

LUISS



Corso di Laurea in Law, Digital Innovation and Sustainability

Cattedra di Data Protection

DATA FLOWS IN THE AI ECOSYSTEM

THE RELATIONSHIP BETWEEN DEVELOPERS, PROVIDERS, AND CLIENTS UNDER GDPR AND THE AI ACT

Relatore

Chiar.mo Prof. Filiberto Emanuele Brozzetti

Correlatore

Chiar.ma Prof.ssa Fernandes Da Silva Ranchordas Sofia Hina

Candidata

Nicole Brolpito (n. mat. 632053)

Anno Accademico 2024/2025

*Con riconoscenza al Professor Brozzetti, per la costante guida, il prezioso supporto
e gli stimoli intellettuali che hanno reso possibile la realizzazione di questa tesi*

Abstract:

The governance of artificial intelligence in enterprise environments emerges at the intersection of technology, law and fundamental rights. Microsoft 365 Copilot, presented as a productivity enhancer, is examined as a case study to explore the challenges of compliance with the General Data Protection Regulation (GDPR), the forthcoming Artificial Intelligence Act (AI Act) and Italian labor law. Copilot's architecture, built on the partnership between Microsoft and OpenAI, illustrates how distributed ecosystems complicate the allocation of accountability and obscure the transparency of data flows. While technical measures such as retrieval-augmented generation and zero-retention policies are framed as safeguards, they fail to dispel uncertainties around indirect data transfers and the role of third parties in processing. The thesis investigates these dynamics through two guiding questions: whether Microsoft indirectly exposes enterprise data to OpenAI during inference, and whether Copilot functions primarily as a tool of efficiency or as an instrument of workplace surveillance. The first line of inquiry reveals how legal responsibility is fragmented, weakening the accountability principle at the core of GDPR. The second highlights the risk of function creep, where insights generated from employee communications blur the boundary between assistance and monitoring, raising concerns under Article 22 GDPR, the AI Act's risk-based framework, and Article 4 of the Workers' Statute. Copilot thus stands as a test case for the European regulatory response to enterprise AI. It embodies the tension between innovation and protection, productivity and dignity, efficiency and autonomy. Its significance lies not only in its technical design but in the broader question of whether legal frameworks can adapt quickly enough to ensure that the digital workplace evolves without eroding fundamental rights.

Keywords:

Artificial Intelligence; AI-as-a-Service (AIaaS); Microsoft 365 Copilot; Open AI; GDPR; AI Act; Labor Law; Workplace Surveillance; Retrieval-Augmented Generation (RAG).

Index:

INTRODUCTION	7
1.1 THE RISE OF ARTIFICIAL INTELLIGENCE IN ENTERPRISE SERVICES	7
1.2 LEGAL AND ETHICAL CHALLENGES OF DATA FLOWS IN THE AI ECOSYSTEM	8
1.3 RESEARCH OBJECTIVES AND GUIDING QUESTIONS	9
1.4 METHODOLOGICAL APPROACH AND ANALYTICAL TOOLS	11
1.4 THE ARGUMENTATIVE PATH OF THE THESIS	12
CHAPTER 1: THE ARCHITECTURE OF AI SERVICES AND DATA RELATIONSHIPS	14
2.1. KEY STAKEHOLDERS IN AI SERVICES: DEVELOPERS, PROVIDERS, AND CLIENTS	14
2.1.1. <i>AI Developers</i>	14
2.1.2. <i>AI Service Providers</i>	15
2.1.3. <i>Enterprise Clients and End-Users</i>	16
2.2. GENERAL ARCHITECTURE OF AI SERVICES AND CLOUD INFRASTRUCTURE	17
2.2.1. <i>Cloud-Based AI Infrastructure</i>	17
2.2.2. <i>AI Model Lifecycle in Cloud Services</i>	18
2.2.3. <i>Security Mechanisms in AI Services</i>	20
2.3. DATA FLOW IN CLOUD-BASED AI: FROM USER INPUT TO MODEL PROCESSING	22
2.3.1. <i>Stages of Data Flow in AI Services</i>	22
2.3.2. <i>Data Transmission in AI Systems</i>	23
2.3.3. <i>Risks and Mitigation Strategies in AI Data Processing</i>	24
2.4. RETRIEVAL-AUGMENTED GENERATION (RAG) AS A PRIVACY-ENHANCING TECHNOLOGY (PET)	25
2.4.1. <i>Introduction to Retrieval-Augmented Generation (RAG)</i>	25
2.4.2. <i>Privacy Benefits of RAG</i>	26
2.4.3. <i>Challenges and Risks of RAG</i>	27
2.4.4. <i>Use Cases of RAG in Enterprise AI</i>	28
2.5. CONTRACTUAL AND REGULATORY FRAMEWORK GOVERNING AI DATA INTERACTIONS	29
2.5.1. <i>Overview of AI Governance Frameworks</i>	29
2.5.2. <i>Compliance Challenges in AI Data Processing</i>	31
CHAPTER 2: ANALYZING THE FIRST RESEARCH QUESTION – DOES MICROSOFT, AS THE PROVIDER OF COPILOT, INDIRECTLY TRANSFER USER DATA TO OPENAI?	33
3.1. THE ARCHITECTURE OF MICROSOFT 365 COPILOT AND OPENAI’S ROLE	33
3.1.1. <i>Technical Integration of Copilot and OpenAI</i>	33
3.1.2. <i>Data Flow in Copilot’s AI Processing</i>	35
3.2. MICROSOFT AZURE OPENAI SERVICE, RETRIEVAL-AUGMENTED GENERATION (RAG), AND ENTERPRISE DATA PROCESSING	36
3.2.1. <i>Microsoft Azure OpenAI Service and Enterprise Data Handling</i>	36
3.2.2. <i>Retrieval-Augmented Generation (RAG) in Microsoft 365 Copilot</i>	37
3.2.3. <i>Risks and Compliance Challenges in Enterprise AI Data Processing</i>	38
3.3. TRANSPARENCY IN MICROSOFT’S DOCUMENTATION AND COMPLIANCE WITH GDPR ARTICLES 12 TO 15 AND THE AI ACT	39
3.3.1. <i>Transparency and AI Governance Under GDPR</i>	40
3.3.2. <i>AI Act Compliance and High-Risk AI Classification</i>	41
3.3.3. <i>Evaluating Microsoft’s AI Governance Strategies</i>	42
3.4. POSSIBLE TEMPORARY EXPOSURE OF DATA TO OPENAI: LEGAL INTERPRETATION OF REAL-TIME PROCESSING WITHOUT RETENTION AND COMPARISON WITH CASE LAW ON INSTANTANEOUS DATA ANALYSIS	43
3.4.1. <i>Understanding the Risk of Instantaneous Data Exposure</i>	44
3.4.2. <i>Legal Precedents on Real-Time Data Processing Without Retention</i>	45
3.4.3. <i>Regulatory Gaps in Defining AI Model Inference Risks</i>	46
3.5. OPENAI’S MODEL TRAINING AND THE POTENTIAL INDIRECT USE OF ENTERPRISE DATA	47

3.5.1. Does OpenAI Indirectly Benefit From Copilot Data?	48
3.5.2. GDPR's Purpose Limitation Principle and AI Training Risks	49
3.5.3. Compliance Strategies to Ensure Data Containment in AI Systems	50

CHAPTER 3: ANALYZING THE SECOND RESEARCH QUESTION – IS MICROSOFT 365 COPILOT A TOOL FOR PRODUCTIVITY OR SURVEILLANCE? 52

4.1. THE PROCESSING OF EMPLOYEE DATA THROUGH COPILOT AND AI'S ROLE IN PERFORMANCE MONITORING	52
4.1.1. Microsoft Productivity Score and Its Relevance to Copilot's Workplace Analytics	52
4.1.2. Copilot's Role in Workplace Analytics: Productivity Enhancement or Indirect Monitoring?	53
4.1.3. Copilot and the Expansion of AI-Powered Performance Monitoring	54
4.2. APPLICABILITY OF ARTICLE 4 OF THE WORKERS' STATUTE: IS COPILOT A NECESSARY WORK TOOL OR A SURVEILLANCE MECHANISM?	56
4.2.1. Legal Framework of Article 4: Key Provisions and Protections	56
4.2.2. Is Copilot a Legitimate Work Tool or a Compliance Risk?	58
4.2.3. Legal Precedents on AI and Workplace Surveillance	60
4.3. GDPR AND AI ACT COMPLIANCE: THE REGULATORY FRAMEWORK FOR AI-DRIVEN WORKPLACE TOOLS	61
4.3.1. GDPR Compliance: Transparency and Employee Data Protection	61
4.3.2. The AI Act and Workplace AI Tools: Is Copilot a High-Risk System?	62
4.3.3. Compliance Challenges and Legal Uncertainties	63
4.4. AI AS A TOOL FOR PRODUCTIVITY OR A RISK FOR EMPLOYEE AUTONOMY	64
4.4.1. The Rise of Algorithmic Management in Workplaces	64
4.4.2. Ethical Considerations: Worker Privacy vs. Employer Control	65
4.4.3. The Role of Labor Unions and Worker Protections in AI Regulation	67
4.5. CONSIDERATIONS AND FINAL ANSWER: COPILOT AS A PRODUCTIVITY TOOL OR SURVEILLANCE RISK?	69
4.5.1. Key Findings from the Analysis	70
4.5.2. Determining Factors: Implementation, Transparency, and Governance	72
4.5.3. Final Verdict: Is Microsoft 365 Copilot a Tool for Productivity or Surveillance?	73

CONCLUSION: 75

5.1 OBJECTIVE, RESEARCH QUESTIONS, AND METHOD RECAP	75
5.2 CRITICAL AND STRUCTURED SUMMARY OF FINDINGS	75
5.2.1 Chapter 1 – AI Service Architecture and Distributed Responsibility	75
5.2.2 Chapter 2 – Indirect Data Transfer to OpenAI and Lack of Transparency	76
5.2.3 Chapter 3 – Algorithmic Surveillance and Labor Law	77
5.3. LIMITATIONS	79
5.4. POLICY AND RESEARCH RECOMMENDATIONS	79

BIBLIOGRAPHY: 81

Introduction

1.1 The Rise of Artificial Intelligence in Enterprise Services

Over the past decade, artificial intelligence (AI) has shifted from being an experimental research field to a transformative force across industries. In the enterprise sector in particular, AI has evolved from discrete applications, such as automated chatbots or predictive analytics, into fully integrated systems embedded within everyday business tools. This integration is epitomized by the growing phenomenon of *Artificial Intelligence as a Service (AIaaS)*, where AI functionalities are seamlessly delivered through cloud infrastructures and incorporated into standard productivity platforms. The spread of AIaaS illustrates how the technology is no longer confined to specialist teams but has become a ubiquitous feature of the digital workplace, directly shaping how knowledge is produced, shared, and consumed within organizations. This development coincides with a broader process of digital transformation. Companies are under increasing pressure to streamline workflows, enhance efficiency, and remain competitive in markets that demand constant innovation. AI systems, particularly those that exploit large-scale language models, promise to automate routine tasks, reduce administrative burdens, and unlock new forms of insight from vast pools of corporate data. By embedding such capabilities within enterprise software ecosystems, firms like Microsoft position AI not merely as a technical add-on, but as an indispensable layer of the digital infrastructure that supports daily work activities. The implications of this shift are profound. On the one hand, the adoption of AI in enterprise services is expected to boost productivity and enable organizations to harness the full potential of their data assets. On the other hand, the same systems raise pressing questions about transparency, accountability, and the safeguarding of fundamental rights. When AI systems operate as “black boxes” within workplace tools, their influence on decision-making, communication, and performance evaluation can blur the line between technological assistance and hidden forms of surveillance. This tension becomes particularly evident when analyzing Microsoft 365 Copilot, a generative AI assistant marketed as a productivity enhancer but capable of producing outputs that may indirectly affect how employees are monitored and assessed. Against this background, the European regulatory landscape takes on a decisive role. The General Data Protection Regulation (GDPR) has established robust principles of accountability, transparency, and proportionality in data processing. The

forthcoming Artificial Intelligence Act (AI Act) aims to complement this framework by introducing a risk-based regulatory model specifically tailored to AI. Together, these instruments seek to balance the benefits of technological innovation with the imperative to protect individuals' rights in digital environments. Enterprise AI services like Copilot thus stand at the intersection of innovation and regulation, serving as a test case for the effectiveness of Europe's evolving legal frameworks.

1.2 Legal and Ethical Challenges of Data Flows in the AI Ecosystem

The integration of artificial intelligence into enterprise environments introduces not only technological opportunities but also complex legal and ethical challenges, particularly in relation to the governance of data flows. Enterprise AI systems, such as Microsoft 365 Copilot, operate within multilayered ecosystems that involve developers, providers, and client organizations. This distributed structure complicates the allocation of responsibilities and creates zones of ambiguity regarding compliance with existing regulatory frameworks. A first challenge concerns the fragmentation of accountability across multiple actors. Developers such as OpenAI create and maintain the generative models; providers like Microsoft embed these models into enterprise software and cloud infrastructures; and corporate clients rely on the tools for daily operations. Each of these actors is involved in different stages of data processing, ranging from collection and storage to inference and output generation. However, the lack of a unified framework to clarify their respective legal roles risks diluting the accountability principle enshrined in Article 5(2) GDPR, which requires data controllers to both ensure and demonstrate compliance. When responsibility is distributed without clear boundaries, it becomes difficult for organizations to determine which actor is liable for which aspect of data handling, and for data subjects to exercise their rights effectively. A second challenge lies in the opacity of data flows. Although Microsoft asserts that enterprise data processed within Copilot does not leave its Azure environment, questions remain about whether transient or indirect interactions with OpenAI's models occur during inference. Under Article 4(2) GDPR, even temporary operations, such as data held momentarily in memory during processing, qualify as "processing." This implies that any exposure of enterprise data to OpenAI infrastructure, however short-lived, could fall under the scope of European data protection law. The absence of explicit technical proof or contractual clarification

regarding OpenAI's role in the inference process raises significant transparency concerns. Without clear documentation and auditability, enterprise clients are left unable to verify whether regulatory safeguards are consistently respected. Beyond legal uncertainty, these dynamics also raise important ethical issues. AI systems embedded in workplace environments often process communications, documents, and metadata that reflect employees' behaviors and interactions. While intended for productivity gains, such processing can inadvertently enable profiling and monitoring practices. The boundary between assistance and surveillance is blurred when AI outputs allow managers to infer patterns of performance, productivity, or collaboration. This tension is particularly acute in jurisdictions such as Italy, where Article 4 of the Workers' Statute prohibits remote monitoring of employees unless specifically authorized. Even absent direct surveillance mechanisms, the possibility that AI systems generate insights with evaluative implications underscores the risk of "function creep," where tools designed for efficiency acquire unintended monitoring capacities. The broader societal implications of these challenges are equally significant. For enterprises, uncertainties regarding data governance increase the risk of non-compliance and potential sanctions. For employees, opaque data flows raise concerns about workplace autonomy, privacy, and dignity. For policymakers and regulators, the rise of enterprise AI highlights the urgent need to adapt existing legal frameworks to account for new modalities of data processing. The GDPR and AI Act offer complementary approaches—one focused on data protection principles, the other on risk-based governance—but the rapid pace of technological deployment exposes the limitations of both regimes in addressing emerging scenarios of accountability and surveillance.

1.3 Research Objectives and Guiding Questions

The rapid adoption of artificial intelligence in enterprise settings requires not only technical adaptation but also a rigorous assessment of the legal frameworks that govern its deployment. This thesis positions Microsoft 365 Copilot as a case study to investigate how existing and forthcoming European regulations respond to the challenges posed by AI-as-a-Service (AIaaS). Against this background, the research pursues a dual objective: to clarify the allocation of legal responsibilities within complex AI ecosystems and to assess the potential consequences of AI-driven tools

for both organizational governance and the protection of individual rights. The first objective is to analyze the extent to which Microsoft, as the provider of Copilot, indirectly transfers enterprise user data to OpenAI in the process of AI inference. While Microsoft publicly claims that no personal data leaves its Azure environment or is shared with OpenAI, the architecture of Copilot reveals points of ambiguity concerning transient data handling. This raises fundamental questions about the applicability of the GDPR, particularly its provisions on accountability, transparency, and data transfers. The inquiry aims to determine whether Microsoft's contractual assurances are sufficient to meet regulatory obligations, or whether the involvement of OpenAI creates risks of hidden or indirect processing that may undermine compliance. The second objective is to examine whether Microsoft 365 Copilot, beyond its marketed role as a productivity enhancer, functions in practice as a tool of workplace surveillance. Because the system processes vast amounts of enterprise data—emails, documents, meeting notes, and communication metadata—its outputs can indirectly enable managers to infer patterns of behavior, performance, and collaboration. Even if such monitoring is not an explicit design feature, the risk of *function creep* places Copilot at the intersection of data protection law, the AI Act, and Italian labor law. The research therefore investigates whether Copilot's functionalities align with GDPR provisions on automated decision-making (Article 22), the AI Act's risk-based classification framework, and the protections against unauthorized monitoring established under Article 4 of the Workers' Statute. These objectives are pursued through two guiding research questions:

1. Does Microsoft, as the provider of Copilot, indirectly transfer user data to OpenAI, and if so, is this practice compatible with the GDPR and the AI Act?
2. Should Microsoft 365 Copilot be considered primarily a tool for productivity, or does it also function as a mechanism of workplace surveillance, with implications under both data protection and labor law?

By addressing these questions, the thesis seeks to contribute to the broader academic and policy debate on the governance of enterprise AI. It provides insights not only for legal scholars but also for enterprises navigating compliance obligations and for regulators tasked with safeguarding fundamental rights in the digital workplace.

1.4 Methodological Approach and Analytical Tools

The complexity of enterprise artificial intelligence services requires a methodological approach that integrates technical, legal, and doctrinal perspectives. This thesis adopts a mixed-method strategy designed to capture both the operational mechanisms of Microsoft 365 Copilot and the normative implications of its deployment under European and national law. The methodology combines documentary analysis of technical sources, doctrinal interpretation of legal frameworks, examination of regulatory practice, and engagement with academic literature. Technical-documentary analysis forms the foundation of the inquiry. Official Microsoft documentation, white papers, and transparency notes were examined to reconstruct the architecture of Copilot, focusing on the movement of data through Microsoft Graph, Azure OpenAI Service, and the GPT-4 inference process. Special attention was given to privacy-enhancing technologies (PETs), such as retrieval-augmented generation (RAG) and zero-retention policies, which are promoted as mechanisms to reduce compliance risks. This analysis provides the necessary basis for assessing whether enterprise data remains fully contained within Microsoft's infrastructure or whether technical ambiguities create the possibility of indirect exposure to OpenAI. Doctrinal and legal analysis complements this technical reconstruction. The research evaluates Copilot's compliance with the General Data Protection Regulation (GDPR), particularly its provisions on accountability (Article 5(2)), transparency (Articles 12–15), and automated decision-making (Article 22). Parallel consideration is given to the forthcoming AI Act, with a focus on its risk-based classification system and obligations for high-risk AI systems, and to Italian labor law, especially Article 4 of the Workers' Statute, which prohibits unauthorized monitoring in the workplace. These legal frameworks are interpreted in light of their objectives of safeguarding fundamental rights while enabling responsible innovation. Examination of regulatory practice and case law further grounds the analysis in real-world enforcement. Decisions from European data protection authorities, such as the Italian DPA's stance on biometric surveillance and the Swedish DPA's ruling on school-based facial recognition, demonstrate that even transient or indirect processing of personal data falls under GDPR obligations. Case law concerning algorithmic management, including rulings against Deliveroo, highlights the relevance of transparency and accountability principles in employment contexts. These precedents are used to assess whether Copilot's architecture aligns with existing regulatory expectations. Finally, a

critical review of academic and policy literature contextualizes the findings within broader debates on AI governance. Scholarships on distributed accountability, the opacity of data flows, and the ethical risks of algorithmic management provides conceptual tools for understanding the structural issues identified in the Copilot case. Policy reports and position papers from European institutions and expert groups complement this analysis by offering insights into the evolving regulatory landscape and highlighting areas where existing law may be insufficient. By combining these approaches, the thesis establishes a robust analytical framework that bridges the gap between technical realities and legal principles. This interdisciplinary methodology not only allows for a detailed evaluation of Copilot's compliance with GDPR and the AI Act but also provides a basis for reflecting on the broader societal implications of embedding generative AI into enterprise environments.

1.4 The Argumentative Path of the Thesis

The structure of this thesis reflects its dual objective of clarifying the allocation of responsibilities within enterprise AI ecosystems and assessing the implications of Microsoft 365 Copilot for privacy and labor rights. The argument unfolds progressively, beginning with the technical and architectural foundations of AI services, moving through the legal and ethical challenges they generate, and concluding with an assessment of regulatory adequacy. The first chapter lays the conceptual groundwork by examining the architecture of enterprise AI and the roles of the key actors—developers, providers, and client organizations—across the data processing lifecycle. Particular attention is given to the infrastructures that enable AI-as-a-Service and to the privacy-enhancing technologies, such as retrieval-augmented generation (RAG), that are presented as safeguards but raise questions of effectiveness. This contextual framework is essential to understanding how data circulates within Copilot and how accountability becomes fragmented across different stakeholders. Building on this foundation, the second chapter addresses the first research question: whether Microsoft, as the provider of Copilot, indirectly transfers enterprise user data to OpenAI. By reconstructing Copilot's data flows and analyzing Microsoft's contractual assurances and technical safeguards, the chapter evaluates their compatibility with the GDPR and the forthcoming AI Act. The analysis highlights unresolved ambiguities regarding OpenAI's role in the inference process, the

interpretation of “processing” under Article 4(2) GDPR, and the sufficiency of transparency obligations. These findings point to the persistence of significant compliance risks, even in the presence of strong corporate commitments to privacy. The third chapter shifts the focus to the second research question: whether Copilot functions purely as a productivity enhancer or also as a mechanism of workplace surveillance. Here the discussion explores how the system’s capacity to process and recombine enterprise communications and documents may allow managers to infer patterns of employee behavior and performance. This dual functionality raises concerns under GDPR Article 22 on automated decision-making, the AI Act’s high-risk classification criteria, and Article 4 of the Italian Workers’ Statute, which strictly regulates employee monitoring. The chapter argues that, depending on its deployment, Copilot can reasonably be regarded as an algorithmic surveillance tool, with profound implications for workplace autonomy and dignity. The conclusion synthesizes these findings and situates them within the broader debate on AI governance. It emphasizes that Copilot illustrates the structural challenges of applying existing legal frameworks to distributed AI ecosystems, where accountability is fragmented and data flows remain opaque. In particular, it underlines the dilution of the GDPR’s accountability principle, the absence of clarity regarding OpenAI’s involvement in inference, and the potential for productivity tools to assume monitoring functions in practice. At the same time, it acknowledges the limitations of the research, which relies on publicly available documentation, and it advances recommendations aimed at strengthening regulatory clarity, contractual transparency, and policy oversight. Through this argumentative path, the thesis develops a layered understanding of how enterprise AI services like Copilot test the resilience of European data protection law, challenge established labor protections, and reveal gaps in the emerging regulatory framework of the AI Act.

Chapter 1: The Architecture of AI Services and Data Relationships

2.1. Key Stakeholders in AI Services: Developers, Providers, and Clients

The development and deployment of artificial intelligence (AI) services involve multiple stakeholders, each playing a distinct role in shaping the AI ecosystem. These stakeholders include AI developers, AI service providers, enterprise clients, and regulatory bodies, all of whom share varying levels of responsibility regarding data governance, security, and regulatory compliance. Understanding their interactions is essential to evaluating how AI services are structured, deployed, and monitored within legal and ethical frameworks.

2.1.1. AI Developers

AI developers are the entities responsible for designing, training, and optimizing machine learning models. They include independent research organizations, corporate AI divisions, and academic institutions engaged in the advancement of AI technologies. They mainly define AI systems' architecture, select training data, and ensure that models perform well and comply with ethical and regulatory standards¹. This surge is led by developers like OpenAI, DeepMind, Microsoft Research, Google AI, and IBM Watson through proprietary and open-source models. One of the most critical responsibilities of AI developers is ensuring data governance and security. They must comply with legal frameworks such as the General Data Protection Regulation (GDPR) and the proposed EU AI Act, which impose strict requirements on data processing, algorithmic transparency, and risk mitigation². Developers also need to directly tackle bias and fairness challenges by performing audits that identify and prevent unfair treatments through AI-generated decisions. Due to the increasing importance of explainability, developers are advised to incorporate approaches such as Explainable AI (XAI)

¹ U.S. Department of Homeland Security, *Roles and Responsibilities Framework for Artificial Intelligence in Critical Infrastructure*, In Consultation with The Artificial Intelligence Safety and Security Board, November 14, 2024

² Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>

into their products, allowing users to gain insight into the output of AI models³. Besides compliance with applicable regulations, developers of data sharing platforms must implement technical privacy-enhancing technologies (PETs), including differential privacy and federated learning, in their platforms. These techniques aim to reduce exposure to unauthorized data access and model inversion attacks, all while maintaining the efficiency of AI systems. With the advancement of AI, developers can no longer be just techies as they are expected to incorporate robust security that can stand up against adversarial attacks and ethical governance of AI systems at every step of the life cycle⁴.

2.1.2. AI Service Providers

AI service providers deploy and manage AI models as cloud-based services, facilitating their accessibility and scalability for enterprise use. Key providers such as Microsoft Azure OpenAI Service, Google Cloud AI, and Amazon AWS AI operate large-scale infrastructure that supports AI model hosting, computation, and real-time data processing. These entities serve as intermediaries between AI developers and enterprise clients, ensuring that AI models function reliably while maintaining compliance with data security regulations⁵. Cloud-based AI infrastructure is a core component of AI service providers' operations. These providers offer computer resources on demand, enabling the smooth execution of AI workloads. It is not enough to simply secure data storage, implement encryption protocols, and enforce multi-layered authentication mechanisms to prevent unauthorized access to enterprise data. For example, Microsoft has built in real-time monitoring tools into Azure OpenAI Service for detecting anomalous efforts in AI-generated interactions, enabling enterprise clients to exercise control over

³ Hoffman Robert R. , Mueller Shane T. , Klein Gary , Jalaeian Mohammadreza , Tate Connor, *Explainable AI: roles and stakeholders, desirements and challenges*, Frontiers in Computer Science Vol.5, 2023-
<https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2023.1117848> DOI=10.3389/fcomp.2023.1117848

⁴ Deshpande, A., & Sharp, H. (2022). Responsible AI systems: Who are the stakeholders? Paper presented at the 227-236.
<https://doi.org/10.1145/3514094.3534187>

⁵ Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021).
<https://doi.org/10.1007/s12599-021-00708-w>

proprietary data while benefiting from AI-driven automation⁶. AI service providers not only manage infrastructure but also have data processing, storage, and compliance-related legal obligations. They should be in compliance with industry standards such as ISO 27001, SOC 2, and GDPR, to ensure that the deployment of AI is compliant with the law in different jurisdictions. In addition, AI service providers are also embracing more privacy-enhancing technologies like Retrieval-Augmented Generation (RAG) and federated learning to mitigate the possibility of data leakage while powering AI workloads⁷.

2.1.3. Enterprise Clients and End-Users

Enterprise clients represent businesses and organizations that integrate AI services into their workflows to enhance decision-making, automation, and operational efficiency. These clients range from multinational corporations to government agencies that deploy AI for predictive analytics, customer engagement, and data-driven insights. While AI adoption offers significant advantages, it also raises concerns related to data privacy, security, and regulatory compliance⁸. To address such risks, enterprise customers need to take control of data controls, data handling, and AI security configurations. It is the responsibility of organizations to enforce stringent access controls to the interaction (certainly AI-generated outputs) so that only accredited personnel can interact with it. In short, enterprises will need to ensure AI service providers comply with data sovereignty laws (which are much stricter in the European Union and China regarding cross-border data transfers). While this issue is compounded for enterprises with third-party AI solutions and AI as a service, enterprises must perform thorough due diligence to ensure that third-

⁶ U.S. Department of Homeland Security, *Roles and Responsibilities Framework for Artificial Intelligence in Critical Infrastructure*, In Consultation with The Artificial Intelligence Safety and Security Board, November 14, 2024

⁷ Hoffman Robert R. , Mueller Shane T. , Klein Gary , Jalaiean Mohammadreza , Tate Connor, *Explainable AI: roles and stakeholders, desirements and challenges*, *Frontiers in Computer Science* Vol.5, 2023-
<https://www.frontiersin.org/journals/computerscience/articles/10.3389/fcomp.2023.1117848> DOI=10.3389/fcomp.2023.1117848

⁸ Deshpande, A., & Sharp, H. (2022). Responsible AI systems: Who are the stakeholders? Paper presented at the 227-236.
<https://doi.org/10.1145/3514094.3534187>

party AI solutions align with their internal security policies⁹. However, despite these protective measures, the adoption of AI comes with inherent risks. Sensitivity leaks are among the top concerns, such as when AI insights accidentally expose sensitive corporate data or IP. Another major concern is bias in AI models; enterprises need to ensure fair and non-discriminatory AI-driven decision-making processes. Then the regulatory uncertainty imposes difficulties for the companies, the kaleidoscopic modalities of legality come with other clefts, the agency that needs to answer it¹⁰.

2.2. General Architecture of AI Services and Cloud Infrastructure

The architecture of AI services is deeply interconnected with cloud infrastructure, enabling organizations to leverage artificial intelligence at scale while ensuring computational efficiency, data security, and regulatory compliance. Cloud-based AI models rely on distributed computing frameworks to train, deploy, and operate machine learning algorithms, which require vast amounts of processing power and structured data flows. Understanding this architecture is essential for assessing the implications of AI adoption on enterprise systems, particularly concerning data handling practices and security risks. This section explores the fundamental components of cloud-based AI infrastructure, the lifecycle of AI models within cloud environments, and the security mechanisms implemented to protect AI-driven applications.

2.2.1. Cloud-Based AI Infrastructure

Cloud-based AI infrastructure consists of a set of interconnected computing resources, including data centers, virtual machines, APIs, and containerized AI models that

⁹ U.S. Department of Homeland Security, *Roles and Responsibilities Framework for Artificial Intelligence in Critical Infrastructure*, In Consultation with The Artificial Intelligence Safety and Security Board, November 14, 2024

¹⁰ Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>

enable businesses to deploy AI-driven applications seamlessly¹¹. These components facilitate AI model inference, real-time decision-making, and large-scale automation. Leading cloud providers such as Microsoft Azure, Amazon Web Services (AWS), and Google Cloud offer AI-as-a-Service (AIaaS) solutions, allowing enterprises to integrate pre-trained AI models into their workflows without requiring in-house model development¹². One of the key differentiators in cloud AI deployment is the choice of cloud models—public, private, and hybrid cloud infrastructures. Public clouds offer scalable AI services on multi-tenant, shared platforms that reduce operational costs but come with risks to data sovereignty and access control¹³. In contrast, they are deployed in dedicated enterprise environments, providing more granular control over security policies and regulatory compliance. Interest in the hybrid cloud model among enterprises is spiking, as organizations can train AI models on private infrastructure but then use public cloud resources for high-computation tasks. It manages to strike a balance between efficiency and privacy, especially for sectors that have strict compliance obligations, such as healthcare and finance¹⁴. A notable trend in real-world AI deployment has been edge computing, where models are deployed closer to data sources such as IoT devices, mobile applications, and on-premise servers rather than exclusively in cloud environments¹⁵. With Edge AI, latency is reduced, real-time decision-making is improved, and the risks of transmitting data to centralized cloud servers are minimized. Yet, deploying AI inference at the edge brings new and different security challenges around data synchronization and model updates.

2.2.2. AI Model Lifecycle in Cloud Services

Inside cloud environments, AI models adopt a structured lifecycle of development, training, deployment, and continuous optimization. This has been used to ensure that the AI platform is accurate, efficient, and compliant with the latest regulatory requirements. ASOM describes the different stages of an AI

¹¹ Allam, Karthik. (2023). Adoption of Artificial Intelligence in Cloud Computing. *International Journal of Computer Trends and Technology*, 71, 91-95. 10.14445/22312803/IJCTT-V71I6P116.

¹² Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>

¹³ van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241232630>

¹⁴ Meurisch, C., & Mühlhäuser, M. (2022). Data Protection in AI Services, *ACM Computing Surveys*, 54(2). <https://doi.org/10.1145/3440754>

¹⁵ Andrew, James. (2025). AI Model Lifecycle Management: Strategies for Scalable Deployment and Maintenance.

lifecycle, their respective roles, and how they complement each other in ensuring the reliability of an AI-based service (in particular, when an AI model is working with sensitive enterprise data). During the model development phase, training data are selected, algorithms are designed, and hyperparameters are tuned, usually in high-performance cloud environments that utilize GPUs (Graphics Processing Units) and TPUs (Tensor Processing Units) to speed up the training of deep learning models¹⁶. GPUs, originally designed for graphical rendering, are optimized for parallel processing, making them well-suited for complex machine learning tasks. TPUs, developed by Google, specialize in tensor operations, enabling faster and more efficient AI model execution—particularly in TensorFlow-based applications¹⁷. Nevertheless, despite the cloud-based AI infrastructures delivering more raw computing power than any individual could ever manage, they also bring risks related to data ownership, intellectual property rights, and compliance. Using third-party cloud providers like Microsoft Azure, Google Cloud, and AWS comes with the question of who is in charge of your training data and how it is being processed. This problem is especially critical for regulated industries that are bound by data sovereignty regulations like the GDPR and the U.S. Cloud Act that mandate stringent data residency, encryption, and access control¹⁸. AI models are subsequently trained and then enter a stage of fine-tuning, being adapted to specific business use cases with the help of domain-specific data. This process is important to increase model accuracy, reduce bias, and maintain relevance in practical scenarios. This forms a compelling argument for enterprises deploying any such AI services to adopt data minimization strategies, limiting data exposure via ensuring only data datasets required are processed. Doing so minimizes the risks tied to excessive data collection and inadvertent data sharing¹⁹.

The deployment stage involves putting AI models into production environments via containerized services like Docker and Kubernetes that improve scalability, fault tolerance, and resource handling. Many cloud providers use an API-centric model to deliver their AI models, enabling enterprises to build AI features into

¹⁶ Spjuth, O., Frid, J., & Hellander, A. (2021). The machine learning life cycle and the cloud: implications for drug discovery. *Expert Opinion on Drug Discovery*, 16(9), 1071–1079. <https://doi.org/10.1080/17460441.2021.1932812>

¹⁷ Andrew, James. (2025). AI Model Lifecycle Management: Strategies for Scalable Deployment and Maintenance.

¹⁸ Meurisch, C., & Mühlhäuser, M. (2022). Data Protection in AI Services, *ACM Computing Surveys*, 54(2). <https://doi.org/10.1145/3440754>

¹⁹ van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241232630>

applications without requiring a complete rework of current infrastructures. But a key challenge remains in the phenomenon of model drift, where the performance of AI gets worse over time as patterns in data evolve. Organizations are turning to MLOps (Machine Learning Operations) pipelines, which automate the retraining, versioning, and monitoring of models in production environments, to stave off model drift and maintain the reliability and accuracy of AI systems²⁰. As AI models evolve, the necessity of continuous monitoring, bias mitigation, and compliance assessments becomes more pressing. Organizations must balance efficiency gains from cloud-based AI models with stringent security measures that safeguard proprietary data, comply with legal frameworks, and prevent unauthorized access. By adopting privacy-preserving AI techniques such as federated learning and differential privacy, enterprises can enhance AI security while maintaining computational performance²¹. The AI model lifecycle thus represents a delicate equilibrium between innovation, security, and regulatory adherence, ensuring that AI-powered enterprise solutions remain effective, ethical, and resilient in dynamic digital environments.

2.2.3. Security Mechanisms in AI Services

As AI adoption increases, robust security mechanisms are critical for protecting AI models from data breaches, adversarial attacks, and unauthorized access. AI security strategies focus on three core areas: encryption, authentication, and compliance. Encryption is the process that converts data into an unreadable format using cryptographic algorithms, ensuring that only authorized parties with the correct decryption key can access the original information it is fundamental in securing AI-driven applications, particularly in cloud environments where sensitive data is stored and transmitted across multiple systems. End-to-end encryption ensures that AI queries and responses remain protected, preventing unauthorized access. To address this challenge, a technology called homomorphic encryption has recently emerged, enabling AI models to perform over encrypted data without decrypting

²⁰ Andrew, James. (2025). AI Model Lifecycle Management: Strategies for Scalable Deployment and Maintenance.

²¹ Sangwan RS, Badr Y, Srinivasan SM. Cybersecurity for AI Systems: A Survey. *Journal of Cybersecurity and Privacy*. 2023; 3(2):166-190. <https://doi.org/10.3390/jcp3020010>

it, thus providing security and privacy in cloud-based inference of AI²². Authentication and access control, security mechanisms that verify a user's identity and enforce permissions to ensure that only authorized individuals can access, are essential to restrict unauthorized interactions with AI models. Using role-based access control (RBAC) and multi-factor authentication (MFA), we can restrict access to only those users and applications that are authorized to obtain AI-generated insights. Most cloud AI platforms follow zero-trust security models, requiring continuous validation of access credentials to mitigate unauthorized activity. Additionally, audit logging and real-time anomaly detection help organizations monitor AI interactions and detect potential security threats²³. In terms of compliance, cloud AI providers need to comply with industry standards like ISO/IEC 27001 (Information Security Management), SOC 2 (Service Organization Control), and GDPR compliance. Such frameworks actually set strict data protection standards for organizations to follow while processing AI-driven data lawfully and transparently²⁴. Under GDPR regulations, AI providers are required to implement privacy-by-design principles, ensuring that AI models handle data with minimal exposure and maintain accountability for automated decisions. Despite these security measures, cloud-based AI services are still prone to adversarial strikes. Yet, there are risks for confidential information with model inversion attacks which can reconstruct training data from AI-generated outputs. Another threat comes from data poisoning attacks, where attackers insert corrupted inputs into AI training datasets, manipulating the AI behavior in counterproductive and unintended ways²⁵. To combat these threats, organizations are utilizing differential privacy techniques, which add controlled noise to the AI training data to protect sensitive information without impacting the accuracy of the model.

²² Sangwan RS, Badr Y, Srinivasan SM. Cybersecurity for AI Systems: A Survey. *Journal of Cybersecurity and Privacy*. 2023; 3(2):166-190. <https://doi.org/10.3390/jcp3020010>

²³ Andrew, James. (2025). AI Model Lifecycle Management: Strategies for Scalable Deployment and Maintenance.

²⁴ Meurisch, C., & Mühlhäuser, M. (2022). Data Protection in AI Services, *ACM Computing Surveys*., 54(2). <https://doi.org/10.1145/3440754>

²⁵ Sangwan RS, Badr Y, Srinivasan SM. Cybersecurity for AI Systems: A Survey. *Journal of Cybersecurity and Privacy*. 2023; 3(2):166-190. <https://doi.org/10.3390/jcp3020010>

2.3. Data Flow in Cloud-Based AI: From User Input to Model Processing

The flow of data in cloud-based AI services is a critical aspect that determines the efficiency, security, and compliance of AI-driven solutions. AI models rely on structured data pipelines that process user input, refine it for inference, and ensure secure transmission and storage. These stages introduce various risks, including unintended data retention, model inversion attacks, and cross-tenant data leakage. This section examines the lifecycle of data in AI services, highlighting security protocols and mitigation strategies that organizations employ to ensure compliance and protect sensitive information.

2.3.1. Stages of Data Flow in AI Services

The lifecycle of data in AI services consists of multiple interconnected phases, from data collection to post-processing and storage. Each phase presents potential vulnerabilities that require proactive security measures. The **data collection phase** represents the entry point where AI models receive input from users through various means, including API calls, software-as-a-service (SaaS) platforms, and enterprise systems. It is in having access to an API endpoint in the cloud that allows for this interaction, where input data is sent over the wire through encrypted communication channels like HTTPS and gRPC²⁶. Without a properly set measurement plan in place, the data collected may not accurately reflect the goal information needed to drive a business-influencing decision. Once collected, data undergoes **pre-processing**, a crucial step involving data anonymization, filtering, and transformation. This is a step to ensure that the AI model does not work on personally identifiable information (PII) directly but only works on identifiable data. For instance, various techniques like tokenization, differential privacy, and the use of synthetic data have been widely used in an effort to secure user privacy while striving to preserve model accuracy²⁷. During **inference**, AI models generate predictions based on processed

²⁶ Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>

²⁷ Witanto, E. N., Oktian, Y. E., & Lee, S.-G. (2022). Toward Data Integrity Architecture for Cloud-Based AI Systems. *Symmetry*, 14(2), 273. <https://doi.org/10.3390/sym14020273>

input data. This phase is where cloud-based machine learning systems derive value from the collected information. Model execution takes place in virtualized cloud environments, where pre-trained AI models process input data to produce responses in real time. Inference can, however, lead to specific threats occurring while using AI, including model inversion attacks in which adversaries attempt to reverse engineer the input data (based on the outputs generated by the AI)²⁸. Once inference is done, post-processing refines, validates, and formats the AI-generated results for use. Human interventions (e.g., human-in-the-loop (HITL)) may be used at this step when these classification outputs made by the AI can be reviewed and corrected to increase accuracy. Also, AI interactions are typically logged and stored for auditing and compliance purposes. Such logs are crucial for regulatory compliance, especially for GDPR and ISO/IEC 27001 requirements²⁹.

2.3.2. Data Transmission in AI Systems

The security of AI data transmission is a key concern, as data moves between clients, cloud models, and storage environments. AI services employ a range of communication protocols to ensure secure and efficient data exchange. The most common channels for data transmission between client devices and AI models as a service on the cloud are HTTPS (Hypertext Transfer Protocol Secure), WebSockets, and gRPC. These protocols create encrypted channels which protect data from unauthorized interception during transfer³⁰.

Moreover, secure API interactions are crucial to deploying the AI because they help ensure any data that goes between applications and the AI service is authenticated and restricted for access.

To protect against risks of data interception, data transit encryption frameworks like TLS (Transport Layer Security) and end-to-end encryption are largely adopted. These standards protect AI interactions by keeping transferred data

²⁸ Rani, P., Kavita, Verma, S., Kaur, N., Wozniak, M., Shafi, J., & Ijaz, M. F. (2022). Robust and Secure Data Transmission Using Artificial Intelligence Techniques in Ad-Hoc Networks. *Sensors*, 22(1), 251. <https://doi.org/10.3390/s22010251>

²⁹ Allam, Karthik. (2023). Adoption of Artificial Intelligence in Cloud Computing. *International Journal of Computer Trends and Technology*. 71. 91-95. 10.14445/22312803/IJCTT-V71I6P116.

³⁰ Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>

unreadable to third parties. Certain AI models additionally exploit homomorphic encryption, which is a cryptographic method that enables calculations on encrypted data without needing it to be decrypted, thus increasing security.

2.3.3. Risks and Mitigation Strategies in AI Data Processing

The processing of data in cloud-based AI services introduces various security risks that organizations must address to maintain compliance and protect user privacy. One of the primary concerns is **unintended data retention**, where AI systems inadvertently store input data beyond the necessary duration. Given strict data protection regulations such as GDPR, this is a compliance risk. Cloud AI reference providers restrict retention, stateless processing of AI queries and responses beyond the immediate lifecycle³¹.

Data minimization is also a strategy employed by organizations, whereby only the information needed to enhance AI models is retained while all data that is excess and does not add further value is disposed of. Model inversion attacks, in which an attacker tries to recreate the original input data from the AI's generated output, are another major threat. This attack uses the patterns in the model predictions to learn sensitive information about the training data. To protect against model inversion, adversarial training, differential privacy techniques, and controlled model access can be employed, making it difficult for outside users to extract significant information from AI models³². A third and final concern in AI data processing is cross-tenant data leakage, which appears in shared cloud environments where disparate organizations utilize a common AI infrastructure. Improperly configured access controls or multi-tenancy vulnerabilities may lead to unauthorized access of sensitive data for other clients. Cloud AI providers mitigate this risk using data isolation technologies like role-based access control (RBAC) and attribute-based access control (ABAC) that ensure the separation of data between tenants.

³¹ van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241232630>

³² Rani, P., Kavita, Verma, S., Kaur, N., Wozniak, M., Shafi, J., & Ijaz, M. F. (2022). Robust and Secure Data Transmission Using Artificial Intelligence Techniques in Ad-Hoc Networks. *Sensors*, 22(1), 251. <https://doi.org/10.3390/s22010251>

Furthermore, there are containerized AI environments like Kubernetes, allowing for the execution of an AI model in isolation as a container, preventing cross-tenant interference³³.

2.4. Retrieval-Augmented Generation (RAG) as a Privacy-Enhancing Technology (PET)

As enterprise artificial intelligence systems continue to evolve, balancing data accessibility and privacy remains a key challenge. Retrieval-Augmented Generation (RAG) has emerged as a technique for improving the accuracy and relevance of AI-generated outputs while reducing reliance on sensitive training data. Unlike conventional language models that rely solely on pre-trained knowledge, RAG dynamically retrieves relevant documents or structured data from external sources to generate responses grounded in verifiable information. This method is particularly beneficial in environments requiring factual accuracy and up-to-date information, as it allows AI to reference external knowledge bases rather than depending entirely on memorized data. The adoption of RAG is increasingly recognized as a privacy-enhancing strategy since it reduces the need for long-term data retention, minimizing risks associated with model memorization and data leakage³⁴. However, while RAG strengthens data control, it also introduces privacy risks related to retrieval security, information exposure, and vulnerabilities in third-party data repositories³⁵.

2.4.1. Introduction to Retrieval-Augmented Generation (RAG)

RAG is a hybrid AI architecture that combines retrieval-based information sources with generative AI models to create more accurate and contextually grounded responses. Unlike traditional models, which derive their content generation purely from pre-trained parameters, RAG proactively retrieves

³³ Allam, Karthik. (2023). Adoption of Artificial Intelligence in Cloud Computing. *International Journal of Computer Trends and Technology*. 71. 91-95. 10.14445/22312803/IJCTT-V71I6P116.

³⁴ Koga, Tatsuki & Wu, Ruihan & Chaudhuri, Kamalika. (2024). Privacy-Preserving Retrieval Augmented Generation with Differential Privacy. 10.48550/arXiv.2412.04697.

³⁵ Anderson, M., Amit, G., & Goldsteen, A. (2024). Is my data in your retrieval database? membership inference attacks against retrieval augmented generation. *arXiv preprint arXiv:2405.20446*.

externally relevant datasets for its functioning, allowing it to adapt to the nature of real-time knowledge extraction. This methodology notably enhances factual accuracy and reduces the likelihood of AI hallucinations, a phenomenon where models generate plausible but inaccurate content. With a retrieval step, RAG makes sure that the responses are grounded by real and reliable sources and prevents speculative outputs, making the AI-generated text more auditable and explainable³⁶. Also, the retrieval mechanism enables AI systems to better cope with new knowledge sources in enterprises, as knowledge source retraining can be done without a full retraining process, making an enterprise system more efficient in operations and response.

2.4.2. Privacy Benefits of RAG

RAG has a major privacy advantage where it does not require sensitive user data to be embedded within model weights. Large-scale traditional AI (especially deep learning) models naturally learn from extensive amounts of training data, increasing the probability that sensitive, proprietary, or insider information may be leaked in their outputs. This approach contrasts with RAG, which retrieves relevant data at generation, eliminating sensitive inputs persisting in sources from which adversaries can extract data via attacks³⁷. Furthermore, because RAG pulls in information only when necessary, it suits the vacated requirement for information minimization showcased in several privacy regulations, including the GDPR. Such a design preserves the interaction of users with the model from being part of a continuous learning chain of the model, preventing the model from memorizing any user-inputted information (personal or corporate)³⁸. A second privacy advantage of RAG is to provide transparency for AI-generated responses. Unlike traditional generative models, which function as a black box system, the RAG service gives users the ability to trace what it

³⁶ Zeng, S., Zhang, J., He, P., Liu, Y., Xing, Y., Xu, H., Ren, J., Chang, Y., Wang, S., Yin, D., & Tang, J. (2024). "The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG)." Findings of the Association for Computational Linguistics: ACL 2024, 4505–4524.

³⁷ Zeng, S., Zhang, J., He, P., Liu, Y., Xing, Y., Xu, H., Ren, J., Chang, Y., Wang, S., Yin, D., & Tang, J. (2024). "The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG)." Findings of the Association for Computational Linguistics: ACL 2024, 4505–4524.

³⁸ Koga, Tatsuki & Wu, Ruihan & Chaudhuri, Kamalika. (2024). Privacy-Preserving Retrieval Augmented Generation with Differential Privacy. 10.48550/arXiv.2412.04697.

produces back to the sources of information that were included in its results, providing greater interpretability and explainability in AI-driven decision-making. This capability is especially beneficial in regulated sectors like finance, healthcare, and legal services, where organizations need to prove compliance with data protection regulations and be responsible for the insights generated by AI³⁹. This allows enterprises to have a lot more control over the data provenance, generating content only if their generative AI systems are drawing from pre-approved knowledge bases with retrieval mechanisms.

2.4.3. Challenges and Risks of RAG

While RAG has privacy-enhancing qualities, it also poses security concerns, especially about the reliability of the retrieval sources. Because RAG models draw on knowledge repositories, the quality and neutrality of these sources affect the integrity of generated outputs. When the retrieval corpus is plagued with misinformation or biases, the AI might propagate incorrect or biased outputs without intention, presenting ethical challenges related to fairness and accuracy⁴⁰. Maintaining curation, openness, and adherence to organization policies for retrieval sources is still a major challenge in deploying RAG-based AI systems. Another restriction of RAG Rise is latency and calculation overhead. However, the retrieval step adds extra computation time and consequently increases AI inference time compared to fully pre-trained versions. For example, in real-time applications, such as automated financial trading or security monitoring, delays in AI response time can have an effect on decision-making effectiveness.

RAG generating organizations must therefore find a balance between the advantages of retrieval-based augmentation and the need for rapid, scalable operations of the AI⁴¹. RAG also poses security risks when it retrieves data from third-party sources. For AI systems whose processes pull from external

³⁹ Zeng, S., Zhang, J., He, P., Liu, Y., Xing, Y., Xu, H., Ren, J., Chang, Y., Wang, S., Yin, D., & Tang, J. (2024). "The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG)." Findings of the Association for Computational Linguistics: ACL 2024, 4505–4524.

⁴⁰ Zeng, S., Zhang, J., He, P., Liu, Y., Xing, Y., Xu, H., Ren, J., Chang, Y., Wang, S., Yin, D., & Tang, J. (2024). "The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG)." Findings of the Association for Computational Linguistics: ACL 2024, 4505–4524.

⁴¹ Koga, Tatsuki & Wu, Ruihan & Chaudhuri, Kamalika. (2024). Privacy-Preserving Retrieval Augmented Generation with Differential Privacy. 10.48550/arXiv.2412.04697.

repositories, sensitive queries used in the enterprise could be seen by untrusted data providers, posing information leakage risks. Evidence exists that attackers reach to know if asked content in a retrieval database⁴². To address these risks, enterprises should enact robust access controls and ensure retrieval sources are secure and authorized for use in AI-driven workflows.

2.4.4. Use Cases of RAG in Enterprise AI

RAG is widely used in enterprise applications where factual accuracy, privacy, and transparency are essential. AI service providers such as Microsoft Azure OpenAI Service, Google Cloud AI, and Amazon AWS AI implement RAG to enable real-time knowledge retrieval across corporate databases without embedding sensitive data into model training. This approach is particularly valuable in customer service automation, where AI chatbots dynamically access company knowledge bases to provide accurate, up-to-date responses while minimizing the risk of AI hallucinations⁴³. Within legal and compliance research, RAG-based AI tools retrieve relevant laws, policies, and case law, tailoring the output of AI according to existing regulatory standards. This reduces dependency on old training data and enhances the verifiability of AI-generated legal summaries. Likewise, in enterprise AI copilots, researchers use RAG to improve document search, workflow automation, and knowledge management by extracting critical insights from structured and unstructured internal data sources. This enables workers to quickly summarize reports and distill important information, all while maintaining proprietary business content within enterprise security boundaries⁴⁴. With the rise of RAG, security issues in retrieval are popping up for AI providers who must ensure that any enterprise AI applications are operating under the same constraints of access controls and governance frameworks. RAG enables a privacy-preserving approach to enterprise AI, allowing organizations to reference the most up-to-date knowledge without needing to store sensitive business data in model weights. However, the

⁴² Anderson, M., Amit, G., & Goldstein, A. (2024). Is my data in your retrieval database? membership inference attacks against retrieval augmented generation. *arXiv preprint arXiv:2405.20446*.

⁴³ Zeng, S., Zhang, J., He, P., Liu, Y., Xing, Y., Xu, H., Ren, J., Chang, Y., Wang, S., Yin, D., & Tang, J. (2024). "The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG)." Findings of the Association for Computational Linguistics: ACL 2024, 4505–4524.

⁴⁴ Koga, Tatsuki & Wu, Ruihan & Chaudhuri, Kamalika. (2024). Privacy-Preserving Retrieval Augmented Generation with Differential Privacy. 10.48550/arXiv.2412.04697.

implementation of it has to be controlled very well to ensure that no unauthorized data will be exposed and the corporate security policies will be followed.

2.5. Contractual and Regulatory Framework Governing AI Data Interactions

The deployment of AI technologies in enterprise environments necessitates a comprehensive regulatory and contractual framework to ensure compliance with data protection laws, ethical AI principles, and evolving governance standards. Microsoft 365 Copilot operates within a legal landscape shaped by the EU AI Act, GDPR, and various global AI governance guidelines, which impose obligations on AI providers regarding transparency, risk management, and data protection. This section explores the legal instruments that regulate AI-driven data processing, focusing on the contractual agreements between Microsoft and enterprise clients, regulatory compliance challenges, and the broader implications for AI governance.

2.5.1. Overview of AI Governance Frameworks

The rapid advancement of AI has necessitated the development of robust governance frameworks to ensure its ethical and secure deployment. Central to these efforts are the European Union's Artificial Intelligence Act (AI Act), the General Data Protection Regulation (GDPR), and various global AI ethics guidelines. The AI Act seeks to establish harmonized rules for AI within the European Union, aiming to ensure that AI systems are safe and respect existing laws on fundamental rights and Union values. It introduces a risk-based classification system, categorizing AI applications based on their potential impact. Prohibited practices include AI systems that deploy subliminal techniques beyond a person's consciousness to materially distort behavior in a manner that causes or is likely to cause physical or psychological harm. High-risk AI systems, such as those used in critical infrastructure, education, or employment, are subject to strict obligations before they can be placed on the market. These obligations encompass rigorous risk assessments, data governance measures, and human oversight mechanisms. The Act also emphasizes transparency, mandating that users be informed

when they are interacting with an AI system, unless this is obvious from the circumstances and the context of use⁴⁵. By setting these standards, the AI Act aims to foster public trust and acceptance of AI technologies. The GDPR, effective since May 2018, serves as a cornerstone for data protection within the EU. It stipulates that personal data must be processed lawfully, fairly, and in a transparent manner in relation to the data subject. Individuals are granted rights such as access to their data, rectification of inaccuracies, and the erasure of data under certain conditions. For AI systems that process personal data, compliance with GDPR is imperative. This includes ensuring data minimization, purpose limitation, and obtaining explicit consent when required. The regulation also addresses automated decision-making, granting individuals the right not to be subject to decisions based solely on automated processing, including profiling, which produces legal effects concerning them or similarly significantly affects them⁴⁶. Thus, the GDPR plays a pivotal role in safeguarding individual privacy in the realm of AI. Beyond legislative measures, ethical guidelines provide a foundational framework for the responsible development and deployment of AI. The EU's Ethics Guidelines for Trustworthy AI, formulated by the High-Level Expert Group on AI, assert that trustworthy AI should be lawful, ethical, and robust. The guidelines delineate seven key requirements: (1) Human agency and oversight: AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights. (2) Technical robustness and safety: AI systems need to be resilient and secure, ensuring a fallback plan in case something goes wrong. (3) Privacy and data governance: Full respect for privacy and data protection must be ensured, along with adequate data governance mechanisms. (4) Transparency: The data, system, and AI business models should be transparent, with traceability mechanisms and explainability of AI decisions. (5) Diversity, non-discrimination, and fairness: Unfair bias must be avoided, ensuring accessibility to all and involving relevant stakeholders throughout the AI system's lifecycle. (6) Societal and environmental well-being: AI systems should benefit all

⁴⁵ European Commission. (2021). Proposal for a Regulation laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. Brussels, Belgium. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>

⁴⁶ European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation - GDPR). Official Journal of the European Union, L 119, 1–88. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

human beings, including future generations, and be sustainable and environmentally friendly. (7) Accountability: Mechanisms should be in place to ensure responsibility and accountability for AI systems and their outcomes⁴⁷. These guidelines serve as a compass for developers and users alike, promoting AI that aligns with societal values and ethical principles. In summary, the confluence of the AI Act, GDPR, and global ethical guidelines establishes a comprehensive governance framework. This framework seeks to balance innovation with fundamental rights, ensuring that AI technologies are developed and utilized in a manner that is both responsible and aligned with societal values.

2.5.2. Compliance Challenges in AI Data Processing

The intersection of AI technologies with existing data protection laws presents several challenges. Ambiguities arise due to the dynamic nature of AI, where traditional data processing definitions may not adequately cover AI's capabilities, leading to potential compliance gaps. The GDPR, for instance, was enacted before the widespread adoption of AI, resulting in interpretative uncertainties regarding automated decision-making and profiling⁴⁸. Although laws such as the EU AI Act and GDPR are designed to protect individual rights and promote ethical AI use, they can also stifle innovation. The challenge is that strict compliance requirements can slow the development of artificial intelligence, forcing a painful cost-benefit analysis between innovation and regulation (Article 42, Artificial Intelligence Act). The AI Act imposes strict conditions to be met for these high-risk AI systems, including a burden for enterprises to collect and provide adequate test data to assess conformity⁴⁹. Enterprises need to adopt strong data governance frameworks, regularly conduct compliance audits, and continuously educate stakeholders to deal with these challenges. Moreover, it is imperative that data

⁴⁷ European Commission. (2021). Ethics guidelines for trustworthy AI. High-Level Expert Group on Artificial Intelligence (HLEG). Retrieved from <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

⁴⁸ European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation - GDPR). Official Journal of the European Union, L 119, 1–88. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

⁴⁹ European Commission. (2021). Proposal for a Regulation laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. Brussels, Belgium. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>

quality and integrity are high, as expressed in the Ethics Guidelines for Trustworthy AI⁵⁰. Routine evaluations can help uncover possible weaknesses and ensure that AI systems comply with new legal benchmarks. Providers of high-risk AI systems must have post-market monitoring systems in place to ensure ongoing compliance⁵¹. Communicating regulatory changes and their impact allows for informed compliance and responsible use of AI tools by all parties involved. Ensure participation and pluralism: during the entire lifecycle of any AI system, stakeholders should be involved in their design and implementation to take in different perspectives⁵². By tackling these outliers proactively, organizations can still leverage the gains afforded by AI innovation while balancing the need to be compliant with the regulators, thereby aligning responsible and lawful AI use.

⁵⁰ European Commission. (2021). Ethics guidelines for trustworthy AI. High-Level Expert Group on Artificial Intelligence (HLEG). Retrieved from <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

⁵¹ European Commission. (2021). Proposal for a Regulation laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. Brussels, Belgium. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>

⁵² European Commission. (2021). Ethics guidelines for trustworthy AI. High-Level Expert Group on Artificial Intelligence (HLEG). Retrieved from <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

Chapter 2: Analyzing the first research question – Does Microsoft, as the provider of Copilot, indirectly transfer user data to OpenAI?

3.1. The Architecture of Microsoft 365 Copilot and OpenAI's Role

The integration of advanced artificial intelligence technologies within enterprise applications represents a significant evolution in workplace productivity. Understanding the architecture of Microsoft 365 Copilot requires examining the longstanding partnership between Microsoft and OpenAI, which began with Microsoft's substantial \$1 billion investment in OpenAI in 2019⁵³, aimed at supporting OpenAI's AI research and securing exclusive cloud hosting rights for OpenAI's models on Azure. This partnership allowed for rapid innovation, as Microsoft received exclusive access to cutting-edge AI technology, including OpenAI's GPT models. Microsoft 365 Copilot lies at the heart of this evolution after integrating OpenAI's GPT-4 Large Language Model (LLM) into Microsoft's ecosystem to create an environment that is much more pro-user in terms of productivity and speed operations⁵⁴. A thorough understanding of this architecture is essential to assess whether, and under what conditions, indirect data transfers could occur between Microsoft and OpenAI, particularly within the stringent frameworks of GDPR and the AI Act.

3.1.1. Technical Integration of Copilot and OpenAI

Microsoft 365 Copilot is integrated into the Microsoft 365 ecosystem through Azure OpenAI Service, a dedicated infrastructure that ensures adherence to enterprise-grade security and compliance standards. OpenAI is the company bringing the sophisticated generative AI capabilities via its GPT-4 model and Microsoft partners with the company to build those into an enterprise use model. Microsoft Graph serves as the

⁵³ Microsoft. (2019, July 22). *Microsoft invests in and partners with OpenAI to support us building beneficial AGI*. <https://news.microsoft.com/2019/07/22/microsoft-invests-in-and-partners-with-openai-to-support-us-building-beneficial-agi/>

⁵⁴ Microsoft. (2025, February 13). *Microsoft 365 Copilot overview*. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-overview>

central store, aggregating organizational data- such as emails, documents, meetings and communication- to provide context for AI interactions and responses⁵⁵. Copilot is completely separate from OpenAI's public models, like ChatGPT, because it is deployed behind Azure OpenAI Service, configured for enterprise environments. This approach ensures compliance with strict enterprise data governance policies and regulatory requirements, especially GDPR and the AI Act⁵⁶. At the core of Copilot's integration in this workflow is the application of Microsoft Graph to get and enhance user queries to form a grounded prompt. This grounded prompt is then processed by GPT-4 within the Azure OpenAI Service infrastructure, following strict data security and governance practices⁵⁷. Microsoft documentation highlights that no data gets shared outside of Microsoft nor gets used for training models, which reinforces strong commitments to data containment and privacy⁵⁸. While the clear boundaries above exist, we arrive at another consideration on how inferences are handled and data might either in-directly pass through an inference processing pipe-line: While Microsoft does say unequivocally that OpenAI does not have access to enterprise data, a detailed examination of the architecture and details of how data moves through Copilot would be required to confirm that there are no potential avenues where the data could be exposed, even if only when a model is being inferred. It's important to stress that Microsoft is the Data Controller, accountable for compliance obligations under the GDPR and AI Act, to handling securely, protecting and enabling compliance for all data processed in Copilot. Microsoft has responsibility for the entirety of enterprise data processing including security and compliance, and transparency responsibilities about AI-driven interactions⁵⁹.

⁵⁵ Microsoft. (2024, March 6). *Azure OpenAI Service powers the Microsoft Copilot ecosystem*. Microsoft Azure Blog. <https://azure.microsoft.com/en-us/blog/azure-openai-service-powers-the-microsoft-copilot-ecosystem/#:~:text=Microsoft%20originally%20introduced%20the%20concept,Edge%2C%20Microsoft%20365%2C%20and%20Windows>

⁵⁶ Microsoft. (2025, January 15). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁵⁷ Microsoft. (2025, March 6). *Semantic indexing for Microsoft 365 Copilot*. Microsoft Learn. <https://learn.microsoft.com/en-us/microsoftsearch/semantic-index-for-copilot>

⁵⁸ Microsoft. (2025, January 15). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁵⁹ OpenAI. (n.d.). *Data processing addendum*. OpenAI. Retrieved March 8, 2025, from <https://openai.com/policies/data-processing-addendum/>

3.1.2. Data Flow in Copilot's AI Processing

The processing of user queries in Microsoft 365 Copilot involves several clearly defined steps, beginning when a user initiates a query or task in a Microsoft 365 application. First, Copilot accesses the Microsoft Graph for the exact enterprise data the user requested, by using advanced semantic indexing methods to ensure the accuracy and relevance of the retrieved data⁶⁰. The returned set of enterprise data is merged with the user's prompt to provide a grounded prompt with added context data which is then sent to the AI model for processing. Once the data is retrieved and grounded, the enriched prompt is sent to the Azure OpenAI Service. Microsoft explains that this inference are run inside the Azure Security environment by secure zero-retention policies meaning that neither the prompts nor the generated outputs are retained beyond the immediate processing session⁶¹. So once the GPT-4 model comes up with the response, that goes back to Copilot, and then there's some processing that has to happen for it to be actually delivered in a compliant, contextual way back to the user. However, despite clear statements in official documentation regarding data containment, there remains a potential point of ambiguity related to whether and how OpenAI's proprietary models indirectly interact with enterprise data during this inference phase. With the intricate relationships between Microsoft's own infrastructure and OpenAI's generative capabilities, some specific questions regarding where transient data might be exposed come into play, especially in light of regulatory scrutiny under the GDPR and the AI Act⁶². Comprehending these subtleties is essential for evaluating if Microsoft's safeguards are sufficient to completely eradicate all possible indirect interaction of data with OpenAI, or if there still exist oblique points of contention in this integration.

⁶⁰ Microsoft. (2025, March 6). Semantic indexing for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/microsoftsearch/semantic-index-for-copilot>

⁶¹ Microsoft. (2025, January 15). Data, privacy, and security for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁶² *ibidem*

3.2. Microsoft Azure OpenAI Service, Retrieval-Augmented Generation (RAG), and Enterprise Data Processing

This section examines how Microsoft Azure OpenAI Service facilitates enterprise AI deployment, focusing on data handling, retrieval-augmented generation (RAG), and transient data exposure during inference. It explores how enterprise data is processed within Copilot, the security measures in place, and the potential risks associated with real-time AI interactions.

3.2.1. Microsoft Azure OpenAI Service and Enterprise Data Handling

The Azure OpenAI Service is designed to enable secure integration of OpenAI's models into business environments with its hosting on Microsoft infrastructure. By running those models only on Azure, Microsoft makes it possible to keep enterprises safe from interaction with OpenAI's public AI services. The service runs in a logically isolated environment and enforces role-based access control, data encryption, and jurisdictional data boundaries to allow global compliance with data protection laws. Make sure that you're processing data only in regions that correspond to regions that satisfy regulatory cross-border data flow requirements, such as GDPR, using Azure OpenAI Data Zones⁶³. One of the core attributes of the Azure OpenAI Service is the zero retention, meaning that a user's prompts, AI generated answers, and embeddings will not be remembered once that immediate inference session ends. This ensures that prompts, responses generated by the AI model, and embeddings themselves are not retained beyond the duration of the inference session, greatly reducing the potential of any data persistence, access, or misuse for AI model refinement⁶⁴. Nonetheless, with real-time AI ideation adding layers of complication, the question stands how exposure of data in temporary creation of inference should be vehemently interpretable against known compliance postures. Although Microsoft upholds tight governance over the flow of data into and interaction with models, it is not necessarily clear to what extent these leave the company exposed to the risks of short-lived or ephemeral copies of

⁶³ ibidem

⁶⁴ Sweetman, S. (2024, September 24). *Enterprise trust in Azure OpenAI Service strengthened with Data Zones* azure.microsoft.com. Microsoft Azure Blog.

data⁶⁵. Microsoft, as the Data Controller under GDPR, holds the primary responsibility for ensuring that enterprise data remains contained within its infrastructure. This role includes enforcing compliance policies that govern data residency, security configurations, and enterprise-specific processing controls. The need for continuous oversight is particularly relevant in scenarios where AI models interact with structured enterprise queries, requiring careful evaluation of the extent to which Microsoft's governance measures mitigate risks of data exposure during processing⁶⁶.

3.2.2. Retrieval-Augmented Generation (RAG) in Microsoft 365 Copilot

Retrieval-Augmented Generation (RAG) is a key mechanism in Microsoft 365 Copilot that enhances AI-generated responses by dynamically retrieving enterprise-specific information from Microsoft Graph. This allows Copilot to ground its answers in the knowledge of an organization, so that the AI outputs can be based on current, contextually relevant information rather than just relying on model knowledge that is static. RAG is especially beneficial in enterprise setups where AI workflows are required to adapt with new datasets and internal knowledge sources⁶⁷. Some of the main advantages of RAG include less memorization of AI model and lower exposure to accidental data retention and leakage. By dynamically retrieving enterprise data instead of embedding it into the model, Copilot can generate task-specific responses without storing sensitive information⁶⁸. Not only does this approach lead to more accurate AI, but it is also more aligned with data governance best practices that inform control over how information is retrieved and used. Users could benefit from being served real-time data, but it comes at the trade-off of exposing volatile data. Though Copilot operates within the secure Azure OpenAI Service provided by Microsoft, the interaction of enterprise data retrieved within Copilot and OpenAI's models during inference has raised questions. Microsoft uses heavy guardrails on data flow and model

⁶⁵ Zeng, S., Zhang, J., He, P., Xing, Y., Liu, Y., Xu, H., Ren, J., Wang, S., Yin, D., Chang, Y., & Tang, J. (2024). *The good and the bad: Exploring privacy issues in retrieval-augmented generation (RAG)*. Proceedings of the Association for Computational Linguistics (ACL 2024), 4505–4524. Retrieved from <https://github.com/psychology/RAG-privacy>

⁶⁶ Microsoft. (2025, January 15). Data, privacy, and security for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁶⁷ Microsoft. (n.d.). *Retrieval Augmented Generation (RAG) in Azure AI Search*. Retrieved from <https://learn.microsoft.com/en-us/azure/search/retrieval-augmented-generation-overview>

⁶⁸ Koga, T., Wu, R., & Chaudhuri, K. (2024). *Privacy-preserving retrieval augmented generation with differential privacy* (Preprint). University of California, San Diego. Retrieved from <https://arxiv.org/abs/2412.04697>

interaction in order to ensure that enterprise queries remain in its systems⁶⁹. But this needs a deeper analysis to check whether they lead to layers of vast complexity in AI enterprise applications, if instantaneous data retrieval and augmentation make way for it⁷⁰. To mitigate these concerns, Microsoft has implemented security policies regulating RAG's access to corporate data sources including query isolation, encryption, and enterprise access controls. As organizations continue to adopt AI-driven productivity tools, maintaining transparency in retrieval processes and ensuring clear governance over data augmentation mechanisms remains an important factor in mitigating potential risks⁷¹.

3.2.3. Risks and Compliance Challenges in Enterprise AI Data Processing

The integration of OpenAI's generative models into Microsoft 365 Copilot introduces compliance challenges that require ongoing evaluation. While Microsoft has applied stringent governance controls that manage enterprise AI interactions, real-time inference processing begs the question of how much scope and applicability current data protection law has⁷². Ensuring that AI-driven workflows align with enterprise security policies is critical, particularly when considering the implications of AI-assisted decision-making in business environments. To reduce the surface of attack from the inference of an AI model, Microsoft adopts methodologies, such as logical isolation, access control, and encryption to Azure OpenAI Service. However, in light of retrieval-based AI interactions gaining popularity, there are ongoing discussions on the importance of short-lived AI interactions in the context of formal data processing episodes⁷³. With Microsoft's zero-retention policy for Copilot interactions, organizations need to evaluate whether such security claims meet their enterprise data

⁶⁹ Microsoft. (2025, January 28). *How does Microsoft 365 Copilot work?* Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-architecture>

⁷⁰ Zeng, S., Zhang, J., He, P., Xing, Y., Liu, Y., Xu, H., Ren, J., Wang, S., Yin, D., Chang, Y., & Tang, J. (2024). The good and the bad: Exploring privacy issues in retrieval-augmented generation (RAG). Proceedings of the Association for Computational Linguistics (ACL 2024), 4505–4524. Retrieved from <https://github.com/psychology/RAG-privacy>

⁷¹ Microsoft. (2025, January 15). Data, privacy, and security for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁷² *ibidem*

⁷³ Zeng, S., Zhang, J., He, P., Xing, Y., Liu, Y., Xu, H., Ren, J., Wang, S., Yin, D., Chang, Y., & Tang, J. (2024). The good and the bad: Exploring privacy issues in retrieval-augmented generation (RAG). Proceedings of the Association for Computational Linguistics (ACL 2024), 4505–4524. Retrieved from <https://github.com/psychology/RAG-privacy>

governance and compliance requirements⁷⁴. A related issue is how enterprise-specific interactions with Copilot could taint model outputs over time. While Microsoft maintains that Copilot customers' data is not used for model training, it remains unclear how structured inputs are interacting with OpenAI's models during inference, given ongoing AI response refinement in the enterprise context. While Microsoft's security framework is designed to restrict external data access, further clarity is needed to ensure that enterprise AI workflows maintain a clear distinction between inference-based interactions and long-term model adaptation⁷⁵. To help address these threats, Microsoft has bolstered its enterprise AI governance approach with contractual protections, access controls, and enterprise-class compliance monitoring. However, as LLM technology and regulatory guidelines evolve, organizations must continue validating that their use of Microsoft 365 Copilot meets shifting legal and practical AI data processing expectations⁷⁶.

3.3. Transparency in Microsoft's Documentation and Compliance with GDPR Articles 12 to 15 and the AI Act

In this section, we look at Microsoft's transparency measures and regulatory compliance obligations relating to Copilot's AI processing. It investigates if Microsoft's disclosures are compliant with GDPR's transparency requirements, considers how Copilot would be classified under the AI Act, and assesses the effectiveness of Microsoft's AI governance strategies. While Microsoft asserts compliance with data protection principles, regulatory analyses reveal gaps in data processing disclosures, explainability, and oversight mechanisms, raising questions about whether Copilot's AI inference processes align with evolving legal frameworks.

⁷⁴ Microsoft. (2025, January 15). Data, privacy, and security for Microsoft 365 Copilot. Microsoft Learn. <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

⁷⁵ Koga, T., Wu, R., & Chaudhuri, K. (2024). Privacy-preserving retrieval augmented generation with differential privacy [Preprint]. arXiv. <https://arxiv.org/abs/2412.04697>

⁷⁶ Sweetman, S. (2024, September 24). *Enterprise trust in Azure OpenAI Service strengthened with Data Zones* azure.microsoft.com. Microsoft Azure Blog.

3.3.1. Transparency and AI Governance Under GDPR

The GDPR mandates that organizations provide data subjects with clear and comprehensive explanations of how their personal data is processed, particularly under Articles 12 to 15, which establish the right to be informed, access, rectify, and restrict personal data processing. Microsoft states that Copilot does not store enterprise data beyond its immediate processing session and does not use customer data for model training, emphasizing that all interactions occur within Azure's secure infrastructure⁷⁷. Nonetheless, the Data Protection Impact Assessment (DPIA) of Microsoft 365 Copilot raises questions in relation to Microsoft's disclosures, namely in relation to Required Service Data, which also includes metadata collected within the context of operations to monitor system performance. The DPIA shows that Microsoft does not provide a disaggregated view over what kind of metadata it collects, for which purpose, and for how long it stores it; that there is a difficulty for an enterprise to understand whether tracking of telemetry data is GDPR-compliant concerning requirements of purpose limitation and data minimization⁷⁸. Another concern is whether momentary inference processing by OpenAI models constitutes a regulated data transfer under GDPR. Although Microsoft maintains that Copilot's AI interactions stay within the enterprise environment, it does not directly address whether transitive inference by OpenAI's models amounts to data processing subject to supplemental regulatory safeguards⁷⁹. Additionally, Microsoft's Data Processing Agreement (DPA) lists Copilot as a data processor, so compliance risk does land on enterprise customers' shoulders. However, the DPA does not provide a clear description of the part that OpenAI plays in the inference process, and therefore it is difficult for both companies to draw any conclusions about whether the Generative AI outputs of Copilot are consistent with the GDPR transparency and accountability principles⁸⁰. To enhance compliance, Microsoft needs to establish what episodic inference processing is for its systems; make telemetry data retention more transparent; and allow enterprise users to have greater granular control over AI-generated outputs.

⁷⁷ Microsoft. (2025, January 15). *Data, Privacy, and Security for Microsoft 365 Copilot*. learn.microsoft.com. Microsoft Learn.

⁷⁸ Ministry of Justice and Security. (2024). *DPIA Microsoft 365 Copilot: Data Protection Impact Assessment on the Processing of Personal Data with Microsoft 365 Copilot*. Retrieved from <https://www.rijksoverheid.nl/jenv>

⁷⁹ Microsoft. (2025). *The EU AI Act: A Microsoft Overview*. Retrieved from Microsoft documentation

⁸⁰ Microsoft. (2024). *GDPR & Generative AI: A Guide for the Public Sector*. Retrieved from Microsoft documentation

3.3.2. AI Act Compliance and High-Risk AI Classification

The EU AI Act introduces a risk-based regulatory approach, imposing the most stringent obligations on high-risk AI systems, which are those that significantly impact fundamental rights, safety, and societal interests⁸¹. While Microsoft 365 Copilot is not designated as a high-risk AI system per se, its capabilities to generate content, handle corporate data, and support enterprise decision-making also introduce new concerns around transparency, explainability, and human oversight⁸². According to the AI Act, high-risk AI systems should make intelligible documentation available that describes the functioning of the AI models and the flow of data through the enterprise so that users can trust in the results generated by AI systems and can adjust and intervene where necessary. Microsoft has implemented traceability solutions like audit trails and sensitivity labels to increase the accountability around Copilot's outputs⁸³. Nonetheless, ongoing assessments by regulators highlight that Copilot's documentation fails to adequately clarify the way in which responses to enterprise queries are formatted through OpenAI's models, leaving OpenAI inference mechanisms uncertain regarding if future outputs are impacted at all from an immediate processing session⁸⁴. The AI Act also specifies that high-risk AI systems must introduce safety measures that ensure human oversight to address the risk of automation bias. More Microsoft describes Copilot as optimized for human-in-the-system behaviors where an AI-generated suggestion needs to be reviewed and refined before it is applied⁸⁵. However the DPIA of Microsoft 365 Copilot notes that enterprise users are not given sufficient guidance on how they can raise queries or override AI, leading to speculation over whether Copilot meets the explainability and risk mitigation requirements of the AI Act⁸⁶. Another important requirement of the AI Act is transparency in AI-driven decision-making. Microsoft has Transparency Notes

⁸¹ European Union. (2024). Artificial Intelligence Act (Regulation EU 2024/1689). <http://data.europa.eu/eli/reg/2024/1689/oj>

⁸² Microsoft. (2025, January). *The EU AI Act: A Microsoft overview*. <https://www.microsoft.com/en-us/ai/tools-practices>

⁸³ Microsoft. (2025). *Enterprise data protection in Microsoft 365 Copilot and Microsoft 365 Copilot Chat*. Retrieved from <https://learn.microsoft.com/en-us/copilot/microsoft-365/enterprise-data-protection>

⁸⁴ Microsoft. (2024). *GDPR & Generative AI: A Guide for the Public Sector*. Retrieved from Microsoft documentation

⁸⁵ Microsoft. (2025, January). *The EU AI Act: A Microsoft overview*. <https://www.microsoft.com/en-us/ai/tools-practices>

⁸⁶ Ministry of Justice and Security. (2024). *DPIA Microsoft 365 Copilot: Data Protection Impact Assessment on the Processing of Personal Data with Microsoft 365 Copilot*. Retrieved from <https://www.rijksoverheid.nl/jenv>

and governance tools, but their docs do not provide enough details about how Copilot fetches, processes, and integrates enterprise-specific data in its AI-generated responses⁸⁷. If Copilot fails to meet the AI Act's transparency principles, this could lead to risks of opacity in AI-generated decision-making, making it difficult for organizations to verify whether AI-generated outputs align with their corporate governance and regulatory compliance policies. To address these concerns, Microsoft should enhance its AI transparency measures by providing greater clarity on inference mechanisms, user control options, and OpenAI's involvement in enterprise AI processing.

3.3.3. Evaluating Microsoft's AI Governance Strategies

Microsoft's AI governance framework is built on transparency, compliance and risk management, and is consistent with the requirements of GDPR and AI Act. The company has undertaken various transparency initiatives, including Privacy Statements, Data Processing Agreements (DPAs), and Transparency Notes, in order to offer enterprise users insights into how their data is processed by Copilot, how AI inference is involved and how users can participate in control in the product⁸⁸. However, a comparison of Microsoft's governance practices and regulatory expectations reveals major disparities in disclosure, user explainability, and the oversight of AI. Where GDPR mandates clear explanations of how personal data is processed by organizations, the AI Act broadens these to include AI decision-making logic and to address potential risk. Microsoft claims that Copilot operates on enterprise data in a controlled environment, and that it does not use any customer data to train OpenAI's underlying models⁸⁹; however, the DPIA on Microsoft 365 Copilot reveals that Microsoft's documentation lacks detailed disclosures on telemetry data, OpenAI's role in inference processing, and the traceability of AI-generated outputs⁹⁰.

⁸⁷ Microsoft. (2024). *Data, privacy, and security for Azure OpenAI Service*. Retrieved from <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

⁸⁸ Microsoft. (2025). *Enterprise data protection in Microsoft 365 Copilot and Microsoft 365 Copilot Chat*. Retrieved from <https://learn.microsoft.com/en-us/copilot/microsoft-365/enterprise-data-protection>

⁸⁹ Microsoft. (2024). *Data, privacy, and security for Azure OpenAI Service*. Retrieved from <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

⁹⁰ Ministry of Justice and Security. (2024). *DPIA Microsoft 365 Copilot: Data Protection Impact Assessment on the Processing of Personal Data with Microsoft 365 Copilot*. Retrieved from <https://www.rijksoverheid.nl/jenv>

The AI Act further requires that providers of the AI allow users interrogate and challenge AI generated content, but the Copilot documentation guidance is also found to be inadequate on detailing how organizations can intervene in AI decision-making or, guard against automation bias⁹¹. In addition, Microsoft's approach to AI governance is also largely geared towards internal compliance rather than helping companies with external tools to track the transparency of AI inference. To bridge this gap, Microsoft should improve its AI documentation, including making more granular disclosures about the inference capabilities of Copilot, be more transparent about the OpenAI partnership and provide more control to enterprise customers over the AI-generated content. While Microsoft has made significant progress in AI compliance, continuous refinements are necessary to ensure that Copilot aligns with evolving transparency, accountability, and risk mitigation expectations under GDPR and the AI Act, while avoiding concerns related to real-time AI inference or model training, which will be examined in later sections.

3.4. Possible Temporary Exposure of Data to OpenAI: Legal Interpretation of Real-Time Processing Without Retention and Comparison with Case Law on Instantaneous Data Analysis

The increasing adoption of AI-driven enterprise tools such as Microsoft 365 Copilot introduces new challenges in data protection and regulatory compliance, particularly regarding real-time data processing and Retrieval-Augmented Generation (RAG). While Microsoft has implemented zero-retention policies and security measures, the temporary nature of AI inference raises concerns about momentary exposure of enterprise data during processing. GDPR's broad definition of data processing means that even transient interactions with personal data require compliance, necessitating clear regulatory guidance and oversight. This section examines the risks, legal precedents, and regulatory gaps associated with temporary data access in AI inference, focusing on how RAG impacts data protection obligations.

⁹¹ Microsoft. (2024). *GDPR & Generative AI: A Guide for the Public Sector*. Retrieved from Microsoft documentation

3.4.1. Understanding the Risk of Instantaneous Data Exposure

The use of real-time AI inference in Microsoft 365 Copilot, particularly through Retrieval-Augmented Generation (RAG), raises concerns about momentary data exposure during processing. RAG improves AI generated answers by dynamically fetching organizational relevant context from Microsoft Graph, anchoring the responses in organizational knowledge rather than just being based on pre-trained model's memory⁹². This decreases the risk of AI memorization, but it introduces another potential attack vector, since data pulled from enterprise sources will need to be temporarily held as it is processed to churn out a reply. GDPR defines processing broadly, covering retrieval, consultation, and use of personal data, regardless of storage duration⁹³. This means that AI inference via RAG, even if ephemeral, falls within GDPR's regulatory scope. The UK Information Commissioner's Office (ICO) has reinforced that even transient access to personal data constitutes processing, necessitating a lawful basis and compliance with transparency obligations⁹⁴. Microsoft asserts that Azure OpenAI Service is stateless, ensuring that no prompts, enterprise data, or generated outputs are retained beyond the immediate processing session⁹⁵. Also, because Copilot is built on RAG-based retrieval, enterprise data isn't injected into the AI model but is queried only on-the-fly, for each question. Nevertheless, even with these guarantees, brief access to enterprise data at inference is not risk-free. Even though RAG eliminates data retention risks associated with long-term storage, it still requires data to be loaded into system memory (RAM/GPU memory) during inference⁹⁶. Even though the AI model itself does not store this data, the underlying cloud infrastructure supporting Copilot's inference process may temporarily cache data, generate logs, or perform real-time filtering for security and abuse detection⁹⁷. For organizations using Copilot and Azure OpenAI Service, this means that even ephemeral RAG-based interactions should be treated as GDPR-regulated processing.

⁹² Microsoft. (n.d.). *Retrieval Augmented Generation (RAG) in Azure AI Search*. Retrieved from <https://learn.microsoft.com/en-us/azure/search/retrieval-augmented-generation-overview>

⁹³ European Data Protection Supervisor, European Union Agency for Fundamental Rights, & Council of Europe. (2018). *Handbook on European Data Protection Law*.

⁹⁴ The UK Information Commissioner's Office (ICO) has reinforced that even transient access to personal data constitutes processing, necessitating a lawful basis and compliance with transparency obligations

⁹⁵ Microsoft. (2024). *Data, privacy, and security for Azure OpenAI Service*. Retrieved from <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

⁹⁶ van Binsbergen, L. T., Steketeer, M. C., Kebede, M. G., Janssen, H. L., & van Engers, T. M. (2025). Lawful and Accountable Personal Data Processing with GDPR-based Access and Usage Control in Distributed Systems. *arXiv preprint arXiv:2503.07172*.

⁹⁷ Jisc Involve. (n.d.). *National Centre for AI*.

Enterprises must assess the full data flow involved in retrieval, inference, and response generation, ensuring that contractual safeguards are in place and users are transparently informed about how AI-driven retrieval affects their data privacy. While zero-retention policies, encryption protocols, and strict access controls mitigate risks, the transient exposure of enterprise data via RAG retrieval still requires regulatory scrutiny.

3.4.2. Legal Precedents on Real-Time Data Processing Without Retention

Legal precedent confirms that even short-lived interactions with personal data qualify as processing under GDPR. A key case illustrating this is the Milan Central Station biometric advertising case (Provvedimento n. 551 del 21 dicembre 2017), in which the Italian Data Protection Authority (Garante per la protezione dei dati personali) ruled that an advertising system performing real-time facial analysis constituted personal data processing, despite no images being retained beyond a fraction of a second⁹⁸. The Milan advertising system used built-in cameras and facial detection software to estimate demographic attributes such as age and gender, adjusting displayed advertisements accordingly. The video stream has been processed in real-time and every frame has been erased, so no sequence has been memorized or stored; but the Garante held that this did in fact amount to processing of personal data, since the system processed biometric details, albeit for a fraction of time. The DPA stipulated that visible public notices should be used, that immediate frame deletion should be applied and processing should be limited to non-identifying features⁹⁹. This ruling applies to Copilot's RAG-based retrieval, where enterprise data is dynamically retrieved, processed in real-time, and discarded after response generation. While Copilot does not store or transmit identifiable user data, its retrieval, contextualization, and inference process still constitutes GDPR-regulated activity¹⁰⁰. This reading is confirmed by additional legal decisions, that even momentary access to personal data is considered processing. A fine by the Swedish Data Protection Authority against the use of live facial recognition by a school for school attendance management, even if no image of students was recorded, shows that processing is not in the clear when it's

⁹⁸ Garante per la Protezione dei Dati Personali. (2017). *Provvedimento n. 551 del 21 dicembre 2017*. ([Source: Italian DPA]).

⁹⁹ *ibidem*

¹⁰⁰ Microsoft. (2024). *Data, privacy, and security for Azure OpenAI Service*. Retrieved from <https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

transient¹⁰¹. These precedents confirm that even RAG-enhanced real-time AI inference must be treated as GDPR-regulated processing. While Microsoft enforces zero-retention policies and encryption, the retrieval process itself, particularly how long enterprise data remains in active memory before being discarded, warrants closer scrutiny. Organizations must assess whether Copilot's retrieval and inference process aligns with GDPR principles of transparency, necessity, and proportionality.

3.4.3. Regulatory Gaps in Defining AI Model Inference Risks

GDPR does not explicitly define how transient AI processing should be regulated, as the regulation was drafted before widespread enterprise AI adoption. While GDPR applies broadly to personal data processing, including real-time inference, regulators have only recently begun clarifying how these rules apply to momentary AI interactions that do not involve persistent storage¹⁰². According to the European Data Protection Board (EDPB) 2024 Opinion on AI and Data Protection, AI models cannot be assumed to be automatically anonymous just because they do not store data. Where there is a risk that personal data may be re-identified from the AI model outputs, the GDPR is triggered¹⁰³. This directly affects AI such as Copilot's RAG-based retrieval, which temporarily processes enterprise data to improve response accuracy, even though the content is not logged. To address these risks, AI governance frameworks must evolve to regulate momentary AI interactions, particularly in cloud-based enterprise tools. The **CNIL has proposed new privacy-by-design approaches**¹⁰⁴, advocating for:

- Greater transparency on AI training data sources, ensuring users are informed of potential data exposure risks.
- Privacy-enhancing techniques (PETs) such as output filtering, differential privacy, and secure multiparty computation to reduce risks during AI inference.

¹⁰¹ European Data Protection Board. (2019, August 22). *Facial recognition in school renders Sweden's first GDPR fine*. https://www.edpb.europa.eu/news/national-news/2019/facial-recognition-school-renders-swedens-first-gdpr-fine_sv

¹⁰² European Data Protection Board. (2024, December 17). *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*.

¹⁰³ *ibidem*

¹⁰⁴ Commission Nationale de l'Informatique et des Libertés (CNIL). (2025, February 7). *AI and GDPR: The CNIL publishes new recommendations to support responsible innovation*

- Mandatory Data Protection Impact Assessments (DPIAs) for AI deployments to assess transient data exposure risks before implementation.

However, the EU AI Act does not go far enough in addressing transient inference risks, and GDPR remains the principal applicable law for AI-based enterprise software applications¹⁰⁵. Responsible AI governance by Microsoft, such as the use of retrieval mechanisms based on RAG, access controls, and encryption methods, help reduce some of the risk, but enterprises need to be vigilant in staying compliant with contractual safeguards, transparency policies, and keeping an eye on the regulatory landscape. As AI inference models continue to advance, organizations need to remain vigilant in monitoring new regulatory clarifications all while implementing privacy-by-design protections in enterprise AI deployment.

3.5. OpenAI's Model Training and the Potential Indirect Use of Enterprise Data

As artificial intelligence systems become increasingly integrated into enterprise environments, concerns arise over how user interactions may influence the development of foundational models. And while Microsoft has put in place measures to prevent Copilot data being used to train OpenAI's own models, it remains unclear whether indirect learning effects could still take place. In particular, retrieval-augmented generation (RAG) presents issues of being able to infer what data was exposed during inference (I.e., are enterprise interactions contributing to OpenAI's global model evolution). This section examines the possibility of indirect data influence, the regulations on AI training and training limitations, and the compliancy inspired interventions to keep enterprise data warehousing from contaminating foundation models.

¹⁰⁵ Microsoft. (2025). *The EU AI Act: A Microsoft Overview*. Microsoft. Retrieved May 7, 2025, from <https://www.microsoft.com/en-us/security/blog/2024/02/22/announcing-microsofts-open-automation-framework-to-red-team-generative-ai-systems/>

3.5.1. Does OpenAI Indirectly Benefit From Copilot Data?

Microsoft says Copilot does not use any user data to train OpenAI's models and everything stays within the enterprise customer's environment, but there are still questions about whether OpenAI's models at the very bottom might benefit indirectly from seeing interactions in the enterprise. Even without retaining data, patterns of user engagement, language structures or enterprise-specific customizations could influence OpenAI's models, he cautioned. A notable point of concern is the Retrieval-Augmented Generation (RAG), which could let Copilot fetch enterprise-specific data from Microsoft Graph in real-time just before executing the response. Microsoft maintains the processing is entirely within the Azure OpenAI Service with no storage but it's this lack of clear visibility of what is done with data grabbed into the inference process that can lead to concerns over compliance and data protection. If enterprise data is used for the purpose of generating AI responses temporarily, it is unclear if OpenAI's systems could lend queries or data to this temporarily stored data during inference. The lack of technical details about how Microsoft prevents leakage of transient data is unclear as to whether RAG interactions still remain completely isolated from OpenAI's base models. Academic research on AI model refinement and data reuse suggests that even when models do not store user-provided data, repeated interactions can still influence reinforcement learning-based improvements¹⁰⁶. This is a concern of GDPR compliance, because the indirect model adaptation can be considered as an implicit type of training. Some legal scholars argue that machine learning models can derive statistical inferences from user interactions, making a complete separation between enterprise AI use and foundation model evolution difficult to guarantee¹⁰⁷. Although Microsoft has data containment policies such as zero-retention for Copilot interactions, it is uncertain whether RAG-based S/R could contribute insights to OpenAI's models. The EDPB has further stressed that AI providers must perform compatibility checks (Article 6(4) GDPR) before reusing company data in ways that may impact model training. This means any processing of this nature will be expected to align with what the user reasonably expects and with

¹⁰⁶ Baheri, A. (2023). Towards Theoretical Understanding of Data-Driven Policy Refinement. *arXiv preprint arXiv:2305.06796*.

¹⁰⁷ McIntosh, T. R., Susnjak, T., Liu, T., Watters, P., Xu, D., Liu, D., & Halgamuge, M. N. (2025). From Google Gemini to OpenAI Q* (Q-Star): A Survey on Reshaping the Generative Artificial Intelligence (AI) Research Landscape. *Technologies*, 13(2), 51. <https://doi.org/10.3390/technologies13020051>

the purpose limitation principle of GDPR¹⁰⁸. For as AI governance continues to develop, the implications of Copilot's on-the-fly parsing on OpenAI's larger AI ecosystem will be a key aspect of compliance and transparency.

3.5.2. GDPR's Purpose Limitation Principle and AI Training Risks

The purpose limitation principle under GDPR (Article 5(1)(b)) mandates that personal data must be processed strictly for its stated purpose and not repurposed without a new legal basis¹⁰⁹. This principle is especially important to apply with respect to AI systems, as base models change over time, it is not clear whether Copilot's enterprise interactions could inform the broader AI model improvement at OpenAI later even in the absence of explicit data retention. A case in point is the Italian Data Protection Authority's (Garante) 2023 decision against OpenAI's ChatGPT that held OpenAI didn't have a legal basis for training models on scraped personal data and didn't disclose how previous user input impacted the model's future behavior¹¹⁰. This decision makes clear that, if AI providers purport to not store user data, they still need to make the commitment that insights drawn from interactions will not be leveraged for purposes that would go against the purpose limitation principle under the GDPR. Further academic analysis suggests that AI inference itself may constitute a secondary use of data if its results influence model optimization, meaning that enterprise interactions with Copilot could still contribute to OpenAI's learning process¹¹¹. The EDPB has emphasized that any reuse of data for AI model training must pass a compatibility assessment (Article 6(4) GDPR), ensuring that data subjects can reasonably anticipate how their interactions may impact future AI model evolution¹¹². Despite Microsoft's claims that Copilot data is never used in training the OpenAI model, questions have been raised about retrieval-augmented generation (RAG) and how it may inform model outputs. As RAG dynamically ingests domain-specific data during inference, policymakers and enterprises need to consider whether these

¹⁰⁸ Data Protection Report. (2025, January). *The EDPB Opinion on training AI models using personal data and recent Garante fine – lawful deployment of LLMs*.

¹⁰⁹ European Data Protection Supervisor. (2018). *Handbook on European data protection law*. Publications Office of the European Union

¹¹⁰ Clifford Chance. (2023, April 1). *The Italian Data Protection Authority halts ChatGPT's data processing operations*.

¹¹¹ Zhu, Z., Li, Y., & Zhang, H. (2024). Optimizing large language models for OpenAPI code completion. *arXiv*.

¹¹² European Data Protection Supervisor. (2018). *Handbook on European data protection law*. Publications Office of the European Union

interactions establish implicit learning paths that then have an impact on OpenAI's base models. The changing regulatory environment underscores the importance of robust AI governance frameworks to ensure that enterprise AI deployments adhere to GDPR's purpose limitation requirements and to stop unintentional AI model enrichment.

3.5.3. Compliance Strategies to Ensure Data Containment in AI Systems

To avoid accidentally spreading enterprise data through into OpenAI's foundational models, Microsoft applied tight data containment policies that ensure data used by Copilot does not spill over into OpenAI's other AI training processes¹¹³. These include zero-retention policies that stop Copilot-generated prompts and responses from being retained, and encryption mechanisms that protect users' interactions within the enterprise-managed Microsoft 365 environment. Microsoft also has contractual provisions that prevent sharing data between enterprise customers and OpenAI without clear authorisation by the customer¹¹⁴. The Azure OpenAI Service, which powers Copilot's AI features, is a "bring-your-own-data" deployment of OpenAI's models, keeping enterprise content safe by never logging or sending it to OpenAI's servers. Microsoft clearly states that user prompts, completions, and interactions are "NOT provided to OpenAI" and "NOT used to train OpenAI's models"¹¹⁵. This architecture allows to ensure that there is no unintended feedback loop, where the data that makes OpenAI's model a great model is in fact proprietary data that the company owns. Data sequestration as a compliance criterion for providers of AI applications has been getting more consideration from regulatory organizations. The ICO's clear guidance emphasises that every step of processing of an AI model- from training, finetuning to inferencing - must be purpose-bound and that data cannot be 'repurposed' from one to the next, in the absence of a clear legal basis¹¹⁶. The EDPB also suggests that companies should be using enterprise versions of AI services and not use organizational personal data to train the provider's models when warning that personal data transfers for AI training cannot as a general rule be justified on the basis of

¹¹³ Microsoft. (2024). *Responsible AI Transparency Report*. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-brand/documents/responsible-ai-transparency-report-2024.pdf>

¹¹⁴ Microsoft. (2024, March 28). Protecting the data of our commercial and public sector customers in the AI era. *Microsoft On the Issues*. <https://blogs.microsoft.com/on-the-issues/2024/03/28/data-protection-responsible-ai-azure-copilot/>

¹¹⁵ Microsoft. (2025, April 15). *Network and access configuration for Azure OpenAI On Your Data*. <https://learn.microsoft.com/en-us/azure/ai-services/openai/how-to/use-your-data-securely>

¹¹⁶ Information Commissioner's Office. (2024). *Guide to the General Data Protection Regulation (GDPR)*.

‘legitimate interests’ under the GDPR¹¹⁷. In the future, AI providers will also probably be subject to additional regulations on the use of enterprise data across foundation models. The EU AI Act (in force in 2025) defines obligations for foundation model providers and AI deployers with additional transparency and data governance¹¹⁸. For example, in Article 28b it is required that AI providers carry out an assessment of risks to personal data in training data, that they take steps to mitigate such risks and to provide summaries of copyrighted or personal material that was used for training, or that they take privacy-by-design to account. Non-compliance with these evolving regulations could lead to significant financial penalties, with the AI Act proposing fines of up to 6-7% of a company’s global turnover for violations¹¹⁹. As AI governance frameworks develop, enterprises must ensure that their AI deployments align with GDPR’s purpose limitation principle, the AI Act’s transparency requirements, and broader compliance obligations. Strengthened contractual commitments, technical safeguards, and regulatory oversight will be essential in maintaining strict data isolation between enterprise AI use and foundation model training.

¹¹⁷ Nance, R., Evans, M., & Gelmetti, F. (2025, January 2). The EDPB Opinion on training AI models using personal data and recent Garante fine – lawful deployment of LLMs. *Data Protection Report*

¹¹⁸ Holistic AI. (2023, July 1). Regulating foundation models and generative AI: The EU AI Act. *Holistic AI*.

¹¹⁹ Nance, R., Evans, M., & Gelmetti, F. (2025, January 2). The EDPB Opinion on training AI models using personal data and recent Garante fine – lawful deployment of LLMs. *Data Protection Report*

Chapter 3: Analyzing the Second Research Question – Is Microsoft 365 Copilot a Tool for Productivity or Surveillance?

4.1. The Processing of Employee Data Through Copilot and AI's Role in Performance Monitoring

The above discussion about the role of Microsoft 365 Copilot in workplace analytics also relates to the previously discussed privacy concerns in Chapter 2, regarding the indirect transfer of enterprise data to OpenAI. Copilot was built to stay within Microsoft's secure infrastructure, but the technology's capability to access, summarize, and analyze large volumes of employee-generated content could challenge how these insights get applied inside an organization. Just as transparency and governance are key to mitigating AI processing's exposure risk, the same is true for Copilot's straddle of workplace monitoring. The idea that analytics driven by AI could go from a productivity-enhancing tool to a de facto surveillance tool requires a close look at how it's implemented, regulated, and that it complies with employment regulations.

4.1.1. Microsoft Productivity Score and Its Relevance to Copilot's Workplace Analytics

Microsoft Productivity Score was designed to offer insights into how employees use Microsoft 365 tools, tracking email activity, meeting participation, and collaboration patterns to help organizations optimize workflows¹²⁰. While its purpose was to support digital transformation, its capability to monitor individual employee behavior sparked concerns about workplace surveillance. Managers were initially able to receive detailed reports on employee activities, including the number of emails sent and time spent in meetings, but there were concerns it could force people to feel as if they had to stay constantly engaged¹²¹. Facing a shower of criticism, Microsoft

¹²⁰ Microsoft. (n.d.). *Adoption Score in Microsoft 365*. Retrieved June 23, 2025, from <https://learn.microsoft.com/en-gb/microsoft-365/admin/adoption-score?view=o365-worldwide>

¹²¹ Sandler, R. (2020, 26 novembre). *C'è un nuovo strumento Microsoft che controlla a distanza i lavoratori. Ed è bufera sulla privacy*. Forbes Italia. <https://forbes.it/2020/11/26/microsoft-productivity-score-il-nuovo-strumento-per-il-controllo-a-distanza-dei-lavoratori/>

removed the ability to track by user, showing data only in aggregate form. The company was quick to stress that this wasn't a tool to monitor the rank and file, but rather one for IT and leadership teams who want to optimize their business operations¹²². But the uproar over its Productivity Score is illustrative of a broader issue in AI-driven workplace analytics, even if a tool is not intended for surveillance, it may be perceived, or deployed, as one¹²³. Microsoft 365 Copilot is not another Productivity Score but rather a generative AI assistant as opposed to an analytics dashboard. But, because it's reading, summarizing your emails, meetings, and documents, people are still concerned that its findings might be used to indirectly assess your contributions as a monkey and an activity tracker¹²⁴. According to studies, AI-based workplace analytics may raise anxiety and opposition when employees are aware of being constantly watched¹²⁵. While Microsoft asserts that Copilot is designed for productivity, not monitoring, organizations must set clear governance policies to prevent its use for performance tracking¹²⁶.

4.1.2. Copilot's Role in Workplace Analytics: Productivity Enhancement or Indirect Monitoring?

Microsoft 365 Copilot integrates with Outlook, Teams, Word, and other applications, providing AI-powered assistance through Retrieval-Augmented Generation (RAG). It also introduces new levels of automation to tasks, such as writing emails, creating meeting minutes, and finding and pulling in the most up-to-date information into documents, delivering to save time and foster efficiency¹²⁷. Microsoft emphasizes that Copilot enhances productivity while operating within an organization's secure cloud environment without training on user data¹²⁸. Despite these advantages, there are lingering concerns as to whether Copilot's

¹²² Microsoft. (2020, December 1). *Our commitment to privacy in Microsoft Productivity Score*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2020/12/01/our-commitment-to-privacy-in-microsoft-productivity-score/>

¹²³ Carter, C. (2024). AI surveillance: Reclaiming privacy through informational control. *European Labour Law Journal*, 0(0). <https://doi.org/10.1177/20319525241306327>

¹²⁴ Microsoft. (2023, March 16). *Introducing Microsoft 365 Copilot*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot/>

¹²⁵ Zitek, E., & Brekken, K. (n.d.). *AI surveillance in the workplace linked to employee resistance, turnover*. SHRM.

¹²⁶ AI Now Institute, april 2023, *Algorithmic management: Restraining workplace surveillance*.

¹²⁷ Microsoft. (2023, March 16). *Introducing Microsoft 365 Copilot*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot/>

¹²⁸ Microsoft. (2025, May). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft Learn. Retrieved from <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>

summarization of talks and highlighting of important contributions would make it easier for employers to monitor their workers indirectly. Summaries processed by AI might indicate the degree of involvement in meetings or group projects and their influence on a manager's perception¹²⁹. While Copilot does not explicitly provide managers with engagement metrics, its insights could still be interpreted as a measure of individual contributions¹³⁰. The debate over AI-driven workplace analytics is not new. Microsoft's Productivity Score initially faced scrutiny for tracking employee activity at an individual level before shifting to aggregated insights¹³¹. Copilot does not offer direct dashboards for monitoring employee engagement, but the ability to process vast amounts of communication data could still raise concerns about implicit performance assessments¹³². Studies indicate that when employees feel monitored by AI, they experience stress and may alter their behavior to appear more engaged¹³³. This has been seen in labour settings that deploy AI-powered monitoring systems as workers change their communicative strategies or push back from the AI to evade being under surveillance¹³⁴. To address this, Microsoft and organizations using Copilot must establish clear governance policies, ensuring transparency and limiting AI-driven insights to productivity enhancement rather than performance evaluation¹³⁵.

4.1.3. Copilot and the Expansion of AI-Powered Performance Monitoring

AI-powered workplace analytics have evolved from simple productivity tracking to sophisticated algorithmic management tools that influence workflow optimization and performance evaluations¹³⁶. Microsoft 365 Copilot is part of this trend, integrating AI into enterprise applications to automate tasks and streamline

¹²⁹ Zitek, E., & Brekken, K. (n.d.). *AI surveillance in the workplace linked to employee resistance, turnover*. SHRM.

¹³⁰ AI Now Institute, april 2023, *Algorithmic management: Restraining workplace surveillance*.

¹³¹ Hern, A. (2020, November 26). *Microsoft productivity score feature criticised as workplace surveillance*. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2020/nov/26/microsoft-productivity-score-feature-criticised-workplace-surveillance>

¹³² Mantelero, A. (2022). *AI surveillance: Reclaiming privacy through informational control*. In S. Nemitz & A. L. R. Zanatta (Eds.), *AI and Democracy* (pp. 123–135). Springer. https://doi.org/10.1007/978-3-030-96185-0_8

¹³³ Davis, J. (2023, October 17). *AI surveillance in the workplace linked to employee resistance*.

SHRM. <https://www.shrm.org/topics-tools/news/employee-relations/ai-surveillance-in-the-workplace-linked-to-employee-resistance->

¹³⁴ Ball, K., & Francis, G. (2023). *Private eyes, they see your every move: Workplace surveillance and worker well-being*. University of Essex / European Commission. Retrieved from <https://www.essex.ac.uk/-/media/documents/research/centre-for-work-organization-and-society/private-eyes-they-see-your-every-move.pdf>

¹³⁵ Mantelero, A. (2022). *AI surveillance: Reclaiming privacy through informational control*

¹³⁶ AI Now Institute. (2023, April 11). *Algorithmic management: Restraining workplace surveillance*. <https://ainowinstitute.org/publication/algorithmic-management>

communication¹³⁷. Although it presented as a software tool for its possible efficiency, fears around its performance-monitoring application linger¹³⁸. The lesson learned from Microsoft's Productivity Score serves as a cautionary precedent. Initially designed to assess organizational adoption of Microsoft 365, it faced criticism for allowing managers to view granular employee activity data, prompting Microsoft to remove individual-level tracking¹³⁹. This history is relevant to Copilot, as its ability to summarize meetings and analyze workplace interactions could similarly be repurposed to evaluate employee engagement¹⁴⁰. The growth of algorithmic management adds to these fears. AI-powered programs are increasingly being used to set work schedules, rate productivity, and even hire or fire workers¹⁴¹. Amazon's use of AI to manage its warehouse workers' productivity serves as a case in point, workers have said they have been penalised for automatically-generated AI recommendations with little human oversight. There was also backlash around Microsoft's Productivity Score, which allowed for workplace surveillance before it was updated¹⁴². Employers claim AI-driven analytics improves efficiency and accountability; workers often see them as intrusive¹⁴³. A further example is that research indicates algorithmic tracking is linked to more stress and job dissatisfaction. Employees fear that AI insights could be used unfairly, leading to micromanagement or punitive measures¹⁴⁴. Regulatory initiatives are focused on reining in these risks. EU AI Act places workplace AI monitoring tools in "high-risk" category with rigorous transparency and oversight requirements. The unions have been looking for collective agreements to add proper limitations to how AI-based analytics are being used. These developments reflect growing recognition that AI's role in

¹³⁷ Microsoft. (2023, March 16). *Introducing Microsoft 365 Copilot*. <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot/>

¹³⁸ Davis, J. (2023, October 17). *AI surveillance in the workplace linked to employee resistance*. SHRM. <https://www.shrm.org/topics-tools/news/employee-relations/ai-surveillance-in-the-workplace-linked-to-employee-resistance->

¹³⁹ Hern, A. (2020, 26 novembre). *Microsoft productivity score feature criticised as workplace surveillance*. *The Guardian*. <https://www.theguardian.com/technology/2020/nov/26/microsoft-productivity-score-feature-criticised-workplace-surveillance>

¹⁴⁰ Carter, C. (2024). *AI surveillance: Reclaiming privacy through informational control*. *European Labour Law Journal*, 0(0). <https://doi.org/10.1177/20319525241306327>

¹⁴¹ AI Now Institute. (2023, April 11). *Algorithmic management: Restraining workplace surveillance*. <https://ainowinstitute.org/publication/algorithmic-management>

¹⁴² pataro, J. (2020, December 1). *Our commitment to privacy in Microsoft Productivity Score*. Microsoft 365 Blog. <https://www.microsoft.com/en-us/microsoft-365/blog/2020/12/01/our-commitment-to-privacy-in-microsoft-productivity-score/>

¹⁴³ Hern, A. (2020, 26 novembre). *Microsoft productivity score feature criticised as workplace surveillance*. *The Guardian*. <https://www.theguardian.com/technology/2020/nov/26/microsoft-productivity-score-feature-criticised-workplace-surveillance>

¹⁴⁴ Ball, K. (2021). *Private eyes, they see your every move: Workplace surveillance and worker well-being*. European Trade Union Institute (ETUI). <https://www.etui.org/publications/private-eyes-they-see-your-every-move>

workplace management must be carefully regulated to prevent privacy violations and protect worker autonomy. In conclusion, AI-powered performance monitoring presents both opportunities and risks. While tools like Copilot can streamline workflows and reduce administrative burdens, they must be implemented responsibly to avoid enabling workplace surveillance¹⁴⁵. Transparency, governance, and clearly defined policies are essential to ensuring that AI enhances productivity without infringing on employee rights¹⁴⁶.

4.2. Applicability of Article 4 of the Workers' Statute: Is Copilot a Necessary Work Tool or a Surveillance Mechanism?

4.2.1. Legal Framework of Article 4: Key Provisions and Protections

Article 4 of the Italian Workers' Statute regulates the use of workplace monitoring technologies, ensuring that employee surveillance is restricted to specific, justified purposes and subject to strict procedural safeguards. The provision establishes three key principles: the permissible use of surveillance technologies, the distinction between work tools and monitoring systems, and the requirement for transparency and worker notification.

1. Permissible Use of Surveillance Technologies

"Audiovisual equipment and other instruments from which also derives the possibility of remote monitoring of workers' activities may be used exclusively for organizational and production requirements, work safety and protection of company assets and may be installed subject to a collective agreement entered into by the unitary trade union representation or company trade union representatives. Alternatively, in the case of enterprises with production units located in different provinces of the same region or in several regions, such agreement may be entered into by the comparatively most representative trade union associations nationwide. In the absence of an agreement, the equipment and instruments referred to in the first sentence may be installed subject

¹⁴⁵ AI Now Institute. (2023, April 11). *Algorithmic management: Restraining workplace surveillance*. <https://ainowinstitute.org/publication/algorithmic-management>

¹⁴⁶ Carter, C. (2024). AI surveillance: Reclaiming privacy through informational control. *European Labour Law Journal*, 0(0). <https://doi.org/10.1177/20319525241306327>

to the authorization of the territorial headquarters of the National Labor Inspectorate or, alternatively, in the case of enterprises with production units located in the areas under the jurisdiction of several territorial headquarters, of the headquarters of the National Labor Inspectorate. The measures referred to in the third sentence are final."

This section sets out that technology that has the capacity for the remote surveillance of workers (like cameras, tracking devices, or an AI-powered monitoring system) can only be used for specific and reasonable purposes:

- Organizational and production needs (e.g., ensuring efficient workflow management).
- Workplace safety (e.g., monitoring hazardous work environments).
- Protection of company assets (e.g., preventing theft or fraud).

Even with these, a company can't unilaterally implement monitoring. The implementation of a monitoring system is a matter of consensus with the trade union delegates. In the absence of a mutual agreement, the employer must obtain authorization from the National Labor Inspectorate, whose decision is final¹⁴⁷. This provision ensures that employees are protected from indiscriminate surveillance, reinforcing that monitoring tools should not be used as a means of controlling worker productivity or behavior. Instead, they should be deployed for legitimate business functions, subject to oversight from trade unions or labor authorities.

2. Exemptions for Work Tools and Attendance Recording

"The provision referred to in paragraph 1 does not apply to instruments used by the worker to render work performance and instruments for recording access and attendance."

This clause clarifies a crucial distinction: work tools that are necessary for employees to perform their job (such as computers, email accounts, and digital communication platforms) and systems used only for attendance tracking (such as badge readers) are not considered surveillance tools and do not require prior approval¹⁴⁸. This distinction is essential in assessing whether Microsoft 365 Copilot falls under Article 4

¹⁴⁷ Repubblica Italiana. (1970, 20 maggio). Legge 20 maggio 1970, n. 300, art. 4 (Statuto dei lavoratori). Norme sulla tutela della libertà e dignità dei lavoratori, della libertà sindacale e dell'attività sindacale nei luoghi di lavoro. Gazzetta Ufficiale della Repubblica Italiana. articolo 4.

¹⁴⁸ Tiraboschi, M. (2021). *Unità produttiva digitale. Perché riformare lo Statuto dei lavoratori*. ADAPT University Press. https://moodle.adaptland.it/pluginfile.php/28957/mod_resource/content/1/Unit%C3%A0%20produttiva%20digitale.%20Perch%C3%A9%20riformare%20lo%20Statuto%20dei%20lavoratori.pdf

restrictions. As a digital assistant that comes integrated into Microsoft's array of workplace tools, this tool can be considered as a work instrument. On the other hand, if enterprises would start to rely on Copilot's AI-powered insights to monitor their employees' productivity, its role might evolve from being a mere work tool to a surveillance device and should be legally compliant¹⁴⁹.

3. Use of Collected Data and Employee Notification Requirements

"The information collected under paragraphs 1 and 2 may be used for all purposes related to the employment relationship provided that the worker is given adequate information on how to use the tools and carry out the checks and in compliance with the provisions of Legislative Decree No. 196 of June 30, 2003."

Workers need to be fully aware of how such systems work and what data is being collected, as well as how that data may then be used. Compliance with privacy laws, particularly Legislative Decree No. 196/2003 (the Italian Personal Data Protection Code), which integrates GDPR principles regarding transparency, proportionality, and necessity¹⁵⁰. That means that whether or not monitoring is permissible, employees should be formally notified of what kind of data Copilot is sifting through, how it stores it, and if and how it could be used in workplace evaluations. In case information is not transparent, then the use of Copilot's insights could be illegal according to both labor and privacy laws¹⁵¹.

4.2.2. Is Copilot a Legitimate Work Tool or a Compliance Risk?

Microsoft 365 Copilot is described as an AI-driven assistant aimed at boosting productivity in the workplace by automating administrative tasks, distilling communication, and surfacing critical information across Microsoft 365 apps¹⁵².

¹⁴⁹ Barozzi, A. (2023). *Obiettivi, strumenti e metodi dell'intelligenza artificiale nella tutela della salute e della sicurezza dei lavoratori*. INAIL. <https://www.inail.it/cs/internet/docs/algoritmi-sicurezza-lavoro-barozzi.pdf>

¹⁵⁰ Panzavolta, M. (2023). *Hic sunt leones! La piramide del rischio costruita dalla proposta di Regolamento sull'intelligenza artificiale (emendata)*. *Diritto Industriale*, 1, 18–47.

¹⁵¹ Tullini, P. (2021). *Data-Driven management e strategie collettive di coinvolgimento dei lavoratori per la tutela della privacy*. In ADAPT University Press (Ed.), *Labour & Law Issues*, 7(2), 1–

20. https://moodle.adaptland.it/pluginfile.php/30133/mod_resource/content/1/Tullini%20-%20Data%20Driven%20Management.pdf

¹⁵² Microsoft. (2023, March 16). *Introducing Microsoft 365 Copilot*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot/>

But its processing of vast amounts of workplace data is problematic under Article 4 of the Workers' Statute, given the doubts over whether it constitutes a legitimate work tool or an indirect form of surveillance. According to Italian labor legislation, instruments used to perform one's job, such as e-mail systems, software applications, and communication tools, do not need to be approved beforehand by the unions¹⁵³. Strictly used as a workplace assistant, Copilot complies with this definition as it functions much like current automation features in Microsoft Office, automating the drafting of employee emails, retrieval of documents, and workflows in general¹⁵⁴. However, compliance risks emerge if Copilot's AI-generated insights are used for performance tracking. The Italian Supreme Court has ruled that even when a system is not explicitly designed for surveillance, it falls under Article 4 restrictions if it allows employers to infer employee behavior¹⁵⁵. Copilot's ability to summarize meetings, highlight key contributors, and retrieve insights from workplace interactions could lead to indirect performance monitoring, making compliance with Article 4 a necessity if these functionalities are leveraged for employee assessments¹⁵⁶. Organizations need to make sure Copilot doesn't become a monitoring tool to keep out of compliance. Employers should:

- Define Copilot's role in workplace policies.
- Restrict access to AI-generated insights to prevent indirect employee evaluations.
- Engage with trade unions to establish safeguards and transparency measures.

If Copilot is perceived as a monitoring tool, its implementation may trigger Article 4 compliance obligations, requiring union agreements or regulatory approval¹⁵⁷.

¹⁵³ Repubblica Italiana. (1970, 20 maggio). Legge 20 maggio 1970, n. 300, art. 4 (Statuto dei lavoratori). Norme sulla tutela della libertà e dignità dei lavoratori, della libertà sindacale e dell'attività sindacale nei luoghi di lavoro. Gazzetta Ufficiale della Repubblica Italiana. articolo 4.

¹⁵⁴ Microsoft. (2024, April 3). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft

Learn. <https://learn.microsoft.com/en-us/microsoft-365-copilot/privacy-and-security-for-microsoft-365-copilot>

¹⁵⁵ Cleary Gottlieb Steen & Hamilton LLP. (2023, November 13). *Defensive controls: The Italian Supreme Court outlines the employer's duties*. <https://www.clearygottlieb.com/news-and-insights/publication-listing/defensive-controls-the-italian-supreme-court-outlines-the-employers-duties>

¹⁵⁶ Zhou, L., White, R. W., Horvitz, E., & Kamar, E. (2024). *AI agents from copilots to coworkers: Historical context, challenges, limitations, implications, and practical guidelines*. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/ai-agents-from-copilots-to-coworkers-historical-context-challenges-limitations-implications-and-practical-guidelines/>

¹⁵⁷ Tullini, P. (2021). *Data-Driven management e strategie collettive di coinvolgimento dei lavoratori per la tutela della privacy*. In ADAPT University Press (Ed.), *Labour & Law Issues*, 7(2), 1–

20. https://moodle.adaptland.it/pluginfile.php/30133/mod_resource/content/1/Tullini%20-%20Data%20Driven%20Management.pdf

4.2.3. Legal Precedents on AI and Workplace Surveillance

Legal precedents have established strict limitations on the use of AI-driven workplace analytics to monitor employees, reinforcing labor protections against disproportionate surveillance. Microsoft 365 Copilot was not specifically built for supervision, but has to be viewed in the context of previous lawsuits around similar technologies. The 2021 Deliveroo decision of the Labour Court of Bologna has been an important experimental precedent, recognizing discrimination against riders through the platform's AI-based scheduling algorithm by punishing them for absences, regardless of legitimate reasons such as illness or labor strikes. This case underscores that AI-driven workforce management tools must comply with labor protections and avoid opaque decision-making that affects employees' rights. If Copilot were used to rank employees based on AI-generated summaries, it could face similar scrutiny. Similarly, Amazon received a 32 Million Euro sanction from France's CNIL in the year 2023 for breach of law concerning worker privacy as a result of excessive AI monitoring (Euronews.com). The inquiry determined that Amazon did not inform its employees about how their data was being collected and used. This is consistent with Article 4 principles that AI tools should not be used for long-term performance monitoring without supervision. These cases demonstrate that AI workplace analytics must be managed carefully. In both of these two legal precedents AI was first used as a tool to work with, then as a monitoring tool. Though Copilot alone doesn't break labor law, previous legal challenges have shown that AI-enhanced tools can potentially serve as illegal surveillance if misused. Employers implementing Copilot must ensure transparency, define access restrictions, and engage worker representatives to prevent potential legal disputes¹⁵⁸.

¹⁵⁸ Aloisi, Antonio and Gramano, Elena, Artificial Intelligence Is Watching You at Work. Digital Surveillance, Employee Monitoring, and Regulatory Issues in the EU Context (June 10, 2019). Special Issue of Comparative Labor Law & Policy Journal, "Automation, Artificial Intelligence and Labour Protection", edited by Valerio De Stefano, Vol. 41, No. 1, pp. 95-121, Available at SSRN: <https://ssrn.com/abstract=3399548>

4.3. GDPR and AI Act Compliance: The Regulatory Framework for AI-Driven Workplace Tools

4.3.1. GDPR Compliance: Transparency and Employee Data Protection

The General Data Protection Regulation (GDPR) establishes stringent requirements for processing employee data, particularly in contexts where artificial intelligence is involved. Microsoft 365 Copilot, as an AI-driven workplace tool, must comply with GDPR principles, ensuring transparency, lawfulness, and accountability in data processing. Under Article 6, employers must have a lawful basis for processing employee data, as consent is rarely considered valid in employment relationships due to the inherent power imbalance¹⁵⁹. Rather, it is generally the processing of such data on the legal grounds of legitimate interests, fulfilment of a contract, or compliance with legal obligations by which organisations must carry out balancing tests to avoid privacy invasion being disproportionate¹⁶⁰. Transparency requirements implemented under Articles 12-15 require that employees are provided with clear information about how AI systems such as Copilot process their data and explanations as to what data is used, how it is processed, and whether it is used in making workplace decisions¹⁶¹. Lack of clarity or transparency can result in regulatory scrutiny, as we have seen in situations where companies have been penalized for not making known the rationale for AI-generated decisions¹⁶². Additionally, GDPR's Article 22 prohibits decisions based solely on automated processing that significantly impact employees unless specific safeguards are in place, such as human oversight or explicit employee consent¹⁶³. Employers rolling out Copilot must make sure that AI-powered insights are not the only factor in the decision and underline the importance of

¹⁵⁹ Smit, D., & Vries, S. Y. (2019). *GDPR and personal data protection in the employment context*. In T. B. Z. Akriopoulou (Ed.), *Personal data protection and legal developments in the European Union* (pp. 146–170). IGI Global. <https://doi.org/10.4018/978-1-5225-8106-2.ch008>

¹⁶⁰ De Stefano, V., & Wouters, M. (2023). *GDPR-compliant AI-based automated decision-making in the world of work*. European Trade Union Institute (ETUI). <https://www.etui.org/publications/gdpr-compliant-ai-based-automated-decision-making-in-the-world-of-work>

¹⁶¹ European Agency for Safety and Health at Work (EU-OSHA). (2022). *AI and digital tools in workplace management and evaluation: A review of the literature and policy options*. <https://osha.europa.eu/en/publications/ai-and-digital-tools-workplace-management-and-evaluation-review-literature-and-policy-options>

¹⁶² Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

¹⁶³ Menegatti, E. (2022). *Artificial intelligence in the work process: A reflection on the proposed European Union regulations on artificial intelligence from an occupational health and safety perspective*. *Italian Labour Law e-Journal*, 15(1), 113–126. <https://doi.org/10.6092/issn.1561-8048/15323>

human intervention to any outcome of any AI-augmented evaluation of employee performance¹⁶⁴.

4.3.2. The AI Act and Workplace AI Tools: Is Copilot a High-Risk System?

The European Union's AI Act adopts a risk-based regulatory approach, classifying certain AI applications as high-risk when they pose potential threats to fundamental rights, safety, or employment conditions. Worker management, performance tracking, hiring, and firing decisions based on AI systems fall under the highly harmful category under Annex III of the regulation, necessitating specification of risk mitigation measures¹⁶⁵. Microsoft 365 Copilot is not inherently categorized as a high-risk system since it functions as a workplace assistant rather than a dedicated HR decision-making tool. However, whether it should be considered one is contingent on its use, as its classification is as an instrument of organization. While those uses of Copilot might not meet the high-risk AI system definition, limited use to monitor productivity trends, employee communications, or inform performance evaluations could invite regulatory scrutiny under the AI Act¹⁶⁶. The regulation places obligations on both AI providers and deployers, meaning that while Microsoft ensures Copilot aligns with compliance standards, businesses using it for high-risk purposes must implement additional safeguards, such as human oversight, transparency, and fairness assessments¹⁶⁷. Employers adopting Copilot into HR procedures need to be careful if it evolves into an implicit performance management device, as failure to comply with the high-risk AI requirements may lead to regulatory intervention and potential enforcement action¹⁶⁸.

¹⁶⁴ Pessi, R. (2022). *Le sfide del giurista nell'era del lavoro digitale: L'applicazione dell'IA per il controllo dei lavoratori*. *Diritti Lavori Mercati*, 1, 13–26.

¹⁶⁵ Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

¹⁶⁶ European Agency for Safety and Health at Work (EU-OSHA). (2022). *AI and digital tools in workplace management and evaluation: A review of the literature and policy options*. <https://osha.europa.eu/en/publications/ai-and-digital-tools-workplace-management-and-evaluation-review-literature-and-policy-options>

¹⁶⁷ Menegatti, E. (2022). *Artificial intelligence in the work process: A reflection on the proposed European Union regulations on artificial intelligence from an occupational health and safety perspective*. *Italian Labour Law e-Journal*, 15(1), 113–126. <https://doi.org/10.6092/issn.1561-8048/15323>

¹⁶⁸ Mökander, J., & Floridi, L. (2022). *Metrics, explainability and the European AI Act proposal*. *Minds and Machines*, 32, 629–645. <https://doi.org/10.1007/s11023-022-09612-9>

4.3.3. Compliance Challenges and Legal Uncertainties

Despite the structured frameworks provided by the GDPR and the AI Act, significant legal uncertainties remain regarding the classification, use, and oversight of AI-driven workplace analytics. Algorithmic management has no specific definition that is standardized across the EU, and creates uncertainty about the regulatory treatment of workplace AI tools such as Copilot, particularly when these tools are used to monitor employees or track performance¹⁶⁹. While the GDPR does grant some umbrella protection against data abuse, the Article 22 restrictions on automated decision-making are reserved for wholly automated processes with significant consequences, placing AI tools that provide advisory guidance out of range for direct regulatory classification¹⁷⁰. Similarly, while the AI Act explicitly designates certain HR-related AI applications as high-risk, it remains unclear whether AI-driven analytics that provide managerial recommendations, rather than making direct decisions, fall within this category¹⁷¹. Microsoft has framed Copilot as an adherent AI tool, with guardrails in place to achieve GDPR & AI governance, like preventing user data from being used for AI model training and abiding by current data protection laws¹⁷². However, responsibility for regulatory adherence ultimately depends on how businesses deploy Copilot, particularly if it is used in HR decision-making or productivity assessments¹⁷³. Regulatory developments will deliver more clarity in the next few years, notably when the AI Act becomes enforceable, and the European Data Protection Board's position on AI-driven workplace tools develops. It's not hard to imagine that businesses who are using Copilot will have to expect ever-evolving compliance requirements and be prepared to enforce governance practices to minimize legal risks should AI

¹⁶⁹ European Agency for Safety and Health at Work (EU-OSHA). (2022). *AI and digital tools in workplace management and evaluation: A review of the literature and policy options*. <https://osha.europa.eu/en/publications/ai-and-digital-tools-workplace-management-and-evaluation-review-literature-and-policy-options>

¹⁷⁰ Menegatti, E. (2022). *Artificial intelligence in the work process: A reflection on the proposed European Union regulations on artificial intelligence from an occupational health and safety perspective*.

¹⁷¹ Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

¹⁷² TUC – Trades Union Congress. (2023). *ChatGPT and beyond: Exploring the responsible use of generative AI in the workplace*. <https://www.tuc.org.uk/research-analysis/reports/chatgpt-and-beyond-exploring-responsible-use-generative-ai-workplace>

¹⁷³ European Parliament. (2023). *Artificial Intelligence 2050: Predictions, challenges, and innovations*. Scientific Foresight Unit (STOA). [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751471/EPRS_STU\(2023\)751471_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751471/EPRS_STU(2023)751471_EN.pdf)

analytics begin to shape... Automatic decision-making tools in the workplace must not end up resulting in errant employee rights violations¹⁷⁴.

4.4. AI as a Tool for Productivity or a Risk for Employee Autonomy

The use of AI-based technologies in workplace management is redefining traditional organizations by creating new efficiencies but also presenting new challenges in relation to transparency, privacy, and the autonomy of workers. Algorithmic management solutions like Microsoft 365 Copilot have the potential to make decision-making and workflows more efficient; however, they must be administered in a way that ensures we do not fall back to traditional, less effective ways of working. This section looks at the consequences of management driven by AI, the ethical issues at stake, and the need for organized labor to ensure that AI is kept as a tool for productivity, rather than a tool for control.

4.4.1. The Rise of Algorithmic Management in Workplaces

The rise of AI-based algorithmic management in the workplace is recasting the traditional role of managers and delegating tasks such as task assignment, performance measurement, and even decision-making to automated systems. AI tools such as Microsoft 365 Copilot will help to enable and drive these shifts in the way that they use large amounts of data to streamline business processes and provide insights for action, often cutting out the need for human involvement in purely managerial tasks¹⁷⁵. While AI may lead to improvements in labor productivity in terms of facilitating administrative activity, it also poses serious questions in terms of worker discretion and fairness. Algorithmic management depends on predictive analysis and historical data patterns, which may embed systemic biases that hurt certain employees more than others¹⁷⁶. AI-driven performance monitoring, particularly when used to quantify productivity metrics, can

¹⁷⁴ De Stefano, V., & Wouters, M. (2023). *GDPR-compliant AI-based automated decision-making in the world of work*. European Trade Union Institute (ETUI). <https://www.etui.org/publications/gdpr-compliant-ai-based-automated-decision-making-in-the-world-of-work>

¹⁷⁵ Rallo, A. (2023). *Socio-economic and ethical impact of artificial intelligence on organizations and the future of work*. In European Parliament (Ed.), *Panel for the Future of Science and Technology (STOA)*. [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU\(2023\)751458_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU(2023)751458_EN.pdf)

¹⁷⁶ Ponce, A. (2023). *AI employment decision-making: Integrating the equal opportunity merit principle and explainable AI*. AI & Society. <https://doi.org/10.1007/s00146-023-01786-6>

create an environment of continuous evaluation that erodes workplace trust and increases psychological stress¹⁷⁷. The black box problem of algorithmic decision-making is another reason to make it challenging to place them in the hierarchy. Workers often have little understanding of the basis for AI-based decisions, so it is hard to challenge evaluations of performance, promotion, or job assignments, which may be generated based on incomplete or incorrect information¹⁷⁸. To address these risks, European regulations emphasize transparency and oversight in AI-driven workplace management. The EU Directive (EU) 2019/1152 on Transparent and Predictable Working Conditions, transposed into Italian law through Legislative Decree No. 104/2022, mandates that employers disclose the use of automated decision-making systems in workplace governance, ensuring that employees are aware of how algorithmic tools influence their roles and conditions¹⁷⁹. But as we move to full-automation requirements, there have been certain amendments where the transparency requirements will not cover semi-automated decision-making as labor organizations have unease with not having to disclose all that is related to making decisions on the basis of data¹⁸⁰. As the AI frontier continues to reconfigure daily management practices, regulatory attempts will likely seek to ensure algorithmic tools more effectively remove waste from the process than rights and autonomy from workers¹⁸¹.

4.4.2. Ethical Considerations: Worker Privacy vs. Employer Control

The adoption of AI in workplace management presents a fundamental ethical challenge: balancing increased productivity and operational efficiency against the potential for invasive surveillance and diminished worker autonomy. AI-based platforms, including Microsoft 365 Copilot, can track employee activity, workflow processes, and even facilitate corporate decision-making and can

¹⁷⁷ Gün, D. (2023). *The impact of artificial intelligence in the workplace and its effect on the digital wellbeing of employees* (Master's thesis). Malmö University. <https://urn.kb.se/resolve?urn=urn:nbn:se:mau:diva-63760>

¹⁷⁸ Parent-Rocheleau, X., & Parker, S. K. (2023). *The rise of algorithmic work: Implications for organizational control and worker autonomy*. *Current Opinion in Psychology*, 56, 101695. <https://doi.org/10.1016/j.copsyc.2023.101695>

¹⁷⁹ Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

¹⁸⁰ Ciešlik, M. (2023). "Negotiating the algorithm": *Automation, artificial intelligence, and labor protection*. *International Journal of Comparative Labour Law and Industrial Relations*, 39(2), 111–138.

¹⁸¹ Moore, P. V., & Woodcock, J. (2023). *A policy primer and roadmap on AI worker surveillance and productivity scoring tools*. The Future of Work(ers) project. <https://www.futureofworkers.org/publications/a-policy-primer-and-roadmap-on-ai-worker-surveillance>

autonomously more effectively delegate tasks and assess performance¹⁸². At the same time, with reason, these very features can foster an atmosphere of constant surveillance and worker anxiety, where the degree of professional autonomy has declined across industries¹⁸³. Surveillance underpinned by AI is no longer limited to the classic image of monitoring but can also intercept keys, speech patterns, and attitudes, bringing an excess of control by the boss into view. Among the biggest ethical issues with workplace AI is its adoption isn't transparent. Few workers have a sense of how much AI affects what they do each day or how algorithmic assessments mold how managers approach their duties¹⁸⁴. The opacity of AI models exacerbates this problem, the reason behind algorithmic decisions is often a black box, preventing workers from disputing or even understanding automated determinations. For instance, the Italian Workers' Statute prohibits covert surveillance and mandates that employees be informed when AI-based performance assessments are in place. However, studies indicate that compliance with these transparency obligations remains inconsistent, with many organizations failing to provide employees with clear disclosures about AI monitoring practices¹⁸⁵. Another pressing ethical issue is the phenomenon of "function creep," where AI systems originally introduced for operational optimization gradually expand into more invasive surveillance tools. AI functionalities that monitor workflow performance could potentially be re-purposed to encode and enforce labor standards for employee behavior beyond work-related tasks, with privacy implications that survive legal compliance¹⁸⁶. European and Italian legislations focus on the assessment of proportionality in workplace monitoring, but ethical AI governance is not mere compliance with regulations; organizations should also act to guarantee that AI respects human dignity of workers and autonomy. Ethical deployment of AI in the workplace should prioritize meaningful human

¹⁸² Gün, D. (2023). *The impact of artificial intelligence in the workplace and its effect on the digital wellbeing of employees* (Master's thesis). Malmö University. <https://urn.kb.se/resolve?urn=urn:nbn:se:mau:diva-63760>

¹⁸³ Parent-Rocheleau, X., & Parker, S. K. (2023). *The rise of algorithmic work: Implications for organizational control and worker autonomy*. *Current Opinion in Psychology*, 56, 101695. <https://doi.org/10.1016/j.copsyc.2023.101695>

¹⁸⁴ Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). *The ethical implications of algorithms: Predictive analytics and their impact on human behaviour*. *Big Data & Society*, 3(2), 1–12. <https://doi.org/10.1177/2053951716679679>

¹⁸⁵ Ciešlik, M. (2023). "Negotiating the algorithm": Automation, artificial intelligence, and labor protection. *International Journal of Comparative Labour Law and Industrial Relations*, 39(2), 111–138.

¹⁸⁶ Moore, P. V., & Woodcock, J. (2023). *A policy primer and roadmap on AI worker surveillance and productivity scoring tools*. The Future of Work(ers) project. <https://www.futureofworkers.org/publications/a-policy-primer-and-roadmap-on-ai-worker-surveillance>

oversight, clear communication with employees, and strict limitations on the scope of AI-driven data collection¹⁸⁷.

4.4.3. The Role of Labor Unions and Worker Protections in AI Regulation

Labour unions around Europe have been appointed as central advocates for worker rights in the face of the growing prevalence of AI-fuelled management systems, calling for transparency, fairness, and employee representation in establishing AI policies. Unions understand AI can make things more productive and effective, but they demand any implementation remain free from the rights of workers or their job security. Italian labor organizations, including CGIL, CISL, and UIL, have been at the forefront of campaigns demanding stronger legal safeguards against the potential misuse of algorithmic management tools¹⁸⁸. In collective bargaining contracts, unions have worked to set the parameters around how AI-driven surveillance can be used, focusing on enabling the tools to support, not control, employees. For example, sectoral agreements in Italy have addressed the use of AI-powered tracking applications, setting strict conditions to prevent their misuse for disciplinary purposes¹⁸⁹. However, recent regulatory changes have weakened some of these protections. In 2023, Italy revised its transparency requirements, limiting disclosure obligations to fully automated decision-making systems while exempting AI tools that involve human oversight. Unions have criticized this policy shift, arguing that even semi-automated systems can significantly impact workers and should be subject to transparency rules¹⁹⁰. At the European Union level, trade unions have been instrumental in shaping regulations around AI to ensure that AI workplace technologies are not being abused. The European Trade Union Confederation (ETUC) wants tougher rules for AI-based hiring, monitoring and disciplining systems, notably in the drafting of the EU AI Act and the Platform Work

¹⁸⁷ Rallo, A. (2023). *Socio-economic and ethical impact of artificial intelligence on organizations and the future of work*. In European Parliament (Ed.), *Panel for the Future of Science and Technology (STOA)*. [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU\(2023\)751458_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU(2023)751458_EN.pdf)

¹⁸⁸ Moore, P. V., & Woodcock, J. (2023). *A policy primer and roadmap on AI worker surveillance and productivity scoring tools*. The Future of Work(ers) project. <https://www.futureofworkers.org/publications/a-policy-primer-and-roadmap-on-ai-worker-surveillance>

¹⁸⁹ Cieřlik, M. (2023). "Negotiating the algorithm": Automation, artificial intelligence, and labor protection. *International Journal of Comparative Labour Law and Industrial Relations*, 39(2), 111–138.

¹⁹⁰ Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

Directive. The new Platform Work Directive provides important safeguards against algorithmic decision-making in gig work, mandating human intervention in essential employment decisions and forbidding automated dismissals. And these protections are also likely to shape other workplace AI policies, having strengthened the principle that AI should not diminish worker rights. Europe's labour unions are militant advocates for protecting workers against the growing reliance on AI-based management systems, arguing for the imperatives of transparency, fairness and employee-contributed input in the creation of AI policies. Unions support the use of AI to improve productivity and efficiency, they say, but want government to ensure that its introduction does not undermine workers' rights and job security. Italian trade unions, such as CGIL, CISL e UIL, have taken the lead in carrying out actions to fight for stronger legislative protection of workers from potential abuse of algorithmic management tools¹⁹¹. In collective bargaining contracts, unions have worked to set the parameters around how AI-driven surveillance can be used, focusing on enabling the tools to support, not control, employees. For instance, Italian sectoral agreements for the COVID-19 emergency have looked into the use of AI-driven tracking apps imposing stringent conditions to prevent them from being used in a disciplinary manner¹⁹². But some of these protections have been weakened through new regulations. Italy updated its transparency obligations in 2023, obliging only fully automated decision-making systems to disclose themselves or themselves, exempting AI tools with a human-in-the-loop. Unions have criticized this policy change, contending that even semi-automated systems could have far-reaching impacts upon workers and therefore ought to be covered by transparency rules¹⁹³. At the European Union level, trade unions have been instrumental in shaping regulations around AI to ensure that AI workplace technologies are not being abused. The European Trade Union Confederation (ETUC) wants tougher rules for AI-based hiring, monitoring and disciplining systems, notably in the drafting of the EU AI Act and the Platform Work Directive. The new Platform

¹⁹¹ Moore, P. V., & Woodcock, J. (2023). *A policy primer and roadmap on AI worker surveillance and productivity scoring tools*. The Future of Work(ers) Project. <https://www.futureofworkers.org/publications/a-policy-primer-and-roadmap-on-ai-worker-surveillance>

¹⁹² Cieřlik, M. (2023). "Negotiating the algorithm": Automation, artificial intelligence, and labor protection. *International Journal of Comparative Labour Law and Industrial Relations*, 39(2), 111–138

¹⁹³ Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

Work Directive provides important safeguards against algorithmic decision-making in gig work, mandating human intervention in essential employment decisions and forbidding automated dismissals. And these protections are also likely to shape other workplace AI policies, having strengthened the principle that AI should not diminish worker rights. European workers' organizations have also called for compulsory algorithmic impact assessments, under which companies would be obligated to assess the impact of AI systems on employees and intervene if bias or discrimination was detected¹⁹⁴. Policy recommendations have focused on the necessity of enforceable transparency requirements that guarantee that workers can access substantive information about how AI systems will impact their work¹⁹⁵. Unions argue that AI governance frameworks must include mechanisms for workers to contest algorithmic decisions, a principle already enshrined in the GDPR but often difficult to enforce in practice¹⁹⁶. Furthermore, unions advocate for AI to be leveraged in ways that improve working conditions rather than serving as tools of surveillance and micromanagement¹⁹⁷. With AI increasingly shaping the realities of the modern workplace, the role of unions has never been more critical in grappling with the challenges and opportunities it presents, underscoring the urgent need for policies that put human oversight, worker empowerment, and ethical AI deployment front and centre.

4.5. Considerations and Final Answer: Copilot as a Productivity Tool or Surveillance Risk?

When deployed in the workplace, Microsoft 365 Copilot can have a two-edged effect: It can operate as an efficiency tool that supports workflows or, in the wrong circumstances, turn into a cudgel that is used as a surveillance and control

¹⁹⁴ Bilgihan, A., Berezina, K., & Okumus, F. (2023). *An artificial intelligence algorithmic approach to ethical decision-making in human resource management processes*. *AI and Ethics*, 3(3), 777–788. <https://doi.org/10.1007/s43681-023-00268-0>

¹⁹⁵ Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). *The ethical implications of algorithms: Predictive analytics and their impact on human behaviour*. *Big Data & Society*, 3(2), 1–12. <https://doi.org/10.1177/2053951716679679>

¹⁹⁶ Smit, D., & Vries, S. Y. (2019). *GDPR and personal data protection in the employment context*. In T. B. Z. Akrivopoulou (Ed.), *Personal data protection and legal developments in the European Union* (pp. 146–170). IGI Global. <https://doi.org/10.4018/978-1-5225-8106-2.ch008>

¹⁹⁷ Parent-Rocheleau, X., & Parker, S. K. (2023). *The rise of algorithmic work: Implications for organizational control and worker autonomy*. *Current Opinion in Psychology*, 56, 101695. <https://doi.org/10.1016/j.copsyc.2023.101695>

mechanism. The main factor will be how organizations implement, control, and oversee its use, in complying with legislation and responsible AI principles. This part focuses on key factors that contribute to Copilot's impact, weighing the potential benefits of a more optimized workplace with the need for transparency, oversight, and worker protections. Even as the sun rises over the horizon, however, there may be a few minutes or more in which its light fails to cut through the bulk of the darkness, and you are able to notice the little things.

4.5.1. Key Findings from the Analysis

Microsoft 365 Copilot offers significant potential to enhance workplace efficiency by automating routine administrative tasks, summarizing information, and assisting employees in managing their workload more effectively. Built into Microsoft 365 apps like Word, Outlook, Teams, and Excel, Copilot is capable of creating reports, composing emails, surfacing relevant stats, and delivering analysis in real-time. These features will enable employees to concentrate on more important and valuable work rather than having to waste time on mundane activities, thus increasing overall productivity and work processes¹⁹⁸. It is revealed that AI-based psychological distance in workplace assistance systems can lower cognitive load, reinforce decisions, and the ability to work together by facilitating the ease of access to company-wide data¹⁹⁹. From a business perspective, Copilot's potential to boost productivity and lower bureaucracy is consistent with larger narratives about the spread of AI applications in the workplace²⁰⁰. Yet while these advantages underpin Copilot's place as a useful workplace tool, questions over transparency, privacy, and potential abuse still loom large in the debate over its deployment. One of the primary concerns surrounding Copilot is its capacity to be repurposed as a tool for performance monitoring, raising significant transparency and privacy issues. Copilot has access to large quantities of workplace data such as emails, chats, documents, and summaries of meetings and thus provides the opportunity for this information to be abused for tracking or analyzing the employee's behavior and work

¹⁹⁸ Noy, S., Gao, R., & Zhang, P. (2024). *Early LLM-based tools for enterprise information workers likely provide meaningful boosts to productivity*. arXiv. <https://arxiv.org/abs/2403.09662>

¹⁹⁹ Reinsel, D., Gantz, J., & Rydning, J. (2023). *Generative artificial intelligence in support of analytics: Copilot 365*.

IDC. <https://www.idc.com/getdoc.jsp?containerId=US51327523>

²⁰⁰ World Economic Forum. (2024). *Managing workplace AI risks and the future of work*. Centre for the New Economy and Society. <https://www.weforum.org/publications/managing-workplace-ai-risks-and-the-future-of-work/>

patterns and making a conjecture about employees' productivity²⁰¹. If Copilot-derived analytics are subsequently used to assess employees without appropriate protections in place, it will effectively give rise to the lawful permissibility of surveillance by stealth capable of invading privacy and curtailing the autonomy of workers²⁰². The problem of AI-based workplace monitoring and surveillance has been associated with greater stress, lower levels of job satisfaction, and worries of micromanagement, particularly when workers are not informed what their data is being used for²⁰³. Transparency concerns remain, however, including the fact that employees don't understand exactly what AI Copilot does with their data, or how data-generated insights may affect their bosses' decisions²⁰⁴. This ambiguity can undermine trust, especially if workers believe they are being evaluated by algorithms without the means of effectively challenging AI-driven assessments. The regulatory landscape for AI workplace tools is the crucial need to follow existing data protection laws, from a legal and regulatory standpoint, under the umbrella of GDPR and the AI Act. This is reflected in GDPR Article 22, which protects individuals against decisions made only by automated processing²⁰⁵ and experts argue that AI-based workplace analytics should not be a substitute for human judgment when it comes to employment decisions as it could infringe Article 22 GDPR protections. Employers have to remember the outputs of Copilot are advisory, not determinative, underlining the insistence on human input in any substantial employment decision²⁰⁶. Furthermore, data governance protections should be established to guard against function creep, the tendency of AI systems introduced for productivity to be coopted as mechanisms for surveillance and behavioral control²⁰⁷. Organizations adopting Copilot should adhere to AI ethics principles and promote transparency to employees about how AI insights are

²⁰¹ Leone, R. (2023). *Privacy and beyond: Socio-ethical concerns of 'on-the-job' surveillance*. *AI & Society*. <https://doi.org/10.1007/s00146-023-01815-4>

²⁰² Kim, P. T. (2023). *The eavesdropping employer: A twenty-first century framework for employee monitoring*. *Yale Law Journal Forum*, 133, 1–24. <https://www.yalelawjournal.org/forum/the-eavesdropping-employer>

²⁰³ Ball, K., & Earl, G. (2023). *Can workers meaningfully consent to workplace wellbeing technologies?* *AI & Society*. <https://doi.org/10.1007/s00146-023-01795-5>

²⁰⁴ Kaminski, M. E. (2023). *Lessons from GDPR for AI policymaking*. *Communications of the ACM*, 66(7), 30–33. <https://doi.org/10.1145/3603829>

²⁰⁵ De Luca, S., & Federico, M. (2025). *Algorithmic discrimination under the AI Act and the GDPR*. European Parliamentary Research Service (EPRS), PE 769.509. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2025\)769509](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2025)769509)

²⁰⁶ World Economic Forum. (2024). *Managing workplace AI risks and the future of work*. Centre for the New Economy and Society. <https://www.weforum.org/publications/managing-workplace-ai-risks-and-the-future-of-work/>

²⁰⁷ Leone, R. (2023). *Privacy and beyond: Socio-ethical concerns of 'on-the-job' surveillance*. *AI & Society*. <https://doi.org/10.1007/s00146-023-01815-4>

being used, trying to ensure best practice through internal protocols that commercially bind them against the use of AI analytics for direct performance evaluation²⁰⁸. In the end, whether Copilot revolutionizes the workplace for the better is down to tight governance models to make sure AI is more a powerful tool than some covert means of surveilling and controlling workers.

4.5.2. Determining Factors: Implementation, Transparency, and Governance

The extent to which Microsoft 365 Copilot functions as a workplace productivity tool or a compliance risk depends on how employers implement and govern its use. Copilot and other systems that rely on AI to drill information out of data sets promise to automate, streamline, and data check, in ways that must be cautiously managed to prevent risks of privacy violations and AI oversight. Employers who implement Copilot with lines in the sand, establishing that it's a tool to back up rather than replace human decision-making, can see to it that it adds productivity without taking it away from workers²⁰⁹. But running Copilot to track employee interactions, measure performance metrics, or even to make employment-related decisions, in such a way that lacks transparency and human intervention, could prompt serious regulatory concerns, especially under GDPR and the AI Act²¹⁰. Employers must determine whether Copilot remains an advisory tool, providing assistance in drafting content and retrieving data, or whether it becomes a de facto management system capable of shaping employee evaluations. The key distinction lies in implementation: when AI-generated insights are leveraged to track employee behavior or performance without disclosure and safeguards, Copilot shifts from an efficiency tool to a compliance liability²¹¹. Transparency is a condition precedent to verify Copilot's deployment is legal and ethical. Employees must be properly informed as to what the Copilot is doing with workplace data, such as whether the Copilot is analyzing internal communications, generating reports on user

²⁰⁸ Kim, P. T. (2023). *The eavesdropping employer: A twenty-first century framework for employee monitoring*. *Yale Law Journal Forum*, 133, 1–24. <https://www.yalelawjournal.org/forum/the-eavesdropping-employer>

²⁰⁹ Noy, S., Gao, R., & Zhang, P. (2024). *Early LLM-based tools for enterprise information workers likely provide meaningful boosts to productivity*. arXiv. <https://arxiv.org/abs/2403.09662>

²¹⁰ Kaminski, M. E. (2023). *Lessons from GDPR for AI policymaking*. *Communications of the ACM*, 66(7), 30–33. <https://doi.org/10.1145/3603829>

²¹¹ World Economic Forum. (2024). *Managing workplace AI risks and the future of work*. Centre for the New Economy and Society. <https://www.weforum.org/publications/managing-workplace-ai-risks-and-the-future-of-work/>

activity, or impacting managerial decisions²¹². Under GDPR, Articles 12–15 mandate that workers have access to information about AI-driven data processing, while Article 22 restricts fully automated decision-making without human review²¹³. Transparency also means making sure employees understand the protections that have been put in place to prevent AI-driven surveillance or biased evaluations. The uncertainty surrounding Copilot’s operations could erode employee trust in the tool and generate legal risks if data usage is subsequently challenged under privacy legislation²¹⁴. Beyond disclosure, meaningful governance mechanisms must be established to monitor Copilot’s use, ensuring that AI-driven recommendations do not become determinative factors in professional assessments. Employers should engage in continuous evaluation, conducting audits and Data Protection Impact Assessments (DPIAs) to assess whether Copilot’s implementation aligns with compliance obligations and organizational policies²¹⁵.

4.5.3. Final Verdict: Is Microsoft 365 Copilot a Tool for Productivity or Surveillance?

Microsoft 365 Copilot has the potential to serve as a highly effective workplace productivity tool, but its ultimate impact depends entirely on how it is implemented, governed, and monitored. When utilized responsibly, Copilot can enhance productivity by automating mundane and repetitive tasks, generating summaries, assisting in report writing, and facilitating data collection. Employees may leverage AI-generated recommendations to optimize task efficiency, reduce administrative burdens, and focus on value-added work, all while maintaining autonomy²¹⁶. From a security and compliance perspective, Microsoft has designed Copilot within the Microsoft 365 framework, ensuring that user data is not employed for AI training and that access controls are aligned with existing organizational protocols. This approach reflects a commendable commitment to data privacy²¹⁷. However, responsible deployment requires strict adherence to legal and ethical guidelines, including transparency in data processing,

²¹² Leone, R. (2023). *Privacy and beyond: Socio-ethical concerns of ‘on-the-job’ surveillance*. *AI & Society*. <https://doi.org/10.1007/s00146-023-01815-4>

²¹³ De Luca, S., & Federico, M. (2025). *Algorithmic discrimination under the AI Act and the GDPR*. European Parliamentary Research Service (EPRS), PE 769.509. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2025\)769509](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2025)769509)

²¹⁴ Ball, K., & Earl, G. (2023). *Can workers meaningfully consent to workplace wellbeing technologies? AI & Society*. <https://doi.org/10.1007/s00146-023-01795-5>

²¹⁵ Kim, P. T. (2023). *The eavesdropping employer: A twenty-first century framework for employee monitoring*. *Yale Law Journal Forum*, 133, 1–24. <https://www.yalelawjournal.org/forum/the-eavesdropping-employer>

²¹⁶ Noy, S., Gao, R., & Zhang, P. (2024). *Early LLM-based tools for enterprise information workers likely provide meaningful boosts to productivity*. arXiv. <https://arxiv.org/abs/2403.09662>

²¹⁷ Reinsel, D., Gantz, J., & Rydning, J. (2023). *Generative artificial intelligence in support of analytics: Copilot 365*. IDC. <https://www.idc.com/getdoc.jsp?containerId=US51327523>

clear communication with employees, and safeguards to prevent AI-generated insights from being used as covert performance monitoring tools²¹⁸. Conversely, if employers misuse Copilot as a means of silently tracking employee activity, monitoring communications, or making HR-related judgments without transparency and oversight, it risks becoming a workplace surveillance tool. Algorithm-driven workplace instruments may create a misleading impression of objectivity, wherein quantitative assessments of employee performance are regarded as unequivocal evidence, devoid of necessary contextual understanding or human judgment²¹⁹. In the absence of appropriate checks and balances, such data collection capabilities could be employed by organizations to monitor productivity, analyze communications, or even to flag employees based on algorithmic assumptions, thereby raising concerns regarding workplace surveillance and inequitable evaluations²²⁰. The distinction between utilizing Copilot as a tool for efficiency and as an instrument for surveillance ultimately hinges on the governance exercised by employers and the regulatory frameworks established to safeguard worker rights. Organizations should implement explicit guidelines regarding the use of Copilot, ensuring that the tool serves as a source of guidance rather than directive mandates for employees²²¹. Transparent communication with employees is essential, providing them with accessible explanations of how Copilot interacts with their data²²². Ultimately, the ethical and legal implications of Copilot depend on the choices made by employers.

²¹⁸ Kaminski, M. E. (2023). *Lessons from GDPR for AI policymaking*. *Communications of the ACM*, 66(7), 30–33. <https://doi.org/10.1145/3603829>

²¹⁹ De Luca, S., & Federico, M. (2025). *Algorithmic discrimination under the AI Act and the GDPR*. European Parliamentary Research Service (EPRS), PE 769.509. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2025\)769509](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2025)769509)

²²⁰ Kim, P. T. (2023). *The eavesdropping employer: A twenty-first century framework for employee monitoring*. *Yale Law Journal Forum*, 133, 1–24. <https://www.yalelawjournal.org/forum/the-eavesdropping-employer>

²²¹ World Economic Forum. (2024). *Managing workplace AI risks and the future of work*. Centre for the New Economy and Society. <https://www.weforum.org/publications/managing-workplace-ai-risks-and-the-future-of-work/>

²²² Kaminski, M. E. (2023). *Lessons from GDPR for AI policymaking*. *Communications of the ACM*, 66(7), 30–33. <https://doi.org/10.1145/3603829>

Conclusion:

5.1 Objective, Research Questions, and Method Recap

This thesis set out to examine the normative challenges posed by the deployment of Microsoft 365 Copilot in enterprise settings, with particular attention to its alignment with the General Data Protection Regulation (GDPR), the forthcoming Artificial Intelligence Act (AI Act), and Italian labor law. The investigation focused on two key questions:

- Does Microsoft, as the provider of Copilot, indirectly transfer user data to OpenAI?
- Is Microsoft 365 Copilot a Tool for Productivity or Surveillance?

To address these questions, the research employed an integrated methodology. The initial phase involved a thorough analysis of technical documents released by Microsoft, focusing on data flows, system architecture, and the role of third parties in the operational framework. Subsequently, a doctrinal analysis was conducted to elucidate the relevant legal standards, including the pertinent legal provisions, enforcement policies, and prevailing interpretations by regulatory authorities, particularly the European Data Protection Board. Finally, a critical review of existing legal literature and case law provided further insights into how architectural design choices interact with fundamental legal principles such as accountability, transparency, and purpose limitation. This mixed-methods approach facilitated a robust analytical framework necessary for understanding the extent to which emerging AI-as-a-Service (AIaaS) platforms align, or fail to align, with the evolving European regulatory landscape, particularly concerning the rights and obligations of both individuals and organizations.

5.2 Critical and Structured Summary of Findings

5.2.1 Chapter 1 – AI Service Architecture and Distributed Responsibility

The first key finding of this thesis is that the architectural configuration of enterprise AI services such as Microsoft 365 Copilot fragments legal responsibility across multiple actors, undermining the effectiveness of core data protection principles. These services rely on a multilayered ecosystem, in which developers (Such as OpenAI), service providers (As Microsoft), and corporate clients each contribute to

different parts of the data processing process, from data acquisition and preprocessing through to model inference and potential logging. Every stage poses unique risks, from model inversion and data leakage to unauthorized cross-tenant access. While Microsoft has implemented certain privacy-enhancing technologies (PETs), such as retrieval-augmented generation (RAG) and federated learning, these technical measures may not be sufficient for Microsoft to constitute itself as meeting its GDPR responsibilities. Notably, there is no unified framework ensuring accountability across the entire lifecycle of data processing. As a consequence, the accountability principle set out in Article 5(2) GDPR, in accordance with which data controllers have to demonstrate not only compliance with but also accountability under the law, is watered down in a substantial manner. The current architecture, by design, dilutes legal responsibility, making it difficult to establish which actor is accountable for which processing operation and under what legal basis. This opacity has real consequences since it interferes with the capacity of enterprise clients to confirm compliance and also with the ability of data subjects to effectively apply their rights.

5.2.2 Chapter 2 – Indirect Data Transfer to OpenAI and Lack of Transparency

The second core finding concerns the legal ambiguity surrounding the involvement of OpenAI in the operation of Microsoft 365 Copilot, particularly during the inference stage. Microsoft claims contractually that no personal data is shared with OpenAI, but this claim rests on a narrow reading of “transfer,” and there is no transparent, technical proof that this is the case. According to article 4(2) of the GDPR, any accessing of personal data, regardless if of a nature that is of transitory character during processing (“such as temporary copying in the RAM of a computer”), does constitute processing. Given that Copilot utilizes the GPT-4 model to generate responses to prompts containing potentially sensitive enterprise information, that inference process is squarely an object within the “material scope” of the Regulation. Microsoft’s documentation, though, does not specifically address whether OpenAI, as a model developer and infrastructure provider, is technically prohibited from touching these inputs, even for a moment. The absence of such a definition is a serious risk of infringement (also with reference to the data subjects and enterprise clients who seek to comprehend the real entity of data processing processes). This opacity is

compounded by the absence of a clearly defined legal role for OpenAI in Microsoft's Data Processing Addendum (DPA). Whether OpenAI acts as a processor, sub-processor, or third party service provider is still up in the air, and this completely contradicts the transparency requirements pursuant to Articles 13 and 14 GDPR. These provisions require that data subjects be informed about all parties involved in processing their data, the purposes of the processing, and the applicable safeguards. When these roles are not made explicit, transparency becomes illusory, and the capacity to attribute responsibility is lost. Precedents from European data protection authorities, including the Italian DPA's position on biometric surveillance at Milan's Central Station and the Swedish DPA's assessment concerning school-based facial recognition, demonstrates that even when data processing is transient or indirect, full regulatory compliance is needed. The temporality does not make these contributions any less relevant to the law. And finally, Microsoft's use of zero-retention policies and non-persistence assurances is not enough protection against that. The inference process is a black box: it is not auditable, traceable, nor verifiable by an enterprise customer. Without the support of technical means or contractual entitlements that make it possible for the customers to verify the processing of data during this phase, undisclosed processing or indirect access cannot be excluded from the picture. As a result, this thesis finds that Copilot's design creates a substantive risk of unauthorized or unsupported data processing with OpenAI. The current governance does not address this risk, and tighter regulation and greater clarity of roles may be required.

5.2.3 Chapter 3 – Algorithmic Surveillance and Labor Law

The final substantive finding addresses the implications of Copilot's deployment in the workplace, where its technical functionalities may produce surveillance-like effects, even if not explicitly intended as monitoring tools. While Microsoft presents Copilot as a neutral productivity enhancer, automating tasks such as drafting emails, summarizing meetings, or retrieving files, the system inherently processes and generates outputs based on workers' interactions with enterprise software. These outputs, recommended links or references, summaries, or edits, are not inherently anonymous, but rather they mirror and reformulate the digital

behavior of identifiable users. Even without direct performance dashboards, this capability allows for indirect ‘profiling’ of behaviour and attributing contributions in ways that may be leveraged (whether intentionally or not) for evaluative or disciplinary purposes. Used as such, Copilot’s tools fall afoul of the circumference of Article 4 of the Italian Workers’ Statute, which prohibits remote monitoring tools unless introduced through a pre-established trade union agreement or labor inspectorate authorization. The risk is not merely theoretical. According to Italian consolidated case law, even instruments that do not seem at first sight likely to infringe, through such a range of safe operational uses, the competing rights of workers can be unlawful by their clandestine use or by the absence of safety measures. What the Italian Supreme Court has clarified is the principle that its purpose for being cannot decide whether the technology at stake violates the law, but how the technology operates and what its effect is. If Copilot enables managers to reconstruct workflows, identify individual bottlenecks, or infer performance levels based on AI-generated summaries, it effectively becomes a monitoring tool. In this case, its use had to comply with collective bargaining processes and the obligation to transparency towards employees. To not do so would not only violate national labor law but also fundamental rights to dignity and workplace autonomy. Further, Copilot is problematic from a regulatory perspective, both under the GDPR and the AI Act. Under article 22 of GDPR, automated processing that produces effects on individuals, especially in employment contexts, requires explicit safeguards, including the right to obtain human intervention and contest the decision. While Copilot may not make direct decisions, the insights it generates could feed into automated evaluation processes, triggering Article 22 protection. The AI Act also deems employee management and monitoring systems as high-risk, with more stringent conformity assessment and transparency requirements applying to them. To the extent that Copilot is used in manners that impact decision-making about tasks, performance, or behavior, the model must be considered in assessment and treatment planning within the risk framework presented. This thesis thus concludes that Copilot, depending on its configuration and use, can reasonably be considered an algorithmic surveillance tool, requiring enhanced safeguards at the intersection of data protection and labor law.

5.3. Limitations

Despite the depth of its legal and technical analysis, this thesis is subject to certain limitations, primarily stemming from its reliance on publicly accessible sources. Microsoft 365 Copilot's adherence to data protection and employment laws was evaluated based on documentation, white papers, and contractual representations from Microsoft and OpenAI. However, such sources do not offer complete visibility into real-time data flows, backend infrastructure, or short-lived processing activities, in particular those associated with the inference stage. Without access to proprietary data such as system logs, model telemetry, and internal audit reports, it is simply not possible to verify beyond any doubt that the provider's processing assertions are an accurate reflection of how the system was behaving. This lack of evidence limits our ability to reach firm conclusions about technical compliance and management of information. Furthermore, the regulatory landscape remains in flux: the final implementation guidance for the AI Act, as well as evolving positions from the European Data Protection Board (EDPB) may impact how inference is legally defined, and how workplace-use AI engines are categorized by degrees of risk. As such, some of the normative implications set out herein, most notably the implications of inference as processing and the threshold for high-risk designation, should be understood as having been derived from the current legal doctrine and may need to be modified as the jurisprudential and regulatory landscape develops.

5.4. Policy and Research Recommendations

Building on the findings of this research, several policy and research recommendations emerge as necessary steps to close the regulatory gaps identified in the governance of enterprise AI services. First, EU regulators need to expressly acknowledge inference as a regulated processing activity both under the GDPR and the AI Act (even where no data is kept). The occasional use of personal or corporate data in conjunction with personal or corporate information to produce an output is processing in substance and is to be regarded as such in order to promote legal certainty. Secondly, the application of general purpose foundation models in B2B environments requires significantly more transparency. Service providers like Microsoft should be legally compelled to exactly state what the particular aspect of data processing, say inference, has played in third party-model developers such as OpenAI. Without this, clients and data subjects are left in

the dark, unable to verify compliance or enforce their rights. In parallel, it is essential to introduce independent audit mechanisms to embedded AI systems: traceability, human oversight and non-retention guarantees should not be based solely on provider self-certification. In the area of labour, national laws will need to be modernised to keep up with the line blurring personal productivity tools and surveillance devices. As demonstrated in this thesis, tools like Copilot may have unintended monitoring effects, activating legal protections that are currently outdated or too narrowly framed. Collective oversight mechanisms, such as prior agreements with unions or supervisory authorities, should be adapted to the digital context to prevent covert data exploitation. Lastly, this thesis calls for more interdisciplinary studies that center around compliance regimes extending across data protection, AI governance, and labour law. And the space between these two fields is where regulatory tension is most pronounced, in particular in difficult-to-map-out AI-as-a-Service models of AI. The need for comprehensive, legally defensible and practically applicable assessment tools could present an obstacle in a world where AI becomes a pervasive component of the enterprise lifestyle. Ultimately, the issues raised in this thesis speak to a deeper structural challenge: contemporary regulatory systems struggle to keep pace with the opacity and distributed responsibility of modern AI architectures. When providers rely on external model developers without full disclosure, when inference is treated as a legal grey zone, and when workers are exposed to automated analysis without safeguards, the foundational principles of transparency, accountability, and fairness are at risk. Addressing these concerns is not merely a matter of compliance but of democratic integrity. If left unresolved, enterprise AI systems may erode the very legal and ethical structures meant to protect individuals in the digital age.

Bibliography:

AI Now Institute. (2023, April 11). *Algorithmic management: Restraining workplace surveillance*. <https://ainowinstitute.org/publication/algorithmic-management>

Allam, Karthik. (2023). Adoption of Artificial Intelligence in Cloud Computing. *International Journal of Computer Trends and Technology*, 71, 91-95. 10.14445/22312803/IJCTT-V71I16P116.

Aloisi, A., & Gramano, E. (2023). *Between risk mitigation and labour rights enforcement: Assessing the transatlantic race to govern AI-driven decision-making through a comparative lens*. *Computer Law & Security Review*, 51, 105783. <https://doi.org/10.1016/j.clsr.2023.105783>

Anderson, M., Amit, G., & Goldstein, A. (2024). Is my data in your retrieval database? membership inference attacks against retrieval augmented generation. *arXiv preprint arXiv:2405.20446*.

Andrew, James. (2025). *AI Model Lifecycle Management: Strategies for Scalable Deployment and Maintenance*.

Baheri, A. (2023). Towards Theoretical Understanding of Data-Driven Policy Refinement. *arXiv preprint arXiv:2305.06796*.

Ball, K., & Francis, G. (2023). *Private eyes, they see your every move: Workplace surveillance and worker well-being*. University of Essex / European Commission. Retrieved from <https://www.essex.ac.uk/-/media/documents/research/centre-for-work-organization-and-society/private-eyes-they-see-your-every-move.pdf>

Barozzi, A. (2023). *Obiettivi, strumenti e metodi dell'intelligenza artificiale nella tutela della salute e della sicurezza dei lavoratori*. INAIL. <https://www.inail.it/cs/internet/docs/algoritmi-sicurezza-lavoro-barozzi.pdf>

Bilgihan, A., Berezina, K., & Okumus, F. (2023). *An artificial intelligence algorithmic approach to ethical decision-making in human resource management processes*. *AI and Ethics*, 3(3), 777–788. <https://doi.org/10.1007/s43681-023-00268-0>

Carter, C. (2024). AI surveillance: Reclaiming privacy through informational control. *European Labour Law Journal*, 0(0). <https://doi.org/10.1177/20319525241306327>

Cieślak, M. (2023). "Negotiating the algorithm": Automation, artificial intelligence, and labor protection. *International Journal of Comparative Labour Law and Industrial Relations*, 39(2), 111–138.

Cleary Gottlieb Steen & Hamilton LLP. (2023, November 13). *Defensive controls: The Italian Supreme Court outlines the employer's duties*. <https://www.clearygottlieb.com/news-and-insights/publication-listing/defensive-controls-the-italian-supreme-court-outlines-the-employers-duties>

Clifford Chance. (2023, April 1). *The Italian Data Protection Authority halts ChatGPT's data processing operations*.

Commission Nationale de l'Informatique et des Libertés (CNIL). (2025, February 7). *AI and GDPR: The CNIL publishes new recommendations to support responsible innovation*

Data Protection Report. (2025, January). *The EDPB Opinion on training AI models using personal data and recent Garante fine – lawful deployment of LLMs*.

Davis, J. (2023, October 17). *AI surveillance in the workplace linked to employee resistance*. SHRM. <https://www.shrm.org/topics-tools/news/employee-relations/ai-surveillance-in-the-workplace-linked-to-employee-resistance-->

De Luca, S., & Federico, M. (2025). *Algorithmic discrimination under the AI Act and the GDPR*. European Parliamentary Research Service (EPRS), PE 769.509. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2025\)769509](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2025)769509)

De Stefano, V., & Wouters, M. (2023). *GDPR-compliant AI-based automated decision-making in the world of work*. European Trade Union Institute (ETUI). <https://www.etui.org/publications/gdpr-compliant-ai-based-automated-decision-making-in-the-world-of-work>

- Deshpande, A., & Sharp, H. (2022). Responsible AI systems: Who are the stakeholders? Paper presented at the 227-236. <https://doi.org/10.1145/3514094.3534187>
- European Agency for Safety and Health at Work (EU-OSHA). (2022). *AI and digital tools in workplace management and evaluation: A review of the literature and policy options*. <https://osha.europa.eu/en/publications/ai-and-digital-tools-workplace-management-and-evaluation-review-literature-and-policy-options>
- European Commission. (2021). Ethics guidelines for trustworthy AI. High-Level Expert Group on Artificial Intelligence (HLEG). Retrieved from <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- European Commission. (2021). Proposal for a Regulation laying down harmonized rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. Brussels, Belgium. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- European Data Protection Board. (2019, August 22). *Facial recognition in school renders Sweden's first GDPR fine*
- European Data Protection Board. (2024, December 17). *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*.
- European Data Protection Supervisor, European Union Agency for Fundamental Rights, & Council of Europe. (2018). *Handbook on European Data Protection Law*.
- European Data Protection Supervisor. (2018). *Handbook on European data protection law*. Publications Office of the European Union
- European Parliament & Council of the European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation - GDPR). Official Journal of the European Union, L 119, 1–88. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- European Parliament. (2023). *Artificial Intelligence 2050: Predictions, challenges, and innovations*. Scientific Foresight Unit (STOA). [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751471/EPRS_STU\(2023\)751471_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751471/EPRS_STU(2023)751471_EN.pdf)
- European Union. (2024). Artificial Intelligence Act (Regulation EU 2024/1689). Garante per la Protezione dei Dati Personali. (2017). *Provvedimento n. 551 del 21 dicembre 2017*. ([Source: Italian DPA]).
- Gün, D. (2023). *The impact of artificial intelligence in the workplace and its effect on the digital wellbeing of employees* (Master's thesis). Malmö University. <https://urn.kb.se/resolve?urn=urn:nbn:se:mau:diva-63760>
- Hern, A. (2020, 26 novembre). *Microsoft productivity score feature criticised as workplace surveillance*. *The Guardian*. <https://www.theguardian.com/technology/2020/nov/26/microsoft-productivity-score-feature-criticised-workplace-surveillance>
- Hoffman Robert R. , Mueller Shane T. , Klein Gary , Jalaiean Mohammadreza , Tate Connor, *Explainable AI: roles and stakeholders, desirements and challenges*, *Frontiers in Computer Science* Vol.5, 2023-
<https://www.frontiersin.org/journals/computerscience/articles/10.3389/fcomp.2023.1117848>
DOI=10.3389/fcomp.2023.1117848
- Holistic AI. (2023, July 1). Regulating foundation models and generative AI: The EU AI Act. *Holistic AI*.
Information Commissioner's Office. (2024). *Guide to the General Data Protection Regulation (GDPR)*.
Jisc Involve. (n.d.). *National Centre for AI*.
- Kaminski, M. E. (2023). *Lessons from GDPR for AI policymaking*. *Communications of the ACM*, 66(7), 30–33. <https://doi.org/10.1145/3603829>
- Kim, P. T. (2023). *The eavesdropping employer: A twenty-first century framework for employee monitoring*. *Yale Law Journal Forum*, 133, 1–24. <https://www.yalelawjournal.org/forum/the-eavesdropping-employer>

- Koga, T., Wu, R., & Chaudhuri, K. (2024). *Privacy-preserving retrieval augmented generation with differential privacy*(Preprint). University of California, San Diego.
- Leone, R. (2023). *Privacy and beyond: Socio-ethical concerns of 'on-the-job' surveillance*. *AI & Society*. <https://doi.org/10.1007/s00146-023-01815-4>
- Lins, S., Pandl, K.D., Teigeler, H. et al. Artificial Intelligence as a Service. *Bus Inf Syst Eng* 63, 441–456 (2021). <https://doi.org/10.1007/s12599-021-00708-w>
- Mantelero, A. (2022). *AI surveillance: Reclaiming privacy through informational control*. In S. Nemitz & A. L. R. Zanatta (Eds.), *AI and Democracy* (pp. 123–135). Springer. https://doi.org/10.1007/978-3-030-96185-0_8
- McIntosh, T. R., Susnjak, T., Liu, T., Watters, P., Xu, D., Liu, D., & Halgamuge, M. N. (2025). From Google Gemini to OpenAI Q* (Q-Star): A Survey on Reshaping the Generative Artificial Intelligence (AI) Research Landscape. *Technologies*, 13(2), 51.
- Menegatti, E. (2022). *Artificial intelligence in the work process: A reflection on the proposed European Union regulations on artificial intelligence from an occupational health and safety perspective*. *Italian Labour Law e-Journal*, 15(1), 113–126. <https://doi.org/10.6092/issn.1561-8048/15323>
- Meurisch, C., & Mühlhäuser, M. (2022). Data Protection in AI Services, *ACM Computing Surveys*., 54(2). <https://doi.org/10.1145/3440754>
- Microsoft. (2019, July 22). *Microsoft invests in and partners with OpenAI to support us building beneficial AGI*.
- Microsoft. (2020, December 1). *Our commitment to privacy in Microsoft Productivity Score*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2020/12/01/our-commitment-to-privacy-in-microsoft-productivity-score/>
- Microsoft. (2023, March 16). *Introducing Microsoft 365 Copilot*. Retrieved June 23, 2025, from <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot/>
- Microsoft. (2024, April 3). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft Learn. <https://learn.microsoft.com/en-us/microsoft-365-copilot/privacy-and-security-for-microsoft-365-copilot>
- Microsoft. (2024, March 28). Protecting the data of our commercial and public sector customers in the AI era. *Microsoft On the Issues*
- Microsoft. (2024, March 6). *Azure OpenAI Service powers the Microsoft Copilot ecosystem*. Microsoft Azure Blog.
- Microsoft. (2024). *Data, privacy, and security for Azure OpenAI Service*.
- Microsoft. (2024). *GDPR & Generative AI: A Guide for the Public Sector*.
- Microsoft. (2024). *Responsible AI Transparency Report*
- Microsoft. (2025, April 15). *Network and access configuration for Azure OpenAI On Your Data*.
- Microsoft. (2025, February 13). *Microsoft 365 Copilot overview*. Microsoft Learn.
- Microsoft. (2025, January 15). *Data, privacy, and security for Microsoft 365 Copilot*. Microsoft Learn.
- Microsoft. (2025, January 28). *How does Microsoft 365 Copilot work?* Microsoft Learn.
- Microsoft. (2025, January). *The EU AI Act: A Microsoft overview*.
- Microsoft. (2025, March 6). *Semantic indexing for Microsoft 365 Copilot*. Microsoft Learn.
- Microsoft. (2025). *Enterprise data protection in Microsoft 365 Copilot and Microsoft 365 Copilot Chat*.

Microsoft. (n.d.). *Adoption Score in Microsoft 365*. Retrieved June 23, 2025, from <https://learn.microsoft.com/en-gb/microsoft-365/admin/adoption/adoption-score?view=o365-worldwide>

Microsoft. (n.d.). *Retrieval Augmented Generation (RAG) in Azure AI Search*.

Ministry of Justice and Security. (2024). *DPIA Microsoft 365 Copilot: Data Protection Impact Assessment on the Processing of Personal Data with Microsoft 365 Copilot*.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). *The ethical implications of algorithms: Predictive analytics and their impact on human behaviour*. *Big Data & Society*, 3(2), 1–12. <https://doi.org/10.1177/2053951716679679>

Mökander, J., & Floridi, L. (2022). *Metrics, explainability and the European AI Act proposal*. *Minds and Machines*, 32, 629–645. <https://doi.org/10.1007/s11023-022-09612-9>

Moore, P. V., & Woodcock, J. (2023). *A policy primer and roadmap on AI worker surveillance and productivity scoring tools*. The Future of Work(ers) project. <https://www.futureofworkers.org/publications/a-policy-primer-and-roadmap-on-ai-worker-surveillance>

Nance, R., Evans, M., & Gelmetti, F. (2025, January 2). The EDPB Opinion on training AI models using personal data and recent Garante fine – lawful deployment of LLMs. *Data Protection Report*

Noy, S., Gao, R., & Zhang, P. (2024). *Early LLM-based tools for enterprise information workers likely provide meaningful boosts to productivity*. arXiv. <https://arxiv.org/abs/2403.09662>

OpenAI. (n.d.). *Data processing addendum*. OpenAI. Retrieved March 8, 2025

Panzavolta, M. (2023). *Hic sunt leones! La piramide del rischio costruita dalla proposta di Regolamento sull'intelligenza artificiale (emendata)*. *Diritto Industriale*, 1, 18–47.

Parent-Rocheleau, X., & Parker, S. K. (2023). *The rise of algorithmic work: Implications for organizational control and worker autonomy*. *Current Opinion in Psychology*, 56, 101695. <https://doi.org/10.1016/j.copsyc.2023.101695>

Pataro, J. (2020, December 1). *Our commitment to privacy in Microsoft Productivity Score*. Microsoft 365 Blog. <https://www.microsoft.com/en-us/microsoft-365/blog/2020/12/01/our-commitment-to-privacy-in-microsoft-productivity-score/>

Pessi, R. (2022). *Le sfide del giurista nell'era del lavoro digitale: L'applicazione dell'IA per il controllo dei lavoratori*. *Diritti Lavori Mercati*, 1, 13–26.

Ponce, A. (2023). *AI employment decision-making: Integrating the equal opportunity merit principle and explainable AI*. *AI & Society*. <https://doi.org/10.1007/s00146-023-01786-6>

Rallo, A. (2023). *Socio-economic and ethical impact of artificial intelligence on organizations and the future of work*. In European Parliament (Ed.), *Panel for the Future of Science and Technology (STOA)*. [https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU\(2023\)751458_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2023/751458/EPRS_STU(2023)751458_EN.pdf)

Rani, P., Kavita, Verma, S., Kaur, N., Wozniak, M., Shafi, J., & Ijaz, M. F. (2022). *Robust and Secure Data Transmission Using Artificial Intelligence Techniques in Ad-Hoc Networks*. *Sensors*, 22(1), 251. <https://doi.org/10.3390/s22010251>

Reinsel, D., Gantz, J., & Rydning, J. (2023). *Generative artificial intelligence in support of analytics: Copilot 365*. IDC. <https://www.idc.com/getdoc.jsp?containerId=US51327523>

Repubblica Italiana. (1970, 20 maggio). Legge 20 maggio 1970, n. 300, art. 4 (Statuto dei lavoratori). Norme sulla tutela della libertà e dignità dei lavoratori, della libertà sindacale e dell'attività sindacale nei luoghi di lavoro. *Gazzetta Ufficiale della Repubblica Italiana*. articolo 4.

Sandler, R. (2020, 26 novembre). *C'è un nuovo strumento Microsoft che controlla a distanza i lavoratori. Ed è bufera sulla privacy*. Forbes Italia. <https://forbes.it/2020/11/26/microsoft-productivity-score-il-nuovo-strumento-per-il-controllo-a-distanza-dei-lavoratori/>

Sangwan RS, Badr Y, Srinivasan SM. Cybersecurity for AI Systems: A Survey. *Journal of Cybersecurity and Privacy*. 2023; 3(2):166-190. <https://doi.org/10.3390/jcp3020010>

Smit, D., & Vries, S. Y. (2019). *GDPR and personal data protection in the employment context*. In T. B. Z. Akrivopoulou (Ed.), *Personal data protection and legal developments in the European Union* (pp. 146–170). IGI Global. <https://doi.org/10.4018/978-1-5225-8106-2.ch008>

Spjuth, O., Frid, J., & Hellander, A. (2021). The machine learning life cycle and the cloud: implications for drug discovery. *Expert Opinion on Drug Discovery*, 16(9), 1071–1079. <https://doi.org/10.1080/17460441.2021.1932812>

Sweetman, S. (2024, September 24). *Enterprise trust in Azure OpenAI Service strengthened with Data Zones* azure.microsoft.com. Microsoft Azure Blog.

Tiraboschi, M. (2021). *Unità produttiva digitale. Perché riformare lo Statuto dei lavoratori*. ADAPT University Press. https://moodle.adaptland.it/pluginfile.php/28957/mod_resource/content/1/Unit%C3%A0%20produttiva%20digitale.%20Perch%C3%A9%20riformare%20lo%20Statuto%20dei%20lavoratori.pdf

TUC – Trades Union Congress. (2023). *ChatGPT and beyond: Exploring the responsible use of generative AI in the workplace*. <https://www.tuc.org.uk/research-analysis/reports/chatgpt-and-beyond-exploring-responsible-use-generative-ai-workplace>

Tullini, P. (2021). *Data-Driven management e strategie collettive di coinvolgimento dei lavoratori per la tutela della privacy*. In ADAPT University Press (Ed.), *Labour & Law Issues*, 7(2), 1–20. https://moodle.adaptland.it/pluginfile.php/30133/mod_resource/content/1/Tullini%20-%20Data%20Driven%20Management.pdf

U.S. Department of Homeland Security, *Roles and Responsibilities Framework for Artificial Intelligence in Critical Infrastructure*, In Consultation with The Artificial Intelligence Safety and Security Board, November 14, 2024

Van Binsbergen, L. T., Steketee, M. C., Kebede, M. G., Janssen, H. L., & van Engers, T. M. (2025). Lawful and Accountable Personal Data Processing with GDPR-based Access and Usage Control in Distributed Systems. *arXiv preprint arXiv:2503.07172*.

van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, 11(1). <https://doi.org/10.1177/20539517241232630>

Witanto, E. N., Oktian, Y. E., & Lee, S.-G. (2022). Toward Data Integrity Architecture for Cloud-Based AI Systems. *Symmetry*, 14(2), 273. <https://doi.org/10.3390/sym14020273>

World Economic Forum. (2024). *Managing workplace AI risks and the future of work*. Centre for the New Economy and Society. <https://www.weforum.org/publications/managing-workplace-ai-risks-and-the-future-of-work/>

Zeng, S., Zhang, J., He, P., Xing, Y., Liu, Y., Xu, H., Ren, J., Wang, S., Yin, D., Chang, Y., & Tang, J. (2024). *The good and the bad: Exploring privacy issues in retrieval-augmented generation (RAG)*. Proceedings of the Association for Computational Linguistics (ACL 2024), 4505–4524.

Zhou, L., White, R. W., Horvitz, E., & Kamar, E. (2024). *AI agents from copilots to coworkers: Historical context, challenges, limitations, implications, and practical guidelines*. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/ai-agents-from-copilots-to-coworkers-historical-context-challenges-limitations-implications-and-practical-guidelines/>

Zhu, Z., Li, Y., & Zhang, H. (2024). Optimizing large language models for OpenAPI code completion. *arXiv*.

Zitek, E., & Brekken, K. (n.d.). *AI surveillance in the workplace linked to employee resistance, turnover*. SHRM.