Beyond Machine Learning: Autonomous Learning

Frédéric Alexandre

Beyond Machine Learning : Autonomous Learning

Frédéric ALEXANDRE[1][2][3]

[1]*Inria Bordeaux Sud-Ouest, 200 Avenue de la Vieille Tour, 33405 Talence, France*
[2]*LaBRI, Université de Bordeaux, Institut Polytechnique de Bordeaux, CNRS, UMR 5800, Talence, France*
[3]*Institut des Maladies Neurodégénératives, Université de Bordeaux, CNRS, UMR 5293, Bordeaux, France*
Frederic.Alexandre@inria.fr

Keywords: Machine Learning, Computational Neuroscience, Autonomous Learning.

Abstract: Recently, Machine Learning has achieved impressive results, surpassing human performances, but these powerful algorithms are still unable to define their goals by themselves or to adapt when the task changes. In short, they are not autonomous. In this paper, we explain why autonomy is an important criterion for really powerful learning algorithms. We propose a number of characteristics that make humans more autonomous than machines when they learn. Humans have a system of memories where one memory can compensate or train another memory if needed. They are able to detect uncertainties and adapt accordingly. They are able to define their goals by themselves, from internal and external cues and are capable of self-evaluation to adapt their learning behavior. We also suggest that introducing these characteristics in the domain of Machine Learning is a critical challenge for future intelligent systems.

1 INTRODUCTION

Machine Learning has achieved impressive results recently, for a variety of applications ranging from text and natural scene labeling (Farabet et al., 2013) to games as difficult as the Go (Clark and Storkey, 2015). Notably, these results have been mentioned as surpassing human performances but this kind of statements deserves some comments. It is true that the AlphaGo computer program by Google DeepMind was the first to beat a professional player and that some software tools developed by Facebook are able to analyze a thousand of posts per second, clearly beyond human performances. These undoubtly impressive results mainly rely on one major characteristic of their underlying algorithms like Deep Learning (LeCun et al., 2015): their learning is based on a combinatorial analysis of cases extracted from huge databases, specifically dedicated to their narrow domain of expertise. AlphaGo can only play Go. A deep learning system trained to extract faces from pictures can only identify faces.

In contrast, one remarkable characteristic of human learning is that, whereas we are in general not excellent in one specific domain, we are quite good in most of them. In addition, we are able to adapt when a new problem appears. This high level of adaptability can be seen by other signs. With neither explicit la-

bels, nor data preprocessing or segmentation, we are able to pay attention to important information and neglect noise. We define by ourselves our goals and the means to reach them, we self-evaluate our performances and re-exploit previously learned knowledge and strategies in some different context. All these characteristics are notably absent from classical Machine Learning approaches.

In summary, whereas Machine Learning demonstrates an impressive brute force of learning in specific domains, humans are versatile and adaptable and can learn in a changing and uncertain world: We are good at autonomous learning. Comparing both kinds of learning is difficult because they address different characteristics; nevertheless, it can be said that autonomous learning is probably an important characteristic, as one want to embed intelligent modules in robots or in interfaces dedicated to act in the real world. It is consequently important to wonder how Machine Learning could integrate more autonomy, instead of just developing more power as it is mainly proposed in most research programs.

The goal of this position paper is not to compare the performances of two approaches, bio-inspired and purely based on mathematics. Neither is it to give precise recipes to improve existing techniques. Instead, we would like to convince the reader that autonomy is a primary property for learning and to in-

troduce two kinds of information that might be useful to develop more autonomous machine learning. One the one hand, we would like to define more precisely what being autonomous means: even if this property may seem obvious, it is sometimes difficult to concretely explain which characteristics are the bases of an autonomous analysis of our situation and of the best decision to make. One the other hand, we propose to mention the main cerebral structures and circuits involved in these facets of autonomous behavior. This might be useful to decide for future research topics. To do so, we propose to describe here a series of characteristics of the human brain that participates in our capability to learn autonomously, with the idea that introducing these characteristics in artificial neural systems could orient Machine Learning toward Autonomous Machine Learning.

2 AN INTERACTING SYSTEM OF MEMORIES

It is known for a long time (Squire, 1992) that specific circuits in the brain are mobilized to learn explicit knowledge and others to learn procedures. These functions have been respectively addressed by recurrent (Hopfield, 1982) and layered (Rosenblatt, 1958) neural networks, the latter being the ancestor of the deep learning networks. It has also been advocated (Alexandre, 2000) that realistic models of memory should include both kinds of networks to be able to learn by heart specific events as well as generalize some skills from a set of experiences.

Besides modeling these circuits, studying their interactions is also crucial to understand how one system can compensate for or supervise another system, resulting in a more autonomous learning. In particular, interaction between systems can lead to situations where one system can propose an answer (possibly of lower quality) if the other one is too specific or is not trained enough or even if it has been damaged. It can consequently give more time to the other one to become more general or more mature, or to recover. As we will exemplify below, interaction between systems can also more directly give the opportunity to one system to send well selected cases to train the other system. In both cases, the mechanisms are internal and do not require help from the external world, hence an increased autonomy.

For example, in the domain of perceptual learning in the medial temporal lobe, models of the hippocampus can store in episodic memories important events in one trial (Kassab and Alexandre, 2015; O'Reilly and Rudy, 2001). This neural structure is also known

(McClelland et al., 1995) to form later, by consolidation in other circuits, new semantic categories. In the domain of decision-making in the loops between the prefrontal cortex and the basal ganglia, models of cerebral mechanisms are currently developed (Piron et al., 2016), by which goal-directed behavior relying on explicit evaluation of expected rewards can later become habits, automatically triggered with less flexibility but increased effectiveness.

In both cases (either perceptual or motor), the strategy of learning is first to store some specific cases of interest and to recall them if necessary. Then if it appears that similar cases frequently occur, the strategy will be to find some generality and build a generic rule or procedure to deal with such cases. Building such rules from the initially stored cases has several advantages, among which autonomy is not the least. Understanding how both kinds of memory cooperate can lead to an autonomous learning system, able to cope with facts and rules, to elaborate rules from selected facts and to decide which kind of memory is the most adapted to the current situation.

3 COPING WITH UNCERTAINTY

We learn the rules that govern the world and consider it uncertain for two main reasons: it can be predictable up to a certain level (stochastic rules) or non-stationary (changing rules). Whereas standard probabilistic models are rather good at tackling the first kind of uncertainty, non-stationarity in a dynamic world raises more difficult problems (Cohen et al., 2007). Learning in autonomy in the real (and hence uncertain) world consequently implies to be able to autonomously give the best explanation to the fact that a previously valid rule has given an unsatisfactory result: Is it just noise or has the rule changed ? It also implies of course to trigger the corresponding best answer (respectively modifying the level of stochasticity associated to the rule or, more critically, selecting another rule in the set of previously designed rules or elaborating a new rule).

Concerning the first point, we are studying how regions of the medial prefrontal cortex are detecting and evaluating the kind and the level of uncertainty by monitoring recent history of performance at managing correctly incoming events (Carrere and Alexandre, 2016a). Concerning the second point, these prefrontal regions and other cerebral regions sensitive to reward prediction errors are also reported to activate the release of neuromodulators like monoamines, known to play a central role in adaptation to uncertainties (Doya, 2002; Alexandre and Carrere, 2016).

Acetylcholine has been reported to be an important factor in case of stochasticity (Yu and Dayan, 2005) and has been modeled as increasing the signal-to-noise ratio in the sensory cortex (Pauli and O'Reilly, 2008) and promoting learning about the context (Carrere and Alexandre, 2015). We have also recently studied the role of noradrenaline in unstationary environments (Carrere and Alexandre, 2016b) and have proposed a biologically-inspired model, proposing that the balance between exploration and exploitation of sensory cues associated to a rule can be modified by the action of noradrenaline on critical cerebral regions. Similarly, the tonic level of dopamine has been reported as increasing when an unstationary environment is detected and as modifying the balance between exploration and exploitation of motor aspects of the rule (Humphries et al., 2012). Altogether, the brain can be presented as a system able to autonomously detect the level of uncertainty and to autonomously modify the way to exploit and update previously acquired rules or to design new rules, with the help of neuromodulators, seen as a meta-learning system modifying hyperparameters of learning algorithms (Doya, 2002).

4 DEFINING GOALS BY EMBODIMENT AND EMOTIONAL LEARNING

One major difference between artificial and natural learning systems is that the latter ones can autonomously detect and define their goals in the surrounding environment. This is important to choose what to learn and to orient attentional systems accordingly. Instead of learning from a corpus prepared offline, learning is made online and adapted to what has happened during the behavior. In addition, if learning is centred on cues that are important for the agent, like mates, preys or predators, the agent will be probably more efficient than if learning is made from a randomly sampled corpus.

The ability to detect ones goals is due to several ingredients. First our body itself tells us by interoception (Craig, 2003) what is good or bad for us; what must be searched or avoided. It is consequently important for a really autonomous learning agent not only to have a model of the brain with classical cognitive functions related to perception, learning, attention or deliberation but also to feed that model with information coming from a substrate corresponding to the body, including sensors for pain and pleasure.

In the cerebral system, the perceptive system is pre-wired to automatically detect biologically-significant aversive and appetitive (emotional) stimuli and to trigger pavlovian reflexes (Kim et al., 2013). Subsequently, pavlovian learning will allow to anticipate these stimuli by the detection of predictive stimuli that will in turn trigger preparatory behavior (Cardinal et al., 2002). All these stimuli are key targets for attentional processing and correspond to the main goals organizing the behavior. Learning to detect them automatically is consequently important for autonomy, since attentional and learning systems will be fed by an over-representation of these meaningful examples, in contrast to artificial systems that only learn from a stereotyped and artificially prepared corpus. It is consequently important to propose models implementing these pavlovian mechanisms (Krasne et al., 2011; Carrere and Alexandre, 2015) and also the effects of Pavlovian responses onto the body and the neuromodulatory system (Carrere and Alexandre, 2016b).

5 MOTIVATION AND SELF-EVALUATION

Specifying relations between the brain and the body is also an opportunity to introduce physiological needs, fundamental to consider internal goals in addition to external goals evoked above. Indeed, it is important for an agent to learn autonomously, as one of its major constraints is to survive by monitoring some internal variables within vital bounds. One of its primary goals will be to elaborate and select behaviors that help controlling these variables, which can be done in autonomy and not by obeying a supervising system. It is consequently required that the critical internal variables be perceived, carefully processed and participate to decision making: They are undoubtly key cues for an organization of behavior decided in full autonomy.

Such a consideration is also the basis for renewed approaches regarding reinforcement learning. Indeed an important aspect of autonomous decision making is to be able to adapt the behavior as a compromise between external goals (what should be searched and avoided) and internal goals (what are the current needs), whereas classical reinforcement learning generally relies on optimizing a simple scalar representing an abstract reward, artificially given in some situations. To progress in that aim, it is fundamental to better understand how internal and external goals (motivational and emotional cues) are combined in decision making (Zahm, 1999; Mannella et al., 2013; Kolling et al., 2016).

In humans, another important source of information for learning autonomously is based on self-evaluation of the performances. Not only this information can recommand what to learn but, more importantly, it can orient the behavior to explore situations that are best adapted to current expertise and give what is called intrinsic motivations (Oudeyer et al., 2007). It is noticeable that both motivation and self-evaluation processing are central in cognitive control (Koechlin et al., 2003) and reported to be located in the anterior part of the prefrontal cortex, that should be consequently critical targets for future researches, particularly to understand how behavioral rules are elaborated and selected from self-evaluation in the prefrontal cortex (Badre, 2008; Donoso et al., 2014).

6 DISCUSSION

In this position paper, we have first noticed that natural and artificial learning systems differ because, whereas the latter ones are high level specialists in a restricted and artificially sampled domain, the former ones are rather characterized by their intrinsic adaptability to any situation and their capacity to update their knowledge and skills accordingly. Basically, this refers to the function of learning in living systems. Their primary goal is to survive and breed in an unknown environment. As their environment is generally too complex and changing to only consider pre-wired behavioral rules, their knowledge and skills must evolve and adapt to what is perceived internally and externally.

Of course, even in living agents, a part of the adaptation can be dictated during epigenesis by external resources, like genes or social and educational environments, but, most of the time, this adaptation must be done in autonomy and is the main goal of learning processes. In this paper, we have consequently set the focus on autonomy, seen as a primary ingredient of learning and we have explored more precisely several characteristics that, we believe, are the way autonomy can be expressed during learning. In short, we propose that autonomous learning is made easier because our different systems of memory can interact and exchange information, because we are able to estimate the kind of uncertainty we are facing and to propose the suitable adaptation, because we can perceive important noxious and positive events in the environment and inside our body, learn to predict them and build the underlying behavioral rules to optimize in some way their occurrence or avoidance.

We have also mentioned the hypothetized underlying cerebral circuitry for each of these mechanisms and have reported modeling efforts to better understand them. Importantly, we think that these models are important not only for computational neuroscience but also for Machine Learning. Transfering these principles to classical learning algorithms would endow them with more autonomy, which is critical in a context where it is more and more aimed at integrating an Artificial Intelligence in autonomous agents and interfaces.

Up to now, we have only enumerated a list of characteristics, whereas an essential goal for a real autonomy would be to integrate all of them in an agent, corresponding to a physically identified and separated entity, thrown in an unknown environment with the recommendation to survive and no subsequent help.

In this perspective, we have recently designed a software platform (Denoyelle et al., 2014) where an autonomous agent with an artificial body can explore an unknown virtual world. The platform is designed in such a way that the characteristics of the agent's body and of the environment can be easily specified and long lasting experiments can be run to evaluate its survival performances. The main challenge is now to integrate more and more sophisticated versions of the mechanisms evoked above to better understand how they interact and how a viable Autonomous Machine Learning framework can be defined. We are also presently experimenting that, beyond Machine Learning, this numerical testbed is also a precious simulation tool for our medical and neuroscientist colleagues.

REFERENCES

Alexandre, F. (2000). Biological Inspiration for Multiple Memories Implementation and Cooperation. In *In V. Kvasnicka P. Sincak, J. Vascak and R. Mesiar, editors, International Conference on Computational Intelligence.*

Alexandre, F. and Carrere, M. (2016). Modeling neuromodulation as a framework to integrate uncertainty in general cognitive architectures. In Steunebrink, B., Wang, P., and Goertzel, B., editors, *Artificial General Intelligence: 9th International Conference, AGI 2016, New York, NY, USA, July 16-19, 2016, Proceedings*, pages 324–333. Springer International Publishing.

Badre, D. (2008). Cognitive control, hierarchy, and the rostrocaudal organization of the frontal lobes. *Trends in Cognitive Sciences*, 12(5):193–200.

Cardinal, R. N., Parkinson, J. A., Hall, J., and Everitt, B. J. (2002). Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience & Biobehavioral Reviews*, 26(3):321–352.

Carrere, M. and Alexandre, F. (2015). A pavlovian model of the amygdala and its influence within the medial temporal lobe. *Frontiers in Systems Neuroscience*, 9(41).

Carrere, M. and Alexandre, F. (2016a). A system-level model of noradrenergic function. In *The Sixth Joint IEEE International Conference on Developmental Learning and Epigenetic Robotics, September 19th 22nd 2016 Cergy-Pontoise*. IEEE.

Carrere, M. and Alexandre, F. (2016b). A system-level model of noradrenergic function. In Villa, E. A., Masulli, P., and Pons Rivero, J. A., editors, *Artificial Neural Networks and Machine Learning – ICANN 2016: 25th International Conference on Artificial Neural Networks, Barcelona, Spain, September 6-9, 2016, Proceedings, Part I*, pages 214–221. Springer International Publishing.

Clark, C. and Storkey, A. J. (2015). Training Deep Convolutional Neural Networks to Play Go. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 1766–1774.

Cohen, J. D., McClure, S. M., and Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 362(1481):933–942.

Craig, A. D. (2003). Interoception: the sense of the physiological condition of the body. *Current Opinion in Neurobiology*, 13(4):500–505.

Denoyelle, N., Pouget, F., Viéville, T., and Alexandre, F. (2014). VirtualEnaction: A Platform for Systemic Neuroscience Simulation. In *International Congress on Neurotechnology, Electronics and Informatics*, Rome, Italy.

Donoso, M., Collins, A. G., and Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science (New York, N.Y.)*, 344(6191):1481–1486.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, 15(4-6):495–506.

Farabet, C., Couprie, C., Najman, L., and LeCun, Y. (2013). Learning hierarchical features for scene labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1915–1929.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. In *Proceedings of the National Academy of Sciences, USA*, pages 2554–2558.

Humphries, M. D., Khamassi, M., and Gurney, K. (2012). Dopaminergic control of the exploration-exploitation trade-off via the basal ganglia. *Frontiers in Neuroscience*, 6(9).

Kassab, R. and Alexandre, F. (2015). Integration of exteroceptive and interoceptive information within the hippocampus: a computational study. *Frontiers in Systems Neuroscience*, 9(87).

Kim, D., Paré, D., and Nair, S. S. (2013). Mechanisms contributing to the induction and storage of Pavlovian fear memories in the lateral amygdala. *Learning & memory (Cold Spring Harbor, N.Y.)*, 20(8):421–430.

Koechlin, E., Ody, C., and Kouneiher, F. (2003). The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science*, 302(5648):1181–1185.

Kolling, N., Behrens, T. E. J., Wittmann, M. K., and Rushworth, M. F. S. (2016). Multiple signals in anterior cingulate cortex. *Current Opinion in Neurobiology*, 37:36–43.

Krasne, F. B., Fanselow, M. S., and Zelikowsky, M. (2011). Design of a Neurally Plausible Model of Fear Learning. *Frontiers in Behavioral Neuroscience*, 5:41+.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

Mannella, F., Gurney, K., and Baldassarre, G. (2013). The nucleus accumbens as a nexus between values and goals in goal-directed behavior: a review and a new hypothesis. *Frontiers in behavioral neuroscience*, 7.

McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419–457.

O'Reilly, R. C. and Rudy, J. W. (2001). Conjunctive Representations in Learning and Memory: Principles of Cortical and Hippocampal Function. *Psychological Review*, 108(2):311–345.

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265–286.

Pauli, W. M. and O'Reilly, R. C. (2008). Attentional control of associative learning–a possible role of the central cholinergic system. *Brain Research*, 1202:43–53.

Piron, C., Kase, D., Topalidou, M., Goillandeau, M., Rougier, N. P., and Boraud, T. (2016). The globus pallidus pars interna in goal-oriented and routine behaviors: Resolving a long-standing paradox. *Movement Disorders*.

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain.

In Anderson, J. A. and Rosenfeld, E., editors, *Neurocomputing: Foundations of Research (1989)*, pages 89–92. The MIT Press.

Squire, L. R. (1992). Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. *J Cogn Neurosci*, 4(3):232–243.

Yu, A. J. and Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron*, 46(4).

Zahm, D. S. (1999). Functional-anatomical implications of the nucleus accumbens core and shell subterritories. *Annals of the New York Academy of Sciences*, 877:113–128.