Trust in the Machine: An Actor-Network Analysis of Privacy in Apple Intelligence

An STS Research Paper Submitted to the
Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

Riley Immel
Spring 2025

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature Riley Immel Date 05/02/2025
Riley Immel

STS Advisor: Richard D. Jacques, Ph.D., Department of Engineering & Society

Introduction

Artificial Intelligence (AI) has become one of the fastest-growing technologies of modern times transforming everyday technology—enabling innovations like image generation and advanced natural language processing—yet its rapid integration into consumer devices brings pressing concerns about data privacy and user trust. As our digital ecosystems expand, the balance between technological advancement and privacy protection, specifically of sensitive user data, becomes increasingly critical. This paper seeks to examine these tensions through the lens of Apple's new AI system, Apple Intelligence, which pushes the boundaries of both innovation and privacy concerns.

Apple Intelligence is engineered to perform most tasks directly on the device—a method that not only helps protect sensitive information by keeping it local but is also considered the gold-standard for data privacy. Apple already uses this paradigm for most of their existing processes that deal with sensitive user data—Apple has what is called the "Secure Enclave" which is separate entirely from the main System On a Chip (SOC) which manages everything else. The Secure Enclave oversees sensitive computations like FaceID which involve 2D and 3D scans of the user's face every time FaceID is used.

However, with Apple Intelligence, when the device lacks sufficient resources for complex processing, the system shifts to a backup mechanism called Private Cloud Compute (PCC). PCC refers to Apple's proprietary cloud server infrastructure that manages these intensive tasks. For instance, if a smartphone needs additional processing power to analyze a high-resolution image, the relevant data (such as text messages, emails, or photos) may be transmitted to the PCC, potentially exposing it to higher risks.

Concerns arise over these design choices thanks to the substantial amounts of user data that Apple Intelligence relies upon. When the PCC needs to be used to manage a request, the necessary data is sent off the user's device and to the PCC over the internet, which is inherently risky as it is more exposed to the outside world and potential attackers.

The significance of these design choices is profound as they extend beyond just the technical questions of privacy but reach further into the realm of user trust. With millions of devices in circulation (Apple sold 217.7 million iPhones in 2018 alone[1]) that feature this recent technology, the impact of any privacy vulnerability extends across a vast user base. This hybrid processing model, while innovative, raises ethical and regulatory questions about consumer trust, corporate accountability, regulatory oversight, and the broader implications of data management in a networked society.

To explore these complexities, this paper employs the Actor-Network Theory (ANT). ANT is a framework that examines the dynamic interactions between human actors (like users and regulatory bodies) and non-human actors (such as devices, cloud servers, and software systems). By applying ANT, the study aims to reveal how these diverse elements work together to shape perceptions of trust and security within Apple's ecosystem.

I argue that although Apple Intelligence's approach of combining on-device processing with a cloud-based fallback is a forward-thinking strategy, it also introduces critical vulnerabilities. These vulnerabilities not only challenge the technical integrity of the system but also disrupt established privacy paradigms—ultimately undermining consumer trust in a society increasingly reliant on interconnected digital technologies.

[1] Apple stopped reporting iPhone sales in Q4 2018, thus all other numbers are simply estimates

Methodology

The items covered in this section are the research approach, data collection process, source selection criteria, analytical framework, and steps taken to ensure credibility and rigor.

Research Approach

This study adopts a qualitative, interpretative approach. By analyzing narratives, design principles, and contextual documentation rather than quantitative data, it captures nuanced socio-technical interactions—how on-device processing and cloud fallback shape user trust and privacy. Moreover, because Apple's AI privacy strategy evolves rapidly, this exploratory framework adapts to emerging patterns and open-ended questions.

Data Collection

Because this research is qualitative and therefore does not involve new experiments or human subjects, it relies on collecting and analyzing existing sources. The "data" for this study consisted of a wide range of documents and literature that shed light on Apple Intelligence and its privacy implications as well as the current state of the art in data privacy. The gathering process began with a comprehensive search of scholarly and industry resources related to Apple's AI, on-device computing, cloud privacy, data privacy, and user trust. Key source types included:

- Academic literature—peer-reviewed journal articles and conference papers on AI privacy, data security, and trust in technology. These sources provide theoretical perspectives and empirical findings from the broader research community.

- Industry and government reports—for example, whitepapers and analysis from tech industry groups or oversight bodies (including a Congressional Research Service report on AI and data privacy) that detail current privacy challenges and regulatory considerations. Such reports offer authoritative insights into the state of the art and policy context.

- Apple's official documentation, the company's published materials (e.g. Apple's *Apple Intelligence* preview webpage and security whitepapers), and developer documentation, which describe how the technology works and the privacy measures in place from Apple's perspective as well as their efforts to assuage user concerns. Reviewing Apple's statements is crucial to understanding the intended design and trust assurances (for instance, Apple's guides on security features like biometric protection and the Secure Enclave).

- Cybersecurity analyses and news articles—expert analyses (such as security blogs or case studies) and reputable tech news coverage examining Apple Intelligence or similar AI features. These sources often highlight real-world implications, vulnerabilities, or public reactions, complementing the academic and official narratives with practical observations.

Source selection criteria

All sources were chosen based on credibility, relevance, and recency. Scholarly works and official publications were prioritized to ensure reliable information. When using news or web sources, only reputable outlets or experts were included to avoid unsubstantiated claims. Each source had to directly address Apple Intelligence or closely related AI privacy issues to stay

on-topic. The research also emphasized recent publications (with almost all in the last five years) to capture the current state of technology and policy, given the fast-paced evolution of AI and privacy norms. By applying these criteria, the collected dataset of documents is both trustworthy and directly pertinent to the research question. Additionally, because this is a document-based study, the methodology is transparent and reproducible.

Analytical Framework: Actor-Network Theory (ANT)

The analysis is grounded in the Actor-Network Theory (ANT), which treats both human and non-human entities as "actors" in a socio-technical network (Gutiérrez, 2023). This framework is appropriate because Apple's privacy ecosystem is shaped by interactions among corporate actors, end users, regulatory bodies, and technological infrastructures. ANT enables tracing how design choices made by Apple—such as on-device processing and cloud offloading—translate into trust relationships among these actors.

In applying ANT to Apple Intelligence, the research identifies and examines the key actors and their interrelations in the network that underpins user privacy and trust. These actors include:

- Apple (the Company)—Apple itself (and its engineers and designers) acts as a central node, making design decisions (like on-device processing and the Private Cloud Compute system) and setting policies that affect privacy. Apple's corporate motivations (innovation, user satisfaction, market trust) influence how the technology is implemented and communicated.
- Consumers (Users)—The individuals using iPhones and Apple services are crucial actors whose trust and behavior are part of the network. Their data is at stake, and their

expectations of privacy and trust in Apple influence how the system succeeds or encounters pushbacks. The study considers user perspectives as seen through consumer reports, public discourse, or usability considerations noted in the literature.

- Regulatory Bodies—Governments and privacy regulators (for example, agencies enforcing General Data Protection Regulation (GDPR) or other data protection laws) form another set of actors. They create the policy environment that Apple must navigate. Through ANT, these entities are seen as shaping the network by exerting requirements or pressures on Apple's design (for instance, pushing for transparency or data safeguards) and can either bolster or challenge the trust network (e.g., by investigating or sanctioning privacy issues).

- Technological Infrastructure—The hardware/software system itself, particularly Apple's on-device machine learning modules and the Private Cloud Compute (PCC) back-end, is treated as an actor in ANT terms. The capabilities and limitations of the iPhone's chips, the algorithmic decision of when to offload tasks to the cloud, and the cloud servers' security measures all actively mediate the privacy outcomes. This non-human actor's role is analyzed by looking at technical reports and documentation of how data flows and where vulnerabilities might arise.

Using ANT as the analytical lens, the methodology involves mapping out how these actors interact and "translate" or influence each other's goals. For example, Apple (through its software design) might translate regulatory requirements into technical features, or consumers' trust concerns might translate into changes in Apple's policies. The actor-network perspective ensures that the analysis stays attentive to connections: how a change in the technical

infrastructure (like a new encryption method) could shift user trust, or how a regulatory guideline could lead Apple to alter the AI's data handling.

By systematically examining the collected documents with these actor relationships in mind, the research uncovers the network of influence and dependency that constitutes privacy in Apple's AI. This ANT-guided analysis ultimately helps explain "trust in the machine" as an emergent property of the entire network, rather than just a factor of user attitude or a technical feature alone.

Integration of Methodological Details and Bias Mitigation

To address under-disclosed corporate evaluation protocols, I triangulated Apple's official whitepapers with independent technical analyses: Trail of Bits blog discussions on PCC verification procedures (Travers, 2024); PCMag prototype threat-model assessments (Rubenking, 2024); and Mozilla Foundation reverse-engineering reports on iOS 18 privacy features (Harding, 2024; Green, 2025).

Potential Bias and Source Diversity

Apple's corporate documentation provides valuable perspectives but may not fully disclose underlying evaluation methods. To mitigate bias, I included a diverse mix of sources: academic literature (Carter et al., 2014; P.s. Dr. V., 2023), nonprofit advocacy (Mozilla Foundation, Busch, 2023), independent security audits (Trail of Bits; PCMag), and regulatory policy whitepapers.

Ensuring Credibility and Rigor

The study employed source triangulation, explicit cross-referencing of conflicting claims, and transparent documentation of analytic steps. Each key finding was confirmed by at least two independent sources to enhance validity.

Overall, this qualitative, literature-based approach—underpinned by the Actor-Network Theory and reinforced by rigorous cross-checking—provides a transparent and robust methodology for exploring privacy in Apple's new AI technology. It enables an in-depth understanding of "trust in the machine" by building on credible evidence and a strong analytical foundation, thereby meeting the expectations of an academic STS investigation.

Analysis

Privacy and Trust Implications of Apple's Hybrid AI Architecture

Apple's claims their design for Apple Intelligence delivers enhanced AI capabilities while maintaining user privacy by processing data on the device whenever possible and routing only intensive tasks to the PCC (Apple Security Engineering & Architecture, 2024). According to Apple, any data sent to the PCC is "not accessible to anyone other than the user – not even to Apple" (Apple Inc., 2024). To strength trust, Apple implements stateless cloud processing, end-to-end encryption in transit, and a transparency mechanism verifying software integrity before any connection (Ray 2024; Federighi, 2024). Moreover, Apple has made server software images available for independent scrutiny and offers substantial bug bounties to encourage external testing of the PCC's security (Rubenking, 2024).

Reactions from the Privacy Community

Digital rights advocates and privacy experts have greeted Apple's PCC with cautious optimism. Analysts at the Mozilla Foundation praise Apple's privacy-by-design intent saying Apple "has clearly thought a lot about privacy" in designing Apple Intelligence, however, they caution that some sensitive data will still inevitably be transmitted to its cloud servers for complex tasks (Harding, 2024). As one Mozilla representative observed, "that is a lot of trust in a company that, while better than many, still has commercial interests that may not always align perfectly with user privacy" (Harding, 2024).

In addition, independent security researchers applaud Apple's decision to publish the PCC software images for verification—a move that sets an interesting precedent in an industry where most cloud services remain opaque (Rubenking, 2024). Nevertheless, organizations such as the Electronic Frontier Foundation have emphasized that time and rigorous independent audits are needed to confirm whether these measures can truly uphold user privacy.

Concerns about PCC Transparency and Data Handling

Despite Apple's robust privacy assertions, several concerns remain regarding the PCC's transparency and data handling. Apple does not open-source the entire PCC platform; rather, it publishes compiled software images and employs a public logging mechanism to verify that running code matches the disclosed version (Green, 2024). While this approach is innovative, some experts argue that it "stops short of being truly open source" and complicates independent auditing (Green, 2024; Harding, 2024).

In this setup, researchers must either reverse-engineer the binaries or rely on Apple's cryptographic proofs, which leaves a margin for uncertainty. As cryptography expert Matthew Green has noted, the PCC's security hinges on Apple "getting a bunch of complicated software

and hardware security features right," rather than on an absolute cryptographic guarantee (Green, 2024). This trade-off—offering strong protections without a full open-source review—requires users to continue trusting Apple's internal processes. Moreover, critics have raised questions regarding whether Apple's documentation fully accounts for all data flows (for example, whether telemetry or error logs might inadvertently contain user information) and emphasize that "trust, but verify" is essential (Harding, 2024).

A specific point of much interest in this debate is Apple's integration of OpenAI's technology for certain Siri functions. Although Apple has built many on-device models, it has opted to partner with OpenAI to offer more advanced generative capabilities (French, 2024). This raises concerns about whether such third-party involvement undermines Apple's "no peeking" guarantee. Apple attempts to mitigate this risk by requiring explicit user consent before sending any prompts to ChatGPT and by ensuring that OpenAI does not store prompts or IP addresses (French, 2024).

Nonetheless, some privacy experts remain uneasy about the arrangement, arguing that the involvement of an external actor (OpenAI) introduces an additional layer of risk that requires further transparency and verification (French, 2024). Adding OpenAI integration is simply adding another link in the chain of trust that users must swallow.

Potential Solutions and Mitigation Strategies

To help close the trust gaps identified in Apple's hybrid AI architecture, the following strategies could be employed. Each approach targets a different aspect of the data-privacy lifecycle, from local model training to cloud-side execution and ongoing verification.

1. Federated Learning

   Securely aggregating encrypted model updates on-device can eliminate the need to send raw user data to the cloud. In a federated learning setup, each device trains the model locally on its personal data and only transmits cryptographically protected parameter updates to a central server (McMahan et al., 2016). A global model is then reconstructed by averaging these updates, ensuring that personal data never leaves the user's device. Case studies have demonstrated that federated learning can achieve near-cloud accuracy while reducing exposure of sensitive inputs.

2. Differential Privacy

   Differential privacy adds calibrated statistical noise to data queries or model gradients, providing provable guarantees that the presence or absence of any single user's data cannot be inferred from the aggregate results. When applied to cloud-side computations—such as analytics on PCC logs or aggregated behavioral signals—this technique ensures utility (e.g., usage statistics) while bounding privacy loss (Dwork et al., 2006).

3. Formal Verification and Independent Audits

   Formal verification uses mathematical proofs to ensure that code adheres precisely to its specification, eliminating entire classes of vulnerabilities before deployment. By extending Apple's current "bug bounty" approach to include scheduled, third-party formal audits—conducted by neutral security firms like Trail of Bits—PCC software can gain stronger assurances of correctness. Formal methods have uncovered subtle flaws in enclave implementations at major cloud providers; applying the same rigor here would reinforce Apple's transparency goals.

4. Hardware-Enforced Enclaves

   Secure enclaves—dedicated hardware regions that isolate code and data in tamper-resistant memory—can ensure that sensitive computations occur exactly as designed. On Apple devices, the Secure Enclave already protects biometric data; extending similar enclave technology to PCC servers would bind cloud execution to signed code images, even if an adversary gains root access. Research in IoT and cloud security demonstrates that combining enclave protections with remote attestation protocols can deliver end-to-end guarantees of code integrity and data confidentiality (Shandilya et al., 2018).

Comparison with Google's Approach to AI Privacy

To contextualize Apple's model, it may be helpful to compare it briefly with Google's approach. Both companies recognize that not all AI tasks can be performed entirely on-device; however, their philosophies diverge significantly.

Apple's model emphasizes a privacy-first mindset by processing nearly all tasks locally and ensuring that any necessary cloud processing is ephemeral and transparent (Apple Security Engineering & Architecture, 2024). In contrast, Google retains more data on its cloud servers for purposes such as context retention and service improvement (Kleidermacher & Hogben, 2024). Google explicitly states that its on-device processing is used "for some of your most sensitive tasks" (Kleidermacher & Hogben, 2024), but when data is sent to the cloud, it is managed entirely in-house.

Although Google asserts that its internal infrastructure is secure and uses state-of-the-art encryption, its approach relies on users trusting Google's internal safeguards rather than external auditability. This contrast illustrates two different strategies in reconfiguring the network of trust:

Apple is actively trying to minimize external access to data, whereas Google's approach continues to depend on established internal protocols and user trust in its corporate data stewardship.

Taken together, the evidence suggests that while Apple's privacy-by-design model shows promise, its reliance on the PCC and third-party partnerships leaves some issues unresolved. The privacy community appears cautiously optimistic but insists that independent verification and ongoing audits are essential for Apple to fully substantiate its claims.

From an Actor-Network Theory perspective, the interplay among the various actors—Apple as the technology provider, consumers as data owners, regulatory bodies, and even third-party entities like OpenAI—collectively shapes the emergent trust in Apple's AI ecosystem. Given the novelty of Apple Intelligence and its rolling release style, more time is needed to fully evaluate what the situation is in regard to user data privacy. Future research should therefore continue to monitor these interactions, testing the resilience of Apple's claims over time and assessing whether its model sets a new standard for privacy in AI.

Conclusion

This study found that Apple's hybrid AI architecture—combining on-device processing with a privacy-preserving cloud component (Private Cloud Compute, or PCC)—offers a promising approach to enhancing AI functionality while striving to protect user privacy. The analysis indicates that by processing most data locally and limiting cloud reliance to intensive tasks, Apple can theoretically reduce the exposure of sensitive information, thereby strengthening the trust relationship between users and the technology. At the same time, critical

vulnerabilities remain, especially concerning the proprietary nature of PCC and the added risks from third-party integrations such as OpenAI.

The research aimed to explore how these technical strategies influence trust within the network of actors—users, Apple, regulatory bodies, and external partners—using the Actor-Network Theory (ANT) as an analytical framework. In doing so, it contributes to a deeper understanding of the complex interplay between technological innovation and privacy, underscoring that robust privacy by design is as much a social and organizational challenge as it is a technical one.

Implications of these findings extend to both industry practice and policy. For practitioners, Apple's model suggests that integrating strict on-device processing with carefully controlled cloud functionalities can mitigate some of the traditional privacy risks associated with AI. For policymakers and regulators, the study highlights the need for ongoing, independent verification of such systems to ensure that public assurances match real-world performance. Ultimately, maintaining user trust in AI systems depends on transparent, verifiable processes and the willingness of companies to subject their architectures to external scrutiny.

However, the study is not without limitations. Its reliance on a qualitative synthesis of limited available scholarly and high-quality non-scholarly sources means that some aspects of the emerging Apple Intelligence architecture remain underexplored, particularly given the novelty of the technology. Future research should include longitudinal studies and technical audits to more comprehensively evaluate the long-term resilience of Apple's privacy safeguards.

In closing, this research reinforces the view that trust in advanced AI systems must be continuously negotiated within a complex socio-technical network. Apple's approach—if

validated through sustained independent verification—could set new standards for privacy in AI. Conversely, any failure to uphold these standards would serve as a cautionary tale, emphasizing that even the most innovative designs must be rigorously and transparently evaluated to secure user data in an increasingly interconnected digital world.

References

Apple Intelligence Preview. (n.d.). Apple. Retrieved October 4, 2024, from https://www.apple.com/apple-intelligence/

Apple Platform Security. (n.d.). Apple Support. Retrieved March 18, 2025, from https://support.apple.com/guide/security/welcome/web

Apple Statistics (2024). (n.d.). Business of Apps. Retrieved November 8, 2024, from https://www.businessofapps.com/data/apple-statistics/

Biometric security. (n.d.). Apple Support. Retrieved March 18, 2025, from https://support.apple.com/guide/security/biometric-security-sec067eb0c9e/1/web/1

Bismi, I. (2023, April 9). Federated Learning by Google: Bringing Privacy to Machine Learning. *Medium*. https://medium.com/@iqra.bismi/federated-learning-by-google-bringing-privacy-to-machine-learning-25f7594849eb

Blog - Private Cloud Compute: A new frontier for AI privacy in the cloud - Apple Security Research. (n.d.). Blog - Private Cloud Compute: A New Frontier for AI Privacy in the Cloud - Apple Security Research. Retrieved March 27, 2025, from https://security.apple.com/blog/private-cloud-compute/

Busch, K. E. (2023). *Generative Artificial Intelligence and Data Privacy: A Primer* (Internet materials; CRS Report). Congressional Research Service. https://purl.fdlp.gov/GPO/gpo213803

Carter, N., Bryant-Lukosius, D., DiCenso, A., Blythe, J., & Neville, A. J. (2014). The use of triangulation in qualitative research. *Oncology Nursing Forum*, *41*(5), 545–547. https://doi.org/10.1188/14.ONF.545-547

Costan, V., & Devadas, S. (2016). *Intel SGX Explained* (No. 2016/086). Cryptology ePrint Archive. https://eprint.iacr.org/2016/086

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating Noise to Sensitivity in Private Data Analysis. In S. Halevi & T. Rabin (Eds.), *Theory of Cryptography* (pp. 265–284). Springer. https://doi.org/10.1007/11681878_14

French, L. (2024, June 13). *Apple's AI debut: Privacy nightmare or FUD?* SC Media. https://www.scworld.com/news/apples-ai-debut-privacy-nightmare-or-fud

Green, M. (2025, March 1). A Few Thoughts on Cryptographic Engineering. *A Few Thoughts on Cryptographic Engineering*. http://matthewdgreen.wordpress.com

Gutiérrez, J. L. M. (2024). On actor-network theory and algorithms: ChatGPT and the new power relationships in the age of AI. *AI and Ethics*, *4*(4), 1071–1084. https://doi.org/10.1007/s43681-023-00314-4

Harding, X. (2024, July 16). *Using Apple's AI Tools? Here's What Features To Expect In iOS 18 Beta And Full Release*. Mozilla Foundation. https://foundation.mozilla.org/en/blog/apple-ai-ios-18/

Hawblitzel, C., Howell, J., Lorch, J. R., Narayan, A., & Zhang, D. (n.d.). *Ironclad Apps: End-to-End Security via Automated Full-System Verification*.

IEEE Electronic Library (IEL) Conference Proceedings. (2021). *2021 International Conference on Artificial Intelligence for Cyber Security Systems and Privacy (AI-CSP)* (Internet materials). IEEE. http://RE5QY4SB7X.search.serialssolutions.com/?V=1.0&L=RE5QY4SB7X&S=JCs&C=TC_046945499&T=marc

IEEE Staff & IEEE Electronic Library (IEL) Conference Proceedings. (2014). *2014 International Conference on Privacy and Security in Mobile Systems (PRISMS)* (Internet materials). IEEE. http://RE5QY4SB7X.search.serialssolutions.com/?V=1.0&L=RE5QY4SB7X&S=JCs&C=TC0001774056&T=marc

iPhone 16 and AI: Everything you missed from Apple's "Glowtime" event. (2024, September 10). Euronews. https://www.euronews.com/next/2024/09/10/new-iphone16-and-apple-intelligence-what-you-missed-from-the-glowtime-event

Jaafar, F. (editor), & Pierre, S. (editor). (2024). *Blockchain and Artificial Intelligence-based Solution to Enhance the Privacy in Digital Identity and IoT* (Internet materials). CRC Press. https://proxy1.library.virginia.edu/login?url=https://www.taylorfrancis.com/books/9781003227656

Kleidermacher, D., & Hogben, G. (n.d.). Private AI For All: Our End-To-End Approach to AI Privacy on Android. *Google Online Security Blog*. Retrieved March 27, 2025, from https://security.googleblog.com/2024/08/android-private-ai-approach.html

Langenfeld, J. (editor), Fagan, F. (editor), & Clark, S. (Data scientist) (editor). (2022). *The Law and Economics of Privacy, Personal Data, Artificial Intelligence, and Incomplete Monitoring* (1st Floor Stacks). Emerald Publishing.

Latour, B. (1996). On actor-network theory: A few clarifications. *Soziale Welt*, *47*(4), 369–381.

Mayo, B. (2024, June 10). *Apple announces "Apple Intelligence": Personal AI models across iPhone, iPad and Mac*. 9to5Mac. https://9to5mac.com/2024/06/10/apple-ai-apple-intelligence-iphone-ipad-mac/

McMahan, H. B., Moore, E., Ramage, D., & Hampson, S. (2016). *Communication-Efficient Learning of Deep Networks from Decentralized Data*.

Moroney, L. & O'Reilly Online Learning: Academic/Public Library Edition. (2021). *AI and Machine Learning for On-device Development: A Programmer's Guide* (Internet materials). O'Reilly Media. http://RE5QY4SB7X.search.serialssolutions.com/?V=1.0&L=RE5QY4SB7X&S=JCs&C=TC0002686552&T=marc

Nasir, N. (2024). *Untangling the Cloud From Edge Computing for the Internet of Things* [University of Virginia, Computer Science - School of Engineering and Applied Science, PHD (Doctor of Philosophy), 2024]. https://doi.org/10.18130/ya3p-vn55

Newman, L. H. (n.d.). Apple Intelligence Promises Better AI Privacy. Here's How It Actually Works. *Wired*. Retrieved March 27, 2025, from https://www.wired.com/story/apple-private-cloud-compute-ai/

P. s. , Dr. V. (2023). How can we manage biases in artificial intelligence systems – A systematic literature review. *International Journal of Information Management Data Insights*, *3*(1), 100165. https://doi.org/10.1016/j.jjimei.2023.100165

Ray, T. (n.d.). *Here's how Apple's keeping your cloud-processed AI data safe (and why it matters)*. ZDNET. Retrieved March 27, 2025, from https://www.zdnet.com/article/heres-how-apples-keeping-your-cloud-processed-ai-data-safe-and-why-it-matters/

Rubenking, N. (2024, June 11). *Does Apple Intelligence Protect Your Privacy?* PCMAG. https://www.pcmag.com/news/does-apple-intelligence-protect-your-privacy

Secure Enclave. (n.d.). Apple Support. Retrieved March 18, 2025, from https://support.apple.com/guide/security/secure-enclave-sec59b0b31ff/1/web/1

Shandilya, S. K., Chun, S. A., Shandilya, S., & Weippl, E. (2018). *Internet of Things Security: Fundamentals, Techniques and Applications* (Internet materials). River Publishers. https://proxy1.library.virginia.edu/login?url=https://www.taylorfrancis.com/books/9781003338642

Sharma, N. (Computer scientist) (editor), Srivastava, D. (Computer scientist) (editor), & Sinwar, D. (editor). (2024). *Artificial Intelligence Technology in Healthcare: Security and Privacy Issues* (Internet materials). CRC Press.

https://proxy1.library.virginia.edu/login?url=https://www.taylorfrancis.com/books/9781003377818

Travers, A. (2024, June 14). *PCC: Bold step forward, not without flaws*. The Trail of Bits Blog. https://blog.trailofbits.com/2024/06/14/pcc-bold-step-forward-not-without-flaws/

Vaidya, J., Gabbouj, M., & Li, J. (Eds.). (2024). *Artificial Intelligence Security and Privacy: First International Conference on Artificial Intelligence Security and Privacy, AIS&P 2023, Guangzhou, China, December 3–5, 2023, Proceedings, Part I* (Vol. 14509). Springer Nature. https://doi.org/10.1007/978-981-99-9785-5

Xu. (n.d.). *Advances in private training for production on-device language models*. Retrieved March 27, 2025, from https://research.google/blog/advances-in-private-training-for-production-on-device-language-models/