

**Comparing Adaptive Control of Thought–Rational (ACT-R)
Baseline Activation Terms for Implementation in the
Symbolic and Subsymbolic Robotic Intelligence
Control System (SS-RICS)**

by Craig T. Lennon

ARL-TN-601

April 2014

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-TN-601

April 2014

Comparing Adaptive Control of Thought–Rational (ACT-R) Baseline Activation Terms for Implementation in the Symbolic and Subsymbolic Robotic Intelligence Control System (SS-RICS)

Craig T. Lennon
Vehicle Technology Directorate, ARL

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
April 2014		Final		July 2013–January 2014	
4. TITLE AND SUBTITLE Comparing Adaptive Control of Thought–Rational (ACT-R) Baseline Activation Terms for Implementation in the Symbolic and Subsymbolic Robotic Intelligence Control System (SS-RICS)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Craig T. Lennon				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: RDRL-VTA Aberdeen Proving Ground, MD 21005				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TN-601	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The U.S. Army Research Laboratory is developing the Symbolic and Subsymbolic Robotic Intelligence Control System (SS-RICS) as a robotic control system based on human cognition. It has, as part of its intelligence, a production system similar to that of Adaptive Control of Thought–Rational (ACT-R). One feature of the ACT-R system is that chunks of information about the surrounding environment have recorded activation values, which increase when the system turns its attention toward them and decrease in the absence of attention. This feature can be beneficial in keeping recently used information readily available to attention. It can also be detrimental when a feedback loop develops, with a chunk being selected for attention because it has the highest activation and then increasing its activation value as a result. This report explores the behavior of the baseline activation equation with different parameter settings, and over different patterns of selection, in order to provide SS-RICS researchers with information for selecting appropriate parameters. It also makes recommendations as to how the equations might be modified to eliminate such feedback loops.					
15. SUBJECT TERMS SS-RICS, ACT-R, cognitive robotics, artificial intelligence, robotic control system					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Craig T. Lennon
Unclassified	Unclassified	Unclassified	UU	46	19b. TELEPHONE NUMBER (Include area code) 410-278-9886

Contents

List of Figures	iv
List of Tables	vi
1. Introduction	1
2. Defining Baseline Activation Terms	2
3. Comparing Baseline Activation Terms	6
3.1 Examples With Sparse Activation.....	7
3.2 Examples of Activation With Fixed Frequency	9
3.3 Examples of Sporadic Activation.....	10
3.4 Examples of Random Sequences	12
4. Conclusions	14
5. References	17
Appendix A. Baseline Activation Terms With Different Parameters	19
Appendix B. Replications of the Examples of Section 3.4 of the Report	25
Appendix C. Code for Reproducing Test Results	33
Distribution List	37

List of Figures

Figure 1. A generic graph demonstrating the concept of refraction (10).....	6
Figure 2. Sparse activation at time periods 3 and 13; equations 2 (red), 9 (blue), and 10 (purple).....	8
Figure 3. Sparse activation at time periods 3 and 23; equations 2 (red), 9 (blue), and 10 (purple).....	8
Figure 4. Plots of equations 2 (red), 9 (blue), and 10 (purple), with activation every 10 time periods.....	9
Figure 5. Plots of equations 2 (red), 9 (blue), and 10 (purple), with activation every 20 time periods.....	10
Figure 6. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30 activations over 150 time steps.	11
Figure 7. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30 activations over 100 time steps.	11
Figure 8. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30 activations over 200 time steps	12
Figure 9. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.	13
Figure 10. Plots of equations 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.	13
Figure 11. Plots of equation 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.	14
Figure A-1. A comparison of parameter values for equation 10 (see table A-1).....	20
Figure A-2. A comparison of parameter values for equation 2 (see table A-2).....	21
Figure A-3. A comparison of parameter values for equation 2 (see table A-3).....	22
Figure A-4. A comparison of parameter values for equation 9 (see table A-4).....	23
Figure A-5. A comparison of parameter values for equation 9 (see table A-5).....	24
Figure B-1. Replication 1 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.	26
Figure B-2. Replication 2 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.	26
Figure B-3. Replication 3 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.	27
Figure B-4. Replication 4 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.	27
Figure B-5. Replication 1 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple),	

randomly activated at each time step with probability 0.05.	28
Figure B-6. Replication 2 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.	28
Figure B-7. Replication 3 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.	29
Figure B-8. Replication 4 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.	29
Figure B-9. Replication 5 of equations 2 (red), 8 (orange), 9 (blue), and 2.10 (purple), randomly activated at each time step with probability 0.05.	30
Figure B-10. Replication 6 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.	30
Figure B-11. Replication 7 of equations 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.	31
Figure B-12. Replication 8 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.	31

List of Tables

Table 1. Baseline activation as a function of n.	4
Table 2. Parameter values and legend for all graphs in section 3.	7
Table A-1. Parameter values for figure A-1.	20
Table A-2. Parameter values for figure A-2.	21
Table A-3. Parameter values for figure A-3.	22
Table A-4. Parameter values for figure A-4.	23
Table A-5. Parameter values for figure A-5.	24

1. Introduction

The U.S. Army Research Laboratory is developing the Symbolic and Subsymbolic Robotic Intelligence Control System (SS-RICS) as a robotic control system based on human cognition (1–4). It has, as part of its intelligence, a production system similar to that of Adaptive Control of Thought–Rational (ACT-R) (5, 6) as well as statistical learning-based modules that process and filter input outside of the production system to reduce the computational load on that system (2). The production system of ACT-R matches “chunks” of information held in buffers (working memory) to achieve goals. Each chunk has a numeric value called activation associated with it, and among matching chunks the one with highest activation is selected. This process may be modified by making the selection process stochastic by randomly perturbing chunk activation values or by accepting partial matches that are close enough to a perfect match, but the essence of the process is selection of a high-activation chunk that constitutes a sufficient match. The ACT-R reference manual (5) contains greater detail. One feature of activation is that regular use of a chunk increases its activation value, which in turn can lead to the chunk being selected for use again. This can be beneficial under some circumstances but could also lead to a system becoming fixated on something in its environment, with this feature resulting in a feedback loop of ever-increasing activation values (7). The purpose of this report is to explore the behavior of ACT-R activation equations, with varying parameters, under varying degrees of stimulation. This will provide SS-RICS developers with better insight into the value of certain parameter choices and the behavior of different terms of the ACT-R baseline activation equation.

The ACT-R activation equation, which determines activation values, can conceptually be broken down in several ways, depending on the degree of detail desired. Using a breakdown that corresponds to the ACT-R reference manual (5, p. 218), we represent the activation of a chunk A as the sum of a baseline component B , a spreading component S , and components representing partial matching P and noise ε :

$$A = B + S + P + \varepsilon \tag{1}$$

The partial match value is used when perfect matching of chunks is not required, and indicates how close the match is. The noise value is a random perturbation of the activation level. The spreading activation value reflects the association of a chunk to other chunks in the buffers and represents the contribution of associative memory to recall efforts. The baseline term consists of a stable component, which is usually set by the modeler to reflect baseline strength, and a function whose value changes while the chunk is in working memory to reflect the recency and frequency of use. This function may include a refraction term that makes a chunk less likely to be used multiple times in quick succession.

Different baseline activation terms have been used with ACT-R (5, 8, 9). These baseline activation terms are developed to mimic the human thought process in terms of task performance. SS-RICS is intended for a robot. The intelligence of a robot should be close enough to that of a human to be understandable but different enough to complement, rather than replicate, human performance. For example, a human would not be expected to repeat complex behavior with great frequency over long periods of time. We are not going to add a number to both sides of an equation once every 3 s for a month. A robotic system, however, might perform such a task with great frequency for a long duration. To assist SS-RICS developers in modifying the structure of ACT-R, this report examines the behavior of the baseline terms of the activation equation given various intensities and patterns of activation. In the next section, we present several baseline activation equations used in ACT-R and a new one developed for comparison. We also present a brief asymptotic analysis of the behavior of these equations under frequent use for long duration. Subsequently, we explore the behavior of these equations on simulated activation timelines. We intend to provide SS-RICS developers with intuition to help choose the appropriate activation equation for different types of memory. In the conclusions, we make observations on and recommendations for the use of different activation terms and make suggestions as to the direction of future research.

2. Defining Baseline Activation Terms

We will consider several alternate forms for the baseline activation term B of equation 1 and examine the following aspects of that baseline activation term:

- How does activation change with frequency of use?
- Is activation value bounded?
- How does activation change with recency of use?
- How does baseline activation decay with no or infrequent use?
- How much memory is required for the equation to be used?

The first baseline activation term presented in the ACT-R Reference Manual (6) is

$$B\left(t, \{t_j\}_{j=1}^n\right) = \ln\left(\sum_{j=1}^n (s_j)^{-d}\right) + \beta; \quad s_j := t - t_j \quad (2)$$

where β is a constant activation for the chunk, d is a decay parameter, t is the present time, and t_j is the j th time the chunk was used, with t_1 being the first use of the chunk and t_n being the most recent. Since the logarithm function is negative for numbers within the interval (0,1), this function will quickly become negative after a single use. This could be prevented by making activation the maximum of the result of equation 2 or zero. Alternatively, it may be desirable for

the term to decay quickly to offset the increase caused by other terms in equation 1, such as spreading activation. Since the ACT-R buffers are small, it may also be desirable to have quick decay to force chunks out of the buffers. In SS-RICS, working memory is larger, and chunks could remain in it for a long time. Thus it may be undesirable for the activation value to decay so quickly.

While the initial behavior of equation 2 is quick decay after one use, the long-term behavior with regular use is different. To gain intuition as to this behavior, we consider what would happen if a chunk is used every T/n time increment for a time period of length T . Under this assumption,

$$\begin{aligned} B_i &:= B_i(t, \{t_j\}_{j=1}^n) = \ln \left(\sum_{j=1}^n (s_j)^{-d} \right) + \beta_i; \quad s_j = T - (j-1) * T/n \\ &= \ln \left(\sum_{j=1}^n (T/n)^{-d} * \left(1 - \frac{j-1}{n} \right)^{-d} \right) + \beta_i \\ &= \ln \left(\left(\frac{n}{T} \right)^d * \sum_{k=1}^n \left(\frac{1}{k} \right)^d \right) + \beta_i \end{aligned} \quad (3)$$

We will estimate the series inside the logarithm in equation 3 using an integral approximation. For $d < 1$,

$$\frac{(n+1)^{1-d}}{1-d} - \frac{1}{1-d} = \int_1^{n+1} x^{-d} \leq \sum_{k=1}^n \left(\frac{1}{k} \right)^d \leq \int_1^{n+1} x^{-d} + 1 = \frac{(n+1)^{1-d}}{1-d} - \frac{d}{1-d} \quad (4)$$

Using $\sim$ to denote a relationship in which the ratio of two terms converges to 1 as n increases, and using $\cong$ to denote a relationship in which the difference of two terms converges to 0 as n increases, equation 4 implies that

$$\sum_{k=1}^n \left(\frac{1}{k} \right)^d \sim \frac{(n)^{1-d}}{1-d} \quad (5)$$

which, in combination with equation 3, implies that

$$B_i \cong \ln \left(\left(\frac{n}{T} \right)^d * \frac{(n)^{1-d}}{1-d} \right) + \beta_i = \ln \left(\frac{n}{T^d(1-d)} \right) + \beta_i \quad (6)$$

Thus, when $n < T^d(1-d)$, baseline activation can be less than β . So for infrequent use, baseline activation values can go below β for some number of uses. Over long periods of time, however, something different happens. Suppose that the frequency n/T is a constant, c . Then equation 6 becomes

$$B_i \cong \ln \left(c^d * \frac{n^{1-d}}{1-d} \right) + \beta_i \sim (1-d) \ln n \quad (7)$$

So a regularly occurring event will have a growth in activation that is proportional to the logarithm of the number of times it occurs. This proportionality can be adjusted by changing the decay parameter d . With regular use, baseline activation increases without bound, but if you stop after a few uses, activation falls below baseline. The rate of increase, however, is slow. Using the

approximation in equation 7 with $d = 0.5$, the baseline activation (approximately) reaches integer values as shown in table 1. The increase is slow and only relevant for tasks regularly repeated for long duration. For values of $d > 1$, equation 2 is bounded and decays much faster.

We consider values of $d < 1$ of primary interest because the results are more interesting, but we provide some numerical examples with $d > 1$ in appendix A (figure A-3).

Table 1. Baseline activation as a function of n .

B_i (from equation 7)	Value of n (rounded to the nearest integer)
1	7
2	54
3	403
4	2,980
5	22,026
6	162,754
7	1,202,604

The estimate in equation 6 is suggested in Avery et al. (5) as a low memory way to keep track of activation values. If the assumption of regular activation is reasonable, it may be a good approximation for equation 2. We suggest, however, a more flexible variation of the baseline activation equation 2, also found in Avery et al., which remembers a fixed number (k) of uses:

$$B_i = \ln \left(\sum_{j=n-k+1}^n (s_j)^{-d} + \frac{(n-k)(s_1^{1-d} - s_{n-k}^{1-d})}{(1-d)(s_1 - s_{n-k})} \right) + \beta_i \quad (8)$$

The first term inside the logarithm tracks the k most recent uses of the chunk, and the second is an approximation similar to those made in equations 4–6. The implicit assumption is that the untracked terms are equally distributed over the untracked portion of time. If the history of the chunk is one of intense but sporadic use, this could be a poor estimate of equation 2. In short-term sporadic use, however, the difference may be minor, and may be negligible after a period of decay. Examples of this behavior appear in sections 3.3 and 3.4. If there are k uses or fewer, then equations 2 and 8 produce the same values.

Refraction is a short-term depression in the value of the baseline equation. In Thomson (8), the authors explain why it may be desirable to have a baseline term with refraction. We will not repeat their reasons but rather present only the baseline activation term they suggest:

$$B_i = \ln \left(\sum_{j=1}^n (s_j)^{-d} \right) - \ln \left(1 + \left(\frac{s_n}{\tau} \right)^{-d_r} \right) + \beta_i \quad (9)$$

Here τ is a time-scaling parameter and d_τ is a short-term decay rate. The term inside the second logarithm depresses the activation by a function of τ , with this depression decreasing as time passes. This version keeps the entire history of use but could easily be altered to keep only the last k terms by including a refraction term in equation 8.

Since these activation terms are all variations of equation 2, we include a distinct alternative with which to compare them. This alternative is an exponential decay baseline activation term with refraction:

$$B_i(t, t_n) = X_i(t, t_n) + \beta_i$$

where

$$X_i(t, t_n) = l \text{ for } t - t_n < r$$

and

$$X_i(t, t_n) = 2^{-(t-t_n)/h} (X_i(t_n) + b * 2^{r/h}) \text{ for } t - t_n \geq r \quad (10)$$

Here l is the refraction level, r is the length of the refraction, h determines the decay rate, and b determines the increase to the baseline after refraction ends. The decay is exponential, so that in the absence of use, the X term will decay by a factor of two every h time units. This exponential decay has a cost and a benefit. The cost is accelerated forgetting and the benefit is a constant upper bound independent of how many times the chunk is activated depending on the parameters h and b . The bound comes from observing that

$$X(t, t_n) \leq b + b * \sum_{i=1}^{\infty} 2^{-i/h} = b + b * \frac{1}{1-2^{-1/h}} \quad (11)$$

The variations of equation 2 do not have a natural upper bound but could be bounded by simply limiting the value to which it can increase. Equation 10 requires remembering only the time of last activation and the value at that activation time. The refraction effect resulting from equation 10 is sharp and sudden and different from the effect caused by the refraction term in equation 9, which is more gradual. This is intentional to provide an alternative version of refraction. We could make the refraction of equation 9 sudden by replacing the refraction term there with the refraction behavior of equation 10. Likewise, we could make the refraction in equation 10 more gradual by replacing its refraction behavior with a term like that in equation 9.

We are aware of another version of an ACT-R baseline activation term (8) that causes baseline activation to decay as a function of the length of time the chunk was held in a buffer. The longer the chunk is in a buffer, the slower the decay rate. This dynamically generated decay strength is not considered here because the reliance on time in buffer is not appropriate criteria for SS-RICS. In ACT-R, buffers (working memory) are small, and retention in working memory could be a reasonable proxy for importance. In SS-RICS, working memory is large, and retention in working memory should not be considered as a proxy for importance.

Finally, the terms activation and refraction come from terms describing the behavior of neurons, which become activated and then experience a refractory period before returning to a resting state. A graph showing such firing is included in figure 1. It is not our intent here to develop equations that reproduce this graph. We are examining the increase in activation and refraction in current ACT-R equations and in an alternative chosen for its contrast with those equations. We will discuss the possibility of developing alternative models in the conclusions section.




Figure 1. A generic graph demonstrating the concept of refraction (10).

In the next section, we compare the activation terms given sequences of activations. That is, we assume that a chunk has been activated at certain times and we observe the activation value. This provides insight into the activation terms for SS-RICS developers, but it is not a simulation of the performance of SS-RICS given these activation terms. Such a simulation would require more time and greater resources and is not justified at this stage.

3. Comparing Baseline Activation Terms

Our goal is to illustrate the behavior of baseline activation terms assuming activation at different time intervals. We will begin with a few examples of sparse activation where only two activations occur. This will give us insight into how each term behaves when activation occurs rarely. Then we will examine cases in which activation occurs frequently. We consider regular sequences in which the activations occur at a fixed frequency, sporadic sequences in which the activations occur clustered together with gaps between clusters, and random (binomial) sequences with a given probability of activation occurring at each time increment.

For each of these cases, we examine sequences where activation occurs with expected proportions ranging from 0.05 to 0.2. We consider proportions much lower than 0.05 to be covered by the sparse examples and values greater than 0.2 are dense enough that behavior is an easily observable increase in the value of the activation equations for all cases. We use a length of time of up to 400 time increments. We find that much beyond that the graphs are difficult to interpret.

Our focus here is on comparing different versions of the baseline activation term, but of course each version has multiple parameters. In appendix A we have included graphs of the activation terms with varying parameters. For all examples, here and in the appendices, the parameter β is set to 0. The other parameters for the examples in this section are in table 2.

Table 2. Parameter values and legend for all graphs in section 3.

Equation Abbreviations	Parameter Values	Color
2 Base	$d = 0.5$	Red (solid)
8 Short Base	$d = 0.5, k = 3$	Orange (dashed)
9 Base w/Ref.	$d = 0.5, dr = 0.5, \tau = 10$	Blue (solid)
10 Exp	$l = 0, r = 5, h = 10, b = 0.2$	Purple (dashed)

3.1 Examples With Sparse Activation

For sparse activation, we examine only three terms: Base, Base w/ Ref., and Exp. For these sparse cases, there is no need to use Short Base since keeping the last three terms would be the same as keeping all terms. The values of the activation equations are set to zero until the first activation occurs. In all figures, the lines are interpolations of the baseline activation values, which are evaluated only at whole number time increments. The patterns are easier to follow with lines than they would be with a discrete plot. After initial use, the Base and Base w/ Ref. activation terms show a decrease, which we can understand by examining equations 2 and 9 and by recalling that the logarithm of number less than 1 is negative. One goal for the activation function could be that, outside of a refractory period, activation should increase and not decrease after use, in which case this behavior would be undesirable. If, however, it is desired that activation decrease after single use but increase with regular use, Base and Base w/ Ref. do perform this function, as is shown comparing figures 2 and 3 here with figure 3 in section 3.2.




Figure 2. Sparse activation at time periods 3 and 13; equations 2 (red), 9 (blue), and 10 (purple).




Figure 3. Sparse activation at time periods 3 and 23; equations 2 (red), 9 (blue), and 10 (purple).

3.2 Examples of Activation With Fixed Frequency

To see how activation values change with regular use, we suppose that activation occurs every 10 or 20 time periods for 200 time units (proportion of activation of 0.1 and 0.05, respectively). From the analysis of section 2, we expect that Base and Short Base will yield similar results when the frequency is regular, so we will again examine the three terms we considered in the last section. The spacing of activation and the length of refraction time r for Exp were chosen to make the trends of the activation value graphs easy to interpret. Activation, which is much more or much less frequent, results in a graph that is difficult to interpret. More-frequent activation would cause the activation equations to increase more quickly and make the dips in the graph of Exp more dramatic, while less-frequent activation would have the opposite effect. All terms exhibit an upward trend with regular use, with refraction having less impact on Base and Base w/ Ref. than on Exp.

With regular activation, it is easy to see the difference in the Base and Base w/ Ref. terms. Looking at figures 4 and 5, we generally see an immediate rise upon activation for Base and an immediate dip for the Base w/ Ref. Both exhibit general upward trends. We can also see how refraction slows the overall rate of growth for Base w/ Ref. but diminishes in significance as the activation value increases. We see that refraction does not slow the rate of increase of the Exp (purple) but always results in a decrease to the same value (by design). For a chunk that is regularly used, refraction will cause a temporary reduction of that activation of up to $\log \tau$ in the Base w/ Ref. and will result in a decrease to a set value for Exp. Either one may be better in a given circumstance. The implications of these differences are discussed in section 4.




Figure 4. Plots of equations 2 (red), 9 (blue), and 10 (purple), with activation every 10 time periods.




Figure 5. Plots of equations 2 (red), 9 (blue), and 10 (purple), with activation every 20 time periods.

3.3 Examples of Sporadic Activation

Having examined how baseline terms behave given regular activation, we now consider their behavior when activation is intense over some intervals and sparse over others. We consider three test sequences—lengths of 100, 150, and 200—each having 30 activations. We found the graphs easier to read when we changed the proportion of activations by altering length of time rather than by changing the number of activations. We cluster the activations into intervals with gaps between to see what differences exist in this case as compared with the fixed frequency case. Sometimes the activations occur with no break between them while at other times they are separated by a number of time steps less than the length of r (table 2). Here we do plot the values of Short Base (equation 8), which we expect to be different from Base in these examples.

Examining the results of figures 6–8, we are surprised how close the values produced by Short Base are to those of Base. Because the estimation in equation 8 assumes a regular distribution of activations, more sporadic activation yields more divergence. Only tracking the three most recent terms, Short Base stayed close to the Base equation and seems like a good substitute when one needs to save memory. The Exp equation, because it refracts to a set value, was fixed at that refraction value during periods of intense use, which could be a benefit or liability depending on the situation. In figure 8, we see that during the time period from 20 to 70, Base w/ Ref. was negative the whole time while Base was negative half the time, which could be also be a benefit or a liability.




Figure 6. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30 activations over 150 time steps.




Figure 7. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30 activations over 100 time steps.




Figure 8. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), with 30. activations over 200 time steps

3.4 Examples of Random Sequences

We considered the fixed-frequency, sporadic, and sparse cases motivated by situations we could imagine occurring during use. To deal with situations that we cannot imagine, we randomly chose times at which to activate the chunk and observe the activation values that result. We used 10 random series of length 200, activated at each time interval with probability 0.1 or 0.05, and 5 series of length 400 activated at each time interval with probability 0.05. Graphs of the activation values for 3 such series are shown in figures 9–11 and the rest are in appendix B.




Figure 9. Plots of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.




Figure 10. Plots of equations 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.




Figure 11. Plots of equation 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.

It may look as though Exp decreases below zero at around time unit 300 in figure 11; this is not the case. It is the case, however, that Exp drops below the value of Base as was seen before in the sporadic case (figure 6). The random activation examples (including those in appendix B) do not show any new behavior of the activation terms. This leaves us feeling comfortable enough to draw some conclusions about when each term is appropriate for use.

4. Conclusions

Upon initial use, the Base term (and its variations) will begin to decay and will go below zero. Unless all other baseline activations are similarly decaying, this will result in lower likelihood of use following initial use. With regular use, the Base term (and variations) will increase without bound. Thus, rarely used chunks become less likely to be used after initial use and regularly used chunks become more likely to be used after a brief time period in which the activation values get large. Presumably there are mechanisms for dealing with the unbounded nature of the terms' behavior in ACT-R, and some mechanism will need to be implemented in SS-RICS. One method would be to program-in upper and lower bounds for the baseline activation values. If the chunk reaches its lower bound, it could stay at that value for a certain length of time and then get reset to its baseline value β . This makes refraction built-in when use is infrequent. For frequent use, one could use an upper bound. Values reaching that upper bound would be kept there until natural decay takes the value below the upper bound.

Refraction (equation 9) seems unnecessary when activation of a chunk is a rare occurrence. When activation is frequent, the activation value will follow a generally increasing trend with or without refraction, so if a chunk is much more commonly used than all other chunks competing (for selection) with it, refraction would not be sufficient to give the competitors a chance. To see the value of refraction, one must imagine two chunks with roughly the same activation values competing to be selected as a match. One is then chosen, and soon after that the same matching criteria appear again. Should the one just chosen be more or less likely to win again as a consequence of having just been selected? If the answer is less likely, then a refraction term (as in equation 9) should be included, and if the answer is more likely, such a term should be omitted. Figure 8 provides a good example of the consequences of including a refraction term.

The refraction of equation 10 is of a different variety. The chunk is essentially eliminated from competition for a period of time and then returns as more likely to be selected. Such refraction would be appropriate for a resource that, once used, is unavailable for some period of time. Such a thing might not happen with human cognition, but one can image a robot appendage that, once ordered to a task, cannot be used again until the time to complete the task is over.

To complete our comparison of equation 10 with the other equations, consider that after refraction the former increases above its initial value after its first use and after every subsequent use. This would mean that as soon as you use a chunk you become more likely to use it as soon as you can. This may not be a good model for mature cognitive function and brings to mind a student who, upon solving one type of problem on a test, assumes that a similar solution would work for every subsequent problem. But there are times when such function is useful, as when you see a child running across the street and subsequently you become more attentive to children in the street. It may be the case that different activation terms should determine the activation value of different chunks.

To suggest a line of future research, we observe that none of the activation equations presented here allow for habituation, which is a decline in stimulus resulting from repeated presentation (11, 12). Including such a feature has been suggested by SS-RICS researchers (13). Such a feature would allow a robot to become habituated to a repeated stimulation and eventually turn its attention to something else. We believe that developing such an equation is a good idea, and intend to do so in collaboration with those researchers. However, we also suggest that multiple activation equations could be a more beneficial way of representing the complete spectrum of cognition. Different types of memories could have different types of activation equations. This would allow us to use human models of cognition for behaviors in which it is superior to that of machines but not in areas in which it is inferior. If our robot is searching a door, we want it to eventually get bored with searching in one area and to try different areas. If our robot is on guard, it should never get bored with looking at its assigned area of responsibility. It may be that the type of activation should be chosen based on the command given by the user. "Find the door.

Start looking over there” would result in selection of an activation equation with habituation. “Watch that path until I tell you to stop” would result in selection of an equation without habituation.

We conclude with some concise observations and recommendations regarding baseline activation terms:

- Equation 8 is a useful memory saving version of equation 2. Unless the entire time series of activations is going to be retained for later learning, we recommend using equation 8 to save memory. The difference between the two may well be insignificant once noise (ϵ) is included.
- Equation 2 and its variations are all unbounded, but this could be easily addressed by putting upper and lower limits on their values. I would not choose a different activation equation just because these are unbounded.
- The difference in refraction between equations 9 and 10 is significant. Equation 10 essentially sets a chunk off limits to search for a period of time no matter how high its activation is, while equation 9 lowers activation value by a predetermined amount. The refraction in equation 9 matters most when a chunk has several similarly valued competitors.
- The refraction properties of equations 9 and 10 could be exchanged if only a change to refraction is desired.
- Equation 10 puts a chunk above its initial value immediately after first use while the other equations allow for quick decay to values below the initial value. The programmer must decide if the likelihood of using a chunk should increase, stay the same, or decrease after a single use.
- We should develop an alternative activation equation with habituation and the type of refraction shown in figure 1.

To replicate the results here, the necessary functions and test cases are included in appendix C.

5. References

1. Kelly, T.; Avery, E.; Wilson, N. U.S. Army Research Laboratory, Aberdeen Proving Ground, MD. Private communication, July 2013.
2. Kelly, T. Developing a Psychologically Inspired Cognitive Architecture for Robotic Control: The Symbolic and Sub-symbolic Robotic Intelligence Control System (SS-RICS). *International Journal of Advanced Robotic Systems* **2006**, 3 (3), 219–222.
3. Kelley, T. D.; Avery, E. *A Cognitive Robotics System: The Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS)*. <http://www.ss-rics.org/Publications/Publications.aspx> (accessed on 2 December 2013).
4. Kelly, T.; Avery, E.; Long, L.; Dimperio, Eric. *A Hybrid Symbolic and Sub-Symbolic Intelligent System for Mobile Robots*. <http://www.ss-rics.org/Publications/Publications.aspx> (accessed on 2 December 2013).
5. Avery, E.; Kelley, T.; Davani, D. *Using Cognitive Architectures to Improve Robot Control: Integrating Production Systems, Semantic Networks, and Sub-Symbolic Processing*. <http://www.ss-rics.org/Publications/Publications.aspx> (accessed on 2 December 2013).
6. Bothell, D. *ACT-R 6 Reference Manual*; Carnegie Mellon University; working draft. <http://act-r.psy.cmu.edu/actr6/> (accessed on 2 December 2013).
7. Anderson, J.; Bothell, D.; Byrne, M.; Douglass, S.; Lebiere, C.; Yulin, Q. An Integrated Theory of the Mind. *Psychological Review* **2004**, 111 (4), 1036–1060.
8. Thomson, R. *Associative Learning and Buffer Decay Modules for ACT-R 6.0 Reference Manual*; version 5.
9. Lebiere, C.; Best, B. Balancing Long-Term Reinforcement and Short-Term Inhibition. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Austin, TX, 2009; Cognitive Science Society: Austin, TX, 2009.
10. “Chris 73.” Action Potential; Wikimedia Commons image. http://commons.wikimedia.org/wiki/File:Action_potential.svg (accessed 28 March 2014).
11. Marsland, S. Using Habituation in Machine Learning. *Neurobiology of Learning and Memory* **2009**, 92, 260–266.
12. Wang, D. A Neural Model of Synaptic Plasticity Underlying Short-Term and Long-Term Habituation. *Adaptive Behavior* **1994**, 2 (2), 111–129.
13. Kelly, T.; Avery, E.; Wilson, N. U.S. Army Research Laboratory, Aberdeen Proving Ground, MD. Private conversation, January 2014.

INTENTIONALLY LEFT BLANK.

Appendix A. Baseline Activation Terms With Different Parameters

In figure A-1 and table A-1 we show the effects of different parameter values for equation 10 of the report. From the beginning, a higher value of b causes the blue curve to increase more than the others, and over time a faster rate of decay causes the purple curve to increase more slowly than the red curve. Higher values of b and slower rates of decay resulting from increased h result in greater increases and a more dramatic effect swing to a refraction value.




Figure A-1. A comparison of parameter values for equation 10 (see table A-1).

Table A-1. Parameter values for figure A-1.

Color	Parameter Values
Red (dashed)	$r = 3, h = 10, l = -0.1, b = 0.1$
Blue (dashed)	$r = 3, h = 10, l = -0.1, b = 0.2$
Purple (solid)	$r = 3, h = 5, l = -0.1, b = 0.1$

In figure A-2 and table A-2 we show the effect of the decay parameter on equation 2 of the report. A larger decay value results in faster decay and slower increase with regular activation.




Figure A-2. A comparison of parameter values for equation 2 (see table A-2).

Table A-2. Parameter values for figure A-2.

Color	Parameter Values
Red (dashed)	$d = .50$
Blue (dashed)	$d = .25$
Purple (solid)	$d = .75$

We have conducted most of our examination of equations 2–9 of the report with values of d less than 1. For completeness, however, we include in figure A-3 and table A-3 a repetition of the activation series of figure A-2 with some d values greater than 1. The baseline activation term is bounded with this decay parameter value because it decays quickly.




Figure A-3. A comparison of parameter values for equation 2 (see table A-3).

Table A-3. Parameter values for figure A-3.

Color	Parameter Values
Red (dashed)	$d = .5$
Blue (dashed)	$d = 1.25$
Purple (solid)	$d = 2.0$

In figure A-4 and table A-4, we can see the effect of the refraction decay parameter on equation 9 of the report. Deeper refraction results from higher values of the refraction decay parameter. When the parameter is set to 0.1, no meaningful refraction is discernible.




Figure A-4. A comparison of parameter values for equation 9 (see table A-4).

Table A-4. Parameter values for figure A-4

Color	Parameter Values
Red (dashed)	$d = .5, dr = 0.1, \tau = 10$
Blue (dashed)	$d = .5, dr = 0.5, \tau = 10$
Purple (solid)	$d = .5, dr = 0.9, \tau = 10$

In figure A-5 and table A-5, we can see the effect of the time scale parameter on equation 9. Deeper refraction results from higher values of the scale parameter.




Figure A-5. A comparison of parameter values for equation 9 (see table A-5).

Table A-5. Parameter values for figure A-5.

Color	Parameter Values
Red (dashed)	$d = .5, dr = 0.5, \tau = 5$
Blue (dashed)	$d = .5, dr = 0.5, \tau = 10$
Purple (solid)	$d = .5, dr = 0.5, \tau = 20$

Appendix B. Replications of the Examples of Section 3.4 of the Report

Since the series used in section 3.4 of the main report were random, we repeated each length and probability of activation with four additional random series, keeping the same parameters as in section 3.4.




Figure B-1. Replication 1 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.




Figure B-2. Replication 2 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.




Figure B-3. Replication 3 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.




Figure B-4. Replication 4 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.1.




Figure B-5. Replication 1 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-6. Replication 2 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-7. Replication 3 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-8. Replication 4 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-9. Replication 5 of equations 2 (red), 8 (orange), 9 (blue), and 2.10 (purple), randomly activated at each time step with probability 0.05.




Figure B-10. Replication 6 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-11. Replication 7 of equations 2 (red), 8 (orange), 9 (blue) and 10 (purple), randomly activated at each time step with probability 0.05.




Figure B-12. Replication 8 of equations 2 (red), 8 (orange), 9 (blue), and 10 (purple), randomly activated at each time step with probability 0.05.

INTENTIONALLY LEFT BLANK.

Appendix C. Code for Reproducing Test Results

The following R code is provided for those wanting to replicate the results in this report. Included are the function definitions, followed by the test cases.

C.1 Function Definitions

Preliminary functions for averaging and refraction:

```
inv.d.pwr<-function(x,d) return(1/(sign(x)*abs(x)^d ));
d.pwr<-function(x,d) return(sign(x)*abs(x)^d);
b.avg<-function(s1,snk,d,n,k){
  tp<-(-n-k)*(d.pwr(s1,1-d)-d.pwr(snk,1-d))
  btm<-(-1-d)*(s1-snk)
  return(tp/btm)
};
b.refract<-function(sn,dr,scale.time){
  if(sn==0){return(0)}
  x<-1+d.pwr(scale.time/sn,dr)
  return(-log(x))
}
```

Baseline activation term definitions:

Base definition

```
base.act.0<-function(time, act.ts,d,beta=0){
  if(length(act.ts)==0){return(beta)}
  x<-sum(inv.d.pwr(time-act.ts,d))
  return(log(x)+beta)
}
```

Short Base definition

```
base.act.1<-function(time,k.act.ts,s1,snk,d,n,k,beta=0){
  if(length(k.act.ts)==0){return(beta)}
  if(snk==0){x<-sum(inv.d.pwr(time-k.act.ts,d));return(log(x)+beta)}
  x<-sum(inv.d.pwr(time-k.act.ts,d))+b.avg(s1,snk,d,n,k)
  return(log(x)+beta)
}
```

Base w/Ref. definition

```
base.act.2<-function(time, act.ts,d,sn,dr,scale.time,beta=0){
  if(length(act.ts)==0){return(beta)}
  x<-sum(inv.d.pwr(time-act.ts,d))
  if(sn==0){return(log(x)+beta)}
  return(log(x)+b.refract(sn,dr,scale.time)+beta)
}
```

Exp definition

```
exp.act<-function(time,sn,r,h,l,c=2,b,Xn){
  if(sn==0){return(list(0,Xn))}
  X.up<-d.pwr(c,-(sn)/h)*(Xn +b*d.pwr(c,r/h))
  if(sn==1){Xn<-X.up}
  if(sn<r){return(list(l,Xn))}
  else{return(list(X.up,Xn ))}
```

```
}
```

A function for preparing an activation time series for Base Short:

```
k.trim.ts<-function(time,act.ts,k){
  n<-length(act.ts)
  if(n==0){return(list("time"= time,"n"=0,"k"=k,"s1"=0,"snk"= 0,"k.act.ts"=NULL ))}
  k.act.ts<-act.ts
  if(n-k<=1){return(list("time"= time,"n"=n,"k"=k,"s1"=0,"snk"= 0,"k.act.ts"=k.act.ts ))}
  k.act.ts<-act.ts[(n-k+1):n]
  s1<- time-act.ts[1]
  snk<- time-act.ts[n-k]
  return(list("time"= time,"n"=n,"k"=k,"s1"=s1,"snk"= snk,"k.act.ts"=k.act.ts ))
}
```

C.2 test cases used in the report

```
beta<-0
d=.5
dr=.5;k.val=3
scale.time=10
r=5;h=10;l=-1;c=2;b=0.2
```

For each instance, define the vector *act.50* as appropriate for the test case, as shown in the following:

Sporadic:

```
#act.50<-c(rep(c(1,0),5),rep(0,10),rep(c(0,1,0,0),5),rep(0,20),rep(c(0,1,0),10),rep(0,25),
rep(1,10),rep(0,25) )
#act.50<-c(rep(c(0,1),5),rep(0,10),rep(c(1),15),rep(0,10),rep(c(0,1,0),5),rep(0,40) )
#act.50<-c(rep(c(0,1),5),rep(0,60),rep(c(0,1,0),20),rep(0,30),rep(c(0,1,0),5),rep(0,25) )
```

Regular:

```
#act.50<-rep(c(0,1,0,0,0,0,0,0,0,0),20)
#act.50<-rep(c(0,1,0,0,0,0,0,0,0,0,rep(0,10)),10)
```

Random:

First set the random seed

```
#set.seed(23)
```

Then execute five of these

```
#act.50<-rbinom(200,1,.1)
```

Now set the second seed

```
#set.seed(27)
```

Five of the first then five of the second

```
#act.50<-rbinom(200,1,.05)
#act.50<-rbinom(400,1,.05)
```

Having defined the test case, run the following code:

```
act<-act.50
act.len<-length(act)
act.val.mtx.0<-matrix(0,ncol=num.curve,nrow=act.len)
```

```

act.val.mtx.1<-matrix(0,ncol=num.curve,nrow=act.len)
act.val.mtx.2<-matrix(0,ncol=num.curve,nrow=act.len)
act.val.mtx.exp<-matrix(0,ncol=num.curve,nrow=act.len)
Run some number of times each case, as defined by num.curve.
num.curve=1
for(li in 1:num.curve){
  act.times<-NULL
  act.val<-rep(0,act.len)
  act.val.base.0<-act.val;act.val.base.1<-act.val;act.val.base.2<-act.val;
  act.val.exp<-act.val
  for(i in 1:act.len ){
    time<-i+1
    if(act[i]==1){act.times<-c(act.times,i)}
    if(length(act.times)==0){sn<-0}else{sn<-time-rev(act.times)[1]}
    a0<-base.act.0(time, act.times,d=d)
    act.val.base.0[i]<-a0
    klst<-k.trim.ts(time,act.times,k=k.val)
    a1<-base.act.1(time=time,k.act.ts=klst[["k.act.ts"]],s1=klst[["s1"]],
                  snk=klst[["snk"]],d=.5,n=klst[["n"]],k=k.val)
    act.val.base.1[i]<-a1
    a2<-base.act.2(time, act.times,d=d,sn=sn,dr=dr,scale.time=scale.time)
    act.val.base.2[i]<-a2
    if(i==1){Xn<-0;X.up<-0}else{
      E.list<-exp.act(time,sn=sn,r=r,h=h,l=l,c=c,b=b,Xn=Xn)
      Xn<-E.list[[2]]; ax<-E.list[[1]]
      act.val.exp[i]<-ax
    }
  }
  act.val.mtx.0[,li]<-act.val.base.0
  act.val.mtx.2[,li]<-act.val.base.2
  act.val.mtx.1[,li]<-act.val.base.1
  act.val.mtx.exp[,li]<-act.val.exp
}

```

Plots are made using a the command with labels altered as appropriate.

```

ts.plot(ts(act.val.mtx.1),ts(act.val.mtx.0),ts(act.val.mtx.2),ts(act.val.mtx.exp),
plot.type="multiple",col=c("orange","red","blue","purple"),
xlab="Time (no units)",ylab="Activation value",lty=c(4,1,1,3),
main="Activation values over time **",type="l",
sub=paste("*** d=",d, ", n=", act.len," p=0.15, sporadic"))

```

NO. OF
COPIES ORGANIZATION

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIRECTOR
(PDF) US ARMY RESEARCH LAB
RDRL CIO LL
IMAL HRA MAIL & RECORDS MGMT

1 GOVT PRINTG OFC
(PDF) A MALHOTRA

1 DIR USARL
(PDF) RDRL VTA
C LENNON

INTENTIONALLY LEFT BLANK.