Lotka-Volterra Dynamics Neural Model for Cyber Risk Assessment*

Serhii Yevseiev[1†], Oleksandr Kushnerov[2†*], Yevhen Melenti[3†], Stanislav Milevskyi[1†] and Andrey Sharapata[4†]

[1] National Technical University "Kharkiv Polytechnic Institute", Kyrpychova 2 61002 Kharkiv, Ukraine
[2] Sumy State University, Kharkivska 116 40007 Sumy, Ukraine
[3] National Academy of Security Service of Ukraine, Maksymovycha 22 03022 Kyiv, Ukraine
[4] Kharkiv National Automobile and Highway University, Yaroslava Mudroho 25 61002 Kharkiv, Ukraine

Abstract

A hybrid model for dynamic cyber risk assessment is proposed that integrates a deep neural network and the Lotka–Volterra model. The model simultaneously classifies network traffic (normal/anomaly) and predicts coefficients (α, β, γ, ϕ) that determine the dynamics of the attack-defence interaction. Trained and tested on NSL-KDD data, the model achieved a classification accuracy of 0.8006, AUC of 0.9016, and MSE of 0.0027 for coefficient prediction. Statistically significant differences in the predicted coefficients for normal and anomalous sessions were found, indicating that the model successfully captures underlying characteristics that differentiate these two classes beyond simple pattern matching. Simulation of the Lotka–Volterra dynamics with predicted parameters demonstrates different patterns for different traffic classes, indicating the approach's potential for deeper risk assessment compared to traditional intrusion detection methods. This ability to forecast interaction dynamics provides a forward-looking view of potential threats, a significant step beyond simple, reactive threat identification.

Keywords

cyber risk assessment, neural networks, Lotka–Volterra dynamics, intrusion detection, NSL-KDD

1. Introduction

Modern cyber threats' growing complexity and dynamic nature necessitate transitioning from traditional, static risk assessment methods to more adaptive and predictive approaches. Classical Intrusion Detection Systems (IDS) and conventional risk assessment methodologies are often limited to real-time threat identification. However, they frequently fail to account for the long-term interaction between attackers and defence systems, which severely limits their ability to predict the evolution of an attack and its potential consequences. In an era where cyberattacks are increasingly automated and operate at machine speed, purely reactive defence mechanisms are fundamentally inadequate.

Applying dynamic systems models has emerged as a promising avenue to address this critical gap, particularly the Lotka–Volterra model [6]. This model, originally formulated to describe predator-prey interactions, is a powerful analogy for modelling cyberspace's adversarial "attack-defence" relationship [6]. However, a significant challenge lies in estimating the parameters of such dynamic models from complex, high-dimensional network data—a task for which modern machine learning approaches are exceptionally well-suited.

Proceedings of the Workshop on Scientific and Practical Issues of Cybersecurity and Information Technology at the V international scientific and practical conference Information security and information technology (ISecIT 2025), June 09–11, 2025, Lutsk, Ukraine
* Corresponding author.
† These authors contributed equally.
Serhii.Yevseiev@gmail.com (S. Yevseiev); o.kushnerov@biem.sumdu.edu.ua (O. Kushnerov); melenty@ukr.net (Y. Melenti); milevskiysv@gmail.com (S. Milevskyi); phd.sharapata@gmail.com (A. Sharapata)
0000-0003-1647-6444 (S. Yevseiev); 0000-0001-8253-5698 (O. Kushnerov); 0000-0003-2955-2469 (Y. Melenti); 0000-0001-5087-7036 (S. Milevskyi); 0000-0003-0823-9262 (A. Sharapata)
© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-01-16

CEUR-WS.org/Vol-4150/paper12.pdf

This research presents an innovative hybrid neural model to integrate two key tasks into a cohesive framework. First, it performs robust network traffic classification for intrusion detection, leveraging the power of modern deep learning approaches that have proven effective in traffic analysis [1; 2; 3]. Second, and what constitutes its key feature, the model simultaneously predicts the parameters for the dynamic Lotka–Volterra model. This dual-purpose approach allows the system to identify an anomaly and quantitatively assess the dynamic potential of the associated risk. This synthesis of a predictive mathematical model with a powerful deep learning engine for parameter estimation is the central contribution of our work.

The primary goal of this work is to develop and thoroughly validate such a model, demonstrating its ability to generate qualitatively different dynamic patterns for normal and anomalous network activity. Thus, we aim to show that the predicted parameters carry essential, actionable information for a deeper, more predictive assessment of cyber risks, moving significantly beyond the capabilities of traditional detection methods. This paper details the model's architecture and training methodology and thoroughly evaluates its performance, demonstrating its prognostic capabilities through simulation.

2. Materials and methods

The empirical foundation of this study is the NSL-KDD dataset [5], which is widely recognised as a standard benchmark for evaluating the performance of intrusion detection systems [4]. A comprehensive series of preprocessing steps was meticulously applied to prepare the data for effective processing by the neural network model.

The byte_ratio feature was created from existing byte counts to provide more relational context. Furthermore, categorical features, specifically protocol_type and flag, which are non-numeric, were converted into a numerical format suitable for the neural network using one-hot encoding. As a final preprocessing step, all numerical features were standardised using Z-score normalisation. This ensured that all features had a mean of 0 and a standard deviation of 1, which is critical for allowing all features to contribute equally to the model's learning and helping to accelerate the convergence of the training process.

The theoretical core of our approach is an adapted Lotka–Volterra model [6], which mathematically describes the dynamic interaction between the level of attack, A(t), and the level of protection, Z(t) . The system is formally defined by the following pair of differential equations:

$$\frac{\mathrm{dA}}{\mathrm{dt}} = \alpha \mathrm{A} - \beta \mathrm{A}\, \mathrm{Z}$$
$$\frac{dZ}{dt} = \gamma AZ - \phi Z$$

Where A(t) represents the aggregate level of attack activity, and Z(t) represents the deployed level of defensive measures at a given time t. The terms in these equations capture the essential feedback loops of the adversarial relationship:

- Attack Dynamics ($\frac{\mathrm{dA}}{\mathrm{dt}}$). Two main forces govern the change in the attack level. The term αA represents the intrinsic growth of the attack, such as the natural rate of malware propagation or scanning for new victims, assuming no defensive opposition. The term - βA Z means the reduction of the attack level due to successful neutralisation by the defence system; it is proportional to the frequency of interactions between attacks and defences.

- Defence Dynamics ($\frac{dZ}{dt}$). Opposing factors similarly drive the change in the defence level. The term γAZ models the reactive growth and adaptation of the defence system in response to

detected attacks, such as deploying new firewall rules or patching vulnerabilities. The term $-\phi Z$ represents the "cost" or natural decay of the defence effort over time, which can be interpreted as maintenance costs, resource depreciation, or the obsolescence of security measures that are no longer effective.

Within this framework, the coefficients are interpreted as follows: α represents the intrinsic growth rate of the attack; β signifies the effectiveness of the defense in neutralising the attack; γ corresponds to the rate at which the defense adapts or grows in response to an attack; and ϕ denotes the cost or natural decay rate of the defense system over time.

A critical step in our methodology was operationalising these abstract coefficients to create trainable targets for the neural network. These coefficients were empirically calculated based on specific, measurable features from the NSL-KDD dataset to generate ground-truth values for training. For instance, metrics such as serror_rate, rerror_rate, and anomaly frequencies across different services were used to derive proxy values for the Lotka-Volterra coefficients. This process allowed us to obtain concrete target values for the neural network's regression task.

We designed a hybrid, multi-task neural network to simultaneously perform two distinct but related tasks: binary classification of network traffic (normal/anomaly) and regression to predict the four coefficients (α, β, γ, ϕ) of the Lotka–Volterra model. The architecture's input layer accepts the preprocessed feature vectors and adds a layer of Gaussian noise, which acts as regularisation to enhance model robustness and prevent overfitting. These inputs pass through two shared, fully connected (dense) layers with 256 and 128 neurons, respectively. These layers utilise the Rectified Linear Unit (ReLU) activation function to introduce non-linearity. Batch Normalisation follows each step to stabilise the training process, and a Dropout layer is used for further regularisation.

Following these shared layers, the architecture splits into two separate output heads, one for each task. The classification head consists of a dense layer with a Sigmoid activation function, which produces a probability score indicating whether the input is an anomaly. The loss for this head is calculated using a binary cross-entropy function, which is standard for binary classification tasks. The regression head employs a dense layer with a linear activation function to output four continuous values corresponding to the Lotka-Volterra coefficients. The loss for this head is measured by the mean squared error (MSE) function, which quantifies the average squared difference between the predicted and actual coefficients. The overall loss for the model is a weighted sum of these two individual losses, with weights of 1.0 for classification and 0.2 for regression, balancing the two tasks during training. Before being used in simulations, the predicted coefficients are clipped to the range [0, 1] to ensure stability. The model was trained using the AdamW optimiser, a robust choice that follows standard deep learning practices [2].

The training process was carefully managed using several control mechanisms to ensure optimal performance and prevent overfitting. To retain the best-performing version of the model, its weights were saved only when the Area Under the Curve (AUC) metric on a separate validation set showed improvement. Additionally, the training employed an adaptive learning rate scheduler, which automatically reduced the learning rate whenever the validation performance plateaued, allowing for finer adjustments and more stable convergence. Finally, an early stopping mechanism was implemented to halt the training process automatically if the validation AUC did not improve for a set number of consecutive epochs, thereby preventing the model from overfitting to the training data and enhancing its generalisation capabilities.

3. Assessment of hybrid model results

The model's training process was monitored to ensure stability and prevent overfitting. The learning curves, depicted in Figure 1, provide a detailed visualisation of the model's performance on both the

training and validation sets across epochs for three key metrics: accuracy, Area Under the Curve (AUC), and the Mean Squared Error (MSE) of the Lotka-Volterra coefficients. As shown in the figure, the performance metrics on the validation set consistently and closely track those on the training set. For instance, both sets' accuracy and AUC curves rise in tandem and stabilise, while the MSE curves decrease sharply and remain low. This parallel progression is strong evidence of stable convergence and indicates that the model did not suffer from significant overfitting.

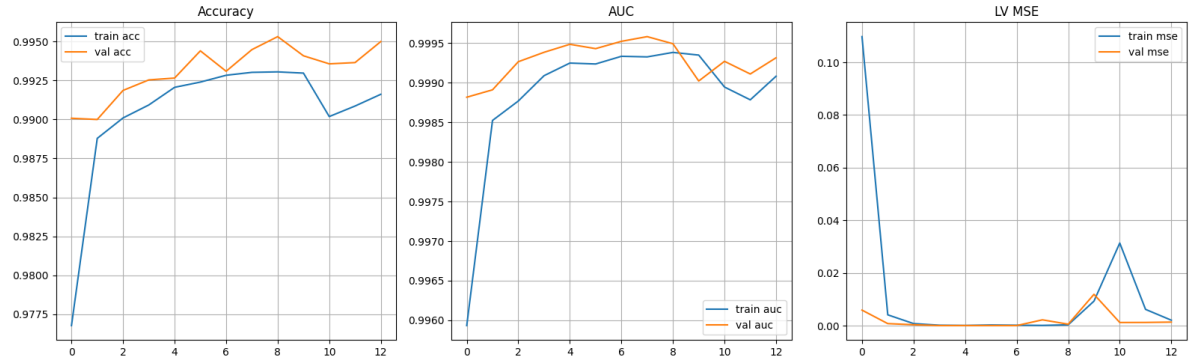

Figure 1: Model learning curves

The learning curves in Figure 1 show that the model's classification metrics (Accuracy, AUC) steadily improved. At the same time, the regression error (LV MSE) rapidly decreased to near-zero for both training and validation sets. Crucially, the validation curves closely track the training curves across all plots. This demonstrates stable convergence and indicates that the model generalises well without suffering from significant overfitting.

Upon completion of training, the model's final performance was evaluated on the unseen NSL-KDD test dataset (Table 1).

Table 1
Model performance indicators on the NSL-KDD test set

Metrics	Value
Classification accuracy	0.8006
Classification AUC	0.9016
MSE for the coefficients of the Lotka–Volterra model	0.0027

The model achieved a classification accuracy of 0.8006, demonstrating a strong capability to identify traffic instances correctly. The Area Under the Curve (AUC) metric reached 0.9016. This high AUC value is significant as it indicates excellent discrimination between the standard and anomalous classes across all classification thresholds, confirming the model's robustness as a classifier. The model demonstrated high fidelity in predicting the dynamic parameters for the regression task, which is central to our hybrid approach. This was evidenced by a very low Mean Squared Error (MSE) of 0.0027, validating the model's ability to learn and predict the Lotka-Volterra coefficients accurately.

A core objective of this study was to determine if the predicted Lotka-Volterra coefficients (α, β, γ, ϕ) capture meaningful, underlying differences between regular and malicious network activity beyond simple classification. To investigate this, an analysis was conducted on the model's predictions for 13,592 normal and 8,952 abnormal sessions from the test set. The descriptive statistics, presented in Table 2, revealed statistically significant differences in the coefficient distributions between the two classes.

Table 2
Descriptive statistics of predicted coefficients

Coefficient	Class	Median	Std.Dev.	Min	Max
α	Normal	0.0603	0.0325	0.0200	1.0000
	Anomaly	0.0764	0.0256	0.0276	1.0000
β	Normal	0.2243	0.0203	0.1730	1.0000
	Anomaly	0.2514	0.0160	0.1941	0.7318
γ	Normal	0.7386	0.0224	0.2732	0.7533
	Anomaly	0.7348	0.0091	0.5205	0.7532
ϕ	Normal	0.8434	0.0225	0.8127	1.0000
	Anomaly	0.8563	0.0159	0.8196	1.0000

Anomalous traffic is characterised by a higher median value for coefficient α (0.0764 vs. 0.0603 for normal), representing the intrinsic potential for attack growth. It also shows a higher median value for β (0.2514 vs. 0.2243), signifying a more intense interaction with the defence system. These statistical differences are visualised in the box plots shown in Figure 2. In the figure, the distributions for the anomaly class are visibly shifted towards higher values for coefficients α and β compared to the regular class, providing strong graphical evidence for the statistical findings. Furthermore, histograms of the coefficients confirm that the very shapes of the distributions differ between the two classes. For instance, the distribution of the α coefficient for the anomalous class is skewed to the right, indicating a prevalence of higher values that correspond to greater attack potential.

Conversely, the coefficients γ (defence adaptation rate) and ϕ (defence cost/decay) show less pronounced, though still informative, differences. The median γ values are nearly identical for both classes, suggesting that the rate of defensive adaptation captured by the model is not a primary distinguishing feature in this dataset. However, the slightly higher median ϕ for anomalous traffic (0.8563 vs. 0.8434) is noteworthy. This could imply that interactions classified as anomalous are associated with a higher 'cost of defence' or a faster rate of obsolescence for the responding security measures.

Taken together, these results demonstrate that the model has successfully learned to assign a distinct 'dynamic signature'—a unique vector of (α, β, γ, ϕ) coefficients—to different classes of network traffic. The systematic variations in these signatures, particularly in the attack-related parameters, provide strong evidence that the model captures the underlying behavioural characteristics of network sessions, moving beyond superficial pattern matching to a more profound, dynamic risk assessment.

To validate the practical utility of these predicted coefficients, we simulated the Lotka-Volterra dynamics using parameter sets generated by the model for both regular and anomalous sessions. The simulations revealed distinctly different behavioural patterns. As shown in Figure 3, parameters predicted for regular traffic tend to generate stable, controlled trajectories where the "attack" and "defence" levels remain in a bounded, cyclical balance. In stark contrast, the simulations for parameters typical of anomalous traffic, shown in Figure 4, more often lead to unstable scenarios characterised by a rapidly increasing "attacks," indicating a system state escalating out of control.

This crucial result confirms that the predicted coefficients carry meaningful information about network activity's potential risk and dynamic behaviour. By translating raw network data into dynamic parameters, the model offers a much deeper insight than traditional intrusion detection methods can provide, enabling a forward-looking assessment of threat evolution.

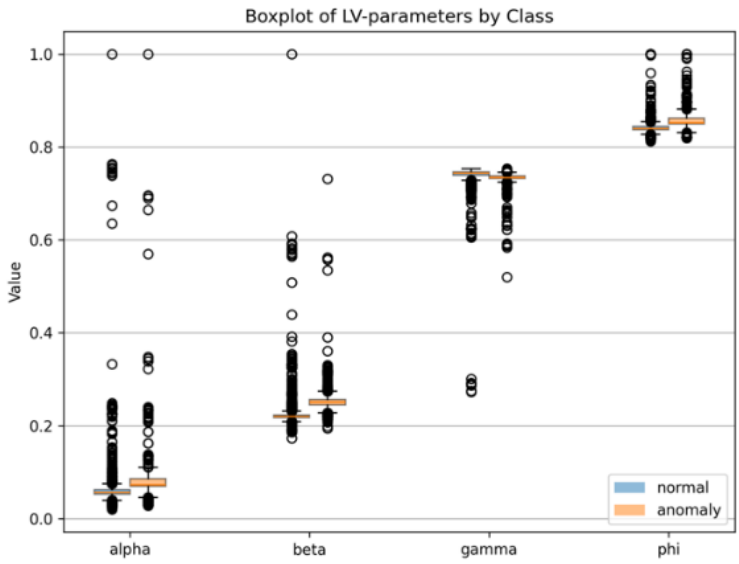

Figure 2: Box plots of predicted

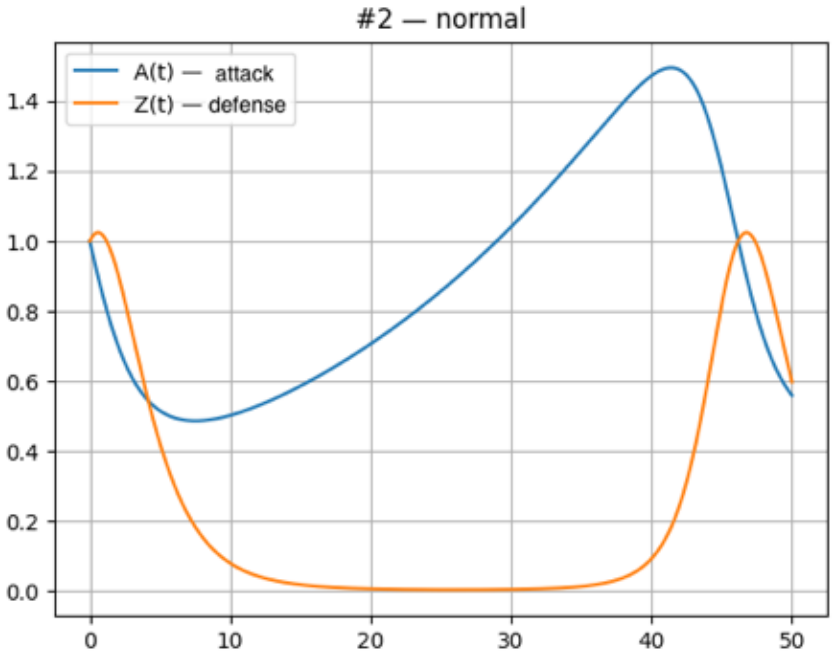

Figure 3: Example of Lotka–Volterra dynamics simulation (normal)

It is essential to acknowledge the study's limitations, which also highlight clear directions for future research. The age of the NSL-KDD dataset is a primary constraint, as it may not fully represent the complexity and signature of modern, sophisticated cyber threats [4; 5]. The Lotka-Volterra model itself, while a powerful analogy, is a simplification of the highly complex, multi-faceted nature of real-world attack-defence interactions, and the specific method for operationalising its coefficients could be further refined [6]. Additionally, while the neural network performs well, the interpretability of its internal decision-making process requires further investigation, which is a common challenge in the field of deep learning [2]. These limitations are not seen as detriments but valuable starting points for improving and extending the proposed approach in future work.

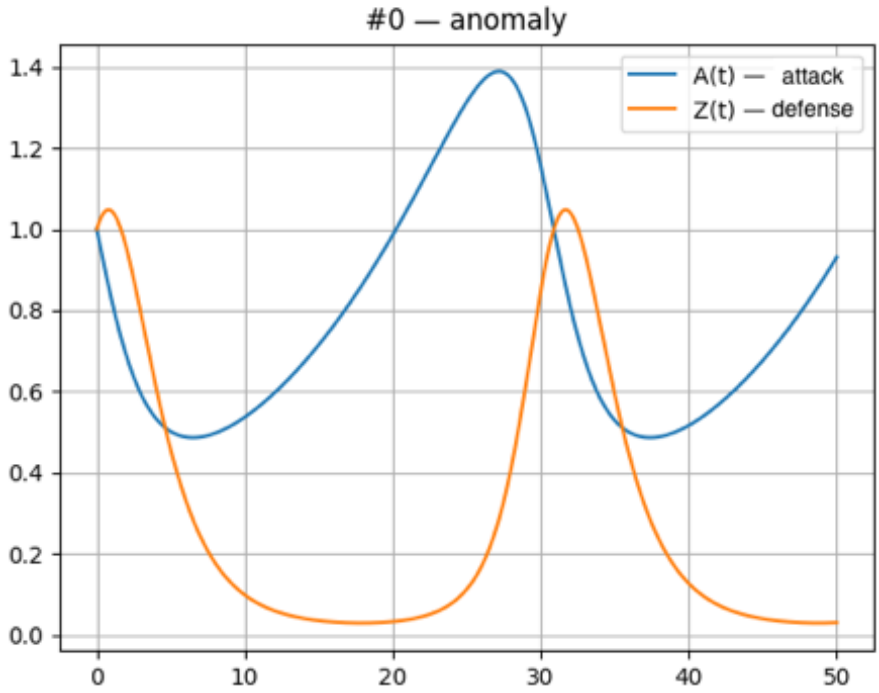

Figure 4: Example of simulation of Lotka–Volterra dynamics (anomaly)

Figure 4 visualises the dynamics for a typical anomalous session. It illustrates an unstable trajectory where the attack level, A(t), shows escalating oscillations, signifying a growing and uncontained threat. This simulation confirms that the model correctly associates anomalous traffic with high-risk, unstable system behaviours, thus validating its ability to forecast the potential evolution of a threat.

4. Conclusion

This paper successfully developed and validated a novel hybrid neural model that integrates deep learning for network traffic classification with predicting parameters for the Lotka-Volterra dynamic model. We demonstrated that the model accurately distinguishes between normal and anomalous traffic and, more importantly, quantifies the underlying dynamics of the "attack-defence" interaction through the predicted coefficients.

The core contribution of this work lies in its departure from traditional, static intrusion detection. The analysis of the predicted coefficients and the subsequent system dynamics simulation confirmed that the model effectively captures the distinct nature of normal and anomalous activity. By predicting dynamic parameters, our approach allows for an assessment that is not limited to identifying a threat's presence but extends to forecasting its potential development. This provides a richer, more nuanced understanding of cyber risks than conventional intrusion detection methods, moving the paradigm from simple detection to prognostic risk assessment.

The practical significance of this research lies in its potential to form the basis for more advanced and informative decision support systems in cybersecurity. However, further development is essential for its practical application. Future work should focus on several key areas: validating the model on larger, more contemporary datasets to ensure its relevance against modern threats; researching alternative or more complex dynamic models beyond the classical Lotka-Volterra framework; enhancing the interpretability of the neural network's predictions; and working towards the integration of this approach into real-time monitoring systems. In conclusion, this study represents a significant step towards creating more intelligent and forward-looking cybersecurity systems that can not only react to threats but also anticipate their evolution.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] [1] Abbasi, M., Shahraki, A., & Taherkordi, A. (2021). Deep learning for network traffic monitoring and analysis (NTMA): A survey. Computer Communications, 170, 19–41. https://doi.org/10.1016/j.comcom.2021.01.021

[2] [2] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press. http://www.deeplearningbook.org

[3] [3] Kushnerov, O., Murr, P., Herasymov, S., Milevskyi, S., Melnyk, M., & Golovashych, S. (2024). Application of neural networks for network traffic monitoring and analysis. In 2024, 8th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT) (pp. 1–8). IEEE. https://doi.org/10.1109/ISMSIT63511.2024.10757251

[4] [4] McHugh, J. (2000). Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations performed by Lincoln Laboratory. ACM Transactions on Information and System Security, 3(4), 262–294.

[5] [5] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. (2009). A detailed analysis of the KDD CUP 99 data set. Second IEEE Symposium on Computational Intelligence for Security and Defence Applications (CISDA).

[6] [6] Yevseiev, S., et al. (2021). Development of a method for assessing the security of cyber-physical systems based on the Lotka–Volterra model. Eastern-European Journal of Enterprise Technologies, 5(113), 30–47.