# Between Equality and Uniformity: Generative AI and the Ethics of Summarising Human Discourse in Organisations

*Miriam Klöpper, Norwegian University of Science and Technology, miriam.kloepper@ntnu.no*
*Isabell Lippert, Dresden University of Technology, isabell.lippert@tu-dresden.de*
*Stefan Klein, University of Münster, stefan.klein@uni-muenster.de*
*Patrick Mikalef, Norwegian University of Science and Technology, patrick.mikalef@ntnu.no*

## Introduction

Organisations are increasingly adopting generative artificial intelligence (GenAI) solutions to summarise and analyse internal and external communications. From customer service calls to employee feedback, these applications promise improvements in productivity and efficiency, even though it is unclear if those promises can be kept (Brynjolfsson et al., 2017). A prominent example of such solutions is Microsoft Copilot (Stratton, 2024), designed to assist employees in analysing and generating textual content, such as meeting summaries or email drafts. Copilot can also suggest improvements to clarity and tone. Some organisations prefer to develop their own tools over the usage of off-the-shelf tools, motivated by concerns over data privacy or dependence on external providers with servers outside Europe. However, even these in-house systems usually rely on pre-existing large language models (LLMs) as their foundational technology. In the context of analysing organisational data, inhouse solutions are argued to be less risky for employees and organisations than commercial mass-market offerings (Jörden et al., 2022).

While GenAI is developing rapidly, research investigating the social and ethical implications of organisational use of GenAI currently remains scarce. As employees increasingly attribute agency and rationality to GenAI, questions about responsibility and control become more prevalent and complex (Leonardi, 2025). Further, employees might feel the need to fact-check or rewrite the content provided by GenAI systems, leading to new, hidden tasks and, therefore, invisible labour (Klein & Watson-Manheim, 2021).

In the context of earlier workplace technologies, it has been highlighted that different groups of employees' also prefer different communication styles (Gefen & Straub, 1997), which might be relevant in the handling of GenAI systems which create textual output. Certain groups of employees might thus be more inclined to edit the text provided by GenAI, as it might not represent their natural communication patterns. GenAI systems have further been criticised for generating biased output, for instance, disadvantaging women (Gross, 2023; Locke & Hodgdon, 2025), which could affect how GenAI summarises employee communication.

However, research investigating how the analysed and summarised communication between employees and customers is perceived by different groups of employees is currently lacking. Both the employees whose communication with customers is summarised and analysed by GenAI, as well as the employees in higher organisational hierarchies who are presented with the GenAI summaries might face new challenges. Understanding the differences in the challenges and the diverging new practices that might emerge can help to support organisations and employees in deploying GenAI in meaningful ways that are in the interest of employees. Against this backdrop, we ask the following research questions:

**RQ1:** *How does the use of GenAI to summarise and analyse customer interactions reshape organisations' internal and external communication across different hierarchies?*

**RQ2:** *Which diverse practices are emerging from the work with generated summaries of customer interactions?*

To investigate these questions, we will conduct semi-structured interviews with 1. customer-facing employees in a Norwegian financial institution, whose interactions with the customers are automatically analysed and summarised by GenAI, and 2. employees in higher levels of the organisational hierarchy who are either reading the generated summaries to conduct an assessment of the customers cases, or who have access to the summaries to assess the work of the customer-facing employees.

**Method and Case**

This study adopts an explorative, interpretive research methodology, aiming to understand how customer-facing employees experience and make sense of the integration, usage and feedback of GenAI tools in organisational communication. As Travers (2001) highlights, interpretivist approaches are especially effective for exploring the complexity of human interactions, the influence of different social settings, and the diverse meanings that individuals attribute to their work. This perspective is particularly useful for studying emerging phenomena like the use of generative AI in organisations, where both practices and interpretations continue to develop. We follow a single-case study design, focusing on a financial institution in Norway that has recently implemented GenAI-based tools to support internal and external communication processes. Norway provides an appropriate context for our study. Compared to many other European countries, the Norwegian public tends to hold fewer biases or fears about AI technologies (Lloyd's Register Foundation, 2022). This cultural backdrop reduces the likelihood of algorithmic aversion (Dietvorst et al., 2015) or other negative perceptions influencing the interviewees. Therefore, we argue that this setting offers particularly rich insights into how GenAI is integrated into organisational communication, as perceived problems by the interviewees are less likely to be based on negative perceptions or prejudices.

We are conducting the study in their headquarters. GenAI is used by multiple groups of employees throughout the organisation, for instance by their data science team who deploy

it as coding assistant, or by employees to summarise internal video-calls with colleagues. The over-all attitude in the organisation is very optimistic towards using GenAI. We interview two groups of employees:

1. Frontline hotline staff. These employees provide customer support via the organisation's phone hotline, assisting customers of the financial institution with processes such as understanding loan applications or directing them to the appropriate specialist for specific inquiries. As the first point of contact for many customers, they traditionally created brief written summaries of each interaction. This task has now been automated through the use of GenAI. The summaries are, for instance, used to inform future phone calls with the same customers, and are made available to employees who work on the respective customers cases. Frontline hotline staff are responsible for verifying the accuracy of the generated summaries and may manually revise them as they see fit.

2. Supervisory and decision-making staff. The second group includes employees higher in the organisational hierarchy who use the generated summaries in their work. The summaries can be read by supervisory staff of the hotline workers, currently mainly in attempts to understand how precise the GenAI tool works and if workers are checking and correcting the generated summaries. Further, decision-making staff can access the summaries to better understand the nature of customers inquiries or problems, to see the history of interactions with the customer, and to generally access information they need to proceed with their own tasks.

**Preliminary Insights**

First, we have conducted two preliminary interviews with senior members of the organisations data science team to inform and refine the guidelines for our semi-structured interviews. The data science team has built the GenAI tool used by the frontline workers, and they have been responsible for educating them about potential risks. These initial conversations have helped shape a more focused and context-sensitive inquiry. Subsequently, we have begun to conduct interviews with the two targeted groups. Our interviews, which are still in progress, already highlight a strong tension: the manual adjustment of the generated summaries of customer interactions varies drastically among employees.

Previously, employees also varied in how they documented customer interactions: some consistently wrote long and detailed summaries, while others recorded only brief key points. Following the introduction of GenAI-generated summaries, all employees initially engaged in reviewing and correcting the automated texts to some extent. However, after a period of regular use, distinct patterns have emerged. Some employees now make only minor adjustments or none at all, while others continue to extensively edit or completely rewrite the summaries. Interestingly, the extent of these edits does not directly correspond to employees' previous documentation habits. For instance, some who formerly recorded only brief notes now revise nearly the entire GenAI-generated text, whereas others who previously wrote detailed summaries tend to leave the automated versions largely unchanged.

Additionally, the introduction of the tools highlighted employees' existing skills in conducting consistently flowing and logically structured conversations with customers. Conversations guided by employees who followed a clear structure were summarised more accurately by the GenAI. Some employees struggled with this, which became apparent through the summaries. This caused some early frustration, but it was resolved as employees proactively supported and educated one another, fostering new social interactions and dynamics within the team.

For the second group of employees, two new social patterns have emerged from our ongoing interviews. First, they have had to adjust their trusting behaviour towards the hotline staff. It is no longer possible for them to determine whether a summary has been manually edited or not. As a result, they must rely on the assumption that the summary has been checked and that all important information is correct and complete. Second, previously existing negative social interactions are decreasing. Caseworkers often had to contact the hotline staff about summaries, especially when only key points were provided. This frequently caused frustration. While this frustration has largely been eliminated, the reduction in such interactions has also led to fewer social exchanges between employees.

**Expected Contribution**

We expect to contribute theoretically by developing a framework for understanding how the use of GenAI to summarise customer interactions reshapes internal communication practices and knowledge work in organisations. The shift from employee-authored summaries—previously varying in tone and detail—to standardised machine-generated outputs introduces new dynamics around agency, authorship, control, and responsibility. Our study thereby focuses on the emerging divergence in employee behaviours: some increasingly accept and use GenAI-generated summaries with minimal or no changes, while others consistently rewrite them, maintaining a strong sense of ownership, which could, for instance, be grounded in fears of replacement or employees' wish to make meaning of their daily work. This enables us to explore how employees negotiate professional standards and their own role in meaning-making.

Secondly, we contribute by examining how different forms of interaction with GenAI influence social relations across workplace hierarchies, with lower organisational levels potentially experiencing greater mistrust from others. These insights will inform broader discussions on how algorithmic systems shape workplace practices and perceptions of legitimacy both among colleagues and in interactions with customers. We thereby aim to theorise how the tonal uniformity of GenAI outputs may influence downstream communication, particularly regarding its effects on organisational trust. Moreover, subtle shifts in tone or emphasis can shape how customers are perceived, raising important questions about how GenAI-mediated communication—such as customer files and profiles—affects organisational narratives and decision-making.

Practically, our study sensitises organisations and policymakers who are navigating the

adoption of GenAI in communication-heavy settings. While standardised summaries may appear efficient, they risk marginalising interpretive labour and embedding inconsistent practices across teams, as employees react differently to the summaries and thus deploy different levels of effort in adjusting and working with them. Yet, managers may assume uniformity in summary quality, when in reality, employees trust in GenAI varies significantly, impacting how and whether corrections are made. This has implications for internal coordination and the reliability of customer records. For policymakers, these findings point to the need for clearer guidelines around transparency, oversight, and linguistic fairness in the deployment of GenAI. Standards should address not only data privacy and bias, but also the more subtle ways in which AI may shape workplace communication norms and reinforce existing inequalities.

**References**


Brynjolfsson, E., Rock, D., & Syverson, C. (2017). *Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics*. National Bureau of Economic Research.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114.

Gefen, D., & Straub, D. W. (1997). Gender differences in the perception and use of e-mail: An extension to the technology acceptance model. *MIS Quarterly*, *21*(4), 389–400.

Gross, N. (2023). What ChatGPT tells us about gender: A cautionary tale about performativity and gender biases in AI. *Social Sciences*, *12*(8), 435.

Jörden, N. M., Sage, D., & Trusson, C. (2022). 'It's so fake': Identity performances and cynicism within a people analytics team. *Human Resource Management Journal*, *32*(3), 524–539.

Klein, S., & Watson-Manheim, M. B. (2021). The (re-) configuration of digital work in the wake of profound technological innovation: Constellations and hidden work. *Information and Organization*, *31*(4), 100377.

Leonardi, P. M. (2025). Homo agenticus in the age of agentic AI: Agency loops, power displacement, and the circulation of responsibility. *Information and Organization*, *35*(3), 100582.

Lloyd's Register Foundation. (2022). *World Risk Poll 2021: A Digital World Perceptions of risk from AI and misuse of personal data* (pp. 1–46). https://www.lrfoundation.org.uk/sites/default/files/2024-06/LRF_2021_report_a-digtial-world-ai-and-personal-data_online_version.pdf

Locke, L. G., & Hodgdon, G. (2025). Gender bias in visual generative artificial intelligence systems and the socialization of AI. *AI & SOCIETY*, *40*(4), 2229–2236.

Stratton, J. (2024). An introduction to microsoft copilot. In *Copilot for Microsoft 365: Harness the Power of Generative AI in the Microsoft Apps You Use Every Day* (pp. 19–35). Springer.

Travers, M. (Ed.). (2001). *Qualitative research through case studies*. SAGE.