The Dynamics of Anchoring in Bidirectional Associative Memory Networks

Sudeep Bhatia (sudeepb@andrew.cmu.edu)
Department of Social & Decision Sciences, Carnegie Mellon University,
5000 Forbes Ave. Pittsburgh, PA 15232 USA

Shereen J. Chaudhry (sjchaudh@andrew.cmu.edu)
Department of Social & Decision Sciences, Carnegie Mellon University,
5000 Forbes Ave. Pittsburgh, PA 15232 USA

Abstract

We formalize the biased activation theory of anchoring using a bidirectional associative memory network. Anchors determine the starting state of this network. As the network settles, we show that the nodes representing numerical responses activate and deactivate consecutively, generating sequential adjustment. By demonstrating that anchoring as adjustment emerges naturally from the dynamics of the biased activation process, we are able to unify the two main theories of the anchoring effect, and subsequently provide a parsimonious explanation for a large range of findings regarding anchoring, and its determinants. Although we focus largely on phenomena related to anchoring, the results of this paper apply equivalently to all judgments under the influence of bidirectional processing, including those involving constraint satisfaction.

Keywords: Decision Making, Neural Networks, Dynamic Processes, Anchoring Effect, Constraint Satisfaction

Introduction

Anchors have a powerful effect on human judgment. Responses to simple questions involving magnitude or time are systematically affected by uninformative numbers, known as anchors, displayed to the decision maker prior to the judgment task. High anchors generate high responses, low anchors generate low responses, and final judgments can be manipulated by selecting the appropriate anchor.

The anchoring effect has been shown to emerge in a large number of domains, and is one of the best studied judgment biases in psychology. Yet despite its importance, the cognitive mechanisms responsible for the anchoring effect are still being debated. In their seminal paper on heuristic choice, Tversky and Kahneman (1974) proposed that anchoring is caused by an imperfect sequential adjustment process. At each step in this process, decision makers evaluate the validity of a particular response. The judgment process terminates if the response in consideration is adequate; otherwise it moves on to the next feasible value. Anchors determine the starting point in this process, and adjustment is insufficient. Subsequently responses are closer to the anchor than optimal.

This explanation for the anchoring effect has been popular for many decades, and formal models of the anchoring effect have assumed that anchoring operates through sequential adjustment (Johnson & Busemeyer, 2005, but see also Choplin & Tawney, 2010). A more recent approach, however, claims that anchoring is the product of biased activation (Chapman and Johnson, 1994, 1999; Mussweiler & Strack, 1999). Anchors, according to this view, increase the accessibility of cues supporting the anchor. This evidence subsequently generates final responses that are closer to the anchor than optimal.

Is anchoring caused by sequential adjustment or biased activation? Both theories are supported by a large number of empirical findings (discussed in later sections), but neither is able to predict all of these findings by itself. In this paper we provide a simple answer to this question. We show that these processes are not necessarily distinct: sequential adjustment emerges from the dynamics of biased activation. Anchoring, thus, is caused by both these mechanisms simultaneously, and a large range of findings regarding anchoring and its moderators, can be explained within a unitary, parsimonious, theoretical framework.

Bidirectional Associative Memory

Consider a very simple judgment task. The decision maker is asked to select one of N responses based on M cues stored in memory. We assume, for simplicity, that the relationship between the responses and the cues is binary, with each cue either supporting or opposing each response. We can write a response i as r_i, and a cue j as c_j. If c_j supports r_i then we can write s_{ij}=+1, and if it opposes r_i then we can write s_{ij}=-1.

These responses can be numeric, as in typical anchoring tasks, or non-numeric as in more general judgment tasks. For numeric responses, we assume that the N nodes are ordered in a sequence r_1, r_2, …, r_N, corresponding to the sequence of available responses. For example, when considering the percentage of African countries in the United Nations, with responses in intervals of 1%, r_1, r_2, …, r_{100} correspond to the responses 1%, 2%, …, 100%.

We can implement this structure in a two layer neural network, with the first layer consisting of M nodes representing the M different cues, and the second layer consisting of N nodes representing the N response options. The activation of the node corresponding to c_j, at time t, can be written as $C_j(t)$, and the activation of the node corresponding to r_i, at time t, can be written as $R_i(t)$.

The connections from the cue layer to the response layer are equal to the strength of support provided by the cues to the responses. As activated response options (such as anchors) also affect the activation of the available cues, these connections can be assumed to be recurrent. Hence the connections from c_j to r_i and from r_i to c_j are both simply s_{ij}.

At a given time t, the activated nodes in the response layer first send inputs, weighted by s_{ij}, into the cue layer. This affects the activation of the nodes in the cue layer. The activated nodes in the cue layer subsequently send inputs weighted by s_{ij} into the response layer, affecting the activation of the response nodes at t+1, at which point the process repeats.

In addition to the inputs from the response layer, we assume that the nodes in the cue layer receive constant exogenous inputs with strength I=1. These inputs ensure that evidence nodes are activated even when none of the response nodes are active, and that the judgment process can begin in the absence of a response bias. We also assume that all of the nodes in our network have the same binary activation function, with a threshold at zero. With this assumption we can write the activation functions of any r_i as $R_i(t)=H[q_i]$, and any c_j as $C_j(t)=H[b_j]$ such that $q_i=\sum s_{ij}\cdot C_j(t-1)$, $b_j=\sum s_{ij}\cdot R_i(t) + 1$, and H as the unit step function with $H[x]=1$ for $x>0$ and $H[x]=0$ for $x\leq 0$.

We can now formalize the effect anchors have on the judgment process. We assume that anchors determine the starting state of the network. Hence if r_i is the anchor, then at t=1, we have $R_i(1)=1$, and $R_k(1)=0$ for $k\neq i$. In the absence of an anchor, the network begins with $R_k(1)=0$ for all k. Finally, we assume that responses active once the network stabilizes are the ones that are selected, and that the response time is proportional to the time it takes for the network to settle.

The proposed network is motivated primarily by the memory structure assumed to be at play in anchoring and related judgment tasks: indeed, it is one of the simplest possible cognitive instantiations of the biased activation theory of anchoring, which posits a recurrent relationship between cues and responses. That said, this network is ultimately a special case of the bidirectional associative memory (BAM) network, introduced in Kosko (1988). BAM itself generalizes the Hopfield network, which BAM resembles when node updating is asynchronous.

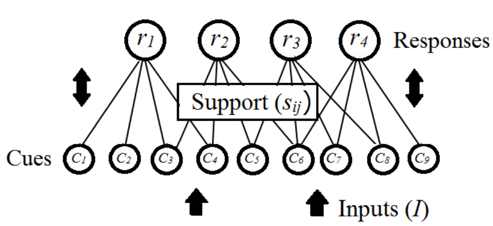

Figure 1: The BAM network.

Activation and Stability

What determines the responses that get activated at any time period, in the BAM network? The answer is cue overlap. Assume that only r_i is activated at time t. This activation causes only the cues that support r_i to be activated at t. Intuitively, the decision maker focuses on the cues that support the activated response and suppresses the cues that oppose the activated response. Once these cue nodes are activated, the activation pattern in the response layer changes. At t+1, responses supported by most of the cues activated at t turn on. These include r_i, but also other novel responses, that overlap sufficiently with r_i in cue support. Eventually at t+2 these responses activate other responses that they overlap with, and this process continues until the network stabilizes. Stability is always guaranteed: any BAM network with any memory structure, starting at any point, will stabilize in a finite number of time steps (Kosko, 1988).

Defining Sequential Adjustment

We hope to show that this settling process of the BAM network in the presence of anchors resembles sequential adjustment. Before we can do this, however, we need to understand what sequential adjustment really is. Sequential adjustment is generally defined as the successive movement through the range of responses available to the decision maker. In the simplest case, this definition imposes a form of serial processing, according to which only one response is considered at any given time. For example, when judging the proportion of African countries in the U.N., decision makers may first consider 1%. After rejecting this response they would consider 2%. If this too is inadequate they would move on to 3%, and so on. We consider the more general (and more realistic) case in which multiple responses can be considered at the same time. This allows decision makers to focus on all the responses within a particular interval, such as 1-10%, simultaneously, before moving on to the next interval in the sequence.

Such a dynamic is compatible with the general idea underlying sequential adjustment, as long as the responses activated are contiguous. Sequential adjustment does not permit the simultaneous consideration of different, non-neighboring responses. For example decision makers who consider both 1% and 99% simultaneously, without considering the responses between these two numbers, would not appear to be displaying sequential adjustment.

This then allows us to formalize the first requirement for sequential adjustment. This requirement, titled *contiguous activation*, states that sequential adjustment must not involve the simultaneous activation of multiple non-neighboring responses. Responses must be considered individually or in contiguous intervals.

Settling dynamics that display contiguous activation do not necessarily resemble sequential adjustment. It is possible for the decision maker to consider responses in contiguous intervals at any given time, but transition across different intervals in a non-sequential manner. For example, when evaluating the proportion of African countries in the U.N., decision makers could begin by considering the interval 1-10%, and then move to the interval 20-30%, without considering the interval 10-20%.

We thus need an additional requirement for our definition of sequential adjustment, in order to rule out these types of dynamics. This requirement, titled *sequential transitions*, states that sequential adjustment must not involve changes in activation that skip over a set of responses. Changes to response activation must be successive.

Connected Memory

Do the dynamics of the anchored BAM network satisfy contiguous activation and sequential transition? Not necessarily. However with a simple assumption about the underlying memory structure, these requirements can indeed be satisfied. This assumption relates to the distribution of cue support for the responses. In numeric judgments, cues can seldom support two disparate responses without supporting intermediate responses. For example, when judging the proportion of African countries in the UN, any cue that supports the 10% response, and the 12% response, should, in general, support the intermediate 11% response. This property, titled *connectedness*, more formally requires that a cue that supports r_i and r_k, also supports r_l for $i<l<k$. Memory structures displaying this property involve cues with a single, connected, interval of supported responses, where as those that do not display this property have cues with multiple, fragmented, intervals of supported responses.

While connectedness may not be satisfied in all judgment tasks, it is certainly a reasonable assumption when responses are ordered, as with the numerical scales used in anchoring tasks. Cues in these settings generally provide support for "large" responses, or "small" responses, or "medium" responses, or some other connected interval of responses. Very few cues provide support for a set of non-neighboring responses, distributed sporadically across the response scale. Indeed it is quite difficult to think of memory structures with diagnostic cues for numerical responses that do not satisfy the connectedness property.

The Emergence of Sequential Adjustment

When memory structures satisfy connectedness, then the resulting BAM network, with the anchored response activated at the start of the decision process, satisfies both contiguous activation and sequential transition. Of course, satisfying these properties does not imply that the decision maker necessarily adjusts away from the anchor. It may be the case that the anchor is stable. If there is adjustment, however, the adjustment is guaranteed to be sequential. Anchors trigger a cascade of activation in the response layer: Neighboring responses activate and deactivate consecutively. There are no jumps in response activation, nor do multiple non-neighboring responses activate, without the activation of the intermediate responses.

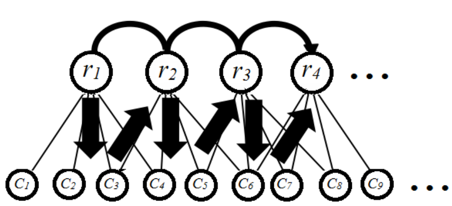

Figure 2: The emergence of sequential adjustment.

How does the connectedness property satisfy contiguous activation and sequential transitions? While the proof of this claim is in the appendix, the intuition for it is as follows. Due to connectedness, cues that support both the anchor and a non-neighboring non-anchored response must also support any intermediate responses, lying between the anchor and the non-neighboring response. Thus if the activation of the anchor activates cues that subsequently activate non-neighboring responses, these cues must also activate all of these intermediate responses. Subsequently, response activation at t=2 must be contiguous, and any transitions that may have happened at t=1 must be sequential. This intuition however also applies for the contiguous interval of responses activated at t=2, implying that any further changes to activation after t=2 must be sequential. Additionally, once a contiguous interval of responses is activated, we can show that connectedness implies that this interval cannot splinter into smaller, non-contiguous activated intervals, implying that contiguous activation must also be satisfied after t=2. Mathematical induction shows that these properties then hold at all times.

Connected BAM memory structures guarantee sequential activation. But can they generate insufficient adjustment? Let us consider the case with one correct response. When the memory structure is such that two nodes lying between the anchor and the correct response do not overlap on an appropriate number of cues, the sequential adjustment process described above will be insufficient: it will stabilize with the activation of response values closer to the anchor than the correct response.

The intuition for this is fairly straight forward. If, for a low anchor, there exist two response nodes between the anchor and the correct response, whose cue support does not overlap sufficiently, then the activation of the lower response node will not lead to the activation of cues that activate the higher response node. As activation must be contiguous and transitions must be sequential, no higher nodes can be activated, the network will stabilize with the activation of incorrectly low responses, and the correct response will remain turned off. The same intuition holds for tasks involving a high anchor, in which the network will stabilize with the activation of incorrectly high responses, and the correct response will remain turned off.

Demonstrations

The above sections have shown that the BAM network with connected memory structures satisfies contiguous activation and sequential transition, and can generate insufficient adjustment. While this is an analytical result, proved in the appendix, and guaranteed to hold regardless of any underlying parameters, demonstrations of the types of sequential adjustment generated by connectedness can provide important insights regarding the behavior emerging from the BAM network.

Figure 3 provides one such demonstration. It shows a hypothetical distribution of cue support for a sequence of responses, and the settling dynamics of the corresponding BAM network with a high anchor, low anchor, and without any anchor. The correct response in this network is r_4, and this is the stable response in the absence of an anchor. When anchored at r_6 (a high anchor), however, the network stabilizes at r_5. Similarly when anchored at r_1 (a low anchor), the network stabilizes at r_3. These behaviors indicate the presence of the anchoring effect. Additionally, the settling dynamics with these anchors display sequential adjustment: response nodes activate and deactivate consecutively until the network stabilizes.

Why do we observe these behaviors? r_1, r_2, and r_3 overlap on the component cues in such a way that the set of cues supported by r_i also on average support r_{i+1}, for i=1, 2. This means that activating r_1 leads to the activation of r_2, which then activates r_3. The set of cues supporting r_3 and r_4 do not however overlap in this way, implying that the cascade of activation begun by anchoring the network at r_1 ends with the stable activation of r_3. A similar property holds for r_5 and r_6. Also note that the network satisfies connectedness, which implies that the activation dynamics generated by the anchor display contiguous activation and sequential transitions, leading to sequential adjustment.

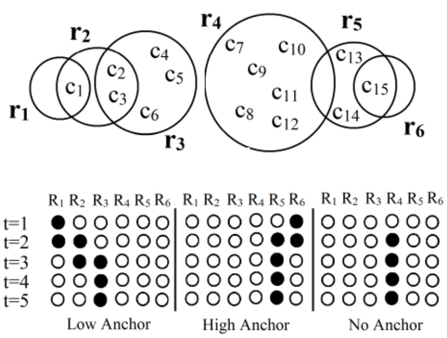

Figure 3: Distribution of cue support, and resulting network dynamics for low, high and no anchors.

These dynamics also emerge with larger, randomly generated memory structures. Consider a setting with N=100 responses and M=1000 cues. Let us randomly generate support or opposition between these cues and these responses. For each cue we can pick a number from the normal distribution with mean 50 and variance 25, and round it to its nearest integer. We can subsequently take an interval of length 20 around this integer, to generate the set of responses supported by the cue. All other responses are opposed by the cue. Taking an interval of responses around the randomly chosen number generates a “blurring” in the underlying memory structure: it is seldom the case that individual cues support point estimates; rather their support is distributed across an interval of responses.

As the randomly generated memory structure satisfies connectedness, it should be able to generate sequential adjustment. Figure 4 displays the dynamics of the BAM network instantiating this randomly generated memory structure, with a high anchor, r_{100} and a low anchor, r_1. Note that the stable responses for the two anchors are different, with the stable responses for the low anchor lower than the stable responses for the high anchor. Additionally, activation at all points of time is contiguous, and all transitions are sequential: we can observe a cascade of activation in the response layer over time, with intervals of responses activating and deactivating consecutively before finally stabilizing.

Note that the dynamics observed in figure 4 also emerge with alternate parameters in the model. In general, however, increasing the ratio of total responses to total cues and increasing the blurring in the cue support for the responses generates a higher likelihood of adjustment, as well as longer sequences of adjustment. This subsequently leads to weaker anchoring effects. Overall the anchoring bias is strongest when there are many relevant cues, and each cue supports few neighboring responses.

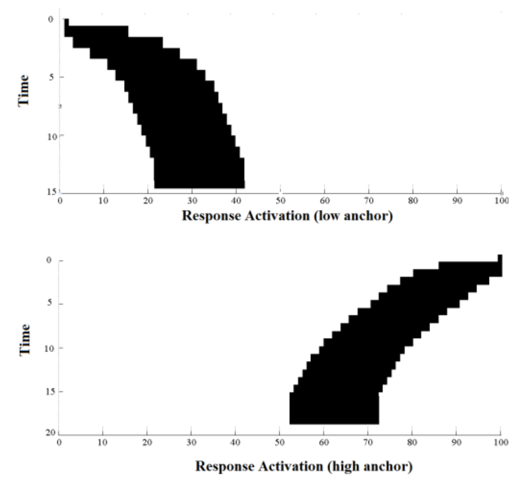

Figure 4: Network dynamics for high and low anchors, with randomly generated memory.

Explaining Anchoring Phenomena

Anchoring is a well-studied phenomenon and the sequential adjustment and biased activation theories of anchoring have a large range of behavioral findings that they must be able to account for. The above sections have shown that these theories are almost identical: the process assumed by one, emerges directly from the process assumed by the other. This section shows how this result can explain most of the findings documented in anchoring research.

Using a lexical decision task, Mussweiler and Strack (2000) find that decision makers identify “cold” related words quicker and more accurately after temperature judgments with low anchors, and identify “hot” related words quicker and more accurately after temperature judgments with high anchors. Sequential adjustment theory is unable to account for this finding, however, the BAM framework allows for both sequential adjustment and anchor dependent cue accessibility biases to emerge simultaneously: once the network settles, the cues that support the stable responses are themselves stable. If the judgment began with a low anchor then stable cues are more likely to support the low anchor than the high anchor. The opposite holds if the judgment began with a high anchor.

The biased activation theory of anchoring also predicts that exogenous factors influencing cue accessibility can affect anchoring. This has been verified by Chapman and

Johnson (1999) and Mussweiler et al. (2000). Unlike sequential adjustment theory, the BAM model can explain these findings. If we assume that exogenous influences on cue attention affect the inputs, I, into the cue layer, then directing attention towards cues that oppose the anchored response r_i, leads to stronger inputs, $I>1$, into these cues. Due to these inputs, these cues are not inhibited by feedback from the activated anchor in the response layer. Subsequently all cues are activated at the start of the decision process, the pattern of activation on the cue layer resembles the pattern observed in the absence of an anchor, and the network stabilizes without an anchoring bias.

According to the traditional sequential adjustment theory, all types of anchors, regardless of underlying cue support, should lead to the anchoring effect. Research by Chapman and Johnson (1994), however, finds that implausible anchors (anchors that are not supported by any cues) have a much weaker effect than plausible anchors. BAM provides a simple explanation for this result. When implausible anchors are activated at the start of the decision process, all cues are suppressed (as these anchors are not supported by any cues). Subsequently none of the response nodes activate in the next time period. This leads the network to a state identical to the starting state of the network in the absence of an anchor. Implausible anchors thus do not generate an anchoring effect.

A fourth finding supporting the biased activation theory of anchoring pertains to the effect of multiple anchors. Sequential adjustment theory predicts that the decision maker adjusts sequentially away from the one anchor presented in the decision task. This theory cannot make predictions for settings with multiple anchors. Switzer and Sniezek (1991) and Whyte and Sebenius (1997), however, demonstrate that multiple anchors affect judgment differently relative to single anchors. Single anchors paired with more extreme anchors generate a stronger anchoring effect than the single anchors alone, whereas single anchors paired with less extreme anchors generate a weaker anchoring effect than the single anchors alone.

BAM can account for the effect of multiple anchors. The activation of multiple response nodes at the start of the judgment process leads to the activation of all the cues supporting these anchors. When a single anchor is paired with a more extreme anchor then the set of cues activated are more likely to support extreme responses, relative to when the single anchor is activated by itself. This can lead to the stable activation of responses close to the extreme anchor, generating a stronger anchoring effect. The opposite happens when a single anchor is paired with a less extreme anchor. Here the activated cues are less likely to support extreme responses. This can lead to the ultimate stable activation of responses close to the moderate anchor, generating a weaker anchoring effect.

The cue accessibility, exogenous attentional influence, implausible anchor and multiple anchor results discussed above present strong evidence for the biased activation theory of anchoring. The standard biased activation theory cannot however provide a comprehensive account of all the moderators of the anchoring effect. Research by Reitsma-van Rooijen and Daamen (2006), for example, finds that time pressure increases the anchoring effect. This has traditionally seen as providing evidence for the sequential adjustment theory of anchoring, according to which time pressure limits the number of adjustments possible, thereby increasing the strength of the anchoring effect. As the BAM network proposed in this paper generates sequential adjustment, it is able to provide an explanation for these results as well. The BAM network often does not settle at its stable response in one time step; rather its response nodes activate and deactivate consecutively over time, before stabilizing at the final response (as in e.g. figure 4). When the decision maker is faced with time pressure, the network is not allowed to stabilize and the adjustment process generated in this network is curtailed, generating a stronger anchoring effect.

Another finding providing evidence for sequential adjustment theory relates to the role of incentives on anchored judgment. Particularly, Simmons et al. (2010) find that financial incentives reduce the anchoring effect. This cannot be explained by biased activation theory. If, however, we assume that incentivized decision makers send stronger inputs into the cue activation layer (perhaps due to increased attention towards all cues relevant to the decision task) then the BAM network can in fact explain this effect. As discussed above, when $I>1$, the exogenous inputs override the inhibitory feedback from the anchor in the response layer. Cue activation subsequently resembles the unbiased decision process, and the anchoring effect disappears.

Anchoring as Constraint Satisfaction

The bidirectionality assumed in this paper is a property of a general class of models that have been used to explain findings on inference across a variety of domains. These are models of constraint satisfaction (see e.g. Holyoak & Simon, 1999 for a review). Constraint satisfaction models provide a powerful approach to studying the interrelationships between cues and responses, and the ways that these relationships affect the dynamics of the decision process. Indeed, the anchoring effect can be seen as just a specific instantiation of the general type of starting point sensitivity displayed by these models: if the memory structures in these models satisfy connectedness then these models will also generate sequential adjustment. In this light, the BAM network is not just a model of anchoring, but rather a model of constraint satisfaction; one which provides a tractable framework with which to understand the cognitive dynamics that constraint satisfaction entails, and the behaviors that these dynamics can generate.

Conclusion

We have used the bidirectional associative memory network to study the anchoring effect. The BAM network provides a simple model for the biased activation theory of anchoring.

We have shown that the settling dynamics of this BAM network generate sequential adjustment. Anchors trigger a cascade of activation in the response layer of the BAM network, with nodes in this layer activating and deactivating consecutively. This progression of activation is generally insufficient and final responses depend critically on starting anchor values. By reconciling two contrasting theories within one framework, the BAM network is able to provide a parsimonious explanation for a wide range of findings regarding anchoring and its moderators.

APPENDIX

Here we shall show that BAM networks with connected memory structures satisfy contiguous activation and sequential transition. Let us define $\boldsymbol{D_i}$ to be the set of cues supporting r_i, $\boldsymbol{D^t}$ to be the set of cues activated at t, $\boldsymbol{E_j}$ to be the set of responses supported by c_j and $\boldsymbol{E^t}$ to be the set of responses activated at t. $|\boldsymbol{X}|$ shall indicate set $\boldsymbol{X}$'s cardinality. Now consider the following propositions:

Proposition 1a: If a contiguous interval of responses, r_i, r_{i+1}, … r_k is activated at t (and all other responses are deactivated at t), and for $l>k$, r_l is activated at t+1, then it is the case that r_k, r_{k+1} … r_{1-1} are activated at t+1. *Proof*: $c_j \in \boldsymbol{D^t}$ implies $c_j \in \boldsymbol{D_i} \cup \boldsymbol{D_{i+1}} \ldots \cup \boldsymbol{D_k}$. Since $r_l \in \boldsymbol{E^{t+1}}$, we have $|\boldsymbol{D_l} \cap \boldsymbol{D^t}| > |\boldsymbol{D^t}|/2$. Connectedness implies that if $c_j \in \boldsymbol{D_i} \cup \boldsymbol{D_{i+1}} \ldots \cup \boldsymbol{D_k}$ and $c_j \in \boldsymbol{D_l}$ then $c_j \in \boldsymbol{D_{l'}}$ for $l>l' \geq k$. Hence if $|\boldsymbol{D_l} \cap \boldsymbol{D^t}| > |\boldsymbol{D^t}|/2$ we also have $|\boldsymbol{D_{l'}} \cap \boldsymbol{D^t}| > |\boldsymbol{D^t}|/2$ for all $l>l' \geq k$, which means that $r_l \in \boldsymbol{E^{t+1}}$ implies $r_{l'} \in \boldsymbol{E^{t+1}}$ for $l>l' \geq k$.

Proposition 1b: If a contiguous interval of responses, r_i, r_{i+1}, … r_k is activated at t (and all other responses are deactivated at t), and for $l<i$, r_l is activated at t+1, then it is the case that r_{l+1}, r_{l+2}, … r_i are activated at t+1. *Proof*: The proof for this is identical to that for proposition 1a.

Proposition 2: If a contiguous interval of responses, r_i, r_{i+1}, … r_k is activated at t (and all other responses are deactivated at t), then for any p and q with $k>p>q>i$, if r_q and r_p are activated at t+1 then so is any r_l for $p>l>q$. *Proof*: $c_j \in \boldsymbol{D^t}$ implies $|\boldsymbol{E_j} \cap \boldsymbol{E^t}| \geq |\boldsymbol{E^t}|/2$. As $\boldsymbol{E_j}$ is contiguous (by connectedness), and $\boldsymbol{E^t}$ is contiguous, $\boldsymbol{E_j} \cap \boldsymbol{E^t}$ is also contiguous. Hence if $c_j \in \boldsymbol{D^t}$ it supports at least $|\boldsymbol{E^t}|/2=(k-i+1)/2$ contiguous responses in $\boldsymbol{E^t}$. Assume that $q<(k+i)/2$. If $c_j \in \boldsymbol{D_q} \cap \boldsymbol{D^t}$ then as c_j supports at least $(k-i+1)/2$ neighboring responses in $\boldsymbol{E^t}$, we must also have $c_j \in \boldsymbol{D_{q+1}}$. Hence if $|\boldsymbol{D_q} \cap \boldsymbol{D^t}| > |\boldsymbol{D^t}|/2$, as is implied by $r_q \in \boldsymbol{E^{t+1}}$, then we have $|\boldsymbol{D_{q+1}} \cap \boldsymbol{D^t}| > |\boldsymbol{D^t}|/2$, which implies that $r_{q+1} \in \boldsymbol{E^{t+1}}$. Now we can use this method again to show that $r_{q+2} \in \boldsymbol{E^{t+1}}$, and keep iterating it to show that $r_l \in \boldsymbol{E^{t+1}}$ for all $(k+i)/2 \geq l \geq q$. Now if $(k+i)/2 \geq p$ then our proof is done. If not then note that we can use the same logic as above to show that $r_l \in \boldsymbol{E^{t+1}}$for $p \geq l \geq (k+i)/2$. This then gives us our result.

Now, propositions 1a and 1b show that if a contiguous interval of responses is activated at time t then a response that does not neighbor this contiguous interval, cannot be activated at t+1 without activating all intermediate responses. Proposition 2 shows that if a contiguous interval of responses is activated at time t then this interval cannot splinter into two or more non-contiguous intervals of activated responses at t+1. Together these results imply both contiguous activation and sequential transitions.

References

Chapman, G. B., & Johnson, E. J. (1994). The limits of anchoring. *Journal of Behavioral Decision Making*, 7(4), 223-242.

Chapman, G. B., & Johnson, E. J. (1999). Anchoring, activation, and the construction of values. *Organizational Behavior and Human Decision Processes*, 79(2), 115-153.

Choplin, J. M. & Tawney, M. W. (2010). Mathematically modeling anchoring effects. *Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society*.

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, *128*(1), 3.

Johnson, J. G., & Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Review*, 112(4), 841.

Kosko, B. (1988). Bidirectional associative memories. *Systems, Man and Cybernetics, IEEE Transactions on*, 18(1), 49-60.

Mussweiler, T., & Strack, F. (1999). Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model. *Journal of Experimental Social Psychology*, 35(2), 136-164.

Mussweiler, T., & Strack, F. (2000). The use of category and exemplar knowledge in the solution of anchoring tasks. *Journal of Personality and Social Psychology*, 78(6), 1038.

Mussweiler, T., Strack, F., & Pfeiffer, T. (2000). Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility. *Personality and Social Psychology Bulletin*, 26(9), 1142-1150.

Reitsma-van Rooijen, M., & L Daamen, D. D. (2006). Subliminal anchoring: The effects of subliminally presented numbers on probability estimates. *Journal of Experimental Social Psychology*, 42(3), 380-387.

Simmons, J. P., LeBoeuf, R. A., & Nelson, L. D. (2010). The effect of accuracy motivation on anchoring and adjustment: Do people adjust from provided anchors?. *Journal of Personality and Social Psychology*, 99(6), 917.

Switzer, F. S., & Sniezek, J. A. (1991). Judgment processes in motivation: Anchoring and adjustment effects on judgment and behavior. *Organizational Behavior and Human Decision Processes*, 49(2), 208-229.

Whyte, G., & Sebenius, J. K. (1997). The effect of multiple anchors on anchoring in individual and group judgment. *Organizational Behavior and Human Decision Processes*, 69(1), 74-85.