

ANALYSIS PROTOCOL

Environmental Scan and Pro-Worker Classification of AI-Enabled Health Technologies in the United States

Protocol Version 3.0, March 2026

Blythe Adamson, PhD, MPH

Document Purpose. This protocol pre-specifies the procedures for a structured environmental scan of AI-enabled health technologies across the US healthcare system. The analysis proceeds in three phases: (1) **Data Collection**—factual extraction of technology characteristics requiring no expert judgment; (2) **Analysis of Labor Market Consequences**—structured expert assessment of each technology’s effects on labor productivity, the value of human expertise, and labor’s share of income, following the dimensions of Acemoglu, Autor, and Johnson (2026); and (3) **Synthesis**—aggregation of per-technology assessments into tables, figures, and summary findings suitable for peer-reviewed publication.

1. STUDY DESIGN AND OBJECTIVES

1.1 Study Design

Cross-sectional structured environmental scan with expert classification. Observational and descriptive. Not a systematic review; uses a convergent multi-source search strategy rather than PRISMA-compliant exhaustive database screening.

1.2 Primary Objective

To identify and classify AI-enabled health technologies using the five-category framework of Acemoglu, Autor, and Johnson (2026), assessing each technology’s predicted implications for

labor productivity, the value of human expertise, labor’s share of national income, labor demand, and wages.

1.3 Secondary Objectives

(a) To characterize the distribution of AI technologies across the five categories and across sectors of the healthcare system. (b) To conduct a pre-specified subanalysis of the administrative friction-cost layer—defined as healthcare functions that generate system friction rather than directly producing health—examining whether its category distribution differs systematically from clinical care delivery. (c) To identify occupational categories facing concentrated displacement risk or demand for new expertise.

1.4 Analytical Framework

All classifications use the five-category framework from Acemoglu, D., Autor, D., & Johnson, S. (2026), “Building Pro-Worker Artificial Intelligence,” The Hamilton Project, Brookings Institution. The framework distinguishes: (1) labor-augmenting, (2) capital-augmenting, (3) automating, (4) new task-creating, and (5) expertise-leveling technologies. Only Category 4 is unambiguously pro-worker. The analyst must read this paper in its entirety before beginning any classification work.

2. THREE-PHASE ANALYTICAL ARCHITECTURE

The analysis is structured as three sequential phases, each with distinct outputs and different skill requirements. This separation ensures that factual data extraction (Phase 1) is complete and verifiable before expert judgment is applied (Phase 2), and that per-technology assessments are stable before aggregation and synthesis (Phase 3). All three phases are recorded on a single spreadsheet tab, with 16 columns per technology: 7 data collection columns and 9 analysis columns. Visual column-header groupings distinguish the phases.

2.1 Phase 1: Data Collection

Factual extraction. No expert judgment required. Any trained RA can execute.

#	Column	Extraction Instructions
1	Company	Parent company or organization marketing the product in the US. Use the entity name as it appears on the product’s US-facing website.
2	Tool / Technology	Specific product or platform name. If a company has multiple distinct AI products with different functional mechanisms, each is a separate row.
3	Sector	Full sector label from the taxonomy in Section 4 (e.g., “Diagnostic imaging and radiology”).
4	Domain Code	Single letter: A, B, C, D, or E, from the taxonomy in Section 4.
5	Subdomain Code	Alphanumeric code from the taxonomy in Section 4 (e.g., A1, B3, C2, D1, E3).
6	Summary of Technology	2–4 sentences describing: (a) what the technology does; (b) deployment scale (number of sites, users, claims processed, or equivalent); (c) regulatory status (FDA clearance/authorization with date, if applicable). Factual, verifiable claims only. No evaluative language.
7	Source URL	Primary verifiable source. Prefer company product page or FDA database entry. Must be a clickable hyperlink. Verified accessible at the scan date.

Phase 1 completion criterion: Phase 1 is complete when all technologies identified through the search strategy (Section 3) that pass inclusion screening (Section 5) have been entered with all seven columns populated and source URLs verified. Phase 2 should not begin until Phase 1 is complete for the full dataset.

2.2 Phase 2: Analysis of Labor Market Consequences

Expert judgment required. Analyst must have read the AAJ (2026) paper in its entirety.

For each technology in the Phase 1 dataset, the analyst assesses the technology’s labor market consequences along the three dimensions that define the Acemoglu-Autor-Johnson framework (Table 1 of the paper): labor productivity, value of human expertise, and change in labor’s share of national income. The analyst then assigns the AAJ category via the decision tree (Section 6) and verifies consistency between the assigned category and the three-dimensional

rating pattern. Finally, the analyst provides plain-language assessments of estimated impacts on labor demand and wages.

Columns 8–13: Three-Dimensional Labor Market Assessment

These six columns (three rating–explanation pairs) correspond to the three rows of Table 1 in Acemoglu, Autor, and Johnson (2026). Together they produce a “Table 1 entry” for each technology in the scan.

#	Column	Rating Scale	Instructions
8	Labor Productivity	+, +/–, –, ≈ 0	Rate the technology’s expected effect on output per unit of labor input. Most AI technologies will receive “+” because they increase throughput or reduce time-per-task. Assign “+/–” if effects are heterogeneous across worker groups. Assign “ ≈ 0 ” only if the technology has no measurable productivity effect.
9	Labor Productivity — Explanation	Free text	1–2 sentences describing the mechanism by which the technology affects labor productivity. Focus on what changes in the production process.
10	Value of Human Expertise	+, +/–, –, ≈ 0	Rate the technology’s expected effect on the scarcity value of existing human expertise. “+” if expertise becomes more valuable/relevant. “–” if existing expertise is commodified or rendered obsolete. “+/–” if some workers’ expertise is enhanced while others’ is devalued (e.g., entrants benefit, incumbents lose). “ ≈ 0 ” if expertise value is unchanged.
11	Value of Expertise — Explanation	Free text	1–2 sentences describing the mechanism. Which expertise is made more or less valuable? Whose expertise?
12	Change in Labor’s Share of National Income	+, +/–, –, ≈ 0	Rate the expected effect on the allocation of tasks between labor and capital. “+” if more tasks are done by labor (new tasks created). “–” if fewer tasks are done by labor (tasks shifted to capital/software). “ ≈ 0 ” if no task reallocation occurs. “+/–” if ambiguous or offsetting.
13	Labor’s Share — Explanation	Free text	1–2 sentences describing the task reallocation mechanism. Are tasks moving from humans to software, or are new human tasks being created?

Column 14: AAJ Pro-Worker Category

#	Column	Instructions
14	AAJ Pro-Worker Category	Step A: Apply the classification decision tree (Section 6) based on the technology’s primary mechanism of labor market impact. Record the category: 1. Labor-Augmenting, 2. Capital-Augmenting, 3. Automation, 4. New Task-Creating, or 5. Expertise-Leveling. Step B: Verify consistency with the canonical rating pattern from Table 1 of the paper (see Section 2.3 below). If the ratings in Columns 8, 10, and 12 are inconsistent with the assigned category, the analyst must resolve the inconsistency by either (a) revising the ratings, (b) revising the category assignment, or (c) documenting the inconsistency as a substantive finding with explanation. Unresolved inconsistencies are flagged for senior reviewer adjudication.

Columns 15–16: Plain-Language Impact Assessments

#	Column	Instructions
15	Estimated Impact on Labor	Qualitative directional assessment of labor demand effects. Specify which occupational categories are affected (using SOC titles where identifiable). Distinguish displacement effects from creation effects. State the direction for each affected group. Cite the evidence tier (Section 8.2) supporting the assessment.
16	Estimated Impact on Wages	Qualitative directional assessment of wage effects. Specify affected occupations separately where heterogeneous effects are expected. Note tensions between worker groups (e.g., “positive for clinicians, negative for scribes”). Where possible, note whether the effect operates through supply-side (more workers can do the task) or demand-side (fewer tasks require humans) channels.

2.3 Canonical Rating Patterns for Consistency Verification

After assigning the AAJ category (Column 14, Step A), verify that the ratings in Columns 8, 10, and 12 are consistent with the canonical pattern from Table 1 of Acemoglu, Autor, and Johnson (2026):

AAJ Category	Col 8: Labor Productivity	Col 10: Value of Expertise	Col 12: Labor’s Share	Pro-Worker?
1. Labor-Augmenting	+	+/-	≈ 0	Ambiguous
2. Capital-Augmenting	+	+/-	≈ 0	Ambiguous
3. Automation	+	–	–	Not pro-worker
4. New Task-Creating	+	+	+	Pro-worker
5. Expertise-Leveling	+	+/-	+/-	Ambiguous

Source: Adapted from Acemoglu, Autor, and Johnson (2026), Table 1.

Inconsistency handling. If the analyst’s ratings deviate from the canonical pattern for the assigned category, this may indicate: (a) a classification error (revise the category); (b) a rating error (revise the rating); or (c) a technology that genuinely combines elements of multiple categories (document as a substantive finding). Option (c) is expected for some technologies and is analytically informative—the paper itself notes that “many real-world technologies combine elements of multiple categories.” Flagged inconsistencies should be reported in the methods section of any resulting publication.

2.4 Phase 3: Synthesis

Aggregation, visualization, and interpretation for publication.

Phase 3 produces the tables, figures, and narrative summaries required for peer-reviewed publication. All Phase 3 outputs are derived from the completed Phase 1–2 dataset. Phase 3 should not begin until Phase 2 is complete, including resolution of all flagged inconsistencies.

Required outputs:

- Table: Distribution of technologies across the five AAJ categories (count, percentage)
- Table: Distribution across the 24 subdomains (A1–E3)
- Table: Friction-cost layer (Domains B+C) vs. clinical delivery (Domain A): comparative category distribution
- Table: Occupational categories facing concentrated displacement risk, with SOC codes where identifiable
- Figure: Stacked or grouped bar chart of AAJ category distribution by domain (A–E)
- Narrative summary addressing the pre-specified analytical questions (Section 10.3)

- Sensitivity analysis: reclassify all flagged-inconsistent entries to their next-best alternative category; report whether primary findings are robust

3. SEARCH STRATEGY

3.1 Overview

Convergent multi-source strategy across six phases. No single source is sufficient. The union of all sources produces the final dataset.

3.2 Phase 1: Regulatory Database

Query the FDA AI/ML-Enabled Medical Devices database. Download the complete device list. Extract companies with multiple authorized devices or Breakthrough Device Designation. Primarily surfaces diagnostic, imaging, cardiology, neurology, and monitoring technologies.

3.3 Phase 2: Industry and Market Intelligence

Query venture capital databases for healthcare AI companies with significant US funding. Review annual industry analyst reports covering healthcare AI market sizing and vendor landscapes. Review healthcare IT trade publication compilations of AI, RCM, and digital health vendors. Review annual award and ranking lists for relevant categories.

3.4 Phase 3: Payer and Administrative Layer

Query insurance industry regulatory body surveys on AI adoption among health insurers. Search for academic analyses of AI in utilization review and prior authorization. Review healthcare financial management publications on AI in revenue cycle. Conduct targeted searches for companies operating in each of: claims processing, prior authorization, utilization management, payment integrity, fraud detection, risk adjustment, medical coding, patient billing, provider credentialing, pharmacy benefit management, and member services.

3.5 Phase 4: Life Sciences and Drug Discovery

Search academic databases and biotechnology trade press for AI drug discovery companies with clinical-stage assets or major pharmaceutical partnerships. Search for AI clinical trial and AI genomics/precision medicine platforms.

3.6 Phase 5: Academic Literature

Search PubMed and Google Scholar with combinations of AI/ML terms, healthcare terms, and labor market terms. Date range: 36 months prior to scan date. Focus on empirical studies of AI deployment effects on clinical workflows, staffing, and documentation.

3.7 Phase 6: Conference and Trade Press

Review proceedings and news from major healthcare conferences held within 12 months preceding the scan. Captures product launches and deployment milestones not yet in peer-reviewed literature.

4. SECTORAL TAXONOMY

4.1 Design Rationale

The system that produces health spans multiple NAICS sectors: clinical delivery (NAICS 62), health insurance (NAICS 524114), pharmacy benefit management (NAICS 524292), pharmaceutical manufacturing (NAICS 3254), scientific R&D (NAICS 5417), health IT (NAICS 5112), and outsourced billing/staffing (NAICS 561). This taxonomy is anchored in NAICS codes and the CMS National Health Expenditure Accounts but organized around the analytical question: *where in the health production system does AI affect labor?* The friction-cost layer (Domains B+C) is defined separately because these administrative functions generate system friction rather than directly producing health, and the labor market dynamics differ from clinical care delivery.

4.2 Taxonomy

Assign each technology to exactly one subdomain. Domains B and C constitute the friction-cost layer.

Domain A: Clinical Care Delivery

NAICS 621 (Ambulatory Health Care Services), 622 (Hospitals), 623 (Nursing and Residential Care)

Code	Subdomain	Scope
A1	Diagnostic imaging and radiology	AI for interpretation, triage, or enhancement of radiological images. Includes AI embedded in imaging hardware.
A2	Pathology and laboratory medicine	AI for digital pathology slide analysis, biomarker quantification, and laboratory test interpretation.
A3	Clinical documentation and EHR workflow	Ambient clinical documentation, AI scribes, note generation, in-basket management, and EHR-embedded workflow tools.
A4	Clinical decision support and predictive analytics	Sepsis prediction, deterioration alerts, treatment recommendation engines, and clinical risk models at point of care.
A5	Surgical and procedural care	Robotic-assisted surgery, AI-enhanced surgical planning, intraoperative imaging AI, and procedural guidance.
A6	Remote monitoring, telehealth, and virtual care	AI-powered virtual care platforms, remote patient monitoring with AI triage, AI-native primary care. Classified here because the displaced or augmented labor is clinical.
A7	Digital therapeutics and behavioral health	AI-powered DTx, AI-enabled physical therapy, AI care companions for mental health, and continuous measurement tools.

Domain B: Healthcare Administration — Provider Side [FRICTION-COST LAYER]

Administrative functions within NAICS 621/622; outsourced services in NAICS 561

Code	Subdomain	Scope
B1	Revenue cycle management and medical coding	AI for autonomous coding, CDI, claim generation/scrubbing, denial management, and end-to-end RCM platforms.
B2	Patient access, scheduling, and eligibility	AI for insurance eligibility verification, appointment scheduling optimization, and patient registration automation.
B3	Patient billing and collections	AI for patient financial engagement, personalized billing, AI voice agents for inquiries, cost estimation, and payment plans.

B4	Provider credentialing and enrollment	AI for primary source verification, license tracking, payer enrollment automation, and credentialing workflows.
B5	Hospital operations and capacity management	AI for surgical scheduling, patient flow, bed management, and operational resource allocation.

Domain C: Healthcare Administration — Payer Side [FRICTION-COST LAYER]

NAICS 524114 (Direct Health and Medical Insurance Carriers), 524292 (Third Party Administration)

Code	Subdomain	Scope
C1	Claims processing and adjudication	AI for automated claims adjudication, real-time provider-payer benefit queries, and straight-through processing.
C2	Prior authorization and utilization management	AI for automated PA decisioning, clinical review acceleration, medical necessity determination, and utilization review.
C3	Payment integrity and fraud detection	AI for claims audit, overpayment identification, fraud ring detection, upcoding detection, and pre-payment risk scoring.
C4	Risk adjustment and quality measurement	AI for HCC coding, risk score optimization, chart review automation, and quality measure reporting.
C5	Member services and benefits administration	AI chatbots and voice agents for member inquiries, benefits explanation, and plan navigation.
C6	Pharmacy benefit management	AI for medication prior authorization, formulary management automation, and PBM operational workflows.

Domain D: Life Sciences and Biomedical R&D

NAICS 3254 (Pharmaceutical Manufacturing), 5417 (Scientific R&D Services)

Code	Subdomain	Scope
D1	Drug discovery and computational chemistry	AI for target identification, molecule generation, protein structure prediction, and synthesis planning.
D2	Clinical trial design and recruitment	AI for patient-trial matching, synthetic control arms, trial protocol optimization, and site selection.
D3	Genomics and precision medicine	AI integrating clinical and genomic data for treatment selection, companion diagnostics, and precision oncology.
D4	Real-world evidence and health data platforms	AI-powered de-identified health data platforms, RWE generation, and health data analytics infrastructure.

Domain E: Public Health, Population Health, and Health Workforce

NAICS 621, 624 (Social Assistance), 923 (Government HR Programs), 561320 (Temporary Help)

Code	Subdomain	Scope
E1	Disease screening and surveillance	AI for population-level screening, point-of-care diagnostics by non-specialists, and public health surveillance.
E2	Population health analytics and SDOH	AI for SDOH identification, community health needs assessment, and social care referral coordination.
E3	Health workforce and staffing	AI for healthcare worker matching, staffing optimization, and workforce planning platforms.

5. INCLUSION AND EXCLUSION CRITERIA

5.1 Inclusion

All four: (a) **AI-core:** AI/ML/NLP/CV as core functional component; (b) **US-deployed:** active US deployment, FDA authorization, or defined US pathway; (c) **Commercially mature:** commercially deployed, FDA-authorized, or Phase II+ with commercialization pathway; (d) **Health system scope:** operates within Domains A–E.

5.2 Exclusion

Consumer wellness apps (without clinical evidence/FDA/clinical integration); general-purpose LLMs (as standalone products); pre-commercial research (without commercialization pathway, excepting drug discovery with Phase I+ assets); EHR systems per se (without distinct AI module); non-US-deployed technologies.

6. CLASSIFICATION DECISION TREE

Q1. Does this technology primarily replace human task execution with automated processes, rendering specific human expertise less necessary?

Indicators: Vendor language about replacing FTEs, autonomous execution, “turn labor into software,” documented headcount reductions, straight-through processing, ROI models based on FTE elimination.

YES → Category 3: Automation. NO → Q2.

Q2. Does this technology create demand for genuinely new human expertise or tasks that did not previously exist?

Indicators: Novel job titles, entirely new workflows, new data/analysis categories generating new work, new occupational specialties.

YES → Category 4: New Task-Creating. NO → Q3.

Q3. Does this technology enable less-specialized workers to perform tasks previously requiring advanced expertise?

Indicators: Task-shifting specialists → generalists, point-of-care tools bypassing referrals, democratizing specialist interpretation, lowering occupational entry barriers.

YES → Category 5: Expertise-Leveling. NO → Q4.

Q4. Does the efficiency gain derive primarily from improving capital equipment rather than enhancing worker capability?

YES → Category 2: Capital-Augmenting. NO → Category 1: Labor-Augmenting (residual).

6.1 Tie-Breaking Rules

Rule 1: Primary mechanism prevails, based on (a) regulatory filings, (b) company materials, (c) peer-reviewed evidence, in that order.

Rule 2: Vendor language explicitly describing FTE replacement is strong evidence for Category 3.

Rule 3 (1 vs. 5): *Who can now do the work?* Same workers, faster → 1. Different, less-specialized workers → 5.

Rule 4 (3 vs. 5): *Is a human still required for the primary task?* No human → 3. Human needed but less expert → 5.

Rule 5: Document genuine ambiguity; flag for adjudication.

7. QUALITY ASSURANCE

7.1 Inter-Rater Reliability

A second analyst independently classifies $\geq 20\%$ of technologies using only the decision tree and Column 6 summaries. Calculate Cohen's κ for the five-category classification. Minimum: $\kappa \geq 0.60$. If below, jointly review and re-classify.

7.2 Consistency Audit

After Phase 2 completion, audit the full dataset for: (a) Column 8/10/12 rating patterns inconsistent with Column 14 category assignments (per Section 2.3); (b) similar technologies receiving different classifications without documented reasoning; (c) empty subdomains that may indicate search gaps.

7.3 Source Verification

All Column 7 URLs verified accessible on scan date. Substitute most authoritative accessible alternative for inaccessible URLs.

8. EVIDENCE STANDARDS

8.1 Phase 1 Standard

All Phase 1 entries (Columns 1–7) must be factual and attributable to a retrieved source. No evaluative or speculative language.

8.2 Phase 2 Evidence Hierarchy

When assessing labor market consequences (Columns 8–16), prefer evidence in this order:

- Tier 1: Peer-reviewed empirical studies with quantified workforce effects
- Tier 2: Vendor-reported metrics with named institutional partners
- Tier 3: Industry surveys with disclosed sample sizes
- Tier 4: Vendor marketing claims (note as “vendor-reported, unverified”)
- Tier 5: Theoretical prediction from the AAJ framework based on category

9. SUPPLEMENTARY INSTRUCTIONS FOR LLM-BASED RESEARCH AGENTS

9.1 Factual Discipline. Every claim in Columns 1–7 must be attributable to a retrieved source. Do not fabricate names, funding, metrics, or FDA clearances. A smaller dataset with verified entries is preferable.

9.2 Classification Consistency. Re-read the five-category table and decision tree before classifying. Apply mechanically. Reconcile inconsistencies before proceeding.

9.3 Search Discipline. 3–5 distinct queries per search phase. Cross-reference after each search.

9.4 Human Review. LLM-generated dataset must be reviewed by a human domain expert: (a) verify 10+ source URLs; (b) independently classify 15+ technologies; (c) review labor/wage assessments for plausibility; (d) flag inaccurate product descriptions.

10. PRE-SPECIFIED ANALYSES

10.1 Primary Analysis

Distribution of technologies across the five AAJ categories (count, percentage).
Distribution across 24 subdomains.

10.2 Friction-Cost Layer Subanalysis

For Domains B+C combined: count and proportion of total; category distribution compared to Domain A; specific occupational categories (SOC codes where identifiable) facing displacement; any technologies classified as Category 4 within the friction-cost layer, with explanation.

10.3 Pre-Specified Analytical Questions

- Is AI deployment disproportionately concentrated in any single AAJ category?
- Does the friction-cost layer exhibit a systematically different category distribution than clinical care delivery?
- Which occupational categories face the most concentrated displacement risk? Which face the greatest demand for new expertise?
- What is the relationship between the pro-worker character of AI deployment and the wage level of the affected occupational category?