

Policy note

Interpreting AI Exposure Estimates: Implications for Low- and Middle-Income Countries

Gayani Hurulle and Rojje Ekanayake

February 2026

About LIRNEasia

LIRNEasia is a pro-poor, pro-market regional policy think tank. Its mission is catalysing policy change and solutions through research to improve the lives of people in the Asia and Pacific using knowledge, information and technology.

Address: 9A 1/1, Balcombe Place, Colombo 8, Sri Lanka

Telephone: +94 11 267 1160

Email: info@lirneasia.net

Website: <https://lirneasia.net/>

LinkedIn - <https://lk.linkedin.com/company/lirneasia>

About Futureworks Asia

FutureWORKS Asia is a multidisciplinary research network led by LIRNEasia focused on understanding and responding to the policy challenges posed by the changing landscape of work across Asia. The network currently consists of 11 think tanks, and brings together researchers and policy actors to generate evidence, strengthen regional research capacity, and engage in public policy processes related to skills, labour markets, and decent work in the future of work.

FutureWORKS Asia forms part of the broader FutureWORKS Collective, a global initiative funded by the International Development Research Centre (IDRC) of Canada. Regional FutureWORKS networks are being established in Asia, Africa, the Middle East and North Africa, and Latin America and the Caribbean, together forming the FutureWORKS Collective.

Funding

This work was carried out with the aid of a grant from the International Development Research Centre, Ottawa, Canada. The views expressed herein do not necessarily represent those of IDRC or its Board of Governors



Canada

Suggested Citation

Hurulle, G., & Ekanayake, R. (2026) *Interpreting AI Exposure Estimates: Implications for Low- and Middle-Income Countries*. LIRNEasia



1 Context

Concerns about the impact of technological change on jobs are not new. From early industrialisation and the Luddite movement to more recent debates around computerisation, periods of rapid technological development have often been accompanied by predictions of labour displacement, task reconfiguration, and productivity gains through automation and augmentation. Less than a decade ago, Frey and Osborne (2017) estimated that a substantial share of occupations in the United States was susceptible to computerisation, particularly in transportation and logistics, office and administrative support, and production roles. Similar estimates were made for other parts of the world, including in Asia. While such estimates were influential, realised labour-market outcomes were more uneven than predicted.

Similar discussions have re-emerged in recent years around artificial intelligence. The International Monetary Fund's Managing Director, Kristalina Georgieva, for example, has warned of an 'AI tsunami' (Wearden & Stewart, 2026). Occupations now considered potentially affected include writers, editors, graphic designers, and administrative support workers, with some overlap but also important differences from earlier automation debates. Concerns have intensified in response to the speed and scale of recent advancements. Unlike earlier industrial transformations that were geographically concentrated and diffused gradually, contemporary AI systems are digitally deployable and accessible across countries almost simultaneously. Generative AI tools in particular, became available worldwide at the point of release. This creates the perception of more immediate and widespread labour-market effects, even as the depth and direction of those effects continue to depend on local economic structures, infrastructure, regulatory environments, and the prevalence of informality.

Against this backdrop, a growing body of research seeks to measure workers' exposure to AI and estimate its potential labour-market impacts. Some studies analyse ex post data, examining observed changes in job postings, wages, employment, or task composition. Others are forward-looking, estimating potential exposure or risk, often producing macro-level projections for countries or regions. These approaches echo earlier efforts to quantify susceptibility during previous technological transitions.

This note reviews the methodology of forward-looking AI exposure literature, examining which occupational data are used, and how AI exposure is measured. It does so to discuss what these measures imply for labour-market analysis in Asian and other low and middle

income country contexts, where occupational structures, task composition, and data systems may differ from settings in their higher income counterparts.

Box 1: Which technology's impact are we measuring?

Artificial intelligence is commonly used as an umbrella term across the literature, but it is not a monolith. In practice refers to a range of different technologies, shaped by differences in applications, capabilities, and the evolution of models over time. Understanding how AI is defined and applied in each study is therefore important for making sense of its empirical findings.

Early empirical studies employ relatively broad definitions of AI, focusing on a range of capabilities rather than specific platforms. For instance, Webb (2020) defines AI use in patent texts by identifying key terms and concepts such as supervised learning, reinforcement learning, neural networks, and deep learning. Similarly, early work by Felten et al. (2018, 2021) focuses on broad AI abilities, including abstract strategy, real-time video game play, image recognition and generation, language modelling, and speech recognition.

Over time, however, the literature has increasingly shifted its focus toward large language models (LLMs) (Eloundou et al., 2023; Gmyrek et al., 2021; Felten et al., 2023; Gmyrek et al., 2025). This shift is consistent with the growing ubiquity of LLMs. The public release of ChatGPT by OpenAI in November 2022 is often treated as a turning point. For example, Kirby et al. (2025) employ difference-in-differences models to assess labour-market impacts associated with the public release of ChatGPT. Some recent studies explicitly address heterogeneity across LLMs by examining multiple models. Dominski and Lee (2025), for instance, compare the capabilities of OpenAI's GPT-4 and Anthropic's Claude 3.5 Sonnet, illustrating how findings may differ across models even within the same class of technology.

Separately, the rapid pace at which LLM capabilities evolve over time raises concerns about the temporal stability of empirical results. As models are frequently updated and improved, estimated effects may depend on the specific model version and period studied. Dominski and Lee (2025) explicitly engage with this issue by distinguishing model capabilities across time, identifying different stages of model development, including distinctions multimodal input and output capabilities in late 2023, and advanced reasoning capability in late 2024. They illustrate this by highlighting that a *customer* service representative's task exposure was about 50% in the early release LLMs around late 2022, but rose to about 80% two years later. Taken together, this highlights an important consideration when interpreting the literature: estimated effects may be tightly linked to particular models and time periods, and therefore relevant only for a limited duration given the rapid evolution of capabilities.

2 Data sources - occupational information

At the core of macro-level AI exposure predictions are occupation-based datasets describing factors such as tasks and abilities. Because AI exposure is typically inferred from the content of work rather than job titles alone, these datasets provide the building blocks for translating technological capabilities into occupational risk or complementarity measures. In many countries, however, detailed task-level occupational data are limited. As a result, researchers rely on large-scale occupational databases and international classification systems as the empirical foundation for exposure modelling.

O*NET is a popular source of information, due to the wealth of data. Much of the seminal literature is focused on the US job market context and specifically rely on the O*NET database developed by the National Centre for O*NET Development, funded by the United States Department of Labour. O*NET provides detailed 8-digit occupational profiles covering tasks, work activities, skills, knowledge, abilities, work context, and technology use, updated on a rolling basis. For example, at the 8-digit occupational level, 10–30 task statements are listed per occupation. This breadth and granularity have made it the dominant source for task-based empirical analysis. Furthermore, the database is updated regularly on a rolling basis using worker-reported information and expert review, contributing to its widespread use in task-based empirical analysis (Gregory et al, 2019). (Table 1).

Table 1: Sources of O*NET Occupational Information for AI Exposure Metrics in Select Studies

Paper	Type of occupation specific information used				
	Tasks	Abilities	Work Activities	Job zone	Work context
Felten et al (2018)		✓			
Felten et al (2021)		✓			
Brynjolfsson et al. (2018)	✓		✓		
Webb (2020)	✓				
Pizzinelli (2023)		✓		✓	✓

Increasingly, O*NET data is being used to make inferences beyond the US context. Kirby et al. (2025), for example, apply the AI Occupational Exposure (AIOE) measure developed by Felten et al. (2021), alongside the complementarity metric proposed by Pizzinelli et al. (2023), to South Asian countries in the World Bank’s South Asia Development Update,

drawing on ONET occupational and task descriptions. Similarly, Copestake et al. (2024) examine demand for AI-related skills

in the Indian labour market using online job vacancy data and employ Webb’s (2020) AI exposure measure, which also relies on ONET-derived occupational information. In both cases, O*NET data plays a central role in constructing exposure indices that are subsequently mapped onto national occupational classifications.

A small but growing and influential set of studies seeks to move beyond exclusive reliance on O*NET. For example, Gmyrek et al. (2023), published by the ILO, construct an exposure metric for 117 countries using the International Standard Classification of Occupations (ISCO-08). In a later study, Gmyrek et al. (2025) also draw on a detailed Polish occupational classification (Table 2). Compared to O*NET, ISCO-08 contains fewer occupational categories and provides far less task-level detail. However, it has the important advantage of explicit alignment with the occupational categories used by most national statistical offices, which facilitates cross-country application. The Polish classification includes a larger number of occupations and task descriptors than ONET, offering greater granularity. This is promising for exposure modelling. However, as shown in Table 3, some occupations, such as orthopaedic surgeons, are classified less precisely than in ONET. Gmyrek et al. (2025) argue that Poland represents an economy with characteristics of both higher- and lower-income countries, which may make its occupational structure more comparable to a broader range of contexts than that of the United States.

Table 2: Data sources for papers focusing on majority world contexts

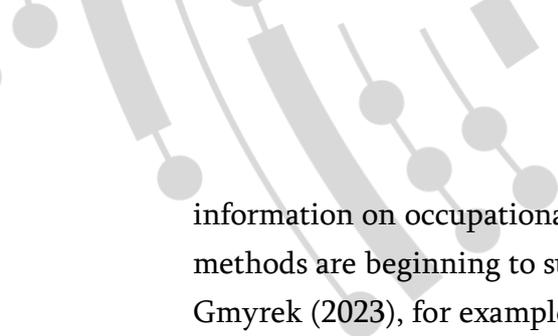
Data source	Examples	Country context
O*NET	Pizzinelli et al. (2023)	United States, United Kingdom, Brazil, Colombia, India, and South Africa
	Kirby et al. (2023)	South Asia (Bangladesh, Bhutan, India, Maldives, Nepal, Sri Lanka)
	Copestake et al. (2023)	India
International Standard Classification of Occupations (ISCO)	Gmyrek et al. (2023)	117 countries globally
ISCO, Polish data	Gmyrek et al. (2025)	n/a – likely similar to above

Table 3: Categories and tasks across data sources

Data source	No. of categories	Total tasks	Tasks per category	Example: Classification of Orthopaedic Surgeons	
				Classification	Implication
ISCO – 08	436	3,123	Range: 4-14	Specialist Medical Practitioners (2212)	Orthopaedic surgeons grouped with all specialist physicians
Polish data	~2,500	29,753	n/a	Doctor – specialist in general surgery (221209)	No distinct orthopaedic category; aggregated with surgical specialists
O*NET	~900	~19,000	Range: 10-30	Orthopedic Surgeons (29-1242.00)	Separate occupation with dedicated task profile

The trade-off involved in using datasets that do not fully reflect the lived realities of the majority world is one that researchers must weigh carefully. In the absence of nationally grounded task-level datasets, scholars often rely on sources such as O*NET to construct exposure measures and map them onto domestic labour market classifications. Yet this approach rests on a strong assumption of cross-country comparability in occupational task and skill content. To what extent do individuals performing the same nominal occupation engage in similar tasks across different economic and institutional settings? For example, is the task profile of a paediatrician in the United States comparable to that of a paediatrician in an Asian developing-country context, where resource availability, patient loads, and scopes of practice may differ substantially? Is task composition more uniform across contexts for certain occupations than others? It may be more defensible to assume comparability for software engineers than for farmers, whose work is deeply shaped by agroecological conditions, technology access, and local production systems. These distinctions matter for how confidently exposure metrics can be transferred across countries.

We must also be cognizant of new types of data sources that are emerging such as tasklists that are LLM generated and drawn from online job-portals. Traditional sources such as O*NET and ISCO rely on expert input and worker surveys to generate structured



information on occupational tasks, skills, and abilities. However, newer data sources and methods are beginning to supplement or, in some cases, challenge these approaches.

Gmyrek (2023), for example, uses a

large language model to generate task descriptions aligned with ISCO-08 occupations and compares these LLM-predicted task lists with official ISCO descriptors. This approach may offer potential advantages in generating context-specific insights, particularly if carefully prompted, but it remains highly dependent on the data quality of the model's training data. In parallel, a growing literature uses online job postings to identify occupational skills and tasks. These data can capture shifts in labour demand in near real time, yet they tend to over-represent formal-sector, white-collar, and digitally mediated occupations. The relative merits, biases, and limitations of these emerging data sources, especially in reflecting the lived realities of countries in the majority world, require systematic and comparative study.

3 Measuring AI exposure

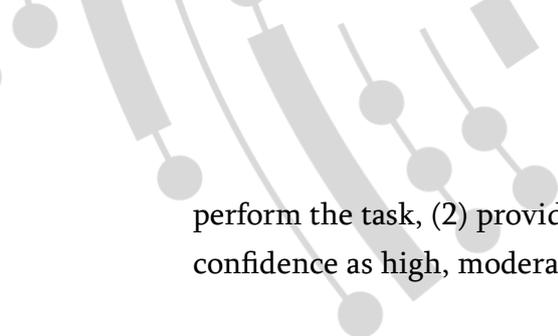
How AI exposure is assigned to different occupations is another important question, and one where variation is emerging in the literature. Much of the early work relied on human judgement to assess exposure. Brynjolfsson et al. (2018), for example, used a crowdsourcing platform to obtain human input on the extent to which occupational tasks and activities could be automated using machine learning. Similarly, Felten et al. (2018, 2021, 2023) linked AI abilities to occupations through structured human assessment, initially relying on computer science PhD students and later through surveys administered on a crowdsourcing platform.

Perhaps unsurprisingly, more recent work increasingly relies on large language models to estimate task-level exposure to AI, moving away from approaches based solely on human judgement (Table 4). Eloundou et al. (2023), for example, use structured prompting to elicit model assessments of whether specific occupational tasks can be performed by generative AI systems, supplementing these outputs with expert review. Many subsequent studies adopt similar hybrid approaches, combining LLM-based evaluations with human validation. In some cases, researchers explicitly differentiate exposure estimates across high- and low-income contexts, as seen in Gmyrek et al. (2025).

Table 4: Use of human and AI input in select literature

Author (year)	Expert/crowdsourced input used	LLM prompting used
Brynjolfsson et al. (2018)	Yes	No
Webb (2020)	No	No
Felten et al. (2018, 2021, 2023)	Yes	No
Eloundou et al. (2023)	Yes	Yes
Gmyrek et al. (2023)	Yes	Yes
Gmyrek et al. (2025)	Yes	Yes
Dominski and Lee (2025)	No	Yes

At the same time, a smaller but growing number of papers rely exclusively on model-based estimation, foregoing expert input altogether, as in Dominski and Lee (2025). Dominski and Lee (2025) compare the estimates produced by two LLMs, OpenAI's ChatGPT and Anthropic's Claude, assessing the percentage of occupational tasks that each model could perform. The authors implement this using a three-part prompt in which the model is asked to (1) reason step by step about the extent to which generative AI could



perform the task, (2) provide a numerical estimate between 0% and 100%, and (3) rate its confidence as high, moderate, or low.

LLM-based approaches may introduce new efficiencies in scaling and updating assessments, but also generate results that are sensitive to model choice and specification.

Even when identical prompts are used, different models can produce systematically different results, as shown in Dominski and Lee (2025). Comparing the top 40 occupations by change in AI exposure between the pre-November 2022 period and late 2023, as estimated by ChatGPT and Claude, the authors find that the magnitude of change is generally slightly higher under Claude. While writers and analysts appear near the top of both models' rankings in terms of change, differences emerge further down the list. For example, Interior Designers appear among the top 40 occupations by change in exposure under Claude but do not feature in ChatGPT's corresponding ranking. Although the broad pattern of substantial increases among language-intensive occupations is similar across models, variation in the size of estimated changes and occupational inclusion underscores the sensitivity of exposure metrics to model selection. Furthermore, variation at the AI exposure estimation stage cannot be separated from earlier decisions about how occupational tasks are defined, and how underlying biases or a lack of context specificity may bleed through.

4 Looking ahead

Forward-looking AI exposure metrics can serve as a useful starting point for planning, but their conclusions are highly contingent on how work is defined and measured. As this review has shown, these measures depend on the underlying occupational data used, the way tasks are defined and aggregated, which are often context-specific, and the methods through which exposure is assigned. Policymakers and practitioners should therefore approach these estimates with caution and remain mindful of the assumptions embedded within them when drawing conclusions or designing interventions.

Moving forward, building stronger local occupational datasets, including details on factors such as tasks and skills, will enhance the robustness and policy usefulness of AI exposure metrics. Ideally, this would be undertaken by national statistical organisations and integrated into regular data collection efforts. While such initiatives can be costly or time-consuming, smaller-scale sectoral studies could provide a pragmatic starting point. Comparative research in selected industries or occupations could document locally observed task composition and test how these profiles differ from O*NET-based descriptions, LLM-generated task lists, or job-posting-derived skill data. Such validation exercises would not only improve the interpretability of exposure metrics in specific contexts, but also help identify where international classifications meaningfully approximate local realities and where they do not.

Attention should turn to how AI exposure estimates relate to actual workplace adoption, labour-market adjustment, and the constraints that influence implementation. Even once local task composition is incorporated, technical exposure may diverge from realized adoption and labour-market impact. Where such gaps emerge, further analysis is needed to understand why. Factors such as contractual arrangements, risk aversion, regulatory uncertainty, capital constraints, infrastructure limitations, and organizational culture may all shape uptake. While a small but growing body of empirical work is emerging on observed productivity and adoption outcomes, this evidence appears to remain relatively limited, particularly in Asian and majority-world contexts. Continued efforts of this kind are especially important given the rapid pace of technological development, which may alter both task composition and adoption dynamics over time. Regular, context-sensitive reassessment would therefore increase relevance of AI exposure metrics as technology, and consequently, labour markets evolve.

5 References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), 293–340.

Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation, and work (NBER Working Paper No. 24196). National Bureau of Economic Research.

https://www.nber.org/system/files/working_papers/w24196/w24196.pdf

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

<https://doi.org/10.1257/pandp.20181019>

Copestake, A., Marczynek, M., Pople, A., & Stapleton, K. (2024). AI and services-led growth: Evidence from Indian job adverts. World Bank.

<https://thedocs.worldbank.org/en/doc/79e00908bd79275e6c22be34483dd91b-0050022023/original/akai.pdf>

Dominski, J., & Lee, Y. S. (2025). Advancing AI capabilities and evolving labor outcomes. arXiv. <https://arxiv.org/abs/2507.08244>

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308.

<https://doi.org/10.1126/science.adj0998>

Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108, 54–57.

<https://doi.org/10.1257/pandp.20181021>

Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3281>

Felten, E. W., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.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.

<https://doi.org/10.1016/j.techfore.2016.08.019>

Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality (ILO Working Paper No. 96). International Labour Organization. https://www.ilo.org/sites/default/files/2024-07/WP96_web.pdf

Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K., & Troszyński, M. (2025). Generative AI and jobs: A refined global index of occupational exposure (ILO Working Paper No. 140). International Labour Organization. https://www.ilo.org/sites/default/files/2025-05/WP140_web.pdf

Gregory, C., Lewis, P., Frugoli, P., & Nallin, A. (2019). Updating the ONET-SOC taxonomy: Incorporating the 2018 SOC structure—Summary and implementation*. National Center for O*NET Development. http://onetcenter.org/dl_files/Taxonomy2019_Summary.pdf

Kirby, P., Rexer, J., & Sharma, S. (2025). Artificial intelligence, real impact: Labor market implications of AI adoption in South Asia. In *Jobs, AI, and trade: South Asia development update* (October 2025). World Bank. <https://doi.org/10.1596/978-1-4648-2291-9>

Pizzinelli, C., Panton, A., Tavares, M. M., Cazzaniga, M., & Li, L. (2023). Labor market exposure to AI: Cross-country differences and distributional implications (IMF Working Paper). International Monetary Fund. <https://www.imf.org/en/publications/wp/issues/2023/10/04/labor-market-exposure-to-ai-cross-country-differences-and-distributional-implications-539656>

Wearden, G. & Stewart, H. (2026, January 23) *Young will suffer most when AI tsunami hits jobs, says head of IMF*. The Guardian. <https://www.theguardian.com/technology/2026/jan/23/ai-tsunami-labour-market-youth-employment-says-head-of-imf-davos>

Webb, M. (2020). The impact of artificial intelligence on the labor market. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.3482150>