Right Sizing AI at Work:

Workforce Transformation and Human–Machine Partnerships in Ireland

CONTENTS

Artificial Intelligence is reshaping work in Ireland primarily through changes in tasks, work-flows, skills, and organisational practices rather than through immediate large-scale job displacement.


This report examines workforce transformation through a mixed-methods design combining a structured literature review, a cross-sector workforce survey of 487 workers, embedded behavioural vignette experiments involving 421 participants, 32 semi-structured interviews with senior organisational leaders, and comparative analysis of 48 policy documents from Ireland, the European Union, and the United Kingdom.

Foreword

Artificial Intelligence is rapidly emerging as one of the most transformative forces shaping the future of work, business competitiveness and economic growth. For Ireland, this represents both an opportunity and an imperative. Our future competitiveness will depend increasingly on our ability to harness the potential of AI while ensuring that workers remain at the centre of transformation.

As this report demonstrates, the most significant impact of AI is not the immediate replacement of jobs, but the reshaping of work itself. Increasingly, people are working alongside intelligent systems, using AI to enhance productivity, improve decision-making and unlock innovation. The organisations that thrive will be those that successfully combine technological innovation with investment in talent and leadership.

The report also reinforces a critical lesson: technology alone does not drive productivity. Lasting impact comes when AI adoption is supported by upskilling effective governance, quality data and redesigned workflows. This is particularly important for SMEs, which form the backbone of the Irish economy and must be supported to fully participate in the opportunities presented by AI.

Meeting this challenge requires collaboration. Industry has a vital role in helping to shape the learning and development solutions that will prepare workers for the future. Employers are uniquely positioned to identify emerging skills needs and, working alongside Government and the education and training sector, can help ensure that learning provision remains relevant, agile and aligned with workplace realities.

At Skillnet Ireland this commitment to industry-led talent development is at the heart of everything we do. Through close engagement with enterprise and alignment with national policy priorities, we are investing in the next generation of skills and capabilities needed to succeed in a rapidly changing world. This report provides important insights to guide that work and to support Ireland's ambition to build a highly skilled, innovative and AI-ready workforce.

I would like to acknowledge all of those who contributed to this research report. Particular thanks are due to the many enterprises, employees and stakeholders who contributed their time and efforts. I would also like to express my thanks to Professor Taha Yasser at the Centre for Sociology of Humans and Machine at Trinity College and the team at Workday. Finally, I would like to thank the Technology Ireland DIGITAL Skillnet for leading this innovative study to a successful conclusion.

Mark Jordan
Chief Executive Officer
Skillnet Ireland

Artificial Intelligence represents a profound shift in the way organisations create value. Like previous technologies, its true significance lies not in the technology itself but in its ability to transform business processes and create entirely new opportunities for growth. For Irish businesses, AI is quickly moving from an emerging technology to a strategic capability.

The evidence presented in this report arrives at an important moment. Across industries, leaders are actively exploring how AI can improve performance, strengthen customer experiences, enhance decision-making and unlock new sources of productivity. Yet the research also confirms that the journey from experimentation to impact is far from straightforward. The organisations achieving the greatest returns are not simply deploying AI tools- they are fundamentally rethinking how work is organised, how decisions are made and how talent is utilised.

This distinction is significant. AI is not merely a technology adoption challenge- it is an organisational transformation challenge. Successfully integrating AI requires businesses to redesign processes, develop trust in human-machine collaboration and support staff with the capabilities needed to adapt as technologies evolve.

For Ireland, this presents a strategic opportunity. Our reputation as a global technology hub, combined with a highly educated workforce and a dynamic enterprise base, positions us well within an increasingly AI-driven global economy. However, maintaining that position will require sustained investment in both technological capability and the broader skills system that supports innovation.

The findings of this report underline too the growing importance of a dual capability agenda. Businesses need access to advanced AI expertise, but they also require the transversal skills that enable organisations to realise real value from the technologies. Critical thinking, problem solving, creativity, communication and sound judgement are increasingly becoming competitive assets in their own right. As AI takes on more routine and analytical tasks, these uniquely human capabilities will become even more important in driving innovation and growth.

This is a priority shared across Europe. Through our engagement with Digital Europe and with our European partners, we see a growing recognition that competitiveness will increasingly depend on the ability of businesses to adopt and scale AI effectively. The challenge is no longer simply access to technology. It is ensuring that organisations of all sizes have the skills and organisational readiness to deploy AI in ways that deliver measurable business value.

For many companies, particularly SMEs, this remains a complex undertaking. While the opportunities are significant, successful implementation requires access to expertise and trusted pathways for capability development. Building these pathways will be essential if Ireland is to maximise the economic and societal benefits of AI.

It is in this context that the work of Technology Ireland DIGITAL Skillnet continues to make a substantial contribution. Through a long-standing commitment to industry-informed innovation, DIGITAL Skillnet has consistently anticipated emerging capability needs and translated them into practical learning programmes for businesses. Its ability to convene companies, academic partners and technology experts has created a powerful platform for accelerating digital transformation across the Irish tech sector.

Particularly noteworthy has been DIGITAL Skillnet's leadership in developing new initiatives focused on AI readiness and adoption. Programmes such as the AI Accelerator are helping organisations move beyond awareness and exploration towards practical implementation, providing enterprises with the knowledge, tools and confidence required to embed AI effectively within their operations.

I would like to thank Professor Taha Yasseri and his team at the Centre for the Sociology of Humans and Machines, Workday, the Technology Ireland DIGITAL Skillnet team and the many employers, employees, policymakers and experts who contributed to this research. Their collective insights have produced a valuable resource that will help inform future action across Ireland's business and skills landscape.

Una Fitzpatrick
Director
Technology Ireland, Ibec

Abstract

Across the evidence base, AI is found to automate selected routine activities, augment professional work, and redistribute responsibilities between people and intelligent systems. Workers are increasingly required to review, verify, coordinate, and exercise judgement over AI-supported outputs.

> The findings indicate that the principal skills challenge is not basic awareness of AI, but the development of practical, role-specific capabilities combining technical literacy with critical evaluation, adaptability, communication, and accountability.

The study also finds that productivity gains are not produced by technology alone. They depend on the **redesign of work-flows and complementary investments in leadership,** workforce development, data quality, secure infrastructure, and clear organisational governance. Adoption is uneven: larger and digitally mature organisations are generally better positioned to scale AI, while SMEs face greater constraints in finance, expertise, training, and implementation capacity.

Comparative policy analysis suggests that Ireland is relatively well positioned internationally through a whole-of-system approach linking education, skills development, and employer engagement. The report concludes, however, that **Ireland's central challenge is to move from fragmented tool adoption towards deliberately designed human–machine partnerships.** It proposes a right-sizing framework focused on task and workflow fit, human capability, organisational readiness, inclusion, and accountability, with practical recommendations for employers, SMEs, education and training providers, and policymakers.

Research Team

The Centre for the Sociology of Humans and Machines (SOHAM) is a joint research centre at Trinity College Dublin and TU Dublin, dedicated to understanding how AI and automation are reshaping society. Through interdisciplinary research, international collaboration, and public engagement, SOHAM investigates the social, ethical, and technological dimensions of human-machine interactions.

The research team was led by Taha Yasseri who is the Workday Full Professor and Chair of Technology and Society at Trinity College Dublin and Technological University Dublin. He is also an adjunct Full Professor at the School of Mathematics and Statistics at University College Dublin. Taha was a Professor and the Deputy Head at the School of Sociology and a Geary Fellow at the Geary Institute for Public Policy at University College Dublin, Ireland. Previously, Taha was a Senior Research Fellow in Computational Social Science at the University of Oxford, a Turing Fellow at the Alan Turing Institute for Data Science and Artificial Intelligence, and a Research Fellow in Humanities and Social Sciences at Wolfson College. Taha has a PhD in Complex Systems Physics from the University of Göttingen, Germany.

He was supported by a multidisciplinary team comprising Dr Marina Schenkel as Lead Researcher, Wagner Guilherme Alves da Silva as Qualitative Researcher, Dr Anna Bertani and Dr Hao Cui as Quantitative and Behavioural Scientists, and Dr Ruiqing Han as Psychological and Experimental Researcher. Administrative support for the project was provided by Helen Stanton.

Introduction

01 Introduction

Artificial intelligence is no longer primarily a future prospect for the Irish workforce.

Generative AI tools are already being used across occupations and organisational functions, while emerging agentic systems are beginning to move beyond content generation towards the execution of actions across digital environments. Recent evidence from 177,436 publicly available tools for AI agents shows that such systems are increasingly being **designed not only to retrieve and analyse information, but also to modify files, communicate through external platforms and perform other actions within operational environments.** This development expands the relevant question from whether workers use AI tools to how responsibilities, decisions and sequences of work are distributed between people and intelligent systems (Stein, 2026).

Public debate about this transformation is frequently organised around two competing narratives. **One anticipates large-scale displacement of human labour;** the other presents AI as a broadly **productivity-enhancing technology that will augment workers** and create new opportunities. Both capture important possibilities, but neither adequately describes the uneven and organisationally mediated nature of the changes currently taking place. Recent UK labour-market analysis finds no detectable aggregate employment effect in occupations most exposed to AI since the release of widely available generative AI systems. At the same time, it finds substantial variation among occupations with similar estimated exposure and shows that actual AI use remains concentrated in a relatively narrow set of tasks. This suggests that technical exposure alone cannot determine whether AI will complement, restructure or replace work in a particular setting (Serôdio, 2026).

Forward-looking assessments nonetheless indicate that substantial labour-market change is plausible. The World Economic Forum's Future of Jobs Report 2025, based on the expectations of more than 1,000 major employers, estimates that the combined creation and displacement of jobs associated with technological, economic, demographic and environmental trends could affect 22% of current formal employment by 2030.

Employers anticipate 170 million new roles and 92 million displaced roles, alongside changes to approximately 39% of existing skill sets.

These figures **reflect employer expectations rather than observed labour-market outcomes,** but they demonstrate the scale of organisational change for which firms and institutions are preparing. The same survey identifies skills shortages as the most frequently reported barrier to business transformation and highlights the **continuing importance of analytical thinking, creativity, resilience and adaptability alongside technical AI capabilities** (World Economic Forum, 2025).

The potential consequences are also distributional. Recent modelling by the Economic and Social Research Institute **examines a range of short- to medium-term scenarios for Ireland.** In its central scenario, approximately 7% of current jobs are displaced, while workers who remain employed experience modest productivity-related wage gains. The model produces an average reduction in household disposable income and a moderate increase in income inequality because employment losses are not fully offset by wage gains and increases in capital income. These are simulated scenarios derived from international evidence, not predictions of an inevitable outcome.

They nevertheless underline that aggregate productivity growth does not guarantee that the benefits and costs of AI adoption will be distributed evenly across workers or households (Doorley et al., 2026).

In Ireland, the immediate challenge is increasingly one of organisational maturity rather than simple access to AI.

A 2026 survey of 250 Irish organisations found that day-to-day use had become widespread, but that many organisations remained at an early stage in translating individual use into sustained operational and strategic transformation. The report identifies a gap between employees' adoption of AI and organisations' redesign of workflows, processes and decision-making structures. It concludes that future value will depend not solely on whether AI is adopted, but on investments in skills, governance and the redesign of how work is organised (Trinity College Dublin, 2026).

This gap is **especially consequential for small and medium-sized enterprises.** SMEs may possess shorter decision chains and greater flexibility to test new practices, but they generally have fewer **resources for data infrastructure,** specialist expertise, training, governance and regulatory compliance. The Irish evidence therefore presents a more complex picture than a simple divide between agile small firms and capable large organisations. **SMEs that invest seriously in AI can report substantial gains, but many remain concentrated at earlier stages of adoption, while larger organisations are better positioned to achieve consistent time savings,** establish formal training and integrate AI across workflows. The relevant policy challenge is to enable SME experimentation without allowing resource constraints to turn organisational flexibility into insecure, fragmented or permanently low-maturity adoption.

These developments require a shift in the unit of analysis. Early discussions of workplace AI often focused on whether individual tasks could be automated. **Task-level analysis remains essential** because occupations consist of activities with different technical, social and legal characteristics. However**, the effects of changing one task cannot always be understood in isolation.** Tasks form part of workflows involving sequencing, handovers, communication, verification, exception handling and accountability. Automating the production of a first draft, for example, may **increase the importance of reviewing, validating and approving outputs.** Delegating part of a client-onboarding process to an agentic system may change who gathers information, who detects errors, when human intervention occurs and who remains answerable for the final decision.

The effects of AI must therefore be examined across several connected levels:

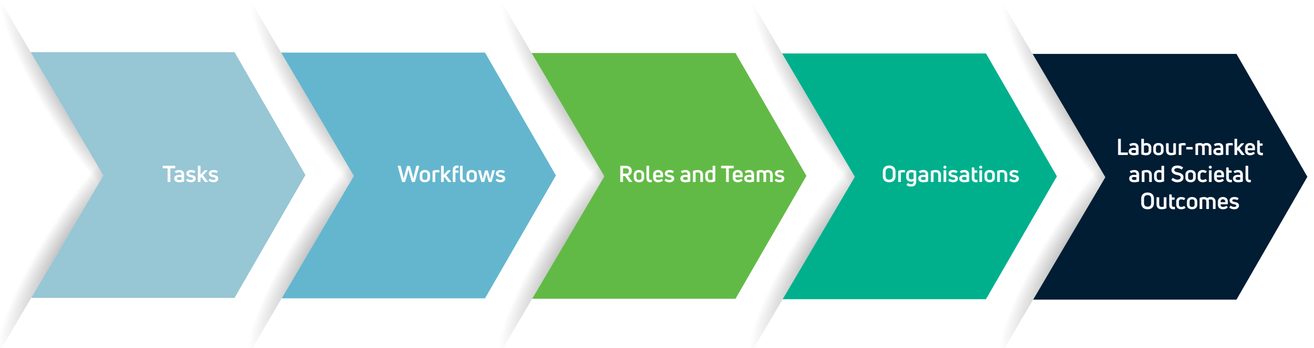

At the task level, AI may automate, augment or support a particular activity. At the workflow level, it may alter the order and allocation of activities and introduce new requirements for oversight. At the level of roles and teams, it may shift workers from direct production towards evaluation, coordination and exception handling. At the organisational level, outcomes depend on leadership, data quality, internal policy, training, technical integration and change management. At the wider societal level, these organisational choices influence employment, skill formation, productivity, inequality and the distribution of opportunity.

This report uses the term human–machine partnership to describe these deliberately structured arrangements between human workers and AI systems.

Partnership does not imply equality between people and machines, nor does it attribute human agency or moral responsibility to AI. It refers to the practical allocation of capabilities, authority and accountability within a combined work system. A productive partnership requires decisions about what AI should do, what humans should retain, how outputs should be verified, when systems should be overridden, and who is responsible for consequences.

This emphasis is important because immediate performance is not the only relevant outcome of AI adoption. Recent experimental evidence indicates that AI assistance can improve performance during supported work while, under some conditions, reducing subsequent persistence and unassisted performance. The negative effects were concentrated particularly among participants who used AI to obtain direct solutions rather than hints or clarification. A related theoretical model describes an "augmentation trap" in which organisations rationally pursue short-term productivity gains while gradually weakening the worker expertise on which longer-term performance depends. These studies are recent preprints and should not be treated as definitive evidence of general workplace de-skilling. They nevertheless demonstrate why workforce strategy must consider learning, retained capability and professional development alongside immediate efficiency (Liu et al., 2026; Caosun and Aral, 2026).

The central question is therefore not simply whether AI raises productivity or removes jobs. **It is whether organisations can design human–machine work systems that improve productivity and innovation while preserving judgement, accountability, learning, trust and job quality.** This requires a whole-system approach. AI tools introduced without clear workflows, secure data, appropriate training or organisational policy may produce isolated time savings while also generating rework, uncertainty, shadow AI and unmanaged risks. Conversely, governance should not function only as a restriction. Clear and proportionate policies can create approved pathways for experimentation by specifying which systems may be used, what data may be entered, where human oversight is required and who is accountable for final outputs.

Against this background, the present study examines how AI is affecting work and workforce adaptation in Ireland. It combines a structured literature review with a **cross-sector survey of 487 workers, behavioural vignette experiments involving 421 participants, 32 semi-structured interviews with senior organisational leaders, and computational analysis of 48 policy documents from Ireland, the European Union and the United Kingdom.** This mixed-methods design connects workers' reported experiences and behavioural responses with organisational perspectives and the wider institutional environment.

The study adopts a right-sizing perspective. Right-sizing does not mean slowing adoption or minimising the role of AI. It means matching the form and degree of AI involvement to the task, workflow, organisational capacity and level of risk. It also means evaluating success through more than adoption rates or time saved. Relevant outcomes include the quality of work, employee capability, trust, accountability, inclusion, innovation and organisational resilience.

The report addresses the following objectives:

Research Objectives

01	Identify and analyse the technical, cognitive, socio-emotional and organisational capabilities required in AI-influenced work environments.
02	Analyse how AI, machine learning, generative AI and agentic AI are affecting existing tasks, workflows and roles through automation, augmentation and delegation.
03	Identify new and hybrid responsibilities emerging from AI adoption, including oversight, orchestration, assurance and human–AI coordination.
04	Examine the behavioural and organisational conditions shaping human–AI collaboration, including trust, transparency, verification, delegation and accountability.
05	Assess the organisational capabilities associated with effective AI adoption, including leadership, culture, data governance, change maturity and workforce readiness.
06	Provide practical recommendations for employers, SMEs, education and training providers, and policymakers.
07	Develop a strategic framework for right-sizing effective and inclusive human–machine partnerships in Irish workplaces.

Research Design and Methodology

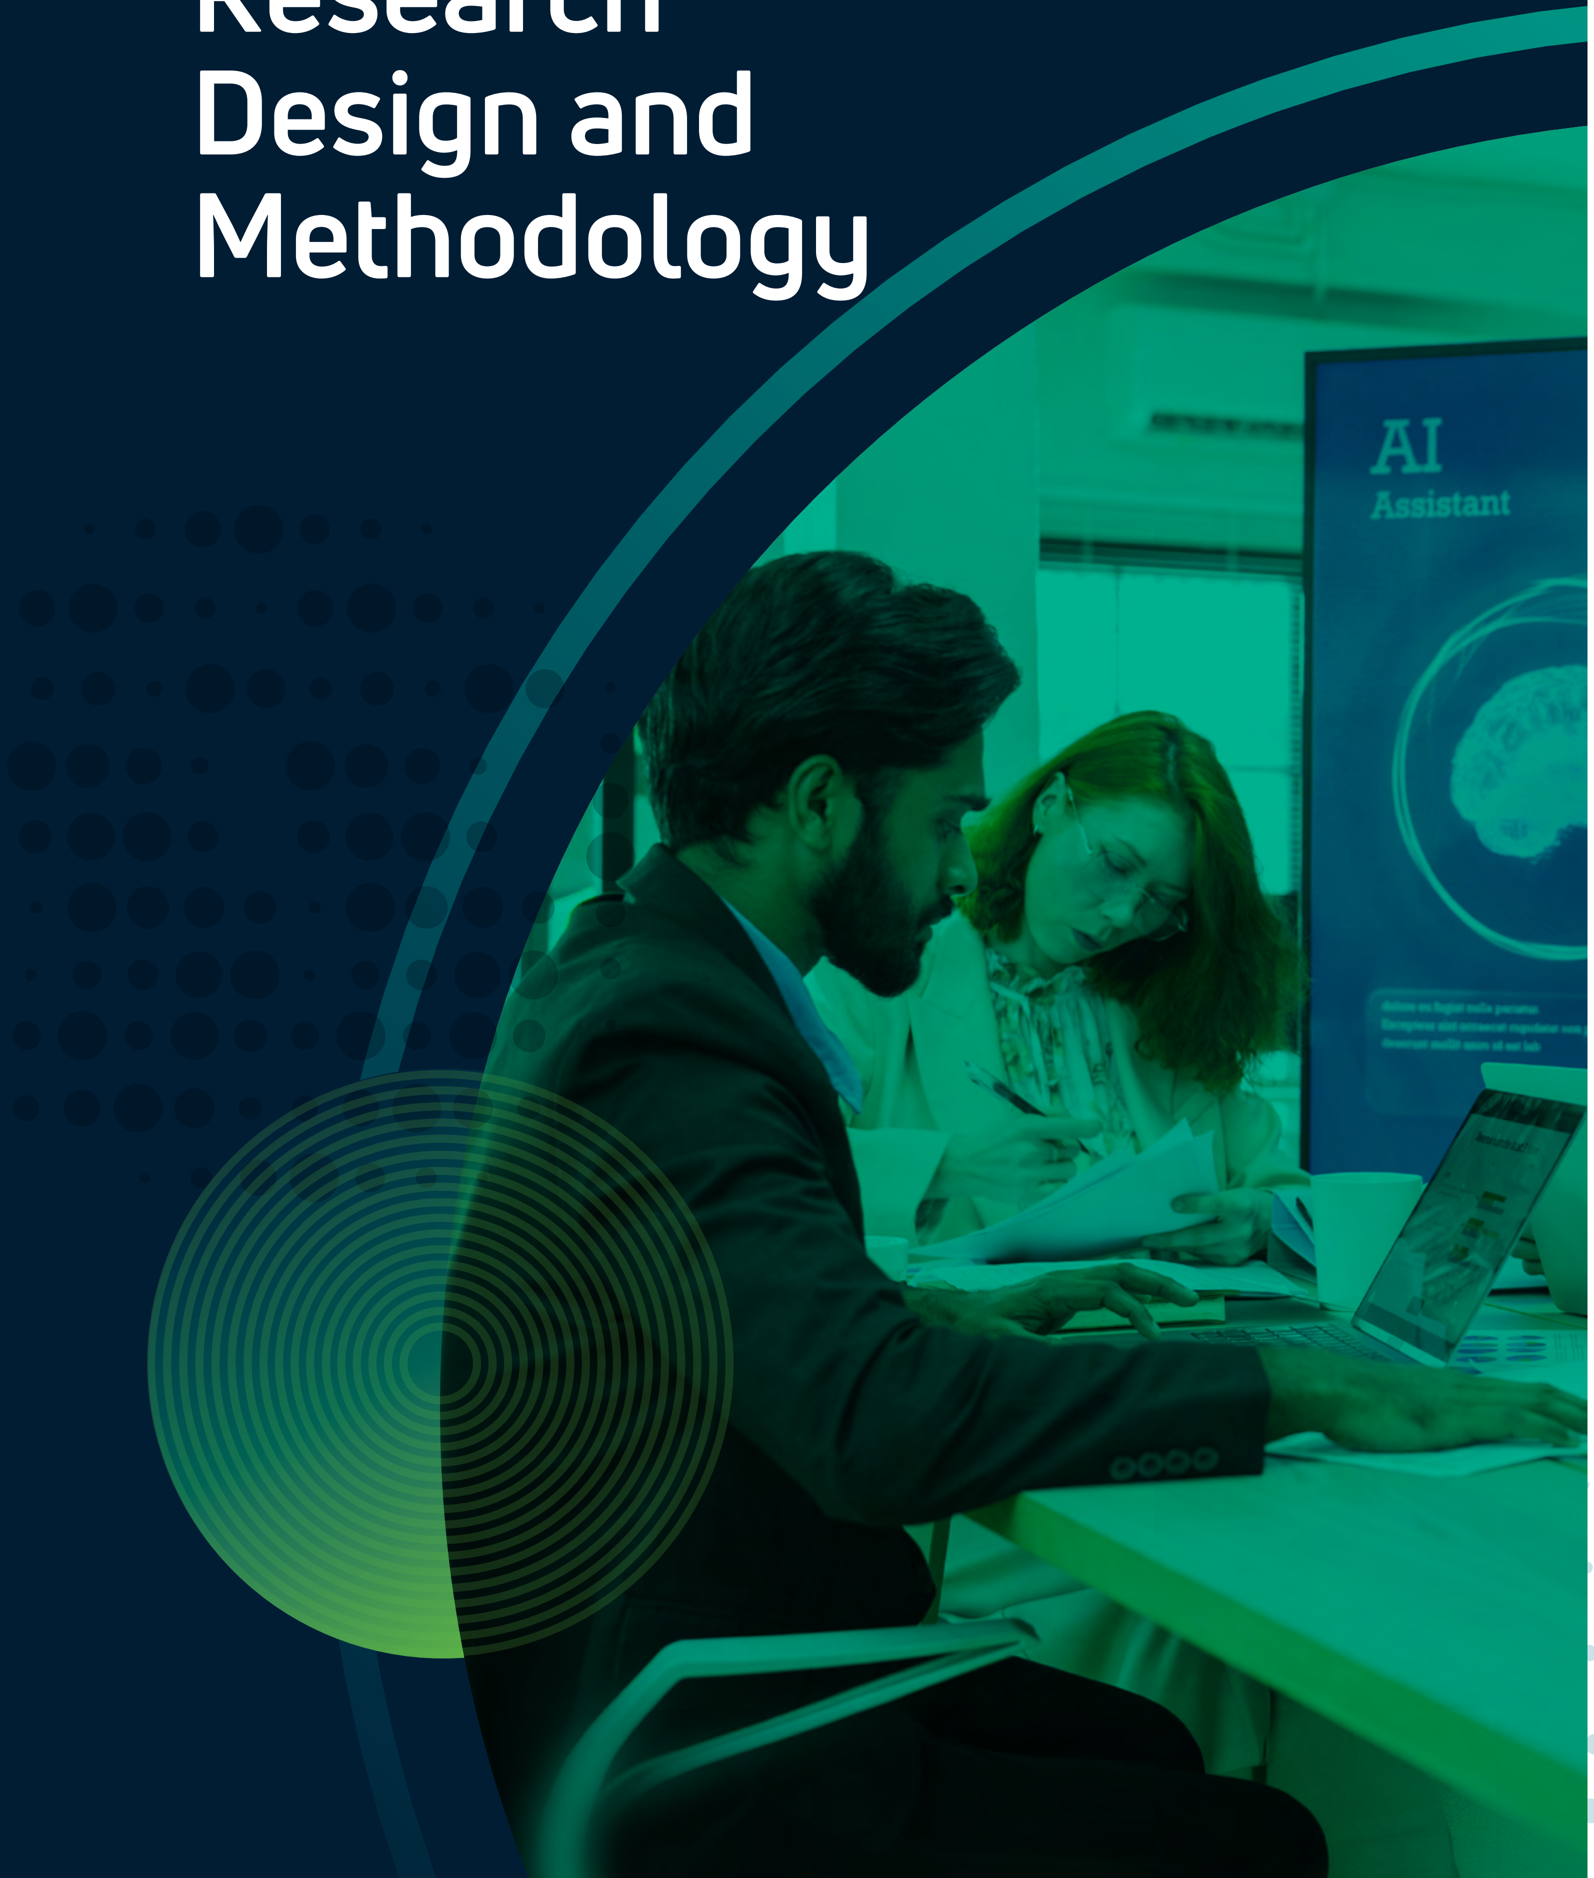

02 Research Design and Methodology

This study uses a mixed-methods research design combining five complementary forms of evidence: a structured literature review, a cross-sector workforce survey, two embedded randomised behavioural experiments, semi-structured interviews with senior organisational leaders, and comparative computational analysis of policy documents. Together, these components enable the report to **examine AI-related workforce transformation at multiple levels, from individual experiences and behavioural judgements to organisational practices and the wider policy environment.** The different components are used for triangulation rather than treated as directly interchangeable sources of evidence

2.1. Literature Review

A structured state-of-the-art literature review synthesised peer-reviewed research, policy documents, and labour-market evidence on how Artificial Intelligence (AI), Machine Learning, and emerging generative and agentic systems are reshaping skills, tasks, workflows, job design, and organisational structures. The review also examined human–machine partnership, trust and accountability, organisational maturity, SME adoption, job quality, and the distributional consequences of AI. Its findings establish the theoretical and empirical foundation for the subsequent components and informed the design of the survey instrument, experimental treatments, interview guide, and computational policy analysis.

2.2. Workforce Survey

The core quantitative component consists of a cross-sector **online survey of 487 workers in Ireland,** administered through Qualtrics and incorporating two embedded randomised behavioural experiments. The survey measures AI exposure and use, confidence in relevant skills and adaptability, perceptions of risk and opportunity, organisational reskilling supports, reported productivity effects, and perceived changes in tasks and job roles. The full questionnaire is provided in Appendix A. All participants provided informed consent, and ethical approval was granted by the Trinity College Dublin School of Social Sciences and Philosophy Research Ethics Committee (REAMs No. 5591).

The survey design targeted five economic sectors central to AI adoption: Industry (B-E), Professional, Scientific and Technical Activities (M), Information and Communication (J), Construction (F), and Financial, Insurance and Real Estate Activities (K and L), while also maintaining broader workforce representation. The achieved sample is weighted towards Information and Communication (30.4%) and Professional, Scientific and Technical Activities (16.8%), reflecting the concentration of AI-intensive activity within the sample. Respondents are drawn from organisations of varying sizes. **The largest share is employed in enterprise-scale firms (39.0%), followed by medium-sized (18.9%) and large organisations (16.4%), with additional representation from small (14.6%) and micro enterprises (10.9%).** The sample covers a wide range of organisational functions, **with strong representation in IT and engineering (18.9%) and operations (17.5%),** alongside finance, sales, human resources, services, and a dedicated cohort in data and AI roles (5.7%), as well as respondents working in education, healthcare, and executive positions. Geographically, respondents are **predominantly based in Dublin (57.3%),** although all regions of Ireland are represented. **The gender distribution is broadly balanced, with 53.2% female and 44.1% male respondents.** Further details on the sample characteristics are provided in Appendix C.

2.3. Vignette Experiments

Two randomised vignette experiments were embedded within the workforce survey and completed by 421 participants. **The experiments examine trust, reliance, accountability, and evaluations of human–AI collaboration under controlled conditions**. All participants provided informed consent, and ethical approval was granted by Trinity College Dublin (REAMs No. 5591). The study was preregistered before data collection on the Open Science Framework (https://osf.io/2xh9v). The scenarios used in the vignette experiments are described in Appendix B.

2.4. Senior Leadership Interviews

A total of **32 semi-structured interviews were conducted** with senior leaders across key industry sectors in Ireland. Using purposive sampling, this component captured a wide range of perspectives on AI adoption, workforce transition, governance challenges, and training infrastructure. **The sample focused on individuals in decision-making roles** with direct experience of implementing AI in business and enterprise contexts, and included participants from different sectors and organisational types.

Participants were identified through professional networks and partner organisations. All participants provided informed consent before taking part. Ethical approval was granted by the Trinity College Dublin School of Social Sciences and Philosophy Research Ethics Committee (REAMs No. 5623). The semi-structured interview guide is presented in Appendix D.

The final sample reflects a diverse cross-section of the Irish economy, spanning Information and Communication; Professional, Scientific and Technical Activities; Industry, including manufacturing and energy; Financial Services; and Construction. Although the sample is weighted towards Information and Communication firms, reflecting the concentration of AI development and deployment in that sector, it also includes organisations applying AI in more traditional industries such as manufacturing, food production, and energy optimisation.

In organisational terms, the sample includes large multinational corporations, Irish firms, Irish-founded scale-ups, and smaller enterprises. Large organisations with 250 or more employees are strongly represented, particularly among multinational technology and financial-services firms, while small and micro enterprises include emerging AI start-ups and specialist consultancies. Medium-sized firms are also represented, particularly in industry, energy, and enterprise software. Ownership structures vary across the sample and include foreign multinationals, primarily headquartered in the United States and the United Kingdom; Irish-owned private firms; Irish-founded multinational scale-ups; and a small number of public and non-profit organisations.

This composition supports comparative analysis of AI adoption by firm size, ownership type, and sector, while also capturing differences between organisations that develop AI technologies and those that adopt AI within operational and organisational settings.

2.5. Policy Document Analysis (EU, Ireland, and UK)

To examine the institutional arrangements that enable or constrain AI skills development, **the project applied computational text analysis to AI-related policy documents from the European Union, Ireland, and the United Kingdom.** The analysis systematically compared how regulations, strategies, and governance frameworks frame workforce transformation and skills development. **A structured corpus of 48 official documents was constructed and analysed to identify recurring topics, similarities, divergences, and key concepts across jurisdictions, over time, and by document type.**

The analysis combined word co-occurrence networks, based on adjacent terms, with topic modelling to identify dominant themes and patterns of policy discourse across Ireland, the European Union, and the United Kingdom. The complete list of documents included in the analysis is provided in Appendix E.

The experiments examine trust, reliance, accountability, and evaluations of human-AI collaboration under controlled conditions.

Literature Review Outcome

03 Literature Review Outcome

3.1. Global and national evidence on AI-driven workforce transformation

Artificial Intelligence (AI) is reshaping work, but the emerging evidence does not support a simple choice between mass technological unemployment and frictionless augmentation. A more defensible interpretation is that AI is changing the content, sequencing and organisation of work unevenly across occupations, firms and sectors. Its effects depend not only on the technical capabilities of particular systems, but also on the tasks to which they are applied, the workflows in which those tasks are embedded, and the organisational arrangements governing their use. AI-driven workforce transformation is **therefore better understood as a process of work-system reconfiguration than as a uniform technological shock.**

Concerns about technological unemployment have accompanied successive waves of mechanisation, computerisation and automation. Autor (2015) argues that technologies rarely remove occupations as indivisible units; instead, they substitute for some activities while complementing others and changing the relative value of different forms of human labour. Frey and Osborne's (2017) influential estimate that 47% of US employment was at high risk of computerisation drew attention to the scale of potential exposure, but later task-based approaches cautioned against treating technical feasibility as a forecast of realised job loss. Acemoglu and Restrepo (2019) distinguish the displacement effects of automation from the reinstatement effects created when technologies generate new tasks for labour.

This distinction remains central in the generative-AI era: exposure identifies where change may occur, but organisational and institutional choices shape whether the result is substitution, augmentation, task reallocation or the creation of new work.

Forward-looking employer evidence nevertheless indicates that **substantial restructuring is expected.** The World Economic Forum's Future of Jobs Report 2025, based on more than **1,000 employers across 55 economies,** estimates that labour-market transformation could affect **22% of current formal employment by 2030.** Employers anticipate 170 million roles being created and 92 million displaced, producing a net increase of approximately 78 million roles. They also expect 39% of workers' existing skill sets to be transformed or become outdated, while 63% identify skills gaps as the leading barrier to organisational transformation. These figures represent employer expectations rather than observed labour-market outcomes, but they show the scale of adaptation for which organisations are preparing (World Economic Forum, 2025).

Recent UK evidence provides an important caution against translating exposure directly into aggregate employment effects. The British Progress review of labour-market evidence finds no detectable overall employment effect to date in occupations most exposed to generative AI, while also documenting substantial variation among occupations with similar measured exposure. Actual use remains concentrated in a narrower set of tasks than headline exposure measures imply. This reinforces the need to distinguish what AI could technically perform from what organisations are currently willing and able to redesign, delegate and govern (Serôdio, 2026).

The relevant unit of analysis is therefore expanding. Occupations remain useful labour-market categories, and task analysis remains indispensable, but contemporary **AI systems increasingly affect sequences of interdependent activities rather than isolated tasks.** A workflow includes the ordering of work, handovers between people and systems, data inputs, communication, verification, exception handling and final accountability. **Automating the production of a draft, for example, may reduce initial production time while increasing the importance of checking sources**, correcting errors, reconciling outputs with organisational standards and obtaining approval. The net effect cannot be inferred from the automated task alone.

The development of agentic AI strengthens the case for workflow-level analysis. A recent study of 177,436 publicly available tools designed for AI agents distinguishes tools that perceive information, reason over it and act on external environments.

Software development currently dominates the ecosystem, but action-oriented tools rose sharply during the study period.

Such tools can modify files, send communications and initiate transactions rather than merely generate content. Although public tool repositories do not provide a representative measure of workplace deployment, they illustrate a movement from assistance within a task towards the delegation and coordination of multi-step digital activity. **This raises questions about authority, monitoring, escalation and responsibility that are fundamentally organisational** (Stein, 2026).

Exposure also differs from earlier automation waves. Large language models can affect cognitive and communicative activities performed by professional and highly educated workers, including drafting, summarisation, analysis, coding and information retrieval (Eloundou et al., 2024). The OECD's employer and worker surveys across seven countries, including Ireland, **found broadly positive perceptions of AI's effects on performance and working conditions among users,** alongside continuing concern about job loss and unequal access to training (Lane, Williams and Broecke, 2023). **The emerging pattern is not a simple shift from low-skill to high-skill work**. Instead, AI can alter activities across the occupational distribution, while the consequences depend on workers' expertise, the quality of implementation and the availability of complementary support.

Productivity research demonstrates both the potential and the limits of augmentation. Brynjolfsson, Li and Raymond (2025) found that generative AI increased productivity among customer-support agents, with the largest gains among less experienced workers. Other field and experimental studies have reported gains in writing, software development and professional knowledge work (Dell'Acqua et al., 2023; Peng et al., 2023). **These studies show that AI can codify practices, reduce search costs and provide scaffolding.** They do not, however, establish that equivalent gains will arise across all tasks, occupations or organisational settings. **Performance can decline when systems operate outside their effective capability boundaries,** when users cannot evaluate outputs, or when time saved in production is absorbed by verification and rework.

Irish evidence similarly points to widespread adoption combined with limited organisational maturity. The AI Economy in Ireland 2026 report, based on a survey of 250 senior leaders, concludes that day-to-day AI use is now common but that many organisations remain at an **early stage in converting individual productivity improvements into sustained operational and strategic value.** It identifies a gap between employees' adoption of AI and organisations' redesign of workflows, processes and decision structures. The report therefore shifts attention from whether organisations have adopted AI to whether they possess the skills, governance and operating models required to integrate it effectively (Trinity College Dublin, 2026).

This maturity gap is particularly important for small and medium-sized enterprises (SMEs) Contemporary Irish evidence suggests that SMEs which invest seriously in AI can report substantial gains, reflecting their capacity for rapid experimentation and shorter decision chains.

At the same time, **SMEs are more concentrated at earlier stages of adoption and face greater constraints** in specialist expertise, formal training, data readiness, governance and implementation finance. **Large organisations are more likely to achieve consistent time savings and scaled integration,** while many SMEs remain dependent on fragmented or individual-led use. The implication is not that SMEs are inherently less capable, but that **flexibility must be matched by accessible infrastructure,** practical support and proportionate governance (Trinity College Dublin and Microsoft Ireland, 2026).

The distributional effects extend beyond differences between firms. ESRI modelling of short- to medium-term scenarios for Ireland shows that productivity growth need not translate automatically into broadly shared household gains. In the central scenario, around 7% of current jobs are displaced, retained workers receive modest productivity-related wage gains, and average disposable household income falls moderately because employment losses are not fully offset by wage and capital-income gains. Inequality also rises. These are simulated scenarios rather than forecasts, and their results depend on modelling assumptions. Their importance lies in demonstrating that the distribution of ownership, employment risk and productivity gains matters alongside aggregate economic performance (Doorley et al., 2026).

Overall, the evidence supports a shift from a binary **debate about jobs towards a multi-level account of transformation:** AI changes particular tasks; these changes alter workflows; workflow redesign reshapes roles and teams; organisational capabilities condition whether the resulting system performs effectively; and the cumulative choices of organisations affect employment, job quality, skill formation and inequality. The dominant near-term pattern is therefore better described as uneven task and workflow restructuring than as immediate, economy-wide replacement

> **It is not that SMEs are inherently less capable, but that flexibility must be matched by:**
>
> - **Accessible Infrastructure**
> - **Practical Support**
> - **Proportionate Governance**

(Trinity College Dublin and Microsoft Ireland, 2026)

3.2. Skills evolution, capability formation and human–AI partnership

AI and automation alter labour demand by changing both the tasks performed within occupations and the capabilities required to perform them. Skill-biased technological change provides an important account of how technological innovation can increase demand for highly educated and cognitively skilled workers while contributing to wage inequality (Autor, Katz and Krueger, 1998; Violante, 2018). However, it treats skill levels relatively broadly. Task-based models provide a more precise account by conceptualising occupations as bundles of activities whose allocation between workers and technologies can change over time (Acemoglu and Autor, 2011).

For AI-influenced work, capability requirements are multidimensional. Technical capabilities include basic AI literacy, data literacy, prompt and tool use, security awareness and, for specialist roles, model development, integration and assurance. **Cognitive capabilities include problem formulation,** critical evaluation, domain judgement and the ability to detect weak or inappropriate outputs. Socio-emotional and relational capabilities include communication, collaboration, negotiation and the management of client or colleague expectations. **Organisational capabilities include** understanding approved systems, data-handling rules, escalation procedures, accountability and how AI use fits within a broader workflow. The growing importance of hybrid profiles means that **workforce development cannot be restricted to recruiting technical specialists.**

The language of AI literacy is useful but can be too generic. Knowing broadly what AI is does not ensure competent use within a particular occupation. Effective capability is role-specific and situated: a finance professional, engineer, human-resources manager and public-service administrator require different knowledge about appropriate delegation, relevant evidence, risk, verification and accountability. **Practical training should therefore connect general concepts to authentic tasks and workflows, rather than treating prompting or tool familiarity as sufficient evidence of readiness.**

Recent research complicates the assumption that access to AI necessarily builds capability. Experiments on AI assistance and persistence suggest that brief support can improve immediate performance while reducing subsequent independent performance and persistence, particularly when users obtain direct solutions rather than hints or clarification. Related work on an '**augmentation trap'** develops the longer-term organisational implication: managers may rationally prioritise short-term gains even where **repeated cognitive offloading weakens the worker expertise on which future productivity depends.** These studies are recent and should not be treated as definitive evidence of general workplace deskilling, but they show why immediate output is an incomplete measure of successful adoption.

The timing and form of assistance also matter. Recent human-computer interaction research indicates that the effect of AI support on critical thinking depends on when assistance is introduced and how much time users have available. **AI can support reasoning when it is used to provide scaffolding, alternative perspectives or feedback, but it can displace effort when it supplies ready-made answers before users have engaged with the problem.** Training design should therefore **teach workers not only how to obtain outputs, but how to preserve active engagement,** verify claims and decide when unaided reasoning is important (Shen and Tamkin, 2026).

Existing expertise remains valuable in AI-augmented work. Studies comparing experts and non-experts indicate that domain knowledge transfers into AI-supported settings and improves users' ability to formulate tasks, recognise errors and select among outputs. **AI may narrow some performance gaps by diffusing codified expertise, but it does not remove the need for judgement.** This creates a central tension for workforce development: organisations need to broaden access to AI while also maintaining the expertise required to supervise it (Eisenmann et al., 2025).

> **Organisations need to broaden access to AI while also maintaining the expertise required to survive it**
>
> **(Eisemann et al., 2025)**

Human–AI partnership is therefore best understood as a designed allocation of capabilities, authority and accountability rather than as a simple interaction between a user and a tool. Productive arrangements require decisions about what AI should do, what humans should retain, how outputs should be verified, when systems should be overridden, and who remains responsible for consequences. Partnership does not imply equality between people and machines, nor does it attribute moral agency to AI. It describes the practical organisation of combined human and technical work.

Social perception can also shape whether AI-supported work is accepted. Experimental research on disclosure identifies an 'AI-use penalty': workers who disclose using AI can be evaluated less favourably, including through lower trust, legitimacy and, in some studies, compensation-related judgements. These effects are not uniform; they vary with attitudes towards technology, expectations about task appropriateness and beliefs about AI accuracy. The implication is that transparency should be designed around accountability and material relevance, rather than assuming that indiscriminate disclosure of every low-level use will necessarily increase trust.

Anthropomorphic design introduces a further layer of complexity. Cross-national evidence suggests that humanlike design reliably increases perceptions of humanness, but does not produce uniform increases in trust or engagement. **Effects vary across cultural contexts and according to the social characteristics attributed to the system.** Research on gendered AI similarly shows that assigned or perceived gender can shape cooperation and evaluation, while visual anthropomorphism can reproduce biases already present in judgements of human managers (Bazazi, Karpus and Yasseri, 2025; Cui and Yasseri, 2026; Han, Cui and Yasseri, 2026). Human–machine partnerships are therefore not socially neutral: they can inherit and amplify existing categories, expectations and inequalities.
Overall, the skills challenge is not simply a shortage of technical expertise. It is a capability-design problem involving role-specific knowledge, critical evaluation, learning, collaboration, governance and retained human judgement. Effective upskilling must enable workers to use AI productively without becoming unable to perform, assess or take responsibility for the work that AI supports.

> **AI may narrow some performance gaps by diffusing codified expertise, but it does not remove the need for judgement.**
>
> **This creates a central tension for workforce development:**

3.3. Organisational adaptation, workflow redesign and productivity

Research in economics and management consistently characterises AI as a general-purpose technology whose value depends on complementary organisational investments. Brynjolfsson, Rock and Syverson (2017) attribute **the modern productivity paradox surrounding digital technologies partly to implementation lags: firms must redesign workflows, develop skills and reorganise production processes before gains appear in measured output.** These adjustments require intangible assets such as data infrastructure, organisational routines, managerial capability and workforce knowledge.

Policy research similarly frames AI as part of a production system combining computing, software, data and human expertise to produce prediction, optimisation and content generation (Filippucci et al., 2024). **Productivity effects therefore depend on how effectively organisations integrate AI into operational workflows and decision-making** rather than on tool availability alone. Firm-level studies generally find positive but heterogeneous associations between AI adoption and productivity, with stronger effects where AI is used across multiple functions and linked to complementary organisational change (Czarnitzki et al., 2023).

Evidence from Ireland highlights the importance of timing and application. Analysing firm data from 2013 to 2024, Siedschlag and Duran (2026) find that early productivity gains associated with AI adoption are generally modest and delayed, with stronger effects in particular applications such as marketing, administration and ICT security. Applications supporting workflow automation and decision-making appear especially associated with productivity improvements. This is consistent with the broader view that value emerges when AI is incorporated into a coherent operating model rather than layered onto existing work without redesign.

Cross-country executive evidence further illustrates the gap between adoption and realised productivity. Yotzov et al. (2026) report that while a large majority of surveyed firms use AI in some form, most report no measurable effect on labour productivity over the previous three years. Such findings should not be interpreted as proof that AI lacks value; rather, they suggest that adoption is often shallow, experimental or disconnected from core workflows. The relevant distinction is between access, use, integration and organisational transformation.

Workday's Beyond Productivity report adds an important operational qualification. Employees may save time during initial production but lose part of that saving to correcting errors, verifying outputs and rewriting generic content. The report describes this as **a hidden 'AI tax'** and argues that organisations achieve better returns when they **reinvest saved time in capability development** and role redesign rather than simply increasing workloads. Although these estimates are based on a corporate survey rather than audited productivity records, they usefully draw attention to rework and verification costs that are often omitted from headline productivity claims (Workday, 2026).

Organisational learning and agility provide a mechanism through which AI adoption can translate into effectiveness. Recent research on **generative-AI infusion suggests that tool adoption improves organisational outcomes when it strengthens learning,** knowledge exchange and the capacity to reconfigure processes. This reinforces the steering principle that AI adoption should be treated as a **whole-system change** involving technology, people, process and governance rather than as a collection of isolated use cases (Koh and Yuen, 2026).

Human capital investment is a central adaptation channel. Muehlemann (2025) finds that sustained AI adoption is associated with increased apprenticeship provision and modest increases in training directed towards higher-skilled workers. Earlier-stage adjustment may be more disruptive, including the reallocation of training resources or the replacement of some skill profiles through hiring. These patterns suggest that organisational adaptation **can involve a period of uncertainty before more stable capability-building arrangements emerge.**

AI tends to amplify existing organisational capabilities: strong organisations may use it to deepen advantage, while weak foundations can convert apparent efficiency into rework, risk and consistency

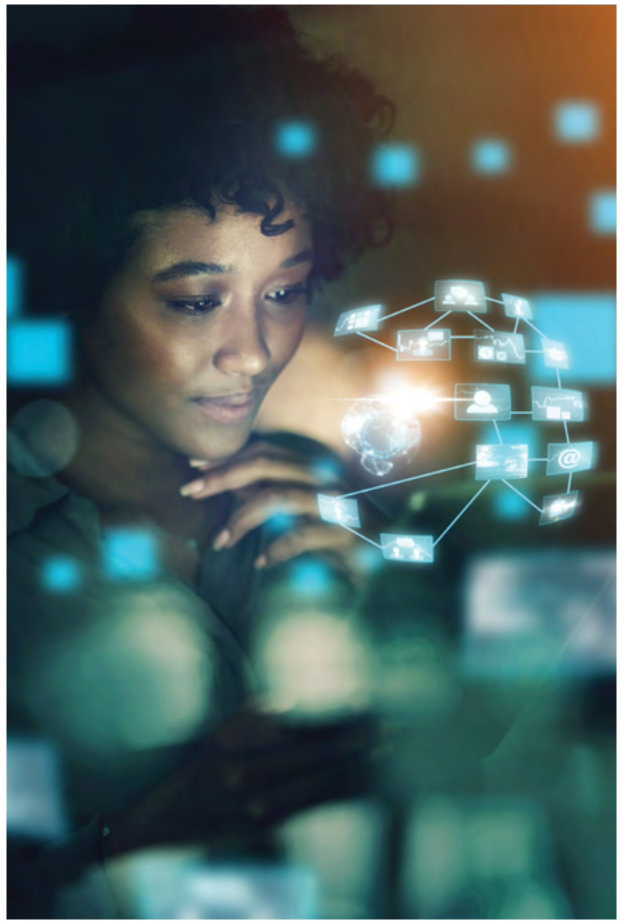

Trust and governance are also productive assets. Workday's Closing the AI Trust Gap report (2024a) documents a **difference between leadership enthusiasm and employee acceptance,** underlining the importance of transparency, consultation and credible safeguards. Clear organisational policy can **reduce shadow AI by defining** approved tools, permissible data use, verification requirements and accountability. Governance should therefore not be understood only as restriction. When it is proportionate and operationally clear, it can **create safe pathways for experimentation and more consistent adoption.**

Firm size and maturity strongly condition these processes. Larger and more productive firms are generally more likely to adopt AI because they possess stronger data infrastructure, specialist expertise and change-management capacity. SMEs can benefit from shorter decision chains and closer integration between leadership and operational work, but often face resource constraints that limit formal training, governance and scaling. **The policy challenge is to support experimentation without allowing organisational flexibility to become fragmented, insecure or permanently low-maturity adoption.**

Taken together, the literature indicates that AI is best understood as an **organisational transformation technology** rather than a standalone productivity tool. Gains are delayed, uneven and conditional on workflow redesign, skills investment, data readiness, leadership, learning and governance. AI tends to amplify existing organisational capabilities: strong organisations may use it to deepen advantage, while weak foundations can convert apparent efficiency into rework, risk and inconsistency.

3.4. Job quality, meaningful work and responsible adoption

Productivity is not the only relevant outcome of workplace AI. **Adoption also affects autonomy, learning, workload, professional identity, trust and the experience of meaningful work.** A system that produces more output while degrading judgement, increasing surveillance or transferring risk to workers cannot be evaluated adequately through time savings alone. **Responsible adoption therefore requires attention to both the quantity and quality of work.**

The Workday Global Workforce Report is useful in this respect because it **links meaningful work strongly to accomplishment** and recommends combining the offloading of repetitive tasks with collaborative job audits. This approach suggests that organisations should examine which activities workers find **administratively burdensome, which provide learning and professional meaning,** and which should remain under human control. Automation can improve job quality when it **removes low-value friction,** but it may reduce it when it removes opportunities for mastery, social contribution or visible accomplishment (Workday, 2024b).

The cognitive-offloading literature reinforces this point. When AI performs too much of the reasoning process, workers may complete tasks more quickly but **have fewer opportunities to practise judgement and build expertise**. Conversely, **well-designed assistance can create room for reflection, feedback and more complex work.** The appropriate objective is therefore not maximum automation, but a **right-sized allocation** of work that considers immediate performance, retained capability and longer-term development.

Responsible adoption also requires recognising that workers are evaluated socially for their use of AI. AI-use penalties may discourage transparent use, while vague organisational rules can push experimentation into informal channels. Organisations should distinguish between uses that materially affect decisions, safety, authorship or accountability and low-risk assistance that does not warrant burdensome disclosure. This allows transparency to serve a clear institutional purpose rather than becoming a symbolic requirement that workers experience as punitive.

Questions of anthropomorphism and perceived consciousness introduce additional risks. Conceptual work on seemingly conscious AI highlights possible emotional dependence, erosion of autonomy and over-attribution of understanding or intent. These concerns are not evidence that current systems are conscious; they concern how users interpret humanlike behaviour. In workplace design, interfaces should communicate capabilities and limitations clearly and avoid encouraging misplaced reliance (Bariach et al., 2026).

A responsible human–machine partnership therefore **combines productivity with safeguards for learning, autonomy, fairness and meaningful work**. It requires ongoing evaluation of who benefits, who bears additional verification or monitoring work, and whether AI-generated efficiencies are used to improve jobs or merely intensify them.

Organisations should distinguish between uses that materiality affect decisions, safety, authorship or accountability and low-risk assistance that does not warrant burdensome disclosure.

3.5. Comparative AI strategies and labour-market policies

Governments are adopting different strategies to address AI-related workforce change, particularly through education, training, diffusion support and regulation. OECD evidence indicates that **skills shortages, adoption costs and organisational uncertainty remain major barriers across advanced economies.** Internal barriers include difficulty identifying relevant skills, recruiting suitable staff and adapting operational processes. **External barriers include uncertainty about liability, data standards and regulation** (European Commission, 2020; Lane et al., 2023).

Most policy frameworks now recognise the importance of AI literacy: the ability to understand, use, monitor and critically evaluate AI without necessarily developing AI systems. However, policy increasingly needs to **move beyond generic literacy** towards differentiated capabilities for particular occupations, sectors and risk contexts. Adult education, vocational provision, workplace learning and flexible short courses are especially important because much of the existing workforce will encounter AI through changing jobs rather than through entry into new specialist occupations.

International practice illustrates several approaches. Finland's Elements of AI initiative focuses on broad public understanding. Norway supports enterprise-led continuing education in AI and data analytics in partnership with employers and employee organisations. Italy's Training 4.0 tax credit incentivises firm-level investment in advanced digital capabilities. Industry-linked institutions such as Canada's Vector Institute and the United States' MxD support diffusion through practice-oriented training and technical assistance. **These examples suggest that effective provision combines public coordination with employer participation and sector-specific application.**

Equity requires particular attention. Workers with lower formal qualifications, those in smaller firms and those whose employers lack training capacity are often less likely to participate in continuing education. Public intervention should therefore address both the supply of training and informational barriers that prevent workers and SMEs from identifying relevant opportunities. **Reskilling policy should also distinguish workers whose tasks are changing within existing roles** from those who face genuine displacement and may require more substantial occupational transition.

National AI strategies vary in the extent to which they integrate technical, transversal and complementary skills. **Ireland has been identified in comparative policy work as adopting a relatively whole-system perspective,** connecting digital and technical capability with communication, resilience, creativity, problem-solving and collaboration across education and workforce provision (Lassebie, 2023; Shi, 2025). **Ireland is also noted for industry-specific in-service training** delivered through mechanisms including SOLAS programmes and apprenticeships. The challenge is less the absence of strategic recognition than consistent implementation across sectors, regions and firm sizes.

Within this landscape, **Skillnet Ireland provides an important enterprise-led mechanism for workplace upskilling and reskilling.** Through its Skillnet Business Networks, training provision is developed in collaboration with employers and tailored to sectoral and organisational needs. Its Ireland's Talent Landscape 2025 report includes a dedicated examination of AI-related workforce and skills challenges, further illustrating the role of enterprise-led skills intelligence in shaping workforce-development provision. **Skillnet Ireland's model was also ranked first among 47 upskilling and reskilling initiatives from 36 countries assessed** in a European Commission-commissioned international study, including for its impact on businesses and its wider contribution to the economy (Skillnet Ireland, 2025; European Innovation Council and SMEs Executive Agency, 2024).

Technology Ireland Digital Skillnet provides a more focused example within the technology sector. Its AI-related initiatives combine technical capability with strategic, leadership and transversal skills. These have included a Generative AI masterclass delivered with London Business School in 2024 and the AI Capability Accelerator for senior leaders, commencing in March 2026. This approach is consistent with the evidence reviewed in this report: **effective AI readiness depends not only on familiarity with individual tools, but also on judgement, leadership and the organisational capacity to redesign work** (Technology Ireland Digital Skillnet, 2024, 2026).

The Irish policy environment must also respond to uneven organisational maturity. Large firms often possess the internal capability to interpret regulation, procure systems and develop training, while SMEs need more accessible guidance, shared infrastructure and trusted intermediaries. Ibec's AI Hub and related employer initiatives illustrate growing demand for practical implementation support (Ibec, 2025).

> Comparative analysis of national strategies suggests three recurring priorities: capturing economic and technological benefits, strengthening national capacity to develop and deploy AI, and establishing governance and oversight

(Guenduez et al., 2025).

Public policy can complement such initiatives through sectoral demonstrators, management training, diagnostic tools, data-governance support and co-funded workforce development.

The strongest strategies treat these priorities as interdependent. Skills without organisational adoption produce limited economic value; adoption without governance creates risk; and regulation without implementation support may entrench advantages among already capable firms.

The policy implication is that AI workforce strategy should not be reduced to producing more technical specialists. It must build broad and role-specific capability, support workflow and organisational redesign, protect workers facing displacement, and ensure that SMEs can adopt safely and productively. **It should also evaluate outcomes beyond adoption rates, including learning, job quality, inclusion, productivity and the distribution of gains.**

Workforce Survey Findings

04 Workforce Survey Findings

This section presents findings from a cross-sector survey of 487 workers in Ireland, examining patterns of AI use, perceived impacts on work, skills confidence, organisational readiness, trust, and training preferences. The survey provides an employee-level perspective on how AI is being experienced among respondents, complementing the leadership interviews, experiments, and policy analysis presented elsewhere in the report. Overall, the findings indicate that AI use is already common within this sample, while organisational governance and formal workforce-development provision have not always developed at the same pace.

4.1. AI Use is Widespread, but Reported Dependency Remains Limited

Survey responses indicate that AI tools are already embedded in everyday work practices for a substantial share of workers. Regular use is common, with many respondents reporting daily or frequent engagement with AI-based tools. Respondents also report that AI is increasingly visible within their occupational areas, suggesting that AI use is becoming normalised rather than confined to specialist technical functions. Generative AI tools appear to be the dominant entry point for most users, reflecting the accessibility and versatility of recent consumer-facing systems.

Figure 1. How often do you use AI-based tools in your work activities?

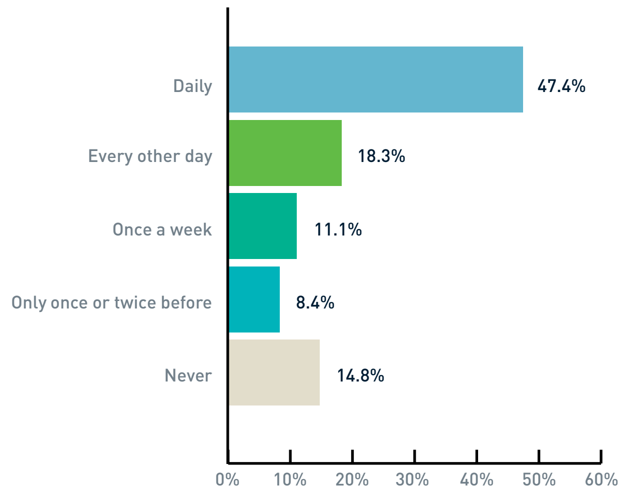

Figure 2. To the best of your knowledge, are AI tools being used by professionals in your area of work?

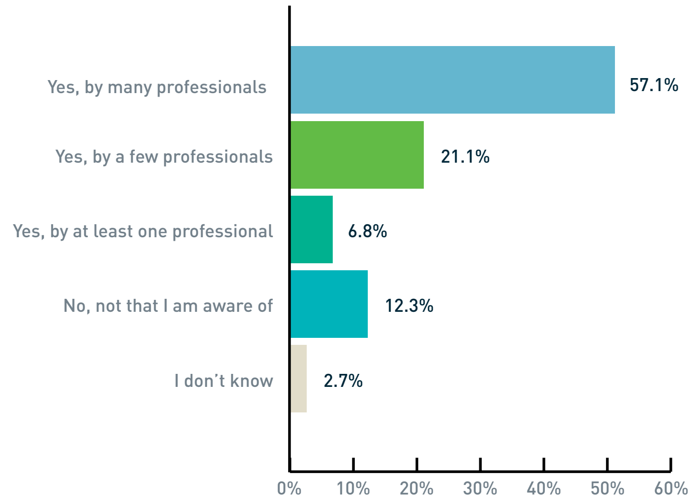

Figure 3. Which of the following AI-based tools do you currently use as part of your work? Select all that apply. (multiple responses)

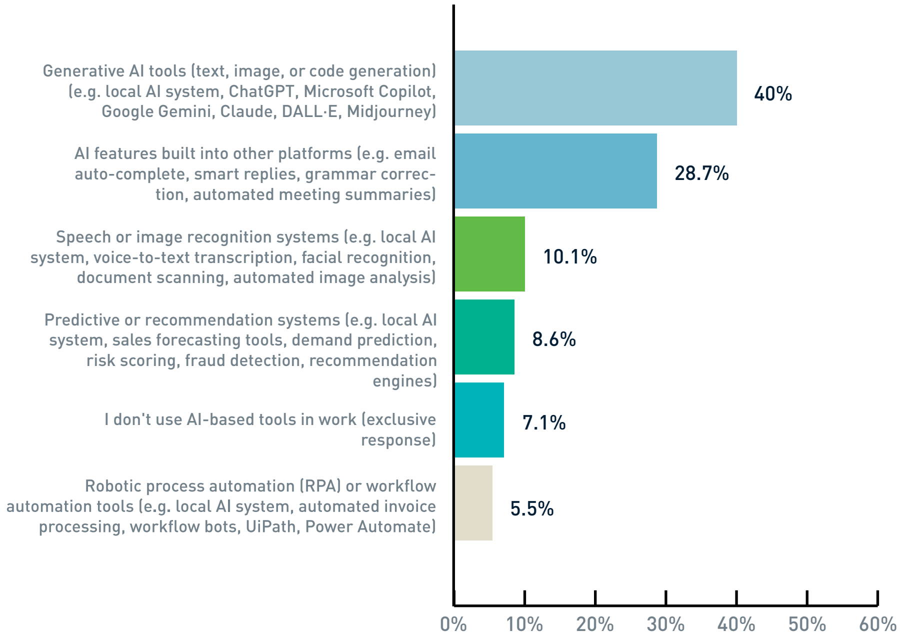

However, this widespread use is not matched by a strong sense of dependency. Many respondents do not perceive themselves as reliant on AI tools, and a substantial share report that their ability to complete tasks without AI has not significantly changed. The combination of frequent use and relatively limited self-reported dependency suggests that many respondents view AI as a useful support rather than as indispensable to completing their work. It may also reflect the incremental integration of AI into existing practices, although the survey cannot determine whether respondents underestimate their reliance on these tools.

Figure 4. To what extent do you feel dependent on AI tools to perform certain tasks in your role?

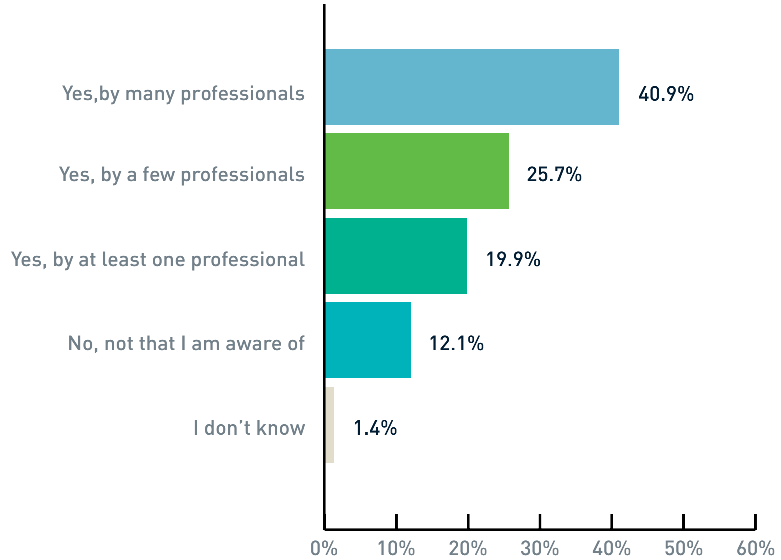

4.2. Confidence is High, but Formal Training is Limited

Respondents report relatively high levels of confidence in their ability to use AI tools effectively in their work. Many also indicate that they feel prepared to adapt their skills to work alongside AI in the future. This suggests a broadly positive orientation towards technological change among respondents, rather than widespread resistance to AI adoption.

Figure 5. Overall, how much has the use of AI affected your ability to complete work-related tasks on your own (without AI)?

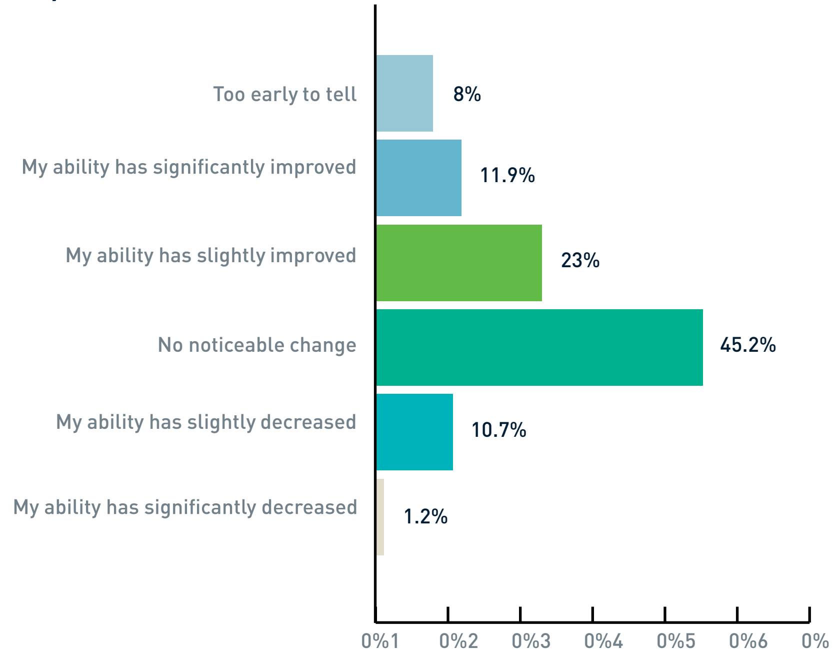

Figure 6. How confident are you in your ability to use AI tools effectively in your working tasks?

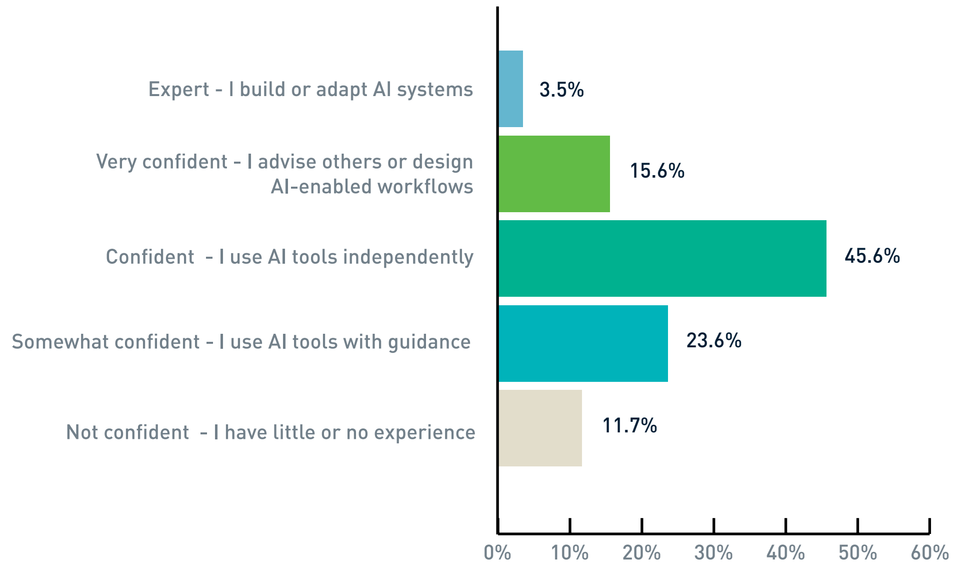

Figure 7. How much do you agree with the following statements: - I feel prepared to adapt my skills to work alongside AI.

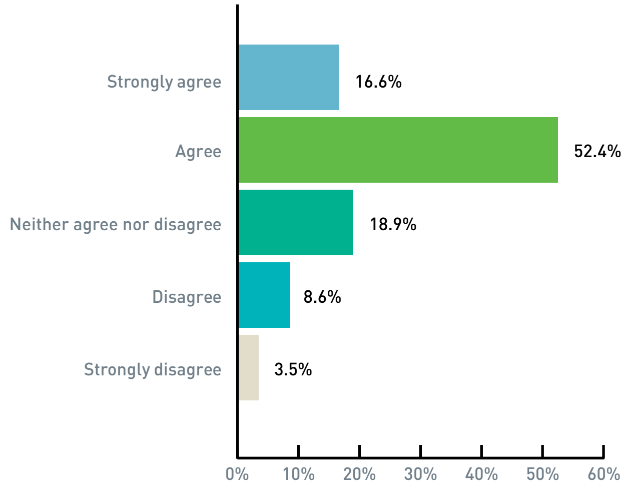

At the same time, a sizeable proportion of respondents report having taken no formal steps to develop AI-related skills. This reveals a gap between self-reported confidence and participation in structured learning. It does not necessarily indicate that respondents lack practical capability, as some may have developed skills informally, but it suggests that confidence is not consistently supported by formal training or recognised development pathways. The challenge for employers is therefore to convert informal experimentation into sustained, role-specific capability building.

Figure 8. Have you taken any steps to adapt your skills in relation to use of AI in your work?

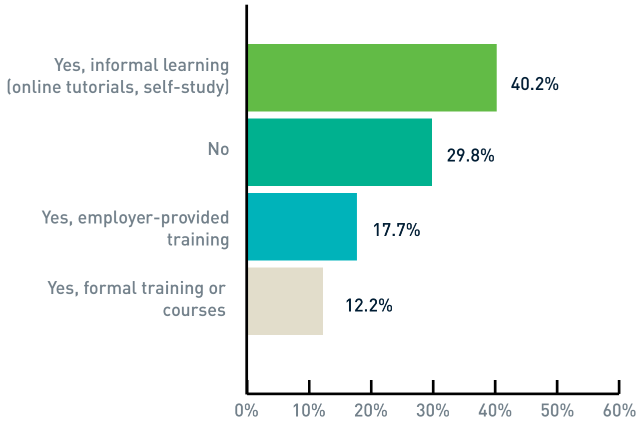

Training preferences provide further insight. Respondents show stronger interest in practical, self-paced, and hands-on formats than in theoretical or generic instruction. This suggests that effective workforce-development strategies should be aligned closely with real job tasks, occupational contexts, and immediate workplace applications.

Figure 9. What type of AI-related training would you prefer? Select all that apply.

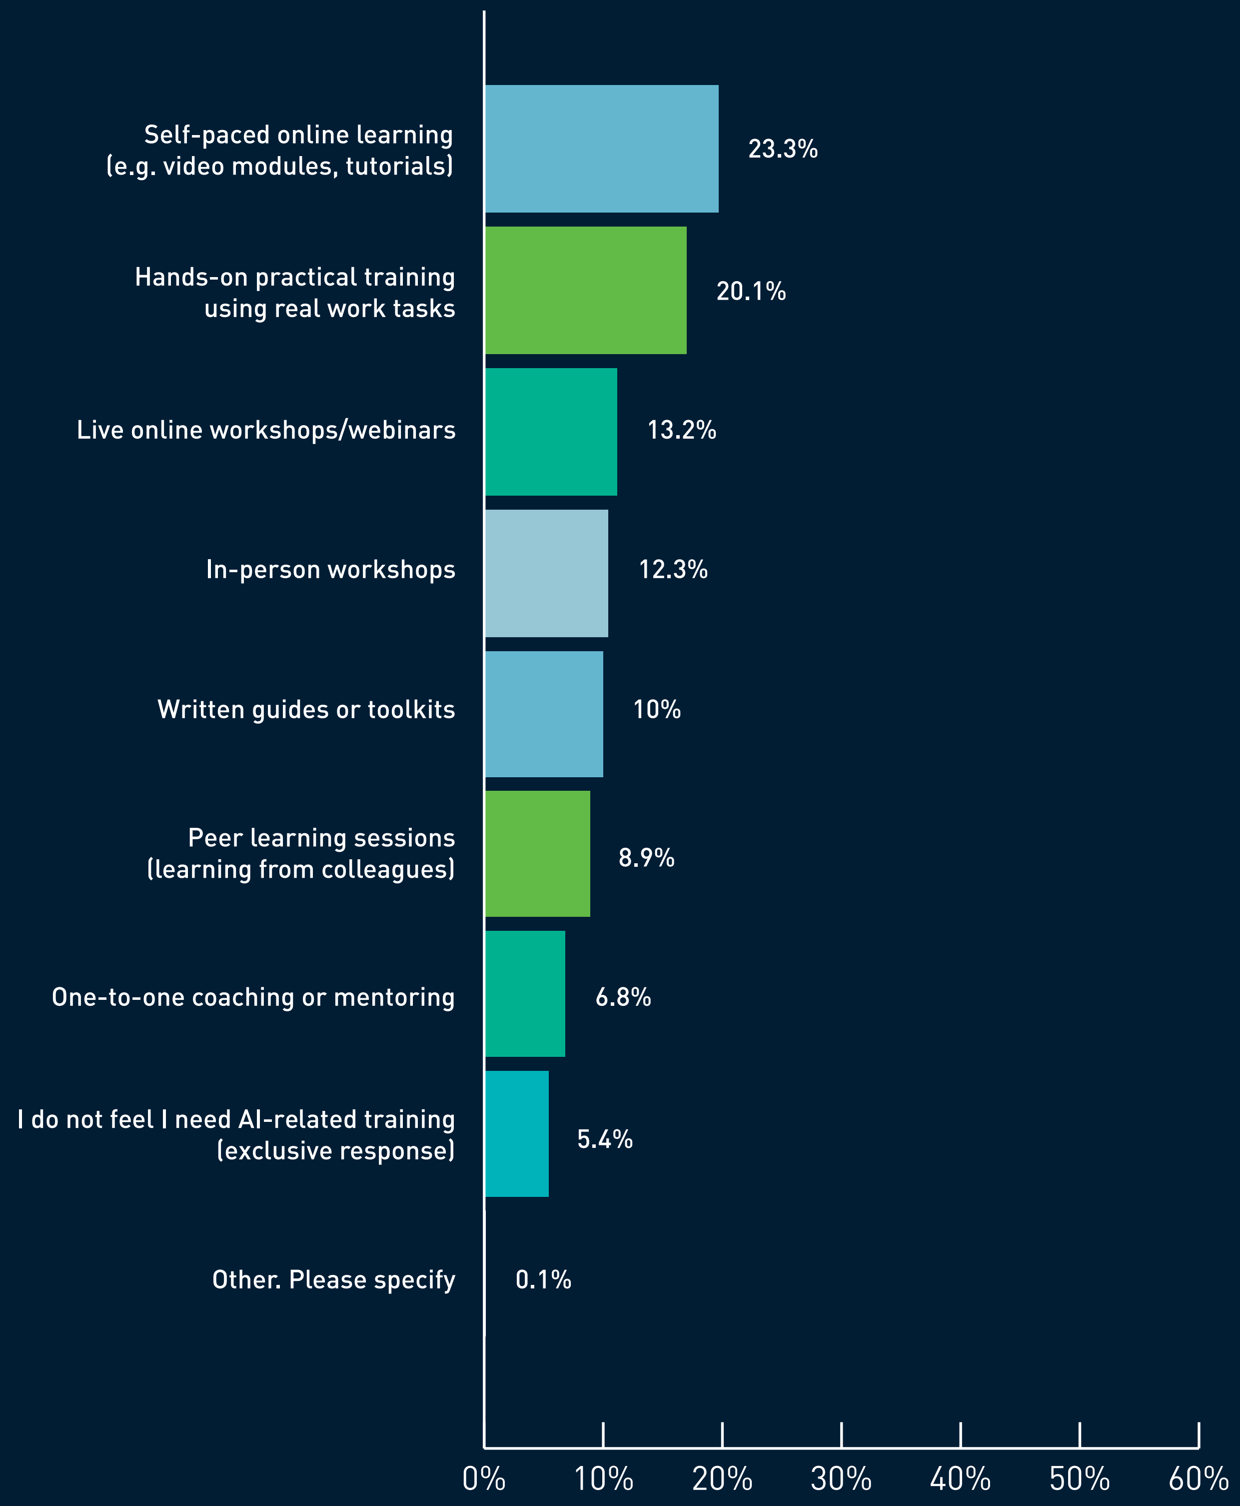

4.3. Workers See Opportunity More Than Threat

Despite prominent public narratives about job displacement, respondents do not appear uniformly fearful that AI will replace significant parts of their roles in the near future. Concerns are present, but the overall pattern is mixed rather than uniformly alarmist.

Figure 10. How much do you agree with the following statement: I am concerned that AI could replace significant parts of my job in the future.

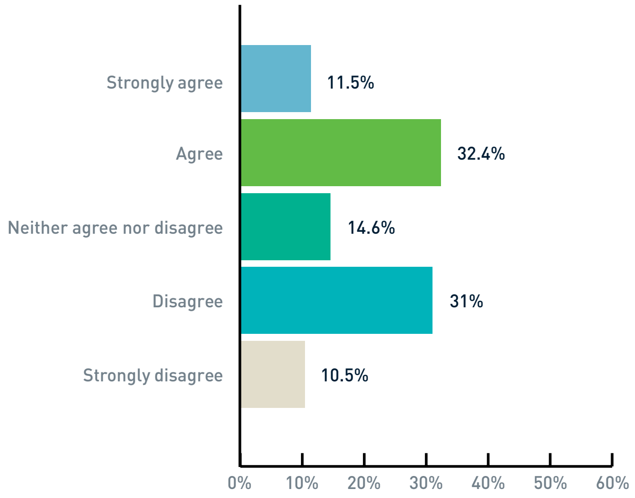

Alongside relatively measured concern about displacement, many respondents also believe that their workplaces may be missing opportunities to use AI more effectively. Taken together, these responses suggest that employees perceive both risks and unrealised organisational opportunities, rather than viewing AI exclusively as a threat.

The responses are consistent with a generally pragmatic orientation towards adoption, particularly where workers perceive potential improvements in efficiency, work quality, or the reduction of repetitive tasks.

Figure 11. How much do you agree with the following statement: I think my workplace is missing opportunities to use AI to improve how it operates.

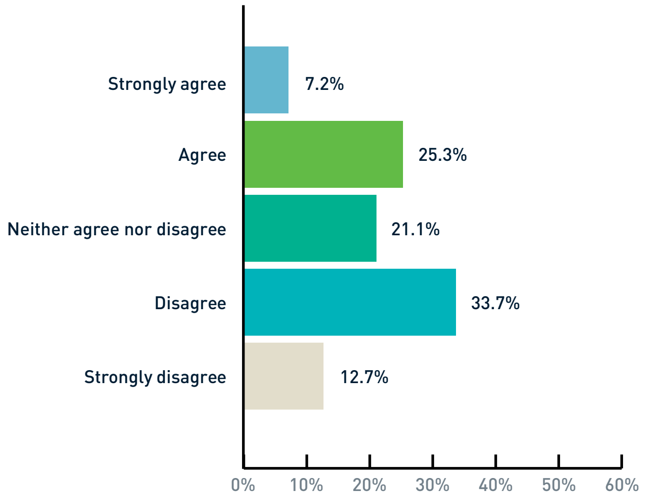

4.4. Trust, Risk and Governance Shape Adoption

Trust in AI-generated content is conditional rather than absent. Respondents do not appear uniformly willing to rely on AI outputs without review, highlighting the continuing importance of human verification. The survey does not establish whether this caution is excessive or well calibrated, but it indicates that perceived reliability, accountability, and risk remain relevant to workplace use.

Figure 12. How much do you trust content that is generated by AI for work tasks?

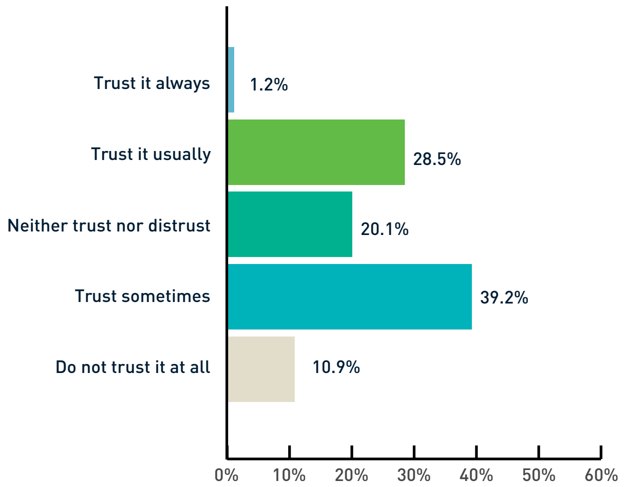

When asked about barriers to greater use of AI in daily work, respondents identify concerns including accuracy, privacy, misuse, and uncertainty about appropriate use. These findings suggest that the challenge is not simply access to tools, but confidence in how they should be used safely, appropriately, and effectively. A substantial proportion of respondents either report an absence of clear organisational policy or uncertainty about whether such a policy exists. This suggests that formal guidance is not consistently visible to employees, although the survey cannot distinguish between organisations lacking policies and those communicating them ineffectively.

Figure 13. Which of the following are the main obstacles to your potential use of AI in your daily work activities? Select all that apply.

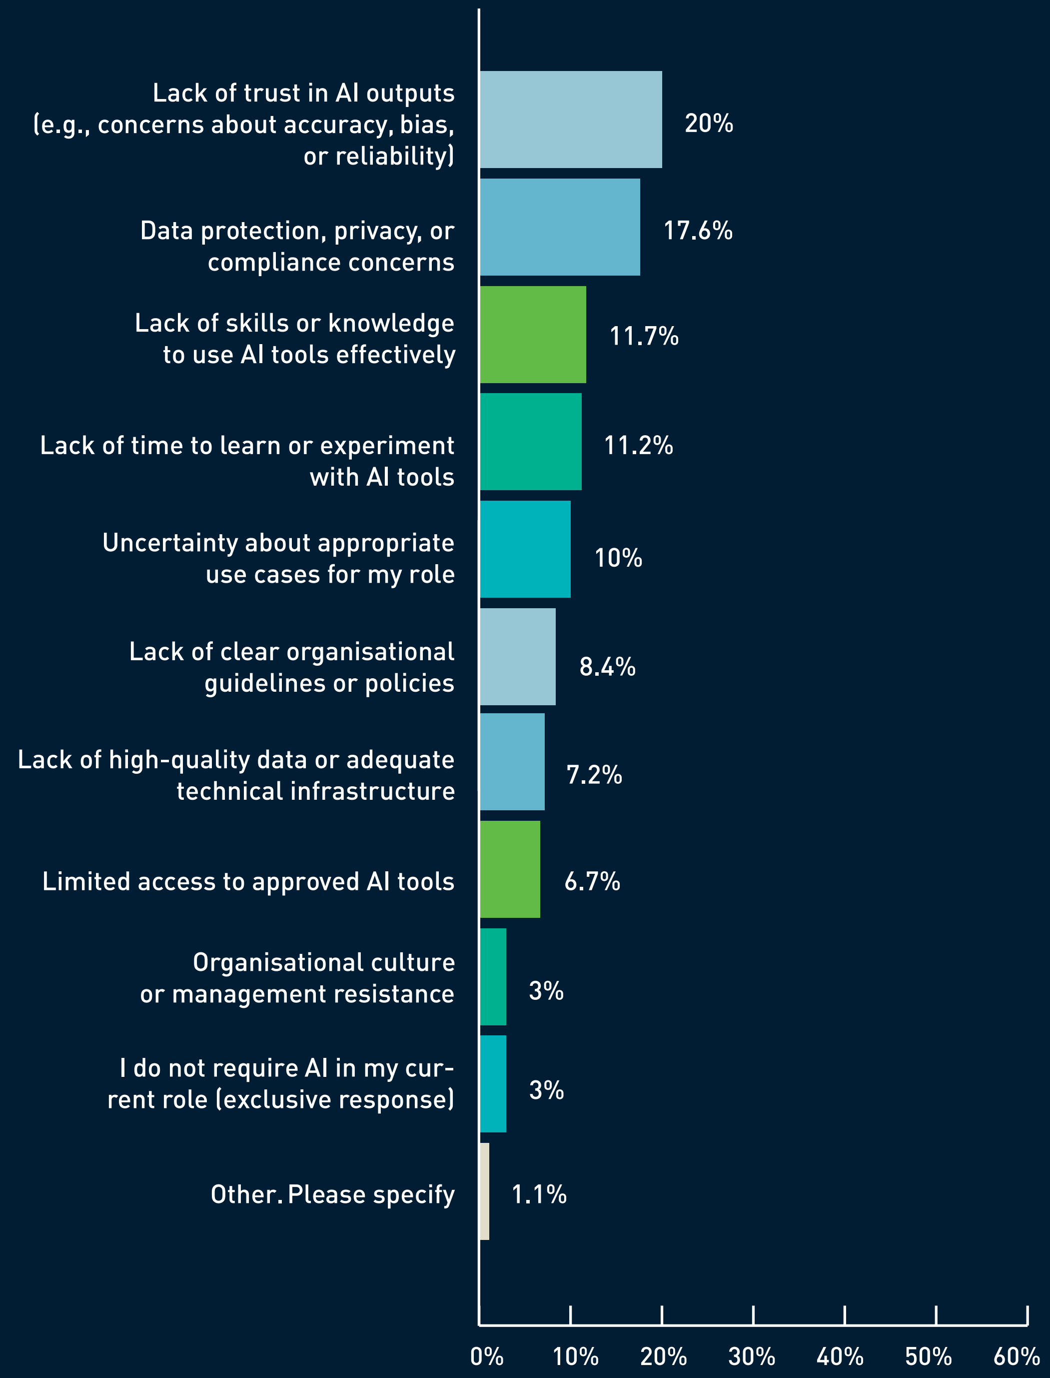

Figure 14. Which of the following best describes your organisation's policy on using AI tools at work?

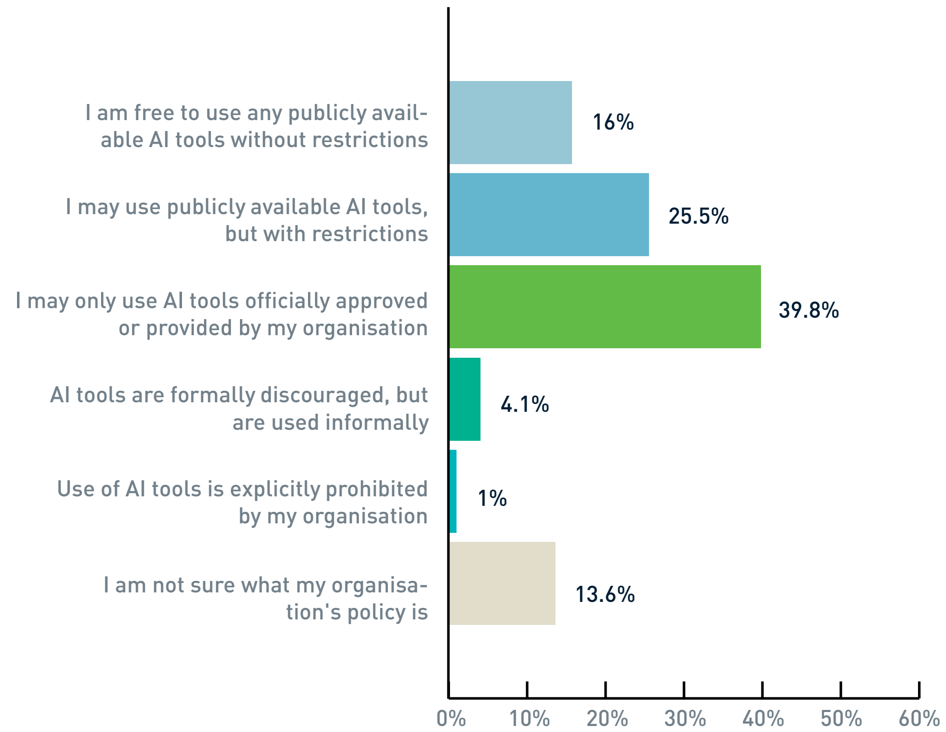

Figure 15. What type of internal AI policy would you recommend for organisations in your sector?

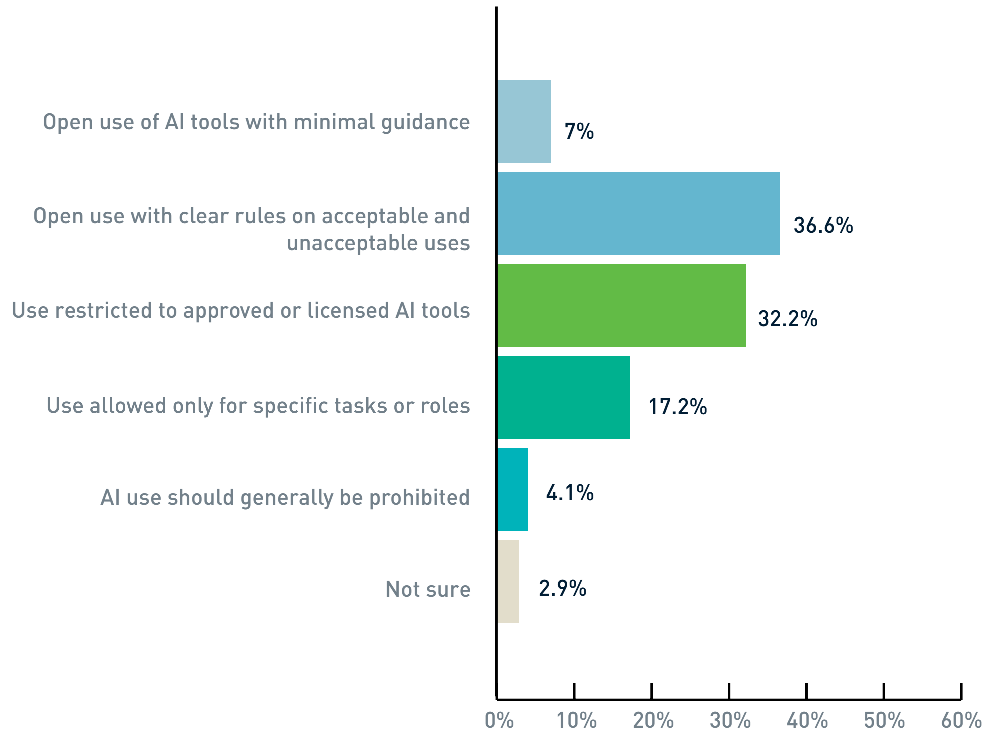

4.5. Uneven Adoption Across Sectors and Roles

Descriptive comparisons indicate that reported AI use, confidence, and participation in AI-related development vary across sectors, roles, and organisation sizes. Higher levels are more visible among respondents in knowledge-intensive sectors and technical occupations, while adoption appears more uneven in traditional sectors and smaller organisations. These patterns should be interpreted cautiously because subgroup sizes vary and the sample was not designed to produce precise population estimates for every sector or organisational category.

Figure 16. How often do you use AI-based tools in your work activities?

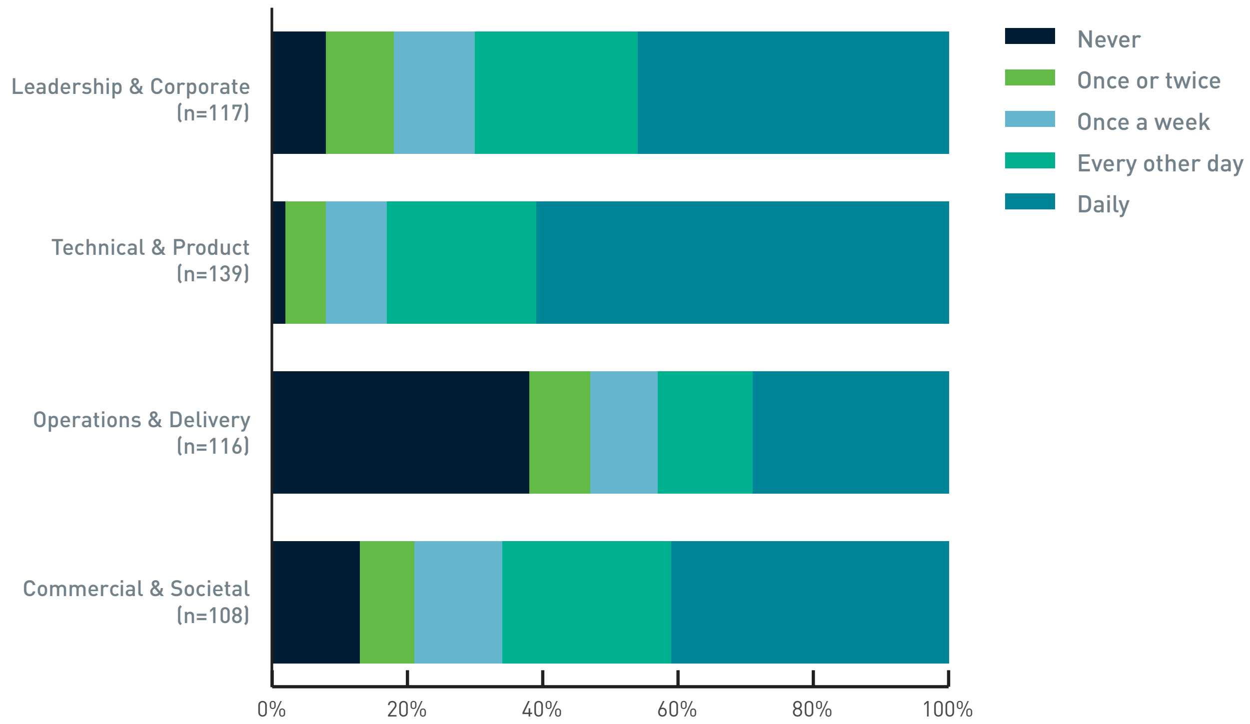

Figure 17. How confident are you in your ability to use AI tools effectively in your working tasks?

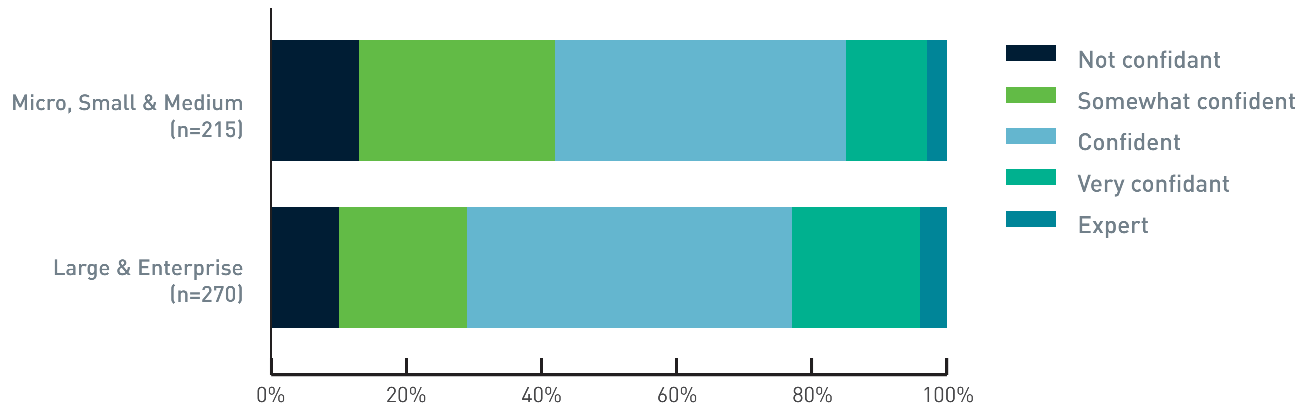

Figure 18. Have you taken any steps to adapt your skills in relation to use of AI in your work?

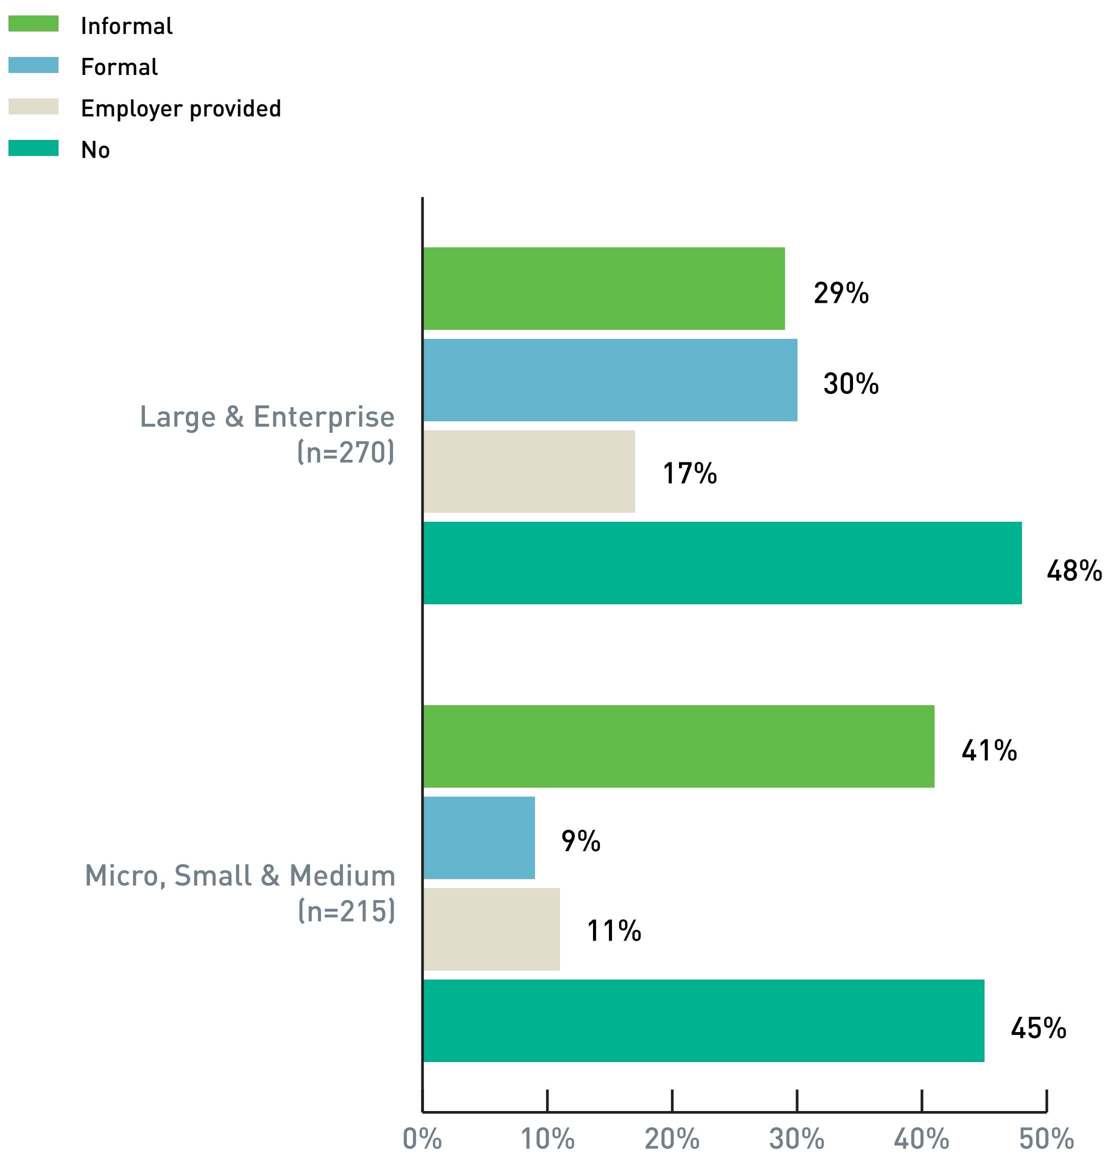

Training participation also varies by organisation type and sector. This suggests that access to development opportunities may be uneven and may reflect differences in organisational resources, digital maturity, and the availability of formal learning pathways.

Figure 19. Have you taken any steps to adapt your skills in relation to use of AI in your work?

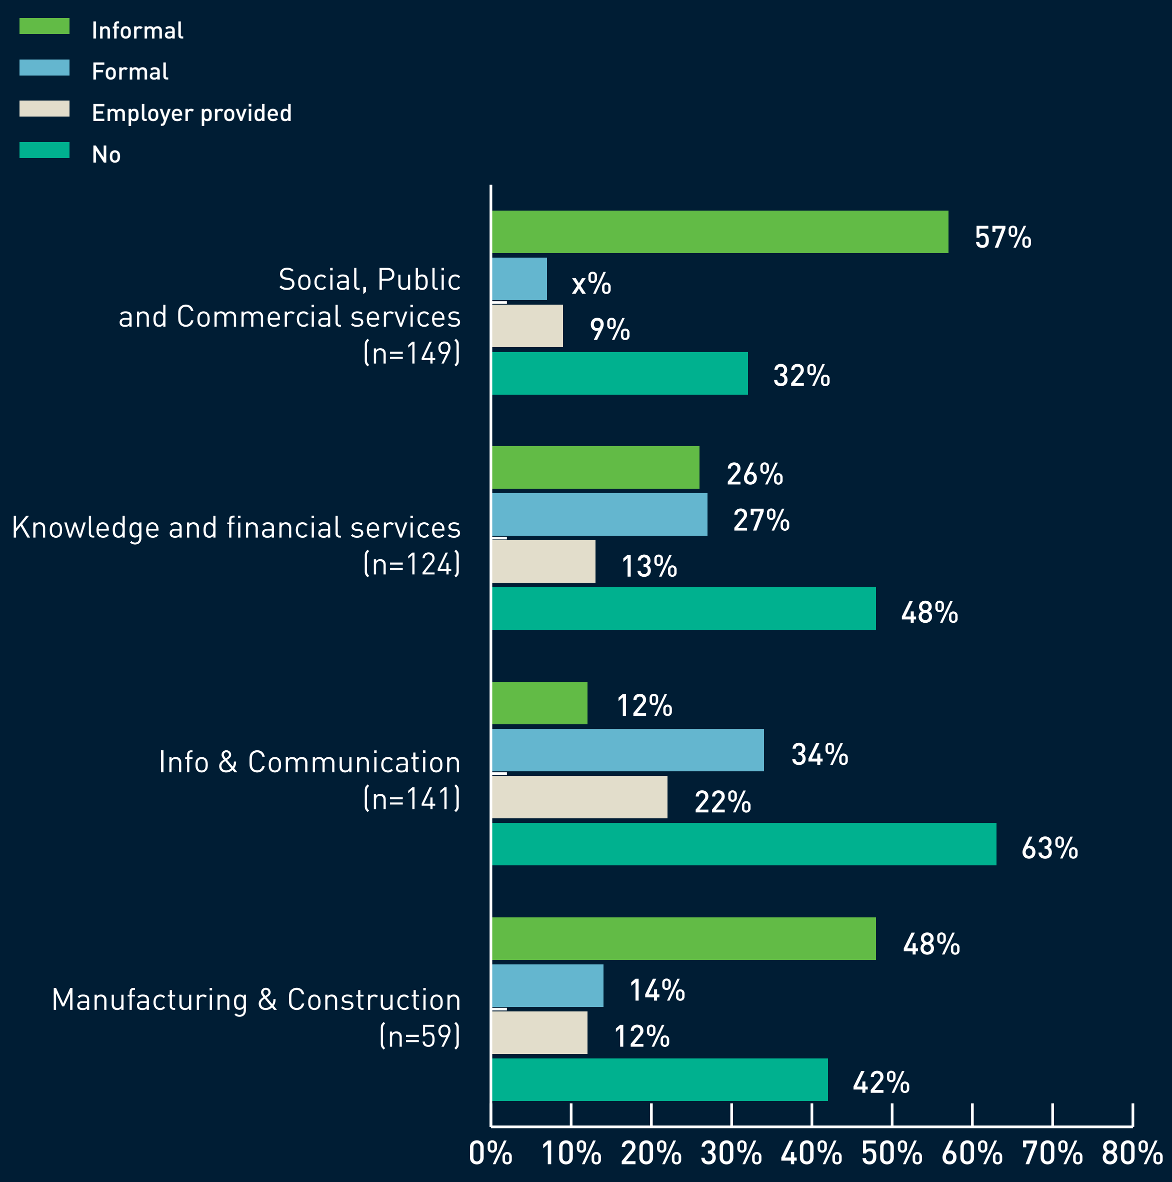

These differences point to a potential risk of uneven capability development, in which workers in better-resourced organisations and AI-intensive sectors may gain access to tools and learning opportunities more quickly than others.

4.6 Key Survey Conclusions

Taken together, the survey findings indicate that AI is already a common feature of working life for many respondents.

Reported use appears to be developing more rapidly than formal training and clearly communicated organisational guidance.

Respondents are generally confident about using AI and adapting their skills, but confidence is not consistently accompanied by participation in structured learning.

They also express conditional trust in AI-generated content and identify accuracy, privacy, misuse, and uncertainty about appropriate use as continuing concerns.

The findings therefore support practical, role-specific capability development, clearer organisational policies, and targeted assistance for sectors and organisations with fewer resources. Because the achieved sample is concentrated in AI-intensive sectors, larger organisations, and Dublin, these patterns should be interpreted as evidence from a diverse cross-sector workforce sample rather than as precise estimates for the entire Irish workforce.

Insights from the Experimental Vignettes

05 Insights from the Experimental Vignettes

5.1. Experiment 1: Conditions Shaping Intended AI Use

Participants (N = 421) were randomly assigned to one of six experimental conditions in a 2 × 3 between-subjects design. Each participant read a hypothetical workplace scenario involving one of two tasks: writing an internal email summarising a meeting or preparing a product description for a client that could inform an important decision. The scenario also specified one of three forms of AI involvement: AI-assisted editing, an AI-generated final version, or unauthorised AI use. Participants then rated (1) how likely they would be to use AI for the task, (2) how likely they would be to verify AI-generated content before using it, and (3) how responsibility should be ranked if the AI-supported output were incorrect or caused harm. Because these are responses to hypothetical scenarios, the results capture stated intentions and judgements rather than observed workplace behaviour.

Figure 20 and 21 How likely are you to use AI for this task?

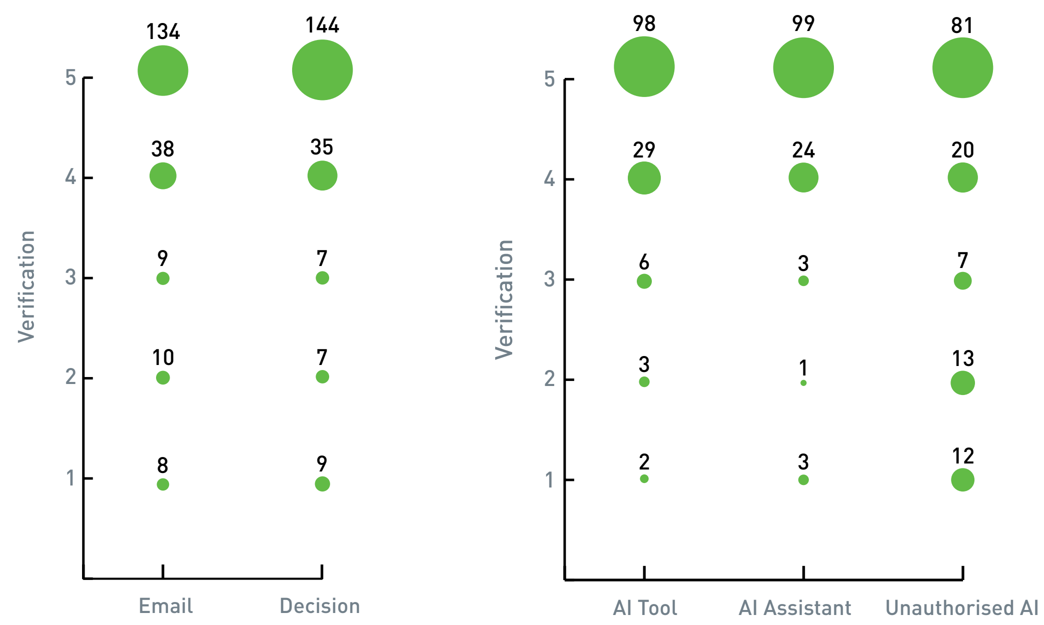

Note: Y-axis represents a five-point Likert scale ranging from 1 = Extremely unlikely to 5 = Extremely likely.

Figures 20–21 present participants' stated likelihood of using AI across the randomised conditions. Figure 20 compares responses across the two task types, while Figure 21 compares the three forms of AI involvement. The descriptive pattern shows little difference between the internal-email and client-facing product-description tasks. By contrast, participants reported a lower likelihood of using AI when its use was described as unauthorised than when AI was authorised either for editing or for generating a final version. This suggests that the organisational legitimacy of AI use may matter more for stated willingness than the distinction between the two tasks examined here.

Figure 22 and 23 How likely are you to check or verify the AI output in case you decide to use it?

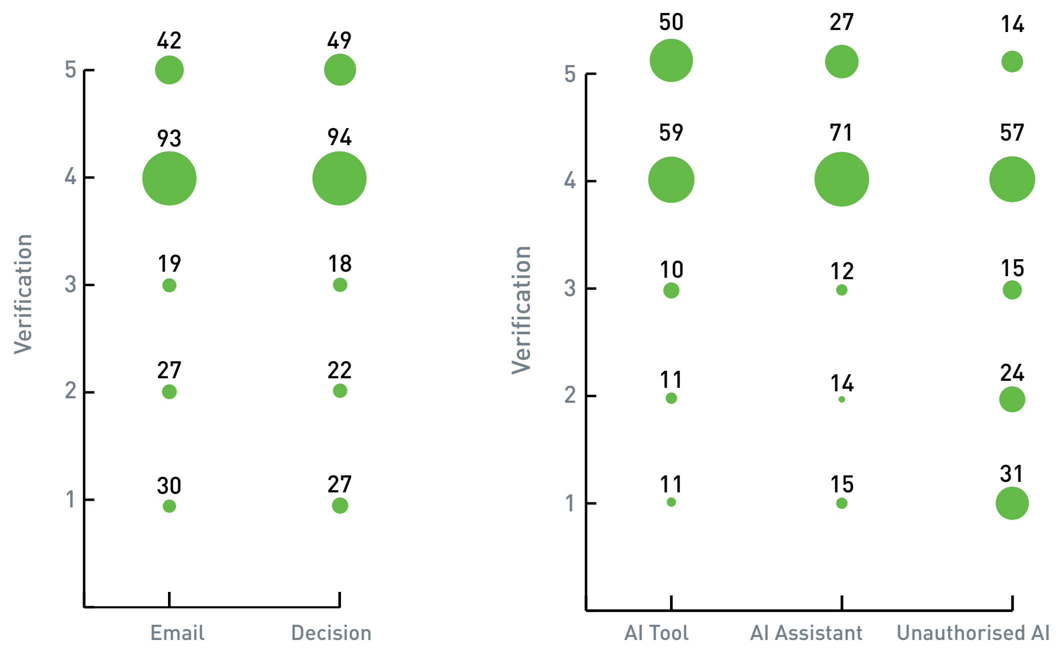

Note: Y-axis represents a five-point Likert scale ranging from 1 = Extremely unlikely to 5 = Extremely likely.

Figures 22–23 show participants' stated likelihood of checking or verifying AI-generated outputs, again by task type (Figure 22) and form of AI involvement (Figure 23). Twenty participants indicated that they would not use AI under any circumstances. Among those open to using it, intended verification was consistently high across conditions, producing a ceiling effect. The experiment therefore provides limited leverage for identifying differences in verification intentions. It also cannot establish whether respondents would verify outputs equally carefully in real workplace settings.

Figure 24 and 25. In case you use the AI output and the outcome is wrong or causes harm, how would you assign responsibility?

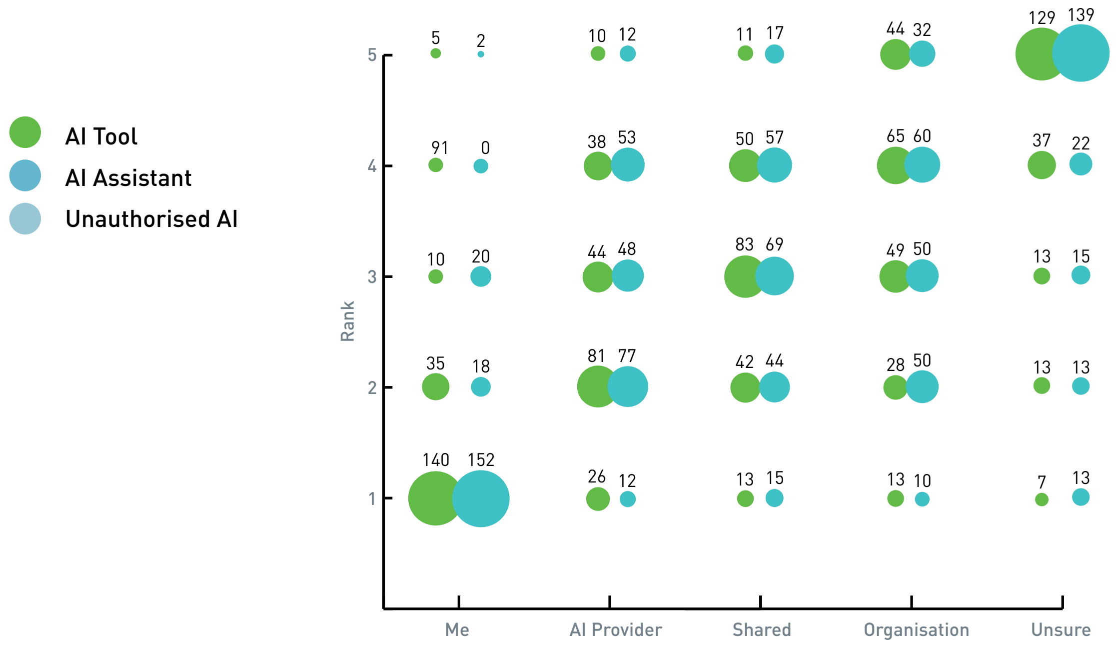

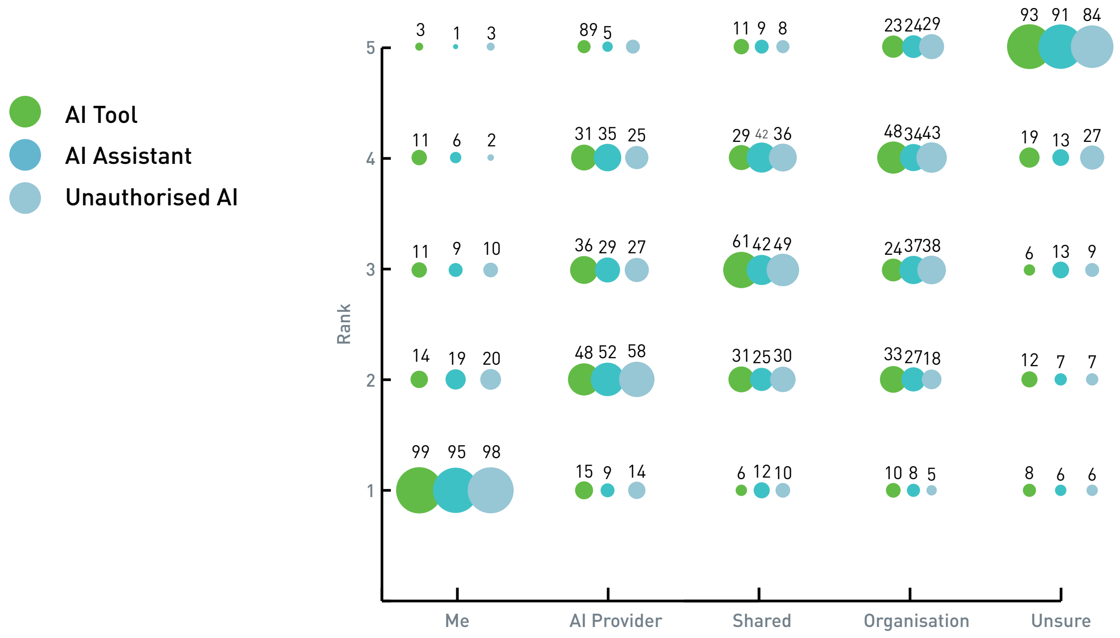

Figures 24–25 present the ranking of responsibility across the experimental conditions. Participants ranked five options: themselves, the AI provider, shared responsibility, the organisation, and an 'Unsure' category. Across conditions, respondents generally assigned the greatest responsibility to themselves, followed by the AI provider and then shared responsibility, while the organisation received the lowest attribution. The presence of 'Unsure' responses also indicates that responsibility was not straightforward for all participants to allocate. These findings concern perceived responsibility in hypothetical scenarios and should not be interpreted as a statement about the legal allocation of liability.

5.2. Experiment 2: Evaluating Performance and AI Use

Participants (N = 421) were randomly assigned to one of eight conditions in a 2 × 4 between-subjects design. Each participant read a hypothetical workplace scenario describing a colleague whose performance was either high or low and whose approach to AI use varied across four conditions. Participants evaluated the colleague's competence, the extent to which the approach supported long-term professional growth, the appropriateness of the AI use, and their willingness to delegate the task to that colleague. As in Experiment 1, these measures capture evaluations of a hypothetical colleague rather than behaviour towards an actual co-worker.

Figure 26. How competent do you consider Person A to be in their role?

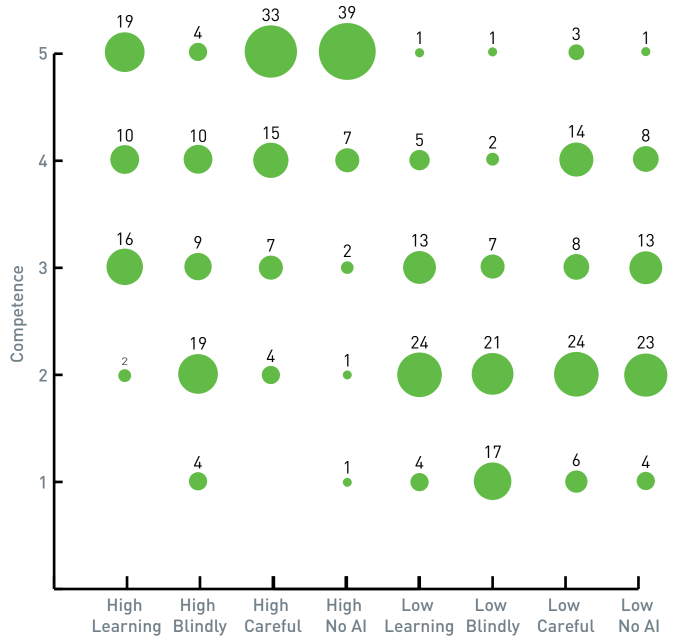

Figure 27. How well does Person A's approach support long-term professional growth?

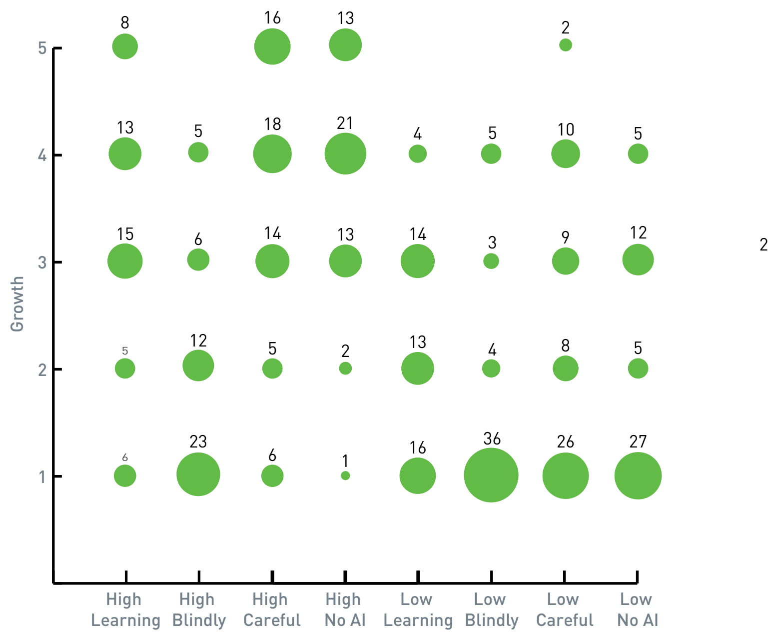

Figures 26–29 summarise participants' evaluations across the eight experimental conditions. Figure 26 presents perceived competence; Figure 27 assesses whether the colleague's approach supports long-term professional growth; Figure 28 shows the perceived appropriateness of the colleague's AI use; and Figure 29 reports participants' stated willingness to delegate the task to the colleague.

Figure 28. How appropriate is Person A's use of AI tools in their situation?

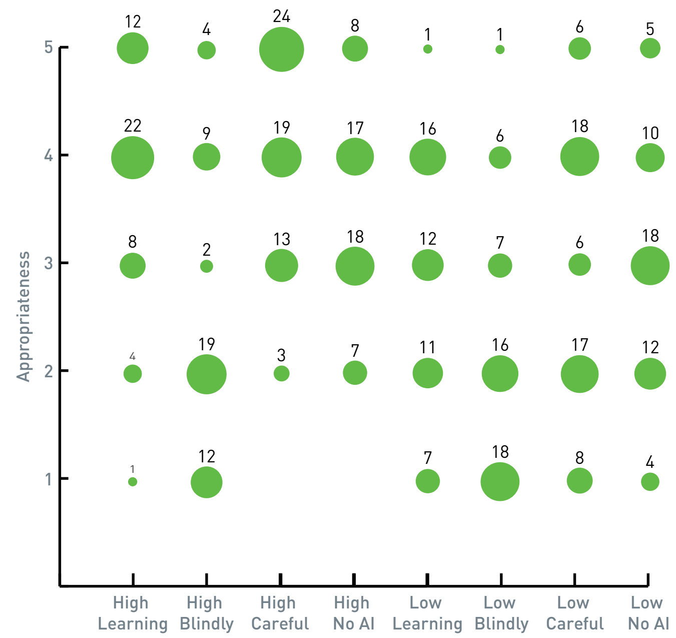

Figure 29. If this task were yours, how likely would you be to assign it to Person A instead of doing it yourself?

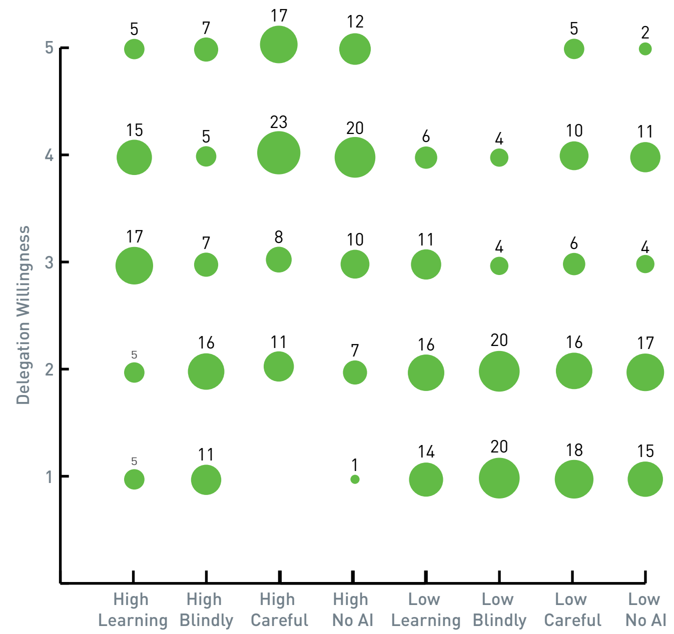

Across all four outcomes, the clearest descriptive pattern is that the colleague was evaluated more positively when performance was described as high.

Ratings were also comparatively favourable when AI was used thoughtfully or not used, whereas 'blind' AI use received lower ratings.

Participants therefore judged AI use together with observed performance and the manner of use, rather than in isolation.

The figures show descriptive patterns; statistical claims should be based on the corresponding inferential analysis, where reported.

Senior Leadership Interview Findings

06 Senior Leadership Interview Findings

6.1. Changing skill requirements under AI

Interviewees described AI adoption as changing the balance of capabilities required at work. Across sectors, they emphasised a growing need for critical evaluation, systems-level thinking, communication, and the ability to review and refine AI-generated outputs. In their accounts, workers are increasingly expected to combine domain expertise with the capacity to formulate problems, assess outputs, and decide when human intervention is required. Several participants also **described adaptation as involving the revision of established routines and assumptions, rather than the acquisition of technical skills alone**.

The interviews did not point to a single, economy-wide shortage of advanced technical talent. Instead, participants described **uneven readiness across roles**, teams, and organisations. They frequently distinguished specialist technical expertise from a broader AI-ready mindset: curiosity, adaptability, informed experimentation, and critical judgement. **Most interviewees currently framed AI as a complement to human work and a means of increasing capacity**, while recognising that its effects vary by task, organisational context, and implementation quality.

6.2. AI adoption patterns and impact on tasks

Interviewees described substantial variation in organisational maturity. In many organisations, adoption remains concentrated in routine and administrative activities through general-purpose or co-pilot tools. Examples included drafting, report generation, information retrieval, data extraction, and support across functions such as human resources, finance, and customer service. **Participants commonly viewed these uses as freeing time for more complex work**, although the interviews did not independently measure the productivity effects.

A smaller group of participants, particularly in **technology and engineering, also discussed movement towards agentic systems capable of carrying out connected sequences of activity.** Examples included elements of client onboarding and software-development workflows. By contrast, participants from more traditional sectors more often described exploratory, decentralised, or narrowly focused use. **These accounts indicate an uneven transition** from individual task assistance towards workflow-level integration, rather than a uniform progression across the economy.

6.3. Changing skill requirements under AI

Participants generally described AI as redesigning roles through the reallocation of tasks rather than immediately replacing whole occupations. Routine and entry-level activities, including basic coding and first-line technical support, were frequently cited as areas of automation or augmentation. Interviewees suggested that this can **shift human work towards problem solving, validation, exception handling, coordination, and oversight**. At the same time, they raised **concerns that removing entry-level tasks** may weaken established pathways through which junior workers acquire experience.

The interviews also identified **emerging hybrid responsibilities** involving the coordination of human colleagues and AI systems. Some technical professionals were described as moving towards a **"player-coach" role** in which they supervise AI-supported work while remaining directly involved in delivery. Participants reported isolated job reductions, but often linked these to wider restructuring, outsourcing, or changing business models rather than to AI alone. The evidence therefore supports cautious conclusions about role redesign and emerging oversight functions, not a general claim that displacement risks have been eliminated.

6.4. Worker responses, behavioural dynamics, and human–AI collaboration

Interviewees described varied worker responses, ranging from enthusiasm and experimentation to anxiety, uncertainty, and disengagement. Some participants perceived older workers, part-time staff, and employees with limited access to training as being at greater risk of exclusion, although the interview sample does not allow population-level conclusions about particular demographic groups. Resistance was often attributed to poor implementation, inadequate integration with existing work, unclear benefits, or the absence of contextualised and role-specific training.

A recurring governance concern was "shadow AI": informal or invisible use of unapproved tools. Participants described employees using publicly available systems, sometimes on personal devices, to complete tasks they regarded as repetitive or inefficient. These accounts suggest a tension between bottom-up experimentation and organisational requirements for data security, privacy, and accountability. Clear approved pathways for experimentation may therefore help organisations retain employee initiative while reducing unmanaged use.

6.5. Transparency, governance, and responsible use

Participants consistently identified governance, accountability, and appropriate transparency as important conditions for responsible adoption. They placed particular emphasis on external or client-facing contexts in which AI use may materially affect a decision, service, or representation of professional work. Internally, the appropriate level of disclosure was viewed as more context-dependent. Across these discussions, interviewees generally maintained that organisations and **human decision-makers must retain responsibility for reviewing outputs and managing consequences**; the interviews should not be interpreted as determining legal liability.

The EU AI Act generated mixed responses. Participants recognised the value of a common framework for higher-risk applications and responsible use, while some expressed concern about complexity, compliance costs, and uncertainty. **SMEs and third-sector organisations were seen as particularly likely to need accessible guidance and implementation support because they may lack specialist legal and governance capacity**. The interviews therefore suggest demand for regulation accompanied by practical, proportionate compliance assistance.

6.6. Organisational enablers and barriers to AI adoption

Data and organisational infrastructure were among the most frequently discussed barriers to scaling AI. Participants referred to fragmented systems, inconsistent data quality, legacy silos, and difficulties linking data across functions. They also highlighted the **need to reconcile AI deployment with data-protection requirements.** These constraints were described as limiting the movement from pilots to reliable operational use, particularly where organisations lacked the resources to redesign systems or develop tailored solutions.

Interviewees also described differences in leadership capacity and financial resources between large organisations and SMEs. Large multinationals were seen as better able to invest in infrastructure, specialist personnel, governance, and ongoing operating costs. Smaller firms could sometimes move more quickly, but were also more exposed to cost, uncertainty, and limited internal expertise. Participants consequently emphasised shared learning, practical advisory support, and targeted financial assistance as possible ways to prevent capability gaps from widening.

6.7 Effective implementation practices

Participants strongly favoured practical, continuous, and role-specific learning over generic or one-off training. Hands-on sessions, including guided experimentation with relevant tasks, were described as helping employees understand both the usefulness and limitations of generative AI. **The interviews suggest that training is most credible when it is linked to actual workflows**, includes verification and data-handling practices, and is supported by opportunities for continued learning.

Leadership and organisational culture were also described as important enablers. Several participants highlighted internal AI champions, communities of practice, and centres of excellence as mechanisms for sharing knowledge and supporting experimentation. Clear security guidance and access to approved tools were viewed as creating a safer environment for learning. These practices were presented as promising approaches reported by interviewees, rather than as experimentally established solutions.

6.8. Conclusions from the Interviews

Across the interviews, AI was described less as a standalone software upgrade than as a change requiring organisations to reconsider how work is designed and coordinated. Participants frequently described a shift in some activities from first-draft production towards reviewing, editing, validating, and integrating machine-generated material. **This does not imply that human creation is disappearing; rather, it indicates that evaluation and judgement are becoming more visible components of AI-supported workflows**. Adaptability, domain expertise, and critical assessment were therefore emphasised alongside technical familiarity.

Interviewees often distinguished an initial augmentation phase, centred on co-pilot tools for common tasks, from more advanced workflow integration involving agentic systems. **They anticipated that wider delegation could shorten some production cycles and shift attention towards outcomes, verification, and exception handling**. Participants also raised concerns about entry-level development where foundational tasks are automated, and about business models such as **billable-hour services where faster delivery may change how value is defined and priced**. These observations represent informed expectations and early organisational experiences, not established economy-wide outcomes.

The interviews also show that adoption is a behavioural and organisational process, not simply an IT deployment. **Participants described considerable bottom-up experimentation, including unapproved use, alongside differences in confidence and perceived benefit**. Some saw AI as creating opportunities for more challenging work, while others described anxiety or exclusion. **These accounts reinforce the need for inclusive training**, employee involvement, and clear organisational guidance. They do not support the stronger inference that apparent acceptance necessarily reflects a failure to understand AI's potential effects.

Structural constraints were repeatedly linked to data quality, legacy systems, governance capacity, and regulatory uncertainty. Participants viewed these conditions as shaping whether AI could move beyond isolated pilots. They also perceived a risk that better-resourced organisations could progress more quickly than SMEs and non-profit organisations. The scale of such a divide cannot be established from the interview sample, but the recurring concern is consistent with a broader maturity gap in adoption.

Taken together, the interviews support a targeted rather than generic approach to implementation. Participants favoured identifying specific workflow problems, providing protected opportunities for hands-on learning, establishing approved tools and clear accountability, and developing internal networks that can diffuse knowledge. **Their accounts suggest that organisational advantage will depend less on access to a particular model than on the capacity to prepare data**, redesign workflows, develop people, and govern human–AI collaboration responsibly.

6.9. Senior Leadership Overall Analysis – Themes and Subthemes

In addition to the qualitative analysis of the interviews, we conducted a complementary exploration using **large language models (LLMs)**. Initially, the interview data were manually reviewed, leading to the identification of seven overarching themes along with several sub-themes. For each theme and sub-theme, we defined specific indicators to support and structure the qualitative coding process. Figure (across) presents the analytical framework developed to guide the LLM-based categorisation of the interview data.

Subsequently, we applied a prompt-based classification approach using the GPT-5.4-mini model to analyse the interview transcripts (excluding the interview questions). The model was instructed to code the data across multiple sub-themes in a single pass, following a set of strict coding rules designed to ensure consistency, exclusivity, and precision in segment identification. The prompt required the model to extract verbatim segments that clearly matched predefined sub-themes, while applying conservative inclusion criteria to minimise false positives.

Figure 30. Overview of the analytical framework outlining the main themes and corresponding sub-themes identified through qualitative analysis.

The LLM-assisted coding provides a structured, complementary overview of how predefined themes and sub-themes appear across the interviews. Because the procedure depends on the initial qualitative framework, prompt design, and model classification, the **resulting counts should be interpreted as descriptive indicators rather than definitive measures** of prevalence or importance.

Across firm-size and sectoral groupings, several recurring patterns are visible. **AI was commonly framed as a complement to productivity rather than a direct substitute** for labour, alongside frequent discussion of routine-task automation and workflow change. These patterns are consistent with the manual thematic analysis, but the figures should not be treated as estimates for the wider population of Irish organisations.

Larger organisations show more coded material relating to governance, oversight, and structured implementation, including formal mechanisms and leadership-led change. Micro, small, and medium-sized organisations show relatively more discussion of experimentation, strategic gaps, and resource constraints. These descriptive differences may reflect the composition and content of the interview sample as well as genuine organisational variation.

The sectoral breakdown similarly suggests that organisational context shapes adoption. Information and communication and professional-services interviews contain more references to advanced applications, hybrid roles, and continued training, while manufacturing, construction, and public-facing services show more limited or uneven patterns. **Subgroup sizes differ, so these comparisons should be read as exploratory**.

Across groups, the analysis repeatedly identifies changing skill requirements, including the combination of technical familiarity, critical judgement, communication, and systems-level thinking. References to resistance, unequal access, and shadow AI further indicate that adoption is organisational and social as well as technical.

Figure 31. Distribution of LLM-coded interview segments across themes and sub-themes by company size. The figures provide a descriptive overview of how themes are expressed across interviews and organisational contexts.

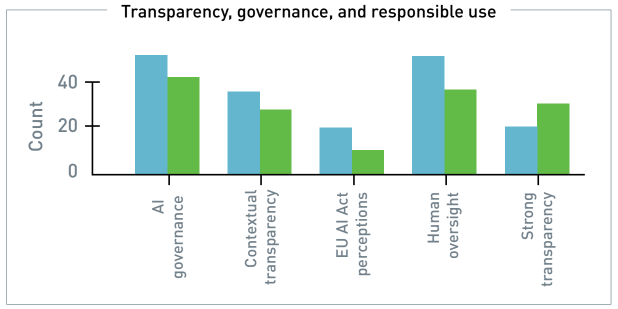

Company Size: Micro, Small, Medium | Large

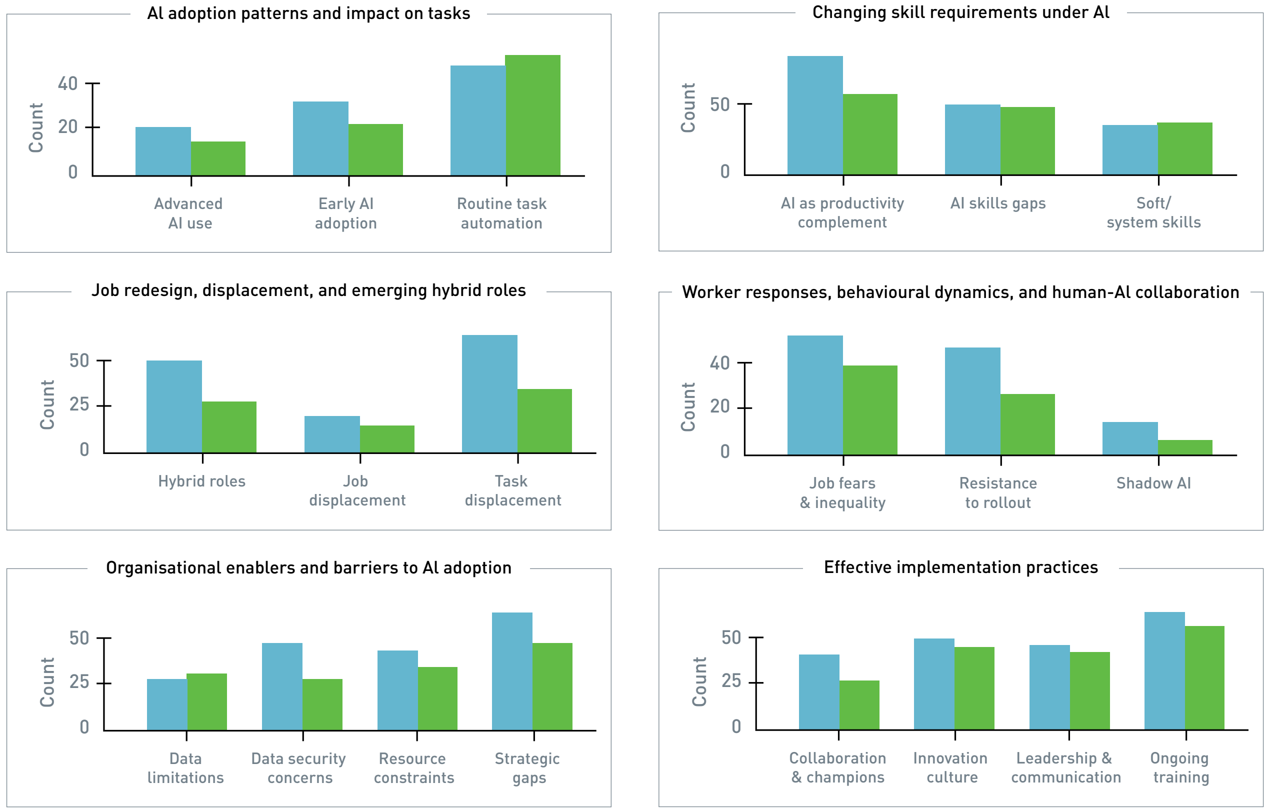

Figure 32. Distribution of LLM-coded interview segments across themes and sub-themes by sectors. The figures provide a descriptive overview of how themes are expressed across interviews and organisational contexts.

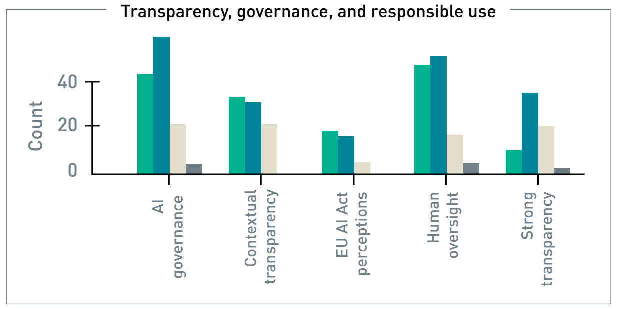

Sector

- Finance & Professional Services
- Info & Communication
- Manufacturing & Construction
- Retail, Education, Health & Public

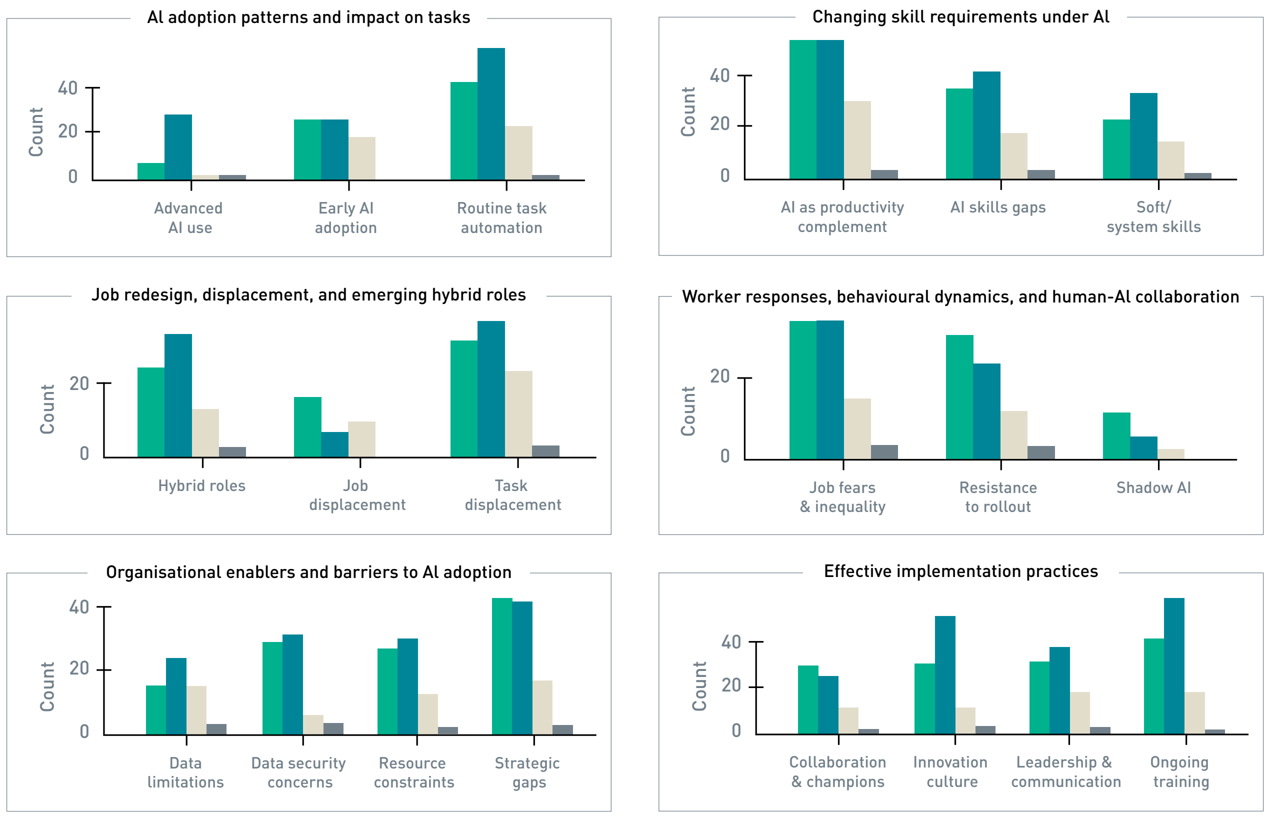

Policy Documents Analysis

07 Policy Documents Analysis

7.1. Overview of Documents Included

This analysis examines the institutional and policy context shaping the development and use of AI in Ireland, situating the Irish case within the European Union framework and comparing it with the United Kingdom.

The corpus was constructed from official government, parliamentary and supranational publications, including regulations, strategy documents, advisory reports and implementation plans addressing AI, automation, digital transformation, or related institutional conditions such as data protection, copyright and digital markets. Searches were conducted on official government and agency websites between January and February 2026, with the final search completed on 26 February 2026. The corpus contains 48 documents, listed in Appendix E.

For descriptive purposes, the documents were grouped into three broad categories according to their formal status and degree of enforceability:

- **Regulation:** legally binding legislative or regulatory instruments, representing the highest level of formal enforcement in the corpus.
- **Strategy:** formal policy strategies, plans or frameworks that establish objectives and direction but are not necessarily legally binding.
- **Guideline:** non-binding reports, guidance documents and advisory materials that provide recommendations or general direction.

Figure 33 shows the distribution of the documents by publication year, covering 2016 to early 2026. The apparent rise in the number of documents over time should be interpreted partly in light of the rapid expansion of AI policy activity and partly in light of the corpus-construction process. No document published in 2017 was included, and only three documents from 2026 appear because the search was completed in February of that year.

Figure 34 presents the distribution by document type and jurisdiction. Within this curated corpus, EU documents are more strongly represented among regulations, while Irish documents are more frequently classified as strategies or guidelines. The UK accounts for a smaller share of the corpus; this reflects the study's comparative focus and document-selection criteria and should not be interpreted as a direct measure of the overall volume of UK policy activity.

The documents were obtained in PDF format, after which their text was extracted and normalised. For the similarity analysis, the corpus was cleaned and tokenised using spaCy (English-language model; Honnibal et al., 2020). Text was tokenised, lemmatised and converted to lower case. Stop words, punctuation and numeric tokens were removed, including a custom list of frequently occurring policy and legal terms such as 'shall', 'section' and 'article'. Domain-relevant terms, including 'EU' and 'AI', were retained.

These preprocessing decisions reduce the influence of high-frequency functional language and allow comparisons to focus more closely on substantive vocabulary. They also shape the resulting similarity scores and should therefore be understood as analytical choices rather than neutral transformations. The cleaned tokens were aggregated by jurisdiction and by document type, so each group was represented by a combined text for comparison.

Figure 33. Number of documents by Year.

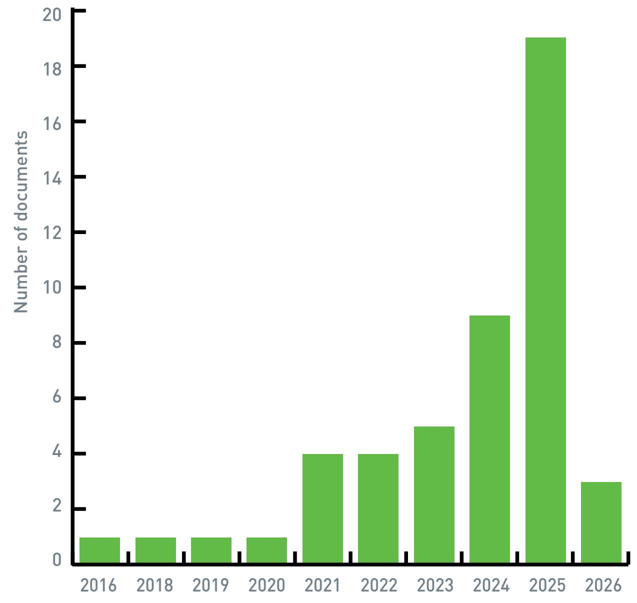

Figure 34. Number of documents by type and jurisdiction.

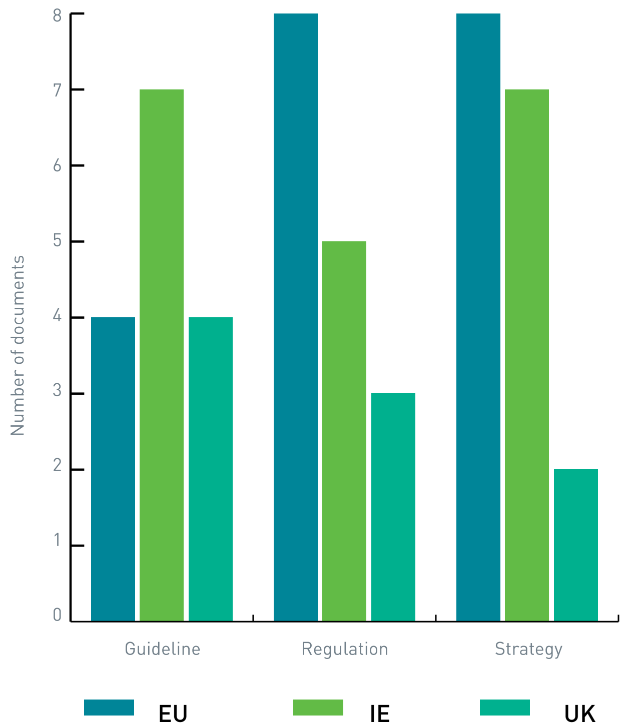

Figure 35. Similarity of documents categorised by type.

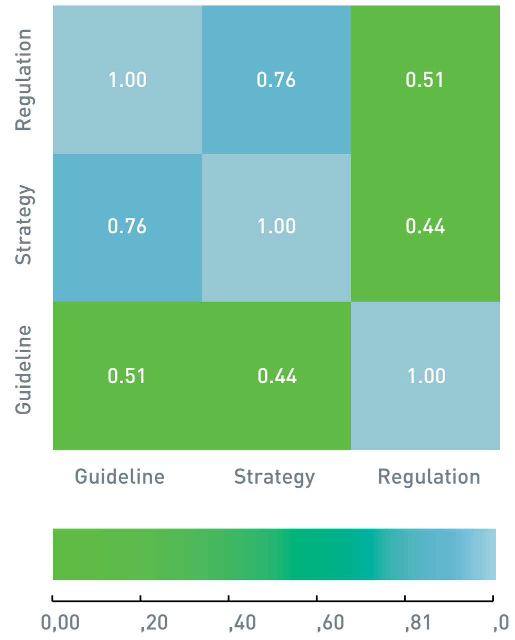

Figures 35 and 36 present cosine similarity between TF-IDF representations of the aggregated policy texts. Figure 35 compares document types: strategies and guidelines are the most similar pair (0.76), followed by regulations and guidelines (0.51), and regulations and strategies (0.44). Figure 36 compares jurisdictions. The highest similarity is between the EU and Irish corpora (0.83), followed by Ireland and the UK (0.65) and the EU and UK (0.59). These values indicate similarity in vocabulary within this corpus; they do not demonstrate equivalence in policy substance, implementation or regulatory effect.

Figure 36. Similarity of documents categorised by jurisdiction.

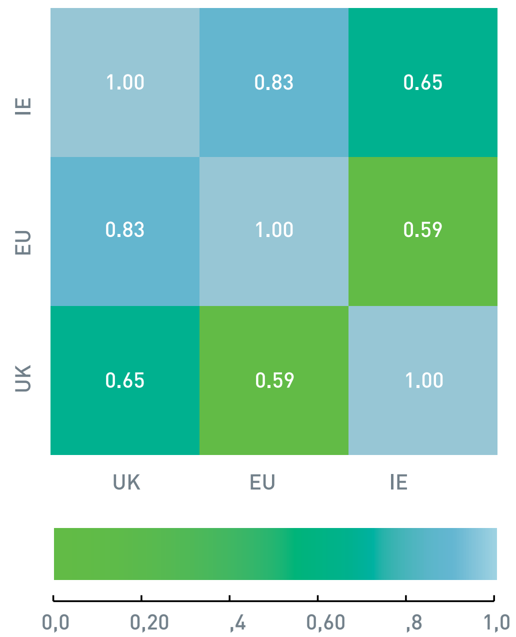

Note: Pairwise cosine similarity was computed between the TF-IDF (term frequency–inverse document frequency) vectors for each grouping. Each cell is the cosine of the angle between the corresponding group vectors (range 0–1; 1 indicates identical direction). The diagonal is 1.0 (each group compared with itself). Each group (e.g. EU, IE, UK, or Guideline, Regulation, Strategy) was represented by a single TF-IDF vector over a shared vocabulary (up to 5,000 terms).

7.2. Network Analysis of AI Discourse in Policy Documents

Figures 37-39 present co-occurrence networks constructed from adjacent words in AI-related passages within policy documents from Ireland, the EU and the UK. The networks provide a visual representation of which terms appear near one another and how clusters of associated vocabulary are structured within each jurisdiction's corpus. They should be interpreted as exploratory maps of discourse rather than direct measures of policy priority or implementation.

To construct the networks, a predefined list of AI-related keywords was used, including 'artificial intelligence', 'machine learning', 'deep learning', 'neural networks', 'generative AI', 'large language models', 'data-driven', 'AI', 'ML', 'LLM', 'LLMs', 'algorithm' and 'automation'. For each occurrence, a contextual window of 10 words on either side was extracted. Networks were then constructed by linking adjacent words, with nodes representing words and edges representing adjacency. Edge weights reflect the frequency with which word pairs co-occurred. For visual clarity and comparability, low-frequency edges were removed using thresholds selected so that each network contained approximately 70 to 90 nodes. Because the threshold is partly determined by visualisation needs, comparisons should focus on broad structural patterns rather than small differences in individual nodes or links.

Community detection was performed using a modularity-based method to identify clusters of terms that frequently co-occurred. In the visualisations, nodes with the same colour belong to the same detected community. These clusters can help identify recurring semantic groupings, but their interpretation remains dependent on the underlying corpus, preprocessing choices and selected threshold.

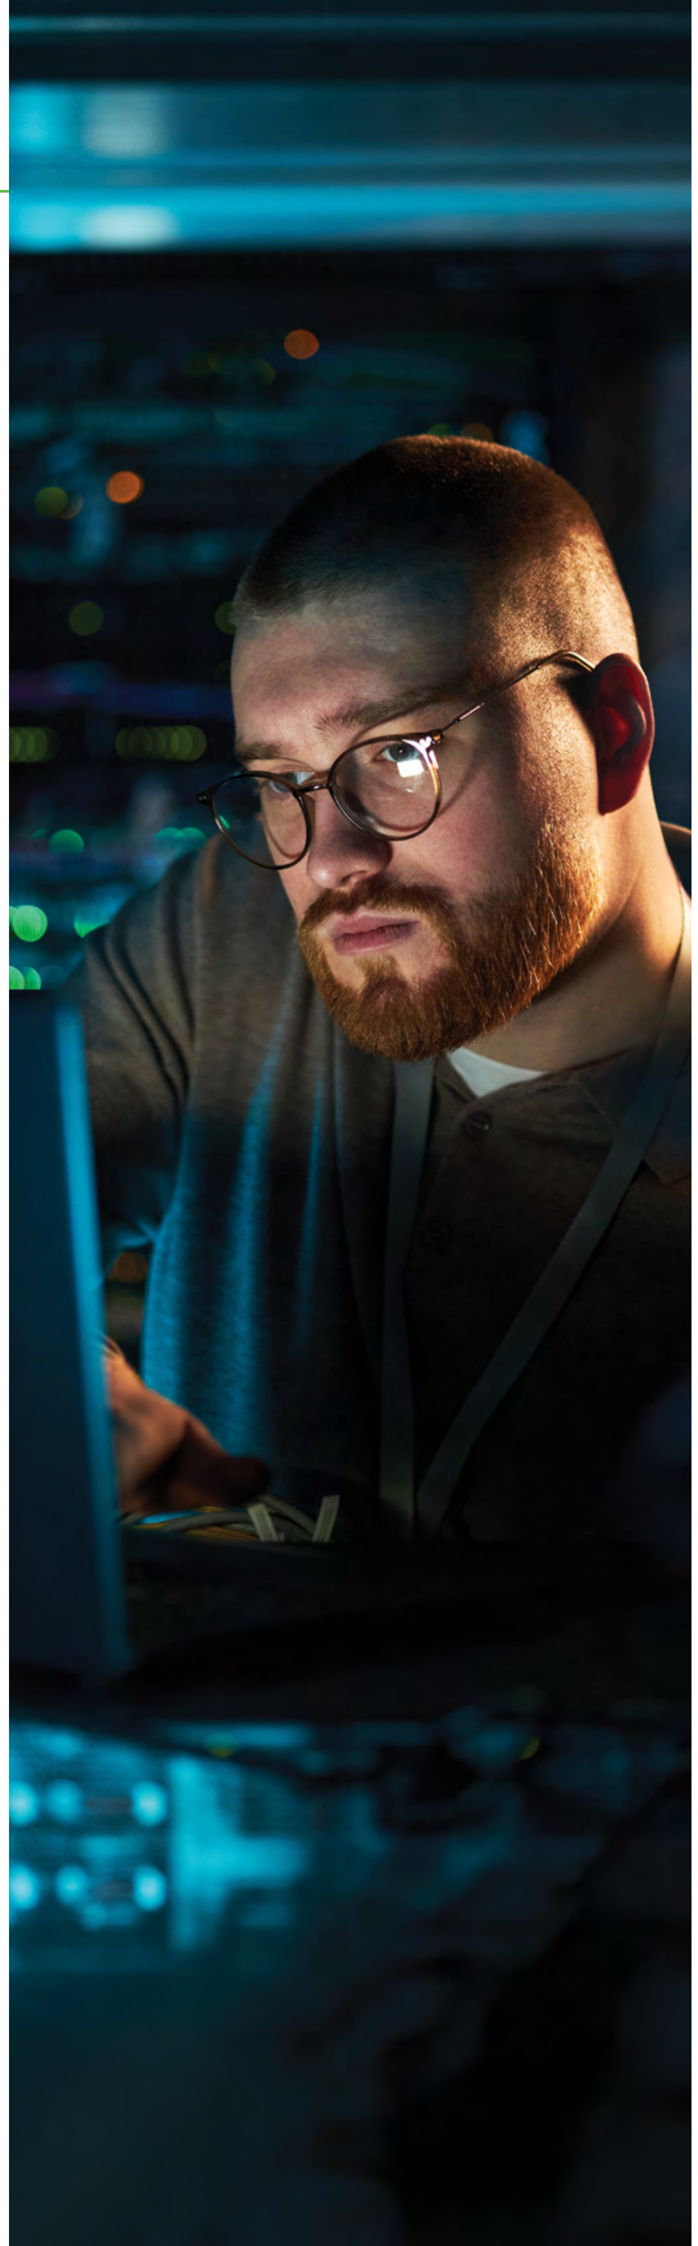

Figure 37. Co-occurrence network of adjacent words in policy documents from Ireland, constructed using a specified co-occurrence threshold. Edge thickness reflects the strength of co-occurrence between word pairs, and colours denote the detected communities.

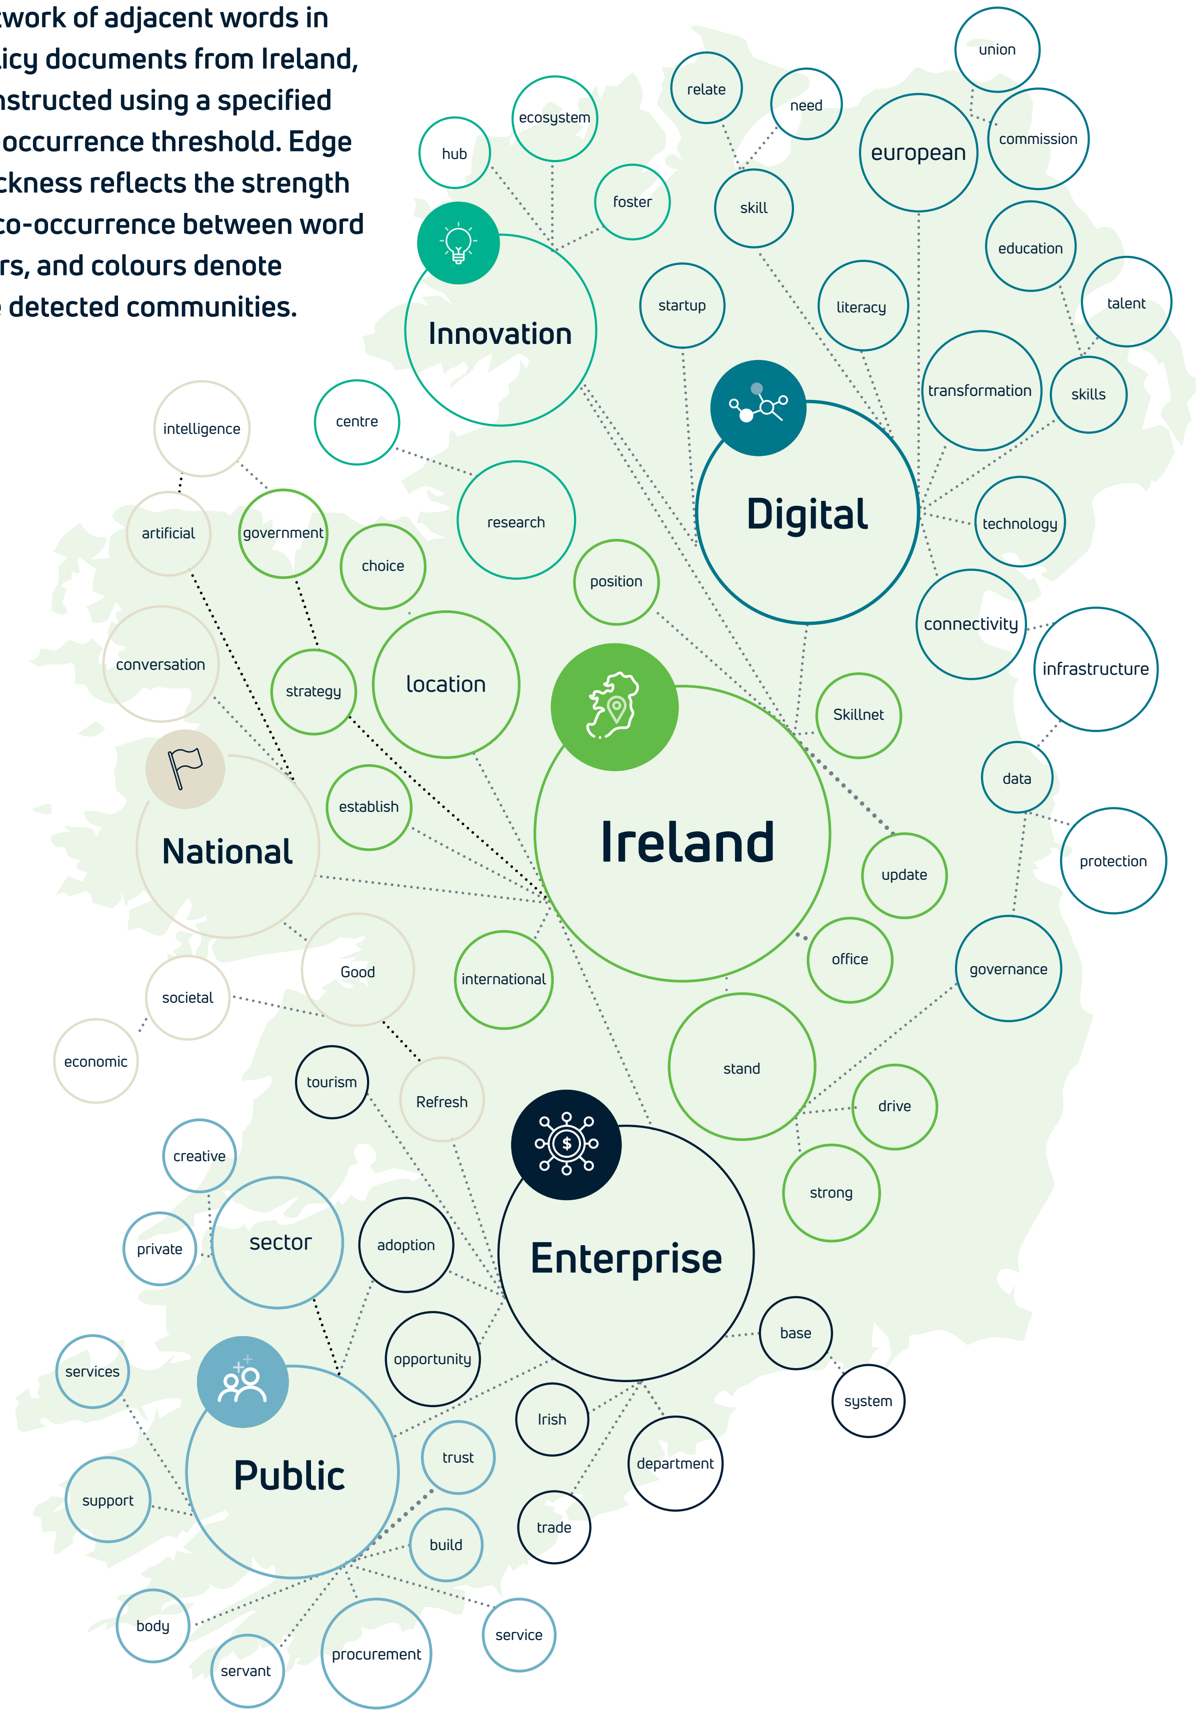

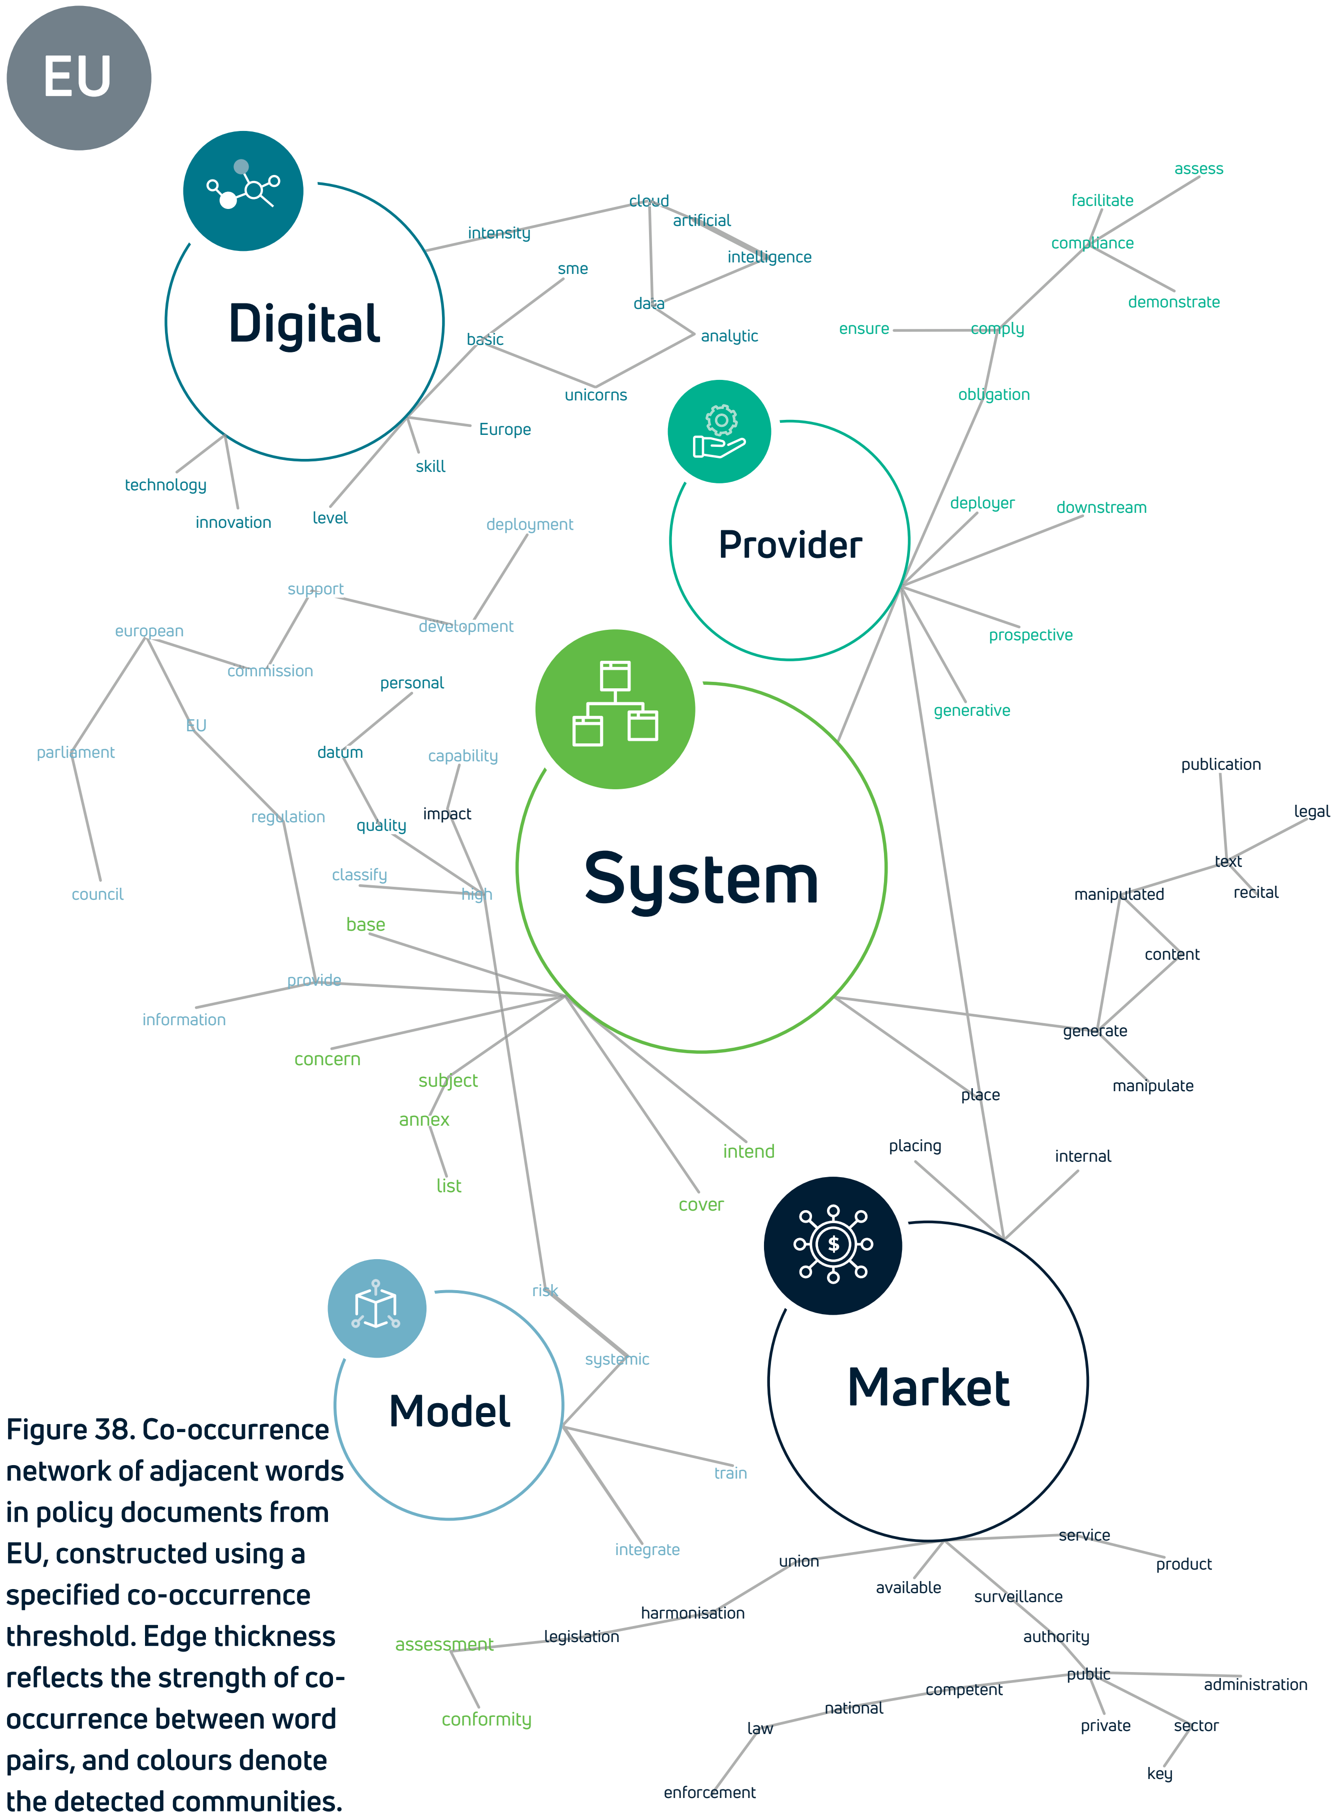

Figure 38. Co-occurrence network of adjacent words in policy documents from EU, constructed using a specified co-occurrence threshold. Edge thickness reflects the strength of co-occurrence between word pairs, and colours denote the detected communities.

Figure 39. Co-occurrence network of adjacent words in policy documents from UK, constructed using a specified co-occurrence threshold. Edge thickness reflects the strength of co-occurrence between word pairs, and colours denote the detected communities.

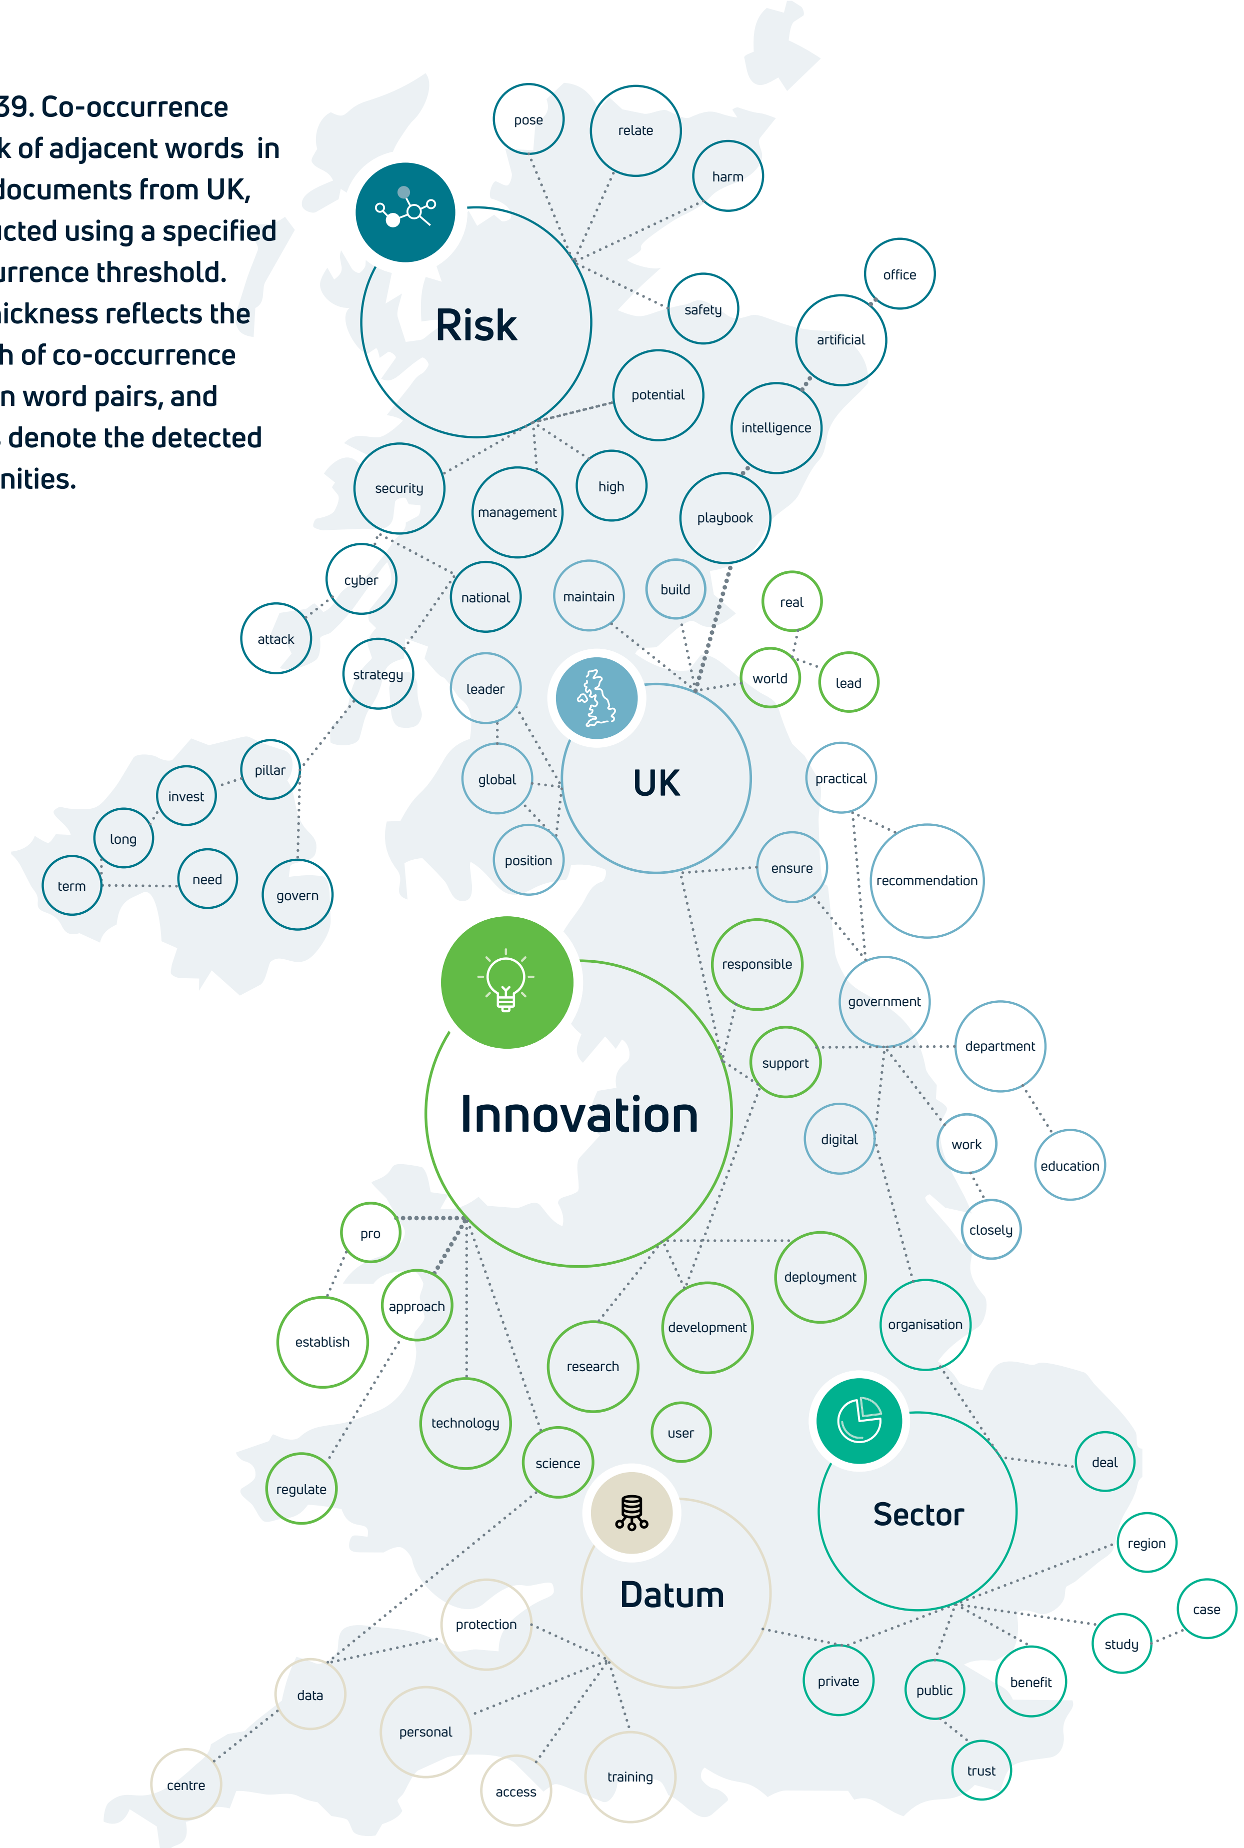

7.3. Topic Modelling of Policy Documents by Jurisdiction

To complement the network analysis, Latent Dirichlet Allocation (LDA) topic modelling was applied separately to the Irish, EU and UK document sets. Three topics were generated for each jurisdiction, and each topic is represented below by its most probable terms. The model is exploratory: topic labels are not produced automatically, and the listed words should be interpreted as indications of recurring lexical patterns rather than definitive or mutually exclusive policy themes. The relatively small and heterogeneous corpus also limits the stability of topic solutions.

IRELAND

- Topic 1: service, meaning, selection, video, jurisdiction, determine, provider, control, medium, effective
- Topic 2: ai, ireland, national, digital, public, strategy, use, intelligence, artificial, good
- Topic 3: ai, use, regulation, council, ensure, office, public, people, advisory, provide

EUROPEAN UNION

- Topic 1: algorithm, ai, digital, technology, objective, specific, intelligence, artificial, cybersecurity, european
- Topic 2: ai, eu, generative, support, european, digital, development, sector, generate, public
- Topic 3: ai, model, purpose, risk, general, provider, high, regulation, market, obligation

UNITED KINGDOM

- Topic 1: ai, innovation, approach, regulation, risk, use, regulator, regulatory, pro, uk
- Topic 2: ai, use, government, model, generative, datum, uk, artificial, intelligence, risk
- Topic 3: ai, uk, government, national, sector, strategy, use, technology, innovation, support

7.4. Comparative Insights

Within the Irish corpus, the network analysis shows a prominent 'digital' cluster connected to terms including 'skill', 'literacy', 'transformation' and 'education'. This is consistent with the whole-system emphasis attributed to Ireland's national AI strategy in OECD analysis and comparative research, which highlight the integration of digital, technical and complementary skills across education levels and institutional actors (Lassebie, 2023, p. 171; Shi, 2025, p. 462). **The computational results are therefore broadly compatible with the qualitative reading of the Irish policy framework**, although the topic model's word lists alone do not provide a strong basis for assigning precise substantive labels.

The Irish network also links 'digital' with 'European', which is consistent with the high lexical similarity between the Irish and EU corpora. By comparison, the UK network gives greater visual prominence to clusters organised around 'innovation', 'government', 'risk' and 'security'. These differences suggest distinct emphases in the language represented in the selected documents, but they should not be interpreted as evidence that one jurisdiction is more innovative or more risk-focused in practice without further qualitative and institutional analysis. The Irish network is the only one of the three visualisations in which 'enterprise' appears as a prominent node, alongside connected terms relating to the public sector, innovation and national strategy. Within this corpus, this pattern suggests that enterprise development and public-sector transformation are closely linked in Irish AI policy discourse. It does not, by itself, establish the extent or effectiveness of public-private partnership in implementation.

The EU network is more strongly organised around legal and regulatory vocabulary, including 'provider', 'deployer', 'obligation', 'compliance', 'high', 'risk', 'classify', 'market' and 'authority'. This is consistent with the EU's supranational regulatory role and the risk-based structure of recent AI legislation. The Irish corpus, by contrast, places relatively greater visible emphasis on national strategy, enterprise, digital transformation and skills, while the UK corpus combines innovation-oriented language with risk, security and government coordination. Taken together, the findings indicate broad lexical alignment between Ireland and the EU alongside differences in emphasis across jurisdictions. Because the analysis aggregates heterogeneous documents and measures language rather than policy outcomes, these results should be treated as descriptive evidence of policy discourse rather than an evaluation of regulatory effectiveness or implementation capacity.

Integration, Strategic Framework and Recommendations

08

Integration, Strategic Framework and Recommendations

This report set out to examine how Artificial Intelligence (AI) is reshaping work, skills, organisational practices, and workforce policy in Ireland. Drawing on a mixed-methods design that combined a structured literature review, a cross-sector survey of 487 workers in Ireland, embedded behavioural experiments involving 421 participants, 32 senior leadership interviews, and comparative analysis of 48 policy documents, the evidence points to a consistent overall pattern: AI use is already influencing work, but workers, organisations, sectors, and institutions are adapting at different speeds.

Across the evidence base, the clearest near-term pattern is the restructuring of tasks and workflows rather than immediate large-scale job displacement. AI is being used to automate selected routine and administrative activities, support professional work, and redistribute responsibility for drafting, checking, coordination, and decision support between workers and digital systems. The central challenge is therefore not only how many jobs may ultimately be created or displaced, but how organisations redesign work, preserve human capability, establish accountability, and distribute the benefits and risks of adoption.

8.1. Integrated Findings Across the Evidence Base

I. AI use is widespread, while reported dependency remains limited

Survey respondents report frequent AI use across a range of workplace activities, while many describe themselves as only minimally dependent on these tools and report limited change in their ability to work without them. **This combination suggests that AI is often being incorporated as a support within existing practices rather than experienced as indispensable**. Interviewees also described bottom-up experimentation, widespread use of generative-AI assistants, and early exploration of agentic systems. The evidence does not establish that workers underestimate their dependency, but it indicates that day-to-day use can develop before organisations have fully redesigned workflows or formalised governance.

II. Work is being restructured at task and workflow level more than jobs are being eliminated

Across the literature review, survey, experiments, and interviews, the most consistent pattern is the redistribution and reorganisation of activities within roles. AI is used to automate selected routine tasks, support drafting and analysis, and assist decision processes. These changes can also alter handovers, verification, exception handling, and accountability across wider workflows. Human work therefore increasingly includes reviewing, validating, coordinating, and taking responsibility for AI-supported outputs. The study does not provide direct longitudinal evidence of net job creation or displacement, but it indicates that current organisational change is more visible within work processes than in wholesale job replacement.

III. The capability challenge is practical, role-specific, and continuous

Survey respondents generally report confidence in using AI and openness to adapting their skills, but participation in formal training remains limited. The survey does not measure objective competence, so this should be understood as a gap between self-reported confidence and structured learning rather than as proof of inadequate capability. **Interviewees repeatedly emphasised that generic awareness training has limited value unless connected to real tasks**, workflows, and occupational responsibilities. Effective workforce development therefore requires practical, role-specific, and continuous learning that combines AI literacy with critical evaluation, domain judgement, data awareness, communication, and accountability.

IV. Productivity gains depend on organisational readiness & workflow design

The combined evidence indicates that AI does not generate productivity gains automatically. Benefits are more likely when organisations combine appropriate tools with leadership commitment, reliable data, workflow redesign, secure infrastructure, clear governance, and workforce capability. **Time saved in initial production may also be partly absorbed by checking, correction, and rework**. Productivity should therefore be assessed at workflow and organisational level rather than inferred from the speed of a single AI-supported task.

V. Adoption may widen differences across firms and workers

Larger and digitally mature organisations are generally better positioned to invest in infrastructure, governance, specialist expertise, and formal training. SMEs may benefit from shorter decision chains and greater flexibility, but often face tighter financial, technical, and managerial constraints. At worker level, access to tools and development opportunities is also uneven. These patterns point to a risk that AI adoption could widen capability, productivity, and opportunity gaps, although the present study does not measure such gaps longitudinally. **Targeted support is therefore important to prevent existing organisational advantages from becoming self-reinforcing.**

VI. Ireland is strategically well positioned but implementation remains uneven

The policy review indicates strong lexical alignment between Irish and EU policy documents and highlights Ireland's emphasis on digital skills, education, enterprise, and public-sector development. Comparative literature also identifies Ireland's relatively integrated approach to technical and complementary skills. However, the policy-text analysis does not itself measure implementation quality or policy outcomes. The principal challenge is therefore **to translate strategic ambition into accessible training, practical guidance, organisational support**, and proportionate compliance pathways across sectors and firm sizes.

8.2. Right-Sizing Framework for Human–Machine Partnerships in Irish Workplaces

The evidence suggests that successful AI adoption should not be treated as a technology deployment exercise alone. It requires coordinated decisions about tasks, workflows, people, organisational foundations, inclusion, and accountability. The proposed **right-sizing framework** is intended to help organisations match the type and degree of AI involvement to the work being performed, the capabilities of workers, the maturity of the organisation, and the level of risk.

I. Task and Workflow Fit: Automate, Augment, or Delegate Deliberately

Organisations should assess not only individual tasks but also the workflows in which they are embedded. Routine activities may be suitable for automation; tasks requiring human judgement may benefit from augmentation; and some multi-step processes may be delegated to agentic systems only where monitoring, escalation, and accountability are clearly defined. High-risk or consequential decisions should retain meaningful human oversight and an identifiable accountable decision-maker.

II. Workforce Capability: Build Applied AI Readiness and Preserve Expertise

Training should focus on practical use cases relevant to particular occupations and workflows. Core capabilities include AI literacy, critical evaluation of outputs, data awareness, secure tool use, judgement, communication, collaboration, and adaptability. Organisations should also monitor whether AI assistance supports learning or encourages excessive cognitive offloading, and preserve the expertise required to verify outputs, manage exceptions, and work effectively when AI is unavailable.

III. Organisational Foundations: Lead, Govern, Integrate, and Learn

Effective adoption requires **executive sponsorship, clear ownership, structured experimentation, reliable data, secure infrastructure, and operationally useful governance**. Policies should specify approved tools, permissible data use, verification requirements, escalation routes, and accountability. Governance should create safe and proportionate pathways for experimentation rather than merely restricting use. Internal AI champions, communities of practice, and centres of excellence can support organisational learning where they are connected to real operational needs.

IV. Inclusion and Fair Transition: Support Workers, SMEs, and Lower-Capacity Sectors

Support should be targeted towards SMEs, sectors with limited adoption capacity, workers with fewer training opportunities, and people whose roles are substantially restructured or displaced. Measures should recognise that SMEs can be flexible and innovative while still lacking specialist expertise, data infrastructure, and compliance resources. Fair-transition strategies should broaden access to capability building and **ensure that productivity gains do not depend on transferring hidden costs or risks to workers.**

V. Outcomes and Accountability: Measure What Matters

Success metrics should extend beyond tool adoption and headline time savings. Organisations should assess work quality, verification and rework, retained human capability, job quality, employee trust, innovation, fairness, data protection, resilience, and the distribution of benefits. **Transparency and disclosure should be proportionate to the material importance of AI use**, particularly where decisions affect clients, workers, safety, rights, or accountability.

8.3. Practical Recommendations

For Employers

1. Redesign work, not just pilot

Move beyond generic AI pilots by redesigning real tasks, handovers, and workflows around clearly defined organisational problems.

2. Build practical capability in the workflow

Provide practical, role-specific training embedded in daily work, including verification, safe data use, and the limits of AI systems.

3. Set clear policies and guardrails

Establish clear internal policies covering approved tools, permissible data, human oversight, accountability, and proportionate disclosure.

4. Create pathways for experimentation

Create approved pathways for experimentation through AI champions, communities of practice, or centres of excellence linked to operational needs.

5. Measure what truly matters

Measure productivity at workflow level, including time spent on verification, correction, rework, escalation, and capability development.

For education providers

Equip people for an AI-enabled future through relevant responsible and real-world learning

1. Integrate AI literacy across disciplines

Integrate AI literacy across disciplines while adapting content to the tasks, risks, and professional standards of different fields..

2. Expand flexible work-integrated pathways

Expand modular, flexible, and work-integrated learning pathways for people already in employment.

3. Set clear policies and guardrails

Combine technical instruction with critical thinking, domain judgement, communication, creativity, ethics, and accountability.

4. Create pathways for experimentation

Co-design provision with employers, workers, and professional bodies so that training reflects changing workflows rather than generic tool use.

For Policy Makers

1. Expand learning and reskilling

Expand lifelong learning and reskilling supports targeted at workers and occupations experiencing substantial task or role change.

2. Provide practical SME support

Provide practical advisory support, co-funded training, shared infrastructure, and proportionate compliance guidance for SMEs.

3. Promote trustworthy AI use

Support trustworthy and productive use through clear public guidance on capabilities, limitations, data protection, and accountability..

4. Ensure clear proportionate regulation

Ensure regulation and implementation guidance are clear, proportionate, and usable for smaller organisations and lower-capacity sectors.

5. Monitor labour-market impacts

Monitor labour-market, job-quality, and distributional outcomes to identify groups and regions at risk of exclusion or displacement.

8.4 Conclusion

AI is already reshaping work in Ireland, primarily through changes in tasks, workflows, skill requirements, and organisational practices rather than immediate wholesale job replacement.

The central question is no longer simply whether AI will affect the workforce, but whether organisations and institutions can design its use in ways that improve productivity and innovation while preserving judgement, learning, accountability, fairness, and job quality.

Ireland enters this transition with significant strengths, including a strong skills base, an internationally connected economy, and a policy framework that recognises the importance of technical and complementary capabilities. Success, however, will depend on implementation: moving from informal use to supported capability, from isolated pilots to workflow redesign, and from generic principles to clear operational governance. The organisations and institutions most likely to benefit will not necessarily be those using the most advanced systems, but those best able to right-size AI involvement and combine human expertise with intelligent tools in productive, trusted, and inclusive work systems.

References

1. Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), Handbook of labor economics (Vol. 4B, pp. 1043–1171). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

2. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3 |

3. 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

4. Autor, D. H., Katz, L. F., & Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? Quarterly Journal of Economics, 113(4), 1169–1213. https://doi.org/10.1162/003355398555874

5. Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics, 118(4), 1279–1333. https://doi.org/10.1162/003355303322552801

6. Babashahi, L., Barbosa, C. E., Lima, Y., Lyra, A., Salazar, H., Argolo, M., Almeida, M. A. de, & Souza, J. M. de. (2024). AI in the workplace: A systematic review of skill transformation in the industry. Administrative Sciences, 14(6), Article 127. https://doi.org/10.3390/admsci14060127

7. Bariach, B., Schoenegger, P., Bhaskar, M., & Suleyman, M. (2026). Seemingly conscious AI risks [Preprint]. SSRN. https://doi.org/10.2139/ssrn.6588659

8. Bazazi, S., Karpus, J., & Yasseri, T. (2025). AI's assigned gender affects human–AI cooperation. iScience, 28(12), Article 113905. https://doi.org/10.1016/j.isci.2025.113905

9. Broecke, S. (2023). Artificial intelligence and the labour market: Introduction. In OECD employment outlook 2023: Artificial intelligence and the labour market. OECD Publishing. https://doi.org/10.1787/08785bba-en

10. Brown, R., Sun, H., Cowling, M., & Dowell, D. (2025). Can we see it in the numbers yet? Artificial intelligence, firm-level productivity and SMEs [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5210948

11. Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

12. Brynjolfsson, E., Rock, D., & Syverson, C. (2017). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics (NBER Working Paper No. 24001). National Bureau of Economic Research. https://doi.org/10.3386/w24001

13. Caosun, M., & Aral, S. (2026). The augmentation trap: AI productivity and the cost of cognitive offloading [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2604.03501

14. Cui, H., & Yasseri, T. (2026). Gender bias in perception of human managers extends to AI managers. Computers in Human Behavior Reports, 21, Article 100984. https://doi.org/10.1016/j.chbr.2026.100984

15. Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. Journal of Economic Behavior & Organization, 211, 188–205. https://doi.org/10.1016/j.jebo.2023.05.008

16. Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality (Harvard Business School Working Paper No. 24-013). https://doi.org/10.2139/ssrn.4573321

17. Doorley, K., O'Connor, S., O'Shea, R., & Tuda, D. (2026). Artificial intelligence and income inequality in Ireland (Jointly-published Reports No. JR16). Economic and Social Research Institute and Department of Finance. https://doi.org/10.26504/jr16

18. Eisenmann, T. F., Karjus, A., Canet Sola, M., Brinkmann, L., Supriyatno, B. I., & Rahwan, I. (2025). Expertise elevates AI usage: Experimental evidence comparing laypeople and professional artists [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2501.12374

19. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), Article eadj0998. https://doi.org/10.1126/science.adj0998

20. European Commission. (2020). European enterprise survey on the use of technologies based on artificial intelligence. Publications Office of the European Union.

21. European Commission. (2023). Employment and social developments in Europe 2023: Addressing labour shortages and skills gaps in the EU. Publications Office of the European Union.

22. European Innovation Council and SMEs Executive Agency. (2024). Pact for Skills: Analysing of up- and reskilling policy initiatives and identifying best practices. Publications Office of the European Union.

23. Expert Group on Future Skills Needs. (2022). AI skills: A preliminary assessment of the skills needed for the deployment, management and regulation of artificial intelligence. Department of Enterprise, Trade and Employment.

24. Expert Group on Future Skills Needs. (2025). Skills insights note 2025-2: How AI is transforming the Irish labour market. Department of Enterprise, Trade and Employment.

25. Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). The impact of artificial intelligence on productivity, distribution and growth: Key mechanisms, initial evidence and policy challenges (OECD Artificial Intelligence Papers No. 15). OECD Publishing. https://doi.org/10.1787/8d900037-en

26. Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

27. Galindo, L., Perset, K., & Sheeka, F. (2021). An overview of national AI strategies and policies (OECD Going Digital Toolkit Notes No. 14). OECD Publishing. https://doi.org/10.1787/c05140d9-en

28. Green, A. (2024). Artificial intelligence and the changing demand for skills in the labour market. OECD Publishing. https://doi.org/10.1787/88684e36-en

29. Guenduez, A. A., Fuchs, S., & Demircioglu, M. A. (2025). AI in a comparative perspective: Linking policy formulation and implementation. Journal of Comparative Policy Analysis: Research and Practice, 27(5–6), 647–670. https://doi.org/10.1080/13876988.2025.2543984

30. Han, R., Cui, H., & Yasseri, T. (2026). Visual anthropomorphism shifts evaluations of gendered AI managers [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2602.17919

31. Holm, J. R., & Lorenz, E. (2022). The impact of artificial intelligence on skills at work in Denmark. New Technology, Work and Employment, 37(1), 79–101. https://doi.org/10.1111/ntwe.12215

32. Hong, S., Zhong, D., & Um, K. H. (2026). The impact of artificial intelligence adoption on operational performance in manufacturing. Journal of Manufacturing Technology Management, 37(2), 295–317. https://doi.org/10.1108/JMTM-03-2025-0227

33. Honnibal, M., Montani, I., Van Landeghem, S., & Boyd, A. (2020). spaCy: Industrial-strength natural language processing in Python [Computer software]. https://spacy.io

34. Ibec. (2025). AI Hub. https://www.ibec.ie/influencing-for-business/ibec-campaigns/ai-hub

35. IDA Ireland. (2023). AI talent is on the rise: New skills required to meet the demands of digital transformation. Labour Market Pulse.

36. Institute of International and European Affairs. (2024). Skills in the digital age: Adapting Ireland's workforce for the future.

37. Kim, J., Schweitzer, S., Riedl, C., & De Cremer, D. (2025). People reduce workers' compensation for using artificial intelligence (AI) [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2501.13228

38. Koh, L. Y., & Yuen, K. F. (2026). Infusion of generative artificial intelligence: Boosting organizational agility and effectiveness. Information, Communication & Society. Advance online publication. https://doi.org/10.1080/1369118X.2026.2648697

39. Kopka, A., & Fornahl, D. (2024). Artificial intelligence and firm growth: Catch-up processes of SMEs through integrating AI into their knowledge bases. Small Business Economics, 62(1), 63–85. https://doi.org/10.1007/s11187-023-00754-6

40. Lane, M., Williams, M., & Broecke, S. (2023). The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers (OECD Social, Employment and Migration Working Papers No. 288). OECD Publishing. https://doi.org/10.1787/ea0a0fe1-en

41. Lassebie, J. (2023). Skill needs and policies in the age of artificial intelligence. In OECD employment outlook 2023: Artificial intelligence and the labour market. OECD Publishing. https://doi.org/10.1787/08785bba-en

42. Liu, G., Christian, B., Dumbalska, T., Bakker, M. A., & Dubey, R. (2026). AI assistance reduces persistence and hurts independent performance [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2604.04721

43. Mila, Vector Institute, & Amii. (2025). Artificial intelligence adoption by small- and medium-sized enterprises: Insights from G7 case studies and Canada's experience.

44. Muehlemann, S. (2025). Artificial intelligence adoption and workplace training. Journal of Economic Behavior & Organization, 238, Article 107206. https://doi.org/10.1016/j.jebo.2025.107206

45. OECD, Boston Consulting Group, & INSEAD. (2025). The adoption of artificial intelligence in firms: New evidence for policymaking. OECD Publishing. https://doi.org/10.1787/f9ef33c3-en

46. OECD. (2023). OECD employment outlook 2023: Artificial intelligence and the labour market. OECD Publishing. https://doi.org/10.1787/08785bba-en

47. OECD. (2025). Bridging the AI skills gap: Is training keeping up? OECD Publishing. https://doi.org/10.1787/66d0702e-en

48. Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2302.06590

49. Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. Organizational Behavior and Human Decision Processes. Advance online publication. https://doi.org/10.2139/ssrn.5205850

50. Schimmelpfennig, R., Díaz, M., Prabhakaran, V., & Davani, A. (2025). Humanlike AI design increases anthropomorphism but yields divergent outcomes on engagement and trust globally [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2512.17898

51. Serôdio, P. (2026). AI and the UK labour market: The evidence so far. British Progress.

52. Shen, J. H., & Tamkin, A. (2026). How AI impacts skill formation [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2601.20245

53. Shi, L. (2025). Global perspectives on AI competence development: Analyzing national AI strategies in education and workforce policies. Human Resource Development Review, 24(4), 447–476. https://doi.org/10.1177/15344843251332360

54. Siedschlag, I., & Duran, J. (2026). Artificial intelligence and firm-level productivity: Early evidence from a small open economy (Productivity Institute Working Paper No. 63). https://doi.org/10.5281/zenodo.18256767

55. Skillnet Ireland. (2025). Ireland's talent landscape 2025: Future skills challenges of Irish business. https://www.skillnetireland.ie/insights/irelands-talent-landscape-2025

56. Stein, M. (2026). How are AI agents used? Evidence from 177,000 MCP tools [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2603.23802

57. Technology Ireland Digital Skillnet. (2024). Generative AI masterclass launch with London Business School. https://digitalskillnet.ie/news/cat-1/102/generative-ai-masterclass-launch-with-london-business-school-/

58. Technology Ireland Digital Skillnet. (2026). Technology Ireland Digital Skillnet launches AI Capability Accelerator for senior leaders. https://www.digitalskillnet.ie/news/cat-1/105/technology-ireland-digital-skillnet-launches-ai-capability-accelerator-for-senior-leaders-/

59. Trinity College Dublin, & Microsoft Ireland. (2026). Trinity College Dublin and Microsoft Ireland research shows a widening AI maturity gap between SMEs and large organisations. Microsoft Source EMEA.

60. Trinity College Dublin. (2026). AI economy in Ireland 2026: AI adoption index—Benchmarking and impact. Trinity Business School and Microsoft Ireland.

61. Violante, G. L. (2018). Skill-biased technical change. In The new Palgrave dictionary of economics (pp. 12389–12394). Palgrave Macmillan. https://doi.org/10.1007/978-1-349-58802-2_1534

62. Wang, Y., Zhong, M., Sharma, P., & Niu, D. F. (2025). The impact of human-centred artificial intelligence on firms' workforce demand in Industry 5.0 transition. Enterprise Information Systems. Advance online publication. https://doi.org/10.1080/17517575.2025.2595706

63. Workday. (2024a). Closing the AI trust gap.

64. Workday. (2024b). Global workforce report: Restoring meaning and connection at work.

65. Workday. (2025). Elevating human potential: The AI skills revolution.

66. Workday. (2026). Beyond productivity: Measuring the real value of AI.

67. World Economic Forum. (2025). The future of jobs report 2025.

68. Wu, Q., Qalati, S. A., Tajeddini, K., & Wang, H. (2025). The impact of artificial intelligence adoption on Chinese manufacturing enterprises' innovativeness: New insights from a labor structure perspective. Industrial Management & Data Systems, 125(3), 849–874. https://doi.org/10.1108/IMDS-06-2023-0378

69. Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., & Wang, B. Z. (2026). Firm data on AI (NBER Working Paper No. 34836). National Bureau of Economic Research. https://doi.org/10.3386/w34836

Appendices

A. Survey Questions

Employment Sector

Sector

Which sector best describes your main occupation?

- ○ Manufacturing (e.g., food products, chemicals, pharmaceutical products, basic metals, computers, etc)
- ○ Construction (e.g., construction of buildings, civil engineering)
- ○ Information and communication (e.g., publishing activities, telecommunications, computer programming, etc)
- ○ Financial and insurance activities (e.g., insurance, pension funding, trust, funds, etc.)
- ○ Real estate activities
- ○ Professional, scientific and technical activities (e.g., legal and accounting activities, consultancy activities,
- ○ scientific research and development, etc)
- ○ Education
- ○ Healthcare and Social Work Activities
- ○ Wholesale and Retail; Accommodation and Food Service activities (e.g., hospitality, tourism)
- ○ Public Administration and Defence
- ○ Other. Please specify: ______________________

Organisation Size

What is the size of the organisation where you work?

- ○ Micro: 1–9 people
- ○ Small: 10–49 people
- ○ Medium: 50–249 people
- ○ Large: 250–999 people
- ○ Enterprise / Corporate: 1,000+ people

Role

Which organisational function best describes your primary role?

- ○ Executive leadership (e.g. CEO, COO, Founder)
- ○ Human Resources / People
- ○ Finance / Accounting
- ○ Legal / Compliance / Risk
- ○ Operations
- ○ IT / Engineering
- ○ Data / AI / Machine Learning
- ○ Product / Innovation
- ○ Sales / Marketing
- ○ Research / Policy
- ○ Healthcare/ Clinical
- ○ Education/ Teaching
- ○ Retail/ Frontline Service
- ○ Other. Please specify: ______________________

Task

Which types of tasks make up most of your day-to-day work

Physical/ manual tasks

- ○ Administrative/ clerical tasks
- ○ Analytical/ problem-solving tasks
- ○ Creative/ design tasks
- ○ Managerial/ supervisory tasks
- ○ Interpersonal/ service/communication tasks
- ○ Technical/ operational tasks (e.g., machinery, IT systems)
- ○ Other. Please specify: ______________________

A. Survey Questions

AI Usage

AI Use

How often do you use AI-based tools in your work activities?

- Daily
- Every other day
- Once a week
- Only once or twice before
- Never

AI Awareness Sector

To the best of your knowledge, are AI tools being used by professionals in your area of work?

- Yes, by many professionals
- Yes, by a few professionals
- Yes, by at least one professional
- No, not that I am aware of
- I don't know

AI Policy

Which of the following best describes your organisation's policy on using AI tools at work?

- I am free to use any publicly available AI tools without restrictions
- I may use publicly available AI tools, but with restrictions (e.g., no confidential, personal, client, or commercially
- sensitive data; limited to certain tasks)
- I may only use AI tools officially approved or provided by my organisation (e.g., Microsoft Copilot, internal LLMs)
- AI tools are formally discouraged, but are used informally in practice
- Use of AI tools is explicitly prohibited by my organisation
- I am not sure what my organisation's policy is

AI Oversight

What type of internal AI policy would you recommend for organisations in your sector?

- Open use of AI tools with minimal guidance
- Open use with clear rules on acceptable and unacceptable uses
- Use restricted to approved or licensed AI tools
- Use allowed only for specific tasks or roles
- AI use should generally be prohibited
- Not sure

AI Proficiency

How confident are you in your ability to use AI tools effectively in your working tasks?

- Not confident – I have little or no experience
- Somewhat confident – I use AI tools with guidance
- Confident – I use AI tools independently
- Very confident – I advise others or design AI-enabled workflows
- Expert – I build or adapt AI systems

AI Adaptability

When a new AI tool becomes available, how likely are you to try it?

- Extremely unlikely
- Somewhat unlikely
- Neither likely nor unlikely
- Somewhat likely
- Extremely likely

AI Trust

How much do you trust content that is generated by AI for work tasks?

- Do not trust it at all
- Trust it sometimes
- Neither trust nor distrust
- Trust it usually
- Trust it always

A. Survey Questions

AI Actual Use

AI Purposes

For which purposes do you use AI tools at work? Select all that apply

- ○ Writing or editing text
- ○ Data analysis
- ○ Coding or technical tasks
- ○ Creative tasks (design, images, music)
- ○ Decision support (strategy, innovation)
- Learning or tutoring
- I don't use AI-based tools in work (exclusive response)

AI Use types

Which of the following AI-based tools do you currently use as part of your work? Select all that apply.

- ○ I features built into other platforms (e.g. email auto-complete, smart replies, grammar correction, automated meeting summaries)
- ○ Generative AI tools (text, image, or code generation) (e.g. local AI system, ChatGPT, Microsoft Copilot, Google Gemini, Claude, DALL·E, Midjourney)
- ○ Predictive or recommendation systems (e.g. local AI system, sales forecasting tools, demand prediction, risk scoring, fraud detection, recommendation engines)
- ○ Speech or image recognition systems (e.g. local AI system, voice-to-text transcription, facial recognition, document scanning, automated image analysis)
- ○ Robotic process automation (RPA) or workflow automation tools (e.g. local AI system, automated invoice processing, workflow bots, UiPath, Power Automate)
- ○ I don't use AI-based tools in work (exclusive response)

AI exposure, perceived risks and opportunities

AI Statements

How much do you agree with the following statements:

Strongly disagree/Disagree/Neither agree nor disagree/Agree/Strongly agree

- I am concerned that AI could replace significant parts of my job in the future.
- Decisions made by AI at work (e.g. hiring, promotion, task allocation) are more fair than human decisions.
- AI could disadvantage certain groups of workers more than others.
- I feel prepared to adapt my skills to work alongside AI.
- I don't know enough about the AI tools available that could improve my work.
- I think my workplace is missing opportunities to use AI to improve how it operates.
- Using AI has changed the scope of my day-to-day work.
- My employer does not provide adequate support to upskill employees for AI-related changes.
- AI can improve productivity in my work.
- Colleagues who use AI tools are perceived as less competent at work.
- We do not currently have the data quality and organisation needed to fully benefit from AI.

AI Obstacles

Which of the following are the main obstacles to your potential use of AI in your daily work activities? Select all that apply.:

- ○ Lack of skills or knowledge to use AI tools effectively
- ○ Lack of trust in AI outputs (e.g., concerns about accuracy, bias, or reliability)
- ○ Lack of high-quality data or adequate technical infrastructure
- ○ Lack of clear organisational guidelines or policies
- ○ Data protection, privacy, or compliance concerns
- ○ Limited access to approved AI tools
- ○ Lack of time to learn or experiment with AI tools
- ○ Uncertainty about appropriate use cases for my role
- ○ Organisational culture or management resistance
- ○ Other. Please specify: ____________________
- ○ I do not require AI in my current role (exclusive response)

A. Survey Questions

AI exposure, perceived risks and opportunities

AI Dependency

To what extent do you feel dependent on AI tools to perform certain tasks in your role?

- ○ Not at all dependent
- ○ Minimally dependent
- ○ Slightly dependent
- ○ Moderately dependent
- ○ Extremely dependent

Skills Affected

Overall, how much has the use of AI affected your ability to complete work-related tasks on your own (without AI)?

- ○ My ability has significantly decreased
- ○ My ability has slightly decreased
- ○ No noticeable change
- ○ My ability has slightly improved
- ○ My ability has significantly improved
- ○ Too early to tell

AI skills and training

Training Past

Have you taken any steps to adapt your skills in relation to use of AI in your work? (multiple responses)

- ○ ▫Yes, informal learning (online tutorials, self-study)
- ○ ▫Yes, formal training or courses
- ○ ▫Yes, employer-provided training
- ○ ▫No (exclusive response)

Training Purpose

In which work activities would you most like training or support to use AI tools? Select all that apply.

- ○ ▫▫Writing or editing text
- ○ ▫Data analysis or data interpretation
- ○ ▫Improving data quality to enable effective AI use
- ○ ▫Coding or technical development
- ○ ▫Creative work (design, images, video, music)
- ○ ▫Using AI for decision-making or strategic support
- ○ ▫Automating routine tasks or workflows
- ○ ▫Evaluating AI outputs (accuracy, bias, reliability)
- ○ ▫Understanding risks, ethics, or responsible AI use
- ○ ▫Improving productivity using AI tools
- ○ ▫Other. Please specify: ______________________
- ○ I do not require further support with AI (exclusive response)

Training Tools

For which types of AI tools used in your daily work activities would you like more training or support? Select all that apply.

- ○ AI features built into existing software (e.g., email auto-complete, smart replies, grammar correction, automated meeting summaries)
- ○ Generative AI tools (text, image, or code generation) (e.g. local AI system, ChatGPT, Microsoft Copilot, Google Gemini, Claude, DALL·E, Midjourney)
- ○ Predictive or recommendation systems (e.g. local AI system, sales forecasting tools, demand prediction, risk scoring, fraud detection, recommendation engines)
- ○ Speech or image recognition systems (e.g. local AI system, voice-to-text transcription, facial recognition, document scanning, automated image analysis)
- ○ Robotic process automation (RPA) or workflow automation tools (e.g. local AI system, automated invoice processing, workflow bots, UiPath, Power Automate)
- ○ Other. Please specify: ______________________
- ○ I do not require further support with AI tools (exclusive response)

A. Survey Questions

Training Format

What type of AI-related training would you prefer? Select all that apply.

- ◦Self-paced online learning (e.g. video modules, tutorials)
- ◦In-person workshops
- ◦Live online workshops/webinars
- ◦Hands-on practical training using real work tasks
- ◦One-to-one coaching or mentoring
- ◦Written guides or toolkits
- ◦Peer learning sessions (learning from colleagues)
- ◦Other. Please specify: ______________________
- I do not feel I need AI-related training (exclusive response)

Attention Check

Attention Check

Please select "Weekly" from the options below.

- ◦Daily
- ◦Weekly
- ◦Monthly
- ◦Never

Demographics

Region

Where is your organisation or primary workplace based? (If you work remotely or from home, please select the location of your organisation's main office or base).

- ◦Dublin
- ◦Border (Cavan, Donegal, Leitrim, Monaghan, Sligo)
- ◦Mid-East (Kildare, Meath, Wicklow, Louth)
- ◦Mid-West (Clare, Tipperary, Limerick)
- ◦Midland (Laois, Longford, Offaly, Westmeath)
- ◦South-East (Carlow, Kilkenny, Wexford, Waterford)
- ◦South-West (Cork, Kerry)
- ◦West (Mayo, Roscommon, Galway)

Age

What is your age?

- ◦18 – 20 years
- ◦21 – 30 years
- ◦31 – 40 years
- ◦41 – 50 years
- ◦51 – 60 years
- ◦61 – 64 years
- ◦65 years and over
- ◦Prefer not to say

Gender

What is your gender?

- ◦Female
- ◦Male
- ◦Non-binary
- ◦Prefer not to say

A. Survey Questions

Demographics

Education

What is the highest level of education you have completed?

- ▫Lower secondary or less (e.g. Junior Certificate or below)
- ▫Upper secondary (e.g. Leaving Certificate or equivalent)
- ▫Post-secondary (non-tertiary) (e.g. PLC, apprenticeship, traineeship)
- ▫Bachelor's degree
- ▫Master's degree
- ▫Doctorate (PhD or equivalent)

End of Survey

Thank you very much for participating in this study. Your contribution is greatly appreciated and valuable to our research. If you have any questions about the study, or if you would like more information about the research findings once the project is completed, please contact us at SOHAM@tcd.ie.

Below you will find your unique, randomly assigned participant code. Please keep a record of this code. If you wish to withdraw your data at any time, you can contact the research team and provide this code. It allows us to locate and permanently delete your responses without collecting or using any personally identifiable information. Once deleted, your data cannot be recovered.

B. Experimental Vignettes

Experiment 1 - Conditions for AI usage

Please read the following hypothetical workplace scenario carefully. On the next screen, you will then be asked three questions, and you cannot return. Imagine you are working in a professional role at an organisation that uses digital tools to support everyday work.

[One of the following task types is randomly presented]

- You are asked to prepare an internal email summarising key points from a meeting and proposing next steps for your team.
- You are asked to prepare a product description for a client or customer that could influence an important decision or outcome.

[And one of the following AI situations (hierarchy of assistance) is randomly presented]

- An AI tool helps refine your text by correcting grammar, improving wording, and suggesting edits based on your notes.
- An AI assistant generates a recommended final version, and highlights key points.
- Your organisation does not use any AI system for this task and the policy is for the employees to use their own skills, however you know of an AI system that you can use to complete the task.

You are under some time pressure and have limited time to finalise your output, similar to a typical busy workday.

Reliance: How likely are you to use AI for this task?

- ▫Extremely unlikely
- ▫Somewhat unlikely
- ▫Neither likely nor unlikely
- ▫Somewhat likely
- ▫Extremely likely

Verification: How likely are you to check or verify the AI output in case you decide to use it?

- Extremely unlikely
- ▫Somewhat unlikely
- ▫Neither likely nor unlikely
- ▫Somewhat likely
- ▫Extremely likely
- ▫I would not use AI under any circumstances.

B. Experimental Vignettes

Experiment 1 - Conditions for AI usage

Accountability: In case you use the AI output and the outcome is wrong or causes harm, how would you assign responsibility? (Please drag and drop to rank, with 1 = most responsible.) Me

- My organisation
- The AI provider
- Shared responsibility
- Not sure

Experiment 2 - Evaluating Performance in the Presence of AI

Please read the following workplace scenario carefully. On the next screens, you will then be asked several questions, and you cannot return.
Imagine a colleague, Person A, who works in a role similar to yours.

[One of the following employee performances is randomly presented]

- Person A consistently delivers high-quality work, meets deadlines, and responds effectively to feedback.
- Person A often delivers work that requires revisions, misses deadlines, or struggles to respond to feedback.

[And one of the following learning orientation in AI use randomly presented]

- Person A uses AI tools to support learning and improvement—for example, asking for explanations, exploring alternative approaches, and using feedback to improve their own skills over time.
- Person A uses AI tools blindly to complete tasks and does not actively focus on understanding or improving the underlying skills involved.
- Person A does not use AI tools for this work. They complete tasks using their own judgement and skills, drawing on standard non-AI resources when needed (e.g., prior examples, templates, guidelines, documentation, or input from colleagues).
- Person A uses AI tools thoughtfully to complete tasks and carefully reviews the output before relying on it.

Competence: How competent do you consider Person A to be in their role? Extremely unlikely

- ◦Extremely incompetent
- ◦Somewhat incompetent
- ◦Neither competent nor incompetent
- ◦Somewhat competent
- ◦Extremely competent

Growth: How well does Person A's approach support long-term professional growth?

- ◦Not well at all
- ◦Slightly well
- ◦Moderately well
- ◦Very well
- ◦Extremely well

Appropriateness: How appropriate is Person A's use of AI tools in their situation?

- ◦Very inappropriate
- ◦Somewhat inappropriate
- ◦Neither appropriate nor inappropriate
- ◦Somewhat appropriate
- ◦Very appropriate

Delegate If this task were yours, how likely would you be to assign it to Person A instead of doing it yourself?

- ◦Very unwilling
- ◦Somewhat unwilling
- ◦Neither willing nor unwilling
- ◦Somewhat willing
- ◦Very willing

C. Survey Sample Overview

Figure 40. Which sector best describes your main occupation?

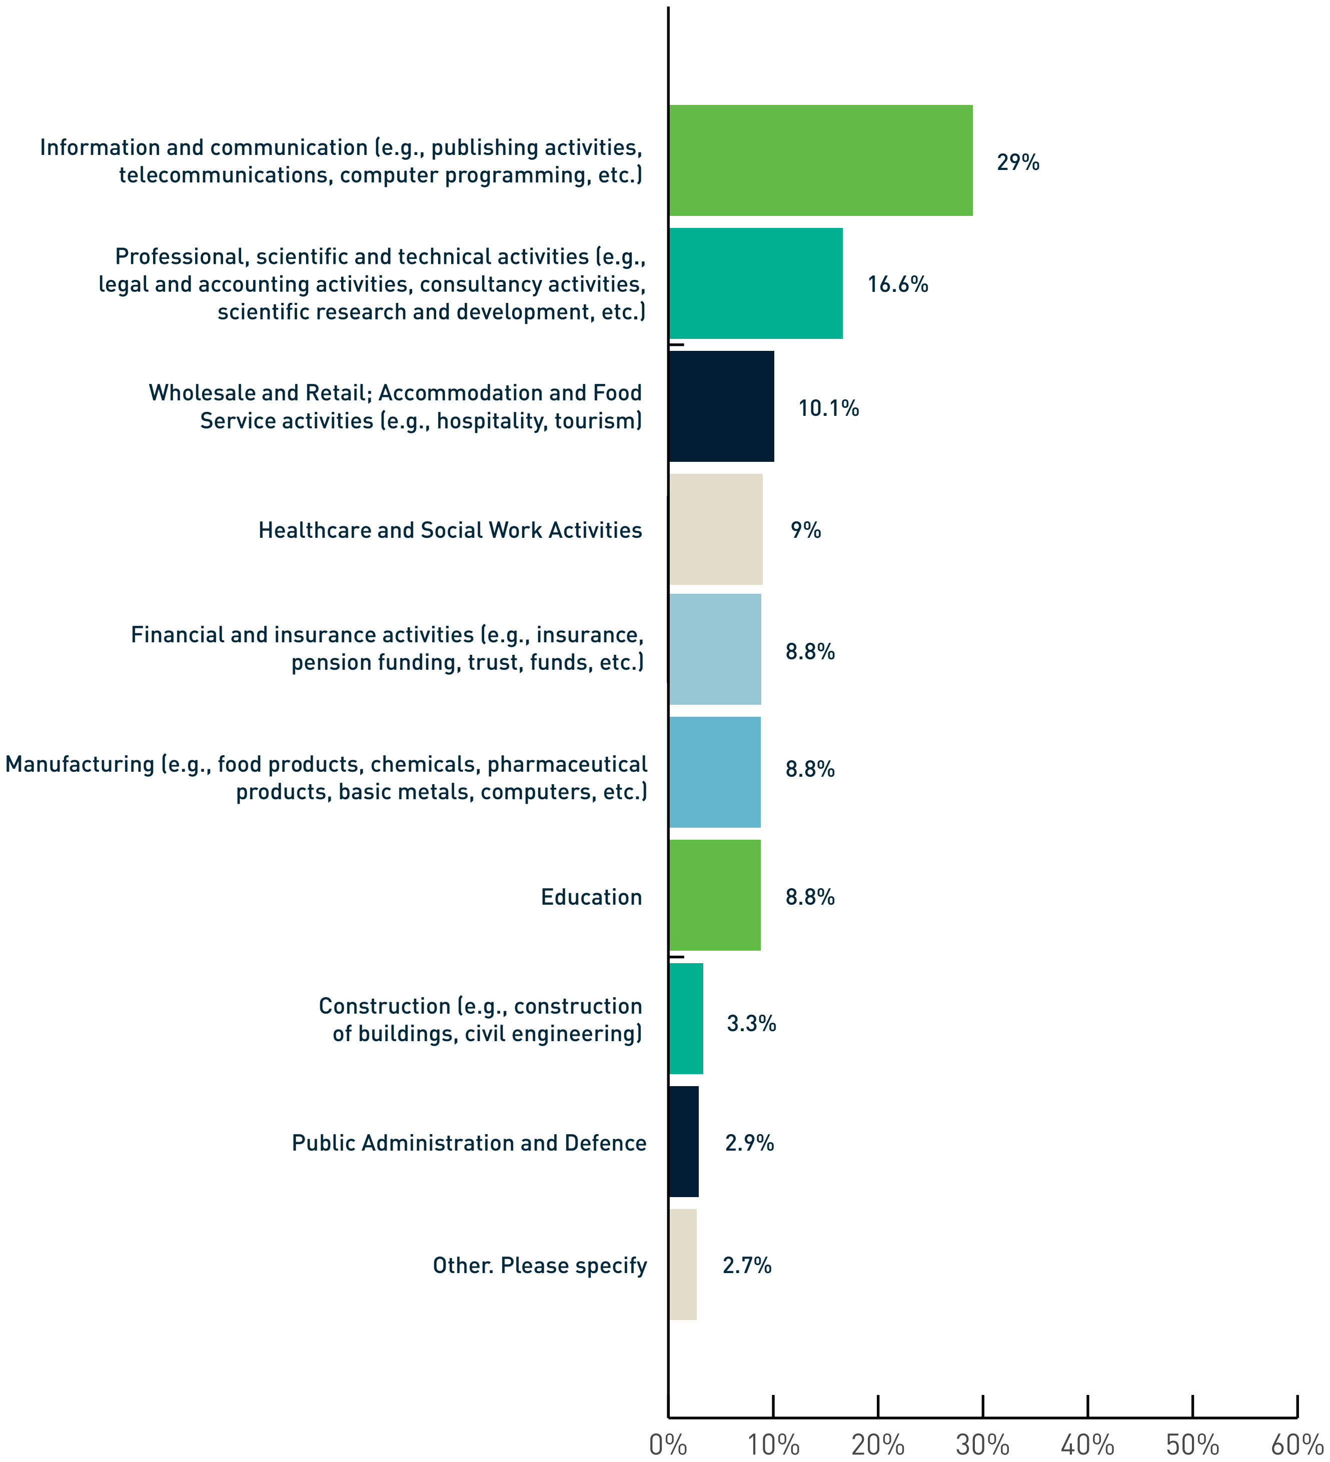

Figure 41. What is the size of the organisation where you work?

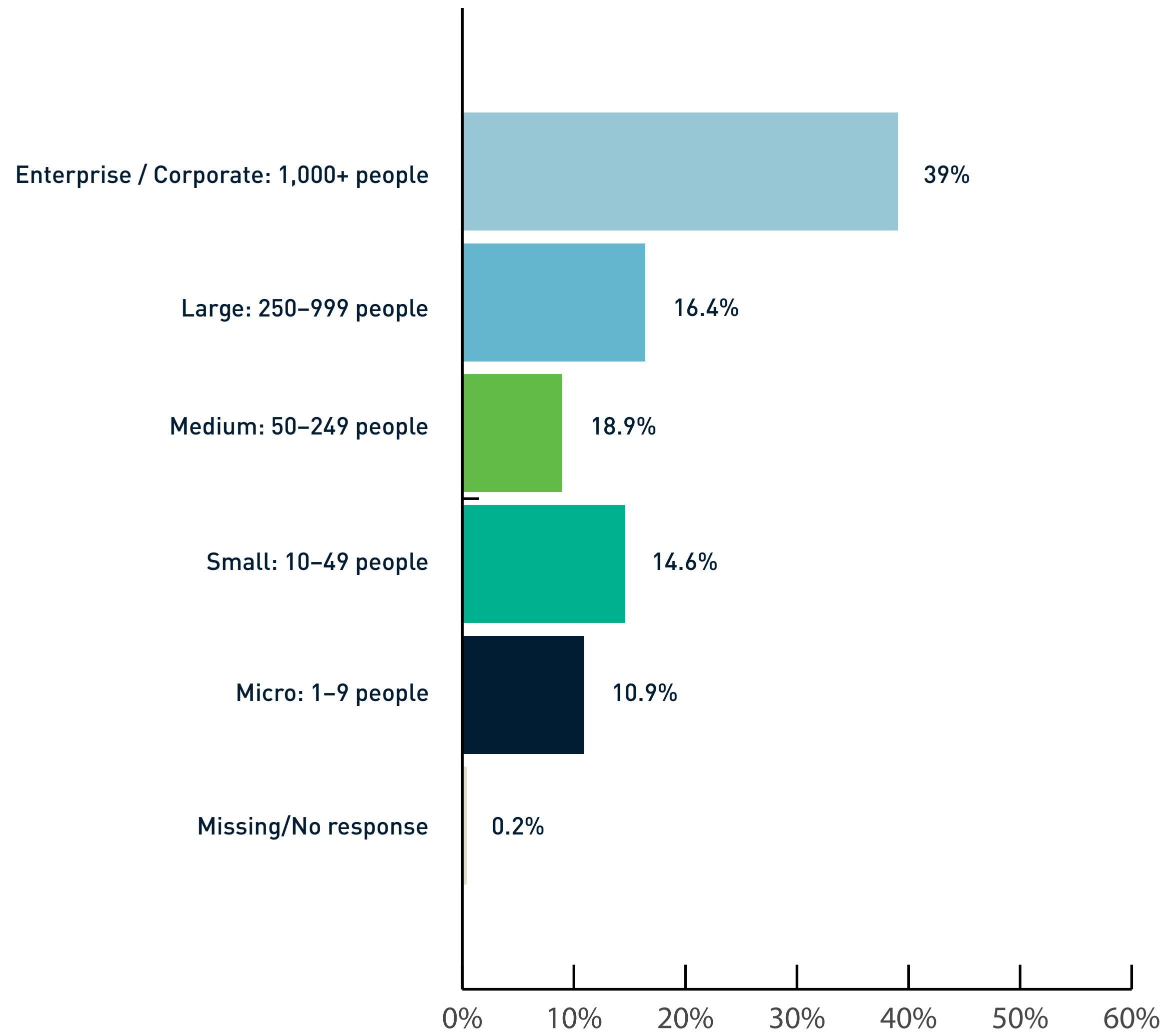

Figure 42. Which organisational function best describes your primary role?

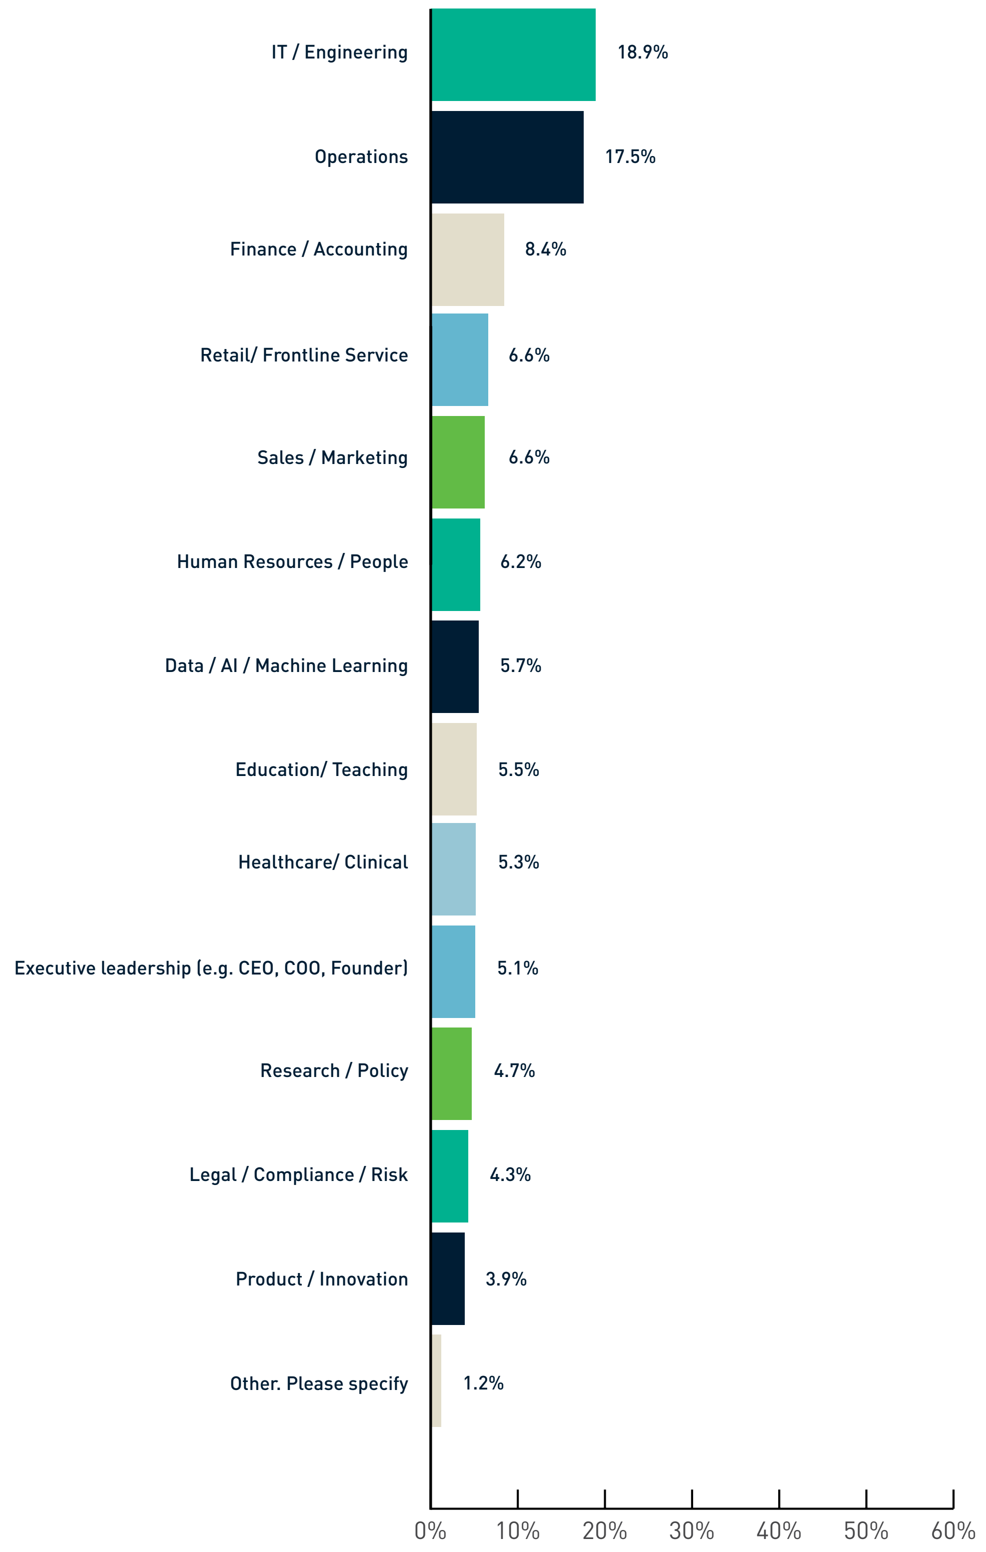

Figure 43. Which types of tasks make up most of your day-to-day work?

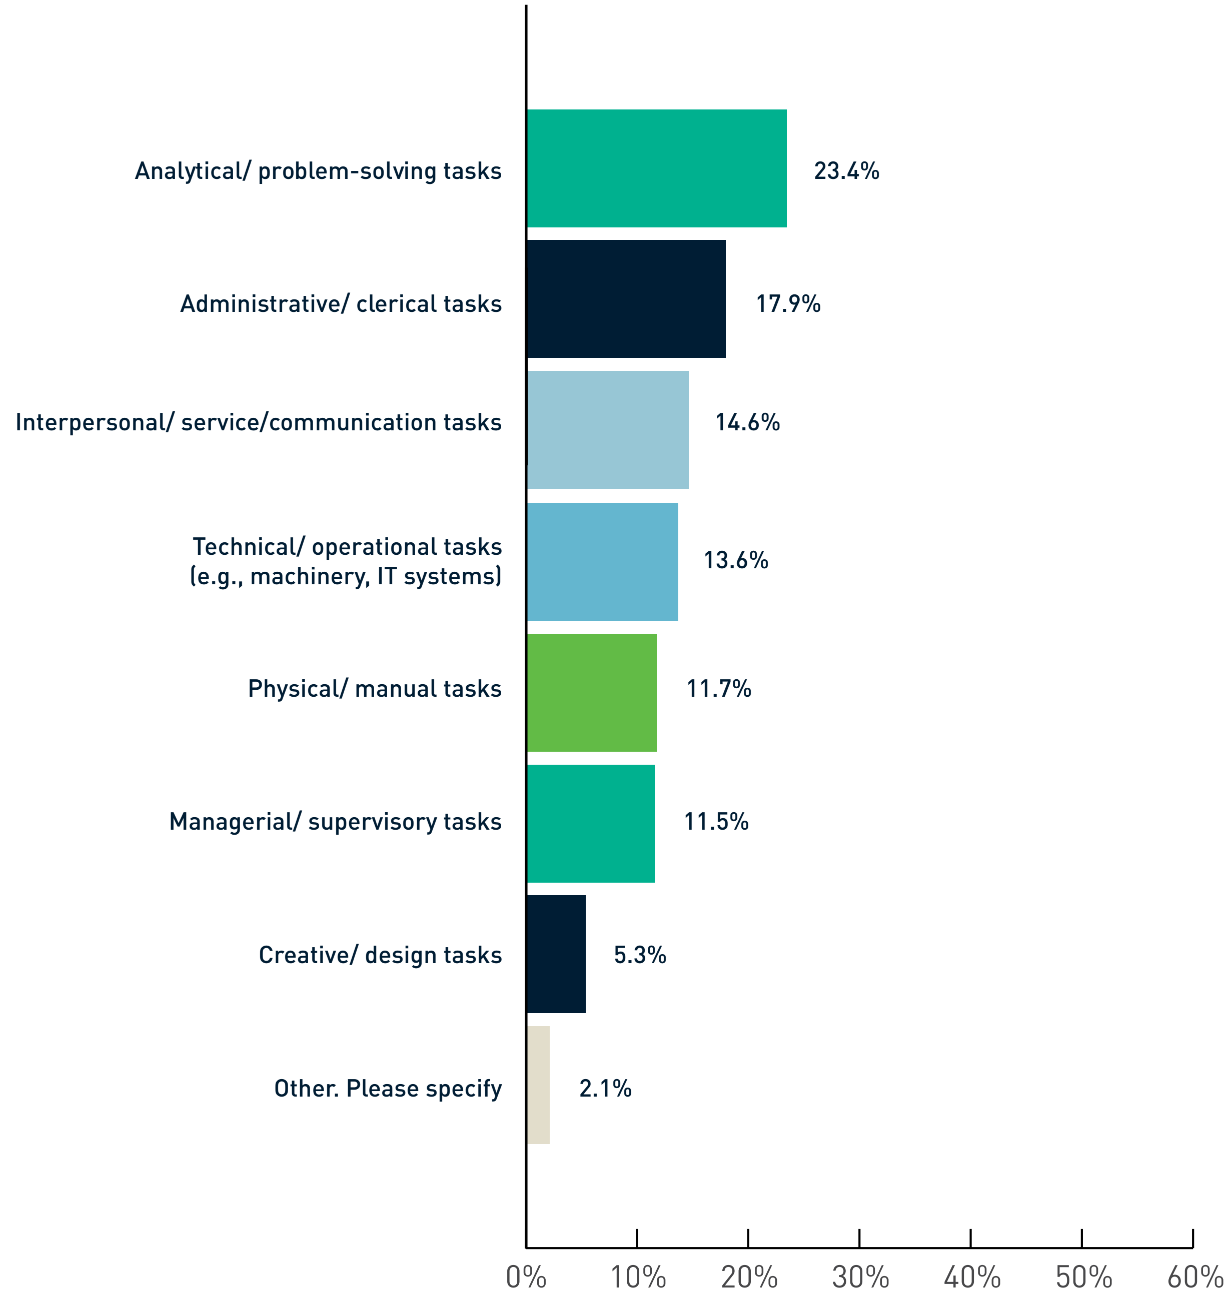

Figure 44. Where is your organisation or primary workplace based? (If you work remotely or from home, please select the location of your organisation's main office or base)

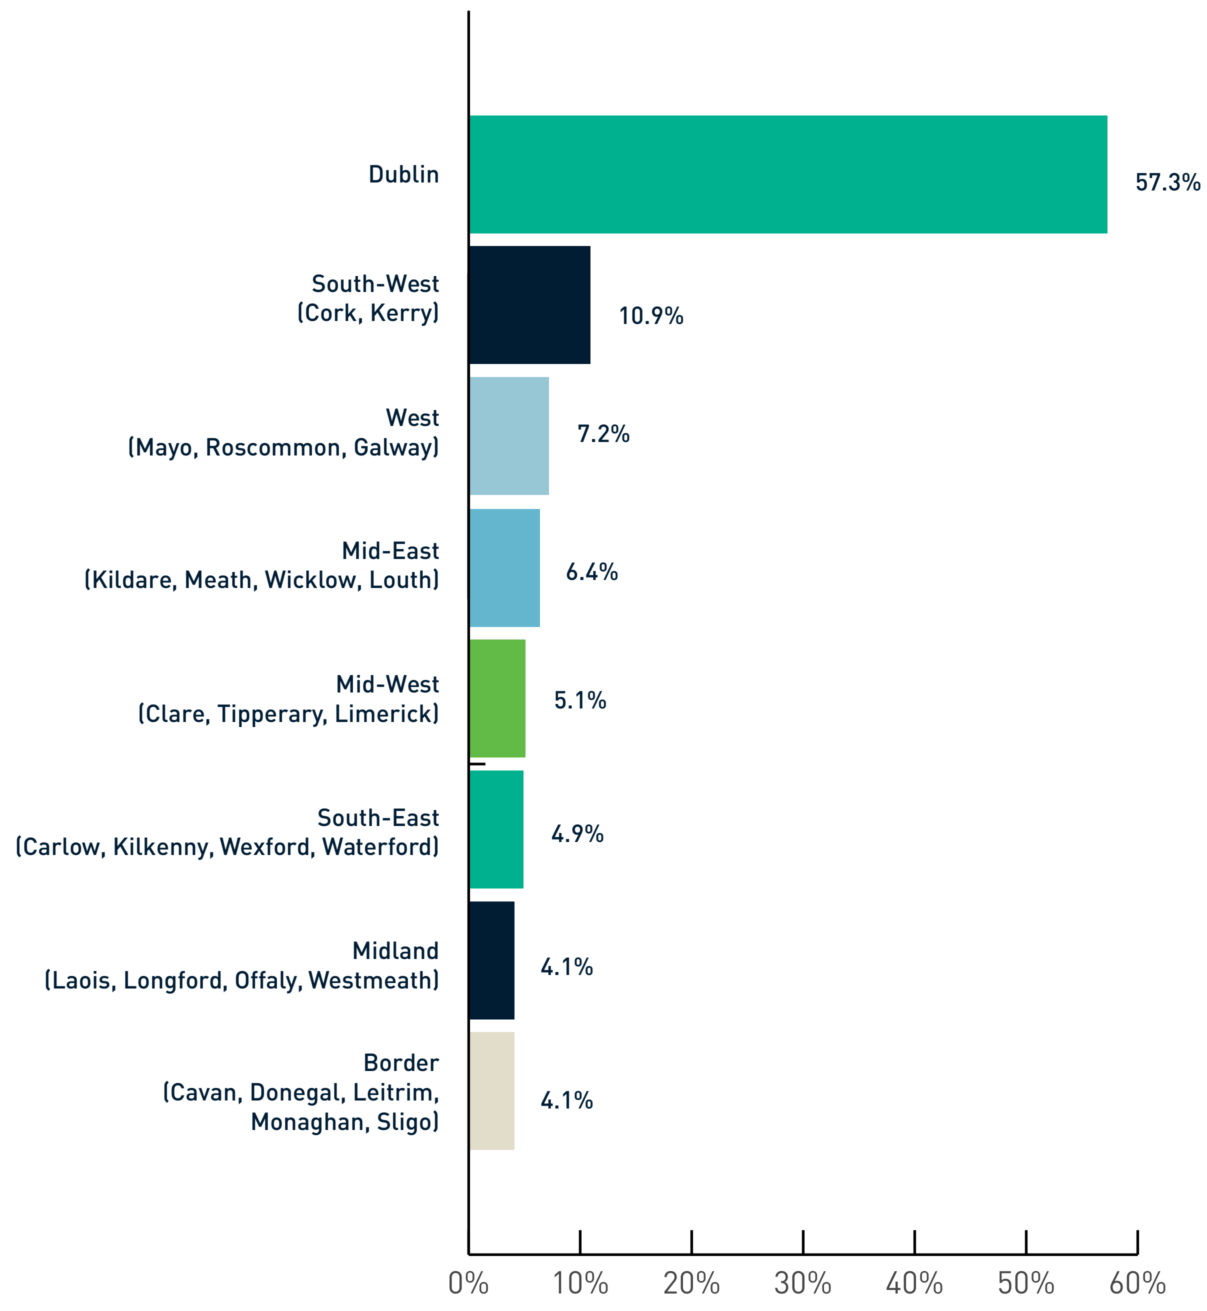

Figure 45. What is your age?

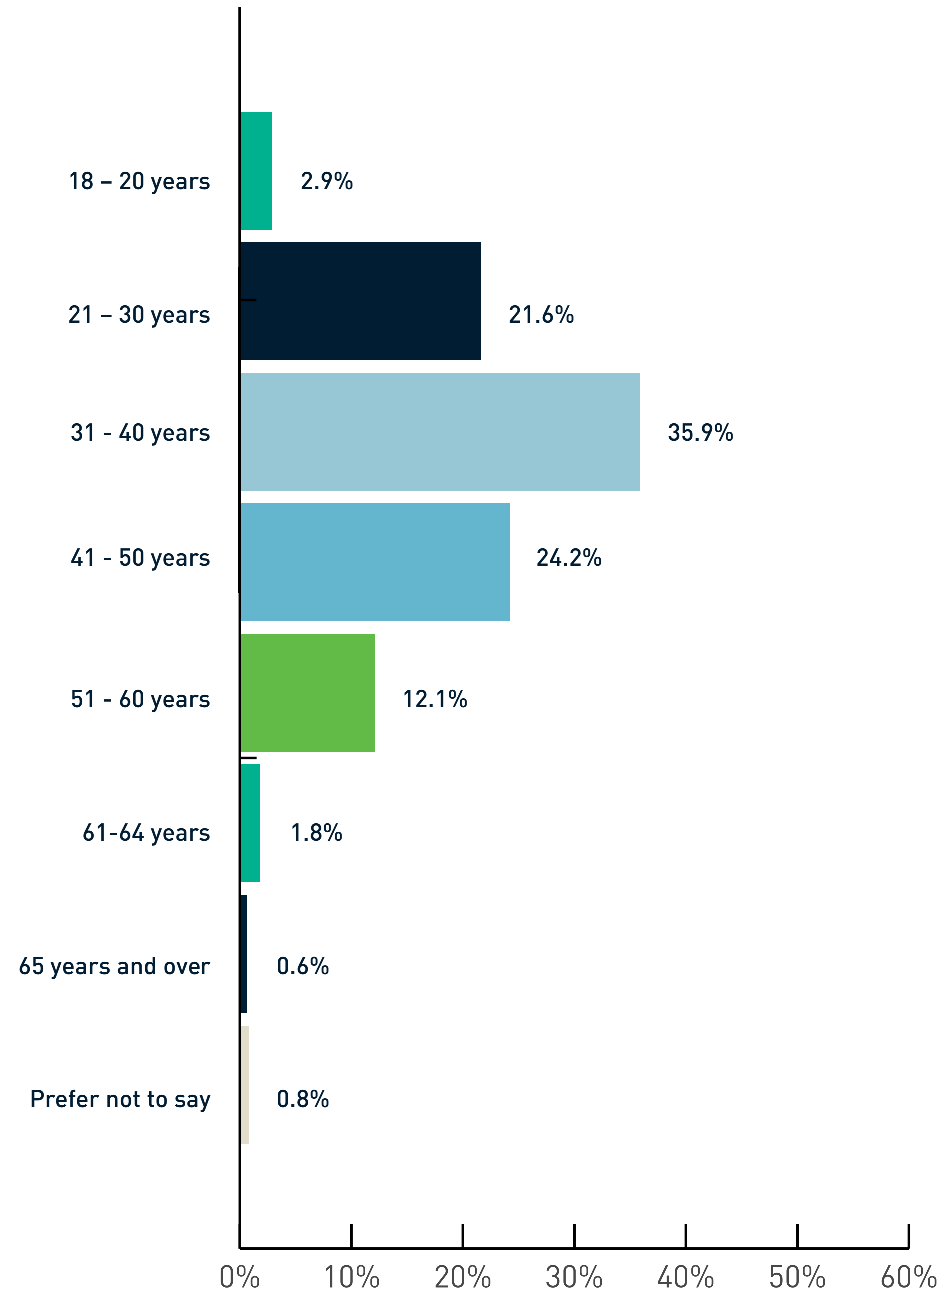

Figure 46. What is your gender?

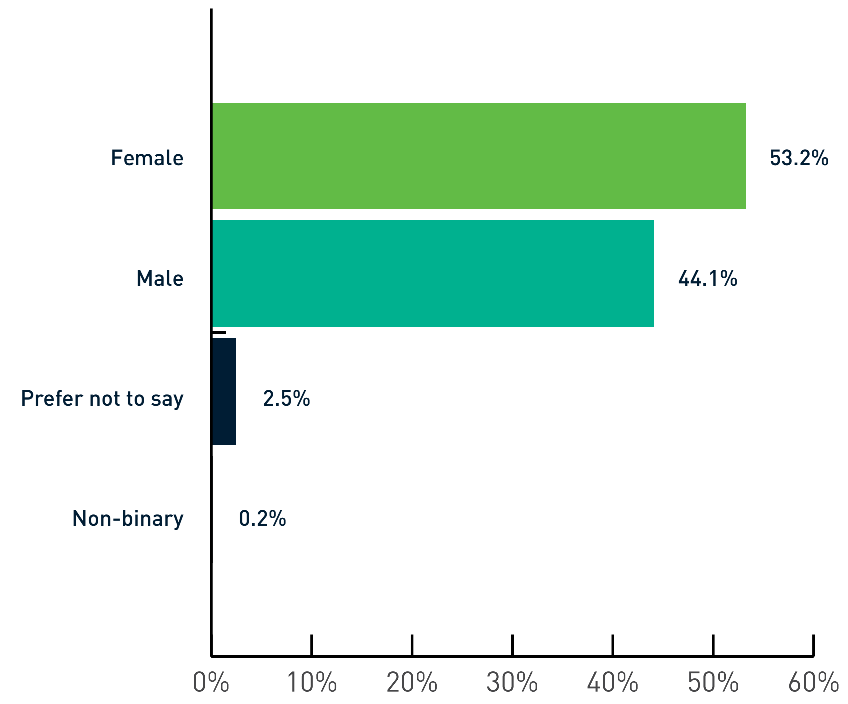

Figure 47. What is the highest level of education you have completed?

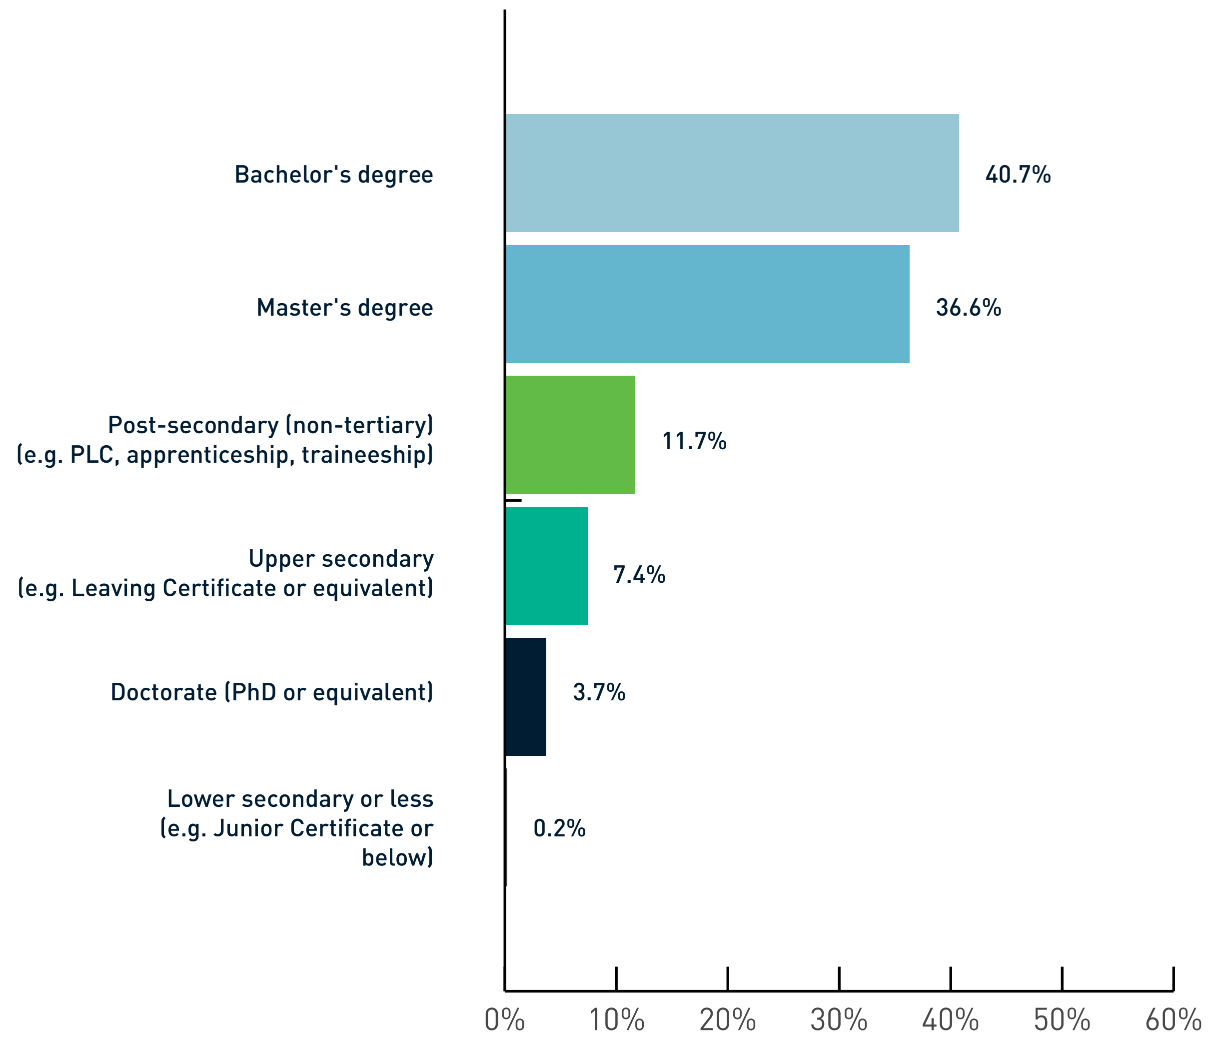

D. Demi Structural Interview Guide

Script to Be Read to Participant (Start of Interview)

Thank you again for agreeing to take part in this interview. As outlined in the information sheet, this interview will take about 30 minutes and focuses on organisational perspectives on AI adoption and its implications for skills and the workforce in Ireland.

With your permission, we're recording the interview for accurate transcription and analysis. The recording will be used only for research purposes, stored securely, and not shared outside the research team.

Before we begin, can I confirm that you are happy for the interview to be recorded? You are also free to skip any question or stop the interview at any time.

If you're happy to proceed, we'll begin.

Role & Organisational Context (3–5 minutes)

Purpose: situates the firm on the AI maturity spectrum.

Ice-break question: Could you briefly tell me about your role and how you came to be working in this area [AI/complement]?

Q1. Can you briefly describe how AI or automation features in your organisation's strategy or operations?
Probes:
• Are you using traditional ML, generative AI, agentive systems?
• Is adoption centralised or decentralised?
• Early-stage experimentation or scaled deployment?

AI Adoption & Business Strategy (5 minutes)

Purpose: extracts good practices.

Q2. How is AI being integrated into your core business functions?
Probes:
• Efficiency vs innovation vs cost reduction?
• Customer-facing vs internal productivity?

Q3. What has worked particularly well in your AI implementation so far?
Probes:
• Organisational structures?
• Leadership support?
• External partnerships (e.g., training bodies, vendors, consultant)?

Workforce Implications (5–6 minutes)

Purpose: identifies workforce bottlenecks.

Q4. How has AI adoption affected roles, tasks, or skill requirements in your organisation?
Probes:
• Within existing roles, what day-to-day tasks have changed because of AI?
• Job redesign vs job displacement?
• Creation of new roles?
• Resistance from staff?

Q5. Where have you experienced the biggest capability or skills gaps?
Probes:
• Technical skills?
• Data literacy?
• Structural problems (e.g. data quality, access to tools)
• Management capability to lead AI projects?

Skills Development & Reskilling (5–6 minutes)

Purpose: extracts replicable training practices and limitations of current infrastructure.

Q6. How is your organisation addressing AI-related skills development?
Probes:
• Internal training vs external providers?
• Collaboration with Irish universities or training bodies?
• Formal qualifications vs micro-credentials?
• Learning-by-doing?

D. Demi Structural Interview Guide

Follow-up: What has been effective — and what hasn't?

Barriers & Organisational Constraints (4–5 minutes)

Purpose: generates insights into structural barriers, Irish ecosystem constraints, Organisational readiness gaps.

Q7. What have been the main practical challenges in implementing AI at scale?

Probes:

• Ireland's education/ training, and your AI workforce needs?

• Talent shortages in Ireland?

• Cost?

• Cultural resistance?

• Data readiness?

•Regulatory uncertainty affecting hiring/investment?

Forward-Looking & Practical Advice (4–7 minutes)

Purpose: directly quotable recommendations.

Q8. Based on your experience, what practical advice would you give to Irish employers beginning to adopt AI?

Probes:

• Where should they start? Where is the low-hanging fruit? (easiest quick wins)

• Biggest mistake to avoid?

• Skills investments to prioritise?

• Organisational structures needed?

Q9. What advice would you give to organisations about transparency when introducing or using AI in the workplace?

Probes:

• How much should employees be told about how AI is used?

• Should companies disclose when decisions involve AI?

• What risks arise if organisations are not transparent?

Optional final prompt (if there is time):

Before we finish, could you share a specific example or case where AI or automation has already changed how work is done in your organisation? What happened and what did you learn from it?

E. Documents Included in the Policy Analysis

	DOCUMENT TITLE	YEAR	SOURCE
EU	Ethics Guidelines for Trustworthy AI	2019	High-Level Expert Group on Artificial Intelligence
EU	White Paper on Artificial Intelligence: A European approach to excellence and trust	2020	European Commission
EU	REGULATION (EU) 2021/694 of 29 April 2021 establishing the Digital Europe Programme	2021	European Parliament and Council
EU	EU AI Act – Regulation (EU) 2024/1689	2024	European Parliament and Council
EU	Draft Code of Practice on Transparency of AI Generated Content	2025	European Commission
EU	The General-Purpose AI (GPAI) code of practice	2025	European Commission
EU	General Data Protection Regulation 2016/679	2016	European Parliament and Council
EU	Communication on 2030 Digital Compass: the European way for the Digital Decade.	2021	European Commission
EU	Digital Services Act 2022/2065	2022	European Parliament and Council
EU	Proposal for a REGULATION laying down additional procedural rules relating to the enforcement of Regulation (EU) 2016/679	2023	European Commission
EU	The Data Act 2023/2854	2023	European Parliament and Council
EU	Artificial Intelligence in the European Commission (AI@EC) A strategic vision to foster the development and use of lawful, safe and trustworthy Artificial Intelligence systems in the European Commission	2024	European Commission
EU	Communication on boosting startups and innovation in trustworthy artificial intelligence	2024	European Commission
EU	Proposal for a REGULATION as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)	2024	European Commission
EU	Communication 'State of Digital Decade 2025: Keep building the EU's sovereignty and digital future'	2025	European Commission
EU	Communication on AI Continent Action Plan	2025	European Commission
EU	Communication on Apply AI Strategy	2025	European Commission
EU	Data Union Strategy – Unlocking Data for AI.	2025	European Commission

E. Documents Included in the Policy Analysis

	DOCUMENT TITLE	YEAR	SOURCE
EU	First Draft Code of Practice on Transparency of AI-Generated Content	2025	Working groups where Members States, industry, academia, and civil society collaborated.
EU	Guidelines on the scope of the obligations for general-purpose AI models established by AI Act	2025	European Commission
IE	Data Protection Act 2018	2018	Oireachtas
IE	AI – Here for Good: A National Artificial Intelligence Strategy for Ireland	2021	Department of Enterprise, Trade and Employment
IE	2022 Progress Report: Harnessing Digital: The Digital Ireland Framework	2022	Department of the Taoiseach
IE	Harnessing Digital: The Digital Ireland Framework	2022	Department of the Taoiseach
IE	Online Safety and Media Regulation Act 2022	2022	Oireachtas
IE	2023 Progress Report: Harnessing Digital: The Digital Ireland Framework	2023	Department of the Taoiseach
IE	2024 Progress Report: Harnessing Digital: The Digital Ireland Framework	2024	Department of the Taoiseach
IE	AI Advisory Council Advice to Government: FRT Use by An Garda Síochána	2024	AI Advisory Council
IE	Digital Services Act 2024	2024	Oireachtas
IE	Ireland's National AI Strategy: AI – Here for Good – Refresh 2024	2024	Department of Enterprise, Trade and Employment
IE	AI Advisory Council Advice to Government: AI and Education	2025	AI Advisory Council
IE	AI Advisory Council Advice to Government: The Impact of AI on Ireland's Creative Sector	2025	AI Advisory Council
IE	Ireland's AI Advisory Council Recommendations: Helping to Shape Ireland's AI Future	2025	AI Advisory Council
IE	Joint Committee on Artificial Intelligence – First Interim Report	2025	Houses of the Oireachtas
IE	S.I. No. 366/2025 - AI Designation Regulations	2025	Minister for Enterprise, Tourism and Employment,

	DOCUMENT TITLE	YEAR	SOURCE
IE	Skills Insights Note 2025-2: How AI is transforming the Irish Labour Market	2025	Expert Group on Future Skills Needs
IE	AI Advisory Council recommendations to Government regarding alleged creation and public dissemination of AI-generated non-consensual intimate images, including child sexual abuse material ("CSAM")	2026	AI Advisory Council
IE	General Scheme of the Regulation of Artificial Intelligence Bill 2026	2026	Department of Enterprise, Tourism and Employment
IE	National Digital & AI Strategy 2030	2026	Department of the Taoiseach
UK	National AI Strategy	2021	Secretary of State for Digital, Culture, Media and Sport
UK	A pro-innovation approach to AI regulation	2023	Department for Science, Innovation & Technology
UK	Online Safety Act 2023	2023	UK Parliament
UK	Digital Markets, Competition and Consumers Act	2024	UK Parliament
UK	AI Opportunities Action Plan	2025	Department for Science, Innovation & Technology
UK	Artificial Intelligence Playbook for the UK Government	2025	Government Digital Service/Department for Science, Innovation and Technology
UK	Data (Use and Access) Act	2025	UK Parliament
UK	Frontier AI Trends Report	2025	UK AI Security Institute
UK	State of Digital Government Review	2025	Department for Science, Innovation & Technology
UK	Crime and Policing Bill	2025	Government Bill

Technology Ireland Digital Skillnet

Technology Ireland
84/86 Lower Baggot Street
Dublin 2, Ireland D02H720

W: www.digitalskillnet.ie

Technology Ireland DIGITAL Skillnet is co-funded by Skillnet Ireland and network companies. Skillnet Ireland is funded from the National Training Fund and the European Union through the Department of Further and Higher Education, Research, Innovation and Science.

Rialtas na hÉireann
Government of Ireland

Co-funded by the European Union