American Psychological Association

Journal of Experimental Psychology:
Animal Learning and Cognition

ISSN: 2329-8456

2020, Vol. 46, No. 3, 205–214
http://dx.doi.org/10.1037/xan0000251

Can the Stimulus Processing Assumptions of the Sometimes-Opponent-Process (SOP) Model Explain Instances of Contextual Learning?

Edgar H. Vogel and Fernando P. Ponce
University of Talca

Susan E. Brandon
New York, New York

One of the most persisting assertions in Allan Wagner's view of conditioning is that the environment or context in which significant events occur can develop an association with these events, more or less in the same way as conditioned and unconditioned stimuli become associated with each other. He was drawn to this idea by evidence of contextual fear conditioning, contingency effects, some instances of context-specificity of long-term habituation, and latent inhibition. From a theoretical point of view, however, homologizing contexts to conditioned stimuli is not as simple as it seems, especially when quantitative theories are involved, as is the case of Wagner's work. It might be, for instance, that contexts cannot be represented merely as long-duration conditioned stimuli, in which case, no net contextual learning can occur due to the context being less correlated with reinforcement than with nonreinforcement. In this article, we use Wagner's sometimes-opponent-process model of conditioning to comment on the quantitative nature of this challenge. Also, based on an idea sketched by Mazur and Wagner, we describe a set of quantitative strategies that might be usefully considered to solve this dilemma within the general framework of Wagner's theory.

Keywords: SOP model, priming theory, Pavlovian conditioning, context learning

Supplemental materials: http://dx.doi.org/10.1037/xan0000251.supp

Perhaps the core and most influential principle in Allan Wagner's theorizing is that expected events are less processed than unexpected events. An early development of this notion appeared in two papers that he published with Robert Rescorla in 1972 (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972). They emphasized that the effectiveness of an unconditioned stimulus for reinforcing a conditioned stimulus (CS)–unconditioned stimulus (US) association on a given trial depends on whether or not the US is already signaled or predicted by the set of CSs that are present in the trial. Although much of their analysis was devoted to explicit experimental cues, they indicated that the expectation of the US could come not only from conditioned stimuli, but also situational cues or the context.

In concluding their seminal paper, Rescorla and Wagner (1972) presented a detailed demonstration of how their model could explain a series of findings that came to be known as "contingency effects." The prototypical data were provided by Rescorla (1968), who reported several experiments in fear conditioning of rats where the addition of some US-alone trials to regular CS–US trials reduced the terminal level of conditioned responding to the CS in direct proportion to the number of unsignaled USs. Rescorla and Wagner argued that these experiments might be theoretically described as instances where the context acts as a conditioned stimulus, which is reinforced when presented in compound with the target CS but also reinforced alone in some occasions. They proceeded to show that, according to their newly presented competitive learning rule, the more associative strength acquires to the context when reinforced alone, the less the associative tendencies accrued to the CS when reinforced in compound with the context. This interpretation was further encouraged by studies demonstrating that CS–US pairings with longer intertrial intervals produced better conditioning to the CS than when the intertrial interval was shorter, which may be seen as due to greater occasions for the extinction of the context–US association in the former case (e.g., Dweck & Wagner, 1970).

Of course, apart from the subtle and arguable participation of the context in contingency effects, there is also direct empirical evidence suggesting that the context may indeed acquire associations with the stimuli it embraces during conditioning. Contextual fear conditioning is an example of this, in which animals learn to freeze when placed in a context that has been paired with footshock, whereas no such freezing occurs when the context is modified (e.g., Blanchard & Blanchard, 1969; Bolles & Collier, 1976; Fanselow, 1990).

In the late 1970s, Wagner extended the core premises of the Rescorla–Wagner model with a theory of priming, which holds that the processing of events in active memory, and therefore their capability to acquire associations and to generate a response, is

---

Edgar H. Vogel and Fernando P. Ponce, Faculty of Psychology, University of Talca; Susan E. Brandon, Consultant, New York, New York.

This work was supported by grant from Fondecyt (1160601) to Edgar H. Vogel and by PIA Ciencia Cognitiva, Centro de Investigación en Ciencias Cognitivas, Facultad de Psicología, Universidad de Talca.

Correspondence concerning this article should be addressed to Edgar H. Vogel, Facultad de Psicología, Universidad de Talca, 2 Norte 685, Talca 3460000, Chile. E-mail: evogel@utalca.cl

inversely proportional to the degree to which the events are expected (Wagner, 1976, 1978, 1979). A quantitatively refined version of the theory, the Sometimes-Opponent-Process (SOP) Model, came shortly afterward (Mazur & Wagner, 1982; Wagner, 1981). In this account, the association of the context with explicit experimental cues played a major role.

The sort of data that further encouraged Wagner to view contexts as conditioned stimuli came from the so-called "latent inhibition" effect, which refers to the observation that prior presentations of a CS alone in the experimental context cause a subsequent decrease in the ability of the CS to form CS–US associations when the pairings occur in the same context (Lubow & Moore, 1959; Reiss & Wagner, 1972; Rescorla, 1971). Wagner (1976) supposed that the same kind of mechanism might operate in some instances of context-dependent decrements in unconditioned responding after repeated stimulation or long-term habituation (Mertl, 1977; Montgomery, 1953; Peeke & Veno, 1973; Shalter, Fentress, & Young, 1977).

The assumption that context conditioning underlies phenomena as diverse as contingency effects, latent inhibition, and habituation is, in many ways, an appealing idea due to its simplicity. However, it has been criticized on empirical and theoretical grounds (e.g., Fanselow & Tighe, 1988; Hall, 1991; Hall & Channell, 1985; Hall & Honey, 1989; Mackintosh, 1987; Marlin & Miller, 1981). In this article, we focus on one fundamental theoretical challenge, which is how to reconcile the assumption (and evidence) that long CSs are expected to develop no net association with the US with the supposition that the contexts should be represented, in principle, as long-duration CSs.

The extinction of context–US associations is a general problem for theories that, like the Rescorla–Wagner and SOP models, compute acquisition during reinforcement and extinction during nonreinforcement. Consider, for instance, an experiment conducted by Hinson (1982), who reported a diminished acquisition of eyeblink conditioned responses to a 600-ms tone in a group of rabbits that previously had 10 sessions in which they received 20 presentations of a 100-ms US with a mean US–US interval of 3 min. Hinson's interpretation was that preexposure to the US in the context resulted in the development of a context–US association that subsequently blocked CS–US learning in that context. This interpretation requires accepting that the very long periods of nonreinforcement of the context (i.e., every 3 min for 20 trials over 10 sessions) survived extinction.

Although in some experimental protocols the exposure of the context to nonreinforcement is relatively minor in relationship to reinforcement (see, e.g., Rescorla and Wagner's simulations of Rescorla's [1968] data on fear conditioning, in which the total exposure to the context was only five times the total exposure to the CS), in other protocols, like that of Hinson (1982), the period of time in which the context is nonreinforced exceeds by far what would be needed to attain conditioning according to the standard assumptions of the theories (see also Lubow & Moore, 1959; Reiss & Wagner, 1972). In this article, we examine this issue quantitatively and propose some solutions within the framework of the SOP model.

Theoretical Simulations of Contextual Learning

The SOP Model and the Context as a Static CS

In the SOP model (Wagner, 1981), priming was rationalized by assuming that stimulus representation comprises a set of stochastic elements whose transition from inactivity to a primary activity state is prevented by their passing through a secondary state of activity, at which they arrive either via previous stimulus presentation or through associated sources. It is supposed that the capability of a stimulus to enter into associations and to produce its primary response depends on the proportion of its elements residing in the primary state. Therefore, the arrival of elements to the secondary state is an indication of refractoriness or priming.

The formal aspects of the model are simple. It is assumed that upon presentation of a stimulus, a proportion of elements, p1, move from inactivity (I) to primary activity (A1) and produce a response of a certain amplitude. Once in the A1 state, some elements move, first to the secondary state (A2) with a probability of pd1, and then back to inactivity with a probability of pd2, where they remain unless a new presentation of the stimulus occurs. Apart from the stimulus itself, other stimuli, via their association with the target stimulus, can promote elements directly from the inactive state to the secondary state of activity of the target with a probability of p2. The parameters p1, pd1, and pd2 and the variable p2 range between 0 and 1 and can be regarded as "transition probabilities." In that respect, their meaning and values are critical for the predicted temporal course of stimulus representations.

The parameter p1 dictates the rate at which elements are transported from I to A1, and its value is zero when the stimulus is off and positive when the stimulus is on. The magnitude of p1 can be seen as the intensity of the stimulus; therefore, we assume that $p1_{US} > p1_{CS} > p1_{context}$. The parameters pd1 and pd2 govern the passage of elements from A1 to A2, and from A2 to I, respectively; they are independent on whether the stimulus is on or off, and they influence how long the stimulus will be processed after its offset. In line with Wagner's work, we assume that $pd1 = 5pd2$ and set the same values for all cues. Finally, SOP states that any stimulus influences p2 of a target stimulus as a function of its A1 activity multiplied by the net value of its associative connection with the target.

In Pavlovian conditioning, changes in the association between a CS and a US (ΔV_{CS-US}) are assumed to be the result of excitatory minus inhibitory associations that develop simultaneously depending on the respective states of activity of the CS and the US. The acquisition of an excitatory CS–US association is the product of concurrent A1 activity of the CS and the US multiplied by a learning rate parameter, L^{+}. In contrast, acquisition of an inhibitory CS–US association is the product of concurrent A1 activity of the CS and A2 activity of the US multiplied by a learning rate parameter, L^{-}. These acquisitions rules can be summarized as follows:

$$\Delta V_{CS-US} = A1_{CS}(A1_{US} \times L^{+} - A2_{US} \times L^{-}) \qquad (1)$$

According to SOP, then, the relative distribution of elements in the three states of activity over time and across the experimental cues is critical for what is learned in an episode of conditioning. Figure 1A presents an example of an episode of 600 moments of duration in which a 10-moment CS is paired at its termination with

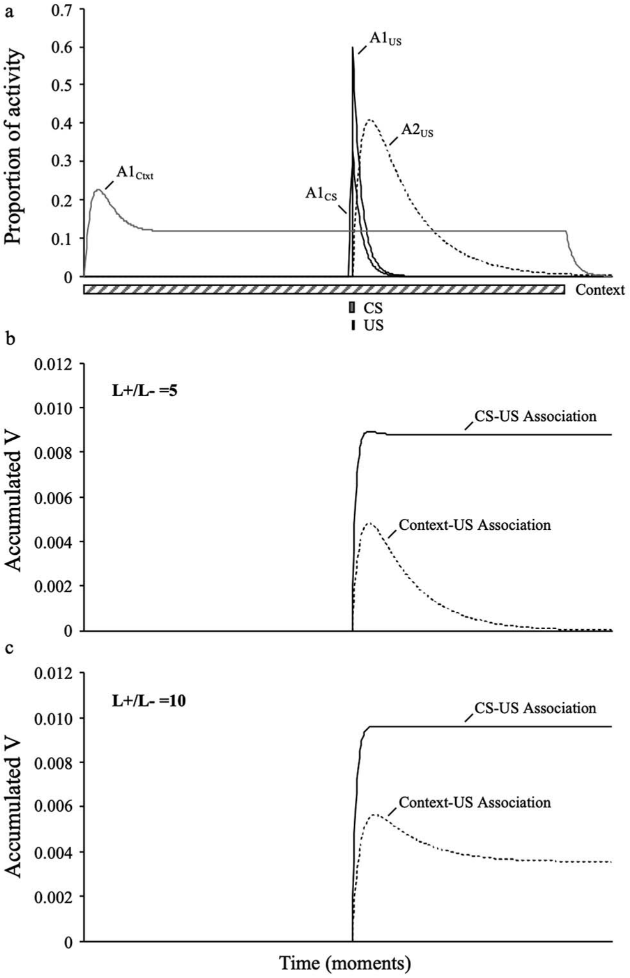

Figure 1. Simulations of the theoretical processes involved in one 600-moment episode with a single presentation of a 10-moment conditioned stimulus (CS) followed by a 1-moment unconditioned stimulus (US) in the midst of a given context. The context was assumed to be on for the entire episode. Panel A shows the courses of A1 activity of the CS and the context and the A1/A2 activity of the US during the entire episode. Panels B and C depict the accumulated CS–US and context–US associative strength as a function of time according to Equation 1. The parameters were $p1_{US} = 0.6$, $p1_{CS} = 0.1$, $p1_{Context} = 0.04$; $pd1_{US} = pd1_{CS} = pd1_{Context} = 0.1$; $pd2_{US} = pd2_{CS} = pd2_{Context} = 0.02$, $L^{+} = 0.01$ and $L^{-} = 0.002$ (B) and 0.001 (C).

a 1-moment US in a given context. The context is assumed to be identical to a CS that is turned on at the beginning of the episode and off a few moments before its termination. The figure displays the A1 activity of the three cues over time and the A2 activity of the US, which are the relevant theoretical processes for computing $V_{CS–US}$ and $V_{Context–US}$. As can be seen, the episode begins with an exponential growth of $A1_{Context}$ activity toward a peak, followed by a decline to a stable value where it remains until a final decay to inactivity due to the termination of the episode. The $A1_{CS}$ activity follows a similar pattern with the difference that it reaches a higher peak and terminates sooner, which is simply a reflection of the higher p1 and shorter duration of the CS. Finally, the presentation of the brief US leads to a very rapid increase and decay in $A1_{US}$ activity, followed by a slower $A2_{US}$ activity that persists for a considerable period of time after the offset of the US.

The theoretical patterns of activity displayed in Figure 1A show how learning would accrue differentially to the CS and the context, depending on the temporal contiguity of their respective A1 activities with $A1_{US}$ activity (excitatory learning) and $A2_{US}$ activity (inhibitory learning). In the case of the CS, notice that because the $A1_{CS}$ activity rises during the duration of the CS and falls gradually after its termination, it overlaps substantially more with $A1_{US}$ activity than with $A2_{US}$ activity, which leads to more excitatory than inhibitory learning at the end of the episode. On the contrary, because $A1_{Context}$ remains at a constant value during the entire duration of US activity, it overlaps equally with $A1_{US}$ and $A2_{US}$, leading to both incremental and decremental learning.

In this scenario, it is clear that the context–US association would depend importantly on the L^{+}/L^{-} ratio. If we assume that, as Wagner (1981) did, $L^{+}/L^{-} = pd1/pd2$, which, in the present example, means that L + should be five times L-, then the model predicts equal excitatory and inhibitory learning for a contextual cue of constant A1. This prediction can be appreciated in Figure 1B, which shows the accumulated V values for context–US and CS–US over time in the episode of Figure 1A. As can be seen, immediately after the US is turned on and for a few moments thereafter, both the CS and the context get some increments in V since $A1_{US}$ is high and $A2_{US}$ is still low. As time progresses and $A2_{US}$ grows, inhibitory learning occurs, producing decrements V_{CS} and $V_{Context}$. However, because the CS is turned off at the time of the US, $A1_{CS}$ activity decays rapidly toward zero, and no further learning occurs during the time of predominant $A2_{US}$ activity. For the context, which remains on until the end of the episode, learning is notably different: $V_{Context}$ begins to decrease when $A2_{US}L^{-}$ eventually surpasses $A1_{US}L^{+}$. As seen in the figure, the initial increments in $V_{Context}$ are extinguished at the end of the episode.

In the face of this problem, one may consider setting the $L+/L-$ ratio greater than the pd1/pd2 ratio, and thus, decreasing the impact of inhibition in the model. Figure 1C depicts the same information as Figure 1B, but this Time $L+/L- = 10$. As expected, by the end of the episode, $V_{Context}$ is now greater than zero. Although this simple solution seems effective in principle, it is still unsatisfactory. The problem now is that the net association acquired by the context during US episodes will provoke some $A2_{US}$ activity during the intertrial intervals, which would lead to sufficient inhibitory learning to overcome most of the acquired excitation (i.e., extinction). In the next section, we discuss a different approach to the problem.

The Context as a CS Susceptible to Disruption

Mazur and Wagner (1982) speculated that the only way of getting context–US associations is by assuming that the representation of the context changes systematically in the presence of the experimental cues. Vogel, Ponce, and Wagner (2019) instantiated this idea by proposing that the presentation of experimental cues, CSs or USs, suppresses the activity of the context by setting its p1 value to zero for some period of time. The suppressive effect of the experimental cues was assumed to be not immediate but occasioned only as the aggregated A1 activity of the experimental stimuli (disruptors), accumulated over time, reached an upper threshold, Td. Once the momentary A1 activity of the disruptors has decayed below some lower threshold, the accumulation is reset, and the context becomes again activated (i.e., $p1_{context} > 0$).

Figure 2A presents a pattern of activity similar to that of Figure 1, but this time only the US was included, and the context was assumed to be disrupted by the US when its accumulated activity reached a value of 3.2. As can be seen, shortly after the presentation of the US (approximately 8 moments), the activity of the context is suppressed. This fact biases learning toward producing more excitatory learning (when $A1_{US}$ activity is substantial, and the disruption is still ineffective) than inhibitory learning (when the $A2_{US}$ activity is substantial, but the disruption process is already effective). Since the context remains inactive for an extended period after the presentation of the US, no sufficient extinction occurs. Figures 2B and 2C show the net $V_{Context}$ accumulated over the episode for simulations with $L+/L- = 5$ (Figure 2B) and $L+/L- = 10$ (Figure 2C). For comparative purposes, in each figure, we copied the respective results of learning without the disrupted context. As can be appreciated, in both cases, the assumption of disrupted context leads to greater context–US learning.

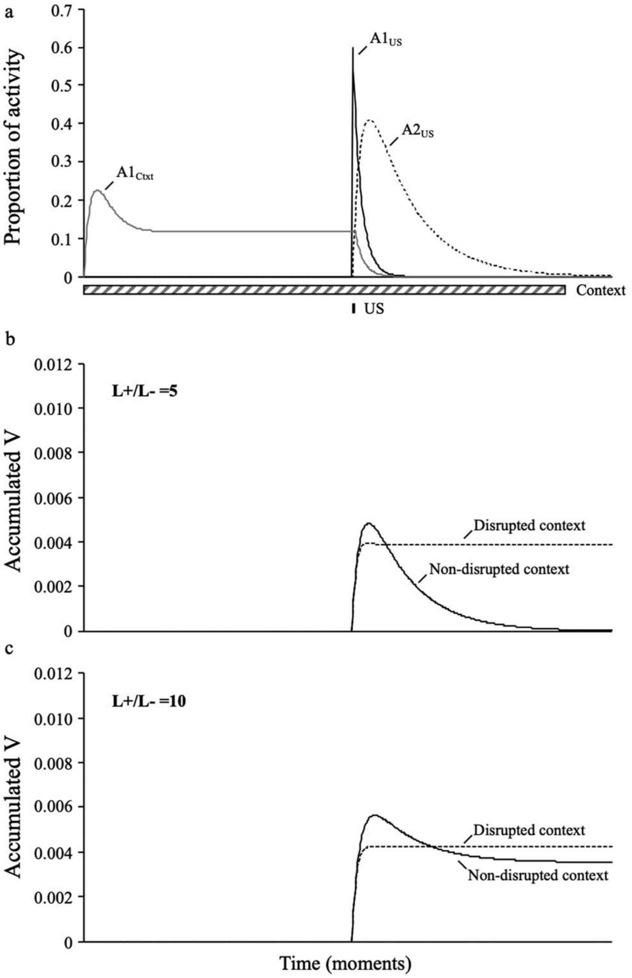

Figure 2. Simulations of the theoretical processes involved in one 600-moment episode with a single presentation of a 1-moment unconditioned stimulus (US) in the midst of a given context. It is assumed that the processing of the context was disrupted by the presentation of the US ($p1_{Context}$ was set to zero when the accumulated A1 activity of the US reached a threshold value of 3.2). Panel A shows the courses of A1 activity of the context and the A1/A2 activity of the US during the entire episode. Panels B and C contrast the accumulated context–US associative strength as a function of time for the disrupted and not disrupted representation of the context. The parameters were the same as those of Figure 1.

Finally, Figure 3 presents the predicted $V_{US\text{–}Context}$ over 50 episodes or trials like that of Figure 2 from the condition in which $L+/L- = 10$. The disrupted-context notion predicts a progressive increase in the context–US association over US episodes despite the prolonged period of nonreinforcement of each episode. On the contrary, when no disruption is assumed, the small association that is acquired in each trial is then overcome by extinction during the intertrial interval, leading to not much progress in V despite extensive training.

In sum, our assumption makes the context more like a CS being trained in a delay conditioning procedure where it is programmed to terminate with the appearance of the US. By being less salient than a standard CS, the context terminates "naturally" with the presentation of the US. Alternatively, contexts and CSs differ, according to our notion, in that CSs are not assumed to be so disrupted by other experimental cues. Thus, if a salient CS was programmed to continue beyond US presentation, we suppose that it would not be disrupted and will develop little conditioned responding, which is a well-documented fact (e.g., Smith, 1968).

A Componential Representation of the Context

Although the disruption of the context by experimental cues approach is, in itself, sufficient to address contextual learning, as one considers the nature of contexts, as they have been experimentally implemented and discussed in the literature, it is clear that a single molar component does not naturally emulate situational cues. Indeed, some authors have suggested that the representation of the context should involve a number of different aspects, perhaps in different modalities, and that these different aspects are less well-controlled than are the CSs and USs that may be administered by the experimenter (e.g., Estes, 1955; Fanselow, 2010; Rudy, 2009). One can think, for instance, that these components vary randomly with the subject's scanning of the environment. Here, we propose that this notion can be instantiated in SOP by assuming that the context is represented by a small number of components of independent activity that follow SOP's rules of activation and that form separate associations with the experimental cues. We suppose that a random process governs the activation of these elements, which are susceptible to disruption by explicit cues. Thus, the A1 processing of the context, rather than being mostly static, would change randomly over time in the absence of experimental cues and be disrupted in their presence.

We instantiate this idea as follows: First, let the context be represented by five components designated as c1, c2, c3, c4, and c5. Suppose that at any time, one and only one component can have a p1 value greater than zero (i.e., each element has a likelihood of 1/5 of being activated at any moment). Once a component is selected for activation, its p1 is set to 0.04 and keeps this value for a duration that varies according to a normal distribution with a mean of 70 moments and a standard deviation of 5. Thus, each

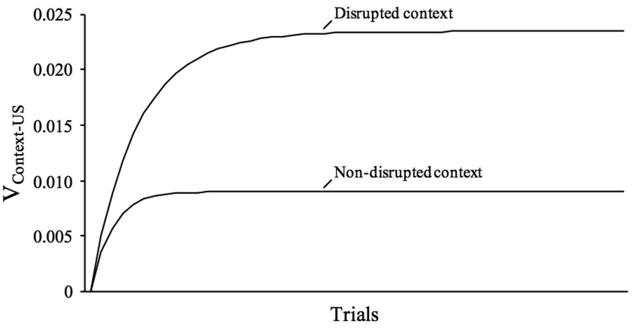

Figure 3. Predicted change in context–unconditioned stimulus (US) association over 50 presentations of a US in one 600-moment episode. The figure contrasts the results of a “disrupted context” representation, in which the $p1_{\text{Context}}$ value was set to zero when the accumulated A1 activity of the US reached a threshold value of 3.2, with the results of a “nondisrupted context” representation, in which $p1_{\text{Context}}$ was unaffected by the US and that, therefore, stayed at its static level throughout. The parameters were $p1_{\text{Context}} = 0.04$, $pd1_{\text{Context}} = 0.1$, $pd2_{\text{Context}} = 0.02$; $p1_{\text{US}} = 0.6$; $pd1_{\text{US}} = 0.1$; $pd2_{\text{US}} = 0.02$; $L+ = 0.01$ and $L- = 0.001$.

component will be processed with variable durations across time. Second, in the presence of experimental cues, the p1 value of any active component is set to zero for some period of time with the same rule as that described in the preceding section. This process causes the inactivation of the already active element and precludes any further activation in the presence of the US.

Thus, over some period of time, each element might be activated several times, in different orders and with different durations. This pattern of activation is exemplified in Figure 4, which displays the opportunities for the activation of each contextual component in eight episodes similar to that of the previous figures. The raster plots simply depict the moments during the episode in which the A1 of each component is greater than zero. Because a random process controls the activity of the components, we plotted the data of four exemplars during eight identical episodes or trials. It is clear that in the absence of the US, each exemplar and each episode display a very distinctive pattern of component activation. This variability can be regarded as within-subjects variability in the pattern of activity across time or as between-subjects variability in a single episode of time. The figure also shows that upon presentation of the US, the opportunity for activation of all components decays uniformly toward zero in every case.

Apart from offering a realistic representation of situational cues, the componential representation of the context also provides a means of obtaining levels of conditioning similar to that using a molar representation. This fact is illustrated in Figure 5, which depicts the acquisition curve for the componential representation of the context, along with the acquisition curve for the molar representation that was already depicted in Figure 3. The curve of the componential representation is based on the mean association of the five random components and across 30 exemplars or “subjects.” As expected, because of the random occurrence of each component, the componential representation displayed a more irregular pattern and a lower slope in the acquisition function than the molar representation. The critical aspect, however, is the fact that both functions reached similar terminal levels of conditioning.

It should be further emphasized that there are subtle but not trivial differences in the way in which the two types of representations acquire those levels of conditioning. On the one hand, in the componential case, each component is activated for a short period of time, which implies fewer extinction episodes than in the molar representation, in which extinction affects the context as a whole over the entire duration of the context. On the other hand, because there is some degree of overlap in the activity of contiguous components, they overshadow each other to some degree. On balance, contextual learning accrued to a molar representation is approximately equivalent to that accrued to a single component in the componential representation. Nonetheless, consider also that the overlap between components would lead to some degree of summation of the $A2_{\text{US}}$ activity generated by each of them. Remember that associatively generated $A2_{\text{US}}$ activity is, according to SOP, an indication of conditioning. Figure 6 exemplifies this fact by plotting the peak anticipatory $A2_{\text{US}}$ activity generated by the context over training. As can be seen, the consequences of the association (measured as the anticipatory $A2_{\text{US}}$ activity) is substantially greater for the componential than the molar representation in the majority of trials. This result is a further advantage of the componential representation over the molar representation: the componential representation predicts greater priming effects of the context.

One might wonder whether this sort of componential representation of the context can cope with the extinction problem by itself, without introducing the particular processing assumption that experimental events suppress contextual cues. Unfortunately, this is not the case. Of course, the randomness of the assumed contextual elements allows them to avoid many extinction episodes that they would suffer if they were active throughout the session. However, their randomness also precludes them from receiving all the available reinforcement. That is, random elements acquire and extinguish associative strength at slower rates than if they had a fixed pattern of activity, but this does not mean that they are biased toward acquisition. Ultimately—if a sufficiently large number of trials occurs—the expected time in which any given random element will have its A1 activity overlapping with $A2_{\text{US}}$ activity (evoked by itself, by the US, or even by the trace of another component) will be sufficient to produce enough extinction to overcome the expected opportunities for acquisition. Suppressing the activity of the context during the predominant $A2_{\text{US}}$ activity is essential to bias learning toward excitation.

Simulating Some Facts Attributed to Contextual Learning

In this section, we present a few simulations illustrating how SOP, with the additional assumptions presented above, namely a componential and disruptable representation of the context, describes some empirical regularities of contextual learning. Because the model comprises a probabilistic representation of the context, the results are presented as the mean across several simulations.

In all simulations, the context was represented by five serial components which were processed for variable durations (randomly selected from a normal distribution with $M = 70$ and $SD = 5$) before being replaced by a next component (that could be another element or the same). In this context, a 1-moment US, a 5-moment CS, or both, depending on the particular phenomena of

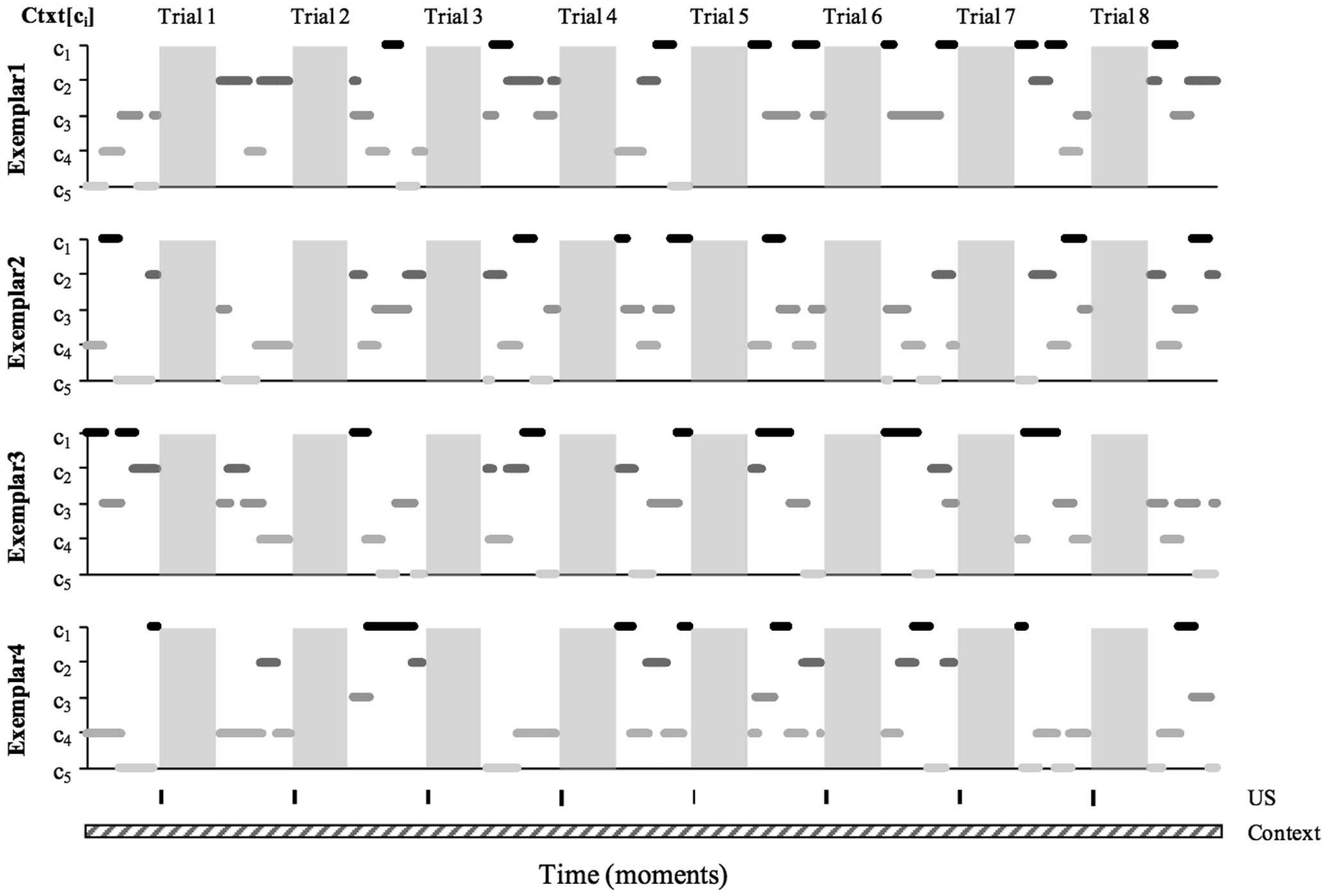

Figure 4. Simulations of the period of time in which each contextual component has an A1 value greater than zero over eight presentations of an unconditioned stimulus (US) in eight 600-moment episodes. The figure depicts the results of four independent computer simulations or exemplars.

interest, were presented a certain number of times at fixed intervals. The parameters for activation of the contextual components were identical to those used for explicit CSs except for the use of a lower p1 value ($p1 = 0.04$). Whenever the aggregated accumulated A1 activity of the conditioning cues reached a threshold of 3.2, the p1 values of all contextual components were set to zero. When the momentary aggregated A1 activity of the experimental cues reached a lower bound close to 0, accumulation was reset, and the contextual components resumed their dynamic of activation and replacement.

One-Trial Context Conditioning

It has been shown that contextual fear conditioning can be accomplished with a single presentation of a shock US (e.g., Fanselow, 1990). The magnitude of this effect, however, depends dramatically on the interval between placement in the context and US presentation, following an approximately increasing linear relationship (Blanchard, Fukunaga, & Blanchard, 1976; Fanselow, 1990). In the extreme case, if the US is delivered simultaneously with placement in the context, no learning occurs, in what has been called “the immediate shock freezing deficit” (Fanselow, 1990).

In principle, SOP can explain these observations through the nonlinear pattern of the A1 activity of the theoretical contextual components. Given the standard transition probabilities assumed by SOP, any contextual component, when activated, follow a characteristic pattern with a slow and exponential growth of primary activity toward a peak, followed by a decay toward an intermediate value, and a final decay toward zero when the component is replaced by another component. Of course, because the activation and deactivation of the components depends on a random process, there is some variability in the pattern of activity of each component across samples. However, the aggregated trace (i.e., the sum of A1 activity across components) is expected to be quite regular. With this pattern, it is intuitive to conceive that if the US occurs at very early stages of processing of the context when the expected trace is still weak, less learning will occur than if the US is placed at a later stage, when the expected trace is stronger. This prediction is demonstrated in Figure 7, which presents the results of 10 simulations of a 600-moment episode in which the US was presented with a delay of 0, 1, 2, 4, 8, 16, or 32 moments after the initiation of the context. As shown in the figure, the mean context–US association increases as a function of the delay up to

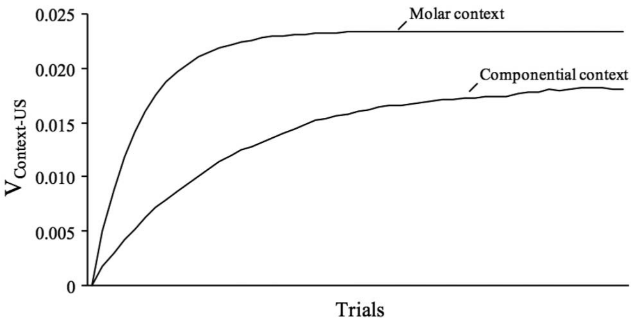

Figure 5. Predicted change in context–unconditioned stimulus (US) association over 50 presentations of a US at intervals of 600 moments. It was assumed that the processing of the context was disrupted by the presentation of the US. The figure contrasts the results of the "molar context" (with the same assumptions as those of Figure 3B) and the "componential context" (with the same assumptions of Figure 4). In the case of the componential context, the figure displays the mean association between each contextual component and the US averaged across 30 independent simulations. The parameters were the same as those of Figure 3.

a maximal at an intermediate interval with some decay at longer intervals.

Although SOP's account of one-trial context conditioning might seem simple and reasonable at first, there is some evidence suggesting that things may be more complicated. For instance, it has been shown that the immediate shock freezing deficit can be alleviated if animals are preexposed to the context (Fanselow, 1990). This observation seems to be calling for a more complex and dynamic representation of contexts, which is outside SOP's principles discussed here (for a theoretical alternative that focuses on this specific topic, see Krasne, Cushman, & Fanselow, 2015).

Multiple-Trials Context Conditioning

Fanselow and Tighe (1988) provided evidence of a sort of "trial spacing" effect in contextual learning. They presented rats with three shocks separated by 3, 16, or 60 s and found that animals in the 60-s group exhibited more fear in a testing session than their counterparts in the 3- and 16-s groups. Subsequent studies, using a broader range of intershock intervals, have shown that, indeed, there is an optimal interval of intermediate duration in which the degree of learning is maximal, suggesting that the relationship between the strength of conditioning and the intershock interval follows an inverted-U function (e.g., Barela, 1999).

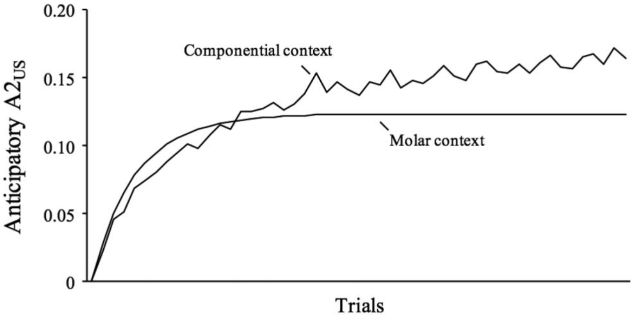

Figure 6. Predicted change in peak anticipatory $A2_{US}$ activity over 50 presentations of the unconditioned stimulus (US) at intervals of 300 moments. It is assumed that the processing of the context was disrupted by the presentation of the US. In the case of the componential context, the figure displays the mean across 30 independent simulations. The parameters were the same as those of Figure 3.

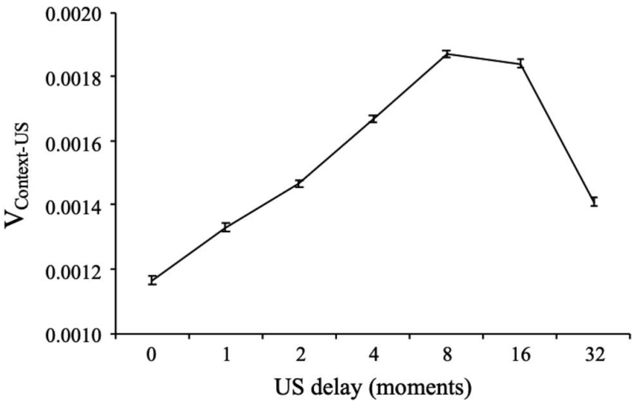

Figure 7. Predicted context–unconditioned stimulus (US) association according to the componential representation of the context after a single presentation of the US as a function of the delay between initiation of the context and US presentation. In all conditions, the episode ended 300 moments after the presentation of the US. The figure displays the mean association between each contextual component and the US averaged across 10 simulations of each condition. Error bars represent the standard error of the means. For more details, see the introduction to the "Simulating Some Facts Attributed to Contextual Learning" section.

SOP's account of this fact can be stated in the following terms: (a) very short US–US intervals result in very little processing of the context, which suffers from being almost permanently disrupted by each presentation of the US, so, very little learning occurs; (b) very long intervals allow the context to recover after disruption and, because of that, to suffer from prolonged periods of extinction, leading very little learning; and (c) intermediate US–US intervals allow for optimal processing of the context and, hence, for learning to occur. These facts are quantitatively demonstrated in Figure 8, which presents the results of 10 simulations involving 50 trials with US–US intervals of 100, 200, 300, 400, 500, and 600 moments.

Effects of Context Conditioning on CS–US Associations

According to our suppositions, the processing of the context is disrupted by the experimental cues. In a somehow paradoxical manner, this disruption is precisely what makes possible the development of associations between the context and the cues that are embedded on it. As a result, the context develops the capacity to prime the experimental cues affecting, for instance, CS–US associations. To illustrate this, we simulated CS and US preexposure effects. The simulations involved two phases and three conditions. In Phase 1, there were 50 presentations of either the US alone in the context (US-preexposure condition) or the CS alone in the context (CS-preexposure condition) at an interval of 600 moments. In a third condition, there was no Phase 1 training (no-

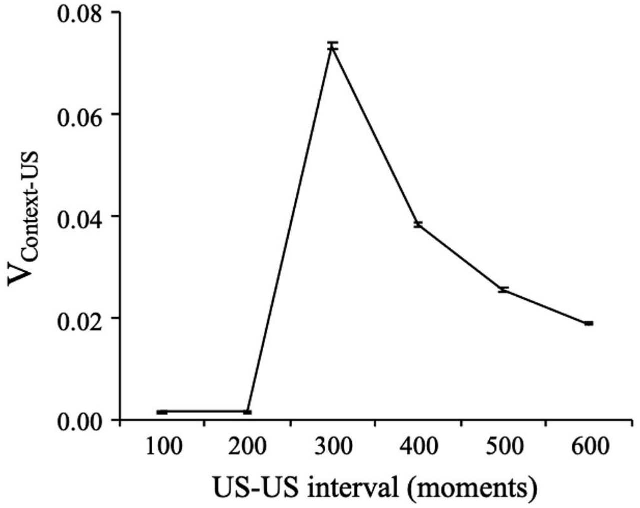

Figure 8. Predicted context–unconditioned stimulus (US) association after 50 presentations of the US as a function of the US–US interval according to a componential representation of the context in the Sometimes-Opponent-Process (SOP) Model. The figure displays the mean association between each contextual component and the US averaged across 10 simulations of each condition. Error bars represent the standard error of the means. For more details, see the introduction to the "Simulating Some Facts Attributed to Contextual Learning" section.

preexposure condition). The second phase was identical for the three conditions and involved 50 CS–US pairings in the same context as that of Phase 1. Figure 9 presents the results of this procedure, expressed as the CS–US association accrued by the end of Phase 2 and averaged across 10 simulations of each condition. The figure indicates that the CS–US association developed in Phase 2 depended on the association that the context developed with the experimental cues in Phase 1. In the no-preexposure condition, since there was no Phase 1, neither the US nor the CS were primed by the context, leading to "normal" CS–US acquisition in Phase 2. In the CS–preexposure condition, the CS came to be associatively primed by the context at the end of phase 1, which prompted a diminished $A1_{CS}$ processing, leading to reduced excitatory learning when the CS then was paired with the US in Phase 2. Likewise, in the US–preexposure condition, the US became primed by the context, which causes decreased $A1_{US}$ and increased $A2_{US}$ activity, leading to reduced excitatory and increased inhibitory learning, respectively. These results agree with the demonstrations of latent inhibition (Lubow & Moore, 1959; Reiss & Wagner, 1972) and contextual blocking (Hinson, 1982).

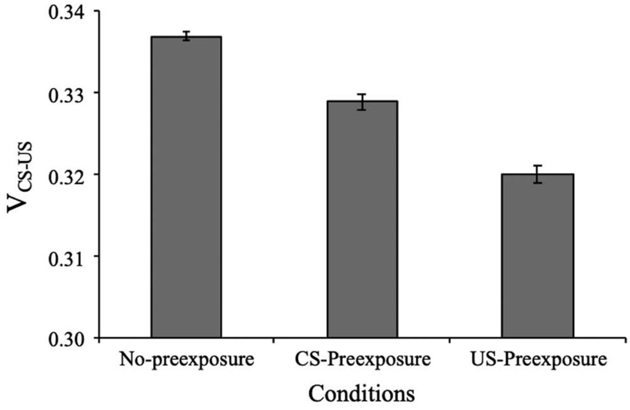

Figure 9. Simulation of preexposure effects according to the componential representation of the context in the Sometimes-Opponent-Process (SOP) Model. In the conditioned stimulus (CS–) and unconditioned stimulus (US)–preexposure conditions, the CS or the US, respectively, was presented 50 times in the context followed by 50 CS–US pairings. In the no-preexposure condition, there were only 50 CS–US pairings. The figure displays the predicted CS–US association averaged across 10 simulations of each condition. Error bars represent the standard error of the means. For more details, see the introduction to the "Simulating Some Facts Attributed to Contextual Learning" section.

Figure 10 illustrates how, by making the same assumptions as those made for preexposure effects, SOP appropriately describes contingency effects. The simulation involved 10 replications of each of two conditions. In the positive-contingency condition, a 5-moment CS terminated with a 1-moment US on 25 occasions, which were separated by an interval of 1,200 moments. In the low-contingency condition, there were 25 CS–US episodes and 25 US-alone episodes, separated by 600 moments. In both conditions, the experimental cues were embedded in a componential context. As seen in the figure, the mean CS–US association developed in each condition agreed with Rescorla's (1968) observation of greater conditioned responding to the CS in the positive- than in low-contingency condition, despite the fact that the two conditions received the same number of CS–US pairings.

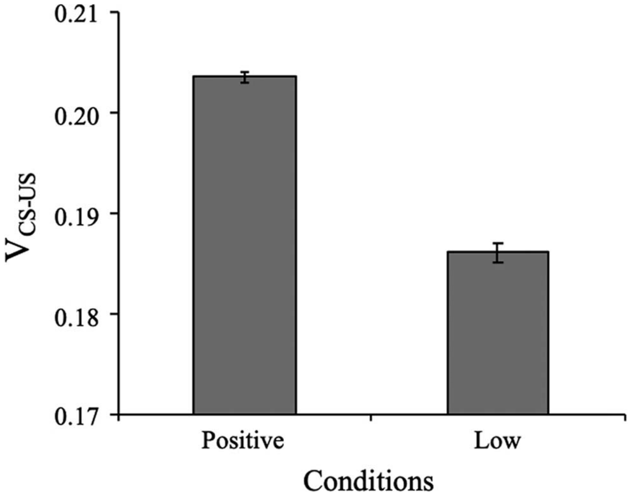

Figure 10. Simulation of contingency effects according to the componential representation of the context in the Sometimes-Opponent-Process (SOP) Model. In the positive-contingency condition, the conditioned stimulus (CS) was paired with the unconditioned stimulus (US) on 25 occasions, which were separated by intervals of 1,200 moments, and in the low-contingency condition, there were 25 CS–US episodes and 25 US–alone episodes, separated by 600 moments. The figure displays the predicted CS–US association averaged across 10 simulations of each condition. Error bars represent the standard error of the means. For more details, see the introduction to the "Simulating Some Facts Attributed to Contextual Learning" section.

Concluding Comments

In this article, we have focused on the questionable ability of the SOP model to address contextual learning. We hope that the basic demonstrations that we have provided will stimulate thinking of contexts as special types of conditioned stimuli. We propose that two features make the useful distinction: A componential representation of the context that is susceptible to disruption by explicit cues.

As shown in our simulations, the view that contexts are susceptible to disruption by experimental events is a simple and powerful strategy to predict contextual learning. It provides a plausible and testable explanation of why a context can develop associations with experimental cues despite its prolonged exposure to nonreinforcement (see the online supplemental material for some examination of the parameter dependence of this solution). There are several possibilities for how this disruption might take place within SOP's framework. Here, for illustrative purposes, we followed the scheme of Vogel et al. (2019) of disrupting the context as a function of the accumulated A1 activity of the experimental cues. Nevertheless, we could equally well have chosen to use a weighted function of the momentary A1 and A2 activity of the experimental cues. In the latter case, context–US learning might benefit from a sort of "auto-disruption" caused by the $\mathrm{A2}_{\mathrm{US}}$ activity that the context itself evokes during periods of extinction. The choice of a rule of this type or another is not inconsequential. However, given that the goal of this article was merely to show the suitability of the idea of disruption, we dismiss the temptation of being too speculative and stay with the simpler rule of using only A1 activity.[1]

Regardless of the preferred rule for contextual disruption, our general supposition is that different experimental stimuli of varied intensity and duration (e.g., CSs and USs), presented alone, in sequence, or in compound, should affect the processing of the context differentially; hence, the association of the context with these cues is also predicted to depend on these parameters. It would be interesting, for instance, to compare the strength of the association between the context and a given cue when that cue is followed by stimuli of different intensities or by expected versus unexpected events. In the absence of a significant corpus of data like those, the rule and parameters for the disruption that we presented in this article should remain as tentative.

In this article, we take one step further and also propose that contexts can be represented as comprising a set of components. For the sake of parsimony, we presumed the existence of only five components that compete with each other for processing through a random process of activation. This sort of representation seems to be a natural way of representing how animals scan multisensorial environments and, as shown in this article, it is entirely consistent with the idea that the context is relatively labile in its processing. Another advantage of multicomponent representations of the context, which we did not consider here but which has been advocated by others, is that the components and their dynamics can be easily used to describe the fact that the context can alternatively behave as an elemental as well as a configural stimulus (e.g., Fanselow, 2010; Rudy, 2009; Rudy, Huff, & Matus-Amat, 2004).

Finally, we would like to emphasize that we do not claim that all instances of Pavlovian conditioning in which contexts have been demonstrated to play a role are to be uniquely explained by the direct association of the context with the experimental events. Indeed, there is a lively and enduring debate on this matter (e.g., Bouton, 1993, 2004; Fanselow & Tighe, 1988; Hall, 1991; Hall & Rodríguez, 2019; Lubow & Gewirtz, 1995; Mackintosh, 1987; Roberts, 2019; Rosas, Todd, & Bouton, 2013; Simon-Kutscher, Wanke, Hiller, & Schwabe, 2019). Our purpose was to make the conceptual point that if an association between the context and the experimental cues is going to be theoretically presumed, some special assumptions about its representation may be useful.

---

[1] Wagner (1981) proposed that the experimental cues may interfere with each other's processing by mutually increasing their decay probabilities, pd1 and pd2 (see distractor rules in Wagner, 1981, p. 23). This notion allows that the contiguous presentation of several experimental cues, which have a delimited duration (i.e., CSs and USs), might result in the cues being processed for a shorter time than they otherwise would be. Unfortunately, this stratagem would not be effective to disrupt the processing of the context, because the context, in contrast to the experimental cues, is permanently "on" (i.e., $\mathrm{p1context} > 0$ throughout the session). Merely increasing the decay probabilities of the context without decreasing its p1 value will only produce a relatively flat pattern of activity of the context.

References

Barela, P. B. (1999). Theoretical mechanisms underlying the trial-spacing effect in Pavlovian fear conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 25,* 177–193. http://dx.doi.org/10.1037/0097-7403.25.2.177

Blanchard, R. J., & Blanchard, D. C. (1969). Crouching as an index of fear. *Journal of Comparative and Physiological Psychology, 67,* 370–375. http://dx.doi.org/10.1037/h0026779

Blanchard, R. J., Fukunaga, K. K., & Blanchard, D. C. (1976). Environmental control of defensive reactions to a cat. *Bulletin of the Psychonomic Society, 8,* 179–181. http://dx.doi.org/10.3758/BF03335118

Bolles, R. C., & Collier, A. C. (1976). The effect of predictive cues on freezing in rats. *Animal Learning & Behavior, 4,* 6–8. http://dx.doi.org/10.3758/BF03211975

Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin, 114,* 80–99. http://dx.doi.org/10.1037/0033-2909.114.1.80

Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning & Memory, 11,* 485–494. http://dx.doi.org/10.1101/lm.78804

Dweck, C. S., & Wagner, A. R. (1970). Situational cues and correlation between CS and U.S. as determinants of the conditioned emotional response. *Psychonomic Science, 18,* 145–147. http://dx.doi.org/10.3758/BF03332345

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review, 62,* 145–154. http://dx.doi.org/10.1037/h0048509

Fanselow, M. S. (1990). Factors governing one-trial contextual conditioning. *Animal Learning & Behavior, 18,* 264–270. http://dx.doi.org/10.3758/BF03205285

Fanselow, M. S. (2010). From contextual fear to a dynamic view of memory systems. *Trends in Cognitive Sciences, 14,* 7–15. http://dx.doi.org/10.1016/j.tics.2009.10.008

Fanselow, M. S., & Tighe, T. J. (1988). Contextual conditioning with massed versus distributed unconditional stimuli in the absence of explicit conditional stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 14,* 187–199. http://dx.doi.org/10.1037/0097-7403.14.2.187

Hall, G. (1991). *Perceptual and Associative Learning*. Oxford, UK: Clarendon Press. http://dx.doi.org/10.1093/acprof:oso/9780198521822.001.0001

Hall, G., & Channell, S. (1985). Differential-effects of contextual change on latent inhibition and on habituation of an orienting response. *Journal* *of Experimental Psychology: Animal Behavior Processes, 11,* 470–481. http://dx.doi.org/10.1037/0097-7403.11.3.470

Hall, G., & Honey, R. C. (1989). Contextual effects in conditioning, latent inhibition, and habituation: Associative and retrieval functions of contextual cues. *Journal of Experimental Psychology: Animal Behavior Processes, 15,* 232–241. http://dx.doi.org/10.1037/0097-7403.15.3.232

Hall, G., & Rodríguez, G. (2019). Attention to perceive, to learn and to respond. *The Quarterly Journal of Experimental Psychology, 72,* 335–345. http://dx.doi.org/10.1080/17470218.2017.1339719

Hinson, R. E. (1982). Effects of UCS preexposure on excitatory and inhibitory rabbit eyelid conditioning: An associative effect of conditioned contextual stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 8,* 49–61. http://dx.doi.org/10.1037/0097-7403.8.1.49

Krasne, F. B., Cushman, J. D., & Fanselow, M. S. (2015). A Bayesian context fear learning algorithm/automaton. *Frontiers in Behavioral Neuroscience, 9,* 112. http://dx.doi.org/10.3389/fnbeh.2015.00112

Lubow, R. E., & Gewirtz, J. C. (1995). Latent inhibition in humans: Data, theory, and implications for schizophrenia. *Psychological Bulletin, 117,* 87–103. http://dx.doi.org/10.1037/0033-2909.117.1.87

Lubow, R. E., & Moore, A. U. (1959). Latent inhibition: The effect of nonreinforced pre-exposure to the conditional stimulus. *Journal of Comparative and Physiological Psychology, 52,* 415–419. http://dx.doi.org/10.1037/h0046700

Mackintosh, N. J. (1987). Neurobiology, psychology and habituation. *Behaviour Research and Therapy, 25,* 81–97. http://dx.doi.org/10.1016/0005-7967(87)90079-9

Marlin, N. A., & Miller, R. R. (1981). Associations to contextual stimuli as a determinant of long-term habituation. *Journal of Experimental Psychology: Animal Behavior Processes, 7,* 313–333. http://dx.doi.org/10.1037/0097-7403.7.4.313

Mazur, J. E., & Wagner, A. R. (1982). An episodic model of associative learning. In M. Commons, R. Herrnstein, & A. R. Wagner (Eds.), *Acquisition: Quantitative analyses of behavior* (pp. 3–39). Cambridge, MA: Ballinger.

Mertl, A. S. (1977). Habituation to territorial scent marks in the field by *Lemur catta*. *Behavioral Biology, 21,* 500–507. http://dx.doi.org/10.1016/S0091-6773(77)92154-X

Montgomery, K. C. (1953). Exploratory behavior as a function of similarity of stimulus situations. *Journal of Comparative and Physiological Psychology, 46,* 129–133. http://dx.doi.org/10.1037/h0055101

Peeke, H. V. S., & Veno, A. (1973). Stimulus specificity of habituated aggression in the stickleback (Gasterosteus aculeatus). *Behavioral Biology, 8,* 427–432. http://dx.doi.org/10.1016/S0091-6773(73)80083-5

Reiss, S., & Wagner, A. R. (1972). CS habituation produces a "latent inhibition effect" but no active conditioned inhibition. *Learning and Motivation, 3,* 237–245. http://dx.doi.org/10.1016/0023-9690(72)90020-3

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66,* 1–5. http://dx.doi.org/10.1037/h0025984

Rescorla, R. A. (1971). Summation and retardation tests of latent inhibition. *Journal of Comparative and Physiological Psychology, 75,* 77–81. http://dx.doi.org/10.1037/h0030694

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

Roberts, W. A. (2019). The role of context in animal memory. *Learning & Behavior, 47,* 117–130. http://dx.doi.org/10.3758/s13420-019-00380-x

Rosas, J. M., Todd, T. P., & Bouton, M. E. (2013). Context change and associative learning. *WIREs: Cognitive Science, 4,* 237–244. http://dx.doi.org/10.1002/wcs.1225

Rudy, J. W. (2009). Context representations, context functions, and the parahippocampal-hippocampal system. *Learning & Memory, 16,* 573–585. http://dx.doi.org/10.1101/lm.1494409

Rudy, J. W., Huff, N. C., & Matus-Amat, P. (2004). Understanding contextual fear conditioning: Insights from a two-process model. *Neuroscience and Biobehavioral Reviews, 28,* 675–685. http://dx.doi.org/10.1016/j.neubiorev.2004.09.004

Shalter, M. D., Fentress, J. C., & Young, G. W. (1977). Determinants of response of wolf pups to auditory signals. *Behaviour, 60*(1–2), 98–114. http://dx.doi.org/10.1163/156853977X00298

Simon-Kutscher, K., Wanke, N., Hiller, C., & Schwabe, L. (2019). Fear without context: Acute stress modulates the balance of cue-dependent and contextual fear learning. *Psychological Science, 30,* 1123–1135. http://dx.doi.org/10.1177/0956797619852027

Smith, M. C. (1968). CS-US interval and U.S. intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology, 66,* 679–687. http://dx.doi.org/10.1037/h0026550

Vogel, E. H., Ponce, F. P., & Wagner, A. R. (2019). The development and present status of the SOP model of associative learning. *The Quarterly Journal of Experimental Psychology, 72,* 346–374. http://dx.doi.org/10.1177/1747021818777074

Wagner, A. R. (1976). Priming in STM: An information-processing mechanism for self-generated or retrieval-generated depression in performance. In T. J. Tighe & R. N. Leaton (Eds.), *Habituation: Perspectives from child development, animal behavior and neurophysiology* (pp. 95–128). Hillsdale, NJ: Erlbaum.

Wagner, A. R. (1978). Expectancies and the priming of STM. In S. H. Hulse, H. Fowler, & W. K. Honig (Eds.), *Cognitive processes in animal behavior* (pp. 177–209). Hillsdale, NJ: Erlbaum.

Wagner, A. R. (1979). Habituation and memory. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume for Jerzy Konorski* (pp. 53–82). Hillsdale, NJ: Erlbaum Associates.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–48). Hillsdale, NJ: Erlbaum.

Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In M. S. Halliday & R. A. Boakes (Eds.), *Inhibition and learning* (pp. 301–336). San Diego, CA: Academic Press.

Received November 30, 2019
Revision received April 17, 2020
Accepted April 19, 2020 ■