Estimating Labor Market Exposure to Artificial Intelligence in Indonesia

Rachmat Reksa Samudra[ϕ, 1]
Ratu Silfa Addiba Nursahla[Ψ]
Izyan Pijar Bungabangsa Satyagraha[Ψ]
Dewita Nanda Prastiwi[ϕ]

[ϕ]Universitas Islam Internasional Indonesia, West Java, Indonesia
[Ψ]Institute for Economic and Social Research (LPEM), Faculty of Economics and Business, Universitas Indonesia, Jakarta, Indonesia

ABSTRACT

This study provides an estimation of labor market exposure to artificial intelligence in Indonesia by adopting a task-based approach used by Autor et al. (2003) and Acemoglu and Autor (2011). We employ six Generative Artificial Intelligence (GenAI) models available as per June 7th 2026 and three methods, namely engine-rule, large language model (LLM), and hybrid, to classify which task is exposed to AI, and, if exposed, whether the task is easy-to-learn by AI, using Indonesia Standard of Occupation Classification (KBJI) 2014. Using the median estimate result generated by Google Gemini 3.5 Flash with hybrid method, we find that 46.05% of KBJI 2014 tasks are exposed to AI. If aggregated into occupation level, defined as an occupation with at least one task exposed to AI, 81.3% of all occupations in KBJI 2014 are exposed, where white-collar occupations are more likely to be exposed to AI than blue-collar ones. Applying task-occupation classification into Indonesia Labor Force Survey (SAKERNAS) in 2022, this study finds that 68.1% of workers are employed in an AI-exposed occupation, and 43.9% in an occupation that is also easy for AI to learn. However, once digital access and the affordability of AI subscriptions are accounted for, these shares fall sharply, to 26.5% or lower for all exposed occupations and 21.8% or lower for easy-exposed occupations, and the realistically exposed workers are concentrated among urban, formal-sector, and more educated workers. This suggests that the near-term labor market impact of AI in Indonesia is likely to be both limited and unequally distributed.

Keywords: Artificial intelligence, labor market, task, occupation, Indonesia
JEL Codes: J24, O33, O53

[1] Corresponding author. E-mail: rachmat.reksa@uiii.ac.id. The view expressed in this paper are those of the authors and do not necessarily reflect the view of author's affiliation. This study is conducted without any financial/research support except from authors' own out of pocket. Further, there is no potential conflict of interests that the authors know in conducting this study. This study uses various generative artificial intelligence models as a tool for task classification. The concepting, planning, and execution stages of this study, including data analysis, are done by the authors. Samudra is in charge for methodology and data, Nursahla is in charge for literature review section, Satyagraha and Prastiwi are in charge for findings section. All authors are contributing equally for introduction and conclusion sections. All errors found in this paper remain authors'.

1. Introduction

An Indonesian proverb said, “every person has their time, every time has its person”, meaning that humans are constantly adapting with different issues they face each generation. Artificial Intelligence (AI) is the latest technology that the current generation needs to face and adapt. The rapid diffusion of AI, including Generative AI (GenAI), has reignited a longstanding debate in labor economics about whether technological change complements or displaces human workers, a question that the task-based framework pioneered by Autor et al. (2003) and extended by Acemoglu and Autor (2011) has proven uniquely well-suited to address by decomposing occupations into their constituent tasks rather than treating workers as homogeneous skill bundles. Unlike prior waves of computerization, which primarily threatened routine cognitive and manual tasks concentrated among middle-skill workers, large language models expose a structurally distinct set of tasks suggesting that the distributional consequences of AI may invert the polarization dynamic that defined the previous automation era (Brynjolfsson et al., 2018; Eloundou et al., 2023). Yet, whether this new exposure frontier reaches Indonesian workers equally, and through which occupations and task types it does so, remains an open empirical question. This study therefore examines how large AI exposure is across Indonesian occupations when measured through the task composition of the Indonesia Standard of Occupation Classification (*Klasifikasi Jabatan Baku Indonesia/*KBJI) 2014.

To do so, we apply the task-based framework to 433 occupations and 3,244 tasks in KBJI 2014, an adaptation of the ILO's International Standard Classification of Occupations 2008 (ISCO-08), excluding armed forces occupations from the analysis. We utilize six GenAI models available as per June 7th, 2026, which are Claude Fable 5.0 (which later banned for use outside U.S. and/or non-U.S. citizen on June 12th, 2026), Claude Opus 4.8, Claude Opus 4.7, Claude Sonnet 4.6, OpenAI GPT 5.4, and Google Gemini 3.5 Flash to classify which task is exposed to AI. Further, we also employ three methods of classification for comparison and robustness, namely engine-rule (derived from Acemoglu and Autor literature family), large language method (LLM) derived from Eloundou et al. (2023), and hybrid, i.e. combination of engine-rule and LLM methods. However, this study’s AI exposure definition is still agnostic, meaning that this study neither investigates whether the exposure will provide productivity gains for exposed workers nor displace the exposed workers due to data limitation. We find that all method-models have medium to strong positive correlation between each other, suggesting that AI-exposed task classification result is converging. To simplify our estimates, we use AI-exposed task classification based on Google Gemini 3.5 Flash with hybrid method as it provides the median estimates compared to all model-methods.

In the task-level, 46.05% of them are exposed to AI. If we narrow the definition of AI-exposed task to only easy-to-execute (as defined in Acemoglu, 2025), the share of tasks exposed to AI is

decreasing to only 15.84%. If the tasks are disaggregated into KBJI 2014 major occupation groups, this study finds that administrative workers, managers, professionals, and technicians and assistant professionals have the highest AI-exposed task (84.38%, 79.79%, 71.02%, 55.90%, respectively), or white-collar occupations. Meanwhile, tasks contained in blue-collar occupations, such as craft workers, operators and machine assemblers, and laborers have the lowest share of AI-exposed tasks (16.07%, 14.13%, and 6.54%). A similar pattern is also observed when we narrow down the definition to only consider easy tasks.

We then aggregate this task-level analysis into occupation level, where an occupation is AI-exposed if at least one task is exposed to AI. We follow this type of occupation classification that was once done for green jobs in Indonesia, pioneered by Granata and Posadas (2024) and adopted by Bappenas (2025) for green jobs skills development and policymaking purposes. This study finds that 81.30% of 433 occupations in KBJI 2014 are exposed to AI. If the definition narrows down to easy-to-execute tasks, then only 59.10% of occupations are exposed to AI. Disaggregated by major occupation groups, we also observe that white-collar occupations are more likely to be exposed to AI (as high as 100% for administrative workers) than blue-collar occupations (as low as 27.3% for laborers). This pattern also emerges if we only consider easy-to-learn AI-exposed tasks.

We then merge AI-exposed classification from KBJI 2014 into the Indonesia Labor Force Survey (*Survei Angkatan Kerja Nasional*/SAKERNAS) year 2022 in August to reveal how exposure is distributed across Indonesian workforce. We exclude armed forces, workers with missing wage data (such as unpaid/family workers), and workers with wage less than Rp 0 from the analysis. In total, there are 116,995,400 workers (weighted) analyzed in this paper to know characteristics of workers exposed to AI.

Weighting tasks by number of employment gives a more conservative picture than the raw classification, with more than 27.5% of all tasks performed in the economy exposed to AI and 13.6% both exposed and easy-to-learn. At the worker level, 68.1% of workers are employed in an occupation with at least one exposed task and 43.9% in an easy-exposed occupation, and exposed work is geographically concentrated on Java, which alone accounts for 56% of all exposed workers. Exposure is consistently uneven across workers. It is higher for women than men (32.4% versus 25.0% of tasks exposed), for urban than rural workers (32.5% versus 21.5%), and for formal than informal workers (35.3% versus 21.5%), and it rises steeply with education, from 17.4% among workers with primary schooling to 55.9% among those with tertiary education. White-collar occupations, led by administration, managerial, professional, and technical workers, are far more exposed than blue-collar workers in agriculture, crafts, and laborers, the same gradient found in the task classification.

Finally, because holding an AI-exposed occupation does not mean a worker can actually use AI, we assess realistic exposure by imposing two additional constraints, namely digital access (use of the internet and a digital device) and affordability (a monthly AI subscription costing no more than 20% of monthly labor income, priced at USD 20, USD 100, and USD 200 at current pricing tier based on token usage). While 68.1% of workers hold an AI-exposed occupation, the realistically exposed share falls to 32.9% once digital access is required, and then to 26.5% under the USD 20 subscription, 2.7% under USD 100, and 0.8% under USD 200. Affordability, rather than capability, is the binding constraint, and the realistically exposed workers are disproportionately urban, formal-sector, and younger. The same pattern holds at uniformly lower levels when exposure is restricted to easy-to-learn tasks. These findings suggest that the near-term labor market disruption from AI in Indonesia is likely to be far more limited, and far more concentrated, than the headline exposure figures imply.

This paper is organized as follows. Section two discusses past studies. Section three discusses methodology to classify AI-exposed task, occupation, and worker. Section four describes data that is used for this study. Section five reveals the findings. Section six concludes this study.

2. Literature Review

2.1 The Task-Based Framework

The early labor economics literature treated technological change as an exogenous force that mechanically shifted the demand curve for skills, conceptualizing the labor market through a canonical two-factor model in which technology takes a factor-augmenting form that raises the productivity of one group relative to the other (Katz and Murphy, 1992; Acemoglu, 2002). This view was addressed by Acemoglu (1998), who demonstrated that the skill-complementary or labor-displacing character of technology is not predetermined but endogenously shaped by profit incentives, whereby a larger relative supply of skilled workers expands the market for skill-complementary innovations and induces firms to develop machines that reinforce rather than displace them. Yet the more foundational methodological problem lay in the canonical model's conflation of workers' skills with the tasks they perform. Autor et al. (2003) established that it is tasks, not occupations or skill categories, that are the proximate objects of technological substitution, showing that computer adoption is associated with reduced labor input in routine cognitive and routine manual tasks and increased labor input in non-routine cognitive tasks. This routinization hypothesis has since been confirmed across multiple countries and empirical contexts (Goos and Manning, 2007; Autor and Dorn, 2013; Goos et.al., 2014).

Building on this foundation, Acemoglu and Autor (2011) mounted a formal critique of the canonical two-factor model, proposing a Ricardian task-based framework in which no skill group

holds absolute advantage across all tasks and technologies can directly substitute for workers in specific task ranges, generating the observed pattern of employment and wage polarization that a simple skilled/unskilled dichotomy cannot produce. A subsequent strand of research translated this theoretical apparatus into empirical predictions about automation risk at the occupation level (Frey and Osborne, 2017), though this approach was criticized for treating occupations as monolithic units rather than bundles of heterogeneous tasks (Arntz et. al., 2016; Brynjolfsson et. al., 2018). When Arntz et al. (2016) disaggregated to the task level and accounted for within-occupation variation, the estimated share of employment at high automation risk fell from approximately 47% to around 9% across OECD countries, underscoring that the choice of unit of analysis carries profound implications for both measurement and policy.

2.2 Beyond Exposure: The Wage and Welfare Consequences of Automation

The welfare consequences of automation depend critically on the dynamic interplay between task displacement and task reinstatement. Acemoglu and Restrepo (2018a) show that in a fully dynamic general equilibrium model with endogenous capital accumulation and directed innovation, a stable balanced growth path exists when automation and new task creation advance in tandem, with productivity gains from automation transmitted to wages over time. When this balance breaks down and automation persistently outpaces task creation, however, displaced workers crowd into a narrowing range of remaining tasks, depressing task-specific wages and generating rising inequality even alongside productivity growth (Acemoglu and Restrepo, 2018b, 2019). Empirical evidence confirms these mechanisms: Acemoglu and Restrepo (2020) find using a Bartik-style instrumental variable strategy that each additional robot per thousand workers reduces local employment-to-population ratios by 0.39 percentage points and wages by 0.77%, with effects concentrated in routine manual and blue-collar occupations. Extending this analysis to the aggregate economy, Acemoglu and Restrepo (2022) document that task displacement explains between 16% and 38% of changes in employment and wages across demographic groups between 1980 and 2016, with automation-driven labor share declines accounting for the bulk of rising wage inequality over this period.

These distributional concerns, historically concentrated among middle-skill routine workers (Autor et al., 2003; Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014), take a qualitatively different form with the advent of machine learning. Unlike rule-based automation, ML systems learn statistical patterns from data and generalize to new inputs, enabling them to penetrate tasks previously considered non-automatable under the standard task-based framework (Brynjolfsson and Mitchell, 2017). Brynjolfsson et. al. (2018) formalize this distinction, demonstrating that ML exposure is distributed across the entire wage spectrum rather than concentrated among middle-skill workers, threatening both high-wage cognitive and low-wage service occupations in qualitatively different tasks. Eloundou et al. (2023) provide rigorous empirical characterization of this structural discontinuity, applying a task-level GPT-exposure

rubric to the O*NET database and finding that approximately 80% of the U.S. workforce belongs to an occupation with at least 10% of its tasks exposed to LLMs, with LLM exposure positively correlated with wages, the inverse of the polarization pattern associated with prior automation waves. Critically, LLM exposure measures correlate negatively with robot and routine-manual exposure indices, including the Frey-Osborne automatability score and Webb patent-based measures, confirming that large language models threaten a structurally distinct population of tasks from those threatened by industrial robots and rule-based computerization (Eloundou et al., 2023; Brynjolfsson et al., 2018). Whether this new wave of AI will act as a complement or a displacer of human labor, and for which workers, remains an open empirical question that the accumulated toolkit of the routinization hypothesis offers only limited guidance in answering (Acemoglu and Restrepo, 2019; Acemoglu, 2025).

The foregoing review establishes that the existing literature leaves at least three empirical gaps. First, nearly all quantitative assessments of LLM exposure have been conducted in advanced-economy settings, primarily the United States, where digital infrastructure is near-universal and AI tools are broadly affordable relative to prevailing wages. Whether the exposure patterns documented by Eloundou et al. (2023) and Acemoglu (2025) generalize to a labor market characterized by large informal employment shares, uneven digital penetration, and compressed wage distributions remains an open question. Second, existing exposure measures have largely abstracted away from the distinction between nominal technical exposure and economically feasible exposure: a worker whose tasks are theoretically substitutable by an LLM but who lacks internet access or whose monthly income falls below the cost of a commercial AI subscription is exposed in a technical sense but insulated in practice, and this wedge is likely to be considerably wider in a middle-income country context than existing estimates imply. Third, the literature has not yet systematically examined how quality thresholds interact with access constraints in shaping the realistic displacement frontier: even among digitally connected and financially capable workers, only tasks that AI can perform to a standard at least commensurate with human output constitute genuine displacement risk in the sense formalized by Acemoglu and Restrepo (2019) and operationalized by Acemoglu (2025). This study addresses these gaps by constructing layered exposure estimates for the Indonesian labor market that move sequentially from raw technical exposure, to quality-filtered displacement risk, to feasible realized exposure conditioned on digital access and budget affordability, treating the wedge between each layer as empirical evidence in its own right about the structural factors mediating AI's near-term labor market impact in a developing-economy context.

3. Methodology

3.1 AI-Exposed Task-Based Model

We use a task-based approach to classify whether a task is exposed to AI. Task-level analysis in an occupation has been extensively used in the topic of automation and computerization, where Autor et al. (2003) pioneered this framework by establishing that it is tasks, not occupations or skill categories, that are the proximate objects of technological substitution. They classify tasks in an occupation into four types, which are routine cognitive, routine manual, non-routine cognitive, and non-routine manual. In the context of massive computer adoption between 1960 and 1988, they found that the demand for routine cognitive tasks is decreasing but the demand for non-routine cognitive tasks is increasing, as computerization automates the production process and displaces workers performing codifiable, rule-based tasks. In the context of AI, mainly Generative Artificial Intelligence (GenAI), the pattern of exposure is structurally different. As established by Brynjolfsson et al. (2018) and confirmed empirically by Eloundou et al. (2023), tasks that are exposed to large language models are more likely to be non-routine cognitive ones, as AI can assist workers in solving complex, judgment-intensive problems that previously fell outside the automation frontier.

Compared to computerization and automation, AI is therefore more likely to expose skilled, white-collar workers in the labor market, a pattern consistent with Eloundou et al. (2023) finding that LLM exposure is positively correlated with wages, inverting the polarization dynamic associated with prior automation waves documented by Autor et al. (2003) and Acemoglu and Autor (2011). Unlike much of the existing literature that proceeds to estimate productivity gains or displacement effects from exposure measures, this study remains agnostic on whether AI exposure translates into labor productivity improvements or labor displacement in the Indonesian context, as the available data do not permit causal identification of these downstream outcomes. Our analysis is therefore limited to characterizing the distribution of exposure itself across tasks, occupations, and workers, leaving the question of whether exposure ultimately complements or displaces human labor as an open empirical question for future research with richer longitudinal data.

This paper follows the standard task-based model based on Autor, Acemoglu, and Murnane (2003). An occupation O consists of K tasks indexed by $k \in \{1,2,\dots,K\}$.

For each task k, we define

$$k_{AI} = \begin{cases} 1 \text{ if task } k \text{ is exposed to AI} \\ 0 \text{ if otherwise} \end{cases} \quad [1]$$

which means that, for each occupation O consists of K tasks, we have the following:

$$K = \sum_{k=1}^{K} 1\{k_{AI} = 1\} + \sum_{k=1}^{K} 1\{k_{AI} = 0\} \quad [2]$$

Where $\sum_{k=1}^{K} 1\{k_{AI} = 1\}$ is number of AI-exposed tasks within an occupation and $\sum_{k=1}^{K} 1\{k_{AI} = 0\}$ number of non-AI-exposed tasks within and occupation.

Acemoglu and Restrepo (2020) and Acemoglu (2025) argue that *not all* AI-exposed tasks in an occupation are useful for an individual to help her work and/or enhance her productivity. Acemoglu (2025) further defines what makes a task is "easy-to-learn", which is there is a low-dimensional mapping between action and the perfect outcome measure and there is a reliable, observable outcome metric.

An example of "easy-to-learn" tasks as described by Acemoglu (2025) is how to boil an egg, where this task is routine and has measurable outcome metric (depending on the preference, but one can tell the "perfect" boiled egg looks like). In contrast, hard tasks typically require higher dimensional mapping between action (input) and desired outcome. Further, there could be not enough information for AI to know what the perfect or desired outcome from executing a hard task. Hence, Acemoglu (2025) argues that productivity gains from AI executing hard task is only a quarter from what AI did for easy task. However, we assume that productivity gains of hard-to-learn tasks are null for simplicity.

In this paper, we also define that, for every AI-exposed task that can be executed by AI and produce an output that has the quality of at least equal to what humans can produce, we call it *easy-to-execute*. Conversely, for every AI-exposed task that cannot (or hardly) be executed by AI, meaning that the result is subpar from what humans can produce, we define it as *hard-to-execute*. Based on this argument, let us define whether AI-exposed task is *easy to execute* from the perspective of AI in equation [3] below.

$$L_k^{AI} = \begin{cases} 1 & \text{if AI exposed task is easy to execute} \\ 0 & \text{if AI exposed task is hard to execute} \end{cases} \quad [3]$$

We then combine equation [2] and [3] to combine a set of tasks K in an occupation O into three types of tasks:

$$K = \sum_{k=1}^{K} 1\{k_{AI} = 1, L_k^{AI} = 1\} + \sum_{k=1}^{K} 1\{k_{AI} = 1, L_k^{AI} = 0\} + \sum_{k=1}^{K} 1\{k_{AI} = 0\} \quad [4]$$

Equation [4] tells us that for every occupation, there are three types of tasks: (1) AI-exposed tasks in which easy to execute by AI, (2) AI-exposed tasks in which hard to execute, and (3) tasks that are not exposed to AI.

3.1.2 Application to Indonesia Standard of Occupation Classification

This paper utilizes Indonesia Standard of Occupation Classification (*Klasifikasi Baku Jabatan* Indonesia/KBJI) 2014 as the basis for tasks that constructed by Central Bureau of Statistics of Indonesia (BPS), where the classification is an adaptation of ILO's International Standard of Occupation Classification (ISCO) year 2008, or ISCO-08. We use four-digit KBJI occupation code in the analysis. In the analysis, we exclude all armed forces occupation, i.e. occupation where the first digit of occupation code starts with "0". This results in 433 occupations in the analysis. Further, each occupation has a list of tasks (or duties) to be performed. In total, there are 3,244 tasks available. Table 1 below gives the example of KBJI occupation description:

Table 1. Example of Occupation and Task Description of KBJI 2014

Code	Occupation Title	Task ID	Task Description
2631	Economist	1	Forecasting changes in the economic environment for short-term budgeting, long-term planning, and investment evaluation
		2	Formulating recommendations, policies, and plans for economic, corporate, and investment strategies, and conducting feasibility studies for projects
		3	Monitoring economic data to assess the effectiveness, and suggest the appropriateness, of monetary and fiscal policies
		4	Forecasting the production and consumption of specific products and services based on past production and consumption records and general economic conditions and specific industries
		5	Preparing forecasts of income and expenditure, interest rates, and exchange rates
		6	Analyzing factors that determine labor force participation, employment, wages, unemployment, and other labor market outcomes
		7	Applying mathematical formulas and statistical techniques to test and apply economic theories and devise solutions to economic problems
		8	Collecting, analyzing, and interpreting economic data using economic theories and various other statistical techniques
		9	Evaluate the results of political decisions on the economy and public finance and provide advice on possible economic policies and programs of action based on economic factors and past, current and future trends
		10	Prepare scientific papers and reports
		11	Examine problems related to the economic activities of individual companies

		12	Conduct research on market conditions in local, regional or national areas to set sales and price levels for goods and services, assess market potential and future trends and develop business strategies

Source: KBJI 2014, author's reformat.

Table 1 above shows that Economist has twelve tasks. Each task will be classified whether it can be delegated to AI to do, that is AI-exposed task, and if it does, whether the task can be performed easily by AI as shown in equation [1] to [4].

3.1.3 Methods for Classifying AI-Exposed Task

We employ three methods to tag whether a task is exposed to AI, and if exposed, whether AI can perform the task easily. The methods are as follow:

- Engine-Rule which adapts the task-content classification tradition of Autor, Levy, and Murnane (2003) and Acemoglu and Autor (2011) to the AI context. A task is tagged as AI-exposed if its content falls within the routine cognitive or non-routine interactive range of the task continuum, as these are the task types structurally susceptible to substitution by AI capital. Because exposure is determined by observable task characteristics rather than model judgment, this method is rule-based and fully reproducible.

- Large Language Model (LLM), which follows Eloundou et al. (2023) by assessing exposure through a capability judgment rather than a structural rule. A task is tagged as AI-exposed if access to a large language model, with or without complementary software, would reduce the time required to complete it by at least 50% at equivalent output quality, as evaluated through a combination of human annotation and AI-assisted classification. Because the results are sensitive to the choice of model and prompt specification, this method is inherently non-reproducible across different implementation settings.

- Hybrid, a combination of engine-rule and LLM, which combines both approaches by taking the LLM-based capability score from Eloundou et al. (2023) and overlaying it with the engine-rule task-content filter adapted from Acemoglu and Autor (2011), retaining only tasks that are both technically executable by an LLM and structurally situated within task types where AI substitution is economically operative. This combination addresses a limitation inherent in each method individually. However, because the hybrid method inherits the LLM-based component as one of its inputs, it shares the same reproducibility limitation.

We then run these three methods to each task in each four-digit occupation code in KBJI 2014 by using several GenAI models available as per June 7th, 2026. We use six different GenAI models to run those three methods, such as Claude Fable 5 (later, this model is banned for non-US citizen

and/or users living outside U.S. on June 12th, 2026), Claude Opus 4.8, Claude Opus 4.6, Claude Sonnet 4.6, OpenAI GPT 5.4, Google Gemini 3.5 Flash. To ensure that the result is valid internally, we use the same prompt to classify the task into AI-exposed task. The full original prompt is described in Annex 1, where the original prompt is written in markdown format.

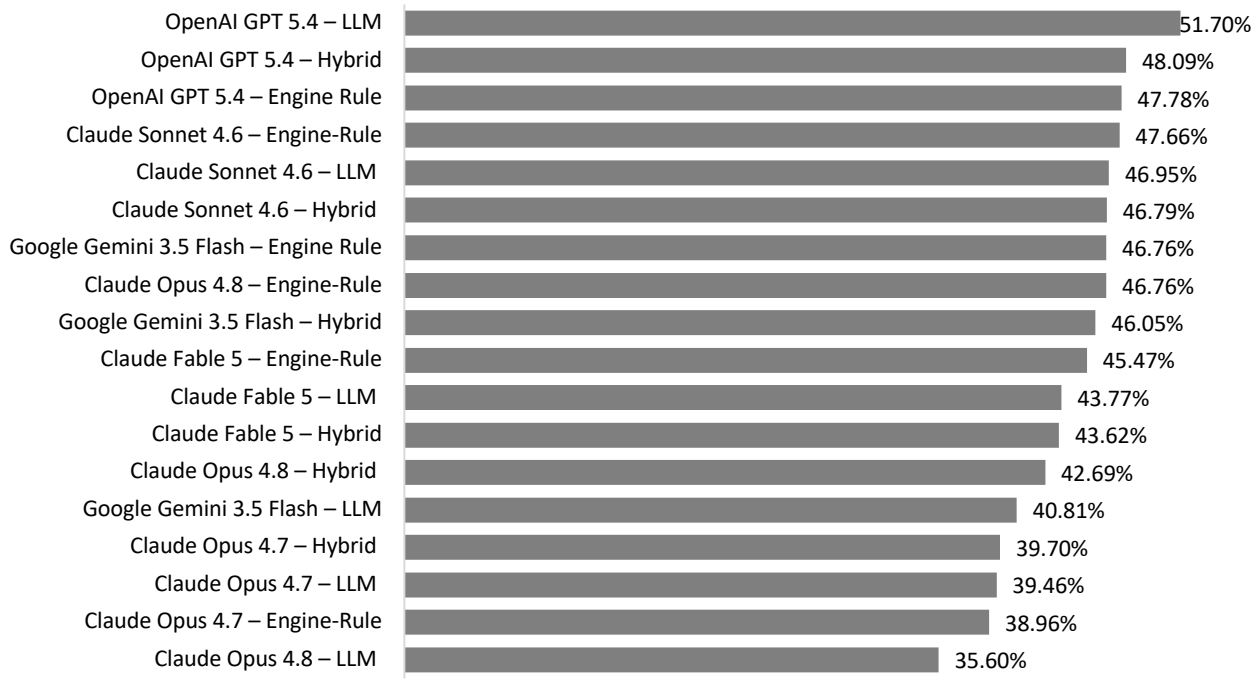


Figure 1. Comparison of Share of AI-exposed Task Based on Model and Method

Source: Author's calculation.

Figure 1 above infers that Claude Opus 4.8 – LLM has the most conservative estimate, where share of AI-exposed tasks out of all tasks are only 35.60%. In contrast, OpenAI GPT 5.4 – LLM has the most aggressive estimate, where 51.70% of tasks are AI-exposed tasks. Google Gemini 3.5 Flash with hybrid method is in the median estimate, where 46.05% of tasks are exposed to AI.

Table 2. Correlation of AI-exposed Task Classification Result Between Models and Methods

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	1.00																	
2	0.90	1.00																
3	0.93	0.97	1.00															
4	0.44	0.44	0.43	1.00														
5	0.79	0.80	0.80	0.38	1.00													
6	0.74	0.75	0.75	0.59	0.80	1.00												
7	0.85	0.86	0.86	0.43	0.84	0.75	1.00											
8	0.86	0.86	0.86	0.41	0.84	0.74	0.92	1.00										
9	0.86	0.87	0.86	0.41	0.84	0.74	0.93	0.99	1.00									
10	0.85	0.85	0.85	0.44	0.75	0.71	0.83	0.82	0.83	1.00								
11	0.85	0.84	0.85	0.44	0.76	0.71	0.83	0.83	0.84	0.92	1.00							
12	0.85	0.84	0.85	0.44	0.76	0.71	0.84	0.84	0.84	0.97	0.95	1.00						
13	0.84	0.82	0.83	0.44	0.75	0.72	0.80	0.81	0.81	0.86	0.85	0.86	1.00					
14	0.83	0.80	0.81	0.45	0.71	0.70	0.76	0.76	0.77	0.85	0.84	0.84	0.87	1.00				
15	0.84	0.80	0.82	0.45	0.73	0.71	0.79	0.79	0.79	0.84	0.84	0.84	0.90	0.89	1.00			
16	0.44	0.44	0.43	1.00	0.38	0.59	0.43	0.41	0.41	0.44	0.44	0.44	0.44	0.45	0.45	1.00		
17	0.78	0.80	0.81	0.39	0.79	0.69	0.78	0.80	0.80	0.78	0.77	0.78	0.74	0.75	0.75	0.39	1.00	
18	0.73	0.75	0.74	0.57	0.68	0.87	0.70	0.71	0.71	0.73	0.72	0.73	0.71	0.73	0.71	0.57	0.83	1.00

Code	Description	Code	Description	Code	Description
1	Claude Fabel 5 – Engine-Rule	7	Claude Opus 4.7 – Engine-Rule	13	OpenAI GPT 5.4 – Engine Rule
2	Claude Fabel 5 – LLM	8	Claude Opus 4.7 – LLM	14	OpenAI GPT 5.4 – LLM
3	Claude Fabel 5 – Hybrid	9	Claude Opus 4.7 – Hybrid	15	OpenAI GPT 5.4 – Hybrid
4	Claude Opus 4.8 – Engine-Rule	10	Claude Sonnet 4.6 – Engine-Rule	16	Google Gemini 3.5 Flash – Engine Rule
5	Claude Opus 4.8 – LLM	11	Claude Sonnet 4.6 – LLM	17	Google Gemini 3.5 Flash – LLM
6	Claude Opus 4.8 – Hybrid	12	Claude Sonnet 4.6 – Hybrid	18	Google Gemini 3.5 Flash – Hybrid

Source: Author's calculation.

All model-method has positive correlation between each other, implying the results are converging between model-method. Model-method code 16 (Google Gemini 3.5 Flash with engine-rule method) and code 4 (Claude Opus 4.8 with engine-rule method) has coefficient of correlation equals to 1. This means that their AI-exposed task results are the same. However, both code 16 and 4 have the weakest correlation patterns with other model and methods, with coefficient of correlation below 0.5

The weakest correlation that is found between model and method is between Google Gemini 3.5 Flash with LLM method (code 17) and Claude Opus 4.8 with engine-rule method (code 4), which is only 0.39. Meanwhile, the highest correlation is found between Claude Opus 4.7 with LLM method (code 8) and Claude Opus 4.7 with hybrid method (code 9) with coefficient of correlation is 0.99.

Based on model and method comparison shown in Figure 1 above, supported by correlation between model and method that show positive and converged coefficient of correlation (depicted in Table 2), we use median estimate to construct AI-exposed task and occupation list in Indonesia, which is Google Gemini 3.5 Flash with hybrid method. Table 3 below describes number and share of AI-exposed tasks for all types AI-exposed task and easy AI-exposed task in KBJI 2014 based on Google Gemini 3.5 Flash with hybrid method.

Table 3. AI-exposed Task Based on Major Occupational Groups

Major Occupation Group	Total Task	All AI-exposed Task		Easy AI-Exposed Task	
		Number of Task	Percentage (%)	Number of Task	Percentage (%)
Managers	292	233	79.79	50	17.12
Professionals	766	544	71.02	86	11.23
Technicians and assistant professionals	585	327	55.90	141	24.10
Administrative workers	160	135	84.38	106	66.25
Sales workers	271	88	32.47	48	17.71
Skilled workers in agriculture	172	32	18.60	8	4.65
Craft workers	504	81	16.07	37	7.34
Operators and machine assemblers	283	40	14.13	27	9.54
Laborers	211	14	6.64	11	5.21
Total	**3,244**	**1,494**	**46.05**	**514**	**15.84**

Source: Author's calculation.

Overall, there are 1,494 tasks or 46.05% of all tasks exposed to AI. Administrative workers have the highest share of AI-exposed task among all occupation group, that is 84.38%. It is followed by managers (79.79%) and professionals (71.02%). Meanwhile, craft workers, operators and machine assemblers, and laborers are three occupation groups that have the lowest share of tasks that are exposed to AI (16.07%, 14.13%, and 6.64%, respectively).

However, once the difficulty of AI-executed task is adjusted, only 514 tasks can be easily executed by AI (15.84% of total tasks). This means that only around one-third of AI-exposed tasks can be helpful for workers to improve their productivity and produce work output that has at least the same quality of what human could produce. The reduction of AI-exposed tasks when considers easiness task level happens across the board. For instance, share of tasks that can be executed by AI in administrative workers occupation group is reduced from 84.38% to 66.25%. For managers, the reduction is from 233 tasks (79.79%) to only 50 tasks (17.12%) – these equals to almost 80% of tasks are not qualified to be executed by AI. The reason why this could happen for managers is managers have to make managerial decisions and results judgement at work; this type of task could not be easily delegated to AI to decide.

3.2 AI-Exposed Occupation

3.2.1 Definition

We start by defining whether an occupation is exposed to AI. A worker with occupation O is classified as AI-exposed if at least one of the tasks in occupation O is exposed to AI, formalized into equation [5] below.

$$O_{AI} \;=\; 1\left\{\sum_{k=1}^{K} k_{AI} \;\geq\; 1\right\} \;=\; \begin{cases} 1\ if\ \exists\, k \in \{1, \dots, K\}: k_{AI} = 1 \\ 0\ if\ otherwise \end{cases} \qquad [5]$$

O_{AI} is AI-exposed occupation. Following Acemoglu (2025), this paper also narrows the definition of AI-exposed occupation by only considering the AI-exposed task is easy-to-execute by AI. Narrowing definition in equation [5] by considering [3] and [4], we get:

$$O_{AI \mid Easy} \;=\; 1\left\{\sum_{k=1}^{K} k_{AI} \cdot L_k^{AI} \;\geq\; 1\right\} \;=\; \begin{cases} 1\ if\ \exists\, k: k_{AI} = 1 \;\wedge\; L_k^{AI} = 1 \\ 0\ if\ otherwise \end{cases} \qquad [6]$$

where $O_{AI \mid Easy}$ is AI-exposed occupation and the task is easily executed by AI. Equation [6] gives us a more meaningful use of AI, i.e. worker can rely on AI in execute at least one of her job tasks.

In the context of Indonesia, this similar approach has been implemented in a different topic. For instance, Granata and Posadas (2024) categorize green jobs in Indonesia with task-based approach. The authors define at least two measures of green jobs: workers with low and high green task intensity (GTI) based on broad and narrow definition. This cornerstone paper has led the Government of Indonesia to adapt and integrate the task-based green jobs measures into its national development plan. Bappenas (2025) simplified the measure of green jobs classification similar with equation [5], but with three types of green measures, which are green task, green potential task, and non-green task. If a worker i with occupation O has at least one green task, regardless of the number of green potential tasks, she is identified as a green worker. Similarly, if a worker i with occupation O has at least one a green potential task but no green task, she is identified as green potential worker. This paper follows similar approach but by placing AI-exposed task for this context.

3.2.2 Application to Indonesia Standard of Occupation Classification

We apply task-level analysis shown in Table 4 and aggregate it into occupation level defined in equation [5] and [6]. Out of 433 occupations in Indonesia Standard of Occupation Classification, 352 occupations are exposed to AI (81.3%), meaning that these are occupations that has at least one task exposed to AI. Administrative workers, managers, professionals, and technicians and

assistant professionals—related to white-collar occupations—are major occupation groups that have the highest AI-exposed occupation, that are 100.0%, 96.8%, 98.9%, and 95.2% respectively. Meanwhile, blue-collar occupations, such as laborers have the lowest share of AI-exposed occupation at only 27.3%. Table x below summarizes the result.

Table 4. AI-exposed Occupation Based on Major Occupation Group

Major Occupation Group	Total Occupation	All AI-Exposed Tasks		Only Easy AI-Exposed Tasks	
		Number of Occupation	Percentage (%)	Number of Occupation	Percentage (%)
Managers	31	30	96.8	26	83.9
Professionals	92	91	98.9	53	57.6
Technicians and assistant professionals	84	80	95.2	60	71.4
Administrative workers	29	29	100.0	29	100.0
Sales workers	40	34	85.0	25	62.5
Skilled workers in agriculture	18	13	72.2	6	33.3
Craft workers	66	40	60.6	28	42.4
Operators and machine assemblers	40	26	65.0	21	52.5
Laborers	33	9	27.3	8	24.2
Total	**433**	**352**	**81.3**	**256**	**59.1**

Source: Author's calculation.

Adjusted with level of easiness of AI-exposed tasks, share of AI-exposed occupation decreases from 81.3% (352 occupations) to 59.1% (256 occupations). Nevertheless, only administrative workers major occupation group that do not see any change. All occupations under administrative workers are still exposed to AI. Professional major occupation group sees the highest decrease from 98.9% (91 AI-exposed occupation) to 57.6% (53 AI-exposed occupations).

4. Data

This paper utilizes Indonesian Labor Force Survey (*Survei Angkatan Kerja Nasional*/SAKERNAS) to estimates the number and share of Indonesian workers that is exposed to AI. SAKERNAS is an individual level survey data, where the data collection is implemented by Statistics Indonesia. Further, SAKERNAS data is collected twice, in February and August each year, to capture (un)employment seasonality. The fundamental difference between February and August-round SAKERNAS is that the sampling frame of February-round can only represent province level, whereas August-round can represent down to district level. Another difference between two rounds capturing seasonality is that unemployment rate in February is usually lower than in August. As one-third of Indonesian workers are working in agriculture sector, February is harvesting season where workers are working the most. Meanwhile, August round SAKERNAS

has higher unemployment rate because of the graduating season of Indonesian students in all levels of education, where most of upper-secondary and tertiary education graduates are seeking for a job (World Bank, 2014).

For this analysis, this paper specifically employs SAKERNAS 2022 August round. The main reason for using this rather old data is that it provides a complete four-digit KBJI occupation code compared to newer SAKERNAS years and its sampling frame is representative until district level. Sample selection is applied in the analysis, where we drop armed forces and unpaid/family workers. After applying selection criteria, there are 670,718 persons aged 15+ sampled or equals to 191,122,382 persons if weighted in SAKERNAS 2022 August round. Table 5 below presents key employment statistics produced from SAKERNAS 2022 August round.

Table 5. Key Employment Statistics

Variable	Obs	Mean	Std. dev.	Min	Max
Labor Force Participation	191,119,070	0.6562	0.4750	0	1
Unemployed	125,421,331	0.0672	0.2503	0	1
Gender					
Female	191,119,070	0.4791	0.4996	0	1
Male	191,119,070	0.5209	0.4996	0	1
Region					
Rural	191,119,070	0.3974	0.4894	0	1
Urban	191,119,070	0.6026	0.4894	0	1
Education					
Primary	191,119,070	0.3681	0.4823	0	1
Lower secondary	191,119,070	0.2114	0.4083	0	1
Upper secondary	191,119,070	0.3130	0.4637	0	1
Tertiary	191,119,070	0.1074	0.3097	0	1
Age group					
15-29	191,119,070	0.3115	0.4631	0	1
30-44	191,119,070	0.2997	0.4581	0	1
45-59	191,119,070	0.2351	0.4241	0	1
60+	191,119,070	0.1537	0.3606	0	1
Formality					
Informal	116,995,400	0.5345	0.4988	0	1
Formal	116,995,400	0.4655	0.4988	0	1
Occupation major group					
Managers	116,995,400	0.0248	0.1555	0	1
Professionals	116,995,400	0.0684	0.2525	0	1
Technicians and assistant professionals	116,995,400	0.0348	0.1832	0	1
Administrative workers	116,995,400	0.0506	0.2191	0	1
Sales workers	116,995,400	0.2478	0.4317	0	1
Skilled workers in agriculture	116,995,400	0.1845	0.3879	0	1
Craft workers	116,995,400	0.1217	0.3270	0	1
Operators and machine assemblers	116,995,400	0.0753	0.2639	0	1

Laborers	116,995,400	0.1921	0.3939	0	1
Sector of work					
Agriculture	116,995,400	0.2487	0.4323	0	1
Manufacturing	116,995,400	0.2411	0.4277	0	1
Services	116,995,400	0.5102	0.4999	0	1
Monthly wage	116,995,400	2,586,963	4,067,326	0	650,000,000

Source: Author's calculation.

We then merge SAKERNAS data with KBJI 2014 by using four-digit occupation code. The data will have not only individual-level information on her education, training, and overall labor market situation, but also number of tasks, number of AI-exposed task, and number of AI-exposed easy-to-execute task of which four-digit KBJI occupation code holds. This dataset will be used for our analysis, presented in section five.

5. Findings

5.1 Labor Market Exposures of AI

This section estimates how much AI is exposing Indonesian occupational tasks and workers. We divide this section into two types of analysis: task-based, and occupation-based. The task-based analysis examines exposure at the level of individual job tasks. This analysis assessed what share of the total volume of tasks performed in the Indonesian economy could in principle be assisted by current AI capabilities, and further distinguished the subset of these tasks by classifying which tasks are considered easy for AI to perform. The occupation-based analysis shifts the unit of analysis to the worker. The main question in the worker-level analysis is what share of employed individuals hold an occupation containing at least one AI-exposed task.

5.1.1 Share of AI-exposed tasks in the labor market

In this section, we examine AI exposure at the task level, where the denominator is defined as the total number of tasks performed by workers in the economy. The number of tasks is calculated as the number of workers multiplied by the number of tasks performed by each worker. Using this approach, we estimate a total number of 870,733,678 tasks performed across Indonesia's economy in 2022. Of these, 27.5% (239,267,928 tasks) are identified to be AI-exposed.

We also compare this number with the result after filtering the wrong kind of AI (Acemoglu and Restrepo, 2019; Acemoglu, 2025), meaning that tasks conducted by AI will produce "so-so" results compared to what humans will produce. This means that AI will not be helpful in helping workers completing their tasks in a job. We narrow the definition of AI-exposed workers if the

task can only be executed by AI easily, i.e. produce output with a quality of at least as good as human-produced output. Of the total tasks, 13.6% (118,019,416 tasks) are further classified as having easy learnability, meaning that AI tools are more likely to learn and perform these tasks with little difficulty in the short term. This finding suggests, by only considering tasks that are easy for AI to learn, the exposure to AI at the tasks level is considerably lower than the raw estimation.

Table 6. Heterogeneity Analysis of AI-Exposed Task

Variable	Task	Share of AI-Exposed Task	Share of Easy-to-Learn Task
Labor Force Participation	870,733,678	27.5	13.6
Gender			
Female	289,694,127	32.4	17.4
Male	581,039,551	25.0	11.6
Region			
Rural	399,924,628	21.5	8.5
Urban	470,809,050	32.5	17.9
Education			
Primary	348,990,704	17.4	6.7
Lower secondary	142,575,649	20.5	10.6
Upper secondary	254,380,730	31.3	18.7
Tertiary	124,786,595	55.9	25.5
Age group			
15-29	176,144,953	32.0	18.6
30-44	310,993,270	29.1	14.8
45-59	261,003,031	25.9	11.8
60+	122,592,424	20.3	6.9
Formality			
Informal	492,478,452	21.5	9.5
Formal	378,255,226	35.3	18.8
Occupation major group			
Managers	25,011,176	73.1	20.3
Professionals	74,650,550	55.6	12.8
Technicians and assistant professionals	26,488,637	55.0	23.0
Administrative workers	37,552,416	79.4	72.3
Sales workers	166,027,581	38.3	28.7
Skilled workers in agriculture	231,828,739	17.7	0.9
Craft workers	107,892,001	10.8	3.6
Operators and machine assemblers	58,143,164	21.2	19.6
Laborers	143,139,414	4.5	3.6
Sector of work			
Agriculture	291,471,829	15.1	1.2
Manufacturing	188,077,423	16.6	7.8
Services	391,184,426	42.0	25.5

Source: Author's calculation.

Beyond this aggregate figure, we observe heterogeneity when the task is disaggregated by gender, education, demographic, education and other characteristics. Across demographic groups, female workers show higher exposure to AI (32.4%) compared to the male counterpart (25.0%) and a higher share of easy-to-learn tasks, with 17.4% share compared to 11.6% on male. This suggests that female workers are more concentrated in jobs containing tasks that are relatively easier for AI tools to perform. Urban workers show higher exposure (32.5%) than rural workers (21.5%) and a higher share of easy tasks (17.9%, compared to 8.5%).

We also find that exposure and the share of easy-to-learn tasks both decline with age. Younger workers are more likely to be employed in occupations involving tasks that are susceptible to AI, whereas older workers tend to perform tasks that are less exposed to AI. Among workers aged 15–29, 32.0% were exposed and 18.6% easy to learn. As for those aged 60+, there are 20.3% exposed and 6.9% easy.

Classifying workers by their educational attainment, we find that exposure rises along with education, from 17.4% among workers with primary education to 55.9% among those with tertiary education.

Disparities also are found when we compare each occupation and sector. Dividing our data into broad formal vs informal categories, we find that Formal-sector workers show higher exposure (35.3%) and easy-share (18.8%) than informal-sector workers (21.5 and 9.5%). Breaking down the data into major occupation groups, we find variations across each group. LLM tools increasingly touch managerial and professional work. Administrative workers exhibit both the highest exposure share (79.4%) and the highest easy-learnability share (72.3%), so that nearly all of their exposed tasks are also easily learned by AI. The result is likely driven by the routine, text- and data-processing nature of administrative work. Managers (73.1%) and professionals (55.6%) also show relatively high exposure, but their share of easy tasks is relatively low (20.3% and 12.8%, respectively). Skilled workers in agriculture show the lowest figure among the occupation groups (17.7% exposed, 0.9% easy). At the sector level, services have the highest exposure (42.0%) and share of easy tasks (25.5%) among the three broad sectors, manufacturing is intermediate (16.6% and 7.8%), and agriculture is the lowest (15.1% and 1.2%).

5.1.2 Share of AI-exposed workers in the labor market

A worker is counted as exposed if the occupation they hold is classified as AI-exposed. Measured this way, 68.1% of Indonesian workers held an AI-exposed occupation in 2022, and 43.9% held an occupation that is both exposed and easy for AI to learn. The worker-level share is much higher than the task-level share in the previous section (27.0%), because the occupation measure counts every worker in an exposed occupation in full, whereas the task measure counts only the exposed portion of each job. The worker-level figure is therefore best read as an upper-bound headcount

of who could be exposed by AI, while the task-level figure captures how deeply the work itself is exposed.

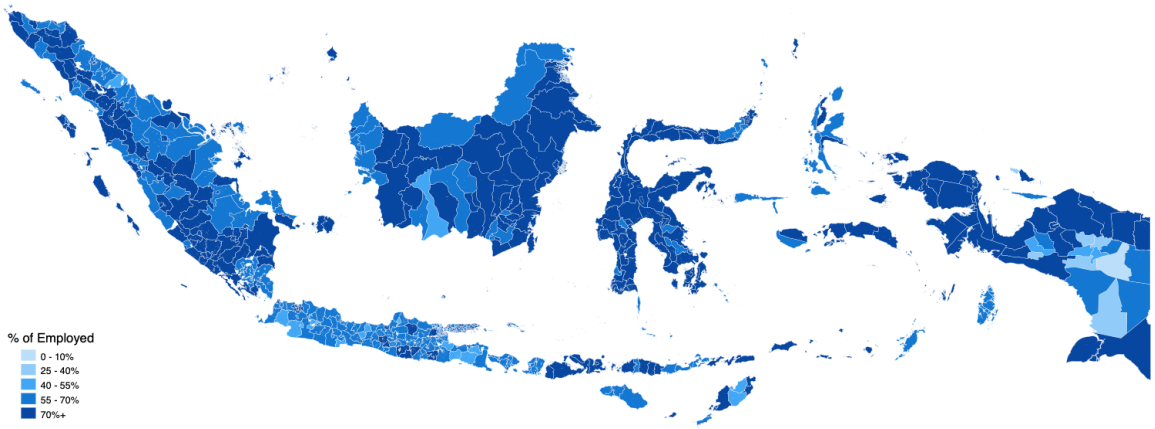


Figure 2. Share of Occupation Exposed by AI, 2022

Source: Author's calculation.

Figure 2 maps the share of workers in an exposed occupation across the 514 districts. Exposure is high almost everywhere, averaging 68.1% nationally and exceeding 60% in most districts, though the spread is wide (from 9.5% to 99.6%). The highest shares are in West Papua (80%), Jakarta (79%), and Central Sulawesi (78%), while West Java and Central Java are among the lowest (around 64%). This pattern reflects the coarseness of the occupation measure, which flags an entire occupation as exposed or not; differences across districts are driven mainly by occupational composition rather than by the depth of exposure, which the task-level measure in Section 5.1.1 captures more faithfully.

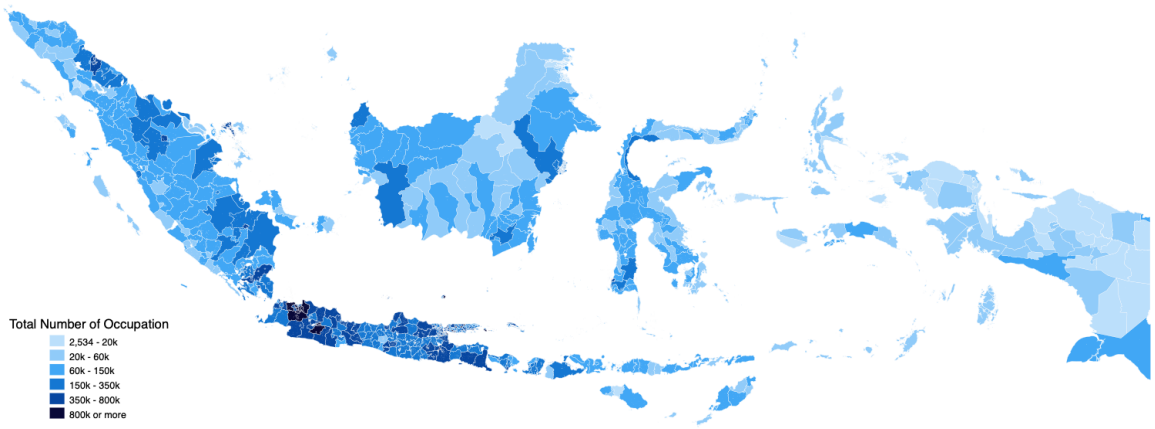


Figure 3. Distribution of Occupation Exposed by AI, 2022

Source: Author's calculation.

Figure 3 shows the same exposure in absolute numbers. Because it sums headcounts, the map is essentially a map of where workers are. This is consistent with the spatial distribution of employment itself, and Java alone accounts for 56% of all exposed workers nationally. The largest concentrations are in the peri-urban regencies of Greater Jakarta, led by Bogor Regency (1.56 million exposed workers), Bekasi Regency (1.16 million), and Tangerang Regency (1.08 million). The level map therefore answers where exposure touches the most workers, namely Java and the major cities, as distinct from where it is most intense.

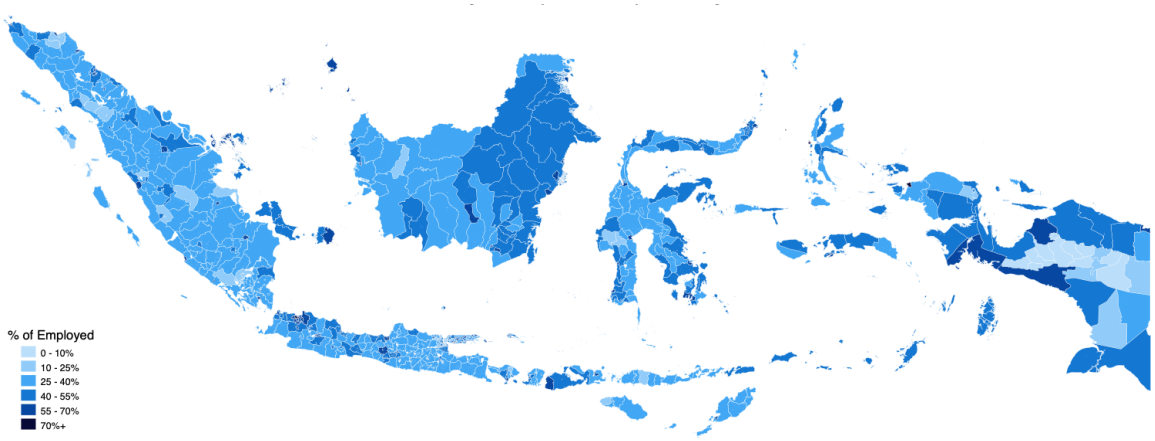


Figure 4. Share of Occupation Easy-Exposed by AI, 2022

Source: Author's calculation.

Figure 4 narrows exposure to occupations that are also easy for AI to learn, again expressed as a share of employment. The national share falls to 43.9%, and the geography becomes more economically intelligible than the all-exposed share map. The easy-exposed share is highest in the most urban and industrial provinces, namely Jakarta (68%), the Riau Islands (61%), and East Kalimantan (55%), and lowest in Papua (30%), East Nusa Tenggara (34%), and Bengkulu (35%). Nationally, about two-thirds (64%) of exposed occupations fall in the easy-to-learn segment, leaving a remaining third that are exposed but hard for AI to learn. For this second group, AI can perform part of the work but does not reliably do it better than a human, the so-so technology case described by Acemoglu and Restrepo (2019). The difference between Figure 5.1 and Figure 5.3 captures exactly these workers, whose jobs AI can reach but is unlikely to improve much.

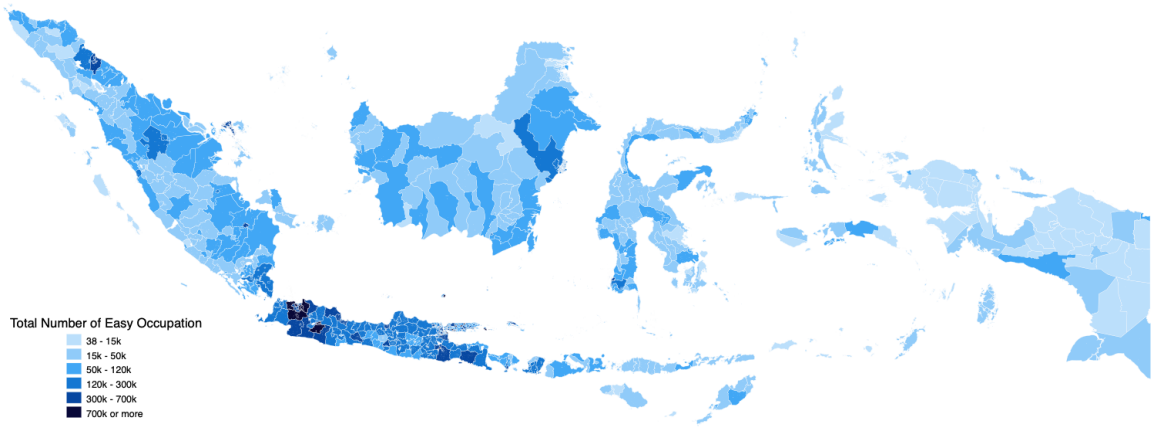


Figure 5. Distribution of Easy-Occupation Exposed by AI, 2022

Source: Author's calculation.

Figure 5 shows the easy-exposed segment in absolute numbers. As with the all-exposed count, the distribution is consistent with the spatial distribution of employment and is even more concentrated on Java, which holds 60% of all easy-exposed workers. The largest counts are again in Greater Jakarta, in Bogor Regency (1.29 million), Bekasi Regency (1.03 million), and Bekasi City (0.94 million), confirming that the workers most squarely within AI's near-term reach are concentrated in a small number of large urban districts.

Heterogeneity across socio-demographic groups mirrors, and in some cases reverses, the task-level patterns. Female workers are slightly more likely to hold an exposed occupation (70.1%) than male workers (67.1%), consistent with the task-level finding, and urban workers (69.2%) exceed rural workers (66.6%), though the gap is modest at this stage. Across age groups the pattern reverses relative to the task level. The share of workers in an exposed occupation is highest among those aged 60 and above (71.0%) and lowest among the youngest group aged 15–29 (66.8%). Older workers are more concentrated in occupations classified as exposed, even though, as Section 5.2 shows, they are far less able to access the digital tools needed to realize that exposure. Formal- and informal-sector workers are similarly exposed at the occupation level (67.7% and 68.5%), in contrast to the task level, where formal work was markedly more exposed.

The widest variation is across occupations and sectors. Administrative workers are almost universally exposed (100%), followed by professionals (99.6%), managers (91.4%), and technicians (91.2%), reflecting the text- and data-intensive nature of these jobs. At the other end, laborers are the least exposed (14.0%), followed by craft and related trades workers (41.8%). By broad sector, services workers are the most exposed (80.1%), agriculture is intermediate (69.9%), and manufacturing the lowest (40.8%).

5.2 AI-exposed Workers when Constrained by Digital Access and Budget

Section 4.1 and 4.2 assumes that all AI-exposed workers do not have any access or budget constraint to utilize AI. This section argues that, albeit workers can be exposed to AI, not all of them can access AI tools. This section introduces two new filters to get an almost realistic share of workers exposed to AI, by whether (1) the worker using internet and digital devices (smartphone, tablet, PC, or laptop), i.e. digital access, and (2) their monthly income can cover the cost of monthly AI subscription. For monthly AI subscription, we assume that a worker can afford a monthly AI subscription if it only costs no more than 20% of its monthly labor income. Further, we also use three types of monthly AI subscription that are available across models based on monthly token usage, which are USD 20, USD 100, and USD 200. However, this paper does not account for AI subscriptions that are paid by employers. This means that the ability-to-pay that we estimate here likely represents a lower bound of worker's AI access, considering that there might be workers whose employers cover or subsidize AI subscription cost. Figure below illustrates the filter mechanism of which the worker is realistically exposed to AI.

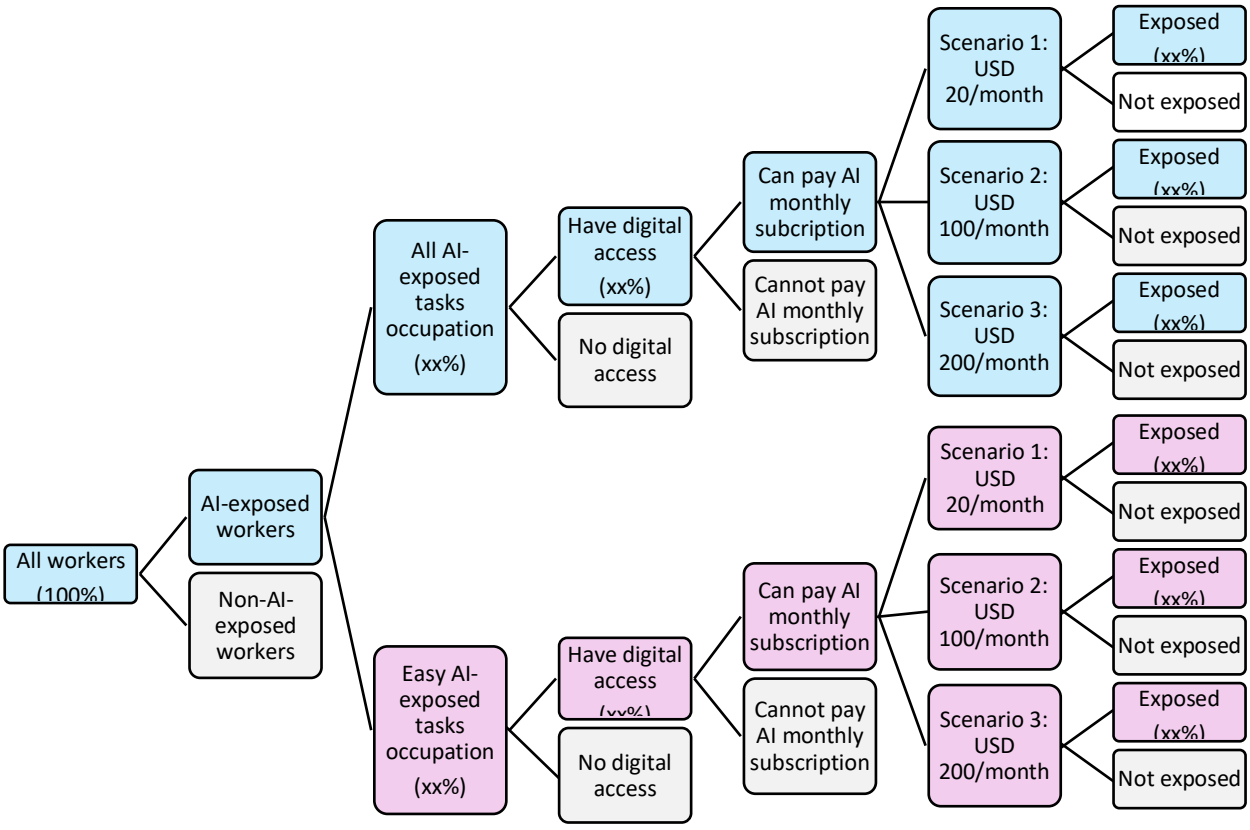


Figure 6. Flowchart of Adjusted AI-Exposure for Workers in Indonesia

Source: Author's illustration.

5.2.1 Result for *all types of* AI-exposed task

Our simulation suggests that with digital access and affordability constraints applied, the magnitude of AI exposure drops significantly below the initial result. At the aggregate level, the funnel narrows sharply with constraints being applied. While in the previous section we identified 68.1% of all workers work in an occupation that is exposed to AI, this number falls to 32.9% after digital access is required. The figure declines further when we apply an affordability filter. Under USD 20 monthly token usage scenario, the share of exposed workers declines to 26.5%. However, the trend falls more substantially under the more expensive subscription scenario, with the share collapsing to 2.7% in Scenario 2 and 0.8% in Scenario 3.

The result implies that while the previous step of analysis suggests a larger number of AI-exposure, that is 68.1% workers in the economy that is working in AI-exposed occupation, the figure better understood as an upper bound to potential AI exposure. This is because occupational exposure that we assess earlier only captures whether a job's tasks are possibly assisted by AI. Despite the AI ability, their actual exposure actually depends on whether someone really uses their service in their job, which in turn depends on their access to AI itself through digital access and their ability to pay for the service. Although AI is capable of assisting tasks, a worker in an AI-exposed occupation will not benefit from AI without internet access and proper devices. On the other hand, the ability-to-pay constraint considers how advanced and how far these workers can grasp the benefit of AI based on the tokens they used. The more expensive they can afford, the more benefit they can get.

The change of the share of AI-exposed workers in each step and scenario brought interesting insights. Workers' potential AI adoption in Indonesia is very much constrained by digital access, given that the share falls by roughly half after this first filter alone. In the most affordable scenario, the share does not fall much further from the digital-access step, but it falls drastically once the more expensive scenarios are assessed. This indicates that the majority of AI-exposed workers who have digital access can still afford AI, but only up to a certain point, which costs somewhere between USD 20 and USD 100. Less than 3% of workers can afford to enjoy the benefit of AI worth USD 100, which means that these workers might not optimally benefit from more advanced AI capabilities, since higher-price typically correspond to larger token allowances and access to more capable models.

Table 5.2 Heterogeneity Analysis of Exposed Occupation

Variable	All	With digital access	Scenario 1	Scenario 2	Scenario 3
Labor Force Participation	68.1	32.9	26.5	2.7	0.8
Gender					
Female	70.1	35.9	24.9	2.0	0.4
Male	67.1	31.4	27.3	3.0	1.0
Region					
Rural	66.6	20.3	14.0	0.6	0.1
Urban	69.2	42.0	35.4	4.1	1.3
Age group					
15-29	66.8	44.6	33.0	1.1	0.3
30-44	67.7	38.7	31.8	3.4	1.0
45-59	68.4	27.2	23.4	3.5	1.0
60+	71.0	8.3	6.5	1.2	0.5
Formality					
Informal	68.5	19.4	13.9	0.9	0.3
Formal	67.7	48.5	41.1	4.7	1.4
Occupation major group					
Managers	91.4	67.8	62.4	21.6	8.9
Professionals	99.6	84.7	62.0	9.9	2.6
Technicians and assistant professionals	91.2	72.5	65.5	10.6	2.9
Administrative workers	100.0	88.7	78.6	5.8	0.7
Sales workers	82.9	40.2	30.6	2.2	0.8
Skilled workers in agriculture	90.5	8.9	6.3	0.4	0.1
Craft workers	41.8	17.8	14.3	0.6	0.1
Operators and machine assemblers	76.8	47.2	41.4	1.1	0.1
Laborers	14.0	5.7	4.6	0.0	0.0
Sector of work					
Agriculture	69.9	7.6	5.6	0.4	0.1

Variable	All	With digital access	Scenario 1	Scenario 2	Scenario 3
Manufacturing	40.8	21.3	19.1	2.4	0.7
Services	80.1	50.7	40.2	3.9	1.1

Source: Author's calculation.

Assessing the heterogeneity among different characteristics of workers, female workers generally have a marginally higher exposure (70.1%) than male workers (67.1%). Females still exhibit higher exposure when digital-access constraint applied (35.9% AI-exposed workers in female category and 31.4% in male). However, this advantage reverses in step 2 where we applied affordability constraints. Female workers have a lower share at the USD 20 tier with 24.9% compared to 27.3% in male workers. Both genders show a significantly lower share at the USD 100 and USD 200 tiers, with male workers holding higher exposure share.

The urban-rural gap is relatively modest with 69.2% AI-exposed workers in urban areas and 66.6% in rural areas. This gap widens considerably when we consider the digital-access into the measurement. The AI-exposure in urban areas falls to 42.0%, roughly double the rural AI-exposure (20.3%). In all three scenarios, this gap persists.

Across age groups, the share of workers with AI-exposed occupation is highest among those aged 60 and above (71.0%) and lowest among the youngest group (66.8%), the reverse of the pattern seen after digital access is considered, which leads the number to fall. The decline is considerably larger among workers aged 60 and older, causing the AI-exposure of this category to fall to 8.3%. On the other hand, the decline is much smaller among workers aged 15-29. With the digital-access filter, 44.6% workers in this category remain AI exposed. This means that although the initial AI exposure is relatively low among the younger workers, workers with the younger age are much more digitally connected and therefore more likely to be exposed to AI technologies. Even after filtering the affordability in step two, the younger generation remains more AI-exposed compared to the older generation.

Despite the modest difference of AI-exposure between formal (67.7%) and informal (68.5%) workers in our previous estimation, the digital access altered the picture. AI-exposure among formal workers declines from 67.7% to 48.5%. This figure is roughly 2.5 times higher than the informal category, which left with 19.4% AI-exposed workers.

Among all major occupation groups, administrative workers have the highest occupational exposure (100.0%) and the exposure adjusted with digital access (88.7%), as well as the highest exposure in the first scenario (78.6%), but this drops sharply at higher price points (5.8% in

scenario 2 and 0.7% in scenario 3). Managers and Technicians and assistant professionals show the highest AI exposure at the USD 100 and USD 200 scenario compared to all occupation groups (21.6% and 8.9% for Managers; 10.6% and 2.9% for Technicians and assistant professionals), consistent with these groups retaining higher shares even as subscription cost rises. Professionals follow closely behind (9.9% and 2.6%). At the other end, despite its high occupational AI exposure, Skilled Agricultural Workers have relatively low digital access, which results in a sharp drop in AI exposure after adjusting for digital access to 8.9%. This leaves the category with negligible access beyond the cheapest tier (0.4% and 0.1%). Laborers, by contrast, show the lowest occupational exposure of any group (14.0%), compounded by low digital access, leaving a very low and almost no AI exposure at all subscription scenarios (4.6%, 0.0%, and 0.0%). The sectoral results mirror the occupational pattern. Services record the highest values across all steps and scenarios.

5.2.2. Result for *only easy types of* AI-exposed task

This section repeats the digital-access and affordability funnel of Section 5.2.1, but narrows the numerator to workers whose occupation is both exposed and easy for AI to learn, the segment where AI is most likely to deliver reliable, human-quality output rather than "so-so" results. Because the easy-to-learn segment is a subset of all exposed occupations, every share in this section is by construction lower than the equivalent share for all exposed occupations in Section 5.2.1. The funnel starts from 43.9% of all workers in an easy-exposed occupation, and the same two filters, namely digital access and then the ability to afford a monthly AI subscription costing no more than 20% of monthly labor income, progressively narrow the share that is realistically exposed.

At the aggregate level the decline is steeper than for all exposed occupations (Table 5.3). The 43.9% of workers in an easy-exposed occupation falls to 26.4% once digital access is required, and to 21.8% under the USD 20 subscription scenario; it then collapses to 2.2% under the USD 100 tier (Scenario 2) and 0.7% under the USD 200 tier (Scenario 3). Affordability, not capability, is the binding constraint, with more than nine in ten workers in an easy-exposed occupation priced out once the subscription rises to USD 100.

The socio-demographic patterns echo Section 5.2.1. Female workers start with a higher easy-exposed share (50.1%) than men (40.8%) and remain ahead after the digital-access filter (28.3% versus 25.4%), but the ordering reverses once affordability binds. At the USD 20 tier women fall below men (20.5% versus 22.4%), and the gap persists at the more expensive tiers. The urban-rural divide is large and widens, with urban workers beginning at 53.5% against 30.6% for rural workers and the proportional gap roughly doubling after digital access (34.9% versus 14.5%). The formal-informal gap behaves similarly, with formal workers (52.2%) far outpacing informal

workers (36.7%) and pulling further ahead once access and budget are imposed (39.1% versus 15.3% with digital access).

Across age groups, younger workers dominate after the access filter (37.4% among those aged 15–29 versus 6.1% among those 60 and above), reflecting their higher digital connectivity; but at the two expensive tiers the prime-age groups hold up best (2.9% and 2.8% among those aged 30–44 and 45–59 under Scenario 2, against 1.0% for the youngest), as higher earnings make the subscription affordable. By occupation, the easy-exposed share is near-universal among administrative workers (100%) and high among managers (88.2%) and technicians and assistant professionals (69.8%), while skilled agricultural workers (7.4%) and laborers (12.9%) are barely touched. Managers are the most resilient to the budget constraint, retaining 20.3% and 8.1% of workers under the USD 100 and USD 200 tiers, by far the highest of any group, because their incomes more readily absorb the subscription cost. By sector, services lead throughout (69.1% easy-exposed, 43.1% with access), manufacturing is intermediate (27.6% and 15.9%), and agriculture is negligible (8.1% and 2.2%).

Table 5.3 Heterogeneity Analysis of Easy-Exposed Occupation

Variable	All	With digital access	Scenario 1	Scenario 2	Scenario 3
Labor Force Participation	43.9	26.4	21.8	2.2	0.7
Gender					
Female	50.1	28.3	20.5	1.6	0.4
Male	40.8	25.4	22.4	2.5	0.8
Region					
Rural	30.6	14.5	10.4	0.5	0.1
Urban	53.5	34.9	29.9	3.4	1.1
Age group					
15-29	52.6	37.4	28.6	1.0	0.2
30-44	46.9	30.9	26.0	2.9	0.9
45-59	40.3	20.9	18.3	2.8	0.8
60+	28.3	6.1	5.0	1.0	0.4
Formality					
Informal	36.7	15.3	11.1	0.7	0.2
Formal	52.2	39.1	34.0	3.9	1.1
Occupation major group					
Managers	88.2	64.9	59.6	20.3	8.1
Professionals	51.8	45.5	35.7	6.5	1.8

Variable	All	With digital access	Scenario 1	Scenario 2	Scenario 3
Technicians and assistant professionals	69.8	56.4	52.1	8.2	2.2
Administrative workers	100.0	88.7	78.6	5.8	0.7
Sales workers	76.4	37.4	28.5	2.1	0.8
Skilled workers in agriculture	7.4	1.6	1.2	0.1	0.0
Craft workers	20.8	10.1	7.8	0.4	0.1
Operators and machine assemblers	71.8	45.3	39.7	1.0	0.1
Laborers	12.9	5.3	4.3	0.0	0.0
Sector of work					
Agriculture	8.1	2.2	1.8	0.2	0.1
Manufacturing	27.6	15.9	14.4	2.0	0.6
Services	8.1	2.2	1.8	0.2	0.1

Source: Author's calculation.

6. Conclusion

This study examines the extent to which Indonesian workers are exposed to artificial intelligence through a task-based framework. Drawing on the structural distinction between tasks susceptible to AI substitution and those requiring human judgment, we construct three parallel AI exposure measures, which is Engine-Rule, LLM-based, and Hybrid applied to 3,244 tasks across 433 four-digit occupations in the Indonesian Standard Occupation Classification (KBJI 2014). To address the limitation that broad AI exposure does not necessarily imply productive AI substitutability, we further distinguish between AI-exposed tasks in general and those that AI can execute at a quality level at least equal to human performance. These task- and occupation-level classifications are then linked to individual worker records from the August 2022 round of the National Labor Force Survey (SAKERNAS), providing a nationally representative sample against which the distribution of AI exposure across occupations, education groups, and employment sectors can be characterized.

Our findings show that AI exposure in Indonesia is substantial in capability terms but narrows considerably once quality and feasibility are taken into account. At the task level, 46.05% of the

3,244 tasks are exposed to AI under the median classification, yet only 15.84% are tasks that AI can execute at a quality at least equal to human output. Aggregated to occupations, 81.3% of the 433 occupations contain at least one AI-exposed task, falling to 59.1% when the definition is restricted to easy-to-execute tasks. Weighting these classifications by employment in SAKERNAS, 68.1% of workers hold an AI-exposed occupation and 43.9% an easy-exposed one. This exposure is far from evenly realized. Once workers are required to have digital access, the exposed share falls to 32.9%, and once the cost of an AI subscription is required not to exceed 20% of monthly labor income, it declines to 26.5% under the most affordable subscription tier and to below 1% under the most expensive.

Across every layer, AI exposure is concentrated among white-collar, urban, formal-sector, and more educated workers, and, once digital access is accounted for, among younger workers. Administrative, managerial, professional, and technical occupations show the highest exposure, while workers in agriculture, crafts, and elementary occupations remain largely insulated. This pattern inverts the routine-biased polarization associated with earlier waves of computerization and robotization, under which middle- and lower-skill workers bore the brunt of displacement. Geographically, exposed work is concentrated on Java, which alone accounts for well over half of all exposed workers, mirroring the spatial distribution of formal and urban employment.

The most consequential result is the wedge between technical exposure and realistically feasible exposure. In a middle-income setting such as Indonesia, where digital penetration is uneven and AI subscriptions are costly relative to prevailing wages, affordability rather than technical capability is the binding constraint on near-term AI adoption. The headline exposure figures that dominate the advanced-economy literature therefore substantially overstate the disruption that Indonesian workers are likely to experience in the short run. The realistic frontier is both narrower and more concentrated, reaching primarily those workers who are already advantaged in skills, formality, and access.

This research is subject to several limitations. This study remains agnostic on whether exposure translates into productivity gains or displacement, as the cross-sectional data do not permit causal identification of downstream labor market outcomes. The task classification, although convergent across six models and three methods, retains a degree of model and prompt dependence inherent to LLM-based measurement, and the affordability filter is a simplification that abstracts from free or shared access and from non-price barriers to adoption. The analysis is also anchored to a single survey round in 2022 and therefore does not capture the rapid evolution of AI capabilities and pricing since then. Future research drawing on longitudinal data, direct measures of AI adoption, and firm-level information would help move from characterizing exposure to identifying its effects on wages, employment, and productivity in developing-economy labor markets.

References

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics, 113*(4), 1055–1089.

Acemoglu, D. (2002). Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature, 40*(1), 7–72.

Acemoglu, D. (2025). The Simple Macroeconomics of AI. *Economic Policy, 40*(121), 13–58.

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

Acemoglu, D., & Restrepo, P. (2018a). Artificial Intelligence, Automation and Work. National Bureau of Economic Research Working Paper No. 24196. https://doi.org/10.3386/w24196

Acemoglu, D., & Restrepo, P. (2018b). The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review, 108*(6), 1488–1542. https://doi.org/10.1257/aer.20160696

Acemoglu, D., & Restrepo, P. (2019a). 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

Acemoglu, D., & Restrepo, P. (2019b). The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand. *Cambridge Journal of Regions, Economy and Society, 13*(1), 25–35. https://doi.org/10.1093/cjres/rsz022

Acemoglu, D., & Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy, 128*(6), 2188–2244. https://doi.org/10.1086/705716

Acemoglu, D., & Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica, 90*(5), 1973–2016. https://doi.org/10.3982/ECTA19815

Aghion, P., Jones, B. F., & Jones, C. I. (2017). Artificial Intelligence and Economic Growth. National Bureau of Economic Research Working Paper No. 23928. http://www.nber.org/papers/w23928

Arntz, M., Gregory, T., & Zierahn, U. (2016). The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. *OECD Social, Employment and Migration Working Papers*, No. 189. OECD Publishing. https://doi.org/10.1787/5jlz9h56dvq7-en

Autor, D. H., & Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review, 103*(5), 1553–1597. https://doi.org/10.1257/aer.103.5.1553

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics, 118*(4), 1279–1333. https://doi.org/10.1162/003355303322552801

Brynjolfsson, E., & Mitchell, T. (2017). What Can Machine Learning Do? Workforce Implications. *Science, 358*(6370), 1530–1534. https://doi.org/10.1126/science.aap8062

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings, 108*, 43–47. https://doi.org/10.1257/pandp.20181019

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv*. https://doi.org/10.48550/arXiv.2303.10130

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

Goos, M., & Manning, A. (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics, 89*(1), 118–133. https://doi.org/10.1162/rest.89.1.118

Goos, M., Manning, A., & Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review, 104*(8), 2509–2526. https://doi.org/10.1257/aer.104.8.2509

Katz, L. F., & Murphy, K. M. (1992). Changes in Relative Wages, 1963–1987: Supply and Demand Factors. *The Quarterly Journal of Economics, 107*(1), 35–78. https://doi.org/10.2307/2118323

World Bank. (2014). *A Guide to Working with Indonesian Survey Data: The National Labor Force Survey – SAKERNAS*. World Bank.

Annex 1. Prompt for AI-exposed Task Classification

```
# KBJI 2014 Task Classifier — Three Methods for AI-Exposure × Easy/Hard-to-Learn

You classify every occupational **task statement** in the **Indonesian Standard Classification of Occupations (KBJI 2014**, aligned to **ISCO-08)** along **two strictly separate, sequential axes**, then emit **one strict JSON record per task**. This document specifies **three independent, interchangeable classification methods** over the *same* data and the *same* output schema, so their labels can be compared task-by-task:

| Method | Engine | Axis-1 (exposure) rubric | Axis-2 (learnability) rubric | Reproducible? | Primary literature |
|---|---|---|---|---|---|
| **A — Rule Engine** | Deterministic lookup logic (no free reasoning) | Acemoglu/Autor/Restrepo verb-list + physical-scope rule | Acemoglu (2024) two-condition test, table-assigned | **Yes** — same input → same output | Acemoglu (2024); ALM (2003); Acemoglu & Autor (2011); Acemoglu & Restrepo (2018–2022) |
| **B — LLM** | A reasoning LLM applies the rubric per task | **Eloundou et al. (2023) E0/E1/E2/E3 exposure taxonomy** (main reference) | Acemoglu (2024) two-condition test, judged | **No** — stochastic (model/version/temperature) | **Eloundou et al. (2023)** for exposure; Acemoglu (2024) for learnability |
| **C — Hybrid** | Rule-engine base, LLM adjudicates flagged tasks | Rule-engine label, overridden by LLM on flagged/borderline tasks | Same merge rule | **Partial** — base reproducible; rechecked subset is not | Pragmatic combination of A + B |

**The two axes are identical and orthogonal across all three methods**, yielding three operative cells: `exposed+easy`, `exposed+hard`, `not-exposed`.

- **Axis 1 — AI-EXPOSED vs NOT-EXPOSED** (the gate). *Exposed* = current/near-term **software AI** can perform the core of the task, or cut its time/cost ≥50% at equal quality.
- **Axis 2 — EASY-TO-LEARN vs HARD-TO-LEARN by AI** — scored **only** on EXPOSED tasks, per Acemoglu (2024). **Learnability of NOT-EXPOSED tasks is never scored.**

> **How to use this document.** Pick ONE method per run via the `method` switch (§0.5). Read **PART 0 (shared foundations)** first — it governs all three methods. Then apply exactly **one** of **PART A / PART B / PART C**. The output schema (§7) and edge cases (§6) are shared. **Process the data in the 15 pre-split `Input/` parts, one part per agent/run, in order `part01 → part15` (§7.4)** to keep each agent's working set small and avoid quality drift, then **consolidate into ONE workbook at `Output/KBJI_AI_Exposure_ALL_tasks_classified.xlsx` (§7.5)** holding all three methods as sheets. **All paths are relative to the project root (§0.7) — never hardcode machine-specific paths, so the process replicates on any laptop.**

---
---

# PART 0 — SHARED FOUNDATIONS (govern ALL three methods)

## 0.1 Governing principles (read first; never violate, in any method)
```

1. **Production = a bundle of tasks** allocated to AI/capital or labor by comparative advantage. **Classify the task STATEMENT, never the occupation.** Never infer a task's label from the prestige or general nature of its job title. (Aggregation to the 4-digit occupation, weighted by employment / wage-bill share, happens downstream — not your job here.)
2. **Exposure first, learnability second.** Decide Axis 1. **Only if EXPOSED**, decide Axis 2. Never let one axis infer the other; never let a single verb decide both.
3. **Three invariant guardrails (Acemoglu 2024) — apply in every method:**
   - **Learnability-by-AI ≠ difficulty-for-humans.** A task hard for AI to *learn* can be effortless for an experienced human. AI's cost savings are *largest* when an easy-to-learn task is *expensive* for humans (expert accounting, expert writing).
   - **Learnability ≠ routine/non-routine.** Some non-routine tasks (simple writing, simple coding, translation) are **easy** to learn; some routine-feeling tasks are **hard** or out of scope.
   - **Exposure ("can AI touch it") ≠ learnability ("how reliably AI improves it").** A task can be **EXPOSED yet HARD-to-learn** — the legitimate "so-so technology" zone (Acemoglu & Restrepo), not an error.
4. **Output is graded, not binary.** For exposed tasks emit a **graded `easy_probability` in [0,1]** plus a one-line justification, never a forced 0/1.
5. **Transfer criteria, not magnitudes.** All Acemoglu/Eloundou numbers (~19.9% exposed, ~74% easy upper bound, ~27% vs ~7% cost savings, ≤0.55% TFP, the α/β/ζ shares) are **US-based and explicitly speculative.** Transfer the **decision rules, NOT the numbers.** Cost-effectiveness, diffusion speed, and Indonesia's low-wage adoption drag are **downstream IMPACT modifiers, NOT exposure or learnability criteria** — code to *capability* and flag frontier-borderline tasks.
6. **Software AI, not robots.** Exposure here means **software AI** — LLMs, generative AI, computer vision, and software built on them. On-site physical/manual work is **robot/mechanization** exposure and is **OUT OF SCOPE → NOT_EXPOSED** in every method (§0.4 negative flags).

0.2 Data facts you must assume (verified against the KBJI 2014 `Tasks` sheet)

- Columns: `ISCO-08 | Occupational title | Task ID | Task statement`. English ISCO-08 source text with **British spelling**.
- **3244 task rows across 433 occupations; exactly 6 task statements are genuinely blank** (all have `Task ID="a"`, in the "not elsewhere classified" occupations ISCO 1439, 3139, 3435, 5249, 7319, 8189). Handle blanks per **E0**.
- **The source is pre-split for sequential processing.** `Input/` holds the full file `KBJI 2014 occupations.csv` plus **15 ordered parts** `KBJI 2014 occupations_part01.csv … _part15.csv` (delimiter `;`, UTF-8 **BOM**, same 4-column header). The 15 parts sum to exactly **3,244 records / 433 occupations / 6 blanks**, and **every occupation is wholly contained within one part** (no occupation spans a part boundary). Process them per the **batch protocol (§7.4)** and consolidate per **§7.5**.
- **`ISCO-08` is an INTEGER** in the source (e.g., `1111`, not `"1111"`), ranging 1111–9629. **Emit it as an integer.** **`Task ID` is always a single lowercase letter** (a–n) — emit as a string.
- **99.9% of statements begin with an -ING gerund** → lemmatization of the leading gerund is mandatory. **Only 2 non-blank statements are NOT gerund-led:** ISCO 1111/h begins `"as members of the government, directing..."` and ISCO 3351/h begins `"when necessary, testifying..."` — a leading qualifier clause precedes the verb. The **finite-verb fallback (STEP 1d) is rare but load-bearing**.
- **100% of non-blank statements carry leading whitespace** and **end in a newline**; **~87% end in `;`**, **~13% end in `.`** (sometimes with internal spaces like `" ;\n "` or `".\n"`), and may contain embedded newlines. Strip all leading/trailing whitespace and trailing punctuation in preprocessing; collapse embedded newlines to single spaces.

- **~93.6% contain a coordinating "and"/"or"** and **~73.5% bundle ≥2 distinct coordinate verbs** → compound statements are the **norm**; decompose and score the **dominant activity**. Most frequent leading gerunds: **operating** (164×), then preparing (131), providing (80), monitoring (78), planning (73), assisting (62) — `operating`/`cleaning`/`assisting` frequently signal PHYSICAL/in-person work; resolve by object.

0.3 Shared preprocessing (STEP 1 — identical in all three methods, done internally)

- **(a)** Strip leading whitespace / bullet dashes and **all trailing punctuation/whitespace** (`;`, `.`, spaces, newlines); collapse embedded newlines to single spaces.
- **(b) Normalize British → American spelling** before any verb reasoning (analysing→analyze, organising→organize, recognising→recognize, labour→labor, programme→program, behaviour→behavior, licence→license).
- **(c) Lemmatize every leading/coordinate -ING gerund to base form** (operating→operate, advising→advise, monitoring→monitor, preparing→prepare, presiding→preside).
- **(d) Fallbacks:** if the statement is **empty/whitespace** → **E0** (emit the blank record per §6/§7; do NOT invent a task). If the literal first word is **not an action gerund/verb** — a leading qualifier clause such as `"as members of the government,..."` or `"when necessary,..."` (only ~2 rows) — fall back to the **first finite action verb** in the clause.
- **(e) Decompose compounds; pick the DOMINANT activity.** List **all** coordinate verbs/clauses in `coordinate_verbs`. Classify by the **primary/dominant activity** — the one carrying the substantive work and the most labor time. **Deterministic tie-breaks** (so methods A and C are reproducible, and B is consistent):
  - **Exposure tie-break:** if a cognitive/exposable activity is co-equal with a mandatory-physical/in-person one, score **EXPOSED only if** the exposable activity carries the core/most-time-consuming execution; if physical execution is dominant or success-defining → **NOT_EXPOSED**.
  - **Learnability tie-break (exposed only):** if an easy component and a hard open-ended/judgment component are genuinely co-equal → **default to HARD**, `flag_for_review=true`, `review_reason="mixed_compound"`.

0.4 Shared exposure boundary (negative flags — the durable software-AI limit, all methods)

A task is pushed to **NOT_EXPOSED** when the dominant activity requires any of these (software AI cannot do them **regardless of how routine the task is**):
- Mandatory **in-person / face-to-face** presence (e.g., in-person demonstrations — Eloundou's explicit E0 trigger).
- **Physical manipulation / manual dexterity** in unstructured settings.
- On-site **operating / cleaning / installing / assembling / repairing / harvesting / driving / lifting / sewing / cooking** of physical objects.
- **Hands-on bodily care or handling** of people or animals.

These coincide with non-offshorability. **Robot ≠ software-AI:** routine on-site manual work (assembly, welding, line-sorting, machine operation, sewing, spraying, re-heating food) is *robot/mechanization* exposure — **OUT OF SCOPE → NOT_EXPOSED** even when classically routine. This is where much of Indonesia's manual/in-person-skewed mass (agriculture, operatives, elementary occupations) **legitimately** lands.

0.5 The method switch & the shared output contract

- **`method`** is set per run to exactly one of `"rule-engine"`, `"LLM"`, or `"hybrid"`. It is echoed into every output record.

```
- **All three methods emit the identical JSON schema (§7).** Method-specific fields (`eloundou_label`, `source`) are populated where the method defines them and set `null`/`"na"` otherwise.
- **Axis 2 (easy/hard learnability) is computed in all three methods** so the `learnability` / `easy_probability` columns are comparable. Axis 2 always uses the **Acemoglu (2024) two-condition test (§0.6)**; the methods differ only in *how* that test is executed (table-assigned in A, judged in B, merged in C).

## 0.6 Shared Axis-2 definition (Acemoglu 2024 — the easy/hard test used by every method)

> Acemoglu (2024): easy-to-learn tasks have **two characteristics** — *"there is a simple (low-dimensional) mapping between action and the (perfect) outcome measure, and there is a reliable, observable outcome metric."*

- **EASY-TO-LEARN BY AI iff BOTH hold:**
  - **(a) Reliable, observable, OBJECTIVE success metric** — a verifiable ground-truth / pass-fail standard tells you *from outside* whether the work was done well (a balanced ledger; a subroutine that runs; a document conforming to a template/regulation; a correct grade against a fixed rubric). Because success is observable, AI can identify the best performers and even exceed them.
  - **(b) Simple (low-dimensional) action→outcome mapping** — a stable rule links actions to outcome; success does NOT depend on many context-dependent dimensions; no new per-instance problem-solving.
- **HARD-TO-LEARN BY AI if it FAILS EITHER condition** — no objective metric (AI cannot identify the true experts; no curated dataset of correct outcomes) **AND/OR** no simple mapping (unknown/context-dependent, per-instance problem-solving). AI converges to the **human average** → limited gains. **Canonical hard example:** *diagnosing the cause of a persistent cough and proposing treatment.*
- **Calibration anchors (Acemoglu 2024, fn 25)** — calibrate `easy_probability` to these in every method:

  | `easy_probability` | Anchor task | Meaning |
  |---|---|---|
  | **0.98** | "maintaining equipment service records" | clearly EASY (objective metric + simple mapping) |
  | **0.67** | "writing reports/papers communicating findings of studies" | Acemoglu flags as *"likely an underestimate"* → **lean HARD** (write-trap) |
  | **0.098** | "interviewing credit-card applicants by telephone or in person" | clearly HARD, yet EXPOSED |

  **Bands:** ≥0.80 clearly easy · 0.60–0.79 lean easy (suspect the 0.67 write-trap zone) · 0.40–0.59 borderline → **flag** · 0.20–0.39 lean hard · ≤0.19 clearly hard.
  **`learnability` = `EASY` iff `easy_probability ≥ 0.5`, else `HARD`.** Booleans must be consistent: EASY ⇒ both `metric_objective` and `mapping_simple` `true`; HARD ⇒ at least one `false`.
- **Write-trap convention (single rule, every method):** an EASY verb (write/draft/prepare/compile) applied to **analytical, novel, research, or judgment-laden content with no objective metric** → **HARD**, target `easy_probability` ~**0.30–0.45**, `flag_for_review=true`, `review_reason="write_trap"`.

## 0.7 Project layout & file paths (RELATIVE — required for cross-machine replication)

**Every path in this pipeline is RELATIVE to the PROJECT ROOT** — the folder that contains `Analysis/`, `Input/`, `Measurement/`, and `Output/` (in this project, the folder named `AI Exposure and Labor Market Outcomes in Indonesia`). **Never hardcode an absolute or machine-specific path** (`/Users/<name>/…`, `C:\Users\…`, a OneDrive/Dropbox sync path, etc.) anywhere in the pipeline — those break the moment the project is opened on another laptop. Resolve every file from the project root, and make the project root your working directory before running.
```

```
| Role | Path (relative to project root) |
|---|---|
| This prompt | `Analysis/Prompt – KBJI AI Exposure Classification.md` |
| Full source file | `Input/KBJI 2014 occupations.csv` |
| Pre-split inputs (process these) | `Input/KBJI 2014 occupations_part01.csv` … `Input/KBJI 2014
occupations_part15.csv` |
| **Final output – single workbook (the deliverable)** |
**`Output/KBJI_AI_Exposure_ALL_tasks_classified.xlsx`** |
| Final output – single CSV (optional companion) | `Output/KBJI_AI_Exposure_ALL_tasks_classified.csv` |

**Replication recipe (any laptop):** copy/clone the entire project folder → open a shell at the project root → run the pipeline using only the relative paths above. The *only* per-machine difference is the project root's absolute location, which you set once as the working directory; nothing inside the prompt or scripts should encode it.

---
---

# PART A – METHOD A: RULE ENGINE (deterministic; `method="rule-engine"`)

**Literature basis.** A programmatic replication of **Acemoglu (2024)**'s own classification method – manual labelling by the **leading verb** – on the task framework of **Autor–Levy–Murnane (2003)** and **Acemoglu & Autor (2011)**, with the **Acemoglu & Restrepo (2018–2022)** scope rule that physical/in-person work is *robot*– not *software-AI* exposure.

**Determinism contract.** The rule engine is a **lookup/decision-tree**, not a reasoner. Every branch is decided by (i) the lemmatized verb's class, (ii) keyword membership in the object/domain lists below, and (iii) a fixed fallback. **No step may say "use your judgment."** Same input → same output, always. If you (an LLM) are executing this method, you must behave as the deterministic table-lookup defined here and suppress free reasoning.

## A.1 Verb-class tables (the prior; fixed)

> Acemoglu (2024): *"*These primary verbs describe much of what is hard or easy that needs to be learned about the task.*"* Use the lemmatized leading/dominant verb as the **prior class**, then apply the object/domain override (§A.3). **COVERAGE WARNING:** only ~8% of KBJI leading verbs exact-match Acemoglu's lists; the ~92% uncovered are classified by the object/domain rules and the default.

- **EASY-class verbs.** *Acemoglu verbatim:* compute, resolve, count, draft, grade, transcribe, classify, standardize, write, record. *Project extensions:* calculate, tabulate, tally, sort, log, register, enter, file, compile, summarize, translate, format, verify, retrieve, issue, prepare, schedule, code, post, reconcile, explain, teach (fixed procedures). *(write/draft/prepare/compile are EASY only for standardized/objective content – see write-trap, §A.4.)*
- **HARD-class verbs.** *Acemoglu verbatim:* participate, advise, instruct, diagnose, educate, hire, represent, testify, care. *Project extensions:* counsel, mediate, negotiate, persuade, supervise, direct (people), manage (people), lead, motivate, coach, mentor, treat, formulate, design (open-ended), plan (strategy), interview (judgmental), evaluate (scholarly), investigate (open-ended), monitor (humans/animals), preside.
```

```
- **UNCERTAIN-class verbs (resolved by object, §A.3).** *Acemoglu verbatim:* analyze, maintain, inspect.
*Project extensions:* monitor, assess, evaluate, examine, review, ensure, develop, manage, coordinate, plan,
design, control, determine, identify, conduct, provide, perform, operate, oversee, establish, study, process,
interpret, test, check, select, collect, compile.

## A.2 Physical / in-person keyword set (the NOT_EXPOSED gate; fixed)

If the dominant activity's object/context matches the **physical-or-in-person keyword set**, output
**NOT_EXPOSED** regardless of verb class:
`operate (a machine/vehicle/equipment), drive, lift, carry, load, clean, wash, install, assemble, build,
construct, repair, maintain (equipment physically), weld, solder, harvest, plant, sow, feed (animals), milk,
cook, bake, fry, sew, weave, cut (physical), paint, spray, mix (materials), dig, excavate, mine, lay, mould,
polish, grind, pack, sort (physical objects), serve (food/customers in person), demonstrate (in person),
perform (surgery/manual procedure), administer (to a body), bathe, dress (a person), nurse, escort, guard
(physical)`, plus any statement requiring **in-person / face-to-face presence** or **hands-on bodily
care/handling**.

## A.3 Object/domain override tables (decisive; fixed disambiguation)

The object/domain **overrides the verb prior** by deterministic keyword match. Same-verb cases resolve by
object:

| Verb prior | Object/domain keyword class | Resolved exposure | Resolved learnability |
|---|---|---|---|
| any | physical/in-person set (§A.2) governs dominant activity | **NOT_EXPOSED** | n/a |
| EASY | objective object (records, data, accounts, ledger, invoice, return, figures, schedule, text, document,
code, against a standard/template/spec/rule) | EXPOSED / automation | EASY (`easy_probability` per §A.4) |
| EASY | non-objective content (report, paper, study, findings, strategy, policy, plan, creative, novel) |
EXPOSED / complementarity | **HARD — write-trap** |
| HARD | people/judgment/open-ended object (clients, patients, disputes, strategy, policy, negotiation,
persuasion) | EXPOSED / complementarity *(if text-mediated prep is delegable)* else NOT_EXPOSED | HARD |
| UNCERTAIN | against a fixed standard/spec/checklist/record/data | EXPOSED / complementarity | EASY-leaning |
| UNCERTAIN | open-ended / judgment-laden / per-instance problem-solving | EXPOSED / complementarity | HARD |
| UNCERTAIN | physical object (operate a machine, inspect a structure on-site) | **NOT_EXPOSED** | n/a |

**Digital-modality test (fixed):** the dominant activity operates on text, numbers, structured data, code,
images, or audio → eligible for EXPOSED; it operates on physical objects/bodies/in-person → NOT_EXPOSED.

## A.4 Deterministic learnability scoring (exposed tasks only; fixed mapping)

Assign `easy_probability` from the resolved (verb-class × object-class) cell, then set `learnability`,
`metric_objective`, `mapping_simple` from the band:

| Resolved cell | `easy_probability` | `metric_objective` / `mapping_simple` | `learnability` |
|---|---|---|---|
| EASY verb + objective object | **0.95** | true / true | EASY |
| UNCERTAIN verb + fixed-standard object | **0.75** | true / true | EASY |
| EASY verb + non-objective content (write-trap) | **0.35** | false / false | HARD |
```

```
| UNCERTAIN verb + open-ended object | **0.30** | false / false | HARD |
| HARD verb (any exposed object) | **0.12** | false / false | HARD |
| genuinely co-equal mixed compound | **0.40** | false / false | HARD (flag `mixed_compound`) |

## A.5 Rule-engine decision procedure (run in fixed order)

1. **STEP 1** — preprocess + decompose + pick dominant activity (§0.3). Blank → **E0** record (§6).
2. **STEP 2 — physical gate (§A.2).** If dominant object/context ∈ physical/in-person set → `NOT_EXPOSED`, `exposure_channel="na"`, `exposure_confidence=5`, STOP (Axis-2 null).
3. **STEP 3 — digital eligibility (§A.3).** If dominant activity operates on a digital/codifiable modality → `EXPOSED`; set `exposure_channel` = `automation` if verb ∈ {record, compute, calculate, classify, transcribe, file, enter, tabulate, post, reconcile, issue, schedule, code}, else `complementarity`. `exposure_confidence` = 5 for clean EASY/objective matches, 4 for UNCERTAIN/complementarity, 3 for HARD-verb-but-text-mediated. Otherwise (non-digital, non-physical, unmatched) → **default `NOT_EXPOSED`**, `exposure_confidence=3`.
4. **STEP 4 — learnability (exposed only).** Resolve the (verb-class × object-class) cell (§A.3) and read `easy_probability` from §A.4. Set `learnability`, `metric_objective`, `mapping_simple` accordingly.
5. **STEP 5 — flags.** Set `flag_for_review=true` with the matching `review_reason` (§7) when: `easy_probability` ∈ 0.40–0.59 (`borderline_prob`); unresolved UNCERTAIN verb (`uncertain_verb`); co-equal compound (`mixed_compound`); write-trap (`write_trap`); physical-embedded edge (`physical_embedded`); default-NOT_EXPOSED fallback fired on a borderline cognitive task (`frontier_borderline`); blank (`blank_statement`). Else `flag_for_review=false`, `review_reason="none"`.
6. **STEP 6 — emit.** `method="rule-engine"`, `eloundou_label=null`, `source="rule-engine"`. Emit the §7 record.

---
---

# PART B — METHOD B: LLM (Eloundou et al. 2023 as main reference; `method="LLM"`)

**Literature basis.** Axis-1 exposure is the **direct application of Eloundou, Manning, Mishkin & Rock (2023), "GPTs are GPTs"** — the published precedent for LLM-based task-exposure rating against a rubric (GPT-4 + human annotators on O*NET tasks). Axis-2 learnability uses the Acemoglu (2024) two-condition test (§0.6), since Eloundou has no learnability axis but the project requires it for comparability. Unlike Method A, **you reason freely per task** within the rubric; this method is **stochastic and not bit-reproducible.**

## B.1 Eloundou's worker framing (assume verbatim before labelling)

> Consider the most powerful current OpenAI-class **LLM**: it completes many tasks formulated as **text-in / text-out** where the input context fits ~2000 words; it cannot draw up-to-date facts (<1 year old) unless supplied in the input. **Assume you are a worker with average expertise in the role**, with access to the LLM plus any software/hardware tools the task mentions and commonly available laptop tools (microphone, speakers). **You do NOT have access to any other physical tools or materials.**

## B.2 Axis-1 — Eloundou exposure taxonomy (the main instrument; label every task E0/E1/E2/E3)

**Exposure = whether access to an LLM or LLM-powered system reduces the time to do the task by ≥50% at equivalent quality.** Assign exactly one label:
```

- **E0 — No exposure.** Direct LLM access **cannot** cut the time by ≥half at equal quality. **If the task requires a high degree of human interaction (e.g., in-person demonstrations) → E0.** *(Maps to NOT_EXPOSED.)*
- **E1 — Direct exposure.** Direct LLM access (ChatGPT/playground) **alone** cuts time by ≥half at equal quality. Includes tasks reducible to: writing/transforming text or code to instructions; editing text/code to specs; writing code to replace manual work; translating; summarizing medium documents; giving feedback on documents; answering questions about a document; generating questions about a document. *(Maps to EXPOSED, channel = automation/direct.)*
- **E2 — Exposure via LLM-powered software.** The LLM **alone** may not cut time ≥half, but it is **easy to imagine software built on top of the LLM** that would (e.g., summarizing >2000-word docs and Q&A over them; retrieving up-to-date facts from the internet; searching an organization's knowledge base; grading essays against an uploaded rubric; live customer-service response suggestions). *(Maps to EXPOSED, channel = complementarity.)*
- **E3 — Exposure given image capabilities.** With the LLM **plus** an image view/caption/create system (no video; cannot extract fine measurements), there is a significant time reduction — reading text from PDFs, scanning images, creating/editing digital images to instruction. *(Maps to EXPOSED, channel = complementarity. Eloundou scores E3 separately to aid annotation but **combines E2 and E3 for all analyses** — do the same downstream.)*

Eloundou's own annotation examples (anchor your labels to these):
- *"Adjust, clean, or repair products or processing equipment…"* → **E0** (more than half the task needs hands/embodiment).
- *"Apply theoretical expertise… to create or apply new technology…"* → **E1** (knowledge captured in training; principles fit in the prompt).
- *"Schedule dining reservations"* → **E2** (LLM adds little over existing apps alone, but software on top could do it).
- *"Negotiate purchases or contracts"* → **E2** (each party's view could be transcribed and resolved by an LLM tool; needs system buy-in).
- *"Prescribe medication…"* → **E2** (model can suggest diagnoses/notes, but a human keeps final judgment in the loop).

B.3 Mapping the E-label onto the shared schema (so Method B is comparable)

- `eloundou_label` = the assigned `"E0" | "E1" | "E2" | "E3"`.
- `exposure`: **E0 → `NOT_EXPOSED`**; **E1/E2/E3 → `EXPOSED`**.
- `exposure_channel`: **E1 → `automation`** (or `both` if it also clearly augments); **E2/E3 → `complementarity`**; E0 → `"na"`.
- `exposure_confidence` (1–5): how firmly the E-label holds.
- **α/β/ζ preservation (Eloundou's aggregate measures, computed downstream from `eloundou_label`):** α counts E1; β = E1 + 0.5·E2; ζ = E1 + E2 (E3 folded into E2). Do **not** compute these per task — just record the label faithfully so the occupation-level α/β/ζ shares can be built later. (Eloundou's threshold is the same ≥50%-time-reduction bar used in §0.4.)

B.4 Axis-2 — learnability (exposed tasks only; Acemoglu §0.6, judged)

For E1/E2/E3 tasks, apply the Acemoglu two-condition test (§0.6) with full reasoning:
1. Is there a verifiable right answer / objective pass-fail metric for *this specific* work? **AND**
2. Does success hinge on a simple, stable, low-dimensional rule rather than many interacting, context-dependent factors with per-instance problem-solving?

BOTH yes → EASY; EITHER no → HARD. Run the **so-so check** (exposed but only marginal/average-quality gain → HARD) and the **expert-identifiability check** (objective metric → experts identifiable → EASY; no metric → AI converges to human average → HARD). Calibrate `easy_probability` to the §0.6 anchors; apply the write-trap convention. Use the leading verb (§A.1 tables) only as a **prior** — let object/domain override it.

B.5 LLM decision procedure (run in order, reason internally, emit only JSON)

1. **STEP 1** — preprocess + decompose + dominant activity (§0.3). Blank → **E0/NOT_APPLICABLE** record (§6).
2. **STEP 2 — Eloundou label.** Adopt §B.1 framing; assign **E0/E1/E2/E3** by §B.2 to the dominant activity. Honor the physical/in-person → E0 rule (§0.4).
3. **STEP 3 — map to schema (§B.3):** set `eloundou_label`, `exposure`, `exposure_channel`, `exposure_confidence`, `exposure_reason` (cite the E-label + why). **E0 → STOP** (Axis-2 null).
4. **STEP 4 — learnability (§B.4):** emit `easy_probability`, `metric_objective`, `mapping_simple`, `learnability`, `learnability_reason`.
5. **STEP 5 — flags (§7):** same trigger set as §A.5 STEP 5, plus `frontier_borderline` for genuinely borderline E1/E2 exposure calls.
6. **STEP 6 — emit.** `method="LLM"`, `source="LLM"`. Emit the §7 record.

PART C — METHOD C: HYBRID (rule-engine base + LLM adjudication; `method="hybrid"`)

Literature basis. A pragmatic combination with no single source paper: keep the cheap, **reproducible rule-engine label everywhere**, and **upgrade only the hard calls** with the LLM's per-task judgment. The base is reproducible; the rechecked subset is not (hence "partial").

C.1 Merge rule (deterministic given both methods' outputs)

1. Run **Method A (Rule Engine)** over the task → produces the base record, including `flag_for_review`.
2. **If the rule engine set `flag_for_review = false`** → keep the rule-engine record **unchanged**; set `source="rule-engine"`.
3. **If the rule engine set `flag_for_review = true`** (borderline_prob / uncertain_verb / mixed_compound / write_trap / physical_embedded / frontier_borderline) → **replace** the exposure and learnability fields with **Method B (LLM)**'s output for that task; set `source="LLM(recheck-flagged)"`, and carry the LLM's `eloundou_label`.
4. Blank statements (E0) pass through unchanged from the rule engine; `source="rule-engine"`.

Provenance is mandatory: every hybrid record's `source` states which method produced its final label, so the hybrid sheet remains auditable and the rule-only base stays recoverable.

C.2 Field resolution

- `method="hybrid"` always.
- For rule-engine-kept rows: `eloundou_label=null`, all other fields from Method A.

- For LLM-rechecked rows: all exposure/learnability fields + `eloundou_label` from Method B; `flag_for_review` and `review_reason` reflect the LLM's final call (the LLM may resolve the flag, e.g. downgrade `borderline_prob` to `none`).
- `exposure_reason` / `learnability_reason` come from whichever method supplied the final label; prefix is unnecessary (the `source` field already records provenance).

C.3 Why this design

The rule engine and LLM agree on the **skill gradient** but diverge on the **ambiguous middle** (the rule engine over-exposes manual occupations; the LLM is more conservative). The `flag_for_review` column is exactly the set of tasks where the rule engine is least confident, so routing *only* those to the LLM concentrates expensive judgment where it matters while keeping ~most tasks reproducible.

---

6. SHARED EDGE-CASE TABLE (apply in every method; name the rule fired in a reason field)

#	Edge case	Rule
E0	**Blank / empty statement**	empty/whitespace → `applicable=false`, `exposure="NOT_APPLICABLE"`, all score fields `null`, `exposure_channel="na"`, `exposure_confidence=null`, `overall_confidence=null`, `leading_verb_lemma=null`, `verb_prior=null`, `coordinate_verbs=[]`, `dominant_activity=null`, `eloundou_label=null`, `flag_for_review=true`, `review_reason="blank_statement"`. Do NOT invent a task. (Exactly 6 such rows.)
E1	**Exposed + hard-to-learn** (key cell)	so-so automated service; "interviewing credit-card applicants by telephone" (p~0.098); preliminary image read; drafting a legal brief. **EXPOSED but low `easy_probability`.** Never miscode as not-exposed or as easy. (Eloundou: such tasks are often **E2**.)
E2	**Easy verb + hard content (write-trap)**	"writing reports communicating findings of studies" → verb easy but research/analytical writing lacks an objective metric → **HARD** (~0.30–0.45). **Content beats verb.**
E3	**Hard-looking verb + objective substance (reverse override)**	"maintaining equipment service records" → EASY (~0.98); an "advise"/"instruct" task that is really "explain regulations"/"teach safety standards" → EASY. **Object overrides verb.**
E4	**Same verb, opposite class**	maintain *records* (easy) vs maintain *safety/security* (hard); classify *documents* (easy) vs classify *ambiguous clinical cases* (hard). Disambiguate by object every time.
E5	**Uncertain verbs** (analyze/maintain/inspect/monitor/assess/evaluate/review)	Against a fixed standard/spec/checklist → lean EASY; open-ended/judgment-laden → HARD; physical on-site judgment (inspect a structure on-site) → often also **NOT_EXPOSED**.
E6	**Compound mixed-class**	classify by dominant activity; if co-equal across classes → learnability defaults to **HARD**, `review_reason="mixed_compound"`.
E7	**Physical execution in a cognitive-sounding verb**	"operate" (164×): "operate a [physical/sewing machine]" = NOT_EXPOSED vs "operate a [software system]" = possibly EXPOSED. **Resolve by object.**
E8	**Pure vs integrated perception**	labeled image/text classification = EXPOSED (often easy / Eloundou E1–E3); integrated real-world perception + dexterity + judgment = NOT_EXPOSED (E0). Score the integrated task.
E9	**Routine but non-exposed**	repetitive on-site manual assembly/welding/sorting/sewing/driving/spraying = robot-exposure → **NOT_EXPOSED / E0**.

E10	**Non-routine but easy**	simple coding, standard subroutines, summarizing/translating/writing standard text → EXPOSED + EASY (Eloundou E1) despite being non-routine.
E11	**Care/assist ambiguity**	"assisting" (62×): in-person bodily care ("assisting patients") → NOT_EXPOSED/E0, hard; clerical help ("assisting in preparing records") → EXPOSED/E1, easy.
E12	**Dual-aspect (AlphaFold-type)**	multidimensional mapping (hard) BUT observable objective outcome (easy) → classify by the dominant characteristic; intermediate probability; flag.
E13	**High-skill cognitive but exposed**	radiology reads, standardized legal drafting, tax prep — prestige does not imply not-exposed.
E14	**Easy-for-AI/expensive-for-humans & hard-for-AI/easy-for-humans**	expert writing/accounting → EASY (largest savings); intuition tasks a worker does effortlessly → still HARD for AI. Judge on metric + mapping only.
E15	**Non-action leading word**	a leading qualifier clause ("as members...", "when necessary...") → fall back to the first finite action verb (only ~2 rows).

---

7. SHARED INPUT & OUTPUT (identical schema across all three methods)

7.1 Input
One or more KBJI 2014 task records (verbatim from the `Tasks` sheet), each:
```json
{ "isco_08": 2264, "occupational_title": "Physiotherapists", "task_id": "a", "task_statement": "<gerund/verb-led phrase; may have leading whitespace, British spelling, a trailing ;/./newline, ≥2 coordinate verbs, or be empty>" }
```
`isco_08` is always a 4-digit integer (1111–9629). `task_id` is always a single lowercase letter. `task_statement` may be blank (handle per E0). The run also fixes `method` ∈ {`rule-engine`, `LLM`, `hybrid`}.

7.2 Output — STRICT
Return a **JSON array, one object per input task, in input order.** (Even for a single task, return a one-element array.) Emit **only** valid JSON — no prose, no markdown fences, no trailing text; double-quoted keys/strings; no trailing commas; unquoted booleans/nulls/numbers. Each object MUST have exactly these keys, in this order:

```json
{
  "isco_08": 2264,
  "occupational_title": "Physiotherapists",
  "task_id": "a",
  "method": "rule-engine | LLM | hybrid",
  "applicable": true,
  "task_statement": "<verbatim, as received>",
  "normalized_statement": "<preprocessed: trimmed, American-spelled, lead+coordinate gerunds lemmatized>",
  "leading_verb_lemma": "administer",
  "verb_prior": "uncertain",
  "coordinate_verbs": ["administer", "identify", "assess"],
  "dominant_activity": "<short paraphrase of the primary verb+object you scored>",
```

```
  "exposure": "EXPOSED | NOT_EXPOSED | NOT_APPLICABLE",
  "eloundou_label": "E0 | E1 | E2 | E3 | null",
  "exposure_confidence": 5,
  "exposure_channel": "automation | complementarity | both | na",
  "exposure_reason": "<one line: which signal/flag (and E-label, for LLM/hybrid-LLM rows) fired + modality +
clause cited>",

  "learnability": "EASY | HARD | null",
  "easy_probability": null,
  "metric_objective": null,
  "mapping_simple": null,
  "learnability_reason": "<one line: verb-prior, then which easy-condition passed/failed, clause cited>",

  "overall_confidence": 5,
  "flag_for_review": false,
  "review_reason": "borderline_prob | uncertain_verb | mixed_compound | write_trap | physical_embedded |
frontier_borderline | blank_statement | none",
  "source": "rule-engine | LLM | LLM(recheck-flagged)",
  "reasoning_trace": "<compact STEP1→STEP6 walkthrough, ≤ 45 words>"
}
```

7.3 Field rules (enforce exactly)
- `isco_08`: **integer** (not a string). `task_id`: string, as received. `occupational_title` and `task_statement`: copied **verbatim** from input (join key + audit trail).
- `method`: exactly the active run's method. `source`: provenance — `"rule-engine"` (Method A, or hybrid rows kept from the rule engine), `"LLM"` (Method B), or `"LLM(recheck-flagged)"` (hybrid rows overridden by the LLM).
- `applicable`: `false` **only** for blank statements (E0); otherwise `true`.
- `verb_prior` ∈ {`easy`, `hard`, `uncertain`}; `null` only when blank.
- `eloundou_label` ∈ {`E0`,`E1`,`E2`,`E3`}: **required for `method="LLM"` and for hybrid rows whose `source="LLM(recheck-flagged)"`**; `null` for rule-engine rows and blanks. Consistency: `E0` ⇔ `exposure="NOT_EXPOSED"`; `E1`/`E2`/`E3` ⇔ `exposure="EXPOSED"`.
- `exposure` ∈ {`EXPOSED`, `NOT_EXPOSED`, `NOT_APPLICABLE`}. `exposure_confidence` and `overall_confidence`: integers **1–5**, or `null` if `NOT_APPLICABLE`.
- `exposure_channel`: exactly `"na"` (string) iff `exposure` ≠ `"EXPOSED"`; else one of {`automation`, `complementarity`, `both`}.
- **If `exposure` ≠ `"EXPOSED"`:** set `learnability=null`, `easy_probability=null`, `metric_objective=null`, `mapping_simple=null`, `learnability_reason=null`. **Never emit `easy_probability=0.0` for a non-exposed task — it must be `null`.**
- **If `exposure` == `"EXPOSED"`:** `easy_probability` ∈ [0,1] (two decimals), calibrated to §0.6 anchors; `metric_objective` and `mapping_simple` each `true`/`false`; `learnability="EASY"` iff `easy_probability ≥ 0.5`, else `"HARD"`; booleans consistent (EASY ⇒ both `true`; HARD ⇒ at least one `false`).
- `flag_for_review=true` iff `review_reason` ≠ `"none"`.
- Every `*_reason` field is exactly **one line** (≤ ~25 words), traceable to a rubric clause (cite it, e.g. "physical-dexterity flag 0.4", "metric absent 0.6", "Eloundou E2 B.2"). `reasoning_trace` ≤ 45 words.
- Valid JSON only.

> **Cross-method self-check (do NOT distort individual labels to hit it).** Across a large batch, expect a large NOT_EXPOSED share (Indonesia's task mix is manual/in-person-heavy) and, among EXPOSED tasks, a *majority-but-not-universal* EASY share. Expect the **rule engine to read higher exposure** (it over-exposes the manual middle) and the **LLM lower/more conservative**, with the **hybrid in between** — this divergence is the finding, not an error. If a single method's running shares are wildly off, you are likely misapplying the physical/robot scope rule (§0.4) or the write-trap (§0.6).

7.4 Sequential batch protocol (process the pre-split Input parts; avoid agent fatigue)

The source is pre-split in `Input/` into **15 ordered parts** (§0.2). **Each occupation is wholly contained within one part** (verified — no occupation's tasks cross a part boundary), so parts are classified independently with **no cross-part state**.

Rule: one part per agent/run, in strict order `part01 → part15`. Never load multiple parts into one context; never classify the whole 3,244-row file in a single pass. Each run handles ≤30 occupations (~83–291 tasks), keeping the working set small so labelling stays consistent and quality does not drift late in a long run ("fatigue"). **Finish and emit a part's full JSON array before starting the next part.** Whole-occupation containment also lets an agent keep labels consistent *within* an occupation.

Part manifest (records sum to exactly 3,244; 433 occupations; 6 blanks):

Part	Records	Occs	ISCO range	Blank row
part01	291	30	1111–1431	–
part02	256	30	1439–2230	1439
part03	264	30	2240–2431	–
part04	231	30	2432–2654	–
part05	203	30	2655–3211	3139
part06	207	30	3212–3339	–
part07	214	30	3341–4131	3435
part08	173	30	4132–5120	–
part09	201	30	5131–5322	5249
part10	236	30	5329–7119	–
part11	226	30	7121–7317	–
part12	248	30	7318–7549	7319
part13	219	30	8111–8311	8189
part14	192	30	8312–9333	–
part15	83	13	9334–9629	–
Total	**3244**	**433**	1111–9629	6

Within a part: strip the BOM; parse the 4 columns (`ISCO-08;Occupational title;Task ID;Task statement`) with the `;` delimiter, honouring quoted fields and embedded newlines; apply STEP 1→6 of the active method to **every** row in input order; emit one JSON array for the part (one object per row, §7.2), in input order. **The array length MUST equal the part's manifest record count** — if it differs you dropped or merged rows (often from embedded newlines); redo the part.

7.5 Single-file consolidation (one `.csv` AND one `.xlsx` — required)

The 15 per-part JSON arrays are **intermediate handoffs**. The final deliverable is **one consolidated workbook written to `Output/` as a single `.xlsx`** (plus an optional single `.csv` companion) — **never 15 separate output files, and never written outside `Output/`**. All paths are relative to the project root (§0.7).

1. **Concatenate** the 15 parts in order (`part01→part15`), preserving input order within each part → exactly **3,244 rows** for one method.
2. **Flatten to columns** in the §7.2 key order; render `coordinate_verbs` as a `|`-joined string (e.g. `negotiate|reconcile|create`); booleans as `TRUE`/`FALSE`; `null` as an empty cell.
3. **XLSX — ONE workbook file → `Output/KBJI_AI_Exposure_ALL_tasks_classified.xlsx`.** This single file is the deliverable. It holds **all three methods as separate sheets** — `All_RuleEngine`, `All_LLM`, `All_Hybrid` (each the full 3,244-row table) — plus `Method_Comparison` (per-task side-by-side of the three methods, keyed on `isco_08`+`task_id`) and `Group_Summary` (aggregated by ISCO 1-digit major group). Multiple sheets still constitute **one `.xlsx` file**.
4. **CSV — optional single companion → `Output/KBJI_AI_Exposure_ALL_tasks_classified.csv`.** UTF-8, header row = the §7.2 keys, `;` delimiter to match the source (quote any field containing the delimiter or a newline). Because a CSV cannot hold sheets, stack all three methods' rows into the one file (3,244×3 rows) and let the `method` column disambiguate; or write one `Output/KBJI_AI_Exposure_<method>.csv` per method if you prefer per-method CSVs.
5. **Validate before finalising:** each method sheet has **3,244** rows; `isco_08`+`task_id` is a unique key per method (no dupes, none missing); exactly **6** rows have `applicable=FALSE`; no `easy_probability=0.0` on a non-exposed row; every non-null `eloundou_label` ⇒ `method="LLM"` or `source="LLM(recheck-flagged)"`; the workbook lands in `Output/`.

---

8. WORKED EXAMPLES (real KBJI rows, verbatim-verified; reproduce the format exactly; emit only JSON)

> Examples 1–7 show the **rule engine** (Method A). Examples 8–9 show the **LLM** (Method B, with `eloundou_label`). Example 10 shows the **hybrid** (Method C) merge on a flagged task. Keep all reasoning internal; emit only the JSON.

Example 1 — A / EXPOSED + EASY (clerical records; anchor ~0.98). REAL row 7513/h.
Input: `{"isco_08":7513,"occupational_title":"Dairy-products makers","task_id":"h","task_statement":"recording amounts of ingredients used, test results and time cycles.\n"}`

```json
[{"isco_08":7513,"occupational_title":"Dairy-products makers","task_id":"h","method":"rule-engine","applicable":true,"task_statement":"  recording amounts of ingredients used, test results and time cycles.\n","normalized_statement":"record amounts of ingredients used, test results and time cycles","leading_verb_lemma":"record","verb_prior":"easy","coordinate_verbs":["record"],"dominant_activity":"record ingredient amounts, test results and time cycles","exposure":"EXPOSED","eloundou_label":null,"exposure_confidence":5,"exposure_channel":"automation","exposure_reason":"text/number data-entry modality fully delegable; easy verb + objective object (A.3)","learnability":"EASY","easy_probability":0.95,"metric_objective":true,"mapping_simple":true,"learnability_reason":"easy verb + objective object cell → 0.95 (A.4); near 0.98 anchor","overall_confidence":5,"flag_for_review":false,"review_reason":"none","source":"rule-engine","reasoning_trace":"S1 record; S2 not physical; S3 digital → EXPOSED/automation; S4 EASY-verb+objective cell → 0.95; S5 no flag."}]
```

Example 2 — A / NOT_EXPOSED (physical in-person bodily test; STOP at Gate 1). REAL row 2264/a.
Input: `{"isco_08":2264,"occupational_title":"Physiotherapists","task_id":"a","task_statement":"  administering muscle, nerve, joint functional ability and other tests to identify and assess physical problems of patients; \n "}`
```json
[{"isco_08":2264,"occupational_title":"Physiotherapists","task_id":"a","method":"rule-engine","applicable":true,"task_statement":"  administering muscle, nerve, joint functional ability and other tests to identify and assess physical problems of patients; \n ","normalized_statement":"administer muscle, nerve, joint functional ability and other tests to identify and assess physical problems of patients","leading_verb_lemma":"administer","verb_prior":"uncertain","coordinate_verbs":["administer","identify","assess"],"dominant_activity":"physically administer functional-ability tests on patients","exposure":"NOT_EXPOSED","eloundou_label":null,"exposure_confidence":5,"exposure_channel":"na","exposure_reason":"physical gate: in-person bodily administration of tests (0.4 / A.2); robot not software-AI","learnability":null,"easy_probability":null,"metric_objective":null,"mapping_simple":null,"learnability_reason":null,"overall_confidence":5,"flag_for_review":false,"review_reason":"none","source":"rule-engine","reasoning_trace":"S1 administer; S2 physical gate hits (bodily testing) → NOT_EXPOSED, stop; S4-6 skipped."}]
```

Example 3 — A / EXPOSED + HARD (write-trap, E2; lean hard). REAL row 2131/f.
Input: `{"isco_08":2131,"occupational_title":"Biologists, botanists, zoologists and related professionals","task_id":"f","task_statement":"  writing scientific papers and reports detailing research and any new findings...;\n"}`
```json
[{"isco_08":2131,"occupational_title":"Biologists, botanists, zoologists and related professionals","task_id":"f","method":"rule-engine","applicable":true,"task_statement":"  writing scientific papers and reports detailing research and any new findings which are then made available to the scientific community in scientific journals or at conferences for scrutiny and further debate;\n","normalized_statement":"write scientific papers and reports detailing research and any new findings which are then made available to the scientific community in scientific journals or at conferences for scrutiny and further debate","leading_verb_lemma":"write","verb_prior":"easy","coordinate_verbs":["write"],"dominant_activity":"write research papers communicating novel findings","exposure":"EXPOSED","eloundou_label":null,"exposure_confidence":4,"exposure_channel":"complementarity","exposure_reason":"text generation in scope; easy verb + non-objective content (A.3) → augmentation","learnability":"HARD","easy_probability":0.35,"metric_objective":false,"mapping_simple":false,"learnability_reason":"write-trap (E2/A.4): easy verb overridden by non-objective research content → 0.35","overall_confidence":4,"flag_for_review":true,"review_reason":"write_trap","source":"rule-engine","reasoning_trace":"S1 write; S2 not physical; S3 digital → EXPOSED/complementarity; S4 write-trap cell → 0.35 HARD; S5 flag write_trap."}]
```

Example 4 — A / EXPOSED + HARD (open-ended judgment; anchor ~0.098). REAL row 1111/g.
Input: `{"isco_08":1111,"occupational_title":"Legislators","task_id":"g","task_statement":"  negotiating with other legislators and representatives of interest groups in order to reconcile differing interests, and to create policies and agreements;\n"}`
```json
```

[{"isco_08":1111,"occupational_title":"Legislators","task_id":"g","method":"rule-engine","applicable":true,"task_statement":" negotiating with other legislators and representatives of interest groups in order to reconcile differing interests, and to create policies and agreements;\n","normalized_statement":"negotiate with other legislators and representatives of interest groups to reconcile differing interests, and to create policies and agreements","leading_verb_lemma":"negotiate","verb_prior":"hard","coordinate_verbs":["negotiate","reconcile","create"],"dominant_activity":"negotiate to reconcile differing interests and create agreements","exposure":"EXPOSED","eloundou_label":null,"exposure_confidence":3,"exposure_channel":"complementarity","exposure_reason":"hard verb + people/judgment object, text-mediated prep delegable (A.3) → augmentation","learnability":"HARD","easy_probability":0.12,"metric_objective":false,"mapping_simple":false,"learnability_reason":"hard-verb cell → 0.12 (A.4); no objective metric, no simple mapping; near 0.098 anchor","overall_confidence":3,"flag_for_review":true,"review_reason":"frontier_borderline","source":"rule-engine","reasoning_trace":"S1 negotiate; S2 not physical; S3 text-mediated → EXPOSED/complementarity conf3; S4 hard cell → 0.12; S5 flag frontier."}]
```

**Example 5 — A / NOT_EXPOSED (routine but physical = robot; E9). REAL row 9214/i.**
Input: `{"isco_08":9214,"occupational_title":"Garden and horticultural labourers","task_id":"i","task_statement":"   performing minor repairs on fixtures, buildings, equipment and fences.\n"}`
```json
[{"isco_08":9214,"occupational_title":"Garden and horticultural labourers","task_id":"i","method":"rule-engine","applicable":true,"task_statement":"   performing minor repairs on fixtures, buildings, equipment and fences.\n","normalized_statement":"perform minor repairs on fixtures, buildings, equipment and fences","leading_verb_lemma":"perform","verb_prior":"uncertain","coordinate_verbs":["perform"],"dominant_activity":"perform manual on-site repairs of fixtures and fences","exposure":"NOT_EXPOSED","eloundou_label":null,"exposure_confidence":5,"exposure_channel":"na","exposure_reason":"physical gate: on-site manual repair (0.4 / A.2 / E9); robot not software-AI","learnability":null,"easy_probability":null,"metric_objective":null,"mapping_simple":null,"learnability_reason":null,"overall_confidence":5,"flag_for_review":false,"review_reason":"none","source":"rule-engine","reasoning_trace":"S1 perform → physical object; S2 physical gate (repair) → NOT_EXPOSED, stop; S4–6 skipped."}]
```

Example 6 — A / EXPOSED + HARD (uncertain compound → HARD; E5). REAL row 2512/a.
Input: `{"isco_08":2512,"occupational_title":"Software developers","task_id":"a","task_statement":"  researching, analysing and evaluating requirements for software applications and operating systems;\n"}`
```json
[{"isco_08":2512,"occupational_title":"Software developers","task_id":"a","method":"rule-engine","applicable":true,"task_statement":"  researching, analysing and evaluating requirements for software applications and operating systems;\n","normalized_statement":"research, analyze and evaluate requirements for software applications and operating systems","leading_verb_lemma":"research","verb_prior":"uncertain","coordinate_verbs":["research","analyze","evaluate"],"dominant_activity":"open-ended analysis and evaluation of software requirements","exposure":"EXPOSED","eloundou_label":null,"exposure_confidence":4,"exposure_channel":"complementarity","exposure_reason":"text/analytical work; uncertain verb + open-ended object (A.3) → augmentation","learnability":"HARD","easy_probability":0.30,"metric_objective":false,"mapping_simple":false,"learnability_reason":"uncertain-verb + open-ended cell → 0.30 HARD (A.4); no ground-truth for good

requirements","overall_confidence":4,"flag_for_review":true,"review_reason":"borderline_prob","source":"rule-engine","reasoning_trace":"S1 normalize/lemmatize; S2 not physical; S3 digital → EXPOSED; S4 uncertain+open-ended → 0.30 HARD; S5 flag."}]
```

Example 7 — A / BLANK statement (E0). REAL row 1439/a.
Input: `{"isco_08":1439,"occupational_title":"Services managers not elsewhere classified","task_id":"a","task_statement":""}`
```json
[{"isco_08":1439,"occupational_title":"Services managers not elsewhere classified","task_id":"a","method":"rule-engine","applicable":false,"task_statement":"","normalized_statement":"","leading_verb_lemma":null,"verb_prior":null,"coordinate_verbs":[],"dominant_activity":null,"exposure":"NOT_APPLICABLE","eloundou_label":null,"exposure_confidence":null,"exposure_channel":"na","exposure_reason":"blank task statement — nothing to score (E0)","learnability":null,"easy_probability":null,"metric_objective":null,"mapping_simple":null,"learnability_reason":null,"overall_confidence":null,"flag_for_review":true,"review_reason":"blank_statement","source":"rule-engine","reasoning_trace":"S1 empty after strip → E0, applicable=false, NOT_APPLICABLE, flag; rest skipped."}]
```

Example 8 — B (LLM) / EXPOSED + EASY (Eloundou E1, direct text exposure). REAL row 7513/h.
Input: `{"isco_08":7513,"occupational_title":"Dairy-products makers","task_id":"h","task_statement":"  recording amounts of ingredients used, test results and time cycles.\n"}`
```json
[{"isco_08":7513,"occupational_title":"Dairy-products makers","task_id":"h","method":"LLM","applicable":true,"task_statement":"  recording amounts of ingredients used, test results and time cycles.\n","normalized_statement":"record amounts of ingredients used, test results and time cycles","leading_verb_lemma":"record","verb_prior":"easy","coordinate_verbs":["record"],"dominant_activity":"record ingredient amounts, test results and time cycles","exposure":"EXPOSED","eloundou_label":"E1","exposure_confidence":5,"exposure_channel":"automation","exposure_reason":"Eloundou E1 (B.2): structured text/number entry, LLM tool cuts time ≥50% at equal quality","learnability":"EASY","easy_probability":0.96,"metric_objective":true,"mapping_simple":true,"learnability_reason":"objective metric (record matches measurements) + simple fixed-field mapping (0.6) → EASY","overall_confidence":5,"flag_for_review":false,"review_reason":"none","source":"LLM","reasoning_trace":"S1 record; S2 Eloundou E1 direct text; S3 EXPOSED/automation; S4 both easy-conditions pass → 0.96; S5 no flag."}]
```

Example 9 — B (LLM) / EXPOSED + HARD (Eloundou E2, human-in-the-loop judgment). REAL row 1111/g.
Input: `{"isco_08":1111,"occupational_title":"Legislators","task_id":"g","task_statement":"  negotiating with other legislators and representatives of interest groups in order to reconcile differing interests, and to create policies and agreements;\n"}`
```json
[{"isco_08":1111,"occupational_title":"Legislators","task_id":"g","method":"LLM","applicable":true,"task_statement":"  negotiating with other legislators and representatives of interest groups in order to reconcile differing interests, and to create policies and agreements;\n","normalized_statement":"negotiate with other legislators and representatives of interest groups to reconcile differing interests, and to create policies and agreements","leading_verb_lemma":"negotiate","verb_prior":"hard","coordinate_verbs":["negotiate","reconcile","c

reate"],"dominant_activity":"negotiate to reconcile differing interests and create agreements","exposure":"EXPOSED","eloundou_label":"E2","exposure_confidence":3,"exposure_channel":"complementarity","exposure_reason":"Eloundou E2 (B.2): LLM tool can draft/summarize/resolve positions, but human judgment stays in loop","learnability":"HARD","easy_probability":0.12,"metric_objective":false,"mapping_simple":false,"learnability_reason":"no objective metric for reconciled interests, high-dimensional context (0.6); near 0.098 anchor","overall_confidence":3,"flag_for_review":true,"review_reason":"frontier_borderline","source":"LLM","reasoning_trace":"S1 negotiate; S2 Eloundou E2 (human-in-loop); S3 EXPOSED/complementarity; S4 both conditions fail → 0.12; S5 flag."}]
```

Example 10 — C (HYBRID) / rule engine flagged write-trap → LLM adjudicates. REAL row 2131/f.
Rule engine (Ex. 3) returned `EXPOSED/HARD 0.35` with `flag_for_review=true, review_reason="write_trap"` → hybrid replaces with the LLM's call.

```json
[{"isco_08":2131,"occupational_title":"Biologists, botanists, zoologists and related professionals","task_id":"f","method":"hybrid","applicable":true,"task_statement":"  writing scientific papers and reports detailing research and any new findings which are then made available to the scientific community in scientific journals or at conferences for scrutiny and further debate;\n","normalized_statement":"write scientific papers and reports detailing research and any new findings which are then made available to the scientific community in scientific journals or at conferences for scrutiny and further debate","leading_verb_lemma":"write","verb_prior":"easy","coordinate_verbs":["write"],"dominant_activity":"write research papers communicating novel findings","exposure":"EXPOSED","eloundou_label":"E2","exposure_confidence":4,"exposure_channel":"complementarity","exposure_reason":"LLM recheck: Eloundou E2 (B.2) — drafting tool augments, but novel-findings communication needs the author","learnability":"HARD","easy_probability":0.38,"metric_objective":false,"mapping_simple":false,"learnability_reason":"LLM recheck: no objective metric for communicating novel findings (0.6); write-trap confirmed HARD","overall_confidence":4,"flag_for_review":false,"review_reason":"none","source":"LLM(recheck-flagged)","reasoning_trace":"Rule engine flagged write_trap → hybrid routes to LLM; LLM confirms E2/HARD 0.38 and resolves the flag."}]
```

---

9. FINAL INSTRUCTIONS

- **Set `method` first** and apply exactly one of PART A / B / C. Always run STEP 1 preprocessing (§0.3), decide **exposure first**, score learnability **only** if EXPOSED, never let one axis infer the other.
- **Process sequentially, one `Input/` part per run, in order `part01 → part15` (§7.4)** — never load multiple parts at once, never classify the whole file in one pass; each part's emitted JSON array length must equal its manifest record count.
- **Consolidate into ONE workbook → `Output/KBJI_AI_Exposure_ALL_tasks_classified.xlsx` (§7.5)** holding all three methods as sheets (+ `Method_Comparison`, `Group_Summary`); optional single CSV companion in `Output/`. Concatenate the 15 parts in order to exactly 3,244 rows per method; never leave the output as 15 separate files; validate row count, unique key, and 6 blank rows before finalising.
- **Use only project-root-relative paths (§0.7)** — `Input/…`, `Output/…`, `Analysis/…`. Never hardcode an absolute or machine-specific path, so the pipeline replicates unchanged on any laptop.
```

```
- **Method A (rule engine):** behave as a deterministic lookup — verb-class table (§A.1) + physical gate (§A.2) + object override (§A.3) + fixed learnability mapping (§A.4). Same input → same output. No free reasoning.
- **Method B (LLM):** label Axis-1 with Eloundou's **E0/E1/E2/E3** taxonomy (§B.2) under the worker framing (§B.1), map to the schema (§B.3), record `eloundou_label`, then judge Axis-2 with Acemoglu's two-condition test (§0.6).
- **Method C (hybrid):** rule-engine label everywhere; **replace with the LLM label wherever the rule engine set `flag_for_review=true`** (§C.1); always record `source`.
- **Physical / on-site / in-person / bodily = NOT_EXPOSED (E0)** for software AI in every method (§0.4, E7/E8/E9/E11), even when routine (robot ≠ software-AI).
- Use the leading verb as a **prior only**; let object/domain override it (E2/E3/E4). Judge learnability on **objective metric + simple mapping only** — never on human difficulty (E14).
- **Calibrate** `easy_probability` to the **0.98 / 0.67-lean-hard / 0.098** anchors; apply the single write-trap convention; **flag** the 0.40–0.59 band, unresolved uncertain verbs, mixed compounds, write-traps, physical-embedded edges, frontier-borderline calls, and blanks.
- **Emit `isco_08` as an integer**; copy `occupational_title` and `task_statement` verbatim; emit `easy_probability=null` (never 0.0) for non-exposed tasks; populate `eloundou_label` for LLM/hybrid-LLM rows and `null` otherwise; keep every reason field to one line; `reasoning_trace` ≤ 45 words.
- All Acemoglu/Eloundou US magnitudes (~19.9% exposed, ~74% easy, 27% vs 7% savings, ≤0.55% TFP, α/β/ζ shares) are **speculative** — transfer the **criteria and decision rules, NOT the numbers**; Indonesia's low-wage adoption drag is a downstream IMPACT caveat, not a relabeling rule.
- Output **only** the JSON array — one object per input task, in input order. No commentary, no markdown fences, no trailing text.
```