Self-Adaptable Learning

Ben Goertzel

March 18, 2011

Abstract

The term "higher level learning" may be used to refer to learning how to learn, learning how to learn how to learn, etc. If an agent is good at ordinary everyday learning, but also at *learning about which learning strategies are most amenable to higher-level learning*, and does both *in a way that is amenable to higher level learning.*– then it may be said to possess **self-adaptable learning**. Goals and environments in which higher-level learning is a good strategy for intelligence, may be called **adaptationally hierarchical** – a property that everyday human environments are postulated to possess. These notions are carefully articulated and formalized; and a concept of **cognitive continuity** is also introduced, which is argued to militate in favor of self-adaptability in a learning system.

1 Introduction

Gregory Bateson, in *Mind and Nature* [Bat80], articulated the appealing idea that human-level intelligence involves multiple levels of learning, which he characterized as

- learning
- learning how to learn (aka second-order learning)
- learning how to learn how to learn (aka third-order learning)
- ...

He also conjectured that human beings very rarely proceed past third-order learning.

A related idea that has achieved some currency in cognitive science and AI lately is the concept of "metacognition" – thinking about thinking [MS94, Shi00]. Metacognition has been proposed as a critical aspect of human-level thinking. Batesonian second-order learning would be a prime example of metacognition.

The concept of "metalearning" in machine learning is closely related – this comprises, for instance, learning which machine learning algorithms or parameters or features are most likely to be appropriate for various specific types of learning problems [BGCSV08]. So one can transfer knowledge about "what works for what sort of problem" from old problems to new ones, thus enabling more rapid learning for new problems, avoiding "reinventing the wheel." Metalearning as commonly conceived is an instance of Batesonian second-order learning; and one could use the term "metametalearning" to refer to Batesonian third-order learning, etc.

Taking a cue from these ideas, it's interesting to think about the prerequisites *enabling* higher order learning.

Starting at the lowest level, some modes of acting are more amenable to learning than others. For instance, if an agent keeps repetitively doing the same things over and over again, that agent's opportunities for learning are going to be rather limited. Whereas if the agent tries a diverse variety of things, the opportunities for learning will be much greater. Trying a variety of things doesn't guarantee learning, but it sets the stage for learning, more so than trying a narrow scope of things.

Similarly, some modes of learning are more amenable to *learning how to learn* than others. This is commonsense among folks involved with using metalearning for supervised categorization or other practical machine learning applications. Suppose you have two different machine learning algorithms A and B, and they both solve individual problems with equal accuracy, but A has simpler dependencies on its parameters than B does. Then, A is more amenable to metalearning. Doing metalearning with B may still be possible, but may require more processing power or memory, or a smarter learning algorithm at the meta level.

You could argue that, in this case, A is simply doing better learning than B, because it is simpler, and according Occam's razor, a simple solution is better. But there are a lot of ways of measuring simplicity, and B might be more favorable than A according to some of them. What I'm pointing out here, though, is a specific way of measuring simplicity of learning algorithms, which correlates well with "amenability to metalearning."

And what about on the higher levels? Similarly, an algorithm with simpler parameter dependencies, used for second-order learning, will be more amenable to third-order learning.

But one can go deeper than this. These observations view the matter from the outside, from the point of view of an algorithm designer, choosing one or another learning algorithm to utilize for first, second or third level learning. But, what if we look at the same issues from the *inside*, from the point of view of an agent that is learning in the world and trying to adapt itself to become smarter and smarter (as part of achieving its basic goals).

In this case, it seems clear that *one* of the agent's goals should be to *learn (on every level) in a way that is amenable to higher-level learning.*

But of course, it's not going to be obvious which learning strategies are most amenable to higher-level learning.

So what needs to be done is to *learn about which learning strategies are most amenable to higher-level learning,* **in a way that is amenable to higher level learning**.

If an agent is good at ordinary everyday learning, but also at *learning about which learning strategies are most amenable to higher-level learning*, and does both **in a way that is amenable to higher level learning**.– then it may be said to possess **self-adaptable learning**.

2 Adaptationally Hierarchical Environments

One may wonder whether it's really worth separating out "self-adaptable learning" as a separate concept from, simply, learning. From a practical standpoint, in terms of studying human intelligence and designing artificial intelligence, it does seem a worthwhile distinction. But it's worth reflecting on why. Bateson's basic observation, in articulating the levels of learning, was that in practical situations, the breakdown of learning into multiple levels appears to be a valuable heuristic. But why is this? It seems to be a property of the specific goals and environments that humans (and our AI systems) operate in. Perhaps this property would be shared by any intelligent agents in the physical universe, but it's less clear that the property must in principle be shared by *any* computationally possible agents in any possible computable environments with any computable goals.

We may define a set of goals (in a set of environments) as *adaptationally hierarchical* if, given a fixed set of computational resources, the most intelligent agents relative to those goals and that environment, are ones that carry out a lot of self-adaptable learning. This is different from simply saying an environment has a compositionally hierarchical structure, with entities built from entities built from entities – although, quite clearly, there are linkages between the two types of hierarchy.

A nontrivial hypothesis is that the environments humans, and reasonably near-term AGI systems, are going to have to deal with, are fairly strongly adaptationally hierarchical. This hypothesis may be considered alongside related hypotheses in [Goe09] regarding the Embodied Communication Prior and the Naturalness of Knowledge Categories. Adaptational hierarchy, ECP and NKC are all examples of *general abstract properties* of goals and environments, which appear to strongly bias the set of agents that display efficient pragmatic general intelligence (see [Goe10] for a formal definition of this) with respect to those goals and environments.

3 Formalizing Self-Adaptable Learning

We now give a mathematical expression to the above informal notion of self-adaptable learning.

Where Φ is a learning strategy (e.g. a learning algorithm with certain parameters) capable solving learning problems in class C, let

- $L(\Phi, C)$ denote the effectiveness of Φ at solving learning problems in C
- $A(\Phi, C)$ denote the "amenability to higher level learning" of the strategy Φ in the context of the problem-class C
- $L(\Phi, C_A)$ and $A(\Phi, C_A)$ denote the learning capability and "amenability to higher level learning" of the strategy Φ in the context of that subclass C_A of the problem-class C that relates to amenability to higher-level learning

Then what is needed is to find a strategy Φ for learning A in the contexts C relevant to the agent – that is both effective at learning, and also amenable to higher-order learning. I.e., we need a strategy Φ, so that

- $L(\Phi, C)$
- $A(\Phi, C)$
- $L(\Phi, C_A)$
- $A(\Phi, C_A)$

are all reasonably large.

Of course, if $L(\Phi, C)$ were high enough, the agent wouldn't need to worry about these other quantities! You don't need meta-learning if your learning algorithm is good enough. But in reality, it seems this isn't always a practical strategy. As noted above, this appears to be a consequence of the particular sorts of learning problems C provided by practical situations, in combination with the realities of real-world resource limitations.

Now, if we understood A, we could find such an Φ by applying our knowledge of A. And if we had such an Φ, we could use it to understand A.... This is either a vicious, or virtuous, cycle for the intelligent agent, depending on how smart the agent is. If the agent figures out *enough* about Φ and A, then it can use this cycle to learn more and more about both – and get smarter and smarter generically. On the other hand, if it's too dim-witted to get an initial clue about Φ and A, then it will flail around, perhaps able to carry out various isolated acts of learning, but unable to

Let $q(w, x, y, z) : R^4 \to R$ be monotone increasing in all its arguments; then we may use

$$Q(\Phi, C) = q(L(\Phi, C), A(\Phi, C), L(\Phi, C_A), A(\Phi, C_A))$$

as a composite measure of quality, combining learning capability and amenability to higher-level learning. A simple example would be a weighted sum. This seems sensibly referred to as the degree of **self-adaptable intelligence** possessed by the learning strategy Φ. A learning strategy with high Q may be said to constitute $\mathbf{self-adaptablelearning}$.

An agent doesn't need to find the absolute maximum of $Q(\Phi, C)$, to achieve a high level of intelligence. What it needs is to find a way to make this quantity reasonably large.

Given a strategy Φ with $q_\Phi = Q(\Phi, C)$, one can identify the amount of resources R_Φ required to apply Φ to find some Φ_1 so that $q_{\Phi_1} \geq kq_\Phi$, where $k > 1$. It's interesting to chart the likely trajectory of R_Φ as an agent develops over time. On average, it would seem that, for fixed k,

1. R_Φ should be quite large when q_Φ is small, as the system is not very smart at self-improvement, and needs to do a lot of work to figure out small tricks for self-improvement
2. R_Φ should decrease as q_Φ begins to increase, as the system gets smarter about self-improvement
3. R_Φ should be fairly low when the system is in a phase of rapid increase in its self-improvement capability
4. R_Φ may increase again once the system gets near the global maximum q_Φ that is possible given its fixed architecture or computational resources ... or if it gets stuck in a local maximum

5. R_Φ is large again when the system is (loosely speaking) as good at self-improvement as it's going to get

Given a maximization problem Ψ, a current best solution X_0, and a learning strategy Φ deployed using computational resources R, we may write $\Phi(\Psi, X_0)$ to denote the probability distribution over possible solutions – i.e. the distribution giving the probability that applying Φ to Ψ, given resources R and current best solution X_0, produces X as the best approximation to the maximizer of Ψ. Then, the solution-distribution resulting from application of Φ to the maximization problem $\Psi(*) = Q(*, C)$ [using the wild-card $*$ to denote the variable being varied to achieve the maximum], given that Φ itself is currently the best known solution to the problem, would be denoted

$$\Phi(Q(*, C), \Phi)$$

If one replaces the single Φ and x_0 with probability distributions, then one still gets a probability distribution for an answer. So, one may talk about an iteration

$$\Phi_n = \Phi_{n-1}(Q(*, C), \Phi_{n-1})$$

potentially ultimately converging toward a fixed point so that

$$\Phi^* = \Phi^*(Q(*, C), \Phi^*)$$

Supposing such a fixed point is found, which constitutes a fairly narrowly-concentrated probability distribution in the space of learning strategies. Then, at this fixed point, one has a strategy where using this strategy as a starting point for thinking about how to find a good strategy for self-improvement, doesn't lead to any improvement. This could be a dead-end arrived at due to stupidity or unfortunate cognitive choices, or it could be an apex of perfection in the art of self-adaptable intelligence.

4 Self-Adaptable Learning and Cognitive Continuity

We have formalized the notion of self-adaptable learning, but in itself, that doesn't tell us much about how to achieve it.

My view is that there is no simple trick providing the sole key to real-world self-adaptable learning. The human brain is complex for good reason – achieving a reasonable level of self-adaptable learning in a complex environment given severely limited computing resources, requires a complex architecture that to some extent mirrors the complexity of the environment and the agent's goals within it.

But nevertheless, some general principles may be discoverable. Here I'll suggest one that seems promising based on prior work in metalearning.

So far one of the most powerful approaches in metalearning for categorization problems, has been the humble "nearest neighbor" approach. Basically, in this approach, when one confronts a new learning problem, one asks what methods have worked for previous similar learning problems, and tries similar methods.

In the present context, suppose that $Q(\Phi, C)$ does not vary too sensitively, on average, when one varies Φ and C. We may call this property of Φ **cognitive continuity**. To formalize this, of course, one needs a metric on the spaces that Φ and C live in. For instance, if Φ is represented as a program in some language, one may achieve this via a metric on program space. If C is a set of goals, these may be represented as mathematical functions (from time-series of observations into real numbers) and may then be metrized by any of the standard metrics on program space.

Once one finds a reasonably good Φ_0 that also has cognitive continuity – it follows that there are likely to be other similar Φ that are also reasonably good. And if the world turns out to be a little different than one thought (i.e. the C one was assuming turns out to be slightly unrealistic), then Φ_0 will probably still be pretty good on the new improved estimate of C. All this makes meta-learning much easier, i.e. it improves amenability to adaptation.

On the other hand, one doesn't want $Q(\Phi, C)$ to be constant either – one wants to be able to vary Φ and get better results.

For instance, consider the case where the behavior of Φ depends on a body of knowledge that the agent has gathered about the world, which is viewable as "biasing" learning in a certain way. Then cognitive continuity means that minor changes to this body of knowledge are only likely to cause modest changes to the learning capability.

With this in mind, I suggest that one way to help oneself along the path toward self-adaptable learning, is to choose learning strategies Φ so that the variation of $Q(\Phi, C)$ near Φ is reasonably but not exorbitantly variant. The quality near Φ should vary, but reasonably continuously.

One could formalize this by formalizing the "reasonableness of variability" of Φ as $V(\Phi, C)$, and then looking at

$$Q(\Phi, C) = q(L(\Phi, C), A(\Phi, C), L(\Phi, C_A), A(\Phi, C_A), V(\Phi, C))$$

(with the weighting function q extended to another argument).

References

[Bat80] Gregory Bateson. *Mind and Nature: A Necessary Unity.* Bantam, 1980.

[BGCSV08] Pavel B. Brazdil, Christophe Giraud-Carrier, Carlos Soares, and Ricardo Vilalta. *Metalearning: Application to Data Mining.* Springer, 2008.

[Goe09] Ben Goertzel. The embodied communication prior. 2009.

[Goe10] Ben Goertzel. Toward a formal definition of real-world general intelligence. 2010.

[MS94] J Metcalfe and A P Shimamura. *Metacognition: knowing about knowing.* MIT Press, 1994.

[Shi00] A P Shinamura. Toward a cognitive neuroscience of metacognition. *Consciousness and Cognition*, 9:313–323, 2000.