An Exploration of Restructuring Learning Objectives in China's Higher Vocational Business Education in the Era of Multi-scenario Agents

Feng Han[1, *], Tianjing Xin[2]
[1]Zhejiang Business Technology Institute, Ningbo 310053, China
[2]Chiang Mai University, Chiang Mai 50200, Thailand
*Corresponding email: hanfeng@zbti.edu.cn

Abstract

Artificial intelligence (AI) agents are spreading quickly and are changing what many jobs consist of, and education has been slower to respond than most sectors. Agent products are also getting easier to use, so the technical set-up that once fell to the user is now built into the product. This paper asks what vocational education should teach once that shift has happened. It draws on Chinese national and provincial policy texts and on all 106.00 scenarios in the first batch of the Ningbo "AI + Manufacturing" scenario list. Two findings follow. What makes an agent act on its own is not the model but the program that runs it, so teaching one platform has a short useful life. And 31.10 per cent of the listed scenarios sit in business functions rather than engineering ones, where the scarce ability is to state a task clearly. It is also to judge whether the result is usable. Both rest on business knowledge. Three measures follow: restructure learning objectives across knowledge, skill, and literacy; embed AI literacy in existing professional courses instead of adding a separate one; and, once a curriculum updating mechanism is in place, change assessment first, training tasks second, and industry links last.

Keywords

Multi-scenario agent, Learning objectives, Vocational education, Business programmes

Introduction

National policy: Three documents give agents an official status

In August 2025, China's State Council issued the Opinions on Deepening the Implementation of the "AI + Initiative". It set a target: By 2027, more than 70 per cent of users should have adopted new-generation intelligent terminals and agents. Agents thus entered national policy as a measured target.

In April 2026, China's Ministry of Education and four other departments issued the Action Plan for "AI + Education". It asks institutions to "scientifically design the "AI + programme" system, curriculum system, and instructional system" and to "teach students to actively embrace, correctly understand, and appropriately apply artificial intelligence".

In May 2026, the Cyberspace Administration of China, the National Development and Reform Commission, and the Ministry of Industry and Information Technology jointly issued the Implementation Opinions on the Standardised Application and Innovative Development of Intelligent Agents. This is the first national document devoted to agents alone. It defines an agent as "an intelligent system possessing the capabilities of autonomous perception, memory, decision-making, interaction, and execution" [1].

This paper uses that definition. It narrows the object further to the multi-scenario agent, meaning one set of agent capabilities used across several business scenarios such as customer service, process handling, and data analytics. This is a working definition adopted for the present analysis.

Provincial policy: From principle to target

Local government followed quickly. Several Chinese provinces issued their own "AI + plans", and Beijing, Shanghai, Guangdong and others issued separate documents on AI in education. In the province relevant to this study, policy has centred on manufacturing. In July 2026, the Zhejiang Provincial Department of Economy and Information Technology and the Zhejiang Communications Administration jointly issued the Zhejiang Action Plan for Integrating the Industrial Internet with Artificial Intelligence. What provincial documents have in common is that agents are no longer a technical idea but a work task with numbers and deadlines attached.

Municipal policy: A sample of demand that can be checked

At the municipal level, Ningbo organised supporting measures around an open innovation hub for “AI + manufacturing”. In May 2026, its Bureau of Economy and Information Technology published the first batch of an application scenario opportunity list, covering 106 proposed scenarios declared by 47 firms [2].

This list is useful here for one reason. It is not a questionnaire written by researchers. Firms declared these scenarios themselves, based on their own problems, and a government department then reviewed and published them. The list is public, and every entry carries four fields: scenario category, subcategory, name, and description. Each entry can therefore be read and checked.

Where the demand sits: Three tenths of the scenarios are business functions

The list already places each scenario in one of seven categories. This paper groups those seven into three sets. Operations management, supply chain management, and product service are treated as business functions. Production and manufacturing, together with research and design, are treated as engineering functions. Industry governance and other innovation are kept separate. Each description was then read against its category label, and no entry needed to be moved.

The test for grouping is whether the job behind the scenario is a management job. This includes operational decision making, finance, procurement, supplier management, human resources, customer relations, warehousing and logistics, or general operations. Where the work happens, on the shop floor or in an office, is not the test. Five warehousing and logistics entries sit on the line between the two, since they involve both floor work and management. They are grouped as business functions because the job that owns them is a management job. The first author carried out the grouping, and both the original fields and the grouped result can be traced back to the published list.

The result is this. The document is titled as a manufacturing list, yet 33 of the 106 scenarios, or 31.10 per cent, are business functions: 17 in operations management, 11 in supply chain management, and 5 in product service. Production and manufacturing is still the largest single group, with 60 scenarios or 56.60 per cent. Business functions are therefore not the majority. The point is a different one. In a list about manufacturing, declared by manufacturing firms, close to a third of the stated demand still falls on business work.

Policy points the same way. The Implementation Opinions on Accelerating the Development of “AI + Consumption”, issued by China’s Ministry of Commerce and seven other departments in June 2026, call for deep integration of e-commerce with artificial intelligence across “operations, customer service, design, marketing, and livestreaming” [3]. These are the jobs that graduates of finance and commerce programmes in Chinese higher vocational colleges go into.

The lag in the educational response

Progress in education has lagged behind policy-making and industry development. Digital and Intelligent Marketing (programme code 530606) was added to China’s national vocational education programme catalogue in December 2025 and enrolled its first intake of students in autumn 2026, nine months after its formal inclusion.

During those nine months, the Action Plan for “AI + Education”, the Implementation Opinions on Intelligent Agents, and the Ministry of Commerce opinions on consumption were all issued. In other words, the first version of the training plan for this programme was written before the industrial and policy setting it was meant to serve had settled [4].

Rebuilding a curriculum takes longer than policy and industry take to change. This gap matters because of an imbalance. Generative AI bears most directly on the operating layer of competence, and the operating layer is where most higher vocational teaching content sits. Labour economists have started to examine how exposure to generative AI relates to wages by occupation, though the evidence so far comes from the United States.

Chinese work notes that generative AI reshapes not only how teaching is done but what the training is for [5].

What a multi-scenario agent is, and where its capacity to act comes from

The model does one thing

A large language model does one thing. It takes in a sequence of tokens and gives back a sequence of tokens. It does not browse the web, open files, or run code, and it remembers nothing between calls. When one says the model "uses a tool", what happens is that the model writes out a piece of text in an agreed format, and a separate program reads that text and carries out the action. The model makes a request. Another program does the work.

The capacity to act comes from the program that runs the model

That separate program treats the model's output as an instruction, carries it out, and feeds the result back to the model. It repeats this until the model stops asking for anything more. The program also decides how much material to keep in view, what the model is allowed to do, and when to stop. Together these functions turn a text predictor into a system that acts. Reviews describe the same shift, from static text-completion engines to agents that plan, reason, and carry out tasks in several steps. They locate the mechanism outside the model itself [6].

The central claim follows. An agent's capacity to act comes from the program around the model, not from the model. Give the same model no such program and it can speak but cannot act. This matches the standard definition quoted above, in which autonomous decision making and execution are part of the definition. Neither is something the model holds by itself.

One current practice is worth noting, because it is often misdescribed. Where several agents work on one task, they are usually not negotiating with each other. A coordinating agent hands out work and receives results, and the sub-agents do not talk among themselves. The main reason for creating a sub-agent is to keep material out of the coordinator's view: The sub-agent reads a large volume of text and returns a short conclusion. This matters for teaching. If students are taught to picture several agents talking things over, they are being taught a system that does not exist. What they need to understand is how a task is split up and how information is condensed as it passes between the parts.

What industry now needs: From operating tools to defining tasks

Two sides of an agent system, and the moving line between them

Building an agent system involves a set of technical arrangements. These include writing the loop that runs the model, managing what stays in view, setting permissions, deciding when to stop, describing the tools the model may call, assigning roles, and specifying how one step hands over to the next. All of this can be called orchestration.

In practice, this work falls on two sides, not on many separate parties. On the supply side are the firms that build agent products and platforms. They write the loop, manage the context, set the permissions, describe the tools, and define the handovers. In China, as elsewhere, the same firm often supplies the model, the end product, and the interface standard, so the supply side is frequently a single company. On the user side are the people who use the system: teachers, students, and staff in business roles. Their work is to state what the task is, judge whether the result can be used, and turn what works into a repeatable method.

The line between the two sides is moving upward. More of what used to be set up by the user is now built into the product, and shared interface standards are pushing the same way. A mature agent product can be recognised by the fact that its users rarely have to think about orchestration at all. The history of computing offers comparable cases, such as query optimisers taking over most hand tuning of database queries. In those cases jobs moved up rather than disappeared, and it is that upward move that concerns this paper.

The consequence for teaching is direct. A course built around describing tools and configuring roles is teaching work that the supply side is steadily absorbing. A course with a longer useful life has to aim at something else: How such systems behave. Model output is a proposal, not an action. The system will keep going until something stops it. And the quality of the whole depends on how well the task was stated at the start.

From single applications to several scenarios at once

What firms ask for is also changing. Two entries in the Ningbo list show this clearly. One asks for a single workbench that brings together three functions: a personal knowledge assistant, automatic meeting minutes, and searchable internal rules. Another asks for a sales and customer relationship agent that brings together five: meeting minutes, work logs, client record retrieval, draft proposals, and contract review.

National technical policy points the same way. The Implementation Opinions on the Special Action for "AI + Manufacturing", issued by China's Ministry of Industry and Information Technology and seven other departments in December 2025, call for work on "task planning and group collaboration" for industrial agents. They also call for work on "open and collaborative agent protocols and interfaces to improve inter-connection and interoperability" [7].

Demand is therefore moving from single tools toward sets of functions used together. Firms do not yet use the term multi-scenario agent, but what they describe already has that shape, and policy has turned openly toward getting agents to work with one another. Education needs to prepare for the form that is coming, not only the form already here.

What stays scarce is business ability

If the supply side keeps absorbing the technical set-up, what is left to the user is stating the task and judging the result. How far this holds depends on whether the two-sided account above is right, so the paper offers it as a proposition to be tested rather than as a settled finding.

Two things support it. First, both abilities rest on knowing what a correct outcome looks like in a particular business process. To state a task well, you must know what the right answer would look like. To judge a result, you must know the field well enough to see that an output is wrong. That knowledge comes from working in the field, and it does not expire when the tools change. Second, the grouping reported in this paper shows that close to a third of demand already sits in management work such as operations, finance, procurement, and customer relations. The core of such work is to set objectives and judge results, not to operate equipment.

Research on AI capability in leadership points the same way. A study of leaders in higher education finds that effective practice in AI-enabled settings "requires more than technical competence, extending to strategic awareness, ethical judgement, and the ability to critically evaluate the broader implications of AI technologies" [8]. That study is about managers, not students. What is taken from it here is its account of what competence consists of, not its target group.

This suggests something that still needs testing. As agents spread, a new kind of role appears: Someone who turns a business need into a clearly stated task and then decides whether the output can be used. That role calls for people who know the business, not people who can write code. Business programmes may therefore supply this role rather than merely absorb the disruption. The material used here is a list of scenarios, not a list of jobs, so nothing can yet be said about how many such posts exist or what they will require.

Restructuring learning objectives in three dimensions

How learning objectives are written is not a neutral matter. The wording carries a judgement about what knowledge is worth teaching [9]. This paper therefore starts from objectives rather than from teaching method. The Action Plan for "AI + Education" asks institutions to "teach students to actively embrace, correctly understand, and appropriately apply artificial intelligence". These three verbs match the three dimensions of learning objectives used in Chinese vocational curriculum documents: embracing maps to literacy, understanding to knowledge, and applying to skill. The restructuring below therefore follows the policy text rather than setting up a new framework. Similar three-part formulations have been used in Chinese higher vocational settings, where a progression from information literacy to digital literacy to AI literacy serves as a competence model.

Table 1 sets out the change. The left column is the authors' summary of objectives commonly found in tool-oriented courses in Chinese higher vocational business programmes. It is used here as a baseline rather than quoted from any one course standard. The right column is the position taken in this paper.

Table 1. Restructured learning objectives for higher vocational business programmes in China.

Dimension	Existing objective	Restructured objective
Knowledge	Memorise software functions, operating steps, and process rules	Understand what agents can and cannot do, and how they fail; understand how a business process can be divided and where it hands over
Skill	Operate tools and produce a deliverable	State a task clearly; judge a result; set boundaries and handle failure
Literacy	General professional norms and teamwork	Work with machines; take responsibility for AI-assisted output; keep learning; know when not to use AI

Knowledge objectives

What changes is the kind of knowledge required. Knowledge of procedures is losing value, while another kind, rarely taught until now, has become necessary. Remembering which menu performs which function stops being useful at the next version. Knowing that an agent gives different answers across runs, that it can only hold so much material in view, and that fluent writing is not the same as correct content remains useful after product changes. These are features of the technology itself, not of one product.

Another kind is practical knowledge of how business processes are built. To use an agent across a process, a student has to see where the process can be divided. They need to understand what passes between the parts, and where an error will likely slip through. It belongs in the courses of marketing, e-commerce, international trade, and business data analytics, not in a standalone technology course.

Skill objectives

(1) Stating a task

This means turning a vague request into a task with limits and a standard for acceptance. A common failure will be familiar to teachers. A student asks for “an analysis of this data”, gets a result, and then says it is not what was wanted. The fault is not in the tool. The request was never specified. Assessment should therefore look at the task statement itself, not only at the final product.

(2) Employer-side gap in the Ningbo list

Within one document, requirements quality varies greatly. Some entries give exact targets, such as cutting missed detections below a set rate, shortening inventory turnover by a set percentage, or reducing a decision cycle from days to hours. Others say only that automation is expected to improve efficiency. That spread suggests the ability to turn a business problem into a workable specification is uneven among firms.

(3) Judging a result

This means seeing that an output is wrong and being able to say why. It is the most knowledge-dependent of the three skills, and that dependence leads to a conclusion about curriculum design. Judgement comes from business knowledge, business knowledge comes from professional courses, and AI literacy should therefore sit inside professional courses rather than in a course of its own. One workable format is to give students AI-generated outputs with errors planted in them and ask them to find the errors and give reasons. Such items can be standardised and produced in quantity. This matches the stage in existing AI literacy frameworks at which the student acts as a synthesiser and evaluator rather than as an observer or operator [10].

(4) Setting boundaries and handling failure

This means deciding which steps need a person to confirm and how an error can be undone. Two opposite failures appear in practice: trusting everything and refusing everything. Both show a lack of judgement rather than caution. Assessment can present a business process and ask the student to mark points where a person must check, giving a reason for each.

Literacy objectives

This dimension moves from general professional norms toward working with machines. Its most practical part is responsibility for output. A student who hands in AI-assisted work signs their own name to it and answers for every figure in it. Put this way, technology ethics stops being an abstract topic and becomes a condition of assessment that can be enforced. The dimension also covers knowing when not to use AI at all, which is the other face of setting boundaries.

How to put this into practice

Curriculum

AI literacy should be built into professional courses rather than taught separately. Among programmes added to China's vocational catalogue in 2026, the two combining business with AI were placed at the vocational undergraduate level. No matching programme was added to finance and commerce at the associate degree level. New programmes at that level reach only newly admitted students. Students already enrolled still need their curriculum rebuilt. Rebuilding within existing programmes is therefore the most widely applicable path.

A mechanism for updating the curriculum

The nine-month gap described in this paper shows that revising a training plan once is not enough to keep up. One workable approach is to move the fastest-changing content down from the training plan to the level of course standards and training task sheets. This ensures that revising it does not have to wait for the whole programme to be adjusted. Teaching content can then be checked once each semester against a public industry document, such as the scenario list used here. This mechanism is the condition for the three changes that follow, as Figure 1 shows.

Assessment

Assessment should look at the process rather than the product, because the product can now be produced by the very technology being examined. Things that can be observed include the quality of the task statement, how many planted errors a student finds, and whether the checkpoints a student sets are the right ones.

Training tasks

Training tasks should move from closed problems to tasks with imperfect data. A task with clean data, one objective, and a fixed answer does not need an agent. An agent earns its place where the data are messy, the goal is unclear, and the requirement has to be settled over several rounds. Unless the tasks change, students cannot see what an agent actually solves.

Industry links

Computing resources and scenario resources should be kept apart. Computing power is hardware and can be bought. Real business processes and records of failure are scarce and held only by firms. Schools therefore need anonymised work orders, real breakpoints in a process, and honest reviews of what went wrong. Provincial evidence from vocational practice in China describes responses that combine programme upgrading, digital platforms, AI-enabled teaching, and training in teachers' digital skills. This suggests that infrastructure is one part of a response rather than the whole of it [11].

Order of implementation

Once the updating mechanism is in place, the three changes are best made in order, as set out in Figure 1. This sequence prevents resource fragmentation.

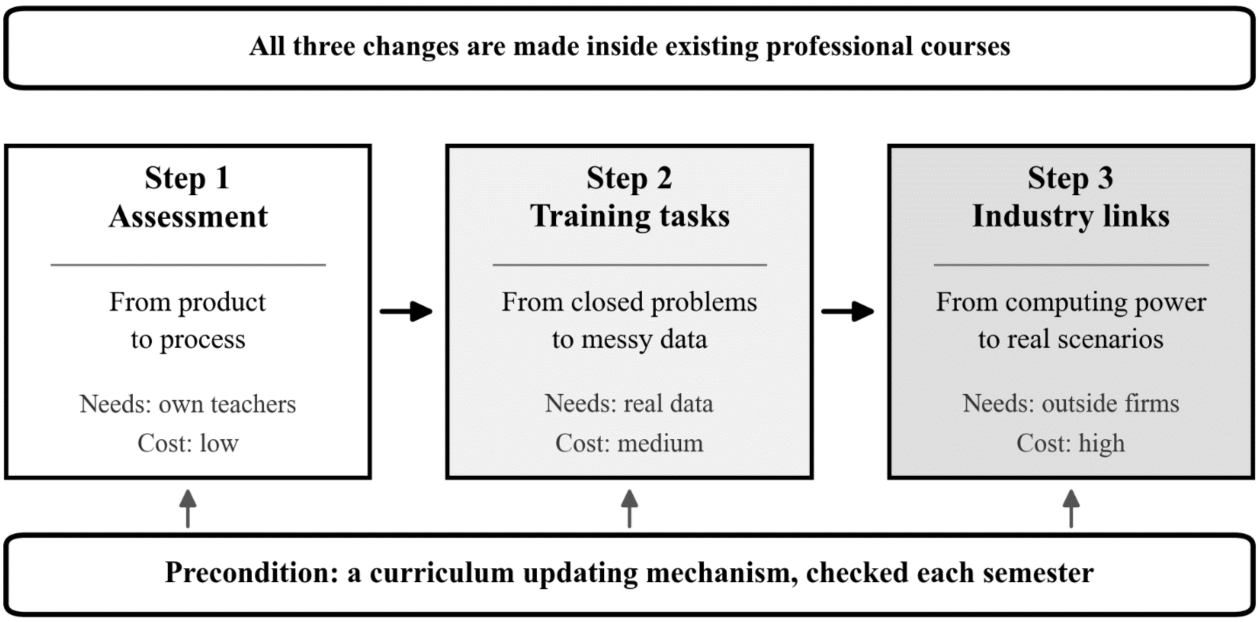

Dependence on parties outside the school rises from left to right

Figure 1. Implementation: scope, precondition, and order.

Note: All three changes are made inside existing professional courses, and the updating mechanism is the precondition for each. Dependence on outside parties rises from left to right. Source: drawn by the authors.

Assessment comes first. It decides what students and teachers pay attention to, and the formats proposed here need no new equipment and no outside agreement. Training tasks come second, since they need imperfect data but not an external partner. Industry links come last, because they depend on parties the school does not control. Doing all three at once spreads resources too thin, and starting from an industry agreement tends to produce resources that no assessment framework is ready to use.

Limits and risks

(1) Teachers

Most teachers have little practical experience of using agents, so training them is a condition for everything above.

(2) Handing over the thinking

Students may hand judgement itself to the technology, and judgement is the very ability this paper argues for. Assessment has to tell apart a student who can judge from one who has handed the judging over, and that problem is not yet solved.

(3) Data ethics

Real business data brought into teaching must be anonymised, and students must be told how it will be used.

(4) Limits of the evidence

The evidence here is one descriptive grouping of one list, from one city at one point in time. It was carried out by the first author alone, with no second coder and no consistency check. The restructuring in Table 1 has not been tried out in a pilot or reviewed by experts, so it is a position supported by material rather than a tested conclusion. Later work should add a second coder and consult programme leaders and firms. The list was also declared mainly by manufacturing firms, so what it says about the shape of demand should not be carried over to services as a whole.

(5) How far this travels

What makes the gap measurable here is a national catalogue that records the date a programme is added and the date students arrive. This is a feature of the Chinese system. Where each college sets up its own programmes, the same gap cannot be dated so precisely. The three-dimensional restructuring, by contrast, does not depend on that feature and can travel. The nine-month gap recorded here is a single case, and how such gaps vary across fields and levels is still an open question. A national catalogue makes them measurable, and measuring them would give the debate about how quickly vocational education responds an evidence base it now lacks.

Conclusion

What agents ask of vocational education is not one more technology on the syllabus. It is that being able to operate is worth less each year, while being able to state a task, judge a result, and set a boundary is worth more. The first can be bought in a product. The second cannot. Restructuring learning objectives is therefore more urgent than adding courses, and for business programmes at the associate degree level in China it is also the more realistic place to start.

For a college looking at any proposed AI curriculum, the question to ask is not only which platform it teaches. It is whether the content sits on the supply side or the user side, because content on the supply side has a date by which it will expire.

Funding

This work was supported by the Fundamental Research Funds Project of Zhejiang Business Technology Institute.

Acknowledgements

The authors would like to show sincere thanks to those technicians who have contributed to this research.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] Ferrag, M. A., Lakas, A., Tihanyi, N., Debbah, M. (2026) LLM and AI agents for autonomous systems: a survey of applications, datasets, and security challenges. *IEEE Open Journal of Intelligent Transportation Systems*, 7, 615-657.

[2] Lai, Y., Luo, J., He, X. (2026) Innovative industrial cluster policies and enterprise digital transformation: Does FinTech matter. *Applied Economics*, 1-21.

[3] Du, S., Ricafort, R. (2026) Ethical AI practices in E-commerce companies: evidence from Henan toward, China toward a policy guide. *International Journal of Sustainability and Advanced Integrated Research*, 2(2), 960-966.

[4] Li, Y., Cheng, H., Qin, Q. (2025) Evaluations and improvement methods of deep learning ability in blended learning. *International Journal of e-Collaboration (IJeC)*, 21(1), 1-17.

[5] Zhang, Y. (2025) AI-driven transformation of vocational education: opportunities, challenges, and future paths. *International Journal of Knowledge Management (IJKM)*, 21(1), 1-20.

[6] Alva, L. R., Pandey, B. (2026) Agentic AI systems in the age of generative models: architectures, cloud scalability, and real-world applications. *Artificial Intelligence Review*, 59(3), 88.

[7] Herrera, M., Pérez-Hernández, M., Kumar Parlikad, A., Izquierdo, J. (2020) Multi-agent systems and complex networks: review and applications in systems engineering. *Processes*, 8(3), 312.

[8] Bollaert, L. (2025) Artificial intelligence: Objective or tool in the 21st-century higher education strategy and leadership? *Education Sciences*, 15(6), 774.

[9] Hamdoun, W. M. A. (2023) Constructive alignment approach: enhancing learning and teaching. *British Journal of Multidisciplinary and Advanced Studies*, 4(2), 162-170.

[10] Wu, Y., Huang, Y. (2026) Faculty AI literacy in higher education systems: a systematic review and future-adaptive framework. *Interactive Learning Environments*, 1-31.

[11] Liu, S., Cheng, H. (2024) Manufacturing process optimization in the process industry. *International Journal of Information Technology and Web Engineering (IJITWE)*, 19(1), 1-20.