Artificial Intelligence, Machine Learning, & Robotics in Business

FROM CONCIERGE TO CODE: ARCHITECTING GUEST TRUST VIA PRIVACY IN AN AI WORLD

Drew Bjerken, DBA

***Correspondence: a.bjerken@vikes.csuohio.edu Data Privacy Global Privacy Office**

Abstract: This editorial explores a strategic vision for the integration of artificial intelligence (AI), machine learning (ML), and robotics in privacy and cybersecurity practices across the global hospitality industry. Drawing on multidisciplinary sources, cutting-edge academic theory, and applied business strategy, this article argues that these technologies represent a paradigm shift in data protection. The essay contextualizes the risks and rewards of advanced technologies and calls for the design of trustworthy digital infrastructures embedded with ethics, human oversight, and regulatory intelligence. This piece outlines a roadmap for industry transformation that respects hospitality's humanistic core while embracing algorithmic innovation.

Keywords: Smart Infrastructure, Cybersecurity, Ethics, Machine Learning

1 Introduction: A Historical Context for a Digital Transformation

The hospitality industry, historically grounded in personal service and high-touch interactions, is poised to undergo an epochal transformation. As the hospitality industry strives to better understand its clientele in a digital age they have transitioned into the business of trust. At its core, this timeless industry is about service, safety, and personal attention; values that seem at first glance, far removed from the cold logic of machine algorithms. Just as electrification, telephony, and the internet once redefined service models and guest engagement, the advent of artificial intelligence (AI), machine learning (ML), and robotics will reformulate the very architecture of hospitality by augmenting not only front-line operations but back-end processes as well, especially those concerning the protection of guest data. Unlike previous industrial revolutions, this moment demands new tools and new epistemologies, a rethinking of data custodianship in an age of algorithmic intermediation.

Global hospitality chains and boutique operations alike sit at the intersection of data abundance and human vulnerability. Booking systems, loyalty platforms, guest biometrics, and in-room IoT environments create continuous data exchanges. Against this backdrop, the ethical, legal, and technical imperatives of privacy cannot remain static. They must evolve alongside the technologies they aim to govern.

2. Data-Centricity and the Algorithmic Shift in Hospitality Risk Models

The hospitality industry is no longer merely customer-centric, it is profoundly data-centric. Every transaction, every behavioral cue, and every digital interaction becomes a data point: something to be categorized, analyzed, compared, monetized, and, most importantly, secured. This shift reflects not just an operational transformation, but also a philosophical one. Where hotels once competed on service quality alone, they now compete on their ability to manage the digital lives of their guests while striving to build trust and loyalty through transparency.

AI and ML technologies provide unprecedented capabilities across multiple domains, pattern recognition, behavioral inference, threat modeling, predictive analytics, and real-time compliance monitoring. These tools allow for the creation of dynamically responsive data protection environments, where systems adapt in real-time to evolving risk profiles and guest behavior patterns. For example, ML algorithms can autonomously identify login anomalies across interconnected hotel systems, automatically escalate alerts, and trigger protective protocols without human intervention. Likewise, natural language processing (NLP) enables multilingual handling of data subject requests (DSRs) in real-time, an especially vital capability for global hotel chains dealing with multilingual and multicultural guest populations.

More transformative are federated learning models, which allow hospitality chains to develop sophisticated AI insights across properties and jurisdictions without ever aggregating raw guest data into a centralized system, dramatically reducing cross-border data risks. But these gains come with sharp edges. The same opacity that makes AI "intelligent" also makes it opaque. The so-called "black box" nature of complex algorithms where inputs cannot easily be linked to outputs poses serious challenges for accountability, auditability, and guest redress mechanisms. Doshi-Velez and Kim (2017) characterize this as a fundamental barrier to interpretability in high-stakes applications like privacy and cybersecurity. This is where governance must evolve in lockstep with engineering. As Brundage et al. (2018) emphasize, AI's dual-use nature makes it a tool of both resilience and exploitation. The hospitality industry must, therefore, apply a layered governance approach rooted in risk-based assessment frameworks, stakeholder impact models, and legal auditability.

Trust is the currency of loyalty in hospitality, and to maintain guest trust, privacy and cybersecurity can no longer remain siloed functions or reactive stances. They must become architectural principles infused throughout the design, implementation, and continuous optimization of intelligent systems. Floridi et al. (2018) underscore the importance of explicability as an ethical imperative especially in environments like hotels where guests surrender deeply personal data with the expectation of discretion and confidentiality. In this emerging landscape, the competitive edge will belong not just to those who use data effectively, but to those who can explain how, why, and to what end that data was used. Ethical AI is no longer optional; it is the foundation of reputational resilience.

3. Robotics and Smart Infrastructure: Autonomous Agents and Ambient Risk

In the early 2020s, robotics in hospitality shifted from futuristic novelty to operational necessity. Smart concierges, autonomous cleaning units, robotic bellhops, and AI-integrated kiosks are now routine fixtures in upscale hotel chains and tech-forward resorts. Initially adopted in response to post-pandemic hygiene expectations and labor shortages, these robotic systems now play vital roles in sustaining efficiency, standardizing service delivery, and managing escalating guest expectations.

Yet, these physical embodiments of automation are far more than functional endpoints. They are intelligent nodes within vast and increasingly complex digital ecosystems. Each robot is not just a mechanical service agent; it is a sensor, a data processor, and often, a transmitter. This makes them uniquely vulnerable to a spectrum of risks: unauthorized access, malware injection, firmware tampering, and data exfiltration. As the number of network-connected robotic devices increases, so too does the digital surface area susceptible to attack.

Overlaying this robotic layer is the rise of ambient computing, the seamless integration of intelligent systems into physical environments that anticipate and respond to human presence. In-room sensors, smart thermostats, voice assistants, occupancy detectors, and even smart mirrors collect a constant stream of behavioral data designed to tailor guest experiences in real time. The goal is fluid, frictionless service but the result is a sensor-rich environment that continuously collects, analyzes, and stores highly sensitive behavioral and biometric data.

This convergence of robotics and ambient AI infrastructure raises profound privacy implications. These systems blur the lines between service and surveillance, between personalization and profiling. As a result, the hospitality industry must implement rigorous safeguards: A) Embedded security-by-design principles must be implemented at the firmware and hardware layers, ensuring resilience against both digital intrusions and physical tampering; B) Dynamic, context-aware consent mechanisms should adapt data collection permissions based on real-time guest behaviors and preferences; for example, disabling voice assistants when minors are present or allowing room-level data control; and C) Digital sovereignty safeguards must align local processing requirements with regional data transfer laws, particularly in multinational hospitality conglomerates with cloud-based systems spanning continents.

This prompts an urgent discussion around AI agency, namely, what decisions should be delegated to machines and which require human involvement. For example, robotic security systems equipped with facial recognition may autonomously restrict access to high-risk areas. Automated dispute resolution bots might issue refunds or deny requests without human validation. These developments require robust accountability frameworks to ensure such automation adheres to due process, fairness, and the procedural rights of both guests and employees.

Ultimately, the hospitality sector must evolve toward robotic governance models that mirror those developed for critical infrastructure: risk classification, AI model certification, post-deployment auditing, and ethics review boards. In an industry where trust is inseparable from reputation, the intelligent systems that operate behind the scenes must be just as hospitable, ethical, and accountable as the staff they increasingly augment or replace.

4. Regulating the Irregular: Navigating the Legal Fragmentation Of AI in Hospitality

The regulatory landscape surrounding AI, privacy, and cybersecurity remains deeply fragmented posing a significant challenge for hospitality enterprises operating across borders. While some jurisdictions have embraced a rights-based framework centered on data subject empowerment, others maintain minimalist, sector-specific, or even contradictory regulatory models. This global asymmetry complicates the development and deployment of AI systems, especially those that traverse physical and jurisdictional boundaries inherent to global hotel chains.

The European Union's GDPR stands as the gold standard, mandating Data Protection Impact Assessments (DPIAs) for high-risk automated processing (Article 35) and articulating clear rights to explanation, rectification, and erasure. The California Consumer Privacy Rights Act (CPRA) provides opt-out rights for automated profiling and targeted advertising, reinforcing the primacy of consumer choice in algorithmic contexts. Meanwhile, China's Personal Information Protection Law (PIPL) mandates explicit consent and imposes strict data localization and cross-border transfer conditions, signaling a growing emphasis on digital sovereignty and state-centric data ethics.

In this increasingly complex regulatory environment, the hospitality sector must pivot toward a model of embedded regulatory intelligence. This entails proactively integrating jurisdiction-specific requirements into every phase of the AI lifecycle from design and data acquisition to training, deployment, and decommissioning. The emerging concept of compliance-as-code exemplifies this direction: encoding legal constraints, audit triggers, and accountability mechanisms directly into the logic of AI models. This approach, however, necessitates a new interdisciplinary cadre of professionals fluent in both law and software architecture; what some are now calling "legal engineers."

For AI to be deployed responsibly at scale in hospitality, organizations must build internal capacity to test and validate regulatory alignment before systems go live. To that end, industry leaders should consider establishing: A) AI Law Clinics or Regulatory Sandboxes: Controlled environments where emerging AI systems can be tested in compliance with data protection laws, ethical benchmarks, and guest-centric standards; B) Global Algorithm Registries: Repositories of deployed models annotated with metadata concerning their purpose, version history, accuracy benchmarks, training datasets, and fairness evaluations;, and C) Sector-Specific Privacy Assurance Schemas: Frameworks modeled on recognized standards such as SOC 2, ISO/IEC 27001/27701, or NIST RMF, but tailored for AI-specific governance, such as model explainability, data minimization, and bias mitigation.

Rather than treating compliance as a post-deployment audit concern, forward-thinking hospitality companies must embed legal fidelity and ethical guardrails directly into the codebase of intelligent systems. This not only protects guest rights and meets regulatory expectations but also reinforces brand equity in an era when data responsibility is inseparable from corporate reputation.

5. Strategic Privacy Architectures: From Static Controls to Adaptive Governance

Traditional privacy frameworks were built for static environments characterized by predictable workflows, fixed compliance checklists, and perimeter-based defenses. In that context, privacy policies were largely declarative: codified on paper, enforced through audit, and updated on an annual cycle. But AI-infused systems operate in a different reality. They learn, they evolve, and they respond dynamically to real-time inputs, including guest behavior, environmental data, and emerging threat vectors. In this setting, adaptive governance becomes not a luxury, but a necessity.

Privacy in the AI era must transcend policy-based abstraction and embrace engineering-based design. What was once “Privacy by Policy” must become “Privacy by Architecture.” Periodic audits must give way to continuous, embedded compliance systems. Risk assessment must shift from static PIA forms to AI-assisted Risk Engines capable of evaluating model behavior, training data lineage, and user impact at runtime.

This transition demands a structural shift in how hospitality organizations govern the lifecycle of data and the systems that process it. Techniques such as differential privacy and homomorphic encryption are emerging as the mathematical vanguards of privacy-preserving data science. Differential privacy enables aggregate analysis while masking individual identities, and homomorphic encryption allows computation on encrypted data without ever exposing its raw form. These tools are not turnkey; they require enterprise-wide architectural alignment, algorithmic transparency, and deep collaboration between privacy officers, data scientists, and legal teams.

To operationalize this vision, a future-ready privacy program in hospitality must be architected around the following foundational elements:

- Algorithmic Transparency: All data collection and decision-making algorithms must be disclosed in terms and conditions in language that is accessible and actionable. Black-box reasoning should be treated as unacceptable in environments where service depends on guest trust.

- Dynamic Data Classification Systems: These systems must automatically label and categorize guest data in real time based on context, usage purpose, sensitivity level, and applicable regulatory frameworks. Metadata should travel with the data, enforcing retention rules and access controls dynamically across systems.
- Interpretable AI Toolkits: All algorithms used in personalization, fraud detection, or surveillance contexts should be subject to mandatory explainability standards. Black-box models should be discouraged in high-stakes decisions, and tools like LIME or SHAP should be integrated into model testing pipelines.
- Granular Guest Control: Empower guests to make real-time, contextual choices about what is collected, when, and by whom. This includes the ability to disable or opt out of facial recognition, location tracking, and behavioral profiling without losing access to core services.
- Participatory Ethics Governance: Prior to deploying AI systems in sensitive spaces (e.g., guest rooms, wellness centers, or conference halls), properties should convene ethics panels that include internal stakeholders and independent experts. Where appropriate, community consultations should be conducted to understand and reflect regional privacy expectations and cultural sensitivities.

This adaptive model of privacy governance is not theoretical. It is already emerging in leading-edge institutions. What remains is the widespread adoption of a mindset that sees privacy not as an obstacle to innovation, but as its foundation. When done well, privacy architecture becomes a source of competitive advantage strengthening brand trust, enhancing user control, and mitigating long-tail regulatory risk in an algorithmically mediated world.

6. Conclusion: Building a Hospitable Digital Future

AI, ML, and robotics are not merely technological enhancements; they are structural forces capable of reshaping the foundational values of industries. Nowhere is this more acute than in hospitality, a sector rooted in human connection, built on trust, and sustained through discretion. As digital transformation accelerates, the imperative is clear: the industry must define its own terms of engagement with intelligent technologies before those technologies define the industry.

This editorial has charted a vision that integrates advanced technology with ethical stewardship. We have explored how algorithmic systems can enhance privacy by architecture, how robotics must be treated as embedded nodes in cyber-physical environments, and how compliance must evolve from documentation to embedded governance. We've argued that ethical AI deployment requires participatory governance, that future cybersecurity must model threats that have yet to materialize, and that workforce transformation must reflect the interdisciplinary nature of modern privacy leadership.

In each domain, the same tension emerges between innovation and restraint, between scale and sensitivity. The hospitality industry stands uniquely poised to show the world how AI can serve—not replace—human care. But this demands action: A) To build intelligent systems that are not just efficient, but explainable; B) To architect privacy not as an afterthought, but as infrastructure; and C) To train a new generation of professionals who are fluent in both empathy and algorithms.

We stand at a historical inflection point. The decisions made in boardrooms, product labs, and regulatory agencies today will shape the ethical terrain of hospitality for decades to come. Will AI be a tool of surveillance or a mechanism of empowerment? Will robotics depersonalize service, or extend the reach of personalized care? Let us choose wisely. Let us lead deliberately. Let us commit to a digital hospitality that amplifies the virtues of our analog past. Because in the end, the most advanced system we can build is one that protects what matters most: the humanity of the guest experience.

REFERENCES

Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., ... & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. arXiv preprint arXiv:1802.07228.

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.

European Commission. (2016). *General Data Protection Regulation (GDPR)*. https://eur-lex.europa.eu/eli/reg/2016/679/oj

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707.

Raji, I. D., Smart, A., White, R., Mitchell, M., Gebru, T., Hutchinson, B., ... & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.

Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2), 76–99.

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.

Dr. Andrew Bjerken, FIP, CISSP, CIPP/E, CIPM, CIPT, C/CISO. Andrew "Dr. Drew" Bjerken is an award-winning privacy and cybersecurity executive with over two decades of experience at the intersections of law, technology, cybersecurity, and risk. He brings 25+ years of experience in privacy, cybersecurity, and intelligence. He currently serves as the Global Chief Privacy Officer for a leading hospitality company, where he leads global privacy, data governance, and AI governance programs across a portfolio of brands spanning more than 90 countries. In this role, Drew is responsible for enterprise-wide privacy strategy, regulatory compliance, privacy operations, privacy architecture and engineering, and the operationalization of data ethics and responsible AI practices. He currently serves on IAPPs Diversity in Privacy Advisory board and was a founding member of The Privacy Board. Previously, Drew served as Chief Information Security & Privacy Officer within the marketing and academic sectors where he led enterprise cybersecurity and privacy programs. He also served as Director of Cyber Threat Intelligence at a Fortune 200 company, overseeing threat management and fraud response operations. Drew is a proud U.S. Air Force veteran and former cyber operations commander; he began his career flying information warfare missions before leading elite red team operations that simulated advanced threats against Department of Defense networks and facilities. He commanded the joint red team for the inaugural Cyber Flag exercise, involving NSA, DOJ, DHS, academia, and all military branches. Drew has presented as a subject matter expert in privacy, AI governance, cybersecurity, and information warfare at RSA, Black Hat, IAPP, and other venues. He holds a B.A. in Criminal Justice, a master's in forensic science, and a Doctorate in Business Administration with a specialization in Computers & Information Security. He hold multiple certifications from IAPP, ISC2, ISACA, and Northwestern University Pritzker School of Law.