Metadata of the chapter that will be visualized online

Chapter Title	Contingency in Learning	
Copyright Year	2011	
Copyright Holder	Springer Science + Business Media, LLC	
Corresponding Author	Family Name	**Gallistel**
	Particle	
	Given Name	**Charles R.**
	Suffix	
	Division	Cognitive Science and Behavioral Neuroscience
	Organization	Rutgers University
	Address	152 Frelinghuysen Rd, New Brunswick, NJ, 08854-8020, USA
	Email	galliste@ruccs.rutgers.edu

C

Contingency in Learning

Charles R. Gallistel
Cognitive Science and Behavioral Neuroscience,
Rutgers University, New Brunswick, NJ, USA

Synonyms

Assignment of credit; Association; Contingency; Correlation; Dependence; Prediction; Retrodiction

Definitions

Contingency: the extent to which knowledge of one event reduces uncertainty about another. *Prediction*: the extent to which knowledge of one event's occurrence enables one to anticipate whether and/or when another event will occur. *Assignment of credit*: determining to which past event an outcome event should be attributed (retrodiction). *Association*: perceived contingency. *Instrumental conditioning*: a learning protocol in which a desired or undesired is contingent on an action of the subject (or agent). *Pavlovian conditioning:* a learning protocol in which the contingency between two events is varied. *Entropy*: the measure of amount of uncertainty, aka the amount of information available in a probability distribution. *Mutual information:* the sum of the entropies of the marginal distributions minus the entropy of their joint distribution. *Uncertainty coefficient*: the percent reduction in uncertainty about whether and/or when a predicted event will occur that is produced by the occurrence of a predictor event. A broadly useable measure of contingency or association.

Theoretical Background

The concept of contingency plays a central role in the analysis of commonly studied learning paradigms and also in research on human judgments of dependence, contingency, and causality. Despite its conceptual importance, there is surprisingly ~~small~~ psychological literature focusing on the following question: What is the proper definition or measure of contingency?

In Instrumental/Operant Conditioning

The concept of contingency is important in the study of instrumental conditioning, because the reinforcing event only reinforces the instrumental response if it is contingent on that response. In the operant conditioning literature, the concept has often been treated as unproblematic, perhaps because the experimenter specified the contingencies that were taken to be of interest. However, implicit in many treatments of reinforcement – and explicit in discussions of the role of delay of reinforcement – is the assumption that what really matters is not contingency per se but rather the close *temporal pairing* of response and reinforcement. This makes the question of the role of contingency in instrumental conditioning the same as the question of its role in Pavlovian conditioning. The challenge in both cases is to specify what constitutes "close."

As the study of reinforcement learning from a computational perspective has become a significant focus of research in computer science and cognitive neuroscience, there has been a greater realization that it was not obvious which aspects of a sequence of actions should be regarded as the aspect on which the feedback-providing outcome was contingent. How to determine this is the assignment of credit problem. It is the contingency problem seen from the other end. It can be reformulated as: What aspect or aspects of an action sequence is an outcome contingent on? One wants a measure of contingency or dependency that is mathematically well grounded and lends itself to the apportionment of contingency or dependency among possible predictors.

In Pavlovian/Classical Conditioning

The concept of contingency became important in the study of Pavlovian conditioning in the late 1960s when a series of experiments from different laboratories called into question the assumption that the temporal pairing was what drove the formation of an association between two stimuli or events (hereafter called the CS and US, with the CS being the predictor and the US the predicted event or stimulus). Rescorla (1968) posed the question whether it was the temporal pairing of CS and US or the CS–US contingency that led to the emergence of a conditioned

Norbert Seel (ed.), *Encyclopedia of the Sciences of Learning*, DOI 10.1007/978-1-4419-1428-6,

response (a response to the CS that anticipates the US). He fixed the number of co-occurrences (temporal pairings) of the CS and US and varied the contingency by varying the frequency of US during intervals when the CS was absent. When there were no US in the absence of the CS, a strong conditioned response was seen on the post-conditioning test trials, even when $p(US|CS)$, the probability of the US given the CS, was as low as 0.1. Regardless of the value of $p(US|CS)$, as the frequency of the US in the absence of the CS increased, the strength of the conditioned response on test trials diminished (see Fig. 1). When $p(US|\sim CS) = p(US|CS)$, that is, when the contingency was eliminated, there was no conditioned response. Thus, it is predictive (and retrodictive) power or contingency rather than temporal pairing that drives conditioning. That is also the implication of the phenomena of *blocking*, *overshadowing*, and *relative validity*, which were discovered at about the same time. All of these phenomena imply that the critical aspect of a conditioning protocol is the predictive power of CS (or of the response in instrumental conditioning), the extent to which it improves the subject's ability to anticipate when the US will occur.

Measures of Contingency

Most measures of contingency in the psychological literature derive from the numbers in a 2×2 contingency table (Table 1). Several have been used, but only two have suitable mathematical properties, such as ranging from 0 to 1 and not depending on N. Both of these are properties of the correlation coefficient, but that measure cannot be computed for dichotomous variables. For dichotomous variables in psychological experiments, Pearson's mean square coefficient of contingency

$$\phi = \sqrt{\chi^2/N} = \sqrt{\frac{(ad-bc)^2}{(a+b)(c+d)(a+c)(b+d)}}$$

is recommended by Gibbon et al. (1974), while the difference in the conditional probabilities of the US,

$$\Delta P = p(US|CS) - p(US|\sim CS) = \frac{a}{a+b} - \frac{c}{c+d}$$

has been used extensively in studies of human contingency and causality judgment (see, e.g., Allan et al. 2008).

Table-based measures are, however, problematic when applied to instrumental and Pavlovian conditioning experiments, which do not reliably have a definable trial structure (Gallistel and Gibbon 2000). This is apparent when one considers how to construct the contingency table for Rescorla's experiment. In that experiment, the CS always lasted 2 min. The interval between CS varied around an average of 10 min. There is no doubt about how many CS and US there were, so the first cell (a in Table 1) is readily determined. All the other cells are problematic, because there is no objectively justifiable answer to the question: How many not-US and how many not-CS were there? The values of contingency underlying Fig. 1 were obtained by following the common practice of assuming that the intervals between CS presentations are composed of "trials" of 2-min durations each, during which a US either occurs or does not. The number of ∼CS is taken to be the number of such arbitrary subdivisions. The number of ∼US is the total number of 2-min intervals, including those when the CS was present, minus the number in which a US occurred. However, the 2-min "trials" during the intervals between CS are a fiction, as is the number of not-US. Absent objectively defined trials, not-US and not-CS have no objectively definable relative frequency, so one cannot construct a contingency table. This problem is acute in the instrumental conditioning case, because there are no trials in those protocols.

A second problem with measures based on a contingency table, and with the correlation coefficient as well, is that they take no account of time. The contingencies of ordinary experience are defined over time, and the temporal intervals between the events are centrally relevant to the psychological perception of contingency and causality. The importance of "close" temporal pairing – of response and reinforcer, or of CS and US – has always been stressed in the conditioning literature. However, attempts to specify what constitutes "close" have never succeeded. Clearly, a psychologically useful measure of contingency must take time into account.

A measure that does this is the uncertainty coefficient, also known as the entropy coefficient. It is the percent reduction in uncertainty about when (or whether) a predicted event (US) will occur gained from knowledge of the times at which (or trials on which) the predictor event (CS) occurred:

$$UC = I(CS;US)/H(US). \quad (1)$$

$I(CS;US)$ is the mutual information between CS and US. $H(US)$ is the US entropy of the US distribution. It is also called the amount of "available" or "source" information. It is the information-theoretic measure of the uncertainty regarding when and/or whether a US will occur. In the case of atemporal dichotomous variables, where there are objectively definable trials, hence objective probabilities for the failure of a US to occur,

$$H(US) = \sum p_i \log(1/p_i) = p(US)\ln(1/p(US)) + p(\sim US)\ln(1/p(\sim US)). \quad (2)$$

The $\ln(1/p_i)$ is the amount of information provided by the occurrence of the ith event in the set of possible events over which a probability distribution is defined (e.g., the US and ∼US events). It is also called the surprisal. Intuitively, the less probable the event, the more unexpected or surprising it is, the more we are informed by its occurrence – but, by the same token, the less often we are so informed. As may be seen from Eq. 2, the entropy of a distribution, ~~H, of a distribution~~ is simply the average surprisal, that is, the amount of information provided by each of the possible events weighted by its relative frequency.

Entropy is the technical term for the amount of uncertainty in a probability distribution, which is the same as the amount of information available from that distribution, because information reduces uncertainty. The mutual information between two events with observed or experimenter-defined probability distributions is:

$$I(CS; US) = H(CS) + H(US) - H(CS, US),$$

where $H(US, CS)$ is the entropy of the joint CS–US distribution. In the case where a contingency table can be constructed, the US distribution is given by the normalized column totals, that is, the column totals in Table 1 divided by N; the CS distribution is given by the normalized row totals; and the joint distribution is given by the normalized cell values (a/N, b/N, c/N, d/N). For each distribution, the entropy is: $H = \sum p_i \ln(1/p_i)$.

The UC measure applies to temporal uncertainty as well (Balsam and Gallistel 2009). If US (or reinforcers) occur at random times, then the uncertainty regarding when the next US will occur is the entropy of an exponential distribution, which depends only on the average US–US interval (the reciprocal of the base rate). This entropy is the basal uncertainty about when the next US will occur. It is the amount of available information. If a CS always precedes a US and always tells us exactly when to expect the US, then there is no residual *objective* uncertainty about when the next US will occur once the CS has occurred. In that case, the UC is 1, that is, the CS reduces the uncertainty about when the next US will occur by 100%. However, humans and other common laboratory animals can only estimate the duration of an elapsing interval with about +/− 15% accuracy. To be useful, the CS must precede the US by some interval. Our residual uncertainty about when exactly to expect the US is then determined by our imprecision in estimating when the remembered CS–US interval has elapsed. Thus, the effective percent reduction in our uncertainty depends on the ratio between the basal interevent interval (the average US–US interval) and the CS–US interval (the delay of reinforcement). The greater this ratio is, the greater the percent reduction in our uncertainty. Thus, this way of measuring contingency explains why "close" temporal pairing is important. However, "close" is relative (to the basal interevent interval), not absolute; there is no critical interval that defines whether two events are or are not temporally paired.

Important Scientific Research and Open Questions

The UC measure of contingency provides a rationale for the two ideas in the famous Rescorla–Wagner model of association formation: $\Delta V = \alpha(\lambda - \sum V)$, where V is associative strength. This formula rests on two assumptions: (1) The sum across all the associations from different CS to one US cannot exceed some limit, which is represented by the asymptote parameter, λ. (2) Associative strengths are additive; their sum is subtracted from λ in determining the amount by which any associative strength is to be incremented, ΔV. The entropy of the US distribution, which determines the amount of available information, puts an upper limit on the amount of information that all predictors combined can provide. Moreover, the entropies of independent events (and independent conditional entropies) are additive. An open question is how far this can take us in understanding the objective basis for the phenomena of cue competition (blocking, overshadowing, relative validity) – see Balsam and Gallistel (2009).

Another open question is whether and how the brain can compute the uncertainties on which the UC measure of contingency depends.

Cross-References

- ▶ Association Learning
- ▶ Associationism
- ▶ Bayesian Learning
- ▶ Communication Theory
- ▶ Connectionist Theories of Learning
- ▶ Formal Learning Theory
- ▶ Human Contingency Learning
- ▶ Law of Effect
- ▶ Reinforcement Learning
- ▶ Temporal Learning in Humans and Other Animals

References

Allan, L. G., Hannah, S. D., Crump, M. J., & Siegel, S. (2008). The psychophysics of contingency assessment. *Journal of Experimental Psychology: General, 137*(2), 226–243.

Balsam, P., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences, 32*(2), 73–78.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review, 107*(2), 289–344.

Gibbon, J., Berryman, R., & Thompson, R. L. (1974). Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior, 21*(3), 585–605.

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), 1–5.

Contingency in Learning. Table 1 2×2 contingency table

	#US	#~US	Row totals
#CS:	*a*	*b*	*a*+*b*
#~CS:	*c*	*d*	*c*+*d*
Col totals:	*a*+*c*	*b*+*d*	

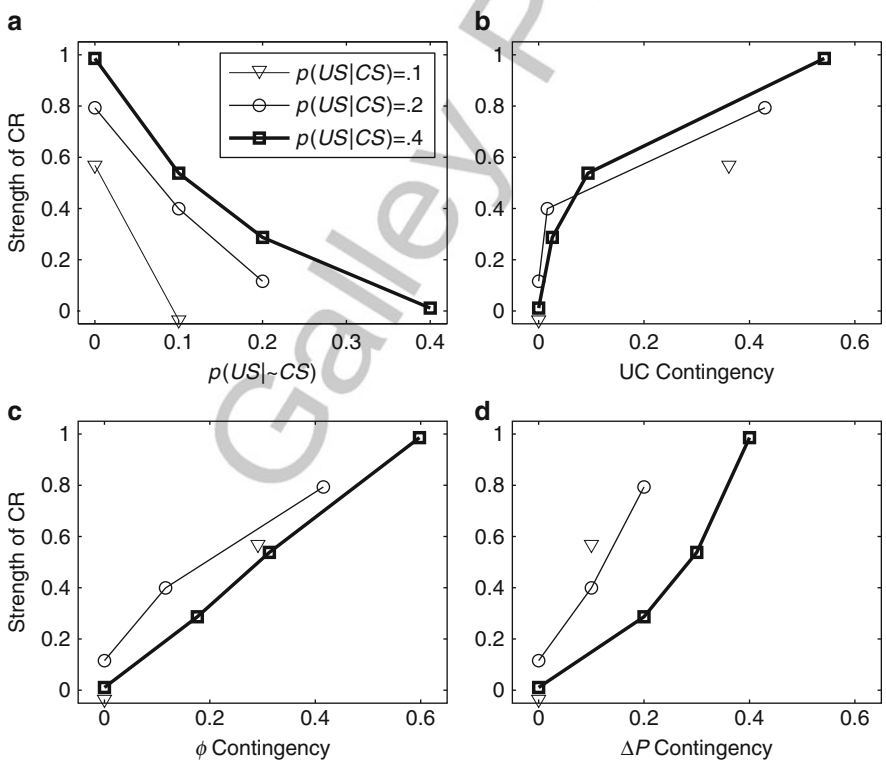

Contingency in Learning. Fig. 1 (**a**). The strength of the CR on first test trial as a function of *p(US|CS)* and *p(US|~CS)* in Rescorla's (1968) experiment on the role of CS–US contingency as against temporal pairing. Although in each of the three conditions, the temporal pairing of US and CS [hence *p(US|CS)*] was held constant, the strength of the CR declined to zero as the contingency was degraded by increasing *p(US|~CS)*. (**b**) Performance data in **a** plotted against the uncertainty coefficient (UC) measure of contingency. (**c**) Performance data in **a** plotted against the ϕ measure of contingency. (**d**) Performance data in **a** plotted against the ΔP measure of contingency

Author Query Form

Encyclopedia of the Sciences of Learning
Chapter No: 135

Query Refs.	Details Required	Author's response
AU1	Please check sentence starting "A broadly useable measure…" for completeness.	it was incomplete because improperly punctuated; I have corrected it