

2025

Hon Hai Research Institute Yearbook

2025 Hon Hai Research Institute Yearbook

2025

Hon Hai Research Institute
Yearbook



2025

Hon Hai Research Institute Yearbook



CONTENTS

Message from the Executive Director

- 04
- From Lab to Industrial Foresight: Powering Hon Hai’s Transformation into a Technology Engine

Message from the CEO

- 06
- Blazing the Trail: Building the Foundation for Hon Hai’s Technological Future

From the Advisory Committee

- 10
- When AI Truly Understands the World
From Technical Evolution to Outcome-Driven Organizational Transformation
Kai-Fu Lee
- 12
- Beyond Moore’s Law: A New Semiconductor Landscape Shaped by Glass Substrates and Heterogeneous Integration
Mau-Chung Frank Chang
- 14
- The NISQ Era: Driving a Practical Revolution with Qubits
Ching-Ray Chang
- 16
- Trust and Resilience in the AI Era Building Zero-Trust Defenses from Edge to Supply Chain
Tsung-Cheng Wu
- 18
- Building Intelligent, Energy-Efficient, and Resilient Networks in the 6G era
Wan-Jiun Liao

Powering the Future of Innovation

- 22
- Artificial Intelligence Research Center
Mastering Technological Autonomy and Sustainable Computing to Drive Transformation in EVs and Smart Manufacturing
- 26
- Semiconductor Research Center
Advancing Integration and Novel Materials to Reshape Next-Generation Standards
- 30
- Next-generation Communications Research Center
From PEARL to 6G Satellite Networks
Establishing Strategic Leadership in Next-Generation Communications
- 34
- Trapped-Ion Quantum Computing Laboratory
From Precise Control to High-Efficiency Readout
Paving the Way Toward a Scalable Commercial Future
- 38
- Quantum Computing Research Center
Transcending Fault-Tolerant Bottlenecks:
Architecting the Commercial Future via First Principles
- 42
- Information Security Research Center
Forging a Global Collaborative Defense Deepening AI Security and Post-Quantum Resilience

Cross-Domain Intelligence Power of Strategic Convergence

- 48
- City University of Hong Kong × Artificial Intelligence Research Center
Advancing Cross-Domain AI Reasoning and Learning
- 49
- NYCU × Semiconductor Research Center
Reshaping Data Transmission Architecture for the AI Era
- 50
- National Central University × Next-generation Communications Research Center
Building Taiwan’s Space Autonomy
- 51
- NYCU × Trapped-Ion Quantum Computing Laboratory
Overcoming Optical Limits with Non-local Metasurfaces
- 53
- University of Tokyo × Quantum Computing Research Center
Defying the Resource Curse: Architectural Innovations in Quantum Computing
- 55
- City University of Hong Kong × Information Security Research Center
Engineering a Resilient Ecosystem for AI Security

Moments of Synergy

- 58
- 2025 Chronology

About Hon Hai Research Institute

The establishment of Hon Hai Research Institute marks a key milestone in Hon Hai Technology Group's strategic transformation toward the F3.0 era. The Institute comprises five research centers and one laboratory, each staffed by an average of 40 top-tier researchers dedicated to exploring cutting-edge technologies poised to shape the next three to seven years. Its mission is to strengthen Hon Hai's innovation in both technology and product development, advancing the company from a labor-intensive to a knowledge-intensive enterprise and enhancing the Group's core competencies under its 3+3+3 development strategy.

Publisher: Young Liu	Editor-in-Chief: Wei-Bin Lee
Published by: Hon Hai Research Institute	Executive Editors: Vanessa Shih, Victor Lai
Address: 32 Jihu Road, Neihu District, Taipei City, Taiwan	Website: https://hhri.foxconn.com

Planning and Production: Commonwealth Publishing Co., Ltd., a division of Global Views - Commonwealth Publishing Group
Written by: Linden Chen
Revised by: Albert L. Chang
Translated by: Howard Lu
Photo Credits: Shutterstock; Hon Hai Research Institute
Project Managers: Nomo Hu, Mavis Luo, Ethel Lo, Alex Lee
Address: 1, Lane 93, Songjiang Road, Zhongshan District, Taipei City, Taiwan
Website: <https://bookzone.cwgv.com.tw>
Publication Date: September 7, 2026

From Lab to Industrial Foresight: Powering Hon Hai's Transformation into a Technology Engine

Reflecting on the past few years, Hon Hai has arrived at a critical inflection point. We are rapidly transforming from a globally recognized manufacturing company into a leading, technology-driven enterprise. At the very heart of this metamorphosis is the Hon Hai Research Institute, the catalytic engine propelling this transformation.

Rooted in Technology: The Institute's Strategic Mission

The mission of the Hon Hai Research Institute is to build the core technological capabilities that the Hon Hai Group will require over the next three to seven years. Our research directly supports the Group's forward-looking strategies, spanning artificial intelligence, semiconductors, next-generation communications, quantum technology, and cybersecurity. The Institute not only supports internal R&D but also serves as an important platform for global collaboration and consultation, ensuring that Hon Hai continues to exert influence on the international technology stage.

In 2025, the Institute achieved a series of milestones. In artificial intelligence, we launched FoxBrain, a large language model (LLM) purpose-built to integrate industrial expertise, manufacturing practices, and engineering contexts. As the world's first AI LLM optimized for traditional Chinese, FoxBrain was developed through a more efficient and cost-effective training approach, setting a new milestone for AI development in Taiwan.

In quantum computing, I regard quantum technology as one of our most vital long-term investments. Quantum communications and quantum sensing are already demonstrating early value, while quantum computing holds the potential to reshape existing computing paradigms. Today, the global quantum landscape remains in the Noisy Intermediate-Scale Quantum (NISQ) era, where systems are constrained by noise and scalability challenges. The Institute is actively advancing ion-trap platforms to develop high-fidelity quantum operations and scalable architectures while also exploring hybrid quantum-classical frameworks to enable large-scale, hybrid quantum-supercomputing (or quantum-HPC) integrated applications.

Semiconductors remain at the core of the technology industry. The Institute is actively advancing heterogeneous integration and optoelectronic integration technologies while also exploring advanced packaging substrate materials to maximize chip integration density. By evolving semiconductors from simple, single-material computing components into system cores that integrate computation, photonics, and communications, we are strengthening our competitiveness in the global semiconductor value chain.

For next-generation communications, we are strategically advancing 6G Non-Terrestrial Network (NTN) technologies. Our initiatives span ground terminal equipment, satellite-to-ground testbeds,

Young Liu

Executive Director
Hon Hai Research Institute



and advanced simulation software. Looking ahead, we will continue to validate Inter-Satellite Link (ISL) capabilities and deepen collaboration with international partners. These efforts lay the foundation for Hon Hai to become an important partner to satellite constellation service providers, building a complete communications technology chain that extends seamlessly from the ground to space.

Information security is the foundation of trust in the digital era. As AI becomes deeply embedded across industries, we are advancing research in both AI for Security and Security for AI, not only mitigating emerging risks but also unlocking AI's full strategic value. Additionally, we are developing Post-Quantum Cryptography (PQC) to proactively address the cryptographic challenges that quantum computing will inevitably bring. Most importantly, we view security governance as inseparable from research and practice. From strategic design to Group-wide implementation, we are building the next generation of defense frameworks.

Our Three-Tier R&D Ecosystem

To maximize execution efficiency, Hon Hai organizes its R&D into three collaborative tiers:

1. Hon Hai Research Institute: Responsible for the Proof of Concept (POC), dedicated to solving fundamental scientific challenges and validating conceptual feasibility.
2. Group R&D Center: Translates scientific breakthroughs into viable product prototypes.
3. Business Units (BUs): Scale, commercialize, and manufacture these prototypes for the global market.

Take FoxBrain for example: The Institute demonstrated the scientific feasibility of the model; the Group R&D Center transformed it into a product-ready version; and finally, the Business Units carried out fine-tuning, post-training, and market commercialization.

Technological breakthroughs do not happen overnight; they are the result of compounding patience, long-term accumulation, and strategic foresight. The Hon Hai Research Institute exists to bridge the gap between forward-looking research and real-world industrial demand, transforming theoretical exploration into engineering capabilities and concrete applications. This is what fuels the Group's continuous innovation, sharpens our global competitiveness, and drives the entire industry forward. We take pride in seeing technology evolve into products and products mature into successful businesses. In creating new opportunities, we also hope that the benefits of technological breakthroughs are shared broadly so that collective effort translates into collective reward.

Blazing the Trail: Building the Foundation for Hon Hai's Technological Future

Since its inception, the Research Institute has carried the core mission of driving the Group's strategic transformation and evolution. Beyond its function as a technical research unit, the Institute stands as a cornerstone of the Group's 3+3+3 strategy. In an era defined by rapid technological progress and industry changes, our role is to begin exploration before technologies mature and to identify new directions before markets take shape. As our corporate flagship advances, the Institute serves as its vanguard, scouting the technological frontier and paving the way for the industrial landscapes of tomorrow.

Over the past few years, we have organically fostered an interdisciplinary research environment. Ongoing cross-fertilization and integration across fields have forged a closer link between theoretical research and system implementation. This collaborative approach enables the Institute to lay the groundwork for future industries even while key technologies are still in their infancy.

The Institute enjoys a structural advantage that is rare on the global stage. The Group's vast industrial ecosystem offers invaluable testing grounds and robust data, enabling us to validate technologies in real-world settings. This not only allows us to explore forward-looking scientific challenges but also gives innovative technologies the opportunity to enter industry, thus creating a sustainable competitive edge.

On the international front, we continue to strengthen ties with leading academic institutions worldwide and actively engage with the global scientific community. Through these collaborations, we not only drive technological innovation and knowledge exchange but also attract elite talent from across the globe. At present, about one-fifth of our researchers hail from overseas or collaborate remotely. This is a testament to the truly global research hub we are building, a shared platform where leading minds from around the world can come together to shape the future of technology.

In research, we work closely with the Chairman and advisory board to chart long-term technological roadmaps and strategic directions. Once priority areas are defined, researchers are given the autonomy to pursue their work. We firmly believe that smart people don't need to be told how to act; they only need clarity on what must be achieved. Guided by a culture that blends trust with discipline, the Institute continues to produce high-quality research outcomes to distinguish itself on the global stage.



CEO, Hon Hai Research Institute



Looking ahead, the Research Institute will continue to serve as a pathfinder, laying the groundwork before industry shifts occur, exploring directions before technologies take shape, and forging connections among global talent. Through sustained investment in R&D and cross-disciplinary collaboration, we are not only opening new technological frontiers for the Group but also guiding the broader industry toward future development. Amid the accelerating tide of technological change, we strive to deliver lasting value to both industry and society.

From the Advisory Committee

Let ideas gradually take shape, let technology push boundaries, together they shape the year's most powerful moments.



When AI Truly Understands the World From Technical Evolution to Outcome- Driven Organizational Transformation

Kai-Fu Lee

Chairman of Sinovation Ventures and CEO of 01.AI

In its early days, robotics advanced slowly because machines relied on rote commands and often failed when confronted with unfamiliar situations. In recent years, however, breakthroughs in world models have equipped AI with physical intuition and common sense. By simulating the outcomes of actions in advance, AI now functions as the brain of a robot, enabling it to navigate real-world environments with human-like adaptability. This breakthrough paves the way for fully autonomous production and personal assistants to become a reality.

Until world models are perfected to handle all unknown scenarios, factories remain the most practical entry point for AI into the physical world. Unlike the unpredictability of open environments, a factory can be designed as a setting optimized for robotic operation. This means AI need not contend with every variable at once but can instead achieve mastery within controlled contexts, specializing in the production of specific components with precision.

Although world models remain central to AI, developing these foundational systems requires immense capital and computing power, much like the winner-takes-all dynamics of the semiconductor industry. I anticipate that, in the long run, only about six global giants will have the resources to compete. Without massive user scale and sustainable returns, most companies will struggle to keep investing.

This division of labor mirrors the PC and mobile eras. While chipmakers and OS incumbents continue to enjoy substantial margins, the application layer will capture the lion's share of both profit and reach. Consequently, the future of AI will not lie

in superficial API calls. Success will require deep integration, where models are fine-tuned and technologies adapted to specific scenarios, fusing applications with models at a fundamental level.

Another important AI trend is that the smartphone will not be the ultimate form of AI-first hardware. The current phone architecture, which requires unlocks, app launches, and taps, creates seconds of delay that break the natural rhythm of human communication. Therefore, I believe the next generation of AI devices must include five defining attributes:

- 1 Voice-Driven: Voice is the most natural user interface, and the technology is now precise enough for everyday use.
- 2 Always-On: No unlocking or launching required.
- 3 Ambient Listening: Continuous environmental awareness, without reliance on wake words.
- 4 Infinite Memory: A digital second brain that logs and synthesizes everything seen and heard.
- 5 Near-Invisible: Ultra-miniaturized wearables: brooches, glasses, or bracelets. AI will shift from a passive executor of commands into an active analyst and a personal assistant capable of reasoning through the nuances of daily life.

However, the deeper impact lies in the logic of enterprise operations. I foresee 2026 as the inaugural year for large-scale adoption of multi-agent AI in corporate settings. Once AI combines physical intuition with executorial capability, its role within an organization will fundamentally change. AI will no longer be merely assistive software but an



employee in its own right. Its value will be measured by the tangible output it generates rather than API calls or subscription fees. For instance, at 01.AI, 80% of code is now AI-written, with human-authored code accounting for only 20%. While this has increased developer productivity by several-fold, the true value is not in cost savings but in the dramatic acceleration of time-to-market to maintain competitive leadership.

Under this trend, the enterprise model for AI adoption will formally move beyond the single-agent era and enter a new epoch of multi-agent collaboration.

Historically, companies have assigned single AI agents to isolated, repetitive tasks such as routine administrative work or expense reporting. However, the future battleground will be an enterprise's ability to build a dedicated silicon team optimized for its core business. Within a multi-agent architecture, specialized AIs with diverse expertise will communicate, reason together, and autonomously deconstruct the company's most complex cross-departmental operations.

This transition from an assistive tool to a virtual employee is, at its core, a restructuring of organizational interests. As AI agents inevitably replace certain human roles, they will unsettle the vested interests and job security of middle management, who may slow adoption out of fear of being replaced. To fully unlock AI's comprehensive potential, top decision-makers must dismantle organizational barriers from the top down. In doing so, they can elevate the enterprise from relying merely on tool intelligence to achieving higher-order competitiveness through organizational intelligence and industry intelligence.

Amid this wave of transformation, the Hon Hai Research Institute has been entrusted with a core

mission: to define the future. Its role must extend beyond technical support to pioneering the next generation of AI-native hardware. Apart from developing miniaturized devices with five key attributes and achieving a seamless three-dimensional integration of software, hardware, and AI, the Institute must leverage the Group's vast factories as the ultimate proving ground for Embodied AI. By developing world models that grasp physical laws and deploying pragmatic, high-efficiency AI agents such as wheeled chassis paired with robotic arms, the Institute will lead the push to carry AI from the digital into the physical world. This is not merely forward-looking R&D; it is the vital strategic fulcrum that ensures Hon Hai can leap from contract manufacturing to defining product standards and mastering the industry's core rules in the AI era.

Beyond Moore's Law A New Semiconductor Landscape Shaped by Glass Substrates and Heterogeneous Integration

Mau-Chung Frank Chang

Distinguished Professor, Department of Electrical and Computer Engineering, UCLA


With Moore's Law reaching its limits, heterogeneous integration has become mainstream in the semiconductor industry. By combining materials with vastly different properties into a single package, engineers are overcoming the physical limits inherent in monolithic substrates. Notably, the most significant breakthroughs are expected to come from material innovations in packaging substrates.

Currently, the mainstream CoWoS (Chip-on-Wafer-on-Substrate) packaging uses silicon wafers as interposers. However, the physical boundaries and circular geometry of a standard 12-inch wafer constrain the total number of dies and layout flexibility within a single package. In addition, silicon interposers are limited by the reticle size of photolithography tools, a constraint that prevents unlimited expansion on a single silicon substrate.

Within the next one to two years, packaging substrates are expected to shift toward glass.

Drawing on display panel technology, glass substrates can scale to 65 inches, 85 inches, or even larger. This breakthrough overcomes the 12-inch wafer limitation, enabling higher-density routing and larger-scale integration. The critical challenge for glass substrates, however, lies in coefficient of thermal expansion (CTE) matching. Specialized glass must be developed with thermal expansion rates aligned to both silicon and compound semiconductors to ensure structural stability and prevent warpage.

The role of silicon is shifting in the era of heterogeneous integration. It is becoming a primary mechanical backbone rather than the sole driver of component performance. Most notably, silicon cannot emit light. By using heterogeneous integration to embed light-emitting materials such as indium phosphide (InP) or gallium nitride (GaN) directly into silicon-based circuits, chips can replace electrons with photons to enable high-speed optical communication.



Heterogeneous integration has emerged as a defining semiconductor trend, with major breakthroughs expected in packaging substrate materials.



To achieve optoelectronic integration and address the efficiency and speed requirements of future AI computing, light-emitting materials such as indium phosphide (InP) or gallium nitride (GaN) will be integrated with CMOS logic circuits at the wafer level. This integration is expected to become mainstream within the next two to three years.

Currently, most packaging technologies place the laser source off-chip. However, lasers are extremely temperature-sensitive; a fluctuation of just 2°C to 3°C can cause severe wavelength drift and lead to communication failure. Solving this requires true on-chip optoelectronic integration, with the light source built directly into the digital circuitry. This requires embedding on-die micro-thermometers and temperature controllers to achieve precise, localized thermal regulation, rather than relying on ambient cooling.

In the 6G and sub-terahertz era, optoelectronic integration will be pivotal. Future data transmission will move away from traditional copper interconnects, which cause severe loss and generate heat, toward lighter and more efficient media. This involves using silicon photonics for intra-chip signaling and polymer waveguides for inter-system sub-THz transmission. Although the materials used and frequencies differ, the fundamental task of data transmission remains the same. In the future, this will be achieved by replacing electrical signal transmission with optical transmission, thus harnessing the speed of light to deliver the extreme bandwidth and low latency required in the 6G era.

Traditional copper wiring is also prohibitively heavy. In robots or electric vehicles, carrying dozens of pounds of copper cabling would severely restrict mobility and efficiency. Coupled with rising

copper costs driven by geopolitical pressures, replacing copper with advanced waveguides for sub-THz communication to achieve system lightweighting is the inevitable path forward.

However, silicon wafers remain highly capable of computation. Therefore, to produce the miniaturized, high-performance chips required for the 6G era, it is necessary to integrate compound semiconductors optimized for high-frequency signals with silicon logic chips that excel at logic and computation. Because 6G signals are extremely sensitive to interference and temperature variation, glass substrates provide more than just additional space. Their superior electrical insulation and precise thermal control are critical to ensuring signal stability and preventing drift.

In the semiconductor industry, failing to sustain high yields is tantamount to failing to be profitable. While technological leadership is essential, the definitive competitive edge lies in achieving commercially viable yields amid the complexity of heterogeneous integration. This helps ensure an unassailable position in the face of competition.

It is recommended that the Semiconductor Research Center adopt a best junction for the function approach. By focusing on sensors and light-emitting components and boldly integrating compound semiconductors with optical technologies, the Center can leverage its superior system design capabilities to build a sustainable technological advantage. This approach will transform complex technical challenges into high-margin commercial opportunities, enabling it to secure an advantageous position within the global semiconductor value chain.

The NISQ Era: Driving a Practical Revolution with Qubits

Ching-Ray Chang

Distinguished Professor, Department of Physics, National Taiwan University

Director, Quantum Information Center, Chung Yuan Christian University

Chairman, Taiwan Association of Quantum Computing and Information Technology

Quantum technology is at a critical juncture, transitioning from scientific research into engineering development and industrial application. While businesses widely expect 2027 to be a pivotal commercial turning point, progress must be assessed based on hardware maturity.

The current quantum technology is supported by three pillars. The first is quantum communication, which has advanced most rapidly and has begun commercial deployment. The second is quantum sensing, which continues to deliver niche, high-precision utility in specialized sectors. The third, quantum computing, remains the most

broadly transformative yet technically challenging. Its promise lies in its capacity to transform the way traditional industries compute and unlock deeper mysteries of nature. Yet, because qubits are inherently fragile and achieving scalable quantum control is difficult, fault-tolerant quantum computation (FTQC) remains one of the most formidable bottlenecks of the modern era.

We currently operate within the NISQ (Noisy Intermediate-Scale Quantum) era, characterized by qubits that are noisy and error-prone, with errors accumulating as circuits deepen. While physicists strive for perfection, investing heavily to raise single-qubit reliability from 99.99% to 99.9999% or beyond, the R&D costs are prohibitively high. By contrast, the engineering community is taking a pragmatic path. This involves combining existing imperfect qubits to address problems solvable with today's technology, bringing products to market for feedback, and thereby sustaining industry growth.

In the NISQ era, quantum computers cannot handle mathematical or physics problems that demand 100% precision. Consequently, R&D focus has shifted toward application areas that are highly error-tolerant and where quantum computers can operate more efficiently than classical computers.

For instance, in the biopharmaceutical sector, quantum computers can help search through tens of thousands of molecular combinations to identify potential drugs. The objective is not necessarily to find a single optimal solution; even a second-best, usable formula carries significant R&D value. In financial investment, quantum computers need only generate a stable, profitable portfolio within an extremely short time frame, even if it does not



yield the maximum possible return, to realize a clear commercial advantage.

Quantum and classical computers, by complementing each other's strengths, can generate even greater benefits.

When the scale and complexity of computational problems exceed a certain threshold, classical computers falter under exponential growth in workload. By contrast, quantum computers possess the capacity to process massive datasets almost instantaneously. Once the two are seamlessly combined, it is akin to equipping a precise navigation blueprint with a jet engine. Enterprises will be empowered to tackle complex challenges in drug development or financial risk management ahead of full quantum maturity.

To realize this fusion of classical and quantum computing, a number of enabling technologies have emerged. One example is the development of quantum algorithms, such as quantum annealing or the quantum approximate optimization algorithm (QAOA), designed to search for optimal solutions within a large number of possible configurations. In these hybrid designs, a classical computer iteratively adjusts parameters to guide the quantum computer in finding correct answers despite noise interference, forming an archetypal model of quantum-classical convergence.

While this approach can readily handle small-scale problems, achieving enterprise-level applications requires the ability to process large-scale data.


Currently, academic researchers in Taiwan have successfully applied these algorithms to portfolio analysis using three decades of financial data.

Research shows that as portfolio complexity exceeds 20 assets, traditional computers struggle to perform real-time risk management. By contrast, quantum algorithms have demonstrated in historical simulations the ability to generate portfolios with higher returns and lower risk than the market. Financial institutions are already in discussions to commit capital for pilot testing, and if quantum computing proves stable, the prospect of automated quantum trading could become a reality.

Furthermore, spintronics, leveraging the properties of electron spin to develop low-power, highly stable hardware, can provide energy-efficient, cryogenic-resilient control systems. These systems help quantum hardware and classical computers integrate and operate efficiently, making spintronics a key enabler for large-scale quantum-classical convergence.

Although mainstream hardware has yet to fully adopt spintronic technology and commercialization costs remain high, a growing number of startups have begun to explore this path.

Against this backdrop, the Quantum Computing Research Center and Trapped-Ion Quantum Computing Laboratory have overcome geographical barriers to recruit top global talent and strengthen their international competitiveness. Looking ahead, I recommend deeper dialogue with Hon Hai Technology Group's business units so that research directions align more closely with corporate strategy. This would enable theoretical research to be transformed into the Group's competitive edge.



By finding optimal solutions within vast combinatorial spaces, quantum algorithms have made quantum-supercomputing convergence, the synergy of classical and quantum computing, a reality.

Trust and Resilience in the AI Era Building Zero-Trust Defenses from Edge to Supply Chain

Tsung-Cheng Wu

Distinguished Professor, Department of Information Management, National Taiwan
University of Science and Technology

At this historic turning point marked by the rapid rise of generative AI and the advancement of quantum computing, deepfake technologies and automated attacks have left traditional perimeter defenses increasingly inadequate. Breakthroughs in quantum technology also directly threaten the existing foundations of digital trust. As a member of the advisory committee, I must emphasize that trust and resilience are the absolute cornerstones for enterprises to survive amid these dual threats of disruption. In this context, building a Zero Trust defense that spans edge computing, supply chains, and post quantum cryptography (PQC) has evolved from a technical option into an unavoidable

operational strategy. Zero Trust is not merely a pursuit of rigid, absolute security; rather, it is a high-level risk-management framework whose core lies in empowering enterprises with highly sensitive anomaly detection and robust risk-control capabilities.

Traditional defense mechanisms risk becoming ineffective in the AI era. Because AI operates like a black box, hackers can easily bypass firewalls through logical manipulation, so defense must shift from static blocking to dynamic assessment. However, if every system behavior is subjected to exhaustive monitoring, enterprise efficiency will



A decentralized defense network prevents threats from proliferating across vast supply chains, establishing true operational resilience.



inevitably collapse. Administrators therefore need to introduce decision support based on behavioral analysis, elevating simple connection monitoring into intelligent risk inference that can precisely detect anomalies without disrupting operations. This is a concrete example of effectively reframing complex technical protocols into management friendly language, bridging communication gaps across departments.

To ensure that the Zero Trust framework is lightweight and streamlined, defense must be advanced to the front, making edge computing critical. In the past, endpoint devices were passive clients with limited computing resources. Today, with increasing processing capacity across smart manufacturing and vehicle-to-everything (V2X) systems, every endpoint can evolve into an autonomous, self-defending node. When every link in the supply chain is equipped with real-time edge defense, a breach at a single node can be instantly detected, effectively blocking further lateral attacks. This decentralized defense network prevents risks from spreading unchecked across vast supply chains, thereby achieving true operational resilience.

While building upper-layer applications, we must not overlook the disruptive challenges facing base-layer cryptography. Confronted with the harvest now, decrypt later (HNDL) cybersecurity threat, current encryption algorithms risk sudden collapse. Preemptively deploying post-quantum cryptography (PQC) is not merely an R&D task but a top-level corporate risk-mitigation strategy. We must seamlessly integrate PQC into edge devices, and this is a challenge that demands performance breakthroughs under limited resources. I am pleased to see our team's recent success in optimizing PQC computational efficiency. This

achievement demonstrates the robustness of our cryptographic foundation. Only with a secure base in cryptography can upper-layer AI applications and Zero Trust architectures deliver genuine defensive power.

Ultimately, technology without top-tier cybersecurity governance and supply chain trust is nothing more than a castle in the air. Drawing inspiration from Michael Porter's value chain theory, we must treat cybersecurity as a value chain in its own right. Defense mechanisms must be integrated as a standard feature of every production phase. Automated governance processes can then transform security from a mere cost burden into a core competitive advantage that earns trust from global clients. Leadership must craft a precise investment portfolio that balances baseline protection with advanced defense for sensitive data. In technology R&D, a vast gap often separates academic innovation from commercial deployment. Academia pursues single-point excellence, while industry demands cost-effective, practical solutions that resolve pain points. I highly commend the Hon Hai Research Institute's unique strength. Its team not only advances PQC and AI security but also wields the industrial capability to translate complex academic research into impactful, market-ready solutions.

I particularly admire the Chairman's strategic vision. He shares some of the research outcomes with the industry to build an ecosystem while keeping the most critical, proprietary technologies as strategic assets. This dual track strategy of open collaboration while consolidating the company's core competencies is the key to sustaining global leadership and building the strongest defense against the security challenges posed by AI and quantum computing.

Building Intelligent, Energy-Efficient, and Resilient Networks in the 6G era

Wan-Jiun Liao

Distinguished Professor, Department of Electrical Engineering, National Taiwan University

As global communications transition from 5G toward 6G, the core network design philosophy is undergoing a fundamental paradigm shift. 6G is no longer defined by a mere linear performance improvement in bandwidth, speed, or latency. Instead, it represents an overall evolution toward networks that are intelligent, resilient, and sustainable. Within this framework, key technologies such as AI-native networks, Non-Terrestrial Networks (NTN), semantic communications, and Integrated Sensing and Communications (ISAC) are collectively shaping the development of next-generation communication systems.

In the 5G era, artificial intelligence primarily functioned as an add-on optimization tool, mainly used in resource allocation and traffic prediction. At the core of 6G, however, lies the concept of being AI-native, where AI is embedded directly into the network architecture as the central decision-making engine. Future networks will feature continuous learning and self-evolution, enabling them to sense environmental changes in real time, acquire and track Channel State Information (CSI) and traffic patterns, and make proactive adjustments. This transformation advances networks from traditional reactive control to proactive control, enabling self-optimization, self-healing, and early risk detection, thereby significantly enhancing system resilience and service reliability.

Building upon this foundation, 6G further equips networks with sensing capabilities. By sharing hardware and spectrum resources, communication and sensing are deeply fused into an Integrated Sensing and Communications (ISAC) architecture, enabling simultaneous data transmission and environmental sensing. For instance, Intelligent Transportation Systems (ITS) can track real-time vehicle speed and

direction, dynamically adjust beams, and reallocate resources to ensure connection persistence under high-mobility conditions. In smart manufacturing, multiple robots can collaborate through shared sensory data to maximize both efficiency and safety. This convergence transforms the network from a simple transport layer into an intelligent system with contextual awareness.

AI deployment is also redefining the essence of communication. With the rise of generative AI, reasoning-based AI, and multimodal applications, communication is increasingly shifting toward information exchange led by AI agents. Against this backdrop, semantic communication has emerged as a critical technology. Unlike traditional communication models that prioritize bit-level accuracy, semantic communication focuses on transmitting only the information that is relevant to a specific task. By enabling AI to interpret data content and context, the system can extract the core semantics required for decision-making. This task-oriented communication model significantly reduces data volume and spectrum usage while directly improving energy efficiency.

Apart from the transformation of communication models, the network domain itself is expanding. One of the major developments in 6G is Non-Terrestrial Networks (NTN), which integrate low earth orbit (LEO) satellites, unmanned aerial platforms, and terrestrial base stations into a Space-Air-Ground Integrated Network (SAGIN). This design not only overcomes the limitations of terrestrial coverage but also supports high-mobility and wide-area connectivity. However, the highly dynamic topology introduced by NTN renders traditional connection management and handover mechanisms inadequate. As a result, AI-native mechanisms become

the critical foundation, requiring the system to instantly synthesize satellite orbits, user mobility, and spectrum conditions to enable cross-layer and cross-domain decision-making. This is especially critical between LEO satellites and ground terminals, where both sides are in rapid motion and the connection window is extremely short. By employing predictive AI models, the system can anticipate the optimal connection timing and execute path reconfiguration and handovers within milliseconds, achieving near-seamless cross-network connectivity.

Driven by these technological advances, the energy footprint of 6G systems has become a critical concern, making energy-efficient design a paramount priority. Broadly, this can be divided into two approaches: AI for Green and Green for AI. AI for Green utilizes AI to optimize network operations to reduce aggregate energy consumption, such as dynamically shutting down idle resources or adjusting power distribution based on real-time traffic loads. Green for AI focuses on reducing the computational cost of AI itself, including the development of lightweight models and high-efficiency inference architectures. In addition, offloading energy-intensive model training to the cloud or core nodes while deploying ultra-low-latency inference at edge devices is an important strategy for maximizing energy efficiency.

Driven by rising energy and computational demands, extending AI data centers into outer space has emerged as a new frontier. Space AI data centers harness solar energy and the naturally low temperatures of space for passive cooling, thereby reducing the energy burden on terrestrial facilities while integrating seamlessly with the NTN framework. Looking ahead, a space-ground collaborative computing model may evolve, in which specific model training or inference tasks are executed on

satellite nodes and then relayed back to terrestrial systems via high-efficiency communication links. Nevertheless, this architecture faces significant challenges that require further research and validation. These include computing offload, data synchronization, and coordination over high-latency links.

Ultimately, the evolution of 6G is not defined by a single, isolated technological breakthrough but by the highly synergistic advancement of multiple critical technologies. AI-native networks provide intelligent decision-making capabilities; NTN expands connectivity beyond terrestrial limitations; semantic communication reshapes data transmission; ISAC equips networks with environmental sensing; and space AI datacenters redefine computing architecture. Together, these technologies converge on a single, unified objective, which is to build a self-learning, autonomous, and energy-efficient intelligent resilient network. In the future, communication systems will no longer serve merely as information-delivery infrastructure but as the core platform for the operations of a highly intelligent society, striking a masterful balance between extreme performance, unwavering reliability, and long-term sustainability.

Amid these trends, the Hon Hai Research Institute has already established a robust foundation in AI, LEO satellite communications, and electric vehicles. With a distinct advantage in cross-domain integration and system-level deployment, the Institute is well positioned to play a pivotal role in the development and application of key 6G technologies. In doing so, it will make substantive contributions to advancing human society toward a new stage of heightened intelligence and cyber-physical convergence.



Powering the Future of Innovation

Driven by a commitment to innovation, we drive continuous technological advancement, turning laboratory breakthroughs into real-world solutions and building momentum that accelerates industry transformation and shapes the future.



Artificial Intelligence Research Center

Mastering Technological Autonomy and Sustainable Computing to Drive Transformation in EVs and Smart Manufacturing

The AI revolution is sweeping across the globe. It serves as a core driver of Hon Hai Technology Group's 3+3+3 Strategy, which encompasses three emerging industries (Electric Vehicles, Digital Health, and Robotics), three key technologies (Artificial Intelligence, Semiconductors, and Next-Generation Communications), and three smart platforms (Smart Manufacturing, Smart EV, and Smart City).

In 2025, the Artificial Intelligence Research Center achieved landmark breakthroughs across three major research frontiers: FoxBrain, an internally developed reasoning engine attuned to Taiwan's sociolinguistic context; EV sensing technologies capable of precise prediction under extreme conditions; and Spiking Neural Networks (SNNs) that push the limits of energy efficiency. These three technologies not only underscore the Institute's competitiveness in leading international academic circles but also highlight the Group's deep commitment to pioneering frontier AI.


Developing FoxBrain: A Traditional Chinese LLM Capable of Reasoning

Amidst the fierce global competition in Large Language Models (LLMs), the Hon Hai Research Institute has chosen an ambitious yet strategically vital path: independently developing and training its own LLM, FoxBrain.

Behind this decision lies a highly pragmatic consideration: linguistic and cultural autonomy. Over-reliance on global models risks marginalizing Taiwan's unique linguistic and cultural nuances. Currently, most Chinese-language corpora used in international models are based on Simplified Chinese, resulting in insufficient coverage and frequent biases regarding Taiwan's localized knowledge, distinct vernacular, and core values. In addition, cybersecurity risks are a paramount concern. As enterprises seek to deeply integrate AI into their production workflows, relying on foreign cloud services necessitates transmitting core data to overseas servers. It is an unacceptable risk for a manufacturing sector that prioritizes patents and privacy.

The development of FoxBrain addresses the growing demand for on-premise deployment. Capable of operating entirely offline, FoxBrain safeguards core technologies while building unique, domain-specific datasets during training.

What truly sets FoxBrain apart from pre-2025 LLMs is its distinct reasoning capability. In the field of autonomous driving, traditional algorithms were largely rule-based and often defaulted to rigid emergency stops when faced with undefined scenarios. FoxBrain, however, functions much like a human brain, equipped with Chain-of-Thought (CoT) reasoning. For instance, if a vehicle en-



Launched in March 2025, FoxBrain, the first traditional Chinese AI large language model, excels in mathematics and reasoning.

counters temporary roadwork where traffic cones direct it across a double yellow line, a traditional model might freeze in place to avoid a traffic violation. FoxBrain, however, uses scene-based logical deduction to flexibly navigate around the obstacle, just as a human driver would. When predicting pedestrian behavior, it goes beyond merely tracking paths; it perceives pedestrian intent and preemptively executes defensive maneuvers such as slowing down.

Currently, FoxBrain's objective is to become the model that best understands manufacturing. By collecting images and operational data from physical robots on factory floors, FoxBrain is evolving into a Vision-Language Model (VLM). In the future, it will be capable of directly interpreting defect images on production lines or roadway visuals during vehicle operation.

In the medical field, our research team has leveraged Retrieval-Augmented Generation

(RAG) and Knowledge Graph technologies to successfully mitigate AI hallucination issues. This ensures that generated medical advice strictly adheres to clinical guidelines, enabling the model to serve as an intelligent co-pilot for physicians and helping bridge the information gap between doctors and patients.

Multimodal Sensing Breakthroughs for Autonomous Driving's Long-Tail Problems

Although AI has become a prominent field of study, the greatest hurdle to deploying autonomous driving in the real world lies in long-tail problems, which are extremely rare but potentially fatal extreme scenarios. In 2025, the AI Research Center redefined the boundaries of vehicle perception and prediction through two technologies that captured first place in international competitions.

The first breakthrough was demonstrated at the Waymo Open Motion Prediction Challenge. The Center's key to winning this competition was its multi-trajectory prediction capability. Human driving intentions are highly complex; at an intersection, a driver might proceed straight, turn left, or wait. The Center's algorithms can simultaneously predict multiple potential paths. If a sudden event disrupts the primary choice, the system can instantly switch to an alternative route to avoid danger.

An extended application of this technology is the creation of highly realistic virtual environments. By simulating a diverse array of dangerous driving behaviors, autonomous models can undergo tens of thousands of hours of rigorous training within a virtual world. This approach is not only safer and more efficient than physical road testing, but it also teaches the AI how to handle unpredictable traffic scenarios and chaotic road conditions.

To overcome poor visibility scenarios such as heavy rain, dense fog, and low-light nighttime conditions, the Center developed Guided Thermal Super-Resolution (GTSR), a breakthrough that earned first place in the Thermal Image Super-Resolution Challenge at CVPR 2025. This core technology fuses RGB imagery with thermal sensing to leverage the strengths of both modalities. Thermal imaging captures heat signatures, while RGB imagery provides fine details and textures, such as leaf edges and traffic light contours, enabling vehicles to accurately

detect and recognize targets even in completely dark environments.

In addition, the MARE technology, recognized at AAAI 2025, was specifically developed to solve the long-tail object problem that traditional perception models struggle to process. The Center pioneered the transformation of Tesla's Occupancy Network concept into the MARE algorithm, which features 3D reconstruction capabilities. By breaking the reliance on expensive LiDAR systems, the Center proved that a single camera, combined with advanced semantic scene understanding, can accurately determine drivable space in complex environments, creating a new paradigm for low-cost autonomous perception.

Furthermore, MARE technology can reconstruct full 3D scenes in real-time. By introducing the concept of a memory bank, the system can rebuild scenes from stored information even when objects are temporarily obscured. This single-camera Semantic Scene Completion (SSC) approach enables Hon Hai to deliver navigation performance comparable to high-end sensors at a fraction of the hardware cost, giving electric vehicle development a significant competitive advantage.

SNNs Grant Edge AI "Biomimetic" Endurance

While the world worries about AI becoming a power-hungry monster, the Center is

Quick Glossary

LiDAR

LiDAR is a remote sensing technology that uses pulsed laser light to measure the time of flight of reflected signals. It is widely used in autonomous vehicles, forestry, and archaeology.



MARE (right) combines single-image semantic understanding and 3D reconstruction to estimate drivable space, overcoming long-tail object detection challenges and reducing autonomous driving costs.

dedicated to developing Spiking Neural Networks (SNNs) that are as energy-efficient as the human brain. SNNs transmit signals using discrete spikes, making them far more power-efficient and biomimetic than traditional deep learning models such as Convolutional Neural Networks (CNNs) and Transformers.

In the past, efforts to shorten simulation times to lower energy consumption often led to performance degradation due to insufficient information and vulnerability to noise. Research published by the Center in *IEEE Transactions on Pattern Analysis and Machine Intelligence* introduces a fundamental solution: Dynamic Capacity Experience Replay. If the AI brain is likened to a learner, this mechanism acts as an intelligent memory manager, automatically filtering and retaining high-quality experiences while discarding noisy, low-quality data.

This innovation shatters a long-standing trade-off in the AI field: the notion that making hardware smaller and more energy-efficient inevitably leads to a severe compromise in learning performance. Through this mechanism, AI can learn stably under extremely low power consumption, which is crucial for Edge AI. Although local edge computing power cannot currently compete with the massive scale of tens of thousands of GPUs used by cloud giants, the Center is seeking new breakthroughs through a paradigm shift in algorithms. This means that future factory inspection robots, agricultural drones, or in-vehicle chips will be able to sustain longer periods of auto-

nous learning and decision-making without relying on cloud computing resources.

It is worth noting that the parameters of existing AI models typically become fixed once training is complete, meaning that learning new tasks requires retraining. However, the goal of SNN research is to achieve continuous learning and lifelong learning. In this approach, the model can evolve during inference, much like a human who continues to absorb new knowledge long after finishing a book.

Once SNNs can be scaled up to hundreds of billions of parameters and paired with dedicated neural network chips, the next generation of AI systems will become both intelligent and energy-efficient.

These three major technologies are deeply interconnected: FoxBrain provides advanced reasoning and decision-making, multimodal perception technology grants vehicles high-precision vision, and SNN technology ensures that all computations run efficiently under low energy consumption. Yung-Hui Li, Director of the Artificial Intelligence Research Center, emphasized that this research is not only aimed at enhancing the competitiveness of Hon Hai Technology Group's products and solutions but also at strengthening Taiwan's overall AI capabilities. It is about safeguarding technological autonomy, cultivating cutting-edge local talent, and collaborating across Taiwan's industries to secure a decisive voice in the global AI landscape.

Semiconductor Research Center

Advancing Integration and Novel Materials to Reshape Next-Generation Standards

As global demand for AI computing grows exponentially and electric vehicles and renewable energy grids become ubiquitous, semiconductor innovation is no longer merely about boosting raw performance; it now requires breakthroughs in system-level integration and energy efficiency. In 2025, the Semiconductor Research Center achieved major advances across three major domains: high-speed silicon photonics, high-voltage silicon carbide (SiC), and nanophotonic integrated circuits.

Silicon Photonics: Building a Light-Speed Network in the AI Era

Today, as AI servers drive an explosive demand for bandwidth, traditional electrical signal transmission is approaching its physical limits. Once single-channel electrical transmission exceeds 100 Gbps, signal loss and interference drastically reduce communication ranges, potentially rendering data unreadable. Silicon photonics technology is the key to overcoming this barrier as it converts electrical signals into optical signals for efficient, high-performance transmission.

Developing silicon photonic chips rests on three key factors: design feasibility, scalable mass production, and rigorous testing and validation. To accelerate development, the Center has strengthened collaborations with simulation software providers and

semiconductor foundries, successfully bringing its proprietary silicon photonic chip designs to fruition. In 2025, it adopted the Multi-Project Wafer (MPW) model for its initial validation phase. By allowing multiple designers to share a single wafer space area, this approach cuts development costs to just 10% of a full wafer's price.

In addition, in 2025, the Center established an optoelectronic measurement platform supporting the 112 GBaud (224 Gb/s 4-level Pulse Amplitude Modulation, or PAM4) standard. In high-speed transmission, exceeding 200 Gbps on a single channel represents an exceptionally high threshold, which involves complex conversions from electrical to optical signals, optical transmission, and conversion back to electrical signals at the receiving end. However, this new platform can validate compliance with the PAM4 modulation standard, resolving issues of signal distortion and compensation. The establishment of this platform allows the research team to use empirical data to fine-tune chip design algorithms in real-time, effectively bridging the gap between design and manufacturing. Only after repeated iterations confirm that virtual simulations align consistently with real-world data can the entire design process be deemed accurate and reliable. The Center anticipates moving toward full-wafer tape-outs between late 2025 and 2026, offering greater latitude for design experimentation and pilot production.



The Semiconductor Research Center has established a comprehensive R&D framework for 1.2–6.5 kV SiC devices, encompassing simulation, fabrication, and system testing, thereby enhancing the Group's technological competitiveness.

In addition to silicon-based technologies, the Center has forged a deep partnership with Delft University of Technology in the Netherlands to explore indium phosphide (InP) chips. Because silicon approaches its physical limits as single-channel speeds near 200 Gbps, achieving cleaner signals and broader bandwidths necessitates the introduction of novel physical mechanisms or materials. The jointly developed InP chips can reach a maximum single-channel transmission rate of 320 Gb/s using PAM4 modulation. Ultimately, these high-performance InP modulators and lasers will be combined with silicon photonic chips through advanced integration techniques, providing next-generation data centers with solutions that surpass the constraints of Moore's Law.

Silicon Carbide: Driving a High-Voltage Power Revolution

If silicon photonics accelerates the communication speed of AI, silicon carbide (SiC) addresses the energy efficiency challenges of both AI and electric vehicles. As EVs migrate to 800V high-voltage platforms and the energy consumption of AI data centers soars, third-generation semiconductors, led by SiC, have become critical for energy

conservation. In 2025, the Center not only achieved breakthroughs in component voltage-tolerance specifications but also took a major step forward in integration.

The 1,700V trench-type SiC devices developed by the Center are designed to meet the demanding requirements of EVs and high-power industrial applications, with primary applications in EVs, AI data centers, and high-voltage direct current (HVDC) transmission. Compared to traditional planar components, the trench design significantly reduces on-resistance. For a chip of the same area, trench devices can handle higher current densities, thereby lowering overall energy consumption.

However, the trench structure is not without its challenges. Etching a trench into the component can cause excessive electric field concentration at the bottom, which risks damaging the gate oxide layer and compromising reliability. In response, the research team developed a dual-protection approach. They created a buffer zone at the bottom using ion implantation techniques while simultaneously thickening the bottom oxide layer to withstand high voltages and keeping the current-conducting sidewalls thin. In terms of on-resistance, this achievement far exceeds the technical standards

set by the National Science and Technology Council and establishes an internationally leading benchmark.

As for ultra-high-voltage scenarios, the Center has successfully developed 6,500V planar SiC devices. This technology is primarily targeted at national-level power grids and large-scale AI energy systems. In the future, hyperscale AI data centers may require dedicated substations to convert ultra-high-voltage power into medium or low voltage. During this process, the 6,500V SiC planar devices are exceptionally well-equipped to perform high-efficiency conversion and switching functions, thereby minimizing transmission losses.

In addition, traditional silicon chips fail at 150°C due to current leakage, necessitating complex heat dissipation and protection circuits. The SiC monolithic integrated circuit, developed in collaboration with National Yang Ming Chiao Tung University, overcomes this limitation to achieve stable operation in environments up to 300°C.

The core advantage of this technology lies in simplifying complexity. By directly integrating the driver circuit and power components onto the same SiC chip, it eliminates the impedance and inductance associated with traditional wire-bonding packaging. This results in faster system response times and a smaller overall footprint.


These SiC chips do not require complex cooling mechanisms and can maintain an operating frequency of 200 kHz even at high

temperatures. This not only reduces the chip area by 20% to 30% but also opens entirely new frontiers for aerospace applications, deep-well drilling, and next-generation industrial control. This achievement has already been presented at the International Symposium on Power Semiconductor Devices and ICs (ISPSD), a premier academic conference in the field of power electronics.

Nanophotonics: Reshaping 3D Sensing Technology

Having mastered speed and power, the Center turned its focus toward perception. In the field of smart devices and robotics, 3D sensing technology is central to a machine's spatial understanding. However, existing systems are bulky and power-hungry, making further miniaturization exceedingly difficult. Defying these constraints, the Center achieved a world first in 2025 with the monolithic integration of a metasurface and a Photonic Crystal Surface-Emitting Laser (PCSEL).

The current mainstream sensing light source in the market is the Vertical-Cavity Surface-Emitting Laser (VCSEL), but it suffers from inherent bottlenecks such as a large divergence angle and difficulties in developing longer wavelengths. Consequently, the Center opted for the more technically demanding PCSEL. By utilizing the nano-hole structure of photonic crystals to generate resonance, the PCSEL produces a small divergence angle to yield almost perfectly collimated light. Its spectrum is also exceptionally narrow, which eliminates



The world's first monolithic integration of a metasurface with a Photonic Crystal Surface-Emitting Laser (PCSEL) can be applied to monocular stereo vision and free-space optical transmission.

the need for an additional collimating lens and thereby addresses the volume issue at its source.

Traditional optical systems require collaboration among laser manufacturers, lens makers, and coating factories in order to stack multiple layers of optical components to generate a specific beam (such as the structured light dot projector used in Face ID). Through the Center's research, all these functions have been integrated into a single chip. This reduces the overall module volume by a factor of 2,480 and cuts energy consumption to just 70% of conventional solutions. In addition, the entire process is completed within a semiconductor fabrication plant, eliminating lens procurement and the complexities of optical alignment and packaging.

This technology provides robust momentum for monocular stereo vision. Currently, smartphone sensors, such as the front camera notch, remain on the millimeter scale. However, the integrated single-chip component measures only about 300 micrometers, roughly the width of three human hairs. If adopted in the future, the sensor could be completely hidden within

the screen bezel. This breakthrough was featured in leading international journals, including *ACS Photonics* and *Nano Letters*, and it received the Gold Award in the I-Zone Innovation Display Zone at Touch Taiwan, one of the world's most influential display industry exhibitions. Looking ahead to future wavelength selections, the Center has also developed 1,350-nm and 1,550-nm long-wavelength components that are eye-safe. Long-wavelength lasers are easily absorbed by water and will not penetrate the eye to damage the retina, safely allowing for longer sensing distances and higher power outputs.

Looking back at the Center's R&D advancements in 2025, despite spanning different domains, their core logic remains highly consistent. Through advanced integration and material innovation, they addressed the fundamental limitations of traditional systems: large size, high energy consumption, and slow transmission. Moving forward, the silicon photonics platform will enter the full-wafer tape-out phase, and SiC chips and PCSEL technology will undergo further optimization. The Center will continue to serve as a pivotal force, driving holistic progress in the era of AI and electric vehicles.

Quick Glossary

Monocular Stereo Vision

A computer vision technology that uses a single camera image and AI algorithms to estimate object distances and spatial structures, thereby reconstructing a 3D scene.

Next-generation Communications Research Center

From PEARL to 6G Satellite Networks

Establishing Strategic Leadership in Next-Generation Communications

In 2025, the global communications industry stood at a critical turning point. There was the transition from 5G to 6G and the expansion from terrestrial networks into Non-Terrestrial Networks (NTN). At this critical juncture, the Next-generation Communications Research Center achieved significant milestones. From the successful decommissioning of the first-generation PEARL satellites to advancements in second-generation inter-satellite links (ISL) and the international recognition of its core 6G algorithms, these accomplishments went beyond technical progress. They embodied the strategic foundation through which Taiwan reclaimed a defining role in shaping the future of global communications standards.

Building a Spaceflight Heritage for Taiwanese Enterprises

Since its inception in October 2021, the primary objective of the PEARL Project has been to build a spaceflight heritage for Taiwanese enterprises. To gain the trust of the international aerospace supply chain, it is essential to show that domestic satellites can operate reliably in orbit.

The first-generation PEARL satellites were 6U XL CubeSats, each about the size of a backpack and weighing approximately 8.8 kilograms. Launched aboard a SpaceX rocket on November 11, 2023, they executed

a series of in-orbit missions, from Launch and Early Orbit Phase (LEOP) operations to the successful deployment of solar panels. Both PEARL-1H and PEARL-1C remained in orbit for about two years and two months. Notably, PEARL-1H completed 12,151 orbits around the Earth while the ground station successfully sent and executed 4,701 commands.

These two satellites successfully demonstrated attitude control, subsystem health monitoring, and the collection of ionospheric parameters. These data are critical for understanding how equipment ages and evolves under the extreme conditions of space. Influenced by Earth's gravity, atmospheric drag, and heightened solar activity during the solar maximum, PEARL-1C re-entered the atmosphere on December 9, 2025, while PEARL-1H completed its mission over the South Pacific on January 12, 2026. Together, they enabled Hon Hai Research Institute to accumulate valuable end-to-end experience that spanned pre-launch integrated testing, in-orbit operations, and orbital decay.

Building upon the success of the first generation satellites, the Center pivoted its entire focus to the second-generation PEARL project in 2025. While the first generation focused on Earth-to-satellite communication, the core mission of its successor is the Inter-Satellite Link (ISL). By upgrading from the first generation's 1T1R (one-transmit,

Timeline: PEARL-1H / 1C Development and Milestones

Time	Mission
Phase 1: Project Initiation and System Development	
Oct. 2021	Project officially initiated; entered the System Design Review (SDR) phase.
Oct. 2021 – June 2023	June 2023: Completion of Preliminary Design Review (PDR), Critical Design Review (CDR), and subsystem module development.
Apr. 2023	Launch service contract signed, confirming the launch schedule and service provider.
May 2023	Communication license applications submitted to the National Communications Commission (NCC) and other agencies.
June 2023	Satellite assembly and testing; completed integration testing from individual units to the flight model.
Phase 2: Testing, Integration, and Rocket Launch	
July 2023	Thermal vacuum and vibration testing conducted at the Taiwan Space Agency (TASA) to simulate the harsh space environment.
Sep. 2023	Mechanical integration of the satellite and deployer completed.
Nov. 11, 2023	Key launch date (UTC).
Nov. 13, 2023	Successful ground communication; confirmed stable command transmission and reception by the ground station.
Nov. 28, 2023	Obtained NORAD ID 58265 from the North American Aerospace Defense Command, officially entering the space tracking registry.
Phase 3: Mission Execution and Completion	
Dec. 2023	Launch and Early Orbit Phase (LEOP); detumbling control and health status monitoring conducted.
March 14, 2024	Solar panels deployed; full-power operation mode initiated.
March 15, 2024	First photography mission; Earth imaging and selfie imaging verification completed.
2024 - 2025	In-orbit missions executed; ionosphere parameter collection and high-frequency communication tests conducted.
Jan. 12, 2026	Orbital decay: The satellite re-entered the atmosphere due to atmospheric drag, successfully completing its mission.

one-receive) architecture to a Ka-band 2T2R system, the satellite is transformed from an isolated node into an active network participant. It can now connect directly with satellites in similar or differing orbits, laying the critical groundwork for future mega-constellations.

The second-generation PEARL also integrates an S-band control system. By leveraging the broader bandwidth and enhanced stability of the S-band, it significantly improves the reliability of control commands. In addition, it is equipped with advanced perovskite solar cells and a Compact Ionosphere Probe (CIP). Currently, the PEARL-1A and PEARL-1B satellites have entered the functional integration testing phase, with launches planned for the second quarter of 2026.

Defining the Strategic ISL Architecture

In addition to physical satellites, the Center also developed two simulation software tools in 2025: OrbitCraft for top-level constellation architecture and BeamCraft for underlying network layout.

Low Earth Orbit (LEO) satellites travel at speeds of up to 7.5 kilometers per second, meaning a single satellite passes over Taiwan in a mere 6 to 8 minutes. How, then, can uninterrupted, full-time coverage be achieved? The answer lies in precise constellation orbit design. OrbitCraft simulates how various orbital parameters affect coverage.

For low-latitude regions such as Taiwan, designing constellations with specific orbital inclinations can optimize satellite contact frequencies. In addition, OrbitCraft can calculate the number of visible satellites from mega-constellations such as OneWeb or Starlink as they pass over Taiwan, providing a scientific basis for evaluating the island's communication resilience during wartime or natural disasters.

The distance between a satellite and the ground changes continuously. Because it is shortest at the zenith and longest near the horizon, satellite spot beams must be dynamically adjusted. BeamCraft addresses these variations by optimizing coverage based on orbital altitude and beam width, making it particularly well-suited for the high-gain phased array antennas demanded by future 6G networks.

A Three-Pronged Approach for 6G and LEO Satellite Technologies


In the field of 6G research and development, the Center has generated 45 high-impact research outcomes, including five top-tier journal papers and 34 conference papers. The Center has strategically positioned itself across three focus areas:

First, Integrated Sensing and Communication (ISAC) is widely regarded as the most anticipated feature of 6G. The Center has pioneered a Vehicle-to-Everything (V2X) sensing architecture based on Orthogonal

Quick Glossary

Satellite Contact Frequency

Satellite-to-ground connectivity depends on orbital altitude and type. Given their high speeds, LEO satellites pass over a single geographic point several times daily, with contact windows lasting from a few minutes to more than ten.



The second-generation Pearl satellite focuses on inter-satellite communication, enabling direct cross-orbit links to build future mega-constellations. Pictured: an on-orbit photo by Pearl.

Time Frequency Space (OTFS) modulation. By utilizing existing communication signals to detect objects, Roadside Units (RSUs) can monitor vehicle movement without the need for additional hardware. Furthermore, the Center's proprietary A2RMMP algorithm reduces computational time by roughly 30% to 40% compared to existing estimators.

Second, the Center actively promotes the concept of AI-enabled communications, embedding cutting-edge artificial intelligence algorithms into satellite systems and 6G development. By implementing federated learning, individual satellite nodes can collaboratively train models without exchanging raw data. This approach safeguards data privacy and security, conserves costly inter-satellite communication bandwidth, and optimizes energy efficiency across the constellation, ultimately extending the operational lifespan of the network.

Third, to address the physical layer challenges posed by the high-speed transit of LEO satellites, the team has applied deep reinforcement learning techniques to develop a self-learning synchronization and coding mechanism. Even in complex multi-user communication environments, this mechanism effectively mitigates frequency offsets and maintains precise signal synchronization.

In satellite communications, Doppler shift remains a formidable challenge. Given the extreme orbital velocities of satellites, high-frequency signals undergo severe drift. To counter this, the Center has developed a high performance Doppler compensation algorithm that stabilizes downlink signals while simultaneously enhancing positioning accuracy.

Trapped-Ion Quantum Computing Laboratory


From Precise Control to High-Efficiency Readout Paving the Way Toward a Scalable Commercial Future

With the global quantum computing race intensifying, the Trapped-Ion Quantum Computing Laboratory has, since its founding in 2021, focused on building a high-quality qubit control environment, with stable ion arrays as its central research objective. By 2025, the Laboratory achieved significant breakthroughs in hardware architecture, ion cooling, single-shot readout fidelity, and preliminary quantum manipulation. These milestones not only validate key technologies but also demonstrate that Hon Hai Technology Group now possesses the independent capability to seamlessly integrate quantum computing hardware and software.

As a hardware platform for quantum computing, trapped-ion systems utilize natural atoms as qubits. Because these atoms have inherently identical physical

properties, they do not depend on the precision manufacturing processes required by artificial qubits (such as superconductors or semiconductors) to improve yield. This eliminates the need for costly calibration of minute component variations and underscores a unique advantage in qubit quality.

Among the many quantum computing platforms, including superconducting qubits, photonics, and neutral atoms, what truly distinguishes trapped-ion systems is their exceptional coherence times and low error rates. Qubit coherence is a critical metric used to measure how long quantum information can be preserved without disruption from environmental noise. While superconducting qubits typically exhibit coherence times on the order of milliseconds, ion qubits can remain coherent for several seconds and, under specific



As a quantum computing hardware platform, trapped-ion systems demonstrate unparalleled technical advantages in qubit fidelity.

conditions, for several hours. This long-term stability is essential for executing quantum algorithms with complex structures and deep computational requirements.

Ytterbium-171 Ions and the Linear Paul Trap

One of the Laboratory's primary research achievements in 2025 was establishing a technological pathway that pairs a linear Paul trap with Ytterbium-171 (Yb-171) ions.

Ytterbium is a rare-earth metal in the lanthanide series, with two electrons in its outermost shell. The team vaporizes bulk ytterbium metal to form an atomic gas and, through photoionization, strips away one of

these outer electrons. This process produces a positively charged ion, which is then confined within the trap by an alternating electric field applied to the electrodes. At this stage, the ion retains only a single outer electron, simplifying its energy-level structure. This reduction in complexity lowers the demands and cost of laser control while enhancing precision in qubit manipulation.

The main reason for selecting the Yb-171 ion is its low sensitivity to external magnetic field interference. This level of technological precision parallels high-fidelity experiments such as atomic clocks, ensuring that the qubits remain stable and resistant to information errors caused by minute environmental fluctuations during computation.

Quick Glossary

Linear Coulomb Crystal

A stable physical structure formed by cooling multiple ions to extremely low temperatures. In this state, mutual Coulomb repulsion balances with the external electric field of the ion trap, holding ions at fixed intervals to serve as fundamental quantum computing units.

Doppler Cooling and Linear Coulomb Crystal

For an ion to serve as a viable qubit, it must first be effectively “frozen” in space. To achieve this, the Laboratory uses Doppler laser cooling technology to bring ions to near absolute zero. In this extreme ultra-cold environment, the Coulomb repulsion between ions balances against the external trapping fields, forming a stable linear Coulomb crystal structure.

Under an optical microscope, these trapped ion arrays resemble strings of pearls, with ions aligning sequentially and separated by several micrometers. This system currently demonstrates extraordinary stability. The Laboratory estimates that the ions can remain confined in the trap for several days or even longer, providing a robust platform for extended quantum experiments.

Achieving 99.67% High-Fidelity Readout

Readout is the final and most critical step in quantum computing. If computational results cannot be accurately extracted, even the most precise prior operations become meaningless. Experimental data shows that the Laboratory has achieved a single-shot readout fidelity of 99.67%. This indicates that State Preparation and Measurement

(SPAM) errors have been suppressed to exceptionally low levels.

In fluorescence detection histograms, the Laboratory observed a clear separation in photon count distributions representing the dark state (State 0) and the bright state (State 1). Initially, background system noise limited the optimal threshold to around eight counts. More recently, however, through optical system optimization, the Laboratory successfully reduced the optimal threshold to between three and four counts, enabling more reliable state discrimination.

This improvement stems from several technical advances. First, the signal-to-noise ratio (SNR) was enhanced by employing a high-numerical-aperture (high-NA) lens system, which increased the efficiency of capturing ion fluorescence. Second, scattered light was reduced by refining shielding and optical designs to prevent stray light caused by lasers striking trap hardware from reaching the detector. Finally, the use of photomultiplier tubes (PMTs) for high-sensitivity signal acquisition further strengthened the system’s performance.

Precise Coherent Manipulation and Rabi Oscillations

In addition to readout, the ability to

command ions to perform calculations is equally crucial. The Laboratory has successfully observed clear Rabi oscillations with exceptionally low decay rates. A Rabi oscillation is the process by which a qubit evolves between states 0 and 1 in response to the duration of a driving pulse. When this oscillation appears as a stable, low-decay sine wave pattern, it shows that the research team has achieved highly precise control over the evolution of the quantum state.

Successfully generating Rabi oscillations provides the critical foundation for precise single-qubit operations (such as the X gate and H gate). It demonstrates that the experimental system possesses long coherence times and provides a more generous fault-tolerance margin for implementing increasingly complex single-qubit logic gates in the future.

Charting a Course Toward Scalability and Chip Integration

In anticipation of the growing demand for increased qubit counts, the Laboratory has proactively designed multiple generations of hardware systems. The first-generation hardware architecture employed a single-zone blade design, primarily for basic ion capture. The second-generation system currently under development introduces a multi-zone structure built on silicon carbide (SiC) wafers with precision circuit designs. This innovative 2.5D structure, bridging conventional bulky 3D setups and fully planar chip integration, enables ions to shuttle between different zones. Ions designated for entanglement can be moved into dedicated computation zones to execute logic gates, while idle ions can be relocated to reduce interference with active qubits.

Looking forward, the objective for the third generation is a true chip-scale ion trap. Conventional systems rely on large optical tables to guide lasers, which severely limits commercialization. The Laboratory is integrating silicon photonics technology,

aiming to embed waveguides directly onto silicon substrates. This architecture would allow light to propagate internally within the chip and be precisely focused upward onto suspended ions.

Currently, in partnership with National Yang Ming Chiao Tung University, the Laboratory has successfully developed an on-chip integrated control light source. This technology employs silicon nitride waveguides to realize a non-local visible-light metasurface equipped with both phase and amplitude control. Ultimately, it is expected to replace bulky free-space optical setups, enabling high-precision, addressable qubit manipulation and laying the core technological foundation for the large-scale expansion and modular commercialization of trapped-ion quantum computers.

For the Laboratory, 2025 marked a pivotal foundational year, establishing a clear development roadmap. The primary focus for 2026 will be the deployment of a Raman laser system to implement critical two-qubit logic gates and ion entanglement experiments. The overarching goal is to unveil a programmable quantum computer prototype featuring 5 to 10 qubits by 2027.

“In 2024, we succeeded in trapping ions, but trapping alone does not grant control of the ions’ quantum states. In 2025, our efforts focused on mastering precise state control, ensuring that each ion functions as a high-quality qubit,” notes Guin-Dar Lin, Director of the Trapped-Ion Quantum Computing Laboratory. By continuously optimizing optical paths, rigorously suppressing environmental noise, and integrating semiconductor manufacturing processes, the Laboratory is advancing steadily toward the commercialization of quantum technology.

Quick Glossary

Optimal Threshold

The optimal threshold is the benchmark used to distinguish between a qubit’s bright and dark states. By analyzing the statistical distribution of photon counts, researchers determine the precise cut-off point that minimizes readout error. When the detected photon count exceeds this threshold, the qubit is identified as being in the bright state. By accurately setting this photon-count threshold, the Laboratory has minimized readout errors and achieved a single-shot fidelity of 99.67%.

Quantum Computing Research Center

Transcending Fault-Tolerant Bottlenecks: Architecting the Commercial Future via First Principles

As technology evolves, quantum computing is heralded as the next paradigm-shifting wave following artificial intelligence. However, while quantum algorithms promise exponential acceleration, the physical realization of quantum computers is hindered by noise and decoherence. Laboratory qubits remain inherently volatile; even infinitesimal environmental perturbations can lead to computational collapse. To achieve commercially viable universal quantum computing, the only path forward is fault-tolerant computing.

Historically, the primary bottleneck has been the prohibitive resource overhead required for error correction, which often exceeded the computational workload itself. If resource consumption scales exponentially with precision demands, quantum computers risk losing their advantage due to inefficiency. Addressing this critical challenge, the Quantum Computing Research Center unveiled a series of breakthroughs in 2025. These results not only theoretically demonstrate the feasibility of efficient error correction but also lay the groundwork for deploying large-scale quantum hardware.

Eliminating the Resource Escalation Nightmare in Magic State Distillation

A cornerstone of the Center's 2025 research

was resolving a fundamental challenge in quantum information theory: the efficiency ceiling of magic state distillation (MSD).

In classical computing, basic arithmetic like addition or basic logic gates is straightforward, whereas complex operations, such as matrix multiplication, require resources that scale quadratically or cubically with increasing bit count. Quantum computing faces a similar challenge. In a fault-tolerant framework, the standard solution is to deconstruct complex operations into a combination of simple operations and magic states. A magic state can be conceptualized as a pre-prepared, high-level resource that facilitates exceptionally costly operations within a quantum circuit, such as the T gate and the CCZ gate. Under a fault-tolerant architecture, these operations are prohibitively costly; without optimized magic state distillation, the hardware overhead required for error correction would become unsustainable.


However, magic states inevitably carry errors during preparation. They must undergo a distillation process in which multiple noisy magic states are purified into a single, high-fidelity resource. Earlier distillation methods suffered from a critical limitation. As the target error rate (ϵ) decreased, the resource overhead scaled as $\log(1/\epsilon)$ raised to the power γ . In prior state-of-the-art

techniques, the γ value was approximately 0.678, meaning that the pursuit of extreme precision triggered an explosive growth in the number of error-correction steps required.

The Center's research team, together with collaborators, successfully developed a completely novel magic state distillation protocol that reduces the γ value to precisely zero. Mathematically, when this exponent is zero, resource consumption theoretically approaches a constant. In other words, regardless of how high the required output precision is, the input resources needed to produce a single purified magic state will not grow infinitely. Since any value raised to the power of zero equals one, this breakthrough

demonstrates that distillation depth can remain within a controllable range, effectively removing a long-standing bottleneck that has constrained the scalability of quantum computing.

The key to this breakthrough lies in the principle of algebraic geometry codes. By harnessing their powerful algebraic structures, the research team constructed operations that support transversal logic gates in high-dimensional spaces and then mapped them back onto standard two-dimensional qubit architectures. This ensures that even in large-scale fault-tolerant quantum computing, the most energy-intensive components remain highly efficient. The research has generated significant impact



With efficient quantum error-correcting codes, quantum computers can sustain high-performance operations even when noise rates fall below specific thresholds.

in the international academic community, drawing citations from world-leading research teams such as IBM and Google.

Securing Quantum Advantage through Specialized Circuits

Beyond optimizing foundational resources, the Center has also investigated whether shallow quantum circuits can retain their computational superiority over classical computers even in noisy environments.

The Center's findings show that with high-efficiency quantum error-correcting codes, high-performance operations remain achievable as long as the noise rate stays below a critical threshold. More importantly, this research highlights the unconditional superiority of certain quantum circuits when compared with classical threshold circuits. This shows that while quantum computing requires additional steps for error correction, these measures do not erase its inherent advantage.

Assisted Verification: Surmounting the 100-Qubit Simulation Barrier

Verification is another major challenge in quantum computing. As the number of qubits increases, the state space expands exponentially. While 10 qubits correspond to 1,024 states, 100 qubits already exceed

the total memory capacity of all modern supercomputers. In this situation, how can we know whether the quantum computer's output is correct?

To address this challenge, the Center developed an AI-assisted circuit verification technique. Rather than brute-force simulation of all quantum states, the approach uses machine learning to capture the characteristics of quantum circuits. Using a small number of quantum computational samples, the research team trained AI models to learn the linear characteristics of the circuit. Once the AI mastered this structural logic, the model could predict future outputs on classical computers.

This method has its limitations; it is currently focused on verifying only certain circuit characteristics. However, it does provide an exceptionally efficient surrogate model. This model not only helps researchers more easily understand quantum circuit behavior but also offers a partial solution to the shortage of verification tools in hardware development.


This approach of overcoming hardware limitations through theoretical innovation is one of the core objectives of the Center's R&D, a philosophy rooted in First Principles.

In quantum computing, First Principles dictate that quantum computers must first be proven theoretically feasible and

Quick Glossary

Constant Overhead

Constant overhead refers to supplementary resources (such as time, space, or magic states) required during system operation, optimization, or fault-tolerant processes, where consumption is constrained within a fixed range regardless of problem scale, data volume, or target precision.



AI-assisted circuit verification trains AI on quantum circuit logic, allowing classical computers to directly predict quantum outputs.

cost-effective. If theory shows that error-correction resources are unsustainable, then no amount of capital invested in hardware development is likely to succeed. Currently, quantum hardware development remains in its infancy. Even with nearly a thousand physical qubits, the number of logical qubits actually available for computation is often in the single digits, leaving a significant gap between current capabilities and the scale required for commercial applications.

The Center's current research focus is to pave the way for future hardware development. It aims to prove that resource consumption can be maintained at a constant overhead, and this will give hardware manufacturers confidence that the success of quantum computers is mathematically feasible. Within just one year of publication, this series of breakthroughs has already accumulated over 40 citations, most from the leading

research teams in the field, demonstrating its profound impact on the industry's future trajectory.

The year 2025 was a landmark year for the Quantum Computing Research Center. From breakthroughs in constant-overhead distillation protocols to the implementation of AI-assisted verification, every milestone directly addressed the field's most critical challenges. Min-Hsiu Hsieh, Director of the Center, summarized, "The core significance of our work lies in deepening our understanding of fault-tolerant computing and proving to the world that quantum error correction does not erase the quantum advantage. We have theoretically cleared the obstacles to efficient quantum computing, ensuring this path can continue forward." As these theoretical foundations mature, quantum computing technology continues to take giant strides forward.

Information Security Research Center

Forging a Global Collaborative Defense Deepening AI Security and Post-Quantum Resilience


In the complex cybersecurity landscape of 2025, human-centric defense alone was no longer a viable safeguard to counter escalating threats. The Information Security Research Center is focused on mitigating risks from the malicious exploitation of AI and failures in AI decision-making. The Center's strategic initiatives range from pioneering Computing-in-Memory (CIM) architectures to overcome the performance barriers of Post-Quantum Cryptography (PQC) to exposing collision-detection vulnerabilities in autonomous driving simulators that had previously gone unnoticed even by official developers.

Safeguarding AI in Mission-Critical Environment

As generative AI and autonomous vehicle

systems move from experimental stages into mission-critical applications such as writing code, planning vehicle routes, and even making decisions, their inherent vulnerabilities are magnified. The Center has systematically revealed AI security risks in real-world scenarios and developed robust defense methods supported by quantifiable indicators.

First, the Center instituted a continuous defense cycle through rigorous red and blue team exercises. Traditional red teaming is often manual, cost-prohibitive, and limited in scope. In response, the Center developed an automated red-teaming framework driven by intrinsic motivation and autonomous exploration. This system can automatically identify vulnerabilities in black-box environments without requiring access to a model's



Hardware and software implementation of Post-Quantum Cryptography (PQC) Key Encapsulation Mechanisms (KEM) and Computing-in-Memory (CIM) technologies.

internal parameters. Research confirms that this approach increases the thematic diversity of adversarial datasets by 41%. In addition, to counter cross-modal jailbreak attacks, the Center engineered a blue-team suffix-based prompt defense module. This tool automatically generates protective text, reducing attack rates by 50% to 73%. These innovations have already been partially integrated into the FoxBrain platform as a frontline security layer.

In smart manufacturing, engineers use AI to assist in code generation. However, if AI-generated code contains hidden vulnerabilities or backdoors, it could lead to severe industrial equipment failures and major financial losses. Traditionally, patching such vulnerabilities requires costly model retraining, which may also degrade the model's original capabilities. To address this, the Center developed a training-free model

patching technology. By precisely adjusting internal safety parameters, this method enhances security by 10% without affecting existing functionality. In addition, to prevent attackers from embedding malicious triggers in training data, the research team developed a new algorithm that analyzes the unnaturalness and abruptness of text passages to detect anomalies. This approach successfully raised defense accuracy to 87%, providing more comprehensive cybersecurity protection for AI-assisted development.

Apart from smart manufacturing, the Center has invested heavily in the safety-critical field of autonomous driving, conducting comprehensive research into both perception and planning layers. Investigating system vulnerabilities, the research team revealed security concerns in online high-precision mapping models. Experiments showed that targeted lighting or physical patch attacks

Quick Glossary

F1 Score

The F1 Score is a comprehensive metric used to evaluate the classification performance of an AI model. It combines precision (avoiding false positives) and recall (avoiding false negatives), requiring both metrics to be high to achieve a strong overall score.

could reduce mapping accuracy by 9.9% and cause up to 44% of route-planning failures. To strengthen system robustness, the Center proposed a root-cause analysis method for multi-sensor fusion systems integrating LiDAR, radar, and cameras. When multiple perception modules fail simultaneously, this technique can pinpoint the fault source with an F1 score above 95%, enabling engineers to debug quickly and effectively.

In addition, the team conducted an in-depth analysis of CARLA, the industry-standard autonomous driving simulator. By developing a cuboid-based modeling detection method, they successfully identified 20 times more undetected collision scenarios than earlier approaches. This breakthrough was later confirmed by the official CARLA development team and assigned a CVE vulnerability ID (CVE-2024-33903). This achievement marks a major contribution to enhancing the safety of international autonomous driving simulation technologies.

PQC + CIM: Eliminating Memory and Energy Bottlenecks

Parallel to the rapid evolution of AI, quantum computing technology is also advancing at an accelerated pace. As quantum computing develops, traditional public-key cryptosystems face the risk of being compromised, making Post-Quantum cryptography (PQC) a critical safeguard for information security. The Center has focused its R&D on the latest

FIPS-203 ML-KEM standard released by the U.S. National Institute of Standards and Technology (NIST). However, the heavy reliance on high-degree polynomial multiplications within ML-KEM has become one of its key performance bottlenecks.

To overcome these constraints, the research team performed deep, hardware-oriented optimizations of the core Number Theoretic Transform (NTT) algorithm. By restructuring the underlying mathematical logic, the team successfully reduced computational complexity, laying a high-efficiency foundation for subsequent hardware integration.

In traditional von Neumann architectures, frequent data movement between the processor and memory causes latency and high energy consumption, a phenomenon commonly referred to as the memory wall. To address this, the Center collaborated extensively with the Industrial Technology Research Institute (ITRI) to introduce a disruptive Computing-in-Memory (CIM) architecture.

The core of CIM lies in embedding processing units directly within the memory array, enabling in-situ computation and eliminating the overhead of data movement. While CIM is an established architectural concept, the Center's application of it to accelerate PQC systems represents a pivotal transition from theoretical research to engineering implementation. Experimental results show that, within the same chip area, combining

NTT optimization with CIM architecture delivers a 14.22x performance increase over traditional software-based implementations.

This technology has successfully transitioned from theory to practical implementation. A Proof-of-Concept (PoC) of the ML-KEM algorithm was completed on the development board by SiliconAuto, an IC design subsidiary of Hon Hai Technology Group. This demonstrated the technology's ability to directly empower smart in-vehicle environments. In addition, FPGA validation confirmed that while pursuing high performance, the correctness of the cryptographic operations fully complies with the NIST FIPS-203 international standard. Both of these R&D achievements were officially unveiled at Hon Hai Tech Day 2025, garnering high praise from industry experts.

Global Cybersecurity Governance: Implementing a Collaborative Defense System

To address Hon Hai Group's cross-disciplinary and multi-site global footprint, the Center is committed to transforming cybersecurity defense into a driver of corporate governance, establishing information security as a pivotal element in the Group's sustainable development.

The foremost objective in implementing cybersecurity governance is cultivating executive accountability. Through the operations of the Cybersecurity Governance Committee, chairmen and presidents of the Group's subsidiaries have actively participated in core cybersecurity decision-making. Their involvement not only strengthens strategic execution but also ensures top-down alignment, effectively mitigating the communication friction and risks caused by inconsistent standards.

This year, the Center officially integrated the Cyber Defense Matrix (CDM) into the Group's operations, aligning it with the NIST CSF (Cybersecurity Framework) to

construct a comprehensive, visualized defense roadmap. This structured approach allows management to intuitively identify protection gaps and resource allocation status. By translating technical metrics into executive-level decision logic, the Center bridges the gap between engineering and management teams to ensure seamless communication. This mechanism has not only refined cybersecurity budget auditing but also catalyzed centralized procurement strategies, optimizing investment returns and enhancing defense resilience while reducing costs.

Regarding cybersecurity incident response, the Center has moved beyond traditional documentation and drills by co-developing a Cybersecurity Incident Response Dashboard with IT departments. This embodies the Code as SOP philosophy, transforming complex standard operating procedures into code-driven, automated workflows. This initiative significantly improves execution efficiency, prevents the omission of critical steps, and ensures a consistent, high-standard baseline throughout the response process.

Reflecting on 2025, the Information Security Research Center demonstrated a comprehensive strategy ranging from low-level structure to high-level strategy. Key breakthroughs were achieved in PQC, AI security, and cybersecurity governance. These advances not only boosted PQC performance but also safeguarded smart manufacturing and autonomous driving systems against AI threats while establishing a robust cross-regional governance framework. Wei-Bin Lee, CEO of the Hon Hai Research Institute and Director of the Information Security Research Center, emphasized that future strategic development will continue to focus on these three pillars. By accelerating research into PQC standards, strengthening automated AI defense, and ensuring global governance consistency, the Center aims to build a resilient security perimeter for the Group in an era shaped by the twin forces of AI and quantum technology.

Quick Glossary

FPGA

A Field-Programmable Gate Array (FPGA) is a semiconductor component that, even after being manufactured, allows users to repeatedly reconfigure its internal hardware circuit structure using code.

Cross-Domain Intelligence Power of Strategic Convergence

By partnering with academic institutions,
Hon Hai Research Institute demonstrates how cross-disciplinary
collaboration can move from theory to practice.
We bring AI, semiconductors, communications, and quantum
technologies into real-world settings.
Through repeated validation and refinement, we foster scalable,
deployable technologies to shape Taiwan's innovation for the future.



City University of Hong Kong × Artificial Intelligence Research Center

Advancing Cross-Domain AI Reasoning and Learning

As artificial intelligence moves beyond the laboratory to the unpredictable realities of autonomous driving and flexible manufacturing, it no longer faces standardized datasets. In the real world, AI must navigate an environment defined by sensory noise, extreme weather, and unpredictable behavior. To overcome these challenges, the City University of Hong Kong (CityU) and the AI Research Center have engaged in deep collaboration to develop next-generation agents capable of uncertainty reasoning and low-energy continual learning.

Equipping AI with Defensive Driving Logic

A core breakthrough of this collaboration lies in the field of uncertainty reasoning. Conventional AI models, when faced with adverse weather or sensor interference, often produce seemingly confident yet incorrect predictions that can lead to safety incidents. Our objective is to enable AI to adopt defensive driving logic similar to that of human drivers.

To achieve this, our team has integrated explicit uncertainty modeling at the perception layer. By dynamically weighting inputs from various sensors such as LiDAR and cameras, the system can autonomously determine which sensor is more reliable under extreme weather or during sensing failure. At the decision-making layer, the system no longer forces precise positional outputs but instead generates probability distributions with variance. As noise increases, the AI deliberately expands prediction variance, just as human drivers instinctively slow down and increase following distance in poor visibility. Through this mechanism, the agent tightens safety margins in the face of chaotic data, realizing decision logic closer to that of humans. Yet for AI to truly advance toward autonomy, it must also achieve continual learning.

Traditional models typically suffer from catastrophic forgetting, losing prior knowledge (such as basic traffic rules) when learning new skills (such as adapting to new road conditions). To address this challenge, the team has developed a growth strategy based on network positional importance. By adding new connections at highly important yet previously unlinked locations within the neural network, fresh knowledge can be stored without overwriting existing knowledge.

In addition, to reduce AI's massive energy footprint, the research team implemented Spiking Neural Networks (SNNs), a next-generation architecture that computes only when signals are received and remains silent otherwise, fundamentally eliminating redundant operations. This means that future edge devices will possess an energy-efficient brain, awakening to think only when necessary and making real-time decisions at ultra-low power.

In addition, the team has developed ModeSeq technology, which elevates trajectory prediction from mere stochastic selection to sequential reasoning across multiple potential behaviors. This gives AI a step-by-step cognitive ability, enabling it to interpret pedestrian gestures, understand the meaning of a horn, and transform raw data into contextual scene anticipation.

Laying the Bedrock for General-Purpose Agents

Hon Hai Technology Group is currently making a full-scale push into electric vehicles (EVs) and robotics. Professor Jianping Wang of City University of Hong Kong notes that the ModeSeq system can achieve high-frequency reasoning on vehicles at low power consumption, truly moving AI from the laboratory to production lines and mass-market cars. In addition, this methodology will also benefit factory robots, transforming them from passive executors into safe, trustworthy, and intelligent collaborators.

Moving forward, the partnership aims to extend the research outcomes into industrial applications. The objective is to enable robots in complex warehouses or campuses to autonomously understand tasks, identify risks, and collaborate with humans.

Professor Wang concludes, "If general intelligence is a skyscraper, our current work represents the bedrock." Through continual learning and robust reasoning, the collaboration between CityU and the AI Research Center is unlocking limitless possibilities for AI systems designed to operate long-term and evolve continuously.

NYCU × Semiconductor Research Center

Reshaping Data Transmission Architecture for the AI Era

As demand for AI computing power grows at an almost exponential rate, the semiconductor industry faces a formidable physical bottleneck. While computing engines (GPUs/XPUs) accelerate like high-performance supercars, the data transmission pathways connecting them remain as congested as narrow one-way streets.

To break through this bandwidth wall, tech giants globally have placed high expectations on silicon photonics and co-packaged optics (CPO). In a landmark collaboration, Distinguished Professor Chi-Wai Chow of the Department of Photonics at National Yang Ming Chiao Tung University (NYCU) and the Semiconductor Research Center have successfully established an optoelectronic measurement platform supporting single-channel speeds up to 112 GBaud (224 Gb/s PAM4). This platform lays the groundwork for Taiwan's supply chain to evolve from a follower to a leader in heterogeneous integration and talent cultivation.

"The computing power of AI servers is growing exponentially. If transmission speeds cannot keep up, chip performance cannot be fully realized," explained Professor Chow. A speed of 112 GBaud or 224 Gb/s means a single optical channel can transmit 224 billion bits of data per second, roughly equivalent to transferring the contents of six single-layer DVDs in just one second. Current mainstream 1.6T optical modules are built from eight channels, each operating at 200 Gb/s.

However, achieving such extreme speeds presents two major challenges for Taiwan. First, the industry has long depended on overseas validation resources; only by establishing a domestic measurement platform can it independently develop and verify next-generation components. Second, at ultra-high frequencies, signal attenuation and noise interference remain the developer's worst nightmare.

Silicon photonics measurement is a high-barrier integration effort. Professor Chow acknowledged that building a 112 GBaud platform requires high-frequency probes, high-speed signal generators, oscilloscopes, and mechanically stable optical tables. In this process, the Semiconductor Research Center acted as a crucial industry-grade technology catalyst, supplying international-standard high-speed measurement equipment. By combining advanced digital signal processing (DSP) with upgraded hardware, the team successfully over-

came challenges such as uneven channel frequency response and system noise, ensuring the reliability of measurement data.

Strategic Deployment of Heterogeneous Integration and Virtual Fabs

In addition, the Center and NYCU partnered with the Eindhoven University of Technology (TU/e) in the Netherlands to develop indium phosphide (InP) chips, driving transmission speeds to 320 Gb/s. Because silicon cannot efficiently emit light, introducing InP, a direct bandgap material, through heterogeneous integration is seen as the key to meeting the demands for high-efficiency light source and low-power operation in AI data centers. NYCU has also linked with international foundries to create a virtual fab environment, positioning itself as a hub for innovation.

"We have accumulated 20 years of tape-out experience worldwide, allowing us to understand the design tolerances of different foundries," Professor Chow explained. This expertise allows NYCU to serve as a bridge, translating Hon Hai's industrial needs into specifications that a foundry can actually fabricate and refining design models through measurement feedback.

The most lasting impact of the collaboration is on talent development. Students participating in this project gained direct exposure to industry standards and cutting-edge design platforms. Professor Chow emphasizes that Taiwan's semiconductor success stems from a generation of engineers fluent in processes, devices, and systems. Through the collaboration with the Hon Hai Research Institute, the students developed the ability to solve real-world engineering problems even before entering the workforce. This enduring advantage secures Taiwan's long-term global leadership.

We are witnessing the third wave of the AI revolution. The role of AI and optical communication is shifting from long-haul cables that connect cities to short-reach interconnects within data centers and ultimately to direct chip-to-chip links. The academic-industrial collaboration between NYCU and the Semiconductor Research Center represents Taiwan's decisive advantage at this pivotal turning point.

National Central University × Next-generation Communications Research Center

Building Taiwan's Space Autonomy

This is a story where academic vision meets industrial ambition. Inside the laboratories of the Department of Space Science and Engineering at National Central University (NCU), the flickering red and green lights of precision instruments illuminate the faces of young students. Dozens of kilometers away, researchers at the Hon Hai Research Institute are strategizing ways to advance Taiwan's communications capability from terrestrial networks to the Low Earth Orbit (LEO) environment, 500 kilometers above the Earth's surface.

Professor Chi-Kuang Chao of NCU's Department of Space Science and Engineering recalls that the Hon Hai Research Institute reached out to the university early on. The two parties engaged in two years of extensive communication and meshing. This process not only enabled the Next-generation Communications Research Center to grasp the complexities of space missions but also confirmed the private sector's commitment to the space industry.

With government support for Low Earth Orbit (LEO) communications satellites, Professor Chao proposed the PEARL project. In the past, satellites often weighed hundreds of kilograms, cost billions of NT dollars, and required five to ten years of development. Instead of taking on the immense financial risks and spending years to develop large satellites, Professor Chao advised the Communications Research Center to adopt the CubeSat model. While this low-cost and short-cycle approach carries its own risks, it enables experimentation and rapid iteration essential for modern innovation.

Project PEARL utilizes a modular CubeSat architecture. During the first-generation mission, although the communication payload test results were not flawless, the team collected a vast amount of on-orbit data. They precisely recorded the temperature differentials between the satellite's sun-facing and shaded sides, as well as the thermodynamic changes before and after solar panel deployment. This invaluable real-world experience was subsequently transformed into technical assets that fed directly into the second-generation mission.

Compared with the exploratory nature of the first generation, the primary objective of the second-generation PEARL is to enhance inter-satellite links. Professor Chao explained that at high orbital speeds,

aligning two CubeSats for broadband data exchange is far more complex than traditional satellite-to-ground communication.

To manage increasingly complex payload integration and orbital environments, the NCU team implemented more rigorous testing standards and simulation processes for the second generation. Whereas subsystems had previously been assigned to individual students, they are now handled by teams of two or more, ensuring cross-verification and redundancy. By enhancing the stability of the Flight Software (FSW) to minimize human error and adjusting the orbital altitude, the team expects to significantly extend the mission's operational lifespan.

Building a LEO Satellite Supply Chain

"The Next-generation Communications Research Center's commitment to space communications is rare and exceptional," Professor Chao observed. "Through deep integration with academia, they have established a complete pipeline spanning design, manufacturing, and verification."

He further pointed out that starting with CubeSats provides companies a pragmatic path to assess domestic and international supply chain requirements at realistic costs. Only when firms understand how even a single screw or a communication payload performs in the harsh environment of space can Taiwan truly establish an independent LEO satellite supply chain.

Perhaps the most enduring impact lies in talent development. Within Project PEARL, NCU students work under strict launch deadlines, communicating and troubleshooting alongside Hon Hai engineers and researchers. The convergence of industry pace and academic rigor teaches students how to meet mission demands in real time.

Ultimately, the collaboration between NCU and Hon Hai serves as a definitive blueprint for building Taiwan's autonomous LEO satellite supply chain. This vision, fueled by the synergy of academic depth and industrial strength, is being realized through these small CubeSats, lighting a path forward in the darkness of space.

NYCU × Trapped-Ion Quantum Computing Laboratory

Overcoming Optical Limits with Non-local Metasurfaces

Step into a traditional quantum computing laboratory, and you are immediately met with a dense jungle of steel and glass, an optical table crowded with precision lenses, mirrors, and laser arrays.

However, this familiar scene is undergoing transformation. Through the collaboration between Associate Professor You-Chia Chang's team at the National Yang Ming Chiao Tung University (NYCU) Department of Photonics and the Trapped-Ion Quantum Computing Laboratory, massive optical systems are being miniaturized onto a silicon nitride (Si₃N₄) chip with a centimeter-scale footprint. This represents not merely the downsizing of optical components but also a strategic milestone in Taiwan's transition from electronic integrated circuits (EICs) to quantum photonic integrated circuits (QPICs).


From Theoretical Exploration to Killer Applications

This cross-disciplinary collaboration stems from the precise alignment of academic foresight with industrial application needs. NYCU has long maintained a leading position in metasurfaces and nanophotonics, while the Hon Hai Research Institute focuses on building scalable ion-trap quantum computing architectures.

Ion trap quantum computers require highly collimated and precise laser beams to manipulate ion qubits, often necessitating the simultaneous control of multiple qubits. While traditional optical tables offer flexibility, they lack the spatial efficiency required for future demands involving thousands of qubits. The Laboratory provides a system-level integration perspective, enabling NYCU's research to move from discrete components to the system level, thereby helping overcome the challenges of hardware implementation.

Technological Breakthrough: The Non-local Advantage

The centerpiece of this collaboration is the development of a non-local visible-light metasurface integrated onto a silicon nitride waveguide. Professor Chang explains that maintaining high-fidelity beam profiles at the 532 nm green light wavelength presents a formidable engineering challenge. Traditional



Condensing massive optical systems onto a fingernail-sized chip marks Taiwan's pivotal leap from electronic integrated circuits to quantum photonic chips.

metasurfaces rely on ultra-fine, high-aspect-ratio sub-wavelength structures to modulate optical phase; however, as wavelengths shorten, the fabrication complexity of these microstructures grows exponentially.

To overcome this challenge, Professor Chang's team introduced a non-local design approach. Instead of confining light strictly within a single nanoscale structure, the design disperses optical fields across multiple resonators. This approach is significantly more compatible with semiconductor fabrication processes, reducing manufacturing complexity while still enabling highly precise local phase control to generate high-quality beams.

In addition, the collaboration moved away from silicon, commonly used in photonics, opting instead for silicon nitride (Si₃N₄). This is because silicon absorbs visible light, while silicon nitride offers excellent transparency and durability across the visible to near-infrared spectrum. Moreover, its high compatibility with existing semiconductor processes makes it the preferred material internationally for advancing ion-trap quantum computing.

Defining the Quantum Chip Ecosystem


Within this collaboration, the Hon Hai Research Institute plays an important role in defining specifications and acting as a bridge between academia and industry.

In the early stages, the teams had to mesh, having come from different fields. NYCU specialized in optical benchmarks such as waveguide loss and phase control, while the Trapped-Ion Laboratory focused on quantum metrics like qubit fidelity. By mapping these complex quantum requirements onto nanophotonic specifications, the researchers gained a clear understanding of how even minor optical deviations, such as beam misalignment or mode mismatch, directly degrade the performance of the quantum system.

The current research results, successfully demonstrated at 532 nm green light, demonstrate exceptional versa-

tility. Because the design principles of the metasurface inherently support cross-wavelength scalability, the team plans to extend integration to red and blue light and ultimately tackle the more lossy ultraviolet regime. Their roadmap envisions consolidating multiple laser wavelengths onto a single chip platform, paving the way for comprehensive on-chip spectroscopic functionality.

The collaboration sends a clear signal to the global technology sector. Leveraging its robust semiconductor supply chain, Taiwan is now equipped to enter the high-value domain of quantum photonic chips. With a fully integrated ecosystem spanning chip design, fabrication, and advanced packaging, Taiwan is positioned not only to provide foundry services but also to command the core delivery technologies of quantum systems in the coming global race.



Leveraging metasurface cross-wavelength scalability, NYCU and the Trapped-Ion Quantum Computing Laboratory plan to monolithically integrate red, blue, and eventually high-loss UV lasers onto a single chip.

University of Tokyo × Quantum Computing Research Center

Defying the Resource Curse: Architectural Innovations in Quantum Computing

In the race toward quantum computing, scientists face formidable challenges. Although quantum computers theoretically surpass classical systems in computational power, the hardware resources and computational overhead required for fault-tolerant quantum computing often grow exponentially with increasing precision requirements.


However, a recent collaboration between the University of Tokyo (UTokyo) and the Quantum Computing Research Center has successfully overcome this seemingly insurmountable physical barrier. Their findings were published in the leading journal *Nature Physics* and received the Best Student Paper Award at QIP 2025, the annual flagship conference for quantum information theory.

To appreciate the significance of this research, one must understand magic state distillation in quantum error correction. Within fault-tolerant architectures, Clifford gates are relatively straightforward to implement, but universal quantum computation requires the addition of non-Clifford gates. These, in turn, rely on high-fidelity magic states as essential resources.

The historical bottleneck lay in the iterative nature of purification. To achieve higher fidelity, scientists previously had to perform multiple successive rounds of distillation. Associate Professor Hayata Yamasaki of the Graduate School of Information Science and Technology at UTokyo explained that traditional quantum error-correcting codes cannot simultaneously optimize code rate, code distance, and symmetry. As precision requirements increase, the number of physical qubits and resource demands expand rapidly, becoming the primary obstacle to realizing a universal quantum computer.

The Power of Algebraic Geometric Codes: Achieving Constant Overhead

The key breakthrough of the UTokyo-Hon Hai partnership is the proposal of a constant-overhead protocol, where the ratio of input magic states remains fixed regardless of the target precision. By harnessing algebraic geometric codes, its research developed a scheme that achieves arbitrary precision in a single round of distillation, fundamentally eliminating the



Eliminating the prohibitive resource costs of the iterative distillation of traditional quantum codes may be the key to unlocking universal quantum computing.


massive resource overhead associated with repeated iterations.

As a world-leading university, UTokyo has no shortage of international partners. Yet Professor Yamasaki chose to establish a close partnership with the Hon Hai Research Institute, citing its unique operational approach as a fertile ground for academic exploration.

Professor Yamasaki observed that the Hon Hai Research Institute grants researchers ample freedom and flexibility, complemented by strong administrative support. More importantly, the Institute demonstrates a long-term commitment to the field by, for instance, funding postdoctoral and doctoral collaborations that sustain the growth of leading top-tier global academic teams. This philosophy of scholarly community building has earned the Institute an outstanding reputation and visibility among pure theorists.

This achievement sends a clear signal to the global quantum community. In the pursuit of fault-tolerant quantum computing (FTQC), theoretical breakthroughs are just as vital as advances in hardware.

From an industrial standpoint, the technology introduces a novel pathway to scalability. Professor Yamasaki likens a magic state distillation factory to a large water filtration plant; it is indispensable to



Through the constant-overhead protocol, future quantum hardware designers can achieve high-quality computation with a far more streamlined architecture.

system operation yet highly space-intensive. However, with the constant-overhead protocol, future quantum hardware designers will be able to deliver high-fidelity computation within a more compact and streamlined architecture.

Co-Creating Quantum Commercialization

The current results remain in the proof-of-concept stage. Although it is still facing engineering hurdles such as large constant factors in hardware implementation, the research has established an entirely new direction. Professor Yamasaki emphasized that this is the true value of industry-academia collaboration. Academia defines physical limits and theoretical foundations, while industry contributes insights into scalability and practical challenges, helping translate abstract theories into concrete technical benchmarks.

The publication in *Nature Physics* and the recognition at QIP show that Taiwan and Japan have reached world-class standing in quantum theory research.

This cross-border, interdisciplinary model of industry-academia co-creation not only strengthens their voice in the international quantum community but also provides a powerful impetus for the practical advancement of quantum technology.

As classical communications such as 5G and 6G approach their theoretical limits, quantum technology now stands at the dawn of its development. Drawing on UTokyo's academic depth and Hon Hai's industrial foresight, this constant-overhead breakthrough may well serve as the key that unlocks the door to universal quantum computing.

City University of Hong Kong × Information Security Research Center

Engineering a Resilient Ecosystem for AI Security

As AI and electric vehicle (EV) technologies advance at lightning speed, cybersecurity threats have spread from the virtual into the physical world. To confront increasingly complex attack methods, the theoretical depth of academia must be tightly integrated with the real-world practice of industry.

Against this backdrop, the Information Security Research Center and the Department of Computer Science at City University of Hong Kong are engaged in close collaboration. As a world-class technology R&D institution, Hon Hai offers its internally developed large language models, such as FoxBrain, as a practical platform, enabling academic research to directly enter real-world industrial applications.

To address the pressing issue of Generative AI security, the joint research team has put forward a Red and Blue Teaming framework. The Red Team uses the intrinsically motivated red teaming methods developed by CityU to identify a target model's vulnerabilities. This approach employs deep reinforcement learning to give the system self-evaluation capabilities, allowing it to explore previously unknown attack spaces and significantly expanding attack surface coverage. On the defensive side, the joint research team developed a suffix-based defense module for multimodal (cross-modal) attacks, pioneering the use of reinforcement learning to train defensive models. The result is a plug-and-play effect without altering the original model. In addition, the team has further optimized detection latency to extremely low levels, ensuring that FoxBrain can effectively resist jailbreak attacks while achieving a 28% improvement in security against specific attack methods.

Deepening Security Resilience

The research team also identified new types of threats in the physical world. For instance, in online HD map construction models, researchers demonstrated that attackers need only use a flashlight or small physical patches to create disturbances. The disturbances trigger amplification effects in deep learning networks, causing the model to misinterpret traffic lights or stop signs. This study indicates that future sensor-fusion designs for intelligent electric vehicles (EVs) cannot rely solely on visual models; they must proactively address physical-world attack vectors.

In testing standards, the joint team identified significant flaws in collision detection within existing autonomous driving simulators. By developing debugging techniques targeted at the simulators themselves, the team exposed corner cases where collisions went unreported. These findings were shared with suppliers, enabling virtual testing to more accurately reflect the complexities of real-world physical environments.

Beyond technical achievements, the collaboration has also yielded vital digital assets, such as the high topic-diversity red team dataset. For academia, the dataset provides a framework for fair competition; for industry, it can be used to align internal models (such as FoxBrain) for security, making them more resilient against adversarial or malicious inputs.

Looking ahead, cybersecurity will be defined by AI versus AI engagements. The joint team has already begun deploying post-quantum cryptography to develop long-term defense mechanisms for multi-robot collaboration systems, laying a secure foundation for future smart factories.

As we transition to the next phase of agent security, Dr. Xiang Zheng, postdoctoral fellow at CityU, notes that AI tools will significantly lower the threshold for attackers. Consequently, the collaboration has expanded into agent privacy protection, aiming to prevent the unlawful extraction of proprietary prompts or corporate memory from core systems.

Professor Cong Wang of the CityU Department of Computer Science emphasized that academic research risks remaining purely theoretical if it does not engage with industry practice. The Information Security Research Center provides the real-world application scenarios, performance benchmarks, and rigorous testing requirements necessary for research teams to innovate while managing cost and risk. This collaboration has not only yielded security breakthroughs but has also fostered a responsible, evolutionary AI security ecosystem.

Moments of Synergy

Every milestone marks the convergence of collective wisdom and professional strength.


Breakthroughs are never sudden flashes of inspiration; they are the cumulative result of unwavering, relentless focus.



2025 Chronology

January 7

The Chairman's Office delegation visited the Trapped-Ion Quantum Computing Laboratory.



February 12

The Girls in CyberSecurity (GiCs) delegation visited HHRI.



March 7


HHRI-NYCU co-hosted an industrial-academic innovation presentation.

March 10

HHRI announced the launch of Fox Brain, the first traditional Chinese large language model with advanced reasoning capabilities.

March 11

A delegation from the University of Illinois Urbana-Champaign (UIUC) visited HHRI.



The Semiconductor Research Center, in collaboration with NYCU, NTU, and Japan's NICT, achieved a breakthrough in asynchronous DPS-QKD technology, enhancing the stability and security of quantum optical communication systems. The results were published in *APL Photonics*.

March 22

HHRI showcased its latest research at NVIDIA GTC 2025.



March 27

HHRI held its Q1 General Assembly.

April 7

The Semiconductor Research Center, in collaboration with NTU and Japan's NICT, developed a Photonic Crystal Surface-Emitting Laser (PhCSEL) featuring an embedded 2D hollow-pyramid structure that advances inter-satellite optical communication technology. The results were published in *Photonics Research* (March 2025).

April 10

The CEA-Leti (France) delegation visited HHRI.



April 17


HHRI, the Taiwan XR Smart Industry Association, and the SID Taipei Chapter co-hosted the Beyond Reality AR/AI Exploration Forum.



The Semiconductor Research Center, NYCU, and the University of Texas at Austin achieved breakthrough in 4th-gen semiconductor applications for satellite communications and high-power devices. The results were recognized by both academia and industry, and published in *Applied Surface Science Advances* and *ACS Applied Electronic Materials*.

April 22

IEK of ITRI visited the Trapped-Ion Quantum Computing Laboratory.



April 29

The Quantum Computing Research Center demonstrated that the capabilities of shallow quantum circuits far exceed existing understanding. The results were accepted for publication in *Nature Communications*, establishing Taiwan's influence in quantum computing.

May 28

The City University of Hong Kong (CityUHK) delegation visited HHRI.



June 26

HHRI held its Q2 General Assembly.

July 10

The AI Research Center, in collaboration with CityUHK, released the multimodal trajectory prediction model ModeSeq, which was accepted at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). The improved Parallel ModeSeq model won first place in the Waymo Open Dataset (WOD) Challenge - Interaction Prediction at CVPR WAD Workshop.

July 22

The Institute for National Defense and Security Research (INDSR) delegation visited HHRI.



July 28

Malaysia's PTPKM, ITRI, III, and PQC-CIA industry alliance delegation visited HHRI.



July 30

HHRI held a team-building event.




July 31

HHRI conducted its Mid-year Research Performance Review.

August 5

Danish physicist Klaus Mølmer visited HHRI.



August 21

Sharp delegation visited HHRI.



August 26

HHRI held its International Internship Research Presentation.

August 27

FIT and FII visited the Ion Trap Lab.



September 9

HHRI and SEMI co-hosted the *NEt Forum* (Power & Optronics).




September 25

- HHRI held its Q3 General Assembly.
- The National Ilan University delegation visited HHRI.

October 13

Taiwan-Netherlands Quantum Information Processing Workshop held at HHRI.



October 16

Regional government IT departments (Taoyuan, Hsinchu, and Miaoli) visited HHRI.

November 21

Hon Hai Tech Day (HHTD).



December 4

HHRI held its Q4 General Assembly.

December 9

HHRI exhibited at IEEE GLOBECOM.



December 18

HHRI announced its support for the Australian Quantum Software Network (AQSN) Micro & MesoGrant program, committing over NT\$3 million to fund early- and mid-career researchers across Australia. The initiative aims to drive collaborative advancements in quantum software and algorithms with HHRI.