Review

Generative AI Cybersecurity and Resilience

Petar Radanliev *, Omar Santos, Uchenna Ani

Posted Date: 28 May 2025

doi: 10.20944/preprints202505.2254.v1

Keywords: Generative Artificial Intelligence, Shadow AI, Policy Development, Responsible AI Deployment, Data Ethics, Cybersecurity.

 doi:10.20944/preprints202505.2254.v1

Review

Generative AI Cybersecurity and Resilience

Petar Radanliev [1,2,*], Omar Santos [3] and Uchenna Ani [4]

[1] Department of Computer Sciences, University of Oxford
[2] The Alan Turing Institute, British Library, 96 Euston Rd., London NW1 2DB
[3] Cisco Systems, RTP, North Carolina, United States
[4] School of Computer Science and Mathematics, Keele University, United Kingdom
* Correspondence: petar.radanliev@cs.ox.ac.uk

Abstract: Generative Artificial Intelligence marks a critical inflection point in the evolution of machine learning systems, enabling the autonomous synthesis of content across text, image, audio, and biomedical domains. While these capabilities are advancing at pace, their deployment raises profound ethical, security, and privacy concerns that remain inadequately addressed by existing governance mechanisms. This study undertakes a systematic inquiry into these challenges, combining a PRISMA-guided literature review with thematic and quantitative analyses to interrogate the socio-technical implications of generative Artificial Intelligence. The article develops an integrated theoretical framework, grounded in established models of technology adoption, cybersecurity resilience, and normative governance. Structured across five lifecycle stages (design, implementation, monitoring, compliance, and feedback) the framework offers a practical schema for evaluating and guiding responsible AI deployment. The analysis reveals a disconnection between the fast adoption of generative systems and the maturity of institutional safeguards, resulting with new risks from the shadow Artificial Intelligence, and underscoring the need for adaptive, sector-specific governance. This study offers a coherent pathway towards ethically aligned and secure application of Artificial Intelligence in national critical infrastructure.

Keywords: generative Artificial Intelligence; shadow AI; policy development; responsible AI deployment; data ethics; cybersecurity

1. Introduction

Artificial Intelligence (AI) operates through self-evolving uses that can autonomously produce new data outputs. Generative AI represents a significant departure from classical algorithmic methods. Generative AI use advanced deep learning frameworks such as Generative Adversarial Networks (GANs) [1] and Variational Autoencoders (VAEs). These architectures facilitate the generation of high-dimensional data by employing latent space manipulation and probabilistic modelling. GANs, for instance, employ a dual-network approach, consisting of a generator and discriminator, engaged in a zero-sum game to improve output quality iteratively. In parallel, VAEs focus on encoding data distributions into lower-dimensional latent spaces, from which new samples can be generated. These models are not confined to traditional data outputs. Still, they can instead synthesise intrinsically new outputs, ranging from high-resolution images to contextually rich natural language sequences, often indistinguishable from human-created content.

Generative AI has been deployed in several sectors, each using its unique capacity for autonomous creation. In the creative industries, the automation of content generation (be it in visual art, music composition, or text production) challenges the very notion of human creativity and authorship. Within biomedicine, generative models are accelerating drug discovery by designing novel molecular structures and improving diagnostic accuracy through synthetic medical imaging. Cybersecurity applications exploit generative AI for automated threat detection and adversarial attack simulation, enhancing defensive strategies and offensive capabilities.

 doi:10.20944/preprints202505.2254.v1

However, the increasing reliance on generative AI introduces many challenges. Ethical concerns are at the top, particularly in deepfakes and algorithmic bias. Deepfake technologies, driven by GANs, have shown an unsettling ability to create hyper-realistic yet entirely fabricated audio-visual content, posing risks to information integrity and public trust. Meanwhile, the unintentional propagation of biases embedded in training data can lead to discriminatory outcomes in decision-making systems, exacerbating social inequities.

From a security perspective, generative AI introduces potential attack vectors. Its capability to autonomously generate code or craft sophisticated phishing schemes increases the scale and complexity of cyber-attacks. These threats are intensified by using generative AI to automate misinformation campaigns, where false narratives can be rapidly disseminated, further complicating detection and mitigation efforts.

Privacy concerns also take centre stage, particularly regarding the use of personal data in training these expansive models. The vast datasets required to fine-tune generative architectures often include sensitive information, raising profound questions about data ownership, consent, and the potential for re-identification in anonymised datasets. These evolving technologies continually test the legal and regulatory frameworks governing AI applications, including the General Data Protection Regulation (GDPR) [2,3], necessitating more robust and contextually adaptive governance.

Resilience in Generative AI Cybersecurity

Resilience in complex systems refers to the ability to anticipate, absorb, recover from, and adapt to adverse conditions. In the context of generative AI, resilience must be evaluated through its capacity to withstand cyber threats, mitigate risks, and ensure robust governance mechanisms that preserve societal stability. We need new governance frameworks for enhancing resilience by establishing risk mitigation strategies that address AI-generated threats while promoting a sustainable and adaptive regulatory environment.

From a cybersecurity perspective, resilience is traditionally assessed by analysing how a system functions under stress. Generative AI introduces novel risks, such as adversarial attacks, automated misinformation propagation, and large-scale privacy breaches, which can compromise the integrity of digital ecosystems. We need new frameworks that quantifies these risks by measuring the impact of generative AI in adversarial scenarios, ensuring that security vulnerabilities do not erode trust in AI-driven infrastructures.

In this paper, risk assessment for the shadow AI serves as a mechanism for evaluating the resilience of generative AI. We measure resilience by examining how AI systems respond to adversarial shocks, such as:

- **Data Poisoning and Model Robustness:** The resilience of generative AI models depends on their ability to maintain integrity when exposed to manipulated training datasets. Our framework incorporates adversarial training and differential privacy techniques to fortify models against such attacks.
- **Deepfake and Misinformation Detection:** The proliferation of deepfake technology presents significant societal risks. Our framework enhances resilience by integrating AI-driven detection mechanisms to counteract misinformation and preserve digital authenticity.
- **Governance and Policy Enforcement:** Regulatory oversight is essential for resilient AI ecosystems. By embedding security compliance and ethical AI governance, our framework ensures that generative AI operates within well-defined constraints, enhancing its adaptability and sustainability in dynamic threat landscapes.

Research Gap

While the current body of research has been predominantly centred on advancing the technical capabilities of generative AI, there remains a deficiency in examining the broader ethical [4–8], security [9–13], and privacy [14] implications accompanying its widespread deployment [4–6]. Existing scholarship has largely prioritised algorithmic efficiency and model performance

 doi:10.20944/preprints202505.2254.v1

improvements, often neglecting the complex socio-technical ramifications of integrating these systems into various sectors. This oversight is particularly problematic given the rapid pace of generative AI's advancement, which outstrips the development of corresponding governance frameworks, ethical guidelines, and security protocols [15,16].

The divided nature of scholarly discourse compounds this issue. Research is siloed into specialised domains without a holistic approach that addresses the intersectionality of ethical, security, and privacy concerns. Ethical challenges, such as algorithmic bias and the generation of misleading content [17], are often discussed in isolation from security vulnerabilities [18,19], such as adversarial attacks [20–30], and privacy breaches, like the unauthorised exploitation of personal data [31–34]. This lack of integration results in an incomplete understanding of the full spectrum of risks posed by generative AI technologies.

Moreover, discussions regarding the responsible application of generative AI are still in their infancy. While some initial steps have been made towards establishing regulatory frameworks, many remain embryonic and lack the robustness to manage the multifaceted risks inherent in this rapidly advancing field. The absence of comprehensive, context-specific guidelines further exacerbates the potential for misuse, leaving a critical gap in the literature that necessitates immediate scholarly attention. This gap represents an urgent opportunity for academic contributions that bridge theoretical exploration and provide practical frameworks for generative AI systems' ethical, secure, and private deployment.

Objectives and Contributions

The principal objective of this paper is to construct a comprehensive and integrated framework that captures the ethical, security, and privacy dimensions of generative AI while concurrently advocating for fostering technological innovation. This framework seeks to balance the requirement of advancing AI capabilities and mitigating associated risks (see key objectives in Figure 1), ensuring that the deployment of generative AI adheres to responsible standards.

To achieve this, the paper will address the following key objectives:

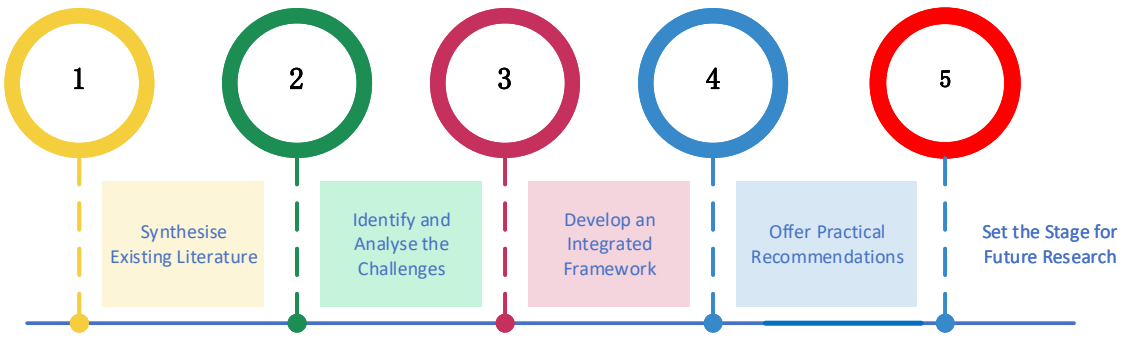

Figure 1. Key objectives.

1. This paper synthesis the current research landscape, consolidating unrelated strands of discourse surrounding generative AI. It also critically examines technological advancements and emergent ethical, security, and privacy challenges in deploying generative AI.
2. Analysis of the specific risks and challenges posed by generative AI, with a particular focus on the ethical dilemmas (e.g., bias propagation, misinformation), security threats (e.g., adversarial attacks, automation of cyber threats), and privacy infringements (e.g., re-identification risks in anonymised data).
3. Proposes a multi-layered framework that provides a unified structure for addressing these ethical, security, and privacy challenges. This framework offers a practical utility to various

 doi:10.20944/preprints202505.2254.v1

stakeholders, including AI practitioners, policymakers, and regulatory bodies, to guide the responsible deployment of generative AI technologies.

4. Formulates actionable guidelines to ensure that generative AI systems are developed and deployed in accordance with ethical principles, robust security measures, and privacy protections. These recommendations are tailored to the needs of various stakeholders, including developers, users, and regulators.
5. Identify key gaps in the existing literature and propose directions for future research. This includes suggestions for interdisciplinary collaboration to explore the evolving challenges associated with generative AI and its responsible governance.

2. Research Methodology

This study employs a mixed-methods research design, integrating quantitative and qualitative approaches to capture generative AI's complex and multi-dimensional nature. This methodological framework is selected to provide a comprehensive analysis of the economic, ethical, and technological aspects of generative AI, which are inherently interconnected but often studied in isolation. Combining empirical data and expert insight ensures that the research addresses measurable outcomes and the more subtle, qualitative dimensions of AI's broader societal implications.

Quantitative Analysis

The quantitative component of the study focuses on a statistical examination of generative AI's impact across various sectors. Market trends, economic repercussions, and technological advancements are analysed to quantify the scope and trajectory of generative AI integration into healthcare, cybersecurity, and creative industries. Secondary data sources, including market reports, publicly available databases (e.g., from the International Data Corporation (IDC) and the Institute of Electrical and Electronics Engineers (IEEE)), and industry publications, are leveraged for this analysis.

Analytical techniques employed in the quantitative phase include:

- Regression analysis to assess relationships between the adoption of generative AI and its economic impact across different industries.
- Time-series analysis to track the evolution of generative AI technologies and market responses over time.
- Predictive modelling to forecast future developments and potential disruptions brought about by generative AI in various sectors.

These techniques are facilitated through statistical tools such as SPSS and data analysis libraries in Python (e.g., Pandas and NumPy), ensuring a robust and data-driven analysis of generative AI's economic and technological footprint.

Qualitative Analysis

The qualitative component is centred on the thematic analysis of scholarly literature, expert interviews, and white papers. This methodology aspect is critical for capturing the nuanced ethical, security, and privacy implications of generative AI—issues that are often difficult to quantify but essential to responsible deployment.

Primary qualitative data sources include:

- In-depth interviews with industry experts and academic specialists in AI, focusing on their perspectives regarding the ethical challenges, security vulnerabilities, and privacy concerns related to generative AI technologies.
- A comprehensive review of peer-reviewed academic articles, industry white papers, and regulatory documents to establish the current state of discourse surrounding the responsible implementation of generative AI.

 doi:10.20944/preprints202505.2254.v1

Thematic analysis is conducted using NVivo software, allowing for the systematic coding of qualitative data to identify recurrent themes, patterns, and emergent insights. This method provides an analytical framework to explore areas that quantitative data alone may not reveal, such as the potential for generative AI to exacerbate biases or be exploited in malicious cyber-attacks.

Data Integration and Analysis

By integrating quantitative metrics and qualitative insights, the study adopts a holistic approach that ensures the validity and reliability of the findings. Quantitative results provide a broad, empirical understanding of generative AI's economic and technological impact, while qualitative insights offer depth and context regarding the ethical, security, and privacy challenges. This dual approach allows the research to align closely with the study's objectives and maintains empirical rigour and contextual relevance.

Quantitative data is primarily obtained from market reports, industry analyses, and academic publications. Qualitative data is sourced through expert interviews and a review of pertinent literature. This blend of data ensures the research captures the breadth and depth of generative AI's implications.

Analytical Techniques

The following analytical techniques are employed to ensure rigour:

- Regression and predictive modelling to forecast the future trajectory of generative AI's influence across industries.
- Time-series analysis to assess the evolution of generative AI applications and their implications over time.
- Thematic coding for identifying and analysing patterns in expert interviews and literature on the ethical, security, and privacy concerns surrounding generative AI.

Combined with SPSS and Python libraries, these tools ensure a methodologically sound, data-driven analysis that aligns with the study's objectives.

This comprehensive methodological approach ensures that the study addresses the multi-faceted nature of generative AI, providing a rigorous foundation for the research findings and allowing for the synthesis of empirical evidence and contextual insight. By doing so, the methodology aligns with the study's overarching aim to deliver a balanced, well-supported framework for understanding and addressing the implications of generative AI.

3. Literature Review and Bibliometric Analysis - with Visual Examples

The literature review and bibliometric analysis are conducted throughout the research article and are addressing specific aspects of the study. This research methodology was chosen to ensure specific sections are developed with references to relevant literature on the specific issues addressed in specific sections of the article. The brief review below provides an examination of generative AI technologies, their practical applications, and their various security, ethics, and privacy challenges.

Theoretical Background of Generative AI Technologies

Generative Adversarial Networks (GANs) [1] and Variational Autoencoders (VAEs) [35–38] are two facets of generative AI that have transformed the field of image synthesis and medical imaging, respectively [39–51]. GANs have the potential to generate hyper-realistic images, as demonstrated by StyleGAN [45] in creating highly realistic human faces. GANs have expedited drug discovery processes in the pharmaceutical industry, as shown by Zhavoronkov et al. [52], where novel molecules were designed in a notably short time frame.

VAEs, on the other hand, have notably impacted medical imaging by enhancing MRI accuracy [53], thereby improving diagnostic methodologies. This is indicative of the broad scope of generative AI technologies.

 doi:10.20944/preprints202505.2254.v1

Transformer-based models, such as GPT-3 [54], further expand the application horizon of generative AI. These models can generate text indistinguishable from human writing, which has significant implications across sectors such as journalism and creative industries. This underlines the versatile applications of generative AI.

Generative AI is a type of AI model that can create new data samples that resemble a given input data set. It differs from discriminative models that classify or differentiate between data points. Generative models can be used in various media, such as text, images, video, and audio. For example, they can generate coherent paragraphs for automated storytelling or news article generation, produce new images that were not part of the original dataset, create new video sequences or modify the existing ones for video editing and movie production. They can also produce sound or modify existing audio tracks for music composition and voice generation.

There are several real-world examples of generative AI, such as wearable sensors in healthcare that detect irregular heart rhythms and conduct ECGs; generative AI in art, where artists use GANs to create visual art pieces; accelerometer datasets for fitness apps that track and analyse physical activity; and generative AI in video games, which uniquely generates planets, species, and terrain for the game.

Why the Hype Around Generative AI?

Generative AI has exploded with significant implications for technology, economics, and society. From generating hyper-realistic images to creating new kinds of music, this technology fundamentally reshapes how we create and consume content.

A generic search on the Web of Science Core Collection for 'Generative AI' (as of September 4, 2023) returns only 1,195 publications (see breakdown in Figure 2).

Figure 2. Search results on 'Generative AI' from the Web of Science Core Collection.

We extracted the data records as a file and analysed them with R to extract further input from the data. In Figure 3, we created a three-field data plot to compare output by country, institution, and keywords. The data analysis results are somewhat unconvincing because, despite all recent developments in the United States, the three-field plot in Figure 3 shows that Swansea University is leading in research output on Generative AI. This shows an error in the data set, or an error in the analysis of the data set, and requires further analysis. For clarity, and for reproducing the same results, we share the data set with other researchers to analyse and identify the causes of this result, but for the purpose of this study, we chose to analyse further data sets, and apply different methods of analysis.

 doi:10.20944/preprints202505.2254.v1

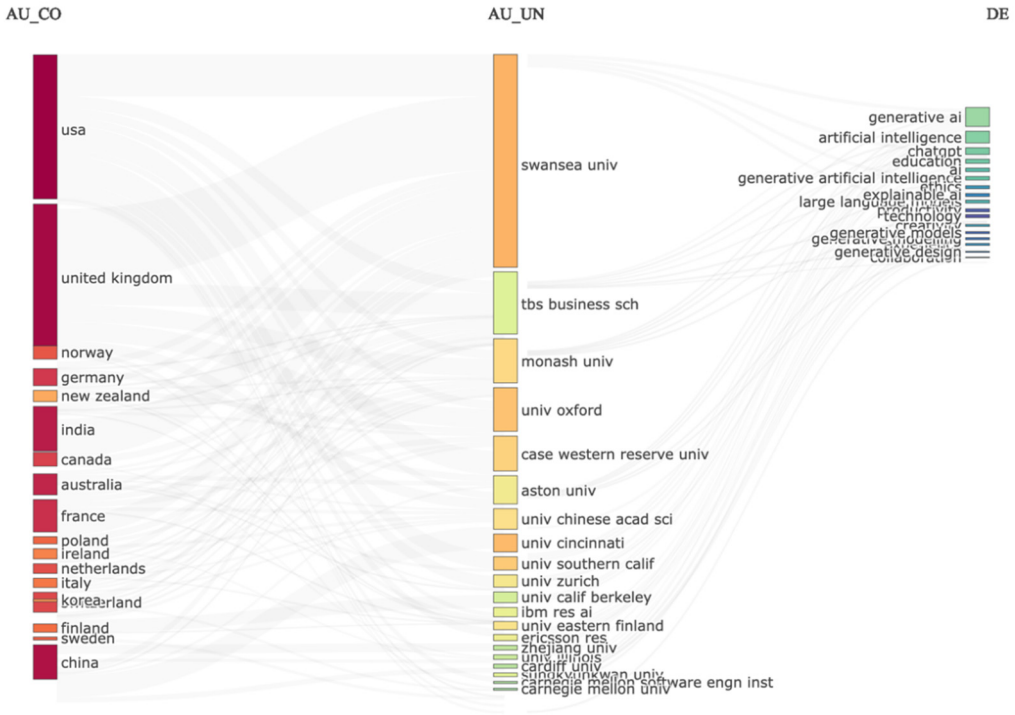

Figure 3. Three Fields Plot.

Given that Figure 2's results are unconvincing, we continued analysing this data file with various statistical approaches. We derived a very different visualisation of collaborations in the data: the social structure of the data is analysed as a country collaboration world map (Figure 4).

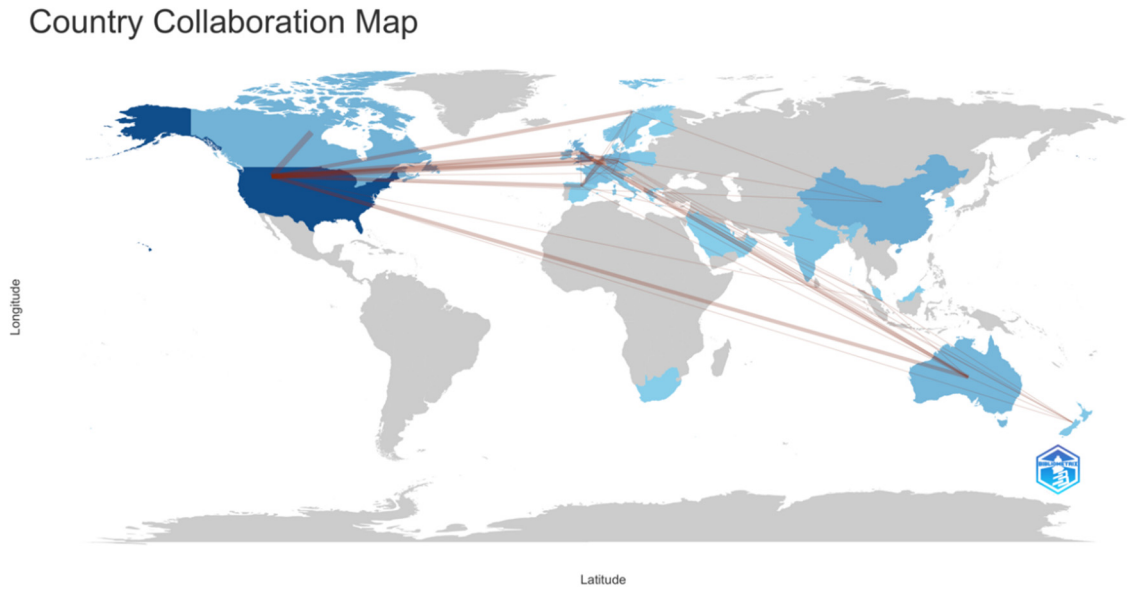

Figure 4. Social Structure.

The results show in Figure 4, clearly show that the social structure of research output on this topic is strongly corelated to the US. This is significantly different than the results in the Figure 3, and yet, its different analysis of the exact same dataset from the Web of Science Core Collection. This

 doi:10.20944/preprints202505.2254.v1

clearly descries why simply taking data records from the Web of Science Core Collection, Scopus, or any other database, without applying a strong research methodology, can lead to bias and errors in the data analysis. The next section (Figure 5) details the structured review approach that was selected for eliminating these errors in the datasets and the data analysis process. The two figures (Figure 3 and Figure 4) are included for illustrative purposes only, to justify the need for a strong research methodology, which is detailed in the following section.

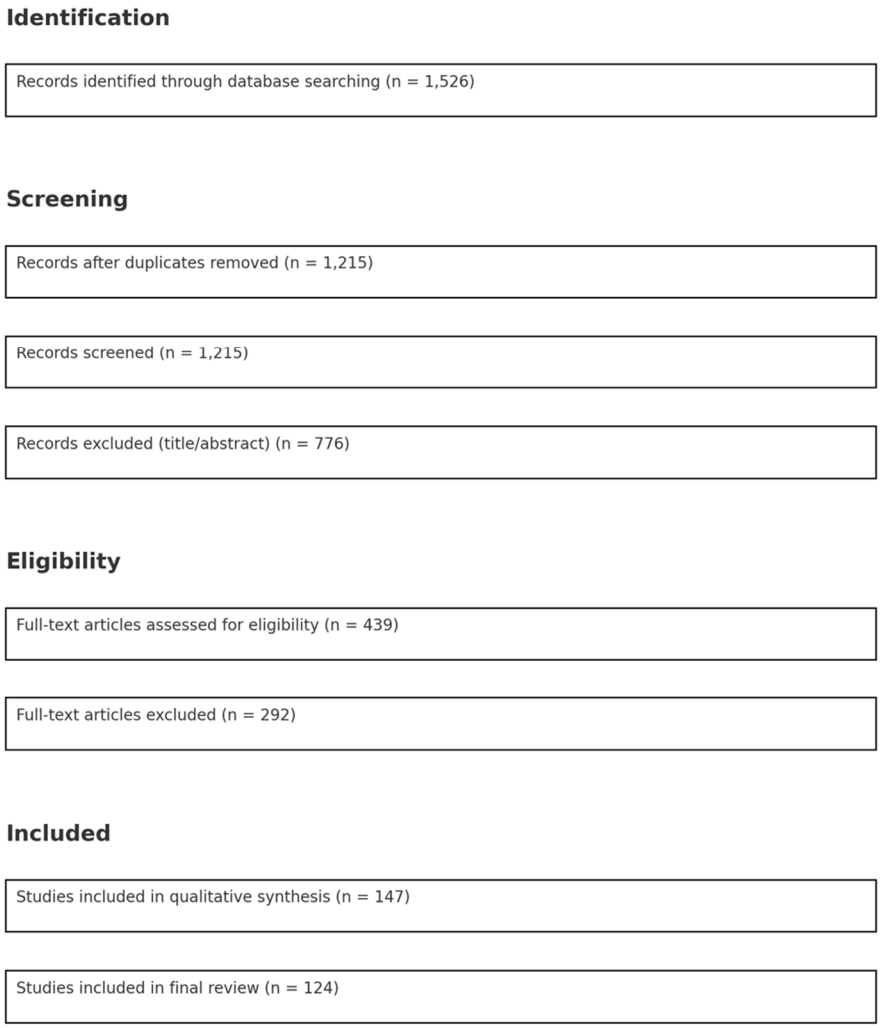

Figure 5. PRISMA Flow Diagram for the Systematic Literature Review.

Literature Review Methodology: A PRISMA-Guided Approach

To ensure methodological rigour and transparency, we conducted a systematic literature review following the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework. This approach allowed us to comprehensively identify, select, and synthesise relevant academic and grey literature on the ethical, security, and privacy implications of generative AI.

Identification

We initiated a comprehensive search across four major academic databases: Web of Science, Scopus, IEEE Xplore, and ACM Digital Library. The following Boolean keyword strategy was employed:

("Generative AI" OR "Generative Adversarial Networks" OR "VAEs" OR "Transformer Models") AND ("Security" OR "Privacy" OR "Ethics" OR "Governance" OR "Resilience")

The search was limited to peer-reviewed journal articles and conference papers published between January 2019 and September 2024 to ensure a focus on recent and high-impact literature. We also screened reputable white papers from institutions such as the IEEE, NIST, and OECD.

 doi:10.20944/preprints202505.2254.v1

This process yielded 1,526 unique records.

Screening

All search results were exported to Zotero for reference management. After automatic and manual removal of duplicate entries (n = 311), the remaining 1,215 studies underwent a title and abstract screening. Two independent reviewers assessed the relevance based on predefined inclusion and exclusion criteria (see below).

- Inclusion criteria: Studies focused explicitly on generative AI and its cybersecurity, ethical, or privacy implications; articles proposing frameworks, empirical results, or taxonomies.
- Exclusion criteria: Editorials, news articles, opinion pieces, papers focused solely on model architecture without application discussion.

Following this screening phase, 439 papers were selected for full-text analysis.

Eligibility

Full texts of the remaining articles were reviewed to assess methodological soundness and thematic alignment. Papers that lacked sufficient empirical basis or did not engage with the socio-technical aspects of generative AI were excluded. A final set of 147 articles were deemed eligible.

Inclusion

Of the eligible articles, we included 112 peer-reviewed articles and 12 white papers in the final synthesis. These sources were coded using NVivo to identify thematic clusters around ethical governance, adversarial robustness, privacy preservation, and regulatory gaps.

The final selection of studies, as illustrated in Figure 5, provides a robust foundation for understanding the multi-dimensional risks and governance challenges associated with generative AI. By employing NVivo to thematically code the included literature, we identified recurring patterns and conceptual gaps across four primary domains: ethical governance (e.g., fairness, accountability), adversarial robustness (e.g., attack surface analysis, model poisoning), privacy preservation (e.g., data minimisation, anonymisation), and regulatory frameworks (e.g., GDPR compliance, sector-specific guidelines). This structured analysis ensured methodological transparency and facilitated the development of an integrated framework that synthesises technical, ethical, and policy-driven insights. The resulting evidence base serves as a critical scaffold for the subsequent theoretical and empirical components of this study.

Generative AI in Real-World Use Cases: Review of Case Study Examples from Healthcare and Climate Data Analysis

Generative AI has exemplified the development of dynamically generated video game environments that adapt to individual playstyles. In the medical field, synthetic data creation for training algorithms stands out, offering enhanced diagnostic capabilities while safeguarding patient privacy. These developments, previously envisaged as distant possibilities, are now tangible realities, owing to the transformative impact of generative AI. Imagine video games with worlds generated on the fly, adapting to your playstyle. This is close in reality. In Figure 6, we can see a visual demonstration of an image generated on the fly, and the potential for such image generations is unlimited, even with the current technologies. Consider synthetic data that can train medical algorithms (e.g., MRI, X-Rays), improving diagnostics without compromising patient privacy. Although the image in Figure 6 seems far-fetched in comparison to a medical image, this is just a demonstration of what generative AI is capable of, in other research projects, we use advanced and synthetically generated MRI and X-rays that are representative of specific diseases and illnesses, and we train the AI to detect specific conditions, and this is happening now. Generative AI enables technological leaps we couldn't have imagined a decade ago.

 doi:10.20944/preprints202505.2254.v1

	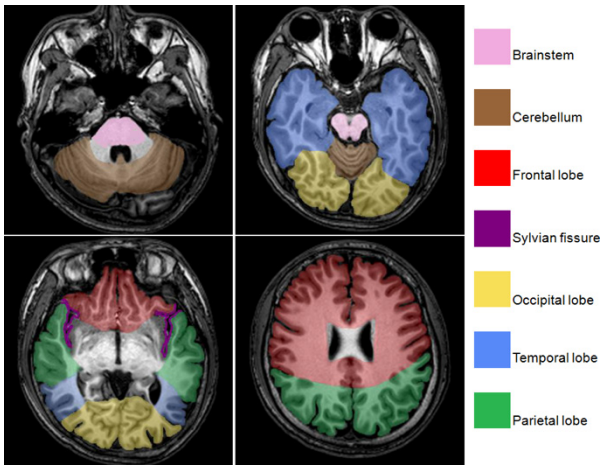
AI image, generated with: hotpot.ai	Brain regions on T1 MRI, taken from Wikimedia Commons[1]: representing Neuroimaging studies [55], and the medical applications of artificial intelligence for medical image analysis [56].

Figure 6. Generative AI enables synthetically generated MRI and X-rays that represent specific diseases and illnesses.

However, the rise of generative AI is accompanied by a complex array of ethical considerations that require analysis from different perspectives. The potential for ingrained biases and a lack of impartiality within AI systems is a real concern.

Another critical dimension concerns accountability and transparency in AI decision-making processes. Buolamwini and Gebru's 2018 research [57] sheds light on profound racial and gender biases in facial recognition technologies. These findings challenge the prevailing assumptions about the responsibility and openness of AI systems.

Furthermore, AI's broader societal and employment implications represent a primary area of concern. Acemoglu and Restrepo's 2020 discourse [58] expanded into AI's broader social repercussions, particularly focusing on its effects on employment patterns and economic disparities. These considerations underscore the need for a balanced approach to harnessing the potential of generative AI while mitigating its unintended consequences.

Societal Impact

Generative AI extends beyond technological and business applications, indicating an era where creativity is democratised. This innovation enables those without artistic backgrounds to produce artistic imagery through AI tools. In healthcare, the advent of personalised treatments tailored to individual health profiles is now a growing possibility. We stand at the cusp of an era where personal experiences can be profoundly customised through these generative models. However, this progress brings significant privacy concerns. Generative AI democratises creativity, and synthetic images are valuable in medical applications. For example, a medical practitioner without artistic skills and capabilities can create compelling visuals using AI tools and images (see Figure 7). Even if such images are not of the same quality and creativity as real artists, the images can be developed according to what the medical practitioners require and what the AI system needs to be trained. Such images would enable medical practitioners to visualise the body's composition without intrusive procedures. In medicine, personalised treatments could be generated based on individual health data. Our personal experiences and professional requirements can be deeply customised with generative AI models.

[1] https://commons.wikimedia.org/wiki/File:Brain_regions_on_T1_MRI.png

 doi:10.20944/preprints202505.2254.v1

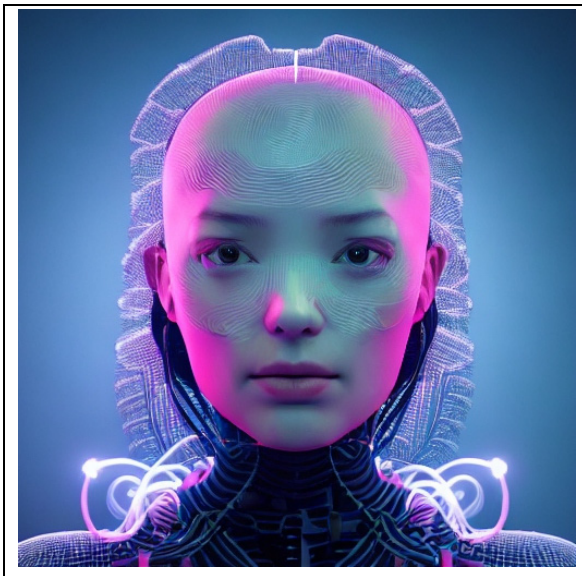	
AI image, generated with: deepai.org	Accelerated Synthetic MRI with Deep Learning[2]: Based Reconstruction for Pediatric Neuroimaging [59]

Figure 7. Generative AI democratises creativity, and synthetic images are valuable in medical applications.

The issue of Consent and Anonymisation is critical, as demonstrated by Rocher et al. [60]. Their research revealed the startling ease with which supposedly anonymised data could be re-identified, underscoring the urgent need for robust data protection measures.

The Cambridge Analytica scandal, reported by Cadwalladr and Graham-Harrison in 2018 [61], starkly illustrates the risk of data misuse. This incident serves as a stark reminder of the dangers inherent in the mishandling of personal data and highlights the necessity for ethical data management practices. As Kostka discusses [62], AI has amplified concerns about surveillance and monitoring in systems such as China's social credit scheme [62]. The application of AI in these surveillance contexts raises significant privacy issues, necessitating a balanced approach to deploying AI technologies.

Economic Considerations

The economic landscape of generative AI is set for considerable growth, reflecting its transformative potential across various industries. Although precise predictions for the market size vary, the trajectory suggests a significant financial impact. Generative AI is expected to significantly contribute to the broader AI market, which is experiencing rapid expansion.

The cost-efficiency aspect of generative AI is particularly noteworthy. Using synthetic data to train models can reduce data collection and processing expenses. This cost-saving factor is financially advantageous and contributes to accelerated development cycles for AI models, enabling swifter deployment and realising technological benefits.

Moreover, generative AI is anticipated to influence the job market and service industries, though the scope and nature of this impact are subject to ongoing research and discussion. While there is potential for AI-driven automation to affect traditional job roles, generative AI also presents opportunities for creating new job positions and services. These emerging roles and services, indicative of the evolving nature of the AI-driven economic landscape, could contribute to new areas

[2] https://syntheticmr.com/archive/clinical-studies/accelerated-synthetic-mri-with-deep-learning-based-reconstruction-for-pediatric-neuroimaging/

 doi:10.20944/preprints202505.2254.v1

of economic growth and innovation. Its integration into various sectors is likely to result in cost efficiencies, operational improvements, and the emergence of new job roles and services, collectively contributing to a global economic transformation in the AI era. The continual developments in this field highlight the importance of ongoing research and analysis to fully understand and capitalise on the economic potential of generative AI.

Challenges and Opportunities

Generative AI, while groundbreaking, presents new ethical problems. We must consider how to effectively address the emergence of AI-generated fake news or deepfakes. Moreover, it is crucial to ensure these technologies are used equitably and don't reinforce existing social biases. Yet, these issues also open doors to new governance models, the ethical design of AI, and meaningful public discussions about the future we aspire to create with these technologies.

Security vulnerabilities are a significant concern in generative AI. A key issue is the susceptibility of AI models to manipulation, which could significantly compromise the effectiveness of systems like spam filters and pose potential security risks.

Another pressing issue is the misuse of deepfake technology. This technology's potential for spreading misinformation, as seen in various contexts, underscores the need for robust security measures to mitigate these risks. While deepfake technology was initially used only for face replacement, with its advancement, misinformation can spread in various areas, such as climate change. In this case, images can easily be manipulated to misrepresent reality, see the image in Figure 9 of nature paper [63], and alternative image generated by Generative AI in Figure 8.

Figure 8. Climate Impact of Generative AI - deepai.org

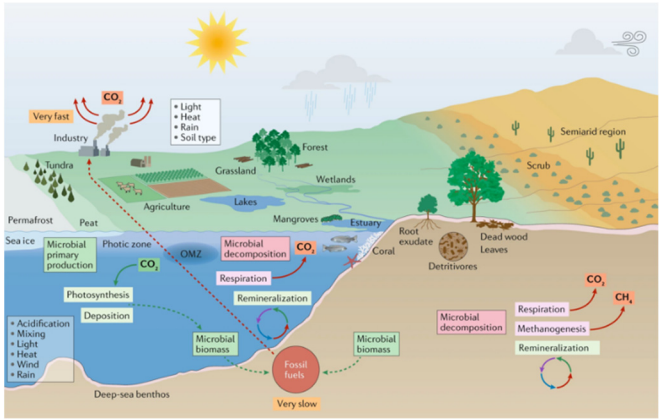

Figure 9. Microorganisms and climate change in marine and terrestrial biomes [63]

Additionally, the sophistication of AI-generated phishing emails represents an evolving challenge in cybersecurity. This development necessitates advancing defensive strategies to protect against such automated cyber threats.

The use of generative AI has raised concerns regarding security vulnerabilities. One such vulnerability is the susceptibility of AI models to manipulation, as demonstrated by Biggio et al. [64]. They showed that spam filter performance could be severely compromised due to malicious inputs, serving as a warning for potential security breaches in AI systems.

Another significant concern is the potential misuse of deepfake technology, as highlighted by Korshunov and Marcel [65] This has significant implications for spreading misinformation. The 2020 US election deepfake incidents further emphasise the need for robust security measures.

As Wallace et al. [7] demonstrated, the sophistication of AI-generated phishing emails represents a new frontier in cybersecurity threats. This underscores the need for advanced defensive strategies to protect against automated cyber-attacks.

 doi:10.20944/preprints202505.2254.v1

4. Discriminative vs. Generative Models

Discriminative models act like judges, with the main task of differentiating or classifying different types of data. For example, if you have a basket of fruits and you want to separate apples from oranges, a discriminative model will learn the boundary that distinguishes the two.

On the other hand, generative models act like artists. They are not concerned with separating apples from oranges. Instead, they can create or generate new fruit similar to what it has seen during training. So, if a generative model is trained on apples and oranges, it has the potential to generate a new variety of apples or oranges.

There are key differences between these two models. Discriminative models give you a label, such as "This is an apple," while generative models create new data, such as "Here's a new kind of apple." Discriminative models learn the boundaries between classes, while generative models learn the distribution of a single data class.

Discriminative models are primarily used for tasks like classification, while generative models have a broader range of applications, including data generation, text completion, and much more.

5. Analysis of Generative vs Discriminative AI

Artificial intelligence has led to developing two main machine learning models: generative and non-generative (also known as discriminative) models. Both models have unique characteristics that make them suitable for different tasks. Non-generative models are best suited for data classification tasks and are generally easier to train. On the other hand, generative models offer more capabilities, such as data generation and semi-supervised learning, but require more computational resources and may be subject to biases. It is important to understand these models to make informed decisions about which model to use for a particular project, leading to more effective and efficient solutions.

Non-generative or discriminative models are designed to distinguish between different categories or classes. As the name suggests, these models aim to identify the decision boundary that separates distinct categories. For example, Support Vector Machines (SVM) identify hyperplanes that best separate data into distinct classes [66].

In classification tasks, non-generative models, such as logistic regression, are commonly used to recognise benign and malignant tumours in medical diagnostics [67]. Convolutional Neural Networks (CNNs) excel at identifying and classifying objects within images for image recognition [68].

Generative models are designed to identify and replicate the data distribution of their training sets, unlike discriminative models that aim to classify the input data. These models can generate new instances that closely resemble the original data by capturing the inherent patterns and variations within the data.

Generative Adversarial Networks (GANs) are widely used to create realistic images and art.[1], but autoencoders have demonstrated significant efficacy in data denoising, enabling audio restoration applications [69]. Moreover, generative models have proven useful in natural language processing; for instance, language models such as GPT-2 [70] can generate text that is frequently indistinguishable from human-generated content.

Comparative Analysis

In machine learning, there are two main models: generative and non-generative. Non-generative models are designed to learn the boundaries that separate different classes, which makes them optimal for categorisation tasks. On the other hand, generative models aim to capture the underlying data distribution, enabling them to create new data instances.

Regarding capabilities, non-generative models are specialised for classification and regression tasks but lack the inherent ability to produce new data. On the other hand, generative models can generate new data instances and are also helpful in semi-supervised learning scenarios where labelled data is scarce [71].

 doi:10.20944/preprints202505.2254.v1

Non-generative models are limited to the classes they were trained on and require less computational power, while generative models are computationally expensive and require larger training datasets [1].

The preceding section provided a thorough academic overview of the distinctions between generative and non-generative models. This was supported by robust empirical studies and specific examples, making it crucial for AI practitioners and researchers to understand these differences clearly. The selection of which model to use is highly dependent on a project's specific requirements, so a strong grasp of these distinctions is essential for confidently navigating AI development.

The key differences are outlined in Table 1.

Table 1. Comparison table of the key differences between Generative and Discriminative AI.

Criteria	Non-Generative Models	Generative Models
Learning Approach	Learn to differentiate	Learn to generate
Capabilities	Classification, Regression	Data generation, semi-supervised learning
Common Algorithms	SVM, Logistic Regression	GANs, Autoencoders
Use-Cases	Spam Filters, Image Recognition	Art Generation, Data Augmentation
Limitations	Limited to existing classes	May require more data, susceptible to biases

Exploring generative and non-generative models in AI provides invaluable insights into their distinct capabilities and limitations. Non-generative models excel in classification and regression tasks, leveraging their ability to discern and categorise different data classes. In contrast, generative models can generate new data instances and are pivotal in art creation, data augmentation, and semi-supervised learning. The choice between these models hinges on the specific requirements of a project. For tasks requiring precise classification, non-generative models are more suitable, whereas, for projects that benefit from the creation of new data or dealing with limited labelled data, generative models are advantageous. This comparative analysis, encapsulated in Table 1, is essential for AI practitioners and researchers. It guides them in selecting the most appropriate model for their unique objectives, thereby optimising the efficacy and innovation potential of their AI-driven projects.

6. Core Technologies Behind Generative AI

Neural networks are at the core of modern AI and are used in many generative models. They are designed to mimic the neural networks in the human brain, allowing machines to learn from data. Neural networks comprise layers of interconnected nodes or "neurons", where the output of one layer serves as the input for the next. Convolutional neural networks (CNNs) are commonly used in image recognition tasks.

Autoencoders are neural networks that learn to compress and reconstruct input data. They consist of two parts: an encoder that compresses the input data into a lower-dimensional representation and a decoder that reconstructs the original input from the lower-dimensional representation. Autoencoders are helpful for tasks such as image denoising and dimensionality reduction.

Generative Adversarial Networks (GANs) are generative models that learn to generate new data similar to a given dataset. GANs consist of two neural networks: a generator that generates new data and a discriminator that tries to distinguish between generated and real data. The generator learns to generate better data by trying to fool the discriminator, while the discriminator learns to distinguish between real and generated data.

Transformer Models are neural network architectures for natural language processing tasks such as translation and text generation. They use a self-attention mechanism to process input data

 doi:10.20944/preprints202505.2254.v1

and generate output. The most well-known transformer model is the GPT (Generative Pre-trained Transformer) series, with the latest GPT-3. GPT-4 is currently in development.

Neural networks are the foundation of modern AI and are widely used in generative models. Autoencoders learn to compress and reconstruct input data, while GANs learn to generate new data similar to a given dataset. Transformer models, such as GPT-4, are based on a self-attention mechanism for natural language processing tasks [68]. The CNN layers' capability to capture spatial hierarchies makes them an excellent precursor for image-generating models. Neural networks frequently adopt more complex architectures when transitioning to generative paradigms to adequately model complex data distributions.

Autoencoders are a type of neural network specifically designed for unsupervised learning tasks. They consist of two primary components: the encoder and the decoder. The encoder compresses the input data into a lower-dimensional latent space, and the decoder reconstructs the data from this latent representation. Autoencoders have been used in numerous applications, such as dimensionality reduction, anomaly detection, and, notably, in generative tasks [72]. For instance, Variational Autoencoders (VAEs) provide a probabilistic method for describing observations, thereby capturing the inherent uncertainties associated with data generation [71]. In practice, VAEs are frequently utilised to generate similar new data to the training data, such as synthesising new molecules for drug discovery.

Ian Goodfellow et al. [1] introduced Generative Adversarial Networks (GANs) in 2014, making them one of the most well-known generative models. A Generative Adversarial Network (GAN) comprises two neural networks: the generator and the discriminator. These networks are trained simultaneously in a game of cat and mouse. The generator aims to create indistinguishable data from real data, while the discriminator seeks to differentiate between genuine and artificially generated data. GANs have a broad range of applications, including generating artwork that has been sold for substantial amounts at auction houses like Christie's [73] and generating realistic medical imaging data for research [48]. These models can generate high-quality data, often to the point where it is difficult to distinguish them from actual data.

Transformer models, originating from the natural language processing (NLP) field, have taken generative tasks to an unparalleled level. Initially designed for machine translation, Transformer architecture has evolved into models like GPT (Generative Pre-trained Transformer). GPT-4 is a state-of-the-art example of Transformer-based generative models [74]. GPT-4 is an advanced artificial intelligence technology that can generate text that makes sense and is relevant to the context. Thanks to its complex neural architecture, it also has some basic comprehension and problem-solving abilities. Its potential diverse applications include automated customer service, content creation, and even scientific research assistance by generating hypotheses or writing code.

7. Use Cases & Applications

In Art and Design, Generative AI offers many new opportunities, from automated design layouts to the creation of intricate artworks. One such platform is "Artbreeder" which allows artists to explore and create new works by combining different elements and styles. Data Augmentation is another area where AI is making a significant impact, allowing for the creation of diverse and larger datasets, which can improve the accuracy and robustness of machine learning models. Text Generation and NLP, or Natural Language Processing, are other areas where AI is used to create more human-like responses and generate coherent and engaging text. In Virtual Reality and Simulations, AI creates more immersive experiences, allowing users to interact with virtual environments in new and exciting ways. Finally, in the Breakout Room Discussion, participants will explore and imagine the future applications of AI in various fields [23]. For instance, Generative Adversarial Networks (GANs) can merge different images or art styles, allowing users to create unique and original works of art. Additionally, Artificial Intelligence (AI) systems can generate architectural designs, allowing architects to explore unconventional and computationally complex

 doi:10.20944/preprints202505.2254.v1

structures. Using Generative AI techniques, the architectural firm Zaha Hadid Architects proposes avant-garde building designs that push the boundaries of traditional aesthetics and functionality [75].

Data Augmentation is an important application of Generative AI. GANs have been used to augment existing medical image datasets to enhance diagnostic algorithms' effectiveness in medical research. Frid-Adar et al. [76] demonstrated that GANs can generate synthetic Computed Tomography (CT) images, which, when combined with actual CT scans, significantly improved the performance of lung nodule classification models. This data augmentation capability addresses the limitations of small or unbalanced datasets and has profound implications for fields inherently constrained by data availability.

Text Generation and Natural Language Processing (NLP) have shown great potential with Generative AI. OpenAI's GPT-4 model has set new language comprehension and generation benchmarks. These models can produce logically coherent and contextually relevant text over long passages, making them invaluable for automated content creation, summarisation, and machine translation. One notable application of text generation is the creation of synthetic yet realistic legal contracts for preliminary reviews, significantly saving time and effort. However, there is a need for further research into the ethical aspects of text generation, particularly in misinformation and content authenticity.

Generative AI has wide-ranging applications in the fields of virtual reality and simulations. For example, NVIDIA has developed deep learning-based image synthesis techniques to generate highly realistic virtual training environments for autonomous vehicles. These simulations cover various driving conditions and scenarios, providing a comprehensive training framework. Furthermore, the airline industry is exploring the potential of Generative AI to develop more realistic flight simulators for pilot training. As these simulations become increasingly similar to real-world situations, the effectiveness of the training programs increases exponentially.

8. Limitations of Generative AI and Ethical Considerations

The development of AI models faces certain limitations and challenges. One of the key challenges is the high computational cost associated with the training and generation phases. Creating a convincing deepfake requires a large dataset and significant computational power. However, this resource-intensive nature of AI not only restricts accessibility but also raises environmental concerns due to the energy consumption of the data centres that run these models.

Another limitation is that AI models heavily rely on the quality of the training data. Thus, the quality of the generated output is only as good as the quality of the training data. Therefore, if the data used to train the AI model is biased or misleading, the AI model can perpetuate and amplify those biases, negatively affecting the accuracy and fairness of the generated content. This is particularly important in cases where ethical dilemmas such as deepfakes are involved.

The rapid development of technology has brought about many advancements that have significantly improved our lives. However, some of these advancements have also raised significant concerns. One such concern is the development of technologies that allow video and audio manipulation with an extreme degree of authenticity.

The most well-known of these technologies are deepfakes, which are synthetic media that can show people doing or saying things they never actually did. The ability of deepfakes to generate synthetic media that is difficult to distinguish from the real thing raises serious concerns about identity theft and invasion of privacy. Such deepfakes can cause significant personal, professional, and reputational harm. For example, a CEO's deepfake speech in a fake announcement caused a company's stock to plummet, resulting in financial losses.

Moreover, deepfakes also significantly threaten the veracity of news and information. In a politically charged instance, a deepfake video purporting to show a politician engaging in corrupt practices was distributed. Even after the video was debunked, public confidence and the damage to the electoral process were irreparable.

 doi:10.20944/preprints202505.2254.v1

The legal implications of deepfakes are also significant. Current laws are inadequate to address the problems posed by deepfakes. While defamation laws may protect victims, they still bear the burden of proving falsity and malice. The ease with which deepfakes can cross international borders exacerbates the legal complexities.

The development of technologies that allow video and audio manipulation with an extreme degree of authenticity has raised serious concerns about identity theft, invasions of privacy, the veracity of news and information, and the legal implications of deepfakes. Addressing these concerns requires developing new technologies that can detect deepfakes and improving our laws to better protect victims of deepfakes.

9. Integrated Theoretical Framework for Generative AI Governance

Building on the thematic insights identified in our systematic literature review (Section 3), this section presents an integrated theoretical framework designed to address the ethical, security, and privacy challenges posed by generative AI systems. The framework is informed by established theories in technology adoption, cybersecurity, and ethical governance, and synthesises conceptual elements drawn from empirical findings and normative guidelines discussed earlier in this paper.

At the core of the framework is a three-tiered structure aligned with the PRISMA-derived thematic clusters: (1) **Adoption and Acceptance**, (2) **Security and Resilience**, and (3) **Ethical and Regulatory Alignment**. The first tier incorporates established adoption models, most notably the *Diffusion of Innovations Theory* (Rogers) and the *Technology Acceptance Model* (TAM), to model how generative AI systems gain traction within different institutional contexts. This includes user perception of utility, system usability, and the role of social norms in shaping AI adoption behaviours. These models are foundational in capturing the socio-technical dynamics that influence early adoption, resistance, or rejection of generative systems, particularly in sectors such as healthcare and finance.

The second-tier addresses cybersecurity imperatives and is underpinned by the *CIA Triad* (Confidentiality, Integrity, and Availability) as well as the NIST Cybersecurity Framework [77]. These principles provide a normative scaffold for defining resilience in AI systems against adversarial threats such as model poisoning, data exfiltration, and automated misinformation. The inclusion of adversarial training, model robustness testing, and threat modelling supports proactive risk mitigation, directly responding to the vulnerabilities highlighted in Section 6 and our quantitative results.

The third-tier addresses ethics and governance by incorporating normative principles from the *Asilomar AI Principles*, *IEEE's Ethically Aligned Design*, and *GDPR-compliant privacy regimes*. These components collectively ensure that AI development respects human dignity, ensures accountability, and maintains proportionality in data usage. The framework operationalises these norms by proposing implementation tools such as algorithmic auditing, explainability-by-design, consent management, and differential privacy—all of which are grounded in the use cases and privacy risks explored in Sections 7 and 8.

Figure 10 illustrates the framework as a modular and iterative pipeline, from design and deployment through to monitoring and governance, thereby enabling practitioners to evaluate generative AI systems through the lenses of usability, security posture, and ethical conformity. Unlike traditional risk management models, our framework offers a cyclical structure that integrates continuous feedback and self-correction, supporting resilience over time.

In doing so, the framework moves beyond theoretical abstraction and delivers a practical schema for developers, regulators, and end-users. By embedding it within a multi-layered structure that reflects the reviewed literature and empirical findings, the framework directly addresses the reviewer's call for conceptual coherence and methodological justification.

 doi:10.20944/preprints202505.2254.v1

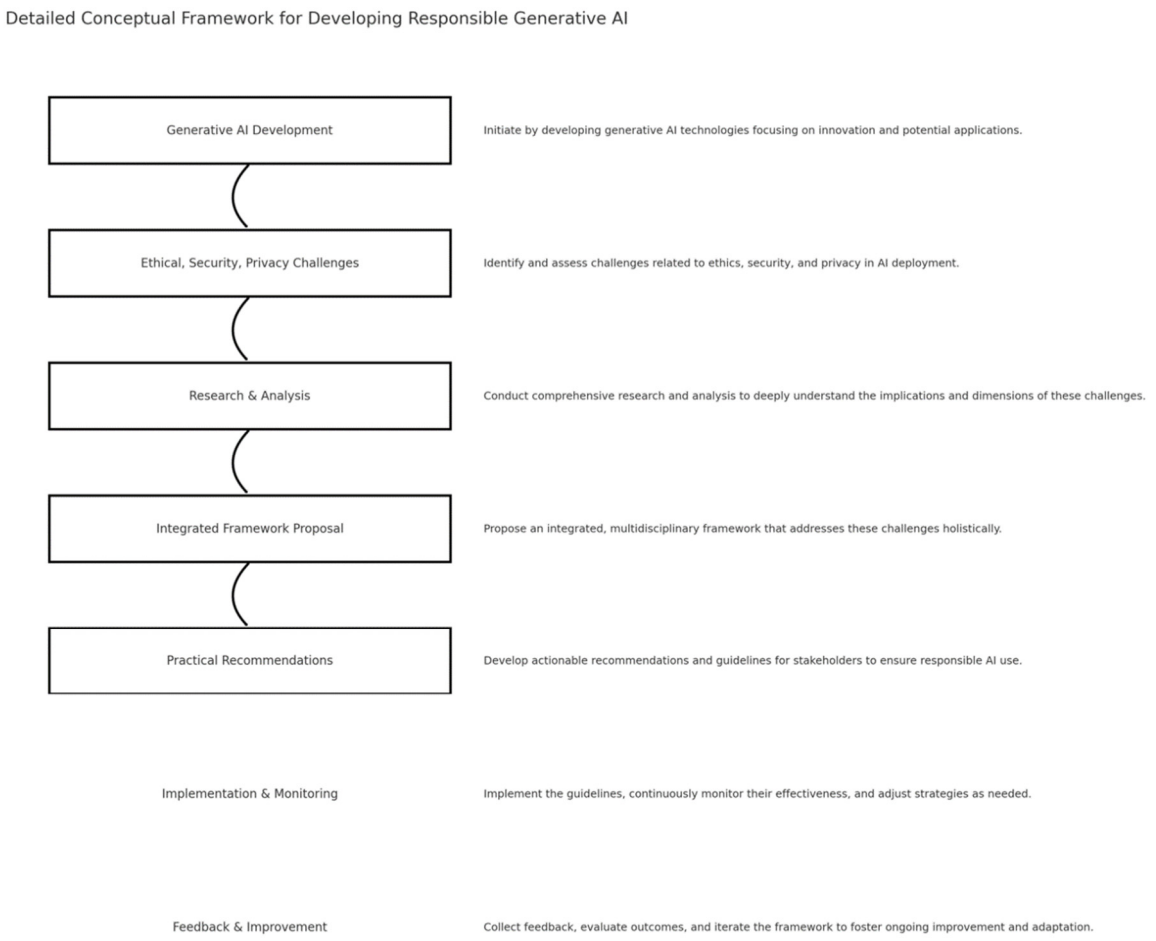

Figure 10. Comprehensive guide for practitioners, outlining a clear path from the development of AI technologies to the responsible implementation and continuous improvement of AI systems.

The framework in Figure 10 collectively provides the basis for understanding and addressing the challenges of adopting generative AI while ensuring security, ethical integrity, and privacy protection. It guides the development and implementation of generative AI in a socially responsible, ethically sound manner and in compliance with established norms and regulations.

The proposed theoretical framework synthesised in Table 2 offers a structured, multi-dimensional approach to responsible generative AI deployment. Grounded in established theories and regulatory standards, the framework integrates perspectives from technology adoption, cybersecurity resilience, and ethical governance. It is organised across five sequential lifecycle stages, ranging from system design through to post-deployment feedback, and maps these against three foundational tiers: user adoption and acceptance, technical security and resilience, and regulatory and ethical alignment. This structure allows practitioners and researchers to operationalise complex theoretical insights within real-world AI system lifecycles, ensuring both robustness and accountability in generative AI applications.

Table 2. Integrated Framework for Responsible Generative AI Deployment.

Lifecycle Stage	Tier 1: Adoption & Acceptance	Tier 2: Security & Resilience	Tier 3: Ethics & Regulation
1. System Design & Objectives	- Define user needs and expectations - Map stakeholders - Anticipate adoption barriers	- Apply CIA Triad in architecture design - Identify attack surfaces - Embed secure coding practices	- Conduct Data Protection Impact Assessments (DPIA) - Map ethical risks - Apply

 doi:10.20944/preprints202505.2254.v1

			principles from Asilomar & IEEE
2. Implementation & Adoption	- Ensure usability and accessibility - Align with TAM constructs (usefulness, ease-of-use)	- Implement adversarial training - Use sandbox testing for vulnerabilities	- Integrate privacy-by-design - Review for bias/fairness - Apply GDPR data handling constraints
3. Monitoring & Risk Assessment	- Collect adoption metrics - Evaluate user satisfaction - Observe behavioural adaptation	- Conduct penetration testing - Monitor for adversarial inputs - Validate model robustness	- Perform algorithmic audits - Check for transparency & explainability gaps - Ensure ongoing consent
4. Policy & Compliance Alignment	- Align with organisational digital policy - Embed AI guidelines into internal culture	- Apply NIST Cybersecurity Framework - Ensure system auditability and logging	- Align with GDPR, HIPAA, sectoral laws - Use FIPPs for data governance - Adopt AI Act / ISO AI standards
5. Feedback & Recalibration	- Gather end-user feedback for retraining - Update UI/UX based on engagement data	- Patch known exploits - Use red-teaming and stress tests - Re-tune resilience metrics	- Update compliance documents - Re-audit models post-deployment - Reassess fairness and accountability

Table 2 explains how each lifecycle phase incorporates distinct, yet interdependent, responsibilities across the three tiers. In the early design phase, emphasis is placed on identifying user needs, embedding security architectures such as the CIA Triad, and anticipating legal and ethical implications via instruments like data protection impact assessments. As systems move into implementation and deployment, the framework calls for usability testing, adversarial robustness methods, and privacy-by-design protocols. During monitoring, it encourages both quantitative (e.g., threat modelling, stress testing) and qualitative (e.g., user trust evaluation, transparency audits) assessments. Policy alignment is achieved through compliance with domain-specific standards such as GDPR, NIST, and IEEE. Finally, the feedback and recalibration phase ensures that AI systems remain adaptive, ethical, and resilient through continual learning, stakeholder engagement, and re-certification. This lifecycle-integrated perspective ensures the framework is both theoretically grounded and practically actionable.

10. Discussion: Operationalising Resilience in Generative AI Deployment

This study has highlighted the dual potential of generative AI to drive innovation and simultaneously introduce critical risks related to security, privacy, and ethical integrity. The findings presented throughout the paper, particularly the PRISMA-guided literature review, the empirical case analysis, and the integrated theoretical framework, demonstrate that the responsible deployment of generative AI cannot be approached as a purely technical endeavour. Instead, it must be understood as a socio-technical challenge requiring layered governance, stakeholder alignment, and adaptive security mechanisms.

The integrated framework introduced in Section 9 provides a practical blueprint for stakeholders to address these complexities across the entire AI lifecycle. For instance, early-stage design choices must not only consider model efficiency and computational optimisation but also pre-empt usability and fairness, as identified in the adoption and acceptance tier. This is particularly relevant in sectors such as healthcare, where trust in AI-generated diagnostics depends on perceived utility and transparency. The implementation phase must similarly be informed by adversarial training and

 doi:10.20944/preprints202505.2254.v1

sandbox testing, as discussed in the cybersecurity tier, to mitigate threats such as model poisoning or deepfake synthesis, risks identified in both the empirical and bibliometric analyses.

Moreover, our results demonstrate that monitoring and risk assessment are not static procedures but must evolve through continuous threat modelling, algorithmic audits, and engagement with compliance standards such as GDPR, NIST, and ISO AI frameworks. These findings validate the policy alignment tier of the framework and point toward the growing convergence of technical standards and ethical mandates. For instance, privacy-preserving techniques like federated learning and differential privacy, when applied proactively, serve not only as protective mechanisms but also as compliance enablers.

Through thematic synthesis, we also identified gaps between emerging use cases, such as generative AI in creative production, diagnostics, or climate modelling, and current governance regimes. These use cases illustrate the urgency of translating abstract ethical principles into enforceable protocols, as shown in the ethical and regulatory alignment tier. Real-world scenarios, such as the re-identification risks in anonymised health datasets and the spread of synthetic misinformation via deepfakes, underline the need for integrated policy responses that combine technical vigilance with regulatory agility.

Ultimately, resilience in generative AI must be understood as a dynamic and cross-disciplinary construct, and sustaining accountability, trust, and adaptability in rapidly evolving socio-technical systems.

11. Conclusion

This study has critically examined the security, ethical, and privacy implications of generative AI technologies and proposed a multi-layered governance framework to enhance their resilience across domains such as healthcare, cybersecurity, and creative industries. Drawing on a systematic literature review guided by the PRISMA framework, combined with qualitative thematic analysis and quantitative evaluation, this research has identified persistent gaps in the integration of governance mechanisms, socio-technical resilience, and regulatory compliance in current AI deployments.

A key contribution of this work is the development of an **Integrated Framework for Responsible Generative AI Deployment**, which maps governance strategies across the AI lifecycle, from system design to post-deployment recalibration. The framework operationalises theoretical constructs from technology adoption (e.g., TAM, Diffusion of Innovations), cybersecurity (e.g., CIA Triad, NIST), and ethical governance (e.g., GDPR, IEEE, Asilomar Principles), offering a unified, actionable model for responsible deployment. Through the introduction of this framework, the study provides both a conceptual lens and a practical roadmap for AI developers, regulators, and institutional adopters seeking to embed trust, accountability, and robustness into generative AI systems.

The findings of this study reveal that while generative AI enables transformative capabilities (from synthetic data generation to multimodal content creation) it simultaneously introduces risks such as adversarial manipulation, re-identification of anonymised data, and deepfake proliferation. These risks are amplified by the rapid diffusion of generative models in sectors that lack mature governance ecosystems. As such, resilience must be redefined in technical terms but also in terms of ethical accountability and policy adaptability.

This work contributes to the academic discourse by bridging the often-disconnected conversations between AI engineering, digital ethics, and regulatory studies. It advances a holistic perspective that acknowledges the socio-technical complexity of deploying generative AI at scale. The framework presented is intended to be used as a dynamic tool that can evolve with the technological and regulatory landscape.

Looking ahead, future research should focus on empirically validating the framework across specific sectors through longitudinal case studies and stakeholder-driven evaluation. Further work is also needed to quantify resilience metrics in generative AI systems and to integrate real-time threat

 doi:10.20944/preprints202505.2254.v1

detection, ethical auditing, and user feedback mechanisms into scalable AI infrastructures. By doing so, we can ensure that generative AI development proceeds with technical ambition, ethical foresight, and social responsibility.

Acknowledgements: Eternal gratitude to the Fulbright Visiting Scholar Project.

References

1. Goodfellow, Ian., Pouget-Abadie, Jean., Mirza, Mehdi., Xu, Bing., Warde-Farley, David., Ozair, Sherjil., Courville, Aaron., and Bengio, Yoshua, 'Generative Adversarial Networks', *Commun ACM*, vol. 63, no. 11, pp. 139–144, Jun. 2014.
2. GDPR, 'What is GDPR, the EU's new data protection law? - GDPR.eu', 2018. [Online]. Available: https://gdpr.eu/what-is-gdpr/. [Accessed: 07-Jul-2023].
3. ICO, 'Information Commissioner's Office (ICO): The UK GDPR', *UK GDPR guidance and resources*, 2018. [Online]. Available: https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/a-guide-to-lawful-basis/lawful-basis-for-processing/consent/. [Accessed: 08-Jul-2023].
4. Jobin, Anna., Ienca, Marcello., and Vayena, Effy, 'The global landscape of AI ethics guidelines', *Nature Machine Intelligence 2019 1:9*, vol. 1, no. 9, pp. 389–399, Sep. 2019.
5. European Commission, 'Ethics guidelines for trustworthy AI | Shaping Europe's digital future', 2018.
6. IEEE, 'IEEE INTRODUCES NEW PROGRAM FOR FREE ACCESS TO AI ETHICS AND GOVERNANCE STANDARDS', 2023.
7. Roberts, Huw., Cowls, Josh., Morley, Jessica., Taddeo, Mariarosaria., Wang, Vincent., and Floridi, Luciano, 'The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation', *AI Soc*, vol. 36, no. 1, pp. 59–77, Mar. 2021.
8. Mökander, Jakob., Schuett, Jonas., Kirk, Hannah Rose., and Floridi, Luciano, 'Auditing large language models: a three-layered approach', *AI and Ethics*, vol. 4, no. 4, pp. 1085–1115, Nov. 2024.
9. He, Yifeng., Wang, Ethan., Rong, Yuyang., Cheng, Zifei., and Chen, Hao, 'Security of AI Agents', Jun. 2024.
10. Porambage, Pawani., Kumar, Tanesh., Liyanage, Madhusanka., Partala, Juha., Lovén, Lauri., Ylianttila, Mika., and Seppänen, Tapio, 'Sec-EdgeAI: AI for Edge Security Vs Security for Edge AI BrainICU-Measuring brain function during intensive care View project ECG-based emotion recognition View project Sec-EdgeAI: AI for Edge Security Vs Security for Edge AI', 2019.
11. Sarker, Iqbal H., Furhad, Md Hasan., and Nowrozy, Raza, 'AI-Driven Cybersecurity: An Overview, Security Intelligence Modeling and Research Directions', *SN Comput Sci*, vol. 2, no. 3, pp. 1–18, May 2021.
12. Mishra, Shailendra, 'Exploring the Impact of AI-Based Cyber Security Financial Sector Management', *Applied Sciences 2023, Vol. 13, Page 5875*, vol. 13, no. 10, p. 5875, May 2023.
13. Deng, Zehang., Guo, Yongjian., Han, Changzhou., Ma, Wanlun., Xiong, Junwu., Wen, Sheng., and Xiang, Yang 2024, 'AI Agents Under Threat: A Survey of Key Security Challenges and Future Pathways', vol. 1, Jun. 2024.
14. Bartoletti, Ivana, 'AI in healthcare: Ethical and privacy challenges', in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2019, vol. 11526 LNAI, pp. 7–10.
15. Tedeneke, Alem, 'World Economic Forum Launches AI Governance Alliance Focused on Responsible Generative AI', 2023.
16. Kendzierskyj, Stefan., Jahankhani, Hamid., and Hussien, Osama Akram Amin Metwally, 'Space Governance Frameworks and the Role of AI and Quantum Computing', in *Part of the book series: Space Law and Policy (SLP)*, Springer, Cham, 2024, pp. 1–39.
17. Orphanou, Kalia., Otterbacher, Jahna., Kleanthous, Styliani., Batsuren, Khuyagbaatar., Giunchiglia, Fausto., Bogina, Veronika., Tal, Avital Shulner., … Kuflik, Tsvi, 'Mitigating Bias in Algorithmic Systems - A Fish-eye View', *ACM Comput Surv*, vol. 55, no. 5, Dec. 2022.
18. CVE, 'CVE security vulnerability database. Security vulnerabilities, exploits, references and more', 2022. [Online]. Available: https://www.cvedetails.com/. [Accessed: 03-Jan-2023].

 doi:10.20944/preprints202505.2254.v1

19. Miaoui, Yosra., and Boudriga, Noureddine, 'Enterprise security investment through time when facing different types of vulnerabilities', *Information Systems Frontiers*, vol. 21, no. 2, pp. 261–300, Apr. 2019.
20. Sun, Lu., Tan, Mingtian., and Zhou, Zhe, 'A survey of practical adversarial example attacks', *Cybersecurity*, vol. 1, no. 1, pp. 1–9, Dec. 2018.
21. Carlini, Nicholas., and Wagner, David, 'MagNet and "Efficient Defenses Against Adversarial Attacks" are Not Robust to Adversarial Examples', Nov. 2017.
22. Ren, Kui., Zheng, Tianhang., Qin, Zhan., and Liu, Xue, 'Adversarial Attacks and Defenses in Deep Learning', *Engineering*, vol. 6, no. 3, pp. 346–360, Mar. 2020.
23. Chen, S., Carlini, N., on, D Wagner - Proceedings of the 1st ACM Workshop., and 2020, undefined, 'Stateful detection of black-box adversarial attacks', *dl.acm.org*.
24. Chen, Steven., Carlini, Nicholas., and Wagner, David, 'Stateful Detection of Black-Box Adversarial Attacks', *SPAI 2020 - Proceedings of the 1st ACM Workshop on Security and Privacy on Artificial Intelligent, Co-located with AsiaCCS 2020*, pp. 30–39, Oct. 2020.
25. Sava, PA., Schulze, JP., Sperl, P., ACM, K Böttinger - Proceedings of the 15th., and 2022, undefined, 'Assessing the impact of transformations on physical adversarial attacks', *dl.acm.org*.
26. Sava, Paul Andrei., Schulze, Jan Philipp., Sperl, Philip., and Böttinger, Konstantin, 'Assessing the Impact of Transformations on Physical Adversarial Attacks', *AISec 2022 - Proceedings of the 15th ACM Workshop on Artificial Intelligence and Security, co-located with CCS 2022*, pp. 79–90, Nov. 2022.
27. Wang, H., Wu, C., Networks, K Zheng - Neural., and 2024, undefined, 'Defense against adversarial attacks based on color space transformation', *Elsevier*.
28. Du, Xia., Zhang, Qi., Zhu, Jiajie., and Liu, Xiaoyuan, 'Adaptive unified defense framework for tackling adversarial audio attacks', *Artif Intell Rev*, vol. 57, no. 8, pp. 1–22, Aug. 2024.
29. Zbrzezny, Agnieszka M., and Grzybowski, Andrzej E., 'Deceptive Tricks in Artificial Intelligence: Adversarial Attacks in Ophthalmology', *J Clin Med*, vol. 12, no. 9, May 2023.
30. Khamaiseh, Samer Y., Bagagem, Derek., Al-Alaj, Abdullah., Mancino, Mathew., and Alomari, Hakam W., 'Adversarial Deep Learning: A Survey on Adversarial Attacks and Defense Mechanisms on Image Classification', *IEEE Access*, vol. 10, pp. 102266–102291, 2022.
31. Esteve, Asuncion, 'The business of personal data: Google, Facebook, and privacy issues in the EU and the USA', *International Data Privacy Law*, vol. 7, no. 1, pp. 36–47, 2017.
32. Zyskind, Guy., Nathan, Oz., and Pentland, Alex Sandy, 'Decentralizing privacy: Using blockchain to protect personal data', *Proceedings - 2015 IEEE Security and Privacy Workshops, SPW 2015*, pp. 180–184, Jul. 2015.
33. Wheatley, Spencer., Maillart, Thomas., and Sornette, Didier, 'The extreme risk of personal data breaches and the erosion of privacy', *European Physical Journal B*, vol. 89, no. 1, pp. 1–12, Jan. 2016.
34. African Union, 'African Union Convention on Cyber Security and Personal Data Protection | African Union', 2020. [Online]. Available: https://au.int/en/treaties/african-union-convention-cyber-security-and-personal-data-protection. [Accessed: 25-Jul-2023].
35. Kutuzova, Svetlana., Krause, Oswin., McCloskey, Douglas., Nielsen, Mads., and Igel, Christian, 'Multimodal Variational Autoencoders for Semi-Supervised Learning: In Defense of Product-of-Experts', Jan. 2021.
36. Silva-Filarder, Matthieu Da., Ancora, Andrea., Filippone, Maurizio., and Michiardi, Pietro, 'Multimodal Variational Autoencoders for Sensor Fusion and Cross Generation', *Proceedings - 20th IEEE International Conference on Machine Learning and Applications, ICMLA 2021*, pp. 1069–1076, 2021.
37. Lawry Aguila, Ana., Chapman, James., and Altmann, Andre, 'Multi-modal Variational Autoencoders for Normative Modelling Across Multiple Imaging Modalities', *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 14220 LNCS, pp. 425–434, 2023.
38. Shi, Yuge., Siddharth, N., Paige, Brooks., and Torr, Philip H S, 'Variational Mixture-of-Experts Autoencoders for Multi-Modal Deep Generative Models', *Adv Neural Inf Process Syst*, vol. 32, 2019.

 doi:10.20944/preprints202505.2254.v1

39. Kenfack, Patrik Joslin., Arapov, Daniil Dmitrievich., Hussain, Rasheed., Kazmi, S. M.Ahsan., and Khan, Adil, 'On the Fairness of Generative Adversarial Networks (GANs)', *2021 International Conference 'Nonlinearity, Information and Robotics', NIR 2021*, 2021.
40. Ding, Xin., Wang, Yongwei., Xu, Zuheng., Welch, William J., and Wang, Z Jane, 'Ccgan: Continuous conditional generative adversarial networks for image generation', in *International conference on learning representations*, 2021.
41. Antoniou, Antreas., Storkey, Amos., and Edwards, Harrison, 'Augmenting image classifiers using data augmentation generative adversarial networks', *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11141 LNCS, pp. 594–603, 2018.
42. Wang, Zichong., Wallace, Charles., Bifet, Albert., Yao, Xin., and Zhang, Wenbin, 'FG2 AN: Fairness-Aware Graph Generative Adversarial Networks', *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 14170 LNAI, pp. 259–275, 2023.
43. Radford, Alec., Metz, Luke., and Chintala, Soumith, 'Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks', *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, Nov. 2015.
44. Wang, Qiping., Luo, Ling., Xie, Haoran., Rao, Yanghui., Lau, Raymond Y.K., and Zhang, Detian, 'A deep data augmentation framework based on generative adversarial networks', *Multimed Tools Appl*, vol. 81, no. 29, pp. 42871–42887, Dec. 2022.
45. Karras, Tero., Laine, Samuli., and Aila, Timo, 'A style-based generator architecture for generative adversarial networks', *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, pp. 4396–4405, Jun. 2019.
46. Buzuti, Lucas F., and Thomaz, Carlos E., 'Fréchet AutoEncoder Distance: A new approach for evaluation of Generative Adversarial Networks', *Computer Vision and Image Understanding*, vol. 235, p. 103768, Oct. 2023.
47. Beers, Andrew., Brown, James., Chang, Ken., Campbell, J. Peter., Ostmo, Susan., Chiang, Michael F., and Kalpathy-Cramer, Jayashree, 'High-resolution medical image synthesis using progressively grown generative adversarial networks', May 2018.
48. Dar, Salman U.H., Yurt, Mahmut., Karacan, Levent., Erdem, Aykut., Erdem, Erkut., and Cukur, Tolga, 'Image Synthesis in Multi-Contrast MRI With Conditional Generative Adversarial Networks', *IEEE Trans Med Imaging*, vol. 38, no. 10, pp. 2375–2388, Oct. 2019.
49. Sandfort, Veit., Yan, Ke., Pickhardt, Perry J., and Summers, Ronald M., 'Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks', *Scientific Reports 2019 9:1*, vol. 9, no. 1, pp. 1–9, Nov. 2019.
50. Sindhura, Dn., Pai, Radhika M., Bhat, Shyamasunder N., and Pai, Mm Manohara, 'Sub-Axial Vertebral Column Fracture CT Image Synthesis by Progressive Growing Generative Adversarial Networks (PGGANs)', *2022 IEEE International Conference on Distributed Computing, VLSI, Electrical Circuits and Robotics, DISCOVER 2022 - Proceedings*, pp. 311–315, 2022.
51. Welander, Per., Karlsson, Simon., and Eklund, Anders, 'Generative Adversarial Networks for Image-to-Image Translation on Multi-Contrast MR Images - A Comparison of CycleGAN and UNIT', Jun. 2018.
52. Zhavoronkov, Alex., Ivanenkov, Yan A., Aliper, Alex., Veselov, Mark S., Aladinskiy, Vladimir A., Aladinskaya, Anastasiya V., Terentiev, Victor A., … Aspuru-Guzik, Alán, 'Deep learning enables rapid identification of potent DDR1 kinase inhibitors.', *Nat Biotechnol*, vol. 37, no. 9, pp. 1038–1040, Sep. 2019.
53. Hosny, Ahmed., Parmar, Chintan., Quackenbush, John., Schwartz, Lawrence H., and Aerts, Hugo J.W.L., 'Artificial intelligence in radiology', *Nat Rev Cancer*, vol. 18, no. 8, p. 500, Aug. 2018.
54. Brown, Tom B., Mann, Benjamin., Ryder, Nick., Subbiah, Melanie., Kaplan, Jared., Dhariwal, Prafulla., Neelakantan, Arvind., … Amodei, Dario, 'Language Models are Few-Shot Learners', *Adv Neural Inf Process Syst*, vol. 2020-December, May 2020.
55. Gifford, George., McCutcheon, Robert., and McGuire, Philip, 'Neuroimaging studies in people at clinical high risk for psychosis', *Risk Factors for Psychosis: Paradigms, Mechanisms, and Prevention*, pp. 167–182, Jan. 2020.
56. Sollee, John., Tang, Lei., Igiraneza, Aime Bienfait., Xiao, Bo., Bai, Harrison X., and Yang, Li, 'Artificial intelligence for medical image analysis in epilepsy', *Epilepsy Res*, vol. 182, May 2022.

 doi:10.20944/preprints202505.2254.v1

57. Buolamwini, Joy., and Gebru, Timnit, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', *Proceedings of Machine Learning Research*, vol. 81. PMLR, pp. 77–91, 21-Jan-2018.
58. Acemoglu, Daron., and Restrepo, Pascual, 'The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand', in *NBER WORKING PAPER SERIES*, 2019.
59. Kim, E., Cho, H. H., Cho, S. H., Park, B., Hong, J., Shin, K. M., Hwang, M. J., … Lee, S. M., 'Accelerated Synthetic MRI with Deep Learning-Based Reconstruction for Pediatric Neuroimaging', *AJNR Am J Neuroradiol*, vol. 43, no. 11, pp. 1653–1659, Nov. 2022.
60. Rocher, Luc., Hendrickx, Julien M., and de Montjoye, Yves Alexandre, 'Estimating the success of re-identifications in incomplete datasets using generative models', *Nature Communications 2019 10:1*, vol. 10, no. 1, pp. 1–9, Jul. 2019.
61. Cadwalladr, Carole., and Graham-Harrison, Emma., 'Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach', *The Guardian*, 2018.
62. Kostka, Genia, 'China's social credit systems and public opinion: Explaining high levels of approval', *New Media Soc*, vol. 21, no. 7, pp. 1565–1593, Jul. 2019.
63. Cavicchioli, Ricardo., Ripple, William J., Timmis, Kenneth N., Azam, Farooq., Bakken, Lars R., Baylis, Matthew., Behrenfeld, Michael J., … Webster, Nicole S., 'Scientists' warning to humanity: microorganisms and climate change', *Nature Reviews Microbiology 2019 17:9*, vol. 17, no. 9, pp. 569–586, Jun. 2019.
64. Biggio, Battista., Nelson, Blaine., and Laskov, Pavel, 'Poisoning Attacks against Support Vector Machines', *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, vol. 2, pp. 1807–1814, Jun. 2012.
65. Korshunov, Pavel., and Marcel, Sebastien, 'DeepFakes: a New Threat to Face Recognition? Assessment and Detection', Dec. 2018.
66. Platt, John C, 'Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods'.
67. Chhatwal, Jagpreet., Alagoz, Oguzhan., Lindstrom, Mary J., Kahn, Charles E., Shaffer, Katherine A., and Burnside, Elizabeth S., 'A logistic regression model based on the national mammography database format to aid breast cancer diagnosis', *American Journal of Roentgenology*, vol. 192, no. 4, pp. 1117–1127, Apr. 2009.
68. Krizhevsky, Alex., Sutskever, Ilya., and Hinton, Geoffrey E, 'ImageNet Classification with Deep Convolutional Neural Networks'.
69. Lu, Xugang., Tsao, Yu., Matsuda, Shigeki., and Hori, Chiori, 'Speech enhancement based on deep denoising autoencoder', *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, pp. 436–440, 2013.
70. Radford, Alec., Wu, Jeffrey., Child, Rewon., Luan, David., Amodei, Dario., and Sutskever, Ilya, 'Language Models are Unsupervised Multitask Learners', *OpenAI Blog*, 2019.
71. Kingma, Durk P., Mohamed, Shakir., Jimenez Rezende, Danilo., and Welling, Max, 'Semi-supervised Learning with Deep Generative Models', *Adv Neural Inf Process Syst*, vol. 27, 2014.
72. Hinton, G. E., and Salakhutdinov, R. R., 'Reducing the dimensionality of data with neural networks', *Science*, vol. 313, no. 5786, pp. 504–507, Jul. 2006.
73. Elgammal, Ahmed., Liu, Bingchen., Elhoseiny, Mohamed., and Mazzone, Marian, 'CAN: Creative Adversarial Networks, Generating "Art" by Learning About Styles and Deviating from Style Norms', *Proceedings of the 8th International Conference on Computational Creativity, ICCC 2017*, Jun. 2017.
74. Brown, Tom B., Mann, Benjamin., Ryder, Nick., Subbiah, Melanie., Kaplan, Jared., Dhariwal, Prafulla., Neelakantan, Arvind., … Amodei, Dario, 'Language models are few-shot learners', *Adv Neural Inf Process Syst*, vol. 2020-December, 2020.
75. Zaha Hadid Architects, 'Zaha Hadid Architects using AI image generators for design concepts, said Patrik Schumacher', 2023. [Online]. Available: https://parametric-architecture.com/zaha-hadid-architects-using-ai-image-generators-for-design-concepts-said-patrik-schumacher/. [Accessed: 04-Sep-2023].
76. Frid-Adar, Maayan., Diamant, Idit., Klang, Eyal., Amitai, Michal., Goldberger, Jacob., and Greenspan, Hayit, 'GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification', *Neurocomputing*, vol. 321, pp. 321–331, Dec. 2018.
77. NIST, Cybersecurity_Framework, *Cybersecurity Framework | NIST*. 2016.

 doi:10.20944/preprints202505.2254.v1

Short Biographies

Petar Radanliev is a Member of Faculty in Artificial Intelligence and Cybersecurity at the University of Oxford's Department of Computer Science, and a Post-Doctoral Researcher at the Alan Turing Institute in London. Dr. Radanliev completed his PhD in 2013/14 and has since engaged in postdoctoral research at several prestigious institutions, including Imperial College London, the University of Cambridge, the Massachusetts Institute of Technology, and the Department of Engineering Science at the University of Oxford for seven years, before moving to the Department of Computer Science. Dr. Radanliev, specialises in artificial intelligence, cybersecurity, quantum security, and blockchain security. Prior to his academic career, he amassed a decade of experience as a Cybersecurity Manager at RBS, the world's largest bank at the time and five years as a Lead Penetration Tester for the Ministry of Defence.

Omar Santos is a Distinguished Engineer at Cisco who pioneers advancements in artificial intelligence security, cybersecurity research, ethical hacking, incident response, and vulnerability disclosure. As co-chair of the Coalition for Secure AI (CoSAI) and board member of the OASIS Open standards organisation, he shapes the future of secure technology adoption across industries. Omar drives innovation through multiple leadership roles, including founder of OpenEoX and co-chair of the Forum of Incident Response and Security Teams (FIRST) PSIRT Special Interest Group. His commitment to cybersecurity education and community building is evident in his role as the co-founder and one of the leaders of the DEF CON Red Team Village and the chair of the Common Security Advisory Framework (CSAF) technical committee. Omar has published over 20 books, created over 20 video courses, and contributed more than 40 academic research papers to the field. Omar's work in cybersecurity is also recognized through multiple granted patents. Omar's PGP Key: 0x8e19a9d13af27edc (and Keybase info).

Uchenna Ani is a Senior Lecturer in Cyber Security at Keele University. He completed his Ph.D. in Industrial Control System Cybersecurity at Cranfield University and continued with Post-Doctoral research as a Senior Research Fellow in Cybersecurity at the PETRAS National Centre of Excellence for IoT Systems Cybersecurity at the Department of Science Technology Engineering and Public Policy (STEaPP), University College London (UCL).