

Social Turing Tests: Crowdsourcing Sybil Detection

Gang Wang, Manish Mohanlal, Christo Wilson, Xiao Wang,
Miriam Metzger, Haitao Zheng, Ben Y. Zhao



Crowdsourcing vs. Sybil Identities

- Social networks losing the battle against fake accounts
 - Measurements show Sybils do not form clusters, target insertion into specific communities instead (IMC 2011)
- Idea: build a crowdsourced Sybil detector
 - Leverage human intelligence and intuition
 - Resilient to changing attacker strategies

Crowdsourcing: a process that enlists many people to do **small jobs** to solve **problems** that software cannot

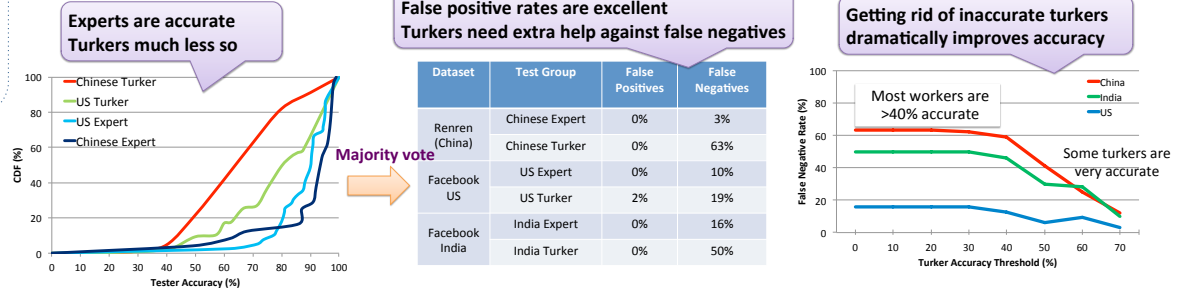


- Open Questions
 - How **accurate** is human based detection?
 - What **factors** affect detection accuracy?
 - Is this approach scalable, i.e. **cost effective** for large systems?

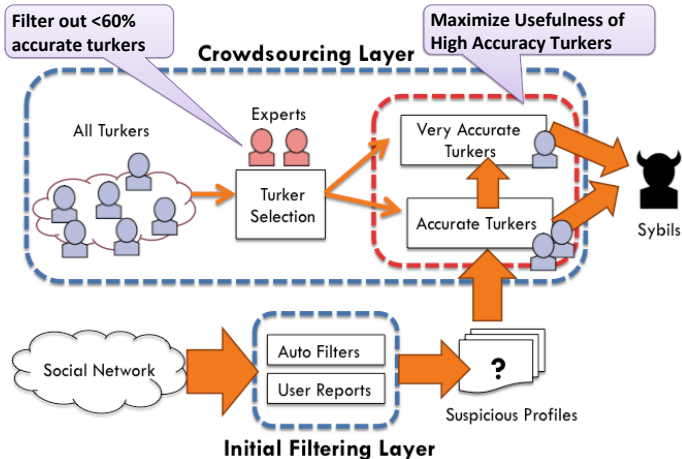
Large User Study

- Two groups of users
 - Experts – CS professors, masters, and PhD students
 - Turkers – crowdworkers from Mechanical Turk and 猪八戒
- Three ground-truth datasets of full user profiles
 - Both fake (Sybil) and legitimate user profiles
 - 人人 – given to us by Renren Inc.
 - Facebook US and India (Crawled (only publicly accessible data))
 - Legitimate** profiles – 2-hops from our own profiles
 - Suspicious** profiles – generic profile images
 - Sybil** profiles – Banned suspicious profiles

Dataset	# of Profiles			Test Group	# of Testers	Profile per Tester
	Sybil	Susp.	Legit.			
Renren (China)	100	0	100	Chinese Expert	24	100
				Chinese Turker	418	10
Facebook US	32	117	50	US Expert	40	50
				US Turker	299	12
Facebook India	50	101	49	India Expert	20	100
				India Turker	342	12



A Crowdsourced Sybil Detection System




Advantages

- Scale to many millions of users, low relative cost
- Extremely high accuracy
- Limit information exposure when giving data to turkers

Cost Estimation

- Estimated cost in a real-world social networks: Tuenti
 - 12,000 profiles to verify daily
 - 14 full-time employees
 - Minimum wage (\$8 per hour) → **\$890** per day
- Crowdsourced Sybil Detection
 - 20sec/profile, 8 hour day → 50 turkers
 - Facebook wage (\$1 per hour) → **\$400** per day
- Cost with malicious turkers
 - Estimate that 25% of turkers are malicious
 - 63 turkers
 - \$1 per hour → **\$504** per day



Supported by
CNS-0916307
CNS-1224100