AI and Employment

This lesson examines the multifaceted relationship between artificial intelligence and employment. It covers labor market impacts, future work scenarios, ethical concerns, algorithmic management in gig economies, emotional AI and surveillance, displacement versus transformation debates, and worker responses including unions and regulation.

Ghassem Tofighi

Contents

Overview 3
Learning Objectives 3
Motivation 3
The Impact of AI on Various Labour Markets 3
- Task-Level Determinants of Exposure 3
- Sectoral and Occupational Patterns 4
- Manufacturing 4
- Logistics and Transportation 4
- Retail and E-commerce 5
- Finance and Insurance 5
- Healthcare 5
- Legal Services 6
- Media and Content Production 6
- Cross-cutting Occupational Shifts 7

The Future of Work in the AI Era 7
- Technological unemployment 7
- Augmentation and complementarity 7
- Polarization and contingent work 7
- Structural Trends Already Observable 8

Ethical Issues Related to AI and Employment 8
- Principal Ethical Conflicts 8
- Illustrative Ethical Cases 9

Algorithmic Management and Gig Work 9
- Functional Components 9
- Worker Conditions in Platform-Mediated Labor 10
- Platform-specific Practices 10

Emotional AI and Workplace Surveillance 10
- Sensing and Inference Modalities 10
- Deployment Contexts 11
- Documented Outcomes 11

Labour Displacement vs Job Transformation 11
- Displacement Pathways 11
- Transformation / Augmentation Pathways 12
- Transitional Implications 12

Worker Resistance, Unions, and Regulation of Algorithmic Management 12
- Modes of Worker Response 12
- Regulatory Directions 13
- Concrete Regulatory Examples 13

Summary 13

Overview

This lesson examines structural transformations in employment relations and labour markets caused by widespread artificial intelligence integration. It covers differential exposure and impact across task types, occupations, sectors and firm sizes. It compares alternative long-term scenarios of work organization. It identifies major ethical conflicts generated by AI-mediated decision processes. It describes the technical architecture, operational logic and lived consequences of algorithmic management in platform economies. It maps modalities, sensing technologies and psychosocial-organizational effects of emotional AI in workplaces. It differentiates mechanisms, evidence base and transitional costs of direct labour displacement versus job augmentation and transformation. It documents evolving patterns of individual and collective worker resistance together with the scope, instruments and limitations of regulatory and policy responses to algorithmic control of labour.

Learning Objectives

- Differentiate AI exposure levels by task codifiability, predictability, data structure, environmental variability, social interaction, and ethical judgment requirements.
- Compare optimistic, pessimistic, and polarized scenarios of AI's employment impact.
- Identify mechanisms of bias, opacity, surveillance, power concentration, and unequal benefit distribution in AI employment systems.
- Describe the architecture and effects of algorithmic management on worker autonomy, income, and subordination.
- Distinguish sensing modalities of emotional AI and their psychosocial and organizational consequences.
- Compare pathways, evidence, and costs of labour displacement versus job augmentation.
- Assess individual, collective, and regulatory responses to algorithmic management.

Motivation

The accelerating incorporation of machine learning models, computer vision, natural language processing, predictive analytics and automated decision systems into private-sector production, service delivery, public administration and platform-mediated work simultaneously removes certain tasks from human execution, augments the productivity of others, reconfigures supervisory and disciplinary authority, intensifies performance monitoring, alters skill obsolescence rates, modifies occupational demand structures, and redistributes economic returns between labour and capital. Rigorous analysis of these concurrent processes is necessary to understand the direction and magnitude of labour-market restructuring, to anticipate pressure points in skill formation systems and social-protection arrangements, to evaluate competing claims about aggregate employment levels and inequality trajectories, and to identify institutional design choices that can channel technological productivity gains toward broadly shared improvements in working conditions, income security and occupational opportunity rather than toward greater precarity and polarization.

The Impact of AI on Various Labour Markets

Task-Level Determinants of Exposure

The probability and speed of AI adoption at the task level depend on several measurable characteristics. Tasks that are highly predictable, follow explicit rules or statistical patterns, operate on fully structured

digital inputs, require minimal adaptation to novel physical environments, and involve limited interpersonal negotiation or ethical discretion exhibit the highest exposure.

Example: Scanning barcodes on a conveyor belt and matching them against a database is highly predictable and structured.

Tasks that demand fine-grained physical dexterity in unpredictable settings, real-time tacit coordination with other humans, creative synthesis of disparate information sources, or morally complex trade-off judgments remain substantially more resistant to full automation in the current technological paradigm.

Example: A nurse repositioning an immobile patient while simultaneously assessing skin integrity and responding to subtle signs of discomfort combines physical skill, situational awareness and empathy.

Between these poles lie tasks amenable to partial automation through pattern recognition, classification, recommendation or generation within relatively well-defined domains.

Example: Suggesting the next best action in a customer-support chat based on previous messages and product database falls in this intermediate category.

Sectoral and Occupational Patterns

Manufacturing

High-exposure tasks include repetitive component assembly on fixed lines, automated visual quality inspection of standardized parts, robotic packaging of uniform products, and CNC machining of parts with pre-programmed tool paths.

Example: Placing identical screws into pre-drilled holes on a circuit board assembly line.

Medium-exposure tasks encompass routine machine setup for batch production, basic process-parameter adjustment based on sensor feedback, and predictive maintenance scheduling using historical failure data.

Example: Recalibrating a welding robot after a batch change using pre-set parameters and vibration sensor readings.

Low-exposure tasks involve custom one-off fabrication, on-site repair of legacy equipment with unknown failure modes, complex root-cause troubleshooting across mechanical, electrical and software subsystems, and real-time adaptation to unexpected production variances.

Example: Diagnosing why a legacy hydraulic press intermittently fails to reach pressure when the error code is ambiguous and no recent maintenance log exists.

Logistics and Transportation

High-exposure tasks comprise centralized route optimization for delivery fleets, automated parcel sorting in high-throughput warehouses, predictable point-to-point driving on highways with clear lane markings, and inventory tracking via RFID and computer vision.

Example: Sorting packages by destination zip code using conveyor belts and overhead cameras in a distribution centre.

Medium-exposure tasks include semi-structured warehouse picking in environments with partial automation, last-mile coordination that combines algorithmic dispatch with human driver input, and basic vehicle diagnostics using telematics data.

Example: A driver receiving a re-optimized route mid-trip after a traffic jam is reported.

Low-exposure tasks cover long-haul driving through dense urban traffic or adverse weather, handling of hazardous or fragile materials requiring special care, customer-facing delivery interactions, and emergency roadside interventions.

Example: Delivering medical oxygen tanks to a residential address during a snowstorm while communicating with a home-care patient.

Retail and E-commerce

High-exposure tasks involve self-checkout systems, continuous inventory monitoring through shelf cameras and weight sensors, automated basic customer-service responses via chat interfaces, and routine pricing adjustments based on demand signals.

Example: A shopper scanning groceries at a self-checkout terminal with computer-vision-assisted theft detection.

Medium-exposure tasks include algorithmic product recommendations, dynamic assortment planning, and fulfilment-centre order batching.

Example: Displaying "customers also bought" suggestions on an online clothing website.

Low-exposure tasks encompass personalized in-store styling advice, resolution of emotionally charged or legally complex customer complaints, and management of high-value or bespoke merchandise.

Example: Helping a customer select a wedding dress that matches cultural traditions and personal body-image preferences.

Finance and Insurance

High-exposure tasks cover routine credit-worthiness scoring using structured data, automated claims processing for low-complexity incidents, know-your-customer identity verification, basic compliance checks against regulatory rule sets, and high-frequency algorithmic trading execution.

Example: Approving a small personal loan based solely on credit-bureau data and income verification.

Medium-exposure tasks include anomaly detection in transaction streams, portfolio rebalancing under defined risk constraints, and preliminary underwriting assessments.

Example: Flagging a credit-card transaction as potentially fraudulent because it deviates from the cardholder's usual spending pattern.

Low-exposure tasks involve relationship management with high-net-worth clients, interpretation of ambiguous or rapidly changing regulatory environments, and strategic negotiation of large syndicated loans or insurance placements.

Example: Structuring a $500 million syndicated loan facility for a multinational corporation during a regulatory transition period.

Healthcare

High-exposure tasks include medical coding and billing classification, preliminary screening of radiological images for obvious abnormalities, appointment scheduling optimization, and transcription of structured clinical notes.

Example: Assigning ICD-10 codes to a discharge summary that describes a straightforward appendectomy.

Medium-exposure tasks encompass diagnostic suggestion engines, evidence-based treatment-pathway recommendations, and triage prioritization in emergency departments.

Example: Suggesting possible differential diagnoses after a patient presents with chest pain and abnormal ECG readings.

Low-exposure tasks cover integration of ambiguous clinical findings into holistic diagnoses, shared decision-making with patients who have complex values or comorbidities, hands-on surgical procedures, and end-of-life care discussions.

Example: Discussing palliative care options with the family of a patient with advanced metastatic cancer who has conflicting treatment preferences among relatives.

Legal Services

High-exposure tasks consist of electronic discovery document review, extraction of standard clauses from contracts, basic due-diligence checks against public records, and summarization of large case-law corpora.

Example: Identifying all documents containing the phrase “indemnification” in a 500,000-document merger discovery set.

Medium-exposure tasks include initial drafting of routine agreements, legal research memo generation, and predictive case-outcome modelling.

Example: Producing a first draft of a non-disclosure agreement using a template and party-specific information.

Low-exposure tasks encompass high-stakes litigation strategy, multi-party negotiation, courtroom advocacy requiring real-time adaptation to judicial temperament, and advising on novel regulatory frontiers.

Example: Cross-examining an expert witness in a patent-infringement trial when the judge interrupts with unexpected questions.

Media and Content Production

High-exposure tasks cover routine copyediting of formulaic text, automated caption and subtitle generation, selection and basic editing of stock footage or images, and production of standardized news summaries from structured data feeds.

Example: Generating captions for a weather forecast video using speech-to-text and template formatting.

Medium-exposure tasks include initial article or script drafting from prompts, automated video highlight reel creation, and basic graphic-layout suggestions.

Example: Producing a first draft of a company earnings-report press release from financial data tables.

Low-exposure tasks involve original investigative journalism, long-form narrative construction, editorial judgment on controversial or culturally sensitive material, and creative direction of visual storytelling.

Example: Writing an in-depth feature on labour conditions in a global supply chain that requires confidential source protection and cross-verification.

Cross-cutting Occupational Shifts

Routine middle-skill cognitive occupations (data-entry clerks, basic bookkeeping, paralegal document review) and routine manual occupations (assembly-line work, warehouse picking, basic sorting) have experienced the most significant contraction.

Example: Replacement of manual invoice processing clerks with optical character recognition and rule-based extraction systems.

Relative employment growth concentrates at two poles: high-skill analytical, creative, strategic, and interpersonal roles that require advanced judgment, originality, or human connection, and low-skill in-person service roles that demand physical presence and adaptability in variable social contexts.

Example: Increase in demand for AI ethics consultants and home health-care aides.

Demand increases sharply for hybrid occupations that combine deep domain knowledge with the ability to critically evaluate, supervise, fine-tune, and audit AI system outputs.

Example: Clinical decision-support validators who assess whether an AI radiology tool misclassifies edge cases.

The Future of Work in the AI Era

Technological unemployment

Rapid advances in general-purpose AI and robotics enable substitution across a broad range of cognitive and manual tasks. Displacement occurs faster than new labour-absorbing activities can emerge. Aggregate demand contracts due to insufficient wage income among displaced workers. Sustained structural unemployment rises and remains elevated unless counteracted by large-scale redistribution mechanisms.

Example: Widespread replacement of truck drivers by autonomous vehicles leading to persistent joblessness in logistics-dependent regions without sufficient new employment creation.

Augmentation and complementarity

AI systems automate narrow, well-defined sub-tasks within broader occupations. Human labour shifts toward activities that require contextual understanding, creative synthesis, interpersonal trust, moral reasoning, and adaptive problem-solving in ill-structured environments. Productivity gains increase real incomes, expand consumption possibilities, and generate demand for new goods, services and cultural products. Net employment grows, although the composition of jobs changes significantly.

Example: Software developers using AI code assistants to produce higher-quality code faster, leading to expansion of software-intensive industries and new roles in system integration and user-experience design.

Polarization and contingent work

Middle-skill routine jobs vanish across both cognitive and manual domains. Labour demand concentrates at the high-skill end (requiring advanced education and continuous learning) and the low-skill end (requiring physical presence and direct human interaction). Earnings inequality widens markedly. Platform-

mediated, short-term, and on-demand work arrangements expand, reducing access to stable benefits, career ladders, and collective representation.

Example: Growth of high-paid AI specialists alongside expansion of precarious gig roles in food delivery and ride-hailing with stagnant real wages.

Structural Trends Already Observable

Platform-mediated gig and contingent work grows rapidly in transportation, delivery, microtasking, content moderation and freelance professional services.

Example: Increase in the number of workers registered on multiple food-delivery and ride-hailing apps simultaneously.

Average job tenure within single employers shortens, particularly among younger cohorts.

Example: Technology workers changing employers every 1.5–2 years on average.

Skill obsolescence accelerates in mid-career technical and administrative occupations.

Example: Legacy COBOL programmers facing reduced demand as financial institutions migrate to modern languages.

Continuous learning shifts from episodic employer-sponsored training toward individual responsibility and market-mediated credentialing platforms.

Example: Widespread use of online platforms offering micro-credentials in machine-learning operations.

Policy experimentation includes trials of shorter statutory working weeks, time-limited income-support schemes, and portable benefits tied to workers rather than employers.

Example: Four-day workweek pilots in several European countries showing maintained output with reduced hours.

Ethical Issues Related to AI and Employment

Principal Ethical Conflicts

Training datasets drawn from historical human decisions frequently encode patterns of discrimination based on gender, race, ethnicity, age, disability or socioeconomic origin; machine-learning models propagate and often amplify these patterns in automated screening, scoring and ranking.

Example: Resume-screening tools downgrading candidates who attended historically Black colleges due to correlation patterns in past hiring data.

Individuals subject to significant automated decisions frequently receive no meaningful, accessible explanation of the criteria, data sources, or model logic that determined outcomes affecting their livelihood.

Example: A driver deactivated from a ride-hailing platform without knowing which specific metric or incident triggered the decision.

Continuous, multi-modal monitoring of location, communication, keystroke dynamics, webcam feeds and physiological signals shrinks the sphere of worker autonomy and personal dignity.

Example: Warehouse workers required to wear devices that track every movement and rest break.

Algorithmic management systems concentrate effective authority in platform owners and opaque code while diffusing responsibility across layers of software, data providers and automated rules.

Example: No single person can be held accountable when an algorithm reduces a worker's pay through dynamic pricing changes.

Productivity and efficiency gains generated by AI deployment are disproportionately captured by capital owners, platform intermediaries and high-skill technical roles rather than distributed across the broader labour force.

Example: Shareholder returns increasing substantially while median worker wages stagnate despite productivity growth.

Illustrative Ethical Cases

Recruitment systems trained on past hiring patterns systematically disadvantage applicants whose resumes omit elite-university attendance, specific extracurricular keywords or certain employment-gap explanations that correlate with protected characteristics.

Example: An applicant with caregiving responsibilities penalized for resume gaps.

Performance-evaluation algorithms penalize delivery drivers for taking short breaks to manage chronic health conditions or penalize customer-service agents for longer call durations when resolving complex emotional complaints.

Example: A driver receiving lower priority dispatch after taking necessary restroom breaks during a long shift.

Automated disciplinary systems deactivate workers based on aggregate behavioural scores without disclosing the relative weight of individual metrics or allowing effective appeal.

Example: A delivery courier permanently removed from a platform after receiving several low ratings from customers unhappy with delivery timing during bad weather.

Mandatory deployment of emotion-recognition software in remote-work settings links inferred frustration or boredom scores to productivity evaluations and bonus calculations.

Example: A remote call-centre agent receiving lower performance ratings due to detected neutral facial expressions during routine calls.

Algorithmic Management and Gig Work

Functional Components

Automated task allocation matches workers to jobs using real-time location, availability, historical ratings, predicted demand, and estimated completion time.

Example: A ride-hailing platform assigns the nearest driver with a high acceptance rate to a new ride request.

Dynamic pricing adjusts compensation according to instantaneous supply-demand conditions, time of day, weather, and local events.

Example: Surge pricing increases pay rates during a major concert ending.

Multi-dimensional performance scoring combines customer ratings, acceptance rates, cancellation rates, speed metrics, and behavioural telemetry.

Example: A delivery platform calculates a composite score from five-star ratings, on-time percentage, and low cancellation rate.

Automated enforcement applies warnings, temporary suspensions, or permanent deactivations when thresholds are crossed.

Example: Automatic deactivation after falling below a 4.6 average rating for 30 days.

Behavioural nudging uses differential incentives to influence worker choices.

Example: Offering priority dispatch to drivers who maintain 95% acceptance rate.

Worker Conditions in Platform-Mediated Labor

Near-complete opacity exists regarding weights, thresholds, and logic of scoring algorithms.

Example: Workers cannot see how much a single low rating affects their overall score.

Limited or no effective recourse exists for contesting penalties or deactivations.

Example: Appeal processes rarely reverse decisions even with evidence of external factors.

Pressure exists to accept low-value tasks to preserve platform access and future dispatch priority.

Example: Accepting distant, low-paying deliveries to avoid being ranked lower.

Income exhibits high volatility from unpredictable surge patterns and algorithmic re-ranking.

Example: Earnings vary 40–60% week-to-week depending on algorithmically determined demand.

Effective subordination occurs without standard employment protections, benefits, or collective bargaining rights.

Example: No guaranteed minimum wage or paid sick leave despite full-time hours.

Platform-specific Practices

Ride-hailing platforms deactivate drivers below hidden star-rating or acceptance-rate thresholds.

Example: Deactivation triggered at 4.6 stars or below 85% acceptance rate.

Delivery platforms prioritize dispatch for couriers with near-perfect on-time completion and low cancellations.

Example: High-performing couriers receive first access to high-tip orders.

Microtask platforms sort available work into reputation-tiered queues.

Example: Workers with 98%+ approval rate see higher-paying HITs first.

Emotional AI and Workplace Surveillance

Sensing and Inference Modalities

Video-based facial action unit detection processes webcam feeds to identify micro-expressions, brow position, lip curvature, and eye changes.

Example: Software detects furrowed brows and downward mouth corners as indicators of frustration.

Audio analysis extracts prosody features including pitch variability, speech rate, amplitude, and pause patterns.

Example: Increased speech rate and higher pitch detected as signs of stress during calls.

Keystroke and mouse dynamics capture typing rhythm, error rates, click frequency, and movement trajectories.

Example: Sudden increase in typing errors and slower mouse movements flagged as fatigue.

Wearable or ambient sensors measure heart-rate variability, galvanic skin response, respiratory rate, and posture.

Example: Elevated heart rate and reduced heart-rate variability interpreted as anxiety.

Deployment Contexts

Contact-centre platforms flag negative sentiment or stress for supervisor intervention.

Example: Real-time alert sent when a customer-service agent shows rising voice pitch and negative facial cues.

Remote-work monitoring suites infer engagement or frustration from combined signals.

Example: Dashboard shows “low engagement” when facial cues indicate distraction and typing slows.

Productivity dashboards correlate inferred affective states with output metrics.

Example: Linking detected boredom scores to periods of lower keystroke activity.

Documented Outcomes

Chronic stress activation occurs from awareness of continuous affective monitoring.

Example: Elevated cortisol levels reported among monitored remote workers.

Self-censorship of natural emotional expression emerges to avoid negative flagging.

Example: Workers deliberately maintaining neutral facial expressions regardless of actual mood.

Erosion of trust and psychological safety develops within teams.

Example: Employees hesitate to express frustration in meetings knowing it may be logged.

Risk of misclassification arises due to cultural variation in emotional display or neurodiverse presentation.

Example: Reserved communication style misinterpreted as disengagement.

Narrow normative definition of acceptable workplace affect becomes linked to performance evaluation.

Example: Only calm and positive expressions rewarded with higher scores.

Labour Displacement vs Job Transformation

Displacement Pathways

Full substitution occurs when AI or robotics replaces human execution of an entire task sequence or occupation.

Example: Self-service kiosks reduce the number of airline counter staff required at check-in.

Example: Automated high-volume document classification eliminates entry-level paralegal review positions.

Example: Robotic case-picking systems decrease manual warehouse-selector roles.

Displacement tends to be rapid once cost-effective reliability is achieved and is often localized within specific firms or regions.

Transformation / Augmentation Pathways

AI removes routine, time-intensive or error-prone sub-tasks, allowing humans to focus on higher-complexity judgment, creative synthesis, interpersonal coordination, or exception handling.

Example: AI-supported diagnostic imaging enables radiologists to review more cases while concentrating on ambiguous findings.

Example: Code-completion and refactoring assistants allow developers to allocate greater effort to architecture and novel problem-solving.

Example: Automated legal-research engines free lawyers to devote more time to client counselling and negotiation strategy.

Transformation usually proceeds gradually and remains embedded within existing occupational identities.

Transitional Implications

Direct displacement frequently generates sharp, localized job loss, significant wage depreciation, and high retraining costs when switching occupations.

Example: Former truck drivers re-entering the labour market at lower wages in unrelated sectors.

Job transformation generally involves incremental skill deepening, lateral movement into adjacent roles, or upward progression within evolving occupations.

Example: Radiologists upskilling in AI tool validation and complex case interpretation.

The incidence of underemployment, wage scarring or successful recovery depends on the quality, accessibility, and timeliness of retraining, active labour-market policies, and employer redeployment incentives.

Example: Workers with access to subsidized retraining programs show higher rates of re-employment at similar wage levels.

Worker Resistance, Unions, and Regulation of Algorithmic Management

Modes of Worker Response

Individual-level tactics include selective task acceptance, multi-platform work, profile optimization, and deliberate rating inflation.

Example: Drivers accepting only high-paying long rides to improve earnings while maintaining high ratings.

Informal collective action encompasses coordinated slowdowns, media exposure of opaque practices, mutual-aid groups, and workaround sharing.

Example: Delivery workers simultaneously refusing low-paying orders during peak hours to force rate increases.

Formal organizing efforts involve platform-worker unions, worker associations, cooperatives, and strategic litigation on misclassification.

Example: Union campaigns demanding algorithmic transparency and collective bargaining rights for gig workers.

Regulatory Directions

Legislation extends employment classification and protections to platform-mediated labour.

Example: Laws granting minimum wage, overtime, and sick leave to reclassified platform workers.

Transparency of algorithmic decision criteria, weights, and thresholds is mandated.

Example: Platforms required to disclose how performance scores are calculated.

Rights to explanation and human review of adverse automated decisions are established.

Example: Workers granted written explanations and appeal hearings before deactivation.

Limits are placed on continuous behavioural, location-based, and emotional monitoring.

Example: Bans on mandatory real-time emotion tracking without explicit consent.

Requirements exist for worker consultation on system design, updates, and performance metrics.

Example: Mandated consultation with worker representatives before introducing new scoring algorithms.

Concrete Regulatory Examples

Platforms must provide advance notification and meaningful appeal processes before deactivation or suspension.

Example: 14-day notice and independent review required before permanent account closure.

Workers gain enforceable rights to download comprehensive performance data logs and scoring logic.

Example: Access to 12 months of raw ratings, acceptance rates, and derived scores.

Restrictions are placed on use of opaque behavioural proxies as sole determinants of compensation or access.

Example: Acceptance rate cannot be the only factor determining dispatch priority.

Summary

Artificial intelligence systems simultaneously automate routine tasks, augment complex work, intensify algorithmic control over labour processes, expand surveillance into behavioural and affective domains, accelerate skill obsolescence in mid-range occupations, and reshape the distribution of economic value between labour and capital. Possible trajectories span productivity-led job evolution with broadly shared gains, persistent structural unemployment, and deepening polarization characterized by precarious

platform work and rising inequality. Outcomes depend on institutional choices concerning continuous education and retraining ecosystems, worker representation and voice mechanisms, transparency and contestability obligations for automated decisions, constraints on surveillance intensity, classification of platform labour, and fiscal and regulatory instruments that influence the allocation of productivity dividends. Sustained focus on fairness, autonomy, explainability, human oversight, and equitable transition support remains essential to orient technological change toward inclusive and democratically governed labour-market outcomes.