ARTIFICIAL INTELLIGENCE

IBM has been a leader in artificial intelligence (AI) since the field's earliest days, when Arthur Samuel developed a checkers-playing program that learned from experience in the early 1950s. Later, Gerald Tesauro developed TD-Gammon, a self-teaching program that learned to play backgammon at the level of the best human players in the world, which textbooks define as a classic success story in machine learning. More recently, IBM made history when Deep Blue, IBM Research's chess-playing program, defeated world chess champion Garry Kasparov in 1997. Today, IBM's AI research is focused on developing new approaches to deal with the increasing complexity of large-scale computer and software systems. To this end, IBM Research is pursuing a wide variety of methodologies, including reinforcement learning, Bayesian reasoning, preference elicitation, knowledge representation, and planning. In addressing the key challenges in complex systems, research activities center on solving real-world problems faced by IBM and its clients, as well as on advancing the state-of-the-art of AI on many different fronts, including autonomic computing, intelligent interfaces, and modeling the brain.

IBM studies the brain at many levels of abstraction , from individual neurons *(left)* and self-organized cortical maps, to human brain structure and processes *(below)*.

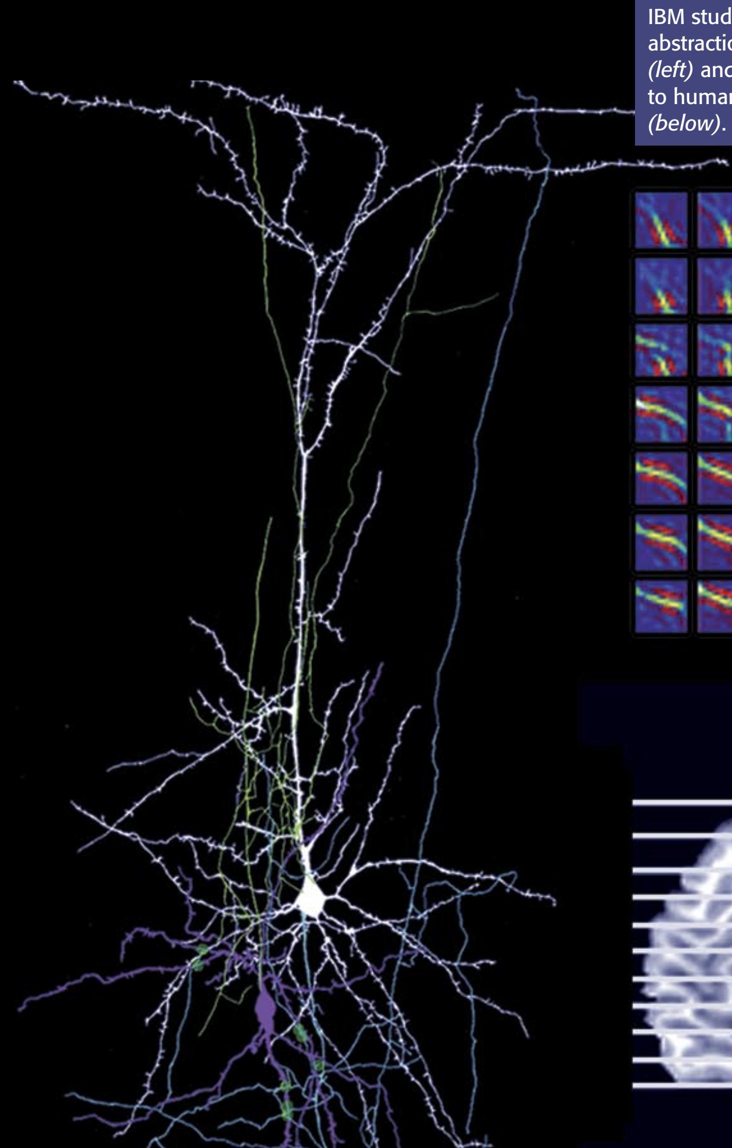

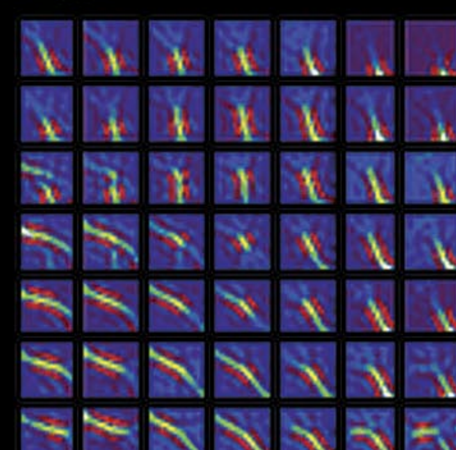

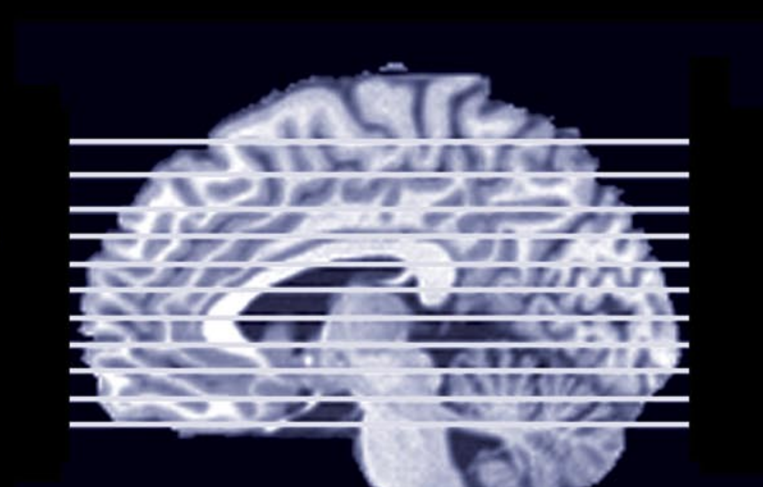

AUTONOMIC COMPUTING

In autonomic computing, a fertile area for AI, researchers are developing novel machine-learning approaches to automatically determine self-management policies in large, complex systems, when given only high-level business objectives. Using a model-free reinforcement learning approach, the system tries various management decisions in many different system states and automatically learns a policy for maximizing cumulative reward. Efficient techniques for resolving resource conflicts via cooperative negotiation are being explored by employing intelligent incremental utility elicitation strategies that quickly find near-optimal allocations with respect to a minimax regret criterion.

Automated, real-time diagnostic techniques for large, distributed computer systems are also being developed that require fast and accurate inferences from huge volumes of data. By combining Bayesian inference and learning with information-theoretic approaches to optimal test selection, researchers have designed novel cost-efficient, adaptive diagnostic algorithms. The key efficiency issues include both the cost of obtaining the information and the computational complexity of inference and learning in approaches such as dynamic Bayesian networks.

INTELLIGENT INTERFACES

Programming-by-demonstration and machine learning methodology is being applied to automatically capture and generalize procedural knowledge (know-how). IBM's system observes users performing complicated procedures (such as configuring servers), combines multiple observations, and produces executable procedures. This human-readable and editable procedure serves as "follow-me" documentation, meaning that it can guide a new user through a similar task by providing in-context help as the procedure is performed. Because the executable procedure can be dynamically updated and disseminated, the need for actual programming or writing documentation is minimized, which can reduce the cost of acquiring and maintaining procedural knowledge.

MODELING THE BRAIN

Researchers are studying the vertebrate brain at multiple levels of abstraction to develop computing algorithms that assimilate knowledge and solve problems as animals do. This work is grounded in computational models of neocortical minicolumns, which realize key emergent phenomena of cortical maps. The goal is to understand and model the structures of vertebrate nervous systems and the integrated processes operating over them that give rise to the capabilities of perception, behavior generation, and learning.