How to Learn a Program: Optimal Universal Learners & Goedel Machines

Jürgen Schmidhuber JUERGEN@IDSIA.CH
IDSIA
Manno-Lugano, Switzerland

Rational embedded agents should try to maximize future expected reward. In general, this requires learning an algorithm that uses internal states to remember relevant past sensory inputs. Most machine learning algorithms, however, just learn reactive behaviors, where output depends only on the current input. Is there an optimal way of learning non-reactive behaviors in general, unknown environments? Our new insights affirm that the answer is yes. I will discuss both theoretical results on optimal universal reinforcement learning & Goedel Machines, and mention applications of the recent Optimal Ordered Problem Solver.

Links related to this talk:

Cogbotlab:
`http://www.idsia.ch/~juergen/cogbotlab.html`

Universal learning machines / Goedel Machines:
`http://www.idsia.ch/~juergen/unilearn.html`
`http://www.idsia.ch/~juergen/goedelmachine.html`

Optimal Ordered Problem Solver:
`http://www.idsia.ch/~juergen/oops.html`