Learning and Selection Processes

Marc ARTIGA

ABSTRACT: In this paper I defend a teleological explanation of normativity, i. e., I argue that what an organism (or device) is supposed to do is determined by its etiological function. In particular, I present a teleological account of the normativity that arises in learning processes, and I defend it from some objections.
Keywords: normativity, function, selection, learning, operant conditioning.

1. Introduction

Learning is the process by means of which systems formed by the action of natural selection can acquire new information and skills. Not every living being is able to adopt new knowledge or behaviors; evolution has gifted only some organisms with the capacity to learn from others and from their environment, in order for them to be more flexible and adapt their behavior to their actual circumstances (Godfrey-Smith, 1996). Learning is an adaptation. If this is true, theories seeking to account for the complexity of current living beings should pay special attention to it.

There are many complex aspects of learning mechanisms that deserve careful examination. Here, I am going to focus on a special difficulty that puzzles philosophers: normativity. Consider a common case; every time John the Hare goes to area A, he bumps into Mike the Wolf. Fortunately for him, John has managed to escape every time and after some episodes he learns to avoid going to A. Thereafter, he *is not supposed* to go to A. If John goes there, he would be making a *mistake*. How can a naturalist theory explain the appearance of these normative notions?

Of course, this problem not only affects learning processes but also inborn devices. For many years, philosophers have been struggling with questions such as why hearts are supposed to pump blood or livers are supposed to filter wastes from blood. However, I think that this difficulty has been overcome, at least in innate mechanisms. Teleonomic theories have explained how the process of natural selection determines what a device is supposed to do, relying on the concept of 'function' (Wright, 1973; Millikan, 1984). Hearts are supposed to pump blood (roughly) because the fact that past hearts pumped blood explains why hearts were favored by natural selection; pumping blood is what explains why hearts exists at all. But how does learning fit in this schema? To what extent does natural selection determine what an organism is able to learn? Can learning be considered a process analogous to natural selection? These are the questions I am going to tackle in this short essay.

Before starting, it is worth pointing out that, as many other common words, 'learning' is used in many contexts and in a wealth of different senses. Here I am going to

focus exclusively on the learning of new skills —in particular, I am going to leave aside the acquisition of novel knowledge. There are two reasons for that. First, I think the latter fits more easily in the teleonomic (and teleosemantic) theory, while the former may raise some difficulties. Furthermore, most naturalist theories have been attracted by the question of the acquisition of new forms of behavior. Both reasons suggest a preliminary discussion of the origin of novel skills in front of new knowledge.

2. Function and Normativity

For a good understanding of the problems that learning generate in teleonomic theories we need to set up some basic concepts. Undoubtedly, 'function' is the most central one.

2.1 Function

Function is the pivotal notion that enabled philosophers to explain how normativity can arise from biological processes, and a great deal of the debate we are going to take part on is going to revolve around this concept. So we will try first to explain this entangled and widely discussed notion.

Philosophers working on this tradition usually rely on an etiological definition of function (EF). Even if the details vary among different philosophers, let me try to put forward a definition that I think many of them would accept:

> **(EF)** A member d of the kind[1] D has an etiological function F[2] iff:
>
> (1) Past members of D performed F
>
> (2) Ds that performed F were selected over some competitors (say, Ds that performed G). The selection process had the following features (Hull et al., 2001):
>
> > (a) *Replication:* Members of D are to a large extent copies of each other.
> > (b) *Variation:* The replication included some changes that gave rise to F and F's competitors.
> > (c) *Environmental interaction*: D's interaction with external circumstances caused differential replication.
>
> (3) 1 and 2 (partially) explain why d exists. (Ayala, 1970; Millikan, 1984, p. 28)

(1) claims that it is not the first time that a member of D performs F (what implies that some time is required). (2) states that some kind of selection process has favored

---

[1] Strictly speaking, D is a 'reproductive established family', i. e. a group of individuals that resemble each other in important ways because they are the result of some process of copy.

[2] In order to avoid unnecessary terminological complexities, in the rest of the paper I will usually dispense with the adjective 'etiological' and just talk of 'function'. That will simplify considerably the terminology of the exposition, especially later on when I introduce direct and relational functions. Because our discussion focuses on evolutionary biology and animal learning, I do not think this simplification will cause any confusion. Nevertheless, it is important to keep in mind that when I make claims about functions, I am just referring to etiological functions. In particular, they are not claims about functions attributed to artefacts or social practices.

things that performed F over things that performed G because of some advantage of F over G. Finally, (3) claims that a partial explanation of why d exists refers to the fact that its ancestors performed F and they were selected because of this performance. The paradigmatic case that exemplifies EF is the heart: the function of my heart is pumping blood because (1) past hearts pumped blood, (2) hearts were selected by natural selection because they pumped blood (and not, for instance, because they made characteristic noises) and (3) the fact that past hearts pumped blood and were selected for it partially explains why my heart exists (and why it is supposed to pump blood).

2.2 Kinds of Functions

Notice that F may be very complicated. A simple interpretation of F in EF falls short of explaining the workings of current living beings such as insects or squirrels, let alone human minds.

To account for the complexity of current living beings, at least two different effects may play the role F in EF:

1. F is some concrete activity: Hearts pumping blood, livers filtering wastes from blood or lungs extracting oxygen from air and expulsing CO_2 are particular performances in the intended sense. These are the actions that were selected in past hearts, livers and lungs. When F corresponds to such a particular activity, I will say that a device has a *direct* function (following Millikan, 1984, ch.2).[3]

2. F is a rule that yields different activities in accordance with diverse circumstances: the paradigmatic case is the chameleon's pigment-rearranging device. The chameleon changes the color of his skin so as to match the color of whatever he is sitting on. The pigment-rearranging device acts (roughly) in accordance with the following rule: *adopt the color of whatever the chameleon is sitting on.* Notice that the historical cases that account for the current existence of the device may vary dramatically from occasion to occasion; some colors may have happened just once. What all chameleons (past and present) share is a common rule that generates different skin colors in accordance with diverse external circumstances. When F falls under this description, I will say that a device has a *relational* function (Millikan, 1984, ch.2).

An important thing to notice is that a device can have many different functions at the same time. In particular, a mechanism can have two functions, one direct and one relational. Consider again the chameleon's example. On a more proximal explanation, the pigment-rearranging device acts according to the rule *adopt the color of whatever the chameleon is sitting on.* The fact that the mechanism has historically acted in accordance with this rule explains why it has been selected. So the mechanism has it as a relational function. A the same time, a more complete explanation should mention the fact that the colors adopted helped the chameleon to camouflage from predators. This is an ef-

---

[3] Millikan's expression is 'direct proper function'

fect of the mechanisms that should be mentioned in a full account of why the device proliferated. So a distal explanation of why the proximal rule has been fitness-enhancing appeals to the fact that the pigment-rearranging device has produced a color that helped the chameleon to camouflage from predators. Consequently, *camouflaging the chameleon* is a concrete effect (like pumping blood) that should be considered a direct function of the mechanism. In this case, the direct function lies, so to say, *beyond* the relational function.[4]

3. Selection in species and selection in individuals

From EF, it follows that a device can have an etiological function only if it has been recruited by a selection process that fulfills condition 2 above. In that respect, the pre-eminent process that has attracted the philosopher's attention has traditionally been natural selection. There is a good reason for that: in nature, most of the devices that have some function evolved under the pressure of natural selection. The fact that it has been the major process of evolution on Earth gives it a crucial role in explaining many systems and organisms.

However, if it is a platitude that natural selection has been the principal motor of change, it is not less obvious that it alone cannot account for all kinds of animal behavior. When seeking to understand living beings, individual experience can be as important as genetic endowment. Consequently, we need to consider whether the process of learning can be incorporated into our teleonomic account.

3.1 Is learning a selection process?

As we said before, we will focus on the learning process by means of which systems acquire novel skills. The question now is how a system produced by natural selection can incorporate abilities by experience. Some people have endorsed the following claim:

> (T1) Some learning consist in a selection process by means of which systems acquire new direct functions.

I take it that Papineau (1987, 1993, 1998) and probably Dretske (1986, 1988) would accept T1. On the other hand, Price (2001) and Kingsbury (1998, 2008) present some problems for T1. Without departing from a teleonomic theory but in opposition to Papineau and Dretske, they lend support to the alternative view:

> (T2) All kinds of learning consist in developing already existing relational functions (acquired by natural selection).

We are going to see which are the reasons for accepting T1 and then which objections Price and Kingsbury raise against this view. We can advance that our conclusion will be to accept T1 and T2. We will see how T1 *and* T2 can be held at the same time without contradiction.

---

[4] That is what Millikan (1984, pp. 41-2) calls 'invariant derived proper functions'. I think it is better to treat them as a special kind of direct functions.

It is worth emphasizing that the principal learning process discussed in the literature on animal psychology is operant conditioning. For this reason, the controversy as to whether learning can be considered a proper selection process focuses on this kind of learning (Walker, 1987; Kingsbury, 1999, p.78).

3.1.1 Acquiring new functions

By operant conditioning we (roughly) mean the process of increasing the probability of the occurrence of a given behavior by means of the appearance of a reinforcer every time the behavior is produced (Shettleworth, 2010; Walker, 1987). Through operant conditioning an association is made between a behavior and its consequences.

The main idea leading to T1 is that learning (in the sense of operant conditioning) can be considered a selection process parallel to natural selection (Papineau, 1987). The striking similarities between the two are apparent; in operant selection the organism performs (more or less) random activities until one of them gets rewarded. When that happens, it enhances the tendency to reproduce the same behavior and every time the organisms gets a beneficial effect the probability of repetition increases. On the contrary, every time that an expected reward does not show up, the likelihood of the relevant behavior being reproduced decreases, a process known as 'extinction' (Hull et al., 2001).

Therefore, it seems that operant conditioning fulfills condition (2) in EF: there is variation (diverse behaviors), replication (reproduction of the same action) and environmental interaction (which causes differential reproduction). So at first glance it seems that we have every reason to regard operant learning as the right kind of case in which new functions may arise.

It must be pointed out that in operational conditioning selection occurs both at the behavioral and neuronal level (Hull et al., 2001). I am going to talk indifferently of behavior or neural connections getting fixed by experience. For our purposes, they are two ways to refer to the same process.

A more complicated issue is whether representation takes place here and which is its role. This is an open and controversial question that still deserve a satisfactory answer. Generally, naturalist theories of representation confront the problem of how content can be relevant at all when explaining behavior, because they want their explanations to rely exclusively on physical causal relationships. Dretske (1988) offered an attractive answer, but some philosophers have pointed out several difficulties (Jaegwon, 1991; Saidel, 2001). This is a very interesting and important question very closely related to the subject of this essay, but unfortunately following this path would lead us too far from our discussion. In this essay, I will take for granted that a purely physical description can explain an organism's behavior, which is fully compatible with the existence of parallel descriptions at other levels that appeal to contentful representations. Therefore, in my account I will dispense with representations (as such).

3.1.2 Different ways of learning

Encouraged by the similarities between operant conditioning and natural selection, some people have tried to extend this result to other kinds of learning. For instance, Papineau claims that when we acquire new beliefs they typically undergo a kind of selection process that fulfills condition 2 above (Papineau, 1987, 1998). This is probably an excessive extrapolation. As Kingsbury (2008) and Price (2001) point out, 'learning' comprehends a great amount of different processes: imitation, language learning, acquisition of new recognitional capacities,... and, even in case learning could provide an organism with new functions by means of an idiosyncratic selection process, many novel abilities and knowledge that we include under the label 'learning' cannot be explained in the same way.

Consider language learning. Kingsbury (2008) claims that for a proper selection process to take place in the acquisition of language, the child should have to receive positive or negative feedback depending on whether the sentence he produces is well or ill-formed (this is one of Papineau's examples (1987, p. 67)). Thus, every time the child says 'I goed' instead of 'I went' the parents should stop him and correct his usage, and they should do that with every word and structure learned. Unfortunately, Kingsbury reminds us this model of language acquisition was discredited long time ago. Since Chomsky (1959) published his review of Skinner's *Verbal Behavior* and developed his alternative model, it has been widely assumed that negative evidence does not play any significative role in language acquisition, and positive evidence is so scarce and defective that it alone cannot account for the quick and almost perfect learning of grammar and lexicon. So to the extent that language is learned (and, of course, even in a Chomskian paradigm, there is something learned) we cannot consider it a selection process that fits in EF.

Price's (2001, p. 129) example concerns learning by observation. A lion cub may learn how and what to prey by observing older lions hunting. In this way, he may learn to hunt long before he is ready to do it by himself, and hence before any selection process may take place. In a similar vein, Kingsbury (2008, p. 504) points out that concept learning occurs when the baby is still not able to perform any behavior that may yield a reward or a punishment. Thus, in both lion cubs and babies, new abilities are acquired 'blindly', without feedback. In those cases learning lacks the required environmental interaction (condition 2c in EF). In view of this evidence, they conclude that most learning cannot be regarded as a selection process, but just as a development of already existing devices that have been recruited by natural selection.

I agree that learning can take place without the action of any selection process. As we said in the introduction, 'learning' includes many different kinds of processes, that range from the simple acquiring of new knowledge to the more complex development of novel and intricate skills.

Consequently, I grant that not all episodes of learning are selection processes. In particular, some examples like language learning clearly do not consist in selection processes. And nonetheless, this contention does not exclude that *some* cases of learning actually operate in a parallel fashion as natural selection. This is the reason we constrained T1 to just *some* learning processes. In my opinion, Kingsbury and Price con-

vincingly show it is wiser not to take natural selection as a general model for learning. I think this is the right conclusion we should keep in mind, and it still leaves us with the question of whether some learning (for instance, operant conditioning) may be better understood as a selection process.

3.1.3 Mislearning

However, Price (2001, p. 128) has an argument against the claim that some kind of learning provide new functions. According to her, the problem with T1 is that 'it fails to accommodate the possibility that an organism makes a mistake in learning'. Consider the following case:

> **(Frederick)** Frederick the mouse got sick after eating a leaf of *Aconitum napellus*, a toxic plant. As a consequence, he learned to avoid aconite. After this experience, a new connection has been wired in Frederick's nervous system linking the olfactory sensation of aconite and the avoidance response. The acquisition of this learned behavior can be readily explained by the fact that in past occasions eating aconite leafs had undesirable consequences, and as a result the avoidance behavior has been incorporated in Frederick's repertoire.

Now consider a parallel case:

> **(Natasha)** Natasha the mouse got sick after eating blackberries from a nearby bramble bush. As a consequence, she learned to avoid blackberries. However, it turns out that blackberries she ate made her sick because it rained acid just that morning. Unfortunately for her, blackberries are the most nutritious fruits in Natasha's diet.[5]

Both Frederick and Natasha learned to avoid something particular through a similar experience. Nevertheless, we want to say that something went *wrong* in Natasha's case. It seems that Natasha's learning was a *mistake*. But certainly, this rule cannot derive from the learning process itself, because the selection process occurred in the same fashion in both cases: in both Natasha and Frederick's examples there was repetition, variation and environmental interaction. How can there be a normative difference if the selection process is what determines what an organism *is supposed* to do?

Two important assumptions are made here. On the one hand, by means of a learning process the organism can acquire a disposition that does not contribute to its welfare. Indeed, the new ability may run against other biological purposes, as we saw in Natasha's example. Learning the wrong thing might be worse that not learning anything at all. This point has been worked out by Godfrey-Smith (1996), who suggests that the capacity of learning enhances dramatically the organism's adaptability to a complex environment, but only at the high cost of making possible the acquisition of the wrong knowledge or skill which eventually may have fatal consequences.

Secondly, according to teleonomy the only way an organism can be supposed to do something is by having a function. Normativity arises because there are functions, and functions solely emerge in mechanisms that have the right kind of selective history.

---

[5] Similar examples can be found in Millikan (1993, p. 227) and Price (2001, p. 128).

For this reason, Natasha cannot learn the wrong thing unless there is a biological function that determines which should be the result of the learning process. From this argument, Price (2001, p. 130) concludes that the capacity of acquiring new skills by the process of learning must be considered a way of developing the biological functions recruited by the organism through natural selection. Thus,

> The most plausible response to cases of this kind is to insist that these devices can be ascribed intentional content, even before they give rise to useful behaviour, *and that their content is fixed by the design function of the mechanisms that govern the process of learning.* (Price, 2001, p. 29. Italics are mine)

Price suggestion raises some questions. How can the fact that Frederick is supposed to avoid aconite be determined by nature? How do the innate functions of the device explain the possibility of mislearning? We need to answer these questions before we can fully understand Price's position. Even if she does not develop in more detail her own view on learning, we will try to make sense of it using some of the tools we have already presented.

3.1.4 Taste Aversion Learning

To understand Price's example, we need first to know a bit more about Taste Aversion Learning, the process involved in Frederick's and Natasha's examples.

Biologists tell us that there are innate and learned food aversions. Arguably, both have the same goal: trying to avoid that organisms ingest poisonous food. Among the inborn mechanisms that serve this purpose, there is aversion to sour and bitter tastes, which warn of sour or spoilage. Taste Aversion Learning is a way to discover new dangerous substances by means of unpleasant experiences.

The capacity to acquire taste aversions from unpleasant experiences is present in an enormous range of animal species (humans included), and it works approximatively in the same way in all of them, what suggests that this system is quite primitive. Moreover, some evidence shows that it is an unconscious process, taking place quite automatically, despite the conscious efforts of preventing it (Bernstein, 1999).

There are some features that make Taste Aversion Learning quite special. First of all, in contrast to other learning mechanisms found in animals, it is a very quick acquisition process. Usually, one trial is enough for establishing a connection between a particular cue and an aversion. Secondly, this mechanism is specially sensitive to flavours (odor and taste). If a rat eats something that makes him sick, he will avoid things that *smell* alike, not things with the same color, shape or location. (Shettleworth, 2010). So, in Taste Aversion Learning mice act according to the following proximal rule:

(P) If ingesting x makes you sick, do not eat anything that smells like x

This is a proximal explanation of the mice's past performances that accounted for the selection of the mechanism. On the other hand, this proximal rule was beneficious for mice because usually enough past mice avoided the ingestion of poison by following it. So there is a more distal explanation that appeals to the fact that this mechanism has helped mice to avoid poisonous food. Consequently, the device has a direct (and distal) function that goes beyond the relational (and proximal) function:

(D) To avoid poisonous food.

Now we can understand how this theory is supposed to explain mislearning. An organism mislearns when a particular behavior that satisfies the proximal function does not satisfy the distal function. In Frederick both rules fit together: the proximal rule tells him no to eat aconite because it made him sick, and the direct function that lies beyond the proximal rule tells him not to eat aconite as well because it is something poisonous. So this is not a case of mislearning.

However, Natasha's case is different. The proximal rule tells her not to eat blackberries, because she got sick after eating some. But it does not fulfill the direct (distal) function because it will not make her avoiding poisonous food and, in fact, it will reduce her fitness. This is a case of mislearning because what she learns does not fulfill its distal function: avoiding poisonous food and thus enhancing her fitness.

Why is that a problem for T1? Crucially, in this example both P and D are *innately* determined. That is: what mice are supposed to do is determined by the inborn mechanisms that have been shaped by natural selection. In Natasha's case, mislearning is possible because the result of Taste Aversion Learning does not comply with the more distal biological function. So, mislearning is possible because fulfilling the (innate) function of the learning mechanism precludes satisfying another (innate) function of the device. This is how we can make sense of Price's objection. If all normativity we find in learning processes is explained by innate functions, there is no place for new acquired functions. This is a serious challenge for T1.

3.1.5 How mislearning is possible

In the discussion of Frederick and Natasha's examples, we have seen that the normativity that arises in learning can be fully explained by innate functions. This is possible because two different kinds of functions can coexist in a single mechanism: a relational function (in a proximal explanation) and a direct function (in a distal explanation). The fact that we can make a mistake by learning shows that learning is not completely independent from history and the general workings of the system. If we want to make room for mislearning, we cannot detach learning processes from the function that systems have acquired by natural selection.

Now, our project is to explain normativity; this is why we appealed to functions in the first place. But if all normativity we find in learning is explained by innate functions, can we then conclude that no learning can incorporate new functions? I do not think so. This conclusion would be precipitated. The fact that natural selection determines to some extent what should be learned, does not imply that no other selection processes can exist. Price has shown that the possibility of mislearning is linked to innate functions. This idea brings two challenges for T1: one is to explain which sort of normativity (if any) is generated in learning processes. The second is to show how T1 can be true and mislearning be possible.

I think both challenges can be met. There is a way in which both natural selection and the selection process taking place in learning can determine what the system is supposed to do. The remainder of the paper is devoted to explaining this account.

3.2 How biology incorporates learning

Our present task is to present a case in which new functions arise (and normativity is generated) and nevertheless mislearning is possible. So we need to find a case in which the learned behavior cannot be fully accounted for just appealing to innate functions. Let us consider a classical case of operant behavior:

> **(Eric)** Eric the rat was set in a chamber (a Skinner box) with a light and a lever. The chamber was so designed that when the light was on, pressing the lever yielded food. Eric was first roaming and inspecting the chamber. By chance, he pressed the lever while the light was on and obtained the reward. After some more or less fortuitous repetitions, Eric's pressing the lever became correlated with the light. The probability of the behavior occurring when the light was on and followed by the ingestion of food increased.

After some trials, Eric learned to press the lever when the light was on. Given the light, he was supposed to behave in a certain way. But, in this case, there is no particular rule that explains Eric's action. We cannot explain Eric's behavior appealing to a rule that tells him something like *If pressing the lever while a light is on yields food, do it*. First, because levers and lights (or a correlation between levers and lights) have played no role in the evolutionary history of rats. Secondly, because any positive reward (not just food) would have produced the learning of the behavior. In this case, the proximal explanation for the proliferation of the mechanism should appeal to the fact that past mice performed whatever behavior yielded them a reward. Thus, the only proximal rule we can appeal to is something like *do whatever gets reinforced* (reinforcers being the production of pleasure, the avoidance of pain,.. or maybe: the segregation of certain chemical substance in the brain). Of course, 'whatever' has some natural constraints: the rat is not going to fly. But the point is that, in contrast to cases of Test Aversion Learning, there is no particular action the rat is supposed to carry out.

3.2.1 Producing mechanisms

So I take it that the proximal rule that determines what Eric the Rat is supposed to do is something like *do whatever gets reinforced*. Conceivably, this is a especial sort of relational function, but one that offers an outstanding variability. This is why animals can learn to execute some activities that otherwise they would hardly ever perform. The fact that dolphins jump out of the water and carry out extraordinary leaps or the fact that pigeons learn to press a lever every time a certain sound is produced can only be explained if we assume a fundamental flexibility of the natural device that makes learning possible. And, probably, the main rule dolphins and pigeons are following must be something close to *do whatever gets reinforced*. I think this is the principle created by the process of natural selection that enables that kind of learning. Notice that this rule is remarkably flexible because many different behaviors can fill in the variable 'whatever' and also because the possible reinforcers are multiple.

The point I want to make here is that *do whatever gets reinforced* is just another form to say *acquire a new function*. A reinforced behavior is a behavior that has acquired a new function. As we saw, in operant conditioning reinforcement is a process analogous to

natural selection, one that satisfies all conditions required for giving rise to new functions (EF, 2). If a repeated behavior is appropriately rewarded enough amount of time, the organism can adopt it as the activity that fulfills the biological function of the learning mechanism. The particular behavior that gets reinforced (for instance, by avoiding pain) may be fixed in the organism by a selection process that confers new functions: this is what explains why Eric is supposed to press the lever. Natural selection equips some organisms with a device that generates diverse behaviors according to different selection processes.

Basically, the idea is that the function of certain learning devices is that the organism acquire new functions if a selection process that fulfills condition 2 in EF takes place. This is their biological function. Consider again the example of the rat; depending on his experience, Eric will learn to press the lever or to somersault. Those are behaviors he will learn only if they got reinforced by the appropriate selective story. Let us see now how in this case new normativity arises and, at the same time, mislearning is possible.

3.2.2 Operant conditioning and mislearning

The final picture we get is the following. In some learning situations, mice (and many other animals) act following the proximal rule: *do whatever gets reinforced.* As we just saw, this is another way to say *acquire a new function.* So, there is a biological (proximal) rule that drives animals to behave in certain ways if they are appropriately rewarded. By means of that process, they acquire new functions, and hence they are supposed to do novel things (unexpected for biology). So the new selection process is what determines which particular action the animal is supposed to do.

On the other hand, all these learning mechanisms have had the same distal effect, and having this effect explains why the mechanisms were selected at all. Some behaviors helped the animals getting food, others avoiding cold, etc... what all have in common is that they increased the fitness of the organism. So the proximal rule has been incorporated into the mouse's repertoire because often enough following it has increased the fitness of the organism. Probably this is the most detailed description we can get that accounts for all the positive effects of the different behaviors in the evolution of the animals. So something like *increasing the fitness of the organism* is a function (in a distal explanation) of this learning mechanism.

Now, we can understand how mislearning is possible even in operant learning. Even if what an animal is supposed to do derives from the particular history of reinforcement, there is always a more distal function that the mechanism is supposed to satisfy, *namely increasing the fitness of the organism.* When a particular learned behavior fails to fulfill this function, mislearning occurs. This is how mislearning and operant conditioning are both possible.

4. The complexity of learning

I think now we have the principal elements required to understand the increasing complexity generated by learning processes and the normativity we find in there. In

most cases (for instance, Taste Aversion Learning) appealing to natural selection will be enough. However, in other cases the innate mechanism just determines that there must be a selection process (say, operant conditioning) that gives raise to new functions. Moreover, I have argued that even if operant conditioning is a selection process that fulfills condition 2 in EF, mislearning is also possible.

At this point, I think we can conclude that T1 and T2 can be held at the same time. T2 is true because, at least in all cases we have considered, learning consists in developing already existing relational functions. Even operant conditioning consists in developing a quite general relational function, *namely do whatever gets reinforced*, On the other hand, T1 is also true because in operant conditioning what a system is supposed to do is determined by a selection process that confers new functions. So, in some cases, learning a new skill or ability may consist in fulfilling a function recruited by natural selection and, at the same time, it may include a selection process that generates new functions. Biology and experience complement each other.

I hope I have also been able to show how the capacity to learn a new skill can be included in a teleonomic framework. Previous analysis either have decoupled learning processes from the organism's evolutionary history or have thought that besides natural selection no other selection process could give rise to new functions. I have tried to show that natural history and personal experience may be unified in a teleonomic account.

In a more general perspective, I think learning is the key process that will give us the clue for understanding complex systems in nature. Of course, it is a triviality that we are part of nature, but filling out this idea requires, among other things, explaining the process of learning. This essay is intended to be a short step in this direction.

REFERENCES

Ayala, F. J., 1970. Teleological Explanation in Evolutionary Biology. *Philosophy of Science*. 37 (1), pp. 1-15

Bernstein, I. L., 1999. Taste Aversion Learning: A Contemporany Perspective. *Nutrition, Immunology, Neuroscience, and Behavior*. 15 (3), pp. 229-234.

Chomsky, N., 1959. Review of B. F. Skinner's Verbal Behavior. *Language*. 35, pp. 26-58

Dretske, F., 1986. Misrepresentation. In Bodgan, R., ed. *Belief*. Oxford: Oxford University Press.

Dretske, F., 1988. *Explaining Behavior*. 1st ed. Cambridge (MA): MIT Press.

Godfrey-Smith, P., 1996. *Complexity and the Function of Mind in Nature*. 1st ed. Cambridge. Cambridge University Press.

Hull, D. Lanfman, R. and Glenn, S., 2001. A General Account of Selection: Biology, Immunology and Behavior. *Behavioral and Brain Sciences*. 24 (2), pp. 511-527

Jeagwon, K., 1991. Dretske: On How Reasons Explain Behavior. In McLaughlin, B., ed. *Dretske and his Critics*. 1st ed. Cambridge (MA): Basil Blackwell.

Kingsbury, J., 1999. *Content and Function: A Defense of Millikanian Teleosemantics*. PhD. Rutgers University.

Kingsbury, J., 2008. Learning and Selection. *Biology and Philosophy*. 23, pp. 493-507

Millikan, R. G., 1984. *Language, Thought and Other Biological Categories*. 1st ed. Cambridge (MA): MIT Press.

Millikan, R. G., 1993. *White Queen Psychology and Other Essays for Alice*. 1st ed. Cambridge (MA): MIT Press.

Papineau, D., 1987. *Reality and Representation*. 1st ed. Cambridge (MA): Basil Blackwell.

Papineau, D.,1993. *Philosophical Naturalism*. 1st ed. Cambridge (MA): Basil Blackwell.

Papineau, D., 1998. Teleosemantics and Indeterminacy. *Australasian Journal of Philosophy*. 76 (1), pp. 1-14

Price, C., 1998. Determinate Functions. *Nous*. 32 (1), pp. 54-75

Price, C., 2001. *Functions in Mind*. Oxford: Oxford University Press.

Saidel, E., 2001. Teleosemantics and the Epiphenomenality of Content. In Scott McIntosh, J., ed., *Naturalism, Evolution and Intentionality*, Calgary (Canada): University of Calgary Press.
Shettleworth, S., 2010. *Cognition, Evolution and Behavior*. Oxford: Oxford University Press.
Walker, S., 1987. *Animal Learning: Introduction*. London: Routledge and Kegan Paul.
Wright, L., 1973. Functions. *Philosophical Review*. 82, pp. 139-168
Wouters, A., 2005. The Function Debate in Philosophy. *Acta Biotheoretica*. 53 (2), pp. 123-151.

Marc ARTIGA is a member of the LOGOS Group and a PhD candidate at the Universitat de Girona. He is working on the Philosophy of Mind and the Philosophy of Biology and his main topic of research is on naturalistic theories of mental representation.

ADDRESS: Departament de Filosofia, Facultat de Lletres, Universitat de Girona. Plaça Ferrater i Mora, 3 Office 4.03. 17071 Girona (Spain). Email: marc.artiga@udg.edu