Technological Forecasting & Social Change 210 (2025) 123836

Contents lists available at ScienceDirect

Technological Forecasting & Social Change

journal homepage: www.elsevier.com/locate/techfore

# Forecasting Cyber Threats and Pertinent Mitigation Technologies

Zaid Almahmoud [a], Paul D. Yoo [a,*], Ernesto Damiani [b,d], Kim-Kwang Raymond Choo [c], Chan Yeob Yeun [d]

[a] School of Computing and Mathematical Sciences, University of London, Birkbeck College, London, WC1E 7HX, UK
[b] Department of Computer Science, University of Milan, Milan, 20122, Italy
[c] Department of Information Systems and Cyber Security, University of Texas at San Antonio, San Antonio, TX, 78249, USA
[d] Department of Electrical Engineering and Computer Science, Center for Cyber-Physical Systems (C2PS), Khalifa University, Abu Dhabi, 127788, United Arab Emirates

ARTICLE INFO

Dataset link: https://github.com/zaidalmahmoud/Cyber-trend-forecasting

*Keywords:*
Cyber threat trend forecasting
Mitigation technology trend forecasting
Alleviation technology trend forecasting
Proactive approach
Big data analytics
Graph machine learning
Technology cycle

ABSTRACT

Geopolitical instability is exacerbating the risk of catastrophic cyber-attacks striking where defences are weak. Nevertheless, cyber-attack trend forecasting predominantly relies on human expertise, which is susceptible to subjectivity and potential bias. As a solution, we have recently presented a novel study that harnesses machine learning for long-term cyber-attack forecasting. Building upon this groundwork, our research advances to the next level, by predicting the disparity between cyber-attack trends and the trend of the relevant alleviation technologies. The proposed approach applies key constructs of Protection Motivation Theory while introducing a proactive version of the theory. Our predictive analysis aims to offer strategic insights for the decision of investment in cyber security technologies. It also provides a sound foundation for the strategic decisions of national defence agencies. To achieve this objective, we have expanded our dataset, which now encompasses records spanning 42 distinct cyber-attack types and various related features, alongside data concerning the trends of 98 pertinent technologies, dating back to 2011. The dataset features were meticulously curated from diverse sources, including news articles, blogs, government advisories, as well as from platforms such as Elsevier, Twitter, and Python APIs. With our comprehensive dataset in place, we construct a graph that elucidates the intricate interplay between cyber threats and the development of pertinent alleviation technologies. To forecast the graph, we introduce a novel Bayesian adaptation of a recently proposed graph neural network model, which effectively captures and predicts these trends. We further demonstrate the efficacy of our proposed features in this context. Furthermore, our study extends its horizon by generating future data projections for the next three years, encompassing forecasts for the evolving graph, including predictions of the gap between cyber-attack trends and the trend of the associated technologies. As a consequential outcome of our forecasting efforts, we introduce the concept of "alleviation technologies cycle", delineating the key phases in the life cycle of 98 technologies. These findings serve as a foundational resource, offering valuable guidance for future investment and strategic defence decisions within the realm of cyber security related technologies.

## 1. Introduction

Over the past decade, there have been numerous waves of cyber-attacks with varying extent of damages to governments, organisations, and enterprises (Ghafur et al., 2019; Linkov et al., 2023). While reactive cyber-defences (Gaurav et al., 2022) may not prevent these incidents from happening, they can possibly reduce the resulting damage. By leveraging the availability of data, a proactive defence strategy can be implemented to address potential cyber threats before they escalate into actual incidents (Goel, 2011). Similar proactive approaches also have proven effective in mitigating non-cyber threats like terrorism and military attacks. For instance, with the aid of advanced software programs, it is now possible to assess the intentions, potential damages, attack methods, and alternative options associated with terrorist attacks (Kebir et al., 2022). We posit that the same proactive approach can be applied to cyber-attacks.

The primary motivation of this work lies in the critical need for accurate and objective forecasting of cyber threats over the foreseeable future. In today's digital landscape, cyber-attacks are becoming

* Corresponding author.
*E-mail addresses:* p.yoo@bbk.ac.uk (P.D. Yoo), ernesto.damiani@unimi.it (E. Damiani), raymond.choo@fulbrightmail.org (K.-K.R. Choo).

https://doi.org/10.1016/j.techfore.2024.123836
Received 20 November 2023; Received in revised form 29 May 2024; Accepted 14 October 2024
Available online 30 October 2024

increasingly sophisticated and prevalent. Hence, the ability to anticipate and prepare for future threats ahead of time is paramount for effective cyber security management. The long-term forecast of cyber threats provides cyber security agencies sufficient time to assess existing defence measures and identify areas where preventive solutions can be developed proactively. However, existing methods for long-term prediction of cyber threats often rely on subjective assessments made by human security experts (Stephens, 2008; Adamov and Carlsson, 2017). This human-centric approach is inherently prone to bias influenced by individual perspectives and lacks the scalability required to handle the growing complexity of cyber threats. Moreover, relying solely on human expertise offers no guarantee of accuracy or reliability, as there is no "safety in numbers" when it comes to expert judgement (Shoufan and Damiani, 2017). Consequently, there is a pressing need for a fully automated procedure that can provide scientifically objective predictions based on quantitative metrics (Cha and Hao, 2022). Such an approach enhances the credibility and reliability of threat forecasts.

In our previous work (Almahmoud et al., 2023), we conducted the first study on the long-term cyber-attack trend forecasting, employing a fully automated approach. Nonetheless, achieving an informed decision-making process in the context of technology investment for the mitigation of these forthcoming threats necessitates a more forward-looking perspective (Kuwahara, 1999). This perspective encompasses the discernment of the divergence between the trajectory of cyber threats and Pertinent Alleviation Technologies (PATs). For instance, being able to predict a substantial future disparity between the trajectory of an attack and the corresponding technical solution empowers us to make judicious choices regarding our investments. Consequently, we can prioritise our defence strategies based on these predictive insights.

Our work advances the Protection Motivation Theory (PMT) (Rogers, 1975), by addressing specific gaps in the literature related to its application in cyber security forecasting frameworks. While PMT has been extensively applied in various domains to understand individuals' responses to threats and the adoption of protective measures (Tsai et al., 2016; Ruthig, 2016; Alshammari et al., 2024), its integration into proactive cyber security methodologies has been limited. This study fills this gap by directly incorporating PMT principles into the development of a proactive framework for cyber threat prediction and mitigation.

One of the core components of PMT is threat appraisal, which involves assessing the severity and vulnerability associated with a threat (Rogers, 1975). In the context of cyber security, our study employs a data-driven approach to analyse and forecast trends in cyber threats, thereby providing a quantitative and proactive assessment of the evolving landscape of cyber-attacks. Also, by leveraging big data analytics, we move beyond the traditional reliance on subjective expert assessments. This unbiased and systematic threat appraisal aligns with PMT's emphasis on understanding perceived threats and enhances our ability to anticipate future cyber threats more accurately.

Another essential aspect of PMT is the appraisal of coping strategies (Rogers, 1975), which evaluates the efficacy of available measures to mitigate perceived threats. Our study extends this principle by forecasting the trends of PATs and assessing their disparity with relevant cyber threats. By analysing and predicting the development and effectiveness of these technologies, we provide insights into how well current and emerging solutions can address future cyber threats. This proactive appraisal of coping strategies ensures that organisations can strategically invest in and develop technologies that are most likely to be effective, thus aligning with PMT's focus on motivating protective behaviours based on the perceived efficacy of coping mechanisms.

Moreover, our work bridges the gap between psychological theories and technical methodologies in cyber security. By integrating PMT principles into a machine-based forecasting framework, we demonstrate the applicability of PMT beyond its traditional domains, such as health psychology, into the domain of cyber security. This interdisciplinary approach enriches the theoretical framework of PMT and highlights its practical relevance in guiding proactive cyber defence strategies. Our study is guided by the principle that understanding psychological motivations and perceptions can significantly enhance the effectiveness of technical solutions in cyber security, thereby contributing to the broader application and advancement of PMT.

Overall, we advance PMT by proposing a proactive version that extends the traditional PMT's focus on current threats and coping strategies to include future forecasting. By evaluating anticipated cyber threats and relevant alleviation technologies, our approach enhances individuals' and organisations' confidence and preparedness. This proactive stance fosters sustained protection motivation and facilitates strategic investment and planning, ensuring long-term resilience against cyber threats. This contribution not only enriches the theoretical framework of PMT but also provides practical implications for developing more effective and forward-looking protective behaviours.

Out of 42 attack types analysed in our previous work, we classified 26 as emerging or rapidly increasing (Almahmoud et al., 2023). These classifications were made based on the findings of our prior research, indicating the urgency and significance of these threats in the cyber security landscape. In this study, we have selected these 26 threats for detailed examination due to their heightened importance and potential impact on cyber security. Emerging and rapidly increasing threats pose the highest risk as they have the propensity to escalate quickly and cause significant harm to individuals, organisations, and critical infrastructure. Compared to other categories, such as declining threats, these emerging and rapidly increasing threats are more critical because they represent ongoing challenges that demand immediate attention. By focussing our analysis on these critical threats, we aim to provide actionable insights that enable policymakers to prioritise their cyber security efforts effectively and timely.

For instance, emerging threats include adversarial attacks and deepfakes, which exploit vulnerabilities in machine learning (ML) algorithms to manipulate data or deceive systems for malicious purposes. Adversarial attacks involve the deliberate manipulation of input data to fool ML models, leading to misclassification or incorrect decisions (Zeng et al., 2020). Deepfakes, on the other hand, utilise artificial intelligence to generate highly realistic but fabricated images, videos, or audio recordings, often for spreading disinformation or conducting fraud (Chadha et al., 2021). Another emerging threat is ransomware attack, which encrypts critical data or systems and demand payment for their release, causing substantial financial losses and operational disruptions to targeted organisations (Ghafur et al., 2019). Addressing such emerging threats is crucial as failure to do so could precipitate their proliferation, which leads to severe consequences such as compromised data integrity, reputational damage, financial losses, and disruptions to essential services.

Rapidly increasing threats include a wide array of attacks, such as Distributed Denial of Service (DDoS), and insider threats, which have demonstrated a notable escalation in frequency or severity (Wueest, 2014; Almahmoud et al., 2023). DDoS attacks flood targeted systems or networks with an overwhelming volume of traffic, rendering them inaccessible to legitimate users and disrupting services (Wueest, 2014). Insider threats involve individuals within an organisation exploiting their access privileges or knowledge to compromise security, steal sensitive information, or sabotage operations (Yuan and Wu, 2021). Being prepared to counter these rapidly increasing threats before they escalate is imperative as neglecting to do so may exacerbate their prevalence, potentially resulting in significant harm such as prolonged service disruptions, compromised data confidentiality, and loss of trust in organisational security measures.

The PATs identified for each of these threats are essential components of comprehensive cyber security defence strategies. For instance, technologies such as Anomaly Detection, ML/DL, and Intrusion Detection/Prevention Systems (IDS/IPS) are instrumental in detecting and mitigating adversarial attacks and deepfakes by identifying anomalous patterns or behaviours indicative of malicious activity (National

Academies of Sciences, Engineering et al., 2019; Shao et al., 2022). Similarly, measures such as Access Control, Data Loss Prevention, and User Behaviour Analytics are crucial for addressing insider threats by monitoring and controlling access to sensitive resources and detecting aberrant behaviours or unauthorised activities (Singh et al., 2020).

By forecasting the trend of the emerging and rapidly increasing threats, our study aims to provide actionable insights into evolving cyber threats and inform proactive risk management strategies. By adopting this approach, organisations can effectively prioritise their cyber security efforts, judiciously allocate resources, and implement targeted measures to mitigate emerging risks before they escalate into significant threats. Furthermore, by predicting the trend of PATs tailored to specific attack vectors, organisations can anticipate future gaps between each threat and its corresponding PATs. This foresight enables them to make informed investment and strategic defence decisions, ultimately strengthening their resilience against evolving cyber threats and more effectively safeguarding their assets, data, and operations.

In this work, we construct a comprehensive graphical representation known as the Threats and Pertinent Technologies graph (TPT). This graph links cyber threats with their respective PATs. The connections between threats and PATs are established through edges, with the weight of each edge quantifying the gap between the trend of these interconnected nodes. To accomplish the construction of this graph, we employ a semi-automated methodology, utilising the capabilities of the Generative Pre-trained Transformer (GPT) model (GPT, 2023), in conjunction with Elsevier Application Programming Interface (API) (API, 2023). This approach facilitates the extraction of PATs associated with each threat. Furthermore, we acquire the monthly trend data for each threat node by leveraging news, blogs, and government advisories data, allowing us to tally the number of monthly incidents. Also, for each PAT node, we use Elsevier platform to retrieve the monthly mentions of that PAT, thereby augmenting the dataset proposed in our prior work (Almahmoud et al., 2023). Our methodology extends to the development of a new Bayesian variation of the Graph Neural Network (GNN) model, building upon the framework introduced in the study by Wu et al. (2020). This enhanced model is deployed for the purpose of forecasting the TPT graph over a forthcoming 3-year period, while addressing inherent model uncertainties. The ultimate goal of this endeavour is to provide insightful recommendations for future investments in the cyber threat landscape. In addition to the aforementioned contributions, our analysis extends to the introduction of a novel concept called the Alleviation Technologies Cycle (ATC), which delineates the principal phases within the life cycle of 98 PATs. The contributions of this paper are highlighted below:

- We constructed the graph of 26 emerging and rapidly increasing threats and their PATs, through a semi-automated approach using GPT-3 model and Elsevier API. A novel algorithm called Extractive GPT (E-GPT) which prompts GPT-3 to extract PATs from Elsevier research documents is presented.
- We used big data sources, such as news, blogs, government advisories (Passeri, 2022), Elsevier research documents (Visser et al., 2021), Twitter tweets (Twitter, 2023), and the Python Holidays package (holidays, 2022), to expand upon the dataset introduced in Almahmoud et al. (2023). This expansion includes incorporating monthly trends of 98 PATs from Elsevier, covering the years 2011 to 2022. Additionally, we included recent trends in cyber threats from news articles and blogs, as well as other features related to wars and conflicts from Twitter and public holidays up to the end of 2022.
- We built a novel Bayesian variation of the Multivariate Time-series Graph Neural Network model (B-MTGNN) proposed in Wu et al. (2020) to forecast the graph while addressing the epistemic uncertainty.
- We provided 3 years forecast for the TPT graph, followed by an analysis and categorisation of future gaps, along with recommendations for future investment and defence strategies.
- We proposed the first ATC, illustrating the state of 98 PATs in the coming 3 years, and identifying the key phases in the life cycle of these PATs.
- We provided comparative analysis to show the effectiveness of the proposed model over traditional models and the importance of the features in our dataset.

The remainder of this paper is organised as follows. Section 2 provides a comprehensive overview of existing literature on the topic. Section 3 describes the framework design and development and the building of the proposed model. Section 4 describes our results including the forecast of the graph, our future recommendations, and the proposal of the ATC. Comparative analysis that illustrates the effectiveness of the proposed model and the features in our dataset is provided in Section 5. Section 6 discusses the implications of this work for both research and practice, and addresses its limitations. Finally, Section 7 concludes the paper and suggests directions for the future work.

## 2. Literature review

To the best of our knowledge, this study represents the pioneering exploration of a machine-based approach to forecast the disparity between cyber threats and their PATs. It marks the inaugural application of the MTGNN model to address this challenge. In the subsequent sections, we provide an overview of relevant research in this domain.

### 2.1. Theoretical framework

The PMT serves as a robust theoretical framework for understanding individuals' responses to perceived threats and their adoption of protective measures (Norman et al., 2015). Originating from the field of health psychology, PMT has been widely applied in various domains, including cyber security (Tsai et al., 2016), to elucidate the cognitive processes underlying risk perception and risk management strategies.

Proposed by Rogers in 1975, PMT was initially developed to explain how individuals respond to health-related threats, such as illness or disease (Rogers, 1975). Building upon earlier theories of fear appeals and cognitive appraisal, PMT posits that individuals engage in protective behaviours when they perceive a threat to be sufficiently severe and when they believe that recommended actions are effective in reducing that threat. Over the years, PMT has evolved to encompass a broader range of threats, including those posed by cyber-attacks and online security breaches (Loukaka and Rahman, 2017).

At the core of PMT lies the concept of threat appraisal, wherein individuals assess the severity and vulnerability associated with a threat. In the context of cyber security, this involves analysing historical data and current trends to evaluate the evolving landscape of cyber threats. Unfortunately, experience has shown that human experts tend to show poor inter-rater agreements when exposed to raw data (Shoufan and Damiani, 2017). On the other hand, leveraging big data analytics enables the quantification of the severity and likelihood of various cyber-attacks, ranging from malware infections to sophisticated phishing campaigns (Almahmoud et al., 2023; Werner et al., 2017). Through this comprehensive threat assessment, emerging patterns can be identified, facilitating anticipation of future cyber threat trends.

PMT emphasises individuals' evaluation of the efficacy of available coping strategies in mitigating perceived threats. In the field of cyber security, coping strategies include a wide array of technological measures and behavioural interventions. Existing literature explores the effectiveness of cyber security technologies, such as IDS and encryption protocols, in combating identified cyber threats (Oggier and Mihaljević, 2013; Vinayakumar et al., 2017). Additionally, research investigates the role of user education and training programmes in enhancing cyber security awareness and promoting safe online behaviours (Sharma and Thapa, 2023). By examining the perceived effectiveness of these coping

strategies, insights can be gained to inform the development of targeted interventions to enhance cyber-defence mechanisms.

Research in the field of cyber security has extensively utilised the PMT to understand and improve individuals' adherence to security protocols. One work investigated employees' compliance with organisational security policies (Loukaka and Rahman, 2017), emphasising the significance of self-efficacy and response efficacy in motivating protective actions. Results indicated that employees are more likely to adhere to security policies when they perceive the outcomes and procedures as rewarding and convenient. Another work demonstrated that PMT-based training effectively increased students' threat knowledge and self-efficacy, consequently influencing their cyber security behaviour (Khan et al., 2023). These findings highlight the importance of emphasising self-efficacy in cyber security training programmes, suggesting a promising avenue for educators to develop cyber security practices among students and employees.

Recent studies have further explored the application of PMT in diverse cyber security contexts. One work addressed the motivations of organisational insiders for self-protection (Vrhovec and Mihelič, 2021), emphasising the mediating role of perceived threats in shaping protection motivation. Another work investigated the motivations of entrepreneurs in adopting preventive measures against cyber threats (Bekkers et al., 2023), revealing the complex relationship of factors such as perceived severity, vulnerability, and subjective norms. Other research underscored the importance of user-centred approaches in enhancing cyber security practices, advocating for tailored interventions and acknowledging individual differences among users (Dodge et al., 2023). Moreover, recent research has explored the influence of employees' emotions on cyber security motivation, integrating PMT with other frameworks to offer new insights into the role of emotions in cyber security (Alshammari et al., 2024). These studies collectively highlight the multifaceted nature of protection motivation in cyber security and underscore the importance of tailored interventions and collaborative approaches to effectively improve cyber security practices.

Despite the richness of existing literature in applying PMT to cyber security, a critical gap emerges regarding the direct integration of PMT into proactive cyber threat forecasting methodologies. None of the existing studies explicitly link PMT to forecasting threats, thus limiting our understanding of how psychological motivations can inform predictive models and algorithms for anticipating emerging cyber threats. This gap hampers progress in enhancing cyber security preparedness and resilience, as proactive approaches to threat forecasting remain underexplored. Therefore, there is a pressing need to bridge this gap by incorporating PMT into forecasting frameworks, thereby providing a holistic understanding of individuals' responses to cyber threats.

There is also a lack of application of key principles of the PMT, including threat appraisal and coping strategies appraisal, in the design and development of proactive frameworks aimed at predicting and mitigating cyber threats. While past research has extensively investigated these PMT principles and their application in various problem domains, they have yet to be systematically applied in the context of cyber security trend forecasting. Our work addresses this gap by directly integrating PMT principles into the development of a proactive framework for cyber threat prediction and mitigation. As a result, the proposed approach is expected to enhance the effectiveness and success of cyber security strategies.

Importantly, traditional PMT focusses on current perceptions of threat severity and vulnerability. Our approach enhances this by forecasting future cyber threats, enabling individuals and organisations to anticipate and prepare for these threats ahead of time. This foresight improves the perceived relevance and urgency of potential threats, maintaining a high level of perceived vulnerability and severity over time. Also, while current PMT evaluates present coping mechanisms, a proactive PMT includes the assessment of emerging and future alleviation technologies. By identifying and investing in these technologies in advance, individuals and organisations can ensure they are equipped with the most effective tools and strategies to counteract anticipated threats. This proactive stance boosts response efficacy and self-efficacy, as confidence in future protective measures is reinforced.

Adopting a data-driven version of PMT as the basis for our research framework provides a comprehensive understanding of individuals' responses to cyber threats. By assessing the intensity of future threats and bridging gaps with relevant mitigation technologies through continuous evaluation and informed investments, individuals will have confidence in the effectiveness of security measures and will engage in proactive actions to improve cyber security preparedness and resilience. This integration can lead to the development of more robust cyber security strategies that not only address current threats but also anticipate and mitigate future risks based on psychological motivations.

### 2.2. Cyber threat forecasting

We categorised cyber threat forecasting based on the prediction timeframe into three main categories. These are long-term (years ahead) (Almahmoud et al., 2023), midterm (months ahead) (Okutan et al., 2019; Liu et al., 2015), and short-term (hours ahead) (Husák and Kašpar, 2019; Husák et al., 2021). The practicality and usefulness of each category depend on the specific objectives and the context of cyber security efforts. Each category has its own set of advantages and challenges as outlined below:

- Long-term Predictions (Years Ahead): Long-term predictions are crucial for strategic planning, policy formulation, and setting cyber security standards. They help organisations and governments anticipate major trends in cyber threats, such as the rise of new types of malware or attack vectors, enabling proactive development of defence mechanisms (Almahmoud et al., 2023). By understanding potential future threats, organisations can allocate resources more effectively, investing in the development of new technologies, training, and infrastructure improvements that will be most relevant in the face of anticipated threats. Yet, a primary challenge with long-term predictions is the high level of uncertainty. The cyber threat landscape evolves rapidly due to technological advancements, changes in attacker tactics, and geopolitical developments (Okutan et al., 2019). Long-term predictions may become outdated quickly, requiring continuous monitoring and adjustment.
- Midterm Predictions (Months Ahead): Midterm predictions are valuable for operational planning, including the deployment of specific security measures, conducting targeted training sessions, and performing security drills or simulations based on anticipated attack scenarios. Organisations can adjust their security postures based on midterm predictions, fine-tuning firewalls, intrusion detection systems, and response protocols to guard against expected threats (Bilge et al., 2017). However, the accuracy of midterm predictions can be affected by sudden changes in attacker behaviour or the emergence of unforeseen vulnerabilities (Almahmoud et al., 2023). These predictions require a balance between specificity and adaptability.
- Short-term Predictions (Hours Ahead): Short-term predictions are critical for immediate threat detection and response. They can enable real-time security measures, such as blocking an imminent attack or isolating affected systems to prevent the spread of malware. Predictions over shorter timeframes can be more precise and actionable, leveraging real-time data analytics and machine learning models to identify and respond to threats as they emerge (Husák and Kašpar, 2019). Nevertheless, short-term predictions require extensive monitoring and data analysis capabilities. The high volume of false positives and the need for rapid, automated decision-making systems can be challenging to manage (Almahmoud et al., 2023).

**Table 1**
Cyber-attack forecasting - Literature review summary and our contribution.

| Ref. | Problem domain | Forecast period | Forecast coverage | Methods |
|---|---|---|---|---|
| Werner et al. (2017, 2018) | Forecast attack count | 1–7 days | Multiple targets | ARIMA model |
| Okutan et al. (2019) | Forecast attack count | Months | Organisation | Unconventional signals, lagged feature selection, concept drift training |
| Munkhdorj and Yuji (2017) | Forecast attack motivation and opportunity | 1 week | 1 target | Social media analysis, SVM, CNN |
| Goyal et al. (2018) | Forecast attack count | 1 week or month | Organisation | Digital traces, ARIMA, ARIMAX, LSTM |
| Qin and Lee (2004) | Predict next attack in the chain | N/A | 1 target | Bayesian network |
| Husák and Kašpar (2019) | Predict intrusion detection alerts | Minutes or hours | Organisation | Stream processing, sequential rule mining |
| Liu et al. (2015) | Forecast if a data breach will occur | Months | Organisation | Externally measurable features, Random Forest |
| Malik et al. (2020) | Reconnaissance detection | N/A | N/A | LSTM, CNN |
| Bilge et al. (2017) | Forecast if a machine will be infected | Months | Machine | Binary file analysis, semi-supervised learning |
| Husák et al. (2021) | Forecast if an IP address will attack | 24 h | N/A | Entity reputation and scoring, decision trees |
| **Ours** | **Forecast cyber-attack trends** | **3 years** | **36 countries** | **Big data, multivariate time series analysis, graph neural network** |

Most of the existing studies on cyber threat forecasting focus on predicting the attacks in the short and midterm (Werner et al., 2017, 2018; Okutan et al., 2019; Munkhdorj and Yuji, 2017; Goyal et al., 2018; Qin and Lee, 2004; Husák and Kašpar, 2019; Liu et al., 2015; Malik et al., 2020; Bilge et al., 2017; Husák et al., 2021), such as predicting the expected number of attacks within a few hours, days, or months. Most of this research is conducted within limited settings, such as against a specific entity or organisation (Okutan et al., 2019; Bilge et al., 2017; Munkhdorj and Yuji, 2017). Some approaches utilise statistical methods assuming parametric data distributions (Werner et al., 2017, 2018), while others employ ML models (Goyal et al., 2018). Compared to statistical methods, ML can capture complex relationships to provide more accurate predictions. It is also possible to utilise a hybrid-approach to improve the prediction performance (Athanasopoulou et al., 2021). Bayesian methods have also been employed, constructing event graphs to estimate the conditional probability of an attack based on a given chain of events (Qin and Lee, 2004). However, these techniques rely on predefined attack graphs and are incapable of addressing previously unseen attacks. Other approaches aim to predict the source of the attack using network entity reputation and scoring (Husák et al., 2021). A growing body of research focusses on the utilisation of external features or warning signals to forecast cyber threats using ML. These features include the number of mentions of a victim on Twitter (Okutan et al., 2019) or in the news articles (Munkhdorj and Yuji, 2017), or represent digital traces from dark web forums (Goyal et al., 2018). Table 1 provides a summary for the related work on cyber-attack forecasting and highlights our contribution.

Given the reactive nature of past cyber defence approaches and the limited predictive capabilities of existing tools, long-term predictions offer significant value in shifting towards a more proactive cyber security posture. Long-term predictions are essential for strategic decision-making and guiding the development of future-proof security technologies.

### 2.3. Technology forecasting

Many of the existing work on technology forecasting rely on the judgement of human experts or adopt semi-automated approach (GRAY, 2001; Adomavicius et al., 2008; Li et al., 2019; Chandra and Collis, 2021). The early work forecasted future generations of tools and technologies based on human imagination and creativity (GRAY, 2001). This is possible by exploring the idea that many of the high technology products we use today were once conceptualised in science fiction before becoming a reality through technological advancements. In Adomavicius et al. (2008), a technology ecosystem model was introduced, which offers analysts a tool to navigate the intricate relationships among technologies. The model aids in dissecting the interplay of various factors influencing technological change, enhancing technology forecasts, investments, and development decisions. More recently, Li et al. (2019) proposed a framework that utilises scientific papers and patents as data sources, while incorporating text mining and expert judgement techniques to predict technology trends.

The Gartner Hype Cycle (GHC) is a graphical representation and methodology that helps organisations understand the maturity and adoption of technologies over time (Chandra and Collis, 2021). It was developed by the research and advisory firm Gartner, Inc (Gartner, 2023). The GHC is based on the premise that technologies go through predictable stages of enthusiasm, disillusionment, and eventual adoption. It was derived from observing and analysing the patterns of technology adoption and understanding how people perceive and adopt new technologies. Gartner's analysts study the life cycle of various technologies, their visibility, and their market expectations to position them on the GHC. Table 2 summarises the existing work on technological forecasting compared to our approach.

### 2.4. Time series forecasting with graph neural networks

Time series forecasting with GNNs has been heavily applied in the domain of traffic prediction (Yu et al., 2017; Guo et al., 2019). In Yu et al. (2017), a deep learning framework called Spatio-Temporal Graph Convolutional Networks (STGCN) was developed for learning spatio-temporal correlations by modelling multi-scale traffic networks. In Guo et al. (2019), dynamic spatial temporal correlations were studied through the use of spatial–temporal attention mechanism. Other methods jointly learn inter-series correlations and temporal dependencies in the spectral domain, by combining Graph Fourier Transform (GFT) and Discrete Fourier Transform (DFT) (Cao et al., 2020).

A recent study introduced a generic GNN model for forecasting multivariate time series data (Wu et al., 2020), applicable across various domains. The model includes a graph learning layer capable of learning

**Table 2**
A comparison between existing approaches to technology forecasting and our approach.

| Ref. | Problem domain | Approach | Methods |
|---|---|---|---|
| GRAY (2001) | Forecast future generations of tools and technologies | Human-based | Human imagination and creativity |
| Adomavicius et al. (2008) | Understand evolution in technology ecosystems | Human-based | Navigating the complex relationships among technologies |
| Li et al. (2019) | Forecast technology trends | Semi-automated | Text mining and expert judgement |
| Chandra and Collis (2021) | Understand the maturity and adoption of technologies over time | Human-based | Observation and analysis by human expert |
| **Ours** | **Forecast the trend of cyber threat related technologies** | **Machine-based** | **Big data, multivariate time series analysis, graph neural network** |

the hidden adjacency matrix in the graph using latent representation of nodes. In addition, the model includes temporal convolution modules and graph convolution modules interleaved with each other for learning both the temporal and the spatial dependencies in the graph. The model was evaluated on multiple datasets and was shown to be effective compared to the state-of-the-art baselines. In our work, we propose the Bayesian variation of this model which expresses the model uncertainty, and apply the model in the cyber security domain. Table 3 summarises the existing GNN models for time series forecasting and highlights our contribution.

**Table 3**
A comparison between existing GNN models for time series forecasting and our model.

| Ref. | Model | Description | Domain |
|---|---|---|---|
| Yu et al. (2017) | Spatio-Temporal Graph Convolutional Network (STGCN) | Learns spatio-temporal correlations by modelling multi-scale traffic networks. | Traffic |
| Guo et al. (2019) | Attention based Spatial–Temporal Graph Convolutional Network (ASTGCN) | Learns dynamic spatial temporal correlations using spatial–temporal attention mechanism. | Traffic |
| Cao et al. (2020) | Spectral Temporal Graph Neural Network (StemGNN) | Learns inter-series and temporal dependencies in the spectral domain using GFT and DFT. | Traffic, energy, ECG |
| Wu et al. (2020) | Multivariate Time-series Graph Neural Network (MTGNN) | Jointly learns the adjacency matrix and the spatial and temporal dependencies. | Traffic, energy, exchange |
| **Ours** | **Bayesian Multivariate Time-series Graph Neural Network (B-MTGNN)** | **Jointly learns the adjacency matrix and the spatial and temporal dependencies, and expresses model uncertainty.** | **Cyber security** |

## 3. Methods

### 3.1. Research model

Our research model includes a comprehensive framework designed to forecast the trend of cyber threats and PATs in the long-term, providing actionable insights for proactive risk management strategies. The model integrates various components, including data collection, graph construction, forecasting methodology, and analysis of future trends. Here, we outline the key elements of our research model:

#### 3.1.1. Data collection

The foundation of our research model lies in the extensive collection and expansion of data related to cyber threats and PATs. We utilise diverse big data sources such as news articles, blogs, government advisories, Elsevier research documents, Twitter data, and public holidays to gather information spanning from 2011 to 2022. This dataset includes monthly trends of 26 cyber threats and 98 PATs in addition to external features, providing a robust basis for analysis.

#### 3.1.2. Graph construction

We construct the TPT graph, linking emerging and rapidly increasing cyber threats with their corresponding PATs. The graph facilitates visual representation of the relationships between threats and technologies, with edges quantifying the gap in trend trajectories. The identification of the edges involves leveraging the GPT model to identify the set of PATs relevant to each threat. The identification of the gap value involves labelling each node with its trend value from our dataset and labelling each edge with the trend difference between its connected nodes.

#### 3.1.3. Forecasting methodology

To forecast the TPT graph over a forthcoming 3-year period, we develop a novel Bayesian variation of the MTGNN model (B-MTGNN). This enhanced model addresses inherent uncertainties and leverages the framework introduced by Wu et al. (2020). By incorporating both historical trend data and current observations, the forecasting methodology provides insights into future cyber threat landscapes and identifies potential gaps between threats and PATs.

#### 3.1.4. Analysis and recommendations

The final stage of our research model involves analysis of forecasted trends and formulation of actionable recommendations. We categorise future gaps between threats and PATs, prioritise defence strategies, and suggest investments based on predictive insights. Additionally, we introduce the Alleviation Technologies Cycle (ATC), delineating the life cycle phases of PATs and offering strategic guidance for cyber security management.

Our research model contributes significantly to the field of cyber threat forecasting and proactive risk management. By integrating automated data collection, graph-based representation, advanced forecasting techniques, and strategic analysis, the model offers a holistic approach to addressing evolving cyber security challenges.

### 3.2. Forecasting framework

The framework's architecture for forecasting cyber threats and PATs is shown in Fig. 1. As illustrated in the figure, our framework leverages a variety of unstructured data sources to gather all relevant information and extract valuable insights. Among these sources, the news, blogs,

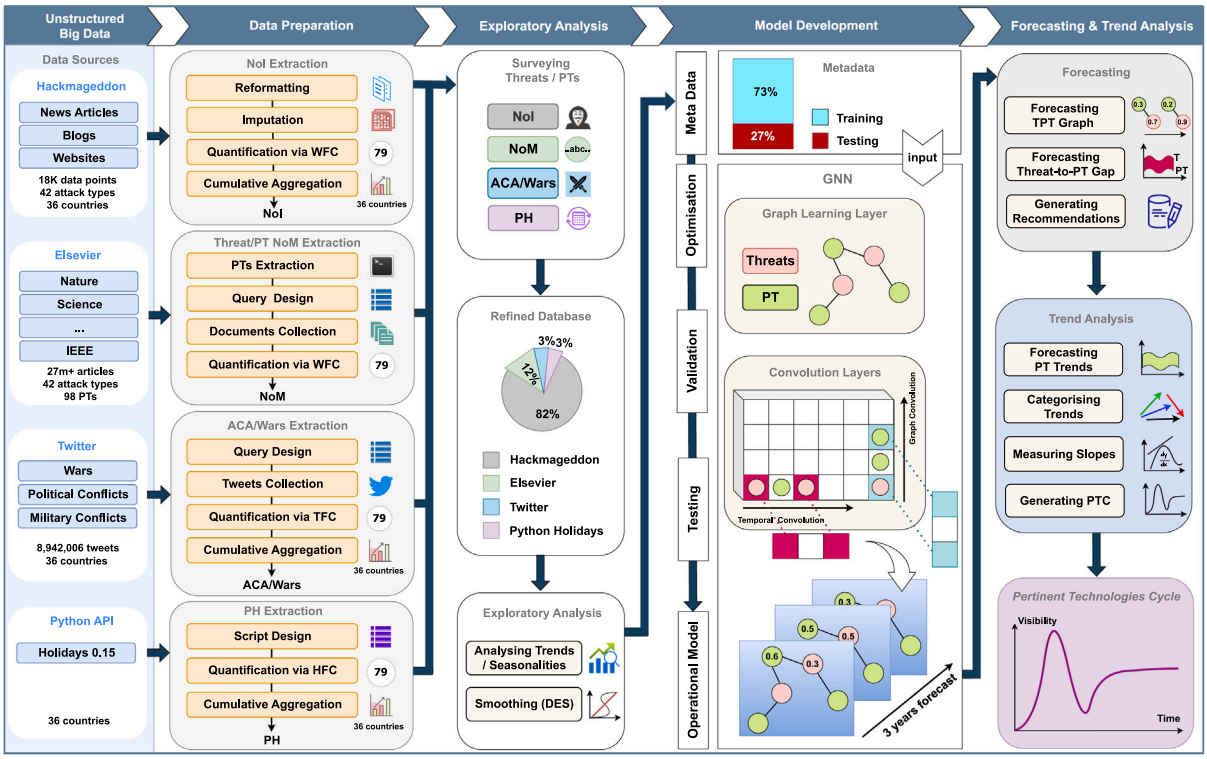


**Fig. 1.** The workflow and architecture of forecasting cyber threats and pertinent alleviation technologies. NoI: Number of Incidents, NoM: Number of mentions, ACA: Armed Conflict Areas, PH: Public Holidays, PT: Pertinent Technology, WFC: Word Frequency Counter, TFC: Tweet Frequency Counter, HFC: Holiday Frequency Counter, DES: Double Exponential Smoothing, TPT: Threats and Pertinent Technologies, PTC: Pertinent Technologies Cycle.

and government advisories' websites play a crucial role, providing an extensive collection of textual data on major cyber-attacks (approximately 18,000 incidents) since July 2011. The monthly count of attacks represents the ground truth of the attacks' trend, and is denoted as the *Number of Incidents* (NoI). Furthermore, by utilising Elsevier API, we gained access to a vast repository of scientific articles from numerous sources. Through this API, we acquired the *Number of Mentions* (NoM) for each attack type and each PAT, which indicates their frequency in scientific publications, typically on a monthly basis. This NoM feature is particularly significant as it serves as a reliable reference for attack types that may not be present in other sources and also represents the ground truth of the PATs' trend. During the initial research phase, we thoroughly examined all potential features and identified a strong correlation between wars and political conflicts and the occurrence of cyber-events. To capture this information, we extracted relevant tweets using the Twitter API, specifically focussing on the number of tweets about *Armed Conflict Areas/Wars* (ACA). Finally, considering that cyber-attacks often coincide with holidays, we employed Python's Holidays package to obtain the count of public holidays per month for each country, denoted as *Public Holidays* (PH).

For the extraction of NoI, the data preparation phase, as illustrated in Fig. 1, starts by collecting and arranging all incidents in a tabular format including the date and attack type in addition to the description and country. This is followed by noise reduction through the handling of missing values, particularly in the earlier years. Here, imputation techniques were employed, utilising information from the description column or external sources such as reliable articles found through Google searches to supplement missing country data. Next, quantification of the textual data involved the implementation of a *Word Frequency Counter*, tallying the occurrences of each attack type per month for each country. Finally, cumulative aggregation facilitated the calculation of attack counts per month for all countries collectively (36 countries).

Prior to extracting NoM, we extracted the PATs by prompting GPT to extract relevant technologies to each attack type from Elsevier abstracts and also through a direct prompt to GPT. We then queried Elsevier API to collect the research documents relevant to each attack type and each extracted PAT. This was followed by running a Python script to obtain NoM for each attack type and each PAT per month within the collected documents.

The extraction of ACA from Twitter involved designing a script that included a query for collecting all tweets about wars and political conflicts relevant to each of the 36 countries in the study and during each month. A *Tweet Frequency Counter* was then used to count the number of such tweets for each individual country per month, followed by a cumulative aggregation to obtain the total number of tweets per month. The extraction of PH was done by writing a Python script including a *Holiday Frequency Counter* to obtain the number of public holidays per month for each country followed by a cumulative aggregation to obtain the total number of holidays per month.

While we focus in this study on 26 emerging and rapidly increasing threats, our monthly dataset includes the trend of 42 attack types in 36 countries, in addition to 98 PATs. Based on the above, we obtain the following columns for each month:

- NoI_C: The number of incidents for each attack type in each country (42 × 36 columns) [News, blogs, government advisories].
- NoI: The total number of incidents for each attack type (42 columns) [News, blogs, government advisories].
- NoM_A: The number of mentions of each attack type in research articles (42 columns) [Elsevier].
- NoM_P: The number of mentions of each alleviation technology in research articles (98 columns) [Elsevier].
- ACA_C: The number of tweets about wars and conflicts related to each country (36 columns) [Twitter].
- ACA: The total number of tweets about wars and conflicts (1 column) [Twitter].
- PH_C: The number of public holidays in each country (36 columns) [Python].

• PH: The total number of public holidays (1 column) [Python].

In the aforementioned list of columns, the name enclosed within square brackets denotes the source of data. By matching and combining these columns, we derive our monthly dataset, wherein each row represents a distinct month. Further details about the data acquisition process can be found in our previous work (Almahmoud et al., 2023).

To gain insights into the dataset's main characteristics, an exploratory analysis was conducted. This analysis involved visualisations to identify key patterns such as trends, seasonality, correlated features, missing data, and outliers. Seasonal data was smoothed to unveil underlying trends while mitigating noise, employing double exponential smoothing (Lai et al., 2006).

In terms of modelling, B-MTGNN was constructed. The MTGNN model has been successfully applied to traffic prediction among other problems (Wu et al., 2020). The model captures both temporal and spatial dependencies in the graph through temporal convolution and graph convolution layers and can additionally learn hidden relationships between nodes using a graph learning layer. Learning such relationships is useful for improving prediction performance. This is in contrast to relying on fixed, pre-assumed relationships between the nodes. We demonstrate this improvement experimentally in Section 5.

The proposed Bayesian variation of the MTGNN model treats the model weights as random variables, allowing for the quantification of epistemic uncertainty through approximate Bayesian inference. Epistemic uncertainty quantifies the prediction error resulting from insufficient information (Mae et al., 2021) and can be reduced by acquiring more samples and informative features. The overall model development phase produced an operational model that can be readily used for forecasting the TPT graph. The performance of this model in predicting the trends up to 36 months in advance was evaluated. The model was ultimately used to forecast future trends and provide investment and strategic defence recommendations based on the predicted disparities between threats and PATs. Moreover, the analysis of past and future trends facilitated the development of the ATC. This was achieved through the categorisation of the trends and the analysis of their slope and direction.

### 3.3. Graph construction

The TPT graph consists of nodes representing the threats and PATs, supplemented by other feature nodes during the modelling step. The value of the node represents the trend level (NoI for threats and NoM for PATs). The edges link each threat to its PATs. The edge weight represents the gap between the threat trend and the connected PAT's trend. Formally, we define $T$ as the total number of rows or months in the dataset, $N$ as the total number of columns or features, and $D$ as the feature dimension, which is set to 1 in our case. Let $t$ denote the threat and $p$ denote the PAT. The gap between $t$ and $p$ in a given month $m$ is given by the following formula,

$$\mathrm{G}_{t,p}(m) = \frac{NoI_{m,t}}{\max_{i\in\mathcal{M},u\in\mathcal{T}} NoI_{i,u}} - \frac{NoM_{m,p}}{\max_{i\in\mathcal{M},v\in\mathcal{P}} NoM_{i,v}} \tag{1}$$

where $\mathbf{NoI} \in \mathbb{R}^{T\times N_{\text{threats}}}$ and $NoI_{m,t}$ represents the trend of threat $t$ in month $m$. Similarly, $\mathbf{NoM} \in \mathbb{R}^{T\times N_{\text{pats}}}$ and $NoM_{m,p}$ represents the trend of PAT $p$ in month $m$. Also, $\mathcal{T}$ is the set of all threats, $\mathcal{P}$ is the set of all PATs, and $\mathcal{M}$ is the set of all months.

The formula normalises the node value over the maximum NoI in the dataset in the case of threats, and over the maximum NoM in the dataset in the case of PATs. This normalisation approach is crucial as it effectively bridges the significant scale disparities between NoI and NoM. The resulting gap value falls within the range −1 to 1, with a positive gap denoting a relatively lower research effort compared to the number of incidents, while a negative gap signifies a higher research effort. Ideally, the gap value should approach zero to indicate a balanced alignment between research efforts and incident occurrences.

In our study, we focus on the emerging and rapidly increasing threats identified in Almahmoud et al. (2023) based on past and future analysis, since these threats require the highest attention when investing in related technologies, compared to the other declining threats. The threats in our study are shown in Table 4.

To extract the PATs of each threat, we propose the E-GPT algorithm shown in Algorithm 1. Given a threat $t$, the algorithm starts by collecting relevant abstracts from Elsevier database. These abstracts include technology related keywords along with $t$. The second step is to iteratively prompt the GPT model to extract PATs from each abstract. The prompt to GPT contains an example for an expected answer in order to improve the performance. Given that there are many abstracts and many keywords that could be returned, the ranking of the PATs is then performed to obtain the top $n$ PATs. In our study, we set $n$ to 10. The ranking is done by considering the frequency defined as the number of times the PAT was returned by GPT. Intuitively, we give higher priority to the PATs with higher frequency. Within the same frequency groups, we perform a secondary ranking that prioritises the PATs that appear closely to technology-related keywords in the abstract, such as the word "solution". This is done by computing the minimum distance within the abstract between the PAT and any of the keywords that belong to a predefined list of keywords $S$. The average distance is kept track of since a PAT can be returned multiple times by GPT. This secondary ranking is motivated by the fact that technology terms are frequently mentioned in close proximity to other keywords in the text. When they appear further in the text, they are more likely to be irrelevant.

One important benefit of the extractive method is to ensure that the returned PATs reflect the state-of-the-art, since GPT can be outdated. Another benefit is controlling the source of information to ensure data reliability. However, to obtain more general answers and improve the accuracy, the list of PATs for each threat is further appended with an additional list that we obtain by asking GPT a direct question (*e.g.*, What are the PATs to $t$?). Finally, manual adjustment by human experts is performed to filter out irrelevant terms or add missing PATs. The final list of threats and PATs in the graph is shown in Table 4. The PATs abbreviations table can be found in Fig. 5.

**Algorithm 1:** Extractive GPT

**input :** Threat $t$, number of PATs $n$, number of abstracts $b$, list of technology related keywords $S$

**output:** top $n$ PATs to $t$

$\mathcal{P} = \{\}_{\text{set}}$, $frequency = \{\}_{\text{dict}}$, $m\_distance = \{\}_{\text{dict}}$
$\mathcal{A}$ = Query_Elsevier_for_PATs_Abstracts($t, b$)
**for** $a$ *in* $\mathcal{A}$ **do**
  $\mathcal{U}$ = Prompt_GPT_to_Extract_PATs($t, a$)
  **for** $p$ *in* $\mathcal{U}$ **do**
    $frequency[p]$++
    $m\_distance[p]$ = avg_min_distance($p, S, a$)
  $\mathcal{P} = \mathcal{P} \cup \mathcal{U}$
sort $\mathcal{P}$ by $frequency$ in descending order
sort by $m\_distance$ in ascending order within the same frequency groups
return $\{p_1, p_2, \ldots, p_n\}$ where $p_i \in \mathcal{P}$ for $i = 1$ to $n$

### 3.4. Model development

To forecast the TPT graph, we developed a Bayesian variation of the MTGNN model proposed by Wu et al. (2020). This model was originally introduced as a general framework for forecasting multivariate time

**Table 4**
Threats and pertinent alleviation technologies in our study.

| Threat | Type | Pertinent alleviation technologies |
|---|---|---|
| Account Hijacking | RI | AC, AD, CAPTCHA, CR, IDS/IPS, IdM, LP, MFA, ML/DL, NLP/LLM, PT, SM |
| Adversarial Attack | E | AD, AdT, BN, DA, DD, DP, DR, DS, ML/DL, NI, NLP/LLM, OD, RRAM, SS, TAI |
| APT | RI | AC, DLP, DRM, DT, GT, IDS/IPS, LP, MFA, ML/DL, NLP/LLM, NS, PT, RA, UBA |
| Backdoor | RI | AD, DAS, IDS/IPS, ML/DL, PT, SA |
| Botnet | RI | AD, BC, BH, BT, CAPTCHA, GM, GT, HP, IDS/IPS, ML/DL, NLP/LLM, PF, PT, RC, RL, SDN, TS |
| Brute Force Attack | RI | CAPTCHA, CR, DBI, IDS/IPS, MFA, ML/DL, OTP, PH, PT |
| Cryptojacking | E | BT, ML/DL, PT, TA |
| DDoS | RI | BC, BH, BT, IDS/IPS, ML/DL, NLP/LLM, PF, PT, RC, RL, TS |
| Data Poisoning | E | AD, AdT, BN, DP, DS, ML/DL, NLP/LLM, OD, TAI |
| Deepfake | E | 3DFR, AD, BO, DW, LD, ML/DL, NLP/LLM |
| Disinformation | RI | BC, CA, DLT, DP, DT, GT, HG, IR, ML/DL, NLP/LLM, SI |
| DNS Spoofing | RI | BC, CR, DNSSEC, ML/DL, PT, RA |
| Dropper | RI | AW, CS, FIM, IDS/IPS, ML/DL, NLP/LLM, PT, SBX |
| Insider Threat | RI | AC, AD, AM, AT, CR, DLD, IDS/IPS, KD, LP, ML/DL, MTD, NLP/LLM, PT, UBA |
| IoT Device Attack | E | AD, BC, CR, IDS/IPS, IdM, MFA, ML/DL, MS, PT, SB |
| Malware | RI | AC, AD, AW, BBD, BC, CR, CS, DAS, DB, DM, DT, FIM, FV, GT, HP, IDS/IPS, ML/DL, NLP/LLM, PMT, PT, SA, SB, SBX, SHMM, SMF, VK |
| MITM | RI | BC, CAPTCHA, CP, CR, ML/DL, PKI, PT, SSL/TLS, SSP, VPN |
| Password Attack | RI | CAPTCHA, CR, GA, IDS/IPS, MA, MFA, ML/DL, NLP/LLM, OTP, PH, PM, PP, PSM, PT |
| Phishing | RI | AC, BT, CR, DT, MA, MFA, ML/DL, NLP/LLM, PKI |
| Ransomware | E | AC, AD, AW, BC, CR, DAS, DB, DT, IDS/IPS, ML/DL, NLP/LLM, PMT, PT, SA, SHMM |
| Session Hijacking | RI | AD, CA, CR, Https, IBE, ML/DL, PT, SAT, SM, SSL/TLS |
| Supply Chain Attack | RI | AC, AD, BC, CR, IdM, ML/DL, NLP/LLM, PT, SCRM |
| Targeted Attack | RI | AC, DRM, DT, GT, IDS/IPS, LP, MFA, ML/DL, NLP/LLM, NS, PT, RA, UBA |
| Trojan | RI | AD, BBD, CR, FV, GT, IDS/IPS, ML/DL, NLP/LLM, PT, SMF |
| Vulnerability | RI | CFI, IDS/IPS, ML/DL, NLP/LLM, PMT, PT, SC, SIEM, VA, VM, VS |
| Zero-day | RI | AD, DT, FIM, GT, IDS/IPS, ML/DL, NLP/LLM, PrP, VM, VPN |

The list of attack types in the **Threat** column are the emerging and rapidly increasing threats identified in Almahmoud et al. (2023) based on past and future analysis. These threats require the highest attention when investing in related technologies, compared to the other declining threats. In the **Type** column, RI refers to the rapidly increasing threats and E refers to the emerging threats. The list of PATs for each attack type was extracted using Algorithm 1. The PATs abbreviations table can be found in Fig. 5.

series, while leveraging state-of-the-art graph neural network components. The model's efficacy was extensively examined and validated across various datasets from different domains (Wu et al., 2020).

Within the context of our research, we apply the aforementioned model to the task of TPT graph forecasting. Furthermore, we enhance its capabilities by addressing the epistemic uncertainty inherent in the model's forecasts. This augmentation allows the model to articulate and quantify its uncertainty during the prediction process, a valuable asset when confronted with limited data or when seeking a measure of the model's confidence in its predictions (Gal and Ghahramani, 2016).

The developed model is depicted in Fig. 2. The first component in the model is the graph learning layer, which aims to adaptively learn the adjacency matrix in the graph. The learning process is designed in such a way that the resulting adjacency matrix leads to more accurate predictions in terms of node values. This approach is more effective than assuming predefined relationships since these can be hidden, unclear, or difficult to quantify. In our scenario, there are additional feature nodes beyond the threats and PATs, including NoM of the threats, ACA, and PH. The connections between these nodes and the threat/PAT nodes are not predefined. Therefore, we opt to let the model learn these hidden links and their weights within the graph.

Given randomly initialised node embeddings $\mathbf{E}_1$, $\mathbf{E}_2 \in \mathbb{R}^{N\times V}$, where $V$ is a hyper-parameter denoting the node dimension, the graph learning layer extracts uni-directional relationships by computing the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N\times N}$ as follows,

$$\mathbf{M}_1 = \tanh(\alpha \mathbf{E}_1 \mathbf{\Theta}_1) \tag{2}$$

$$\mathbf{M}_2 = \tanh(\alpha \mathbf{E}_2 \mathbf{\Theta}_2) \tag{3}$$

$$\mathbf{A} = \mathrm{ReLU}(\tanh(\alpha(\mathbf{M}_1\mathbf{M}_2^T - \mathbf{M}_2\mathbf{M}_1^T))) \tag{4}$$

$$\mathbf{A}[i, -\mathrm{argtopk}(\mathbf{A}[i,:])] = 0, \forall i \in [N] \tag{5}$$

where $\mathbf{\Theta}_1$, $\mathbf{\Theta}_2 \in \mathbb{R}^{V\times V}$ are model parameters, $\mathbf{M}_1$, $\mathbf{M}_2 \in \mathbb{R}^{N\times V}$, $\alpha$ is a hyper-parameter for controlling the saturation rate of the activation function, and argtopk(.) returns the index of the top $k$ closest nodes

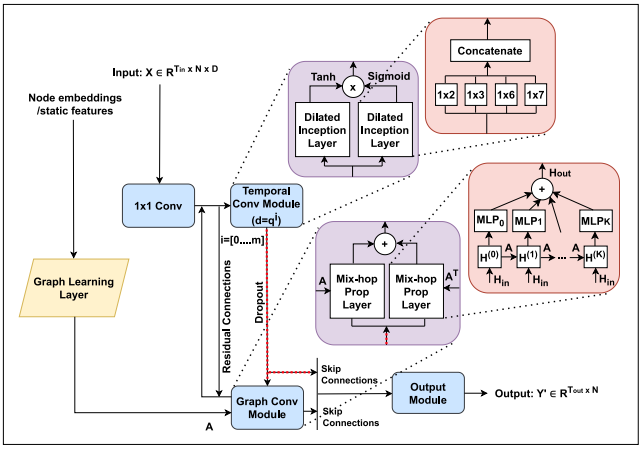


**Fig. 2.** The B-MTGNN model learns the adjacency matrix of the graph through the graph learning layer, while capturing temporal and spatial dependencies using temporal and graph convolution modules. The dilation factor $d$ increases exponentially with the increase in the number of layers $m$ at the rate of $q$. The red arrows indicate the use of dropout during inference to approximate a Bayesian model. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

to be selected as neighbours. This selection strategy makes the adjacency matrix sparse while reducing the computation cost of the graph convolution (Wu et al., 2020).

The graph convolution module aims to fuse a node's information with its neighbours' information to capture the spatial dependencies. As shown in Fig. 2, it consists of two mix-hop propagation layers for processing inflow and outflow information for each node. The mix-hop propagation layer mainly consists of two steps. The first step is the information propagation step defined as follows,

$$\mathbf{H}^{(k)} = \beta \mathbf{H}_{in} + (1-\beta)\tilde{\mathbf{A}}\mathbf{H}^{(k-1)} \tag{6}$$

where $\mathbf{H}^{(k)} \in \mathbb{R}^{B\times C\times N\times O}$. Here, $B$ is the batch size, $C$ is the number of convolution channels, and $O$ is the last dimension of the output from

the previous layer. $\beta$ is a hyper-parameter for controlling the amount of information to be retained from the root node's original states, and $\mathbf{H}_{in} \in \mathbb{R}^{B\times C\times N\times O}$ denotes the input hidden states from the previous layer. The second step is the information selection step given by the following formula,

$$\mathbf{H}_{out} = \sum_{k=0}^{K} \mathbf{H}^{(k)}\mathbf{W}^{(k)} \tag{7}$$

where $\mathbf{H}_{out} \in \mathbb{R}^{B\times I\times N\times O}$ denotes the output hidden states of the current layer, where $I$ is a hyper-parameter that denotes the number of residual channels. $K$ is the propagation depth, and $\mathbf{W}^{(k)} \in \mathbb{R}^{I\times C}$ is a feature selector for controlling what to be retained from the original node's information. Further details about these steps can be found in Wu et al. (2020).

As shown in Fig. 2, the temporal convolution module captures the temporal dependencies by utilising dilated inception layers. Given that the receptive field increases exponentially with the increase in the number of layers, the dilation strategy is employed to handle large sequences while reducing the model complexity (Oord et al., 2016). The inception strategy is used to handle temporal patterns with different ranges by using filters with multiple sizes (Szegedy et al., 2015). Formally, given a sequence input $\mathbf{z} \in \mathbb{R}^{T_{\text{in}}}$ and four filters of the form $\mathbf{f}_{1\times 2} \in \mathbb{R}^2$, $\mathbf{f}_{1\times 3} \in \mathbb{R}^3$, $\mathbf{f}_{1\times 6} \in \mathbb{R}^6$, and $\mathbf{f}_{1\times 7} \in \mathbb{R}^7$, the dilated inception layer takes the following form,

$$\mathbf{z} = \text{concat}(\mathbf{z} \star \mathbf{f}_{1\times 2}, \mathbf{z} \star \mathbf{f}_{1\times 3}, \mathbf{z} \star \mathbf{f}_{1\times 6}, \mathbf{z} \star \mathbf{f}_{1\times 7}) \tag{8}$$

Let $d$ denote the dilation factor. The dilated convolution denoted by $\mathbf{z} \star \mathbf{f}_{1\times k}$ is defined as follows,

$$\mathbf{z} \star \mathbf{f}_{1\times k}(t) = \sum_{s=0}^{k-1} \mathbf{f}_{1\times k}(s)\mathbf{z}(t - d \times s) \tag{9}$$

Since we have relatively short time series within our refined data (*i.e.*, 138 monthly data points between July 2011 and December 2022), it is vital to extract the model's uncertainty. Deterministic neural network models that do not involve randomness are insufficient for this task, since they offer single-point predictions of model parameters. Instead, we employ a Bayesian approach to capture epistemic uncertainty. Specifically, we employ the Monte Carlo dropout method proposed by Gal and Ghahramani (2016), who showed that the use of dropout neurons during inference provides a Bayesian approximation of the deep Gaussian processes. The use of dropout mask in our model during inference is highlighted in red arrows (Fig. 2). Therefore, during the prediction phase, the trained model runs multiple times, which results in a distribution of prediction (representing the uncertainty) rather than a single point (Fig. 3).

### 3.5. Experimental settings

In our experimental setup, we partitioned the dataset into three distinct subsets: 43% for training, 30% for validation, and 27% for testing. This allocation was carefully chosen to ensure that ample data was available for rigorous testing of the model's performance. Specifically, our model's input comprises 10 months of historical data, corresponding to 10 time steps, while the output encompasses forecasts for the subsequent 36 months. This forecasting framework constitutes a multi-horizon approach, wherein predictions are made for multiple future time steps simultaneously.

Our experimental findings support the utilisation of a non-autoregressive approach in our forecasting methodology. Training the model to predict multiple time steps concurrently, without dependence on previously generated predictions, yielded higher accuracy and more comprehensive pattern capture. Unlike autoregressive models, our approach solely utilises past observed values for forecasting the subsequent months. By avoiding reliance on prior predictions, our model mitigates the error propagation problem, leading to enhanced forecasting accuracy and efficacy (Taieb et al., 2010).

### 3.6. Hyper-parameter optimisation

We performed a random search with 60 iterations, in order to find the set of hyper-parameters that produces the model with the lowest validation error. Random search is a simple method for hyper-parameter optimisation, with several advantages including efficiency, flexibility, and robustness. Extensive research in the literature has demonstrated that this method outperforms grid search in numerous cases (Bergstra and Bengio, 2012). For each set of hyper-parameters, we trained the model using the mean absolute error (MAE) as the loss function, and while using ADAM as the optimisation algorithm (Kingma and Ba, 2014). The model then was validated by forecasting the graph 3 years in advance, and the average performance was recorded. Once the set of hyper-parameters with the minimum error was found, we assessed the model's performance on the testing set and recorded the corresponding error. As a last step, we employed the optimal hyper-parameter settings to train the model using the entire dataset, followed by generating forecasts for the forthcoming three years, extending up to December 2025.

The first group of hyper-parameters includes the learning rate with values that range from $1 \times 10^{-4}$ to $1 \times 10^{-2}$, the number of epochs with values up to 200, the number of layers in the range 1 to 2, and the dropout value between 0.2 and 0.7. Other hyper-parameters are specific to the graph neural network including the graph convolution depth in the range 1 to 3, the convolution channels in the range 4 to 16, the activation function controller $\alpha$ (see Eq. (4)) within the range of 0.05 to 9, and the information propagation controller $\beta$ (see Eq. (6)) ranging from 0.05 to 0.8. The range of these values was obtained from the literature and online code repositories (Wu et al., 2020; Kim et al., 2022).

### 3.7. Model evaluation

In the evaluation phase (validation and testing), we used two evaluation metrics namely the Root Relative Squared Error (RSE) and the Relative Absolute Error (RAE) (Lai et al., 2018). These metrics compute the model's error relative to the error of a simple model that can predict the average trend of each node. Formally, let $Y_{j,m}$ denote the actual value in the test set of node $j$ during month $m$, and $\hat{Y}_{j,m}$ denote the predicted value, where $\mathbf{Y}, \hat{\mathbf{Y}} \in \mathbb{R}^{N\times T_{\text{test}}}$. Then, RSE and RAE are given by the following formulas,

$$RSE = \frac{\sqrt{\sum_{(j,m)\in\Omega_{Test}} (Y_{j,m} - \hat{Y}_{j,m})^2}}{\sqrt{\sum_{(j,m)\in\Omega_{Test}} (Y_{j,m} - \text{mean}(\mathbf{Y}_j))^2}} \tag{10}$$

$$RAE = \frac{\sum_{(j,m)\in\Omega_{Test}} |Y_{j,m} - \hat{Y}_{j,m}|}{\sum_{(j,m)\in\Omega_{Test}} |Y_{j,m} - \text{mean}(\mathbf{Y}_j)|} \tag{11}$$

These metrics provide readable evaluation, regardless the scale of the data. For both metrics, the lower value is better.

The model validation results are provided in Fig. 3. As shown in the figure, the predicted data points are aligned with the ground truth, and the model is able to capture the time series patterns effectively. For some nodes (*e.g.*, NLP/LLM), we notice a slight increase in the confidence interval as we move towards the later years, suggesting less certainty about the prediction in those years. This increase in the uncertainty can be reduced with more knowledge in terms of new features or more samples (Almahmoud et al., 2023). Overall, in terms of validation error, the average RSE computed over 142 nodes is 0.52, and the average RAE is 0.66, which provides a noticeable improvement over the benchmark model.

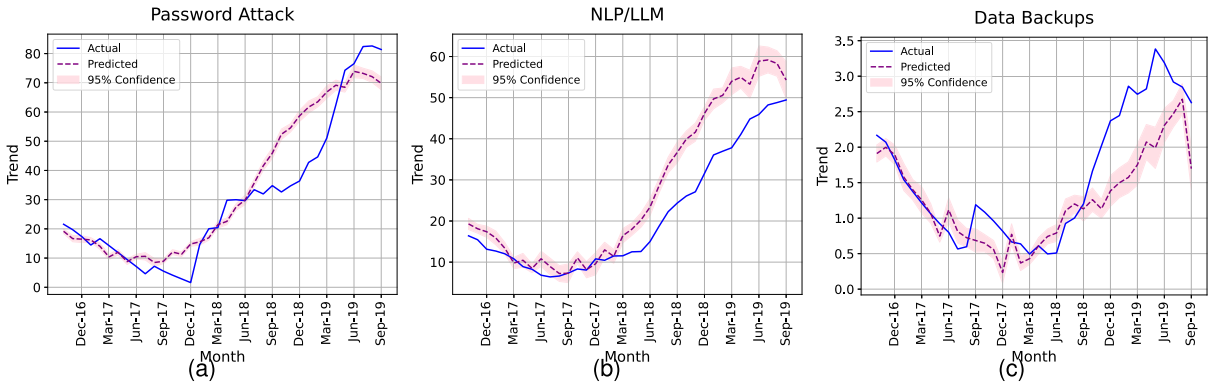


**Fig. 3.** The B-MTGNN validation results of predicting threats and PATs from October, 2016 to September, 2019. (a) Password Attack with RAE = 0.37. (b) NLP/LLM with RAE = 0.53. (c) Data Backups with RAE = 0.51. The 95% confidence interval of the predicted distribution using the Bayesian approach is shown in pink colour. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

## 4. Results

### 4.1. Trend forecast

The forecast of the cyber threats and their PATs in the coming 3 years is provided in Fig. 4. Here, we focus on the most important threats for which there will be a significant gap in the future with the respective PATs based on the forecast, while including threats from both categories (the rapidly increasing and emerging threats). We also focus on the PATs that will likely have a positive gap with the relevant threat. In other words, the PATs shown are those for which the trend was forecasted to be below the trend of the relevant threat. In Fig. 4, the gap area is visually represented using the same colour as the corresponding PAT curve.

The malware attack stands out for having the most significant gaps with respect to its PATs compared to other types of attacks. The forecast illustrated in Fig. 4(a) indicates a considerable disparity reaching a value of 0.8, and expected to persist over the next three years between malware and various PATs, including Application Whitelisting, File Integrity Monitoring, and Darknet Monitoring. Other PATs such as Blockchain, Anomaly Detection, and ML/DL are also expected to trail behind malware. However, the gaps of these PATs with respect to malware are comparatively smaller, narrowing clearly in the case of Blockchain, thanks to the recent growing body of research in these fields (Kosmarski, 2020; Shaukat et al., 2020; Dwivedi et al., 2023).

The next concern is the vulnerability related attacks shown in Fig. 4(b). Here, we observe a consistently widening gap with some PATs including Standardised Communication, Security Information and Event Management (SIEM), and Control Flow Integrity. Compared to these PATs, Vulnerability Assessment and NLP/LLM are expected to be more visible, even though the anticipated gaps are still large, exceeding a value of 0.2.

Concerning the more recently emerging threats (Figs. 4(c) and 4(d)), ransomware will likely exhibit gap values above 0.1 with respect to several PATs including Application Whitelisting, Deception Technology, and Data Backups, while having relatively smaller gaps with Access Control and Anomaly Detection (below 0.05). The adversarial attack is expected to have the largest gap of 0.09 with respect to Spatial Smoothing, Defensive Distillation, and Noise Injection, and the smallest gap with respect to NLP/LLM (around 0.05).

### 4.2. Trend categories

Our analysis for the future gaps between the threats and PATs allowed us to categorise the gap trend into four main categories, as shown in Tables 5 and 6. In these tables, PATs are listed in descending order of the gap, while considering different types of threats and threat categories. Here, we computed the average gap in each year and recorded the result for each of the three years (2023 to 2025).

**Table 5**
Widening gaps.

| Strictly widening gaps | | | | | |
|---|---|---|---|---|---|
| Threat | PAT | Gap forecast | | | GD |
| | | 2023 | 2024 | 2025 | |
| Vulnerability | SC | 0.202 | 0.218 | **0.244** | ↑ ↑ |
| Vulnerability | SIEM | 0.201 | 0.217 | **0.241** | ↑ ↑ |
| Vulnerability | CFI | 0.200 | 0.216 | **0.241** | ↑ ↑ |
| Account Hijacking | LP | 0.186 | 0.199 | **0.229** | ↑ ↑ |
| Account Hijacking | SM | 0.186 | 0.199 | **0.229** | ↑ ↑ |
| Account Hijacking | MFA | 0.182 | 0.195 | **0.226** | ↑ ↑ |
| Ransomware | AW | 0.146 | 0.149 | **0.170** | ↑ ↑ |
| Ransomware | DT | 0.146 | 0.149 | **0.169** | ↑ ↑ |
| Ransomware | DB | 0.146 | 0.148 | **0.169** | ↑ ↑ |
| IoT Device Attack | MS | 0.043 | 0.050 | **0.055** | ↑ ↑ |
| IoT Device Attack | SB | 0.043 | 0.049 | **0.054** | ↑ ↑ |
| IoT Device Attack | MFA | 0.039 | 0.046 | **0.052** | ↑ ↑ |
| **Overall widening gaps** | | | | | |
| Threat | PAT | Gap forecast | | | GD |
| | | 2023 | 2024 | 2025 | |
| Malware | AW | 0.766 | 0.763 | **0.837** | ↓ ↑ |
| Malware | FIM | 0.766 | 0.763 | **0.836** | ↓ ↑ |
| Malware | DM | 0.766 | 0.763 | **0.836** | ↓ ↑ |
| Ransomware | NLP/LLM | 0.116 | 0.114 | **0.131** | ↓ ↑ |
| Adversarial Attack | SS | 0.080 | 0.079 | **0.088** | ↓ ↑ |
| Adversarial Attack | DD | 0.080 | 0.079 | **0.088** | ↓ ↑ |
| Adversarial Attack | NI | 0.079 | 0.078 | **0.087** | ↓ ↑ |
| Account Hijacking | AC | 0.074 | 0.068 | **0.086** | ↓ ↑ |
| Phishing | AC | 0.062 | 0.049 | **0.068** | ↓ ↑ |
| Ransomware | AD | 0.049 | 0.046 | **0.051** | ↓ ↑ |
| Deepfake | 3DFR | 0.047 | 0.046 | **0.051** | ↓ ↑ |
| Deepfake | DW | 0.046 | 0.045 | **0.049** | ↓ ↑ |

Items are displayed in descending order of the gap. GD refers to the Gap Directions. Please refer to Fig. 5 for the PAT abbreviations.

The first category is the Strictly Widening Gaps (SWG) shown in the first half of Table 5. These are the gaps that are predicted to be consistently increasing between the years 2023 and 2025. Examples of such gaps include the gaps between the vulnerability related threats and each of Standardised Communication (SC), SIEM, and Control Flow Integrity (CFI). Similarly, the gaps for IoT Device Attack with respect to Merkle Signature (MS), Secure Boot (SB), and Multi-Factor Authentication (MFA) are consistently widening, even though they exhibit smaller values.

The second category is the Overall Widening Gaps (OWG) shown in the second half of Table 5. These gaps are anticipated to increase in the year 2025 compared to 2023, with expected fluctuations in between. Among the top in the list are the gaps between malware

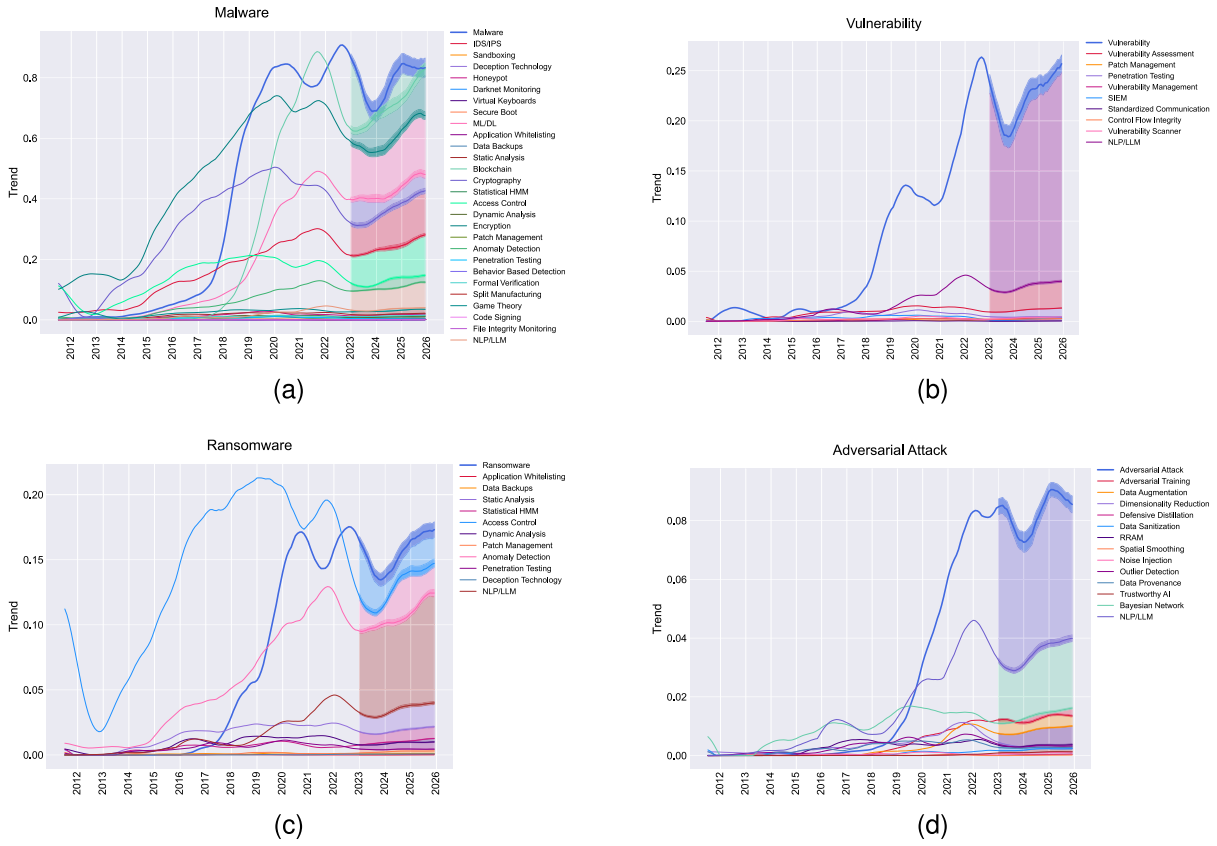


**Fig. 4.** The forecast of the trend for four threats and their PATs. The period of the trend plots is between July, 2011 and December, 2025, with the period between January, 2023 and December, 2025 forecasted using B-MTGNN. The shown PATs are those for which the trend is predicted to be lower than the trend of the corresponding threat. The gaps are highlighted in the same colour as the corresponding PAT curve. (a) Malware (b) Vulnerability (c) Ransomware (d) Adversarial Attack. The curves are smoothed using exponential smoothing with $\alpha = 0.1$ to reduce the noise and capture the trend. The 95% confidence interval is shown for each trend prediction. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

**Table 6**
Narrowing gaps.

| Overall narrowing gaps | | | | | |
|---|---|---|---|---|---|
| Threat | PAT | Gap forecast | | | GD |
| | | 2023 | 2024 | 2025 | |
| Malware | CR | **0.449** | 0.401 | 0.429 | ↓ ↑ |
| Ransomware | AC | **0.033** | 0.018 | 0.027 | ↓ ↑ |
| Deepfake | NLP/LLM | **0.017** | 0.011 | 0.012 | ↓ ↑ |
| MITM | SSP | **0.014** | 0.013 | 0.013 | ↓ → |
| MITM | PT | **0.010** | 0.009 | 0.009 | ↓ → |
| MITM | VPN | **0.007** | 0.006 | 0.006 | ↓ → |
| APT | UBA | $\mathbf{14.2 \times 10^{-4}}$ | $7.1 \times 10^{-4}$ | $9.5 \times 10^{-4}$ | ↓ ↑ |
| APT | NS | $\mathbf{12.9 \times 10^{-4}}$ | $6.5 \times 10^{-4}$ | $7.8 \times 10^{-4}$ | ↓ ↑ |
| APT | DLP | $\mathbf{11.2 \times 10^{-4}}$ | $5 \times 10^{-4}$ | $8.2 \times 10^{-4}$ | ↓ ↑ |
| Disinformation | CA | $\mathbf{7.2 \times 10^{-4}}$ | $4.1 \times 10^{-4}$ | $5.5 \times 10^{-4}$ | ↓ ↑ |
| APT | DT | $\mathbf{8.1 \times 10^{-4}}$ | $2 \times 10^{-4}$ | $4.1 \times 10^{-4}$ | ↓ ↑ |
| APT | LP | $\mathbf{7.8 \times 10^{-4}}$ | $2.2 \times 10^{-4}$ | $3.9 \times 10^{-4}$ | ↓ ↑ |
| **Strictly narrowing gaps** | | | | | |
| Threat | PAT | Gap forecast | | | GD |
| | | 2023 | 2024 | 2025 | |
| Malware | EN | **0.199** | 0.184 | 0.174 | ↓ ↓ |
| Malware | BC | **0.129** | 0.064 | 0.057 | ↓ ↓ |
| MITM | PKI | **0.006** | 0.004 | 0.003 | ↓ ↓ |

and Application Whitelisting (AW), File Integrity Monitoring (FIM), and Darknet Monitoring (DM). Other examples with smaller gap values include the gaps between deepfake and each of 3 Dimensional Face Reconstruction (3DFR) and Digital Watermark (DW).

Third is the Overall Narrowing Gaps (ONG) illustrated in the upper part of Table 6. These gaps are likely to decrease in the year 2025 compared to their values in 2023, despite the expected fluctuations in between. Among the top in the list are the gaps between malware and Cryptography (CR) and between ransomware and Access Control (AC). Examples with much smaller gap values include the gaps between Advanced Persistent Threat (APT) and Deception Technology (DT), as well as between APT and Least Privilege (LP).

The fourth and last category is the Strictly Narrowing Gaps (SNG). As shown in the lower part of Table 6, these gaps are consistently decreasing between the years 2023 and 2025. Examples include the gaps between malware and each of Encryption (EN) and Blockchain (BC), and the gap between Man-In-The-Middle attack (MITM) and Public Key Infrastructure (PKI). It is worth noting that this category comprises the fewest items, indicating the rarity of these gaps.

### 4.3. Recommendations

Our recommendations include the investment in the research and development of the PATs with widening gaps with respect to the relevant threats, which are listed in Table 5. These PATs can be prioritised in the order of the table so that the PATs with wider gaps are given higher priority. Similarly, the PATs in the SWG group should receive higher attention compared to the PATs in the OWG group, since they are more likely to persist this widening trend. It follows that the investment in Standardised Communication (SC), SIEM, Control Flow Integrity (CFI), Least Privilege (LP), and Session Management (SM) is highly recommended (top five PATs in the SWG group). At the same time, it is also important to consider the significant gap values observed

in the OWG group, hence to invest in Application Whitelisting (AW), File Integrity Monitoring (FIM), Darknet Monitoring (DM), NLP/LLM, and Spatial Smoothing (SS). We note that the decision to invest in the top five technologies in each category is only an example. Policymakers may adjust this number according to their capacity and resources.

On the other hand, it is recommended that the PATs in the ONG and SNG groups be given less priority when making an investment decision, especially if they did not appear in the SWG or OWG groups. Here, less priority can be given to the PATs with smaller gap values and PATs with gaps that are consistently narrowing (SNG). Examples include Encryption (EN), Blockchain (BC), and Public Key Infrastructure (PKI). While these PATs play an important role in cyber security, the forecast suggests that they are catching up with the trend of relevant threats and it is time to consider additional technologies to effectively combat evolving cyber threats.

### 4.4. Alleviation technologies cycle

Our large scale analysis for the PATs' historical data and future predictions spanning three years facilitated the development of a generalisable model that provides a comprehensive understanding of the progression of these PATs as they transition through 5 phases, namely the launch, growth, maturity, trough, and stability. This model is referred to as the Alleviation Technologies Cycle (or ATC), which is depicted in Fig. 5. During the launch phase, a new technology emerges and is adopted by few agencies for a brief period. Subsequently, there is a rapid surge in both the frequency and prominence of the technology as more security agencies become acquainted with and adopt the new PAT. Typically, PATs exhibit numerous variations in terms of speed of progression. For most of the PATs, we observe a slow progression during the growth phase compared to other types of technologies. This is due to the presence of various challenges in the world of cyber security including the resistance of attackers to the new security solution (Reddy and Reddy, 2014). As the visibility reaches its peak, the PAT enters the maturity phase, characterised by a sustained and stable pattern for a short period of time. This is followed by a temporary decline into the trough where enthusiasm diminishes as trials and executions fall short of expectations. Based on the forecast, we identified two possible troughs that the PAT can reach. One of these troughs is deeper than the other, depending on the usability of the PAT and the demand for it. Eventually, the PAT recovers and moves to either a higher or lower plateau, depending on which trough it originated from. This recovery takes place as additional examples showcasing how the technology can advantage the organisation begin to solidify and gain broader comprehension. Within the plateaus, mainstream adoption accelerates as the criteria for evaluating viability become more distinct, showcasing the technology's widespread market utility and effectiveness (Dedehayir and Steinert, 2016).

As depicted in Fig. 5, the positioning of the PATs on the cycle is determined by analysing their current trend slope, their historical patterns, and their future projections. During the trough phase, PATs exhibit either a trajectory towards the upper plateau or the lower plateau. By leveraging the predicted trends, illustrated in Fig. 4, we were able to indicate the future destination for some PATs near the trough using distinct colours (blue or purple). For instance, Distributed Ledgers Technology (DLT), Resistive Random-Access Memory (RRAM), and Virtual Private Network (VPN) are displayed in blue colour, indicating their likelihood of transitioning towards the upper plateau. It is important to note that during the initial three phases, the ultimate destination of a particular PAT, whether it will reach the upper or lower plateau, often remains uncertain and challenging to predict, thus denoted in grey. In addition, we distinguish each PAT by employing distinct shapes, indicating their relevance to either rapidly increasing or emerging threats (or possibly to both categories).

The ATC is similar to the well-known Gartner Hype Cycle (GHC) (Dedehayir and Steinert, 2016), with some important differences. The ATC has a slower rate of growth compared to GHC given that it is specific to the challenging field of cyber security, as previously mentioned. Another notable distinction is the presence of two distinct troughs (and two plateaus) in the ATC instead of a single trough observed in GHC. This difference arises because the ATC is a specialised variant of GHC designed specifically for the cyber security domain.

In the early stage of the growth phase, PATs are mostly related to the emerging threats, as can be observed in Fig. 5. These PATs include Defensive Distillation (DD), Deception Technology (DT), Trustworthy AI (TAI), and Adversarial Training (AdT). In the later stages of the growth phase, different types of PATs can be observed including those relevant to the rapidly increasing threats. Examples include NLP/LLM, Split Manufacturing (SMF), Certificate Pinning (CP), and Continuous Authentication (CA). After the peak, and into the upper trough, we find a combination of PATs (relevant to threats from different categories) sliding down, including Distributed Ledgers Technology (DLT), Control Flow Integrity (CFI), Static Analysis (SA), Dynamic Analysis (DAS), and Data Augmentation (DA). On the upper plateau, most of the PATs are relevant to the rapidly increasing threats including Session Management (SM), Rate Limiting (RL), Activity Monitoring (AM), Rank Correlation (RC), and Password Policy (PP). Many PATs are falling into the lower trough, and those are mostly relevant to the rapidly increasing threats. They include Supply Chain Risk Management (SCRM), One Time Password (OTP), Domain Name System Security Extensions (DNSSEC), and File Integrity Monitoring (FIM). On the lower plateau, most of the PATs are relevant to the rapidly increasing threats. These include Password Management (PM), Code Signing (CS), Data Loss Prevention (DLP), Identity-based Encryption (IBE), and Behaviour-based Detection (BBD).

## 5. Comparative analysis

### 5.1. Ablation study

In this section, we show experimentally the effect of our proposed external features (NoM, ACA, and PH) on the performance of the MTGNN model. In addition, we demonstrate the effectiveness of the graph convolution layers and the graph learning layer. To this end, we conducted multiple experiments to evaluate the performance of eight different variations of the MTGNN model in predicting the trends up to 3 years in advance, while using unseen data. For each model variant, we split the dataset into 70% training/validation and 30% testing. Each model undergoes random search with 60 iterations to optimise the set of hyper-parameters, and the final testing errors RSE and RAE are averaged over 10 experiments.

The first four models (Table 7) do not utilise our external features during the prediction and rather rely on the ground truth. The first model does not include any graph convolution layer and only performs temporal convolution. In the next three variations, we experimented with models that utilise graph convolution layers including two models that use a predefined adjacency matrix (uni-directional and bi-directional variants), and one model that uses the adaptively learned adjacency matrix through the graph learning layer. Intuitively, within the uni-directional predefined adjacency matrix, the threat node points to the relevant PAT node, since the threat often precedes the security measure. In the case of bi-directional adjacency matrix, both types of nodes point to each other. In the case of adaptive learning, we allow the model to learn these relationships. We note that in the case of predefined adjacency matrix, the edge weight is set to 1 or 0 (depending on whether two nodes are connected), since it is challenging to identify the level of relationship, which can be rather learned adaptively. This relationship weight is only used during model training, and not to be confused with the edge weight in the original graph, which represents the gap (Eq. (1)).

The rest of four models utilise the external features with the following variations. The first model does not include any graph convolution layer and only performs temporal convolution. The second

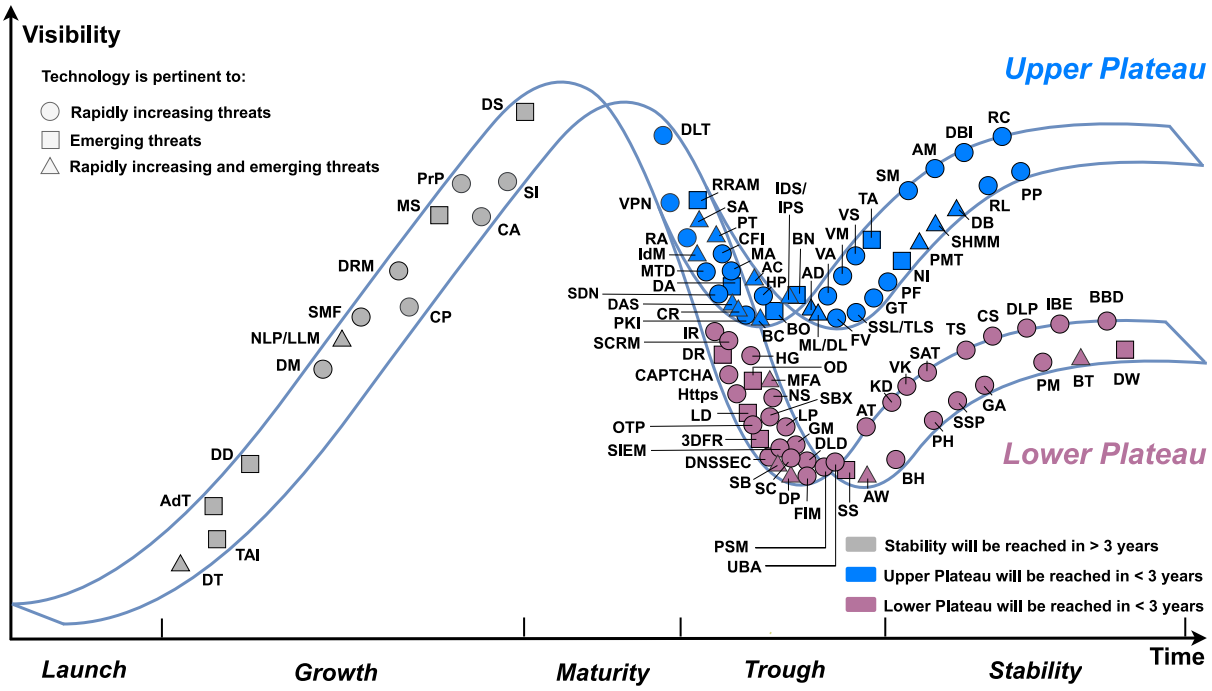


| A – C | | D – I | | K – R | | S – V | |
|---|---|---|---|---|---|---|---|
| 3DFR | 3D Face Reconstruction | DA | Data Augmentation | KD | Keystroke Dynamics | SA | Static Analysis |
| AC | Access Control | DAS | Dynamic Analysis | LD | Liveness detection | SAT | Strong Authentication |
| AD | Anomaly Detection | DB | Data Backups | NLP/LLM | Natural Language Processing/ Large Language Models | SB | Secure Boot |
| AM | Activity Monitoring | DBI | Dynamic Binary Instrumentation | LP | Least Privilege | SBX | Sandboxing |
| AT | Attack Tree | DD | Defensive Distillation | MA | Mutual Authentication | SC | Standardised Communication |
| AW | Application Whitelisting | DLD | Data Leakage Detection/Prevention | MFA | Multi-factor Authentication | SCRM | Supply Chain Risk Management |
| AdT | Adversarial Training | DLP | Data Loss Prevention | ML/DL | Machine Learning/Deep Learning | SDN | Software Defined Network |
| BBD | Behaviour-based Detection | DLT | Distributed Ledgers Technology | MS | Merkle Signature | SHMM | Statistical Hidden Markov Model |
| BC | Blockchain | DM | Darknet Monitoring | MTD | Moving Target Defence | SI | Source Identification |
| BH | Blackholing | DNSSEC | Domain Name System Security Extensions | NI | Noise Injection | SIEM | Security Information and Event Management |
| BN | Bayesian Network | DP | Data Provenance | NS | Network Segmentation | SM | Session Management |
| BO | Biometrics | DR | Dimensionality Reduction | OD | Outlier detection | SMF | Split Manufacturing |
| BT | Blacklisting | DRM | Dynamic Resource Management | OTP | One Time Password | SS | Spatial Smoothing |
| CA | Continuous Authentication | DS | Data Sanitisation | PF | Packet Filtering | SSL/TLS | Secure Sockets Layer/Transport Layer Security |
| CAPTCHA | Public Turing Comp Human | DT | Deception Technology | PH | Password Hashing | SSP | Secure Simple Pairing |
| CFI | Control Flow Integrity | DW | Digital Watermark | PKI | Public Key Infrastructure | TA | Taint Analysis |
| CP | Certificate Pinning | FIM | File Integrity Monitoring | PM | Password Management | TAI | Trustworthy AI |
| CR | Cryptography/Encryption | FV | Formal Verification | PMT | Patch Management | TS | Traffic Shaping |
| CS | Code Signing | GA | Graphical Authentication | PP | Password Policy | UBA | User Behaviour Analytics |
| | | GM | Graphical Model | PSM | Password Strength Meters | VA | Vulnerability Assessment |
| | | GT | Game Theory | PT | Penetration Testing | VK | Virtual Keyboards |
| | | HG | Hypergame | PrP | Privacy Preserving | VM | Vulnerability Management |
| | | HP | Honeypot | RA | Risk Assessment | VPN | Virtual Private Network |
| | | Https | Hypertext Transfer Protocol Secure | RC | Rank Correlation | VS | Vulnerability Scanner |
| | | IBE | Identity-based Encryption | RL | Rate Limiting | | |
| | | IDS/IPS | Intrusion Detection/Prevention System | RRAM | Resistive Random Access Memory | | |
| | | IR | Image Recognition | | | | |
| | | IdM | Identity Management | | | | |

**Fig. 5.** The Alleviation Technologies Cycle (ATC). The PATs go through 5 stages, namely, launch, growth, maturity, trough, and stability. ATC captures the state of each PAT in 2023, where the colour of the PAT indicates which slope it would follow based on the model prediction until 2025 (*e.g.*, blue: upper plateau or purple: lower plateau). The PATs with unknown final destination are coloured in grey. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and third models utilise graph convolution layers with a predefined adjacency matrix (uni-directional and bi-directional variants). In the uni-directional variant, the feature node, such as ACA points to the threat nodes (*e.g.*, wars and conflicts precede the attack), and the threat node points to the relevant PAT node. The fourth model employs the graph learning layer along with the graph convolution layers to adaptively learn the relationships in the graph.

The evaluation results are presented in Table 7. The use of the external features made a significant difference, reducing the relative error to a value below 1, which provides an improvement over the simple model. The results also show that using graph convolution leads to a lower error compared to relying solely on the temporal convolution. In addition, we observe that the use of uni-directional predefined adjacency matrix consistently resulted in a better performance compared to the use of bi-directional variant. This is consistent with the findings in Wu et al. (2020). However, the use of graph learning layer to learn the adjacency matrix resulted in a better performance than using any predefined adjacency matrix. This is explained by the fact that the graph structure is not optimal and should be updated during training (Wu et al., 2020). Overall, the best performance was obtained when combining the graph convolution layers (in addition to the temporal convolution layers), the graph learning layer, and the external features. This justifies the use of these layers along with our proposed features in our future forecast.

### 5.2. Comparative evaluation

#### 5.2.1. MTGNN

We conducted a comprehensive comparative evaluation to assess the performance of MTGNN against four established baseline models. These are AutoRegressive Integrated Moving Average (ARIMA), Vector AutoRegression (VAR), Long Short-Term Memory (LSTM), and Transformer Encoder-Decoder. Both ARIMA and VAR are statistical models commonly used for time series analysis and forecasting (Thomakos and Guerard, 2004). However, ARIMA is a univariate model, while VAR operates in a multivariate context. LSTM and Transformer are ML models

**Table 7**
Comparative evaluation for 8 variations of MTGNN.

| Model | RSE | RAE |
|---|---|---|
| TCN | 3.75 | 3.31 |
| TCN, GCN with predefined adj. matrix (bi-directional) | 3.25 | 2.98 |
| TCN, GCN with predefined adj. matrix (uni-directional) | 3.25 | 2.97 |
| TCN, GCN with adaptively learned adj. matrix | 3.20 | 2.89 |
| TCN, external features | 0.83 | 0.93 |
| TCN, GCN with predefined adj. matrix (bi-directional), external features | 0.76 | 0.88 |
| TCN, GCN with predefined adj. matrix (uni-directional), external features | 0.75 | 0.88 |
| **TCN, GCN with adaptively learned adj. matrix, external features** | **0.73** | **0.85** |

TCN stands for temporal convolution and GCN stands for graph convolution.

commonly used for sequence-to-sequence prediction (Sutskever et al., 2014; Song et al., 2021), and were evaluated both as univariate and multivariate models. In contrast, MTGNN inherently operates as a multivariate model, leveraging its capacity to capture spatial relationships among all features and adaptively learn their hidden relationships.

For each of the four baseline models, we trained separate models to predict each feature in the dataset. This method aimed to facilitate easier learning and convergence by reducing dimensionality. Univariate models relied solely on the ground truth data for prediction (the single feature at hand), while multivariate models integrated additional features. Here, we employed a domain-driven feature selection approach, leveraging prior knowledge and assumptions to determine which features to include in addition to the ground truth. For instance, in models predicting NoI, additional features included external factors (NoM, ACA, PH) alongside pertinent technologies (PATs). Conversely, models predicting PATs incorporated relevant attack types as additional features. In our experiment, each model underwent standardised data partitioning, with approximately 70% allocated for training/validation and 30% for testing. Model performance was assessed on the testing set (unseen data). Random search with 30 iterations was employed to optimise hyper-parameters for each model, and the final performance was averaged over 5 experiments.

Analysis of the results, as depicted in Table 8, reveals MTGNN as the top performer in terms of both RSE and RAE. With an RSE of 0.77 and RAE of 0.83, MTGNN demonstrates superior forecasting accuracy compared to ARIMA, VAR, LSTM, and Transformer models, across both univariate and multivariate settings. The notable performance enhancement of MTGNN can be primarily attributed to its ability to adaptively learn and capture intricate spatial relationships among features, while effectively leveraging external information. While ARIMA and VAR models display reasonable performance, LSTM and Transformer models exhibit comparatively higher errors, indicating challenges in capturing the underlying temporal dependencies. This underscores the advantage of incorporating graph-based adaptive learning mechanisms, particularly in MTGNN, for time series forecasting tasks. Moreover, the integration of external features further enhances MTGNN's predictive capabilities, underscoring its versatility and effectiveness in real-world forecasting scenarios.

A notable trend observed from the comparative evaluation is the consistent outperformance of univariate approaches over their multivariate counterparts across the four baseline models, as evident in Table 8. With the exception of the MTGNN model, multivariate models, including multivariate LSTM and Transformer, consistently exhibited higher RSE and RAE compared to their univariate counterparts, and the univariate model ARIMA outperformed the multivariate model VAR. This discrepancy is attributed to pre-assumed feature relationships that may not necessarily be optimal, underscoring the importance of learning these interdependencies among multiple variables for accurate

**Table 8**
Comparative evaluation for MTGNN and 4 baseline models.

| Model | RSE | RAE |
|---|---|---|
| LSTM (M) | 1.42 | 1.38 |
| Transformer Encoder-Decoder (M) | 1.40 | 1.39 |
| Transformer Encoder-Decoder (U) | 1.40 | 1.36 |
| LSTM (U) | 1.40 | 1.34 |
| VAR (M) | 1.20 | 1.32 |
| ARIMA (U) | 1.00 | 0.87 |
| **MTGNN (M)** | **0.77** | **0.83** |

U stands for univariate model and M stands for multivariate model.

time series forecasting. MTGNN explicitly learns and quantifies these relationships. Additionally, representing these features as nodes in a graph provides the opportunity to capture hierarchical relationships. Moreover, in cases where no such relationships exist between nodes, the graph convolution layer can adapt and preserve the original node's self-information (Wu et al., 2020).

#### 5.2.2. B-MTGNN

We additionally conducted a quantitative evaluation to justify the inclusion of the Bayesian module. Here, we evaluated the performance of the MTGNN model compared to five variations of the B-MTGNN model, where each variation uses a different number of iterations in the range 10–50 to approximate a Bayesian model. The number of iterations is denoted as $it$, where $it > 1$. Similar to our previous experiment, we divided the dataset into 70% for training/validation and 30% for testing. Additionally, we employed random search with 30 iterations to optimise the hyper-parameters of each model, and the final performance was averaged over 5 experimental runs.

The evaluation results are illustrated in Table 9. The results indicate that the inclusion of the Bayesian module significantly impacts the model's performance, particularly as the number of iterations increases. Specifically, the B-MTGNN model with 30 iterations ($it = 30$) outperforms all other models including the MTGNN model, achieving the lowest RSE of 0.67 and the lowest RAE of 0.78. This suggests that a higher number of iterations in the Bayesian approximation improves the model's accuracy and generalisation capability. However, it is also noteworthy that increasing the iterations beyond 30 does not yield further improvements, as observed with the B-MTGNN models having 40 and 50 iterations, where the performance slightly declines. This phenomenon highlights the presence of an optimal range for the number of iterations, beyond which the model's accuracy may not continue to increase and may even decrease due to factors such as computational inefficiencies or diminishing returns in model complexity. Therefore, the B-MTGNN model with 30 iterations stands out as the most effective configuration for balancing performance and computational cost, underscoring the value of Bayesian methods in enhancing predictive accuracy in complex models like MTGNN.

The superior performance of the Bayesian model (B-MTGNN) compared to its deterministic counterpart (MTGNN) can be attributed to its ability to aggregate predictions from multiple iterations. By taking the mean of the distribution as the prediction, the Bayesian model leverages the collective knowledge encoded in these iterations, resulting in a more comprehensive and stable forecast. This approach helps mitigate the effects of overfitting and variability, leading to improved generalisation ability and enhanced predictive accuracy. We note that the benefits of the Bayesian model are not limited to improved accuracy; it also provides a measure of uncertainty, offering confidence in its predictions. Overall, the Bayesian model's capacity to capture uncertainty information and its robust averaging mechanism enable it to outperform its deterministic counterpart in terms of both performance and reliability.

**Table 9**
Comparative evaluation for MTGNN and 5 variations of B-MTGNN.

| Model | RSE | RAE |
|---|---|---|
| MTGNN | 0.77 | 0.83 |
| B-MTGNN ($it = 10$) | 0.75 | 0.85 |
| B-MTGNN ($it = 20$) | 0.73 | 0.81 |
| **B-MTGNN ($it = 30$)** | **0.67** | **0.78** |
| B-MTGNN ($it = 40$) | 0.72 | 0.82 |
| B-MTGNN ($it = 50$) | 0.71 | 0.82 |

*it* stands for the number of iterations in the Bayesian model.

## 6. Discussion

### 6.1. Highlights and contributions

This work pioneers a proactive approach in cyber security using ML for long-term prediction of cyber threats and the PATs. It represents a step forward in the field of cyber security, aligning with the growing body of literature advocating for proactive defence strategies (Anticipating, 2015; Husák et al., 2018; Okutan et al., 2019). By proposing the long-term prediction of cyber threats and PATs, this research addresses a critical gap identified in prior studies (Almahmoud et al., 2023). The integration of advanced ML techniques, particularly Bayesian graph learning, builds upon existing literature on predictive modelling approaches from different domains such as traffic forecasting (Wu et al., 2020; Guo et al., 2019). Furthermore, the improved model's performance when using the proposed features echoes findings from prior research (Okutan et al., 2019; Munkhdorj and Yuji, 2017; Goyal et al., 2018), highlighting the crucial role of feature engineering in enhancing predictive models' performance.

The implications of this work on research include advancing the research on proactive cyber security. It sets a precedent for future research to explore and refine predictive models, incorporating evolving ML techniques to foresee cyber threats effectively as well as the relevant technologies. The demonstrated improved performance when using the Bayesian model indicates a potential shift towards employing advanced techniques in graph analytics. Future research may focus on optimising and customising graph-based algorithms for cyber threat prediction, thereby enhancing the accuracy and efficiency of predictive models. The proposed effective data features can be also utilised and extended to further improve the performance. Furthermore, by highlighting the use of extensive global data and coverage of 36 countries, this work underlines the importance of comprehensive data analysis for a more holistic understanding of the cyber threat landscape. Future research could explore further enhancements in data collection, analysis, and representation for an even broader international scope.

In practice, this work enhances cyber security preparedness and planning. The proactive approach advocated in this work emphasises the need for organisations to establish early-stage communication with potential cyber threats and the PATs. This suggests that real-world applications should invest in proactive planning, enabling them to develop optimal defensive measures well in advance. This optimality results from the reduced uncertainty which leads to the prioritisation of the security measures by considering future threat gaps. Furthermore, this shift towards automated, data-driven methodologies aims to minimise subjective biases. In practice, this implies a transition towards quantitatively-driven decisions, reducing reliance on human judgement. Organisations should consider integrating automated, data-centric approaches to ensure consistency and impartiality in threat analysis and decision-making processes. Finally, the noted improvement in performance using Bayesian GNN and the proposed features suggests that incorporating advanced ML techniques, especially those suited for graph-based data, can significantly enhance predictive capabilities. Organisations should explore and implement such techniques to improve the accuracy and efficacy of their cyber threat prediction systems.

### 6.2. Results analysis

The forecast analysis in Fig. 4 enabled us to identify technologies worthy of investment by visualising projected gaps between each threat and its PATs. However, we recognise that incorporating gap categorisation and tabulation (Tables 5 and 6) enhances this process by introducing a systematic approach to prioritise investments more effectively. This approach considers not only the magnitude of the gap but also its category, leading to more informed decision-making. It follows that categorising gaps into four distinct categories (SWG, OWG, ONG, and SNG) enables policymakers to prioritise investments in mitigation technologies more efficiently.

The ATC provides insights into the progression of PATs and their relevance to emerging and rapidly increasing threats. This understanding allows agencies to anticipate trends and prioritise resources accordingly. For example, during the growth phase, where PATs are often related to emerging threats, agencies can focus on early adoption and experimentation. As PATs mature and reach stability, agencies can assess their effectiveness and make informed decisions about long-term integration. Furthermore, the identification of trough phases in the ATC highlights potential challenges and areas for improvement in PAT deployment. Agencies can use this information to proactively address issues such as declining enthusiasm or performance gaps. By recognising these patterns, agencies can better navigate the complexities of cyber security technology adoption and ensure continuous improvement in their defence strategies.

Furthermore, the ATC presents policymakers with a comprehensive framework to strategically allocate resources and align defence mechanisms with the evolving landscape of cyber threats. For instance, we showed in our previous work that malware is currently peaking (Almahmoud et al., 2023), while in Fig. 5, the PAT *File Integrity Monitoring* (FIM) is situated in the lower trough. In response, policymakers should prioritise advancing File Integrity Monitoring to the plateau swiftly. This action would help bridge the gap between this technology and the evolving trend of malware, potentially facilitating a decline in malware incidents. Similarly, ransomware exhibits rapid growth, while Application Whitelisting (AW) is in the process of recovering from a trough phase. To address this gap, policymakers should focus on elevating the trend of Application Whitelisting to the plateau, thereby aligning its efficacy with the escalating trend of ransomware.

We advocate for policymakers to prioritise advancing the PATs towards the upper plateau rather than the lower plateau. PATs positioned on the upper plateau offer greater visibility and are better aligned with relevant threats, reducing the likelihood of significant gaps. Achieving this entails closely monitoring the trend of PATs and enhancing their usability as they enter the trough phase. By encouraging investment in these technologies during this phase, increased effort and experimentation can raise awareness and illustrate how the technology benefits organisations, facilitating a quicker recovery from the trough. This concerted effort ultimately propels the PATs towards the upper plateau, where they are better positioned to effectively address emerging cyber threats.

### 6.3. Findings in light of protection motivation theory

Our findings are well aligned with key constructs and principles of PMT, specifically threat appraisal and coping appraisal (Rogers, 1975). The long-term prediction of cyber threats and PATs directly relates to the threat and coping appraisal components of PMT, where individuals and organisations assess the severity and vulnerability associated with specific cyber threats and evaluate the effectiveness of cyber security measures. By utilising ML for predictive modelling, our research provides a data-driven version of PMT to evaluate the evolving landscape of cyber threats, thereby enhancing the accuracy of threat appraisal and enabling more informed decision-making. The coping appraisal aspect of PMT is addressed through forecasting the trend of alleviation

technologies and identifying their future disparity with cyber threats. By identifying and validating the efficacy of these coping strategies, our research underscores the importance of perceived response efficacy and response costs in shaping individuals' and organisations' protective behaviours.

Moreover, self-efficacy, a critical component of PMT, is reflected in our emphasis on automated, data-driven methodologies to reduce subjective biases and enhance predictive capabilities. By demonstrating the effectiveness of these advanced techniques, our research strengthens the confidence of organisations in their ability to implement and benefit from such approaches. This, in turn, can lead to increased motivation to adopt proactive measures and engage in continuous improvement of cyber defence strategies. Also, understanding the predicted disparity highlighted in this work will empower and motivate staff to take the initiative in developing and refining these technologies, presenting findings to the team, suggesting improvements, and mentoring junior team members. This increased confidence and proactive engagement are crucial for the successful implementation of relevant security measures.

In terms of advancing PMT knowledge, introducing a proactive dimension to PMT enriches the theoretical framework, providing a more comprehensive understanding of how protection motivation can be sustained and strengthened over time. This new dimension bridges the gap between immediate threat responses and long-term preparedness strategies. With a proactive PMT, interventions can be designed to not only address current risks but also prepare for future ones, ensuring that protective measures remain relevant and effective as new threats emerge. Policymakers can use insights from a proactive PMT to develop regulations and standards that encourage forward-looking security practices, enhancing overall societal resilience to cyber threats.

A proactive PMT encourages strategic investment in research and development of new technologies to address predicted future gaps. This not only ensures preparedness but also helps in lowering perceived response costs by planning and budgeting for necessary resources in advance. By assessing and addressing both current and future threats and coping strategies, individuals and organisations are better prepared and more confident in their ability to protect themselves. This enhanced confidence leads to a stronger and more sustained protection motivation, fostering long-term behavioural changes rather than reactive responses to immediate threats. This shift in mindset can lead to more robust and enduring protective behaviours.

Our results also suggest potential refinements to the PMT framework in the context of cyber security. For instance, the categorisation of threat gaps and the systematic prioritisation of investments based on these categories introduce a more structured approach to coping appraisal. This suggests that incorporating additional layers of analysis and decision-making criteria could enhance the applicability of PMT to complex and dynamic domains like cyber security. Moreover, the integration of extensive global data highlights the need for a broader perspective in threat appraisal, suggesting that PMT could benefit from incorporating more comprehensive data sources to enhance the accuracy and relevance of threat assessments. This aligns with the evolving nature of cyber threats, where global trends and patterns play a significant role in shaping local security landscapes.

Overall, our findings not only support the core constructs and principles of PMT but also offer insights into potential extensions and refinements of the framework. By linking our results to PMT, we provide an understanding of how psychological motivations can inform and enhance predictive models and algorithms for anticipating and mitigating cyber threats. This holistic approach is crucial for developing robust cyber security strategies that address both current and future risks based on a comprehensive understanding of threat perception and response mechanisms.

### 6.4. Limitations

One limitation of our approach is the absence of a real-time feedback mechanism in the proposed framework. Integrating this mechanism would enhance the framework based on PMT, as it provides continuous threat evaluation and assessment of the alleviation technologies. By incorporating real-time feedback from various data sources, such as news and social media, the system can dynamically adjust its threat perception and evaluation of the security measures. This continuous feedback assists organisations in updating their threat and PAT forecasting models and response strategies, thereby increasing perceived vulnerability and demonstrating the effectiveness of adaptive defence mechanisms. Staying informed of the latest and predicted threat trends and adjusting defence technologies accordingly ensures that organisations maintain motivation and confidence in the alleviation technologies to achieve robust security posture against evolving cyber threats.

With regards to the data, the current dataset provides valuable insights into high-level attack types and PATs, offering a foundational understanding of cyber security threats and mitigation strategies. The applicable scope of the dataset primarily includes strategic and tactical analysis for cyber security professionals, serving as a basis for developing broad-based defence mechanisms against a spectrum of cyber threats. It is particularly valuable for organisations seeking to establish a foundational cyber security posture by understanding prevalent threats and corresponding preventive technologies. However, as the cyber security landscape continues to evolve rapidly, there is a growing need to explore the possibility of extending the dataset to encompass more fine-grained attack types. This expansion presents an opportunity to delve deeper into specific attack vectors and enhance the effectiveness of cyber security defences.

For example, consider the category of "Malware" within the current dataset. While it provides a broad overview of malicious software threats, including viruses, worms, and ransomware, a more granular approach could distinguish between different variants and functionalities of malware. By categorising malware based on behaviour, propagation methods, and targeted platforms, organisations could tailor their defence mechanisms more precisely to combat specific threats. For instance, distinguishing between fileless malware (Sudhakar and Kumar, 2020), which operates solely in memory, and traditional file-based malware could inform strategies for endpoint detection and response.

Similarly, the category of "Adversarial Attack" highlights the diverse range of techniques employed by threat actors to subvert ML models and AI systems. However, a finer-grained classification could differentiate between adversarial attacks targeting image recognition systems, natural language processing models, and reinforcement learning algorithms. Indeed, the profoundly different nature of the training algorithms applicable to these categories of AI systems suggests differentiation among their adversarial attacks. Finer-grained attack classification would enable researchers and practitioners to develop specialised countermeasures, such as robustness enhancements (Tong et al., 2021), data augmentation techniques (Zeng et al., 2020), and adversarial training strategies (Zhang and Wang, 2019), tailored to each specific threat context.

Incorporating more fine-grained attack types into the dataset also opens avenues for exploring emerging threats and vulnerabilities. For instance, the rise of deepfake technology poses novel challenges in detecting and mitigating manipulated media content. By analysing different types of deepfake attacks, such as facial manipulation (Shao et al., 2022), voice synthesis (Bilika et al., 2023), and video impersonation (Gerstner and Farid, 2022), cyber security professionals can develop innovative detection algorithms and authentication mechanisms to combat the spread of disinformation and fraudulent content.

Moreover, extending the dataset to include fine-grained attack types facilitates cross-domain analysis and correlation studies. For example,

correlating specific malware families with targeted industries or geographical regions could reveal patterns of cybercriminal activity and inform proactive defence strategies (Mezzour et al., 2015). Similarly, identifying commonalities between adversarial attacks in different domains, such as image recognition and natural language processing, could lead to the development of holistic defence frameworks that address underlying vulnerabilities across diverse application areas.

Other limitations of this work include its reliance on a limited dataset that encompasses data since 2011 only. This is due to the challenges encountered in accessing confidential and sensitive information. Extending the prediction period necessitates the model to forecast further ahead into the future, requiring increased data samples and informative features. Also, a notable limitation stems from the lack of a systematic approach for the evaluation of the E-GPT algorithm, which is instrumental in extracting the PATs and constructing the graph. Moreover, such evaluation often depends on subjective and potentially biased human judgement. As a result, ensuring an optimal graph structure becomes challenging, particularly in the absence of a mechanism to quantify the assumed relationships between nodes in the graph. The subjectivity issue is also observed in the placement of PATs on the cycle, where a fully automated approach would lead to a more efficient process and more reliable results.

## 7. Conclusion and the road ahead

In this work, we introduced a proactive approach based on machine learning for long-term prediction of cyber threats and PATs. The goal is to establish an effective communication with the future disparity between the potential attacks and relevant security measures at an early stage, enabling proactive planning for the future. By adopting this approach, there is an increased chance to prevent incidents by allowing more time for the development of optimal defensive actions and tools, thereby bridging the gap between cyber threats and PATs. Moreover, our automated approach shows promise in addressing the widely recognised challenges associated with human-based analysis. By eliminating the reliance on human judgement and adopting a purely quantitative methodology driven by data, our approach aims to minimise subjective biases and promote consistency within the subject matter. With access to extensive data sources encompassing a vast volume of information and global geographic coverage, our study contributes to the construction of a comprehensive dataset encompassing different cyber security trends, which can be utilised for various purposes. We used this dataset to construct a novel Bayesian GNN model which was utilised to provide 3 years forecast for the future gaps between several cyber threats and PATs. Based on the future forecast, we categorised the gap trends, and recommended future investment decisions accordingly. Following a large-scale analysis for past and future trends, we proposed the alleviation technologies cycle (ATC) identifying the life cycle phases in the trend of 98 alleviation technologies. This cycle serves as a robust foundation for raising awareness when investing in security measures aimed at preventing cyber-attacks. It presents policymakers with a comprehensive framework to strategically allocate resources and align defence mechanisms with the evolving landscape of cyber threats. Furthermore, we have demonstrated the efficacy of our Bayesian model, outperforming several baseline models while also providing a measure of confidence through the articulation of epistemic uncertainty. Additionally, our incorporation of external features has demonstrated tangible improvements in model performance, further enhancing the reliability and utility of our predictive framework. Overall, our work not only advances the theoretical understanding of cyber threat prediction but also furnishes practical insights for managerial decision-making. By offering a proactive, data-driven approach to cyber security planning, we aspire to equip policymakers with the foresight and tools necessary to navigate an increasingly complex threat landscape with confidence and efficacy.

In future research, it is recommended to integrate real-time feedback mechanisms into the developed framework to enhance its effectiveness in terms of PMT. Additionally, while the ATC model provides a new dimension to PMT by highlighting the stages of technology adoption, future research could explore how this cycle influences both individual and organisational motivations to adopt new security measures. Specifically, methods such as surveys, longitudinal studies, and case analyses can investigate how each stage of the cycle affects the willingness to invest in and implement new PATs. This understanding can guide the development of strategies that enhance protection motivation and ensure a proactive approach to cyber security. Moreover, it is recommended to expand the dataset by including more fine-grained attack types, which enables more specific predictions and allows the development of tailored security measures. The dataset can be further expanded by incorporating more samples and informative features. This augmentation will lead to improved performance of the model and enable more accurate long-term trend forecasting. We additionally suggest the establishment of a systematic approach for the evaluation of the E-GPT algorithm, used for the extraction of the PATs (*i.e*, the graph construction). This can be done by interviewing and consulting security experts, in order to validate the algorithm's outputs, and possibly to contribute to the adjustment of the graph connectivity (semi-automated approach). This may require careful design of expert panels to tune inter-rater variance of evaluations. Alternatively, and perhaps more promisingly, a fully automated approach is feasible through the utilisation of the graph learning layer. This layer adaptively learns and quantifies the relationships between the threats and PATs. The obtained relationship values can then be leveraged, in conjunction with the predicted gaps, to prioritise the PATs effectively. Finally, the positioning of the PATs on the ATC can also be automated by computationally measuring the slope of the PAT curves and placing each PAT on the cycle accordingly. Overall, the automation of these phases maximises the machine's involvement in the process, thereby reducing the reliance on human bias and subjectivity.

## CRediT authorship contribution statement

**Zaid Almahmoud:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Paul D. Yoo:** Writing – review & editing, Visualization, Validation, Supervision, Project administration. **Ernesto Damiani:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Funding acquisition. **Kim-Kwang Raymond Choo:** Writing – review & editing, Supervision, Project administration, Investigation, Conceptualization. **Chan Yeob Yeun:** Writing – review & editing, Project administration.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

The authors extend their heartfelt gratitude to the DASA Machine Learning Team for their invaluable discussions and feedback. Special thanks are also extended to the EBTIC Cyber security Team at British Telecom (BT) for their constructive criticism of this work. Furthermore, the authors express their appreciation to O. Alhussein at Huawei, Ottawa, for his valuable contributions and unwavering support in this research.

## Data availability

The dataset and source code used in this work can be accessed at: https://github.com/zaidalmahmoud/Cyber-trend-forecasting.

## References


Adamov, A., Carlsson, A., 2017. The state of ransomware. Trends and mitigation techniques. In: EWDTS. pp. 1–8.

Adomavicius, G., Bockstedt, J., Gupta, A., Kauffman, R.J., 2008. Understanding evolution in technology ecosystems. Commun. ACM 51 (10), 117–122.

Almahmoud, Z., Yoo, P.D., Alhussein, O., Farhat, I., Damiani, E., 2023. A holistic and proactive approach to forecasting cyber threats. Sci. Rep. 13 (1), 8049.

Alshammari, A., Benson, V., Batista, L., 2024. The influences of employees' emotions on their cyber security protection motivation behaviour: A theoretical framework. In: 26th International Conference on Enterprise Information Systems.

2015. Anticipating cyber attacks: There's no abbottabad in cyber space. Infosecurity Magazine. https://www.infosecurity-magazine.com/white-papers/anticipating-cyber-attacks.

2023. Elsevier research products APIs. Elsevier Developer Portal. https://dev.elsevier.com.

Athanasopoulou, M.E., Deveikyte, J., Mosca, A., Peri, I., Provetti, A., 2021. A hybrid model for forecasting short-term electricity demand. In: Proceedings of the Second ACM International Conference on AI in Finance. pp. 1–6.

Bekkers, L., van't Hoff-De Goede, S., Misana-ter Huurne, E., van Houten, Y., Spithoven, R., Leukfeldt, E.R., 2023. Protecting your business against ransomware attacks? Explaining the motivations of entrepreneurs to take future protective measures against cybercrimes using an extended protection motivation theory model. Comput. Secur. 127, 103099.

Bergstra, J., Bengio, Y., 2012. Random search for hyper-parameter optimization. J. Mach. Learn. Res. 13 (2).

Bilge, L., Han, Y., Dell'Amico, M., 2017. Riskteller: Predicting the risk of cyber incidents. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 1299–1311.

Bilika, D., Michopoulou, N., Alepis, E., Patsakis, C., 2023. Hello me, meet the real me: Audio deepfake attacks on voice assistants. arXiv preprint arXiv:2302.10328.

Cao, D., Wang, Y., Duan, J., Zhang, C., Zhu, X., Huang, C., Tong, Y., Xu, B., Bai, J., Tong, J., et al., 2020. Spectral temporal graph neural network for multivariate time-series forecasting. Adv. Neural Inf. Process. Syst. 33, 17766–17778.

Cha, Y.-O., Hao, Y., 2022. The dawn of metamaterial engineering predicted via hyperdimensional keyword pool and memory learning. Adv. Opt. Mater. 10 (8), 2102444.

Chadha, A., Kumar, V., Kashyap, S., Gupta, M., 2021. Deepfake: an overview. In: Proceedings of Second International Conference on Computing, Communications, and Cyber-Security: IC4S 2020. Springer, pp. 557–566.

Chandra, R., Collis, S., 2021. Digital agriculture for small-scale producers: challenges and opportunities. Commun. ACM 64 (12), 75–84.

Dedehayir, O., Steinert, M., 2016. The hype cycle model: A review and future directions. Technol. Forecast. Soc. Change 108, 28–41.

Dodge, C.E., Fisk, N., Burruss, G.W., Moule, Jr., R.K., Jaynes, C.M., 2023. What motivates users to adopt cybersecurity practices? A survey experiment assessing protection motivation theory. Criminol. Public Policy 22 (4), 849–868.

Dwivedi, Y.K., Sharma, A., Rana, N.P., Giannakis, M., Goel, P., Dutot, V., 2023. Evolution of artificial intelligence research in technological forecasting and social change: Research topics, trends, and future directions. Technol. Forecast. Soc. Change 192, 122579.

Gal, Y., Ghahramani, Z., 2016. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. arXiv preprint arXiv:1506.02142v6.

2023. Gartner website. https://www.gartner.co.uk/en. (Accessed 17 October 2023).

Gaurav, A., Gupta, B.B., Panigrahi, P.K., 2022. A novel approach for ddos attacks detection in COVID-19 scenario for small entrepreneurs. Technol. Forecast. Soc. Change 177, 121554.

Gerstner, C.R., Farid, H., 2022. Detecting real-time deep-fake videos using active illumination. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 53–60.

Ghafur, S., Kristensen, S., Honeyford, K., Martin, G., Darzi, A., Aylin, P., 2019. A retrospective impact analysis of the WannaCry cyberattack on the NHS. NPJ Digit. Med. 2 (1), 1–7.

Goel, S., 2011. Cyberwarfare: connecting the dots in cyber intelligence. Commun. ACM 54 (8), 132–140.

Goyal, P., Hossain, K., Deb, A., Tavabi, N., Bartley, N., Abeliuk, A., Ferrara, E., Lerman, K., 2018. Discovering signals from web sources to predict cyber attacks. arXiv preprint arXiv:1806.03342.

2023. GPT-3 model. OpenAI Platform. https://platform.openai.com/docs/models/gpt-3.

GRAY, J., 2001. Futuristic forecast of tools and technologies. Commun. ACM 44 (3), 29.

Guo, S., Lin, Y., Feng, N., Song, C., Wan, H., 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 33, pp. 922–929.

2022. Holidays 0.15. PyPI · The Python Package Index. https://pypi.org/project/holidays/.

Husák, M., Bartoš, V., Sokol, P., Gajdoš, A., 2021. Predictive methods in cyber defense: Current experience and research challenges. Future Gener. Comput. Syst. 115, 517–530.

Husák, M., Kašpar, J., 2019. AIDA framework: real-time correlation and prediction of intrusion detection alerts. In: Proceedings of the 14th International Conference on Availability, Reliability and Security. pp. 1–8.

Husák, M., Komárková, J., Bou-Harb, E., Čeleda, P., 2018. Survey of attack projection, prediction, and forecasting in cyber security. IEEE Commun. Surv. Tutor. 21 (1), 640–660.

Kebir, O., Nouaouri, I., Rejeb, L., Said, L.B., 2022. ATiPreTA: AN analytical model for time–dependent prediction of terrorist attacks. Int. J. Appl. Math. Comput. Sci. 32 (3), 495–510.

Khan, N.F., Ikram, N., Murtaza, H., Javed, M., 2023. Evaluating protection motivation based cybersecurity awareness training on Kirkpatrick's model. Comput. Secur. 125, 103049.

Kim, J., Lee, G., Lee, S., Lee, C., 2022. Towards expert–machine collaborations for technology valuation: An interpretable machine learning approach. Technol. Forecast. Soc. Change 183, 121940.

Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Kosmarski, A., 2020. Blockchain adoption in academia: Promises and challenges. J. Open Innov.: Technol. Mark. Complex. 6 (4), 117.

Kuwahara, T., 1999. Technology forecasting activities in Japan. Technol. Forecast. Soc. Change 60 (1), 5–14.

Lai, G., Chang, W.-C., Yang, Y., Liu, H., 2018. Modeling long-and short-term temporal patterns with deep neural networks. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. pp. 95–104.

Lai, K.K., Yu, L., Wang, S., Huang, W., 2006. Hybridizing exponential smoothing and neural network for financial time series predication. In: International Conference on Computational Science. Springer, pp. 493–500.

Li, X., Xie, Q., Daim, T., Huang, L., 2019. Forecasting technology trends using text mining of the gaps between science and technology: The case of perovskite solar cell technology. Technol. Forecast. Soc. Change 146, 432–449.

Linkov, I., Ligo, A., Stoddard, K., Perez, B., Strelzoffx, A., Bellini, E., Kott, A., 2023. Cyber efficiency and cyber resilience. Commun. ACM 66 (4), 33–37.

Liu, Y., Sarabi, A., Zhang, J., Naghizadeh, P., Karir, M., Bailey, M., Liu, M., 2015. Cloudy with a chance of breach: Forecasting cyber security incidents. In: 24th USENIX Security Symposium. USENIX Security 15, pp. 1009–1024.

Loukaka, A., Rahman, S., 2017. Discovering new cyber protection approaches from a security professional prospective. Int. J. Comput. Netw. Commun. (IJCNC) 9.

Mae, Y., Kumagai, W., Kanamori, T., 2021. Uncertainty propagation for dropout-based Bayesian neural networks. Neural Netw. 144, 394–406.

Malik, J., Akhunzada, A., Bibi, I., Imran, M., Musaddiq, A., Kim, S.W., 2020. Hybrid deep learning: An efficient reconnaissance and surveillance detection mechanism in SDN. IEEE Access 8, 134695–134706.

Mezzour, G., Carley, K.M., Carley, L.R., 2015. An empirical study of global malware encounters. In: Proceedings of the 2015 Symposium and Bootcamp on the Science of Security. pp. 1–11.

Munkhdorj, B., Yuji, S., 2017. Cyber attack prediction using social data analysis. J. High Speed Netw. 23 (2), 109–135.

National Academies of Sciences, Engineering, M., et al., 2019. Robust Machine Learning Algorithms and Systems for Detection and Mitigation of Adversarial Attacks and Anomalies: Proceedings of a Workshop. National Academies Press.

Norman, P., Boer, H., Seydel, E.R., Mullan, B., 2015. Protection motivation theory. In: Predicting and Changing Health Behaviour: Research and Practice with Social Cognition Models. Vol. 3, Open University Press Maidenhead, pp. 70–106.

Oggier, F., Mihaljević, M.J., 2013. An information-theoretic security evaluation of a class of randomized encryption schemes. IEEE Trans. Inf. Forensics Secur. 9 (2), 158–168.

Okutan, A., Yang, S.J., McConky, K., Werner, G., 2019. Capture: cyberattack forecasting using non-stationary features with time lags. In: 2019 IEEE Conference on Communications and Network Security. CNS, IEEE, pp. 205–213.

Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K., 2016. Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499.

Passeri, P., 2022. Hackmageddon data set. Hackmageddon. https://www.hackmageddon.com.

Qin, X., Lee, W., 2004. Attack plan recognition and prediction using causal networks. In: 20th Annual Computer Security Applications Conference. IEEE, pp. 370–379.

Reddy, G.N., Reddy, G., 2014. A study of cyber security challenges and its emerging trends on latest technologies. arXiv preprint arXiv:1402.1842.

Rogers, R.W., 1975. A protection motivation theory of fear appeals and attitude change1. J. Psychol. 91 (1), 93–114.

Ruthig, J.C., 2016. Health risk perceptions and exercise in older adulthood: an application of protection motivation theory. J. Appl. Gerontol. 35 (9), 939–959.

Shao, R., Wu, T., Liu, Z., 2022. Detecting and recovering sequential deepfake manipulation. In: European Conference on Computer Vision. Springer, pp. 712–728.

Sharma, R., Thapa, S., 2023. Cybersecurity awareness, education, and behavioral change: strategies for promoting secure online practices among end users. Eigenpub Rev. Sci. Technol. 7 (1), 224–238.

Shaukat, K., Luo, S., Varadharajan, V., Hameed, I.A., Xu, M., 2020. A survey on machine learning techniques for cyber security in the last decade. IEEE Access 8, 222310–222354.

Shoufan, A., Damiani, E., 2017. On inter-rater reliability of information security experts. J. Inf. Secur. Appl. 37, 101–111.
Singh, M., Mehtre, B., Sangeetha, S., 2020. Insider threat detection based on user behaviour analysis. In: Machine Learning, Image Processing, Network Security and Data Sciences: Second International Conference, MIND 2020, Silchar, India, July 30-31, 2020, Proceedings, Part II 2. Springer, pp. 559–574.
Song, X., Wu, Y., Zhang, C., 2021. Tstnet: a sequence to sequence transformer network for spatial-temporal traffic prediction. In: Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14–17, 2021, Proceedings, Part I 30. Springer, pp. 343–354.
Stephens, G., 2008. Cybercrime in the year 2025. Futurist 42 (4), 32.
Sudhakar, Kumar, S., 2020. An emerging threat fileless malware: a survey and research challenges. Cybersecurity 3 (1), 1.
Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks. Adv. Neural Inf. Process. Syst. 27.
Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–9.
Taieb, S.B., Sorjamaa, A., Bontempi, G., 2010. Multiple-output modeling for multi-step-ahead time series forecasting. Neurocomputing 73 (10–12), 1950–1957.
Thomakos, D.D., Guerard, Jr., J.B., 2004. Naive, ARIMA, nonparametric, transfer function and VAR models: A comparison of forecasting performance. Int. J. Forecast. 20 (1), 53–67.
Tong, L., Chen, Z., Ni, J., Cheng, W., Song, D., Chen, H., Vorobeychik, Y., 2021. Facesec: A fine-grained robustness evaluation framework for face recognition systems. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13254–13263.
Tsai, H.-y.S., Jiang, M., Alhabash, S., LaRose, R., Rifon, N.J., Cotten, S.R., 2016. Understanding online safety behaviors: A protection motivation theory perspective. Comput. Secur. 59, 138–150.
2023. Twitter API v2. Developer Platform. https://developer.twitter.com/en/docs/twitter-api.
Vinayakumar, R., Soman, K., Poornachandran, P., 2017. Evaluation of recurrent neural network and its variants for intrusion detection system (IDS). Int. J. Inf. Syst. Model. Des. (IJISMD) 8 (3), 43–63.
Visser, M., van Eck, N.J., Waltman, L., 2021. Large-scale comparison of bibliographic data sources: Scopus, web of science, dimensions, crossref, and microsoft academic. Quant. Sci. Stud. 2 (1), 20–41.
Vrhovec, S., Mihelič, A., 2021. Redefining threat appraisals of organizational insiders and exploring the moderating role of fear in cyberattack protection motivation. Comput. Secur. 106, 102309.
Werner, G., Yang, S., McConky, K., 2017. Time series forecasting of cyber attack intensity. In: Proceedings of the 12th Annual Conference on Cyber and Information Security Research. pp. 1–3.
Werner, G., Yang, S., McConky, K., 2018. Leveraging intra-day temporal variations to predict daily cyberattack activity. In: 2018 IEEE International Conference on Intelligence and Security Informatics. ISI, IEEE, pp. 58–63.
Wu, Z., Pan, S., Long, G., Jiang, J., Chang, X., Zhang, C., 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 753–763.
Wueest, C., 2014. The continued rise of DDoS attacks. In: White Paper: Security Response, Symantec Corporation.
Yu, B., Yin, H., Zhu, Z., 2017. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. arXiv preprint arXiv:1709.04875.
Yuan, S., Wu, X., 2021. Deep learning for insider threat detection: Review, challenges and opportunities. Comput. Secur. 104, 102221.
Zeng, Y., Qiu, H., Memmi, G., Qiu, M., 2020. A data augmentation-based defense method against adversarial attacks in neural networks. In: Algorithms and Architectures for Parallel Processing: 20th International Conference, ICA3PP 2020, New York City, NY, USA, October 2–4, 2020, Proceedings, Part II 20. Springer, pp. 274–289.
Zhang, H., Wang, J., 2019. Defense against adversarial attacks using feature scattering-based adversarial training. Adv. Neural Inf. Process. Syst. 32.

**Zaid Almahmoud** received the bachelor's degree in software engineering from Khalifa University in 2013 and the master's degree in computing and information science from Masdar Institute of Science and Technology in 2016 through collaboration with MIT. He is currently a Ph.D. student in the University of London, Birkbeck College. His research interests include applied machine learning, forecasting, and time-series analysis. He is a Python, Java, and Android programmer, with several programming contests awards. He was a recipient of the Second Place in the Gulf Programming Contest, UAE, in 2013, and the First Place in Khalifa University Programming Contest, UAE, in 2011.

**Paul D. Yoo** (Senior Member, IEEE) is currently with the School of Computing and Mathematical Sciences, Birkbeck College, University of London, London, U.K. He held academic/research posts in Sydney, Cranfield, and the Korea Advanced Institute of Science and Technology (KAIST). He is the Editor of the IEEE Sustainable Computing and ACM Computing Surveys and holds more than 100 prestigious journal and conference publications in highly regarded IEEE/ACM journals. He is affiliated with the University of Sydney, Australia, and KAIST as a Visiting Professor. His research addresses two broad topics: advanced data analytics and technological forecasting.

**Ernesto Damiani** (Senior Member, IEEE) received the Honorary Doctorate degree from the Institute National des Sciences Appliquees de Lyon, France, 2017, for his contributions to research and teaching on big data analytics. He is currently a Full Professor with the Department of Computer Science, Universita degli Studi di Milano. He has published over 600 peer-reviewed articles and books. His research interests include cyber security, big data, and cloud/edge processing. He was a recipient of the 2017 Stephen Yau Award. He serves as an Editor-in-Chief for IEEE Transactions on Services Computing. He is a Distinguished Scientist of ACM.

**Kim-Kwang Raymond Choo** (Senior Member, IEEE) received the Ph.D. degree in information security from the Queensland University of Technology, Brisbane, QLD, Australia. He currently holds the Cloud Technology Endowed Professorship with The University of Texas at San Antonio, TX, USA. He is the Founding Co-Editor-in-Chief of ACM's Distributed Ledger Technologies: Research & Practice (from June 2021), Founding Chair of IEEE Technology and Engineering Management Society's Technical Committee (TC) on Blockchain and Distributed Ledger Technologies, and Department Editor of IEEE Transactions on Engineering Management, and Associate Editor for IEEE Transactions on Dependable and Secure Computing, and IEEE Transactions on Big Data.

**Chan Yeob Yeun** (Senior Member, IEEE) holds M.Sc. and Ph.D. degrees in information security from the University of London. Formerly Vice President at LG Electronics in Seoul (2005), he later joined ICU, South Korea, until August 2008, then Khalifa University, UAE, from September 2008. Presently, he serves as a cyber security researcher, Associate Professor in the Department of Electrical Engineering and Computer Science, and Cyber security Leader at the Center for Cyber-Physical Systems (C2PS). With over 140 research papers, nine book chapters, and ten international patent applications, Yeun also works on various international journal editorial boards and is on the steering committee of international conferences.