Technology and the Shifting Architecture of Occupational Skills

Nikhil George, Christophe Combemale, Ramayya Krishnan, Rahul Telang

November 2025

Abstract

Technological transformation during the 2010s reshaped U.S. skill demands, but workers experienced divergent outcomes. Understanding this heterogeneity requires identifying not just whether skills changed, but how they changed. We develop a measurement framework from job postings in 2010 and 2018, constructing aligned skill spaces for 514 occupations covering over 75% of U.S. employment. Two geometric dimensions (transformation magnitude and shifts in occupational distinctiveness) jointly reveal three distinct processes. Cognitive distance arises when new requirements conflict with existing expertise, creating heterogeneous adaptation. Scaffolded learning occurs when new capabilities build cumulatively on existing foundations, enabling collective upgrading. Simplification emerges when technology eliminates specialized requirements. Validating these interpretations against within-occupation wage dispersion, we find each mechanism produces its predicted pattern. Among high-wage occupations with substantial transformation, cognitive distance generates 4.8-percentage-point wage spread widening while scaffolded learning maintains stability. Among low-wage occupations with minimal transformation, simplification produces 4.2-percentage-point compression. The same geometric feature (declining distinctiveness) produces opposite wage effects depending on transformation context. The qualitative nature of skill change determines worker outcomes, not magnitude alone.

Independent Researcher. Email: nikhilgeorgect@gmail.com.
Carnegie Mellon University. Email: ccombema@andrew.cmu.edu.
Carnegie Mellon University. Email: rk2x@cmu.edu.
Carnegie Mellon University. Email: rtelang@andrew.cmu.edu.
We thank Gavin Wang for generous advice on the vacancy posting data. We are grateful to TCS whose gift to Carnegie Mellon University supports work on AI and the Future of Work. This work was also supported by generous grants from the Advanced Robotics Manufacturing Institute and Richard King Mellon Foundation. The analysis and conclusions are those of the authors.

1. Introduction

Technological transformation during the 2010s fundamentally reshaped the U.S. economy. Workers did not experience this change uniformly: some acquired new capabilities and advanced, others lost value as old expertise eroded, and some adapted collectively within their fields. Yet existing frameworks focus primarily on whether technology substitutes for or complements tasks within fixed occupational boundaries, missing the mechanisms that generate these divergent outcomes.

Task-based and exposure-index approaches explain aggregate wage polarization but cannot distinguish among transformation processes. These approaches link technology to tasks (Autor and Dorn, 2013; Acemoglu and Restrepo, 2019) or trace disruption from trade and automation (Autor, Dorn, and Hanson, 2015; Cortes, Jaimovich, and Siu, 2017; Webb, 2019; Felten, Raj, and Seamans, 2021). Recent work shows labor market structure matters: technology clusters in skill-adjacent occupations, constraining mobility within a given occupational structure (Fan, 2025). Yet how technology transforms occupational structure itself (which boundaries shift, how skill relationships reconfigure) remains unmeasured. Without observing how occupations move in skill space, both internally and relative to others, we cannot infer the underlying dynamics driving change or predict their impact on workers.

We introduce a new measurement infrastructure that overcomes three obstacles. First, aggregating noisy job postings requires filtering firm-specific terminology while retaining true occupational signals. Second, embeddings must capture functional similarity so that geometric distances represent genuine skill relationships. Third, comparing across time requires aligning independently trained embeddings without assuming skill persistence. Our two-stage neural architecture resolves these challenges using contrastive learning and Orthogonal Procrustes alignment on 21 stable occupations, reducing anchor distance error by 82%. The resulting dataset spans 191,919 job postings across 514 occupations covering 75% of U.S. employment.

This framework produces two diagnostic metrics: Skill Evolution Magnitude (SEM), the geometric distance between 2010 and 2018 embeddings, representing how much an occupation's skill profile changed; and shifts in distinctiveness rankings (ΔRANK), the change in relative uniqueness within the occupational structure, capturing whether an occupation became more or less geometrically distinct from neighboring occupations. Neither dimension alone suffices: many occupations show similar magnitudes of change. Measuring distinctiveness alongside magnitude enables characterizing transformation qualitatively—whether occupations specialized, converged, or simplified—distinctions invisible in magnitude alone.

Critically, neither metric alone determines economic outcomes. The same distinctiveness decline can signal boundary erosion toward cross-domain capabilities or simpli-

fication of specialized requirements. These are opposing processes requiring different worker responses. Only the full geometric constellation (transformation magnitude, distinctiveness shifts, and neighbor configurations) reveals which process operates.

Measuring transformation requires both dimensions because a single magnitude metric would obscure the qualitative character of skill evolution. An occupation can transform by deepening its specialized expertise, thereby reinforcing its unique professional boundaries, or by converging toward capabilities common to other professions, thereby eroding them. Measuring shifts in distinctiveness alongside magnitude is essential because the combination reveals qualitatively different transformation processes. Whether boundaries reinforced, converged, or eroded through simplification proves decisive for labor market outcomes. These represent fundamentally different economic processes invisible in a scalar measure of change. Jointly analyzing these metrics reveals three distinct processes of skill transformation, each defined by whether new capabilities proved conflicting, cumulative, or eliminative in relation to existing expertise.

Cognitive Distance arises when new skill requirements conflict with established expertise. High-wage occupations like Financial Managers experienced major skill shifts toward analytical and computational capabilities: quantitative modeling, predictive analytics, and data-driven decision-making. By 2018, Financial Managers' nearest neighbors in skill space included Economists and Statisticians—analytical professions requiring fundamentally different cognitive paradigms than traditional finance. Acquiring these capabilities required reconciling conflicting approaches: algorithmic reasoning versus relationship-based judgment, statistical modeling versus experiential assessment. Some workers successfully integrated both domains, commanding premiums for rare dual expertise. Others struggled with cognitive conflicts. This heterogeneous adaptation resulted in widened wage dispersion.

Scaffolded Learning arises when new capabilities build cumulatively on established expertise. Occupations such as Architects evolved through additive skill development. Emerging tools like Building Information Modeling (BIM) built directly upon core architectural knowledge. The occupation's neighbors remained within the built environment ecosystem (Landscape Architects, Civil Engineers), not converging toward generic analytical work. Every digital feature mapped to concepts architects already mastered: spatial relationships, structural systems, material properties. Learning BIM meant extending familiar mental models rather than integrating conflicting paradigms. Work structure proved indivisible: contemporary projects cannot accommodate practitioners using different systems, forcing collective adoption. This enabled collective upgrading with stable wage structures.

Simplification arises when technology eliminates specialized requirements. Low-wage occupations such as Library Technicians saw declining distinctiveness as automation elim-

inated specialized procedural tasks. Neighbors shifted from specialized roles (pharmacy technicians, office specialists) to generic service work (vehicle cleaners, dining attendants). Digital catalog systems automated classification procedures, barcode scanning eliminated transaction recording expertise, integrated software reduced routing specialization. The procedural expertise that would have generated experience premiums became less critical. Job posting language in 2018 emphasized nearly identical competencies as 2010, yet the occupation's position in skill space shifted dramatically. Their roles converged with generic service work, eroding wage premiums and compressing pay spreads.

Empirically, occupations undergoing cognitive distance transformations exhibit 4.8-percentage-point wage spread widening, scaffolded learning contexts maintain stable dispersion, and simplification experiences show 4.2-point compression. The validation lies in correctly predicting directional effects from geometric signatures: declining distinctiveness predicts widening when paired with high transformation magnitude and cross-domain neighbors (cognitive distance), but compression when paired with low transformation and shift to generic neighbors (simplification). This demonstrates both dimensions jointly determine outcomes—boundary erosion alone cannot predict direction without considering transformation magnitude and neighbor context. A critical finding: the same geometric feature (declining distinctiveness) produces opposite wage effects depending on transformation context—demonstrating that no single metric captures mechanism type. Matching our 514 occupations to Bureau of Labor Statistics wage data, we achieve coverage of 432 occupations representing 72% of U.S. employment, examining within-occupation wage spreads (P75-P25) from 2010 to 2018.

Our approach follows an arc from representation to validation. We first develop infrastructure using contrastive learning to construct skill embeddings where geometric relationships capture occupational similarity. We then query this representation to identify systematic geometric patterns across 514 occupations. From these measurements, we develop economic interpretations for these signatures, grounded in theory of human capital and occupational boundaries, and generate falsifiable predictions about their labor market consequences. The validation of these predictions (each mechanism predicts specific dispersion dynamics, and observed patterns align) confirms that our interpretations correctly identify genuine economic mechanisms.

This framework provides the first systematic method to diagnose distinct mechanisms of occupational transformation using language-based skill embeddings. The qualitative form of skill change (whether conflicting, cumulative, or eliminative) proves decisive in determining economic outcomes. Existing exposure indices cannot distinguish these mechanisms: an occupation with high AI exposure might face cognitive distance, scaffolding, or simplification. Exposure alone cannot reveal which. Task-based frameworks characterize occupational content but not relational dynamics (whether boundaries shift

or persist). Our geometric approach captures both internal transformation and external repositioning. Recognizing these mechanisms enables targeted policy responses: bridging support for cognitive distance, domain-specific training for scaffolded learning, and transition aid for workers displaced by simplification. Beyond policy, understanding transformation types informs firm adoption strategies and explains rising within-occupation inequality as technology creates heterogeneous adaptation pressures.

Our analysis establishes occupation-level patterns while leaving micro-mechanisms for future investigation. Section 2 positions our work within literature. Section 3 details measurement infrastructure. Section 4 develops the analytical framework. Section 5 presents empirical patterns and validation. Section 6 concludes with implications.

2. Related Literature

Our work builds on and extends four distinct literature streams that have each provided crucial insights into how technology affects labor markets, while leaving the measurement of occupational boundary dynamics as an unaddressed frontier.

Task-Based Approaches and Their Limitations

The foundational insight of the task-based approach recognized that technology's impact on labor operates heterogeneously through specific tasks that define occupational content. This framework successfully explained broad patterns of wage polarization by demonstrating how technology substitutes for routine tasks while complementing non-routine cognitive work, with recent extensions distinguishing displacement from reinstatement (Acemoglu and Restrepo, 2019) and identifying which specific skills drive outcomes (Deming, 2017; Autor and Dorn, 2013). Studies document occupation-specific disruption from trade shocks (Autor, Dorn, and Hanson, 2015) and routine job displacement (Cortes, Jaimovich, and Siu, 2017), with recent work developing exposure indices that predict technological impacts through patent-task overlap (Webb, 2019) and AI capability assessments (Felten, Raj, and Seamans, 2021). However, by design, these approaches focus on the internal composition of occupations—the tasks within existing boundaries—rather than how technology might reshape the boundaries themselves. While this literature excels at measuring whether an accountant's tasks are becoming more analytical or routine, it cannot capture whether the distinction between "accountants" and "financial analysts" is becoming sharper or more blurred.

Recent work demonstrates that labor market structure determines technology's distributional impact. Fan (2025) shows that automation clusters in skill-adjacent occupations, creating "mobility traps" where workers' most accessible transitions lead to equally threatened jobs, revealing that adjustment depends on distance-dependent transferability across

the occupational network. Bocquet (2024) demonstrates that the network structure of occupations, particularly "bridge occupations" linking distinct clusters, significantly affects reallocation speed following technology shocks, with slow reallocation creating substantial adjustment costs. Rio-Chanona et al. (2021) model how workers navigate an occupational mobility network in response to automation, finding that network structure plays a crucial role in determining unemployment levels, with occupations in particular areas having few transition opportunities. These studies advance our understanding of how occupational proximity and network position shape worker adjustment, taking the structure of skill relationships as given to analyze flows between occupations.

Our contribution is complementary: we study how technology transforms occupational structure itself. Rather than analyzing worker flows between occupations with fixed boundaries, we measure how those boundaries shift—whether requirements extend toward cross-domain capabilities, build on existing expertise, or erode through simplification. These boundary reconfigurations determine adaptation pressures for incumbent workers within their current occupations, operating through mechanisms distinct from the mobility constraints emphasized in network-based approaches. By measuring occupational boundary dynamics directly, we provide the foundation for understanding both internal pressures on workers and the changing landscape through which displaced workers must navigate.

Text-as-Data and Embedding Models in Economic Analysis

The emergence of large-scale text data has revolutionized the measurement of economic phenomena, enabling researchers to move beyond static surveys to dynamic, real-time observation of market demands (Gentzkow, Kelly, and Taddy, 2019). In labor economics, analysis of job posting data reveals systematic variation in employer skill demands and how requirements respond to labor market conditions (Modestino, Shoag, and Ballance, 2020). Vacancy text analysis demonstrates effects on matching efficiency (Marinescu and Wolthoff, 2020), while skill complementarities can be measured through posting co-occurrence patterns (Atalay et al., 2020).

Recent advances in representation learning (Dell, 2025) enable conversion of unstructured data into structured embeddings where geometric relationships capture economic similarity. These methods have demonstrated power across diverse applications: leveraging portfolio holdings to construct asset representations for pricing (Gabaix et al., 2025) and using foundation models to capture rich career history representations for wage decomposition (Vafa, Athey, and Blei, 2024). However, applications to labor markets have focused primarily on measuring skill inventories or content similarities within occupations rather than the relational architecture between economic categories. Our work addresses the challenge of measuring how the distinctiveness between occupations

evolves—the boundary dynamics that determine labor market competition.

Occupational Boundaries and Labor Market Dynamics

A parallel literature has long recognized that occupational boundaries play a crucial role in determining labor market outcomes, affecting everything from wage setting to worker mobility patterns. Occupational boundaries are reinforced by licensing requirements affecting 29% of U.S. workers (Kleiner and Krueger, 2013), while skill transferability governs transition costs between occupations (Gathmann and Schönberg, 2010; Lise and Postel-Vinay, 2020). Research on occupational mobility has demonstrated that transitions between occupations involve substantial human capital costs, while declining occupational tenure suggests that these boundaries may be becoming more porous (Kambourov and Manovskii, 2009). Theoretical work on monopsony power emphasizes how occupational specificity affects firms' wage-setting ability (Manning, 2003). Despite the theoretical importance of occupational boundaries, this literature has been limited to inferring boundary changes from their consequences—mobility patterns or wage differentials—rather than directly measuring the boundary erosion process itself.

Innovation Economics and Technological Reorganization

The innovation economics literature has long emphasized that the most significant impacts of new technologies often come not from direct task substitution, but from the reorganization of economic activity they enable. General purpose technologies fundamentally reshape production processes through complementary organizational innovations (Bresnahan and Trajtenberg, 1995), suggesting that technological change should affect not just the content of individual jobs but the entire architecture of how work is organized. Our framework provides the first systematic approach to measuring this reorganization at the occupational level, quantifying how technology reshapes the competitive landscape between related professions.

Addressing these measurement limitations requires moving from static taxonomies to direct observation of firms' revealed skill demands. Recent advances in representation learning (Dell, 2025) enable conversion of unstructured text into structured embeddings where geometric relationships capture economic similarity. Contrastive learning methods have proven effective for economic applications requiring nuanced similarity measurement. We adapt these techniques to construct occupational skill spaces, developing the first comprehensive framework for measuring both internal skill transformation and boundary dynamics simultaneously.

Our contribution bridges these literature streams by developing the measurement infrastructure necessary to observe occupational boundary dynamics directly. This enables

us to move beyond inferring boundary changes from their consequences to measuring the boundary evolution process itself, providing new insights into how technological transformation reshapes the fundamental architecture of labor market competition.

3. A New Engine for Measuring Skill Change

3.1. Data and Economic Motivation

The central empirical challenge of this paper is to move beyond indirect proxies and directly measure the changing skill content of occupations. This creates an unexplored opportunity for systematic occupational analysis, but one that requires solving three interconnected measurement challenges that reveal why existing approaches prove insufficient.

First is the fundamental challenge of quantifying occupational skill relationships in a way that reflects genuine economic costs. Understanding technological change requires measuring how similar or different occupations are in terms of their skill requirements, where similarity serves as a proxy for the real economic cost of worker transitions—the time, effort, and lost productivity involved in acquiring new expertise. The distance from Financial Analyst to Quantitative Analyst may seem small given their shared domain, yet it requires substantial investment in mathematical and programming skills. Conversely, the distance from Financial Analyst to Credit Analyst, while perhaps seeming larger due to different specializations, may actually involve highly transferable analytical capabilities.

Second, beyond measuring static relationships, we must characterize how occupations themselves evolve over time. Rather than relying on expert judgment about which occupations face technological disruption, we need systematic measures derived directly from job posting text that capture both the magnitude of internal skill reconfiguration and changes in occupational boundaries with related professions. This characterization must be sensitive to the emergence of new skill combinations and the obsolescence of established expertise.

Third, any temporal comparison requires solving a fundamental measurement problem: ensuring that changes we observe reflect genuine occupational transformation rather than methodological artifacts. When comparing skill requirements across years, we need a framework that preserves each period's internal logic while enabling valid measurement of change across time.

Job postings represent firms' real-time, revealed preferences in competitive labor markets. When firms craft job descriptions, they define precise skill boundaries to attract specific human capital while distinguishing roles from potential substitutes. This makes job posting text powerful for observing not just skill inventories within occupations, but the architecture of professional identity as employers perceive and value it. Unlike

static classification systems, job postings capture the dynamic reality of what skills firms actually demand when competing for workers, immediately reflecting new technological requirements and evolving skill combinations.

Traditional data sources cannot address these measurement challenges adequately. The O*NET database represents a systematic skill documentation effort through expert review, cataloging hundreds of detailed occupational characteristics. Our measurement objectives require observing skill demand at high temporal resolution across the full heterogeneity of specific competencies—from Python programming to regulatory compliance to robotic surgery. Expert review systems like O*NET provide valuable standardized taxonomies but update more slowly than real-time firm requirements.

However, this direct approach presents measurement challenges that reveal why naive text analysis fails. Individual job postings contain substantial firm-specific jargon, boilerplate language, and idiosyncratic variations that can obscure genuine occupational signals. Simple averaging of job postings would muddy the occupational signal, giving equal weight to well-written postings with clear skill requirements and boilerplate-heavy postings with generic language that provides little occupational insight. The semantic complexity of skill descriptions means that simple keyword-based approaches fail to capture functional relationships between different skill terms—they cannot understand that Python and R represent functionally similar analytical capabilities for a data scientist, or that stakeholder management represents a more sophisticated manifestation of communication skills. Most fundamentally, any measurement framework must produce representations that are simultaneously temporally stable within occupations and maximally discriminative between similar occupations, as these boundary distinctions carry the greatest economic significance for understanding technological change.

Our solution begins with two comprehensive corpora of online job vacancy postings that capture the economy's skill demands at two critical time points: 2010, representing the immediate aftermath of the Great Recession, and 2018, after nearly a decade of technological advancement and economic recovery. These postings were systematically collected from major job boards and company career pages, providing comprehensive coverage across all skill levels and industries.

The transformation from raw vacancy text to analysis-ready data requires a multi-stage processing pipeline designed to ensure both quality and representativeness. Each posting undergoes systematic quality assessment using automated algorithms that evaluate semantic coherence, grammatical structure, and content relevance, with scores ranging from 2–7. Text length filtering removes postings with insufficient information (fewer than 50 words) or likely errors (exceeding 5,000 words), while duplicate detection employs both exact text matching and semantic similarity to identify and remove near-duplicate postings that could artificially inflate certain occupational signals.

To balance computational feasibility with analytical coverage, we implement stratified sampling that caps each occupation at between 50 and 200 postings per year, ensuring adequate sample sizes for stable prototype generation while preventing large occupations from dominating the analysis. When occupations exceed the 200-posting threshold, we employ quality-weighted sampling to ensure our final samples capture representative skill requirements rather than being skewed toward particular employers or regions. This systematic processing yields two high-quality analytical corpora containing 90,545 job postings for 2010 and 101,374 job postings for 2018, spanning 514 distinct occupations that appear in both years. This sample covers approximately 72% of total U.S. employment, ranging from highly specialized professions to widespread service occupations.

The temporal comparison reveals patterns that foreshadow our findings: average word count per posting increases from 348 words in 2010 to 468 words in 2018, reflecting growing sophistication in job descriptions as firms recognize the importance of precise skill specification. Detailed data processing procedures, including quality assessment algorithms, duplicate detection methods, and stratified sampling implementation, are provided in Appendix 6.

3.2. Model Architecture and Training

Our measurement approach adopts a "Build, Then Align" philosophy that resolves a fundamental tension in temporal analysis. Our approach creates the most accurate possible map of each era using its own information, then carefully aligns them to enable rigorous comparison. This preserves fidelity to each period's skill relationships while maintaining the ability to measure change systematically.

Our architecture is guided by three fundamental requirements that emerge from our economic objectives. The first requirement is *semantic richness*: our representations must capture nuanced functional relationships between skills that drive economic outcomes rather than mere textual similarity. This extends beyond keyword matching to understand that Python programming and R statistical analysis serve similar analytical functions for data scientists, while stakeholder management represents a more sophisticated manifestation of communication skills.

The second requirement is *robust centrality*: each occupation's skill signature must represent stable, core requirements rather than being distorted by firm-specific jargon or sampling artifacts. Different firms may use varying terminology for identical capabilities, and our system must identify consistent occupational signals that persist across surface variations.

The third requirement is *maximal separation*: the framework must excel at distinguishing between similar, potentially substitutable occupations, as these boundary distinctions carry the greatest economic significance for understanding technological change.

This measurement challenge requires a system that captures functional skill relationships determining worker transition costs, identifies stable occupational signals persisting across employer-specific and idiosyncratic posting variations, and distinguishes between occupations where boundary differences matter most for understanding skill transferability. We construct a custom two-stage architecture addressing these requirements sequentially. The first stage aggregates job posting text to extract occupation-level skill signatures while filtering both firm-specific language and idiosyncratic noise. The second stage positions these signatures in a geometric space where distances reflect skill similarity as revealed through employer hiring requirements.

Our solution implements a three-stage architecture that addresses these requirements sequentially. The first stage leverages the E5-large embedding model (Wang et al., 2022) to provide semantic richness, converting each job posting into a 1024-dimensional vector that captures sophisticated semantic relationships between skill terms.[1] E5's transformer architecture understands that when a Database Developer posting lists Oracle PL/SQL, Exadata, Informatica, Tableau, and AutoSys alongside monitoring ETL loads and automating manual processes, these represent a coherent data engineering workflow rather than isolated requirements.

The second stage employs a custom Attention Network (AttnNet) designed to address robust centrality by learning to identify genuine occupational signals while filtering firm-specific noise. The architecture consists of a single linear transformation layer ($1024 \to 1$) that processes each of 32 randomly sampled job postings independently to generate attention logits, followed by softmax normalization to create importance weights that sum to one. This differs fundamentally from transformer self-attention, which computes relationships between postings. Our mechanism asks a more focused question: How representative is this specific posting of the occupation's core skill requirements? For occupations with fewer than 32 postings, the system implements zero-padding with binary masking to ensure padded positions receive effectively zero attention weights.

The third stage implements a Compression Network (CompressNet) that transforms clean 1024-dimensional representations into discriminative 128-dimensional prototypes through a two-layer multilayer perceptron ($1024 \to 512 \to 128$) with ReLU activations and dropout regularization. The network includes an additive skip connection providing direct linear projection ($1024 \to 128$) that ensures stable gradient flow during training and allows the model to learn both complex non-linear transformations and simpler linear projections of the skill space. This stage maps cleaned vectors into a well-separated layout where similar occupations remain appropriately close while dissimilar occupations

[1] Embedding in machine learning refers to mapping discrete objects (words, documents, occupations) into a continuous vector space where geometric properties—distances, angles—correspond to meaningful relationships. In our application, occupations are embedded into $\mathbb{R}^{128}$ such that proximity reflects skill similarity as revealed through employer hiring requirements.

achieve clear separation.

The compression to 128 dimensions learns a task-specific subspace that filters irrelevant variation. The original 1024-dimensional E5 embedding space captures semantic information including tone, style, and contextual markers reflecting firm-specific, industry-specific, or regional linguistic patterns. A job posting for an Accountant at a tech startup may differ linguistically from one at a manufacturing firm despite identical core skill requirements. For constructing a stable occupational map, this contextual variation represents noise rather than signal. The CompressNet learns to preserve only dimensions of variation useful for distinguishing occupations while collapsing dimensions reflecting irrelevant linguistic artifacts. This dimensionality reduction acts as regularization, producing representations that capture transferable skill requirements stripped of confounding context-specific language, while also providing computational efficiency that makes economy-wide measurement tractable. Complete implementation details, including hyperparameters, initialization schemes, and convergence monitoring procedures, are provided in Appendix 6.

Our approach leverages a fundamental property of occupational classifications: job postings from the same occupation should require similar skill profiles, while postings from different occupations should be distinguishable. This assumption—that occupational categories (SOC codes) reflect genuine skill boundaries—provides the supervisory signal for learning. Our data includes SOC code assignments generated by Lightcast's expert-curated algorithms, which map job postings to standardized occupational classifications. This enables systematic supervision at scale while maintaining consistency with official taxonomies.

Our architecture processes occupations rather than individual job postings. During training, each sample consists of $k = 32$ randomly drawn postings from a single occupation. The model receives three such samples simultaneously: an anchor occupation, a positive sample (32 different postings from the same occupation), and a negative sample (32 postings from a different but similar occupation). All three mini-batches flow through the shared two-stage architecture in parallel. AttnNet processes each batch independently, learning attention weights that aggregate the 32 postings into a single 1024-dimensional vector representing that occupation's skill signature. These three occupation-level vectors then pass through CompressNet, which maps them to final 128-dimensional positions. The triplet margin loss compares the three outputs, encouraging the model to position the anchor closer to its positive match (same occupation, different postings) than to the negative alternative (different occupation). Backpropagation adjusts both AttnNet and CompressNet weights jointly, teaching the attention mechanism which postings carry genuine occupational signal and teaching the compression network which semantic features distinguish similar occupations.

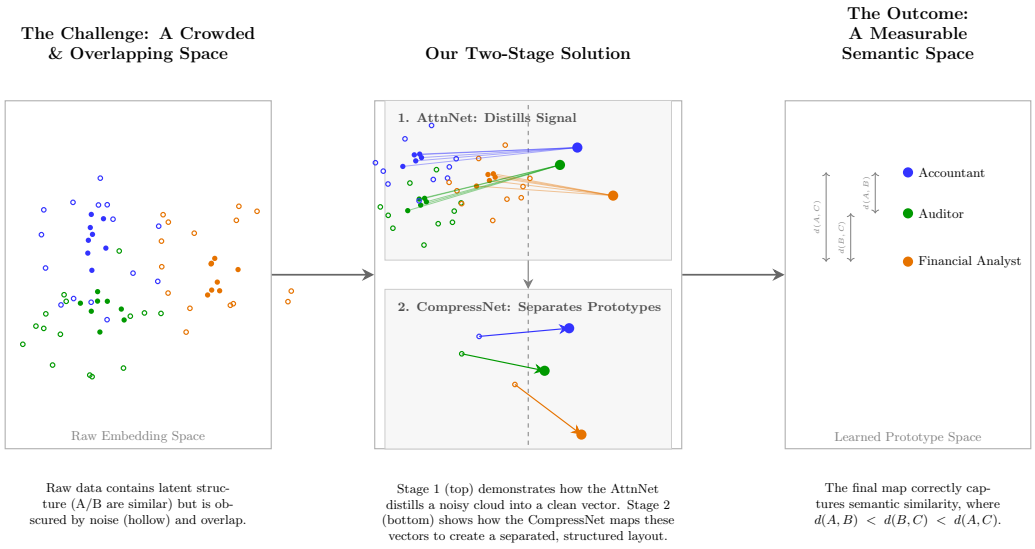


FIGURE 1. A Two-Stage Architecture for Measuring Occupational Skill Content. Panel 1 illustrates the measurement challenge: raw embeddings contain latent occupational structure (blue and green groups share similar skills) but individual job postings include firm-specific noise (hollow circles) and spatial overlap. Panel 2 demonstrates our solution: AttnNet learns to aggregate each occupation's postings by weighting genuine skill signals (filled circles) over noise, producing clean occupation-level vectors (top); CompressNet then positions these vectors to maximize separation between similar occupations (bottom). Panel 3 shows the outcome: a learned prototype space where geometric distances between occupations (Accountant, Auditor, Financial Analyst) reflect skill similarity as revealed through employer hiring requirements. The distance ordering $d(A, B) < d(B, C) < d(A, C)$ captures the economic intuition that Accountants and Auditors require more similar skill portfolios than either shares with Financial Analysts.

Our training approach employs contrastive learning (Chen et al., 2020; Dell, 2025), specifically supervised contrastive methods (Khosla et al., 2020), to teach the system skill relationships through systematic comparisons, which proves essential for learning discriminative representations where geometric distances reflect genuine skill substitutability. Each training step processes triplets: an anchor occupation, a positive sample (different postings from the same occupation), and a negative sample (postings from a similar but distinct occupation). All three samples flow through the complete architecture, with both AttnNet and CompressNet processing each sample in sequence—the attention mechanism first aggregates posting clouds into clean intermediate vectors, then the compression network maps these to final 128-dimensional prototype positions.

Crucially, both networks train jointly through the same contrastive objective. This end-to-end learning approach means the attention mechanism learns to filter noise specifically in ways that help discrimination, while the compression network receives clean, de-noised inputs that make boundary learning easier. This joint optimization proves superior to the obvious alternative of first cleaning data with predetermined criteria, then learning occupational boundaries separately.

The training proceeds in two phases that address distinct optimization challenges sequentially. Phase 1 trains AttnNet using mean squared error loss to reproduce pre-computed occupation centroids. This provides a dense, stable learning signal teaching the attention mechanism to identify postings reflecting an occupation's central tendency—a straightforward aggregation task establishing a sensible initial filtering policy. Phase 2 then trains the full architecture jointly using triplet margin loss (Schroff, Kalenichenko, and Philbin, 2015), introducing a sparser and more challenging signal that carves out fine-grained decision boundaries between similar occupations. This curriculum prevents forcing a randomly initialized attention mechanism to simultaneously learn basic aggregation and complex discrimination, which can produce unstable training dynamics. The sequential approach allows the system to master representational stability before tackling the harder metric learning objective.

Formally, during the warm-up phase, the model is trained using mean squared error (MSE) loss against pre-computed occupation centroids:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^{N} \|\hat{\mathbf{h}}_i - \mathbf{c}_i\|^2 \tag{1}$$

where $\hat{\mathbf{h}}_i$ denotes the predicted representation of occupation i and $\mathbf{c}_i$ is the corresponding centroid.

In the main training phase, we adopt the triplet margin loss:

$$\mathcal{L}_{\text{triplet}} = \frac{1}{M}\sum_{j=1}^{M} \max\left(0,\ \|\mathbf{h}_j^a - \mathbf{h}_j^p\|^2 - \|\mathbf{h}_j^a - \mathbf{h}_j^n\|^2 + \alpha\right) \tag{2}$$

where each triplet consists of an anchor vector $\mathbf{h}_j^a$, a positive sample $\mathbf{h}_j^p$, and a negative sample $\mathbf{h}_j^n$, with margin parameter $\alpha = 0.5$.

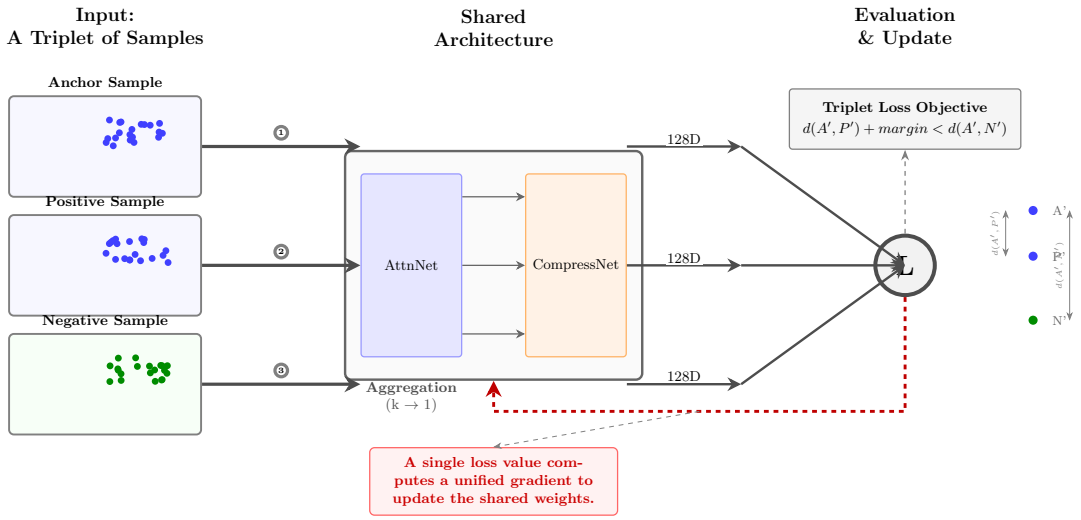


FIGURE 2. Contrastive Learning with Occupational Category Supervision. The training process uses triplets of occupation samples, where each sample contains $k = 32$ randomly drawn job postings from a single occupation. The anchor and positive samples draw from the same occupation (both blue, indicating shared SOC code), while the negative sample draws from a different occupation (green). All three samples pass through the shared architecture simultaneously: AttnNet aggregates each set of 32 postings into a single occupation-level vector (the $k \to 1$ aggregation), filtering noise through learned attention weights; CompressNet then maps these three vectors to final 128-dimensional positions (A', P', N'). The triplet margin loss evaluates whether $d(A', P') + \text{margin} < d(A', N')$, with backpropagation updating both networks to satisfy this constraint. This joint optimization teaches AttnNet which postings are representative and CompressNet which features distinguish occupations.

The effectiveness of contrastive learning depends critically on negative sample selection—choosing which different occupation to contrast against the anchor. Random negatives produce trivial learning signals because obviously distinct occupations (e.g., Accountants versus Construction Workers) require no sophisticated discrimination. We instead employ semi-hard negative mining, restricting negative samples to each occupation's most similar neighbors based on pre-computed embeddings.[2] This focuses the model on learning fine-

[2]We identify similar neighbors using cosine similarities of mean E5 embeddings computed before training,

grained distinctions between genuinely similar occupations—such as Software Engineers versus Data Scientists—ensuring our measurements capture boundary dynamics that determine skill transferability.

The model successfully learns genuine skill similarity: decreasing triplet loss and nearest neighbor analysis show that learned representations cluster occupations according to skill similarity rather than superficial characteristics. Technical occupations cluster appropriately with programming-focused roles grouping together while remaining distinct from hardware-oriented positions, while healthcare occupations cluster by functional care delivery requirements, with Registered Nurses closest to Licensed Practical Nurses based on patient care skills while remaining distinct from Physicians despite shared workplace environments.

3.3. Aligning Vector Spaces for Temporal Comparison

We align the two vector spaces using Orthogonal Procrustes analysis, following established methods for aligning embedding spaces across time periods (Hamilton, Leskovec, and Jurafsky, 2016), with 21 stable occupations as anchor points. This approach rests on an economic assumption: while individual occupations may undergo substantial skill changes between 2010 and 2018, the overall architecture of skill relationships exhibits relative stability over this period. Our contrastive learning process creates an alignment challenge: the triplet margin loss optimizes only relative relationships within each year, not absolute positions. The constraint that occupation A should sit closer to P than to N remains satisfied under any rotation or reflection of the entire space. Independent training on two datasets thus produces two embedding spaces with arbitrarily different coordinate systems—dimensions in 2010 have no inherent correspondence to dimensions in 2018. Without alignment, we cannot distinguish genuine occupational transformation from coordinate system artifacts. An occupation moving from position (0.5, 0.3) in 2010 to (0.3, 0.5) in 2018 might reflect skill change, arbitrary rotation, or both.

This stability assumption reflects our understanding that technological change primarily affects the content of occupational skill requirements rather than completely reorganizing the fundamental structure of skill relationships across the entire economy. Most fundamental relationships characterizing the 2010 labor market should persist in 2018 despite technological advancement. Technology may alter specific programming languages required for software roles or introduce new analytical tools for financial positions, but it rarely transforms the basic architecture of how different professional domains relate to each other.

This stability assumption enables our alignment by providing anchor points—occupational

creating restricted sampling pools that exclude obviously dissimilar occupations while maintaining sufficient variation for effective learning.

relationships that should remain approximately constant across time periods. We selected 21 stable occupations to serve as anchors based on three criteria: licensing requirements that constrain skill evolution, specialized training that creates stable professional boundaries, and fundamental service needs that remain consistent over time. These anchors span five conceptual categories: Healthcare Practitioners (including Dentists, Pharmacists, Physical Therapists), Skilled Trades (including Electricians, Plumbers, Carpenters), Protective and Civil Services (including Police Officers, Firefighters, Elementary School Teachers), Transportation (including Airline Pilots, Truck Drivers), and Human Services (including Hairdressers, Chefs, Medical Assistants). The verification process confirmed that 21 of our 23 candidate anchors appeared in both datasets with sufficient data quality. The complete list of anchor occupations with SOC codes, selection criteria, and verification procedures are detailed in Appendix 6.

We implement temporal alignment using Orthogonal Procrustes analysis, which finds the optimal rigid transformation to align two point clouds while preserving their internal geometric structure:

$$\mathbf{Q}^* = \arg\min_{\mathbf{Q}} \sum_{o \in \text{anchors}} \|\mathbf{p}_{o,2010} - \mathbf{Q} \cdot \mathbf{p}_{o,2018}\|^2 \tag{3}$$

subject to the constraint that $\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}$. The algorithm treats our 21 stable occupations as anchor points, identifying the orthogonal transformation matrix that minimizes the sum of squared distances between corresponding prototypes. The orthogonality constraint ensures the transformation preserves distances and angles within the 2018 space, maintaining the skill relationships we've carefully learned.

The alignment achieves an 82% reduction in anchor pair distances (from 1.17 to 0.21), confirming the transformation captures genuine correspondences rather than artifacts. This dramatic improvement in anchor alignment provides strong evidence that our linear transformation assumption is appropriate and that the method successfully establishes a common coordinate system for temporal comparison. Mathematical implementation details, including the SVD solution and validation metrics, are provided in Appendix 6.

The aligned 2018 prototypes complete our measurement infrastructure, enabling reliable quantification of occupational change. With both years embedded in the same coordinate system, we can construct specific metrics of occupational transformation that capture the economic phenomena we seek to understand rather than methodological noise. This measurement engine represents a novel methodological contribution, providing the first systematic approach for quantifying occupational skill change at economy-wide scale using job posting text. With this measurement engine constructed and its internal components validated, we can now turn to the decisive test: whether these new measures reveal meaningful, previously unobservable patterns in real-world wage and

employment outcomes.

4. Measuring Occupational Transformation

4.1. Two Geometric Dimensions

Section 3 developed measurement infrastructure for embedding occupations in aligned skill spaces, positioning each occupation based on the capabilities employers seek when hiring. This enables observing how 514 occupations repositioned between 2010 and 2018 as technological change reshaped skill requirements across the economy. The natural next question: how do we characterize this movement systematically?

Any movement in geometric space possesses two fundamental properties that provide complementary information. First, magnitude captures how far an occupation's skill requirements changed—the distance its prototype vector traveled in the embedding space between time periods. Second, relative positioning captures whether the occupation moved toward or away from other occupations—changes in its proximity to neighboring occupations. These correspond to scalar information (distance traveled) and relational information (proximity to neighbors), both independently informative about the nature of transformation.

Magnitude alone provides incomplete characterization because it discards directional information. Two occupations can travel identical distances to completely different destinations. One may move to an isolated region where its skills remain unique, maintaining distance from all potential substitutes. Another may move toward a crowded area where its skills now resemble many others, with proximity to multiple neighbors. Magnitude alone cannot distinguish these geometrically opposite situations—one occupation preserved its distinctiveness despite transformation, while the other converged toward neighbors. This motivates measuring relative positioning to capture changes in occupational distinctiveness that magnitude misses.

We formalize these properties through two measures. Section 4.2 defines Skill Evolution Magnitude (SEM), quantifying transformation magnitude through geometric distance traveled. Section 4.3 defines changes in distinctiveness rankings (ΔRANK), capturing whether an occupation became more or less distinctive relative to all others through proximity patterns to nearest neighbors. Critically, proximity patterns—asymmetric versus uniform—form distinct geometric signatures, which we define formally below. Section 4.4 shows these measures combine in systematically different ways across occupations. Section 5 develops the mechanism explaining why these combinations predict opposing economic outcomes.

Our empirical approach combines three elements: occupational wage levels from BLS data, which provide baseline context; transformation magnitude (SEM) from embedding

movements; and distinctiveness shifts (ΔRANK) from proximity patterns. Neighbor identity—which occupations occupy adjacent skill space—provides evidence for interpreting what transformation and boundary shifts represent economically.

4.2. Measuring Transformation Magnitude (SEM)

Skill Evolution Magnitude quantifies the intensity of an occupation's internal skill reconfiguration between 2010 and 2018. We measure this using cosine distance between an occupation's prototype vectors across time periods:

$$\text{SEM} = 1 - \frac{\mathbf{p}_{2010} \cdot \mathbf{p}_{2018}}{\|\mathbf{p}_{2010}\| \, \|\mathbf{p}_{2018}\|} \tag{4}$$

where $\mathbf{p}_{2010}$ and $\mathbf{p}_{2018}$ represent an occupation's prototype vectors in the respective years. This scalar measure ranges from 0 (identical skill requirements) to 2 (completely reversed skill profiles), providing continuous quantification of transformation intensity.

Geometrically, SEM measures the distance an occupation's prototype vector traveled in the embedding space. Large SEM values indicate substantial repositioning—the occupation's 2018 skill requirements differ markedly from its 2010 requirements. Small SEM values indicate relative stability—the occupation maintained similar competency profiles despite the surrounding technological transformation. This geometric displacement directly captures the magnitude of change in what employers sought when hiring for this occupation.

Economically, SEM captures the internal skill reconfiguration that occurred within occupational boundaries. When an occupation's prototype vector moves substantially, this indicates employers demanded different capabilities in 2018 than in 2010. Incumbent workers face human capital updating pressure—the skills that made them valuable in 2010 may have partially depreciated, while entirely new competency domains became essential. SEM quantifies the degree to which technological change rewrote the competency requirements for successful job performance within the occupation's traditional scope. High SEM occupations experienced fundamental reconceptualization of role requirements, while low SEM occupations experienced incremental evolution within stable competency frameworks.

Consider Financial Managers (SEM = 0.67) versus Elementary School Teachers (SEM = 0.12). Financial Managers' prototype vectors show substantial displacement, indicating employers sought markedly different capabilities in 2018—quantitative analytics, predictive modeling, data-driven forecasting—compared to traditional relationship-based financial planning in 2010. Teachers show modest evolution, with 2018 requirements emphasizing technology integration within stable pedagogical competencies. Both adapted to technological change, yet transformation magnitudes differed substantially.

Empirically, SEM exhibits substantial variation across our 514 occupations. Values range from 0.076 to 0.751 with a positively skewed distribution (mean = 0.229). This pattern reveals that most occupations experienced modest skill evolution consistent with incremental technological adaptation. However, a significant tail underwent intensive transformation, reflecting fundamental shifts in their human capital requirements. The positive skew indicates that while technological change during 2010–2018 was not uniformly disruptive, a meaningful subset of the economy experienced substantial restructuring warranting systematic measurement.

SEM quantifies transformation magnitude but not the nature of that transformation. Two occupations with identical SEM may have absorbed fundamentally different types of capabilities with completely different economic implications. An occupation might add sophisticated new tools that integrate deeply with existing domain expertise, requiring all practitioners to master capabilities building on their specialized knowledge. Another might adopt broadly-applicable analytical skills used across many professions, enabling some workers to develop competencies valuable beyond their original occupational context. Both show high SEM—substantial skill reconfiguration occurred—yet the geometric character of change differs. Magnitude alone provides incomplete information about transformation's character.

This limitation motivates our second measure, which captures changes in occupational distinctiveness by examining how an occupation repositioned relative to neighboring occupations in skill space.

4.3. Measuring Distinctiveness Shifts (ΔRANK)

Distinctiveness captures whether an occupation's skill requirements remain geometrically separated from neighboring occupations. We measure distinctiveness through proximity patterns to nearest neighbors in skill space.

The Neighborhood Separation Index (NSI) quantifies the relationship between an occupation's closest neighbor and its broader local neighborhood:

$$\text{NSI} = \frac{d_1}{\bar{d}_5} \tag{5}$$

where d_1 represents the cosine distance to an occupation's nearest neighbor and $\bar{d}_5$ represents the average cosine distance to its five nearest neighbors. This ratio detects local geometric asymmetry in proximity patterns. Low NSI values indicate the nearest neighbor sits substantially closer than other local neighbors—an asymmetric configuration. High NSI values indicate all nearby neighbors maintain similar distances—uniform spacing.

When NSI is low ($d_1 \ll \bar{d}_5$), the nearest neighbor sits substantially closer than others, creating asymmetric proximity. This geometric configuration means one specific neigh-

boring occupation sits unusually close. Consider a geometric example: if an occupation's nearest neighbor is at distance 0.05 while the next four neighbors average 0.20, NSI = 0.25. This pronounced asymmetry shows one specific neighbor converged substantially more than others. Section 5 examines whether asymmetric proximity patterns predict distinct labor market outcomes.

When NSI is high ($d_1 \approx \bar{d}_5$), all nearby neighbors maintain comparable distances, creating uniform spacing. This geometric configuration means no single neighbor sits substantially closer than others. Consider a contrasting geometric example: if an occupation's five nearest neighbors are all at distances 0.14–0.18, NSI = 0.90. No single neighbor became unusually close—the local geometry stayed symmetric. The occupation maintained balanced separation from its entire neighborhood. Section 5 examines whether uniform spacing patterns predict distinct labor market outcomes.

NSI levels alone prove insufficient because they capture static positioning rather than dynamic change. An occupation might maintain high NSI simply because it always operated in a sparse region, not because its boundaries strengthened during this period. We require temporal measurement: $\Delta\text{NSI} = \text{NSI}_{2018} - \text{NSI}_{2010}$ isolates changes in distinctiveness from starting conditions. However, even this faces limitations because the entire occupational landscape reorganized simultaneously—an occupation might experience NSI changes simply because surrounding occupations transformed, not because its own boundaries shifted economically.

The solution employs percentile rank changes, creating our final measure: $\Delta\text{RANK} = \text{NSI_rank}_{2018} - \text{NSI_rank}_{2010}$. We convert each occupation's NSI to its percentile rank within the economy-wide distribution for each year, then measure how this rank shifted. This percentile-based approach captures whether an occupation became more or less distinctive relative to how all other occupations simultaneously evolved. Falling ΔRANK indicates an occupation lost distinctiveness compared to the broader economy—it moved toward lower NSI percentiles as either its own nearest neighbor converged or it fell behind others' increasing separation. Rising ΔRANK indicates an occupation gained distinctiveness—it moved toward higher NSI percentiles by maintaining or increasing separation while others' boundaries eroded. ΔRANK isolates relative positioning dynamics, providing clean measurement of competitive environment changes independent of internal evolution.

4.4. Systematic Geometric Patterns

When we compute SEM and ΔRANK across 514 occupations, systematic variation emerges in how transformation magnitude and distinctiveness shifts combine. Different patterns appear with distinct geometric signatures. These combinations vary in ways that prove economically meaningful, though Section 5 will develop the mechanism explaining why.

Some occupations exhibit substantial internal change combined with falling distinctiveness. Among occupations with high SEM (above median transformation magnitude), a subset shows declining ΔRANK with mean value of -0.26. These occupations show asymmetric proximity patterns—their NSI declined as nearest neighbors became substantially closer relative to local neighborhood averages. The geometric signature indicates transformation coincided with asymmetric convergence to nearest neighbor.

Other occupations show comparable transformation magnitude yet stable distinctiveness. Among high-SEM occupations, another subset maintains stable ΔRANK with mean value of $+0.04$. These preserve uniform spacing patterns—distances to all neighbors remain proportional as NSI maintained levels despite internal skill reconfiguration. The geometric signature indicates transformation occurred without creating particular proximity to any specific alternative occupation.

A third group combines minimal internal change with falling distinctiveness. Among occupations with low SEM (below median transformation magnitude), many show declining ΔRANK with mean value of -0.23. These display NSI decline without substantial prototype vector movement. The geometric signature indicates distinctiveness eroded despite limited internal skill evolution.

These patterns also differ in the identity of nearest neighbors observed in 2018. Among occupations with falling distinctiveness and high transformation, many find 2018 neighbors in analytical professions requiring statistical modeling, quantitative forecasting, or data-driven decision making. Among those with stable distinctiveness despite transformation, neighbors remain in original professional clusters—design professions, healthcare specialties, legal services—maintaining coherent domain boundaries. Among those with falling distinctiveness but low transformation, neighbors shift toward general service roles. The neighbor configurations vary systematically, though what this variation means for workers requires mechanism development.

The geometric patterns from Section 4.3 suggest these combinations may reflect different transformation processes. Asymmetric proximity describes convergence toward specific neighbors, while uniform spacing describes maintained separation from all neighbors. If this geometric intuition holds, these signatures should predict different economic outcomes. High transformation with asymmetric proximity may create different worker pressures than high transformation with uniform spacing. Minimal transformation with proximity changes may operate through entirely different mechanisms than substantial transformation with boundary stability. Section 5 examines whether combinations of transformation magnitude, distinctiveness shifts, and neighbor configurations indeed predict opposite wage dynamics within occupations experiencing different transformation patterns.

4.5. The Complementarity of SEM and ΔRANK

SEM and ΔRANK capture complementary aspects of occupational reorganization, providing distinct information about transformation character. Across our 514 occupations, the correlation between SEM and ΔRANK is 0.08, indicating the measures identify largely orthogonal dimensions. This low correlation demonstrates that transformation magnitude and distinctiveness shifts vary independently—high transformation can occur with falling, stable, or rising distinctiveness; distinctiveness can shift substantially regardless of transformation magnitude.

Yet Section 4.4 showed certain combinations appear systematically rather than randomly. This reveals that while the measures are mathematically independent, they combine in patterns suggesting underlying processes structure how occupations reorganize. SEM alone tells transformation magnitude but not the nature of change—whether transformation moved the occupation toward capabilities found in other professions or occurred within professional boundaries. ΔRANK alone tells boundary shifts but not whether substantial capabilities changed—distinctiveness can erode without transformation or maintain while capabilities evolved substantially. Together, the measures enable characterizing transformation comprehensively.

This complementarity validates the two-dimensional framework introduced in Section 4.2. Magnitude alone provides incomplete information about transformation because identical distances can reflect movement toward isolation or convergence. Distinctiveness shifts alone miss whether internal capabilities evolved substantially. The combination reveals transformation's geometric signature—how occupations repositioned both internally and relationally—providing the foundation for mechanism development.

Section 5 leverages this complementarity to develop a formal mechanism explaining why geometric combinations predict opposing economic outcomes. The measures jointly identify transformation patterns; Section 5 explains what these patterns mean for workers facing technological change.

4.6. From Measurement to Mechanism

Section 4 established two geometric properties characterizing occupational movement in skill space—transformation magnitude (SEM) and distinctiveness shifts (ΔRANK)—that vary systematically across occupations. These patterns are systematic but require interpretation: What do different combinations mean economically? Do they predict distinct labor market outcomes? Section 5 investigates by developing economic interpretations of geometric signatures and testing whether they predict opposing wage dynamics.

5. Patterns, Mechanisms, and Validation

5.1. The Concentration of Technological Transformation

We construct occupational skill embeddings and compute transformation measures (SEM, ΔRANK) for 514 occupations appearing in both our 2010 and 2018 job posting samples. To validate these measures and examine labor market outcomes, we merge with Bureau of Labor Statistics Occupational Employment and Wage Statistics (OEWS) data, which provides occupation-level employment counts and wage distributions at the 25th, 50th, and 75th percentiles for May 2010 and May 2018.

OEWS uses Standard Occupational Classification (SOC) codes to define occupations. Between 2010 and 2018, BLS revised the SOC system—some occupations were redefined, merged, or split. Of our 514 occupations with embeddings, we achieve clean 1-to-1 matches for 432 occupations where SOC definitions remained consistent, enabling reliable wage comparisons. We exclude 82 occupations with classification changes to ensure observed wage evolution reflects actual dynamics rather than definitional artifacts. The 432-occupation sample represents 91 million workers in 2010 (72

TABLE 1. Workforce Composition and Baseline Economic Context, 2010–2018

Wage Group	N Occupations	Employment 2010 (millions)	Employment 2018 (millions)	Employment Growth (%)	Wage Growth P50 (%)
Low-wage (bottom 75%)	324	74.5	81.7	+9.6	+16.1
High-wage (top 25%)	108	16.5	20.3	+23.0	+11.9
Total	**432**	**91.0**	**102.0**	**+12.1**	**+14.3**

Note: Wage groups defined by 2010 median wage. Employment figures represent sample coverage (72% of total OEWS employment). Growth rates are employment-weighted means.

TABLE 2. Skill Transformation Intensity by Wage Group

Wage Group	Mean SEM	Mean Boundary Erosion	% High-SEM Occupations	% Eroding Boundaries
Low-wage	0.247	−0.087	22.6%	31.3%
High-wage	0.449	−0.053	74.3%	32.0%

Note: SEM measures occupational skill reconfiguration magnitude. Boundary Erosion = −ΔRANK. High-SEM classification determined by Jenks natural breaks algorithm. Eroding defined as ΔRANK < −0.1. All statistics employment-weighted.

Table 2 reveals transformation concentrated in high-wage work. High-wage occupations show mean SEM of 0.449 compared to 0.247 for low-wage occupations—an 81.9% difference. Yet boundary erosion shows no wage gradient: 31.3% of low-wage versus 32.0% of high-wage occupations experienced falling distinctiveness. This confirms SEM and

ΔRANK measure distinct dimensions. The technological disruption narrative accurately describes high-wage work but misrepresents both broader workforce experience and critical heterogeneity within high-wage transformation itself—occupations showing similar transformation magnitudes faced fundamentally different mechanisms.

TABLE 3. Empirical Classification Structure and Analytical Focus

HIGH-WAGE OCCUPATIONS (>P75, N=108)		
	High SEM (74.3%)	**Low SEM (25.7%)**
Boundary Erosion (ΔRANK < -0.1)	**Type 1: Cognitive Distance** [FOCUS]	*Included in regression*
Stable Boundaries (ΔRANK ≥ -0.1)	Type 2: Scaffolded Learning [REFERENCE]	*Included in regression*

LOW-WAGE OCCUPATIONS ($\leq$P75, N=324)		
	High SEM (22.6%)	**Low SEM (77.4%)**
Boundary Erosion (ΔRANK < -0.1)	*Included in regression*	**Type 3: Simplification** [FOCUS]
Stable Boundaries (ΔRANK ≥ -0.1)	*Included in regression*	[REFERENCE] *(for Type 3)*

Note: This table illustrates the conceptual structure of the empirical analysis. Occupations are stratified by 2010 median wage. Within each stratum, they are classified by transformation magnitude (High/Low SEM) and boundary dynamics (Erosion/Stable). Percentages are employment-weighted and show the concentration of high-SEM transformation in high-wage work. The analysis focuses on comparing the [FOCUS] cells to their [REFERENCE] counterparts to identify the distinct effects of boundary erosion in high- versus low-transformation contexts. All four cells within each stratum are included in the regression analysis.

This concentration reflects that high transformation operates differently across wage levels. High-wage occupations hold specialized expertise where boundary erosion signals movement toward other domains; low-wage occupations hold generic capabilities where distinctiveness loss signals commodification. Wage level thus reveals whether transformation operates from specialized or generic human capital. We focus on dominant configurations—high-SEM among high-wage occupations (74.3%) and low-SEM among low-wage occupations (77.4%)—examining mechanisms where they operated most intensively rather than imposing uniform treatment effects across heterogeneous contexts.

Our mechanism labels—Cognitive Distance, Scaffolded Learning, Simplification—describe specific cells in the wage-stratified design shown in Table 3. Cognitive Distance (high-wage, high-SEM, eroding boundaries) versus Scaffolded Learning (high-wage, high-SEM, stable boundaries) is identified through the High-SEM $\times$ Boundary_Erosion interaction in the high-wage sample. Simplification (low-wage, low-SEM, eroding boundaries) is identified through the Low-SEM $\times$ Boundary_Erosion interaction in the low-wage sample. Section 5.2

examines each mechanism in detail before validating their distinct wage effects.

5.2. Three Distinct Processes Driven by Different Mechanisms

Section 4.3 identified two geometric patterns in how occupations repositioned: asymmetric proximity (one neighbor substantially closer than others) and uniform spacing (all neighbors at comparable distances). These patterns prove diagnostic of underlying mechanisms. We now examine specific occupations' neighbor configurations and what they reveal about the relationship between new skill requirements and existing expertise. Three systematic patterns emerge, each combining transformation magnitude and distinctiveness shifts in ways that create fundamentally different economic pressures for incumbent workers.

TABLE 4. Three Processes of Technological Transformation with Representative Occupations

Process	Characteristics	Representative Occupations
Type 1: Cognitive Distance Creating Heterogeneous Adaptation High-wage + High-SEM + Eroding	High SEM (0.546) High Neighbor SEM (0.562) Boundary Erosion (−0.263)	Financial Managers, Computer and Information Systems Managers, Medical and Health Services Managers, Chemical Engineers, Economists
Type 2: Scaffolded Learning Enabling Collective Upgrading High-wage + High-SEM + Stable	High SEM (0.524) High Neighbor SEM (0.523) Stable Boundaries (+0.042)	Actuaries, Aerospace Engineers, Statisticians, Financial Analysts, Materials Scientists
Type 3: Barrier Reduction Through Simplification Low-wage + Low-SEM + Eroding	Low SEM (0.250) Low Neighbor SEM (0.259) Boundary Erosion (−0.233)	Shipping, Receiving, and Traffic Clerks, Library Technicians, Printing Press Operators, Couriers and Messengers

Note: Processes identified by transformation intensity (SEM), boundary dynamics (ΔRANK), and coordination (Neighborhood_SEM). Neighborhood_SEM is the mean SEM of an occupation's five nearest neighbors in skill space, measuring the transformation intensity of proximate occupations. Characteristics show mean values. Representative occupations selected as clear exemplars.

Type 1: Cognitive Distance and Heterogeneous Adaptation

Type 1 occupations exhibit high SEM (mean 0.546) combined with falling ΔRANK (mean −0.263), indicating substantial skill reconfiguration alongside boundary erosion. Financial Managers exemplify this pattern. The neighborhood evolution provides decisive evidence for skill content. In 2010, Financial Managers' five nearest neighbors were Agents and Business Managers of Artists/Performers/Athletes, Compensation and Benefits Managers, Insurance Underwriters, Aerospace Engineers, and Market Research Analysts—reflecting

business administration, risk management, and project coordination capabilities. By 2018, the nearest neighbors became Economists, Statisticians, Civil Engineers, Purchasing Managers, and Agents/Business Managers—a dramatic shift toward analytical and quantitative professions. The convergence with Economists and Statisticians reveals employers sought analytical capabilities—data analytics, statistical modeling, quantitative reasoning—representing fundamentally different competencies from traditional financial management expertise.

By 2018, Financial Managers absorbed predictive modeling, quantitative risk assessment, and data-driven strategic planning. The wage separation pattern we observe suggests this work proved organizationally flexible—firms could accommodate practitioners with varying skill bundles. Some positions emphasized analytical capabilities while others maintained traditional relationship management and qualitative judgment, both coexisting under the same occupational title. This organizational flexibility enabled heterogeneous adaptation across the incumbent workforce.

These analytical capabilities proved cognitively distant from traditional financial management expertise, creating high-variance acquisition costs across workers. Traditional Financial Managers developed relationship-based judgment, industry knowledge, and qualitative reasoning over careers. The new capabilities—statistical modeling, computational thinking, data science frameworks—require fundamentally different cognitive approaches. Acquiring them demanded paradigm integration: reconciling conflicting ways of approaching problems rather than building additively on existing knowledge.

This cognitive distance created heterogeneous adaptation patterns. Some workers successfully integrated both paradigms, bridging the conceptual gap between traditional financial judgment and quantitative analytics. Others struggled with the cognitive conflicts, maintaining traditional approaches. The variance in acquisition success generated internal scarcity—only a subset of Financial Managers developed the analytical capabilities employers increasingly valued.

This internal scarcity generates the wage separation we observe. While occupational skill requirements evolved (boundary erosion), heterogeneous adaptation created winners and losers within the occupation. Workers who successfully bridged the cognitive distance acquired analytical capabilities that proved valuable within and beyond finance. These capabilities created external options—statistical modeling and quantitative analytics commanded value across research, policy analysis, and strategic consulting, enabling workers who mastered them to capture premiums. Workers maintaining traditional approaches lacked such leverage.

The result: within-occupation divergence driven by differential acquisition success. Workers who mastered cognitively distant capabilities pulled away at the distribution's top while those who couldn't adapt experienced slower growth. Internal scarcity from high-

variance acquisition costs generated wage separation despite—indeed because of—similar technological exposure across the occupation.

Type 2: Scaffolded Learning and Collective Upgrading

Type 2 occupations show comparable transformation magnitude (mean SEM 0.524) yet stable distinctiveness (mean ΔRANK $+0.042$). Architects exemplify this pattern. Technology transformed practice through Building Information Modeling, integrated design systems, and advanced environmental modeling. The neighborhood evolution reveals skill content remained domain-bound. Architects' 2018 neighbors stayed within the built environment ecosystem—Landscape Architects, Civil Engineers, Soil Scientists—not converging toward generic analytical capabilities. New tools required deep architectural principles: spatial relationships, structural systems, environmental factors, building codes specific to design practice. These capabilities built directly on knowledge architects already possessed, maintaining occupational distinctiveness while transforming practice.

The stable wage pattern we observe suggests work organization required uniform capabilities across all practitioners. The integrated nature of modern design workflows appears to have precluded mixing traditional and digital approaches within project teams—contemporary projects required all participants work within the same digital environment. This organizational constraint meant all architects practicing contemporary work faced the same adoption imperative.

New capabilities scaffolded onto existing architectural knowledge rather than requiring cognitive paradigm shifts. Building Information Modeling maps directly to concepts architects already mastered: spatial relationships, structural systems, material properties, building codes. Every digital feature corresponds to traditional architectural principles practitioners spent careers developing. Learning BIM meant extending existing mental models—applying familiar spatial reasoning to digital environments, translating established design principles into new software interfaces.

This scaffolding created uniform acquisition costs across incumbent practitioners. While learning curves existed, they followed predictable paths leveraging accumulated domain expertise. The cognitive distance between traditional architectural practice and digital tools proved minimal—architects could build new capabilities on established foundations rather than integrating conflicting paradigms. All architects practicing contemporary work possessed the prerequisite conceptual knowledge enabling successful adaptation.

This uniformity in acquisition costs generates the stable wage distribution we observe. While occupational skill requirements evolved substantially, scaffolded learning enabled collective upgrading. All practitioners could successfully adapt by building on existing expertise, precluding the internal scarcity that generates wage separation.

The capabilities architects developed remained valuable primarily within architecture, limiting external options. Enhanced proficiency in integrated design systems commanded value for architectural practice but transferred poorly to other professional contexts. Without external opportunities creating differential premiums for early adopters, practitioners at all levels upgraded together through uniform acquisition, maintaining existing wage hierarchies rather than creating separation. Stable wages reflect collective adaptation through scaffolded learning, not absence of transformation. Aerospace Engineers and Actuaries show parallel patterns—advanced tools requiring deep domain knowledge (aerodynamics, mortality modeling) with neighborhoods reflecting coordinated advancement within specialized corridors.

Type 3: Barrier Reduction Through Simplification

Type 3 combines minimal transformation (mean SEM 0.250) with boundary erosion (mean ΔRANK -0.233). Library Technicians exemplify this pattern. Job posting language in 2018 emphasized nearly identical competencies as 2010: patron assistance, circulation management, basic reference support. Yet distinctiveness fell substantially through simplification—technology eliminated specialized knowledge rather than demanding new capabilities.

The neighborhood evolution provides stark evidence. In 2010, Shipping Clerks' nearest neighbors were functionally related logistics roles—a coherent professional community. By 2018, neighbors became unrelated production and service work: retail salespersons, machine operators, general laborers. Library Technicians show sharper dissolution: 2010 neighbors were skilled support roles (pharmacy technicians, office specialists). By 2018, neighbors included vehicle cleaners and dining room attendants. Digital catalog systems eliminated specialized classification knowledge previously distinguishing Library Technicians from general service workers.

Technology eliminated specialized procedural knowledge rather than demanding cognitively distant new capabilities. Barcode scanning automated transaction recording, digital inventory platforms eliminated coordination knowledge, integrated software reduced routing specialization. The specialized expertise that would have generated experience premiums became less critical. Meanwhile, entry-level wages across these occupations and newly-proximate neighbors remained similar—basic requirements became comparable, creating limited variance at the bottom. This predicts wage compression: top wages rise modestly as specialization premiums erode while bottom wages remain stable, narrowing overall spread.

Three distinct mechanisms link technological change to worker outcomes through different acquisition cost structures. Type 1: New requirements proved cognitively distant from existing expertise, creating high-variance acquisition costs that generated heteroge-

neous adaptation and internal scarcity. While occupational boundaries eroded, not all workers successfully bridged the cognitive gap. Workers who did acquire the capabilities gained external options (skills proved portable across analytical roles), amplifying premiums through differential compensation. Type 2: New requirements scaffolded onto existing knowledge, creating uniform acquisition costs that enabled collective upgrading. All practitioners could adapt by building on established expertise, maintaining stable wage structures despite substantial transformation. Enhanced capabilities remained valuable primarily within the occupation, limiting external options for differential premiums. Type 3: Technology eliminated specialized knowledge requirements rather than demanding new capabilities, eroding both occupational distinctiveness and the experience premiums protective expertise once commanded. The diagnostic framework (SEM $\times$ ΔRANK) identifies which mechanism operated based on geometric signatures in skill space.

5.3. Validating Mechanisms Through Wage Dispersion

We test predictions using OEWS percentile wage data (P25, P50, P75) to construct within-occupation wage spreads (P75−P25) and examine evolution between 2010 and 2018. We estimate separate models for higher-wage and lower-wage occupations because boundary erosion operates differently across these contexts—Type 1 dynamics (heterogeneous adaptation from cognitive distance) versus Type 3 dynamics (commodification through simplification).

The higher-wage sample comprises occupations above the 75th percentile of 2010 median wage distribution (N=108), capturing professional work where transformation concentrated. The lower-wage sample comprises occupations below this threshold (N=324). We use the 75th percentile cutoff because transformation concentrated in upper wage tiers (Table 2), providing clean separation between contexts. We estimate the following specification separately for each sample:

$$(6)\quad Y_i = \beta_0 + \beta_1 \cdot 1\{\text{SEM}_i > \alpha\} + \beta_2 \cdot (-\Delta\text{Rank}_i) + \beta_3 \cdot [1\{\text{SEM}_i > \alpha\} \times (-\Delta\text{Rank}_i)] + \beta_4 X_i + \varepsilon_i$$

where Y_i is occupation i's wage outcome,[3] $1\{\text{SEM}_i > \alpha\}$ indicates occupations above the skill evolution threshold α identified by Jenks natural breaks,[4] and $-\Delta\text{Rank}_i$ measures declining distinctiveness (positive values indicate convergence toward other occupations' skill profiles).[5] The coefficient β_3 captures how declining distinctiveness affects wages

[3] We estimate four specifications with different dependent variables: (1) P75 wage growth = $(\text{Wage}_{75,2018} - \text{Wage}_{75,2010})/\text{Wage}_{75,2010}$, (2) P50 wage growth, (3) P25 wage growth, and (4) wage spread change = P75 growth − P25 growth.

[4] The Jenks algorithm partitions the SEM distribution to minimize within-group variance, identifying a natural threshold separating substantial from minimal skill change occupations.

[5] ΔRank_i is the percentile rank change in the NSI distribution from 2010 to 2018, where NSI (Nearest

differently depending on the magnitude of skill change. The vector X_i includes employment growth and log median wage in 2010 as controls. For lower-wage occupations, we use $1\{\text{SEM}_i \leq \alpha\}$ to identify occupations with minimal skill change. Both specifications use robust standard errors (HC1).

TABLE 5. Within-Occupation Wage Dynamics in Higher-Wage Occupations

	(1)	(2)	(3)	(4)
	P75 Growth	P50 Growth	P25 Growth	Spread Change
$1\{\text{SEM} > \alpha\}$	0.0188	0.0165	0.0124	0.0064
	(0.0140)	(0.0133)	(0.0131)	(0.0111)
$-\Delta$Rank	-0.0259	-0.0083	0.0174	-0.0434**
	(0.0270)	(0.0308)	(0.0234)	(0.0218)
$1\{\text{SEM} > \alpha\} \times (-\Delta\text{Rank})$	0.0463	0.0160	-0.0017	0.0480*
	(0.0298)	(0.0332)	(0.0291)	(0.0257)
Employment Growth	-0.0534***	-0.0469**	-0.0138	-0.0396*
	(0.0205)	(0.0225)	(0.0300)	(0.0223)
Log(Median Wage 2010)	0.0140	0.0324	0.0036	0.0104
	(0.0315)	(0.0279)	(0.0293)	(0.0267)
Constant	-0.023	-0.236	0.086	-0.109
	(0.352)	(0.311)	(0.327)	(0.298)
N	108	108	108	108
R^2	0.106	0.088	0.021	0.095

Note: OLS with robust standard errors. Sample: occupations above 75th percentile of 2010 median wages. Spread Change = Δ(P75$-$P25). $1\{\text{SEM} > \alpha\}$ indicates occupations above the Jenks threshold for skill evolution magnitude. $-\Delta$Rank measures declining occupational distinctiveness (positive values = convergence toward other occupations). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 examines higher-wage occupations. Focus on the final column (Spread Change). The interaction term $1\{\text{SEM} > \alpha\} \times (-\Delta\text{Rank})$ shows a coefficient of $+0.048$ (SE=0.026, $p = 0.061$). Among occupations experiencing substantial skill change, those whose skill requirements became less distinct from other occupations saw wage spreads widen 4.8 percentage points more than occupations whose distinctiveness remained stable. This compares occupations with similar amounts of skill change but different patterns: some converged toward capabilities used in other professions, others evolved within their specialized domains.

The wage separation concentrates at the top. High earners (75th percentile) in occupations with declining distinctiveness saw wages grow 4.6 percentage points faster, while low earners (25th percentile) saw essentially no difference (-0.2pp). The top pulled away

neighbor Similarity Index) measures euclidean distance to the closest occupation in embedded skill space. Multiplying by -1 makes positive values correspond to declining distinctiveness.

TABLE 6. Within-Occupation Wage Dynamics in Lower-Wage Occupations

	(1)	(2)	(3)	(4)
	P75 Growth	P50 Growth	P25 Growth	Spread Change
$1\{\text{SEM} \leq \alpha\}$	−0.0077	−0.0018	0.0048	−0.0125*
	(0.0092)	(0.0106)	(0.0088)	(0.0073)
$-\Delta\text{Rank}$	0.0307	0.0170	0.0050	0.0258*
	(0.0232)	(0.0203)	(0.0202)	(0.0150)
$1\{\text{SEM} \leq \alpha\} \times (-\Delta\text{Rank})$	−0.0304	−0.0121	0.0117	−0.0422**
	(0.0249)	(0.0221)	(0.0227)	(0.0172)
Employment Growth	−0.0198**	−0.0302**	−0.0258***	0.0060
	(0.0077)	(0.0120)	(0.0061)	(0.0049)
Log(Median Wage 2010)	−0.0459***	−0.0540***	−0.0585***	0.0126
	(0.0102)	(0.0083)	(0.0092)	(0.0104)
Constant	0.686***	0.794***	0.851***	−0.165
	(0.108)	(0.091)	(0.099)	(0.110)
N	324	324	324	324
R^2	0.086	0.130	0.138	0.034

Note: OLS with robust standard errors. Sample: occupations at or below 75th percentile of 2010 median wages. Spread Change = Δ(P75−P25). $1\{\text{SEM} \leq \alpha\}$ indicates occupations below the Jenks threshold for skill evolution magnitude. $-\Delta$Rank measures declining occupational distinctiveness (positive values = convergence toward other occupations). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

from the bottom, creating the spread widening we observe.

Declining distinctiveness reveals what kind of skill change occurred. These occupations converged toward analytical and computational capabilities used across professional domains—skills that conflict with traditional expertise rather than building on it. Integrating quantitative modeling, data analysis, or algorithmic reasoning into professions built on relationship-based judgment or experiential knowledge creates cognitive distance. Some workers successfully acquired these cross-domain capabilities, developing expertise valuable beyond their original occupation. Others maintained traditional approaches. This uneven adaptation generated internal scarcity: only a subset mastered the new analytical skills employers increasingly valued, commanding wage premiums. In contrast, occupations maintaining stable distinctiveness experienced similar skill change magnitudes but through domain-specific tools that scaffolded onto existing expertise, enabling more uniform adaptation.

Table 6 examines lower-wage occupations. The interaction term $1\{\text{SEM} \leq \alpha\} \times (-\Delta\text{Rank})$ shows a coefficient of −0.042 (SE=0.017, $p = 0.014$). Among occupations experiencing no striking skill change, those whose skill requirements became less distinct from other occupations saw wage spreads compress 4.2 percentage points more than

occupations whose distinctiveness remained stable. This compares occupations with similar amounts of skill change but different boundary dynamics.

The compression concentrates at the top. High earners (75th percentile) in occupations with declining distinctiveness saw wages grow 3.0 percentage points slower, while low earners (25th percentile) saw essentially no difference (+1.2pp). The top compressed toward the bottom, narrowing the wage distribution.

Declining distinctiveness without substantial skill change signals erosion of specialized requirements rather than adoption of new capabilities. The convergence toward less distinct skill profiles indicates that occupation-specific expertise became less valuable. This eroded the knowledge barriers that previously protected workers with experience and specialized procedural knowledge, compressing returns to occupation-specific human capital. High earners who commanded premiums for occupation-specific expertise saw those advantages diminish as distinctiveness declined, while low earners—whose compensation never relied heavily on specialized knowledge—experienced minimal change.

The key validation comes from opposite effects of similar distinctiveness decline. Among high-wage occupations with substantial skill change, declining distinctiveness associates with 4.8 percentage point wage spread widening ($p = 0.061$). Among low-wage occupations with no striking skill change, similar magnitude distinctiveness decline associates with 4.2 percentage point compression ($p = 0.014$). The same geometric pattern—falling distinctiveness—produces large effects in opposite directions depending on transformation context and the relationship between new requirements and existing expertise.

This validates that our geometric measurement framework successfully identifies qualitatively distinct processes—skill change requiring integration of cognitively distant capabilities creating high-variance acquisition costs and internal scarcity, versus simplification eliminating protective knowledge barriers. The mechanism matters: declining distinctiveness signals entirely different worker pressures depending on whether it accompanies substantial skill change or minimal change, and whether baseline occupation-specific human capital is high or low.

The results prove robust to alternative specifications. Using a 70th percentile cutoff (N=130 higher-wage, N=302 lower-wage) yields consistent patterns: 4.7 percentage point widening for higher-wage occupations ($p = 0.041$) and 3.9 percentage point compression for lower-wage occupations ($p = 0.020$). The qualitative pattern—opposite effects of declining distinctiveness across transformation contexts—persists across specifications, confirming results capture genuine processes rather than artifacts of particular sample definitions.

5.4. Contributions and Limitations

Direct measurement using job posting embeddings reveals structured heterogeneity with clear economic significance. We provide systematic observation of how employer skill demands evolved across 514 occupations during 2010–2018, with labor market validation across 432 occupations representing 72% of U.S. employment.

The measurement infrastructure reveals that occupational transformation operates through qualitatively distinct mechanisms linked to acquisition cost structures. Cognitively distant capabilities create high-variance acquisition costs generating heterogeneous adaptation and wage separation through internal scarcity. Scaffolded capabilities create uniform acquisition costs enabling collective upgrading and maintaining stable wage structures. Technological simplification eliminates specialized knowledge, producing commodification and wage compression. These processes are validated through opposite wage dispersion effects from similar boundary characteristics.

These findings carry limitations inherent to occupation-level analysis. We cannot observe individual worker trajectories, distinguish selection from adaptation, or identify firm-specific strategies. The patterns provide strong evidence for qualitatively distinct processes while leaving micro-mechanisms for investigation using worker-level and firm-level data. Our contribution is demonstrating that comprehensive skill evolution measurement is feasible at economy-wide scale, that distinct mechanisms explain diverse patterns, and that the measurement apparatus successfully identifies occupations experiencing qualitatively different technological processes with different consequences for workers.

6. Conclusion

When technology reshapes skill demands, the nature of required capabilities varies fundamentally. Some new requirements prove cognitively distant from occupational expertise—demanding integration of conflicting paradigms about how to approach work. Others scaffold onto existing knowledge—extending capabilities within established professional frameworks. Still others represent simplification rather than enhancement—technology eliminating specialized knowledge requirements altogether. These differences create qualitatively distinct pressures for incumbent workers attempting to adapt. We identify three systematic processes through which skill change operates during technological transformation and validate that these processes generate divergent wage patterns, demonstrating they represent genuine mechanisms rather than taxonomic convenience.

The framework emerges from measuring skill evolution relationally across 514 occupations. We developed aligned embedding spaces from 191,919 job postings capturing employer requirements in 2010 and 2018, enabling observation of how occupations repositioned in skill space. Combining transformation magnitude with distinctiveness dynamics

reveals geometric signatures diagnostic of underlying skill content.

Our approach represents measurement-first discovery. We build infrastructure to observe skill transformation independently of labor market outcomes, identify systematic geometric patterns in how 514 occupations evolved, develop economic interpretations grounded in acquisition cost theory, and validate interpretations through wage dynamics. This arc—independent measurement enabling discovery and validated interpretation—demonstrates that comprehensive skill evolution analysis is feasible at economy-wide scale and reveals structured heterogeneity invisible in exposure-based approaches.

Three distinct processes emerge. Type 1 features high transformation alongside eroding distinctiveness—occupations converged toward capabilities valued across professional contexts, creating adoption challenges from cognitive distance between new requirements and existing expertise. Type 2 shows comparable transformation magnitudes but stable distinctiveness—new capabilities remained embedded within occupational domains, enabling scaffolded learning. Type 3 involves minimal transformation despite distinctiveness erosion—technology simplified rather than enhanced requirements, eliminating protective specialized knowledge. The critical insight: transformation operates through qualitatively distinct mechanisms, not scalar shifts. Technology doesn't uniformly increase or decrease skill requirements—it transforms them in fundamentally different ways. Capabilities can conflict with existing expertise (cognitive distance), build cumulatively on domain knowledge (scaffolded learning), or eliminate specialized requirements (simplification). Whether skills became less occupationally distinct tells us nothing about internal dynamics among workers without knowing what kinds of capabilities changed and their relationship to existing expertise. The qualitative nature of transformation—how new requirements relate to existing expertise—determines worker outcomes, not transformation magnitude alone.

The empirical validation confirms these represent economically meaningful processes. Distinctiveness decline in Type 1 contexts—high transformation toward cross-domain analytical capabilities—associates with 4.8 percentage point wage spread widening, while distinctiveness decline in Type 3 contexts—minimal transformation toward generic service work—associates with 4.2 percentage point compression. Validation comes from prediction-outcome alignment. The diagnostic framework—combining wage level, transformation magnitude, boundary dynamics, and neighbor identity—classifies occupations by mechanism type. Each mechanism generates distinct predictions: cognitive distance predicts widening through heterogeneous acquisition costs, scaffolded learning predicts stability through uniform costs, simplification predicts compression through eroded premiums. All three predictions materialize as expected, confirming the diagnostic framework successfully identifies which transformation process operated in each occupation.

Effects concentrate at the 75th percentile where theory predicts—internal scarcity generating premiums in one case, experience premiums eroding in the other. This pattern replicates across specifications and proves economically substantial, affecting occupations employing millions of workers. The mechanisms require continued investigation—our occupation-level analysis establishes systematic processes operate, but understanding individual trajectories, firm strategies, and temporal dynamics demands further study. We demonstrate these patterns link to wage dispersion through economically reasoned mechanisms generating predictions empirical evidence validates, not merely statistical associations between technological change and compensation.

These findings have implications for how policy approaches technological change. Skill building occurs through multiple channels—firm investments, public subsidies, individual worker decisions—and uniform policies assuming homogeneous technological pressure will be inefficient. Our diagnostic framework reveals distinct adaptation challenges: some occupations require incumbent workers to acquire capabilities within their roles, and for some (cognitive distance) this proves more challenging than others (scaffolded learning). Other occupations (simplification) may benefit more from policies enabling pathways to roles with better prospects than from retraining in eroded capabilities. The contribution is diagnostic not prescriptive—we identify which occupations face which pressures, enabling future work to evaluate targeted interventions rather than assuming uniform skill-biased technical change.

This measurement infrastructure enables research previously infeasible. Researchers can identify which occupations experienced which mechanism, enabling studies evaluating whether different policy approaches work for different transformation types. Firms can assess whether their workforce faces adaptation challenges requiring within-occupation skill building or whether mobility pathways might better serve workers. The framework provides diagnostic capacity—identifying which pressures workers face—that future work can use to evaluate intervention effectiveness rather than assuming technological change operates uniformly.

Fundamental questions about skill evolution lack systematic empirical investigation. Do skill transformations diffuse geographically, radiating from innovation centers, or emerge simultaneously? When do processes accelerate—gradually or through discrete shifts? How do industry contexts shape adaptation patterns—manufacturing versus services, regulated versus unregulated sectors? Do institutional arrangements mediate which mechanisms dominate—labor protections, educational systems, employment regulations? How do transformations vary within occupational hierarchies—senior versus junior practitioners, specialists versus generalists? Rich longitudinal data capturing skill requirements exist across countries, time periods, and contexts yet remain underutilized for studying these dynamics. Job postings, resumes, training records, professional certifications

combined with advancing methods for extracting semantic structure from text enable investigating skill evolution at economy-wide scale. Our analysis demonstrates this research direction proves productive, revealing patterns invisible to static task taxonomies or exposure-based approaches. The field would benefit from more systematic empirical investigation of skill change dynamics using rich observational data.

Understanding how technological transformation affects workers requires diagnosing what types of skill changes occur in specific contexts—whether capabilities prove cognitively distant or scaffolded, requirements enhanced or simplified, tasks divisible or indivisible. Three qualitatively distinct processes create systematically different adaptation pressures. The diagnostic framework we developed identifies which process operates through geometric signatures in how skill requirements evolve relationally, validated across 432 occupations representing 72% of the U.S. workforce. This enables mechanism-specific understanding rather than uniform treatment of technological change. Rich representations from large-scale empirical data provide foundation for systematic study of skill evolution as economies transform.

References

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum. 2020. "The Evolution of Work in the United States." *American Economic Journal: Applied Economics* 12 (2): 1–34.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *The American Economic Review* 103 (5): 1553–1597.

Autor, David H., David Dorn, and Gordon H. Hanson. 2015. "Untangling Trade and Technology: Evidence from Local Labour Markets." *The Economic Journal* 125 (584): 621–646.

Bocquet, Léonard. 2024. "The Network Origin of Slow Labor Reallocation.". `https://www.parisschoolofeconomics.eu/app/uploads/2024/10/BOCQUET-Leonard-JMP.pdf`. Job Market Paper, Paris School of Economics.

Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65 (1): 83–108.

Chen, Ting, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. "A Simple Framework for Contrastive Learning of Visual Representations." In *Proceedings of the 37th International Conference on Machine Learning*, vol. 119 of ICML'20: 1597–1607: PMLR.

Cortes, Guido Matias, Nir Jaimovich, and Henry E. Siu. 2017. "Disappearing Routine Jobs: Who, How, and Why?" *Journal of Monetary Economics* 91: 69–87.

Dell, Melissa. 2025. "Deep Learning for Economists." *Journal of Economic Literature* 63 (1): 5–58.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics* 132 (4): 1593–1640.

Fan, Tianyu. 2025. “The Labor Market Incidence of New Technologies.”. https://ssrn.com/abstract=5160195. Available at SSRN.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses.” *Strategic Management Journal*.

Gabaix, Xavier, Ralph S. J. Koijen, Robert J. Richmond, and Motohiro Yogo. 2025. “Asset Embeddings.” Working Paper No. 2025-51, Becker Friedman Institute for Economics, University of Chicago. Chicago, IL. Available at https://bfi.uchicago.edu.

Gathmann, Christina, and Uta Schönberg. 2010. “How General Is Human Capital? A Task-Based Approach.” *Journal of Labor Economics* 28 (1): 1–49.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. “Text as Data.” *Journal of Economic Literature* 57 (3): 535–574.

Hamilton, William L., Jure Leskovec, and Dan Jurafsky. 2016. “Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change.” In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,: 1489–1501: Association for Computational Linguistics.

Kambourov, Gueorgui, and Iourii Manovskii. 2009. “The Rising Return to Occupational Experience.” *The Review of Economic Studies* 76 (2): 731–759.

Khosla, Prannay, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. “Supervised Contrastive Learning.” In *Advances in Neural Information Processing Systems*, vol. 33: 18661–18673: Curran Associates, Inc.

Kleiner, Morris M., and Alan B. Krueger. 2013. “Analyzing the Extent and Influence of Occupational Licensing on the Labor Market.” *Journal of Labor Economics* 31 (S1): S173–S202.

Lise, Jeremy, and Fabien Postel-Vinay. 2020. “Multidimensional Skills, Sorting, and Human Capital Accumulation.” *American Economic Review* 110 (8): 2328–2376.

Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.

Marinescu, Ioana, and Ronald Wolthoff. 2020. “Opening the Black Box of the Matching Function: The Power of Words.” *Journal of Labor Economics* 38 (2): 535–568.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2020. “Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?” *Review of Economics and Statistics* 102 (4): 793–805.

del Rio-Chanona, R. Maria, Penny Mealy, Mariano Beguerisse-Díaz, François Lafond, and J. Doyne Farmer. 2021. “Occupational mobility and automation: a data-driven network model.” *Journal of the Royal Society Interface* 18 (183).

Schroff, Florian, Dmitry Kalenichenko, and James Philbin. 2015. “FaceNet: A Unified Embedding for Face Recognition and Clustering.” In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*,: 815–823.

Vafa, Keyon, Susan Athey, and David M. Blei. 2024. “Estimating Wage Disparities Using Foundation Models.”. https://arxiv.org/abs/2409.09894.

Wang, Liang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. “Text Embeddings by Weakly-Supervised Contrastive Pre-training.”

Webb, Michael. 2019. “The Impact of Artificial Intelligence on the Labor Market.” SSRN 3482150,

Social Science Research Network (SSRN). Available at: https://ssrn.com/abstract=3482150.

Appendix A: Technical Implementation Details

A.1. Data Processing and Corpus Construction

Quality assessment algorithms evaluate postings using semantic coherence metrics, grammatical structure analysis, and content relevance scoring. Specific filtering criteria include:

- text length bounds (50–5,000 words),
- quality score thresholds (minimum 4.0 on a 7-point scale), and
- duplicate detection using both exact string matching and cosine similarity thresholds.

Stratified sampling employs quality-weighted selection when occupations exceed 200 postings, with sampling probabilities proportional to quality scores and inverse document frequency (IDF) weights to ensure representative coverage across firm types and geographic regions.

The semantic coherence evaluation uses automated linguistic analysis to identify postings with clear job descriptions, coherent skill requirements, and minimal formatting errors. Grammatical structure analysis employs syntactic parsing to ensure postings contain complete sentences and logical organization of requirements. Content relevance scoring validates that postings contain substantive skill and qualification information rather than purely administrative or legal text.

Duplicate detection operates at two levels: exact string matching identifies identical postings across different sources or time periods, while semantic similarity detection using cosine similarity thresholds (> 0.95) identifies near-duplicates with minor variations in wording or formatting. The stratified sampling algorithm ensures balanced representation across firm sizes, geographic regions, and industry sectors when occupations exceed the 200-posting threshold.

A.2. Model Architecture and Training Details

The two-phase training process begins with a warm-up stage using mean squared error (MSE) loss for 2 epochs at learning rate $\eta = 1 \times 10^{-4}$, training only AttnNet parameters against pre-computed occupation centroids. This initialization ensures attention weights begin from meaningful baselines rather than random values.

Formally, the MSE loss is defined as:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^{N} \left\| \hat{\mathbf{h}}_i - \mathbf{c}_i \right\|^2, \tag{A.1}$$

where $\hat{\mathbf{h}}_i$ denotes the predicted representation of occupation i and $\mathbf{c}_i$ is the corresponding centroid.

The main contrastive training employs triplet margin loss with margin $\alpha = 0.5$ for 25 epochs at learning rate $\eta = 3 \times 10^{-5}$, with batch size 32 triplets and weight decay $\lambda = 1 \times 10^{-5}$, using the Adam optimizer with default momentum parameters. The triplet loss is given by:

$$\mathcal{L}_{\text{triplet}} = \frac{1}{M} \sum_{j=1}^{M} \max\left(0,\ \|\mathbf{h}_j^a - \mathbf{h}_j^p\|^2 - \|\mathbf{h}_j^a - \mathbf{h}_j^n\|^2 + \alpha\right), \tag{A.2}$$

where each triplet consists of an anchor vector $\mathbf{h}_j^a$, a positive vector $\mathbf{h}_j^p$ (same occupation), and a negative vector $\mathbf{h}_j^n$ (different but related occupation).

Padding and masking implementation uses zero-vector padding for occupations with fewer than 32 postings, with binary masks ensuring padded positions receive attention logits of -1×10^9 before softmax normalization, effectively zeroing their contribution. CompressNet skip connections add the direct linear projection ($1024 \to 128$) to the main pathway output using element-wise addition. Both AttnNet and CompressNet components use Xavier initialization for linear layers, with bias terms initialized to zero.

Training convergence monitoring tracks triplet loss reduction, attention weight entropy (ensuring attention does not collapse to single postings), and validation performance on held-out occupation pairs. Early stopping employs a patience of 5 epochs if validation loss stops improving.

Appendix B: Anchor Occupation Selection and Methodology

B.1. Selection Criteria and Conceptual Categories

We selected anchor occupations based on three criteria designed to identify occupations with inherently stable skill requirements: (1) licensing requirements that legally constrain skill evolution, (2) specialized training programs that create stable professional boundaries, and (3) fundamental service needs that remain consistent despite technological change.

Our 21 verified anchors span five conceptual categories:

- **Healthcare Practitioners:** Dentists (29-1021), Pharmacists (29-1051), Physical Therapists (29-1123), Veterinarians (29-1131), Speech-Language Pathologists (29-1127).
- **Skilled Trades:** Electricians (47-2111), Plumbers, Pipefitters, and Steamfitters (47-2152), Carpenters (47-2031), HVAC Mechanics (49-9021), Brickmasons and Blockmasons (47-2021).
- **Protective and Civil Services:** Police and Sheriff's Patrol Officers (33-3051), Firefighters (33-2011), Elementary School Teachers (25-2021), Postal Service Clerks (43-5051).
- **Transportation:** Heavy and Tractor-Trailer Truck Drivers (53-3032), Airline Pilots,

Copilots, and Flight Engineers (53-2011), Railroad Conductors and Yardmasters (53-4031).

- **Human and Personal Services:** Hairdressers, Hairstylists, and Cosmetologists (39-5012), Chefs and Head Cooks (35-1011), Medical Assistants (31-9092), Security Guards (33-9032).

B.2. Verification Process and Data Quality

From 23 candidate anchors, verification confirmed that 21 appeared in both 2010 and 2018 datasets with sufficient data quality for reliable prototype generation. Two candidates (Dental Hygienists, 29-2021; Postal Service Mail Carriers, 43-5052) were excluded due to insufficient representation in one or both periods.

The verification process required each anchor occupation to have at least 50 high-quality postings per year and quality scores above the dataset median. The final anchor set represents approximately 4% of our total sample (21 of 514 occupations), providing sufficient constraint for alignment without over-determining the transformation. This proportion follows established practices in computer vision and signal processing for landmark-based registration, where 3–5% of features typically serve as reliable anchors.

B.3. Mathematical Implementation of Alignment

Orthogonal Procrustes analysis minimizes the objective function:

$$\mathbf{Q}^* = \arg\min_{\mathbf{Q}} \|\mathbf{X}\mathbf{Q} - \mathbf{Y}\|_F^2 \quad \text{subject to} \quad \mathbf{Q}^\top\mathbf{Q} = \mathbf{I}, \tag{B.1}$$

where $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{21\times 128}$ represent the matrices of anchor occupation prototypes from 2010 and 2018, respectively, and $\mathbf{Q}$ is the orthogonal transformation matrix.

The solution computes the singular value decomposition

$$\mathbf{X}^\top\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top,$$

and constructs

$$\mathbf{Q} = \mathbf{V}\mathbf{U}^\top,$$

ensuring $\det(\mathbf{Q}) = 1$ for proper rotation without reflection.

The algorithm centers both point clouds at their centroids before computing the transformation, then applies the same translation to all occupations. This ensures the alignment preserves both the relative positions and the absolute positioning of the occupational clusters.

B.4. Alignment Implementation and Validation

The Procrustes alignment process operates on the 21 verified anchor occupations, computing the optimal orthogonal transformation matrix $\mathbf{Q}^*$ to minimize distances between corresponding anchor prototypes across 2010 and 2018. Implementation uses singular value decomposition to ensure the transformation preserves distances and angles within the 2018 space while reorienting it to match the 2010 coordinate system.

Alignment validation employs distance-based metrics to confirm successful transformation. The disparity between anchor sets after alignment is 0.4946, indicating the root mean square distance between corresponding anchor pairs in the aligned space. More critically, the average distance between anchor pairs before alignment (1.1663) reduces substantially to 0.2120 after standardization and alignment, representing an 82% reduction in anchor disparity.

This dramatic reduction in distances between theoretically stable occupations provides strong evidence that the linear alignment assumption is appropriate and that the transformation captures genuine occupational correspondences rather than arbitrary coordinate matching.

The successful alignment enables reliable temporal comparison by establishing a common coordinate system where changes in occupational positions represent genuine skill requirement transformations rather than methodological artifacts. All 514 occupations in the 2018 dataset are transformed using the same rotation matrix computed from the anchor occupations, preserving the internal geometric structure of the 2018 skill space while making it directly comparable to the 2010 reference frame.