SaTC 2019 PI Meeting Breakout Session

AI, ML, and NLP for Personalized Privacy and Security Assistants

Co-leads: Norman Sadeh (CMU) and Shomir Wilson (Penn State)

Problem/Domain Summary

- **Expertise and effort required to manage one's security and privacy far exceed people's ability and time**
- **Lack of awareness and understanding**
  - People fail to follow security and privacy best practices – they often don't know what they are and don't appreciate the consequences
  - Too much to learn, too little time
- **Security and privacy as secondary tasks**
- **Why is it important to society?**
  - People unknowingly put themselves, their families, and their organizations at risk
- **Long-term significance**: **These challenges will only get worse over time**

  - Similar challenges exist for other stakeholders (e.g. developers/publishers, app stores/ecosystems, regulators): All these stakeholders would benefit from the availability of shared services (e.g. shared repositories of vulnerabilities, breaches, descriptions of new technologies and new data practices, shared models of privacy preferences such as preferences organized in the form of clusters, etc.)

Key Research Challenges

- Each individual is different
  - Can we build models of what people know, what they care about, what they want to be informed about and how/when, how much assistance they need, how, etc.?
- How to configure assistants to
  - Reduce user burden – being helpful rather than annoying
  - Without taking away sense of control/autonomy – some differences between security and privacy
  - Different people might prefer different interaction styles
- Understanding what are the main threats
  - Constantly evolving, etc.
  - Many sources of information – including many textual ones
- Building shared user models while respecting people's privacy
- Evaluation "metrics"/dimensions – multi-faceted

Promising Avenues

- **Privacy Assistants**
  - Selectively notify us about data practices we may not expect/are uncomfortable with, help us configure privacy settings, nudge us to more carefully consider privacy decisions, etc.
  - Requires user models (what people expect, what they know, what their preferences are, what they would want to be notified about, how often and when, etc.) – personalization aspect
  - NLP work to understand privacy policies and other relevant documents and supporting dialog functionality
- **Security Assistants**
  - Could also selectively notify us about risks we may not be aware of or may be overlooking, help us carry out some tasks securely, nudge us to adopt best practices
  - Again based on user models – be helpful rather than annoying, assist us with tasks we (end-users) are not good at (from detecting phishing emails to helping us carry out everyday tasks more securely)
  - Requires ML and NLP but also AI planning, etc.

- **Multi-party computation & differential privacy** – to crowdsource user models (e.g. privacy preferences)