HYBRID MACHINE LEARNING-BASED DIGITAL TWINS FOR CRITICAL INFRASTRUCTURE SYSTEMS: MODELING, ANOMALY DETECTION AND RISK ASSESSMENT

Olena Moiseienko[1]

DOI: https://doi.org/10.30525/978-9934-26-673-7-3

Abstract. Modelling of complex information systems, processes, and cyber-physical interactions, i.e. their study through the construction of analytical and computational representations that reproduce their essential properties, is a powerful tool for understanding and managing modern digital infrastructures. In conditions of increasing cyber threats and hybrid warfare, particularly affecting critical infrastructure systems, modelling becomes a key instrument for ensuring system resilience, adaptability, and security. *The purpose* of the paper is to develop a hybrid digital twin framework for cybersecurity-oriented infrastructure systems, which integrates multiple analytical paradigms into a unified model capable of capturing dynamic, heterogeneous, and adversarial environments. The solution of this research problem determines the logic of the presentation of the material: systematisation of existing modelling approaches and identification of their limitations; theoretical substantiation of hybrid modelling based on the integration of time-series analysis, event-driven stochastic processes, and machine learning; development of a formal model of system behaviour; and construction of an architectural framework for the implementation of the proposed digital twin.

The methodology of the study is based on general scientific methods of analysis and synthesis, abstraction and formalisation, as well as computational methods, including statistical modelling, time-series analysis, and machine learning techniques. Special attention is given to stochastic modelling of event streams using self-exciting processes, which allow capturing temporal dependencies and cascading effects characteristic of cyberattacks. The integration of heterogeneous data sources is achieved

[1] Candidate of Technical Sciences, Associate Professor,
Department of Computer Systems and Networks,
Ivano-Frankivsk National Technical University of Oil and Gas, Ukraine

through feature transformation and hybrid fusion mechanisms with adaptive weighting, ensuring the consistency and flexibility of the model in non-stationary environments.

The results of the study demonstrate that the proposed hybrid digital twin provides a coherent theoretical and architectural framework for modelling cybersecurity-oriented systems. The model enables the integration of continuous system states and discrete event streams into a unified representation, supporting anomaly detection and risk estimation. The analysis shows that combining multiple modelling paradigms improves the ability to capture complex system behaviour compared to traditional approaches.

Practical implications. The proposed framework can be applied in real-world IT infrastructures, including server systems, network environments, and distributed client-server architectures, for continuous monitoring, anomaly detection, and risk assessment. It is particularly relevant for critical infrastructure systems operating under conditions of uncertainty and evolving cyber threats.

Value/originality. The originality of the study lies in the development of a unified hybrid digital twin framework that integrates time-series modelling, event-driven stochastic processes, and machine learning within a single adaptive architecture. The proposed approach extends the concept of digital twins to cybersecurity applications and provides a foundation for the development of intelligent, resilient, and scalable systems for critical infrastructure protection.

1. Introduction

The rapid digitalization of critical infrastructure systems, including energy networks, telecommunications, cloud platforms, and governmental information systems, has fundamentally transformed modern societies. However, this transformation has simultaneously introduced new vulnerabilities, particularly in the context of cyber warfare, hybrid threats, and large-scale distributed attacks. For countries operating under conditions of ongoing military conflict, such as Ukraine, ensuring the resilience and security of cyber-physical and information systems has become a strategic priority.

Recent studies on the resilience of critical infrastructure under hybrid and cyber warfare conditions highlight the increasing vulnerability of energy systems, governmental platforms, and communication networks to coordinated attacks [1; 2]. In particular, research on cyber resilience of power grids and national digital infrastructures demonstrates that modern threats combine cyber, physical, and informational components, requiring integrated modeling approaches. These findings emphasize the necessity of adaptive and hybrid analytical frameworks capable of capturing both system dynamics and adversarial behavior.

Traditional approaches to monitoring and managing infrastructure systems rely on static models, rule-based detection mechanisms, or isolated machine learning algorithms [2]. While such approaches can be effective in controlled environments, they exhibit significant limitations in real-world scenarios characterized by non-stationarity, adversarial behavior, and high-dimensional heterogeneous data streams. In particular, modern cyber threats demonstrate adaptive behavior, evolving attack patterns, and temporal dependencies that cannot be adequately captured by conventional models.

In recent years, the concept of the digital twin has emerged as a promising paradigm for modeling complex systems. A digital twin represents a dynamic virtual replica of a physical or cyber system, continuously updated using real-time data. This paradigm enables predictive analytics, anomaly detection, and decision support by simulating system behavior under various conditions. However, existing digital twin implementations are predominantly focused on industrial and engineering applications and often rely on either physics-based models or purely data-driven approaches, without sufficiently addressing the stochastic and event-driven nature of cyber threats [2; 3].

Simultaneously, advances in machine learning and artificial intelligence have enabled the development of sophisticated models for anomaly detection, classification, and prediction [4; 5; 6]. Techniques such as ensemble learning, deep neural networks, and probabilistic modeling have demonstrated high performance across a range of domains, including cybersecurity. Nevertheless, most existing machine learning approaches treat data as independent observations or rely on simplified temporal dependencies, neglecting the intrinsic structure of event-driven processes.

In the context of cybersecurity, the analysis of event streams–such as system logs, network traffic, and user behavior–is of critical importance. These streams exhibit complex temporal patterns, including clustering, self-excitation, and cascading effects, where one event increases the probability of subsequent events. Such behavior is well captured by point process models, particularly self-exciting processes, which have been successfully applied in domains such as finance and seismology but remain underutilized in cybersecurity applications [7].

Another important challenge is the integration of heterogeneous data sources. Infrastructure systems generate multiple types of data, including continuous time-series measurements, discrete event logs, and contextual information. Existing approaches typically process these data types separately, leading to suboptimal performance and limited interpretability.

To address these limitations, this chapter proposes a unified methodological framework for constructing hybrid digital twins of critical infrastructure systems, combining machine learning, stochastic process modeling, and event-driven analytics. The proposed approach integrates:

– dynamic system modeling for continuous state variables;

– event-based modeling using self-exciting stochastic processes;

– machine learning models for high-dimensional feature extraction and prediction.

A particular emphasis is placed on cybersecurity applications, including anomaly detection in system logs, identification of malicious activity in network traffic, and adaptive risk assessment under evolving threat conditions.

The proposed framework builds upon and extends previous research conducted by the author in the areas of phishing detection, anomaly detection in log data, and ensemble-based reliability assessment. In particular, prior work has demonstrated the effectiveness of hybrid feature selection methods (combining mutual information, principal component analysis, and genetic algorithms), as well as adaptive ensemble models incorporating probabilistic filtering techniques. These results indicate that combining statistical and machine learning approaches can significantly improve detection accuracy and robustness compared to individual models [8–12].

The main contributions of this chapter are as follows:

1. A formalized definition of a hybrid digital twin model for critical infrastructure systems that integrates continuous dynamics, event-driven processes, and machine learning components.

2. A unified mathematical framework for anomaly detection and risk assessment based on heterogeneous data sources.

3. An architecture for implementing digital twins in cybersecurity applications, including log analysis and network traffic monitoring.

4. An empirical validation of the proposed approach using real and synthetic datasets, demonstrating improved performance compared to conventional methods.

To position the proposed approach within the existing body of research, a critical analysis of current modeling paradigms is presented in the following section.

2. Analysis of Existing Approaches

2.1. Data-Driven Approaches Based on Machine Learning

Data-driven approaches constitute one of the most widely used paradigms for modeling complex systems and detecting anomalies. These approaches rely on machine learning algorithms trained on historical data to identify patterns, classify system states, and predict future behavior.

Classical machine learning models, such as logistic regression, decision trees, random forests, and support vector machines, have been extensively applied in cybersecurity tasks, including intrusion detection, malware classification, and phishing detection [13]. For instance, ensemble-based models, particularly random forests and gradient boosting methods, have demonstrated strong performance in detecting phishing websites and malicious network activity due to their ability to capture nonlinear relationships and handle high-dimensional feature spaces [5; 6].

Previous studies conducted by the author have shown that the use of hybrid feature selection techniques – combining mutual information, principal component analysis, and genetic algorithms–can significantly reduce feature dimensionality while preserving predictive performance [13; 14]. In experiments on benchmark datasets, such approaches achieved comparable or improved classification accuracy with reduced computational

cost, highlighting the importance of feature engineering in data-driven models.

Deep learning approaches, including recurrent neural networks (RNNs), long short-term memory (LSTM) networks, and convolutional neural networks (CNNs), have further extended the capabilities of data-driven modeling [15]. These models are particularly effective in capturing temporal dependencies and complex patterns in sequential data. In cybersecurity, LSTM-based models have been applied to detect anomalies in network traffic and system logs, while autoencoders have been used for unsupervised anomaly detection.

Despite their advantages, data-driven approaches exhibit several limitations:

– Lack of interpretability, particularly in deep learning models;

– Dependence on labeled data, which is often scarce or expensive to obtain in cybersecurity contexts;

– Limited robustness to distribution shifts, especially under adversarial conditions;

– Inability to explicitly model event dependencies, such as cascading attacks or correlated anomalies.

These limitations motivate the integration of machine learning with other modeling paradigms.

2.2. Physics-Based and Rule-Based Models

Physics-based models represent one of the earliest and most theoretically grounded approaches to modeling complex systems. Their development is rooted in classical control theory, where system behavior is described using differential equations derived from physical laws. In industrial applications, such models have proven effective for simulation, prediction, and optimization, particularly in systems with well-defined physical properties [1].

In the context of digital twins, physics-based models are frequently used to replicate the behavior of mechanical, energy, or transport systems, allowing for high interpretability and explainability of results. Their main advantage lies in the explicit representation of causal relationships between system variables, which makes them suitable for safety-critical applications.

However, when considering modern IT infrastructures and cybersecurity environments, the applicability of physics-based models becomes

significantly limited. Unlike physical systems, information systems are characterized by logical interactions, user-driven behavior, and external adversarial influence, which cannot be easily expressed in terms of deterministic equations. As a result, constructing accurate physics-based models for such systems requires extensive domain knowledge and often becomes impractical.

An alternative class of approaches is formed by rule-based and signature-based systems, which have historically been widely used in cybersecurity. These systems rely on predefined rules or patterns to identify known threats and anomalies. Intrusion detection systems (IDS), for example, commonly employ signature matching techniques to detect malicious activities [7].

While rule-based approaches offer high interpretability and efficiency in detecting known attack patterns, they suffer from a fundamental limitation: their inability to generalize beyond predefined scenarios. In rapidly evolving threat landscapes, where attackers continuously modify their strategies, such systems quickly become obsolete.

Consequently, both physics-based and rule-based approaches exhibit limited adaptability and scalability in dynamic and adversarial environments. This limitation necessitates the exploration of more flexible modeling paradigms capable of handling uncertainty, heterogeneity, and evolving system behavior.

2.3. Event-Driven and Stochastic Process Models

Event-driven models represent an important class of approaches for describing systems in which behavior is governed not only by continuous dynamics but also by discrete events occurring over time. Such models are particularly relevant in domains where system evolution is determined by irregular and often unpredictable interactions.

In cybersecurity-oriented environments, the role of event-driven modeling becomes especially significant. Unlike traditional engineering systems, where changes occur gradually, many cybersecurity phenomena manifest as sequences of discrete events. These include login attempts, network requests, error messages, and malicious activities, which collectively form complex temporal patterns.

A key characteristic of such event streams is that events are rarely independent. In practice, they often exhibit temporal clustering and causal

dependencies. For instance, a series of failed authentication attempts may indicate a brute-force attack, while bursts of network traffic may signal coordinated intrusion attempts. This behavior suggests that the occurrence of one event can increase the likelihood of subsequent events, forming cascades.

To capture these dependencies, stochastic point processes provide a suitable mathematical framework. Among them, self-exciting processes, such as Hawkes processes, are particularly well suited for modeling systems with memory and temporal causality [7].

The intensity of event occurrence is then modeled as:

$$\lambda(t) = \mu + \sum_{t_i < t} \alpha e^{-\beta(t-t_i)},$$

where λ(t) – the intensity function;

μ – the baseline intensity, and the summation term represents the influence of past events.

This formulation enables the model to capture not only the frequency of events but also their temporal structure and interdependence. As a result, it becomes possible to distinguish between random fluctuations and structured patterns indicative of coordinated behavior.

To further justify the selection of self-exciting processes for modeling cybersecurity events, it is important to compare them with classical stochastic models traditionally used in temporal analysis. In particular, Poisson processes and Markov models are widely applied in various domains due to their mathematical simplicity and well-established theoretical foundations.

However, the applicability of these models to cybersecurity scenarios remains limited, primarily due to their underlying assumptions regarding event independence and memory constraints. A comparative analysis of these stochastic modeling approaches is presented in Table 1.

As shown in Table 1, classical stochastic models such as Poisson processes assume independence between events, which makes them unsuitable for capturing the cascading nature of cyberattacks. Markov models partially address this limitation by introducing state transitions; however, they are constrained by limited memory and often fail to capture long-range dependencies.

In contrast, Hawkes processes explicitly model the influence of past events on future occurrences, allowing the representation of temporal clustering and cascading effects. This property makes them particularly suitable for modeling coordinated cyberattacks, intrusion attempts, and other complex event-driven phenomena.

Table 1

Comparison of stochastic processes for modeling cybersecurity events

Model	Assumptions	Memory	Event Dependency	Ability to Model Attack Cascades	Suitability for Cybersecurity
Poisson Process	Independent events	None	No	No	Low
Markov Model	State-based transitions	Limited (short memory)	Partial	Limited	Medium
Hawkes Process	Self-exciting events	Long-term dependency	Yes	Yes	High

Therefore, the use of self-exciting processes provides a theoretically justified and practically effective foundation for the event-driven component of the proposed hybrid digital twin model.

Hawkes processes explicitly model event interdependence and temporal clustering, making them particularly effective for representing coordinated and evolving cyberattacks.

Despite their strong theoretical foundation, event-driven stochastic models remain underutilized in cybersecurity applications. One of the main reasons is the difficulty of integrating such models with high-dimensional feature spaces and machine learning techniques. Most existing approaches either simplify event dependencies or ignore them altogether.

This limitation creates a gap between theoretical capabilities and practical implementations. Addressing this gap requires the integration of event-driven models with complementary approaches, such as time-series analysis and machine learning, which is further developed in the subsequent sections.

2.4. Digital Twin Paradigm

The concept of the digital twin has emerged as a transformative paradigm in the modeling and management of complex systems. Originally introduced in the context of industrial engineering, the digital twin is commonly defined as a virtual representation of a physical system that is continuously updated using real-time data [1].

Over time, this concept has evolved beyond its initial scope and is now applied in a wide range of domains, including manufacturing, healthcare, transportation, and smart cities. The core idea underlying digital twins is the integration of data, models, and analytics into a unified framework that enables monitoring, simulation, and prediction of system behavior.

To better position the proposed approach within the existing body of research, it is necessary to analyze the architectural characteristics of current digital twin implementations. Although numerous frameworks have been developed in industrial and cyber-physical domains, they differ significantly in terms of data integration, event handling capabilities, and adaptability to dynamic environments.

In particular, most existing digital twin architectures are primarily designed for physical systems and rely on continuous sensor data, with limited support for discrete event modeling and cybersecurity-oriented analysis. Furthermore, the ability of these systems to operate under non-stationary and adversarial conditions remains constrained.

A comparative analysis of representative digital twin approaches is presented in Table 2, highlighting their key characteristics and limitations in the context of cybersecurity applications.

Table 2

Comparison of Digital Twin Architectures

Approach	Data Type	Event Handling	Adaptivity	Cybersecurity Suitability
Industrial DT	Sensor	No	Low	Low
Physics-based	Continuous	No	Low	Low
ML-based	Mixed	Limited	Medium	Medium
Proposed	Hybrid	Yes	High	High

As shown in Table 2, existing digital twin architectures demonstrate strong performance in modeling physical processes and continuous system

dynamics. However, they exhibit significant limitations when applied to cybersecurity-oriented environments, where system behavior is driven by discrete events, adversarial actions, and heterogeneous data streams.

In particular, the lack of explicit event-driven modeling and limited adaptability to evolving conditions restrict the applicability of traditional digital twin frameworks for detecting complex cyber threats. While machine learning-based approaches partially address these issues, they often lack interpretability and fail to capture temporal causality.

These observations confirm the necessity of developing a hybrid digital twin model that integrates continuous dynamics, event-driven processes, and machine learning techniques within a unified and adaptive framework. The proposed approach aims to address these limitations by combining multiple modeling paradigms and introducing an adaptive integration mechanism.

From a theoretical perspective, a digital twin can be viewed as a dynamic mapping between the real system and its virtual counterpart, where the state of the virtual model is continuously synchronized with incoming data streams. This synchronization enables not only descriptive analytics but also predictive and prescriptive capabilities.

Despite these advantages, existing digital twin implementations exhibit several limitations when applied to cybersecurity-oriented systems. First, the majority of current digital twin architectures are designed for physical systems and rely heavily on sensor data and physics-based models [2]. Such approaches are not directly transferable to IT infrastructures, where system behavior is driven by logical operations and user interactions rather than physical processes.

Second, traditional digital twin models typically focus on continuous data streams while neglecting the role of discrete events. However, in cybersecurity, event-driven behavior is of primary importance, as attacks and anomalies often manifest through sequences of discrete actions rather than gradual changes in system metrics.

Third, the integration of heterogeneous data sources remains an open challenge. Infrastructure systems generate diverse data types, including logs, network traffic, performance metrics, and contextual information. Existing digital twin frameworks often process these data streams independently, resulting in fragmented system representations.

These limitations highlight the need for a redefinition of the digital twin concept in the context of cybersecurity. In particular, there is a need for hybrid digital twins that combine continuous system modeling, event-driven processes, and machine learning techniques within a unified framework.

2.5. Hybrid Approaches

Hybrid approaches have emerged as a natural response to the limitations of single-paradigm models. The fundamental idea behind hybrid modeling is to combine multiple analytical techniques in order to leverage their complementary strengths while mitigating their individual weaknesses.

In the context of cybersecurity and complex system analysis, hybrid approaches often involve the integration of statistical methods, machine learning algorithms, and domain-specific knowledge. Such integration enables the construction of models that are both expressive and robust, capable of capturing nonlinear relationships while maintaining interpretability [5].

One important direction in hybrid modeling is the use of ensemble methods, where multiple models are combined to produce a final prediction. Ensemble techniques, such as Random Forest and Gradient Boosting, have demonstrated strong performance in various classification and anomaly detection tasks due to their ability to reduce variance and improve generalization [6].

The author's previous research has further extended this idea by introducing adaptive ensemble models with dynamic weighting mechanisms [8; 9; 10; 11; 12]. In particular, the use of Kalman filtering for adaptive weight adjustment has shown promising results in non-stationary environments, where the relevance of individual models may change over time. This approach allows the system to dynamically emphasize the most informative components, improving both accuracy and stability.

Another important aspect of hybrid modeling is feature engineering. The combination of mutual information, principal component analysis, and genetic algorithms enables effective feature selection, reducing dimensionality while preserving essential information. Experimental results have demonstrated that such hybrid feature selection techniques can significantly improve model performance in cybersecurity tasks.

Despite these advances, existing hybrid approaches are often developed in an ad hoc manner, without a unified theoretical foundation. In many cases, the integration of different models is performed heuristically, without a formal framework that defines how these components interact.

This lack of formalization limits the scalability and generality of hybrid methods. In particular, there remains a need for a systematic approach that integrates time-series modeling, event-driven processes, and machine learning into a coherent structure. Addressing this gap is essential for the development of next-generation digital twin systems capable of operating in complex and adversarial environments.

An essential aspect of hybrid modeling is the strategy used to combine heterogeneous data sources and model outputs. The effectiveness of hybrid systems largely depends on how information from different components is integrated, particularly in environments characterized by high dimensionality and non-stationarity.

In the context of anomaly detection and cybersecurity, two principal fusion paradigms are commonly distinguished: feature-level fusion and decision-level fusion. Each of these approaches offers specific advantages and limitations depending on the nature of the data and the modeling objectives.

To clarify these differences, a comparative analysis of fusion strategies is presented in Table 3.

As shown in Table 3, feature-level fusion enables the integration of heterogeneous data sources into a unified representation, allowing the model to capture complex interactions between variables. However, this approach often leads to high-dimensional feature spaces and increased computational complexity.

Decision-level fusion, on the other hand, provides a modular framework in which individual models can be developed and optimized independently. This makes it particularly suitable for real-time systems and environments with evolving data distributions. Nevertheless, it may fail to capture deeper relationships between features.

The proposed hybrid approach combines the advantages of both strategies by integrating feature-level representations with decision-level aggregation. In particular, the use of adaptive weighting mechanisms allows the system to dynamically adjust the contribution of individual components based on current conditions.

Table 3

Comparison of fusion strategies in hybrid modeling

Fusion Level	Description	Advantages	Limitations	Applicability in Cybersecurity
Feature-level fusion	Integration of raw or preprocessed features into a single feature space before modeling	Captures interactions between heterogeneous features; enables unified modeling	High dimensionality; risk of overfitting; requires careful feature engineering	Effective for structured data and combined datasets
Decision-level fusion	Combination of outputs from multiple models (e.g., voting, weighting)	Modular; flexible; robust to noise; allows independent model optimization	May lose fine-grained feature interactions; requires calibration of outputs	Suitable for multi-model and real-time systems
Hybrid fusion (proposed)	Combination of feature-level and decision-level integration with adaptive weighting	Balances expressiveness and robustness; supports dynamic environments; improves adaptability	Increased complexity; requires adaptive mechanisms	Highly suitable for cybersecurity and non-stationary environments

This integration strategy is especially relevant for cybersecurity applications, where system behavior is influenced by multiple heterogeneous factors and where adaptability to changing threat patterns is essential.

2.6. Research Gap

Based on the analysis presented above, the following research gaps can be identified:

– The absence of a unified framework that integrates continuous system dynamics, event-driven processes, and machine learning.

– Limited application of stochastic point processes in cybersecurity modeling.

– Insufficient adaptability of existing models to non-stationary and adversarial environments.

– Lack of scalable and interpretable models for real-time risk assessment in critical infrastructure systems.

These gaps motivate the development of a hybrid digital twin framework, which is presented in the subsequent sections.

Addressing the identified limitations requires the development of a fundamentally new modeling approach, which forms the basis of the proposed framework.

Scientific Novelty of the Proposed Approach

The proposed model introduces several novel contributions:

1. A unified hybrid formulation integrating time-series, event-driven, and machine learning components.

2. The application of self-exciting processes for cybersecurity-oriented digital twins.

3. An adaptive weighting mechanism for dynamic environments.

4. A generalized framework applicable across multiple domains.

To formally define the proposed approach, it is necessary to establish a theoretical foundation that integrates system dynamics, event-driven processes, and anomaly modeling.

3 Theoretical Foundations of Digital Twin Modeling for Cyber-Physical and Information Systems

3.1. Formal Representation of Infrastructure Systems

The modeling of modern infrastructure systems, particularly those operating in cybersecurity-sensitive environments, requires a generalized mathematical framework capable of capturing both deterministic and stochastic behavior. Unlike classical engineering systems, IT infrastructures are characterized by high variability, partial observability, and continuous interaction with external agents, including users and potential attackers.

In this context, the system cannot be adequately described using static models or purely deterministic formulations. Instead, it is necessary to adopt a dynamic representation that reflects the temporal evolution of system states under the influence of internal processes and external inputs.

Let the system state at time t be defined as a multidimensional vector:

$$X(t) = \{x_1(t), x_2(t), \ldots, x_n(t)\}, \quad (1)$$

where each component represents a measurable system characteristic, such as CPU load, memory utilization, request rate, latency, or error frequency.

The representation of infrastructure systems as dynamic processes is rooted in the theory of nonlinear dynamical systems. However, unlike classical physical systems, IT infrastructures are characterized by high variability, partial observability, and external stochastic influences.

This necessitates the use of generalized state-space models that can incorporate both deterministic and stochastic components [16; 17]. In this context, the system evolution is described as:

$$X(t+1) = f(X(t), U(t), \theta) + \varepsilon(t),$$

where $U(t)$ denotes control inputs (e.g., configuration changes, user actions);

θ represents system parameters;

$\varepsilon(t)$ – stochastic noise.

This formulation provides a flexible framework capable of representing a wide range of system behaviors. At the same time, it highlights an important limitation: the model captures continuous dynamics but does not explicitly account for discrete events, which play a crucial role in cybersecurity scenarios.

3.2. Event-Driven Dynamics and Temporal Dependencies

A defining characteristic of cybersecurity-oriented systems is the presence of discrete events that significantly influence system behavior. These events include login attempts, network requests, system errors, and malicious activities. Unlike continuous state variables, such events are inherently irregular and often exhibit strong temporal dependencies.

In practice, many cyber incidents are not isolated but occur in bursts or sequences, where one event increases the likelihood of subsequent events. For example, repeated failed login attempts may indicate a brute-force attack, while a sudden increase in network requests may signal a distributed denial-of-service attack.

To capture these phenomena, it is necessary to move beyond independent event models and consider stochastic processes with memory. A suitable framework for this purpose is provided by self-exciting point processes, which explicitly model the influence of past events on future occurrences.

Let the event stream be defined as:

$$E = \{(t_i, c_i, a_i)\}_{i=1}^{N}, \quad (2)$$

where t_i – the occurrence time;

c_i – the event class;

a_i represents event attributes.

A key property of cybersecurity event streams is self-excitation, where the occurrence of an event increases the likelihood of subsequent events. This phenomenon can be modeled using a self-exciting point process:

$$\lambda(t) = \mu + \sum_{t_i < t} \alpha e^{-\beta(t - t_i)},$$

where $\lambda(t)$ – the event intensity;

μ – the baseline intensity;

α and β control the influence of past events.

This formulation allows the system to represent temporal clustering and cascading effects, which are typical of coordinated cyberattacks. Importantly, it introduces a notion of temporal causality, enabling the model to distinguish between isolated events and structured attack patterns.

3.3. Anomaly Representation in Hybrid Systems

Anomaly detection is a central task in monitoring infrastructure systems. However, in hybrid environments, anomalies cannot be defined solely in terms of deviations in continuous variables or irregularities in event sequences. Instead, anomalies often emerge from the interaction between these two domains.

For example, a moderate increase in CPU load may be considered normal under high demand, but it may indicate an attack if accompanied by abnormal login patterns or network activity. This illustrates the need for a unified representation that incorporates both state and event information.

Let the anomaly indicator be defined as:

$$A(t) \in \{0,1\},$$

where $A(t)$=1 denotes an anomalous state.

The anomaly detection function is defined as:

$$A(t) = g(X(t), E(t)),$$

where g(·) – a decision function integrating state and event information.

As an example, the function g(·) can be implemented using a logistic regression model, where the anomaly score is computed as a weighted combination of features derived from both system state and event data.

Traditional approaches typically consider only *X(t)* or *E(t)*, leading to incomplete representations of system behavior.

This formulation generalizes traditional anomaly detection approaches by transforming the problem into a multi-source inference task. As a result, the detection process becomes more robust to noise and more sensitive to complex patterns that cannot be captured by single-source models.

3.4. Risk Modeling in Cybersecurity Systems

While anomaly detection provides a binary classification of system states, practical cybersecurity applications require a more nuanced understanding of system behavior. In particular, it is essential to estimate the likelihood of adverse events and to quantify the associated risk.

Risk modeling enables the transition from reactive detection to proactive decision-making. By assessing the probability of system compromise or failure, it becomes possible to prioritize responses and allocate resources more effectively.

The risk function is defined as:

$$R(t) = P\left(A(t) = 1 \mid X(t), E(t)\right).$$

This probabilistic formulation reflects the inherent uncertainty of cybersecurity environments, where incomplete information and adversarial behavior are common.

Importantly, the integration of risk estimation into the digital twin framework allows the system to operate not only as a monitoring tool but also as a predictive and decision-support system.

3.5. Limitations of Existing Models

The analysis presented above highlights several fundamental limitations of existing approaches to system modeling and anomaly detection.

First, models based solely on continuous state variables fail to capture the discrete and often abrupt nature of cyber events. Second, event-driven models, while effective in representing temporal dependencies, typically lack the ability to incorporate high-dimensional feature spaces. Third, machine learning models, although powerful, often operate as black boxes and do not explicitly model temporal causality.

As a result, each class of models provides only a partial view of system behavior. This fragmentation leads to reduced accuracy, limited interpretability, and decreased robustness in real-world applications.

These limitations underscore the need for a unified framework that integrates multiple modeling paradigms into a coherent system. The development of such a framework is the focus of the next section.

These considerations provide the theoretical basis for the development of a hybrid digital twin model, presented in the next section.

4. Proposed Hybrid Digital Twin Model for Cybersecurity-Oriented Infrastructure Systems:

4.1. Concept of a Hybrid Digital Twin

The concept of a digital twin has traditionally been associated with physical systems, where a virtual model replicates the behavior of a real-world object using sensor data and physical laws. However, in the context of cybersecurity and IT infrastructures, such an interpretation is insufficient. Unlike purely physical systems, IT environments are characterized by high-dimensional data, discrete event streams, and adversarial dynamics.

Therefore, in this work, a digital twin is reinterpreted as a hybrid analytical construct that combines continuous system dynamics, event-driven processes, and data-driven intelligence. This reinterpretation is motivated by the observation that cybersecurity phenomena cannot be adequately captured using a single modeling paradigm.

In particular, system behavior is influenced simultaneously by:

gradual changes in system metrics (e.g., load, latency);

discrete events (e.g., login attempts, errors, attacks);

complex nonlinear dependencies captured by machine learning models.

To account for these aspects, we define the digital twin as a mapping:

$$Z(t) = \Phi\left(X(t), E(t), \mathcal{M}\right),$$

where $X(t)$ – represents system state variables;

$E(t)$ – denotes the event stream;

M – a set of heterogeneous models.

The output $Z(t)$ represents a comprehensive system assessment, including anomaly and risk levels.

This formulation emphasizes that the digital twin is not merely a passive representation but an active analytical system capable of interpreting system behavior, detecting anomalies, and estimating risk.

While the proposed model defines the analytical structure of the digital twin, its practical implementation requires a well-defined system architecture.

4.2. Architecture of the Proposed Model

The architecture of the proposed digital twin is derived from the limitations identified above. Existing approaches typically rely on a single type of model – either machine learning, rule-based, or statistical–which leads to incomplete system representation.

To overcome this limitation, we adopt a modular hybrid architecture consisting of three complementary components:

1. Time-Series Component.
2. Event-Driven Component.
3. Machine Learning Component.

The rationale behind this decomposition is that each component captures a distinct aspect of system behavior:

– the time-series component models continuous dynamics;

– the event component captures temporal dependencies and cascades;

– the machine learning component identifies complex nonlinear relationships.

This separation allows each model to specialize while enabling integration at a higher level. Importantly, such an architecture supports extensibility, allowing additional components to be incorporated without altering the overall structure.

These components are integrated within a unified architecture.

Time-Series Modeling Component. The time-series component is responsible for capturing gradual changes in system state variables. In IT infrastructures, many critical indicators – such as CPU usage, memory consumption, request rate, and latency–evolve continuously over time and exhibit temporal correlations.

Ignoring these dependencies leads to loss of important information about system behavior. For example, a sudden increase in latency may not be anomalous if it follows a known pattern, but it may indicate a problem if it deviates from expected temporal dynamics.

To address this, the time-series component is defined as:

$$S_{ts}(t) = f_{ts}\left(X(t), X(t-1), \ldots, X(t-k)\right), \tag{3}$$

where k – the window size.

The choice of this formulation is motivated by classical time-series analysis and modern sequence modeling techniques. Depending on implementation, f_{ts} may represent:

- statistical models (e.g., moving averages);
- autoregressive models;
- recurrent neural networks.

The key requirement is that this component captures temporal continuity and trends, providing a baseline for normal system behavior.

Event Modeling Component. While time-series models capture continuous dynamics, they are insufficient for representing discrete events, which are central to cybersecurity. Events such as login failures, error bursts, or suspicious requests often occur in clusters and exhibit causal relationships.

To model these phenomena, we employ an event-driven approach based on self-exciting processes. The choice of this model is grounded in the observation that cyberattacks frequently exhibit cascading behavior: one event increases the likelihood of subsequent events.

The event component is defined through an intensity function:

$$S_{event}(t) = f_{event}(E(t)) = \lambda(t), \tag{4}$$

where $\lambda(t)$ – derived from a self-exciting process

$$\lambda(t) = \mu + \sum_{t_i < t} \alpha e^{-\beta(t-t_i)}. \tag{5}$$

This formulation allows the model to capture:

– temporal clustering;
– dependency between events;
– escalation patterns typical of coordinated attacks.

Unlike conventional anomaly detection methods, which treat events independently, this approach introduces temporal causality, making it particularly suitable for cybersecurity applications.

Machine Learning Component. The machine learning component complements the previous two by capturing complex nonlinear relationships that cannot be explicitly modeled using analytical or statistical methods.

In high-dimensional environments, such as network traffic analysis or log classification, the relationships between features are often intricate and

context-dependent. Machine learning models are well suited to handle such complexity.

The ML component is defined as:

$$S_{ml}(t) = f_{ml}\left(X(t), E(t)\right), \quad (6)$$

where f_{ml} – may represent Random Forest, Gradient Boosting, Neural Networks.

The choice of model is not restricted; however, ensemble methods such as Random Forest and Gradient Boosting are particularly suitable due to their robustness and interpretability [5; 6].

This component is also informed by prior research results [8-12], which demonstrated that:

– hybrid feature selection improves model performance,

– ensemble methods provide stable predictions in noisy environments.

Thus, the machine learning component serves as a data-driven inference layer, extracting latent patterns from combined state and event information.

4.3. Hybrid Integration Mechanism

The central novelty of the proposed approach lies in the integration of heterogeneous components into a unified representation.

Rather than relying on a single model, we construct a composite representation:

$$Z(t) = w_1 S_{ts}(t) + w_2 S_{event}(t) + w_3 S_{ml}(t), \quad (7)$$

where w_1, w_2, w_3 are adaptive weights.

To ensure normalization:

$$w_1 + w_2 + w_3 = 1, \quad w_i \geq 0. \quad (8)$$

Weights can be fixed, dynamically updated, learned via optimization.

The justification for this linear combination is twofold:

1. It provides a transparent and interpretable integration mechanism, allowing the contribution of each component to be analyzed.

2. It enables flexible adaptation, as weights can be adjusted dynamically depending on system conditions.

In contrast to static models, the proposed approach allows the system to emphasize different components under different conditions. For example:

– during stable operation, the time-series component may dominate;

– during attack bursts, the event component becomes more important;

– in complex scenarios, the ML component captures hidden dependencies.

This adaptive integration ensures robustness in non-stationary environments, which is critical for cybersecurity applications.

The resulting function $Z(t)$ represents a comprehensive system state indicator, which can be further used for anomaly detection and risk estimation. By integrating multiple sources of information, the digital twin provides a more accurate and robust representation compared to traditional approaches.

4.4. Adaptive Weighting Mechanism

A key limitation of static hybrid models is their inability to adapt to changing system conditions. In cybersecurity environments, system behavior is inherently non-stationary: attack patterns evolve, system loads fluctuate, and normal behavior may shift over time.

To address this challenge, the proposed model incorporates an adaptive weighting mechanism that dynamically adjusts the contribution of each component of the digital twin.

The weights $w_i(t)$ are updated based on the recent performance of individual components, allowing the system to prioritize the most informative signals under current conditions. Formally, the update rule is defined as:

$$w_i(t+1) = w_i(t) + \eta \cdot \Delta_i(t),$$

where η – the learning rate;

$\Delta_i(t)$ – reflects model performance.

This formulation is inspired by adaptive filtering and ensemble learning theory, where model weights are adjusted based on prediction error. In particular, similar ideas have been successfully applied in Kalman filtering-based ensemble models, where dynamic weighting improves robustness under uncertainty [15].

The adaptive mechanism enables the digital twin to respond to sudden attack bursts (increasing w_2), gradual system degradation (increasing w_1), complex nonlinear patterns (increasing w_3).

Unlike classical boosting methods, which iteratively reweight training samples in a static dataset, the proposed adaptive weighting mechanism operates in an online setting, dynamically adjusting model contributions

based on real-time performance. This allows the system to respond to non-stationary environments and evolving threat patterns, which are not adequately addressed by traditional boosting algorithms [10; 11].

Thus, the model remains effective under varying operational conditions.

4.5. Anomaly Detection in the Hybrid Model

Anomaly detection is one of the primary functions of the digital twin. Unlike traditional approaches that rely on a single source of information, the proposed model integrates multiple signals into a unified anomaly score.

The anomaly indicator is defined as:

$$A(t) = \begin{cases} 1, & Z(t) > \tau \\ 0, & \text{otherwise} \end{cases}, \tag{9}$$

where τ – a predefined threshold.

The use of the aggregated score *Z(t)* provides several advantages. First, it combines complementary information from time-series dynamics, event patterns, and machine learning predictions. Second, it reduces sensitivity to noise in individual components. Third, it enables more stable detection in non-stationary environments.

From a theoretical perspective, the anomaly detection problem is transformed from a single-model classification task into a multi-source inference problem, which improves both accuracy and robustness.

4.6. Risk Estimation

While anomaly detection provides a binary decision, practical systems require a continuous measure of system vulnerability. For this purpose, the digital twin estimates the risk level associated with the current system state.

The risk function is defined as:

$$R(t) = \sigma\left(Z(t)\right), \tag{10}$$

where $\sigma(\cdot)$ – a sigmoid function:

$$\sigma(x) = \frac{1}{1+e^{-x}}.$$

This transformation maps the aggregated score *Z(t)* into a probability-like value in the range [0,1], facilitating interpretation and decision-making.

The use of a probabilistic risk measure allows prioritization of threats, early warning systems, integration into decision support frameworks.

Moreover, risk estimation provides a bridge between anomaly detection and operational decision-making, which is essential for real-world applications.

4.7. Algorithm of the Digital Twin

To operationalize the proposed model, we define the algorithm of the digital twin as a sequence of steps performed at each time instance.

At time t, the system performs:

1. Collect data $X(t)$, $E(t)$;
2. Compute $S_{ts}(t), S_{event}(t), S_{ml}(t)$;
3. Integrate components into $Z(t) = \sum_{i=1}^{3} w_i(t) S_i(t)$;
4. Detect anomaly $A(t)$;
5. Estimate risk $R(t)$;
6. Update weights $w_i(t)$.

This algorithm reflects the real-time operation of the digital twin and emphasizes its dynamic and adaptive nature.

Importantly, the modular structure of the algorithm allows it to be implemented in both offline and streaming environments, making it suitable for practical deployment.

4.8. System Architecture of the Hybrid Digital Twin

1. General Architectural Principles. The architecture of the proposed hybrid digital twin is designed to provide a unified representation of complex infrastructure systems operating under conditions of uncertainty, heterogeneity, and potential adversarial influence. In contrast to traditional architectures that rely on a single modeling paradigm, the proposed system adopts a layered and modular structure, enabling the integration of multiple analytical approaches within a coherent framework.

Such a design is motivated by the intrinsic nature of cybersecurity-oriented environments, where system behavior cannot be adequately described using either purely continuous models or purely discrete representations. Instead, these systems exhibit a combination of continuous

state evolution, discrete event occurrences, and nonlinear dependencies. Consequently, the architecture must support the simultaneous processing of heterogeneous data streams and the interaction of different model types.

At a high level, the digital twin is structured as a sequence of interconnected layers, each responsible for a specific stage of data transformation and analysis. This layered organization ensures both conceptual clarity and practical scalability, allowing the system to be adapted to different types of infrastructure while preserving its analytical integrity.

2. Data Acquisition and Representation. The first stage of the architecture is responsible for acquiring data that reflects the current state and behavior of the monitored system. In cybersecurity-oriented IT infrastructures, such data is inherently heterogeneous and includes both continuous measurements and discrete events.

Continuous data is represented through system state variables, forming a vector (1), which may include metrics such as processor load, memory utilization, request rate, and latency. These variables capture the operational characteristics of the system and evolve over time.

In parallel, discrete events are represented as a sequence (2), where each event is defined by its occurrence time, type, and associated attributes. This representation is particularly important for cybersecurity applications, as it allows the system to capture abnormal activities such as repeated login failures, bursts of network requests, or error cascades.

The joint consideration of $X(t)$ and $E(t)$ forms the basis of the digital twin input:

$$D(t) = \{X(t), E(t)\}.$$

This formulation reflects the dual nature of the system and provides a foundation for subsequent analysis.

3. Data Processing and Feature Transformation. Raw data obtained from infrastructure systems is typically noisy, incomplete, and heterogeneous in structure. Therefore, a critical component of the architecture is the transformation of this data into a consistent and informative representation suitable for modeling.

This transformation is formalized as:

$$F(t) = \psi(X(t), E(t)),$$

where $\psi(\cdot)$ denotes a feature extraction and transformation operator. The role of this operator is to map heterogeneous inputs into a structured

feature space that captures both statistical properties and contextual relationships.

The necessity of this stage is supported by empirical evidence from prior studies, which demonstrate that carefully designed feature representations significantly improve the performance of machine learning models. In particular, the combination of statistical feature selection methods with dimensionality reduction techniques allows the system to retain informative features while reducing redundancy.

Importantly, the transformation process is not purely technical but also conceptual: it defines how the system "interprets" raw data, thereby influencing all subsequent analytical steps.

4. Analytical Modeling Layer. Following data transformation, the architecture employs a set of analytical models that operate in parallel. This design choice reflects the understanding that no single model is sufficient to capture all aspects of system behavior.

The first component focuses on time-series modeling and captures temporal dependencies in system metrics. It is defined as (3).

This component provides a baseline representation of normal system dynamics, allowing the detection of deviations from expected temporal patterns.

The second component is based on event-driven modeling and addresses the stochastic nature of discrete events. It is defined through an intensity function (4) and (5).

This formulation enables the system to capture temporal dependencies between events, including clustering and cascading effects, which are characteristic of cyberattacks.

The third component employs machine learning techniques to model complex nonlinear relationships within the data. It is defined as (6).

This component leverages the expressive power of data-driven models to identify patterns that are not explicitly represented in the other components.

The parallel operation of these models ensures that different aspects of system behavior are captured simultaneously, providing a richer and more robust representation.

5. Hybrid Integration Layer. The outputs of the analytical models are combined within the integration layer, which constitutes the core of the digital twin. This layer performs the fusion of heterogeneous information

sources into a unified representation (7), subject to the normalization condition (8).

The choice of a weighted combination is motivated by the need to balance the contributions of different components while maintaining interpretability. Unlike black-box integration methods, this formulation allows explicit analysis of how each component influences the final output.

Furthermore, the integration mechanism reflects a deeper conceptual principle: the system state is not determined by a single source of information but emerges from the interaction of multiple perspectives. By combining time-series dynamics, event dependencies, and machine learning predictions, the digital twin achieves a more comprehensive understanding of system behavior.

6. Decision and Interpretation Layer. The final stage of the architecture translates the integrated representation $Z(t)$ into actionable outputs. This involves both anomaly detection and risk estimation.

The anomaly indicator is defined as (9), while the risk level is computed as (10).

This layer plays a crucial role in bridging the gap between analytical modeling and practical application. By transforming numerical outputs into interpretable indicators, the system enables real-time monitoring, alert generation, and decision support.

Importantly, the separation of this layer ensures that the analytical core of the digital twin remains independent of specific application contexts, facilitating integration with external systems such as security monitoring platforms.

7. Architectural Implications. The proposed architecture represents a departure from conventional approaches by explicitly incorporating multiple modeling paradigms within a unified framework. Its layered structure ensures modularity and scalability, while the integration mechanism provides robustness against noise and uncertainty.

From a cybersecurity perspective, the architecture is particularly well suited to environments characterized by adversarial behavior and non-stationarity. By combining continuous and event-driven representations, the digital twin captures both gradual system changes and abrupt disruptions, enabling more effective detection and prediction of anomalous states.

To provide a clearer representation of the interaction between the components of the proposed hybrid digital twin, the overall system architecture is illustrated in Figure 1.

The architecture consists of several interconnected layers. The input layer includes heterogeneous data sources, combining continuous system state variables X(t) and discrete event streams E(t). These data are transformed into a unified feature representation F(t) through a preprocessing and feature extraction module.

The analytical core of the digital twin is formed by three parallel components: the time-series modeling module, the event-driven module based on stochastic processes, and the machine learning module. Each component produces a partial representation of system behavior, denoted as Sts(t), Sevent(t), and Sml(t), respectively.

These outputs are integrated within the hybrid fusion layer using an adaptive weighting mechanism, forming a unified system representation *Z(t)*. The weights are dynamically adjusted based on model performance, enabling the system to adapt to changing conditions.

The final layer performs anomaly detection and risk estimation, producing outputs A(t) and R(t). A feedback loop connects the output layer with the integration mechanism, allowing continuous adaptation of model weights.

This architecture reflects the multi-level and adaptive nature of the proposed hybrid digital twin and demonstrates how heterogeneous data and models are combined within a unified framework.

5. Data Requirements and System Implementation

5.1. General Considerations on Data in Cybersecurity-Oriented Digital Twins

The effectiveness of a digital twin model is fundamentally determined by the quality, structure, and representativeness of the data used for its construction. In cybersecurity-oriented infrastructure systems, data is inherently heterogeneous, multi-source, and often noisy. Unlike traditional engineering systems, where measurements are obtained from well-defined sensors, IT infrastructures generate complex data streams that reflect both system behavior and external interactions.

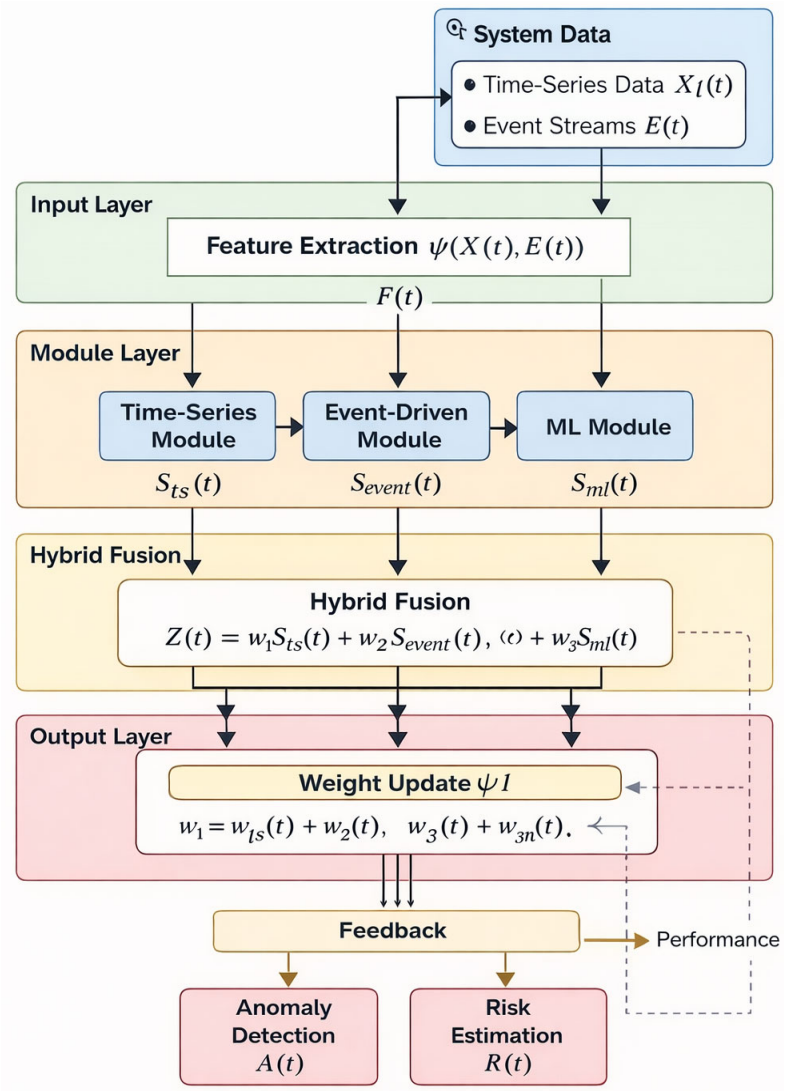

Figure 1. Architecture of the hybrid digital twin for cybersecurity-oriented infrastructure systems

Source of the figure: author's development

A key characteristic of such data is its hybrid nature. On the one hand, system operation is described by continuous metrics that evolve over time, such as resource utilization, response latency, and throughput. On the other hand, system behavior is strongly influenced by discrete events, including user actions, system errors, and potential attack activities. This duality necessitates a unified data representation capable of capturing both continuous and event-driven aspects.

In the context of the proposed hybrid digital twin, the input data is formalized as:

$$D(t) = \{X(t), E(t)\},$$

where *X(t)* represents system state variables and *E(t)* denotes the event stream. This formulation ensures consistency with the theoretical framework introduced in the previous sections and enables the integration of heterogeneous data sources within a single analytical model.

5.2. Sources of Data and Their Characteristics

To ensure both theoretical validity and practical applicability, the proposed approach relies on a combination of publicly available datasets and synthetically generated data. This strategy allows for controlled experimentation while maintaining a connection to real-world scenarios.

Public log datasets provide a valuable source of event-driven information [8], reflecting real system behavior under normal and anomalous conditions. In particular, large-scale log collections from distributed computing systems capture temporal patterns of events, including failures, warnings, and operational messages. Such datasets are especially suitable for modeling event streams and evaluating anomaly detection performance [14; 18].

In addition to log data, network traffic datasets play a crucial role in representing cybersecurity scenarios [9; 10]. These datasets typically include labeled instances of normal and malicious activities, enabling the evaluation of classification and detection models. Their feature spaces often contain both statistical and behavioral characteristics, which are essential for machine learning components.

However, relying solely on real datasets presents limitations, particularly in terms of completeness and controllability. Therefore, synthetic data is introduced to complement real-world observations. In this work, synthetic system metrics are generated to simulate continuous system behavior,

including resource usage, request intensity, and latency dynamics. This allows the digital twin to incorporate a state representation layer that may not be directly available in log-based datasets.

The combination of real and synthetic data enables the construction of a comprehensive experimental environment, where both event-driven and continuous aspects of system behavior are represented.

5.3 Data Preprocessing and Feature Transformation

Raw data obtained from heterogeneous sources cannot be directly used for modeling due to inconsistencies in format, scale, and structure. Therefore, preprocessing and feature transformation play a critical role in the proposed system.

The transformation process is formalized as:

$$F(t) = \psi\left(X(t), E(t)\right),$$

where $\psi(\cdot)$ represents a mapping from raw data to a structured feature space.

For continuous variables, preprocessing involves normalization and temporal aggregation. These operations ensure that features are comparable across different time intervals and reduce the impact of noise. Temporal aggregation, in particular, allows the extraction of meaningful trends and patterns from high-frequency data.

For event data, preprocessing includes parsing, categorization, and encoding. Log messages and network events are transformed into structured representations that capture both their semantic meaning and temporal properties. This step is essential for integrating event-driven information into machine learning models.

Feature engineering further enhances the representational capacity of the system. Building upon previous research, hybrid feature selection techniques are employed to identify the most informative features while reducing dimensionality [10]. This includes the combination of statistical measures, dimensionality reduction, and heuristic optimization methods.

The resulting feature space $F(t)$ provides a unified representation of system behavior, enabling consistent input for all modeling components.

5.4. Implementation of Model Components

The implementation of the hybrid digital twin follows the architectural principles defined in Section 4, where multiple analytical components operate in parallel and are subsequently integrated.

The time-series component is implemented using statistical or lightweight sequence modeling techniques, which capture temporal dependencies in system metrics. Given the focus on scalability and interpretability, preference is given to models that balance accuracy and computational efficiency.

The event-driven component is implemented using an intensity-based approach inspired by self-exciting processes. In practice, the exact formulation may be simplified to ensure computational feasibility while preserving the key property of temporal dependency. This component estimates the likelihood of future events based on historical event patterns.

The machine learning component is implemented using ensemble-based models, which have demonstrated strong performance in cybersecurity tasks. These models are particularly suitable due to their robustness to noise and ability to handle high-dimensional feature spaces.

The outputs of these components are computed independently and represent different perspectives on system behavior. This modular implementation ensures flexibility and allows individual components to be replaced or extended without affecting the overall system.

5.5. Integration and System Workflow

The integration of model components is performed within the digital twin core, where individual outputs are combined into a unified representation:

$$Z(t) = \sum_{i=1} w_i(t) S_i(t).$$

This integration step reflects the fusion strategy discussed in previous sections and serves as the basis for anomaly detection and risk estimation.

The overall system workflow follows a sequential process, where data acquisition, preprocessing, modeling, and integration are performed iteratively over time. At each time step, the system updates its internal representation and produces outputs that reflect the current state of the infrastructure.

Importantly, the workflow supports both offline and real-time operation. In offline mode, historical data is used to train and evaluate models, while

in real-time mode, the system processes incoming data streams and updates predictions dynamically.

5.6. Implementation Considerations

From a practical perspective, the implementation of the proposed system does not require specialized hardware or large-scale computational resources. The selected models are designed to operate efficiently on standard computing platforms, including local environments and cloud-based notebooks.

This design choice ensures accessibility and reproducibility, allowing the proposed approach to be implemented and validated without significant infrastructure requirements. At the same time, the modular architecture enables future scaling to more complex environments if needed.

Furthermore, the use of publicly available datasets and well-established modeling techniques enhances the transparency of the experimental setup and facilitates comparison with existing approaches.

6. Experimental Design and Evaluation Framework

6.1. Problem Formulation

The evaluation of the proposed hybrid digital twin model is formulated as a cybersecurity-oriented anomaly detection and risk estimation problem in IT infrastructure systems. The objective is to assess the ability of the model to accurately identify anomalous system states and estimate the associated risk under heterogeneous data conditions.

Formally, the problem can be defined as the mapping:

$$\left(X(t), E(t)\right) \to \left\{A(t), R(t)\right\},$$

where *R(t)* denotes the anomaly indicator and *A(t)* represents the risk level.

The evaluation focuses on assessing how effectively the hybrid model integrates multiple sources of information and whether this integration leads to improved detection performance compared to individual modeling approaches.

6.2. Experimental Scenario

The experimental scenario is designed to simulate the operation of a cybersecurity-oriented digital twin for an IT service infrastructure.

The system is assumed to operate under normal conditions interspersed with anomalous events, including potential cyberattacks and system failures.

The dataset used for evaluation combines event-driven and feature-based information. Event data is derived from system logs and network activity, while continuous features represent aggregated system metrics. Synthetic data generation is employed to complement real-world datasets, ensuring the presence of both normal and anomalous patterns.

This experimental setup reflects realistic conditions in which system behavior is influenced by both internal dynamics and external interactions. It also allows controlled variation of anomaly frequency and intensity, which is essential for evaluating model robustness.

6.3. Evaluation Metrics

To ensure a comprehensive assessment of model performance, multiple evaluation metrics are considered [14; 18]. These metrics capture different aspects of detection quality, including accuracy, robustness, and stability.

The primary metric used for anomaly detection is the F1-score, which balances precision and recall and is particularly suitable for imbalanced datasets commonly encountered in cybersecurity applications.

In addition, the receiver operating characteristic (ROC) curve and the corresponding area under the curve (ROC-AUC) are used to evaluate the discriminative ability of the model across different threshold settings.

To account for class imbalance and provide a more informative evaluation, the Matthews correlation coefficient (MCC) is also considered. This metric provides a balanced measure even when the classes are highly skewed.

From a system perspective, additional metrics are introduced to assess operational performance. Detection delay is used to measure the time required to identify anomalies after their occurrence, while the false alarm rate quantifies the frequency of incorrect anomaly detections.

Finally, stability metrics are considered to evaluate model robustness under non-stationary conditions, where data distributions may change over time.

6.4. Baseline Models

To evaluate the effectiveness of the proposed hybrid digital twin, its performance is compared against a set of baseline models representing different modeling paradigms.

The first group of baselines consists of classical machine learning models, including Random Forest and Gradient Boosting. These models are widely used in cybersecurity applications and provide strong performance in classification tasks.

The second group includes anomaly detection methods based on unsupervised learning, such as Isolation Forest. These approaches are particularly relevant when labeled data is limited.

The third group represents time-series-based models, which rely on temporal patterns in system metrics. These models provide a baseline for evaluating the contribution of temporal dynamics.

Finally, an event-driven baseline is considered, where anomaly detection is performed using event intensity measures derived from stochastic processes.

By comparing the hybrid model with these baselines, it becomes possible to isolate the contribution of each component and evaluate the benefits of their integration.

6.5. Evaluation Strategy

The evaluation is conducted using a controlled experimental setup that allows for systematic comparison between models. The dataset is divided into training and testing subsets, ensuring that model performance is assessed on unseen data.

To improve robustness, cross-validation techniques may be applied, particularly in cases where data availability is limited. This ensures that the evaluation results are not biased by specific data splits.

The performance of each model is assessed using the metrics defined above, and results are analyzed in terms of both accuracy and robustness. Special attention is given to scenarios involving non-stationary data and evolving anomaly patterns, which are characteristic of real-world cybersecurity environments.

6.6. Expected Outcomes and Validation Perspective

Although the primary objective of this study is to develop and formalize a hybrid digital twin framework, the experimental design provides a foundation for empirical validation. It is expected that the integration of time-series, event-driven, and machine learning components will lead to improved detection performance and increased robustness compared to individual models.

In particular, the hybrid model is expected to demonstrate superior performance in scenarios involving complex event dependencies and non-stationary behavior, where traditional approaches are less effective.

The presented evaluation framework can be directly extended in future work to include full-scale empirical validation, including real-time implementation and large-scale testing.

Conclusions

The rapid evolution of cyber threats and the increasing complexity of critical infrastructure systems necessitate the development of advanced analytical frameworks capable of capturing dynamic, heterogeneous, and adversarial environments. This work has addressed these challenges by proposing a hybrid digital twin framework tailored for cybersecurity-oriented infrastructure systems.

The study has demonstrated that traditional approaches to system modeling, including purely data-driven, rule-based, or physics-based methods, are insufficient when applied in isolation. Each of these paradigms captures only a limited aspect of system behavior and fails to provide a comprehensive representation required for reliable anomaly detection and risk assessment.

To overcome these limitations, a unified theoretical framework has been developed, integrating continuous system dynamics, event-driven stochastic processes, and machine learning techniques. The formalization of system behavior as a combination of state variables and event streams provides a consistent basis for modeling complex infrastructure environments. In particular, the incorporation of self-exciting stochastic processes enables the representation of temporal dependencies and cascading effects characteristic of cyberattacks.

A key contribution of this work is the development of a hybrid digital twin model that combines multiple analytical components within a unified architecture. The proposed integration mechanism, based on adaptive weighting, allows the system to dynamically adjust to changing conditions and evolving threat patterns. This feature is particularly important in non-stationary environments, where static models often fail to maintain performance.

The architectural design of the digital twin emphasizes modularity, scalability, and adaptability. By separating data acquisition, processing, modeling, and decision-making layers, the system ensures flexibility and extensibility, allowing it to be applied to different types of infrastructure systems without significant structural modifications.

The proposed experimental framework provides a structured approach for evaluating the effectiveness of the hybrid model. Although full-scale empirical validation remains a subject for future work, the defined evaluation methodology establishes a solid foundation for systematic comparison and performance assessment.

From a practical perspective, the proposed approach is particularly relevant in the context of critical infrastructure protection under conditions of hybrid warfare. The ability to detect anomalies, model risk, and adapt to evolving threats makes the digital twin a promising tool for enhancing the resilience of national infrastructure systems.

The results of this study contribute to the development of next-generation cybersecurity systems by bridging the gap between theoretical modeling and practical implementation. The integration of multiple analytical paradigms into a single framework represents a step toward more intelligent, adaptive, and robust infrastructure monitoring solutions.

Future research directions include the implementation of real-time digital twin systems, the integration of additional data sources such as threat intelligence feeds, and large-scale empirical validation using real-world infrastructure data. Further investigation is also required to optimize adaptive weighting mechanisms and to explore advanced learning techniques for dynamic environments.

References:

1. Grieves, M., & Vickers, J. (2016). Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. Transdisciplinary Perspectives on Complex Systems (c. 85–113). Springer International Publishing. DOI: https://doi.org/10.1007/978-3-319-38756-7_4

2. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital Twin in Industry: State-of-the-Art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. DOI: https://doi.org/10.1109/TII.2018.2873186

3. Lu, Y., Liu, C., Wang, K. I.-K., Huang, H., & Xu, X. (2020). Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing, 61, 101837. DOI: https://doi.org/10.1016/j.rcim.2019.101837

4. Lee, J., Bagheri, B., & Kao, H.-A. (2015). A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3, 18–23. DOI: https://doi.org/10.1016/j.mfglet.2014.12.001

5. Zhou, Z.-H. (2012). Ensemble Methods. Chapman and Hall/CRC. DOI: https://doi.org/10.1201/b12207

6. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. doi: 10.1023/A:1010933404324

7. Scarfone, K. A., & Mell, P. M. (2007). Guide to Intrusion Detection and Prevention Systems (IDPS). National Institute of Standards and Technology. DOI: https://doi.org/10.6028/nist.sp.800-94

8. Moiseienko, O. V., & Shyrmovska, N. H. (2026). Vyiavlennia anomalii u zhurnalakh podii serveriv na osnovi analizu shabloniv ta yikh statystychnoi informatyvnosti [Detection of anomalies in server event logs based on pattern analysis and their statistical informativeness]. *Naukovi pratsi DonNTU* [Scientific Papers of DonNTU], 1(23), 56–67. DOI: https://doi.org/10.31474/2074-7888-2026-1-23-56-67 (in Ukrainian)

9. Moiseienko, O. V., & Harasymiv, V. M. (2026). Kombinovanyi pidkhid vyiavlennia fishynhovykh saitiv iz vykorystanniam ansamblevykh modelei mashynnoho navchannia [A combined approach to phishing site detection using ensemble machine learning models]. *Komunalne hospodarstvo mist* [Municipal Economy of Cities], 1(196), 13–21. DOI: https://doi.org/10.33042/3083-6727-2026-1-196-13-21 (in Ukrainian)

10. Moiseienko, O. V. (2024). Adaptatsiia alhorytmu vyiavlennia anomalii u chasovykh riadakh dlia nestatsionarnykh potokovykh danykh [Adaptation of an anomaly detection algorithm in time series for non-stationary streaming data]. *Komunalne hospodarstvo mist* [Municipal Economy of Cities], 3(184), 16–22. DOI: https://doi.org/10.33042/2522-1809-2024-3-184-16-22 (in Ukrainian)

11. Moiseienko, O. V. (2024). Prohnozuvannia kiberatak na osnovi monitorynhu intensyvnosti trafiku v kompiuternykh merezhakh [Predicting cyberattacks based on traffic intensity monitoring in computer networks]. *Vcheni zapysky TNU im. V.I. Vernadskoho* [Scientific Notes of Taurida National V.I. Vernadsky University], 35(74), 136–143. DOI https://doi.org/10.32782/2663-5941/2024.3.1/21 (in Ukrainian)

12. Moiseienko, O. V., Zaiachuk, Y. I. Prohnozuvannia ryzykiv kiberatak na osnovi monitorynhu trafiku v kompiuternykh merezhakh z vykorystanniam alhorytmu analizu chasovykh riadiv [Predicting cyberattack risks based on traffic monitoring in computer networks using time series analysis algorithms]. *Vcheni zapysky TNU im. V.I. Vernadskoho* [Scientific Notes of Taurida National V.I. Vernadsky University]. Series: Technical Sciences. 2025. vol. 2, № 3. 268–275. DOI: https://doi.org/10.32782/2663-5941/2025.3.2/36 (in Ukrainian)

13. Ahmed, M., Naser Mahmood, A., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. DOI: https://doi.org/10.1016/j.jnca.2015.11.016

14. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection. *ACM Computing Surveys*, 41(3), 1–58. DOI: https://doi.org/10.1145/1541880.1541882

15. Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1), 35–45. DOI: https://doi.org/10.1115/1.3662552

16. Bishop, C. M., & Nasrabadi, N. M. (2006). Pattern recognition and machine learning (Vol. 4, No. 4, p. 738). New York: springer.

17. Box, George E. P. Time series analysis: forecasting and control. Fifth edition / George E.P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, Greta M. Ljung. 2015. Published by John Wiley and Sons Inc., Hoboken, New Jersey, p. 712.

18. Hawkes, A. G. (1971). Point Spectra of Some Mutually Exciting Point Processes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 33(3), 438–443. DOI: https://doi.org/10.1111/j.2517-6161.1971.tb01530.x