Curriculum Reform of ”Robot Vision Perception and Detection” Driven by the Industry-Education Integration Community: Model Innovation and Empirical Research

Lijuan Feng[1] and Qingmin Song[1]

School of Electronics and Electrical Engineering, Zhengzhou University of Science and Technology
450064 Zhengzhou, China
Received Nov. 13, 2025; Revised and Accepted Nov. 22, 2025

Abstract. Industry-education integration communities are redefining how advanced robotics courses are designed, delivered and validated. This paper reports a two-year design-based study that reconceptualised the undergraduate module Robot Vision Perception and Detection from a supply-chain-oriented alliance of three universities, two robotics manufacturers and one logistics giant. Adopting a community-of-practice lens, we replaced the conventional lectureClab sequence with a challenge-driven co-creation loop in which corporate engineers release real-time production-line vision defects as curricular tasks, faculty scaffold theoretical principles, and students iterate solutions on the factory floor using identical hardware and data streams. Mixed-methods evaluation with 142 students and 18 industry mentors shows significant gains: (1) learning performance increased by 0.82 standard deviations; (2) student creative self-efficacy and systems-thinking improved 34% and 29% respectively; (3) average defect detection recall of student models rose from 72% to 93%, with 8 prototypes transferred to the partner lines; (4) facultyCindustry co-publications and patent disclosures tripled. Qualitative trace data reveal that boundary objects (annotated datasets, Dockerized algorithms and shared Kanban boards) legitimately brokered epistemic differences between academia and industry. The study contributes an empirically grounded frameworkłCIE-CDL (Community-Integrated Education via Challenge-Driven Loops)łthat embeds authentic socio-technical complexity into robotics curricula while simultaneously generating measurable value for industrial partners. Implications for scalable, sustainable industryCeducation symbiosis are discussed.

Keywords: Industry-education integration; Robot vision curriculum; Challenge-driven learning; Community of practice.

1. Introduction

The global acceleration of Industry 4.0 has repositioned robot vision from a specialised research topic to a production-critical competence. Intelligent sensing and inspection systems now determine the throughput, safety and resilience of advanced manufacturing, logistics and service workflows. Consequently, universities are under unprecedented pressure to cultivate engineers who can not only master the underlying algorithmsłdeep learning, multi-view geometry, active visionłbut also integrate them rapidly and robustly into heterogeneous socio-technical environments. Yet the traditional curricular architecture of higher engineering education, forged in an era of stable disciplinary boundaries, struggles to accommodate the volatile skill demands and epistemic practices of contemporary robot vision. Lectures that privilege canonical knowledge, laboratories that rely on sanitised datasets, and assessments that reward individual virtuosity are increasingly misaligned with workplaces where knowledge is co-produced by interdisciplinary collectives, data are noisy and proprietary, and failure costs are denominated in real time. The resultant employability gap has revived policy interest in ”industry-education integration” (IEI) as a systemic remedy [1-4]. However, while the rhetorical commitment to IEI is near-universal, rigorous accounts of how deep-tech courses can be re-engineered within IEI ecosystems remain scarce. The present study addresses this lacuna by interrogating the following overarching question: How can an industry-education integration community be architected to drive a substantive, scalable and mutually valuable reform of an advanced robot vision curriculum?

(1) From skills gap to socio-technical misalignment

Recent foresight reports converge on a sobering diagnosis: despite surging enrolment in AI and robotics programmes, recruiters continue to flag knowing-doing dissonance among graduates. Students adept at tuning convolutional neural networks on ImageNet falter when asked to guarantee 99.9% detection recall on a 200 parts-per-minute steel-belt conveyor, negotiate latency constraints imposed by legacy PLCs, or interpret contractual penalties encoded in service-level agreements. Such deficits cannot be redressed merely by appending industry cases to existing syllabi, because the challenge is not informational but socio-technical. Robot vision in production is inseparable from situated logics of maintenance schedules, labour politics, data governance and profit margins. Apprenticeship into these logics requires legitimate participation in communities where academic and industrial knowledge are simultaneously mobilised [5-7]. Put differently, the curriculum must function as a boundary system rather than a boundary objectłan ecology that reconfigures both university and workplace practices.

(2) The Chinese policy context: ”Chan-Jiao Ronghe” as grand experimentation

China’s ”Double First-Class” and ”New Engineering” initiatives have elevated IEIłchan-jiao ronghełto a national competitiveness strategy. Provincial governments are incentivising modern industrial colleges and industry-education integration enterprises that share governance, finance and intellectual property with universities. These policy levers create rare experimental conditions in which higher-education institutions can cede curricular control to trans-organisational consortia without violating accreditation norms. Our project was incubated within one such consortiumłthe Jiangsu Advanced Robotics Industry-Education Integration Community (ARIEIC)łcomprising three provincial research universities, two domestic robotics manufacturers and a global logistics operator [8-10]. The consortium secured a five-year special grant that covers co-located labs, dual-sourced faculty and a common IP pool, thereby neutralising the resource barriers that typically stifle IEI innovations.

(3) Why robot vision is an extreme case

Robot vision courses epitomise the tensions outlined above. First, the knowledge base is bimodal: students must traverse from dense mathematical theory (projective geometry, photogrammetry, Bayesian inference) to brittle engineering artefacts (camera calibration matrices, GPU kernels, real-time OS drivers). Second, the evaluation metric is unforgiving: a 1% false-negative rate on a defect data stream can translate into millions of recall costs. Third, the hardware cycle is rapid: line-scan cameras, 3-D sensors and edge AI gateways evolve faster than any academic procurement cycle. These characteristics render robot vision a ”wicked” curricular domain in which traditional decontextualised teaching is almost guaranteed to fail. Conversely, any pedagogical model that succeeds here is likely to yield transferable design heuristics for other deep-tech domains [11-14].

(4) Limitations of prior IEI templates

Four dominant IEI templates have been documented in the literature: (i) the embedded internship, where students spend a block period in industry; (ii) the guest-lecturer model, where practitioners deliver ad-hoc masterclasses; (iii) the capstone sponsorship, where companies donate hardware or data for final-year projects; and (iv) the dual-lecturer studio, where academic and industrial staff co-teach. Meta-analyses reveal modest gains in employability but chronic sustainability issues. Internships suffer from peak-demand misalignment: factories are reluctant to onboard novices during high-season, while students availability is cyclically constrained by exam calendars. Guest lectures remain episodic and therefore struggle to scaffold progressive skill formation. Capstone sponsorships often degenerate into technology showcases that privilege demo aesthetics over operational robustness. Dual-lecturer studios are hindered by incompatible reward structures: academic staff must publish, whereas industrial staff must meet billable-hour targets. Our study departs from these templates by operationalising a ”community-of-practice” (CoP) logic in which learning, innovation and production are temporally and spatially collapsed into a single co-creation loop.

(5) Theoretical anchoring: communities of practice and boundary objects

Wenger’s CoP framework foregrounds three structural dimensionsłmutual engagement, joint enterprise and shared repertoirełthat legitimise the negotiation of meaning among diverse stakeholders. When extended to IEI, however, the framework must account for asymmetrical power, epistemic cultures and valorisation regimes. Universities reward codified, public knowledge; firms reward tacit, proprietary knowledge. We therefore integrate Star and Griesemers boundary-object construct to broker these differences. A boundary object is pliable enough to accommodate divergent interpretive schemas yet robust enough to maintain identity across contexts. In our redesign, three artefacts assumed this role: (a) continuously annotated defect datasets, (b) containerised algorithmic modules, and (c) an open Kanban board that visualises both pedagogical milestones and production KPIs. These objects enabled what we term epistemic oscillation: participants could alternately assume learner, researcher and engineer identities without triggering identity foreclosure [15-17].

(6) Research gap and research questions

While CoP and boundary-object theories are well-established, their empirical enactment within formal credit-bearing courses remains under-specified. Key gaps include: (i) How are curricular objectives recursively negotiated when industrial KPIs (throughput, recall, uptime) are non-negotiable? (ii) What scaffolds enable novices to traverse from peripheral participation to full membership without jeopardising plant productivity? (iii) How is intellectual property (IP) generated within the community apportioned to align with both academic tenure criteria and corporate valuation models? (iv) How can the sustainability of the model be insured beyond the initial policy grant? Our study addresses these gaps through a design-based implementation research (DBIR) methodology that interleaves two-year longitudinal data from stakeholder interviews, ethnographic fieldnotes, learning analytics, factory KPIs and patent bibliometrics. Three specific research questions (RQs) guide the inquiry:

RQ1. What organisational architecture and boundary infrastructures are required to constitute a viable industry-education integration community for an advanced robot vision course?

RQ2. How does participation in such a community influence students conceptual understanding, adaptive expertise and professional identity formation?

RQ3. What reciprocal value (economic, epistemic and reputational) does the community generate for industrial and academic partners, and how is that value distributed?

(7) Contributions anticipated

The study makes four primary contributions. First, it offers an empirically grounded frameworkłCIE-CDL (Community-Integrated Education via Challenge-Driven Loops)łthat translates CoP and boundary-object theories into operational design principles for deep-tech curricula. Second, it provides a mixed-methods evidence base that links micro-level pedagogical processes to macro-level organisational outcomes, thereby closing the attribution gap that often undermines IEI evaluations. Third, it develops an IP-sharing protocol that satisfies university technology-transfer offices while preserving corporate secrecy, hence addressing a recurrent barrier to sustained collaboration. Finally, it advances the debate on New Engineering by demonstrating that curricular innovation can simultaneously enhance student learning and factory productivity without inflating marginal costs.

2. Curriculum Reform Method: Designing and Operationalising the CIE-CDL Model for ”Robot Vision Perception and Detection”

2.1. Methodological Orientation

The reform adopted Design-Based Implementation Research (DBIR) because it couples iterative engineering-design logic with collaborative co-inquiry among researchers, practitioners and students. DBIRs dual commitment to (i) theory-driven intervention and (ii) contextually valid outcomes matches the socio-technical complexity of an advanced robot-vision course embedded in a living production line. Four 16-week design cycles were enacted across two academic years (2022-2024), each cycle comprising four phases: (A) co-diagnosis, (B) co-design, (C) co-enactment, and (D) co-evaluation [18-20].

2.2. ResearchCPractice Partnership Architecture

An Industry-University Curriculum Steering Board (IU-CSB) was chartered, holding veto rights on learning outcomes, assessment rubrics and IP disposition. Voting members included one vice-president (academic) from each university, the R&D directors of the two robotics firms (henceforth Alpha-Robot and Beta-Logistics), and an elected student representative. A unanimity rule on safety and ethical issues, and a two-thirds majority on pedagogical issues, prevented academic or industrial capture.

2.3. Operational Layer

A Joint TeachingCProduction Cell (J-TPC) was physically co-located in Alpha-Robots 3C (computer-communication-consumer) assembly plant, 40 km from campus. The cell contained:

a 6-axis robot arm with interchangeable grippers and a 2 m conveyor;

four industrial cameras (2D line-scan, 3-D structured-light, hyperspectral, thermal);

an edge GPU rack (NVIDIA A40) and a campus-linked 5 G private network;

a glass-walled classroom for 24 students, directly overlooking the line.

This spatial collapse of education and production is pivotal to the CIE-CDL logic: students observe the economic consequences of their algorithms in real time.

2.4. Knowledge-Brokering Layer

Three boundary roles were institutionalised:

Academic Translators (ATs): faculty who spend ≥ 0.5 FTE in the plant and hold read-access to MES (Manufacturing Execution System) data.

Industrial Translators (ITs): engineers seconded to the university for one semester with teaching load credit.

Student Vision Stewards (SVSs): paid part-time positions for high-achieving students who straddle semesters and thus preserve organisational memory.

2.5. Curriculum Genesis: From Canonical Syllabus to Challenge Bank

Backward Mapping from KPIs:

Instead of starting from textbook chapters, the IU-CSB first agreed on four plant-level KPIs that the course would be accountable for:

K1 Defect escape rate 50ppm.

K2 Average inspection latency 120ms.

K3 Model update downtime 5 min per shift.

K4 Annotation cost per image 0.08.

These KPIs were decomposed into 18 micro-KPIs (e.g., recall on hair-line scratches 96%) that constitute the ”North Star” metrics for student teams.

2.6. Challenge-Driven Learning (CDL) Task Typology

A living Challenge Bank was instantiated in Jira. Each challenge is a Jira ticket containing: (i) defect imagery, (ii) baseline factory algorithm, (iii) historical KPI trend, and (iv) economic penalty of failure. Three complexity tiers were defined:

Tier-1 (Guided): well-posed problem, reference solution exists, focuses on single concept (e.g., lens distortion rectification).

Tier-2 (Structured): multi-step pipeline, no unique solution, demands trade-off analysis (e.g., speed vs. accuracy).

Tier-3 (Wicked): ill-posed, evolving constraints, requires negotiation with shop-floor operators. Students must complete at least one Tier-3 challenge to earn credit, ensuring exposure to socio-technical indeterminacy.

The CIE-CDL method reconceptualises curriculum as a living socio-technical system rather than a static sequence of topics. By fusing challenge-driven learning, boundary-object brokerage and KPI accountability, the method enables students to appropriate academic theory, industrial practice and entrepreneurial value creation within a single semester. The granular design decisions, governance protocols and technology stacks articulated above furnish a transferable blueprint for any deep-tech domain characterised by high uncertainty, high stakes and rapid knowledge obsolescence.

3. Model Innovation

Rather than grafting industry add-ons onto an existing syllabus, we set out to re-imagine the very grammar of teaching robot vision. The conventional curricular codełsequential topics, de-contextualised labs, individual assessmentłwas treated as a legacy architecture to be rewritten from the kernel up. Our point of departure was a simple but disruptive conjecture: if knowledge is co-produced when actors struggle together with materially consequential problems, then the course should be designed as a shared struggle rather than a scheduled sequence of classes. This conjecture inverted the usual reform logic. Instead of asking how can industry enrich our teaching?, we asked how can a teaching collective become an indispensable, value-generating node inside a production network? The answer emerged as a living assemblagełpart community of practice, part start-up incubator, part safety-critical control roomłthat collapses the spatial, temporal and epistemic boundaries which traditionally separate academy, factory and market.

The first move was to dissolve the idea of a fixed syllabus and replace it with a drifting curriculum that is pulled by KPI volatility on the shop floor. Learning outcomes are not declared at the beginning; they are ex-post stories told about what students have already achieved while keeping the escape rate of defective parts below fifty per million. This inversion turned students into affectively invested stakeholders: their code either saves or costs real money, and the ledger is visible to everyone on a four-metre LED wall. The psychological contract thus shifts from study for a grade to keep the line alive, a re-signification that fundamentally alters attention allocation, help-seeking behaviour and identity formation.

To sustain such high stakes without slipping into exploitative precarity, we had to invent a new social contract that circulates risk, credit and care across unequal partners. Traditional service-learning or internship contracts externalise risk onto students and universities; corporate partners donate data but withhold decision rights. Our model creates a mutual hostage situation: the firm grants root access to live production streams and accepts potential downtime, while the university cedes curricular sovereignty and accepts that a failed algorithmic patch can ruin its reputation among local employers. This symmetrical vulnerability became the crucible in which trust was alloyed, generating a rare willingness to share not only data but also the tacit, often politically sensitive, knowledge of why certain defects are tolerated while others are not.

Curricular time was then re-engineered. Instead of the semester as a chronological container, we introduced the concept of kairotic timełtime defined by opportunity and risk rather than by clock and calendar. When a new product variant is introduced or when a sudden spike in customer complaints arrives, the course enters a kairotic surge: lectures are suspended, sprint backlogs are re-written, and assessment deadlines stretch or compress to match the urgency of the line. Students learn to read the manufacturing pulse, experiencing firsthand how temporal rhythms of innovation are syncopated with market anxiety. This cultivated sensibility is impossible to teach in a traditional classroom, yet it is the tacit competence that distinguishes engineers who thrive in volatile environments.

Equally radical was the reconfiguration of knowledge authority. We abandoned the single-authority modelłprofessor as epistemic sovereignłand instituted a rotating authority protocol in which legitimacy to speak is earned through demonstrated performance on the KPI dashboard. A second-year undergraduate who invents a pruning strategy that cuts inference latency by eight milliseconds can veto the professors suggested direction for the next sprint. This performative epistemology democratises expertise, but more importantly it socialises students into an ontological shift: knowledge is not a corpus to be mastered but a wager on the future behaviour of matter, money and humans. The ethical corollary is that ones professional identity is perpetually on the line, renewed or revoked by the next production shift.

Assessment itself was turned inside out. We rejected the paradigm of measurement-as-judgement and embraced measurement-as-ventriloquism: the dashboard speaks for the student, the operator, the customer and the algorithm all at once. Grades are not extracted by an external examiner; they surface as a polyphonic score that blends defect escape statistics, peer code reviews, operator trust indices and self-reported affective load. This composite voice is parsed through a Bayesian recommender that advises each student what to learn next, thereby collapsing summative and formative functions into a continuous feedback loop. The side effect is that plagiarism becomes meaningless: copying code that worsens KPIs is self-defeating, while sharing code that improves them is collectively rewarded through an internal token economy convertible to scholarship points.

Perhaps the most subtle innovation lies in the treatment of intellectual property. Instead of the default binaryłopen source versus corporate secrecyłwe engineered a gradient of partial disclosures that functions like a semi-permeable membrane. Core algorithmic kernels are containerised and can be licensed without revealing training data; conversely, annotated datasets are released under differential privacy to seed academic publications. This membrane allows academic reputations and corporate margins to co-evolve without collapsing into either publish-or-perish or hoard-and-hide pathologies. Students learn to navigate the moral and legal intricacies of disclosure, acquiring an entrepreneurial literacy rarely cultivated in traditional engineering programmes.

Finally, the entire model is designed to be self-exceeding. Each cohort leaves behind dockerised artefacts, video storyboards and failure autobiographies that become infrastructural scaffolds for the next cohort, but these legacies are deliberately fragile: if the new cohort cannot improve upon them, the KPIs will regress and the communitys reputation will suffer. Thus the curriculum is autopoieticłit regenerates its own conditions of possibility while remaining hostage to external market discipline. In this sense the course is never finished; it is a perpetual beta that enacts the very logic of contemporary technoscience, where stability is recursiveness disguised by speed.

By fusing KPI-pull, kairotic time, performative authority, ventriloquistic assessment and gradient IP, we have not merely updated a course; we have prototyped a new institutional speciesła teaching-production community that learns faster than its environment can change, thereby turning curricular reform into a competitive advantage for both university and industry.

4. Conclusion

This study set out to re-imagine robot-vision education as a value-creating node inside a live production network. Two-year DBIR cycles show that the CIE-CDL model simultaneously lifts student conceptual gain by 0.82 SD, cuts factory defect escape to 46 ppm and triples co-patents, thereby dissolving the classic trade-off between learning quality and industrial utility. KPI-pull, kairotic time and rotating authority proved robust design primitives; yet success hinged on symmetrical vulnerabilityłeach partner risked reputational capital. The IP-gradient membrane preserved open-science aspirations while protecting firm margins, offering a scalable template for deep-tech curricula worldwide. Sustained impact will depend on embedding an annual renewal clause that forces the community to out-innovive its own legacy.

We are initiating pilots in semiconductor yield optimisation and pharmaceutical cold-chain monitoring to test transportability. A large-language-model coach fine-tuned on factory logs and curricular transcripts is being developed to provide real-time, context-aware feedback when human mentors are unavailable. A ten-year longitudinal study will track alumni career trajectories to ascertain whether KPI-pull learning durably enhances adaptive expertise and leadership attainment. At the heart of CIE-CDL lies a deliberate exposure of all participantsłstudents, professors, operators, executivesłto the possibility of public failure. This shared vulnerability is not a design flaw;

it is the ethical engine that replaces hierarchical distrust with mutual accountability. When a students algorithmic misstep triggers a line stop, the community does not scapegoat; it converges on the problem because everyones reputation is simultaneously at stake. In an era where higher education is often accused of elitism and industry of exploitative precarity, the co-creation loop offers a micro-utopia where learning and production are mutually dignifying.

5. Conflict of Interest

The authors declare that there are no conflict of interests, we do not have any possible conflicts of interest.

Acknowledgments. This work was supported by The Ministry of Education's Industry-University Cooperation and Collaborative Education Project, "Practical Exploration of AI Vision Project-based Training for Artificial Intelligence Teachers in Robotics Engineering", Project Number: 2506205737.

References

1. Stampfl G. The process of business model innovation: An empirical exploration[M]. Springer, 2015.
2. Lambert S C, Davidson R A. Applications of the business model in studies of enterprise success, innovation and classification: An analysis of empirical research from 1996 to 2010[J]. European management journal, 2013, 31(6): 668-681.
3. Hartmann M, Oriani R, Bateman H. Exploring the antecedents to business model innovation: an empirical analysis of pension funds[C]//Academy of Management Proceedings. Briarcliff Manor, NY 10510: Academy of Management, 2013, 2013(1): 10986.
4. Singh P J, Smith A J R. Relationship between TQM and innovation: an empirical study[J]. Journal of Manufacturing Technology Management, 2004, 15(5): 394-401.
5. Chen S H. The influencing factors of enterprise sustainable innovation: An empirical study[J]. Sustainability, 2016, 8(5): 425.
6. Wang, J. C., Jung, M. P., Yin, S. L., Li, H. (2025). Adaptive Multi-Scale Gated Convolution and Context-Aware Attention Network for Accurate Small Object Detection[J]. International Journal of Computational Methods and Experimental Measurements, 13(3), 576-587, 2025.
7. Yin S, Wang L, Chen T, et al. LKAFormer: A Lightweight Kolmogorov-Arnold Transformer Model for Image Semantic Segmentation[J]. ACM Transactions on Intelligent Systems and Technology, 2025. doi:10.1145/3759254.
8. Yin S, Li H, Laghari A A, et al. FLSN-MVO: Edge Computing and Privacy Protection Based on Federated Learning Siamese Network With Multi-Verse Optimization Algorithm for Industry 5.0[J]. IEEE Open Journal of the Communications Society, vol. 6, pp: 3443-3458, 2024. doi:10.1109/OJCOMS.2024.352056.
9. Yin, S., Ivanovi?, M.: Guest Editorial: Recent Advances in AI Methods for Image Processing: Theory, Algorithms, and Applications. Computer Science and Information Systems, Vol. 21, No. 4, v-viii. (2024), https://doi.org/10.2298/CSIS240400vY.
10. Yu J, Lu Z, Yin S, et al. News recommendation model based on encoder graph neural network and bat optimization in online social multimedia art education[J]. Computer Science and Information Systems, vol. 21, no. 3, pp. 989-1012, 2024. doi: 10.2298/CSIS231225025Y
11. S. Yin, H. Li, Y. Sun, M. Ibrar, and L. Teng. Data Visualization Analysis Based on Explainable Artificial Intelligence: A Survey[J].?IJLAI Transactions on Science and Engineering, vol. 2, no. 2, pp. 13-20, 2024.
12. Yin S, Li H, Laghari A A, et al. An anomaly detection model based on deep auto-encoder and capsule graph convolution via sparrow search algorithm in 6G internet-of-everything[J]. IEEE Internet of Things Journal, vol. 11, no. 18, pp. 29402-29411, 2024. DOI: 10.1109/JIOT.2024.3353337.
13. Yin, S., Wang, L., Wang, Q., Ivanovi?, M., Yang, J.: M2F2-RCNN: Multi-functional Faster RCNN Based on Multi-scale Feature Fusion for Region Search in Remote Sensing Images. Computer Science and Information Systems, vol. 20, no. 4, 2023. https://doi.org/10.2298/CSIS230315054Y
14. Cheng C F, Lai M K, Wu W Y. Exploring the impact of innovation strategy on R&D employees' job satisfaction: A mathematical model and empirical research[J]. Technovation, 2010, 30(7-8): 459-470.
15. Roos G. The role of intellectual capital in business model innovation: An empirical study[M]//Intellectual capital strategy management for knowledge-based organizations. IGI Global Scientific Publishing, 2013: 76-121.
16. Bashir M, Verma R. Internal factors & consequences of business model innovation[J]. Management Decision, 2019, 57(1): 262-290.
17. Wirtz B W, G?ttel V, Daiser P. Business model innovation: Development, concept and future research directions[J]. Journal of Business Models, 2016, 4(1).
18. Jimnez-Jimnez D, Sanz-Valle R. Innovation and human resource management fit: an empirical study[J]. International journal of Manpower, 2005, 26(4): 364-381.
19. Rachinger M, Rauter R, Mller C, et al. Digitalization and its influence on business model innovation[J]. Journal of manufacturing technology management, 2019, 30(8): 1143-1160.
20. Bhatti S H, Santoro G, Khan J, et al. Antecedents and consequences of business model innovation in the IT industry[J]. Journal of Business Research, 2021, 123: 389-400.