such adaptation is based purely on learning, it could easily overfit, and catastrophic forgetting could result. However, if pattern-generator-based learning continues together with learning from the environment, it can serve a stabilizing effect. Adaptation to new inputs is combined with continual adaptation to the fundamental patterns in the domain. Such learning could occur e.g. during REM sleep. This mechanism could potentially explain why animals learn altered environments only partially, and why they spend much time on REM sleep when their neural structures are most plastic. Evolved pattern generators can thus provide a mechanism for continual genetic influences on behavior. It could similarly be instrumental in keeping artificial systems both adaptive and stable.

A further aspect of the synergy between evolution and learning is that evolution can discover the actual learning mechanisms. For instance in the task of discovering repeated patterns in an input sequence with a spiking neural network, evolution discovered plasticity rules that made the task possible in three different settings (Jordan et al. 2021): with reward feedback (reinforcement learning), error feedback (supervised learning), and without feedback (correlation-based unsupervised learning). With cartesian genetic programming as the evolution method (Julian F. Miller 2011), the system discovered symbolic expressions for such plasticity, making it possible to interpret the underlying physical factors, such as homeostasis in the well-known spike-timing-dependent plasticity mechanisms (STDP; Song, Miller, and Abbott 2000).

Many of the metalearning methods reviewed in Chapter 11 and others optimize different aspects of the learning mechanisms (Confavreux et al. 2020; Najarro and Risi 2020a; Tyulmankov, Yang, and Abbott 2022; Gonzalez and Miikkulainen 2021b; Bingham and Miikkulainen 2022; Elsken, Metzen, and Hutter 2019). While often the goal is to simply improve machine learning performance, such methods can also lead to insights into the learning algorithms themselves. For instance, in an experiment where agents needed to adapt to changing reward locations in a Minecraft navigation task, evolution discovered innate reward neurons that made the search for the reward effective even without an explicit reward signal (Ben-Iwhiwhu et al. 2020). Neuroevolution thus discovered structures that facilitated learning during the lifetime of the agent. Such synergies result in more powerful machine learning, but also help us formulate speific hypotheses about biological adaptation.

14.5 Constrained Evolution of Behavior

Section 7.1 illustrated an important principle in evolution of complex behavior: It does not exist in a vacuum, but is constrained and guided by interactions with the environment and with other agents. Simulations of cooperative evolution can thus help us understand the origins of biological behaviors as well. Section 7.1 already demonstrated several such opportunities, including how role-based cooperation may emerge, how adaptive teams can evolve, and how an evolutionary arms race may result in sophisticated herding and hunting behaviors.

This section further expands and generalizes that principle. The guidance may originate not only from complex interactions with the environment, but from general constraints on what the agent can do. For instance, a physical body imposes limits on what movements are possible. Sensory perception is limited, and processing power in decision-making is finite. If the goal is to build capable artificial agents, it makes sense to furnish them with as few such constraints as possible. Evolution can then be the most creative, and the agents most powerful in their task. However, if the goal is to create agents that are believable, for instance as simulated intelligent agents in a virtual environment, such constraints constitute an important guide: Evolution under constraints observed in nature leads the optimization process to discover behaviors that are natural, believable, and human-like. In other words, it explains the observed behaviors as optimal under the constraints seen in nature.

These effects can be observed most clearly in simulations of virtual creatures. Both the bodies and the brains of simulated physical creatures are evolved simultaneously in a simulated physical medium, such as a terrain or water (Sims 1991, 1994). With even a simple fitness reward such as getting close to a target, they develop both body structures and ways of moving their body that look remarkably animate.

Such target-following behaviors have been evolved in multiple experiments, with increasingly complex body structures and environments, and modes of locomotion such as running and swimming. However, evolving more complex behaviors has turned out significantly more challenging. For instance, it has been difficult to evolve creatures that would be able to employ different behaviors at different times, and make intelligent decisions between them.

One possible approach is to design a syllabus, i.e. a hierarchy of increasingly complex behaviors, and evolve them incrementally (Dan Lessin, Don Fussell, and Miikkulainen 2014; Daniel Lessin, Donald Fussell, and Miikkulainen 2013). The bodies in this experiment consisted of cylinders of different shapes, connected through muscles and attached through different kinds of joints, as well as sensors for threatening and attractive targets. The brains were neural networks containing some higher-level nodes such as those generating oscillation. At the lowest level, bodies and brains were evolved to move as fast as possible, to turn left and right, and to exert as strong a strike on the ground as possible. These behaviors were then encapsulated, i.e. the evolved neural network structures frozen and a trigger node added in order to activate and deactivate them. A second layer of behaviors was then evolved as neural networks that could activate the low-level behaviors as their output; they included moving or following a target, as well as running away from a target, both as as a combination of turning and locomotion. These behaviors were similarly encapsulated, and at the next level, combined with the strike behavior to establish an attack behavior. At the highest level, then, the attack and the running away were combined into “fight-or-flight”: if the object was sensed as threatening, run away—if it was sensed as attractive, attack.

The behavior that evolved was indeed highly believable, at least in a subjective sense. Several different kinds of bodies evolved at the lowest level, and behaviors were natural to them. For instance, some creatures had multiple legs and moved them rhythmically in order to advance. One agent consisted of simply two blocks, and was jumping forward one block by shaking the other block up and down. In order to create a strike, an agent with two side blocks acting as weights evolved to jump and land hard. Another one with a long arm evolved to hit the ground hard with it. In all these cases, the behaviors that evolved made sense in that particular body—it was also fascinating to see that there was no one solution, but many quite different solutions that were successful.

The behavior was also believable at the higher levels, including fight or flight. After watching the simulation for a while, it is easy to anthropomorphize the agent: It seems to

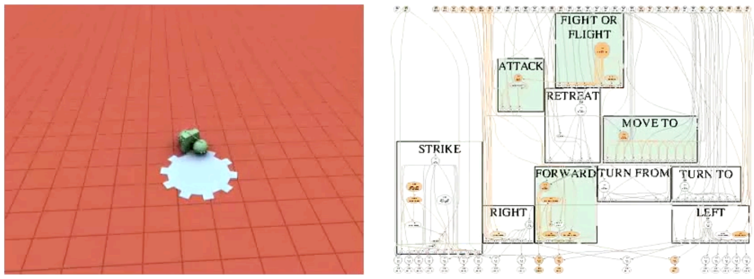

(*a*) Fight (i.e. attack) activated when sensing a good object

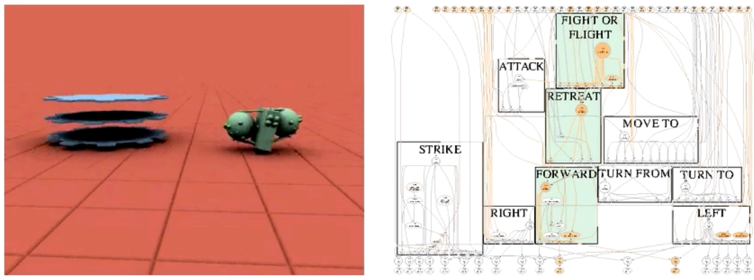

(*b*) Flight (i.e. retreat) activated when sensing a bad object

Figure 14.6: **Neuroevolution of complex behavior in evolved virtual creatures.** The bodies and the brains of simulated creatures were evolved together, thus providing constraints on what kind of movements were possible. As a result, they appear natural and therefore believable. The low-level behaviors such as locomotion, turning right and left, and strike were encapsulated and formed sub-behaviors to more complex behaviors turn-from, turn-to, retreat, and attack. At the highest level, the creature chooses between (*a*) fight and (*b*) flight depending on the object, as seen in this pair of figures. Such believability makes it natural to anthropomorphize the agents, which can be appealing in constructing virtual worlds. For animations of these behaviors, see https://neuroevolutionbook.com/neuroevolution-demos.

have a purpose when it chases a moving target, and when it changes to a threatening one, it seems scared reacting to the change and running away. And if the threatening object catches up with it and destroys it, you feel sorry for it. It is this kinds of agents that we can identify with and anthropomorphize that we would like to inhabit virtual worlds that are now being constructed. Constrained body-brain evolution may be a good way to get there. It is also a possible way to demonstrate why and how such a diversity of bodies and behaviors have evolved in nature—as different possible solutions to the same survival challenges.

Whether a behavior is believable or not is highly subjective and difficult to evaluate. In order to do that, several blind human judgments need to be collected under controlled conditions. It is of course possible to conduct such a study in the laboratory with human subjects. However, observing and interacting with virtual creatures is a lot of fun, and the evaluation can be as well. What if we turn the evaluation into a competition, and in addition to that, run it as an event at a conference where the audience consists of intelligent agent researchers and people interested in bringing AI into games?

This was indeed the goal of the Botprize competition, ran at the Computational Intelligence in Games conference in 2007-2012 (Hingston 2012). In essence, the competition was a Turing test for game bots: In the Unreal 2004 video game, there were both agents controlled by AI and agents controlled by human players. Some of the humans were playing the game as usual, trying to win. The AI agents were trying to play the same way as the humans did, and therefore be indistinguishable from human players. Some of the humans acted as judges, playing the game and interacting with the other players in order to decide whether those were controlled by humans or AI. They made the judgment about the other agents at the end of each game: The objective for the AI was to garner as many "human" judgments than "bot" judgments across several games with several different human players and judges.

Unreal is a representative of a multiplayer first-person-shooter game genre. Human players control their avatars who roam multiple levels in the game, gather possessions, and attack other players with different weapons. The game moves fast and requires quick control and decision making, however, it does not require linguistic communication. Therefore, to appear human, the AI-controlled bots would have to react, move, and make decisions similarly to the human players.

Indeed at the time it was not clear whether it was possible to capture such behavior. AI bots were routinely easy to identify in games in general: they behaved mechanically and repetitively, and the players often learned strategies that made it easy to defeat the AI bots. In many cases the game play consisted of figuring out the AI, and then moving on to other games. On the other hand, part of the reason for multiplayer games was to keep the game more interesting. It is always fun to beat your friends, but the friends also provided more interesting challenges. Therefore, being able to construct bots that behave indistinguishably from humans is not only important scientific question, but also has great value for game development in general.

It was also not clear what human-like behavior even was. In a human-subject study in the lab, Botprize games were captured on video, and the judges interviewed afterwards, trying to understand how they made their decisions, i.e. what constituted human-like behavior to them. Very little came out from that study. It turns out that humans are not very good at explaining what they do, and they may not even understand how they do it. More precisely, they are very good at constructing explanations when prompted to do so, but the explanations may have little to do with their actual process. In several occasions the judges gave fluent and logical explanations for why they judged the opponent as a bot, for example, because they moved in a certain way, or reacted in a certain way—not realizing that in the game, they actually judged this opponent as a human.

Yet the human judges were quite reliable in making those distinctions, at least in the beginning of the Botprize competition. Remarkably accurate, as a matter of fact. Sometimes the opponent jumped in front of them, interacted with them for a few seconds only, and ran away—and still the judges were able to make decisions well above chance. So there appears to be a quality in the behavior that humans have but bots at the time lacked. What is it?

In the first several years there was a significant and consistent gap between the humans and AI: While the human players were judged as human 60-70% of the time, the bots were mistook for humans only 20-30% of the time. Part of the problem turned out to be network latency—when the games were played over the internet, a time lag was introduced, and

Figure 14.7: **Neuroevolution of human-like behavior in the Botprize competition.** The competition is essentially a Turing test for game bots. The judge in this screenshot is player 443, and is interactive with another player, 932, in order to determine whether it is an AI-controlled bot or a human player. When neuroevolution was used to maximize the game score of the bot, the behavior was too systematic, repetitive, and effective to be human. Instead, when various constraints were imposed on accuracy, behavior selection, and multitasking, behavior become eventually indistinguishable from human behavior. Thus, the simulation demonstrated how even complex behavior can be seen as emerging from evolutionary optimization under environmental constraints. For animations of these behaviors, see https://neuroevolutionbook.com/neuroevolution-demos.

the humans dealt with that issue better than the bots. However, there were also significant differences in the behavior that gave the bots away. The bots were constructed to play well: for instance in evolution, the fitness early on was simply the final game score (Schrum, I. Karpov, and Miikkulainen 2012; Karpov, Schrum, and Miikkulainen 2012). Therefore, they evolved behaviors that were highly effective–but not necessarily human-like. For instance, they would run at full speed, and at the same time, shoot at maximum accuracy. If the judge did something unexpected, e.g. run straight into them, they would react immediately and perform the same behaviors as always when close to the opponent. Humans rarely do that. They get startled when something unexpected happens, and need to process it before they can react. Their performance varies, and becomes less accurate and effective under load. They do not perform multiple behaviors well at the same time. This was a fundamental difference between bots and humans.

However, when such performance constraints were imposed on the bots during evolution, their behavior changed significantly. They were no longer able to simply optimize the game score, but had to do it while limited in their accuracy, choice of actions, and ability to multitask (Figure 14.7; Schrum, I. V. Karpov, and Miikkulainen 2011). In essence, they got tired and distracted, and performed inconsistently. In other words, they become more human-like. In the last Botprize competition in 2012, they were indeed mistook for humans more than 50% of the time. Not only that, they were judged as humans more often than half of the human players!

Therefore, Botprize was a remarkable success in three ways: (1) it demonstrates how even complex behavior seen in nature can be seen as optimization under constraints; (2) It demonstrated how neuroevolution can be similarly constrained to discover more believable, more human-like behavior; and (3) it showed how a scientific evaluation can be turned into a fun and interesting event, i.e. a competition that promotes innovation and focus to the entire area of research.

This success by no means suggests that the work on evolving human-like behavior is now concluded. While it was successful at the low levels, there is an entire cognitive level that is not yet captured. For instance, human players lay traps such as running around the corner and waiting for the opponent there in order to ambush them. A human player may fall for that trap once or twice, but will learn very quickly to avoid it. In contrast, the bots will fall for it over and over again. In order to play like a human more comprehensive, the bots will need to learn and adapt. They need to adjust their play depending on the opponent. Moreover, there are challenges in playing with multiple other agents, especially in coordinating a team play. And of course, such coordination will ultimately require communication, which was not addressed in Botprize at all. Some of these issues will be addressed in the remaining two sections of this chapter.

14.6 Understanding Evolutionary Breakthroughs

As discussed above, neuroevolution experiments have demonstrated how competition, cooperation, environmental constraints, diversity, effective encodings, and many other ingredients can give rise to intelligent behavior. However, they are very general, and rarely address a specific research question in biology, i.e. how a particular behavior in a particular species may have evolved.

Such simulations are possible as well, especially in cooperation with evolutionary biologists. One promising opportunity is to understand evolutionary origins of the behaviors seen in hyenas, particularly the spotted hyena *crocuta crocuta*. A group of biologists led by Kay Holekamp has maintained a research station in Masai Mara since 1988, and have chronicled much of the hyena behaviors as well as their biology (Smith et al. 2017). These observations have been a motivation for several of the experiments already discussed, including those of role-based cooperation (Section 7.1.2 and the evolutionary arms race (Section 7.2.2), as well as others such as the tradeoffs between cooperative vs. individual hunting (Rajagopalan et al. 2011).

However, one of the behaviors of *crocuta crocuta* is particularly interesting: hyenas can team up to steal a kill from lions (Lehmann et al. 2016). Lions are much bigger and stronger predators and can easily kill hyenas. The Holekamp team has observed hundreds of interactions between them; usually hyenas stay out of their way, but there are many cases where they seem to employ a sophisticated cooperative strategy in order to drive the lions away from their kill. For example some 2-3 lions may have caught a zebra, and are feasting on it, when a few hyenas wonder by. The hyenas do not get close, but appear careful and even fearful, as they should be in the presence of such a predator threat. Instead, they start vocalizing loudly. Other hyenas within hearing distance are attracted to these vocalizations, and soon a large number of them, e.g. 20-30, starts to gather around the lions. Their behavior changes to that of strong interactions: their vocalizations change, they rub against each other, they make fast moves, and generally excite each other. As the excitement builds, they get less fearful, push each other closer to the hyenas, and make threatening gestures towards them, until (it seems) they cannot hold back their aggressive behavior any longer. In a dramatic, highly coordinated, and precisely timed move, they form a wall around the hyenas and attack them simultaneously. Typically they approach from three sides, leaving the lions a way out. If there are enough hyenas, typically four times the lions, and they are coordinated enough, the lions are overwhelmed and simply escape, leaving the kill to the hyenas.

How can such mobbing behavior have emerged in evolution? It is even more mysterious because hyenas, as effective as they are as hunters, are not that sophisticated in other ways. They live in clans and have a strict matriarchal hierarchy–perhaps because they have teeth and jaws that can crack bones, so that any disputes between them could be fatal. They do have territories, and vicious clan wars where they are sometimes disputed. They can hunt small pray individually, and team up to hunt larger prey, such as zebras. They also collaborate to take care of their young. But compared to other species that live in the same environment, such as baboons, these behaviors are less advanced. In particular, whereas baboons are good at learning new behaviors and coping with new situations, hyenas are not very flexible in their ways, and they do not learn as easily (Benson-Amram and Holekamp 2012). Stealing a kill from lions appears unusually sophisticated for them, and it is likely not a behavior they have learned—instead, it appears to be innate, i.e. an immediate product of evolution. Moreover, other hyena species that live nearby in Eastern Africa do not exhibit the mobbing behavior. Therefore, this behavior seems to be a breakthrough for the species—evolution of intelligence in action.

Computational simulations thus offer a potentially powerful way to gain insights into the mobbing behavior and its origins. Indeed, several such simulations have been built, focusing on game-theoretic as well as evolutionary computation aspects of it (Rajagopalan, Holekamp, and Miikkulainen 2019; Jahns and Hintze 2018). One such simulation suggested that a leading bold individual might evolve, making the cooperative behavior more likely to emerge (Solomon, Soule, and Heckendorn 2012; Fairey and Soule 2014). However, such individuals are not clearly identifiable in biology. The hyenas do indeed differ in how bold they are—some get closer sooner, and others hang back—but eventually they act primarily as a homogeneous team. Their behavior is associated with strong emotions, with fear competing with affiliation and aggression. While the behaviors themselves suggest emotions, it is also possible to measure them quantitatively, albeit coarsely, by analyzing the hormones in the stool samples they leave behind. The analysis indeed reveals elevated levels of the signature hormones for these emotions after such a lion encounter. The emotions may thus play a crucial role in allowing the team to form and to act cohesively.

Based on these observations, a neuroevolution simulation was set up to study how the mobbing behavior might emerge (Rajagopalan, Holekamp, and Miikkulainen 2020, Figure 14.8;). Ten hyenas and one lion were placed randomly in a 100×100 toroidal grid world. The hyenas could move at each timestep, and the lion was stationary (with a kill). If a hyena came within 20 steps of the lion, i.e. inside an “interaction circle”, it was likely to get killed, but if there were four or more hyenas within the interaction circly at any time, the lion got mobbed. The hyenas sensed the distance and direction to the lion, whether there were at least three other hyenas within the interaction circle, and whether the lion had

(a) A crucial moment in the interaction (b) Simulation setup

Figure 14.8: **Complex coordinated behavior of hyenas mobbing lions.**. In this behavior, hyenas form a mob that attacks a group of lions, gaining possession of their kill. (a) A screen capture of a video documenting a mobbing event. Lions are much stronger than hyenas, but if the hyenas are much more numerous and coordinate their attack well, they can drive the lions away from the kill. This behavior is more complex than others the hyenas exhibit, largely hereditary, and may represent an evolutionary breakthrough. (b) A simulation of mobbing. A lion and several hyenas are placed in a 100×100 grid world. If four or more hyenas enter the interaction circle simultaneously, they get a high reward; if fewer than four, they get killed. Neuroevolution simulations suggest that mobbing can arise from the simpler stepping stones of attacking, waiting at a distance, and waiting at the circle. These behaviors persist even in prolonged evolution, making the mobbing behaviors more robust. (Figure (b) from Rajagopalan, Holekamp, and Miikkulainen 2020) For videos and animations of these behaviors, see https://neuroevolutionbook.com/neuroevolution-demos.

already been mobbed. The hyenas that participated in the mobbing event receive a full fitness, those that stepped into the circle after mobbing had already happend receive an 80% fitness, and others received no fitness at all. Thus, the ideal hyena would approach the lion until it was just outside the interaction circle, wait there until at least three other hyenas made it there as well, and then step inside the circle at the same time as those other hyenas. However, for this behavior to be successful, at least three other hyenas needed to be able to perform it as well, and also time it just right. This required cooperation and timing makes mobbing very difficult to evolve.

Neuroevolution was based on NEAT, and as usual, started with random small networks. Over 1000 generations four mian behaviors were observed, differing based on how bold they were: (1) risk-takers ran straight to the lion regardless of other hyenas, and were usually killed quickly; however, they were sometimes successful if other hyenas joined them at the right time. (2) Risk-evaders-outside-circle hanged back and only approached the lion after it had been killed, receiving lower rewards with little risk, but also sometimes running out of time and not receiving any rewards. (3) Risk-evaders-at-circle approached the lion but stopped at the circle, and only stepped in after the lion had been killed, receiving low rewards reliably; and (4) mobbers behaved successfully as described above.

At the start of the simulation the networks were random and their actions were random as well, which amounted to imperfect and inconsistent risk-taking and risk-evasion. Both of these behaviors quickly became more consistent. The number of risk-takers increased quickly because such a rushing behavior is easy to construct. On the other hand, risk-evading hyenas are more likely to survive, and they thus persisted in the population as well, establishing the opposite behavior, i.e. waiting. These two behaviors constituted the first two stepping stones.

Over a few generations, mobbing events started to happen by accident, and such events increased gradually with an increasing number of risk-takers. Risk-takers were occasionally recombined with risk-evaders, bringing them closer to the circle without crossing it. This progress led to the discovery of the circle, and thus the third stepping stone of risk-evaders-at-circle. Mobbing was happening still largely by accident, but frequently enough so that eventually it was possible for evolution to discover precise timing for it. As a result, in approximately 10 generations, 90% of the hyenas were mobbers, and succesful 90% of the time.

Thus, each of the stepping stones played a role in discovering mobbing behavior. Because of them, it was possible to overcome the deceptive fitness landscape and develop the precise coordination required. Interestingly, even in prolonged evolution over 1000 generations, these stepping stones still existed in the population in low numbers. Evolution reached a dynamic equilibrium where some of the mobbers had risk-taker or risk-evader offspring, who again may have mobber offspring. The teams were robust enough to tolerate such diversity: as long as at least six of the 10 hyenas were mobbers, they successfully mobbed most of the time. However, the teams were even more successful with more mobbers, so why did such diversity persist?

As has been observed in prolonged evolution experiments in general, if evolution is continued after solutions have been discovered, the solutions often become more robustly encoded, and less likely to break in crossover and mutation Watson et al. 2011; Rajagopalan, Holekamp, and Miikkulainen 2014. However, the behavior itself may become more robust as well: In this case, the mobbers can be successful with more challenging initial states, and able to work with teammates with more varied behavior. Thus, diversity is important not only in discovering novel solutions, but also in refining the solutions so that they are more effective in complex, uncertain environments, i.e. in the real world. It is interesting that in such environments, evolutionary pressures exist that promote diversity automatically.

Thus, the simulation demonstrated how the mobbing behavior could have emerged, and in particular the stepping stones required. A most interesting observation is that it does require individuals that are extremely bold, even to their own detriment. If some of them survive and reproduce, the offspring may discover a moderation that it successful in a surprising way. There has, of course, been a long debate on the role of such behaviors in evolutionary biology, and many efforts to explain e.g. altruism (where individuals sacrifice themselves for the common good) have been developed (Kay, Keller, and Lehmann 2020). The simulation suggests that altruism may not be necessary, but instead simply variation in how bold the individuals are in trying to achieve their goals. Such variation may be implemented through different emotional balance, e.g. less fear and more affiliation and aggression.

In a broader sense, such variation in boldness may be crucial for innovation more generally. Even in humans there are always individuals who are willing to take more risks, and it is often those individuals that drive innovation. Indeed, individuals may simply wonder what's on the other side of those mountains, what's on the other side of the ocean, and such somewhat irrational wonderlust may have allowed humans to spread over the entire globe. Even today, thousands of people have already signed up for chance to get a one-way ticket to Mars, even though colonies or even the technology to get there does not exist. Such individuals are fascinated by the novelty and the unknown. Being the first there is a reward of itself. We still share a lot of the boldness of the first hyenas who wondered "what happens if I just ignore the lions and run straight towards the kill?"

Further, such simulations may be a way to look into the future as well, i.e. to predict how the hyenas are likely to evolve from their current state. Could this synchronized cooperative behavior serve as a foundation for developing more sophisticated communication? Or perhaps higher functions that could be useful in it as well, such as learning and memory? Other simulations suggests that discovering such functions require overcoming deceptive fitness (Lehman and Miikkulainen 2014)—very much like the immediate disadvantage of being too bold in the kill capture. Eventually, it may be possible to simulate major transitions as well, as discussed in Section 9.1.5. One of them is the evolution of language, which may already be within reach of neuroevolution simulations, as will be discussed next.

14.7 Evolution of Language

The last major transition in biology is the evolution of language (Maynard Smith and Szathmary 1997; Szathmáry 2015). It made cooperation possible more broadly and at a more sophisticated level: It allowed individuals to define roles and make them flexible, reason with hypotheticals and counterfactuals, and ultimately record knowledge and build on prior knowledge. Language is the ingredient that made it possible to construct complex societies. After a brief review of biological theory of language, neuroevolution approaches to evolving communication and structured language are reviewed in this section.

14.7.1 Biology of language

Language can be defined as the ability to generate an unlimited number of meanings from a finite set of symbols using grammatical rules. Although many animal species communicate using signals (essentially single words), language is unique to humans; therefore, some crucial aspects of the language ability must be genetically encoded. However, every human still needs to learn the specifics of their language through interaction with the environment. Such interactions also need to take place at a precise time during development (Friedmann and Rusou 2015). If a child does not get proper linguistic input when they are one-to-five years old, they do not develop full language abilities. The urge to develop language at that age is so great that groups of children in a linguistically poor environment may develop their own language systems, or enhance the existing ones. For instance, pidgin languages, or incomplete communication systems between adults who do not share a common language, become creole languages, i.e. fully formed languages of the next generation. It is also not tied to the verbal modality: deaf children of hearing parents can develop a fully formed sign-language system (Singleton and Newport 2004). Language learning is thus biologically programmed into humans. It can be seen as an example of both an expressive encoding and of synergetic development (Sections 9.1.4 and 14.4): Evolution specifies a learning mechanism that constructs the final complex system.

The degree of genetic determination has been up to a debate for decades. Chomsky and others have argued that the entire structure of language, a universal grammar, is genetically coded, and language learning consists of simply observing and setting the parameters of the grammar to obtain any specific language (Chomsky 1986). On the other hand, there are now large language models that learn perfectly good language simply by observing large amounts of text (Ouyang et al. 2022). If the model is large enough, and there's enough data to train it, the simple task of predicting the next word results in a model that can generate grammatical and even meaningful text.

Large language models still need to see much more language examples than humans do during development. It is thus likely that genetic influences play a larger role in biasing the learning system towards the right kind of structures. What exactly these constraints are and how evolution discovered them is a fascinating question. Given the progress in evolution of cooperation and intelligent behavior described above, it may be a question that we may be able to answer soon with neuroevolution simulations.

There are also clues from biology beyond just observations of current human language abilities. Earlier hominid species such as *homo erectus* are thought to have developed protolanguage abilities. They were able to cooperate more generally e.g. in scavenging that required competing with other species, and such cooperation may have required rudimentary language (Bickerton and Szathmáry 2011). Several current higher species, such as dolphins and apes, communicate regularly through vocalizations and gestures. Moreover, it is possible to train them to extend these abilities to structures similar to human language, even when they do not spontaneously utilize them in the wild (Herzing and Johnson 2015; Bindra et al. 1981). It is therefore possible to see these species as intermediate stages on the evolution of language, potentially constraining simulations.

The next two subsections review work done so far in this area, from early emergence of a communication code to multitask codes and cultural transmission. They also outline possible avenues for evolving language and uncovering the ingredients that make it possible.

14.7.2 Evolving Communication

Communication in artificial agents has been an active area of research for a long time (Wagner et al. 2003). Several experiments, many of them using neuroevolution, demonstrate emergence of communication codes for fundamental tasks such as mating, hunting, herding, and fighting. They are usually composed of symbols with simple meaning, although sometimes contextualized, rather than full language systems with grammatical structure. Nevertheless, they help us understand some of the conditions for communication and language to emerge.

One challenge is that it is difficult for the population in evolutionary simulations to converge on a common code. It is more likely to emerge within genetically related groups where selection operates at the group level (Floreano et al. 2007). It may also emerge more readily when the population is asymmetric, with clearly delineated roles. For instance, an influential early experiment focused on the simple but compelling problem of evolving a code for a cooperative task (Werner and Dyer 1992). In a simulated grid world there were males and females, both controlled through neural networks. The females were stationary but could sense the males location and emit 3-bit signals to them; the males could move and could perceive the signals but could not see the females. If a male entered the same location as a female, they would create offspring through genetic algorithms. Thus, in order to mate, the females needed to send instructions to the males, guiding them step by step to find the females. Initially the males would wonder around randomly; however, guidance on their last step would soon emerge, and gradually the symbols and their interpretation from further away. Eventually, a common code evolved that was effective and reliable in most situations. The simulation thus demonstrated that an effective communication code emerges when it enables effective evolution, and that asymmetric roles can make it easier to discover.

Since mating is a fundamental constituent in evolution, an interesting question is whether it is a indeed a possible origin for communication. In particular, proper mate selection may guide evolution towards more effective mating and higher quality offspring. In the simplest case, mate selection may be based on direct visible features and displays such as size, color, or strength. In higher animals, it is often based on communication, i.e. vocalizations or ritualized movements and gestures. Such signals can be interpreted as indicators of traits, making it possible to decide whether the potential mate is compatible. Once communication evolved to serve mate selection, it may have been exapted, or reused and adapted, for other tasks, eventually forming a basis for protolanguage (Bickerton 1990).

Such a possibility can be investigated in neuroevolution simulations (Rawal, Boughman, and Miikkulainen 2014). In a simulated world, individuals were controlled by neural networks, and they each had a 2-bit trait encoding that determined their compatibility with other individuals (Figure 14.9). The network output a 2-bit message, as well as a control signal on whether to mate or not, and whether to move or not. As their input, they received a 2-bit message, the distance to a prey, and a bit indicating whether they were in a mate or hunt situation. They were then paired up in both of these tasks. In mating, they communicated their trait to their partner and upon receiving the trait message from their partner, decided whether to mate; if they mated when the traits were compatible, they received a high fitness. In hunting, they had to move closer to the prey at each step, and also communicate to their partner whether they were one step away from the prey; if they entered the prey location at the same time, they received a high fitness.

In a series of experiments, it turned out that if mate selection was evolved first, and hunting was then added as a second task, the agents evolved successful behavior in both tasks much faster than when the tasks were introduced in the opposite order, or both at once. In other words, the code evolved for mate selection served as a better foundation for a code needed for hunting, than the other way around. The mate-selection code was simpler, and it was possible to complexify it to add hunting. Such incremental evolution was also more efficient than trying to evolve both behaviors at once. The final code used fewer symbols, and for instance the message to indicate readiness to mate was often reused to indicate readiness for prey capture. It thus served as an effective stepping stone for evolving complex behavior. The simulations thus suggest that communication may have evolved incrementally through stepping stones, and mate selection is a plausible origin for that process.

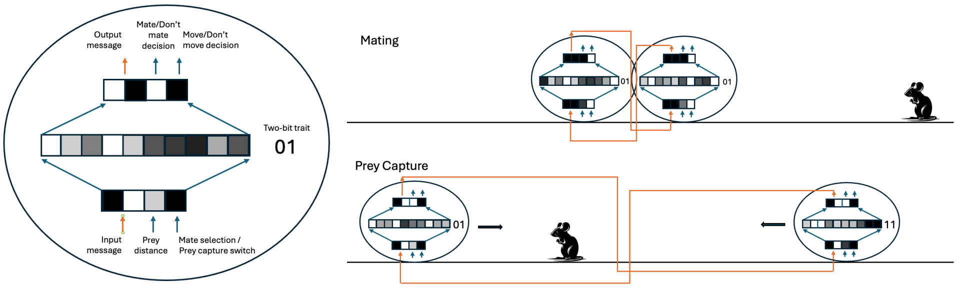

Figure 14.9: **Evolution of communication code for mate selection and hunting.** The agents were able to move in a simulated 1-D world where their fitness depended on successful mating and hunting. (*a*) Each agent in the population is controlled by an evolved neural network that receives the current task (either mate selection or hunting), the distance to the prey, and the message from the other agent as its input. At its output it decides to mate or move and generates a message that the other agents can use to decide whether to mate or whether to coordinate prey capture. For mating to be successful, the agents need to be compatible; compatibility is determined by an inherited 2-bit trait. For prey capture to be successful, they need to step on it at the same time. (*b*) Over evolution, the agents discover a messaging code that allows them to communicate their trait and their current distance to the prey effectively to other agents. It turns out that if mate selection is evolved first, instead of evolving prey capture first or at the same time, the agents develop a more effective and parsimonious code for both task. This result suggest that communication may have originally evolved for mate selection, and later adapted to other uses.

One fundamental aspect that is missing from such simulations is that the communication codes in nature are usually not innate, but are learned during the early life of the individual. That is, it is the ability for learning the code that is evolved. It is possible to extend language evolution simulations to such a setting as well (Li and Miikkulainen 2016). As in prior simulations, the agents were paired up in trials, and had to cooperate in order to hunt or mate successfully. Each generation began with a parenting phase: The newly generated offspring were paired up with their parents, and learned to be successful in the necessary communication through reinforcement learning. Next, all agents were paired up randomly in a socializing phase, and their overall fitness measured. Finally, the most successful agents became parents for the next generation. In this manner, it was possible to evolve successful behavior for both tasks, through a communication code that was evolved over multiple generations, and learned by each individual in each generation.

The simulation could then be used to further understand the pressures that cause communication to evolve. For the hunting and mating to be successful both partners had to be ready for it. The agents could either sense that readiness directly, or communicate it. By enabling and disabling such sensing and communication channels was is possible to make communication necessary or optional.

It turned out that if the agents could sense readiness directly, communication did not evolve, even when communication channels were available. Evolution thus discovered the simplest and most reliable way to be successful. However, if one or both readiness senses were disabled, communication did evolve. This result makes sense: without communication they would be successful only randomly, and there was thus a strong pressure to take advantage of communication-based coordination. Most interestingly, if communication was evolved for one of the tasks, it was also utilized in the other, even if it was not necessary in it. That is, if a communication ability is available, evolution will utilize it.

Evolution of communication and language may thus follow a similar process as many other innovations: evolution is a tinkerer, and will adapt whatever abilities there exist to other uses. Communication may be one such general ability that originated from a fundamental need e.g. for mate selection, and was then exapted to others. Would it be possible to make the transition for signaling with single symbols to communication with linguistic structures in this way? Possibilities are discussed in the next section.

14.7.3 Evolution of Structured Language

Evolution of language is difficult to study in biology because there is no fossil record and little other clues on how human ancestors communicated. Consequently, there are many theories about it, and they tend to be philosophical in nature. However, one significant tool we have at our disposal is computational modeling. It may be possible to gain insight into the conditions under which language evolves by building simulations.

Many computational approaches have indeed been developed using different techniques (Wagner et al. 2003). Rather than evolution, many of them focus on emergence of language. That is, they do not aim to model multiple generations of agents, but rather how communication can emerge in small groups of agents—sometimes even just two. They do, however, demonstrate discovery of some linguistic structure, not simply signaling between agents.

One approach is agent-based modeling, which may even involve physical robots (Steels 2016; Kirby, Griffiths, and Smith 2014). They take on roles of a teacher and learner and language emerges in order to perform a joint task. The signals not only combine into larger structures, but they also have a grounding, i.e. a semantic system emerges. In a larger group, iterated learning may be established, where the language is taught by individuals that learned it themselves earlier.

Mathematical modeling based on game theory has also provided interesting insights (Nowak and Krakauer 1999). When the game focuses on establishing reliable computation, it turns out words emerge from signaling, and grammar emerges from words, as a way to compensate for errors that are likely to arise in the communication medium.

Neural networks have also been used as an implementation for language agents in many studies (Batali 1998; Galke, Ram, and Raviv 2022). Most often they use recurrency or LSTM to input and output language, and a reinforcement learning mechanism such as REINFORCE to adapt. While compositional structures do emerge, they still do not match human languages well. It is possible that further cognitive constraints such as memory and alternation of speaker and listener roles are needed.

Evolutionary computing models constitute a fourth category of approaches. For instance, grammars can be evolved directly and compositionality discovered in service of a task (Zuidema and Hogeweg 2000). It is also possible to apply evolution to neural networks that generate the language. This kind of approach fits the problem most naturally: The ability for language is evolved over generations of a large number of individuals, and each individual learns the particular language during their lifetime.

While it is easy to discover communication through signaling in this manner (as was reviewed above), it is much harder to discover compositionality, i.e. linguistic structure. However, there has been some progress even early on. For instance, in an artificial environment with poisonous and edible mushrooms, neuroevolution discovered a signaling system that allowed the individuals to guide others to edible ones while avoiding poisonous ones (Cangelosi and Parisi 1998; Cangelosi 1999). Significantly, the system consisted of pairs of symbols signifying action and object. The offspring then learned the particular symbols through backpropagation. In this manner, a rudimentary grammatical structure evolved, and it is strikingly similar to the structures that can be taught to e.g. chimpanzees. Perhaps such a capability is the first step towards evolution of human language?

From such a starting point, why did language evolve only in humans? It is possible that the origin of language is not in communication, but in cognition. That is, while it is possible to build such a simple action-object protolanguage by complexifying signaling, perhaps true linguistic structure was discover as an exaption of other cognitive functions?

One theory is that language emerged as a useful tool in society, making it possible to coordinate actions such as group hunting, and group caring for the young when mothers were needed in foraging and other activities. As these activities became more sophisticated, it was necessary to understand that different individuals could take on different roles at different times, and how these roles might relate—in other words, flexible relational structures similar to grammatical structures. Once this structure was in place in the brain, it was exapted to enhance communication, and eventually, structured language emerged.

However, many other animals live in societies as well, and hunt in groups, and care for the young together (for instance, the hyenas discussed above). There was something different about human societies that serve as a stepping stone—and, again due to lack of any kind of direct evidence, there are many theories about what that might be (Bickerton 2007; Knight and Power 2012; Corballis 2011). One theory is that as humans became the apex scavenger, they needed to communicate the type and location of the kill. First this would be done iconically, but gradually with a displacement in time and space, which may have led to the abstraction needed for language. Another is that alliances and cliques formed in the societies when members wanted to dominate other members, and their maintenance required language. Gossip has also been indicated as a potential source, replacing or adding to physical grooming. A plausible explanation is that language emerged as a result (or together with) symbolic culture, for which there is some evidence in early objects and paintings (Figure 14.10). As societies grew more complex, rules were established for them to function better; symbolic representations and displacement made them possible, forming an impetus for language.

The time may now be right to start evaluating these hypotheses in computational neuroevolution simulations. There is enough computing power and sophistication to create virtual worlds where many of these conditions and constraints can be simulated. The neural networks would have to be much more complex, and able to perform many different tasks, but it is also an ability that is now emerging, as reviewed in this book. It is also possible to build up the simulations and hypotheses gradually from simple to more complex ones, and gain insight along the way. Neuroevolution is uniquely well suited to meeting these challenges, and may form a crucial ingredient in developing a theory of how language evolved, which is one of the most fascinating and perplexing questions in science.

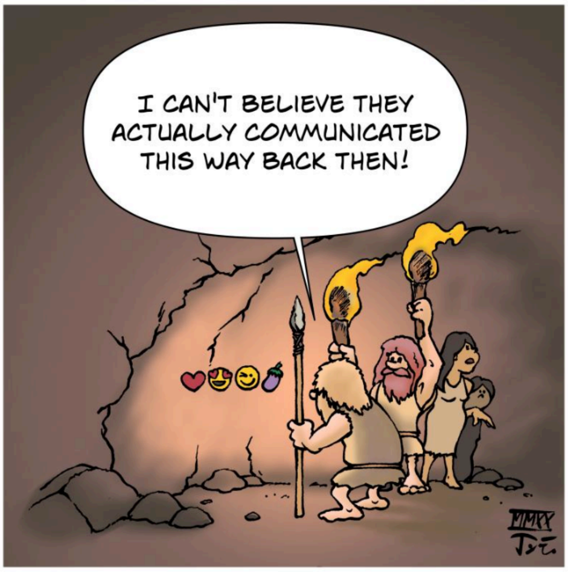

Figure 14.10: **A primary hypothesis is that language emerged at the same time as symbolic culture.** Then again, we don't really know. (Cartoon by Timo Essner, cartoonmovement.com/cartoon/emojis-0)

14.8 Chapter Review Questions

1. **Neural Structure and Evolutionary Origins:** How can neuroevolution simulations help us understand the evolutionary origins of specific neural structures, such as command neurons, and their role in behaviors like navigation and foraging?
2. **Central Pattern Generators (CPGs):** What are central pattern generators (CPGs), and how have neuroevolution experiments been used to model their role in controlling locomotion in animals, such as lampreys and salamanders?
3. **Modularity and Wiring Length:** How does the principle of minimizing wiring length contribute to the evolution of modular neural networks? Why does modularity lead to better performance and adaptability in evolving neural systems?
4. **Neuromodulation:** What role does neuromodulation play in adapting neural behavior? How does neuroevolution demonstrate its utility in tasks like the T-maze navigation?
5. **Synergetic Development:** How does the concept of synergetic development explain the interplay between genetic biases and lifetime learning? How have neuroevolution experiments demonstrated this principle in tasks such as foraging or pattern recognition?
6. **Constrained Evolution of Behavior:** How do body and environmental constraints influence the evolution of believable and natural behaviors in simulated agents, as demonstrated in fight-or-flight behavior evolution?
7. **Human-like Behavior in AI:** What role did performance constraints (e.g., limited accuracy, multitasking, and behavioral variability) play in evolving AI bots that were indistinguishable from human players in the Botprize competition?

8. **Evolutionary Breakthroughs in Social Behavior:** How did neuroevolution simulations model the emergence of mobbing behavior in hyenas, and what stepping stones contributed to the evolution of this complex coordinated strategy?
9. **Origins of Communication:** In simulations of mate selection and hunting, how did evolving communication for one task (e.g., mating) serve as a foundation for communication in another task (e.g., hunting)?
10. **Evolution of Language:** What theories exist about the origins of language, and how might neuroevolution simulations contribute to understanding the conditions and stepping stones that enabled its emergence?