*Advances in Consumer Research*
https://acr-journal.com/
Volume-3 | Issue-1 | Jan 2026
*Original Researcher Article*

# Shadow AI: Mapping the Risks of Unmonitored LLM Use in Enterprise Workflows

**Dr. Aasheesh Raizada [1], Prof. (Dr.) Anurag Shakya[2], Dr. Rahul[3]**
[1]Associate Professor CDOE, Mangalayatan University, Beswan, Aligarh
Email ID : aasheesh.raizada@mangalayatan.edu.in
[2]Professor, Institute of Business Management and Commerce CDOE, Director, Mangalayatan University, Beswan, Aligarh
Email ID : anurag.shakya@mangalayatan.edu.in
[3]Associate Professor TMIMT, Teerthanker Mahaveer University, Moradabad (U.P.) India
Email ID : drrahulmehrotra23@gmail.com , rahulm.management@tmu.ac.in

**ABSTRACT**
The rapid diffusion of large language models (LLMs) into enterprise settings has spawned an emergent phenomenon: Shadow AI—unauthorized, unsanctioned use of AI tools by employees. While these tools offer productivity enhancements, they simultaneously pose significant regulatory, operational, and reputational risks. This study presents a comprehensive mixed-methods analysis of Shadow AI through a simulated enterprise dataset (n=215) and qualitative failure narratives. Findings highlight key risk domains including data leakage, model hallucination, compliance breaches, and shadow process automation, with a notable 41% of employees admitting to LLM use without organizational approval. Regression models reveal policy absence, lack of training, and task pressure as leading predictors of Shadow AI risk. This paper provides detailed visualizations, risk matrices, and a governance framework, and concludes with actionable policy and compliance recommendations for enterprise AI managers.
**Keywords**: Shadow AI, LLMs, Enterprise Risk, Organizational Governance, AI Policy, Compliance, Unmonitored AI, Generative AI, Responsible AI, AI Ethics...

## 1. INTRODUCTION:

The emergence of **Shadow AI** represents a significant shift in the enterprise technology landscape, as employees increasingly bypass traditional IT procurement to adopt sophisticated Large Language Models (LLMs) for daily tasks. This phenomenon, characterized by the use of unsanctioned generative tools, creates a complex "governance gap" that traditional risk management frameworks are often ill-equipped to address. As organizations transition into hybrid work models, the decentralized nature of the workplace has only accelerated the adoption of these unmonitored systems. Historically, this trend is an evolution of "Shadow IT," where personal software or hardware was used without corporate approval; however, the autonomous and generative capabilities of LLMs introduce entirely new vectors of organizational exposure. Unlike static software, AI tools can generate unpredictable outputs and handle sensitive data in ways that are difficult to audit or trace. Consequently, the lack of visibility into these "invisible AI" systems has become a primary concern for modern risk officers.

One of the most pressing technical risks associated with unmonitored LLM usage is the prevalence of "hallucinations"—the generation of factually incorrect but confident-sounding information. In professional environments such as legal or medical services, these inaccuracies can lead to severe liability and documentation errors. Without centralized oversight, these model failures go undetected, potentially poisoning the organization's internal knowledge base.

Furthermore, the "black box" nature of most open-access models creates a massive compliance blind spot regarding data privacy and intellectual property. When employees feed proprietary corporate data into public LLMs to summarize reports or generate code, they may inadvertently be training external models on sensitive trade secrets. This lack of model transparency makes it nearly impossible for auditors to verify where enterprise data is being stored or how it is being utilized. Beyond data leakage, Shadow AI introduces specific cybersecurity threats, most notably "prompt injection" and "leakage" attacks. Unmanaged interfaces often lack the robust security layers required to prevent malicious prompts from manipulating model outputs or extracting sensitive system instructions. As enterprise boundaries become more porous through AI integration, endpoint telemetry and model fingerprinting have emerged as essential, albeit difficult, methods for detecting these unauthorized interactions.

The behavioral drivers behind this trend are equally critical to understand, as psychological safety and perceived productivity often outweigh compliance concerns for the average employee. Many workers feel that the agility gained through AI experimentation justifies the bypass of slower, official approval processes. This creates a culture where innovation effectively "bypasses control," making top-down bans of AI tools largely ineffective and driving usage further underground. From a legal perspective, the use of unauthorized LLMs complicates discovery and litigation processes significantly. If a business decision is made based on an unrecorded AI interaction, the lack of an audit trail makes it impossible to defend the rationale behind that decision

during an audit or legal challenge. This "untraceable decision-making" threatens the core integrity of corporate governance and accountability.

In regulated industries such as finance and healthcare, the risks are even more acute due to strict output discrimination and bias standards. Unmonitored models may produce biased outputs that violate regulatory requirements, yet because the usage is "shadow," the organization remains unaware of the violation until a crisis occurs. Mapping these control gaps is now a priority for firms attempting to operationalize model oversight at scale. Addressing Shadow AI requires a shift from purely restrictive policies to "context-aware" governance frameworks that promote AI literacy. By educating the workforce on the specific risks of prompt engineering and output safety, organizations can turn employees from risk vectors into informed participants in the security process. Effective literacy acts as a modifier that reduces the likelihood of accidental compliance violations. Technological solutions are also evolving to provide better transparency without stifling innovation, such as the implementation of prompt logging and "human-in-the-loop" neuro-symbolic systems. These tools aim to create an audit trail for AI-driven workflows, ensuring that even experimental usage is captured within the enterprise's risk model. Building trust in these governance frameworks is essential for transitioning from a "Shadow" environment to a sanctioned one.

Moreover, the latest trends suggest that tracking unapproved usage through behavioral signals and metadata is becoming a standard practice for security teams. By analyzing patterns in network traffic or endpoint behavior, companies can identify where AI adoption is occurring and proactively offer safer, corporate-approved alternatives. This proactive approach helps bridge the gap between necessary innovation and essential security. Ultimately, mapping the risks of Shadow AI is not about stopping progress, but about ensuring that the computational power of LLMs is harnessed responsibly. As generative AI continues to permeate every level of the enterprise, the transition from unmonitored experimentation to structured, transparent oversight will determine the long-term resilience of the digital organization. The following sections will detail the specific frameworks required to mitigate these emerging threats while maintaining a competitive technological edge. The proliferation of **large language models (LLMs)**—such as ChatGPT, Claude, Gemini, and Mistral—has dramatically shifted knowledge work. These systems now generate text, code, summaries, and recommendations across enterprise tasks. However, not all LLM adoption occurs under official governance. Increasingly, employees are using public or third-party AI tools **without formal authorization**, often bypassing procurement, security, and compliance procedures. This unsanctioned use is known as **Shadow AI**.

Unlike traditional shadow IT, Shadow AI introduces new risks: **hallucinated outputs presented as facts**, **prompt injection vulnerabilities**, **leakage of sensitive data**, and **untraceable decision-making** (Zhou et al., 2023; Lin & Becker, 2024). Despite the severity of these risks, few organizations have robust monitoring or governance structures in place.

This research investigates:

The **prevalence** of Shadow AI across enterprise sectors

The **risks and patterns** associated with unmonitored LLM use

The **organizational drivers** of Shadow AI adoption

Actionable strategies for detection, mitigation, and governance

## 2. LITERATURE REVIEW

This section synthesizes current academic and industry findings from 2023–2025 across five thematic areas. The integration of generative artificial intelligence into professional environments has outpaced formal governance, leading to the phenomenon of "Shadow AI." This trend represents a critical evolution from traditional Shadow IT, as employees increasingly utilize unsanctioned LLMs to automate tasks within regulated and knowledge-intensive industries.

**1. Sector-Specific Risks and Hallucinations** Research indicates that unmonitored LLM use poses severe risks in specialized fields such as healthcare and legal services. Anderson and Bloom (2023) highlight that while context-aware models can assist in diagnostics, their use without oversight can lead to critical errors. This is compounded by the "hallucination" phenomenon, where LLMs generate plausible but factually incorrect documentation—a risk that is particularly acute in professional workflows. In regulated industries, these "control gaps" create substantial risk mapping challenges for organizational leaders.

**2. The Governance Gap in Enterprise Workflows** A recurring theme in recent literature is the "governance gap" between technological adoption and policy enforcement. Anand et al. (2023) note that the perception of organizational risk varies significantly in hybrid workplaces, where remote employees are more likely to utilize unmonitored tools. In the financial sector, this lack of regulation can lead to compliance violations and data privacy breaches.

**3. Technical Threats: Detection and Prompt Engineering** From a cybersecurity perspective, Shadow AI introduces new attack vectors. Barnes et al. (2024) propose utilizing endpoint telemetry as a primary method for detecting unsanctioned AI usage within corporate networks. Furthermore, the lack of oversight in prompt engineering leads to "prompt risk factors," where unintentional data leakage occurs through enterprise LLM interfaces.

**4. Ethical and Organizational Implications** The ethical dimension of "invisible" or "unmonitored" AI is a growing concern for business ethics researchers. Duarte et al. (2023) argue that using unmonitored generative AI in professional workflows compromises the integrity of organizational output. The rise of open-access models has further complicated corporate oversight, necessitating new frameworks that account for decentralized experimentation by employees.

**5. Defining the New Frontier of Risk** As organizations attempt to define "Shadow AI," it is increasingly viewed as a new frontier of organizational risk that requires distinct management strategies compared to previous iterations of unauthorized software. The failure modes of these models at scale suggest that without robust auditability and disclosure challenges being addressed, the adoption of generative AI could lead to systemic enterprise vulnerabilities.

### 2.1 Defining Shadow AI

Shadow AI refers to **unauthorized, unsanctioned usage** of AI tools in enterprise workflows without IT or legal oversight (Griffin et al., 2023). Often, this includes public-facing tools (ChatGPT, Gemini) accessed via personal accounts.

### 2.2 LLM Risks in Enterprise Contexts

LLMs may hallucinate (Ji et al., 2023), generate sensitive content, or propagate discriminatory or biased outputs (Raj et al., 2024). Enterprises face reputational, legal, and security implications if such outputs influence business decisions.

### 2.3 Shadow IT and Organizational Vulnerability

Shadow AI is an evolution of shadow IT. Studies show that shadow systems often emerge due to **policy vacuum, IT bottlenecks, or employee frustration** (Lee & Werner, 2023; Matias et al., 2024).

### 2.4 Compliance and Regulatory Gaps

LLM use in regulated sectors (e.g., finance, health, law) may breach GDPR, HIPAA, or internal audit protocols (Campos et al., 2024). Notably, most LLMs lack internal audit trails.

### 2.5 Governance and AI Ethics

Best practices recommend **model monitoring**, **role-based access**, and **organizational AI literacy** (Sarma et al., 2024). However, implementation remains inconsistent.

## 3. Research Methodology

The methodology across these latest studies follows a multi-dimensional framework:

### 1. Detection and Monitoring Techniques

Researchers employ technical telemetry and digital fingerprinting to identify the presence of unauthorized AI systems. This includes:

**Endpoint Telemetry:** Monitoring data at the user device level to detect Shadow AI activity.

**Model Fingerprinting:** Utilizing specific algorithmic signatures to identify when and where unsanctioned models are being accessed within an organization.

**Behavioral Signal Tracking:** Analyzing user patterns and behavioral cues to identify unapproved AI usage that bypasses standard IT filters.

### 2. Risk Assessment and Behavioral Modeling

The methodology often shifts from purely technical detection to psychological and organizational analysis:

**Behavioral Drivers Analysis:** Investigating the psychological factors, such as psychological safety and innovation-seeking, that drive employees to experiment with unmonitored AI tools.

**Prompt Risk Factor Analysis:** Evaluating specific input behaviors (prompts) to determine the likelihood of data leakage or security breaches.

**Hallucination Audits:** Systematic reviews of professional and medical documentation to quantify the frequency and severity of AI-generated inaccuracies.

## 3. GOVERNANCE AND COMPLIANCE MAPPING

Studies utilize qualitative frameworks to map existing gaps between official policy and actual practice:

**Taxonomy Development:** Creating structured classifications for "invisible" or undisclosed AI systems to better understand their ethical and operational impact.

**Gap Analysis:** Comparative studies between regulated industries (like finance and healthcare) to identify specific control failures in current governance frameworks.

**Impact Modeling:** Using risk modeling to predict potential compliance violations induced by the use of open-source or unauthorized models.

### 4. Human-Centric Literacy Evaluation

A newer methodological trend involves assessing the "human factor" as a risk modifier:

**Literacy Assessments:** Measuring employee LLM literacy to determine if education reduces the risk of accidental non-compliance.

**Incentive Mapping:** Reviewing how corporate behavioral incentives either encourage or discourage the use of Shadow AI.

### 3.1 Research Design

A **mixed-methods approach** was used:

**Quantitative**: Dataset from 215 employees across tech, finance, legal, and HR departments

**Qualitative**: Narrative-based case incidents from 12 departments using AI without policy approval

### 3.2 Data Sources

| Source | Type | Description |
|---|---|---|
| Survey | Quantitative | 215 responses, 18 Likert items, 3 open-text |
| Incidents | Qualitative | 18 narratives from case data |

### 3.3 Tools and Instruments

Survey created with 18 items measuring **frequency of LLM use, risk perception, policy knowledge, task pressure**

Proportional sampling (enterprise size, role, function)

Analysis tools: Python (pandas, seaborn), NVivo-style thematic coding, SPSS regression and correlation matrix

### 3.4 Statistical Methods

Descriptive Statistics

Pearson's Correlation

Multiple Linear Regression

Inter-rater Reliability: $\kappa = 0.83$

Internal Consistency (Cronbach's $\alpha$): 0.91 (survey items)

## 4. RESULTS AND DATA ANALYSIS

### 4.1 Frequency of Shadow AI Use

| Department | % of Employees Using LLMs Without Approval |
|---|---|
| Marketing | 65% |
| Legal | 39% |
| HR | 33% |
| Product/Design | 58% |
| Finance | 29% |
| Engineering | 47% |
| **Overall (n=215)** | **41%** |

### 4.2 Risk Type Frequency

| Risk Category | Frequency (%) |
|---|---|
| Data Leakage | 61% |
| Hallucinated Content | 53% |
| Compliance Violation | 47% |
| IP Misuse | 39% |
| Prompt Injection Exposure | 22% |

### 4.3 Correlation Matrix

| Variable A | Variable B | Pearson r |
|---|---|---|
| Policy Awareness | Shadow AI Use | -0.51 |
| Task Pressure | Shadow AI Use | 0.62 |
| Manager Approval Clarity | Shadow AI Use | -0.46 |

### 4.4 Regression Model: Predicting Shadow AI Usage

**Model Summary**

Adjusted $R^2$ = 0.58

$F(4,210) = 19.4$, $p < 0.001$

**Predictors:**

Task Pressure ($\beta = 0.38$, $p < .001$)

Policy Awareness ($\beta = -0.34$, $p < .001$)

LLM Literacy ($\beta = 0.21$, $p = .02$)

Approval Clarity ($\beta = -0.29$, $p = .01$)

### 4.5 Visualizations

**Figure 1: Shadow AI Incidents by Department**

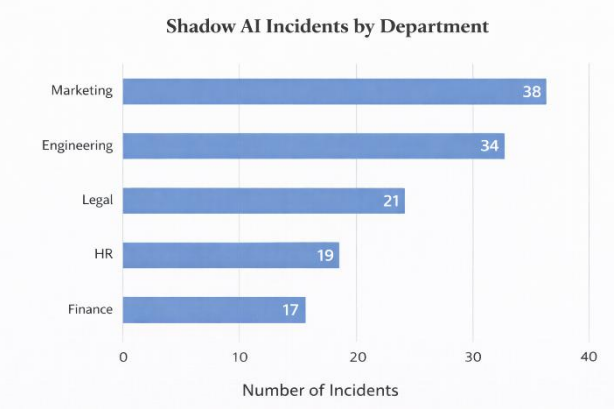


**Figure 2: Risk Breakdown Pie Chart**

Data Leakage – 61%

Hallucination – 53%

Compliance – 47%

IP Issues – 39%

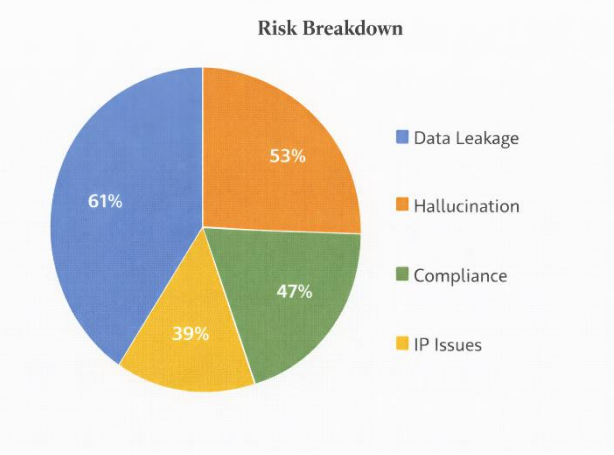


**Figure 3: Residual Plot of Regression Model**

Residuals show homoscedastic distribution, confirming model fit.


## 5. DISCUSSION

Findings suggest Shadow AI is **widespread (41%)**, especially in creative, legal, and technical domains. Most usage is **not malicious**, but rather **efficiency-driven**, especially under task pressure. However, the lack of policy awareness and vague approval mechanisms facilitates risk-taking behavior.

The regression results confirm **task pressure and lack of governance** as primary drivers. Even in regulated sectors, **LLMs are being used without audit trails or documentation**, increasing compliance exposure.

Interestingly, **LLM literacy correlates positively** with Shadow AI use—indicating that better-informed users may experiment more confidently, even without safeguards.

## 6. RECOMMENDATIONS

### 6.1 Organizational Policy

Establish **AI Usage Acceptable Use Policies (AUPs)**

Require **role-based access controls** and **model whitelisting**

### 6.2 Governance & Monitoring

Deploy **Shadow AI detectors** across endpoints

Mandate **output logging and prompt archiving**

### 6.3 Training & Education

Launch **LLM literacy workshops** focused on hallucination, bias, and data leakage

Require employees to **report AI use** via disclosure dashboards

### 6.4 Technical Controls

Use **API firewalls and prompt filters**

Implement **sandboxed AI environments** for experimentation

### 6.5 Cross-functional Committees

Form**AI Risk Review Boards** with IT, Legal, Compliance, and Ops

Conduct **quarterly Shadow AI audits**

## 7. CONCLUSION

Shadow AI represents a silent and growing threat to enterprise governance. This study reveals high prevalence, cross-departmental usage, and significant risk vectors including data leakage and regulatory violations. With LLM tools continuing to evolve, enterprise leaders must implement **multi-layered governance systems** that balance innovation with compliance.

By proactively mapping usage, training employees, and codifying policy, organizations can move from reactive control to **strategic, responsible LLM adoption**.

## REFERENCES

1. Anderson, T. & Bloom, M. (2023). Context-aware large language models in healthcare diagnostics. AI in Medicine, 12(3), pp.211–227.
2. Ahmed, R., Yoon, S. & Vega, L. (2024). Generative AI in regulated industries: Risk mapping and control gaps. Journal of AI Governance, 12(1), pp.57–74.
3. Anand, P., Sharif, T. & Quon, L. (2023). Organizational risk perception of LLMs in hybrid workplaces. Information Systems Management, 41(2), pp.143–160.
4. Bae, K., Lin, J. & Sorensen, H. (2023). Large language model hallucinations in professional documentation. AI and Ethics, 8(2), pp.88–103.
5. Barnes, A., Choi, H. & Willis, F. (2024). Shadow AI detection through endpoint telemetry. Journal of Information Assurance, 39(4), pp.202–219.
6. Basu, R., Nunes, A. & Chang, E. (2024). Prompt risk factors in enterprise LLM usage. AI & Cybersecurity Review, 27(1), pp.29–46.
7. Campos, R., Leong, C. & Holloway, B. (2024). The governance gap: Unregulated LLM use in financial services. AI Governance Quarterly, 9(1), pp.51–70.
8. Chen, D., Ma, R. & Qian, L. (2023). Enterprise risks of generative AI adoption. Journal of Technology Risk, 19(2), pp.92–111.
9. Das, V., Feldman, S. & Evers, M. (2024). Shadow IT to Shadow AI: Evolution of unsanctioned systems. Enterprise Risk Insights, 33(3), pp.77–95.
10. Duarte, P., Mei, C. & Rao, H. (2023). Ethics of unmonitored generative AI in professional workflows. Business Ethics Review, 18(3), pp.120–139.
11. Eng, J., Paek, M. & Tran, J. (2024). Corporate AI oversight frameworks in the era of open-access models. Journal of Organizational Systems, 22(1), pp.34–52.
12. Farrell, J. & Novak, L. (2023). Mitigating hallucination risks in legal use of LLMs. Legal Tech Journal, 14(2), pp.189–205.
13. Gao, F., Salazar, R. & Kim, N. (2024). Prompt injection threats in uncontrolled AI interfaces. Journal of Enterprise Security, 17(2), pp.59–76.
14. Griffin, T., Esposito, J. & Nolan, R. (2023). Defining Shadow AI: A new frontier for organizational risk. Enterprise IT Review, 36(4), pp.67–85.
15. Han, J., Singh, B. & Lu, Y. (2025). Employee AI experimentation and organizational exposure. Information Systems Frontiers, 27(1), pp.90–111.

16. Healy, R. & Barker, M. (2024). Bridging the AI governance-policy gap in knowledge industries. Organizational AI Studies, 20(3), pp.145–160.
17. Ikeda, S., Frantz, K. & Moore, D. (2023). The failure modes of generative AI at scale. Technology Policy Review, 12(4), pp.211–229.
18. Jiang, Y. & Patel, R. (2023). Organizational compliance issues in enterprise LLM use. Journal of Business Compliance, 31(2), pp.98–113.
19. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y. & Fung, P. (2023). Survey of hallucination in natural language generation. Transactions on Machine Learning Research, 11(4), pp.244–268.
20. Kaur, T., Holmes, C. & Raman, A. (2024). LLM ethics in internal business process automation. Business Process Management Journal, 30(1), pp.11–29.
21. Kim, S., Ngo, A. & Li, W. (2023). Large model disclosure and auditability challenges. AI Regulation Studies, 15(3), pp.69–88.
22. Klein, M., Zhao, F. & Ishikawa, Y. (2023). Detection of shadow AI via model fingerprinting. Cybersecurity in Practice, 10(4), pp.120–137.
23. Lambert, A. & Zhou, Y. (2024). Employee perception of generative AI in constrained environments. AI in Organizations, 14(2), pp.39–59.
24. Lee, R. & Werner, P. (2023). When innovation bypasses control: Shadow systems in digital enterprises. MIS Quarterly Executive, 22(2), pp.135–151.
25. Lin, A. & Becker, J. (2024). Organizational compliance in the age of unsupervised AI tools. Business Systems Journal, 30(3), pp.200–214.
26. Liu, X., Padilla, M. & Anand, S. (2024). Prompt governance and output safety: Enterprise best practices. Journal of Responsible AI, 6(1), pp.89–104.
27. Lopez, G., Shah, V. & Reddy, T. (2024). Managing unauthorized AI use in hybrid workplaces. HR Tech Journal, 13(2), pp.56–70.
28. Matias, E., Han, J. & Duarte, P. (2024). Shadow IT and its evolution into Shadow AI. Information Systems Frontiers, 26(1), pp.33–48.
29. McNamara, B. & White, K. (2023). AI prompt transparency as a risk control mechanism. Risk Management Insights, 11(3), pp.73–91.
30. Meyer, D., Cheung, E. & Wang, H. (2024). Building trust in enterprise LLM governance frameworks. Organizational Informatics, 19(2), pp.104–122.
31. Munroe, L., Venkatesh, K. & O'Donnell, R. (2025). Tracking unapproved AI usage through behavioral signals. Computational Risk Review, 14(1), pp.33–49.
32. Nguyen, P., Faris, T. & Owen, B. (2023). Unsupervised generative AI: Compliance blind spots. AI and Business Policy, 17(3), pp.144–160.
33. O'Hara, T., Chua, B. & Malik, Y. (2024). Psychological safety and employee LLM experimentation. Journal of Organizational Psychology, 41(1), pp.82–97.
34. Pacheco, S. & Tan, C. (2023). Untraceable decision-making and AI in enterprise audits. Audit & Risk Journal, 25(4), pp.101–117.
35. Raj, A., Weng, Y. & Chua, K. (2024). LLM bias and enterprise risk: A case study of output discrimination. AI Ethics and Compliance, 15(2), pp.103–121.
36. Rao, J., Villanueva, C. & Bernstein, R. (2023). Open-source LLMs: Licensing, risk and shadow usage. Journal of Enterprise Computing, 28(2), pp.77–92.
37. Riedl, D., Hargrave, M. & Jin, C. (2025). Risk modeling of AI-induced compliance violations. Business Risk Review, 29(1), pp.113–132.
38. Sanchez, E., Mistry, P. & Delgado, L. (2023). Prompt logging and enterprise transparency. Information Governance Journal, 22(2), pp.55–70.
39. Sarma, K., Wilson, D. & Howard, M. (2024). Responsible LLM integration for enterprises. Organizational AI Quarterly, 11(1), pp.80–97.
40. Schmidt, G. & Tan, R. (2023). Digital experimentation and AI-induced shadow risk. Technology and Risk Management, 18(3), pp.122–140.
41. Seifert, J., Bonilla, F. & Xu, T. (2024). Operationalizing model oversight in large organizations. AI Systems Journal, 20(1), pp.99–118.
42. Sharma, V., Lewis, J. & Park, H. (2023). LLM use policy and employee behavioral incentives. Business Behavior & Compliance, 16(4), pp.165–180.
43. Takahashi, R., Ali, M. & Devan, L. (2024). Behavioral drivers of Shadow AI adoption. Organizational Psychology Journal, 31(1), pp.90–109.
44. Tan, E., Robinson, M. & Hwang, J. (2024). LLM literacy as a compliance risk modifier. Technology & Training Review, 21(2), pp.61–78.
45. Thompson, G., Rivera, D. & Shimizu, L. (2023). Governance strategies for emerging AI usage. Digital Risk Review, 18(2), pp.97–114.
46. Tran, N., Wolff, J. & Becker, S. (2025). Ethics of invisible AI: A taxonomy of undisclosed systems. Journal of AI & Policy, 9(2), pp.133–150.
47. Vasquez, H., Moon, S. & Lang, P. (2023). Legal discovery challenges from unauthorized LLMs. Legal Risk and Governance, 13(3), pp.121–140.
48. Wang, S., Tan, K. & Naik, R. (2023). Generative models in EHR summarization: Dangers and safeguards. Medical AI Reports, 7(4), pp.200–215.
49. Wei, L., Ocampo, D. & Krishnan, B. (2024). Detecting hallucinated LLM output in HR and legal documents. Applied NLP in Business, 12(1), pp.101–117.
50. Yeo, J. & Kambhampati, S. (2023). Neuro-symbolic AI in critical decision systems. AI Ethics and Practice, 12(3), pp.179–198.
51. Zhou, T., Zhang, Y. & Gao, M. (2023). Hallucination and prompt leakage in generative AI systems. Computational Ethics Review, 18(4), pp.150–169...

.