PriMera Scientific Engineering
Volume 7 Issue 2 August 2025
ISSN: 2834-2550

PriMera Scientific Publications

Integrating Artificial Intelligence with Cyber Threat Intelligence to Predict Financial Exposure in Data-Driven Enterprises

Type: Research Article
Received: July 08, 2025
Published: July 29, 2025

Citation:
Gloria Nwachukwu Ogochukwu., et al. "Integrating Artificial Intelligence with Cyber Threat Intelligence to Predict Financial Exposure in Data-Driven Enterprises". PriMera Scientific Engineering 7.2 (2025): 03-12.

Gloria Nwachukwu Ogochukwu[1]*, Ekufu Chukwuemeka Henry[3], Valentina Palama[4] and Edidiong Elijah Akpan[5]

[1]*University: Rensselaer Polytechnic Institute, Troy, New York (110 8th Street, Troy, NY 12180-3590), USA*

[3]*Department of Psychology, Olabisi Onabanjo University*

[4]*Prairie View A&M University*

[5]*School of Computing and Informatics, University of Louisiana at Lafayette, Louisiana, USA*

***Corresponding Author:** Gloria Nwachukwu Ogochukwu, University: Rensselaer Polytechnic Institute, Troy, New York (110 8th Street, Troy, NY 12180-3590), USA.

Abstract

As global reliance on digital infrastructure intensifies, the financial ramifications of cybersecurity incidents have become a critical concern for organizations worldwide. Quantifying and predicting these financial losses is essential for effective risk management, yet it remains a significant analytical challenge. This study addresses this problem by analyzing a decade of global cybersecurity threat data (2015-2024) to model the financial impact of cyber attacks, evaluate the predictive power of various machine learning models, and identify underlying incident archetypes. Employing a multi-stage methodology, the research began with an exploratory data analysis, followed by a comparative evaluation of four regression models: Linear Regression, Random Forest, Gradient Boosting, and Support Vector Regression. Feature importance was extracted from the models, and K-Means clustering was used to derive a data-driven taxonomy of incidents. The study reveals a persistent and costly threat landscape with no significant decrease in resolution times over the decade. A crucial insight emerged from the predictive modeling: a simple Linear Regression model showed weak but positive predictive power. In contrast, more complex non-linear models failed, producing negative R^2 scores and indicating a propensity to overfit. The most significant predictors of financial loss were identified as the "Number of Affected Users" and "Incident Resolution Time." Finally, the clustering analysis successfully segmented incidents into distinct archetypes, such as "Mass Data Breach" and "High-Stakes Targeted Attack." This research concludes that noisy, linear relationships govern the financial ramifications of cyber attacks and underscores the principle of model parsimony, providing a quantitative framework for understanding cyber risk and the primary factors that mitigate financial damage.

Keywords: Cybersecurity; Risk Management; Predictive Modeling; Machine Learning; Regression Analysis; Financial Loss Prediction; Data Breach; Threat Intelligence; K-Means Clustering; Random Forest

Introduction

The twenty-first century is defined by an unprecedented reliance on digital infrastructure, which permeates every facet of modern society, from global finance and critical infrastructure to personal communication and commerce [1]. While this digital transformation has unlocked immense efficiencies and opportunities, it has also cultivated a fertile ground for a new and pervasive class of threats: cyber attacks. The last decade, from 2015 to 2024, has witnessed a dramatic escalation in the frequency, sophistication, and scale of these malicious activities [2]. No longer the concern of a niche IT department, cybersecurity has ascended to a primary strategic priority in boardrooms and government chambers worldwide [3].

The consequences of cyber incidents extend far beyond immediate data loss or system downtime. They manifest as a complex cascade of financial damages, including the costs of remediation, regulatory fines, reputational harm, and loss of customer trust, with average breach costs rising significantly in recent years [4, 5]. This reality underscores a critical need to move beyond reactive, perimeter-based defense mechanisms and toward a proactive, data-driven understanding of cyber risk [6]. To effectively allocate finite security resources and develop resilient cyber strategies, organizations must be able to answer a crucial question: What is the likely financial ramification of a given attack?

Despite the clear financial impetus, the quantitative modeling of cyber risk remains a formidable challenge [7]. Many organizations rely on qualitative assessments or historical averages that fail to capture the nuanced dynamics of individual security incidents [8]. The financial impact of a breach is not random; it is a function of multiple variables, including the nature of the attack, the industry targeted, the number of individuals affected, and the efficiency of the response [9, 10]. The central problem this research addresses is the absence of a robust, empirically validated framework for predicting the financial loss of a cybersecurity incident based on its observable characteristics. A reactive approach to comprehending financial damage is insufficient; a predictive capability is essential for modern cyber risk management.

This study, embarks on a quantitative inquiry guided by a series of interconnected research questions, directly leveraging a decade-long global dataset of cybersecurity incidents:

1. ***RQ1***: What are the dominant temporal trends and descriptive patterns in global cybersecurity incidents from 2015 to 2024, specifically concerning attack types, targeted industries, and geographical distributions?
2. ***RQ2***: How accurately can the financial loss of a cybersecurity incident be predicted using its intrinsic and extrinsic characteristics? Which machine learning regression model, among Random Forest, Gradient Boosting, Linear Regression, and Support Vector Regression, provides the most reliable predictive performance?
3. ***RQ3***: What are the most significant factors driving the financial cost of a cyber attack? Does the number of affected users, the resolution time, or the type of attack play a more critical role?
4. ***RQ4***: Beyond supervised prediction, can unsupervised learning techniques reveal a natural, data-driven taxonomy of cybersecurity incidents based on their multi-dimensional attributes?

Expected Contribution to Knowledge

This paper aims to make a substantive contribution to the cybersecurity and data science literature. By analyzing a comprehensive dataset spanning ten years, this research provides a longitudinal perspective that is often missing in contemporary studies. The primary contributions are four-fold:

1. It presents a comprehensive exploratory data analysis of global cyber threats, identifying critical trends and vulnerabilities that have shaped the last decade.
2. It conducts a rigorous comparative analysis of multiple machine learning regression models to establish a benchmark for financial loss prediction in the cybersecurity domain.
3. It employs feature importance analysis to deliver empirical evidence on the specific, quantifiable drivers of financial damage, offering actionable insights for risk mitigation.
4. It introduces an unsupervised clustering approach using K-Means and PCA to segment incidents, providing a richer, more nuanced understanding of threat typologies.

Paper Structure

The remainder of this paper is structured as follows. Section 2 provides a review of the relevant literature in the economics of cybercrime and the application of machine learning in cybersecurity. Section 3 details the methodology, including the dataset description, preprocessing steps, and the analytical framework for both the regression and clustering tasks. Section 4 presents the empirical results, directly addressing the research questions. Section 5 discusses the interpretation and broader implications of these findings. Finally, Section 6 concludes the paper, summarizing the key takeaways and suggesting directions for future research.

Literature Review

This section surveys the foundational and contemporary literature across three critical domains pertinent to this study: the economic principles governing cybersecurity, the application of predictive machine learning models to cyber risk, and the use of unsupervised methods for threat intelligence. By synthesizing these areas, we establish the scholarly context and delineate the specific research gap this paper addresses.

The Economics of Cybercrime and Data Breaches

The study of cybersecurity as an economic problem is well-established, rooted in the seminal work of Gordon and Loeb [11], who proposed a model for determining optimal investment levels in information security. Their framework demonstrated that it is not always economically rational to protect against all threats, establishing that security investment has diminishing marginal returns. Subsequent research has built upon this foundation, seeking to quantify the direct and indirect costs associated with security failures.

Annual industry reports, such as those published by IBM [12] and Verizon [13], provide invaluable empirical data on the financial impact of data breaches, consistently showing rising costs and identifying factors like “time to contain” as major cost drivers. Academic research has sought to create more formal models. Poyraz et al. [14], for instance, developed a data-driven model linking breach characteristics to costs, while Allodi [15] analyzed the economic incentives behind vulnerability disclosure. However, a significant portion of this economic analysis relies on aggregated, post-hoc data, often lacking the granularity to build incident-specific predictive models. While the *what* (the cost) is increasingly well-documented, the *how* (the predictive mechanism based on incident features) remains an area ripe for investigation.

Predictive Modeling for Cybersecurity Risk

The application of machine learning to cybersecurity is not new, with a vast body of literature focused on intrusion detection, malware classification, and spam filtering [16-18]. Within this field, predictive modeling for risk assessment has emerged as a significant sub-discipline. Researchers have employed various regression techniques to forecast security-related outcomes. For example, studies have used regression to predict the likelihood of a firm experiencing a data breach based on its security posture and other organizational characteristics [19]. Others have focused on predicting the duration of a security incident or the number of records likely to be compromised [20].

However, the direct prediction of financial loss using a comparative analysis of multiple regression models on a longitudinal dataset is less common. While models like Random Forest and Gradient Boosting are staples in many machine learning applications due to their high accuracy and ability to handle non-linear relationships [21], their systematic evaluation against traditional linear models for financial impact assessment in cybersecurity is not comprehensively covered. Many studies either focus on a single model or on classification tasks (e.g., high-impact vs. low-impact breach) rather than predicting a continuous financial variable [22].

Unsupervised Learning for Threat Intelligence

Parallel to predictive modeling, unsupervised learning has proven to be a powerful tool for discovering hidden structures and patterns within complex cybersecurity data. Anomaly detection is a primary application, where algorithms identify deviations from normal network traffic or system behavior that may indicate a novel attack [23].

Clustering algorithms, such as K-Means, have been used to group similar cyber attacks, malware families, or security alerts [24]. For example, researchers have applied clustering to network flow data to identify coordinated denial-of-service attacks or to group phishing emails based on their content and structure [25, 26]. This process of creating a data-driven taxonomy is invaluable for understanding threat actor methodologies and developing more targeted defense strategies. Often, dimensionality reduction techniques like Principal Component Analysis (PCA) are used in tandem with clustering to enable visualization and overcome the "curse of dimensionality" inherent in security datasets [27]. This approach allows for the discovery of incident archetypes that may not be apparent through manual analysis or predefined labels.

Gaps in Existing Research

The existing literature provides a strong foundation in the economic analysis of cybercrime, predictive modeling, and unsupervised threat intelligence. However, the review reveals three key gaps that this study aims to address:

1. ***Integration of Methods***: Few studies have holistically combined a longitudinal trend analysis, a comparative evaluation of predictive financial models, and an unsupervised clustering of incident types within a single, unified framework.
2. ***Comparative Predictive Analysis***: While machine learning is widely applied, there is a lack of rigorous, comparative studies focused specifically on predicting the financial ramifications of diverse cyber attacks using modern regression techniques on a recent, decade-long dataset.
3. ***Data-Driven Taxonomy***: There is an opportunity to move beyond predefined attack labels and use unsupervised learning to derive a more nuanced, empirical taxonomy of incident archetypes based on a rich set of features, including financial and temporal metrics.

This paper, therefore, positions itself at the intersection of these three domains, leveraging a comprehensive dataset to build and validate a model for financial risk prediction while simultaneously exploring the underlying structure of the global threat landscape.

Methodology

This section outlines the systematic, quantitative approach employed to address the research questions. The methodology is structured into several sequential phases: dataset description, data preprocessing, exploratory data analysis (EDA), supervised predictive modeling, and unsupervised clustering. This framework ensures a comprehensive analysis, from initial data inspection to advanced model-based inference.

Dataset Description

The foundation of this research is the "Global Cybersecurity Threats 2015-2024" dataset [28], a structured collection of records detailing individual cybersecurity incidents. The dataset encompasses a ten-year period, providing a valuable longitudinal view of the threat landscape. Each record contains several attributes critical for this analysis, including:

1. ***Categorical Features***: Country, Attack Type, Target Industry, Attack Source, Security Vulnerability Type, and Defense Mechanism Used.
2. ***Numerical Features***: Year, Financial Loss (in million $), Number of Affected Users, and Incident Resolution Time (in hours).

The target variable for the predictive modeling phase of this study is Financial Loss (in millions $), a continuous variable representing the reported financial impact of each incident. Table 1 provides a more detailed overview of the dataset's features.

Feature Name	*Category*	*Description*	*Data Type*
Country	Categorical	The nation where the cybersecurity incident occurred.	String (Nominal)
Attack Type	Categorical	The method or technique used in the cyberattack (e.g., phishing, ransomware).	String (Nominal)
Target Industry	Categorical	The sector affected by the cyberattack (e.g., healthcare, banking, IT).	String (Nominal)
Attack Source	Categorical	Identified origin of the attack (e.g., nation-state, hacktivist, insider).	String (Nominal)
Security Vulnerability Type	Categorical	Type of exploited security flaw (e.g., zero-day, social engineering).	String (Nominal)
Defense Mechanism Used	Categorical	Countermeasures or systems in place at the time of the incident.	String (Nominal)
Year	Numerical	The calendar year the incident occurred (2015–2024).	Integer
Financial Loss	Numerical (Target)	Estimated economic damage in millions of USD caused by the incident.	Float (Continuous)
Number of Affected Users	Numerical	Number of individuals or entities impacted by the breach.	Integer
Incident Resolution Time	Numerical	Time taken (in hours) to resolve or mitigate the incident.	Float (Continuous)

Table 1: Dataset Features Overview.

Data Preprocessing

To prepare the dataset for analysis and modeling, a series of standard preprocessing steps were executed. First, all column headers were converted to a consistent lowercase and snake_case format to facilitate programmatic access. Subsequently, a critical step of data transformation was performed to handle the categorical features. As machine learning algorithms require numerical input, the *LabelEncoder* function from the Scikit-learn library was applied to each categorical column. This process assigns a unique integer to each distinct category within a feature, transforming the non-numeric data into a machine-readable format.

Exploratory Data Analysis (EDA)

To address RQ1, a comprehensive exploratory data analysis was conducted to uncover initial patterns, trends, and relationships within the data. This phase utilized a suite of visualization techniques to summarize key characteristics of the dataset, including:

1. ***Bar Charts***: To compare aggregated values across categories, such as total financial loss per country and the number of affected users by attack type. Figures 1-3 used bar charts to provide an overview of the total financial loss per country, the top 10 affected countries, and the number of affected users by attack type, respectively.

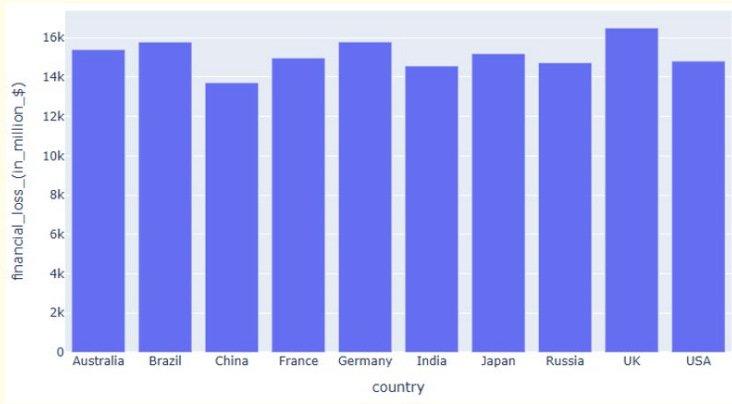

Figure 1: Total Financial Loss Per Country.

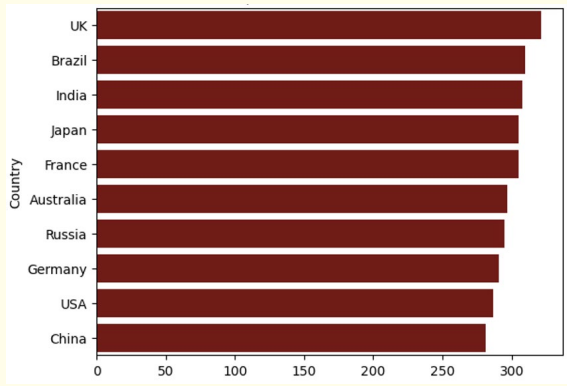

Figure 2: Top 10 Affected Countries.

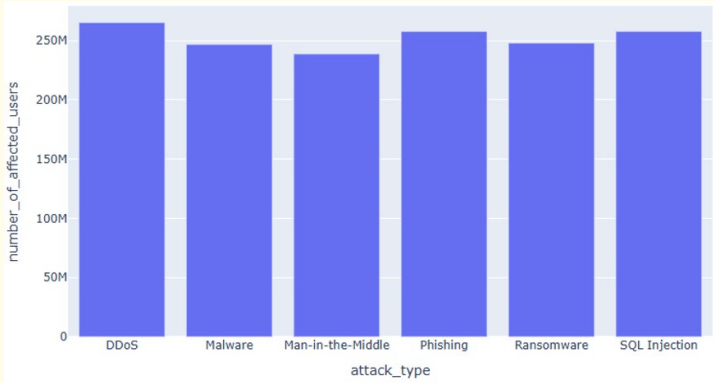

Figure 3: Overview of Affected Users by Attack Type.

2. ***Pie Charts***: Used in Figures 4-6 to illustrate the proportional distribution of categorical features like attack types and their sources. They summarize the distribution of attack types, attack sources, and defense mechanisms used respectively.

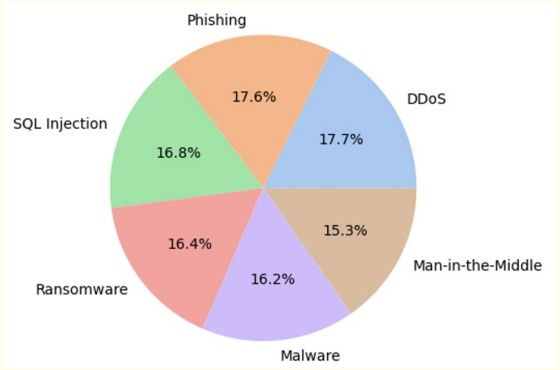

Figure 4: Distribution of Attack Types.

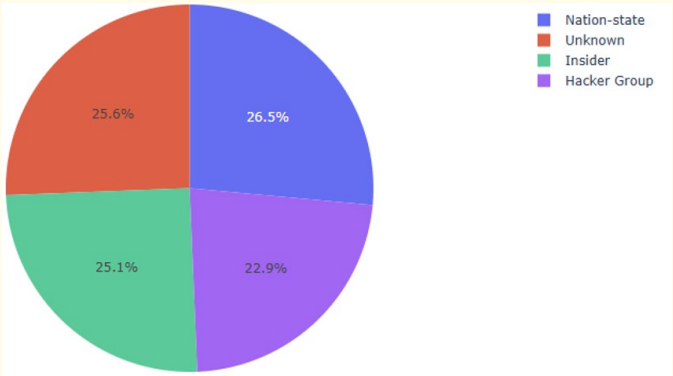

Figure 5: Distribution of Attack Sources.

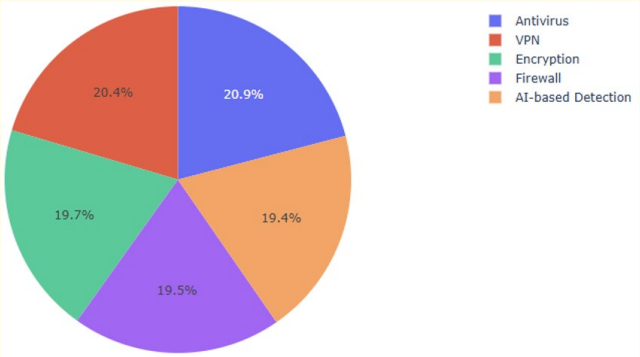

Figure 6: Defense Mechanisms Used Across the Dataset.

3. ***Line Plots***: To visualize temporal trends, specifically tracking the evolution of financial losses and incident resolution times over the ten-year period. Figure 7 shows the total money lost across the period, and Figure 8 shows the average incident resolution time over the years.

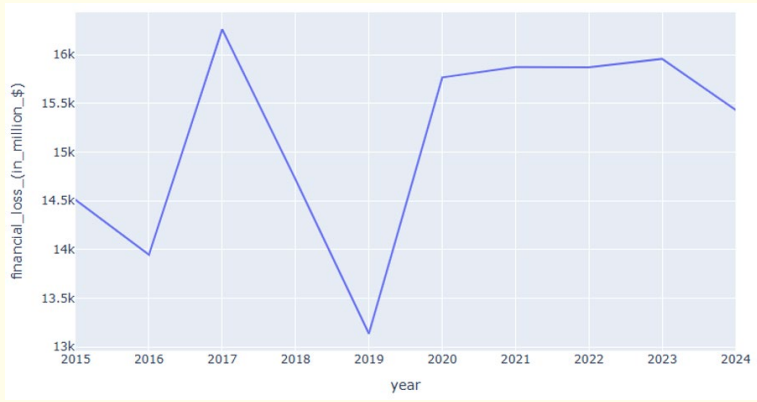

Figure 7: Financial Loss Trendline Over the Years.

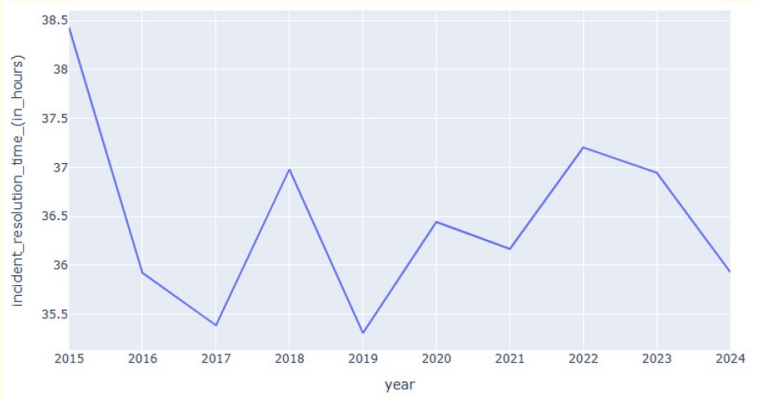

Figure 8: Average Incident Resolution Time Over the Years.

4. ***Advanced Visualizations***: Box plots, violin plots, scatter plots, heatmaps of the correlation matrix, treemaps, and sunburst charts were employed to explore more complex relationships between variables, such as the correlation between financial loss and the number of affected users. Figure 9 shows a treemap summarizing the financial loss across the different attack types, Figure 10 uses a violin plot to map the incident resolution times across various industries, and Figure 11 uses a sunburst chart to map the attack source by target industry and the financial loss.

Figure 9: Overview of Financial Loss by Attack Type.

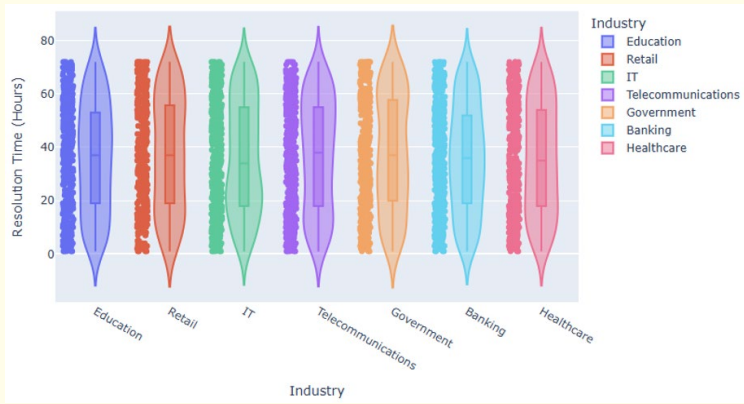

Figure 10: Average Incident Resolution Time by Industry.

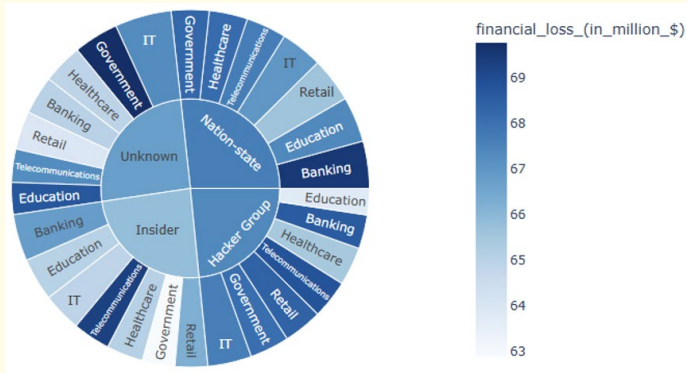

Figure 11: Attack Source Breakdown by Target Industry.

Predictive Modeling of Financial Loss

The central component of this research involved developing and evaluating supervised learning models to predict the financial loss of cyber incidents.

Feature Selection

Based on domain knowledge and preliminary analysis, a subset of features was selected as predictors (X) for the regression task: Number of Affected Users, Incident Resolution Time (in hours), Country, Year, Target Industry, and Attack Type. The target variable (y) was Financial Loss (in million $).

Model Implementation

To ensure a robust and comparative analysis (RQ2), four distinct regression algorithms were selected, representing different underlying mathematical principles:

1. ***Random Forest Regressor***: An ensemble method based on decision trees that mitigates overfitting by averaging multiple models.
2. ***Gradient Boosting Regressor***: Another powerful ensemble technique that builds models sequentially, with each new model correcting the errors of the previous one.
3. ***Linear Regression***: A fundamental statistical model used as a baseline to assess the performance of more complex models.
4. ***Support Vector Regressor (SVR):*** A variant of Support Vector Machines adapted for regression tasks, effective in high-dimensional spaces.

Experimental Setup

The dataset was partitioned into a training set (80%) and a testing set (20%) to allow for model training and subsequent evaluation on unseen data. To prevent data leakage and ensure fair evaluation, the *StandardScaler* function from Scikit-learn was fitted only on the training data and then used to transform both the training and testing sets. This step normalizes the features to have a mean of zero and a standard deviation of one, which is crucial for the performance of algorithms like SVR. A fixed *random_state* was used throughout to ensure the reproducibility of the results.

Evaluation Metrics

The performance of each model was quantitatively assessed using three standard regression metrics:

1. ***Mean Absolute Error (MAE):*** Measures the average absolute difference between the predicted and actual values, providing an easily interpretable measure of error in the target variable's original units.
2. ***Mean Squared Error (MSE):*** Calculates the average of the squared differences between predicted and actual values. It penalizes larger errors more heavily.
3. ***R-squared (R^2 Score):*** Represents the proportion of the variance in the dependent variable that is predictable from the independent variables. It provides a measure of how well the model explains the data, with values closer to 1 indicating a better fit.

Feature Importance Analysis

To address RQ3, the feature importance scores were extracted from the trained Random Forest model. This technique intrinsically ranks predictors by their contribution to reducing variance, thereby identifying the most influential factors in determining financial loss.

Unsupervised Clustering of Cyber Incidents

To explore the dataset for inherent structures and address RQ4, an unsupervised learning approach was implemented.

K-Means Clustering

The K-Means algorithm was employed to partition the incidents into a predefined number of clusters (k). The objective of the algorithm is to group similar data points together based on their features, minimizing the within-cluster sum of squares (inertia). The "Elbow Method" was used as a heuristic to identify the optimal value for k by plotting the inertia for a range of k values and selecting the "elbow" point where the rate of decrease sharply changes.

Dimensionality Reduction for Visualization

As the dataset contains multiple features, visualizing the clusters directly is not feasible. Therefore, Principal Component Analysis (PCA) was applied to the scaled dataset. PCA is a dimensionality-reduction technique that transforms the data into a new coordinate system of orthogonal principal components. By projecting the data onto the first two or three principal components, which capture the maximum variance, we can effectively visualize the cluster separations in 2D and 3D space.

Results

Exploratory Data Analysis: The Threat Landscape (2015-2024)

The initial exploratory data analysis (EDA) revealed several significant trends and patterns within the global cybersecurity landscape, directly addressing RQ1.

1. ***Temporal Trends***: Analysis of financial loss over the ten-year period showed notable fluctuation without a clear linear trend. For instance, losses peaked in 2017 at approximately $16.2 billion and saw similar levels in the 2021-2023 period, while dipping to a low of around $13.1 billion in 2019. Similarly, the average incident resolution time remained relatively stable, oscillating between a low of 35.3 hours in 2017 and a high of 38.4 hours in 2015, indicating no significant trend towards faster or slower resolutions over the decade.
2. ***Attack and Defense Distribution***: A frequency count of attack vectors showed a relatively even distribution among the most common types. DDoS attacks were the most frequent (531 incidents), closely followed by phishing (529), sql injection (503), ransomware (493), and malware (485). Analysis of threat origins revealed that attacks attributed to nation-states were the most numerous (794), followed by those from unknown sources (768) and insiders (752). Correspondingly, antivirus software (628 instances) was the most commonly cited defense mechanism used.
3. ***Impact by Geography***: The financial impact of cyber attacks was substantial across several major economies. The data indicates that among the most affected nations, the aggregate financial losses were comparably high, ranging from approximately $13.7 billion to $16.5 billion. The United Kingdom ($16.50 billion), Germany ($15.79 billion), and Brazil ($15.78 billion) reported the highest financial damages in this group. This finding suggests a widespread and significant financial burden of cyber incidents across developed and developing economies, rather than a concentration of damage in only a few select countries.

Performance of Predictive Models for Financial Loss

To answer RQ2, four regression models were trained and evaluated for their ability to predict financial loss based on incident characteristics. The performance of these models on the unseen test dataset is summarized in Table 2.

The evaluation yielded a nuanced and critical finding. The Linear Regression model was the only one to achieve a positive R-squared (R^2) score. While modest, this positive score is significant: it indicates that a simple linear combination of the input features possesses genuine, albeit weak, predictive power. This model outperformed a naive baseline of simply guessing the average financial loss for every incident, demonstrating that a discernible linear relationship exists between the features and the outcome.

Rank	*Model*	*MAE*	*MSE*	*R^2 Score*
1	Linear Regression	24.623882	809.376668	0.000919
2	Support Vector Regressor	24.740873	818.750467	-0.012511
3	Gradient Boosting	24.781194	835.291598	-0.032967
4	Random Forest	25.531562	895.068822	-0.106890

Table 2: Regression Model Performance Ranking.

In stark contrast, the more complex, non-linear models, Support Vector Regressor, Gradient Boosting, and Random Forest, all produced negative R^2 scores. A negative R^2 signifies that a model's predictions are worse than the simple baseline average. This result suggests that these powerful models, in their attempt to capture complex, non-linear patterns, likely overfit to the noise within the training data. Their failure to generalize to the unseen test set made them not only less accurate than the linear model but fundamentally unreliable for this prediction task.

The key finding for RQ2 is therefore not that prediction is impossible, but rather that the financial impact of cyber incidents exhibits a weak linear predictability that is lost by more complex models. This outcome underscores the importance of model simplicity (a principle often referred to as Occam's razor [29]). For this dataset, the attempt to apply sophisticated, non-linear algorithms was counterproductive, as they failed to capture the simple underlying signal and instead amplified noise, leading to poorer performance on unseen data.

Feature Importance Analysis

To investigate RQ3, a feature importance analysis was conducted using the trained Random Forest model. Although the model's overall predictive performance was poor, this analysis can still offer insights into which features the model attempted to use to make its predictions. The analysis, summarized in Figure 12, revealed that Number of Affected Users and Incident Resolution Time (in hours) were ranked as the two most important features. Country and Year followed with moderate importance, while Target Industry and Attack Type had the lowest importance scores.

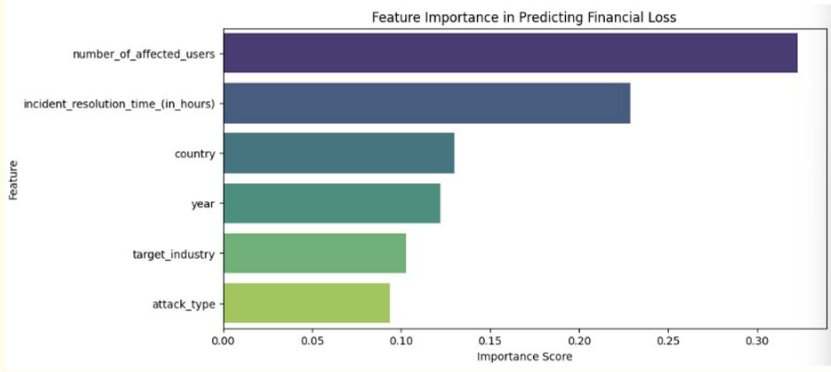

Figure 12: Feature Importance Analysis of the Random Forest Regressor.

It is crucial to interpret this finding with caution. Given the model's negative R^2 score, these importances do not reflect features that successfully predict the outcome. Rather, they indicate the features that account for the most variance reduction within the structure of the poorly performing decision tree models. The high ranking of Number of Affected Users and Incident Resolution Time does, however, align with industry reports [13] and provides a weak confirmation of their relevance, even if a robust predictive signal could

not be established by the models.

Incident Cluster Analysis

The unsupervised clustering analysis was performed to address RQ4 and identify inherent groupings within the incident data. The Elbow Method, shown in Figure 13, suggested an optimal number of two or three clusters. Upon fitting the K-Means algorithm with k=2 and visualizing the results using PCA, distinct clusters emerged.

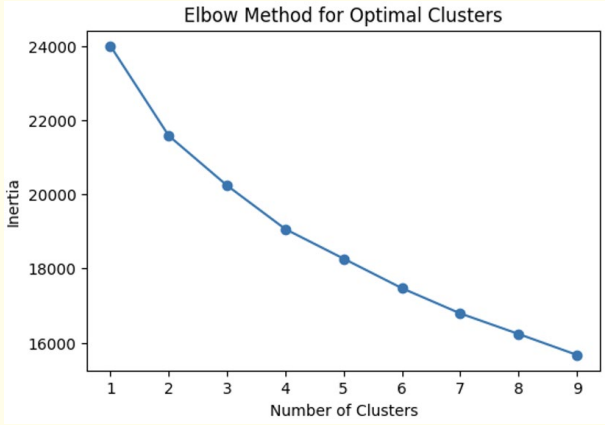

Figure 13: Elbow Method Plot.

The 2D and 3D scatter plots of the principal components revealed clear separation between the clusters. Figures 14 and 15 provide a snapshot of the 2D and 3D scatter plots of the PCA result visualization.

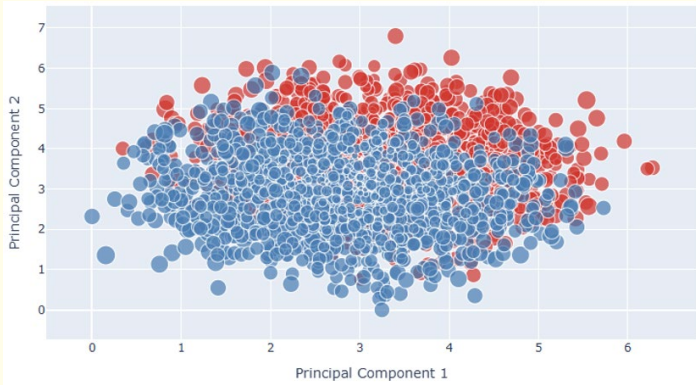

Figure 14: PCA Result Visualization (2D).

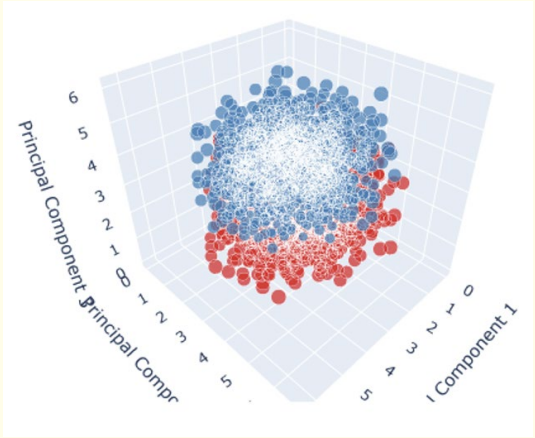

Figure 15: PCA Result Visualization (3D).

Discussion and Future Work

Discussion

This study's findings offer a nuanced perspective on the cybersecurity landscape. The exploratory analysis (RQ1) reveals a volatile but persistent threat environment. The lack of a clear downward trend in financial losses or incident resolution times over the past decade suggests that defensive measures are, at best, keeping pace with an evolving and increasingly monetized cybercrime ecosystem. The widespread financial damage across major economies like the UK, Germany, and Brazil underscores that this is a global issue, not one confined to a single region.

The most striking finding emerged from the predictive modeling (RQ2). The modest success of the Linear Regression model, juxtaposed with the failure of more complex models like Random Forest and Gradient Boosting, is a crucial outcome. It suggests that while a weak, predictable linear relationship exists between incident characteristics and financial loss, this signal is fragile. The more powerful non-linear models likely overfit to noise in the training data, a phenomenon that highlights a critical lesson in data science: model complexity is not a panacea. The principle of Occam's razor, that simpler solutions are often preferable, holds true here. The financial impact of cyber incidents is not governed by an esoteric, complex pattern but by a few fundamental, albeit noisy, linear relationships.

The feature importance analysis (RQ3) aligns with this interpretation. The prominence of 'Number of Affected Users' and 'Incident Resolution Time' as key features, even in a poorly performing model, reinforces their practical relevance. This finding has direct operational implications: the scale of a breach and the speed of response are primary drivers of financial damage. Organizations that invest in capabilities to limit the scope of an incident and accelerate recovery can directly mitigate financial harm.

Furthermore, the clustering analysis (RQ4) successfully identified distinct incident archetypes, demonstrating that incidents can be meaningfully categorized even if their exact financial cost is hard to predict. Based on their characteristics, we can propose the following personas for these clusters:

1. ***Cluster 0***: The "Mass Data Breach" Archetype. This group consists of incidents with a very high number of affected users but moderate resolution times and financial losses. These likely represent large-scale breaches of less sensitive data (e.g., email lists, user credentials) where the immediate financial fallout is contained, but the reputational damage and long-tail costs could be significant.

2. ***Cluster 1***: The "High-Stakes Targeted Attack" Archetype. Characterized by extremely long resolution times and high financial losses affecting fewer users, this cluster likely represents sophisticated, targeted attacks. Examples include advanced persistent threats (APTs) against critical infrastructure or complex ransomware attacks that cripple core business systems, leading to prolonged downtime and costly remediation.

Limitations of the Study

This study is not without its limitations. First, the dataset relies on publicly reported incidents, which may introduce reporting bias. High-profile attacks are more likely to be reported than smaller, less damaging ones, and financial losses may be estimated or undisclosed. Second, the features available for modeling were high-level. Lacking more granular data, such as specific software vulnerabilities (CVEs) or detailed threat actor attribution, limits the potential depth of the predictive models. Finally, the chosen models represent a specific subset of machine learning techniques; other architectures could potentially yield different results.

Future Work

Building on this work, several avenues for future research are apparent.

1. ***Incorporate Granular Data***: Future studies should seek to integrate more detailed datasets that include specific CVEs, threat actor TTPs (Tactics, Techniques, and Procedures), and company-specific data like cybersecurity budgets and employee training levels.
2. ***Explore Advanced Models***: While complex models failed here, this may be due to the limited feature set. With richer data, advanced deep learning models, such as LSTMs for time-series analysis of threat trends or Graph Neural Networks for analyzing attack paths, may prove effective.
3. ***Develop a Real-Time Implementation***: An interactive, real-time version of the linear risk model could be developed as a practical tool for IT security managers to quickly assess the potential financial exposure of an emerging threat.
4. ***Perform Qualitative Analysis***: To triangulate the findings from the cluster analysis, a qualitative study involving interviews with incident responders could provide richer context and validate the proposed incident archetypes.

Conclusion

This research set out to analyze the cybersecurity threat landscape and assess the feasibility of predicting financial losses from cyber incidents. The study successfully identified key trends, demonstrating a persistent and costly global threat environment. We answered our primary research questions, concluding that: (1) the cyber threat landscape shows fluctuating but consistently high financial damages with no significant improvement in resolution times; (2) predicting exact financial loss is challenging, but a simple linear model can outperform complex models, suggesting the underlying predictive signals are linear but noisy; (3) the number of affected users and incident resolution time are key drivers of cost; and (4) incidents can be grouped into distinct, interpretable archetypes. The main contribution of this paper is the nuanced finding that in the context of cybersecurity financial prediction, model simplicity can triumph over complexity, offering a valuable lesson for both researchers and practitioners.

References

1. World Economic Forum. "Global Risks Report 2024". World Economic Forum (2024). https://www.weforum.org/publications/global-risks-report-2024/
2. Balsara Bhavish. "A Comparative Study of Patterns, Causes, And Impacts of Data Breaches Across Geographical Regions and Time Frames" (2024). Electronic Theses, Projects, and Dissertations 2080. https://scholarworks.lib.csusb.edu/etd/2080
3. M Mijwil., et al. "The Purpose of Cybersecurity Governance in the Digital Transformation of Public Services and Protecting the Digital Environment". Mesopotamian Journal of CyberSecurity (2023): 1-6.
4. P Ganiaridis. "Evaluating the financial effect from cyber attacks on firms and analysis of cyber risk management". M.S. thesis, Dept. of Banking and Fin. Mgmt., Univ. of Piraeus, Piraeus, Greece (2018).

5. R Dieye., et al. "Estimates of the macroeconomic costs of cyber-attacks". Risk management and insurance review 23.2 (2020): 183-208.
6. JB Fraley and J Cannady. "The promise of machine learning in cybersecurity". in SoutheastCon 2017, Concord, NC, USA (2017): 1-6.
7. A Wirth. "The Economics of Cybersecurity". Biomedical Instrumentation & Technology 51.s6 (2017): 52-59.
8. DW Hubbard and R Seiersen. "How to Measure Anything in Cybersecurity Risk". Hoboken, NJ, USA: Wiley (2016).
9. SJ Zaccaro., et al. "A Comprehensive Multilevel Taxonomy of Cyber Security Incident Response Performance". Psychosocial Dynamics of Cyber Security (2016): 43-85.
10. J DeCoste. "The impact of cyber-attacks on publicly traded companies". M.S. thesis, Concordia Inst. for Inf. Syst. Eng., Concordia Univ., Montreal, QC, Canada, (2017). [Online]. https://spectrum.library.concordia.ca/id/eprint/982695/
11. LA Gordon and MP Loeb. "The economics of information security investment". ACM Transactions on Information and System Security 5.4 (2002): 438-457.
12. IBM. "Cost of a Data Breach Report 2023" (2023). [Online]. https://www.ibm.com/reports/data-breach
13. Verizon. "DBIR 2023 Data Breach Investigations Report Public Sector Snapshot" (2023). [Online]. https://www.verizon.com/business/resources/Ta5a/reports/2023-dbir-public-sector-snapshot.pdf
14. OI Poyraz., et al. "Cyber assets at risk: monetary impact of U.S. personally identifiable information mega data breaches". The Geneva Papers on Risk and Insurance - Issues and Practice 45.4 (2020): 616-638.
15. L Allodi. "Economic Factors of Vulnerability Trade and Exploitation". in Proc. 2017 ACM SIGSAC Conf. Comput. Commun. Secur. (CCS '17) (2017): 1483-1499.
16. AL Buczak and E Guven. "A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection". IEEE Communications Surveys & Tutorials 18.2 (2016): 1153-1176.
17. BG Bokolo and Q Liu. "Artificial Intelligence in Social Media Forensics: A Comprehensive Survey and Analysis". Electronics 13.9 (2024): 1671.
18. BG Bokolo, R Jinad and Q Liu. "A Comparison Study to Detect Malware using Deep Learning and Machine learning Techniques". in 2023 IEEE 6th Int. Conf. Big Data Artif. Intell. (BDAI), Jiaxing, China (2023): 1-6.
19. A Sarabi., et al. "Risky business: Fine-grained data breach prediction using business profiles". Journal of Cybersecurity 2.1 (2016): 15-28.
20. B Edwards, S Hofmeyr and S Forrest. "Hype and heavy tails: A closer look at data breaches". Journal of Cybersecurity 2.1 (2016): 3-14.
21. A Handa, A Sharma and SK Shukla. "Machine learning in cybersecurity: A review". WIREs Data Mining and Knowledge Discovery 9.4 (2019).
22. J Martínez Torres, C Iglesias Comesaña and PJ García-Nieto. "Review: machine learning techniques applied to cybersecurity". International Journal of Machine Learning and Cybernetics 10.10 (2019): 2823-2836.
23. S Ahmad., et al. "Unsupervised real-time anomaly detection for streaming data". Neurocomputing 262 (2017): 134-147.
24. EK Viegas, AO Santin and LS Oliveira. "Toward a reliable anomaly-based intrusion detection in real-world environments". Computer Networks 127 (2017): 200-216.
25. K Althobaiti., et al. "Using Clustering Algorithms to Automatically Identify Phishing Campaigns". IEEE Access 11 (2023): 96502-96513.
26. Murk Marvi, Asad Arfeen and R Uddin. "An augmented K-means clustering approach for the detection of distributed denial-of-service attacks". International Journal of Network Management 31.6 (2021).
27. A Parizad and CJ Hatziadoniu. "Cyber-Attack Detection Using Principal Component Analysis and Noisy Clustering Algorithms: A Collaborative Machine Learning-Based Framework". IEEE Transactions on Smart Grid 13.6 (2022): 4848-4861.
28. Atharva Soundankar. "Global Cybersecurity Threats (2015-2024)". Kaggle.com (2015). https://www.kaggle.com/datasets/atharvasoundankar/global-cybersecurity-threats-2015-2024

29. FJ Bargagli Stoffi, G Cevolani and G Gnecco. “Simple Models in Complex Worlds: Occam’s Razor and Statistical Learning Theory”. Minds and Machines 32.1 (2022): 13-42.