

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Interaction of goal-directed and Pavlovian systems in aversive domains

Permalink

<https://escholarship.org/uc/item/99v709sm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN

1069-7977

Authors

Rigoli, Francesco
Pavone, Enea Francesco
Pezzulo, Giovanni

Publication Date

2011

Peer reviewed

Interaction of goal-directed and Pavlovian systems in aversive domains

Francesco Rigoli (francesco.rigoli@istc.cnr.it)

University of Siena, Piazza S. Francesco 7, 53100 Siena, Italy
ISTC-CNR, via San Martino della Battaglia, 44, Roma, Italy

Enea Francesco Pavone (eneafrancesco.pavone@uniroma1.it)

University of Rome "La Sapienza", Roma, Italy

Giovanni Pezzulo (giovanni.pezzulo@istc.cnr.it)

ILC-CNR, via Giuseppe Moruzzi, 1, 56124 Pisa, Italy
ISTC-CNR, via San Martino della Battaglia, 44, 00185 Roma, Italy

Abstract

Recent neuroscientific models of human behavior distinguish between different cognitive controllers: two instrumental systems (goal-directed and habitual) that maximize utility through learned actions, and a so-called Pavlovian system, which implements innate reactive responses. Although the interaction between instrumental and Pavlovian controllers has been suggested as a key process underlying emotional phenomena and surprising forms of misbehavior, few is known about it, especially in the sensorimotor aversive domain. With a combined experimental and computational approach, we study the interactions between instrumental (goal-directed) and Pavlovian processes in the aversive domain. First, we present a human experiment in which goal-directed and Pavlovian systems compete in order to control responses. The results indicate that Pavlovian processes can significantly interfere with goal-directed behavior. Second, we compare four alternative Bayesian models for their accuracy in modeling human performance. The results indicate a better fit for an architecture in which the Pavlovian controller can use both model-based and model-free features.

Keywords: Goal-directed system, Pavlovian system, Bayesian model, implicit classical conditioning.

Introduction

Recent models of human and non-human cognitive control assume that behavior depends on the joint activity of multiple cognitive systems (Dayan and Seymour, 2008): a so-called "multicontroller" framework. The existence of interacting behavioral controllers is supported by numerous data in the animal literature (Adams and Dickinson, 1981). For instance, in experimental paradigms such as devaluation and contingency degradation, Dickinson (1985) showed that instrumental behavior of rats usually depends on associations between actions and outcomes, and on associations between outcomes and their current assigned values. However, in some circumstances, namely when action is simple and overtrained, rats follow simpler stimulus-response associations. The former pattern of behavior was thought as being controlled by a *goal-directed system* (GDC), the latter by a *habitual system* (HAB). Although they adopt different mechanisms, both these instrumental controllers aim at maximizing utility.

Moreover, a third type of controller has been proposed: a *Pavlovian system* (PAV) (Dayan and Seymour, 2008). Animal studies have demonstrated the existence of complex behavioral patterns that are rigid and inbuilt in nature, as they are not learned through the ontogenetic development (Fendt

and Fanselow, 1999). These so-called Pavlovian responses involve either the avoidance or the approach of salient biological stimuli (unconditioned stimuli, US) and cues associated to them (conditioned stimuli, CS). Since the rapid activation of these behavioral patterns has increased the fitness of the organism during evolution, Pavlovian mechanisms are inherited. The importance of PAV in affecting behavior has been shown in animal experimental paradigms in which the activation of Pavlovian responses dramatically impaired goal-directed performance. For instance, Hershberger (1986) showed that chickens do not learn to run away from a source of food that they can see through a "looking glass" in order to get access to it. A similar paradoxical effect was observed in an experimental procedure called *negative automaintenance* (Williams and Williams, 1969). Food was given to pigeons when they abstained to peck a CS; nonetheless, these animals were unable to learn this instrumental contingency and continued to peck it.

In this article, we address two open issues within this literature, both relative to the nature of PAV. The first one concerns the interactions between PAV and instrumental controllers, in particular GDC, that has been considered a key process underlying many emotional phenomena (Seymour and Dolan, 2008) and a source of maladaptive behavior such as impulsivity, dread and framing effects (Dayan and Seymour, 2008). While Daw et al. (2005) proposed that GDC and HAB compete on the basis of their uncertainty, less is known about the GDC-PAV and HAB-PAV interactions.

In relation to the second controversial issue, it is unclear which are the computational mechanisms adopted by PAV. From a computational perspective, GDC and HAB have been associated to model-free and model-based methods of reinforcement learning, respectively (Daw et al., 2005). Model-free models learn and use cached values of future states or action-value associations, without explicitly representing action-outcome and outcome-value relationships of the environment, as model-based models do. Because PAV interacts more with HAB than GDC (Holland, 2004), it might have a model-free architecture similar to HAB (Dayan et al., 2006). Alternatively, PAV could exploit the predictions done by both GDC and HAB (Dayan and Seymour, 2008).




Figure 1: The experimental paradigm. In blocks 1 and 2, the shock is associated to one of two colors (CS), yellow or red. In block 3, the shock is never delivered.

GDC vs. PAV: an experimental study

The aim of this study is investigating the interactions between GDC and PAV in humans. In keeping with the aforementioned literature, we hypothesized that GDC and PAV could be (partially) distinct processes, which can compete by triggering two distinct motor commands: the former instrumental to goal achievement, and the latter triggered by the detection of a CS. When the two motor commands are in conflict, GDC and PAV may compete to control behavior, and hence instrumental performance may be impaired.

We tested this hypothesis by comparing the performance of human subjects in a sensorimotor task consisting in pressing a button at the right time: when a ball moving from one side of a screen to the opposite one is in the target (see fig. 1). In the experimental condition, a CS (a color, either yellow or red) anticipating the delivery of an US (an electric shock to the pressing finger) was presented, while in the other condition a neutral stimulus (NS, the other color) was presented. Subjects were informed that they would have received the electric shock if they were unable to press the button at the right time, but only in the presence of CS, not NS. We hypothesized that, since in the first (but not the second) condition GDC and PAV compete for generating conflicting motor responses (“push” for winning the game vs. “raise the finger” as an innate avoidance response, respectively), performance should be worse in the first condition. Note that this would have the paradoxical consequence that a Pavlovian response actually determines more electric shock, making our study similar in spirit to the aforementioned “negative automaintenance”.

To study the interaction of HAB vs. PAV (rather than GDC vs. PAV) in the same framework, we designed a second experimental block of trials (block 2) identical to the first one, with the hypothesis that after a sufficient number of trials the response becomes habitual.

Finally, we were interested in studying if, after a CS-US association has been repeated for many trials, CS activates the conditioned response (CR) even in the absence of an explicit expectancy of US (manipulated through verbal instruc-

tion). Although some findings suggest that CS can trigger directly CR (without the US expectancy), however this hypothesis is controversial (Lovibond and Shanks, 2002). To this aim, we designed a final experimental block of trials (block 3) in which US (the electric shock) was never delivered and participants were informed that they would have never received it in that block.

Methods Eight participants (5 females, 3 males; average age: 27) were tested in one session. During the session, each participant was sitting in front of a computer screen. The task consisted of 3 blocks, 80 trials for each block (fig. 1). At every trial, a colored circle appeared in the center of the screen. The circle was red for half of the trials and yellow for the other half. After two seconds, a ball of the same color of the circle appeared on one side of the screen and moved towards the opposite side, passing through the circle. The speed of the ball varied randomly trial by trial in all three blocks. The participants had to press a button when the ball was in the circle. Participants were divided into two groups. Participants from group A received an electric shock when they made a mistake (they either pressed too slowly or too fast) in the presence of the red circle (CS), and never in the presence of the yellow circle (NS). Conversely, participants from group B received the electric shock when making a mistake in the presence of the yellow circle (CS), and never in the presence of the red circle (NS). The intensity of the shock was established by each participant before the experiment following the same procedure as in (Ohman and Soares, 1998). In block 3, participants of both groups never received an electric shock. At the beginning of block 3, participants were informed of the absence of shock delivery.

Results A repeated-measure ANOVA was conducted (factors: Block, 3 levels: 1, 2, 3. Stimulus, 2 levels: CS and NS). The dependent variable was task performance (percentage of winnings). The main effect of block was not significant, $F(2,14) = 3,011$, $p = 0,089$ ($p > 0,05$), $\eta_p^2 = 0,3011$. On the contrary, the presence of CS significantly impaired performance, $F(1,7) = 25,94$, $p = 0,001$ ($p < 0,05$), $\eta_p^2 = 0,787$. Interaction between Block and Stimulus is not significant, $F(2,14) = 0,139$, $p = 0,890$ ($p > 0,05$), $\eta_p^2 = 0,019$. This result indicates that the main effect of CS is stable in the three blocks. Fig. 2 plots performance in *Block x Stimulus* conditions.

Discussion We reasoned that, during the task, the subjects’ goal was to win the game and/or to avoid the shock (both give rise to the same pressing action). Therefore, if behavior was controlled only by GDC, performance should have been the same in both CS and NS conditions in the case goal was winning the game (and perhaps even better in CS condition in the case the subject’s goal was that of avoiding the shock). Both hypotheses are contrary to our findings. Indeed, we found that the presence of a CS significantly impaired performance. In keeping with the multicontroller perspective, we interpreted this result as follows: in the presence of CS, a Pavlovian con-




Figure 2: Results: interaction *Block x Stimulus*.

troller triggered a specific CR of finger withdrawal (as the US associated to the CS corresponded to a shock delivered to the finger), which interfered with the more adaptive pressing response selected by the GDC¹.

A second result is that we did not find habituation in the three blocks (probably for weak statistical power). This does not allow us to study interactions between HAB and PAV. To this aim, it is necessary to test more subjects.

Finally, we found that the effect of CR activation on performance was present in all of the three blocks. This result raises interesting issues on the computational nature of PAV, that we discuss in the following section through a computational theoretical proposal. Our model is based on the reported data and on two further assumptions that, although being controversial, have received some empirical support. First, we assume that verbal instructions about CS-US contingency modify model-based causal representations in the brain, and that those representations can be flexibly updated by subsequent verbal instructions (Lovibond and Shanks, 2002). Therefore, when the CS-US association is verbally induced (first and second blocks, but not the third one), PAV might involve a model-based component. Second, we assume that learning the CS-CR association requires a slow trial-and-error procedure (as for HAB) (Konorski, 1967), hence in the first block that association should be absent or weaker. Since our data suggest that the effect of PAV on performance is stable along the three blocks, it might have a double nature, both model-free and model-based.

Computational model of GDC vs. PAV

To further investigate the issue of a model-based, model-free of hybrid nature of PAV, we designed four alternative computational models of the interaction between GDC and PAV, and compared them in a simulated scenario equivalent to the human study described so far.

We implemented the four models (shown in figs. 4, 5, 6, and 7) using the Dynamical Bayesian Network (DBN) formalism. Bayesian networks represent relevant variables as nodes, and probabilistic relations between variables as edges.

¹As discussed in the conclusions, we are conducting a control experiment to rule out the possibility that impaired performance was due to unspecific effects such as shivering or anxiety.




Figure 3: Illustration of state and action transitions




Figure 4: Model 1. Three time steps are reported. Bold lines represent the GDC; dotted lines represent the environment.

DBNs add to this an explicit representation of temporal dynamics; indeed, in the figures, three temporal time steps are shown (e.g., the three S in the figures are relative to time t , $t + 1$ and $t + 2$). While most Bayesian models in cognitive science incorporate only cognitive components, our models integrate environmental dynamics (dotted lines) and dynamics of multiple cognitive controllers, in our case GDC and PAV (bold lines and thin lines, respectively), with their associated cognitive variables. Variables used in the models are shown in tab. 1 and explained in the main text.

Table 1: Variables used in the models, see figs. 4, 5, 6, and 7

var	where	explanation	values
E	world	environmental state	1-18
B	world	agent behavior	press, no press
S	GDC	state as cognitively represented	1-18
O	GDC	observations	1-18
A	GDC	selected action	press, no press
US	PAV	unconditioned stimulus	yes, no
CS	PAV	conditioned stimulus	yes, no
CR	PAV	conditioned response	press, no press

Model 1, shown in fig. 4, includes only GDC and its interaction with the environment (ENV). *States S* represent the position and speed of the ball in the screen. This variable can assume 18 values, resulting from the combination of 9 spatial positions and 2 speeds ("still" and "moving"); see fig. 3. *Actions A* represent the goal-oriented action and can assume two values: "press" and "no press". For instance, pressing in S1 leads to S10; nor pressing leads in S2, and so on. (Note that, once the button has been pressed, the ball remains in the




Figure 5: Model 2. Bold lines represent the GDC; dotted lines represent the environment; thin lines represent the PAV.




Figure 6: Model 3. Bold lines represent the GDC; dotted lines represent the environment; thin lines represent the PAV.

same position for all of the following time steps and the trial ends.) In Model 1 (and all the other models) the probability of pressing is maximal when $S = S5$ and it is inversely proportional to the distance from this value. Behaviors B represent the effective behavior executed by the efferent motor system and they can assume the same values as A : "press" or "no press". Noteworthy, in this model B depends only on A . Environments E represent the real state of the world and depend both on the previous E and on B . Observations O represent the perceptive processes and depend both on E and S . Both E and O can assume the same 18 values as S .

Model 2, shown in fig. 5, represents the interaction between GDC and a model-based PAV (plus the environment). PAV is model-based in that CR is triggered only if the presence of CS recalls the explicit expectancy of the US . This version adds three variables to the previous model. The first variable is the *Conditioned Stimulus* CS and detects the presence or absence of CS . The second variable is the *Unconditioned Stimulus* US and detects the presence or the absence of US . It depends on CS . Specifically, the conditional probability of US depends on the explicit belief that CS and US are associated. Finally, the *Conditioned Response* CR represent the activation of CR and can assume the same values as A ("press" or "no press"). CR is activated every time US is expected. Crucially, in this model, as well as in the following ones, B depends jointly on A and CR .

Model 3, shown in fig. 6, implements a model-free PAV. This model is similar to the second one, except that CS and



Figure 7: Model 4. Bold lines represent the GDC; dotted lines represent the environment; thin lines represent the PAV.

US are not connected. They are both directly connected to CR . Noteworthy, in this case, US cannot be anticipated by CS , in that they are not linked. At the beginning of the task, CR dependence on CS is null and it is strengthened every time the US and the CS are presented together. In other words, after a sort of trial-and-error learning of the US - CS association, the CS can trigger the CR even if presented alone.

Model 4, shown in fig. 7, implements both model-based and model-free PAV. This was obtained by linking both CS and US , and CS and CR . The first link is active when there is an explicit belief of the association between the US and the CS ; the second link is active after that the US and the CS have been repeatedly paired.

Note that all models, transitions $S \rightarrow A$, $A \rightarrow B$, and $CR \rightarrow B$ are probabilistic. Rather, for the sake of simplicity, we designed the less relevant transitions $S \rightarrow A$, $B \rightarrow E$, $A \rightarrow S$, and $S \rightarrow S$ as deterministic (however, equivalent results are obtained by making them probabilistic).

Simulations and results

We tested the four models in a simulated scenario equivalent to the human experiment described so far. Testing consists in computing the marginal probability of the node E at time step $t = 10$. A trial is "won" if $E = S14$ (this corresponds to executing the action "press" in state $S5$); otherwise it is "lost". In other words, we recorded the probability distribution of the final position of the ball.

We set the value of variables as follows (see also tab. 2). In conditions in which CS was present, in models implementing PAV (models 2, 3, and 4), the value of the node CS was set to "yes". In blocks 1 and 2, in models where US depends on CS if they are linked (models 2 and 4), the conditional probability of the node US was set so as when $CS = "yes"$, then $US = "yes"$. At the same time, in models where CS is directly associated to CR (models 3 and 4), the conditional probability of the node CR was set so as when $CS = "yes"$, however $CR = "no press"$. Conversely, in blocks 2 and 3, in models where US depends on CS if they are linked (models 2 and 4), the conditional probability of the node US was set so as when $CS = "yes"$, however $US = "no"$. In models where CS is directly associated to CR (models 3 and 4), the conditional probability of the node CR was set so as when $CS = "yes"$,

Table 2: Value of variables and conditioned distributions in the four models (for each block and trial)

block	trial	M1	M2	M3	M4
1	CS	-	CS=yes \rightarrow US=yes	CS=yes \rightarrow CR=no press	CS=yes \rightarrow US=yes
1	NS	-	CS=no \rightarrow US=no	CS=no \rightarrow CR=no press	CS=no \rightarrow US=no
2	CS	-	CS=yes \rightarrow US=yes	CS=yes \rightarrow CR= press	CS=yes \rightarrow US=no, CR=press
2	NS	-	CS=no \rightarrow US=no	CS=no \rightarrow CR=no press	CS=yes \rightarrow US=no, CR=no press
3	CS	-	CS=yes \rightarrow US=no	CS=yes \rightarrow CR= press	CS=yes \rightarrow CR= press
3	NS	-	CS=no \rightarrow US=no	CS=no \rightarrow CR=no press	CS=no \rightarrow CR=no press




Figure 8: Performance of the four Bayesian models.

CR = "press".

Results of the simulation are shown in fig. 8. Plots of the *Block x Stimulus* interaction show that only model 4 compares well with human data (fig. 2) in all three conditions. Model 1 cannot account the finding that the performance is affected by the presence of CS. Model 2 cannot account for data relative to block 3, since the belief that US follows the CS was not present, and model 3 cannot account for data relative to block 1, since CS still was not able to directly trigger CR.

General Discussion

Goal-directed (GDC) and Pavlovian (PAV) controllers may interact at many levels. For instance, animal studies have investigated the likely influence of PAV over the values assigned to outcomes. This phenomenon emerges in paradigms such as *Pavlovian instrumental transfer*, *conditioned reinforcement* and *conditioned suppression*, in which the presentation of a CS during an instrumental task may influence the values assigned to outcomes, which, in turn, may influence frequency and vigor of instrumental actions (Holland, 2004). (Note that, in these paradigms, the influence of PAV over the goal-oriented action is indirect.)

Although it has been suggested that interaction between GDC and PAV could be at the origin of important phenomena like emotional reactions and paradoxical misbehavior (Dayan and Seymour, 2008), so far it has not been studied in humans, nor it has been studied in animals in the aversive sensorimotor case (but see a report that squirrel monkeys, punished for biting on a restraining leash, tend to increase their biting

Morse et al., 1967). This study investigates the direct influence of PAV over goal-oriented actions in humans. Specifically, we were interested in the sensorimotor conflict between GDC and PAV in the aversive domain, namely when US corresponds to a punishment. We found that, in a task where a conditioned response and an instrumental command impose opposite constraints to action selection, PAV can exert an influence over GDC, even if this paradoxically determines delivery of more punishment rather than its avoidance. In our discussion and computational model, we described interaction between GDC and PAV at the motor level, or in other words as the competition of CR and A for determining the behavioral response.

This claim is in agreement with a previous model of interaction between PAV and instrumental processes, proposed by Dayan et al. (2006). In line with many proposals framed in the reinforcement learning paradigm, this model considers PAV as a model-free process similar to HAB. In this perspective, PAV would learn to anticipate the cached values of future states through a trial-and-error procedure, and not the future states themselves, as typical of model-based systems. However, contrary to HAB, PAV would not associate a learned instrumental action to those anticipations, but rather an unlearned CR. Although this point of view is common in animal research, nonetheless it has been questioned by human studies, which suggest that CS may trigger CR only if US is explicitly expected (suggesting a model-based component of PAV) (Lovibond and Shanks, 2002). Relative to this debate, our experimental and computational study suggests that PAV may involve both model-free and model-based processes; indeed, PAV seemed to impair performance both in block 1, when subjects expected US, and in block 3, when subjects had learned a CS-CR association. However, our results are also in line with an alternative interpretation, according to which PAV is exclusively a model-free process that, conversely to HAB, establishes CS-CR associations very quickly. Thus, just few trials could be sufficient for CS to trigger CR directly (without model-based CS-US associations). Further investigations are necessary to disentangle these two alternative hypotheses, for example by varying the amount of trials in which the shock is delivered before the extinction phase.

A final issue concerns the interaction between GDC and PAV on the one hand, and between HAB and PAV on the other hand. In relation to this, Dayan et al. (2006) hypothesized that, as GDC is based on a tree-search (for the prediction of future prospects), it can anticipate the consequences of

Pavlovian responses, which can therefore be easier to eliminate. Conversely, PAV and HAB should more plausibly compete, because they are at the same level of processing. Consequently, PAV might interact mostly with HAB, rather than with GDC.

In our experiment, to distinguish GDC vs. PAV from HAB vs. PAV interactions, we designed two blocks. We hypothesized that, in block 1 the absence of previous experience leads to recruitment of GDC, while in block 2 the task can be routinizing, possibly leading to recruitment of HAB. Since we did not find the main effect of block (plausibly due to weak statistical power), we were not allowed to study the HAB vs. PAV interaction and thus we delimit our discussion to the GDC vs. PAV one. Our results indicate that GDC, at least in this circumstance, cannot “compensate” for the effects of PAV, although it is possible that with more trials a more accurate predictive model of the PAV influence could allow so.

In sum, the present study suggests that, at least in aversive domains, model-free and possibly also model-based Pavlovian processes interfere with GDC directly at the motor level. Further investigations are necessary to clarify crucial aspects of this phenomenon, such as in what conditions GDC prevail on PAV, and whether the interaction is due to specific motor responses (e.g., raising the finger), as we hypothesized, or to more broad motor processes, such as shivering, or even to processes that lie outside motor performance, such as anxiety and attention (Mobbs et al., 2007). Furthermore, although our results point to a double nature of PAV (both model-free and model-based mechanisms seem to be involved, at least during the interaction with GDC), however they cannot rule out the possibility that PAV is a “fast-learning” model-free system. Moreover, the neural bases of such interactions are largely unknown. Our ongoing work regards the investigation of hypothesis that the competition between the goal-oriented action and CR, and the surprise of GDC (generated by the unexpected motor activation due to the Pavlovian influence), may be related, respectively, to the cingulate cortex, involved in conflict monitoring (Botvinick et al., 2001), and the parietal cortex, involved in the sense of agency (Blakemore and Frith, 2003). Additional important issues for future research are studying the mechanisms through which cognitive agents predict increasingly distal consequences of actions (Pezzulo, 2011; Pezzulo and Castelfranchi, 2007), learn to inhibit prepotent responses and exert cognitive control over extended periods of time (Pezzulo and Castelfranchi, 2009).

Acknowledgments Research funded by the EU’s FP7 under grant agreement no FP7-ICT-270108 (Goal-Leaders).

References

Adams, C. and Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, 33(2):109–121.

Blakemore, S. and Frith, C. (2003). Self-awareness and action. *Current Opinion in Neurobiology*, 13(2):219–224.

Botvinick, M., Braver, T., Barch, D., Carter, C., and Cohen, J. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3):624.

Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12):1704–1711.

Dayan, P., Niv, Y., Seymour, B., and Daw, N. D. (2006). The misbehavior of value and the discipline of the will. *Neural Netw*, 19(8):1153–1160.

Dayan, P. and Seymour, B. (2008). Values and actions in aversion. In *Neuroeconomics: decision making and the brain*, page 175. Academic Pr.

Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philos Trans R Soc London*, B308:67–78.

Fendt, M. and Fanselow, M. (1999). The neuroanatomical and neurochemical basis of conditioned fear. *Neuroscience & Biobehavioral Reviews*, 23(5):743–760.

Hershberger, W. (1986). An approach through the looking-glass. *Learning & behavior*, 14(4):443–451.

Holland, P. (2004). Relations between Pavlovian-instrumental transfer and reinforcer devaluation. *J Exp Psychol*, 30(2):104–117.

Konorski, J. (1967). *Integrative activity of the brain*. Chicago.

Lovibond, P. and Shanks, D. (2002). The role of awareness in pavlovian conditioning: Empirical evidence and theoretical implications. *Animal Behavior*, 28(1):3–26.

Mobbs, D., Petrovic, P., Marchant, J., Hassabis, D., Weiskopf, N., Seymour, B., Dolan, R., and Frith, C. (2007). When fear is near: threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science*, 317(5841):1079.

Morse, W., Mead, R., and Kelleher, R. (1967). Modulation of elicited behavior by a fixed-interval schedule of electric shock presentation. *Science*, 157(3785):215.

Ohman, A. and Soares, J. (1998). Emotional Conditioning to Masked Stimuli: Expectancies for Aversive Outcomes Following Nonrecognized Fear-Relevant Stimuli. *Journal of Experimental Psychology: General*, 127(1):69–82.

Pezzulo, G. (2011). Grounding procedural and declarative knowledge in sensorimotor anticipation. *Mind and Language*, 26(1):78–114.

Pezzulo, G. and Castelfranchi, C. (2007). The symbol detachment problem. *Cognitive Processing*, 8(2):115–131.

Pezzulo, G. and Castelfranchi, C. (2009). Thinking as the control of imagination: a conceptual framework for goal-directed systems. *Psychological Research*, 73(4):559–577.

Seymour, B. and Dolan, R. (2008). Emotion, decision making, and the amygdala. *Neuron*, 58(5):662–671.

Williams, D. and Williams, H. (1969). Auto-maintenance in the pigeon: sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior*, 12(4):511.