FUJITSU

2026 Predictions

Towards Thinking Machines

Dr. Manohar Kaul, Head of Next Generation AI Research Center, Fujitsu Limited

The progress of artificial intelligence (AI) is entering its most exciting stage. Large Language Models (LLMs) have shown impressive pattern recognition abilities, making AI a transformative force across industries. However, this success also exposes fundamental limits.

Current models, trained increasingly on their own outputs, face what researchers call "model collapse": A gradual decline in quality, like a photocopy of a photocopy losing sharpness with each new copy.

The scaling laws mapped by Kaplan and others in 2020 showed that making models larger improves their performance, but there is a catch: the cost increases much faster than the benefit. This is the classic law of diminishing returns. DeepMind's Chinchilla study reinforced this reality: The gains are real but get smaller each time.

No amount of computing power or data will turn a text predictor into a reasoning mind. As researchers like Yann LeCun have noted, simply adapting LLMs will not achieve even the intelligence of a cat. This is not dismissive but an acknowledgment of what true intelligence requires. A cat understands the physical world through mental models of reality, grounding its behavior in memory, prediction, and cause-and-effect. Until AI can do the same, scaling alone will not take us further.

The real question is whether AI can go beyond prediction and start learning by interacting with reality itself.

1. From Text to Physical Intelligence

The next decade will be defined not by chatbots but by robots that can reason and remember. Companies like Figure AI and Agility Robotics already have humanoids in warehouses, but the main challenge is in the brain, not the body.

In 2026, we will begin to see robots equipped with memory, reasoning and basic world models. They will start to recall past interactions, adapt in real-time and exhibit the first signs of planning. These early steps will be fragile, but they will mark the shift from locomotion to cognition: the beginning of Physical AI.

Vision-language-action (VLA) models show progress by linking instructions to sensor data. However, they often remain fragile. Move an object behind another, and the system fails.

Rearrange the room, and it becomes confused. True intelligence requires more than just following instructions; it requires building on experience.

Imagine a maintenance robot that not only repairs a leaking joint but also remembers that the same pipe has failed twice before and suggests replacing the entire section. This shift from simply reacting to anticipating signals marks a step toward true intelligence.

2. Memory: The Foundation of Intelligence

Today's robots resemble athletes who are strong and agile, yet they often forget their past performances. They can move but cannot build on experience. The real breakthrough will occur when machines begin to carry forward what they have learned, rather than starting over each time.

Memory is the foundation of intelligence. Without it, there is no continuity and no understanding of the world. A cat navigating a cluttered space does not calculate every path. It remembers obstacles it previously encountered and knows which can be safely ignored.

In 2026, we anticipate the introduction of the first robots that operate in this manner. They will not just sense and react; they will begin to recall, form simple hypotheses, and plan. These systems will process sensor data, visual information, and temporal sequences as naturally as current models handle text—building toward multimodal intelligence that seamlessly combines vision, sensing, language, and environmental understanding.

3. World Models: Understanding Physics and Causality

The most significant breakthrough, starting in 2026, will be AI systems that build world models: Digital representations of physical reality that enable rapid adaptation to new environments. These systems will develop an intuitive understanding of physics, akin to biological intelligence, grasping concepts such as weight, balance, structural integrity, and coping with spatial relationships without explicit mathematical programming.

Unlike current systems that require extensive task-specific training, these AI systems will be able to reason about how objects can be manipulated and how spaces can be navigated. Industrial applications will drive this change. Manufacturing systems will adapt to new production demands in hours, rather than weeks. Warehouse robots will optimize logistics in real time by considering physical constraints, efficiency trade-offs, and safety issues.

The breakthrough goes beyond robotics to any AI system that operates in physical environments. Smart building systems will analyze occupancy patterns, energy flows and structural dynamics. Agricultural systems will combine soil conditions, weather patterns and crop biology to make advanced growing decisions.

4. Safety and Ethics: Intelligence with Integrity

As these systems become more advanced, safety cannot be an afterthought. It must be integrated into their core design. Intelligence and safety are not mutually exclusive; they can and should develop in tandem. Post-quantum security measures will be seamlessly incorporated into core AI frameworks, safeguarding against both present and future threats.

Ethical reasoning will become a fundamental aspect of intelligence, with AI systems showing a sophisticated understanding that surpasses simple rule adherence and involves genuine contextual judgment. Democratic governance mechanisms will be embedded in AI architectures, enabling rapid adaptation to changing regulations while maintaining consistent ethical standards.

This forward-thinking approach shows how responsible AI development can promote, rather than hinder, innovation.

5. Democratized Intelligence: Efficient AI for Everyone

Major advancements in resource-efficient AI will make advanced intelligence accessible to organizations of all sizes. These efficient systems will learn continuously from streaming data rather than relying on large pre-training phases, enabling real-time adjustments to changing conditions.

Small manufacturers will adopt AI systems comparable to those used by large corporations. Service organizations will tailor AI behavior to suit their specific operational needs. The integration of edge computing will enable advanced AI reasoning to occur locally, thereby enhancing response times while reducing bandwidth usage and addressing privacy concerns.

This democratization will drive innovation in previously underserved sectors. Agricultural cooperatives will implement advanced crop management systems. Small clinics will gain access to modern diagnostic tools. Educational institutions will provide personalized learning experiences regardless of resource limitations.

Conclusion: The Next Chapter of Intelligence

The last decade was focused on text and LLMs. The upcoming decade will be about Physical AI: not just bodies that move but minds that think. A cat doesn't predict the next word. Instead, it navigates the world by remembering where it hid yesterday, recognizing which patterns matter and which can be ignored, and planning how to move through space without disruption. The question is whether we can build machines capable of doing the same. Can AI form mental maps, hold them in memory and use them to plan its next move?

At Fujitsu's Next Generation AI Research Center, we focus on developing AI that safely manages the complexity of real-world scenarios. The transition from pattern recognition to genuine understanding represents both a technological breakthrough and a fundamental broadening of what machine intelligence can achieve.

By 2026, we will have made decisive progress beyond the current chapter toward systems that demonstrate genuine intelligence, ethical behavior, and practical capabilities. The future belongs not to larger models but to smarter ones: AI that truly understands the world and acts within it with wisdom and integrity.

The true milestone may not be when AI passes a benchmark test, but when it finally exhibits the adaptive intelligence that we observe in the natural world around us.

Dr. Manohar Kaul, Head of Next Generation AI Research Center, Fujitsu Limited

Dr. Manohar Kaul leads Fujitsu's Next Generation AI Research

Center in Tokyo, where he directs advanced research in machine learning, robotics, and secure AI systems. His work focuses on geometric and topological methods, optimal-transport modeling, and representation-efficient architectures; he aims to develop robust and globally deployable AI technologies that support Fujitsu's long-term innovation strategy.

Before joining Fujitsu Japan, he headed the AI and AI Security laboratories at Fujitsu Research India, guiding research programs that combined fundamental science with enterprise-scale applications. Earlier, he served as a tenured Associate Professor at IIT Hyderabad, where he founded the Krama Lab and established cross-industry and academic research collaborations. His earlier industry experience includes nine years as a systems and database architect at Oracle and other global technology companies.

Dr. Kaul has lived and worked across multiple countries, including India, Sweden, Denmark, Germany, Japan, and Australia. This international background shapes his multicultural research leadership and global perspective in AI development. He holds an M.Sc. in Computer Science from Uppsala University in Sweden and a Ph.D. from Aarhus University in Denmark, followed by a postdoctoral appointment at TU Berlin's DIMA group on scalable machine learning and high-performance data systems.