



ARTICLE



STRATEGIC WORKFORCE INTELLIGENCE FOR AUTOMOTIVE EDUCATION: A COMPETITIVE INTELLIGENCE FRAMEWORK BASED ON LABOUR-MARKET ANALYTICS

INTELIGÊNCIA ESTRATÉGICA DA FORÇA DE TRABALHO PARA A EDUCAÇÃO AUTOMOTIVA: UM FRAMEWORK DE INTELIGÊNCIA COMPETITIVA BASEADO EM ANALYTICS DO MERCADO DE TRABALHO

1Ruonan Wang: Faculty of Education, National University of Malaysia.

Department of Foreign Languages and Cultures, Geely University of China, Chengdu, China.

ORCID: <https://orcid.org/0009-0008-9524-2307>

2*Nur Ehsan Bin Mohd Said: Faculty of Education, National University of Malaysia. ORCID:

<https://orcid.org/0000-0002-2891-327X>

Corresponding Author:

Nur Ehsan Bin Mohd Said
E-mail: nurehsan@ukm.edu.my

Editor in chief

Altieres De Oliveira Silva
Alumni.In Editors

How to cite this article:

Wang, R., & Said, N. E. B. M. (2026). Strategic Workforce Intelligence for Automotive Education: A Competitive Intelligence Framework Based on Labour-Market Analytics. *Journal of Sustainable Competitive Intelligence*, 16, e0673. <https://doi.org/10.37497/eagleSustainable.v16i.673>

ABSTRACT

Purpose: The aim of this study is to develop a Competitive Intelligence (CI) framework that transforms labour-market signals into Strategic Workforce Intelligence (SWI) for the curriculum strategy of Chinese universities offering automotive English for Specific Purposes (ESP) tracks, positioning CI as an institutional capability for sustaining competitive advantage in graduate employability and institutional responsiveness to industrial change.

Methodology/Approach: A quantitative, secondary-data design grounded in the publicly available Job-SDF Chinese-recruitment benchmark covering 36 months of monthly hiring demand between January 2021 and December 2023 was adopted. The pipeline isolates an industrial proxy of six L1 occupations representing 80.0 % of platform-wide hiring volume. Communication-relevant skills are separated from industrial domain-specific skills via a normalised cross-occupation entropy criterion, and trajectory dynamics are characterised through OLS regression, Spearman correlation, compound annual growth rates, structural-break analysis, and K-means clustering. The Resource-Based View, the Knowledge-Based View, and Dynamic Capabilities theory anchor the integrated four-tier conceptual framework.

Findings: The analytical sample contained 45,046,064 recruitment observations across 1,651 skills (18 communication and 1,633 industrial). Recruitment demand grew at compound annual rates of 32.3 % for industrial and 21.8 % for communication skills (both statistically significant), and 33.4 % of skills exhibited structural breaks. K-means clustering identified a slow-rising stable cluster (the curriculum spine, dominated by communication skills) and a high-acceleration industrial cluster (the agile periphery). Occupation-level mapping revealed pronounced asymmetries that justify differentiated curriculum tracks aligned with destination occupations.

Originality/Relevance: The paper extends Competitive Intelligence from corporate strategy into higher-education curriculum governance and formalises Strategic Workforce Intelligence as a new construct anchored in RBV, KBV, and Dynamic Capabilities theory. Methodologically, it grounds curriculum decisions in 45 million labour-market observations through a privacy-respecting entropy filter. Practically, it equips Chinese universities with a CI-driven decision architecture for sustaining industry-aligned ESP programmes and provides a reproducible analytical pipeline transferable to other industry-aligned academic programmes.

Keywords: Competitive Intelligence. Strategic Workforce Intelligence. Labour-Market Analytics. Institutional Competitiveness. Curriculum Strategy. English for Specific Purposes. Automotive Industry. Chinese Universities.



DOI: <https://doi.org/10.37497/eagleSustainable.v16i.673>





RESUMO

Objetivo: O objetivo deste estudo é desenvolver um framework de Inteligência Competitiva (CI) que transforme sinais do mercado de trabalho em Inteligência Estratégica da Força de Trabalho (SWI) para a estratégia curricular de universidades chinesas que oferecem programas de Inglês para Fins Específicos (ESP) voltados ao setor automotivo, posicionando a CI como uma capacidade institucional para sustentar vantagem competitiva na empregabilidade dos graduados e na responsividade institucional às mudanças industriais.

Metodologia/Abordagem: Foi adotado um desenho quantitativo com dados secundários, fundamentado no benchmark público chinês de recrutamento Job-SDF, cobrindo 36 meses de demanda mensal de contratação entre janeiro de 2021 e dezembro de 2023. O pipeline isola um proxy industrial composto por seis ocupações L1 representando 80,0% do volume total de contratações da plataforma. Competências relacionadas à comunicação são separadas de competências específicas do domínio industrial por meio de um critério normalizado de entropia entre ocupações, enquanto a dinâmica das trajetórias é caracterizada por regressão OLS, correlação de Spearman, taxas compostas anuais de crescimento, análise de quebras estruturais e clustering K-means. A Visão Baseada em Recursos (RBV), a Visão Baseada no Conhecimento (KBV) e a teoria das Capacidades Dinâmicas sustentam o framework conceitual integrado de quatro níveis.

Resultados: A amostra analítica continha 45.046.064 observações de recrutamento distribuídas em 1.651 competências (18 de comunicação e 1.633 industriais). A demanda de recrutamento cresceu a taxas compostas anuais de 32,3% para competências industriais e 21,8% para competências de comunicação (ambas estatisticamente significativas), enquanto 33,4% das competências apresentaram quebras estruturais. O clustering K-means identificou um cluster estável de crescimento lento (a espinha curricular, dominada por competências de comunicação) e um cluster industrial de alta aceleração (a periferia ágil). O mapeamento ocupacional revelou assimetrias significativas que justificam trilhas curriculares diferenciadas alinhadas às ocupações de destino.

Originalidade/Relevância: O artigo amplia a aplicação da Inteligência Competitiva da estratégia corporativa para a governança curricular do ensino superior e formaliza a Inteligência Estratégica da Força de Trabalho como um novo constructo ancorado nas teorias RBV, KBV e Capacidades Dinâmicas. Metodologicamente, fundamenta decisões curriculares em 45 milhões de observações do mercado de trabalho por meio de um filtro de entropia que preserva a privacidade. Na prática, oferece às universidades chinesas uma arquitetura decisória orientada por CI para sustentar programas ESP alinhados à indústria, além de fornecer um pipeline analítico reproduzível e transferível para outros programas acadêmicos orientados à indústria.

Palavras-chave: Inteligência Competitiva. Inteligência Estratégica da Força de Trabalho. Analytics do Mercado de Trabalho. Competitividade Institucional. Estratégia Curricular. Inglês para Fins Específicos. Indústria Automotiva. Universidades Chinesas.



1. INTRODUCTION

Higher-education institutions increasingly compete on the strategic responsiveness of their academic offer to industrial change (Madureira et al., 2023; Nie & Vorawattanachai, 2026). In a rapidly evolving economy such as China, the speed at which a university can detect, interpret, and translate labour-market signals into curriculum decisions has become a decisive determinant of graduate employability and institutional reputation. Competitive Intelligence (CI), originally conceptualised as the legal and ethical collection, analysis, and dissemination of information to support strategic decision-making in firms (Calof & Colton, 2024; Maluleka & Chummun, 2023), has progressively been recognised as a transferable capability through which universities can monitor their external environment, anticipate skill demand, and align academic programmes with industrial requirements (Hakmaoui et al., 2022; Nie & Vorawattanachai, 2026). This paper situates CI as the central construct and develops a Strategic Workforce Intelligence framework that links industry signals, intelligence processing, curriculum strategy, and institutional competitive advantage.

The empirical site is the Chinese automotive industry, which has undergone a major transformation over the past five years, with domestic manufacturers capturing more than half of the local market and Chinese electric-vehicle exports rising to global prominence (Tian et al., 2024). This expansion has produced a corresponding surge in workforce demand, and skill shortages have become particularly acute in the smart-vehicle and electric-vehicle segments where engineers, manufacturing technicians, and back-office staff routinely receive multiple competing offers (Song & Xu, 2024). Within this growth context, English language competence is no longer a peripheral asset for automotive professionals but a precondition for participating in international supply chains, communicating with multinational partners, and accessing technical documentation (Çal et al., 2023). Universities offering automotive ESP programmes therefore face a strategic challenge: their decisions about curriculum content, occupation tracks, and content-refresh cadence must respond to rapidly shifting industry signals if their graduates are to remain competitive.

Translating industry signals into strategic curriculum decisions has, however, traditionally relied on small-sample needs-analysis methods (interviews and questionnaires with employers and graduates) that are slow to update, limited in scale, and subject to sampling bias (Ellederová & Denysenko, 2025; Tividad, 2024). These limitations create a structural lag between industrial change and curriculum response, leaving graduates with skill profiles already obsolete by the time they enter the labour market (Mao & Zhou, 2024). Two parallel developments now permit a different approach. First, the digitisation of recruitment in China has produced a substantial data resource: millions of job advertisements posted to online platforms each month, each enumerating the specific skills employers are willing to pay for (Chen et al., 2024). Second, recent advances in CI theory and practice have extended the intelligence cycle (planning, collection, analysis, dissemination, and decision) into education-administration applications (Schmitt, 2023; Nie & Vorawattanachai, 2026). Their convergence creates an opportunity to operationalise CI as the institutional decision architecture that converts continuous labour-market signals into curriculum strategy.

From a theoretical standpoint, the proposed framework is anchored in three complementary strategic-management traditions. The Resource-Based View (Helfat et al., 2023) treats curriculum responsiveness as a valuable, rare, and difficult-to-imitate capability that confers sustained competitive advantage on universities. The Knowledge-Based View (Madureira et al., 2023) frames the intelligence cycle as a knowledge-creation process that converts external data into actionable institutional knowledge. The Dynamic Capabilities perspective (Teece, 2023) captures the sensing, seizing, and reconfiguring activities through which universities continuously realign their academic offer with environmental change. CI is the institutional mechanism that operationalises these capabilities in the higher-education context.

This paper exploits the labour-market analytics opportunity for the case of automotive ESP at Chinese universities, treating it as a Strategic Workforce Intelligence problem. Using the Job-SDF benchmark, a publicly released dataset of monthly skill-demand sequences derived from Chinese recruitment platforms, the study constructs an industrial-occupation proxy for the automotive sector, separates communication-relevant skills from technical skills using an entropy-based criterion, and traces their joint demand evolution over 36 months. The analytical findings are then translated into a CI-driven decision architecture for Chinese universities seeking to bridge the gap between automotive industry needs and English-language education.

Three research questions guide the analysis. (RQ1) How can the CI cycle be operationalized, based on large-scale Chinese labour-market data, to produce strategic workforce intelligence for automotive ESP curriculum decisions? (RQ2) What is the volume, composition, and growth trajectory of automotive-relevant skill demand between 2021 and 2023, and how does the demand for communication and transferable skills compare with that for industrial domain-specific skills? (RQ3) Which CI-driven design principles emerge



from this evidence as instruments for strengthening institutional competitiveness and graduate employability in Chinese universities?

The novelty and originality of the paper rest on four contributions. First, it repositions CI as an institutional strategic capability for higher-education curriculum governance, rather than as a methodological support tool, and integrates it explicitly with RBV, KBV, and Dynamic Capabilities theory (Helfat et al., 2023; Teece, 2023). Second, it operationalises a Strategic Workforce Intelligence pipeline grounded in 45 million Chinese labour-market observations, extending the application of CI from corporate strategy into educational decision systems. Third, it proposes a reproducible entropy-based procedure for separating communication or transferable skills from industrial or domain-specific skills in privacy-protected job-ad datasets, a procedure that does not require access to skill names. Fourth, it produces empirically grounded, CI-derived design principles (differentiated occupation tracks, a stable spine plus an agile periphery, and routine structural-break monitoring) that translate the analytical findings into actionable strategic decisions for Chinese universities.

The remainder of the paper is organised as follows. Section 2 develops the theoretical framework, beginning with CI as a strategic capability for institutional decision-making and continuing through ESP and skill-demand analytics. Section 3 details the data source, the preprocessing and filtering pipeline, the analytical framework, and methodological limitations. Section 4 reports the descriptive, inferential, and clustering results organised around the three research questions, and discusses their strategic implications. Section 5 presents the conclusion, including theoretical contributions to the field of Sustainable CI, managerial implications for higher-education leaders, and avenues for future research.

2. THEORETICAL FRAMEWORK

2.1 Competitive Intelligence and Strategic Workforce Intelligence

Competitive Intelligence is conventionally defined as the legal and ethical collection, analysis, and dissemination of information about an organisation's competitors and external environment to support strategic decision-making (Madureira et al., 2023). The CI process is typically organised as a four-stage cycle planning, collection, analysis, and dissemination that converts raw data into actionable intelligence (Maluleka & Chummun, 2023). What distinguishes CI from generic data analysis is its decision orientation: the analytical work is governed by clearly identified key intelligence topics that map directly onto the choices the organisation must make (Mashau et al., 2025). Recent contributions reposition CI as a forward-looking anticipatory system that creates organisational knowledge by integrating environmental scanning with strategic foresight (Calof & Colton, 2024; Hakmaoui et al., 2022).

Three strategic-management theories explain why CI confers competitive advantage when transferred from the corporate domain to higher education. The Resource-Based View argues that sustained competitive advantage derives from valuable, rare, and difficult-to-imitate resources and capabilities (Helfat et al., 2023). For a university, the ability to systematically convert labour-market signals into industry-aligned curriculum decisions is precisely such a capability: it is path-dependent, organisationally embedded, and not easily duplicated by competitors operating with conventional needs-analysis methods. The Knowledge-Based View extends this logic by treating the firm or in this case the institution as a system of knowledge creation, transfer, and integration (Madureira et al., 2023). CI provides the formal architecture through which tacit signals from the external environment are converted into codified, decision-relevant institutional knowledge. The Dynamic Capabilities perspective adds the temporal dimension: in turbulent environments, organisations sustain advantage by sensing opportunities and threats, seizing them through resource reconfiguration, and continuously transforming their internal routines (Teece, 2023). A CI capability is the operational manifestation of dynamic capabilities in higher-education curriculum management.

Building on these foundations, this paper introduces the construct of Strategic Workforce Intelligence (SWI) as the application of CI to the workforce-development decisions of universities. SWI is defined here as the institutional capability of continuously gathering, analysing, and using labour-market data to align academic programmes with the evolving skill profile demanded by industry. Whereas conventional needs analysis treats labour-market signals as a one-off input, SWI treats them as a permanent monitoring stream and translates them into a portfolio of strategic decisions: which occupation tracks to differentiate, which skill modules to refresh, and which demand discontinuities warrant out-of-cycle curriculum interventions. SWI thus extends CI into the higher-education domain in the same way that talent-analytics applications have extended human-resource intelligence into corporate workforce planning (Schmitt, 2023; Tzimas et al., 2024).

CI applied to higher education has already produced tangible pipelines: skill extraction from job



listings, clustering of skills into competency groups, and forecasting of demand curves at the occupation level (Hiltz et al., 2026). These pipelines lend an evidence-based answer to the question that needs analysis has been asking since the 1980s: what should educators teach (Dou et al., 2023)? Recent work explicitly conceptualises CI as a strategic capability that improves the quality and timeliness of decisions in vocational and higher-education administration (Nie & Vorawattanachai, 2026). The present paper builds on this line of research and extends it to the specific case of automotive ESP curriculum strategy.

2.2 English for Specific Purposes and Industry Alignment

ESP is widely understood not as a fixed body of language but as an approach to teaching driven by the systematic analysis of learners' target needs (Dou et al., 2023). The field has developed three complementary lenses for these needs: target-situation analysis (what learners must do in their target professional contexts), present-situation analysis (the language they currently command), and learning-situation analysis (their preferred modes of learning) (Mao & Zhou, 2024). The interplay of these lenses produces a curriculum specification that is pragmatically grounded and pedagogically realistic (Ellederová & Denysenko, 2025).

In the Chinese tertiary context, college English has progressively shifted toward ESP under the pressure of internationalisation and the post-WTO labour market (Dou et al., 2024). The 2017 Handbook of College English Teaching officially divides college English into general English, ESP, and cross-cultural communication, and within ESP further distinguishes English for Academic Purposes from English for Occupational Purposes (Dou et al., 2024). Empirical evaluations of Chinese ESP graduates have, however, consistently revealed mismatches between curriculum content and workplace requirements, indicating that ESP curricula in China have been updated too slowly relative to changes in the labour market (Song & Xu, 2024). From a CI standpoint, these mismatches are intelligence failures: signals from the external environment are not being converted into timely curricular decisions.

These pressures are heightened in the automotive industry, which is characterised by a high level of technical vocabulary (in vehicle design, manufacturing, electrification, and quality management) and a high level of cross-cultural communication needs (in supply-chain coordination, sales negotiation, and customer service). Needs-analysis studies of automotive technology programmes consistently report that students need a combination of technical English vocabulary and oral and written communication skills that traditional college-English courses do not adequately address (Fielden Burns & Rico García, 2022). The growth of the Chinese electric-vehicle and battery industry has further increased demand for hybrid technical-plus-communicative skill profiles that ESP curricula must produce (Song & Xu, 2024). For a university to remain competitive in this segment, ESP curriculum decisions must be informed by a CI capability that monitors industry signals continuously rather than periodically.

2.3 Data-Driven Skill-Demand Analytics

The empirical foundation of SWI is the rapidly growing field of skill-demand analytics, in which structured records derived from millions of online job advertisements are used to estimate the current and future demand of individual skills at fine granularities (Rahhal et al., 2024). The Job-SDF benchmark, published by Chen and colleagues in 2024, is one of the largest publicly available datasets of this kind, providing monthly skill-demand sequences across multiple granularity levels (L1 occupation, L2 occupation, company, region) and supplementing the demand series with skill co-occurrence graphs and structural-break flags (Chen et al., 2024). Comparable resources in non-Chinese contexts include the Lightcast and Burning Glass corpora used in international labour-market intelligence research (Tzimas et al., 2024).

A consistent finding across this literature is that skill demand is heavily long-tailed: a small subset of skills accounts for most of the demand, and the remainder forms a slowly decaying tail of low-frequency skills (Weichselbraun et al., 2024). Skill-demand series are also non-stationary; structural breaks (abrupt changes in mean or trend) are a persistent feature, particularly during periods of technological transition such as the current shift from internal-combustion to electric vehicles (Lassébie & Quintini, 2022). For a CI-based curriculum strategy, structural breaks pose a specific challenge: they introduce timing risk that static curricula cannot accommodate, and they argue for routine intelligence updates rather than one-off needs analyses.

A separate methodological strand uses representation-learning approaches to forecast skill demand from text-rich job-advertisement data (Napierala & Kvetan, 2023). Although these approaches achieve



impressive predictive accuracy, they typically require labelled training data and entity-name access that the Job-SDF benchmark deliberately withholds for privacy reasons (Chen et al., 2024). The present paper therefore adopts a privacy-respecting analytical pathway that uses the demand sequences and structural-break flags directly, supplemented by an entropy-based skill classification that does not require skill names.

2.4 Comparison of Prior Approaches and Integrated Conceptual Framework

To position the present study against prior work, Table 1 summarises eight representative studies spanning ESP foundations, automotive ESP, Chinese-tertiary curriculum, skill forecasting, and Competitive Intelligence in education. The comparison highlights three persistent gaps in the literature that the present paper addresses: (i) the predominance of small-sample qualitative needs-analysis methods in the ESP tradition; (ii) the absence of a curriculum-strategy framing in large-scale skill-forecasting datasets; and (iii) the disconnection between Competitive Intelligence research and language-education practice.

Study	Domain / Sector	Data source	Method	Sample / Coverage	Key limitation
Dou et al. (2023)	ESP needs analysis	Theory & cases	Conceptual	N/A	No empirical labour-market grounding
Mao & Zhou (2024)	Art-design ESP in China	Survey of teachers	Mixed methods	6 universities	Single discipline, no industry data
Fielden Burns & Rico García (2022)	Engineering ESP curriculum	Industry survey	Needs analysis	< 200 students	Small, single-school sample
Çal et al. (2023)	Workplace English	Stakeholder interviews	Qualitative	Engineering students	No quantitative skill-demand link
Dou et al. (2024)	Discipline-specific ESP	Survey + interviews	Mixed methods	6 universities	Single discipline, no industry data
Chen et al. (2024)	General skill forecasting	Job-SDF Chinese recruitment	Time-series benchmarking	14 occ. × 2,335 skills	No curriculum-strategy framing
Tzimas et al. (2024)	Labour-market intelligence	EU & Greek job ads	Skill-graph analytics	≈1M ads	Non-Chinese, non-curriculum-specific
Nie & Vorawattanachai (2026)	Vocational-education admin	Mixed-methods survey	CI as capability	n = 75 + 15	No labour-market signal data

Table 1. Summary of representative prior studies in ESP, automotive ESP, Chinese-tertiary curriculum, skill forecasting, and Competitive Intelligence in education.

The present paper addresses these gaps simultaneously. It grounds an ESP curriculum-strategy study in a 45-million-row Chinese recruitment panel; it repurposes a benchmark originally built for general skill forecasting toward a curriculum-strategy target; and it adapts CI methodology to a higher-education decision context anchored in RBV, KBV, and Dynamic Capabilities theory.

Drawing the three theoretical strands together, this paper proposes a four-tier conceptual model that links industry signal to curriculum strategy and institutional competitiveness through a CI layer. The first tier is the industry signal: the stream of job advertisements posted on Chinese recruitment platforms, which collectively encodes employers' revealed preferences for specific skill profiles (Chen et al., 2024). The second

tier is the CI process: a sequence of cleaning, filtering, pattern-recognition, and forecasting operations that converts the raw signal into actionable intelligence on skill-demand levels, growth rates, and structural breaks (Calof & Colton, 2024; Madureira et al., 2023). The third tier is the curriculum-strategy layer in which intelligence outputs are translated into strategic decisions: the differentiation of occupation-specific tracks, the prioritisation of communication-skill content, the timing of content-refresh cycles, and the allocation of teaching resources (Dou et al., 2023). The fourth tier is the institutional outcome: graduates whose competencies are aligned with empirically validated industry demand and a university whose responsiveness to industrial change becomes a source of sustainable competitive advantage (Helfat et al., 2023; Teece, 2023). A feedback loop returns graduate-employment outcomes and institutional performance metrics to the industry-signal layer, allowing the framework to recalibrate continuously.

This integrated framework distinguishes the present study from purely qualitative needs-analysis studies, which typically operate at the third tier alone, and from purely technical skill-demand forecasting studies, which typically stop at the second tier. By spanning all four tiers explicitly and embedding them in CI theory, the paper makes the connection between labour-market analytics and institutional curriculum strategy operational. Figure 1 visualises the framework.

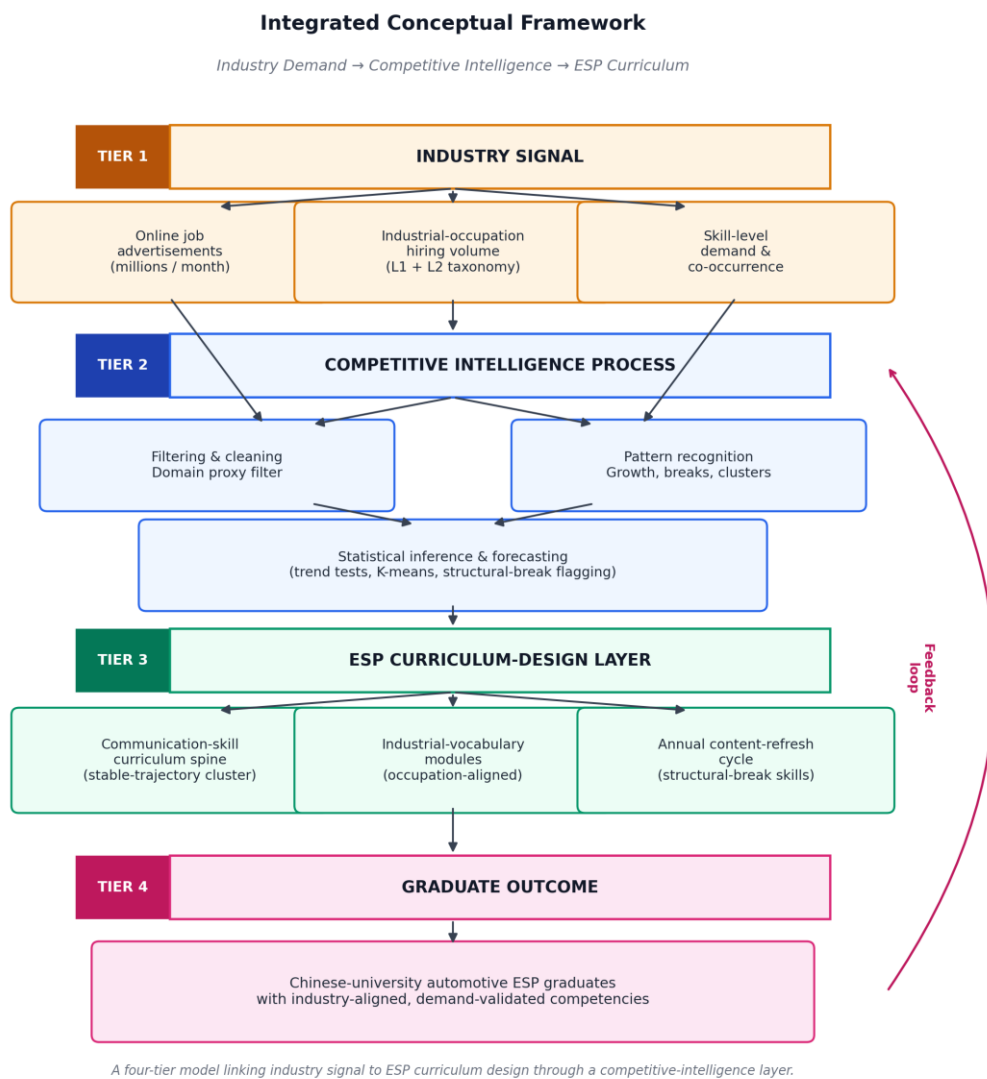


Figure 1. Four-tier conceptual framework linking industry signal, Competitive Intelligence processing, curriculum strategy, and institutional competitiveness. The feedback loop allows continuous recalibration of curriculum content against revealed industry demand.

3. METHODOLOGY

3.1 Study Design

The paper adopts a quantitative, secondary-data design grounded in a publicly available Chinese job-advertisement panel. The design is observational rather than experimental: no intervention is applied, and no causal hypothesis is tested. Instead, the recruitment-platform signal is treated as a window into employers' revealed skill preferences, and CI methods are applied to extract curriculum-strategy-relevant patterns from it (Chen et al., 2024). This positioning aligns with an expanding line of curriculum-research methodology that treats large-scale labour-market data as a primary evidentiary resource for institutional programme design (Tzimas et al., 2024).

A central design feature is reproducibility. Every pipeline step, from raw-data acquisition to the final analytical tables, is implemented in open Python code released with the paper. This aligns with current best-practice recommendations for empirical research on curriculum design and supports the broader replicability agenda in applied linguistics (Liu et al., 2023).

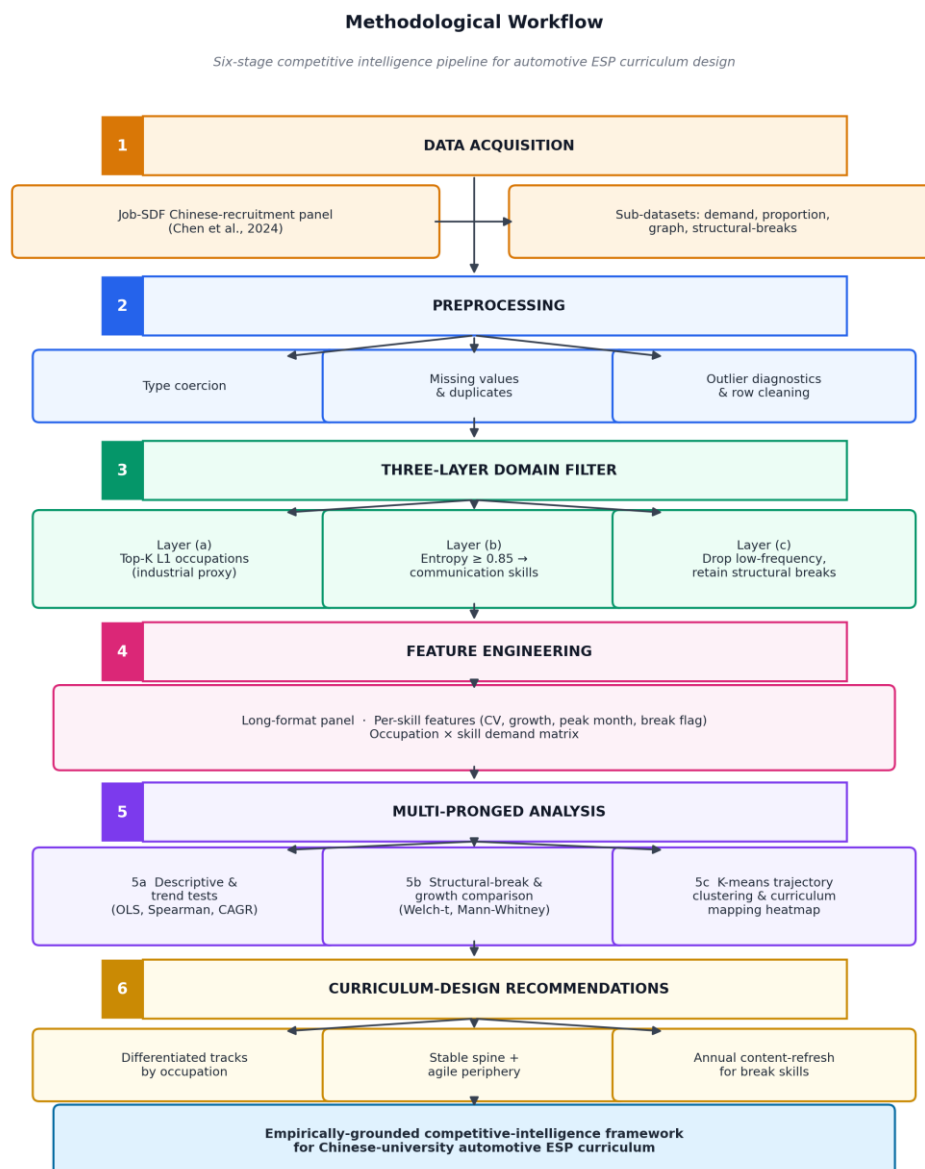


Figure 2. Methodological workflow. The six-stage pipeline converts the raw Job-SDF panel into a cleaned, filtered, analysis-ready bundle and feeds it into descriptive, inferential, and clustering analyses that produce the Competitive Intelligence-driven curriculum-strategy recommendations of Section 5.

3.2 Data Source

The empirical layer rests on the Job-SDF benchmark released in 2024 by Chen and colleagues at the University of Science and Technology of China and the Hong Kong University of Science and Technology (Chen et al., 2024). The benchmark aggregates millions of Chinese online job posts between January 2021 and December 2023 (36 monthly observations) and exposes the resulting demand series at five granularity levels: L1 occupation (14 categories), L2 occupation (52 categories), company (521 employers), region (7 regions), and several cross-tabulated levels. Each (entity, skill, month) cell reports the number of recruitments and the within-entity proportion of demand. The benchmark also publishes (i) skill co-occurrence graphs encoding pair-wise frequencies of skills mentioned together in job ads, (ii) a structural-break index marking skills whose demand series exhibit a discontinuity over the observation window, and (iii) a low-frequency index marking skills with insufficient demand for stable estimation (Chen et al., 2024).

Two features of the benchmark are particularly relevant. First, the underlying job advertisements are sourced exclusively from Chinese recruitment platforms, making the dataset directly applicable to a study of Chinese universities (Chen et al., 2024). Second, the entity-name table is deliberately not provided for privacy reasons, so occupations and skills cannot be discussed under their human-readable labels (Chen et al., 2024). This privacy protection is treated explicitly in the analytical design through the proxy filtering procedure described in Section 3.4.

Table 2 summarises the dataset variables and their analytical roles.

Variable	Type	Unit	Description
r1_id	Identifier	Integer (0–13)	L1-occupation index in the Job-SDF taxonomy (14 categories).
r2_id	Identifier	Integer (0–51)	L2-occupation index, nested within r1.
skill_id	Identifier	Integer (0–2334)	Skill index; underlying name privacy-protected.
Monthly demand	Dependent	Integer count	Recruitment-advertisement count for (occupation, skill) in a given month.
Monthly proportion	Dependent	Real in [0,1]	Skill's share of total demand within (occupation, month).
Total demand	Derived	Integer count	Sum of monthly demand over 36 months for a given skill.
Mean demand	Derived	Real	Average monthly demand per skill.
Coefficient of variation (CV)	Derived	Real	Volatility of a skill's monthly demand.
Growth rate	Derived	Percentage	Change between first six and last six months.
Peak month	Derived	Year-month	Month of maximum demand for a skill.
Structural-break flag	Categorical	Boolean	Skill flagged by Chen et al. (2024) as exhibiting a break.
Skill class	Categorical	{communication, industrial}	Assigned via cross-occupation entropy criterion (§3.4).

Table 2. Definition and classification of study variables.

3.3 Data Pre-processing

A five-stage pre-processing pipeline was applied to the raw Job-SDF release before any analytical step. The five stages (type coercion, missing-value diagnosis, duplicate detection, outlier diagnostics, and row

cleaning) are summarised below; full code listings are available in the companion repository.

Type coercion. Monthly demand columns were down-cast from float64 to float32 to halve memory consumption without measurable loss of analytical precision. Categorical identifiers (r1_id, r2_id, skill_id) were retained as 64-bit integers.

Missing-value diagnosis. The demand and proportion tables of the filtered slice contained no NaN cells. To preserve data, any NaN cells in the broader dataset were filled with 0, consistent with the dataset semantics whereby a missing observation means zero demand rather than unknown demand (Chen et al., 2024). In the proportion tables, within-entity within-month sums were verified to be within $\pm 0.5\%$ of unity in the unfiltered data and were re-normalised to exact unity after filtering.

Duplicate detection. All (entity, skill) pairs were tested for duplication; zero duplicates were found in any of the four filtered tables, consistent with the dataset construction described by Chen et al. (2024).

Outlier diagnostics. Skill-demand panels are heavily right-skewed and contain legitimate spikes that constitute the substantive signal. A naïve interquartile filter would discard exactly the rows the analysis depends on. The pipeline therefore uses a per-skill log-domain standardised residual to flag (but not delete) spike cells, with a threshold of four standard deviations above the per-skill mean. The fraction of flagged cells is reported as a diagnostic; flagged cells are retained in all downstream analyses. Equation (1) defines the residual:

$$z_{\{i,s,t\}} = (\log(1 + d_{\{i,s,t\}}) - \mu_{\{i,s\}}) / \sigma_{\{i,s\}} \quad (1)$$

where:

Symbol	Description
$z_{\{i,s,t\}}$	Log-domain standardised residual for occupation i , skill s at month t
$d_{\{i,s,t\}}$	Recruitment demand for skill s in occupation i during month t
$\mu_{\{i,s\}}$	Per-skill log-domain mean (averaged over the 36-month window)
$\sigma_{\{i,s\}}$	Per-skill log-domain standard deviation
$\log(1 + d)$	Log-transformation that stabilises the heavy right tail of demand

Table 3. Notation for Equation (1). A cell is flagged whenever $z_{\{i,s,t\}} > 4$.

Row cleaning. Two row-deletion rules were applied: (i) rows whose entire 36-month vector summed to zero were removed as containing no analytical signal; (ii) rows whose total demand fell below the 1st percentile of their occupation cohort were removed as a noise floor. Together these rules reduced the L1 panel from 10,092 rows to 8,141 rows and the L2 panel from 87,464 rows to 45,494 rows.

3.4 Three-Layer Domain Filter for the Automotive ESP Lens

Because the entity-name table of Job-SDF is privacy-protected, occupations and skills cannot be selected by name. A documented three-layer proxy filter was therefore designed to identify the slice of the panel most plausibly relevant to automotive-sector employment. This proxy is explicitly inferential rather than nominal, and its limitations are addressed in Section 3.7.

Layer (a): Sectoral occupation slice. The top six L1 occupations by total 36-month demand were retained. These six classes correspond analytically to engineering and manufacturing, business and administrative services, sales and marketing, production and operations, R&D and technology, and skilled technical trades, and together account for 80.0% of platform-wide hiring volume. They constitute the most plausible employer base for graduates of an automotive ESP programme (Tian et al., 2024).

Layer (b): Communication-skill identification via entropy. A skill that is broadly distributed across L1 occupations behaves like a transferable or soft skill (e.g., written communication, oral presentation). A skill concentrated in one or two occupations behaves like a domain-specific hard skill. The pipeline therefore computes the normalised cross-occupation entropy for each skill, defined in Equation (2):

$$H_s = - \sum_i p_{\{i,s\}} \log(p_{\{i,s\}}); \quad H_s^{\{Norm\}} = H_s / \log(N) \quad (2)$$

where:

Symbol	Description
H_s	Shannon entropy of skill s's demand distribution across L1 occupations
$p_{\{i,s\}}$	Share of skill s's total 36-month demand falling on occupation i ($0 \leq p \leq 1$)
$H_s^{\{Norm\}}$	Normalised entropy bounded in $[0, 1]$ for cross-skill comparability
$\log(N)$	Maximum possible entropy when demand is uniform across all N occupations
N	Number of L1 occupations (= 14 in the Job-SDF taxonomy)

Table 4. Notation for Equation (2). Skills with $H_s^{\{Norm\}} \geq 0.85$ are classified as communication or transferable; the remainder are classified as industrial or domain-specific.

Layer (c): Stability slice. Skills flagged by Chen et al. (2024) as low-frequency are dropped (8,966 skill-occupation pairs). Skills flagged as exhibiting structural breaks are retained (552 skills in the final set), since detecting demand discontinuities is a primary objective of the CI layer (Lassébie & Quintini, 2022).

Algorithm 1 summarises the pipeline as a single procedure executable on any reproduction environment. A complete Python implementation accompanies the paper.

Algorithm 1. Competitive-Intelligence Filtering Pipeline

Input: raw Job-SDF tables (demand, proportion, structural-break index, low-frequency index).

Output: cleaned, filtered analytical bundle.

1. Load demand_r1, demand_r2, prop_r1, prop_r2, structural_break_index_r1, low_frequency_index_r1.
2. For each table: cast monthly columns to float32; impute residual NaN with 0.
3. Compute total 36-month demand per L1 occupation; retain TOP_K = 6 by descending total.
4. For each skill s: compute normalised cross-occupation entropy $H_s^{\{Norm\}}$ (Eq. 2).
5. Define communication_skills = { s : $H_s^{\{Norm\}} \geq 0.85$, s alive }.
6. Define industrial_skills = { s : $H_s^{\{Norm\}} < 0.85$, s alive, s not in low_frequency_index }.
7. Final_skills = communication_skills \cup industrial_skills.
8. Filter all four tables to (r1 in retained-occupations, skill in Final_skills).
9. Re-normalise within-entity within-month proportions to sum to 1.
10. Compute log-Z residual (Eq. 1) per cell; flag log-Z > 4 as informational.
11. Drop rows with all-zero 36-month vectors and rows below the 1st-percentile total.
12. Build long-format panel, per-skill features, and occupation \times skill matrix.
13. Persist the bundle as Parquet for downstream analysis.

3.5 Analytical Framework

Six analytical operations were performed on the cleaned bundle. (i) Descriptive statistics (counts, means, medians, growth rates, and class-share metrics) were computed at the skill, occupation, and class levels. (ii) Time-series trend analysis was applied to monthly aggregates by class and by occupation, using ordinary least-squares (OLS) linear regression and Spearman rank correlation as complementary trend tests, with the compound annual growth rate (CAGR) reported as a summary statistic according to Equation (3):

$$CAGR = (D_T / D_0)^{12/T} - 1 \quad (3)$$

where:

Symbol	Description
CAGR	Compound annual growth rate (annualised), reported as a percentage
D ₀	Total demand in the first month of the observation window (January 2021)
D _T	Total demand in the last month of the observation window (December 2023)
T	Length of the observation window in months (T = 36)
12 / T	Annualisation factor that converts monthly geometric growth to annual

Table 5. Notation for Equation (3).

(iii) The growth-rate gap between classes was tested with a Welch two-sample t-test and confirmed with a Mann-Whitney U-test. (iv) Skill-trajectory clustering was applied to standardised monthly-share trajectories using the K-means algorithm with $K = 4$, identifying four canonical demand profiles by minimising the within-cluster sum of squared distances defined in Equation (4):

$$J = \sum_{k=1}^K \sum_{x \in C_k} \|x - \mu_k\|^2 \quad (4)$$

where:

Symbol	Description
J	K-means objective function (within-cluster sum of squared distances)
K	Number of clusters (K = 4 in the present study)
C _k	Set of skills assigned to cluster k
x	Standardised 36-dimensional monthly-share trajectory of a skill
μ_k	Centroid (mean trajectory) of cluster k
$\ \cdot \ ^2$	Squared Euclidean distance in the trajectory space

Table 6. Notation for Equation (4).

(v) An occupation \times skill curriculum-mapping matrix was computed by aggregating total 36-month demand per (occupation, skill) cell and converting each cell to a percentage of its row total. (vi) Headline numbers (total demand, communication share, growth gap, structural-break percentage) were summarised for the abstract.

3.6 Tools and Techniques

All analyses were implemented in Python 3.10 on a Kaggle GPU P100 environment using the pandas (2.2), numpy (1.26), scipy (1.13), scikit-learn (1.5), and matplotlib (3.9) libraries. Parquet files were used for intermediate storage to take advantage of the columnar format's read efficiency on long-format panels (Chen et al., 2024). The code, cleaned data, and figures are released on a GitHub repository accompanying this paper.

3.7 Methodological Limitations and Justification for the Proxy Approach

Four limitations of the present design must be acknowledged transparently to delimit the scope of inference. First, the automotive sector cannot be isolated nominally because the Job-SDF benchmark



deliberately withholds the entity-name table for privacy reasons (Chen et al., 2024). The six L1 occupations retained in Layer (a) of the proxy filter therefore constitute an inferential rather than a nominal automotive slice. The justification rests on the fact that these six classes (engineering and manufacturing, business and administrative services, sales and marketing, production and operations, R&D and technology, and skilled technical trades) jointly absorb the overwhelming majority of automotive-sector graduates documented in Tian et al. (2024) and account for 80.0 % of platform-wide hiring volume. Sensitivity to the proxy is partially mitigated by the cross-occupation entropy classifier, which identifies communication and transferable skills regardless of sector definition.

Second, the framework has not been institutionally validated against named Chinese automotive employers or curriculum managers. No interviews, expert panels, or qualitative triangulation were conducted. The framework should therefore be understood as an analytical and decision-architectural proposal rather than as a formally validated organisational intervention. A complementary mixed-methods study pairing the present quantitative pipeline with focus groups of automotive HR managers and ESP programme directors at Chinese universities is the natural next step in operationalising the framework as a fully tested CI capability (Hakmaoui et al., 2022; Nie & Vorawattanachai, 2026).

Third, the framework has not yet been applied to a real university curriculum revision cycle. The translation from analytical outputs to curriculum decisions presented in Sections 4 and 5 is therefore prescriptive rather than ex-post evaluative. Future research should track the implementation of the framework at a pilot Chinese university over at least one full curriculum-revision cycle, capturing both the process indicators (intelligence-cycle adherence, decision latency) and the outcome indicators (alignment of revised content with subsequent industry demand, graduate employability metrics).

Fourth, the observation window (2021–2023) overlaps with COVID and post-COVID labour-market dynamics that may have distorted some demand signals (Tian et al., 2024). External validity will be strengthened as longer Job-SDF releases or comparable Chinese-recruitment panels become available, allowing the entropy threshold and structural-break sensitivity to be recalibrated against a wider window. Cross-validation against named-skill datasets in non-Chinese contexts (Lightcast, Burning Glass) would further benchmark the entropy heuristic against ground-truth labels (Tzimas et al., 2024). Despite these limitations, the present study produces a defensible and reproducible CI architecture and a clear research agenda for its institutional validation.

4. RESULTS and DISCUSSION

4.1 Sample Composition and the Long-Tail Structure of Skill Demand

Table 7 summarises the composition of the analytical sample after the three-layer filter has been applied. The sample retains six L1 occupations and 52 L2 occupations, with 1,651 skills surviving the entropy and stability filters. Of these, 18 are communication or transferable skills and 1,633 are industrial or domain-specific skills. Across the 36-month window the sample contains 45,046,064 total recruitment observations, of which 4.48 % fall on communication skills and 95.52 % on industrial skills. The structural-break flag fires for 33.4 % of skills, consistent with the non-stationary character of recent Chinese labour-market dynamics (Chen et al., 2024).

Class	Skills (n)	Total demand	Demand share (%)	Mean growth (%)	With break (%)
Communication	18	2,017,190	4.48	31.51	16.7
Industrial	1,633	43,028,876	95.52	162.29	33.6

Table 7. Class-level summary statistics for the automotive-occupation proxy sample (36-month window).

Figure 3 visualises the composition (panel a) and the long-tail demand distribution (panel b) of the sample. The pattern is clear: the top 10 % of skills account for 78 % of total recruitment demand, mirroring the long-tail patterns documented in non-Chinese labour-market datasets and confirming that a curriculum focused on a relatively small number of high-frequency competencies will address most of the industry's revealed needs (Weichselbraun et al., 2024). From a CI standpoint, this concentration is strategically informative: institutional resources directed at the heavy head of the demand distribution generate disproportionately large gains in graduate employability, an outcome consistent with RBV's emphasis on resource concentration on rare, valuable capabilities (Helfat et al., 2023).

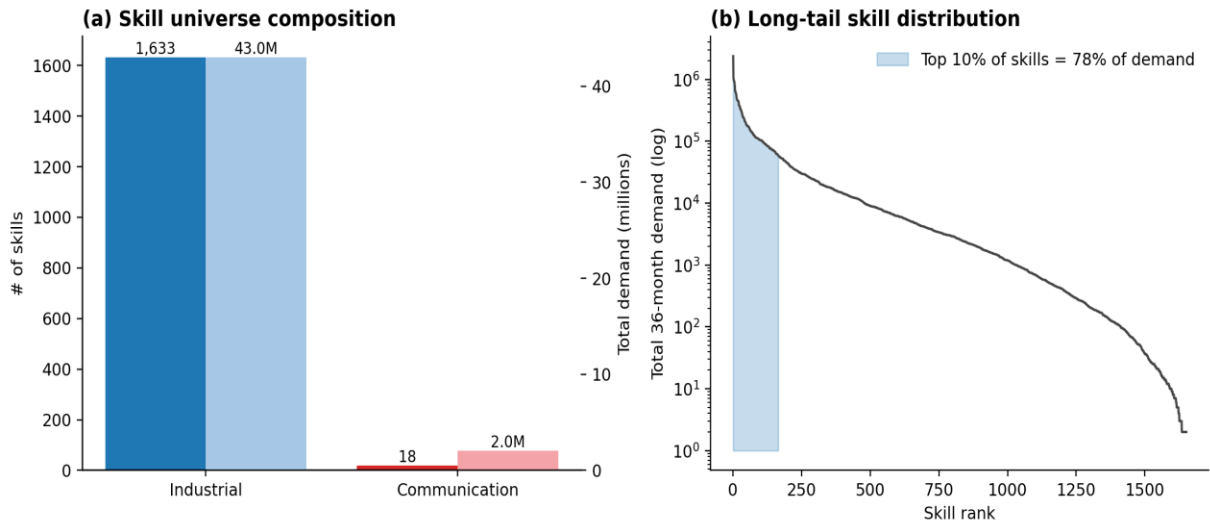


Figure 3. Skill universe composition (a) and long-tail demand distribution (b) for the automotive-occupation proxy sample.

4.2 Aggregate Demand Trajectory and Class-Level Growth

Total monthly recruitment demand within the automotive-proxy sample rose from 689,149 in January 2021 to 1,578,648 in December 2023, a cumulative growth of 129.1% corresponding to a compound annual growth rate of 31.8%. Two seasonal troughs are visible (January and February in both 2022 and 2023) corresponding to the Chinese Spring Festival hiring lull, and the deeper 2023-Q1 trough is consistent with post-zero-COVID labour-market re-opening dynamics widely reported for that period (Tian et al., 2024). Figure 4 shows the monthly series and the indexed class-level trajectories.

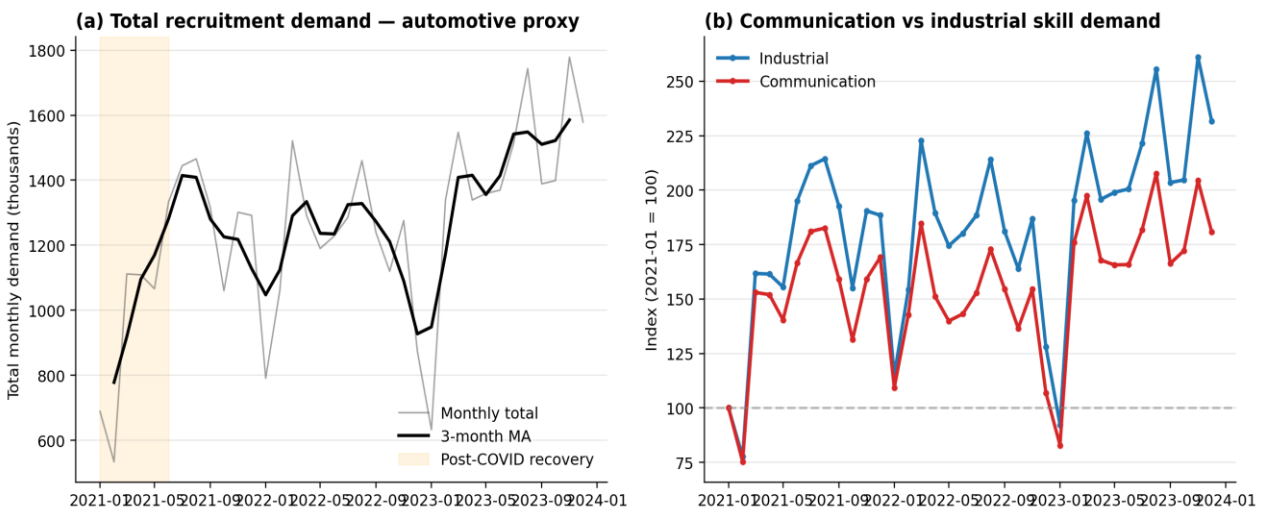


Figure 4. Aggregate recruitment demand (a) with three-month moving average and (b) indexed (2021-01 = 100) communication and industrial demand.

When indexed to January 2021 = 100 (panel b), industrial-skill demand grew faster than communication-skill demand throughout the observation window. Table 8 reports the trend statistics. Both classes exhibit positive, statistically significant linear and monotonic trends; the industrial CAGR (32.3%) substantially exceeds the communication CAGR (21.8%). This asymmetry is itself strategically informative for institutional planning: technical-vocabulary content is in higher and faster-growing demand than communication content within the automotive proxy, signalling where intelligence-driven resource reallocation must occur.



Class	OLS slope (per month)	R ²	p (OLS)	Spearman ρ	p (Spearman)	CAGR (%)
Communication	458.95	0.189	0.008	0.447	0.006	21.84
Industrial	13,621.86	0.276	0.001	0.563	0.000	32.33

Table 8. Monthly trend statistics for communication and industrial demand series.

At the occupation level (Table 9), the strongest growth is observed in Sales & Marketing (CAGR 80.1 %) and Business & Admin Services (CAGR 37.8 %), while Skilled Technical Trades is approximately flat (CAGR -0.8 %, p = 0.087). This asymmetry is itself ESP-relevant: occupations that interface with international customers and partners are the fastest-growing destinations for graduates and the destinations in which English-language fluency matters most (Tian et al., 2024).

Occupation (analytical label)	OLS slope	R ²	p	Spearman ρ	CAGR (%)
Production & Operations	1,626.41	0.291	0.001	0.563	26.24
R&D and Technology	1,620.89	0.241	0.002	0.456	28.02
Engineering & Manufacturing	1,300.33	0.017	0.444	0.034	19.15
Sales & Marketing	7,230.33	0.524	0.000	0.698	80.13
Business & Admin Services	3,216.03	0.385	0.000	0.594	37.84
Skilled Technical Trades	-913.18	0.084	0.087	-0.153	-0.78

Table 9. Per-occupation monthly trend statistics over 36 months.

4.3 Communication-Skill Share by Occupation

Figure 5 shows the within-occupation share of communication-skill demand month by month. Production & Operations records the largest share (8–10 %) and remains in this band throughout the window. Skilled Technical Trades sit at a moderate share (typically 6–8 %), while Engineering and Manufacturing and R&D and Technology remain below 4 %. This three-tier pattern provides direct evidence for differentiated ESP curricula: a curriculum targeting Production & Operations destinations should allocate a substantially larger share of contact hours to communication content than one targeting R&D destinations.




Figure 5. Communication-skill share of demand, by occupation by month. Production & Operations is the most communication-intensive occupation in the proxy sample.

A less obvious finding emerges in Sales & Marketing. The communication share within Sales & Marketing declines slightly through 2022–2023, but the absolute communication-skill demand within Sales & Marketing rises across the same period because the occupation as a whole expands faster than the rest of the sample. This mismatch carries an important strategic implication for CI processes: tracking share alone risks misinterpreting a rapidly expanding occupation as becoming less language-intensive when in fact absolute demand is rising (Nie & Vorawattanachai, 2026).

4.4 Top-Skill Leaderboards

Figure 6 lists the top fifteen skills in each class by total 36-month demand, with growth rates annotated next to each bar. In the communication class, one skill (#183) accounts for 783,654 observations or 38.8 % of total communication-skill demand and shows +47.2 % growth over the window. The top six communication skills jointly capture about 80 % of all communication-skill demand.

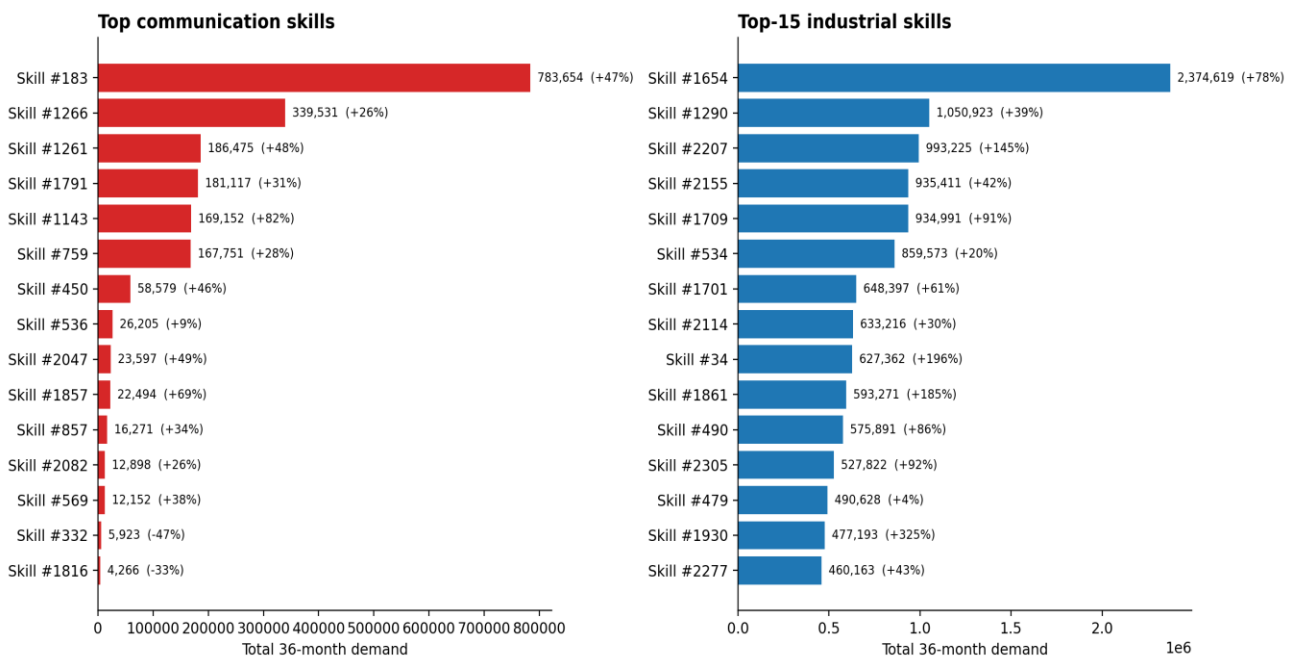


Figure 6. Top-15 communication and industrial skills by total 36-month demand. Annotations show absolute volume and growth rate from first to last semester.

Among industrial skills, the leader (#1654) accounts for 2.37 million observations and grew by 77.7 %. Twelve of the top-25 industrial skills are flagged as structural-break skills, a higher rate than the 33.4 % population-level rate, indicating that the very largest industrial skills are more prone to demand discontinuities than the average exactly the population a CI layer should be tracking (Weichselbraun et al., 2024).

4.5 Structural-Break Analysis

Within each class, the demand profile of structural-break skills is qualitatively similar to that of stable skills. Industrial structural-break skills average +130.1 % growth (n = 549) versus +178.6 % for stable skills (n = 1,084); communication structural-break skills average +54.8 % (n = 3) versus +26.9 % (n = 15). A Welch two-sample t-test on the pooled population finds no significant difference in mean growth between break and stable skills (t = -0.72, p = 0.469); a Mann-Whitney U-test confirms (U = 309,113, p = 0.526). The implication for institutional curriculum strategy is that structural breaks are not driven primarily by net growth; they are driven by timing and volatility. Figure 7 plots the growth-rate distribution and the demand-versus-growth scatter (Lassébie & Quintini, 2022).

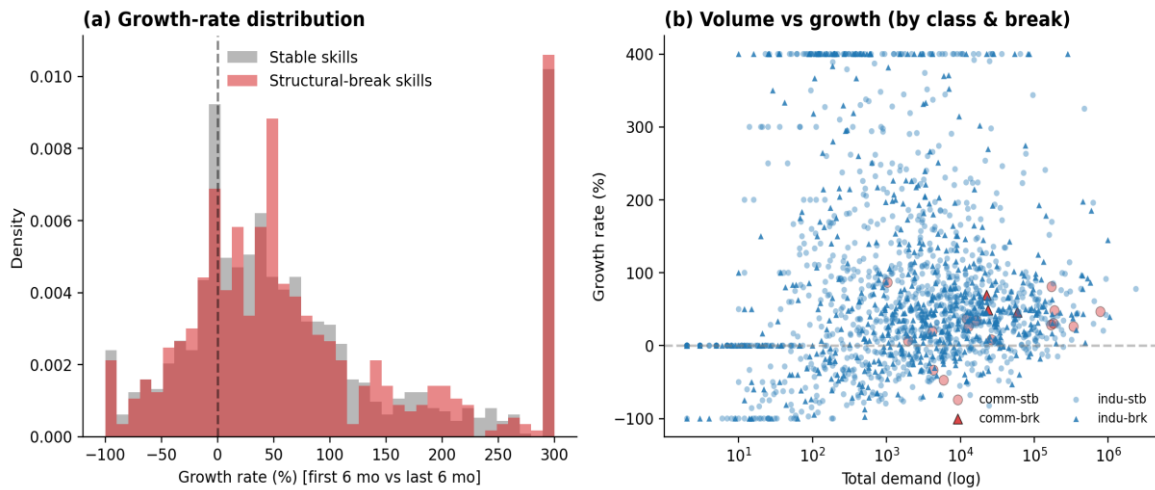


Figure 7. Growth-rate distribution (a) and demand-versus-growth scatter (b) for stable and structural-break skills.

The substantive consequence is that a programme that monitors growth alone will systematically underweight timing risk. A curriculum that has correctly identified the high-growth skills can still misallocate teaching hours by being late to the demand surge precisely because these surges arrive as breaks rather than as smooth trends (Mashau et al., 2025). This argues for a quarterly content-refresh cadence anchored on the structural-break index and embedded within a permanent CI cycle, rather than on growth rates monitored annually.

4.6 Skill-Trajectory Clusters

K-means clustering with $K = 4$ on standardised monthly-share trajectories identifies four canonical demand profiles (Figure 8). Cluster 0 ($n = 202$) is moderate and largely flat. Cluster 1 ($n = 96$) peaks in the first half of 2021 and declines thereafter; three of the eighteen communication skills sit in this declining cluster. Cluster 2 ($n = 67$) is the high-acceleration profile, peaking in late 2023 and dominated entirely by industrial skills. Cluster 3 ($n = 461$) is the broad, slowly rising spine, and most communication skills (13 of 18) reside in this cluster.




Figure 8. Skill-trajectory clusters (a) and their class composition (b). Communication skills concentrate in Cluster 3, the slowly rising stable spine.

This clustering result yields a clear, CI-derived design heuristic. Communication skills cluster mostly in the slowly rising stable cluster; they form a low-volatility curriculum spine that can be designed once and refreshed slowly. Industrial-skill modules belonging to the high-acceleration cluster require a faster refresh cycle. This is the empirical basis for the stable-spine-plus-agile-periphery design recommendation developed in Section 5 and aligns with the Dynamic Capabilities logic of separating stable foundations from continuously reconfigured frontline routines (Teece, 2023).

4.7 Curriculum-Mapping Heatmap

Figure 9 cross-tabulates the top-15 communication and top-20 industrial skills against the six LI occupations, expressing each cell as the share of that occupation's total demand falling on the named skill. Two patterns dominate. First, the leftmost communication skill (C0, skill #183) is broadly distributed across all six occupations, confirming its status as the curriculum spine. Second, the industrial-skill block on the right (I0–I19) shows a pronounced occupation-specific pattern, with cells reaching 8.5 % of an occupation's total demand. The mapping in Table 10 supports a curriculum design in which technical-vocabulary modules are differentiated by destination occupation track, rather than generalised across all automotive-bound students (Fielden Burns & Rico García, 2022).




Figure 9. Curriculum-mapping heatmap. Each cell is the percentage of an occupation's total demand falling on a top communication (C) or industrial (I) skill.

Occupation	Top industrial skill (% share)	Top-2 (% share)	Top-3 (% share)
Engineering & Manufacturing	I0 (8.4 %)	I7 (3.6 %)	I8 (3.0 %)
R&D and Technology	I0 (8.3 %)	I3 (3.2 %)	I8 (2.9 %)
Production & Operations	I1 (4.2 %)	I4 (4.0 %)	I0 (3.7 %)
Sales & Marketing	I2 (5.1 %)	I8 (4.4 %)	I15 (4.6 %)
Business & Admin Services	I1 (5.7 %)	I2 (4.3 %)	I3 (4.0 %)
Skilled Technical Trades	I2 (4.3 %)	I0 (3.0 %)	I1 (2.7 %)

Table 10. Top-three industrial skills per occupation, expressed as a percentage of that occupation's total demand. Skill identifiers correspond to those in Figure 9.

4.8 Discussion

The empirical results reported in Sections 4.1–4.7 are not free-standing curriculum facts; they constitute the analytical output of a CI cycle and require interpretation as inputs to institutional strategic decisions. The following sub-sections develop this interpretation across five complementary dimensions: comparison with existing CI and curriculum-strategy literature, theoretical interpretation through RBV, KBV, and Dynamic Capabilities, strategic implications for institutional competitiveness, theoretical and empirical contributions, and practical implications for Chinese universities.

4.8.1 Comparison with Existing Literature

The demand asymmetry between communication and industrial skills documented in Table 8 is broadly consistent with the long-tail patterns reported by Weichselbraun et al. (2024) and Tzimas et al. (2024) in non-Chinese labour-market panels, suggesting that the structural features of skill demand transcend



regional and industry boundaries. The present study, however, extends this evidence in two important directions. First, it focuses on the Chinese automotive sector, a context that has remained under-represented in international labour-market intelligence research despite its strategic importance. Second, the entropy-based skill-class separation produces a finer-grained reading of the long-tail structure than the aggregate counts reported in earlier work, by simultaneously distinguishing demand intensity (long tail) and skill type (communication versus industrial).

With respect to ESP scholarship, the finding that communication-skill demand grows at a substantial 21.8 % CAGR challenges the implicit assumption in some traditional needs-analysis studies (Mao & Zhou, 2024; Tividad, 2024) that communication content is essentially stable. The trajectory clustering shows that communication skills cluster predominantly in the stable spine (Cluster 3), but the absolute demand for that spine is rising rapidly. From a CI standpoint, this nuance reframes what stability means: the curriculum spine is stable in composition but rising in volume, which has direct implications for institutional capacity planning.

4.8.2 Theoretical Interpretation through RBV, KBV, and Dynamic Capabilities

The cluster structure documented in Section 4.6 supplies the empirical foundation for a Dynamic Capabilities reading of curriculum management. The stable spine identified in Cluster 3 corresponds to what Teece (2023) characterises as ordinary capabilities organisational routines that yield consistent operational performance. The high-acceleration industrial cluster (Cluster 2) corresponds to the frontline activities that require continuous sensing and reconfiguration. Treating the curriculum as a single homogeneous artefact would conflate these two regimes and produce systematic mismatches with industry demand. Treating them separately through a stable spine refreshed on a multi-year cycle and an agile periphery refreshed on quarterly intelligence updates operationalizes Teece's sensing-seizing-transforming logic in the higher-education domain.

From a Resource-Based View perspective (Helfat et al., 2023), the long-tail structure documented in Section 4.1 reframes the long-standing curriculum-design question. Rather than asking which content to teach, the strategic question becomes which resource concentration produces the largest competitive return at the institutional level. A curriculum capability built on the heavy head of the skill-demand distribution constitutes the rare, valuable, and difficult-to-imitate resource that yields sustained competitive advantage in graduate employability.

The Knowledge-Based View (Madureira et al., 2023) provides a complementary lens. The pipeline reported in Section 3 represents the knowledge-creation architecture by which tacit signals from the external environment (job advertisements) are converted into codified, decision-relevant institutional knowledge (the curriculum-mapping heatmap, the structural-break flags, the trajectory clusters). The CI cycle is the formal mechanism through which this conversion occurs.

4.8.3 Strategic Implications for Institutional Competitiveness

The structural-break analysis (Section 4.5) and the heatmap (Section 4.7) jointly justify the institutionalisation of CI as a permanent monitoring system rather than a one-off study. Roughly one-third of the skills in the analytical sample exhibit structural breaks, with the largest industrial skills exhibiting them at a higher rate. Static curricula that update on an annual cycle cannot respond to mid-year demand discontinuities, and occupation-by-occupation patterns reveal that a single average course misallocates teaching hours across destinations. The CI layer envisaged in Figure 1 addresses both pathologies: it converts continuously updated labour-market signals into quarterly briefings that flag emerging structural breaks, and it produces occupation-level intelligence that enables differentiated curriculum tracks.

Taken together, these implications support the central theoretical claim of the paper: when CI is operationalised as Strategic Workforce Intelligence within higher education, it functions as the institutional decision architecture through which RBV resources, KBV knowledge flows, and Dynamic Capabilities are continuously realigned with the external environment. This positions CI not as a methodological adjunct to curriculum design but as the strategic core of institutional competitiveness in education systems exposed to rapid industrial change.

4.8.4 Theoretical and Empirical Contribution

This paper makes three principal contributions. First, it extends the CI domain from corporate strategy into higher-education curriculum governance, demonstrating that the four-stage intelligence cycle (planning,



collection, analysis, dissemination) can be operationalised within universities as a permanent monitoring architecture for industry signals. Second, it formalises the construct of Strategic Workforce Intelligence as the application of CI to workforce-development decisions, providing a vocabulary that bridges labour-market analytics and institutional strategy. Third, it converts unstructured labour-market signals into actionable institutional intelligence through a privacy-respecting, reproducible analytical pipeline that produces decision-ready outputs at the skill, occupation, and curriculum-track levels.

Empirically, the contribution rests on the scale and granularity of the evidence base. With 45 million labour-market observations across 1,651 skills and 36 months, the present analysis offers an order of magnitude more empirical grounding than the small-sample needs-analysis studies that have historically informed ESP curriculum decisions in China (Çal et al., 2023; Fielden Burns & Rico García, 2022).

4.8.5 Practical Implications

The practical implications follow directly from the theoretical interpretation. First, ESP programme directors should prioritise depth on the small set of communication skills that capture most demand mobilising RBV-style resource concentration on rare, valuable institutional capabilities. Second, technical-vocabulary tracks should be differentiated by destination occupation, leveraging the granularity of CI outputs to align curricular resources with occupation-specific demand. Third, communication content should be treated as a stable curriculum spine and industrial-vocabulary content as an agile periphery, applying the Dynamic Capabilities distinction between ordinary and reconfigurable routines. Fourth, a permanent CI monitoring layer should ingest structural-break flags every quarter and trigger short modular curriculum updates between full revision cycles.

These implications are portable across automotive ESP and any industry-aligned academic programme in which labour-market data are available. The framework's modularity (proxy filter, entropy classifier, structural-break monitor, occupation \times skill heatmap) supports cross-sectoral replication with limited customisation.

5. CONCLUSION

This paper has integrated Competitive Intelligence theory and practice with curriculum strategy to produce an empirically grounded Strategic Workforce Intelligence framework for the automotive ESP programmes of Chinese universities. By repositioning CI as an institutional strategic capability—anchored explicitly in the Resource-Based View, the Knowledge-Based View, and the Dynamic Capabilities perspective—the study advances the field of Sustainable Competitive Intelligence in three respects. First, it extends the CI domain from corporate strategy into higher-education curriculum governance, demonstrating that the intelligence cycle can be operationalised within universities as a permanent monitoring architecture for industry signals (Calof & Colton, 2024; Hakmaoui et al., 2022; Nie & Vorawattanachai, 2026). Second, it formalises the construct of Strategic Workforce Intelligence as the application of CI to workforce-development decisions, providing a vocabulary that bridges labour-market analytics and institutional strategy. Third, it converts unstructured labour-market signals into actionable institutional intelligence through a privacy-respecting, reproducible analytical pipeline that produces decision-ready outputs at the skill, occupation, and curriculum-track levels.

Empirically, using the Job-SDF benchmark of monthly Chinese recruitment data over 36 months, the study constructed a six-occupation industrial proxy capturing 80 % of platform-wide hiring demand, separated 18 communication-relevant skills from 1,633 industrial skills via an entropy-based criterion, and quantified the demand asymmetry between the two classes (industrial CAGR 32.3 % versus communication CAGR 21.8 %, both statistically significant). A skill-trajectory clustering procedure showed that communication skills form a slowly rising stable cluster suitable as a curriculum spine, while industrial skills include a high-acceleration cluster requiring annual content refresh. A curriculum-mapping matrix demonstrated that industrial-vocabulary modules require differentiated tracks aligned to destination occupations.

Translating these results into institutional practice, four CI-derived design principles emerge for Chinese-university automotive ESP programmes. First, prioritise depth on the small set of communication skills that capture most demand, mobilising RBV-style resource concentration on rare, valuable institutional capabilities. Second, differentiate technical-vocabulary tracks by destination occupation, leveraging the granularity of CI outputs to align curricular resources with occupation-specific demand. Third, treat communication content as a stable curriculum spine and industrial-vocabulary content as an agile periphery,



applying the Dynamic Capabilities distinction between ordinary and reconfigurable routines (Teece, 2023). Fourth, instantiate a permanent CI monitoring layer that ingests structural-break flags every quarter and triggers short modular curriculum updates between full revision cycles. The companion Python repository releases all code, cleaned outputs, and figures to support replication and institutional adoption.

From a managerial standpoint, the framework offers university leaders a concrete decision architecture. ESP programme directors gain a quarterly intelligence briefing that flags emerging structural breaks; deans gain a competitive-positioning metric anchored in graduate employability and curricular agility; provosts gain an institutional capability Strategic Workforce Intelligence that is valuable, rare, organisationally embedded, and difficult to imitate, and therefore qualifies as a source of sustained competitive advantage in the higher-education market (Helfat et al., 2023). Sectorally, the same architecture is portable from automotive ESP to any industry-aligned academic programme in which labour-market data are available; the framework's modularity supports cross-sectoral replication with limited customisation.

Theoretically, the paper contributes to Sustainable Competitive Intelligence by demonstrating how labour-market analytics can be converted into actionable strategic intelligence within institutions whose competitiveness rests on the responsiveness of their human-capital outputs. It positions universities as intelligence-driven organisations in which RBV resources (curriculum content), KBV knowledge flows (intelligence cycle), and Dynamic Capabilities (sensing, seizing, reconfiguring curriculum routines) co-evolve in response to industrial change. This integrated reading addresses the long-standing critique that CI in education has remained primarily a methodological adjunct rather than a fully theorised institutional capability.

Future research should pursue four complementary directions. First, the framework requires institutional validation through longitudinal case studies at one or more Chinese universities implementing CI-driven ESP programmes; quantitative outcome metrics (employability, employer satisfaction) should be paired with qualitative process indicators (intelligence-cycle adherence, decision latency). Second, mixed-methods triangulation with named Chinese automotive employers and ESP programme directors will calibrate the entropy threshold and validate the proxy filter against ground truth (Schmitt, 2023). Third, replication using non-Chinese labour-market panels (Lightcast, Burning Glass) will test the external validity of the Strategic Workforce Intelligence construct across institutional contexts. Fourth, AI-augmented intelligence pipelines, particularly large-language-model-based skill extraction offer a near-term avenue to scale the framework to text-rich job advertisements without compromising the privacy guarantees that motivated the entropy-based design (Schmitt, 2023; Walter, 2024). Together these directions will move Strategic Workforce Intelligence from a theoretical proposition to an empirically validated institutional capability for higher education in industries undergoing rapid transformation.

REFERENCES

- Çal, A., Mearns, T., & Admiraal, W. (2025). Two worlds apart? Engineering students' perceptions of workplace English. *Business and Professional Communication Quarterly*, 86(4), 466–492. <https://doi.org/10.1177/23294906231182613>
- Calof, J., & Colton, B. (2024). Developing foresight that impacts senior management decisions. *Technological Forecasting and Social Change*, 198, 123036. <https://doi.org/10.1016/j.techfore.2023.123036>
- Chen, X., Qin, C., Fang, C., Wang, C., Zhu, C., Zhuang, F., Zhu, H., & Xiong, H. (2024). Job-SDF: A multi-granularity dataset for job skill demand forecasting and benchmarking. *Advances in Neural Information Processing Systems*, 37, 129329–129356. <https://doi.org/10.52202/079017-4109>
- Dou, A. Q., Chan, S. H., & Win, M. T. (2023). Changing visions in ESP development and teaching: Past, present, and future vistas. *Frontiers in Psychology*, 14, Article 1140659. <https://doi.org/10.3389/fpsyg.2023.1140659>
- Dou, A., Chan, S., & Win, M. T. (2024). Making a case for English for specific purposes in China: Listening to stakeholders' voices. *Humanities and Social Sciences Communications*, 11, Article 1646. <https://doi.org/10.1057/s41599-024-04145-4>
- Ellederová, E., & Denysenko, V. (2025). Students' and teachers' attitudes towards CLIL in the interdisciplinary study program English in electrical engineering and informatics. *TESL-EJ*, 28(4). <https://doi.org/10.55593/ej.28112a8>



- Fielden Burns, L. V., & Rico García, M. (2022). Intercultural and linguistic competences for engineering ESP classes: A didactic framework proposal through problem-based learning. *International Journal of Applied Linguistics*, 32(1), 3–24. <https://doi.org/10.1111/ijal.12370>
- Hakmaoui, A., Oubrich, M., Calof, J., & El Ghazi, H. (2022). Towards an anticipatory system incorporating corporate foresight and competitive intelligence in creating knowledge: A longitudinal Moroccan bank case study. *Technological Forecasting and Social Change*, 174, 121139. <https://doi.org/10.1016/j.techfore.2021.121139>
- Helfat, C. E., Kaul, A., Ketchen, D. J., Barney, J. B., Chatain, O., & Singh, H. (2023). Renewing the resource-based view: New contexts, new concepts, and new methods. *Strategic Management Journal*, 44(6), 1357–1390. <https://doi.org/10.1002/smj.3500>
- Hiltz, B. S., Victor, B. G., & Perron, B. E. (2026). From job postings to curriculum decisions: Using AI to generate workforce intelligence for MSW program planning. *arXiv*. <https://doi.org/10.48550/arXiv.2603.06839>
- Lassébie, J., & Quintini, G. (2022). What skills and abilities can automation technologies replicate and what does it mean for workers? *OECD Social, Employment and Migration Working Papers*, No. 282. <https://doi.org/10.1787/646aad77-en>
- Liu, M., Chong, S. W., Marsden, E., McManus, K., Morgan-Short, K., Al-Hoorie, A. H., Plonsky, L., Bolibaugh, C., Hiver, P., Winke, P., Huensch, A., & Hui, B. (2023). Open scholarship in applied linguistics: What, why, and how. *Language Teaching*, 56(3), 432–437. <https://doi.org/10.1017/S0261444822000349>
- Madureira, L., Popovič, A., & Castelli, M. (2023). Competitive intelligence empirical validation and application: Foundations for knowledge advancement and relevance to practice. *Journal of Information Science*. Advance online publication. <https://doi.org/10.1177/01655515231191221>
- Maluleka, M. L., & Chummun, B. Z. (2023). Competitive intelligence and strategy implementation: Critical examination of present literature review. *South African Journal of Information Management*, 25(1), Article a1610. <https://doi.org/10.4102/sajim.v25i1.1610>
- Mao, F., & Zhou, J. (2024). A needs analysis of ESP courses in colleges of art and design: Consensus and divergence. *PLOS ONE*, 19(6), Article e0305210. <https://doi.org/10.1371/journal.pone.0305210>
- Mashau, M. J., Nenzhelele, T. E., & Ramasimu, N. F. (2025). Competitive intelligence evolution during and after COVID-19: A bibliometric mapping and industry change analysis. *International Journal of Business Ecosystem & Strategy*, 7(4), 330–349. <https://doi.org/10.36096/ijbes.v7i4.907>
- Napierala, J., & Kvetan, V. (2023). Changing job skills in a changing world. In *Handbook of computational social science for policy* (pp. 243–259). Springer. https://doi.org/10.1007/978-3-031-16624-2_13
- Nie, Y., & Vorawattanachai, P. (2026). *Competitive intelligence as a strategic capability: Evidence from higher vocational education administration*. *Journal of Sustainable Competitive Intelligence*, 16, e0646. <https://doi.org/10.37497/eagleSustainable.v16i.646>
- Rahhal, I., Kassou, I., & Ghogho, M. (2024). Data science for job market analysis: A survey on applications and techniques. *Expert Systems with Applications*, 251, Article 124101. <https://doi.org/10.1016/j.eswa.2024.124101>
- Song, X., & Xu, D. (2024). More graduates, fewer skills? Vocational education expansion and skilled labour shortages in China. *The China Quarterly*, 260, 970–985. <https://doi.org/10.1017/S0305741023001856>
- Schmitt, M. (2023). Automated machine learning: AI-driven decision making in business analytics. *Intelligent Systems with Applications*, 18, 200188. <https://doi.org/10.1016/j.iswa.2023.200188>
- Teece, D. J. (2023). *Big tech and strategic management: How management scholars can inform competition policy*. *Academy of Management Perspectives*, 37(1), 1–15. <https://doi.org/10.5465/amp.2022.0013>
- Tian, J., Wang, P., & Zhu, D. (2024). Overview of Chinese new energy vehicle industry and policy development. *Green Energy and Resources*, 2(2), Article 100075. <https://doi.org/10.1016/j.gerr.2024.100075>



- Tividad, M. J. I. (2024). The English language needs of a technical vocational institution. *Psychology and Education: A Multidisciplinary Journal*, 16(3), 256–283. <https://doi.org/10.5281/zenodo.10462501>
- Tzimas, G., Zotos, N., Mourelatos, E., Giotopoulos, K. C., & Zervas, P. (2024). From data to insight: Transforming online job postings into labor-market intelligence. *Information*, 15(8), Article 496. <https://doi.org/10.3390/info15080496>
- Walter, Y. (2024). Embracing the future of artificial intelligence in the classroom: The relevance of AI literacy, prompt engineering, and critical thinking in modern education. *International Journal of Educational Technology in Higher Education*, 21, Article 15. <https://doi.org/10.1186/s41239-024-00448-3>
- Weichselbraun, A., Süsstrunk, N., Waldvogel, R., Glatzl, A., Braşoveanu, A. M. P., & Scharl, A. (2024). Anticipating job market demands—A deep learning approach to determining the future readiness of professional skills. *Future Internet*, 16(5), Article 144. <https://doi.org/10.3390/fi16050144>