Neural Network Models of Decision Making:

Past, Present, and Future

Eric Dimperio

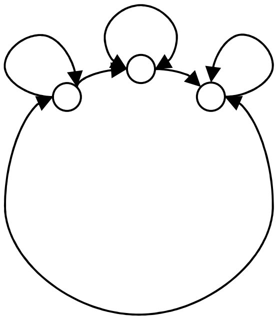

Qualifying examination essay 2 of 3.

Advisory committee:

Professor Jerome Busemeyer, Ph.D.
Professor Robert Goldstone, Ph.D.
Professor John K. Kruschke, Ph.D.
Professor Larry Yaeger

Neural Network Models of Decision Making: Past, Present, and Future

Eric Dimperio

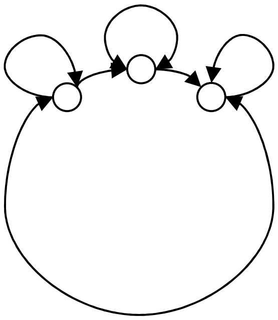

The use of neural networks as a computational tool has had a long and rocky past. When McCulloch & Pitts (1943) first laid out the calculus of the neural network, it was somewhat revolutionary in that it was a completely different method of computation than any developed to that point. The computation was not explicitly stated in axiomatic rules, but was instead imbedded in a distributed set of weights connecting nodes. Though these nodes were certainly inspired by neurons, they were not exact models of them. Just as interest was growing, Minsky and Papert (1969) proved the failure of single-layer perceptron networks to classify linearly separable patterns. Their conclusions that neural networks could not be a useful computational device even with more layers (due to added complexity), lead to a drastic reduction of interest in the subject. Over the next twenty years, significant advances in the understanding of such networks were made, but it wasn't until the publication Rumelhart & McClelland's *Parallel Distributed Processing* books (1986) that connectionist models gained renewed acceptance as a valid computational tool.

In the field of psychological decision making, neural networks have been used in a limited fashion to model human behaviors for the last two decades. By focusing on a particular field, we can explore some of the trends and developments seen in the science of artificial neural networks. The rest of this article will be organized as follows. First, I shall review the most significant neural network models developed to explain various aspects of human decision making. Next I shall look at some of the network structures used within these models and identify the benefits of such structures. I will then identify portions of neural network theory that have been left out in the decision making models. Finally, I will look toward the future by exploring how other types of neural network models can be incorporated into existing models to help develop better theories of more complex decision making.

Neural Networks in human decision making

Decision making in itself is quite a broad topic. Decisions are made to accomplish most tasks studied by cognitive scientists. However, what unites models of decision making is the description of how one action is taken over another possible action. This choice may happen under a variety of contexts. Traditional decision theory, as formalized by von Neumann & Morgenstern (1947), and Savage (1954) describes decision making as the act of maximizing the utility of a decision. The utility of an option is a sum of the values of uncertain events that may occur if chosen weighted by the expected probability of the event occurring. The values are based on subjective function of the actual outcome. These functions provide an explanation as to why we see a more significant difference between $10 and $100 than we do between $1010 and $1100. The first class of models to be examined was created to explain choice behaviors observed under such conditions of uncertainty.

One form of uncertainty arises when making risky decisions. A decision is risky if some action has multiple possible outcomes described by some probability of occurring if that action is selected. When presented with risky decisions, people do not always make rational choices as described by utility theory (Allais 1953, Tversky & Kahneman 1981). The observations of such violations of have lead to the development of new models. One significant model is Prospect theory (Kahneman & Tversky 1979). Prospect theory relies decomposes a decision into distinct editing and evaluation stages. The editing phase first places the decisions into a standard form by utilizing several heuristics. In the evaluations stage, outcomes are not seen in terms of absolute wealth, but as gains and losses from a neutral reference point. The function used to calculate utilities is convex in the region of gains and concave in the region of losses such that losses are seen as greater than an equivalent gain. In addition to the value function, a function is used to determine subjective probabilities where small probabilities are over estimated, while mid-level and high probabilities are underestimated. Prospect theory is capable of explaining why people appear risk averse when presented with a choice of gamble in the context of gains, but risk seeking when the equivalent choice is framed in terms of losses. Prospect theory has remained a dominant theory for explaining risky decision making.

Grossberg and Gutowski (1987)

Grossberg and Gutowski (1987) point out that prospect theory has two serious shortcomings and offer their own model of risky decision making. First off, although prospect theory provides a static algebraic description of choice distributions, it does not give any insight into the underlying processes involved in making a decision. Second, it was noted that prospect theory fails to account for preference reversals observed in decision experiments. An example of a preference reversal would be when presented a series of individual items, some object A was valued by the subject at $70 and object B was valued at $55, yet when presented with binary choices, object B was chosen by the subject over object A. Grossberg and Gutowski introduced their affective balance theory to account for such behaviors.

The affective balance theory can be formulated as a neural network that provides a description of the affective and cognitive events that occur throughout a risky choice. Within the network, emotional or affective states are regulated by a network structure called a *gated dipole*. The gated dipole is an opponent processing model where two alternatives compete for activation. Gated dipoles are useful for creating a habituating ON-response in the presence of a signal as well as a temporal contrast effect, where removing the signal triggers the opponent process.

Figure 1 provides a description of how a gated dipole operates. A signal is presented to an ON channel, but not to the OFF-channel. Both channels trigger the release of a transmitter whose levels decay with sustained activation. The transmitter causes habituation which establishes a reference point from which later events are compared. The channels are then summed with the inverse of the opposing channel.

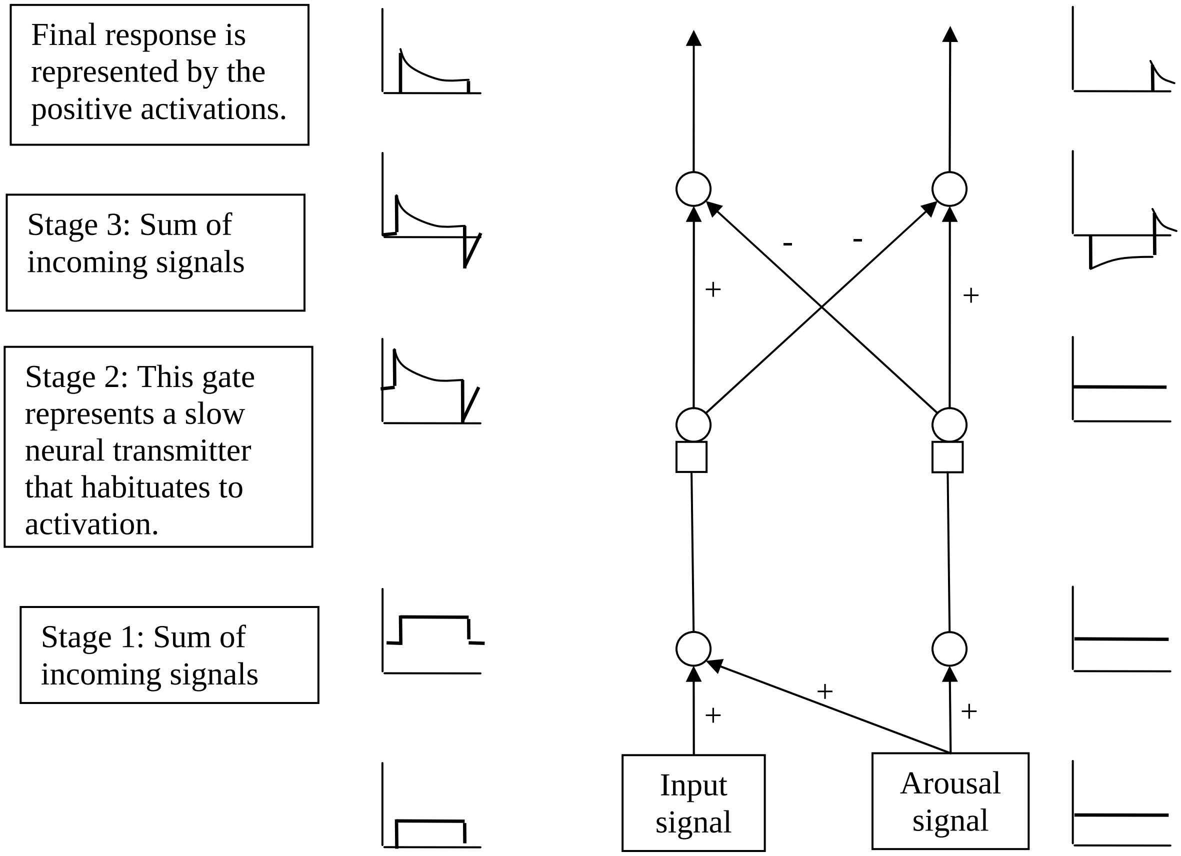

Figure 1: The structure and function of a gated dipole

Grossberg and Gutowski utilize frequency of outcomes to determine the affective processing in the model. They propose that this influence comes from both short-term and long-term effects. The long term effects represent the underlying beliefs about the values of each option. Short-term effects represent the recent experience with the options. Grossberg and Gutowski modified the basic gated dipole by having the long-term trace of the cue influence the gated signal. This influence is based on a function used in prospect theory where negatively valenced cues follow a concave curve and positively valence cues follow a convex curve. Also, the output of stage 3 is fed back as input to stage 1 to act as a short-term trace of activation as seen in figure 2.

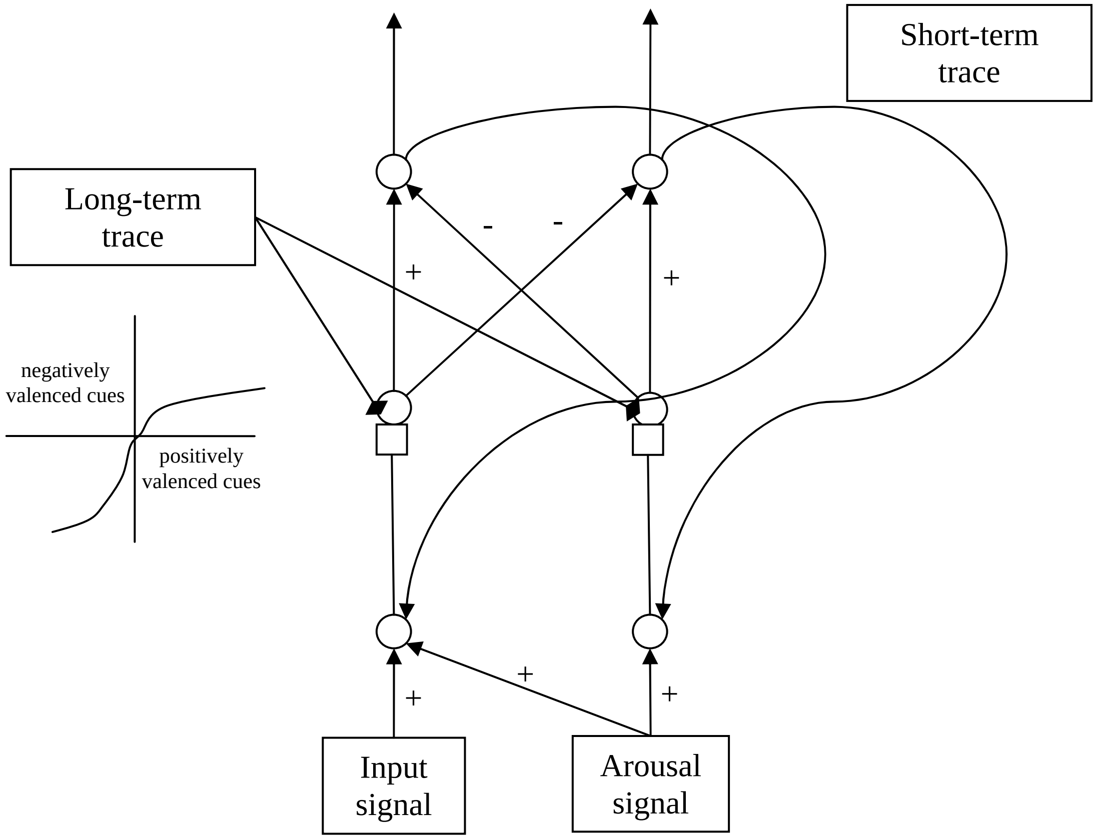

Figure 2: The Grossberg and Gutowski model of decision making under risk. It is a gated dipole with recurrent connections representing short-term and long-term memories.

Grossberg and Gutowski go on to describe the mechanisms in the model that allow for an explanation of preference reversals, framing effects, and the gambler's fallacy. Unfortunately, these effects seem to be highly dependent on the representation of the long-term effects which were not formalized in detail.

Leven and Levine (1996)

Leven and Levine (1996) also used gated dipoles in an explanation of decision making. They were specifically trying to explain the unintuitive results of their Coke example. In the mid 1980’s, the Coca-Cola tried to replace their Coca-Cola product with a new cola drink called New Coke. Market tests showed that the New Coke flavor outperformed the traditional cola as well as its competitors. Few surveyed consumers said they would be opposed to a new flavor with the same name. However, when introduced, the product was so unpopular that it had to be recalled and the original product was put back on the market as Classic Coke.

Leven and Levine took the unique approach of identifying specific properties of neural network organization and utilizing them to build a complex network that would qualitatively explain the results of the Coke example. Moreover, their model serves to be descriptive rather than prescriptive, explaining not just the results, but identifying the underlying cognitive processes that led to the rejection of New Coke. They identified the five following principles:

1) *Associative learning* – The updating of weights to make closely associated (by contiguity or probable causality) events have a stronger connection
2) *Lateral inhibition* - The activation of some event suppresses the activation of competing events
3) *Opponent processing* –The turning off of activity of some concept leads to transient activation of the opposite concept (stopping pain isn’t just neutral, but is actually pleasurable in itself)
4) *Neuromodulation* – the external manipulation of activation level based on context
5) *Interlevel resonant feedback* – the indirect association of concepts in a single level to some property via the association of one concept in that level to a property.

As mentioned, the gated dipole was utilized a key network structure in the full model. Instances of gated dipoles were utilized to represent soft drink categories (Coke, Pepsi), motivations (excitement, security), and sensory attributes (Coke label, Pepsi label, familiarity, taste). These were then connected to each other using the five principles listed above. The final model is a large network that tries to explain how the drinking of the original coke product enhances feelings for not only the coke label, but also familiarity. In the context of a taste test, the role of taste if given a higher weight and produces a high degree of excitement. However, in a context of consuming, familiarity is given greater weight to encourage strong feeling of security. This allowed New Coke to come out ahead in taste tests, but fail against competitors in the actual market.

Usher & Zakay (1993)

In 1993, Usher & Zakay introduced a new type of network into the decision making arena. Instead of having the activation of a particular node indicate a decision, they used a style of network model where states change until the network's activity converges toward a stable state for some time. This type of network is called an attractor neural network (ANN). The Usher & Zakay framework is actually a variation of the ANN called a transient attractor neural network (TANN) because the dynamics allows for transitions among multiple attractors. By using this particular structure, Usher and Zakay were able to open up neural network models to a much broader range of decision behaviors. The Grossberg and Gutowski model was able to make decisions between two opposite choices given a single input signal. The Leven and Levine model was specifically designed to choose between two cola options based on taste and familiarity. The

Usher & Zakay model framework allows for decisions between *n* actions that differ along *m* attributes. The dynamic nature of the TANN allows one to observe how the activations of different options change over time. It should be noted that unlike the previous models, the Usher & Zakay model does not explain decisions under uncertainty.

Usher & Zakay, like Leven & Levine, were motivated to create a model that did more to explain the internal dynamic processes of decision making than previous static models. They were, however, highly inspired by two previous models. The first is Tversky's elimination by aspects (EBA) model (1972) and the second is Audley's (1960) model. In fact, the Usher & Zakay framework is a general case of the elimination by aspects model. EBA is a conceptually simple model where all alternatives in a choice are described by a set of overlapping aspects. An aspect is chosen randomly according to a probability distribution determined by the strength of that aspect, and any action that is not related to the aspect is eliminated from the possible actions. This process continues until a single action is left. Audley's model allows for alternative options to be partially chosen. In Usher and Zakay's model, when a set of alternatives is active for some number of steps, a choice is made. This model was used to help account for reaction times in certain decision tasks and relates to 'degree of confidence' measures when making decisions.

According to Usher & Zakay, a decision network is made up of two parts. The first is a subnetwork of aspects and the second is a subnetwork of alternatives. The subnetworks are connected via projections from the aspects to the alternatives. A connection is made between a particular aspect / alternative pair if the alternative is characterized by that particular aspect. Each subnetwork has recurrent connections through a collection of inhibitory neurons (see figure 3). The feedback from the inhibitory assemblies allows the network to operate without input once it is activated. To allow transient states, the nodes in the aspects assemblies are provided with a dynamic threshold mechanism that acts to mimic neural fatigue. Because the alternatives do not have such a mechanism, once activated, they tend to stay activated unless the inputs from the aspects network cause strong conflicts.

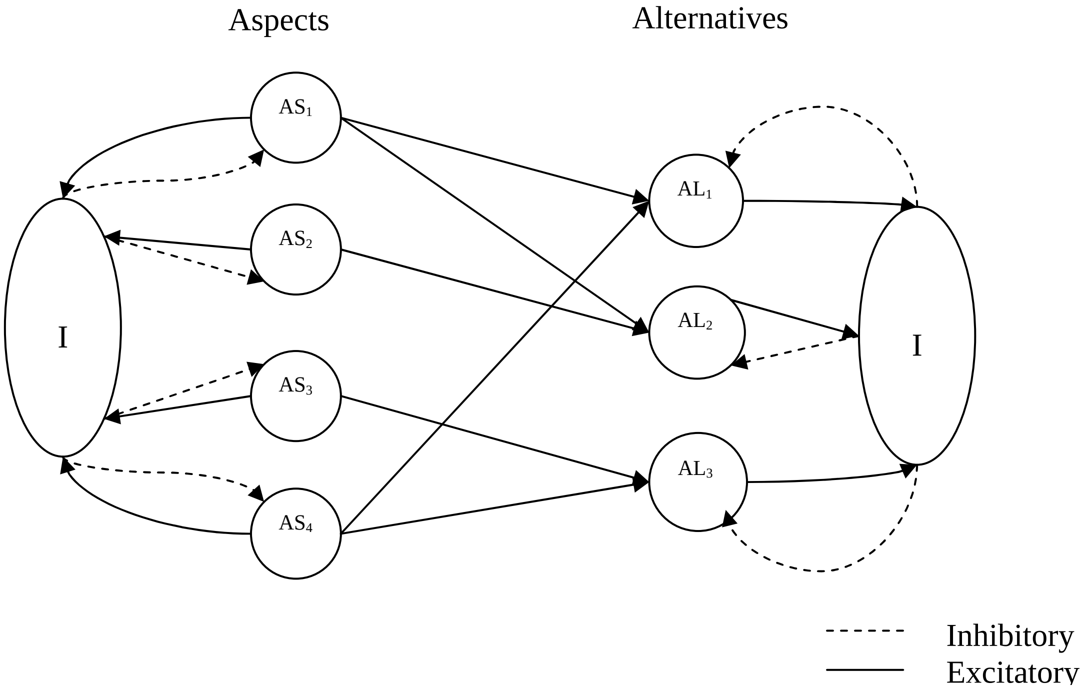

Figure 3: Usher & Zakay's connectionist decision framework. The subnetwork on the left respresents the aspects (properties) of the choices. The subnetwork on the right represents the alternatives themselves. Each subnetwork is connected to an inhibitory assembly through recurrent connections. Information flows from the aspects network to the alternatives network

The versatility of this model framework was shown by demonstrating that altering the values of the parameters permits the model to mimic specific decision strategies. Table 1 shows the significant parameters to be manipulated.

Table 1: The significant parameters that determine the heuristic functionality of the Usher & Zakay framework.

Parameter	Name	Purpose
T1	Noise factor of aspects network	Determines the competition and spread among aspects
T2	Noise factor of alternatives network	Determines the inertia of alternatives
B1	Inhibitions in aspects network	Determines competition between aspects
B2	Inhibitions in alternatives network	Determines competition between alternatives
O2	Threshold in alternatives network	Determines the external criterion for disjunctive and conjunctive rules

The following are brief descriptions of the decision strategies mimicked by the model and how the network is able to act in such a manner.

> *Focused attention on aspects* – This situation is very analogous to EBA. If the inhibition levels in the aspects subnetwork (B1) are very high, then it leads to the consideration of a single aspect at any given time. Only those alternatives with the activated aspects remain active themselves. Because of the inertia involved, no other alternatives will become active as activity progresses. Once the current aspect fatigues, a new aspect will be selected. In elimination by aspects, the probability of choosing an aspect remains constant, where fatigue lowers the chances of an aspects being chosen consecutively in the network model. Also, since mixed states are less stable than pure states, selecting an aspect that is not related to any remaining alternative will lead to collapsing on a single choice. In EBA, the alternatives would remain and a new aspect would be chosen.

> *Broad attention on aspects* – To simulate this strategy, the inhibition levels (B1) are set low such that multiple aspects are active at one time. All alternatives stimulated by the active aspects become turn on. This set up will lead to more transient states since multiple activated aspects may influence a non-active alternative. Unlike EBA, alternatives would not be permanently

eliminated. Instead, a criterion for decision must be created such that a decision is final when a single alternative is consecutively active for some time. This resemble Audley's model.

Dominance rule – Using the dominance rule, and option is chosen over a competitor if it is greater on at least one aspect and not worse on any other aspects. In the previous two decision rules, the connections between aspects and alternative were all set to 1 (all or nothing). In this situation, varying values are allowed to determine the degree to which an alternative is characterized by a particular aspect. To follow the dominance rule, a single aspect must be considered at any given time (high B1). Also, there must be greater competition between alternatives. This can be achieved by increasing the inhibition among the alternatives (B2). Finally, the inertia of the alternatives must be reduced (increase T2). When no one alternative is dominant, the network will vacillate between multiple alternatives.

Conjunctive models – A criterion (O2) is selected such that an alternative is eliminated if it does not meet the criterion on all any one aspect. This rule operates similar to the dominance rule, but instead of alternatives competing with each other (B2=0), they compete against the threshold (O2). The criterion for decision is remaining active after all aspects have been scanned.

Disjunctive models – The disjunctive models require that an option is chosen if the threshold is met by at least one aspect, while all other alternatives fail to meet the threshold on all aspects. This can be satisfied in the network by retaining the same parameters as in the conjunctive model, except for the inertia which must be increased (reduce T2). This allows an alternative to stay activated once any activated aspect meets threshold.

Elimination by least attractive aspect rule – This decision rule sequentially eliminates alternatives that have the worst over all aspect. This can be thought of as a generalization of the conjunctive rule where the threshold is gradually increased until a single alternative remains.

> *Choice by most attractive aspect rule* – Such a rule chooses the alternative with the highest valued aspect. This is a generalized version of the disjunctive rule where the threshold is initially set high and gradually reduced. The first alternative to become activated is chosen.
>
> *Lexicographic decision rule* – The lexicographic rule involves looking at the aspects in a certain prioritized order. Alternatives that are not contributed to by the active aspect are eliminated. This is just like the focused attention on aspects model, except where the self excitation of aspects is manipulated such that they are scanned in a specific order.
>
> *Addition of Utilities Rule* – This rule makes a decision by selecting the option with the highest weighted utility which is calculated by summing the contribution of each aspect to an alternative. This can be achieved in this framework by starting the network with uniform activation of the aspects, decreasing the noise in the aspects assembly (T1) and having a high competition among alternatives (B2).

The value of Usher & Zakay's model is that is shows that a single, simple framework can be responsible for varying decision strategies. These strategies are merely dependent on parameter values that may be set by personal characteristics of the decision maker as well as constraints set by environment and context.

Attractor networks are essentially a mechanism for solving constraint satisfaction problems (Read, Vanman & Miller 1997). If we treat the connections between nodes in the network as both positive and negative constraints (excitatory & inhibitory connections), the network settles into states in which any neighboring states would satisfy less constraints than the active one. Hopfield (1982) made an analogy to physical models and created an 'energy' measure on such network. The process of satisfying constraints, always tries to minimize the energy in the system. This type of process seems well suited for decision making. Much of the motivation for abandoning traditional decision theories comes from observations of people making seemingly irrational decisions. The network model developed by Usher and Zakay was highly influenced by Tversky's EBA model, and the EBA decision rule was developed to explain violations of independence from irrelevant alternatives (IIA). This 'rational' assumption holds that if A is preferred over B when C is present, A should also be preferred over B if D is present since C and D do not change the properties of A and B. People have been shown to regularly violate this assumption (Tversky 1972). On one hand, the model framework that Usher & Zakay developed is able to make 'irrational' decisions. On the other hand, it can be thought of as going through the very rational process of best satisfying the constraints on a very complex system.

Thagard & Millgram – DECO (1995)

The attractor network ECHO was developed by Thagard (1989) to explain how hypotheses are accepted in the presence of evidence. ECHO was based on the idea of achieving *coherence* in finding an explanation of observations. ECHO has been used as the inspiration for a couple of neural network models of decision making.

Thagard and Millgram's (1995) variation of the ECHO model was called the DEliberative COherence (DECO) model. Like many of those who previously used connectionist methods to model decision making, they were motivated by a desire to develop a descriptive model of the decision making process. Thagard and Millgram do take a difference approach to understanding the mechanisms in decision making than previous models. Traditional decision theory is based on the idea of preferences. People select the item they prefer the most. In decision theory, these preferences come from utilities which are determined by the alternatives presented. Thagard and Millgram make the observation that explanations of decisions that utilize utilities are less than helpful when utilities are just a summary of preferences. What does a utility mean, and where does it come from? Instead of claiming that decisions are made based on the most preferred actions, they claim that actions are chosen to meet coherence among the actions and goals involved in the development of a plan. Any decision is made to try to satisfy some underlying goal. Between the decision at hand and the final goal, there may be subgoals. For instance, if my goal is to become famous, I may be able to reach that goal if I achieve certain subgoals such as starring in a movie, or winning a political office. Those subgoals are facilitated by the choice I have at hand between moving to Los Angeles, Washington, D.C, or to Boise. Since subgoals, can in fact also be actions, they are treated the same in the model. Just as in any ANN, constraints among the factors (actions and goals) are used to determine the values of weights connecting them. DECO has six assumptions that set is apart from general attractor networks.

1) *Symmetry* – all connections between factors, whether they represent coherence or incoherence, are symmetrical relations
2) *Facilitation* (excitatory connections) – If a set of factors facilitate some goal then each factor coheres with the goal and each other factors in that set. The greater the number of factors in the set, the less total coherence there is.
3) *Incompatibility* – factors that are difficult to achieve or perform together are incoherent.
4) *Goal priority* – some goals are intrinsically desirable regardless of coherence. Such goals have the highest priority
5) *Judgment* – Coherence with judgments of factual beliefs can influence facilitations and competition relations.

6) *Decision* – A decision is made to support the overall coherence of a set of actions and goals.

The goals with the highest priority are stimulated by a constant external drive. These are the only asymmetric connections in the network. Even though actions and goals are treated the same mathematically, they have a different placement in the structure of the network. The intrinsic goals of the system are fed by the external drive. All other goals are often referred to as subgoals. And those actions which are of interest in the decision at hand are labels actions. The decision process allows one to observe the temporal dynamics of what actions are weighed back and forth. In addition, one can observe the evolution of the coherence of different plans. For instance, if the example of moving given above is a subnetwork in a more complex decision, I may begin to decide to move to Boise because it facilitates my love of potatoes which in turn makes me happy. However, once concerned with making myself happy, I may notice that becoming famous has a stronger coherence than simply loving potatoes. Moving to Los Angeles has a high coherence with making me famous and so I may settle into that choice.

Thagard and Millgram make reference to simulations they ran to demonstrate how DECO can model how human behavior regularly diverges from the axioms of utility theory. They, however, have yet to publish a description of those simulations.

Guo & Holyoak (2001)

In 2001, Guo & Holyoak developed a model of decision making based on Thagard’s ECHO. Although not expressed in the same terminology, the Guo & Holyoak model follows the same principles as those in DECO. However, Guo and Holyoak developed their model to specifically replicate some curious results in decision experiments. Specifically, they account for two affects where a decoy is added to a pair of alternatives leading to inconsistencies of choice. Let us imagine that two alternatives are presented to a decision maker where the alternatives differ along two dimensions. They are presented such that neither alternative is preferred over the other. We can expect that on about half of the trials option X is chosen while option Y is selected for the other half. The first effect of interest, the similarity effect, occurs when a third option S is presented and is very similar one option, but no option dominates any other (See figure 4). In this case, the third option highly decreases the probability of selecting X, while minimally affecting the selection of Y. Choices that originally would have gone to X seem to be split between X and S. The second effect, called the attraction effect, occurs when the decoy is again similar to but this time dominated by one of the options. In this case, the third option A is rarely selected, but it does serve to increase the choice probability of X and reduce that of Y. It seems as if X looks more attractive overall when placed next to something similar that it clearly dominates.

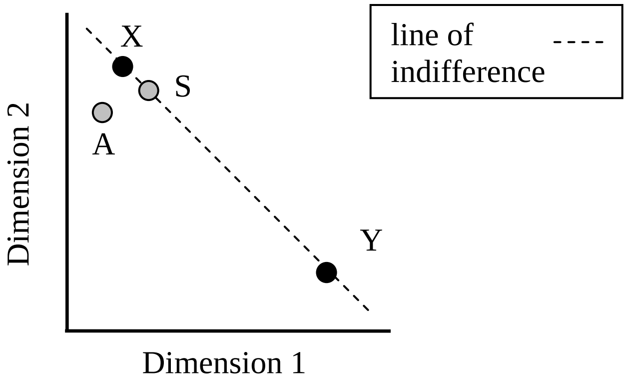

Figure 4: Several choice alternatives are displayed according to their values along two descriptive dimensions. The decoy S shown has an appropriate relationship to the alternative to demonstrate the similarity effect while A is appropriate for the attraction effect.

Using some terminology from DECO, Guo and Holyoak developed a network with two intrinsic goals (satisfy dimension 1; satisfy dimension 2) and three actions (select option X, select option Y, select option Z). Each action facilitates each goal by some degree, depending on where it lies in the aspect space. Only one option can be selected and so actions are highly incompatible with each other. This setup led to the network shown in figure 5.

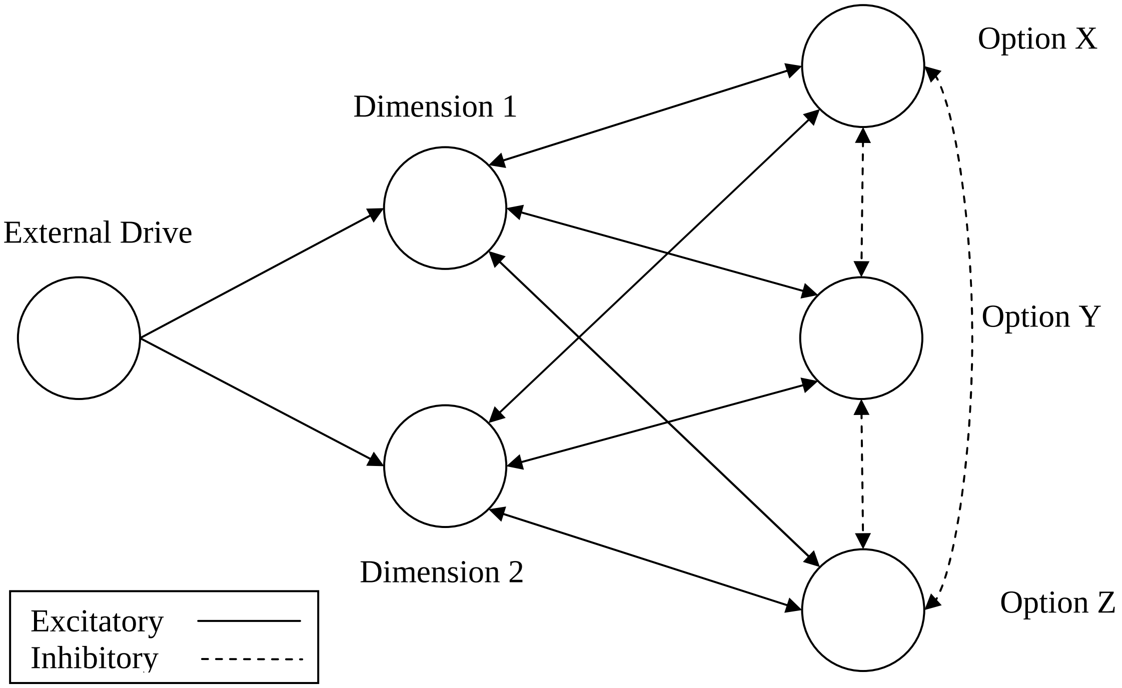

Figure 5: Guo & Holyoak's decision network based on the ECHO framework.

All activations in the network are limited between 0 and 1 and initialized to 0.5. The weights connecting the external drive to the goals, determine the priority of satisfying each dimension. The incompatibility of the alternatives is made possible through lateral inhibitory connections. Decisions are made when the system maintains a particular pattern of activation exceeds a threshold. In order to obtain the similarity and attraction affects, Guo and Holyoak impose a two part decision process. It is postulated that similar items are first compared and decided upon. Once that decision is made, then the third option is introduced into the network. The two similar items continue at their current activation levels while the third option enters the competition at the initialized value of 0.5.

To explain the similarity effect, the model first compares the similar options (X and the decoy). Since neither one is dominated, they maintain the same relative activation. However, do to the strong inhibitory links in the systems, the actual activation levels fall. Once they stabilize, the third option enters with an initial activation level of 0.5. Again, all options lie along a line of indifference (what they are lacking in dimension 1 is made up for in dimension 2) and so they maintain the same relative position (see figure 6).

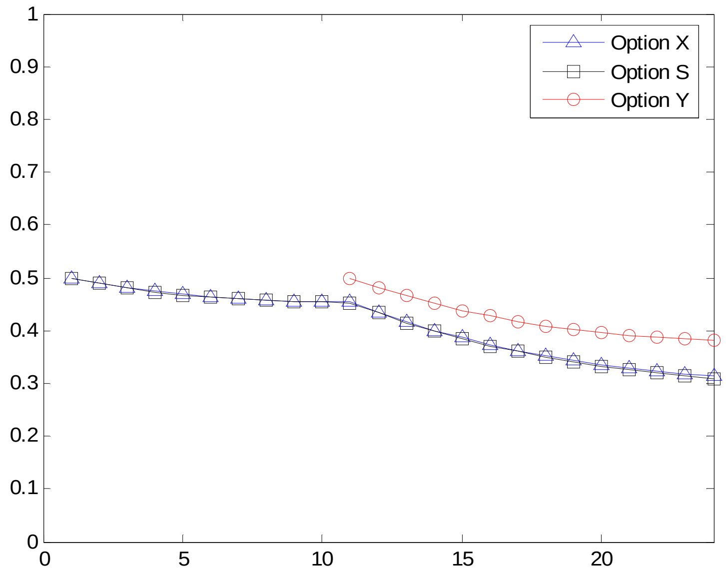

Figure 6: Activations representing the similarity effect in a 2 stage decision process

This account does not match common theories to explain the similarity effect. It is commonly treated as if the similar decoy and original option are grouped and this group is compared to the third option. Within the grouping, preference is split between the similar options. The Guo & Holyoak model suggests that some type of initial fatigue from comparing the two options reduces their relative activation strength. The third option only has a higher choice probability because is it more ‘fresh’ and has had less time to decay (see figure 7).

The explanation of the attraction effect more closely matches common theories. During the first stage when only the similar options are being compared to one another, option X dominates the decoy A. Therefore, option X receives a stronger activation and inhibits A which in turn further boosts the activation of X. Once it is already decided that X is a strong option, option Y enters the process. Unfortunately, it can never make up for the praise X received while being compared to an obviously less desirable option.

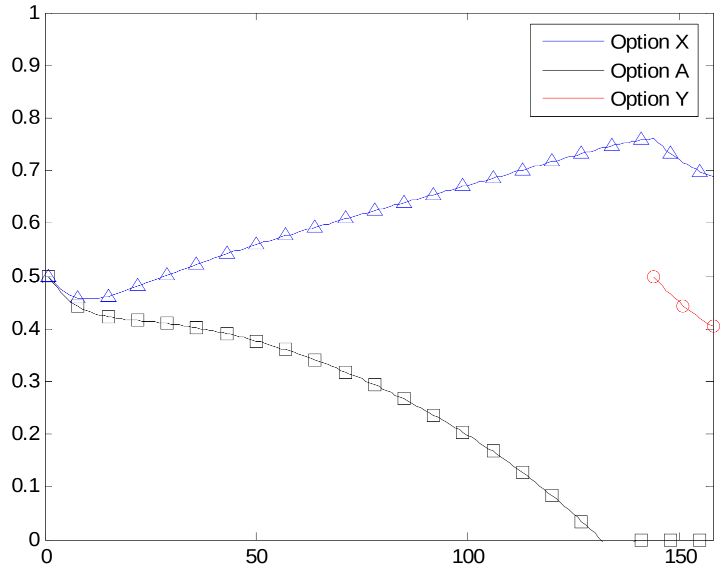

Figure 7: Activations representing the attraction effect in a 2 stage decision process

It should be noted here that there is a third effect that has appeared in decision studies (Roe, Busemeyer, & Townsend, 2001) called the compromise effect. The compromise effect arises when the decoy falls between the two options on both dimensions. Again, when X and Y are presented as a binary choice, they are equally likely to be chosen. A compromise is presented that also lies along the line of indifference. In a binary choice between C and one of the original options, no preference for any of the alternatives is shown (see figure 8). When the compromise C is presented along with X and Y, the compromise is preferred over the two original options.

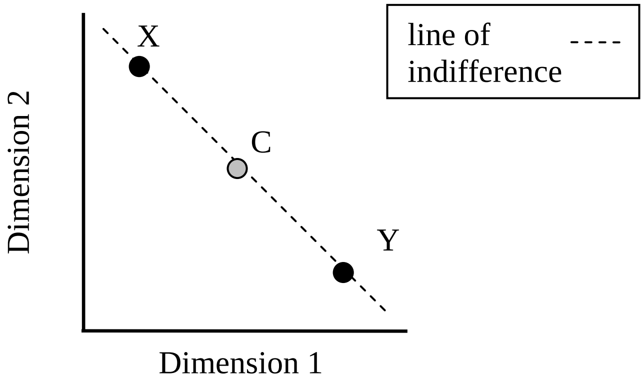

Figure 8: Several choice alternatives are displayed according to their values along two descriptive dimensions. The decoy C shown has an appropriate relationship to the alternative to demonstrate the compromise effect.

Guo & Holyoak did not mention this effect in their paper, though they must have been aware of it. This effect can not be achieved using the two stage process suggested because no pair stands out as being very similar and therefore subject to an initial comparison. However, if you alter the requirements for initial deliberation to be a comparison of the pair with the most extreme similarity or difference, the two options X and Y would be subject to the stage 1 comparison. The process would follow just as it did in the similarity effect and yield the compromise effect (see figure 9).

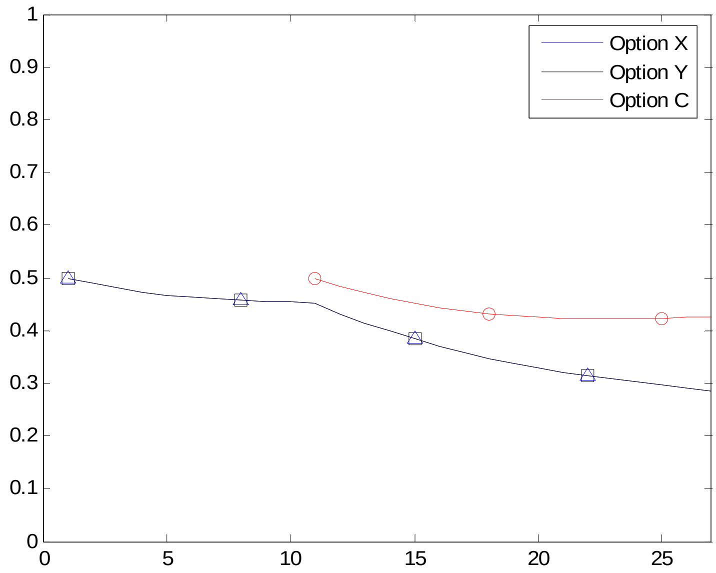

Figure 9: Activations representing the compromise effect in a 2 stage decision process

Although the Guo & Holyoak model seems to explain the similarity effect, the attraction effect, and possibly the compromise effect, the timing of the dynamics have not been supported by measures of human reaction times. This becomes especially difficult when decision makers are forced to make decisions on a very short time scale since the two stage process seems to eliminate one of the options from early consideration.

Roe, Busemeyer, & Townsend (2001)

Roe, Busemeyer, & Townsend (2001) developed a neural network instantiation of the decision field theory (Busemeyer & Townsend, 1993) model of decision making. Decision field theory was originally formulated to provide accurate predictions of reaction times as well as choice behaviors. Nonetheless, within this essay, references to MDFT are specifically referring to the neural model described by Roe, Busemeyer, & Townsend (2001). Just as with the Guo & Holyoak model, MDFT provides a model of making decisions among multiple alternatives defined along multiple attributes.

I shall first describe MDFT and then compare it to networks previously discussed. Although, visually the networks appear quite similar (see figures 5 & 10), there are some significant differences between MDFT and the ECHO inspired model of Guo and Holyoak. First of all, MDFT does not act as an attractor network. Instead of settling into a state, the activations of the option nodes represent preference states. These preferences continue to change throughout the decision until one of them crosses some threshold. At that point, the action with the preference above threshold is chosen. Unlike the ECHO inspired models, connections in MDFT are not always symmetrical. In fact the only recurrent connections are those present among the lateral inhibition of the preference states. Moreover, rather than assuming uniform inhibition, MDFT relies on inhibition strengths that are dependent on the similarity of the alternatives. Alternatives that are quite similar will compete with each other more and will therefore strongly inhibit each other.

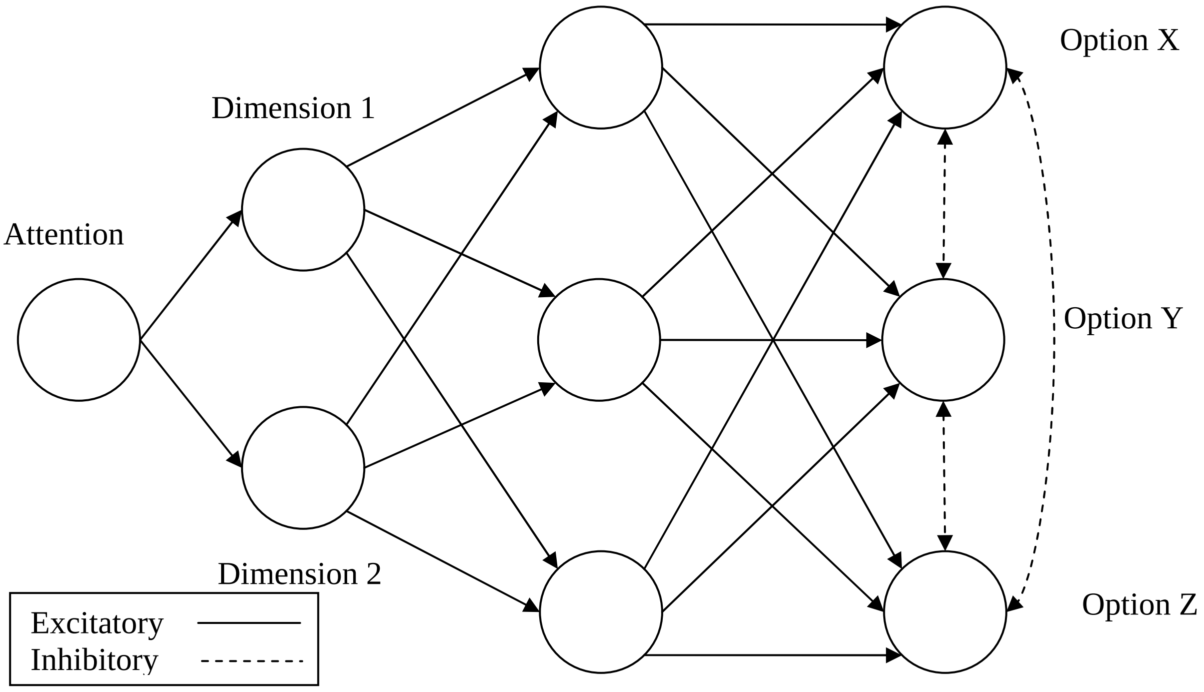

Figure 10: The multiatribute decision field theory

Thagard & Millgram’s DECO model as well as Guo & Holyoak’s model relies on an external drive to provide constant stimulation of the attributes. The strength of this activations depended on the *priority* or the relative importance of satisfying that attribute over the others. In MDFT, the external driver is seen as an attention mechanism. Rather than provide a constant activation, the attention mechanism stochastically selects a single attribute to activate. The probability of an attribute being selected on any iteration of the network is similarly dependent on the relative importance of satisfying that attribute.

The final difference is the introduction of an additional layer in the network. The DECO model claimed that judgments regarding the beliefs of factual statements will influence the coherence in a model. The nodes in the additional layer act as a set of judgments of how much one believes a particular action is going to lead to a desirable outcome. Roe et al. (2001) refers to these instantaneous judgments as momentary valuations. Even though the MDFT model can be described in comparison to the DECO model, it is more true to the history of the development of the model to look at the equations underlying the network structure.

The decision field theory model describes a stochastic diffusion process. The values of the preferences stochastically shift values until one exceeds a threshold. The linearity of the system imposes the restriction that all preference values must sum to zero. The benefits of the linearity of the system are 1) the simplicity of mathematically expressing the model and 2) the solvability of the equations. The momentary valuations give rise to valences that are calculated as

$$V(t) = CMW(t)$$

Where $W(t)$ is the attention vector (remember that at any point in time, attention is focused on a particular attribute). The matrix M represents valuations of how each attribute contributes to each alternative. Finally C is a contrast matrix that is used to normalize the values such that they sum to zero. The final preferences are calculated utilizing the following equation:

$$P(t+1) = SP(t) + V(t+1)$$

The matrix S contains both the values for the distance dependent linear inhibition and the self decay.

Even though this network contains stochastic processes, determining its behavior under certain parameters does not require calculating statistics based on many different simulations. This linear system allows for the distributions of both choice behaviors and reaction times to be calculated directly. Beyond ease of use, MDFT adds more explanatory value as compared to the Guo & Holyoak model. MDFT does not require a complex two-stage process. Its single process predicts the dynamics of shifting preferences from beginning to end. This allows the testing of effects brought about by forcing decisions after a short amount of time. If fact, MDFT has been used to model such results (Diederich, 2003). The model induces a specific speed / accuracy trade off (Busemeyer & Townsend, 1993) primarily due to the all-or-nothing aspect of the attention mechanism. As time increases, the proportion of times any attribute is sampled approaches the proportion of attention weight given to that attribute. However, early in deliberation when attention has shifted relatively few times, the stochastic nature of the model may cause one attribute to dominate the decision choice. When the decision stopping rule is not forced, people vary in how long they take to make decisions. Nonetheless, MDFT has been able to accurately predict distributions of time taken to make decisions along with the choice probabilities for individual decision makers.

The Guo & Holyoak model was designed to give a functional description of processes that lead to the similarity effect and the attraction effect. With a slight change in how the options are paired for the first stage of the model, the ECHO based model can also provide an explanation of the compromise effect. Decision field theory also can be used to explain all three of those effects. However, it uses different mechanisms to explain the effects.

Similarity effect – With the similarity effect, a decoy is introduced into a pair of competing, non-similar alternatives (X & Y) such that the decoy is similar to, but not dominated by one of the original alternatives (X). Although the two alternatives may have been equally preferred (each chosen 50% of the time) when given a binary choice, when the decoy is present, alternative Y is chosen more often than X. let’s assume the alternatives can be described along 2 dimensions and alternative X is strong on the second dimension and weak on the first dimension, while option Y is strong along dimension 1 and weak on the second dimension. The decoy, therefore, is also strong on dimension 2 and weak on the first (see figure 4). While the attention mechanism is focusing on the second dimension, then option X and the decoy are seen as advantageous (their valence values are greater) and must compete with each other. When attention is on dimension 1, alternative Y is seen as advantageous, but does not have to compete with any alternatives with a similar strength.

Attraction effect – The attraction effect occurs when the decoy is similar to alternative X, but is dominated by X. Let us assume the same major strengths and weaknesses as those listed above. However, in this case, let us assume the decoy is obviously weaker than X along dimension 2 and matches X on dimension 1 (see figure 4). Here, the decoy serves to actually increase the probability of choosing option X. This violates the regularity principle of many random utility models. According to the principle of regularity, adding options to an existing set should only allow choice probabilities to decrease or remain the same. Remember that MDFT normalizes the valences to sum to zero. As attention shifts back and forth between the dimensions, the relative strengths and weaknesses of X and Y even out (assuming equal attention weights), but the decoy remains weaker than the other two. This leaves the decoy with a negative valence (on average). MDFT utilizes distance dependent strengths on the lateral inhibition strengths. The negative activation of the decoy filtered through the relatively strong ‘inhibition’ of option X actually boosts the preference of the X alternative (negative x negative = positive). Here the distance dependent inhibitory connections are responsible for the production of the attraction effect

Compromise effect – We observe the compromise when the decoy is between the two original alternatives on all dimensions (see figure 5). All three options are considered equally attractive. When provided a binary choice between any two alternatives from the set, decision makers would have a 50% probability of choosing any particular option. Nonetheless, when all three options are presented, the compromise (in this case the decoy) is chosen more than either of the two original alternatives. This means that during the trinary choice, the mean valences are equal to zero. The distance dependent inhibitory connections in MDFT are responsible for the production of the compromise effect through their influence on the momentary fluctuations of the valences. The shorter distance between the compromise and the two extremes provides a stronger inhibition than that which exists between the two extremes. This leads to the compromise being negatively correlated with the other two alternatives while X and Y are correlated with each other. That means when the compromise is momentarily seen as advantageous, it is not competing with any alternatives, while X & Y compete with each other. This leads to the higher probability of choosing the compromise.

Usher & McClelland (2004)

In 2004, Usher & McClelland published a model that also represents a diffusion process and is based on their leaky competing accumulator (LCA) model of perceptual identification. The Usher & McClelland model is very similar in spirit to MDFT, but addresses several assumptions. These addresses lead to three significant differences between the two models. The first difference is the lack of negative activations. The linearity imposed by the MDFT model allows nodes to become negatively activated. Usher & McClelland (2004) feel that limiting the model to have only positive activations lends itself to greater biological plausibility. The tradeoff is the need to run Monte Carlo simulations to determine choice probabilities and decision times since the nonlinear system cannot be solved directly. The second major difference is the lack of distance dependent lateral inhibition. In the LCA model, all lateral inhibition is of constant value. The distance dependent lateral inhibition was responsible for both the attraction and compromise effects. The LCA model is able to account for these effects using the third major difference: a loss aversion function that weights losses as being greater than gains. This type of function was utilized in Kahneman & Tversky’s (1979) prospect theory. Usher & McClelland point out that there is such a long list of research supporting loss aversion effects in decision making that it should remain an integral part of decision models. The LCA model is able to account for the three major effects that MDFT was able to explain although it has slightly different mechanisms for attaining these effects.

Structural analysis

Table 2 summarizes some of the components and mechanisms used in neural networks and how the models we have looked at use such mechanisms.

Network Features		Network creators						
		Grossman & Gutowski	Leven & Levine	Usher & Zakay	Thagard & Millgram	Guo & Holyoak	Roe, Busemeyer, & Townsend	Usher & McClelland
General	Attractor network			x	x	x		
	Stochastic attention to attributes			x			x	x
	Opponent Processing	x	x					
Learning	Associative learning	x	x					
Recurrent Connections	symmetric connections				x	x		
	lateral inhibition			x	x	x	x	x
	recurrent excitations	x	x		x			
External Controls	Neuro-modulation	x	x					x
	Multi-stage processing					x		
Structural	Distinct levels of processing	x	~	x		x	x	x
	Interlevel Resonant feedback		x	x	?			
	Network of networks		x		x			

The earlier models discussed that were based on the gated dipole used associative learning mechanisms to interact with some external signal. In the general field of neural networks, a large portion of the research has been focused on developing mechanisms that allow connectionist structures to learn strengths of connections between nodes (Haykin, 1998). Learning processes are generally separated into supervised and unsupervised learning. Supervised learning involves providing some sort of feedback regarding desired output in a given context. Unsupervised learning involves the adaptation to some stimulus without any feedback regarding the correctness of that adaptation.

Cognitive psychology has made significant strides in understanding human learning. Learning researchers are also using connectionist frameworks to design models of human learning. In particular, the field of category learning has had some significant successes using neural network learning mechanisms to classify stimuli. The ALCOVE model (Kruschke, 1992) uses backpropagation to learn which attributes of a stimulus to attend to in order to make a proper classification. ALCOVE’s success has led to many variations that have been adapted to a myriad of related tasks. The category tasks are those that provide trial by trial feedback just like the algorithms used for supervised learning.

In many real world dynamic tasks, feedback may either be delayed or it may be based on a history of responses instead of just a single response. One experimental task that is often used to represent such a dynamic decision making situation is referred to as the sugar production factory (SPF) task (Gibson, Fichman, & Plaut, 1997). In each trial of the task, subjects are asked to select the number of workers that should go to work at the factory on a given day. They are provided feedback as to what the factory’s production level was for the day and their goal is to try to maximize sugar production. The feedback is based on not only the current response, but also on the previous day’s production level and is calculated as $P_{t+1} = 2W_{t+1} - P_t + \varepsilon$. Gibson Fichman & Plaut (1997) derived a unique two level feed forward network where the forward model learns how actions taken influence actual outcomes and the action model learns which actions to take to meet desired outcomes. Unfortunately, this type of structure does not solve a general problem. The extra level allows the decision process to consider effects that go back 1 time step. In the real world we make choices based on actions and feedback that span a very long history.

Most of the decision making models described in this essay utilize some sort of recurrent connection. Recurrent connections can provide a memory of any n-steps back (Haykin, 1998), or even an infinitely long memory (Doya, 2002). It hasn’t been until recently, however, that both methods of performing gradient decent learning on recurrent networks and the computing power to perform such algorithms have been available. Although most progress has been in the supervised learning arena where temporal arrays of input-output mappings act as training sets, there are algorithms for reinforcement learning (single reward signal) and unsupervised learning (learns the statistics of the input signal) (Doya, 2002). A general outline of how the MDFT network may get feedback to engage in some sort of supervised learning has been proposed (Busemeyer, Dimperio, & Jessup, 2005) but not formalized. Integrating formal models of learning into complex, dynamic models of decision making will permit the modeling of how people develop the knowledge that is used in making judgments. Although people do show poor performance in short-term dynamic decision making tasks where feedback is delayed, we know that people with specialized skills learn under such conditions (Gibson Fichman & Plaut, 1997). Exploring models in this domain will help expand our knowledge of how skill learning and decision making are intertwined.

Functional analysis

In addition to understanding the differing structures of these models, it is important to compare the functional aspects of these models. I will revisit some of the behavioral effects that some models have been shown to explain. One of the first effects mentioned was that of observed preference reversals when comparing prices offered to buy alternatives and choices between alternatives. The Grossberg & Gutowski (1987) model of affective balance theory was designed to explain such an effect. Their model was based on Prospect theory (Kahneman & Tversky, 1979) which is not capable of explaining such reversals. However, since their model was designed to explain binary, single attribute decisions between risky alternatives, other context effects that involve trinary choices cannot be explained. The Leven and Levine model (1996) is a complex variation of the gated-dipole network used by Grossberg & Gutowski. It was used specifically to explain results in their New Coke example and was not expressed as a general decision making network. It is likely that it is limited to merely explaining preference reversals like that observed in the New Coke example. Because the other models were designed to make explicit choices and not provide pricing values, they cannot explain these preference reversals. It should be noted, however, that the Decision Field Theory models was used as a component in a model of price selection (Johnson & Busemeyer, 2005) that does explain such preference reversals.

The three context effects discussed were the similarity, attraction, and compromise effects. Although the Usher & Zakay (1993) model was not specifically used to explain such effects, we can infer certain behaviors. The model was highly influenced by the EBA model (1972) and under certain parameter values mimic the operation of EBA. EBA was able the explain the similarity effect but held the property of regularity and could not explain attraction or compromise effects which represent violations of regularity. The Usher & Zakay model should have no problems producing a similarity effect, but it is not clear if regularity holds. Future simulations will be required to make a conclusive argument.

Guo & Holyoak's extension of the ECHO model (2002) was explicitly shown to produce similarity and attraction effects. It utilized a multistage method where similar items were first compared and then the third option of introduced. As mentioned previously, as the model stands it is unable to explain the compromise effect since no pair stands out as being most similar. If, however, the first stage is based on absolute values of differences from the average similarity, then the extremely different options in the compromise situation will be evaluated first. This will lead to the compromise effect using the same mechanisms that lead to a similarity effect.

Both the DFT and LCA models have been shown to explain the three context effects, but they rely on different mechanisms to do so. To yield a similarity effect, DFT and LCA rely on an attention switching mechanism that causes similar items to compete when their most significant attributes are attended. When less significant attributes receive attention, the preference for non-similar alternative is augmented and does not have to compete with other alternatives.

The attraction effect is reached in DFT by relied on the lateral inhibition. As in the similarity effect, similar items compete, but in the case of the attraction effect, one alternative dominates the other. When the significant attribute receives attention, the dominant alternative will have a greater valence and will yield a greater preference. The lateral inhibition from the dominant option will drive the competing option to a negative value. The lateral inhibition from the dominated option will enhance (negative x negative = positive) the preference for the dominant option, thereby producing the attraction effect. The LCA model produces the attraction effect from the direct comparison of gains and losses. When the decoy A (figure 4) is added, it is adding two small gains to option X and increasing its choice probability while adding a large gain and a large loss to Y. The loss function is steeper than the gain function and the drop in Y’s choice probability comes from the greater effects of loss. The same mechanism explains the compromise effect in LCA. When the decoy C is added (figure 8), significant gains and losses are felt by X & Y equally. Again the losses dominate and choice probabilities for both X & Y decrease. DFT relies on its distance dependent lateral inhibition to negatively correlate the compromise and the two extreme options. Although the valence for the compromise will have a mean of zero, the stochastic error will tip the compromise positive about half of the time. When it goes positive, it beats out the other two options, but when it goes negative, the other two must compete with each other for the race to threshold. For this reason, the compromise is chosen more often than the other two.

References

Allais, P. M. (1953). The behavior of rational man in risky situations - A critique of the axioms and postulates of the American School. *Econometrica, 21*, 503-546.

Audley, R. J. (1960). A stochastic model for individual choice behavior. *Psychological Review, 67*, 1-15.

Busemeyer, J. R., Dimperio, E., & Jessup, R. K. (2005). Integrating emotional processes into decision making models. To appear in W. D. Gray (Ed.), *Integrated Models of Cognitive Systems*. Oxford University Press.

Busemeyer, J.R., Jessup, R. K., Johnson, J.G., & Townsend, J. T. (2006) Building bridges between neural models and complex human decision making behavior. To appear in *Neural Networks*.

Busemeyer, J. R. & Townsend, J. T. (1993). Decision Field Theory: A dynamic cognition approach to decision making. *Psychological Review*, 100, 432-459.

Diederich, A. (2003). MDFT account of decision making under time pressure. *Psychonomic Bulletin and Review, 10*(1), 157-166.

Doya, K. (2002). Recurrent networks: learning algorithms. In M. A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks,* Cambridge, MA: MIT Press.

Gibson, F. P., Fichman, M., & Plaut, D.C. (1997). Learning in Dynamic Decision Tasks: Computational Model and Empirical Evidence. *Organizational Behavior and Human Decision Processes, 71*(1), 1-35.

Grossberg, S. & Gutowski, W. E., (1987). Neural dynamics of decision making under risk: Affective balance and cognitive-emotional interactions. *Psychological Review, 94*(3), 300-318.

Guo, F. Y., & Holyoak, K. J. (2002). Understanding similarity in choice behavior: A connectionist model. *Proceedings of the Cognitive Science Society Meeting*.

Haykin, S. (1998). *Neural Networks: A Comprehensive Foundation*. Upper Saddle River, NJ: Prentice Hall.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*, 2554–2558.

Johnson, J. G. & Busemeyer, J. R. (2005) A dynamic, computational model of preference reversal phenomena. *Psychological Review, 112*(4), 841-861

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, *47,* 263-291.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22-44.

Leven, S. J., & Levine, D. S. (1996). Multiattribute decision making in context: A dynamic neural network methodology. *Cognitive Science*, *20*, 271-299.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in neural nets. *Bulletin of Mathematical Biophysics*, *5*, 115-137.

Minsky , M. L., & Papert, S. (1969). *Perceptrons: an introduction to computational geometry*. Cambridge, MA: MIT Press.

Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionism, parallel constraint satisfaction and gestalt principles: (Re)introducing cognitive dynamics to social psychology. *Personality and Social Psychology Review*, *1*, 26-53.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multi-alternative decision field theory: A dynamic connectionist model of decision-making. *Psychological Review, 108*, 370-392.

Rumelhart. D. E. & McClelland, J. L (1986). *Parallel Distributed Processing,* Vols. *1 and 2.* Cambridge, MA: MIT Press.

Savage, L. J. (1954). *Foundations of Statistics*. Oxford: Wiley.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*, 435–467.

Thagard, P., & Millgram, E. (1995). Inference to the best plan: A coherence theory of decision. In A. Ram & D. B. Leake (Eds.), *Goal-driven learning* (pp. 439-454). Cambridge, MA: MIT Press.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79*(4), 281-299.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*(4481), 453-458.

Usher, M. & Zakay, D. (1993). A neural network model for attribute-based decision processes. *Cognitive Science, 17*(3), 349-396.

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review, 111*(3), 757-769.

von Neumann, J., & Morgenstern, O. (1947) *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.