

University of Sopron
Széchenyi István Doctoral School of Management and Organizational
Sciences

**AI AT WORK: PROMISES, PERILS, AND PARADOXES OF
ADOPTION**

Doctoral (PhD) Dissertation

Prepared by:
Levente Szabados

Supervisor(s):
Dr. habil. Ferenc Kiss
Dr. Andrea Kópházi

Sopron
2025

AI AT WORK: PROMISES, PERILS AND PARADOXES OF ADOPTION

Dissertation to obtain a PhD degree

Written by:

Levente Szabados

Prepared by the University of Sopron

István Széchenyi Economics and Management Doctoral School

within the framework of the Business Administration and Management Programme

Supervisor(s): Dr. habil. Ferenc Kiss

Dr. Andrea Kópházi

The supervisor(s) has recommended the evaluation of the dissertation be accepted: yes / no

supervisor(s) signature

The evaluation has been recommended for approval by the reviewers (yes/no):

1. judge: Dr. _____ yes / no _____
(signature)

2. judge: Dr. _____ yes / no _____
(signature)

Result of the public dissertation defence: _____ %

Sopron, 20____ year _____ month _____ day

Chairperson of the Judging Committee

Qualification of the PhD degree: _____

UDHC Chairperson

Contents

1	Abstract	1
2	Introduction and hypotheses	2
2.1	Analytical Premise	2
2.2	Research Questions	3
2.3	Methodological Innovations	3
2.4	Key Hypotheses	4
2.5	Structure of the Dissertation	4
2.6	Historical context	4
3	Literature review	6
3.1	Topic and motivation	6
3.2	Background	6
3.3	Methodology	6
3.3.1	Literature Search and Collection	7
3.3.2	Screening and Selection	7
3.3.3	Thematic Analysis	7
3.3.4	Citation Analysis	7
3.3.5	Quantitative and Content Analysis	7
3.3.6	Identification of Research limitations	8
3.4	General trends	8
3.4.1	Why the trend?	9
3.5	Detailed analysis: The "Curated list"	12
3.5.1	General topics present in the "Curated set"	12
3.5.2	What is the consensus about the effects on unemployment specifically?	15
3.6	Establishing a Foundational Subset through Citation Graph Analysis	19
3.6.1	Assessing the Network Topology	19
3.6.2	Identifying Central Articles with PageRank	20
3.6.3	Curating a Shortlist of Influential Articles	20
3.6.4	The shortlist	21
3.6.5	What are the main conclusions of the "shortlist"?	22

3.6.6	What methods and data are utilized to study the unemployment question?	23
3.6.7	Quantification of AI exposure: a subliteration with great potential	24
3.7	Threats from an emerging paradigm	26
3.8	Conclusions	26
3.9	Theoretical framing - Book review: Prediction and Power	27
3.9.1	Importance and broader context	27
3.9.2	Narrower context: Prediction machines	28
3.9.3	Theory of Prediction Machines:	28
3.9.4	Power and prediction: going beyond the basic theory	29
3.9.5	More detailed relationship between prediction and judgement	30
3.9.6	Evaluation of the book's contributions	31
3.9.7	Takeaways and recommendation	32
4	Interviews	34
4.1	Topic and motivation	34
4.2	Literature Context	34
4.3	Definition of Research Topic	35
4.4	Quantitative Background	35
4.4.1	General Demographics of Respondents	36
4.4.2	Current AI Usage	38
4.4.3	Prospects	41
4.5	Qualitative Examination	42
4.5.1	Method	42
4.5.2	Interview Structure	43
4.5.3	General Observations	43
4.5.4	Basic Findings	45
4.5.5	Impact on Employment	47
4.5.6	Junior vs. Senior Programmers	48
4.5.7	Views on the Role of AI	49
4.5.8	Summary	50
4.5.9	The Speed of Spread - "In-Between Times"	52
4.5.10	Policy Recommendations	53
5	Large scale analysis of online job market data	54
5.1	Topic and motivation	54
5.2	Methodological Considerations	54
5.2.1	Data Source	55
5.2.2	Text Processing Methodology	56
5.3	Analysis	61

5.3.1	Exploratory Data Analysis	61
5.3.2	General Trends	64
5.3.3	Where Are the Latest Technologies?	69
5.3.4	Sectoral Distribution	69
5.3.5	Company-Level Analysis	70
5.3.6	Attempts at Group Formation	74
5.3.7	Job Role Analysis	75
5.4	Conclusions	76
5.5	Limitations	78
5.6	Policy Recommendations	79
6	"Trust my AI" - Survey on the hurdles of AI adoption in business contexts	80
6.1	Topic and motivation	80
6.2	Introduction and context	80
6.2.1	Emerging trust problems with AI systems	80
6.3	Research question	81
6.4	Methodology	81
6.4.1	Survey Design and Distribution	81
6.4.2	Sample Characteristics	82
6.4.3	Data Collection Period	82
6.4.4	Limitations	82
6.4.5	General description of survey participants	83
6.5	Survey results	87
6.5.1	AI usage in the organizations	87
6.5.2	AI's Importance and Impact on the Organization	88
6.5.3	Trust vs. usage	92
6.5.4	Organizational readiness and fears	95
6.5.5	Tools for Establishing Trust	98
6.5.6	Value of trust tools	102
6.6	Discussion of Results	104
6.6.1	Policy Implications	105
6.6.2	Building AI Trust in practice: Organizational strategies	107
6.6.3	Key takeaways: AI Trust for competitive advantage	109
7	Conclusions and Policy Recommendations	111
7.1	Conclusions	111
7.2	Policy Recommendations	112
7.2.1	Company Level	112
7.2.2	Country Level	112
7.3	Epilogue - Sense of urgency	113

Bibliography	114
A DOIs of the articles analysed	124
B Structure of the Interviews (Outline)	138
C Demographics data of Respondents	140
D Stack Overflow 2023 Hungarian Senior Respondents' Opinions Count	141
E Specific questions of the survey	142
F Regular Expression for Detecting Technological Terms	147
G DSPy Optimized Prompt for Semantic Classification	149

List of Figures

3.1 Broadly relevant articles	8
3.2 Compute usage of SOTA AI models in time	9
3.3 Automated speech recognition history (modified to show approximate human level performance)	10
3.4 Image classification on ImageNet	10
3.5 Geoffrey Hinton's lecture on the history of Deep Learning	10
3.6 Google Trends Search results for (Global) for the search term Deep Learning Google, 2023a	11
3.7 Google Trends results (Global) for the search term generative AI Google, 2023b	12
3.8 Time distribution of the publications in the curated set (n=250)	13
3.9 LDA topic ratios, temporal distribution in the curated set and their human interpretation	15
3.10 Ratio of articles having overall positive vs. negative outlook	17
3.11 Ratio of total mentions of positive vs negative phenomena in the corpus	17
3.12 Ratio of articles being more positive, negative or undecided in time throughout the corpus	18
3.13 Ratio of mentions for positive vs negative phenomena in time throughout the corpus	18
3.14 Citation graph of the "curated sample"	19
3.15 PageRank score of the papers in the "curated sample"	20

3.16	Correlation of original and "generative AI" adjusted AIOE scores	25
4.1	Age distribution of respondents (n=57240)	36
4.2	Region distribution of respondents (n=57240)	37
4.3	Organizational size distribution of respondents (n=57240)	37
4.4	Current AI usage of respondents (n=57240)	38
4.5	Current AI usage by age	39
4.6	Current AI usage by region	39
4.7	Current AI usage by organizational size	40
4.8	How will your work change in a year?	41
4.9	Age distribution of interview subjects (n=20)	42
5.1	Temporal distribution of job advertisement publications.	61
5.2	Spatial distribution of job advertisements. (Budapest scaled down by 1/3) .	62
5.3	Distribution of job advertisement sources.	63
5.4	Top 15 settlements by frequency of job advertisements.	63
5.5	Confusion matrix of agreement/disagreement between RegEx and DSPy solutions	64
5.6	Spatial distribution of AI-related (classified based on DSPy) job advertisements - top 10 cities	65
5.7	Frequency of AI-related job advertisements (classified based on DSPy) . . .	66
5.8	Google Trends data for "AI jobs" search term from searches initiated in the United States, measured in Google's own relative metric	66
5.9	Google Trends data for "artificial intelligence" search term from searches initiated in Hungary, measured in Google's own relative metric	67
5.10	Percentage of AI-related job advertisements (classified based on DSPy) - Total population: 1382827	67
5.11	Forecast for AI-related job advertisements based on DSPy – Model ARIMA 0,1,1 with seasonal order 0,0,12	68
5.12	Forecast for the percentage of job advertisements mentioning AI-related terms based on DSPy - Model ARIMA 0,0,4 with seasonal order 0,1,12	68
5.13	Sectoral distribution - based on TEÁOR main activity - of the top 100 companies posting AI advertisements (based on DSPy)	70
5.14	Sectoral distribution of the top 50 companies posting AI advertisements (based on DSPy)	71
5.15	Number of AI-related advertisements by Bosch Group	72
5.16	Percentage of AI-related advertisements by Bosch Group	73
5.17	Percentage of AI-related advertisements by IBM Group	73
5.18	Percentage of advertisements mentioning AI based on RegEx by IQVIA Group	74
5.19	Top 50 job titles by frequency	75

6.1	Country distribution of respondents	83
6.2	Industry of respondents	84
6.3	Organizational Sizes of Respondents	85
6.4	Organizational role of respondents	85
6.5	Organizational roles of respondents - Merged	86
6.6	How frequently do you or your organization use AI technologies in your professional activities?	87
6.7	Strategic Importance of AI by Organization Size - Percentage	88
6.8	AI's impact on your industry by Organization Size	90
6.9	Distribution of Categories per Use Case - Internal Use Cases	91
6.10	Distribution of Categories per Use Case - Customer-Facing Use Cases	91
6.11	Percent Distribution of AI Usage by Role	93
6.12	Normalized relationship between role at work and trust in AI systems	94
6.13	Relationship between AI usage at work and trust in AI systems	95
6.14	Respondent's agreeing that AI poses risk for a given issue - Yes: Dark Grey, No: Light Grey	96
6.15	Respondents perception of safety measures in their organizations	98
6.16	Participants perception of the value of an LLM assessment tool	103
6.17	Percent Distribution of LLM Assessment Tool Value by Role	103
6.18	Value of an LLM Assessment Tool by Organization Size	104

List of Tables

3.1	Shortlist of influential articles	21
C.1	Roles, Company Sizes, and Ages of Interviewees	140
D.1	Assessment of Software Development Tasks by Similarity	141

Chapter 1

ABSTRACT

This dissertation investigates the multifaceted impacts of artificial intelligence (AI) adoption on labor markets, with a focus on Central and Eastern Europe. Treating AI as a general-purpose technology, the study explores substitution and augmentation effects on employment, evolving skill demands, organizational responses, and policy frameworks required for equitable technological transitions.

A mixed-methods approach underpins the analysis. First, a scientometric meta-analysis of 250 core publications employs Latent Dirichlet Allocation and PageRank citation networks to identify prevailing themes—such as job polarization, productivity gains, digital transformation of occupations, algorithmic management, and robotics in industry. Sentiment analysis of these works reveals a modest pessimistic bias regarding AI's effects on unemployment. Second, qualitative interviews with twenty professionals from Hungary and surrounding countries, contextualized by global developer survey data, uncover differential experiences of AI usage: senior staff report efficiency gains, while junior workers express displacement concerns. Third, large-scale text classification of 1.3 million Hungarian online job advertisements quantifies AI exposure across regions, sectors, and roles, uncovering leading sectors as well as a general lag in AI adoption using a hybrid of regular expressions and prompt-optimized language models. Fourth, an international survey examines organizational trust barriers and enablers, identifying governance, transparency, and quality assurance tools as critical factors for AI integration.

Key findings confirm that AI induces both job creation and displacement, driving wage polarization and skill mismatches. Regional and firm-size disparities persist, and trust deficits impede broader adoption.

By combining bibliometric synthesis, empirical NLP analysis, fieldwork, and survey data, this work advances theoretical understanding and highlights the main issues—such as upskilling, trust-building, and governance—that must be addressed to foster inclusive, sustainable AI-driven labor transformations.

Chapter 2

INTRODUCTION AND HYPOTHESES

The rapid proliferation of artificial intelligence (AI) systems across a wide spectrum of economic domains has reignited fundamental questions about technological change, labor dynamics, and societal transformation. While public debate often oscillates between utopian promises and dystopian fears, this dissertation approaches AI as an economic force - a general-purpose technology with deep implications for the organization of labor and capital. From a pragmatic and empirically grounded perspective, we focus on AI's observed and projected effects on employment, productivity, skill demands, and trust in technological adoption.

The following introductory chapter outlines the analytical scope, research questions, and methodological innovations of this dissertation. Building upon a synthesis of consensus of over 2,000 academic sources - narrowed to 250 core publications through a rigorous screening and citation analysis pipeline - subsequent qualitative and large scale quantitative as well as survey based own research this work develops a structured and evidence-based view on how AI is transforming labor markets, particularly in Central and Eastern Europe.

2.1 Analytical Premise

AI is not a single technology but a bundle of advancing capabilities - from predictive modeling to robotic automation and natural language processing. These capabilities affect labor markets through multiple channels:

- Substitution of human labor by automation
- Transformation of occupational skill requirements
- Creation of new job roles and augmentation of existing ones
- Sectoral shifts in labor demand and organizational redesign
- Sociotechnical frictions such as trust, bias, and decision opacity

This dissertation investigates these dynamics through a multi-faceted approach, combining large-scale quantitative analyses, qualitative fieldwork, and a firm theoretical framing.

2.2 Research Questions

The central inquiries guiding this dissertation are:

1. What are the dominant themes and concerns in the academic literature on AI and labor markets?
2. How is AI adoption reshaping employment structures and skill demands in various economic sectors?
3. How is the Central European labor market (especially Hungary) responding to AI-induced pressures?
4. What forms of organizational adaptation - in particular, judgment integration and trust governance - enable successful AI implementation?
5. What are the policy and governance tools available to mitigate risks and foster equitable technological transitions?

2.3 Methodological Innovations

This research adopts a multi-method design characterized by:

- scientometric meta-analysis of 250 key papers using Latent Dirichlet Allocation (LDA) and PageRank-based citation network modeling
- qualitative interviews with 20 Central European professionals, contextualized by the Stack Overflow global developer survey (90,000+ respondents)
- empirical NLP-based classification of 1.3 million online Hungarian job advertisements using regex heuristics and LLM prompt optimization (via DSPy)
- design and execution of the *Trust My AI* international survey on organizational AI trust factors
- synthesis through the theoretical lens of the "Power and Prediction" framework, with emphasis on system change, judgment bottlenecks, and adoption hurdles

2.4 Key Hypotheses

Based on preliminary analyses and theoretical grounding, the dissertation explores the following working hypotheses:

- **H1:** The academic literature is skewed towards cautious pessimism regarding AI's labor market impact, with a stable majority of publications emphasizing risks over benefits.
- **H2:** Senior professionals experience measurable efficiency gains from AI use, while junior professionals face displacement risks - undercutting future junior employee supply through career blocking ("junior paradox").
- **H3:** AI adoption remains uneven across regions, firm sizes, and professional seniority levels—with trust, skill mismatch, and governance gaps as major barriers.
- **H4:** Organizational readiness for AI correlates with levels of internal trust and availability of quality assurance and evaluation-support tools.
- **H5:** Without proactive policy intervention, AI will likely amplify existing labor market inequalities rather than reduce them.

2.5 Structure of the Dissertation

The work proceeds as follows:

- chapter 3: a systematic literature review of the economic and labor market impacts of AI, as well as a theoretical framing based on economics of AI
- chapter 4: fieldwork-based analysis combining quantitative survey data and qualitative interviews
- chapter 5: large-scale empirical analysis of job advertisements in Hungary
- chapter 6: survey-based exploration of organizational trust in AI adoption
- chapter 7: synthesis and policy recommendations

2.6 Historical context

The development of Artificial Intelligence (AI) systems has gained significant momentum over the past decade and a half, particularly with the latest wave of deep neural network technology, known as *Deep Learning* (for the origin of the term, see also: Dechter, 1986,

Fradkov, 2020), and the so-called *foundational models* (Bommasani et al., 2022a), especially the *Large Language Models* (LLMs; Brown et al., 2020) with their wide availability and low cost (for the decrease in consumer prices of LLM usage, see also: Appenzeller, 2024).

However, the concept of Artificial Intelligence is not new: the development of AI can be divided into several phases, during which the technology attempted to simulate certain aspects of human thinking in different forms and approaches. AI systems that emerged in the 1950s were primarily based on symbolic logic, fitting into the *expert system* paradigm: they worked with strict rule systems, relying on manually codified knowledge bases. These systems were functional in certain environments, such as medical diagnostics or financial analyses, but could not dynamically adapt to new situations.

In the early 2000s, the *Machine Learning* (ML) paradigm took the lead, enabling learning from data without explicitly coding rules. ML methods, particularly supervised learning, were successfully applied in a broad range of fields from image recognition to market predictions. The technology reached a breakthrough point in 2012 when deep neural networks (*Deep Learning*) demonstrated their dominance in computer vision and other AI applications. Nevertheless, AI remained strictly *narrow AI*, operating with unique, task-specific models optimized for specific tasks.

After 2020, AI research underwent another paradigm shift, with the emergence of *Large Language Models* (LLMs) marking the breakthrough. LLMs were not trained for a single specific task but demonstrated general problem-solving capabilities, capable of synthesizing human language and solving complex tasks through it. The release of OpenAI's *ChatGPT* model in November 2022 widely demonstrated the capabilities of these models, resulting in global media attention and unprecedented industrial and economic interest.

In light of this, it is not surprising that the application of AI in the world of work has also undergone a rapid development Anthropic, 2025, a McKinsey survey Singla et al., 2024 from early 2024 found that 65% of respondents' organizations use generative AI, mainly LLMs in some capacity, and it has appeared in Hungarian business practices as well, see also Szabados, 2025.

For a more detailed analysis of the historical developments in AI, its timeline and reception by the public see the Literature Survey chapter.

Chapter 3

LITERATURE REVIEW

3.1 Topic and motivation

In this chapter, based on large scale literature and bibliometric studies I present the results of a deep literature review on the scientific consensus and topics about AI's impact on the job market. (This chapter was published in Szabados, 2024a.)

3.2 Background

Over and beyond the prevalence of recent discussion about potential existential risks of AI technology, from an economist's perspective, the broad set of technological capabilities that are collectively referenced by the term "Artificial Intelligence" (further on: AI) represents a new kind of force in the economic life of society, that bears huge promises, as well as hold significant risks for the prosperity and wellbeing of society at large. Since there is a significant amount of work dedicated already to the topic of "economic effects of AI", especially it's aspects pertaining to the labour markets, a synthesis and overview of the broader literature can be useful to provide some orientation about the current state of discussion, as well as to shed light on areas of further potential research, thus, with a broad scope analysis we endeavor to provide such an overview.

3.3 Methodology

This chapter aimed to conduct a thorough analysis of the existing literature on the impact of artificial intelligence (AI) on the labor market. Given the extensive research activity in this area, our methodological framework was designed to systematically evaluate a broad array of academic publications to gain comprehensive insights into the topic.

3.3.1 Literature Search and Collection

Our literature search began with a targeted strategy using Google Scholar to ensure access to a wide range of academic journals and conference proceedings. We complemented this online search with a manual examination of the citation networks from key papers, which allowed us to capture additional relevant studies that might not have been indexed or immediately apparent in digital search results.

3.3.2 Screening and Selection

The search strategy resulted in an initial pool of over 2000 publications. To determine relevance, we applied a two-stage screening process. First, an automated filtering system was used to perform an initial screening based on keywords and abstract content. This reduced the pool to approximately 725 potentially relevant articles. Second, we engaged in a manual curation step narrowing down the number of immediately relevant papers for the topic of AI in the labor market. This process led to a curated set of 250 articles for in-depth analysis.

3.3.3 Thematic Analysis

To uncover the thematic structure within the curated set, we utilized Latent Dirichlet Allocation (LDA) Blei et al., 2003, a technique utilized in Natural Language Processing which allowed us to identify meaningful topics, assess their contribution to the given documents, thus to track the prevalence of these topics over time, which informed us about the focus areas within the field and their evolution.

3.3.4 Citation Analysis

With the aim of identifying the most influential studies, we conducted a citation graph analysis which helped us isolate a shortlist of 13 papers that were not only frequently cited but also held significant sway in shaping the discourse in the field. These papers received a detailed examination, wherein we scrutinized their hypotheses, methodologies, datasets, and findings for getting a more in depth view on the consensus in the field.

3.3.5 Quantitative and Content Analysis

Each paper on the shortlist underwent a rigorous quantitative and content analysis. We examined the papers' predictions and conclusions about the impact of AI on the labor market. We also identified a strand of research specifically focusing on the quantitative estimation of AI's effects on different occupations based on their associated skill sets.

3.3.6 Identification of Research limitations

Throughout our analysis, we remained cognizant of the limitations inherent in the existing body of work. We documented these limitations and proposed directions for future research to address these gaps and to advance the understanding of AI's role in the labor market.

In the following sections of this paper, we will provide detailed descriptions of the methods used at each stage of our analysis.

3.4 General trends

As a first level of analysis, we looked at the broadest set of 725 "generally relevant" articles identified by filtering the results (>2000 entries) of automated Google Scholar searches carried out for such expressions as "effects of Artificial Intelligence on the labour market".

We analyzed the temporal distribution of the broad set of 725 to get a grasp on the general trend of interest in the field.




Figure 3.1: Broadly relevant articles

Image Source: Author's illustration

As we can see on Figure 3.1, there is a strong increasing trend visible in the number publications, with an inflection point at 2015. (For better illustration, we extrapolated - based on the proportion of days remaining in the year at the time of writing - the number of publications in 2023.)

This begs the question, why did this trend appear just in this time period, what are the prime events moving it?

3.4.1 Why the trend?

Though the long term progress of some classical areas in AI research, like automated speech recognition or image classification showed steady progress through the decades, the gains in performance (or decrease in error rate, which can be considered the same) were achieved at a cost of investing considerable amounts of engineering manpower. As for example the case of "AlexNet" Krizhevsky et al., 2012, the first really successful Deep Learning model to beat "conventional" (non-neural network based statistical models mainly relying on manual, expert driven feature engineering) illustrates, a paradigm shift appeared in the form of the new "end-to-end learning" paradigm. This paradigm enabled the application of single models without extensive manual feature design on large scale datasets (eg. ImageNet Deng et al., 2009 in case of image classification), letting the learning procedure "figure out" the necessary features, essentially trading human expert engineering hours for computation.

The increased "compute" utilization in what Sevilla-Lara and Learned-Miller, 2022 call the "Deep Learning era" is quite visible in the chart below (Figure 3.2).




Figure 3.2: Compute usage of SOTA AI models in time

Image Source: Sevilla-Lara and Learned-Miller, 2022

This - combined with the wide availability of computation resources and larger datasets lead to a breakthrough in performance on common AI benchmarks as illustrated by Figure 3.3 and Figure 3.4

As one of the "founding fathers" of Deep Learning, the Nobel prize laureate and Turing Award winning computer scientist Geoffrey Hinton pointed out in his public lecture (Hinton, 2017): The theoretical advancements of the late 90s did not bear fruit until enough data and computing power became available (to a suitable model structure, that is Deep Learning, since other previous modeling architectures did not directly benefit from such a boost in data / computing power)

As the notes of Hinton illustrate, from a decades long held substantially lower performance (26% error rate) in image classification, in 2012 the first well trained Deep Learning model cut the error in nearly half, then in the span of approximately 3 years, on this task




Figure 3.3: Automated speech recognition history (modified to show approximate human level performance)

Image Source: Greene, 2017




Figure 3.4: Image classification on ImageNet

Image Source: Paperswithcode.com




Figure 3.5: Geoffrey Hinton’s lecture on the history of Deep Learning

Image Source: Hinton, 2017

human level performance was achieved (see Figure 3.5 and Figure 3.3).

As the knowledge about this breakthrough - and the rapid increase in performance - became more commonplace (as evidenced by e.g. the Google Search Trends for the term "Deep Learning" - Figure 3.6) this coincides with the increase in interest by the economic sciences in the effect of AI.




Figure 3.6: Google Trends Search results for (Global) for the search term Deep Learning Google, 2023a

Image Source: Author's illustration

Looking closer to the number of AI related economic publications in time, there is also a small "saturation" effect visible in 2020-21, until the concept of generative AI burst into consciousness with the release of ChatGPT, and gave a new push to the discussion.

It is worth noting, that technology again rapidly and qualitatively changed with the advent of "generative AI" (which itself is somewhat of a misnomer, covering the combined advancements in Deep Learning based Large Language Models - what Stanford researchers Bommasani et al., 2021 in their paper "On the Opportunities and Risks of Foundation Models" call "foundational models", with (mainly) diffusion based image generation models (see Ho et al., 2020), and their potential combination under the umbrella of "multi-modality").

The sudden gain in capabilities for "foundational models" re-sparked the interest of economic research in the effects of AI. (see Figure 3.7)

With this interest in mind, it is essential to point out, that the change in capabilities in this case was not just "quantitative", the set of capabilities for the foundational models expanded rapidly, thus, their practical applications changed in at least three meaningful ways:

- the accuracy of ML models on some tasks increased
- the scope of their applicability exploded
- the technological threshold for their application dropped




Figure 3.7: Google Trends results (Global) for the search term generative AI Google, 2023b

Image Source: Author's illustration

Based on this, we can conclude, that the discussion about AI currently must aim at a more broad set of capabilities. Essentially: talking about AI's impact before November 2022 is different than after it.

3.5 Detailed analysis: The "Curated list"

For detailed analysis, a manually curated list of 250 articles were selected, that focused more on the direct labor market influences, and at least touched upon the pertinent question of technology-induced unemployment (see Figure 3.8). The main selection criterion for this manual curation was that the stated focus of the article had to have direct relevance for the theme of AI's effect on the labor market and had to at least partially concern the question of technology-induced unemployment.

During the curation process, we took care to ensure that the temporal distribution of articles reflects that of the broader set, thus ensuring that the general dynamics of topics remain discernable.

3.5.1 General topics present in the "Curated set"

For a more quantitatively driven - thus hopefully more objective - analysis of the topics prevalent in the "Curated set," we decided to utilize Natural Language Processing-based techniques, especially Latent Dirichlet allocation (see Blei et al., 2003), which is a technique for document topic analysis. At its core, the approach assumes that the documents present in a corpus of text (in our case the 250 hand-selected articles) came to existence as a probabilistic mixture of latent topics as a generative distribution, thus every document can be




Figure 3.8: Time distribution of the publications in the curated set (n=250)

Image Source: Author's illustration

interpreted as a kind of mixture of these not directly observed thematic causes.

For the training procedure of the model, we utilized the OpenSource "Gensim" library (see Rehurek and Sojka, 2010), and after experimentation, we settled to utilize 5 topics for modeling. With this technique, we created a topic model of this subliteration, and we identified the following topics (by interpreting the low-level LDA results). The topic labels themselves were created by the interpretation of the keyphrase probabilities of the LDA model and confirmed by manual inspection of the broader set of 750 articles.

1. topic 1: "job dynamics and wage variations" ("gigification", wage change and modes of changes in the work)
2. topic 2: "technological innovation and labor productivity (the effects on productivity gain)"
3. topic 3: "occupational skills and digital transformation" (the "unemployment question")
4. topic 4: "human interaction with emerging technologies" (including AI hiring, bias, the effects of algorithmic management)
5. topic 5: "robotics and labor share in industries" (Physical robotization of manufacturing as a separate topic)

3.5.1.1 Some takeaways

The presence of this topic distribution as it stands can hint to some interesting associations and structure inside the literature.

1. Job Dynamics and Wage Variations: The presence of this more distinguishable topic suggests that the literature is not just concerned with the employment question in the frames of "traditional" employment relations, but the differentiation between gig work and traditional employment is getting emphasized, hinting at potential variance in job stability and wage trajectories. This seems to be a noteworthy concern for the literature, raising questions about how these trends may diverge from established labor market behaviors and whether these differences could lead to alternative employment models becoming more normative in certain sectors.
2. Technological Innovation and Labor Productivity: The presence of technology as a driver of productivity is well-documented, but its intersection with AI introduces new considerations. There appears to be a possibility that AI and automation are altering the established dynamics between capital and labor, with open questions on whether these changes lead to complementarities or new forms of substitution.
3. Occupational Skills and Digital Transformation: The observed association between the demand for new skills and unemployment rates seems to be the main area of concern, which suggests a shifting landscape. Though the causal pathways are not completely clear, this seems to be one of the main areas of interest for research, so quantifying and analyzing the changes in the skill sets demanded by employers driven by AI adoption remains one of the main areas of concern.
4. Human Interaction with Emerging Technologies: The inclusion of AI in HR processes points to an emergent field of inquiry within labor economics, where technology is not just a backdrop but a participant. The considerations around bias and algorithmic decision-making indicate a nascent concern with profound implications, which invite further examination into how these technologies reshape labor market practices.
5. Robotics and Labor Share in Industries: Robotics' specific mention reflects its visible impact on manufacturing, but also its slight distinctiveness in the area of study, since arguably it impacts a very specific subset of job areas, mainly in manufacturing and the agriculture sector, so its effects, though maybe more easily measurable, are focused on a very specific area.

In general, we can argue (giving additional legitimacy to the choice of LDA as a model specifically) that there is a deeply interwoven set of relations between the different areas. Hence, when analyzing the effects of AI on the labor economy, the picture is more like overlaying effects on top of each other, rather than some nicely separable set of distinct mechanisms.

3.5.1.2 Topics in time

To further examine the dynamics of the topic distribution over time, getting a sense of its stability, as well as ratio of relative frequency, we analyzed the topics' prevalence by averaging their presence through the years with respect to the articles in the "curated set" (see Figure 3.9). For this, we utilized the property of LDA models to represent documents as a weighted mixture of topics. Thus, we could easily take the mean of the topic presences for a subset of articles published in a given year.




Figure 3.9: LDA topic ratios, temporal distribution in the curated set and their human interpretation

Image Source: Author's illustration

The resulting chart can lead us to the conclusion that the topic distribution over time is generally stable. The interest in AI's effect on hiring and the fairness question increases around 2016, presumably with the more widespread adoption of AI/ML solutions in HR processes, thus increased scrutiny from the scientific community. Beyond this, the skills and employment question is pretty important.

3.5.2 What is the consensus about the effects on unemployment specifically?

3.5.2.1 Quantitative Analysis

Our aim was to establish an estimate of the consensus in the selected literature (250 papers) about the beneficial or harmful effects of AI on employment. For this, we endeavored to assess each paper's overall sentiment concerning AI's impact on unemployment by counting the number of mentions of "positive" or "negative" phenomena related to employment within the text. The count was strictly numerical, irrespective of the length or depth of the discussion surrounding each mention.

3.5.2.2 Classification Criteria

The classification of papers was contingent on a direct comparison between the counts of positive and negative mentions. We operationalized our assessment using the following criteria:

1. Class A - Positive Outlook: A paper was categorized as having a positive outlook if the count of positive mentions outnumbered the negative mentions.
2. Class B - Ambiguous Outlook: A paper was deemed to have an ambiguous outlook if the difference between positive and negative mentions was within a 2.5% margin, reflecting an almost equal consideration of both types of phenomena.
3. Class C - Negative Outlook: Conversely, a paper was classified as having a negative outlook if negative mentions surpassed positive mentions in number.

3.5.2.3 Positive and Negative Phenomena

Positive phenomena were defined as any mention of trends or effects where AI resulted in:

- the creation of new job types
- the simplification of job tasks, potentially requiring less skill
- a shift towards more fulfilling tasks within existing jobs

Negative phenomena included mentions where AI was associated with:

- a decreased need for human workers, potentially resulting in layoffs
- the competitive displacement of human labor across various sectors

3.5.2.4 Threshold for Classification

To classify the outlook of each paper, a threshold was set. If a paper had more than a 2.5% greater count in either direction (positive or negative mentions), it was classified accordingly. For instance, if positive phenomena were mentioned with a frequency greater than 2.5% compared to negative mentions, the paper was categorized under Class A.

3.5.2.5 Results

If we utilize the classification framework, and then count the articles accordingly, the ratio of positively and negatively dominated articles overall is as follows (Figure 3.10):

If we would like to paint a more nuanced picture and calculate the ratio of individual positive or negative mentions (Figure 3.11), the consensus is maybe less bleak, but nonetheless remains negatively biased.




Figure 3.10: Ratio of articles having overall positive vs. negative outlook

Image Source: Author’s illustration




Figure 3.11: Ratio of total mentions of positive vs negative phenomena in the corpus

Image Source: Author’s illustration

As an added exercise, we analysed the temporal distribution of the "overall positive", "overall negative" and "undecided" articles in time (as in: year of publication) (Figure 3.12), as well as the ratio of total individual mentions (Figure 3.13), and we found, that the not dominating, but none the less very perceptible negative dominance is generally stable.




Figure 3.12: Ratio of articles being more positive, negative or undecided in time throughout the corpus

Image Source: Author's illustration




Figure 3.13: Ratio of mentions for positive vs negative phenomena in time throughout the corpus

Image Source: Author's illustration

Overall, we can conclude, that the surveyed literature can be characterized by a certain "cautious pessimism", so stopping short from being completely dominated by the outlook of negative effects of AI on the labour market, but being generally slightly pessimistic in outlook.

3.6 Establishing a Foundational Subset through Citation Graph Analysis

As a methodology for identifying the core literature, to distill a foundational subset of articles having the dominant influence on the field, from the collection of 250 academic papers we implemented a citation graph analysis. This involved constructing a directed graph where the nodes represent the articles within our curated set, and the edges represent citations between these articles. For this phase of analysis, we focused exclusively on the internal citation dynamics, intentionally omitting external citations – those references pointing outwards from the curated subset. This allowed for a more concentrated examination of the discourse and intellectual lineage within the scope of our research question.

3.6.1 Assessing the Network Topology

Preliminary observations suggest that the citation graph (Figure 3.14) displays characteristics reminiscent of a 'small-world network', a concept rooted in the field of network theory. Small-world networks are marked by high clustering and short path lengths between nodes. In the context of our citation graph, this could manifest in a pattern where a limited number of highly-cited papers form the backbone of the research area, with a multitude of less-cited works branching off from these central nodes. These foundational works typically serve as keystones in the construction of the field's academic edifice, setting the research agenda and framing the scientific discourse.




Figure 3.14: Citation graph of the "curated sample"

Image Source: Author's illustration

To substantiate this observation, we used numeric methods for examining the "centrality" of given articles.

3.6.2 Identifying Central Articles with PageRank

Utilizing the PageRank (Page et al., 1999) algorithm, an approach famously employed by Google for ranking web pages, but also widely used as a general measure of network centrality, we utilized to identify the most 'central' articles within our citation graph. PageRank serves as a measure of node influence in a network, based on the notion that connections to high-scoring nodes contribute more to the score of a node than equal connections to low-scoring nodes. By adapting this algorithm to our citation network, we were able to objectively quantify the influence of each paper within our subset, beyond mere citation counts, taking into account the 'quality' of citations in terms of the influence of the citing papers. (see Figure 3.15)




Figure 3.15: PageRank score of the papers in the "curated sample"

Image Source: Author's illustration

Looking at the PageRank values, we indeed see signs of a handful of "core" papers dominating the citation graph, so we feel vindicated in creating a "shortlist" of most influential papers.

3.6.3 Curating a Shortlist of Influential Articles

To further refine our analysis and extract a 'shortlist' of the most influential articles, we employed the technique outlined in "Finding a Needle in a Haystack: Detecting Knee Points in System Behavior" by Satopää et al., 2011. This method, designed to identify points of interest in a system's behavior, was adeptly repurposed to determine a threshold for influence within the citation graph. By plotting the distribution of PageRank scores and identifying

the 'knee point' — the point where the curve sharply changes — we were able to delineate a natural cutoff. Papers above this threshold are considered as having a disproportionately large influence on the research field and are therefore included in our shortlist.

3.6.4 The shortlist

As final representatives of the "shortlist", we included:

Table 3.1: Shortlist of influential articles

Author(s)	Title	Year	DOI
Autor	Why Are There Still So Many Jobs? The History and Future of Workplace Automation	2015	10.1257/jep.29.3.3
Acemoglu and Restrepo	Robots and Jobs: Evidence from US Labor Markets	2020	10.1086/705716
Acemoglu and Restrepo	Robots and Jobs: Evidence from US Labor Markets	2017	10.3386/w23285
Huang and Rust	Artificial Intelligence in Service	2018	10.1177/1094670517752459
Agrawal et al.	Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction	2019	10.1257/jep.33.2.31
Brynjolfsson and Mitchell	What can machine learning do? Workforce implications	2017	10.1126/science.aap8062
Frank et al.	Toward understanding the impact of artificial intelligence on labor	2019	10.1073/pnas.1900949116
Brynjolfsson et al.	What Can Machines Learn and What Does It Mean for Occupations and the Economy?	2018	10.1257/pandp.20181019
Degryse	Digitalisation of the Economy and its Impact on Labour Markets	2016	10.2139/ssrn.2730550
Felten et al.	A Method to Link Advances in Artificial Intelligence to Occupational Abilities	2018	10.1257/pandp.20181021
Furman and Seamans	AI and the Economy	2019	10.1086/699936
DeCanio	Robots and humans – complements or substitutes?	2016	10.1016/j.jmacro.2016.08.003
Brougham and Haar	Smart Technology, Artificial Intelligence, Robotics, and Algorithms (STARA): Employees' perceptions of our future workplace	2017	10.1017/jmo.2016.55

Source: Author's data

Based on this analysis, we can conclude that by far the most "central" work of this curated set is "Why Are There Still So Many Jobs? The History and Future of Workplace Automation" by Autor, 2015. This is a seminal work that "kickstarted" the study of AI's

effects on the labor market early on in the "revolution".

Also noteworthy, but somewhat misleading with two entries, is Acemoglu and Restrepo's paper "Robots and Jobs: Evidence from US Labor Markets" (see Acemoglu and Restrepo, 2017). It first appeared as a report from the National Bureau of Economic Research and then later got published in the *Journal of Political Economy* in 2020 (see Acemoglu and Restrepo, 2020). Combining its citations makes it also a remarkable cornerstone of the topic.

Below, we endeavor to summarize the key findings of the centrality measure based "shortlist".

3.6.5 What are the main conclusions of the "shortlist"?

Research represented by the "shortlist" literature has focused on quantifying and illustrating the impact of automation, robotics, and artificial intelligence (AI) on the labor market and the economy. Amongst the various papers exploring the implications, some common hypotheses and conclusions emerge:

One common hypothesis is that automation and technology do not necessarily lead to job loss but result in job displacement and polarization. This is supported by research papers such as Autor, 2015 and Acemoglu and Restrepo, 2017, which suggest that routine tasks are being automated, leading to growth in high-education, high-wage jobs and low-education, low-wage jobs at the expense of middle-wage, middle-education jobs such as bank tellers and brokers. It is also proposed that the main economic issue will be one of distribution, rather than scarcity.

Another common finding is that automation both eliminates and displaces jobs while also raising the value of tasks that are uniquely supplied by humans. Agrawal et al., 2019 conclude that automation has ambiguous labor market impacts, as it can automate prediction tasks but may also create new opportunities. Brynjolfsson and Mitchell, 2017 suggest that AI can lead to job displacement in specific occupations such as legal research and gaming but create new opportunities and complement human skills in others. The key factor for workers to benefit from automation is to engage in tasks that are complemented by AI, and require some additional human skills, such as intuitive and empathetic abilities.

The impact of robotics and AI on employment and wage inequality is another recurring theme in the literature. Acemoglu and Restrepo, 2020 find that exposure to industrial robots in the US labor market negatively affects employment and wages, particularly in low- and medium-skill occupations such as manufacturing. This finding is supported by other studies, such as Frank et al., 2019, which suggest that the rise of AI could lead to job polarization and income inequality. Furman and Seamans, 2019 investigate the impact of AI on occupational abilities and conclude that certain occupations such as drivers and retail workers are more susceptible to advances in AI technology.

There is also a growing consensus that the impact of technology on the labor market

depends on factors such as task characteristics, contextual factors, and skill requirements. Brynjolfsson et al., 2018 argue that the successful application of AI depends on a variety of task characteristics and contextual factors, and that job bundling of tasks can offer diversification with respect to machine learning exposure. E. W. Felten et al., 2018 suggest that jobs that can be broken down into homogeneous tasks are more likely to be replaced by AI, even if they require higher intelligence. They also emphasize the need to acquire intuitive and empathetic skills as a strategy to counteract large-scale displacement due to AI replacing lower-skilled jobs.

In conclusion, the "shortlist" literature suggests that automation, robotics, and AI have varying impacts on the labor market and the economy. While the main concern still is job displacement and polarization, the potential for job creation and complementarity also exists. The distributional effects and implications for wage inequality represent the most important consideration.

Regarding the specific impact of technology on a given field, the consensus is that it depends on factors such as task characteristics, skill requirements, and contextual factors typical in a given area of economic activity.

On the level of broader policy, to navigate the challenges and opportunities posed by automation and technological advancements, investment in human capital and the development of skills that are complemented by technology is crucial. Additionally, policy responses and governance frameworks that address distributional challenges and ensure broad-based benefits are important for inclusive economic growth.

3.6.6 What methods and data are utilized to study the unemployment question?

Looking from a more methodological angle, it is interesting to take note of the different approaches the articles in the "shortlist" take from the angle of data collection and quantitative analysis.

One group of articles, including Autor, 2015 and Acemoglu and Restrepo, 2020, uses empirical data and statistical analysis to investigate the relationship between robot adoption and employment. These articles utilize data on robot usage, employment rates, and wages to estimate the impact of robots on labor markets. They employ statistical models, such as instrumental variable (IV) estimates and regression analysis, to provide quantitative evidence on the relationship between robot adoption and employment outcomes.

Another group of articles, such as M.-H. Huang and Rust, 2018 and Brynjolfsson et al., 2018, explores the implications of AI and machine learning on workforce dynamics. These articles utilize empirical data and statistical analysis to investigate the impact of AI on various sectors, tasks, and occupations. They employ statistical methods, such as regression analysis and correlation analysis, to examine the relationship between AI adoption and employment

outcomes, as well as the effect of AI on productivity and wages.

Other articles, including Frank et al., 2019 and Furman and Seamans, 2019, focus on the impact of AI on the skill requirements of occupations. These articles utilize data on occupation-level skill requirements, automation risk, and AI advancements to analyze the changing skill demands in the labor market. They employ statistical methods, such as correlation analysis and regression analysis, to examine the relationship between AI technologies and occupational abilities.

All in all, one of the main challenges in the field seems to be the fact that hard data quantifying the phenomena of "AI adoption" in itself is hard to come by, so linking changes in general observable measurements (like wage inequality) to the effect of AI adoption is extremely challenging.

3.6.7 Quantification of AI exposure: a subliteration with great potential

While surveying the general consensus of the literature regarding AI's direct unemployment effects, as already stated, it is apparent that the question of skills from the human side is of paramount importance. It is equally important though for the task of quantifying the risk of structural unemployment to analyze the quantifiable metrics of different AI model's performance in given skills, thus bridging the gap between human job taxonomies (for example O*NET - see U.S. Department of Labor and (USDOL/ETA), 2023 or ISCO - see International Labour Office, 2012) showing the necessary skills for certain occupations, and AI's specific automation risk regarding these.

In frames of our investigation, we identified a very promising direction of research, and a corresponding small set of articles trying to carry out exactly this task. If we would have to summarize their methodology, it can be roughly sketched as follows:

1. Take a taxonomy (like O*NET or ISCO) that defines a decomposition of jobs in terms of skills required for carrying them out;
2. take a source describing AI capabilities (like Electronic Frontier Foundation's AI Progress Measurement - see Foundation, 2023), that endeavors to quantify AI "skills" in different taxonomy domains;
3. create a mapping (by typically manual labor) matching the AI "capabilities" to the human skills;
4. estimate progress in AI skills by projecting progress into the future;
5. project this progress via the skill mapping back to the jobs as an "automation risk" measure.

This strand of research is most clearly characterized by the series of papers by Edward W. Felten and colleagues in E. W. Felten et al., 2018, E. W. Felten et al., 2019a, and E. Felten et al., 2021a, which have the same basic pattern: utilizing O*NET for labor taxonomy, EFF's measurement for AI progress, and a mapping created by dedicated labor (acquired via Amazon's Mechanical Turk - see Turk, 2023). They call the resulting metric the "AI occupational exposure" (AIOE) score.

Beyond this series, though, very similar patterns appear in Colombo et al., 2019, though it utilizes semantic embedding methods for mapping, as well as replaces O*NET and ISCO with ESCO; Paolillo et al., 2022 focusing on robotics; but also in such recent works as Pizzinelli et al., 2023 (which explicitly references the Felten group's work) and Gmyrek et al., 2023, and finally an "update paper" from the Felten team E. W. Felten et al., 2023 which even took into account the recent advancements in generative AI.

This latter work is all the more important, since it not just "corrects" for the effects of "generative AI" (so basically "large language models" or "foundational models - see Bommasani et al., 2021 - and diffusion based image generation models - see Ho et al., 2020), but they carry out a kind of "stability analysis", so analyse the correlation between the AIOE score before and after this correction. (Figure 3.16)




Figure 3.16: Correlation of original and "generative AI" adjusted AIOE scores

Image Source: E. W. Felten et al., 2023

The analysis suggests, that the AIOE score is pretty robust.

3.6.7.1 Merits and limitations of this approach

The merits of this research methodology have to be emphasized since it "opens up" the black box of AI and does not treat it as a single set of capabilities – which, in our view, is all the more important since the term "AI" does not even represent a single set of technologies and is not static in time, meaning: newer and newer methods appear under its umbrella with their specific characteristics.

That said, since some of this research is coming from time periods representing previous iterations of technology, some empirical testing of the actual realized automation could be fruitful to reach a more precise estimate between the automation potential and its actual application.

As a general remark, it has to be pointed out that this "shifting of the goalpost," with which we "normalize" or "commoditize" technology (like the Optical Character Recognition, which once was an area at the forefront of computer vision, hence AI research, and now is treated as a solved problem that has "nothing to do with real AI"), makes it in general more difficult to systematically investigate the effects of "AI" (whatever technology it currently means) on the economy. As a more nuanced ground for observation, the different *types of technologies* (like classical expert systems, non-neural machine learning models, "classical" Deep Learning models, and now "foundational" models) would be more appropriate as a subject of study. In this sense, the approach to decompose AI – via possible task and skill investigations – to different aspects is all the more laudable.

3.7 Threats from an emerging paradigm

It is important to note, that recent work on foundational level language models (or simply "Large Language Models", LLMs, as they are called in common parlance) shows, that these models possess - beyond their "mere" language generation skills - strong reasoning and planning abilities (see eg: Kojima et al., 2023 and W. Huang et al., 2022). Reasoning and planning on large scale, in a learned manner was thought this far to be a pretty uniquely human capability, but it seems, that the emergence of these models, and the "Autonomous Agents" and "Agent societies"(Minsky, 1987) built upon them represent a whole new level of capability, that's release as available technological solutions is imminent (see eg. the "Autogen" Framework of Microsoft Research Microsoft Research, 2023), and will dramatically increase the impact of AI automation possibilities, given it is deployed at scale.

This basically means, that it is reasonable to assume, that the AIOE like scores will have to be recalculated pretty soon, or even structural changes might be necessary.

3.8 Conclusions

After conducting a thorough literature review on the impact of artificial intelligence (AI) on the labor market and the economy, several key findings and trends have emerged.

Firstly, the overall consensus is that the effects of AI on employment are complex and nuanced. While there is concern about job displacement and polarization, there is also recognition that automation and AI can create new job opportunities and complement human skills. The distributional effects and implications for wage inequality are significant considerations in the discussion.

The literature highlights the importance of task characteristics, contextual factors, and skill requirements in understanding the impact of AI on the labor market. Different studies have explored various aspects of AI's effects, including job dynamics, labor productivity, occupational skills, human interaction with technology, and robotics in industries. These studies provide valuable insights into the specific ramifications of AI in different fields and shed light on the challenges and opportunities posed by automation and technological advancements.

Quantitative analysis plays a crucial role in studying the unemployment question. Some papers have used empirical data and statistical analysis to investigate the relationship between AI adoption and employment outcomes. Others have focused on the quantification of AI exposure, mapping AI capabilities to human skills and estimating the risk of automation for different occupations. These methodologies provide valuable frameworks for assessing the impact of AI on employment and should be further developed and refined in future research.

It is important to note the limitations of the existing literature. The definition of AI itself is subject to concept drift, and the rapid advancements in technology require ongoing updates and refinements in research methodologies. Additionally, the field would benefit from incorporating a broader range of AI technologies beyond machine learning and taking into account the shifting ground of AI capabilities.

In conclusion, the literature on AI's effect on the labor market and the economy presents a multifaceted picture of both challenges and opportunities. The discussion around job displacement, job creation, wage inequality, and skill requirements provides a foundation for further research and policy development. Future work should focus on refining research methodologies, incorporating a broader range of AI technologies, and addressing the evolving nature of AI capabilities. Additionally, policy responses and governance frameworks that consider the distributional challenges and ensure inclusive economic growth are important to navigate the challenges and opportunities posed by AI.

3.9 Theoretical framing - Book review: Prediction and Power

3.9.1 Importance and broader context

Beyond the previously conducted literature survey - and its slightly pessimistic consensus on the impacts of AI on the labour markets - some more elaborate theoretical frameworks are needed to be able to adequately capture in form of a more compact theory the mechanisms behind the adoption and economic impact of AI technologies. In this regard, the research trio of Agrawal, Gans and Goldfarb made seminal contributions by outlining a general framework about the economic mechanisms of AI application. Thus, I will adopt this framework as a general viewpoint of theory, and below proceed to elaborate it via a detailed

review of the latest iteration of the work "Power and Prediction" - see Agrawal et al., 2022.

3.9.2 Narrower context: Prediction machines

To be able to make sense of this work's contribution, it is essential to take into account the previous, fundamental book from the same authors from 2018, since "Power and prediction" (Agrawal et al., 2022) can be considered a kind of reflection, and further development of the theories laid down in the previous work titled "Prediction machines: The simple economics of Artificial Intelligence" (Agrawal et al., 2018). None the less this is not just a mere "second edition" to the previous, there are substantial modifications, in case some strong revisions of the core theory. With that in mind, the short summary of "Prediction machines" as a foundational piece must be done to give context.

3.9.3 Theory of Prediction Machines:

The authors tried to interpret AI as a *"general purpose technology"*, much in line with the broader literature, where a general purpose technology is not constrained to be applied in a single economic domain, but can have applications across all possible fields. A good - and often used - analogy in the broader literature is electricity (or the steam engine), which had profound and general impact, reshaping all industrial activity. This viewpoint ties the author's theories strongly to the broader literature of economic innovation (and as we will see, that of "creative destruction" Ulgen, 2017 or *"disruption"* Agrawal et al., 2022)

The most concise way to look at AI - according to Prediction Machines - is, that it is *"cheap, good quality prediction"*, that is, the capability to utilize computers to learn and obtain scalably and cheaply the ability to get predictions based on data about phenomena of interest. For example, the capability to obtain a model trained on previously annotated images of tumor X-ray images, and to use it to produce a probabilistic prediction for a new, yet unseen patient's X-ray. Or to build a model on some time series data about customer demand for a given article, and use it to forecast future demand etc.

The argument of the authors was: if a given commodity (in this case prediction) becomes cheap, it's *consumption increases* and it is starting to be utilized as *substitute for other goods and services*. We use more prediction.

Second corollary to this is, that we start to use prediction as a substitute for other goods and services (since it's cost is low and quality is good), so we substitute for example car driving as an activity with a predictive system for steering decisions, that we call self driving in vehicles. We substitute things with prediction.

If we can rely on good quality predictions, it is worth *changing the business processes* accordingly. The example they proposed: if we have good enough prediction for what people will buy, we can start shipping before they order. (In fact, this became reality with a patent of Amazon about "Anticipatory Shipping" Kusnitz, 2014, as well as the case of the company

Stichfix's "Ship-then-shop" - see Choi et al., 2024 - model, that the authors also mention in the later book.) Short: if prediction is good enough, business processes and models can get radically more efficient / overhauled ("disrupted").

Since prediction is based on good quality data, that is *data is "the new oil"* (see Talagala, 2022) for the "engines" of prediction, it's value will go up and become a strategic asset for organizations

Maybe most importantly, the authors introduce the *distinction between "prediction" and "judgement"* in a decision. Former (eg. "this is a tumor with 75 percent probability") can be of great service for the latter ("it is worth it to immediately do an exploratory surgery, even when it is risky for the patient"), but the second is uniquely human and valuable. Gaining information about the future (prediction) is not enough, since decisions entail value judgement and analysis / taking of risk. This means, that the value of human judgement - considered to be a "complementary good" for the commodity of prediction goes up.

This latter will be the key, the relationship between prediction and judgement in the insights of the next book of the authors.

3.9.4 Power and prediction: going beyond the basic theory

As an honest and laudable starting point: the authors admit, that they were wrong. They predicted the rapid spread of AI and profound transformation of nearly all business domains. They were wrong, in the sense that AI had great successes in given domains, and is obviously on the agenda of nearly all companies, but with all successes, it was more deployed in narrow niches, and did not usher in the great wave of change they predicted.

To elaborate on this, and in strong connection to a broader set of concepts about the *"life-cycle of innovation"* (see David, 1989), the authors introduce three concepts of technology application: *point solutions, applications and system solutions*.

Point solutions can be considered technologies that are replacing a single step in the current business process, making it more efficient without any modification to the process itself (doing the same drilling at the same spot of the factory, but instead of steam, with an electric drill), while applications are local solutions where the reorganization of single processes become feasible (a whole drilling / milling station electrified, doing things differently), though keeping the general system unchanged. System changes (like reorganizing the whole layout and processes of the factory to accommodate electric equipment all around) are profound changes of how we do things, having the greatest potential impact - albeit with the greatest investment.

With this in mind, it has to be noted, though, that the authors were not wrong in the sense that they misjudged technology or the main economic drivers, but in the sense of not considering other important factors, and were surprised to see, that point solutions are popping up everywhere, applications are rarer, system change is largely absent yet with

respect to AI. In a sense, their second book is dedicated to elaborate, why.

3.9.5 More detailed relationship between prediction and judgement

After introducing the dilemma of point solutions vs. system change, the authors recur to elaborate the concept of judgement and prediction more in detail. Their argument is, that in all cases of decisions, the two elements are necessarily present, albeit two more observations are made:

1. Decision is "costly", in the sense of cognitive load and responsibility (which by the way ties in well with the concept of "*decision fatigue*" - see Tierney, 2011 - developed in the psychology literature, partly also quoted by the authors), thus human beings must necessarily "economize" the usage of their judgement as a precious, constrained resource. (We can take in only so much information, and form only a handful of well thought out value judgements about what should be done as an action.) Because of this, there is a natural tendency to forego decisions, especially in the form of fixed rules that are helping to avoid decision situations but introducing *hidden uncertainty*, as well as rigidity into organizations' behavior.

On the other hand, the usage of rules is a key contributing factor for stability, since it allows for the efficient coordination of activities across an organization by foregoing the necessity for constant communication, reducing the coordination overhead.

In this sense, AI technologies - so prediction - is basically reintroducing the necessity for applying judgement, thus on the one hand causing extra cognitive load, as well as making coordination more complicated (by challenging the rules), but also holds a great promise, if "well oiled", so flexible processes can dominate the business, or if decision processes can be dynamically adjusted.

2. The separation of the aspects of "prediction" and "judgement", which were practiced together by humans previously is the second major effect the authors study. (It is worth noting, that this is not a mere conceptual construct from the authors, but a tenet in decision literature, see: "Forecasting is about estimating the likelihood of future events. Judgment is about deciding what those forecasts mean for us and what we should do in response" Tetlock and Gardner, 2015, p. 20.) With respect to this distinction, the authors intend to show, that this new division is even more challenging, since it shifts decision power from the people, organizational units or groups that had control over them previously. As they demonstrate it with several community and company examples, the expertise needed for jointly doing prediction and judgement exercised by experts is now "factored" into scalable machine prediction, and either a new set of experts specializing in judging the probabilistic predictions of AI models, or to committees that set up decision criteria (basically pre-judgements of value). In this latter case, judgement can also be scalable, but requires a complete overhaul of the responsibility and decision system.

It is worth noting, that in this light, much of the literature on organizational change, social

epistemology and decision making comes into play, so this represents a deep connection point to other literature that would definitely be worth exploring more in detail.

3.9.6 Evaluation of the book's contributions

With that in mind, let us now move to a more critical evaluation of the contributions of the book.

On the positive side, it is worth strongly emphasizing, that the book is filling a large gap in the understanding of AI's economic potential, and is especially useful for those practitioners that want to use it as a kind of thought framework for implementing AI solutions into organizations (this is a clearly intended audience of the book, demonstrated by the introduction of such conceptual aids as the "AI Systems Discovery Canvas" and so on.)

On the exceptionally positive side it is also clear, that the book is invaluable in providing a thoroughly thought-out model that can guide the inquiries into the economic impact of AI, also connecting it to the broader field of organizational behavior, and behavioral economics in general.

On the criticism side - maybe not completely fairly - one can mention the fact, that due to some recent events, namely the introduction, and subsequent rapid mass adoption of Large Language Model based solutions (notably OpenAI's ChatGPT and alike), some of the insights of the book became nearly instantly a little dated.

To be more specific: the book does not concern itself with the possible broad (and currently rapidly advancing) adoption of AI tools in the "business to customer" domain, so the whole slew of changes (in personal productivity, consumption of information, cognitive labour etc.) that are coming from the individual's side, and enabling new business models and huge competitive pressure on the labour market. One could argue, that this is just one (or a few) company's success with business model innovation enabled by AI, which is true, but gives less credit to these phenomena than they deserve.

Going forward: the strong push towards more generally applicable AI models ("foundational models" - see Bommasani et al., 2022b - would merit some new considerations. In many of the expositions of the book, AI solutions, even when system like, are focusing on (from the technological sense) narrow, domain specific AI capabilities, thus positing competitive advantage in a single domain for a well "oiled" "feedback flywheel". In practice, though, the training and deployment (and often even Open Source release) of generally capable AI solutions and agents (see Xi et al., 2023 and Durante et al., 2024) are questioning the defensibility of solutions, systems and companies that have a narrow AI solution, in stead foreshadowing, that the real advantage may lie in the hands of "AI first" organizations, that come from purely the domain of AI research and innovation, and intend to disrupt many business areas simultaneously. This is a kind of phenomena that the authors simply did not have the opportunity yet to study.

Another - maybe less severe - challenge towards the book's theories comes from the apparent ability and usage of current state-of-the-art Large Language Model based systems in planning and task decomposition (see Kojima et al., 2023), in which they are autonomously decomposing general goals given by humans into actionable tasks and carry out their execution with a broad sense of "self control". One could argue, that this case does not violate the premise of the book about human judgement, since still, the setting of goals and the evaluation of the final results is still in human hands, but none the less, this phenomena deeply alters the meaning, or the type of judgement that is necessary in these cases.

Also, it is worth noting that throughout the book, the topic of execution automation is not reflected upon clearly. If we take such passages as: "AI provides predictions, and hence, value is created by improved decision-making. Point solutions allow prediction to improve existing decisions." We can feel, that there is another aspect which is not sufficiently emphasized: much of automation, that is, automatic execution has a non negligible cognitive, hence decision element in it, (see Varela et al., 1991) thus: good quality prediction is not just an input for decisions, but also a vital prerequisite of automated execution! The real breakthrough might come when automated execution of tasks (sometimes in a digital, some times in an "embodied" physical modality) is perfected!

The complex problems and necessary skills with the widespread adoption of human value judgement based automated execution are still to be investigated, and loom large on the near horizon.

3.9.7 Takeaways and recommendation

The main dilemma of the book is extremely timely, the contributions of it in building the thinking framework invaluable.

The question is, from the business practice:

"Large companies rarely find it worth it to transform the way their industry operates, especially if their industry is currently profitable. The risk of getting it wrong is too high.", so while "system change" and "disruption" is the promise (and also the challenge), maybe incumbent organizations are at an inherent disadvantage...

If the hypothesis is true, that economic power does not disappear, just shifts, where will it shift to? Towards the organizations having the largest models? With AI training know how? With application know how? With data? With capital? With judgement surplus? With ones that build new economic / governance systems for old business problems? ("However, those who do not have a stake in the current system are often best poised to reap the reward from creating a new one.")

And where does the economic power of labour shift?

"As Christensen saw it, the problem was the velocity of history" - how can organizations speed up? By automation / computerization / agentification?

But how can individuals speed up??? Mental training? Using AI to speed up individual learning? With human self augmentation?

As the book argues: “Automation requires codifying judgment. A human must specify judgment when the machine is deployed, rather than upon receiving a prediction. That means that judgment has to be useful for a large number of decisions, and it has to be described in a way that can be coded. That isn’t so easy.” So maybe abstract, structured thinking is the key trait for human employability?

If the fact that the amount of judgement (if scales in pre-judgement) needed to operate our business processes require way smaller number of way higher qualified specialists than the number of people working now. Comparatively speaking judgement will take on a small amount of working time, thus, at current demand levels, it is not needed in such quantities to mandate mass employment.

Finding answers to these questions, and empirically validate the claims of the theoretical framework proposed is thus an interesting area to work on.

Chapter 4

INTERVIEWS

4.1 Topic and motivation

In this chapter, based on the previously identified research directions, hypotheses and theoretical framework I present the results of a quantitative and qualitative analysis of AI's impact on the jobs of senior programmers. (This chapter was published in Szabados, 2025.)

4.2 Literature Context

Based on my investigations of economic literature presented before (and published in Szabados, 2024a) two important research trends are to be highlighted: firstly, in the field of technological adaptation, the framework of Agrawal et al., 2022, which sharply distinguishes between technological possibilities and organizational realities that limit their practical application. As we will see in later chapters, the experiences I outline strongly confirm the authors' theses regarding the delay in adaptation (for a more detailed review of the book, see the previous chapter or: Szabados, 2024b). Secondly, in terms of the impact of AI on labor economics, particularly from a quantitative perspective, the work of E. Felten et al., 2021b is considered a prominent research direction, in which they assign a specific metric called the "AI occupational exposure score" to various occupations, aiming to make the relative vulnerability of different professions comparable. The metric is based on a classification of the skills defined by the US Department of Labour Statistics for each profession (see e.g.: U.S. Department of Labor, Employment and Training Administration, 2024) in terms of their substitutability by AI solutions, and then aggregates these to create a unified indicator of the automation potential of a given profession. In this light, it seems particularly noteworthy that, somewhat contrary to our basic intuitions, the authors feel that "white-collar" tasks are more automatable by modern AI systems compared to "blue-collar", physically demanding (unstructured) jobs. In this regard, the position of programmers on the list created by Felten and colleagues is particularly noteworthy, as those working in the software field received a very high AIOE score between 1.201 and 1.283 (depending on their exact job), which is

very close to the maximum of 1.528 (top 15%), indicating that, for example, dancers (-2.67) are much less likely to be affected by any changes due to AI technologies. This strongly motivates my investigation into whether 1. the adaptation of technology has already begun in the context of Hungarian programmers, and at what stage it is on the path described by Agrawal et al., 2022, 2. whether the exposure is indeed as strong as implied by E. Felten et al., 2021b's AIOE value.

4.3 Definition of Research Topic

The subject of my investigation is the application of AI solutions, as understood above, in the field of software development work, touching on all its aspects. While it is obvious that the central element of this work is the creation of new software source code (or simply "code" in professional parlance), it is important to emphasize that a software developer's work also integrally and indispensably includes testing the code, documenting it, setting up its runtime environment, as well as communicating with "clients", i.e., those who articulate demands towards the development process, and other developers, just as much as the professional textual recording of requirements ("specification"). Since my respondents come from a highly heterogeneous organizational environment (from large corporations to freelancers), separate specializations (software development, testing, business analysis, data science, etc.) are often present in mixed roles (see Appendix C), so I generally speak of "programmers", examining the impact of AI solutions on any element of software development work, including all these aspects.

A further narrowing of the topic is that my analysis mainly focuses on "senior" software developers. In my definition of "seniority", a person is considered senior if they have been engaged in programming activities in a full-time, work-related capacity for at least three years. I took this definition into account both in filtering quantitative data and in selecting interviewees. Statements about "junior" software developers (i.e., those who have been practicing their profession for less than three years) are also recorded from the perspective of senior interviewees, reflecting their opinions.

4.4 Quantitative Background

Although my analysis is primarily qualitative in nature, many insights can be contextualized by using conclusions derived from quantitative data as a kind of thinking framework. Since AI solutions themselves are software products, it is not surprising that key players in the software development world, including the Stack Overflow Inc., 2024 internet portal, which plays a central role in the profession, pay strong attention to the impact of AI solutions on the field. The Stack Overflow portal itself, among other content, operates primarily as a professional question-and-answer community site, holding a dominant position in developer work

due to the high-quality community communication that takes place on it. Posting on the site and providing good quality answers to questions is, in some sense, a matter of professional prestige, so the site can be considered one of the focal points of the IT profession's public sphere. In this spirit, Stack Overflow Inc., leveraging its position as a central professional forum visited by developers practically every day and an indispensable platform for mutual assistance, conducts thorough questionnaire research among its visitors on various current topics every year. Their 2023 survey, covering about 90,000 respondents Stack Overflow, 2023a, in this spirit, devoted a significant portion to questions related to AI technologies, providing an in-depth analysis Stack Overflow, 2023b of the topic. The data is publicly accessible, so I also processed it to provide a broader context for the lessons of the in-depth interviews. It is noteworthy that the same data source has already been processed by Kovács and Vastag, 2023 and they found it to yield significant information regarding the income differences of Hungarian programmers, reinforcing confidence in the relevance of the dataset.

4.4.1 General Demographics of Respondents

The Stack Overflow survey, due to its impressive number of participants (>90k), has significant descriptive power, but it should be noted that since it is a voluntary completion, its distribution may not be fully representative of the entire examined population, essentially all programmers worldwide. Nevertheless, due to Stack Overflow's central role, it can be considered a highly relevant sample. It is worth highlighting that the questionnaire is voluntary, so mainly those programmers who are proud of their profession and motivated take the time to complete it. This may possibly be accompanied by a kind of professional openness, even towards new technologies, although we do not have direct data on this. In the following, I always focus on the subgroup of "senior" developers from the entire data.




Figure 4.1: Age distribution of respondents (n=57240)

Image Source: Author's illustration

In terms of age (Figure 4.1), respondents are noticeably from the younger age group, but

this may easily coincide with the overall age distribution of the programming field (see also: Developer Nation, 2024).




Figure 4.2: Region distribution of respondents (n=57240)

Image Source: Author's illustration

In terms of geographical region (Figure 4.2), we see a strong emphasis on "western" areas, with Western Europe being particularly dominant and, together with Eastern Europe, representing a greater weight than North America. Since North America (specifically the USA) is a dominant force in the software development field, and Asia is also significant (for example, the most populous area in terms of mobile developers, see also: Statista Research Department, 2016), the data is not entirely representative from this perspective, likely slightly overrepresenting Europe.




Figure 4.3: Organizational size distribution of respondents (n=57240)

Image Source: Author's illustration

Regarding "organizational size" (Figure 4.3), or the number of employees in the economic organization serving as the context for the programmer's work, we see a much more

balanced picture, with the data significantly including individual ("freelancer") and large organizational contexts (for freelancers, for example, an approximately 10% ratio can be assumed in the USA, see also: U.S. Bureau of Labor Statistics, 2023 and Zippia, 2024, so the global 6.6% value is not unrealistic).

4.4.2 Current AI Usage

The main question the survey sought to answer was to what extent AI-based tools have become integrated into the everyday practice of software development by 2023, following their "explosion" in the 2020-2022 period.




Figure 4.4: Current AI usage of respondents (n=57240)

Image Source: Author's illustration

From this perspective, it is a significant result that nearly 40% (sic!) of respondents (Figure 4.4) stated that they are currently actively using AI-based solutions in their work, and more than a quarter plan to do so, with only 32% believing that there is no place for AI solutions in their work. It can be stated that the vast majority of professionals have integrated AI, making it particularly interesting to see what opinions they have about the impact of these tools on work in such a high "AI saturation" area (the responses can be considered significant, as the Chi-square statistical value associated with the question – 1311.22 – has an empirical significance level that is negligible - $p < 0.0000000001$). Therefore, we can reject the null hypothesis that all categories occur in the same proportion, meaning the differences observed in the sample are significant.

It is noteworthy that there is indeed an observable age effect in the use of AI by professionals (Figure 4.5), although its strength is weak (Cramér's V coefficient: $\varphi_c = 0.1041$), as it is visibly more likely that higher AI usage (or willingness to use) is observed in younger age groups. This seems to somewhat reinforce the (often strongly stereotyped) relationship between age and the adaptation of new technologies.




Figure 4.5: Current AI usage by age

Image Source: Author's illustration




Figure 4.6: Current AI usage by region

Image Source: Author's illustration

However, the distribution by regions shows less stereotypical results (Figure 4.6) (although the effect is also weak here, Cramér's V coefficient: $\varphi_c = 0.1069$): the higher AI adaptation characteristic of African areas, due to the relatively low proportion of respondents, may be a result of self-selection (more innovative, open workforce uses Stack Overflow, thus responding to the survey), but we can suspect a kind of technological "leapfrogging" (for example: Lee, 2021) in the background, where a developing area skips elements of the developed ones' previous path and directly uses the more mature technology. In this respect, Africa, South America, and Asia's openness towards AI adoption can be a significant advantage, but more relevant to us is that Eastern Europe's relatively higher usage rate also holds potential. Therefore, we expect to observe a kind of more open attitude during the qualitative analysis as well.




Figure 4.7: Current AI usage by organizational size

Image Source: Author's illustration

It is also an interesting observation that organizational size (Figure 4.7) influences the adaptation of AI solutions, albeit to a very slight extent (Cramér's V coefficient: $\varphi_c = 0.0721$). Ideally, in organizations with more than one person (where the possibility of group learning is already meaningfully present), slightly higher adaptation can be observed, which continuously decreases in proportion to the size of the organization (and thus the expected level of regulation) (this effect, however, seems weaker than the previous ones based on Cramér's V). During the interviews, we will attempt to explore whether there is indeed such a correlation between organizational size and adaptation.

4.4.3 Prospects

Perhaps the most important element of the Stack Overflow survey concerns developers' expectations. In Figure 4.8, I attempted to summarize the opinions of senior developers on which areas they expect the most change in the upcoming period.




Figure 4.8: How will your work change in a year?

Image Source: Author's illustration

Although it is visible (Figure 9) that every area seems significantly affected, according to seniors' expectations, the main process of writing code will undergo the most significant change, thus being exposed to potential efficiency gains. Additionally, the planning/specification phase of projects, and even learning about different technologies, receive significant emphasis. The latter is particularly interesting, as it constitutes a frequently underemphasized element of developer work, where the professional must invest significant time in getting to know a new solution or tool before even starting to solve a problem.

In the qualitative examination, I directed attention to these areas with specific questions regarding their daily workflows, time requirements, and automation potential, as well as attempted to inquire about expectations (for the next maximum three-year period). It is worth noting that while further quantitative insights could be gained by filtering the Stack Overflow data to Hungarian respondents only, unfortunately, among the respondents, a total of 231 Hungarian senior programmers provided information on any question in this topic, and their response distribution is quite fragmented (see Appendix D), so we cannot

draw significant quantitative conclusions. Accordingly, I will explore the Hungarian results along the lines of qualitative analysis.

4.5 Qualitative Examination

4.5.1 Method

Taking into account the broader context above, I sought to answer, using in-depth interview methods, how Hungarian programmers' practical experiences with the application of AI solutions compare to the results of the Stack Overflow survey, and how their prospects – particularly their expectations regarding changes in the labor market – are influenced by the spread of AI technologies.

During my investigation, I conducted **20 in-depth interviews** to cover the above general age distribution (see Appendix Appendix C), thus I specifically interviewed Hungarian senior programmers in 40-60 minute individual online interviews about their insights on the topic. The selection of interviewees was based on professional groups ("meetup") and collegial relationships, using a "snowball" method, through recommendations, with the selection criteria being compliance with the age distribution in addition to seniority (Figure 4.9 compared to Figure 4.1).




Figure 4.9: Age distribution of interview subjects (n=20)

Image Source: Author's illustration

As Figure 10 shows, my respondents fit well with the distribution of demographic variables presented in the previous chapter. Respondents included those under 30 and over 60, as well as those working in large organizations (over 1000 employees) and individual "freelancers" (see Appendix C). Despite all my respondents being of Hungarian origin and the interviews conducted in Hungarian, due to the nature of the software market, both their place of residence (Spain, Germany, Austria, Switzerland, Hungary) and the organizations

or clients employing them resulted in a mix that included both Eastern and Western Europe, which is actually characteristic of the industry, especially in the post-COVID period, where remote work is significant.

Since my respondents considered their statements about their work to be sensitive from a competitive market perspective, I only process their responses in anonymized and aggregated form.

4.5.2 Interview Structure

The interviews conducted during the research followed a comprehensive and detailed outline (see Appendix B), aiming to deeply examine the integration of artificial intelligence solutions in the workplace and their impact on the labor market and programmers' careers. In the introductory section of the interviews, we discussed the participant's background in entering the programming field and their current role, as well as their work environment, focusing on the methodological and technological maturity of the environment. The main section explored the integration of AI into daily work, the potentials of automation, the analysis of daily schedules and tasks, and the risks and problems associated with AI application. This was followed by a discussion on the impact on careers and the labor market, while the closing section provided an opportunity for clarification, sharing further insights, and concluding thoughts. The interview outline was designed to provide a thorough insight into the current state and future prospects of AI technology application in the workplace, deliberately not addressing its broader societal impacts.

4.5.3 General Observations

4.5.3.1 Factors Behind Openness/Application Level

Although a significant portion of the interviewees reported that a kind of career orientation developed quite early (for some in childhood), leading most to participate in specialized higher education, several (6) stories emerged during the conversations showing sharp career shifts (e.g., from teaching or managerial positions) towards programming. One might assume that such strong transitions influence the degree/speed of openness towards new technology, but both among those who "inherently" chose the career and among the "career changers", higher levels of experience in applying AI technologies could be observed, as well as lower levels among those who did not. It seems that it is not a deeper attitudinal difference, but rather the level of professional interest and the early investment of effort in testing technologies that are the main determining factors in the level of technological adaptation observed in a given professional.

It can also be established that in cases where the level of use is not high, this is not linked to a lack of early experience. Professionals who do not apply AI solutions in their

everyday work also reported experiences of trying/testing (e.g., in the form of a "hackathon" - see Lexikon, 2024 - type event), but on the one hand, they report mixed impressions ("it doesn't work independently yet"). It should be noted that in framing their experiences, the "autonomous" application of AI solutions was much more present, which did not show satisfactory performance, rather than a kind of "assistance" type of use. This could easily lead to these professionals not making the significant effort – which all speakers consistently reported – to make the technology "usable", customize it, and integrate it into their individual workflows.

A strikingly consistent position among the respondents was that, in addition to shaping the technology, individual learning, trying, and experimenting were dominant aspects. They made it clear that it is not a kind of given, fixed source (e.g., curriculum) knowledge that is needed, but experiential, experimental knowledge that the individual acquires through open and unstructured experimentation with the technology. The extent of this experimentation and the knowledge of the exact limits of the technology, as well as the diversity and efficiency of its application, were – unsurprisingly – closely correlated.

Another, although not dominant or necessary, facilitating factor was if the given professional had previously gained experience with an earlier generation of AI technologies during their career or studies. If they moved in a related field (data engineer, machine learning engineer roles), the adoption of the latest technologies was almost a natural, expected step for them, but even those who moved entirely in "traditional" roles (e.g., "backend" engineer) and had no prior AI exposure were also present in significant numbers among the early experimenters. Overall, it can be said that it is a facilitating but by no means necessary factor. (For example, there was a specific "data scientist" specialist who still chose the lower usage level because, in their opinion, validating the solutions proposed by AI systems requires a lot of effort.)

A factor hindering application was the issue of data protection. Several expressed concerns that corporate protected information could leak or be accessed by unauthorized "third parties" with the currently dominant solutions. Two solutions are conceivable to address these problems: the creation of stricter corporate policies and the development of secure solutions (e.g., dedicated corporate subscriptions). However, several other, more innovative speakers highlighted that they were "lucky" that the development of corporate guidelines took place slowly because it allowed them to experiment freely (and, in their opinion, with full consideration of data protection aspects), gaining the essential experiences that now make them routine users. This highlights the innovation-inhibiting effect of overly bureaucratic corporate guidelines and emphasizes that ideally, companies themselves proactively provide the opportunity for free experimentation in a protected environment.

4.5.4 Basic Findings

The interviews revealed a wide range of opinions on the impact of AI-driven automation on the employment of junior and senior programmers. There was a high degree of agreement among participants on certain core themes, but contradictions occasionally arose.

Two products/services, "Github Copilot" and "ChatGPT", were mentioned by almost every speaker, with their experiences largely revolving around these two solutions. Among those with more serious experience, several showed strong commitment to these tools. "I would pay for these subscriptions before a bus pass," said one of the most senior developers, working in a very closed, financial environment, indicating that they had to fight to even access these tools in a work environment.

4.5.4.1 Efficiency Increase

All interviewed software developers agreed that the application of AI technologies brings small or large (according to many, fundamental) changes to the work of software developers. Even the most pessimistic interviewees highlighted the possibility of **efficiency increase**, while more experienced ones noted its fact and estimated **10-30%** magnitude. Accordingly, it can be concluded that there is a **consensus among professionals regarding the relevance of the technology**. This strongly confirms the findings of other studies, such as Dell'Acqua et al., 2023, who measured a 43% performance increase for junior "knowledge workers" and a 17% increase for seniors in knowledge-intensive activities, highlighting the fact that the automation of junior-type activities is more feasible with the help of modern large language model technology. Additionally, the quantitative results (Stack Overflow questionnaire) can also be well supported by the picture emerging from the interviews: a smaller proportion of developers, roughly a third, do not have significant experience in the professional application of AI technologies and do not feel it particularly emphasized, while two-thirds, with varying degrees of experience, are open and agree that significant impact is expected.

4.5.4.2 Outsourcing Monotonous Tasks

There is a strong consensus that, at their current level of knowledge, AI systems are much more effective at replacing more monotonous and schematic tasks in a software project, primarily by writing and documenting a specific module or even a "function" within it. This was typically framed as a lower value-added, "junior" task, but several added that in terms of quantity, such tasks can also make up a substantial part of a senior's work, so the currently observed efficiency increase mainly results from the automation of these tasks.

It is important to highlight the limitation mentioned by almost every respondent: current systems are only limitedly capable of creating complete projects or comprehensive solutions, so human supervision is always necessary. Opinions were more divided on whether a significant advancement in this AI capability is expected in the near future.

An interesting statement was, for example: "I can precisely specify what needs to be done, direct the AI's attention to how I want to solve it. I converse with it, enjoy it, use an iterative development process." Reading between the lines, on the one hand, it is strongly felt that from a senior's perspective, they can easily judge when a system makes a mistake, as they have a clear picture of what they want, and on the other hand, it is clearly visible that outsourcing more monotonous work elements makes the work atmosphere more enjoyable. Behind this is obviously the sense of control from seniority and the sense of efficiency from AI solutions.

Related to this, a fundamental and constantly recurring insight was that the errors or "hallucinations" (see more in L. Huang et al., 2023) made by AI systems require two important skills for detection and correction, both of which deserve discussion in themselves.

4.5.4.3 Necessary Skills

One of the skills is essentially seniority itself, or **experience**. Almost every respondent with AI system experience stated that they rely on their professional experience and insight when judging whether a solution generated by AI "looks good", i.e., meets professional requirements, "does what we want it to do". This ability (being seniors) was given among the conversation partners, but several expressed doubts about how a junior, who "doesn't have an eye" for such things, would cope with this task (see below: junior paradox).

Several felt that the task of the programmer of the very near future will no longer be writing code, but much more clarifying requirements and judging the result, ensuring quality. This deeply resonates with Agrawal et al., 2018's theory that the power of AI solutions lies in cheap "prediction" (in this case, predicting "what should be written here in code"), but as a complementary good, the demand for human "judgment" will increase significantly.

Another significant skill was – surprisingly – a kind of doubt or **reflection**, which several mentioned as a crucial element for effectively working with AI systems as a programmer (or more broadly as an informed, not misinformed, independent citizen). Since the possibility of "hallucination" always exists, almost without exception, every speaker saw the greatest danger in "overconfidence" and "lack of reflection", where we simply surrender ourselves to AI solutions without critically evaluating their responses. Several also spoke of a kind of "crutch effect", where we become too accustomed to using a mental aid (in this case, AI), and thus lose our independent abilities.

4.5.4.4 Other Areas of Application

Two other areas highlighted by the majority of speakers well reflect the conclusions of the Stack Overflow survey: the AI-based automation of writing software tests and the use of AI solutions for learning purposes.

Regarding testing, several mentioned that low-level software component tests are them-

selves source codes, and a significant portion of them are quite schematic, making them particularly suitable targets for AI-based automation. Quickly and efficiently writing the "framework" of tests alone can significantly shorten workflows (see also: Naimi et al., 2024), as one respondent noted: only the more interesting, "exotic" test cases need to be written by hand. However, another respondent specifically drew attention to the fact that if both the code and the test are written by a machine system, the likelihood of catastrophic errors increases, so in their opinion, either one or the other should be written by a human. A respondent working as a lead programmer instead sees the usual social quality assurance processes ("code review"/"peer review") as effective, again highlighting the importance of reflection and judgment. It should be noted that in the field of testing, the Stack Overflow survey measured lower usage intensity than what can be inferred from the statements of current respondents (see also Quantitative Background – Prospects section above).

In terms of "learning together", it is crucial to highlight that practically every speaker confirmed the possibility of using AI solutions as a kind of "oracle" (a term used in one interviewee's workplace group) or "trainer" to learn much faster (this aligns with Petrovska et al., 2024's findings, who examined the integration of "generative AI" solutions into software engineering education). This can be good news in terms of counterbalancing the increased learning demand caused by accelerated technological innovation affecting programmers (and can help solve the junior paradox as well). However, several emphasized that, on the one hand, learning together with AI systems is also a skill to be practiced, and on the other hand, filtering out potential hallucinations in responses intended to aid learning is also a serious human task (again requiring experience and reflection).

Interestingly, even when asked directly, several had negative opinions about the area of client communication and specification writing, which they still consider a dominantly human area – not confirming the strong emphasis observed in the Stack Overflow survey (the discrepancy remains even if the Stack Overflow survey is filtered to Hungarian respondents only). Only a freelancer raised the possibility that communication with clients and the proper "facilitation" of specifications from clients could also be an AI-supported area.

4.5.5 Impact on Employment

Asking a bit more broadly and shifting from individual work to the organizational/economic level, the following conclusions could be drawn:

Uncertainty and Negative Consequences: There is uncertainty regarding the impact of artificial intelligence, with some expressing concerns about the potential for AI to automate jobs, particularly those of junior developers, and the possibility of increased unemployment in the programming sector due to the reduced demand for human programmers in light of automation tools and no-code platforms.

Potential for Efficiency and Competitive Advantage: Others (and often the same in-

dividuals as above) see AI-driven automation as a promising development that could lead to more efficient processes, provide a competitive advantage to those who integrate it into their workflows early, and potentially significantly change the programming field. This view suggests that while the demand for developers may decrease, the efficiency and effectiveness of project completion would increase. Several added that it seems the application of AI will soon become an essential foundation that competitive organizations simply cannot ignore.

Regarding productivity growth, however, one of the interviewees in a leadership position drew attention to an interesting phenomenon: if the use of AI tools is not part of the corporate "policy" (i.e., not centrally introduced, supported, and transparent) in an organization, then in a peculiar way, employees take the initiative themselves and start using AI tools in their work. This raises several issues: on the one hand, the phenomenon of "shadow IT", as Ogedengbe et al., 2024 reinforces, i.e., the use of IT tools without the knowledge and approval of decision-makers, which can pose serious cybersecurity and data protection problems, and on the other hand, it can also reveal a kind of conflict of interest. If the system of workplace incentives is not appropriate, then the efficiency increase of employees does not appear in the overall performance of the organization, as employees "keep the freed-up time and capacity for themselves". This highlights the challenge that requires the earliest and clearest AI strategy development from organizations.

Shift Towards More Abstract Work: There is a strong belief that AI-driven automation will lead to job transformation, and programmers will need to focus more on business processes and abstract tasks rather than low-level coding. The work essentially "abstracts". This includes a shift in the talent market towards less coding and more use of AI tools. One interviewee put it that this has always been the case in the history of IT, as they themselves do not know exactly how the computer works at the electronic level, they only formulate tasks on an "abstract" language layer (programming language). In this sense, AI tools only provide the possibility of formulating in a more human-friendly, natural language. The opinions of my interviewees strongly align with Agrawal et al., 2022's view that the basic unit of business, practical life is not the isolated solution, but the "system", which consists of the joint operation of several solutions and is based on the deterministic operation of solutions. AI-based solutions will only have a path to wider adoption if they can reliably integrate with other solutions in this sense, forming a "system".

4.5.6 Junior vs. Senior Programmers

Decrease in Junior Programmers' Employment (the "Junior Paradox"): Many believe that AI-based automation will unfavorably affect junior programmers, as the replacement of tasks typically performed by them will result in higher unemployment for them, and companies will prefer to hire experienced programmers skilled in handling complex tasks. Most pointed out the paradoxical situation that becoming a senior programmer can only be

achieved by passing through the junior programmer status, but due to automation, junior positions become scarce. This places a significant burden on higher education, as much more "mature" students need to be released into the market if we want to see a realistic chance of their employment (as mentioned earlier, this fits well with Dell'Acqua et al., 2023's findings).

Several speakers also raised the possibility that learning programming and during the junior level, workers will be forced to acquire sufficient proficiency in the form of quasi self-education or "hobby", i.e., unpaid work, until they can apply for a senior position. One even stated that this phenomenon endangers the sustainability of the entire field, as the replacement of senior professionals is not ensured, so some form of intervention or investment will be necessary, either from organizations or the state.

However, other respondents were much more optimistic, highlighting the possibility of "increasing abstraction", where today very few, possibly maximally senior programmers are forced to use "low-level" languages or concepts close to machine code. AI solutions will only represent an additional level of abstraction, and juniors can start learning this level right away. (However, other speakers saw serious dangers and "general skill decline" in such a leap over the "basics".) Based on these, several outlined a kind of "junior paradox".

Opportunities for Senior Programmers: Older programmers are considered better prepared to effectively use AI systems and lead projects, meaning their employment prospects will not be as negatively affected. Their experience and deeper understanding of the field may enable them to navigate the changes brought about by artificial intelligence more successfully.

Critical Skills and Continuous Learning: The necessity of critical skills such as systems thinking, analytical skills, and a strong understanding of business processes is emphasized for all programmers. Continuous learning and adaptation to new technologies are highlighted as essential for maintaining employability in the face of AI-driven changes. Fortunately, as we have seen, AI systems themselves can help with this.

4.5.7 Views on the Role of AI

Concerns About Reliability and Quality: There are concerns about the quality of work and the reliability of AI-generated code, with some fearing that automation could lead to lower quality outcomes, more project failures, and a lack of critical thinking and work verification, but there was also a counter-opinion: well-functioning automation could lead to more stable, reliable solutions according to some. The issues of quality concerns and "over-reliance" fit well with Dell'Acqua et al. (2023)'s findings, who demonstrated that in tasks where current AI systems are still operating at the limits of their performance, and their users do not reflect on this, merely accepting their solutions, they may even reach significantly worse outcomes than if they had not applied any AI solution at all. Therefore, a critical attitude and

proficiency are particularly essential.

Optimism About Support and Learning: Several are optimistic about the educational and guiding role of artificial intelligence for programmers, helping them learn and develop their skills, as well as increasing work efficiency and project completion rates.

4.5.8 Summary

As a fundamental statement, it is acceptable that there is almost complete consensus in the profession regarding the possibility of a significant, at least 10-30% performance increase, confirming both Dell'Acqua et al. (2023)'s results and Felten, Raj & Seamans (2021)'s earlier estimates in terms of the significant exposure of programmers' work to modern AI-based automation. Everything else can be interpreted based on this. If not in full magnitude, but in terms of significance, this aligns with several measurements, such as the conclusions drawn from Google's own methodology-based estimates - see DeepMind, 2024.

Additionally, it can be inferred from the statements that while the market value of junior software development work is expected to decrease, the value of senior work (after a smaller temporary decrease, perhaps) will increase in the long term, so we can speak much more of a kind of polarization, which aligns well with E. W. Felten et al., 2019b's description.

In terms of consequences, while there are concerns about the potential increase in unemployment, particularly among beginner programmers, and the impact on work quality, there is also a strong belief in the possibilities of efficiency, learning, and the transformation of jobs towards more abstract and conceptual work driven by artificial intelligence. Continuous learning, adaptation, and a focus on high-level skills are key strategies for programmers to navigate the changing environment, and experience and reflection become fundamental values.

It is worth noting that the above conclusions receive explicit reinforcement when considering the opinions of those creating AI tools (e.g., the most frequently mentioned and used Github Copilot), as the head of Github Scheffler, 2023 also sees the future of programming work moving towards similar areas as my interviewees.

A recurring dilemma, however, is that if the efficiency of development work increases significantly, i.e., more software can be created with the same number of people, or the same performance can be achieved with fewer employees, which scenario will occur. In the first case, a significant boom in the software business area is expected, while in the other case, a wave of layoffs that seriously suppresses employment is possible, even if it increases profitability.

One decision point that most professionals mentioned multiple times is the following: if the "order flow" or "innovation potential" is healthy in a given business, then these organizations are more likely to aim for increased production rather than reduced employment, as "there is still so much to do" (especially since the unit cost of programming tasks decreases,

allowing new clients, e.g., SMEs, to enter the market with realistic demand, for whom "custom software" was previously too expensive).

However, if there is no significant demand for innovation in the given company, or if there is no significant potential for market expansion ("we've already written all the software we needed") in a given product/service area, then while the broader introduction of AI solutions may have a favorable impact on the future profitability of the organization, it will only come at the cost of significant employment reduction.

The above question, whether the efficiency increase can lead to market expansion with maintained workforce or even serious unemployment, aligns with Bessen, 2019's model, who modeled the impact of previous waves of AI technologies on the manufacturing sector. His model, which aims to address the above two scenarios uniformly, emphasizes the elasticity of demand, i.e., it points out – as implied above – that a key question when automating a given industry is whether there is such unmet demand in the background that can keep pace with the supply increase caused by efficiency growth. The key question, therefore, is what our expectations are regarding the demand for software solutions. Do we think that we want to perform a finite amount of tasks with software – thus demand can be easily satisfied – or do other economic effects prevail in the software field? I share von Engelhardt, 2008's view in *The Economic Properties of Software*, which states that a "network effect" can be observed in the software field, i.e., the more software we use, the more valuable newly created software products can be, simply by enriching the functionality and thus the usefulness of the existing ecosystem (unfortunately, there is currently very little literature available to provide a definitive assessment of whether it is indeed elastic or inelastic demand, for an overview of the question, see also Peng et al., 2023. Additionally, the opinions of interviewees were divided on the issue, perhaps with a slightly higher proportion – 12:20 – being present in "optimism"). Based on this, in my opinion, we should not be so concerned about the finiteness of demand, but rather strongly support technological adaptation so that businesses can adapt with sufficient dynamism, and the "cultural lag" (see Laato et al., 2023) does not become a dominant hindering force. The "cultural lag" theory describes that technological and material culture, such as software and digital technologies, develop faster than non-material culture, which includes social norms, beliefs, and organizational practices; this dissonance results in society and businesses having difficulty adapting to new technologies, as the skills, knowledge, and attitudes necessary for the introduction of new technological developments develop more slowly, causing significant delays in effective adaptation and the full potential of technological innovations.

Therefore, it is interesting that during my interviews, there was a (strongly illustrating the above "cultural lag") opinion that due to the insufficient number of sufficiently innovative people and the delay or risk aversion observed on the part of businesses, the corporate adaptation of AI technologies is proceeding much more slowly than would be possible, giving people time to retrain themselves. The result of this is a kind of labor market stability, but

also the absence of growth potential. A calm before the storm.

For now.

4.5.9 The Speed of Spread - "In-Between Times"

The above tension, i.e., the existing strong potential, the performance increase reported by early users, but the inertia (due to lack of knowledge or institutional/personal professional inertia) mentioned by pioneer professionals, deserves special mention, as it fundamentally supports Agrawal et al., 2022's ideas. The authors' concept regarding any general-purpose technology – in this case, AI – is that even in the case of a new, radically more efficient technique, its spread can be divided into well-defined stages: initial experiments leading to "point solutions", an emerging phase characterized by the application of new complete solutions ("application solutions") implying the reorganization of certain processes, and then the full saturation phase, in which entire systems are transformed in light of the new technology ("system solutions"). The area between point solutions and the growth phase is referred to by the authors as "in-between times", characterized by the enormous technical potential alongside its frustrating underutilization, due to the high inertia of existing systems, but also their good functioning, as if there is no sufficient external pressure, investing in radical transformation simply does not pay off, as the sunk cost in the previous system is significant, it works satisfactorily, and investing in the adaptation of the newer technology is costly. This is currently true for AI systems, although the costs surprisingly do not stem from the technology (several interviewees highlighted that they can even pay for the use of AI tools themselves), but rather from organizational costs and the cost of acquiring knowledge. Almost all interviews described the phenomena of "in-between times". Nevertheless, the technological transition, and thus investment in it, is considered inevitable by the unanimous opinion of the respondents, so it is important to consider the questions arising from the timing of the investment. As we can infer from Bughin et al., 2018's simulation results, there can be significant differences in the "cash flow" implications of technology adaptation (where the authors use "cash flow" in the sense of investment in technology and its return). Their conclusion (which, although a 2018 estimate, predates the latest AI wave) is that considering a roughly 15-year spread time, companies investing in AI technology in the first 5-7 years can expect a 122% return by the end of the period, late adopters only a 10% profit, and those not adapting at all a 23% loss. Although their estimates are by no means limited to the IT sector or even the Central and Eastern European region, they are worth considering, as both in direction and magnitude, they align with the accounts of my interviewees. Based on these, it can be stated that exploiting the currently available potential is a fundamental interest, thus requiring action from a broader policy perspective.

4.5.10 Policy Recommendations

During conversations with professionals, several topics or strategies were mentioned repeatedly, which I formulate into policy recommendations below:

Resolving the Junior Paradox:

Paradoxically, due to the spread of higher levels of automation, acquiring professional experience is expected to require greater time and financial effort from beginners, so supporting targeted entry-level jobs (e.g., through tax policy tools) and creating longer and more deeply integrated training programs with professional practice may be essential for the growth of future generations of experienced professionals. This is a particularly important area for intervention, as without it, a kind of sustainability crisis may develop.

Supporting Start-Up Enterprises:

It seems that several statements confirm Agrawal et al., 2022's (and even Scheffler, 2023's) theory that established companies with existing processes and traditions have a significant innovation disadvantage compared to newly established companies with a "clean slate", as transforming existing processes and retraining their employees is a severely costly process. Accordingly, supporting newly established companies that inherently incorporate innovation can be a key development tool at the national economic level.

Programmatic Development of Skills:

Most respondents highlighted the fundamental importance of "interpersonal" (human-to-human, human-to-machine) communication skills, as well as the ability for "reflection" and the ability to absorb new information, so systematic training/development in the form of programs to support these "soft skills" should be considered. It is particularly noteworthy that the automation of programming tasks by AI solutions exerts additional pressure on specialized higher education systems, as they must produce even more "mature", "marketable" workforce if they wish to maintain the competitiveness of degrees. This competitiveness, as well as the "competence disharmony" between education and market requirements (see Nagyné Halász, 2023), was thoroughly examined in the context of Hungarian IT higher education, and the findings, which reflect the strong demand for "soft skills" (see also Galster et al., 2022) from market players must receive even more decisive emphasis in light of the automation phenomena discussed above.

Broader Information Dissemination:

Both individual career decisions and societal/political decisions surrounding technological regulation and the redistribution of economic goods require much broader information to be available regarding the nature, mechanisms, possibilities, and dangers of AI technology, so that responsible citizens can make informed decisions about both their individual and collective future, as Brauner et al., 2023's research also implies, strong misunderstandings and prejudices currently hinder enlightened public discourse.

Chapter 5

LARGE SCALE ANALYSIS OF ONLINE JOB MARKET DATA

5.1 Topic and motivation

In this chapter, based on the previously identified research directions, hypotheses and theoretical framework I present the results of a large scale quantitative analysis of 1.3 million online job advertisements posted online, targeting Hungary, to give a detailed analysis of possible measures for real life AI adoption. (This chapter is accepted for publication as a separate article, and is currently being published).

5.2 Methodological Considerations

The application of AI solutions in the world of work is the result of a market process, where on one side is the availability of technology (its price, the level of capital investment required for its use), and on the other side is its usefulness (for example, the productivity increase achievable through AI solutions). However, any widespread use also requires the presence of employees in an organization who are capable of effectively applying the technology, which can be strongly influenced by the presence of knowledge about AI-based job opportunities, i.e., the fact that an employee encounters job opportunities where AI capabilities are a significant element. To examine this - and perhaps as a "proxy" for AI application - we believe that analyzing the text of public job advertisements is very suitable. In our work, we analyzed more than 1.3 million online job advertisements relevant to Hungary (in terms of the location of work or targeting of the advertisement) between 2019 and 2024 using computational text analysis methods.

5.2.1 Data Source

The professional large-scale extraction of internet content has become a significant business area, so filtered and pre-processed job advertisements are also available in "turnkey" database form. Among these, we used the services of the *Jobspikr, 2024* company, which provided Hungarian, English, and German job advertisements collected in the aforementioned time frame and quantity, on which we could apply natural language processing tools.

The most important columns of the database are as follows:

- **crawl_timestamp**: the time of data extraction
- **url**: the web address of the job advertisement source.
- **job_title**: the title of the job.
- **category**: the category of the job advertisement.
- **company_name**: the name of the employer company.
- **city, state, country**: the geographical location of the job (city, state, country).
- **post_date**: the publication date of the job advertisement.
- **job_description**: the job description in free text form.
- **job_type**: the nature of the work (e.g., full-time, part-time).
- **salary_offered**: the offered salary, if available.
- **job_board**: the source of the job advertisement (e.g., job portal).
- **contact_email, contact_phone_number**: the contact details of the job advertisement publisher.
- **apply_url**: the link to apply.
- **logo_url**: the availability of the company logo.

Furthermore, the database contains several "inferred" columns generated by machine, such as *inferred_city*, *inferred_state*, *inferred_country*, *inferred_job_title*, which were produced from the text of the job advertisement or other metadata.

5.2.1.1 Used and Omitted Data

During the research, the following columns were found to be relevant and were used:

- **job_description**: the full text content of the job advertisements, which formed the basis of the analyses.
- **city**: the city name used to determine the geographical location.
- **post_date**: the publication date of the advertisement, which supported the time series analyses.

In contrast, columns with the *inferred* prefix – such as *inferred_city*, *inferred_job_title*, *inferred_salary* – and data describing the category, salary level, etc., were not used in the analysis. These data were based on automatic classifications, which seemed to have highly variable accuracy and could have resulted in excessive noise during data analysis. Although these columns have potential for further processing, they were not considered reliable for the current research purposes.

For sectoral classification, company names were first manually normalized (merging name variations and treating, for example, corporate groups as a single entity even if they have different subsidiaries), and then the TEÁOR 2025 classification system was applied, which was done with a separate program. This ensured that job advertisements were accurately classified into industry sectors, regardless of the accuracy of the data provided by JobsPikr.

5.2.2 Text Processing Methodology

Our analysis is based on processing the *job_description* text field. The text contained job advertisement texts largely stripped of HTML elements, which we primarily wanted to classify in terms of:

1. Whether it explicitly mentions AI or related terms in any language?
2. Whether its semantic content implies that the person filling the position works with AI solutions, either explicitly or implicitly?

While the first approach involves complex keyword-based matching and aims to capture any advertisement that mentions AI in some form, the latter involves deeper, meaning-based, *semantic processing* to classify advertisements based on whether the job itself involves AI-related tasks. The technical aspects of these two approaches are detailed below.

5.2.2.1 Processing Environment

Our entire analysis was conducted in a Python 3.10.13 based development environment, which allowed scalable processing of large input data using the Pandas framework, execution of language processing steps, and subsequent general and time series-specific analyses.

5.2.2.2 Processing with Regular Expressions

The most efficient way to perform complex keyword-based pattern matching is by using regular expressions (*Regular Expressions*, *RegEx*), based on the foundations laid by Kleene, 1951. RegEx expressions allow matching of keywords and keyword combinations (even in inflected forms). In this spirit, we created a regular expression that matches various linguistic variations of terms related to "AI," "machine learning," "computer vision," "natural language processing," "data science," and other AI-related terms, including English, German, and Hungarian languages, as well as some common misspellings, enabling binary (Yes/No) classification of topic mentions.

Regular expressions identify and classify textual data based on a predefined pattern. They examine a given text character by character and determine whether it contains expressions that match the specified pattern matching rules. RegEx engines support various character groups, metacharacters (e.g., `.` for any character, `*` for repetition), as well as alternatives and groups, allowing for the recognition of complex linguistic structures. In application, for example, AI-related keywords and their variations can be identified in a job advertisement text, enabling automated determination of the advertisement's connection to AI. However, due to the rigidity of rule-based matching, false positives can often occur, especially if a word or phrase has multiple meanings or if the mention is merely formal.

The exact regular expression used for classification can be found in Appendix Appendix F.

The regular expression can be efficiently applied to the entire dataset using the Pandas framework's `.apply()` solution.

5.2.2.3 Processing with LLM Methods

The significant disadvantage of the above approach is that it can give "false positive" signals in situations where the given job advertisement matches the regular expression merely due to other orthographic peculiarities (unfortunately, a completely accurate match cannot be provided due to the frequent use of the ai letter combination) or, worse, merely mentions AI solutions as a stylistic element (see below), or - more relevant to our topic - as part of corporate branding. (We will address this phenomenon later.)

Example: AI as a misleading stylistic element in a job advertisement

I thought about using ChatGPT to write a job ad for me and its answer was actually... Nothing. Because I decided not to use ChatGPT... Not everything can be driven by A.I., sometimes we need to speak to one another, human to human to create real connections, even in business...

We are currently looking for someone who can lend their communication and interpersonal skills to our MSP admin team to work closely with candidates and clients to build meaningful and fruitful relationships...

Nothing Artificial about this role ;) At XXX we believe working as part of great teams is the key driver to our success! Our amazing teams provide solutions and expertise that help companies position for growth, execute on strategy, and improve business agility.

Work with us and you'll be joining the market leader in Recruitment Outsourcing.

So if you are someone who: Is experienced in an administrative role or commercial support environment, preferably in staffing or recruitment...

Accordingly, deeper, meaning-based, *semantic processing* is needed for content-based classification of job advertisements. Considering the scaling challenges posed by over 1.3 million data points, human semantic evaluation is not feasible, leaving only a machine-based approach. Approaching the problem as a binary classification task, we could mobilize the full toolkit of machine learning and likely achieve acceptable results with some form of pre-trained and fine-tuned neural network system used for semantic classification on our data. However, the widespread adoption of LLM models allows for their effective and scalable use as semantic classifiers even in a *zero-shot* sense, as demonstrated by Wang et al., 2023. This aligns well with the trend outlined by Ziems et al., 2024, which shows the effective use of LLMs for social science research purposes.

In such cases, we would only need to write an appropriate instruction, i.e., a *prompt* (for the origin of the term, see also McCann et al., 2018), or perhaps provide multiple examples using the *few-shot* method to help maximize performance, as demonstrated by Liu et al., 2024. The problem-solving ability of LLMs largely depends on how we provide instructions to them, i.e., how we prompt them. Designing the appropriate prompt is crucial, as LLM models, due to their generative nature, do not necessarily respond in a structured manner but often elaborate on their answers, ask follow-up questions, or provide explanations. Therefore, in a binary classification task, it is particularly important that the prompt clearly defines the desired output: the model should respond with only one of the predetermined two categories (e.g., "True" or "False") without any additional explanation. Several methods exist to enforce structured responses, such as explicit instructions on the response format, e.g., "Please respond with only 'True' or 'False' and do not provide further explanation!" or similar instructions that increase model consistency.

In the binary classification task, we need to set the LLM prompt so that it can distinguish between job advertisements that are actually related to AI in content and those that only mention it superficially or are completely irrelevant. This can be done in a completely "zero-shot" manner, i.e., without showing examples, or with "few-shot" learning, where a few examples are provided to better tune the model to the task. The prompt can be fine-tuned by directly including clear instructions on what the model should look for in the text. For example, we can specify that only advertisements that explicitly mention machine learning, deep learning, data science, or other AI-related keywords should be classified into the AI-related category. Alternatively, the LLM can be used for semantic classification, where it not only examines explicit keyword matches but also considers the content and context of the given text.

Interpreting the output of the LLM model is also crucial, as the output can often contain variability. For example, the model may generate formats like "True." or "True!", which are easily manageable, but it may also provide responses that do not clearly meet the expected binary classification criteria, such as: "This advertisement is likely related to AI, so yes." To avoid this, in addition to the clarity of the prompt formulation, post-processing steps are often necessary to ensure that the output is processed in the correct format. This can be done using regular expressions or other text analysis methods, or by using the "structured output" option available in the application interface of the given provider (in our case, OpenAI).

To avoid the uncertainties of LLM-based classifiers, we opted for a stricter approach: we manually created a training and test dataset (fortunately small in size due to the nature of the solution), and then used automatic prompt optimization procedures to create a highly efficient LLM-based classifier prompt. With the optimized prompt, we ensured that the model not only performed accurate binary labeling but also achieved nearly flawless performance on the manually validated test dataset. In the next step, we applied this optimized prompt to the entire dataset to automatically identify AI-related advertisements in a large-scale job advertisement dataset.

5.2.2.4 Prompt Optimization with DSPy Package

A known issue with LLM-based task-solving (as demonstrated by Si et al., 2023) is that the model's performance can vary greatly in task-solving depending on the exact wording of the instructions (prompts), even if they may seem equivalent to a human reader. Accordingly, a wide range of automatic prompt optimization procedures have been developed (for an overview, see Li et al., 2024), which treat the question of fine-tuning a given prompt as an unsupervised or supervised learning task.

A particularly useful tool for this purpose is *Khattab et al., 2023*'s open-source software package, which enables automatic prompt optimization.

DSPy (Declarative Self-improving Python) is an open-source framework developed by the Stanford NLP group with the goal of enabling the programming of language models

(LLMs) in a declarative, modular manner instead of traditional prompting. Instead of using manually written prompts, DSPy allows users to describe the desired behavior as Python code, which the system automatically optimizes based on specified metrics. This approach is particularly useful for complex tasks such as chain-of-thought reasoning (see Wei et al., 2023), retrieval Augmented Generation (RAG) (see Lewis et al., 2021), or multi-step "agent systems." DSPy's modules and optimizers allow for the automatic fine-tuning of prompts (and even model weights), thereby reducing the need for manual intervention and increasing the system's reliability and performance.

In our work, using this package, we formulated a supervised learning task with a training dataset consisting of 150 non-AI-related and 50 AI-related job advertisements, and a test dataset of 10 examples, ensuring that misleading advertisements similar to those mentioned above and other challenging data points were included in the test. Using the DSPy framework, we created a classifier with 100% test F1 metric, i.e., one that solved the task flawlessly on the test data, using the *COPRO (Cooperative Prompt Optimization)* optimization procedure.

To create our solution, we used the COPRO optimization procedure. COPRO (Cooperative Prompt Optimization) is an optimization algorithm available in the DSPy framework that performs iterative fine-tuning of prompts using cooperative learning strategies. A key element of the method is that during a joint optimization process, it not only modifies the prompts but also could adaptively tunes the model, ensuring maximum alignment with the target task. In our case, no model fine tuning was necessary. COPRO learns from feedback on performance achieved on the given dataset through multiple rounds of iteration and continuously refines the optimized prompts. This allows the resulting prompts to make the best use of the base model's capabilities and minimize errors such as irrelevant or misleading classification in the given context.

During the procedure, starting from the original prompt, several (by default 10) alternative instructions are generated, to which the system automatically attaches a so-called *prefix* part corresponding to the last output field. Then, each alternative prompt is evaluated individually with the specified metric by rerunning the entire program. Based on the feedback obtained, the best-performing instructions are selected, and new instruction candidates are generated from them by providing examples (in a few-shot style). This process is repeated through multiple (e.g., 3) depth iterations. During selection, not only the absolute best but also the average results of the latest prompt generation rounds are considered. The code allows for tracking statistical indicators (minimum, maximum, average, standard deviation) during the optimization process.

The optimized prompt expression resulting from the procedure can be found in Appendix Appendix G.

Subsequently, we used the prompt with the *OpenAI GPT-4o-mini* model via the OpenAI Platform API (see OpenAI, 2025), and with approximately 200 hours of runtime (due to Tier

4 limitations), we ran it for binary semantic classification on the entire 1.3 million dataset.

In our subsequent analyses, we will refer to the keyword-based solution as RegEx and the above LLM-based solution as DSPy.

5.3 Analysis

5.3.1 Exploratory Data Analysis

As the first step in data analysis, we conducted exploratory data analysis to examine the spatial and temporal distribution of job advertisements.

The database contained a total of 1382827 job advertisements, distributed by publication date as follows (Figure 5.1):




Figure 5.1: Temporal distribution of job advertisement publications.

As can be seen, the number of data points in 2019 and 2020 is not entirely balanced compared to other years (which may cause problems in trend forecasting), but it is still a large number, and from 2021 onwards, it is a particularly robust dataset.

The spatial distribution of the data shows significant imbalance with the absolute dominance of Budapest. (For graphical representation, Budapest's data had to be scaled down to 1/3.) (see Figure 5.2)




Figure 5.2: Spatial distribution of job advertisements. (Budapest scaled down by 1/3)

Image Source: Author's illustration

This clearly shows the limitation of the data in that it is only from online sources, with foreign-owned online job portals dominating. On one hand, it is important to highlight that even large Hungarian portals like Profession.hu are missing, and on the other hand, this orientation may shift the data towards more internationally oriented jobs. Fortunately, upon closer examination, there were plenty of jobs requiring lower qualifications as well as those outside Budapest (see Figure 5.3).




Figure 5.3: Distribution of job advertisement sources.

Image Source: Author's illustration

The analysis of the 15 most common cities clearly shows Budapest's dominance, but it also shows generally good coverage of the entire country. (see Figure 5.4)




Figure 5.4: Top 15 settlements by frequency of job advertisements.

Image Source: Author's illustration

5.3.2 General Trends

The main focus of our analysis is the sectoral distribution of AI-related job advertisements and the analysis of their general and specific temporal movements. However, to form a meaningful picture of these, it is first worth comparing the results of the two developed methods, i.e., the RegEx and DSPy-based classifications, and their differences, which may shed light on the unique advantages of the methods.

5.3.2.1 Differences Between RegEx and DSPy Methods

As mentioned earlier, the DSPy-based classifier performed very well on the selected test data, and it also performed well in random checks after applying it to the entire dataset, so - unsurprisingly - we will rely more heavily on it in our investigations.

The "confusion" (or rather "disagreement") matrix between the RegEx and DSPy-based solutions is as follows (see Figure 5.5):




Figure 5.5: Confusion matrix of agreement/disagreement between RegEx and DSPy solutions

Image Source: Author’s illustration

As can be seen, there are over 3,000 entries that the RegEx solution simply did not identify, for example, because they are related to AI based on the meaning of the text, not keyword matching, or because they gave a positive signal due to superficial (sometimes purely random) word matching. As such, we might feel that the results of the RegEx solution are not worth using, but as discussed later, they may highlight a special case where companies deliberately use AI-related terms in their advertisements even if the advertised position is not related to the topic.

Overall, in the dataset (based on DSPy), there are **18,496 AI-related advertisements**, which represent a negligible **1.34% of the total advertisement inventory**. From this, it can be concluded that direct work with AI technologies is present in a very small percentage of Hungarian job advertisements.

5.3.2.2 General and Temporal Distribution

First, considering the spatial distribution of AI-related job advertisements, it is clear that Budapest's general dominance is even stronger. This is greatly influenced by the fact that the base data itself is strongly Budapest-focused, although the conclusion that companies in the capital are ahead in AI usage does not seem unfounded, especially since, compared to the general frequency of advertisements, AI-related ones are disproportionately concentrated in Budapest (see Figure 5.6).




Figure 5.6: Spatial distribution of AI-related (classified based on DSPy) job advertisements - top 10 cities

Image Source: Author's illustration

Examining the temporal distribution of all AI-related job advertisements based on DSPy (see Figure 5.7) highlights important results: Considering only the number, two significant "waves" can be observed in the time series: in 2020 and 2023. While the first wave is likely related to "classic," non-"foundational model" (or more commonly known as "generative AI") solutions and can be explained by the boom in the automotive sector, the 2023 wave is likely influenced by the announcement of OpenAI's ChatGPT model in November 2022 and the subsequent "interest tsunami" phenomenon that permeated public discourse. A good

illustration of this is the frequency of Google search terms both overseas (see Figure 5.8) and in Hungarian searches (see Figure 5.9) (Source: Google, 2025)




Figure 5.7: Frequency of AI-related job advertisements (classified based on DSPy)

Image Source: Author's illustration

However, it should be noted that there is an important difference that we also considered when choosing the specific search term: Searching for the AI topic in general, the overseas interest is huge, but the rise related to ChatGPT coincides with the very active use of "AI jobs" terms. In the case of the chart for Hungary, we see searches for the "artificial intelligence" topic in general, because search terms like "artificial intelligence job" and similar result in negligible, uninterpretable small data.




Figure 5.8: Google Trends data for "AI jobs" search term from searches initiated in the United States, measured in Google's own relative metric

Image Source: Author's illustration




Figure 5.9: Google Trends data for "artificial intelligence" search term from searches initiated in Hungary, measured in Google's own relative metric

Image Source: Author's illustration

This phenomenon may suggest that while overseas, the arrival of the ChatGPT model (and with it the entire "generative AI" trend) was considered by the public as directly related to the world of work, focusing on its "practical application," in Hungary, the public discourse viewed these models more as a "curiosity" distant from work, and in the Hungarian world of work, the "classic," task-specific AI concept still dominates. This effect can also be observed if we examine Hungarian AI job advertisements not in their absolute numbers but relatively, as a percentage of all advertisements (see Figure 5.10). In this view, the 2023 wave, while noticeable, is significantly smaller in magnitude.




Figure 5.10: Percentage of AI-related job advertisements (classified based on DSPy) - Total population: 1382827

Image Source: Author's illustration

This phenomenon may explain the somewhat contradictory result that, when analyzing

the frequency and percentage time series with the ARIMA (see Box and Jenkins, 1970), specifically AutoARIMA (see Hyndman and Khandakar, 2008) method and applying a one-year further prediction horizon, the total frequency forecast (see Figure 5.11) shows an upward trend, while the percentage data forecast (see Figure 5.12) shows a downward trend. This may indicate that, as discussed above (at least at the moment limited by our analysis dataset), the "generative AI" wave has not yet "arrived" in Hungary in 2024.




Figure 5.11: Forecast for AI-related job advertisements based on DSPy – Model ARIMA 0,1,1 with seasonal order 0,0,12

Image Source: Author's illustration




Figure 5.12: Forecast for the percentage of job advertisements mentioning AI-related terms based on DSPy - Model ARIMA 0,0,4 with seasonal order 0,1,12

Image Source: Author's illustration

5.3.3 Where Are the Latest Technologies?

The question arises as to what extent the latest incarnations of AI technologies, commonly referred to by the public and media as "Generative AI" solutions (where it would be more accurate to generally speak of "foundational models" in the sense of Bommasani et al., 2022a, or specifically most often "large language models"), appear in advertisements. After a thorough search, it can be determined that "ChatGPT" as a specific solution and LLM technologies in general are meaningfully mentioned in 128 advertisements, and "Generative AI" in 106 advertisements in total, indicating that the labor market penetration of this technology is still very weak.

It should be noted here that if these are considered general-purpose technologies, it is possible that many advertisers may think that mentioning ChatGPT or other LLM technologies is "not even worth mentioning" in terms of the job. We cannot rule out this possibility, but it is worth mentioning that the number of advertisements containing the term Excel (based on RegEx, only as a standalone, capitalized word) is 168,348, which is still 12.17% of all advertisements, despite the fact that we often feel this technology is almost part of the "mandatory office minimum."

From this perspective, it will be worth conducting further research on the integration of the latest AI technologies, and we can conclude that the above-mentioned 100 mentions strongly suggest that these technologies have simply not yet reached the everyday operations of Hungarian companies.

5.3.4 Sectoral Distribution

A particularly interesting aspect of the analysis is the sectoral distribution of organizations most frequently advertising AI-related topics. To make this examinable, we manually processed the majority of company names in the data (all companies that posted more than 5 AI-related advertisements), merging all known variations of their names, ensuring that in the case of certain corporate groups, their subsidiaries were also grouped under one name. Subsequently, we looked up the TEÁOR'25 classification of the top 100 companies most frequently posting AI-related advertisements (based on both DSPy and RegEx), and using the first two digits of their main activity, we conducted a sectoral analysis. (see Figure 5.13)




Figure 5.13: Sectoral distribution - based on TEAOR main activity - of the top 100 companies posting AI advertisements (based on DSPy)

Image Source: Author's illustration

The sectoral distribution reveals interesting insights, although it also contains distortions: First, it is quite striking, though not surprising, that the IT sector dominates, with a 24% share in the top 100, far exceeding the number of AI-related advertisements from the information technology sector. This can be considered an expected result, as this sector partly facilitates AI usage in other sectors - for example, as a supplier - and is naturally "closer to the topic," but it can also be considered an "early adopter" (see Rogers, 2003). This implies that in future focus analyses, it is worth concentrating on this sector in terms of AI adoption.

The second-place "labor market services" sector is quite misleading: our analysis suggests that it reached such a prominent position because all advertisements posted by recruitment agencies or "headhunters" that do not wish to publicly disclose their client's name had to be attributed to the recruitment agency itself in the database. Thus, the 14% represented by this sector should actually represent a kind of "unknown" category.

In third place is clearly the consulting sector, not only because it can be considered a progressive, "technology-close" sector, but again because consulting firms are likely building AI capacity to develop solutions for their clients as part of their services. In our opinion, this sector, along with IT, strongly represents the *AI supplier layer*.

Other noteworthy areas include the banking, wholesale (mainly energy trading), and research (mainly pharmaceutical research) sectors, partly because their financial strength allows them to invest in forward-looking technologies, thus, in addition to "ordering" AI solutions, they have the opportunity and need to build their own specialist capacity.

5.3.5 Company-Level Analysis

Examining in more detail - and considering the limits of visualization - the companies most frequently associated with AI-related job advertisements (see Figure 5.14), we can draw

interesting conclusions: Among the most dominant companies, two, Bosch and Continental, are associated with the automotive sector, and in both cases, the Hungarian development centers focused on self-driving cars may contribute to their strong position. This theme is reinforced by the presence of AIMotive (formerly known as ADAS-Works), an innovative Hungarian company exclusively focused on self-driving. We can conclude that the field of self-driving is one of the most significant sectors employing AI specialists in Hungary, collectively accounting for nearly 20% of all AI advertisements. This area can be considered a key sector for Hungarian AI development.




Figure 5.14: Sectoral distribution of the top 50 companies posting AI advertisements (based on DSPy)

Image Source: Author's illustration

While Randstad's presence is still explained by the "headhunting" phenomenon, IBM and Deutsche Telekom (mainly through its former "IT Services" division) primarily belong to the field of AI development services provided to corporate clients, reinforcing our previous assumption that most companies rely on large software suppliers or consulting firms to meet their AI-related needs instead of developing individual AI specialist capacity.

5.3.5.1 Unique Analyses, Temporal Dynamics

"One Wave" Strategy Our hypothesis about the decisive role of the automotive industry is further strengthened when we closely examine the temporal distribution of Bosch Group's AI advertisements, both in absolute terms and as a percentage (see Figures 5.15 and 5.16): while in absolute terms, 2021 can be considered outstanding, it is also striking that during a period in 2020, nearly 20% of all Bosch Group advertisements were AI-relevant, coinciding with the scaling up of the group's efforts in self-driving, and well before the November 2022 breakthrough of ChatGPT. Although the topic of self-driving played a prominent role in Bosch's Hungarian activities earlier, making exact timing not easily justifiable, the strength of the topic is well illustrated by, for example the press release from early 2021 by Robert Bosch Kft., 2021.




Figure 5.15: Number of AI-related advertisements by Bosch Group

Image Source: Author's illustration




Figure 5.16: Percentage of AI-related advertisements by Bosch Group

Image Source: Author's illustration

This "wave-like" demand for AI, which saturates within a short period, suggests a kind of fixed-frame competency building.

"Transition" Strategy However, it is worth comparing this with, for example, IBM's strategy, which is particularly striking from the percentage-normalized time series data (see Figure 5.17).




Figure 5.17: Percentage of AI-related advertisements by IBM Group

Image Source: Author's illustration

IBM's strategy seems to be maintaining a consistently high percentage of AI-capable workforce after an initially lower starting level, i.e., a kind of continuous hiring (or at least continuous talent search), which may indicate that the company perceives sufficient market demand from mid-2021 onwards to strategically build and currently expand/refresh its AI-

capable workforce capacity. In essence, we could say that after a turning point, the company "transitioned" to a certain more or less constant level of hiring (or searching) in the topic.

"Company Branding" Strategy However, if we expand the scope of the investigation beyond the strictly (DSPy-based) top companies, we can notice an interesting case: Although IQVIA is not among the most frequent seekers based on DSPy (though it represents a "middle ground"), it is very interesting to observe that based on RegEx, i.e., in terms of keyword appearance, it is the most dominant company (top 1). Its RegEx-based temporal, percentage data (see Figure 5.18) shows a short run-up followed by maximum, 100% saturation.




Figure 5.18: Percentage of advertisements mentioning AI based on RegEx by IQVIA Group

Image Source: Author's illustration

Interpreting the observations: Although in terms of content, it is far from the case that IQVIA is only looking for AI specialists in the labor market (it is not among the most frequent based on DSPy), the company consistently mentions the AI topic in all its advertisements as part of corporate branding, i.e., it feels that this is a positive, attractive message to potential employees and an integral part of the corporate image.

The above individual examples illustrate well that there are various different strategies available for companies to incorporate AI capabilities into their workforce.

5.3.6 Attempts at Group Formation

As a complement to the above manual "strategy search," we also conducted a thorough quantitative-based search to see if any further informative and interpretable grouping could be identified in the company-specific percentage-normalized time series data of AI advertisements determined based on DSPy. To do this, we first performed clustering with the DBSCAN algorithm (see Ester et al., 1996) on the "raw" monthly percentage time series, where the determination of the optimal ϵ hyperparameter was entrusted to the "Kneedle"

algorithm (see Satopää et al., 2011), and then repeated this analysis on the "Dynamic Time Warping" (see Sakoe and Chiba, 1978) pre-processed pairwise similarity matrix, followed by the "Matrix Profile" (see Yeh et al., 2016) pre-processed data, as well as its quarterly smoothed version. However, we were unable to identify any significant new strategies beyond the above findings.

5.3.7 Job Role Analysis

To further explore the specific roles behind the advertisements, we analyzed the distribution of specific roles mentioned or inferred from the texts. For this, we again used the set of AI-related advertisements identified by the DSPy method (a total of 18,496 advertisements) and performed a two-step LLM-based processing to extract 4,574 unique job titles from the texts by instructing the LLM to extract the specific position if it appears in the text, or to provide the most likely job title based on the text if it does not. This set was then transformed into a frequency list of 983 unique "aggregated" or "abstracted" elements through another LLM processing step, where, for example, *ML Ops Software Engineer* is replaced with *Machine Learning Engineer*, *EMEA Payroll Accountant* is simply replaced with *Accountant*, *Project Manager - Deep Learning* is replaced with *Project Manager*, etc. Finally, each advertisement was assigned to one of these elements. The top 50 elements of the resulting frequency statistics are shown in Figure 5.19.




Figure 5.19: Top 50 job titles by frequency

Image Source: Author's illustration

Based on a closer examination of the figure, it can be observed that a significant portion of AI-related positions are software development-focused, as evidenced by the mention of *software engineer* alongside all AI and data-oriented software specialist mentions. This suggests that the majority of job advertisements focus not only on the application of AI

but also on its development and integration. The most common positions are as follows:

- **software developers:** key players in AI development and implementation processes, operating in various corporate environments
- **specialists:** a significant number of **data scientist** and **machine learning developer** positions appear, focusing on the development and fine-tuning of AI models
- **development support roles:** positions such as **architect** and **project manager** clearly indicate the complexity of AI projects and the structural support needed for their successful implementation
- **managerial roles:** the presence of **product manager** positions suggests a trend where more companies are investing in the development and market introduction of AI-driven products or internal product development

It is also interesting to note that artificial intelligence does not only appear in IT-centric areas. Traditionally non-IT-focused positions such as **consultant** or even **accountant** also appear. This may suggest that the integration of artificial intelligence into broader work processes - albeit slowly - may have begun.

Finally, the presence of **internship levels** observed in job advertisements indicates that companies are actively seeking to nurture AI-related talent and ensure succession. This can contribute to the continuous development of the AI ecosystem and the maintenance of the industry's innovation potential in the long term.

The systematic establishment of the hierarchy of positions and potential in-depth analysis could be the subject of further research.

5.4 Conclusions

Summarizing our findings, we can state: In this research, we examined the labor market penetration of artificial intelligence (AI) in Hungary through text analysis of publicly available online job advertisements. Based on the results obtained, the following key conclusions can be drawn:

- **The proportion of AI-related job postings is low:** In the full database, the share of advertisements related to AI was only 1.34%, indicating that AI-related jobs currently represent a small proportion of the Hungarian labor market.
- **Geographic concentration:** AI-related job advertisements show a strong concentration in the capital. While Budapest already has a dominant position in the full dataset, this dominance is even more pronounced in the case of AI-related postings.

- **Temporal trends:** We identified two main waves in the temporal distribution of AI-related job advertisements:
 - a peak in 2020–2021, likely linked to developments in the automotive industry
 - a smaller increase in 2023, probably reflecting the impact of OpenAI’s ChatGPT launch, but without showing as sharp a rise as seen in international trends
- **Sectoral distribution:** AI-related postings are mainly concentrated in the IT sector (24%), consulting (11%), as well as the banking and wholesale sectors.
- **Dominant companies:** Among the companies most actively posting AI-related jobs, Bosch, Continental, and AIMotive play a leading role, reinforcing the dominance of the automotive industry in this domain.
- **Distinct strategies:** Companies follow different strategies regarding AI adoption:
 - Some companies follow a “wave-like” AI recruitment strategy.
 - Others show a more “gradual transition,” where the proportion of AI-related postings remains consistently high after an initial increase.
 - Finally, there are cases supporting a “company branding” strategy, where AI is mentioned in every job advertisement regardless of whether the role is directly related to AI.
- **Embeddedness of generative AI:** The presence of the latest AI technologies (such as ChatGPT) in job postings is minimal (only 128 explicit mentions), suggesting that these solutions have not yet been integrated into Hungarian labor market practices.
- **Methodological validation:** The LLM (DSPy)-based semantic classification proved more accurate than the simple keyword-based RegEx solution. While the RegEx method can highlight the brand-building use of AI references, it also generates many false positives.

The results show that AI’s presence in the Hungarian labor market is sectorally concentrated, geographically centered around Budapest, and that the adoption strategy varies from company to company. Although AI technologies are gaining popularity, their widespread application has not yet reached the level of mass adoption—particularly regarding newer generative AI solutions.

5.5 Limitations

Our research has several significant limitations, which we detail below.

One of the most important constraints is the increasingly general nature of AI solutions. Previously, AI solutions were task-oriented, while in recent years, large language models (LLMs) have emerged, functioning as universal assistants and problem solvers. This change affects the methods for measuring AI exposure as well.

In their works, E. W. Felten et al., 2018, E. W. Felten et al., 2019a, E. Felten et al., 2021a, as well as subsequent updated and expanded works such as E. W. Felten et al., 2023, examine the relationship between AI and occupations using an approach based on the O*NET taxonomy. They use EFF AI progress measurements and data collection processes conducted on Amazon Mechanical Turk, resulting in an "AI occupational exposure" (AIOE) score. (Similar research programs are also followed by Pizzinelli et al., 2023 and Gmyrek et al., 2023.)

Our current research examined job advertisements not assuming universal LLM "assistants," but specifically in a way that AI is explicitly mentioned in the advertisement text (RegEx), or the use of AI targeted systems is implicitly (semantically) assumed (DSPy). As AI technologies evolve, they become more general, and at some point, they may become as naturally integrated into our daily workflows as "Office" or other "office software packages" are today: widely used but rarely mentioned tools, which are simply not mentioned in a job advertisement unless it is for a very low-skilled or junior-level position (not mentioned literally even then). This means that in the future, analyzing AI exposure based on explicit mentions may become meaningless, and we will have to somehow revert to indirect analysis of job descriptions, making assumptions about what current technology usage they imply.

To this end, scaling up Felten's group's aforementioned research method with LLM-based processing (even using DSPy) can be applied to the analysis of the advertisement corpus. This represents a strong future research direction, following a method very similar to Eloundou et al., 2023, but it carries the risk that we subjectively determine the capabilities of the technology and their mapping, thus "embedding" our biases into the LLM instructions. Despite all reservations, due to the significant reduction in resource requirements for LLM-based processing, this remains a promising research direction.

Another development opportunity to deepen our results could be the potential noise-filtered processing of the mainly automatically generated columns contained in the original dataset, which could make the seniority of job roles and the possible level of compensation associated with them analyzable.

A further fruitful research direction could be the expression-level semantic analysis of texts to uncover which concepts or business areas are most closely related to the AI topic in the advertisement texts, potentially shedding light on the most affected business areas/processes.

5.6 Policy Recommendations

Based on the analysis results, it can be supported that the adoption of artificial intelligence (AI) in the Hungarian labor market is low and mainly linked to the IT sector. Several indications suggest that businesses do not consider AI, especially generative AI solutions, as general-purpose technology, and their application in everyday business processes is not yet widespread. In light of all this, several well-defined policy measures can be proposed to promote stronger AI adoption in the Hungarian labor market:

- **Information and training programs:** Launching widespread information campaigns to showcase AI business and industrial applications. Workshops, educational materials, and seminars organized through chambers of commerce could help business leaders and employees understand how to effectively use AI technologies.
- **“Generative AI User License”:** Introducing a national-level “generative AI user license” modeled after the ECDL (European Computer Driving License), certifying the effective application of AI-based tools and platforms in various business and industrial environments.
- **AI adoption-promoting vocational training programs:** Supporting training for non-programming-level but professional qualification roles related to AI introduction (e.g., AI integration consultants, AI coordinators), launching professional courses and adult education programs.
- **Monitoring AI adoption:** Continuously monitoring the spread of AI, especially generative AI, including tracking job advertisements, to achieve penetration similar to Excel in everyday workflows in the long term.
- **Funding sources supporting AI application:** Introducing targeted financial incentives to facilitate the corporate introduction of AI technologies. Providing grant funding for SMEs to integrate AI solutions into their processes through pilot projects.
- **Making state IT resources available:** Making IT infrastructure provided by the public sector (e.g., computing capacities, databases) accessible to companies and research institutions to promote AI research and development.
- **Targeted support for automotive AI developments:** The automotive industry is one of the leading application areas for AI in Hungary. Introducing targeted industrial policy programs and support to encourage the development of self-driving technologies and AI-based mobility solutions.

These measures can facilitate the broader application of AI in the Hungarian labor market, increase the level of AI technology adaptation, and ensure that generative AI solutions are integrated into the everyday operations of companies.

Chapter 6

"TRUST MY AI" - SURVEY ON THE HURDLES OF AI ADOPTION IN BUSINESS CONTEXTS

6.1 Topic and motivation

In this chapter I present the results of the "Trust My AI" survey, an international questionnaire measuring the relationship and factors influencing the perception and adoption of AI technologies by business in the Central-European area. (This chapter is submitted for publication as a separate article, and is currently in peer review).

6.2 Introduction and context

6.2.1 Emerging trust problems with AI systems

Parallel with the drastic increase in capabilities of foundational models in general and LLMs in particular some problems seemed to rise quite naturally from the fact that these models are in a sense "too good" at playing their roles, so even in their early work Radford et al., 2019 mentioned the capability of models to produce convincing but factually incorrect answers (later on popularly called "hallucinations"), that undermine the trust of their users in their abilities, and especially their usefulness in business relevant applications.

Current broad consensus from the technical community seems to have settled (for now) on the solution framework proposed by Lewis et al., 2020, namely "Retrieval Augmented Generation" (or RAG for short), that is basically constituting a hybrid solution, where the LLM is paired with a knowledge retrieval / search capability that provides relevant (often company specific or even internal) sources to "ground" the behaviour of LLM models, thus enabling them to substantially mitigate the problems caused by hallucination.

Thus said, even with these efforts in place, pretty impactful - and reputationally as well as

monetarily damaging - cases appeared in the news that report failure cases of LLM based solutions in business applications, like the case Ea, 2024 when Air Canada was ordered by the British Columbia Civil Resolution Tribunal to honor a policy created by its chatbot, which incorrectly promised a bereavement fare refund, after the airline argued the chatbot was a separate legal entity, or the case when a ChatGPT-powered AI chatbot at a Chevrolet dealership was tricked into agreeing to sell a 2024 Chevy Tahoe for \$1 after a buyer manipulated it into accepting any customer request as a legally binding offer (see AIAAIC, 2024).

These - and multiple similar - events were quickly disseminated by the media, and contributed to a substantial degree of (quite justified, one might argue) mistrust towards the application of AI in the business community. With our survey we specifically set out to investigate this trust relationship, and the possible venues for mitigation of distrust.

6.3 Research question

With the previous context in mind, we set out to investigate the following research questions:

- What is the current state of adoption of AI systems in the Central European region?
- What is the perceived level of trust in AI systems?
- How does this influence the decisions to apply AI in business contexts?
- What can be done methodically to increase trust in AI systems?

6.4 Methodology

This study employs a quantitative approach to assess perceptions of trust in AI systems and their influence on business adoption. The methodology primarily revolves around a structured survey distributed to professionals across various industries and organization sizes.

6.4.1 Survey Design and Distribution

The survey was designed to capture respondents' attitudes toward AI adoption, trust levels, and factors influencing AI integration. Key components of the questionnaire included:

- demographic information (region, industry, organizational role, organization size)
- AI usage within the organization (frequency, scope, and strategic importance)
- trust in AI systems (perceived risks, confidence in AI safety measures)
- evaluation of AI trust-establishing tools

The survey was distributed using a **snowball sampling** technique, leveraging social media platforms (e.g., LinkedIn, industry forums) and business networks. This approach ensured the recruitment of professionals actively engaged in AI-related discussions.

6.4.2 Sample Characteristics

- The survey received responses from **150 participants**.
- The regional distribution was predominantly from Hungary (58%) and Germany (30%), ensuring a strong representation from Central Europe.
- Respondents represented various industries, with a strong presence from the technology (50%) and financial sectors (12.7%).
- Participants came from organizations of various sizes, ensuring a balanced view of AI adoption across different business scales.
- Respondents held diverse organizational roles, including IT, executive management, and research and development.

6.4.3 Data Collection Period

The survey was conducted over a defined period, concluding on **August 25, 2024**. Responses were collected, anonymized, and analysed to ensure data integrity and respondent privacy.

6.4.4 Limitations

While the survey methodology provides valuable insights, several limitations must be acknowledged:

- **Self-selection bias:** Participants who are already engaged with AI topics may be over-represented.
- **Geographic limitation:** The majority of respondents are from Hungary and Germany, which may limit global generalizability.
- **Industry concentration:** The predominance of technology sector respondents could skew results toward AI-intensive industries.

Despite these limitations, the methodology provides a robust foundation for understanding AI trust dynamics within business contexts.

6.4.5 General description of survey participants




Figure 6.1: Country distribution of respondents

Image Source: Author's illustration

The survey collected responses from a **total of 150 participants**. The geographical distribution of respondents can be studied on Figure 6.1 which illustrates the country-wise representation, with the majority of respondents originating from Hungary (58%). This indicates that our insights will be heavily weighted toward Hungarian perspectives on AI adoption and related issues.

Germany accounts for the second-largest share of respondents (30%), followed by Austria (3.33%), North America (6%), and the Netherlands (2.67%). The strong representation from Hungary and Germany ensures that the findings are most reflective of AI adoption trends in Central Europe. The inclusion of Austria adds further context within the DACH region, though Switzerland is not represented.

The smaller number of responses from North America and the Netherlands suggests that the survey results may not fully capture perspectives from Western Europe or the broader international AI landscape. As a result, while the findings are relevant for understanding AI deployment in Hungary and Germany, they may not be as generalizable to other regions with different economic, regulatory and technological environments.




Figure 6.2: Industry of respondents

Image Source: Author's illustration

Figure 6.2 provides an overview of the industry distribution of respondents. The majority of participants (50%) are from the technology sector, reflecting the strong interest of professionals in AI-related developments. The financial sector represents the second-largest group (12.7%), followed by manufacturing (8%), education (6.7%), and retail (5.3%). Other industries, such as consultancy (2.7%), government (2.0%), research (1.3%), healthcare (1.3%), and public agencies (1.3%), are represented to a lesser extent. The remaining respondents are distributed across a variety of sectors, including logistics, transportation, construction, and legal research, each contributing a marginal share. This industry distribution suggests that the survey findings primarily reflect the perspectives of professionals working in AI-intensive sectors, particularly technology and finance, while industries with lower AI penetration, such as healthcare and public administration, are less represented.




Figure 6.3: Organizational Sizes of Respondents

Image Source: Author’s illustration

Figure 6.3 presents the distribution of organizational sizes among respondents. The largest share of participants (36%) work in large organizations with more than 251 employees, while medium-sized organizations (51-250 employees) and micro-enterprises (1-10 employees) each account for 26% of respondents. Small organizations (11-50 employees) make up 12% of the sample. The dataset primarily represents companies with a sufficient scale to consider AI adoption as a strategic initiative, avoiding a strong bias toward very small businesses that may lack the resources or incentive to engage in AI-driven transformation. The relatively even distribution of organizations across different size categories ensures that the survey captures insights from both well-established enterprises and smaller, more agile businesses, offering a balanced perspective on AI adoption across various operational scales.




Figure 6.4: Organizational role of respondents

Image Source: Author’s illustration

Figure 6.4 presents the detailed breakdown of organizational roles among survey respondents. The data highlights a diverse representation across multiple functions, with the largest proportion belonging to IT (26.7%), followed by Research and Development (24.0%) and Executive Management (14.7%). Other roles such as Finance, Marketing, Sales, Procurement, and Operations are represented to a lesser extent, each accounting for a few percentage points of the total. The long-tail of roles, including Data Science, Organizational Development, and Management Consultancy, suggests that AI adoption is relevant across a wide spectrum of business functions, albeit with varying degrees of involvement.




Figure 6.5: Organizational roles of respondents - Merged

Image Source: Author’s illustration

To facilitate the analysis of AI usage an views across broad functional categories, the roles were semantically grouped into key domains, as depicted in Figure 6.5. The classification consolidates various functions into four primary categories: IT (26.67%), Research and Development (27.33%), Executive Management (14.67%), and Other roles (31.33%). This grouping helps to better understand how different functional areas engage with AI adoption and trust-building efforts. The high representation of IT and Research and Development indicates a strong involvement of technical and innovation-driven teams in AI deployment, while the presence of executive management underscores the strategic importance of AI initiatives. The “Other roles” category captures the diversity of respondents whose roles may not fit neatly into the primary groupings but still contribute to AI-related decision-making and implementation.

6.5 Survey results

6.5.1 AI usage in the organizations




Figure 6.6: How frequently do you or your organization use AI technologies in your professional activities?

Image Source: Author's illustration

Our survey results (Figure 6.6) indicate a high level of AI adoption among respondents:

- **Frequently:** 35.3% of respondents use AI technologies regularly in their professional activities.
- **Occasionally:** 38.0% use AI from time to time.
- **Rarely:** 22.7% report minimal AI usage.
- **Never:** Only 4.0% of respondents state that they do not use AI at all.

These figures suggest that AI is widely integrated into the professional activities of our survey participants, with more than **73% using AI at least occasionally**.

6.5.1.1 Comparison with Eurostat AI Adoption Data

To put these results into perspective, we compare them with the latest AI adoption statistics from **Eurostat, 2024**, which reports the percentage of enterprises using at least one AI technology:

- **Germany:** 19.75% of enterprises used AI in 2024, up from 11.55% in 2023.
- **Hungary:** 7.41% of enterprises used AI in 2024, up from 3.68% in 2023.
- **EU Average:** 13.48% of enterprises used AI in 2024, up from 8.03% in 2023.

Compared to these figures, our survey shows a much higher rate of AI adoption among respondents. The gap between our results and the Eurostat data suggests that our sample is **not representative of the general enterprise landscape in Germany, Hungary, or the EU**. Instead, our dataset is likely biased toward professionals who are already engaged in AI discussions and implementation.

6.5.1.2 Interpretation of the Results

This response bias is expected and does not diminish the relevance of our findings. Since our research focuses on **AI trust and adoption concerns**, we are primarily interested in the perspectives of those who use AI, rather than those who do not. The high level of AI adoption in our dataset ensures that our analysis captures the opinions of those actively working with AI technologies, making our insights more relevant to understanding the challenges, risks, and trust factors associated with AI adoption in business contexts.

6.5.2 AI’s Importance and Impact on the Organization




Figure 6.7: Strategic Importance of AI by Organization Size - Percentage

Image Source: Author’s illustration

Figure 6.7 aims to examine how organizational size correlates with the strategic importance attributed to AI by our respondents. A general observation is that across all organization

sizes, a considerable proportion of respondents regard AI as either important or critical, indicating its growing relevance in business strategy. Large enterprises tend to recognize AI as highly significant, with 25.9% of respondents considering it critical and an additional 29.6% rating it as important. This trend is expected, as larger organizations typically have more resources, structured innovation strategies, and established AI-driven initiatives, making AI integration a key component of their operations.

An intriguing aspect of the data is that micro-enterprises report the highest percentage of respondents perceiving AI as critical, at 43.6%. This finding is somewhat unexpected, as smaller organizations often lack the same level of financial and technological resources as larger enterprises. However, it is likely influenced by the composition of survey respondents, many of whom may come from technology-centric micro enterprises, consulting firms, or other AI-driven sectors where the adoption and reliance on AI are fundamental to business operations. This suggests that in certain industries, particularly in technology and professional services, even very small organizations may consider AI to be a crucial element of their strategic approach.

In contrast, small and medium-sized enterprises exhibit a more balanced distribution of responses. For small enterprises, 22.2% of respondents classify AI as critical, while 44.4% consider it important, indicating that while AI holds value, it is not yet universally seen as an indispensable strategic asset at this scale. Similarly, medium-sized enterprises have a strong presence in the important category (48.7%), with a lower percentage (12.8%) identifying AI as critical. This suggests that while AI is widely recognized as an important tool for competitiveness and innovation, its absolute necessity may be less pronounced compared to micro and large enterprises.

Overall, the data highlights that AI is increasingly viewed as a strategic priority across businesses of all sizes, though the degree to which it is considered critical varies. Large enterprises predictably show a strong inclination toward integrating AI into their strategic framework, but the particularly high importance placed on AI by micro-enterprises points to industry-specific factors at play. This finding underscores the fact that while AI adoption is often associated with large organizations, smaller firms, especially those in AI-intensive industries, may be just as, if not more, reliant on it for their survival and growth.




Figure 6.8: AI’s impact on your industry by Organization Size

Image Source: Author’s illustration

The above mentioned effects are also reinforced by the respondents’ views about the impact of AI on their general industries, so as Figure 6.8 shows, they feel, that not only for their specific company setups is the adoption of AI a necessary component, but also estimate the impacts of it for their whole industries to be considerable. As visible on the chart, the effect of organizational size is also very similar on the respondent’s views, so large, and micro organizations (with sectoral caveats) might be more prone to feel the need for AI adoption.

6.5.2.1 Most frequent use cases of AI

As a deeper layer of analysis for the actual usage of AI (with emphasis on most recent AI capabilities) we included questions to assess the areas and use cases that respondents encounter the most frequently in their organizational contexts.




Figure 6.9: Distribution of Categories per Use Case - Internal Use Cases

Image Source: Author’s illustration




Figure 6.10: Distribution of Categories per Use Case - Customer-Facing Use Cases

Image Source: Author’s illustration

Figures 6.9 and 6.10 illustrate the most frequent AI use cases in organizations, categorized into internal and customer-facing applications. The data highlights that AI adoption is relatively higher for internal use cases, with an average of **54.5%** of respondents indicating at least minimal usage. In contrast, external AI adoption stands at **43.95%**, reflecting both interest and caution in deploying AI in customer-facing roles.

For internal applications, AI is most frequently used for accessing policies (71.3% indicate it is not used, while 28.7% report some level of adoption), generating text and charts (54.7% adoption, with 20.0% extensively using it and 2.7% marking it as a core function), and generating code (54.0% adoption, including 15.3% extensively using it and 6.7% consid-

ering it a core function). These results suggest that AI plays a significant role in knowledge retrieval and automation of documentation, as well as in content generation where Large Language Models (LLMs) are often implemented using Retrieval-Augmented Generation (RAG) techniques. However, challenges such as hallucinations—incorrect AI-generated outputs—remain an issue, which we will explore in further detail later in the article.

Externally, AI is primarily applied in **FAQ access** and **sales chatbots**. The survey data indicates that FAQ access is actively planned by **44% of organizations**, while sales chatbots are even more widely adopted, with **51.3% of respondents planning for at least minimal integration**. Notably, sales chatbots have a **higher proportion of respondents (9.3%) identifying them as a core function**, indicating that businesses see AI-driven customer interactions as a crucial part of their digital strategy.

A key observation is that **external AI use cases carry a higher risk compared to internal applications**. Since customer-facing AI directly interacts with consumers, errors in chatbot responses, bias in FAQ generation, or misinterpretations could lead to reputational damage, diminished brand trust, and even financial losses. This is particularly relevant for organizations where AI-generated interactions influence customer decision-making and support experiences.

The comparison between internal and external AI adoption suggests a **more cautious approach toward customer-facing AI applications** due to their potential risks. However, the fact that nearly **44% of organizations are already integrating AI externally** reflects the growing confidence in AI's capabilities and its increasing role in business automation. In the following sections, we will further analyse how organizations manage AI-related risks and trust-building mechanisms, especially in high-stakes external use cases.

6.5.3 Trust vs. usage

The core subject of our investigation is the trust (and the associated concerns / possible solutions) in AI systems applied in the professional business context. To analyse the attitudes and perceived trust in AI systems we decided to conduct a detailed evaluation of the usage patterns of said technologies with respect to the given organizational roles of the respondents (in a "merged" manner, as defined already above).




Figure 6.11: Percent Distribution of AI Usage by Role

Image Source: Author's illustration

But before we go deeper into the levels of trust and it's relationship to organizational role, as a prerequisite (and in a sense a possible confounder) we have to ponder the general usage patterns of AI for the given roles.

As Figure 6.11 demonstrates, though in our "tech savy" respondent pool, the utilization of AI tools is generally high (especially in comparison to the Eurostat, 2024 data quoted before), so on average over the roles 35.0% report frequent usage, 37.15% report occasional usage, or in sum a total 72.15% report noteworthy usage, the levels vary considerably with respect to their individual roles.

Beside the maybe more expected result that IT related roles are reporting the most dominant frequent usage of AI systems, research and development staff seems to rely in total (so frequent and occasional together) the most (at 80.5%) on AI, significantly higher than the total of "other roles" standing at 74.5% (which is in itself impressive). Also noteworthy is the observation, that though the AI usage of executive and management personnel does not reach the average, but with it's 63.% total still can be considered highly relevant.




Figure 6.12: Normalized relationship between role at work and trust in AI systems

Image Source: Author's illustration

With that in mind, looking at the data with respect to the connection between the roles at work and the respective trust in AI systems (illustrated by the Figure 6.12) an interesting picture emerges.

As a starting point, we have to observe, that if we average across roles the fully, strongly and limitedly trusting individuals, and sum these average proportions, we can state, that 53.4% of respondents place some level of trust in AI systems, which on the one hand can be considered promising (it is a slim majority after all), but prominently points to the fact, that **establishing and enhancing trust in AI systems is a crucial challenge in their adoption**, and most likely some **systematic approaches, on the policy and toolset level have to be implemented**.

But going deeper into the details reveals more interesting insights: generally, IT roles are the very strongly trusting 62.5% if we sum fully, strongly and limitedly (or with the highest proportion of fully, 32.5%), which is in a sense as expected. We could theorize, that beyond the general openness towards technology in these roles the possibly deeper knowledge / understanding of the way of operation of AI systems can play an important role here.

It is also noteworthy, that management related roles exhibit a remarkably strong trust, 63.7% (summed), so the highest of all groups! This is especially important, since it means that we do not just see a kind of "grassroots" level of naive trust emerging, but IT (as the "topic owners") and management (as the ones giving directions to the organization) jointly exhibit positive attitudes about AI systems, that can foster faster adoption that if only the "general" employees would turn to AI in their work (as represented by either the 53% average or the 36.2% of the "other" roles).

Further unpacking the notion of "more knowledge about AI's workings leads to higher trust" - as we theorized in case of IT, we can endeavour to directly quantify this relationship by examining the connection between usage frequency and trust directly, irrespective of roles.




Figure 6.13: Relationship between AI usage at work and trust in AI systems

Image Source: Author's illustration

Figure 6.13 tries to make this connection tangible by visualizing the dependence of the level of subjective trust and the frequency patterns of usage. Clearly, we can see a considerable relationship emerging, or to quantify specifically: a Spearman correlation of 0.371 with a P-value of <0.00001 , that we can definitely consider statistically significant. We can thus conclude, that one of the essential aspects of trust towards AI systems can be derived from specific knowledge, since **the more people use AI solutions, the more experience they have with them, hence the more they trust them**, which is in a sense extremely encouraging with respect to the further perspectives of adoption, but on the other hand points us towards the questions: what can we do, to enhance this trust even further?

On one hand: early and frequent exposure to AI technologies, so a culture and business investment that **enables employees to gather first hand, real life experience** is one of the most important factors we can think of. On the other hand, we can investigate some policy actions and practical technological solutions that can potentially boost trust, thus the level of adoption. In the remaining sections we focus our attention on these topics.

6.5.4 Organizational readiness and fears

If we want to make systematic progress in deploying AI systems in the broadest set of business applications possible, we have to also take a sharp look at the flip side of the above detailed phenomena, namely the topic of fears.

The first topic that we tried to quantify is what exactly are the areas - phrased in quite broad strokes so as to allow more respondent opinions to converge - that professionals are concerned about with respect to the broad scale adoption of AI technologies throughout business life.




Figure 6.14: Respondent's agreeing that AI poses risk for a given issue - Yes: Dark Grey, No: Light Grey

Image Source: Author's illustration

Figure 6.14 summarizes the fears, and makes it evident, that there is a clear hierarchy in the perceived level of risk with respect to topics. Most interestingly, the direct labour market impacts, so the potential of AI for causing job displacement lands squarely at the last place, so an 58.7% majority of respondents consider it as "non-issue", strongly contradicting the general economic narratives of the current day (though one could argue, the approx. 40% agreeing is by no means negligible), followed only closely by the topic of AI's inherent bias with 46.7% agreement, which might point to the fact that this topic has been already worked upon extensively in the last decade of AI progress. (It is well worth mentioning as a contrast, though, that the Eurostat, 2024 dataset reports for the proportion of "Enterprises which have measures to check the results generated by AI technologies for possible biases towards individuals" in the few EU member states which have at all this data sub 1% frequency, so it is quite questionable, if the problem is solved or simply neglected...)

The question of erosion of human skills is the first issue where the respondents seem to be ever so slightly (50.7%) more concerned than not, which also points towards the perceived potential of the technology: if the general consensus would point towards AI as being "incompetent" or "useless", the issue with over-reliance and thus, the diminishing of human

capabilities would not even emerge in the first place.

The topic of security vulnerabilities and of data privacy issues move hand in hand, a 60% majority consider them as noteworthy concerns, thus pointing to the fact, that clear policy guidelines and internal company regulations must be established to mitigate risks, which might in turn help establish and enhance trust.

Despite the overall trend of increasing trust in AI systems, and our reasoning above, that there does exist a high level of baseline trust in AI capabilities, concerns about performance persist, particularly regarding the misinterpretation of user queries. **64.7%** of respondents highlighted this issue, indicating that AI systems, while powerful, still struggle to always accurately understand and respond to complex input. This suggests that improving AI comprehension and contextual reasoning remains a key challenge for trust-building in professional environments.

Legal and compliance concerns are even more pronounced, with **70.0%** of respondents identifying them as significant barriers to AI adoption. Notably, security risks themselves were not perceived as a primary issue; instead, regulatory compliance — especially with GDPR — seems to be the dominant concern. Though we do not have direct evidence, we can hypothesize, that many respondents view data related regulations not necessarily as a safeguard but as an external burden that complicates AI deployment, suggesting that organizations may benefit from clearer regulatory guidance and standardized compliance frameworks to mitigate these concerns.

It is also notable for this point to ponder, that a general trend seems to exist in AI adoption representing the growing concern regarding data protection, privacy, and legal uncertainty, which has acted as a cooling factor on business enthusiasm for AI integration. Recent Eurostat, 2024 data demonstrates a notable increase in enterprises citing privacy and legal concerns as barriers to AI adoption. Specifically, the percentage of EU enterprises refraining from AI use due to data protection and privacy concerns rose from **2.74% in 2023 to 4.89% in 2024**, reflecting a heightened sensitivity to regulatory compliance. This trend is even more pronounced in Germany, where **7.6% of enterprises** in 2024 reported privacy concerns as a primary deterrent to AI implementation. Similarly, the proportion of enterprises avoiding AI due to a lack of clarity about legal consequences increased from **2.94% to 5.26%** across the EU, with Germany again exhibiting a significantly higher share at **8.25%**. Hungary, while demonstrating a lower overall prevalence of these concerns, still experienced a measurable increase, with privacy-related worries growing from **1.54% in 2023 to 2.56% in 2024** and legal uncertainty concerns rising from **1.84% to 2.92%** in the same period. These statistics suggest that businesses are not only struggling with the technical implementation of AI but are also facing escalating regulatory and compliance-related hesitations. The increased apprehension surrounding data protection laws, particularly the GDPR, underscores the necessity for clearer regulatory guidance, standardized compliance frameworks, and organizational policies that facilitate responsible AI deployment while maintaining legal certainty.

Without targeted interventions to address these concerns, such as government-backed support programs, streamlined compliance mechanisms, or industry-wide best practices for AI governance, regulatory uncertainty may continue to hinder AI adoption at a critical juncture of technological advancement.

Perhaps the most striking result is the widespread apprehension regarding AI's role in spreading misinformation. **70.7%** of respondents regarded this as a critical issue, reflecting not only concerns about AI-generated errors but also a broader awareness of the risks associated with deliberate misuse of AI technologies. This highlights the necessity for stringent verification mechanisms, transparency in AI-driven content generation, and policies to prevent the amplification of misinformation at both the corporate and societal levels.




Figure 6.15: Respondents perception of safety measures in their organizations

Image Source: Author's illustration

To round out the picture, we have to mention the question of confidence by the respondents, that the safety measures utilized by their respective organizations are sufficient in counteracting the possible dangers analysed above. As Figure 6.15, the opinions are quite split: in total 48% of the respondents is slightly or not at all confident that the policies and measures are adequate, so we can see a pretty obvious gap and need for serious work on the internal policy and trust-enabling technology area. These topics should merit considerable investment!

6.5.5 Tools for Establishing Trust

As discussed above, there are two main avenues for enabling trust in AI systems: organizational policies and technical evaluation tools. Here we focus on the latter – the recent

tools and methodologies that allow systematic auditing and error correction of AI solutions, thereby establishing a technical basis for trust. These solutions have advanced rapidly and now make the abstract concept of “AI trust” measurable and actionable in practice. By deploying the right evaluation frameworks, organizations can obtain objective evidence of their AI systems’ trustworthiness, which supports informed decisions about AI use and provides assurance to stakeholders.

6.5.5.1 What kind of tools exist for establishing trust?

Several evaluation frameworks have emerged to quantify the performance and trustworthiness of AI systems. Notably, RAGAS, ARES and TrustMyAI are three recent methodologies developed for this purpose. All three target retrieval-augmented generation (RAG) scenarios – for example, Large Language Model (LLM) chatbots that retrieve and cite external documents – but their principles apply broadly to AI assistants based on large language models. Each framework shares a common goal: to assess how well an AI model uses retrieved context to produce relevant, accurate, and trustworthy answers. They typically evaluate along key dimensions such as **context relevance (does the model appropriately use the provided information?)**, **answer faithfulness (is the answer grounded in the source data?)**, **answer relevance (does it address the user’s query fully and correctly?)**, and sometimes **transparency (does the system indicate where its information comes from?)**. These tools thus move beyond ad-hoc, human gut-checks toward formalized scoring of AI outputs. In effect, frameworks like RAGAS and ARES (see Chen et al., 2023; Mehrotra et al., 2023) enable a quantifiable trust score for AI responses, providing a structured way to judge whether an AI’s output can be trusted in a business context. This represents a significant recent development in AI evaluation: trust, which was once hard to measure, can now be tracked with metrics.

6.5.5.2 State-of-the-art methods for AI trust

RAGAS RAGAS (Retrieval-Augmented Generation Assessment) is an open-source framework introduced in 2024 to streamline RAG evaluation (see Mehrotra et al., 2023). It provides a suite of automatic metrics that jointly analyse the retrieval and generation components of an AI system. Rather than relying on manual annotation or predefined answers, RAGAS uses reference-free checks – for instance, scoring how well the retrieved context matches the query and how closely the AI’s answer stays true to that context. It measures attributes like relevance and factual consistency using language model scoring and similarity measures. In practical terms, RAGAS allows developers to quickly identify if a chatbot or QA system is hallucinating or diverging from its source material. Businesses have found this useful for rapid iteration: by running RAGAS metrics on test queries, teams can pinpoint weaknesses (e.g. irrelevant citations or unsupported statements) and improve the system be-

fore deploying it to real users. This improves the system's reliability and reduces the chance that end-users encounter incorrect information, directly boosting trust in the AI.

ARES ARES (Automated RAG Evaluation System) is another state-of-the-art methodology, emerging from the latest research to further automate trust evaluations (see Chen et al., 2023). ARES introduces AI-based judges into the evaluation loop: it generates synthetic Q&A pairs and fine-tunes lightweight language models to act as specialized reviewers for each component of a RAG pipeline. These AI "judges" score the context relevance of retrieved documents, the faithfulness of the answer to those documents, and the overall answer quality. By training on synthetic data (with a small set of human-verified examples for calibration), ARES minimizes the labor required for evaluation while still aligning with human judgment on correctness. The result is a largely automated scoring system that can be run continuously as an AI application evolves. In a business setting, ARES can serve as an ongoing monitor of an AI assistant's trustworthiness – flagging when the model's answers might be straying from verifiable facts or when retrieval quality drops. This proactive feedback loop helps organizations ensure their AI systems remain reliable over time, catching issues early (before users do) and maintaining consistency even as content or user queries change. Such capabilities are invaluable in high-stakes industries (finance, healthcare, etc.), where continuous validation of AI outputs is required for compliance and risk management.

TrustMyAI TrustMyAI is Neuron Solutions' proposed evaluation framework, which builds on concepts from RAGAS and ARES while tailoring them to practical enterprise needs. Like the others, TrustMyAI combines rule-based metrics with AI-driven assessment to produce an overall "trust score" for model outputs. In essence, it uses RAGAS-style scoring rubrics (checking context usage, factual alignment, relevance) and augments them with an LLM-based evaluator that judges answer quality in context – similar to ARES's approach. The aim is to provide a comprehensive yet easy-to-use toolset for organizations to audit their AI systems. For example, TrustMyAI can be integrated into a company's chatbot platform to automatically rate each answer given to customers, highlighting low-trust responses for human review or retraining. By doing so, it operationalizes trust metrics: rather than trust being a vague notion, it becomes a dashboard item that product managers and AI developers can monitor and improve. The business impact is significant – TrustMyAI offers a repeatable process to ensure AI deployments meet predetermined trust thresholds. This means companies can set standards (e.g. "the answer must be 90% faithful to provided documents") and regularly verify compliance, which is especially important for meeting internal quality benchmarks or external regulatory requirements.

Overview of RAG evaluation solutions In sum, frameworks like TrustMyAI, RAGAS, and ARES exemplify the cutting edge of AI evaluation. They turn qualitative concerns (like

honesty or relevance) into quantitative scores, making it feasible to validate AI systems at scale and iterate on them much like one would refine any other critical business process. Organizations adopting these tools gain a measure of control and confidence over AI behavior that previously was difficult to attain.

6.5.5.3 Platform integration of AI trust solutions

Importantly, the concepts behind these AI trust evaluation tools are being embraced and built into major AI platforms, reflecting an industry-wide move towards “responsible AI” by design.

Microsoft Azure For instance, Microsoft Azure now provides an integrated RAG evaluation framework as part of its AI offerings. Azure’s architecture guidelines for LLM applications recommend measuring metrics such as groundedness, relevancy, and correctness of responses – essentially similar criteria to the frameworks above – to ensure that generated answers remain faithful to the retrieved data (see Microsoft, 2024). Microsoft has even introduced GPT-4-based evaluation metrics within Azure Machine Learning Studio, allowing developers to automatically grade their chatbot responses on these trust dimensions. This means that for example, if you would build a RAG-based Q&A system on Azure, you can leverage built-in tools to calculate how well your model is using its knowledge sources and where it might be making unfounded claims. By embedding such evaluation at the platform level, Azure simplifies the adoption of trust measures: even non-technical teams can get feedback on AI output quality in real time, without needing to assemble a custom evaluation pipeline from scratch. This streamlined evaluation process helps businesses quickly identify issues (like a wrong citation or an incomplete answer) and improve their models, ultimately leading to more trustworthy AI services deployed on Azure.

Amazon Web Services Similarly, Amazon Web Services (AWS) has invested in AI trust and monitoring tools through its SageMaker platform. Amazon SageMaker Clarify, for example, is a tool that helps developers detect bias in models and understand the reasoning behind predictions, providing transparency into how AI decisions are made. (see Amazon Web Services, 2025a) Clarify produces detailed reports on potential biases in datasets and model outputs (e.g. checking if a model’s answers favour certain groups), and it offers feature importance analysis to explain why the model arrived at a given answer. These capabilities are crucial for building fair and transparent AI systems, which are key pillars of user trust. In addition, AWS offers SageMaker Model Monitor Amazon Web Services, 2025b, a managed service that continuously watches deployed models for quality issues. Model Monitor can automatically detect data drift, concept drift, and performance degradation in real time – for instance, alerting if users start asking questions outside the AI’s training distribution or if the accuracy of responses declines over time. By receiving such alerts, businesses can take

corrective action (retraining models, adjusting data inputs, etc.) before small issues snowball into major failures.

Other key players in RAG Evaluation Beyond Azure and AWS, other leading platforms are integrating trust and explainability features into their AI ecosystems. Google's Vertex AI platform offers tools for evaluating model fairness and explainability. For instance, Vertex Explainable AI provides feature-based and example-based explanations to enhance understanding of model decision-making processes. (see Google Cloud, 2025)

IBM's Watsonx platform supports the development of Retrieval-Augmented Generation (RAG) applications, enabling enterprises to build AI solutions that generate factually accurate outputs grounded in information from knowledge bases. In addition, IBM provides some of the aforementioned frameworks for evaluating and monitoring RAG pipelines, such as the Ragas framework, assessing the performance of RAG systems. (see Gutowska and Lukashov, 2024)

Overview of solutions for AI trust The convergence of these platform-level integrations from the aforementioned global software companies amplify the impact of frameworks like RAGAS, ARES, and TrustMyAI. By embedding evaluation and monitoring features directly into their products, cloud providers operationalize key concepts such as measuring groundedness, faithfulness, and bias at scale. This alignment with business demands for trusted AI facilitates the deployment of reliable and transparent AI solutions across various industries. For enterprises, this means adopting AI trust tools is becoming easier and more cost-effective. A team using Microsoft or AWS can leverage built-in evaluators and dashboards as a starting point, then extend with specialized frameworks for additional rigor if needed. The result is a robust end-to-end approach to AI trust – from development to deployment – that aligns with both technical best practices and emerging regulatory standards. In practice, organizations that embrace these integrated trust solutions are better positioned to deploy AI innovations with confidence. They can demonstrate to clients and regulators that measurable safeguards are in place, and they gain a competitive advantage by offering AI products and services distinguished by their reliability and accountability. In short, the convergence of new AI trust evaluation frameworks and their adoption by major platforms is making trustworthy AI a tangible, achievable goal for businesses today.

6.5.6 Value of trust tools

After his short survey of trust enabling technologies, the question naturally arises, what is the perceived value of these approaches in the eyes of the practitioners, so the professionals in the field who responded to our survey?

We directly stated this question, with some short explanation about what we think such a trust enabling technology would look like, trying to estimate the attitudes of respondents

towards them.




Figure 6.16: Participants perception of the value of an LLM assessment tool

Image Source: Author’s illustration

As Figure 6.16 illustrates, there is broad consensus about the value of trust enabling tools: in total **72% of respondents agree, that a tool that for example assesses the output of Large Language Model based solutions is at least moderately, very or extremely valuable**, pointing to the very specific felt need for some more systematic evaluations than just purely "manual" evaluation by humans. **We consider this finding very significant, and would thus highly encourage business decision makers to adopt such technologies / methodologies.**




Figure 6.17: Percent Distribution of LLM Assessment Tool Value by Role

Image Source: Author’s illustration

As for a more detailed, role based analysis of this value perception we can see by studying

Figure 6.17 that generally the more "hands on" roles are laying emphasis on the usage of LLM evaluation methods, but even in case of management and executive roles, more than 66% of the respondent agree, that such solutions carry at minimum moderately value for the organizations, making this an unanimously perceived need on the market.




Figure 6.18: Value of an LLM Assessment Tool by Organization Size

Image Source: Author's illustration

As a further important remark, by studying the organizational size distribution of this value judgement (illustrated by Figure 6.18) we can see, that especially large organizations have particularly high perceived need for evaluation methodologies in case of AI solutions. The relationship between this need and organizational size can be described by a Cramér's V: 0.3595 (moderate association) with Chi-squared: 55.82 and a p-value of: 0.00000013, thus we can conclude that the relationship between organization size and perceived value of an LLM assessment tool is statistically significant.

6.6 Discussion of Results

Our findings underscore that **trustworthiness is a cornerstone of AI adoption** in business. Even in our tech-savvy sample, only about **53.4% of respondents expressed at least some trust in AI systems**, a slim majority that highlights a significant trust gap. This gap matters because trust (or the lack thereof) directly influences the degree to which organizations embrace AI. Notably, we observed a virtuous cycle between **AI experience and trust**: respondents **who use AI more frequently tend to trust it more** (Spearman $\rho \approx 0.37$, $p < 0.00001$). This correlation suggests that familiarity breeds confidence – an encouraging sign that **hands-on exposure can build trust**.

However, experience alone may not bridge the trust gap for everyone. The key question that emerges is **what can organizations do** to enhance trust further and ensure that skepti-

cism about AI does not hinder its potential benefits? To answer this, we place our results in the context of the evolving AI governance landscape and derive practical strategies for businesses.

6.6.1 Policy Implications

The AI governance landscape is rapidly maturing, setting clear expectations for organizations to prioritize trustworthiness. Four major international governance or regulatory frameworks shape this environment, as we will discuss below.

6.6.1.1 EU AI Act

The European Union's Artificial Intelligence Act is the first comprehensive legal framework for AI, adopting a risk-based approach to ensure trustworthy AI.(see European Commission, 2021). The purpose of the AI Act was to build a comprehensive risk-based regulatory framework to ensure that AI systems are safe, transparent, and accountable.

It categorizes AI systems into four risk levels: unacceptable, high, limited, and minimal. High-risk systems (such as those used in critical domains like healthcare or finance) must satisfy stringent criteria including thorough risk assessments, data quality control, human oversight, and transparency measures. The Act also mandates clear documentation and establishes mechanisms for auditing and tracing AI-driven decisions, thereby aligning AI development with public trust and safety requirements.

The AI Act entered into force on 1 August 2024, and will be fully applicable 2 years later on 2 August 2026, with some exceptions:

- prohibitions and AI literacy obligations entered into application from 2 February 2025
- the governance rules and the obligations for general-purpose AI models become applicable on 2 August 2025
- the rules for high-risk AI systems - embedded into regulated products - have an extended transition period until 2 August 2027 (see European Commission, n.d.).

In practice, the AI Act means for boards and executives that **they must treat AI risks much like financial or legal risks**, with formal processes to audit and trace AI decisions so that systems remain aligned with public safety and trust. Non-compliance could **not only invite regulatory penalties but also erode public trust**, a risk few businesses can afford in an era where nearly 44% of organizations are already integrating AI into customer-facing functions, based on our survey data.

The EU's regulatory push thus serves as a wake-up call: organizations must be proactive in building trustworthy AI before mishaps lead to reputational damage or legal liability. As

our respondents noted, errors in consumer-facing AI (e.g. chatbots giving wrong answers or exhibiting bias) can diminish brand trust and even cause financial losses.

As we see in this context, the high-level message of the EU AI Act for business leaders is that **AI trust is now a governance imperative**. The Act requires companies to institute robust oversight for higher-risk AI applications, for example, by requiring thorough risk assessments, documentation, and human oversight for AI used in sensitive domains.

6.6.1.2 OECD AI Principles and NIST AI Risk Management Framework

Importantly, this European approach aligns with global trends. The OECD AI Principles provide international, consensus-based guidelines emphasizing human-centric and ethical use of AI. The principles include promoting inclusive growth, human-centered fairness, transparency, robustness, security, and accountability. They act as foundational guidelines influencing national and industry-level standards, contributing significantly to global alignment around trustworthy AI (see OECD, 2019).

Likewise to the OECD AI Principles, the **U.S. NIST AI RMF**, while voluntary, gives a practical blueprint to identify, assess, and mitigate AI-related risks through its core functions: **Govern, Map, Measure and Manage**. By operationalizing AI principles into actionable strategies, the RMF helps businesses achieve compliance and cultivate trust in AI deployment (see NIST, 2023).

These aforementioned frameworks translate abstract principles into a risk management process that companies can adopt internally, bridging the gap between lofty ideals and day-to-day practice. In essence, there is an emerging global alignment on what “trustworthy AI” entails – **transparency, fairness, accountability, and security are recurring pillars**. Businesses operating across borders can take note that adhering to these principles will not only ease compliance with any one regulation but also bolster their credibility in the eyes of customers and partners.

6.6.1.3 ISO/IEC 38507:2022 (AI Governance Standard)

Another key piece of the governance puzzle is corporate-level standards, such as **ISO/IEC 38507:2022 on AI governance**. ISO/IEC 38507:2022 provides detailed guidance for organizational governance of AI technologies. It reinforces corporate accountability for AI systems and stresses the importance of governance oversight at the board level. According to the standard, boards must set policies and oversee AI deployment processes to ensure ethical use and minimize risks associated with AI deployment (see ISO/IEC, 2022).

Boards are expected to set clear AI policies and oversee their implementation, ensuring ethical use and risk mitigation as part of corporate governance. This also corresponds to our previous findings that trust in AI is now a strategic issue, demanding executive attention. Our survey results show that trust in AI is not just a concern of the technical staff. In fact, we

found that management respondents exhibited among the highest trust levels in AI (nearly 64% reporting some degree of trust).

This suggests that leaders are increasingly optimistic about AI's potential – an optimism that needs to be backed by solid governance. By formally integrating AI oversight into corporate governance, organizations signal both internally and externally that they are serious about managing AI's risks and rewards. Such signals can enhance stakeholder confidence, aligning with what the emerging regulations and standards seek to achieve.

6.6.2 Building AI Trust in practice: Organizational strategies

Translating these policy insights into practice, our study points to several strategies that organizations can pursue to foster trust in AI.

6.6.2.1 Trust Evaluation in AI Processes

First, businesses should embed systematic trust evaluation into their AI development and deployment processes. Rather than relying on ad-hoc or purely manual assessments of whether an AI system is trustworthy, organizations can adopt structured evaluation methodologies as a routine "quality control" for AI.

Our respondents clearly see the merit in this approach: an overwhelming **72% indicated that tools for assessing AI outputs (e.g. evaluations of LLM responses) would be at least moderately valuable**. This consensus signals a felt need in the industry for more **rigorous, repeatable validation of AI systems** beyond what human intuition can provide. In line with this, we recommend that companies establish regular AI performance audits focused on trust metrics like accuracy, fairness, transparency, and robustness.

Emerging evaluation frameworks make this feasible. For example, methodologies such as **RAGAS** and **ARES** – originally developed for **retrieval-augmented generation (RAG) systems** – use a combination of automated metrics and AI-based judges to rate how well a model's answers align with provided context and factual correctness. Our own proposed framework, **TrustMyAI**, builds on similar concepts, combining RAGAS-style scoring with LLM-based judging to systematically measure an AI system's context relevance, answer faithfulness, and overall quality.

The specifics of these tools aside, the broader point is that **quantitative trust assessments are becoming accessible**. By adopting such tools (or others like them), organizations can obtain objective evidence of their AI systems' trustworthiness. This has two major business implications: internally, it supports more informed decision-making about where AI can be safely deployed; externally, it provides **assurance to regulators and clients**. In fact, as AI regulations tighten, having **quantifiable trust metrics** can demonstrate compliance with requirements for transparency and risk management, as envisioned by frameworks like the EU AI Act and NIST RMF.

In short, systematic evaluation methodologies offer a practical bridge between high-level governance mandates and on-the-ground AI operations, making the abstract concept of “AI trust” measurable and actionable.

6.6.2.2 Internal AI Policies and Culture

Second, organizations should **align their internal policies and culture with the emerging best practices on AI ethics and governance**. This involves training and awareness to ensure that everyone from developers to executives understands AI’s limitations and risks. Our results show that **security and data privacy are top-of-mind concerns**, with about 60% of respondents citing these as noteworthy issues with AI.

To address such concerns, companies need clear internal guidelines – for example, data handling rules for AI, bias mitigation checklists, and incident response plans – that translate external principles into daily workflow checkpoints. By instituting these policies, firms create an environment where compliance and ethical considerations are second nature, not afterthoughts.

Moreover, **transparency with users and stakeholders is paramount**. Organizations should openly communicate what their AI systems can and cannot do, including known limitations or error rates. This kind of honesty was echoed in our recommendations for enhancing AI literacy and transparency.

When users are educated about how an AI-driven decision is made (and what safeguards are in place), **they are more likely to trust the outcome**. Such openness is not only a good practice but may soon be a legal requirement for high-impact AI systems (as transparency is a common thread in AI regulations).

6.6.2.3 AI Certification and Labelling

Third, there is value in exploring **certification and labelling mechanisms for AI trustworthiness**. Just as industries have certifications for quality management or cybersecurity, AI is likely to move in this direction. Policymakers and industry groups are considering **trust seals** or standardized **trust scores** for AI.

Our study’s context suggests that tools like TrustMyAI, RAGAS, ARES and similar evaluation frameworks could contribute to these certification programs by providing consistent criteria to score AI systems. For organizations, participating in such certification schemes could offer a **competitive advantage** – a certified AI product may inspire greater confidence among customers and partners. It’s a way to **signal trustworthiness** in a crowded marketplace.

Furthermore, establishing industry-wide benchmarks for trust can help avoid a scenario where each company invents its own metrics (leading to confusion and “trust washing”). Instead, with common standards, businesses can compete on delivering genuinely trustworthy

AI, not just marketing claims. Given that **larger organizations in our survey expressed a particularly high need for AI evaluation methodologies**, we anticipate that leading firms will push for these sorts of standardized trust evaluations, which smaller enterprises might then adopt as they become more accessible.

6.6.2.4 AI Literacy

Finally, fostering trust also means addressing the **human element** within organizations. Trust in technology often hinges on how well the people using or affected by that technology understand and accept it. Our results hinted that **when management and IT teams jointly trust AI** (as we saw with their relatively high trust levels), **AI initiatives can gain momentum**.

To replicate this, companies should invest in **AI literacy programs** and cross-functional dialogues about AI. By educating employees on AI's capabilities and pitfalls, and by engaging stakeholders in setting AI ethics guidelines, organizations create a culture where building trustworthy AI is a shared mission. This cultural alignment complements the technical and policy measures discussed above, ensuring that trustworthiness isn't just a box to tick, but a core value driving AI adoption.

6.6.3 Key takeaways: AI Trust for competitive advantage

Stepping back, **why do these findings matter?** In practical terms, they illuminate a path for organizations to navigate the **dual challenge of leveraging AI and managing its risks**. Our research confirms that businesses see immense potential in AI – for instance, a majority of large enterprises in our survey regard AI as important or even critical to their strategy.

However, realizing this potential hinges on trust: without stakeholder trust, AI projects may face resistance, underutilization, or backlash when things go wrong. The discussion above outlines how organizations can respond. By **anticipating regulations** like the EU AI Act and embracing the spirit of those rules proactively, companies can avoid playing catch-up and instead position themselves as **trust leaders**. By adopting **rigorous evaluation tools and best practices**, they can directly address the concerns that make people wary of AI, whether it's an employee unsure about an algorithmic recommendation or a customer hesitant about an AI-driven service. The payoff is twofold: internally, smoother AI adoption and innovation, and externally, enhanced reputation and **regulatory compliance**.

In essence, our results and analysis can serve as a **note for organizations to convert AI trustworthiness into a strategic asset**. Firms that invest in AI governance and trust-building not only **mitigate risks** – such as biases, errors, or security breaches – but also create an environment where AI can be deployed with confidence and creativity. Trust in AI is not just a moral or compliance issue, but a **business enabler**. It determines which companies will successfully scale their AI deployments and win public trust, and which might stumble

due to unseen pitfalls. As AI technologies continue to evolve, those organizations that have woven trust into their AI strategy will be best placed to harness AI's transformative power responsibly and sustainably.

Our study contributes to this forward-looking perspective, showing that while the challenges of AI governance are significant, the tools, frameworks, and insights to tackle them are increasingly within reach. By combining high-level governance foresight with on-the-ground practical measures, organizations can confidently navigate the emerging era of AI with both **innovation and integrity**.

Chapter 7

CONCLUSIONS AND POLICY RECOMMENDATIONS

7.1 Conclusions

Based on the in depth analysis of the previous chapters, the following main conclusions can be drawn with respect to the effects of AI on the labour market, and especially the adoption of this technology in the Hungarian / Central-European context:

- The slight pessimism detected in the literature survey is warranted, based on the problems reported by interview subjects in Chapter 4.
- The theoretical constructs of Agrawal et al., 2022 on prediction and judgment are strongly reflected in the seniority levels required for successful AI applications (see Chapter 4), as well as some of the hurdles of adoption (like trust / lack of judgement - see Chapters 4 and 6) seem to be empirically detectable.
- There is a strong promise and lived experience in certain business areas - particularly software development - pointing to major potential productivity gains (see Chapter 4).
- The benefits of such productivity gains manifest in both positive phenomena (e.g., increased job stability, job satisfaction and productivity for employees; enhanced productivity and scaling for employers) and negative effects (e.g., layoffs due to inability to scale market demand, “shadow” IT and “shadow” productivity where employees conceal gains). Organizational culture and incentive systems are essential to mitigate these harms (see Chapter 4).
- The “junior paradox” - automation of junior-level tasks combined with the need for senior judgment and a shrinking supply of juniors - is a critical and dangerous phenomenon. It leads not only to immediate job polarization but also undermines the sustainability of highly skilled work, potentially resulting in workforce shortages and structural unemployment (see Chapter 4).

- AI adoption in Hungary remains immature overall; although some sectors are more advanced, AI skills are not yet demanded in adequate quantities (see Chapter 5).
- Adoption of generative AI (“GenAI”) is even less mature (see Chapter 5).
- The general-purpose nature of AI may mask the perception of skill requirements, as future hiring may assume AI proficiency as background knowledge - analogous to general computer and office skills - highlighting the need for broad baseline education programs (see Chapter 5).
- Trust in AI systems correlates directly with usage experience and the presence of systematic procedures and tooling for result validation. There is strong demand for such trust-building solutions, which represent a major adoption bottleneck assuming willingness and informedness are present (see Chapter 6).
- The risks of skill mismatches, the junior paradox, misaligned organizational incentives, and structural unemployment necessitate coordinated organizational and national policy interventions to mitigate risks and unlock productivity gains (see Chapter 4).

7.2 Policy Recommendations

Based on the above, the following broad policy recommendations can be given:

7.2.1 Company Level

- Invest in trust-building tools such as RAGAS and ARES to ensure reliable AI deployment.
- Recognize and address the junior paradox by maintaining clear career pathways and continual talent development.
- Develop education and tooling strategies to avoid “shadow IT” and surface hidden productivity gains.
- Realign incentives to reduce employee fear of automation-driven layoffs and encourage transparent reporting of productivity improvements.

7.2.2 Country Level

- Systematically prioritize AI skill development alongside critical thinking and judgment training in national education curricula.

- Prepare incentive interventions—such as tax credits for employers hiring juniors and enhanced unemployment benefits for reskilling—to offset the longer learning curves and unpaid experience requirements associated with higher judgment levels.
- Double down on supporting key industries in AI adoption while broadening the scope of application through infrastructure investments, grants, and other incentives.

7.3 Epilogue - Sense of urgency

Though the adoption rate of AI technologies - as we saw - in the Central-European region is still in it's infancy, the urgency for policy action is all the more given, since in case of the much more advanced markets, namely the United States, a huge wave of layoffs already hit the technology sector - according to data compiled by Roger Lee, 2020– at least 145000 workers were laid off from the US technology sector in the first half of 2025 alone - which point to major disruptive changes ahead (partly fuelled by AI), and thus point to a unique window of opportunity for policy action to avoid drastic consequences.

BIBLIOGRAPHY

- Acemoglu, D., & Restrepo, P. (2017). Robots and jobs: Evidence from us labor markets. *NBER Working Paper No. 23285*. <https://doi.org/10.3386/w23285>
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and prediction: The disruptive economics of artificial intelligence*. Harvard Business Review Press. <https://books.google.de/books?id=86lYEAAAQBAJ>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31–50. <https://doi.org/10.1257/jep.33.2.31>
- AIAAIC. (2024, January). Driver persuades chatbot to sell car for usd 1 [Accessed: [Insert Date]]. <https://www.aiaaic.org/aiaaic-repository/ai-algorithmic-and-automation-incidents/driver-persuades-chatbot-to-sell-car-for-usd-1>
- Amazon Web Services. (2025a). Detect bias and explain predictions with amazon sagemaker clarify [Accessed: 2025-03-12]. <https://aws.amazon.com/sagemaker/clarify/>
- Amazon Web Services. (2025b). Monitor models for quality with amazon sagemaker model monitor [Accessed: 2025-03-12]. <https://aws.amazon.com/sagemaker/model-monitor/>
- Anthropic. (2025). The anthropic economic index [Accessed: 2025-02-15. Full paper available at: https://assets.anthropic.com/m/2e23255f1e84ca97/original/Economic_Tasks_AI_Paper.pdf]. <https://www.anthropic.com/news/the-anthropic-economic-index>
- Appenzeller, G. (2024). Welcome to llmflation – llm inference cost is going down fast [Accessed: 2025-02-15]. <https://a16z.com/llmflation-llm-inference-cost/>
- Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- Bessen, J. E. (2019). Artificial intelligence and jobs: The role of demand [Accessed: 2025-04-30]. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial*

- intelligence: An agenda* (pp. 291–307). University of Chicago Press. <https://doi.org/10.3386/w24235>
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022. <https://www.jmlr.org/papers/volume3/blei03a/blei03a.pdf>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N. S., Chen, A. S., Creel, K. A., Davis, J., Demszky, D., . . . Liang, P. (2021). On the opportunities and risks of foundation models. *ArXiv*. <https://crfm.stanford.edu/assets/report.pdf>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., . . . Liang, P. (2022a). On the opportunities and risks of foundation models.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., . . . Liang, P. (2022b). On the opportunities and risks of foundation models.
- Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control* (1st). Holden-Day.
- Brauner, P., Hick, A., Philipsen, R., & Ziefle, M. (2023). What does the public think about artificial intelligence? — a criticality map to understand bias in the public perception of ai. *Frontiers in Computer Science*, 5, 1113903. <https://doi.org/10.3389/fcomp.2023.1113903>
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6370), 1530–1534. <https://doi.org/10.1126/science.aap8062>
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47. <https://doi.org/10.1257/pandp.20181019>
- Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018, September). *Notes from the ai frontier: Modeling the impact of ai on the world economy* (tech. rep.) (Accessed: 2025-04-30). McKinsey Global Institute. <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>

- Chen, Y., Yang, J., Yu, Y., & Luo, J. (2023). Ares: An automated evaluation framework for retrieval-augmented generation systems [Accessed: 2025-03-05]. *arXiv preprint arXiv:2311.09476*. <https://arxiv.org/abs/2311.09476>
- Choi, W. J., Liu, Q., & Shin, J. (2024). Predictive analytics and ship-then-shop subscription. *Management Science*, 70(2), 1012–1028. <https://doi.org/10.1287/mnsc.2023.4723>
- Colombo, E., Mercurio, F., & Mezzanzanica, M. (2019). Ai and the changing skill requirements of work organization. *Information Economics and Policy*. <https://doi.org/10.1016/j.infoecopol.2019.05.003>
- David, P. (1989). *Computer and dynamo: The modern productivity paradox in a not-too distant mirror* (The Warwick Economics Research Paper Series (TWERPS)). University of Warwick, Department of Economics. <https://EconPapers.repec.org/RePEc:wrk:warwec:339>
- Dechter, R. (1986). Learning while searching in constraint-satisfaction-problems. *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, 178–183. <https://www.aaai.org/Papers/AAAI/1986/AAAI86-029.pdf>
- DeepMind. (2024). *Gemini v1.5: Technical report* (tech. rep.). DeepMind. https://storage.googleapis.com/deepmind-media/gemini/gemini_v1_5_report.pdf
- Dell’Acqua, F., III, E. M., Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R. (2023, September). *Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality* (tech. rep. No. 24-013) (Working Paper). Harvard Business School. https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- Developer Nation. (2024). State of the developer nation 26th edition: Developer trends for q3 2024 [Accessed: 2025-04-30].
- Durante, Z., Huang, Q., Wake, N., Gong, R., Park, J. S., Sarkar, B., Taori, R., Noda, Y., Terzopoulos, D., Choi, Y., Ikeuchi, K., Vo, H., Fei-Fei, L., & Gao, J. (2024). Agent ai: Surveying the horizons of multimodal interaction.
- Ea, H. (2024). Air canada forced to honour chatbot offer. *Law Society Journal*. <https://lsj.com.au/articles/air-canada-forced-to-honour-chatbot-offer/>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. <https://arxiv.org/abs/2303.10130>
- Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, 226–231.

- European Commission. (n.d.). Ai act - shaping europe's digital future [Accessed: 2025-03-05]. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- European Commission. (2021). Proposal for a regulation laying down harmonised rules on artificial intelligence (ai act) [Accessed: 2025-03-05]. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- Eurostat. (2024). Artificial intelligence use in enterprises - eurostat [Accessed: 2025-03-05]. https://doi.org/10.2908/ISOC_EB_AI
- Felten, E., Raj, M., & Seamans, R. (2021a). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/https://doi.org/10.1002/smj.3286>
- Felten, E., Raj, M., & Seamans, R. (2021b). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42. <https://doi.org/10.1002/smj.3286>
- Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108, 54–57. <https://doi.org/10.1257/pandp.20181021>
- Felten, E. W., Raj, M., & Seamans, R. (2019a). The occupational impact of artificial intelligence: Labor, skills, and polarization. *NYU Stern School of Business*. <https://doi.org/10.2139/ssrn.3368605>
- Felten, E. W., Raj, M., & Seamans, R. (2019b, September). *The occupational impact of artificial intelligence: Labor, skills, and polarization* (tech. rep.) (Available at SSRN: <https://ssrn.com/abstract=3368605>). NYU Stern School of Business. <https://doi.org/10.2139/ssrn.3368605>
- Felten, E. W., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative ai [<https://ssrn.com/abstract=4414065>]. *Available at SSRN*. <https://doi.org/10.2139/ssrn.4414065>
- Foundation, E. F. (2023). *Ai progress measurement (archived)* [Accessed: November 2, 2023]. <https://www.eff.org/ai/metrics>
- Fradkov, A. L. (2020). Early history of machine learning [21st IFAC World Congress]. *IFAC-PapersOnLine*, 53(2), 1385–1390. <https://doi.org/https://doi.org/10.1016/j.ifacol.2020.12.1888>
- Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E., Wang, D., Youn, H., & Rahwan, I. (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, 116(14), 6531–6539. <https://doi.org/10.1073/pnas.1900949116>
- Furman, J., & Seamans, R. (2019). Ai and the economy. *Innovation Policy and the Economy*, 19(1), 161–191. <https://doi.org/10.1086/699936>

- Galster, M., Mitrovic, A., Malinen, S., & Holland, J. (2022). What soft skills does the software industry really want? an exploratory study of software positions in new zealand. *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '22)*, 272–282. <https://doi.org/10.1145/3544902.3546247>
- Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative ai and jobs: A global analysis of potential effects on job... *International Labour Office*. <https://doi.org/10.54394/fhem8239>
- Google. (2023a). *Trend analysis for deep learning* [Accessed: 2023-11-02]. <https://trends.google.com/trends/explore?date=all&q=%2Fm%2F0h1fn8h&hl=en>
- Google. (2023b). *Trend analysis for deep learning* [Accessed: 2023-11-02]. <https://trends.google.com/trends/explore?date=2019-01-01%202023-10-27&q=generative%20AI&hl=en>
- Google. (2025). Google Trends [Accessed: 2025-02-19]. <https://trends.google.com/>
- Google Cloud. (2025). Introduction to vertex explainable ai [Accessed: 2025-03-12]. <https://cloud.google.com/vertex-ai/docs/explainable-ai/overview>
- Greene, L. (2017, May). *Language: Finding a voice: Computers have got much better at translation, voice recognition and speech synthesis* [Technology Quarterly]. Retrieved November 2, 2023, from <https://www.economist.com/technology-quarterly/2017-05-01/language>
- Gutowska, A., & Lukashov, V. (2024). Evaluate rag pipeline using ragas in python with watsonx [Accessed: 2025-03-12]. <https://www.ibm.com/think/tutorials/ragas-rag-evaluation-python-watsonx>
- Hinton, G. (2017). *Neural networks for language and understanding* [Creative Destruction Lab Machine Learning]. <https://www.youtube.com/watch?v=o8otywnWwKc>
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2023). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions.
- Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172. <https://doi.org/10.1177/1094670517752459>
- Huang, W., Abbeel, P., Pathak, D., & Mordatch, I. (2022). Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *CoRR*, abs/2201.07207. <https://arxiv.org/abs/2201.07207>
- Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for r. *Journal of Statistical Software*, 27(3), 1–22.
- Inc., S. E. (2024). Stack overflow - where developers learn, share, & build careers [Accessed: 2024-03-24]. <https://stackoverflow.com/>

- International Labour Office. (2012). *International standard classification of occupations 2008*. International Labour Organisation (ilo). <https://www.ilo.org/public/english/bureau/stat/isco/docs/publication08.pdf>
- ISO/IEC. (2022). Iso/iec 38507:2022 - governance implications of the use of artificial intelligence by organizations [Accessed: 2025-03-05]. <https://www.iso.org/standard/56641.html>
- Jobspikr. (2024). Jobspikr job market dataset [Accessed: 2025-02-15]. <https://www.jobspikr.com/contact-us/>
- Khattab, O., Singhvi, A., Maheshwari, P., Zhang, Z., Santhanam, K., Vardhamanan, S., Haq, S., Sharma, A., Joshi, T. T., Moazam, H., Miller, H., Zaharia, M., & Potts, C. (2023). Dspy: Compiling declarative language model calls into self-improving pipelines. <https://arxiv.org/abs/2310.03714>
- Kleene, S. C. (1951, December). *Representation of events in nerve nets and finite automata* (Research Memorandum No. RM-704). RAND Corporation. Santa Monica, CA. https://www.rand.org/content/dam/rand/pubs/research_memoranda/2008/RM704.pdf
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2023). Large language models are zero-shot reasoners.
- Kovács, L., & Vastag, E. (2023). Magyar szoftverfejlesztők bérkülönbségeinek elemzése a stack overflow kérdőíve alapján. *Statisztikai Szemle*, 101(11), 947–977. <https://doi.org/10.20311/stat2023.11.hu0947>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25, 1097–1105. https://papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c4-Abstract.html
- Kusnitz, S. (2014). Why amazon’s anticipatory shipping is pure genius. *Forbes*. <https://www.forbes.com/sites/onmarketing/2014/01/28/why-amazons-anticipatory-shipping-is-pure-genius/>
- Laato, S., Mäntymäki, M., Islam, A. N., Hyrynsalmi, S., & Birkstedt, T. (2023). Trends and trajectories in the software industry: Implications for the future of work. *Information Systems Frontiers*, 25(2), 929–944. <https://doi.org/10.1007/s10796-022-10267-4>
- Lee, K. (2021). Economics of technological leapfrogging. In J.-D. Lee, K. Lee, D. Meissner, S. Radošević, & N. Vonortas (Eds.), *The challenges of technology and economic catch-up in emerging economies* (pp. 123–159). Oxford University Press. <https://doi.org/10.1093/oso/9780192896049.003.0005>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kumar, H., Bhooshan, S., Gray, I., Power, R., Power, A., Scao, A., & Wang, Y. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*.

- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2021). Retrieval-augmented generation for knowledge-intensive nlp tasks. <https://arxiv.org/abs/2005.11401>
- Lexikon, L. (2024). Hackathon jelentése [Accessed: 2024-03-24].
- Li, H., Leung, J., & Shen, Z. (2024). Towards goal-oriented prompt engineering for large language models: A survey. <https://arxiv.org/abs/2401.14043>
- Liu, H., Zhao, S., Zhang, X., Zhang, F., Wang, W., Ma, F., Chen, H., Yu, H., & Zhang, X. (2024). Liberating seen classes: Boosting few-shot and zero-shot text classification via anchor generation and classification reframing. <https://arxiv.org/abs/2405.03565>
- McCann, B., Keskar, N. S., Xiong, C., & Socher, R. (2018). The natural language decathlon: Multitask learning as question answering. <https://arxiv.org/abs/1806.08730>
- Mehrotra, S., Bhatia, K., Sharma, R. A., & Gupta, P. (2023). Ragas: Retrieval-augmented generation assessment for open-domain question answering [Accessed: 2025-03-05]. *arXiv preprint arXiv:2309.15217*. <https://arxiv.org/abs/2309.15217>
- Microsoft. (2024). Assess ai systems by using the responsible ai dashboard [Accessed: 2025-03-12]. <https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai-dashboard?view=azureml-api-2>
- Microsoft Research. (2023). Autogen [Accessed: 2023-11-23]. *Microsoft*.
- Minsky, M. L. (1987, February). *Society of mind*. Simon & Schuster.
- Nagyné Halász, Z. (2023). Az it szakember-képzés és a piaci kompetencia-elmvárások közötti diszharmonia feloldásának lehetőségei. *GIKOF Journal*. <https://doi.org/10.14267/GIKOF.2024.02>
- Naimi, L., Bouziane, E. M., Manaouch, M., & Jakimi, A. (2024). A new approach for automatic test case generation from use case diagram using llms and prompt engineering. *2024 International Conference on Circuit, Systems and Communication (ICCSC)*, 1–5. <https://doi.org/10.1109/ICCSC62074.2024.10616548>
- NIST. (2023). Ai risk management framework (ai rmf 1.0) [Accessed: 2025-03-05]. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
- OECD. (2019). Recommendation of the council on artificial intelligence [Accessed: 2025-03-05]. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Ogedengbe, F. A., Talib, Y. Y. A., & Rusly, F. H. (2024). Influence of structural factors on employee cloud shadow it usage during covid-19 lockdown: A strain theory perspective. *Cognition, Technology and Work*, 26(1), 63–81. <https://doi.org/10.1007/s10111-023-00748-0>
- OpenAI. (2025). Openai platform [Accessed: 2025-02-16]. <https://platform.openai.com/>
- Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). The pagerank citation ranking: Bringing order to the web. *Technical report, Stanford InfoLab*. <https://citeseerx.ist.psu.edu/doc/10.1.1.31.1768>

- Paolillo, A., Colella, F., Nosengo, N., Schiano, F., Stewart, W., Zambrano, D., Chappuis, I., Lalive, R., & Floreano, D. (2022). How to compete with robots: Assessing job automation risks and resilient alternatives. *Science Robotics*, 7. <https://doi.org/10.1126/scirobotics.abg5561>
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of ai on developer productivity: Evidence from github copilot. <https://arxiv.org/abs/2302.06590>
- Petrovska, O., Clift, L., Moller, F., & Pearsall, R. (2024, January). Incorporating generative ai into software development education [Conference date: January 5, 2024]. In J. Waite & R. Crosby (Eds.), *Cep '24: Proceedings of the 8th conference on computing education practice* (pp. 37–40). Association for Computing Machinery. <https://doi.org/10.1145/3633053.3633057>
- Pizzinelli, C., Shibata, I., & Stavrev, E. (2023). Labor market exposure to ai: Cross-country differences and implications. *IMF Working Papers*, 2023(216). <https://doi.org/10.5089/9798400254802.001>
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Amodei, C. (2019). Language models are unsupervised multitask learners. *arXiv preprint arXiv:1901.11504*.
- Rehurek, R., & Sojka, P. (2010). Software framework for topic modelling with large corpora. *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50. https://radimrehurek.com/lrec2010_final.pdf
- Robert Bosch Kft. (2021, March). Driver assistance systems and automated driving function development at bosch [Press release].
- Roger Lee. (2020–). Layoffs.fyi: Tech layoff tracker and startup layoff lists [Founded by Roger Lee in 2020; site had aggregated over 450000 layoffs as of early 2025].
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Sakoe, H., & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1), 43–49.
- Satopää, V., Albrecht, J., Irwin, D., & Raghavan, B. (2011). Finding a "kneedle" in a haystack: Detecting knee points in system behavior. *arXiv preprint arXiv:1108.3085*. <https://raghavan.usc.edu/papers/kneedle-simplex11.pdf>
- Scheffler, I. (2023). Github ceo says copilot will write 80% of code “sooner than later”. *Freethink*. <https://www.freethink.com/robots-ai/github-copilot>
- Sevilla-Lara, L., & Learned-Miller, E. G. (2022). Robustness and plasticity of object categorization. *arXiv preprint arXiv:2202.05924*. <https://arxiv.org/pdf/2202.05924.pdf>
- Si, C., Gan, Z., Yang, Z., Wang, S., Wang, J., Boyd-Graber, J., & Wang, L. (2023). Prompting gpt-3 to be reliable. <https://arxiv.org/abs/2210.09150>
- Singla, A., Sukharevsky, A., Yee, L., & Chui, M. (2024, May). The state of ai in early 2024 [Global Survey on AI]. [https://www.mckinsey.com/~media/mckinsey/business%](https://www.mckinsey.com/~media/mckinsey/business%20)

- 20functions/quantumblack/our%20insights/the%20state%20of%20ai/2024/the-state-of-ai-in-early-2024-final.pdf
- Stack Overflow. (2023a). 2023 stack overflow developer survey [Accessed: 2024-03-24].
- Stack Overflow. (2023b). Developer sentiment around ai/ml [Available online at <https://stackoverflow.co/labs/developer-sentiment-ai-ml/>].
- Statista Research Department. (2016). Share of mobile app developers worldwide in 2014, by country [Accessed: 2025-04-30].
- Szabados, L. (2024a). Ai's effect on labour: What does economic literature say? *The International Journal of Engineering and Science*, 13(5), 328–350. <https://doi.org/10.9790/1813-1305328350>
- Szabados, L. (2024b). The dynamics of ai adoption: Book review: *power and prediction: The disruptive economics of artificial intelligence* [megjelenés alatt]. *Alma Mater*, 1.
- Szabados, L. (2025). The impact of ai tools on the work and labour market prospects of programmers. *Vezetéstudomány / Budapest Management Review*, 56(4), 42–56. <https://doi.org/10.14267/VEZTUD.2025.04.04>
- Talagala, N. (2022). Data as the new oil is not enough: Four principles for avoiding data fires. *Forbes*. <https://www.forbes.com/sites/nishatalagala/2022/03/02/data-as-the-new-oil-is-not-enough-four-principles-for-avoiding-data-fires/>
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Crown.
- Tierney, J. (2011). Do you suffer from decision fatigue? *New York Times Magazine*. <https://www.nytimes.com/2011/08/21/magazine/do-you-suffer-from-decision-fatigue.html>
- Turk, A. M. (2023). Amazon mechanical turk [Accessed: 2023-11-02]. <https://www.mturk.com/>
- Ulgen, F. (2017). Creative destruction. In *Encyclopedia of creativity, invention, innovation and entrepreneurship* (pp. 1–8). Springer New York. https://doi.org/10.1007/978-1-4614-6616-1_407-2
- U.S. Bureau of Labor Statistics. (2023). *Occupational employment and wages, may 2022: Software developers (15-1252)* (tech. rep.) (Accessed: 2025-04-30). U.S. Department of Labor. <https://www.bls.gov/oes/current/oes151252.htm>
- U.S. Department of Labor, E., & (USDOL/ETA), T. A. (2023). *O*net online* [Accessed on November 3, 2023]. National Center for O*NET Development. <https://www.onetonline.org/>
- U.S. Department of Labor, Employment and Training Administration. (2024). 17-2021.00 - agricultural engineers [O*NET OnLine. Creative Commons Attribution 4.0 International License.].
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press.

- von Engelhardt, S. (2008). *The economic properties of software* (Jena Economic Research Papers No. 2008-045) (Accessed: 2025-04-30). Friedrich Schiller University Jena and Max Planck Institute of Economics. <https://doi.org/10.2139/ssrn.1430885>
- Wang, Z., Pang, Y., & Lin, Y. (2023). Large language models are zero-shot text classifiers. <https://arxiv.org/abs/2312.01044>
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2023). Chain-of-thought prompting elicits reasoning in large language models. <https://arxiv.org/abs/2201.11903>
- Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., ... Gui, T. (2023). The rise and potential of large language model based agents: A survey.
- Yeh, C.-C. M., Zhu, Y., Ulanova, L., Begum, N., Ding, Y., Dau, H. A., Silva, D. F., Mueen, A., & Keogh, E. J. (2016). Matrix profile i: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets. *2016 IEEE 16th International Conference on Data Mining (ICDM)*, 1317–1322.
- Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., & Yang, D. (2024). Can large language models transform computational social science? <https://arxiv.org/abs/2305.03514>
- Zippia. (2024). Freelance software developer demographics and statistics in the us [Accessed: 2025-04-30].

Chapter A

DOIS OF THE ARTICLES ANALYSED

ID	DOI	Author(s)	Published
1	10.1111/j.1121-7081.2004.00281.x	Behtoui, Alireza	2004
2	10.1353/eca.0.0026	Jan Hatzius,	2009
3	10.1080/09766634.2010.11885540	Chandra, Rajasree	2010
4	10.1016/s0169-7218(11)02410-5	Acemoglu, Daron; Autor, David	2011
5	10.1080/02529203.2012.702944	Shouhai, Ding	2012
6	10.2139/ssrn.2403824	Alvarez-Cuadrado, Francisco; Van Long, Ngo; Poschke, Markus	2014
7	10.1515/9781400823130.20	nan	2014
8	10.1257/jep.29.3.3	Autor, David H.	2015
9	10.1016/j.jmacro.2016.08.003	DeCanio, Stephen J.	2016
10	10.1515/9789048526352-006	nan	2016
11	10.2139/ssrn.2744714	Falck, Oliver; Heimisch, Alexandra; Wiederhold, Simon	2016
12	10.2139/ssrn.2845762	Jerbashian, Vahagn	2016
13	10.1111/j.1564-913x.2015.00051.x	AKÇOMAK, Semih; KOK, Suzanne; ROJAS-ROMAGOSA, Hugo	2016
Continued on next page			

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
14	10.2139/ssrn.2730550	Degryse, Christophe	2016
15	10.1080/00213624.2017.1391582	Santos, Marcelo; Sequeira, Tiago Neves; Ferreira-Lopes, Alexandra	2017
16	10.1109/indin.2017.8104891	Hamid, Oussama H.; Smith, Norris Lee; Barzanji, Amin	2017
17	10.1109/ecai.2017.8166487	Tudor, Sofia Loredana	2017
18	10.1007/s11948-017-9901-7	Cath, Corinne; Wachter, Sandra; Mittelstadt, Brent; Taddeo, Mariarosaria; Floridi, Luciano	2017
19	10.1017/9781316761380.002	nan	2017
20	10.1080/10301763.2017.1397258	Healy, Joshua; Nicholson, Daniel; Parker, Jane	2017
21	10.2139/ssrn.2931339	Petit, Nicolas	2017
22	10.2139/ssrn.3021135	Thierer, Adam D.; Castillo, Andrea; Russell, Raymond	2017
23	10.1111/irel.12193	Guery, Loris; Stevenot, Anne; Wood, Geoffrey T.; Brewster, Chris	2017
24	10.2139/ssrn.3015350	Calo, Ryan	2017
25	10.1080/0023656x.2016.1242716	Virgillito, Maria Enrica	2017
26	10.4337/9781785369070.00011	McIntosh, Steven	2017
27	10.1111/ecin.12412	Morikawa, Masayuki	2017
28	10.1093/qje/qjx032	Kleinberg, Jon; Lakkaraju, Himabindu; Leskovec, Jure; Ludwig, Jens; Mullainathan, Sendhil	2017
29	10.2139/ssrn.2934610	Oschinski, Matthias; Wyonch, Rosalie	2017
30	10.1109/sisy.2017.8080580	Rajnai, Zoltan; Kocsis, Istvan	2017
31	10.1126/science.aap8062	Brynjolfsson, Erik; Mitchell, Tom	2017
32	10.3386/w23285	Acemoglu, Daron; Restrepo, Pascual	2017
33	10.1257/pandp.20181019	Brynjolfsson, Erik; Mitchell, Tom; Rock, Daniel	2018

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
34	10.1108/ijoem-02-2017-052	Awdeh, Ali	2018
35	10.2139/ssrn.3112877	Wyonch, Rosalie	2018
36	10.1016/j.jebo.2017.11.014	Shim, Myungkyu; Yang, Hee-Seung	2018
37	10.3390/su10051661	Vermeulen, Ben; Kesselhut, Jan; Pyka, Andreas; Saviotti, Pier	2018
38	10.2139/ssrn.3178233	De Stefano, Valerio	2018
39	10.1007/s00146-017-0736-1	Caruso, Loris	2018
40	10.1177/1094670517752459	Huang, Ming-Hui; Rust, Roland T.	2018
41	10.1111/ntwe.12124	Upchurch, Martin	2018
42	10.1177/1470785318797810	nan	2018
43	10.1257/pandp.20181021	Felten, Edward W.; Raj, Manav; Seamans, Robert	2018
44	10.2139/ssrn.3322306	Ooi, Vincent; Goh, Glendon	2018
45	10.1145/3278721.3278738	Kalyanakrishnan, Shivaram; Panicker, Rahul Alex; Natarajan, Sarayu; Rao, Shreya	2018
46	10.1515/bis-2018-0018	Bruun, Edvard P.G.; Duka, Alban	2018
47	10.2139/ssrn.3286084	Brynjolfsson, Erik; Liu, Meng; Westerman, George F.	2018
48	10.2139/ssrn.3290708	Martens, Bertin; Tolan, Songül	2018
49	10.1257/mac.20150258	Bárány, Zsófia L.; Siegel, Christian	2018
50	10.1017/9781108669016.015	nan	2018
51	10.1017/jmo.2016.55	Brougham, David; Haar, Jarrod	2018
52	10.1108/ejtd-03-2018-0030	Chuang, Szufang; Graham, Carroll Marion	2018
53	10.1016/j.tele.2018.05.013	Garcia-Murillo, Martha; MacInnes, Ian; Bauer, Johannes M.	2018
54	10.1038/s41562-019-0670-y	Granulo, Armin; Fuchs, Christoph; Puntoni, Stefano	2019

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
55	10.1145/3349341.3349438	Astafurova, Olga	2019
56	10.1145/3345252.3345261	Andreeva, Andriyana; Yolova, Galina; Dimitrova, Diana	2019
57	10.1257/jep.33.2.31	Agrawal, Ajay; Gans, Joshua S.; Goldfarb, Avi	2019
58	10.3233/jifs-179127	Wang, Haibo; Li, Hua	2019
59	10.2139/ssrn.3328877	Bessen, James E.; Goos, Maarten; Salomons, Anna; Van den Berge, Wiljan	2019
60	10.3386/w25619	Agrawal, Ajay; Gans, Joshua; Goldfarb, Avi	2019
61	10.31337/oz.74.3.2	González Fabre, Raúl	2019
62	10.1108/jeas-04-2018-0049	Gomes, Orlando; Pereira, Sónia	2019
63	10.5465/ambpp.2019.140	Felten, Edward; Raj, Manav; Seamans, Robert Channing	2019
64	10.1177/1077699019859901	Broussard, Meredith; Diakopoulos, Nicholas; Guzman, Andrea L.; Abebe, Rediet; Dupagne, Michel; Chuan, Ching-Hua	2019
65	10.1016/j.biosystemseng.2019.06.013	Marinoudi, Vasso; Sørensen, Claus G.; Pearson, Simon; Bochtis, Dionysis	2019
66	10.1109/ptc.2019.8810516	Flores, David Rodriguez; Markovic, Uros; Aristidou, Petros; Hug, Gabriela	2019
67	10.1108/978-1-78756-687-320191001	Ivanov, Stanislav; Webster, Craig	2019
68	10.1145/3349341.3349439	Astafurova, Olga; Zapryagaylo, Valeriy; Kulagina, Irina	2019
69	10.1177/2378023119846249	Dahlin, Eric	2019
70	10.1016/j.trpro.2019.07.139	Chinoracký, Roman; Čorejová, Tatiana	2019
71	10.1007/s40812-019-00121-1	Tubaro, Paola; Casilli, Antonio A.	2019
72	10.1080/1331677x.2019.1661788	Arendt, Łukasz; Grabowski, Wojciech	2019
73	10.1086/699935	Agrawal, Ajay; Gans, Joshua; Goldfarb, Avi	2019

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
74	10.1080/17496535.2019.1574253	Bellucci, Sergio	2019
75	10.1016/j.infoecopol.2019.05.003	Colombo, Emilio; Mercorio, Fabio; Mezzanica, Mario	2019
76	10.1111/ntwe.12149	Lloyd, Caroline; Payne, Jonathan	2019
77	10.1145/3325730.3325746	Astafurova, Olga A.; Borisova, Anna S.; Kulagina, Irina I.	2019
78	10.1016/j.procs.2019.09.093	Kurt, Resul	2019
79	10.2139/ssrn.3368605	Felten, Edward W.; Raj, Manav; Seamans, Robert	2019
80	10.1073/pnas.1900949116	Frank, Morgan R.; Autor, David; Bessen, James E.; Brynjolfsson, Erik; Cebrian, Manuel; Deming, David J.; Feldman, Maryann; Groh, Matthew; Lobo, José; Moro, Esteban; Wang, Dashun; Youn, Hyejin; Rahwan, Iyad	2019
81	10.1080/22243534.2019.1689701	Drexler, Nadine; Beckman Lapré, Viyella	2019
82	10.1086/699936	Furman, Jason; Seamans, Robert	2019
83	10.1177/1024258919876416	Todolí-Signes, Adrián	2019
84	10.1080/17538963.2019.1681201	Zhou, Guangsu; Chu, Gaosi; Li, Lixing; Meng, Lingsheng	2020
85	10.1108/er-07-2019-0274	Bejaković, Predrag; Mrnjavac, Željko	2020
86	10.1109/access.2020.3000505	Zhang, Yingying; Xiong, Feng; Xie, Yi; Fan, Xuan; Gu, Haifeng	2020
87	10.1109/acit49673.2020.9208838	Rozum, Daryna; Grazhevskaya, Nadiya; Virchenko, Volodymyr	2020
88	10.1109/icaica50127.2020.9182667	Duan, Chenchen; Wei, Qingjie	2020
89	10.1016/j.labeco.2020.101885	de Vries, Gaaitzen J.; Gentile, Elisabetta; Miroudot, Sébastien; Wacker, Konstantin M.	2020
90	10.1080/21582041.2020.1806346	Rapanyane, M. B.; Sethole, F. R.	2020
91	10.1093/ej/ueaa044	Feng, Andy; Graetz, Georg	2020

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
92	10.1007/s43253-020-00022-3	Bertani, Filippo; Raberto, Marco; Teglio, Andrea	2020
93	10.1093/cjres/rsz022	Acemoglu, Daron; Restrepo, Pascual	2020
94	10.3390/ijfs8030045	Mhlanga, David	2020
95	10.1016/j.eap.2020.07.008	Fatima, Samar; Desouza, Kevin C.; Dawson, Gregory S.	2020
96	10.3390/su12177168	Yu, Jongsik; Ariza-Montes, Antonio; Giorgi, Gabriele; Lee, Aejo; Han, Heesup	2020
97	10.1088/1742-6596/1629/1/012034	Liu, Renbao; Zhan, Yige	2020
98	10.1088/1757-899x/806/1/012004	Pan, Xiaodie; Zhong, Hongsen	2020
99	10.1108/978-1-80043-380-920201006	Erer, Elif; Erer, Deniz	2020
100	10.1016/j.techfore.2020.120302	Lingmont, Derek N.J.; Alexiou, Andreas	2020
101	10.1016/j.jmoneco.2019.01.004	vom Lehn, Christian	2020
102	10.1109/itms51158.2020.9259295	Grodek-Szostak, Zofia; Siguencia, Luis Ochoa; Szelag-Sikora, Anna; Marzano, Gilberto	2020
103	10.1093/cjres/rsz019	Leigh, Nancey Green; Kraft, Benjamin; Lee, Heonyeong	2020
104	10.1016/j.econlet.2020.109032	Gardberg, Malin; Heyman, Fredrik; Norbäck, Pehr-Johan; Persson, Lars	2020
105	10.1186/s12651-020-00275-9	Dauth, Wolfgang; Eppelsheimer, Johann	2020
106	10.1093/inthealth/ihaa007	Hazarika, Indrajit	2020
107	10.1108/978-1-83867-663-620201019	Mahmoud, Ali B.; Tehseen, Shehnaz; Fuxman, Leonora	2020
108	10.1016/j.techfore.2020.120276	Brougham, David; Haar, Jarrod	2020
109	10.1504/ijtg.2020.104905	Jang, Ha Yeon; Lee, Young Min	2020
110	10.1016/j.techsoc.2020.101256	Novakova, Lucia	2020

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
111	10.1145/3375627.3375831	Martínez-Plumed, Fernando; Tolan, Songül; Pesole, Annarosa; Hernández-Orallo, José; Fernández-Macías, Enrique; Gómez, Emilia	2020
112	10.1093/cjres/rsz026	Brooks, Chay; Gherhes, Cristian; Vorley, Tim	2020
113	10.1086/705716	Acemoglu, Daron; Restrepo, Pascual	2020
114	10.4337/roke.2020.01.11	Cashman, Kevin	2020
115	10.1109/iccea50009.2020.00077	Yang, Yu	2020
116	10.1016/j.jbusres.2020.05.019	Rampersad, Giselle	2020
117	10.3390/su12104035	Sima, Violeta; Gheorghe, Ileana Georgiana; Subić, Jonel; Nancu, Dumitru	2020
118	10.1109/icaie50891.2020.00078	Yin, XIA; Junwei, WANG	2020
119	10.1080/10438599.2019.1643557	Nomaler, Önder; Verspagen, Bart	2020
120	10.2478/cejpp-2021-0006	Fatun, Martin; Pazour, Michal	2021
121	10.1093/oxrep/grab012	Crafts, Nicholas	2021
122	10.1016/j.jmacro.2021.103317	Cavenaile, Laurent	2021
123	10.1080/00343404.2021.1928041	Crowley, Frank; Doran, Justin; McCann, Philip	2021
124	10.1016/j.rssm.2021.100623	Haslberger, Matthias	2021
125	10.1016/j.ijhm.2020.102763	Koo, Bonhak; Curtis, Catherine; Ryan, Bill	2021
126	10.21272/mmi.2021.4-12	Jazdauskaite, Jorune; Prívarova, Magdalena; Baranskaite, Edita; Juscius, Vytautas; Kelemen-Henyel, Nikoletta	2021
127	10.3233/faia210170	Virgilio, Gianluca P.M.; Paz López, Manuel Ernesto	2021
128	10.1007/s41347-020-00153-8	Bhargava, Amisha; Bester, Marais; Bolton, Lucy	2021
129	10.1787/7c895724-en	nan	2021
Continued on next page			

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
130	10.1016/j.respol.2020.104064	Falck, Oliver; Heimisch-Roecker, Alexandra; Wiederhold, Simon	2021
131	10.1787/3ed32d94-en	nan	2021
132	10.3390/soc11030093	Illéssy, Miklós; Huszár, Ákos; Makó, Csaba	2021
133	10.1136/bmj.n367	Korinek, Anton; Stiglitz, Joseph E	2021
134	10.31410/limen.2021.61	Sokolic, Danijela, ,	2021
135	10.11648/j.ijefm.20210906.16	Gomez-Mejia, Alberto	2021
136	10.1108/aea-11-2020-0154	Dolado, Juan J.; Felgueroso, Florentino; Jimeno, Juan F.	2021
137	10.1080/1226508x.2020.1867610	Yi, Hye Rim; Shim, Myungkyu; Yang, Hee-Seung	2021
138	10.1016/j.eurocorev.2021.103808	Schmidpeter, Bernhard; Winter-Ebmer, Rudolf	2021
139	10.36689/uhk/hed/2021-01-089	Wu, Chenzi	2021
140	10.1016/j.econmod.2021.01.009	Xie, Mengmeng; Ding, Lin; Xia, Yan; Guo, Jianfeng; Pan, Jiaofeng; Wang, Huijuan	2021
141	10.1016/j.labeco.2021.102002	Alekseeva, Liudmila; Azar, José; Giné, Mireia; Samila, Sampsa; Taska, Bledi	2021
142	10.1080/10438599.2020.1839173	Naudé, Wim	2021
143	10.1177/20539517211003120	Stephany, Fabian	2021
144	10.51593/20200086	Gehlhaus, Diana; , ; Rahkovsky, Ilya	2021
145	10.1613/jair.1.12647	Tolan, Songül; Pesole, Annarosa; Martínez-Plumed, Fernando; Fernández-Macías, Enrique; Hernández-Orallo, José; Gómez, Emilia	2021
146	10.1002/smj.3286	Felten, Edward; Raj, Manav; Seamans, Robert	2021
147	10.2139/ssrn.3957858	Stapleton, Katherine; Copestake, Alex; Pople, Ashley	2021
148	10.1016/j.respol.2021.104289	Ciarli, Tommaso; Kenney, Martin; Massini, Silvia; Piscitello, Lucia	2021
149	10.14488/bjopm.2021.010	Reis, João; Santo, Paula Espírito; Melão, Nuno	2021

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
150	10.32996/jefas.2021.3.2.13	Bonsay, Jamielyn; Cruz, Abigail P.; Firozi, Homa C.; Camaro, And Peter Jeff C.	2021
151	10.24818/basiq/2021/07/018	Bănescu, Carmen-Elena; , ; Țițan, Emilia; Manea, Daniela	2021
152	10.1787/2e278150-en	nan	2021
153	10.1108/ejtd-10-2019-0183	Chuang, Szufang	2021
154	10.1080/13678868.2020.1818513	Su, Zhan; Togay, Guillaume; Côté, Anne-Marie	2021
155	10.1007/s41027-021-00340-y	Oware, Kofi Mintah; Mallikarjunappa, Thathiah	2021
156	10.1177/2041386621992105	Alcover, Carlos-María; Guglielmi, Dina; Depolo, Marco; Mazzetti, Greta	2021
157	10.1016/j.jbusres.2019.09.019	Fossen, Frank M.; Sorgner, Alina	2021
158	10.1016/j.respol.2020.104002	Goos, Maarten; Rademakers, Emilie; Röttger, Ronja	2021
159	10.1016/j.respol.2021.104269	Foster-McGregor, Neil; Nomaler, Önder; Verspagen, Bart	2021
160	10.3390/su13147703	Szabó-Szentgróti, Gábor; Végyvári, Bence; Varga, József	2021
161	10.54204/taji/vol41202201	Laura, Claudia; , ; Choi, Chaem; ,	2022
162	10.1590/0101-31572022-3320	ACYPRESTE, RAFAEL DE; PARANÁ, EDEMILSON	2022
163	10.21916/mlr.2022.21	Handel, Michael	2022
164	10.1186/s12651-022-00319-2	Gries, Thomas; Naudé, Wim	2022
165	10.1108/manm-02-2022-0034	Mukherjee, Arunava Narayan	2022
166	10.1007/s40888-022-00272-w	Hidalgo-Pérez, Manuel A.; Molinari, Benedetto	2022
167	10.3389/frai.2022.832736	Georgieff, Alexandre; Hyee, Raphaela	2022
168	10.7202/1094209ar	Perez, Fabienne; Conway, Neil; Roques, Olivier	2022

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
169	10.1080/09585192.2021.1891114	Jaiswal, Akanksha; Arun, C. Joe; Varma, Arup	2022
170	10.4324/9781003275534-5	Furusawa, Taiji; Kusaka, Shoki; Sugita, Yoichi	2022
171	10.1111/ntwe.12215	Holm, Jacob Rubæk; Lorenz, Edward	2022
172	10.1108/rep-11-2019-0145	Aly, Heidi	2022
173	10.2139/ssrn.4165159	Li, Boshuo; Huang, Ni; Shi, Wei	2022
174	10.17700/jai.2022.13.1.638	Munnisunker, Shivaan; Nel, Lyndre; Diederichs, Dawid	2022
175	10.30525/2500-946x/2022-3-6	Cheromukhina, Olha	2022
176	10.3386/w30074	Autor, David	2022
177	10.1126/scirobotics.abg5561	Paolillo, Antonio; Colella, Fabrizio; Nosengo, Nicola; Schiano, Fabrizio; Stewart, William; Zambrano, Davide; Chappuis, Isabelle; Lalive, Rafael; Floreano, Dario	2022
178	10.1108/ITP-06-2019-0304	Goethals, Frank; Ziegelmayr, Jennifer L.	2022
179	10.1007/s42001-021-00134-8	Chen, Haohui Caron; Li, Xun; Frank, Morgan; Qin, Xiaozhen; Xu, Weipan; Cebrian, Manuel; Rahwan, Iyad	2022
180	10.1371/journal.pone.0277280	Chen, Ni; Li, Zhi; Tang, Bo	2022
181	10.1080/09537287.2021.1882692	Braganza, Ashley; Chen, Weifeng; Canhoto, Ana; Sap, Serap	2022
182	10.1016/j.respol.2021.104446	Duch-Brown, Néstor; Gomez-Herrera, Estrella; Mueller-Langer, Frank; Tolan, Songül	2022
183	10.1016/j.jrurstud.2021.12.012	Rijnks, Richard Henry; Crowley, Frank; Doran, Justin	2022
184	10.1016/j.labeco.2022.102146	Guimarães, Luís; Mazedo Gil, Pedro	2022
185	10.3390/informatics9010019	Gonçalves, Maria José Angélico; da Silva, Amélia Cristina Ferreira; Ferreira, Carina Gonçalves	2022

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
186	10.3390/app12126014	Stoica (Răpan), Ivona; Zaman, Gheorghe; Suciu, Marta-Christina; Purcărea, Victor-Lorin; Jude, Cornelia-Rodica; Radu, Andra-Victoria; Catană, Aida; Radu, Anamaria-Cătălina	2022
187	10.1016/j.procs.2022.04.029	Yan-ping, Li; Ai-qin, Qi	2022
188	10.1007/s11205-021-02768-7	Fernández-Macías, Enrique; Bisello, Martina	2022
189	10.1007/s00146-022-01496-x	Deranty, Jean-Philippe; Corbin, Thomas	2022
190	10.1155/2022/7406716	Sen, Wang; Xiaomei, Zhu; Lin, Deng	2022
191	10.1080/07421222.2022.2096542	Xue, Mei; Cao, Xing; Feng, Xu; Gu, Bin; Zhang, Yongjie	2022
192	10.1080/09585192.2020.1871398	Vrontis, Demetris; Christofi, Michael; Pereira, Vijay; Tarba, Shlomo; Makrides, Anna; Trichina, Eleni	2022
193	10.1080/08839514.2022.2080336	Gu, Ting-Ting; Zhang, San-Feng; Cai, Rongrong	2022
194	10.2196/preprints.38558	Al-Ansari, Ahmed Malalla; Al-Medfa, Mohammed Khalid	2022
195	10.1177/0950422221990990	Verma, Amit; Lamsal, Kamal; Verma, Payal	2022
196	10.2139/ssrn.4060233	Babina, Tania; Fedyk, Anastassia; He, Alex Xi; Hodson, James	2022
197	10.1504/ijatm.2022.124366	Moniz, António B.; Candeias, Marta; Boavida, Nuno	2022
198	10.3389/fsoc.2022.959091	Bousrih, Jihen; Elhaj, Manal; Hassan, Fatma	2022
199	10.20965/jaciii.2022.p0655	Shen, Yang; , ; Zhang, Xiuwu	2022
200	10.2478/tjeb-2022-0006	Cismas, Laura Mariana; Dumitru, Cornelia; Negrut, Teodor	2022
201	10.1371/journal.pone.0263704	Phiromswad, Piyachart; Srivannaboon, Sabin; Sarajoti, Pattarake	2022
Continued on next page			

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
202	10.20944/preprints202309.0193.v1	Tachicart, Ridouane	2023
203	10.2139/ssrn.4385988	Faioli, Michele	2023
204	10.1007/s10639-023-12203-8	Hetmańczyk, Piotr	2023
205	10.1177/10242589221147228	Hassel, Anke; Özkiziltan, Didem	2023
206	10.1016/j.jik.2023.100407	Ključnikov, Aleksandr; Popkova, Elena G.; Sergi, Bruno S.	2023
207	10.15240/tul/009/lef-2023-24	Šášky, Michael	2023
208	10.1057/s11369-023-00337-z	Bresnahan, Timothy F.	2023
209	10.1007/s10799-023-00408-9	Wang, Haonan; Qiu, Fangjuan	2023
210	10.1038/s41598-023-28874-9	Scher, Sebastian; Kopeinik, Simone; Trügler, Andreas; Kowald, Dominik	2023
211	10.2139/ssrn.4366560	Reljic, Jelena; Cirillo, Valeria; Guarascio, Dario	2023
212	10.1007/s00191-023-00809-7	Carbonero, Francesco; Davies, Jeremy; Ernst, Ekkehard; Fossen, Frank M.; Samaan, Daniel; Sorgner, Alina	2023
213	10.22371/05.2023.023	, ; Call, Greg	2023
214	10.54394/fhem8239	Gmyrek, Pawel; Berg, Janine; Bescond, David;	2023
215	10.2139/ssrn.4412505	Campello de Souza, Bruno; Andrade Neto, Agostinho Serrano de; Roazzi, Antonio	2023
216	10.2139/ssrn.4337611	Bell, Stephanie	2023
217	10.3390/jintelligence11100194	Dumitru, Daniela; Halpern, Diane F.	2023
218	10.4018/jdm.318455	Peng, Gang; Bhaskar, Rahul	2023
219	10.1016/j.hrmr.2021.100857	Pereira, Vijay; Hadjielias, Elias; Christofi, Michael; Vrontis, Demetris	2023
220	10.1016/j.labeco.2023.102456	Javed, Mohsin	2023

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
221	10.6007/ijarbss/v13-i3/16546	Yusoff, Yusri Hazrol; Johari, Anis Syahirah; Mohd Rahmatullah, Dina Adriana; Zainal, Nur Amirah; Tajuddin, Nur Atiqah; Thilaiampalam, N. T. Shanmuganathan	2023
222	10.28945/5078	Morandini, Sofia; Fraboni, Federico; De Angelis, Marco; Puzzo, Gabriele; Giusino, Davide; Pietrantoni, Luca	2023
223	10.1111/ntwe.12240	Howcroft, Debra; Taylor, Phil	2023
224	10.5089/9798400254802.001	Pizzinelli, Carlo	2023
225	10.36348/sijlcj.2023.v06i10.001	Nnamdi, Nmesoma; Ogunlade, Babafemi Zachaeus; Abegunde, Babalola	2023
226	10.1016/j.wds.2023.100107	Guliyev, Hasraddin; Huseynov, Natiq; Nuriyev, Nasimi	2023
227	10.32388/3bwnxg	Ekwueme, Francis Okechukwu; Areji, Anthony C.; Ugwu, Anayochukwu	2023
228	10.2139/ssrn.4534294	Huang, Qiwen; Shen, Yifan; Sun, Yuanchi; Zhang, Qingquan	2023
229	10.26855/acc.2023.06.006	u, Yulin; Meng, Xiangtao; Li, Anqi	2023
230	10.36948/ijfmr.2023.v05i03.3133	-, Subharun Pal	2023
231	10.1007/s43681-023-00263-y	Gomes, Orlando	2023
232	10.2139/ssrn.4527336	Hui, Xiang; Reshef, Oren; Zhou, Luofeng	2023
233	10.2991/978-94-6463-142-5/37	Yue, Qiran	2023
234	10.5937/turpos0-43739	Štilić, Anđelka; Nicić, Miloš; Puška, Adis	2023
235	10.1007/s10663-023-09571-2	Lorenz, Hanno; Stephany, Fabian; Kluge, Jan	2023
236	10.1007/s00168-023-01234-1	Yang, Seongjun; Kim, Donghyun	2023
237	10.32388/4hasum	Biswas, Som	2023
238	10.1007/s12122-023-09346-5	Jacobs, Arthur; Verhofstadt, Elsy; Van Ootegem, Luc	2023
239	10.2139/ssrn.4339329	van der Kooij, Bouke J.G.	2023

Continued on next page

Table A.1 – continued from previous page

ID	DOI	Author(s)	Published
240	10.2139/ssrn.4592960	Caselli, Mauro; Fracasso, Andrea; Marcolin, Arianna; Scicchitano, Sergio	2023
241	10.2139/ssrn.4350925	Zarifhonarvar, Ali	2023
242	10.58445/rars.337	Pillai, Raaghav	2023
243	10.2139/ssrn.4529739	Liu, Jin; Xu, Xingchen (Cedric); Li, Yongjun; Tan, Yong	2023
244	10.2139/ssrn.4535534	Tiku, Sarvesh	2023
245	10.1016/j.heliyon.2023.e18670	Poláková, Michaela; Suleimanová, Juliet Horváthová; Madzík, Peter; Copuš, Lukáš; Molnárová, Ivana; Polednová, Jana	2023
246	10.3390/systems11030114	Li, Chao; Zhang, Yuhang; Niu, Xiaoru; Chen, Feier; Zhou, Hongyan	2023
247	10.4108/eetsis.3841	Ruiz-Talavera, Doris; Cruz-Aguero, Jaime Enrique De la; García-Palomino, Nereo; Calderón-Espinoza, Renzo; Marín-Rodríguez, William Joel	2023
248	10.1111/1748-8583.12524	Budhwar, Pawan; Chowdhury, Soumyadeb; Wood, Geoffrey; Aguinis, Herman; Bamber, Greg J.; Beltran, Jose R.; Boselie, Paul; Lee Cooke, Fang; Decker, Stephanie; DeNisi, Angelo; Dey, Prasanta Kumar; Guest, David; Knoblich, Andrew J.; Malik, Ashish; Paauwe, Jaap; Papagiannidis, Savvas; Patel, Charmi; Pereira, Vijay; Ren, Shuang; Rogelberg, Steven; Saunders, Mark N. K.; Tung, Rosalie L.; Varma, Arup	2023
249	10.1016/j.ceqi.2023.06.001	Du, Yang; Jia, Peng; Park, Albert	2023
250	10.47772/ijriss.2023.7487	Yeboah, Frederick Forkuo	2023

Source: Author's data

Chapter B

STRUCTURE OF THE INTERVIEWS

(OUTLINE)

The interviews followed the following structural outline:

B.0.0.1 Introduction (5 minutes)

B.0.0.1.1 Background and Experience (3 minutes):

- Could you describe how you entered the field of programming and what your current role is?
- How long have you been working in this capacity and what are your main responsibilities?

B.0.0.1.2 Work Environment (2 minutes):

- Could you briefly describe the nature of your company or freelance work, focusing on its methodological and technological maturity?

B.0.0.2 Main Section (45 minutes)

B.0.0.2.1 Integration of AI Solutions into Daily Work (10 minutes):

- How are AI technologies currently applied in your daily work? Could you provide examples?
- Which specific tasks have improved or changed due to AI integration?

B.0.0.2.2 Automation Potential (10 minutes):

- Based on your experience, which work areas do you see the greatest promise for AI-based automation?
- How do you see AI's impact on these areas in the near future?

B.0.0.2.3 Analysis of Schedule and Tasks (10 minutes):

- Could you outline a typical day, focusing on the types of tasks you perform?
- For each type of task, could you estimate the potential for AI-based automation?
- What would this mean for your role?

B.0.0.2.4 Risks and Problems Associated with AI Solutions (5 minutes):

- In your opinion, what is the biggest risk or problem with using AI in your work?

B.0.0.3 Impact on Career and Labor Market (10 minutes):

- How do you see the broader impact of AI on programmers' job opportunities and working conditions?
- What do you think about the broader societal impact of AI solutions, especially in the near future (up to 3 years)?

B.0.0.4 Closing (10 minutes)**B.0.0.4.1 Clarification and Further Insights (5 minutes):**

- Is there anything you would like to clarify or add to your previous answers?
- Do you have any personal anecdotes or experiences with AI that you find relevant to this discussion?

B.0.0.4.2 Closing Thoughts (5 minutes):

- What advice would you give to a recent graduate, junior programmers regarding AI, considering our discussion?
- How do you plan to adapt or change your skills in line with the development of AI in programming?

Chapter C

DEMOGRAPHICS DATA OF RESPONDENTS

Identifier	Company Size	Role	Age
B.A.	10–100	Backend, ML	32
B.G.	>1000	ML	33
B.N.	>1000	ML	34
B.T.	>1000	Full stack	49
C.T.	>1000	Backend, ML	43
C.T.	10–100	Backend, ML	28
D.S.	1–10	Backend	35
D.T.	Freelancer	Full stack	39
F.Á.	10–100	Backend, ML	32
F.P.	100–500	Frontend	38
H.A.	100–500	Backend	29
H.J.	100–500	Tester, Backend	42
I.T.	10–100	Backend	27
L.A.	1–10	Backend	34
M.K.	100–500	Backend	24
N.G.	100–500	Backend	41
P.C.	100–500	ML	33
S.A.	Freelancer	ML	53
S.B.	10–100	Backend	43
S.Z.	>1000	Backend, ML	62

Table C.1: Roles, Company Sizes, and Ages of Interviewees

Source: Author's data

Chapter D

STACK OVERFLOW 2023

HUNGARIAN SENIOR

RESPONDENTS' OPINIONS COUNT

	Very Similar	Somewhat Similar	Neither Different nor Similar	Somewhat Different	Very Different
Understanding Codebase	3	5	7	32	13
Project Planning	3	7	3	19	7
Writing Code	17	25	26	105	38
Code Documentation	2	5	5	30	17
Debugging and Assistance	8	12	13	54	30
Code Testing	5	5	10	24	17
Code Quality Assurance	0	1	4	12	5
Deployment and Monitoring	2	0	5	3	5
Collaboration with Teammates	0	0	2	2	0
Other (Free Text)	0	0	0	0	0

Table D.1: Assessment of Software Development Tasks by Similarity

Source: Stack Overflow 2023 Developer Survey

Chapter E

SPECIFIC QUESTIONS OF THE SURVEY

Section 1: AI and Your Organization

1. Country of Operations

Please write the name of the country where your organization operates.

2. Your Role in the Organization

Please select the role that best describes your position within the organization:

- IT
- Finance
- Human Resources
- Operations
- Sales
- Marketing
- Research and Development
- Executive Management
- Research
- Executive Management
- Don't know / Don't want to answer
- Other

3. Organization Size

Please select the size of your organization based on employee count:

- Micro (1-10 employees)

- Small (11-50 employees)
- Medium (51-250 employees)
- Large (251+ employees)
- Don't know / Don't want to answer

4. Organization Industry

Please select the industry your organization operates within:

- Technology
- Finance
- Healthcare
- Manufacturing
- Education
- Retail
- Don't know / Don't want to answer
- Other

Section 2: Awareness of Artificial Intelligence (AI)

5. Experience in interacting with AI in professional life

How experienced are you in interacting with AI in your professional life?

- Not experienced at all
- Slightly experienced
- Moderately experienced
- Very experienced
- Extremely experienced

6. Frequency of AI Use

How frequently do you or your organization use AI technologies in your professional activities?

- Never
- Rarely
- Occasionally
- Frequently

- Always

7. Experiences with Inappropriate AI Behaviors

Have you experienced or heard about inappropriate behaviours from AI models in a professional context?

- Never
- Rarely
- Occasionally
- Frequently
- Always

8. Current AI Use Cases

What are the top 5 use cases for AI within your organization? Please list them.

9. Business Use Cases for AI

For each scenario below, indicate if your organization currently uses AI and to what extent:

- Internal Use Case: Accessing internal policies
 - Not used
 - Minimally used
 - Moderately used
 - Extensively used
 - Core function
- Customer-Facing Use Case: Interactive access to FAQs
 - Not used
 - Minimally used
 - Moderately used
 - Extensively used
 - Core function
- Customer-Facing Use Case: Sales chatbot
 - Not used
 - Minimally used
 - Moderately used
 - Extensively used

- Core function

10. **Planned Business Use Cases for AI**

For each scenario below, indicate your organization's intended level of AI integration:

- Internal Use Case: Enhancing internal policy access
 - Not planned
 - Planned as a minor feature
 - Planned as a moderate feature
 - Planned as a major feature
 - Planned as a core function
- Customer-Facing Use Case: Enhancing interactive FAQ access
 - Not planned
 - Planned as a minor feature
 - Planned as a moderate feature
 - Planned as a major feature
 - Planned as a core function
- Customer-Facing Use Case: Enhancing sales chatbot functionality
 - Not planned
 - Planned as a minor feature
 - Planned as a moderate feature
 - Planned as a major feature
 - Planned as a core function

Section 3: Risks and Challenges of LLMs in Business

11. **Potential harm from LLMs in business**

In your view, what are the potential risks or ways that Large Language Models (LLMs) like GPT could cause harm or pose challenges in a business setting? Please select all that apply, and feel free to add any additional concerns.

- Data privacy breaches
- Spread of misinformation
- Bias in decision-making processes
- Job displacement due to automation
- Security vulnerabilities

- Misinterpretation of complex queries
- Legal and compliance issues
- Erosion of human skills
- Other (please specify)

Section 4: Assessment of AI Trust

12. AI Initiative Budget

If your company has AI initiatives, how much do you plan to spend on these initiatives annually? (Specify in Euros)

13. Company Competence in AI Safety

How confident are you in your company's competence in ensuring the safety and ethical use of AI technologies?

- Not confident
- Slightly confident
- Moderately confident
- Very confident
- Extremely confident

14. Value of an LLM Assessment Tool

How valuable would you find a third-party assessment tool for ensuring the appropriate behavior of your AI systems, including LLMs?

- Not valuable
- Slightly valuable
- Moderately valuable
- Very valuable
- Extremely valuable

15. Willingness to Pay for AI Assessment

How much would you be willing to pay annually for a third-party service that assesses and ensures the appropriate behavior of your AI systems? (Please specify in Euros. Write 0 if you are not willing to pay.)

Chapter F

REGULAR EXPRESSION FOR DETECTING TECHNOLOGICAL TERMS

The following regular expression (RegEx) aims to identify technological terms occurring in different languages (English, German, Hungarian) in textual content. The expression supports different spelling variations of certain terms and also considers certain grammatical suffixes in the case of the Hungarian language. The following expression was used in Python to detect technological keywords and was applied to the dataset using the efficient execution tools of the Pandas framework. The runtime remained within approximately 5 minutes.

```
pattern = r"""\b(AI|  
Artificial\s+Intelligence|  
Machine\s+Learning|  
Computer\s+Vision|  
Large\s+Language\s+Model(s)?|  
Vision\s+Systems|  
ChatGPT|  
Natural\s+Language\s+Processing|  
Data\s+Science|  
Data\s+Sciences|  
Künstliche\s+Intelligenz|  
Maschinelles\s+Lernen|  
Computersehen|  
Große\s+Sprachmodelle|  
Natürliche\s+Sprachverarbeitung|  
Sprachverarbeitung|  
Datenwissenschaft|  
Datenwissenschaften|
```

```
mesterséges\s+intelligencia|
gépi\s+tanulás|
számítógépes\s+látás|
gépi\s+látás|
vision\s+rendszer(ek)?|
nagy\s+nyelvi\s+modell(ek)?|
természetes\s+nyelv(i?)\s+feldolgozás|
nyelv(i?)\s+feldolgozás|
nyelvfeldolgozás|
adattudomány(ok)?|
adat\s+tudomány(ok)?|

Aritficial\s+Intelligence|
Machin\s+Lerning|
NLP|
LLMs?
)\b(?:val|vel|ban|ben|nak|nek|
bó|ből|hoz|hez|höz|
tól|től|ra|re|al|el|s(ok)?"""
```

This regular expression has the following properties:

- It can recognize relevant terms in multiple languages (English, German, Hungarian).
- It considers spelling variations and abbreviations specific to each language (e.g., AI, NLP, LLMs).
- In the case of the Hungarian language, it recognizes inflected forms related to technological terms.
- Optimized for use in Python, applicable with `re.IGNORECASE` and `re.VERBOSE` flags.

The regular expression was successfully applied to various texts to detect technological terms, enabling their systematic identification and analysis.

Chapter G

DSPY OPTIMIZED PROMPT FOR SEMANTIC CLASSIFICATION

The final prompt resulting from the optimization procedure, which achieved flawless performance on the test data, was as follows:

```
Conduct a thorough examination of the provided job description
to determine its association with AI and ML.
Classify it as "True" if any explicit mentions
are noted regarding fields such as Artificial Intelligence,
Machine Learning, Data Science, Deep Learning,
or related concepts such as
Neural Networks, Natural Language Processing.
Also consider subtle hints or phrasing
that suggest involvement with such technologies,
and be mindful of variations in language from across the globe.
If these elements are absent, unequivocally
mark it as "False." Please respond with only
the two words "True" or "False," maintaining
accuracy and clarity while avoiding additional commentary.
```

DECLARATION ON IDENTITY

I, the undersigned Levente Szabados, declare that the printed and electronic versions of the doctoral dissertation and thesis booklet are identical in all respects .

Sopron, 20 25 year July month 5 day



signature of PhD candidate

LEGAL DECLARATION

I, the undersignedLevente Szabados....., by signing this declaration declare that ~~AI at work: promises, perils and paradoxes of adoption~~..... my PhD dissertation was my own work; during the dissertation I complied with the Act LXXVI of 1999 on the rules of copyright and the rules of the doctoral dissertation prescribed by the Doctoral School, especially regarding references and citations.¹

Furthermore, I declare that I did not mislead the supervisor (s) or the program leader with the dissertation.

By signing this declaration, I acknowledge that if it can be proved that the dissertation is not self-made or the author of a copyright infringement is related to the dissertation, the University of Sopron is entitled to refuse the acceptance of the dissertation.

I further declare that I am not in the process of obtaining a doctoral degree in the same discipline, that I am not in the process of having my doctoral degree revoked, and that I have not had a previously awarded doctoral degree revoked within 5 years.

Refusing to accept a dissertation does not affect any other (civil, legal, criminal) consequences of copyright infringement.

Sopron, 20 25 year July month 05 day



signature of PhD candidate

¹ LXXVI. TV. 1999 Section 34 (1) Any person may quote the details of the work, to the extent justified by the nature and purpose of the receiving work and in the original, by the name of the source and the author designated there.

Article 36 (1) Details of public lectures and other similar works, as well as political speeches, may be freely used for information purposes, within the scope justified by the purpose. For such use, the source, along with the author's name, should be indicated, unless this is impossible.

CO-AUTHOR'S DECLARATION

(each article must be declared on a separate page)

I/We, the undersigned Ádám Buza co-author(s) agree that
..... Levente Szabados doctoral student may use the results of the joint
publication titled "Trust My AI" in his/her
doctoral dissertation titled AI at work: promises, perils and paradoxes of adoption

Sopron, 20 25 year July month 05 day



signature of co-author(s)