Cybersecurity Public Health: Institutions and Disciplines

Adam Shostack
The CyberGreen Institute
adam@cybergreen.net

Revised draft, submitted to NY Academy of Science, special issue on Public Health Methodologies for Societal-Scale Digital Platforms

Abstract: Cybersecurity makes extensive use of a small set of paradigms, including warfare, criminology, and to a lesser extent, medicine (such as computer viruses or hygiene.) The admirable goal of adopting public health methodologies for societal scale digital platforms can be helpfully accelerated and enabled by a broader movement for institutions and disciplines of cybersecurity public health. We outline some of the steps that will be needed.

Introduction

The lone defender stands on a parapet, deploying the tools of their profession against the raging hordes - armed not worth swords, but with DNS blocklists and anti-virus software. This “hero” model of security is popular with many defenders, despite the correlation with burnout. But even more important may be that the many talented and enthusiastic defenders are not clearly keeping attackers at bay. Cybersecurity issues are impactful enough that the US has both an agency (CISA) and a National Cyber Director in the White House.

The current dominant paradigms for cybersecurity are war and crime. Both capture aspects of today’s conflicts, with nation states using offensive capabilities apparently indiscriminately, and criminals extorting billions of dollars via ‘ransomware’ and threats to knock organizations offline with massive “denial of service” attacks. These paradigms tend to focus on individual platforms and threat

reduction, with less thought given to populations. By analogy, current paradigms focus on treating the victim, without consideration of the underlying causes or the incidence across the population.

The Academy's dedication of a special issue is an exciting development for how we think about societal scale digital platforms. Those platforms have many documented negative effects on society ranging from enabling bullying and harming the mental health of young people to state-based, state-sponsored, and criminal groups tearing apart societies and causing pogroms and genocide. I expect these are covered well elsewhere in this volume. Some of these problems appear prominently on digital platforms, and some definitions ("social media bullying") constrain the problems to the platforms, but bullying predates computers, never mind social media. Therefore, it makes sense to have fuller context around these problems, and that context includes cybersecurity problems. For example, phishing attacks that lead to stolen credentials are a classic security problem. Some phishing attacks will steal credentials so that attackers can steal information to use in bullying. What fraction of bullying involves that step, or what fraction of phishing has that goal? Without context, these questions are hard to answer. Similarly, cybersecurity mechanisms are used to constrain research. These platforms have threatened to sue researchers, blocked programmatic access by data-gathering code, and otherwise interfered with attempts to understand their impacts on society.[1]

The harms carried or even facilitated by social digital platforms are clearly of global concern, and many of these harms manifest in the physical realm. (There is a community who studies cyber-physical systems, which can amplify these issues, for example, shutting down pipelines or other critical infrastructure.) There are additional types of harms which are being studied by a nascent community under the rubric of Cyber Public Health, in which we seek to bring the tools and methods of the public health discipline to cybersecurity.

Applying a public health approach to cybersecurity will usefully expand other studies on effects of social media on mental health from platform-centered to population-centered. While there are clear advantages to single-platform research, including not needing to validate cross-system identity, many people are on many

platforms, and studying the impact of those platforms on society requires society or population level views.

Crossing these research streams will enable a fuller view of populations and harms and enable better interventions.

A discipline of Cyber Public Health will orient around measurement and interventions to demonstrably reduce harms to populations. Those populations might be computers or a subset such as routers, mobile phones, or home IoT sensors, and they might be accounts, such as gmail accounts or work or school accounts. They might be centered on people or institutions. (For example, do organizations that celebrate Cybersecurity Awareness Month have fewer incidents in the next quarter than those who do not? See for example Woods and Seymor, 2023.[2]) Harms might include that a computer is unusable because of ransomware, or that an account is inaccessible because it has been locked out.

Doctors and public health specialists use a frame of 'activities of daily living' to consider the impact of disease on daily activities such as preparing meals or bathing. Being locked out of one's Facebook account has an impact on both digital activities of daily living and also on social connection and community. The growing use of social platforms for commerce means that a lockout from an account can severely damage small businesses as well. In computer security terms, both ransomware and account lockouts are availability problems, and correcting either is "toil," meaningless tactical work which is devoid of enduring value.[3]

It may appear that cybersecurity is awash in useful data. However, that data is rarely grounded in populations. Grounding data in a populations is an excellent tool for contextualization. It enables comparison between various datasets. The difference between knowing there were 2 deaths from measles and 2 deaths from measles in a population of 100,000 puts the first number in context. Similarly, while it's tragic in either context, there's a difference between 2 measles deaths in Canada (pop 40m) and 2 in the US (pop 347m). That difference shows the relevance of "incidence," which is a measurement of new cases in a given population.

By way of example, we can examine a recent, typical report from a cybersecurity vendor. They present data on incidents which seems quite interesting. To keep attention on the data presented, the vendor has been lightly anonymized.[5]

Data point	Questions prompted by a public health analysis
"84% of the dataset was derived from organizations with fewer than 1000 employees";	What fraction of organizations have fewer than 1000 employees? How does this relate to the distribution of their customers? Perhaps they have more customers in countries with smaller firms?
the manufacturing sector was the most likely to request [our] response services, though the percentage of customers hailing from Manufacturing decreased from 25% in 2023 to 16% in 2024. Education (10%), Construction (8%), Information Technology (7%), and Healthcare (6%) round out the top five.	How do these compare to the vendor's customer base? Do manufacturers represent 32% of their customers or 8%? One would indicate good security, the other poor. How do these numbers relate to the size of sector by count of firms? For example, does their data under-represent in healthcare, where most organizations are large?
(Not provided)	What fraction of their customers experience the sort of problems which result in the report data?

The lack of grounding in a measured population limits our ability to ask questions like "is this getting better or worse?" The vendors might argue that such data is confidential for commercial reasons. Assuming that's true, that leaves many scientific questions outside the space of what can be addressed by these published data. Additionally, a great deal of the published data is surveys, and surveys focused on opinions, such as "are attacks getting more sophisticated" (without defining that) or "is your budget sufficient?" This is not to argue that those surveys are worthless. Rather for important public health data like the spread of communicable disease, we have reporting mechanisms, not surveys. For a longer discussion of the value of published data, see Chapter 3 of my *The New School of Information Security* (Addison Wesley, 2008).

There are also many tools, including Shodan, Shadowserver, and more that show vulnerability data, but they don't share data on a "per population" level.

Effective interventions are a key goal of the public health discipline. From campaigns to stop smoking or to vaccinate against a disease, is what we're doing having the effects we want? If not, what can we change? Similarly, with Cyber Public Health, are the domain takedowns, prefix blocklists, and other interventions effective?

Assessing interventions is easier when we have population data and incidence data because we can assess if the change is due to a change in population or our intervention. If we have several populations, and we intervene with one, then changes that happen in several are less likely to be a result of our intervention. Obviously, we can get more nuanced with statistical techniques that are beyond the scope of this essay.

Important questions to be answered

In the world of human health, it's reasonably clear what longer, healthier lives are and that we want them. Measuring deaths in and across populations has shown to be a staggeringly useful tool for helping people live longer and healthier. (In the 20th century, average lifespans in industrialized countries roughly doubled as a result of better understanding of why they were dying.) We treasure each life and seek to ensure that we all can pursue happiness. When a life ends, there is loss. The technological equivalent is obviously less meaningful, but not without impact. Irreplaceable photographs, messages and more are stored in devices and accounts. But that usual state doesn't mean that cybersecurity problems never represent threats to life, liberty or pursuit of happiness. McGlave and colleagues found that ransomware attacks increase in-hospital mortality for patients[4]. Many news reports have covered how people have been arrested as a result of bad matches by AI systems. And the impact of social media on self-esteem and health, especially for young women, are well-known.

Public health relies on measurements of populations and what happens to them, with a particular interest in what happens after an intervention.

Cyber Populations

The range of questions that might be better answered with population forces the question of what is population in a cybersecurity sense?

Measuring the number of people in a place is easier than measuring population in a technological realm. This is an area where research will be needed. For example, we could measure devices? That leads to the question of what qualifies as a device now? "General purpose computers with keyboards" used to be the vast majority, but those have been replaced by the computers which fit in our pockets; both are outnumbered by special-purpose computers ranging from lightbulbs to jet engines. These are often lumped under an "internet of things" label. Some of those things run Linux or even Windows, others run special purpose operating systems. In what ways does grouping enable or inhibit our understanding of the world?

Devices and accounts enable us to consider harms. With modern cloud systems storing our photographs and documents, replacing a Chromebook is a matter of cost and e-waste, not an occasion for grief. But not all machines are Chromebooks: some have local data that we don't share or even want to share. So in another sense, the lifetime of those devices may be worth measuring. Are they lasting longer? If not, what are the causes of them being retired? Is it a cybersecurity consideration? For example, Microsoft has announced an "end of support" date for Windows 10, while Windows 11 requires better hardware. So many people and businesses will face a dilemma of either replacing those computers or accepting a degree of insecurity. And while that may seem bad, we have few ways to gather data on its impact.

Measuring by accounts has similar problems, and space limits our ability to discuss those. See the report from the Inaugural Workshop On Cyber Public Health.[6]

IP addresses might seem like, or even be a good proxy for population, but most homes are behind network address translation (NAT) services. Understanding the impact of that layer of indirection is not obvious. For some definitions of population, it is possible to write multiple queries in the Shodan search engine (For example, http.title:"hacked by" and then a second query with city:"Seattle" appended). That contextualization is rare today.

Cyber Harms

Accounts are locked or eliminated as we leave school or a job. But other accounts are the key to those photographs and documents, and so it being compromised, or even us being locked out of them can be deeply impactful. This is similar to a health outcome we seek to avoid, and understanding how often it happens may be an interesting measure of the 'health' of a system provider. To illustrate with specific examples, does Facebook or Twitter have higher rates of lockout? Overall? Due to password stuffing? What about Yahoo Mail vs Outlook.com vs Gmail?

And so while, as people, we don't care about these events in the same ways, these 'deaths' of systems are frequently an effect of security problems which we hope to prevent via interventions. More generally, while considering harms to humans isn't simple, it is somewhat easier than considering harms in the technical realm. Which of these do we care about: "phishing," unauthorized access to a bank account, or unauthorized transfers of funds? It may be more than one. For example, many people probably care about the privacy of their balance and the counterparties even if theft is prevented.

Developing definitions and ontologies is work that informs and is informed by data collection mechanisms. Those data collection mechanisms enable certain analyses, and those analyses we can't do inform what data we might want to collect. We hope that our analyses are worth publishing, and that those publications drive further questions. A visualization of this idea is shown in Figure 1.

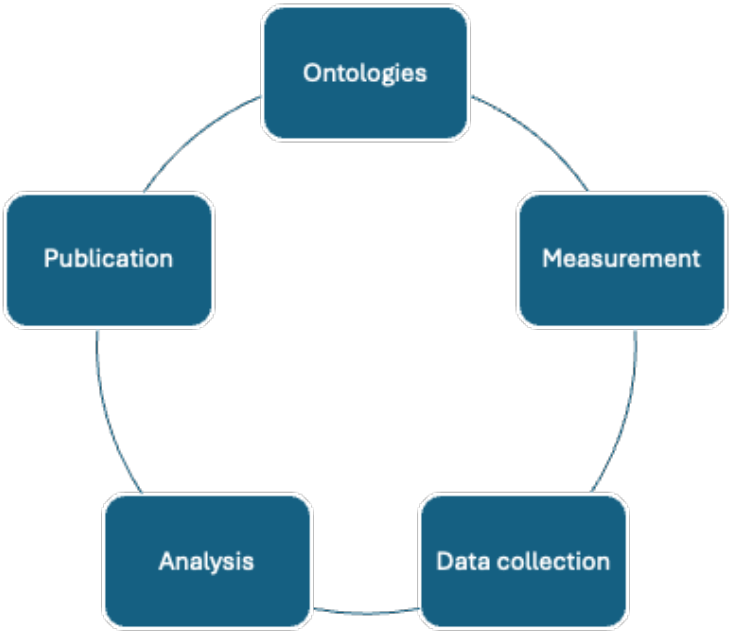

The infrastructure of Public Health

The work of public health is distributed through our society, from local health departments inspecting restaurants or engaging with communities, through state departments, all engaged with the National Center for Disease Control and National Institutes of Health. At the international level, the World Health Organization coordinates between national authorities, and there are international scientific meetings to share research and encourage collaboration. Similar concepts may be helpfully applied in the context of Cyber Public Health.

Institutions

Much of the data that will help us understand societal scale digital problems will come from companies. Contention will come not only because it's proprietary, but also over standardization of terms, collection and dissemination. Commercial companies may have their commercial interests at heart. By way of example, Microsoft may want to downplay malware, while anti-malware companies may want to exaggerate it. Both the possibility and perception may make a neutral party helpful. So aån institution that advocates for the public interest will be a useful balance. For example, in 2021, Facebook sued New York University researchers studying political ads and more on its platform, and blocked both their software

and their ability to use the platform. (Bond, 2021) Professor Damon McCoy said "they're trying to send a message to other independent researchers that are trying to study their platform." A Bureau of Cyber Statistics might be more willing and able to protect research and researchers. (This is not to judge NYU's response, but to offer the idea that a public University might have more competing priorities than a dedicated institution.)

Public Health Institutions bring together data, expertise, infrastructure and more. We need them for Cyber Public Health as well. In particular, a statistical office like a Bureau of Cyber Public Health Statistics would be able to collect and distribute authoritative population data, and be a home for other data collection work including ontologies and the forms which they support.

Disciplines and Measurements

There are degrees in, and conferences for, public health specialists. There are also many disciplines that support or engage with public health, including epidemiology and medicine. We will need to establish such disciplines for Cyber Public Health.

Specific metrics that could be used to measure the effectiveness of Cyber Public Health interventions include primary measurements of the things which most concern us, and supplemental measurements. Additionally, we can think of measures that are directly informed by public health, such as infections, and those that come from other disciplines, or the unique aspects of technology and cybersecurity, such as our ability to use logs to measure 'time to detect.' Examples include:

- Prevalence and incidence of cyber threats: Tracking the number of newly successful attacks (e.g., malware infections, data breaches) within a given population over time, as well as the prevalence: systems that are never "cleaned up," or continue sending spam.
- Severity and impact of cyberattacks: Assessing the financial loss, downtime, data loss, and other consequences of successful cyberattacks.
- Adoption of preventive measures: Monitoring the percentage of individuals and organizations implementing recommended security practices (e.g.,

using strong passwords, enabling multi-factor authentication, applying software updates).
- Time to detection and response: Measuring the time it takes to identify and mitigate cyber threats, including the time to recover from attacks.
- User awareness and knowledge: Assessing the level of understanding and awareness of cyber threats and preventive measures among the general public.
- Resilience of critical infrastructure: Evaluating the ability of essential services and systems to withstand and recover from cyberattacks.

By tracking these and other relevant metrics, we can gain valuable insights into the effectiveness of Cyber Public Health interventions and identify areas for improvement. A great deal of measurement for cybersecurity seems to have elements of looking for car keys under the streetlight, and a value we can take from public health is to assess what value we get from the measures. This data-driven approach will enable us to develop more targeted and effective strategies for mitigating cyber risks and protecting individuals and society.

A movement

The chronic problems of security represent opportunities to learn from more disciplines, and expand our array of tools for protecting people, organizations and societies. Public health offers a framework that focuses on what matters, and allows us to build proactive, preventative and measurable strategies. By establishing a dedicated Cyber Public Health discipline, complete with robust institutions and interdisciplinary collaboration, we can begin to address cybersecurity threats at a population level. This will involve the development of standardized metrics, comprehensive data collection mechanisms, and dedicated research initiatives. Through a concerted and coordinated effort, we can foster a more secure and resilient digital ecosystem, safeguarding individual and societal well-being in the face of evolving cyber threats. The time to act is now.

Acknowledgements

David Conrad, Josiah Dykstra, Yurie Ito, Arastoo Taslim, and Bill Reid provided helpful comments.

Author contribution: The lead author did the primary drafting, with support and help from those acknowledged.

Competing Interest Statement: The author has no competing interest.

References

1. Bond, Shannon, NYU Researchers Were Studying Disinformation On Facebook. The Company Cut Them Off, NPR, August 4, 2021 https://www.npr.org/2021/08/04/1024791053/facebook-boots-nyu-disinformation-researchers-off-its-platform-and-critics-cry-f
2. Woods, Daniel W., and Sezaneh Seymour. 2023. "Evidence-Based Cybersecurity Policy? A Meta-Review of Security Control Effectiveness." *Journal of Cyber Policy* 8 (3): 365–83. https://doi.org/10.1080/23738871.2024.2335461.
3. Rau, Vivek, Eliminating Toil, in Building Secure & Reliable Systems (O'Reilly, 2020)
4. McGlave, Claire C., Hannah Neprash, and Sayeh Nikpay. 2023. "Hacked to Pieces? The Effects of Ransomware Attacks on Hospitals and Patients." SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.4579292.
5. Shier, John, Angela Gunn, Hilary Wood, It takes two: The 2025 Sophos Active Adversary Report, April 02, 2025, https://news.sophos.com/en-us/2025/04/02/2025-sophos-active-adversary-report/
6. Shostack, Adam, Inaugural Workshop On Cyber Public Health, CyberGreen Institute Tech Report 24-01, 2024. https://cybergreen.net/workshop-report-24-01-inaugural-workshop-on-cyber-public-health/