The Growing Reality of AI:

Toward a Natural Intelligence Machine

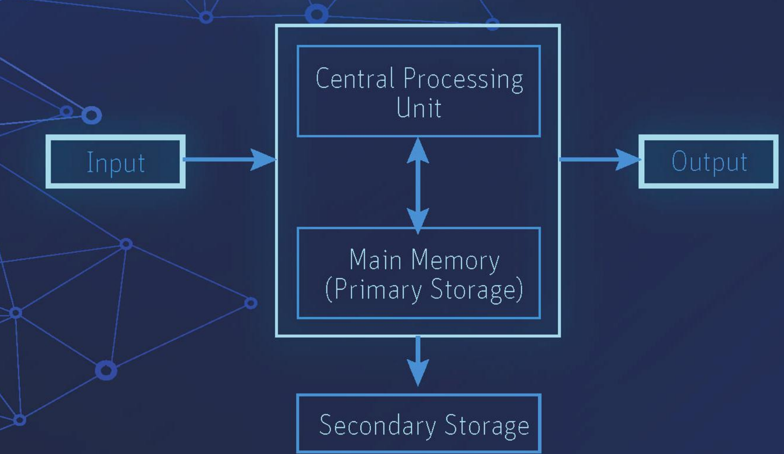

Computation machine
is a Von Neumann model, and it processes function.

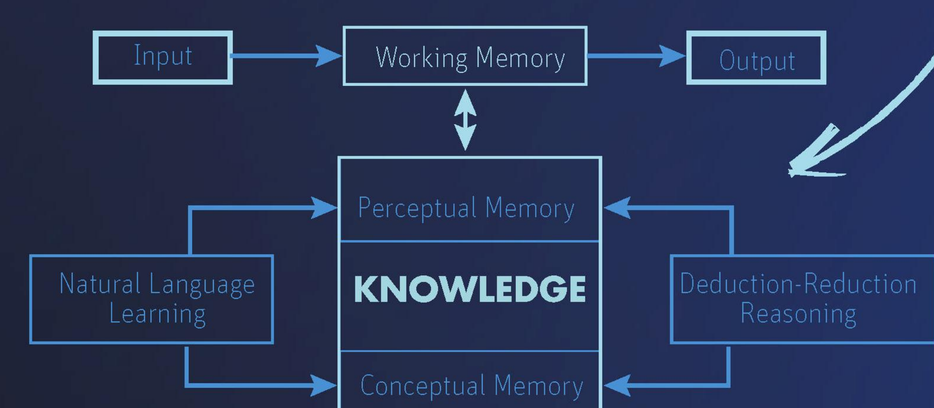

Knowledge Machine
is a non-Von Neumann model, and it maps knowledge.

Introduction to Artificial Intelligence

For all its remarkable power, most of the software people use today treats computers like very capable calculators. These programs do what they were programmed to do–solve a math problem, draw a picture, communicate with another computer–no more, no less. This strictly logic-based model, in which programmers essentially detail the precise sequence of manipulations required to complete some task, has proven remarkably enabling, influencing most aspects of modern life (top left diagram). Yet, there is much such software cannot do. It cannot, for instance, learn. Nor can it adapt to and act on new information.

Learning and adaptation are traditionally the purview of the human brain, and for nearly 60 years computer scientists have been building and enhancing computational systems that can do *in silico* what the brain does *in vivo*. Today, they have, at least in part, succeeded.

Using "artificial intelligence" (AI), sometimes called "cognitive computing," researchers are at last able to build true–albeit still somewhat crude–computational assistants, tools that can learn, adapt, and even "think," sifting through reams of data to identify hidden relationships and unexpected connections. This is not the AI of movies. Yet in fields from medicine to the mundane, the results promise to be transformative.

How AI Works

The questions, "Who won tonight's football match between Spain and Italy?" and "Did Spain win their game tonight?" are similar, but different; the former asks for the name of the winning team, while the latter is a yes/no question. Answering them requires understanding their grammar, the intent of the question, and finally, the requested information.

Modern speech recognition and natural language processing software can easily tackle such questions. But more sophisticated queries–which medical treatment will be most effective for a given patient, for instance–require more complex systems, the realm of AI.

AI combines natural language processing–the magic behind Apple's Siri, among other applications–and machine learning to make computers better able to interact with their human operators. Rather than being hard-wired to do *X* in response to *Y*, these systems absorb knowledge, from spoken language, text, or image data, and make connections based on what they "read." Using the text "John is Robert's brother," for instance, such systems can extrapolate the appropriate players and relationships, and apply and integrate those with other data, such as "Sally is John's daughter." Thus, they "learn," and as new information is added, the system grows ever more adept in its work, making better and better decisions.

These features are not, per se, new: Systems have had some capacity for learning–for instance, using artificial neural networks to drive facial recognition–and recapitulating human dialog since the dawn of the computer era. Yet the confluence of software engineering and computational muscle has at last begun to yield some truly remarkable results.

IBM's Watson supercomputer represents a deep statistical approach to AI. Comprising some 2,900 processors and 16 terabytes of RAM, the system uses a training set of books, research articles, videos, and images, as well as a knowledge of grammar and vernacular, to first interpret a query and then answer it. Those skills allowed one offline Watson implementation to defeat two human *Jeopardy!* game show champions in 2011, and since then, the system has only gotten more sophisticated.

An example of a different approach is Weiming's Knowledge Machine (see bottom diagram, above). In contrast to the AI approach used by IBM, it is based on "natural intelligence" (NI), which attempts to capture how knowledge is learned, organized, and logically retrieved in a biological system. The Knowledge Machine implements a "dual-memory" system that mimics the innate logic structure of the human brain, which allows for both the acquisition of information (perception) and its integration with existing knowledge (conception). As knowledge is added to and retrieved from memory, it is processed both deductively and reductively, just as a human would, the "natural" part of NI.

Applications

AI and NI promise to redefine what computers can do, and indeed, they already have done so. Google, for instance, has developed systems that can master a video game or drive a real-world car. Others have developed systems that can automatically and in an unsupervised manner analyze images to see whether they contain, for instance, adult content.

More importantly, AI is expanding human-machine interaction. Simple Boolean keyword searches can be replaced by sophisticated queries that allow users to ask more nuanced questions and better leverage the systems' stored knowledge. The results, in turn, are graded not just on how often, say, a keyword appears, but on how well the information reflects the question asked, and they often include supporting evidence.

One application involves sifting through vast knowledge bases to derive novel conclusions. In the life sciences, for instance, where there are tens of thousands of research journals and millions upon millions of publications, it is almost impossible to stay abreast of the research literature, much less synthesize it effectively. But cognitive computing can, and in 2014 researchers at Baylor College of Medicine worked with IBM to develop a system that could do just that. The Baylor team trained a custom Watson implementation on some 70,000 scientific publications to identify proteins likely to be involved in the control of the cancer-related protein, p53. The system found 74, many of them representing new potential research directions.

AI likewise can be used to provide more intuitive interactions between man and machine, such as the ubiquitous computer agents that answer customer service calls, and even as a form of "virtual psychologist"–a concept that has been in development since at least the ELIZA system of 1965. Today, Watson is being used to intelligently answer questions on a financial services website, and to help an insurance company better weigh requests for medical procedure preapproval.

Unlike AI, NI is not based on a statistical (or stochastic) model–rather, it depends on a more mathematically consistent, deterministic methodology. This removes some of the randomness, with the intention of bringing a higher level of accuracy to the process.

Using NI and powered by a simple "innate processor" and "simulated 3D memory," the Knowledge Machine is being applied to formulate and process natural language database queries and to drive more accurate speech recognition algorithms. The system can even be used to drive genetics research, using its "innate logic" to infer relationships between DNA and protein sequences, as well as their three-dimensional structure.

Produced by the *Science*/AAAS Custom Publishing Office
Science AAAS
Sponsored by
未名集团 SINOBIOWAY