RobocallGuard: Fighting Voice Spam with a Virtual Assistant

Sharbani Pandit

Georgia Institute of Technology

Joint research done with Jienan Liu (UGA), Roberto Perdisci (UGA) and Mustaque Ahamad (Georgia Tech).

Internet Threats Are Moving to Telephony

- The phone voice channel is largely unprotected (phone spam via robocalls, phishing etc.).
- YouMail estimates that March 2019 saw 5.23 billion robocalls in the United States.
- Do call blocking applications (True Caller, YouMail etc.) work?

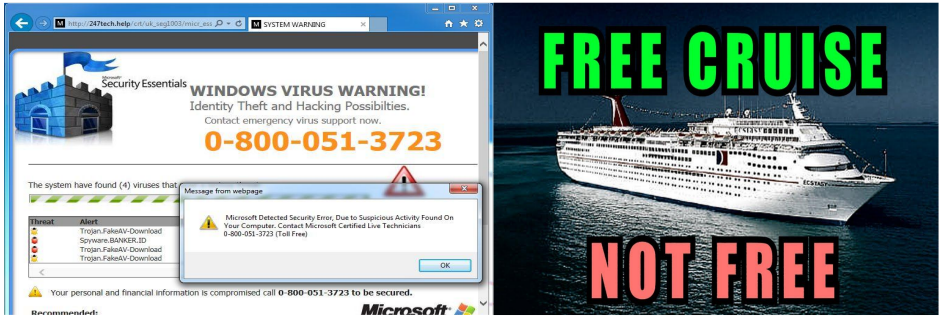

Motivation

- Blacklists rely on historical data such as user complaints or honeypot generated information.
- Blacklists can only be somewhat effective against scam calls.
- Effectiveness further degrades with caller ID spoofing.
- Hiya reported that 56.7% scams reported by their users relied on neighbor spoofing.

SHAKEN/STIR

- FCC proposed caller ID authentication protocol.
- Caller ID authentication will likely reduce the number of robocalls.
- All carriers need to support SHAKEN/STIR for it to work on a large scale.
- What if most of the scammers are operating outside of US?
- What if victims continue to fall for scams that don’t use caller ID spoofing?
- Not a cure-all solution.

Research Goal

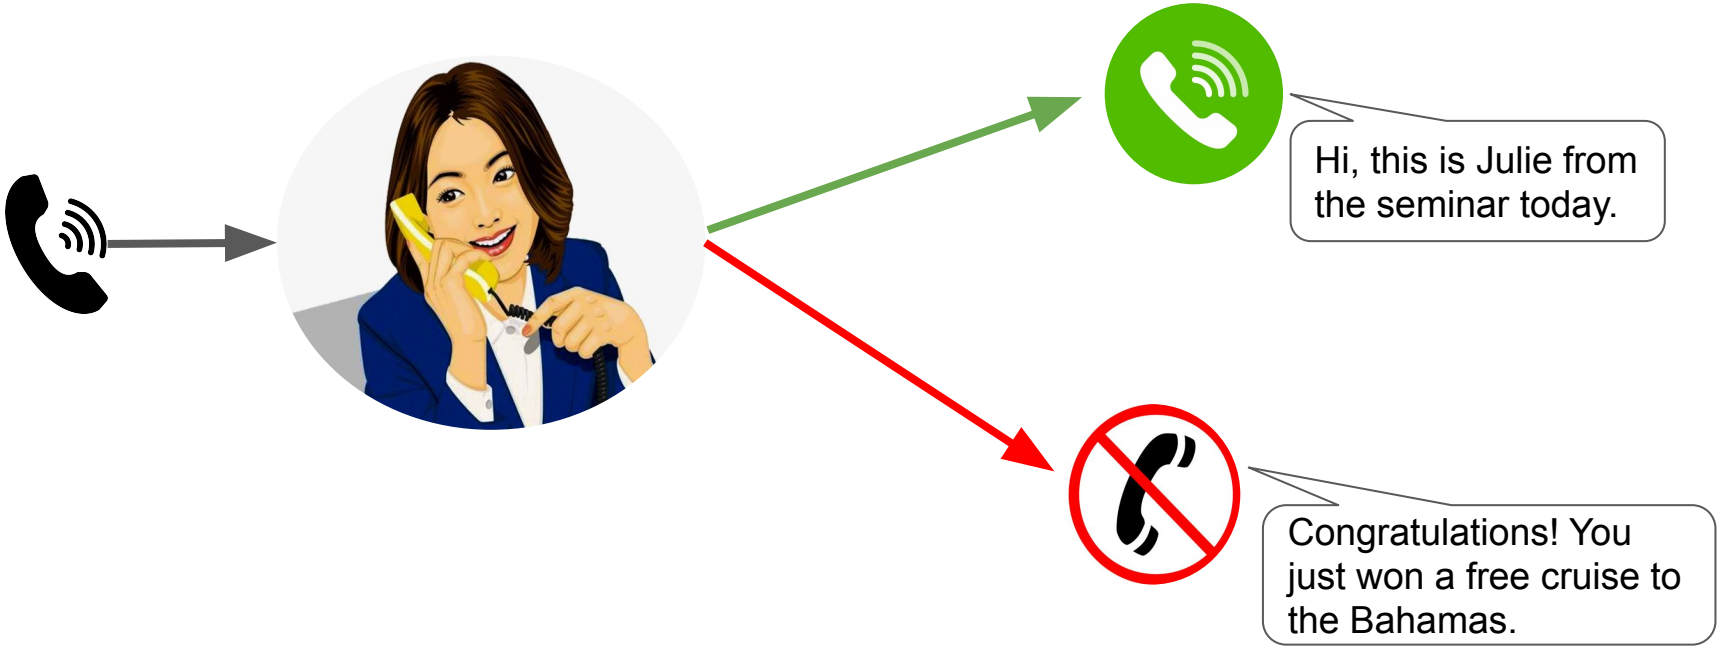

Current call screening methods

- **Google call screen**

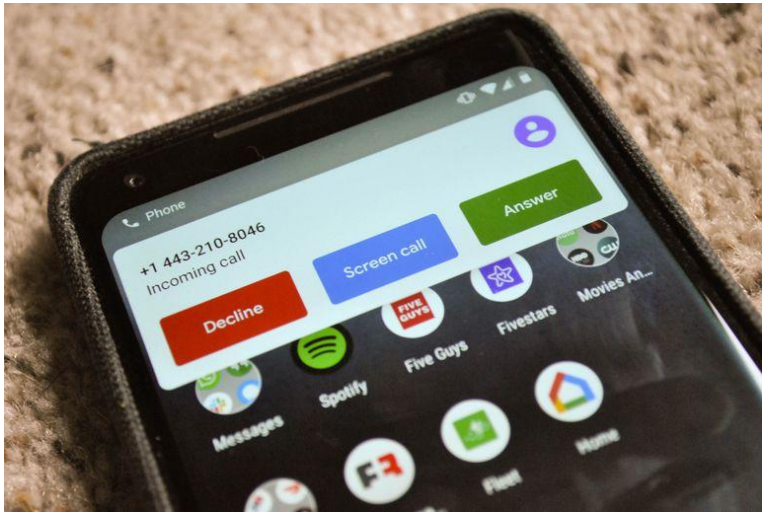

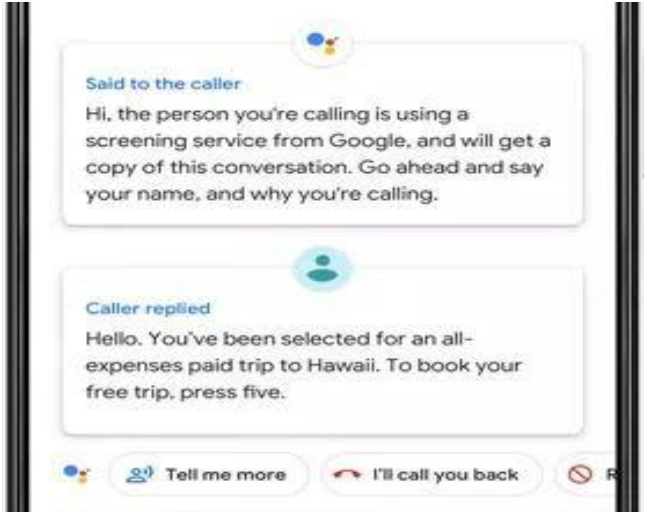

Current call screening methods

Robokiller

- All calls are forwarded to a centralized server.
- Performs audio analysis to detect robocallers.
- All human callers are forwarded back to the user.
- All rocallers are passed to a honeypot.

Contributions

- We are the first to evaluate a call screening virtual assistant that defends spam calls with caller ID spoofing.
- Our virtual assistant aims to provide a similar, albeit much less “sophisticated”, personal secretary.
- We have developed a proof-of-concept smartphone app named “RobocallGuard” and blocked 100% robocalls.
- We conducted an IRB approved user study to assess the usability of our virtual assistant.

Design Goals

- Add an extra layer of security between the caller and the recipient of the call.
- Provide additional information to the recipient about the content of the call.
- Preserve user experience in regards of latency and accuracy.
- Ensure privacy when an incoming call is handled by the VA.

System Overview

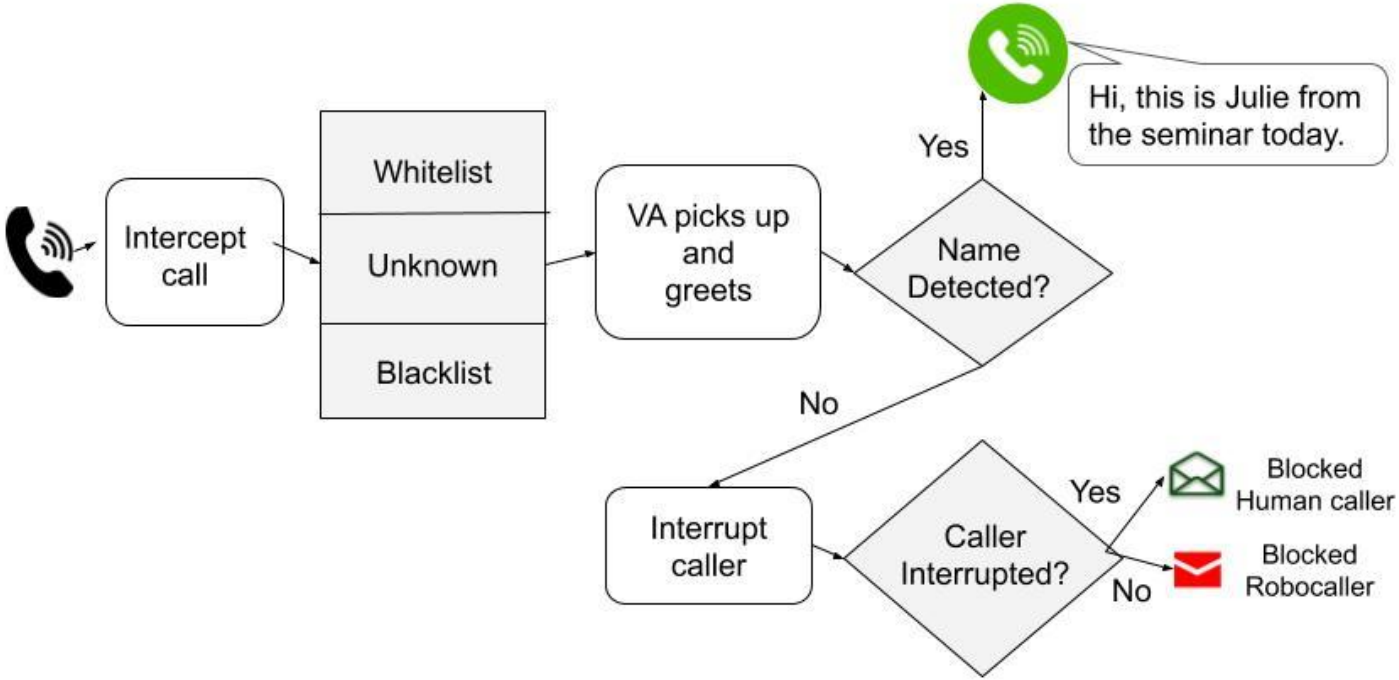

Threat Model

- **In scope threats**
  - Mass robocalls
  - Mass unwanted live calls from humans (telemarketers, debt collectors etc.)
  - Spoofed calls
  - AI equipped attack
- **Out of scope threats**
  - Targeted attack
  - Landline calls

System Architecture

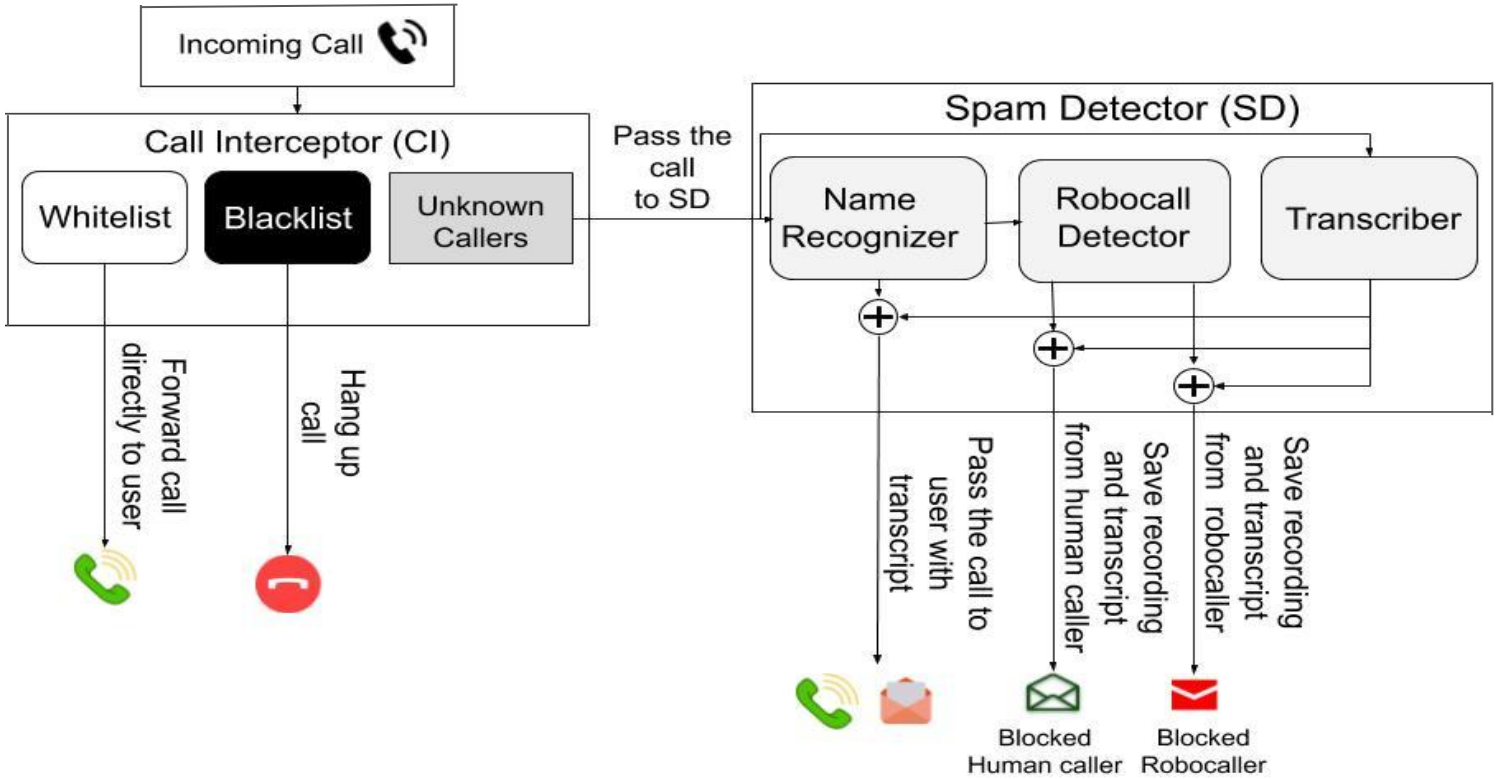

Call Interceptor

- Initial decision based on the caller ID.
- Whitelisted calls → passed ; Blacklisted calls → dropped
- Intercepts calls to acquire the incoming audio stream.
- Injects recorded voice messages by the VA into the outgoing audio stream.

Name Recognizer

- Caller knows the name of the callee → wanted; Caller does not know the name of the callee → unwanted.
- User is allowed to set the correct name(s).
- A keyword spotting algorithm to detect the correct name(s).
- **Snowboy**: light weight DNN-based customizable keyword detection engine.
- Can be trained with limited audio recordings.

Robocall Detector

- Calls from unwanted callers are handled by the Robocall Detector(RD) module.
- VA interrupts the caller at a predefined timestamp (20th second).
- Uses Voice Activity Detection(VAD) to determine if the caller became silent.
- Caller interrupted → Human ; Caller not interrupted → Robocaller

Transcriber

- Provide context for both wanted and unwanted calls.
- Google Cloud Speech API: very high accuracy and transcription can be performed from a mobile device.
- Privacy issues.
- Android system restrictions.

Implementation

- We envision our VA to be embedded with the Phone app.
- Current Android system restrictions do not allow injecting voice messages in the outgoing audio stream without OS modifications.
- We implemented a VoIP app to conduct our user study experiment.

Usability Study

User Study Setup

- Experiment 1: Caller was provided with the correct name.
    - Caller is given a script: “Hello, can you please forward my call to Taylor?"
    - Caller is instructed to make a call to their friend Taylor to make movie plans.
- Experiment 2: Caller was provided with incorrect name/no name.
    - Caller is given a script: “Hello, can you please forward my call to Robert? We met at a seminar today."
    - Caller instructed to make a call to an office trying to sell a computer.

User study results

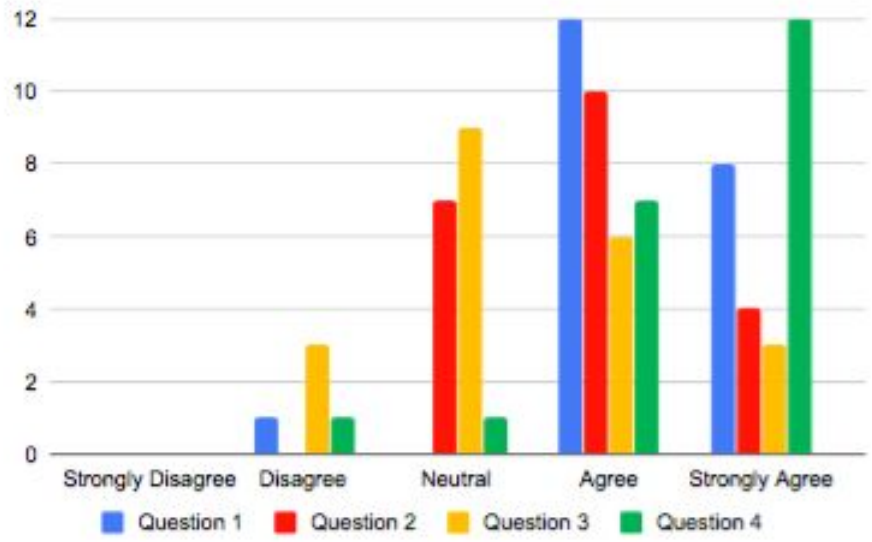

Caller and Callee Response

- Q1 : It was easy to interact with the VA.
- Q2 : The delay you experienced before the other person responded to the call is acceptable.
- Q3 : The transcript was able to provide sufficient information to infer the topic of the incoming calls.
- Q4: The transcript was able to provide sufficient information about the content of the blocked calls.

User study results (cont.)

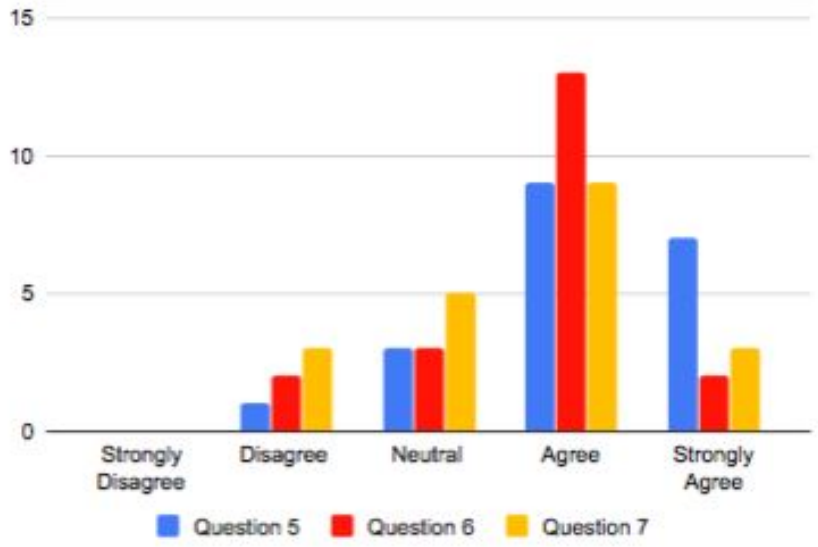

Generic Response

- Q5 : I found the app beneficial to me as it provides prior knowledge about the incoming calls.
- Q6 : I think I would like to use an app equipped with a VA frequently.
- Q7 : I felt comfortable with the Virtual Assistant intervening in the phone calls.

VA Performance

Interaction with Human Callers

- **Legitimate Callers:**
  - Four experiments with each of the 21 callers.
  - In the first two experiments, the callers were given the correct name.
  - All of them detected as legitimate callers.
- **Unwanted Callers:**
  - Callers were given an incorrect name/no name at all.
  - All the calls were stopped.

Interaction with Robocallers

- 8081 calls coming into a telephony honeypot during April 23 and May 6, 2018.
- Topic modeling + DBSCAN to create 79 clusters of robocalls.
- One call from each cluster.
- All calls are detected as unwanted.

Robocall Labeling

- Robocalls > 20s labeled accurately.
- Short robocalls were mislabeled.
- 86% of the short robocallers ask the callee to press or enter a digit in the phone
- If a call contains the keywords "press" or "enter", we label it as a robocall.
- The VA is able label 97.8% of all robocalls correctly.

Comparison with other apps

- We used a Twilio phone number to make 10 robocalls to RoboKiller.
- Robokiller was not able to block any of the robocalls, RobocallGuard blocked all 10 robocalls.

- We spoofed 10 most popular phone numbers from FTC (August 3-5, 2019).
- On Sept. 10 Robokiller blocked 9/10 calls.

- We spoofed 10 least popular phone numbers from FTC (Sept. 9, 2019).
- On Sept. 10 Robokiller blocked 2/10 calls.

Related Work

- Telephony Honeypots [1]
- Domain and IP blacklisting [2]
- Email spam filtering [3]
- Phoneprinting clusters of campaigns [4]
- Tech Support campaign [5]

Discussion

- It only works when the callee has a smartphone.
- RobocallGuard is not designed to protect users against targeted attacks.
- Bad actors can obtain names associated with phone numbers from leaked data.
- RobocallGuard is effective against robocallers currently making cheap mass robocalls.
- It can stop unwanted live calls that come from human callers.

Conclusion

- We have proposed RobocallGuard, a virtual assistant system which can be effective against mass robocalls including spoofed calls.
- Our proposed VA handles all incoming calls and users are no longer interrupted by unwanted calls.
- We have developed an Android prototype app hosting the VA.
- The VA is able to block all robocalls and label 97.8% of the robocalls correctly.
- The user study shows that most users are comfortable using the app and believe that the application is beneficial to them.

Thank You