SAND2008-7134C

BECCA:

A Brain-Emulating Cognition and Control Architecture

Brandon Rohrer

Intelligent Systems, Robotics, and Cybernetics Group
Sandia National Laboratories

S-Learning

S-Learning is a solution to some Dynamic Reinforcement Learning problems:

- Finite state and action spaces
- No a priori system model.
- Reward is a known, *time-varying* function of state.
- Describes many real-world interaction problems.

S-Learning allows BECCA to drive entirely unmodeled hardware in unknown environments.

Context-Based Similarity

Context-Based Similarity (CBS) provides a solution to the problem of semantically clustering symbols:

Given an ordered set of symbols, determine the semantic distance (similarity) between any two.

CBS provides BECCA the capability for fully autonomous concept generation.

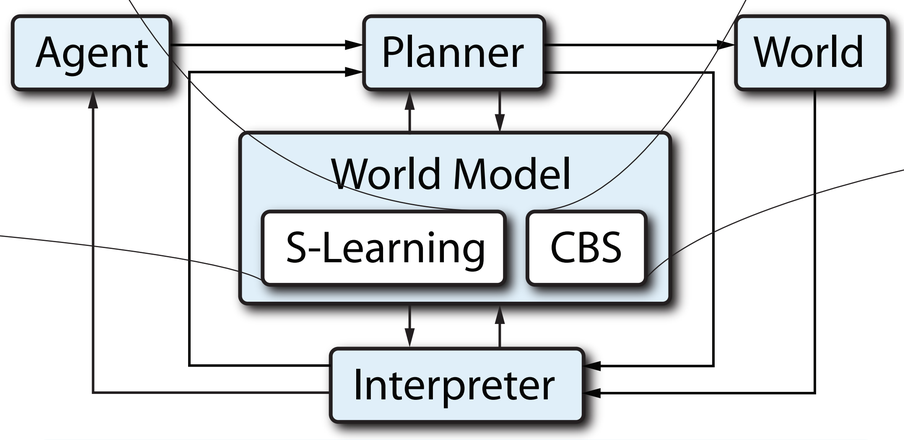

BECCA is a biomimetic approach to achieving human-like reasoning, perception, learning, and movement control in machines.

World - provides symbolized sensor state.

Interpreter - based on existing concepts and predicted state, modifies the incoming state.

World Model - stores repeated sequences from the interpreted state history.

Agent - determines the reward function.

Planner - uses the current state and reward function to select action plans from past experience.

Robot control: 1 dof

- 10° position bins.
- 10° movements.
- Degenerate case prediction task and control task.

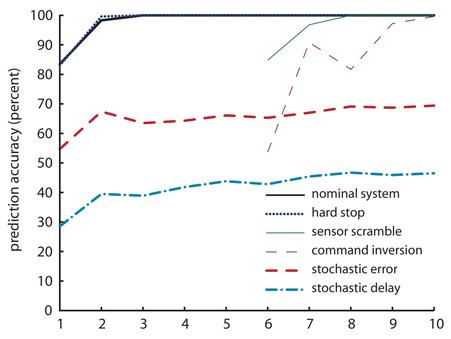

S-Learning approached optimal prediction performance in all cases.

In control task, achieved perfect performance.

Solved a Dynamic Reinforcement Learning problem.

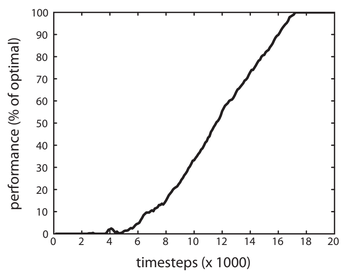

Natural Language Processing

Textual analog of semantic clustering is identifying related terms.

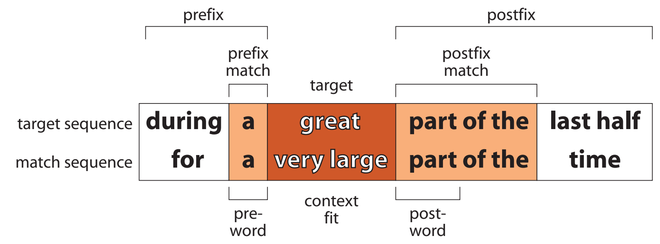

- World Model trained using 25M-word tokenized English corpus.
- No tags or annotations.
- Seed words produced related terms.

seven	large	sugar	feet	father	road
five	great	flour	face	mother	river
two	small	fruit	heart	wife	street
four	considerable	butter	head	son	table
three	certain	salt	side	head	fire
ten	very large	water	house	life	hill
twelve	good	mace	lips	voice	head
fifteen	very small	mear	work	face	house
twenty	vast	cream	hands	heart	room
fifty	larger	brandy	back	name	lake

Robot control: 7 dof

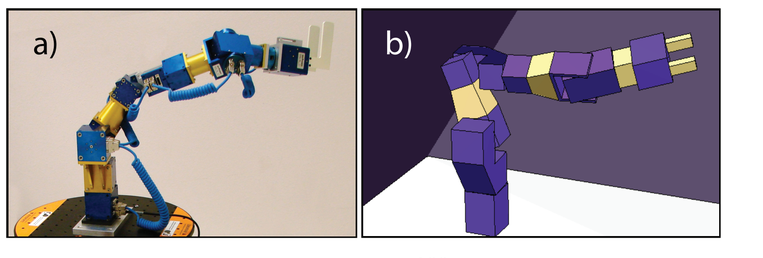

Robot and reaching task mimicked real-world. Joint positions, primitive vision, and grip sensor. No kinematic or spatial model.

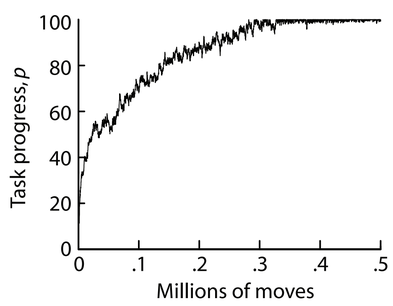

- Early attempts appear unguided, later attempts are more direct.
- Learned reaching task in 4000+ dimensional state-space.

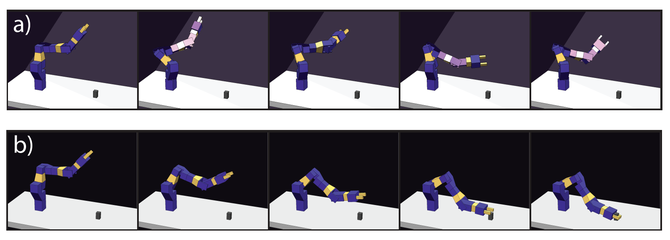

Robot control: 2 dof

- Planar reaching task.
- Coarse and fine joint position.
- Primitive vision.

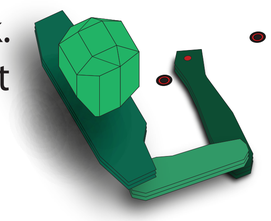

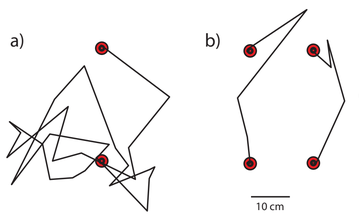

Early attempts appear unguided, later attempts more direct.

Demonstrated generalization: practice on task A improved performance on task B.

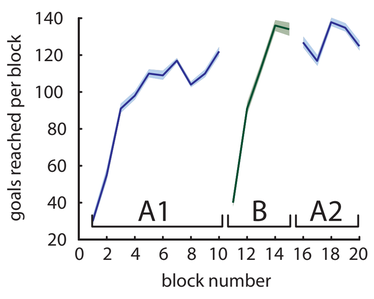

S-Learning and CBS in BECCA

- S-Learning provides CBS with a library of sequences.
- CBS clusters symbols into concepts, reducing the dimensionality of S-Learning's state space.
- S-Learning interacts with the World based on CBS' concepts, providing CBS rich data with which to refine them.
- CBS-interpreted data creates a new symbol stream which can be combined and passed to yet other instances of BECCA in a hierarchy.

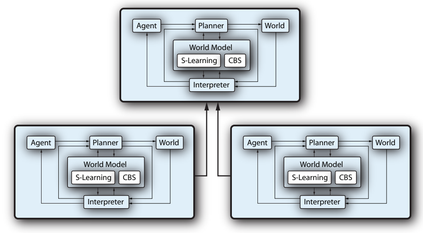

Correspondence welcome: brrohre@sandia.gov 505-844-6569 MS 1010, PO Box 5800 Albuquerque, NM 87185-1010

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.