TALLINN UNIVERSITY OF TECHNOLOGY
School of Information Technologies

Uku Joosep Palu 242792IVCM

An analysis of the actionability of predictive cyber threat intelligence

Master's thesis

Supervisor: Margus Ernits
MSc

Tallinn 2026

TALLINNA TEHNIKAÜLIKOOL
Infotehnoloogia teaduskond

Uku Joosep Palu 242792IVCM

Ennustava küberohuluure rakendatavuse analüüs

Magistritöö

Juhendaja: Margus Ernits
MSc

Tallinn 2026

Author’s declaration of originality

I hereby certify that I am the sole author of this thesis. All the used materials, references to the literature and the work of others have been referred to. This thesis has not been presented for examination anywhere else.

Author: Uku Joosep Palu
20.05.2026

Abstract

Cyber threat intelligence (CTI) has increasingly been relied upon as a tool for enabling proactive cybersecurity and defense. In recent years, advances in AI and machine learning have increased the prevalence of predictive CTI, CTI that aims to forecast the likelihood, timing, or characteristics of future cyber threats. Current research into predictive CTI is dominated by discussion of the tools and techniques that enable it and improve its accuracy, and, despite the vast amount of threat intelligence produced globally, there is little study into the quality of past predictions and their influence on the defensive behaviour of organisations. This paper bridges this gap by investigating academic and industry literature as well as first-hand accounts from cybersecurity practitioners to determine what characteristics make predictive CTI actionable and how it has been used and should be used in defensive cybersecurity practices. This insight will be used to create a set of recommendations that aim to increase the actionability of predictive CTI by enabling better alignment of intelligence outputs to consumer needs and by identifying implementations and processes that make use of it the best. Using the outcomes of this thesis, producers of predictive CTI can find ways in which to improve the value of their outputs, and consumers can gain insight into how to initially set-up, expand, and optimise their predictive CTI related capabilities.

The thesis is in English and contains 68 pages of text, 8 chapters, 5 figures, 6 tables.

Table of Contents

List of figures

List of tables

1 Introduction

The introduction and proliferation of artificial intelligence (AI) and machine learning (ML) methods have fundamentally changed the cyber threat landscape [1]. These technologies have enabled a boom in cyber attacks, providing benefits in attack speed, scale, coverage, and sophistication [2]. The capabilities of AI and the malicious actors using it are constantly evolving [3], which can make it hard for organisations to keep pace and remain vigilant against them.

To alleviate this, many organisations turn to cyber threat intelligence (CTI), which can be described as the proactive gathering, analysing, and sharing of information related to current and emerging cyber threats [4]. It can be a useful tool for organisations to anticipate potential attack vectors and allocate defensive resources accordingly [5]. CTI often looks at the data and trends of current threats and attacks [4]. Predictive CTI involves making predictions about what tactics, techniques, and procedures (TTPs) are likely to be used by threat actors in the future [6].

Despite the uptick in calls for predictive CTI use, it remains under-researched how useful predictive CTI differs from traditional CTI. How, or whether, these predictions meaningfully improved organisational resilience is also under-researched. Without these kinds of analyses, there lacks an evidence-based understanding of which forecasting approaches are reliable, which types of predictions offer the most value, and where current predictive CTI falls short in improving defensive cybersecurity. An analysis into predictive CTI and how it is currently being used is therefore important for those who use it in their decision-making process to better determine the quality and actionability of it, and better make use of it. This thesis can also be important in aiding producers of CTI make more actionable predictions of future attacks and novel threats.

1.1 Goal and Scope

The overarching goal of this thesis is to find out how to make predictive CTI more actionable for the practitioners that produce and consume it. This includes the intelligence analysts and vendors that produce and disseminate it, those who oversee and design its consumption from the strategic level, and those who are tasked with ingesting and acting upon it. This overarching goal can be split into three smaller, more focused, aims. The first is to determine the applicability of traditional CTI actionability research to predictive CTI. The second is to analyse how predictive CTI is currently used in practice . The third is to determine the ways producers of predictive CTI can make it more actionable and the ways consumers of predictive CTI can make better use of it. The outcomes of this thesis will therefore be:

- The creation of a set of evaluation criteria for determining the actionability of predictive CTI.
- An analysis of how predictive CTI is used in practice and why.
- Recommendations for improving the actionability of future predictive CTI.

The scope of this thesis partly depends on the definition of the key terms used in the title, namely actionability and predictive CTI. Actionability, in the context of this thesis, is based on the European Union Agency for Cybersecurity (ENISA) definition for actionable information, which is information that "is used to take actions that mitigate against future threats, or help address existing compromises" [7]. Actionability is therefore the efficacy or ability to facilitate decision-making processes and guide implementation of controls to mitigate existing and future threats. Cyber threat intelligence in this thesis is defined as the "obtaining, processing, evaluating, and disseminating information about potential risks and opportunities inside the cyber domain" [8]. Predictive CTI differs from this definition slightly by being more "probabilistic, [time-aware and] based on historical patterns and behavioural signals, rather than descriptive statistics or correlational analysis" [6]. As a main driver for the push for predictive CTI and CTI in general is the transformative impact of AI on the threat landscape [2, 3], focus will be on sources and data from this 'era of AI', which is

loosely defined in this thesis as the period of time from 2018 to the present. This is to limit the predictive corpus while keeping a timeline that includes early adoption of ML and AI and its rise to widespread use [9].

This study is also built on a few key assumptions. One is that there are a sufficient number of historical predictive CTI reports that exist within the defined timeframe and are publicly accessible. Another assumption is that intelligence analysts, and those who consume predictive CTI in their decision-making process, are willing and able to provide insights into concrete measures taken in response to predictive intelligence.

1.2 Research Questions

The lack of analysis of past predictive CTI leaves a gap in the understanding of how much and well predictive CTI has been used, the reasons why, and how to improve future predictive CTI. By addressing this problem, this study will create insight into how predictive CTI has shaped the cybersecurity resilience of organisations and the ways to increase its effect.

In order to better approach this problem, and to achieve the previously outlined aims of this thesis, smaller research questions have been devised to guide and focus the work. Each research question tackles different aspects of the problem and is listed below:

- RQ1: How do current actionability criteria for CTI apply to predictive CTI?

Answering this research question will involve determining the current criteria used to evaluate the actionability of CTI. For each criterion, the extent to which predictive CTI has that quality will be investigated. The scope of this research question also involves investigating whether there are any criteria that apply to predictive CTI that may not be as relevant for non-predictive CTI, and the relative importance of each criterion.

- RQ2: How is predictive CTI used in practice?

Answering this research question will look at the extent to which various organisations make use of cyber threat intelligence. This includes if they use it at all, how they use it, and what cybersecurity decisions or controls result from it. In the scope of this research question also lies the perception of predictive CTI and how it compares to non-predictive CTI. This research question is therefore split into the following subquestions:

RQ2.1: To what extent is predictive CTI relied on in cybersecurity related decision-making?

RQ2.2: What decisions are influenced using predictive CTI data?

RQ2.3: What tools, platforms, and processes are used with predictive CTI?

- RQ3: How can predictive CTI be better utilised?

By analysing the trends in the data gathered when answering the first two research questions, relationships between certain characteristics of predictive CTI and their influence on its actionability can be discovered. From this analysis, this research question can be answered through two lenses: how producers of predictive CTI can improve it, and how consumers of predictive CTI can better make use of it. This research question is therefore split into the following subquestions:

RQ3.1: What are the main barriers to utilising predictive CTI?

RQ3.2: How can predictive CTI be improved?

RQ3.3: What are the benefits and drawbacks to current predictive CTI related tools, platforms, and processes?

1.3 Structure of the Thesis

The thesis will be structured as follows. The ‘Background’ chapter will cover the existing literature regarding CTI, predictive CTI, and actionability. The ‘Methodology’ chapter will explain and justify how the research questions were answered, including how and from where the data was collected and how the data was analysed. The ‘Results’ chapter will cover the findings from the research, and the ‘Discussion’ chapter will analyse these findings, relate them to the research questions, and extract

evidence-backed conclusions. The ‘Evaluation’ chapter will view the insights gained in the previous chapters in relation to similar work, and discuss the limitations of the work and its results. Finally, the ‘Conclusion’ chapter will summarise the work, provide answers to the research problem and questions, and highlight the key insights, findings, and conclusions of the thesis.

2 Background

The body of literature involving CTI, its justification in today's cyber threat landscape, and its actionability will be examined in this chapter. Through accomplishing this, the necessary foundations will be established from which this work's contributions can stand. Academic literature was primarily collected through databases and search engines including Google Scholar, IEEE Xplore, ACM Digital Library, and SpringerLink. Additional authoritative sources like the National Institute of Standards and Technology (NIST) and the European Union Agency for Cybersecurity (ENISA) were used when appropriate. Searched keywords included combinations of "predictive", "cyber threat intelligence", "actionability", "proactive cybersecurity", and "artificial intelligence". Common synonyms and acronyms (e.g. CTI, AI, operationalisation) were also searched for. The free model of ChatGPT[1] was also used to find relevant papers, where it was prompted to find papers that related to the topics present in this thesis.

Sources were initially selected based on their relevance to the research questions of this thesis, determined from reading their title, abstract, and conclusion. In some cases the results section was also factored into the initial selection process. Due to this thesis focusing on a post generative AI cybersecurity landscape, preference was made to more recent studies, especially those within the last 5 years. However, exceptions were made for papers highly relevant to CTI and its actionability. Inclusion criteria consisted of peer-reviewed academic publications and recognised industry reports that address CTI, predictive intelligence, AI-assisted cybersecurity, or CTI operationalisation. Sources unrelated to cybersecurity or lacking relevance to intelligence-based security were excluded. Industry and security vendor sources were included selectively after critical evaluation of any potential bias. Remaining sources were then read through in their entirety. Backwards and forwards snowballing was done by looking at the papers that

[1] https://chatgpt.com/

each source referenced and were referenced by. This was done to find additional relevant sources and papers found this way underwent the same selection process.

2.1 The AI-Enabled Threat Landscape

The rapid advancement and accessibility of AI and ML technologies have significantly altered the cybersecurity threat landscape. Applications of AI in cyber attacks are a commonly explored topic in literature. The most discussed malicious uses of AI include its application in phishing and other forms of social engineering [1, 10, 11, 12], deepfakes [11, 13, 14], and malware [15, 16, 17, 18, 19].

Across the literature, a recurring theme is that AI enables cyber attacks of increased speed, scale, coverage, and sophistication [2, 3, 20, 21, 22]. AI technologies can reduce the cost and effort of conducting an attack, and can enable more targeted operations. They can also be used to bolster defensive cybersecurity posture. Therefore, several authors describe a sort of arms race between attackers and defenders [20, 21]. However, there is some disagreement about whether AI-enabled threats are enhancements of existing attack methodologies or should be viewed as fundamentally different. This distinction is particularly important when evaluating predictive CTI and its use, as perceptions of threats play a massive role in influencing organisational security priorities and investments. It is also important to note that most literature surrounding AI-enabled threats relies heavily on conceptual analysis rather than empirical evidence.

2.1.1 Implications for Cyber Threat Intelligence

The increasing sophistication and speed of AI-enabled threats have significant implications for CTI. Traditional cybersecurity approaches have historically focused on reactive detection and response, where defensive actions are taken after indicators of compromise or other malicious activities are identified. However, the speed and adaptability of modern attacks increasingly challenge the effectiveness of reactive security models [8, 23, 24].

Recent literature also increasingly positions predictive CTI as a necessary evolution of traditional intelligence practices in response to the highly dynamic threat environment [6, 24, 25]. There is discussion in the literature about whether current CTI models are sufficient. Traditional CTI frameworks are largely designed around observable indicators and historical incident analysis [26, 27, 28]. This is because these forms of intelligence are empirically verifiable, operationally actionable, and suitable for standardised and structured sharing mechanisms and protocols, characteristics that allow it to support important cybersecurity activities like intrusion detection and incident response. While these frameworks are valuable for situational awareness and operational response, some argue that they are insufficient for dynamic and rapidly evolving threat environments [6, 24]. This discussion highlights that, while predictive CTI is increasingly being promoted as a key component of proactive cybersecurity, there is still a limited empirical understanding of how actionable and operationally useful predictive intelligence is in practice.

2.2 Cyber Threat Intelligence

Cyber threat intelligence is commonly defined as evidence of, and evidence-based knowledge of, threats, threat actors, vulnerabilities, and adversarial behaviour that supports informed decision-making [8, 23, 28]. The exact definition varies slightly across sources, but the general idea remains constant.

A distinction between raw threat data and actionable intelligence is often made. NIST defines CTI as information that has been analysed to support decision-making regarding cybersecurity threats [28]. ENISA emphasises CTI as contextualised and actionable knowledge rather than isolated technical indicators [7].

Cyber threat intelligence is commonly categorised into (from higher level to lower level) strategic, operational, tactical, and technical level intelligence [23, 26, 29]. Each level serves different audiences and supports different forms of decision-making. The technical level is often focused on indicators of compromise, malware signatures, and observable attack patterns [23, 26], while the strategic level focuses on things like

broader threat trends and adversarial intent. What actionability might mean therefore differs depending on which level you approach CTI from, and this discussion is important to have for predictive CTI as well, as the uncertainty inherent with predictive CTI might complicate notions about actionability applicable to traditional intelligence.

2.2.1 Actionability of Cyber Threat Intelligence

Actionability, or the capacity to enable and support concrete defensive action, is widely viewed as key to making CTI valuable [4, 7, 29, 30]. It is therefore important to investigate and ensure that predictive CTI is maximally actionable as well. The literature commonly identifies relevance, timeliness, accuracy, completeness, and contextualisation as key components of actionable intelligence [4, 7, 31]. How important each component is varies from one author's opinion to another. ENISA emphasises contextual relevance and operational usability [7], while others [31] focus on ingestibility.

While these frameworks demonstrate substantial agreement regarding the core ideas, inconsistencies and differing views exist regarding how CTI should be measured and operationalised. These frameworks also focus primarily on traditional CTI and therefore do not fully address aspects more unique to predictive CTI like predictive uncertainty, the temporal dimension, and AI-assisted intelligence generation. Therefore it is important that research is carried out to see how well these existing frameworks apply to predictive CTI and what changes need to be made to them to better incorporate it.

2.2.1 Organisational Use of Cyber Threat Intelligence

It was found that organisations utilise CTI in different ways depending on factors like maturity, resources, and threat exposure [32, 33, 34]. Large-scale surveys investigating the actionability of CTI by the SANS Institute [5, 32-37] identified rates of adoption of CTI as well as its sources, main use cases, and implementation details. While these surveys and other relevant literature [38] provide a strong foundation and starting-off point for this thesis, the existing research largely views CTI as a whole, lumping together traditional and predictive threat intelligence. Due to the unique characteristics

and challenges of predictive cyber intelligence, similar investigation into specifically predictive CTI is warranted.

2.3 Predictive Cyber Threat Intelligence

Predictive CTI refers to intelligence capabilities for the purpose of anticipating future threats, adversarial behaviour, or attack trends before they materialise [6, 39, 40]. Unlike traditional CTI, which focuses on concrete incidents and indicators, past and present, predictive CTI attempts to forecast future risks based on historical data, behavioural patterns, and contextual analysis [6]. It is important to note that both predictive and traditional CTI enable a more proactive approach to cybersecurity, but predictive CTI differentiates itself by being validated in a "time-aware, forward-looking setting" [6]. A defining characteristic of predictive intelligence is its probabilistic nature. Predictions are inherently uncertain and cannot guarantee future outcomes. This differentiates predictive CTI from traditional intelligence approaches and creates significant challenges regarding trust, validation, and operational decision making [6].

Literature about predictive CTI is dominated by research into improving the methods and technologies that enable it. Research in this area tends to focus on how AI and ML technologies can be applied in things like behavioural analytics and threat analysis [39, 40, 41]. AI assistance in analytics and its impact on scalability and automation within CTI workflows, especially when dealing with large amounts of data, are also present in the literature [40, 41]. Past studies have focused on improving the content and technology behind predictive intelligence, so an evaluation of the predictions themselves and what has come from them and why is a gap in the literature. There is also a limited amount of research into how specifically predictive CTI is currently used in practice.

2.4 Novelty of the Study

This study attempts to fill in a gap in cybersecurity research by looking at the actionability of predictive CTI and how it can be improved. By analysing existing literature on AI-enabled threats, traditional and predictive CTI, and actionability, several gaps were revealed. First, research into the actionability of predictive CTI is limited. Frameworks exist for traditional CTI but they do not account for the unique challenges and characteristics of predictive CTI, and as such their applicability to predictive CTI needs to be investigated. Furthermore, research into organisational use of CTI similarly does not differentiate between traditional and predictive CTI, so there lacks an evidence based view of the state of predictive intelligence in practice. Research into predictive CTI is predominantly related to its technical feasibility and the techniques and tools to improve its content. Therefore, a non-technical approach to improving predictive CTI and its actionability provides novel insight to the field. These gaps all inform the chosen research questions and methods of this thesis.

3 Methodology

This thesis utilises a primarily qualitative mixed-methods approach for gathering the relevant data. It is a survey based approach that involves a literature review, a questionnaire, and interviews. The focus on the qualitative was due to many reasons. First, the thesis deals with concepts for which there is no standardised method for translating into a series of metrics and numbers. For example, actionability as a concept can not be measured objectively as it is heavily context dependent [42], relying on things like organisational practices and human judgement. Actionability might also mean different things to different organisations or stakeholders [42], which further decreases the suitability of quantitative approaches when working with this topic. Qualitative methods are therefore ideal for capturing the nuanced experiences, interpretations, and constraints that influence how CTI is used in practice. In addition, many other papers looking at similar topics related to cybersecurity on the organisational level also employ qualitative methods [43], suggesting that this approach is an effective tool in this domain.

Focusing on the qualitative aspects of the topic also allows the analysis of past predictive CTI to focus on general trends and patterns rather than specific statistics. One reason why this is a benefit is that it avoids potential numerical inaccuracies that may arise from the under-reporting of incidents. Another reason is that determining causal links between a specific prediction and a specific decision done by an organisation is difficult, and the necessary information in sufficient detail to make that link is not often made readily accessible. Therefore such attribution can only come from looking at trends across multiple organisations, or from confirmation during interviews with decision-makers.

In the context of this thesis, a questionnaire was used to gain qualitative and quantitative insights related to the role of predictive CTI in cybersecurity decision-making. This questionnaire, alongside research into available policy/strategy documents, informed the development of semi-structured interviews of experts that are responsible for cybersecurity decision-making or are operating in industry. The interviews provided deeper insights about how predictive CTI is viewed, evaluated, and operationalised, and expanded on any insights or patterns discovered in the previous research. A breakdown of the different methods used for data collection and analysis and what research questions they answer is as follows.

3.1 Questionnaire

An online questionnaire was conducted to gather empirical data on the use and perceived actionability of predictive CTI. The questionnaire was designed to capture patterns across organisations of differing levels of cybersecurity maturity, and included questions regarding the extent of CTI adoption, its role in decision-making, and the factors that influence its practical utility. As such, the questionnaire serves a descriptive and exploratory function, providing a wider contextual basis for answering all three research questions and informing the direction of the interviews, which went into more depth on the same topics.

The target population for the questionnaire consisted of cybersecurity professionals involved in decision-making processes, particularly those with exposure to, or responsibility for, handling CTI. This includes roles such as security analysts, threat intelligence analysts, security engineers, incident responders, and chief information security officers (CISOs). Individuals in these positions are most likely to interact with CTI in a practical context, whether through direct analysis, operational use, or strategic decision-making. The questionnaire also targeted a range of organisations across a wide range of cybersecurity maturity levels. This was to ensure the questionnaire captured as accurate and holistic a snapshot of CTI use in practice as possible, and to enable questions concerning the use of predictive CTI (for organisations with a higher level of

cybersecurity maturity) and the barriers to using predictive CTI (for organisations with a lower level of cybersecurity maturity).

The questionnaire employed non-random sampling methods including purposive, snowball, and convenience sampling. Purposive sampling was used as a clear profile of who has the capacity to properly answer the questions was known. Snowball sampling was used to leverage the tight-knit nature of many cybersecurity communities. Cybersecurity professionals are also less likely to click on a random questionnaire link, so having it come from a trusted source (i.e. a colleague or trusted acquaintance) can increase the rate at which the questionnaire is engaged with. Convenience sampling was used to maximise reach, and establish as wide a base as possible for snowball sampling to occur.

3.1.1 Question Design

The questionnaire was structured into three main sections, and consisted of both open- and closed-ended questions. A copy of it is included in Appendix 2. To maximise engagement, questions were generally designed to avoid asking for information that a questionnaire participant may not be willing to give (e.g. sensitive operational data). In addition, all questions were designed so that participants could provide answers that were as general or specific as they were comfortable with, and no question was required. This was to mitigate response bias and to maximise questionnaire engagement.

The first section was for collecting background information on the respondent, including their role or job title, their organisation and the industry it would fall under, and the size of their organisation's dedicated security team. These questions were used to ensure that respondents were part of the target population and that their responses were relevant. They were also used to provide context for analysis. Dedicated security team size was used as one of the metrics for determining the cybersecurity maturity of an organisation instead of a self-report of cybersecurity maturity on a Likert scale. This was because cybersecurity maturity isn't a properly defined term that can be evaluated objectively and consistently across different organisations, so a more objective heuristic (security team size) would enable better comparison between responses. Using an objective heuristic as such also leads to less bias in survey responses that could arise

from participants misrepresenting themselves or evaluating cybersecurity maturity using different criteria.

The second section focused on CTI usage and its objective was to determine the extent of predictive CTI usage. It provided answers related to RQ2. It asked whether respondents' organisations use CTI and how much predictive CTI compared to non-predictive CTI is used. Terms like predictive CTI were defined and examples of predictive CTI and non-predictive CTI were provided to ensure conceptual clarity and consistency across participants by reducing ambiguity. This was to lower response bias and ensure questionnaire participants without exposure to such terminology (e.g. from organisations with lower cybersecurity maturity) but who still engage with it are not excluded. Leading to increased validity and making sure that any observed patterns reflect differences in practice rather than differences in understanding.

The third section focused on the actionability of CTI: how it is used and operationalised, the qualities which make it actionable, the barriers to use it, its perceived benefit, and how it could be improved. It provided insights related to RQ1, RQ2, and RQ3. The section started with a definition of actionability to remove ambiguity and increase conceptual consistency across survey participants and responses. The first question in this section was a multiple choice question asking for the ways predictive CTI is used. The possible answers were based on the most popular use cases for CTI found in literature and other surveys on CTI [5, 32-37]. Multiple answers could be selected to get the most holistic view of predictive CTI use and allow maximum flexibility for participants. Participants also had the option to add additional, custom answers to make sure that responses could be as accurate to real life operations as possible. The second question asked for the (types of) decisions influenced by predictive CTI data. Similar to the previous question, it was multiple choice, allowed multiple and custom answers, and the answers choices were based on past surveys and literature about CTI. The next question asked participants to rank on a Likert scale how influential predictive CTI was in the decision process for them. This was done on a scale from 1 to 7 where 1 meant "Not at all" and 7 meant "Extremely". This question provides insight to how predictive CTI is perceived. The following question asked participants to rank the factors that make predictive CTI most actionable in order of importance. The factors chosen were based on existing literature about evaluating the actionability of CTI. Participants could

also add any additional factors that they deemed to be important, striking a balance between building on existing literature and capturing new practitioner-based perspectives.

3.1.2 Distribution Methods

The questionnaire was distributed across many different channels using purposive, snowball, and convenience sampling. The shared questionnaire always came with a message emphasising who the target population is, that participation is voluntary, and that independently sharing the questionnaire is encouraged. In the message there also were assurances about how the data gathered for the questionnaire would only be analysed in aggregate and that participants and the organisations they represent won't be identifiable in the final thesis and any publications. It was also promised that all questionnaire participants would receive the findings of this thesis upon its completion, and this was done to encourage participation.

The questionnaire was shared via mail with the Tallinn University of Technology Master's cybersecurity cohort and with the individuals who teach cybersecurity related subjects at the Tallinn University of Technology and Tartu University (totalling 64 recipients). This sample was chosen due to the ease of assembling the mailing list (convenience sampling) and the fact that many are working in the cybersecurity field or have connections in industry to which they could forward the questionnaire. The questionnaire was also shared via a LinkedIn post, and reshared by multiple people in the cybersecurity industry to expand its reach. These posts combined had hundreds of impressions. The questionnaire was also shared to the mailing list of an Estonian academic corporation which contained over three hundred people [44]. I also attended cybersecurity events in person to promote the questionnaire, including a RIA CyberMeetUp organised by the Republic of Estonia Information Security Authority. By contacting the Cybersecurity Community Coordinator of that event, the questionnaire was also shared to the Estonian-speaking cyber expert community. The questionnaire was also directly shared with individuals who fit the target population whose information and contact details were found on sites like LinkedIn. Contacted individuals included chief information security officers (CISOs) and cybersecurity alumni.

While the sharing of the questionnaire was encouraged, the full extent to which it was forwarded and shared by those who initially received it is unknown. It is known, however, that some amount of snowballing did occur, as some of those who were contacted reported sharing the questionnaire with colleagues or industry group chats and channels.

3.2 Interviews

To complement the questionnaire and provide deeper insight into how predictive CTI is perceived and used, a series of semi-structured interviews was conducted. The interviews also served to gain understanding of the patterns observed in the questionnaire responses, exploring the underlying reasoning and contextual factors behind the data. Interviews are a common research method in similar cybersecurity research [43] and are a good fit for obtaining the qualitative data needed to answer the research questions of the thesis. Interview participants were chosen from questionnaire respondents who indicated a willingness to engage in follow-up discussions. This choice was opt-in and completely voluntary. As interview participants make up a subset of the questionnaire participants, the target population for the interviews also included security analysts, threat intelligence analysts, security engineers, incident responders, and CISOs.

3.2.1 Interview Design

A semi-structured interview format was chosen to ensure consistency and enable comparisons across interviews. Semi-structured interviews also have the flexibility necessary to explore topics in depth and get the most out of each interview participant. This is especially important as there is a lot of disparity between different organisations when it comes to predictive CTI adoption and overall cybersecurity maturity, [32-37] meaning not every interviewee can answer every question related to the topic of predictive CTI actionability.

Each interview followed a guide that was generally split into two categories of questions. The category of questions mainly used during each interview depended on the interviewee and their cybersecurity maturity and adoption of predictive CTI. These data points were probed for early in each interview to form an initial impression that was clarified and refined as each interview went on. The two categories are defined by their target populations, with one category meant for organisations with lower cybersecurity maturity or predictive CTI adoption and the other meant for organisations with higher cybersecurity maturity or predictive CTI adoption. The benefits of this distinction are that it accounts for the wide spectrum of predictive CTI usage in industry, enables gathering of different yet relevant types of insights, and allows for more nuanced analysis

For those with limited adoption of predictive CTI, questions focused on understanding how cybersecurity decisions are made in the absence of, or with minimal reliance on, predictive CTI. This category of questions also attempted to identify the conditions under which such CTI could become actionable. The questions explored the interviewees' existing prioritisation processes, information sources, and past real-world decision-making scenarios, establishing a baseline for how organisations operate without predictive CTI. The questions also investigate the perceived value of predictive CTI, the types of predictions that would be most useful, and the barriers preventing adoption. Interview participants were also asked to describe what characteristics predictive CTI would need to have in order to be actionable, explicitly probing criteria such as relevance, timeliness, accuracy, completeness, and ingestibility if necessary. This allows for the identification of the preconditions for actionability and use of predictive CTI, directly informing RQ1 by examining how established CTI actionability criteria apply in contexts where predictive CTI is not yet operationalised. Furthermore, by exploring perceived benefits and adoption challenges, this track contributes to RQ3 by highlighting gaps between current capabilities and desired future use.

For those with high adoption of predictive CTI, the questions primarily examined how predictive CTI is currently used and evaluated in practice. Questions covered defining actionable CTI from a practitioner perspective, distinguishing between actionable and non-actionable intelligence, and understanding how predictive CTI differs from non-predictive CTI in real-world contexts. Participants are asked to provide concrete

examples of decisions influenced by predictive CTI and to describe how such intelligence is internally evaluated, probing for key criteria such as relevance, timeliness, accuracy, completeness, and ingestibility. This enables a direct evaluation of how existing actionability criteria manifest in practice, addressing RQ1 from an applied perspective. At the same time, by analysing how predictive CTI informs operational and strategic decisions, this track directly addresses RQ2. Finally, the identification of limitations, evaluation practices, and requirements for taking action contributes to RQ3 by revealing where predictive CTI falls short and how it could be improved to better support cybersecurity decision-making.

Some questions and topics were covered regardless of cybersecurity maturity of predictive CTI adoption. These included the core characteristics of actionability, desired improvements and changes to predictive CTI, and perceptions of ideal CTI capabilities. This captures a wide range of perspectives on both the intrinsic qualities of predictive CTI and the ways in which it is produced and ingested. All interview participants were also asked about the future trajectories of their organisations and predictive CTI in order to gain insight into future adoption trends and the evolving role of predictive CTI.

3.2.2 Interview Process

The interviews were conducted remotely via video conferencing tools such as Microsoft Teams. The interviews themselves took the form of video calls and lasted around 30-40 minutes. To allow for participants to talk more freely about the details of what they do, what they think, and past incidents, the interviews were not recorded or transcribed. Notes were taken manually during the interview, with clarification of participants' responses being asked for as necessary.

Interviews began with a quick introduction to and reminder of the thesis topic. Interviewees were also made aware that the interview wouldn't be recorded and that information that could personally identify them or their organisation would not be included within the thesis or any publications. The first questions of each interview were about the interviewee's role in their organisation and how it relates to cybersecurity. This was done to ensure that interviewees and their responses were valid for use in this thesis. Interviewees were also questioned on their organisations'

cybersecurity approaches, values, and resources, alongside questions about their overall use of predictive CTI. These questions served to establish the interviewees' cybersecurity maturity and adoption of predictive CTI in order to determine which category of questions to focus on. During the main part of each interview, which category of questions to draw from was adaptable and changed depending on the interviewee's answers to previous questions. The flow of the interview and the topics that were perceived as most likely to result in new insights were also factored in when choosing which questions to ask and in which order to ask them. In addition, when an interviewee's response seemed implicitly codeable to a specific theme or actionability criterion, follow-up and probing questions were used to clarify meaning, reduce ambiguity, and ensure accurate coding and analysis. This was done to ensure that the notes taken during the interview and the future analysis of them reflected the interviewee's intended meaning and perspective, rather than researcher interpretation.

4 Results

This chapter will introduce the findings from the questionnaire and interviews. First, an overview of the questionnaire respondents and the interviewees will be given, then the results of the questionnaire and interviews will be presented. The results will be organised by which of the three main research questions they relate to in order to improve the readability of the chapter. In some cases, exact details included in respondents' answers in the questionnaire and interviews have been redacted in this publication in order to prevent identification of the individuals or organisations that provided them.

4.1 Questionnaire Respondent Overview

The total number of respondents to the questionnaire was 21. Of those 21, one response was excluded from analysis due to it being empty and thus possessing no analytical value. The remaining twenty responses were determined to be legitimate and were used for analysis. In order to maximise response rates, no question was made mandatory. This ensured that nobody would be deterred by having to provide potentially sensitive information, but means that not every question was answered by each respondent. Therefore the numbers may not add up exactly to 20 on every question.

The questionnaire was responded to by mostly cybersecurity professionals at the executive or managerial level, and respondents were classified into 6 main roles, with similar job titles classified together. The exact roles and counts are presented in Table I. The most commonly reported role was Chief Information Security Officer (CISO), with five respondent's responses being classified as such. Three of the respondents were cyber or threat intelligence analysts, three were security specialists, and two were IT

managers. In addition, one respondent was a data protection officer (DPO) and four were classified as cybersecurity experts, which includes red/blue teamers and SOC analysts.

TABLE I

Roles among questionnaire respondents

Role	Count
CISO	5
Intelligence Analyst	3
Security Specialist	3
IT Manager	2
DPO	1
Cybersecurity Expert	4

The most represented industry sector amongst the respondents was the government and public sector, with seven followed by cybersecurity service providers and healthcare. The exact industries and counts can be read from Fig. 1 below.

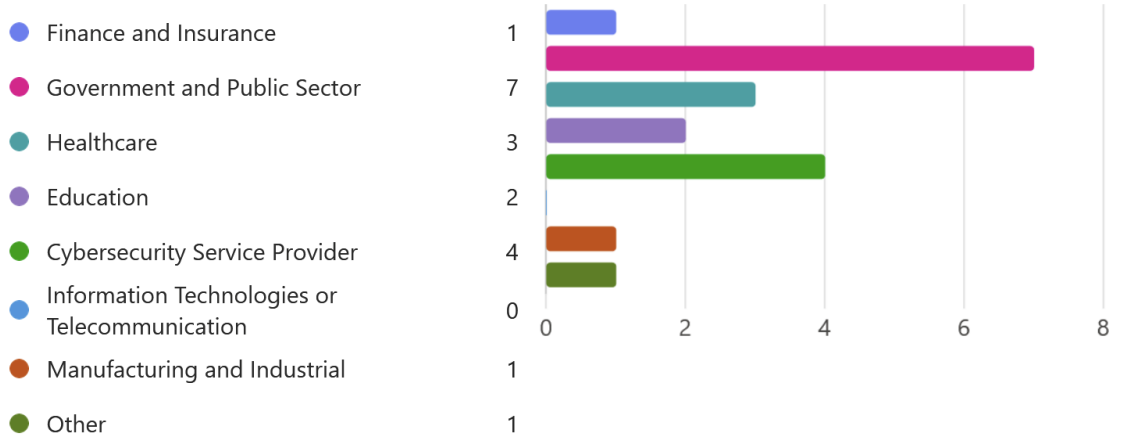

Fig. 1. Industries in which questionnaire respondents work in.[2]

To aid in determining the cybersecurity maturity of each organisation, respondents were asked for the size of their dedicated security team. The most commonly chosen option was between 10-20 people. One respondent said their organisation had a dedicated

[2] The industry listed as 'Other' is sales.

security team of one person, while the highest reported dedicated security team size was 500 (under other in Fig. 2). Other answers in the 'Other' category include a security team size of a few hundred, and a comment that the security team is outsourced.

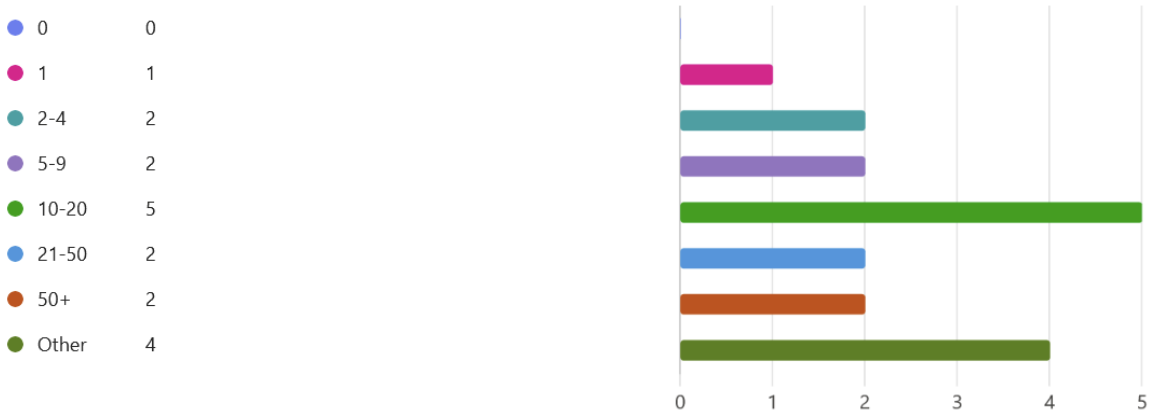

Fig. 2. Responses to the question about dedicated security team size.

In response to the question "Does your organisation use CTI?" 15 respondents said their organisation does and one said that their organisation does not. One respondent answered "I don't know" to the question and three did not answer. This information is displayed below in Fig. 3.

Fig. 3. Respondents' organisation's use of CTI.

4.2 Interviewee Overview

In total, six interviews were conducted. Interviewees were chosen from the subset of questionnaire respondents who indicated they were willing to further participate in the

study and ready for an interview. To protect the anonymity of those that were interviewed, information that may identify the interviewee or the organisation they work for has been generalised or excluded from this publication. To make referencing the conducted interviews and the insights gained from them easier, they have been labeled interviews 1-6.

The roles of each of the interviewees, as well as some information about their organisations, are displayed below in Table II. Overall, the interviewees represented a range of sectors including finance, healthcare, and government. A wide range of experience in the cybersecurity field was also observed in the interviewees, ranging from a couple years to over a decade. In addition, interviewees represented organisations of various sizes and cybersecurity maturity.

TABLE II

Interviewee overview

ID	Role	Organisation Info
Interviewee 1	Cybersecurity specialist, security team lead	Public sector
Interviewee 2	CISO	Large healthcare provider
Interviewee 3	CISO	Finance
Interviewee 4	Intelligence analyst	Cybersecurity service provider
Interviewee 5	Information security specialist	Public sector
Interviewee 6	CISO, former intelligence analyst	Public sector

4.3 Results to Research Question 1

This subchapter covers the results and findings of the questionnaire and interviews that relate to RQ1, “how do current actionability criteria for CTI apply to predictive CTI?” The survey questions that provide insights related to this research question are questions

11 and 12. Question 11 asked respondents to rank factors that make predictive CTI actionable while question 12 asked respondents to list any additional factors that may influence this actionability. The results to question 11 are shown below in Table III.

TABLE III

Rank distribution for factors that make predictive CTI actionable.

Factor	Percentage of responses by rank						Mean Rank
	1st	2nd	3rd	4th	5th	6th	
Relevance to organisation and context	55.6%	16.6%	22.2%	5.6%	0%	0%	1.8
Accuracy and confidence level	16.6%	33.4%	22.2%	22.2%	5.6%	0%	2.9
Timeliness	11.1%	33.3%	22.2%	22.2%	5.6%	5.6%	3.2
Ingestibility (ease of use)	11.1%	11.1%	5.6%	22.2%	16.6%	33.4%	4.2
Specificity	0%	5.6%	27.8%	5.6%	22.2%	38.8%	4.6
Completeness	5.6%	0%	0%	22.2%	50.0%	22.2%	4.8

From the above table it can be read that the factors in order of importance are relevance, accuracy and confidence level, timeliness, ingestibility, specificity, and completeness. This ordering is generally agreed upon across all responses. Relevance to organisation and context was ranked the most important by the majority of responses. Accuracy and timeliness were generally evaluated similarly, being deemed the second or third in over half of the responses. The last three factors were all ranked similarly and much lower than the top three factors.

Survey responses to question 12 indicate that additional factors influencing the actionability of predictive CTI include the provision of recommended mitigation controls, contextual risk analysis, and evidence that a human analyst was involved in the intelligence production process.

Across the interviews, several evaluation criteria emerged. Interviewees 1 and 5 both explicitly emphasised the importance of contextual and organisation fits, stating that

predictive CTI must be about similar organisations, sectors, or infrastructures in order to be actionable. Sector and geographically relevant intelligence was also said to be valued. Interviewees also said that without this relevance it is a lot harder to map predictions to their concrete environments and act on them.

Another evaluation criteria extracted from the interviews was timeliness. Interviewee 1 mentioned that occasionally they would receive intelligence reports about a threat after already coming in contact with it, rendering the report not very useful. Conversely, they also said receiving intelligence about potential threats, say, months in advance might also lead to minimal actions being taken due to other, more current issues being more of a priority.

The accuracy and confidence of predictions was also brought up as an important evaluation criterion. Multiple interviewees presented a skepticism about the accuracy of predictive CTI data, with interviewee 1 even comparing it to a weather forecast. This comparison also highlights the fact that, when dealing with a potential threat, you have to consider the possibility that the prediction is a false alarm. A consideration that isn't required when dealing with non-predictive CTI based on real data and IOCs. Interviewees 3 and 4 explicitly mentioned that reducing false positives and managing signal volume are a priority, furthering the notion that accuracy is important for the actionability of predictive CTI. Interviewee 2 also mentioned that high-risk environments (like in the healthcare context) demand high levels of certainty and caution, especially if the data source is predictive in nature.

In the same vein as accuracy and confidence, interviewee 4 said that reliability and trustworthiness of the source of predictive CTI data is the factor that they consider first. To add to this, interviewee 6 said the two main factors to look out for when evaluating predictive CTI sources are the presence of non-public data and high quality analysis. Having access to data others do not have access to makes a source of CTI much more valuable, and analysis. Interviewee 6 also added that, while ML can do decently well at analysis, analysis from a human with lots of experience, knowledge, and context is superior and preferred.

All interviewees mentioned the inclusion of specific mitigation measures as an important feature of actionable predictive CTI. In addition to controls, several interviewees also mentioned a preference for data with clearly defined impacts for potential threats and a recommended time window for taking action. Interviewees 5 and 6, who are both currently engaged with raising their organisations' cybersecurity maturity, stated that they are striving towards better integration of CTI into their workflows. Interviewee 2 also brought up that sometimes traditional predictive CTI reports may be too technical for smaller organisations who may not have a lot of in-house technical expertise. These remarks suggest that the ease of ingestion and integration of predictive CTI data should be considered when evaluating it.

4.3 Results to Research Question 2

This section covers the results and findings of the questionnaire and interviews that relate to RQ2, "How is predictive CTI used in practice?" and its subquestions. This research question correlates to the questionnaire questions numbered 7, 8, 9, and 10. The results to these five questions are visualised in Table IV, Fig. 4, Fig. 5, and Table V, respectively.

Table IV

Responses to question 7.

	Percentage of responses							Mean
	1	2	3	4	5	6	7	
How much of the CTI that your organisation uses is predictive?	22.2%	11.1%	16.7%	44.4%	0%	0%	5.6%	3.1

For question 7, an answer of 1 meant "none of the CTI we use is predictive", while an answer of 4 corresponded to "we use a balanced mix of predictive and non-predictive CTI", and an answer of 7 "all of the CTI we use is predictive". Almost every respondent said that they use less predictive CTI than non-predictive CTI, with one respondent

claiming that all the CTI they use is predictive. Roughly one fifth of respondents who use CTI do not use predictive CTI.

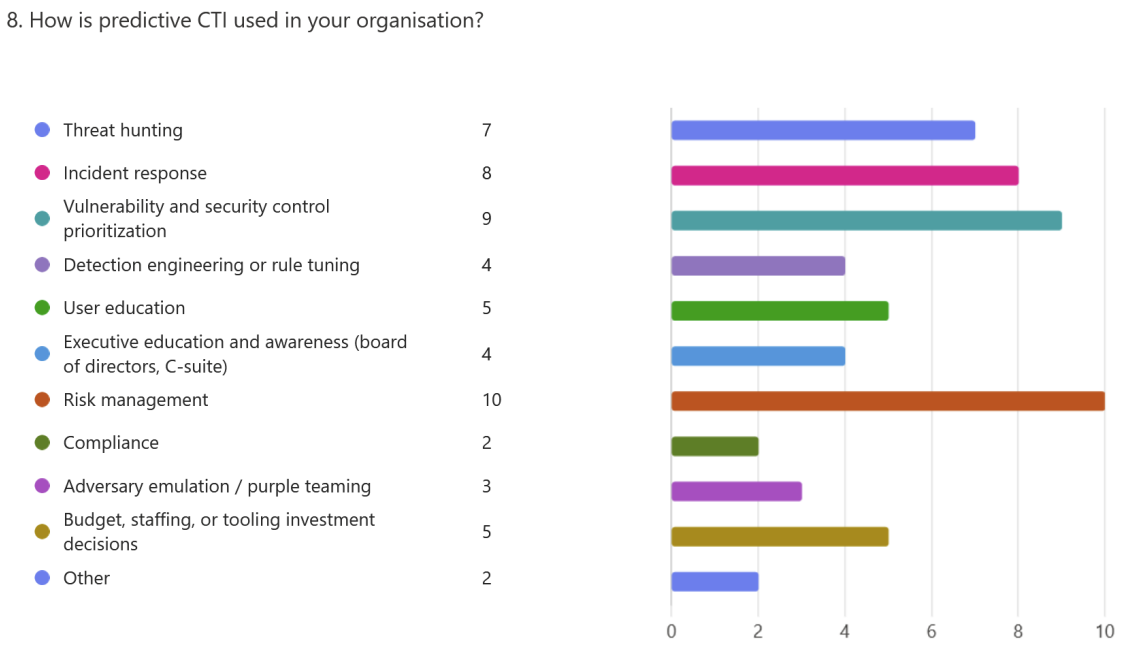

Fig. 4. Response data for how predictive CTI is used.

Question 8 of the questionnaire asked respondents to select the use cases of predictive CTI that reflect their operations. As seen in the above figure, predictive CTI is used in many different ways, with the most popular use cases being for risk management, vulnerability prioritisation, incident response, and threat hunting, in that order. The two responses that selected the “Other” option used it to say that they did not use predictive CTI.

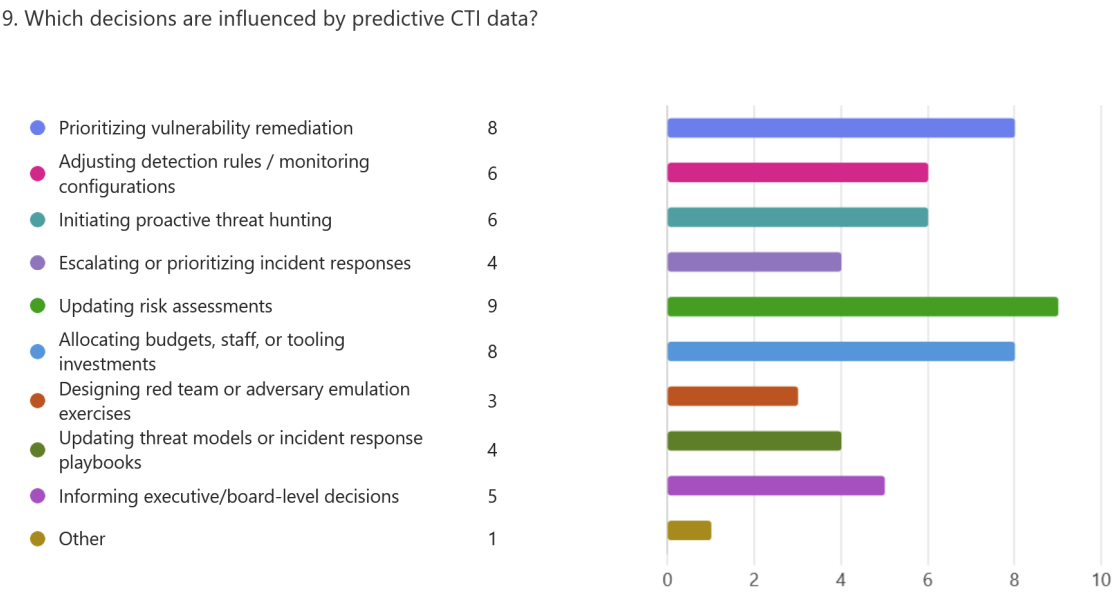

Fig. 5. Response data for the decisions influenced by predictive CTI data.

Question 9 asked respondents about the decisions they apply predictive CTI data to. As seen in Fig. 5, predictive CTI can be used to inform a variety of decisions. The three most common answers among respondents were updating risk assessments, allocating budgets and other investments, and prioritising vulnerability remediation.

TABLE V

Responses to question 10.

	Percentage of responses							Mean
	1	2	3	4	5	6	7	
How influential is predictive CTI in the decision making process?	11.1%	16.7%	27.7%	22.2%	16.7%	0%	5.6%	3.39

For question 10, an answer of 1 meant “not at all” and an answer of 7 meant “extremely”. As seen in Table V, predictive CTI data in most cases plays a relatively minor role in the decision making process.

Among the interviewees there was a range of predictive CTI usage and experience. Interviewees 1, 3, 5, 6 worked for organisations with little or no predictive CTI use. Of these, interviewee 5’s organisation had little use of non-predictive CTI, while

interviewees 1 and 6 had moderate use of non-predictive CTI and interviewee 3 had high use of traditional CTI. Interviewee 2 had moderate use of predictive and non-predictive CTI while interviewee 4 had high use of predictive and traditional CTI. It was found that CTI usage has a positive correlation with the cybersecurity maturity of an organisation, with those with lower maturity levels generally employing less CTI and vice versa. It was also found that all interviewees used less predictive CTI than its non-predictive counterpart.

Interviewees with a moderate use of predictive CTI mainly used it for risk and vulnerability management. Using it mainly at the strategic level to guide prioritisation of which threats to focus on and which mitigating measures to implement. It can also be used at the operational level when it comes to guiding threat hunting and updating detection rules. At a moderate level of maturity, predictive CTI integration can be quite informal, but generally follows a three step workflow where predictive CTI about a potential future threat is consumed by a human who then sends it to leadership or someone in a technical role who then checks if the threat is relevant to their systems. If it is relevant, then the threat is addressed. Interviewee 3 mentioned using platforms such as Microsoft Sentinel[3], MISP[4], and Cortex[5] for IOC gathering and enrichment, and TheHive[6] for case management. They also mentioned platforms such as DarkTrace[7] for the automatic integration of predictive CTI, but cited the platform's automatic implementation of controls based on these metrics as a reason why they couldn't use it. Sources of CTI tended to come from bilateral agreements with other organisations, and public threat reports from agencies such as ENISA or the Estonian Information System Authority (RIA).

Interviewee 4, as someone working for an organisation with high maturity of predictive CTI usage, was able to share some insights into what the workflow and use cases of predictive CTI are at that level. As a former intelligence analyst for producers and consumers of CTI, interviewee 6 also contributed to these insights. For interviewee 4, sources of predictive CTI included finished intelligence from security vendors (either

[3] https://www.microsoft.com/en-us/security/business/siem-and-xdr/microsoft-sentinel/
[4] https://www.misp-project.org/
[5] https://www.paloaltonetworks.com/cortex
[6] https://strangebee.com/thehive/
[7] https://www.darktrace.com/

open source or from strategic partners), proprietary tools, and CTI feeds. Incoming predictions were always analysed, never automatically ingested, with a focus on establishing the reliability of the source and credibility of the information. IOC related data was rotated according to a calculated IOC lifecycle and algorithms and algorithms are used to automatically assign risk and priority to the incoming signal volume. The most proactive of intelligence is gathered in order to answer a specific question deemed important enough to ask. Predictive CTI, especially at the strategic level, primarily takes the form of written reports, which was also brought up by interviewee 6.

Interviewee 6 also mentioned using social media and RSS feeds as predictive CTI sources, and that predictive CTI also comes in the form of technical IOCs, emerging exploit code, and risk scores for vulnerabilities. The CTI related tools and platforms mentioned by them include Trellix (formerly FireEye)[8], Crowdstrike[9], Recorded Future[10], proofpoint[11], and CISCO Vulnerability Management (formerly Kenna VM)[12].

4.4 Results to Research Question 3

This section covers the results and findings of the questionnaire and interviews that relate to RQ3, “how can predictive CTI be better utilised?” and its subquestions. Findings in this section will form the basis for insight into how producers of predictive CTI can improve its actionability, and how consumers of predictive CTI can better utilise it. This research question correlates to the questions of the questionnaire numbered 13, 14, and 15. The results to question 13, which asked respondents to rank barriers to using predictive CTI in order of significance, are visualised below in Table VI.

[8] https://www.trellix.com/
[9] https://www.crowdstrike.com
[10] https://www.recordedfuture.com
[11] https://www.proofpoint.com
[12] https://www.cisco.com/site/us/en/products/security/vulnerability-management

TABLE VI

Rank distribution for barriers to using predictive CTI

Barrier	Percentage of responses by rank							Mean Rank
	1st	2nd	3rd	4th	5th	6th	7th	
Low confidence or perceived accuracy	16.7%	16.7%	38.8%	16.7%	0.0%	11.1%	0.0%	3.0
Lack of relevant predictive CTI	11.1%	33.3%	22.2%	11.1%	16.7%	0.0%	5.6%	3.1
Poor integration with tools and workflows	22.2%	11.1%	11.1%	5.6%	27.7%	16.7%	5.6%	3.8
Not timely	11.1%	5.6%	16.7%	38.8%	11.1%	11.1%	5.6%	3.9
Lack of resources	27.7%	16.7%	5.6%	5.6%	0.0%	11.1%	33.3%	4.0
Too vague	5.6%	16.7%	5.6%	5.6%	11.1%	27.7%	27.7%	4.9
Incomplete	5.6%	0.0%	0.0%	16.7%	33.3%	22.2%	22.2%	5.3

As seen in Table VI, there is less of a unified opinion among respondents about the barriers to using predictive CTI. This can be seen from the more even spread of rankings for each barrier. According to the mean rank assigned to each barrier, the barriers in order of significance are: low confidence or perceived accuracy, a lack of relevant predictive CTI, poor integration with tools and workflows, predictive CTI being not timely, a lack of resources, predictive CTI being too vague, and predictive CTI being incomplete. The top two ranks by mean rank, low confidence and a lack of relevance, had very similar mean ranks. Lack of resources, which by mean rank was the fifth most significant barrier, was chosen most commonly to be the most significant barrier and also least significant barrier. Exactly half of all responses to this question had lack of resources as a top 3 barrier, while the other half had it in the bottom 4.

Respondents, when asked to list or talk about the main barriers to using predictive CTI, raised many ideas that can be classified into two main points. The first main point is a low confidence of predictive CTI data. Many respondents said that predictions do not consistently prove accurate leading to uncertainty about the return of investment. Others also mentioned the likelihood of false positives as a main barrier to its use. The other main point is a lack of cybersecurity maturity to properly employ predictive CTI. Many respondents listed that limited in-house expertise and other higher priorities were some of the main barriers to them employing predictive CTI. Outside of these two main points, one respondent brought up the fact that most CTI is mainly used operationally and reactively as a barrier to predictive CTI use.

Question 15 asked respondents to suggest what would make predictive CTI more useful to their organisations. One common theme in responses to this question was a desire for better integration of predictive CTI with the organisational context and security stack. Respondents were looking for data more relevant to their specific clients, inventory, infrastructure, and supply chains. Respondents also desired better integration between CTI tools and frameworks and the rest of the security stack.

Another common theme in the responses was that there is a difficulty to implementing predictive CTI due to a lack of knowledge about where to find it and how to use it. Respondents desired for it to be easier to use in practice, citing a need for better availability and broader know-how about predictive CTI. Some respondents also noted that it may become more useful and influential once they achieve greater maturity in their cybersecurity posture. One thing several respondents brought up was a desire for increased formalisation of predictive CTI and its use. Respondents brought up wanting clear courses of action tied to CTI products, and a consistent approach to confidence level and attribution of source. Stakeholder buy-in and formal commitment to regular CTI tasking and ingestion was brought up as well. One more theme that was brought up across multiple responses was a change in the regulations surrounding predictive CTI, especially in the law enforcement and gaming sectors, where use of predictive data, especially if it is machine generated, may not be allowed.

In the interviews, a common barrier to using predictive CTI, especially for those who did not utilise it, was a lack of resources, including time, personnel, and budget. These

claims imply that other aspects of upgrading the cybersecurity maturity of an organisation are valued more highly than predictive CTI. Lack of knowledge and expertise regarding where to obtain and how to ingest predictive CTI data was also a common theme in interviewee responses. Interviewees mentioned a lack of clear use cases for predictive CTI, citing missing playbooks and formal decision frameworks.

The accuracy of predictive CTI was also another concern. Multiple interviewees expressed fears of high false positive rates and the potential blocking of legitimate activity when acting on predictions. Among those with higher utilisation of CTI, integration challenges including restricted automation capabilities and a lack of standardisation across sources also hindered predictive CTI adoption.

5 Discussion

This chapter will be a synthesis of the collected and analysed data into the key outcomes of the study. The combination of findings from the literature, survey, and interviews will guide the creation of a set of evaluation criteria for determining the actionability of predictive CTI, an analysis of how it is used in practice and why, and recommendations for improving its actionability in the future. By accomplishing these outcomes, a better understanding of the intricacies of predictive CTI use and its evaluation will be achieved. This knowledge will provide value to producers of predictive CTI by providing the empirical evidence needed to better align outputs with real organisational needs. Insight into how different organisations use predictive CTI and why can help shape strategic decision-making of producers of predictive CTI to shape their output and services in the future. This knowledge also provides value to consumers of CTI, through the enabling movement towards more empirically grounded evaluation of incoming intelligence and away from more informal, ad-hoc methods of judgement. Organisations of all maturity levels and CTI usage can also use this knowledge to determine where they stand in terms of maturity, and determine what tools, platforms, and processes involving predictive CTI are best for them.

5.1 Evaluation Criteria for Predictive Cyber Threat Intelligence Actionability

One contribution of this thesis is the development of a set of criteria for evaluating the actionability of predictive CTI. The questionnaire and interviews identified that pre-existing criteria for evaluating the actionability of non-predictive CTI, namely relevance, timeliness, accuracy, completeness, and ingestibility, still very much apply

when the nature of intelligence is predictive. Of these, questionnaire responses identified relevance, accuracy, and timeliness as the most relatively important criteria, with strong agreement across respondents. These criteria also all came up during interviews as well. However, it was also found that when CTI becomes predictive additional factors have to be considered. The biggest of which are predictive uncertainty and the temporal dynamic, which change how organisations view and interact with predictive CTI. This change in behaviour when dealing with predictive intelligence necessitates modifications and clarifications to the criteria.

The new set of criteria developed from the synthesis of the results has five main components based on the existing criteria for traditional CTI. The first and most important criterion, relevance, demands contextual alignment of the intelligence and the consuming organisation. Respondents and interviewees repeatedly emphasised that predictive intelligence should ideally be relevant to the specific geography, sector, and technical environment in order to be actionable. Without this alignment, consuming organisations must spend additional, already stretched thin resources in order to attempt to map predictions to concrete systems and controls or just ignore the intelligence entirely. This significantly reduces the likelihood of action.

The next criterion is predictive quality. This criterion accounts for the accuracy and confidence levels of predictions, and the reliability of the source. The inherently uncertain nature of predictive CTI introduces a level of skepticism not typically associated with traditional, especially IOC based intelligence. Concerns around false positives and the credibility of sources were frequently raised, indicating that predictive CTI is held to a higher evidentiary threshold. Indicators that consuming organisations look out for that are more relevant in the predictive context include explainability and transparency of the source, and proof of human oversight in the prediction.

Temporal relevance was also identified as a critical but nuanced factor. Predictive reports that arrive too late lose value while intelligence that comes in too far in advance might be deprioritised in favor of more immediate concerns. This suggests that there is a ‘goldilocks period’ within which predictive CTI is most useful. While a repeated theme in questionnaire responses and interviews was a preference for finished, processed, intelligence with high quality analysis, that kind of intelligence takes time. This means a

balance must be struck between the quality of analysis of predictive CTI and its timing. There was a strong agreement among questionnaire responses and interviewees that completeness being the least significant factor affecting actionability of predictive intelligence (relative to the other criteria), suggesting that consumers might stomach incomplete intelligence as long as it arrives at the right time.

The fourth and fifth components of the new set of criteria for predictive CTI are both derived from the concept of ingestibility, and are decision support capability and operational usability. Decision support capability refers to features of predictive intelligence that reduce the cognitive and operational burden placed on the consumer if they want to ingest predictive CTI data. Examples of such features include risk ratings, impact assessments, recommended mitigation measures, and guidance on when to act. Operational usability refers to whether predictive CTI can be effectively applied, and includes the ease of integration and interpretation. This criterion is difficult to measure and relies heavily on the context of the consuming organisation. Elements of it include integrability with the rest of the security stack, relevance to the technological context, and application in workflows. Challenges related to technical complexity and ingestion were particularly significant for organisations with lower levels of cybersecurity maturity.

Together, these five criteria extend traditional CTI actionability research and incorporate the uncertainty, temporal dynamic, and other challenges specific to predictive CTI. This set of criteria allows for more nuanced understanding of actionability in predictive contexts.

5.2 Analysis of Predictive Cyber Threat Intelligence Use in Practice

The second major contribution of this thesis is an analysis of how predictive CTI is currently used in practice and the factors influencing its adoption.The findings indicate that predictive CTI is underutilised relative to traditional CTI, with most organisations relying more heavily on non-predictive intelligence. Questionnaire responses also show

that predictive CTI generally has a limited influence on the decision-making process, suggesting that its potential is not fully realised.

In practice, predictive intelligence is applied to many types of decisions, but mainly informs vulnerability and risk management. Operational decisions like threat hunting and detection rule updating were also identified. On a more strategic level, predictive CTI can be used during threat landscape mapping to inform security posture and priorities.

A key finding is that use of predictive CTI is strongly influenced by organisational cybersecurity maturity. Interview data revealed a clear pattern where organisations with higher levels of maturity made greater use of both traditional and predictive CTI. Furthermore, use of traditional CTI was always more significant than the use of predictive intelligence, suggesting that predictive CTI is a capability that builds upon existing CTI practices, rather than a standalone solution.

Another insight about predictive CTI use in practice is that related workflows are all human-centric. Across all levels of maturity, incoming predictive intelligence was not automatically acted upon but instead underwent human analysis which looked at things like relevance, credibility, and validity before informing decision-makers. In most cases, this is a reflection of the uncertainty inherent in predictive intelligence and the risks associated with acting on predictions too hastily. It is important to note that in some sectors and industries, regulations require provable human involvement, which can also explain this phenomenon.

Several factors that contribute to the limited use of predictive CTI among questionnaire respondents and interviewees were identified. These include a lack of trust in predictive outputs, resource constraints, integration challenges, and organisational constraints. The lack of trust towards predictive intelligence comes from concerns of accuracy and skepticism of its predictive nature, as well as fears of false positives and of potential collateral harm from reacting to it. Constraints in resources like time, budget, and expertise can arise from a lack of understanding of predictive CTI and its use-cases but mainly arise from insufficient baseline maturity. Integration challenges faced by organisations include high cognitive and operational burden of predictive CTI

implementation, poor compatibility with existing tools and workflows, and a lack of formalisation or standardisation of intelligence requirements and source reliability. Organisational constraints, which include regulatory constraints, were found to occasionally limit automation of predictive CTI, even in environments where integration is technically possible.

Overall, it was found that predictive CTI currently serves as a supporting capability, primarily used in the guidance of some strategic and operational decisions. Direct driving of automated response actions based on predictive intelligence was rare. This shows a gap between the theoretical potential of predictive CTI and how it is currently implemented in practice.

5.3 Recommendations for Improving Production and Consumption of Predictive Cyber Threat Intelligence

Building on the findings and previous synthesis, this subchapter will provide a set of recommendations for improving the actionability and utilisation of predictive CTI. These recommendations address both producers and consumers of CTI, reflecting the different approaches present in addressing the identified challenges.

5.3.1 Improving Predictive Cyber Threat Intelligence Actionability from the Producer Side

Common barriers identified in the questionnaire and interviews by consumers of predictive CTI were a lack of sufficient resources and cybersecurity maturity, and challenges interpreting and integrating predictive intelligence. It is therefore suggested that measures taken to improve the accessibility and ease of use of predictive outputs will lead to increased adoption and utilisation of predictive CTI, especially in less mature organisations. One way of improving this is through making predictive outputs more prescriptive, including explicit, ideally context-specific recommendations for controls and mitigation actions. Another means to achieving better accessibility is through providing more decision support, combining mitigation measures with a

recommended time window to act and impact and risk assessments. Increasing contextual relevance of predictive intelligence, especially to clients' sector, location, tools, and supply chain, can further improve accessibility of predictive outputs, and is something that many organisations demand and look for. In addition, better compatibility between predictive CTI outputs and existing security tools and workflows is something that many consumers are looking for, suggesting that improving format and delivery of predictions, not just content, can yield benefits. It is also important to consider catering towards organisations of varying maturity levels through reducing technical complexity and providing tiered reporting, in order to increase accessibility and usability, especially for less mature organisations.

Another common theme among the consumers of predictive CTI who responded to the questionnaire was that they required higher trust and confidence thresholds in order to act upon predictive data. It is therefore especially important for producers of predictive CTI to present themselves and their outputs as reliable and trustworthy. Transparency and explainability of intelligence outputs, especially if they involved machine generated outputs, were identified as something consumers value. Establishing reliability and trust is important but may be difficult, as the vast majority of respondents got their predictive intelligence through bilateral agreements with partners, and many rarely attempted to consider other sources. The relatively high scrutiny of predictive outputs also suggests that focus should be placed on lowering signal to noise ratios and reducing false positive rates in order to increase public trust in predictive CTI and the likelihood of it being operationalised.

There are many features that can increase the perceived value of predictive outputs for consumers. In interviews it was revealed that access to non-public data, including dark web and other closed sources, was viewed as valuable. So were insights into publicly available or otherwise circulating exploit code. From interviews with people involved in the production of predictive CTI, it was revealed that a lot of unique insight can come from malware, exploit, and killchain analysis. It was also revealed that predictive reporting benefits from both reactive and proactive methods.

Another insight that may benefit producers of predictive CTI is a shift in demand towards more strategic level outputs, especially among higher maturity organisations.

These kinds of organisations are beginning to demand reports more focused on trends concerning emerging threats and technologies, as well as how (certain) threat actors think.

5.3.2 Improving Predictive Cyber Threat Intelligence Actionability from the Consumer Side

A key insight was the vast difference in predictive CTI adoption and utilisation across organisations of differing cybersecurity maturity levels. Therefore, recommendations for improving predictive CTI utilisation from the consumer side are split into two types: recommendations for developing or enabling predictive CTI use, and recommendations for improving existing predictive CTI use.

In order to effectively ingest and act upon predictive intelligence, capability development must be properly sequenced. It was determined that predictive CTI utilisation is dependent on other foundational security capabilities, so organisations must first achieve some level of maturity in detection, protection, and awareness capabilities before predictive intelligence can be incorporated into workflows for positive effect. Proper processing of predictive CTI also requires resources, especially in the form of human capital, as many workflows involving predictive intelligence are human-centric. It is important to consider if undertaking predictive CTI implementation is feasible. Many respondents utilised predictive intelligence in unstructured, informal, processes, suggesting that value can still be gleaned from it without full formal commitment. However, interviews also revealed that defining clear priority intelligence requirements and formalising processes and playbooks are paramount for organisations with higher levels of cybersecurity maturity.

When predictive CTI ingestion capability already exists, recommendations to improve its actionability include embedding it into SOC and other operational workflows, as interviews revealed that standalone CTI functions are less effective. Another avenue for improving predictive intelligence use is the prioritisation of efficient internal and external communication channels to ensure effective coordinated response between different teams or partnering organisations, at least to the extent of enabling direct

messaging and calling capability. Findings also showed that predictive CTI is most commonly effectively utilised when it is being applied to proactive threat hunting and vulnerability management, and is less effective in incident response situations. The final recommendation is to consider the actionability criteria discussed earlier in the chapter when analysing or otherwise evaluating incoming predictive CTI data.

5.4 Synthesis of Outcomes

Bringing the outcomes together, the central finding of this thesis is that the actionability of predictive CTI emerges from the alignment between intelligence characteristics and organisational capabilities, rather than from technical sophistication alone. The first two outcomes of the thesis, the evaluation criteria and analysis of the use of predictive cyber intelligence in practice, contribute to research by extending predictive CTI literature beyond its technical focus. Introducing the idea that improving actionability for predictive intelligence comes from both producers and consumers. The recommendations, the third outcome of the thesis, provides concrete practical groundwork for practitioners involved in, or looking to get involved in, predictive CTI, whether as a producer or consumer.

6 Evaluation

When considering the findings of this thesis, it is important to be aware of the limitations and assumptions present in the research design, methodology, and results. Regarding the review of literature, one major limitation was the reliance on publicly available data. Proprietary or classified intelligence could differ significantly in quality and usage, and many organisations do not, or are otherwise unwilling to, describe their cybersecurity decision processes and workflows. As a result, the findings may not fully represent the practices of organisations that rely heavily on private or internal intelligence sources.

The questionnaire and interviews were voluntary, leading to self-selection bias as factors like interest in CTI or increased cybersecurity awareness could make individuals more likely to respond. This means that respondents and their beliefs may not be fully representative of the cybersecurity field. The questionnaire and interviews were also prone to self reporting bias, as participants may have been inclined to report a higher cybersecurity maturity than reality in an effort to appear more socially acceptable.

Additionally, when looking at the survey, the questionnaire and interviews both had a relatively small sample size. This limits the robustness of the findings and raises doubts about the generalisability of the data to the broader cybersecurity community. Other factors that may limit the generalisability of the data include the fact that the vast majority of respondents were based in Estonia, and a large portion of respondents were in the public sector.

Despite being performed in Estonia, the questionnaire was only available in English, which could have led to the exclusion of some relevant respondents. This choice was made as the cybersecurity community in Estonia is predominantly English speaking,

meaning any potential respondents that would be excluded by the choice of language of the questionnaire would more than likely not be part of the target demographic of the questionnaire.

Other limitations of the questionnaire arose from compromises made to encourage responses. One such limitation was that the questionnaire answers were anonymous and questions were not mandatory. These choices encouraged truthful responses to the questions, but made it harder to determine the quality and validity of each response. Although responses were audited with obviously non-serious ones being removed, the risk that some responses were not authentic can not be ignored. The questionnaire was also intentionally kept short (with an estimated time to complete of around 5 minutes) to encourage responses, which led to some questions being removed or simplified. To mitigate the impact of this, the questions went through an iterative process where questions were pruned to maximise relevance to the research goals and minimise redundancy.

A limitation of the interviews was that they were not recorded or transcribed verbatim. Instead, notes were taken during and immediately after each interview. This approach introduces the possibility of inaccuracies, omissions, or unintended interpretation in the captured data. While efforts, like requests for confirmation or repetition of answers, were made, this approach's effect on the reliability of the findings must be acknowledged.

However, a strength of the study is the diversity in traditional CTI and predictive CTI adoption levels among respondents. The questionnaire and interviews included organisations with varying degrees of maturity and experience with CTI, contributing to a more holistic understanding of the current state of CTI. Another strength to note is the independent, outcome-neutral nature of the study. Many similar studies are often done in affiliation with certain organisations with the purpose of justifying or selling a specific product or solution, introducing a bias from the vested interest. This study only approaches the topic for the purpose of increasing knowledge which lowers the risk of intentional bias and increases the credibility of the findings. There was also corroboration between the findings of each element of the survey and existing research

into predictive CTI and the actionability of traditional CTI, which suggests that the results of this thesis possess a level of generalisability and veracity.

7 Future Directions

This thesis and its findings present new questions which could be the focus of future research. One possible direction of research could be the development of quantitative metrics and data for the concepts discussed in this paper. Ways to quantify things like the accuracy or return on investment of predictive CTI can further extend the findings of this paper into more practical solutions for practitioners. Other directions of research that can provide more practical solutions for practitioners would be investigations into the possible ways to increase integration of predictive CTI with existing security tools and platforms, and ways to safely introduce automation into the process. Effort into the formalisation and standardisation of predictive CTI production and ingestion is likely to be valuable as well.

One insight from the findings of this paper was the existence of a so-called “goldilocks window” where predictive CTI isn’t too early or too late and actionability is maximised. A possible research direction involving this would be attempting to define this window and find the optimal timing of predictive intelligence, and how this timing may change depending on the type of intelligence and other factors. Future research can also involve replicating this study in specific regions and sectors, to determine the generalisability of the findings of this thesis. The trends and challenges reported in this thesis’s survey are quite similar to historical reports of (traditional) CTI from a couple years ago. Therefore, investigation into the organisational adoption of predictive CTI and how closely it mirrors the history of organisational adoption of traditional CTI may also be valuable.

8 Conclusion

The goal of this thesis was to investigate the actionability of predictive CTI and ways to improve it, while also analysing how predictive CTI is currently used in practice. The main objectives were to develop a set of criteria for the evaluation of predictive cyber intelligence, analyse how predictive CTI is currently used in practice, and create a set of recommendations for improving the actionability of future predictive cyber intelligence. Accomplishing these objectives would enable movement towards a more formal, evidence-backed evaluation of predictive CTI, as well as aid organisations of varying levels of cybersecurity maturity in developing their predictive intelligence capabilities or extracting more value from their existing predictive intelligence.

For the purpose of achieving these objectives, three main research questions were asked. The first research question was about the applicability of actionability criteria for traditional CTI to predictive CTI. It was found that the commonly used actionability criteria devised by ENISA [7] (relevance, timeliness, accuracy, completeness, ingestibility) do generally apply when the intelligence is predictive, with relevance and accuracy being the criteria that consumers were most sensitive to. However, some changes need to be made to the criteria in order to account for the unique challenges of predictive CTI and its implementation. These insights informed the development of a set of evaluation criteria tailored specifically to predictive intelligence which extended existing models by accounting for predictive uncertainty and the temporal dynamic.

The second research question focused on how predictive CTI is used in practice. This question was broken down into smaller subquestions which each looked at different aspects such as the extent of predictive CTI use, its weight in decision making, and the tools, platforms, and processes involved with it. It was found that predictive CTI is currently underused compared to its traditional, non-predictive counterpart. In addition,

it is most commonly seen in areas such as risk management, vulnerability prioritisation, and threat hunting, though its influence is often limited. It was also found that predictive intelligence is not a priority for organisations with lower cybersecurity maturities. Across all levels of maturity, however, predictive CTI workflows were found to be predominantly human-centric, with little reliance on automation due to concerns about trust, accuracy, and risks associated with acting hastily on predictions. The key finding for this research question is therefore that there is a gap between the theoretical potential and the current implementation of predictive CTI, with the gap growing more apparent at lower levels of cybersecurity maturity.

The third research question asked how predictive CTI can be better utilised. The scope of this question includes how actionability of predictive intelligence can be taken to the next level, and how its producers and consumers can get it there. Its subquestions therefore focused on the barriers to utilising predictive CTI, the ways it can be improved, and the benefits and drawbacks to the currently used processes and tools used by organisations when dealing with it. It was identified that the barriers to better utilisation are drastically different between organisations of high and low cybersecurity maturity. The most commonly identified barriers were insufficient contextual relevance, integration challenges with existing tools and workflows, and constraints related to resources and cybersecurity maturity. To address these challenges, a set of recommendations for both producers and consumers of predictive CTI was developed. For producers, it was identified that improving actionability involves making predictive intelligence more accessible and easier to use, especially for any lower maturity clients. For consumers, it was identified that better utilisation of predictive CTI involves properly sequencing cyber capability development and the formalisation of intelligence requirements and analytical expectations. Overall, it was found that the actionability of predictive intelligence relies heavily on the alignment of intelligence characteristics with organisational capabilities, not just the improvement of the content of predictions and the technology behind them.

Reference List

[1] P. V. Falade, “Decoding the threat landscape : CHATGPT, FRAUDGPT, and wormgpt in social engineering attacks,” *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, vol. 9, no. 5, pp. 185–198, Oct. 2023. doi:10.32628/cseit2390533

[2] M. Aksela, S. Marchal, A. Patel, L. Rosenstedt, and WithSecure, 31/2022 ed. National Cyber Security Centre Finland, 2022.

[3] N. Kaloudi and J. Li, “The AI-based Cyber Threat Landscape,” *ACM Computing Surveys*, vol. 53, no. 1, pp. 1–34, Feb. 2020. doi:10.1145/3372823

[4] A. Dimitriadis et al., “EVACTI: Evaluating the Actionability of Cyber Threat Intelligence,” International Journal of Information Security, vol. 24, no. 3, May 2025. doi:10.1007/s10207-025-01033-z

[5] R. Brown and R. M. Lee, “The Evolution of Cyber Threat Intelligence (CTI): 2019 SANS CTI survey,” SANS Institute, https://www.sans.org/white-papers/38790 (accessed May 18, 2026).

[6] M. Barrios-González, J. M. Aguiar-Pérez, M. Á. Pérez-Juárez, and E. Castañeda-de-Benito, “Redefining cyber threat intelligence with artificial intelligence: From Data Processing to predictive insights and human–AI collaboration,” Applied Sciences, vol. 16, no. 3, p. 1668, Feb. 2026. doi:10.3390/app16031668

[7] P. Pawliński et al., Actionable Information for Security Incident Response. ENISA, 2014.

[8] S. Saeed, S. A. Suayyid, M. S. Al-Ghamdi, H. Al-Muhaisen, and A. M. Almuhaideb, “A systematic literature review on Cyber Threat Intelligence for Organizational Cybersecurity Resilience,” *Sensors*, vol. 23, no. 16, p. 7273, Aug. 2023. doi:10.3390/s23167273

[9] A. Sfakianakis, C. Douligeris, L. Marinos, M. Lourenço, and O. Raghimi, “ENISA Threat Landscape Report 2018: 15 Top Cyberthreats and Trends.” European Union Agency For Network and Information Security, Jan. 28, 2019

[10] R. Jabir, J. Le, and C. Nguyen, “Phishing attacks in the age of Generative Artificial Intelligence: A systematic review of human factors,” AI, vol. 6, no. 8, p. 174, Jul. 2025. doi:10.3390/ai6080174

[11] M. Schmitt and I. Flechais, “Digital deception: Generative artificial intelligence in Social Engineering and phishing,” Artificial Intelligence Review, vol. 57, no. 12, Oct. 2024. doi:10.1007/s10462-024-10973-2

[12] F. Heiding et al., “Evaluating large language models’ ability to automate spear phishing,” Expert Systems with Applications, vol. 314, p. 131546, Jun. 2026. doi:10.1016/j.eswa.2026.131546

[13] M. Brundage et al., “The malicious use of artificial intelligence: Forecasting, prevention, and mitigation,” Apollo Repository, Feb. 2018. doi:https://doi.org/10.17863/CAM.22520

[14] T. King, N. Aggarwal, M. Taddeo, and L. Floridi, “Artificial Intelligence Crime: An interdisciplinary analysis of foreseeable threats and solutions,” SSRN Electronic Journal, vol. 26, no. 2020, pp. 89–120, Feb. 2019. doi:10.2139/ssrn.3183238

[15] C. T. Thanh and I. Zelinka, “A survey on Artificial Intelligence in malware as next-generation threats,” MENDEL, vol. 25, no. 2, pp. 27–34, Dec. 2019. doi:10.13164/mendel.2019.2.027

[16] Y. Tong, H. Liang, H. Ma, S. Zhang, and X. Yang, “A survey on reinforcement learning-driven adversarial sample generation for PE malware,” Electronics, vol. 14, no. 12, p. 2422, Jun. 2025. doi:10.3390/electronics14122422

[17] L. Fritsch, A. Jaber, and A. Yazidi, “An overview of artificial intelligence used in malware,” Communications in Computer and Information Science, pp. 41–51, 2022. doi:10.1007/978-3-031-17030-0_4

[18] A. Hernandez-Suarez et al., “REINFORSEC: An Automatic Generator of synthetic malware samples and denial-of-service attacks through reinforcement learning,” Sensors, vol. 23, no. 3, p. 1231, Jan. 2023. doi:10.3390/s23031231

[19] O. Kubovič, P. Košinár, and J. Jánošík, “Can Artificial Intelligence Power Future malware?,” ESET, https://web-assets.esetstatic.com/wls/en/papers/white-papers/Can_AI_Power_Future_Malware.pdf (accessed May 17, 2026).

[20] A. Lohn, “Anticipating AI’s impact on the cyber offense-defense balance,” Center for Security and Emerging Technology, https://cset.georgetown.edu/publication/anticipating-ais-impact-on-the-cyber-offense-defense-balance/ (accessed May 17, 2026).

[21] G. WAIZEL, “Bridging the AI divide: The evolving arms race between AI-driven cyber attacks and AI-powered cybersecurity defenses,” International Conference on Machine Intelligence & Security for Smart Cities (TRUST) Proceedings, https://www.scrd.eu/index.php/trust/article/view/554 (accessed May 18, 2026).

[22] K. Zdrojewski, “AI - Powered Cyber attacks: A Comprehensive Review and Analysis of Emerging Threats,” Advances in IT and Electrical Engineering, vol. 50, no. 2025, pp. 55–71. Accessed: May. 17, 2026. [Online]. Available: https://journals.prz.edu.pl/aitee/article/view/2031/1507

[23] N. Sun et al., “Cyber Threat Intelligence Mining for Proactive Cybersecurity defense: A survey and new perspectives,” IEEE Communications Surveys & Tutorials, vol. 25, no. 3, pp. 1748–1774, 2023. doi:10.1109/comst.2023.3273282

[24] L. Hassan and R. Al-Mazrouei, “Evolving paradigms and future trajectories in Cyber Threat Intelligence,” International Journal of Cyber Threat Intelligence and Secure Networking, vol. 2, no. 06, pp. 1–7, Jun. 2025. doi:10.55640/ijctisn-v02i06-01

[25] L. H. Bennett, “Proactive Cyber Threat Intelligence: Integrative Frameworks for Detection, Attribution, and Predictive Defense,” The American Journal of Engineering and Technology, vol. 7, no. 7, pp. 218–222, Jul. 2025. Accessed: 2026. [Online]. Available: https://www.theamericanjournals.com/index.php/tajet/article/view/7101

[26] W. Tounsi and H. Rais, A survey on technical threat intelligence in the age of sophisticated cyber attacks, https://doi.org/10.1016/j.cose.2017.09.001 (accessed May 18, 2026).

[27] G. Sakellariou, P. Fouliras, I. Mavridis, and P. Sarigiannidis, “A reference model for Cyber Threat Intelligence (CTI) systems,” Electronics, vol. 11, no. 9, p. 1401, Apr. 2022. doi:10.3390/electronics11091401

[28] C. S. Johnson, M. L. Badger, D. A. Waltermire, J. Snyder, and C. Skorupka, “Guide to cyber threat information sharing,” NIST Special Publication, Oct. 2016. doi:10.6028/nist.sp.800-150

[29] S. Ainslie, D. Thompson, S. Maynard, and A. Ahmad, “Cyber-threat Intelligence for Security Decision-making: A review and research agenda for practice,” Computers & Security, vol. 132, p. 103352, Sep. 2023. doi:10.1016/j.cose.2023.103352

[30] M. Sahrom Abu, S. Rahayu Selamat, A. Ariffin, and R. Yusof, “Cyber threat intelligence – issue and challenges,” Indonesian Journal of Electrical Engineering and Computer Science, vol. 10, no. 1, pp. 371–379, Apr. 2018. doi:10.11591/ijeecs.v10.i1.pp371-379

[31] D. Schlette, F. Böhm, M. Caselli, and G. Pernul, “Measuring and visualizing cyber threat intelligence quality,” International Journal of Information Security, vol. 20, no. 1, pp. 21–38, Mar. 2020. doi:10.1007/s10207-020-00490-y

[32] R. Brown and K. Nickels, “Sans 2023 CTI survey: Keeping up with a changing threat landscape,” SANS Institute, https://www.sans.org/white-papers/2023-cti-survey-keeping-up-changing-threat-landscape (accessed May 18, 2026).

[33] R. Brown and A. Sfakianakis, “Sans 2024 CTI survey: Managing the evolving threat landscape,” SANS Institute, https://www.sans.org/white-papers/2024-cti-survey-managing-evolving-threat-landscape (accessed May 18, 2026).

[34] R. Brown and A. Sfakianakis, “Sans 2025 CTI survey: Navigating uncertainty in today’s threat landscape,” SANS Institute, https://www.sans.org/white-papers/2025-cti-survey-webcast-forum-navigating-uncertainty-todays-threat-landscape (accessed May 18, 2026).

[35] R. M. Lee, “2020 sans cyber threat intelligence (CTI) survey,” SANS Institute, https://www.sans.org/white-papers/39395 (accessed May 18, 2026).

[36] R. Brown and R. M. Lee, “2021 SANS Cyber Threat Intelligence (CTI) survey,” SANS Institute, https://www.sans.org/white-papers/40080 (accessed May 18, 2026).

[37] R. Brown and P. Stirparo, “SANS 2022 cyber threat intelligence survey,” SANS Institute, https://www.sans.org/white-papers/sans-2022-cyber-threat-intelligence-survey (accessed May 18, 2026).

[38] P. Santos, R. Abreu, M. J. Reis, C. Serôdio, and F. Branco, “A systematic review of Cyber Threat Intelligence: The effectiveness of technologies, strategies, and collaborations in Combating modern threats,” Sensors, vol. 25, no. 14, p. 4272, Jul. 2025. doi:10.3390/s25144272

[39] M. Wilson, P. Peace, and J. Owen, Predictive Analytics for Cyber Threat Intelligence, Dec. 2024. Accessed: 2026. [Online]. Available: https://www.researchgate.net/publication/387243092_Predictive_Analytics_for_Cyber_Threat_Intelligence

[40] V. V. Vigesna and A. Adepu, “Leveraging Artificial Intelligence for Predictive Cyber Threat Intelligence ,” International Journal of Creative Research In Computer Technology and Design, vol. 6, no. 6, pp. 1–19, 2024. Accessed: 2026. [Online]. Available: https://jrctd.in/index.php/IJRCTD/article/view/64

[41] C. Gilbert and M. A. Gilbert, “Artificial Intelligence (AI) and Machine Learning (ML) for Predictive Cyber Threat Intelligence (CTI),” International Journal of Research Publication and Reviews , vol. 6, no. 3, pp. 584–617, Mar. 2025. Accessed: 2026. [Online]. Available: https://www.researchgate.net/publication/389554836_Artificial_Intelligence_AI_and_Machine_Learning_ML_for_Predictive_Cyber_Threat_Intelligence_CTI

[42] T. D. Wagner, K. Mahbub, E. Palomar, and A. E. Abdallah, “Cyber threat intelligence sharing: Survey and research directions,” Computers & Security, vol. 87, p. 101589, Nov. 2019. doi:10.1016/j.cose.2019.101589

[43] D. Fujs, A. Mihelič, and S. L. Vrhovec, “The power of interpretation,” Proceedings of the 14th International Conference on Availability, Reliability and Security, pp. 1–10, Aug. 2019. doi:10.1145/3339252.3341479

[44] “Korporatsioon Rotalia on Levinud üle Kogu Maailma,” Korporatsioon Rotalia, https://rotalia.ee/ (accessed Apr. 16, 2026).

Appendix 1 - Non-exclusive licence for reproduction and publication of a graduation thesis[13]

I, Uku Joosep Palu

1 grant Tallinn University of Technology free licence (non-exclusive licence) for my thesis "An analysis of the actionability of predictive cyber threat intelligence in the era of AI" , supervised by Margus Ernits

1.1 to be reproduced for the purposes of preservation and electronic publication of the graduation thesis, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright;

1.2 to be published via the web of Tallinn University of Technology, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright.

2 I am aware that the author also retains the rights specified in clause 1 of the non-exclusive licence.

3 I confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights, the rights arising from the Personal Data Protection Act or rights arising from other legislation.

[13] The non-exclusive licence is not valid during the validity of access restriction indicated in the student's application for restriction on access to the graduation thesis that has been signed by the school's dean, except in case of the university's right to reproduce the thesis for preservation purposes only. If a graduation thesis is based on the joint creative activity of two or more persons and the co-author(s) has/have not granted, by the set deadline, the student defending his/her graduation thesis consent to reproduce and publish the graduation thesis in compliance with clauses 1.1 and 1.2 of the non-exclusive licence, the non-exclusive license shall not be valid for the period.

Appendix 2 - Questionnaire

Consent Page

This survey is part of a master's thesis titled: “An Analysis on the Actionability of Predictive Cyber Threat Intelligence in the Era of AI.” The purpose of this study is to understand how predictive cyber threat intelligence is used in practice and whether it supports cybersecurity decision-making. The survey focuses on how cybersecurity professionals use predictive threat intelligence and how actionable they consider it to be. Participants should be individuals involved in cybersecurity-related work or decision-making.

Participation in this study is voluntary. You may skip any question you do not wish to answer. The survey takes approximately 5-7 minutes to complete. The information gathered will be used solely for academic purposes. Responses will be analyzed in aggregate form and individual participants and organisations will not be identifiable in the thesis or any publications. The data collected will be retained for research purposes and stored for a maximum of 2 years after which it will be securely deleted.

Contact Information

The following information will only be used for the purpose of sharing with you the insights and conclusions of this study.

1. Please enter your email address

Background Information

Individual participants or organisations will not be identifiable in the thesis or any publications.

2. What is your role / job title?

3. What is the name of your organisation?

4. What industry sector does your organisation primarily reside under?

- ☐ Finance and Insurance
- ☐ Government and Public Sector
- ☐ Healthcare
- ☐ Education
- ☐ Cybersecurity Service Provider
- ☐ Information Technologies or Telecommunication
- ☐ Manufacturing and Industrial
- ☐ Other

5. What is the size of your organisation's dedicated security team?

- ☐ 0
- ☐ 1
- ☐ 2-4
- ☐ 5-9
- ☐ 10-20
- ☐ 21-50
- ☐ 50+
- ☐ Other

Cyber Threat Intelligence (CTI) Usage

6. Does your organisation use CTI?

- ☐ Yes
- ☐ No
- ☐ Other

7. How much of the CTI your organisation uses is predictive?

Predictive CTI refers to threat intelligence that attempts to forecast the likelihood, timing, or characteristics of future cyber threats in order to inform proactive security decisions.
Examples of predictive CTI include:
- Predictions of future attack trends targeting your sector.
- AI/ML based threat and risk forecasts.
- Threat actor activity forecasting.

Examples of CTI that are NOT predictive include:
- IOC feeds.
- Incident reports.
- Alerts about vulnerabilities that do not forecast exploitation.

- ☐ None of the CTI we use is predictive
- ☐
- ☐
- ☐ We use a balanced mix of predictive and non-predictive CTI
- ☐
- ☐
- ☐ All of the CTI we use is predictive

Actionability of Predictive CTI

The section title can be defined as the extent predictive CTI supports concrete cybersecurity decisions or operational actions within an organisation.

8. How is predictive CTI used in your organisation? Select all that apply.

- ☐ Threat hunting
- ☐ Incident response
- ☐ Vulnerability and security control prioritisation
- ☐ Detection engineering or rule tuning
- ☐ User education
- ☐ Executive education and awareness (board of directors, C-suite)
- ☐ Risk management
- ☐ Compliance
- ☐ Adversary emulation / purple teaming
- ☐ Budget, staffing, or tooling investment decisions
- ☐ Other

9. Which decisions are influenced by predictive CTI data? Select all that apply.

- ☐ Prioritising vulnerability remediation
- ☐ Adjusting detection rules / monitoring configurations
- ☐ Initiating proactive threat hunting
- ☐ Escalating or prioritising incident responses
- ☐ Updating risk assessments
- ☐ Allocating budgets, staff, or tooling investments
- ☐ Designing red team or adversary emulation exercises
- ☐ Updating threat models or incident response playbooks
- ☐ Informing executive/board-level decisions
- ☐ Other

10. How influential is predictive CTI in the decision making process? (1 = Not at all, 7 = Extremely)

11. What factors make predictive CTI most actionable? Rank in order of importance.

- ☐ Relevance to organisation and context
- ☐ Timeliness
- ☐ Accuracy and confidence level
- ☐ Completeness
- ☐ Ingestibility (ease of use)
- ☐ Specificity

12. What factors make predictive CTI most actionable? Please include any factors you think are important but were not mentioned in the previous question here.

13. What are the main barriers to using predictive CTI? Rank in order of importance.

- ☐ Lack of relevant predictive CTI
- ☐ Low confidence or perceived accuracy
- ☐ Not timely
- ☐ Incomplete
- ☐ Poor integration with tools and workflows
- ☐ Too vague
- ☐ Lack of resources

14. What are the main barriers to using predictive CTI? Please include and barriers you think are important but were not mentioned in the previous question here.

15. What would make predictive CTI more useful to your organisation? Describe briefly.

Future Correspondence

16. Are you open to have a follow up meeting/interview about this topic in the future? If yes, please also provide some contact information.