VB2021
localhost

7 - 8 October, 2021 / vblocalhost.com

WHAT CYBER THREAT INTELLIGENCE ANALYSTS CAN LEARN FROM SHERLOCK HOLMES

Selena Larson
Proofpoint, USA

slarson@proofpoint.com

ABSTRACT

In 1887, physician and crime fiction writer Sir Arthur Conan Doyle introduced readers to the brilliant, arrogant, and cocaine-addicted consulting detective Sherlock Holmes. One of the most beloved characters in literary history, Holmes' unbelievable adventures reported by his trusty sidekick Doctor John Watson introduced Victorian popular culture to the capabilities of forensic science and analytical techniques that would become the foundations of modern detecting. And these tools can be applied to cyber threat intelligence, too.

'In solving a problem of this sort, the grand thing is to be able to reason backward', Holmes tells Watson in *A Study in Scarlet*. This puzzle-solving technique, though presented as a work of fiction, is a reliable method for cyber threat intelligence analysts and forensic cyber investigators. Modern crimes perpetrated by cybercriminals and state-backed actors have things in common with Victorian-era murderers: they leave evidence behind. In cyber threat intelligence, these are known as 'threat behaviours', or the tactics, techniques and procedures executed by adversaries. Each of these behaviours is a clue to identifying cyber attackers' motives and methods.

In his debut story, Conan Doyle sums up what it means to think like a detective – or, in our case, a cyber threat analyst: 'There are few people, however, who, if you told them a result, would be able to evolve from their own inner consciousness what the steps were which led up to that result,' Holmes says. 'This power is what I mean when I talk of reasoning backward, or analytically.'

This piece will dig into the investigation and forensic techniques Sherlock Holmes first introduced to mainstream readers, as well as modern interpretations of the detective's analytical methods. Additionally, analysts will learn how to apply those concepts to modern cyber investigations and understand how critical thinking, analytical puzzle solving, and historic forensic sciences can apply to their current careers.

SO IT'S A MURDER

Before an investigation begins in earnest, cyber threat intelligence analysts must look inward. That is, think about how you think. Every person in the world has inherent biases based on their life, work, and general human experiences. In *The Boscombe Valley Mystery*, Holmes and Watson are on a train to investigate a crime, and the following exchange occurs:

> Watson: So, it's a murder then.
>
> Holmes: At least, that's how it has been explained to me, but I will not conjecture until I look into it myself.

Before he begins the investigation, Holmes ensures his judgement remains unclouded by assuming something occurred before he knows the truth. As intelligence analysts, it is important to be self-aware of prejudices that exist and the conclusions we want to be true. For example, assuming an organization is targeted by Chinese threat actors because the adversaries used PlugX malware [1] is an example of biases in action.

Cognitive biases – or our preconceived notions – impact our objectivity and make it more difficult to think critically and consider alternative hypotheses. This can result in poor analysis, misrepresentation of data, such as attribution to the wrong threat actor, and conclusions that are unable to be verified by other researchers or customers. Katie Nickels, the Director of Intelligence at *Red Canary*, effectively describes a common thought process called confirmation bias in her 2019 talk 'The Cycle of Cyber Threat Intelligence' [2]. Analysts frequently fall into cognitive bias behaviours by selectively supporting one hypothesis, looking for information and evidence that supports their preconceived notions while rejecting information that refutes them, and giving greater weight to data that supports their hypothesis than to data that contradicts it.

CIRCUMSTANTIAL EVIDENCE

One effective method of combating biases in intelligence analysis is to conduct analyses of competing hypotheses (ACH), a useful structured analytic technique that can help analysts look at multiple possible explanations for observed activity.

The Boscombe Valley Mystery again provides us with a prescient Holmes quote about jumping to conclusions without objectively reviewing all evidence:

> 'Circumstantial evidence is a very tricky thing. It may seem to point very straight to one thing, but if you shift your own point of view a little, you may find it pointing in an equally uncompromising manner to something entirely different.'

In cyber threat intelligence, much of the evidence we collect and observe is circumstantial. Unless we can infiltrate a threat actor's machine, hijack their front-facing camera, and take a photo of the operator with 'hands on keyboard', assessments of activity generally rely on educated inferences made based on all available data. In educating Watson about the importance of perspective when it comes to assessing a suspect's guilt, Holmes presents readers with a definition of ACH wrapped in prose – shifting your point of view often provides different answers.

ACH helps analysts navigate the cognitive roadblocks inherently present due to bias. In the U.S. Central Intelligence Agency report *A Tradecraft Primer: Structured Analytic Techniques for Improving Intelligence Analysis* [3], ACH is described as using all the available evidence to disconfirm, rather than confirm, hypotheses. Analysts should explicitly list all possible hypotheses and evidence, cross-referencing each piece of evidence with the conclusion it supports or disproves.

For example, in a 2021 SANS presentation [4] detailing espionage activity targeting medical professionals from the Iranian threat group TA453 [5], *Proofpoint* researchers and analysts Josh Miller and Crista Giering illustrated the ACH methodology used while writing their report to ensure thorough analysis. The analysts identified multiple hypotheses to explain the activity and threat actors' possible motivation and objectives, and how evidence might impact each conclusion.

Their process can be used as a blueprint to conduct structured analysis:

Hypothesis 1a

- This campaign demonstrates an intelligence requirement to collect specific medical information related to genetic, oncology, or neurology research.

Hypothesis 1b

- This campaign may demonstrate an interest in the patient information of the targeted medical personnel.

Hypothesis 1c

- This campaign may demonstrate an aim to use the recipients' accounts in further phishing campaigns.

Hypothesis 2a

- This campaign may represent a shift in TA453 targeting overall.

Hypothesis 2b

- This campaign may be an outlier, reflective of a specific priority intelligence tasking given to TA453.

Figure 1: List of available hypotheses used in investigating TA453 activity targeting medical professionals.

	H1a	H1b	H1c	H2a	H2b
E1. Historically has targeted think tanks, dissidents, academics, etc.	-	-	+	+	+
E2. Spoofed Gmail account of a prominent Israeli physicist	-	-	NA	-	-
E3. Subject & lures of emails all on Israel nuclear weapons	-	-	NA	-	-
E4. Prior campaigns by TA453 + other IR APT have used compromised accounts for further phishing	NA	NA	+	NA	NA
E5. Targeted less than 25 individuals at variety of medical orgs in US & Israel	+	+	NA	+	+
E6. All of the 25 researching genetics, oncology, or neurology	+	+	NA	+	+

Figure 2: List of available evidence supporting or not supporting available hypotheses.

Ultimately, there is not enough data in the observed campaign to narrow down ACH into one high-confidence assessment, but it does provide investigative questions to guide future research.

Understanding the limitations of data is fundamentally important as it helps analysts scope confidence levels, generate new questions and routes of investigation, and effectively and comprehensively report findings to customers.

Much like contemporary intelligence analysts, Sir Arthur Conan Doyle put ACH into practice beyond his fictional narratives – and saved people's lives. A real-life doctor, detective, and perpetual justice-seeker, Conan Doyle helped free two wrongfully convicted men [6, 7] – who had been sentenced for crimes based on circumstantial evidence – by assessing available evidence and providing alternative conclusions, proving their innocence.

WRITING ON THE WALL

But what exactly can evidence tell us? Depending on the type of crime, evidence can provide details on the who, what, when, why, and how – the key questions we must answer to provide effective analysis and recommendations to stakeholders and prevent similar crimes from occurring.

Sherlock Holmes used methods of deduction relying on indicators that have parallels to what cyber threat analysts call threat behaviours, or tactics, techniques, and procedures (TTPs). The private detective studied human anatomy and the

environmental and physical indicators people left behind to paint a picture of a criminal. In *A Study in Scarlet,* Holmes describes a pattern that all people adhere to, which can provide insight into a suspects' height:

'When a man writes on a wall, his instinct leads him to write above the level of his own eyes.'

Holmes used footprints and the length of a stride to deduce the height of, or possible injury to, a suspect; cigarette ash to determine how long a person stayed at a particular location; characteristics of fingerprints to identify their likely owners (even before Scotland Yard officially adopted the scientific technique); and the wear patterns in a shoe's tread to identify a man's gait and possible occupation. In all his adventures, Holmes provided a critical look at the impact individuals' behaviour has on their surrounding environment, shocking both fictional characters and the public with the amount of information one can glean if one pays close attention.

Deduction based on human conduct is so effective because habits are difficult to change. The way humans write, speak, hold their cigarette, or wear down their shoe is fundamentally habitual, and without consciously thinking about modifications, can be used to identify them as individuals. For cyber threat actors, the methods they use to infiltrate a target, move laterally within an environment, and exfiltrate data or wipe hosts are not easily changed, thus can be used to paint a full picture of an adversary.

What's more, as David Bianco explains with his Pyramid of Pain [8], identifying and blocking adversary behaviours instead of static indicators is more effective because it forces threat actors to change their behaviour.

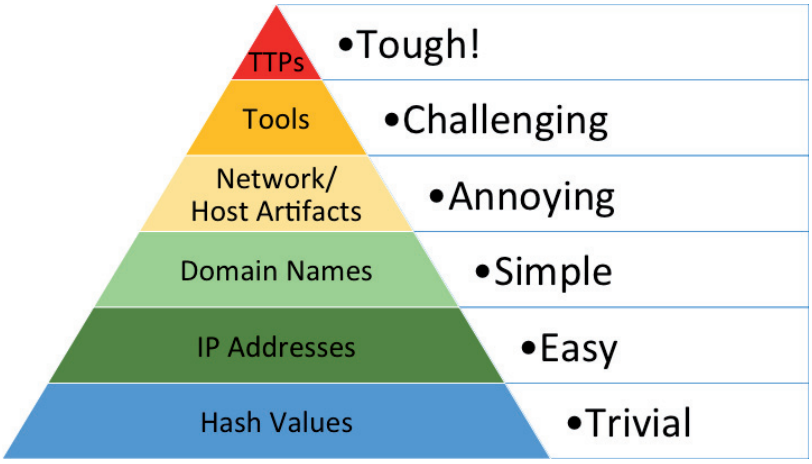

Figure 3: Pyramid of Pain.

By identifying an adversary's habits, analysts can more accurately track, identify, and block the methods and operational requirements used by adversaries that impact their organizations or customers.

MATHEMATICAL CERTAINTIES

In *The Sign of Four*, Holmes notes that although it is difficult to predict what one person will do, it is much easier to assume the behaviours and intent of groups of people.

'...while the individual man is an insoluble puzzle, in the aggregate he becomes a mathematical certainty.'

Holmes, paraphrasing philosopher William Winwood Reade's book *The Martyrdom of Man*, believes humans are inherently predictable. And while Holmes finds human nature's general sameness extremely boring, it does become useful when grouping and tracking adversaries in the cyber realm. When extrapolating the above assumption as it relates to human behaviour, and specifically activities in cyberspace, one can deduct it is likely that groups of people who exhibit the same behaviours over and over will continue to do so in the future.

For example, *Proofpoint* tracks TA406, a threat actor associated with the Democratic People's Republic of Korea, based on multiple threat behaviours observed in almost all its campaigns. This actor frequently targets political, foreign policy, and non-profit organizations, especially those working with or experts on activities impacting the Korean Peninsula. TA406 uses its own registered and controlled infrastructure to host credential capture web pages and malicious documents that it distributes via phishing. The actor uses different registrars and hosting providers in various geographies including Eastern Europe and Southeast Asia, and uses *Gmail*, *Yandex* and *Mail[.]ru* email accounts masquerading as legitimate government or non-profit entities to distribute its lures. TA406 may leverage URLs linking to the *SendGrid* email delivery service that redirect to an attacker-controlled domain that hosts the malicious payload.

Independently, these behavioural indicators are not indicative of a specific adversary, but in aggregate they can be used to identify a likely threat actor – in this case, TA406.

TYPEWRITTEN INDICATORS

Language, writing style, and filenames threat actors use are all brush strokes on the canvas of attribution [9]. Like their Victorian-era counterparts, cybercriminals' words and how they are communicated can provide investigators with useful details.

Sherlock Holmes solved *A Case of Identity* – the catfishing story of its time – by looking at the small details in a typewritten note in the same way investigators today look at the arches, loops, and whorls of fingerprints.

> 'It is a curious thing,' remarked Holmes, 'that a typewriter has really quite as much individuality as a man's handwriting.'

He traced the identity of the letter writer in the same way threat intelligence analysts look at file-naming conventions, strings in code, or language used in forums.

For instance, two 2020 reports from *NCC Group* [10] and *Clearsky Cyber Security* [11] detailed separate threat activities, but both included web shell reuse with the string `citrix@kharpedar`. The Farsi words in the string, which translate to 'donkey father' [12], helped analysts attribute the activities to an Iranian state threat actor known as Pioneer Kitten.

Modern cybercriminal investigations are known to use language to trace anonymous postings on web forums back to their real-world authors. According to an April 2021 report from *48 Hours*, a teenager confirmed [13] the identity of an anonymous account that tried to have her killed based on the grammar and punctuation he used on his dark web forum posts.

WOULD YOU LIKE TO PLAY A GAME?

Sherlock Holmes is perhaps best known for the phrase elicited in multiple stories and throughout contemporary retellings of the detective's adventures:

> 'The game is afoot!'

It can be disconcerting to think about hunting down criminals this way, but the mathematical models and strategy that apply to board and video games like Chess or *League of Legends* are undeniably effective in intelligence analysis and detecting.

Modern intelligence operators, including those working for the Central Intelligence Agency, use games to hone their investigative and decision-making skills. In 2017, I attended [14] an event featuring David Clopper, senior collection analyst with 16 years' experience at the CIA, and the CIA's game maker. Clopper described how the agency trains its officers via board games that are designed to mimic real-world situations. In *Collection* [15], officers must work together to solve major international crises via intelligence gathering, with players representing military, economic and political analysts.

Cyber defenders also use games to assess outcomes and resilience in the face of cyber attacks. The Departments of Energy and Defense, for example, hosts cyber attack exercises [16] to determine possible impact to the electric grid system if a threat actor can compromise a utility's operations network. The event is part of a program run by the Defense Advanced Research Projects Agency called Rapid Attack Detection, Isolation and Characterization Systems (RADICS).

Holmes indirectly advocates for the power of games when he says in *Scarlet*, 'This power is what I mean when I talk of reasoning backward, or analytically.' This idea is called backward chaining [17], or using strings of 'if, then' clauses to make an inference based on the previous known facts. When an 'if, then' statement is found to be true, it is added to the string of clauses to help further the analysis and is used to remove untrue statements. Some researchers have applied this theory to gaming via backward induction [18]. Basically, analysts can learn to look at the solution or outcome of a series of actions and reason backward to determine what each moment in the series looks like, thereby identifying a suspects' behaviours.

IT'S ELEMENTARY

Sherlock Holmes provides us with numerous lessons in deduction and analysis, but perhaps the most important lesson is one of humanity. Holmes, for all his genius, is not a kind person. He uses Watson as a sounding board rather than a collaborator in crime solving. He looks down upon Scotland Yard's Detective Inspector G. Lestrade, and makes fun of criminals, suspects, bereaved clients, and law enforcement directly to their faces.

Conan Doyle's Sherlock was less of an abrasive sociopath than Benedict Cumberbatch's interpretation in the BBC's *Sherlock* or Robert Downey Jr.'s in Guy Ritchie's gritty retelling of Holmes' adventures. However, he was always alone in his genius, and codependent in his emotional wellbeing, neither of which are healthy.

Whether your expertise is in strategic analysis, malware reversing, threat hunting or detection development, threat intelligence is a team sport. Holmes's holier-than-thou attitude was harmful, because despite the public treating him like a rock star and his ability to solve the unsolvable, he stepped on people's feelings (including Watson's), and created a toxic space where people were afraid to share their own ideas or solutions to problems. He failed to make the people around him better and therefore limited the team's overall success.

That type of environment is not conducive to effective research and intelligence activities. It prevents innovation and collaboration and promotes exclusivity within an organization and the information security community at large. We can embrace Holmesian critical thinking and investigative techniques while rejecting the absurd idea that egotistical genius is more valuable than thoughtful collaboration.

FINDING THE TRUTH

At his core, Sherlock Holmes teaches us to rely on data, logic, and sound reasoning, even if the outcome is unexpected or undesirable. And this lesson serves as a reminder of our mission as cyber threat intelligence analysts: Find the truth.

'Once you eliminate the impossible, whatever remains, no matter how improbable, must be the truth.'

REFERENCES

[1] MITRE ATT&CK. PlugX. https://attack.mitre.org/software/S0013/.

[2] Nickels, K. The Cycle of Cyber Threat Intelligence. https://www.youtube.com/watch?v=J7e74QLVxCk.

[3] US Government. A Tradecraft Primer: Structured Analytic Techinques for Improving Intelligence Analysis. https://www.cia.gov/static/955180a45afe3f5013772c313b16face/Tradecraft-Primer-apr09.pdf.

[4] Nickels, K.; Miller, J.; Giering, C. STAR Webcast: Dissecting BadBlood: an Iranian APT Campaign. https://www.sans.org/webcasts/star-webcast-dissecting-badblood-iranian-apt-campaign-119545/.

[5] Miller, J. BadBlood: TA453 Targets US and Israeli Medical Research Personnel in Credential Phishing Campaigns. Proofpoint. https://www.proofpoint.com/us/blog/threat-insight/badblood-ta453-targets-us-and-israeli-medical-research-personnel-credential.

[6] National Records of Scotland. The Case of Oscar Slater. https://www.nrscotland.gov.uk/research/learning/features/the-case-of-oscar-slater.

[7] Heydt, B. Sir Arthur Conan Doyle and the case of George Edalji. British Heritage Travel. https://britishheritage.com/history/sir-arthur-conan-doyle-george-edalji.

[8] Bianco, D. The Pyramid of Pain. Enterprise Detection & Response. https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html.

[9] Rod, T.; Buchanan, B. Attributing Cyber Attacks. The Journal of Strategic Studies, 2015 Vol. 38, Nos. 1–2, 4–37. https://ridt.co/d/rid-buchanan-attributing-cyber-attacks.pdf.

[10] NCC Group. RIFT: F5 Networks K52145254: TMUI RCE vulnerability CVE-2020-5902 Intelligence. https://research.nccgroup.com/2020/07/05/rift-f5-networks-k52145254-tmui-rce-vulnerability-cve-2020-5902-intelligence/.

[11] ClearSky Cyber Security. Pay2Kitten Pay2Key Ransomware – A New Campaign by Fox Kitten. https://www.clearskysec.com/wp-content/uploads/2020/12/Pay2Kitten.pdf.

[12] https://twitter.com/irsdl/status/1280439017755086850?lang=en.

[13] CBS News. Link between dark web murder-for-hire plot and gamer may be bad grammar. https://www.cbsnews.com/news/dark-web-murder-for-hire-gamer-grammar-48-hours/.

[14] Larson, S. Why the CIA uses board games to train its officers. CNN Business. https://money.cnn.com/2017/03/13/technology/cia-board-games-training/index.html.

[15] Machkovech, S. The CIA uses board games to train officers – and I got to play them. Ars Technica. https://arstechnica.com/gaming/2017/03/the-cia-uses-board-games-to-train-officers-and-i-got-to-play-them/.

[16] Lyngaas, S. How the US military used a creepy island to test cyberattacks on the grid – in the middle of a pandemic. Cyberscoop. https://www.cyberscoop.com/plum-island-darpa-cyber-exercise-grid/.

[17] Great Learning. Chaining Techniques in Artificial Intelligence. https://www.mygreatlearning.com/blog/chaining-techniques-artificial-intelligence/.

[18] Economic Applications of Game Theory: Backward Induction. https://ocw.mit.edu/courses/economics/14-12-economic-applications-of-game-theory-fall-2012/lecture-notes/MIT14_12F12_chapter9.pdf.