A

Avoidance Learning

Tiago V. Maia
Department of Psychiatry, Columbia University and New York State Psychiatric Institute, New York, NY, USA

Synonyms

Learning to avoid aversive outcomes

Definition

Learning to avoid aversive outcomes is crucial for survival, and it is ubiquitous in everyday life. This entry focuses on ► active avoidance, in which performing a certain behavior prevents an aversive outcome that would otherwise occur. ► Avoidance learning in that context consists of learning to perform the appropriate behaviors in the appropriate circumstances to prevent aversive outcomes. Active avoidance should be distinguished from ► passive avoidance, often called *punishment*, in which performing a behavior results in an aversive outcome, and learning therefore consists of suppressing, rather than strengthening, the behavior. Punishment is discussed in the "Punishment and Reward" entry. Henceforward, whenever the term "avoidance" is used, it is intended to mean active avoidance.

Theoretical Background

The earliest avoidance-learning experiments were conducted at the beginning of the twentieth century by Bekhterev, using what would become one of the most common paradigms in research on avoidance: the signaled avoidance procedure. In this procedure, a trial starts with the presentation of a neutral warning stimulus, often referred to as the conditioned stimulus (CS). Unless the subject performs a pre-specified response, an aversive unconditioned stimulus (US), such as a shock, occurs at a given time following the onset of the CS. If, however, the subject performs the pre-specified response – the avoidance response – after the onset of the CS but before the US occurs, the US is omitted. Subjects – both humans and nonhuman animals – typically learn to perform the avoidance response, thereby avoiding the US.

Most experiments in avoidance learning have been conducted with rats, but a wide variety of other species (e.g., monkeys, dogs, guinea pigs, pigeons, chicks, and goldfish) have also been shown to learn to avoid. Although there are some differences between species in avoidance learning – specifically, animals learn the avoidance response more easily when it is closely related to one of their species-specific defense reactions – many findings in avoidance learning have been shown to generalize across species, and theories of avoidance learning are generally assumed to apply to a wide variety of species (including humans). Henceforward, the term "subject" will be used instead of "animal," to highlight the fact that the findings and ideas discussed are generally assumed to apply to both humans and nonhuman animals.

One of the earliest explanations for avoidance learning was based on Pavlov's ► stimulus-substitution theory. In Pavlov's experiments, a neutral stimulus such as the ringing of a bell (the CS) was paired with food (the US). After a few such pairings, the CS would elicit the same response (salivation) as the US. Pavlov proposed that the CS becomes a substitute for – that is, starts eliciting the same responses as – the US. This idea also provided an elegant explanation for avoidance learning, because in the early experiments on avoidance, the avoidance response was typically a fleeing behavior that was spontaneously elicited by the shock. Avoidance learning was therefore assumed to be due to the pairings of the CS with the US early in avoidance training (before the avoidance response is learned); such pairings would make the CS elicit the same fleeing response that the US elicited. Experiments in the late 1930s, however, contradicted this Pavlovian account, suggesting that avoidance learning depends instead on instrumental learning. Explaining avoidance learning in instrumental terms raised a theoretical conundrum, though. The main account of instrumental learning was Thorndike's ► law of effect. Thorndike proposed that instrumental learning consists of learning stimulus-response (or situation-response; S-R) associations. The association between situation S and response R is strengthened when the subject is in situation S, performs response R, and that response is "accompanied or closely followed by satisfaction" (in Thorndike's own words). As Mowrer (1947) noted, however, applying the law of effect to avoidance learning raised a difficult question: "how can a shock which is *not experienced*, i.e., which is avoided, be said to provide [. . .] a source of [. . .] satisfaction" (p. 108)?

Mowrer's (1947) two-factor theory – which remains one of the main theories of avoidance learning – provided an answer to this puzzle. Mowrer was strongly influenced by Hull's ► drive-reduction theory. Hull suggested that responses are learned via drive reduction, where *drives* are states such as hunger, the reduction of which acts as a reward. Mowrer suggested that fear is an acquired drive, and that reductions in fear are therefore reinforcing. He proposed that avoidance learning involves two processes (hence the name "two-factor theory"). Early in training, before the subject consistently avoids, the CS is often followed by the US, and this produces fear of the CS via Pavlovian conditioning. Subsequently, a response that terminates the CS reduces fear, and that fear reduction reinforces the response. For Mowrer, the avoidance response is therefore reinforced by CS termination, not by the actual shock avoidance. He suggested that "the *avoidance* of the shock is a [. . .] by-product."

In experiments conducted before the development of two-factor theory, the avoidance response both terminated the CS and avoided the US, making it impossible to determine the relative contributions of these two events to learning. Following the development of two-factor theory, however, several experiments tried to disentangle the differential effects of these events. One of the most influential such experiments was the ► acquired-drive experiment of Brown and Jacobs (1949). The experiment consisted of two phases. In the first phase, rats underwent Pavlovian conditioning of a CS to a shock. In the second phase, the CS was presented and the rats could terminate it by performing a response. The US was never presented during this phase, regardless of the rats' actions, so the response did not play any role in avoiding the US. The rats learned to perform the response, suggesting that, as predicted by two-factor theory, CS termination was sufficient to support avoidance learning.

Although this finding was widely replicated, subsequent experiments demonstrated that, contrary to the predictions of two-factor theory, US avoidance also plays a role in reinforcing the avoidance response. These experiments typically used a 2 × 2 factorial design, in which the two factors were CS termination (yes or no) and US avoidance (yes or no). Generally, both factors were found to be effective in reinforcing the response, with the highest level of avoidance learning occurring when the response produced both effects.

Another prediction of two-factor theory was also found to be at odds with experimental findings. According to the theory, the avoidance response is reinforced by reduction of the fear associated with the CS; fear of the CS and the strength of the avoidance response should therefore correlate. Several studies, however, found that long after fear of the CS had virtually extinguished (due to consecutive avoidance responses that ensured that the CS was not followed by the US), the avoidance response still persisted.

These difficulties with two-factor theory prompted the development of alternative theories, the most prominent of which is the cognitive theory of Seligman and Johnston (1973). The main tenet of Seligman and Johnston's theory is that the avoidance response is driven not by S-R associations, but by expectancies about response-outcome (R-O) contingencies. The theory suggests that during avoidance learning subjects develop two R-O expectancies: (1) if they perform the avoidance response, no shock will occur; and (2) if they do not perform the avoidance response, shock will occur. Subjects perform the avoidance response because they prefer no shock to shock.

This theory naturally explains the effectiveness of US avoidance in supporting learning. It also explains the persistence of avoidance after fear is extinguished, because avoidance is assumed to be driven by R-O expectancies rather than being reinforced by fear reduction. The theory may even explain learning in acquired-drive experiments: In the first phase of such experiments, subjects may develop the expectancy that if they do not perform the avoidance response (which is typically blocked in this phase), shock will occur. In the second phase, they may develop the expectancy that if they perform the avoidance response, no shock will occur, while failing to disconfirm the expectancy that shock will occur if they do not perform the response. The theory, however, has difficulty explaining avoidance learning when the response terminates the CS but is followed by the US.

Important Scientific Research and Open Questions

The past decade has witnessed a revolution in our understanding of conditioning, brought about by the use of computational models from the field of ► reinforcement learning to explain a myriad of behavioral and neural findings in conditioning (Maia 2009).

Recent work has shown that one such model – the so-called ► actor-critic – explains a wide variety of findings in avoidance learning (Maia 2010). Remarkably, the model is closely related to two-factor theory (Maia 2010). The model consists of two components: the *critic* and the *actor*. The critic implements Pavlovian conditioning by learning the *values* of stimuli or situations (i.e., the future reinforcements predicted by those stimuli or situations). In fear conditioning, the critic learns a negative value for the CS, because the CS predicts an aversive US. The actor implements S-R learning. Unlike in the law of effect, however, the change in strength of the S-R association is *not* determined simply by whether the response is followed by a positive or negative outcome. Instead, it is determined by whether the response is followed by an outcome that is *better* or *worse* than expected. Specifically, a ► prediction error is calculated by subtracting the value that was expected from the actual outcome. If the prediction error is positive, the outcome was better than expected, and the S-R association is strengthened. If the prediction error is negative, the outcome was worse than expected, and the S-R association is weakened.

The actor-critic's explanation for many of the findings in avoidance learning is similar to the explanation of two-factor theory. Consider, for example, the learning of the avoidance response. Early in training, when the CS is often followed by the US, the CS acquires a negative value. This corresponds to Pavlovian fear learning in two-factor theory. Subsequently, when the CS is presented and the avoidance response terminates it, the model goes from a situation with a negative value to a situation with a value of 0 (because in the absence of the CS, no shock is predicted). This produces a positive prediction error, which reinforces the response. The positive prediction error is caused by a reduction in fear, so the avoidance response is reinforced by fear reduction, as in two-factor theory.

Despite these similarities between the actor-critic and two-factor theory, the fact that in the model the S-R strength is changed on the basis of prediction errors rather than on the basis of external outcomes allows the model to explain findings that two-factor theory cannot explain (Maia 2010). Consider, for example, the persistence of the avoidance response after fear of the CS has extinguished. When fear of the CS extinguishes, the value of the CS becomes 0. When the CS is subsequently presented and the model performs the avoidance response, the model goes from a situation with a value of 0 to another situation with a value of 0 (because the US is not predicted in either the presence or absence of the CS). The prediction error is therefore 0, and the strength of the S-R association remains unchanged. The response therefore persists after fear has extinguished. In fact, in the model, the response persists perpetually, unless responding has some cost (representing the effort of responding).

With certain extensions, the model can also explain the effects of US avoidance on learning (Maia 2010). Furthermore, the model explains some findings that no other theory can explain – for example, the reduction in avoidance latencies that occurs with extended training (Maia 2010). The model may therefore provide the most comprehensive theory of avoidance to date. Furthermore, two aspects of this theory seem quite satisfactory. First, the model was developed in machine learning and is independently motivated on computational grounds. The theory makes no assumptions specific to avoidance; it simply shows that this general-purpose ► reinforcement-learning system explains a variety of findings in avoidance. Second, the model maps closely to the brain (Maia 2009), so it offers the prospect for an integrated neurobehavioral theory of avoidance.

Despite the strengths of this approach and the appeal of explaining a broad range of findings using a simple model, the fact that the actor-critic does not implement R-O contingencies may be a weakness of this approach. Substantial evidence suggests that in other instrumental-conditioning paradigms, subjects can learn both S-R and R-O contingencies, and the same may apply to avoidance learning. Other reinforcement-learning models learn R-O contingencies (Maia 2009), and such models have been used to explain the effects of antipsychotic drugs on avoidance (Smith et al. 2004). A comprehensive theory of avoidance may have to include both S-R and R-O learning, as well as their interactions. Developing such a theory should be a major focus for future research.

Cross-References

- ► Abnormal Avoidance Learning
- ► Behaviorism and Behaviorist Learning Theories
- ► Conditioning
- ► Fear Conditioning in Animals and Humans
- ► Law of Effect

► Operant Behavior
► Pain-Relief Learning
► Punishment and Reward
► Reinforcement Learning
► Reinforcement Learning in Animals

References

Brown, J. S., & Jacobs, A. (1949). The role of fear in the motivation and acquisition of responses. *Journal of Experimental Psychology, 39*, 747–759.

Maia, T. V. (2009). Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cognitive, Affective, & Behavioral Neuroscience, 9*, 343–364.

Maia, T. V. (2010). Two-factor theory, the actor-critic model, and conditioned avoidance. *Learning & Behavior, 38*, 50–67.

Mowrer, O. H. (1947). On the dual nature of learning – a reinterpretation of conditioning and problem solving. *Harvard Educational Review, 17*, 102–148.

Seligman, M. E. P., & Johnston, J. C. (1973). A cognitive theory of avoidance learning. In F. J. McGuigan & D. B. Lumsden (Eds.), *Contemporary approaches to conditioning and learning* (pp. 69–110). Washington, DC: Winston.

Smith, A., Li, M., Becker, S., & Kapur, S. (2004). A model of antipsychotic action in conditioned avoidance: A computational approach. *Neuropsychopharmacology, 29*, 1040–1049.

Avoidant

► Fear of Failure in Learning

Avoidant Oriented

► Fear of Failure in Learning

Aware

► Consciousness and Emotion: Attentive vs. Pre-attentive Elaboration of Face Processing

Axiom Schema

Schemas are used in formal logic to specify rules of inference, in mathematics to describe theories with infinitely many axioms, and in semantics to give adequacy conditions for definitions of truth. Accordingly, an axiom schema is a well-formulated formula in the language of an axiomatic system, in which one or more schematic variables may appear. Well-known examples of axiom schemas are the induction schema as part of Peano's axioms for the arithmetic of the natural numbers, and the axiom schema of replacement as part of the Zermelo–Fraenkel set theory.