Open-World Learning for Radically Autonomous Agents

Pat Langley

Information and Technology Systems Division

Institute for Defense Analyses

Alexandria, Virginia

This work was supported by the Institute for Defense Analyses. Many of the ideas were inspired by DARPA's SAIL-ON Program, but it is not responsible for them. Thanks to Ted Senator, David Aha, and others for useful discussions.

Autonomous Systems: Progress and Limits

Autonomous agents are becoming more common and impressive in the form of:

- Self-driving cars
- Delivery drones
- Military robots
- Planetary rovers

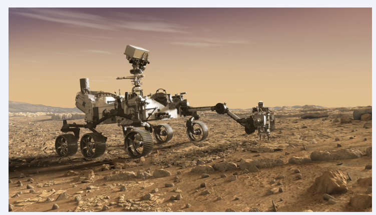

However, these systems depend on two critical assumptions:

- The environment will not change in substantial ways
- Their initial expertise will remain correct and accurate

These postulates will not hold in many real-world settings.

A Radically Autonomous Agent

Consider an unmanned underwater vehicle in a coastal area.

The system's expertise is accurate and its behavior good until:

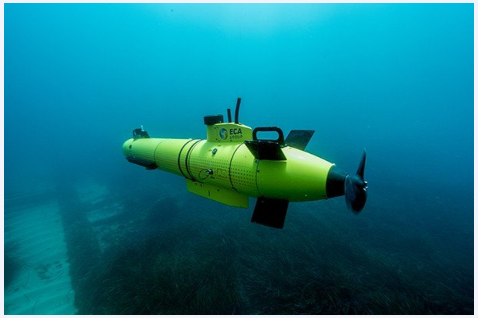

- Unfamiliar kelp fouls its propellers
- A large unknown predator attacks it
- A mysterious current drags it off course
- Sonar becomes distorted in a low-visibility area
- A nearby volcanic causes novel corrosive reactions

A radically autonomous system would realize these fall outside its expertise and learn rapidly enough to continue its mission.

A Radically Autonomous Agent

Consider an unmanned underwater vehicle in a coastal area.

The system's expertise is accurate and its behavior good until:

- Unfamiliar kelp fouls its propellers
- A large unknown predator attacks it
- A mysterious current drags it off course
- Sonar becomes distorted in a low-visibility area
- A nearby volcanic causes novel corrosive reactions

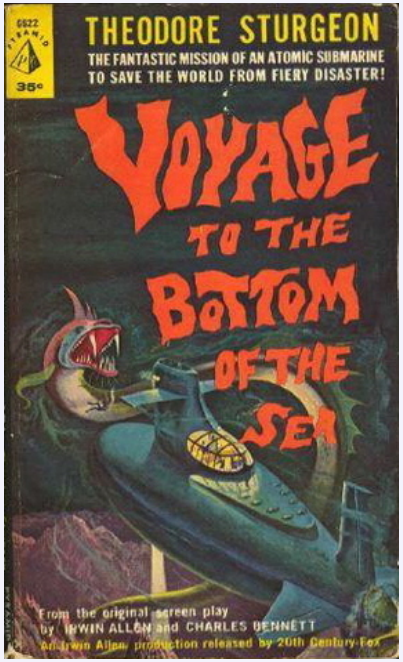

The agent would be as responsive and adaptable as the crew of the *Seaview* in *Voyage to the Bottom of the Sea*.

Open-World Learning

We will refer to such autonomous agents as *open-world learners*.

We can specify the task of open-world learning as:

- *Given:* An agent architecture that can operate in some class of environmental settings
- *Given:* Expertise that supports acceptable performance in these environments
- *Given:* ***Limited*** experience after ***sudden***, ***unannounced*** changes to the environment degrade performance
- *Find:* ***When*** the environmental change occurs and ***what*** revised expertise gives acceptable performance

This formulation applies to many agents and situations, whether initial expertise is handcrafted or learned.

Components of Open-World Learning

This problem statement suggests four main component abilities:

- A *performance* element that uses available expertise to pursue tasks and achieve goals
- A *monitoring* element that compares its observations with expectations to detect anomalies
- A *diagnostic* element that localizes expertise faults, generates hypotheses, and evaluates candidates
- A *repair* element that revises problematic expertise to correct agent behavior

These components must be embedded in an agent architecture that interfaces with the environment.

The Technical Challenge

But why is this a challenge? Modern learning techniques can do almost anything, right? ***Unfortunately, no***.

Remember environmental shifts are *sudden* and *unannounced*, and expertise repair must be *rapid*.

- Statistical supervised induction? X
  - *Nonincremental, requires too many labeled cases*
- Reinforcement learning? X
  - *Requires far too many trials, no simulator available*

Mainstream approaches are ill suited for open-world learning.

Some Relevant Paradigms

However, as Senator (2019) notes, some other paradigms are far more relevant:

- Model-based diagnosis / repair
- Metacognitive processing
- Problem reformulation
- Theory / model revision
- Computational scientific discovery
- Structural transfer of learned expertise

How does these conceptual frameworks differ from modern statistical learning? What do they offer?

The Need for Theories

To make open-world learning tractable, we must constrain its operation without overly restricting it.

In classic ML terms, we must provide a strong *inductive bias* that limits search through the model space.

- Hypothesis: *A theory of physical environments and transforms over them is necessary for effective open-world learning*.

Such a theory would reduce the time needed to detect changes and to repair the agent's expertise in response.

Note: A theory can be *domain independent* and very general yet still provide strong constraints.

Theories of Environmental Change

What form should a theory of environment change take?

It might be procedural and implicit, but this has disadvantages; in contrast, a *declarative* theory would explicitly specify:

- Initial environments an agent may encounter
- Transformations that can alter these environments
- Distributions of such transformations over time

Each such specification will define a space of environments and possible trajectories through the space.

In classic machine learning, this is known as a *declarative bias*.

A Framework for Environmental Change

Any theory of environmental novelty requires that we make some limiting assumptions.

Here we posit that an environment includes four constituents:

- A *spatio-temporal* matrix
- *Structures* that can occur in this matrix
- *Processes* that can operate on these structures
- *Constraints* on these structures and processes

Other frameworks are possible, but this should let us describe environmental change in many single-agent settings.

Note: The extension to multiple agents is a more complex story.

Spatio-Temporal Changes

The spatio-temporal matrix provides the basic physical setting in which an agent operates.

Some possible spatio-temporal transformations include:

- Increase in spatial extent (e.g., a larger undersea cavern)
- New attributes added to spatial field (e.g., salinity, viscosity)
- Altered distribution of values for spatial field
  - Shift from constant to location-dependent viscosity
  - Direction of water or air currents suddenly reversed

Such changes can have a pervasive influence on the agent’s and other entities’ behaviors.

Changes to Structures

Structures comprise the entities that populate the world, along with their associated characteristics.

Some possible structural transformations include:

- Adding / removing object categories (e.g., rock forms, species)
- Introducing new attributes for objects (e.g., texture, odor)
- Altering distributions of attribute values (e.g., color, height)
- Shifting composite object categories (e.g., bond strength)

These changes affect entities that the agent observes and manipulates to achieve its goals.

Changes to Processes

Processes describe behavior that causes the agent's environment to evolve over time.

Some possible process-related transformations include:

- Introducing new *physical* processes (e.g., hail, reactions)
- Adding new *control* processes (e.g., jumping, grasping)
- Removing *perceptual* processes (e.g., losing vision, hearing)
- Altering process parameters (e.g., reaction rate, turn speed)

These changes affect the dynamics of the natural word and how the agent interacts with it.

Changes to Constraints

Constraints specify what situations and events occur, including conditions on them.

Some possible constraint-related transformations include:

- Changing entities involved in processes (e.g., in digestion)
- Altering conditions on physical processes (e.g., temp limits)
- Revising conditions on control processes (max weight liftable)

In multi-agent settings, constraints also encode social norms that describe other agent's goals and activities.

Altering these norms can shift social behavior in drastic ways.

Experiments in Open-World Learning

Machine learning has been an empirical discipline for decades, but open-world learning raises special challenges.

We need new experimental methods to let us study agents that exhibit this ability.

- Our goal should *not* be to show that one system is superior to others using mindless bakeoffs.
- We should aim instead to understand *when* and *why* open-world learning succeeds or fails.

For the latter, we need systematic experiments that test specific hypotheses of interest.

Domains and Testbeds

Natural domains for research on open-world learning include:

- Underwater exploration
- Aerial reconnaissance
- Self-driving taxis
- Planetary rovers
- Robotic spelunkers

But real-world environments offer only limited ability to inject change; they provide *relevance* but little *insight*.

This makes *high-fidelity simulators* more suitable for controlled experiments and scientific understanding.

Dependent Measures

Scientific experiments collect *dependent measures* that examine some aspect of behavior.

- In AI, these focus on system performance (e.g., goals achieved)
- In ML, they may be higher-order measures (e.g., learning rate)

For open-world learning, we are especially interested in:

- The time needed for an agent to detect environmental change
- The time needed for the agent to recover from such a change

Experiments should also average across multiple runs to ensure trustworthy results.

Independent Variables

Classic experiments on learning alter two types of independent variables:

- Features of the system (e.g., lesion / parametric studies)
- Features of the domain (e.g., irrelevant attributes, noise level)

The first is still relevant to studies of open-world learning, but the latter is replaced by:

- Features of environmental change (e.g., type, frequency)

These independent factors should be elements in one's theory of environmental novelty.

Hypotheses About Open-World Learning

Good AI experiments are designed to answer specific questions.

For open-world learning, we want to know what types of change are difficult and what agent features can handle them.

We can specify and test specific experimental hypotheses, like:

- Agents will recover from changes to structures more easily than to changes in processes (Senator, 2019).
- Increasing perceptual noise will slow ability to detect change, bit it will not effect on diagnosis and repair.

Researchers should think about which questions they want to answer and design experiments in response.

Key Ideas of the Talk

Open-world learning to detect *unannounced*, *sudden* changes and revise their expertise *rapidly* in response.

This talk defined and analyzed open-world learning, including:

- *Why it poses a challenge to mainstream methods*
- *What components it needs (perform, monitor, diagnose, repair)*
- *How a theory of environmental change makes it tractable*
- *Elements of a theory (structure, process, constraint changes)*
- *How to study this ability with controlled experiments*

The ability to learn in open worlds is essential for radically autonomous agents that operate in remote settings.

Radical Autonomy

Future intelligent agents will need to pursue extended missions in unfamiliar environments without assistance.

Thus, they must be *radically* autonomous in that they cannot:

- Rely on the environment to remain unchanged
- Assume that their initial expertise will remain correct

To operate effectively in such settings, the agents must detect changes and respond rapidly.

In other words, they must exhibit *open-world learning*.

OWLs from Media

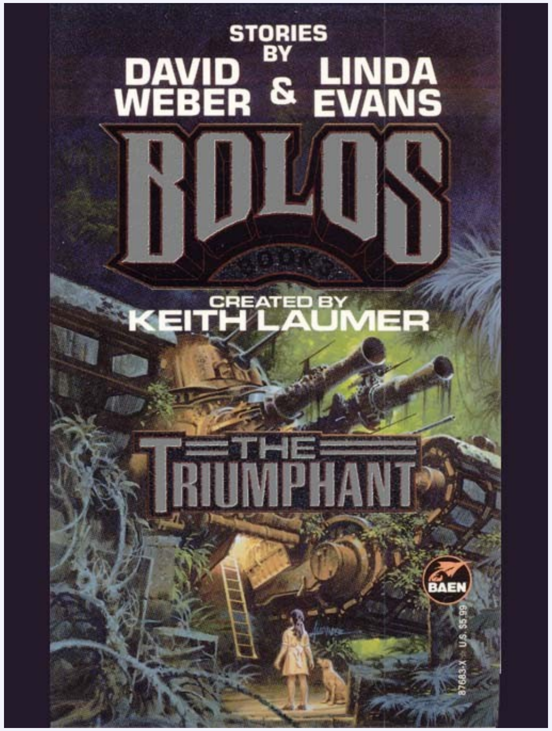

Bolos

WALL-E

Bubo

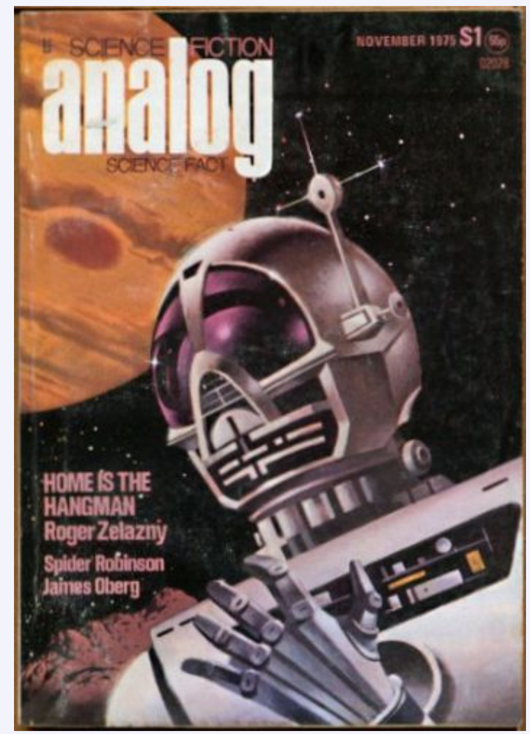

Hangman