



UNIVERSIDADE CATÓLICA PORTUGUESA

Exploring the Ethical Implications of Leveraging NotebookLM to Reduce Turnover:

An Exploratory Conceptual Study of Learning
Processes within Healthcare

by

Ana Catarina da Silva Vieira

Católica Porto Business School

July 2025



UNIVERSIDADE CATÓLICA PORTUGUESA

Exploring the Ethical Implications of Leveraging NotebookLM to Reduce Turnover:

An Exploratory Conceptual Study of Learning Processes within Healthcare

Master's Final Assignment

Presented to *Universidade Católica Portuguesa*
to obtain a Master's Degree in Management

by

Ana Catarina da Silva Vieira

Under supervision of
Dr. Professor Andrea Thrope

Católica Porto Business School
July 2025

Agradecimentos

À minha mãe,
Que me ensinou sobre resiliência.

Aos meus irmãos, Nelson e Paulo,
Que me permitiram sonhar, mesmo quando dormir era desafiante.

Ao meu primo, 'Jokinha',
Com quem partilho as memórias mais bonitas da minha infância, adolescência e idade adulta – desde o luto e o medo à felicidade e sonhos. Hoje, acompanhamo-nos na aventura da vida – ainda que a 2000km de distância – e celebramos cada objectivo concretizado.

Ao meu namorado, Manuel,
Que acredita no meu potencial e me mostra, todos os dias, que sou do tamanho dos meus sonhos.

À Catarina, Cristiana, Cristina e Oriana,
Que se tornaram família e me acolheram.

Ao meu 'Francis',
A quem prometi aguentar o mundo nas costas para que o pudesse ver crescer.

A todos os meus 'piquis', Beatriz, Carol, Alice e José,
Que tornam a minha vida mais feliz.

A todos os meus amigos, Ana, Bárbara, Beatriz, Dara, David, Diogo, Edu, Maria V. e Maria J., Rui, Wei.
Que me acompanharam na escola e, agora, na vida.

And to my amazing advisor, Andrea, to whom I thank for the weariless support throughout this process,

Obrigada. Esta conquista também é vossa.
Thank you. This achievement is yours too.

Resumo

A Inteligência Artificial tem demonstrado fortes capacidades para a otimização de operações. Este trabalho reflete um enquadramento conceptual baseado nas três principais lentes éticas – consequencialismo, deontologia e ética das virtudes – para responder à questão ‘É ético utilizar o NotebookLM para reduzir a taxa de rotatividade no seio das organizações de saúde?’. Fundamentada na literatura disponível sobre Inteligência Artificial e Educação, começo por discutir as possibilidades de utilizar a plataforma para o desenvolvimento de programas personalizados para aprendizagem e formação em organizações de saúde, objetivando fomentar a motivação e, assim, reduzir a rotatividade – questão crítica em ambientes de alta pressão e que resulta em grandes perdas monetárias para as organizações. Compreendendo os potenciais riscos e benefícios da implementação sob as lentes mencionadas, este estudo fornece, em última análise, recomendações práticas para a mitigação de riscos. No final, conclui sugerindo que as empresas adotem uma abordagem híbrida – tirando partido da tecnologia, sem desvalorizar o raciocínio humano – prestando foco na intenção, conforme é proposto pela ética das virtudes.

Key-words: Aprendizagem, Consequencialismo, Deontologia, Educação, Ética, Ética das Virtudes, Inteligência Artificial, Inteligência Artificial Generativa, NotebookLM
Palavras: 7033

Abstract

Artificial Intelligence has demonstrated strong capabilities for operations' optimization. This paper builds a conceptual framework grounded on the three main ethical lenses – consequentialism, deontology, and virtue ethics – to address the question 'is it ethical to leverage NotebookLM to reduce turnover in healthcare organizations?'. Drawing on the available literature on Artificial Intelligence and Education, I begin by discussing the possibilities of leveraging NotebookLM to develop customized learning and training programs within healthcare organizations, aiming to foster motivation and, consequently, reduce turnover – a critical issue in high-pressure environments that also reflects great financial losses for organizations. Understanding the potential risks and benefits through the mentioned ethical standpoints, the study ultimately provides actionable insights for risk mitigation. Finally, concludes by suggesting firms opt by a hybrid training approach – leveraging from technology, whilst respecting human reasoning – simultaneously focusing on intention, as proposed by virtue ethics.

Key-words: Artificial Intelligence, Consequentialism, Deontology, Education, Ethics, Generative Artificial Intelligence, Learning, NotebookLM, Virtue Ethics
Word Count: 7033

Contents

Resumo.....	vi
Abstract.....	viii
Contents.....	x
Figure Index.....	xii
Table Index.....	xiv
Introduction.....	17
Artificial Intelligence and People Management: Understanding Human Learning Processes to Foster Organizational Growth.....	20
1. Learning Mechanisms and Continuous Learning.....	20
1.1 Neural vs Cognitive Learning Mechanisms.....	20
1.1.1 Neural Mechanisms.....	20
1.2 Continuous Learning.....	23
2. Enhancing Learning with GAI.....	28
2.1 NotebookLM: User Guide.....	31
2.2 Lessons from Education.....	35
Exploring Ethics: NotebookLM in Healthcare Organizations.....	38
1. Moral Compass:Unpacking Consequentialism, Deontology & Virtue Ethics	38
2. Ethical Crossroads: NotebookLM in Healthcare.....	40
2.1. The Ethical Dilemma of Applying NotebookLM for Training in Healthcare.....	41
Leveraging NotebookLM to Reduce Turnover: Is It Ethical?.....	45
Final Thoughts.....	50
Bibliography.....	51
Appendices.....	59
1. VARK Questionnaire.....	59

Figure Index

Figure 1: Home Page	31
Figure 2: Source Upload	32
Figure 3: Interface – Sources, Chat and Studio	32
Figure 4: Studio	33
Figure 5: Chat	34
Figure 6: Sources	34
Figure 7: Saved Notebook	35

Table Index

Table 1: Principles of Consequentialism, Deontology and Virtue Ethics	38
Table 2: Practical Implications of Implementing NotebookLM for Training Purposes	48

Acronym List

AI – Artificial Intelligence

DL – Deep Learning

FAQ – Frequent Asked Questions

GAI – Generative Artificial Intelligence

HR – Human Resources

LLM – Large Language Models

ML – Machine Learning

REK – Reliable External Knowledge

Introduction

Artificial Intelligence has demonstrated strong capabilities for operational improvements across industries. Particularly for education, Generative Artificial Intelligence has been increasingly utilized, mostly motivated by the COVID-19 pandemic, which required re-shaping the way education is delivered (Baidoo-Anu and Owusu Ansah, 2023). Even though there is no consensus on the neural mechanisms behind human learning (Nayebi et al., 2020, Hebb, 1949 Hubert et al., 2013, O'Rilley et al., 2021), literature broadly agrees that, in terms of cognitive mechanisms, or 'learning styles', as defined by Sintia and Samsudin et al. (2019), learners can be divided into four categories: Visual, Aural, Read/Write (Flemming 1987). Thus, tailoring teaching materials and strategies can positively impact learning outcomes by fostering student engagement, improving motivation and supporting teachers to develop educational materials (Bernard et al. 2022, Latif et al., 2021). This hypothesis was later empirically corroborated by Huffman and Hutson (2024) at Lindenwood University. The authors showed that NotebookLM, an online research assistant powered by Generative Artificial Intelligence, could be used to transform static historical documents into dynamic learning resources, offering a cost-effective strategy to enhance the learning process.

However, whilst the use of Artificial Intelligence is widely accepted in education, it is not as well understood in healthcare. Despite fostering a learning environment is equally as critical, healthcare is an intrinsically complex and ethically sensitive domain (Trocin et al., 2023) where 'learning' also implies ensuring the safety of patients and workers, understanding procedures and cultivating empathy when delivering care. The tool I discuss in this piece,

NotebookLM, has been shown to have a transformative potential in medical operations (Chan and Hu, 2023), e.g., by supporting patient diagnosis through image analysis and treatment recommendations (Rigby, 2019, Keshavarz et al., 2024 and Kadoya et al., 2024, Gaonkar et al., 2020, Yelne et al., 2023). Thus, it is reasonable to assume there is further potential for application in matters like training and onboarding, typically managed by Human Resources (Nguyen et al., 2023, Sivasubramanian et al., 2023). Currently, this has not been explored. As such, I aim to explore the ethical implications of implementing NotebookLM to enhance organizational learning processes. Additionally, understand its potentially positive implications for employee motivation, which in turn can result in turnover decreasing, i.e., lowering the rate at which employees leave organizations and generating a need for their replacement. This is relevant because Cascio's (1991) model, later refined by Tziner and Birati (1996), brought to attention that financial losses associated with dysfunctional turnover significantly impacted companies. For healthcare in particular, the funds 'saved' by reducing turnover through motivation increase could be allocated to enhance patient care. Nevertheless, ethical implications emerge as healthcare looks to expand its use of Artificial Intelligence, which will likely involve ethical dilemmas for managers and key decision makers seeking out efficiency for their organizations.

In order to address this gap in knowledge, I ask, "Is it Ethical to Leverage NotebookLM to Reduce Turnover in Healthcare Organizations?". Hence, as a starting point, I focus on explaining the mechanisms which play a part in human learning. Drawing on the discussion of how NotebookLM has been leveraged in education to increase student engagement and motivation, I theorize on the potential for the tool to develop customized learning and training programs within healthcare organizations. Consequently, fostering motivation and reducing turnover.

In addition, I explore the ethical implications posed by using AI-based technology, particularly NotebookLM, on medical training programs by applying the three main ethical frameworks – consequentialism, deontology, and virtue ethics. Despite their theoretical differences, these approaches actually provide complementary insights for morality assessment, which allowed me to develop a strategic mitigation framework applicable not only to healthcare organizations, but across various industries as well.

This thesis aims to contribute to furthering the current understanding of how Artificial Intelligence can support learning processes within healthcare organizations – including hospitals, clinics, rehabilitation centers, laboratories – especially those facing higher turnover rates. Additionally, drawing a conceptual comparative analysis through the three major ethical frameworks. Practically, as aforementioned, it culminates in a guiding framework for the practical implications of implementing NotebookLM to enhance onboarding and training processes in healthcare settings. Ultimately, intending to offer key decision-makers the necessary knowledge on how to balance technology integration and ethical responsibility. Finally, proposing the use of a hybrid training model grounded on intention.

Literature Review

Artificial Intelligence and People Management: Understanding Human Learning Processes to Foster Organizational Growth

1. Learning Mechanisms and Continuous Learning

The purpose of this section is to elucidate on the human learning experience. Particularly, how neural, and cognitive mechanisms work together to allow individuals to acquire and retain knowledge, as well as highlighting the importance of continuous learning within organizations to foster growth.

Therefore, I will dedicate this part to exploring these mechanisms by understanding how they operate together to enable human learning. Additionally, explaining how Large Language Models mimic neural mechanisms in their training.

1.1 Neural vs Cognitive Learning Mechanisms

Researchers recognize that human learning occurs through a series of mechanisms. Neural and cognitive based phenomena are especially prevalent in literature and give an insight into mechanisms that are intrinsic to individual learning, i.e., rather than those relating purely to extrinsic or 'social' influences.

I will, now, outline these.

1.1.1 Neural Mechanisms

The first approach within literature concerns neural mechanisms, which relate to the biological bases of learning, i.e., the tangible aspects of the brain. More

specifically, the idea that “the brain modifies its synaptic strengths during learning in order to better adapt to its environment.” (Nayebi et al., 2020). This concept, known as ‘neuro plasticity’, describes how the brain responds and adapts when faced with new information in a process of ‘learning’.

The ‘factual’ rules governing the neural plasticity process are relatively unknown. However, both Hebbian Learning and Error Backpropagation offer insights into the mechanisms underlying synaptic adjustment, which are outlined below.

Hebbian Learning is defined by the principle of "cells that fire together, wire together" (Hebb, 1949). In short, it “transcribes the correlations between couple of neurons within their connecting synapse” (Hubert et al., 2013). Thus, suggesting that learning is an associative, subconscious process, deeply embedded in neurological systems. Interestingly, the “Hebb Learning Rule” stands as an unsupervised learning rule, capable of extracting features from training sets and classifying data according to similarity, proposing the earliest idea of Machine Learning. (Jiang et al., 2022).

Alternatively, Error Backpropagation is error-driven. The theory presumes that “learning always happens whenever something unexpected occurs, at any point, and drives the development of predictions immediately prior, to the extent such predictions are possible to generate.” (O’Rilley et al., 2021). Briefly, this means that human learning happens whenever an individual is faced with unforeseen circumstances. ‘Surprise’ triggers the brain to adjust, improve, and predict similar events. Similar to Hebbian Learning, error backpropagation has been adapted in Machine Learning, demonstrating great effectiveness during the training step of neural networks, particularly in Deep Learning (DL), since it reduces the gap between predicted and actual outcomes.

1.1.2 Cognitive Mechanisms

The second approach slightly moves the focus from the tangible biological aspects of the brain, and onto cognition. Thus, the primary concern of the research is exploring how learning takes place in both human and non-human animals. Primarily in terms of the former, the cognitive mechanisms which reflect the preferences in learning and information processing of individuals are defined by “learning styles” (Sintia and Samsudin et al., 2019).

Flemming (1987) introduced a categorizing model to label individuals based on their preferred ‘learning style’. The model consists of a questionnaire that differentiates Visual, Aural, Read/Write and Kinesthetics (VARK¹) learners. Essentially, visual learners thrive when information is presented in graphs, flowcharts, and images; aural learners prefer listening to lectures, participating in discussions or explaining information; read/write learners excel by reading texts multiple times and taking notes; and kinesthetics perform better through hands-on experiences; Individuals who blend multiple learning styles are classified as “multimodal” learners (Marcy, 2001 & Fitvok-Norris & Yeghiazarian, 2015).

Understanding the nuances of this classification is essential, as “learning style-based adaptive educational systems can be designed, which have been shown to increase student satisfaction or learning gain, while reducing time needed to learn.” (Bernard et al. 2022). Moreover, Latif et al. (2021) revealed that educational contents transformed into multimedia format, i.e., podcasts, videos, and interactive dialogue, enhance information retention in students when compared to traditional teaching methods. Hence, it is logical that such digital means have evolved into the territory of AI. As Huffman and Hutson (2024)

¹ Can be consulted in the *Appendences*

argue, “the use of AI-generated multimedia materials has been gaining attention for its potential to transform educational engagement”.

Regardless of the neural mechanisms underlying human learning, which machine learning mimics, cognitive mechanisms remain critical, as they highlight the various stimuli that effectively facilitate neural processes in different individuals, which leads to learning and information retention.

1.2 Continuous Learning

Juxtaposed to individual learning is the idea that organizations themselves also have the potential to become ‘learning organizations’ as they evolve and adapt through changes in the status quo. Somunoglu and Ofluoglu (2012) propose that efficiency, competitiveness, and turnover are directly related to a firm’s adaptability to a fast-changing environment, which is currently exacerbated by technology. The authors emphasize that “apprehending the value of knowledge, accessing knowledge, and enhancing knowledge” is critical as ‘learning organizations’ are particularly more resilient (and so, superior) than those which diminish the importance of continuous learning. Thus, arguing that prioritizing knowledge-sharing is key to long-term success. Atak and Atik (2007) further reflect that providing education to stakeholders plays an important part in the “effective implementation of the mission and vision of an organization”.

1.2.1 The High Price of Turnover

‘Learning organizations’ build a foundation for motivated stakeholders which, in theory, contributes to lower turnover. Given the high price associated with staff loss, retention should be a strategic priority for companies.

Cascio (1991) significantly contributed to the assessment of turnover costs by developing a mathematical model capable of estimating turnover financial loss. The author proposed an arithmetic sum of separation, replacement, and training

costs², as he believed these constituted the “de facto outlay” of personnel loss. Furthermore, he emphasized the cost of reduced productivity, defined by the period in which the replacement worker would reach the same level of performance as the previous employee.

Tziner and Birati (1996) furthered the discussion by affirming that “although Cascio’s model constitutes an important contribution to a previously neglected issue, it seems to suffer from several flaws”. Namely, the lack of accountability regarding the differences between functional and dysfunctional turnover. Factually, the costs associated with letting high and low performers go are fundamentally different and should not be evaluated equally: “If poor performers are encouraged to leave voluntarily or are laid off, turnover produces beneficial results by furthering the attainment of a functional goal. (...) However, if good performers choose to quit or are terminated due to downsizing, the turnover will engender negative effects”.

The authors then addressed the latter proposal’s flaws, suggesting a revised model for turnover cost estimation. The consolidated model considers both direct and indirect costs, while also introducing the turnover multiplier. Essentially, direct costs encompass the immediate expenses associated with employee replacement, i.e., difference in compensation between the new hire and the veteran, recruitment costs, selection and training outlays, as well as socialization costs (related to welcoming and integrating the new hire). Indirect costs include, in the tangible spectrum, the overtime pay for those temporarily assuming the responsibilities of the prior worker, monetary loss as production decreases and potential service disruptions (e.g. delivery issues). Additionally, considering the

² Cascio (1991) defined separation, replacement and training costs as follows:

Separation Costs: total financial value associated with an employee departing. Composed by the value of time spent both by the interviewer and the employee during the exit interview, the administrative expenses of removing him from the payroll and severance pay.

Replacement Costs: include the cost of job advertising and all recruitment and selection activities.

Training Costs: the required investment for onboarding a new hire, such as cultural integration (socialization) and technical training.

impact on morale of the remaining staff, which, although we assume as a particularly intangible aspect, can affect productivity. Thus, arguably causing a measurable impact. The turnover multiplier refers to the assumption that new hires are often more prone to leave the organization than veterans, as corroborated by Kurnat-Thoma et al. (2017) who investigated a Washington DC's medical facility.

Cascio's (1991) initial model and the subsequent developments from Tziner and Birati (1996) were important contributions in that they were the first to theorize turnover costs into coherent models. Later, empirical research was added to the conversation.

Waldman et al. (2004), for instance, developed an experiment in a major healthcare center aiming to accurately measure the operational burden brought by dysfunctional turnover. By observing the available databases with information on selection, hiring, training, working and termination, the authors realized that turnover represented between 3.4 % and 5.8 % of the annual operating budget for the facility assessed, which translated to \$17 – 29 million solely for turnover expenses. The analysis also uncovered that different work groups behaved differently (cost-wise) when evaluating turnover. Physicians, nurses, and the remaining support staff do not incur the same turnover costs, hence, the width of the calculated range.

1.2.2 Continuous Organizational Learning and AI

Promoting a culture of continuous improvement and innovation, as well as earning the 'badge' of 'learning organization', requires that companies prioritize investing in training programs. Individuals can only do so much to further develop their education as continuous improvement through formal (paid) programs is often too expensive for them to enroll in merely to fit a job's

requirements. Therefore, the most effective way to ensure businesses have highly qualified staff is to invest in their employees' development.

Fortunately, as AI becomes more sophisticated, it is now possible, for example, to create simulated environments for developing managerial skills. As Mulongo (2024) describes, "investing in AI-driven management training boosts individual competencies and fortifies organizational resilience and adaptation". Hence, using AI in this context has the potential to contribute to an organization's competitive advantage.

1.3 Defining Generative Artificial Intelligence

Generative Modeling Artificial Intelligence (GAI) encompasses an unsupervised (or partially supervised) machine learning framework capable of generating manmade relics through statistics and probabilities. (Hu, 2023 & Jovanic, 2023). Put simply, GAI is a subtype of AI that 'learns' by recognizing patterns in existing data through the application of a mathematical model and uses the gathered knowledge to generate new content (images, text).

Large Language Models (LLMs) are a key application of this technology as they are "designed to understand natural language and generate human-like responses" (Alec et al., 2018). Operating through Hebbian Learning or Error Backpropagation mechanisms (similar to human cognition, as discussed above), LLMs have achieved "considerable success in a variety of domains." (Alec et al., 2018).

Building on the theory discussed, this chapter begins exploring how AI is currently used across various operations within organizations. It then emphasizes the particular applications for onboarding and training, i.e. learning' processes, which are typically managed by HR. Furthermore, assesses the capabilities of NotebookLM by illustrating a successful real-world application in the Education sector.

1.3.1 Artificial Intelligence in Large Organizations

Alsheibani et al. (2020) consider that AI positively impacts an organization's business value by increasing revenue, reducing costs, and improving efficiency. AI receptivity is, however, directly linked to a firm's functional structure. "Departments who keep relevant data to themselves, an overreliance on the status quo, as well as slow and bureaucratically shaped corporate structures will have a negative effect on the adoption of AI in companies" (Pumplum et al., 2019). Considering this, Enholm et al. (2021) suggest that it is likely that agile structured companies would adapt better to AI implementation in operations due to their flexible nature.

Lacking coherent knowledge on how AI can generate value, most companies fail to effectively exploit AI benefits and actually implement it within their daily operations. (Enholm et al., 2021). Nevertheless, AI applications across various fields, such as medicine, engineering, and data analysis, have gathered significant attention.

1.3.2 Artificial Intelligence in Human Resources

AI can be immensely valuable in an HR context. As Sivasubramanian et al. (2023) argue that "the application of AI helps to enhance the quality of employees through training and development at par with industrial needs (...) in the functional areas of the human resource segment such as recruitment, training, performance assessment, and employee retention".

Particularly in the training context, during the onboarding stage, standardized practices may discourage curiosity and diminish employees' motivation to learn. Assessing the individual stakeholders' learning styles, for example, through the VARK questionnaire, allows companies to design captivating and, therefore, more effective programs. Potentially, boosting productivity and decreasing turnover, especially among new hires.

Nguyen et al. (2023) leveraged AI to develop flight crew training materials. The experiment demonstrated that programs developed through AI offered a more cost-efficient approach to knowledge sharing, as they required less time and fewer human resources to design. Thus, although it may seem utopian, the study showed that it is possible for organizations to use AI to develop learning content customized to each individual's learning style. It also triggered a reflection on the democratization of the competencies among practitioners, i.e., the possibility of making the required skills more accessible to everyone in the field.

Naturally, transferring the responsibility of designing the learning programs from humans to AI can lead to a replacement of 'traditional' knowledge. Thus, introducing the need for a skill set focused on AI monitoring and output checking. In the long run, this may require a complete organization restructuring of organizational roles and personnel, potentially affecting overall AI acceptance and adoption.

2. Enhancing Learning with GAI

Due to GAI's "ability to handle complex prompts and produce human-like outputs", effectively integrating AI tools can positively impact "various fields such as healthcare, medicine, education, media, and tourism." (Chan and Hu, 2023). For education, for example, Ayala-Pazmiño (2023) advocates for its "potential to revolutionize teaching and learning, offering new ways to enhance personalized learning". Understanding this potential, and driven by the growing GAI revolution, technology-based enterprises competed to develop increasingly innovative and high-performing tools aiming to offer users a great experience in a variety of applications, e.g., ChatGPT, DALL-E, Stable Diffusion, Bard, and

NotebookLM. Typically, these offer specialized features for paying clients – constituting a “freemium” business model to secure both competitive advantage and a revenue stream.

ChatGPT, for instance, provides an online canvas for writing, coding, editing, and revising. Although the training data fed into the model is static, and “as a result, ChatGPT doesn’t know everything about specialized topics and current events” (OpenAI, n.d.), paying users can enable the deep research mode, specifically designed to “perform in-depth, multi-step research using data on the public web” (OpenAI, n.d.). Additionally, by incorporating image processing and creation (powered by Dall-E), the tool proves to be useful for a wide range of projects. Nevertheless, because of the nature of the information fed into it, the outputs need to be fact-checked; Similarly to Dall-E, Stable Diffusion “is a latent text-to-image diffusion model capable of generating photo-realistic images given any text input” (Stable Diffusion Web, n.d.). Depending on the user’s subscription, these images’ styles can also be adapted, making it a great tool for educational and media purposes; Bard, later rebranded as Gemini, “handles logical reasoning, creative writing, and code creation. The chatbot can generate text that resembles human writing in various styles, such as stories, poems, and scripts, making it a valuable tool for creativity.” (Perplexity AI, n.d.). Building on these capabilities, Google launched NotebookLM – a virtual research assistant capable of generating summaries, timelines, FAQs, study guides and audio overviews from a variety of sources (PDFs, YouTube URLs, Google Docs, Texts...). NotebookLM automatically gains expertise on the uploaded sources, making it a valuable tool for reading, taking notes and refining/organizing ideas (Google, n.d., para. 1-3). Moreover, the platform can be prompted to generate a podcast-like version of the content provided as input.

Due to its versatility, NotebookLM supports users’ knowledge retention independently of their preferred learning style, whether Visual, Aural, or

Read/Write, by creating tailored materials for everyone's needs. Particularly for kinesthetic learners, NotebookLM can be prompted to elaborate hypothetical scenarios, providing an adequate framework and/or stimulating reflection regarding the different courses of action. Later, as in any other professional environment, participating in shadowing moments to observe the ways of working of more experienced personnel would be crucial for kinesthetic workers to become familiar with the institution's routines. In corporate settings, this could translate into attending meetings to understand team dynamics or become familiar with the decision-making process. In healthcare settings, it could mean accompanying nurses/doctors/physicians during appointments to practice interactions, learning the facility's clinical practices and procedures. Put simply, gather hands-on experience through observation.

Now, we narrow the focus towards NotebookLM, exploring its capabilities for learning processes' enhancement by outlining the available features, and presenting the aforementioned successful application in Education.

2.1 NotebookLM: User Guide

Similarly to the priorly discussed tools, NotebookLM's paid and free versions offer the same set of features at different levels of performance. Figures 1-8 display the functionalities available for non-paying users.

Arriving at the Home Page (Figure 1), the user sees an overview of the core features and can proceed by clicking in "Create new notebook".

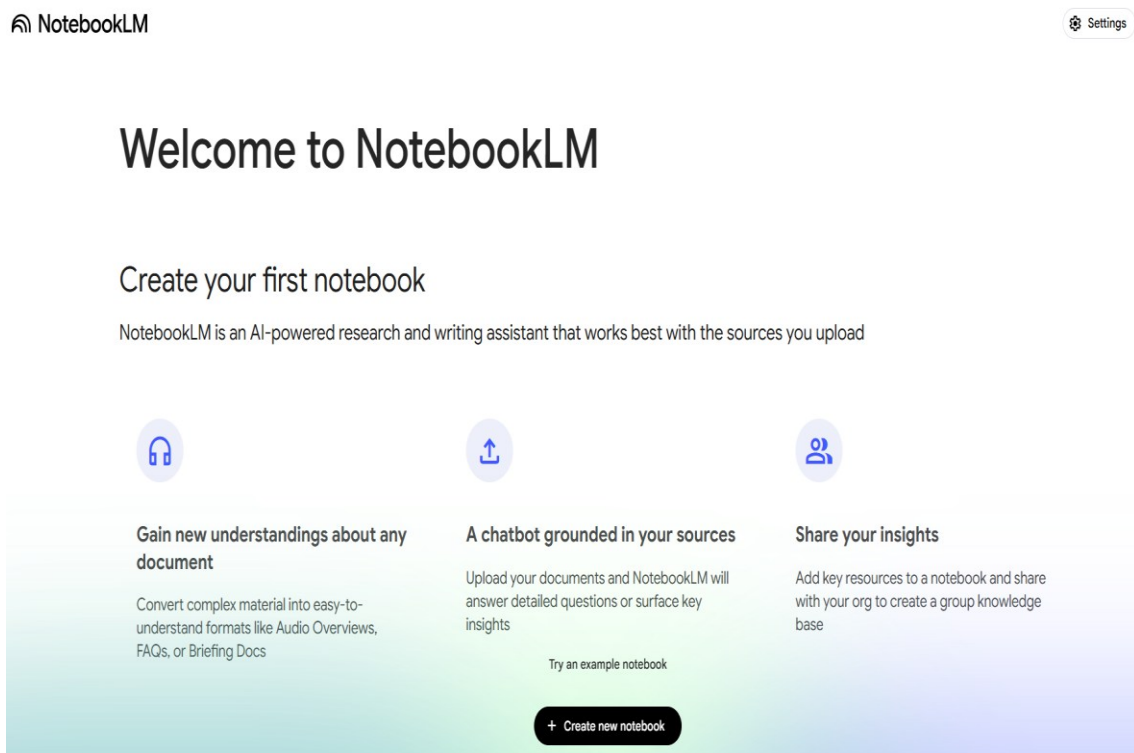


Figure 1: Home Page

A new tab appears (Figure 2) prompting the user to upload up to 30 sources in various formats, as above-mentioned. The researcher is then redirected to an environment composed of three panels – Source, Chat and Studio (Figure 3).

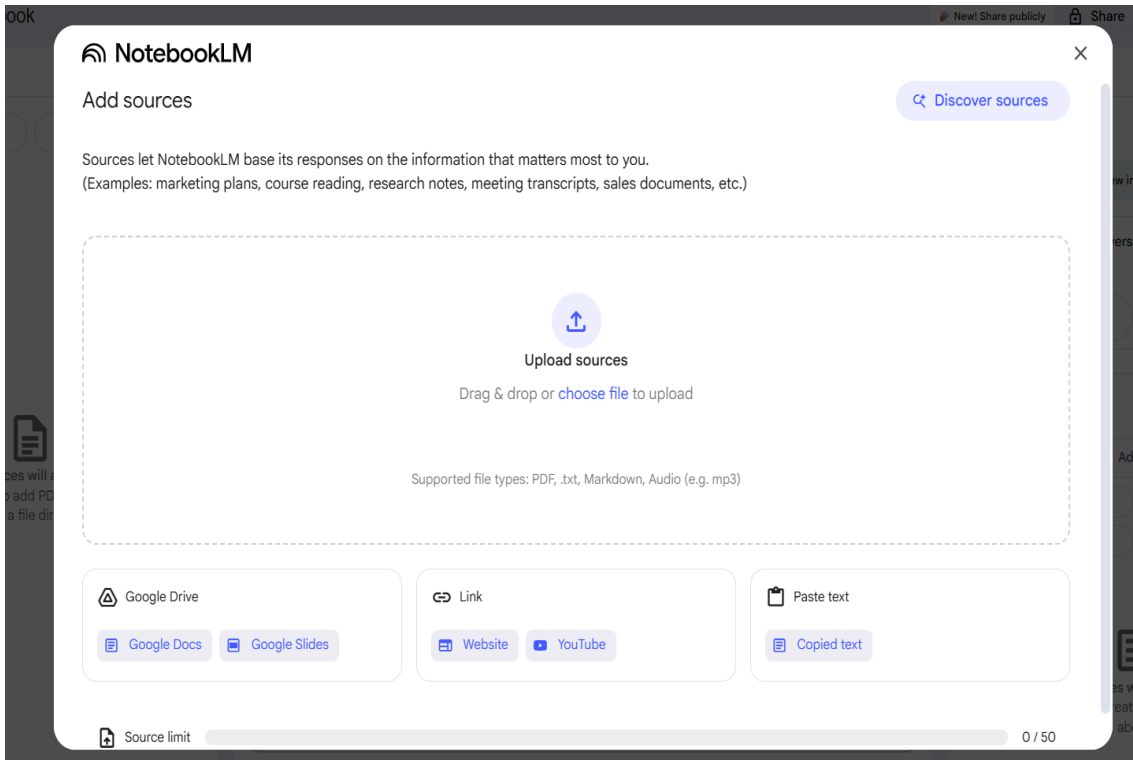


Figure 2: Source Upload

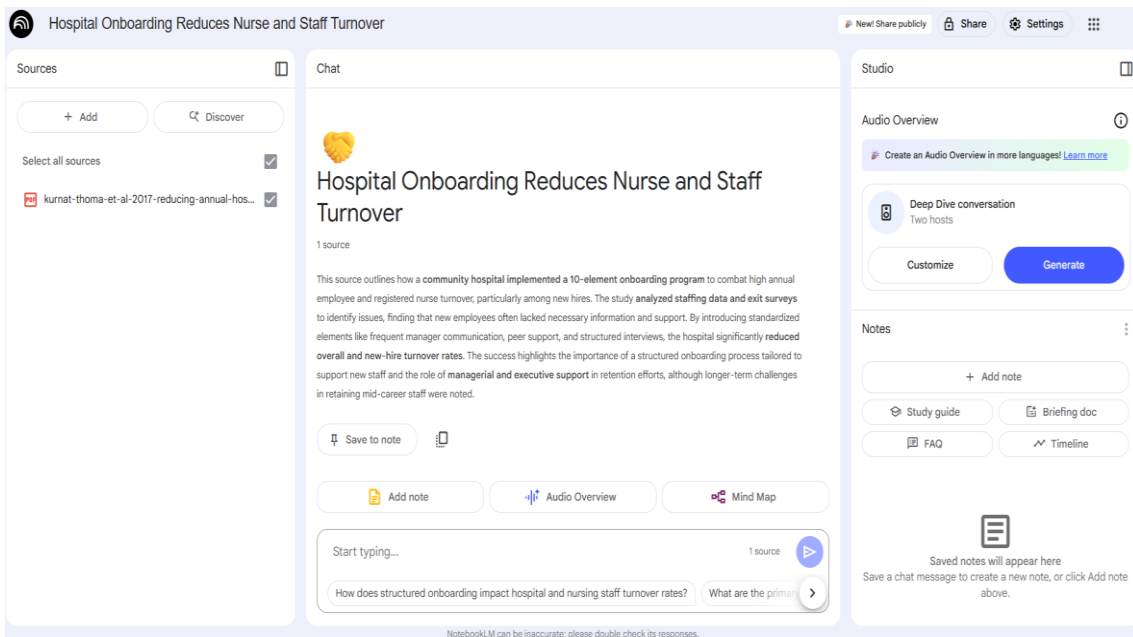


Figure 3: Interface – Sources, Chat and Studio

The “Studio” panel (Figure 4) stores the notes created by the user and allows the development of study guides, briefing docs, FAQ, and a timeline of the source’s events. This section further allows users to create a podcast based on the uploaded documents.

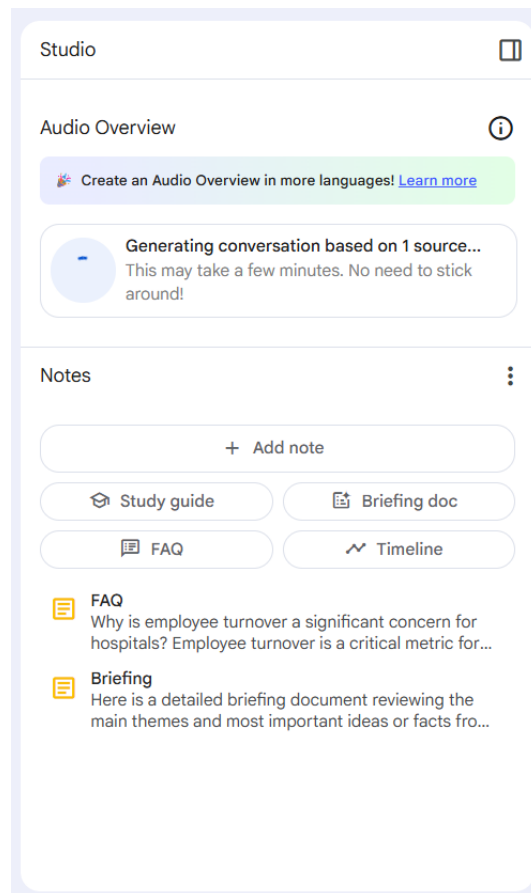


Figure 4: Studio

The “Chat” section (Figure 5) works as a typical chatbot, allowing the user to clarify questions, which can be saved as notes and directly stored under the latter panel. The answers can be fact-checked since their sources are listed and can be easily consulted, as shown in Figure 6.

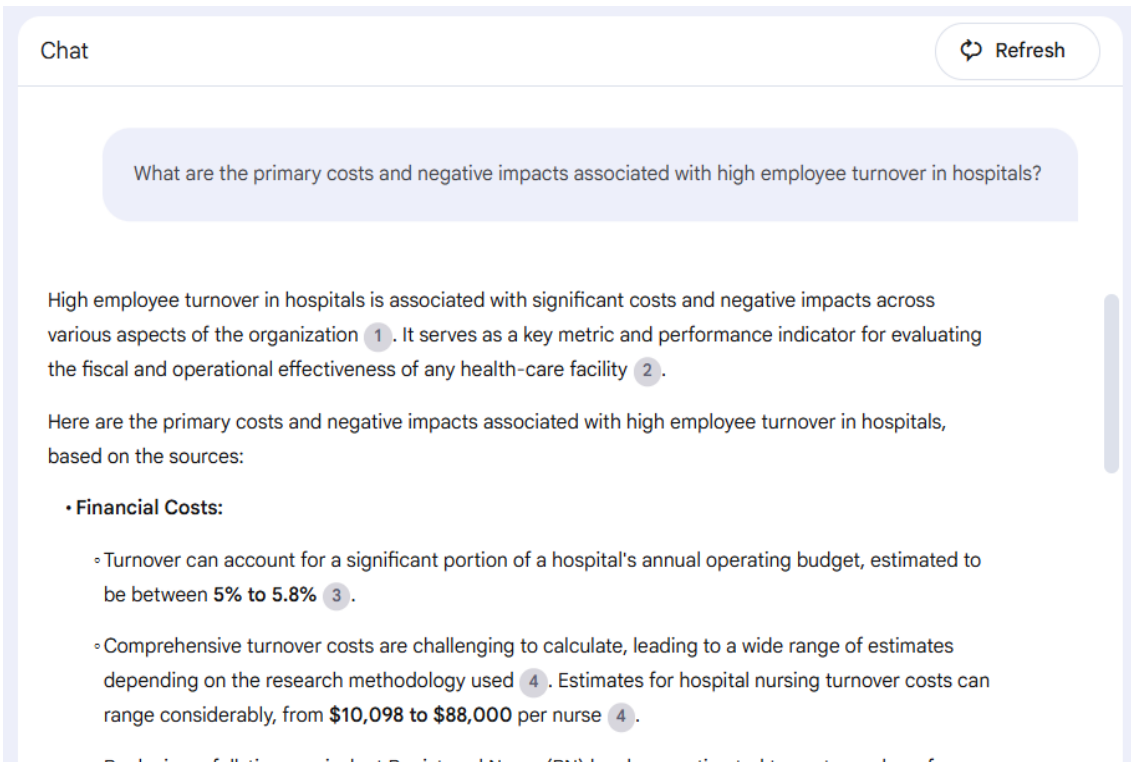


Figure 5: Chat

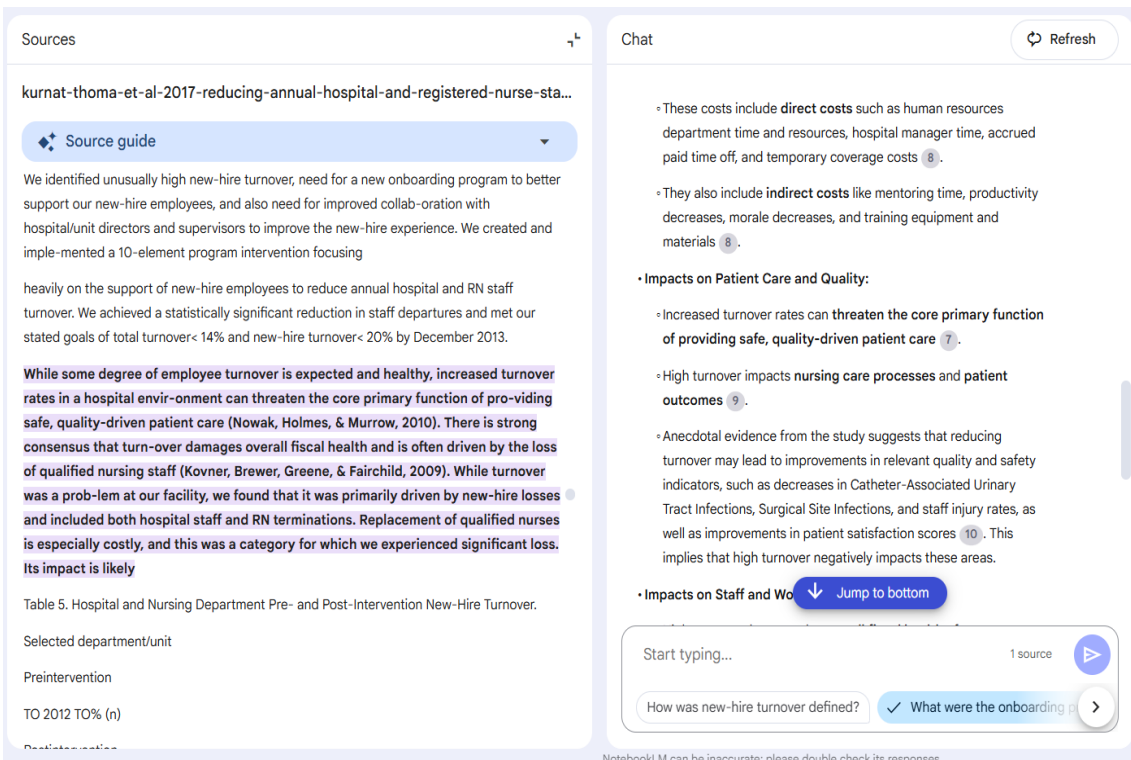


Figure 6: Sources

Finally, the notebook is saved (Figure 7) and can be easily accessed at any time. The user can create multiple notebooks for different investigations.

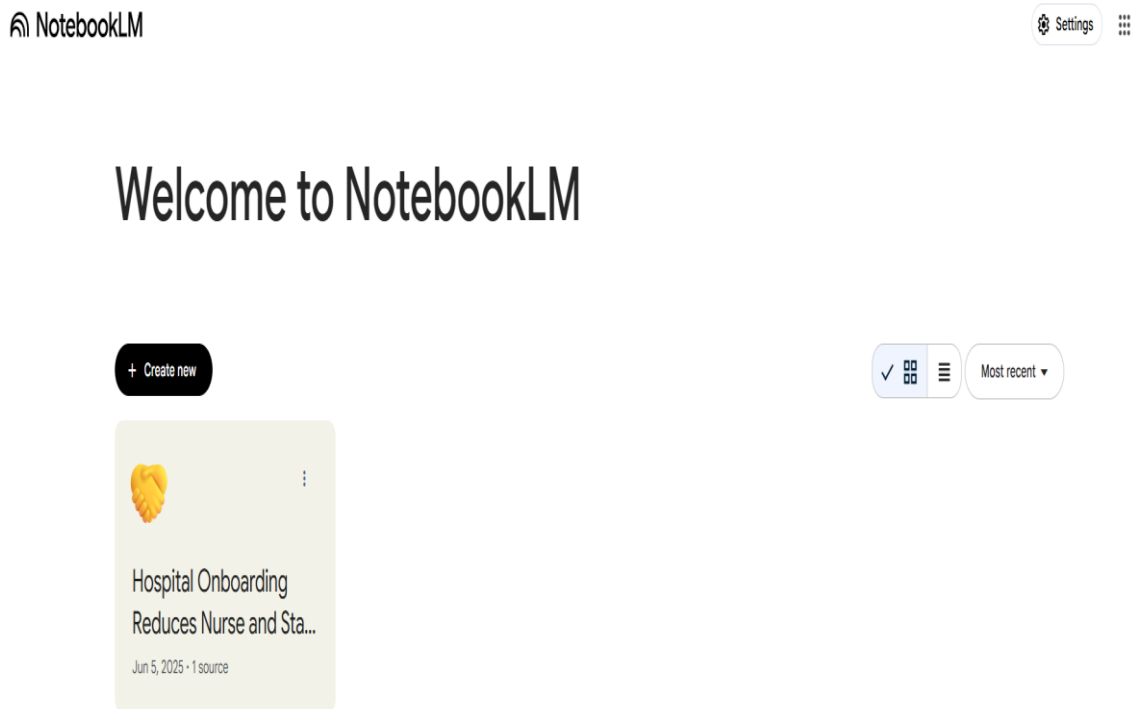


Figure 7: Saved Notebook

2.2 Lessons from Education

In light of the complexity of designing effective training programs tailored to accommodate possible differences in individual learning and information retention within individuals, I believe that large firms can leverage GAI to enhance learning processes based on the success of NotebookLM's prior applications. Literature has recently explored the tool's potential in the context of education and learning. As Ayala-Pazmiño (2023) proposes, "one of the main benefits of AI in education is its potential to personalize learning experiences for individual students." Whilst education can be described as an industry of sorts, although somewhat controversially, schools and universities are still a nuanced

organizational setting. Furthermore, to illustrate an application of GAI in the real-world, 'education' is a robust empirically driven example.

The COVID-19 pandemic greatly impacted the way education is delivered. Schools and universities were obliged to adapt to an online setting in order to comply with social distancing guidelines. (Baidoo-Anu and Owusu Ansah, 2023). Tools like NotebookLM led this innovation by "generating educational podcasts that discuss topics in a natural conversational style" (Pataranitaporn et al., 2021), allowing the production of multimedia materials that can engage students in innovative ways (Huffman and Hutson, 2024).

Artificial Intelligence has "revolutionized the ability of educators, especially those without extensive technical expertise, to produce a wide range of educational materials with minimal effort. LLMs like GPT-3 and other generative AI tools are now capable of generating lesson plans, rubrics, and assessments, all of which can be customized to specific curriculum needs." (Huffman and Hutson, 2024).

Undoubtedly, GAI "is already pushing educational boundaries and initiating a significant paradigm shift in existing educational praxis." (Baidoo-Anu and Owusu Ansah, 2023).

Huffman and Hutson (2024) at Lindenwood University (USA) explored the impact on students' knowledge acquisition of using NotebookLM's features for multimedia content creation. Drawing upon the transcribed diary of Mary Easton Sibley, the authors "effectively generated FAQs, a study guide, a table of contents, a briefing document, and an audio overview in podcast format", successfully "transforming static historical documents into dynamic learning materials", demonstrating that "the advancement of technology has ushered in an era where interacting with the deceased through digital means is no longer a futuristic concept". Ultimately, the experiment highlights the tool's ability to develop a variety of learning materials, tailored for diverse types of learners,

allowing educators to allocate more time enhancing student engagement through classroom active learning.

On the other hand, inherent limitations of the platform were discussed. Particularly, flawed timelines and fictionalized quotes which highlight the necessity of fact-checking materials before using them.

This perhaps suggests that, whilst the potential for GAI seems positive, there are still technological limitations that need to be addressed. Specifically, biases and privacy concerns (Ayala-Pazmiño, 2023). This, in turn, implies, at least in part, some ethical concerns.

Exploring Ethics: NotebookLM in Healthcare Organizations

1. Moral Compass: Unpacking Consequentialism, Deontology and Virtue Ethics

Exploring the impact of implementing NotebookLM in large organizations' learning processes through an ethical lens requires understanding the key viewpoints of consequentialism, deontology, and virtue ethics. Whilst a diverse range of ethical paradigms exist, the majority of these can be collapsed into three main schools. Table 1 (Thorpe, 2014) presents a comparative analysis of the major ethical frameworks.

	Consequentialism	Deontology	Virtue ethics
Sub-schools	Utilitarianism	Kantianism	Aristotle
Examples of key concept	The greatest good for the greatest number (Bentham)	Categorical imperative	The disclosure rule The golden rule
Orientation of analysis	Exogenous to the person or people: the consequence of the act	Exogenous to the person or people: the means to the outcome	Endogenous to the person or people: engrained values
Approach to what is good	Underpinning principle: to achieve maximum happiness	The 'right' action is also what is good	The consequences of the actions of virtuous people
Approach to what is right	The 'right' actions are those which will achieve maximum happiness	Underpinning principle: to engage in ethical action, i.e., to engage moral duties	The action of virtuous people
Approach to what is virtuous	A combination of virtues that enable maximum happiness to be sought	Traits that encourage the carrying out of ethical action	Underpinning principle: internalized value systems

Table 1: Principles of Consequentialism, Deontology and Virtue Ethics (Thorpe, 2014)

Consequentialism, particularly Utilitarianism, assesses the morality of a given action solely based on the anticipated outcomes, regardless of the process in which they are achieved. (Bentham, 1789/1961; Gandz and Hayes, 1988; Kujala and Pietilainen, 2004; Mill, 1861; Sidgwick, 1907; Kagen, 1998). Thus, following

“the greatest good for the greatest number” as a guiding principle, introduced by Bentham’ emphasizing the consequences rather than the intrinsic nature of the act. The action capable of maximizing overall happiness would be, therefore, preferred and considered “right”. Virtue, then, is the combination of attributes that facilitate the pursuit of maximum happiness.

Conversely, deontology assumes ethical rules must be followed throughout every course of action- hence, the outcome is deemed irrelevant (Kamm, 1996; Kant, 1785/2005; Koehn, 1995; Rawwas et al., 2005; Ross, 1930). In this regard, Kant introduced the “Categorical Imperative”, establishing that morality must be universally applicable and performed out of moral obligation. In essence, “Good” and “Right” are indistinguishable independently of the achieved outcome, and “Virtue” designates the traits that encourage adherence to ethical action. Thus, both consequentialism and deontology presume ethics to be external to the individual and propose universal moral guidelines (Waller, 2005).

Virtue Ethics, ingrained in Aristotelian philosophy, underlines the intrinsic values of the individual. Morality is measured by the innate and subjective values that drive actions, disregarding the existence of fixed external rules (Aristotle, as discussed in Broadie, 1991; Hume, 1739/1949; Koehn, 1995, 1998; Murphy, 1999; Stocker, 1976; Whetstone, 2001). In this perspective, “good” is defined by the outcome of actions performed by virtuous individuals, whereas “virtue” is shaped by the internalized value systems that cultivate moral character.

The comprehensive approaches bring complexity to the ethical discussion we are enduring. In this project, I intend to assess the employer’s viewpoint and key considerations in the decision-making regarding the implementation of NotebookLM in large organizations’ learning processes.

2. Ethical Crossroads: NotebookLM in Healthcare

Koski and Murphy (2021) reflect on the mediocre reality of healthcare systems, highlighting the elevated costs and stating that they don't work properly, nor do they serve the population equally, and assume that, although, "AI cannot address all the societal, political and environmental issues" it "(...) can contribute to increasing efficiency, raising standards of care, delivering on the promise of precision medicine and supporting research". Thus, intending to explore the ethical issues that might arise with the application of NotebookLM in large organizations' learning processes, this thesis will focus its analysis on healthcare as one example.

As Koski and Murphy (2021) argue, "the earliest applications of expert systems in healthcare were designed to mimic human reasoning processes, but they were necessarily bound by the limits of established medical knowledge". Recently, Artificial Intelligence has demonstrated a wide variety of applications in medicine (Rigby 2019), in areas such as image diagnostic support, radiology education, and medical physics" (Keshavarz et al., 2024 and Kadoya et al., 2024).

Particularly for Lung Cancer Staging, Tozuca et al. (2024) revealed that NotebookLM outperformed GPT-4o by achieving an 86% versus 39% diagnosis accuracy using REK, and 95% versus 25% without it. The researchers pointed out several limitations for the study, the results were solely based on 100 fictitious objects (rather than clinical data), the sources required translation, and clinical practice would likely pose data privacy concerns. Anyhow, the findings portray a great possibility for further exploration.

Ultimately, LLMs "have been shown to be capable of diagnosing cancer, Alzheimer's disease, and even selecting treatment options" (Gaonkar et al., 2020), posing a promising future for healthcare (Yelne et al., 2023).

Healthcare could be considered a “pro-technology” industry, as it already tends to embrace technological advancements (including NotebookLM) in clinical settings, while also being particularly sensitive to ethical concerns with its use (Trocin et al., 2023). Thus, the possibility of these organizations adopting AI tools in their business functions, including human resources, seems quite likely.

2.1. The Ethical Dilemma of Applying NotebookLM for Training in Healthcare

The tendency for LLM hallucinations posed a great risk for clinical diagnosis. Entering the era of Retrieval-Augmented Generation (RAG) significantly mitigated that risk due to the possibility of feeding tools with Reliable External Knowledge (REK). Nonetheless, data breaches, security concerns, technology-induced errors capable of impacting patient care, provider burnout due to complex systems, overreliance on technology (Gong et al., 2024) and perpetuation of biases posed by the generalization of clinical trials fed into these supervised models can still occur when implementing AI tools in healthcare operations, as well as in human resources. Undoubtedly, “AI can assist in developing personalized learning experiences, allowing healthcare providers to enhance their skills and knowledge more efficiently and effectively” (Gong et al., 2024), even facilitating a better patient experience, improve population health, reduce the costs per capita and ultimately improving the overall satisfaction of health professionals and consequently, increasing employee retention.

Employee retention is a key indicator of an organization’s stability and financial well-being (Griffeth and Hom, 2001). Particularly in intense acute care practice environments, a structured orientation process is essential to prevent employee discouragement and disconnection from peers (Goldschmidt et al., 2011). In a Washington DC healthcare facility, for example, Kurnat-Thoma et al.

(2017) “identified unusually high new-hire turnover”, uncovering a “need for a new onboarding program to better support (...) new-hire employees”.

From a consequentialist perspective, the ability of NotebookLM to develop various training materials, tailored to the different individual’s learning preferences and the impact it could potentially have on both motivation and productivity is one of the strongest arguments in favor of utilizing the tool. Tailoring educational content to each employee’s learning style encompasses both knowledge retention and confidence. Thus, increasing productivity and willingness to accommodate and meet stakeholders’ demands, all while reducing costs of operations (associated with training and turnover) and possibly increasing revenue. In this way, overall happiness is maximized, i.e., as it is manifested in all parts of the process.

From a deontological standpoint, however, patients’ privacy concerns must be addressed above all. As a supervised AI system, the quality of the materials generated is directly linked to the information being fed into the software, which means that a great amount of private data must be used to ensure efficiency. If a proper data management framework is not implemented, confidentiality cannot be guaranteed. Thus, legally, and ethically, security failures may lead to severe consequences. Although an anonymization process could hypothetically solve this concern, the reality is that re-identification through AI is possible (and likely). Reinforcing the role of rules and duties within the deontological lens, the assumption that patient confidentiality has to be upheld as a moral obligation is clear. However, obtaining the patients’ consent before anonymizing their particular case studies and utilizing them to develop medical training in NotebookLM could be an ethical workaround to ensure the ethical responsibility deontology requires.

On the other hand, the primary intention behind the implementation of NotebookLM in healthcare organizations’ learning processes must be addressed

from an Aristotelian viewpoint. If virtuous individuals perform virtuous actions, resorting to NotebookLM exclusively to reduce operational costs, disregarding employee satisfaction or privacy concerns, would certainly not be considered ethical.

Additionally, deontology would suggest feeding this software with clinical data from various patient demographics, aiming to mitigate any biases the platforms could generate that would later influence the courses of action of healthcare professionals. Unique factors of specific groups, whether religious, ethnic, gender or age, can influence the required care in certain situations. Not taking these into account would most likely lead to malpractice. Hence, ensuring training materials represent a wide range of patient profiles is essential to deliver equal, fair, and reliable care. Besides, given its intrinsic preoccupation with moral obligations, deontology corroborates virtue ethics' need for recurrent fact-checking/proofreading of generated materials used throughout the learning processes and, more importantly, patient consent beforehand. Hypothetically, creating an ethical oversight committee within institutions that chooses to integrate NotebookLM appears to be a possible solution to tackle this particular concern.

Moreover, underscoring the importance of human judgment is crucial if we are considering a consequentialist perspective. Training agents lacking critical thinking through overreliance on AI may struggle in complex real-world settings. And although NotebookLM could act as a roadmap, it cannot substitute human reasoning, especially when human lives are at stake. Including an instruction guide on the ways of working with the platform, or regularly assessing employees' knowledge on the matter, could be a viable solution.

Balancing empathy and expertise is vital for treatment outcomes. From a virtue ethics standpoint, fostering compassion through human connection is imperative; as for consequentialism, stakeholders' efficient gains are also to be

considered; deontology incentivizes careful data management and ethical data collection. Ultimately, simultaneously accommodating the needs of each lens would point us in the direction of creating a hybrid training model. Using NotebookLM as a supplementary tool rather than an exclusive source could reduce training costs, empower knowledge retention through tailored programs, reduce turnover, and enhance care.

Discussion

Leveraging NotebookLM to Reduce Turnover: Is It Ethical?

In this paper, it is discussed how standardized learning programs impact motivation while also suppressing curiosity, both in classrooms and in organizational settings. Therefore, suggesting tailoring training materials and programs upon assessing stakeholders' preferred learning style could positively increase knowledge retention and, consequently, lower turnover. Due to the high cost of turnover discussed prior, my hypothesis is that employee retention preserves resources that can be further redirected towards enhancing operations. Particularly for healthcare organizations, these 'preserved' financial resources can be allocated to acquire better equipment and, eventually, improve patient care, ultimately addressing Koski and Murphy's (2011) aforementioned concerns.

Literature has been investigating the benefits and risks associated with leveraging GAI as a means to optimize operations within firms. Specifically, when used for knowledge enhancement, Ayala-Pazmiño (2023) reflected on the dichotomy between the possible positive outcomes for personalized learning and the concerns raised, e.g., regarding biases and the dehumanization of the learning experience. Eager, B. and Brunton, R. (2023) argue that, ideally, AI-generated content should be gradually introduced. The authors propose companies to start utilizing it in basic tasks and progress to more complex ones, transitioning smoothly to optimized operations and enabling feedback and adjustments throughout the process.

Building on this, I propose that organizations use NotebookLM to tailor training programs to diverse learning styles as it demonstrates strong

adaptability to cognitive mechanisms among currently available tools. However, in doing so, a tension is instigated between operational needs and ethical considerations. On one hand, AI, specifically NotebookLM, is clearly attractive for organizations in general, and healthcare providers in particular, as a means to leveraging a competitive advantage. On the other hand, such means invoke ethical dimensions that suggest such organizations must proceed with caution.

I have demonstrated that considering the ethics of organizational actions is somewhat problematic due to the different perspectives that are to analyze it.

For instance, from a deontological standpoint, patient privacy and ethical data handling are of utmost importance. Thus, using private data to build tailored training materials requires an implementation process guided by strict moral obligations. Gong et al. (2024) support, in this context, that “ensuring patient safety in healthcare involves training professionals and implementing clinical decision support systems (CDSS) and health IT solutions”. Consequentialism, on the other hand, highlights the platform’s potential for productivity, motivation, and patient outcome improvement, leading to a maximization of overall happiness (Goldschmidt et al., 2011, Gong et al., 2024, Kurnat-Thoma et al., 2017). Thus, whilst deontology and utilitarianism may appear opposite to each other, they are united by the assumption that objective rules should be used to determine what is ethical, and what is not.

Due to the nature of their operations, organizations vary widely in the way in which they apply either NotebookLM, or other forms of AI into the business. That means that healthcare will inevitably be different from education, despite what we’ve learned from the latter. Regarding the ‘socialization’ factor, for example, Ayala-Pazmiño (2023) postulates that recognizing human connection is vital in the process of learning. And although this should be accounted for in education, lack of ‘socialization’ is not a threat in healthcare facilities.

Considering the ethical lens explored, I would suggest organizations within healthcare approach the implementation of NotebookLM and AI from a virtue ethics position, as it primarily focuses on intention (Aristotle, as discussed in Broadie, 1991; Hume, 1739/1949; Koehn, 1995, 1998; Murphy, 1999; Stocker, 1976; Whetstone, 2001). Suggesting, then, that the tool – particularly the data used to train the models – is handled with integrity and transparency, e.g., through patient anonymization and consent. Most importantly, seeking the patients’ well-being and placing the decision to each individual organization – or, more realistically – the individual manager or key decision maker.

This is relevant since the ethical standpoint of organizations is defined by “the set of principles and standards that guide” the firm’s decisions, values, and stakeholders’ behavior. (Saéz, 2025). More explicitly, “the company’s internal conduct and daily practices,” which includes “compliance with laws and regulations, honesty in communication, and fairness in dealings with stakeholders” (Saéz, 2025). Therefore, the framework that organizations choose to address the risks that arise from NotebookLM’s implementation reveals their ethical standpoint, whether they choose to approach the dilemma by focusing on the outcomes, duties, or universal values.

To aid this process, Table 2 describes (some) reflective questions for organizations in general, and healthcare in particular, to use as a starting point. It presents a multi-dimensional analysis of the described ‘tension’ by outlining the ethical challenges organizations may encounter when implementing NotebookLM. The presented framework integrates actionable courses of action and assigns the responsibilities to the different teams/departments within firms – in summary, categorizes ethical concerns in three domains: Privacy & Security, NotebookLM Dependency and Biases & Errors. The Privacy & Security domain postulates a framework for organizations and facilities to prevent data breaches when feeding NotebookLM with sensitive information, emphasizing the

importance of consent, user and patient anonymization, and overall ethical data treatment. NotebookLM Dependency focuses on ensuring the tool support rather than replacing human judgement. Thus, promoting hybrid training models and material oversight. The Biases & Errors domain guides organizations on how to manage wrongful outputs, whether hallucinations or outdated guidelines, reinforcing the necessity of maintaining reliable data sources.

	Practical Implications (for General Organizations)	Practical Implications (for Healthcare)	Framework	Courses of Action	Team/Department
Privacy & Security (Preventing Data Breaches when feeding NotebookLM with Sensitive Information)	How should organizations deal with user data in training materials?	How should facilities deal with real patient information in training materials?	Anonymization of patient/user data per given consent and data stewardship.	<ul style="list-style-type: none"> Remove personally identifiable information from model training data. Only use data after obtaining explicit (written) consent. Nominate a data stewardship manager to handle privacy practices. 	<ul style="list-style-type: none"> Legal IT
	How does user consent work?	How does patient consent work?	Informed consent regarding risks (e.g. possible re-identification through AI) associated with the use of confidential information for training purposes.	<ul style="list-style-type: none"> Partner HR & Legal to draft transparent consent forms. Educate users on how data might be used. Provide clear opt-out mechanisms. 	<ul style="list-style-type: none"> HR Legal Compliance
	Should users be aware when NotebookLM is being used for training purposes?	Should patients be aware when NotebookLM is being used to train their caretakers?	Develop an experiment where the experimental group of patients receives care from AI trained practitioners and the control group is treated from 'normally' trained employees (upon prior consent). Report on treatment efficiency and provide the results to patients.	<ul style="list-style-type: none"> Design a pilot comparing AI-assisted care vs standard care. Collect consent from participants beforehand. Share insights & results with stakeholders to foster trust. 	<ul style="list-style-type: none"> Clinical Research Ethics Board
NotebookLM Dependency (Assuring the tool supports rather than replacing human judgement)	How can organizations track NotebookLM use and assess employee practices?	How can facilities track NotebookLM use and assess employee practices?	Include an instruction guide on the ways of working with NotebookLM and assess the stakeholders' knowledge on the matter on a given cadence. Foster a hybrid training model - i.e., using NotebookLM as a supplementary tool rather than a substitute for human reasoning. This can reduce costs, empower knowledge retention, reduce turnover and enhance care.	<ul style="list-style-type: none"> Compose and share a guide for responsible use of NotebookLM. Offer periodic workshops and training refreshers. Monitor employee reliance and reinforce human judgment. 	<ul style="list-style-type: none"> Training and Development HR HoD across business
Biases & Errors (Dealing with Wrongful Outputs)	How can organizations handle biases posed by LLM hallucinations?	How can facilities handle biases posed by LLM hallucinations?	Create a feedback mechanism for practitioners to flag wrong/biased outputs early and report.	<ul style="list-style-type: none"> Create a simple feedback system for AI error reporting. Encourage employees to flag odd results immediately. Assign a team to review and address the feedback regularly. 	<ul style="list-style-type: none"> Operations IT Risk Management
			Feed the software with clinical data from various patients' demographics to deliver equal, fair and reliable care.	<ul style="list-style-type: none"> Curate diverse clinical datasets for AI training. Test outputs to ensure consistency across the different patient types. Collaborate with clinical advisors. 	<ul style="list-style-type: none"> IT Data Science DEI
	How can organizations ensure the output follows guidelines and internal policies?	How can facilities ensure the output respects protocol and guidelines are up to date?	The quality of the output is directly related to the information fed into the system. Thus, proofreading the sources on a given decance is of utmost importance.	<ul style="list-style-type: none"> Create an Ethical oversight committee to help integrate NotebookLM. Align training materials with latest policies and protocols. Require critical outputs to be reviewed by a human expert. 	<ul style="list-style-type: none"> Compliance Audit Legal

Legend:

 Specific for *Consequentialism*
 HoD Head of Department
 DEI Diversity, Equity and Inclusion

Table 2: Practical Implications of Implementing NotebookLM for Training Purposes

Next Steps

This project primarily contributes to the critical assessment of NotebookLM's features and capabilities to enhance learning processes within organizations, particularly focusing on its potential benefits to reduce employee turnover. While the ethical discussion emphasized the tool's capacity in the healthcare sector, the discussion brought up the possibility of extending the findings to other organizational contexts, which culminated in Table 2.

Recognizing the potential of NotebookLM across sectors is crucial, as dysfunctional employee turnover represents a significant cost for firms, regardless of the sector. This applies especially to turnover stemming from employees' lack of motivation due to misadjusted training programs, which can be addressed, at least in part, by tailoring organizational learning processes to each stakeholder's learning preferences. NotebookLM has proven to be a versatile and, thus, promising tool for learning engagement, as it facilitates a personalized learning experience, further leading to retention rate increase.

Building on the theoretical insights posed by the present conceptual piece, future research should empirically assess NotebookLM's effectiveness for enhancing learning processes, whether within healthcare or other organizational contexts, and its impact on dysfunctional turnover rates. More precisely, conduct an A/B test capturing employee engagement, knowledge retention, overall satisfaction, and turnover before and after integrating NotebookLM into a firm's operations to derive measurable insights. Additionally, there is a necessity to simultaneously educate stakeholders on ethical approaches to AI leveraging to ensure acceptance and foster trust.

Ultimately, this thesis aims to build a concrete ethical foundation and an action roadmap for companies intending to consider NotebookLM as a strategic resource for learning improvement. However, it is important to recognize the

significance of empirical testing and validation for flagging limitations and measuring effectiveness.

Final Thoughts

In essence, I explored how human learning is processed (both at the neural and cognitive level). Grounded in this knowledge, as well as Huffman and Hutson' (2024) findings on Lindenwood University, I theorized on how healthcare organizations, which are generally pro-technology, could implement NotebookLM to improve training processes.

Throughout this research, I recognized the importance of assessing the concerns derived from this implementation through the lenses of deontology, consequentialism, and virtue ethics. Concluding that organizations themselves should be responsible for deciding the framework that better aligns with the reputation and values they intend to project. Nevertheless, and intending to provide a broader framework roadmap, easily applicable to the different schools, I developed Table 2.

Ultimately, I suggest companies focus mostly on intention, and implement a hybrid model that balances human reasoning and AI capabilities.

Bibliography

Atak, M., & Atik, İ. (2007). The importance of continuing education and its effects on the process of creating a learning organization—Air Force Command, the sample of a continuing education model. *Journal of Aviation and Space Technologies*, 3(1), 63–70.

Alec Radford, K., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). *Improving language understanding by generative pre-training*. OpenAI.

Alsheibani, S., Cheung, Y., Messom, C., & Alhosni, M. (2020). Winning AI strategy: Six steps to create value from artificial intelligence. *Americas Conference on Information Systems*, Online.

Ayala-Pazmiño, M. (2023). Artificial intelligence in education: Exploring the potential benefits and risks. *593 Digital Publisher CEIT*, 8(3), 892–899. <https://doi.org/10.33386/593dp.2023.3.1827>

Baidoo-Anu, D., & Owusu Ansah, L. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1), 52–62.

Bentham, J. (1789/1961). *An introduction to the principles of morals and legislation*. Garden City, NY: Doubleday.

Bernard, J., Popescu, E., & Graf, S. (2022). Improving online education through automatic learning style identification using a multi-step architecture with ant colony system and artificial neural networks. *Applied Soft Computing*, 131. <https://doi.org/10.1016/j.asoc.2022.109779>

Broadie, S. (1991). *Ethics and Aristotle*. New York, NY: Oxford University Press.

Cascio, F. W. (1991). *Managing human resource: Productivity, quality of work life and profits* (3rd ed.). New York, NY: McGraw-Hill.

Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(43). <https://doi.org/10.1186/s41239-023-00411-8>

Defined Learning. (n.d.). What are VARK learning styles. *Defined Learning Blog*. Retrieved January 28, 2025, from <https://blog.definedlearning.com/blog/what-are-vark-learning-styles>

Eager, B., & Brunton, R. (2023). Prompting Higher Education Towards AI-Augmented Teaching and Learning Practice. *Journal of University Teaching & Learning Practice*, 20(5). <https://doi.org/10.53761/1.20.5.02>

Enholm, I. M., Papagiannidis, E., Mikalef, P., et al. (2022). Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, 24, 1709–1734. <https://doi.org/10.1007/s10796-021-10186-w>

Esade. (2025, January). Business ethics, why they are so important? *Esade Beyond*. <https://www.esade.edu/beyond/en/business-ethics>

Fitkov-Norris, E. D., & Yeghiazarian, A. (2015). Validation of VARK learning modalities questionnaire using Rasch analysis. *Journal of Physics: Conference Series*, 588, 012048.

Gandz, J., & Hayes, N. (1988). Teaching business ethics. *Journal of Business Ethics*, 7(9), 657–669.

Gaonkar, B., Cook, K., & Macyszyn, L. (2020). Ethical issues arising due to bias in training A.I. algorithms in healthcare and data sharing as a potential solution. *AI Ethics Journal*, 1(1). <https://doi.org/10.47289/AIEJ20200925>

Goldschmidt, K., Rust, D., Torowicz, D., & Kolb, S. (2011). Onboarding advanced practice nurses: Development of an orientation program in a cardiac center. *Journal of Nursing Administration*, 41(1), 36–40.

Gong, Y., Min, H., Jing, X., & Yu, P. (2024). Challenges and opportunities of artificial intelligence in CDSS and patient safety. In J. Mantas et al. (Eds.), *Digital health and informatics innovations for sustainable health care systems* (pp. 1250–1255). IOS Press. <https://doi.org/10.3233/SHTI240638>

Google. (n.d.). How to use NotebookLM with your Google Account. *Google Support*. Retrieved November 23, 2024, from <https://support.google.com/notebooklm/answer/15724458>

Griffeth, R., & Hom, P. (2001). Chapter 1: How turnover can be costly to organizations. In *Retaining valued employees* (pp. 1–30). Thousand Oaks, CA: Sage Publications.

Hu, L. (2023). Generative AI and future. Retrieved January 23, 2025, from <https://pub.towardsai.net/generative-ai-and-future-c3b1695876f2>

Hubert, J., Matsuda, E., & Ikegami, T. (2013). Hebbian learning in a multimodal environment. In *Proceedings of the 12th European Conference on the Synthesis and Simulation of Living Systems: Advances in Artificial Life, ECAL 2013* (pp. 698–705). <https://doi.org/10.7551/978-0-262-31709-2-ch100>

Hume, D. (1739/1949). *A treatise of human nature: Book 3*. London: J. M. Dent & Sons.

Jovanović, M. (2023). Generative artificial intelligence: Trends and prospects. <https://www.computer.org/csdl/magazine/co/2022/10/09903869/1H0G6xtREk>

Kagen, S. (1998). *Normative ethics*. Boulder, CO: Westview.

- Kamm, E. M. (1996). *Morality, mortality: Vol. II: Rights, duties, and status*. New York, NY: Oxford University Press.
- Kant, I. (1785/2005). *Groundwork of the metaphysics of morals*. London: Routledge.
- Koehn, D. (1995). A role for virtue ethics in the analysis of business practice. *Business Ethics Quarterly*, 5, 533–539.
- Koehn, D. (1998). Virtue ethics, the firm, and moral psychology. *Business Ethics Quarterly*, 8, 497–513.
- Koski, E., & Murphy, J. (2021). AI in healthcare. In M. Honey, C. Westbrooke, & D. Foster (Eds.), *Nurses and midwives in the digital age* (pp. 295–298). IOS Press. <https://doi.org/10.3233/SHTI210726>
- Kujala, J., & Pietiläinen, T. (2004). Female managers' ethical decision-making: A multi-dimensional approach. *Journal of Business Ethics*, 53, 153–163.
- Kurnat-Thoma, E., Ganger, M., Peterson, K., & Channell, L. (2017). Reducing annual hospital and registered nurse staff turnover — A 10-element onboarding program intervention. *SAGE Open Nursing*, 3. <https://doi.org/10.1177/2377960817697712>
- Latif, S., Chen, S., & Beck, F. (2021, June). A deeper understanding of visualization-text interplay in geographic data-driven stories. *Computer Graphics Forum*, 40(3), 311–322.
- Marcy, V. (2001). Adult learning styles: How the VARK learning style inventory can be used to improve student learning. *Journal of Physician Assistant Education*, 12(2), 1–5.
- Mill, J. S. (1861/1998). *Utilitarianism*. Oxford, UK: Oxford University Press.

Murphy, P. E. (1999). Character and virtue ethics in international marketing: An agenda for managers, researchers and educators. *Journal of Business Ethics*, 18, 107–124.

Nayebi, A., Srivastava, S., Ganguli, S., & Yamins, D. L. K. (2020). Identifying learning rules from neural network observables. *arXiv*. <http://arxiv.org/abs/2010.11765>

Nguyen, B., Sonnenfeld, N., Finkelstein, L., Alonso, A., Gomez, C., Duruaku, F., & Jentsch, F. (2023). Using AI tools to develop training materials for aviation: Ethical, technical, and practical concerns. *II: Innovations in Training and Simulation, Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 67(1), 1343–1349. <https://doi.org/10.1177/221695067231192904>

Noriyuki Kadoya, Arai, K., Tanaka, S., Kimura, Y., Tozuka, R., Yasui, K., Hayashi, N., Katsuta, Y., Takahashi, H., Inoue, K., & Jingu, K. (2024). Assessing knowledge about medical physics in language-generative AI with large language model: Using the medical physicist exam. *Radiological Physics and Technology*. Advance online publication.

N. Y., Mulongo. (2024). Navigating the interplay between artificial intelligence and managerial skills. In *2024 International Symposium on Networks, Computers and Communications (ISNCC)* (pp. 1–4). <https://doi.org/10.1109/ISNCC62547.2024.10759025>

OpenAI. (n.d.). ChatGPT capabilities overview. <https://help.openai.com/en/articles/9260256-chatgpt-capabilities-overview>

O'reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrlich, J. (2021). Deep predictive learning in neocortex and pulvinar. *Journal of Cognitive Neuroscience*, 33(6), 1158–1196. https://doi.org/10.1162/jocn_a_01708

Pataranutaporn, P., Danry, V., Leong, J., Punpongsanon, P., Novy, D., Maes, P., & Sra, M. (2021). AI-generated characters for supporting personalized learning and well-being. *Nature Machine Intelligence*, 3(12), 1013–1022.

Pedram Keshavarz, Sara Bagherieh, Seyed Ali Nabipoorashrafi, Hamid Chalian, Amir Ali Rahsepar, Grace Hyun J. Kim, Cameron Hassani, Steven S. Raman, & Arash Bedayat. (2024). ChatGPT in radiology: A systematic review of performance, pitfalls, and future perspectives. *Diagnostic and Interventional Imaging*, 105(7–8), 251–265.

Perplexity AI. (n.d.). Meet AI Bard Gemini: The future. <https://www.perplexity.ai/page/meet-ai-bard-gemini-the-future-GnWnC52XTASgh0ds18RotA>

Pumplun, L., Tauchert, C., & Heidt, M. (2019). A new organizational chassis for artificial intelligence—Exploring organizational readiness factors. In *Proceedings of the 27th European Conference on Information Systems (ECIS)*, Stockholm & Uppsala, Sweden.

Rawwas, M. Y. A., Swaiden, Z., & Oyman, M. (2005). Consumer ethics: A cross-cultural study of the ethical beliefs of Turkish and American consumers. *Journal of Business Ethics*, 57, 183–195.

Rigby, M. J. (2019). Ethical dimensions of using artificial intelligence in health care. *AMA Journal of Ethics*, 21(2), 121–124.

Ross, W. D. (1930). *The right and the good*. Oxford: Clarendon Press.

Sidgwick, H. (1907). *The methods of ethics*. London: Macmillan.

Sintia, I., Rusnayati, H., & Samsudin, A. (2019). VARK learning style and cooperative learning implementation on impulse and momentum. *Journal of*

Physics: Conference Series, 1280(5). <https://doi.org/10.1088/1742-6596/1280/5/052032>

Sivasubramanian, K., Mukthar, K. P. J., Allauca-Castillo, W., Asis-Lopez, M., Cayotopa-Ylatoma, C., & Mory-Guarnizo, S. (2023). Application of artificial intelligence in human resource activities for advanced governance and welfare. In Shakya, S., Balas, V. E., & Haoxiang, W. (Eds.), *Proceedings of Third International Conference on Sustainable Expert Systems* (Vol. 587, pp. 759–770). Springer. https://doi.org/10.1007/978-981-19-7874-6_66

Stable Diffusion Online. (n.d.). *Stable Diffusion Online*. <https://stablediffusionweb.com/>

Stocker, M. (1976). The schizophrenia of modern ethical theories. *The Journal of Philosophy*, 73, 453–466.

Thorpe, A. S., & Roper, S. (2019). The ethics of gamification in a marketing context. *Journal of Business Ethics*, 155(2), 597–609. <https://doi.org/10.1007/s10551-017-3501-y>

Trocin, C., Mikalef, P., Papamitsiou, Z., & Conboy, K. (2023). Responsible AI for digital health: A synthesis and a research agenda. *Information Systems Frontiers*, 25(6), 2139–2157.

Tziner, A., & Birati, A. (1996). Assessing employee turnover costs: A revised approach. *Human Resource Management Review*, 6(2), 113–122.

Waller, B. N. (2005). *Consider ethics: Theory, readings, and contemporary issues*. New York, NY: Pearson Longman.

Waldman, J. D., Kelly, F., Arora, S., & Smith, H. L. (2004). The shocking cost of turnover in health care. *Health Care Management Review*, 29(1), 2–7.

Whetstone, J. T. (2001). How virtue ethics fits within business ethics. *Journal of Business Ethics*, 33, 101–114.

Yelne, S., Chaudhary, M., Dod, K., Sayyad, A., & Sharma, R. (2023). Harnessing the power of AI: A comprehensive review of its impact and challenges in nursing science and healthcare. *Cureus*. <https://doi.org/10.7759/cureus.49252>

Appendices

1. VARK Questionnaire



The VARK Questionnaire (Version 8.01) How Do I Learn Best?

Choose the answer which best explains your preference and circle the letter(s) next to it. Please circle **more than one** if a single answer does not match your perception. Leave blank any question that does not apply.

- I need to find the way to a shop that a friend has recommended. I would:
 - find out where the shop is in relation to somewhere I know.
 - ask my friend to tell me the directions.
 - write down the street directions I need to remember.
 - use a map.
- A website has a video showing how to make a special graph or chart. There is a person speaking, some lists and words describing what to do and some diagrams. I would learn most from:
 - seeing the diagrams.
 - listening.
 - reading the words.
 - watching the actions.
- I want to find out more about a tour that I am going on. I would:
 - look at details about the highlights and activities on the tour.
 - use a map and see where the places are.
 - read about the tour on the itinerary.
 - talk with the person who planned the tour or others who are going on the tour.
- When choosing a career or area of study, these are important for me:
 - Applying my knowledge in real situations.
 - Communicating with others through discussion.
 - Working with designs, maps or charts.
 - Using words well in written communications.
- When I am learning I:
 - like to talk things through.
 - see patterns in things.
 - use examples and applications.
 - read books, articles and handouts.
- I want to save more money and to decide between a range of options. I would:
 - consider examples of each option using my financial information.
 - read a print brochure that describes the options in detail.
 - use graphs showing different options for different time periods.
 - talk with an expert about the options.
- I want to learn how to play a new board game or card game. I would:
 - watch others play the game before joining in.
 - listen to somebody explaining it and ask questions.
 - use the diagrams that explain the various stages, moves and strategies in the game.
 - read the instructions.

©2019 VARK-Learn Limited

<https://vark-learn.com>

- I have a problem with my heart. I would prefer that the doctor:
 - gave me something to read to explain what was wrong.
 - used a plastic model to show me what was wrong.
 - described what was wrong.
 - showed me a diagram of what was wrong.
- I want to learn to do something new on a computer. I would:
 - read the written instructions that came with the program.
 - talk with people who know about the program.
 - start using it and learn by trial and error.
 - follow the diagrams in a book.
- When learning from the Internet I like:
 - videos showing how to do or make things.
 - interesting design and visual features.
 - interesting written descriptions, lists and explanations.
 - audio channels where I can listen to podcasts or interviews.
- I want to learn about a new project. I would ask for:
 - diagrams to show the project stages with charts of benefits and costs.
 - a written report describing the main features of the project.
 - an opportunity to discuss the project.
 - examples where the project has been used successfully.
- I want to learn how to take better photos. I would:
 - ask questions and talk about the camera and its features.
 - use the written instructions about what to do.
 - use diagrams showing the camera and what each part does.
 - use examples of good and poor photos showing how to improve them.
- I prefer a presenter or a teacher who uses:
 - demonstrations, models or practical sessions.
 - question and answer, talk, group discussion, or guest speakers.
 - handouts, books, or readings.
 - diagrams, charts, maps or graphs.
- I have finished a competition or test and I would like to have feedback. I would like to have feedback:
 - using examples from what I have done.
 - using a written description of my results.
 - from somebody who talks it through with me.
 - using graphs showing what I achieved.
- I want to find out about a house or an apartment. Before visiting it I would want:
 - to view a video of the property.
 - a discussion with the owner.
 - a printed description of the rooms and features.
 - a plan showing the rooms and a map of the area.
- I want to assemble a wooden table that came in parts (kitset). I would learn best from:
 - diagrams showing each stage of the assembly.
 - advice from someone who has done it before.
 - written instructions that came with the parts for the table.
 - watching a video of a person assembling a similar table.

©2019 VARK-Learn Limited

<https://vark-learn.com>



The VARK Questionnaire – Scoring Chart

Use the following scoring chart to find the VARK category that each of your answers corresponds to. Circle the letters that correspond to your answers.

e.g. If you answered b and c for question 3, circle V and R in the question 3 row:

Question	a category	b category	c category	d category
3	K	V	R	A

Scoring Chart

Question	a category	b category	c category	d category
1	K	A	R	V
2	V	A	R	K
3	K	V	R	A
4	K	A	V	R
5	A	V	K	R
6	K	R	V	A
7	K	A	V	R
8	R	K	A	V
9	R	A	K	V
10	K	V	R	A
11	V	R	A	K
12	A	R	V	K
13	K	A	R	V
14	K	R	A	V
15	K	A	R	V
16	V	A	R	K

Calculating Your Scores

Count the number of each of the VARK letters you have circled to get your score for each category:

Total number of V s circled =	<input type="text"/>
Total number of A s circled =	<input type="text"/>
Total number of R s circled =	<input type="text"/>
Total number of K s circled =	<input type="text"/>

Fill in the questionnaire online at <https://vark-learn.com/the-vark-questionnaire/> to find out your VARK learning preference.

This document must not be published on the Internet – instead please suggest people download the latest version from the vark-learn.com website.

©2019 VARK-Learn Limited

<https://vark-learn.com>

2. Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of my written thesis, “Exploring the Ethical Implications of Leveraging NotebookLM to Reduce Turnover: An Exploratory Conceptual Study of Learning Processes within Healthcare”, ChatGPT, NotebookLM, and Engineering Village AI were used for the following tasks: proofreading, idea restructuring, and source finding with the prompts used listed at the end of the document in the Prompts List section. After using these tools, I reviewed and edited the content as necessary, and I take full responsibility for the content of the work presented.

I also declare that I am aware of and respect the Artificial Intelligence Rules of Conduct of Católica Porto Business School.

Full name: Ana Catarina da Silva Vieira

Student number: 355423007

Date: July 9th, 2025

Signature:

A handwritten signature in black ink, consisting of a stylized, cursive letter 'A' with a horizontal line crossing it.

3. Prompt List

1. Proofread this section.
2. Give me synonyms for [insert expression].
3. How can I say this more professionally?
4. I want to say [idea 1], [idea 2], [...]. In order for the text to become clear to the reader, what is the best structure to follow?