

COMMENTARY



KRIS OOSTHOEK AND CHRISTIAN DOERR

Cyber Threat Intelligence: A Product Without a Process?

Cyber threats have become a permanent threat to society. Over the last few years, accounts of hacking campaigns into public- and private-sector enterprises have drawn significant attention. In 2017, Yahoo announced that three billion user account details were exposed in a hacking operation dating back to 2013. In 2018, Equifax disclosed that malicious actors had penetrated

Kris Oosthoek is a Senior Cyber Threat Intelligence Analyst with the Dutch government. He is a Researcher with the Cyber Threat Intelligence Lab. His area of expertise includes cyber threat intelligence, network security, malware analysis, and security operations. He received his M.Sc. degree from Erasmus University, Rotterdam, The Netherlands and is currently Ph.D. candidate with Delft University of Technology, The Netherlands. He can be contacted at k.oosthoek@tudelft.nl

Dr. Christian Doerr is Professor of Cyber Security and Enterprise Security and Director of the Cyber Threat Intelligence Lab at the Hasso Plattner Institute in Potsdam, Germany. His research focuses on network security, cyber threat intelligence, and situational awareness. He received a joint Ph.D. in Computer Science and Cognitive Science from the University of Colorado at Boulder, USA. He can be contacted at christian.doerr@hpi.de

© 2020 The Author(s). Published with license by Taylor & Francis Group, LLC. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

its corporate network and exposed sensitive personal data of 143 million U.S. citizens. The same year, Marriott Hotels declared that the records of 383 million guests were exposed to malicious actors. These breaches descended from state-coordinated hacking campaigns. Cybersecurity breaches cause technical catastrophes, but also have significant ramifications at economic, legal, and individual and personal levels.

In many cyberattacks the victim network is breached long before detection. The average time to identify a breach is 206 days, with the mean time to then contain it being 73 days. Cybersecurity incidents caused by malicious actors are the most common and most expensive to solve.¹ The unrecognized presence of malicious actors within the trusted enterprise network boundary effectively signifies an intelligence gap in computer network defense.

In order to close this gap, the cybersecurity community established the field of Cyber Threat Intelligence (CTI). The primary objective of CTI is to realize a knowledge advantage over cyber threat actors. At the tactical and operational levels, CTI expedites early detection of malicious behavior, preferably before a malicious actor gains a foothold in the network. On a strategic level, CTI provides sense-making and insight into the relevant threat environment to decisionmakers. Effectively, CTI is the civilian, private-sector alternative to defensive counterintelligence executed by the established Intelligence Community (IC).

The marketing of CTI-related products and services is an increasingly important revenue-generating asset for many cybersecurity vendors with roots in the production of firewall and antivirus offerings. They have re-branded the commodity activity of providing a blacklist into a “CTI” operation and generate intelligence reports using malicious activity detected in telemetry from sensors placed in customer networks. Many bigger enterprises have in-house CTI teams, focusing on threats to their network and catering to stakeholders, such as executive teams, technical teams (security operations, vulnerability management), and legal and compliance officers. CTI builds on intelligence studies, computer science, and computer security fields, such as malware analysis, computer network security, and intrusion detection. To a lesser extent, it also builds on previous work on situational awareness, risk forecasting, and risk management.

CTI is a flourishing industry. According to the Council on Foreign Relations, 352 unique state-sponsored cyberattacks have been identified from 2005 onward.² A total of 28 countries are suspected of sponsoring cyber operations, ranging from espionage to actual sabotage of critical infrastructure. The majority of these campaigns were exposed by the CTI industry. Many technical experts originally employed in the IC have moved to work for CTI vendors, such as CrowdStrike, FireEye, Talos, and

Kaspersky. Where Bellingcat has open-sourced intelligence analysis related to Malaysia Airlines Flight 17, the Syrian Civil War, and the Skripal poisoning, the CTI field has open-sourced as well as commercialized intelligence analysis on cyber threats. With its deep technical expertise and subject-matter knowledge, the CTI field holds tremendous potential to address cybersecurity threats in the years ahead. As its capabilities equal or in some cases exceed those of government intelligence agencies, CTI is a useful associate in the current cybersecurity cat-and-mouse reality.

However, CTI is also a field in its infancy. This has several consequences that we discuss in this article. We argue that CTI is a product without a process, which has several underlying causes and consequences for the CTI practice. CTI has already contributed toward increased computer security, but it has solved technology problems with technology solutions. This is a logical reflex for a field firmly grounded in computer science. However, its initial innovation has stalled and many of its challenges are not widely recognized, while the field still has problems to solve. We argue that the CTI field should not address its challenges by adding more technology. They rather need to be informed by the work on intelligence analysis and methodology, also referred to as analytical tradecraft, originally cultivated in the field of intelligence studies.

TODAY'S CYBER THREAT INTELLIGENCE FIELD

Much of the cybersecurity debate is grounded on the structural game-theoretic asymmetry between attacking and defending agents. An attacker needs only one weak point to compromise a network, while its defender needs to account for all potential weak points in the security posture of its organization. Even for a highly secured network, the system with the lowest level of defense defines the entire network's actual level of defense.

CTI strives to reduce the knowledge asymmetry between attackers and defenders through the gathering and analysis of defensive counterintelligence on cyber threats. A cyber threat is the composition of an actor with the motivation and intention to exploit vulnerable points in a victim's technology infrastructure, with a set of specific capabilities such as malicious software (malware). The scope of CTI is deliberately limited to defensive purposes, as the majority of CTI analysis output intends to mitigate intrusions or otherwise facilitate early detection.

The CTI movement gained traction with the publication of a report by Mandiant on the activities of Chinese operatives, which they labeled Advanced Persistent Threat 1 (APT-1).³ The report described intrusion operations by Unit 61398 of China's People's Liberation Army, an organization affiliated with the army, with intrusions into more than 150 organizations. As its campaigns persisted over a long time frame, Mandiant

deemed it APT-1. It was one of the first actor groups to be exposed to the public. The report was a breakthrough because it included deep analysis, and not only of the group's technical tradecraft; it was also rich in context, such as its personnel requirements, organizational placement, personal profiles of operatives, and geospatial intelligence. Together with the report, Mandiant released domain names, Internet Protocol (IP) addresses, and file hashes serving as fingerprints of APT-1's behavior.

While it can be argued that the report was put out to attract media attention, it forced other vendors to publish technical details in order to remain relevant. The Mandiant report set the standard for many reports to come. After 2013, many security vendors started conducting similar analysis on other actor groups, resulting in CTI being an industry in its own right. Gartner has predicted that 20% of large enterprises will use commercial services to inform their CTI by 2022.⁴ The commercial CTI market is expected to grow from 5.3 billion USD in 2018 to 12.9 billion in 2023, at a compound annual growth rate of 19.7%.⁵

The CTI community consists of commercial, independent, and private-sector researchers who share their analysis in various ways and forms. Threat artifacts usually come in the form of Indicators of Compromise (IoCs); tactics, techniques, and procedures (TTPs); or qualitative research reports. IoCs are IP network addresses, domain names, and file fingerprints (hashes) associated with malicious activity. IoCs are machine-readable, which enables distribution to network security devices for automated detection. A TTP describes an actor's broader modus operandi. TTPs usually also include IoCs and thus allow for contextualized analysis of malicious artifacts. Reports, like the APT-1 report, usually come in portable document format (PDF), describing threat actor groups, new threat campaigns, or new malware. Security vendors are the biggest producers of these reports, while national intelligence agencies also have released them in matters of national security. Reports tend to be highly detailed, but are not machine-readable and less suited for automated detection as a consequence.

The field has yielded analytical models, such as the Cyber Kill Chain,⁶ the Diamond Model,⁷ the Pyramid of Pain,⁸ and ATT&CK⁹ to standardize the semantics on the description and characterization of cyber threats. These models have contributed to the field's maturity, but also indicate its need for standardization and methodology, which we will demonstrate in the next section.

CYBER THREAT INTELLIGENCE: STILL CONNECTING THE DOTS

In this section, we will consider the current challenges of the CTI field. These are independent challenges that will all benefit from methodological improvement, as methodology is currently largely nonexistent. We argue that

the field of intelligence studies can provide the CTI field the insight it requires to move toward further maturation, as it has put significant effort into improving intelligence analysis.

CTI Is Lacking Methodology

The body of knowledge of intelligence studies builds on the qualitative methods from social science research.¹⁰ This is reflected in the work on cognitive bias and its impact on analysis by authors such as Richards J. Heuer and Sherman Kent. Many CTI community conference talks paraphrase their work, together with Priority Intelligence Requirements and Kent's Words of Estimative Probability.¹¹ Currently, however, the CTI field seems to be stuck at assuming this borrowed terminology. Much of CTI analysis is building on loose concepts to suggest it has analytic rigor, but in fact it does not.

While Heuer's seminal work, *Psychology of Intelligence Analysis*,¹² is often referenced in CTI conference talks,¹³ speakers fail to address how it guides their analysis and how it helps solve challenges in their daily analysis practice. The same fate befalls Heuer and Pherson's Structured Analytic Techniques (SATs).¹⁴ While the SATs form a collection of multiple techniques, many CTI talks only provide lip service to the Analysis of Competing Hypotheses specifically. Here again, accounts of operational implementation and declassified examples are scarce. On a day-to-day basis, most CTI analysis is input-driven by alerts and incoming raw data, rather than predetermined hypotheses. This is an understandable effect of a field still in its infancy and according to some also the case in regular intelligence analysis.¹⁵

Cyber threats, however, have unique native properties that make the absence of methodology painfully manifest. The volume and velocity with which new attacks are reported leads to a high daily influx of many single IoC datapoints that need further triangulation to assess their relevance to the specific threat context. It is, however, unfeasible to perform structural analysis on each single datapoint. On the other hand, the absence of a process can induce analysis paralysis, especially in smaller teams. While the computer science field has offered several machine learning algorithms that support data preprocessing, the translation of tacit knowledge into algorithms will likely remain an unresolved challenge for years to come. The solution lies in introducing process, not more technology.

It is issues like these where intelligence practitioners and scholars can contribute to a solution that has the potential to benefit both fields. The IC has an established reputation of introspection and self-scrutiny when it comes to analytic benchmarks. The work on tradecraft work by Gates¹⁶ and MacEachin¹⁷ are prime examples, as it led to graduate-level training courses on methodology for Central Intelligence Agency analysts. It also led to

analytic standards being formally encoded in Intelligence Community Directives 200 and 203.¹⁸ This work signifies an era of advancement of intelligence analysis with many lessons learned that can illuminate the CTI field.

CTI Is Shared but Hardly Being Used

Where Pearl Harbor and 11 September 2001 are significant events that impacted the IC, CTI has its own intelligence failures. These often result from an inability or unwillingness to share. If CTI is shared, it is usually not picked up, leading to intelligence failures of its own. In this section we consider barriers to sharing. Sharing is caring, but it is also currently scaring a lot of organizations.

In 2015, the Bundestag, the German parliament, discovered a malicious actor had gained access to the Parlakom network. The forensic investigation reported the IP address 176.31.112.10 as being used by the attacker to command and control the malware remotely. Details of the breach were made available to the public in June 2015. The IP address used by the attacker was also made public for use as an IoC.¹⁹ This allowed any organization to detect potential malicious activity coming from that specific IP address. Germany's domestic security agency attributed this operation to Russia's Main Intelligence Directorate.²⁰ A month before, the IP address was also reported in an analysis of malware by U.S. security company Root9B, also referring to Russian actors.²¹

One year later, the U.S. Democratic National Committee,²² the World Anti-Doping Agency,²³ and German Chancellor Angela Merkel's Christian Democratic Union party²⁴ discovered breaches of their networks. The 2016 operations were also attributed to Russian state-sponsored hacking groups. It is less known that the attacks could have been detected earlier if the IoC from the 2015 Bundestag hack would have been used. In any of the cases the attacker established network access long before discovery. However, simple alerting of a sighting of the IP shared by the Bundestag would have resulted in early detection of the malicious activity. This is what is called a Cassandra in intelligence parlance. An outsider warns for an imminent adversary course of events, but they are basically being ignored.²⁵ However, the novelty for CTI in this case was that the Cassandra (Bundestag) did not merely prophesy a "strategic surprise" with circumstantial evidence, but offered actual forensic evidence that was ignored by the organizations that were also hacked.

Sharing of CTI is further complicated by the limitations of the Traffic Light Protocol (TLP). This uses traffic light colors to indicate whether information can be shared across trust boundaries (the organization, the Information Sharing and Analysis Center [ISAC]). Red restricts distribution to direct participants only, whereas green limits disclosure to the community.

White indicates unrestricted sharing. Amber, however, is very ambiguous: share, but only within your organization, where the specific constraints can be designated by the source. Furthermore, TLP is only made for sharing between humans; it does not work for machine-based sharing of threat data. Formal standards for machine-to-machine sharing, such as Structured Threat Information eXpression, exist; however, most CTI is still shared in unstructured means. This might explain why the majority of threat intelligence sharing still takes place via unstructured formats, such as loose comma-separated values and PDF files or no standard at all.²⁶

ISACs facilitate information-sharing across industry verticals and sectors. ISACs can be a good source for free exchange of quality CTI. However, their success tends to uphold only initially, as willingness to share is dependent on ISAC size. As soon as additional participants enter the ISAC, sharing tends to fall, as participants do not want freeloaders. As mentioned earlier, this is not a technology problem, but a problem of trust.

CTI remains largely unshared. We believe that these are indeed important barriers, and each should be addressed properly. However, we would like to argue that these challenges are mostly preceded by a low level of situational awareness about the magnitude of the cyber threat. Awareness of its potential impact on business continuity usually offsets such concerns, hence implementation of cybersecurity is often given free rein after major intrusions. The lack of dissemination needs mending, otherwise CTI is rather a self-licking ice cream cone—a self-perpetuating system feeding only back into itself.

CTI Is Generally of Low Quality

Data that can serve as a potential source of CTI comes in different varieties. Its most common form is the IoC. This is an artifact like an IP address, domain name, or file hash that is related to malicious activity according to the source. IoCs are applied in the preposition that, if encountered on a network, they indicate a potential compromise. IoCs, however, do not indicate an actual compromise, as they can be false positive if a cloud infrastructure IP address is abandoned by its original malicious owner and is now used for benign purposes. Strictly taken, IoCs do not have intelligence value by themselves, as they need to be correlated against network infrastructure logging. IoCs should therefore not be regarded as CTI, as they are not a finished intelligence product. Although many refer to IoCs as being CTI, an IoC is an intermediary product that needs to be evaluated in the context of the relevant threat environment.

The most common form of IoC consumption is through a feed. This is a method of transferring the artifacts in a machine-readable, standardized format that enables them to be “fed” to various network security products to

automate detection. This way, the IoC feed has in essence become a new approach to the classic blacklisting. In configurations like this, intelligence analysis gets effectively outsourced to the CTI provider, as the CTI is not seen by a human analyst. The IT systems of an average medium-sized organization produce millions of system messages daily, of which only a very minor share is investigated by human analysts. IoC-based detection can facilitate risk-based prioritization, but this hinges on IoC quality. If a feed produces too many false positives, this will lead to alert fatigue with analysts.²⁷

Most sources of raw cyber threat data used to inform CTI are undependable. Considered against quality parameters, such as completeness, timeliness, verifiability, data interoperability, false positive ratio, and source similarity, most feeds' content skews toward one of these.²⁸ Providers of timely data tend to have a high false positive ratio. However, providers with high accuracy and less false positives tend to be unverifiable, as the original source of the data provided remains unspecified.

The market, however, has not been able to meet the fast-growing demand. This has resulted in many CTI feeds offering intelligence of which the value is difficult to estimate. A recent study of 24 open source CTI feeds has shown that some feeds report malicious activity months after the first observance and are biased toward specific countries.²⁹ While vendors were observed recycling data from other feeds, the observed overlap in artifacts was low, even for feeds tracking the same threat. This implies that none of the vendors have an acceptable coverage of the threat they are tracking.

The absence of significant overlap between feeds tracking the same threat implies that CTI vendors currently do not only have partial coverage; the slow reporting also implies a relatively limited contribution to timely threat detection. One may argue that this is also the case with the traditional IC, as significant discrepancies between the intelligence estimates of different subcontractors might exist. With CTI vendors this is, however, primarily caused by untransparent methodology and procedures, which are caused by a limited amount of deployed sensors that serve as vantage points. This is important, especially in the case of intelligence on cyber threats, as they can pivot from data centers around the globe. With a broad vector of potential various geographical locations, one must place sensors in a sensible and balanced way, distributed over various geographical locations to ensure the measurement base is adequate. With CTI vendors effectively operating as intelligence contractors, they should be held accountable toward their methodology, analysis practice, and procurement of raw intelligence from sensors.

Many indicators currently shared via CTI feeds are of low value. They are at the bottom of the Pyramid of Pain, a CTI model to indicate the quality of

different data points.³⁰ File hashes acting as a fingerprint to identify malicious files are the most shared IoC type, but have a very short-lived intelligence value. Of malware hashes, 98% are only seen for 58 seconds or less, comparing first and last time observed.³¹ This proves that malware is evolving so fast that sharing this type of IoC is irrational. Where antivirus engines have abandoned hash-based detection mechanisms a long time ago, the CTI field seems to be reiterating them.

CTI feeds can only exist due to a lack of competition, which is a problem of economics. As the demand is bigger than the supply, it is difficult to assess the added value. Currently this prompts many defenders to just ingest as many feeds as possible, creating a signal-to-noise problem that has been covered extensively in intelligence literature. The introduction of a formal CTI methodology will also improve CTI quality. Where completeness and timeliness are objective indicators of IoC value, currently the most-used indicator of quality is true or false positive.

CTI Vendors Are Untransparent in Their Supply

The intelligence value of a raw datapoint depends on its source. To be able to stand on its own it has to meet criteria of completeness, timeliness, and verifiability. As of now, most organizations are “consumers” rather than “customers” of threat data they use to feed their analysis process. Not only is the methodology of their providers unknown, they also remain ignorant about its provenance. The quality of underlying sources and assumptions remains unclear.

From previous research it is known that commercial CTI providers often outsource their CTI data to competitors because of lacking research and development resources.³² One publicly known example of this is the Cyber Threat Alliance, through which 25 member organizations share four million observables on a monthly basis.³³ The coalition-forming of commercial providers can lead to overlap in the reporting on certain threats, which is largely absent in the freely available open source CTI feeds.³⁴ To many practitioners, this overlap is unknown and difficult to recognize in practice due to the high price point of commercial feeds.

In its current form, the CTI supply chain is so deep that it is opaque how individual pieces of intelligence are established. This lacking of a ground truth not only makes it very difficult to establish quality, it is also impossible to judge the relevancy of the data delivered by a provider (e.g., toward a specific industry). In the scarce cases where any source attribution is done, it is not more than a label (e.g., *honeypot*, *The Tor exit*). It can be argued that procuring a CTI feed is effectively partial outsourcing of intelligence analysis, which implies a trade-off in applicability and relevancy of analysis. Therefore,

procured CTI must always be analyzed internally, but many CTI feeds are consumed directly rather than being used as a third-party preprocessed input to an in-house analysis cell.

CTI Is Too Biased

Naturally, bias can never fully be ruled out, and the CTI field is aware of potentially biased analysis. In this section we argue, however, that the field is not aware of the heavy bias of many of the CTI providers. As many organizations consume raw data without any additional in-house analysis, additional measures are required to address bias.

We have already shown that most CTI feeds are biased due to limitations of their sensor base, which is a sampling bias. However from a commercial perspective, for most commercial CTI providers it is very attractive to focus on state actors and sophisticated threat groups. Many commercial providers tend to frequently put out new reporting on the big nation-state actors, because it makes for good marketing. The CTI industry does not seem uncomfortable with alluding to state-sponsored activity constantly, instead of other activity that might seem cyber petty criminal but is more relevant to the average consumer. Also, the activities of lesser-known criminal groups or nation-state actors like Turkey and India receive less coverage. The magnification of actors from the opposing “Bloc” is not strange to the CTI industry. Russia-based CTI vendor Kaspersky, known for uncovering the activities of the Equation Group, the former Tailored Access Operations unit of the National Security Agency,³⁵ has scarce reporting on Russian cyber operations. On the other hand, U.S. CTI firm CrowdStrike is prolific on the reporting of Russian activity but is tongue-tied on activity from its own territory.

In the previous section we have argued providers must provide insight into their supply chain. In addition to this, they should also provide insight into their methodology and analysis process. Only then are customers able to account for the biases introduced by the source and identify relevancy to their own threat environment.

CTI Attribution Is Difficult

The attribution of malicious acts to their effective actors in CTI is problematic. The OlympicDestroyer malware that targeted the Winter Olympic Games in South Korea was only attributed to Russia after being linked to North Korea initially. Some analysts stepped into the deceptive trap deliberately placed by the malware authors, trying to frame North Korea.³⁶ It is perhaps for this reason that, in 2018, the United States Office of the

Director of National Intelligence released its guidelines on the attribution of attacks to specific countries or actor groups.³⁷

The challenge with attribution is further complicated by the myriad of actor group names created by vendors for marketing purposes. Most CTI vendors attribute new campaigns to specific actor groups. This introduces challenges of its own, of which the differences in the naming of Russian military intelligence actors is exemplary. There are known as Fancy Bear (CrowdStrike), APT-28 (Mandiant/FireEye), Sofacy (Kaspersky), STRONTIUM (Microsoft), Sednit (ESET), Tsar Team (iSight), Swallowtail (Symantec), Pawn Storm (Trend Micro), TG-4127 (SecureWorks), and Grizzly Steppe (U.S. government). The proliferation of group names complicates the analysis of actor TTPs.

The above is largely a consequence of attribution being used as a marketing instrument. CTI vendors each use their own naming scheme to identify actors in order to be easily recognizable. While this is sensible from a marketing point of view, it causes confusion with practitioners, as it distorts the threat landscape. Currently actor group names are only useful to identify the CTI vendor that coined their names. A common naming convention does not exist, which leads to two problems. First, the number of threat groups that actually exists is largely overestimated. Second, the absence of a naming convention complicates intelligence sharing. Each different name for the same group is an additional data point that complicates inference, which in turn contributes to the signal-to-noise ratio.

Initial work on this has been performed by Mitre, which has categorized actor entities and associated names in the ATT&CK framework.³⁸ However, the CTI field must reach a common agreement on the naming of threat actor groups. Only through uniformity can practitioners interpret actual new TTPs and campaigns over just rebranded ones.

CYBER THREAT INTELLIGENCE: AN IDEAL LIAISON

Despite its challenges, CTI should not be discarded too quickly. CTI is an emerging field that has potential for the IC of government intelligence agencies and related organizations. With its aim to apply computer science in the defense against the malicious use of communication technology, the field carries important opportunities, which we highlight in this section.

CTI Expands the Collection Spectrum

Commercial providers of CTI usually have sensors placed in customer networks. This provides them with a different resource for CTI collection over government intelligence agencies. While government intelligence agencies may also have sensors in networks of organizations of strategic

importance, data from the sensor base of commercial CTI providers can potentially fill collection gaps and expand the view on the threat landscape. This provides an important intelligence angle in a threat environment where state-sponsored groups have intrusion operations into businesses of strategic importance in other states. Some argue that by adding CTI providers in the collection process, the signal-to-noise ratio is even further increased. However, as pointed out by Wirtz, the biggest risk is not being flooded by raw data, but it not offering valuable information about the most important targets.³⁹

One example of the value added by commercial CTI companies is the technical analysis from FireEye, which was referenced multiple times in a 2018 report ordered by the president of the United States on Chinese industrial espionage. FireEye's intelligence indicated 262 intrusions over a half-year period by 72 unique Chinese actor groups.⁴⁰ The report also recognized that in April and May 2016, four semiconductor and chemical companies from the United States, Europe, and Asia were compromised by the same operatives. It is these cases where CTI firms can add value, as they might provide additional visibility into victim networks through sensor placement, which would have otherwise remained unknown.

CTI Can Bring Artificial Intelligence (AI) to Intelligence

The abundant and growing amount of data on cyber threats needs to be aggregated and data-mined into information ready for human analysis. Intelligence analyst resources are scarce and expensive. Fortunately, AI technology is now becoming mature. AI holds much potential to solve the Big Data problem that many organizations are trying to solve. AI can help to preprocess raw input data to offload human analysts. Various examples of successful application of AI technology that are relevant to intelligence analysis already exist, such as Support Vector Machines, which have been proven successful many times in detecting covert channels on a computer network.⁴¹

With its potential for decision support through the preprocessing of large quantities of input data, we believe that, in the near future, AI will provide useful intelligence use-cases. Especially for intelligence analysis, "artificial" intelligence will never be able to replace the decisionmaking capabilities of a team of human analysts.

CTI Increases the Intelligence Community of Practice

The CTI community did not grow on its own. It has a strong link to the government intelligence agencies in jurisdictions such as the United States and United Kingdom. Many experts formerly employed with these agencies

now fill intelligence roles at commercial CTI providers. The rise of commercial CTI increases the size of the community of intelligence practice as the number of workers with analytical skills and technical skills is increased. Furthermore, commercial CTI providers can help prioritize analysis activities at government agencies, because they are supplied with more quality resources of raw input on various threat actor groups. This can offload resources at government intelligence agencies that historically struggle to attract quality resources. The rise of commercial CTI, however, also increases the funnel of potential hires for government intelligence agencies.

CONCLUSION

In this article we stated the position that the CTI field is often inadequate in its analysis. This is primarily because its methodology is flawed. As a result, CTI is currently delivering a broken product due to a broken process.

This is, however, quite reasonable. While CTI has already contributed to the exposure of many intrusion campaigns by nation-state hacking groups, it is a field in its infancy that needs to advance to a higher level of maturity. We have argued that for this, CTI must deduce from work on methodology encoded in its parent field of intelligence studies.

We have further demonstrated that the field has several challenges that it needs to solve. When the CTI field succeeds at improving its methodology, the impact of the challenges with the quality, supply, bias, and actor naming will be reduced as well. Furthermore, we have demonstrated that the CTI field, through its strong roots in computer science, has important opportunities for the broader IC.

As optimistic practitioners and scholars of CTI, we believe the initiation of this debate is necessary to advance the CTI field to its next era, analogous to the several reformations through which the IC has lived. A field's principles and theories must be falsifiable and able to confront aims of falsification in order to progress its theories. This is basic scientific scrutiny and if CTI is antifragile it will thrive and mature when exposed to such pressures.

With their shared duty of maintaining a strategic advantage in an age of complex threats, further alliance of the IC with the CTI field will drive cyber defense for the years to come.

REFERENCES

¹ Ponemon Institute, "Cost of a Data Breach Report 2019," Michigan, 2019, <https://www.ibm.com/security/data-breach>

² Council on Foreign Relations, "Cyber Operations Tracker," https://www.cfr.org/interactive/cyber-operations/export-incidents?_format=csv

- ³ Mandiant, “APT1 Exposing One of China’s Cyber Espionage Units,” Alexandra, 2013, <https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf>
- ⁴ Gartner, “Gartner Market Guide for Security Threat Intelligence Products,” Stamford, 2019, <https://go.eclecticiq.com/gartner-market-guide-to-cti-2019>.
- ⁵ Shelly Singh, “Threat Intelligence Market Worth \$12.9 billion by 2023,” Markets and Markets, 1 November 2018, <https://www.marketsandmarkets.com/PressReleases/threat-intelligence-security.asp>
- ⁶ Eric M Hutchins, Michael J. Cloppert, and Rohan M. Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains,” *Proceedings of the International Conference on Information Warfare & Security*, 2011, <https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>
- ⁷ Sergio Caltagirone, Andrew Pendergast, and Christopher Betz, “The Diamond Model of Intrusion Analysis,” Center for Cyber Intelligence Analysis and Threat Research, 2013, <https://apps.dtic.mil/docs/citations/ADA586960>
- ⁸ David Bianco, “The Pyramid of Pain,” *Enterprise Detection & Response*, 17 January 2014, <https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>
- ⁹ Blake E. Strom, Andy Applebaum, Doug P. Miller, Kathryn C. Nickels, Adam G. Pennington, and Cody B. Thomas, “MITRE ATT&CK: Design and Philosophy,” McLean, 2018, <https://www.mitre.org/publications/technical-papers/mitre-attack-design-and-philosophy>
- ¹⁰ Stephen Marrin, *Improving Intelligence Analysis: Bridging the Gap between Scholarship and Practice* (Abingdon-on-Thames: Routledge, 2012), p. 21.
- ¹¹ Sherman Kent, “Words of Estimative Probability,” Central Intelligence Agency, 19 March 2007, <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/sherman-kent-and-the-board-of-national-estimates-collected-essays/6words.html>
- ¹² Richards J. Heuer, Jr., *Psychology of Intelligence Analysis* (Washington, DC: Center for the Study of Intelligence, 1999).
- ¹³ See, for example, Erick Mandt and Robert M. Lee, “Leveraging Threat Intelligence in an Active Defense,” SANS DFIR Summit, 2016, <https://www.sans.org/cyber-security-summit/archives/file/summit-archive-1492180823.pdf>; Andreas Sfakianakis, “Stop Tilting at Windmills: 3 Key Lessons that CTI Teams Should Learn from the Past,” SANS CTI Summit, 2020, <https://www.sans.org/cyber-security-summit/archives/file/summit-archive-1579635728.pdf>
- ¹⁴ Richards J. Heuer, Jr. and Randolph Pherson, *Structured Analytic Techniques for Intelligence Analysis* (Washington, DC: CQ Press, 2011).
- ¹⁵ David R. Mandel, “The Occasional Maverick of Analytic Tradecraft,” *Intelligence and National Security*, Vol. 35, No. 3 (2020), pp. 438–443.
- ¹⁶ Robert M. Gates, “Guarding against Politicization,” *Studies in Intelligence*, Vol. 36, No. 5, 1992, pp. 5–13.
- ¹⁷ Douglas J. MacEachin, *The Tradecraft of Analysis: Challenge and Change in the CIA* (Washington, DC: Consortium for the Study of Intelligence 1994).

- ¹⁸ Office of the Director of National Intelligence, “Intelligence Community Directive (ICD) 200: Management, Integration, and Oversight of Intelligence Community Analysis,” January 2007, http://www.dni.gov/files/documents/ICD/ICD_200.pdf, pp. 1–8; Office of the Director of National Intelligence, “Intelligence Community Directive (ICD) 203: Analytic Standards,” June 2007, http://www.dni.gov/files/documents/ICD/ICD_203.pdf, pp. 1–6.
- ¹⁹ Claudio Guarnieri, “Digital Attack on German Parliament: Investigative Report on the Hack of the Left Party Infrastructure in Bundestag,” *Netzpolitik*, 19 June 2015, <https://netzpolitik.org/2015/digital-attack-on-german-parliament-investigative-report-on-the-hack-of-the-left-party-infrastructure-in-bundestag>
- ²⁰ Inlandsnachrichten, “Russland Soll Cyber-Attacke Auf Bundestag Verübt Haben,” Reuters, 13 May 2016, <https://de.reuters.com/article/deutschland-russland-cyberangriff-idDEKCN0Y41D2> (accessed 6 February 2020).
- ²¹ Root9B, “Technical Follow Up—APT28 Malware Analysis,” Colorado Springs, 2015, https://www.root9b.com/sites/default/files/whitepapers/root9b_follow_up_report_apt28.pdf
- ²² Office of the Director of National Intelligence, “Intelligence Community Assessment (ICA) 2017-01D, Background to Assessing Russian Activities and Intentions in Recent US Elections: The Analytic Process and Cyber Incident Attribution,” 6 January 2017, https://www.dni.gov/files/documents/ICA_2017_01.pdf, pp. 1–25.
- ²³ World Anti-Doping Agency, “Cyber Hack Update: Data Leak Concerning 20 Athletes from 14 Countries and 13 Sports,” 3 October 2016, <https://www.wada-ama.org/en/media/news/2016-10/cyber-hack-update-data-leak-concerning-20-athletes-from-14-countries-and-13>
- ²⁴ Zeit Online, “CDU Prüft Möglichen Hackerangriff,” 13 May 2016, <https://www.zeit.de/digital/2016-05/cdu-hacker-angriff-trend-micro-russland-verdacht>
- ²⁵ Milo Jones and Philippe Silberzahn, *Constructing Cassandra: Reframing Intelligence Failure at the CIA, 1947–2001* (Stanford, CA: Stanford University Press, 2013).
- ²⁶ Ponemon Institute, “The Value of Threat Intelligence,” Michigan, 2019, <https://www.anomali.com/resources/whitepapers/2019-ponemon-report-the-value-of-threat-intelligence-from-anomali>
- ²⁷ Neda Afzaliseresht, Yuan Miao, Sandra Michalska, Qing Liu, and Hua Wang, “From Logs to Stories: Human-Centred Data Mining for Cyber Threat Intelligence,” *IEEE Access No. 8*, 2020, <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8960350&isnumber=8948470>
- ²⁸ Thomas Schaberreiter, Veronika Kupfersberger, Konstantinos Rantos, Arnolnt Spyros, Alexandros Papanikolaou, Christos Ilioudis, and Gerald Quirchmayr, “A Quantitative Evaluation of Trust in the Quality of Cyber Threat Intelligence Sources,” *Proceedings of the 14th International Conference on Availability, Reliability and Security, 2019*.
- ²⁹ Harm Griffioen, Tim M. Booi, and Christian Doerr, “Quality Evaluation of Cyber Threat Intelligence Feeds,” *Proceedings of the 19th International Conference on Applied Cryptography and Network Security, 2020*.

- ³⁰ Bianco, “The Pyramid of Pain.”
- ³¹ Verizon, “Data Breach Investigations Report,” New York, 2016, https://conferences.law.stanford.edu/cyberday/wp-content/uploads/sites/10/2016/10/2b_Verizon_Data-Breach-Investigations-Report_2016_Report_en_xg.pdf
- ³² Aviram Zrahia, “Threat Intelligence Sharing between Cybersecurity Vendors: Network, Dyadic, and Agent Views,” *Journal of Cybersecurity*, Vol. 4, No. 1 (2018).
- ³³ Cyber Threat Alliance, “About CTI—Our Sharing Statistics,” <https://www.cyberthreatalliance.org/about-cta/>
- ³⁴ Griffioen, Booij, and Doerr, “Quality Evaluation of Cyber Threat Intelligence Feeds,” p. 1.
- ³⁵ Dave Neal, “Kaspersky: The Equation Group Looks and Smells Even More like the NSA,” *The Inquirer*, 13 March 2015, <https://www.theinquirer.net/inquirer/news/2395638/kaspersky-fingers-nsa-style-equation-group-for-hard-drive-backdoor-epidemic>
- ³⁶ Paul Rascagnères and Warren Mercer, “Who Wasn’t Responsible for Olympic Destroyer,” *Virus Bulletin*, 2018, <https://www.virusbulletin.com/virusbulletin/2018/10/vb2018-paper-who-wasnt-responsible-olympic-destroyer/>
- ³⁷ Office of the Director of National Intelligence, “A Guide to Cyber Attribution,” 14 September 2018, https://www.dni.gov/files/CTIIC/documents/ODNI_A_Guide_to_Cyber_Attribution.pdf, pp. 1–5.
- ³⁸ The Mitre Corporation, “APT28,” Mitre ATT&CK Knowledge Base, 11 October 2019, <https://attack.mitre.org/groups/G0007>
- ³⁹ James J. Wirtz, “The American Approach to Intelligence Studies,” in *Handbook of Intelligence Studies*, edited by Loch K. Johnson (Abingdon-on-Thames: Routledge, 2006).
- ⁴⁰ Office of the United States Trade Representative—Executive Office of the President, “Findings of the Investigation Into China’s Acts, Policies and Practices Related to Technology Transfer, Intellectual Property, and Innovation Under Section 301 of the Trade Act of 1974,” 22 March 2018, <https://ustr.gov/sites/default/files/Section%20301%20FINAL.PDF>
- ⁴¹ Taeshik Sohn, JungTaek Seo, and Jongsub Moon, “A Study on the Covert Channel Detection of TCP/IP Header Using Support Vector Machine,” *Proceedings of the International Conference on Information and Communications Security*, 10–13 October 2003, pp. 313–324.