Citation: Jennings, D., Alonso, E., Mondragon, E. & Bonardi, C. (2010). Temporal uncertainty during overshadowing: A temporal difference account. In: Computational Neuroscience for Advancing Artificial Intelligence: Models, Methods and Applications. (pp. 46-55). IGI Global. ISBN 9781609600211 doi: 10.4018/978-1-60960-021-1.ch003

Computational Neuroscience for Advancing Artificial Intelligence:

Models, Methods and Applications

Eduardo Alonso
City University London, UK

Esther Mondragón
Centre for Computational and Animal Learning Research, UK

MEDICAL INFORMATION SCIENCE REFERENCE
Hershey · New York

Director of Editorial Content:	Kristin Klinger
Director of Book Publications:	Julia Mosemann
Acquisitions Editor:	Lindsay Johnston
Development Editor:	Joel Gamon
Publishing Assistant:	Jamie Snavely
Typesetter:	Michael Brehm
Production Editor:	Jamie Snavely
Cover Design:	Lisa Tosheff

Published in the United States of America by
Medical Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Computational neuroscience for advancing artificial intelligence : models,
methods and applications / Eduardo Alonso and Esther Mondragon, editors.
p. cm.
Summary: "This book argues that computational models in behavioral
neuroscience must be taken with caution, and advocates for the study of
mathematical models of existing theories as complementary to neuro-
psychological models and computational models"-- Provided by publisher.
Includes bibliographical references and index.
ISBN 978-1-60960-021-1 (hardcover) -- ISBN 978-1-60960-023-5 (ebook) 1.
Computational neuroscience. 2. Neurosciences--Mathematical models. 3.
Artificial intelligence. I. Alonso, Eduardo, 1967- II. Mondragon, Esther,
1965-
QP357.5.C634 2011
612.80285'63--dc22
2010018588

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Chapter 3

Temporal Uncertainty during Overshadowing:

A Temporal Difference Account

Dómhnall J. Jennings
Newcastle University, UK

Eduardo Alonso
City University London, UK

Esther Mondragón
Centre for Computational and Animal Learning Research, UK

Charlotte Bonardi
University of Nottingham, UK

ABSTRACT

Standard associative learning theories typically fail to conceptualise the temporal properties of a stimulus, and hence cannot easily make predictions about the effects such properties might have on the magnitude of conditioning phenomena. Despite this, in intuitive terms we might expect that the temporal properties of a stimulus that is paired with some outcome to be important. In particular, there is no previous research addressing the way that fixed or variable duration stimuli can affect overshadowing. In this chapter we report results which show that the degree of overshadowing depends on the distribution form - fixed or variable - of the overshadowing stimulus, and argue that conditioning is weaker under conditions of temporal uncertainty. These results are discussed in terms of models of conditioning and timing. We conclude that the temporal difference model, which has been extensively applied to the reinforcement learning problem in machine learning, accounts for the key findings of our study.

INTRODUCTION

Overshadowing is a procedure in which a target conditioned stimulus (CS1) is conditioned together with a second stimulus (CS2), usually in a simultaneous compound which co-terminates with the delivery of an unconditioned stimulus (US). This training results in attenuated (or *overshadowed*) responding to CS1 when it is subsequently tested alone, compared to the case in which CS1 is con-

DOI: 10.4018/978-1-60960-021-1.ch003

ditioned in isolation. A key variable in determining the degree of overshadowing is the relative salience of the two stimuli: the more salient the overshadowing stimulus relative to the target, the more overshadowing is observed (Mackintosh, 1976). This observation may be interpreted theoretically in terms of learning models such as that proposed by Rescorla and Wagner (1972): the limited amount of associative strength that may be supported by the US must be distributed between the two competing CSs; as speed of acquisition depends on stimulus salience, this distribution is affected by the relative salience of the two stimuli.

However, pairing a CS of fixed duration with a US is a procedure that informs the animal not only *that* the US be delivered, but also *when* it is to be delivered, and there is good evidence to suggest that animals are indeed able to time the delivery of a US (e.g. Kirkpatrick & Church, 2000). This observation makes little contact with standard associative models of conditioning (e.g. Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972) which, despite accounting for the magnitude of conditioned responding to a CS, generally fail to provide a comprehensive account of how animals learn about the temporal properties of the CS. More detailed discussion of theories of timing will be postponed until the general discussion; at this point it is sufficient to note that standard associative theories typically fail to conceptualise the temporal properties of a stimulus, and hence cannot easily make predictions about the effects such properties might have on the magnitude of overshadowing. Despite this, in intuitive terms one might expect these temporal aspects to be important. As noted above, the most important determinant of the degree of overshadowing is usually thought to be the ease with which the overshadowing stimulus can acquire associative strength. One can, however, make a parallel argument about a stimulus' temporal properties; given that better predictors of the US acquire associative strength more easily than unreliable predictors, then one might expect a fixed stimulus to be a better overshadowing stimulus than a variable one. This is because, in some sense, a stimulus of fixed duration is a more reliable predictor of US outcome than one of variable duration, as the fixed stimulus gives precise information about the time of US delivery while the variable one does not.

There is no previous research addressing the way in which fixed and variable CSs can produce overshadowing; however, Kohler and Ayres (1979) examined the ability of fixed and variable duration CSs to produce blocking. They conditioned animals to a tone-light compound; for two groups of animals both tone and light were fixed, and one group was pretrained with a fixed light. For another two groups both tone and light were of variable duration and one group was pretrained with a variable light. They found equal blocking in both groups, suggesting that a fixed stimulus was no more able to produce blocking than a variable one. However, this conclusion is complicated by the fact that the treatment of the pretrained stimulus was confounded with the distribution of the to-be-blocked stimulus: the fixed stimulus blocked a fixed target and the variable stimulus a variable target. Moreover, subsequent work has suggested that a fixed CS may after all be a more effective blocking stimulus than a variable (D. Jennings and K. Kirkpatrick, unpublished data). It should be noted, however, that neither of these studies included any groups trained to the target alone. As a consequence, it is impossible to assess what effect the temporal distribution of the stimuli had on their ability to produce an overshadowing effect. This was, therefore, the purpose of the present experiment.

The experiment employed six groups of Lister hooded rats (Harlan UK, Table 1) maintained on a restricted diet at 80% of their free feeding weight. Since we used a houselight as one of the stimuli the boxes were normally not illuminated. All six groups were trained with a light CS that produced diffuse illumination of the conditioning chamber of about 200 lux when operational; three groups (Fc, VF and FF) experienced the light for a fixed

Table 1. The design of the overshadowing experiment (N= 8 per group).

	Treatment	
Group	Training	Test
Absent-Fixed (Fc)	CS2→US	CS2
Fixed-Fixed (FF)	CS1CS2→US	CS2
Variable-Fixed (VF)	CS1CS2→US	CS2
Absent-Variable (Vc)	CS2→US	CS2
Fixed-Variable (FV)	CS1CS2→US	CS2
Variable-Variable(VV)	CS1CS2→US	CS2

CS2 = Light; CS1 = 70db white noise; F = Fixed 30s; V = Variable 30s

duration and three for a variable duration (Groups Vc, FV and VV). For one of each of these triplets of groups the light was conditioned alone (Groups F and V), for another it was conditioned with a fixed noise (Groups FF and FV) and for the third it was conditioned with a variable noise (Groups VF and VV). In this way we could compare the ability of fixed and variable stimuli to produce overshadowing, independent of whether the overshadowed stimulus was itself fixed or variable.

The basic experimental approach is described in Table 1; training was conducted over six sessions and consisted of 26 trials per session, and each trial comprised presentations of the houselight followed by the delivery of a single food pellet. Each head entry into the food cup was recorded by interruption of a photobeam; therefore, the rate and distribution of responding during the CS constituted the measure of learning (for further details see e.g. Jennings, Bonardi and Kirkpatrick, 2007). In the control groups the light was presented alone, whereas in the experimental groups it was accompanied by CS1, the 70-db white noise. Critical to the present experiment was the distribution form of the stimuli; two different distributions were employed - either a *fixed* stimulus that was 30 s in duration, or a *variable* stimulus, whose duration was drawn from an exponential distribution with a mean and standard deviation of 30 s (so that the average level of exposure to the fixed and variable duration stimuli was the same; Evans, Hastings & Peacock, 1993). The inter-trial interval (ITI), measured from food delivery to onset of the next CS, consisted of a fixed 60 s, plus an additional variable duration period with a mean of 60 s. Following training we conducted ten test sessions during which three reinforced probe trials of the light (all groups) and noise (the four experimental groups only) stimuli alone were presented to assess any overshadowing effect.

Compound training with the six different groups outlined in Table 1 was conducted without incident over six consecutive daily sessions. An initial inspection of responding during training indicated that there was no difference between the groups. We also conducted an examination of the timing functions; groups that were conditioned to a compound containing a fixed duration stimulus displayed an increase in response rate from stimulus onset to the expected time of food delivery, consistent with them having learned the time to reinforcement. Conversely, groups that received only variable duration stimuli during conditioning responded at a steady rate over the course of the CS.

Following training we investigated overshadowing by examining responding during test trials of the light stimulus (CS2) presented alone. Data recorded during this phase are presented for all six groups in Figure 1. It is evident that in general there was more responding to the light when the stimulus was of a fixed duration than when it was variable. Moreover, responding was lower in the overshadowing groups than in the control groups, and this overshadowing effect was larger when the overshadowing stimulus was fixed than when it was variable. Essentially, groups trained with a fixed overshadowing stimulus (FF and FV) responded at a significantly lower rate than the control groups (F and V), whereas groups trained with a variable duration overshadowing stimulus (VF and VV) did not differ from either the control or the other overshadowing groups. Although there

Figure 1. Elevation scores displaying the mean level of responding (± SE) to the light (CS2) dependent on whether the overshadowing stimulus was absent (Groups F and V), variable (Groups VF and VV) or fixed (Groups FF and FV). Groups trained with a fixed light are shown in the left half of the Figure, and those trained with a variable light are shown on the right.

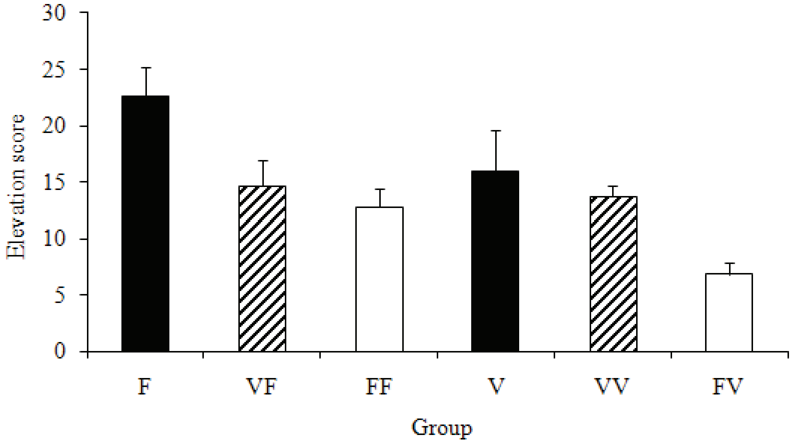

was some indication of higher response rates when the noise was fixed than when it was variable, this was not statistically reliable.

Groups that received the fixed duration probe trials displayed an increase in their rate of responding from stimulus onset to the time when food was delivered at stimulus offset (Figure 2, top panel), whereas those experiencing variable duration probes tended to respond at a steady rate throughout CS presentation (Figure 2, bottom panel). These data are by no means unusual in that it has been shown that fixed and variable distributions tend to generate this form of responding; yet it should be remembered that, at least for the compound groups with mixed distributions (FV and VF), there were two sources of competing temporal information. Our analysis indicates that the rats in these two groups were able to track these different distribution forms despite this competition. We will return to this issue shortly following a theoretical discussion of the data set as a whole.

A brief recap of the main finding of this experiment shows that the degree of overshadowing that was observed depended on the distribution form of the overshadowing stimulus: overshadowing was more profound when the overshadowing stimulus was fixed than when it was variable. At face value the implication of this finding is that a fixed duration stimulus acquires associative strength more effectively than one of variable duration. This suggestion is consistent with our additional observation of greater conditioned responding to the target CS when it was of fixed duration than when it was variable.

It is not immediately clear how these results are to be explained. Standard associative models do not incorporate any mechanism for explaining timing effects, and so it is far from obvious how they could accommodate these data. It should be noted, however, that one interpretation of our results is that a fixed stimulus acquires associative strength more than a variable one because it is a better predictor of the US. Mackintosh (1975) proposed a model of conditioning according to which the associability of a stimulus is directly related to its predictive power, such that better predictors have higher associability. One possibility, therefore, is to argue that the variable stimulus varies from trial to trial, whereas the fixed does

Figure 2. The probability of responding as a function of time since CS onset during light probe trials, for groups that received a fixed target stimulus (top panel) and a variable target stimulus (bottom panel).

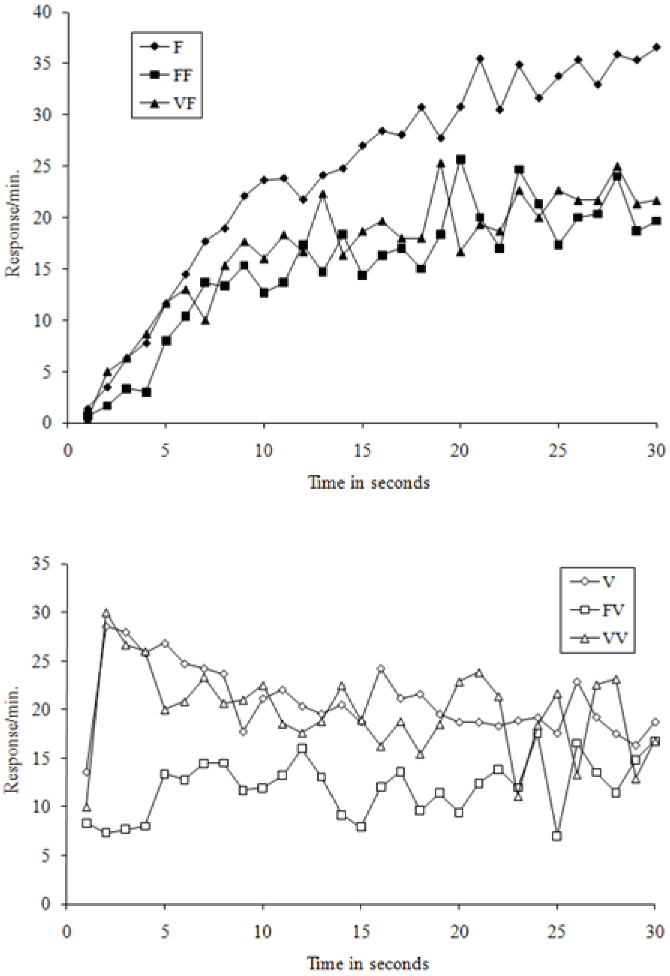

not; the model then predicts that the associability of the variable stimulus should fall below that of the fixed simply because the fixed, being the same from one trial to the next, will acquire more associative strength. His model could thus predict that a fixed stimulus would have greater associability than a variable stimulus, and that this is the source of the results that we observed. Although logical, however, this account fails to capture any of the temporal characteristics of responding and so on these grounds must necessarily be regarded as incomplete.

An alternative approach is to look to hybrid models that have been developed to account for both conditioning and timing, such as rate expectancy theory (RET: Gallistel & Gibbon, 2000, 2002; cf Gibbon, 1977). RET proposes that conditioning and timing are dependent on a timing-like mechanism, with timing occurring after conditioned responding has been established – a serial process account. Conditioning requires a comparison of the rate of reinforcement during the CS with that during the background: when the CS is perceived as elevating the rate of reinforcement over that occurring in the CSs absence, then conditioned responding to the CS emerges. However, RET predicts that responding during the fixed and variable duration CSs of our experiment should be identical, as their average duration, and hence the average reinforcement rate during fixed and variable CS presentation, is identical - and yet in our experiment this was clearly not the case. For the same reason it would predict no difference in the ability of fixed and variable duration stimuli to produce overshadowing, and yet we observed that the fixed CS was the better overshadowing stimulus. In fact RET cannot easily explain overshadowing *per se*; the overshadowing process lies outside the scope of the model, and requires the introduction of additional principles, such as stimulus salience. In short, then, RET cannot provide a satisfactory explanation of our results.

An alternative model of conditioning and timing is the temporal coding hypothesis (Savastano & Miller, 1998). This is a form of single process model which asserts that conditioning and timing arise simultaneously, as a consequence of pairing a CS and a US together. Specifically, when an association is formed, a temporal map detailing the predictive relationship between the CS and US is established, and the match or mismatch of the temporal maps associated with any two stimuli will affect the way in which they interact associatively (Barnet, Grahame & Miller, 1993; Blaisdell, Denniston & Miller, 1998). In the present experiment this account would predict that overshadowing should be most profound between two stimuli that share the same temporal map - specifically in Groups FF and VV (Blaisdell, et al., 1998). However, this was not what was observed; rather, overshadowing was greater in Groups FF and FV than in Groups VV and VF. This suggests that matching temporal maps is not an important determinant of the overshadowing effect, and that our results that do not, therefore, support the temporal coding hypothesis.

One further model of timing and conditioning is the temporal difference (TD) model (Sutton & Barto 1987, 1990), a single process model that emphasises an integrated approach to conditioning and timing. It employs a conditioning mechanism that is a modification of the Rescorla-Wagner (Rescorla & Wagner, 1972) model, but which incorporates a temporal component enabling it to accommodate the fact that delayed reinforcers are less effective than immediate ones. This may be adapted to deal with CSs of varying durations by assuming that a stimulus of given duration is comprised of a series of time steps that condition independently of each other, and that is reinforced by a punctate US (cf. Moore & Choi, 1997); the model then describes how later portions of the CS will condition more effectively than earlier ones. This permits an explanation of how responding can gradually increase over the course of the CS (Moore & Choi, 1997; Sutton & Barto, 1990), and has successfully predicted conditioned performance in serial-compound experiments (e.g. Egger & Miller, 1962) and blocking studies (e.g. Gaioni, 1982; Jennings & Kirkpatrick, 2006). This model can account for our results, because it predicts that a variable stimulus will acquire less associative strength than a fixed, even though the mean duration of the two stimuli is the same. The reason for this lies in the way in which reinforcement accrues to CS time steps that are removed from the US. The model assumes that the time step closest to the US acquires associative strength in the normal way, but that delayed time steps acquire strength *not* through direct association with the US, but recursively via second-order conditioning to the unit closer in time to the US, and the amount of associative strength each delayed time step receives is determined by a discounting parameter *gamma* (γ), which has a value between 0 and 1. For example, let us assume that the final time step acquires an associative strength of 1 unit. The first delayed time step will acquire this strength, discounted by γ, and thus acquire 0.9 units of strength. The second delayed time step will acquire the first step's strength (γ), also discounted by γ, meaning it acquires γ^2 (0.81

units) the third delayed step acquires $\gamma^{3,}$ (.73 units) and so on. Effectively this means that units more contiguous with the US acquire substantially more associative strength than temporally distant units (Sutton & Barto, 1990). This is important because although the variable CS will comprise the same total number of time steps as the fixed, on some trials the variable CS will be either shorter or longer than the fixed stimulus. Consequently, when the CS is shorter, fewer units will be available to receive reinforcement and where longer, some time units will be far removed from the US and will, therefore, receive negligible associative strength.

For example, consider a fixed duration CS of two time steps, and a variable CS drawn from a uniform distribution, that can range between 1 and 3 steps. Over three trials a total of six steps will be reinforced: the fixed CS is two steps long on each trial, whereas the variable (let us assume) is one step on one trial, two on another, and three on a third. Let us again say the final unit on each trial acquires an associative strength of 1. Setting gamma to 0.9, associative strength acquired on the first trial with each delayed time step is effectively directly proportional to gamma. Thus for the fixed CS the first step (defined as that closest to the US) will accrue 1 unit of strength on each trial (3 after three trials), and the second step 0.9 on each trial (2.70 after three trials), yielding a total of 5.70 units after three trials. The first step of the variable stimulus will also acquire 1 unit per trial (3 after three trials). However, the second step is only present on *two* of the three trials, and so will acquire a total of 0.9 x 2 = 1.8 over three trials. On the trial in which the variable stimulus is longer than the fixed, there will be a third time step, which will acquire $(0.9)^2 = 0.81$ units. This yields a total of 5.61 units, which is less than the 5.70 units acquired by the fixed CS. In short, this difference arises because the fixed CS differs from the variable CS only in that a time step from position 2 has been moved to position 3, and time steps further from the US acquire less associative strength. The TD model can, therefore, account for the key findings of the present study.

A further aspect of our data concerns the form of responding (timing) that we observed during the various stimuli. To date, most models of conditioning and timing have tended to focus on conditioned responding during fixed duration stimuli, and as a consequence are silent about performance during variable duration CSs (although see the TD model referred to above). An analysis of the timing functions in the present study indicated that rats responded quite differently to the target stimuli depending on whether they were of fixed or of variable duration. When the target was a fixed 30-s there was a gradually increasing rate of responding which reached a peak at about the time of US delivery (Figure 2a) and a slope that differed significantly from zero; when it was of variable duration, on the other hand, responding was relatively stable and the slope did not differ from zero (Figure 2b). Although the different patterns of responding during stimuli of fixed and variable duration have been demonstrated before (Kirkpatrick & Church, 2003), one further marked feature of the present results was that this pattern was unaffected by whether the target stimulus had suffered overshadowing or not. Whether or not this is a general feature of response form during cue competition will require further investigation; however, these results are consistent with previous reports that demonstrate that stimuli of differing fixed durations were tracked accurately by rat subjects during both blocking and overshadowing (Jennings & Kirkpatrick 2006; Jennings et al., 2007).

Hybrid accounts of timing and conditioning usually implicitly assume that attenuation of responding during cue competition tasks will also lead to disruption of timing. For instance, the temporal coding hypothesis argues that conditioning and timing emerge directly as part of the learning process, whereas according to RET timing may only emerge after conditioning has occurred; according to both these accounts therefore, disruption of conditioning must necessarily be accompanied by a disruption of timing. One account that might not predict such a deficit is the TD model; according to this account cue competition and, therefore, overshadowing occurs between the units that comprise the stimulus, according to Rescorla Wagner principles. Timing occurs because of the asymmetric distribution of associative strength over the course of the CS, with units close in time to the US acquiring the most strength. Even if the total associative strength acquired by the CS were reduced, as long as the proportion acquired by the later time steps relative to the earlier ones were preserved, then one would expect similar timing functions regardless of whether overshadowing was obtained or not.

Both conditioning and timing are known to occur in many basic conditioning procedures, and are often studied using highly similar experimental procedures; it is the dependent measures of responding that differ (Kirkpatrick & Church 1998). Yet surprisingly little attention has been paid to understanding the relationship between measures of the magnitude of responding (conditioning) and the time of responding (timing), and the nature of the relationship between these two forms of learning. What attention there has been has tended to come from students of timing who have sought to model the conditioned response within a variety information processing accounts of learning (e.g. Gallistel & Gibbon 2000). While this approach shows imagination and a willingness to address the thorny issue of how these two core psychological processes interact, such an approach also requires that a great many advances in associative theorising should necessarily be abandoned. The alternative (and arguably more parsimonious approach) is to incorporate timing within a model that is at least, in part, associatively based; based on the experimental results presented above, the TD appears to make a step in this direction.

In summary, we have provided further evidence that the temporal properties of the stimulus can affect the outcome of classical conditioning procedures. Standard associative theories cannot easily accommodate findings of this type, as by and large they fail to conceptualise the temporal properties of the stimulus. Hybrid models, which attempt to accommodate both timing and conditioning within a single framework, fare no better, in part because they have no principled account of associative phenomena such as the overshadowing effect. The account that provides the best explanation of our data is probably the temporal difference (TD) model (e.g. Sutton & Barto, 1987; 1990 cf., Vogel, Brandon & Wagner, 2000) - a model that uses incorporates a modified Rescorla-Wagner conditioning rule within a framework that has a principled conceptualisation of the temporal properties of the stimulus. This type of approach is probably that best suited to account for the phenomena of conditioning and timing within a single theoretical framework.

REFERENCES

Barnet, R. C., Grahame, N. J., & Miller, R. R. (1993). Temporal encoding as a determinant of blocking. *Journal of Experimental Psychology. Animal Behavior Processes*, *19*, 327–341. doi:10.1037/0097-7403.19.4.327

Blaisdell, A. P., Denniston, J. C., & Miller, R. R. (1998). Temporal encoding as a determinant of overshadowing. *Journal of Experimental Psychology. Animal Behavior Processes*, *24*, 72–83. doi:10.1037/0097-7403.24.1.72

Egger, M. D., & Miller, N. E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology*, *64*, 97–104.doi:10.1037/h0040364

Evans, M., Hastings, N., & Peacock, B. (1993). *Statistical Distributions*. New York: Wiley.

Gaioni, S. J. (1982). Blocking and nonsimultaneous compounds: Comparison of responding during compound conditioning and testing. *The Pavlovian Journal of Biological Science*, (January-March): 16–29.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review*, *107*, 289–344.doi:10.1037/0033-295X.107.2.289

Gallistel, C. R., & Gibbon, J. (2002). *The symbolic foundations of conditioned behavior*. Mahwah, NJ: Erlbaum Associates.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*, 279–325.doi:10.1037/0033-295X.84.3.279

Jennings, D. J., Bonardi, C., & Kirkpatrick, K. (2007). Overshadowing and stimulus duration. *Journal of Experimental Psychology. Animal Behavior Processes*, *33*, 464–475.doi:10.1037/0097-7403.33.4.464

Jennings, D. J., & Kirkpatrick, K. (2006). Interval duration effects on blocking in appetitive conditioning. *Behavioural Processes*, *71*, 318–329. doi:10.1016/j.beproc.2005.11.007

Kirkpatrick, K., & Church, R. M. (1998). Are separate theories of conditioning and timing necessary? *Behavioural Processes*, *44*, 163–182. doi:10.1016/S0376-6357(98)00047-3

Kirkpatrick, K., & Church, R. M. (2000). Stimulus and temporal cues in classical conditioning. *Journal of Experimental Psychology. Animal Behavior Processes*, *26*, 206–219.doi:10.1037/0097-7403.26.2.206

Kirkpatrick, K., & Church, R. M. (2003). Tracking of the expected time to reinforcement in temporal conditioning procedures. *Learning & Behavior*, *31*, 3–21.

Kohler, E. A., & Ayres, J. J. B. (1982). Blocking with serial and simultaneous compounds in a trace conditioning procedure. *Animal Learning & Behavior*, *10*, 277–287.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298. doi:10.1037/h0076778

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning & Behavior*, *4*, 186–192.

Moore, J. W., & Choi, J.-S. (1997). The TD model of classical conditioning: Response topography and brain implementation. In Donahoe, J. W., & Dorsel, V. P. (Eds.), *Neural-networks models of cognition* (pp. 387–405). New York: Elsevier Science.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, *87*, 532–552. doi:10.1037/0033-295X.87.6.532

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement. In Black, A. H., & Prokasy, W. F. (Eds.), *Classical conditioning: II. Theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.

Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioural Processes*, *44*, 147–162.doi:10.1016/S0376-6357(98)00046-1

Sutton, R. S., & Barto, A. G. (1987). A temporal difference model of classical conditioning. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society* (pp. 355-378). Erlbaum.

Sutton, R. S., & Barto, A. G. (1990). Time derivative models of Pavlovian reinforcement. In Gabriel, M. R., & Moore, J. W. (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: MIT Press.

Vogel, E. H., Brandon, S. E., & Wagner, A. R. (2000). Stimulus representation in SOP II: An application to inhibition of delay. *Behavioural Processes*, *110*, 67–72.

KEY TERMS AND DEFINITIONS

Conditioning: The mechanism of learning about contingencies between environmental events.

Timing: The ability to accurately anticipate the time at which a predicted outcome will be delivered.

Overshadowing: Interference in learning about the relationship between an event and an outcome due to the presence of a second stimulus that predicts the same outcome.

Temporal Variability: The distribution of a stimulus' temporal properties indicating whether it is fixed or variable.

Temporal Difference Learning: An approach to learning how to predict a quantity that depends on future values of a given signal. The name TD derives from its use of changes, or differences, in predictions over successive time steps to drive the learning process. The prediction at any given time step is updated to bring it closer to the prediction of the same quantity at the next time step.