UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

NEUROMODULATION IN ARTIFICIAL SYSTEMS

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

ESTHER LO
Norman, Oklahoma
2012

NEUROMODULATION IN ARTIFICIAL SYSTEMS

A THESIS APPROVED FOR THE
DEPARTMENT OF ZOOLOGY

BY

_____
Dr. Thomas Ray, Chair

_____
Dr. Randall Hewes

_____
Dr. Dean Hougen

Table of Contents

List of Figures

Introduction

Neuronal communication is most often characterized by a rapid transfer of information between two cells, beginning with the generation and conduction of an electrical impulse in the presynaptic cell, leading up to a release of chemical messengers that traverse the synaptic cleft, and finally culminating with a direct inhibition or excitation of the postsynaptic cell, the activity of which, in turn, is effectively determined by a simple summation of such signals. There exists, however, another form of communication between neurons, one that is often relatively understated and involves not the simple and direct activation of membrane ion channels but rather the activation of elaborate cell signaling pathways, that is equally central to the normal functioning of neural circuits. This type of communication, termed *neuromodulation*, usually involves the activation of G-protein coupled receptors and their associated intracellular signaling cascades rather than the activation of ligand-gated ion channels. Thus, neuromodulatory communication results in a wider range of effects that are generally slower-acting and longer-lasting. This type of signaling, which occurs at every level of the nervous system and shares many of the same chemical messengers with the primary or traditional method of neuronal communication, allows a neuron or neural circuit to take on multiple functional roles such that one type of stimulus may elicit a number of different responses that depend on the state of the neuron, circuit, or animal.

Because the information-processing capacity of an animal's brain emerges from the combination and interaction of both traditional neurotransmission and neuromodulation, an artificial emulation of such a structure for the purpose of data processing should reasonably also involve both modes of communication. Artificial

neural networks (ANNs), which are inspired by their biological counterparts, are computational constructs used for a large variety of computing tasks. For a number of reasons, including practicality, functionality, and the inherent differences between the biological and digital domains, the extent of biological influence on ANNs is unsurprisingly limited, and many aspects of biological networks, including neuromodulation, are omitted from implementations of ANNs. However, because neuromodulation is not a simple or minor enhancement of network function, but is rather an essential component of normal operation, its incorporation into ANNs may impart substantial benefits. Neuromodulation in ANNs, for example, may increase network robustness or solution evolvability, as has been demonstrated by a small number of implementations reviewed here. The purpose of this work is to illustrate ways in which neuromodulation may take place in artificial networks by extracting the fundamental attributes of the biological process and transferring them to an artificial setting. Because biological and artificial networks are essentially distinct, each possessing unique capabilities and constraints, a direct transfer of the neuromodulatory process between the two domains would be impractical. Corollaries of ideas presented here will not necessarily exist in biology, nor would all or even most biological neuromodulatory mechanisms, systems, or effects be entirely imitable in a digital setting. Furthermore, there may arise adverse effects from the introduction of neuromodulation to ANNs that are unanticipated due to the fundamental ways in which they differ from biological networks. For example, neuromodulation may cause instability by altering or interfering with a network's basic mode of function. Also, an excessive number of network components resulting from the addition of an extra

signaling system may complicate network training or evolution. Therefore, implementations of neuromodulation should include some level of control over this added signaling system as well as adhere to simplicity in design, which is one of the most crucial aspects of ANNs.

The contents of this work are as follows. Chapter one presents the concept of neuromodulation in biological systems. The primary mechanisms and functional consequences of neuromodulation at different levels are demonstrated through various examples. Chapter two provides a basic overview of ANNs. Chapter three reviews existing implementations of ANNs that include a neuromodulatory system. The design details and functional significance of each implementation are presented. Chapter four illustrates additional ways in which the process of neuromodulation may be incorporated into ANNs.

Chapter 1: Biological neurotransmission and neuromodulation

1. Overview of neural communication

A biological neuron is a cell that is specialized for information transmission, structurally characterized by extensions from the main cell body called axons and dendrites (Purves *et al.* 2008, p. 1-22). Typically, incoming information is received at multiple dendrites and outgoing signals are dispatched from a single axon (Purves *et al.* 2008, p. 1-22). Signals are electrochemical: neurons accumulate negative charge on the inside by maintaining specific ion gradients across the cell membrane, allowing signals to travel rapidly through the neuron in the form of an electric current (Purves *et al.* 2008, p. 41-60). Communication between a neuron and its target is usually chemical, involving the release of neurotransmitters (NTs) (Purves *et al.* 2008, p. 85-118).

The firing of a neuron begins with a positive change (depolarization) in the neuron's membrane potential triggered by a change in the flow of certain ions across the cell membrane (Purves *et al.* 2008, p. 25-39). Depolarization above a threshold value leads to the generation of an action potential (AP), a wave of depolarizations propagated down the length of the axon (Purves *et al.* 2008, p. 41-60). The terminal end of an axon contains vesicles containing NTs and Ca^{2+} channels that open in response to depolarization (Purves *et al.* 2008, p. 85-118). The AP opens these channels, and Ca^{2+} influx mediates the fusion of the vesicles to the cell membrane, releasing their contents (Purves *et al.* 2008, p. 85-118). NTs diffuse across the synapse and are received by receptor proteins in the postsynaptic cell, triggering a wide range of possible responses (Purves *et al.* 2008, p. 85-118). Even though it is common to categorize a NT as either

predominately excitatory or inhibitory, it is the identity of the receptor that ultimately determines the effects of the NT (Purves *et al*. 2008, p. 119-152).

Receptors can simply be transmembrane proteins that open after the binding of a NT. These ion channels, called *ionotropic* receptors (Trimmer 1999), will increase the permeability of certain ions (and thus raise or lower the membrane potential), directly altering the likelihood of AP generation. Activated receptors can also invoke other kinds of responses within the cell. These receptors, called *metabotropic* receptors (Trimmer 1999), activate molecules called G-proteins that set off long cascades of intracellular signals, which, in turn, can alter membrane permeability directly or give rise to more long-lasting changes. These long-lasting changes, termed *neuromodulation* (Katz 1999), are essential for continual adaptation to the environment and are the focus of the remainder of this chapter.

2. Neuromodulation

The term neuromodulation will be used here to describe a certain variety of changes that a neuron undergoes as a result of communication with another neuron. These changes are set apart from those often described in simplified or classical views of neurotransmission, in which the rapid fluctuations in cell membrane potential, mediated primarily by ionotropic receptors, are underscored (Katz 1999; Trimmer 1999). Neuromodulatory communication, on the other hand, alters longer-lasting properties of the neuron that can affect future signaling events. Neuromodulatory effects often emerge later and last longer than those of neurotransmission due to the biochemical cascades activated by metabotropic receptors.

Metabotropic receptors are coupled to molecules known as G-proteins, which are activated when the receptor binds a NT. An activated G-protein can either directly interact with ion channels (Purves *et al*. 2008, p. 85-118) or activate effector enzymes (such as adenylyl cyclase, guanylyl cyclase, and phospholipase C) (Purves *et al.* 2008, p. 153-176), which in turn produce molecules called second messengers. Common examples of second messengers include cyclic nucleotides (cAMP and cGMP), Ca^{2+}, diacylglycerol, and inositol triphosphate (Purves *et al.* 2008, p. 153-176). Targets of second messengers are most commonly enzymes called kinases and phosphatases that add and remove phosphate groups from their substrates, respectively (Purves *et al.* 2008, p. 153-176). Many proteins, including those associated with ion channels and ion channels themselves, can be regulated by the activity of these enzymes (Jonas and Kaczmarek 1999). Consequently, the long-term effects of neuromodulators are primarily attributed to the activity of the proteins activated by second messengers. The following sections will describe some of the molecular mechanisms of neuromodulation and the effects of neuromodulatory communication on local cell behavior and global circuit function.

2.1 Molecular mechanisms of neuromodulation

2.1.1 Ion channels

Ion channels are proteins that span the cell membrane and allow the passage of certain ions. The behavior of a neuron is shaped by the flow of ions (see section 2.2) and principally altered by the modulation of ion channels via the activity of various kinases

(such as those activated by the biochemical cascades described above). Phosphorylation may directly or indirectly modulate ion channels, subsequently altering the ion currents they mediate. This phenomenon has been widely studied and reviewed (Swope *et al.* 1992; Levitan 1994; Ismailov and Benos 1995; Swope *et al.* 1999).

In a cell, the total current mediated by a single type of ion channel may be represented by $I = nPi$, where n is the number of active membrane channels, P is the probability that a channel is open, and i is the current of a single channel. Thus, ion channel phosphorylation may theoretically modulate ion currents by altering any of these variables (Levitan 1985). A number of underlying biophysical mechanisms that can produce such changes have been reported in studies of both ligand-gated channels (ionotropic receptors) and voltage-gated channels (which activate in response to changes in membrane potential). Examples of these mechanisms include changes in desensitization rates (Downing and Role 1987; Schmidt *et al.* 1994; Hinkle and Macdonald 2003), activation of silent membrane channels (Margiotta *et al.* 1987; Vijayaraghavan *et al.* 1990), regulation of channel trafficking (Lan *et al.* 2001; Wang *et al.* 2003; Misonou *et al.* 2004) or subunit assembly (Ross *et al.* 1991), and changes in voltage dependence (Reinhart *et al.* 1991; England *et al.* 1996; Fitzgerald *et al.* 1999; Park *et al.* 2006) or various gating kinetics (Numann *et al.* 1991; Covarrubias *et al.* 1994; Roeper *et al.* 1997; Chen *et al.* 2006).

Long-term modulation of ion currents can occur through changes in the transcription of ion channel genes. Second messenger cascades can relay signals to the nucleus of the cell to promote RNA synthesis by activating transcription factors (proteins that bind to DNA and enable transcription). The creation of new ion channels

(and other nuclear signaling events in general) results in relatively long-lasting changes in the behavior of the neuron.

2.1.2 Neurotransmitter release

Neuron function may also be modulated presynaptically by controlling the release of NTs. Vesicle exocytosis can be regulated directly. The parts of the exocytotic machinery that may be targeted (by kinases, for example) for modulation include fusion proteins on the vesicles and cell membrane, the protein that detects Ca^{2+}, proteins involved in the linking of Ca^{2+} channels to vesicles, and proteins involved in vesicle trafficking (Miller 1998; Jonas and Kaczmarek 1999). Because exocytosis is dependent upon Ca^{2+} influx into the axon terminal, NT release can also be modulated through the regulation of Ca^{2+} levels and dynamics. Certainly, Ca^{2+} currents may be regulated by modulating Ca^{2+} channels. Ca^{2+} channel modulation by various mechanisms, such as G-protein activation, and its effects on Ca^{2+} currents and neuron function, including NT release, has been reviewed (Dolphin 1990; Miller 1990; Numann *et al.* 1991). The inhibition of Ca^{2+} currents leads to a decrease of NT release and presynaptic inhibition, which can be demonstrated by the concurrent reduction of Ca^{2+} influx and synaptic transmission following the activation of the metabotropic receptors of various NTs such as GABA (Isaacson 1998), glutamate (Takahashi *et al.* 1996), and acetylcholine (ACh) (Qian and Saggau 1997). Ca^{2+} influx may also be attenuated in a less straightforward manner. The depolarization propagated by an AP that opens terminal voltage-gated Ca^{2+} channels may be countered by altering the conductance of other ions (such as K^{+})

such that the cell is repolarized and Ca^{2+} currents are reduced (Miller 1990; Miller 1998).

2.2 Cellular effects of neuromodulation

Neurons show a diverse range of intrinsic firing patterns, and these are subject to regulation by ion channel modulation. While some neurons fire a single AP following depolarization, others can exhibit bursting activity in which APs are fired in rapid succession in between periods of silence (Izhikevich 2000). Furthermore, many neurons can also display spontaneous firing activity in the absence of stimuli (Marder and Calabrese 1996). The state of a neuron (i.e., whether it is at rest or firing) at any moment depends on its membrane potential, and its behavior over time is determined by changes in the membrane potential resulting from the modulation of ion currents. A single neuron may simultaneously express different types of ion channels on its surface (e.g., Na^+, K^+, Cl^-, Ca^{2+}, and nonspecific cation channels), including a number of varied subtypes, which are exquisitely regulated to produce a wide variety of firing patterns. Modulation may also occur presynaptically via the alteration of Ca^{2+} currents, resulting in the facilitation or depression of NT release and altering the degree to which the postsynaptic element is excited or inhibited.

Several neural systems in the well-studied marine mollusk *Aplysia* serve to illustrate the cellular effects of both postsynaptic and presynaptic neuromodulation. Firstly, the neuromodulator serotonin (5-HT) has been shown to have complicated effects on the endogenously (spontaneously) bursting neuron R15. Via the cAMP pathway, low concentrations of 5-HT acts to eliminate bursting (Drummond *et al.* 1980)

by enhancing an inward K^+ current (Benson and Levitan 1983) (which hyperpolarizes the cell), while at higher concentrations, 5-HT enhances the AP frequency during bursts and eventually induces a tonically firing state (in which the neuron fires single spikes periodically) by increasing an inward Ca^{2+} current (which enhances depolarization) in addition to the K^+ current (Levitan and Levitan 1988). R15 modulation shows that one neuromodulator can have opposing, concentration-dependent effects on firing activity.

Secondly, studies of the accessory radula closer (ARC) muscle illustrate that neuromodulators can work in concert (here, by adjusting opposing currents) to produce the range of contractions necessary for a complex motor behavior. Note that this is an example of neuromodulation in non-neuronal excitable cells. The ARC muscle, which is involved in feeding, is innervated by two neurons that release several distinct classes of neuromodulators together with ACh (which induces muscle contractions) (Cropper *et al.* 1987). Additional neuromodulators are released from other neurons in the network (Cropper *et al.* 1994). Acting through second messengers, neuromodulators both enhance and depress contractions by increasing Ca^{2+} influx (Brezina *et al.* 1994b) and activating K^+ channels (Brezina *et al.* 1994a), respectively. Furthermore, the motor neurons may be inhibited presynaptically by modulating the release of ACh (Cropper *et al.* 1988).

Finally, both long- and short-term presynaptic modulation resulting in the facilitation of NT release is demonstrated by the gill and siphon withdrawal reflex (GSWR), in which a touch to the siphon results in gill withdrawal. The GSWR can be sensitized for a short period of time by applying a noxious stimulus elsewhere on the body (Pinsker *et al.* 1970; Kandel and Schwartz 1982), and long-term sensitization can

occur following repeated applications of noxious stimuli (Pinsker *et al.* 1973). Short-term sensitization is attributed to the release of 5-HT from interneurons upon noxious stimulation, triggering the suppression of K^+ channels in siphon sensory neurons via a cAMP cascade and, in turn, an increase in Ca^{2+} influx and NT release (Kandel and Schwartz 1982). The same cAMP pathway, initiated by 5-HT, leads to long-term sensitization by signaling to the nucleus to stimulate protein synthesis by activating transcription factors (Dash *et al*. 1990).

2.3 Functional significance of neuromodulation

From sensory receptors to neuromuscular junctions, neuromodulators can target essentially any type of synapse in the nervous system, effecting the continuous changes necessary for normal function. Neuromodulation can alter the response of primary sensory neurons to both external and internal stimuli, the generation of motor patterns by lower central circuits and motor circuits, and numerous facets of brain function.

2.3.1 Neuromodulation of sensory systems

Sensory receptors transduce chemical, mechanical, light, and other types of stimuli into electrochemical signals that can propagate through the rest of the nervous system. The sensitivity to stimuli of a sensory neuron is commonly increased as a result of neuromodulation. For example, the release of 5-HT from interneurons results in the sensitization of siphon mechanoreceptors in *Aplysia* (see section 2.2). Recently, this type of sensitization in *Aplysia* was observed to be elicited by a naturally occurring stimulus (an attack by the spiny lobster *Panulirus interruptus*), suggesting that this

phenomenon may be an adaption to prevent subsequent predatory attacks (Watkins *et al*. 2010).

Another example of increased receptor sensitivity is the sensitization of nociceptors (pain receptors) that follows tissue injury. Nociceptor sensitization is a result of the release of inflammatory mediators and neuromodulators such as 5-HT, leading to hyperalgesia (an increased response to pain), which probably serves as a protective mechanism to avoid further damage and promote wound healing (Treede *et al.* 1992). Hyperalgesia can also be induced centrally by the neuromodulator-induced sensitization of secondary sensory neurons that synapse with nociceptors in the spinal cord (Schaible *et al*. 2002).

Neuromodulation can also heighten sensitivity to other kinds of environmental stimuli, such as odor. For example, octopamine (OA), an insect neuromodulator, has been shown to increase sensitivity to pheromones in various species of moths (Linn and Roelofs 1986; Pophof 2000; Grosmaitre *et al.* 2001). There is evidence that OA may do so by acting directly on olfactory receptor neurons, either by directly increasing receptor sensitivity (Pophof 2000), or possibly by modifying the adaptation state of the receptor (i.e., reversing receptor desensitization) (Grosmaitre *et al.* 2001).

Sensitivity to internal stimuli may also be modulated. For example, in the locust *Locusta migratoria*, the response of forewing stretch receptors, which are connected to neurons involved in flight rhythm generation and are involved in the control of the wing-beat frequency, is increased by OA, which is released during the onset of flight (Ramirez and Orchard 1990). Here, the neuromodulation of a sensory receptor has direct consequences on motor behavior.

2.3.2 Neuromodulation of motor systems

Animals perform a full spectrum of behaviors using many different sets of motor patterns that, owing to neuromodulation, do not necessarily require a large number of neurons organized in extensive networks. Studies of central pattern generators (CPGs) provide perhaps the best understood effects of neuromodulation of motor behavior. CPGs are usually central circuits that drive motor neurons to produce rhythmic patterns and are capable of sustaining these patterns without sensory feedback.

A system that is classically used to illustrate neuromodulation in CPGs is the crustacean stomatogastric nervous system (STNS). The STNS is composed of four CPGs that produce four distinct rhythms involved in the grinding and filtering of food in the crustacean stomach (Marder and Bucher 2007). Neuromodulation has a number of well-characterized effects on this system. Although distinct neurons have been identified as belonging to specific CPGs, some may switch between different CPGs as a result of neuromodulatory input (Hooper *et al.* 1990; Weimann *et al*. 1991). Furthermore, two whole CPGs can be merged to produce a new rhythm distinct from the patterns produced by either of the original CPGs (Dickinson *et al*. 1990). Finally, neuromodulators can even construct a new network by selecting specific components from existing CPGs (Meyrand *et al*. 1991). In this system, a functional (rather than an anatomical) rewiring of network components by weakening or strengthening specific synapses dramatically demonstrates the flexible nature of motor networks.

Neuromodulation can also temporarily transform a non-rhythmic network into a CPG. For example, in the marine mollusk *Tritonia*, one network underlies reflexive

withdrawal as well as crawling, both of which are non-rhythmic behaviors (Popescu and Frost 2002). However, when presented with an aversive stimulus, the network is transformed, via the action of 5-HT, into a CPG that produces the patterns necessary for escape swimming (Katz *et al.* 1994; Frost *et al.* 2001).

2.3.3 Neuromodulation of brain function

The large variety of psychoactive drugs that target neuromodulatory pathways reflect the important functions of neuromodulation in the mammalian brain. The primary neuromodulators in the brain, dopamine (DA), norepinephrine (NE), 5-HT, histamine, and ACh, all have anatomically distinct pathways and play characteristic roles in greatly diverse aspects of brain function. Some of these aspects include sleep and arousal, attention and alertness, various kinds of learning and synaptic plasticity, memory, cognition, stress and anxiety, motor activity, consciousness, mood, spirituality, general mental state, and a number of neurological disorders such as Parkinson's disease, Alzheimer's disease, and schizophrenia (Smythies 2005).

An example of one of the many crucial functions of neuromodulation comes from the regulation of the gating of sensory information. Almost all inputs received from sensory systems converge in the thalamus, which filters and relays the information to the cerebral cortex. It is necessary to differentially process sensory input depending on the state of an organism. For example, responsiveness to outside stimuli is decreased during sleep compared to the waking state, and the modulation of thalamocortical neurons is thought to be critically involved in the switching between these states (McCormick 1992; McCormick and Bal 1997). In these neurons, the transition from

the rhythmic bursts that characterize slow-wave sleep to the more rapid, tonic firing that characterizes rapid eye movement sleep and the waking state follows the release of ACh, NE, histamine, glutamate, and possibly other modulatory substances from other parts of the brain, and correlates with an increase in the responsiveness of the thalamus to input from the sensory periphery (McCormick 1992; McCormick and Bal 1997).

Another example of a critical process that may be regulated by neuromodulation is synaptic plasticity, which is the inherent ability of synaptic strength to change according to synaptic activity. Perhaps the most widely studied component of synaptic plasticity is long-term potentiation (LTP), which is characterized by a persistent increase in synaptic strength and has long been a prominent candidate for the cellular basis of learning and memory. Therefore, the neuromodulation of LTP has the consequence of affecting the degree to which information is acquired and stored. In the hippocampus (a region of the brain associated with memory processing), for instance, a number of neuromodulators such as ACh (Huerta and Lisman 1993), NE (Izumi and Zorumski 1999), and DA (Otmakhova and Lisman 1996) have been shown to promote LTP. It may be beneficial to enhance learning and memory formation in certain circumstances. For example, in the rat hippocampus, a DA-dependent enhancement of LTP was observed following exposure to a novel environment (Li *et al.* 2003).

2.4 Regulation of neuromodulation

The previous sections have detailed the ways in which classical neurotransmission may be modulated. The ways in which these neuromodulatory processes are regulated are relatively less well known. However, the modulation of neuromodulatory neurons by

other neuromodulatory neurons has been characterized. This phenomenon, in which the primary neuromodulatory neurons that modulate classical neurotransmission are targeted by other (second-order) neuromodulatory neurons, has been termed *metamodulation* (Katz and Edwards 1999). Metamodulation may occur in a serial fashion (e.g., when a neuromodulatory neuron is directly activated by another neuromodulatory neuron) or in an indirect or convergent fashion, in which the effectiveness of the primary neuron is altered (Katz and Edwards 1999). The latter may occur in several ways. Second-order neurons may indirectly modulate primary neurons by altering their output (e.g., by modulating the amount or type of substance released), the responsiveness of the postsynaptic element (e.g., by desensitizing metabotropic receptors of targets of primary neurons), or the concentration of neuromodulators within the synaptic cleft released by primary neurons (e.g., by modulating reuptake transporters or acting on degradative enzymes that remove NTs from the synapse) (Katz and Edwards 1999). These events are accomplished by some of the same mechanisms used in primary neuromodulation.

3. Summary

The classical story of neurotransmission is conveyed by the fast excitatory and inhibitory communications that are predominately mediated by ionotropic receptors. Neuromodulation, characterized by its regulatory effects and a slow time course of action, plays an equally essential role in neuronal communication. At the molecular level, neuromodulatory substances such as 5-HT, DA, NE, and ACh activate metabotropic receptors, setting off a series of G-protein-mediated events within the cell

that eventually affects the phosphorylation states of ion channels or various aspects of vesicle exocytosis and NT release. On a cell-wide level, postsynaptic modulation alters a neuron's firing pattern, while presynaptic modulation alters the amount of NT released to impact synaptic strength. Functionally, neuromodulation augments network flexibility and seemingly reduces complexity by imparting a one-to-many correspondence between input and output. Broadly speaking, the behavior of a neuron or network that is under the influence of neuromodulators depends not only on the eliciting stimulus but also on the current state of the organism. Finally, in addition to granting flexibility to all components of a nervous system, neuromodulators can serve a higher regulatory function as regulators of neuromodulation itself.

Chapter 2: Artificial neural networks

1. Introduction

Artificial neural networks (ANNs) are parallel computational tools inspired by the processing capabilities of animal brains. Biological neurons work together simultaneously and continuously to process an abundance of information. This type of parallel processing is emulated by ANNs, which are composed of simple, interconnected individual processing units called neurons or *nodes* that form larger, complex processing systems (Gurney 1997). Like neurons in a brain, these simple processing elements function independently but work simultaneously to detect patterns in data. Furthermore, connections between nodes have variable synaptic strengths called *weights* (Gurney 1997). Essentially, an ANN is a pattern detection system in which specific input-output mappings are encapsulated in its weights. Brains exhibit synaptic plasticity, or the capacity to perpetually adjust synaptic strength in response to changes in the environment, which allows for the acquisition and storage of information. This type of plasticity also occurs in an analogous fashion in ANNs. ANNs learn specific models of data by undergoing a training phase in which example problems are presented and a training algorithm is used to search for appropriate weights. Because models are learned without explicit programming, extensive prior knowledge of the data is not necessarily required. The simple design of ANNs makes them widely applicable in data processing and modeling (Gurney 1997; Rabuñal and Dorado 2006).

The following sections will briefly describe the basic elements and operation of ANNs. Evolutionary techniques for training ANNs and the use of ANNs in robotics will also be discussed, as the subsequent chapter is concerned with these topics.

2. Artificial neurons

The neuron is the basic processing element in a neural network. Like the biological neuron upon which it is based, an artificial neuron (fig. 1) integrates input signals and transmits the information to other neurons. However, the similarities are essentially limited to this overly simplified view of neurotransmission. Biological neurotransmission involves chemical messengers, receptor proteins, ion channels, and an enormous array of intracellular processes. Thus, the flow and nature of information is more difficult to understand and not easily predictable. Information transfer between artificial neurons is much more direct. An artificial neuron receives one or more inputs, sums them, and then passes the sum of its inputs through an activation function to compute one output (Gurney 1997). Activation functions are commonly nonlinear, bounded, and sigmoidal (Dreyfus 2005). A commonly used activation function is the logistic function: $f(x) = (1/(1+e^{-x}))$. The connections between neurons are not uniform but have associated weights that determine the strength of the connection between two neurons. Connection weights are used to compute the weighted sum of inputs that is passed to the activation function (Gurney 1997). Commonly, thresholds are set by inputs from bias neurons that transmit constant values with associated bias weights (Dreyfus 2005).

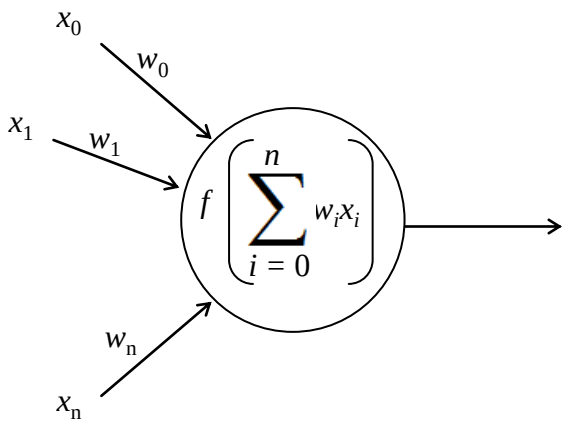

Figure 1. Artificial neuron. This neuron has *n* inputs. The parameter passed to the activation function *f* is a weighted sum: $x_0w_0 + x_1w_1 + ... + x_nw_n$, where x_i is the i^{th} input to the neuron and w_i is the weight of the connection between input *i* and the neuron. The output is then passed to all neurons that have an incoming connection from this neuron.

Because each neuron is basically a nonlinear function, a connected network is a composition of many nonlinear functions that form a system capable of representing nonlinear, diverse, and complex mappings.

3. Network structures

There are many ways in which neurons can be connected to each other in a network, and a network can take on many different configurations. In general, network topologies fall into two main classes: feedforward networks and recurrent (or feedback) networks. A feedforward network is an acyclic network in which information only flows in one direction (Gurney 1997). This kind of network is a direct function of its inputs, since no internal states can be stored. Usually, these networks are graphically represented by layers, beginning with an input layer and ending with an output layer. Any layers in between are termed *hidden* layers. Each layer receives input from the

immediately previous layer and transmits to the next layer. Recurrent networks, on the other hand, have at least one cycle or loop (Gurney 1997). A neuron's output may then be fed back into the network as input. The output depends not only on the given input but also on the network's internal state. Because of feedback connections, recurrent networks are dynamic and can exhibit complicated behaviors.

There are many different types of networks within these two main classes. Networks may take on a number of different configurations. For example, networks do not necessarily have to be arranged in layers. It is also not necessary for networks to be fully connected, and the pruning of connections (when prior knowledge of the problem allows for it) may allow for more efficient learning (Rabuñal and Dorado 2006).

4. Learning

Neural networks must be trained to embody a desired input-output mapping. Training involves the implementation of some learning algorithm that will search for suitable network parameters that will allow the network to produce some desired output. The adaptable parameters are most often the connection weights, which are adjusted within the training algorithm. In instances where correct output values are given to the network, an algorithm computes the error between network output and desired output and attempts to minimize this error by continually adjusting the weights according to some predefined update rule. The error is calculated by comparing the target value and the actual output given by the network and is then used to adjust the weights. The goal of training is to gradually minimize the error as training progresses. This type of training is known as supervised learning (Hassoun 1995), since the desired output is

known and used for comparison. There is a type of learning in which exact desired outputs are not provided, but feedback from the environment still exists (Hassoun 1995). In reinforcement learning, qualitative signals conveying the desirability of an output are presented to the network, which subsequently learns by attempting to maximize a reward function. This kind of learning is comparable to some types of associative learning in animals, such as operant conditioning. Learning may also be unsupervised. In this case, the desired output is not presented to the network (Hassoun 1995). In these cases, the network must find patterns present in the input data in the absence of environmental feedback, and weights must be adjusted using rules that do not rely on error values. An example of such an update rule is the Hebbian learning rule (Hassoun 1995). Here, the weight between two nodes is adjusted according to their simultaneous activation, much like the Hebbian plasticity that occurs at biological synapses.

5. Evolutionary algorithms and evolutionary robotics

A commonly used algorithm for training neural networks (and other types of search problems) is a class of algorithms based on Darwinian evolution and selection. Evolutionary algorithms (EAs) evolve populations of solutions to search and optimization problems (Bäck *et al.* 2000), such as finding optimal neural network weights. In genetic algorithms (GAs) (Holland 1975), a frequently used subclass of EAs, solutions are represented as a string of numbers. For example, in the case of neural network training, these numbers may be weights. Each individual solution competes with the rest of the population to create offspring solutions for the next

generation. Each solution is evaluated over a given period of time (i.e., its lifetime) by a fitness function that calculates how successfully the problem of interest was solved, and solutions with higher fitnesses are more likely to be selected to reproduce for the next generation. In reproduction, crossover between parent solutions, in which parts from each parent solution are used to create new solutions, and mutation of solutions (i.e., random changes in solution strings) are implemented to introduce population variation.

Of particular concern to systems discussed in the next chapter is the application of EAs in the development of autonomous robots, which are robots that can adapt to changes in the environment without outside guidance. Typically, such robots are assigned relatively simple tasks (e.g., maze navigation or locomotion) that require them to adaptively interact with the environment to handle various challenges (e.g., obstacle avoidance or maintaining balance in unsteady terrains). In evolutionary robotics, EAs are used to create robot control systems, which are often ANNs (Nolfi and Floreano 2000). Considering the functions of their biological counterpart, it seems natural to employ ANNs as robot neurocontrollers, which represent low-level mappings of sensory input to motor output. Common targets of evolution in neurocontrollers include connection weights and network structure.

6. Summary

Biological brains have inspired a class of computational systems to perform complex, nonlinear mappings. ANNs may be trained by a variety of methods that are traditionally categorized based on the amount of feedback or guidance received from

the learning environment. Evolutionary computing techniques, such as GAs, use concepts from biological evolution and selection to search for solutions to optimization problems, such as finding suitable weights in a neural network. ANNs are widely applicable in many scientific fields, but the next chapter largely focuses on their role in evolutionary robotics, a field that uses evolutionary computation to develop autonomous robotic agents.

Chapter 3: Review of neuromodulation in artificial neural networks

1. Introduction

A crucial component of biological networks that is conspicuously absent in most implementations of artificial neural networks (ANNs) is neuromodulation, a major aspect of neuronal communication that is as fundamental to biological networks as classical neurotransmission. Since simplicity is an important attribute of ANNs, and tasks that they are designed to accomplish are relatively simple, an attempt to build them progressively more similar to animal brains would be an undesirable (and formidable) endeavor. However, functional benefits can conceivably arise in neuromodulated ANNs without significant sacrifices in simplicity. Furthermore, as is the case in biological systems, neuromodulation may actually allow more simplicity in network design than a functionally equivalent network that is not neuromodulated.

The following sections provide a review of ANNs that include a neuromodulatory component. For each system, a general description of the network architecture and the neuromodulatory mechanism will be given, followed by notable experiments in which they were investigated. This chapter concludes with comments concerning the nature and possible functional roles of neuromodulation in ANNs.

2. Neuromodulation in literature

2.1 GasNets

GasNets are neuromodulated networks based on nitric oxide (NO) signaling (Garthwaite and Boulton 1995). NO is a gaseous, non-traditional signaling molecule used by the nervous system in nonsynaptic neurotransmission and neuromodulation. NO differs from traditional signaling molecules because it is a membrane-permeant molecule that freely diffuses from a cell to multiple targets (Garthwaite and Boulton 1995). GasNets are unconventional networks that operate on a two-dimensional (2D) plane. Instead of incorporating direct, point-to-point links between nodes, network connections are determined by positive and negative circle segments that are associated with each node. Two nodes are linked if one node resides within the segment of the other, and the type of segment (positive or negative) determines the type of connection (excitatory or inhibitory). This arrangement facilitates the diffusion of modulatory signals: each node emits, under a predefined condition, zero or one of a number of distinct and diffusible "gases" (neuromodulators) that disperses radially on the 2D plane according to a mathematical model. If a node emits a gas, it will do so either due to an activity threshold or a gas concentration threshold. The type of gas each node may emit, the conditions for emission from each node, various properties of diffusion (such as the rate of decay), and network topology are all evolved. The transfer function of each node is subject to continuous gaseous modulation according to the type of gas and its concentration.

Two types of GasNets with two different transfer functions have been developed. One type of GasNet uses four distinct gases, and another type uses two gases. In the former version, the transfer function contains two variables that are

modifiable by the gases, while the transfer function in latter version contains one variable that is modifiable.

GasNets were implemented in robot controllers given various tasks such as T-maze navigation (Husbands 1998), in which the robot must decide the correct direction to turn at the T-junction of a T-shaped maze based on a light signal coming from either side of the junction, and target discrimination (Husbands *et al*. 1998a; Husbands *et al.* 1998b), in which the robot must move towards one of two shapes attached to a wall. In both groups of experiments, GasNets were found to be more evolvable than conventional networks that lacked gas diffusion. That is, evolution found successful solutions faster when GasNets were in use. Furthermore, evolved GasNet controllers were generally simpler in design than conventional networks. In the latter set of experiments, GasNets were also found to be more evolvable than the same type of 2D network in which gases were removed (Husbands *et al*. 1998a; Husbands *et al.* 1998b). Furthermore, the two types of GasNets (using four and two gases) produced similarly successful results (Husbands *et al.* 1998b), further suggesting that the process of gaseous modulation itself, rather than a particular type of network or transfer function, was responsible for increasing the rate of successful evolution. The authors have noted that the heterogeneity of the networks emerging from variation amongst transfer functions is important for successful performance. In instances in which gases did not evolve (Husbands 1998) or modulation was turned off (Husbands 1998a), fairly successful networks shared the common feature of heterogeneity. GasNets have also demonstrated success in the more complex robotic task of legged locomotion. GasNet controllers achieved the highest fitness out of more than a dozen types and variants of

dynamic recurrent networks in a simulated bipedal locomotion experiment (McHale and Husbands 2004a). In a related study of quadrupedal locomotion, GasNets were compared to two other recurrent networks (McHale and Husbands 2004b). Here, although simulated robots controlled by GasNets displayed more stable gaits, another type of network achieved similarly high fitness values. These results suggest that the nature of the problem or task is an important factor in evaluating or comparing the success of different networks.

Finally, several variations of GasNets have shown increased performance compared to traditional GasNets. In traditional GasNets, single nodes release modulators with corresponding non-uniform concentration distributions centered on the releasing node. One alternate GasNet version was based on the dynamics of NO release in certain areas of the brain in which the subthreshold concentrations of NO released by individual nerve fibers in a plexus sum to above-threshold levels some distance away from the cell bodies (Philippides and Husbands 2005). In this *plexus* GasNet model, concentration distributions are uniform and may center away from the releasing node. Another version was based on biological receptor proteins (Philippides and Husbands 2005). In this *receptor* GasNet model, nodes only respond to modulation if receptors are present, and only one gas with a single modulatory effect was used. The intensity of a modulatory signal depends on the gas concentration as well as the amount of receptors present. In a target discrimination task (as described above), both models were found to be more evolvable than a traditional GasNet, and receptor GasNet solutions evolved considerably faster than plexus GasNet solutions (Philippides and Husbands 2005). Another GasNet variant, termed a non-spatial GasNet (NSGasNet), does not operate on

a coordinate plane (Vargas *et al.* 2007). Instead, every node in the network is subject to modulation by every gas-emitting node depending on an evolvable parameter (termed *modulator bias*) that determines the intensity of modulation. Each node has a modulator bias associated with every gas-emitting node in the network. In a simple central pattern generator (CPG) task (Vargas *et al.* 2007), in which networks were evolved to cyclically generate a particular bit sequence (mimicking a sort of CPG), NSGasNets were shown to be more evolvable than traditional GasNets. Furthermore, the number of evolved successful NSGasNets was greater than that of traditional GasNets. However, in a more complex, delayed response task (Vargas *et al*. 2008) (similar to the T-maze task described above), although evolution found a greater number of successful NSGasNets than traditional GasNets, the speed of evolution was comparable between the two models. These results suggest that the success of GasNets may not hinge on a spatial model of gas diffusion.

2.2 Dynamically-rearranging neural networks

Another group of neuromodulated systems that were strongly inspired by biology are dynamically-rearranging neural networks (DRNNs) (Kondo *et al.* 1999), which are based on the neuromodulatory system found in the crustacean stomatogastric nervous system (STNS) (see chapter one). Just as neuromodulators can functionally rearrange circuits in the STNS, the synapses of DRNNs are dynamically altered as a result of neuromodulatory input. A neuron may release zero or one of two types of neuromodulators when its activation lies within a certain range. Neuromodulation alters the way in which the connection weight between the two neurons is updated based on

the combination of modulators received by the postsynaptic neuron. The synapse may become Hebbian, anti-Hebbian (in which the weight decreases proportionally to simultaneous presynaptic and postsynaptic activation), or non-learning (in which the weight remains the same). All of these factors (the type of modulator released, the range of activation values that will trigger release, and the postsynaptic response to modulation) are genetically determined.

In a simulated peg-pushing task, in which a robot controlled by a three-layered, fully connected feedforward network must push a peg towards a light source, both DRNNs and conventional networks (in which weights were evolved and fixed) were found capable of evolving successful networks (Kondo *et al.* 1999). However, successful DRNNs consistently outperformed conventional networks when random elements were introduced in the environment by slightly perturbing motor output, peg movement, or peg size. This suggested that neuromodulation increases robustness against environmental fluctuations. The authors proposed that DRNNs can help overcome the difficulty of employing agents that were evolved in simulation in the real world. Indeed, DRNNs implemented in real peg-pushing robots achieved success where conventional networks could not (Eggenberger *et al*. 1999).

The adaptability of DRNNs was further demonstrated in bipedal robot locomotion (Fujii *et al.* 2001; Ishiguro *et al.* 2003), in which a robot with jointed legs must walk across surfaces of varying degrees of smoothness. Here, the network consisted of a set of neural oscillators located at each joint in the body that generated rhythmic movement. Furthermore, four types of neuromodulators were used, and each had associated concentration values that were either proportional to the activation of the

releasing neuron (Fujii *et al.* 2001) or evolved (Ishiguro *et al.* 2003). These values were used to directly modify synaptic weight, and each synapse had an evolved set of receptors specific for each neuromodulator. In a simulation of locomotion on flat terrain, successful neurocontrollers were often not modulated (i.e., weights were fixed), suggesting that modulation is not especially useful or necessary in relatively static environments (Fujii *et al.* 2001). However, modulation was found to be necessary when controllers were evolved in strictly downhill or uphill environments and when evolved controllers were tested in novel environments in which the surface had downhill, uphill, and flat areas (Ishiguro *et al.* 2003). Modulated controllers were also able to withstand unexpected forces applied to the robot (Ishiguro *et al*. 2003). DRNNs have also been successfully evolved for simulated quadrupedal locomotion on even and uneven terrains (Otsu *et al.* 2001).

2.3 Multileveled networks

Extended sequential cascaded networks (ESCNs), which are recurrent networks that use internal activation values to dynamically update weights, have been used to explore several different robotic tasks (Ziemke 1999; Ziemke and Thieme 2002). The network consists of two subnetworks called the *function* network and the *context* network. The function network is the primary network that maps sensory input to motor output. This network also uses its internal state (i.e., its internal activation values) to produce input for the context network, which subsequently updates the function network's weights in the next time step. Furthermore, a decision unit maps internal state to input for the context network, which prevents the context network from updating the function

network when it was not activated. ESCNs are variations of sequential cascaded networks (SCNs), which do not have decision units and, thus, the function network is modulated at every time step. Because function network weights are updated by the context network, only context network weights were evolved in all experiments discussed here.

An ESCN was tested alongside a SCN, other recurrent but non-modulated networks, and a feedforward network in a simulated task in which the robot navigated towards a marked zone in an enclosed area while avoiding obstacles, such as walls (Ziemke 1999). A second scenario incorporated objects inside and outside of the target zone. Here, the goal of the robot was to navigate towards the zone, picking up objects inside of the zone and avoiding objects outside of the zone. In both scenarios, the ESCN outperformed all other networks, while the feedforward network had the worst performance by large margins.

ESCNs have also been used to investigate a possible role for synaptic plasticity in short-term memory (STM) in ANNs (Ziemke and Thieme 2002). The task involved a T-maze in which the simulated robot decided which direction to turn at the T-junction based on a stimulus encountered some distance prior to reaching the junction. The stimulus was a light that appeared on either side of the corridor, indicating the direction that the robot should turn. Because there is a delay between the stimulus and turn, the robot must somehow "remember" that stimulus long enough to make the correct decision. The authors thus investigated the use of the ESCN as a STM mechanism. Tested in a number of different T-mazes of varying degrees of difficulty, the ESCN controllers were able to efficiently utilize neuromodulation (via activation of the

decision unit) to switch between different sensorimotor mappings when it was necessary, while relying on a purely reactive (i.e., static) network otherwise. These results suggest that synaptic plasticity (via weight modification in the function network), which has primarily been considered a long-term memory mechanism in ANNs, may have a useful role in STM tasks.

A similarly structured, two-leveled network has also been designed and tested on the same type of delayed response task described above (Bergfeldt and Linåker 2002). This system consists of two levels, termed *echelons*, that are similar to the function and context networks of the ESCN. The lower level behaves like the function network, consisting of a basic echelon (echelon 1) that maps sensory input to motor output, while the higher echelon (echelon 2) is analogous to the context network, periodically modulating the lower level. An unsupervised algorithm was used to classify sensory input into discrete categories or events (e.g., "junction", "corridor", or "lights") such that echelon 2 was only updated whenever new events occurred. Furthermore, a gating unit (similar to the decision unit of the ESCN) that gated the output from the echelon 2 to echelon 1 was implemented to prevent unnecessary modulation. Outputs from echelon 2 were added to the bias weights of the output nodes in echelon 1. Thus, modulation here is limited compared to modulation in the ESCN, where a new sensorimotor mapping was established every time modulation was triggered.

Modulation by echelon 2 was found to be unnecessary when simulated robots were evolved for a simple delayed response T-maze task. In a related but more complicated maze task, in which the robot can choose to either turn or go straight

depending on the location of the light, echelon 2 was fully utilized. Furthermore, success was also achieved in a novel environment consisting of a maze with multiple corridors, stimuli, and turns.

Another modulated system (unrelated to those described above) that utilizes two levels of information processing was developed by Meng *et al*. (2010). This system uses a "gene regulatory network" (GRN) to continually modulate the primary network. This model was inspired by biological genes and the regulatory functions of their products. Neurons in the primary network contain genes, the products of which are used to modify various parameters of synapses (such as the learning rate). Each modifiable parameter corresponds to the expression of a gene. Gene products could also modify the activities of their own or other genes. Finally, input to a gene consists of two "ion concentrations", which are values proportional to the input and activation threshold of its neuron. Thus, the GRN modulates the primary network, which in turn affects the functioning of the GRN. The system was used to analyze video sequences showing various human behaviors (such as walking or running) and assign each sequence a behavior label. Results showed that the GRN-modulated system was capable of a high rate of behavior recognition, and outperformed several other methods.

2.4 Networks with modulatory neurons

Soltoggio *et al*. (2007) aimed to separate the regulatory function of neuromodulation from the process of neurotransmission by incorporating a special type of modulatory neuron alongside standard neurons. In this system, modulatory neurons carry out the singular role of modifying synaptic plasticity; their outputs do not affect the activation

of postsynaptic neurons. Standard neurons transform inputs via a standard sigmoid function, while modulatory neurons use a slightly altered function with a more limited range. Modulatory neurons exert their effect on the weight update rule for connection weights between traditional neurons. The summed modulatory input for each standard neuron simply acts as a multiplicative factor for the update rule between itself and all associated presynaptic neurons.

This type of network was tested in a simulated foraging task in which a bee explores a 3D environment populated by colored flowers with varying amounts of nectar. Flower colors indicated, with varying degrees of reliability, the amount of nectar the bee received upon reaching the flower. The uncertainty of reward introduced a dynamic component into the environment. This associative learning problem was based on the biological discovery of a neuron that mediates the unconditioned stimulus in the conditioning of a feeding behavior in honeybees (Hammer 1993). The number of each type of neuron and other network attributes, such as connectivity and update rule parameters, were evolved. Each network was evolved in different environmental scenarios that varied in the accuracy of flower color as a reward predictor. Modulated networks were tested against networks with non-modulated update rules (i.e., all summed modulatory inputs were set to one) and networks with fixed weights. Modulation of synaptic plasticity (via alterations of update rules) was found to be crucial to an agent's ability to develop a learning strategy essential for the fluctuating environment. Networks with modulation performed better than both networks without modulation and networks without any synaptic plasticity. Furthermore, solving the foraging task did not necessitate networks with complex designs as long as modulation

was incorporated, since successful networks typically had a small number of both types of neurons.

Interestingly, successful agents seemed to be able to adjust the level of neuromodulation according to the level of learning required by the environment. For example, modulation was low when flowers were outside of the visual field, while modulation was high upon reaching a flower. The ability to turn learning on and off was further demonstrated with the same type of modulated network used in a simulated T-maze-navigation task (Soltoggio *et al.* 2008). Here, a high reward was placed on one end of the T-arm, while a low reward was placed at the other end. The goal of the agent was to collect the high reward and return to its starting position. A double T-maze, which contained more turns and possible rewards, was also implemented. This learning task was similar to the foraging problem in that reward locations were not fixed. Modulated networks were once again tested against fixed-weight and non-modulated (but plastic) networks. Modulated and plastic networks were approximately equally superior to fixed-weight networks in solving the single T-maze task. However, both fixed-weight and plastic networks showed greatly lowered performance in the double T-maze. Successful networks evolved in this dynamic environment exhibited an ability to switch between explorative and exploitative behavior. Although this ability to switch strategies was present in all successful networks (both modulated and non-modulated), evolution with modulation found these strategies more successfully, and removing modulation from successful modulated networks lowered their performance. Thus, these results suggest that neuromodulation does not enable the evolution of successful networks but rather makes the process more efficient. However, a more realistic

extension of the simulated maze-navigation task revealed that it may be desirable to limit neuromodulation to certain parts of the network (Dürr *et al*. 2008). Here, agents aimed to avoid collisions in addition to locating rewards in a T-maze. Networks with modulation restricted to certain neurons achieved higher fitness than networks with unrestricted modulation and fixed networks. The authors proposed that unlimited modulation may disrupt simple behaviors that do not require plasticity, such as obstacle avoidance.

This type of neuromodulation has been used as a mechanism of introducing plasticity into networks in various applications. For example, Arnold (2011) employed plasticity via modulatory neurons to evolve reinforcement-free learning. The author asserted that if the network organization affecting problem-solving behavior is closely matched to the organization of the environment, networks can better tolerate dynamic environments by reducing the number of behavioral adjustments that would need to be carried out as the environment changes. It was further hypothesized that in imposing selection pressure on a network's ability to learn (here, through neuromodulation) by incorporating environmental unpredictability, isomorphism (i.e., similarity in organization) between the network and the environment would simultaneously emerge. Because it was proposed that the use of reinforcements or rewards in machine learning hinders the evolution of such isomorphism, networks were evolved for pursuing prey in a simple grid world using neuromodulation in a reinforcement-free "learning phase" that took place prior to operation in the test environment. In the learning phase, the effect of each action was randomized and actions did not affect fitness, thus networks in this phase learned only the effects of each action without learning how it might aid in

survival in the testing environment. This type of learning was accomplished by incorporating the same type of modulatory neuron described above. Results showed that evolution with neuromodulation during the learning phase led to much higher fitness than evolution without neuromodulation, thus it may be concluded that neuromodulation allowed the network to handle unpredictability in its actions as well as the subsequent emergence of a reinforcement-free learning ability.

Others have used the type of neuromodulation described here to demonstrate the advantages of evolving plastic networks using various non-traditional techniques in neuroevolution. For example, Tonelli and Mouret (2011) evolved neuromodulated networks using a map-based encoding scheme, which uses grids of neurons as building blocks, rather than more traditional methods of genetic encoding, while *Risi et al.* (2012) evolved neuromodulated networks using novelty search, which focuses on rewarding novel behavior rather than reaching the final objective.

2.5 Global neuromodulation and multiple modulatory effects

Inspired by the neuromodulatory processes responsible for emotion (Fellous 1999) and the possible functional roles of emotions in robotic agents (Fellous 2004), Parussel and Smith (2005) designed a neuromodulated neural network tasked to solve a simple action selection problem. The goal of the simulated agent (i.e., the network) was to maximize two resources by selecting from a discrete set of actions (i.e., output neurons) that had direct and predetermined effects on its internal state, which signified the need and satiety level of each resource. There was no tangible environment, thus the agent could only sense and act according to its internal state. In this system, neuromodulatory

signals are time-dependent decaying global signals that correspond to the need of each resource and are increased by the firing of certain neurons. Such neurons have "secretors" for modulators, while other neurons have "receptors". Receptors are either excitatory or inhibitory. Neurons receptive to modulatory signals are subject to one of two types of alterations: a change in input sensitivity or a change in firing probability. The type and extent of modulation depend on the type and level of the modulatory signal as well as the type of receptor. Secretors and receptors are restricted to neurons in certain layers. Modulation parameters, such as signal decay rate and secretion rate, were evolved alongside other network parameters.

Different experiments varied the locations of modulation and the number of types of modulators. Modulation of the outer layer (containing receptors) by the middle layer (containing secretors) by a single modulator outperformed both non-modulated networks and other versions of modulated networks (Parussel and Cañamero 2007). Contrarily, modulation of the middle layer by the input layer caused perturbations in the network, leading to fluctuations in action choices rather than stabilization on an optimal strategy, thus resulting in lowered performance compared to that of non-modulated networks (Parussel and Smith 2005). However, when these networks were tested for an extended period of time following evolution (i.e., in a novel environment), modulation proved to provide a robustness advantage (Parussel and Smith 2005). The primary significance of these findings is that neuromodulation could be used to alter the type of strategy that an agent employs. Depending on which layer was receptive to modulation, one of two possible behaviors was elicited: exploitative behavior or exploratory behavior, both of which are important strategies for learning agents.

French and Cañamero (2005) have also developed a simple neuromodulatory system that comprises multiple types of modulatory effects. A robotic system was implemented for a foraging task, in which the agent navigated an arena containing a variable amount of food. The agent's energy continuously decreased over time, and foraged for food when its energy dropped below a certain amount. The goal of the agent was to learn a foraging strategy that would allow it to survive for as long as possible. The strategy would include the time (in terms of energy level) at which the agent would begin foraging. The network consisted of sensory neurons, motor neurons, and interneurons. The modulatory component resided within the motor neuron that controlled eating. When this neuron fired, all interneurons were differentially affected. Neuromodulation could decrease the activation threshold of the postsynaptic neuron, cause the postsynaptic neuron to fire without input, or alter the response of the postsynaptic neuron to the crossing of its threshold (i.e., whether it becomes excited or inhibited). Because this modulatory motor neuron is not directly synapsed with any interneurons, the modulatory signal here may be considered somewhat global. The effects of modulation on each interneuron were predetermined and the robot was tested with varying amounts of food in the environment and compared to systems lacking neuromodulation. Results showed that the modulated robot survived longer than the non-modulated robot only when there was an intermediate amount of food. When food was scarce, modulation was not an advantage even though it allowed the robot to being foraging sooner. The authors hypothesized that this may be due to the fact that the longer it foraged, the slower it moved owing to inherent characteristics of the robot's architecture.

2.6 External modulatory signals

All of the implementations described thus far share the common feature that all modulatory signals are internal to the network itself or, at the very least, originate from another neuron. That is, the changes that a neuron undergoes result from the activity of other neurons. However, it may be possible for a neuron or synapse to be modulated not as a direct result of the firing of another neuron but as a result of some other type of signal. For example, global reward signals have been used to modulate spike-time-dependent synaptic plasticity (STDP) to achieve reinforcement learning. STDP is a type of synaptic plasticity that depends on the timing of the firing of the presynaptic neuron relative to that of the postsynaptic neuron. Florian (2007) argued that because Hebbian STDP (i.e., the strengthening of synapses when presynaptic activity occurs immediately before postsynaptic activity) could be used to establish particular input-output mappings, its modulation using a reward signal could lead to reinforcement learning. That is, if associations linked to desirable outcomes were strengthened by a positive reward signal and associations linked to undesirable outcomes were weakened by a negative reward signal, reinforcement learning would emerge. This was indeed accomplished by modulating weight update rules with positive as well as negative reward signals, leading to Hebbian and anti-Hebbian STDP, respectively. Farries and Fairhall (2007) also demonstrated similar results using modulatory reward signals. External, global reward signals have also been successfully used to control synaptic plasticity in a network implementing actor-critic learning, a type of reinforcement learning (Potjans *et al*. 2009).

2.7 Models of neuromodulation

Finally, as stated previously, the aim of this work is not to directly model biological neuromodulation, and the systems described thus far are not explicit simulations of biology. However, it may nonetheless be worthy to mention certain implementations of artificial neuromodulation that are designed to be more analogous to specific biological neuromodulatory systems, as they may help bridge the gap between the intricate biological process and an effective, generalizable artificial implementation.

For example, Cox and Krichmar (2009) developed a neuromodulated robotic controller inspired by specific modulatory systems in the animal brain. Their neuromodulatory system includes three centers modeled after the raphe nucleus (RN), the basal forebrain, and the ventral tegmental area (VTA), sites of 5-HT, ACh, and DA release, respectively. The authors asserted that 5-HT, ACh, and DA activity is associated with risk taking, attention, and reward anticipation, respectively, and that the release of these modulators, triggered by specific environmental stimuli, allows an animal to switch between exploitative and exploratory behavior. The artificial model follows this same line of thought. For example, when the robot encounters harmful stimuli, the RN would be activated, resulting in risk-aversion behavior. The neuromodulatory centers reside within the primary network, which contains visuomotor neurons, action neurons that allowed the robot to flee (move away from an object) or find (move towards an object), and behavior driver neurons (consisting of two groups of neurons signaling good and bad events) connected to the neuromodulatory system as well as action neurons. "Good" behavior drivers were set to strongly excite the VTA

and "find" action neurons and inhibit the RN and "flee" neurons, while "bad" drivers were set to carry out the opposite effects. Furthermore, modulatory signals alter synaptic input as well as synaptic plasticity in an additive fashion. The robot was tested in an enclosure with color-changing panels on the floor, and its task was to associate certain colors with certain behaviors (e.g. "red" with "flee"). Results showed that the neuromodulated robot was able to form associations between stimuli and actions, and simulated lesion experiments resulted in lowered performance. The authors attributed the success of this system to the ability of neuromodulation to increase the robot's focus on important stimuli by amplifying certain signals and suppressing others.

A less elaborate model using the same ideas was later developed (Krichmar 2012). In this model, behavior selection is once again driven by neuromodulation, which is in turn activated by certain sensory stimuli. The neuromodulatory system here included a DA neuron, a 5-HT neuron, and several ACh neurons. The rest of the network consisted of neurons signaling certain events (such as "low battery") and neurons for each possible behavior (such as "find home"). The possible events signaled either possible harm or possible reward and activated the appropriate neuromodulatory neurons, which in turn triggered either exploratory or risk-aversion behaviors. The agent was once again shown to be able to switch between exploratory or exploitative behaviors using neuromodulation. Furthermore, because ACh neurons were connected to event neurons through plastic synapses, the ACh system was found to be useful in gating events (i.e., decreasing attention to frequent events and increasing attention to rare or novel events). The latter system is a less detailed, and probably more generalizable, model of neuromodulation. Although both are somewhat literal

translations of biological neuromodulation, their implementation in robotic systems may provide insight on various artificial intelligence and robotics problems, such as action selection and switching between different types of behaviors.

3. Functional significance of neuromodulation in artificial neural networks

Before the implications of the experiments and results of the systems described above are examined, it may be practical to first define or characterize neuromodulation in an artificial setting as presented in this chapter. The neuromodulatory mechanisms described here involved either the modulation of the synaptic plasticity of the network (by targeting the weights or weight update rules) or the modulation of some intrinsic neuron property (by targeting the activation function). Consequently, these mechanisms introduced a dynamic element into the network. This allowed for the agent to alter its sensorimotor mapping in the course of its lifetime according to environmental or internal prompts. Thus, artificial neuromodulation may be broadly characterized here as a mechanism that endows agents with lifetime plasticity, or the ability to learn. Populations, therefore, are evolved not to simply react but also to adapt.

A distinction should be made at this point concerning two different methods of adapting to an environment: learning and evolution. The former occurs over individual lifetimes at the phenotypic level, while the latter occurs over generations at the genotypic level. Nolfi and Floreano (2000) have proposed that because any adaptation discovered through learning could, theoretically, also be discovered by evolution, one of the primary advantages of learning lies in its ability to smoothen the fitness landscape.

That is, the lifetime plasticity imparted by learning or neuromodulation facilitates the climb of an evolving population towards fitness peaks.

Learning agents with evolved local plasticity, in which the continual changes that individual synapses or neurons undergo are due to predefined synaptic or neuron properties rather than to modulatory signals from other neurons, have shown success in adapting during individual lifetimes. For example, Floreano and Mondada (1996) developed networks in which each synapse was individually and continually modified according to its own evolved learning rule, while Di Paolo (Di Paolo 2000) developed a neurocontroller in which each synapse was individually modified according to its own evolved rule whenever the firing rate of associated neurons were outside of a certain range. The plasticity exhibited in these systems, in contrast with the plasticity described above, is controlled at a strictly local level. That is, the changes that occur at a particular synapse or neuron are not directly influenced or evoked by another neuron. Neuromodulatory signals, on the other hand, underlie heterosynaptic regulatory systems that likely allow for the detection of environmental characteristics or changes that cannot be recognized at the local level.

In at least in one instance, the advantage of lifetime learning imparted by neuromodulation cannot be explained in terms of fitness landscapes. Of the neuromodulatory systems described above, the success of GasNets have been studied most extensively. The results of these studies do not support the idea that neuromodulation (as implemented in GasNets) simply modifies the fitness landscape such that good solutions are easier to find (by "sampling" nearby genotypes). One study found no measurable differences between the search spaces of GasNets and the

same type of network with gases disabled (NoGas networks) (Smith *et al.* 2001). Random sampling of solutions from the search spaces of networks with and without gases showed no significant differences in terms of the number of good solutions, landscape ruggedness, or landscape modality (a measure of the number of local optima) (Smith *et al*. 2001). Instead, it has been proposed that a primary reason behind the increased evolvability of GasNets is its temporal adaptivity, or the ability to integrate environmental signals and operate over different timescales (Smith *et al.* 2002). In an analysis comparing GasNets and NoGas networks evolved for the target discrimination task described in section 2.1, the success of GasNets was attributed to its ability to easily generate a timing mechanism that distinguishes between the two different shapes (for which the duration of specific input signals differed) and switches the network between stable states over multiple time steps (Smith *et al.* 2002). GasNets were able to accomplish this through the gradual buildup and decay of emitted gases, while the timing mechanism found in NoGas controllers were more complex and less intuitive. It was argued that the GasNet timing mechanism is more easily modified during evolution through simple adjustments of modulation parameters, and thus GasNets are more adaptable to specific temporal characteristics of the environment. This was further supported in the same study by a considerable increase in evolvability of GasNets over NoGas networks when they were re-evolved in altered environments with different temporal characteristics.

Another possible role of neuromodulation is to allow neurocontrollers to function in novel environments. In biology, neuromodulation confers a greater degree of adaptability to a nervous system by allowing it to change in accordance with the

environment. In artificial systems, adaptability would be desirable in dynamic environments or in instances in which the real-life problems that networks are tasked to solve differ greatly from the example problems that were encountered during evolution or training. Such discrepancies between the training environment and the real environment may significantly limit the functionality of the network. This "reality gap" has been characterized in evolutionary robotics (Jakobi *et al.* 1995), in which neural networks are employed in actual robotic agents that must navigate real-life environments. Thus, aside from increasing evolvability, a function for neuromodulation may be to provide a mechanism by which artificial networks can achieve robustness in novel environments. Researchers of DRNNs have stressed the crossing of the reality gap as a function for neuromodulation (Kondo *et al.* 1999).

4. Conclusions

Because the research studies discussed in this chapter vary widely in many aspects (e.g., the types of learning tasks networks were evolved for, the structures that networks adopted, whether the network-controlled agents were real or simulated, the evolutionary schemes used, how networks were tested, and controls used in the tests), direct comparisons between systems may not be practical. However, a general conclusion that can be drawn here is that the switching between strategies or behaviors in response to environmental input that is facilitated by neuromodulation, here broadly defined as a mechanism for lifetime learning, is central to increases in performance. In GasNets, it was found that switching between stable states over several time steps, which was crucial to its success in the particular task studied, was easily accomplished by diffusing

gases. Switching between different sensorimotor mappings as a result of modulatory signals was noted in ESCNs (Ziemke and Thieme 2002). Soltoggio *et al*. (2008) and Parussel and Smith (2005) both noted that neuromodulation allowed the transition between explorative and exploitative behavior. These strategies are important for learning, presumably (and especially) in unstable environments in which the agent must set out to acquire new information (explore) instead of using old behaviors to maximize rewards (exploit).

In some cases, it was found that the evolution of successful controllers does not necessarily require neuromodulation, since it may well be that neuromodulated networks are computationally equivalent to non-modulated networks. However, neuromodulation allowed for a more rapid or efficient search for good solutions and a reduction in the design complexity of successful networks. Biological neuronal signaling involves not only electrical impulses in the form of action potentials, which is the instantaneous type of communication emulated in traditional, non-modulated neural networks, but also involves chemical signals that, at times, induce long-term changes, thereby introducing an additional component or dimension to neuronal communication. A non-modulated network (having just one dimension in signal transmission) that is functionally equivalent to one that is modulated would likely have to be much more elaborate in design, and thus be more difficult to evolve.

Studies of GasNets have suggested that the degree to which the two signaling systems in a network (normal "electrical" signals and neuromodulatory signals) influence each other may affect its success. Specifically, it has been found that a flexible coupling between neurotransmission and neuromodulation is desirable over a

tight coupling in which altering one system necessarily alters the other. The increased evolvability of the GasNet variants described in section 2.1 was attributed to a decoupling or loosened coupling between the two systems. That is, through evolution, gas modulation could be physically separated by some distance from emitting nodes (in the plexus GasNets), while receptors on a node could limit the amount of modulation it undergoes (in the receptor GasNets). The authors suggested that the loose coupling between interacting systems renders a network phenotypically stable (i.e., genotypic changes do not necessarily lead to phenotypic changes) and thus more evolvable.

Finally, while neuromodulated networks usually outperformed non-modulated networks, neuromodulation was shown to be an undesirable feature in at least one instance in which change and adaption were not necessary (Dürr *et al.* 2008). Neuromodulation seemed to cause disturbances in the system, facilitating alterations between behaviors. Therefore, in certain instances, such as when operating in environments that are relatively static, such alterations may be unfavorable. Because the complexity of biological neural networks is not paralleled by neurocontrollers, which have relatively simple sensorimotor mappings, the continuous modulation of biological networks should perhaps be mirrored only to an extent in artificial systems, and networks should thus have the ability to switch off or adjust any evolved neuromodulatory features. For example, the design of the ESCN explicitly incorporated a switch to turn neuromodulation on and off (Ziemke and Thieme 2002).

5. Summary

Some of the neuromodulated neural networks described in this chapter were inspired by very specific mechanisms or systems of biological neuromodulation (i.e., NO signaling and neuromodulation in the crustacean STNS), while others emulated the general effects of biological neuromodulation. While these systems differed greatly in the details of implementation, the primary effect was the continual adjustment of synaptic plasticity or of transfer functions over the lifetime of the neurocontroller. Explanations for the increased performance demonstrated by neuromodulated systems over their non-modulated counterparts have been given at various levels of abstraction. Although each computational task and the specific type of ANN employed for each task differed, the keys to the success seen in the systems described here may still provide insight on very general elements that may be advantageous or disadvantageous in any implementation of neuromodulation. Here, neuromodulation has been mechanistically and broadly defined as a means by which lifetime plasticity is achieved. Finally, although biological neuromodulatory mechanisms may regulate certain types of synaptic plasticity, neuromodulation encompasses a wider range of effects (see chapter one). That is, there are other ways in which network plasticity could conceivably be achieved. The implementation of these other effects in an artificial setting will be explored in the next chapter.

Chapter 4: Neuromodulation in artificial neural network

1. Introduction

Biological neurons and the circuits that they constitute may adopt different roles in different circumstances due to neuromodulation. Neuromodulation increases the scope of functioning of neurons by granting them the ability to shift between different functional states on distinct neurochemical cues. Artificial neural networks, which mostly do not incorporate neuromodulatory processes, could potentially benefit from similar increases in functionality. In chapter three, neuromodulation within the context of artificial systems was broadly defined as a mechanism to introduce lifetime plasticity. It was shown, albeit through a limited variety of implementations, that the use of a signaling system separate from that of "classical" neurotransmission to introduce this type of plasticity generally led to improvements in evolvability and robustness.

In this chapter, variations on the theme of incorporating an additional signaling system to augment functionality will be explored. Due to inherent differences between biological and artificial neural systems, it is likely not appropriate or feasible to attempt a modeling or direct transfer of specific concepts from a biological to an artificial context. Biological neuromodulation may be generally characterized by signals that are often diffuse and have slow-acting and long-lasting effects (Katz 1999). The importance of these fundamental features suggests that they should serve as the starting point of the development of artificial neuromodulatory systems.

Due to the enormous variety of artificial networks, this chapter will broadly present some general ways in which the neuromodulatory process may be introduced

into an artificial setting, rather than address a particular type of network or emulate a specific type of biological neuromodulatory system.

2. Neuromodulatory signals

2.1 Nature of modulatory signals

In biological systems, neuromodulatory communication is typically not solely distinguishable from neurotransmission by the signals that are involved. Both types of communication involve the release of the same kinds of chemical messengers, while the reception of these messengers determines their effects. Receptor systems can also be implemented such that the effect of a signal is determined by postsynaptic nodes (fig. 2). In such a scenario, one signal may have different effects on different neurons depending on a neuron's collection of receptors. Receptors may determine whether neurons undergo modulation or how they will be affected by modulation.

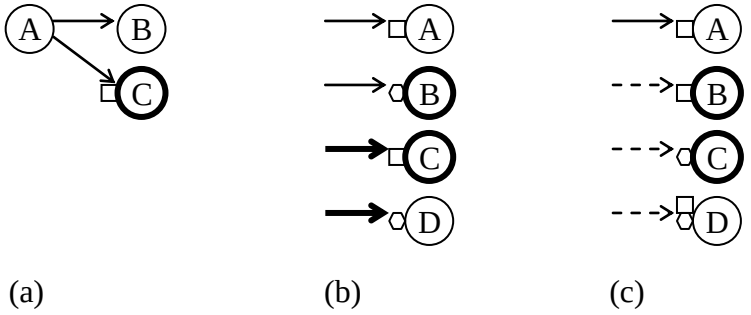

Figure 2. Receptor systems. Neurons with a bold outline are modulated. Arrows in (a) and (b) indicate signals that may be standard or modulatory depending on the postsynaptic neuron. In (c), the solid arrow indicates a standard signal, while dashed arrows indicate modulatory signals. Neurons transmit variable types of signals, while receptors determine the effects. (a) Neuron C, which possesses a receptor (square) that mediates modulatory effects, is modulated by neuron A. Neuron B receives standard input from neuron A. All neurons transmit standard signals, while receptors determine whether such signals also modulate their neurons. This is analogous to a biological

network using neurotransmitters that activate ionotropic receptors and, in some cases, metabotropic receptors. (b) This is an extension of the situation shown in (a). Here, m ultiple receptor and signal types exist; each receptor determines the type of signal that will modulate its neuron. The type of receptor belonging to neurons A and C (square) treats one type of signal (regular arrow) as a modulatory signal, while the receptors belonging to neurons B and D (hexagon) treats the other type of signal (bold arrow) as modulatory. In this case, neurons B and C will undergo neuromodulation. This is analogous to a biological network using neurotransmitters that activate ionotropic receptors in some neurons and metabotropic receptors in other neurons. (c) Neurons transmit either standard signals or modulatory signals; receptors determine the effect of modulation. Neuron A will not undergo neuromodulation, while neurons B and C become modulated with different effects. It may be possible for some neurons (D) to possess more than one type of receptor. Depending on the types of effects, multiple modulatory signals may completely cancel each other out (e.g., increasing and decreasing a neuron's threshold by the same magnitude), result in an additive effect (e.g., multiplying a neuron's output by different values), simultaneously modulate a neuron without any effects on each other (e.g., increasing threshold and amplifying output), or result in an intermediate effect (e.g., reducing a neuron's firing rate between two values set by two different modulatory neurons).

With the absence of receptor systems in artificial networks, signals that modulate neurons must be distinguishable in some manner from other signals. For example, some subset of values within the range of possible node outputs may be set aside for neuromodulation (fig. 3a). The postsynaptic neuron may also determine whether a signal is modulatory (figs. 3b and 3c). Modulatory signals may also be distinguished by their source. That is, certain nodes may be designated as modulatory neurons. In some instances, the numerical value of the signal may be used to modulate the postsynaptic unit (fig. 4a). Signals can also be non-numerical. For example, signals could be flags for some postsynaptic modulatory event to take place (fig. 4b).

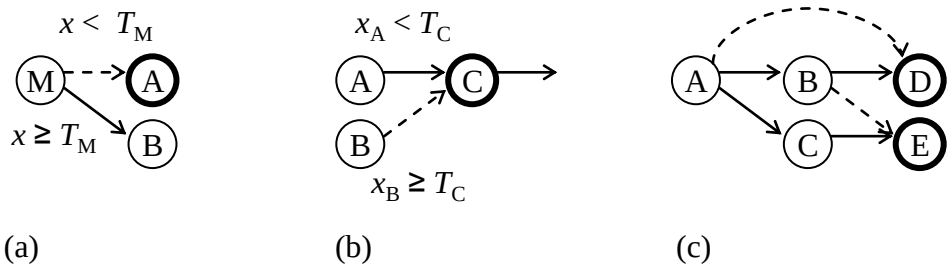

Figure 3. Distinguishing modulatory signals. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. (a) Outputs (x) from neuron M are modulatory until a threshold (T_M) is reached. Thus, neuron M modulates neuron A but not neuron B. (b) A modulatory threshold (T_C) is set by neuron C. This value determines whether incoming signals are modulatory. Values below T_C are not modulatory. Thus, neuron C is modulated by neuron B but not by neuron A. (c) The location of incoming signals determined whether signals are modulatory. Neuron D is modulated by signals more than one layer away, while neuron E is modulated by signals from the previous layer.

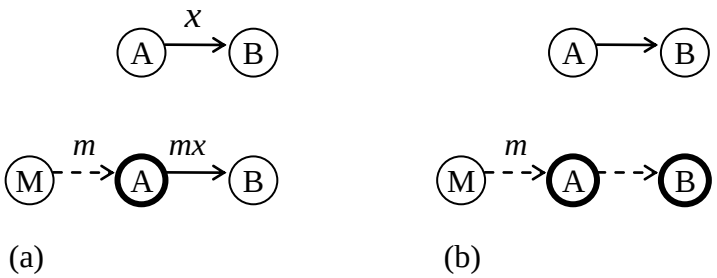

Figure 4. Using modulatory signals. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. (a) Neuron M modulates neuron A by multiplying the output of neuron A (x) by its output value (m). (b) Before (top) and after (bottom) modulation. The output value of the modulatory neuron M (m) is not directly used. Neuron M modulates neuron A by changing the type of connection between neurons A and B from standard to modulatory.

2.2 Source and range of modulatory signals

Intuitively, the most straightforward approach to establishing neuromodulatory sources may be to simply designate certain neurons as neuromodulatory. However, it may detract from design simplicity and be functionally redundant for a system to divide different functions into entirely separate sets of neurons. The alternative to this type of

design entails the delivery of both types of signals from one type of neuron in addition to a system that determines the timing and frequency of neuromodulation. Neurons may be prompted to release modulatory signals under certain conditions or release both types of signals at all times (fig. 5). In the latter case, the degree of modulation may be influenced by other variables such that neurons are not continually modulated at every time step.

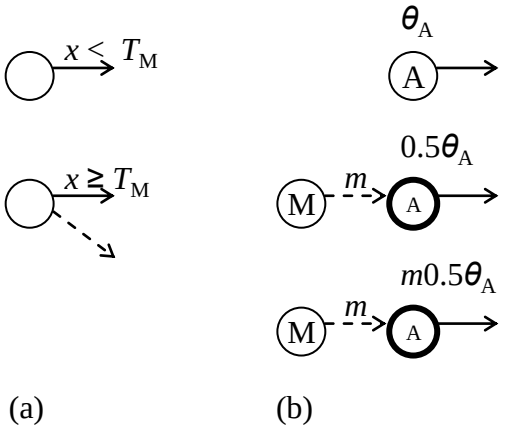

Figure 5. Release of modulatory signals. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. (a) Neurons may release modulatory signals intermittently. The output (x) of this neuron becomes neuromodulatory above a threshold value (T_M). (b) Neurons may continuously release modulatory signals. Before (top) and after (middle and bottom) modulation. Neuron M normally modulates neuron A by reducing its firing threshold (θ_A) by a fixed value (middle). However, to prevent neuron A from undergoing the same degree of modulation at every time step, the modulated threshold is multiplied by the output of neuron M (m) (bottom).

Neurotransmission in ANNs is usually limited to just the presynaptic and postsynaptic nodes. The range of modulatory signals could be similarly restricted (fig. 6a), but its effects on the performance of the whole network would likely be less appreciable than those produced by more diffuse signaling systems, as is seen in biological neuromodulation. Conversely, a fully global modulatory system may generate excessive disruptions in the network such that the results may be equally

ineffective. In such a system, it may be beneficial for each modulatory signal to differentially affect different neurons. For example, a modulatory signal might decrease as it travels away from its source (figs. 6b and 6c). Global signals may also be summed such that their modulatory effects are proportional to the sum, or may only modulate the network when the sum reaches some threshold (fig. 7). The value of the global signal may also determine its target destination (fig. 8).

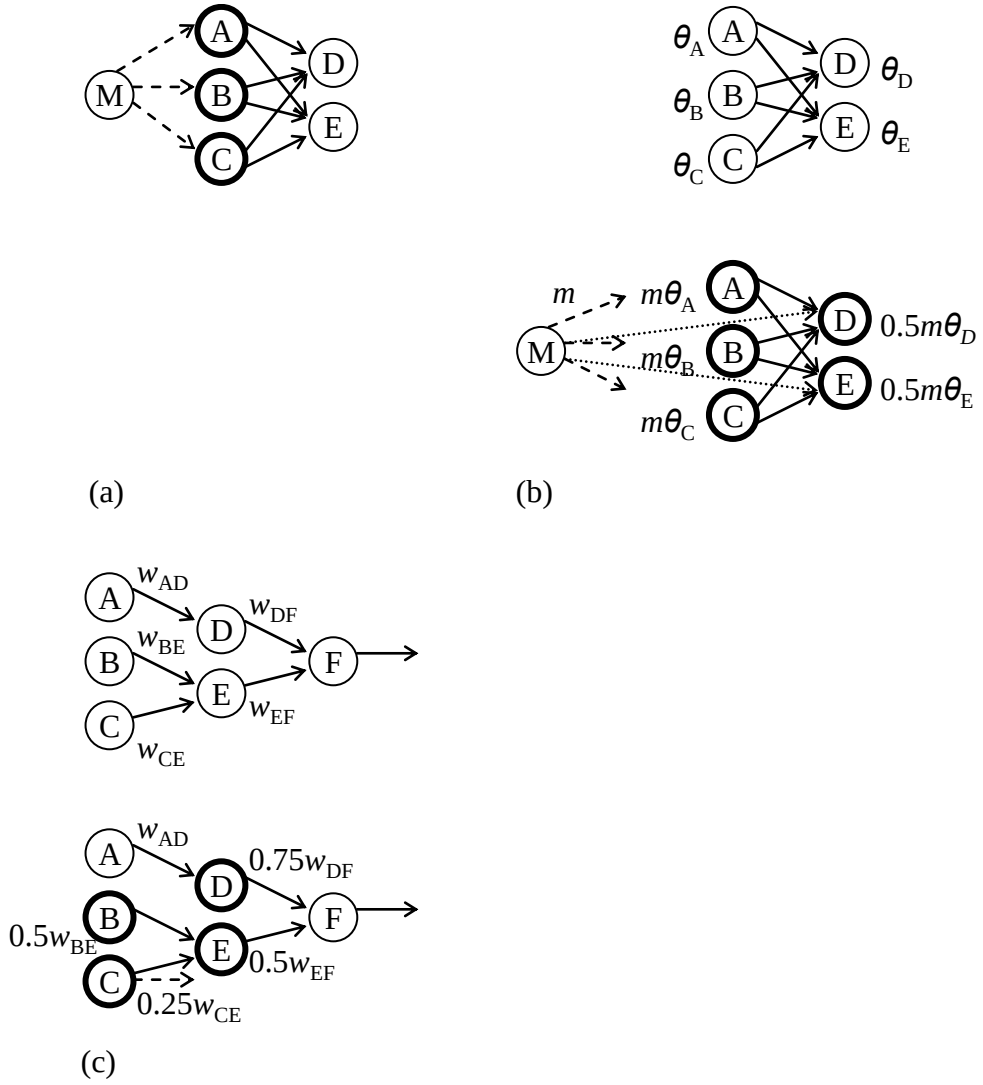

Figure 6. Range of modulatory signals. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. (a) Modulatory signals may only have local effects. Neuron M only modulates neurons with which it is directly synapsed (neurons A-C). (b) Modulatory signals may have a

wider range. Before (top) and after (bottom) modulation. Neuron M modulates neurons A-E but is not directly synapsed with neurons D and E. That is, neuron M transmits standard signals in addition to modulatory signals to neurons A-C, while neurons D and E receive only modulatory signals from neuron M. To prevent a network from becoming overly modulated, the modulatory effect of neuron M on neurons D and E (multiplying thresholds θ_D and θ_E by $0.5m$) is less severe than that on neurons A-C (multiplying thresholds θ_A-θ_C by m). (c) Modulatory signals may decrease as they travel away from its source. Neuron C emits a modulatory signal that decreases synaptic weights. The greatest decrease occurs at the synapse between the emitting neuron and its postsynaptic node (w_{CE}), while the synapse between neurons A and D (w_{AD}) is too far away to be affected. Effects on the other synapses (w_{BE}, w_{EF}, and w_{DF}) are proportional to their distances from neuron C.

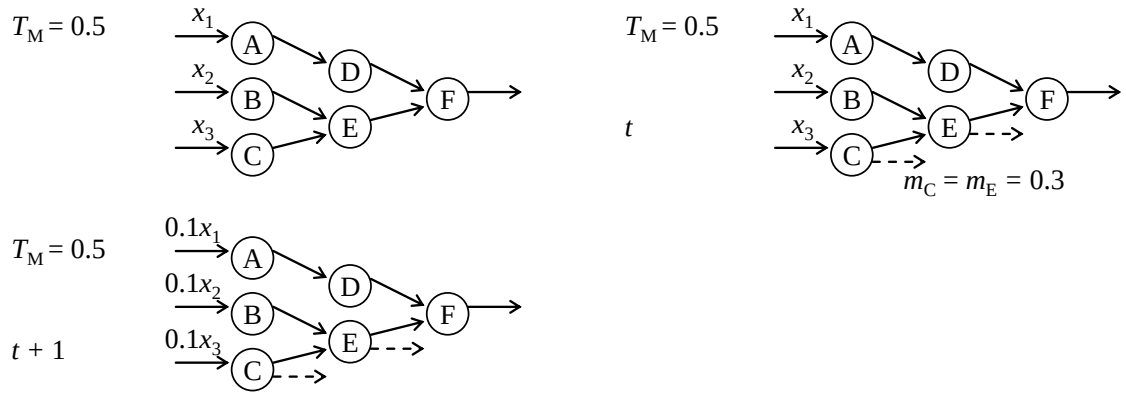

Figure 7. Summing modulatory signals. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top left) and after (top right and bottom) modulation. Global modulatory signals may be summed such that modulation does not occur below a threshold and the effect is proportional to the magnitude of the summed signal. At time t, neurons C and E emit modulatory signals m_C and m_E. The sum of these signals exceeds the modulatory threshold T_M (top right). The effect of modulation here is to multiply external inputs (x_1-x_3) by the difference between the sum of the modulatory value and T_M at time $t + 1$ (bottom left).

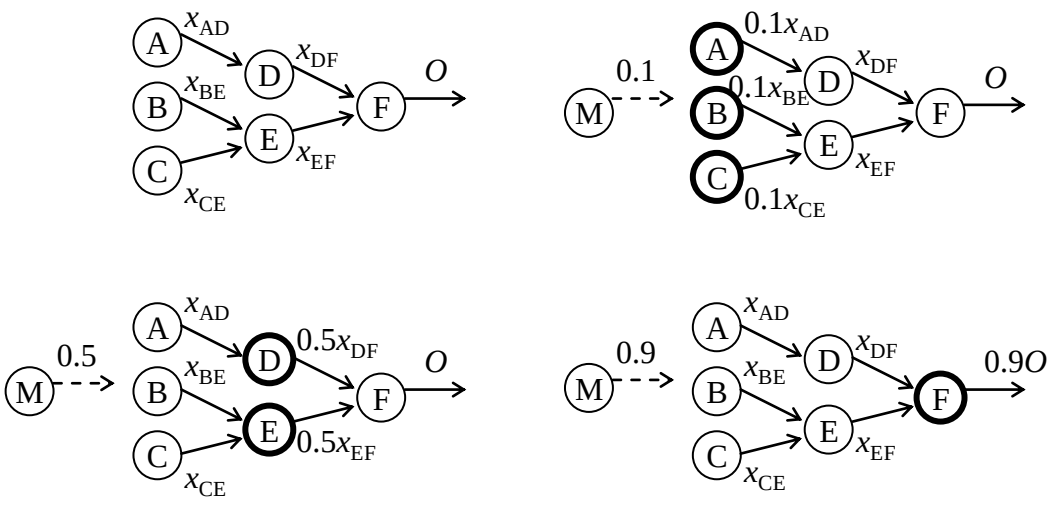

Figure 8. Value of modulatory signal determines target destination. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top left) and after (top right and bottom) modulation. The value of a non-local modulatory signal may determine its target. Here, the effect of the signal emitted by the modulatory neuron M is to multiply output values (x). Outputs from the input layer are affected if the signal is within [0, 0.5), while outputs from the middle layer are affected if the signal is within [0.5, 0.9). The output of the network (O) is affected if the modulatory signal is within [0.9, 1].

3. Neuromodulatory effects

Neuromodulatory signals should alter the functioning of nodes or the network itself in a way that is distinct from an ANN's principal form of communication: the instantaneous reception of input and generation of output at each node. Ultimately, neuromodulation changes the input-output mapping of the network, and such changes may have slow-onsets or long durations.

Figure 9 categorizes possible areas in a network that may be targeted by neuromodulation. These general targets include the operations within a single neuron, the reception and transmission of signals for a single neuron, synapses between neurons, and the network as a whole.

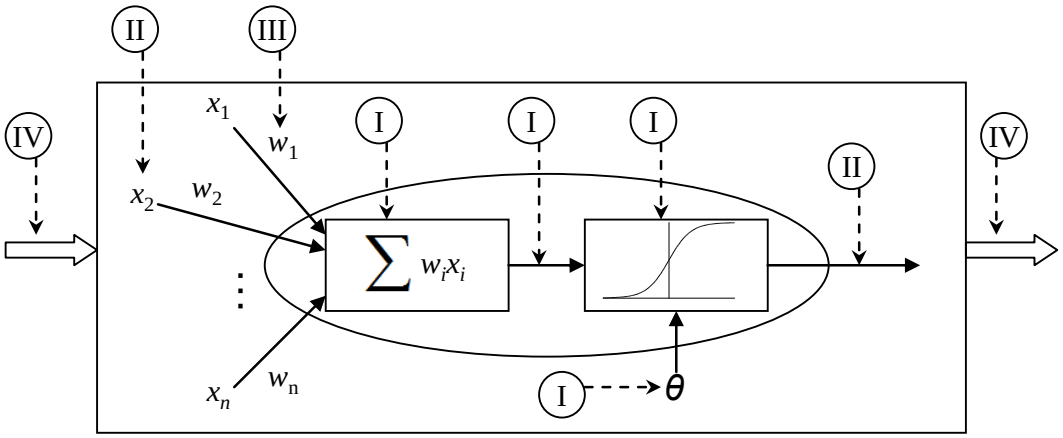

Within neuron (I)	**Neuron (II)**	**Synapse (III)**	**Network (IV)**
Input integration	Input gating	Weights	Neuron groups
Input to activation function	Output pattern	Weight update rules	Network structure
Activation function	Output value	Synapse type	Data flow
Thresholds			Input/output

Figure 9. Possible targets of neuromodulation. The large oval indicates a neuron containing an integration function, an activation function, and a threshold (θ). The integration function sums weighted inputs w_1x_1 to w_nx_n and passes the result to the activation function. This neuron resides within a larger network (outer rectangle), the input and output of which are indicated by hollow arrows. Dashed arrows point to possible areas (I-IV) that may be neuromodulated. Major targetable components are listed under each category.

3.1 Modulation of neuronal properties

The primary ways in which the systems described in chapter 3 altered individual neurons involved changes to connection weights or the way weights were updated at each time step. Modulation in GasNets involved alterations to the activation function of each node by modifying a variable within the function parameter. Another possible method of altering the input used by each function is to modify the method of input integration (fig. 10), which typically involves taking the weighted sum of all of inputs. Modulation could alter the operation performed on the set of inputs, for example, by taking the min, max, or median of the set. Other operations, such as an arithmetic,

quadratic, or geometric mean, or the square root of the sum of squares, could also be computed. Activation functions may also be modulated by directly changing the function itself. The commonly used logistic function, for example, may be transformed into another function within the family of sigmoid functions in which it resides.

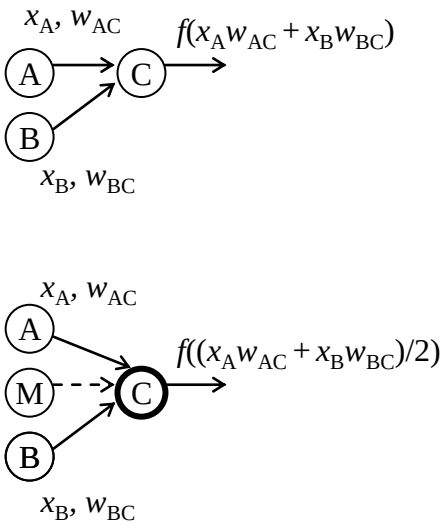

Figure 10. Modulating input integration. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. Neuron M modulates neuron C such that the weighted inputs from neurons A and B ($x_A w_{AC}$ and $x_B w_{BC}$, respectively) is averaged instead of summed.

Other, less direct characteristics of individual neurons may also be targets of neuromodulation. Firing patterns, for example, are not explicit attributes of neurons due to the discrete time steps over which ANNs operate. That is, all output neurons necessarily fire at every time step. Conceivably, this pattern could be modified to introduce periods during which a neuron would not generate outputs (i.e., generate an output of zero). Other possible neuronal properties that may be modifiable involve the gating of inputs (figs. 11a and 11b). The sensitivity of a neuron to input, for example, could be modified through the threshold of a neuron, which is a predetermined value above which the neuron would fire. Other types of thresholds, such as a minimum

number of inputs received simultaneously or over a number of time steps, may also be established and modulated. Input sensitivity may also be altered by limiting the domain of a neuron's activation function (fig. 11c).

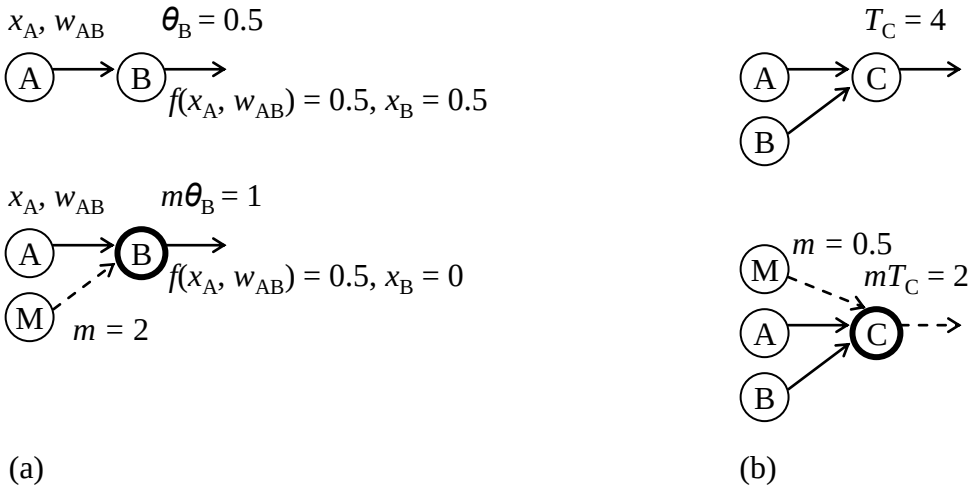

Figure 11. Modulating the gating of inputs. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. (a) Neuron M modulates neuron B by increasing its firing threshold (θ_B) such that an activation value of $f(x_A w_{AB}) = 0.5$ no longer results in neuron B firing ($x_B = 0$). (b) Neuron M modulates neuron C by multiplying the minimum number of inputs it must receive to become modulatory (T_C) by the modulatory signal (m). Following modulation, neuron C becomes a modulatory neuron due to a decrease in T_C. (c) Neuron M modulates neuron B by limiting the domain of its activation function (f) to values less than one. Following modulation, neuron B produces an output of zero.

The output of a neuron may also be directly amplified or attenuated. For example, a neuromodulatory signal may act as a simple multiplicative factor on neuronal outputs (fig. 4a). Output signals may be also be dampened by splitting a neuron's output value among all postsynaptic neurons (fig. 12). Furthermore, a modulated neuron may alter its target to affect unconnected neurons (fig. 13). Normally, a neuron transmits the same value to all of its postsynaptic neurons. A modulated neuron could compute different signals for different postsynaptic neurons (fig. 14).

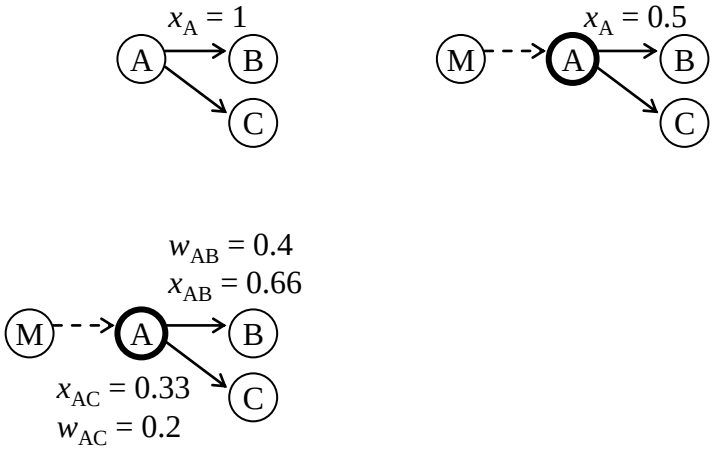

Figure 12. Splitting output values. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top left) and after (top right and bottom) modulation. Neuron M modulates neuron A by splitting its output value (x_A) between its postsynaptic neurons. The output value may be split evenly (top right) or unevenly (bottom). In the latter case, the division of x_A (into x_{AB} and x_{AC})is proportional to the weights of the associated synapses.

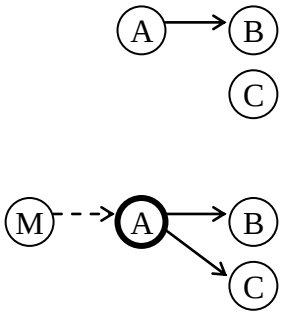

Figure 13. Modulating neuronal targets. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. Neuron M modulates neuron A by adding a connection between neurons A and C.

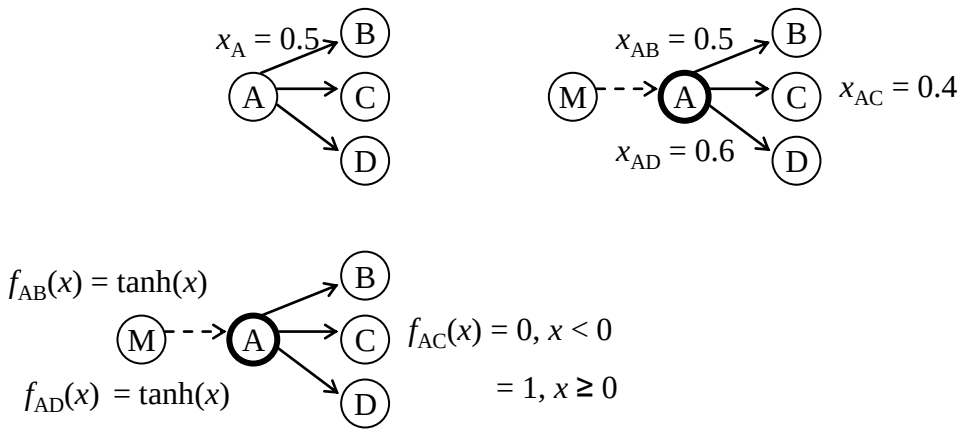

Figure 14. Modulating output values. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top left) and after (top right and bottom) modulation. Neuron M modulates neuron A by altering its output values such that different values are transmitted to different postsynaptic neurons. This may occur by computing a single output value (x_A) and choosing a random value within a range ([0.4, 0.6]) set around the output value for each postsynaptic neuron (top right). A neuron may also transmit multiple values by alternating between different transfer functions (bottom).

Synapses between neurons may also be modulated. Many examples of modulating connection weights, for example, were presented in chapter 3. The type of connection between two neurons is also another possible target for modulation (fig. 15).

Figure 15. Modulating synapses. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (left) and after (right) modulation. Neuron M emits a global signal that acts on neurons with at least two inputs and changes the type of one of these synapses such that neuron B now modulates neuron D. This may be analogous to the neuromodulatory activation of silent receptors or the generation of new receptors in biological neurons.

Certain types of neurons or neurons in certain layers can be specifically targeted. For example, input neurons may be modulated such that environmental signals are filtered or amplified. A neurocontroller with various environmental sensors, for instance, might block light input below a certain level of brightness or ignore objects beyond a certain distance. Furthermore, an input neuron could exchange input values between itself and another neuron, thereby effectively rerouting environmental signals to alternate "sensory neurons". Output neurons could undergo similar types of modulation by allowing only certain neurons to fire or funneling signals to certain "effector neurons".

3.2 Modulation of multiple neurons and network function

Groups of neurons may also be simultaneously modulated. The firing patterns of a group of neurons, for instance, may be synchronized, or connections within the group could be concurrently adjusted (fig. 16). Neurons may be grouped for the purposes of neuromodulation in a number of ways. Groups may be formed by neuron type, such as sensory or motor neurons, neuron location, such as the layer in which a neuron resides, or the proximity of a neuron to the source of a modulatory signal, such as distance measured in terms of the number of neurons or layers.

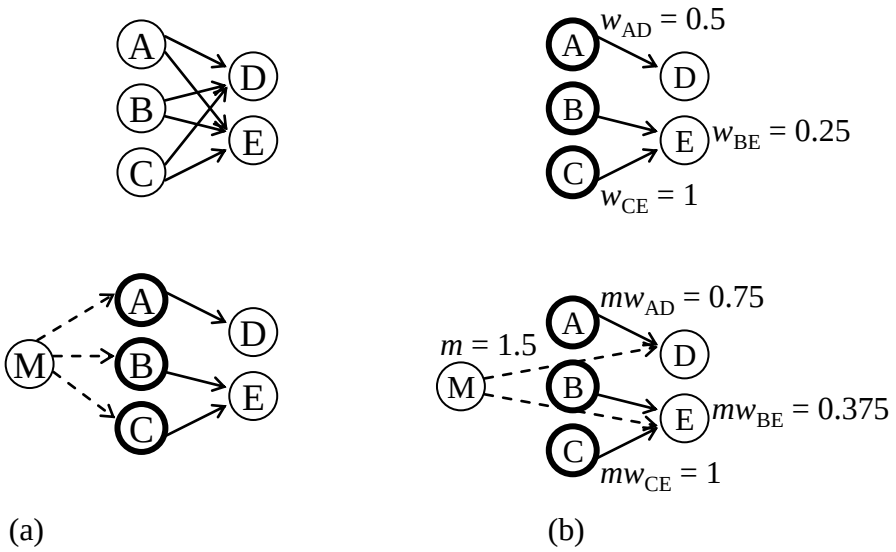

Figure 16. Modulating groups of neurons. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. (a) Neurons immediately postsynaptic to neuron M become modulated. An output connection is removed from neurons in this layer with at least two output synapses. Essentially, this has the same effect of reducing synaptic weights to zero, which may occur naturally over time. However, this allows for a much more rapid functional rearrangement of the network, which may be useful in certain types of environments. (b) Neuron M modulates neurons two layers away (neurons D and E) by multiplying weights between these neurons and their presynaptic nodes (w_{AD} and w_{BE}) by the modulatory output (m). The upper limit on weights is 1, thus w_{CE} is not altered.

Neuromodulatory signals, particularly ones that function globally, may also alter the functioning of the network as a whole. Multiple groups of neurons may be differentially targeted, perhaps in some extension of the manner described above. The network may also be modulated on a more abstract level by altering the flow of data (fig. 17). For example, various temporal dynamics could be introduced to a basic network in which information travels instantaneously by retaining values at synapses to be used at later time steps. Furthermore, as signals flow down the network, the magnitude of their values or the courses by which they travel are also possible targets of modulation. The structure of the network could also be altered, either by local or global

modulatory signals, by adding and removing neurons (fig. 18) or splitting the network into components (fig. 19). For example, a modulated neuron or an adjacent neuron could be temporarily removed from the network, or a global signal could partition the system into functionally distinct subnetworks.

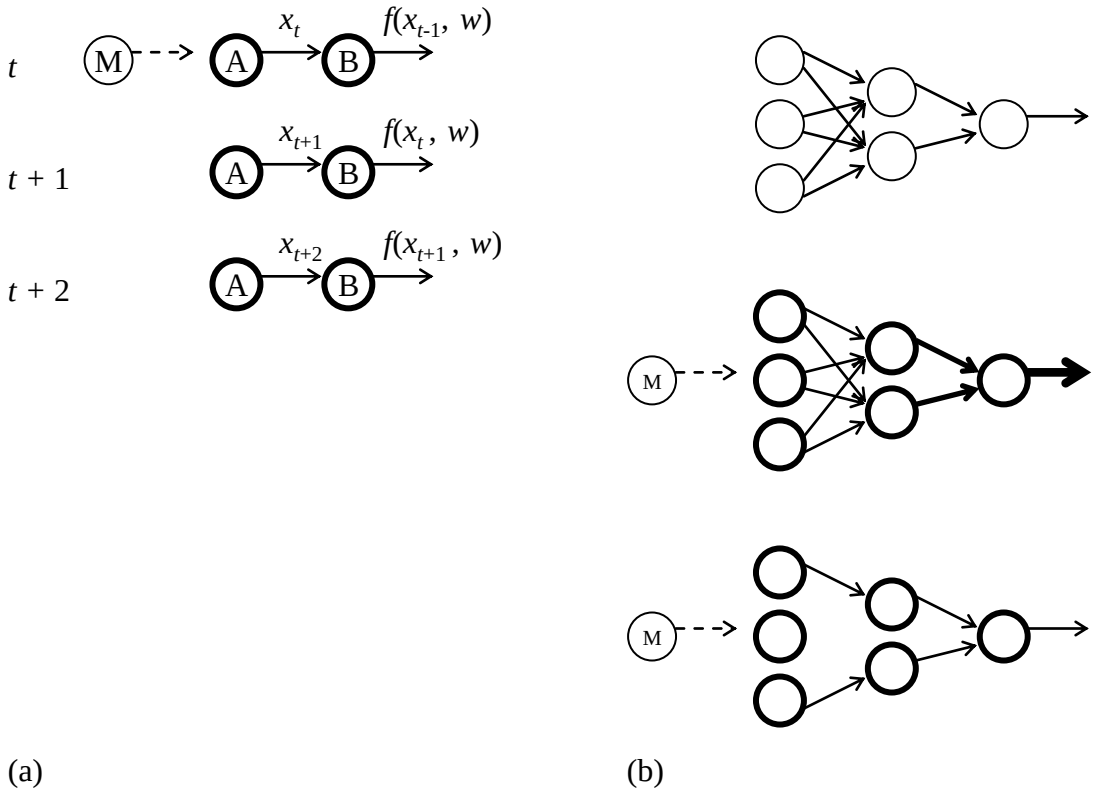

Figure 17. Modulating data flow. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. (a) At time *t*, neuron M modulates the network, causing the output of neuron A to be used at the next time step. At time *t*, the input for neuron B ($f(x_{t-1}, w)$) comes from the output of neuron A at time *t*-1. This design requires the storage of all output values from the previous time step. (b) Before (top) and after (middle and bottom) modulation. Neuron M causes signals to become amplified (indicated by darkened arrows) as data flow down the network (middle). Network dynamics may also be modulated by changing the course of signals by limiting data flow to neurons on the outside "edges" of the network (bottom).

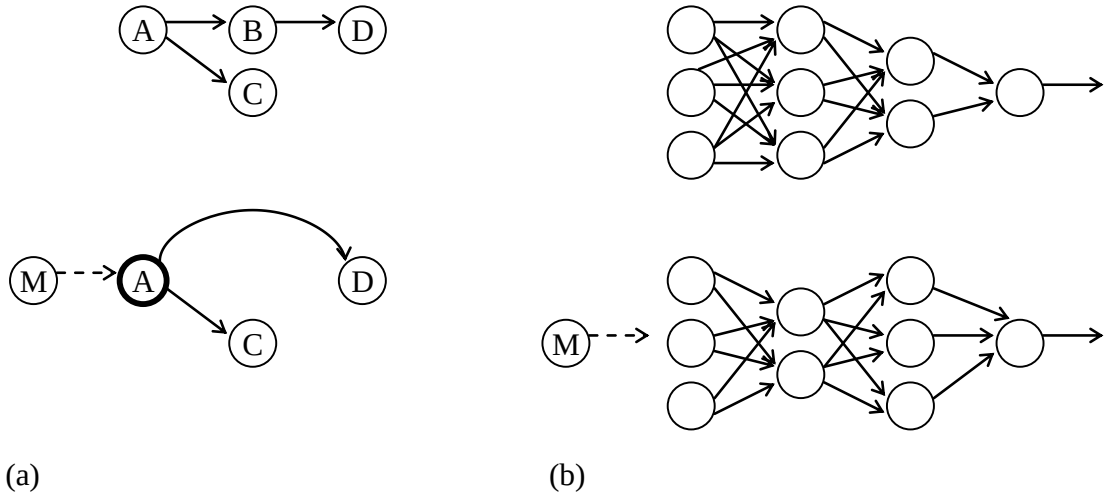

Figure 18. Adding and removing neurons. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Before (top) and after (bottom) modulation. (a) Neuron M modulates neuron A by removing one of its postsynaptic neurons (B). This has the same effect as setting certain weights (w_{AB} and w_{BD}) to zero. Neuron D is reintroduced into the network by adding a synapse between neurons A and D. (b) A modulatory signal removes a neuron from one layer (the layer postsynaptic to the input layer) and adds a neuron to another (the layer presynaptic to the output layer). This type of modulation may be restricted to areas of the network such that the removal of a neuron does not completely halt the flow of data.

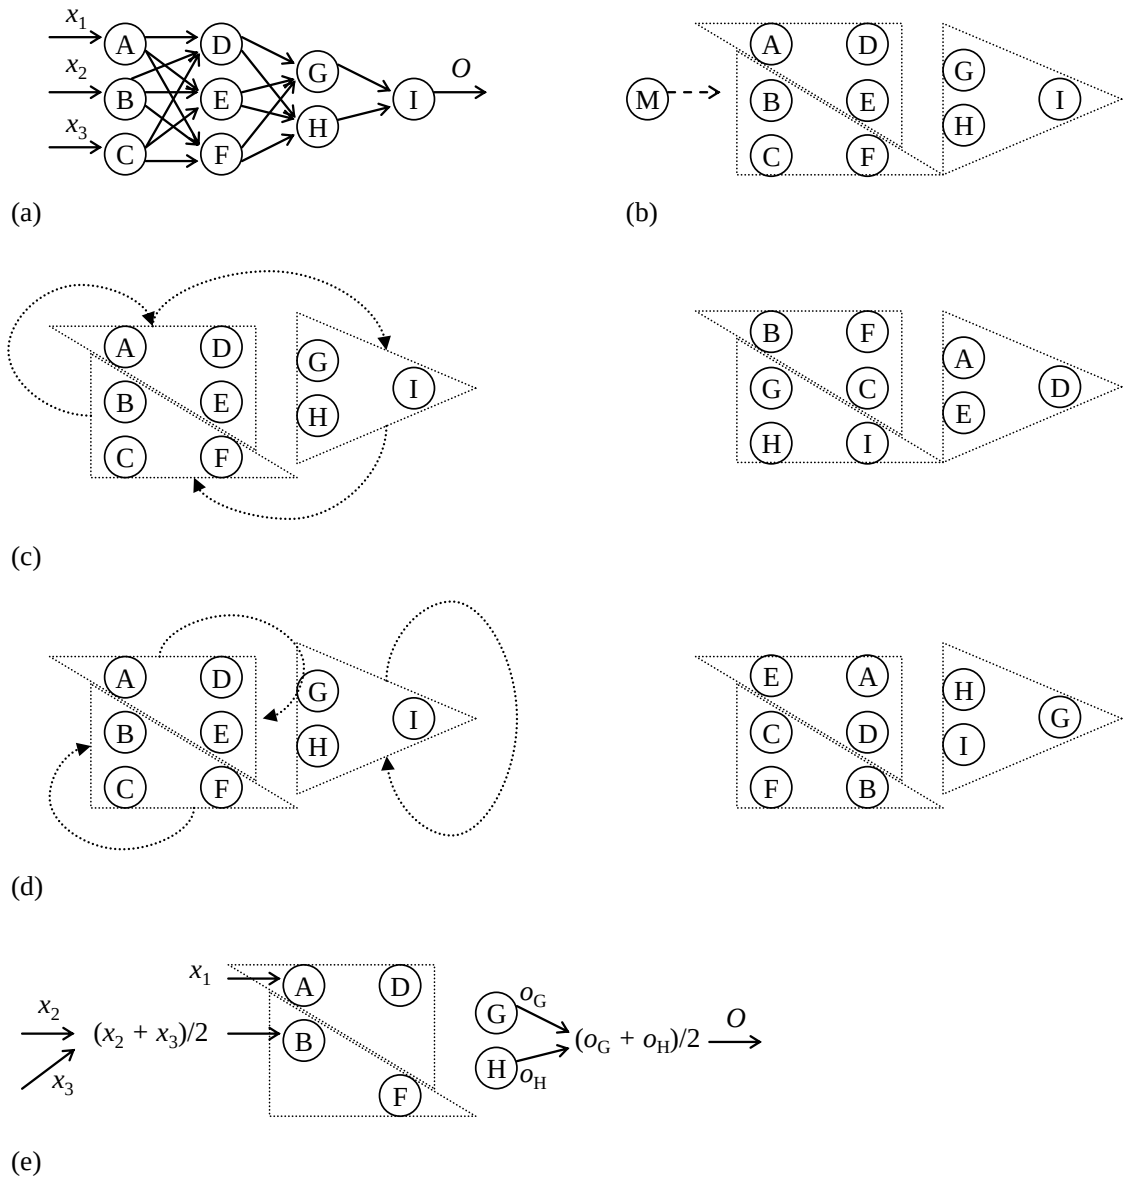

Figure 19. Modulating network structure. Neurons with a bold outline are modulated. Solid arrows and dashed arrows indicate standard and modulatory signals, respectively. Modulation may split the network into distinct subnetworks. (a) A network receiving three external inputs (x_1-x_3) and transmitting one output (O). (b) Neuron M splits the network into three parts (connections have been removed for clarity). (c) Before (left) and after (right) modulation. Assuming that individual neurons have distinct properties (e.g., they are undergoing different modulatory effects or utilize different activation functions), neurons may be relocated. Groups of neurons may be switched such that all of the synapses remain in the same location relative to the network but "reconnect" with new neurons. (d) Neurons within each group may relocate within the group. (e) One neuron is removed from each group. Because the number of external inputs and network outputs must remain unchanged, the input and output values may be combined. x_3, which is normally received by neuron C, is now averaged

with x_2 and sent to neuron B. Similarly, output values from neurons G and H are averaged to produce a single network output.

4. Control of neuromodulation

The slow-onset effects of neuromodulation should not be ignored in designing a neuromodulatory system, as this characteristic may provide some temporal flexibility in the functioning of the network. Modulatory signals, if given a period of time over which they will build up to their full effect, may allow a network to adopt a wider range of input-output mappings or to operate over various timescales. The effect of a modulatory signal could be either entirely delayed or gradually reach its full magnitude by a variable number of time steps. For example, the threshold of a neuron could be gradually decreased or an increasing portion of the domain of an activation function could be limited over several time steps.

Within a network, the function of neuromodulatory communication is ultimately secondary to the standard, instantaneous transmission of data, and although one of the key characteristics of neuromodulation is its long time course of action relative to that of neurotransmission, modulatory effects are not typically permanent. Furthermore, it may be disadvantageous for neuromodulatory systems to operate concurrently with standard neurotransmission at all times. The attenuation of signal intensity over time, eventually leading to a neuron's return to its normal state, as well as the ability to adjust a network's level of neuromodulation or even switch it on and off should be incorporated into implementations of neuromodulatory systems.

The effects of a neuromodulatory signal could be limited simply by a number of cycles or time steps, which may be a set value or be determined by other factors such as

the type of signal, the type of effect, or the postsynaptic node. The value of a neuromodulatory signal or the magnitude of its effect could also decrease continuously over time. Some combination of limiting neuromodulatory signals continuously and limiting signals through fixed time steps could also occur, depending on the nature of the signal and its effect.

Presumably, any unnecessary neuromodulation within a network would be extracted via the training process or evolution. However, it may be beneficial for a network to possess the ability to regulate its modulatory system during its lifetime. Such regulation could occur, for example, by modulating neuronal receptors in systems that utilize receptors, modulating the time it takes for modulatory effects to build up, or modulating the duration or magnitude of modulatory effects. These effects could emerge from the same sources as standard neuromodulatory effects, since a separate "metamodulatory" system may add an unnecessary burden to the network. The regulation of neuromodulation allows for input experienced during a network's lifetime to alter the neuromodulatory process during the same lifetime, which may valuable to the network if sudden changes within its environment occurred over a short period of time.

5. Summary

Ultimately, the definition and significance of neuromodulation are the same in an artificial setting as they are in biological networks. Essentially, neuromodulatory systems alter neurons in such a way that their functions depend not solely on input but also on their current state, increasing the range of operations a network can perform.

The distinguishing properties of biological neuromodulation, the slow-acting, long-lasting effects and the wide range over which signals can have an influence, are likely to be important characteristics that should form the basis of artificial neuromodulatory systems. Furthermore, a neuromodulatory system should allow a network to swiftly modify its input-output mapping, and the network should in turn be able to modify its neuromodulatory components. Because the numerous varieties of existing artificial neural networks open up a considerable number of design possibilities in the implementation of neuromodulatory systems, the search for appropriate components and configurations should presumably be aided by evolutionary computation techniques.

References

Arnold SF (2011). Neuro-cognitive organization as a side-effect of the evolution of learning ability. In *IEEE Symposium on Artificial Life* (*ALIFE*).

Bäck T, Fogel DB, Michalewicz T, eds. (2000). *Evolutionary Computation 1: Basic algorithms and operators.* New York, NY: Taylor and Francis Group.

Benson JA, Levitan IB (1983). Serotonin increases an anomalously rectifying K^+ current in the *Aplysia* neuron R15. *Proc. Natl. Acad. Sci.* 80: 3522–5.

Bergfeldt N, Linåker F (2002). Self-organized modulation of a neural robot controller. In *Proceedings of the International Joint Conference on Neural Networks*, Piscataway, NJ: IEEE Press, p. 495-500.

Brezina V, Evans CG, Weiss KR (1994a). Activation of K current in the accessory radula closer muscle of *Aplysia californica* by neuromodulators that depress its contractions. *J. Neurosci*. 14: 4412-32.

Brezina V, Evans CG, Weiss KR (1994b). Enhancement of Ca current in the accessory radula closer muscle of *Aplysia californica* by neuromodulators that potentiate its contractions. *J. Neurosci.* 14: 4393-411.

Chen Y, Yu FH, Surmeier J, Scheuer T, Catterall WA (2006). Neuromodulation of Na^+ channel slow inactivation via cAMP-dependent protein kinase and protein kinase C. *Neuron* 49: 409-20.

Covarrubias M, Wei A, Salkoff L, Vyas TB (1994). Elimination of rapid potassium channel inactivation by phosphorylation of the inactivation gate. *Neuron* 13: 1403-12.

Cox BR, Krichmar JL (2009). Neuromodulation as a robot controller: a brain-inspired strategy for controlling autonomous robots. *IEEE Robotics and Automation Magazine* 16: 1115-29.

Cropper EC, Brezina V, Vilim FS, Harish O, Price DA, Rosen S, Kupfermann I, Weiss KR (1994). FRF peptides in the ARC neuromuscular system of *Aplysia*: purification and physiological actions. *J. Neurophysiol.* 72: 2181-95.

Cropper EC, Lloyd PE, Reed W, Tenenbaum R, Kupfermann I, Weiss KR (1987). Multiple neuropeptides in cholinergic motor neurons of *Aplysia*: evidence for modulation intrinsic to the motor circuit. *Proc. Natl. Acad. Sci*. 84: 3486-90.

Cropper EC, Miller MW, Tenenbaum R, Kolks MAG, Kupfermann I, Weiss KR (1988). Structure and action of buccalin: a modulatory neuropeptide localized to an identified small cardioactive peptide-containing cholinergic motor neuron of *Aplysia californica*. *Proc. Natl. Acad. Sci.* 85: 6177-81.

Dash PK, Hochner B, Kandel ER (1990). Injection of the cAMP-responsive element into the nucleus of *Aplysia* sensory neurons blocks long-term facilitation. *Nature* 345: 718-721.

Di Paolo EA (2000). Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In Meyer JA, Berthoz A, Floreano D, Roitblat HL, Wilson SW, eds. *From Animals to Animats 6: Proceedings of the VI International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, p. 440-9.

Dickinson PS, Mecsas C, Marder E (1990). Neuropeptide fusion of two motor-pattern generator circuits. *Nature* 344: 155-8.

Dolphin AC (1990). G protein modulation of calcium currents in neurons. *Annu. Rev. Physiol.* 52: 243-55.

Downing JEG, Role LW (1987). Activators of protein kinase C enhance acetylcholine receptor desensitization in sympathetic ganglion neurons. *Proc. Natl. Acad. Sci.* 84: 7739-43.

Dreyfus, G (2005). *Neural Networks, Methodology and Applications*. New York, NY: Springer-Verlag.

Drummond AH, Benson JA, Levitan IB (1980). Serotonin-induced hyperpolarization of an identified *Aplysia* neuron is mediated by cyclic AMP. *Proc. Natl. Acad. Sci.* 77: 5013-5017.

Dürr P, Mattiussi C, Soltoggio A, Floreano D (2008). Evolvability of neuromodulated learning for robots. In *Proceedings of the 2008 ECSIS Symposium on Learning and Adaptive Behavior in Robotic Systems*. Los Alamitos, CA: IEEE Computer Society, p. 41-6.

Eggenberger P, Ishiguro A, Tokura S, Kondo T, Uchikawa Y (1999). Toward seamless transfer from simulated to real worlds: a dynamically-rearranging neural network approach. In *LNCS* 1812:44-60, Berlin: Springer-Verlag.

England S, Bevan S, Docherty RJ (1996). PGE_2 modulates the tetrodotoxin-resistant sodium current in neonatal rat dorsal root ganglion neurones via the cyclic AMP-protein kinase A cascade. *J. Physiology* 495: 429-40.

Farries MA, Fairhall AL (2007). Reinforcement learning with modulated spike timing-dependent synaptic plasticity. *J Neurophysiol* 98: 3648-65.

Fellous JM (1999). The neuromodulatory basis of emotion. *The neuroscientist* 5: 283-94.

Fellous JM (2004). From human emotions to robot emotions. In *Architectures for Modeling Emotions: Cross-Disciplinary Foundations. Papers from the 2004 AAAI Spring Symposium,* AAAI Press, p. 37-47.

Fitzgerald EM, Okuse K, Wood JN, Dolphin AC, Moss SJ (1999). cAMP-dependent phosphorylation of the tetrodotoxin-resistant voltage-dependent sodium channel SNS. *J. Physiology* 516: 433-46.

Floreano D, Mondada F (1996). Evolution of plastic neurocontrollers for situated agents. In Maes P, Mataric M, Meyer JA, Pollack J, Wilson S, eds. *From Animals to Animats 4, Proceedings of the International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.

Florian RV (2007). Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Computation* 19: 1468-1502.

French RLB, Cañamero L (2005). Introducing neuromodulation to a Braitenberg vehicle. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, p. 4188-93.

Frost WN, Hoppe TA, Wang J, Tian LM (2001). Swim initiation neurons in *Tritonia diomedea*. *Am. Zool.* 41: 952-61.
Fujii A, Ishiguro A, Aoki T, Eggenberger P (2001). Evolving bipedal locomotion with a dynamically-rearranging neural network. In *Advances in Artificial Life, LNCS* 2159: 509-18, Berlin: Springer-Verlag.

Garthwaite J, Boulton CL (1995). Nitric oxide signaling in the central nervous system. *Annu. Rev. Physiol.* 57: 683-706.

Grosmaitre X, Marion-Poll F, Renou M (2001). Biogenic amines modulate olfactory receptor neurons firing activity in *Mamestra brassicae*. *Chem. Senses* 26: 653-61.

Gurney K (1997). *An Introduction to Neural Networks*. London: University College London Press.

Hammer M (1993). An identified neuron mediates the unconditioned stimulus in associated olfactory learning in honeybees. *Nature* 366: 59-63.

Hassoun MH (1995). *Fundamentals of Artificial Neural Networks*. Cambridge, MA: MIT Press.

Hinkle DJ, Macdonald RL (2003). β subunit phosphorylation selectively increases fast desensitization and prolongs deactivation of α1β1γ2L and α1β3γ2L $GABA_A$ receptor currents. *J. Neurosci.* 23: 11698-710.

Holland JH (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.

Hooper SL, Moulins M, Nonnotte L (1990). Sensory input induces long-lasting changes in the output of the lobster pyloric network. *J. Neurophysiol.* 64: 1555-73.

Huerta PT, Lisman JE (1993). Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. *Nature* 364: 723-5.

Husbands P (1998). Evolving robot behaviours with diffusing gas networks. In P. Husbands and J.-A. Meyer, eds. *Evolutionary Robotics: Proc. EvoRob'98*, LNCS 1468, Springer Verlag, p. 71-86.

Husbands P, Smith T, Jakobi N, O'Shea M (1998a). Better living through chemistry: evolving GasNets for robot control. *Connection Science* 10: 185-210.

Husbands P, Smith T, O'Shea M, Jakobi N, Anderson J, Phillipides A (1998b). Brains, Gases and Robots. In *Proceedings ICANN '98* Springer-Verlag, *LNCS*, p. 51-63.

Isaacson JS (1998). $GABA_B$ receptor-mediated modulation of presynaptic currents and excitatory transmission at a fast central synapse. *J. Neurophysiol*. 80: 1571-6.

Ishiguro A, Fujii A, Eggenberger P (2003). Neuromodulated control of bipedal locomotion using a polymorphic CPG circuit. *Adaptive Behavior* 11:7-17.

Ismailov II, Benos DJ (1995). Effects of phosphorylation on ion channel function. *Kidney International* 48: 1167-79.

Izhikevich EM (2000). Neural excitability, spiking and bursting. *Int. J. Bifurc. Chaos* 10: 1171-1266.

Izumi Y, Zorumski CF (1999). Norepinephrine promotes long-term potentiation in the adult rat hippocampus in vitro. *Synapse* 31: 196-202.

Jakobi N, Husbands P, Harvey I (1995). Noise and the reality gap: the use of simulation in evolutionary robotics. *LNCS* 929: 704-20.

Jonas EA, Kaczmarek LK (1999). The inside story: subcellular mechanisms of neuromodulation. In Katz PS, ed. *Beyond Neurotransmission: Neuromodulation and its importance for information processing*. New York: Oxford University Press.

Kandel ER, Schwartz JH (1982). Molecular biology of learning: modulation of transmitter release. *Science* 218: 433-443.

Katz PS (1999). What are we talking about? Modes of neuronal communication. In Katz PS, ed. *Beyond Neurotransmission: Neuromodulation and its importance for information processing*. New York: Oxford University Press.

Katz PS, Edwards DH(1999). Metamodulation: the control and modulation of neuromodulation. In Katz PS, ed. *Beyond Neurotransmission: Neuromodulation and its importance for information processing*. New York: Oxford University Press.

Katz PS, Getting PA, Frost WN (1994). Dynamic neuromodulation of synaptic strength intrinsic to a central pattern generator circuit. *Nature* 367: 729-31.

Kondo T, Ishiguro A, Tokura S, Uchikawa Y, Eggenberger P (1999). Realization of robust controllers in evolutionary robotics: a dynamically-rearranging neural network approach. In *Proceedings of the 1999 Congress of Evolutionary Computation* 1:366-73.

Krichmar JL (2012). A biologically inspired action selection algorithm based on principles of neuromodulation. To be presented at: International Joint Conference on Neural Networks, Brisbane, Australia.

Lan J, Skeberdis VA, Jover T, Grooms SY, Lin Y, Araneda RC, Zheng X, Bennett MVL, Zukin RS (2001). Protein kinase C modulates NMDA receptor trafficking and gating. *Nat. Neurosci.* 4: 382-90.

Levitan ES, Levitan IB (1988). Serotonin acting via cyclic AMP enhances both the hyperpolarizing and depolarizing phases of bursting pacemaker activity in the *Aplysia* neuron R15. *J. Neurosci.* 8: 1152-1161.

Levitan IB (1985). Phosphorylation of ion channels. *J. Membrane Biol.* 87: 177-90.

Levitan IB (1994). Modulation of ion channels by protein phosphorylation and dephosphorylation. *Annu. Rev. Physiol*. 56: 193-212.

Li S, Cullen WK, Anwyl R, Rowan MJ (2003). Dopamine-dependent facilitation of LTP induction in hippocampal CA1 by exposure to spatial novelty. *Nature* 6: 526-31.

Linn CE, Roelofs WL (1986). Modulatory effects of octopamine and serotonin on male sensitivity and periodicity of response to sex pheromone in the cabbage looper moth, *Trichoplusia ni. Arch. Insect Biochem. Physiol.* 3: 161-72.

Marder E, Bucher D (2007). Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annu. Rev. Physiol.* 69: 291-316.

Marder E, Calabrese RL (1996). Principles of rhythmic motor pattern generation. *Physiological Reviews.* 76: 687-717.

Margiotta JF, Berg DK, Dionne VE (1987). Cyclic AMP regulates the proportion of functional acetylcholine receptors on chicken ciliary ganglion neurons. *Proc. Natl. Acad. Sci.* 84: 8155-9.

Meng Y, Jin Y, Yin J, Conforth M (2010). Human activity detection using spiking neural networks regulated by a gene regulatory network. *IEEE International Joint Conference on Neural Networks* (*IJCNN 2010*).

McCormick DA (1992). Neurotransmitter actions in the thalamus and cerebral cortex and their role in neuromodulation of thalamocortical activity. *Prog. Neurobiol.* 39: 337-88.

McCormick DA, Bal T (1997). Sleep and arousal: thalamocortical mechanisms. *Annu. Rev. Neurosci.* 20: 185-215.

McHale G, Husbands P (2004a). GasNets and other evolvable neural networks applied to bipedal locomotion. In Schaal S *et al.*, eds. *From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behaviour* (*SAB'2004*), MIT Press, p. 163-172.

McHale G, Husbands P (2004b). Quadrupedal locomotion: GasNets, CTRNNs and hybrid CTRNN/PNNs compared. In Pollack J *et al.*, eds. *Proc. Alife IX*, MIT Press, p. 106-112.

Meyrand P, Simmers J, Moulins M (1991). Construction of a pattern-generating circuit with neurons of different networks. *Nature* 351: 60-3.

Miller RJ (1990). Receptor-mediated regulation of calcium channels and neurotransmitter release. *FASEB J.* 4: 3291-9.

Miller RJ (1998). Presynaptic receptors. *Annu. Rev. Pharmacol. Toxicol.* 38: 201-27.

Misonou H, Mohapatra DP, Park EW, Leung V, Zhen D, Misonou K, Anderrson AE, Trimmer JS (2004). Regulation of ion channel localization and phosphorylation by neuronal activity. *Nat. Neurosci.* 7: 711-8.

Nolfi S, Floreano D (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.

Numann R, Catterall WA, Scheuer T (1991). Functional modulation of brain sodium channels by protein kinase C phosphorylation. *Science* 254: 115-8.

Otmakhova NA, Lisman JE (1996). D1/D5 dopamine receptor activation increases the magnitude of early long-term potentiation at CA1 hippocampal synapses. *J. Neurosci.* 16: 7478-86.

Otsu K, Ishiguro A, Fujii A, Aoki T (2001). Evolving an adaptive controller for a quadruped-robot with dynamically-rearranging neural networks. *Intelligent Robots and Systems, Proceedings*. *IEEE/RSJ International Conference on Intelligent Robots and Systems* 4: 2036-2044.

Park K, Mohapatra DP, Misonou H, Trimmer JS (2006). Graded regulation of the Kv2.1 potassium channel by variable phosphorylation. *Science* 313: 976-9.

Parussel K, Cañamero L (2007). Biasing neural networks towards exploration or exploitation using neuromodulation. In *Proceedings of the 17th International Conference on Artificial Neural Networks, Part II, ICANN* 2007, *4669 of Lecture Notes in Computer Science*. Berlin: Springer-Verlag, p. 889-98.

Parussel K, Smith LS (2005). Cost minimisation and reward maximization. A neuromodulating minimal disturbance system using anti-hebbian spike timing-dependent plasticity. In *Proceedings of the Symposium on Agents that Want and Life: Motivational and Emotional roots of Cognition and Action at the AISB-05 conference*, p.98-100.

Philippides A, Husbands P (2005). Flexible couplings: diffusing neuromodulators and adaptive robotics. *Artif. Life* 11: 139-160.

Pinsker H, Kupfermann I, Castellucci V, Kandel E (1970). Habituation and dishabituation of the gill-withdrawal reflex in *Aplysia. Science* 167: 1740-2.

Pinsker HM, Hening WA, Carew TJ, Kandel ER (1973). Long-term sensitization of a defensive withdrawal reflex in *Aplysia*. *Science* 182: 1039-42.

Popescu IR, Frost WN (2002). Highly dissimilar behaviors mediated by a multifunctional network in the marine mollusk *Tritonia diomedea*. *J. Neurosci*. 22:1985-93.

Pophof B (2000). Octopamine modulates the sensitivity of silkmoth pheromone receptor neurons. *J. Comp. Physiol. A* 186: 307-13.

Potjans W, Morrison A, Diesmann M (2009). A spiking neural network model of an actor-critic learning agent. *Neural Computation* 21: 301-39.

Purves D, Augustine GJ, Fitzpatrick D, Hall WC, LaMantia A, McNamara JO, White LE, eds. (2008). *Neuroscience.* (4th ed.). Sunderland, Massachusetts: Sinauer Associates Inc.

Qian J, Saggau P (1997). Presynaptic inhibition of synaptic transmission in the rat hippocampus by activation of muscarinic receptors: involvement of presynaptic calcium influx. *Br. J. Pharmacol.* 122: 511-9.

Rabuñal JR, Dorado J (2006). *Artificial Neural Networks in Real-Life Applications*. Hershey, PA: Idea Group Publishing.

Ramirez J, Orchard I (1990). Octopaminergic modulation of the forewing stretch receptor in the locust *Locusta migratoria*. *J. Exp. Biol.* 149: 255-79.

Reinhart PH, Chung S, Martin BL, Brautigan DL, Levitan IB (1991). Modulation of calcium-activated potassium channels from rat brain by protein kinase A and phosphatase 2A. *J. Neurosci.* 11: 1627-35.

Risi S, Hughes CE, Stanley KO (2010). Evolving plastic neural networks with novelty search. *Adaptive Behavior* 18: 470-91.

Roeper J, Lorra C, Pongs O (1997). Frequency-dependent inactivation of mammalian A-type K^+ channel $K_v1.4$ regulated by Ca^{2+}/calmodulin-dependent protein kinase. *J. Neurosci.* 17: 3379-91.

Ross AF, Green WN, Hartman DS, Claudio T (1991). Efficiency of acetylcholine receptor subunit assembly and its regulation by cAMP. *J. Cell Biol.* 113: 623-36.

Schaible H, Ebersberger A, von Banchet GS (2002). Mechanisms of pain in arthritis. *Ann. N. Y. Acad. Sci.* 966: 343-54.

Schmidt K, Kruse M, Hatt H (1994). Dopamine alters glutamate receptor desensitization in retinal horizontal cells of the perch (*Perca fluviatilis*). *Proc. Natl. Acad. Sci.* 91: 8288-91.

Smith T, Husbands P, Philippides A, O'Shea M (2002). Neuronal plasticity and temporal adaptivity: GasNet robot control networks. *Adaptive Behavior* 10: 161-84.

Smith TMC, Husbands P, O'Shea M (2001). Not measuring evolvability: initial investigation of an evolutionary robotics search space. *Congress on Evolutionary Computation*, p. 9-16.

Smythies J (2005). *The Neuromodulators (The International Review of Neurobiology, Volume 64)*. San Diego, CA: Elsevier Academic Press.

Soltoggio A, Bullinaria JA, Mattiussi C, Dürr P, Floreano D (2008). Evolutionary advantages of neuromodulated plasticity in dynamic, reward-based scenarios. In *Proceedings of the 11th International Conference on Artificial Life*. Cambridge, MA: MIT Press, p. 569-76.

Soltoggio A, Dürr P, Mattiussi C, Floreano D (2007). Evolving neuromodulatory topologies for reinforcement learning-like problems. In *Proceedings of the 2007 IEEE Congress on Evolutionary Computation*. IEEE Press, p. 2471-8.

Swope SL, Moss SJ, Blackstone CD, Huganir RL (1992). Phosphorylation of ligand-gated ion channels: a possible mode of synaptic plasticity. *FASEB J.* 6: 2514-23.

Swope SL, Moss SJ, Raymond LA, Huganir RL (1999). Regulation of ligand-gated ion channels by protein phosphorylation. *Adv. Second Messenger and Phosphoprotein Res.* 33: 49-78.

Takahashi T, Forsythe ID, Tsujimoto T, Barnes-Davies M, Onodera K (1996). Presynaptic calcium current modulation by a metabotropic glutamate receptor. *Science* 274: 594-7.

Tonelli P, Mouret J (2011). Using a map-based encoding to evolve plastic neural networks. In *IEEE Workshop on Evolving and Adaptive Intelligent Systems* (*EAIS*).

Treede R, Meyer RA, Raja SN, Campbell JN (1992). Peripheral and central mechanisms of cutaneous hyperalgesia. *Prog. Neurobiol.* 38: 397-421.

Trimmer BA (1999). The messenger is not the message; or is it? In Katz PS, ed. *Beyond Neurotransmission: Neuromodulation and its importance for information processing*. New York: Oxford University Press.

Vargas PA, Di Paolo EA, Husbands P (2007). Preliminary investigations on the evolvability of a non-spatial GasNet model. In Almeida e Costa F, ed. *Proc. ECAL'07*. Springer , LNCS 4648, p. 966-75.

Vargas PA, Di Paolo EA, Husbands P (2008). A study of GasNet spatial embedding in a delayed-response task. In Bullock S *et al.*, eds. *Proc. Artificial Life XI*. MIT Press, 640-7.

Vijayaraghavan S, Schmid HA, Halvorsen SW, Berg D (1990). Cyclic AMP-dependent phosphorylation of a neuronal acetylcholine receptor α-type subunit. *J. Neurosci.* 10: 3255-62.

Wang Q, Liu L, Pei L, Ju W, Ahmadian G, Lu J, Wang Y, Liu F, Wang YT (2003). Control of synaptic strength, a novel function of Akt. *Neuron* 38: 915-28.

Watkins AJ, Goldstein DA, Lee LC, Pepino CJ, Tillett SL, Ross FE, Wilder EM, Zachary VA, Wright WG (2010). Lobster attack induces sensitization in the sea hare, *Aplysia californica*. *J. Neurosci.* 30: 11028-31.

Weimann JM, Meyrand P, Marder E (1991). Neurons that form multiple pattern generators: identification and multiple activity patterns of gastric/pyloric neurons in the crab stomatogastric system. *J. Neurophysiol.* 65: 111-22.

Ziemke T (1999). Remembering how to behave: recurrent neural networks for adaptive robot behavior. In Medsker LR, Jain LC, eds. *Recurrent Neural Networks: Design and Applications*. Boca Raton: CRC Press. p. 359-89.

Ziemke T, Thieme M (2002). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior* 10: 185-198.