Reviewer Comments

SBE TTP: Medium: Securing Cyber Space: Understanding the Cyber Attackers and Attacks via Social Media Analytics

NSF Award #1314631

Context Statement

NATIONAL SCIENCE FOUNDATION
Secure & Trustworthy Cyberspace Program
General Information for Secure & Trustworthy Cyberspace Proposal Applicants

This year the Secure & Trustworthy Cyberspace (SaTC) Program (Medium) reviewed 145 proposals. Each proposal was considered by a panel of experts along with other proposals pertaining to similar topics. In all, about 8 separate panels met to consider the SaTC proposals. The proposals were discussed during the panel meetings, and where necessary, ad-hoc reviews were submitted. Panel recommendations to NSF were "Highly Competitive," "Competitive," "Low Competitive," or "Not Competitive." Conflict-of-interests were very carefully monitored so that researchers who had submitted a proposal to a particular Secure & Trustworthy Cyberspace category (Small, Medium or Frontiers activities) were disallowed from participating in the review of any Secure & Trustworthy Cyberspace proposal in that category. Researchers who had a statutory conflict with a lead proposal in a panel reviewing team proposals were disallowed from reviewing any proposals in that panel. Panel members or panel moderators with other conflict-of-interests left the room and did not participate in the discussion when the proposal with which they had a conflict was discussed.

Panel Summary

This proposal focuses on understanding hackers and develops an open-source longitudinal testbed to track online discussion forums that are popular with hackers. It adds to social science knowledge related to computational linguistics and tracking social networks of hackers, particular the "most effective" hackers.

Intellectual Merit
The testbed is likely to be uniquely useful, particularly since it may allow the study of a complicated population that can not be reached any other way. It would also develop a corpus of data, that could prove useful to criminologists among others. There is significant ambition also in looking at a broad number of settings across multiple languages. This internationalization of the scope of the project is a significant strength. The panel generally though the proposal was well thought out, and can be useful.

Looking at the problem from the attacking perspective, with a strong team, is likely to have significant benefits. One weakness is the lack of theory on social interactions. The detail of the

textual data that will be compared, allowing for the potential probing of hacker motivations is unique and important. There was some concern on ethics: will the PIs anonymize data? Some are taken from public chats, but others are password protected. Other panelists suggested that this is likely to be less of a problem if a "cut down" version of the data is made publicly available.

Broader Impact
Can improve security, with a unique perspective from hacker point of view. The benefits of possibly understanding hacker motivations is a key impact.

Relevance to SaTC
Highly relevant.

Suggested Improvements
Could add additional social science, along with a more detailed discussion of IRB-issues outside of data management plan.

Option (TTP)

Conclusion
A very strong proposal with intellectual merit and broad implications. It was very unique and exciting to many panelists.

Panel Recommendation
Highly competitive

Note: The summary was read by/to the panel and the panel concurred that the summary accurately reflects the panel discussion.

PANEL RECOMMENDATION: Highly Competitive

Review #1

Rating: Excellent

Summary

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to intellectual merit.

This proposal has significant intellectual merit in creating a longitudinal archive of hacker discussion forums and documents, along with a tool suite for network and linguistic analysis of the discussions. I can imagine the longitudional research testbed could substantially advance understanding not only of cyberattacks but also of the community and culture of hackers. As well, the proposers are building off their prior Dark Web work of suspected terrorist networks. The research team has substantial prior experience, and a sensible plan for carrying out their vision. The proposal seems appropriately resourced, and the PIs have existing institutional

infrastructure to also rely on.

The concern I have relates to the ethics of archiving public discussions that are inaccessbile to the public without permission of the PIs. One concern is that it may be the case that some of the public discussants are not anonymous, or their identities can be determined through some fairly simple analysis. What will the researchers do with non-anonymous posts in their archive? Second, if the materials were originally public, why would they be protected or withheld from the public in the proposed plan for making the data available? If that is to protect the people whose discussions were archived, how are their identities protected? I noted that in the discussion of collecting materials that the PIs would use connections to gain access to forums that are private or require permissions to gain access. Typically, human subjects review panels view forums that are password protected or require permissions to become a member as not public, and therefore consent from participants must be secured. How will the PIs grapple with these ethical issues?

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to broader impacts.

The proposed archive and tool kit for analysis could be valuable to the practitioners and policy makers. The proposal highlights efforts to train underrepresented minorities, and will actively reach out to practitioner communities.

Please evaluate the strengths and
weaknesses of the proposal with respect to any additional solicitation-specific review criteria, if applicable

The proposed archive supported the goals and objectives of the Secure and Trustworthy Cyberspace initiative.

Summary Statement

Overall, I think this is an important proposal from well-qualified researchers. I wish they had engaged with the ethical questions of the proposed archive, but I will have to trust that their human subjects review boards will help them contemplate the impact on human subjects in creating the archive.

Review #2

Summary

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to intellectual merit.

This proposal is focus on collecting and then data mining information on underground criminal forums and chat networks for the purposes of driving criminology analytics. There are few

publicly available corpus' of this kind and acquiring and curating one would likely spur considerable broader research. The Arizona team has a history in collecting such data (for Jihadist sites) and Holt is fairly well known on the criminology side of cybercrime circles and has sufficient domain expertise to guide them to appropriate forums and analyses.

Much of the infrastructure for this project has already been created (Dark Web Crawler, the basic analytics, etc) funded under previous efforts focus on Jihadist sites and it is not clear if there will be significant new requirements for underground forums. Thus, given the SBE label, it would have been nice to see more of the proposal devoted to what hypotheses/analyses would be explored against this data set.

Also, given the TTP thrust, I would also have liked to see a commitment to share the forum data in a raw forum instead of simply via a portal (the portal interface inherently limits the kinds of questions a researcher can ask). I would have thought this was due to IRB issues, but the proposers seem to imply that there is no IRB restriction on their collection and dissemination (this is particularly surprising given that there is an attribution goal in the research framework).

Finally, I would drop the bit on Malware -- which is extraneous, will not manifest in the forums/IRC channels the propoers suggest, and is easily found via other providers (e.g., threatexpert) -- and focus more clearly on issues like sentiment analysis, social network construction, etc.

These criticisms aside, I suspect that the mere creation of this dataset and sharing it (in any form) would be quite helpful in encouraging broader researcher into the social science aspects of cybercrime.

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to broader impacts.

I think there is a great broader impacts story here. Cybercrime is extremely costly and its organization is largely via underground networks such as those described by the proposers. Better understanding the nature of these networks and how they inform the strengths and weaknesses of cybercrime threat actors is clearly of importance for plicy makers and law enforcement alike.

One note I would make is that law enforcement already has such resources and, to some extent, analytics internally (e.g., SOCA's "Beast", FBI's DB, the NCFTA databases, etc) and it would probably be useful for the proposers to contact such organizations (NCFTA is probably the most open) to understand what can be learned from what is already done and what kinds of research will be most likely to have "broader impacts" on LE.

Please evaluate the strengths and
weaknesses of the proposal with respect to any additional solicitation-specific review criteria, if applicable

Summary Statement

I think this is a fine proposal. It is not high on creativity (i.e., much of the technical work has already been done) but has the potential for high impact because such data is not widely available and there are few (if any) cookbook analysis tools that have been tuned for this kind of data. If successful, I think the broader impacts could be quite significant.

Review #3

Rating: Very Good

Summary

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to intellectual merit.

Very good intellectual merit. Community and network structures among hackers would be studied. These are important topics. The project team proposed a much- needed approach to studying the attacking community. They plan to both investigate attackers' social structures, and to develop analytic tools and data sets that other researchers can employ in the future.

Project P.I. & Team generally very strong; if really wish to serve broader social-science interests in culture, communities, networks, identities, might add team member from more mainstream (i.e., not Criminology) social sciences. There is, for example, substantial amounts of literature on culture in the cross-cultural psychology and the organizational-behavior (OB) literatures on national culture and organizational-culture. OB also has a large literature on influences on and implications of identity, as does the social-psychology social-identity area.

There is a strong institutional support base, along with xcellent application-relevant systems, tools and operations developed from prior projects available for the proposed new work.

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to broader impacts.

Very good broader impacts are proposed. The team will begin transitioning past work (largely on terrorist networks) and proposed new to security practice. They will develop database & analytic tools and make available to social scientists, as well as cybersecurity experts.

Please evaluate the strengths and
weaknesses of the proposal with respect to any additional solicitation-specific review criteria, if applicable

Summary Statement

A consistently strong proposal on all elements.

Review #4

Rating: Excellent

Summary

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to intellectual merit.

The proposed activity has the potential to advance knowledge and understanding across multiple fields, from computational science to criminal justice research. The proposal presents a proactive cyber security attribution strategy, in contrast to existing reactive approaches to infrastructure security.

The proposed activities suggest and explore creative, original, longitudinal data collection and analytical tools The proposed activities include collecting longitudinal data in very creative ways (E.g. developing hacker forums, deploying honeypot platforms.

The plan for carrying out the proposed activities is well-reasoned, well-organized, and based on a sound rationale. The connection to the social science research on the hacker community is very important. The proposed analytical tools, including for instance, peer to peer communication and social networks analysis of participants within hacker networks, are appropriate and valuable.

The plan incorporates mechanisms to assess success, such as annual software releases and feedback from the broader user community.

The individual, team, or institution are well qualified and the resources are adequate to conduct the proposed activities. The project leaders, who are experts in information systems, artificial intelligence, engineering and criminal justice, will capitalize on well published previous work in terrorism informatics through computational Dark Web research and previous projects mining crime data (funded by NSF).

In the context of the five review elements, please
evaluate the strengths and weaknesses of the proposal with respect to broader impacts.

The proposed activity has the potential to advance cybersecurity and other desirable societal and research outcomes. The open source analytical data on hacker communities and virtual communications has great potential to advance the field of cybersecurity and will be made available to other scholars for hypothesis testing and other analyses.

A broad impact will be secured through annual software releases, incorporation of feedback from the broader community, and through user studies at each development phase.

Please evaluate the strengths and
weaknesses of the proposal with respect to any additional solicitation-specific review criteria, if applicable

Summary Statement

This is a very ambitious project with great chances to contribute to research and policies on cyber security, criminal justice, and related fields. The proposed activities aim to explore cyber criminal communities. The plan is to analyze the discussions interactions and exchanges in such communities. The project plans to develop a framework for collecting computational social media communications to analyze communities of cyber attackers, accounting for language and cultural differences. The proposed activities include collecting longitudinal data in very creative ways (E.g. developing hacker forums, deploying honeypot platforms. The project leaders will capitalize on successful previous work in terrorism informatics through computational Dark Web research and previous projects mining crime data (funded by NSF).