

Optimization of Generative AI Efficiency: A Case Study Based on Professional Workshop Questionnaires

Milica Mravik¹, Marko Šarac²

Faculty of Informatics and Computing, Singidunum University¹, Belgrade,

Faculty of Computing and Informatics, Sinergija University², Bijeljina

In this research, we explored the strategic optimization of Microsoft 365 Copilot by examining the transition from theoretical readiness to practical user mastery. This research was based on a structured "Art of the Possible" (AOTP) workshop approach, which was the key method for discovering high-impact business scenarios. By analysing quantitative and qualitative information from an AI Workshop Questionnaire, we were able to determine the impact of specific organizational interventions on user adoption and technical efficiency. Our results showed that successful optimization was more than just a function of technical setup and was instead highly reliant on expertise in prompt engineering and addressing security issues, like data oversharing. In addition, we emphasized the need for a "Crawl-Walk-Run" approach and the critical role of internal "Champions" in maintaining digital transformation. We concluded that a feedback loop between end-users and stakeholders is necessary for aligning generative AI capabilities with organizational goals.

Keywords – *Microsoft 365 Copilot, AI Optimization, Prompt Engineering, Digital Transformation, Art of the Possible (AOTP), User Adoption Framework*

I. INTRODUCTION

The current state of the organizational environment is in the midst of a paradigm shift that involves the integration of large language models (LLMs) into the very heart of workplace productivity [1, 2]. Microsoft 365 Copilot is a more advanced form of this technological shift, providing a seamless integration between generative artificial intelligence and traditional enterprise software [3]. Nevertheless, the literature on digital transformation and technology adoption implies that the mere possession of highly advanced technological assets does not automatically lead to enhanced performance. Rather, we noticed that the difference between technological potential and actual productivity is often bridged by aligning the technological asset with the socio-technical context of the organization [4, 5].

Before embarking on this research, we were cognizant of the fact that the major challenge to the successful implementation of AI technology is the absence of a structured pedagogical framework for end-users. While the adoption of traditional software technology is based on deterministic functional training, the adoption of generative AI technology

is predicated on a heuristic understanding of prompt engineering and feedback. Hence, we aimed to explore the role of a methodology-driven intervention, namely the "Art of the Possible" (AOTP) workshop, in facilitating the shift from passive observation to active, high-value use [6].

One of the major academic concerns in the use of AI in the corporate setting is data governance and security. We highlighted that institutional concern about data oversharing and privacy tends to lead to a conservative adoption rate, which is fundamentally limiting to the tool's potential. To mitigate this, we incorporated a full analysis of the organization's data estate into our design, ensuring that technical setup and security posture was the basis for user engagement. This enabled us to investigate the relationship between perceived security and user willingness to test novel AI-based workflows.

In addition, we investigated the role of organizational change management through the prism of a "Champion" network. Our study assumed that innovation diffusion is most successful when facilitated by peer-to-peer knowledge sharing, rather than top-down directives. By encouraging internal early adopters to identify and record "Quick Wins," we created a knowledge base of localized success stories that appealed to the specific needs of different organizational departments. This enabled a transition from theoretical benefit to concrete task-specific optimization [7].

The primary research objective of this study is to assess the effectiveness of structured adoption processes in improving user expertise and ROI as a measure of organizational return on investment. By conducting a thorough analysis of the data collected from an AI Workshop Questionnaire, we sought to uncover the key success factors that distinguish a mature AI ecosystem. We focused on the degree to which a staged "Crawl-Walk-Run" strategy can counter the cognitive burden on users, thus ensuring the long-term viability of digital transformation.

In conclusion, this paper offers a comprehensive assessment of AI optimization, one that extends from the technical requirements of the software to the human and organizational factors that determine its success. We believe that this research has helped to advance the scientific

understanding of how generative AI can be systematically incorporated into the workplace to facilitate significant cognitive and functional progress [8].

II. METHODS AND MATERIALS

In order to scientifically assess the optimization of Microsoft 365 Copilot in a business setting, we adopted a mixed-methods research design. This allowed us to combine the results of quantitative diagnostic testing with qualitative pedagogical interventions. The research was conducted in two main phases: a diagnostic test using a standardized questionnaire, followed by a structured multi-stage instructional program called the "Art of the Possible" (AOTP) workshop. This two-pronged approach enabled us to measure the technical readiness level while also monitoring user adoption matrices.

A. Participant Selection and Cohort Structuring

The research employed a purposive sampling approach to build the initial deployment pool. Instead of launching the solution on a universal scale, we isolated a localized pool of 80 "Early Adopters" and internal "Champions" from various cross-functional departments, such as Human Resources, Marketing, and Operations. This pool was used as the experimental group to test the effectiveness of prompt engineering and isolate high-friction workflows.

B. Instrument Design: The AI Readiness Questionnaire

Before the intervention in the workshop, we administered an "AI Workshop Questionnaire" to evaluate organizational maturity on various digital fronts. This tool was designed to go beyond the subjective user excitement and instead measure the structural and technical readiness. We broke down the questionnaire into five different dimensions of operations, which became our key independent variables for evaluating the potential for optimization [9].

TABLE I. STRUCTURAL DIMENSIONS OF THE AI READINESS QUESTIONNAIRE

ID	AI READINESS QUESTIONNAIRE		
	Questionnaire Dimension	Primary Research Focus	Key Metrics Evaluated
1	Strategic Drivers	Organizational motivation for AI adoption	Generalized AI interest vs. specific Microsoft 365 integration intent
2	Target Scenarios	Departmental applicability	Expected impact on HR, Marketing, Operations, and Customer Service
3	Technical Architecture	Foundational software readiness	Current licensing, update channels, and device management protocols
4	Data Governance	Risk mitigation and compliance	Utilization of Data Loss Prevention (DLP) and sensitivity labelling
5	User Proficiency	Existing digital literacy	Frequency of cloud collaboration and automated workflow usage

Table 1 - Structural Dimensions of the AI Readiness Questionnaire

C. The "Art of the Possible" (AOTP) Workshop Framework

After the completion of the diagnostic process, we started the AOTP workshop series. The educational goal of these workshops was to move the users from a passive state of theoretical knowledge to an active state of application. The curriculum was designed to gradually decrease the cognitive load while increasing the complexity of AI interactions [10, 11].

The workshop was executed across four structured modules, as detailed in Table 2.

TABLE II. CURRICULUM STRUCTURE OF THE AOTP ENABLEMENT WORKSHOP

Phase	AOTP ENABLEMENT WORKSHOP		
	Module Phase	Pedagogical Objective	Core Activities and Demonstrations
1	Foundations	Establish theoretical literacy	Introduction to Large Language Models (LLMs), AI ethics, and Responsible AI frameworks.
2	Functional Demo	Showcase cross-platform utility	Live demonstrations of Copilot within Word, Excel, PowerPoint, and Teams environments.
3	Scenario Discovery	Contextualize AI to daily tasks	Collaborative mapping of departmental pain points into actionable "Persona" use cases.
4	Prompt Engineering	Optimize user-machine interaction	Practical exercises in constructing structured, context-rich queries to minimize AI hallucination.

Table 2 - Curriculum Structure of the AOTP Enablement Workshop

D. Governance Assessment and Risk Mitigation

One of the most critical aspects of our strategy was the aggressive audit of the data estate of the organization. We recognized that the generative AI model is optimized for data retrieval, which inherently leads to the risk of "oversharing" if the internal documents are not properly restricted. However, since large language models are based on semantic indexing to establish relationships across the entire Microsoft 365 tenant, traditional permission models, which are based on user obscurity rather than explicit restriction, are no longer adequate to protect sensitive information.

From the questionnaire, we were able to establish that the target infrastructure did not have automated Data Loss Prevention (DLP) protocols and sensitivity labels. Therefore, our strategy had to establish a "governance-first" baseline. Before allowing users to query the organizational data graph, we had to enforce the use of basic sensitivity labels. This targeted remediation sprint included the process of classifying existing documents based on clear confidentiality levels, such as public, internal, and highly restricted.

This ensured that the data collected for the subsequent optimization was done in a secure and compliant environment,

and not in an environment where the results were biased by the concerns of data leakage.

E. The Phased Implementation Model

To implement the results of the workshop and questionnaire, we aligned our deployment plan with a staged "Crawl-Walk-Run" approach to adoption.

1. "Crawl" Stage (Assess & Prepare): Primarily aimed at proving the scenarios identified in the AOTP workshops and supporting the initial 80 users.
2. "Walk" Stage (Execute): Primarily aimed at raising awareness, supporting peer-to-peer knowledge sharing through "Champions", and validating the quality of initial semantic indexing.
3. "Run" Stage (Optimize): Primarily aimed at full-scale deployment, where optimization was measured through tangible time-savings in the specific use cases identified (e.g., HR resume summarization, Marketing text creation).

Through the alignment of the quantitative readiness results (Table 1) with the qualitative results of the workshop scenarios (Table 2), this approach offered a comprehensive framework for measuring exactly how technical implementation and human training interact to optimize AI efficiency.

III. RESULTS AND DISCUSSION

The empirical evidence gathered from the diagnostic stage and the "Art of the Possible" (AOTP) workshops provided invaluable information on the organizational dynamics of the adoption of generative AI. The results of our study reveal an important gap between the strong user interest in the operational efficiency of AI and the weaknesses in the data architecture.

A. Quantitative Readiness and Data Governance Posture

The results of the AI Readiness Questionnaire provided a conclusive basis for understanding both the strategic intent and the technical maturity of the organization. Significantly, the results showed that the driving force behind the adoption of generative AI was not a generic, reactive response to the global AI phenomenon. Rather, the strategic intent of the organization was to have a specific, deliberate approach to integrating Microsoft 365 Copilot as a native extension of their existing productivity suite. This is a crucial point, as it indicates a transition from theoretical AI research to practical, workflow-integrated technological adoption with the goal of reducing user context-switching.

When asked about expected operational use cases, the respondents showed a universal acceptance for all business units polled: Customer Service, Sales & Marketing, Research & Development, Supply Chain & Operations, and Internal Finance/HR.

This universal acceptance suggests that the organization's workforce recognizes the value of generative AI not only as a vertical IT solution but also as a horizontal capability that can

be leveraged to drive efficiency in multiple operational areas. For example, while the Marketing function expected rapid content creation, the Supply Chain function recognized the potential for automated logistical forecasting and process documentation.

However, a key dichotomy that emerged was through the Data Governance section of the diagnostic questionnaire. Despite the strong user enthusiasm and readiness for adoption, the technical architecture was found to be extremely vulnerable to severe data protection risks.

The integration of Large Language Models (LLMs) in a corporate setting alters the paradigm of data exposure and synthesis. Microsoft 365 Copilot uses semantic indexing to index the entire data graph of the organization. Hence, if the data is not classified and controlled, the AI system has the potential to bypass the traditional security measures by exposing confidential data to unauthorized personnel [12, 13].

These structural issues were clearly identified through the results of the questionnaire. The organization was failing to protect data at the content level by not using rights management, nor were they using active Data Loss Prevention (DLP) protocols for cloud-based content. More seriously, the organization was failing to use active sensitivity labels to manage document access. In the context of generative AI, the lack of sensitivity labels immediately put the organization at risk of "oversharing." Without these cryptographic controls, a generic user request could easily ask the AI system to analyse and summarize files that were unrestricted but highly confidential, such as internal HR salary structures, unpublished financial data, or proprietary R&D information. Therefore, these results immediately made it necessary that a remediation sprint occur before any kind of deployment was possible.

TABLE III. BASELINE DATA GOVERNANCE AND RISK ASSESSMENT RESULTS

ID	DATA GOVERNANCE AND RISK ASSESSMENT		
	Governance Control Parameter	Current Organizational Status	Identified Optimization Risk
1	Content-Level Encryption	Not implemented	High risk of unauthorized data parsing by LLMs.
2	Data Loss Prevention (DLP)	Not deployed	Inability to track or restrict the external sharing of AI-generated sensitive data.
3	Sensitivity Labeling	Not deployed	High risk of internal "oversharing," allowing unauthorized employees to surface confidential HR or financial data via Copilot prompts.
4	Mobile Device Management	Implemented (Intune)	Baseline security met for external device access, reducing end-point vulnerabilities.

Table 3 - Baseline Data Governance and Risk Assessment Results

These results empirically verified our hypothesis that technical configuration had to occur before widespread user enablement. The absence of sensitivity labelling created an immediate risk of internal data oversharing once the semantic index of Copilot was turned on, requiring a delay in the full organizational rollout until a basic governance structure could be put in place.

B. Scenario Discovery and Operational Optimization

In the AOTP workshops, we were able to successfully translate the theoretical need for AI into actionable, localized workflows. By employing a persona-based framework ("As someone in [Role], I want to use Copilot to [Action], to achieve [Result]"), we were able to identify the precise pain points where AI optimization would provide the greatest Return on Investment (ROI).

The workshops provided us with a very specific set of operational scenarios. These were organized into a functional adoption matrix, as seen in Table 4, which became the roadmap for our prompt engineering modules.

TABLE IV. PRIORITIZED BUSINESS SCENARIOS IDENTIFIED DURING AOTP WORKSHOPS

Business Unit	BUSINESS SCENARIOS IDENTIFIED DURING AOTP WORKSHOPS		
	Target Persona / Role	AI Optimization Action (Prompt Focus)	Strategic Objective / Expected ROI
HR	HR Recruiter	Summarize candidate resumes and applicant profiles.	Shorten the hiring process and facilitate faster callback decisions.
Marketing	Content Creator	Generate advertising texts, relevant metatexts, keywords, and tags.	Rapid and efficient creation of tailored marketing collateral.
Marketing & Supply	Project Lead	Accelerate the creation of complex PowerPoint presentation structures.	Save time and resources while generating "food for thought" frameworks.
Cross-Functional	Meeting Facilitator	Generate meeting agendas, outlining main topics and training content.	Prevent omission of critical topics and achieve measurable time savings.
Cross-Functional	Knowledge Worker	Summarize lengthy PDF documents and convert generated Word text into PowerPoint bases.	Decrease time spent processing documentation and accelerate output formatting.

Table 4 - Prioritized Business Scenarios Identified During AOTP Workshops

The discovery of these scenarios was essential. We found that users who tried to use Copilot for general queries soon suffered from tool fatigue. However, users who used Copilot for the specific scenarios listed in Table 4 reported immediate

cognitive relief and time savings, thus establishing that the discovery of structured scenarios is a necessary antecedent for AI optimization.

C. Evaluating the "Crawl-Walk-Run" Implementation Strategy

To address the governance risks outlined in Section 3.1 while leveraging the high-value scenarios outlined in Section 3.2, we performed the deployment using the "Crawl-Walk-Run" approach.

Rather than a global deployment, we chose to enable Microsoft 365 Copilot for a carefully controlled initial set of 80 early adopters and internal Champions. This "Crawl" phase was closely tracked. We observed that a delayed assignment of licenses directly impacted the quality of Semantic Indexing in the first few weeks, as the AI model needed time to understand the relationships in the data graph of the organizational data.

At the same time, the IT leadership was challenged to complete an aggressive data governance sprint, cleaning up old external access and implementing sensitivity labels, before entering the "Walk" phase. This approach ensured that while the AI model's speed to uncover information increased, the security guardrails kept pace. This study asserts that achieving full user adoption and realizing ROI (the "Run" phase) will only be possible in late Q3 and Q4, as AI optimization is a long-term organizational marathon, not an IT sprint.

IV. CONCLUSION

The inclusion of generative artificial intelligence within the corporate environment is a paradigm shift in the way organization's function. Based on our analysis of the "Art of the Possible" (AOTP) workshop framework and the AI Readiness Questionnaire, we determined that the optimization of Microsoft 365 Copilot cannot be accomplished by technical implementation alone. To achieve true optimization, there must be a harmonious convergence of strong data governance, user enablement, and well-defined operational use cases.

Our research has shown that there is a high level of passion for AI, but there is a lack of the underlying data security infrastructure, such as Data Loss Prevention (DLP) and sensitivity labelling, necessary to implement these solutions in a secure manner. By recognizing this gap early in our diagnostic questionnaire, we have shown that a "governance-first" strategy is a necessary precursor to avoid the potential risk of internal data over-sharing [14].

Moreover, our results underscore the essential role of scenario discovery. Users interacting with Copilot via generic queries saw minimal improvements in productivity. However, when users started applying the tool to very specific, persona-related tasks such as summarizing HR applicant resumes or creating structured marketing materials, the ROI was immediate and quantifiable. This suggests that prompt engineering is more than a technical competence, it is a fundamental competence for the contemporary worker.

In conclusion, this study confirms the effectiveness of the "Crawl-Walk-Run" model of adoption. By limiting initial access to a controlled group of "Early Adopters" and

"Champions," organizations can deliberately observe the quality of semantic indexing and user feedback before rolling out to the broader enterprise. We propose that AI optimization is not a discrete IT initiative but a continuous, socio-technical process. Future studies should investigate the long-term cognitive effects of these tools on the creativity and task fatigue of the workforce over longer fiscal cycles.

ACKNOWLEDGEMENT

We would like to thank the IT leadership in the organizations, the internal network of AI Champions, and the early adopters who took part in the diagnostic questionnaires and workshops. Their open feedback and readiness to refine new operational processes were critical to the development of the findings of this research.

REFERENCES

- [1] E. Brynjolfsson i A. McAfee , The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies, New York: W. W. Norton & Company, 2016.
- [2] M. Corporation, „Will AI Fix Work?“, Microsoft Corporation, May 2023. [Online]. Available: <https://www.microsoft.com/en-us/worklab/work-trend-index/will-ai-fix-work>. [Last access: February 2026].
- [3] C. Vasilescu i M. Gheorghe, „Improving the Performance of Corporate Employees through the Use of Artificial Intelligence: The Case of Copilot Application“, u *Proceedings of the International Conference on Business Excellence*, Warsaw, 2024.
- [4] M. Mravik, T. Vetriselvi, K. Venkatachalam, M. Sarac i N. Bacanin, „Diabetes Prediction Algorithm Using Recursive Ridge Regression L2“, *Computers, Materials & Continua*, vol. 71, no. 1, pp. 457-471, 2022.
- [5] M. Šarac, M. Mravik, J. Dijana, Š. Ivana, Ž. Miodrag i B. D. Nebojša, „Intelligent diagnosis of coronavirus with computed tomography images using a deep learning model“, *Journal of Electronic Imaging*, vol. 32, no., pp. 1-10, 2022. doi: 10.1117/1.JEI.32.2.021406.
- [6] V. Venkatesh, M. G. Morris, G. B. Davis i F. D. Davis, „User Acceptance of Information Technology: Toward A Unified View“, *MIS Quarterly*, vol. 27, no. 3, pp. 425-478, 2003.
- [7] S. Noy i W. Zhang, „Experimental evidence on the productivity effects of generative artificial intelligence“, *Science*, vol. 381, no. 6654, pp. 187-192, 2023.
- [8] M. Corporation, „Microsoft 365 Copilot: Architecture and Data Governance Frameworks“, unpublished internal technical documentation, Belgrade, 2023.
- [9] M. Corporation, „What Can Copilot’s Earliest Users Teach Us About Generative AI at Work?“, Microsoft Corporation, November 2023. [Online]. Available: <https://www.microsoft.com/en-us/worklab/work-trend-index/copilots-earliest-users-teach-us-about-generative-ai-at-work>. [Last access February 2026].
- [10] N. Carlini, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, C. Raffel, D. Song, Ú. Erlingsson i A. Oprea, „Extracting Training Data from Large Language Models“, *USENIX Security Symposium*, 2021.
- [11] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith i D. C. Schmidt, „A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT“, *arXiv preprint*, p. arXiv:2302.11382, 2023.
- [12] I. D. W. Team, „AI Readiness and Copilot Implementation Questionnaire Results“, Quantitative organizational data, unpublished, Belgrade, 2024.
- [13] I. P. M. Office, „Art of the Possible (AOTP): Workshop Methodology and Scenario Discovery Framework“, Corporate adoption guide, unpublished, Belgrade, 2024.
- [14] A. Hasegawa i H. Shimizu, „The Role of Generative AI in Transforming Data Engineering Workflows and Automating Computational Infrastructure Design“, *AIJCS*, vol. 4, no. 6, pp. 1-11, 2022.