

Can Artificial Intelligence Improve Our Health?



Radu Grosu

Abstract Open and adaptive living systems share many of their underlying principles with artificial rational agents. Control theory and artificial intelligence drew and continue drawing, therefore, inspiration from nature. For example, we have shown that a biophysical model of neurons can reduce the size of a lane-keeping neural network by a factor of 100 in autonomous driving. In this paper, I review the main principles of control theory (e.g., how a rational agent uses the states and rewards returned by a plant to learn the actions maximizing its total expected reward), and of artificial intelligence (e.g., feedforward versus recurrent networks, unsupervised, supervised, and reinforcement learning) in an accessible and comprehensive manner. Smart, distributed control and machine learning technologies have currently entered all segments of our civilization (e.g., farming, energy, mobility, city management, and industry) in order to optimize and individualize processes, increase efficiency and security, and reduce the environmental burden. Examples from health care include the detection of novel disease features in medical imaging, prediction of ablation strategies in cardiology, better growth models for improved tumor control, and optimized rescue of victims of traffic accidents.

Keywords Artificial intelligence · Systems theory · Rational agent · Dynamic programming · Health care

1 Living Systems are Open Learning Systems

Health is an indispensable attribute of all living organisms and a healthy biosphere is essential for the survival of mankind. Scientific and technological advances over the past centuries have considerably improved our health. Unfortunately, it achieved this at a great pollution cost that puts at risk the rest of the biosphere [2]. This will inevitably lead to our doom, too. A key question is therefore, if science and

R. Grosu (✉)
TU Wien, Vienna, Austria
e-mail: radu.grosu@tuwien.ac.at

technology can now also come to our rescue. Can artificial intelligence help in this respect?

Health is intimately related to life. To understand health, we therefore need to first understand life itself. A descriptive definition of life, from a biological point of view is as follows [11].

Life is a characteristic of open systems that preserves and reinforces their existence in a given environment. This characteristic exhibits all or most of the following seven traits: Homeostasis, organization, metabolism, growth, adaptation, response to stimuli, and reproduction.

A disease occurs when any of these traits is damaged, and the purpose of health-care is to repair, or at least to control the damage. Repair and control can thus only be conceived after a deeper look into the nature of the seven traits above, themselves. They all involve optimal regulation.

Homeostasis is a dynamic internal equilibrium maintained by living organisms to ensure their optimal functioning. Organization is a structuring principle allowing the emergence of complex organisms from cells, the basic units of life. Metabolism is a set of life sustaining chemical reactions optimally converting food to energy or cellular building blocks, and eliminating waste. Growth and reproduction are responsible for cell increase and division, and the ability to produce new organisms. Adaptation is the ability of an organism to change over time in response to a change in its environment. Finally, response to stimuli is the ability of an organism to optimally interact with its environment.

2 Information Processing is at the Core of Living and Artificial Systems

Systems biology [16] studies the emergence of complexity in functional organisms from the point of view of dynamic systems theory. Control theory [3] and artificial intelligence [14] draw inspiration from living organisms, too, with an emphasis on information processing. The latter is at the heart of all the above traits, through the very concept of optimal regulation.

Open systems use information processing in order to learn a dynamic model of their environment, including their own abilities, with the purpose of acting in an optimal fashion to a possibly unbounded sequence of environmental stimuli and rewards, each associated to their previous actions.

Acting in an optimal fashion refers to choosing the actions that maximize the total expected reward over time, also known as the utility of the system. In biology, rewards play a central role [7], and they range from most primitive ones, such as pain, hunger, and pleasure, facilitating survival and reproduction, to more abstract ones, such as promotion, or simply winning a game in sports, facilitating selection. Reward shaping [15] is also a main concern in artificial intelligence and decision theory [1]. The open system, called a rational agent [14], is assumed to possess finitary means,

only, and key to such means are learning and memory. This is where most of recent advances in artificial intelligence, in particular in machine learning, happened.

The architecture of a rational agent, together with its environment, is shown in Fig. 1. The physical plant represents the environment to be controlled, for example, one of the seven life traits in biology, or a factory, energy system, or even a car in engineering. The plant has a set of sensors, actuators, and rewards. The actuators control the behavior of the plant, whereas the sensors and rewards provide evidence, or feedback, about the state of the plant and also about the desirability of this state. The main goal of the rational agent is to maximize the total reward $r = \sum_{i=1}^{\infty} \gamma^i r_i$ received from the plant over time, in expectation, where $\gamma \in (0, 1)$ is a discount weighting the importance of future rewards.

As the evidence output by the plant might contain incomplete information about the internal plant state, a state estimator is used first, to compute the most likely state of the plant $P(b|e, a)$, also called a belief state, from the past sequence of controller actions $a = a_1 \dots a_n$, and the past sequence of plant evidence $e = e_1 \dots e_n$. For this purpose, the estimator takes advantage of a stochastic model of the plant, consisting of two distributions: the next state and reward probability, given the current state and current action $P(s', r|s, a)$, and the evidence probability, given the current state $P(e|s)$.

The current belief state and the current reward is given to an optimal controller, which computes the current action $a = \pi(b)$, such that the utility $U(b)$ of the agent, that is the rewards it receives over time starting from b , are maximized in expectation. The agent may possess a map, or desired-goals landscape, allowing a planner to compute an optimal path p , that is a sequence of waypoints, or sub goals, which are also given to the optimal controller.

A very important question is how does the rational agent infer the plant model, the current belief state, the optimal action, and the waypoints? In biology, this is achieved through very sophisticated regulatory networks [7], which are either chemical and therefore slow, such as gene regulatory networks, homeostatic regulatory networks, metabolic regulatory networks, and reproduction regulatory networks, or electric

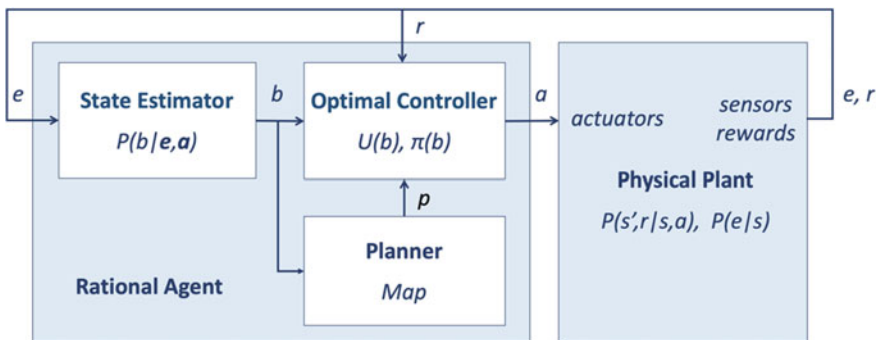


Fig. 1 The architecture of a rational agent together with its environment

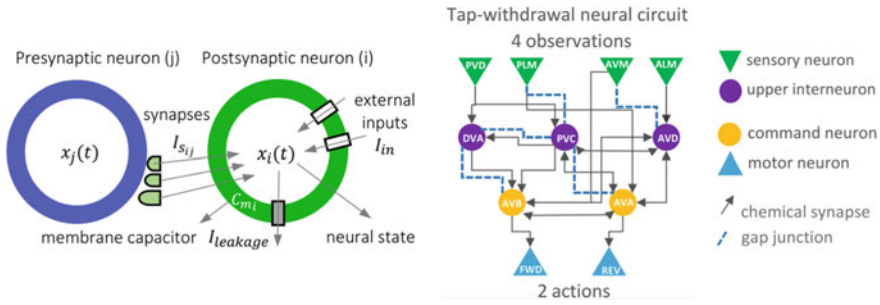


Fig. 2 Simplified architecture of biological neural networks

and therefore fast, such as neural regulatory networks. All these networks possess plasticity, that is, they adapt to environmental changes. In neural regulatory networks [7], the basic information-processing unit is the neuron and its synapses, with which the neuron interacts with the other neurons. A simplified architecture of biological neurons and their associated neural networks is shown in Fig. 2.

Figure 2a shows a biophysical model of biological neurons [7, 9, 17], which focusses solely on the electric interaction between pre- and post-synaptic neurons. Electrically, the fatty membrane of a neuron is a capacitor, and the inside–outside difference of ionic concentrations defines the membrane potentials $x_j(t)$, $x_i(t)$ of the neurons at time t , respectively. The neurons interact with each other through either electric synapses, also called gap junctions, or through chemical synapses, which are either activating or inhibiting. Through each synapse s passes a current $I_{sij}(t)$. Through the membrane also passes a passive leakage current $I_{leakage}(t)$ and possibly an input current $I_{in}(t)$.

Figure 2b shows a simple neural circuit [17] of the *C. elegans* nematode, the so-called tap withdrawal circuit. The circuit has four sensory neurons, responsive to forward, backward, and middle taps, of the Petri dish containing the nematode; two motor neurons, causing either a forward or a backward movement; three interneurons further processing the inputs; and two command neurons. The neurons are recurrently interconnected through either gap junctions or chemical synapses. Originally, the circuit causes the nematode to make a large forward or backward movement, of the size of the nematode, in case the dish is tapped on the back-side or the front side of the nematode, respectively. In case the dish is tapped close to the middle of the nematode, the nematode decides to move backwards. If, however, repeated taps to the dish do not harm the nematode in any way, the nematode learns that there is no harm, and considerably reduces its movement response.

At TU Wien [10], we developed a biophysical model of neurons and synapses, capturing the above electric description with smooth ordinary differential equations. This model is more complex than artificial neural network models, described below, but it is much more succinct, as each neuron can accomplish more complex tasks.

For example, to get a taste of this succinctness, we were able to learn a lane-keeping controller for an autonomous driving car with only 19 neurons, whereas

the corresponding artificial deep neural network learned by our MIT collaborators, required 1900 neurons for the same precision [10]. Both networks were able to drive successfully a car equipped with appropriate sensing and actuation capabilities on the streets of Boston. Owing to its small size, the biophysical network is also explainable [10] and thus certifiable. However, this is not the case for the deep neural network.

3 Principles of Machine Learning

The key property of biological networks, that enables them to learn while interacting with their environment, is the plasticity of synapses, that is, the ability of synapses to change their strength in response to the accumulated rewards, for example, in the form of serotonin [8].

Artificial neural networks [5] take this observation to the limit, by using a model that very remotely resembles biological neurons, only. Each synapse s has an associated weight w_s , which multiplies the value presented by the presynaptic neuron. The neuron itself makes a smoothed, step-like decision, called a sigmoidal output, based on the sum $\sum_{i=1}^n w_i x_i - \mu$ of its weighted inputs and its associated bias μ . A network that has at least two hidden layers and sufficient neurons, can approximate to arbitrary precision any nonlinear function. The more hidden layers are available, the easier it is to learn the function, as each layer acts like one parallel step of a sequential computation. Thus, layers decompose the original function in the sequential composition of simpler functions. If the network has more than two layers, it is called a deep neural network [5]. According to the graph structure of the network, one distinguishes between feedforward and recurrent networks. Figure 3a shows a feed forward network [5], that is, a network with no feedback connections.

This network has one input layer x , one output layer y , here with a single neuron, and l hidden layers of neurons. Like biophysical networks, these networks can be trained to accomplish a particular task, such as lane-keeping in autonomous driving. Their drawback is that they tend to be very large, in both terms of depth and width, and the individual role of each neuron is very hard to explain. If a network allows feedback connections, like in biophysical networks, then it is called recurrent [5].

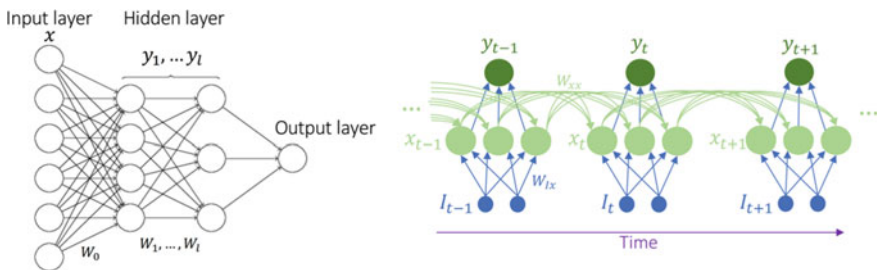


Fig. 3 A feedforward and a recurrent artificial deep neural network

In such cases, the behavior of a network can be explained by its unfolding in time, in a feedforward network, as shown in Fig. 3b. In this unfolding, each slice t computes the output y_t based on the current state x_t . The current state is computed based on the current input I_t and the previous state x_{t-1} . Recurrent networks are especially useful in sequential decision tasks, where the current output, or action, depends on the entire sequence of previous evidence and previous actions. Hence, recurrent neural networks are very well suited for learning the components of a rational agent.

The *learning techniques* are best explained in terms of Fig. 1, giving the architecture of a rational agent. They *can be classified into three categories*: Unsupervised learning, supervised learning, and reinforcement learning. The main goal of unsupervised learning [5] is to uncover the fundamental patterns within the environmental, that is, the plant's evidence.

The basic neural architecture employed for this purpose is a fan-in fan-out feedforward network, called an autoencoder. The fan-in encodes (compresses) the evidence through successively smaller layers of hidden neurons. The fan-out feedforward network decodes (decompresses) the state obtained this way, through successively larger layers of neurons, such that the output state, reproduces the input up to a desired precision. Each neuron of the middle layer, containing the smallest number of neurons, can be understood in this case as a fundamental feature, or pattern of the evidence, as the nonlinear combination of these features is able to accurately reproduce the original evidence. More advanced techniques include variational autoencoders [13] and generative adversarial networks [6]. An example for the use of unsupervised learning is uncovering the main patterns occurring in prostate cancer, for different levels of cancer malignity, as defined by the well-established Gleason score [4].

The main goal of supervised learning [5, 14] is to learn an input–output model, given a set of desired input–output sequences.

For example, given action sequence $\mathbf{a} = a_1 \dots a_n$ and evidence sequence $\mathbf{e} = e_1 \dots e_n$, the goal is to learn a plant model that produces for same actions an evidence $\hat{\mathbf{e}} = \hat{e}_1 \dots \hat{e}_n$, that minimally differs from the plant evidence \mathbf{e} , according to a given cost function. A very popular cost is $\sum_{i=1}^n (e_i - \hat{e}_i)^2 / 2$. Minimization is achieved by successively updating the synaptic weights of the neural network through backpropagation and stochastic gradient descent, until a local minimum, sufficiently close to the global one, is achieved. Supervised learning is not restricted to plant models. One can also apply it to the optimal controller, if one already has a teacher. This can be a human or a program, producing the proper actions given the current states. An example of supervised learning in medicine is prostate cancer segmentation, where the prostate is delineated by an expert physician within the input image [4]. A further example is classification, where the malignity of the prostate carcinoma is labeled by a physician according to the Gleason score [4]. In this case, one obtains superior results, if supervised learning is combined with unsupervised learning. The latter abstracts the pixels in the original image into very useful lower dimensional features, which considerably simplifies the ulterior classification of the carcinoma's malignity.

Finally, the main aim of reinforcement learning [14, 15] is to learn an input–output model, for example, the optimal controller of a rational agent, solely based on rewards.

While supervised learning cannot learn better models than their teacher can, reinforcement learning can, by discovering new actions on their own. Rewards are the most abstract way of teaching, and they may only occur after many actions. Synaptic weights are learned by maximizing the cumulative expected reward. If a finitary plant model is available and the plant is fully observable, then one may employ dynamic programming in order to solve the Bellman equation [1], defining in a recursive fashion the utility of a rational agent: $U(s) = R(s) + \gamma \max_a \sum_{s'} P(s'|s, a)U(s')$. Note that knowing the model is equivalent to knowing the recursion-unfolding tree of this equation, for every possible action and every possible next state. If the model is not known but it is finitary, one can employ adaptive dynamic programming, to learn the model on-the-fly. If the model is not finitary, one can employ Monte-Carlo execution sampling, where a sample corresponds to one path in the tree, or temporal-differencing, which amounts to executing just one step. In this case, the utility of a state is $U(s) \cong U(s) + \alpha(R(s) + \gamma U(s') - U(s))$. The parenthesis contains the prediction error, that is the difference in the utility after and before action a in state s leading to next state s' and reward $R(s)$. The learning rate is tuned with α . In neural networks, the main learning tool is the policy-gradient theorem stating that: one has to update the synaptic weights, proportionally to the sample utility, and to the gradient of the probability of choosing the action actually taken. Learning any kind of sport involves reinforcement learning, and so does the optimal irradiation in prostate carcinoma [4].

4 Applications of Artificial Intelligence

Combining the classic, hypothesis-driven approaches to health care, with the novel, data-driven machine learning techniques, is already unleashing a revolution in health care: *Imaging* mines vast imaging repositories for overlooked disease traits that are thereafter used in better disease classification and control. *Cardiology* is on the brink of creating personalized dynamic heart models that allow to explore the results of various ablation strategies in atrial fibrillation. *Oncology* exploits data-driven techniques to learn better tumor-growth models and develop novel strategies for improved tumor control. *Cities* are learning dynamic models, facilitating optimal first response in case of accidents. This requires traffic, patient, physician, and hospital distributions. However, artificial intelligence technology in health-care is still in its infancy. A better understanding of how to apply it in order to repair or at least control diseases will open—without any doubt—treatment possibilities that we could have never dreamt about today.

Artificial intelligence can also indirectly help to improve our health, by more efficiently regulating our industrial society [12]. This will considerably reduce

pollution or human fatalities. For example, artificial intelligence is expected to play a key role in.

Smart mobility with the grand challenge of zero traffic fatalities. Autonomous cars are expected to reduce pollution and accidents through optimal control.

Smart energy with the grand challenge of black-out free electricity. The energy grids will be operated by smart and adaptive controllers.

Smart buildings with the grand challenge of energy-awareness. Deploying a swarm of sensors and actuators will allow for much better monitoring and control.

Industry 4.0 with the grand challenge of on-the-fly production. Digital twins and the industrial Internet of things will revolutionize the way factories work.

Smart farming with the grand challenge of max-yield agriculture. Sophisticated weather prediction and the agricultural Internet will be a game-changer.

In all the above areas, the Internet of things, akin to our body, is going to play a central role. Its swarm of sensors and actuators, like our skin and muscle cells, the fog, like our spinal cord, and the cloud, like our brain, will all employ machine learning technology, in order to enhance their own abilities.

References

1. Bellman RE (1957) Dynamic programming. Princeton University Press
2. Carrington D (2017) Global pollution kills 9m a year and threatens survival of human societies. The Guardian
3. Franklin GF, Powell JD (1993) Feedback control of dynamic systems. Addison-Wesley Series in Electrical and Computer Engineering. Control Engineering
4. Goldenberg SL, Nir G, Salcudean SE (2019) A new era: artificial intelligence and machine learning in prostate cancer. Nat Rev Urol 16:391–403
5. Goodfellow I, Bengio Y, Courville A (2016) Deep learning. MIT Press
6. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Advances in neural information processing systems. In: Ghahramani Z, Welling M, Cortes C, Lawrence N, Weinberger KQ (eds) Generative adversarial nets, vol 27, pp 1–9, Curran Associates, Inc
7. Kandel ER, Schwartz JH, Jessel TM, Siegelbaum SA, Hudspeth AJ (2013) Principles of neural science. 5 edn, McGraw-Hill Education/McGraw-Hill Medical
8. Kandel ER, Schwartz JH, Jessell TM, Siegelbaum SA, Hudspeth AJ (2012) Principles of neural science. McGraw Hill Medical
9. Lechner M, Hasani RM, Zimmer M, Henzinger TA, Grosu R (2019) Designing worm-inspired neural networks for interpretable robotic control. In: International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, IEEE, May 20–24, 2019, pp 87–94
10. Lechner M, Hasani R, Amini A, Henzinger T, Rus D, Grosu R (2020) Neural circuit policies enabling auditable autonomy. Nat Mach Intel 2:642–652
11. McKay CP (2004) What is life and how do we search for it in other worlds? PLoS Biol 2(9):302
12. Microsoft IoT Signals Report. Summary of research learnings (2019). <https://azure.microsoft.com/mediahandler/files/resourcefiles/iot-signals/IoT-Signals-Microsoft-072019.pdf>
13. Pu Y, Gan Z, Heno R, Yuan X, Li C, Stevens A, Carin L (2016) Variational autoencoder for deep learning of images, labels, and captions. Adv Neural Inform Process Syst 2352–2360
14. Russell S, Norvig P (2016) Artificial intelligence: a modern approach, 3rd edn. Pearson Education
15. Sutton RS, Barto AG (2018) Reinforcement learning: an introduction. 2 edn, The MIT Press

16. Tavassoly I, Goldfarb J, Iyengar R (2018) Systems biology primer: the basic methods and approaches. *Essays Biochem* 62(4):487–500
17. Wicks SR, Roehrig CJ, Rankin CH (1996) A dynamic network simulation of the nematode tap withdrawal circuit: predictions concerning synaptic function using behavioral criteria. *J Neurosci* 16(12):4017–4031