Review

Born to learn: The inspiration, progress, and future of evolved plastic artificial neural networks

Andrea Soltoggio [a,*], Kenneth O. Stanley [b], Sebastian Risi [c]

[a] Department of Computer Science, Loughborough University, LE11 3TU, Loughborough, UK
[b] Department of Computer Science, University of Central Florida, Orlando, FL, USA
[c] IT University of Copenhagen, Copenhagen, Denmark

ARTICLE INFO

ABSTRACT

Biological neural networks are systems of extraordinary computational capabilities shaped by evolution, development, and lifelong learning. The interplay of these elements leads to the emergence of biological intelligence. Inspired by such intricate natural phenomena, Evolved Plastic Artificial Neural Networks (EPANNs) employ simulated evolution in-silico to breed plastic neural networks with the aim to autonomously design and create learning systems. EPANN experiments evolve networks that include both innate properties and the ability to change and learn in response to experiences in different environments and problem domains. EPANNs' aims include autonomously creating learning systems, bootstrapping learning from scratch, recovering performance in unseen conditions, testing the computational advantages of particular neural components, and deriving hypotheses on the emergence of biological learning. Thus, EPANNs may include a large variety of different neuron types and dynamics, network architectures, plasticity rules, and other factors. While EPANNs have seen considerable progress over the last two decades, current scientific and technological advances in artificial neural networks are setting the conditions for radically new approaches and results. Exploiting the increased availability of computational resources and of simulation environments, the often challenging task of hand-designing learning neural networks could be replaced by more autonomous and creative processes. This paper brings together a variety of inspiring ideas that define the field of EPANNs. The main methods and results are reviewed. Finally, new opportunities and possible developments are presented.

© 2018 Elsevier Ltd. All rights reserved.

Contents

* Corresponding author.
  E-mail addresses: a.soltoggio@lboro.ac.uk (A. Soltoggio), kstanley@cs.ucf.edu (K.O. Stanley), sebr@itu.dk (S. Risi).

1. Introduction

Over the course of millions of years, evolution has led to the emergence of innumerable biological systems, and intelligence itself, crowned by the evolution of the human brain. Evolution, development, and learning are the fundamental processes that underpin biological intelligence. Thus, it is no surprise that scientists have tried to engineer artificial systems to reproduce such phenomena (Dawkins, 2003; Sanchez et al., 1996; Sipper et al., 1997). The fields of artificial intelligence (AI) and artificial life (AL) (Langton, 1997) are inspired by nature and biology in their attempt to create intelligence and forms of life from human-designed computation: the main idea is to abstract the principles from the medium, i.e., biology, and utilize such principles to devise algorithms and devices that reproduce properties of their biological counterparts.

One possible way to design complex and intelligent systems, compatible with our natural and evolutionary history, is to simulate natural evolution in-silico, as in the field of evolutionary computation (Eiben & Smith, 2015; Holland, 1975). Sub-fields of evolutionary computation such as evolutionary robotics (Harvey, Husbands, Cliff, Thompson, & Jakobi, 1997; Nolfi & Floreano, 2000), learning classifier systems (Butz, 2015; Lanzi, Stolzmann, & Wilson, 2003), and neuroevolution (Yao, 1999) specifically research algorithms that, by exploiting artificial evolution of physical, computational, and neural models, seek to discover principles behind intelligent and learning systems.

In the past, research in evolutionary computation, particularly in the area of neuroevolution, was predominantly focused on the evolution of static systems or networks with fixed neural weights: evolution was seen as an alternative to learning rules to search for optimal weights in an artificial neural network (ANN). Also, in traditional and deep ANNs, learning is often performed during an initial training phase, so that weights are static when the network is deployed. Recently, however, inspiration has originated more strongly from the fact that intelligence in biological organisms considerably relies on powerful and general learning algorithms, designed by evolution, that are executed during both development and continuously throughout life.

As a consequence, the field of neuroevolution is now progressively moving towards the design and evolution of lifelong learning plastic neural systems, capable of discovering learning principles during evolution, and thereby able to acquire knowledge and skills through the interaction with the environment (Coleman & Blair, 2012). This paper reviews and organizes the field that studies *evolved plastic artificial neural networks*, and introduces the acronym EPANN. EPANNs are *evolved* because parts of their design are determined by an evolutionary algorithm; they are *plastic* because parts of their structures or functions, e.g. the connectivity among neurons, change at various time scales while experiencing sensory-motor information streams. The final capabilities of such networks are autonomously determined by the combination of evolved genetic instructions and learning that takes place as the network interacts with an environment.

EPANNs' ambitious motivations and aims, centered on the autonomous discovery and design of learning systems, also entail a number of research problems. One problem is how to set up evolutionary experiments that can discover learning, and then to understand the subsequent interaction of dynamics across the evolutionary and learning dimensions. A second open question concerns the appropriate neural model abstractions that may capture essential computational principles to enable learning and, more generally, intelligence. One further problem is the size of very large search spaces, and the high computational cost required to simulate even simple models of lifelong learning and evolution. Finally, experiments to autonomously discover intelligent learning systems have a wide range of performance metrics, as their objectives are sometimes loosely defined as the increase of behavioral complexity, intelligence, adaptability, evolvability (Miconi, 2008), and general learning capabilities (Tonelli & Mouret, 2011). Thus, EPANNs explore a larger search space, and address broader research questions, than machine learning algorithms specifically designed to improve performance on well-defined and narrow problems.

The power of EPANNs, however, derives from two autonomous search processes: evolution and learning, which arguably place them among the most advanced AI and machine learning systems in terms of open-endedness, autonomy, potential for discovery, creativity, and human-free design. These systems rely the least on pre-programmed instructions because they are designed to autonomously evolve while interacting with a real or simulated world. Plastic networks, in particular recurrent plastic networks, are known for their computational power (Cabessa & Siegelmann, 2014): evolution can be a valuable tool to explore the power of those computational structures.

In recent years, progress in a number of relevant areas has set the stage for renewed advancements of EPANNs: ANNs, in particular deep networks, are becoming increasingly more successful and popular; there has been a remarkable increase in available computational power by means of parallel GPU computing and

dedicated hardware; a better understanding of search, complexity, and evolutionary computation allows for less naive approaches; and finally, neuroscience and genetics provide us with an increasingly large set of inspirational principles. This progress has changed the theoretical and technological landscape in which EPANNs first emerged, providing greater research opportunities than in the past.

Despite a considerable body of work, research in EPANNs has never been unified through a single description of its motivations and inspiration, achievements and ambitions. This paper aims firstly to outline the inspirational principles that motivate EPANNs (Section 2). The main properties and aims of EPANNs, and suitable evolutionary algorithms are presented in Section 3. The body of research that advanced EPANNs is brought together and described in Section 4. Finally, the paper outlines new research directions, opportunities, and challenges for EPANNs (Section 5).

2. Inspiration

EPANNs are inspired by a particularly large variety of ideas from biology, computer science, and other areas (Downing, 2015; Floreano & Mattiussi, 2008). It is also the nature of inspiration to be subjective, and some of the topics described in this section will resonate differently to different readers. We will touch upon large themes and research areas with the intent to provide the background and motivations to introduce the properties, the progress, and the future directions of EPANNs in the remainder of the paper.

The precise genetic make-up of an organism, acquired through millions of years of evolution, is now known to determine the ultimate capabilities of complex biological neural systems (Deary, Johnson, & Houlihan, 2009; Hopkins, Russell, & Schaeffer, 2014): different animal species manifest different levels of skills and intelligence because of their different genetic blueprint (Schreiweis et al., 2014). The intricate structure of the brain emerges from one single zygote cell through a developmental process (Kolb & Gibb, 2011), which is also strongly affected by input–output learning experiences throughout early life (Hensch et al., 1998; Kolb & Gibb, 2011). Yet high levels of plasticity are maintained throughout the entire lifespan (Kiyota, 2017; Merzenich et al., 1984). These dimensions, *evolution*, *development* and *learning*, also known as the phylogenetic (evolution), ontogenetic (development) and epigenetic (learning) (POE) dimensions (Sipper et al., 1997), are essential for the emergence of biological plastic brains.

The POE dimensions lead to a number of research questions. Can artificial intelligence systems be entirely engineered by humans, or do they need to undergo a less human-controlled process such as evolution? Do intelligent systems need to learn, or could they be born already knowing? Is there an optimal balance between innate and acquired knowledge? Opinions and approaches are diverse. Additionally, artificial systems do not need to implement the same constraints and limitations as biological systems (Bullinaria, 2003). Thus, inspiration is not simple imitation.

EPANNs assume that both evolution and learning, if not strictly necessary, are conducive to the emergence of a strongly bio-inspired artificial intelligence. While artificial evolution is justified by the remarkable achievements of natural evolution, the role of learning has gathered significance in recent years. We are now more aware of the high level of brain plasticity, and its impact on the manifestation of behaviors and skills (Doidge, 2007; Grossberg, 2012; LeDoux, 2003). Concurrently, recent developments in machine learning (Alpaydin, 2014; Michalski, Carbonell, & Mitchell, 2013) and neural learning (Deng, Hinton, & Kingsbury, 2013; LeCun, Bengio, & Hinton, 2015; Silver et al., 2016), have highlighted the importance of learning from large input–output data and extensive training. Other areas of cognition such as the capabilities to make predictions (Hawkins & Blakeslee, 2007), to establish associations (Rescorla, 2014) and to regulate behaviors (Carver

& Scheier, 2012) are also based on learning from experience. Interestingly, skills such as reading, playing a musical instrument, or driving a car, are mastered even if none of those behaviors existed during evolutionary time, and yet they are mostly unique to humans. Thus, human genetic instructions have evolved not to learn specific tasks, but to synthesize recipes to learn a large variety of general skills. We can conclude that the evolutionary search of learning mechanisms in EPANNs tackles both the long-running *nature vs. nurture* debate (Moore, 2003), and the fundamental AI research that studies learning algorithms. This review focuses on *evolution* and *learning*, and less on development, which can be interpreted as a form of learning if affected by sensory-motor signals. We refer to Stanley and Miikkulainen (2003) for an overview of artificial developmental theories.

Whilst the range of inspiring ideas is large and heterogeneous, the analysis in this review proposes that such ideas can be grouped under the following areas:

- *natural and artificial evolutionary processes*,
- *plasticity in biological neural networks*,
- *plasticity in artificial neural networks*, and
- *natural and artificial learning environments*.

Fig. 1 graphically summarizes the topics described in Sections 2.1–2.4 from which EPANNs take inspiration.

2.1. Natural and artificial evolutionary processes

A central idea in evolutionary computation (Goldberg & Holland, 1988) is that evolutionary processes, similar to those that occurred in nature during the course of billions of years (Darwin, 1859; Dobzhansky, 1970), can be simulated with computer software. This idea led to the belief that intelligent computer programs could emerge with little human intervention by means of evolution in-silico (Fogel, 2006; Holland & Reitman, 1977; Koza, 1992).

The emergence of evolved intelligent software, however, did not occur as easily as initially hoped. The reasons for the slow progress are not completely understood, but a number of problems have been identified, likely related to the simplicity of the early implementations of evolutionary algorithms and the high computational requirements. Current topics of investigation focus on levels of abstraction, diversity in the population, selection criteria, the concepts of evolvability and scalability (Lehman & Stanley, 2013; Pigliucci, 2008; Wagner & Altenberg, 1996), the encoding of genetic information through the genotype–phenotype mapping processes (Hornby, Pollack, et al., 2002; Wagner & Altenberg, 1996), the deception of fitness objectives, and how to avoid them (Lehman & Stanley, 2008; Stanley & Lehman, 2015). It is also not clear yet which stepping stones were most challenging for natural evolution (Roff, 1993; Stanley & Lehman, 2015) in the evolutionary path to intelligent and complex forms of life. This lack of knowledge highlights that our understanding of natural evolutionary processes is incomplete, and thus the potential to exploit computational methods is not fully realized. In particular, EPANN research is concerned with those evolutionary algorithms that allow the most creative, open-ended and scalable design. Effective evolutionary algorithms and their desirable features for EPANNs are detailed later in Section 3.3.

2.2. Plasticity in biological neural networks

Biological neural networks demonstrate lifelong learning, from simple reflex adaptation to the acquisition of astonishing skills such as social behavior and learning to speak one or more languages. Those skills are acquired through experiencing stimuli and actions and by means of learning mechanisms not yet fully understood (Bear, Connors, & Paradiso, 2007). The brief overview here
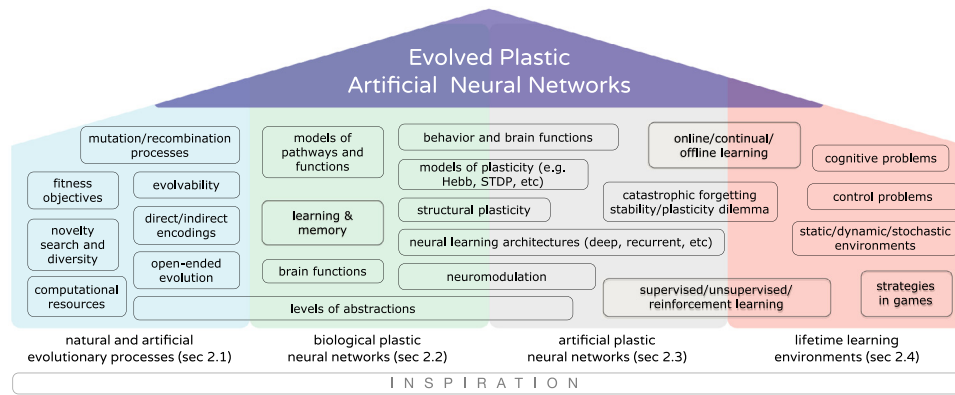
Fig. 1. EPANN inspiration principles described in Section 2. The figure brings together the variety of inspirational topics and areas with no pretense of taxonomic completeness.

outlines that adaptation and learning strongly rely on neural plasticity, understood as "the ability of neurons to change in form and function in response to alterations in their environment" (Kaas, 2001).

The fact that experiences guide lifelong learning was extensively documented in the works of *behaviorism* by scientists such as Hull (1943), Pavlov (1927), Skinner (1938, 1953), and Thorndike (1911) who started to test scientifically how experiences cause a change in behavior, in particular as a result of learning associations and observable behavioral patterns (Staddon, 1983). This approach means linking behavior to brain mechanisms and dynamics, an idea initially entertained by Freud (Køppe, 1983) and later by other illustrious scientists (Hebb, 1949; Kandel, 2007). A seminal contribution to link psychology to physiology came from Hebb (1949), whose principle that neurons that *fire together, wire together* is relevant to understanding both low level neural wiring and high level behaviors (Doidge, 2007). Much later, a Hebbian-compatible rule that regulates synaptic changes according to the firing times of the presynaptic and postsynaptic neurons was observed by Markram, Lübke, Frotscher, and Sakmann (1997) and named Spike-Timing-Dependent Plasticity (STDP).

The seminal work of Kandel and Tauc (1965), and following studies (Clark & Kandel, 1984), were the first to demonstrate that changes in the strength of connectivity among neurons, i.e. plasticity, relates to behavior learning. Walters and Byrne (1983) showed that, by means of plasticity, a single neuron can perform associative learning such as classical conditioning, a class of learning that is observed in simple neural systems such as that of the Aplysia (Carew, Walters, & Kandel, 1981). Plasticity driven by local neural stimuli, i.e. compatible with the Hebb synapse (Brown, Kairiss, & Keenan, 1990; Hebb, 1949), is responsible not only for fine tuning, but also for building a working visual system in the cat's visual cortex (Rauschecker & Singer, 1981).

Biological plastic neural networks are also capable of structural plasticity, which creates new pathways among neurons (Chklovskii, Mel, & Svoboda, 2004; Lamprecht & LeDoux, 2004; Russo et al., 2010): it occurs primarily during development, but there is evidence that it continues well into adulthood (Pascual-Leone, Amedi, Fregni, & Merabet, 2005). Axon growth, known to be regulated by neurotrophic nerve growth factors (Tessier-Lavigne & Goodman, 1996), was also modeled computationally in Roberts et al. (2014). Developmental processes and neural plasticity are often indistinguishable (Kolb, 1989; Pascual-Leone et al., 2005) because the brain is highly plastic during development. Neuroscientific advances reviewed in Damasio (1999), Doidge (2007), Draganski and May (2008), LeDoux (2003) and Pascual-Leone et al. (2005) outline the importance of structural plasticity in learning motor patterns, associations, and ways of thinking. Both structural and functional plasticity in biology are essential to acquiring long-lasting new skills, and for this reason appears to be an important inspiration for EPANNs.

Finally, an important mechanism for plasticity and behavior is neuromodulation (Bailey, Giustetto, Huang, Hawkins, & Kandel, 2000; Gu, 2002; Marder & Thirumalai, 2002). Modulatory chemicals such as acetylcholine (ACh), norepinephrine (NE), serotonin (5-HT) and dopamine (DA) appear to regulate a large variety of neural functions, from arousal and behavior (Harris-Warrick & Marder, 1991; Hasselmo & Schnell, 1994; Katz, 1995; Katz & Frost, 1996; Marder, 1996), to pattern generation (Katz, Getting, & Frost, 1994), to memory consolidation (Hasselmo, 1995, 1999; Kupfermann, 1987; Marder, 1996). Learning by reward in monkeys was linked to dopaminergic activity during the 1990s with studies by Schultz (1998), Schultz, Apicella, and Ljungberg (1993) and Schultz, Dayan, and Montague (1997). For these reasons, neuromodulation is considered an essential element in cognitive and behavioral processes, and has been the topic of a considerable amount of work in EPANNs (Section 4.5).

This compact overview suggests that neural plasticity encompasses an important set of mechanisms, regulated by a rich set of signals and dynamics currently mostly ignored in ANNs. Thus, EPANNs can be used to explore, via evolutionary search, the potential of plasticity and to answer questions such as: (1) How does a brain-like structure form – driven both by genetic instructions and neural activity – and acquire functions and behaviors? (2) What are the key plasticity mechanisms from biology that can be applied to artificial systems such as EPANNs? (3) Can memories, skills, and behaviors be stored in plastic synaptic connections, in patterns of activities, or in a combination of both? Whilst neuroscience continues to provide inspiration and insight into plasticity in biological brains, EPANNs serve the complementary objective of seeking, implementing, and verifying designs of bio-inspired methods for adaptation, learning, and intelligent behavior.

2.3. Plasticity in artificial neural networks

In EPANN experiments, evolution can be seen as a meta-learning process. Thus, established learning rules for ANNs are often used as *ingredients* that evolution uses to search for good parameter configurations, efficient combinations of rules and network topologies, new functions representing novel learning rules, etc. EPANN experiments are suited to include the largest possible variety of rules because of (1) the variety of possible tasks in a simulated behavioral experiment and (2) the flexibility of evolution to combine rules with no assumptions about their dynamics. The following gives a snapshot of the extent and scope of various learning algorithms for ANN that can be used as building blocks of EPANNs.

In supervised learning, *backpropagation* is the most popular learning rule used to train both shallow and deep networks (LeCun et al., 2015; Rumelhart, Hinton, & Williams, 1988; Widrow & Lehr, 1990) for classification or regression. Unsupervised learning is implemented in neural networks with self-organizing maps (SOM) (Kohonen, 1982, 1990), auto-encoders (Bourlard & Kamp, 1988), restricted Boltzmann machines (RBM) (Hinton & Salakhutdinov, 2006), Hebbian plasticity (Cooper, 2005; Gerstner & Kistler, 2002a; Hebb, 1949), generative adversarial networks (Goodfellow et al., 2014), and various combinations of the above. RBM learning is considered related to the free-energy principle, proposed by Friston (2009) as a central principle governing learning in the brain. Hebbian rules, in particular, given their biological plausibility and unsupervised learning, are a particularly important inspirational principle for EPANNs. Variations (Willshaw & Dayan, 1990) have been proposed to include, e.g., terms to achieve stability (Bienenstock, Cooper, & Munro, 1982; Oja, 1982) and various constraints (Miller & Mackay, 1994), or more advanced update dynamics such as dual weights for fast and slow decay (Bullinaria, 2009a; Hinton & Plaut, 1987; Levy & Bairaktaris, 1995; Soltoggio, 2015). Hebbian rules have been recently proposed to minimize defined cost functions (Bahroun, Hunsicker, & Soltoggio, 2017; Pehlevan, Hu, & Chklovskii, 2015), and more advanced systems have used backpropagation as meta-learning to tune Hebbian rules (Miconi, Clune, & Stanley, 2018).

Neuromodulated plasticity (Fellous & Linster, 1998) is often used to implement reward-learning in neural networks. Such a modulation of signals, or gated learning (Abbott, 1990), allows for amplification or reduction of signals and has been implemented in numerous models (Alexander & Sporns, 2002; Baxter, Canavier, Clark, & Byrne, 1999; Birmingham, 2001; Doya, 2002; Fujii, Saito, Nakahira, & Ishiguro, 2002; Krichmar, 2008; Sporns & Alexander, 2003; Suri, 2002; Suri, Bargas, & Arbib, 2001; Ziemke & Thieme, 2002).

Plastic neural models are also used to demonstrate how behavior can emerge from a particular circuitry modeled after biological brains. Computational models of, e.g., the basal ganglia and modulatory systems may propose plasticity mechanisms and aim to demonstrate the computational relations among various nuclei, pathways, and learning processes (Krichmar, 2008; Schroll & Hamker, 2015; Vitay & Hamker, 2010).

Finally, plasticity rules for spiking neural networks (Maass & Bishop, 2001) aim to demonstrate unique learning mechanisms that emerge from spiking dynamics (Izhikevich, 2006, 2007; Markram et al., 1997), as well as model biological synaptic plasticity (Gerstner & Kistler, 2002b).

Plasticity in neural networks, when continuously active, was also observed to cause catastrophic forgetting (Robins, 1995). If learning occurs continuously, new information or skills have the potential to overwrite previously acquired information or skills, a problem also known as plasticity-stability dilemma (Abraham & Robins, 2005; Finnie & Nader, 2012).

In conclusion, a large range of plasticity rules for neural networks have been proposed to solve different problems. In the majority of cases, a careful matching and engineering of rules, architectures and problems is necessary, requiring considerable design effort. The variety of algorithms also reflects the variety of problems and solutions. One key aspect is that EPANN systems can effectively play with all possible plasticity rules to offer a unique testing tool and assess the effectiveness and suitability of different models, or their combination, in a variety of different scenarios.

2.4. Lifelong learning environments

One aspect of EPANNs is that they can continuously improve and adapt both at the evolutionary scale and at the lifetime scale in a virtually unlimited range of problems. Natural environments are an inspiration for EPANNs because organisms have evolved to adapt to, and learn in, a variety of conditions. Fundamental questions are: what makes an environment conducive to the evolution of learning and intelligence? What are the challenges faced by learning organisms in the natural world, and how does biological learning cope with those? How can those challenges be abstracted and ported to a simulated environment for EPANNs? EPANNs employ lifelong learning environments in the attempt to provide answers to such questions.

In the early phases of AI, logic and reasoning were thought to be the essence of intelligence (Cervier, 1993), so symbolic input–output mappings were employed as tests. Soon it became evident that intelligence is not only symbol manipulation, but resides also in subsymbolic problem solving abilities emerging from the interaction of brain, body, and environment (Sims, 1994; Steels, 1993). More complex simulators of real-life environments and body–environment interaction were developed to better represent the *enactivist* philosophy (Varela, Thompson, & Rosch, 2017) and cognitive theories on the emergence of cognition (Butz & Kutter, 2016). Other environments focus on high-level planning and strategies required, e.g., when applying AI to games (Allis et al., 1994; Millington & Funge, 2016) or articulated robotic tasks. Planning and decision making with high bandwidth sensory-motor information flow such as those required for humanoid robots or self-driving vehicles are current benchmarks for lifelong learning systems. Finally, environments in which affective dynamics and feelings play a role are recognized as important for human well-being (De Botton, 2016; Lee & Narayanan, 2005). Those intelligence-testing environments are effectively the "worlds" in which EPANNs may evolve and live in embodied forms, and thus largely shape the EPANN design process.

Such different testing environments have very different features, dynamics, and goals that fall into different machine learning problems. For example, supervised learning can be mapped to a fitness function when a precise target behavior exists and is known. If it is useful to find relationships and regularities in the environment, unsupervised learning, representation learning, or modularity can be evolved (Bullinaria, 2007b). If the environment provides rewards, the objective may be to search for behavioral policies that lead to collecting rewards: algorithms specifically designed to do so are called reinforcement learning (Sutton & Barto, 1998). While reinforcement learning maximizes a reward or fitness, recent advances in evolutionary computation (Lehman & Stanley, 2011; Stanley & Lehman, 2015) suggest that it is not always the *fittest*, but at times it is the *novel* individual or behavior that can exploit environmental niches, thus leading to creative evolutionary processes similar to those observed in nature. Temporal dynamics, i.e. when a system requires to behave over time according to complex dynamics, need different computational structures from functions with no temporal dynamics. This case is typical for EPANN experiments that may exhibit a large variety of time scales in complex behavioral tasks. With traditional approaches, all those different cases require careful manual design to solve each problem. In contrast, the evolution in EPANNs can be designed to address most problems by mapping a measure of success to a fitness value, thus searching for solutions in an increasingly large variety of problems and environments. In conclusion, lifelong learning environments of different types can be used with EPANNs to explore innovative and creative solutions with limited human intervention and design.

3. Properties, aims, and evolutionary algorithms for EPANNs

Having introduced the inspirational principles of EPANNs, we now propose: a list of primary properties that define EPANNs (Section 3.1); the principal aims of EPANN studies (Section 3.2); and a list of desired properties of EAs for EPANNs (Section 3.3).

3.1. EPANN properties

EPANNs, as formalized in this review, are defined as artificial neural networks with the following properties:

Property 1 (*Evolution*). Parts of an EPANN are determined by an evolutionary algorithm. *Inspired by natural and artificial evolution (Section* 2.1*), such search dynamics in EPANNs implement a design process.*

Property 2 (*Plasticity*). Parts of the functions that process signals within the network change in response to signals propagated through the network, and those signals are at least partially affected by stimuli. *Inspired by biological findings on neural plasticity (Section* 2.2*) and empowered by the effectiveness of plasticity in neural models (Section* 2.3*), EPANNs either include such mechanisms or are set up with the conditions to evolve them.*

Property 3 (*Discovery of Learning*). Properties 1 and 2 are implemented to discover, through evolution, learning dynamics within an artificial neural network. *Thus, an EPANN uses both evolution and plasticity in synergy to achieve learning. Such a property can be present in different degrees, from the highest degree in which no learning occurs before evolution and it is therefore discovered from scratch, to the lowest degree in which learning is fine-tuned and optimized, e.g., when evolution is seeded with proven learning structures. Given the very diverse interpretations of* learning *in different domains, we refer to* Michalski et al. (2013) *for an overview, or otherwise assume the general machine learning definition by* Michalski et al. (2013).[1]

Property 4 (*Generality*). Properties 1 to 3 are independent from the learning problem(s) and from the plasticity mechanism(s) that are implemented or evolved in an EPANN. *Exploiting the flexibility of evolution and learning, (1) EPANNs can evolve to solve problems of different nature, complexity, and time scales (Section* 2.4*); (2) EPANNs are not limited to specific learning dynamics because often it is the aim of the experiment to discover the learning mechanism throughout evolution and interaction with the environment.*

In summary, the EPANNs' properties indicate that within simple assumptions, i.e., using plasticity and evolution, EPANNs are set to investigate the design of learning in creative ways for a large variety of learning problems.

3.2. Aims

Given the above properties, EPANN experiments can be set up to achieve the following aims.

Aim 1: *Autonomously design learning systems*: in an EPANN experiment, it is essential to delegate some design choices of a learning system to the evolutionary process, so that the design is not entirely determined by the human expert and can be automated. The following sub-aims can then be identified.

Aim 1.1: *Bootstrap of learning from scratch*: in an EPANN experiment, it may be desirable to initialize the system with no learning capabilities before evolution takes place, so that the best learning dynamics for a given environment is evolved rather than human-designed.

Aim 1.2: *Optimize performance*: as opposed to *Aim 1.1*, it may be desirable to initialize the system with well know learning capabilities, so that evolution can autonomously optimize the system, e.g., for final performance after learning.

Aim 1.3: *Recover performance in unseen conditions*: in an EPANN experiment, the desired outcome may be to enable the learning system to autonomously evolve from solving a set of problems to another set without human intervention.

Aim 2: *Test the computational advantages of particular neural components*: the aim of an EPANN experiment might be to test whether particular neural dynamics or components have an evolutionary advantage when implementing particular learning functions. The presence of particular neural component may be fostered by evolutionary selection.

Aim 3: *Derive hypotheses on the emergence of biological learning*: an aim may be to draw similarities or suggest hypotheses on how learning evolved in biological systems, particularly in combination with *Aim 1.1* (bootstrap of learning).

Aim 1 is always present in any EPANN because it derives from the EPANN properties. The other aims may be present in different EPANN studies and can be expanded into more detailed and specific research hypotheses.

3.3. Evolutionary algorithms for EPANNs

In contrast to parameter optimization (Bäck & Schwefel, 1993) in which search spaces are often of fixed dimension and static, EPANNs evolve in dynamic search spaces in which learning further increases the complexity of the evolutionary search and of the problem itself. Evolutionary algorithms (Holland, 1975; Michalewicz, 1994) for EPANNs often require additional advanced features to cope with the challenges of the evolution of learning and open evolutionary design (Bentley, 1999). The analysis in this review suggests that evolutionary algorithms (EAs) for EPANNs may implement the following desirable properties.

3.3.1. Variable genotype length and growing complexity

For some learning problems, the size and properties of a network that can solve them are not known in advance. Therefore, a desirable property of the EA for EPANNs is that of increasing the length of the genotype, and thus the information contained in it, as evolution may discover increasingly more complex strategies and solutions that may require larger networks (see, e.g., Stanley and Miikkulainen (2002)).

3.3.2. Indirect genotype to phenotype encoding

In nature, phenotypes are expressions of a more compact representation: the genetic code. Similarly, EAs may represent genetic information in a compact form, which is then mapped to a larger phenotype. Although EPANNs do not require such a property, such an approach promises better scalability to large networks (see, e.g., Risi and Stanley (2010)).

3.3.3. Expressing regularities, repetitions, and patterns

Indirect encodings are beneficial when they can use one set of instructions in the genotype to generate more parts in the phenotype. This may involve expressing regularities like symmetry (e.g. symmetrical neural architectures), repetition (e.g. neural modules), repetition with variation (similar neural modules), and patterns, e.g., motifs in the neural architecture (see, e.g., Stanley, 2007).

3.3.4. Effective exploration via mutation and recombination

Genetic mutation and sexual reproduction in nature allow for the expression of a variety of phenotypes and for the exploration of new solutions, but seldom lead to highly unfit individuals (Ay, Flack, & Krakauer, 2007). Similarly, EAs for EPANNs need to be able to effectively mutate and recombine genomes without destroying the essential properties of the solutions. EAs may use recombination to generate new solutions from two parents: how to effectively recombine genetic information from two EPANNs is still an open question (see, e.g., tracking genes through historical marking in NEAT Stanley & Miikkulainen, 2002).

---

[1] A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

3.3.5. Genetic encoding of plasticity rules

Just as neural networks need a genetic encoding to be evolved, so do plasticity rules. EPANN algorithms require the integration of such a rule in the genome. The encoding may be restricted to a simple parameter search, or evolution may search a larger space of arbitrary and general plasticity rules. Plasticity may also be applied to all or parts of the network, thus effectively implementing the evolution of learning architectures.

3.3.6. Diversity, survival criteria, and low selection pressure

The variety of solutions in nature seems to suggest that diversity is a key aspect of natural evolution. EAs for EPANNs are likely to perform better when they can maintain diversity in the population, both at the genotype and phenotype levels. Local selection mechanisms were shown to perform well in EPANN experiments (Soltoggio, 2008a). Niche exploration and behavioral diversity (Lehman & Stanley, 2011) could also play a key role for creative design processes. Low selection pressure and survival criteria might be crucial to evolve learning in deceptive environments (see Section 4.4).

3.3.7. Open-ended evolution

Evolutionary optimization aims to quickly optimize parameters and reach a fitness plateau. On the contrary, EAs for EPANNs often seek a more open-ended evolution that can evolve indefinitely more complex solutions given sufficient computational power (Stanley, Lehman, & Soros, 2017; Taylor et al., 2016).

3.3.8. Implementations

Many EAs include one or more of these desirable properties (Fogel, 2006). Due to the complexity of neural network design, the field of *neuroevolution* was the first to explore most of those extensions of standard evolutionary algorithms. Popular algorithms include early work of Angeline, Saunders, and Pollack (1994) and Yao and Liu (1997) to evolve fixed weights (i.e. weights that do not change while the agent interacts with its environment) and the topology of arbitrary neural networks, e.g., recurrent (addressing 3.3.1 and 3.3.4). Neuroevolution of Augmenting Topologies (NEAT) (Stanley & Miikkulainen, 2002) leverages three main aspects: a recombination operator intended to preserve network function (addressing 3.3.4); speciation (addressing 3.3.6); and evolution from small to larger networks (addressing 3.3.1). Similarly, EPANN-tailored EAs in Soltoggio (2008a) employ local selection mechanisms to maintain diversity. Analog Genetic Encoding (AGE) (Mattiussi & Floreano, 2007) is a method for indirect genotype to phenotype mapping that can be used in combination with evolution to design arbitrary network topologies (addressing 3.3.1 and 3.3.2) and was used with EPANNs. HyperNEAT (Stanley, D'Ambrosio, & Gauci, 2009) is an indirect representation method that combines the NEAT algorithm with compositional patterns producing networks (CPPN) (Stanley, 2007) (3.3.1, 3.3.2 and 3.3.3). Novelty search (Lehman & Stanley, 2008, 2011) was introduced as an alternative to the *survival of the fittest* as a selection mechanism (3.3.6). Initially, the majority of these neuroevolution algorithms were not devised to evolve plastic networks, but rather fixed networks in which the final synaptic weights were encoded in the genome. To operate with EPANNs, these algorithms need to integrate additional genotypical instructions to evolve plasticity rules (3.3.5).

By adding these EA features (3.3.1–3.3.7) to standard evolutionary algorithms, EPANNs aim to search extremely large search spaces in fundamentally different and more creative ways than traditional heuristic searches of parameters and hyper-parameters.

The process of evolving plastic neural networks is depicted in Fig. 2.

4. Progress on evolving artificial plastic neural networks

This section reviews studies that have evolved plastic neural networks (EPANNs). The analysis of the field has led to present studies in six sections, although some studies may be described in more than one section: the evolution of plasticity rules, the evolution of neural architectures, EPANNs in evolutionary robotics, the evolutionary discovery of learning, the evolution of neuromodulation, and the evolution of indirectly encoded plasticity. Accordingly, Fig. 3 provides our perspective on the organization of the field, reflected in the structure of this paper.

4.1. Evolving plasticity rules

Early EPANN experiments evolved the parameters of learning rules for fixed or hand-designed ANN architectures. Learning rules are functions that change the connection weight w between two neurons, and are generally expressed as

$$\Delta w = f(\mathbf{x}, \theta), \tag{1}$$

where $\mathbf{x}$ is a vector of neural signals and θ is a vector of fixed parameters that can be searched by evolution. The incoming connection weights $\mathbf{w}$ to a neuron i are used to determine the activation value

$$x_i = \sigma \left(\sum (w_{ji} \cdot x_j) \right), \tag{2}$$

where x_j are the activation values of presynaptic neurons that connect to neuron i with the weights w_{ji}, and σ is a nonlinear function such as the sigmoid or the hyperbolic tangent. The vector $\mathbf{x}$ may provide local signals such as pre and postsynaptic activities, the value of the weight w, and modulatory or error signals.

Bengio, Bengio, and Cloutier (1990) and Bengio, Bengio, Cloutier, and Gecsei (1992) proposed the optimization of the parameters θ of generic learning rules with gradient descent, simulated annealing, and evolutionary search for problems such as conditioning, boolean function mapping, and classification. Those studies are also among the first to include a modulatory term in the learning rules. The optimization was shown to improve the performance in those different tasks with respect to manual parameter settings. Chalmers (1990) evolved a learning rule that applied to every connection and had a teaching signal. He found that, in 20% of the evolutionary runs, the algorithm rediscovered, through evolution, the well-known delta rule, or Widrow–Hoff rule (Widrow, Hoff, et al., 1960), used in backpropagation, thereby demonstrating the validity of evolution as an autonomous tool to discover learning. Fontanari and Meir (1991) used the same approach of Chalmers (1990) but constrained the weights to binary values. Also in this case, evolution autonomously rediscovered a hand-designed rule, the directed drift rule by Venkatesh (1993). They also observed that the performance on new tasks was better when the network evolved on a larger set of tasks, possibly encouraging the evolution of more general learning strategies.

With backpropagation of errors (Widrow & Lehr, 1990), the input vector $\mathbf{x}$ of Eq. (1) requires an error signal between each input/output pair. In contrast, rules that use only local signals have been a more popular choice for EPANNs, though this is changing with the rise in effectiveness of deep learning using backpropagation and related methods. In the simplest form, the product of presynaptic (x_j) and postsynaptic (x_i) activities, and a learning rate η

$$\Delta w = \eta \cdot x_j \cdot x_i \tag{3}$$

is known as Hebbian plasticity (Cooper, 2005; Hebb, 1949). More generally, any function as in Eq. (1) that uses only local signals is considered a local plasticity rule for unsupervised learning. Baxter (1992) evolved a network that applied the basic Hebbian rule in
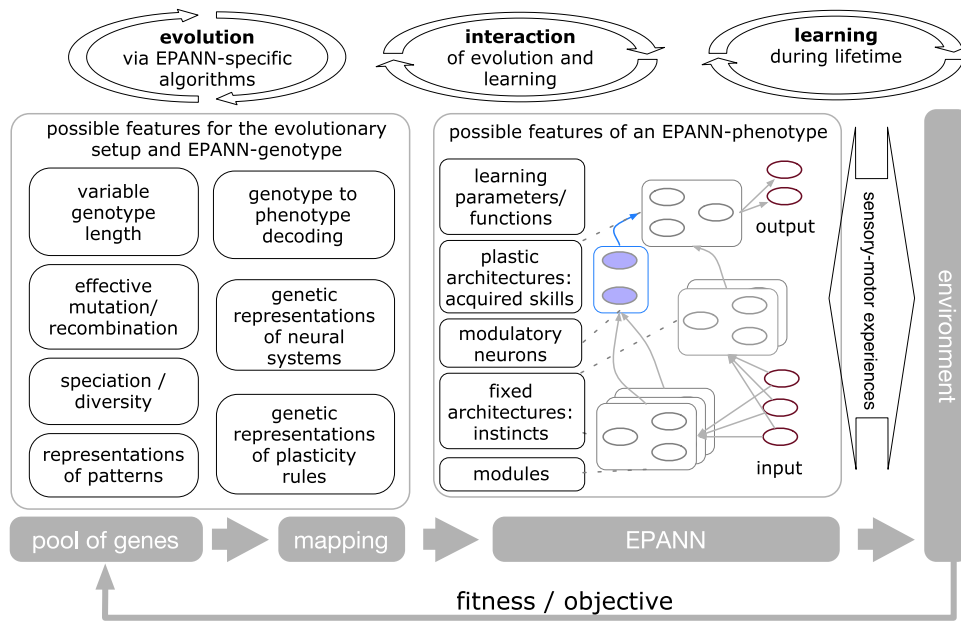
Fig. 2. Main elements of an EPANN setup in which simulated evolution (left) and an environment (right) allow for an EPANN (center) to evolve through generations and learn within a lifetime in the environment. Possible features of the genotype, evolutionary process and phenotype are illustrated as an example.
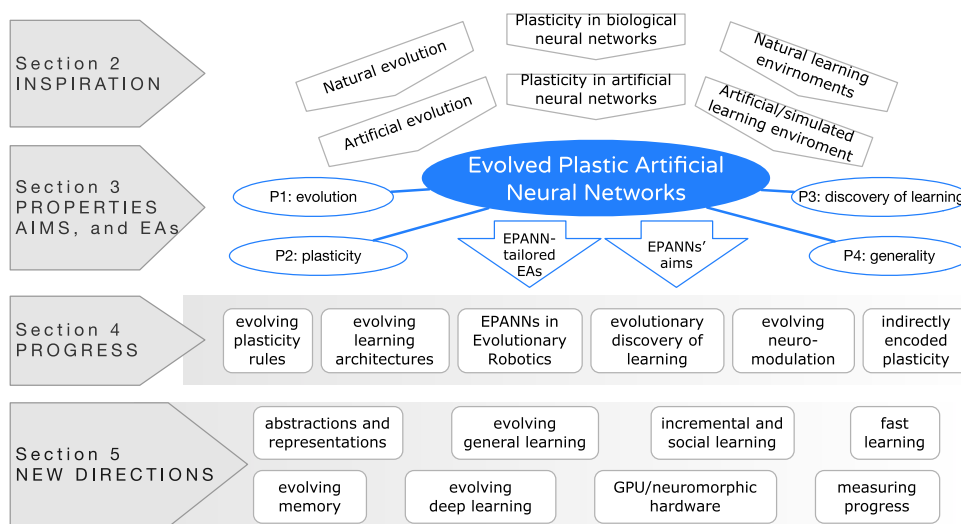

Fig. 3. Organization of the field of EPANNs, reflected in the structure of this paper.

Eq. (3) to a subset of weights (determined by evolution) to learn four functions of one variable. The network, called Local Binary Neural Net (LBNN), evolved to change its weights to one of two possible values (± 1), or have fixed weights. The experiment proved that learning can evolve when rules are optimized and applied to individual connections.

Nolfi and Parisi (1993) evolved networks with "auto-teaching" inputs, which could then provide an error signal for the network to adjust weights during lifetime. The implication is that error signals do not always need to be hand-designed but can be discovered by evolution to fit a particular problem. A set of eight different local rules was used in Rolls and Stringer (2000) to investigate the evolution of rules in combination with the number of synaptic connections for each neuron, different neuron classes, and other network parameters. They found that evolution was effective in selecting specific rules from a large set to solve simple linear problems. In Maniadakis and Trahanias (2006), co-evolution was used to evolve agents (each being a network) that could use ten

different types of Hebbian-like learning rules for simple navigation tasks: the authors reported that, despite the increase in the search space, using many different learning rules results in better performance but, understandably, a more difficult analysis of the evolved systems. Meng, Jin, and Yin (2011) evolved a gene regulatory network that in turn determined the learning parameters of the BCM rule (Bienenstock et al., 1982), showing promising performance in time series classification and other supervised tasks.

One general finding from these studies is that evolution operates well within large search spaces, particularly when a large set of evolvable rules is used.

4.2. Evolving learning architectures

The interdependency of learning rules and neural architectures led to experiments in which evolution had more freedom on the network's design. The evolution of architectures in ANNs may involve searching an optimal number of hidden neurons, the

number of layers in a network, particular topologies or modules, the type of connectivity, and other properties of the network's architecture. In EPANNs, evolving learning architectures implies more specifically to discover a combination of architectures and learning rules whose synergetic matching enables particular learning dynamics. As opposed to biological networks, EPANNs do not have the neurophysiological constraints, e.g., short neural connections, sparsity, brain size limits, etc., that impose limitations on the natural evolution of biological networks. Thus, biologically implausible artificial systems may nevertheless be evolved in computer simulations (Bullinaria, 2007b, 2009b).

One seminal early study by Happel and Murre (1994) proposed the evolutionary design of modular neural networks, called CALM (Murre, 1992), in which modules could perform unsupervised learning, and the intermodule connectivity was shaped by Hebbian rules. The network learned categorization problems (simple patterns and hand written digits recognition), and showed that the use of evolution led to enhanced learning and better generalization capabilities in comparison to hand-designed networks. In Arifovic and Gencay (2001), the authors used evolution to optimize the number of inputs and hidden nodes, and allowed connections in a feedforward neural network to be trained with backpropagation. Abraham (2004) proposed a method called Meta-Learning Evolutionary Artificial Neural Networks (MLEANN) in which evolution searches for initial weights, neural architectures and transfer functions for a range of supervised learning problems to be solved by evolved networks. The evolved networks were tested in time series prediction and compared with manually designed networks. The analysis showed that evolution consistently found networks with better performance than the hand-designed structures. Khan, Miller, and Halliday (2008) proposed an evolutionary developmental system that created an architecture that adapted with learning: the network had a dynamic morphology in which neurons could be inserted or deleted, and synaptic connections formed and changed in response to stimuli. The networks were evolved with Cartesian genetic programming and appeared to improve their performance while playing checkers over the generations. Downing (2007) looked at different computational models of neurogenesis to evolve learning architectures. The proposed evolutionary developmental system focused in particular on abstraction levels and principles such as Neural Darwinism (Edelman & Tononi, 2000). A combination of evolution of recurrent networks with a linear learner in the output was proposed in Schmidhuber, Wierstra, Gagliolo, and Gomez (2007), showing that the evolved RNNs were more compact and resulted in better learning than randomly initialized echo state networks (Jaeger & Haas, 2004). In Khan, Khan, and Miller (2011b), Khan, Miller, and Halliday (2011a) and Khan and Miller (2014), the authors introduced a large number of bio-inspired mechanisms to evolve networks with rich learning dynamics. The idea was to use evolution to design a network that was capable of advanced plasticity such as dendrite branch and axon growth and shrinkage, neuron insertion and destruction, and many others. The system was tested on the Wumpus World (Russell & Norvig, 2013), a fairly simple problem with no learning required, but the purpose was to show that evolution can design working control networks even within a large search space.

In summary, learning mechanisms and neural architectures are strongly interdependent, but a large set of available dynamics seem to facilitate the evolution of learning. Thus, EPANNs become more effective precisely when manual network design becomes less practical because of complexity and rich dynamics.

4.3. EPANNs in evolutionary robotics

Evolutionary robotics (ER) (Cliff, Husbands, & Harvey, 1993; Floreano & Mondada, 1994, 1996; Floreano & Nolfi, 2004; Urzelai & Floreano, 2000) contributed strongly to the development of EPANNs, providing a testbed for applied controllers in robotics. Although ER had no specific assumptions on neural systems or plasticity (Smith, 2002), robotics experiments suggested that neural control structures evolved with fixed weights perform less well than those evolved with plastic weights (Floreano & Urzelai, 2001a; Nolfi & Parisi, 1996). In a conditional phototaxis robotic experiment,[2] Floreano and Urzelai (2001b) reported that networks evolved faster when synaptic plasticity and neural architectures were evolved simultaneously. In particular, plastic networks were shown to adapt better in the transition from simulation to real robots. The better simulation-to-hardware transition, and the increased adaptability in changing ER environments, appeared intuitive and supported by evidence (Nolfi & Parisi, 1996). However, the precise nature and magnitude of the changes from simulation to hardware is not always easy to quantify: those studies do not clearly outline the precise principles, e.g., better or adaptive feedback control, minimization principles, etc., that are discovered by evolution with plasticity to produce those advantages. In fact, the behavioral changes required to switch behaviors in simple ER experiments can also take place with non-plastic recurrent neural networks because evolution can discover recurrent units that act as switches. A study in 2003 observed similar performance in an associative learning task (food foraging) when comparing plastic and non-plastic recurrent networks (Stanley, Bryant, & Miikkulainen, 2003). Recurrent networks with leaky integrators as neurons (Beer & Gallagher, 1992; Funahashi & Nakamura, 1993; Yamauchi & Beer, 1994) were also observed to achieve similar performance to plastic networks (Blynel & Floreano, 2002, 2003). These early studies indicate that the evolution of learning with plastic networks was at that point still a proof-of-concept rather than a superior learning tool: aided by evolutionary search, networks with recurrent connections and fixed weights could create recurrent nodes, retain information and achieve similar learning performance to networks with plastic weights.

Nevertheless, ER maintained a focus on plasticity as demonstrated, e.g., in The Cyber Rodent Project (Doya & Uchibe, 2005) that investigated the evolution of learning by seeking to implement a number of features such as (1) evolution of neural controllers, (2) learning of foraging and mating behaviors, (3) evolution of learning architectures and meta-parameters, (4) simultaneous learning of multiple agents in a body, and (5) learning and evolution in a self-sustained colony. Plasticity in the form of modulated neural activation was used in Husbands, Smith, Jakobi, and O'Shea (1998) and Smith, Husbands, Philippides, and O'Shea (2002) with a network that adapts its activation functions according to the diffusion of a simulated gas spreading to the substrate of the network. Although the robotic visual discrimination tasks did not involve learning, the plastic networks appeared to evolve faster than a network evolved with fixed activation functions. Similar conclusions were reached in Di Paolo (2003) and Federici (2005). Di Paolo (2002, 2003) evolved networks with STDP for a wheeled robot to perform positive and negative phototaxis, depending on a conditioned stimulus, and observed that networks with fixed weights could learn but had inferior performance with respect to plastic networks. Federici (2005) evolved plastic networks with STDP and an indirect encoding, showing that plasticity helped performance even if learning was not required. Stability and evolvability of simple robotic controllers were investigated in Hoinville, Tapia Siles,

---

[2] The fitness value was the time spent by a two-wheeled robot in one particular area of the area when a light was on, divided by the total experiment time.

and Hénaff (2011) who focused on EPANNs with homeostatic mechanisms.

Experiments in ER in the 1990s and early 2000s revealed the extent, complexity, and multitude of ideas behind the evolutionary design of learning neuro-robotics controllers. They generally indicate that plasticity helps evolution under a variety of conditions, even when learning is not required, thereby promoting further interest in more specific topics. Among those are the evolutionary discovery of learning, the evolution of neuromodulation, and the evolution of indirectly encoded plasticity, as described in the following.

4.4. Evolutionary discovery of learning

When evolution is used to search for learning mechanisms, two main cases can be distinguished: (1) when learning is used to acquire constant facts about the agent or environment, and (2) when learning is used to acquire changeable facts. The first case, that of *static* or *stationary environments*, is known to be affected by the Baldwin effect (Baldwin, 1896) that suggests an acceleration of evolution when learning occurs during lifetime. A number of studies showed that the Baldwin effect can be observed with computational simulations (Boers, Borst, & Sprinkhuizen-Kuyper, 1995; Bullinaria, 2001; Hinton & Nowlan, 1987; Mayley, 1996; Smith, 1986). With static environments, learning causes a faster transfer of knowledge into the genotype, which can happen when facts are stationary (or constant) across generations. Eventually, a system in those conditions can perform well without learning because it can be born knowing to perform well. However, one limitation is that the genome might grow very large to hold large amount of information, and might, as a result, become less easy to evolve further. A second limitation is that such solutions might not perform well in non-stationary environments.

In the second case, that of *variable* or *non-stationary environments*, facts cannot be embedded in the genotype because those are changeable as, e.g., the location of food in a foraging problem. This case requires the evolution of learning for the performance to be maximized. For this reason, non-stationary reward-based environments, in which the behaviors to obtain rewards may change, are more typically used to study the evolution of learning in EPANNs.

EPANN experiments have been used to observe the advantages of combining learning and evolution, and the complex interaction dynamics that derives (Nolfi & Floreano, 1999). Stone (2007) showed that distributed neural representations accelerate the evolution of adaptive behavior because learning part of a skill induced the automatic acquisition of other skill components. One study in a non-stationary environment (a foraging problem with variable rewards) (Soltoggio, Dürr, Mattiussi, & Floreano, 2007) suggested that evolution discovers, before optimizing, learning in a process that is revealed by discrete fitness stepping stones. At first, non-learning solutions are present in the population. When evolution casually discovers a weak mechanism of learning, it is sufficient to create an evolutionary advantage, so the neural mechanism is subsequently optimized: Fig. 4 shows a sudden jump in the fitness when one agent suddenly evolves a learning strategy: such jumps in fitness graphs are common in evolutionary experiments in which learning is discovered from scratch (*Aim 1.1*), rather than optimized (*Aim 1.2*), and were observed as early as in Fontanari and Meir (1991).

When an environment changes over time, the frequency of those changes plays a role because it determines the time scales that are required from the learning agent. With time scales comparable to a lifetime, evolution may lead to phenotypic plasticity, which is the capacity for a genotype to express different phenotypes in response to different environmental conditions (Lalejini & Ofria, 2016). The frequency of environmental changes was
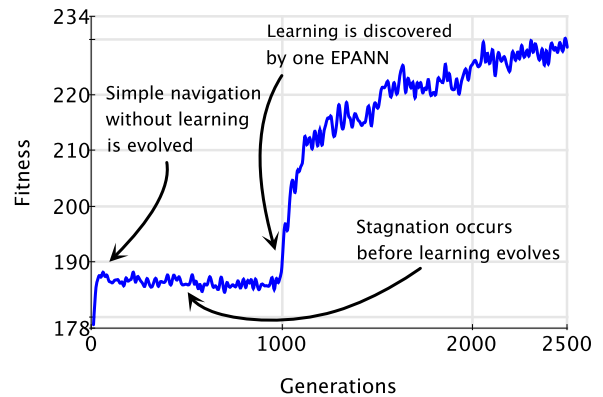

Fig. 4. Discovery of a learning strategy during evolution in a non-stationary reward-based environment, i.e. where learning is required to maximize the fitness because of the variability of the environment. After approximately 1000 generations of 100 individuals, a network evolves to use the reward signal to modify networks weights, and thus implement learning. The graphic is adapted from Soltoggio et al. (2007). This experiment shows that evolution stagnates until learning is discovered. At that point, the evolutionary search hits on a gradient in the search space that improves on the learning as reflected by the fitness values.

observed experimentally in plastic neural networks to affect the evolution of learning (Ellefsen, 2014), revealing a complex relationship between environmental variability and evolved learning. One conclusion is that evolution copes with non-stationary environments by evolving the specific learning that better matches those changes.

The use of reward to guide the discovery of neural learning through evolution was shown to be inherently deceptive in Lehman and Miikkulainen (2014) and Risi, Hughes, and Stanley (2010). In Risi et al. (2010); Risi, Vanderbleek, Hughes, and Stanley (2009), EPANN-controlled simulated robots, evolved in a discrete T-Maze domain, revealed that the stepping stones towards discovering learning are often not rewarded by objective-based performance measures. Those stepping stones to learning receive a lower fitness score than more brittle solutions with no learning but effective behaviors. A solution to this problem was devised in Risi et al. (2010, 2009), in which *novelty search* (Lehman & Stanley, 2008, 2011) was adopted as a substitute for performance in the fitness objective with the aim of finding novel behaviors. Novelty search was observed to perform significantly better in the T-Maze domain. Lehman and Miikkulainen (2014) later showed that novelty search can encourage the evolution of more adaptive behaviors across a variety of different variations of the T-Maze learning tasks. As a consequence, novelty search contributed to a philosophical change by questioning the centrality of objective-driven search in current evolutionary algorithms (Stanley & Lehman, 2015). By rewarding novel behaviors, novelty search validates the importance of exploration or curiosity, previously proposed in Schmidhuber (1991, 2006), also from an evolutionary viewpoint. With the aim of validating the same hypothesis, Soltoggio and Jones (2009) devised a simple EPANN experiment in which exploration was more advantageous than exploitation in the absence of reward learning; to do this, the reward at a particular location depleted itself if continuously visited, so that changing location at random in a T-maze became beneficial. Evolution discovered exploratory behavior before discovering reward-learning, which in turn, and surprisingly, led to an earlier evolution of reward-based learning. Counterintuitively, this experiment suggests that a stepping stone to evolve reward-based learning is to encourage reward-independent exploration.

The seminal work in Bullinaria (2003, 2007a, 2009c) proposes the more general hypothesis that learning requires the evolution
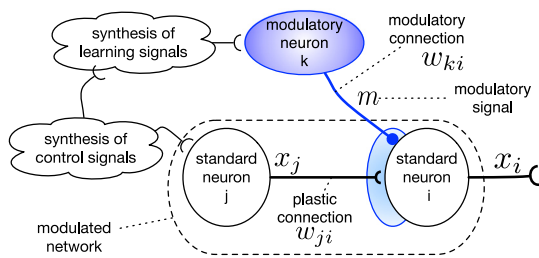
Fig. 5. A modulatory neuron gates plasticity of the synapses that connect to the postsynaptic neuron. The learning is local, but a learning signal can be created by one part of the network and used to regulate learning elsewhere.

of long periods of parental protection and late onset of maturity. Similarly, Ellefsen (2013a, b) investigates sensitive and critical periods of learning in evolved neural networks. This fascinating hypothesis has wider implications for experiments with EPANNs, and more generally for machine learning and AI. It is therefore foreseeable that future EPANNs will have a protected childhood during which parental guidance may be provided (Clutton-Brock, 1991; Eskridge & Hougen, 2012; Klug & Bonsall, 2010).

4.5. Evolving neuromodulation

Growing neuroscientific evidence on the role of neuromodulation (previously outlined in Section 2.2) inspired the design of experiments with neuromodulatory signals to evolve control behavior and learning strategies (Section 2.3). One particular case is when neuromodulation gates plasticity. Eq. (1) can be rewritten as

$$\Delta w = m \cdot f(\mathbf{x}, \boldsymbol{\theta}), \tag{4}$$

to emphasize the role of m, a modulatory signal used as a multiplicative factor that can enhance or reduce plasticity (Abbott, 1990). A network may produce many independent modulatory signals $\mathbf{m}$ targeting different neurons or areas of the network. Thus, modulation can vary in space and time. Modulation may also affect other aspects of the network dynamics, e.g., modulating activations rather than plasticity (Krichmar, 2008). Graphically, modulation can be represented as a different type of signal affecting various properties of the synaptic connections of an afferent neuron i (Fig. 5).

Evolutionary search was used to find the parameters of a neuromodulated Hebbian learning rule in a reward-based armedbandit problem in Niv, Joel, Meilijson, and Ruppin (2002). The same problem was used later in Soltoggio et al. (2007) to evolve arbitrary learning architectures with a bio-inspired gene representation method called Analog Genetic Encoding (AGE) (Mattiussi & Floreano, 2007). In that study, evolution was used to search both modulatory topologies and parameters of a particular form of Eq. (4):

$$\Delta w = m \cdot (A x_i x_j + B x_j + C x_i + D), \tag{5}$$

where the parameters A to D determined the influence of four factors in the rule: a multiplicative Hebbian term A, a presynaptic term B, a postsynaptic term C, and pure modulatory, or heterosynaptic, term D. Such a rule is not dissimilar from those presented in previous studies (see Section 3.2). However, when used in combination with modulation and a search for network topologies, evolution seems to be particularly effective at solving reward-based problems. Kondo (2007) proposed an evolutionary design and behavior analysis of neuromodulatory neural networks for mobile robot control, validating the potential of the method.

Soltoggio, Bullinaria, Mattiussi, Dürr, and Floreano (2008) tested the question of whether modulatory dynamics held an evolutionary advantage in T-maze environments with changing reward locations.[3] In their algorithm, modulatory neurons were freely inserted or deleted by random mutations, effectively allowing the evolutionary selection mechanism to autonomously pick those networks with advantageous computational components (*Aim 2*). After evolution, the best performing networks had modulatory neurons regulating learning, and evolved faster than a control evolutionary experiment that could not employ modulatory neurons. Modulatory neurons were maintained in the networks in a second phase of the experiment when genetic operators allowed for the deletion of such neurons but not for their insertion, thus demonstrating their essential function in maintaining learning in that particular experiment. In another study, Soltoggio (2008c) suggested that evolved modulatory topologies may be essential to separate the learning circuitry from the input–output controller, and shortening the input–output pathways which sped up decision processes. Soltoggio (2008b) showed that the learning dynamics are affected by tight coupling between rules and architectures in a search space with many equivalent but different control structures. Fig. 5 also suggests that modulatory networks require evolution to find two essential topological structures: what signals or combination of signals trigger modulation, and what neurons are to be targeted by modulatory signals. In other words, a balance between fixed and plastic architectures, or selective plasticity (DARPA-L2M, 2017), is an intrinsically emergent property of evolved modulated networks.

A number of further studies on the evolution of neuromodulatory dynamics confirmed the evolutionary advantages in learning scenarios (Soltoggio, 2008a). Silva, Urbano, and Christensen (2012a) used simulations of 2-wheel robots performing a dynamic concurrent foraging task, in which scattered food items periodically changed their nutritive value or became poisonous, similarly to the setup in Soltoggio and Stanley (2012). The results showed that when neuromodulation was enabled, learning evolved faster than when neuromodulation was not enabled, also with multi-robot distributed systems (Silva, Urbano, & Christensen, 2012b). Nogueira, de Brito, Vidal, and Neto (2016); Nogueira, Lenon, Eduardo, Vidal, and Cavalcante Neto (2013) also reported evolutionary advantages in foraging behavior of an autonomous virtual robot when equipped with neuromodulated plasticity. Harrington, Awa, Cussat-Blanc, and Pollack (2013) demonstrated how evolved neuromodulation applied to a gene regulatory network consistently generalized better than agents trained with fixed parameter settings. Interestingly, Arnold, Suzuki, and Arita (2013b) showed that neuromodulatory architectures provided an evolutionary advantage also in reinforcement-free environments, validating the hypothesis that plastic modulated networks have higher evolvability in a large variety of tasks. The evolution of social representations in neural networks was shown to be facilitated by neuromodulatory dynamics in Arnold, Suzuki, and Arita (2013a). An artificial life simulation environment called Polyworld (Yoder & Yaeger, 2014) helped to assess the advantage of neuromodulated plasticity in various scenarios. The authors found that neuromodulation may be able to enhance or diminish foraging performance in a competitive, dynamic environment.

Neuromodulation was evolved in Ellefsen, Mouret, and Clune (2015) in combination with modularity to address the problem of catastrophic forgetting. In Gustafsson (2016), networks evolved with AGE (Mattiussi & Floreano, 2007) for video game playing were shown to perform better with the addition of neuromodulation. Norouzzadeh and Clune (2016) showed that neuromodulation produced forward models that could adapt to changes

---

[3] In reward-based T-Maze environments, it is often assumed that the fitness function is the sum or all rewards collected during a lifetime.

significantly better than the controls. They verified that evolution exploited variable learning rates to perform adaptation when needed. In Velez and Clune (2017), diffusion-based modulation, i.e., targeting entire parts of the network, evolved to produce task-specific localized learning and functional modularity, thus reducing the problem of catastrophic forgetting.

The evidence in these studies suggests that neuromodulation is a key ingredient to facilitate the evolution of learning in EPANNs. They also indirectly suggest that neural systems with more than one type of signal, e.g., activation and other modulatory signals, might be beneficial in the neuroevolution of learning.

4.6. Evolving indirectly encoded plasticity

An indirect genotype to phenotype mapping means that evolution operates on a compact genotypical representation (analogous to the DNA) that is then mapped into a fully fledged network (analogous to a biological brain). Learning rules may undergo a similar indirect mapping, so that compact instructions in the genome expand to fully fledged plasticity rules in the phenotype. One early study (Gruau & Whitley, 1993) encoded plasticity and development with a grammar tree, and compared different learning rules on a simple static task (parity and symmetry), demonstrating that learning provided an evolutionary advantage in a static scenario. In non-static contexts, and using a T-Maze domain as learning task, Risi and Stanley (2010) showed that HyperNEAT, which usually implements a compact encoding of weight patterns for large-scale ANNs (Fig. 6a), can also encode patterns of local learning rules. The approach, called *adaptive HyperNEAT*, can encode arbitrary learning rules for each connection in an evolving ANN based on a function of the ANN's geometry (Fig. 6b). Further flexibility was added in Risi and Stanley (2012) to simultaneously encode the density and placement of nodes in substrate space. The approach, called *adaptive evolvable-substrate HyperNEAT*, makes it possible to indirectly encode plastic ANNs with thousands of connections that exhibit regularities and repeating motifs. Adaptive ES-HyperNEAT allows each individual synaptic connection, rather than neuron, to be standard or modulatory, thus introducing further design flexibility. Risi and Stanley (2014) showed how adaptive HyperNEAT can be seeded to produce a specific lateral connectivity pattern, thereby allowing the weights to self-organize to form a topographic map of the input space. The study shows that evolution can be seeded with specific plasticity mechanisms that can facilitate the evolution of specific types of learning.

The effect of indirectly encoded plasticity on the learning and on the evolutionary process was investigated by Tonelli and Mouret (2011, 2013). Using an operant conditioning task, i.e., learning by reward, the authors showed that indirect encodings that produced more regular neural structures also improved the general EPANN learning abilities when compared to direct encodings. In an approach similar to adaptive HyperNEAT, Orchard and Wang (2016) encoded the learning rule itself as an evolving network. They named the approach *neural weights and bias update* (NWB), and observed that increasing the search space of the possible plasticity rules created more general solutions than those based on only Hebbian learning.

5. Future directions

The progress of EPANNs reviewed so far is based on rapidly developing theories and technologies. In particular, new advances in AI, machine learning, neural networks and increased computational resources are currently creating a new fertile research landscape, and are setting the groundwork for new directions for EPANNs. This section presents promising research themes that have the potential to extend and radically change the field of EPANNs and AI as a whole.

5.1. Levels of abstraction and representations

Choosing the right level of abstraction and the right representation (Bengio, Courville, & Vincent, 2013) are themes at the heart of many problems in AI. In ANNs, low levels of abstraction are more computationally expensive, but might be richer in dynamics. High levels are faster to simulate, but require an intuition of the essential dynamics that are necessary in the model. Research in EPANNs is well placed to address the problem of levels of abstraction because it can reveal evolutionary advantages for different components, structures and representations.

Similarly to abstractions, representations play a critical role. Compositional Patterns Producing Networks (CPPNs) (Stanley, 2007), and also the previous work of Sims (1991), demonstrated that structured phenotypes can be generated through a function without going through the dynamic developmental process typical of multicellular organisms. Relatedly, Hornby et al. (2002) showed that the different phenotypical representations led to considerably different results in the evolution of regular structures with patterns and repetitions. Miller (2014) discussed explicitly the effect of abstraction levels for evolved developmental learning networks, in particular in relation to two approaches that model development at the neuron level or at the network level.

Finding appropriate abstractions and representations, just as it was fundamental in the advances in deep learning to represent input spaces and hierarchical features (Bengio et al., 2013; Oquab, Bottou, Laptev, & Sivic, 2014), can also extend to representations of internal models, learning mechanisms, and genetic encodings, affecting the algorithms' capabilities of evolving learning abilities.

5.2. Evolving general learning

One challenge in the evolution of learning is that evolved learning may simply result in a switch among a finite set of evolved behaviors, e.g., turning left or right in a T-Maze in a finite sequence, which is all that evolving solutions encounter during their lifetime. A challenge for EPANNs is to acquire general learning abilities in which the network is capable of learning problems not encountered during evolution. Mouret and Tonelli (2014) propose the distinction between the evolution of behavioral switches and the evolution of synaptic general learning abilities, and suggest conditions that favor these types of learning. General learning can be intuitively understood as the capability to learn any association among input, internal, and output patterns, both in the spatial and temporal dimensions, regardless of the complexity of the problem. Such an objective clearly poses practical and philosophical challenges. Although humans are considered better at general learning than machines, human learning skills are also specific and not unlimited (Ormrod & Davis, 2004). Nevertheless, moving from behavior switches to more general learning is a desirable feature for EPANNs. Encouraging the emergence of general learners may likely involve (1) an increased computational cost for testing in rich environments that include a large variety of uncertain and stochastic scenarios with problems of various complexity, and (2) an increased search space to explore the evolution of complex strategies and avoid deception.

5.3. Incremental and social learning

An important open challenge for machine learning in general is the creation of neural systems that can continuously integrate new knowledge and skills without forgetting what they previously learned (Parisi, Kemker, Part, Kanan, & Wermter, 2018), thus solving the stability-plasticity dilemma. A promising approach is progressive neural networks (Rusu et al., 2016), in which a new network is created for each new task, and lateral connections
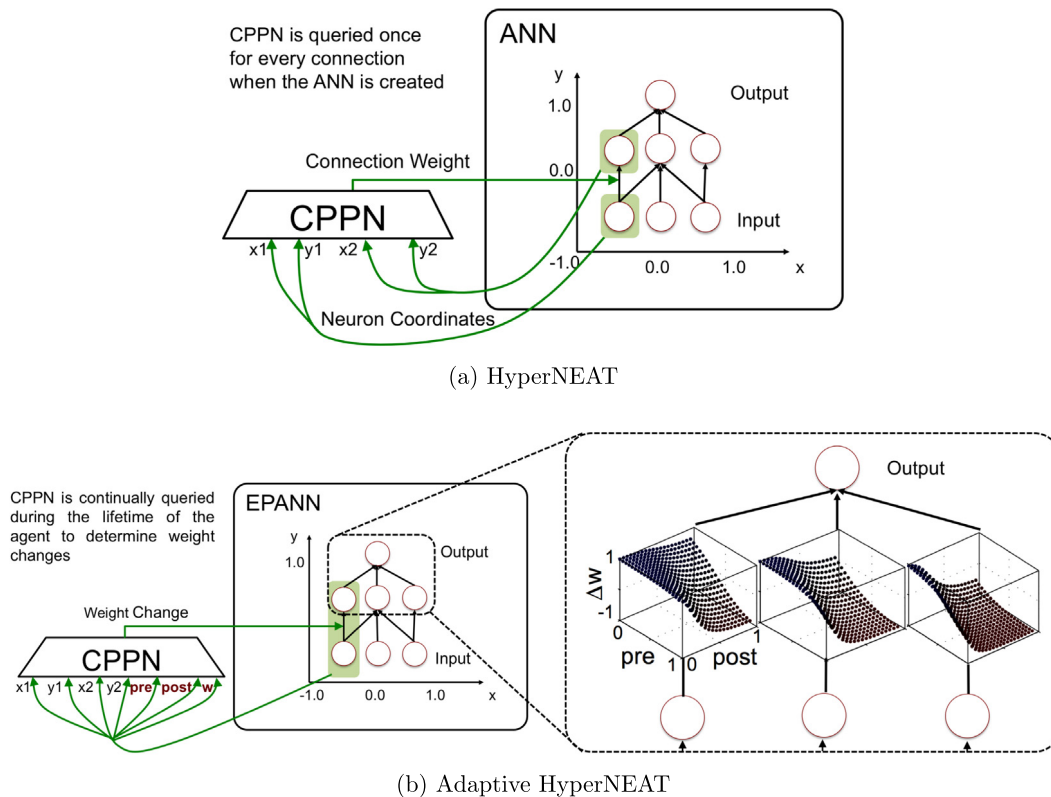
(a) HyperNEAT

(b) Adaptive HyperNEAT

Fig. 6. Example of an indirect mapping of plasticity rules from a compact genotype to a larger phenotype. (a) ANN nodes in HyperNEAT are situated in physical space by assigning them specific coordinates. The connections between nodes are determined by an evolved Compositional Patterns Producing Network (CPPN; Stanley (2007)), which takes as inputs the coordinates of two ANN neurons and returns the weight between them. In the normal HyperNEAT approach (a), the CPPN is queried once for all potential ANN connections when the agent is born. On the other hand, in adaptive HyperNEAT (b), the CPPN is continually queried during the lifetime of the agent to determine individual connection weight changes based on the location of neurons and additionally the activity of the presynaptic and postsynaptic neuron, and current connection weight. Adaptive HyperNEAT is able to indirectly encode a pattern of nonlinear learning rules for each connection in the ANN (right).

between networks allow the system to leverage previously learned features. In the presence of time delays among stimuli, actions and rewards, a rule called hypothesis testing plasticity (HTP) (Soltoggio, 2015) implements fast and slow decay to consolidate weights and suggests neural dynamics to avoid catastrophic forgetting. A method to find the best shared weights across multiple tasks, called elastic weight consolidation (EWC) was proposed in Kirkpatrick et al. (2017). Plasticity rules that implement weight consolidation, given their promise to prevent catastrophic forgetting, are likely to become standard components in EPANNs.

Encouraging modularity (Durr, Floreano, & Mattiussi, 2010; Ellefsen et al., 2015) or augmenting evolving networks with a dedicated external memory component (Lüders, Schläger, & Risi, 2016) have been proposed recently. An evolutionary advantage is likely to emerge for networks that can elaborate on previously learned sub-skills during their lifetime to learn more complex tasks.

One interesting case in which incremental learning may play a role is social learning (Best, 1999). EPANNs may learn both from the environment and from other individuals, from scratch or incrementally (Offerman & Sonnemans, 1998). In an early study, McQuesten and Miikkulainen (1997) showed that neuroevolution can benefit from parent networks teaching their offspring through backpropagation. When social, learning may involve imitation, language or communication, or other social behaviors. Bullinaria (2017) proposes an EPANN framework to simulate the evolution of culture and social learning. It is reasonable to assume that future AI learning systems, whether based on EPANNs or not, will acquire knowledge through different modalities. These will involve direct experience with the environment, but also social interaction, and possibly complex incremental learning phases.

5.4. Fast learning

Animal learning does not always require a myriad of trials. Humans can very quickly generalize from only a few given examples, possibly leveraging previous experiences and a long learning process during infancy. This type of learning, advocated in AI and robotics systems (Thrun & Mitchell, 1995), is currently still missing in EPANNs. Inspiration for new approaches could come from complementary learning systems (Kumaran, Hassabis, & McClelland, 2016; McClelland, McNaughton, & O'reilly, 1995) that humans seem to possess, which include fast and slow learning components. Additionally, approaches such as probabilistic program induction seem to be able to learn concepts in one-shot at a human-level in some tasks (Lake, Salakhutdinov, & Tenenbaum, 2015). Fast learning is likely to derive not just from trial-and-error, but also from mental models that can be applied to diverse problems, similarly to transfer learning (Pan & Yang, 2010; Thrun & Mitchell, 1995; Thrun & O'Sullivan, 1996). Reusable mental models, once learned, will allow agents to make predictions and plan in new and uncertain scenarios with similarities to previously learned ones. If EPANNs can discover neural structures or learning rules that allow for generalization, an evolutionary advantage of such a discovery will lead to its full emergence and further optimization of such a property.

A rather different approach to accelerate learning was proposed in de Vladar and Szathmáry (2015), Fernando, Karishma, and Szathmáry (2008) and called Evolutionary Neurodynamics. According to this theory, replication and selection might happen in a neural system as it learns, thus mimicking an evolutionary dynamics at the much faster time scale of a lifetime. We refer to de Vladar and Szathmáry (2015) and Fernando, Szathmáry, and Husbands

(2012) for an overview of the field. The appeal of this method is that evolutionary search can be accelerated by implementing its dynamics at both the evolution's and life's time scales.

5.5. Evolving memory

The consequence of learning is memory, both explicit and implicit (Anderson, 2013), and its consolidation (Dudai, 2012). For a review of computational models of memory see Fusi (2017). EPANNs may reach solutions in which memory evolved in different fashions, e.g., preserved as self-sustained neural activity, encoded by connection weights modified by plasticity rules, stored with an external memory (e.g. Neural Turing Machine), or a combination of these approaches. Recurrent neural architectures based on *long short-term memory* (LSTM) allow very complex tasks to be solved through gradient descent training (Greff, Srivastava, Koutník, Steunebrink, & Schmidhuber, 2015; Hochreiter & Schmidhuber, 1997) and have recently shown promise when combined with evolution (Rawal & Miikkulainen, 2016). Neuromodulation and weight consolidation could also be used to target areas of the network where information is stored.

Graves, Wayne, and Danihelka (2014) introduced the *Neural Turing Machine* (NTM), networks augmented with an external memory that allows long-term memory storage. NTMs have shown promise when trained through evolution (Greve, Jacobsen, & Risi, 2016; Lüders, Schläger, Korach, & Risi, 2017; Lüders et al., 2016) or gradient descent (Graves et al., 2014, 2016). The *Evolvable Neural Turing Machine* (ENTM) showed good performance in solving the continuous version of the double T-Maze navigation task (Greve et al., 2016), and avoided catastrophic forgetting in a continual learning domain (Lüders et al., 2017, 2016) because memory and control are separated by design. Research in this area will reveal which computational systems are more evolvable and how memories will self organize and form in EPANNs.

5.6. EPANNs and deep learning

Deep learning has shown remarkable results in a variety of different fields (Krizhevsky, Sutskever, & Hinton, 2012; LeCun et al., 2015; Schmidhuber, 2015). However, the model structures of these networks are mostly hand-designed, include a large number of parameters, and require extensive experiments to discover optimal configurations. With increased computational resources, it is now possible to search design aspects with evolution, and set up EPANN experiments with the aim of optimizing learning (*Aim 1.2*).

Koutník, Schmidhuber, and Gomez (2014) used evolution to design a controller that combined evolved recurrent neural networks, for the control part, and a deep max-pooling convolutional neural network to reduce the input dimensionality. The study does not use evolution on the deep preprocessing networks itself, but demonstrates nevertheless the evolutionary design of a deep neural controller. Young, Rose, Karnowski, Lim, and Patton (2015) used an evolutionary algorithm to optimize two parameters of a deep network: the size (range [1,8]) and the number (range [16,126]) of the filters in a convolutional neural network, showing that the optimized parameters could vary considerably from the standard best-practice values. An established evolutionary computation technique, the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) (Hansen & Ostermeier, 2001), was used in Loshchilov and Hutter (2016) to optimize the parameters of a deep network to learn to classify the MNIST dataset. The authors reported performance close to the state-of-the-art using 30 GPU devices.

Real et al. (2017) and Miikkulainen et al. (2017) showed that evolutionary search can be used to determine the topology, hyperparameters and building blocks of deep networks trained through gradient descent. The performance were shown to rival those of hand-designed architectures in the CIFAR-10 classification task and a language modeling task (Miikkulainen et al., 2017), while Real et al. (2017) also tested the method on the larger CIFAR-100 dataset. Desell (2017) proposes a method called evolutionary exploration of augmenting convolutional topologies, inspired by NEAT (Stanley & Miikkulainen, 2002), which evolves progressively more complex unstructured convolutional neural networks using genetically specified feature maps and filters. This approach is also able to co-evolve neural network training hyperparameters. Results were obtained using 5500 volunteered computers at the Citizen Science Grid who were able to evolve competitive results on the MNIST dataset in under 12,500 trained networks in a period of approximately ten days. Liu, Simonyan, Vinyals, Fernando, and Kavukcuoglu (2017) used evolution to search for hierarchical architecture representations showing competitive performance on the CIFAR-10 and Imagenet databases. The Evolutionary DEep Networks (EDEN) framework (Dufourq & Bassett, 2017) aims to generalize deep network optimization to a variety of problems and is interfaced with TensorFlow (Abadi et al., 2016). A number of similar software frameworks are currently being developed.

Fernando et al. (2017) used evolution to determine a subset of pathways through a network that are trained through backpropagation, allowing the same network to learn a variety of different tasks. Fernando et al. (2016) were also able to rediscover convolutional networks by means of evolution of Differentiable Pattern Producing Networks (Stanley, 2007).

So far, EPANN experiments in deep learning have focused primarily on the optimization of learning (*Aim 1.2*) in supervised classification tasks, e.g. optimizing final classification accuracy. In the future, evolutionary search may be used with deep networks to evolve learning from scratch, recover performance, or combining different learning rules and dynamics in an innovative and counter-intuitive fashion (*Aims 1.1, 1.3 or 2* respectively).

5.7. GPU implementations and neuromorphic hardware

The progress of EPANNs will crucially depend on implementations that take advantage of the increased computational power of parallel computation with GPUs and neuromorphic hardware (Jo et al., 2010; Monroe, 2014). Deep learning greatly benefited from GPU-accelerated machine learning but also standardized tools (e.g. Torch, Tensorflow, Theano, etc.) that made it easy for anybody to download, experiment, and extend promising deep learning models.

EPANNs have shown promise with hardware implementations. Howard, Bull, de Lacy Costello, Gale, and Adamatzky (2014) and Howard, Gale, Bull, de Lacy Costello, and Adamatzky (2011, 2012) devised experiments to evolve plastic spiking networks implemented as memristors for simulated robotic navigation tasks. Memristive plasticity was observed consistently to enable higher performance than constant-weighted connections in both static and dynamic reward scenarios. Carlson, Nageswaran, Dutt, and Krichmar (2014) used GPU implementations to evolve plastic spiking neural networks with an evolution strategy, which resulted in an efficient and automated parameter tuning framework.

In the context of newly emerging technologies, it is worth noting that, just as GPUs were not developed initially for deep learning, so novel neural computation tools and hardware systems, not developed for EPANNs, can now be exploited to enable more advanced EPANN setups.

5.8. Measuring progress

The number of platforms and environments for testing the capabilities of intelligent systems is constantly growing, e.g., the Atari or General Video Game Playing Benchmark (GVGAI, 2017),

the Project Malmo (Microsoft, 2017), or the OpenAI Universe (OpenAI, 2017). Because EPANNs are often evolved in reward-based, survival, or novelty-oriented environments to discover new, unknown, or creative learning strategies or behaviors, measuring progress is not straightforward. Desired behaviors or errors are not always defined. Moreover, the goal for EPANNs is often not to be good at solving one particular task, but rather to test the capability to evolve the learning required for a range of problems, to generalize to new problems, or to recover performance after a change in the environment. Therefore, EPANNs will require the community to devise and accept new metrics based on one or more objectives such as the following:

- the time (in the evolutionary scale) to evolve the learning mechanisms in one or more scenarios;
- the time (in the lifetime scale) for learning in one or more scenarios;
- the number of different tasks that an EPANN evolves to solve;
- a measure of the variety of skills acquired by one EPANN;
- the complexity of the tasks and/or datasets, e.g., variations in distributions, stochasticity, etc.;
- the robustness and generalization capabilities of the learner;
- the recovery time in front of high-level variations or changes, e.g., data distribution, type of problem, stochasticity levels, etc.;
- computational resources used, e.g., number of lifetime evaluations, length of a lifetime;
- size, complexity, and computational requirements of the solution once deployed;
- novelty or richness of the behavior repertoire from multiple solutions, e.g., the variety of different EPANNs and their strategies that were designed during evolution.

Few of those metrics are currently used to benchmark machine learning algorithms. Research in EPANNs will foster the adoption of such criteria as wider performance metrics for assessing lifelong learning capabilities (DARPA-L2M, 2017; Thrun & Pratt, 2012) of evolved plastic networks.

6. Conclusion

The broad inspiration and aspirations of evolved artificial plastic neural networks (EPANNs) strongly motivate this field, drawing from large, diverse, and interdisciplinary areas. In particular, the aspirations reveal ambitious and long-term research objectives related to the discovery of neural learning, with important implications for artificial intelligence and biology.

EPANNs saw considerable progress in the last two decades, primarily pointing to the potential of the autonomous evolution and discovery of neural learning. We now have: (i) advanced evolutionary algorithms to promote the evolution of learning, (ii) a better understanding of the interaction dynamics between evolution and learning, (iii) assessed advantages of multi-signal networks such as modulatory networks, and (iv) explored evolutionary representations of learning mechanisms.

Recent scientific and technical progress has set the foundation for a potential step change in EPANNs. Concurrently with the increase of computational power and a resurgence of neural computation, the need for more flexible algorithms and the opportunity to explore new design principles could make EPANNs the next AI tool capable of discovering new principles and systems for general adaptation and intelligent systems.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., & Citro, C. et al., (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467.

Abbott, L. F. (1990). Modulation of function and gated learning in a network memory. *Proceedings of the National Academy of Science of the United States of America*, 87(23), 9241–9245.

Abraham, A. (2004). Meta learning evolutionary artificial neural networks. *Neurocomputing*, 56, 1–38.

Abraham, W. C., & Robins, A. (2005). Memory retention–the synaptic stability versus plasticity dilemma. *Trends in Neurosciences*, 28(2), 73–78.

Alexander, W. H., & Sporns, O. (2002). An embodied model of learning, plasticity, and reward. *Adaptive Behavior*, 10(3–4), 143–159.

Allis, L. V., et al. (1994). *Searching for solutions in games and artificial intelligence*. Ponsen & Looijen.

Alpaydin, E. (2014). *Introduction to machine learning*. MIT press.

Anderson, J. R. (2013). *The architecture of cognition*. Psychology Press.

Angeline, P. J., Saunders, G. M., & Pollack, J. B. (1994). An evolutionary algorithm that constructs recurrent neural networks. *IEEE Transactions on Neural Networks*, 5(1), 54–65.

Arifovic, J., & Gencay, R. (2001). Using genetic algorithms to select architecture of a feedforward artificial neural network. *Physica A: Statistical Mechanics and its Applications*, 289(3), 574–594.

Arnold, S., Suzuki, R., & Arita, T. (2013). Evolution of social representation in neural networks. In *Advances in artificial life, ECAL 2013: proceedings of the twelfth European conference on the synthesis and simulation of living systems* (pp. 425–430). Cambridge, MA: MIT Press.

Arnold, S., Suzuki, R., & Arita, T. (2013). Selection for reinforcement-free learning ability as an organizing factor in the evolution of cognition. *Advances in Artificial Intelligence*, 2013, 8.

Ay, N., Flack, J., & Krakauer, D. C. (2007). Robustness and complexity co-constructed in multimodal signalling networks. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 362(1479), 441–447.

Bäck, T., & Schwefel, H. P. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, 1(1), 1–23.

Bahroun, Y., Hunsicker, E., & Soltoggio, A. (2017). Building efficient deep hebbian networks for image classification tasks. In *International conference on artificial neural networks*.

Bailey, C. H., Giustetto, M., Huang, Y. Y., Hawkins, R. D., & Kandel, E. R. (2000). Is heterosynaptic modulation essential for stabilizing Hebbian plasticity and memory? *Nature Reviews Neuroscience*, 1(1), 11–20.

Baldwin, J. M. (1896). A new factor in evolution (continued). *American Naturalist*, 536–553.

Baxter, J. (1992). The evolution of learning algorithms for artificial neural networks. In D. Green, & T. Bossomaier (Eds.), *Complex systems* (pp. 313–326). Amsterdam, the Netherlands: IOS.

Baxter, D. A., Canavier, C. C., Clark, J. W., & Byrne, J. H. (1999). Computational model of the serotonergic modulation of sensory neurons in Aplysia. *Journal of Neurophysiology*, 82, 1914–2935.

Bear, M. F., Connors, B. W., & Paradiso, M. A. (2007). *Neuroscience, Vol. 2*. Lippincott Williams & Wilkins.

Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1), 91–122.

Bengio, Y., Bengio, S., & Cloutier, J. (1990). Learning a synaptic learning rule. Technical report, Université de Montréal, Département d'informatique et de recherche opérationnelle.

Bengio, S., Bengio, Y., Cloutier, J., & Gecsei, J. (1992). On the optimization of a synaptic learning rule. In *Preprints conf. optimality in artificial and biological neural networks*. Dallas: Univ. of Texas.

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.

Bentley, P. (1999). *Evolutionary design by computers*. Morgan Kaufmann.

Best, M. L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7(3–4), 289–306.

Bienenstock, L. E., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1), 32–48.

Birmingham, J. T. (2001). Increasing sensor flexibility through neuromodulation. *Biological Bulletin*, *200*, 206–210.

Blynel, J., & Floreano, D. (2002). Levels of dynamics and adaptive behavior in evolutionary neural controllers. In *Proceedings of the seventh international conference on simulation of adaptive behavior on from animals to animats* (pp. 272–281). Cambridge, MA, USA: MIT Press.

Blynel, J., & Floreano, D. (2003). Exploring the T-Maze: Evolving learning-like robot behaviors using CTRNNs. In *EvoWorkshops* (pp. 593–604).

Boers, E. J., Borst, M. V., & Sprinkhuizen-Kuyper, I. G. (1995). Evolving neural networks using the "Baldwin effect". In *Artificial neural nets and genetic algorithms* (pp. 333–336). Springer.

Bourlard, H., & Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value decomposition. *Biological Cybernetics*, *59*(4–5), 291–294.

Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapse: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, *13*, 475–511.

Bullinaria, J. A. (2001). Exploring the Baldwin effect in evolving adaptable control systems. In *Connectionist models of learning, development and evolution* (pp. 231–242). Springer.

Bullinaria, J. A. (2003). From biological models to the evolution of robot control systems. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 2145–2164.

Bullinaria, J. A. (2007). The effect of learning on life history evolution. In *Proceedings of the 9th annual conference on genetic and evolutionary computation* (pp. 222–229). ACM.

Bullinaria, J. A. (2007). Understanding the emergence of modularity in neural systems. *Cognitive Science*, *31*(4), 673–695.

Bullinaria, J. A. (2009a). Evolved dual weight neural architectures to facilitate incremental learning. In *IJCCI* (pp. 427–434).

Bullinaria, J. A. (2009). The importance of neurophysiological constraints for modelling the emergence of modularity. *Computational Modelling in Behavioural Neuroscience: Closing the Gap Between Neurophysiology and Behaviour*, *2*, 188–208.

Bullinaria, J. A. (2009). Lifetime learning as a factor in life history evolution. *Artificial Life*, *15*(4), 389–409.

Bullinaria, J. A. (2017). Imitative and direct learning as interacting factors in life history evolution. *Artificial Life*, *23*.

Butz, M. V. (2015). Learning classifier systems. In *Springer handbook of computational intelligence* (pp. 961–981). Springer.

Butz, M. V., & Kutter, E. F. (2016). *How the mind comes into being: Introducing cognitive science from a functional and computational perspective*. Oxford University Press.

Cabessa, J., & Siegelmann, H. T. (2014). The super-Turing computational power of plastic recurrent neural networks. *International Journal of Neural Systems*, *24*(08), 1450029.

Carew, T. J., Walters, E. T., & Kandel, E. R. (1981). Classical conditioning in a simple withdrawal reflex in Aplysia californica. *The Journal of Neuroscience*, *1*(12), 1426–1437.

Carlson, K. D., Nageswaran, J. M., Dutt, N., & Krichmar, J. L. (2014). An efficient automated parameter tuning framework for spiking neural networks. *Neuromorphic Engineering Systems and Applications*, *8*.

Carver, C. S., & Scheier, M. F. (2012). *Attention and self-regulation: A control-theory approach to human behavior*. Springer Science & Business Media.

Cervier, D. (1993). *AI: The tumultuous search for artificial intelligence*. New York: Basic Books.

Chalmers, D. J. (1990). The evolution of learning: An experiment in genetic connectionism. In *Proceedings of the 1990 connectionist models summer school* (pp. 81–90), San Mateo, CA.

Chklovskii, D. B., Mel, B., & Svoboda, K. (2004). Cortical rewiring and information storage. *Nature*, *431*(7010), 782–788.

Clark, G. A., & Kandel, E. R. (1984). Branch-specific heterosynaptic facilitation in Aplysia siphon sensory cells. *Proceedings of the National Academy of Sciences of the United States of America*, *81*(8), 2577–2581.

Cliff, D., Husbands, P., & Harvey, I. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, *2*(1), 73–110.

Clutton-Brock, T. H. (1991). *The evolution of parental care*. Princeton University Press.

Coleman, O. J., & Blair, A. D. (2012). Evolving plastic neural networks for online learning: review and future directions. In *AI 2012: Advances in artificial intelligence* (pp. 326–337). Springer.

Cooper, S. J., & Donald, O. (2005). Hebb's synapse and learning rule: a history and commentary. *Neuroscience & Biobehavioral Reviews*, *28*(8), 851–874.

Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. Houghton Mifflin Harcourt.

DARPA-L2M. (2017). DARPA, Lifelong Learning Machines. https://www.fbo.gov/spg/ODA/DARPA/CMO/HR001117S0016/listing.html (accessed July 2017).

Darwin, C. (1859). *On the origin of species by means of natural selection, or The preservation of favoured races in the struggle for life*. London: Murray.

Dawkins, R. (2003). The evolution of evolvability. *On Growth, Form and Computers*, 239–255.

De Botton, A. (2016). Six areas that artificial emotional intelligence will revolutionise. http://www.wired.co.uk/article/wired-world-2016-alain-de-botton-artificial-emotional-intelligence (Accessed: 12/11/16).

Deary, I. J., Johnson, W., & Houlihan, L. M. (2009). Genetic foundations of human intelligence. *Human Genetics*, *126*(1), 215–232.

Deng, L., Hinton, G., & Kingsbury, B. (2013). New types of deep neural network learning for speech recognition and related applications: An overview. In *2013 IEEE international conference on acoustics, speech and signal processing* (pp. 8599–8603). IEEE.

Desell, T. (2017). Large Scale Evolution of Convolutional Neural Networks Using Volunteer Computing. arXiv preprint arXiv:1703.05422.

Di Paolo, E. (2002). Spike-timing dependent plasticity for evolved robots. *Adaptive Behavior*, *10*(3–4), 243–263.

Di Paolo, E. A. (2003). Evolving spike-timing-depending plasticity for single-trial learning in robots. *Philosophical Transactions of the Royal Society*, *361*(1811), 2299–2319.

Dobzhansky, T. (1970). *Genetics of the evolutionary process, Vol. 139*. Columbia University Press.

Doidge, N. (2007). *The brain that changes itself: Stories of personal triumph from the frontiers of brain science*. Penguin.

Downing, K. L. (2007). Supplementing evolutionary development systems with abstract models of neurogenesis. In Dirk Thierens, et al. (Eds.), *Proceeding of the genetic and evolutionary computation conference* (pp. 990–998). ACM.

Downing, K. L. (2015). *Intelligence emerging: Adaptivity and search in evolving neural systems*. MIT Press.

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

Doya, K., & Uchibe, E. (2005). The Cyber Rodent Project: Exploration and adaptive mechanisms for self-preservation and self-reproduction. *Adaptive Behavior*, *13*(2), 149–160.

Draganski, B., & May, A. (2008). Training-induced structural changes in the adult human brain. *Behavioural Brain Research*, *192*(1), 137–142.

Dudai, Y. (2012). The restless engram: consolidations never end. *Annual Review of Neuroscience*, *35*, 227–247.

Dufourq, E., & Bassett, B. A. (2017). EDEN: Evolutionary Deep Networks for Efficient Machine Learning. arXiv preprint arXiv:1709.09161.

Durr, P., Floreano, D., & Mattiussi, C. (2010). Genetic representation and evolvability of modular neural controllers. *IEEE Computational Intelligence Magazine*, *5*(3), 10–19.

Edelman, G. M., & Tononi, G. (2000). *A universe of consciousness: How matter becomes imagination*. Basic books.

Eiben, A. E., & Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature*, *521*(7553), 476–482.

Ellefsen, K. O. (2013). Different regulation mechanisms for evolved critical periods in learning. In *2013 IEEE congress on evolutionary computation* (pp. 1193–1200). IEEE.

Ellefsen, K. O. (2013b). Evolved sensitive periods in learning. In *Advances in artificial life, ECAL 2013: Proceedings of the twelfth European confer- ence on the synthesis and simulation of living systems, Vol. 12* (pp. 409–416).

Ellefsen, K. O. (2014). The evolution of learning under environmental variability. In *Artificial life 14: Proceedings of the fourteenth international conference on the synthesis and simulation of living systems*.

Ellefsen, K. O., Mouret, J. B., & Clune, J. (2015). Neural modularity helps organisms evolve to learn new skills without forgetting old skills. *PLoS Computational Biology*, *11*(4), e1004128.

Eskridge, B. E., & Hougen, D. F. (2012). Nurturing promotes the evolution of learning in uncertain environments. In *2012 IEEE international conference on development and learning and epigenetic robotics* (pp. 1–6). IEEE.

Federici, D. (2005). Evolving developing spiking neural networks. In *Proceedings of the IEEE congress on evolutionary computation*.

Fellous, J. M., & Linster, C. (1998). Computational models of neuromodulation. *Neural Computation*, *10*, 771–805.

Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., & Rusu, A. A. et al., (2017). PathNet: Evolution Channels Gradient Descent in Super Neural Networks. arXiv preprint arXiv:1701.08734.

Fernando, C., Banarse, D., Reynolds, M., Besse, F., Pfau, D., & Jaderberg, M. et al., (2016). Convolution by Evolution: Differentiable Pattern Producing Networks. arXiv preprint arXiv:1606.02580.

Fernando, C., Karishma, K., & Szathmáry, E. (2008). Copying and evolution of neuronal topology. *PLoS One*, *3*(11), e3775.

Fernando, C., Szathmáry, E., & Husbands, P. (2012). Selectionist and evolutionary approaches to brain function: a critical appraisal. *Frontiers in Computational Neuroscience*, *6*.

Finnie, P. S., & Nader, K. (2012). The role of metaplasticity mechanisms in regulating memory destabilization and reconsolidation. *Neuroscience & Biobehavioral Reviews*, *36*(7), 1667–1707.

Floreano, D., & Mattiussi, C. (2008). *Bio-inspired artificial intelligence: theories, methods, and technologies*. MIT press.

Floreano, D., & Mondada, F. (1994). Automatic creation of an autonomous agent: genetic evolution of a neural-network driven robot. In *Proceedings of the third international conference on simulation of adaptive behavior : from animals to animats 3* (pp. 421–430). Cambridge, MA, USA: MIT Press.

Floreano, D., & Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics–Part B*, *26*(3), 396–407.

Floreano, D., & Nolfi, S. (2004). *Evolutionary robotics*. The MIT Press.

Floreano, D., & Urzelai, J. (2001). Evolution of plastic control networks. *Autonomous Robots*, *11*(3), 311–317.

Floreano, D., & Urzelai, J. (2001). Neural morphogenesis, synaptic plasticity, and evolution. *Theory in Biosciences*, *240*, 225–240.

Fogel, D. B. (2006). *Evolutionary computation: toward a new philosophy of machine intelligence, Vol. 1*. John Wiley & Sons.

Fontanari, J., & Meir, R. (1991). Evolving a learning algorithm for the binary perceptron. *Network. Computation in Neural Systems*, *2*(4), 353–359.

Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, *13*(7), 293–301.

Fujii, A., Saito, N., Nakahira, K., & Ishiguro, A. (2002). Generation of an adaptive controller cpg for a quadruped robot with neuromodulation mechanism. In *Proceedings of the 2002 IEEE/RSJ international conference on intelligent robots and systems* (pp. 2619–2624). Lausanne, Switzerland: EPFL.

Funahashi, K. i., & Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, *6*(6), 801–806.

Fusi, S. (2017). Computational models of long term plasticity and memory. arXiv preprint arXiv:1706.04946.

Gerstner, W., & Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biological Cybernetics*, *87*(5–6), 404–415.

Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.

Goldberg, D. E., & Holland, J. H. (1988). Genetic algorithms and machine learning. *Machine Learning*, *3*(2), 95–99.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672–2680).

Graves, A., Wayne, G., & Danihelka, I. (2014). Neural Turing machines. arXiv preprint arXiv:1410.5401.

Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., et al. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, *538*(7626), 471–476.

Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., & Schmidhuber, J. (2015). LSTM: A search space odyssey. arXiv preprint arXiv:1503.04069.

Greve, R. B., Jacobsen, E. J., & Risi, S. (2016). Evolving neural turing machines for reward-based learning. In *Proceedings of the genetic and evolutionary computation conference 2016*. ISBN: 978-1-4503-4206-3, (pp. 117–124). New York, NY, USA: ACM.

Grossberg, S. (2012). *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control, vol. 70*. Springer Science & Business Media.

Gruau, F., & Whitley, D. (1993). Adding learning to the cellular development of neural networks: Evolution and the Baldwin effect. *Evolutionary Computation*, *1*(3), 213–233.

Gu, Q. (2002). Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. *Neuroscience*, *111*(4), 815–853.

Gustafsson, J. (2016). *Evolving Neuromodulatory Topologies for Plasticity in Video Game Playing* (Master's thesis), Blekinge Institute of Technology, Faculty of Computing, Department of Software Engineering.

GVGAI. The General Video Game AI Competition. Accessed: 21/01/17, http://www.gvgai.net.

Hansen, N., & Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, *9*(2), 159–195.

Happel, B. L., & Murre, J. M. (1994). Design and evolution of modular neural network architectures. *Neural Networks*, *7*(6), 985–1004.

Harrington, K. I., Awa, E., Cussat-Blanc, S., & Pollack, J. (2013). Robot coverage control by evolved neuromodulation. In *Neural networks (IJCNN), The 2013 international joint conference on* (pp. 1–8). IEEE.

Harris-Warrick, R. M., & Marder, E. (1991). Modulation of neural networks for behavior. *Annual Review of Neuroscience*, *14*, 39–57.

Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1997). Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, *20*(2), 205–224.

Hasselmo, M. E. (1995). Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behavioural Brain Research*, *67*, 1–27.

Hasselmo, M. E. (1999). Neuromodulation: acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, *3*, 351–359.

Hasselmo, M. E., & Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: computational modeling and brain slice physiology. *Journal of Neuroscience*, *14*(6), 3898–3914.

Hawkins, J., & Blakeslee, S. (2007). *On intelligence*. Macmillan.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Wiley, New York.

Hensch, T. K., Fagiolini, M., Mataga, N., Stryker, M. P., Baekkeskov, S., & Kash, S. F. (1998). Local GABA circuit control of experience-dependent plasticity in developing visual cortex. *Science*, *282*(5393), 1504–1508.

Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495–502.

Hinton, G. E., & Plaut, D. C. (1987). Using fast weights to deblur old memories. In *Proceedings of the ninth annual conference of the cognitive science society* (pp. 177–186).

Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, *313*(5786), 504–507.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780.

Hoinville, T., Tapia Siles, C., & Hénaff, P. (2011). Flexible and multistable pattern generation by evolving constrained plastic neurocontrollers. *Adaptive Behavior*, *19*(3), 187–207.

Holland, J. H. (1975). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press.

Holland, J. H., & Reitman, J. S. (1977). Cognitive systems based on adaptive algorithms. *ACM SIGART Bulletin*, *63* 49–49.

Hopkins, W. D., Russell, J. L., & Schaeffer, J. (2014). Chimpanzee intelligence is heritable. *Current Biology*, *24*(14), 1649–1652.

Hornby, G., Pollack, J., et al. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, *8*(3), 223–246.

Howard, G., Bull, L., de Lacy Costello, B., Gale, E., & Adamatzky, A. (2014). Evolving spiking networks with variable resistive memories. *Evolutionary Computation*, *22*(1), 79–103.

Howard, G., Gale, E., Bull, L., de Lacy Costello, B., & Adamatzky, A. (2011). Towards evolving spiking networks with memristive synapses. In *Artificial life (ALIFE), 2011 IEEE symposium on* (pp. 14–21). IEEE.

Howard, G., Gale, E., Bull, L., de Lacy Costello, B., & Adamatzky, A. (2012). Evolution of plastic learning in spiking networks via memristive connections. *IEEE Transactions on Evolutionary Computation*, *16*(5), 711–729.

Hull, C. L. (1943). *Principles of behavior*. New-Your: Appleton Century.

Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving gasnet for robots control. *Connection Science*, *10*, 185–210.

Izhikevich, E. M. (2006). Polychronization: computation with spikes. *Neural Computation*, *18*(2), 245–282.

Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, *17*(10), 2443–2452.

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*, *304*(5667), 78–80.

Jo, S. H., Chang, T., Ebong, I., Bhadviya, B. B., Mazumder, P., et al. (2010). Nanoscale memristor device as synapse in neuromorphic systems. *Nano letters*, *10*(4), 1297–1301.

Kaas, J. (2001). Neural Plasticity. In N. J. Smelser, & P. B. Baltes (Eds.), *International Encyclopedia of the Social & Behavioral Sciences*. ISBN: 978-0-08-043076-8, (pp. 10542–10546). Oxford: Pergamon. http://dx.doi.org/10.1016/B0-08-043076-7/03619-6. URL https://www.sciencedirect.com/science/article/pii/B0080430767036196.

Kandel, E. R. (2007). *In search of memory - the emergence of a new science of mind*. New York: WW Norton & Company.

Kandel, E. R., & Tauc, L. (1965). Heterosynaptic facilitation in neurones of the abdominal ganglion of Aplysia depilans. *The Journal of Physiology*, *181*, 1–27.

Katz, P. S. (1995). Intrinsic and extrinsic neuromodulation of motor circuits. *Current Opinion in Neurobiology*, *5*, 799–808.

Katz, P. S., & Frost, W. N. (1996). Intrinsic neuromodulation: altering neuronal circuits from within. *Trends in Neurosciences*, *19*, 54–61.

Katz, P. S., Getting, P. A., & Frost, W. N. (1994). Dynamic neuromodulation of synaptic strength intrinsic to a central pattern generator circuit. *Nature*, *367*, 729–731.

Khan, G. M., & Miller, J. F. (2014). In search of intelligence: evolving a developmental neuron capable of learning. *Connection Science*, *26*(4), 297–333.

Khan, G. M., Miller, J. F., & Halliday, D. M. (2008). Coevolution of neuro-developmental programs that play checkers. In Springer (Ed.), *Proceedings of 8th international conference on evolvable systems (ICES): From biology to hardware, vol. 5216* (pp. 352–361).

Khan, G. M., Miller, J. F., & Halliday, D. M. (2011). Evolution of cartesian genetic programs for development of learning neural architecture. *Evolutionary Computation*, *19*(3), 469–523.

Khan, M. M., Khan, G. M., & Miller, J. F. (2011). Developmental plasticity in cartesian genetic programming based neural networks. In *ICINCO 2011 - Proceedings of the 8th international conference on informatics in control, automation and robotics, Volume 1, Noordwijkerhout, the Netherlands, 28 - 31 July, 2011* (pp. 449–458).

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*.

Kiyota, T. (2017). Neurogenesis and brain repair. In *Neuroimmune pharmacology* (pp. 575–597). Springer.

Klug, H., & Bonsall, M. B. (2010). Life history and the evolution of parental care. *Evolution*, *64*(3), 823–835.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*(1), 59–69.

Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, *78*(9), 1464–1480.

Kolb, B. (1989). Brain development, plasticity, and behavior.. *American Psychologist*, *44*(9), 1203.

Kolb, B., & Gibb, R. (2011). Brain plasticity and behaviour in the developing brain. *Journal of the Canadian Academy of Child & Adolescent Psychiatry*, *20*(4).

Kondo, T. (2007). Evolutionary design and behaviour analysis of neuromodulatory neural networks for mobile robots control. *Applied Soft Computing*, *7*(1), 189–202.

Køppe, S. (1983). The psychology of the neuron: Freud, Cajal and Golgi. *Scandinavian Journal of Psychology*, *24*(1), 1–12.

Koutník, J., Schmidhuber, J., & Gomez, F. (2014). Evolving deep unsupervised convolutional networks for vision-based reinforcement learning. In *Proceedings of the 2014 annual conference on genetic and evolutionary computation* (pp. 541–548). ACM.

Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection, vol. 1*. MIT press.

Krichmar, J. L. (2008). The neuromodulatory system: a framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior*, *16*(6), 385–399.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097–1105).

Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, *20*(7), 512–534.

Kupfermann, I. (1987). Cellular Neurobiology: Neuromodulation. *Science*, *236*, 863.

Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338.

Lalejini, A., & Ofria, C. (2016). The evolutionary origins of phenotypic plasticity. In *Proceedings of the artificial life conference*.

Lamprecht, R., & LeDoux, J. (2004). Structural plasticity and memory. *Nature Reviews Neuroscience*, *5*(1), 45–54.

Langton, C. G. (1997). *Artificial life: An overview*. Mit Press.

Lanzi, P. L., Stolzmann, W., & Wilson, S. W. (2003). *Learning classifier systems: from foundations to applications*. Springer.

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

LeDoux, J. E. (2003). *Synaptic self: How our brains become who we are*. Penguin.

Lee, C. M., & Narayanan, S. S. (2005). Toward detecting emotions in spoken dialogs. *IEEE Transactions on Speech and Audio Processing*, *13*(2), 293–303.

Lehman, J., & Miikkulainen, R. (2014). Overcoming deception in evolution of cognitive behaviors. In *Proceedings of the genetic and evolutionary computation conference*, Vancouver, BC, Canada.

Lehman, J., & Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (pp. 329–336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, *19*(2), 189–223.

Lehman, J., & Stanley, K. O. (2013). Evolvability is inevitable: increasing evolvability without the pressure to adapt. *PloS One*, *8*(4), e62186.

Levy, J. P., & Bairaktaris, D. (1995). Connectionist dual-weight architectures. *Language and Cognitive Processes*, *10*(3–4), 265–283.

Liu, H., Simonyan, K., Vinyals, O., Fernando, C., & Kavukcuoglu, K. (2017). Hierarchical representations for efficient architecture search. arXiv preprint arXiv:1711.00436.

Loshchilov, I., & Hutter, F. (2016). CMA-ES for Hyperparameter Optimization of Deep Neural Networks. arXiv preprint arXiv:1604.07269.

Lüders, B., Schläger, M., Korach, A., & Risi, S. (2017). Continual and one-shot learning through neural networks with dynamic external memory. In *Proceedings of the 20th european conference on the applications of evolutionary computation*.

Lüders, B., Schläger, M., & Risi, S. (2016). Continual learning through evolvable neural turing machines. In *Proceedings of the NIPS 2016 workshop on continual learning and deep networks*.

Maass, W., & Bishop, C. M. (2001). *Pulsed neural networks*. MIT press.

Maniadakis, M., & Trahanias, P. (2006). Modelling brain emergent behaviours through coevolution of neural agents. *Neural Networks*, *19*(5), 705–720.

Marder, E. (1996). Neural modulation: Following your own rhythm. *Current Biology*, *6*(2), 119–121.

Marder, E., & Thirumalai, V. (2002). Cellular, synaptic and network effects of neuromodulation. *Neural Networks*, *15*, 479–493.

Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*(5297), 213–215.

Mattiussi, C., & Floreano, D. (2007). Analog genetic encoding for the evolution of circuits and networks. *IEEE Transactions on Evolutionary Computation*, *11*(5), 596–607.

Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, *4*(3), 213–234.

McClelland, J. L., McNaughton, B. L., & 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.

McQuesten, P., & Miikkulainen, R. (1997). Culling and teaching in neuro-evolution. In *ICGA* (pp. 760–767).

Meng, Y., Jin, Y., & Yin, J. (2011). Modeling activity-dependent plasticity in BCM spiking neural networks with application to human behavior recognition. *IEEE Transactions on Neural Networks*, *22*(12), 1952–1966.

Merzenich, M. M., Nelson, R. J., Stryker, M. P., Cynader, M. S., Schoppmann, A., & Zook, J. M. (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *The Journal of Comparative Neurology*, *224*(4), 591–605.

Michalewicz, Z. (1994). GAs: What are they? In *Genetic algorithms+ data structures= evolution programs* (pp. 13–30). Springer.

Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. (2013). *Machine learning: An artificial intelligence approach*. Springer Science & Business Media.

Miconi, T. (2008). *The road to everywhere: Evolution, complexity and progress in natural and artificial systems* (Ph.D. thesis), University of Birmingham.

Miconi, T., Clune, J., & Stanley, K. O. (2018). Differentiable plasticity: training plastic neural networks with backpropagation. arXiv preprint arXiv:1804.02464.

Microsoft. 2017. Project malmo https://github.com/Microsoft/malmo, accessed Nov 2017.

Miikkulainen, R., Liang, J., Meyerson, E., Rawal, A., Fink, D., & Francon, O. et al. (2017). Evolving Deep Neural Networks. arXiv preprint arXiv:1703.00548.

Miller, J. F. (2014). Neuro-centric and holocentric approaches to the evolution of developmental neural networks. In *Growing adaptive machines* (pp. 227–249). Springer.

Miller, K. D., & Mackay, D. J. C. (1994). The role of constraints in hebbian learning. *Neural Computation*, *6*, 100–126.

Millington, I., & Funge, J. (2016). *Artificial intelligence for games*. CRC Press.

Monroe, D. (2014). Neuromorphic computing gets ready for the (really) big time. *Communications of the ACM*, *57*(6), 13–15.

Moore, D. S. (2003). The Dependent Gene: The Fallacy of "Nature vs. Nurture".

Mouret, J.-B., & Tonelli, P. (2014). Artificial evolution of plastic neural networks: a few key concepts. In *Growing adaptive machines* (pp. 251–261). Springer.

Murre, J. M. (1992). *Learning and categorization in modular neural networks*. Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc..

Niv, Y., Joel, D., Meilijson, I., & Ruppin, E. (2002). Evolution of reinforcement learning in uncertain environments: a simple explanation for complex foraging behaviours. *Adaptive Behavior*, *10*(1), 5–24.

Nogueira, Y. L. B., de Brito, C. E. F., Vidal, C. A., & Neto, J. B. C. (2016). Emergent vision system of autonomous virtual characters. *Neurocomputing*, *173*, 1851–1867.

Nogueira, B., Lenon, Y., Eduardo, C., Vidal, C. A., & Cavalcante Neto, J. B. (2013). Evolving plastic neuromodulated networks for behavior emergence of autonomous virtual characters. In *Advances in artificial life, ECAL 2013: Proceedings of the twelfth european conference on the synthesis and simulation of living systems* (pp. 577–584). Cambridge, MA: MIT Press.

Nolfi, S., & Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, *7*(1), 89–113.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics*. Mit Press Cambridge.

Nolfi, S., & Parisi, D. (1993). Auto-teaching: networks that develop their own teaching input. In *Free university of brussels*. Citeseer.

Nolfi, S., & Parisi, D. (1996). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, *5*(1), 75–98.

Norouzzadeh, M. S., & Clune, J. (2016). Neuromodulation improves the evolution of forward models. In *Proceedings of the genetic and evolutionary computation conference*.

Offerman, T., & Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of Economic Behavior & Organization*, *34*(4), 559–575.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*(3), 267–273.

OpenAI. OpenAI Universe. Accessed: 21/01/17. https://openai.com/blog/universe/.

Oquab, M., Bottou, L., Laptev, I., & Sivic, J. (2014). Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE Conference on computer vision and pattern recognition* (pp. 1717–1724).

Orchard, J., & Wang, L. (2016). The evolution of a generalized neural learning rule. In *Neural Networks (IJCNN), 2016 international joint conference on* (pp. 4688–4694). IEEE.

Ormrod, J. E., & Davis, K. M. (2004). *Human learning*. Merrill.

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *The IEEE Transactions on Knowledge and Data Engineering*, *22*(10), 1345–1359.

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2018). Continual lifelong learning with neural networks: A review. arXiv preprint arXiv:1802.07569.

Pascual-Leone, A., Amedi, A., Fregni, F., & Merabet, L. B. (2005). The plastic human brain cortex. *Annual Review of Neuroscience*, *28*, 377–401.

Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford : Oxford University Press.

Pehlevan, C., Hu, T., & Chklovskii, D. B. (2015). A Hebbian/anti-Hebbian neural network for linear subspace learning: A derivation from multidimensional scaling of streaming data. *Neural Computation*.

Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, *9*(1), 75–82.

Rauschecker, J. P., & Singer, W. (1981). The effects of early visual experience on the cat's visual cortex and their possible explanation by Hebb synapses. *Journal Physiology*, *310*, 215–239.

Rawal, A., & Miikkulainen, R. (2016). Evolving Deep LSTM-based Memory Networks Using an Information Maximization Objective. In *Proceedings of the genetic and evolutionary computation conference 2016*. ISBN: 978-1-4503-4206-3, (pp. 501–508). New York, NY, USA: ACM.

Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., & Le, Q. et al. (2017). Large-Scale Evolution of Image Classifiers. ArXiv e-prints.

Rescorla, R. A. (2014). *Pavlovian second-order conditioning (psychology revivals): Studies in associative learning*. Psychology Press.

Risi, S., Hughes, C. E., & Stanley, K. O. (2010). Evolving plastic neural networks with novelty search. *Adaptive Behavior*, *18*(6), 470–491.

Risi, S., & Stanley, K. O. (2010). Indirectly encoding neural plasticity as a pattern of local rules. In *Proceedings of the 11th international conference on simulation of adaptive behavior*. New York: Springer.

Risi, S., & Stanley, K. O. (2012). A unified approach to evolving plasticity and neural geometry. In *Neural Networks (IJCNN), the 2012 international joint conference on* (pp. 1–8). IEEE.

Risi, S., & Stanley, K. O. (2014). Guided self-organization in indirectly encoded and evolving topographic maps. In *Proceedings of the genetic and evolutionary computation conference* (p. 8). New York, NY: ACM.

Risi, S., Vanderbleek, S. D., Hughes, C. E., & Stanley, K. O. (2009). How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the 11th annual conference on genetic and evolutionary computation* (pp. 153–160). ACM.

Roberts, A., Conte, D., Hull, M., Merrison-Hort, R., al Azad, A. K., Buhl, E., et al. (2014). Can simple rules control development of a pioneer vertebrate neuronal network generating behavior? *Journal of Neuroscience*, *34*(2), 608–621.

Robins, A. (1995). Catastrophic forgetting, rehearsal, and pseudorehearsal. *Connection Science: Journal of Neural Computing, Artificial Intelligence and Cognitive Research*, *7*(123–146).

Roff, D. (1993). *Evolution of life histories: theory and analysis*. Springer Science & Business Media.

Rolls, E. T., & Stringer, S. M. (2000). On the design of neural networks in the brain by genetic evolution. *Progress in Neurobiology*, *61*(6), 557–579.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1988). Learning representations by back-propagating errors. *Cognitive Modeling*, *5*(3), 1.

Russell, S., & Norvig, P. (2013). *Artificial intelligence: a modern approach* (3rd ed.). Pearson.

Russo, S. J., Dietz, D. M., Dumitriu, D., Morrison, J. H., Malenka, R. C., & Nestler, E. J. (2010). The addicted synapse: mechanisms of synaptic and structural plasticity in nucleus accumbens. *Trends in neurosciences*, *33*(6), 267–276.

Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., & Kavukcuoglu, K. et al. (2016). Progressive Neural Networks. arXiv preprint arXiv:1606.04671.

Sanchez, E., Mange, D., Sipper, M., Tomassini, M., Pérez-Uribe, A., & Stauffer, A. (1996). Phylogeny, ontogeny, and epigenesis: Three sources of biological inspiration for softening hardware. In *International conference on evolvable systems* (pp. 33–54). Springer.

Schmidhuber, J. (1991). Curious model-building control systems. In *Neural Networks, 1991. 1991 IEEE international joint conference on* (pp. 1458–1463). IEEE.

Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, *18*(2), 173–187.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, *61*, 85–117.

Schmidhuber, J., Wierstra, D., Gagliolo, M., & Gomez, F. (2007). Training recurrent networks by evolino. *Neural Computation*, *19*(3), 757–779.

Schreiweis, C., Bornschein, U., Burguière, E., Kerimoglu, C., Schreiter, S., Dannemann, M., et al. (2014). Humanized Foxp2 accelerates learning by enhancing transitions from declarative to procedural performance. *Proceedings of the National Academy of Sciences*, *111*(39), 14253–14258.

Schroll, H., & Hamker, F. H. (2015). Computational models of basal-ganglia pathway functions: focus on functional neuroanatomy. *Frontiers in Systems Neuroscience*, 7.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *The Journal of Neuroscience*, *13*, 900–913.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate for prediction and reward. *Science*, *275*, 1593–1598.

Silva, F., Urbano, P., & Christensen, A. L. (2012). Adaptation of robot behaviour through online evolution and neuromodulated learning. In *Ibero-American conference on artificial intelligence* (pp. 300–309). Springer.

Silva, F., Urbano, P., & Christensen, A. L. (2012). Continuous adaptation of robot behaviour through online evolution and neuromodulated learning. In *Fifth international workshop on evolutionary and reinforcement learning for autonomous robot systems (ERLARS 2012) in Conjunction with the 20th European conference on artificial intelligence (ECAI 2012). Montpellier*. Citeseer.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*(7587), 484–489.

Sims, K. (1991). *Artificial evolution for computer graphics, vol. 4*. ACM.

Sims, K. (1994). Evolving 3D morphology and behavior by competition. *Artificial Life*, *1*(4), 353–372.

Sipper, M., Sanchez, E., Mange, D., Tomassini, M., Perez-Uribe, A., & Stauffer, A. (1997). A phylogenetic, ontogenetic, and epigenetic view of bio-inspired hardware system. *IEEE Transactions on Evolutionary Computation*, *1*(1), 83–97.

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York, London: D. Appleton-Century Company, Incorporated.

Skinner, B. F. (1953). *Science and human behavior*. New York: MacMillan.

Smith, J. M. (1986). When learning guides evolution. *Nature*, *329*(6142), 761–762.

Smith, T. (2002). *The evolvability of artificial neural networks for robot control* (Ph.D. thesis), University of Sussex.

Smith, T., Husbands, P., Philippides, A., & O'Shea, M. (2002). Neuronal plasticity and temporal adaptivity: GasNet robot control networks. *Adaptive Behavior*, *10*, 161–183.

Soltoggio, A. (2008). *Evolutionary and computational advantages of neuromodulated plasticity* (Ph.D. thesis), School of Computer Science, The University of Birmingham.

Soltoggio, A. (2008). Neural plasticity and minimal topologies for reward-based learning problems. In *Proceeding of the 8th international conference on hybrid intelligent systems (HIS2008), 10-12 September, Barcelona, Spain*.

Soltoggio, A. (2008). Neuromodulation increases decision speed in dynamic environments. In *Proceedings of the 8th international conference on epigenetic robotics, Southampton, July 2008*.

Soltoggio, A. (2015). Short-term plasticity as cause-effect hypothesis testing in distal reward learing. *Biological Cybernetics*, *109*(1), 75–94.

Soltoggio, A., Bullinaria, J. A., Mattiussi, C., Dürr, P., & Floreano, D. (2008). Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios. In *Artificial life XI: Proceedings of the eleventh international conference on the simulation and synthesis of living systems*. MIT Press.

Soltoggio, A., Dürr, P., Mattiussi, C., & Floreano, D. (2007). Evolving neuromodulatory topologies for reinforcement learning-like problems. In *Proceedings of the IEEE congress on evolutionary computation, CEC 2007*.

Soltoggio, A., & Jones, B. (2009). Novelty of behaviour as a basis for the neuroevolution of operant reward learning. In *Proceedings of the 11th annual conference on genetic and evolutionary computation* (pp. 169–176). ACM.

Soltoggio, A., & Stanley, K. O. (2012). From modulated hebbian plasticity to simple behavior learning through noise and weight saturation. *Neural Networks*, *34*, 28–41.

Sporns, O., & Alexander, W. H. (2003). Neuromodulation in a learning robot: interactions between neural plasticity and behavior. In *Proceedings of the international joint conference on neural networks, Vol. 4* (pp. 2789–2794).

Staddon, J. E. R. (1983). *Adaptive behaviour and learning*. Cambridge University Press.

Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, *8*(2), 131–162.

Stanley, K. O., Bryant, B. D., & Miikkulainen, R. (2003). Evolving adaptive neural networks with and without adaptive synapses. In B. Rylander (Ed.), *Genetic and evolutionary computation conference late breaking papers* (pp. 275–282), Chicago, USA.

Stanley, K. O., D'Ambrosio, D. B., & Gauci, J. (2009). A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life*, *15*(2), 185–212.

Stanley, K. O., & Lehman, J. (2015). *Why greatness cannot be planned, Vol. 10*. Springer Science Business Media, 978–3.

Stanley, K. O., Lehman, J., & Soros, L. (2017). Open-endedness: The last grand challenge youve never heard of. https://www.oreilly.com/ideas/open-endedness-the-last-grand-challenge-youve-never-heard-of, accessed April 2018.

Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, *10*(2), 99–127.

Stanley, K. O., & Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, *9*(2), 93–130.

Steels, L. (1993). The artificial life roots of artificial intelligence. *Artificial Life*, *1*(1_2), 75–110.

Stone, J. V. (2007). Distributed representations accelerate evolution of adaptive behaviours. *PLoS Computational Biology*, *3*(8), e147.

Suri, R. E. (2002). TD models of reward predictive responses in dopamine neurons. *Neural Networks*, *15*, 523–533.

Suri, R. E., Bargas, J., & Arbib, M. A. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience*, *103*(1), 65–85.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA, USA: MIT Press.

Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., et al. (2016). Open-ended evolution: perspectives from the OEE workshop in York. *Artificial Life*.

Tessier-Lavigne, M., & Goodman, C. S. (1996). The molecular biology of axon guidance. *Science*, *274*(5290), 1123.

Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.

Thrun, S., & Mitchell, T. M. (1995). Lifelong robot learning. *Robotics and Autonomous Systems*, *15*(1–2), 25–46.

Thrun, S., & O'Sullivan, J. (1996). Discovering structure in multiple learning tasks: The TC algorithm. In *ICML, Vol. 96* (pp. 489–497).

Thrun, S., & Pratt, L. (2012). *Learning to learn*. Springer Science & Business Media.

Tonelli, P., & Mouret, J.-B. (2011). Using a map-based encoding to evolve plastic neural networks. In *Evolving and adaptive intelligent systems (EAIS), 2011 IEEE workshop on* (pp. 9–16). IEEE.

Tonelli, P., & Mouret, J.-B. (2013). On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PloS One*, *8*(11), e79138.

Urzelai, J., & Floreano, D. (2000). Evolutionary robotics: Coping with environmental change. In *Genetic and evolutionary computation conference (GECCO'2000)*.

Varela, F. J., Thompson, E., & Rosch, E. (2017). *The embodied mind: Cognitive science and human experience*. MIT press.

Velez, R., & Clune, J. (2017). Diffusion-based neuromodulation can eliminate catastrophic forgetting in simple neural networks. CoRR, abs/1705.07241.

Venkatesh, S. S. (1993). Directed drift: A new linear threshold algorithm for learning binary weights on-line. *Journal of Computer and System Sciences*, *46*(2), 198–217.

Vitay, J., & Hamker, F. H. (2010). A computational model of basal ganglia and its role in memory retrieval in rewarded visual memory tasks. *Frontiers in Computational Neuroscience*.

de Vladar, H. P., & Szathmáry, E. (2015). Neuronal boost to evolutionary dynamics. *Interface Focus*, *5*(6), 20150074.

Wagner, G. P., & Altenberg, L. (1996). Perspective: complex adaptations and the evolution of evolvability. *Evolution*, *50*(3), 967–976.

Walters, E. T., & Byrne, J. H. (1983). Associative conditioning of single sensory neurons suggests a cellular mechanism for learning. *Science*, *219*, 405–408.

Widrow, B., & Hoff, M. E., et al. (1960). Adaptive switching circuits. In *IRE WESCON convention record, Vol. 4* (pp. 96–104), New York.

Widrow, B., & Lehr, M. A. (1990). 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proceedings of the IEEE*, *78*(9), 1415–1442.

Willshaw, D., & Dayan, P. (1990). Optimal plasticity from matrix memories: What goes up must come down. *Neural Computation*, *2*(1), 85–93.

Yamauchi, B. M., & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, *2*(3).

Yao, X. (1999). Evolving artificial neural networks. *IEEE Transactions on Evolutionary Computation*, *87*(9), 1423–1447.

Yao, X., & Liu, Y. (1997). A new evolutionary system for evolving artificial neural networks. *IEEE Transactions on Neural Networks*, *8*(3), 694–713.

Yoder, J., & Yaeger, L. (2014). Evaluating topological models of neuromodulation in polyworld. In *Artificial life XIV*. Cambridge, MA: MIT Press.

Young, S. R., Rose, D. C., Karnowski, T. P., Lim, S.-H., & Patton, R. M. (2015). Optimizing deep learning hyper-parameters through an evolutionary algorithm. In *Proceedings of the workshop on machine learning in high-performance computing environments* (p. 4). ACM.

Ziemke, T., & Thieme, M. (2002). Neuromodulation of reactive sensorimotor mappings as short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, *10*, 185–199.