

Welcome to the AI Seminar Series

Today's USC/ISI Speaker:

Pat Langley
Stanford University

**The Computational Gauntlet
of Human-Like Learning**



Upcoming Schedule:

=====

Remainder of December 2022

No Talks Scheduled

=====

Friday, January 13 @ 11 AM PST / 2 PM EST

Speaker: Jeff Krichmar, UCSD

Title: TBA

=====

Friday, January 20 @ 11 AM PST / 2 PM EST

Speaker: Yu-Xiang Wang, UCSB

Title: Deep Learning Meets Nonparametric
Regression: Are Weight-decayed DNNs locally
adaptive?

The Computational Gauntlet of Human-Like Learning

Pat Langley

Institute for the Study of
Learning and Expertise

Center for Design Research
Stanford University

This research was supported by ONR Grant N00014-20-1-2643 and by AFOSR Grant FA9550-20-1-0130. Thanks to many colleagues for discussions that led to these ideas.

The Problem

Despite its modest origins, machine learning has come to play a dominant role in artificial intelligence.

Statistical induction on massive data sets has led to impressive results in multiple areas, including:

- *Computer vision*
- *Natural language*
- *Game playing*

But in the process, the field has lost its intellectual diversity and abandoned its conceptual roots.

Claim: *We can remedy both drawbacks, and devise even more effective systems, by focusing on **human-like learning**.*

Example: Learning Mathematics

Consider how students master mathematics in our educational system by learning, successively, to:

- *Recognize and write digits* $3 \quad 7 \quad 6 \quad 4$
- *Retrieve and use arithmetic tables* $2 \times 1 = 2, 2 \times 2 = 4, 2 \times 3 = 6$
- *Carry out multi-column addition, subtraction*
$$\begin{array}{r} 327 \\ - 64 \\ \hline \end{array}$$
- *Simplify complex fractions* $3/4 + 1/8 = ?$
- *Solve algebraic equations, word problems* $7x - 5 = 2x$

This curriculum takes years, but it does ***not*** require thousands of instances per concept or skill.




The trajectory of human learning here differs drastically from how we currently train machines.

Example: Learning to Drive

Now consider how people – often teenagers – learn to drive an automobile by acquiring:

- *Categories for roads, lanes, intersections, signs*
- *Skills for changing lanes, passing, turning, parking*
- *Social norms for driving, including laws and customs*

Mastering these elements requires training and practice, but most drivers are reasonably good after a short course.



Unlike statistical learners, humans do ***not*** need millions of miles' experience to acquire basic competence.

Machine Learning: A Brief History

Machine learning was founded more than four decades ago as a spinoff of mainstream AI:

- *First workshop 1980, Journal 1986, Conference 1988*
- *Focused initially on acquiring symbolic structures*
- *Concerned with automating creation of expert systems*
- *But also with modeling **high-level learning in humans***

This paradigm was successful, producing demonstrations of new capabilities and deployed systems.

During this early period, links to cognitive psychology played key roles in the field's aims and progress.

Machine Learning: More History

The new discipline of machine learning evolved rapidly and, by the mid-1990s, it had:

- *Redefined learning as improvement of performance*
- *Broadened to include statistical methods and neural nets*
- *Adopted controlled experiments for evaluation purposes*
- *Birthered the closely related discipline of data mining*

Each step seemed a positive one but also took the field further away from its psychological origins.

More recent results on learning with deep neural networks have only worsened the situation.

Constraints on Learning

To develop AI systems that learn like people, we must first identify the core features of human learning:

- *High-level regularities* observed in human cognition
- *Recurring phenomena* that hold across many settings
- *Laws of qualitative structure* (Newell & Simon, 1976)
- **Not** detailed models that fit specific experimental results

Insights about the character of human learning can serve as strong constraints on system design.

But how might researchers use such constraints effectively?

A Computational Gauntlet

A *gauntlet* is a passage, lined with armed adversaries, that one must traverse to survive a trial.

- We can use characteristics of human learning to devise a *computational* gauntlet.
- Each constraint introduces a new threat that AI systems must encounter and overcome.
- To reach the end, they must make it past each obstacle along the dangerous path.



This offers a radical alternative to the performance-oriented ‘bake offs’ that now guide the field.

But what aspects of human learning can serve this purpose?

Modular Structures

One basic feature of human learning (Bower, 1981) concerns the nature of acquired content:

- *Learning involves the acquisition of modular cognitive structures.*

This does not specify the structures' details; only that expertise consists of *discrete* mental elements.

Candidates include *concepts, production rules, cases, chunks*, and even *stimulus-response pairs*.

But each contrasts sharply with the idea that learning produces a large-scale monolithic structure.

Composable Elements

A second characteristic is enabled by the first one and often associated with it closely:

- *Learned cognitive structures can be composed during performance.*

That is, relevant elements of expertise are accessed and then combined *as needed* to produce behavior.

Production systems (Klahr et al., 1987), *logic programs*, and *generative grammars* comprise composable elements.

These differ from the large structures (e.g., neural networks or decision trees) produced by most statistical induction.

Piecemeal Acquisition

Another feature involves how people process experiences and create structures. In particular:

- *Expertise is acquired in a piecemeal manner, with one element added at a time.*

Humans learn one cognitive structure, then another, continuing until they achieve broad coverage.

E.g., they acquire each concept and skill for mathematics and driving in a reasonably independent manner.

They do not create complex models *en masse*, as done by most methods for statistical induction.

Incremental Learning

Another processing constraint focuses not on the knowledge elements but on handling training cases:

- *Learning is an incremental activity that processes one experience at a time.*

This is linked to on-line approaches that interleave learning tightly with performance mechanisms.

E.g., people process the training events for mathematics and driving in an ongoing stream, not all at once.

Incremental and piecemeal learning can co-occur, but they are orthogonal; most rule induction is piecemeal but batch.

Guidance from Knowledge

The sequential nature of human learning also means that later processing builds on previous results:

- *Learning is guided by knowledge that aids interpretation of new experiences.*

Because acquisition is piecemeal and incremental, it occurs in the context of existing mental structures.

E.g., complex skills for both mathematics and driving build on simpler ones acquired earlier in training.

Knowledge is central to human learning but it receives limited attention in data-intensive paradigms.

Rapid Acquisition

A final characteristic of human learning, enabled by piecemeal, incremental, and knowledge-guided processing, is that:

- *Cognitive structures are acquired and refined rapidly, each from small numbers of training cases.*

The claim is *not* that *all* expertise comes from a few instances, but that we learn modular elements this way.

Human learning curves in mathematics and driving, which plot performance vs. training cases, are very steep.

Again, this diverges from statistical induction's dependence on thousands or millions of items.

Critiques and Responses

- Why change paradigms when deep learning works so well?
 - Because it is ***not*** as data efficient as human learners. And we should understand the ***entire space*** of learning methods.
- Why build AI systems that learn like people? (planes \neq birds)
 - Birds offer many insights into flight (e.g., *lift*, *thrust*, and *drag*). And we now have small drones that fly very much like birds.
- Do you mean structure learning vs. parameter estimation?
 - No, neural nets are often used for tasks that involve structure. But human-like learners are likely to fare better on them.
- Are you saying that human learning never involves statistics?
 - No, but the rapid acquisition of new structures is a distinctive feature of human learning; statistics is a *background* process.

Examples of Human-Like Learning

The literature contains some cases of human-like learning that count as positive instances:

- Fisher's (1987) *Cobweb* – constructs a probabilistic conceptual taxonomy from unsupervised training cases
- Minton's (1990) *Prodigy* – acquires control rules from planning traces to guide search on future problems
- McClure et al.'s (2015) *SAGE* – invokes structural analogy to learn relational concepts from training sequences
- Muggleton et al.'s (2018) *meta-interpretive learning* – acquires logic-based concepts very rapidly

These systems run the gauntlet successfully and offer useful role models for the research community.

Probabilistic Concept Hierarchies

Fisher's (1987) Cobweb is a process model of categorization and category learning that:

- Constructs a taxonomy of probabilistic concept descriptions
 - Terminal nodes are cases; nonterminals summarize descendants
- Sorts new cases down the hierarchy guided by *category utility*
 - On halting, uses selected concept node to predict missing values

Cobweb unifies ideas from the *decision tree*, *naive Bayes*, and *nearest neighbor* frameworks.

The system replicates well-known psychological phenomena like basic-level categories and typicality effects.

Learning Concept Hierarchies

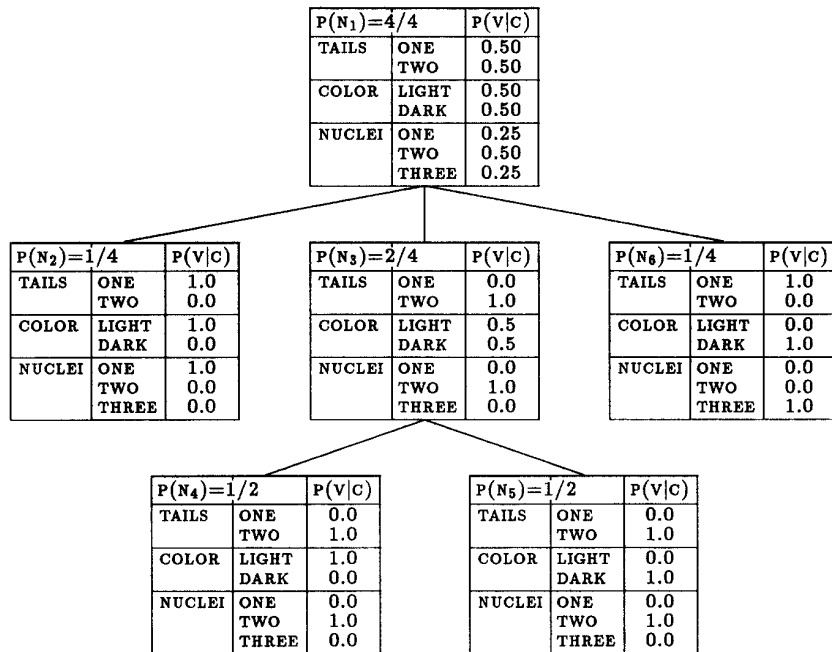
Cobweb interleaves learning with categorization to construct its hierarchy from unsupervised data by:

- Updating the distribution for concepts to which a case is sorted
- Extending the taxonomy downward on reaching a terminal node
- Adding a new branch when no children are similar enough
- Merging / splitting a node's children if category utility improves

The system learns categories very rapidly in an incremental, piecemeal way that builds on prior acquisitions.

Each Cobweb training case leads to both the creation of new cognitive structures and revision of statistical summaries.

Cobweb Examples and Results

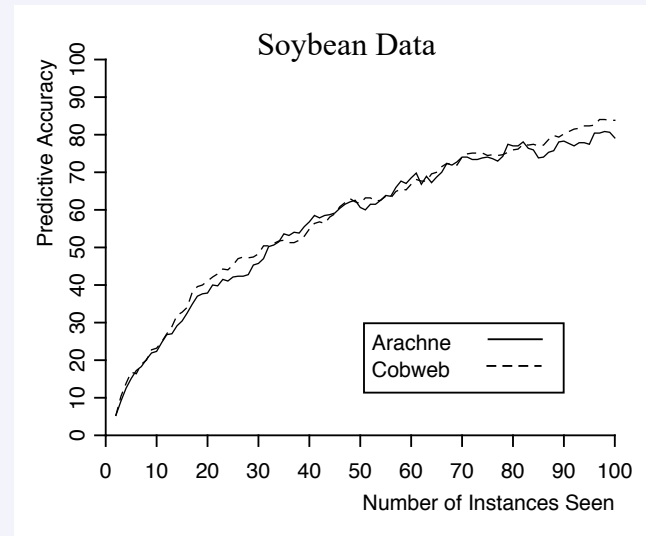
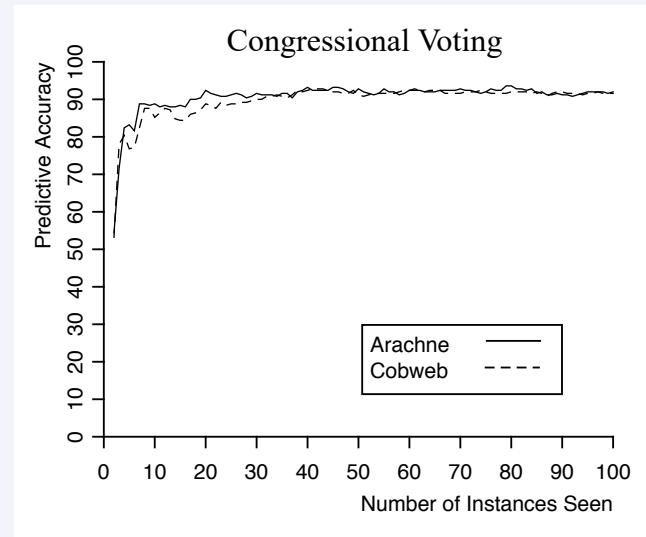


A Cobweb Hierarchy

Category Utility

$$\frac{\sum_{k=1}^K P(C_k) \sum_i \sum_j P(A_i = V_{ij} | C_k)^2 - \sum_i \sum_j P(A_i = V_{ij})^2}{K}$$

Cobweb originally dealt with ‘tabular’ encodings, but recent extensions incorporate *convolutional* processing of images and *context* in language.



Knowledge-Guided Planning

Minton's (1990) Prodigy offers an architecture for knowledge-guided planning that:

- Encodes knowledge as domain operators and control rules for selecting or rejecting goals, operators, or bindings
- Invokes means-ends analysis to carry out goal-directed search in a space of problem decompositions
- Uses control rules to reduce search by blocking poor choices and favoring good ones

Prodigy's reliance on means-ends analysis is consistent with studies of human problem solving.

The system unifies AI's four key ideas: *reasoning*, *heuristic search*, *knowledge*, and *learning*.

Learning Search-Control Knowledge

Prodigy acquires planning expertise from traces of its own search processes by:

- Using a generic theory of problem solving to explain why each choice led to success or failure
- Compiling each explanation into a control rule for selecting or rejecting a goal, operator, or binding
- Collecting statistics on these rules' utilities to determine which ones to retain or abandon

The system substantially reduces both nodes examined and CPU time to solve new problems in many domains.

Prodigy combines rapid generation of new structures through knowledge with their gradual evaluation by statistics.

Prodigy Examples and Results

A Prodigy Rejection Rule

```
(FAILS goal node)
if (AND (CURRENT-GOAL node goal)
        (MATCHES goal (ON x y))
        (OR (AND (KNOWN (ON-TABLE y))
                  (EQUAL x y))
            (AND (KNOWN (ON y z))
                  (EQUAL x y))
            (AND (KNOWN (HOLDING y))
                  (EQUAL x y)))))
```

A Prodigy Selection Rule


```
(SUCCEEDS goal node)
if (AND (CURRENT-GOAL node goal)
        (MATCHES goal (HOLDING x))
        (KNOWN node (IS-BLOCK x))
        (KNOWN node (ON-TABLE x)))
```

Prodigy's Utility Criterion


$$Utility = (AvrSavings \times ApplicFreq) - AvrMatchCost ,$$

Prodigy's many successors supported planning by abstraction, analogical problem solving, and learning for plan quality rather than efficiency.

STRIPS Robot



Scheduling



Analogical Concept Learning

McClure et al.'s (2015) SAGE can acquire complex concept descriptions from labeled training cases by:

- Representing each concept as a set of relational literals with associated probabilities
- For each new training case T:
 - Using structural analogy to retrieve descriptions similar to T and selecting the best candidate C
 - If C and T match well enough, then using T to update C's probabilities and to add new relations
 - Else storing a new disjunctive description based on case T

SAGE learns geographical concepts, musical genres, and object shapes far more rapidly than statistical methods.

Meta-Interpretive Learning

Muggleton et al. (2018) report a new abductive approach to learning relational logic programs by:

- Using domain-independent knowledge stated as logical rules
- Searching for simple explanations that cover each case
- Positing new predicates that may be reused in the explanation
- Transforming the explanations into domain rules for later use

This meta-interpretive scheme induces visual concepts and control programs from much less data than statistical ones.

This work shows that representation learning is not limited to deep neural networks.

Comparison of Characteristics

We can compare how these systems – and the most popular class of statistical learners – fare on the computational gauntlet.

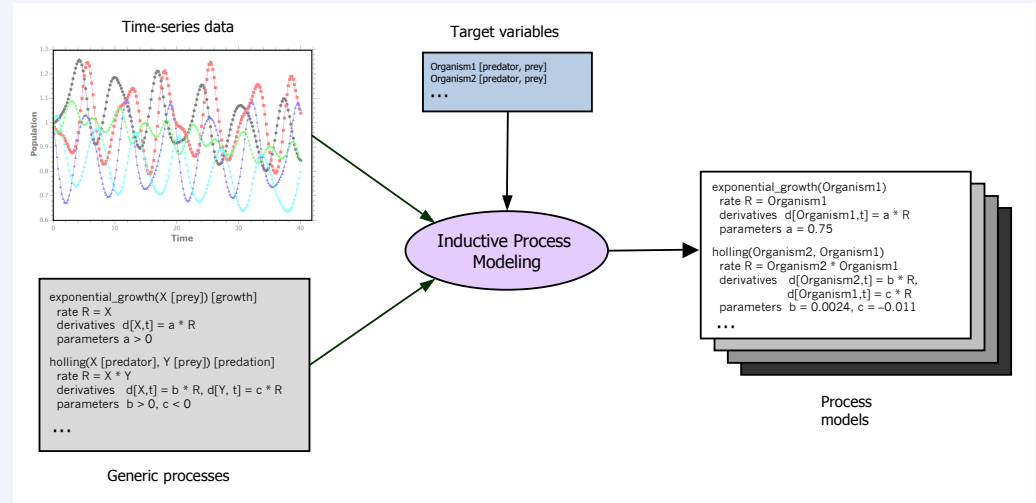
Characteristics	Deep Net	Cobweb	Prodigy	SAGE	Meta-Int
Modular structures	○	●	●	●	●
Composable elements	⊙	○	●	○	●
Piecemeal learning	○	●	●	●	●
Incremental processing	○	●	●	●	○
Knowledge guidance	⊙	●	⊙	●	⊙
Rapid acquisition	○	●	●	●	●

The table shows that Cobweb, Prodigy, SAGE, and meta-interpretive learning all pass most of its challenges.

Inductive Process Modeling (Langley, 2019)

Inductive process modeling constructs explanations of multivariate time series from generic background knowledge.

Discovered models are sets of *differential equations* organized into higher-level *processes*.




A Quantitative Process Model

(a) organism_loss[phyto, detritus]
 rate $r = \text{phyto}$
 equations $d[\text{phyto}, t] = -0.05 * r$
 $d[\text{detritus}, t] = 0.05 * r$
 nutrient_uptake[phyto, nitro]
 rate $r = \text{phyto} * \text{nitro}$
 equations $d[\text{phyto}, t] = 0.5 * r$
 $d[\text{nitro}, t] = -0.005 * r$
 remineralization[detritus, nitro]
 rate $r = \text{detritus}$
 equations $d[\text{detritus}, t] = -0.04 * r$
 $d[\text{nitro}, t] = 0.04 * r$

(b) $d[\text{phyto}, t] = -0.05 * \text{phyto} + 0.5 * \text{phyto} * \text{nitro}$
 $d[\text{nitro}, t] = -0.005 * \text{phyto} * \text{nitro} + 0.04 * \text{detritus}$
 $d[\text{detritus}, t] = 0.05 * \text{phyto} + -0.04 * \text{detritus}$

Trajectories for 20-Organism Food Chain



Hierarchical Problem Networks

(Langley, 2022)

HPNL learns *hierarchical problem networks* from sample solutions, one method for each subproblem.

The system uses *domain constraints* to identify conditions on methods, not classic EBL or ILP techniques.

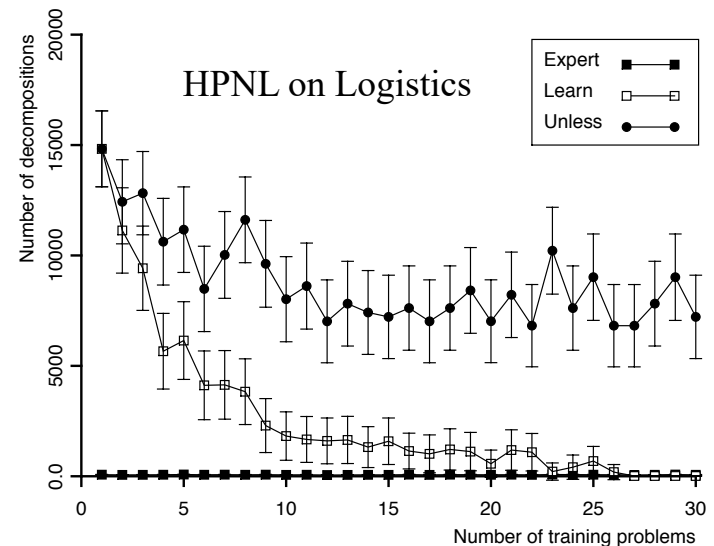
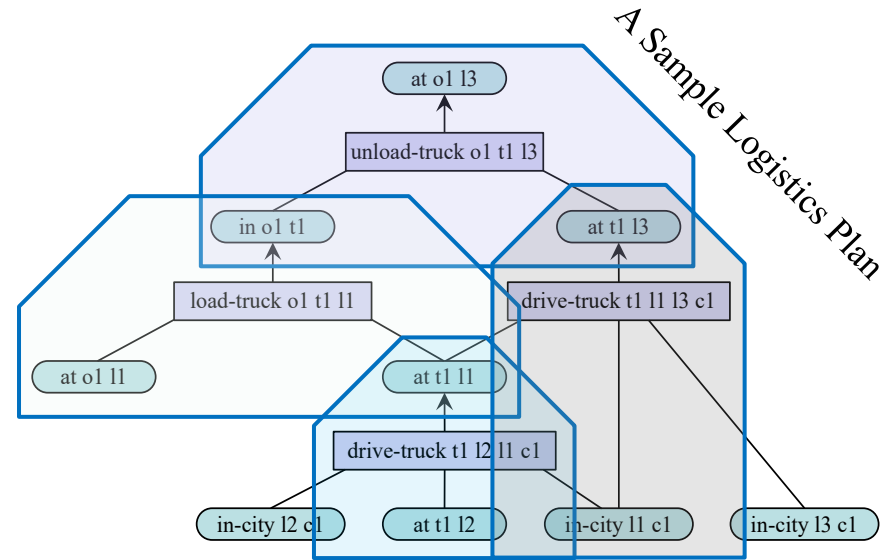
HPN Methods for Logistics

```
((at ?o1 ?l3))
conditions: ((object ?o1) (truck ?t1) (location ?l3) (location ?l1)
              (in-city ?l3 ?c1) (in-city ?l1 ?c1) (at ?t1 ?l3) (at ?o1 ?l1))
operator:   (unload-truck ?o1 ?t1 ?l3)
subproblem: ((at ?t1 ?l3) (in ?o1 ?t1))

((at ?t1 ?l1))
conditions: ((truck ?t1) (location ?l3) (location ?l1)
              (city ?c1) (in-city ?l3 ?c1) (in-city ?l1 ?c1)
              (in-city ?l2 ?c1) (at ?t1 ?l3))
operator:   (drive-truck ?t1 ?l3 ?l1 ?c1)
subproblem: ((at ?t1 ?l3))
unless-goals: ((in ?o ?t1)))

((in ?o1 ?t1))
conditions: ((object ?o1) (truck ?t1) (location ?l1) (location ?l3)
              (in-city ?l1 ?c1) (in-city ?l3 ?c1) (at ?t1 ?l3) (at ?o1 ?l1))
operator:   (load-truck ?o1 ?t1 ?l1)
subproblem: ((at ?t1 ?l1) (at ?o1 ?l1))

((in ?o1 ?t1)
:conditions ((object ?o1) (truck ?t1) (location ?l1) (airport ?l1)
              (location ?l2) (location ?l3) (in-city ?l1 ?c1) (in-city ?l2 ?c1)
              (in-city ?l3 ?c2) (at ?t1 ?l2) (at ?o1 ?l3))
:operator   (load-truck ?o1 ?t1 ?l1)
:subproblem ((at ?t1 ?l1) (at ?o1 ?l1))
```



Fostering Work on Human-Like Learning

Research on human-like learning was once widely accepted by the AI community. How can we restore this vision?

- *Broaden education to cover classic methods*
- *Expand funding to support human-like approaches*
- *Establish publication venues that value such work*
- *Champion evaluation with **computational gauntlets***

Together, these steps can help create a *Zeitgeist* that recaptures the spirit of early AI and machine learning.

This call to arms echoes similar appeals by Marcus and Davis (2021), Fahlman (2012), and others.

Summary Remarks

Machine learning, despite impressive advances, has abandoned many of its early, profound insights.

A promising alternative is to develop AI systems that learn in a more human-like manner by:

- *Acquiring modular, composable structures in a piecemeal, incremental way, aided by knowledge, from little data.*

We can treat these features as design constraints that define a *computational gauntlet* for learning systems.

I call on audacious AI researchers to tackle this challenge.

The Computational Gauntlet

