



AIGW Policy Observatory
Global AI Governance & Workforce Transformation

Global AI Governance and Workforce Transformation Policy Observatory

AIGW FLAGSHIP REPORT

JD.com: From Customer Service to AI Operations

How routine service automation is splitting human work into complex customer care, AI training and new technical roles

October 2026

A decorative graphic in the bottom right corner consisting of several concentric, curved lines in shades of blue and purple, creating a sense of motion or a stylized 'C' shape.

CONTENTS

Contents

Executive Summary	03
-------------------	----

Key Findings	06
--------------	----

Service Automation and Worker Transition	08
--	----

Workforce Architecture and Evidence	25
-------------------------------------	----

Policy, Governance and Durable Work	39
-------------------------------------	----

Measured Evidence	35
-------------------	----

Conclusion	54
------------	----

Notes and References	55
----------------------	----

EXECUTIVE SUMMARY

Executive Summary

JD.com offers one of the clearest public examples of a service workforce moving beyond a simple “human versus chatbot” model.

The company's customer-service automation journey began more than a decade ago. JD's JIMI customer-service robot entered testing in 2012 and launched publicly in 2013. By 2020, JD's Yanxi intelligent customer-service platform was already describing an operating model in which automated systems handled routine inquiries while judging conversation complexity and handing more difficult cases to human representatives. The platform also supported human staff with agent assistance, intelligent training, ticket creation, risk warnings and operational tools.[3][4]

The scale has since expanded dramatically. During JD's 2025 11.11 promotion, the company says its AI customer-service system handled more than 4.2 billion customer inquiries across business segments. By the end of the fourth quarter of 2025, more than 50,000 AI agents had been integrated into JD's internal systems across operations.[5]

Those figures establish AI at scale.

They do not tell us how many human jobs disappeared.

The more important workforce evidence appears in 2026.

Science and Technology Daily profiled a JD employee identified as Yuchen, who had nearly ten years of professional customer-service experience before becoming an AI trainer. His new work is not described as generic data labeling. When he finds gaps in an AI response, he uses product, category and merchant knowledge to diagnose whether the failure comes from system configuration or “intelligent logic”, then updates content, configures agents, or conducts maintenance and tuning.[1]

A later September 2026 report, republished by Xinhua and the Digital China Summit platform, added the scale missing from the individual story. A JD representative said that over roughly two years the customer-service organization had expanded multiple AI-related job categories and helped thousands of people complete job transformation or upgrading. JD described the emerging service model as:

AI handling basic services → humans concentrating on complex scenarios → multiple supporting roles enabling the system.[2]

That formulation captures the most important transition in this case.

Customer service is not disappearing as a single job.

It is decomposing into different forms of work.

EXECUTIVE SUMMARY

At least three human pathways are visible in the public record:

1. Complex-service specialists continue to handle situations that require judgment, domain knowledge, emotional understanding or operational action. 2. AI operations roles such as AI trainers configure, maintain, evaluate and improve the automated service system. 3. Adjacent technical roles are emerging around new technologies, including robot remote operation, data collection and robot maintenance.

JD's own service history supports this division. Its early JIMI documentation explicitly allowed unresolved issues to transfer to human representatives, and its later Yanxi system said it assessed conversation complexity before handing work to human staff. The 2026 workforce description therefore represents an evolution of a long-running human-machine service architecture rather than a sudden switch from "all human" to "all AI".[3][4]

The case also reveals a more subtle transition.

The AI trainer role is itself becoming more automated.

JD's current JoyAI customer-service product says large models can automatically analyze consultation conversations and documents, enhance knowledge and mine frequently asked

questions in order to reduce the cost of AI trainers configuring the knowledge base.[6]

This creates what AIGW calls a second-order automation problem:

"A worker may successfully move from an automatable frontline role into an AI-support role, only for routine parts of the AI-support role to become automatable next."

The durable human value therefore appears to be moving again - away from simple knowledge-base configuration and toward:

- diagnosing system failure; • understanding business context; • defining edge cases;
- evaluating outputs; • configuring workflows and agents; • governing escalation;
- translating domain expertise into system behavior.

This makes JD particularly important for workforce strategy.

It suggests that reskilling should not aim merely to move people into the first new job created by AI. It should build capabilities that remain useful as the AI operating layer continues to advance.

The case also has important evidence limits.

EXECUTIVE SUMMARY

JD disclosed more than 20,000 customer-service representatives at its main centers at the end of 2024.[7] Its 2025 annual filing continued to describe substantial 24/7 human customer-service operations and increased use of AI, but the customer-service section reviewed for this report did not repeat that headcount figure.[8] The public sources do not disclose:

- the current number of customer-service representatives; • how many of the “thousands” who transformed jobs came specifically from customer service; • the destination-role distribution;
- how many former customer-service positions were eliminated; • how many were not backfilled; • whether pay increased after transition;
- how many workers attempted but did not complete a transition.

It would therefore be wrong to describe JD's case as proof either that AI destroyed customer-service

employment or that AI transition protected every worker.

The strongest conclusion is narrower and more useful:

“JD has documented a scaled redesign in which routine service is increasingly automated, human service shifts toward complex scenarios, and some customer-service expertise is converted into AI-training and system-improvement work. The company reports thousands of job transformations or upgrades, but the full workforce reconciliation remains unpublished.”

For employers, this case offers a practical lesson:

When AI begins doing the frontline task, the transition question is not only “what will humans do instead?” It is “which parts of human expertise should move into the design, supervision and improvement of the AI system itself?”

KEY FINDINGS

Key Findings

01

1. JD's customer-service transition is a long-running operating-model change, not a single GenAI event

JD began deploying automated customer service in 2012-2013. Its service architecture evolved from chatbot response toward all-channel human-machine integration, intelligent routing, agent assistance, quality management and, later, large-model and agent capabilities.[3][4][6]

This matters because workforce redesign developed over years rather than after one product launch.

02

2. AI now handles service at massive scale

JD says its AI customer-service system handled more than 4.2 billion customer inquiries during the 2025 11.11 promotion.[5]

This is evidence of process scale.

It is not a count of unique customers, resolved cases, or human jobs substituted.

03

3. The human role is splitting rather than simply disappearing

JD's 2026 description of its service ecosystem is explicit:

- intelligent systems handle basic service; • humans concentrate on complex scenarios; • multiple job categories support the operating system.[2]

That implies task redistribution + role differentiation, not one uniform "customer-service job".

04

4. Domain expertise can become AI-operations expertise

Yuchen's nearly ten years of customer-service experience became useful for diagnosing AI-response failures, updating content, configuring agents and maintaining/tuning the system.[1][2]

The capability shift is not simply "learn prompting."

It is:

domain expertise + error diagnosis + system configuration + AI orchestration.

KEY FINDINGS

01

5. JD reports workforce transition at scale, but the denominator is missing

A JD representative said customer service had created multiple new AI-related job categories and helped thousands of people complete job transformation or upgrading over roughly two years. [2]

The public record does not provide:

- total eligible population; • completion rate; • source-role distribution;
- destination-role distribution; • pay or progression outcomes.

03

7. Human escalation remains part of the service model

From early JIMI to Yanxi and the 2026 service model, public JD materials repeatedly distinguish routine/basic service from cases requiring humans. [3][4]

The reviewed evidence supports human handoff and complex-case routing.

It does not specify formal approval, veto or override rights for every AI-generated response.

02

6. New roles do not automatically equal net job creation

JD's evidence supports new role categories and actual transitions.

It does not establish that the number of new roles exceeded any decline in older customer-service roles.

Role creation is not the same as net employment creation.

04

8. The AI-trainer role is itself a moving target

JoyAI currently advertises automated analysis of service conversations and documents, FAQ mining and knowledge enhancement specifically to reduce AI-trainer knowledge-base configuration cost.[6]

This means the transition pathway can itself be exposed to new automation.

The more durable role may be AI operations, not routine AI training.

CHAPTER 1

Service Automation and Worker Transition

The case begins with routine inquiry automation, then follows human escalation, specialist service and AI-operations work.

SERVICE TRANSITION

01

1. Why This Case Matters

Customer service is often treated as the canonical example of AI job substitution.

The logic seems straightforward.

Customers ask repetitive questions.

Conversational AI can answer many of them.

Therefore fewer humans should be needed per unit of service demand.

That logic is directionally useful but institutionally incomplete.

A real service organization contains much more than question answering.

It includes:

- transaction handling; • exceptions; • complaints;
- emotional recovery; • product expertise; • merchant coordination;
- returns and refunds; • fraud concerns; • logistics escalation;
- regulatory requirements; • knowledge management;
- quality assurance;
- workforce scheduling; • system maintenance.

As AI absorbs one layer of the work, the other layers become more visible.

JD is valuable because the public evidence shows that process happening at scale.

The case lets us distinguish four different questions that are often collapsed:

02

Can AI perform the task?

Yes. JD's AI systems handle enormous volumes of customer inquiries.

03

Does that change the human workflow?

Yes. Basic and routine service shifts toward automation while complex scenarios remain more human-centered.

04

Does that change human roles?

Yes. JD documents movement into AI training, robot operations and other technical/support roles.

SERVICE TRANSITION

01

Does that tell us the net employment effect?

No. The public evidence does not provide a complete headcount or worker-flow reconciliation.

The fourth answer is important.

A sophisticated workforce case should resist the temptation to turn strong process evidence into an unsupported employment conclusion.

02

2. Starting Point: A Large Human Service Organization

JD's customer-service transition matters partly because it began from a substantial in-house workforce.

JD's 2024 Environmental, Social and Governance Report says the company had built a professional customer-service team of more than 20,000 people, providing 24/7, multi-channel and full-scenario support.[7]

JD's 2024 Form 20-F similarly stated that, as of 31 December 2024, the company had more than 20,000 customer-service representatives across centers in Suqian, Wuhan, Chengdu and Datong.[16]

This was not a small outsourced contact-center pilot.

Customer service was a major operational capability and part of JD's customer-experience strategy.

03

Human service was already becoming more specialized

JD's 2024 ESG reporting also shows that the human service model was not simply a generic queue.

The company had established dedicated service teams for specific populations. For example, more than half of the maternal-and-infant service team held nursery-care qualifications in 2024.[7]

By May 2026, Science and Technology Daily reported that JD's self-operated maternal-and-infant team had reached 100% nursery-care certificate coverage. [1]

That detail matters because it illustrates one future direction for human service:

“When routine information can be automated, humans may become more valuable where service requires trusted domain expertise.”

This path is different from becoming an AI trainer.

Both are responses to the same underlying automation pressure.

TRANSITION TIMELINE

3. A Timeline: From Chatbot to Workforce Redesign

JD's service transition is easier to understand as a sequence.

2012-2013: the robot enters the queue

JD's JIMI customer-service robot entered testing in late 2012 and formally launched as a 24-hour service robot on 1 March 2013.[3]

The historical help documentation already reveals two principles that remain relevant today:

- JIMI handled defined business questions and routine service;
- customers could transfer to human service when the issue required operational action or further verification.[3]

The same page also allowed users to “teach” JIMI proposed questions and answers, which would be reviewed by backend staff before use.[3]

Even at this early stage, the system contained:

automation + human review + knowledge maintenance.

2017-2020: customer service becomes a human-machine operating system

JD's intelligent customer-service platform was launched in 2017 and later branded Yanxi.

A 2020 JD Cloud account described the move from a single-channel human model toward an all-channel “human + intelligent” service model.[4]

The system:

- automated routine queries;
- assessed conversation conditions and complexity;
- transferred difficult cases to human staff;
- provided agent assistance;
- delivered intelligent training;
- supported ticket creation and risk warning;
- supported scheduling, dispatch and other management decisions.[4]

During the 2020 11.11 period, JD reported very large interaction volumes through the platform, demonstrating that automation was already becoming operational infrastructure.[4]

TRANSITION TIMELINE

2024: human service remains large

Despite more than a decade of customer-service automation, JD still reported more than 20,000 human customer-service representatives at the end of 2024.[7]

This is an important corrective to simplistic technological determinism.

Large-scale automation and a large human service workforce can coexist.

That does not mean automation has no workforce effect.

It means the effect may appear through:

- task mix; • specialization; • staffing growth avoided;
- new roles; • redeployment; • hiring patterns;

rather than only through immediate elimination of the entire function.

TRANSITION TIMELINE

2025: generative and multimodal AI reaches massive operational scale

JD's 2025 full-year results say its multi-scenario, multimodal AI customer-service system handled more than 4.2 billion customer inquiries during the 11.11 promotion.[5]

The company also reported more than 50,000 AI agents integrated across internal systems by the end of the fourth quarter.[5]

Those figures mark a new scale of AI deployment.

But they remain process/adoption metrics.

They do not tell us how many inquiries would otherwise have required a person.

2026: the workforce transition becomes explicit

In May 2026, Science and Technology Daily reported Yuchen's transition from experienced customer-service worker to AI trainer and described adjacent workers moving into robot remote-operation work.[1]

In September, a JD representative said that during the previous two years the customer-service organization had created multiple new job categories around AI applications and helped thousands of people complete job transformation or upgrading.[2]

At this point the story is no longer only:

“JD automated service.”

It becomes:

“JD automated parts of service and reorganized some of the human workforce around the new system.”

SERVICE TRANSITION

01

4. How the Workflow Changed

The emerging JD service architecture can be summarized as three layers.

03

Layer B - human complex service

Humans concentrate on situations requiring more:

- judgment; • empathy; • domain expertise;
- cross-system coordination; • operational intervention; • ambiguity resolution.

This logic is consistent with JD's historical JIMI and Yanxi materials, which route unresolved or complex cases to human staff.[3][4]

02

Layer A - automated basic service

AI handles large volumes of:

- recurring questions; • standard information; • routine service interactions;
- knowledge retrieval; • some transactional workflows.

JD's own 2026 description says intelligent systems are taking basic service.[2]

04

Layer C - AI operations and support

A third workforce layer appears behind the conversation.

Workers maintain and improve the systems themselves.

Yuchen's tasks include:

- identifying answer gaps; • diagnosing whether the cause is configuration or system logic; • updating content;
- configuring agents; • maintaining and tuning the system.[1]

This creates a service operating model with both a frontstage and a backstage.

SERVICE TRANSITION

01 Layer C - AI operations and support - Frontstage Customer interaction: • AI handles routine work. • Human specialists handle complex work.	02 Layer C - AI operations and support - Backstage System operation: • AI trainers; • knowledge managers; • configuration specialists; • evaluators; • technical support roles; • adjacent robot/data roles.	03 Layer C - AI operations and support - AIGW interpretation The future of service work may not be a smaller version of the old contact center. It may be a different labor architecture: fewer humans doing repetitive first-line interaction per unit of demand + more specialized human service + new human roles managing the automation layer. The public JD evidence supports the architecture. It does not yet quantify the relative size of each layer.
---	---	---

WORKER TRANSITION


5. Worker Transition: Yuchen - Customer Service → AI Trainer

Service expertise becomes AI operations work.

FIGURE 1.1

reported work sequence

From service experience to AI-operations work



Source: Science and Technology Daily, 凝“新”聚力促就业 | AI为他们带来新身份--京东协同推进技术应用与就业转型探索实践, 2026

Note: Interpretive sequence based on the worker-level reporting.

This worker-level example does not establish a representative transition rate.

WORKER TRANSITION

Yuchen's reported pathway

Yuchen is the strongest worker-level evidence in the JD case. Science and Technology Daily reported in May 2026 that he had nearly ten years of professional customer-service experience and knowledge of product categories and merchant cooperation before moving into an AI-trainer role. [1] That prior experience became useful in a different way. Yuchen said that when he sees a gap in an AI-generated answer, he can anticipate what the correct customer experience should look like and diagnose whether the problem comes from:

- configuration; • intelligent logic; • content. He can then: • update content; • configure the agent; • maintain or tune the system.[1] This is a more sophisticated transition than: “service worker → data labeler.” The value of the worker lies partly in the ability to understand why the AI is wrong in business context. The capability chain — Yuchen's transition can be represented as: customer-service experience → domain knowledge → expectation of correct answer / experience → error recognition →

root-cause diagnosis → content or configuration change → AI system improvement

The transition is important because the old job supplies part of the training for the new job. The employee already understands: • customer intent; • product rules; • merchant processes; • common failure modes; • service expectations. AI capability is added on top. The capability chain - AIGW interpretation — The most valuable transition asset is not generic AI literacy. It is: “domain expertise that can be translated into model and workflow operations.”

This is similar to the mechanism seen in LONGi's inspector-to-AI-trainer transition, but in a different domain. At LONGi, defect knowledge moves into the quality system. At JD, customer-service knowledge moves into the conversational service system. The common pattern is: performing the task → teaching / configuring / validating the system that performs the task.

SERVICE TRANSITION

01

6. What an AI Trainer Actually Does

The term AI trainer can be misleading.

It is sometimes understood as low-skill data labeling.

China's national occupational framework is broader.

The Ministry of Human Resources and Social Security and the Ministry of Industry and Information Technology formally issued a national occupational skill standard for Artificial Intelligence Trainer, occupation code 4-04-05-05, in 2021.[9]

The official standard defines the occupation as work involving AI-training software and tasks such as:

- database management; • algorithm-parameter setting; • human-machine interaction design;
- performance testing and tracking; • related supporting work.[9]

The occupation has five skill levels, from junior worker through senior technician / advanced technician.[9]

The standard includes work ranging from:

- data collection and processing; • data labeling; • system operation and maintenance;

toward more advanced responsibilities such as:

- business analysis; • training-set and test-set construction; • algorithm testing;
- analysis of error cases; • human-machine interaction process design; • intelligent-system optimization;
- AI application solution design; • training and guidance.[9]

SERVICE TRANSITION

01

Why the national standard matters for JD

It does not prove that Yuchen holds a national AI-trainer credential.

The reviewed sources do not state that.

It does show that AI trainer is becoming an institutionalized occupation in China, with a skill architecture beyond annotation.

Yuchen's reported work aligns with several higher-value elements of that architecture:

- business-context analysis; • error diagnosis; • content/knowledge maintenance;
- agent configuration; • system tuning.

This makes the transition more credible as an occupational pathway rather than a temporary label created for one story.

02

7. JD Is Also Building an AI-Trainer Learning Infrastructure

JD's merchant learning platform contains a dedicated Jingxiaozi AI Trainer Certification series, including beginner and advanced courses.[10]

The platform is primarily a merchant learning environment, so this evidence should not be described as JD's internal employee certification system.

But it matters for another reason.

It shows that JD's customer-service ecosystem has developed:

- defined AI-trainer skills; • structured learning; • beginner and advanced progression;
- explicit certification language.

That is evidence of role formalization around the service technology.

03

What this does not prove

The public course page does not establish:

- how many JD employees completed the course; • whether Yuchen completed it; • whether certification is required for the internal AI-trainer role;
- whether the curriculum maps directly onto the national occupational standard.

Those distinctions should remain explicit.

SERVICE TRANSITION

01

8. A Broader Set of Worker Pathways

JD's workforce response is broader than one AI-trainer job.

The May 2026 Science and Technology Daily article describes several types of role expansion or transition.[1]

03

Pathway B - customer service → robot remote operations

Yuchen said some customer-service colleagues had moved into work related to remote robot operation.[1]

That work requires learning areas such as:

- robot operating logic; • command response; • movement control processes;
- basic robot repair in some situations; • data collection across scenarios.[1]

This is a larger occupational jump.

The transferable capability is less direct than in the AI-trainer pathway.

It suggests that JD's workforce transition can also involve cross-domain reskilling into new technology-enabled operations.

02

Pathway A - customer service → AI trainer

Yuchen is the clearest example.

The worker moves behind the service interface and helps improve the AI system.

04

Pathway B - customer service → robot remote operations - Evidence boundary

The source does not provide:

- the number of employees in robot remote operations; • training duration; • whether the roles are permanent;
- pay changes; • selection criteria.

The pathway is evidenced.

Its scale is not.

SERVICE TRANSITION

01

Pathway C - technical repair → robot maintenance

Science and Technology Daily also profiled Jin Lei, who had worked at JD for eleven years, beginning with computer repair and later becoming a robot maintenance engineer.[1]

This is not a customer-service-to-AI transition and should not be merged with the customer-service worker's pathway.

It is useful adjacent evidence because it reveals the same enterprise strategy:

“preserve an existing technical foundation, then extend it into a new technology domain.”

JD said it launched a plan to train more than 10,000 professional robot maintenance service workers over three years to support a growing robot repair network.[1]

That is a future training target, not a completed hiring or transition outcome.

02

Pathway D - generic service → certified domain service

JD's maternal-and-infant service team illustrates another direction.

By 2026, the team reportedly achieved 100% nursery-care certificate coverage.[1]

Here the pathway is not:

“customer service → AI job.”

It is:

“customer service → deeper human domain specialization.”

This is important because service automation can increase the relative value of:

- professional knowledge; • trust; • judgment;
 - human interaction;
- even when the employee does not move into a technical role.

03

9. The Service Job Splits Rather Than Simply Disappears

JD's emerging service architecture can be understood as a portfolio of human roles.

SERVICE TRANSITION

01

Role 1 - automated service operator: mostly AI

Handles:

- common questions; • repetitive requests; • high-volume standard interactions.

02

Role 2 - complex-case human specialist

Handles:

- ambiguous issues; • operational exceptions; • emotionally sensitive interactions;
- cases requiring deeper product or policy judgment.

03

Role 3 - domain specialist

Provides expertise where customers need trusted human knowledge, such as maternal-and-infant care.

04

Role 4 - AI trainer / AI operations specialist

Improves:

- knowledge; • agent configuration; • system behavior;
- failure diagnosis; • response quality.

05

Role 5 - quality / evaluation specialist

Potentially monitors:

- response quality; • error patterns; • escalation;
- safety; • service standards.

JD's public product architecture includes intelligent quality inspection and training, although the reviewed workforce sources do not disclose a specific employee transition into a dedicated evaluator role.[6]

06

Role 6 - adjacent technology operator

Examples include:

- robot remote operation; • data collection; • robot maintenance.

SERVICE TRANSITION

01

Role 6 - adjacent technology operator - AIGW interpretation

The customer-service function becomes less like one occupation and more like a service system containing:

AI execution + expert humans + AI operations + quality + technical support.

This matters for workforce planning.

Training every service employee for the same future role would make little sense.

The transition requires pathway differentiation.

02

10. Scale: What Does “Thousands” Actually Tell Us?

The September 2026 Tech Daily report provides the strongest scale claim.

A JD representative said that over the previous two years, JD customer service had:

- continuously expanded new AI-application job opportunities;
- created multiple job categories;
- helped thousands of people achieve job transformation or upgrading.[2]

The statement is unusually valuable because it moves beyond:

- a pilot;
- a single worker;
- a future training commitment.

It describes completed workforce movement at scale.

But the metric remains incomplete.

03

What it establishes

It establishes that:

- JD reports actual job transition/upgrading;
- the activity is not limited to one individual;
- more than one destination job category exists;
- the period is roughly two years.

SERVICE TRANSITION

01

What it does not establish

It does not tell us:

- whether “thousands” means 2,000 or 9,000; • how many were customer-service employees; • whether all transitions were AI-related;
- how many workers were eligible;
- how many applied; • how many completed training;
- how many moved into newly created roles; • how many moved into existing roles with upgraded content; • how many later remained in those roles;
- how compensation changed.

02

The denominator problem

The last clearly disclosed customer-service headcount in the reviewed filings is more than 20,000 at the end of 2024.[7]

That creates a temptation to calculate a transition rate.

We should not.

The definitions are not aligned.

The >20,000 figure covers customer-service representatives at named centers.

The “thousands” figure refers to people completing job transformation or upgrading across multiple job categories over roughly two years.

The source does not state that the numerator is a subset of the exact same denominator.

03

**The denominator problem
- Research rule**

Do not produce a percentage transition rate until JD publishes a comparable base population and methodology.

CHAPTER 2

Workforce Architecture and Evidence

This chapter separates interaction volume, job-transition claims and the workforce data JD has not yet reconciled.

WORKFORCE EVIDENCE

01

11. Process Scale Is Not Employment Scale

The most visually impressive number in the case is:
4.2 billion customer inquiries

JD says its AI customer-service system handled more than that amount during the 2025 11.11 promotion.
[5]

But an inquiry is not a job.

Several distinctions matter.

03

One inquiry is not necessarily a complete case

The public release does not define whether “inquiry” means:

- message; • turn; • session;
- request; • completed interaction.

05

Automated volume does not reveal human escalation

The public 2025 result does not disclose:

- autonomous resolution rate; • handoff rate; • complex-case rate;
- percentage requiring human action.

02

One customer can generate multiple inquiries

The number is not a unique-customer count.

04

Automated handling may absorb growth

During a major commerce festival, service volume can surge.

AI may allow an organization to handle additional demand without proportional human hiring.

That is economically important even if no existing job is removed.

06

Automated volume does not reveal human escalation - AIGW interpretation

The correct use of the 4.2 billion figure is:
“evidence of operational scale.”

The incorrect use is:

“evidence of 4.2 billion units of human labor replaced.”

WORKFORCE EVIDENCE

01

12. Human Escalation: A Persistent Design Principle

JD's service history repeatedly shows that human escalation is not an afterthought.

03

Yanxi

JD's 2020 account says the intelligent service system judges the state and complexity of a conversation and connects it seamlessly to human service when needed.[4]

05

What remains unknown

The reviewed public evidence does not specify:

- escalation thresholds; • whether AI can execute consequential actions without human review; • formal veto rights;
- override logging; • required human approval categories; • accountability when an automated answer causes harm;
- how confidence scores affect routing.

Therefore the case supports:

“human escalation and complexity-based routing”

but not a stronger claim that:

“humans formally approve or can veto all AI decisions.”

02

JIMI

Historical JD help documentation says that when a question requires customer-service action or additional verification, users should switch to human service.[3]

It also says users can transfer to humans when JIMI cannot solve the problem.[3]

04

2026 workforce model

JD's later description summarizes the operating logic as:

intelligent basic service + human complex scenarios + multiple supporting roles.[2]

This consistency across years supports a genuine human-handoff architecture.

06

13. The Capability Shift: From Answering Questions to Operating an AI Service System

The capability change is one of the strongest parts of the JD case.

WORKFORCE EVIDENCE

01

Traditional customer-service strengths

An experienced representative develops:

- product knowledge; • policy knowledge; • merchant-process knowledge;
- customer-intent recognition; • expectation management; • communication;
- experience with unusual edge cases.

03

What is retained?

The old knowledge is not discarded.

It becomes the basis for knowing:

- when the AI is wrong; • why the response will create a bad user experience; • what the correct business behavior should be.

02

New AI-operations strengths

Yuchen's new role adds:

- model/agent error recognition; • configuration diagnosis; • logic diagnosis;
- content maintenance; • agent configuration; • system tuning.[1]

The new job therefore combines:

domain expertise + AI fluency + system operations + judgment + feedback design

04

What is new?

The worker has to learn how that knowledge is represented in the technical system.

This is the essence of AI orchestration.

The worker does not need to become a foundation-model researcher.

The worker needs to become capable of changing how the AI system behaves in a real service environment.

WORKFORCE EVIDENCE

01

14. Domain Expertise as Transition Capital

JD reinforces a pattern already visible across several AIGW cases.

Workers often transition most effectively when the new job preserves some scarce knowledge from the old job.

That knowledge can be:

- customer knowledge; • product knowledge; • process knowledge;
- quality knowledge; • technical troubleshooting; • relationship context.

A useful transition equation is:

existing domain expertise + new digital / AI capability
= adjacent higher-value work

This is different from a generic reskilling model:

old worker → learn entirely new occupation →
compete from zero

The first path has lower transition friction because the employee carries forward an asset.

02

For Yuchen

The asset is nearly a decade of service experience.

03

For JD's maternal-and-infant team

The asset is human service capability combined with certified professional knowledge.

04

For robot repair roles

The asset can be prior device-repair and technical troubleshooting experience.

05

For robot repair roles - AIGW takeaway

The best workforce-transition map may start by asking:

“What knowledge does the worker already possess that the AI system or adjacent business still needs?”

Only then ask:

“What new capability must be added?”

WORKFORCE EVIDENCE

01

15. Second-Order Automation: The AI Trainer Is Not a Permanent Safe Harbor

The JD case contains a particularly important warning.

The new role created by AI may itself become exposed to AI.

JD's current JoyAI customer-service product says that large models can automatically:

- analyze customer-service conversations; • analyze documents; • enhance knowledge;
- mine FAQs;

in order to reduce AI-trainer knowledge-base configuration cost.[6]

That is direct evidence that part of the AI-trainer workflow is already a target for automation or compression.

03

Second transition

Large model:

“automatically analyzes conversations and documents; mines knowledge / FAQs; reduces configuration effort.”

The trainer's role moves again.

02

First transition

Human customer-service worker:

“answers routine customer question.”

AI takes more of that work.

Worker moves to:

“configure and maintain AI knowledge.”

04

What may remain more durable?

Activities requiring:

- business-context diagnosis; • error taxonomy; • edge-case judgment;
- workflow architecture; • evaluation; • risk analysis;
- escalation design; • interaction-policy decisions; • cross-functional coordination.

WORKFORCE EVIDENCE

01

What may remain more durable? - AIGW concept: the moving safe harbor

A reskilling strategy can fail if it moves workers only one step away from the task being automated.

A more durable strategy asks: “Which capabilities become more valuable as the automation system itself improves?”

This is one of the strongest takeaways from the JD case.

02

16. AI Trainer as a Career Ladder, Not a Single Job

China's national occupational standard is useful because it presents AI training as a progression.

At lower skill levels, work can include:

- data collection; • data processing;
- data labeling;
- basic system operations.

At higher levels, the work expands toward:

- business analysis; • intelligent training; • algorithm testing;
- error-case diagnosis; • human-machine interaction design; • solution design;
- training-process optimization; • guidance and training of other workers.[9]

The implication is important.

If routine labeling and knowledge-base maintenance become more automated, the career can move upward toward:

evaluation + system design + business integration + optimization.

Yuchen's reported work already sits closer to this higher-value side than to simple annotation.

03

Workforce design implication

Employers should avoid defining “AI trainer” too narrowly.

A better pathway might have explicit stages:

AI content specialist → AI trainer → AI evaluator / quality specialist → AI operations specialist → AI workflow designer / business translator

The public JD sources do not state that JD uses this exact ladder.

It is an AIGW analytical model derived from the tasks and occupational standard.

WORKFORCE EVIDENCE

01

17. Organizational Infrastructure: Transition Needs More Than a New Job Title

Yuchen's transition did not occur inside an organization with no learning infrastructure.

JD's 2024 ESG report describes a broad employee-development system with:

- structured training; • professional growth maps; • career pathways;
- internal mobility mechanisms; • continuing education programs.[11][12]

In 2024, JD reported:

- 100% employee training coverage; • an average of 41.0 training hours per employee; • more than 23.38 million total training hours;
- the opening of the JINGDONG Global Learning Center in Suqian.[12]

The company also said it upgraded its internal employee-mobility rules in 2024 so high-performing contributors could search for and select longer-term career opportunities internally according to professional ability and personal interests.[11]

These measures provide useful organizational context.

They do not prove that the thousands of AI-related customer-service transitions occurred through those exact mechanisms.

The public evidence does not disclose:

- Yuchen's training curriculum; • whether he used the internal mobility system; • the duration of his transition;
- who selected him; • whether the learning happened during paid work time; • whether the transition was voluntary.

02

AIGW interpretation

The infrastructure is still relevant because large-scale transition normally requires at least four organizational systems:

1. skill diagnosis - what can the worker already do? 2. learning - what new capability is missing? 3. mobility - where can the worker go?
4. work redesign - does a real destination role exist?

Training without mobility produces certificates.

Mobility without role design simply moves workers around.

JD's public evidence is strongest on the existence of destination categories and reported movement. It is weaker on the mechanics that connect each individual worker to a destination.

WORKFORCE EVIDENCE

01

18. Governance: AI Service Quality Is Also a Workforce Question

JD's 2024 ESG report identifies AI-related risks including:

- inaccurate or biased AI decisions; • sensitive-data leakage; • lack of transparency;
- ethical controversy.[13]

The company's listed responses include:

- ongoing technology optimization and monitoring; • data-protection and security controls; • an AI ethics framework and usage policy;
- continuing employee training to support responsible AI use and management.[13]

These are enterprise-level governance measures.

They are not specific proof of customer-service decision rights.

But they matter to the AI-trainer role because the trainer sits close to the boundary between: model behavior and customer experience.

An AI trainer who changes knowledge, configuration or agent behavior is not only performing a technical task.

The worker can influence:

- what the system says; • which information it relies on; • how a customer issue is interpreted;
- when an interaction succeeds or fails.

This creates an emerging governance responsibility.

02

Skills that may become governance skills

The Yuchen case suggests several operational skills with governance implications:

WORKFORCE EVIDENCE

01

Skills that may become governance skills - Error recognition

Can the worker notice when the AI gives a plausible but incorrect answer?

02

Skills that may become governance skills - Root-cause diagnosis

Is the error caused by:

- source content; • configuration; • reasoning / logic;
- workflow integration?

03

Skills that may become governance skills - Corrective action

Should the worker:

- edit content; • change configuration; • escalate to technical teams;
- restrict behavior?

04

Skills that may become governance skills - Service impact judgment

Could the response:

- mislead the customer; • create a financial/service problem; • damage trust;
- violate policy?

The public reporting does not specify whether Yuchen has formal authority over all of these decisions.

But the role is clearly moving toward operational governance of AI behavior.

MEASURED EVIDENCE

19. Measured Evidence: What We Know and What Each Number Means

FIGURE 2.1

reported result by analytical level

Process, workforce and AI-adoption measures

The JD case contains several powerful numbers. They answer different questions; they should not be collapsed into one workforce outcome.

Measure and analytical level	Reported result	Evidence boundary
Human customer-service representatives - Workforce stock	20,000 at end-2024	Last clearly disclosed comparable count in reviewed filings ^{[7] [16]}
Frontline customer-service salary increase - Workforce context	20,000 employees; avg annual pay +30% announced in 2024	Predates 2026 transition evidence; not AI outcome. ^[17]
AI customer-service inquiries - Process / scale	4.2 billion during 2025 11.11	Not unique customers, resolved cases or jobs ^[5]
AI agents in JD internal systems - Enterprise AI adoption	50,000 by Q4 2025	Not customer-service roles ^[5]
Job transition / upgrading - Workforce movement	Thousands over ~2 years	Company-reported via Tech Daily/Xinhua; denominator not disclosed ^[2]
New customer-service AI job categories - Role creation	Multiple	Categories/count not enumerated ^[2]

Source: [7] JD.com, n.d.; [16] JD.com, n.d.; [17] JD Corporate Blog, 2024; [5] JD.com, 2026; [2] Xinhua / Science and Technology Daily, 2026

MEASURED EVIDENCE

Measured Evidence: Training and Domain Context

FIGURE 2.2

reported result by analytical level

Training, specialization and future capability targets

Measure and analytical level	Reported result	Evidence boundary
Employee training coverage - Organizational capability context	100% in 2024	Company-wide; not AI-transition-specific ^[12]
Avg employee training - Organizational capability context	41.0 hours in 2024	Not AI-specific ^[12]
Total training hours - Organizational capability context	23.39m in 2024	Not AI-specific ^[12]
Maternal/infant service certification - Domain specialization	half in 2024; reported 100% by May 2026	Specific team; no team-size disclosure ^{[1] [7]}
Robot maintenance talent plan - Future capability target	10,000 over 3 years	Plan, not completed outcome ^[1]

Source: [12] JD.com, n.d.; [1] Science and Technology Daily, 2026; [7] JD.com, n.d.

Note: Company-wide and team-specific context measures do not establish an AI-transition outcome.

WORKFORCE RECONCILIATION

20. The Missing Workforce Reconciliation

FIGURE 2.3

not publicly reported in sources reviewed

The missing worker-flow accounting

The ideal evidence table for this transition would look something like:

Worker flow	Number
Customer-service employees before transition	? [2] [7] [16]
Remained in human service	? [2] [7] [16]
Shifted to complex-service specialist work	? [2] [7] [16]
Shifted to AI trainer / AI operations	? [2] [7] [16]
Shifted to robot / adjacent technology work	? [2] [7] [16]
Shifted to other JD functions	? [2] [7] [16]
Left voluntarily	? [2] [7] [16]
Positions not backfilled	? [2] [7] [16]
Involuntary exits	? [2] [7] [16]
New external hires into AI-support roles	? [2] [7] [16]

Source: [2] Xinhua / Science and Technology Daily, 2026; [7] JD.com, n.d.; [16] JD.com, n.d.

Note: Question marks in the source table mark values not disclosed in the materials reviewed.

WORKFORCE EVIDENCE

Reading the Workforce Evidence

01

Why the distinctions matter

A company can handle billions of AI inquiries and still employ thousands of humans.

A company can create new AI roles and still need fewer workers in an old role.

A company can train thousands and still fail to provide equivalent jobs.

A company can report “job upgrading” without disclosing pay, grade or long-term retention.

The strongest research product therefore does not ask for one number.

It constructs a measurement stack.

02

What the public record does not establish

JD has not published this reconciliation in the sources reviewed.

That means several tempting conclusions remain unavailable.

We cannot say:

“AI reduced JD customer-service headcount by X%.”

We also cannot say:

“All affected customer-service employees were redeployed.”

And we cannot calculate:

“transition success rate = thousands / 20,000+”

because the numerator and denominator are not defined as the same population.

Workforce transition evidence should eventually move from stories and totals to flow accounting.

The next generation of corporate AI disclosure should report:

where people came from → where they went → how many completed the move → what happened to pay and progression.

JD's evidence is strong enough to show the mechanism.

It is not yet strong enough to quantify the full distribution.

CHAPTER 3

Policy, Governance and Durable Work

The final chapter examines policy context, AI governance, the durability of new roles and choices for workers and leaders.

POLICY & IMPLICATIONS

01

21. Policy Context: The JD Case Sits Inside a New Chinese Employment Architecture

The JD transition is especially relevant in China because Artificial Intelligence Trainer is not only a corporate role.

It is a nationally recognized occupation with an official occupational skill standard.[9]

China's broader employment policy is also increasingly connecting AI adoption to worker transition.

In 2025, a national large-scale vocational skills action called for stronger enterprise in-role and job-transition training.[14]

In 2026, the national Stable Employment, Capacity Expansion and Quality Improvement Action Plan explicitly encouraged companies applying AI to conduct transition and redeployment training at the same time, with the goal of helping workers remain employed in enterprises where possible.[15]

The September 2026 national discussion of emerging AI occupations similarly framed occupations such as AI trainer as part of the response to structural labor-market change.[2]

02

Why this context matters

JD's case provides enterprise-level evidence of a policy problem that national institutions are now trying to solve:

“How can AI adoption increase productivity without making technological change and worker transition two separate processes?”

The company says it has intentionally expanded AI-related job opportunities in customer service and helped workers move toward higher-value positions.[2]

That aligns directionally with the policy agenda.

03

Causal boundary

The public evidence does not establish that JD's transition program was created because of a particular government policy.

The company had been building intelligent customer service for more than a decade, well before the latest policies.

Policy and company practice should therefore be described as converging, not as a proven one-way causal chain.

POLICY & IMPLICATIONS

01

22. External Institutionalization: Why Occupational Standards Matter

New AI jobs often begin with ambiguous titles.

One company may call someone:

- AI trainer; • bot trainer; • knowledge operator;
- conversation designer; • AI quality analyst.

Without a common skill architecture, workers may struggle to understand:

- what counts as competence; • which skills are portable; • how the role progresses;
- whether training has value outside one company.

China's AI-trainer occupational standard addresses part of this problem by defining levels and skill functions.[9]

JD's merchant-side training and certification ecosystem adds another layer of role formalization. [10]

02

AIGW interpretation

A durable workforce transition needs more than a new internal title.

It benefits from:

role definition + skill standard + learning pathway + real job demand + progression route.

This is one reason institutional infrastructure matters in AI labor transitions.

If the role remains company-specific and temporary, worker mobility may be weak.

If the skill can be recognized across employers and sectors, transition investment becomes more portable.

03

23. Comparative Insight: JD, Salesforce and LONGi Reveal the Same Underlying Mechanism

JD adds an important pattern to the AIGW Case Library.

POLICY & IMPLICATIONS

01

Salesforce

Routine customer-support interactions move toward Agentforce.

Some human expertise moves into:

- Success; • Renewals; • Forward Deployed Engineering;[18][19]
- other growth roles.

The organization also reports lower support staffing requirements.[18][19]

02

LONGi

Routine visual inspection moves toward machine vision and AI.

A documented quality inspector becomes an AI trainer, while humans remain in verification.[20]

03

JD

Basic service moves toward AI.

A documented experienced customer-service worker becomes an AI trainer and helps configure/tune the automated system.[1]

04

JD - Common mechanism

Across all three:[1][18][20]

AI takes a repeatable task → human domain expertise remains useful → some people move toward exception handling, training, verification, configuration or higher-context work.

05

JD - Important difference

The workforce outcomes are not identical.

Salesforce has direct evidence of lower support staffing requirements.[18][19]

LONGi has direct evidence of lower labor input in inspection, but incomplete worker-level reconciliation.

JD has direct evidence of thousands of job transformations/upgrades, but no public current headcount decomposition.[2]

This is why the same task-transition mechanism can produce different observable workforce responses.

POLICY & IMPLICATIONS

01

Layer 1 - Task substitution

AI handles more routine/basic customer inquiries.

Measure: AI interaction volume, automation rate, autonomous resolution.

JD provides interaction scale but not a clean current autonomous-resolution rate.

02

Layer 2 - Workflow redistribution

Complexity determines whether work stays with AI or moves to humans.

Measure: handoff rate, complexity, repeat contact, escalation reasons.

JD documents the architecture but not current detailed rates.

03

Layer 3 - Role decomposition

The old customer-service role splits into:

- automated basic service; • complex human service;
- domain-specialist service;
- AI operations; • adjacent technical support.

Measure: role definitions, staffing by role, new job categories.

JD reports multiple new categories but does not enumerate them publicly.

04

Layer 4 - Capability transition

Workers add:

- AI fluency; • configuration; • diagnosis;
- system tuning; • technical skills.

Measure: curriculum, certification, training completion, observed workplace capability.

Yuchen provides observed capability evidence; population-level curriculum data remain limited.

05

Layer 5 - Workforce movement

People move into new or upgraded roles.

Measure: source role, destination role, completion, retention, pay, exits.

JD reports thousands of transformations/upgrades but not full flow accounting.

POLICY & IMPLICATIONS

01

**Layer 5 - Workforce movement -
The central insight**

An enterprise can be highly advanced at Layer 1 while still having weak disclosure at Layer 5.

That distinction should become a standard part of AI workforce analysis.

03

**2. Build more than one transition
pathway**

Not every service worker should become an AI trainer.

Some may be better suited to:

- complex service; • domain specialization; • AI operations;
- technical support; • adjacent roles.

05

**4. Separate process scale from
employment impact**

Billions of AI inquiries show operational adoption. They do not show how many jobs disappeared.

02

**1. Automate the task, map the
expertise**

Before automating a frontline task, identify what expertise is embedded inside it.

The customer-service worker may know far more than the script.

04

**3. Treat domain knowledge as
transition capital**

Yuchen's service experience is useful because it helps him identify when the AI is wrong and why.

Generic AI courses cannot manufacture ten years of domain context.

06

**5. Separate role creation from net
job creation**

A company can create AI-trainer jobs while reducing other roles.

Both can happen at the same time.

POLICY & IMPLICATIONS

01

6. Measure workforce flows

Report:

- source roles; • destination roles; • number transitioned;
- number retained; • number exited; • pay and grade changes.

“Thousands upgraded” is valuable but incomplete.

03

8. Expect the destination role to change

AI trainers themselves face automation.

Reskilling should build higher-order capability, not just move workers into the next routine task.

05

10. Evaluate transition quality, not only transition quantity

A successful transition should eventually answer:

- Is the new role more secure? • Is pay maintained or improved? • Is progression credible?
- Are skills transferable? • Does the worker retain agency?

The public JD evidence does not yet answer these questions.

02

7. Design human escalation explicitly

“Human in the loop” should specify:

- what triggers handoff; • who decides; • what humans can override;
- who is accountable.

04

9. Make skills portable

National standards and recognizable certification can make new AI occupations more durable for workers.

POLICY & IMPLICATIONS

01

26. What the Evidence Establishes

The reviewed public evidence establishes that:

- JD has operated automated customer service since at least 2012-2013. • The customer-service architecture has long included human handoff for unresolved or complex cases. • By 2020, JD publicly described customer service as an integrated “human + intelligent” operating model.
- JD's service technology includes automation, agent assistance, intelligent training, quality inspection, knowledge management and operational decision support. • JD had more than 20,000 professional customer-service representatives at the end of 2024.
- During the 2025 11.11 promotion, JD's AI customer-service system handled more than 4.2 billion customer inquiries.
- More than 50,000 AI agents had been integrated into JD internal systems across operations by the end of Q4 2025. • Yuchen, an employee with nearly ten years of customer-service experience, transitioned into an AI-trainer role. • Yuchen uses domain experience to diagnose AI-response gaps and performs content updates, agent configuration and maintenance/tuning.
- Some customer-service colleagues moved into work related to remote robot operation and data collection.
- JD reported creating multiple AI-related job categories around customer service. • A JD representative reported that thousands of people completed job transformation or upgrading over roughly two years.
- JD describes the emerging service model as AI handling basic service, humans focusing on complex scenarios and multiple supporting roles enabling the system. • JD's current JoyAI product explicitly automates parts of knowledge-base analysis and FAQ mining to reduce AI-trainer configuration effort. • JD operates structured employee training and internal career-development infrastructure at company level.

02

26. What the Evidence Establishes

- Artificial Intelligence Trainer is a nationally recognized occupation in China with an official occupational skill standard. • JD's merchant learning ecosystem includes beginner and advanced AI-trainer certification courses.

POLICY & IMPLICATIONS

01

27. What the Evidence Does Not Establish

The reviewed evidence does not establish:

- the current total number of JD customer-service employees; • a before-and-after customer-service headcount attributable to AI; • the number of customer-service jobs eliminated by AI;
- the number of positions not backfilled because of AI;
- the number of involuntary exits linked to customer-service automation; • that all workers affected by AI were offered a transition pathway;
- that all workers offered a pathway completed it; • the exact number represented by “thousands”; • the precise denominator for the reported job transformations/upgrades;
- how many transitioned workers came from customer service versus other functions; • how many became AI trainers; • how many entered robot remote-operation roles;
- the distribution of workers across destination roles;
- whether transition was voluntary in every case; • training duration for the customer-service-to-AI-trainer pathway;
- whether training occurred during paid work time; • whether Yuchen holds a national AI-trainer skill certificate; • whether JD's merchant AI-trainer certification is used for internal employee transitions;
- whether pay, grade, benefits or career prospects improved after transition; • long-term retention in destination roles; • a current autonomous-resolution rate for JD customer-service AI;
- a current human-escalation rate; • formal human approval, veto or override rights across all AI service decisions; • that 4.2 billion AI-handled inquiries correspond to 4.2 billion human tasks replaced;
- that newly created AI roles produced net employment growth; • that JD's company-wide training statistics caused the reported AI job transitions.

POLICY & IMPLICATIONS

01

28. Research Gaps to Watch

1. What is JD's current customer-service headcount?
2. How has that headcount changed since the end of 2024?
3. How many employees were in the population eligible for AI-related transition?
4. What exact number is represented by “thousands”?
5. How many transitioned specifically from customer service?
6. What are the names and definitions of the multiple new job categories?
7. How many workers entered each destination category?
8. How many became AI trainers?
9. What proportion remained in complex human service?
10. What proportion entered domain-specialist service?
11. What proportion moved to robot remote operations or other technical roles?
12. What is the training curriculum for internal AI trainers?
13. How long does transition training take?
14. Is training conducted during paid work time?
15. Are employees selected, nominated or allowed to self-apply?
16. Can employees decline a proposed transition?
17. What proportion completes the training?
18. What proportion remains in the destination job after 6, 12 and 24 months?
19. How do pay and grade change?
20. What are the promotion paths from AI trainer?
21. Does JD map internal AI-trainer roles to the national occupational skill standard?
22. What share of AI-trainer work is already automated by JoyAI?
23. Which AI-trainer tasks are expected to remain human over the next three years?
24. What is the current AI autonomous-resolution rate?

POLICY & IMPLICATIONS

01

28. Research Gaps to Watch

25. What percentage of interactions escalate to humans? 26. What types of cases are prohibited from autonomous handling? 27. Can human representatives override AI decisions? 28. Are AI errors recorded and fed back into training? 29. How does JD measure model quality and customer harm? 30. How has entry-level customer-service hiring changed? 31. Has AI reduced replacement hiring or allowed service volume to grow without equivalent hiring? 32. How many new AI-related roles are filled internally versus externally? 33. What happens to workers who cannot or do not want to move into AI-related work? 34. Can independent research validate JD's reported workforce-transition scale? 35. How portable are the new skills outside JD?

03

A. Worker-level reporting

Science and Technology Daily's May 2026 reporting provides the strongest detailed evidence about Yuchen's work.

It establishes:

- prior customer-service experience; • transition to AI trainer; • actual tasks in the destination role;
- adjacent worker pathways.

The article contains direct worker statements and JD examples, but it is not an independent longitudinal workforce evaluation.

02

29. Method and Source Quality

This Deep Dive separates five evidence types.

04

B. Company-attributed workforce scale

The September 2026 article, republished by Xinhua and the Digital China Summit platform, contains a direct statement attributed to a JD representative that thousands of people had completed job transformation or upgrading over roughly two years. [2]

This is the strongest scale evidence.

It remains company-reported.

The source does not publish underlying administrative data, methodology or worker-level outcomes.

POLICY & IMPLICATIONS

01

C. First-party technology evidence

JD's JIMI, Yanxi and JoyAI materials provide strong evidence of:

- technology evolution; • human handoff; • system features;
- automated knowledge work; • second-order automation of trainer tasks.

These sources are vendor/company materials and should not be treated as independent effectiveness evaluations.

02

D. Formal corporate reporting

JD's ESG reports, annual results and SEC filings provide stronger evidence for:

- customer-service workforce scale; • AI deployment scale; • training infrastructure;
- internal mobility infrastructure; • enterprise governance.

These sources improve the case because they prevent a single employee profile from carrying the full argument.

They still do not provide a complete AI workforce reconciliation.

03

E. Government occupational and policy evidence

The national Artificial Intelligence Trainer occupational standard establishes that AI training is a formally recognized occupation with a structured skill framework.[9]

National skills and employment policies show that worker transition alongside AI adoption is becoming an institutional policy concern.[14][15]

These sources provide context.

They do not validate JD's company-specific outcome claims.

04

30. The Most Important Interpretation: Service Becomes a Two-Sided Human-AI System

The JD case is often easiest to understand if customer service is divided into two sides.

05

Customer-facing side

AI absorbs more basic service.

Humans become relatively more concentrated in:

- ambiguity; • emotion; • trust;
- domain knowledge; • operational exceptions.

POLICY & IMPLICATIONS

01

System-facing side

Humans increasingly work on:

- knowledge quality; • AI configuration; • failure diagnosis;
- evaluation; • maintenance; • tuning;
- workflow improvement.

The same person may move from one side to the other.

Yuchen is a documented example.

02

System-facing side - Why this matters

The old workforce model asks:

“How many agents do we need to answer demand?”

The new model asks two questions:

“How much demand can AI handle?”

and:

“How many people, with what skills, are needed to operate the AI-enabled service system?”

That second question creates a different workforce architecture.

03

31. The Durability Test for New AI Jobs

JD also suggests a useful test for whether a new AI occupation is likely to be durable.

Ask four questions.

04

1. Is the role based mainly on repetitive configuration?

If yes, the role itself may be highly automatable.

JoyAI's automated FAQ mining is an example.

05

2. Does the role require domain judgment?

If yes, it may remain valuable longer.

Yuchen's ability to distinguish configuration failure from logic failure is a stronger capability than simple content entry.

POLICY & IMPLICATIONS

01

3. Does the role control consequential workflow behavior?

Roles that design routing, escalation, quality and system behavior may become more important as AI scale grows.

03

4. Is the skill portable? - AIGW takeaway

Do not measure reskilling success by whether a worker enters an AI-labeled job.

Measure whether the worker acquires capabilities that remain valuable as AI advances.

05

For customer-experience leaders

Redesign the queue before redesigning the workforce.

Define:

- which interactions AI should own; • which require humans; • what triggers escalation;
- which expertise is needed in complex service; • what roles maintain the automated layer.

Workforce demand follows the workflow design.

02

4. Is the skill portable?

Occupational standards, recognizable competencies and cross-industry applicability can increase worker resilience.

04

For CHROs

Build a transition portfolio, not a single reskilling course.

Map workers into several destination archetypes:

- complex human service; • domain expert; • AI trainer;
- AI evaluator; • AI operations specialist; • adjacent technical roles.

Track actual movement, not training enrollment alone.

POLICY & IMPLICATIONS

01

For CIOs / AI leaders

Treat frontline employees as system-design inputs. The workers closest to customer failure modes can provide:

- edge cases; • error taxonomies; • business rules;
- evaluation criteria; • escalation logic.

Do not separate model development from operating knowledge.

02

For L&D leaders

Teach system reasoning, not only tool operation.

A stronger curriculum for AI service operations would include:

- domain knowledge; • error analysis; • knowledge management;
- agent configuration; • evaluation; • prompt/instruction design where relevant;
- escalation design; • data/privacy rules; • service-risk judgment.

03

For policymakers

Measure both displacement and transition quality.

Useful employer disclosure would include:

- positions automated; • workers affected; • workers trained;
- workers successfully redeployed; • new roles created; • pay changes;
- non-backfill; • exits.

New occupational standards are useful, but employment quality still needs evidence.

04

For workers

Do not abandon domain expertise in order to “learn AI.”

The most defensible position may be:

“know the business deeply enough to tell the AI when it is wrong, and know the system well enough to fix what happens next.”

That is a stronger transition asset than generic familiarity with an AI interface.

CONCLUSION

Conclusion

JD.com's customer-service transition shows what happens when automation becomes infrastructure rather than a pilot. For more than a decade, the company has moved routine service into increasingly capable automated systems while maintaining human escalation and support. By 2025, the AI layer was handling billions of inquiries. By 2026, the workforce consequences were becoming more visible. Experienced customer-service knowledge was no longer valuable only at the point of direct customer interaction. It could be redeployed into:

- AI training;
- agent configuration;
- system tuning;
- complex service;
- domain specialization;
- adjacent technical work.

Yuchen's transition makes that change concrete. But the most important lesson is not that “customer-service workers can become AI trainers.” That is only one stage. JD's own JoyAI product is already automating parts of AI-trainer knowledge configuration. The deeper lesson is that the human

role keeps moving upward through the system. From: answering to handling exceptions to training to diagnosing to configuring to evaluating and governing.

That makes workforce transition a continuous process rather than a one-time reskilling event. JD also exposes the limits of current corporate AI disclosure. The company reports thousands of workers transformed or upgraded, but does not publish the complete workforce flow behind that number. We know that new roles exist and real people have moved into them. We do not know how those successes compare with the full population affected by automation. That distinction matters. A credible AI workforce strategy must eventually answer both:

“What new work did AI create?” and “What happened to everyone whose old work changed?” JD provides unusually strong evidence for the first question. The second remains the research frontier.

SOURCES

Notes and References

A parallel government-hosted republication of reference [2] is available at https://www.digitalchina.gov.cn/2026/szzg/xyzx/202609/t20260907_5368538.htm.

The promulgation notice for the occupational standard in reference [9] is available at https://www.mohrss.gov.cn/xxgk2020/fdzdgknr/rcrs_4225/jnrc/202112/t20211227_431404.html.

1. Science and Technology Daily. 凝“新”聚力促就业 | AI为他们带来新身份--京东协同推进技术应用与就业转型探索实践. 19 May 2026. Attributable national science/technology journalism. https://www.stdaily.com/web/gdxw/2026-05/19/content_518575.html
2. Xinhua / Science and Technology Daily. 未来五年，如何接住AI“新饭碗”. 7 September 2026. National media republication of Science and Technology Daily reporting. <https://www.news.cn/tech/20260907/eb1d91d961c84dd8b417ba41d649b129/c.html>
3. JD Consumer Help Center. *JIMI intelligent customer-service robot*. n.d.. JD first-party product/help documentation. <https://help.jd.com/o/help/question-1024.html>
4. JD Cloud Developer Community. 京东智能客服品牌焕新：“言犀”亮相2020京东JDD大会. 26 November 2020. JD first-party technology account. <https://developer.jdcloud.com/article/1317>
5. JD.com. *Fourth Quarter and Full Year 2025 Results*. March 2026. Company investor reporting. <https://ir.jd.com/news-releases/news-release-details/jdcom-announces-fourth-quarter-and-full-year-2025-results-and>
6. JD JoyAI. *Intelligent Customer Service*. n.d.. JD first-party current product page. <https://jdai.jd.com/portal/ai-products/customer-service>
7. JD.com. *2024 Environmental, Social and Governance Report*. n.d.. Company ESG report. https://ir.jd.com/system/files-encrypted/nasdaq_kms/assets/2025/06/19/15-13-13/2024_jd.com_inc._environmental_social_and_governance_report.pdf
8. JD.com. *2025 Form 20-F*. 16 April 2026. Regulatory filing. <https://www.sec.gov/Archives/edgar/data/1549802/000119312526157870/d53690d20f.htm>
9. Ministry of Human Resources and Social Security / Ministry of Industry and Information Technology. *Artificial Intelligence Trainer National Occupational Skill Standard (2021)*. n.d.. Official national occupational standard. <https://www.mohrss.gov.cn/wap/zc/zqyj/202106/W020210617509883457681.pdf>
10. JD Merchant Learning Center. *Jingxiaozi AI Trainer Certification Series*. n.d.. JD first-party merchant learning platform. <https://xue.jd.com/ondemandCourse/searchOndemandCoursePage.action?classifyId=159>

SOURCES

11. JD.com. *2024 ESG Report, Employee Training and Development / Talent Development*. n.d.. Company ESG report.
https://ir.jd.com/system/files-encrypted/nasdaq_kms/assets/2025/06/19/15-19-22/2024_jingdongjituanhuanjingshehuijizhilibaogao_.pdf
12. JD.com. *2024 ESG Report, Training Performance*. n.d.. Company ESG report.
https://ir.jd.com/system/files-encrypted/nasdaq_kms/assets/2025/06/19/15-13-13/2024_jd.com_inc_environmental_social_and_governance_report.pdf
13. JD.com. *2024 ESG Report, Emerging AI Risks and Responses*. n.d.. Company ESG report.
https://ir.jd.com/system/files-encrypted/nasdaq_kms/assets/2025/06/19/15-13-13/2024_jd.com_inc_environmental_social_and_governance_report.pdf
14. State Council Employment Promotion and Labor Protection Leading Group. *Guidance on Large-Scale Vocational Skills Upgrading*. 2025. National employment / skills policy.
https://www.gov.cn/zhengce/zhengceku/202507/content_7032222.htm
15. State Council Employment Promotion and Labor Protection Leading Group. *Stable Employment, Capacity Expansion and Quality Improvement Action Plan*. April 2026. National employment policy.
<https://chinajob.mohrss.gov.cn/h5/c/2026-05-20/544527.shtml>
16. JD.com. *2024 Form 20-F*. n.d.. Regulatory filing.
<https://www.sec.gov/Archives/edgar/data/1549802/000119312525083473/d871796d20f.htm>
17. JD Corporate Blog. *JD.com Announces Salary Increase for Over 20,000 Customer Service Employees*. 6 February 2024. Company workforce announcement.
<https://jdcorporateblog.com/jd-com-announces-salary-increase-for-over-20000-customer-service-employees/>
18. Salesforce. *How Salesforce Is Reshaping Its Workforce in the Age of AI*. 29 April 2026. Company / first-party.
<https://www.salesforce.com/news/stories/salesforce-reshaping-workforce-in-age-of-ai/>
19. Salesforce. *Today's Hottest Role: Forward Deployed Engineer*. n.d.. Company / role description and examples.
<https://www.salesforce.com/ap/blog/forward-deployed-engineer/>
20. LONGi. *Sustainability Report 2023*. 2023. Company sustainability report.
https://static.longi.com/2023_longi-sustainability-report.pdf