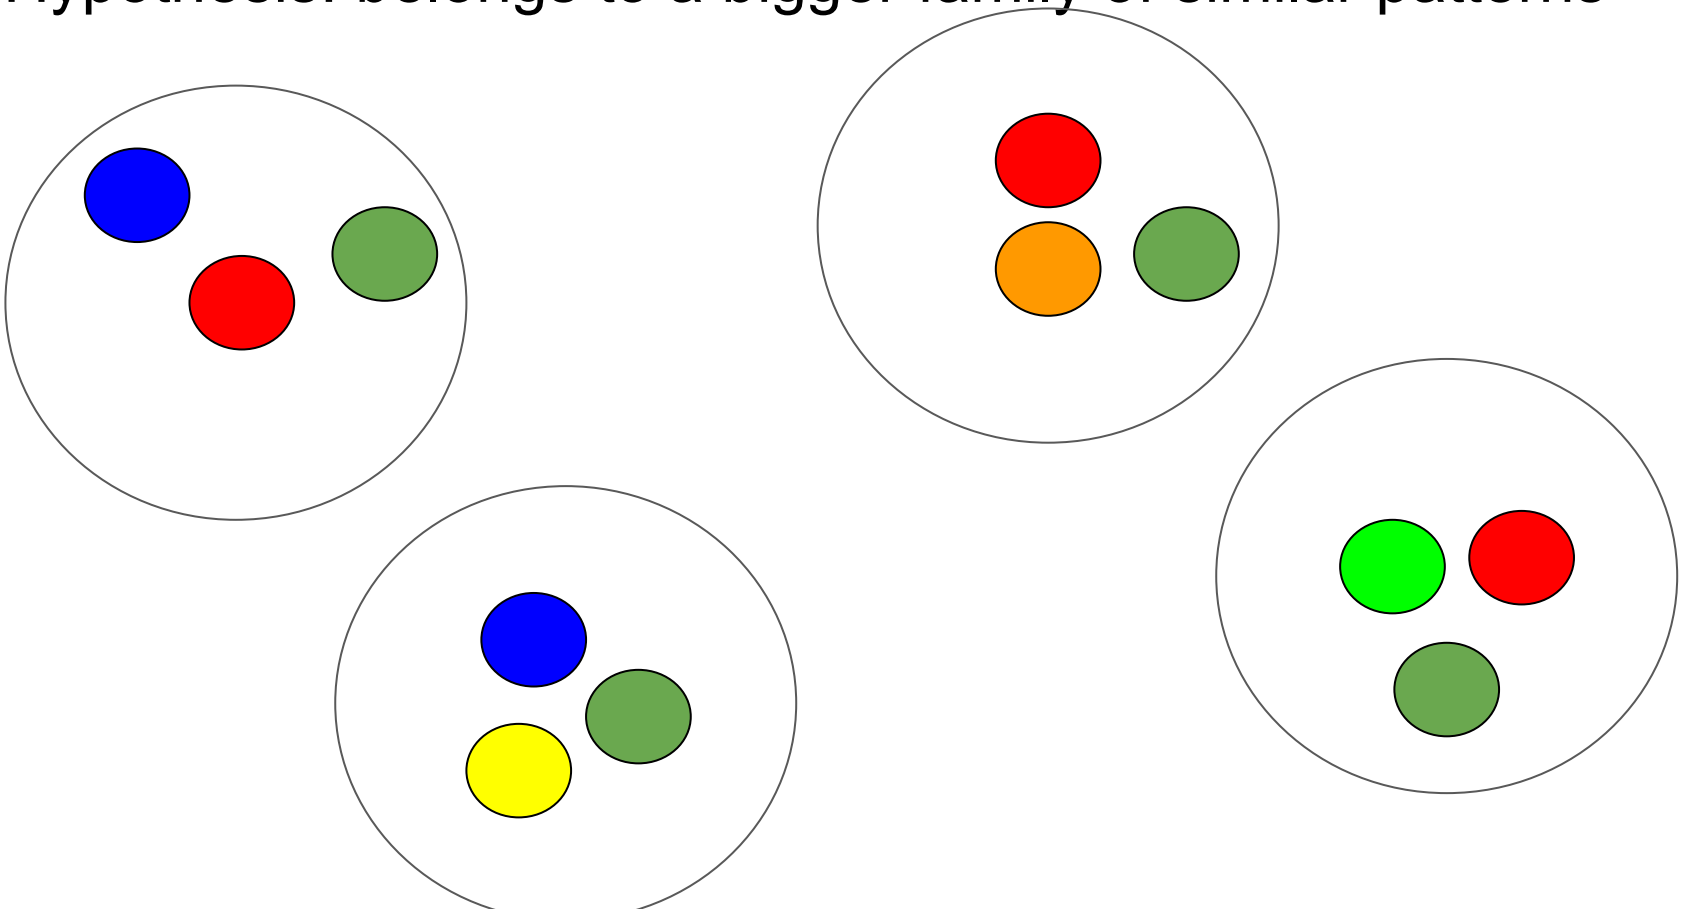

Hypothesis: If we find another user in the same “family”, we can “fool” the classifier

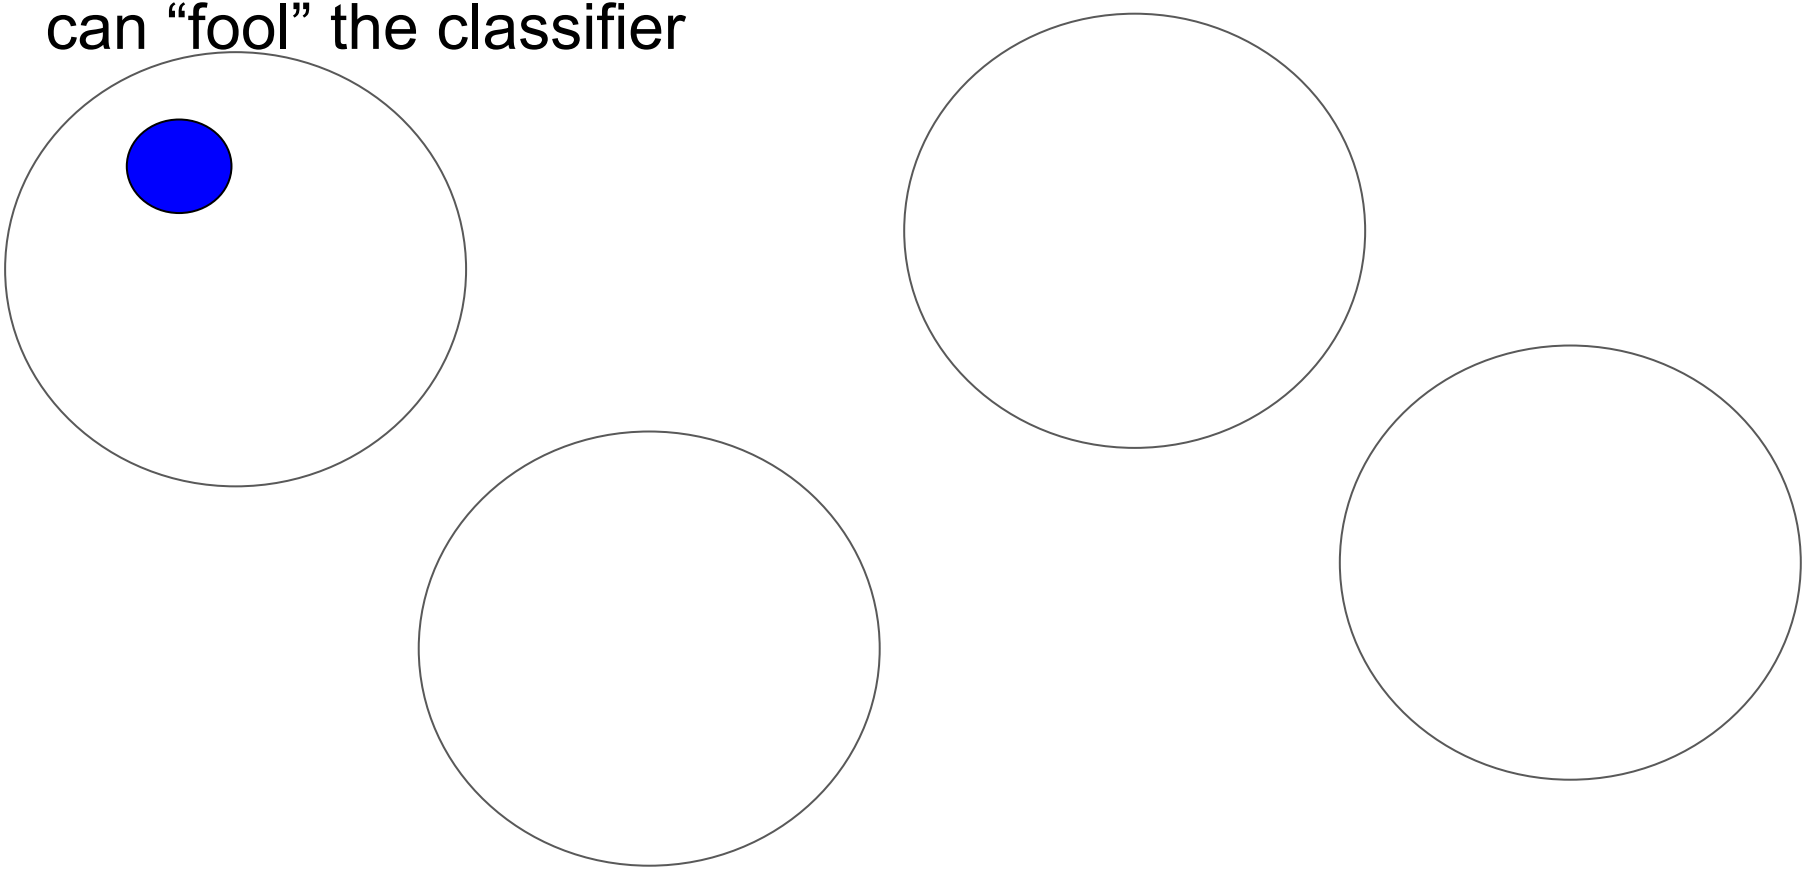

Idealized Algorithm: Randomly Choose first try

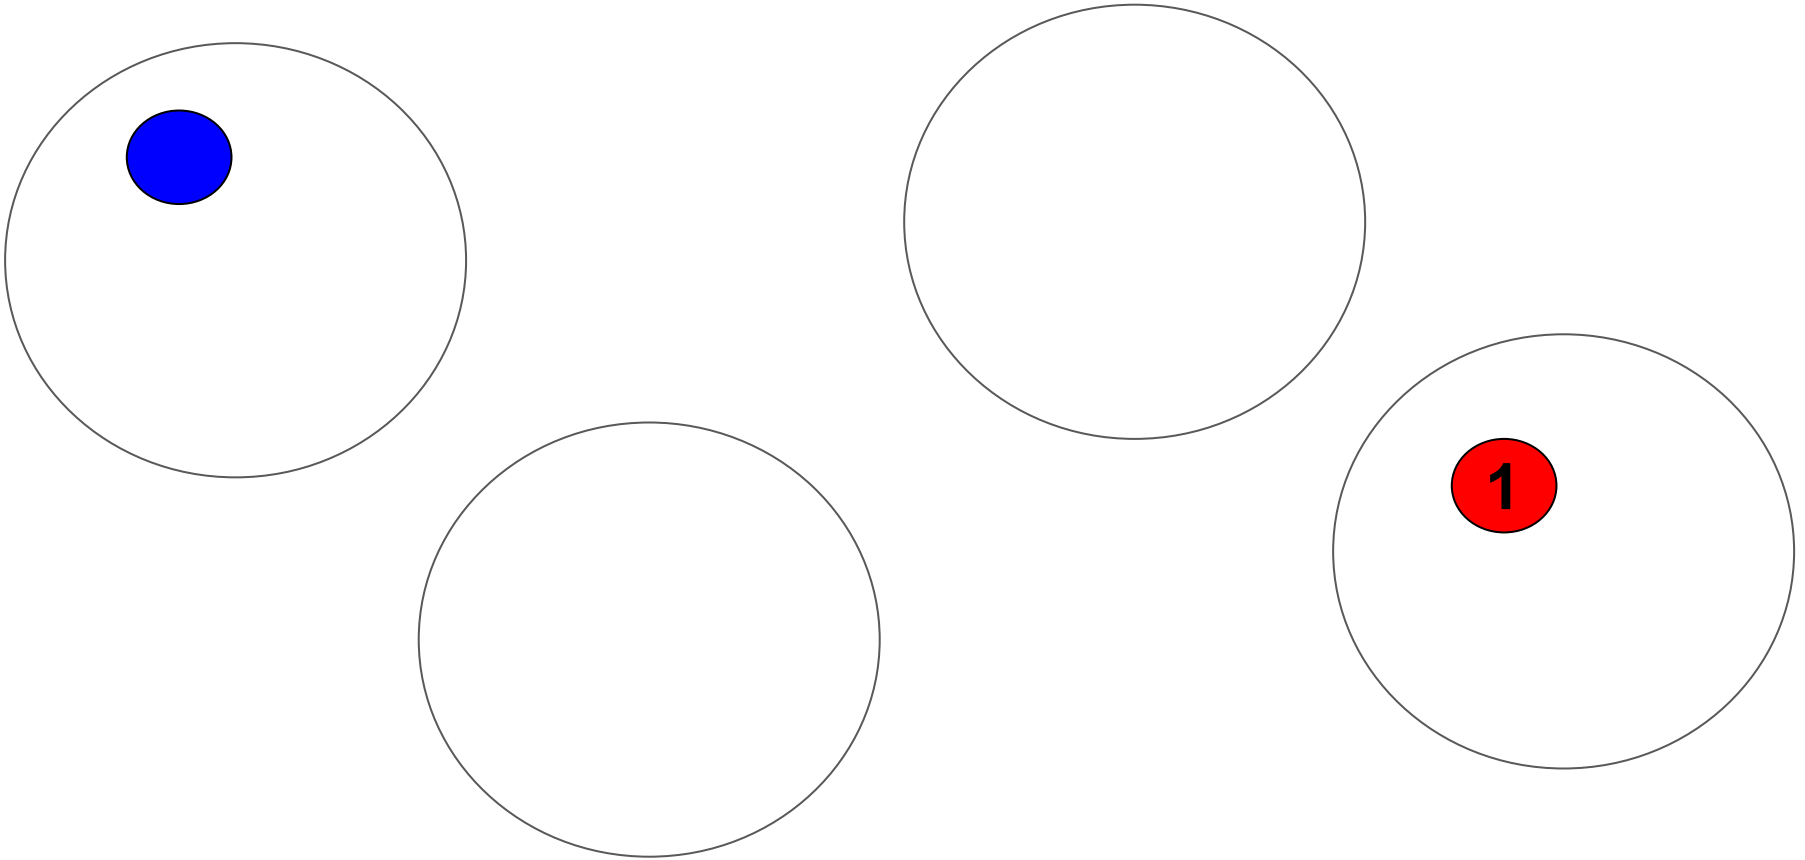

Idealized Algorithm: Choose next try from another cluster

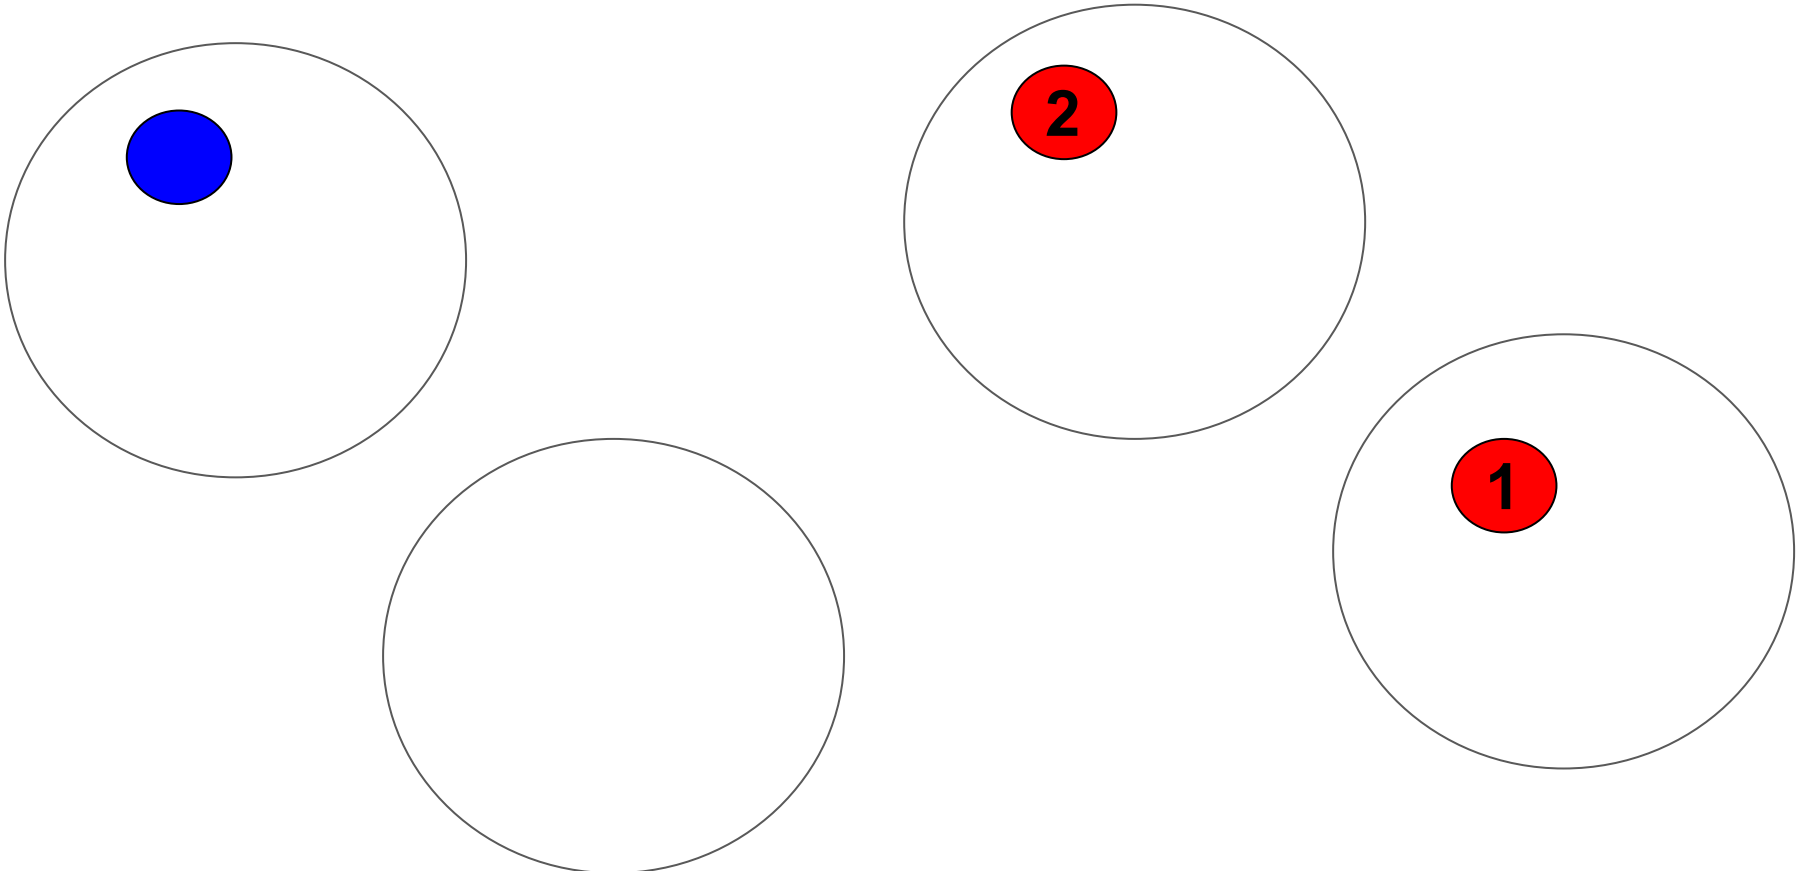

Idealized Algorithm: Choose next try from another cluster

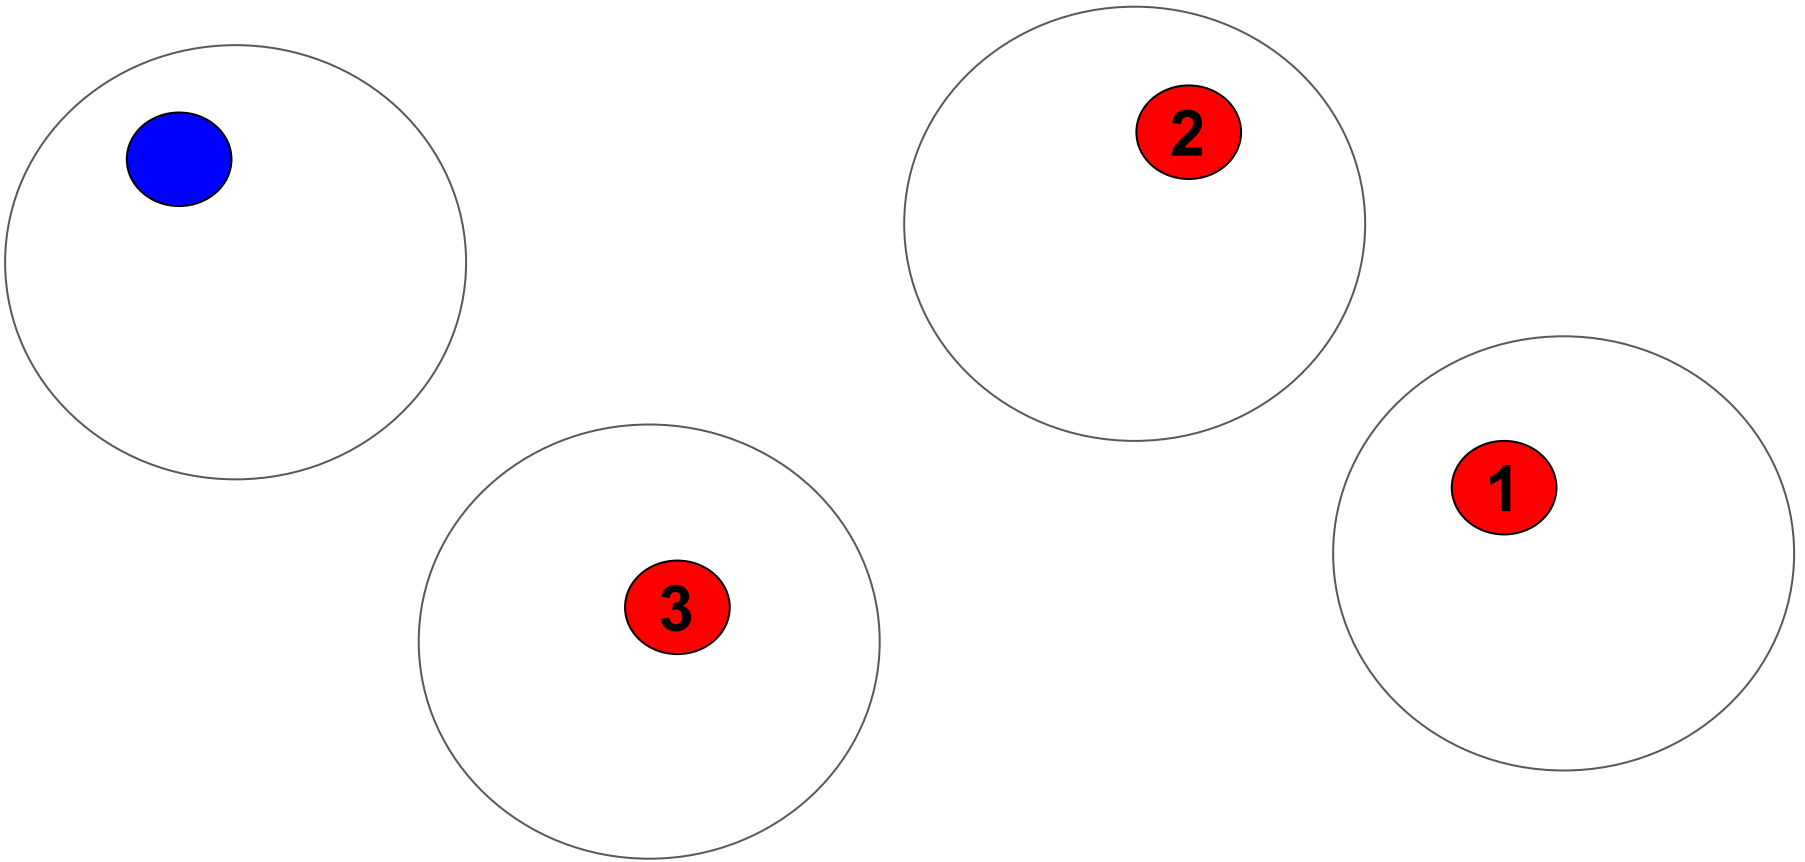

Hypothesis: If we find another user in the same “family”, we can “fool” the classifier

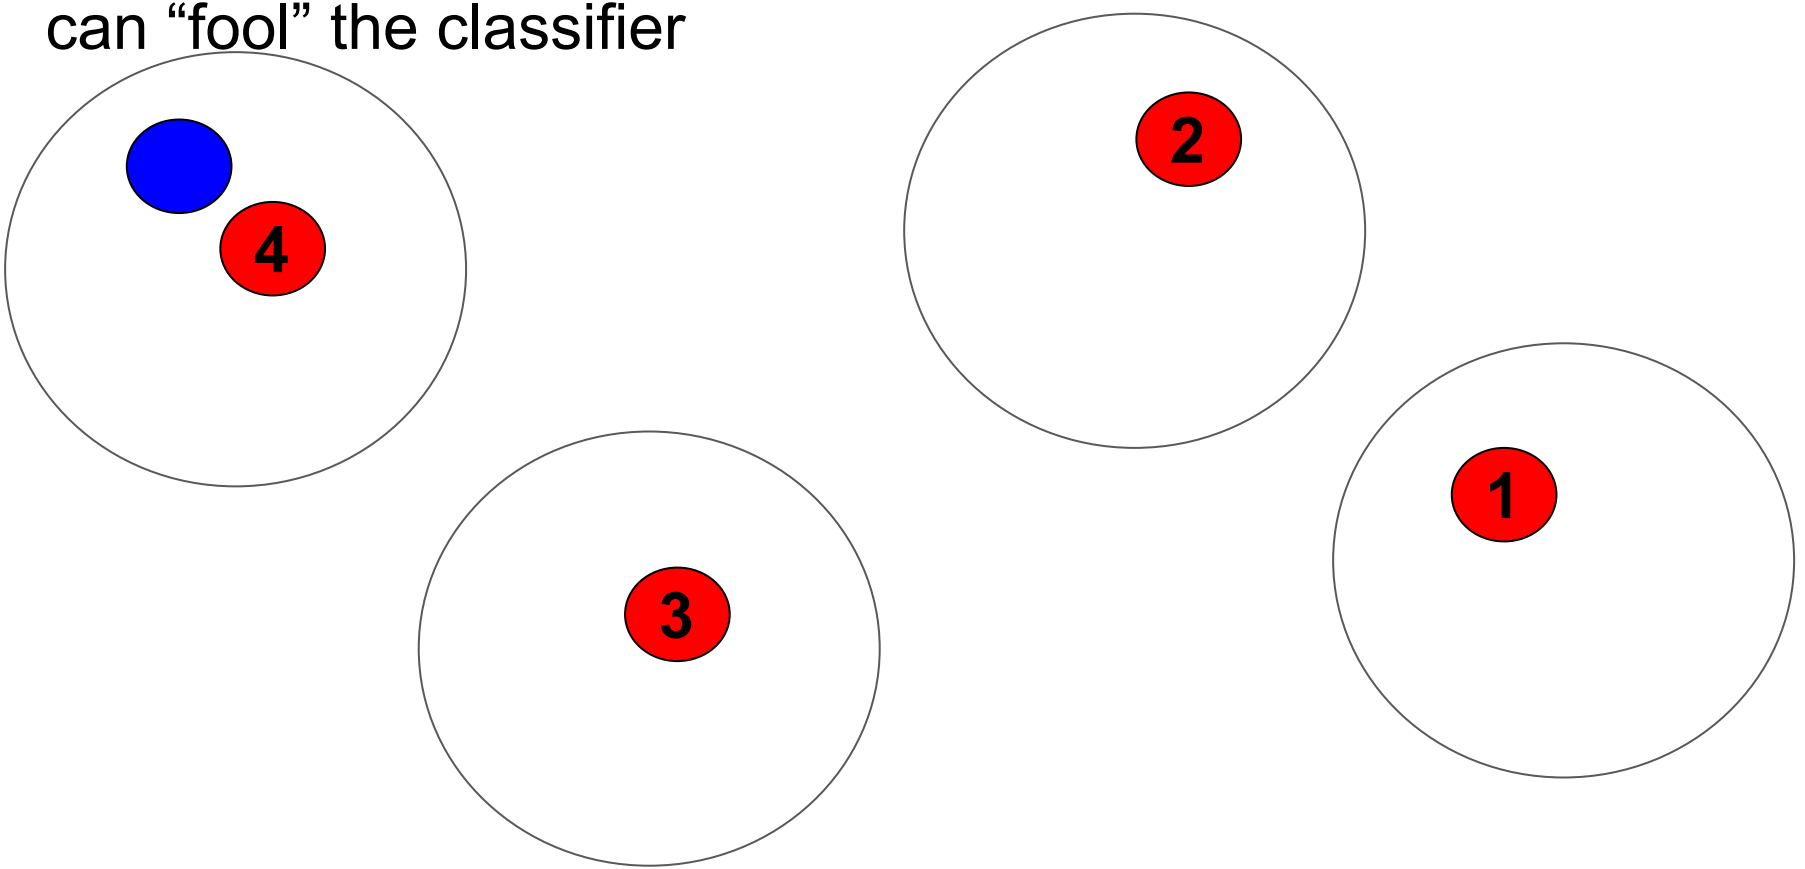

K-means++

- Initialization routine for centroids of K-means clustering
- At each successive iteration, finds centroids that are “far away” from the previous centroid
  - i.e., similar to finding a new try from a different family

Dataset I: DSN

- password: **.tie5Roanl**
- 51 subjects
- 400 repetitions

Dataset II: MTurk

- passwords: **mustang, password, letmein, abc123, 123456789**
- 583 subjects
- ~100 repetitions per password

One Class Classifiers

- Manhattan
- SVM
- Autoencoder
- Contractive Autoencoder
- Gaussian
- Gaussian Mixture

Two Class Classifiers

- Random Forests
- K-Nearest Neighbors
- Fully Connected Neural Network

EER Scores

Name of Classifier	DSN EER	MTurk EER
Manhattan	0.091	0.097
SVM	0.087	0.097
Gaussian	0.121	0.109
Gaussian Mixture	0.137	0.135
Autoencoder	0.099	0.099
Contractual Autoencoder	0.086	0.099
Random Forest	0.08	0.067
k-NN	0.09	0.090
FC Neural Net	0.08	0.091

Results

MTurk Dataset
SVM

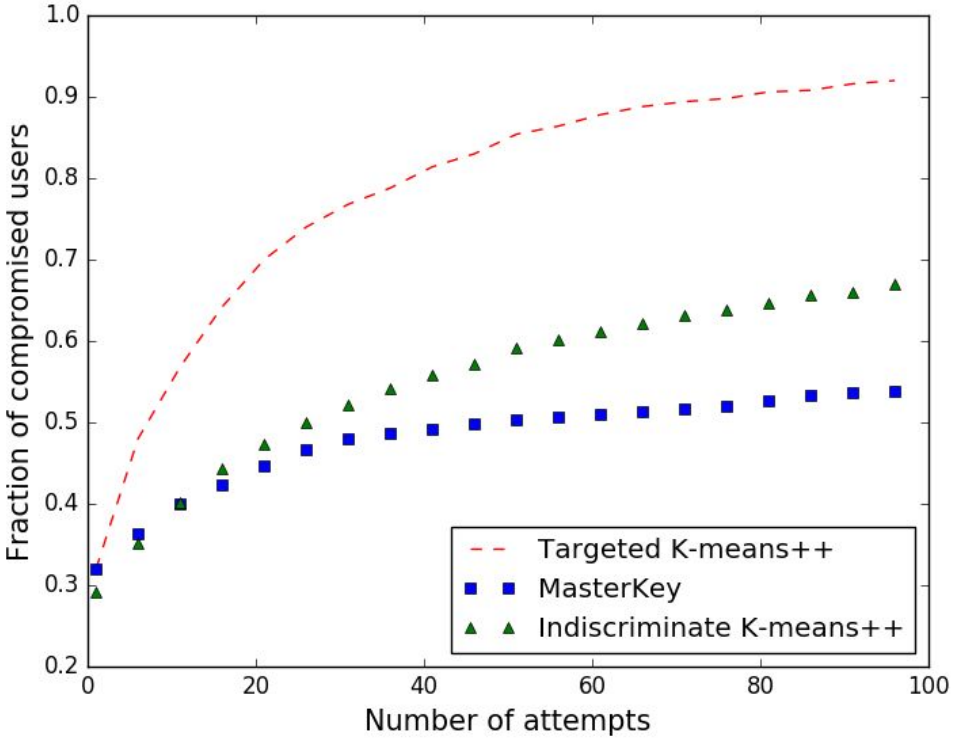

Results

MTurk Dataset
Random Forests

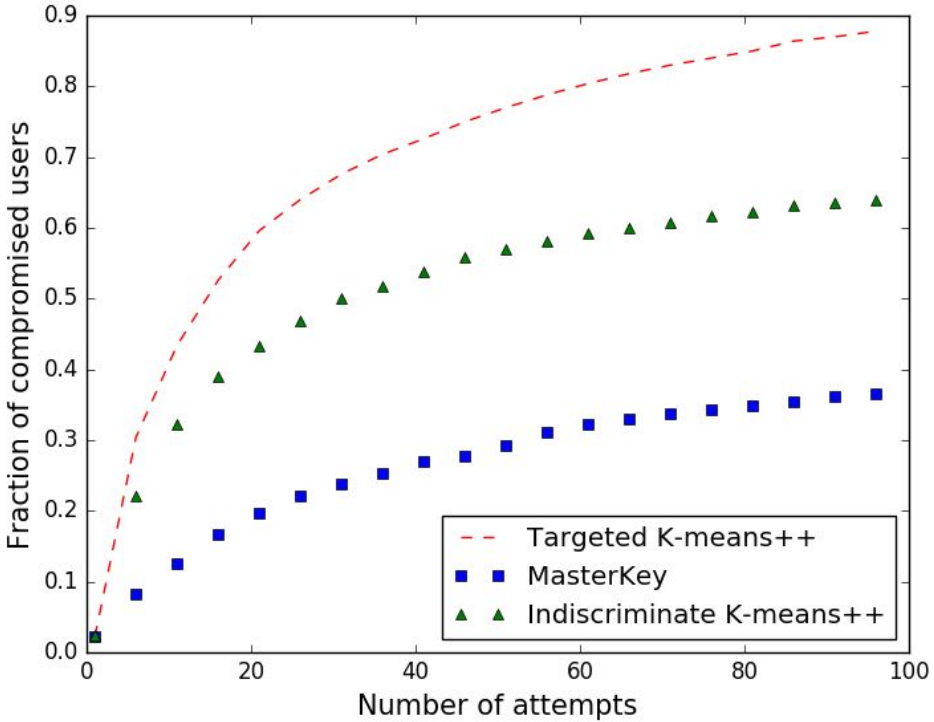

Usual Threshold

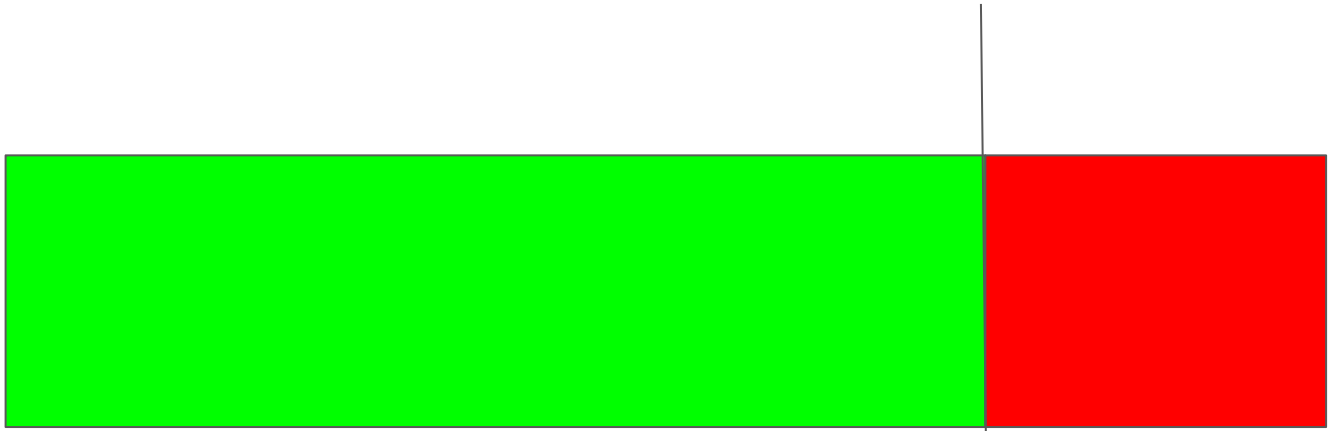

User Scores

EER Threshold

Conservative Threshold

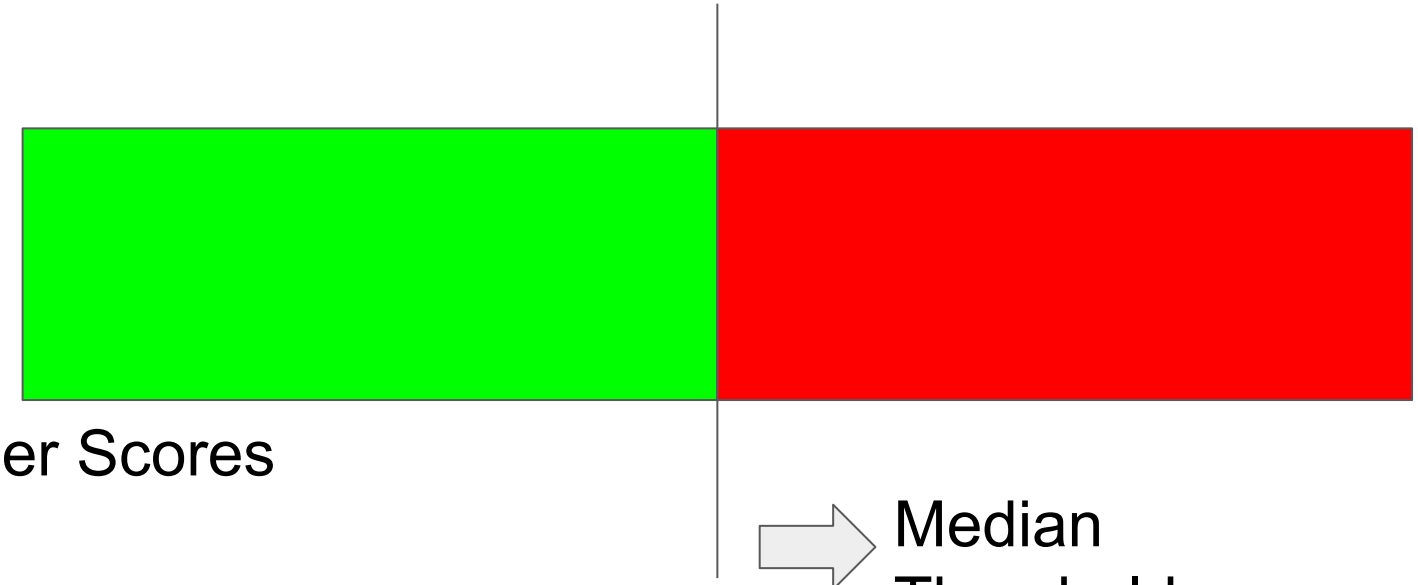

Conservative Thresholds I

Targeted Manhattan

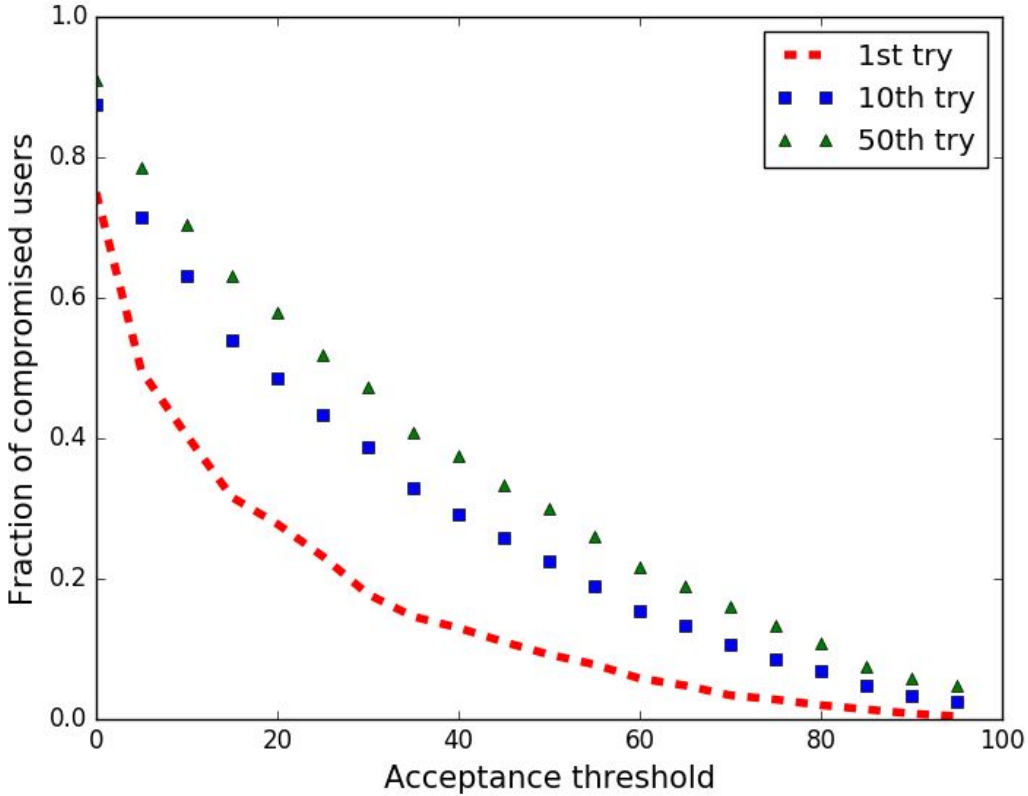

Conclusion

- Behavioral Biometrics are promising but we need to improve them with regards to motivated adversaries
- Classifiers can potentially be made more robust by aiming to thwart such adversarial models
- datasets, code