Building Machines that Learn and Think like People

Current deep learning methods have achieved remarkable achievement in recognition & control, demonstrating the power of ***gradient-based learning (ERM)*** and deep hierarchies of latent variables (many probabilistic machine learning have also developed a lot, but it's not in the scope of this paper). However, truly human-like learning and thinking machines will have to reach beyond current engineering trends in both what they learn and how they learn it, specifically, we need to incorporate perspective from cognitive scientists & psychology researchers:

1. ***Ground learning*** in intuitive theories of physics and psychology to support and enrich the knowledge that is learned.
2. Build ***causal models of the world*** that support explanation and understanding, rather than merely solving pattern recognition problems.
3. Harness ***compositionality and learning-to-learn*** to rapidly acquire and generalize knowledge to new tasks and situations.

There are 2 different approaches towards computational intelligence:

1. The ***statistical pattern recognition*** approach treats prediction as primary, usually in the context of a specific classification, regression, or control task.
    a. Learning is about discovering features that have high-value states in common (a shared label in a classification setting or a shared value in a reinforcement learning setting) across a large, diverse set of training data.
2. The ***cognitive model building*** focuses on the models of the world and explanations where learning is the process of model building.
    a. Cognition is about using these models to understand the world, to explain what we see, to imagine what could have happened that didn't, or what could be true that isn't, and then planning actions to make it so.

The difference between pattern recognition and model building, between prediction and explanation, is central to our view of human intelligence. ***Just as scientists seek to explain nature, not simply predict it, we see human thought as fundamentally a model building activity***.

Assumptions For Computation Intelligence

Any computational model of learning must ultimately be grounded in the brain's biological neural networks. As long as natural intelligence remains the best example of intelligence, the project of reverse engineering the human solutions to difficult computational problems will continue to inform and advance AI.

- Future generations of neural networks will look very different from the current state-of-the-art neural networks. They may be endowed with ***intuitive physics***, ***theory of mind***, ***causal reasoning, and more***.

- More structure and inductive biases could be built into the networks or ***learned from previous experience with related tasks***, leading to more human-like patterns of learning and development.
- ***Learning with few data (learn-to-lean)***: Networks may learn to effectively search for and discover new mental models or intuitive theories, and these improved models will, in turn, enable subsequent learning, allowing systems that learn-to-learn using previous knowledge to make richer inferences from very small amounts of training data.

Reverse engineering human intelligence can usefully inform AI and machine learning, especially for the types of domains and tasks that people excel at, including concept learning, scene understanding, language acquisition, language understanding, speech recognition, creativity, common sense, and general-purpose reasoning.

Three Ingredients for Intelligence

There are a few sets of ingredients that are important to human intelligence & learning. It can be shown that no-pure computation in one area is good enough, a collaboration across domains is what makes intelligence.

1. ***Developmental start-up software***, or cognitive capabilities present early and are fundamental to development. These types of inferences further accelerate the learning of new tasks.
    a. ***Intuitive Physics***: infants know that objects will persist over time and that they are solid and coherent.
    b. ***Intuitive Psychology***: infants understand that other people have mental states like goals and beliefs, and this understanding strongly constrains their learning and predictions.
2. ***Model building*** is the hallmark of human-level learning, or explaining observed data through the construction of causal models of the world.
    a. The early present capacities for intuitive physics and psychology are also causal models of the world.
    b. Primary job of learning is to extend and enrich these models and to build analogous causally structured theories of other domains.
    c. Children come with the ability and the desire to uncover the underlying causes of sparsely observed events and to use that knowledge to go far beyond the paucity of the data.
    d. It might seem paradoxical that people are capable of learning these richly structured models from very limited amounts of experience. We suggest that ***compositionality*** and ***learning-to-learn*** are ingredients that make this type of rapid model learning pos- sible
3. It is remarkable how fast we can perceive, ***think***, and ***put action in real time*** to act.
    a. A ***model-free*** method can accelerate slow ***model-based*** inferences in perception and cognition.

    i. By learning to recognize patterns in these inferences, the outputs of inference can be predicted without having to go through costly intermediate steps (AlphaGo).
    ii. Integrating neural networks that "learn to do inference" with rich model building learning mechanisms offers a promising way to explain how human minds can understand the world so well and so quickly.
b. Once a causal model of a task has been learned, humans can use the model to ***plan action sequences that maximize future reward (RL)***. We review evidence that humans combine model-based and model-free learning algorithms both competitively and cooperatively and that these interactions are ***supervised by metacognitive processes***.
    i. The sophistication of human-like reinforcement learning has yet to be realized in AI systems, but this is an area where crosstalk between cognitive and engineering approaches is especially promising.

Symbolic to Sub-symbolic Computations

Understand the perspective, understand where you stand, and understand the challenges

Symbolic Computations

Alan Turing suspected that it was easier to build and educate a child machine than to try to fully capture adult human cognition (Turing 1950).

- Turing pictured the child's mind as a notebook with "***rather little mechanism and lots of blank sheets," and the mind of a child machine as filling in the notebook by responding to rewards and punishments***, similar to reinforcement learning.
- It is a behavioral psychology perspective and also perspective of ***empiricism of modern connectionist*** models–the idea that we can learn almost everything we know from the ***statistical patterns of sensory inputs***.

A similar sentiment was expressed by Minsky (1974): "I draw no boundary between a theory of human thinking and a scheme for making an intelligent machine; ***no purpose would be served by separating these today since neither domain has theories good enough to explain***—or to produce—enough mental capacity"

Much of this research assumed that human knowledge representation is symbolic and that reasoning, language, planning and vision could be understood in terms of ***symbolic operations***.

Sub-symbolic Computations

Later perspective breaks the symbolic perspective further into sub-symbolic computations, that thoughts about the nature of cognition is a ***parallel distributed processing (PDP)*** where we conduct parallel computation by combining simple units to collectively implement sophisticated computations (NN is the demo of sub-symbolic computations).

Neural network models and the PDP approach offer a view of the mind and intelligence more broadly that is sub-symbolic. The knowledge learned by these networks would be ***distributed across the collection of units rather than localized as in most symbolic data structures***.

- The PDP perspective is compatible with “model building” in addition to “pattern recognition.”
- Very little assumption should be built into the networks.
- Proponents of this approach maintain that many classic types of structured knowledge, such as ***graphs, grammars, rules, objects, structural descriptions, and programs***, can be useful yet ***misleading metaphors for characterizing thought***.
    - These ***structures are more epiphenomenal than real***, ***emergent properties of more fundamental sub-symbolic cognitive processes***.

More Than Distributed Is Needed

Concept learning + Generation + Prior + Higher level understanding + CL flexibility

A different picture has emerged that highlights the importance of early ***inductive biases***, including core concepts such as number, space, agency, and objects, as well as powerful learning algorithms that rely on ***prior knowledge to extract knowledge*** from small amounts of training data. This knowledge is often ***richly organized and theory-like in structure***, capable of the ***graded inferences*** and productive capacities characteristic of human thought. There may be 2 benchmarks in assessing performance.

1. Learning simple visual concepts (Supervised):
    a. Humans learn from ***less examples but form a rich representation***.
    b. Humans ***learn a concept***, that is a model of the class that allows their acquired knowledge to be flexibly applied in new ways, ***generating new examples***.
2. Learning to play the Atari game & Frostbite game (RL Control):
    a. Too expensive learning -> Different learned representation between machines and humans. Sparse feedback?
    b. Huma understand higher-level ***cues*** with small hints, but DQN needs to have sub-goals to actually learn, or it would be just trying random actions.
        i. DQN needs achievement of sub-goals and proceeding to next sub-goals.
    c. No flexibility to changes in game rules. People can learn models and use them for arbitrary new tasks and goals. Although neural networks can learn multiple mappings or tasks with the same set of stimuli (***adapting their outputs depending on a specified goal***) these models require substantial training or reconfiguration to add new tasks (Continual Learning Problem).
        i. Once the environment has been established for AlphaGo or any RL, the environment condition cannot be changed -> need to restart from scratch.

DQN does startes completely from scratch and humans have extensive prior experience before they even starts -> this may be important why we need to have foundation models? No boost in the flavor of large (almost too large set) of prior experiences?

Innate Intuitions For Grounding

Early in development, humans have a foundational understanding of several core domains.

- number (numerical and set operations),
- space (geometry and navi- gation),
- physics (inanimate objects and mechanics),
- psychology (agents and groups).

These c***ore domains cleave cognition at its conceptual joints***, and each domain is ***organized by a set of entities and abstract principles relating the entities to each other***. The underlying cognitive representations can be understood as “intuitive theories,” with a causal structure resembling a scientific theory.

Children seek out new data to distinguish between hypotheses, isolate variables, test causal hypotheses, make use of the data-generating process in drawing conclusions, and learn selectively from others.

Intuitive Physics For Generalization

At the age of 2 months, human infants expect inanimate objects to follow principles of persistence, continuity, cohesion, solidity, and believe objects should move along smooth paths. These expectations would guide later learnings.

People ***reconstruct a perceptual scene*** using ***internal representations*** of the objects and their physically relevant properties (such as mass, elasticity, and surface friction) and forces acting on objects (such as gravity, friction, or collision impulses).

- Relative to physical ground truth, the intuitive physical state representation is ***approximate and probabilistic, and oversimplified and incomplete in many ways***. Still, it is rich enough to support mental simulations that can predict how objects will move in the immediate future, either on their own or in responses to forces we might apply.

The ***intuitive internal physics engine*** approach enables flexible adaptation to a wide range of everyday scenarios and judgments in a way that goes beyond perceptual cues.

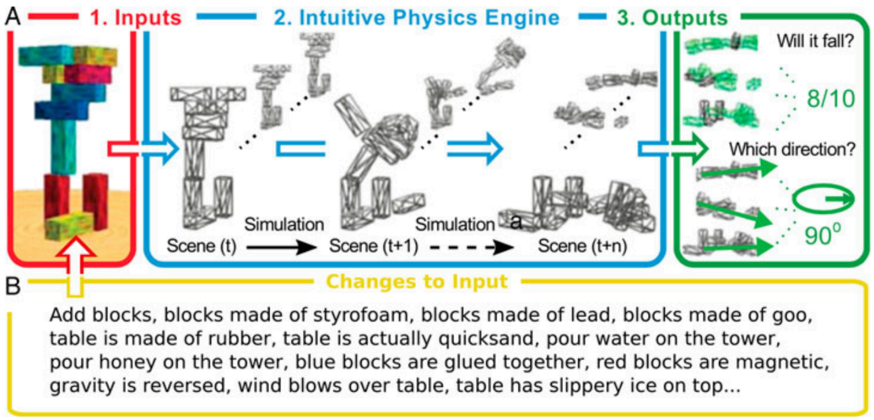

Instead of using a physics simulator, could neural networks be trained to emulate a general-purpose physics simulator, given the right type and quantity of training data? However, it is not sure if higher level would actually encode more generic physics properties instead of just task-specific values.

What is actually learned? is it just some specific things related to the task? Or is it a more generic understanding of the world also captured -> same performance results does not imply same learning.

Intuitive Psychology For Planning

At the very beginning, infants distinguish inanimate objects with animate objects to distinguish who their parents are. Infants also ***expect agents to act contingently and reciprocally, to have goals, and to take efficient actions toward those goals subject to constraints*** (these goals can be socially directed). at around 3 months of age, infants begin to discriminate antisocial agents that hurt or hinder others from neutral agents, and they later distinguish between anti-social, neutral, and pro-social agents.

- Does this give the ability of higher level goal understanding?
- Models formalize explicitly mentalistic concepts such as “goal,” “agent,” “planning,” “cost,” “efficiency,” and “belief,” used to describe core psychological reasoning in infancy.
    - Bayesian inverse planning, or Bayesian theory of mind (ToM)

Planning computations may be formalized as solutions to MDP or POMDP, taking as ***input utility*** and ***belief functions defined over an agent’s state-space*** and the agent’s ***state-action transition functions***, and returning a ***series of actions the agent should perform*** to most efficiently fulfill their goals (or maximize their utility).

- By simulating these planning processes, people can predict what agents might do next, or use ***inverse reasoning*** from observing a series of actions to infer the utilities and beliefs of agents in a scene.
- ***Direct analogous*** to how simulation engines can be used for intuitive physics, to predict what will happen next in a scene or to infer objects’ dynamical properties from how they move.

- simulation-based reasoning in intuitive psychology can be nested recursively to understand social interactions. ***We can think about agents thinking about other agents***.

Any full formal computational account of intuitive psychological reasoning needs to include representations of agency, goals, efficiency, and reciprocal relations.

- Let us ***infer the beliefs, desires, and intentions*** of the experienced player.
- It is an early ***emerging property*** that helps us to ***share with others cognitive ability***.
- Behavior is explained as acting under such belief, once inferred belief is established, no need for actual experience to learn.

Learning As Rapid Model Expansion

Gradient-based methods can be seen as a "gradual adjustment of connection strengths" with a large set of data. However, ***infants can grasp the boundary of the infinite set that defines each concept from the infinite set of all possible objects*** (learning words) without a large set of data. There may be some key concepts that should be considered more than just searching starting from nowhere. The three main objectives are all boosted by each other, increasing performance of one would lead to the increase of performance on the other and BPL is a current model that somewhat does pretty well on all three objectives.

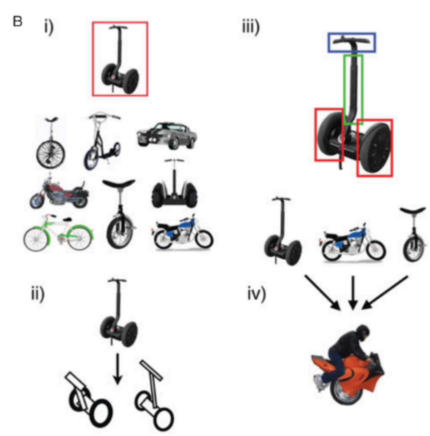

Even with just a few examples, people can learn remarkably rich conceptual models. One indicator of richness is the variety of functions that these models support. Beyond classification, concepts support prediction, action, communication, imagination, explanation, and composition. These abilities are not independent; rather they hang together and interact, coming for free with the acquisition of the underlying concept.

The ability of extending one learned concept or seeing the sub-components of one task to piece it into a sub-component of a new task may be very relevant for continual learning and general intelligence.

This richness and flexibility suggest that ***learning as model building is a better metaphor than learning as pattern recognition***. models are built upon rich domain knowledge rather than starting from a blank slate.

Bayesian Program Learning As Demo of Ideal Models

BPL represents concepts as simple stochastic programs: a structured procedures that generate new examples of a concept when executed. These programs allow the model to ***express causal knowledge*** about how the raw data are formed, and the probabilistic semantics allow the model to handle noise and ***perform creative tasks***. Structure sharing across concepts is accomplished by the compositional re-use of stochastic primitives that can combine in new ways to create new concepts.

A

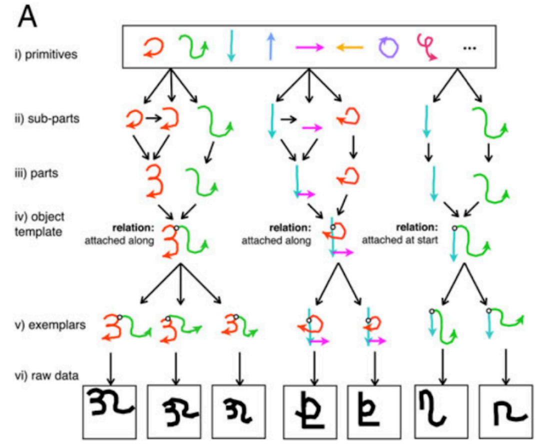

Note that we are overloading the word model to refer to the BPL framework as a whole (which is a ***generative model***), as well as the individual concepts (or ***probabilistic models***) that it infers from images to represent novel handwritten characters, there is a hierarchy of models:

1. A higher-level generative program that generates different types of concepts.
2. A lower-level probabilistic concept, which are themselves programs that can be run to generate tokens of a concept.

Here, describing learning as rapid model building refers to the fact that BPL constructs generative models that produce novel concepts, which generates probabilistic tokens of such concepts. The below is a visual turing test, some data are generated by humans and some by BPL.

B

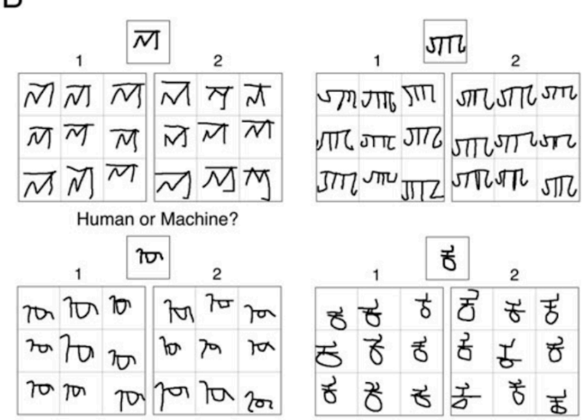

Compositionality

Compositionality is the classic idea that ***new representations can be constructed through the combination of primitive elements***. In computer programming, primitive functions can be combined to create new functions, and these new functions can be further combined to create even more complex functions. This function hierarchy provides an efficient description of higher-level functions, such as a hierarchy of parts for describing complex objects or scenes.

- Parts can themselves be composed of sub-parts, forming a "partonomy" of part-whole relationships.
- The parts and relations can be shared and re-used from existing related concepts.
- Because the parts and relations are themselves a product of previous learning, their facilitation of the construction of new models is also an example of learning-to-learn.

Compositionality is also at the core of productivity: ***an infinite number of representations can be constructed from a finite set of primitives***, just as the mind can think an infinite number of thoughts, utter or understand an infinite number of sentences, or learn new concepts from a seemingly infinite space of possibilities.

- This solves the infinite space exploration problem?
- object-oriented reinforcement learning & other RL methods for compositional learning.

Compositionality is also central to the construction of other types of symbolic concepts beyond characters, where new spoken words can be created through a novel combination of phonemes or a new gesture or dance move can be created through a combina- tion of more primitive body movements.

Casuality

Causal models represent hypothetical real-world processes that produce the perceptual observations. Causality has been influential in theories of perception. "Analysis-by-synthesis" theories of perception maintain that sensory data can be more richly represented by modeling the process that generated it. Relating data to their causal source provides strong priors for perception and learning, as well as a richer basis for generalizing in new ways and to new tasks.

In control and reinforcement learning, causal models represent the ***structure of the environment***, such as modeling state-to-state transitions or action/state-to-state transitions (somewhat model-based ideas).

- Although a generative model describes a process for generating data, or at least assigns a probability distribution over possible data points, this generative process may not resemble how the data are produced in the real world.
- Causality refers to the ***subclass of generative models that resemble, at an abstract level, how the data are actually generated***.
- Deep Belief Networks & VAE are on one spectrum of GAN while BPL is on the other because BPL resembles more to the actual hand-written process.
    - For the BPL of learning handwritten characters, causality is operationalized by treating concepts as motor programs, or abstract causal descriptions of how to produce examples of the concept, rather than concrete configurations of specific muscles.

Learning-to-Learn

When humans or machines make inferences that go far beyond the data, strong prior knowledge (or inductive biases or constraints) must be making up the difference. One way people acquire this prior knowledge is through "***learning-to-learn***", a term introduced by Harlow (1949) and closely related to the machine learning notions of "***transfer learning***", "***multitask learning***", and "***representation learning***". These terms refer to ways that learning a new task or a new concept can be accelerated through previous or parallel learning of other related tasks or other related concepts.

The strong priors, constraints, or inductive bias needed to learn a particular task quickly are often ***shared to some extent with other related tasks***. A range of mechanisms have been developed to adapt the learner's inductive bias as they learn specific tasks and then apply these inductive biases to new tasks.

- BPL transfers readily to new concepts because it learns about object parts, sub-parts, and relations, capturing learning about what each concept is like and what concepts are like in general.
- It is crucial that learning-to-learn occurs at multiple levels of the hierarchical generative process. ***Previously learned primitive actions and larger generative pieces can be re-used and re-combined to define new generative models for new characters***. Further transfer occurs by learning about the typical levels of variability within a typical generative model. This provides knowledge about how far and in what ways to generalize when we have seen only one example of a new character, which on its own could not possibly carry any information about variance.
- Deep reinforcement learning systems for playing Atari games have had some impressive successes in transfer learning. For example, the "actor-mimic" algorithm that first learns 13 Atari games by watching an expert network play and trying to ***mimic the expert network action selection*** and/or internal states.

Thinking Fast

The ***combination of rich models with efficient inference*** suggests another way psychology and neuroscience may usefully inform AI. It also suggests an additional way to build on the successes of deep learning, where ***efficient inference and scalable learning*** are important strengths of the approach. There needs to be a way to resolve the conflict between fast inference and structured representations, a collaboration.

Approximate Inference From Structured Models: AlphaGo

Computing a probability distribution over an entire space of programs is usually intractable, and often ***even finding a single high-probability program poses an intractable search problem***. In contrast, gradient-based learning is very fast even in a vast space. ***A complete account of learning and inference must explain how the brain does so much with limited computational resources***.

Looking at pure probabilistic inference, popular algorithms for approximate inference in probabilistic machine learning have been proposed as ***psychological models***. Most prominently, it has been proposed that humans can approximate Bayesian inference using ***Monte Carlo methods*** (stochastically sample the space of possible hypotheses and evaluate these samples according to their consistency with the data and prior knowledge). We are beginning to understand how such methods could be implemented in neural circuits.

- Although Monte Carlo methods are powerful and come with asymptotic guarantees, it is challenging to make them work on complex problems like program induction and theory learning, ***and it is unlikely that they are the only mechanism we use to process***.
- When the hypothesis space is vast and only a few hypotheses are consistent with the data, how can good models be discovered without exhaustive search (the full combinatorial complexity)?
- Humans use ***high-level abstract features of a domain*** to ***guide hypothesis selection***, by reasoning about distributional properties, dynamical properties, or monotonic relationships between causes and effects. Is there some guidance like that for the machines?

How might efficient mappings from questions to a plausible subset of answers be learned and making inference a smaller problem to tackle?

- One approach is to amortize probabilistic inference computations into an efficient feed-forward mapping (gradually lower the computation cost), somewhat a ***learning to do inference*** (independent from the ideas of learning as model building).
- These feed-forward mappings can be learned in various ways, for example, using paired generative/recognition networks and variational optimization, or nearest-neighbor density estimation.

This trend is an avenue of potential i***ntegration of deep learning models with probabilistic models and probabilistic programming***: training neural networks to help perform probabilistic inference in a generative model or a probabilistic program. Another avenue for potential integration is through ***differentiable programming***, by ensuring that the program-like hypotheses are differentiable and thus learnable via gradient descent.

- Neural networks with “working memories” that augment the shorter-term memory provided by unit activation and the longer-term memory provided by the connection weights.
- These developments are also part of a broader trend toward “differentiable programming,” the incorporation of classic data structures, such as random access memory, stacks, and queues, into gradient-based learning systems such as Neural Turing Machine (NTM) and Differentiable Neural Computer (DNC).

Gradual Model-free to Model-based

Model-based planning is an essential ingredient of human intelligence, ***enabling flexible adaptation to new tasks and goals***; it is where all of the rich model-building abilities discussed in the previous sections earn their value as guides to action.

Once the learned skills become “***habitized***”, a shift from model-based to model-free control can happen. This shift may arise from a rational arbitration between learning systems to ***balance the trade-off between flexibility and speed***.

- plans can be amortized into cached values by allowing the model-based system to simulate training data for the model-free system (Sutton 1990). This process might occur offline (e.g., in dreaming or quiet wakefulness), suggesting ***a form of consolidation in reinforcement learning***.
- Consistent with the idea of cooperation between learning systems in the human brain, a recent experiment demonstrated that model-based behavior becomes automatic over the course of training. Thus, ***a marriage of flexibility and efficiency might be achievable if we use the human reinforcement learning systems as guidance***.

Example: AlphaGo

This is AlphaGo, a really smart system, but it is not robust to the face of variants of the game Go. Does there exist a way in which we can reuse some of the trees that have been explored, to find some robust representation of the game Go that can be carried over even in new variants of the game Go to explore more parts of the new tree.

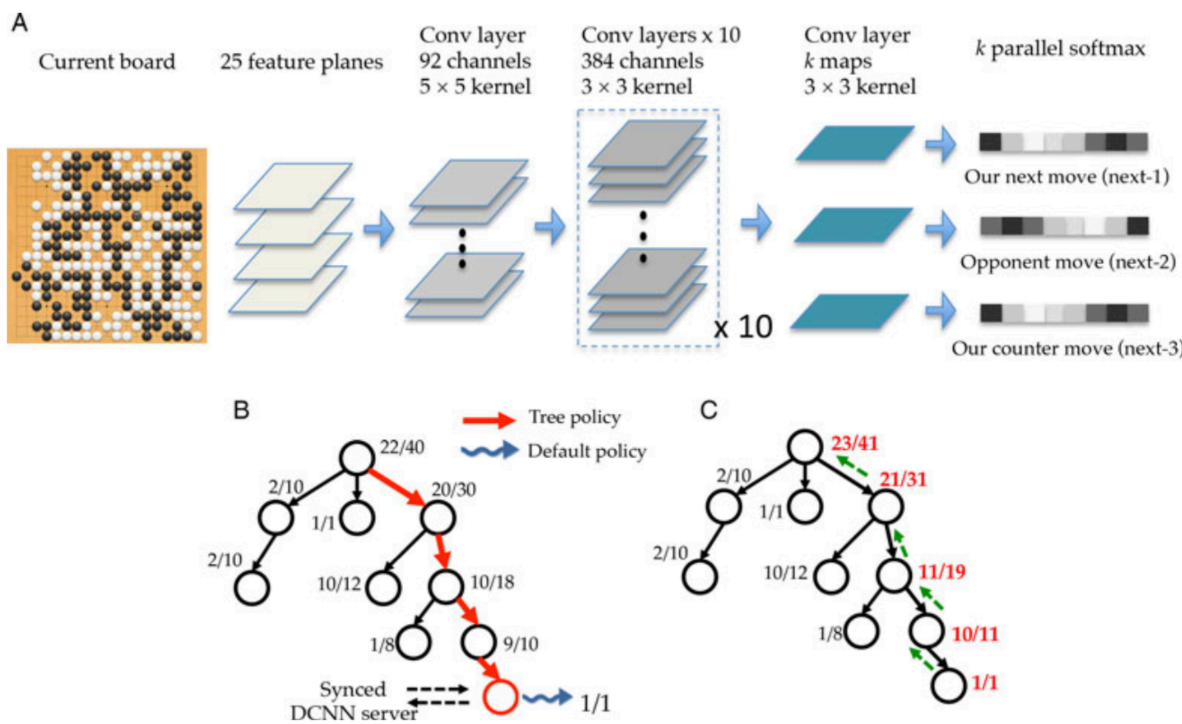

Humans understand these variants and adapt to them because they ***explicitly represent Go as a game***, with a goal to beat an adversary who is playing to achieve the same goal he or she is, governed by rules about how stones can be placed on a board and how board positions are scored. ***Humans represent their strategies as a response to these constraints, such that if the game changes, they can begin to adjust their strategies accordingly***.

- Is it a further constraint solving problem with different constraints?

Go presents compelling challenges for AI beyond matching world-class human performance, in trying to ***match human levels of understanding and generalization, based on the same kinds and amounts of data, explicit instructions, and opportunities for social learning afforded to people***. In learning to play Go as quickly and as flexibly as they do, people are drawing on most of the cognitive ingredients this article has laid out. They are learning-to-learn with compositional knowledge. They are using their core intuitive psychology and aspects of their intuitive physics (spatial and object representations). And like AlphaGo, they are also integrating model-free pattern recognition with model-based search.