AI Adoption Cases Across Sectors

November 05, 2026

As of early November 2026, artificial intelligence technology is being widely integrated into practical operational fields, ranging from defense weapon systems, space radio observation, hospital image reading, manufacturing processes, and corporate IT operations to government administrative networks. As agents (autonomous task-performing AI) that receive user objectives, autonomously call external tools, and continue tasks spread to frontline operations, control mechanisms to prevent IT network damage caused by unauthorized setting changes by external programs or reasoning errors have become essential prerequisites for adoption. In foundational technology, lightweighting research to reduce the computational cost burden of maintaining long conversation contexts and rapidly run small models inside devices stands out, while parliaments in various countries are testing artificial intelligence as a legislative assistance tool while concurrently overhauling safety regulations and international norms.

Contents

1. Defense, Military, and Autonomous Systems 3
- A. Key adoption cases 3
- B. Maritime Uncrewed Combat Systems and Manned-Unmanned Teaming Command 3
- C. Aerial and Ground Autonomous Weapons, Battlefield Surveillance, and Logistics Intelligentization 5

2. Cybersecurity and AI Safety 7
- A. Key adoption cases 7
- B. Agent Supply Chain Threats and Corporate Risk Disclosures 7
- C. Safety Evaluation of Security Patching and Investigative Reasoning 9

3. Scientific Research and Research Infrastructure 10
- A. Key adoption cases 10
- B. Real-Time Astronomical Observation and Particle Physics Simulations 10

4. Healthcare and Life Sciences 12
- A. Key adoption cases 12
- B. Medical Imaging/Pathology Analysis and Disease Prediction 12
- C. Hospital Operations and Care Administration 12

5. Humanoids and Service Robots 14
- A. Key adoption cases 14
- B. Robot Manipulation Data Collection Facilities and Visual Perception Models 14

6. Manufacturing, Factories, and Industry 16
- A. Key adoption cases 16
- B. Industrial Equipment Embedded Vision Models and Hardware Design Experiments 16

7. Transportation and Logistics 17
A. Key adoption cases 17
B. Autonomous Driving Virtual Environment Simulation 17
8. Energy and Environment 18
A. Key adoption cases 18
B. Circular Economy and Waste Sorting Technology 18
9. Finance, Insurance, and Accounting 19
A. Key adoption cases 19
B. Fintech Engineering Infrastructure Cost Reduction 19
10. Video, Content, and Creative Industries 20
A. Key adoption cases 20
B. Enterprise Design Document Generation and Game Development Policy Adjustments 20
11. Software Development and Work Automation 21
A. Key adoption cases 21
B. Enterprise Agent Control Frameworks and Cloud Operation Practical Verification 21
12. Foundation Models and Common Technologies 23
A. Key adoption cases 23
B. Re-evaluating Language Model Limitations and Long-Context Efficiency 23
C. Device Operation Interconnection, Edge Inference Acceleration, and Visual Representation Learning 25
13. Law, Policy, and Public Governance 27
A. Key adoption cases 27
B. Pan-Government Administrative AI Launch and Platform Age Rating Regulations 27
14. AI Discussions and Legislation in Parliaments Worldwide 29
A. Key adoption cases 29
B. Reviewing Field Applications in Legislation, the Judiciary, and Education, and Protecting Electoral Democracy 29
C. International AI Governance Cooperation and the Korean National Assembly's Discussions on Industry, Regulation, and Copyright 31
References 34

1. Defense, Military, and Autonomous Systems

A. Key adoption cases

In the defense sector, starting from the international maritime defense exhibition Euronaval 2026 and combat experiences in Ukraine, the autonomous control of maritime, air, and ground uncrewed systems and the construction of mass-production foundations are progressing in tandem. Moving beyond remotely piloting individual equipment, Manned-Unmanned Teaming (MUM-T)—which teams crewed and uncrewed equipment together—is becoming established as standard design, and the trend of reducing decision-making time across operations and logistics, from satellite reconnaissance to weapon maintenance, continues.

B. Maritime Uncrewed Combat Systems and Manned-Unmanned Teaming Command

Uncrewed Underwater Vehicle and Uncrewed Surface Vessel Development by Major European Nations and Fleet Combat System Modernization

On November 4, 2026, the French Ministry of the Armed Forces announced a development contract with shipbuilding defense contractor Naval Group for a large ocean-going Uncrewed Combat Underwater Vehicle (UCUV) demonstrator. It is designed to handle missions ranging from surveillance and reconnaissance to offensive operations by equipping a controlled autonomous decision-making system that autonomously responds to tactical situations, marine environmental changes, and equipment failures. In the 2026 French defense budget, a commitment authorization of 30 million euros and payment appropriations of 17 million euros were allocated, with sea trials scheduled to begin in 2028 to verify operational suitability.

In surface forces, Sweden's Saab unveiled the CB90 CX on November 2, an 18-meter-class optionally piloted/uncrewed fast combat boat equipped with Autonomous Ocean Core artificial intelligence control software. The first vessel among the 36 ordered by the Swedish Navy is undergoing final acceptance testing prior to delivery, and whether it will be deployed for full uncrewed operations will be verified after passing the tests. On November 3, Saab also announced its next-generation 9LV naval combat system, which applies Zero Trust security—verifying every access privilege each time—and a modular architecture to provide integrated command of manned-unmanned hybrid fleets.

Armed Uncrewed Surface Vessel Swarm Command and Uncrewed Airborne Anti-Submarine Warfare Demonstration

On November 3, Turkish state-owned defense company MKE showcased PIRANA AUSV, a new armed uncrewed surface vessel (USV, a military vessel navigating maritime waters without a crew) featuring a maximum payload of 1,000 kg, a range of 400 nautical miles, and a top speed exceeding 35 knots. Bundling radio frequency, satellite communications, and LTE networks, it operates as a swarm command-and-control hub leading formations of maritime kamikaze uncrewed boats, and the exhibition mock-up housed eight tube-launched uncrewed aerial vehicles of undisclosed models inside. Simultaneous airborne ejection and the full integrity of open-water swarm control remain tasks for additional sea trials.

In NATO exercise REPMUS 2026 held off the coast of Portugal, Tekever and General Dynamics Mission Systems-Canada announced on November 3 that they demonstrated uncrewed airborne anti-submarine warfare (ASW, operations to detect and track enemy submarines). Sonobuoy (buoy-type sensors that detect submarines using underwater acoustics) data dropped into the sea by an AR5 uncrewed aerial vehicle was analyzed in real time with a distributed acoustic processing system to detect underwater contacts and share them on the operational picture, while regular operational fielding will be decided following additional force-integration evaluations.

C. Aerial and Ground Autonomous Weapons, Battlefield Surveillance, and Logistics Intelligentization

Deployment of Machine Vision Interceptor Drones on the Ukrainian Frontline and Unveiling of Medium Uncrewed Ground Vehicles

On November 3, the Ukrainian Serhiy Prytula Charity Foundation delivered 626 ZIRKA interceptor drones and three ground control stations—procured through 51.8 million hryvnias in fundraising—to the National Guard's 11th Brigade (326 units) and the 412th Unmanned Systems Brigade (300 units). Equipped with thermal imaging cameras and machine vision (technology where a computer recognizes camera images to identify objects) automatic target tracking algorithms, they fly at up to 410 km/h and autonomously collide once the pilot approves a target within 1 km. The 412th Brigade reported that it succeeded in intercepting a Russian jet drone, Geran-5, flying at 400 km/h at an altitude of 3,500 meters.

In ground logistics support, Ukraine's NUMO Robotics unveiled LOKO on November 3, a wheeled uncrewed ground vehicle (UGV) designed for transporting up to 500 kg of cargo, evacuating wounded soldiers, and towing damaged uncrewed equipment. Powered by four 1.5-kilowatt (kW) electric motors, it travels up to 80 km at 20 km/h on paved roads and 12 km/h off-road, and is preparing for official Ministry of Defense fielding registration procedures incorporating field feedback from over 60 units.

Rollout of Vertical Takeoff and Landing Uncrewed Strike Aircraft and Military Cargo Drone Mass-Production Partnership

U.S. defense company Shield AI announced on November 4 that it rolled out X-BAT 001, the first full-scale physical prototype of an uncrewed strike aircraft that takes off and lands without a runway by combining Hivemind autonomous flight artificial intelligence with a General Electric (GE) F110-GE-129 jet engine featuring thrust vectoring. Actual flight performance will only be confirmed after the first flight tests scheduled to begin at a Kansas test site in late 2026.

In the manufacturing base, U.S. military vehicle manufacturer Oshkosh Defense partnered with startup Grid Aero on November 2 for the mass production of the large autonomous cargo drone Lifter. The concept involves producing thousands of units similar to a truck assembly line by using commodity metals and piston engines instead of expensive composites for an aircraft weighing over 1,320 pounds (approx. 598.7 kg); the actual ramp-up rate will be determined incrementally based on future military procurement demand.

Ultra-Fast AI Satellite Imagery Supply and Korean Military Weapon Maintenance Agent Development

On November 3, U.S. satellite imagery company Vantor (formerly Maxar Intelligence) announced a capability to deliver satellite imagery to frontline Ukrainian units within 7 minutes of an order request by combining edge processing—processing data immediately near ground stations—with Tensorglobe, an AI terrain change detection platform. It explained that photographing the same location 12 to 15 times a day using about ten optical satellites with roughly 30 cm-class resolution reduced the time from surveillance to strike by up to 90%, although the 7-minute transmission record has not yet undergone independent third-party verification.

On November 4, the Republic of Korea's Ministry of National Defense and the Ministry of Trade, Industry and Energy launched a joint research and development project for artificial intelligence agents that calculate reliability, availability, maintainability, and total cost of ownership (RAM-C) of weapon systems. Investing a total of 30 billion won over four years, with Korea Aerospace Industries (KAI) serving as the lead R&D institution, the project will demonstrate failure prediction and optimal maintenance timing calculation targeting the Army's K-2 tank, the Navy's Maritime Surveillance Radar-II, and the Air Force's FA-50 fighter. As it is currently at the development commencement stage, actual operational availability improvements and cost reduction margins will only be confirmed following future empirical evaluations.

2. Cybersecurity and AI Safety

A. Key adoption cases

As artificial intelligence agents handling coding and general tasks directly manage privileges on developer computers and external servers, supply chain threats—where neighboring software alters agent settings without authorization or plants advertisements inside natural language instructions—have been identified. As model limitations and risks of misjudgment were also numerically exposed in security vulnerability repairs and investigative record analyses, enterprises are responding with cryptographic output proofs and official risk disclosures.

B. Agent Supply Chain Threats and Corporate Risk Disclosures

Development Tool's Unauthorized Setting Registration and the Emergence of Prompt Adware

On November 3, 2026, it was revealed that version 12.33.4 of the Postman desktop app, an API development tool, registered plugins, command-line tools, a remote Model Context Protocol (MCP, a standard specification used when AI agents connect to external tools or data) server (mcp.postman.com), and session-start hooks (scripts that run automatically under specific conditions) into Codex and Claude Code configuration files without the prior consent of existing logged-in users. Although no evidence of external data exfiltration emerged, it demonstrated a structural flaw where other programs can arbitrarily modify agent execution paths.

On November 3, security firm Check Point Research analyzed the Haiou package, which conceals itself within a Codex jailbreak bundle file, implants six skill modules and instructions into global settings, and causes advertising links for a specific relay server to be repeatedly printed. The researchers defined the

technique of achieving commercial objectives by tampering not with executable program code but with the natural language instructions trusted by the agent as Prompt Adware, pointing out that it is distributed in encrypted and compressed form, making prior detection difficult.

Apple's Agent Risk Disclosure and Anthropic's Zero-Knowledge Proof Protocol

In its Annual Report (Form 10-K) reported on November 3, Apple formally specified for the first time as an official business risk factor the risks of artificial intelligence agents behaving autonomously in unintended manners and executing unauthorized system access, code modifications, message transmissions, and unauthorized payments, along with the limitations of overseeing third-party agents.

To enhance trust during model serving, Anthropic unveiled FlashProve on November 2 using zero-knowledge proofs (ZKP, a cryptographic technology that mathematically proves a certain fact is true without revealing internal secret information). It is a method where a service provider proves that the exact model agreed upon in advance generated the response tokens while keeping the model weights (internal numerical values learned by neural networks) private. Based on Llama 4 Maverick (a 4-bit model with 400 billion parameters and a 40-billion-class active architecture), it proved 110 tokens per second and verified in 17 seconds on a single Nvidia H100 graphics processing unit (GPU) (with a 19 MB proof file per 1,024 tokens), but because it takes 70 to 95 times more GPU compute time than standard inference, it remains at the research prototype stage rather than commercial service.

C. Safety Evaluation of Security Patching and Investigative Reasoning

Limitations Identified in Open-Source Security Vulnerability Fixing and Virtual Investigation Evaluations

The CWE-bench v2 evaluation released on November 3 by Hugging Face and security researchers tested agents on finding and fixing weaknesses on their own across 100 real open-source repository tasks and 91 Common Weakness Enumeration (CWE) categories. Gemini 4 Argon achieved the highest rates with a single-attempt success rate of 67.8% and a four-attempt success rate of 81.0%, Claude Opus 5.5 recorded 67.3%, and DeepSeek V4.1 Flash resolved 57.0% at a cost of $0.04 per attempt. However, even the top-performing model failed to repair 19% of the total tasks at all.

KierstonBench, published on November 4 by AI researchers, tested 32 incident investigation tasks based on administrative, telecommunication, and financial records of a virtual city of 125,000 residents. Most evaluated models exhibited hallucination errors (a phenomenon where AI plausibly fabricates information contrary to facts)—identifying innocent citizens with no basis in records as suspects instead of withholding judgment when clues were insufficient—confirming that they cannot replace human investigators in actual investigations or auditing work.

3. Scientific Research and Research Infrastructure

A. Key adoption cases

In basic scientific observation, a shift in research infrastructure is emerging away from storing large volumes of data for later batch analysis toward artificial intelligence processing signals in real time on-site at observation equipment and enhancing measurement resolution.

B. Real-Time Astronomical Observation and Particle Physics Simulations

Radio Telescope Real-Time Signal Analysis Agent and Gamma-Ray Observation Resolution Improvement

The SETI Institute and Breakthrough Listen announced on November 3, 2026, that they applied an artificial intelligence agent based on the Nvidia Holoscan platform to the Allen Telescope Array (ATA). While the hardware processes radio signals at up to 2.4 terabits per second (Tbps), the agent controls telescope pointing directions and analyzes signals upon arrival, speeding up readings that previously took weeks into hours and capturing the Crab Nebula pulsar (a rapidly rotating neutron star periodically emitting radio waves) in real time. Whether any artificial signals from extraterrestrial intelligence were discovered has not been confirmed.

In a study published on October 7, 2026 (included in the comprehensive report on November 4), researchers of the Tibet-III air shower array combined convolutional neural network (CNN, a neural network extracting spatial patterns from images) analysis with time-of-arrival data of particles striking ground detectors. As a result of Monte Carlo simulations calculating physical phenomena using random samples, the angular resolution distinguishing the arrival

direction of cosmic gamma rays improved by about 10–15% compared to conventional techniques, leaving replication and verification in actual observation data as a follow-up task.

4. Healthcare and Life Sciences

A. Key adoption cases

In the medical field, attempts to incorporate structured voice reporting and multi-agent monitoring systems into existing hospital IT environments are continuing in order to alleviate repetitive administrative burdens on radiologists and clinical laboratory personnel.

B. Medical Imaging/Pathology Analysis and Disease Prediction

Mount Sinai Health System in New York Signs Radiology AI Reporting Contract

New York's Mount Sinai Health System in the U.S. announced on November 4, 2026, that it entered into a seven-year contract with medical imaging company Sectra to integrate artificial intelligence report-drafting capabilities into over 1.2 million radiology examinations annually. The architecture assists in organizing the interpreting physician's voice dictation into standardized reports and comparing it side by side with AI-analyzed drafts; implementation is currently underway targeting deployment in spring 2027, so it is not yet in full clinical operation.

C. Hospital Operations and Care Administration

Weill Cornell Medicine Proposes Multi-Agent Monitoring System for Clinical Laboratories

In a perspective paper published on October 28, 2026 (included in the comprehensive report on November 4), researchers at Weill Cornell Medicine in the U.S. proposed LabARIA, a multi-agent architecture designed to support the operation of blood and urine clinical laboratories. Designed so that a central orchestrator distributes information on equipment operating statuses, specimen

quality, and test delays scattered across multiple IT networks to specialized agents for integrated monitoring and reagent depletion forecasting, it is currently at the academic proposal stage and requires empirical evaluation in actual hospital environments.

5. Humanoids and Service Robots

A. Key adoption cases

In the field of humanoids—human-shaped robots—dedicated training facilities to collect action data for dexterously manipulating objects have opened, while perception models that enhance robots' spatial understanding using visual information alone without ground-truth joint control labels and explorations of military remote operation are progressing concurrently.

B. Robot Manipulation Data Collection Facilities and Visual Perception Models

Opening of Daejeon Humanoid Data Factory and Release of Unlabeled Perception Model

On November 3, 2026, the Korea Institute of Machinery and Materials (KIMM) and the Electronics and Telecommunications Research Institute (ETRI) opened an AI Humanoid Data Factory equipped with 12 humanoids in Daejeon. It is a facility where motion data of grasping and moving objects is accumulated when operators move robots using teleoperation devices; whole-body motion imitation technology using Kairos robots was demonstrated, with plans to open it incrementally to external industry, academic, and research organizations starting in late 2027.

On November 4, Chinese electric vehicle and robotics company Xpeng, jointly with Peking University and the University of Hong Kong, open-sourced GroundingPI, an approximately 4-billion-parameter (Parameter, the number of connection weights indicating a model's computational scale) model pretrained solely on visual perception data without robot joint action labels. It covers 11 domains, including object detection and spatial localization, and was reported to have scored an average of 73.68 points across 34 evaluations, surpassing the

comparison model Astra (71.54 points); however, as these figures are centered on virtual environments, they must be assessed separately from actual physical task success rates.

Chinese People's Liberation Army Explores Humanoids for Urban Operations and Rules of Engagement Control Limitations

According to an analysis of over 100 Chinese military bidding, academic, and patent documents reported on September 8, 2026, the Chinese People's Liberation Army (PLA) and the National University of Defense Technology (NUDT) are exploring military applications based on domestic manufacturing infrastructure, which accounted for approximately 95% of global humanoid shipments as of 2025. A robotics tender in May 2025 and a budget allocation of approximately $300,000 for a sensor data classification system in June 2026 were confirmed, and China North Industries Group (NORINCO) showcased the teleoperated robot Fuxi for hazardous area reconnaissance and guarding (the U.S. Army is also pursuing military humanoid testing valued at up to $1.25 million). However, due to mobility limitations on irregular terrain, mass-production unit costs, and international rules of engagement compliance issues such as distinguishing combatants from civilians and identifying signs of surrender, it currently remains at the stage of teleoperation and building-clearing simulations assuming technology 5 to 10 years away.

6. Manufacturing, Factories, and Industry

A. Key adoption cases

In industrial settings, while small vision-language models running directly inside factory equipment without network connections are being deployed for welding inspection, automated electronic device design driven by general-purpose artificial intelligence failed to account for physical manufacturing tolerances, requiring repetitive human corrections.

B. Industrial Equipment Embedded Vision Models and Hardware Design Experiments

Development of Welding-Specialized Lightweight Vision-Language Models and Limitations in Automated Circuit Board Design

NC AI announced on November 4, 2026, that it developed Baeky Vision 3B, a 3-billion-parameter (3B) vision-language model mounted directly on robots and factory equipment, along with Baeky Vision 3B Weld, specialized for welding processes. According to company descriptions, its Riawelk welding classification evaluation score increased from 41.1 points to 80.9 points, and cooperation is being pursued with Hanwha Ocean for autonomous ship welding and with POSCO DX for industrial humanoid application (actual defect rate reduction performance on-site requires separate verification).

On November 3, OpenAI introduced that GPT-6 Astra took charge of circuit board design and embedded hardware programming (firmware) to complete a prototype voice conversation terminal within one week. However, it also revealed that engineers had to modify the design multiple times because Astra failed to consistently align component assembly conditions and tolerances (tolerance, the permissible range of variation between designed dimensions and actual manufactured part dimensions), showing that fully autonomous design has not yet been achieved in physical manufacturing processes.

7. Transportation and Logistics

A. Key adoption cases

In the autonomous driving field, simulation technology that trains robotaxis by replicating sudden hazardous situations in large volumes within virtual spaces—which are difficult to obtain solely through real-world road driving—is advancing.

B. Autonomous Driving Virtual Environment Simulation

Baidu Unveils Apollo World, a World Model for Robotaxi Training

On November 4, 2026, China's Baidu announced Apollo World, a world model (an artificial intelligence model that simulates spatial structures and physical interactions of the real world) supporting the training of its driverless autonomous taxi service Apollo Go. It explained that it generates up to 100,000 km of virtual driving per day by synthesizing vehicle camera footage and LiDAR (sensors that emit laser light to precisely measure distances to surrounding objects) signals in real time and implementing reactions of surrounding vehicles and pedestrians; however, it did not disclose whether it would be opened to external researchers or data on its error rate compared to actual road driving.

8. Energy and Environment

A. Key adoption cases

In the resource circulation sector, artificial intelligence sorting equipment combining optical imagery of different wavelengths to accurately sort recyclable textiles from mixed waste is being tested on-site at processing facilities.

B. Circular Economy and Waste Sorting Technology

Demonstration of Textile and Footwear Identification in Mixed Waste Combining Multi-Optical Imagery

In a field test announced on October 22, 2026 (included in the comprehensive report on November 4), Spanish applied research institute Tecnalia and environmental services company Valoriza conducted an empirical demonstration identifying textiles and footwear in mixed waste by installing CIRTEXCAM technology at an actual waste treatment plant. The method jointly analyzes standard color images, near-infrared that reads material compositions, and stereoscopic images that capture object volumes, and is at a stage of continually refining models before finalizing the ultimate performance of commercial continuous processing lines.

9. Finance, Insurance, and Accounting

A. Key adoption cases

In the fintech industry, internal adoption cases have been reported where open-weight models are hosted on in-house servers and computational lightweighting is applied to reduce the burden of external artificial intelligence usage fees for large-scale software development organizations.

B. Fintech Engineering Infrastructure Cost Reduction

Revolut's Transition of Coding Agents to In-House Servers and Reduction of Inference Unit Costs

On November 3, 2026, British fintech company Revolut announced that it transitioned the primary model of its internal coding agents from external commercial APIs like Claude to Z.ai's open model GLM-5.3 deployed on in-house servers, reducing inference costs in one year to approximately one-tenth, around $0.30 per 1 million tokens (with image inputs and other tasks handled separately by a dispatcher). It stated that by applying quantization (a lightweighting technique that reduces memory usage and computational burden, NVFP4) reducing model numerical precision to 4-bit floating-point, as well as disaggregated execution of input reading (prefill) and output generation (decode), it shortened coding time by 60% and development lead time by 70%. However, in-house GPU facilities capable of continuously running weights reaching 380 gigabytes (GB) are an essential prerequisite, and the cost and time reduction figures are internal tallies that have not undergone independent third-party verification.

10. Video, Content, and Creative Industries

A. Key adoption cases

In the content creation domain, designs that automatically produce business presentation materials by injecting corporate visual design rules into small models, as well as tool policy changes allowing new creations referencing existing game rules, were carried out.

B. Enterprise Design Document Generation and Game Development Policy Adjustments

IBM's Design-Specification-Compliant Slide Generator and Claude Code Game Creation Policy Relaxation

On November 3, 2026, IBM Research unveiled Palette, which creates presentation materials adhering to corporate design languages by applying Low-Rank Adaptation (LoRA, a technique that fine-tunes only small auxiliary modules without retraining the entire model) to a 20-billion-parameter (20B) open-weight model. Through a process of concurrently generating layout designs and slide generation code (pptxgenjs) and measuring and correcting layout errors such as text overlap, it recorded 59.7 points on the PresentBench evaluation; however, the trained model and adapter files were not distributed externally.

In the Claude Code 2.1.317 update reported on November 4, Anthropic lifted a blocking rule that had uniformly rejected development requests asking to reference existing game mechanics. By officially permitting requests to create new original games by referencing play rules and operating principles without directly copying commercial games, it reduced prototype development constraints for small-scale creators.

11. Software Development and Work Automation

A. Key adoption cases

In the enterprise work automation market, data governance tools that unify metric definitions that differ by department and agent malfunction prevention mechanisms are being released in succession. However, in actual cloud server operation tests, instances where model reasoning mistakes led to permanent infrastructure damage were confirmed, demanding a cautious approach to granting automation privileges.

B. Enterprise Agent Control Frameworks and Cloud Operation Practical Verification

Snowflake's Business Concept Unification System and Anthropic's Execution Safeguard Additions

On November 3, 2026, Snowflake announced Business Ontology (a knowledge system established so computers consistently understand data concepts and relationships; private preview), which defines corporate common metrics and entity relationships so multiple agents do not produce differing calculated values for the same business question. It connects with coding agent CoCo and work agent CoWork, and abuse detection, isolated execution environments, and cost control features were added to the Cortex AI Gateway.

On November 3 and 4, Anthropic deployed Claude Code versions 2.1.317 and 2.1.318, tightening enterprise security controls. It added prompts for prior approval when web or shell tools attempt to transmit user account email addresses to external networks, a blocking configuration (onFailure: block) that immediately halts subsequent operations if software development kit (SDK) hook

execution fails, notifications 10 days prior to login expiration, and enhanced protection against overwriting configuration files during non-interactive execution.

Agent Training Environment Synthesis and Reasoning Bottlenecks Revealed in Cloud Operation Benchmarks

On November 4, AI researchers published the ToolSpace study, which automatically synthesized 804 training environments by extracting data and state dependencies between adjacent tools from real MCP services and public call logs. After generating and training on 34,376 supervised fine-tuning samples and 653 reinforcement learning samples, the Qwen3.6-35B-A3B model score rose from 23.0 points to 41.0 points, though the limitation remained that 77.5% of the synthesized environments merely returned pre-stored static results.

CloudOpsBench, released on the same day, tested 45 cloud operational tasks across real Amazon Web Services (AWS) and simulated environments. Claude Fable 5.1 and GPT-6 Astra reached resolutions of only 60% and 53%, respectively; roughly 80% of total failures were reasoning issues—misidentifying task scopes, omitting related services, or giving up midway—while roughly 20% led to erroneous operations, including irreversible damage to server infrastructure.

Huawei and ByteDance Overhaul Chinese Enterprise Agent Services

In the Chinese enterprise market, the retirement of legacy models and service consolidation are progressing in conjunction. On November 5, Huawei announced that it applied the new openPangu-2.1-Pro (detailed specifications undisclosed) for the first time to its enterprise agent app Xiaoyi Work (小艺 Work), and will cease sales of existing 2.0-series models on November 6 before terminating service on November 30. ByteDance will also suspend new sales of the current enterprise edition of its agent development platform Coze (扣子) starting November 30, consolidating features and pricing tiers into the Doubao (豆包) enterprise service ecosystem (support for existing contracted customers will be maintained within their contract terms).

12. Foundation Models and Common Technologies

A. Key adoption cases

In the foundation model field, while it was confirmed that many assertions about language model weaknesses once accepted as academic orthodoxy have already been resolved through rapid model improvements, practical bottlenecks—such as sharp serving cost surges for maintaining long contexts and failures in multi-step variable tracking—have newly come to the fore. In order to solve these, research on session compression, differential transmission, device operation integration, and edge semiconductor acceleration was published side by side.

B. Re-evaluating Language Model Limitations and Long-Context Efficiency

Stanford University's Demonstration of Expired Past Weaknesses and Identification of Silent Training Loss in Gated Models

On November 3, 2026, researchers at Stanford University released results re-evaluating 50 claims of large language model (LLM) weaknesses repeatedly cited in past papers across 20 models from eight developers released between July 2024 and September 2026. With 44% of existing weaknesses effectively gone, 34% weakened, and an average lifespan of weakness claims of 15.7 months, they recommended not uncritically projecting flaws observed in older models onto newer models.

On November 4, researchers from Washington University in St. Louis, Nankai University, and others revealed a phenomenon where, during continual pretraining of gated models such as Mamba or xLSTM on multilingual or conversational data, representation values of tokens appearing rarely in text become excessively small, causing updates across entire batches to vanish to zero without any error warnings during the gradient clipping (a technique

capping update magnitude to prevent training divergence) step. When continually training Nemotron-3-Nano-4B on Telugu, updates vanished in 19.8% of total steps; the researchers resolved this using GroupClip and GateFloor methods.

Serving Cost Surges in Long Contexts, Variable State Tracking Limitations, and API Latency Reduction

On November 4, NVIDIA researchers analyzed that based on H100 server serving, when the input context window expands from 16,000 (16K) tokens to 256,000 (256K) tokens, the actual computational cost per token surges by 3.6 times up to 7 times depending on the model (based on Qwen 3 Coder Next), yet consumer pricing sheets obscure this. They estimated that costs could be reduced by 2.2 to 8.8 times through techniques compressing and summarizing conversations to continue into new sessions, but limitations also emerged where scores on some coding tasks dropped significantly due to the loss of detailed information during the summarization process.

StateBench, announced on November 3 by the Allen Institute for AI (Ai2), evaluated state-tracking abilities to overwrite, add, or copy variable values within long texts; even top cutting-edge models experienced a steep drop in accuracy to below 13% on addition and copy types as document length grew. Meanwhile, on November 4, Anthropic opened a public beta for Message Threads—which sends only newly added deltas rather than resending entire previous conversation histories on each API call—reducing the median time to first token for requests of 500,000 to 1,000,000 tokens by approximately 35%, though the total volume of tokens processed by the server and billing costs do not decrease, and it expires after 5 minutes of inactivity.

C. Device Operation Interconnection, Edge Inference Acceleration, and Visual Representation Learning

Tencent and Apple's Screen and App Action Control, and Nvidia's Database Ontology Release

On November 4, Tencent released the official version of Hunyuan Hy4, incorporating visual image understanding and direct computer screen operation (Computer Use) capabilities for the first time, applying it to productivity tools such as WorkBuddy, CodeBuddy, and Yuanbao. On November 3, Apple released a developer preview of App Intents—enabling third-party app developers to register in-app functions such as delivery tracking, reservations, and ordering with Siri artificial intelligence—in six languages including Korean (excluding the European Union and users under 13 years old; official App Store submissions begin in spring 2027).

On November 3, Nvidia open-sourced Auto Ontology, which automatically generates knowledge graphs from structural schema information in Databricks, PostgreSQL, Snowflake, and DuckDB to verify the consistency of natural language database queries (SQL). If a separate secure vault (Vault) is not connected, connection passwords may be stored in unencrypted plaintext, necessitating verification of security settings when deploying to enterprise intranets.

Edge Device Inference Acceleration, Microcontroller Strategy, and Visual Alignment Theory Research

In small-server acceleration, procedures were documented on November 4 running the Nemotron 3.5 Lightning model on Nvidia DGX Spark (64GB) hardware with TensorRT Edge-LLM (version 0.11.1), NVFP4 quantization (21GB), and DSpark speculative decoding (a technique where a small auxiliary model first predicts multiple next-word candidates and a large model verifies them all at once to boost output speed), increasing generated tokens per second by up to 1.73 times (from 73.4 to 126.9 based on HumanEval prompts) (figures are for single-request environments and cannot be used concurrently with multi-batch processing). On November 4, semiconductor firm STMicroelectronics announced

sample availability of the STM32A microcontroller family, bridging on-device artificial intelligence and real-time control across industrial equipment, automotive, and robotics, with mass-production availability planned for 2027.

In foundational representation learning, researchers led by Yann LeCun and Randall Balestriero announced on November 4 that a 109-million-parameter (109M) Joint Embedding Predictive Architecture (JEPA, a neural network architecture that predicts outcomes in core semantic representation space rather than generating entire pixels), trained on two observation views of the same scene, scored 0.184 on the WIT-1024 evaluation, exhibiting higher representation alignment with a vision model (DINOv2) than 70-billion-class large language models (Qwen2.5-72B at 0.154, Llama-3-70B at 0.151) (does not imply replacing sentence generation or logical reasoning).

13. Law, Policy, and Public Governance

A. Key adoption cases

In public administration and private app marketplaces, while government common artificial intelligence infrastructure that passed security reviews was officially launched for frontline civil servant tasks, regulations uniformly imposing adult age restrictions on emotional and sexual virtual companion apps were newly established to protect minors.

B. Pan-Government Administrative AI Launch and Platform Age Rating Regulations

Upstage Solar Pro 4 Officially Launches on Pan-Government Common Infrastructure and Apple Introduces Adult Age Ratings for Virtual Companions

Upstage officially launched service on November 3, 2026, by deploying its proprietary language model, Solar Pro 4, onto the government-wide AI common platform dedicated to administrative networks (interlinked with Samsung SDS FabriX), which completed a security review by the National Intelligence Service. Because the common platform centrally manages GPU computing resources, central government ministries and local governments can immediately use it for drafting responses to civil petitions, summarizing reports, and reviewing whether local ordinances conflict with superior laws, without having to build or lease separate individual servers.

On November 4, Apple decided to designate artificial intelligence virtual companion features that interact with users through sexual or romantic conversations as an independent review category in App Store age ratings, uniformly assigning them a minimum age rating of 18+. The corresponding

questionnaire item will be added to developer management tools starting November 16, and relevant app developers are required to complete their declarations by January 12, 2027.

14. AI Discussions and Legislation in Parliaments Worldwide

A. Key adoption cases

Parliaments around the world are moving beyond viewing artificial intelligence merely as a subject of legislative regulation to directly examine its practical impacts on legislative drafting as well as judicial and educational settings, with parliamentary activities surrounding leadership in international norms and multilingual safety evaluation cooperation also gaining momentum. Meanwhile, the U.S. Senate, House of Representatives, and major state legislatures were in an official Election Recess due to the November 3, 2026 midterm elections, so no hearings were held during this survey period.

B. Reviewing Field Applications in Legislation, the Judiciary, and Education, and Protecting Electoral Democracy

European Parliament Legislative Support Workshop and Italian Bicameral Fact-Finding Inquiry on Judiciary and Education

On November 4, 2026, the European Parliament's Panel for the Future of Science and Technology (STOA) held a workshop in Brussels on artificial intelligence for legislative work, attended by Chair Christian Ehler and European Parliament President Roberta Metsola. Sharing real-world parliamentary and public deployment cases, such as the Italian Chamber of Deputies' GENAI4LEX and Finnish Sitra's EurLexGPT, the session addressed methods for assisting in gathering citizen opinions, identifying legislative gaps, drafting and amending bills, and conducting ex-post evaluations. Participants clearly highlighted the limitations, emphasizing that while artificial intelligence can serve as a practical tool to enhance legislative clarity, it cannot replace parliament's democratic decision-making and must be subject to human oversight.

The Italian Parliament conducted sector-specific official fact-finding inquiries (Indagine conoscitiva) through standing committees of both houses. On November 4, the Senate's 2nd Standing Committee (Justice Committee) heard testimony from Laura Castelli, professor of civil law at the University of Milan, examining how the introduction of artificial intelligence in civil procedure rules and trial practice impacts legal liability; on November 3, the Chamber of Deputies' 7th Committee (Culture, Science, and Education Committee) examined the classroom application of AI in elementary, middle, and high schools, screen and social media dependency, and students' digital well-being (both are investigative phases aimed at identifying future legislative tasks).

Canadian House of Commons Warning on Election Chatbot Distortions and Peruvian Congress Forum on Public Administration

On November 3, the Canadian House of Commons Standing Committee on Procedure and House Affairs (PROC) held a meeting to study artificial intelligence and Canadian democratic institutions. Stéphane Perrault, Chief Electoral Officer of Canada, who appeared before the committee, warned that voters are increasingly relying on AI chatbots to look up party platforms and voting procedures despite the risk of inaccurate answers, and proposed transparency legislation mandating the labeling of election-related AI-generated content since voters lack tools to determine authenticity (this is currently a policy recommendation by the Chief Electoral Officer, and whether legislation will be passed remains undetermined).

On November 2, the Congress of the Republic of Peru, hosted by the office of Member of Congress Pierangeli Dodero Jovic in partnership with the World Organization for Democracy and a Purdue University research institute, held an Intelligent Governance Forum to discuss the adoption of artificial intelligence for public administration modernization, personal data protection, and responses to algorithmic bias risks, and shared the achievements of a free scholarship program in core technologies and AI for 1,000 youth.

C. International AI Governance Cooperation and the Korean National Assembly's Discussions on Industry, Regulation, and Copyright

UK House of Lords Multilateral Norms Hearing and China NPC's Deliberation on Treaty Establishing International Organization

On November 4, the UK House of Lords International Relations and Defence Committee held an oral evidence hearing on multilateral artificial intelligence governance as the second stage of its multilateralism inquiry. Professor Robert Trager of the University of Oxford and experts from the UN Foundation appeared as witnesses, recommending that the UK contribute to establishing global AI safety standards through international bodies and formulate a cooperation framework encompassing emerging economies (at the House of Lords committee inquiry testimony stage).

On November 2, at the first plenary session of its 25th meeting, the Standing Committee of the 14th National People's Congress (NPC) of China heard an explanatory briefing from Vice Minister of Foreign Affairs Hua Chunying, entrusted by the State Council, on a bill to ratify (a procedure whereby parliament gives final approval to an international treaty concluded by the government) the agreement establishing the World Artificial Intelligence Cooperation Organization (WAICO); on November 4, at the 75th Council of Chairpersons meeting presided over by Standing Committee Chairman Zhao Leji, it reviewed the draft approval decision and decided to submit it to the Standing Committee session (final confirmation of ratification depends on a subsequent plenary vote). Chairman Zhao Leji also held talks on November 2 with Valentina Matviyenko, Speaker of the Federation Council of Russia, who visited China, agreeing to strengthen exchanges of legislative experience in artificial intelligence and the digital economy.

Republic of Korea National Assembly's Discussions on International Evaluation Systems, Prohibitory Regulations, Physical AI Budgets, and Publishing Tax Credits

In the National Assembly of the Republic of Korea, seminars and forums hosted by lawmakers' offices between November 2 and 4 successively addressed agendas concerning international safety evaluation, high-risk regulation, regional robotics industry promotion, and copyright protection for training source materials (all remain at the stage of policy proposals and legislative/budgetary recommendations; whether they result in actual statutory amendments or budget allocations will be determined through subsequent parliamentary deliberations). In international cooperation, at a forum hosted on November 3 by the office of Rep. Kim Gunn, Professor Catherine Régis of the University of Montreal (co-director of the research program at the Canadian AI Safety Institute) proposed creating a multilingual joint AI evaluation framework combining Korea's application and industrialization capabilities with Canada's strengths in socio-technical safety research. In safety legislation, at a forum on revising the Framework Act on Artificial Intelligence hosted on November 4 by the office of Rep. Kim Sun-min and others, participants urged the creation of clauses prohibiting certain AI practices—akin to the European Union AI Act—to explicitly delineate red lines companies must not cross, such as human rights violations.

In industrial budget discussions, at a major forum on November 4 co-hosted by nine lawmakers representing the Jeonbuk region, including Chung Dong-young, An Ho-young, and Han Byung-do, plans were discussed to reflect five core Physical AI projects—valued at 1.3 trillion KRW, which perceive and act in the physical world combined with physical devices such as robots—into the 2027 national budget; at a seminar on November 3 hosted by Rep. Park Jie-won, Jeonbuk State, Hyundai Motor Group, and Daedong discussed an initiative to build an agricultural Physical AI robot farm hub valued at approximately 300 billion KRW on a 21-hectare (ha) site in the Saemangeum area. In addition, at a seminar on November 2 co-hosted by the offices of Reps. Jeong Seong-ho and Yoon Hu-deok and organized by the Korean Publishers Association, participants

voiced concern that the ecosystem of original records is shrinking due to unauthorized book training by generative AI and the proliferation of summarization services, urging the introduction of tax credits under the Restriction of Special Taxation Act for the publishing and web novel sectors, which have been excluded from benefits unlike film, video, and webtoons.

View references · Appendix, page 34

References

Select an entry to open the original source. The original reference order and links are preserved.

001 France tasks Naval Group to develop large oceangoing underwater drone

002 Saab unveils optionally crewed CB90 CX ahead of Euronaval 2026

003 Saab launches next generation 9LV naval combat system for connected crewed and uncrewed fleets

004 Turkiye's MKE debuts PIRANA armed USV variant with drone payload

005 Tekever and General Dynamics demonstrate AR5 drone's role in detecting and tracking underwater targets

006 Ukrainian Armed Forces and National Guard receive 626 Zirka interceptor drones from Prytula Foundation

007 Ukraine's NUMO Robotics unveils LOKO medium unmanned ground vehicle for logistics and evacuation

008 IFC 2026 Shield AI rolls out first X-BAT prototype

009 Oshkosh partners with startup to build military drones

010 US company Vantor demonstrates capability to supply satellite imagery to Ukrainian front line within 7 minutes

011 South Korea's MND and MOTIE initiate development of military RAM-C AI agent

012 Postman desktop app unauthorized Codex and Claude Code configuration issue

013 Free the Model Sell the Gateway Prompt Adware Hidden in Jailbreak Bundles

014 Apple warns of AI agents behaving in unintended or unpredictable ways

015 Provable Inference FlashProve

016 CWE-bench v2 Evaluating AI Agents on Real-World Vulnerability Patching

017 KierstonBench Synthetic City Investigation Benchmark

018 AI agents take over at SETI's Allen Telescope Array scanning for extraterrestrial intelligence

019 Xpeng unveils GroundingPI pre-trained model

020 China explores humanoid robots for future battlefield operations

021 OpenAI says GPT-6 Astra designed a hardware device from scratch in under a week

022 Baidu launches Apollo World model to train robotaxis in simulated environments

023 British fintech Revolut slashes AI coding costs 10x by going self-hosted

024 Palette Deck Building with Small Open-Weight Models

025 Anthropic Claude Code 2.1.317 allows creation referencing game mechanics

026 Snowflake expands agentic AI platform with business context and governance tools

027 Anthropic Claude Code v2.1.317 Release Notes

028 Anthropic Claude Code v2.1.318 Release Notes

029 ToolSpace Synthesizing Agent Learning Environments from Real MCP Tool Calls

030 New open-source benchmark reveals reasoning as the biggest bottleneck for AI in cloud operations

031 Huawei applies openPangu-2.1-Pro to Xiaoyi Work and terminates 2.0 service

032 ByteDance suspends new sales of Coze Enterprise and integrates into Doubao enterprise services

033 Rapid LLM development is outrunning the research meant to explain it

034 Silent Gradient Vanishing in Gated Models Continual Pretraining

035 Long Context Serving Cost and Session Compression Analysis

036 StateBench Tracking Variable State Changes in Long Contexts

037 Tencent Hunyuan Hy4 Official Release with Visual Understanding and Computer Use

038 Apple Developer News Siri App Intents Preview

039 NVIDIA Auto Ontology GitHub Repository

040 Running Nemotron 3.5 Lightning based on TensorRT Edge-LLM in DGX Spark environment

041 Comparison of visual alignment between 109M JEPA and 72B LLM based on identical scene observation

042 STOA Workshop AI for legislative work

043 Italian Senate 2nd Standing Committee agenda on fact-finding inquiry into AI impact in the judicial sector

044 Italian Chamber of Deputies 7th Committee meeting notice on fact-finding inquiry into new technologies and AI in schools

045 Chief Electoral Officer of Canada testifies before House Committee on AI risks to election information

046 Congress of the Republic of Peru official notice on Intelligent Governance and Public Administration Modernization Forum

047 Lords Committee inquiry session on multilateral AI governance

048 25th Session of the 14th NPC Standing Committee of China reviews agreement establishing World AI Cooperation Organization

049 Chairman of China's NPC Standing Committee meets with Chairman of Russian Federation Council

050 National Assembly seminar on AI safety, innovation, and international cooperation proposes Korea-Canada joint evaluation framework

051 National Assembly seminar on revising AI Framework Act to establish regulations on prohibited AI

052 National Assembly grand debate in Jeonbuk on becoming a global leader in Physical AI and robotics

053 National Assembly meeting on establishing Saemangeum Physical AI Robot Farm Factory

054 National Assembly meeting urging tax credits for publishing and web novels serving as AI training sources

Advertisement | **Custom In-Depth Report Writing**

Just provide a topic, and we will research all materials from around the world right up through today and compile them into a report (or book).

Service Information & Inquiries | **Sample Reports**