Intelligent machines that learn unaided

Research on artificial intelligence is moving toward the creation of intelligent systems able to learn by themselves from experience, as is the case with the human mind. This is the target that the research project 'Universal Artificial Intelligence' has set itself. Under the supervision of Marcus Hutter, this research is carried out at the Dalle Molle Institute for Artificial Intelligence (IDSIA), a joint institute of Università della Svizzera italiana (USI) and Scuola Universitaria Professionale della Svizzera Italiana (SUPSI). The funding is provided by the Swiss National Science Foundation.

May 1997: world chess champion Garry Kasparov is defeated by Deep Blue, IBM's supercomputer. At the end of the game, Kasparov admitted to feeling as if what stood there, opposite him, was a kind of intelligence he was not used to. It was an 'artificial mind' out of the ordinary, able to anticipate 200 million moves per second but incapable of coming to grips with new situations. It could not learn from its mistakes and had no way of recognising the weak points of its opponent. It could 'only' obey its algorithm, a program so beautifully conceived as to checkmate the human mind. Nonetheless, unlike Deep Blue, human beings are able to tackle unexpected situations on the strength of their experience and intuition. This is what 'Universal Artificial Intelligence' is steering towards, a project conducted at the Dalle Molle Institute for Artificial Intelligence Studies and headed by Dr Marcus Hutter. Subsidised by the Swiss National Science Foundation, it begins by appraising learning methods in man, and converts them into mathematical algorithms. Car-manufacturing plants employ robots, which perform perfectly well provided the environment is not altered. As soon as a new variable is introduced the robot has to be reprogrammed. The new intelligent systems, however, have to be able to learn from the impact of their actions on their surroundings; for example, because of some manoeuvre a robot may knock against an obstacle, or its battery may go flat if it is asked to move too far. The 'carrot and stick' approach, also known as 'reinforcement learning' is: reward and punishment alternate staking out the path to be followed. Researchers are in no position to tell in advance what sort of surroundings a robot - despatched, say, to Mars - may find itself in. While initially the intelligent system will have very little information and lots of freedom of action at its command, trial and error will teach it how to cope with the unexpected or unpredictable. "*To come back to our example of the game of chess* - says Marcus Hutter - *the algorithm of the new intelligent systems is not going to explain all the rules of the game; it will simply instruct the machine to seek a positive reward. Our task will be to issue a negative feedback if it loses and a positive feedback if it wins. The algorithm will inculcate into the machine that its life-purpose is to maximise its rewards*". This general notion is crucial to the process of endowing the robot with the necessary cognitive skills to tackle complex situations. The algorithms developed by IDSIA's researchers will be instrumental in dealing with real contexts characterised by a strong random component. The machine will be expected to learn by itself.

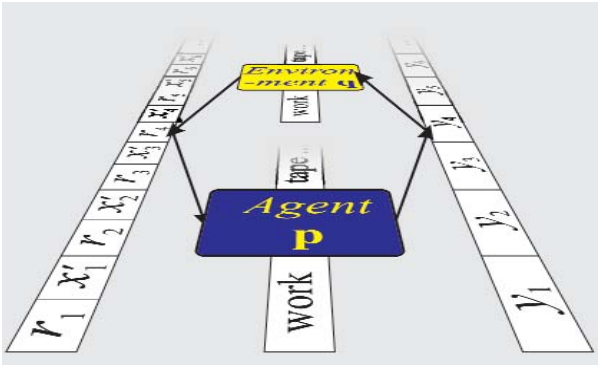

In the picture: agent interacting with its surroundings.

IDSIA: a joint USI / SUPSI institute

IDSIA was set up in Lugano in 1988 by the Dalle Molle Foundation, which promotes quality of life. Since then, IDSIA has put a good deal of resources into various strands of IT research: from computational linguistics to the teaching methods of artificial intelligence, from computers that learn (artificial learning) to optimisation techniques. When USI and SUSPI were founded, IDSIA donned a new role as a bridging institute, affiliated to both institutions of higher education. It is engaged in many activities, carried out in association with universities in Switzerland and abroad. The research work currently in progress is funded by the Swiss National Science Foundation, by the Technology and Innovation Commission, and by the European Union. In se-veral fields, IDSIA has developed successful methods and it has more than once even achieved top scoring.

Artificial learning within an algorithm

Observing human behaviour has proved useful in laying out the terms of the algorithm. *"If one decides to accept a job offer* - explains Marcus Hutter - *one does it because one believes that that job can ensure personal satisfactions, future prospects and a good salary. In other words, decisions try to maximise one's long-term advantage."* Clearly, for a robot, advantage is closely linked to the chosen goal (for example, collecting rocks on Mars). Once the target has been set, the algorithm orders the robot to learn to select the behaviour strategy that maximises the average benefit over its life-span. In other words, the point is not to teach the robot how to carry out a task (pick up rock!) but quite simply to maximise its benefit. In the real world, as we know, there is no way of predicting future events with certainty: building on past experience and on the circumstances we live in, we can nonetheless estimate whether, for example, today it is likely to rain. Similarly, the model formulated by our researchers expects to equip the intelligent machine with a probabilistic (Bayesian) sort of analysis and to combine it with reinforcement learning. Further, it is experience that teaches human beings how to come to terms with a new environment; similarly, by exploring the world around, a robot will bag up new useful knowledge. Basically, the robot gathers data, and this will help it identify the environment. The basic idea is to search for the simplest or shortest algorithm which can reproduce the data. Joining these concepts into one single algorithm (call AIXI, see box) is the conclusive step for creating versatile machines, without any particular biased know-how, but apt to learn and adjust to unforeseen changes in the outside world. At present, the project has completed the theoretical part; so researchers are working at refining the general theory by homing in on some concrete applications. One doubt remains: are we to be afraid of the intelligent machine of the future? "*No* - is Marcus Hutter's firm answer - *this is not a dispute between man and machine. The future of humans depends on their increasing interaction with computers. Even today, a PC is already seen as an extension of our skills. And the same is true of the new intelligent machines: it will be a natural evolution of mankind."*

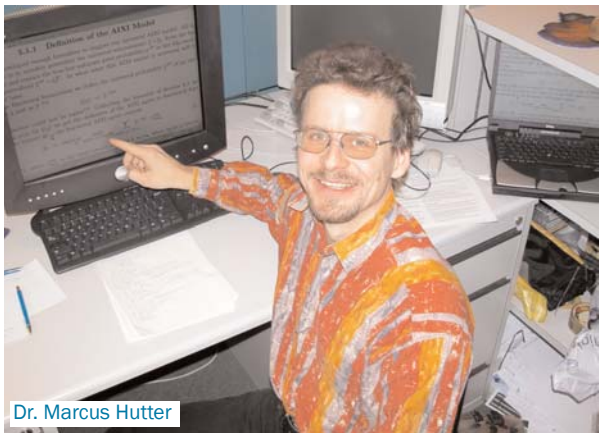

Dr. Marcus Hutter

The universal algorithmic agent AIXI

AIXI is a theory of sequential decision making akin to Solomonoff's universal theory of induction. Prof. Solomonoff, co-founder of algorithmic information theory (and visiting researcher at IDSIA in 2001), derived an optimal way of predicting future data, given previous observations, provided the data is sampled from a computable probability distribution. AIXI is a universal theory without adjustable parameters, making no assumptions about the environment except that it is sampled from an unknown computable distribution. From an algorithmic complexity perspective, the AIXI model generalizes optimal passive universal induction to the case of active agents. AIXI is a suggestion of a new "learning" algorithm, which may overcome all (except computational) problems of previous reinforcement learning algorithms.

Further information

Dr. Marcus Hutter

Dalle Molle Institute for Artificial Intelligence

Galleria 2 - 6928 Manno (Lugano) - Switzerland

Phone: +41 91 6108668 -

E-mail: marcus@idsia.ch

http://www.idsia.ch/~marcus/idsia/

Web addresses:

USI: www.unisi.ch

SUPSI: www.supsi.ch

IDSIA: www.idsia.ch

I M P R E S S U M

Project of the month

No. 9 - April 2004

on-line: http://www.ticinoricerca.ch/progettomese.htm

via Lambertenghi 10a, CH- 6904 Lugano

Tel. +41 91 912 46 17; e-mail: sr@unisi.ch