


## Article

# Advancing Cyber Incident Timeline Analysis Through Retrieval-Augmented Generation and Large Language Models

Fatma Yasmine Loumachi <sup>1</sup>, Mohamed Chahine Ghanem <sup>1,2,\*</sup> and Mohamed Amine Ferrag <sup>3</sup>

<sup>1</sup> Cyber Security Research Centre, London Metropolitan University, London N7 8DB, UK; fal0167@my.londonmet.ac.uk

<sup>2</sup> Cybersecurity Institute Liverpool, University of Liverpool, Liverpool L69 7ZX, UK

<sup>3</sup> Department of Computer Science, Guelma University, Guelma 24000, Algeria; ferrag.mohamedamine@univ-guelma.dz

\* Correspondence: mohamed.chahine.ghanem@liverpool.ac.uk

**Abstract:** Cyber timeline analysis or forensic timeline analysis is critical in digital forensics and incident response (DFIR) investigations. It involves examining artefacts and events—particularly their timestamps and associated metadata—to detect anomalies, establish correlations, and reconstruct a detailed sequence of the incident. Traditional approaches rely on processing structured artefacts, such as logs and filesystem metadata, using multiple specialised tools for evidence identification, feature extraction, and timeline reconstruction. This paper introduces an innovative framework, GenDFIR, a context-specific approach powered via large language model (LLM) capabilities. Specifically, it proposes the use of Llama 3.1 8B in zero-shot, selected for its ability to understand cyber threat nuances, integrated with a retrieval-augmented generation (RAG) agent. Our approach comprises two main stages: (1) Data preprocessing and structuring: incident events, represented as textual data, are transformed into a well-structured document, forming a comprehensive knowledge base of the incident. (2) Context retrieval and semantic enrichment: a RAG agent retrieves relevant incident events from the knowledge base based on user prompts. The LLM processes the pertinent retrieved context, enabling a detailed interpretation and semantic enhancement. The proposed framework was tested on synthetic cyber incident events in a controlled environment, with results assessed using DFIR-tailored, context-specific metrics designed to evaluate the framework's performance, reliability, and robustness, supported by human evaluation to validate the accuracy and reliability of the outcomes. Our findings demonstrate the practical power of LLMs in advancing the automation of cyber-incident timeline analysis, a subfield within DFIR. This research also highlights the potential of generative AI, particularly LLMs, and opens new possibilities for advanced threat detection and incident reconstruction.



Academic Editor: Wenbing Zhao

Received: 28 December 2024

Revised: 28 January 2025

Accepted: 10 February 2025

Published: 13 February 2025

**Citation:** Loumachi, F.Y.; Ghanem, M.C.; Ferrag, M.A. Advancing Cyber Incident Timeline Analysis Through Retrieval-Augmented Generation and Large Language Models. *Computers* **2025**, *14*, 67. <https://doi.org/10.3390/computers14020067>

**Copyright:** © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

**Keywords:** digital forensics; incident response; DFIR; timeline analysis; cyber incident; GenAI; LLM; RAG; cybersecurity

## 1. Introduction

In recent years, a significant rise in cyber incidents has been driven by exposed vulnerabilities affecting a broad range of digital devices such as computers, IoT devices, network hardware (including routers, switches, and IDS), and embedded systems. After a cyber incident, a DFIR investigation is conducted to uncover the complexities of the attack [1]. This process begins with the collection of digital artefacts, followed by the extraction of reliable evidence, and concludes with identifying the incident's root cause.

A crucial aspect of this investigation is timeline analysis, which focuses on examining the temporal sequences and chronological order of events. This involves identifying anomalies and suspicious patterns to reconstruct a comprehensive timeline of the entire incident [2].

The process of timeline analysis has traditionally been time-consuming due to the sheer volume and heterogeneity of collected data, requiring multiple specialised tools.

The DFIR field is witnessing a notable shift towards technological automation, with recent statistics indicating that 40% of practitioners have integrated automated systems into their evidence processing and analytical workflows. Within this trend, AI integration has gained particular momentum, with one-fifth of professionals adopting AI-driven solutions for their operations [3]. However, the sector's modernisation is still incomplete, as a significant proportion of organisations continue to rely on legacy methodologies, such as paper-based investigative procedures and basic digital tools. These traditional practices contribute to operational inefficiencies, with investigators spending an average of 45 h per case on digital evidence examination. This imperative is further emphasised by the 52% of practitioners who have identified a pressing need for more advanced tools to improve workflow efficiency and expedite investigative processes [4].

Digital artefact analysis, as the first step in timeline analysis, is often performed using tools like Velociraptor [5], FTK [6], EnCase [7], Dissect [8], and others. These tools process structured data, such as logs and filesystem metadata, to extract relevant and reliable evidence from many events.

Following this, events are reconstructed to produce a coherent timeline that provides context and meaning to the incident. Specialised tools such as Timesketch [9] and Log2Timeline (Plaso) [10] are widely employed in this phase, as they facilitate the reconstruction and visualisation of timelines, simplifying the correlation of activities and enhancing the interpretability of events. As an example, Splunk [11] represents an advanced platform that offers features powered via AI and ML to assist in detecting anomalies, identifying unusual patterns in large datasets, and providing deep insights for incident investigations [12]. Aside from tools that utilise AI for automation, other research methodologies and approaches to advance timeline analysis have been introduced [13–15], as evidenced in the literature.

Recently, generative AI models (GenAI), such as large language models (LLMs), have emerged as a transformative force, surpassing traditional AI solutions. These models have been integrated and utilised across various fields, processes, and tasks, and they have also been proposed to automate aspects of digital forensics, incident response, and cybersecurity.

LLMs, including GPT [16], Llama [17], and Claude [18], excel in processing data, detecting anomalies, and generating natural language explanations, making them valuable for assisting with DFIR artefact analysis. However, their direct application to cyber incident timeline analysis remains unresolved. Furthermore, according to DFIR standard practices, traditional tools are still the preferred choice for managing structured artefacts such as logs, filesystem metadata, and binary data.

Despite their promise, LLMs face challenges such as hallucinations [19], memory limitations [20], and gaps in context-specific knowledge. Innovations like retrieval-augmented generation (RAG) address these limitations by integrating external knowledge bases, enhancing the contextual accuracy and reliability of LLM outputs [21]. Additionally, their effectiveness can be further improved through prompt engineering [22], especially in decoder-based models [23], and the deployment of task-specific agents.

### *Research Question*

As per the definition of timeline analysis and the tools and techniques present in the practices of DFIR, there is a gap where current methods face many limitations. For instance,

existing solutions often present a final timeline without sufficient semantic context, where only timestamps of events are correlated. Furthermore, the use of multiple tools for analysis and others dedicated to event correlation complicates the process. Another challenge lies in the difficulty of spotting and successfully extracting evidence that could support claims regarding the incident and its root cause. While advancements in GenAI, particularly LLMs, and techniques like RAG, which enhance LLMs with external knowledge, offer promising potential to address these limitations, their application to DFIR timeline analysis remains underexplored. This study seeks to answer the following research questions:

- **RQ1:** How can LLMs and RAG be leveraged to enhance the current automation in cyber incident timeline analysis?
- **RQ2:** How can a framework driven by RAG and LLMs advance incident timeline analysis by seamlessly integrating artefact analysis and event correlation?
- **RQ3:** How can the framework be optimised to produce a reliable, comprehensive, and semantically rich timeline for DFIR investigations?

The paper is organised as follows to address the research questions: Section 1 introduces the field of DFIR and LLMs, establishing a foundational understanding of their core aspects. Section 2 expands on these details, offering deeper insights and a tailored definition relevant to this research. We then review related works and literature to build a robust foundation and adopt additional perspectives. Section 3 describes the methodology of our framework, including its primary functions. Section 4 discusses the implementation and testing of GenDFIR using synthetic scenarios. Section 5 touches on the results and outputs generated via GenDFIR. Due to the size of the produced report, only one scenario is showcased in Appendix A of the paper. This section also evaluates the framework's reliability, functionality, and effectiveness to assess its performance. Section 6 considers the limitations and ethical aspects of developing and deploying this framework, as well as its potential for expansion and adoption in real-world DFIR scenarios. Finally, Section 7 concludes the paper by highlighting its novelty and its potential for integration and application in the field.

## 2. Research Background and Related Works

The proposed GenDFIR framework combines different technologies to leverage generative AI in automating a DFIR task. To provide clarity on the application of RAG and LLMs in cyber incident timeline analysis, this section defines essential concepts and discusses related works. Some definitions have been adapted to suit the specific context of this research.

### 2.1. Timeline Analysis in DFIR

**Digital artefact:** There is no formal or precise definition of this term in the literature [24]. However, within the field of DFIR, artefacts are typically described as processed and relevant data collected and extracted from digital devices. For instance, in operating system (OS) forensics, these artefacts may include the file system, OS executables, network activity, internet history, cache, and other related data.

**Cyber incident event and anomaly:** An event refers to an action that potentially compromises or alters the security state of a system. In the context of a cyber incident, such actions are aimed at breaching the system's security policies [25]. An anomaly refers to deviations in behaviour that differ from established norms or expected patterns within a dataset. In the context of a cyber incident, an anomaly typically represents deviations from the expected normal behaviour of a system [26]. For instance, in Windows Event Logs, this could include unusual login attempts, irregular application activity, or unexpected changes in system configurations.

**DFIR:** is a term that encompasses two essential processes: digital forensics (DF) and incident response (IR).

- *DF* involves the management and analysis of digital evidence from its initial discovery to its presentation in legal contexts. This process includes the identification, collection, and analysis of evidence, with a key component being timeline analysis. Timeline analysis is essential in DF, as it helps reconstruct the sequence of events by establishing the chronological order of actions, which is crucial for uncovering critical details and understanding the flow of the incident.
- *IR:* On the other hand, refers to the set of actions and procedures an organisation follows to detect, manage, and mitigate cyber incidents. It typically starts with preparation, followed by detection, where timeline analysis also plays a role in understanding the sequence of events and determining the scope of the incident. During the analysis phase, establishing a chronological order of actions helps assess how the incident progressed. The process continues with containment, eradication, recovery, and post-incident activities, where insights from timeline analysis can guide decisions and help evaluate the effectiveness of the response.

DF plays a crucial role in IR, providing in-depth technical analysis that supports the overall process. DF helps identify the root cause of an incident, detect hidden access points, and uncover malicious activities. While DF focuses primarily on the technical aspects of an incident, IR addresses the broader scope, including containment, vulnerability remediation, and managing both technical and organisational elements.

Together, DF and IR complement each other in addressing a cyber incident. For instance, while IR may struggle with handling legal and regulatory aspects, DF excels in this area by ensuring the proper collection and presentation of evidence. Thus, timeline analysis not only supports each process individually but also strengthens their collaboration in responding effectively to incidents [2].

**Timeline analysis** involves the presence of a range of sub-activities to construct a coherent timeline of a cyber incident. Many studies in the literature focus on individual tasks, such as artefact analysis, anomaly detection, event correlation, or timeline reconstruction in isolation, with some incorporating modern AI-based solutions. Unfortunately, few works address the integration of all components of timeline analysis, particularly with an emphasis on automation. Table 1 introduces notable works relevant to the field of full/partial timeline analysis automation.

**Table 1.** DFIR timeline analysis.

Finding	Approach	Overview
<i>Tool:</i> Eric Zimmerman's tools [27]	Processing various types of data, including event logs, registry entries, and metadata, to provide detailed insights into incidents.	Beyond the tools discussed earlier, others, [27], have gained recognition for their capabilities in performing timeline analysis at a deterministic forensic level. However, they are not AI-based and lack automation, relying heavily on the expertise of the analyst or investigator.
<i>Study:</i> Chabot et al., 2014 [28]	Data are collected from various sources and analysed using tools like Zeitline and log2timeline. The FORE system manages events, semantic processing converts data into knowledge, and algorithms correlate events. Graphical visualisations then illustrate event sequences, relationships, and patterns, enhancing the clarity and interpretability of the data insights.	This contribution proposes a systematic, multilayered framework focusing on semantic enrichment to tackle challenges in timeline analysis. This approach not only automates timeline analysis but also delivers semantically enriched representations of incident events. However, one apparent limitation is the reliance on multiple standalone tools, which may complicate the workflow.

Table 1. Cont.

Finding	Approach	Overview
<i>Study:</i> Bhandari et al., 2020 [29]	Techniques that primarily involve managing, organising, and structuring temporal artefacts into a more comprehensible timeline. Log2timeline is utilised to extract timestamps from disk image files, while Psort processes the output to further handle the temporal artefacts and generate the final timeline.	A new approach that addresses the complexities and challenges of understanding generated temporal artefacts using abstraction techniques. Artefact analysis is performed manually; although it claims to manage the textual nature of events and produces easily interpretable results, it still relies on manual intervention for analysis.
<i>Study:</i> Christopher et al., 2012 [30]	Achieved by proposing the use of analyser plugins to conduct detailed analysis on raw, low-level events. These plugins extracted relevant data and aggregated it into high-level events. They then used Bayesian networks to correlate and link these high-level events by performing probabilistic inference.	The study focuses on automating event reconstruction and generating a human-understandable timeline. The main advantage of this approach is its ability to successfully handle and process large volumes of data, as well as produce an interpretable timeline.

## 2.2. Large Language Models

LLMs are advanced models of GenAI designed to understand and generate human language. They can predict word sequences and generate new text based on input data. Distinguished by their vast training datasets and sophisticated architectures, LLMs go beyond mimicking human-like creativity. They play a transformative role across various domains by enhancing productivity and automating complex tasks that would traditionally require human ingenuity [31]. However, a key distinction must be kept in mind, as various types of LLMs have been introduced, each optimised for specific tasks. For instance, consider the following.

**Decoder-based Models** such as GPT, LLaMA, Mistral, and Microsoft Phi excel in text generation, where the output is based on user input and the initial prompt.

- *Prompt Engineering* is the art and science of skillfully crafting and designing prompts to maximise the capabilities of a model. In the context of timeline analysis, this involves strategically framing inputs, specifying technical details, and establishing the investigative context to refine the model's output based on its precise and meaningful definition [23]. For example, a DFIR analyst might prompt the system as follows: "Conduct artefact analysis, correlate events, and reconstruct a coherent timeline of the incident." This method aims to ensure that the LLM understands the context, adheres to specific DFIR timeline analysis constraints, and achieves the intended goals of the investigation and analysis.

These models use an autoregressive decoder architecture, which generates one token at a time, predicting the next token based on the previous ones.

- *Token:* In LLMs, a token represents a character, word, subword, symbol, or number [23]. However, in timeline analysis, a token can be represented as follows:

*A non-tokenised Windows event log, see Figure 1:*

```
Event ID: 7124, Details: Unusual network response pattern, Level:
Error, Date and Time: 08/01/2024 15:43:21, Source: Intrusion Detection
System, Task Category: Threat Detection.
```

**Figure 1.** Windows event log sample.

*A tokenised form (using the text-embedding-ada-002—external embedding model), see Figure 2:*

```
Event ID: 7124, Details: Unusual network response pattern, Level:
Error, Date and Time: 08/01/2024 15:43:21, Source: Intrusion
Detection System, Task Category: Threat Detection.
```

**Figure 2.** Text-embedding-ada-002—Windows event log embedding (Hugging Face Tokeniser [32]).

*A tokenised form (using the GPT3—internal embedding model), see Figure 3:*

```
Event ID: 7124, Details: Unusual network response pattern, Level:
Error, Date and Time: 08/01/2024 15:43:21, Source: Intrusion
Detection System, Task Category: Threat Detection.
```

**Figure 3.** GPT3 embedding—Windows event log embedding (Hugging Face Tokeniser [32]).

This process occurs during the tokenisation phase, which is part of the LLM architecture. The specific tokenisation and embedding approach used depends on whether the model relies on its internal system or integrates an external embedding model. Also, each model produces different outputs based on its tokenisation and embedding techniques [23].

- *Embedding* is the process of converting text into numerical representations, often in the form of tensors, suitable for the LLM. It begins with tokenisation, where words or characters within a text are transformed into tokens, representing individual units. These tokens are then mapped to numerical values that capture their semantic representation. Subsequently, additional layer transformations and processing are applied to refine these representations further. The final output is a dense vector where each value corresponds to a specific feature of the text [33].

This makes them particularly suited for generating coherent and contextually relevant text in response to user queries [34].

**Encoder-based models**, such as BERT, excel in tasks like classification and sentiment analysis, where they are commonly used in applications such as text categorisation and emotion detection. In these models, the output is generated based solely on the information present in the input, without relying on external context or previous outputs.

**Encoder–decoder models**, like T5 and BART, excel in tasks such as text translation and summarisation, as well as speech recognition and image recognition. The output of such models is based on both the input and the context, where the output is a transformed version of the input [35].

In the field of cybersecurity and DFIR, several studies have been conducted, and approaches have been proposed that integrate LLMs for various applications, as illustrated in Table 2:

Based on the available findings, the most widely used state-of-the-art models in the scope of this research are GPT and Llama. In the study [36], researchers disclosed the capabilities of Llama 3.1 in performing advanced cybersecurity tasks, trained on publicly validated data, including cybersecurity-related content. An important consideration is the benchmarking over the new CYBERSECEVAL 3 suite to measure cybersecurity risks and capabilities. The model can recognise and identify cyber threats with high accuracy.

In addition, GPT has been used in different studies, as shown in Table 2, for the purpose of identifying and mitigating cyber threats. Furthermore, in the study [37], the authors mentioned its ability to perform even in a zero-shot setting, as it has been extensively trained on massive datasets, such as reports, journal and conference papers, previous DFIR cases, standards, frameworks, guidelines, and conversations related to the field. Examples include IEEE articles, Wikipedia entries, CTDD, IDS2017, and IDS2018 datasets, ADFA (Australian Defence Force Academy) datasets, OpenTitan System-on-Chip (SoC) data, Hack@dac 2021 SoC data, CVE reports, and more [38].

**Table 2.** Large language models.

Finding	Approach	Overview
<i>Study:</i> Ferrag et al., 2024 [38]	Reviewing and examining research studies, published articles, and journals addressing the integration of generative AI and LLMs into cybersecurity. Additionally, the paper discusses features, findings, insights, and theoretical approaches derived from datasets relevant to cybersecurity, including training methods, architectures, and associated mathematical equations.	The authors provide an extensive review of the application of LLMs in the field of cybersecurity, covering its subfields, including intrusion detection, cyber forensics, and malware detection.
<i>Study:</i> Ota et al., 2024 [39], Wang et al., 2024 [40], Fariha et al., 2024 [41], Saha et al., 2024 [42]	<ul style="list-style-type: none"> <li>- [39] LLaMA3 8B and 70B, and Phi3 applied to honeypot systems for conducting advanced malicious activity analysis and detection</li> <li>- [40] GPT and BERT, with LSTM to predict cyberattacks in IoT networks.</li> <li>- [41] GPT-3.5 for log summarisation to analyse and summarise log files and detect specific events.</li> <li>- [42] advanced paradigm utilising GPT for SOC tasks, including vulnerability insertion, security assessment, and security verification.</li> </ul>	All of these studies explored the use of different LLMs to automate specific cybersecurity tasks by embedding them into their workflows.
<i>Study:</i> Wickramasekara et al., 2024 [43]	Theoretically introducing and explaining how LLMs can be utilised in various phases of a DF investigation and how specific models can perform particular tasks. For example, a model like GPT-3.5 can generate textual reports at the conclusion of investigations, while multimodal LLMs, such as GPT-4 and LLaVA with vision assistance, can analyse images and videos, providing contextual outputs for digital forensics.	The paper provides an extensive literature review on the integration of LLMs to advance the DF process.
<i>Study:</i> Scanlon et al., 2023 [44]	The role of ChatGPT in supporting various tasks, including artefact analysis, generating regular expressions and keyword lists, creating scripts for file carving, RAID disk acquisition, and password cracking, identifying IR sources, anomaly detection, and developing detailed forensic scenarios.	Presents a comprehensive study on how ChatGPT can assist during DF investigations, examining this concept from multiple perspectives. The paper also addresses the limitations and strengths of ChatGPT, clearly stating that the model significantly enhances the DF process.
<i>Study:</i> Sakshi et al., 2023 [45]	ChatGPT, powered by GPT-4 and GPT-3.5 models, was used to analyse artefacts (input data) and extract relevant evidence, such as conversations, images, and other information pertinent to the investigation.	The paper proposes using ChatGPT to enhance digital investigations by identifying evidence during the DF process. The paper emphasises that, despite the efficiency of this method, its outputs must always be verified and monitored by humans.
<i>Tool:</i> CADO 2024 [46]	Dedicated to assisting forensic analysts with investigations by providing insights and streamlining the investigation process.	A recent AI-based platform, CADO [46], powered by a local LLM, was developed.

### 2.3. Retrieval-Augmented Generation (RAG) in Information Extraction and Knowledge Expansion

**RAG** is a technique that primarily aims to optimise and enhance the performance of LLMs. It utilises an external knowledge base beyond the LLM's pre-existing knowledge and training datasets to provide additional information during inference. This process is not the same as fine-tuning or training an LLM, but rather, it involves dynamically retrieving, generating, and integrating relevant external textual information from databases (external storage), vault files (internal storage), or cloud pipelines. It is crucial to acknowledge that, while an LLM alone is very powerful, its knowledge may not encompass specific contexts [21,47].

RAG consists of two essential functions, retrieval and generation.

- *Retrieval-augmented*: In the context of RAG, retrieval refers to the process of searching and selecting relevant information from a knowledge base or document dataset to enhance the output of a generative model. This ensures that the model generates more accurate and contextually relevant responses [47]. Retrieval techniques commonly used in information retrieval include cosine similarity, Jaccard similarity, BM25, TF-IDF, latent semantic analysis (LSA), and embedding-based models. These methods evaluate the similarity between the query and documents, retrieving the top k pieces with the highest similarity scores to provide relevant contextual knowledge for the LLM [48].

In GenDFIR, retrieval is applied to extract relevant incident events from documents containing cyber incident events. It helps identify and select the most pertinent ones (top k pieces) based on their relevance to the query or event context.

- *Augmented generation* is the process through which, after retrieval, the extracted information is provided to the LLM for semantic enhancement, enabling it to generate more accurate and contextually enriched outputs [47]. In the case of GenDFIR, this involves extracting evidence (mostly anomalous events) from the incident knowledge base and passing it to the LLM to produce a timeline of the incident, which includes key events, correlations, and an analysis of the evidence.

**Knowledge base:** In LLMs with RAG, a knowledge base refers to a structured repository of factual knowledge and data related to specific domains, which the model may use during inference [21]. In the context of this research (GenDFIR), it will serve as the repository for storing and managing all data generated and collected from a cyber incident. These data are presented in a natural language format, with all data stored in English.

**Top-k pieces/evidence:** Top-k is a hyperparameter in RAG that controls the number of the most relevant pieces of information retrieved from an external knowledge base during the retrieval phase. The value of k determines how many chunks or documents, ranked by relevance, are returned for further processing via the LLM [49].

In GenDFIR, the top-k pieces represent the most relevant retrieved log events (or evidence) with the highest anomaly scores pertinent to the cyber incident. Initially, k refers to all events within the knowledge base, and the top-k represents the most relevant ones selected for further analysis via the LLM.

**Chunk:** A chunk is a manageable segment of text extracted from a document for specific processing or analysis. Chunking involves dividing large texts into smaller segments that can be embedded through an embedding model. In timeline analysis, chunking ensures that critical details, such as timestamps and numerical tokens representing seconds or minutes, are accurately captured. These small details are vital for event correlation, where even slight variations in timing can significantly impact the analysis. By segmenting the text into manageable chunks, these details can be preserved and embedded effectively, enabling precise token calculation and retrieval from an external knowledge base when using RAG with LLMs. The precision of timeline analysis depends on the segmentation method used. Factors like token calculation, and the careful selection of the embedding model, all influence the accuracy and reliability of the process. Several reliable approaches to chunking have been proven effective [50].

- *Token-based*: divides text based on a fixed number of tokens.
- *Paragraph-based*: segments text by paragraphs to maintain context.
- *Semantic-based*: groups text based on meaning or topics.
- *Sentence-based*: segments text into sentences, each of which may have unique semantics.



**DFIR context-specific RAG agent:** A RAG agent is a key component of some modern LLMs-powered AI systems. It acts as an intermediary between the LLM and the user, managing the search process when a query or input is received. The agent retrieves relevant information from a knowledge base by employing retrieval methods that measure data similarity. This context-specific information is then provided to the LLM, which uses it to generate a response tailored to the query. Additionally, as part of the process, the agent can augment generation by refining and modifying human-crafted prompts to enhance the quality of the generated response in line with the specific task. During inference, the agent processes the input, adjusts it based on its designated role, and helps the LLM produce a more relevant and accurate answer [51]. It is important to note that the role of a RAG agent is defined by a human, ensuring that the task aligns with its intended purpose. In the case of timeline analysis, an example could be summarised as follows: a DFIR analyst prompts the system with “Conduct timeline analysis”, while the agent’s prompt (task) would be “You are a DFIR AI assistant, tasked with analysing artefacts, correlating events, and producing a coherent timeline of the incident. Base your answer on the provided context and do not include additional information outside of the context given.”. This allows the RAG agent to extract relevant incident information, such as pertinent event logs, from the knowledge base. Additionally, the system can support the DFIR analyst by refining the prompt if critical details are overlooked, aiming for more optimised results.

Table 3 highlights several studies that have explored the use of RAG for information extraction and knowledge augmentation.

Table 3. Retrieval-augmented generation.

Finding	Approach	Overview
<i>Study:</i> Tihanyi et al., 2024 [52]	RAG to generate high-quality, context-based questions using an external knowledge.	The paper presents RAG usage to advance the process of creating a cybersecurity dataset. Subsequently, CyberMetric was used to benchmark the general cybersecurity knowledge of cybersecurity-oriented LLMs.
<i>Study:</i> Lála et al., 2023 [53]	Employing RAG agents to answer questions based on embedded scientific literature.	The paper explains and showcases the power of RAG agents in addressing some limitations of traditional LLMs, such as hallucinations and lack of interpretability.
<i>Study:</i> Wang et al., 2023 [54]	RAG extraction and retrieval are optimised using chunks and tokens, where chunks are used for text segmentation, and tokens represent units, words, or subwords within a chunk.	The authors propose an approach to enhance the performance of RAG in single LLMs for both contextual ranking and answer generation. The goal is to ensure that each chunk contains sufficient context to answer questions and queries accurately.
<i>System:</i> Chat2Data by Zhao et al., 2024 [55]	A prototype for advanced data analysis using RAG for data retrieval, and a knowledge base where all data are stored. Outputs are shown in a graphical representation.	The authors applied RAG and LLMs to build and introduce an interactive system.
<i>System:</i> BIDARA by Toukmaji et al., 2024 [56]	Employing RAG technologies and LLM agents to address the complexities of biomimicry.	An AI research assistant model was presented as a system.

### 3. Research Methodology

GenDFIR is an LLM-powered framework specifically proposed for timeline analysis in DFIR. However, using an LLM alone presents several challenges, particularly with respect

to its context window limitations. LLMs, such as GPT, process input as large text blocks, which restricts the amount of detailed event data they can accommodate at once [57,58]. This creates a significant issue in DFIR timeline analysis, where precise event-specific details are crucial. When processing DFIR incident documents, the LLM may overlook vital information, as its context window prioritises broader context over the granular details necessary for an accurate timeline reconstruction.

Furthermore, capturing the temporal and logical relationships between events in a cyber incident is a complex task that is often overlooked in traditional methods. These approaches generally require repeatedly inputting artefacts into the LLM, each time referencing previous data. This process is not only labour-intensive but also inefficient, as the model must reprocess prior references without maintaining a coherent understanding of the evolving context of the incident.

To address these limitations, we propose integrating RAG technology within GenDFIR. Rather than enhancing the LLM with an external knowledge base, RAG will treat the DFIR incident reports themselves as the source of knowledge. These reports, containing detailed event data, will be directly uploaded into the framework, with the knowledge representing the interconnected events within the incident. RAG enhances the LLM's ability to correlate events accurately and efficiently, enabling the dynamic extraction of events and providing easy, instant access to the incident data.

In addition to RAG, a RAG agent powered by the LLM will further augment the process by acting as a DFIR timeline analysis context-specific agent. This agent will manage the process of filtering and extracting information from the knowledge base according to its role. It will enable advanced event filtering, ensuring that only the most relevant events are retrieved for analysis in the specific context of the incident.

This approach not only overcomes the limitations of the context window but also surpasses traditional methods by providing interpretable and enriched contextual outputs, ensuring a more comprehensive and efficient timeline analysis.

The following sections further explain and advance the technical methodology of GenDFIR, outlining its development.

### 3.1. Incident Events Preprocessing and Structuring

After a cyber incident, various artefacts are generated and stored in CSV files, forming incident datasets. The preprocessing step involves converting these CSV files into a structured text document.

A single event,  $E$ , is characterised and defined by its attributes,  $e$ , as follows:

$$E = \langle e_1, e_2, \dots, e_n \rangle \quad \text{where } n \text{ is the last number of attributes}$$

and  $e_1$ ,  $e_2$ , and  $e_n$  are the individual attributes of the event, such as level, date and time, source, event id, and task category.

We define a cyber incident,  $I$ , as a set or sequence of events,  $E$ :

$$I = \langle E_1, E_2, \dots, E_m \rangle \quad \text{where } m \text{ is the last number of Events}$$

A cyber incident,  $I$ , represented by a CSV file, contains  $t$  events.

The CSV file is converted into a textual incident document,  $ID$ , preserving the same events, and it is structured as follows:

$$I_{\text{CSV}} = \langle E_1, E_2, \dots, E_t \rangle$$

$$ID = \langle E_1, E_2, \dots, E_t \rangle \quad \text{where } ID = I_{\text{CSV}}$$

- The collected events refer to those that occurred during the timeframe of the cyber incident.

Each row of an event in  $I_{CSV}$  is treated as an individual context, representing a unique event within the timeline of the incident. These rows are then converted into corresponding segments in  $ID$ , maintaining their sequential order and the integrity of their context throughout the transformation process.

The following figures illustrate the transformation of  $I_{CSV}$  to  $ID$ :

*Unprocessed and Unstructured Windows Event Logs stored in CSV format, see Figure 4:*

	A	B	C	D	E	F	G
1	Event ID	Details	Level	Date and Time	Source	Task Category	
2	7124	Unusual network response pa	Error	08/01/2024 15:43	Intrusion Detecti	Threat Detection	
3	8921	Anomalous IP returned in DNS	Warning	08/01/2024 13:12	DNS Server	Query Response	
4	5429	Network packet inconsistency	Critical	08/01/2024 16:55	Network Firewall	Suspicious Traffic	
5	3302	Multiple DNS queries with del	Warning	08/01/2024 10:24	Intrusion Detecti	Network Performance	
6	1203	Outbound traffic to unexpecte	Error	08/01/2024 18:32	DNS Server	Unusual Traffic	
7	6708	Increased latency in DNS req	Warning	08/01/2024 11:46	Web Proxy	Performance Degradation	
8							

**Figure 4.** Windows event log CSV.

*Processed and Structured Windows event logs stored in a document (PDF) format, see Figure 5:*

Event ID: 7124, Details: Unusual network response pattern, Level: Error, Date and Time: 08/01/2024 15:43:21, Source: Intrusion Detection System, Task Category: Threat Detection.

Event ID: 8921, Details: Anomalous IP returned in DNS query, Level: Warning, Date and Time: 08/01/2024 13:12:05, Source: DNS Server, Task Category: Query Response.

Event ID: 5429, Details: Network packet inconsistency, Level: Critical, Date and Time: 08/01/2024 16:55:40, Source: Network Firewall, Task Category: Suspicious Traffic.

Event ID: 3302, Details: Multiple DNS queries with delayed responses, Level: Warning, Date and Time: 08/01/2024 10:24:11, Source: Intrusion Detection System, Task Category: Network Performance.

Event ID: 1203, Details: Outbound traffic to unexpected server, Level: Error, Date and Time: 08/01/2024 18:32:45, Source: DNS Server, Task Category: Unusual Traffic.

Event ID: 6708, Details: Increased latency in DNS requests, Level: Warning, Date and Time: 08/01/2024 11:46:22, Source: Web Proxy, Task Category: Performance Degradation.

**Figure 5.** Windows event log document (knowledge base).

The format in which an event is written in the document plays a crucial role in enhancing retrieval precision and optimising the LLM's understanding.

In GenDFIR, we propose a dynamic structure based on the context of the incident. For instance, in the case of collected event logs from a Windows incident, the structure is as follows: each attribute of an event is separated by a comma (","), and the event concludes with a full stop (".") to signify that each event is separate and represents unstructured data. However, when the format of an event differs from that of Windows event logs (e.g., Linux syslog events, web server logs, firewall logs, or database logs), alternative structural models can be adapted based on the type and nature of the data. In incidents requiring more contextual analysis, such as phishing email incidents, each email is represented as a single event, with a symbol (e.g., "/" ) used to indicate the end of each email.

- A critical consideration at this stage is that GenDFIR, with its core LLM, is capable of recognising a wide range of cyber threats, extending beyond Windows-based systems. The Llama-3.1 LLM integrated within GenDFIR has been trained on diverse publicly available datasets that cover various types of cyber threats. As a result, even when artefacts are gathered from different sources, such as IoT devices or Linux operating

systems, the core LLM, trained to recognise these types of threats through contextual data, can process and analyse relevant cybersecurity incidents within the framework.

### 3.2. Incident Event Chunking

Building on the previously defined concepts of chunking and the structured document containing the incident, we propose a DFIR context-specific chunking method designed to improve retrieval precision, with each event representing a distinct chunk.

1. **Cyber incident event length:** The length of an event is defined as the total sum of the number of characters across its attributes. The characters in a line of events are counted based on their attributes,  $e_{ij}$ , where the following applies:

- $i$  represents the event number (e.g., Event 1 and Event 2), and
- $j$  represents the attribute number within that event (e.g., event ID and level).

The total number of characters in the  $i$ -th event is given as follows:

$$T(E_i) = \sum_{j=1}^n T(e_{ij}), \quad (1)$$

- $T(E_{ij})$  is the total number of characters in the  $i$ -th event.
- $T(e_{ij})$  is the number of characters in the  $j$ -th attribute of the  $i$ -th event.
- $n$  is the total number of attributes in the  $i$ -th event.

2. **Event chunking:** The proposed method extends the general chunking concept by incorporating the calculated lengths of individual events. It enhances retrieval precision through precise semantic embedding, as every event attribute contributes to the chunking process, improving the accuracy of retrieval tasks. The chunking process adheres to the following equations to ensure compliance with both the token capacity of the LLM and the adopted embedding model.

To calculate the number of tokens for an event or chunk, we consider the token capacity of the embedding model, for instance, models that can process up to 512 tokens per input.

For optimisation, we assume an average of 4 characters per token, which is within the standard range of 4 to 6 characters in English, with 4 selected to maximise token capture and improve input accuracy. Based on this, the number of tokens for an event  $T_{\text{tokens}}(E_i)$  is approximately the total number of characters in the event  $T(E_i)$  divided by the average of  $C_{\text{avg}}$  characters per token. This relationship is expressed in Equation (2):

$$T_{\text{tokens}}(E_i) \approx \frac{T(E_i)}{C_{\text{avg}}} \quad \text{where } C_{\text{avg}} = 4, \quad (2)$$

- $i$  represents the index of the  $i$ -th event in the sequence.

Similarly, the number of tokens for a chunk,  $C_k$ , is calculated in the same manner, as in Equation (3):

$$T_{\text{tokens}}(C_k) \approx \frac{T(C_k)}{C_{\text{avg}}}, \quad (3)$$

- $k$  represents the index of the  $k$ -th chunk in the sequence.

This equation follows the same principle as Equation (2) but applies to a chunk,  $C_k$  (a chunk,  $C_k$ , represents a single event,  $E_i$ ).  $T(C_k)$  is the total number of characters in a chunk.

The constraint on the number of tokens in a chunk is introduced in Equation (4). It ensures that the number of tokens for any chunk,  $T_{\text{tokens}}(C_k)$ , does not exceed the model's token capacity,  $TM$ :

- $TM$  represents the model's maximum token capacity, such as 512 tokens.

This helps ensure that a chunk fits within the embedding model's processing limits.

$$T_{\text{tokens}}(C_k) \leq TM, \quad (4)$$

$$M_{\text{max}} \approx TM \times C_{\text{avg}}, \quad (5)$$

- $M_{\text{max}}$  represents the maximum number of characters that the model can embed per chunk.

We define a cyber incident,  $ID$ , in a document as the sum of all chunks,  $C_i$ , where each chunk corresponds to an event  $E_i$ :

$$ID = \begin{cases} \sum_{i=1}^n C_i, \\ \text{where } C_i = E_i, \\ \text{and } E_i = \sum_{j=1}^{m_i} \text{Character}_{ij}, \\ \text{with } m_i \leq M_{\text{max}}. \end{cases} \quad (6)$$

- $n$  is the total number of chunks,
- $i$  represents the index of the  $i$ -th event in the incident,
- $j$  represents the index of the  $j$ -th character in the  $i$ -th event, and
- $m_i$  is the maximum or total number of characters that can be processed in the  $i$ -th event.

$$\text{MaxLength} = \max(\text{Length of } E_1, \text{Length of } E_2, \dots, \text{Length of } E_n). \quad (7)$$

- $\text{MaxLength}$  represents the greatest number of characters found in any single event within the incident, ensuring that the maximum event length is selected as the standard size of a chunk.
- If a chunk does not reach the  $\text{MaxLength}$  limit, it is padded and adjusted to this limit.

Finally, the total length of a cyber incident represents the sum of the lengths of all events or chunks within the incident:

$$T(C_{\text{Incident}}) = \sum_{i=1}^t T(E_i) = \sum_{i=1}^t T(C_i). \quad (8)$$

- $T(\text{Incident})$  represents the total length of the incident, calculated as the sum of the lengths of each event,  $E_i$ , or chunk,  $C_i$ .
- $t$  represents the total number of events or chunks in the incident. This total length determines how many tokens will be used to process the entire incident, ensuring that the chunking and tokenisation process adheres to the embedding model's constraint.

The following figure, Figure 6, illustrates a visual representation of our context-specific chunking method applied to an incident knowledge base (events):

Event ID: 7124, Details: Unusual network response pattern, Level: Error, Date and Time: 08/01/2024 15:43:21, Source: Intrusion Detection System, Task Category: Threat Detection.	Chunk 1 $m_i = 177$
Event ID: 8921, Details: Anomalous IP returned in DNS query, Level: Warning, Date and Time: 08/01/2024 13:12:05, Source: DNS Server, Task Category: Query Response.	Chunk 2 $m_i = 163$
Event ID: 5429, Details: Network packet inconsistency, Level: Critical, Date and Time: 08/01/2024 16:55:40, Source: Network Firewall, Task Category: Suspicious Traffic.	Chunk 3 $m_i = 168$
Event ID: 3302, Details: Multiple DNS queries with delayed responses, Level: Warning, Date and Time: 08/01/2024 10:24:11, Source: Intrusion Detection System, Task Category: Network Performance.	Chunk 4 $m_i = 193$

$MaxLength = 193$

**Figure 6.** DFIR context-specific chunking.

However, the practical application of this is shown in Listing 1, which provides a concise overview of the process and illustrates its implementation:

**Listing 1.** Chunking code snippet.

```

1 # 'events' is a text containing multiple events
2 # Max length of a chunk ( Max number of characters within a chunk)
3 MaxLength = 208 # Set the fixed maximum length for each event/chunk (
4     Maximum number of characters in a chunk)
5
6 # Split the events text by (.) to process each event separately
7 events = events.split(".") # or ("/") - Split the document into events
8
9 # Loop over each split event
10 for event in events:
11     # Remove extra spaces from each event for accurate characters
12     # calculation
13     event = event.strip()
14
15     # Check if the event is < than the max length, and pad if necessary
16     if len(event) < MaxLength:
17         # If the event is < than MaxLength, pad it with spaces
18         event = event.ljust(MaxLength)
19
20     # Add the event as a chunk
21     chunks.append(event) # Event is a chunk with the required length

```

### 3.3. Incident Events and DFIR Query Embedding

1. **Events:** Following the chunking process, this step focuses on embedding the chunks into high-dimensional vector representations. An embedding model is adopted to transform each chunk or event into a dense vector, which captures its semantic meaning and enables efficient similarity search and analysis. The first step involves passing an individual event,  $E_i$ , through the embedding model, which first tokenises the event (i.e., splits it into tokens such as words or subwords) and then generates its vector,  $\mathbf{v}_{E_i}$ :

$$\mathbf{v}_{E_i} = \text{Embed}(\text{Tokenise}(E_i)), \quad (9)$$

$$\mathbf{v}_{E_i} = \text{Embed}(E_i) = \text{Embed}(\langle e_1, e_2, \dots, e_n \rangle). \quad (10)$$

- The event,  $E_i$ , consists of several attributes,  $e_1, e_2, \dots, e_n$ , and each attribute  $e_k$  represents a specific aspect or detail of the event (such as the date, description, etc.).
2. **DFIR query:** The DFIR expert conducts an automated timeline analysis by querying or instructing the framework. The query/instruction,  $Q_{DFIR}$ , undergoes an embedding process similar to that of the event,  $E$ . When the user inputs the query, it is first

tokenised and then embedded, transforming it into a dense vector  $\mathbf{V}_{QDFIR}$ . The process is expressed as follows:

$$\mathbf{v}_{QDFIR} = \text{Embed}(\text{Tokenise}(Q_{DFIR})). \quad (11)$$

### 3.4. Context-Specific LLM-Powered RAG Agent for Timeline Analysis

1. **Event retrieval:** The retrieval phase is initiated when the DFIR expert provides a query related to an incident. The query is embedded into a vector,  $\mathbf{v}_{QDFIR}$ , along with the vectors representing anomalous events,  $\mathbf{v}_{E_i}$ , stored in the external knowledge base,  $KB$ . We represent the retrieval process of GenDFIR as follows:

$$\text{RAG}_{DFIR}(\mathbf{v}_{QDFIR}, \mathbf{v}_{E_i}) = \text{Retrieve}(\mathbf{v}_{QDFIR}, \{\mathbf{v}_{E_i}\}_{i \in \mathcal{R}}, \text{cosine\_similarity}(\mathbf{v}_{QDFIR}, \mathbf{v}_{E_i})). \quad (12)$$

These embeddings allow the framework to semantically compare the input query with historical incident events. The similarity between the query vector,  $\mathbf{v}_{QDFIR}$ , and each event vector,  $\mathbf{v}_{E_i}$ , is computed using the cosine similarity function:

$$\text{cosine\_similarity}(\mathbf{v}_{QDFIR}, \mathbf{v}_{E_i}) = \frac{\mathbf{v}_{QDFIR} \cdot \mathbf{v}_{E_i}}{\|\mathbf{v}_{QDFIR}\| \|\mathbf{v}_{E_i}\|}. \quad (13)$$

- $\mathbf{v}_{QDFIR} \cdot \mathbf{v}_{E_i}$  is the dot product of the query and event vectors, quantifying their alignment in the vector space.
- $\|\mathbf{v}_{QDFIR}\|$  and  $\|\mathbf{v}_{E_i}\|$  are the Euclidean norms (magnitudes) of the respective vectors.

The cosine similarity identifies how closely the query aligns with the context of each anomalous event. This computation forms the basis for selecting the most relevant events. To ensure efficiency, only the top- $k$  events with the highest similarity scores are retrieved, as defined by the following:

$$\mathcal{R} = \text{TopK}\left\{ \frac{\mathbf{v}_{QDFIR} \cdot \mathbf{v}_{E_i}}{\|\mathbf{v}_{QDFIR}\| \|\mathbf{v}_{E_i}\|} \mid i = 1, 2, \dots, t \right\}, k. \quad (14)$$

- $\mathcal{R}$  represents the set of retrieved event vectors ranked by similarity scores. It plays a critical role in narrowing down the large pool of potential event to a focused subset of highly relevant anomalous events.
  - $t$  denotes the total number of events in the knowledge base  $KB$ .
  - $k$  is a hyperparameter that is defined and set manually based on the number of events within the knowledge base that are intended to be considered for analysis. The choice of  $k$  directly impacts the quality and completeness of the DFIR timeline analysis, while an optimally defined  $k$ , equal to the exact number of events in the controlled knowledge base  $KB$ , ensures that all evidence is included.
2. **DFIR context-specific agent workflow:** Once the initial set of relevant events has been retrieved, the DFIR RAG agent focuses on further filtering and refining these results based on contextual relevance. During this phase, the agent focuses on identifying and filtering the most contextually relevant events, ensuring that irrelevant evidence is excluded according to its assigned task, as shown in Listing 2.

**Listing 2.** Code snippet for GenDFIR Agent Prompt and Role.

```

1 DFIR_Agent_Prompt = """
2 You are a DFIR AI assistant, tasked
3 with analysing artefacts, correlating events, and producing a
4 coherent timeline of the incident. Base your answer on the
5 provided context and do not include additional information
6 outside of the context given.
7 """
8
9 DFIR_Agent = {
10     "role": "DFIR Timeline Analysis AI Assistant",
11     "content": DFIR_Agent_Prompt,
12     "maxtokens": "To define"
13 }

```

To augment the retrieval process, the framework adopts a matrix-based representation of event embeddings, defined as follows:

$$\mathcal{R} = \text{TopK}(\mathbf{V}_E \mathbf{v}_{Q_{DFIR}}^\top, k). \quad (15)$$

- $V_E$  is the matrix containing the embedding vectors of all events in the knowledge base,  $KB$ .
- $v_{Q_{DFIR}}$  is the embedding vector of the query.
- $k$  is the predefined number of top events to retrieve, which is equal to the number of events in the controlled knowledge base.

This matrix-based approach allows for the batch processing of embeddings, where similarity scores between the query and each event are computed simultaneously. To compute these scores, the query embedding  $v_{Q_{DFIR}}$  is transposed (i.e.,  $v_{Q_{DFIR}}^\top$ ) so that it can be aligned with the event embeddings in the matrix  $V_E$ . This transposition ensures that the query vector is properly oriented for the dot product calculation with each event's embedding. The events in  $\mathcal{R}$ , being both relevant and contextually aligned with the DFIR query  $Q_{DFIR}$ , are then passed to the next stage of the framework, where they serve as the foundation for generating DFIR timelines and conducting further analysis through the LLM.

At this point, the agent not only retrieves the top- $k$  evidence (relevant events) based on their similarity scores but also refines the results based on DFIR timeline analysis context-specific filtering.

3. **Timeline analysis generation:** In this phase, the framework utilises both the context provided via the retrieval process (relevant evidence,  $V_{RE_i}$ ) and the user's input,  $Q_{DFIR}$ , to generate a timeline of anomalous events. The process involves multiple steps, including attention-based mechanisms, contextual enrichment through the LLM [59], and final timeline generation via a decoder-based LLM model.

After retrieving the relevant evidence, the framework applies attention mechanisms. Attention scores are calculated between the query vector,  $v_{Q_{DFIR}}$ , and the event vectors,  $v_{RE_i}$ , allowing the model to weigh the relevance of each event. The attention score,  $\alpha_i$ , is computed as follows:

$$\alpha_i = \frac{\exp\left(\frac{\mathbf{v}_{Q_{DFIR}}^\top \mathbf{v}_{RE_j}}{\sqrt{d}}\right)}{\sum_{j=1}^k \exp\left(\frac{\mathbf{v}_{Q_{DFIR}}^\top \mathbf{v}_{RE_j}}{\sqrt{d}}\right)}, \quad (16)$$

- $d$  is the dimension of the vectors (the number of features in each vector), and
- $k$  is the number of the retrieved events.



This process amplifies the differences in similarity scores between events and the query, so events that are more similar to the query will have higher attention scores. This attention score is calculated using the dot product between  $v_{QDFIR}$  and  $v_{RE_i}$ , followed by an exponential function. The exponential function serves two primary purposes in this context. First, it emphasises retrieved incident events,  $v_{RE_i}$ , that are more relevant to the query,  $v_{QDFIR}$ , by assigning them higher attention scores, which helps these events contribute more to the final timeline generation. Second, it ensures that the difference between similarity scores is amplified in a controlled manner, preventing events with slightly lower similarity from being overshadowed or overlooked. This balance allows the model to consider both highly relevant events and those with smaller, yet still meaningful, similarity and relevance to  $v_{QDFIR}$ . Once the attention score,  $\alpha_i$ , is calculated, the next step is to enrich the context of each event. The weighted sum of event vectors, adjusted by the attention scores, represents an initial context,  $c$ , for the events, reflecting their weighted relevance to the query  $v_{QDFIR}$ :

$$c = \sum_{i=1}^k \alpha_i v_{RE_i}. \quad (17)$$

The LLM further enriches this context by processing the weighted sum of events,  $c$ , in light of its pre-existing training knowledge. For instance, Llama 3.1 has been extensively trained to recognise anomalous patterns, as well as publicly known cyber incident threats and anomalies. In this case, the context  $c_{\text{final}}$  represents a nuanced, model-based interpretation of anomalous events, incorporating the LLM's understanding,  $c_{\text{extra}}$ , of their significance based on patterns, relationships, and domain knowledge learned during training. This enriched context incorporates event correlation, where the LLM identifies and links related events based on temporal order, causality, or contextual associations.

For example, failed login attempts in a cyber incident might signal unauthorised access. The LLM correlates this event with subsequent related activities, such as privilege escalation or data exfiltration attempts, tying them to the same threat source. By analysing timestamps, anomaly severity, and underlying causes, the LLM constructs a logically sequenced timeline of events, capturing the progression and full scope of the incident. Without such automation, resolving incidents would require significant time and expertise, particularly for analysing complex event relationships and uncovering causal links. This process often demands considerable manual effort from specialists with deep technical knowledge. In contrast, GenDFIR leverages the LLM's extensive contextual understanding to accelerate and improve incident analysis, minimising reliance on manual intervention.

The final enriched context,  $c_{\text{final}}$ , derived from the attention mechanism and LLM-based enrichment, is represented as a vectorised timeline,  $VTA$ . This step involves integrating relevant evidence,  $v_{RE_i}$ , and user inputs,  $v_{QDFIR}$ , to generate the incident's timeline.

At this stage,  $VTA$  has all relevant event information in a numerical form. It reflects the incident's sequence, correlation, and anomaly patterns across multiple dimensions (e.g., event features, timestamps, and anomaly levels).

To make the timeline accessible and interpretable for human users, the structured intermediate representation  $VTA$  is decoded into a natural-language format. This transformation is achieved through a decoder-based LLM, which processes  $VTA$  and generates a human-readable, logically structured timeline of events, denoted as  $\hat{T}_A$ . This output organises forensic knowledge—such as system logs, user activities, and related events—into a chronologically ordered narrative formatted in natural language. This allows investigators

or DFIR experts to analyse temporal relationships, detect anomalies, and derive actionable insights efficiently during investigations.

The GenDFIR workflow is depicted in Figure 7 and is summarised in the following algorithm, Algorithm 1:

---

**Algorithm 1: GenDFIR Core**


---

**Input:** DFIR Expert/Analyst Query or Instruction as input  $Q_{DFIR}$ , Incident Knowledge Base  $ID$

**Output:** Auto-Generated Timeline Analysis Report relevant to the incident

$TA_{Report}$

- 1 Process Incident Document  $ID$ ;
  - 2 Load LLM Model Llama-3.1  $llama-3.1$ ;
  - 3 Load LLM Embedding Model  $mxbai-embed-large$ ;
  - 4  $Q_{DFIR} \leftarrow$   
“Conduct DFIR timeline analysis by examining the artefact, correlating events, and reconstructing the timeline of the cyber incident.”
  - 5
  - 6  $EmbedModel \leftarrow mxbai-embed-large$ ;
  - 7  $LLM \leftarrow llama-3.1$ ;
  - 8 Chunk the Incident Document  $ID$ ;
  - 9 Tokenise  $E$  and  $Q_{DFIR}$ ;
  - 10  $VE \leftarrow Embed(E, EmbedModel)$ ;
  - 11  $VQ_{DFIR} \leftarrow Embed(Q_{DFIR}, EmbedModel)$ ;
  - 12  $VRE \leftarrow Agent_{DFIR}.Process(TopK, VQ_{DFIR})$ ;
  - 13  $VTA \leftarrow Enrich\_Context(VRE, VQ_{DFIR}, LLM)$ ;
  - 14  $TA_{Report} \leftarrow LLM\_Decoder(VTA)$ ;
- 

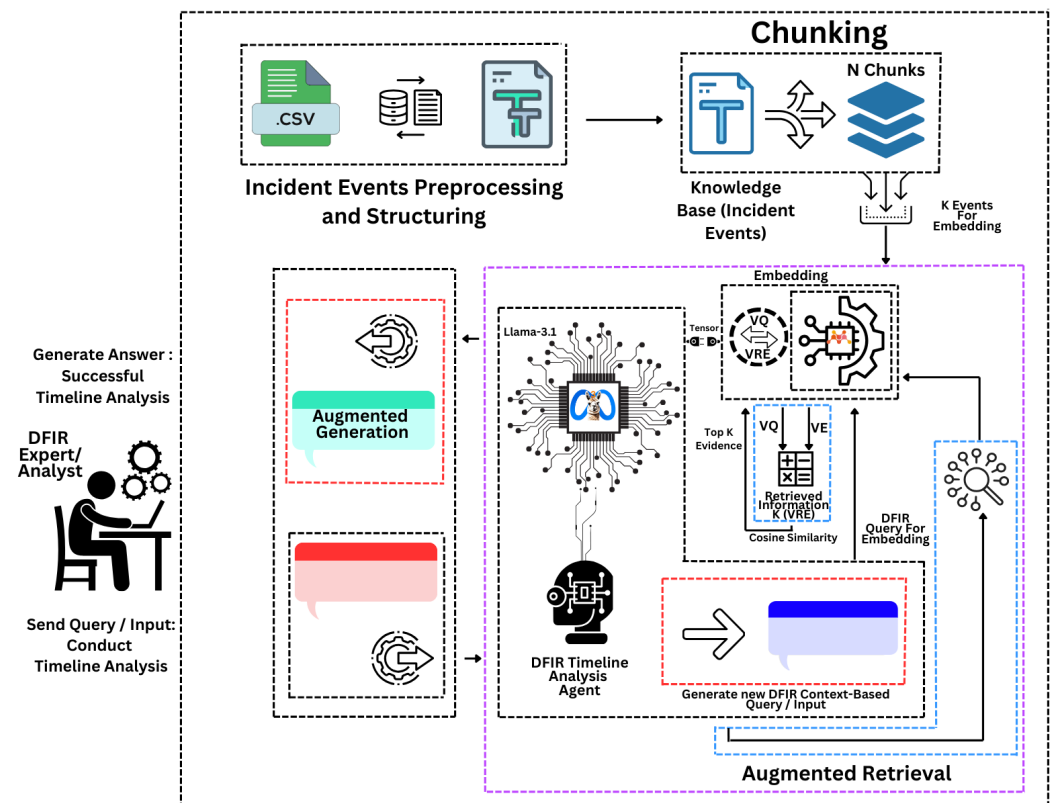


Figure 7. GenDFIR framework workflow.

## 4. GenDFIR Implementation

To properly implement the proposed GenDFIR framework, it is essential to establish clear distinctions regarding its intended purpose and scope. The GenDFIR framework is not designed to function as a conventional forensic tool that strictly adheres to legal evidentiary standards, nor is it intended for direct application in professional forensic cases that require full compliance with legal and regulatory frameworks at this stage of research. Instead, the framework's conceptualisation, development, and design serve as a foundational base, demonstrating the potential use of GenAI, specifically LLMs, within DFIR scenarios.

An important consideration for its correct implementation is the training knowledge of the LLM, especially the enriched context tied to the events in its knowledge base. LLMs raise significant concerns about the origin of their training data, which must come from legally compliant repositories in accordance with regulations such as the GDPR in Europe. For instance, Meta has chosen not to release its latest models like Llama-3.2 [60] in the EU due to regulatory uncertainties, particularly regarding the data used for its training, as stated by *The Guardian* [61]. In this research, Llama 3.1 is implemented in a zero-shot setting, which aligns with UK GDPR and security standards. No fine-tuning has been applied, ensuring the framework relies entirely on the model's original capabilities. Furthermore, the data used in the experiment are entirely synthetic, generated to simulate realistic scenarios while maintaining control over results monitoring and analysis. However, further considerations and concerns are addressed in the following lines, along with the components of the framework, the experimental environment, and the case studies used for testing.

1. **Variability in DFIR Timeline Analysis Practices and GenDFIR:** It is essential to recognise the diversity of methodologies and procedures within the DFIR industry. Each organisation or entity may adopt different investigative workflows and reporting formats, depending on the specific nature of the investigation. Broadly, DFIR processes can be divided into two primary domains: DF, which focuses on the preservation, collection, and analysis of digital evidence, and IR, which centres on the containment, eradication, and recovery from cyber incidents.

Organisations vary in their approaches to timeline analysis during investigations, as enterprise-focused IR differs from legal forensics in terms of the depth of analysis, objectives, and reporting. In enterprise environments, timeline analysis often prioritises speed and the operational impact, emphasising the rapid identification of critical events and resolution steps, sometimes with less emphasis on maintaining a formal chain of custody. In contrast, legal forensic investigations place greater importance on the preservation and reconstruction of timelines, ensuring that all digital events are accurately sequenced and documented to meet evidentiary standards for potential legal proceedings.

The reporting mechanisms for timeline analysis also differ. IR reports typically concentrate on delivering actionable insights—such as the timing of compromise, propagation, and remediation steps—often to guide immediate decision-making. These reports prioritise clarity and immediacy over legal formalities. In contrast, forensic timeline analysis reports must adhere to strict documentation standards, ensuring that every event is precisely tracked and presented in a manner that preserves the integrity of the investigation for legal scrutiny [2]. This involves the rigorous validation of timestamps, metadata, and cross-referencing multiple data sources to reconstruct a legally defensible timeline of events.

The GenDFIR framework is not a one-size-fits-all solution, as DFIR practices differ across organisations. The framework is designed as a flexible approach to integrate GenAI models into DFIR workflows, supporting experts in improving the efficiency

of investigations. The framework focuses on tasks such as log analysis, anomaly detection, and rapid evidence identification, with the purpose of providing a context-rich timeline analysis to help DFIR experts make informed decisions more quickly.

2. **Large language models in GenDFIR:** The practical application of LLMs within DFIR tasks is still novel. Recent advancements in the use of LLMs across various fields have generated interest, but they are far from being fully reliable tools for DFIR investigations [62]. Several challenges must be acknowledged, including issues of hallucination, precision, and limitations related to input/output token lengths, which can restrict the models' ability to process large datasets typical in DFIR cases. LLMs are powerful tools for automating certain aspects of data analysis, summarisation, and even generating incident reports. However, they are not yet capable of performing highly specialised forensic tasks without human oversight.
3. **Data in GenDFIR:** The LLM used in GenDFIR operates in a zero-shot mode, relying solely on its pre-existing training data to generate general-purpose DFIR timeline analysis reports. While this approach works for broad applications, it is insufficient for the specific and dynamic needs of DFIR, which requires the real-time recognition and processing of incident-specific data. This is particularly important during the phase of enriching the context of incident events, where the effectiveness of enrichment and interpretation depends on the LLM's pre-existing knowledge and its ability to process relevant data. Furthermore, each DFIR case involves unique and sensitive information, often shaped by regulatory frameworks, organisational policies, and legal requirements. As a result, incident data are highly diverse, reflecting variations in breach detection and response mechanisms.

This diversity highlights a significant gap between generalised knowledge and the nuanced demands of DFIR. Bridging this gap requires moving beyond generic outputs, such as the timeline analysis report generated via GenDFIR, and integrating incident-specific data into the model. This can be achieved by training LLMs in secure, controlled environments using curated datasets derived from real-world incidents. Such an approach enables LLMs to adapt to the distinct characteristics of various DFIR scenarios, improving their ability to identify patterns and respond effectively across different contexts.

However, the integration of real-world data into LLM training raises important privacy, legal, and ethical considerations. DFIR investigations often involve sensitive information, including personal data and proprietary organisational content. The use of incident-specific data must adhere to data protection laws, organisational consent policies, and robust security measures to safeguard against breaches during both model training and inference. These safeguards are critical to ensuring confidentiality while maximising the value of such data for future applications.

Moreover, current LLMs face limitations in adapting to new or evolving threats outside their original training scope. Without regular updates and retraining on fresh, ethically and legally sourced datasets, their performance in DFIR scenarios may degrade over time.

#### 4.1. GenDFIR Components

The following are the components integrated into GenDFIR:

- **LLM model:** Llama-3.1 a powerful model developed by Meta, was evaluated using a range of standard benchmarks and human evaluations. The rationale behind selecting this model is its open-source nature and training on cybersecurity-related content [37]. This offers significant advantages for GenDFIR, as it suggests that the LLM can analyse and potentially recognise anomalous events, providing insights into cyber incidents.

The model is available in different versions with varying parameter sizes, including 405B, 70B, and 8B parameters. At an early stage, we adopted the 8B model to reduce hallucinations and for system prototyping.

- **Embedding model:** mxbai-embed-large, a Mixedbread release, with a maximum token limit of 512 and a maximum dimension of 1024, performed highly on the Massive Text Embedding Benchmark (MTEB) [63]. The reason for opting for this model is its capabilities in retrieval and semantic textual similarities. Additionally, it is effective for scoring methods during the retrieval process, such as cosine similarity, adopted in GenDFIR.
- **DFIR RAG agent:** the Autonomous Context-Augmented Agent is powered by Llama-3.1. It is employed in GenDFIR for context-based retrieval and output generation. Operating autonomously as an expert in cybersecurity and DFIR, it specialises in DFIR timeline analysis and is tasked with identifying events and information related to cyber incidents from our knowledge base containing cyber incident events.

#### 4.2. Experiment Environment

GenDFIR and the LLM that empowers it are designed to operate locally. The data used for the experiment are synthetic and represent a variety of cyber incidents that mimic the characteristics of Windows event logs, phishing scenarios (emails), and other cyber incidents, all presented in their textual format and in English. The following outlines the details of the experimental environment and conditions in which GenDFIR was tested (Listing 3).

1. **Python libraries:** Our code was written in Python. The primary libraries used are PyTorch for tensor operations, NumPy for numerical computations such as cosine similarity, and pandas for data manipulation and handling embeddings.
2. **LLM response temperature:** The temperature of the core LLM (Llama3.1:8b) is set to a very low value (`temperature = 0.1`) to reduce perplexity and control prediction variability. This limits the generation of unnecessary content, enabling the framework to focus more effectively on the knowledge retrieved from the knowledge base, rather than relying on the LLM's additional and irrelevant internal knowledge.
3. **The max output tokens** are set to (`maxtokens = 2000`) to work conveniently with the embedding model (mxbai-embed-large:335m) and the LLM (Llama3.1:8b) specifications.
4. **The number of completions** is set to a low value (`n = 1`) to make the LLM's generation more deterministic and to select the optimal word predictions.

**Listing 3.** GenDFIR code snippet Visualisation.

```

1 answer = client.chat.
2 completions.create(
3     model=LlmModel,
4     messages=[{"role": "system", "content": prompt}],
5     maxtokens=2000,
6     temperature=0.1,
7     n=1,
8 )

```

5. **Parallel computing:** for optimised tensor computation, CUDA is used to run the tensors on the GPU (Listing 4).

**Listing 4.** CUDA and GPU.

```

1 GPU Model: NVIDIA GeForce RTX 4060
2 CUDA Version: 12.6
3 Total VRAM: 8GB

```

### 4.3. Dataset Elaboration

The experiments in this study were based on synthetic scenarios designed to mimic real-world cyber incidents while preserving privacy, consent, and confidentiality. Each scenario was customised to meet the specific requirements of the experimental setup (maxtokens = 2000) and (llama-3.1:8b) running locally), ensuring compatibility with its design and objectives. The following table, Table 4, outlines the descriptions and details of the incident scenarios used in the experiments:

**Table 4.** Cyber incident scenarios.

Scenario	Description	Number of Events (Chunks)	Max Chunk Length (Characters—MaxLength) per Event	Event Splitter
SYN Flood	<p>This is a SYN flood attack in which unusual network events disrupted standard operations. The anomalies were characterised by a high volume of synchronise (SYN) requests, causing intermittent service degradation across the network. Data were collected from firewalls, network scanners, and intrusion detection systems (IDSs). The analysis focuses on critical attributes such as event ID, details, level, timestamps, source, task category, and affected devices to assess the nature of the attack and identify potential threats or operational issues.</p> <ul style="list-style-type: none"> <li>Each event log represents a chunk.</li> </ul>	30	210	“.”
Rhino Hunt (inspired by [64])	<p>This scenario is inspired by the well-known “Rhino Hunt” incident, but in this case, it involves the illegal exfiltration of copyrighted rhino images. An unauthorised individual accessed the company’s FTP server and stole twelve protected images. The investigation traced the exfiltration to a device within the company’s internal network, with the stolen data directed to an external IP address. Forensic analysis revealed that the user associated with the IP address possessed additional images matching the stolen ones. The collected data included images that met specific metadata criteria, including camera model, artist, and copyright details.</p> <ul style="list-style-type: none"> <li>The images used in this scenario are AI-generated (using DALL·E 3) for the purpose of ensuring copyright compliance and consent. The metadata of the images were modified to align with the scenario description, and they can be extracted and viewed using the Metaminer module found in [65].</li> <li>The events in this scenario represent log entries where the context of the image was added at the beginning of each event to enrich the entry with additional context.</li> </ul>	8	500	“/”
Phishing Email—1	<p>This scenario represents a phishing attack where an employee was targeted by emails impersonating a security service. The organisation’s policy prohibits communication with untrusted domains, permitting only interactions with verified sources. All suspicious emails were collected for analysis, focusing on domains, sender and receiver details, IP addresses, email content, and timestamps to determine the nature of the attack.</p> <ul style="list-style-type: none"> <li>Each single email represents a chunk.</li> </ul>	15	725	“/”
Phishing Email—2	<p>Like in the previous scenario, this is a phishing attack where an employee was targeted by emails impersonating a trusted support service, attempting to deceive the employee into verifying account information. The organisation’s policy restricts communication with unverified domains and permits only trusted sources. All suspicious emails received during the suspected phishing period were collected for analysis, focusing on domains, sender and receiver details, IP addresses, email content, and timestamps to assess the nature of the attack.</p> <ul style="list-style-type: none"> <li>Each single email represents a chunk.</li> </ul>	20	500	“/”
DNS Spoof	<p>This incident involves a DNS spoofing scenario where multiple event logs were collected from various devices, including Windows event logs, DNS server logs, firewall records, and network traffic monitoring tools. Irregularities such as delayed DNS responses, inconsistent resolutions, and unexpected outbound traffic were identified, triggering alerts from the intrusion detection system (IDS) and performance monitoring tools.</p> <ul style="list-style-type: none"> <li>Each event log represents a chunk.</li> </ul>	23	200	“.”
Unauthorised Access	<p>This scenario is an unauthorised access attempt detected via an intrusion detection system (IDS). The system flagged multiple access attempts from a blacklisted IP address, which was not authorised for any legitimate activity within the network. The collected data, including warnings, errors, and critical alerts, provided the basis for further investigation into the potential breach.</p> <ul style="list-style-type: none"> <li>Each event log represents a chunk.</li> </ul>	25	208	“.”

## 5. Results, Evaluation, and Discussion

### 5.1. Results

The results found in [65] contain an analysis of the artefacts, which are the events in the knowledge base, their correlations, and their interpretation in relation to the incident. Additionally, context enrichment was applied to the incident events to enhance their interpretability. The framework also provides additional knowledge, such as identifying anomalous events and trends, root causes, and mitigation solutions. This reflects the role of

the agent acting as a DFIR assistant, which, in real-life scenarios, provides such information during a cyber incident. Part of these results is presented in Appendix A.

However, these results cannot be fully deemed relevant without a further evaluation, which is presented in the following section.

## 5.2. Evaluation

Since GenDFIR is a context-specific framework that employs LLMs in a zero-shot setting, we propose a tailored evaluation approach focused on assessing the framework's outputs, including the effectiveness of retrieval, evidence identification, and extraction. Traditionally, LLMs are evaluated using established benchmarks such as MMLU, which measure general-purpose language understanding and reasoning. For example, Llama-3.1 has been benchmarked by Meta, demonstrating its reliability not only for general-purpose applications but also for general cybersecurity contexts [37]. However, such benchmarks do not adequately reflect the requirements of highly specialised domains like DFIR.

Metrics like F1-score, precision, and recall, originally designed for numerical or classification tasks, fail to capture the complexity of evidence retrieval and context-specific analysis in this context. To address this, we propose and adapt metrics such as accuracy, relevance, exact match, and top-k evidence retrieval to suit the unique demands of this framework. In addition, we briefly highlight the need for DFIR industry benchmarks as a critical requirement for deploying GenDFIR at an organisational scale. The design and application of these metrics are described in the following sections:

1. **Accuracy:** Accuracy in GenDFIR is determined by evaluating the generated timeline analysis reports to identify and verify the factual content, as well as to assess how and to what extent these reports can be useful and reliable in assisting DFIR experts. This process involves cross-referencing the generated information with a verified knowledge base (incident events) to ensure correctness and relevance while confirming that interpretations are logical and reasonable. The reports are divided into two sections:
  - (a) **Knowledge Base (Events) Facts Section:** This section includes artefact analysis, timeline analysis, event correlation, and timeline reconstruction. These components are derived from facts retrieved from the knowledge base and enriched with contextual information via the LLM (Llama-3.1).
  - (b) **Additional Insights Facts Section:** This section provides supplementary information, including mitigation strategies, recommendations, and other relevant insights. These outputs are generated using the general knowledge of the LLM (Llama-3.1) and its contextual application to the knowledge base (incident events).

The accuracy of GenDFIR reports is evaluated using a proposed equation designed to quantify the proportion of verified correct facts relative to the total number of facts generated. The equation is defined as follows:

$$\text{Accuracy} = \frac{\text{Overall Correct Facts}}{\text{Overall (Correct Facts + Incorrect Facts)}} \quad (18)$$

where, according to Table 5, the following applies:


- **Overall correct facts** are the total number of accurate facts obtained by combining the correct facts from the incident knowledge base with the correct facts generated via the LLM, as shown in Table 5
- **Correct facts** are facts that are verified and deemed accurate individually from both the incident knowledge base and the LLM-generated facts.
- **Incorrect facts** are, similarly, facts that are inaccurate and identified separately from each source.

Based on the extensive monitoring and assessment of the timeline analysis reports and results, as illustrated in Table 5, minor incident events found in the generated report, along with their interpretation and retrieval, were identified as incorrect in some scenarios. For instance, in Phishing Email-1, it was reported that “2. Follow-up emails attempting to get Michael to verify his account through a link (10:45 a.m., 03:00 p.m., and 04:30 p.m.)” is incorrect, as the time (04:30 p.m.) does not exist and is not mentioned in the incident events knowledge base. On the other hand, the additional knowledge generated via the LLM was all correct. For example, in the SYN-FLOOD scenario, the key discoveries generated and found in the timeline report were relevant “**Key Discoveries**” \* Multiple SYN flood attacks detected throughout the day. \* High volume of SYN packets and excessive packet volumes exceeded thresholds. \* Regular network performance checks revealed potential issues and were successfully cleared.” Even though these discoveries and knowledge are general knowledge related to any type of SYN flood attack, they were accurate in this context.

The following figure, Figure 8, illustrates the accuracy rate for each scenario, calculated according to the previously given equation and results found in Table 5:

**Table 5.** Timeline analysis report facts.

Scenarios	Overall Timeline Analysis Report Facts	Incident Knowledge Base	LLM Facts (All Correct)
SYN Flood	20	- Correct: 17 - Incorrect: 0	03
Rhino Hunt	16	- Correct: 08 - Incorrect: 0	08
Phishing Email—1	13	- Correct: 09 - Incorrect: 1	03
Phishing Email—2	22	- Correct: 13 - Incorrect: 03	06
DNS Spoof	19	- Correct: 14 - Incorrect: 0	05
Unauthorised Access	22	- Correct: 14 - Incorrect: 01	07



**Figure 8.** Overall Accuracy of GenDFIR timeline analysis reports for different cyber incident scenarios.

Calculating the average accuracy across all scenarios revealed that the overall accuracy of GenDFIR is 95.52%.



2. **Relevance:** In this context, we define the relevance metric as the extent to which the answers provided through the framework are pertinent. To measure this, we crafted 20 DFIR context-specific prompts (available on [65]—240 Prompts in total) in the form of questions for each scenario related to the incident. These prompts are characterised by their focus on various aspects, including sentiment (e.g., “What is the overall sentiment of the image described in Event 2 (Image b)?”—Rhino Hunt Scenario), intention (e.g., “What specific action is requested by GlobalBank in the initial email?”—Email Phishing-1 Scenario), and deep analysis (e.g., “Analyze the severity levels associated with the events involving SYN flood attacks. List the events and their corresponding severity levels.”—SYN-FLOOD Scenario), as well as retrieval, prediction, and insights. According to the generated answers, GenDFIR failed to provide correct responses for certain questions (DFIR prompts). For instance, in the SYN Flood scenario, GenDFIR failed on Prompt 4; in the Rhino Hunt scenario, Prompt 13 was answered partially correctly (50%), and in the Phishing Email 1 scenario, GenDFIR failed at both Prompts 11 and 12. The correctness and incorrectness of these answers were compared and cross-referenced with the incident events knowledge base, which is considered the ground truth data. The detailed failures and successes are shown in the graph in Table 6:

**Table 6.** Prompts’ relevance (in percentages) in GenDFIR timeline analysis reports for different cyber incident scenarios.

Prompt	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
SYN Flood	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Rhino Hunt	1	1	1	1	1	1	1	1	1	1	1	1	0.5	1	1	1	1	1	1	1
Phishing 1	1	1	1	1	1	1	1	1	1	1	0	0	1	1	1	1	1	1	1	1
Phishing 2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	1	1	1	1	1
DNS Spoof	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Unauthorised Access	1	1	1	1	1	1	1	0.25	1	1	0.66	1	1	1	1	1	1	1	1	1

The overall relevance rate of GenDFIR in this case is 94.51%.

3. **Exact match:** The EM metric in GenDFIR is designed to evaluate the framework’s ability to retrieve precise and granular information from the knowledge base. It tests the framework’s performance by posing DFIR questions that require exact matches to the ground truth data. This includes specific details such as timestamps (e.g., “What was the exact time when Michael Davis responded to the first email from the GlobalBank Security Team?”) and descriptions (e.g., “In Event 5, what was the description of the alligator?”), along with other relevant queries. The results found in [65] show that GenDFIR performed well, aligning perfectly with the ground truth. GenDFIR successfully passed all checks without any errors. Therefore, the EM in this case is rated at 100%, as all 20 prompts were answered correctly.
4. **Evidence (Top-k):** DFIR timeline analysis is primarily focused on identifying evidence through the analysis of digital artefacts, with the correlation between events linked to the relevance of evidence deemed pertinent to the incident. To facilitate the real-time visualisation of evidence identification during inference, we developed a script, illustrated in Listing 5, which uses the UMAP library.
  - **UMAP library adaptation for GenDFIR:** The UMAP library contains algorithms used for dimensionality reduction and visualisation, enabling the identification of the closest data points in a 2D vector space. These algorithms work by calculating the distances between data points, selecting the nearest neighbours, and

generating a graphical representation that illustrates how data points are related, preserving both local and global structures [66]. In the context of GenDFIR, we developed a script using this library to visualise the relationship between features from the vectorised user input  $VQ_{DFIR}$  and their corresponding features in the embedded knowledge base (incident events,  $VE$ ). This process helps identify how input data from the DFIR expert align with relevant incident events.

It is important to note that evidence identification in this scenario depends on the input provided by the user, who determines which events or attributes should be considered as evidence. In other words, the evidence is solely based on the query prompts issued by the DFIR expert, whose suspicious assessment guides the process. As discussed in the methodology, GenDFIR integrates algorithms based on cosine similarity and a DFIR-specific agent to assess relevance. In UMAP, we added a heat parameter to visualise the similarity during inference between the query from the DFIR expert and the events embeddings from the knowledge base. This heat parameter ranges from 0 to 1, where a value closer to 1 indicates a stronger similarity between the query and the events embeddings, and a value closer to 0 indicates a weaker similarity. The colour intensity in the visualisation reflects this similarity, with higher values (closer to 1) represented by warmer colours, signifying that the events are more relevant to the query. In addition to this, the script leverages the traditional functionality of the UMAP algorithm, which clusters data points according to how similar they are to each other.

**Listing 5.** Code snippet for Calculating and Visualising Similarity.

```

1 def calculate_cosine_similarity(VQ, VE):
2     """
3     Calculate cosine similarity between the input and Knowledge Base
4     embeddings.
5     Args:
6     - VQDFIR (list or np.array): Embedding of the query.
7     - VE (list or np.array): Embeddings of the Knowledge Base content.
8     Returns:
9     - torch.Tensor: Cosine similarity scores between input and Knowledge
10    Base embeddings.
11    """
12    cos_scores = torch.cosine_similarity(torch.tensor(VQ).unsqueeze(0), VE)
13    return cos_scores
14
15 def visualise_cosine_similarity(cos_scores):
16     # Convert cosine similarity scores to 2D
17     scores_2d = np.array(cos_scores).reshape(-1, 1)
18     # Apply UMAP for dimensionality reduction
19     embedding = umap.UMAP(n_neighbors=5, min_dist=0.1, random_state=42).
20     fit_transform(scores_2d)
21     # Create a scatter plot using Matplotlib
22     scatter = plt.scatter(embedding[:, 0], embedding[:, 1], c=scores_2d.
23     flatten(), cmap='viridis')
24     plt.show()
25
26 cos_scores = calculate_cosine_similarity(VQDFIR, VE)
27 # Visualise the cosine similarity scores
28 visualise_cosine_similarity(cos_scores)

```

To interpret the graphs generated via UMAP, data points that are close to each other indicate higher contextual similarity. During inference, events deemed relevant to the query will be positioned near one another in the visualisation. The heat parameter,

represented by colour variation, reflects the degree of similarity between the data points, with warmer colours indicating a stronger match to the expert's query.

In the code, the `calculate_cosine_similarity` function computes the cosine similarity between the vectorised input  $VQ_{DFIR}$  and the knowledge base embeddings  $VE$ . The `visualise_cosine_similarity` function then applies UMAP to reduce these similarity scores to a 2D space, where proximity indicates stronger similarity between events. The resulting visual representation enables the DFIR expert to better understand the alignment between the query and the knowledge base.

Figures 9–14 illustrate the UMAP evidence visualisations. The number of Top K shown in these figures accurately corresponds to the evidence manually identified in the incident event knowledge base, as well as the interpretation of the similarity between the prompt and the events.




Figure 9. SYN flood—evidence (top K = 5).




Figure 10. Rhino Hunt—evidence (top K = 4).




Figure 11. Phishing Email 1—evidence (top K = 6).




Figure 12. Phishing Email 2—evidence (top K = 7.)


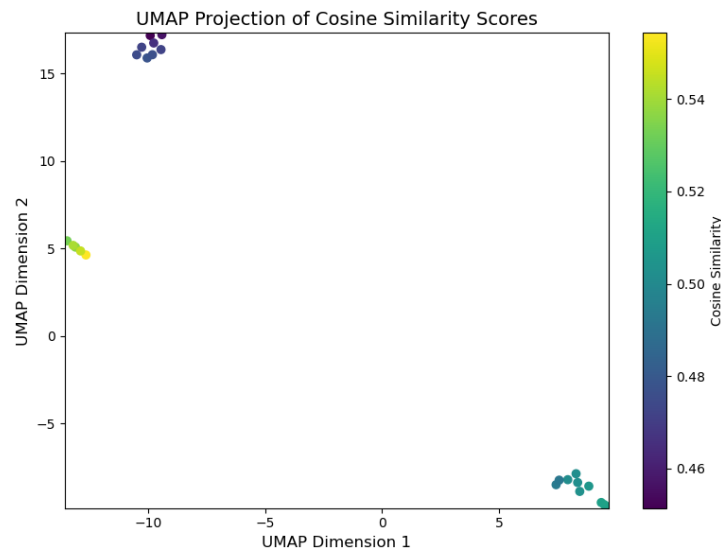


Figure 13. Unauthorised Access—evidence (top K = 25).



**Figure 14.** DNS Spoof—evidence (top K = 6).

Figure 13 provides an example where all events are considered evidence. The knowledge base for this scenario contains events where the attribute “Level” is set to “Warning”. The ground truth data in Table 7 specify a total of 25 pieces of evidence (Unauthorised Access Scenario), based on the user’s input: “Identify all events where the Level is ‘Warning’”. However, Figure 13, generated using UMAP, displays all 25 data points clustered together, indicating that both neighbourhood and context similarity are present. Despite this, the cosine similarity colour in the heat parameter for some data points appears colder, with a cosine similarity score of  $\leq 0.54$ . This discrepancy suggests that, while all events share the same “Level” attribute, other features or contextual factors within the events (such as timestamps, event descriptions, or sources) contribute to their embedding representations, influencing their overall similarity. The colder colours indicate that, despite sharing the same “Level”, these events may differ in other attributes or in how they relate to the query’s context, particularly with the presence of the internal DFIR context-specific agent and its internally assigned role.

**Table 7.** Top K Evidence within the knowledge base (incident events).

Cyber Incident	Criteria of Evidence (Evidence Extraction Prompt)	Incident Events—K	Evidence—Top K
SYN Flood	Identify all events with Level: Critical.	30	05
Rhino Hunt (inspired by [64])	Rhino.	08	04
Phishing Email—1	Identify all events that appear to be phishing.	15	06
Phishing Email—2	Identify all events that appear to be phishing.	20	07
Unauthorised Access	Identify all events with Level: Warning.	25	25
DNS Spoof	Identify all events with Level: Error.	23	06

However, with this metric, we focus solely on evidence retrieval and identification, which has been successfully demonstrated, as all 25 pieces of evidence are clustered

together. Another aspect to consider in the visualisation of Figure 13 is the size of Dimension 2, which ranges from 0 to 3.2, a range that is reduced when compared, for instance, to Figure 14. If increased, the data will appear closer together.

Table 7 contains the number of ground truth evidence (Incident Events—K) for each cyber incident, as determined by the user, the instructions prompted by the DFIR expert, and the Top K evidence generated via GenDFIR:

#### 5. *DFIR Industry Benchmark*

Since industry entities typically rely on a variety of tools tailored to their specific needs and practices, creating a universal benchmark for GenDFIR in DFIR timeline analysis presents a challenge. While established industry standards such as NIST SP 800-86 [67] and ISO/IEC 27037 [68] provide guidelines for the general process, they do not offer a precise or dedicated framework for timeline analysis.

In this paper, we propose a tailored solution for how GenDFIR could potentially be approached for industry benchmarking. This benchmark aims to align the framework's reliability, robustness, and performance with industry standards, particularly those governing forensic integrity and accuracy, with a particular focus on the trustworthiness of the generated outputs in reconstructing accurate incident timelines:

- **Human verification and inherent llm limitations in timeline analysis:** While GenDFIR is designed to assist investigators in timeline analysis, it does not replace human expertise. As shown in previous studies and industry experience, LLMs can produce hallucinations—incorrect or irrelevant information—particularly when generating long sequences of text or handling complex contexts. These limitations are inherent to the LLM architecture, where the models predict sequences of words based on probability, often leading to incoherent or irrelevant outputs. This challenge is particularly critical in timeline analysis, where even small inaccuracies can have significant consequences for the investigation. According to NIST 800-86 and other forensic guidelines, human oversight is essential to validate that the reconstructed timeline is both accurate and legally defensible [67]. Forensic standards consistently emphasise the importance of human expertise in reviewing and verifying automated results, particularly in high-stakes contexts such as legal proceedings or regulatory investigations. By integrating human verification at key stages, GenDFIR enhances investigators' ability to grasp the broader context of an incident. This allows them to concentrate on decision-making and strategic response while automating the more time-consuming aspects of evidence examination.
- **GenAI in DFIR:** As of now, there is no established benchmark or standard that specifically guides the integration of GenAI—particularly LLMs—into forensic or incident response practices. However, the absence of a clear and comprehensive established framework for incorporating GenAI into DFIR leaves many questions about best practices, performance expectations, and the limitations of such technologies unresolved. The use of GenAI and LLMs involves handling contextual data during training, inference, and interaction. This introduces substantial risks related to privacy and confidentiality, arising from the data collection process through to LLM output generation. Solving these challenges could potentially lead to the development of advanced dedicated, industry-specific benchmarks for integrating GenAI into DFIR, providing a structured foundation for its reliable and ethical application.
- **Optimising performance in large-scale deployments:** To ensure the effectiveness of GenDFIR in large-scale timeline analysis, especially in complex incidents, it is essential to optimise the deployment configuration. In alignment with indus-

try standard practices, including NIST standards for computational forensics [67], the performance of LLMs is highly sensitive to parameter tuning, particularly the temperature setting. The temperature controls the LLM's perplexity by influencing how deterministic or diverse the model's predictions are. Large-scale deployments also require a system architecture capable of handling vast amounts of data efficiently. This can be achieved by extending the token limits of both the embedding and LLM models, allowing the framework to process larger outputs. When combined with high-performance GPU configurations, GenDFIR is better equipped to scale for more complex investigations, producing outputs that are both precise and reliable.

- **Knowledge base management and regular output review:** Reliable timeline analysis with GenDFIR also requires the effective management of the knowledge base used in the system. This knowledge base must be both up-to-date and accurate, as it directly impacts the accuracy of evidence extraction and timeline reconstruction. Industry standards, such as ISO/IEC 27037 [68], emphasise that forensic tools must undergo continuous validation to ensure that they produce accurate and reliable results. This ongoing review process aligns with the regular verification of both inputs (evidence embedding from the knowledge base) and generated outputs, ensuring that timeline analysis remains in sync with the latest evidence and incident context.

Additionally, the knowledge base should be handled locally in secure storage vaults, rather than relying on online systems. While modern cloud infrastructures offer advanced security features, local storage is often considered more reliable for forensic purposes. According to NIST Special Publication 800-53 [69] on security and privacy controls for information systems, local storage can offer more control over sensitive data and is less vulnerable to the risks associated with cloud-based systems, such as third-party access or data breaches. Forensic guidelines generally recommend maintaining control over the entire storage environment, particularly when handling sensitive evidence.

A critical aspect of evidence extraction in timeline analysis is its role in supporting legal claims or formal procedures. As such, forensic experts must deploy GenDFIR in an optimised state during evidence extraction tasks. This involves ensuring that the system is configured correctly, that the knowledge base is accurately managed, and that human verification is integrated to safeguard the integrity of the most sensitive aspects of evidence handling.

- **Deployment for legal and formal use:** For timeline analysis to be suitable for legal and formal use, GenDFIR must meet the rigorous standards required by forensic standard practices. Tools used in legal investigations must demonstrate a high level of accuracy, precision, and admissibility in legal contexts, as outlined by forensic guidelines such as ISO/IEC 27037 [68] and ISO/IEC 17025 [70]. Experts must verify that GenDFIR is deployed according to its high-level requirements, affirming that all outputs used for legal defences or formal investigations are reliable and defensible.

### 5.3. GenDFIR Overall Performance

The overall performance of GenDFIR is evaluated by calculating the aggregate accuracy across each metric, as presented in Table 8:

**Table 8.** Overall performance.

Metric	Rate
Accuracy	95.52%
Relevance	94.51%
EM	100%
Top-K	100%
<b>Overall</b>	<b>97.51%</b>

The results demonstrate a high performance rate, suggesting that, at this stage of development, the framework shows promise in terms of reliability and precision for DFIR investigations and timeline analysis. It is important to note, however, that these assessments were conducted under specific configurations and experimental conditions, which may affect their generalisability to other contexts and according to formal DFIR industry benchmarks.

#### 5.4. Comparison with Other DFIR Tools

GenDFIR addresses critical needs in DFIR timelines by responding to current trends and limitations in existing tools. To demonstrate its application, we compare its capabilities, derived from our results, with those of other tools (Table 9):

**Table 9.** Comparison of GenDFIR with other tools.

Tool	Artefact Analysis	Anomaly Detection	Event Correlation	Full Contextual Timeline Report
Splunk	Yes	Yes	Yes	No
FTK, Autopsy	Yes	No (basic detection)	No (manual correlation)	No
Velociraptor	Yes	Yes (basic)	Yes (automated)	No
Timesketch	Yes	No (basic detection)	Yes (timeline-based correlation)	No (visual timeline only)
GenDFIR	Yes (Basic)	Yes (basic anomaly detection for basic incidents)	Yes (automated)	Yes (full report with additional context)

#### 5.5. Discussion

According to our experiment, results, and the overall performance of the framework, and as previously discussed, GenDFIR should not be viewed as a universal solution for all DFIR tasks. Instead, it is designed as a foundational framework to address specific challenges faced by DFIR analysts or experts, particularly in the analysis of artefacts.

The evaluation results, as illustrated in the previous graphs, show a success rate ranging from 90% to 100%, depending on the metrics used. It is important to emphasise that these results are based on simplified experimental scenarios and minimal configuration settings. Moreover, the primary purpose of the generated report is to support the investigation process by providing general contextual knowledge and cybersecurity information related to the incident. In practice, casual and precise DFIR timeline analysis reports differ significantly, as both play a crucial role in documenting the sequence of events in an incident. These reports typically focus on the following:

- *Event chronology*: a detailed timeline of events, presenting a chronological order of activities related to the incident.
- *Evidence correlation*: an analysis of how various pieces of evidence relate to the timeline.



- *Incident overview*: a summary of the incident, including key findings and impacts. Additionally, these reports are often tailored to different audiences [71]:
- *Technical stakeholders*: forensics experts and IT professionals require detailed, technical reports with precise timestamps, technical analysis, and evidence correlation. For example, a technical report may include granular timestamps of system log entries and detailed forensic data.
- *Non-technical stakeholders*: Reports for non-technical audiences, such as senior management or legal teams, present a streamlined timeline with key events highlighted. These reports focus on the overall sequence of events and impacts, formatted to be accessible and understandable to non-experts. An example is an executive summary, followed by a simplified timeline of events.

In the context of GenDFIR, the report templates were generated via the LLM, guided by both human prompts and the internal DFIR agent's prompt. Furthermore, the generated reports were constrained by the token limits of the selected LLM and the chosen embedding model. Consequently, if both the model and the token output size were altered to incorporate a larger model, the timeline analysis report would likely be longer and exhibit differences in content. However, the reports were automatically condensed within the scope of our current configuration, which is limited to a maximum of 2000 tokens due to the LLM's token constraints.

Another key aspect that underscores the utility of GenDFIR is its capability to identify evidence. This framework can readily identify evidence based on suspicions, assumptions, or specific elements intended to be monitored during the investigation. As previously demonstrated, evidence retrieval is facilitated by prompting the framework with context-specific DFIR instructions or queries, allowing for the targeted identification of relevant events.

## 6. Limitations, Ethical Issues, and Future Works

### 6.1. Limitations

Despite the success of bringing the proposed framework to life and successfully generating timeline analysis reports, this research encountered several limitations:

- **Novelty of approach**: The application of LLMs to DFIR timeline analysis is a relatively novel approach. The limited existing research on automating timeline analysis means there are few established guidelines or benchmarks, which has necessitated the development of a new context and methodology for this task.
- **Data volume and variety**: The large volume and heterogeneity of data in cyber incident scenarios present significant challenges for data management and processing [28]. Real-world cyber incidents often introduce unique challenges due to their distinct circumstances, requiring tailored approaches. The data generated during these incidents can vary significantly, and the rules or triggers behind each incident often differ [72], further complicating the implementation of universal solutions.

In addition, the dynamic and evolving nature of cyber threats necessitates a deep understanding of their root causes and the interconnected factors contributing to their emergence. For example, an organisation's infrastructure vulnerabilities and the interplay of interconnected devices can collectively expose weaknesses. Often, vulnerabilities are not confined to a single asset but result from a combination of flaws across interconnected systems, amplifying the complexity of incident response.

In the case of GenDFIR, our experiments utilised synthetic and simplified cyber incident scenarios, intentionally excluding advanced anti-forensics techniques. This approach allowed us to focus on developing and validating the core conceptual

framework in controlled settings. However, this introduces limitations. Real-world cyber incidents typically involve sophisticated adversaries who employ anti-forensics and other evasion techniques, aspects not considered in this study.

Furthermore, transitioning from these experimental scenarios to real-life applications will require access to diverse, realistic datasets, which often involve legal and ethical considerations.

Additionally, the scalability challenges of handling large datasets in incident response will require upgrades to hardware resources (e.g., high-performance GPUs) and adjustments to framework configurations, particularly with respect to LLM capabilities, such as token limits. Employing optimisation techniques like instance parallelism, batch size adjustment, and advanced resource allocation (e.g., distributed processing and load balancing) will further support the framework's expansion.

- **Privacy and data security:** Handling DFIR cases inherently involves processing sensitive, confidential, and often personal data, including internal organisational information, customer records, or technical system logs. In our experiments, we mitigated these privacy risks by using synthetic datasets derived from simulated cyber incidents. This approach eliminates reliance on real-world data, ensuring compliance with strict legal and ethical standards, such as the GDPR in the EU, and the NIS2 Directive, which specifically addresses cybersecurity and data integrity in critical sectors such as finance, energy, and healthcare.

One core challenge of leveraging LLMs like GenDFIR is their dependency on substantial data inputs, including operational and historical data from past incidents, to provide tailored and accurate responses. This poses two critical concerns: maintaining data confidentiality and adhering to the legal frameworks governing data retention and use. For instance, fine-tuning an LLM with detailed historical incident data—such as attack patterns, mitigative responses, or forensic artefacts—may risk breaching GDPR's "purpose limitation" principle if data originally collected for investigation are reused without explicit consent. Similarly, the NIS2 Directive mandates robust safeguards when handling logs and metadata in industries critical to public safety, such as encrypting sensitive logs during transit and at rest.

To address these challenges, organisations adopting GenDFIR should consider implementing advanced privacy-preserving techniques, including the following.

- **Data anonymisation:** Removing identifiable attributes (e.g., IP addresses, usernames, and email domains) from training data to ensure that traceability back to individuals or assets is eliminated while preserving the context of incidents.
  - **Federated learning:** enabling the collaborative training of models directly on distributed, secured systems within organisational premises.
  - **Differential privacy:** applying algorithmic modifications to introduce statistical uncertainty to query results, which prevents the reverse-engineering or exposure of individual records, even in shared environments.
- **Evaluation methods:** Assessing the effectiveness and performance of our framework presents significant challenges at this stage, particularly in the context of DFIR practices. The current research on automated evaluation methods for LLMs and AI systems in domain-specific applications is still in its early stages, and most approaches rely heavily on manual evaluation. In our case, this necessitated the creation of custom DFIR context-based evaluation prompts to measure performance. However, even with the proposed metrics, we contend that they are insufficient. The DFIR field is inherently complex and non-deterministic, with results often influenced by a wide range of unpredictable factors, such as varying incident types, data quality, and contextual nuances. These variables complicate the ability to derive consistent, controlled outcomes,

making automated evaluation. Moreover, the dynamic and rapidly evolving nature of cyber incidents demands evaluation methods that can adapt to real-world scenarios' technical and contextual variables, presenting an ongoing challenge to develop robust, reliable assessment frameworks.

## 6.2. Ethical Issues

Introducing GenAI into the DFIR field presents serious ethical issues that are critically important and require meticulous consideration, given the sensitive and private nature of the data, practices, and activities involved. The following are the most relevant and common concerns [73]:

- **Privacy and confidentiality:** The use of LLMs to advance cyber incident timeline analysis can lead to significant privacy breaches. These technologies often require access to and process vast amounts of sensitive data, including personal, digital, financial, health, and other types of information, all of which must comply with security standards and frameworks. This increases the risk of exposing personal and confidential information without proper safeguards.
- **Accuracy and efficiency:** Ensuring the accuracy and reliability of automated analyses and generated reports from the LLM-powered framework is critical. Inaccurate results could lead to flawed or irrelevant conclusions and decisions by investigators and DFIR experts, potentially affecting legal proceedings and justice.
- **Consent:** Obtaining proper consent from individuals whose data is used for automatic processing in real-life scenario adaptations by the framework is essential. Without explicit consent, using such data would directly violate their privacy rights.
- **Bias, hallucination, and fairness concerns:** LLMs are naturally exposed to the possibility of reinforcing biases hallucination and fairness issues, potentially leading to unfair outcomes during cyber incident analysis. In the context of GenDFIR, the process of context enrichment—where events from a knowledge base are interpreted and connected—can inadvertently propagate existing biases or generate skewed insights. This compromises the impartiality and credibility of investigations, particularly in high-stakes scenarios where decisions must be based on objective evidence.

By design, they generate responses based on patterns and correlations learned from their training data. This predisposition means biases present in the training dataset can directly influence outputs, while hallucinations may result in fabricated details unrelated to the input data or facts. Although these issues cannot be eradicated fully, mitigation strategies can significantly reduce their occurrence.

The success of GenDFIR is largely tied to the environment in which it is deployed and the alignment of its configuration to domain-specific contexts. Several measures can minimise bias and hallucinations:

- **Fine-tuning on domain-specific data:** through the training or fine-tuning of the LLM on curated, unbiased, and context-specific datasets, outputs can become more aligned with the objectives and operational needs of DFIR workflows.
- **Prompt engineering:** carefully structuring prompts can guide the model to provide more accurate, more relevant, and less biased responses.
- **Bias audits and testing:** conducting regular assessments of the model's outputs for indicators of bias and other errors can help in identifying and rectifying problematic patterns.

At this stage, LLMs, including their implementation in GenDFIR, should not be regarded as fully autonomous systems capable of ultimate reliability. The outputs must always be subject to scrutiny, with their limitations clearly acknowledged, as LLMs are not substitutes for expert judgment and careful analysis.

- **Automatically detected evidence:** in cases where the framework is employed to detect and retrieve evidence autonomously, improper internal processing or analysis could compromise its reliability and validity.

### 6.3. Future Works

The future direction of this framework focuses on mitigating its limitations by improving its performance, scalability, and applicability in real-world contexts. A key objective is the integration of real-world cyber incident data to assess the framework's outputs based on credible data, rather than synthetic sources. This will allow for an analysis of a broader range of cyber threats, including those involving anti-forensic techniques, and it will improve the framework's ability to address the evolving nature of complex security incidents. The data used will adhere to DFIR legal standards, ensuring that proper consent is obtained where necessary.

To increase the system's functionality, specialised autonomous LLM agents will be implemented, each designed for specific tasks such as timeline analysis, evidence correlation, and root cause identification. This modular approach will improve operations, particularly in more complex incidents. Traditional AI metrics, such as the F1-score, precision, and recall, are inadequate for evaluating outputs that emphasise contextual and narrative-driven reporting, leading to the development of novel methodologies and metrics tailored to this purpose. The integration of GenAI into DFIR processes presents a significant dilemma. While frameworks such as the NIST AI Risk Management Framework (AI RMF) [74] offer general guidance for AI governance, they fail to address the specific challenges posed by forensic applications. Of particular concern is GenAI's capacity to generate fabricated data based on its training, raising critical questions about the reliability and authenticity of outputs used in investigations. This gap presents an opportunity to develop specialised regulatory frameworks tailored to this domain.

The performance of the framework will also be optimised by deploying larger language models that can handle more tokens, as this will make the framework more efficient in processing large datasets and handling complex tasks in high-demand environments. The scope of the framework will be extended to other sectors, starting with digital forensics. We will explore its application in regulatory compliance and legal proceedings, addressing sector-specific needs. Additionally, the framework will be tested in the context of IoT security, allowing the analysis of data from connected devices and supporting forensics in this growing area. Thus, with these proposed directions, the framework is expected to perform more effectively and reach a new level of capability.

## 7. Conclusions

This research presents a novel framework for automating cyber incident timeline analysis, addressing the longstanding challenge of deriving human-understandable insights from complex digital artefacts. By integrating a DFIR-specific retrieval-augmented generation (RAG) agent with the Llama 3.1 8B model, the framework introduces the unique capability of generating coherent timelines, extracting contextual incident insights, and responding to diverse investigative queries.

The framework solves key challenges in incident event interpretation by combining semantic correlation, anomaly detection, and evidence retrieval. It achieves this through the advanced use of cosine similarity to align investigative queries with domain-specific knowledge, along with an agent-driven filtering mechanism that prioritises relevance. This innovation surpasses conventional tools in its ability to contextualise events dynamically, providing a fresh perspective on AI based automation in DFIR practices.

Despite operating within a proof-of-concept stage, the framework demonstrates significant potential for scalability and adaptability in real-world scenarios. Future efforts will focus on mitigating minor noise observed during large-scale artefact analysis and validating the framework's performance across more complex environments. These advancements aim to establish a robust tool for expert practitioners, contributing to the evolution of AI applications in cyber incident timeline analysis investigations.

**Author Contributions:** Conceptualisation, F.Y.L. and M.C.G.; methodology, F.Y.L. and M.C.G.; software, F.Y.L. and M.C.G.; validation, F.Y.L. and M.C.G.; formal analysis, F.Y.L. and M.C.G.; investigation, F.Y.L. and M.C.G.; resources, F.Y.L. and M.C.G.; data curation, F.Y.L. and M.C.G.; writing—original draft preparation, F.Y.L.; writing—review and editing, F.Y.L. and M.C.G.; visualisation, F.Y.L.; supervision, M.C.G. and M.A.F.; project administration, M.C.G.; funding acquisition, M.C.G. All authors have read and agreed to the published version of the manuscript.

**Funding:** This APC and open access for this article were funded by the University of Liverpool.

**Institutional Review Board Statement:** This research was deemed not to require the University's Ethical Committee approval, as it does not fall under any of the cases requiring ethical approval.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** We acknowledge that all data used in this case study are synthetic and used solely to model and visualise cyber incident scenarios while strictly adhering to privacy, confidentiality, consent, and copyright regulations. All case studies, the knowledge base (incident events—ground truth), the GenDFIR timeline analysis reports, and the evaluation results are available at <https://github.com/GenDFIR/GenDFIR> (accessed on 2 November 2024).

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

## Appendix A. GenDFIR Results

These results consist of the timeline analysis reports generated via GenDFIR. As the reports are quite large in size, we only showcased one example, with others available in [GenDFIR Results](#) [65]. To generate the reports, it is crucial to first visualise the knowledge base. For this paper, we chose to showcase the "Unauthorised Access" scenario.

### Appendix A.1. GenDFIR Knowledge Base

Table A1 illustrates our knowledge base, which also serves as the ground truth. This knowledge base is uploaded to the framework to function as its dynamic memory, enabling the retrieval of events as needed (either simultaneously or at different times), depending on the context.

**Table A1.** Incident knowledge base (incident events—ground truth data).

---

Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:15:05, Source: Windows Security, Task Category: Logon.
Event ID: 5038, Details: Integrity of system file failed, Level: Error, Date and Time: 2016/09/15 02:16:37, Source: System, Task Category: System Integrity.
Event ID: 4723, Details: Attempt to change account password, Level: Warning, Date and Time: 2016/09/15 02:17:05, Source: Windows Security, Task Category: Account Management.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:17:21, Source: Windows Security, Task Category: Logon.
Event ID: 1102, Details: Audit log cleared, Level: Critical, Date and Time: 2016/09/15 02:18:05, Source: Security, Task Category: Audit Logs.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:18:55, Source: Windows Security, Task Category: Logon.
Event ID: 4769, Details: Kerberos service ticket request failed, Level: Warning, Date and Time: 2016/09/15 02:20:12, Source: Security, Task Category: Credential Validation.

---

**Table A1.** *Cont.*


---

Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:21:00, Source: Windows Security, Task Category: Logon.
Event ID: 4771, Details: Pre-authentication failed, Level: Warning, Date and Time: 2016/09/15 02:22:30, Source: Security, Task Category: Credential Validation.
Event ID: 4648, Details: Logon attempt with explicit credentials, Level: Warning, Date and Time: 2016/09/15 02:23:35, Source: Windows Security, Task Category: Logon.
Event ID: 4724, Details: Password reset attempt, Level: Warning, Date and Time: 2016/09/15 02:24:12, Source: Windows Security, Task Category: Account Management.
Event ID: 5031, Details: Firewall service encountered an error, Level: Warning, Date and Time: 2016/09/15 02:24:45, Source: Security, Task Category: Firewall.
Event ID: 4776, Details: Logon failure due to incorrect credentials, Level: Warning, Date and Time: 2016/09/15 02:27:15, Source: Windows Security, Task Category: Credential Validation.
Event ID: 4771, Details: Pre-authentication failed, Level: Warning, Date and Time: 2016/09/15 02:30:01, Source: Security, Task Category: Credential Validation.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:30:45, Source: Windows Security, Task Category: Logon.
Event ID: 4723, Details: Attempt to change account password, Level: Warning, Date and Time: 2016/09/15 02:31:55, Source: Windows Security, Task Category: Account Management.
Event ID: 5038, Details: Integrity of system file failed, Level: Error, Date and Time: 2016/09/15 02:32:07, Source: System, Task Category: System Integrity.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:33:15, Source: Windows Security, Task Category: Logon.
Event ID: 4648, Details: Logon attempt with explicit credentials, Level: Warning, Date and Time: 2016/09/15 02:34:15, Source: Windows Security, Task Category: Logon.
Event ID: 4771, Details: Pre-authentication failed, Level: Warning, Date and Time: 2016/09/15 02:34:30, Source: Security, Task Category: Credential Validation.
Event ID: 1102, Details: Audit log cleared, Level: Critical, Date and Time: 2016/09/15 02:35:00, Source: Security, Task Category: Audit Logs.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:35:12, Source: Windows Security, Task Category: Logon.
Event ID: 4776, Details: Logon failure due to incorrect credentials, Level: Warning, Date and Time: 2016/09/15 02:36:15, Source: Windows Security, Task Category: Credential Validation.
Event ID: 5031, Details: Firewall service encountered an error, Level: Warning, Date and Time: 2016/09/15 02:36:30, Source: Security, Task Category: Firewall.
Event ID: 4625, Details: Logon attempt failed, Level: Warning, Date and Time: 2016/09/15 02:37:45, Source: Windows Security, Task Category: Logon.

---

### *Appendix A.2. GenDFIR Agent*

The DFIR RAG Agent task and prompt are as follows:

```

1 GenDFIR RAG Agent :
2 Role: DFIR Timeline Analysis AI Assistant.
3 Agent Prompt: You are a DFIR AI assistant, tasked with analysing artefacts,
correlating events, and producing a coherent timeline of the incident.
Base your answer on the provided context and do not include additional
information outside of the context given.

```

### *Appendix A.3. GenDFIR Query/Instruction*

The prompt/query/instruction provided via the DFIR expert is as follows:

```

1 DFIR Expert - Analyst Query: Conduct DFIR timeline analysis by examining
the artefact, correlating events, and reconstructing the timeline of
the cyber incident.

```

### *Appendix A.4. GenDFIR Output*

The output is given in [Table A2](#).

**Table A2.** GenDFIR output (powered by Llama-3.1).**\*\*Incident Analysis\*\***

I have analyzed the given event IDs to identify anomalous events, conducted a timeline analysis, correlated between them, reconstructed the events timeline, and provided precise mitigation solutions.

**\*\*Event Timeline Reconstructed\*\***

1. \*\*2016/09/15 02:16:37\*\* : Event ID 5038—Integrity of system file failed (Source: System)
2. \*\*2016/09/15 02:17:05\*\* : Event ID 4723—Attempt to change account password (Source: Windows Security)
3. \*\*2016/09/15 02:17:21\*\* : Event ID 4625—Logon attempt failed (Source: Windows Security)
4. \*\*2016/09/15 02:20:12\*\* : Event ID 4769—Kerberos service ticket request failed (Source: Security)
5. \*\*2016/09/15 02:22:30\*\* : Event ID 4771—Pre-authentication failed (Source: Security)
6. \*\*2016/09/15 02:24:12\*\* : Event ID 4724—Password reset attempt (Source: Windows Security)
7. \*\*2016/09/15 02:24:45\*\* : Event ID 5031—Firewall service encountered an error (Source: Security)
8. \*\*2016/09/15 02:27:15\*\* : Event ID 4776—Logon failure due to incorrect credentials (Source: Windows Security)
9. \*\*2016/09/15 02:30:01\*\* : Event ID 4771—Pre-authentication failed (Source: Security)
10. \*\*2016/09/15 02:31:55\*\* : Event ID 4723—Attempt to change account password (Source: Windows Security)
11. \*\*2016/09/15 02:32:07\*\* : Event ID 5038—Integrity of system file failed [duplicated] (Source: System)
12. \*\*2016/09/15 02:34:30\*\* : Event ID 4771—Pre-authentication failed [duplicated] (Source: Security)
13. \*\*2016/09/15 02:36:15\*\* : Event ID 4776—Logon failure due to incorrect credentials [duplicated] (Source: Windows Security)
14. \*\*2016/09/15 02:36:30\*\* : Event ID 5031—Firewall service encountered an error [duplicated] (Source: Security)
15. \*\*2016/09/15 02:37:45\*\* : Event ID 4625—Logon attempt failed [duplicated] (Source: Windows Security)

**\*\*Anomalous Events and Trends\*\***

1. Multiple attempts to change account password (Event IDs 4723) between 2:17 a.m. and 2:31 a.m.
2. Repeated pre-authentication failures (Event ID 4771) between 2:22 a.m. and 2:34 a.m.
3. Logon attempt failures (Event IDs 4625, 4627) multiple times between 2:17 a.m. and 2:37 p.m.
4. Firewall service encountered errors (Event ID 5031) at different intervals between 2:24 a.m. and 2:36 p.m.

**\*\*Root Cause Analysis\*\***

These events suggest a possible malicious action or an automated script that was attempted to gain unauthorized access. The repeated failures may indicate a faulty password reset attempt or other security measures that could not be bypassed.

**\*\*Mitigation Solutions\*\***

1. **\*\*Enhance Security Measures\*\***: Immediate updates and patches should be applied to counter vulnerabilities.
2. **\*\*Alert System Deployment\*\***: Activate systems capable of detecting such attempts in real time to prevent further occurrences.
3. **\*\*Audit Logging Review\*\***: Review all logging mechanisms to determine if more detailed records are needed.
4. **\*\*Training Session\*\***: Update security training to emphasize common attack vectors and prevention strategies.

**\*\*Recommendations\*\***

1. The security team needs to review the system's configuration for any vulnerabilities that could have been exploited.
2. Investigate these events further with a focus on identifying potential weaknesses the attacker might have used.
3. Conduct in-depth audits of systems, configurations, and account management policies to prevent similar incidents from recurring.

## References

1. Hamouda, D.; Ferrag, M.A.; Benhamida, N.; Seridi, H.; Ghanem, M.C. Revolutionizing intrusion detection in industrial IoT with distributed learning and deep generative techniques. *Internet Things* **2024**, *26*, 101149. <https://doi.org/10.1016/j.iot.2024.101149>.
2. Johansen, G. *Digital Forensics and Incident Response*; Packt: New York, NY, USA 2017.
3. Forensics, M. The 2024 State of Enterprise DFIR Industry Report: DFIR Teams Seek Automation, AI, and New Tools to Meet Rising Demands. 2024 Available online: <https://www.magnetforensics.com/blog/the-2024-state-of-enterprise-dfir-industry-report/> (accessed on 21 October 2024).
4. Cellebrite. Cellebrite's 2024 Industry Trends Survey Reveals Urgent Digital Challenges in Law Enforcement. 2023. Available online: <https://investors.cellebrite.com/news-releases/news-release-details/cellebrites-2024-industry-trends-survey-reveals-urgent-digital> (accessed on 8 August 2024).
5. Velociraptor. Velociraptor. 2024. Available online: <https://docs.velociraptor.app/> (accessed on 11 July 2024).
6. FTK Exterro. FTK Forensic Toolkit. 2024. Available online: <https://www.exterro.com/digital-forensics-software/forensic-toolkit> (accessed on 8 August 2024).
7. EnCase. EnCase Forensic Suite. 2024. Available online: <https://e-forensic.ca/products/encase-forensic-suite/> (accessed on 8 August 2024).
8. Dissect. Dissect. 2024. Available online: <https://docs.dissect.tools/> (accessed on 21 October 2024).
9. Timesketch. Timesketch. 2024. Available online: <https://timesketch.org/> (accessed on 21 October 2024).
10. Palso, L. log2timeline Palso. 2024. Available online: <https://plaso.readthedocs.io/> (accessed on 2 September 2024).
11. Splunk. Splunk. 2024. Available online: <https://www.splunk.com/> (accessed on 2 September 2024).

12. Farzaan, M.A.M.; Ghanem, M.C.; El-Hajjar, A.; Ratnayake, D.N. Ai-enabled system for efficient and effective cyber incident detection and response in cloud environments. *arXiv* **2024**, arXiv:2404.05602. <https://doi.org/doi.org/2404.05602>.
13. Dunsin, D.; Ghanem, M.C.; Ouazzane, K.; Vassilev, V. A comprehensive analysis of the role of artificial intelligence and machine learning in modern digital forensics and incident response. *Forensic Sci. Int. Digit. Investig.* **2024**, *48*, 301675. <https://doi.org/10.1016/j.fsidi.2023.301675>.
14. Dunsin, D.; Ghanem, M.C.; Ouazzane, K. The Use of Artificial Intelligence in Digital Forensics and Incident Response (DFIR) in a Constrained Environment. *Int. J. Inf. Commun. Technol.* **2022**, *16*, 280–285.
15. Dunsin, D.; Mohamed C. Ghanem, K.O.; Vassilev, V. Reinforcement Learning for an Efficient and Effective Malware Investigation during Cyber Incident Response. *High-Confid. Comput.* **2025**. <https://doi.org/10.1016/j.hcc.2025.100299>.
16. OpenAI GPT. GPT-4. 2024. Available online: <https://openai.com/index/gpt-4/> (accessed on 8 August 2024).
17. Meta Llama. Llama. 2024. Available online: <https://llama.meta.com/> (accessed on 8 August 2024).
18. Anthropic Claude. Claude. 2024. Available online: <https://www.anthropic.com/claude> (accessed on 8 August 2024).
19. Perković, G.; Drobnjak, A.; Botički, I. Hallucinations in LLMs: Understanding and Addressing Challenges. In Proceedings of the 2024 47th MIPRO ICT and Electronics Convention (MIPRO), Opatija, Croatia, 20–24 May 2024; pp. 2084–2088. <https://doi.org/10.1109/MIPRO60963.2024.10569238>.
20. Cuskey, C.; Woods, R.; Flaherty, M. The Limitations of Large Language Models for Understanding Human Language and Cognition. *Open Mind* **2024**, *8*, 1058–1083. [https://doi.org/10.1162/opmi\\_a\\_00160](https://doi.org/10.1162/opmi_a_00160).
21. Mansurova, A.; Mansurova, A.; Nugumanova, A. QA-RAG: Exploring LLM Reliance on External Knowledge. *Big Data Cogn. Comput.* **2024**, *8*, 115. <https://doi.org/10.3390/bdcc8090115>.
22. Linzbach, S.; Tressel, T.; Kallmeyer, L.; Dietze, S.; Jabeen, H. Decoding Prompt Syntax: Analysing its Impact on Knowledge Retrieval in Large Language Models. In Proceedings of the Companion Proceedings of the ACM Web Conference 2023, Austin, TX, USA, 30 April–4 May 2023; ACM: New York, NY, USA, 2023; WWW '23 Companion; pp. 1145–1149. <https://doi.org/10.1145/3543873.3587655>.
23. Naveed, H.; Khan, A.U.; Qiu, S.; Saqib, M.; Anwar, S.; Usman, M.; Akhtar, N.; Barnes, N.; Mian, A. A Comprehensive Overview of Large Language Models. *arXiv* **2024**, arXiv: 2307.06435. <https://doi.org/10.48550/arXiv.2307.06435>.
24. Dimitriadis, A.; Ivezic, N.; Kulvatunyou, B.; Mavridis, I. D4I—Digital forensics framework for reviewing and investigating cyber attacks. *Array* **2020**, *5*, 100015. <https://doi.org/10.1016/j.array.2019.100015>.
25. National Institute of Standards and Technologies (NIST). Cybersecurity Incident. Available online: [https://csrc.nist.gov/glossary/term/cybersecurity\\_incident](https://csrc.nist.gov/glossary/term/cybersecurity_incident) (accessed on 8 August 2024).
26. VMware. Anomaly Detection. Available online: <https://www.vmware.com/topics/anomaly-detection> (accessed on 8 August 2024).
27. Zimmerman, E. Eric Zimmerman’s Tools. 2024. Available online: <https://ericzimmerman.github.io/> (accessed on 8 August 2024).
28. Chabot, Y.; Bertaux, A.; Nicolle, C.; Kechadi, M.T. A complete formalized knowledge representation model for advanced digital forensics timeline analysis. *Digit. Investig.* **2014**, *11*, S95–S105. <https://doi.org/10.1016/j.diin.2014.05.009>.
29. Bhandari, S.; Jusas, V. The Phases Based Approach for Regeneration of Timeline in Digital Forensics. In Proceedings of the 2020 International Conference on INnovations in Intelligent SysTems and Applications (INISTA), Novi Sad, Serbia, 24–26 August 2020; pp. 1–6. <https://doi.org/10.1109/INISTA49547.2020.9194649>.
30. Hargreaves, C.; Patterson, J. An automated timeline reconstruction approach for digital forensic investigations. *Digit. Investig.* **2012**, *9*, S69–S79. <https://doi.org/10.1016/j.diin.2012.05.006>.
31. Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; et al. A Survey on Evaluation of Large Language Models. *ACM Trans. Intell. Syst. Technol.* **2024**, *15*, 1–45. <https://doi.org/10.1145/3641289>.
32. Hugging Face. The Tokenizer Playground. Available online: <https://huggingface.co/spaces/Xenova/the-tokenizer-playground> (accessed on 8 August 2024).
33. Ning, L.; Liu, L.; Wu, J.; Wu, N.; Berlowitz, D.; Prakash, S.; Green, B.; O’Banion, S.; Xie, J. USER-LLM: Efficient LLM Contextualization with User Embedding. Technical report, Google DeepMind, Google. *arXiv* **2024**, arXiv:2402.13598. <https://doi.org/10.48550/arXiv.2402.13598>.
34. Dong, Z.; Xiao, Y.; Wei, P.; Lin, L. Decoder-Only LLMs are Better Controllers for Diffusion Models. In Proceedings of the 32nd ACM International Conference on Multimedia, Melbourne, VIC, Australia, 28 October–1 November 2024; ACM: New York, NY, USA, 2024; MM ’24; pp. 10957–10965. <https://doi.org/10.1145/3664647.3680725>.
35. Chen, Z.; Xu, L.; Zheng, H.; Chen, L.; Tolba, A.; Zhao, L.; Yu, K.; Feng, H. Evolution and Prospects of Foundation Models: From Large Language Models to Large Multimodal Models. *Comput. Mater. Contin.* **2024**, *80*, 1753–1808. <https://doi.org/10.32604/cmc.2024.052618>.



36. Wan, S.; Nikolaidis, C.; Song, D.; Molnar, D.; Crnkovich, J.; Grace, J.; Bhatt, M.; Chennabasappa, S.; Whitman, S.; Ding, S.; et al. CYBERSECEVAL 3: Advancing the Evaluation of Cybersecurity Risks and Capabilities in Large Language Models. Release Date: 23 July 2024. 2024. Available online: <https://github.com/meta-llama/PurpleLlama/tree/main/CybersecurityBenchmarks> (accessed on 8 August 2024).
37. López Espejel, J.; Ettifouri, E.H.; Yahaya Alassan, M.S.; Chouham, E.M.; Dahhane, W. GPT-3.5, GPT-4, or BARD? Evaluating LLMs reasoning ability in zero-shot setting and performance boosting through prompts. *Nat. Lang. Process. J.* **2023**, *5*, 100032. <https://doi.org/10.1016/j.nlp.2023.100032>.
38. Ferrag, M.A.; Alwahedi, F.; Battah, A.; Cherif, B.; Mechri, A.; Tihanyi, N. Generative AI and Large Language Models for Cyber Security: All Insights You Need. *arXiv* **2024**, arXiv:2405.12750. <https://doi.org/10.48550/arXiv.2405.12750>.
39. Otal, H.T.; Canbaz, M.A. LLM HoneyPot: Leveraging Large Language Models as Advanced Interactive HoneyPot Systems. In Proceedings of the 2024 IEEE Conference on Communications and Network Security (CNS), Taipei, Taiwan, 30 September–3 October 2024; pp. 1–6. <https://doi.org/10.1109/CNS62487.2024.10735607>.
40. Wang, T.; Zhao, Z.; Wu, K. Exploiting LLM Embeddings for Content-Based IoT Anomaly Detection. In Proceedings of the 2024 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (PACRIM), Victoria, BC, Canada, 21–24 August 2024; pp. 1–6. <https://doi.org/10.1109/PACRIM61180.2024.10690230>.
41. Fariha, A.; Gharavian, V.; Makrehchi, M.; Rahnamayan, S.; Alwidian, S.; Azim, A. Log Anomaly Detection by Leveraging LLM-Based Parsing and Embedding with Attention Mechanism. In Proceedings of the 2024 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), Kingston, ON, Canada, 6–9 August 2024; pp. 859–863. <https://doi.org/10.1109/CCECE59415.2024.10667308>.
42. Saha, D.; Tarek, S.; Yahyaei, K.; Saha, S.K.; Zhou, J.; Tehranipoor, M.; Farahmandi, F. LLM for SoC Security: A Paradigm Shift. *IEEE Access* **2024**, *12*, 155498–155521. <https://doi.org/10.1109/ACCESS.2024.3427369>.
43. Wickramasekara, A.; Breitingner, F.; Scanlon, M. SoK: Exploring the Potential of Large Language Models for Improving Digital Forensic Investigation Efficiency. *arXiv* **2024**, arXiv:2402.19366. <https://doi.org/10.48550/arXiv.2402.19366>.
44. Scanlon, M.; Breitingner, F.; Hargreaves, C.; Hilgert, J.N.; Sheppard, J. ChatGPT for digital forensic investigation: The good, the bad, and the unknown. *Forensic Sci. Int. Digit. Investig.* **2023**, *46*, 301609. <https://doi.org/10.1016/j.fsidi.2023.301609>.
45. K, S.E.; Sakshi.; Wadhwa, M. Enhancing Digital Investigation: Leveraging ChatGPT for Evidence Identification and Analysis in Digital Forensics. In Proceedings of the 2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), Greater Noida, India, 3–4 November 2023; pp. 733–738. <https://doi.org/10.1109/ICCCIS60361.2023.10425000>.
46. Security, C. Cado Security. 2024. Available online: <https://www.cadosecurity.com/> (accessed on 22 July 2024).
47. Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Wang, M.; Wang, H. Retrieval-Augmented Generation for Large Language Models: A Survey. *arXiv* **2024**, arXiv:2312.10997. <https://doi.org/10.48550/arXiv.2312.10997>.
48. Li, Z.; Li, C.; Zhang, M.; Mei, Q.; Bendersky, M. Retrieval Augmented Generation or Long-Context LLMs? A Comprehensive Study and Hybrid Approach. *arXiv* **2024**, arXiv:2407.16833. <https://doi.org/10.48550/arXiv.2407.16833>.
49. Xu, L.; Lu, L.; Liu, M.; Song, C.; Wu, L. Nanjing Yunjin intelligent question-answering system based on knowledge graphs and retrieval augmented generation technology. *Herit. Sci.* **2024**, *12*, 118. <https://doi.org/10.1186/s40494-024-01231-3>.
50. Chroma Research. Evaluating Chunking Strategies for Retrieval. 2024. Available online: <https://research.trychroma.com/evaluating-chunking> (accessed on 3 July 2024).
51. Lu, Y.; Aleta, A.; Du, C.; Shi, L.; Moreno, Y. LLMs and generative agent-based models for complex systems research. *Phys. Life Rev.* **2024**, *51*, 283–293. <https://doi.org/10.1016/j.pprev.2024.10.013>.
52. Tihanyi, N.; Ferrag, M.A.; Jain, R.; Bisztray, T.; Debbah, M. CyberMetric: A Benchmark Dataset based on Retrieval-Augmented Generation for Evaluating LLMs in Cybersecurity Knowledge. In Proceedings of the 2024 IEEE International Conference on Cyber Security and Resilience (CSR), London, UK, 2–4 September 2024; pp. 296–302. <https://doi.org/10.1109/CSR61664.2024.10679494>.
53. Lála, J.; O’Donoghue, O.; Shtedritski, A.; Cox, S.; Rodrigues, S.G.; White, A.D. PaperQA: Retrieval-Augmented Generative Agent for Scientific Research. *arXiv* **2023**, arXiv:2312.07559. <https://doi.org/10.48550/arXiv.2312.07559>.
54. Wang, X.; Wang, Z.; Gao, X.; Zhang, F.; Wu, Y.; Xu, Z.; Shi, T.; Wang, Z.; Li, S.; Qian, Q.; et al. Searching for Best Practices in Retrieval-Augmented Generation. *arXiv* **2024**, arXiv:2407.01219. <https://doi.org/10.48550/arXiv.2407.01219>.
55. Zhao, X.; Zhou, X.; Li, G. Chat2Data: An Interactive Data Analysis System with RAG, Vector Databases and LLMs. *Proc. VLDB Endow.* **2024**, *17*, 4481–4484. <https://doi.org/10.14778/3685800.3685905>.
56. Toukmaji, C.; Tee, A. Retrieval-Augmented Generation and LLM Agents for Biomimicry Design Solutions. *Proc. AAAI Symp. Ser.* **2024**, *3*, 273–278. <https://doi.org/10.1609/aaais.v3i1.31210>.
57. Li, R.; Xu, J.; Cao, Z.; Zheng, H.T.; Kim, H.G. Extending Context Window in Large Language Models with Segmented Base Adjustment for Rotary Position Embeddings. *Appl. Sci.* **2024**, *14*, 3076. <https://doi.org/10.3390/app14073076>.
58. Kucharavy, A., Fundamental Limitations of Generative LLMs. In *Large Language Models in Cybersecurity: “Threats, Exposure and Mitigation”*; Springer Nature: Cham, Switzerland, 2024; pp. 55–64. [https://doi.org/10.1007/978-3-031-54827-7\\_5](https://doi.org/10.1007/978-3-031-54827-7_5).

59. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 4–9 December 2017; NIPS'17; pp. 6000–6010.
60. Meta AI. Llama 3.2: Revolutionizing Edge AI and Vision with Open, Customizable Models. 2024. Available online: [https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/?utm\\_source=chatgpt.com](https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/?utm_source=chatgpt.com) (accessed on 3 July 2024).
61. The Guardian. Meta Releases Advanced AI Multimodal LLaMA Model for EU Facebook Owner. 2024. Available online: [https://www.theguardian.com/technology/article/2024/jul/18/meta-release-advanced-ai-multimodal-llama-model-eu-facebook-owner?utm\\_source=chatgpt.com](https://www.theguardian.com/technology/article/2024/jul/18/meta-release-advanced-ai-multimodal-llama-model-eu-facebook-owner?utm_source=chatgpt.com) (accessed on 3 July 2024).
62. Michelet, G.; Breitingner, F. ChatGPT, Llama, can you write my report? An experiment on assisted digital forensics reports written using (local) large language models. *Forensic Sci. Int. Digit. Investig.* **2024**, *48*, 301683. <https://doi.org/10.1016/j.fsidi.2023.301683>.
63. MixedBread AI. MXBAI-Embed-Large-V1 Documentation. 2024. Available online: <https://www.mixedbread.ai/docs/embeddings/mxbai-embed-large-v1> (accessed on 3 July 2024).
64. National Institute of Standards and Technology. RhinoHunt. 2005. Available online: <https://cfreds.nist.gov/all/NIST/RhinoHunt> (accessed on 3 July 2024).
65. Loumachi, F.Y.; Ghanem, M.C.; Ferrag, M.A. GenDFIR: Advancing Cyber Incident Timeline Analysis Through Retrieval-Augmented Generation and Large Language Models. 2024. Available online: <https://github.com/GenDFIR> (accessed on 3 July 2024).
66. UMAP. How UMAP Works. 2024. Available online: [https://umap-learn.readthedocs.io/en/latest/how\\_umap\\_works.html](https://umap-learn.readthedocs.io/en/latest/how_umap_works.html) (accessed on 3 July 2024).
67. Kent, K.; Chevalier, S.; Grance, T.; Dang, H. *Guide to Integrating Forensic Techniques into Incident Response*; Technical Report SP 800-86; National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2006. <https://doi.org/10.6028/NIST.SP.800-86>.
68. *ISO/IEC 27037:2012(en)*; Information Technology—Security Techniques—Guidelines for Identification, Collection, Acquisition and Preservation of Digital Evidence. ISO: Geneva, Switzerland, 2012.
69. *NIST Special Publication 800-53*; Security and Privacy Controls for Information Systems and Organizations. National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2020. <https://doi.org/10.6028/NIST.SP.800-53r5>.
70. *ISO/IEC 17025:2017*; General Requirements for the Competence of Testing and Calibration Laboratories. ISO: Geneva, Switzerland, 2017.
71. Jiménez, M.B.; Fernández, D.; Eduardo Rivadeneira, J.; Flores-Moyano, R. A Filtering Model for Evidence Gathering in an SDN-Oriented Digital Forensic and Incident Response Context. *IEEE Access* **2024**, *12*, 75792–75808. <https://doi.org/10.1109/ACCESS.2024.3405588>.
72. Ghanem, M.C.; Chen, T.M.; Ferrag, M.A.; Kettouche, M.E. ESASCF: Expertise Extraction, Generalization and Reply Framework for Optimized Automation of Network Security Compliance. *IEEE Access* **2023**, *11*, 129840–129853. <https://doi.org/10.1109/ACCESS.2023.3332834>.
73. Tynan, P. The integration and implications of artificial intelligence in forensic science. *Forensic Sci. Med. Pathol.* **2024**, *20*, 1103–1105. <https://doi.org/10.1007/s12024-023-00772-6>.
74. National Institute of Standards and Technology; National Institute of Standards and Technology (NIST). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; National Institute of Standards and Technology (NIST): Gaithersburg, MD, USA, 2023.

**Disclaimer/Publisher’s Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.