Alibaba

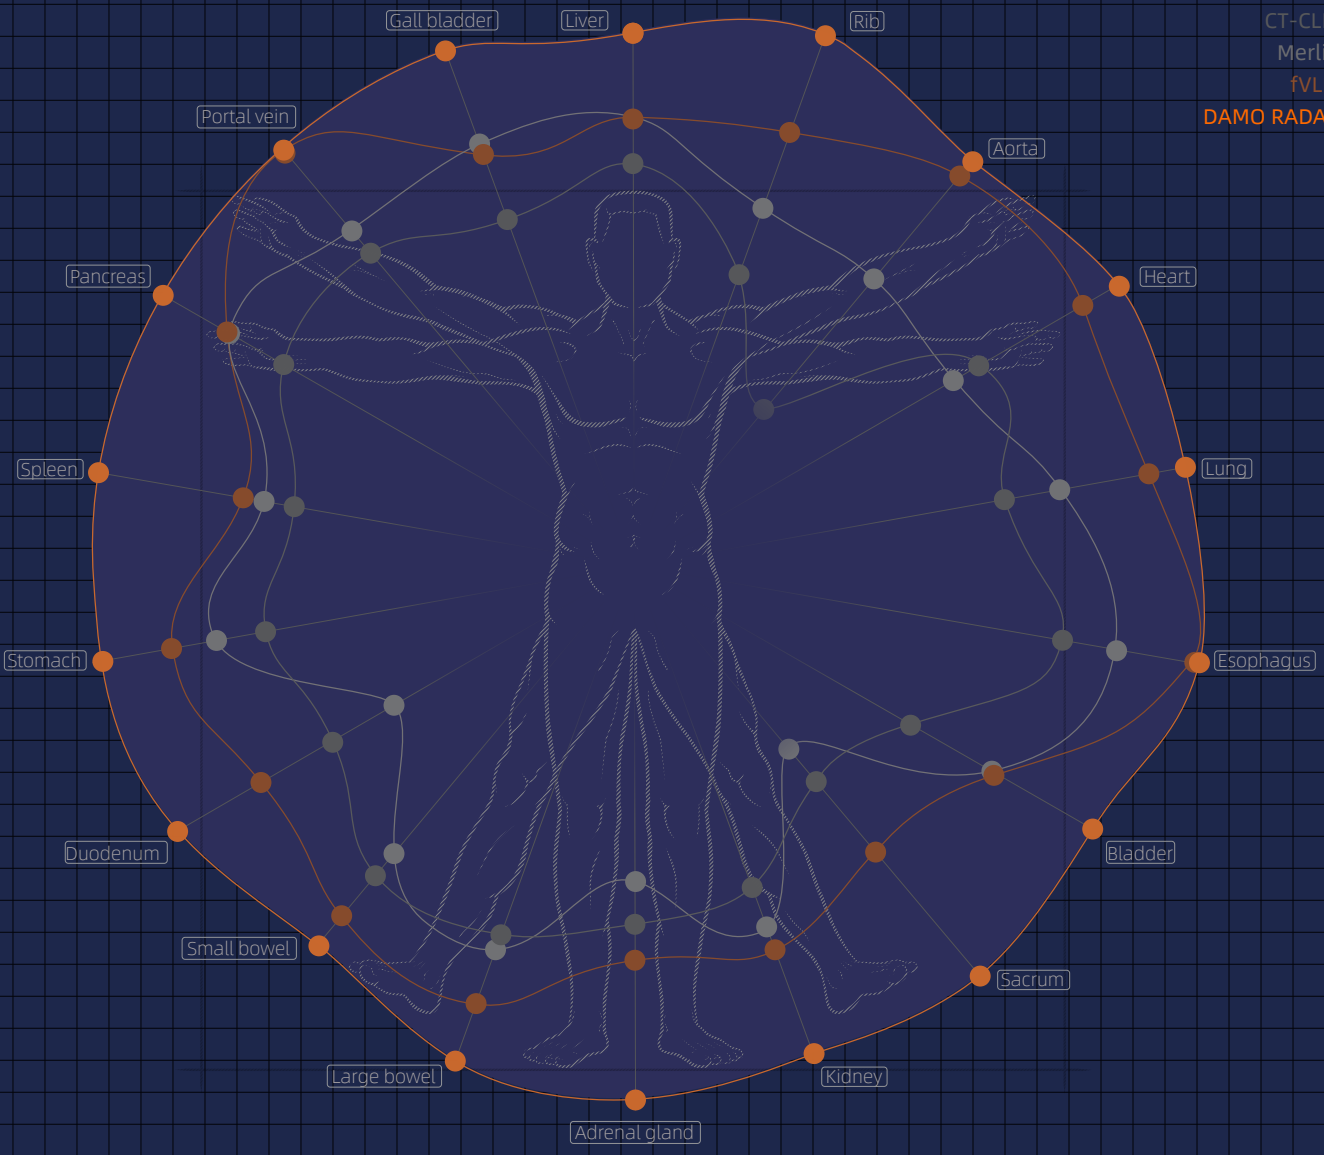

Alibaba
AI for Good
Action Report 2026

Cover Caption:
Alibaba DAMO Academy has open-sourced DAMO RADAR, an expert-level general-purpose medical imaging AI model covering 18 organs and 146 conditions. The findings are published in Science. The cover image is from the paper.

Contents

AI for Good: At the Dawn of the Agentic Era 2

AI for Good: A Shared Journey 4

Chapter 1 · The Co-evolution of AI and Society 6

AI's Rapid Integration into Society 7

How Technology and Society Adapt to Each Other 7

What Sustainable AI Development Requires 8

Chapter 2 · Full-Stack AI: Working Together for Good 9

Defense in Depth for AI Safety 11

Open Source and Openness: Broadening Participation 14

Green Infrastructure and Compute-Power Coordination 16

Chapter 3 · AI for Good in Action 2026 19

AI for Health 22

AI for Education 29

AI for Science 35

AI for Green Development 39

AI for Economic Opportunites 44

AI for Safe and Inclusive Communities 49

Chapter 4 · Looking Ahead 56

How Society Shapes AI's Next Stage of Evolution 57

Seeking Truth, Pursuing Good: Eight Questions for AI's Future 58

AI for Good: At the Dawn of the Agentic Era



Over the past year, AI has entered a new stage of development, moving from answering questions to executing tasks. We now stand at the dawn of the agentic era. AI agents can now autonomously plan, use tools, coordinate across multiple subsystems, and make decisions. They are entering real-world scenarios and gradually becoming systems capable of taking actions. Together, we have witnessed a transformation with significance extending far beyond technology itself.

The development of AI is a process in which stakeholders continuously work together to calibrate its direction. Throughout history, technological development has presented humanity with similar challenges. In the early stages of a technology, its social impact is difficult to predict. Yet once the technology matures and becomes deeply embedded in society, changing its course can be enormously costly. We cannot wait until every question has been addressed before taking action. Instead, we must act with a commitment to AI for Good amid uncertainty, continuously learning and adjusting to ensure that technological development remains aligned with the well-being of humanity and social value creation.

The challenges, of course, are real and concrete. I believe the path forward for AI must meet three conditions: it must be socially acceptable, environmentally sustainable, and economically viable.

Social acceptability means that as AI agents enter real-world environments, their services must be reliable and trustworthy, while delivering tangible benefits to people. In this year's report, we see how Qwen has introduced an AI agent to help students in China in choosing universities and academic programs after completing the Gaokao, China's national college entrance examination. Accio, the AI-powered business agent on Alibaba's international B2B trading platform, is helping 10 million small and medium-sized enterprises engage in cross-border trade. Alibaba Health's Hydrion has begun assisting physicians with clinical decision-making. As AI agents begin helping people make decisions, trust in AI is no longer an abstract concept. It becomes something traceable in every interaction, and progressively built through continuous service improvement.

Environmental sustainability has become a major concern, with AI's consumption of energy and water attracting particular public attention over the past year. In fiscal year 2026, clean electricity accounted for 73.6% of the electricity used by Alibaba's self-built data centers, reaching 100% in Guangdong, Jiangsu, and Sichuan. However, with demand for computing power still growing rapidly, improvements in per-unit efficiency do not necessarily translate into lower absolute consumption. How to expand AI capabilities while controlling carbon emissions and reducing impacts on

water resources remains a question the entire industry must address together.

Economic viability concerns whether AI can create real value commensurate with the resources invested. Capital investment in AI infrastructure has reached unprecedented levels. We expect and believe that the value of these investments will extend far beyond improving efficiency across industries. AI will continue to address important social challenges, benefit a broader range of people, and bring about more lasting changes in society.

As a leading global AI company, Alibaba has long attached great importance to AI safety and ethical governance. At the end of 2022, we established our Committee on Governance of Ethics in Science and Technology and subsequently published the Code of Governance of Ethics in Science and Technology of Alibaba Group. As AI evolves from conversational interfaces into autonomous agents, the implication of safety is expanding accordingly. We must pay attention not only to what models say, but also to what AI agents do. To this end, Alibaba has established a six-layer defense-in-depth safety framework covering the entire model lifecycle.

At the same time, we remain firmly committed to open source, enabling the benefits of AI to reach a much broader developer ecosystem. As of September 2026, Alibaba had open-sourced more than 460 models in the Qwen family, giving rise to over 300,000 derivative models and surpassing 3 billion cumulative downloads worldwide. These models support more than 200 languages and dialects. On Hugging Face, developers around the world have created tens of thousands of locally deployable variants of Qwen models. ModelScope, the open-source community created and operated by Alibaba, has brought together 33 million developers and more than 250,000 open-source models. We believe that an open, community-driven ecosystem offers greater resilience and creativity than a closed one.

The ultimate test of technology is whether it can help solve real social problems. The complexity and diversity of each AI for Good use case are valuable resources for training more capable AI. For example, building on our AI technology for pancreatic cancer detection, DAMO Academy has developed DAMO RADAR, a general-purpose medical imaging AI model that now covers 18 organs and 146 diseases and medical conditions. This demonstrates the enormous potential of AI to continuously improve, adapt, and create greater social value through real-world applications.

Emerging technologies inevitably bring new challenges, and AI is no exception. Some questions still have no well-established answers anywhere in the world. This is both a challenge and an opportunity. It means that we are still in the formative period of the rules governing this technological evolution. Through our thoughts and actions, we have the opportunity to explore a path for AI development that advances well-beings for humanity.

This exploration is embedded in our everyday work and reflected throughout this report: what we have accomplished, what we are working on, and what challenges remain. The report represents our progress to date, as well as an open and candid dialogue with our peers and partners.

The direction is clear. The path must be taken one step at a time, and every step forward calls for more people to join.

Wu Zeming
Alibaba Partner
Alibaba Group Chief Technology Officer
and Chairman, Committee on Governance of Ethics in Science and Technology

AI for Good: A Shared Journey



If last year AI was a powerful tool in the lab, this year it is becoming a practical tool in the hands of many, serving industries and communities across society. AI can now not only converse but also get things done. It is increasingly finding its way into hospitals, schools, communities, and rural areas, becoming part of our everyday work and lives.

When assessing the progress of AI, advances in technical performance within labs certainly matter. But what matters more is whether AI can reliably serve people and deliver tangible benefits to their lives. To truly realize AI for Good, we need an inclusive social ecosystem in which everyone can access and benefit from AI.

Cutting-edge technology does not automatically find its way into ordinary people's lives. To make AI accessible and genuinely useful to everyone, we must create the conditions for technology to become integrated into society. This requires collective efforts across society, and the role of businesses is an important part of that endeavor. This year's report highlights many of Alibaba's efforts and insights in this regard.

First, we must prioritize AI applications based on social impact, exploring where technology can deliver the greatest benefits to the largest population. As we stated in last year's report, we should "direct technological capabilities toward key areas where they can generate the greatest systemic impact." This commitment remains unchanged. Given that AI represents the most resource-intensive technological innovation in human history, we should prioritize its applications in areas where it can make the greatest contribution to solving important social problems.

In healthcare, the industry faces a lot of imbalance between supply and demand, and the potential of AI is immense. DAMO Academy's early-screening technology, combining non-contrast CT scans with AI, is reaching an increasingly wide range of use cases. What began as screening for a single type of cancer has expanded to five cancers of the digestive system this year. Meanwhile, enhanced abdominal CT scans combined with AI can now identify 146 medical conditions and findings in a single examination. Such technology can not only save lives, but also help ease the pressure on public healthcare resources if widely adopted.

In scientific research, AI is also demonstrating enormous potential in frontier fields such as virus identification, materials discovery, and cellular research. A recursive cycle of improvements is beginning to take shape, in which AI accelerates scientific discovery, and scientific discovery in turn advances AI. We prioritize scientific research because every breakthrough has the potential to benefit humanity.

Second, we must create avenues for more people to use AI. The possibilities of the AI era are not simply about what AI can achieve. They are about what people can achieve, or more precisely, what people can achieve with AI.

Through our "Cloud for Youth" initiative, we promote access to AI education because we want children in rural areas to have greater opportunities to learn, grow, and realize their potential in the AI era. Through Hydrion, our evidence-based medical AI model, we support clinical diagnosis because we want China's millions of doctors, especially those in remote regions, to have more opportunities for continued learning and improving their quality of care. Through the ModelScope community, we make open-source models more widely accessible because we want more developers to discover opportunities to address social challenges through affordable computing resources and in an open ecosystem.

Third, we must develop mechanisms that bring technological capabilities and real-world needs together, enabling AI to be deployed in an ever-wider range of scenarios.

This year, one new role has attracted considerable attention across the AI industry: the Forward Deployed Engineer (FDE). What we have observed on the front lines of our philanthropic work confirms the importance of this role. In rural revitalization, for example, bringing technology into villages is not enough. We must understand the real needs of local communities and help integrate technology into existing organizational processes. This year, in Shengsi, Zhejiang Province, Alibaba's Rural Revitalization Commissioner combined local knowledge of a complex ferry ticketing system, route regulations, and changing weather conditions with AI capabilities to create an intelligent ferry ticketing assistant, helping address a major transportation challenge for local residents. Similar initiatives are now being explored in more counties across China. This kind of collaboration is essential to bringing AI into an increasing number of real-world applications.

These on-the-ground experiences remind us that while AI has remarkable capabilities, turning those capabilities into meaningful outcomes depends on people, not technology alone. During one of my field visits, I heard a doctor in a county-level hospital make an observation that has stayed with me. AI-powered early screening can indeed detect cancer risks. But can doctors accurately verify the findings? Can hospitals successfully contact patients and bring them back for follow-up examinations? Will patients agree to further testing? And can they be referred to the right specialists afterward? AI alone is far from sufficient. Much of the work involved, both today and in the future, will continue to require human judgment and participation.

AI is powerful, but it is not omnipotent. Bridging the gap between advanced AI and real social challenges requires human efforts, as well as the institutions and culture that bring people together. If we want to enjoy the benefits of AI, we must work together to build a caring and resilient society.

Over the past year, I have seen Alibaba colleagues pushing the boundaries of technological innovation with conviction. I have also seen their relentless efforts to bring technology into real-world use cases. For example, the "AI for Starry Kids" initiative featured in this report grew out of the spontaneous efforts of Alibaba employees who are parents of children with autism. Alibaba engineers have also devoted their personal time to developing an AI-powered safety solution for older adults living alone in Hangzhou's Yuhang District. And Alibaba technical specialists have designed an AI-powered wildlife identification solution to support conservation efforts at the source of the Yangtze River.

We have often said that it takes people who care to create technology and products with compassion, and that such technology and products are what define a good company. In the age of AI, however far technology advances, this original commitment remains unchanged.

I hope this report will serve as a small signpost, marking the path we have traveled and inspiring us to continue moving forward.

Sun Lijun

President of Alibaba Philanthropy, Alibaba Group
Chair of the Board, Alibaba Foundation

Chapter 1
The Co-evolution of AI and Society



AI's Rapid Integration into Society

In 2026, AI reached a clear turning point in its development, beginning to move from technological breakthroughs toward large-scale adoption. Since the Industrial Revolution, machine power has gradually replaced much of human physical labor. Now, as machine intelligence begins to take on thinking and decision-making tasks at scale, technology's impact on society may only just be beginning in earnest.

A defining feature of this new stage is the changing nature of AI applications. Unlike earlier conversational interactions, AI agents can now plan tasks, use tools, and participate in business processes. They are moving AI applications beyond question-answering assistance toward the execution of complex tasks, gradually entering a wide range of industries and becoming part of people's everyday work and lives. Meanwhile, continued investment in chips, cloud infrastructure, models, applications, and other technologies is further expanding the supply of machine intelligence at scale, enabling more organizations and individuals to access useful AI tools and services.

As AI accelerates its integration into society, it also brings more immediate, real-world challenges. Improvements in model capabilities do not automatically lead to greater reliability, while intense industry competition places additional pressure on safety alignment and risk assessment. The expansion of data centers is increasing energy and water consumption, adding to environmental burdens. Concerns about job displacement and workforce transitions are also attracting greater attention.

The development of technological capabilities does not automatically translate into widespread social benefits. It also requires the right organizational, institutional, and market conditions. Infrastructure development and capital investment help build technological capabilities, but whether these capabilities can generate real social value also depends on how effectively organizations apply them, whether appropriate institutions are in place, and whether markets can support the long-term delivery of services. These two dimensions complement and influence each other. Sound institutional arrangements and governance mechanisms are essential to the reliable operation of technology over the long term.

In this process of transformation, market forces alone inevitably leave certain needs unmet. Some important social needs may be overlooked because the people they serve have limited ability to pay, service costs are high, or returns take a long time to materialize. Addressing this imbalance requires public policies, research support, and sustained corporate investment to preserve room for exploration in these areas. It also depends on technological progress and economies of scale to reduce costs, making viable solutions possible for social needs that were previously difficult to serve.

Therefore, understanding AI for Good in 2026 requires us to look beyond the benefits and risks that emerge after AI is deployed. We must also look upstream, to the frontiers of technology research and development, resource allocation, and application design, making the commitment to AI for Good an integral consideration from the very beginning of technological development.

How Technology and Society Adapt to Each Other

As AI moves from the development of technological capabilities to widespread adoption across society, it must undergo a process of adaptation and integration. Whether this process succeeds depends on two things: whether technology can respond to real needs, and whether society has the capacity to use, evaluate, and govern it.

This means that the direction of technological development is not determined by technology alone. Social needs, institutional choices, and public participation all influence its trajectory, shaping which problems innovation addresses first and who stands to benefit. In some settings, improving organizational processes and institutional arrangements can do more to improve real-world outcomes than simply raising a particular technical performance metric.

In this process, the key is not only to make AI more intelligent, but also to make it better able to function within society. When AI is deployed in specific settings, it must operate under clear rules and responsibilities. Questions of how identities are established, rights and responsibilities assigned, boundaries defined, and accountability enforced when problems arise all require appropriate institutional

arrangements. Without these arrangements, greater technological capabilities may also bring greater potential impacts and more systemic risks to society.

Today, major economies are refining their governance frameworks alongside industrial deployment, seeking to ensure that institutional development keeps pace with technological adoption. Yet AI evolves rapidly, and governance cannot be a one-time exercise in institutional design. It requires a capacity for continuous adaptation. This, in turn, calls for collaboration among multiple stakeholders: laws and regulations establish basic requirements; standards and implementation measures evolve in response to new evidence; businesses document problems encountered during deployment and operation; professional organizations provide independent assessments; and users and affected communities contribute feedback.

At the same time, individuals and organizations must gradually develop the capabilities needed to adapt to AI. Businesses may need to redesign workflows and job roles. Teachers need to use AI for students to support learning while preserving essential opportunities for independent thinking and practice. Physicians must maintain their professional judgment when using AI to assist with diagnosis and treatment. The earlier these adjustments begin, the lower the costs and difficulties are likely to be when adoption reaches a larger scale.

Technological and social systems must co-evolve if new technologies are to benefit more people. Only when AI can be successfully integrated into social systems, and those systems have the capacity to absorb, adapt to, and govern AI capabilities, can AI truly evolve from a frontier technology into a trusted, widely shared form of social infrastructure that creates public value sustainably.

What Sustainable AI Development Requires

Achieving the co-evolution of technological and social systems requires three foundational conditions: social acceptability, environmental sustainability, and economic viability. Together, they define the "safe operating space" within which AI can develop steadily and responsibly. Only when these conditions are in place can the full potential of AI for Good be realized.

From a social perspective, large-scale AI deployment depends on public acceptance and regulatory approval. Safety incidents, a lack of system transparency, and concerns about harmful or disruptive outcomes can all constrain this space. Whether society accepts AI depends on the reliability and trustworthiness of its services, as well as whether different groups genuinely benefit. Ensuring safety, improving transparency, promptly correcting problems, and providing remedies when harm occurs are fundamental to building trust.

Environmental sustainability involves several considerations, including energy consumption, water resources, and carbon emissions. Optimizing models, chips, and operational scheduling can improve resource efficiency, while the use of clean energy can help reduce emissions. However, these improvements may not be sufficient to offset the overall impact of continued growth in scale. Environmental constraints must therefore be incorporated into every stage of AI planning, construction, and operation from the outset, rather than addressed only after deployment has expanded.

From an economic perspective, investment in AI has reached unprecedented levels. Its long-term sustainability depends on whether AI can address sufficiently important problems and create real value commensurate with the resources invested. At the same time, mechanisms are needed to support continued operation, allowing services with limited short-term commercial returns to remain viable for further exploration and ultimately deliver their intended value.

These three dimensions influence and can reinforce one another, creating a virtuous cycle between technological development and social benefit. Maintaining and expanding this "safe operating space" requires collective efforts across society, with institutions that command greater technological capabilities and resources bearing responsibilities commensurate with their influence. AI for Good is not an additional social responsibility initiative separate from core business activities. It is an intrinsic requirement that runs through the entire process of technology research and development, infrastructure construction, and product and service delivery.



Chapter 2
Full-Stack AI: Working Together for Good

> As AI capabilities extend from "speaking" to "acting", safety defenses increasingly need to be organized according to a zero-trust architecture.

In last year's report, we introduced Alibaba's framework for putting AI for Good into practice. Guided by the principle of human-centricity, the framework set out six principles to guide technological development, emphasized empowering individuals and organizations through accessible infrastructure, and encouraged the exploration of solutions to social challenges in specific real-world settings. As AI moves from content generation to task execution, becoming more deeply embedded in organizational workflows and everyday life, coordination across technological capabilities, participation by developers and partners, and validation and feedback from real-world applications have become increasingly important.

This year, drawing on Alibaba's accumulated experience in technology, ecosystems, and business applications, we have further refined the framework, explicitly identifying full-stack AI capabilities, open source and community collaboration, and real-world applications and feedback as the foundations supporting our work. We have also expanded our areas of focus to six fields: health, education, scientific research, green development, economic opportunity, and safe and inclusive communities. These areas are closely connected to several United Nations Sustainable Development Goals. The starting point of our framework remains unchanged: technological capabilities are the means, while delivering tangible benefits to people is the goal.

This refinement of our framework also builds on Alibaba's experience in technology development and real-world applications over the past year. We have improved operational efficiency and reduced usage costs through co-designed optimization of models, chips, and inference systems. By opening models and tools and encouraging community collaboration, we have enabled more developers and partners to participate in innovation. Through real-world applications and feedback, we continue to improve our technologies and services. We have also strengthened safety measures across research and development, deployment, and operations, while exploring ways to improve resource and energy efficiency through resource scheduling and the use of clean electricity.

These efforts require trade-offs among different objectives. Full-stack coordination does not mean that performance, cost, safety, and environmental goals can always improve simultaneously, but it provides a foundation for addressing these requirements together. We remain focused on the real impact of technological development on people and continue to explore how the principles of AI for Good can be integrated into business practices. The following sections examine the foundational work that supports AI for Good applications, presenting our practices and progress over the past year in three areas: AI safety, green infrastructure, and open source and openness.

AI for Good: A Human-centered Action Framework

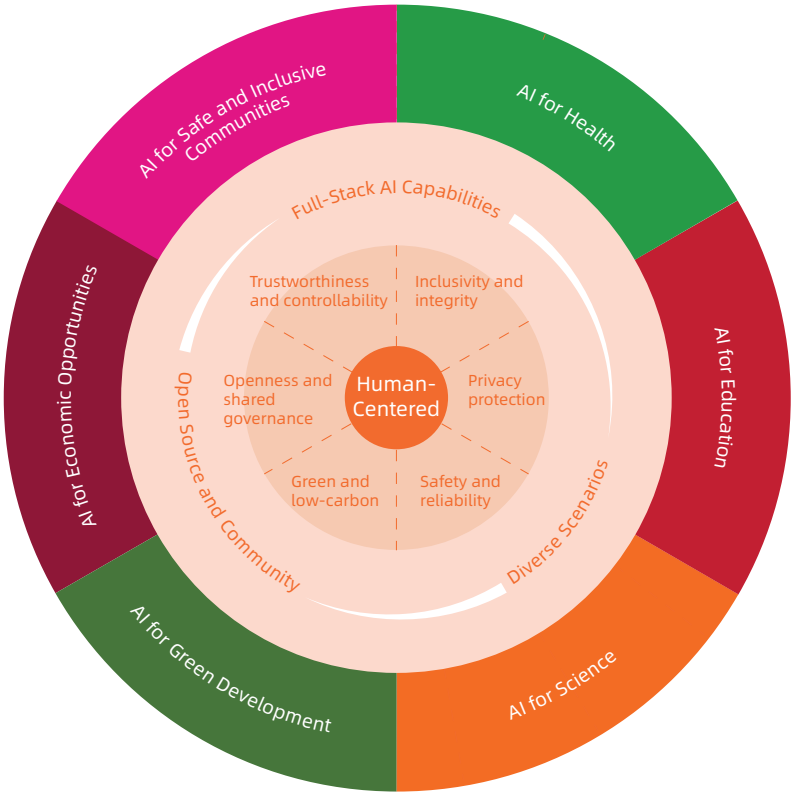

Defense in Depth for AI Safety

Safety and reliability are among Alibaba's six guiding principles for AI for Good. At the heart of our approach to model safety is the concept of "self-discipline". This does not mean conservatively restricting model capabilities. Rather, it means ensuring that models consistently respect the boundaries of safety, compliance, trustworthiness, and controllability, even as their capabilities expand rapidly.

As AI evolves from conversational interfaces into autonomous agents, the meaning of safety is expanding accordingly. We must pay attention not only to what models say, but also to what AI agents do. To this end, Alibaba has established a six-layer defense-in-depth safety framework covering the entire model lifecycle: evaluation-driven safety, intrinsic safety, external guardrails, digital watermarking, cloud-based security, and agent safety.

The first layer is evaluation-driven safety. Alibaba's evaluation framework is designed to measure models' actual safety performance. It covers

Progress in Governance and Disclosure

2021
Established the *Group's Sustainability Committee*; proposed *Cloud Computing Carbon Neutrality Strategy and Targets*.

2022
Established the *Committee on Governance of Ethics in Science and Technology*; released the *Artificial Intelligence Governance & Sustainable Development* Practice White Paper.

2023
Integrated ethics governance into the Group's sustainability work system; released the *AIGC Governance & Practice White Paper.*

2024
Released the *Code of Governance of Ethics in Science and Technology of Alibaba Group*, the *Management Standards for the Ethics Review of Science and Technology of Alibaba Group*, and *The Large Language Model Technology Development and Governance Practice Report*.

2025
Released the first *Alibaba AI for Good Action Report*.

2026
Jointly released the AI Safety White Paper *Self-Discipline and Social Benefit: Developing Responsible Technology in the Age of AI*.

foundation models, reasoning processes, and final outputs, as well as text, images, audio, video, cross-modal interactions, and real-world business scenarios. Alibaba has developed safety evaluation frameworks such as S-Eval, combining conventional evaluations, red teaming, adversarial testing, and neuron-level knowledge probes. These evaluations do more than identify problems. They also translate findings about specific risks into actionable inputs for improving model training, alignment, and guardrails.

The second layer is intrinsic safety: making safety an inherent capability of large AI models. Safety should not depend solely on blocking mechanisms in external guardrail systems. More importantly, safety capabilities should be built into the models themselves through model architecture, training data governance, alignment techniques, and task design, enabling models to better recognize and respond to risks. External guardrails are then used to block risks only where models have safety flaws or weaknesses, and in specific high-risk situations where errors could have particularly serious consequences.

Alibaba's published Model Specification translates abstract safety values into behavioral rules that models can learn, follow, be evaluated against, and continuously improve upon. It functions much like a "constitution" for model behavior. When user requests, developer policies, and fundamental principles come into conflict, the specification establishes the priority of higher-level safety principles while seeking to satisfy legitimate user needs as fully as possible within established safety boundaries.

The specification establishes an instruction hierarchy of Root Principles > System Rules > Developer Rules > User Preferences. This enables models to follow higher-priority safety requirements more reliably when instructions conflict, while preserving developer control over downstream customization and application development.

The third layer is external guardrails. Deployed around the model inference service pipeline, these mechanisms can detect risks and intervene intelligently before user inputs reach the model or before model outputs are delivered to users. Native guardrail technologies such as QwenGuard, PlugGuard, and EDS can strengthen safety defenses in complex interactions without requiring significant modifications to the underlying model, reducing unnecessary refusals while maintaining strong safety performance. Native streaming guardrails, in particular, can use intermediate hidden states during generation to identify potential risks early and block harmful content before it is delivered.

The fourth layer is digital watermarking and AI-generated content detection, which are essential technologies for preventing the misuse and abuse of AI. Alibaba has developed a closed-loop system of

Conflict Resolution Hierarchy

Tier	Set By	Override-ability	Design Intent
Root	Model Spec	Cannot be overridden	Cross-cultural consensus floor
System	Model Provider	By system instructions only	Jurisdictional Compliance & Regulated User Cohorts
Developer	App Creator	By developer instructions	Differentiated AI applications
User	End User	By dev or user input	Personalization & Consistency

Source: Alibaba Group Security Department. The hierarchy turns a value declaration into an engineering-grade target.

Overall Approach to Agent Security

A Three-Pillar Approach to Building a Secure and Trustworthy Agent Ecosystem

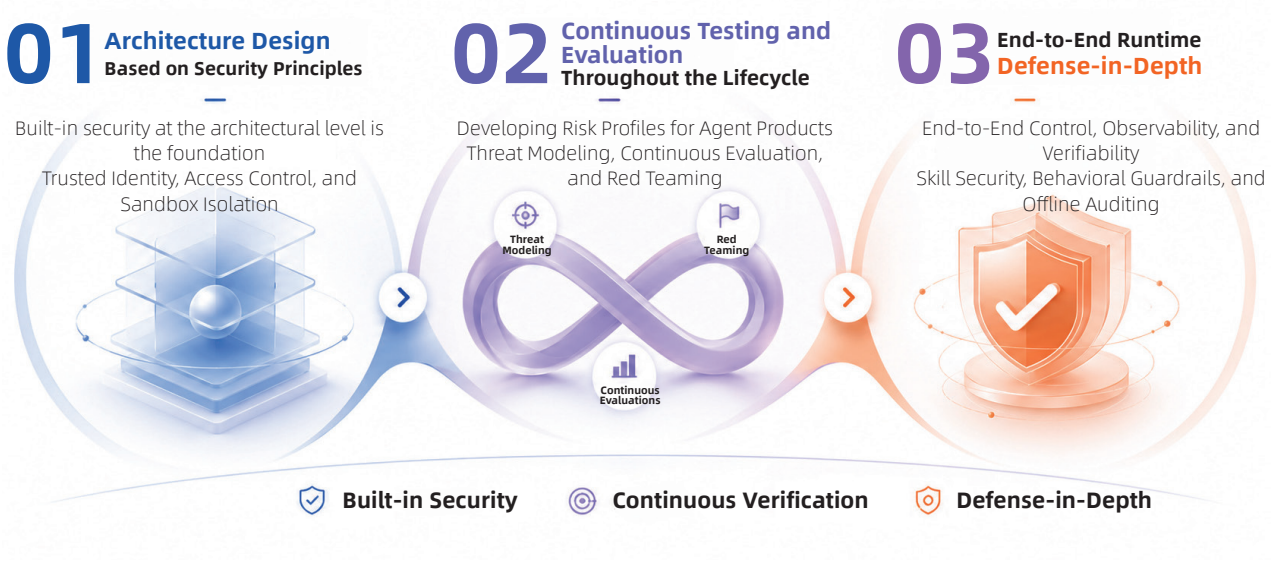

Source: Alibaba Group Security Department

"detection, watermarking, and verification". The system can first identify AI-generated content, then apply identifying markers through visible labels, invisible metadata markers, or watermarks embedded within the content itself, and finally support traceability and compliance verification. It covers multiple modalities, including text, images, audio, and video, and has already been deployed in business scenarios across Amap, Xianyu, Taobao, and Tmall. This demonstrates how model safety now extends beyond generation itself to the circulation, identification, provenance, and governance of AI-generated content.

The fifth layer is the security of cloud-based large-model services. Public cloud platforms must do more than provide models. They must also offer comprehensive security safeguards throughout the lifecycle of model training, inference deployment, tool invocation, service orchestration, monitoring, and operations. For cloud-based large-model services, we emphasize defense in depth, zero trust, and the principle of least privilege. Identity verification, access control, runtime protection, data lifecycle protection, and DevSecOps (the integration of development, security, and operations) are embedded into AI infrastructure.

The sixth layer is agent safety. As AI agents become integrated into business workflows, the risks requiring attention extend beyond model outputs to the reliability of agent perception, the explainability of decision-making, permission management, and multi-agent coordination. To address these challenges, Alibaba has developed dedicated permission management systems and plugin sandboxing mechanisms for AI agents, implementing isolation across networks, execution environments, users, permissions, and data. End-to-end auditing records agent identities, operation types, and key data identifiers. Combined with independent guardrails for agent behavior, these capabilities enable anomaly monitoring, detection of deviations from intended objectives, and tracing and forensic investigation in the event of attacks or loss of control.

Alibaba's approach to model safety can therefore be summarized as follows: rather than adding a temporary safety gate outside the model, we seek to build safety into the model itself, into engineering workflows, and into cloud infrastructure, as well as into the permissions, identities, auditing mechanisms, and accountability structures that govern how AI agents operate in the real world.

Open Source and Openness: Broadening Participation

AI cannot become widely integrated into society through the efforts of a few companies alone. Bringing AI into a wide range of industries requires researchers, developers, businesses, social organizations, and public institutions to draw on their respective expertise and capabilities to build AI applications suited to their specific needs and circumstances.

As a platform company, Alibaba can make AI capabilities available to different groups through multiple channels: open-source models for technically capable developers and research institutions; cloud platforms and API services for businesses and organizations that need AI capabilities but lack the resources to build their own infrastructure; and ready-to-use applications that bring AI capabilities to knowledge workers and ordinary users without requiring a technical background.

- **Open-source models:** lowering technical barriers and accelerating the diffusion and accessibility of AI. As of September 2026, Alibaba had open-sourced more than 460 models in the Qwen family, giving rise to over 300,000 derivative models and exceeding 3 billion cumulative downloads worldwide. Qwen models have ranked first globally in authoritative benchmarks for several languages, including Arabic and Japanese, and have been widely adopted by leading international communities such as Hugging Face. Models of different sizes and capabilities give developers the flexibility to choose according to their tasks, devices, and cost constraints.
- **Community collaboration:** from shared use to shared development. The value of the Qwen ecosystem comes not only from its foundation models, but also from the vitality of community collaboration. Much of Qwen's ecosystem value today has not been created by Qwen alone. It has grown organically through the contributions of its community.

The vast majority of the more than 100,000 derivative models hosted on ModelScope are the results of secondary development, fine-tuning or modifications made by other developers on top of Qwen. On Hugging Face, most of the 28,531 GGUF conversion variants of Qwen models for local deployment have been created by developers around the world building on Qwen. This shows that Qwen's success lies not merely in a model product, but in an open ecosystem driven by the

Derivatives on Hugging Face by Organization

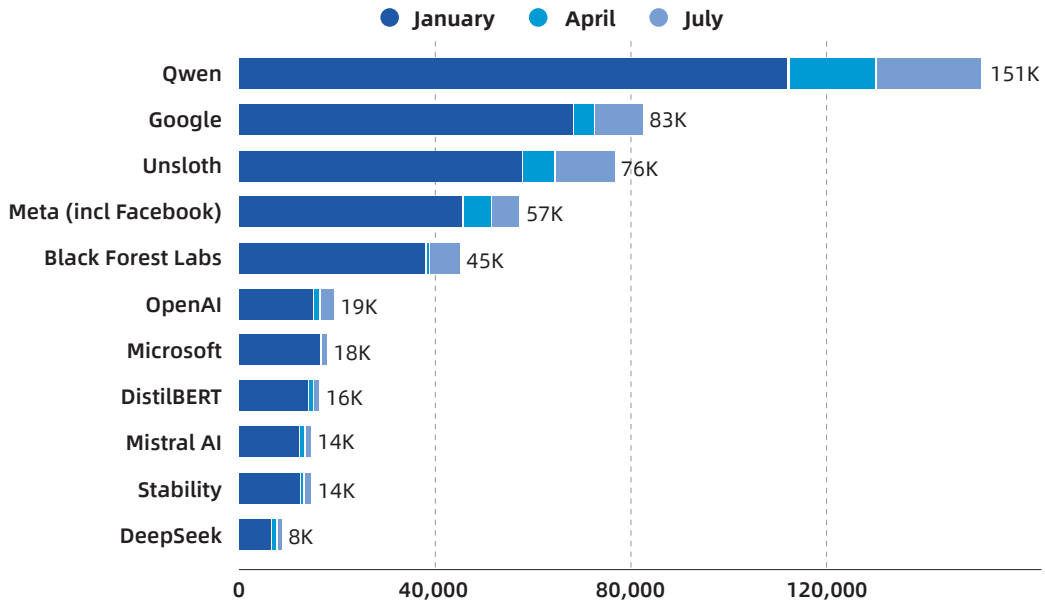

Source: State of Open Models: Summer 2026 Observations, Hugging Face

community and capable of sustaining its own growth.

- **Multilingual deployment:** opening up possibilities for developers across different cultural and linguistic backgrounds. Accessibility is not only about overcoming technical barriers. It is also about overcoming language barriers. Speakers of low-resource languages may struggle to receive services comparable to those available in major languages because of limited training data. Multilingual open-source models have the potential to help address this disparity.

As of September 2026, the Qwen family supported more than 200 languages and dialects, with expanded coverage of low-resource languages in South Asia, Oceania, Africa, and other regions.

The relationship between open source and AI safety has been an important topic of discussion over the past year. Numerous cases have shown that safety does not depend solely on whether a model is open or closed. It is shaped by a range of factors, including model capabilities, deployment boundaries, permission controls, monitoring, and auditing.

At the same time, open source has demonstrated characteristics that can contribute to stronger safety. The transparency made possible by releasing model weights and training methods can make potential risks easier to identify and discuss. Collaborative governance involving multiple stakeholders can help detect errors and biases more promptly. Improvements in model specifications and intrinsic safety mechanisms are also helping strengthen public and regulatory trust in large AI systems.

Open Source, Open Access, Open Possibilities

A survey on "AI Talent Development from the Perspective of Open-Source Communities", covering 200,000 users of ModelScope and its partner communities, found that amid the profound disruption and restructuring brought about by AI, open-source communities are playing an increasingly important positive role. They have become essential platforms for the diffusion of AI technologies, the incubation of applications, and the development of shared governance principles, while also creating new pathways for talent development.

Impact of Open Source Communities and AI on Career Development

(Multiple responses allowed; over 10,000 respondents.)

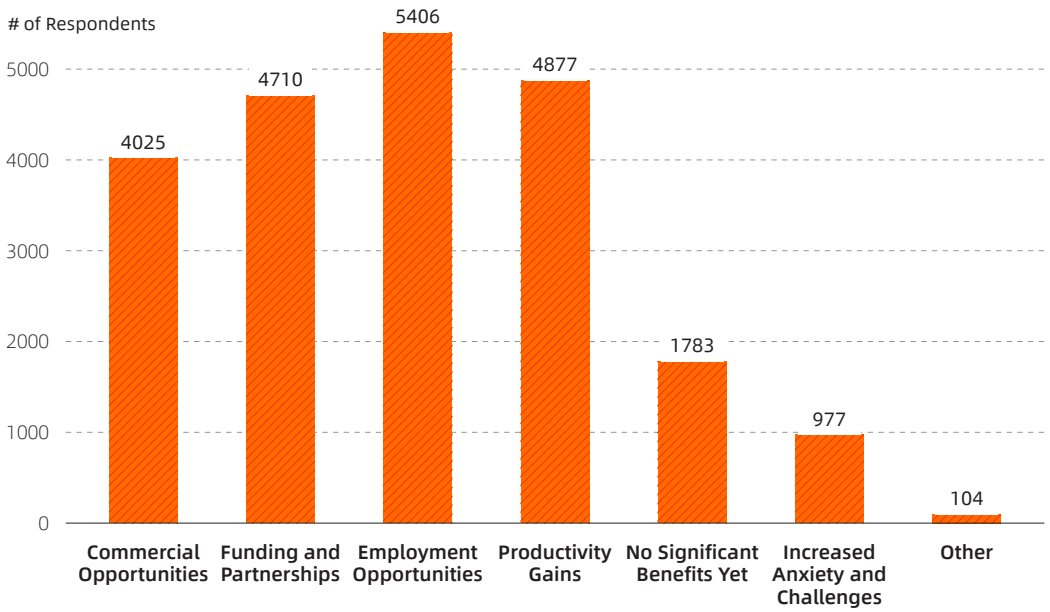

Source: 2026 AI Open Source Talent Report, p. 17

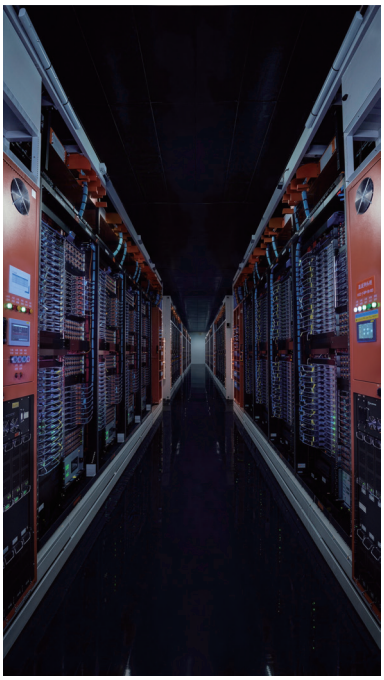

Green Infrastructure and Compute-Power Coordination

Over the past year, the development of AI infrastructure has accelerated worldwide. Rapid growth in demand for computing power has driven a corresponding increase in electricity demand, making energy and power supply major constraints on AI development. At the same time, the global challenge of climate change is becoming more pressing. Even in China, the world's largest electricity producer, the timeline for meeting carbon neutrality commitments is drawing closer.

Green and low-carbon development is one of Alibaba's six guiding principles for AI for Good. In 2021, Alibaba publicly announced its carbon neutrality goals. Under these commitments, Alibaba Cloud, as a provider of digital infrastructure, aims not only to achieve carbon neutrality in Scope 1 and Scope 2 emissions alongside Alibaba Group, but also to take the lead in achieving Scope 3 carbon neutrality and become a green cloud provider.

To advance these goals, we have pursued end-to-end system optimization across data center planning, construction, and operations. Our approach has three components: integrating compute-power coordination—coordination between computing demand and electricity supply—into planning to make the most of energy advantages; adopting better architectures to raise the upper limits of efficiency; and continuously improving operational efficiency through intelligent management and scheduling.



Alibaba Cloud Ulanqab Data Center

- **Compute-power coordination:** making the most of energy advantages through integrated planning. Alibaba Cloud considers the availability, stability, and cost of electricity supply, as well as access to clean electricity, among the key factors in selecting data center locations. This helps align the geographic distribution of computing infrastructure with power supply conditions.

  In fiscal year 2026, clean electricity accounted for 73.6% of the electricity used by Alibaba Cloud's self-built data centers, an increase of 9.6 percentage points from the previous fiscal year. Data centers in Guangdong, Jiangsu, and Sichuan achieved 100% clean electricity coverage. Across leased data centers worldwide, the share of clean electricity increased to 40.8%, up 18.4 percentage points year over year.

  Meanwhile, in regions where energy supply is relatively constrained, Alibaba Cloud is exploring locally adapted approaches to developing large-scale energy supply bases for data centers. These include new models such as direct connections to green electricity sources, advanced nuclear energy, and offshore energy islands.

> “Over the past year, we have also paid increasing attention to water consumption associated with data center cooling. Technologies such as liquid cooling can reduce reliance on evaporative heat dissipation. However, as computing density continues to increase, water resource constraints need to be incorporated earlier into data center site selection and cooling system design.

Annual average PUE across Alibaba Cloud's self-built data centers

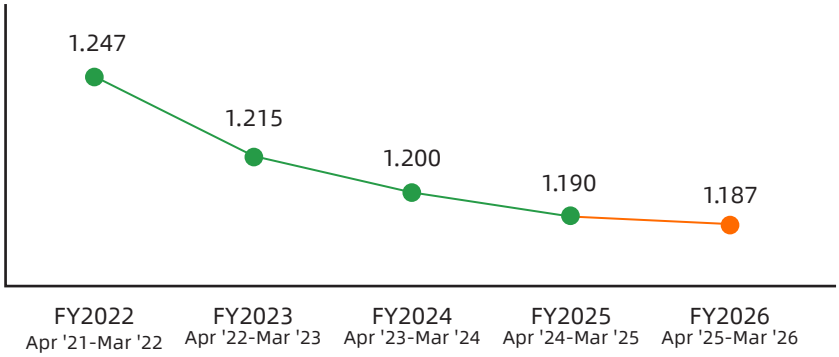

Share of clean electricity at Alibaba Cloud's self-built data centers

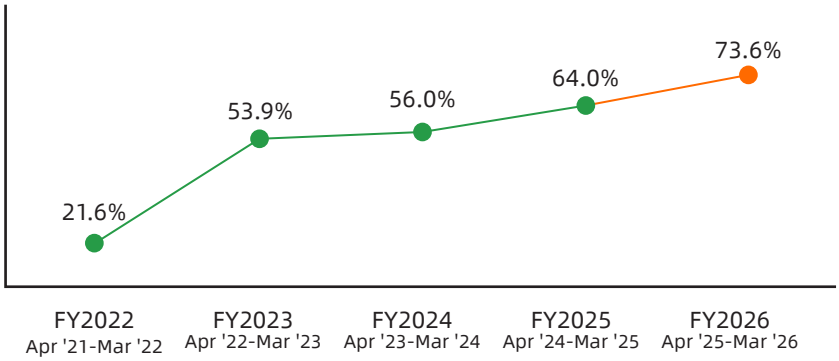

- **Better architectures:** raising the upper limits of efficiency. Once a data center location has been selected, its architectural design determines the upper limits of facility efficiency. In power delivery architecture, Alibaba is exploring ways to reduce the number of conversion stages between the power source and the chip. Building on the high-voltage direct-current power architecture jointly developed with partners, known as the Panama Power Supply, the latest generation of data centers supports 800V DC power delivery. We are also exploring technologies such as solid-state transformers (SSTs), which could provide a technical pathway for directly supplying high-voltage DC power from renewable energy sources.

  In terms of construction, Alibaba Cloud has adopted the CUBE 5.0 architecture and a fully modular design. Five major systems, including power supply, cooling, security, intelligent management, and fire protection, are pre-integrated into modules at the factory. Combined with liquid-cooling equipment, this architecture supports high-density computing deployments while substantially shortening on-site construction time. The overall modularization rate has increased from 30% to 90%, enabling agile data center delivery within 100 days.

- **Operations and scheduling:** continuously improving energy efficiency. During data center operations, Alibaba Cloud combines digital twin models with AI optimization algorithms to simulate and analyze cooling system operating modes and key equipment parameters. The resulting control strategies are validated through simulation before being deployed to the actual systems, enabling continuous improvements in operational energy efficiency.

AI-driven workload optimization and cooling control have been deployed across 24 self-built data centers. In fiscal year 2026, the average annual Power Usage Effectiveness (PUE) of Alibaba Cloud's self-built data centers reached 1.187. PUE measures the ratio of total facility energy consumption to the energy consumed by IT equipment.

Recognizing that AI workloads have flexibility in both time and location, Alibaba Cloud has developed a framework for coordinated computing and power scheduling across regions. Through forecasting and real-time dynamic management, the framework responds to grid capacity and electricity price signals, flexibly shifting computing workloads to times and locations where electricity is more abundant, while also participating in ancillary grid services. Data centers are therefore beginning to evolve into active participants in balancing energy systems.

Over the past year, we have also paid increasing attention to water consumption associated with data center cooling. Technologies such as liquid cooling can reduce reliance on evaporative heat dissipation. However, as computing density continues to increase, water resource constraints need to be incorporated earlier into data center site selection and cooling system design.

Developing green AI infrastructure cannot be accomplished by any single company. It requires close collaboration throughout the industrial value chain. Through the Yungu Token and Watt Forum, Alibaba Cloud has brought together nearly 50 organizations from industry, academia, research, and application sectors to jointly explore frontier issues in "AI + Energy" and compute-power coordination, and to launch collaborative research projects.

In developing overseas infrastructure, we are also working with partners to explore delivery models that combine the fully modular CUBE 5.0 architecture with green energy solutions, adapting them to local energy structures and resource conditions.


Chapter 3
AI for Good in
Action 2026
小雨正站在窗边擦玻璃，水珠顺着窗边滑下来，
像在玩透明的滑滑梯。

> Compared with last year, the cases presented in this report cover a broader range of sustainable development issues and involve more diverse participants and beneficiaries.

AI offers new possibilities for expanding the reach of initiatives for social good, reducing service costs, and responding to individual needs. Building on our full-stack AI capabilities, open-source technologies and communities, and real-world applications and feedback, we continue to explore ways to empower organizations and serve individuals. Our goal is to bring AI beyond the efforts of a small group of early adopters into broader use, translating technological progress into tangible, accessible benefits.

Drawing on our core businesses, Alibaba continues to explore how AI can respond to social needs. Building on the foundational work described in the previous chapter, this chapter brings together new developments in our AI for Good initiatives since last year's report, presenting specific applications across six areas and their connections to the United Nations Sustainable Development Goals (SDGs).

The six areas include both initiatives focused on specific fields and efforts that cut across multiple sectors. Health, education, and scientific research respectively focus on how AI can improve healthcare services, support learning and teaching, and expand the capacity for scientific discovery. These initiatives are primarily associated with Good Health and Well-being (SDG 3), Quality Education (SDG 4), and Industry, Innovation and Infrastructure (SDG 9), among other goals. Economic opportunity, safe and inclusive communities, and green development have broader, cross-sector dimensions. Economic opportunity focuses on employment, entrepreneurship, and productivity growth. Safe and inclusive communities focus on creating the conditions for different groups to access technology, services, and opportunities for participation. Green development addresses energy and resource use, climate action, and ecological conservation. Each of these three areas is connected to multiple SDGs and intersects with our work in health, education, and scientific research.

Compared with last year, the cases presented in this report cover a broader range of sustainable development issues and involve more diverse participants and beneficiaries. They also reveal several emerging approaches. Industry expertise is being transformed through AI into capabilities that are easier to access and reuse. Scientific research is exploring new pathways to knowledge discovery with the help of AI. And AI applications are gradually evolving from standalone tools into components of organizational workflows, becoming integrated into more aspects of production and everyday life.

This chapter presents not only progress achieved to date but also initiatives that remain under development. Some have produced initial results, others are still being tested for feasibility, and still others are adjusting their direction based on practical feedback. We hope to share the lessons learned, opportunities, and challenges, and to work with more partners to turn the potential of AI for Good into sustained action.


3 GOOD HEALTH AND WELL-BEING
AI for Health
4 QUALITY EDUCATION
AI for Education
9 INDUSTRY, INNOVATION AND INFRASTRUCTURE
AI for Science
AI for Green Development
7 AFFORDABLE AND CLEAN ENERGY
12 RESPONSIBLE CONSUMPTION AND PRODUCTION
13 CLIMATE ACTION
15 LIFE ON LAND
AI for Economic Opportunities
8 DECENT WORK AND ECONOMIC GROWTH
9 NDUSTRY. NNOVATION AND NFRASTRUCTUR
10 REDUCED INEQUALITIES
AI for Safe and Inclusive Communities
10 REDUCED INEQUALITIES
11 SUSTAINABLE CITIES AND COMMUNITIES
16 PEACE, JUSTICE AND STRONG INSTITUTIONS

AI for Health

DAMO Academy's Non-Contrast CT + AI

Alibaba Health's Hydrion

AI for Starry Kids

Healthcare systems worldwide face the dual pressures of rapidly growing demand and persistently constrained supply.

In China, across different areas of healthcare, from disease screening and clinical diagnosis to drug development and interventions for people with special needs, challenges such as shortages of professional resources, limited service efficiency, and unequal distribution of high-quality healthcare resources remain widespread.

AI has the potential to improve the efficiency of professional healthcare services and expand overall healthcare capacity. Predictive models, decision-support tools, and generative AI applications are already beginning to enter clinical practice and everyday care. Yet technical feasibility does not automatically translate into practical usability. Healthcare services are highly localized, fragmented, and strictly regulated, which means AI applications in healthcare must be adapted to local circumstances throughout design, training, deployment, and evaluation. Patients, physicians, hospitals, insurers, pharmaceutical companies, regulators, and other stakeholders need to establish a new balance among clinical safety, organizational workflows, and commercial incentives.

DAMO Academy's Non-Contrast CT + AI

A Low-Cost, High-Value Pathway for Advancing Healthcare Infrastructure

Non-contrast CT is among the most widely used and standardized imaging technologies in clinical practice, generating enormous volumes of medical images every year. Yet much of their potential value has long remained untapped. Traditionally, a CT scan has generally served the diagnosis of a particular condition, with physicians focusing on a specific suspected lesion while substantial information about other diseases contained in the same images has not been systematically utilized. Even for a particular lesion, the inherent limitations of visual interpretation have

Alibaba Is Also Exploring AI-Powered Voice Screening for Diseases

In 2026, "Never Forget Me," an AI-powered early screening project for Alzheimer's disease, won first prize in the Inclusive Elderly Care category at the second "Small Ideas, Big Impact" AI for Good Innovation Challenge. Powered by Qwen-Omni's multimodal capabilities, the application uses speech and acoustic features, including speaking rate, pause ratios, semantic coherence, and response patterns, to conduct an initial assessment of cognitive risk during a single phone call.

DAMO Academy AI Detects Subtle Lesions in CT Scans, Tackling Major Challenges in Global Healthcare

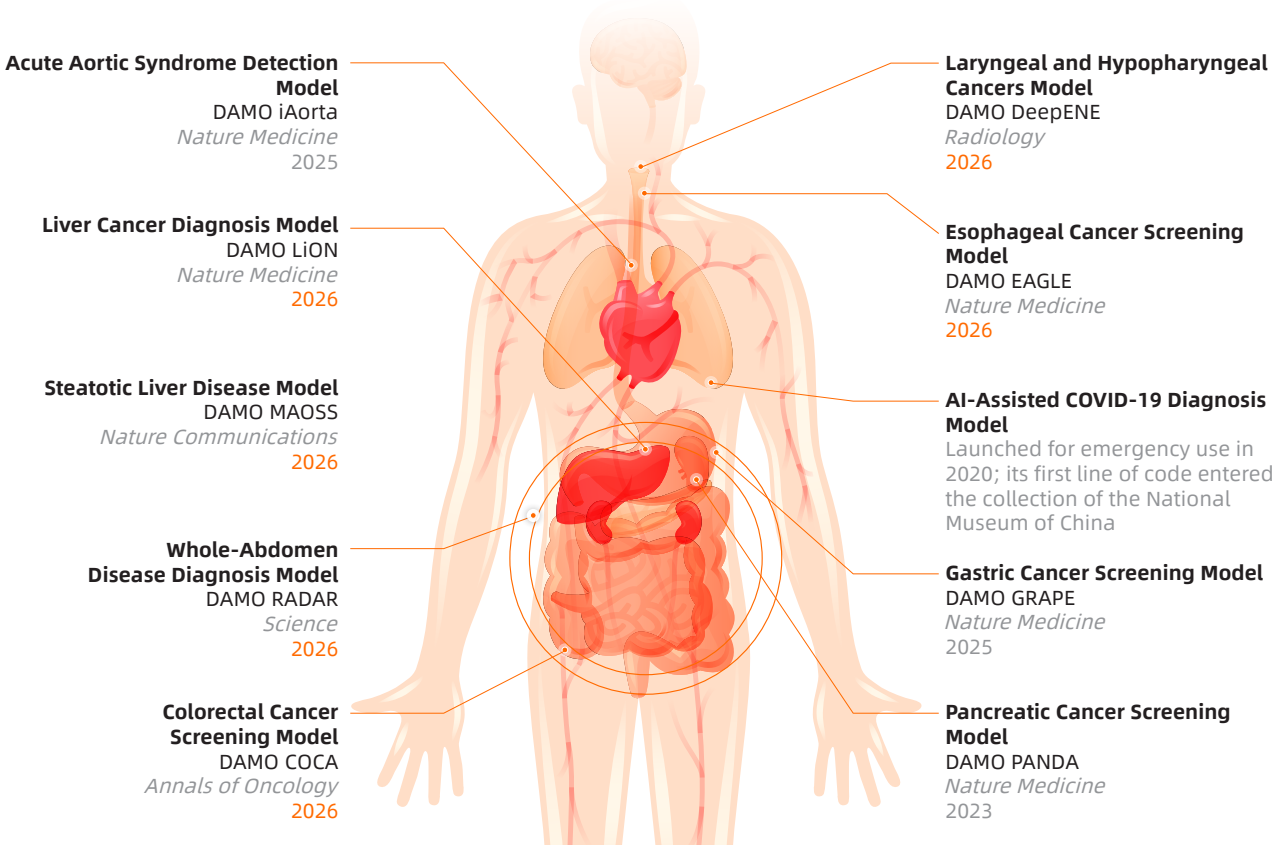

constrained the further extraction of value from medical imaging.

AI is enabling non-contrast CT to evolve continuously. Without requiring changes to existing equipment or clinical workflows, DAMO Academy's "Non-Contrast CT + AI" solution continues to expand the range of questions that can be answered from the same images. A routine examination can gradually become a comprehensive healthcare entry point supporting disease screening, risk assessment, and clinical decision-making. This is helping AI-powered healthcare move from single-disease screening toward screening for multiple conditions in a single scan.

In 2026, the principal advances in Non-Contrast CT + AI have been the continued improvement of multi-cancer screening capabilities. Meanwhile, DAMO Academy has progressively expanded the capabilities built around the same non-contrast CT images to include emergency alerts, chronic disease assessment, and other applications, creating a continuously evolving medical AI capability system.

As early as 2023, DAMO Academy released a pancreatic cancer screening model, demonstrating for the first time internationally the feasibility of early pancreatic cancer screening using non-contrast CT. Nature Medicine described the development as signaling that "a golden era of cancer screening based on medical imaging AI is on the horizon." Over the past year, screening models for gastric, colorectal, liver, and esophageal cancers have been released in succession, providing further evidence that organs traditionally considered difficult to assess using non-contrast CT may nevertheless contain rich and detectable disease signals.

Compared with last year, DAMO Academy is also

exploring additional cancers, including breast and kidney cancers. In emergency care, capabilities for identifying critical conditions such as pulmonary embolism continue to expand. In chronic disease management, models for osteoporosis and fatty liver disease have begun to be used in practice, while exploration has begun in areas such as valvular heart disease. At the same time, the DeepHTD model has been used experimentally to predict the degree of pathological response in lung cancer patients following neoadjuvant therapy, providing additional reference information for clinical decision-making.

Different models target different diseases, but they share the same underlying technological approach. As the range of detectable conditions expands, the marginal cost of adding another screening capability to the same CT scan continues to decline, increasing the value of existing healthcare infrastructure. More importantly, these additional capabilities do not require patients to undergo more examinations or incur additional radiation exposure, nor do they require changes to physicians' established clinical workflows.

In September, DAMO Academy open-sourced DAMO RADAR, a general-purpose medical imaging AI model. Designed for diagnosis using contrast-enhanced abdominal CT, the model covers 18 organs and can identify 146 diseases and medical conditions in a single examination. Its ability to distinguish patients with relevant conditions from those without them, measured by the area under the receiver operating characteristic curve (AUC), reached the level of expert radiologists for the first time. With RADAR's assistance, radiologists' ability to avoid missed diagnoses improved by 10%, while interpretation time was reduced by more than 30%. The research was published in Science.

More general-purpose screening technologies have the potential to reduce the high cost of multiple examinations to a fraction of that cost through a single non-contrast CT scan combined with AI. They also make opportunistic screening for conditions such as low bone density and fatty liver disease possible, creating opportunities for early detection and intervention among people who might otherwise never enter a screening program.

General-Purpose Medical Imaging AI Achieves Expert-Level Performance for the First Time

DAMO RADAR demonstrates superior performance across 146 tasks involving the detection of medical signs and diagnosis of diseases.

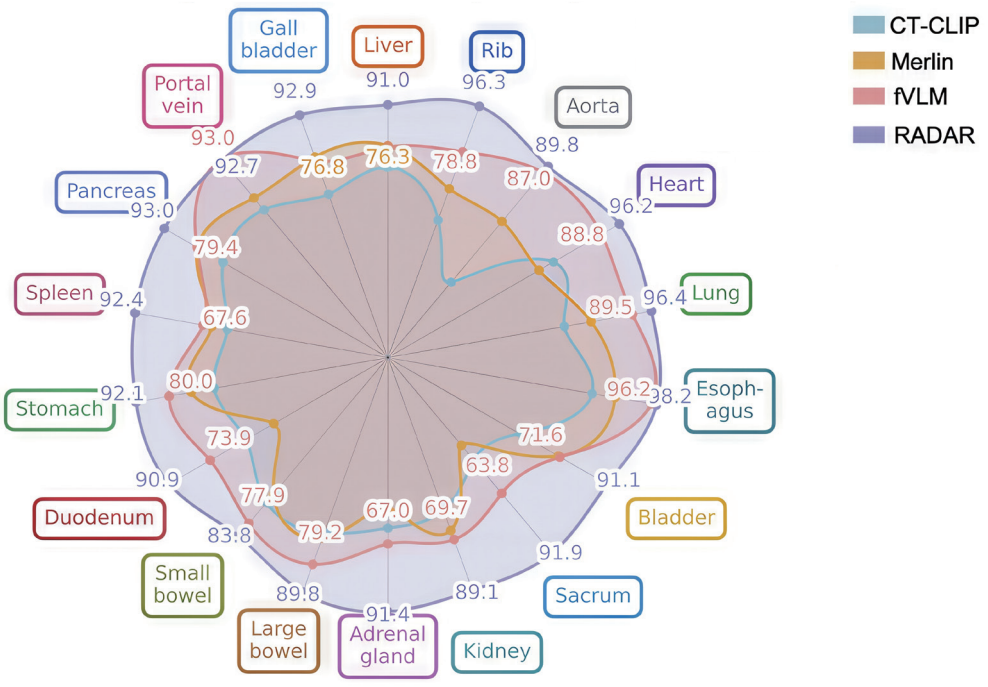

An article published early this year by China's National Cancer Center in a Nature portfolio journal noted that China accounts for approximately 17% of the world's population but bears around 25% of the global cancer burden. The five-year relative survival rate for cancer patients remains lower than in developed countries such as the United States. Early detection is critical to addressing cancer, not only because it can save more lives, but also because it can help reduce healthcare costs.

However, the large-scale deployment of AI in healthcare also depends on establishing a continuously evolving pathway from data acquisition and model training to clinical validation. During the COVID-19 pandemic, hundreds of CT imaging AI tools emerged worldwide, yet most failed to achieve large-scale adoption after the pandemic. A retrospective assessment by the Alan Turing Institute in the United Kingdom identified several common causes: insufficient training data quality, a lack of real-world validation, and failure to integrate with clinical workflows.

China has extensive medical imaging resources and a large patient population, with major medical institutions continuously accumulating real-world clinical cases. At the same time, compared with highly market-driven healthcare systems, China also benefits from relatively lower organizational costs for conducting collaborative clinical research across hospitals, building high-quality datasets, and advancing real-world validation. These conditions provide an important foundation for continuously expanding the capabilities of Non-Contrast CT + AI in China.

Truly effective medical AI should not require physicians to adapt to AI. Instead, AI should adapt to physicians' existing workflows. Acute aortic syndrome provides a representative example. The condition can progress rapidly and carries substantial risks of missed or incorrect diagnosis. More than half of patients undergo non-contrast CT during their initial presentation.

To address this clinical challenge, the DAMO iAorta model is integrated with hospital imaging information systems. It can interpret non-contrast CT images in real time within seconds. When a high-risk case is detected, an alert is pushed to the physician's interface, prompting prioritized review. The model also generates heatmaps identifying suspicious regions, allowing physicians to understand the basis of the AI assessment more intuitively. It can also serve as a tool for physicians to continue learning and accumulating experience, accelerating the development of younger doctors and helping spread clinical expertise across different levels of healthcare institutions.

This value is particularly evident in primary and local healthcare systems. Around this model, a hospital in Shaoxing, Zhejiang Province, piloted a 15-minute emergency response network. Branch hospitals upload CT images, while AI at the central hospital provides real-time alerts. Pilot data showed that the average time to confirmed diagnosis fell from 4.3 hours to 1.7 hours, while diagnostic sensitivity among junior physicians improved by approximately 40%. Healthcare capacity is beginning to expand beyond the limitations imposed by the number of medical specialists, supported by the broader and lower-cost availability of algorithmic capabilities.

AI for health must undergo multiple rounds of rigorous, evidence-based clinical validation. Medical AI needs to accumulate evidence from prospective clinical studies and real-world data, demonstrating that it can improve early detection, support better clinical decisions, and ultimately benefit patients. This is not a one-time process, but one of continuously gathering evidence and refining models through real-world application.

An increasing number of healthcare institutions are participating in the ongoing development of clinical evidence for medical AI. In 2026, DAMO Academy signed agreements with leading domestic and international CT equipment manufacturers, including GE HealthCare and Neusoft Medical Systems, and established partnerships with multiple medical institutions. The first group of 10 hospitals in China has deployed DAMO iAorta, the AI model for aortic emergencies. DAMO Academy's multi-cancer screening models have cumulatively provided more than 20 million AI-assisted service instances across 11 countries and regions worldwide. We see a potential new pathway for supporting medical AI innovation, one that progresses from algorithms to evidence and ultimately to institutional adoption.

Alibaba Health's Hydrion

Earning Physicians' Trust Through Evidence, Advancing Clinical Care Through AI

The essence of medical decision-making lies in organizing and evaluating a vast body of medical evidence within limited time. China has 5 million physicians who collectively deliver approximately 10 billion patient consultations and treatments every year. Every clinical encounter involves an individualized synthesis of medical knowledge, and the quality of these decisions forms the foundation of the healthcare system.

In May 2026, Alibaba Health launched Hydrion, an AI-native application designed for physicians, to explore an evidence-based medical decision-support system. Its goal is not to replace physicians' judgment, but to make existing medical knowledge more reliable and efficient to use in clinical decision-making, helping doctors retrieve evidence, conduct differential analyses, and validate decisions for individual cases.

This initiative addresses a fundamental challenge facing modern healthcare. Medical knowledge continues to expand, disease classifications are becoming increasingly detailed, and rare and less common cases continue to emerge. Yet physicians' ability to apply knowledge remains constrained by time, experience, and access to information. As of September 2026, more than 40 million medical publications were publicly available worldwide, with the volume continuing to grow rapidly. Traditionally, medical knowledge has been disseminated primarily through research papers, clinical guidelines, and academic conferences. Physicians must search for, select, understand, and apply this information within limited time. The pressure is even greater in regions with limited resources. Against this backdrop, diagnostic errors and mistakes in medical decision-making remain important challenges for modern healthcare systems.

Meanwhile, although general-purpose large language models possess strong language understanding and generation capabilities, they cannot readily meet the high standards of accuracy, reliability, and traceability required in medical settings. They may generate content that appears plausible but lacks supporting evidence, contradicts established facts, or is entirely fabricated, a phenomenon commonly referred to as "hallucination." Their performance may also be limited by outdated training data, insufficient specialist knowledge, and unstable reasoning. General-purpose models therefore cannot be used directly as the basis for clinical decisions. This highlights the need for more reliable and verifiable medical AI systems grounded in professional knowledge.

For healthcare systems, this is why trustworthy AI matters more than simply more intelligent AI. Only by establishing evidence and building trust can medical AI

Hydrion Outperforms Peers on Evidence-Based Medicine Benchmarks

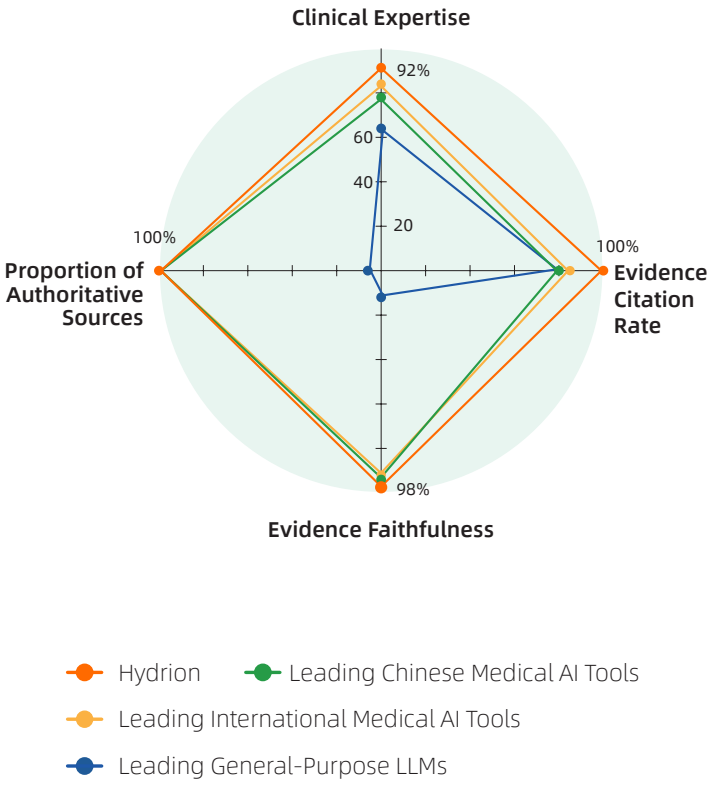

Note: Clinical expertise refers to the accuracy of key claims, supporting arguments, and conclusions; evidence citation rate refers to the proportion of answer paragraphs that cite supporting evidence; evidence faithfulness refers to the consistency of claims, supporting arguments, and conclusions with the cited evidence; and the proportion of authoritative sources refers to the share of cited sources that are authoritative and professionally recognized.

move from the laboratory into the consultation room, transforming technological capabilities into healthcare capabilities. To achieve this, it must provide traceable evidence sources, explainable reasoning processes, and continuously validated clinical feedback.

As of September 30, 2026, Hydrion had established content partnerships with the BMJ Group, NEJM Group, and JAMA Network, representing three of the four leading international medical journal groups, as well as authoritative Chinese academic institutions including the Chinese Medical Association Publishing House, People's Medical Publishing House, and the China Anti-Cancer Association. It has also worked with leading industry experts to establish a Medical AI Expert Committee and jointly develop medical AI datasets.

Authoritative information sources are only the foundation. More important is an evidence-driven reasoning mechanism. Hydrion grounds AI's use of medical evidence within an evidence-based medicine framework, presenting physicians with the sources, levels of evidence, applicable contexts, and limitations alongside its responses. Currently, Hydrion has achieved 92% accuracy, a 100% paragraph-level evidence-citation rate, and 98% evidence fidelity in medical scenarios. It continues to improve the clinical user experience through feedback from physicians' real-world use. Approximately 30% of new physician users come through recommendations from colleagues, further illustrating that medical AI is undergoing a critical transition from technological recognition to clinical trust.

As adoption expands, Hydrion is helping move medical AI beyond isolated technical validation toward broader applications across evidence-based clinical practice. Historically, advances in medical knowledge have primarily reached clinical settings through specialists, journals, and academic conferences, with limitations in timeliness and uneven coverage. In the future, AI has the potential to embed validated medical evidence into everyday clinical workflows more frequently and at lower cost. This can strengthen physicians' continuous learning and clinical decision-making capabilities while improving access to high-quality healthcare resources and advancing health equity.

AI for Starry Kids

Supporting People with Special Needs in Reconnecting with Themselves and the World

Autism spectrum disorder is a neurodevelopmental condition. In China, there are more than 2 million children with autism between the ages of 0 and 14, with approximately 200,000 new cases each year requiring timely intervention and support. However, effective intervention is highly dependent on each child's interests, cognitive development, and behavioral characteristics, creating three major challenges: high development costs, demanding professional requirements, and the need for highly individualized approaches.

With support from Alibaba Philanthropy, Alibaba volunteers and special education experts launched the "AI for Starry Kids" storybook agent in 2024, exploring how generative AI could help overcome this challenging combination of constraints. Parents or teachers need only enter simple instructions. The storybook agent can then use Qwen models to generate corresponding stories, images, and audio, orchestrating these capabilities to create narrated picture books that can be dynamically adjusted.

In 2026, AI for Starry Kids underwent an upgrade at the model level. Qwen-Long, a model designed for long-context understanding and generation, can incorporate complete special education guidance documents, classic picture-book structure templates, materials on emotional development psychology, and other relevant context. It has also been specifically optimized for

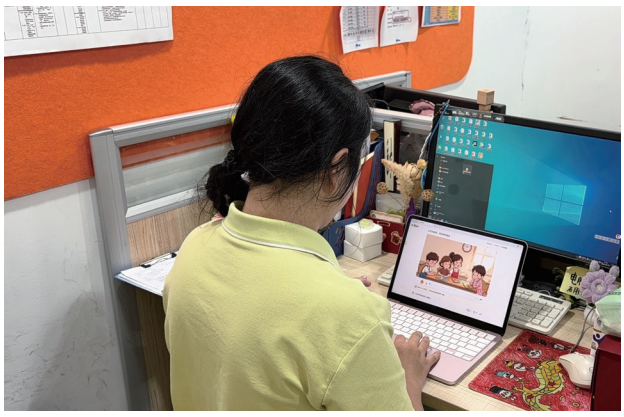

Special Education Teachers Use "AI for Starry Kids" to Support Lesson Planning.

semantic understanding in this specialized domain, helping generated content balance psychological and educational guidance with children's ability to engage with and understand it. The multimodal visual engine Qwen-Image maintains consistent characters across pages. Together, the two models achieve semantic alignment between text and images, ensuring that the words and visuals on each page serve the intended intervention objectives.

This is a representative AI-native application. Improvements in model capabilities have not only enhanced efficiency and quality but also created user experiences that previously did not exist. Over the past year, the amount of sample audio required for voice cloning has fallen from several minutes to just four seconds. This allows even children with autism who have difficulty sustaining verbal expression to participate in content creation, moving from listening to others toward telling their own stories. AI is lowering barriers to equal participation in everyday life for people with special needs, helping them reconnect with themselves and the world.

As of June 2026, AI for Starry Kids had delivered more than 400,000 storybook generation and reading services, twice the number recorded during the same period last year. It had reached more than 20,000 children with autism and special education teachers.

AI-Native Product and Technology Architecture

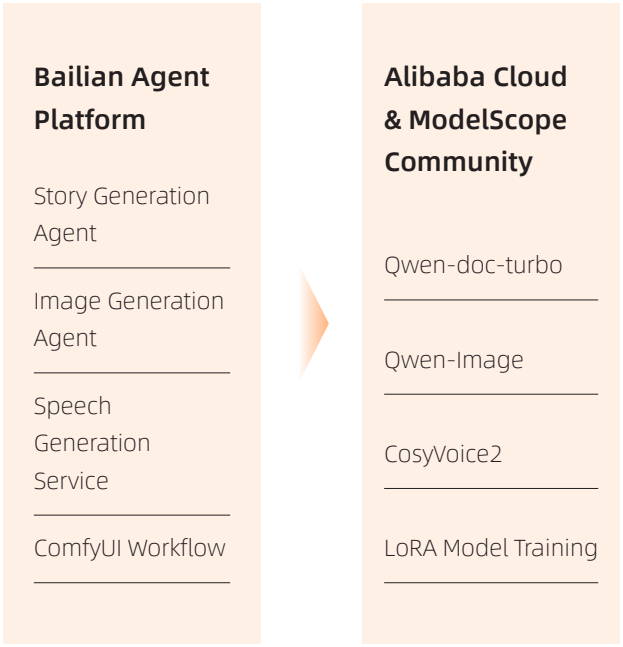

Looking Ahead

We are entering a new stage of large-scale AI adoption in healthcare. Just as autonomous driving, another highly regulated field, has undergone more than two decades of rigorous iteration before gradually demonstrating performance exceeding that of human drivers, healthcare AI must also build trust through continuous validation. The challenges ahead include accelerating adoption in settings already supported by sufficient evidence, while filling critical evidence gaps through clinical trials. Only then can medical AI move from localized validation toward widespread adoption.

High-quality data remains a critical foundation for the large-scale development of medical AI. In the future, medical data needs to evolve from fragmented resources into infrastructure that can be securely shared and sustainably utilized. Through multimodal data integration, privacy-preserving computation, and continuous feedback mechanisms, healthcare systems have the potential to establish a mutually reinforcing cycle of evolution among data, models, and clinical practice.

AI for Education

Cloud for Youth Education Support Program
Yungu School
Qwen App's Gaokao Admissions Agent
University AI Support Program
RynnBot Platform
Alibaba Cloud Science Education Base

After decades of development, China has built the world's largest education system, with substantial achievements in educational quality. Despite this progress, significant challenges remain. Access to high-quality educational resources is still constrained by geography, teacher availability, and infrastructure, leaving room for further improvements in educational equity. Traditional educational models and assessment systems also tend to place excessive emphasis on knowledge acquisition and standardized outcomes, limiting opportunities for personalized learning, critical thinking, and the development of broader capabilities.

As technology evolves rapidly, education systems must also address a new question: how can people of different ages and from different regions acquire the ability to understand and use emerging technologies?

AI is changing how educational resources and capabilities are delivered. From personalized learning in primary and secondary education to intelligent decision support for university applications and teaching and research in higher education, AI makes it possible to combine personalized services with delivery at scale. Yet as AI capabilities themselves become fundamental skills and reshape the talent requirements of future society, integrating AI into education must also empower more people to use and direct AI effectively. Helping teachers, students, and schools proactively redesign the learning process is becoming a new priority for AI in education.

Cloud for Youth Education Support Program

Advancing Educational Equity from Access to Resources to Access to Capabilities

AI has powerful knowledge generation capabilities and can support learners across different educational levels, stages of development, and cultural backgrounds. But whether these capabilities widen educational divides or genuinely promote educational equity depends to a significant extent on how effectively education systems prepare students to use AI. Teachers are evolving from transmitters of knowledge into designers and facilitators of learning, while students are moving from passively receiving answers toward asking questions proactively and making independent judgments.

Alibaba's Cloud for Youth Education Support Program has been operating for five years. As educational equity increasingly extends beyond access to resources to include access to capabilities, the program has continued to evolve. As of March 2026, Cloud for Youth had reached 28 provinces across China, established 403 AI cloud classrooms, and provided 380,000 personal cloud computer accounts for teachers and students.

Five participating schools have been selected by China's Ministry of Education as AI education bases for primary and secondary schools. Participating schools have also received honors from UNESCO and other organizations, including the Global Awards for World Heritage Education Innovative Cases (AWHEIC) and the Global Smart Education Innovation in Practice Prize.

One of the challenges of bringing AI into education is the need to transform teachers' roles, teaching processes, and assessment systems together. Currently, teachers often participate primarily as users of AI tools, without sustained mechanisms for engaging in curriculum design and practical innovation. Cloud for Youth is therefore shifting its focus toward strengthening teachers' capabilities.

In 2026, alongside the addition of 100 cloud classrooms, the program launched the "10,000 AI Pioneer Teachers Initiative" to support AI teaching practices nationwide and establish cross-regional communities for collaborative teaching research. It is also bringing leading Chinese schools into a co-development and resource-sharing model, exploring ways to provide more equitable starting conditions for AI education across different regions.

In classroom practice, AI is changing how learning takes place. Cloud for Youth's AI agents can analyze students' learning trajectories, skill profiles, and recurring errors, making previously invisible learning processes measurable and analyzable and providing data and technical support for personalized instruction. AI agents also offer students accessible interactive environments, encouraging them to ask questions and explore unfamiliar subjects. As answers become increasingly easy to obtain, the ability to ask meaningful questions is itself becoming an essential learning skill.

In curriculum development, Cloud for Youth is collaborating with institutions including Peking University's Graduate School of Education to investigate how students understand and evaluate AI, as well as cognitive biases, excessive reliance, and other issues that may arise during human-AI interaction. The findings are being used to inform improvements in curriculum design.

Yungu School Explores the Frontiers of AI Education

How AI Can Enhance Learning Rather Than Replace It

As AI becomes deeply integrated into education, educators must establish appropriate boundaries and divisions of responsibility between technology and the learning process. The essential questions are whether genuine learning has taken place after AI is introduced, whether students' own capabilities have become stronger, and whether they have developed the ability to understand others and engage with the real world.

Over the past year, Yungu School, founded with investment from Alibaba partners, has explored AI education in greater depth. In doing so, the school has increasingly recognized and sought to address four emerging challenges in primary and secondary education:

- Moving from measuring how much AI increases output to examining what students have genuinely learned, while guarding against the "illusion of competence".
- Moving beyond the pursuit of efficiency to preserving the necessary "inefficiencies" of reading, thinking, trial and error, and avoiding excessive "cognitive outsourcing".
- Moving from evaluating final outcomes to assessing the thinking process, with particular attention to whether students can transfer and apply their capabilities when AI is no longer available.
- Moving beyond discussions of how AI can strengthen knowledge and cognition to also addressing adolescents' social development, particularly their ability to understand real people and navigate, negotiate, and repair relationships.

Building on these considerations, Yungu School is gradually moving beyond fragmented experiments with AI tools toward a systematic redesign of teacher capabilities, student competencies, curricula, and the role of schools in the AI era.

First, the school is developing teachers' capacity for technological judgment. The guiding idea is that "teachers need to move fast", meaning they should quickly understand, experiment with, and identify the risks of new technologies, while assessing whether AI genuinely enhances learning. At the same time, "students need to move slowly", meaning children's developmental pace should not be dictated by the speed of model updates.

Second, the school is cultivating students' initiative in interacting with AI. Before using AI, students are



Yungu School's teachers have published two books: From Tool to Partner: Case Studies of AI in Teaching Practice and Unlocking K-12 STEM Education with AI.

encouraged to think through and break down tasks, decide which parts to delegate to AI and which to complete themselves, and continuously evaluate results, identify problems, and adjust their strategies. The goal is to transform AI from an "answer machine" into a "thinking partner".

Third, the school is redesigning learning tasks and assessments. Students are encouraged to use AI to obtain information, feedback, and learning opportunities, while preserving the essential processes of reading, thinking, writing, hands-on practice, failure, and revision. Assessment standards are being adjusted accordingly, moving beyond final outputs to include necessary evidence of the learning process.

Finally, the school positions itself as a "practice ground for real human relationships". The more virtual technology becomes, the more education must create opportunities for genuine human connection. The better AI understands children, the more important it becomes for children to learn to understand real people.

In 2026, Yungu School's secondary division established Zhejiang Province's first specialized class for AI Innnovation & Application Program. The program seeks to develop an AI talent cultivation framework over the three years of senior secondary education. Rather than starting with available courses, it works backward from the developmental goals of a 15-year-old student three years later, considering the scientific foundations, engineering skills, problem awareness, AI capabilities, and creative experience that the student should acquire.

Qwen App's Gaokao Admissions Agent

Bridging Gaps in Information Resources and Decision-Making Capabilities

Choosing universities and academic majors after the Gaokao, China's national college entrance examination, is a critical second stage in the distribution of educational opportunities and one of the most stressful moments for students and their families. Every year, more than 10 million candidates must make highly complex combinations of choices among nearly 3,000 higher education institutions and over 2,000 academic majors within a short period. Differences in family access to information and decision-making capabilities, which may have been partially mitigated during earlier education, can become magnified again at this stage.

In 2026, Qwen launched China's first full-cycle Gaokao university application planning agent, offering free university application planning and consultation services to candidates nationwide. Compared with Quark's Gaokao AI Model introduced in 2025, the new agent offers greater autonomy and long-term memory capabilities.

Performance Optimization for Low-End Devices, Qwen Gaokao Agent Launch Time Reduced by 50%

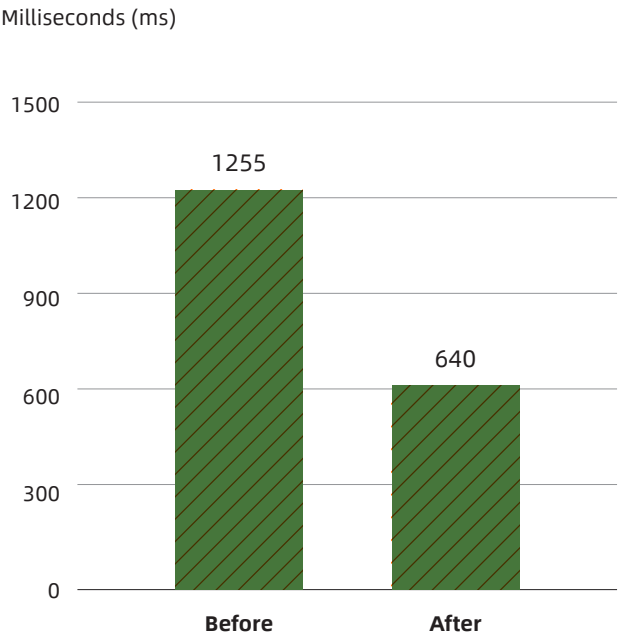

Users can interact with the agent throughout key stages, including post-examination planning, preparation for receiving examination scores, and the formal university application process. It provides continuous support as students make decisions about their educational futures. Comparable commercial services can cost several thousand yuan, while only 5% of candidates nationwide have engaged paid experts for application guidance.

During last year's implementation, the team found that 35% of users provided only four basic pieces of information: province, selected examination subjects, examination score, and ranking. Many lacked sufficient understanding of emerging industries and future career opportunities and were unable to formulate deeper follow-up questions. What many candidates lack is not simply information, but the ability to construct a decision-making framework and identify the questions that matter.

This year, the Qwen App's Gaokao Admissions Agent has begun playing a more substantial role in the application decision-making process. It has strengthened its dynamic adjustment and proactive recommendation capabilities, using multi-turn conversations to help candidates progressively refine their personal profiles, clarify their actual needs, and narrow their options. All recommendations retain traceable supporting evidence, and candidates remain responsible for making the final decisions.

In addition, the Qwen engineering team has carried out targeted optimizations for older devices and rural areas with weaker network connectivity. Working with the Warm Light Program, it is also seeking to narrow the information and resource gaps that candidates in counties and rural communities face outside the school environment.

Using an AI agent to assist with university applications can also help young people systematically understand themselves, academic disciplines, and careers, often for the first time. During the 2026 Gaokao season, users obtained more than 23 million free AI-generated university application planning reports through Qwen, nearly twice the 13 million reports provided last year.

University AI Support Program

Exploring Human-AI Collaboration Skills for the Future

AI is reshaping teaching and research in higher education, while also changing what employers expect universities to prepare students for in the AI era. Universities may become one of the first environments in which young people encounter and adapt to a technological revolution affecting knowledge-based work. However, access to frontier AI tools remains constrained by cost, language, and network availability, potentially widening differences in students' ability to use AI effectively.

In 2026, Alibaba Cloud and Alibaba Philanthropy launched the "University AI Support Program", providing university students with free access to two AI tools: QwenWork and Qoder. Together, these tools offer students AI-powered productivity capabilities for the future of knowledge work.

In its first year, the program plans to reach more than 200,000 undergraduate, master's, and doctoral students across nearly 10 universities participating in China's Double First-Class Initiative, with a target of 100,000 activated user seats.

On September 19, the program was first launched at the University of Science and Technology of China. At the launch event, "QwenWork Campus Ambassadors" and "Qoder Exploration Ambassadors" were appointed to support more students during implementation.



Appointment Ceremony for QwenWork Campus Ambassadors and Qoder Exploration Ambassadors

On September 24, the program was introduced at Shanghai Jiao Tong University. Alibaba provided all registered faculty members and students with one year of free access and corresponding computing quotas, together with product training, technical support, and campus activities. As of the publication of this report, access was being progressively extended to additional universities.

From opening access to tools, to supporting practical experimentation, and then to sharing methods, the University AI for Good Support Program seeks to help students become not only users of AI tools but also creators. Its long-term goal is to enable university students to do more than simply complete additional tasks with AI. By preserving independent thinking and judgment, students can gradually develop the human-AI collaboration capabilities needed for future knowledge work.

RynnBot Platform

Developing Multidisciplinary Talent at the Frontiers of AI

Embodied intelligence is a frontier field where AI and engineering are becoming deeply integrated. Its development involves not only robotic hardware, but also data collection, model training, simulation-based validation, and deployment on physical robots. The technical barriers and hardware investment requirements are substantial, making it difficult for university students to gain practical experience across the complete development process.



RynnBot SparkEdu Campus Outreach Program

To lower these barriers, DAMO Academy launched the RynnBot embodied AI platform in 2025 and open-sourced the Rynn Robot Context Protocol (RynnRCP). In February 2026, it further open-sourced RynnBrain, a foundation model for embodied intelligence. Through coordinated edge-cloud computing, developers can run mainstream embodied AI models without configuring GPUs on local robotic devices.

As of July 2026, the RynnBot SparkEdu training program, designed for educational applications, had established partnerships with 18 universities in China and abroad, expanding its reach to Hong Kong SAR, Singapore, Pakistan, and other locations. It had brought together 17,000 embodied AI developers and supported the creation of more than 100 innovative projects.

Students gain practical experience throughout the entire process, from teleoperation data collection and model training to deployment on physical robots. This provides a practical pathway for universities to cultivate multidisciplinary talent equipped to meet the needs of future industries.

Alibaba Cloud Science Education Base

Using AI to Explain AI

The value of public AI education lies not only in lowering barriers to using AI, but more importantly in giving people of different ages, professions, and regions the opportunity to understand and learn about emerging technologies.

The Alibaba Cloud Science Education Base is an immersive AI exhibition space covering more than 1,000 square meters. Employee volunteers and volunteers from partner universities form the backbone of its public education efforts, offering free exhibition tours, science lectures, AI workshops, and other activities. Through visible, hands-on, and practical experiences, the base helps the public develop a better understanding of AI.

Volunteers design different learning experiences for different audiences:

- For primary and secondary school students in urban areas and their families, free educational visits are designed to inspire curiosity and creativity.

- For young people in remote regions, targeted invitations and special sessions provide opportunities to experience AI firsthand.
- For visually impaired people and others with disabilities, accessible AI interaction experiences are explored to demonstrate the potential of technology to support people with disabilities.
- For technology enthusiasts, makers, and university developers, AI workshops and open experiments provide opportunities for deeper technical engagement.

The base embraces the principle of "using AI to explain AI" throughout its educational activities. Large models assist in generating exhibition scripts, preparing teaching materials, and designing interactive experiences, improving both educational efficiency and visitor engagement.

As of September 2026, the base had more than 80 science education volunteers, welcomed over 100,000 visitors, and organized more than 100 free public education activities. Visitors ranged in age from 6 to 78.

The base has been recognized as a Zhejiang Provincial Science Education Base and a Zhejiang Provincial Science Popularization Experience Center for New Quality Productive Forces.

Looking Ahead

A comprehensive evaluation framework for AI in education has yet to be established. The central question such a framework must answer is what capabilities students should develop. Evaluations should determine whether AI genuinely strengthens students' capacity to learn, rather than focusing only on task completion efficiency or improvements in examination scores. Otherwise, education risks drifting away from its fundamental purpose of cultivating independent thinking and lifelong learning.

An evaluation framework determines what should be assessed, but successful implementation ultimately depends on teachers. Technological development is currently moving far faster than traditional teacher training systems can adapt, making teachers one of the groups facing the greatest pressure in the AI-driven transformation of education.

In the future, while continuing to strengthen their professional expertise, teachers will need not only to master AI tools but also to organize learning processes and guide human-AI collaboration. AI for education therefore requires sustainable and scalable mechanisms for teacher development.

AI for Science

Luca Family of Models
Lingshu-Cell and ReproPerturb Models
ElementsClaw

Scientific problems are becoming increasingly complex, with research moving beyond individual disciplines toward interdisciplinary, multiscale, and multimodal approaches. The volume of knowledge and data that researchers must process is growing exponentially. Globally, despite continued increases in research investment, gains in innovation efficiency have slowed in some fields. China is entering a new stage characterized by more frontier scientific discovery and more fundamental technological research and development, creating an urgent need to explore new research paradigms.

AI is reshaping the process of scientific discovery, lowering barriers to interdisciplinary research and enabling researchers to integrate knowledge and data across different fields more efficiently, accelerating the journey from hypothesis to validation. AI is already making inroads into life sciences, materials science, new energy, climate change, and other fields. Alongside training specialized models for different scientific disciplines, research institutions are developing AI agents for scientific discovery, allowing experimental feedback to inform subsequent research and even drive the iterative improvement of the models themselves. Achieving major breakthroughs in scientific discovery through AI is one of the most promising areas of AI for Good.

Luca Family of Models

Becoming Part of Scientific Discovery Workflows

AI for Science is entering a new stage. General-purpose large models, scientific AI agents, automated code execution, and automated experimental validation are working together with scientists to shape a new paradigm of scientific discovery. Within this emerging paradigm, specialized AI models such as AlphaFold remain indispensable. They will serve as domain-specific capability modules that general-purpose scientific agents can invoke.

In 2026, Alibaba's Qwen-3.8-Max demonstrated the ability to reconstruct from scratch the complete training and evaluation pipeline described in a research paper, reproduce its six principal findings, and independently

Luca Series of Specialized Domain Models Built on LucaOne

LucaProt
RNA Virus Identification Model
Cell, 2024

LucaPCycle
Phosphorus-Cycling Protein Family Discovery Model
Nature Communications, 2025

LucaVirus
Multimodal Viral Foundation Model,
National Science Review, 2026

LucaCell
Single-Cell Foundation Model
GitHub Open Source, 2026

propose and test multiple directions for improvement, much like a human researcher. Ultimately, it identified a new approach that outperformed the method presented in the original paper. This capability builds on the model's understanding of scientific objectives, semantic reasoning across disciplinary knowledge, and an agent architecture capable of decomposing research tasks and coordinating tools.

Alibaba's biological foundation model LucaOne is a specialized tool that can be invoked within this new scientific discovery system. Its training data spans 169,000 species, comprising approximately 1.2 billion nucleic acid sequences and 600 million protein sequences. It is a biological language model built around a unified representation of the central dogma of molecular biology. As of June 2026, LucaOne was receiving more than 5,000 downloads per month on Hugging Face, demonstrating strong activity among models of its kind.

Biological sequences constitute a common language of life. Building on LucaOne's general-purpose representation capabilities, the team has further developed a family of specialized Luca models for different domains of this "language of life". These specialized models are beginning to be orchestrated and invoked in combination with general-purpose large models such as Qwen, becoming integrated into scientific discovery workflows.

In September 2026, LucaCell, the latest member of the Luca family, was open-sourced. Also built on LucaOne's sequence representation capabilities, LucaCell is a cross-species, cross-modal foundation model for single-cell biology. It was pretrained on 85 million human and mouse single-cell transcriptomic profiles, covering 62 tissue sources and 18 disease states. LucaCell has demonstrated strong generalization capabilities in tasks including cross-species cell identification, gene expression reconstruction, and viral load prediction. It has the potential to provide a foundation for future research into interactions between microorganisms and their hosts, as well as the cellular mechanisms underlying disease.

Building on this foundation, Alibaba continues to expand and refine its specialized scientific tools in response to frontier research needs. The next-generation LucaGen will combine biological sequence understanding and generation within a unified framework, moving beyond supporting scientific discovery toward enabling scientific creation. LucaProtGen, meanwhile, is being developed to enable the design of biological proteins through natural-language instructions.

Lingshu-Cell and ReproPerturb Models

Moving Life Sciences from Static Analysis to Dynamic Simulation

AI is transforming the discovery and design of drug molecules. Yet the real bottleneck in drug development often lies not in molecular design itself, but in understanding disease. This requires establishing connections between patients' clinical manifestations and the mechanisms through which drugs act at the cellular level.

Most existing AI models focus on static molecular representations or localized interactions. They struggle to simulate the overall responses of cells to genetic or environmental perturbations across broader spatial and temporal scales. Meanwhile, even after new molecules or targets have been identified, preclinical research remains heavily dependent on animal experiments. These experiments are time-consuming, costly, and limited in throughput, creating a critical bottleneck in the advancement of drug development pipelines.

To overcome these constraints, research institutions worldwide have accelerated their exploration of virtual cells since 2025. DAMO Academy has also been working in this direction. In collaboration with Westlake University and other universities and research institutions, it has released Lingshu-Cell, a generative cellular world model, and ReproPerturb, a stem cell reprogramming prediction model. Together, these efforts are progressively establishing a capability framework that extends from simulating cellular states to predicting cell fate.

Lingshu-Cell can generate cellular transcriptomic states under different genetic perturbations and drug treatments, marking a shift from static cell atlases toward dynamic cellular world models. ReproPerturb goes a step further by integrating multimodal

ReproPerturb Predicts Optimal Conditions for Generating High-Quality Hypoblast-Like Stem Cells

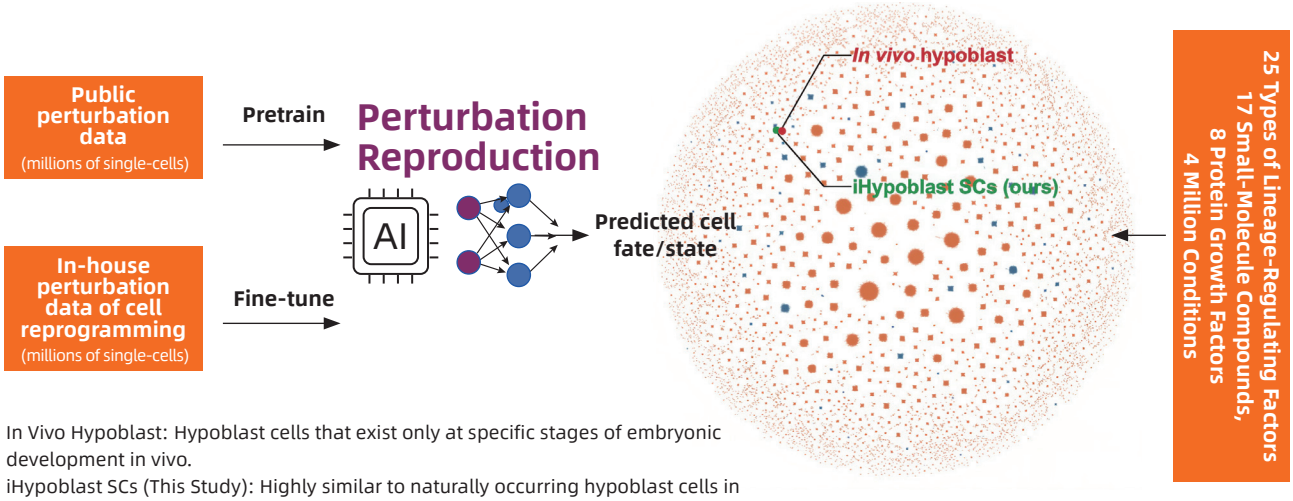

In Vivo Hypoblast: Hypoblast cells that exist only at specific stages of embryonic development in vivo.
iHypoblast SCs (This Study): Highly similar to naturally occurring hypoblast cells in their molecular characteristics, with stable expression of key pluripotency factors. They retain stem cell properties even after 50 passages in vitro.

information, including small molecules and growth factors, to predict how different combinations of perturbations affect the direction of stem cell reprogramming. Through interpretable mechanisms, it also reveals potential biological pathways, providing a basis for experimental design.

As virtual cells become increasingly capable of approximating the behavior of real cells, many experiments in drug screening, disease mechanism research, and regenerative medicine may eventually be conducted first in digital environments before proceeding to laboratory validation. This could not only shorten drug development cycles and reduce research costs, but also help life sciences gradually establish a new research paradigm built around "AI simulation, experimental validation, and model iteration".

ElementsClaw: An AI Agent for Superconducting Materials Discovery

Participating in the Organization of Scientific Discovery

AI for scientific research is moving from a predictive paradigm toward a generative one. Yet both approaches face a common decision-making bottleneck: among millions of candidate materials, which are truly worth synthesizing in the laboratory? Making this decision requires coordinating multiple dimensions, from atomic-scale numerical calculations of structure and stability to higher-level semantic reasoning about novelty and synthesis feasibility.

AI agents offer a promising pathway for moving materials discovery beyond isolated processes toward an integrated framework. In July 2026, DAMO Academy, together with Renmin University of China, the University of Chinese Academy of Sciences, and other institutions, introduced ElementsClaw, the industry's first AI agent for superconducting materials discovery, demonstrating a complete scientific decision-making process.

Elements is a large-scale atomistic model with 1 billion parameters. ElementsClaw serves as the decision-making hub for the entire research workflow, orchestrating tasks,

invoking tools, and coordinating four specialized skills to make a sequence of scientific decisions.

In just 28 GPU-hours, the agent screened 2.4 million stable crystal structures and identified 68,000 candidate superconducting materials. Four materials were ultimately synthesized experimentally, and all four were confirmed to exhibit superconductivity. This scientific agent framework can be applied not only to superconducting materials discovery but also to research and development in other materials fields, including solid-state battery electrolytes, heterogeneous catalysts, and thermoelectric materials.

More importantly, the four newly discovered superconductors did not emerge from a single discovery mechanism. Instead, they represent four distinct pathways through which a scientific agent can conduct scientific exploration.

As scientific agents progressively take on tasks across literature understanding, hypothesis generation, experimental design, tool development, and experimental validation, and become deeply integrated with automated experimental platforms, scientific research may evolve toward an end-to-end, autonomous discovery loop. The value of AI agents could then extend beyond improving the efficiency of individual research steps to becoming engines of innovation that continuously generate new knowledge.

ElementsClaw Demonstrates Four Pathways to Materials Discovery

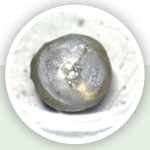


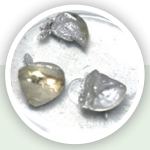

Discovery Pathway	Motif-Guided	De Novo Generated	Structurally Reinterpreted	Database-Latent
Relationship to Existing Knowledge	• Builds on existing structural motifs • Extends knowledge through inheritance from known structures	• Completely independent of existing samples • Generates previously unknown structures	• Revises existing knowledge • Rediscovers long-overlooked materials	• Uncovers latent patterns not yet explicitly identified in existing knowledge • extrapolates based on these patterns
Discovered Materials	Zr_3ScRe_8	$HfZrRe_4$	Zr_4VRe_7	$Hf_{21}Re_{25}$

Looking Ahead

The expansion of scientific AI agents' capabilities depends on continuous feedback and iteration during reinforcement learning, while the efficiency of real-world experimental validation determines the speed of this feedback loop. Compared with tasks in the digital domain, wet-lab experiments remain time-consuming, costly, and relatively limited in automation, making them a major bottleneck in scaling scientific AI agents.

Looking ahead, advances in automated experimental platforms, digital twins, and related technologies could further shorten end-to-end AI research cycles, enabling more frequent and more autonomous research iterations. At the same time, human scientists and AI researchers will increasingly define research problems together, driving further changes in the methodology of scientific research.

AI for Green Development

Xianyu's Circular Economy
High-Resolution Wind and Solar Energy Distribution Map
"Small Ideas, Big Impact" Green Initiatives
The GreenNet Program for Wildlife Protection

Global climate change, biodiversity loss, and pollution constitute three interconnected environmental crises. Addressing these challenges requires not only reducing emissions, but also rethinking how resources are produced, used, and disposed of, while protecting ecosystems and the life they sustain.

AI offers new tools for the green transition. By identifying patterns in complex data, improving forecasting and matching, and supporting monitoring and decision-making, AI has the potential to help people understand environmental changes more accurately, use energy and material resources more efficiently, and detect ecological threats earlier. At the same time, the environmental value of technology does not arise automatically. AI must both manage its own resource consumption and work in conjunction with domain expertise, governance mechanisms, and practical action to translate its analytical capabilities into verifiable environmental improvements.

The four approaches presented here, from extending the useful life of everyday goods and mapping renewable energy resources to protecting wildlife and mobilizing developers through the "Small Ideas, Big Impact" initiative, explore a common question: How can AI capabilities address specific environmental needs while enabling more people and organizations to participate in the green transition?

Xianyu's Circular Economy

Using AI to Help Unused Goods Circulate Again

Helping an item that is still usable find its next owner is one of the most accessible ways for people to participate in the circular economy. By extending product lifespans and, to some extent, reducing purchases of new goods, the reuse of idle items can help reduce demand for new resources and the generation of waste. Yet moving an item from "sitting unused at home" to being used again often involves several steps, including organizing product information, finding buyers, estimating prices, and communicating during transactions.

Digital platforms provide the infrastructure for connecting these fragmented sources of supply and demand. A study of 246 circular economy

Carbon Emissions Avoided Through the Reuse and Recycling of Idle Items

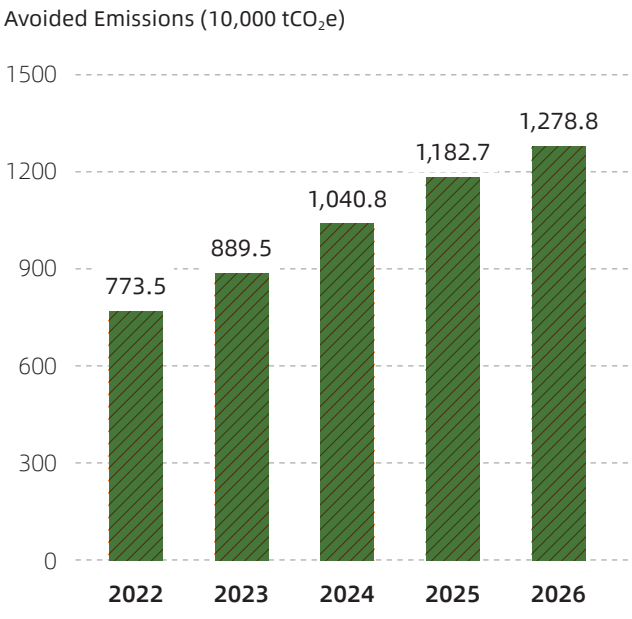

platforms worldwide found that platforms support the recirculation of used goods, repairable products, and materials by facilitating search and matching, reducing transaction costs, and expanding participation networks. The study identified Xianyu as an example of a cross-category circular economy platform.

As one of China's leading secondhand trading platforms and lifestyle communities, Xianyu sees its role as going beyond providing a place for users to list unused items. It also seeks to establish information and transaction connections among individuals, enabling products across different categories and in varying conditions to continue delivering value.

A major obstacle to the circulation of secondhand goods is the information friction created by their non-standardized nature. Even products of the same model can differ in price because of their condition, usage history, or accessories. Sellers may not know how to describe their items, while buyers often need to ask repeated questions. These information gaps leave many useful products sitting idle for extended periods, and some are discarded before finding another user.

Reducing these costs is particularly important for lower-value secondhand items that still require time and effort to organize, describe, and negotiate. AI can help address these frictions by improving product information, supply-demand matching, and transaction responsiveness, making it easier for more usable goods to find suitable new owners.

In 2026, Xianyu deployed AI agents on both the supply and demand sides of its marketplace, exploring how fragmented, non-standardized transactions could be transformed into service processes supported by continuous AI assistance. The goal is to make it easier for users to list, discover, and trade secondhand goods.

On the supply side, the seller agent Xianyu Selling Assistant uses multimodal capabilities to interpret images and text, helping identify product categories, brands, and models, generate titles and descriptions, and prompt sellers to provide additional information about condition, repair history, and accessories for review and completion. Drawing on similar products, historical transaction prices, and market supply and demand, its valuation function provides pricing references, reducing the burden of not knowing how to describe an item or how much to charge. After listing, AI-assisted management can help handle repetitive inquiries and fulfillment reminders, reducing the time sellers need to spend following up on transactions.

On the demand side, the buyer agent Xianyu Buying Assistant seeks to translate buyers' everyday, often vague expressions into specific product search criteria. It interprets requirements such as intended use, budget, model, and condition, asks clarifying questions when necessary, and selects potential matches from a large inventory of non-standardized goods. In this way, AI extends the platform's role in connecting supply and demand to organizing information and facilitating communication, helping users complete transactions that would otherwise be relatively cumbersome.

As of October 2026, Xianyu's seller-side AI agent had achieved a 55% adoption rate for its recommendations, while the proportion of listed products successfully sold increased by 6.5%.

According to a third-party audit, during fiscal year 2026 (April 2025 to March 2026), Xianyu users avoided an estimated 12.788 million metric tons of emissions through secondhand transactions and recycling, based on the platform's existing accounting methodology. This figure reflects the overall impact of circular economy activities on the platform. The specific contribution of AI cannot yet be separately quantified.

As AI becomes increasingly integrated into the platform's services, we hope to make continued use of existing goods a more convenient and trustworthy choice. If secondhand consumption depends solely on people's willingness to conserve resources, barriers to participation remain relatively high. When the transaction experience itself becomes sufficiently seamless, more people can make this choice naturally in their everyday lives.

High-Resolution Wind and Solar Energy Distribution Map

Advancing Green Energy Development Through AI

As wind and solar power expand, coordination between these renewable energy sources becomes increasingly important. Wind and sunlight are not always available at the same time or in the same locations. Cross-regional scheduling can help smooth fluctuations in generation and increase the utilization of renewable electricity. The need to achieve compute-power coordination for gigawatt-scale AI computing centers makes this challenge even more pressing. However, effective cross-regional scheduling requires a high-resolution map of actual renewable energy facilities across the country.

Producing such a map at a national scale is no easy task. China's landscapes vary enormously across regions, from deserts, grasslands, and snow-covered terrain to forests and bodies of water, placing high demands on a model's ability to generalize.

Using a proprietary deep learning model and 7.56 TB of satellite remote sensing imagery covering China at a resolution of 0.5 meters, Alibaba's DAMO Academy mapped, for the first time at high precision, the spatial distribution of 319,000 photovoltaic installations and 91,600 wind turbines nationwide. The research was published in Nature in May 2026.

This work transforms geographically dispersed wind and solar resources from aggregate statistical data into green energy resources that can be analyzed computationally, scheduled, and incorporated into planning. The research found that, during the study period, using data through 2022, cross-provincial coordination could enable the additional utilization of approximately 100 billion kilowatt-hours of renewable electricity, equivalent to around 9.1% of China's total wind and solar electricity generation.

Beyond supporting the real-time scheduling of renewable electricity, this digital base map also provides a foundation for early-stage resource planning, including energy storage deployment and the selection of locations for computing centers.

"Small Ideas, Big Impact" Green Initiatives

Exploring Ways to Coexist with Nature

The 2026 "Small Ideas, Big Impact" AI for Good Innovation Challenge is a public innovation competition jointly launched by Alibaba's ATH Business Group, Alibaba Philanthropy, and Alibaba Cloud Model Studio. It invites participants to develop AI solutions to social and environmental challenges. Now in its second year, the competition introduced a new Green Development track as a major focus.

The second edition attracted 1,111 participating teams and received 558 AI-powered submissions. Building on the first edition's focus on disability inclusion, inclusive elderly care, rural education, and other areas, this year's competition added green development under the theme of "Coexisting with Nature".

Several participating AI applications focused on recognizing natural sounds. Some transformed animal calls into music, others used sound to reveal patterns in bird activity, while still others supported nature field studies conducted by parents and children. Behind these applications were models from Alibaba's open-source ecosystem, including the multimodal Qwen3.5-Omni and the audio generation model ThinkSound. The training dataset for ThinkSound was curated from 2,532 hours of high-quality audio samples, including animal vocalizations, mechanical sounds, and environmental audio.

In addition, three "Zoo Companion Agents" reached the finals. These agents help make educational content scattered across zoo social media accounts, animal care records, and interpretive scripts accessible beyond zoo grounds, transforming it into portable resources for nature education.

The competition also provided participants with 100 million free tokens, offering foundational support for the development and experimentation of these AI applications.

The GreenNet Program for Wildlife Protection

Turning Searches for Prohibited Wildlife Products into Opportunities for Nature Education

The global illegal wildlife trade is estimated to be worth as much as US$23 billion annually, with transactions increasingly moving online. Sellers often use coded language, altered terms, and multiple accounts to conceal their trading intentions and activities.

Consider a search for "ivory". What appears to be an innocent query may be connected to illegal poaching and wildlife trafficking networks, while consumers may have little understanding of the ecological consequences. Platforms operate across search, recommendation, risk control, and transactions, giving them the ability to intervene throughout the process in ways that individual consumers cannot.

The GreenNet Program was launched in 2020, initially focusing on interventions targeting suspicious searches.

To date, the GreenNet Program has expanded its coverage to more than 110,000 high-risk blocked search terms and carried out over 200 million interventions targeting suspicious searches. Of these, more than 25 million interventions have involved blocking or redirecting suspicious searches related to wildlife and plants.

In 2026, the GreenNet Program expanded from individual search interventions toward comprehensive platform governance. Supported by AI capabilities, its traditional approach of "identification, blocking, and intervention" has broadened to include "assistance, education, and collaboration", extending its coverage to protecting lives, ecological conservation, child and youth safety, and public education on internet-related laws.

In June 2026, the initiative was invited to participate in the United for Wildlife Business Forum organized by The Royal Foundation, where it shared the governance practices of Chinese internet platforms in combating illegal wildlife trade.

When users search for sensitive content involving endangered species on platforms such as Taobao, the system no longer simply returns a "page not found" message. Instead, it provides explanations of relevant regulations, information about species, and conservation messages, transforming a potentially prohibited search into an opportunity for nature education. If consumers are simply looking for decorative crafts or collectibles, the platform can also recommend suitable alternatives.

The GreenNet Program: Four Key Areas of Protection

Life Protection

- Suicide Prevention and Intervention
- Mood Guardians Program

Ecological Protection

- Animal Waystation
- One Mu of Land

Youth Development and Protection

- Green Orange Classroom
- Minor Protection Center

Public Legal Education

- Interventions Targeting Suspicious Searches
- Consumer Product Alerts
- Online Legal Education

AI expands the reach of these governance efforts, particularly by understanding more complex search expressions and identifying altered terms, coded language, and potentially risky information embedded in images and videos. AI is also used to generate initial drafts of legal education content and page templates, allowing the platform to produce species-specific educational materials more efficiently. The GreenNet Program has also made these AIGC capabilities available to the public, encouraging people to participate actively in content creation.

Building on the continued validation of these capabilities, the GreenNet Program launched "the One Mu of Land Initiative" in 2026. The initiative focuses on collaborative governance against illegal wildlife trade and alternative approaches to ecological restoration, extending online governance capabilities into ecological restoration, nature education, and public participation.

Looking Ahead

AI is both an additional source of energy demand and an important tool for the green transition. How to balance the growth in AI's energy consumption against the emissions reductions it helps enable is a fundamental question this generation of technology must answer. Accelerating the development of a new power system and making the computing revolution part of the energy revolution will be among the critical technological levers.

Over the longer term, AI is changing how green development can be achieved. It is making resources and ecosystems that were previously difficult to observe increasingly measurable, computable, manageable, and responsive to coordinated planning. It is also enabling more individuals and organizations to participate in these efforts.

Only when technological progress and ecological responsibility advance in the same direction can green development gain lasting momentum.

AI for Economic Opportunities

ModelScope Community
Y/OUR SPACE
Accio: A Global E-commerce AI Workspace
AI-empowered Rural Revitalization Commissioner

The United Nations identifies sustained, inclusive, and sustainable economic growth, full and productive employment, and decent work for all as key objectives of the 2030 Agenda for Sustainable Development. As a general-purpose technology, AI is rapidly entering industries across the economy. Expanding economic opportunities and promoting inclusive growth are fundamental goals of making technology accessible to all.

The rapid advancement of AI and its increasingly affordable diffusion are lowering barriers to innovation and entrepreneurship. More individuals, small and medium-sized enterprises (SMEs), and traditional industries can now participate in market competition. With smaller teams, they can also access a wide range of tools and capabilities to complete complex tasks that previously required substantial resources.

ModelScope Community

Enabling Innovation by Those Closest to the Problems

AI is expanding the community of innovators. Open-source models, low-code development tools, and affordable computing resources are continuously lowering barriers to innovation, enabling people who understand real-world problems best to develop solutions themselves, often for the first time. Frontline practitioners across different industries, as well as people with disabilities and others who have traditionally had limited access to technology development, are moving from being recipients of innovation to creators of solutions.

As China's largest AI open-source community, ModelScope is becoming an important part of the infrastructure supporting this transformation. As of September 2026, the community had brought together 33 million developers and made more than 250,000 open-source models and 40,000 datasets freely available to developers worldwide. The Qwen family of models supports 201 languages and dialects, helping developers from different cultural backgrounds build and deploy localized applications.

In 2026, ModelScope continued evolving into an infrastructure platform for AI agent development. The community introduced a Skills Center, where developers can package complex capabilities into reusable skill modules, with more than 80,000 skills already available. It also launched an MCP Marketplace, bringing together more than 12,000 MCP services and offering free hosting. At the same time, the community opened up core interfaces for models, datasets, and other resources. Together, these capabilities are reducing the barriers to moving from invoking models to building AI agents.

AI is also changing how knowledge and experience create value. Industry expertise that previously existed primarily within individuals can increasingly be transformed into reusable, continuously improvable, and even tradable AI agents or skills. According to the community's report on AI and open-source talent development, 28% of developers do not have a

professional AI background but come from industries such as healthcare, education, and manufacturing. An increasing number of people closest to real-world problems are gaining the opportunity to participate directly in innovation.

This shift is particularly evident in philanthropic initiatives emerging from the community. During the inaugural "Small Ideas, Big Impact" AI for Good Innovation Challenge in 2025, co-organized by ModelScope, a developer with a visual impairment independently trained a traffic light recognition model designed to address the mobility needs of people with visual impairments that general-purpose vision models struggled to meet. Another developer with dysarthria created Ranyan, an application built on Qwen's speech models, and used AI to deliver his first-ever public speech. In both cases, identifying a need, developing a capability, and validating it in the real world formed a complete feedback loop.

ModelScope provides the environment in which such feedback loops can continue to operate. By giving developers access to open models, toolchains, and free GPU computing resources, the community enables them to explore new ideas and move from concept to initial implementation at relatively low cost. Moreover, 96% of developers in the community have explicitly expressed a willingness to open-source their derivative work, allowing these innovations to be reproduced, improved, and disseminated more widely.

From schools to communities and wider society, some projects are gradually achieving commercial adoption or scaling up as philanthropic initiatives. Many others continue to contribute reusable models, workflows, AI agents, and AI-native talent, laying the groundwork for the next generation of innovation.

What is your current field of study/work? (Single choice)

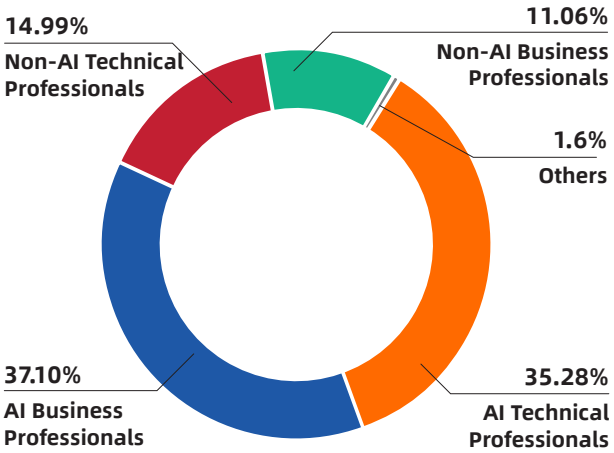

Source: 2026 AI Open Source Talent Report, p. 16

Y/OUR SPACE

An Offline Community for Super Individuals

The smallest viable unit of entrepreneurship is changing in the AI era. Large language models have substantially lowered the technical barriers to starting a business, giving rise to new entrepreneurial forms such as "super individuals" and one-person companies (OPCs). These entrepreneurs represent not only new forms of flexible employment but also some of the most active participants in social innovation.

Yet during a period of rapid technological change, entrepreneurs operating as individuals have an even greater need to rebuild connections among people and businesses. They need opportunities to exchange ideas and learn about the latest AI technologies, as well as access to real business scenarios, industry resources, and collaborative partners to translate AI capabilities into viable entrepreneurial ventures. Much of this connection-building still takes place offline.

To address these needs, Alibaba's Real Estate Management Department launched Y/OUR SPACE, an entrepreneurial community, in May 2025. Rather than simply operating a workspace rental business, the initiative is designed to support microenterprises and OPCs.

Beyond shared office facilities and one-stop entrepreneurial services, Y/OUR SPACE seeks to serve as an "interface" between individual entrepreneurs and the broader industrial ecosystem. It connects ordinary entrepreneurs with Alibaba's AI capabilities, including large models and computing resources, as well as industry resources such as e-commerce. It brings in ModelScope and external communities to organize technical exchanges and industry networking



Scenes from the Y/OUR SPACE Event

events. It also connects entrepreneurs with local government support policies for OPCs, helping facilitate subsidies for computing power, data, application scenarios, tokens, and other resources.

Through these efforts, Y/OUR SPACE aims to help emerging entrepreneurial ventures create social value while building and expanding an ecosystem of shared prosperity.

Currently, approximately 60% of Y/OUR SPACE tenants are former employees of internet companies, with typical teams consisting of two to four people. Shared offices are therefore evolving beyond conventional coworking spaces. They are helping these super individuals rebuild the networks of connections, resources, and collaboration that larger organizations traditionally provided.

Entrepreneurs from Y/OUR SPACE are already beginning to gain recognition. In September, at the inaugural Global Open-source AI Challenge (GOAI), two entrepreneurs from the Hangzhou Yungu location stood out among more than 14,000 developers from 91 countries and regions. Their projects, "Quality Chain Agent" and "Shanmao's Journey", won first place in the Borderless Applications track and second place in the Embodied Future track, respectively.

As of September 2026, beginning with its first location in Xixi, Hangzhou, Y/OUR SPACE had expanded to 31 locations across 24 cities in China, serving more than 4,000 entrepreneurs. Occupancy rates in established communities remained above 90%, with demand at some locations exceeding available space.

Among these locations, the ModelScope Hangzhou AI Developer Center (Yungu YS), jointly established with the Xihu District Government, has become a representative community within Hangzhou's AI open-source ecosystem.

Accio: A Global E-commerce AI Workspace

Lowering Barriers for SMEs to Participate in Global Trade

Entering global trade has traditionally involved substantial barriers. From market research and product sourcing to merchandise operations, cross-border marketing, and customer service, the entire process requires coordination among specialized teams. SMEs and individual entrepreneurs often find it difficult to manage these activities independently. As the global trading system becomes increasingly fragmented, smaller orders, more frequent transactions, and rapid responses are becoming the new competitive norm, further raising the barriers to cross-border business.

In 2026, Alibaba International introduced Accio, a global e-commerce AI workspace designed for international trade. It connects multiple Chinese e-commerce platforms, including Taobao, Tmall, 1688, Douyin, Pinduoduo, and JD.com, as well as international platforms such as Amazon, eBay, Shopify, Temu, and TikTok.

Drawing on 26 years of experience in international trade, Accio translates critical cross-border business processes into more than 100 callable skills. Through multi-agent collaboration, it can execute

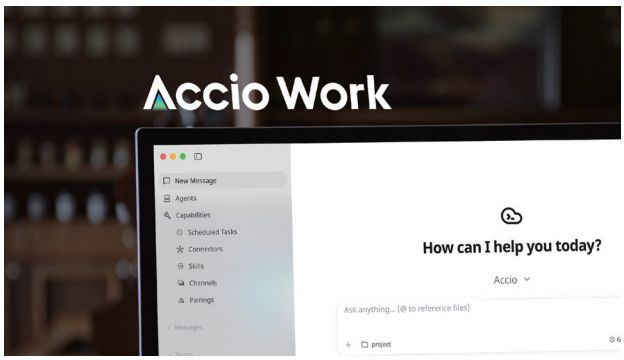

complex business tasks that previously depended on specialized organizational roles and divisions of labor. Alibaba International has also open-sourced CommerceAgentBench, a benchmark designed to help the industry evaluate agent performance in real-world commercial scenarios.

These capabilities substantially lower the barriers to participating in global trade. Previously, setting up a cross-border online store could take several weeks and cost tens of thousands of yuan. Today, the process can be completed in as little as 30 minutes, with token costs of approximately RMB 70. For overseas buyers, the time required to move from product search to purchase intent has fallen by 80%, while the number of new overseas buyers on Alibaba.com has increased by 40% year over year.

As of the end of July 2026, 10 million SMEs worldwide had used Accio to conduct business. Compared with the same period last year, token consumption by SMEs on Alibaba.com had increased tenfold, while Accio's token costs were 50% lower than those of general-purpose AI agents.

As commercial capabilities are increasingly packaged into skills and AI agents, these capabilities, once tied to individuals and organizational structures, may become fundamental units of collaboration and transactions among agents. In the future, more individuals will be able to organize and invoke AI agents around their own core strengths, moving global trade toward a new stage of agent-driven collaboration.

AI-empowered Rural Revitalization Commissioner

Organizational Scaffolding for Grassroots AI Adoption

The greatest challenge in bringing AI to grassroots communities often lies not in the technology itself, but in understanding local needs, integrating AI into organizational workflows, and establishing mechanisms for sustainable service delivery. This requires not only AI to adapt to grassroots organizations, but also these organizations to develop new capabilities for the AI era.

Since 2019, Alibaba has assigned rural revitalization commissioners to less-developed counties to help local communities connect with industry and technology resources. Building on this mechanism, Alibaba launched the "AI Rural Revitalization Commissioner" initiative in October 2025, extending its support from industrial development to the deployment of AI capabilities at the grassroots level.

These commissioners translate fragmented local challenges into product requirements, transform tacit knowledge accumulated by communities over many years into digital capabilities that AI models can understand, and turn dispersed AI capabilities into public services that can operate sustainably. This has become an important pathway for Alibaba's AI capabilities to reach grassroots communities.

Shengsi, Zhejiang Province, offers an example. Its ferry ticketing system, route regulations, and changing weather conditions are complex, and navigating them has long depended on local experience. Visitors face different ferry options depending on whether they are traveling by car and which island they intend to visit, making information difficult and costly to obtain.

The local commissioner brought together this tacit knowledge and existing data systems to create an AI Travel Assistant, making this locally developed expertise available as a public service that all visitors can access. Similar initiatives are being replicated across more counties, helping AI better understand local expressions, grassroots workflows, and real-world circumstances.

As of September 2026, the "AI + County" intelligent

service platform had been launched in 15 regions, covering more than ten grassroots service scenarios, including policy consultation and the business environment, enterprise administrative procedures, and service requests handled through the 12345 government service hotline.

The platform has reduced the workload associated with repetitive inquiries by half, while complaints have declined significantly. This has also freed frontline staff to concentrate on more complex tasks, such as coordinating services across government departments.

AI Rural Revitalization Commissioners function as a form of "organizational scaffolding", helping grassroots organizations gradually establish mechanisms for continuously absorbing, accumulating, and applying AI capabilities. This, in turn, stimulates local communities' own motivation to embrace AI.

In Shengsi, 30 local organizations have developed 60 application scenarios. When these capabilities become embedded within local institutions and remain available to them, AI for Good is no longer simply a one-way transfer of technology into society.

Work Order Processing

Processing Time	Number of Complaints	Repeat Inquiries
Reduced to 1–2 Days	Down 30%	Down 50%

Note: The figures are drawn from selected AI-powered county-level service platforms and do not represent comprehensive statistics. They are intended to illustrate the outcomes of these applications.

Looking Ahead

While AI is expanding economic opportunities, it is also profoundly changing the nature of work and the structure of employment. As AI agents gradually take over large numbers of repetitive tasks, a fundamental question for this generation of technology is how to translate efficiency gains into broader opportunities for development rather than deepen divisions in employment and skills.

The key lies in whether businesses can move beyond simply deploying AI tools toward redesigning their organizations, workflows, and talent development systems. By doing so, they can help employees shift from repetitive work toward creativity, collaboration, and complex decision-making, continuously building the human capital needed in the AI era.

Over the longer term, AI is lowering the barriers to innovation, entrepreneurship, and global trade to unprecedented levels, giving more individuals and smaller businesses the opportunity to participate in global markets, often for the first time. Yet whether these opportunities become genuinely inclusive will continue to depend on sustained openness across models, tools, platforms, and ecosystems.

From lowering barriers to transactions in the e-commerce era to lowering barriers to creation in the AI era, platform companies are not only inventors and users of AI technologies. They are also builders of AI ecosystems and creators of economic opportunity, enabling more people to participate in global innovation and commerce.

AI for Safe and Inclusive Communities

The GreenNet Program and Qwen AI Mental Health Companion
Youku Barrier-Free Theater
Taobao AI Video Calling
DingTalk A1 for Social Workers
Amap's Eagle Eye Guardian
"Reviving Ancient Murals" AI-powered Digital Reconstruction

AI should help make communities more caring. Building safe, inclusive, and accessible human settlements and public spaces is one of the important goals of the United Nations 2030 Agenda for Sustainable Development. This goal has also expanded from physical environments into digital spaces. Yet whether online or offline, public services continue to face challenges including insufficient coverage, limited professional resources, and delayed responses.

AI is providing a new technological foundation for these services. From safety, emotional support, and accessibility in digital spaces, to community services and public safety in physical environments, and cultural heritage conservation in historical spaces, AI is increasingly embedding safety, convenience, and inclusion into everyday life. It is creating opportunities for more needs that might otherwise be overlooked to be recognized, understood, and addressed.

The GreenNet Program and Qwen AI Mental Health Companion

From Intercepting Risks to Addressing Underlying Needs

Mental health has become one of the most representative areas of long-tail demand for public services. Stigma, financial pressures, and the uneven distribution of professional resources make it difficult for some people experiencing potential mental health crises to receive timely support. An even more critical challenge is whether help can be made available at the earliest moment when someone expresses a willingness to seek assistance or experiences suicidal thoughts.

AI still cannot replace psychotherapy, but it can enable platforms to become an early source of support. Building on technological advances in AI companionship and conversational interaction, Qwen's AI Mental Health Companion has partnered with The GreenNet Program to develop a closed-loop support system that extends from risk identification and emotional support to referral for professional assistance.

During fiscal year 2026, the GreenNet program used dynamic lexicon updates and AI-based semantic recognition to add nearly 50,000 keywords to its risk-intervention lexicon, including terms associated with suspicious searches for depression-related product transactions.

When the GreenNet Program identifies high-risk search behavior associated with suicidal or self-harming tendencies, the Qwen AI Mental Health Companion can provide emotional support and practical guidance through conversations initiated by users. Through continued interaction and support, it seeks to help users gradually regain emotional stability and, when necessary, guide them toward professional assistance.

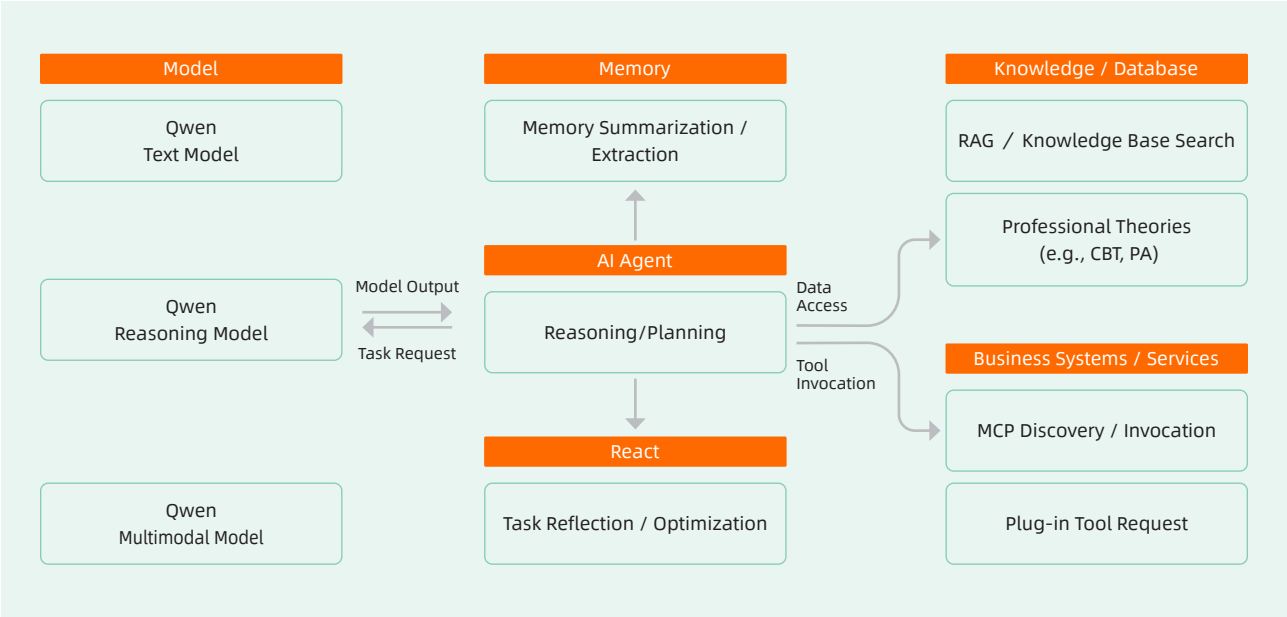

Compared with commercial applications, mental health crises impose higher requirements for safety, reliability, and alignment with appropriate values. They are also more likely to expose the limits of model capabilities, creating a need for continuous improvement in companion agents.

To address these challenges, the team has collaborated with East China Normal University and other institutions to incorporate psychological assessment scales and clinical experience into a structured evaluation framework. For scenarios involving extreme risks, the system can guide users toward offline mental health support services, while strictly adhering to user authorization and compliance requirements throughout the process.

As of August 2026, the mental health support system accessed through the GreenNet Program had facilitated approximately 1.09 million conversational interactions, provided support to around 455,000 people, and reached a peak of 16,000 visits in a single day. It had also successfully referred 2,376 people for further assistance.

Youku Barrier-Free Theater

Creating an Inclusive and Accessible Audiovisual Space

As access to information, cultural entertainment, and public services increasingly shifts online, digital spaces have become important public spaces in everyday life. Yet if a space cannot be used equally by everyone, its accessibility remains limited.

For approximately 28.7 million people with hearing impairments, 17 million people with visual impairments, and many older adults in China, substantial barriers to accessing digital content remain.

Drawing on the capabilities of multimodal large models, Alibaba is developing AI-powered accessible content production. AI can now directly interpret visual scenes, dialogue, environmental sounds, and narrative structures.

At Youku Barrier-Free Theater, AI supports the adaptation of audiovisual content by generating audio-described versions for viewers with visual impairments and enhanced caption annotations for viewers with hearing impairments. While preserving the integrity of the story, these versions use descriptive narration and subtitle annotations to convey characters' actions, spatial relationships, environmental sounds, emotional changes, and other important details.

As a result, the viewing experience can move beyond simply knowing what has happened toward understanding why it happened.

AI has substantially shortened the production cycle for accessible versions of films and television programs, allowing newly popular content to become accessible more quickly. AI-generated narration voice packs also help meet the diverse personal preferences of viewers with disabilities.

Previously, producing an accessible version of a two-hour film, from writing the audio description script and recording narration to completing production and release, often took a month. With AI, the process has been shortened from at least one month to within three days.

AI narration voice packs also allow viewers with visual impairments to choose the voices of their favorite movie stars, such as Hu Ge, to describe the visual details of a film throughout the viewing experience.

As of October 2026, Youku Barrier-Free Theater had made more than 9,500 films and television episodes accessible. Audio-described and enhanced-subtitle versions continue to be updated, with three films added each month and three television episodes each week in each format.

During fiscal year 2026, accessible content on the platform recorded more than 1.21 million plays.



Youku's Barrier-free Theater Volunteers on Campus

Taobao AI Video Calling

Improving the Shopping Experience for Users with Visual Impairments

For people with visual impairments, meaningful participation in digital life depends on the continuous development of accessible technologies.

Taobao has identified 1.51 million users, 980,000 of whom are monthly active users of the Taobao mobile app. These users include not only people with visual impairments, but also people with presbyopia, low vision, or temporary needs for screen-reading assistance arising from other circumstances.

From the web era to the mobile internet era, Taobao has continuously improved accessibility. Screen-reader compatibility helps users listen to page text and identify interactive controls. Duguang OCR (Optical Character Recognition) further enables text embedded in images to be read aloud.

Yet making text readable does not necessarily make products understandable. Details such as color, style, and the location of defects still depend heavily on visual presentation.

How can an e-commerce system that has long relied on visual interfaces and text-based interactions be redesigned to better serve people with visual impairments and other screen-reader users?

Snap&Shop RTC, also known as Taobao AI Video Calling, is exploring a new answer through AI.

Users can activate their cameras, point them toward physical objects, and ask questions. Through image recognition, product search, and voice interaction, the system brings "seeing" and "searching" together in a continuous conversation.

Importantly, this feature is not being developed exclusively for people with visual impairments. It also represents Taobao's effort to extend its core product understanding and search capabilities from visual interaction to voice interaction.

As interfaces expand from "seeing and typing" to "speaking and listening", potential beneficiaries also include older adults experiencing declining vision, parents whose hands are occupied while caring for

children, and users who are less comfortable with text input.

From screen-reader compatibility and Duguang OCR to Snap&Shop RTC, Taobao continues to explore a progression from making interfaces operable and information readable toward enabling users to make independent judgments.

Making AI accessible to all does not mean requiring everyone to adapt to the same interface. It means enabling the same service to adapt to more people.

DingTalk A1: An AI Assistant for Social Workers

A Practical Tool for Frontline Social Workers

Social services depend heavily on sustained human relationships, personal support, and trust. Yet frontline social workers, community grid workers, and volunteers have long faced two major challenges: shortages of professional personnel and high staff turnover.

More seriously, approximately 30% to 50% of their working time is occupied by administrative tasks such as organizing records, filing documents, and entering information into multiple platforms.

As a result, the quality and availability of social services remain constrained by human resources. When experienced workers leave, much of their accumulated knowledge and expertise often leaves with them.

In 2026, Alibaba Philanthropy partnered with DingTalk, with professional support from Fudan University's Department of Social Work, to transform DingTalk's integrated AI software and hardware capabilities into tools designed to empower social workers, initially focusing on philanthropic services for children.

DingTalk A1 combines DingTalk's AI software and hardware capabilities. Through speech transcription, automatic recording, structured information organization, and assisted analysis, it helps social workers complete repetitive tasks such as documentation, filing, and reporting.

AI does not replace social workers. Instead, it gives time back to them.

Home visits, communication, and personal companionship, activities that AI cannot replace and that should never be marginalized, can once again become the central focus of their work.

As AI begins recording service activities, the service process can be preserved more comprehensively. Each visit, intervention, and follow-up no longer exists only in the moment. Instead, these activities gradually become part of a body of practical experience that can be analyzed and reused. Different service approaches can also be compared and improved against a common basis for the first time.

AI is enabling social service organizations to develop a capacity for continuous learning. Although staff turnover in grassroots social services remains relatively high, professional experience can be continuously accumulated through DingTalk's organizational platform rather than disappearing when individual employees leave.

New social workers can therefore build on the experience of their predecessors, while these capabilities can also be replicated across geographic boundaries.

As of June 2026, DingTalk A1 had served more than 150 philanthropic organizations. Applications for the first phase of hardware deployment covered 1,000 social workers, with the initial group of 100 already using the devices.

In the same month, DingTalk A1 introduced the Doukou Doctor Edition, extending methodologies developed in social service settings into healthcare and other professional fields.

Amap's Eagle Eye Guardian System

Moving Public Safety Toward Proactive Early Warnings

Road traffic accidents remain one of the world's major public safety challenges, claiming more than one million lives every year.

Traditional traffic safety management has focused primarily on accident reporting, emergency response, and road control after incidents occur. Yet risks arising during high-speed driving can develop within seconds.

Society's expectations for road safety are therefore shifting toward continuous risk detection, real-time assessment, and proactive warnings before accidents occur.

Amap, together with the China Academy of Safety Science and Technology, developed the Eagle Eye Guardian early warning system.

The system analyzes large volumes of anonymized and de-identified vehicle speed, acceleration, and location data. Combining spatial intelligence with cloud-based large models, it identifies potential hazards in real time, including abnormal deceleration or unexpected stops involving multiple vehicles traveling in the same direction.

It then delivers warnings within seconds to vehicles approaching from behind, using both map-based visual alerts and voice notifications, giving drivers valuable additional reaction time.

In 2026, the system further expanded its capabilities to cover unusual weather conditions, such as snow and ice and localized dense fog, as well as special road environments, extending its reach across more public safety scenarios.

The value of safety lies in preventing accidents before they happen.

As of September 30, 2026, the Eagle Eye Guardian system had issued more than 71.3 billion warnings since its launch, with average daily warnings reaching 190 million.

According to Amap's big data analysis, on pilot road sections including the G2 Beijing-Shanghai Expressway, the daily average number of accidents per 10,000 vehicles during the 2025 National Day holiday declined by approximately 10% year over year following the system's deployment.

AI Is Also Helping Protect Older Adults with Alzheimer's Disease

In 2026, 18-year-old university student Dong Zhicheng developed Angui, a collaborative search assistance system for locating missing people, using Alibaba Cloud to build its backend and integrating Qwen models with Amap. On the family interface, Qwen helps family members provide essential information, including when and where a person went missing and what clothing they were wearing. On the volunteer interface, Amap supports location sharing and marking areas that have already been searched. Finally, on the command interface, human coordinators consolidate updates and organize search operations.

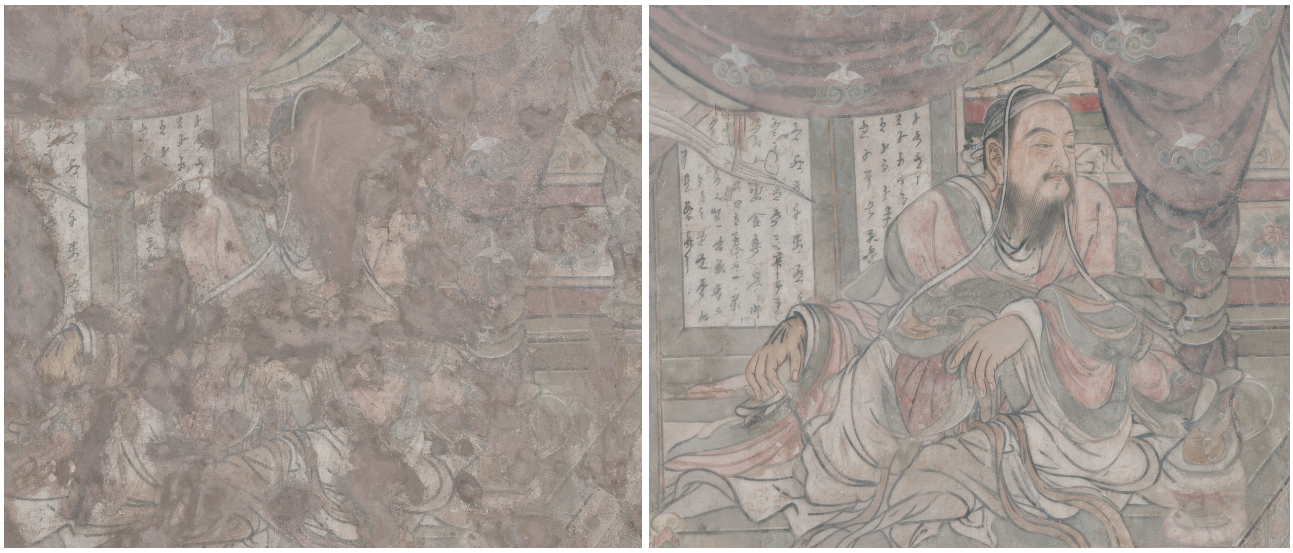

Left: Original faded murals at Dayun Temple in Pingshun, Shanxi.
Right: AI-assisted digital reconstruction version.

"Reviving Ancient Murals": AI-Powered Digital Reconstruction Initiative

Using AI to Preserve Cultural Heritage at Risk of Disappearing

The boundaries of inclusion can extend further, into humanity's shared historical spaces.

For cultural heritage that is gradually disappearing, the greatest threat does not necessarily come from a sudden event. It may come from the passage of time itself.

The Amitabha Hall at Dayun Temple in Pingshun County, Shanxi Province, was built in 940 CE. It is one of the few surviving wooden structures from China's Five Dynasties period, often described as a "living fossil of Five Dynasties architectural art", and is designated as a Major Historical and Cultural Site Protected at the National Level.

On the hall's eastern wall is a mural depicting the Transformation of the Vimalakirti Sutra, covering approximately 21 square meters. It is the only relatively well-preserved surviving temple mural from the Five Dynasties period in China.

Yet the mural is gradually deteriorating.

This is not a challenge unique to Dayun Temple. China has approximately 150,000 square meters of surviving murals, many of which face similar threats.

Mural conservation requires highly specialized knowledge, including an understanding of mineral pigments, traditional artistic techniques, and painting styles from different historical periods. Specialists with this combination of expertise are extremely rare nationwide. In counties and townships, many valuable murals struggle to receive even basic assessments of their deterioration.

Alibaba Philanthropy, together with partners including Xi'an Academy of Fine Arts, Aizhidao, Wan, and Duiyou, launched the "Reviving Ancient Murals" AI Digital Reconstruction Initiative.

The initiative explores a method of digitally reconstructing historical murals through the combination of conservation professionals, large visual models, and specialized AI agent skills.

Built on the Wan-2.7-Image model, the approach uses AI to identify and digitally remove signs of weathering, infer the original composition of missing sections, and generate richly colored digital reconstructions while respecting the artistic style of the original paintings.

Zhang Li, Dean of the School of Design at Xi'an Academy of Fine Arts, explained that the first phase of specialized knowledge training had been completed, saying, "It is now a mural expert."

The resulting AI capabilities have been made available on the Duiyou platform for mural conservation practitioners across China.

Many of China's murals are scattered among small temples in counties and villages, where their caretakers are not necessarily leading conservation specialists.

When AI becomes a tool for identification, assessment, and digital reconstruction, more frontline cultural heritage workers can acquire the capabilities needed to protect these historical spaces.

Looking Ahead

AI is changing not only how information is produced, but also how people form their understanding of the world. In the future, people will increasingly interact with AI agents capable of sustained conversations, continuous feedback, and proactive adaptation to individual preferences. Particularly in sensitive areas such as mental health, financial consumption, and public opinion, the cognitive feedback loops created by AI may also reinforce users' existing biases. The development of safe and inclusive digital capabilities will therefore increasingly extend beyond content safety toward intrinsic model safety.

AI can reduce the costs of acquiring knowledge and taking action, but it can also lower the costs of carrying out harmful activities. As AI increasingly becomes an entry point for search, conversation, shopping, and decision-making, it also begins to influence what users believe, what they choose, and what actions they ultimately take. Looking ahead, the focus of digital safety and inclusion will shift further upstream, with safeguards increasingly embedded throughout human-AI interactions, potentially in the form of dedicated safety agents.

Chapter 4
Looking Ahead



How Society Shapes AI's Next Stage of Evolution

All 26 cases presented in this year's AI for Good Action Report are grounded in real social needs. For the development of AI technology, real-world applications provide the most fertile ground for innovation. Technologies invented in laboratories inevitably undergo a process of being "reinvented" as they enter specific social contexts. This process necessarily involves trial and error, feedback, and adjustment. These very efforts, in turn, provide new sources of learning and improvement that drive technological evolution.

AI for Good is therefore not only the purpose of our efforts to advance AI, but also a structural force driving its continued development. In last year's report, we introduced the idea that "AI for Good is Good for AI". Building a mutually reinforcing flywheel between AI for Good and technological progress requires three interconnected feedback loops:

Loop One: Real-World Feedback. As AI enters complex and demanding real-world settings, such as helping physicians interpret medical images, assisting rural teachers with lesson preparation, or supporting people with visual impairments in shopping, it can receive feedback unavailable in laboratory environments. The complexity and diversity of AI for Good applications are themselves scarce resources for developing more capable AI and strengthening AI systems' ability to act.

Loop Two: Value and support. AI for Good initiatives address genuinely important problems and require sustained investment. As capabilities accumulate and improve, these initiatives may develop into sustainable business models or attract support from public resources. This process requires each cycle to deliver some degree of improvement, making it particularly challenging.

Loop Three: Trust and room to operate. When AI can reliably move beyond isolated trials into sustained service delivery, entering socially valuable fields that were previously difficult to access, it can earn greater trust and gain more room to develop. This creates opportunities for AI to move from providing assistance in specific tasks toward deeper forms of collaboration.

AI for Good is good for AI

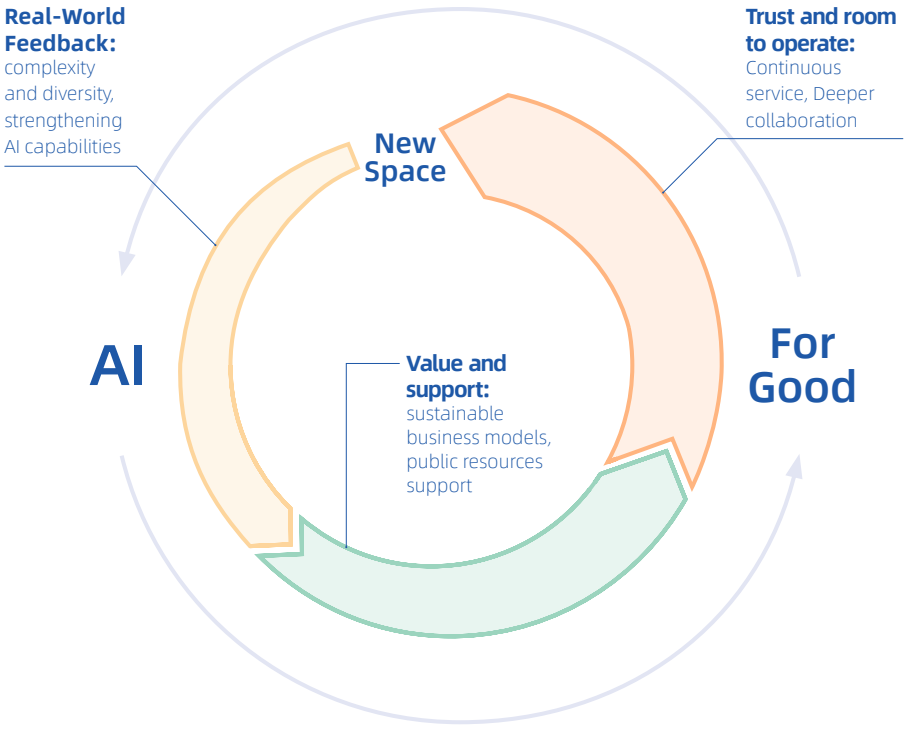

Every real-world deployment is a genuine dialogue between AI and society. For AI companies, only through this flywheel can numerous, dispersed explorations of application contexts be brought together to support systematic capability building and institutional learning, both of which are necessary for AI's next stage of evolution.

Seeking Truth, Pursuing Good: Eight Questions for AI's Future

Over the past year, the pace of change has astonished many of us, often departing sharply from our familiar expectations and judgments. This reminds us of the importance of remaining humble.

In the AI era, we face a shared dilemma: the best window of opportunity to shape the direction of a technology is precisely when our understanding of it is still relatively limited. More specifically, during the early stages of rapid technological development, its social impacts are difficult to predict. Yet once a technology has matured and become deeply embedded in society, controlling its negative consequences or changing its course becomes far more difficult and costly.[1]

The first report published by the United Nations Independent International Scientific Panel on Artificial Intelligence also highlights an evidence dilemma facing policymakers: they must either make consequential AI governance decisions immediately, despite insufficient scientific evidence, or wait until adequate evidence becomes available, by which time intervention may already be too late.

Therefore, as we discuss AI for Good, we must not neglect the pursuit of truth. We need to continuously adjust our approaches in light of the latest evidence and facts. Without developing and using AI, we cannot truly understand it. Without sufficiently examining and understanding the relationship between AI and society, it will be difficult to establish the agile governance mechanisms needed to address emerging challenges.

Throughout the discussions and preparations for this report, we have continued to ask fundamental questions about AI for Good. These questions reflect both the self-examination expected of an AI company and our hope for open and meaningful dialogue with society.

1. Impact Validation: How Can We Determine Whether AI Delivers Real and Lasting Improvements?

How can we distinguish technical performance, service efficiency, and ultimate benefits to people, using context-appropriate evaluations to identify AI's actual contribution? How can we assess long-term outcomes, differences among population groups, and potential adverse effects, rather than focusing solely on the scale of adoption?

2. Safety at Scale: How Can We Detect Risks Earlier and Limit Their Impact?

As AI agents begin executing tasks, how can we connect pre-release evaluations with ongoing operational monitoring, and enable systems to be paused, rolled back, or handed over to human operators when necessary? How can we fulfill applicable disclosure obligations while providing greater clarity about the scope and limitations of safety evaluations, without revealing details that could be misused?

3. Openness and Governance: How Can We Choose Approaches to Openness That Are Appropriate to Model Capabilities?

How can we assess the appropriate scope of openness, release conditions, and accompanying safeguards based on model capabilities, intended uses, and foreseeable misuse? How can external evaluations and community feedback inform subsequent release decisions, supporting innovation without assuming that openness automatically guarantees safety?

4. Resources and the Environment: How Can We Align the Growth of AI Applications with Environmental Goals?

How can we combine efficiency improvements, clean energy, water resource management, and workload scheduling while accounting for both efficiency per unit and absolute resource consumption? How can we clearly define measurement boundaries and improve

[1] This phenomenon is known in the philosophy of technology as the Collingridge dilemma. See Collingridge, D. (1980). The Social Control of Technology. Frances Pinter.

environmental disclosures relating to training, inference, and infrastructure, in accordance with applicable requirements and available data?

5. Responsibility and Collaboration: How Can We Establish Clear Lines of Accountability for AI Services Involving Multiple Parties?

Regardless of whether a model is open or closed, how can we clearly define the responsibilities of different participants in AI service delivery for authorization, oversight, notification, and incident response? When problems arise, how can accountability be enforced through necessary records, investigations, corrective actions, and user remedies, so that a complex technology chain does not obscure lines of responsibility?

6. Human Capabilities and Autonomy: How Can AI Assistance Support Long-Term Human Development?

How can we preserve essential opportunities for thinking, practice, and verification in education and professional services, while assessing the capabilities users retain when AI assistance is no longer available? In companionship and personalized services, how can user controls, clear disclosures, and human support help reduce excessive dependence and preserve people's autonomy in making choices?

7. Workforce Transitions: How Can We Support New Pathways for Growth as the Nature of Work Changes?

How can we monitor changes in tasks and skills across different occupations without directly equating productivity improvements with increases or decreases in employment? How can job redesign, mentorship, workplace learning, and retraining preserve pathways for career development, particularly for new entrants and groups affected by technological change?

8. Inclusive Access: How Can We Move from Making AI Available to Enabling People to Use It Effectively?

How can we identify practical barriers related to language, devices, cost, and accessibility, while involving intended beneficiaries in the design and evaluation of services? How can we combine resource support with training, local services, and ongoing maintenance, and verify whether these investments genuinely improve people's opportunities to benefit?

These questions have no single technological answer, and it would be difficult for any one company to resolve them independently.

Advancing this exploration requires us to incorporate real-world outcomes, environmental impacts, and inclusive access into technological and business decisions, while complying with laws and regulations, protecting fundamental rights and interests, and meeting essential safety requirements.

Improving practices through verifiable evidence and enabling collaboration through clear accountability are central to our understanding of seeking truth and advancing AI for Good.

Selected References

1. Preliminary Report of the Independent International Scientific Panel on AI, United Nations, 2026.
2. Action Plan for Promoting Coordinated Development Among Large, Medium-Sized, and Small Enterprises in the Platform Economy (2026–2028), Ministry of Industry and Information Technology and six other government departments.
3. Implementation Guidelines for the Responsible Application and Innovative Development of AI Agents, Cyberspace Administration of China.
4. Self-Discipline and Social Benefit: Developing Responsible Technology in the Age of AI, Alibaba and Shanghai Artificial Intelligence Laboratory.
5. Alibaba Environmental, Social and Governance (ESG) Report 2026, Alibaba.
6. Alibaba AI for Good Action Report 2025, Alibaba.
7. 2026 AI Open Source Talent Report, ModelScope Community.
8. Research Report Series, Alibaba Research Institute.
9. Technological Revolutions and Financial Capital, Carlota Perez.
10. The Social Control of Technology, David Collingridge.
11. Resilience in the AI Era, OpenAI.
12. Anthropic's Advanced AI Framework, Anthropic.
13. An Economy of AI Agents, Gillian K. Hadfield, Andrew Koh.
14. A.I. and Our Economic Future, Chad Jones.
15. Funding AI for Good: A Call for Meaningful Engagement, Hongjin Lin, et al.
16. A definition, benchmark and database of AI for social good initiatives, Josh Cowls, et al.
17. Agentic AI and the next intelligence explosion, James Evans, et al.
18. Advancing solar and wind penetration in China through energy complementarity, Hu, Y., Jiang, H., Zhang, C. et al.
19. Colorectal cancer detection using noncontrast CT and deep learning: a multicenter and international cohort study, Chen X, Qiu M, Zhang J.
20. Harness Engineering for Self-Improvement, Lilian Weng.
21. Ask an AI expert: What exactly is the full stack? Molly McHugh-Johnson, Google.
22. Learned Hardware/Software Co-Design of Neural Accelerators, Zhang Shi et al.
23. The health of nations: Stronger health, stronger economies, McKinsey Health Institute.
24. Cancer burden and control in China: landscape, trends and challenges, Wanqing Chen, Chen Wu, et al.
25. MM-StoryAgent: Immersive Narrated Storybook Video Generation with a Multi-Agent Paradigm across Text, Image and Audio, Mengyue Wu, Ming Yan.
26. The future of artificial intelligence and the mathematical and physical sciences, NSF, Andrew Ferguson, et al.
27. Agentic Fusion of Large Atomic and Language Models to Accelerate Superconductor Discovery, Wenbing Huang, et al.
28. Oyster-II: Reinforcement Learning for Constructive Safety Alignment in Large Language Models, Alibaba AAIG.
29. State of Open Models: Summer 2026 Observations, Hugging Face.
30. Reasoning Models Generate Societies of Thought, Junsol Kim et al.
31. The Economics of Recursive Self-Improvement, Tom Cunningham et al.


Alibaba