

Please note

The text in this file has been automatically extracted and may contain minor errors. For the original version please consult the paper copy held in the Swinburne Library.

7

STM LEARNING RULES

OVERVIEW	7-1
INTRODUCTION	7-2
EXTINCTION	7-4
REINSTATEMENT	7-7
REINFORCEMENT	7-9
DELAYED AND TRACE CONDITIONING	7-11
ADAPTIVELY REGULATED MAXIMUM ALTM LEVELS	7-14
SUMMARY	7-16

OVERVIEW

This chapter develops the remaining mechanisms required to complete each CS input channel of ACE. Since **all** other aspects of ACE have been developed in the **previous** chapters, this represents the final **stage** in the development of ACE. The emphasis here is upon the qualitative development of the mechanisms which bind together the separately developed subsystems of ACE so that they may **functionally** complement one another. Most of these mechanisms determine how the associative STM in each CS input channel is directly affected by experience, and so are collectively referred to as **the** STM learning rules. Note however, that because of the highly interactive internal **nature** of ACE, these learning rules by no means fully encapsulate the effect that experience has **upon the** strength of associative STM.

INTRODUCTION

The schematic diagram of **ACE** provided below (**Figure 7-1**) shows, with the exception of the internal detail of the **CSTC**, all mechanistic aspects of ACE. In an attempt to clearly explain how ACE functions, an effort was made to **modularise** its internal mechanisms so that they could be developed and discussed in relative isolation. Three **main** subsystems were identified as the Neural **Multiprocess** Memory Model (**NMMM**), the Conditioned Stimulus Trace Circuit (**CSTC**), and the output **stage**, which were discussed separately in Chapters 4, 5, and 6 respectively. All of the mechanisms previously developed in these chapters are shown shaded in **Figure 7-1** in order to highlight the remaining (**unshaded**) mechanisms which complete **ACE**. Each of these remaining mechanisms is addressed in this chapter, in a **mechanistically logical** sequence.

Starting at the upper left portion of the unshaded region of Figure 7-1, and traversing around in an essentially anticlockwise direction, these new mechanisms may be outlined as follows. First, a negative feedback loop is used to make Memory Gating Short Term Memory (**MGM**) value increase towards that of the newly introduced cumulative quantity referred to as Reinforcement Gating Short Term Memory (**RGM**). An increase in **MGM** is required so that the results of recent experience temporarily retained in **STM** may be allowed to transfer via **MTM** to **LTM** so that consolidation may occur.

RGM attempts to follow a Compressed Conditioned Stimulus Trace (**CCST**) of prior CS input activation by another simple negative feedback loop, but has a rate of increase limited to reduce the effectiveness of reinforcement when the Inter-Stimulus Interval (**ISI**) between CS and US onset is very short. **RGM** modulates the reinforcing effect of subsequent US presentation to produce the Reinforcing signal that increases **STM**, which is labelled "R" in Figure 7-1.

Extinction of conditioned excitors is supported by an Extinction signal, labelled "E" in the Figure 7-1. This extinction signal acts to decrease **STM** in response to the compound expectation of reinforcement, which is determined by the output stage.

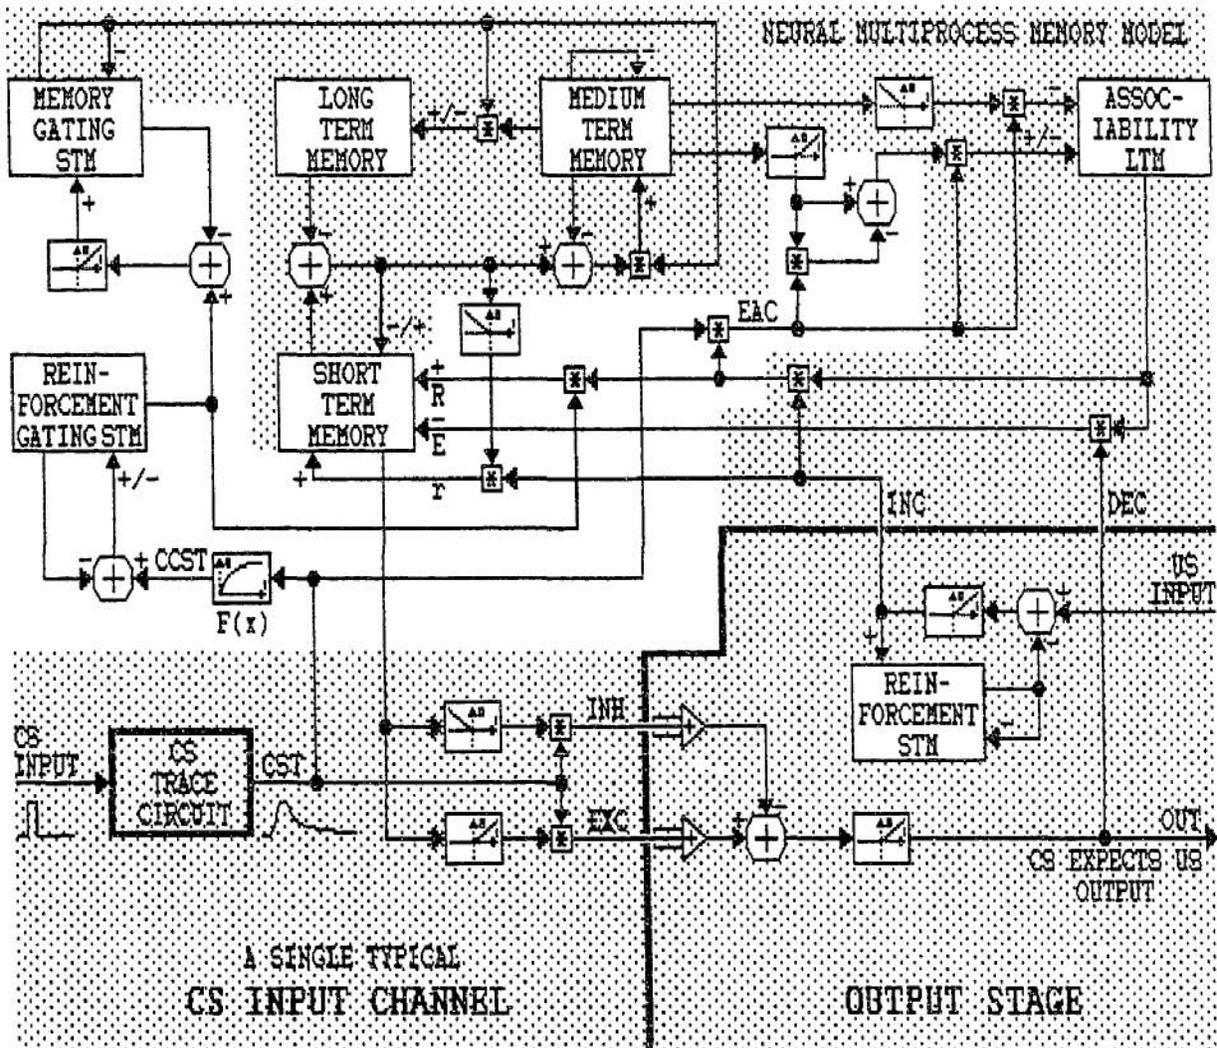


FIGURE 7-1. Schematic diagram of ACE, showing a complete typical CS input channel, and the single output stage. Those mechanisms shown shaded have already been developed and documented in previous chapters, and include the NMMM, the CSTC, and the output stage. Each unshaded mechanism is progressively developed in this chapter.

A special reinstatement signal, labelled "r" in Figure 7-1, is also included to restore STM level back up towards its previous long term value when, following extinction, unpaired US presentations are experienced. Finally, a special mechanism located in the upper right portion of the unshaded region in Figure 7-1 utilises the uncompressed CST signal and the US driven reinforcing feedback signal INC to adaptively regulate the maximum allowable Associability Long Term Memory (ALTM) level.

EXTINCTION

Acquisition of a classical CR is usually achieved over many successive trials in which a CS and then a US is presented, producing a gradual increase in the strength and/or probability of the CR. If after acquisition, the CS is presented alone on each trial, without any subsequent presentation of the US, then the strength of the CR will gradually decline. This procedure, and the result it produces, are both referred to as "extinction". In other words, an extinction procedure leads to the extinction of the CR.

The signal responsible for producing reductions in synaptic STM, and hence extinction of **the** CR, is identified with an "E" in **Figure 7-1**. A positive STM value, previously established by acquisition training, enables a CR to be generated when the CS input is activated, by producing nonzero EXC and OUT signals, leading to generation of the CR. The OUT signal is also relabelled as DEC and fed back to **the** CS input channel, where after being modulated by the current level of associability for that CS input channel (which is retained in ALTM), it acts to decrease SIM value. Modulation of DEC by the current ALTM level enables previous experience to determine the most appropriate rate of extinction, via the adaptive associability mechanism introduced in Chapter 4. This sequence of events ensures that the strength of a CR diminishes towards zero with each nonreinforced CS presentation, thereby supporting basic extinction behavior.

Note that this extinction mechanism is activated whenever a nonzero OUT signal is generated by **activation** of CS inputs, in whatever context, be it extinction or acquisition training. In other words, this extinction mechanism does not identify when extinction training is being experienced, and then come into action. Instead, it operates whenever positive **compound** expectation of reinforcement is conveyed by DEC to the CS input channels, and relies upon the operation of a separate reinforcement mechanism (discussed below) to counter this induced decrease in STM value **when** effective reinforcement is experienced. This basic type of nondiscriminatory **operation** of an extinction mechanism is not new, being conceptually similar to the reactive inhibition

theory usually associated with Hull's name (Mull, 1943, Chapter 16), and the operation of the Rescorla–Wagner model (Rescorla and Wagner, 1972).

However, the extinction mechanism depicted in Figure 7-1 differs radically from all others in the following very basic way - the extinguishing effect of a positive compound expectation of the US is **not** gated by prior CS input activity. This means that presentation of a CS previously trained to expect a US will reduce the **STM** level of not only its **own** CS input channel, **but** will reduce the STM levels in **all** of the CS input channels within ACE. The extent to which the **CRs** produced by each CS input channel are reduced will depend upon the initial **STM** level associated with **the** presented CS, the extent to which it is extinguished, **the** initial STM levels of each CS input **channel** relative to that of the presented CS, and their ALTM levels.

While such behavior seems to have received very little attention in recent times, Pavlov (1927, pp. 54–56) noted behavior consistent with this, referring to extinction of the presented CS as "primary extinction", and that of other CSs generating a similar (or "homogeneous") CR as "secondary extinction". Pavlov found that:

- (i) The extent of **secondary** extinction increases with the extent of primary extinction.
- (ii) **The** weaker the CR generated by **the** primarily extinguished CS compared to that of the secondarily extinguished CSs, the less complete is the secondary extinction. Conversely, "if the stronger CR is subjected to primary experimental extinction the weaker conditioned reflex undergoes complete secondary extinction" (**Pavlov**, 1927, pp. 55).
- (iii) **The same** relationship between extent of secondary extinction and relative CR strength described in (ii) also applies to components of compound **CSs**, when only a component CS is subject to primary extinction, and the remainder undergo secondary extinction.

The above empirically observed relationship between primary and secondary extinction is in accord with that expected from the extinction mechanism implemented within ACE, which it should be noted, is entirely different to that proposed by Pavlov (1927, p. 60). However, an effort needs to be made to

explain how secondary extinction can be considered as appropriate behavior. On the surface, it hardly seems appropriate to extinguish all of the CS-US associations leading to the production of a single CR, because one of those CSs which used to predict reinforcement no longer does. In fact, when considered in isolation, this behavior would **appear** to result in the inappropriate loss of many CS-US associations which would most probably remain useful, and which were acquired after a great **deal** of experience. However, a case will now be made that, when viewed in a specific wider context, secondary extinction behavior may no longer appear inappropriate, and in fact that it may considerably enhance behavior.

Pavlov (1927, p. 58) somewhat boldly stated **that** "all those conditioned reflexes which **have** been fully established invariably **and** spontaneously return **sooner** or later to their full strength" following experimental extinction. Although it is unclear, Pavlov seems to be referring to the spontaneous recovery of primarily extinguished CRs, though he may **also** be including secondarily extinguished CRs. More recent results indicate that such complete spontaneous recovery from primary extinction seems to be the exception, with something like 50% being more typical (Kimble, 1961, p. 284). Empirical results specifically addressing the extent of spontaneous recovery from secondary extinction are apparently unavailable.

The partial spontaneous recovery of a **primarily** extinguished CS is already supported by the NMMM (Chapter 4). The suggestion here, which is claimed to be new and unique, is that unlike primary extinction, secondary extinction routinely undergoes complete spontaneous recovery. This eliminates the inappropriate **long term** secondary extinction of CS-US associations **that** are not actively extinguished. Complete spontaneous recovery of secondarily extinguished CSs occurs because changes in STM are only **able** to be transferred to MTM and LTM when **MGM value** is increased by CS presentation, and STM value spontaneously changes towards the value of **LTM** (Figure 7-1). Although **MGM was** introduced to gate changes in MIM and LIM primarily in preparation for the use of MTM as a suitable **source to** drive changes in adaptive **associability**, MGM also has the effect of isolating STM from both MTM **and** ETM.

Potential behavioral advantages of secondary extinction emerge because STM is now available for use as a temporally sensitive register of ongoing US availability. When the immediate effect of the unexpected omission of an otherwise expected US presentation is not specific to a particular CS, then all of the previously reinforced CS components presented in this temporal context will assist one another to promptly extinguish. This capability takes on a whole new dimension if a new CS (not previously reinforced) is presented only in such a temporal context of unexpected US omission. The SIM level of its corresponding CS input channel will not only be driven to attain negative values - its presentation will also ensure that negative LTM values are attained. In other words, the **development** of conditioned inhibitors is now **also** supported by ACE in successive discrimination procedures. This represents a substantial extension to the simultaneous discrimination capability that most models and theories of associative conditioning are content to achieve.

REINSTATEMENT

When Pavlov (1927, p. 59) noted that the response to an extinguished CS could be largely restored by simply presenting the original US on its own, he described an effect now referred to as reinstatement. Later experiments confirmed the effect (Konorski, 1948, p. 185; Rescorla, 1979), and showed that **the** US is similarly effective whether signalled by another CS or unsignalled, and that its reinstatement effect remains substantial 24 hours after US presentation (Rescorla & Heth, 1975). Rescorla and Heth (1975) also argued that reinstatement is not specific to **CS-US** associations, but instead appears to act upon a US **representation**. This view of reinstatement received further support from Rescorla and Cunningham (1977), who showed that extinction of a second CS is apparently capable of "erasing" the prior reinstatement effect of the US presentation upon the previously extinguished CS, when both CSs are associated with the same US.

A reinstatement effect with the above characteristics can be achieved by enhancing the NMMM with a simple mechanism that generates a reinstatement signal, labelled "r" in Figure 7-1. This reinstatement signal is formed by the

product of the INC feedback signal from the output stage, which is activated primarily by onset of a US presentation, and a rectified version of the amount by which LTM value exceeds STM value. This means that if, as discussed above, recent extinction of a CS has depressed STM value below that of LTM, then an unpaired US presentation will act to rapidly increase SIM value towards that of LTM. Thus, an unpaired US presentation will reinstate previously masked excitatory strength which is still retained in the LTM of each CS input channel.

The extent to which such reinstatement is complete will depend upon the degree to which LTM value has been reduced by the extinguished STM values before the US is presented. This in turn depends upon the elapsed time since the beginning of extinction training, and the extent of SIM extinction. This is consistent with Pavlov's observation that "If the extinction has not been carried very far, a single application of the unconditioned stimulus is often sufficient to restore the reflex to full strength; but if the extinction has been made profound, repeated reinforcements are necessary" (Pavlov, 1927, p. 59).

Even though the associative memory in each CS input channel which mediates this reinstatement effect is specific to both a CS and a US, the resulting reinstatement effect is behaviorally nonassociative. This is because neither the decreases in STM resulting from extinction produced by CS alone presentations, nor the restorative increases in SIM resulting from US alone presentation, require contiguous CS and US presentation.

There are, however, good reasons for utilising the associative NMMM in such a way as to produce an apparently nonassociative reinstatement effect. To begin with, the mechanism mediating nonassociative reductions in STM, which has already been implemented to support extinction, provides at least half of the mechanism required to support reinstatement behavior. The combined short and long term retention characteristics of the NMMM also become an integral part of reinstatement behavior, providing a ready explanation for the dependence upon prior extinction, and the longevity of the reinstatement effect (Rescorla and Heth, 1975). The apparent interaction between reinstatement and spontaneous recovery (Rescorla, 1979), and the observation

that extinction after reinstatement makes subsequent retraining of the CS more difficult (Rescorla and Cunningham, 1977), further implicates the NMMM because of its production of spontaneous recovery behavior, and its pivotal role in the mediation of learning (Chapter 4).

The above reinstatement mechanism also plays an important role when a Partial Reinforcement (PR) schedule is experienced. The effect of Nonreinforced (N) trials to reduce STM level may be largely negated when subsequent Reinforced (R) trials occur, because STM is rapidly elevated back up towards the level of LTM by the US presentation in each R trial. This can substantially reduce the overall impact of the N trials upon the asymptotic strength of association able to be achieved in acquisition, resulting in a higher SIM asymptote than would be expected from the scheduled percentage of R trials. This specific type of behavior is entirely consistent with empirical results (e.g., Gibbs, Latham, and Gormezano, 1978).

REINFORCEMENT

In its most general sense, reinforcement refers to the experimental operation of arranging outcomes contingent upon events to increase the strength and/or probability of a type of behavior. In the context of classical conditioning, the arranged outcome is usually the presentation of a US that is made contingent upon the prior presentation of a CS, and which has the effect of increasing the probability that a CR will be generated. More specifically, reinforcement leads to acquisition of an associative relationship, which in the case of ACE is between a CS input and the US input, that is subsequently available to generate the CR when the CS is presented. Expressed yet **another** way, reinforcement results in ACE acquiring a predictive relationship between a CS input and **the** US input.

Acquisition of this CS-US association requires some type of associative mechanism **that** is capable of selectively producing a persistent increase in the STM level of those CS input channels that were **previously** activated by the CS, when the US is subsequently presented. The associative effect of acquisition **due** to reinforcement is mediated by the Reinforcement signal,

labelled "R" in Figure 7-1. The enduring effect of the increase in SIM produced by this R signal is achieved because prior CS input activation also enables transfer of changes in SIM to MIM and then LTM, via the gating effect of MGM (Chapter 4). The R signal produces increases in SIM only when both of the following conditions are met:

- (i) The corresponding CS input has been activated prior to US presentation.
- (ii) The US is not already fully predicted by the CS input channels.

The mechanism supporting condition (i) utilises a processed version of the CS Trace (CST) signal generated by the CSTC (Chapter 5) to enable increases in SIM when the CS is reinforced by subsequent US presentation. It might be **possible**, depending upon **the** specific application, to directly use the CST to gate the reinforcing effect of paired CS-US presentations upon STM. However, since the reinforcement mechanism also determines precisely how acquisition is affected by ISI, and the use of delayed or trace conditioning procedures, it was considered prudent to provide the facility to separately tailor the effects of both of these operational variables upon acquisition. The mechanism provided for this purpose is detailed below.

Condition (ii) above is supported by the implementation of a modified intratrial version of the Rescorla-Wagner model (Rescorla and Wagner, 1972). The **output** stage determines the compound expectation of reinforcement, which via the "extinction" mechanism discussed above, produces decreases in SIM. The output stage also determines the asymptotic strength of association able to be supported by the US, which via this reinforcement mechanism, selectively produces increases in SIM. SIM therefore becomes the site at which compound expectation of reinforcement is compared with actual reinforcement delivered, in order to determine how the individual CS-US associations maintained in the SIM and LTM of each CS input channel should change. When expectation matches the actual reinforcement (or nonreinforcement) subsequently experienced, the effect of the extinction and reinforcement mechanisms upon **STM** is equal but opposite, resulting in no net change to STM value.

Note that the effect of this reinforcement mechanism upon SIM is also gated by the current ALIM level, Since ALTM also gates the effect of the extinction

mechanism upon STM, ALTM is able to modulate the rate of learning without affecting the asymptotic levels resulting from acquisition or extinction training.

DELAYED AND TRACE CONDITIONING

Delayed and trace conditioning procedures define the two main types of temporal relationship between CS and US presentation in classical conditioning. Figure 7-2 illustrates how in delayed conditioning, the CS is sustained until the US is presented. Normally both the CS and the US terminate together, though in some delayed conditioning procedures the CS may terminate at US onset, or alternatively some time after US offset. In contrast, trace conditioning defines procedures in which the CS terminates before US onset, leaving a distinct interval of time between CS and US presentation. This interval is sometimes referred to as the "trace interval".

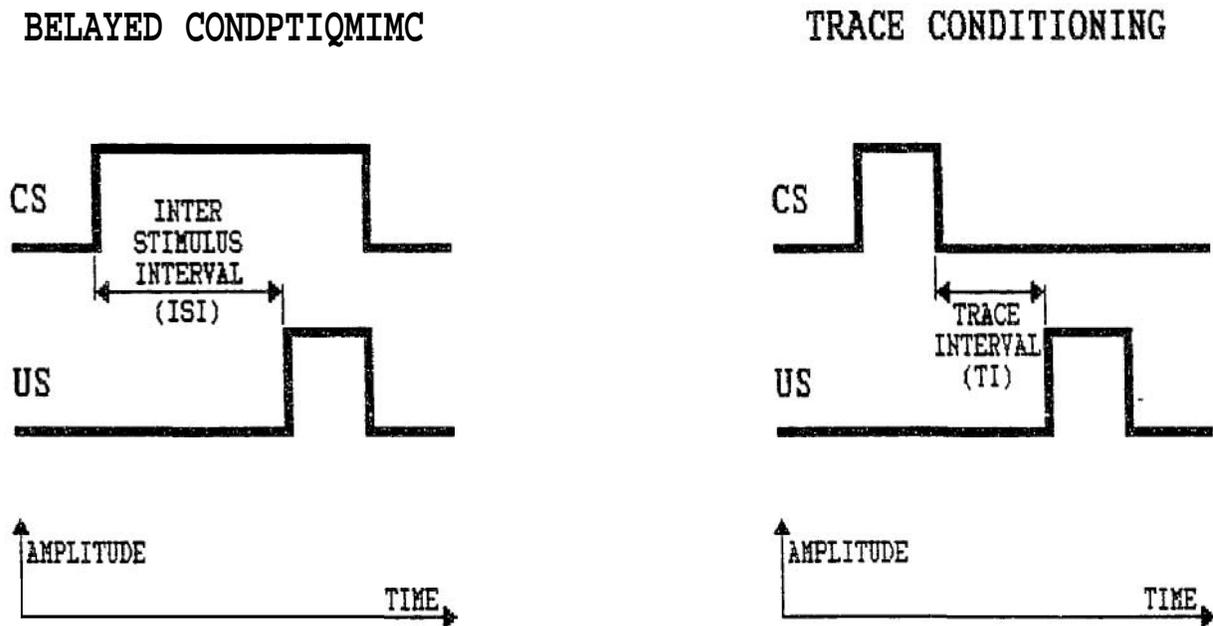


FIGURE 7-2. Temporal relationships between a CS and a US in delayed conditioning and trace conditioning procedures.

In the case of both delayed and trace conditioning, the ISI refers to the interval between CS onset and US onset. Trace conditioning is nearly always less effective than **delayed** conditioning, and becomes increasingly less effective as the trace interval is increased (Mackintosh, 1983, pp. 86-87).

The CSTC developed in Chapter 5 was specifically designed to produce a **CST** output signal **with** an overall amplitude that decreases as the CS terminates further in advance of the time at which the CST **peaks**. Although the specific **issue** then was **how** performance is **modulated** by CS duration, when this same **CST signal** is used in a **micromolar** (differential correlated reinforcement) mode, the rate and asymptote of acquisition **are** similarly affected by CS duration. **This** automatically leads to the **common** empirically observed **result** that delayed conditioning is more effective than **trace** conditioning.

However, the extent of the difference in acquisition between delayed **and** trace conditioning may not be as dramatic as **the** extent of the immediate performance differential **produced** by different CS durations. Furthermore, since in the case of ACE a derivative of the CST **is** used to modulate both increases in SIM due to reinforcement, and increases in MGM in order to gate STM changes through to MTM **and** LTM, and since both jointly determine the rate of acquisition, some additional processing is **also** required to prevent the acquisition rate **from** being excessively **modulated** by CS duration. A function is therefore used to **compress** the range of possible CST signals before it **is** used to modulate both of the above mechanisms. The relationship between the Compressed CST (CCST) signal (Figure 7-1) and **the** CST signal is defined as follows:

$$\text{CCST}[T, N] = F(\text{CST}[T, N]) \quad [7-1]$$

$$F(x) = x/(x + \text{CCSTcon}) \quad 17-21$$

Where: CCST[T, N] = Compressed CS Trace of channel N at time state T.
 CST[T, N] = CS Trace **for** CS input channel N **at** time state T.
 F(x) □ Compression function.
 CCSTcon = Compression function constant, with value of 0.2.

It is apparent that the compression function $F(x)$ defined by Equation [7-2] equals 0.0 when $x = 0.0$, and approaches unity as x approaches infinity. When $x = CCST_{con} = 0.2$, then $F(x) = 0.5$. The most important aspect of $F(x)$ is that it is a monotonically increasing function of x , and so maintains the relative relationship between the amplitudes of the spectrum of CST signals hypothetically generated by many CS input channels. This ensures that the CST signal which most precisely coincides with US presentation still produces the most rapid rate of SIM acquisition, attains the greatest single share of associative strength, and tends to dominate the overall CR topography - which will therefore tend to peak shortly after US onset.

A final cumulative quantity, hereafter referred to as Reinforcement Gating Short Term Memory (RGM), is introduced to provide a CSTC-independent means of tailoring how the extent of acquisition varies with ISI. Figure 7-1 indicates that RGM attempts to follow the CCST signal via the operation of a simple negative feedback loop. However, its rate of increase is intentionally limited so that it is unable to track those short-peaking signals which rapidly rise and then rapidly decay. Consequently, RGM further reduces the rate and extent of acquisition at short ISIs, but has progressively less effect as the ISI increases. The rate of RGM depletion is set to unity, enabling RGM to accurately track the falling edge of each CCST signal. This minimises disruption to the differential correlated reinforcement mechanism, which supports acquisition of that CR topography which coincides approximately with US onset, at all but the shortest ISIs.

RGM is also used to increase MGM, because it is suitably compressed in terms of its variation in amplitude, but is still appropriately responsive to CS duration and amplitude. However, the choice as to which signal drives increases in MGM is one of the less critical aspects of ACE's design.

ADAPTIVELY REGULATED MAXIMUM ALTM LEVELS

Having established above how STM contents are affected by experience, a consideration of the different asymptotic STM limits produced by delayed and trace conditioning procedures leads to the apparent necessity for some means to appropriately regulate maximum ALTM levels. Although the mechanism responsible for this is more directly associated with the adaptive associability mechanism (**Chapter 4**) than the STM learning rules **described** above, it is included here because of its dependence **upon**: the CST output signal from the **CSTC** (**Chapter 5**), the potentially reinforcing effect of US presentation that is mediated by the **INC** feedback signal (**Chapter 6**), and the interaction between the effect experience has **upon STM** and the maximum appropriate **ALTM** level (discussed above).

A given **US** of fixed amplitude and duration, presented at a fixed **ISI**, will support a fixed maximum compound expectation of reinforcement within **ACE**. **When the** duration and amplitude of the **CS** is also fixed, the fixed expectation of reinforcement translates into a fixed asymptotic **STM** limit. However, if the **CS** duration is set to less than the **ISI** to produce a trace conditioning procedure, then the asymptotic **STM** limit will be different to that when the **CS** is extended to produce a delayed conditioning procedure.

More specifically, a trace **CS** produces a small **CST**, and leads to a large **STM** asymptotic limit in order to produce the **total** expectation of **reinforcement** supported by the **US** presentation. Conversely, a delayed **CS** produces a large **CST**, and **leads** to a **small STM** asymptotic limit. If **ALTM** **can** attain only a single fixed **maximum** level, then at this **maximum ALTM** level the absolute **change** in expectation of **reinforcement** due to each reinforced and nonreinforced trial for trace and delayed conditioning will be identical, but the relative change will be dramatically different. This is because the extinction and reinforcement signals, labelled "E" and "R" respectively in **Figure 7-1**, will have a much greater impact upon the **smaller STM** levels achieved in delayed conditioning than upon the larger **STM** levels resulting **from** trace conditioning. **The** specific settings used in the **CSTC** **may**, for

example, lead to **differences** in CST amplitude, and consequently **STM** asymptotic limit, of approximately one order of magnitude (a factor of 10).

While it is normal for delayed conditioning to proceed more rapidly than trace conditioning, the difference in rate tends to be much less than an order of magnitude. Furthermore, most of this difference may be accounted for by assuming that naive subjects usually begin with similarly low initial **ALTM** levels. This automatically produces a more rapid initial learning rate for subjects on a delayed conditioning procedure. Furthermore, this effect is compounded by a more rapid and complete increase in **ALTM** level.

A simple **mechanism** which adaptively regulates the maximum **ALTM** level to maintain a similar maximum learning rate for both trace and delayed conditioning is illustrated in Figure 7-1. The Enable Associability Change (**EAC**) signal, as its name suggests, enables changes in **ALTM** to occur when reinforcement is experienced. **EAC** is formed by the product of the **US** feedback signal **INC**, the current **ALTM** level, and the uncompressed **CST** signal.

Direct use of the **CST** signal means that **EAC** is affected by trace and delayed **CS** inputs in the same **way** as the "CS expects US" output. A special new negative feedback loop, illustrated above **EAC** in Figure 7-1, reduces the extent of the increase in **ALTM** as **EAC** becomes increasingly larger, when **MTM** value **is** positive. **When** $EAC = 1$, no **additional** increase in **ALTM** is possible. If **EAC** exceeds unity amplitude, a negative **increase** (ie; a decrease) in **ALTM** actually results, because the signal **which** normally acts to increase **AETM** now attains a negative value. That this condition is permitted is indicated in Figure 7-1 by **the** use of a "+/-" **label** on the **signal** as it impacts upon **ALTM**. This additional capacity to decrease **ALTM** level back down to a **new** Lower maximum asymptotic **level** permits the **maintenance** of appropriate maximum **ALTM** levels, despite the possibility of arbitrary **transitions** between trace and delayed conditioning procedures. Furthermore, the fact that this can occur when only positive **MTM** values are present, which for **example** is the case during acquisition, ensures that appropriate adjustment of the **maximum** **ALTM** level will occur even during conditions **which** normally increase **ALTM** level.

A mechanism which sets a maximum, albeit **adaptively** regulated, **ALTM** level can **also** ensure that unstable operation during acquisition does not occur. This undesirable condition may otherwise arise **as** increases in **AETM** enable larger increases in **STM**, which produce **larger** **MTM** values, and **which** in turn enable progressively **larger** increases in **ALTM**. This represents one **of** the few instances in which a special effort needs to be made to avoid unstable operation within ACE. The need **arises** because of the positive feedback relationship that exists between **increases** in **ALTM**, and increases in **STM**. Stable operation within the **other** internal mechanisms of ACE is **assisted** by the extensive, indeed almost exclusive, use of simple negative feedback loops.

SUMMARY

A qualitative description of the mechanisms that constitute ACE has **now** been completed by focussing upon those mechanisms which **bind** together all of the subsystems previously developed in Chapters 4, 5, and 6. These remaining mechanisms are collectively referred here as the **STM** learning rules because they describe primarily how **experience** induces relatively rapid changes in associative **STM**. They thus **bear** a superficial **resemblance** to the "learning rules" of ANN terminology. However, because they are embedded within so many other mechanisms which contribute to their proper and complete operation, the mechanisms described in this chapter really only form part of **what** might be regarded as the set of learning rules for ACE.

Most of the mechanisms were identified in terms **of** their primary functional role. A new extinction mechanism was introduced which, **because** of the selective consolidation behavior **already** supported by the NMMM, was able **to** be mechanistically exceptionally simple. This extinction mechanism supports **not** **only** primary extinction of an excitatory CS **when** presented alone, but also secondary extinction of **all** other CS-US associations to the same US. However, while primary extinction is subject to partial spontaneous recovery, the new type of secondary extinction supported here undergoes full spontaneous recovery, thus preventing the inappropriate long **term** loss **of** such a widespread set of associations. However, **in** the short term, **such** secondary

extinction considerably enhances adaptive behavior by facilitating rapid extinction of many **CSs** when presented in a context of unexpected omission of the US, and by supporting a successive discrimination capability.

A new reinstatement mechanism **was** also introduced to directly support the nonassociative restoration of conditioned **exciters** when, shortly following extinction, the US is presented alone. Like the secondary extinction phenomena, reinstatement is an **empirically** observed phenomenon of considerable, though perhaps not readily apparent, behavioral utility. Aside from helping **to** prevent the possibly inappropriate **loss** of excitation during temporary **periods** of unexpected US omission, the **reinstatement mechanism also helps** maintain a **stronger CS-US** association during **partial** reinforcement schedules.

A relatively standard micromolar reinforcement mechanism was then introduced, with some enhancements made **to** reduce the difference between the reinforcing effect of trace and delayed conditioning procedures, and **to** facilitate **isolated tailoring** of the reinforcing effect of US presentations at short **ISIs**. A **micromolar**, or differential correlated reinforcement, mechanism was used to directly support acquisition of appropriately timed and generalised **CRs**.

Finally, a new mechanism **to** adaptively regulate **maximum ALTM levels** was developed to ensure stable and appropriate operation of the adaptive associability mechanism, even when arbitrary transitions between trace and delayed conditioning procedures are experienced.