Theoretical Computer Science 633 (2016) 83–99

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

The principle of least cognitive action

Alessandro Betti [a], Marco Gori [b,*]

[a] *Department of Physics, University of Pisa, Italy*
[b] *Department of Information Engineering and Mathematics, University of Siena, Italy*

A R T I C L E   I N F O

Article history:
Received 3 January 2015
Received in revised form 17 June 2015
Accepted 19 June 2015
Available online 2 July 2015

Keywords:
BG-brackets
Lifelong learning
Natural learning theory
On-line learning
Least cognitive action

A B S T R A C T

By and large, the interpretation of learning as a computational process taking place in both humans and machines is primarily provided in the framework of statistics. In this paper, we propose a radically different perspective in which the emergence of learning is regarded as the outcome of laws of nature that govern the interactions of intelligent agents with their own environment. We introduce a *natural learning theory* based on the principle of *least cognitive action*, which is inspired to the related mechanical principle, and to the Hamiltonian framework for modeling the motion of particles. The introduction of the kinetic and of the potential energy leads to a surprisingly natural interpretation of learning as a dissipative process. The kinetic energy reflects the temporal variation of the synaptic connections, while the potential energy is a penalty that describes the degree of satisfaction of the environmental constraints. The theory gives a picture of learning in terms of the energy balancing mechanisms, where the novel notions of boundary and bartering energies are introduced. Finally, as an example of application of the theory, we show that the supervised machine learning scheme can be framed in the proposed theory and, in particular, we show that the Euler–Lagrange differential equations of learning collapse to the classic gradient algorithm on the supervised pairs.

1. Introduction

Machine Learning is at a lively crossroad of disciplines, where the exploration of neuroscience and cognitive psychology meets computational models mostly based on statistical foundations. For these models to be realistic, one is typically concerned with the acquisition of learning skills on a sample of training data, that is sufficiently large to be consistent with a statistically relevant test set. However, in most challenging and interesting learning tasks taking place in humans, the underlying computational processes do not seem to offer such a neat identification of the training set. As time goes by, humans react surprisingly well to new stimuli, while keeping past acquired skills, which seems to be hard to reach with nowadays intelligent agents. This suggests us to look for alternative foundations of learning, which are not necessarily based on statistical models of the whole agent life.

In this paper, we investigate the emergence of learning as the outcome of laws of nature that govern the interactions of intelligent agents with their own environment, regardless of their nature. The underlying principle is that the acquisition of cognitive skills by learning obeys information-based laws on these interactions, which hold regardless of biology. In this new perspective, in particular, we introduce a *natural learning theory* aimed at discovering the fundamental temporally-embedded

* Corresponding author.
E-mail address: marco@dii.unisi.it (M. Gori).

http://dx.doi.org/10.1016/j.tcs.2015.06.042

mechanisms of environmental interaction. While the role of time has been quite relevant in machine learning (see e.g. the Hidden Markov Models), in some challenging problems, like computer vision, the temporal dimension seems to play quite a minor role. What could human vision had been in a world of visual information with shuffled frames? Any cognitive process aimed at extracting symbolic information from images that are not frames of a temporally coherent visual stream would have been extremely harder than in our visual experience [1]. More than looking for smart algorithmic solutions to deal with temporal coherence, in this paper we rely on the idea of regarding learning as a process which is deeply interwound with time, just like in most laws of nature.

In order to derive the laws of learning, we introduce the *principle of least cognitive action*, which is inspired to the related mechanical principle, and to the Hamiltonian framework for modeling the motion of particles. Unlike mechanics, however, the cognitive action that we define is in fact the objective to be minimized, more than a functional for which to discover a stationary point. In our learning framework, this duality is based on a proper introduction of the "kinetic" and of the "potential energy," that leads to a surprisingly natural interpretation of learning as a dissipative process. The kinetic energy reflects the temporal variation of the synaptic connections, while the potential energy is a penalty that describes the degree of satisfaction of the environmental constraints. The theory gives a picture of learning in terms of the energy balancing mechanisms, where the novel notions of *boundary and bartering energies* are introduced. These new energies arise mostly to model complex environmental interactions of the agent, that are nicely captured by means of time-variant high-order differential operators. When pairing them with the novel notion of BG-bracket, we show how intelligent processes can be understood in terms of energy balance; learning turns out be a dissipative process which reduces the potential energy by moving the connection weights, thus producing a kinetic energy. However, the energy balancing does involve also the boundary and bartering energies. The former arises because of the energy exchange at the beginning and at the end of the temporal horizon, while the last one is deeply connected with the energy exchange during the agent's life. It is worth mentioning that this energy balance, which has a nice qualitative interpretation, is in fact the formal outcome of the main theoretical results given in the paper. In particular, the proposed framework incorporates classic mechanics as a special case, while it opens the doors to in-depth analyses in any context in which there is an explicit temporal dependence in the Lagrangian of the system.

While the paper focuses on laws of learning that are independent of the nature of the intelligent agent, it is shown that these laws also offer a general framework to derive classic on-line machine learning algorithms. In particular, we show that the supervised machine learning scheme can be framed in the proposed theory, and that the Euler–Lagrange differential equations of learning derived in the paper do collapse to the classic gradient algorithm. Interestingly, the theory prescribes the time-variant learning rate, which arises from the given variational formulation.

The paper is organized as follows. In the next section, we introduce natural principles of learning and, in particular, the notion of cognitive action. Its minimization, which leads to the laws of learning, is described in Section 3. In Section 4 we propose a dissipative Hamiltonian framework of learning, which, in Section 5, is completed by the analysis of the BG-brackets. Section 6 relies on the previous results to provide an interpretation of learning dynamics by means of energy balance. In Section 7, we show how the laws of learning can be converted to the gradient algorithm. Finally, some conclusions are drawn in Section 8.

2. Natural principles of learning

The notion of time is ubiquitous in laws of nature. Surprisingly enough, most studies on machine learning have relegated time to the related notion of iteration step.[1] From one side the connection is sound, since it involves the computational effort, that is also observed in nature. From the other side, while time has a truly continuous nature, the time discretization typical in fields like computer vision seems to give up to the challenge of constructing a truly theory of learning. This paper presents an approach to learning that incorporates time in its truly continuous structure, so as the evolution of the weights of the neural synapses follows equations that resemble laws of physics. We consider environmental interactions that are modeled under the recently proposed framework of *learning from constraints* [2]. In particular, we focus attention on the case in which each constraint originates a loss function, that is expected to take on small values in case of soft-satisfaction. Interestingly, thanks to the adoption of the t-norm theory [3], loss functions can be constructed that can also represent the satisfaction of logic constraints.

A lot of emphasis in machine learning has been on *supervised learning* where we are given the collection $\mathcal{L} = \{(t_\kappa, u(t_\kappa)), s_\kappa\}_{\kappa \in \mathbb{N}}$ of supervised pairs $(u(t_\kappa), s_\kappa)$ over the temporal sequence $\{t_\kappa\}_{\kappa\in\mathbb{N}}$. At a given t, the pairs $(u(t_\kappa), s_k)$ that have become available are denoted by $\triangleright\mathcal{L}_t := \{\kappa \in \mathbb{N} : \ t_\kappa \le t\}$. We assume that those data are learned by a *feedforward neural network* defined by the function

$$f(\cdot,\cdot) : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^D,$$

so as the input $u(t)$ is mapped to $y(t) = f(w(t), u(t))$. Classic supervised learning can be naturally modeled by introducing a loss function $L(\cdot,\cdot) : \mathbb{R}^2 \to \mathbb{R}$ along with the given training set $\mathcal{L}$. Examples of different loss functions can be found

[1] This is often named "epoch."