International Labour Organization

ILO brief

Research Brief

February 2026

Workers’ exposure to AI: What indicators tell us – and what they don’t[1]

Key points

- AI exposure indicators estimate the extent to which AI systems can substitute for humans in specific tasks. Available exposure indices vary widely depending on the specific method used. Earlier computerization and automation measures suggested lower paid-workers in repetitive, routine manual or routine cognitive jobs to be more at risk, including some engineering-related occupations. In contrast, more recent AI capability–based indicators point to jobs with more “brain work” with higher exposure scores among cognitive, analytical, administrative and managerial occupations.
- Exposure patterns confirm substantial heterogeneity within occupational groups. Across different exposure measures, higher-skill and higher-wage occupations emerge as the most exposed. Occupations in business, finance, computing, mathematics, and education consistently show the highest exposure scores. Lower-skilled groups such as office and administrative support, and sales, also appear vulnerable, though with greater within-category variation.
- AI exposure extends beyond directly affected jobs via career paths and occupational transitions. Highly exposed jobs tend to occupy central positions in occupational networks—particularly in analytical, administrative, legal, financial and other professional fields. Because these jobs are closely connected to many others through shared skills—and career transitions, shocks affecting them can spill over to related roles, indirectly affecting workers whose own jobs do not appear directly automatable. By contrast, manual, care, and craft occupations lie on the periphery of the network and experience fewer spillovers.
- Limitations affect all exposure measures. They rely on static task lists of existing jobs, omit other adoption constraints, such as economic conditions and institutional barriers, embed subjective judgements (expert, worker, AI-based) and differ conceptually in how “exposure” is defined. As they lack any references to relative wages, economic feasibility and exposure might diverge significantly.
- Exposure indicators reveal technological susceptibility, not labour market outcomes. They capture only what AI could do—under a static view of tasks—not whether firms find it profitable to automate, how workflows will change, or how employment, wages, and demand will adjust. In particular, they do not account for productivity gains that may lower costs, expand demand and, historically, have contributed to net job growth despite task automation. Therefore, exposure measures offer risk assessments about potential job transformations but cannot be interpreted as predictions of job displacement, productivity gains or reskilling needs.

[1] The authors, Rossana Merola (ILO), Ekkehard Ernst (ILO), Daniel Samaan (ILO), Maria del Rio-Chanona (University College London), Ole Teutloff (Oxford University) thank Caroline Fredrickson and Sher Verick for constructive comments. We gratefully acknowledge Uma Rani and Morgan Williams for preparing Table A1 in the Annex.

▶ Introduction

As the development of generative artificial intelligence (GenAI) accelerates in capability and adoption, governments and social partners urgently need tools to anticipate which workers, occupations and sectors are most affected. To this effect, ex-ante exposure measures support policy debates by estimating which tasks or occupations could potentially be automated or transformed by AI. Specifically, they help forecast automation risks, priorities for reskilling investments and assess potential inequalities. Yet these measures vary widely in methodology and interpretation. Some rely on expert judgment, others on patent data or natural-language processing. The most recent approaches even involve using GenAI models themselves to evaluate tasks. Beyond methodological diversity, exposure measures remain subject to conceptual limitations. They measure what could be automated technically, not whether automation is profitable, generates productivity gains, nor how it impacts employment.

This brief draws on a review of the literature by del Rio-Chanona et al. (2025) and explains how exposure measures are constructed, presents a comparative analysis of leading indicators and clarifies what these measure reveal, what they obscure, and how their results can be interpreted and used in labour-market analysis and policy debates.

▶ How are AI exposure measures built?

Most exposure estimates begin with a mapping between tasks and occupations, recognizing that technologies substitute for or complement specific tasks rather than whole jobs. However, extensive information on different tasks carried out in various jobs is available and regularly updated only for a few countries, e.g., the United States, Italy, Germany and several others. Most studies collect task-level information from the Occupational Information Network (O*NET) database, a comprehensive resource developed by the U.S. Department of Labour. O*NET provides detailed information on the tasks, skills, abilities, work contexts, and knowledge required across occupations in the U.S. labour market.

O*NET is particularly valuable because it offers detailed task-level data and task importance weights, enabling researchers to aggregate task information into occupation-level metrics. Researchers aggregate task-level exposure scores (e.g., whether AI could complete a task) into occupational-level exposure scores using O*NET's task importance weights. These weights reflect how essential each task is within a given occupation—for example, "complex problem solving" for software developers, or "oral comprehension" for customer service representatives.

However, the transferability of O*NET to other countries is imperfect. Occupational structures differ, and tasks performed under the same job title can vary widely across countries due to differences in economic structure, production regimes and key labour market characteristics, including the degree of formalization and levels of education. An alternative approach, suitable for a selected number of developing countries, is to use skill surveys, such as the World Bank STEP survey. Studies attempting to globalise exposure measures either build crosswalks using semantic similarity (e.g., between O*NET and the World Bank STEP survey) or use regression-based extrapolations incorporating country characteristics (e.g. Carbonero et al., 2023). These approaches extend coverage but introduce new uncertainties due to cross-country heterogeneity and outdated task data in some surveys. Similarly, the OECD's PIAAC competency framework has been applied to estimate exposure rates for a large number of OECD, emerging and developing countries (Lewandowski et al., 2025).

Ultimately, all ex-ante exposure measures rest on an initial assumption: that current task descriptions remain meaningful in a future where GenAI may fundamentally alter how work is organised. Yet past waves of technological change show that the main impact is often a reconfiguration of work processes and workflows, rather than a one-to-one replacement of existing tasks (Poot and Samaan, 2024). This foundational limitation shapes all subsequent challenges.

Three approaches to estimating exposure

A wide range of methodological approaches has emerged to estimate workers' exposure to artificial intelligence, each reflecting different assumptions about how technology interacts with tasks and occupations. Some studies rely on expert or crowd-based assessments to judge which tasks can be automated; others use patent data to infer exposure from documented technological inventions. Recent studies (e.g., Gmyrek et al., 2025) combine multiple approaches, using external validation together with GenAI-based task assessments.

More recent approaches use natural-language processing on patents and GenAI models themselves to evaluate the automatability or augmentability of tasks. Each method captures different facets of technological potential—feasibility, innovation trends, or modelled capabilities—and therefore produces distinct exposure scores.

Table 1 summarizes strengths and limitations for each approach and reports the most relevant literature.

Expert judgement

The first major approach relies on experts. Frey and Osborne (2017), working in a pre-GenAI context, asked machine-learning experts to evaluate whether entire occupations—rather than their tasks—were automatable. Their judgments were then used to train a classifier to predict the automation risk of all occupations. This approach had two consequences. First, it conflated occupations with tasks: if the number of tasks automatable exceeded a specific (fixed) threshold, the entire occupation was often labelled "high risk." Second, experts at the time assumed that most creative, social, and non-routine cognitive tasks would remain out of reach for AI. In retrospect, these assumptions underestimated the capabilities of GenAI. Their limitations are visible when comparing older exposure scores to more recent ones, which now identify writing, analysis, and reasoning—tasks once thought non-automatable—as highly exposed.

Felten et al. (2018, 2021, 2023) combined expert assessments with occupational data to measure how exposed jobs are to AI. First, AI experts from the Electronic Frontier Foundation evaluated progress in specific AI application areas (e.g., image recognition, language modelling). Second, the authors linked these AI capabilities to occupational abilities in ONET using a crowd-sourced relatedness matrix from Amazon Mechanical Turk. They mapped AI capabilities to abilities such as oral comprehension, inductive reasoning, and arm–hand steadiness and then aggregated these using ONET importance weights to produce the AI Occupational Exposure (AIOE) measure, which captures the overlap between AI applications and job requirements.

This approach has limitations. Experts may be overly optimistic about AI progress, introducing bias. Focusing only on current capabilities may also understate future developments and their impact on work. In addition, the crowd-sourced mapping between AI capabilities and occupational abilities may oversimplify tasks, failing to reflect the full complexity of job requirements and work processes.

Alternatively, some studies (e.g., Gmyrek et al., 2025) ask ordinary workers—not experts—to judge what AI could plausibly do. Crowdsourcing studies like Brynjolfsson et al. (2018) assume that workers understand the practical realities of their own tasks better than external observers. Yet as the report argues, such methods introduce a different kind of subjectivity: workers often misunderstand what technologies can and cannot do, either out of fear or unfamiliarity. As a result, crowd-based measures tend to correlate with perceptions of automation risk rather than objective technological capabilities. Their strength lies in grounding exposure in real work processes, but their weakness is the potential distortion from worker expectations.

Patent-based text analysis

Patent-based approaches seek to overcome the subjectivity of expert judgment by examining documented technological innovations. Researchers identify patents associated with AI or automation and then assess the semantic similarity between patent descriptions and task descriptions in O*NET.

Compared to expert-based measures, patent-based methods focus on technologies that already exist and are potentially adoptable, and they avoid biases inherent in expert judgments.

The literature differs in how it selects patents and how it measures semantic similarity. Mann and Püttmann (2023) classify U.S. patents from 1976 to 2014 as automation-related by manually labelling hundreds of patents and training a machine-learning model to classify the rest, though their process is somewhat opaque and predates the rise of modern AI. Using European patent data, Dechezleprêtre et al. (2020) identify automation-related patents through keyword frequencies, while Gathmann and Grimm (2022) combine patent codes and keyword searches to classify patents from 1990 to 2018, offering more transparent approaches than Mann and Püttmann (2023). A central contribution is Webb (2020), who identifies AI-related patents through keywords such as "neural network" or "unsupervised learning," extracts verb–noun pairs from patent titles and matches them to verb–noun pairs in O*NET task descriptions; tasks and occupations showing greater overlap are considered more exposed. While transparent, this method may miss deeper semantic nuances. More recent studies leverage advanced natural language processing (NLP): Prytkova et al. (2024) use sentence-transformer embeddings to match AI-related patents to O*NET tasks, while Septiandri et al. (2024) employ Bidirectional Encoder Representations from Transformers (BERT) embeddings to perform a similar task. These deep-learning approaches capture richer semantic meaning but reduce interpretability, making it harder to understand how exposure scores are generated.

Still, patent data comes with other limitations. Patents are an imperfect proxy for technological readiness or adoption: many technologies never reach commercial scale, and many transformative applications of AI—especially proprietary or open-source models—are not patent-protected. Emerging NLP-based approaches (e.g., BERT, sentence transformers) improve the semantic matching between patents and occupational tasks, but at the cost of interpretability. These models capture subtle linguistic relationships, but policymakers gain little insight into why certain tasks appear exposed. Such black-box methods limit their usefulness for designing targeted interventions.

Patent-based measures are methodologically rigorous and informative about the overall direction of technological change, but they primarily track innovation, not adoption, feasibility rather than organisational change, and relatively narrow slices of technological development rather than the broader generative capabilities reshaping work today.

GenAI-based self-assessment

The newest wave of exposure studies asks GenAI models themselves—usually GPT-based—to evaluate which tasks can be performed by AI. These approaches have several advantages: they are simple to implement, draw on large corpora of knowledge, and can be rapidly updated as new models are released.

Eloundou et al. (2024) pioneered this method by asking ChatGPT whether it could perform given O*NET tasks "at least as well as a human" and in "half the time." Tasks meeting these criteria were labelled as high exposure. Gmyrek et al. (2023, 2025) extended this approach globally across ISCO-coded tasks, asking the model to justify each decision.[2] Kogan et al. (2023) distinguished between substitution (AI doing the task independently) and complementarity (AI assisting humans).

GenAI-based measures raise new challenges. Models may overestimate their own capabilities, mirroring the optimism embedded in their training data. Conversely, they may downplay tasks requiring embodied skills or tacit judgment. Because these models are not grounded in real-world production processes, their assessments reflect textual patterns rather than operational realities and human roles. Very few studies conduct robustness checks across different models. Among these, Chen et al. (2025) suggest the need for caution when interpreting the findings and highlight the importance of external validation to ensure accurate assessment of AI impact on the labour market.

Thus, although GenAI-based approaches make use of advanced technology and are efficient, they risk reinforcing simplified views of AI capabilities. This underscores the need to use them in conjunction with other data and external validation, rather than as a sole evidential basis for policymaking.

[2] Gmyrek et al. (2023) rely primarily on GPT-4–based AI scoring of task exposure with limited direct involvement of workers or expert panels, while Gmyrek et al. (2025) combine AI predictions with workers' survey data and expert review to refine occupational exposure estimates.

▶ Table 1: Comparison of different approaches to measure AI exposure

Method	Key Papers	Strengths	Limitations
Expert-based assessment	Frey & Osborne (2017); Brynjolfsson et al. (2018); Felten et al. (2018, 2021, 2023)	Provides qualitative insight into AI feasibility; transparent and easy to interpret; captures nuanced expert understanding of technological progress; adaptable to new expert updates.	Possible expert optimism bias; reliance on subjective human perceptions; may over- or underestimate future AI capabilities; crowd-sourced mappings risk oversimplifying tasks; occupation-level assessments (e.g. Frey & Osborne) ignore task heterogeneity.
Patent-based semantic similarity methods	Mann & Püttmann (2023); Dechezleprêtre et al. (2020); Gathmann & Grimm (2022); Webb (2020); Prytkova et al. (2024); Septiandri et al. (2024)	Based on actual technological innovation; avoids expert-judgment biases; large-scale data sources; NLP approaches capture semantic depth; transparent in methods like Webb (2020).	ML classifiers may be opaque; keyword methods miss semantic nuance; patent data may not reflect economic adoption; older datasets may miss modern AI; NLP embeddings reduce interpretability.
GenAI-based assessment	Eloundou et al. (2024); Gmyrek et al. (2023, 2025); Septiandri et al. (2024)	Scalable and low cost; directly reflects frontier AI capabilities; can assess both substitution and complementarity; easily updated as models evolve; flexible for task-level or occupation-level analysis.	LLMs may misinterpret tasks; results are sensitive to prompt design; methods lack transparency; may inherit model biases; usually grounded in U.S. O*NET data, limiting transferability.

Source: del Rio-Chanona et al. (2025)
Note: Several studies listed above employ mixed methodologies, combining expert judgement, worker surveys and AI-based task assessments. The classification reflects the primary methodological emphasis rather than exclusive use of a single approach.

▶ Quantitative analysis and comparison of AI exposure measures

Despite their conceptual and methodological diversity, existing exposure measures ultimately aim to quantify a common underlying phenomenon: the extent to which workers' tasks are technologically susceptible to advances in artificial intelligence. Yet because these measures rely on different data sources—expert judgement, worker surveys, patent text, occupational abilities, or GenAI self-assessments—the degree to which they produce consistent results remains an empirical question. The analysis in Del Rio Chanona et al. (2025) examines whether different methodologies—ranging from earlier automation-focused measures (Frey and Osborne, 2017; Brynjolfsson et al., 2018) to more recent AI-oriented ones (Webb, 2020; Felten et al., 2021; Eloundou et al., 2024; Septiandri et al., 2024; Gmyrek et al., 2023, 2025)—converge toward consistent assessments of occupational exposure.[3]

To enable meaningful comparison across indicators that rest on very different assumptions—not only in terms of methodology, but also in how they conceptualise AI capabilities and their potential impacts on work — all exposure scores are first normalised to a common scale: a composite "Mean Normalised Measure" by scaling each exposure index from 0 to 1 and averaging them.4 This allows for an examination of their relative rankings across occupations irrespective of their original units or distributions. Results are reported in Table 2 and show a sharp divergence between older and newer measures. Frey and Osborne (2017) and Brynjolfsson et al. (2018) exhibit negative correlations with wages, implying that lower-wage occupations are predicted to be most exposed. By contrast, exposure measures explicitly designed to capture AI capabilities—those of Webb (2020), Felten et al. (2021),

[3] For a comparison of AI measures, see also Nurski and Ruer (2024).

[4] This normalisation does not imply that the underlying measures are conceptually equivalent or interchangeable; rather, it provides a transparent way to compare relative occupational rankings across heterogeneous approaches.

and Eloundou et al. (2024)—show positive correlations with wages, between 0.27 and 0.54, with Felten et al. (2021) displaying the strongest positive relationship. The Mean Normalised Measure also correlates positively with wages, but more moderately. Overall, the comparison reveals that different methodological approaches lead to substantially different conclusions about which jobs are most exposed to AI, with recent AI-specific measures indicating that higher-wage, cognitively intensive occupations face greater exposure.

▶ **Table 2: Correlation between wages and exposure measures**

Measure	Correlation
Frey & Osborne (2017)	−0.5576***
Brynjolfsson et al. (2018)	−0.0741**
Webb (2020)	0.2704***
Felten et al. (2018)	0.5375***
Eloundou et al. (2024)	0.4243***
Septiandri et al. (2024)	0.1787***
Gmyrek et al. (2023)	0.2811***
Mean normalised measure	0.3043***

Notes: *** p<0.001; ** p<0.01; * p<0.05.
Source: del Rio-Chanona et al. (2025)

Direct exposure measures: Which occupations are most exposed to AI?

Due to their conceptual differences, the exposure measures we review identify different types of occupations as most affected by AI. Early approaches such as Frey and Osborne (2017), Brynjolfsson et al. (2018) and Webb (2020) point primarily to occupations requiring engineering skills or low-skill manual work as highly exposed. In contrast, capability-based measures—Felten et al. (2018, 2021) and Eloundou et al. (2024)—consistently identify cognitive, analytical, and mathematically intensive professions as most susceptible to AI-driven change. The Mean Normalised Measure combines these perspectives, indicating high exposure for both cognitively demanding and lower-skill occupations.

These patterns are reinforced when exposure measures are aggregated to the one-digit SOC level. Mean exposure and its variation within each of the 22 major occupational groups reveal substantial heterogeneity across methodologies. Frey and Osborne (2017) assign the highest exposure to "Office and Administrative Support" and "Installation, Maintenance, and Repair," with large standard deviations indicating strong within-category dispersion. Brynjolfsson et al.'s (2018) SML scores, by contrast, show relatively uniform exposure across broad categories, suggesting that variation occurs more within occupational groups than between them. Webb's (2020) patent-based measure highlights "Business and Financial Operations," "Computer and Mathematical Occupations," and "Architecture and Engineering" as most exposed, while still identifying elevated exposure among lower-skilled categories such as "Office and Administrative Support" and "Production." Measures by Felten et al. (2021) and Eloundou et al. (2024) similarly concentrate exposure in higher-skilled occupations—including business, finance, computing, mathematics, and education—though "Sales and Related" and "Office and Administrative Support" also show notable sensitivity. When these measures are combined, the Mean Normalised Measure smooths out methodological differences and consistently identifies "Business and Financial Operations" and "Computer and Mathematical Occupations" among the most exposed overall. For a more detailed overview of occupations at high and low risk of being automated according to the literature, see Annex Table 1.

Indirect exposure and occupational networks

Del Rio Chanona et al. (2025) further highlight that occupations connected to highly exposed ones through skill networks may face indirect exposure, as workers displaced from one field increase competition in adjacent occupations—an effect invisible to task-based exposure scores at the individual occupational level but potentially large for policy planning.

Figure 1 visualises the occupational network, showing how jobs cluster based on shared skills and transitions. High-exposure occupations form dense hubs—largely analytical, administrative, legal, financial, and professional roles—where strong interconnections mean that shocks in one occupation can quickly spill over into neighbouring ones. As a result, displacement or task restructuring among highly exposed professions such as accountants, paralegals, financial analysts, or technical writers can indirectly heighten pressure in adjacent roles like auditing, compliance, office administration, or project coordination.

▶ **Figure 1: US occupational networks**

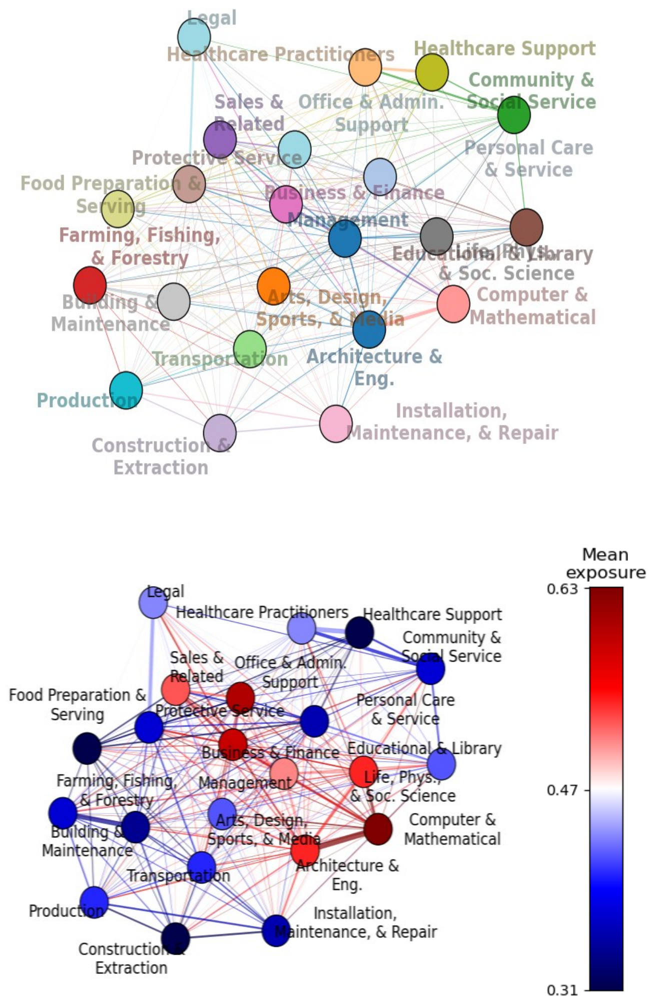

Notes: Network representation of occupation similarity based on the method developed by Mealy et al. (2018) of intermediate work activities. Nodes are occupations, while edges denote the overlap of work activities. The upper panel shows the network with labels, while the lower panel shows the network with nodes coloured by their exposure to automation.
Source: del Rio-Chanona et al. (2025)

By contrast, manual, care-oriented, and craft occupations appear on the periphery of the network, with fewer links to AI-intensive clusters. Their distance from high-exposure hubs limits spillovers, offering some insulation even when direct exposure is low. This asymmetry underscores that AI exposure reflects not only task automatability but also an occupation's position within the broader skill network. Figure 1 highlights how network proximity can broaden the circle of workers affected by generative AI, extending impacts beyond those occupations that appear most exposed in isolation.

▶ Common structural limits of all exposure measures

Despite methodological differences, all exposure measures share five structural limitations that place important limits on how they can be used for policy analysis. Importantly, these are not minor technical caveats—they fundamentally shape how exposure scores should be interpreted in labour market contexts.

The first limitation is the assumption of static tasks. Exposure measures rely on task lists that describe existing work organisation, yet transformative technologies often eliminate or reorganise existing tasks and create entirely new ones. A task labelled "highly automatable" today may disappear tomorrow, not because AI performs it but because the entire workflow is redesigned. This creates an asymmetric bias: exposure measures can overstate threat by evaluating the wrong tasks.[5]

Second, exposure measures provide a partial, task-based view of technological change and do not capture economy-wide productivity effects, which can offset automation by reducing costs, increasing demand and supporting employment growth.

More fundamentally, exposure measures provide only a partial view of the effects of technology on the economy. They focus on the technical susceptibility of tasks to AI, abstracting from the broader economic mechanisms through which technological change affects employment

[5] This could also work in the other direction: Tasks that are considered not automatable with a low AI exposure may nevertheless become obsolete through AI and disappear. For example, workers specialized in tasks such as adjusting and cleaning typebars, springs and levers of typewriters and replacing mechanical components, aligning keys, repairing ribbon mechanisms have lost their jobs. Not that the move to personal computers and word processors created a "typewriter-repair robot", but the need for typewriters altogether in most workplaces disappeared, and with it the specialised skill of typewriter mechanics.

and production. They do not capture productivity effects that can lead to additional jobs creation, which happens when higher efficiency due to automation translates into lower costs that expands output and increase labour demand. Within this partial framework, exposure measures also omit key economic variables. They do not incorporate wages, relative costs of adoption, complementarities, firm strategies, regulation, or productivity effects. A task may be automatable, but not economically viable to automate, for example, if the adoption costs are relatively high and wage costs comparably low. Svanberg et al. (2024) shows that cost-effectiveness of AI models will likely play an important role in the proliferation of the technology. Conversely, a task may appear to have a low exposure, but it may become automated if firms restructure workflows.

The third limitation concerns adoption dynamics. That AI can perform a task does not imply it will be adopted. In addition to economic factors, adoption depends on institutional, ethical, infrastructural, and organisational factors that differ across countries.

The fourth limitation is subjectivity. Whether experts, workers, or GenAI models are used, all exposure estimates embed human (or machine-learned) expectations and hence potential biases. Exposure measures should not be treated as objective forecasts but as informed assessments.

The fifth limitation is that most AI exposure measures explicitly or implicitly rely on U.S. skill and task data, which creates challenges when applying them to other countries. Task content within the same occupations can differ substantially across contexts—especially between developed and developing economies—so U.S.-based measures may not accurately reflect AI exposure elsewhere. As a result, there is a research gap on the impact of AI in developing countries, with Gmyrek et al. (2023, 2025) being a notable exception.

Finally, there is no shared interpretation of what exposure means. Some authors treat exposure as automation risk, others as augmentation potential, and others simply as potential for change. Policymakers frequently treat these as interchangeable, even though they reflect different underlying concepts and assumptions.

Conclusions and policy recommendations

AI Exposure measures offer a valuable but partial view of how AI could affect labour markets. They tell us where (Gen) AI is technically able to perform tasks, but they do not reveal whether adoption will occur, how quickly, or with what economic consequences. They should not be used to forecast job losses. They are best understood as early warning indicators that identify occupations where task content is likely to change. For this reason, the ILO should encourage governments to interpret exposure scores as signals of potential transformation rather than predictors of displacement, and to assess them alongside evidence on actual labour-market developments. In particular, exposure analysis should be combined with observed employment, wage and transition trends from labour force surveys and with data on AI adoption in workplaces, including firm- and worker-level surveys. When embedded in this broader empirical analysis—and supplemented with richer economic and institutional information—exposure measures can support proactive skills, employment and regulatory strategies.

The ILO is uniquely positioned to guide countries in the responsible use of exposure measures. By clarifying their limits, improving underlying data, and linking exposure to adoption pathways, the ILO can help ensure that technological transformation becomes an opportunity for inclusive growth rather than a source of new inequalities.

Annex

Annex Table 1: Top 5 high- and low-risk occupations

Key:
Low-risk
High-risk

	Expert insights			Patent data		AI-generated	Expert insights	
	Felton et al. (2021)	**Felton et al. (2023)**	**Felton et al. (2023)**	**Webb (2020)**	**Septiandri et al. (2024)**	**Eloundou et al. (2024)**	**Georgieff & Hyee (2021)**	**Lassebie & Quintini (2022)**
High-skilled (ISCO Skill levels 3 and 4)	Dancers	Dancers	Dancers	Animal caretakers	Insurance underwriters	Mathematicians	Business professionals	Relgious workers
	Fitness trainers	English language and literature teachers; post-secondary	Fitness trainers and aerobics instructors	Subject instructors, college	Cardiovascular technologists and technicians	Blockchain engineers	Managers	Top executives
	Genetic counsellors	Foreign language and literature teachers; post-secondary	Massage therapists	Art/entertainment performers	Sound engineering technicians	Court reporters and simultaneous captioners	Chief executives	Advertising, marketing, promotions, public relations and sales managers
	Financial examiners	History teachers; post-secondary	Interior designers	Clinical laboratory technicians	Nuclear medicine technologists		Sciences and engineering professionals	Office supervisors
	Actuaries	Law teachers; post-secondary	Architects	Chemical engineers	Air traffic controllers		Business administration, associate professionals	Operations specialities managers
	Purchasing agents		Chemical engineers	Optometrists	Magnetic resonance imaging technologists			
	Budget analysis		Art directors	Power plant operators				
			Astronomers					
Middle-skilled (ISCO Skill level 2)	Reinforcing iron and rebar workers	Reinforcing iron and rebar workers	Orderlies	Mail carriers for postal service	Pile driver operators	Proofreaders and copy markers	Agricultural, forestry, fishery labourers	Fishing and hunting workers
		Fallers		Dispatchers	Dredge operators	Correspondence clerk		Forest, conservation and logging workers
		Telemarketers			Graders and sorters			Rail transportation workers
								Wood workers
Low-skilled (ISCO Skill level 1)	Helpers - painters, paperhangers, plasterers, and stucco masons	Helpers - brick masons, block masons, stonemasons, and tile and marble setters	Dining room and cafeteria attendants and bartender helpers	Food preparation workers	Aircraft cargo handling supervisors		Cleaners, helpers	Extraction workers
	Pressers, textile, garment and related materials	Pressers, textile, garment and related materials					Food preparation assistants	
							Labourers	
							Refuse workers, other elementary workers	
Type of technology	**AI**	**GenAI, language modelling**	**GenAI, image generation**	**AI**	**AI**	**GenAI**	**AI**	**AI**
Country	**United States**	**United States**	**United States**	**United States**	**United States**	**United States**	**OECD (23)**	**OECD (37)**

Source: Prepared by Morgan Williams and Uma Rani

References

Acemoglu, D.; Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In: Ashenfelter, O.; Card, D. (eds), *Handbook of labour economics*, Vol. 4B. Amsterdam, Elsevier.

Autor, D.; Levy, F.; Murnane, R. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4): 1279–1333.

Brynjolfsson, E.; Mitchell, T.; Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108: 43–47.

Carbonero, F., Davies, J., Ernst, E., Fossen, F., Samaan, D. & Alina Sorgner: The impact of artificial intelligence on labor markets in developing countries: a new method with an illustration for Lao PDR and urban Viet Nam. *Journal of Evolutionary Economics* 33, 707–736 (2023). https://doi.org/10.1007/s00191-023-00809-7

Chen, Q.; et al. (2025). Large language models at work in China's labor market. *China Economic Review*, 92.

Dechezleprêtre, A.; et al. (2020). Automating labour: Evidence from firm-level patent data. Discussion Paper 1679, London, Centre for Economic Performance.

del Rio-Chanona, M.; Ernst, E.; Merola, R.; Samaan, D.; Teutloff, O. (2025). AI and jobs: A review of theory, estimates, and evidence. Available at: https://arxiv.org/abs/2509.15265.

Eloundou, T.; Manning, S.; Mishkin, P.; Rock, D. (2024). GPTs are GPTs: Large language models as general-purpose technologies. *Science*, 384 (6702).

Felten, E.W.; Raj, M.; Seamans, R. 2018. A method to link advances in Artificial Intelligence to occupational abilities. *AEA Papers and Proceedings,* 108: 54–57.

—. 2021. Occupational, industry and geographic exposure to Artificial Intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42 (12): 2195–2217.

—. 2023. Occupational heterogeneity in exposure to Generative AI. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4414065

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.

Gathmann, C.; Grimm, F. 2022. The diffusion of digital technologies and its consequences in the labour market, in VfS Annual Conference 2022 (Basel): Big Data in Economics (Verein für Socialpolitik / German Economic Association), pp. 1–35.

Gmyrek, P.; Berg, J.; Bescond, D. 2023. Generative AI and Jobs: A global analysis of potential effects on job quantity and quality. Working Paper 96, International Labour Organization.

Gmyrek, P. et al. 2025. Generative AI and jobs: A refined global index of occupational exposure. Working Paper 140, International Labour Organization.

Kogan, L. et al. 2023. Technology and labor displacement: Evidence from linking patents with worker-level data. NBER Working Paper 31846.

Lewandowski, P. et al. 2025. Worker's exposure to AI across development stages. IZA Discussion Paper, No. 18235.

Mann, K.; Püttmann, L. 2023. Benign effects of automation: New evidence from patent texts. *The Review of Economics and Statistics,* 105 (3): 562–579.

Nurski, L.; Ruer, N. (2024). Exposure to generative artificial intelligence in the European labour market. Bruegel Working Paper 06/2024.

Poot, M., & Samaan, D. (2024). Rethinking production processes with AI and avoiding the innovator's dilemma (Research Brief). International Labour Organization.

Prytkova, N.; et al. (2024). Mapping AI-relevant patents to occupational tasks using sentence transformers. Working paper 10955, CESifo.

Septiandri, A.A.; Constantinides, M.; Quercia, D. 2024. The potential impact of AI innovations on U.S. occupations. Available at: https://arxiv.org/abs/2312.04714.

Svanberg, M.; Wensu, L.; Fleming, M.; Goehring, B.; Thompson, N. (2024). Beyond AI exposure: Which tasks are cost-effective to automate with computer vision? Available at http://dx.doi.org/10.2139/ssrn.4700751.

Webb, M. (2020). The impact of artificial intelligence on the labour market. Stanford University Working Paper

Contact details
International Labour Organization
Route des Morillons 4
CH-1211 Geneva 22
Switzerland

T: +41 22 799 7239
E: RESEARCH@ilo.org

DOI : https://doi.org/10.54394/00033279