ACIG
APPLIED CYBERSECURITY & INTERNET GOVERNANCE

NASK

AI-Powered Pervasive Computing: Discerning Threat Model to Unveil Security and Privacy Challenges

Aditya K. Sood | Security Engineering and AI Strategy, Aryaka, USA | ORCID: 0000-0002-7738-2890

Sherali Zeadally | College of Communication and Information, University of Kentucky, Lexington, KY, USA; Academy of Computer Science & Software Engineering, University of Johannesburg, South Africa | ORCID: 0000-0002-5982-8190

Received: 17.02.2025

Accepted: 12.06.2025

Published: 09.07.2025

Cite this article as:
A.K. Sood, S. Zeadally, "AI-Powered Pervasive Computing: Discerning Threat Model to Unveil Security and Privacy Challenges," ACIG, vol. 4, no. 1, 2025, doi: 10.60097/ACIG/207029

Corresponding author:
Aditya K. Sood, Security Engineering and AI Strategy, Aryaka, United States; E-mail: soodadit.msu@gmail.com
0000-0002-7738-2890

Abstract

The integration of Artificial Intelligence (AI) with pervasive computing is transforming environments into intelligent, context-aware systems capable of seamless and intuitive user interactions. This convergence drives advancements across various domains, including smart homes, healthcare, and industrial automation. Traditional security paradigms often fall short in addressing AI-specific risks, such as adversarial attacks, data inference, and the complexities of autonomous decision-making in highly interconnected systems. However, it also introduces a range of security and privacy challenges that must be addressed. We critically analyse these emerging concerns, delving into the unique vulnerabilities arising from the fusion of AI's inferential capabilities with pervasive computing's ubiquitous data collection. We further explore why existing threat models designed for conventional IT infrastructures prove insufficient in adequately capturing the nuanced and dynamic risks posed by intelligent, autonomous, and deeply integrated pervasive systems. We present a practical threat model specifically tailored to highlight the key risks associated with AI-driven pervasive computing and discuss potential mitigation

strategies to enhance security and safeguard user privacy in these interconnected environments.

Keywords

pervasive computing, cybersecurity, threat model, security, privacy

1. Introduction

Artificial Intelligence (AI) integration in pervasive computing [1] is a transformative leap in daily interactions with technology. Pervasive computing, also known as ubiquitous computing, embeds computational processes into everyday objects and environments, seamlessly integrating technology into the fabric of daily life. When combined with AI, these systems can collect and process vast amounts of data from their surroundings and analyse it in real-time, enabling them to adapt to the user's needs, predict future behaviours, and provide personalised experiences. This symbiosis between AI and pervasive computing creates intelligent context-aware environments that respond dynamically to human interactions.

Integrating AI in pervasive computing drives innovation across various domains, including smart homes, healthcare, urban infrastructure, and industrial automation. In smart homes, for example, AI-powered pervasive systems [2] can learn a user's habits and preferences, automatically adjusting lighting, temperature, and even security settings to create a personalised living environment. In healthcare, wearable devices and home sensors connected to AI systems can monitor patients' health in real-time, offering proactive health management and early detection of potential issues. The convergence of AI with pervasive computing enhances convenience and efficiency, shifting the technology landscape significantly towards more intuitive and human-centric ecosystems.

Integrating AI into pervasive computing significantly amplifies security and privacy concerns because these systems rely on constant data collection and real-time processing to function effectively. With AI analysing large amounts of data from many types of sensors and devices embedded in everyday environments, the potential for security and privacy threats increases. This creates a larger attack surface for cyber threats, where vulnerabilities in the network or devices can lead to widespread privacy violations and potential misuse of personal data.

AI-powered pervasive computing also presents significant challenges for governments in cybersecurity. First, the sheer scale and interconnectedness of AI systems across various sectors, including critical infrastructure, healthcare, and public services, significantly amplify the attack surface, making it challenging for governments to secure every potential entry point. Second, the rapid evolution and complexity of AI technologies enable malicious actors to utilise AI power to conduct sophisticated attacks, including autonomous threats that adapt to defensive measures in real-time. Governments are required not just to respond but also to proactively keep pace with these advancements to develop and implement effective countermeasures. Third, the global nature of AI-powered pervasive computing introduces geopolitical risks. Many AI technologies and Internet of things (IoT) devices are designed and manufactured in different countries, some of which may have conflicting interests with others. This raises concerns about the potential for supply chain attacks or state-sponsored cyber espionage. Additionally, the use of AI in nation-state operations and espionage raises significant ethical and legal considerations, complicating the development of policies that balance security and respect for civil liberties. Governments must navigate these complex international relationships while ensuring that their infrastructure and data remain safe, a challenge that requires a careful balance of diplomacy, trade policies, and cybersecurity strategies.

Moreover, the pervasive nature of these AI-based systems can erode user privacy [3], often without the user's explicit knowledge or consent. AI's ability to cross-reference and analyse data from multiple sources increases the risk of re-identification, even from anonymised datasets. The need for more transparency in data practices and the complexity of privacy policies further complicate users' understanding of how their data is collected and used.

In this paper, we examine the security and privacy threats to pervasive computing in AI-driven environments, and we provide a comprehensive analysis of how to mitigate these threats.

1.1. Contributions of this work

Our main research contributions include the following:

- We present a holistic model of AI integration in pervasive computing.
- We discuss security and privacy challenges associated with integrating AI in pervasive computing environments and present mapping the same into a targeted threat model.

- We propose a threat model related to AI-powered pervasive computing different from traditional threat model frameworks, describing the security and privacy threats associated with it.
- We propose several effective countermeasures to mitigate the security and privacy threats to AI-powered pervasive computing.

2. The Need for AI-powered Pervasive Computing

We first describe the importance of integrating AI with pervasive computing, which leads to more powerful advanced applications and services. A pervasive computing environment refers to an interconnected, intelligent space where computational technologies are seamlessly integrated into everyday objects and surroundings, creating a network of devices that interact with one another and users in an intuitive manner. These environments are context-aware, adaptive, and responsive, leveraging embedded sensors, actuators, and networked devices to enhance functionality and personalised experiences. Typical applications of the pervasive computing environment include the following:

- *Smart homes:* Integrating devices and systems to manage home automation, security, energy efficiency, and user comfort.
- *Smart cities:* Using sensors and data analytics to enhance urban infrastructure, traffic management, public safety, and environmental monitoring.
- *Healthcare:* Wearable devices and remote monitoring systems track health metrics, provide personalised care, and support telemedicine.
- *Industrial automation:* Design intelligent manufacturing, logistics, and process cost systems to optimise operations and improve efficiency.

In recent years, the benefits of AI have led to its adoption in many areas, including pervasive computing environments. Figure 1 illustrates a holistic model that explains the reasons for integrating AI capabilities into pervasive computing.

We discuss the capabilities and benefits of an AI-powered pervasive computing environment next.

- *Context-aware systems using adaptive behaviour:* AI leverages pervasive computing to gather and analyse data, enabling embedded systems to understand the context in which they operate. AI systems adjust their behaviour based on the context provided

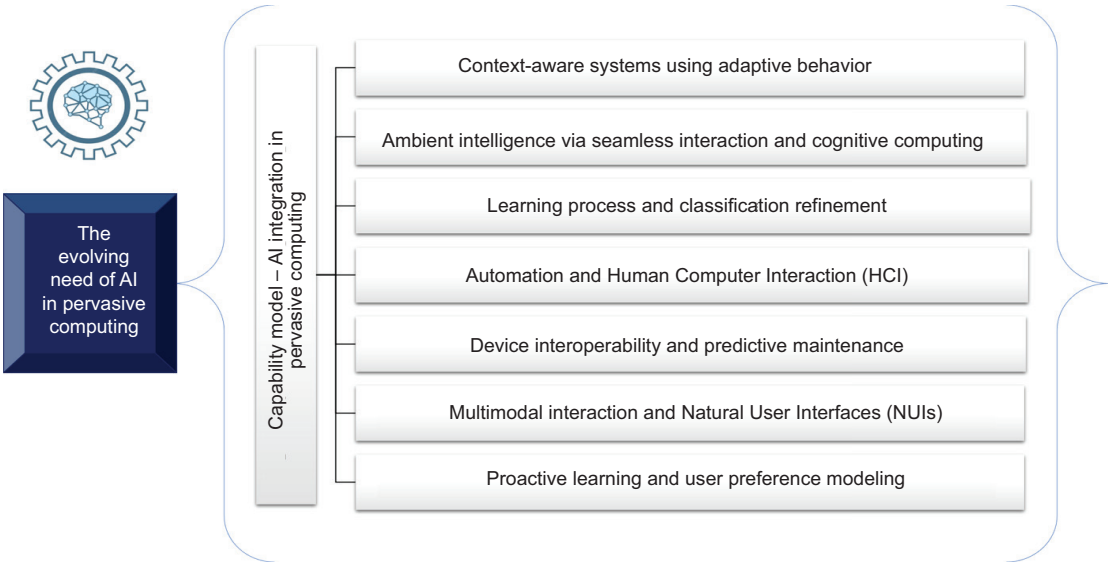

Figure 1. Capability model: AI integration in pervasive computing.

by pervasive computing. For example, a smart thermostat might adjust the temperature based on time, forecast, and the presence of people in the house.

- *Ambient intelligence via seamless interaction and cognitive computing:* AI systems embedded in pervasive computing environments can process and understand large amounts of data in real-time, enabling more intelligent decision-making and interaction. AI enhances pervasive computing by creating environments that respond to people's presence naturally and seamlessly. For example, a room might automatically adjust lighting and temperature based on the preferences of the person who enters.
- *Learning process and classification refinement:* Pervasive computing systems continuously collect data, which AI algorithms can use to learn and adapt over time. This ongoing learning process enables systems to become more accurate and responsive to user needs. AI enables user interfaces to learn and adjust based on the context of use, user behaviour, and preferences, making interactions more intuitive and effective. By integrating feedback, AI can refine its algorithms and enhance accuracy, making the system more resilient and better equipped to classify and predict outcomes.
- *Automation and human–computer interaction (HCI):* Pervasive computing combined with AI can lead to highly automated systems that reduce downtime and increase operational efficiency. For example, AI-powered predictive maintenance systems can

monitor equipment health and proactively predict failures. Additionally, AI enhances HCI in pervasive computing environments by enabling more natural and intuitive interactions. Voice-activated assistants, such as Amazon Alexa or Google Assistant, utilise AI to interpret and respond to voice commands efficiently.

- *Device interoperability and predictive maintenance:* Pervasive computing enables the integration of diverse IoT devices into a cohesive system. AI plays a crucial role in managing these devices, ensuring they work together smoothly and respond effectively to environmental changes. In industrial settings, the combination of pervasive computing and AI can predict when machines are likely to fail by analysing data from embedded sensors, thereby preventing downtime and minimising the impact of machine failures.
- *Multimodal interaction and natural user interfaces (NUIs):* Pervasive computing often involves natural interfaces powered by AI, such as voice and gesture recognition, to enable users to interact with the environment intuitively, for example, by issuing voice commands to control home appliances. AI systems can analyse and integrate multiple input forms, such as combining voice, touch, and gestures from various devices deployed in pervasive computing environments, to create seamless experiences.
- *Proactive learning and user preference modelling:* Pervasive computing enables AI to collect data on user behaviours and preferences over time, facilitating more informed decision-making. This data helps the creation of detailed user models, allowing systems to provide personalised experiences, such as recommending products, adjusting home settings for appliances, or delivering tailored content. AI systems can anticipate user needs based on routines and habits, offering proactive assistance. For instance, a smart home system might automatically start brewing coffee when it senses the user has woken up.

Pervasive computing environments significantly improve how technologies interact with and enhance daily life, creating a more interconnected, responsive, and intelligent world.

3. Security and Privacy Challenges

Next, we discuss several critical security and privacy challenges associated with embedding AI in pervasive computing systems.

- *Excessive data collection and storage risks:* Pervasive computing environments collect data from multiple sources, including

sensors, IoT devices, and user interactions for AI operations. The increasing amount of data extends the attack surface, making these systems prime targets for cyberattacks. In addition, the centralised storage of large datasets creates single points of failure. Breaches in these databases can expose sensitive personal information, impacting millions of users.

- *AI model attacks:* Attackers subvert the outcomes of AI models in pervasive systems by introducing adversarial inputs [4]. This could lead to incorrect or harmful decisions, posing risks in critical applications, such as autonomous vehicles or healthcare. Additionally, attackers can reverse-engineer AI models to extract sensitive information about the training data, leading to potential security and privacy breaches.
- *Attack surface expansion due to connected devices:* Pervasive computing systems heavily rely on networked communication, and many IoT devices have limited processing power, making it challenging to implement robust security protocols. This creates a potential vulnerability that attackers can exploit, using these devices as entry points for attacks on the broader network, including AI systems. For example, attackers target AI-powered systems and connected devices in pervasive environments to form botnets (compromised device networks) [5], which malicious actors can control to launch advanced attacks on connected networks. These attacks can result in service disruptions, financial losses, and even physical damage in scenarios where IoT devices control critical infrastructures.
- *Surveillance risks:* Integrating AI with pervasive computing enhances surveillance capabilities, allowing for the real-time tracking and monitoring of individuals [6]. AI can analyse large datasets to create detailed profiles of individuals, predicting their behaviours and preferences. Attackers can use entity profiling for targeted advertising or more nefarious purposes, such as discrimination or social control.
- *Sensitive information disclosure:* Pervasive systems often collect contextual data that, when combined with other information, can reveal sensitive details about individuals, even without direct identification. AI's ability to identify patterns and correlations in the data collected from a pervasive computing environment increases the risk of re-identification, thereby undermining privacy protections.
- *Privacy risks due to opaque data collection:* Users may need to understand fully how AI systems deployed in pervasive environments collect data using non-transparent practices. This can lead to uninformed consent, ethical concerns, and potential regulatory violations.

- *Bias and fairness in AI decision-making:* AI systems trained on biased or incomplete data can produce discriminatory outputs. These biases may not be immediately apparent and can perpetuate systemic inequalities. In the context of pervasive computing, this could lead to unfair profiling or exclusion of certain groups.
- *System misconfiguration due to human error:* Human error remains a major contributor to security vulnerabilities, especially in complex AI-integrated systems used for pervasive computing. Misconfigured access controls, unpatched services, or improper deployment of AI models can expose sensitive data or enable unauthorised access. As these systems scale, even minor misconfigurations can lead to significant data breaches.
- *Interoperability and standards weaknesses:* The lack of uniform standards across AI and pervasive computing vendors can result in inconsistent security practices. Devices that fail to authenticate, encrypt properly, or update can become weak links in otherwise secure networks. In pervasive environments, this fragmented approach increases attack surfaces and complicates incident response.
- *Insider threats or third-party supply chain risk:* Insiders or third-party vendors often have deep access to systems, making them high-risk vectors for data leaks or malicious activity. In AI-powered systems, this could involve manipulating training data, embedding backdoors, or extracting sensitive models. The complexity of supply chains and outsourced roles magnifies this risk.

Table 1 illustrates how to create a threat model using the challenges mentioned above, which encompasses threat vectors, attacker goals, and system vulnerabilities, using the use case example.

Table 1 presents a practical balance between breadth and depth, making them ideal for most threat-modelling exercises. After discussing the intrinsic challenges related to security and privacy in AI-powered pervasive computing, the following section outlines several reasons that highlight why the existing threat model frameworks are insufficient.

4. Why Existing Threat Model Frameworks Are Not Sufficient?

Existing threat modelling frameworks, such as STRIDE [7], DREAD [8], and MITRE ATT&CK [9], have played a crucial role in identifying and mitigating risks in traditional computing environments. However, these models were primarily designed for conventional software systems and architectures, which are relatively

Table 1. Mapping security and privacy challenges to threat model.

Challenge	Threat vector	Attacker goal	System vulnerabilities	Example scenario
Excessive data collection and storage risks	Unauthorised access to data repositories and poor access control	Harvest personal or behavioural data for resale, surveillance, or blackmail	Overcollection, lack of access controls, and insecure cloud storage	Smart home devices often store voice and audio recordings in unencrypted cloud storage
AI model attacks	Adversarial inputs, model inversion, poisoning, model extraction	Bypass detection systems, reverse-engineer the model, and extract sensitive training data	Lack of input validation, poor model monitoring, and reliance on unverified third-party models	Attackers feed adversarial malware code to bypass machine learning (ML)-based endpoint protection
Attack surface expansion due to connected devices	Exploiting unpatched or poorly configured IoT/edge devices	Pivot into the network, perform lateral movement, and disrupt services	Weak firmware, default credentials, and lack of segmentation	A compromise of smart sensor in a factory leads to the spread of ransomware through the operational technology (OT) network
Surveillance risks	Misuse or hijacking of AI-powered monitoring systems	Mass surveillance, behavioural tracking, and political profiling	Centralised logging, lack of oversight, and missing anonymisation	Hijacked public smart cameras are used to track individuals across city intersections
Sensitive info disclosure	Leakage from logs, model inference, and weak data classification	Identity theft, reputational damage, and financial fraud	Lack of personally identifiable information (PII) classification, exposed logs, and inference attacks	AI assistant logs include PII from user queries, which an unauthorised third-party vendor can access
Privacy risks due to opaque data collection	Covert data harvesting via sensors, apps, or wearables	Behavioural profiling, unauthorised targeting, and data resale	No user consent, hidden telemetry, dark patterns	Health wearable devices collect sleep and location data without users' awareness and consent and share it with advertisers

(*continues*)

Table 1. Continued.

Challenge	Threat vector	Attacker goal	System vulnerabilities	Example scenario
Bias and fairness in AI decision-making	Poisoned training data, opaque model logic, and ineffective feedback loops	Skewed outcomes for specific groups and discriminatory profiling	Lack of diverse training data and non-transparent model behaviour	AI facial recognition misidentifies minority groups more often, leading to biased and unjust enforcement
System misconfiguration due to human error	Misconfigured permissions and incorrect network settings	Unauthorised access and disruption of system operations	Poor DevSecOps hygiene and lack of automated checks or validation	System accidentally exposes IoT dashboard with default credentials to the public Internet
Interoperability and standards weaknesses	Insecure application programming interfaces (APIs) and proprietary protocol mishandling	Exploit vendor-specific gaps and break cross-system trust	Lack of universal security standards and poor validation at integration points	An innovative door system fails to encrypt messages when integrated with a third-party home hub
Insider threats or third-party supply chain risk	Compromised contractors and malicious insider actions	Exfiltrate data and install persistent malicious code	Over-permissive access and lack of activity monitoring	A subcontractor copies internal AI models used in critical infrastructures and sells them externally

ACIG APPLIED CYBERSECURITY & INTERNET GOVERNANCE

deterministic and well understood. In contrast, AI operates in a dynamic, data-driven paradigm, particularly in pervasive computing environments, where system behaviours evolve based on training data, model drift, and external input patterns. These unique attributes, which are not fully addressed by traditional threat models, are a fundamental challenge in AI security because these models often assume static, rule-based systems.

AI in pervasive computing (e.g., smart homes, edge devices, and IoT-enabled healthcare) introduces new classes of threats, such as adversarial inputs, model poisoning, and inference attacks. These threats exploit the statistical nature of AI systems and can subtly alter outcomes without triggering traditional alarms. For example, adversaries may manipulate sensor data to mislead an AI model in an intelligent surveillance system or use membership inference attacks to extract sensitive training data. Such attacks do not fit neatly into categories like spoofing or elevation of privilege from STRIDE, nor are they adequately described in existing risk matrices or impact scoring models used by conventional threat frameworks.

Furthermore, the integration of AI into edge and pervasive devices often lacks centralised oversight and standardisation, complicating threat surface visibility and response coordination. This underscores the need for specialised threat modelling approaches that incorporate AI-specific risks, data-centric threats, and evolving attack vectors unique to pervasive environments. These approaches should consider the distributed, context-aware, and autonomous characteristics of these systems, including the lifecycle of AI models and data pipelines. They should also consider real-time decision-making, continual learning systems, and the implications of AI behaviour unpredictability. Therefore, the current frameworks must evolve or be extended to remain relevant and practical in the fast-changing AI-driven security landscape. As a result, our proposed threat framework is tailored to AI-powered pervasive computing environments, enabling effective handling of security and privacy-related risks.

5. AI-Powered Pervasive Computing Threat Model

This section presents threat actors, including a threat model specific to AI-powered pervasive computing. A good understanding of the different types of threat actors that can target AI-powered pervasive computing environments is essential.

5.1. Threat Actors

Threat actors [10] are individuals, groups, or entities that engage in malicious activities targeting computer systems, networks, and digital information. They vary widely in terms of their motivations, methods, and capabilities. Table 2 presents the threat actors that could target the AI-powered pervasive computing environments.

Next, we discuss the threat model to understand the threats associated with integrating AI in pervasive computing environments.

5.2. Proposed Threat Model

We discuss the real-world threat model [11] in a unified manner, tailored explicitly to security and privacy, as shown in Fig. 2, and describe the impact of integrating AI in pervasive computing environments in Table 3.

Table 3 presents the security and privacy threats, along with their descriptions and impacts.

Security and privacy threats can have a significant impact on individuals and organisations. Mitigating these threats requires a combination of advanced controls, which we discuss in the next section.

5.3. Understanding the Threat Model Workflow

Creating a threat model workflow is crucial for systematically identifying and mitigating security and privacy risks in AI-powered pervasive systems. It helps to anticipate potential attack paths, understand system vulnerabilities, and prioritise defences. This proactive approach strengthens the overall system resilience and ensures compliance with privacy and security standards. We present the workflow of the proposed threat model using the proposed threat modelling framework.

- *Define scope and identify assets:* Identify all AI components (models and training data), pervasive devices (IoT and sensors), data types (PII and telemetry), and stakeholders involved. Clarify boundaries and trust levels across systems and vendors to ensure effective collaboration.
- *Describe the architecture:* Create a visual representation of the system that depicts data flows between devices, cloud/edge processing layers, AI inference locations, and control interfaces. This helps to identify weak points and dependencies.

Table 2. Threat actors targeting AI-powered pervasive computing environment.

Threat actor	Motivation	Capabilities
Nation-state actors	These actors are driven by geopolitical objectives, including espionage, sabotage, and gathering intelligence. Nation-state actors may target AI-powered systems to disrupt critical infrastructures, steal sensitive information, or gain strategic advantages	Nation-state actors typically possess significant resources and advanced capabilities, enabling them to execute sophisticated attacks, including zero-day exploits, advanced persistent threats (APTs), and supply chain attacks
Cybercriminal organisations	Financial gain is the primary motivator for cybercriminal organisations. They may target pervasive computing environments to conduct ransomware attacks, steal personal and financial data, or engage in fraud	These groups often employ many tactics, including phishing, malware distribution, and exploiting vulnerabilities in IoT devices and AI systems. They may also use AI to enhance their attack strategies, such as automating phishing campaigns or evading detection
Competitors and industrial spies	Competitors may use corporate espionage to gain an advantage by stealing proprietary information and trade secrets or disrupting operations.	Industrial spies may employ tactics such as infiltrating organisations, planting insiders, or using advanced malware to exfiltrate sensitive and valuable data. They may also target AI algorithms to manipulate market outcomes or product development.
Malicious insiders	Malicious insiders, such as disgruntled employees or contractors, may harbour personal grievances, financial incentives, or ideological motivations that lead them to sabotage systems or steal data	Insiders have access to critical data and systems, making them uniquely positioned to carry out attacks that may go undetected for extended periods. They can exploit their knowledge of the system to bypass security controls or introduce malicious code
Hacktivists	Ideological or political goals drive hacktivists. They may target AI-powered systems in pervasive computing environments to promote their causes, disrupt services, or draw attention to specific issues	While hacktivists may have less sophisticated capabilities than nation-state actors, they can still carry out significant attacks, such as defacing websites, launching Distributed Denial of Service (DDoS) attacks, or leaking sensitive data
Cybersecurity researchers and ethical hackers	While not malicious, these actors may intentionally target systems to identify vulnerabilities and improve security. Their activities, however, can sometimes be mistaken for malicious intent, mainly if conducted without permission	Ethical hackers and researchers typically use advanced technical skills to enhance system security by conducting penetration testing, vulnerability assessments, and demonstrating their exploits
Script kiddies	These actors have a limited skill set and are typically driven by curiosity, boredom, a desire for peer recognition, or simply the thrill of causing disruption. Their attacks are often opportunistic, rather than targeted	Script kiddies typically lack in-depth technical expertise. Instead of writing original code or discovering novel vulnerabilities, they rely heavily on pre-built exploit kits, automated scripts, and freely available tools downloaded from the Internet or shared within underground forums. While their tools can still cause disruption, script kiddies cannot generally perform advanced, persistent, or stealthy attacks

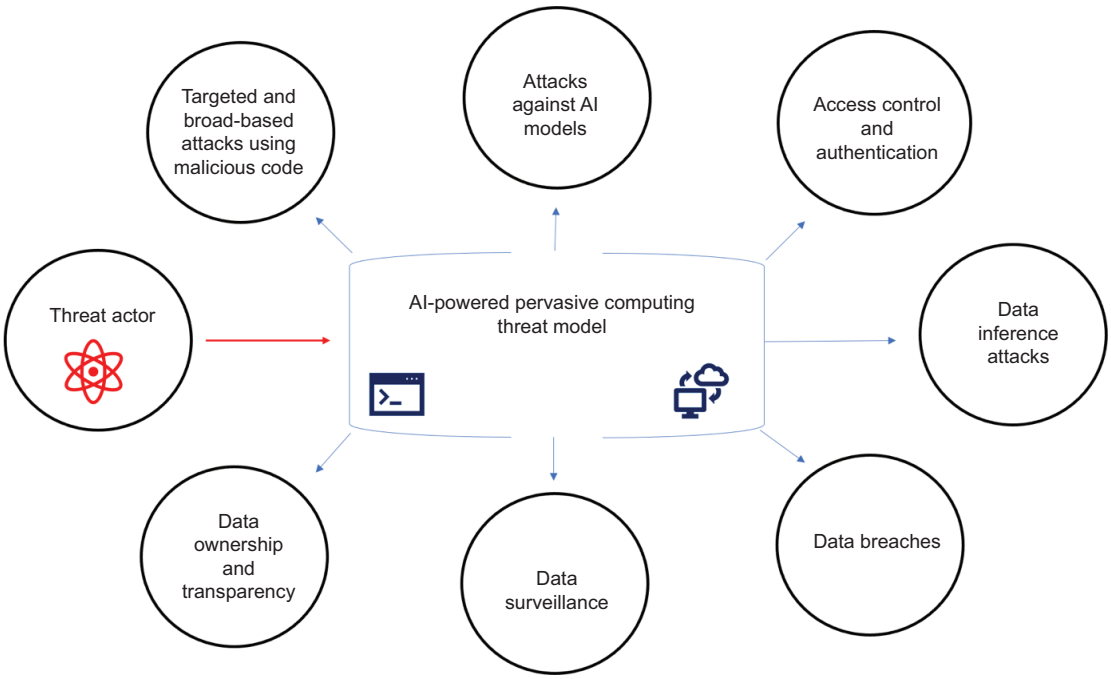

Figure 2. Security and privacy threat model for AI-powered pervasive computing.

- *Identify entry points:* List all physical and logical interfaces, such as API endpoints, bluetooth/near field communication (NFC), admin consoles, and user interfaces. These are potential vectors that attackers may exploit to gain access or manipulate data.
- *Enumerate threats:* Utilise the proposed framework to construct a threat library tailored to one's architecture to ensure a comprehensive risk coverage.
- *Map system vulnerabilities:* Assess components for flaws, including misconfigured APIs, outdated firmware, unsecured AI pipelines, or weak authentication. Identify exploitable weaknesses in both traditional IT and AI subsystems.
- *Define attacker profiles and goals:* Describe potential threat actors: nation-states, cybercriminals, insiders, or malicious AI users. Document their motivations, such as stealing data, disrupting operations, or influencing AI decision-making.
- *Prioritise risks:* Evaluate threats based on impact (e.g., data breach vs. service disruption), likelihood, and detection feasibility. Utilise scoring systems as approved by the organisations to prioritise high-risk threats.
- *Propose mitigating solutions:* Recommend layered defences to address both general cybersecurity and AI-specific risks.

Table 3. Security and privacy threats along with their impact for an AI-powered pervasive computing environment.

Security and privacy threats	Threat description	Impact
Targeted and broad-based attacks using malicious code	AI-powered systems running in pervasive computing environments can be targeted [12] by sophisticated malicious code that spreads across devices, forming botnets that attackers can control and manipulate	The widespread impact of such malicious code can lead to service disruptions, financial losses, and even physical damage in scenarios where IoT devices control critical infrastructures. For example, malicious code enables network-based attacks, such as eavesdropping, man-in-the-middle attacks, and DDoS attacks, which can disrupt operations and compromise the integrity of data
Attacks against AI models	AI models are vulnerable to adversarial attacks, model poisoning, and backdoor attacks, where subtle, crafted inputs cause the AI to make incorrect or harmful decisions. Attackers can reverse-engineer AI models [13] to extract sensitive information about the training data, leading to potential security and privacy breaches.	Exploiting AI model vulnerabilities can lead to scenarios where attackers manipulate innovative systems for malicious purposes, such as causing smart homes to unlock doors or turning off security cameras. For instance, an adversarial attack could cause an AI-powered security camera to misidentify a threat or fail to recognise an intruder.
Access control and authentication	Devices running in AI-integrated pervasive environments require strong access controls to restrict unauthorised access. Many devices, such as IoT sensors and smart appliances, have limited processing power and security features. As a result, attackers gain access to the network and manipulate the functionalities of devices	Weak or compromised access control can allow unauthorised entities to gain control over critical systems, potentially leading to sabotage, espionage, or the theft of sensitive information. For example, if IoT devices have weak or default passwords, attackers can easily guess or brute-force them to gain unauthorised access to these critical devices.
Data breaches	Significant amounts of sensitive data are collected from smart devices, sensors, and wearable devices. Maintaining data integrity, including confidentiality in transit and at rest, becomes increasingly challenging as the number of data points and communication channels increases	Compromised data integrity can lead to incorrect AI decisions, potentially harming critical applications like healthcare or autonomous driving. Breaches of confidentiality can expose sensitive personal information, leading to privacy violations
Data surveillance	Pervasive computing environments collect a lot of personal data without explicit user consent. AI systems process this data to provide personalised services; however, the extensive collection of data raises significant concerns about privacy and data security	Continuous monitoring and data collection can generate detailed profiles of individuals, which attackers can use for unauthorised activities, leading to a significant invasion of privacy
Data inference attacks	AI systems in pervasive environments often infer sensitive information from seemingly benign data. For instance, location data combined with activity data can infer a user's health status, habits, or even relationships. Attackers can exploit these inferences to gain unauthorised access to users' private information	Inference attacks can expose sensitive personal information without direct access to the data, making privacy violations challenging to detect or prevent

(*continues*)

Table 3. Continued.

Security and privacy threats	Threat description	Impact
Data ownership and transparency	Users often lose control over their data once AI systems collect it, raising questions about who owns it and how it can be used for unauthorised purposes. AI systems in pervasive environments often operate in the background, making it difficult for users to understand what data is being collected and how it is used. Obtaining informed consent becomes challenging, especially when the data collection is continuous, and AI must be more transparent	Lack of control over personal data can lead to its misuse by third parties, including the sale of data to advertisers or other organisations without user consent. This can result in privacy breaches and a loss of trust in pervasive systems. The lack of transparency and consent can lead to privacy concerns and legal challenges. Users may feel violated if they discover that their data has been used in ways for which they did not explicitly consent

- *Monitor and update regularly:* Establish processes for periodic threat reviews, incident response, and retraining models securely. Update the threat model as devices are added, software changes, or new AI threats emerge.

Overall, this process is iterative to ensure continued protection in dynamic, intelligent environments.

6. Why is the Proposed Threat Modelling Framework Efficient?

Our proposed threat modelling framework is not just a theoretical concept but a unique and practical tool that can be applied to real-world AI-centric pervasive computing scenarios. It is adaptable and efficient for modelling security and privacy risks associated with AI-powered pervasive computing. The framework handles several practical challenges discussed below that are not addressed in other conventional frameworks, as they don't take AI-centric attacks, including security and privacy issues associated with the AI ecosystem, into consideration. The key features and functionalities of the proposed framework include the following:

- It addresses AI-specific threats, including adversarial AI attacks (e.g., model evasion and poisoning), inference attacks, and the manipulation of training data. These threats are unique to systems leveraging AI solutions that leverage machine learning and deep learning.
- It comprehensively addresses data-centric and model lifecycle risks, as pervasive environments rely on continuous data streams

and often involve real-time learning. In pervasive systems, data is usually crowd-sourced, unverified, or context-sensitive, adding layers of unaccounted risks. It can provide insights into risks across the AI lifecycle, covering attack vectors that target training or inference data streams as well as the risks associated with compromised data pipelines or insecure model updates, all of which provide a comprehensive view of the overall system's security.

- It is designed to help practitioners effectively handle the autonomous and context-aware behaviour that AI-driven systems often exhibit in real-world scenarios. It is not focused on static system exploits and network layer techniques; instead, it covers context-dependent behaviour, decision boundaries, and the probabilistic nature of AI systems, as they often operate in dynamic, constantly changing environments. This makes it a practical and relevant tool for your organisation's security needs.
- It addresses security and privacy challenges around model transparency, explainability, and integrity. Threats such as model stealing, exfiltration of learned data, and a lack of explainability leading to hidden failures are critical in safety-sensitive pervasive systems.
- It handles threats associated with AI model deployment vectors. AI models can be deployed in containers, edge devices, or serverless functions. It succinctly considers attack vectors specific to AI model hosting, API exposure, and secure model delivery (e.g., via AI operation pipelines).
- As AI agents in pervasive environments interact with each other and with users autonomously, the framework provides a mechanism to model trust boundaries between agents, humans, and data sources, making it inadequate for multi-agent or decentralised environments.

It is worth mentioning also that organisations can use our proposed threat model framework in conjunction with the existing frameworks to map traditional infrastructure threats that support AI systems (e.g., cloud platforms, endpoints, and identity management).

7. Proposed Mitigating Solutions

We propose several mitigation strategies to help defend against security and privacy threats specific to using AI in pervasive computing environments. Table 4 presents several technical and non-technical controls that can strengthen the design of an AI-powered pervasive computing environment by making informed decisions upfront, considering both security and privacy risks.

Table 4. Technical controls to improve the security and privacy of an AI-powered pervasive computing environment.

Technical and non-technical control categories	Description
Data encryption and secure communication	• Encrypt all data in transit and at rest to prevent unauthorised access. • Use protocols, such as transport layer security/secure sockets layer (TLS/SSL), to safeguard data exchanges in the pervasive computing environment.
Privacy-preserving technologies	• Apply techniques, such as differential privacy and homomorphic encryption, to minimise the risk of re-identification. • Implement federated learning to keep data localised on user devices while sharing only aggregated model updates.
IoT device security	• Enable secure, over-the-air firmware updates to patch vulnerabilities in IoT devices, ensuring continuous security protection. • Implement robust authentication mechanisms, such as multi-factor authentication (MFA), and enforce strong access controls.
AI-system security	• Train AI models to combat adversarial inputs by including adversarial examples in the training dataset. • Scan AI models [14] hosted in the code repository for embedded malicious code. • Develop explainable AI systems to ensure transparency and trust in AI decision-making processes.
Network security	• Segment networks to isolate IoT devices from systems that are part of the critical infrastructure, limiting the impact of potential breaches and damages. • Deploy advanced intrusion detection and prevention systems to monitor and protect against unauthorised access and attacks.
Continuous monitoring and incident response	• Implement continuous monitoring to detect and respond to threats in real-time. • Develop a robust incident response plan that outlines the steps for addressing security breaches.
AI-powered security	• Use AI-powered tools to detect threats in real-time by identifying abnormal behaviour, detecting anomalies in data, and responding to attacks dynamically. • Utilise AI tools to automatically respond to detected threats by isolating affected devices or network segments to restrict the spread of malware or further unauthorised access.
Regulatory compliance	• Ensure that organisations adhere to the relevant data protection regulations, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), through regular audits and compliance checks. • Establish clear protocols for timely notification to affected entities and regulatory authorities in case of data breaches.
Education and awareness	• Educate users on the security and privacy risks of pervasive computing and how they can protect themselves. • Provide ongoing training on security best practices, emerging threats, and privacy-preserving technologies.
Ethical and legal oversight	• Incorporate ethical considerations into AI development, focusing on bias mitigation, fairness, and user autonomy. • Establish ethical review boards to oversee the deployment and operation of AI-powered systems, assessing potential ethical implications.
Transparency and control	• Communicate how data is collected, processed, and used, ensuring users are informed and can consent to data practices. • Provide granular privacy settings allowing users to manage data collection and usage according to their preferences.

(*continues*)

Table 4. Continued.

Technical and non-technical control categories	Description
Security by design	• Adopt a security-by-design [15] approach, integrating security considerations into the development process. • Conduct regular security audits to discover and fix vulnerabilities in the pervasive computing environment.
Trust and accountability	• Enhance transparency by ensuring AI decisions are explainable and understandable to users. • Implement clear accountability frameworks to ensure that organisations and individuals are held responsible for data breaches or unethical practices.

As AI-driven pervasive computing becomes more integrated into daily life, addressing these security and privacy challenges through robust controls and stringent regulatory frameworks is crucial to protect devices and users in an increasingly connected world. By combining these technical and non-technical measures, organisations can create a more secure and privacy-conscious AI-powered pervasive computing environment, reducing the risks associated with these advanced technologies.

8. Future Research

As designers, researchers, and practitioners, we must strive to address the security and privacy concerns associated with integrating and using AI in pervasive computing. First, we must conduct advanced longitudinal studies to understand better the long-term societal impacts of pervasive computing and AI integration, particularly regarding privacy erosion and surveillance. Second, we must conduct further research on AI-powered data analytics for pervasive computing systems that provide valuable insights without compromising user security and privacy. Third, we must perform investigative and empirical studies to determine how AI can power pervasive computing systems and secure them against evolving threats. These studies will help the field design robust, secure, and privacy-preserving systems that run AI in conjunction with pervasive computing.

9. Conclusion

AI drives the intelligence, adaptability, and responsiveness of pervasive computing environments, making them more effective in seamlessly integrating into our daily lives. However, it also introduces significant security and privacy challenges that organisations

must address to ensure the safe and ethical use of these technologies. Interdisciplinary collaboration between technologists, ethicists, and policymakers is paramount to addressing the real-world challenges posed by the convergence of AI and pervasive computing.

Acknowledgements

We thank the anonymous reviewers for their valuable comments, which helped us improve the content, organisation, and presentation of this work.

References

[1] L. Kagal, T. Finin, A. Joshi. "Trust-based security in pervasive computing environments," *Computer*, vol. 34, no. 12, pp. 154–157, 2001, doi: 10.1109/2.970591.

[2] A. Kumar, T. Braud, S. Tarkoma, P. Hui, "Trustworthy AI in the age of pervasive computing and big data," in *2020 IEEE international conference on pervasive computing and communications workshops* (PerCom Workshops), Austin, TX, USA, 2020, pp. 1–6. Available: https://ieeexplore.ieee.org/document/9156127 [Accessed: Jan. 10, 2025].

[3] J. Hong, "The privacy landscape of pervasive computing," *IEEE Pervasive Computing*, vol. 16, no. 03, pp. 40–48, 2017. Available: https://www.computer.org/csdl/magazine/pc/2017/03/mpc2017030040/13rRUwgQpAi [Accessed: Jan. 10, 2025].

[4] Y. Wang, T. Sun, S. Li, X. Yuan, W. Ni, E. Hossain et al., "Adversarial attacks and defenses in machine learning-empowered communication systems and networks: A contemporary survey," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 4, pp. 2245–2298, 2023. Available: https://ieeexplore.ieee.org/abstract/document/10263803 [Accessed: Jan. 6, 2025].

[5] A.K. Sood, S. Zeadally, R.J. Enbody, "An empirical study of HTTP-based financial botnets," *IEEE Transactions on Dependable and Secure Computing*, vol. 13, no. 2, pp. 236–251, Mar–Apr 2016. Available: https://ieeexplore.ieee.org/document/6991594 [Accessed: Jan. 20, 2025].

[6] A. Oulasvirta, P. Aurora, J. Perkiö, D. Ray, T. Vähäkangas, et al. "Long-term effects of ubiquitous surveillance in the home," in *Proceedings of the 2012 ACM conference on ubiquitous computing* (*UbiComp '12*). New York, NY: Association for Computing Machinery, 2012, pp. 41–50, doi: 10.1145/2370216.2370224.

[7] Learn Microsoft. (2008). STRIDE. [Online]. Available: https://learn.microsoft.com/en-us/azure/security/develop/threat-modeling-tool-threats [Accessed: Dec. 5, 2024].

[8] Learn Microsoft. (2018). DREAD. [Online]. Available: https://learn.microsoft.com/en-us/windows-hardware/drivers/driversecurity/threat-modeling-for-drivers [Accessed: Dec. 8, 2024].

[9] MITRE Corporation. (2013). MITRE ATT&CK. [Online]. Available: https://www.mitre.org/focus-areas/cybersecurity/mitre-attack [Accessed: Dec. 16, 2024].

[10] A. Villalón-Huerta, H. Marco-Gisbert, I. Ripoll-Ripoll. "A taxonomy for threat actors' persistence techniques," *Computers & Security*, Vol. 121, C, p. 102855, 2022, doi: 10.1016/j.cose.2022.102855.

[11] W. Xiong, R. Lagerström, "Threat modeling – A systematic literature review," *Computers & Security*, Vol. 84, C, pp. 53–69, 2019, doi: 10.1016/j.cose.2019.03.010.

[12] A.K. Sood, R.J. Enbody, "Targeted cyberattacks: A superset of advanced persistent threats," *IEEE Security & Privacy*, vol. 11, no. 1, pp. 54–61, 2013 , doi: 10.1109/MSP.2012.90.

[13] K. Li. (Dec 26, 2018). Reverse engineering of AI models. [Online]. Hack in the Box Security Conference. Available: https://conference.hitb.org/hitbsecconf2018dxb/materials/D1T1%20-%20AI%20Model%20Security%20-%20Reverse%20Engineering%20Machine%20Learning%20Models%20-%20Kang%20Li.pdf. [Accessed: Dec. 22, 2024]].

[14] Axios. (2024). New tool will scan AI models for malware. [Online]. Available: https://www.axios.com/2024/01/24/protect-ai-malware-scanning-tool [Accessed: Dec. 1, 2024].

[15] F.M. Awaysheh, M.N. Aladwan, M. Alazab, S. Alawadi, J.C. Cabaleiro, et al., "Security by design for big data frameworks over cloud computing," *IEEE Transactions on Engineering Management*, vol. 69, no. 6, pp. 3676–3693, 2022, doi: 10.1109/TEM.2020.3045661.