Anytime Bounded Rationality

Eric Nivel, Kristinn R. Thórisson, Bas Steunebrink, Jürgen Schmidhuber

Icelandic Institute for Intelligent Machines

Reykjavik University

The Swiss AI Lab IDSIA, USI & SUPSI

Overview

> Motivation

> General Approach

> Deep Pervasive Control

> Simulation & Commitment

> Results

Motivation: functional requirements

> Multi-objective control

> to achieve its mission, a system may have to pursue multiple goals under multiple constraints: overlapping, antagonistic, synchronous, etc.

> some goal states can only be reached using coordinated, co-dependent synchronous courses of action.

> Anticipative control

> sub-goaling depend on predictions, spending resources to predict depends on goals.

> arbitrary time horizons.

> Anytime control

> the value of reaching a goal state is extrinsic and dynamic.

MOTIVATION: REALITY

> Fixed memory budget, CPU cycles/s, # of CPUs

> Computation times/task are variable

> depend on the task, knowledge, available resources, context.

> Knowledge (what, when, how) is limited & variable

> insufficient, incomplete, inconsistent.

> experience-based: comes and go, more or less trustworthy.

> relevancy is context-dependent.

> Deluge of inputs

> variable periodicity and value.

General approach: operational requirements

> Process a deluge of inferences

> variable periodicity, value and trustworthiness.

> Continuous time, asynchronicity

> no ”time steps”.

> Keep latencies under control

> no ”cognitive cycles”.

> Keep expenditures under control

> no ”fast enough” nonsense.

> Goals & predictions must be concurrent

> Keep the system on its toes

> simulate, maintain and revise possible courses of action.

General approach: inferencing

A(X Y) M B(Y Z)

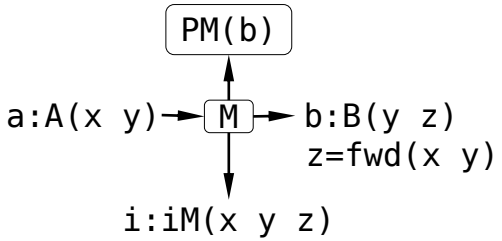

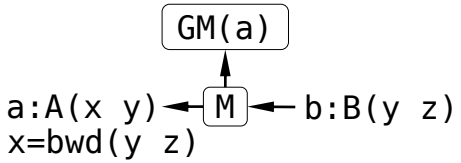

General approach: inferencing

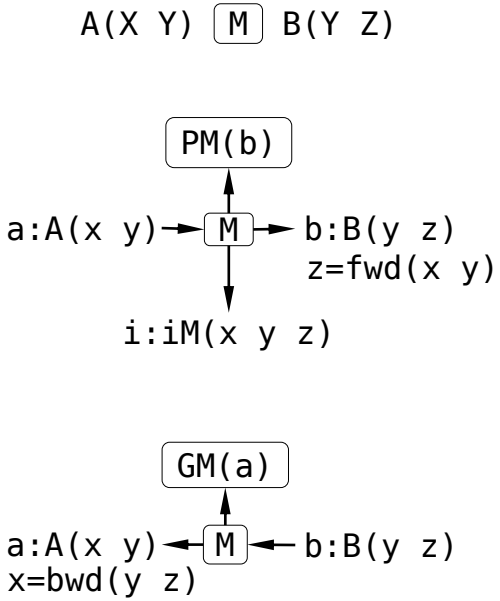

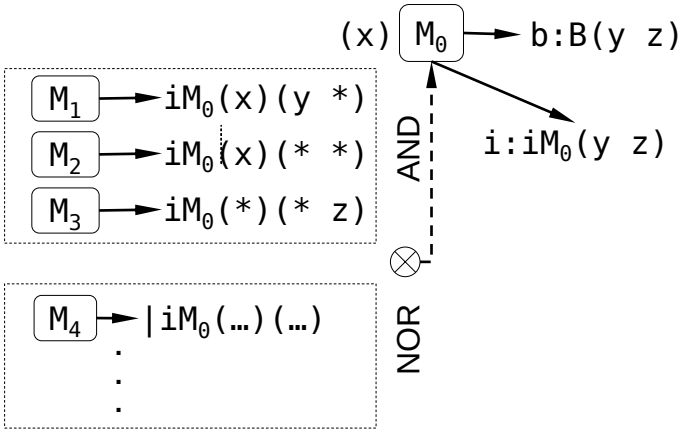

General approach: inferencing

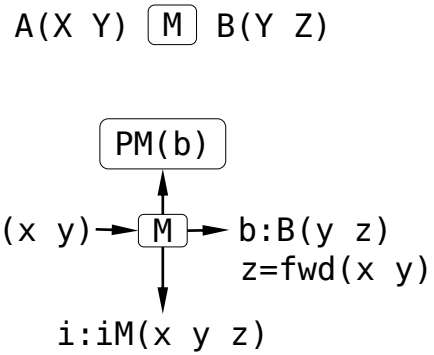

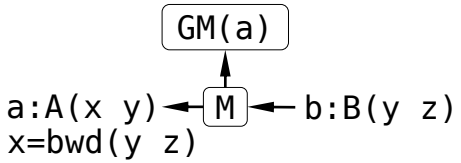

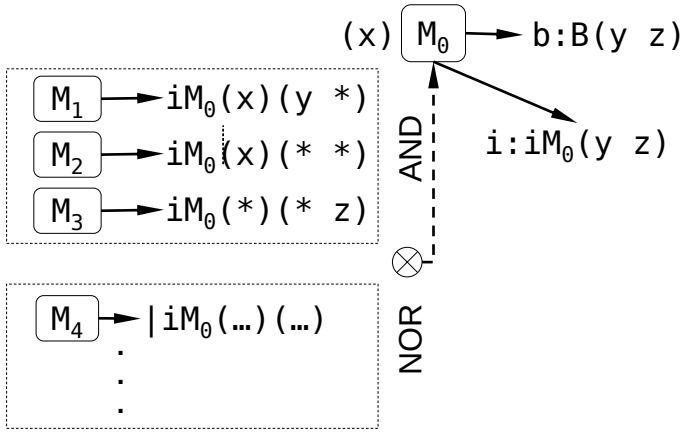

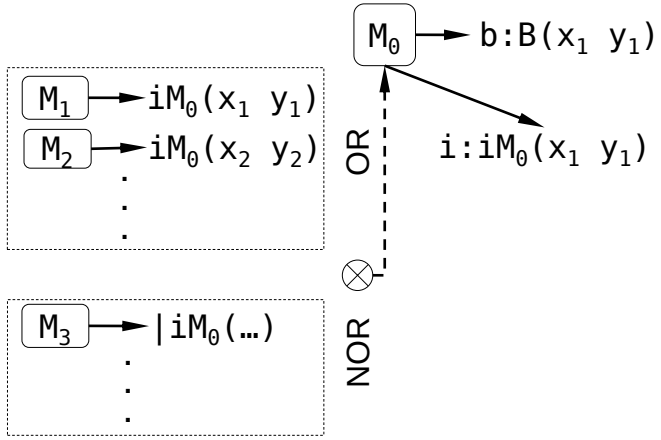

General approach: topology

> Dynamic network of learned ”circuits”

> indirect coupling (data coupling).

> double capability: predict & control, depending on the coupling & chaining direction.

> hierarchies of negative feedback anticipatory controllers (std. CT).

> Continuous

> the quantitiy & quality of evidence matters.

> Persistent

> outputs updated as new evidences are accumulated / considered.

> Pervasive

> control exerted at the lowest level (job).

> from the top (mission) to the bottom (commands).

> from the bottom (percepts) to the top (simulations).

Deep Pervasive Control: structure

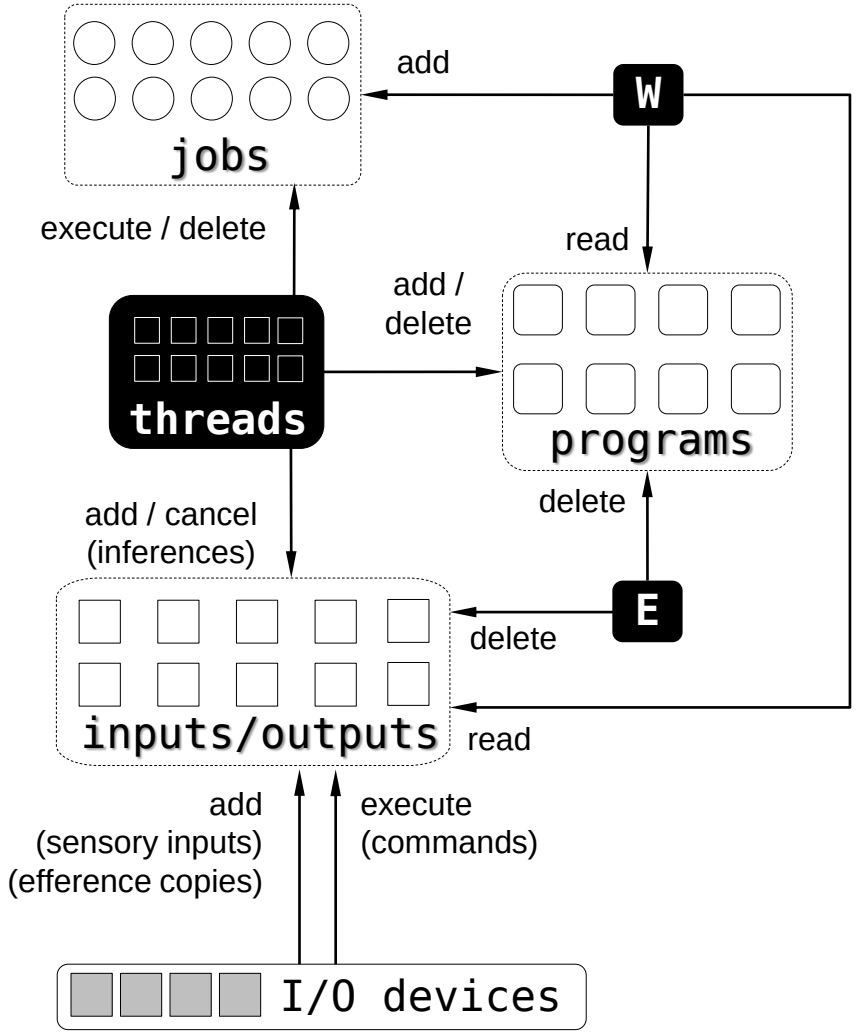

input + program → job

> All computation times bounded

Deep Pervasive Control: value-driven

> Job(i,p).priority(t) = i.tending_value(t) x p.relevance(t)

> i.tending_value

> quality and length of the inference chain.

> decreases with the time horizon (urgency).

> for goals: decreases with predicted achievements.

> p.relevance

> forward: tending values of goals produced by p.

> backward: tending values of predictions produced by p.

> Eraser: value-based

> N.B.: attention = scheduling

Simulation & Commitment: goals

> Continual simulation for readiness

> Anytime commitment: whenever it makes sense WRT planning

> Commitment is revisable

Simulation & Commitment: operation

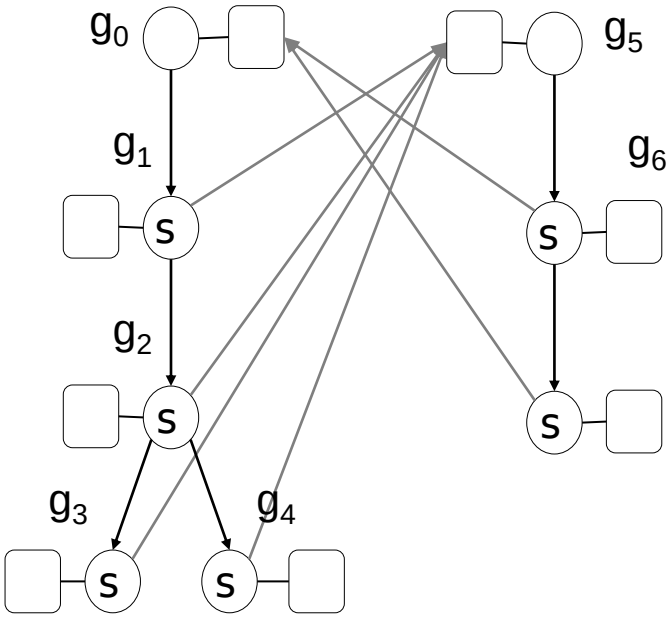

simulated predictions

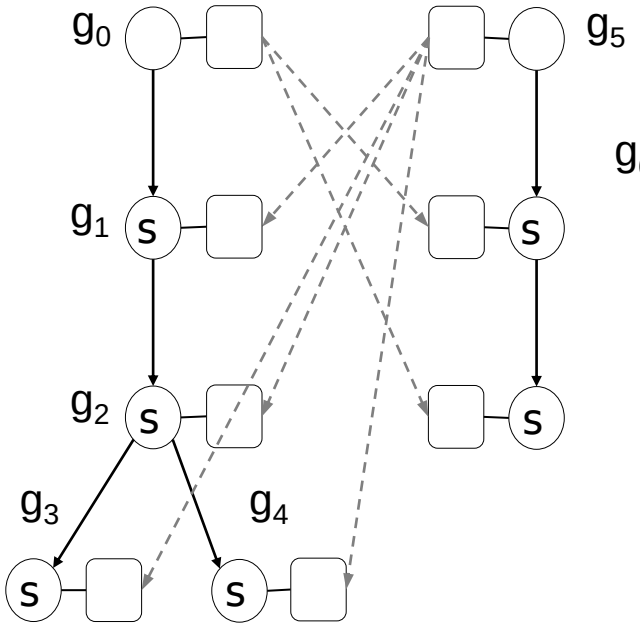

simulated outcomes

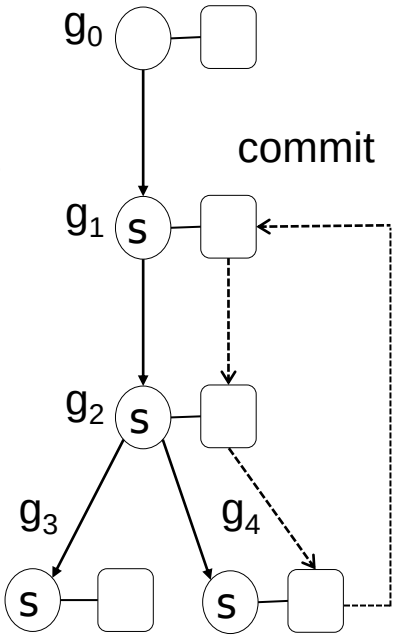

commitment

Results

> Multi-objective control

> coordination (speech, object manipulation, nodding, etc.).

> Anticipative control

> arbitrary time scales (word, sentence, full interview).

> continual predictions of the interviewer/ee's behavior.

> continual simulations.

→ early detection of opportunities and mishaps

- (re)planning questions.
- interrupting the interviewee early to avoid violating constraints later.

> Anytime control

> plans continually revised and kept ready for execution.

> response constrained by the timing/content of human behavior.

→ smooth interaction, in both content and form.