

Human Unsupervised and Supervised Learning as a Quantitative Distinction

Todd M. Gureckis and Bradley C. Love
Department of Psychology
University of Texas at Austin
`{gureckis,love}@love.psy.utexas.edu`

January 31, 2003

Abstract

SUSTAIN (Supervised and Unsupervised STRatified Adaptive Incremental Network) is a network model of human category learning. SUSTAIN initially assumes a simple category structure. If simple solutions prove inadequate and SUSTAIN is confronted with a surprising event (e.g., it is told that a bat is a mammal instead of a bird), SUSTAIN recruits an additional cluster to represent the surprising event. Newly recruited clusters are available to explain future events and can themselves evolve into prototypes/attractors/rules. SUSTAIN has expanded the scope of findings that models of human category learning can address. This paper extends SUSTAIN to account for both supervised and unsupervised learning data through a common mechanism. The modified model, uSUSTAIN (unified SUSTAIN), is successfully applied to human learning data drawn from Love (2002) that compares unsupervised and supervised learning performance.

1 Introduction

Categories provide a crucial function underlying the cognitive abilities of humans. They allow us to generalize our knowledge to novel situations and to infer unknown properties of the environment. These abilities are indispensable to any intelligent system.

Researchers studying human categorization have traditionally focused on human performance in supervised learning tasks (see Ashby, Quellar, and Berretty, 1999; Billman and Knutson, 1996; and Clapper and Bower, 1994, for some exceptions). In this experimental paradigm, subjects learn to classify stimuli as members of contrastive categories through trial by trial learning with corrective feedback. Theories (and models) of learning are favored that can account for the relative difficulty of acquiring different category structures (Nosofsky et al., 1994; Shepard, Hovland, and Jenkins, 1961).

Although classification learning does capture aspects of human learning, others are not addressed by this paradigm. For instance, humans can spontaneously construct categories in the absence of feedback. As an example, many of us have created the categories “interesting” email and “junk” email in the absence of explicit feedback. Such learning is referred to as unsupervised learning. Supervised and unsupervised learning are often seen as being qualitatively different. Supervised learning is characterized as intentional, in that learners actively search for rules (perhaps by hypothesis testing) and are explicitly aware of the rule they are considering (Nosofsky, Palmeri, and McKinley, 1994). On the other hand, unsupervised learning is seen as an incidental, undirected, stimulus driven, and incremental accrual of information (Berry and Dienes, 1993; Cleermans, 1993; Hayes and Broadbent, 1988; Hock, Malcus, and Hasher, 1986; Lewicki, 1986).

In contrast to this view, Love (2002) has found that intentional unsupervised learning performance is more similar to supervised learning performance than it is to incidental unsupervised learning performance. This result suggests that the unsupervised/supervised dichotomy may not be valid. Gureckis and Love (2002b) have argued that unsupervised and supervised learning can be modeled through a common mechanism. Here, we apply Gureckis and Love’s (2002b) variant of the SUSTAIN (Supervised and Unsupervised STRatified Adaptive Incremental Network) model, referred to as uSUSTAIN (unified SUSTAIN), to the Love (2002) data. uSUSTAIN differs from other models that seek to unify unsupervised and supervised learning, such as Anderson’s (1991) rational model, in that uSUSTAIN is applicable to both unsupervised and supervised learning tasks while not predicting that these tasks lead to equivalent performance (which they do not). In the remainder of this paper, we overview SUSTAIN and uSUSTAIN. We then fit uSUSTAIN to the Love (2002) data and consider the implications of the simulations.

2 The Modeling Approach: SUSTAIN and uSUSTAIN

SUSTAIN has been successfully applied to an array of challenging human data sets spanning a variety of category learning paradigms including classification learning (Love and Medin, 1998b), learning at different levels of abstraction (Love and Medin, 1998a), inference learn-

ing (Love, Markman, and Yamauchi, 2000), and unsupervised learning (Gureckis and Love, 2002c; Love, Medin, and Gureckis, 2002).

In the following sections, we discuss SUSTAIN’s operation, its underlying principles, and the mathematical equations that follow from these principles. We then introduce a modification to SUSTAIN that enables it to account for supervised and unsupervised learning data through a single recruitment mechanism. This mechanism makes use of an intuitive and general notion of *surprise* to facilitate learning. This modified version of SUSTAIN is referred to as uSUSTAIN.

2.1 Overview of Model

SUSTAIN is a clustering model of human category learning. The model takes as input a set of perceptual features that are organized into a series of independent feature dimensions. The internal representations in the model consist of a set of clusters. Categories are represented in the model as one or more associated clusters. Initially, the network only has only one cluster that is centered upon the first input pattern. As new stimulus items are presented, the model attempts to assign new items to an existing cluster. This assignment is done through an unsupervised procedure based on the similarity of the new item to the stored clusters. When a new item is assigned to a cluster, this cluster updates its internal representation to become the average of all items assigned to the cluster so far. However, if SUSTAIN discovers through feedback that this similarity based assignment is incorrect, a new cluster is created to encode the exception. Classification decisions are ultimately based on the cluster to which an instance is assigned. In unsupervised learning, all items are assumed to belong to a single global category, thus prediction errors cannot occur. In this situation, SUSTAIN recruits a new cluster when the current stimulus is not sufficiently similar to any existing cluster.

2.2 The Key Principles of SUSTAIN

With this general understanding of the operation of the model in mind, we now examine the five key principles of SUSTAIN. These principles highlight the important features of the model and provide the foundation for the model’s formalism.

2.2.1 Principle 1, SUSTAIN is biased towards simple solutions

SUSTAIN is initially directed towards simple solutions. At the start of learning, SUSTAIN has only one cluster which is centered on the first input item. It then adds clusters (i.e., complexity) only as needed to accurately describe the category structure. Like other models of category learning (e.g. Kruschke, 1992), SUSTAIN learns to selectively attend to stimulus feature dimensions that are most useful for categorization. This focus on a subset of stimulus dimensions also serves to bias SUSTAIN towards simple solutions.

2.2.2 Principle 2, similar stimulus items tend to cluster together

In learning to classify stimuli as members of two distinct categories, SUSTAIN will cluster similar items together. For example, different instances of a bird subtype (e.g., sparrows) could cluster together and form a sparrow cluster instead of leaving separate traces in memory for each instance. Clustering is an unsupervised process because cluster assignment is done on the basis of similarity, not feedback.

2.2.3 Principle 3, SUSTAIN learns in both a supervised and unsupervised fashion

In learning to classify the categories “birds” and “mammals”, SUSTAIN relies on both unsupervised and supervised learning processes. Consider a learning trial in which SUSTAIN has formed a cluster whose members are small birds, and another cluster whose members are four-legged mammals. If SUSTAIN is subsequently asked to classify a bat, it will initially predict that a bat is a bird on the basis of overall similarity (bats and birds are both small, have wings, fly, etc.). Upon receiving feedback from the environment (supervision) indicating that a bat is a mammal, SUSTAIN will recruit a new cluster to represent the bat as an exception to the mammal category. The next time SUSTAIN is exposed to the bat or another similar bat, SUSTAIN will correctly predict that a bat is a mammal. This example also illustrates how SUSTAIN can entertain more complex solutions when necessary through cluster recruitment (see Principle 1).

2.2.4 Principle 4, the pattern of feedback matters

As the example used above illustrates, feedback affects the inferred category structure. Prediction failures result in a cluster being recruited, thus different patterns of feedback can lead to different representations being acquired. This principle allows SUSTAIN to predict different acquisition patterns for different learning modes that are informationally equivalent but differ in their pattern of feedback. The learning conditions in the Love (2002) study considered in this paper are informationally equivalent, but differ in their pattern of feedback.

2.2.5 Principle 5, cluster competition

Clusters can be seen as competing explanations of the input. The strength of the response from the winning cluster (the cluster the current stimulus is most similar to) is attenuated in the presence of other clusters that are somewhat similar to the current stimulus (see Sloman’s (1997), account of competing explanations in reasoning).

2.3 Mathematical Formulation of SUSTAIN

This section of the paper explains how the general principles that govern SUSTAIN’s operation are implemented in an algorithmic model.

2.3.1 Input Representation

Stimuli are represented in the model as vector frames where the dimensionality of the vector is equal to the dimensionality of the stimuli. The category label is also included as a stimulus dimension. Thus, stimuli that vary on three perceptual dimensions (e.g., size, shape, and color) and are members of one of two categories would require a vector frame with four dimensions. A four dimensional binary-valued stimulus (three perceptual dimensions plus the category label) can be thought of as a four character string (e.g., **1 2 1 1**) in which each character represents the value of a stimulus dimension. For example, the first character could denote the size dimension with a **1** indicating a small stimulus and a **2** indicating a large stimulus.

Of course, a learning trial usually involves an incomplete stimulus representation. For instance, in classification learning all the perceptual dimensions are known, but the category label dimension is unknown and queried. After the learner responds to the query, corrective feedback is provided. Assuming the fourth stimulus dimension is the category label dimension, the classification trial for the above stimulus is represented as **1 2 1 ?** \rightarrow **1 2 1 1**.

On every classification trial, the category label dimension is queried and corrective feedback indicating the category membership of the stimulus is provided. In contrast, on inference learning trials, subjects are given the category membership of the item, but must infer an unknown stimulus dimension. Possible inference learning trials for the above stimulus description are **? 2 1 1** \rightarrow **1 2 1 1**, **1 ? 1 1** \rightarrow **1 2 1 1**, and **1 2 ? 1** \rightarrow **1 2 1 1**. Notice that inference and classification learning provide the learner with the same stimulus information after feedback (though the pattern of feedback varies).

Unsupervised learning does not involve informative feedback. In unsupervised learning, every item is considered to be a member of the same global category. Thus, the category label dimension is unitary valued and uninformative for differentiating between stimuli. However, the degree to which any particular stimulus activates this category dimension indicates the degree to which the network recognizes the stimulus.

In order to represent a nominal stimulus dimension that can display multiple values, SUSTAIN devotes multiple input units. To represent a nominal dimension containing k distinct values, k input units are utilized. All the units forming a dimension are set to zero, except for the one unit that denotes the nominal value of the dimension (this unit is set to one). For example, the stimulus dimension of marital status has three values (“single”, “married”, “divorced”). The pattern [0 1 0] represents the dimension value of “married”. A complete stimulus is represented by the vector I^{posik} where i indexes the stimulus dimension and k indexes the nominal values for dimension i . For example, if marital status was the third stimulus dimension and the second value was present (i.e., married), then I^{pos32} would equal one, whereas I^{pos31} and I^{pos33} would equal zero. The “pos” in I^{pos} denotes that the current stimulus is located at a particular position in a multi-dimensional representational space.

2.3.2 Receptive Fields

Each cluster has a receptive field for each stimulus dimension. A cluster’s receptive field for a given dimension is centered at the cluster’s position along that dimension. The position of a cluster within a dimension indicates the cluster’s expectations for its members.

The tuning of a receptive field (as opposed to the position of a receptive field) determines how much attention is being devoted to the stimulus dimension. All the receptive fields for a stimulus dimension have the same tuning (i.e., attention is dimension-wide as opposed to cluster-specific). A receptive field’s tuning changes as a result of learning. This change in receptive field tuning implements SUSTAIN’s selective attention mechanism. Dimensions that are highly attended to develop peaked tunings, whereas dimensions that are not well attended to develop broad tunings. Dimensions that provide consistent information at the cluster level receive greater attention.

Mathematically, receptive fields have an exponential shape with a receptive field’s response decreasing exponentially as distance from its center increases. The activation function for a dimension is:

$$\alpha(\mu) = \lambda e^{-\lambda\mu} \tag{1}$$

where λ is the tuning of the receptive field, μ is the distance of the stimulus from the center of the field, and $\alpha(\mu)$ denotes the response of the receptive field to a stimulus falling μ units from the center of the field. The choice of exponentially shaped receptive fields is motivated by Shepard’s (1987) work on stimulus generalization.

Although receptive fields with different λ have different shapes (ranging from a broad to a peaked exponential), for any λ , the area “underneath” a receptive field is constant:

$$\int_0^\infty \alpha(\mu) d\mu = \int_0^\infty \lambda e^{-\lambda\mu} d\mu = 1. \tag{2}$$

For a given μ , the λ that maximizes $\alpha(\mu)$ can be computed from the derivative:

$$\frac{\partial \alpha}{\partial \lambda} = e^{-\lambda\mu} (1 - \lambda\mu). \tag{3}$$

These properties of exponentials prove useful in formulating SUSTAIN.

2.3.3 Cluster Activation

With nominal stimulus dimensions, the distance μ_{ij} (from 0 to 1) between the *ith* dimension of the stimulus and cluster *j*’s position along the *ith* dimension is:

$$\mu_{ij} = \frac{1}{2} \sum_{k=1}^{v_i} |I^{pos_{ik}} - H_j^{pos_{ik}}| \tag{4}$$

where v_i is the number of different nominal values on the *ith* dimension, I is the input representation (as described in a previous section), and $H_j^{pos_{ik}}$ is cluster *j*’s position on the *ith* dimension for value *k* (the sum of all *k* for a dimension is 1). The position of a cluster

in a nominal dimension is actually a probability distribution that can be interpreted as the probability of displaying a value given that an item is a member of the cluster. For example, a cluster in which 20% of the members are single, 45% are married, and 35% are divorced will converge to the location [.20 .45 .35] within the marital status dimension. The distance μ_{ij} will always be between 0 and 1 (inclusive).

The activation of a cluster is given by:

$$H_j^{act} = \frac{\sum_{i=1}^m (\lambda_i)^r e^{-\lambda_i \mu_{ij}}}{\sum_{i=1}^m (\lambda_i)^r} \quad (5)$$

where H_j^{act} is the activation of the j th cluster, m is the number of stimulus dimensions, λ_i is the tuning of the receptive field for the i th input dimension, and r is an attentional parameter (always nonnegative). When r is large, input units with tighter tunings (units that seem relevant) dominate the activation function. Dimensions that are highly attended to have larger λ s and will have greater importance in determining the clusters' activation values. Increasing r simply accentuates this effect. If r is set to zero, every dimension receives equal attention. Equation 5 sums the responses of the receptive fields for each input dimension and normalizes the sum (again, highly attended dimensions weigh heavily). Cluster activation is bound between 0 (exclusive) and 1 (inclusive). Unknown stimulus dimensions (e.g., the category label in a classification trial) are not included in the above calculation.

2.3.4 Competition

Clusters compete to respond to input patterns and in turn inhibit one another. When many clusters are strongly activated, the output of the winning cluster H_j^{out} is less:

For the winning H_j with the greatest H^{act} ,

$$H_j^{out} = \frac{(H_j^{act})^\beta}{\sum_{i=1}^n (H_i^{act})^\beta} H_j^{act} \quad (6)$$

For all other H_j ,

$$H_j^{out} = 0.$$

where n is the number of clusters and β is the lateral inhibition parameter (always non-negative) that regulates cluster competition. When β is small, competing clusters strongly inhibit the winner. When β is large the winner is weakly inhibited. Clusters other than the winner have their output set to zero. Equation 6 is a straightforward method for implementing lateral inhibition. It is a high level description of an iterative process where units send signals to each other across inhibitory connections. Psychologically, Equation 6 signifies that competing alternatives will reduce confidence in a choice (reflected in a lower output value).

2.3.5 Response

Activation is spread from the clusters to the output units of the queried (the unknown) stimulus dimension z :

$$C_{zk}^{out} = \sum_{j=1}^n w_{j,zk} H_j^{out} \quad (7)$$

where C_{zk}^{out} is the output of the output unit representing the k th nominal value of the queried (unknown) z th dimension, n is the number of clusters, and $w_{j,zk}$ is the weight from cluster j to category unit C_{zk} . A winning cluster (especially one that did not have many competitors and is similar to the current input pattern) that has a large positive connection to a output unit will strongly activate the output unit. The summation in the above calculation is not really necessary given that only the winning cluster has a nonzero output, but is included to make the similarities between SUSTAIN and other models more apparent.

The probability of making response k (the k th nominal value) for the queried dimension z is

$$Pr(k) = \frac{e^{(d \cdot C_{zk}^{out})}}{\sum_{j=1}^{v_z} e^{(d \cdot C_{zj}^{out})}} \quad (8)$$

where d is a response parameter (always nonnegative) and v_z is the number of nominal units (and hence output units) forming the queried dimension z . When d is high, accuracy is stressed and the output unit with the largest output is almost always chosen. The Luce choice rule is conceptually related to this decision rule (Luce, 1959).

2.3.6 Learning

After responding, feedback is provided to SUSTAIN. The target value for the k th category unit of the queried dimension z is:

$$t_{zk} = \left\{ \begin{array}{l} \max(C_{zk}^{out}, 1), \text{ if } I^{pos_{zk}} \text{ equals } 1. \\ \min(C_{zk}^{out}, 0), \text{ if } I^{pos_{zk}} \text{ equals } 0. \end{array} \right\} \quad (9)$$

Kruschke (1992) refers to this kind of teaching signal as a “humble teacher” and explains when its use is appropriate. Basically, the model is not penalized for predicting the correct response more strongly than is necessary.

A new cluster is recruited if the winning cluster predicts an incorrect response. In the case of a supervised learning situation, a cluster is recruited according to the following procedure:

$$\begin{array}{l} \text{For the queried dimension } z, \\ \text{if } t_{zk} \text{ does not equal } 1 \text{ for the } C_{zk} \\ \text{with the largest output } C_{zk}^{out} \text{ of all } C_{z*}, \\ \text{then recruit a new cluster.} \end{array} \quad (10)$$

In other words, the output unit representing the correct nominal value must be the most activated of all the output units forming the queried stimulus dimension.

In the case of an unsupervised learning situation, SUSTAIN is self-supervising and recruits a cluster when the most activated cluster H_j 's activation is below the threshold τ :

$$\text{if } (H_j^{act} < \tau), \text{ then recruit a new cluster.} \quad (11)$$

Unsupervised recruitment in SUSTAIN bears a strong resemblance to recruitment in Adaptive Resonance Theory (Carpenter and Grossberg, 1987), Clapper and Bower's (1991) qualitative model, and Hartigan's (1975) leader algorithm.

When a new cluster is recruited it is centered on the misclassified input pattern and the clusters' activations and outputs are recalculated. The new cluster then becomes the winner because it will be the most highly activated cluster (it is centered upon the current input pattern — all μ_{ij} will be zero). Again, SUSTAIN begins with a cluster centered on the first stimulus item.

The position of the winner is adjusted:

$$\begin{aligned} &\text{For the winning } H_j, \\ &\Delta H_j^{pos_{ik}} = \eta(I^{pos_{ik}} - H_j^{pos_{ik}}) \end{aligned} \quad (12)$$

where η is the learning rate. The centers of the winner's receptive fields move towards the input pattern according to the Kohonen (1984) learning rule. This learning rule centers the cluster amidst its members.

Using our result from Equation 3, receptive field tunings are updated according to:

$$\Delta \lambda_i = \eta e^{-\lambda_i \mu_{ij}} (1 - \lambda_i \mu_{ij}) \quad (13)$$

where j is the index of the winning cluster.

Only the winning cluster updates the value of λ_i . Equation 13 adjusts the peakedness of the receptive field for each input so that each input dimension can maximize its influence on the clusters. Initially, λ_i is set to be broadly tuned with a value of 1. The value of 1 is chosen because the maximal distance μ_{ij} is 1 and the optimal setting of λ_i for this case is 1 (i.e., Equation 13 equals zero). Under this scheme, λ_i cannot become less than 1, but can become more narrowly tuned.

When a cluster is recruited, weights from the unit to the output units are set to zero. The one layer delta learning rule (Widrow and Hoff, 1960) is used to adjust these weights:

$$\Delta w_{j,zk} = \eta(t_{zk} - C_{zk}^{out})H_j^{out}. \quad (14)$$

where z is the queried dimension. Note that only the winning cluster will have its weights adjusted since it is the only cluster with a nonzero output.

2.4 uSUSTAIN: A Unified Approach to Supervised and Unsupervised Learning

SUSTAIN can model both supervised and unsupervised learning, but it relies on different recruitment mechanisms. In both cases, a cluster is recruited in response to a surprising event

(i.e., the existing cluster structure does not properly characterize the current stimulus), but how a surprising event is defined differs. In the supervised case, the surprising event is a prediction error, whereas in the the case of unsupervised learning the surprising event is an unfamiliar stimulus

Although the two separate recruitment procedures has been successful, a single recruitment procedure is preferable. Beyond parsimony, a unified account could prove useful in clarifying the relationship between unsupervised and supervised learning. A simple way to integrate the two recruitment strategies is to generalize the unsupervised procedure so that it is applicable to supervised learning situations. Under this scheme, a new cluster is recruited when the current stimulus is not sufficiently similar to any cluster in its category:

$$\begin{aligned} &\text{For the queried dimension } z, \\ &\text{If } \text{Max}(\{H_j^{act} \mid \mu_{zj} = 0\}) < \tau, \text{ then recruit a new cluster.} \end{aligned} \tag{15}$$

where H_j^{act} is the activation of cluster j , μ_{zj} is the distance (as defined in Equation 4) along the z th dimension of the current stimulus and cluster j 's position along the z th dimension, and τ is a constant between 0 and 1 (a parameter). The requirement that μ_{zj} be zero specifies that only clusters associated with the category of the current stimulus are considered. In unsupervised learning, all items belong to the same global category which represents items the network has seen before. Thus, $\text{Max}(\{H_j^{act} \mid \mu_{zj} = 0\})$ refers to the most activated cluster overall. In supervised learning, the most activated cluster predicting the correct category may not be the most activated cluster overall.

Besides providing a unified framework, this recruitment strategy has a number of other virtues over SUSTAIN's original recruitment rule (Equation 10) for supervised learning. For example, the unified procedure will recruit a new cluster when an unusual item is encountered that does not result in a prediction error whereas the previous error-driven recruitment scheme would not recruit a new cluster to encode the unusual item. Assigning a very unusual item to an existing cluster (a cluster the item is not very similar to) could result in catastrophic interference (see Ratcliff, 1990) as the cluster must undergo radical change to accommodate its newest member.

3 Evaluating uSUSTAIN

In order to evaluate this unified formulation of the model we applied uSUSTAIN to the studies previously accounted for using separate cluster recruitment mechanisms for supervised and unsupervised learning (Gureckis and Love, 2002b). It is important to recognize that the recruitment procedure that uSUSTAIN uses is, in fact, a generalization of unsupervised recruitment procedure used by the original SUSTAIN model. Thus, uSUSTAIN and SUSTAIN provide equivalent accounts of unsupervised learning (Gureckis and Love, 2002a). uSUSTAIN and SUSTAIN have fit a number of unsupervised learning studies and have generated novel predictions that have been subsequently tested and confirmed with human subjects (Gureckis and Love, 2002c).

A true test of generality of the uSUSTAIN approach lies in its ability to fit supervised learning data. Gureckis and Love (2002b) applied uSUSTAIN to a number of supervised learning studies and found that uSUSTAIN approximated SUSTAIN’s successes. Despite its simplicity, the unified recruitment procedure in uSUSTAIN has proven remarkably successful in this domain.

Although uSUSTAIN has demonstrated the ability to account for human learning performance across a wide range of category learning paradigms, it has never been applied to a study specifically designed to compare unsupervised and supervised learning. Given the past successes of the model, it would be informative to apply the model to a direct comparison between supervised and unsupervised learning. The following section examines uSUSTAIN’s account of Love’s (2002) study that compares incidental unsupervised learning, intentional unsupervised learning, and supervised classification learning in a controlled manner.

4 Comparing Supervised and Unsupervised Learning

The Love (2002) study is unique in that it specifically allows for a direct comparison of supervised and unsupervised learning. In supervised learning, the common dependent measure used to assess learning difficulty is training accuracy (Nosofsky et al., 1994; Shepard, Hovland, and Jenkins, 1961). However, there is no measure of training accuracy in unsupervised learning (there is no right or wrong response on each study trial). In order to directly compare learning performance across these two types of learning, a comparable dependent measure was developed.

To accomplish this, stimuli were created by embedding the category label (which is typically a verbal label such as category “A” or “B”) into each stimulus as a fourth binary-valued perceptual dimension (see Table 1). On supervised classification study phase trials, subjects were shown the value of the first three perceptual dimensions and were queried on the fourth. After responding, the correct value of the fourth dimensions was shown. In the Love (2002) study, the fourth dimension (i.e., the “category” dimension) was the border color (either yellow or white) of a geometric figure. Subjects indicated whether they believed the border color (not shown on the display) was yellow or white based on the three other perceptual dimensions (which were shown on the display). After responding, the complete figure was displayed.

On unsupervised study phase trials, all four perceptual dimensions were shown on study phase trials (the fourth dimension was not queried). In the intentional unsupervised learning condition, subjects were aware they were in a learning task and were instructed to actively search for patterns that characterized the training items. In contrast, subjects in the incidental unsupervised learning condition were not aware that they were in a learning task and were instructed to simply rate how pleasant they found each stimulus item.

In each of the three study conditions (supervised classification learning, intentional unsupervised learning, incidental unsupervised learning), subjects were trained on either the Type I, II, IV, or VI category structures (see Table 1) defined by Shepard, Hovland and Jenkins (1961). The Type I problem only requires attention along one input dimension,

Table 1: The logical structure of the Type I, II, IV, and VI classification problems tested in Shepard et al. (1961) are shown.

Stimulus	I	II	IV	VI
1 1 1	1	1	1	1
1 1 2	1	1	1	2
1 2 1	1	2	1	2
1 2 2	1	2	2	1
2 1 1	2	2	1	2
2 1 2	2	2	2	1
2 2 1	2	1	2	1
2 2 2	2	1	2	2

whereas the Type II problem requires attending to two dimensions (Type II is XOR on the first two dimensions with an irrelevant third dimension). The categories in the Type II problem have a highly nonlinear structure. Type IV requires attention along all three perceptual dimensions with each dimension serving as an imperfect predictor. Type IV is notable because it displays a linear category structure. Type VI also requires attention to all three perceptual dimensions and has no regularities across any pair of dimensions. In all conditions, subjects completed ten study blocks (a block consists of the presentation of each stimulus item in a random order).

Category learning performance was measured in a test phase which followed the study phase. Subjects viewed a pair of stimuli that varied only on the fourth dimensions (i.e., the category dimension). Subjects were instructed to choose the item that appeared during the study phase (a familiarity or recognition judgment). As in traditional supervised classification learning studies, subjects could base this judgment on their knowledge of the relationship between the category dimension and other dimensions (e.g., rules, correlations, etc.) as well as on memorized exemplars. Love (2002) verified that this testing procedure yields performance scores that correlate highly with study phase accuracy in the supervised condition. Thus, test phase accuracy can be used to compare the ability of subjects to learn in each of the three study conditions.

The results are shown in Table 2. The acquisition patterns for the three learning conditions differ significantly. Subjects in the unsupervised conditions did not show a preference for the Type II category structure relative to the Type IV structure. This effect was most pronounced in the incidental unsupervised learning condition. One explanation for this difference between the incidental and intentional unsupervised learning conditions is that intentional unsupervised learning task encouraged subjects to form explicit rules which is an efficient strategy for category structures that are describable by a compact rule (Type II has an XOR structure, see Table 1)

Table 2: The study phase and test phase results from Love (2002). uSUSTAIN’s fit is shown in parentheses.

Problem	Study Accuracy	Test Accuracy
Supervised Classification Learning		
Type I	.86 (.74)	.89 (.90)
Type II	.67 (.63)	.73 (.75)
Type IV	.65 (.60)	.70 (.65)
Type VI	.59 (.56)	.61 (.58)
Intentional Unsupervised Learning		
Type I	NA	.84 (.86)
Type II	NA	.64 (.57)
Type IV	NA	.67 (.66)
Type VI	NA	.54 (.50)
Incidental Unsupervised Learning		
Type I	NA	.85 (.81)
Type II	NA	.56 (.51)
Type IV	NA	.67 (.63)
Type VI	NA	.56 (.50)

4.1 Modeling Results

uSUSTAIN was trained in a manner analogous to how subjects were trained using 10 randomly ordered study blocks. In the supervised condition, the model was asked to predict the value of the fourth feature dimension (the category label), while in both unsupervised conditions, the value of all four stimulus dimensions was available to the model.

All study phase stimuli were encoded as belonging to a global category. In order for uSUSTAIN to mimic the forced choice nature of the test phase, the activation of this global category unit was calculated for each of the two items presented as a pair. The degree to which this unit was activated indicates the level of familiarity the model has for the item. The ultimate response of the network was towards the item in the forced choice that had the strongest response activity using a probabilistic decision procedure analogous to Equation 8.

uSUSTAIN’s fit of the data (averaged of 5,000 runs) is shown in Table 2 in parentheses and the best-fit modeling parameters used are shown in Table 3. uSUSTAIN correctly demonstrates different patterns of acquisition for the three study conditions. In the supervised classification condition, uSUSTAIN captures the correct ordering of problem difficulty (Type I is easiest, followed by Type II, then Type IV, and finally Type VI). In the both unsupervised conditions, the model correctly shows a decrease in test phase accuracy for the Type II category structures relative to the supervised classification condition. Note that uSUSTAIN predicts lower accuracy for category structures that are nonlinear (Type II and Type VI) in both unsupervised conditions.

The clusters that uSUSTAIN creates were analyzed. The modal number of clusters recruited in the supervised learning condition exactly matches with previously reported results

Table 3: uSUSTAIN’s best fitting parameters for the Love, 2002 studies.

function/adjusts	symbol	supervised	intentional/incidental	unsupervised
learning rate	η	.0172		.0186
cluster competition	β	2.90		0.608
decision consistency	d	14.474		14.442
attentional focus	r	.475		2.209
threshold	τ	0.568		.553/.487

using the original version of SUSTAIN to account for the Shepard, Hovland, Jenkins (1961) problems (Love, Medin, and Gureckis, 2002). Two clusters were created in the Type I condition, which divided the stimulus items on the basis of the first stimulus dimension. Four clusters were recruited in the Type II condition, which captures the nature of the XOR on the first two dimensions. Six clusters were created in the Type IV problem and eight clusters were recruited in the Type VI condition where each item had to be memorized. In the both of the unsupervised learning conditions, the modal number of clusters was two for all problems except Type II under incidental conditions, which required four clusters.

Using essentially the same set of parameters uSUSTAIN was able to qualitatively approximate performance in both the intentional and incidental unsupervised learning conditions. Only the threshold parameter, τ , was different for the two studies ($\tau = .533$ for intentional unsupervised learning and $\tau = .487$ for incidental unsupervised learning). The increased threshold parameter in the intentional condition allowed uSUSTAIN to create four clusters in the Type II problem which improves its test phase performance for this problem.

5 Discussion and Conclusions

uSUSTAIN holds that supervised and unsupervised human learning engage common mechanisms. In both cases, a new cluster is recruited when a stimulus is not sufficiently similar to an any existing cluster belonging to the appropriate category. In order to evaluate uSUSTAIN’s account of human learning, the model was applied to Love’s (2002) data set that compared human performance in supervised classification learning, incidental unsupervised learning, and intentional unsupervised learning. uSUSTAIN successfully fit these data.

An analysis of uSUSTAIN’s clustering solutions suggests that humans are more likely to aggregate stimulus items in memory in unsupervised learning whereas they are more likely to segregate stimulus items and store separate memory traces in supervised learning. Within unsupervised learning, the tendency to collapse items into common clusters is stronger in incidental learning.

uSUSTAIN’s fit of the Love (2002) data suggests that unsupervised learning, particularly incidental unsupervised learning, is best matched with linear category structures because the optimal clustering solution for a linear category structure involves one cluster per category. On the other hand, nonlinear category structures are not well matched to an unsupervised induction task because nonlinear category structures can only be captured with multiple

clusters per category. While the linear/nonlinear distinction has not proved critical in supervised classification learning (Medin & Schwanenflugel, 1981), Love (2002) suggested that the distinction may be meaningful in unsupervised learning. uSUSTAIN’s account of the data supports this conjecture.

One counterintuitive prediction that uSUSTAIN makes is that incidental unsupervised learning may be the preferred induction task for some tasks. In other words, sometimes humans may be better off not trying to master the learning problem. One such situation is when numerous stimulus dimensions are weakly correlated with one another. Under such circumstances, uSUSTAIN predicts that supervised classification learning and intentional unsupervised learning will lead to clustering solutions that over-differentiate items and therefore do not fully capture the intercorrelated structure of the categories. In contrast, incidental unsupervised learning tends to aggregate items in common clusters and is more likely to capture the underlying category structure. uSUSTAIN’s lower setting of the τ parameter (which increases uSUSTAIN’s tendency to cluster items together) for incidental unsupervised learning drives this prediction.

Despite the apparent differences between supervised classification learning, intentional unsupervised learning, and incidental unsupervised learning, all three induction tasks are modeled through a common mechanism in uSUSTAIN. Beyond the current project, an important goal of our efforts is to model human learning across a range of situations and induction tasks. Doing so highlights theoretical connections across data sets and should lead to a more general understanding of human learning.

References

- Anderson, J., 1991, The adaptive nature of human categorization, *Psychological Review*, 98:409–429.
- Ashby, F., Queller, S., and Berretty, P. M., 1999, On the dominance of unidimensional rules in unsupervised categorization, *Perception & Psychophysics*, 61:1178–1199.
- Berry, D. C., and Dienes, Z., 1993, *Implicit Learning: Theoretical and empirical issues*, Erlbaum, Hillsdale, NJ.
- Billman, D., and Knutson, J., 1996, Unsupervised concept learning and value systematicity: A complex whole aids learning the parts, *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 22(2):458–475.
- Carpenter, G. A., and Grossberg, S., 1987, A massively parallel architecture for a self-organizing neural pattern recognition machine, *Computer Vision, Graphics, and Image Proc.*, 37:54–115.
- Clapper, J. P., and Bower, G. H., 1991, Learning and applying category knowledge in unsupervised domains, *The Psychology of Learning and Motivation*, 27:65–108.

- Clapper, J. P., and Bower, G. H., 1994, Category invention in unsupervised learning, *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 20:443–460.
- Cleermans, A., 1993, *Mechanisms of implicit learning: Connectionist models of sequence processing*, MIT Press, Cambridge, MA.
- Gureckis, T., and Love, B. C., 2002a, Modeling unsupervised learning with sustain, In *Proceedings of the 15th Annual FLAIRS Conference*, pp. 123–456.
- Gureckis, T., and Love, B. C., 2002b, Towards a unified account of supervised and unsupervised learning, *Journal of Experimental and Theoretical Artificial Intelligence*, in press.
- Gureckis, T., and Love, B. C., 2002c, Who says models can only do what you tell them? unsupervised category learning data, fits, and predictions, In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, pp. 399–404, Hillsdale, NJ, Lawrence Erlbaum Associates.
- Hartigan, J. A., 1975, *Clustering Algorithms*, Wiley, New York.
- Hayes, N., and Broadbent, D. E., 1988, Two modes of learning for interactive tasks, *Cognition*, 28:249–276.
- Hock, H. S., Malcus, L., and Hasher, L., 1986, Frequency discrimination: Assessing global and elemental letter units in memory, *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 12:232–240.
- Kohonen, T., 1984, *Self-Organization and Associative Memory*, Springer, Berlin, Heidelberg, 3rd ed. 1989.
- Kruschke, J. K., 1992, ALCOVE: An exemplar-based connectionist model of category learning, *Psychological Review*, 99:22–44.
- Lewicki, P., 1986, *Nonconscious social information processing*, Academic Press, New York.
- Love, B. C., 2002, Comparing supervised and unsupervised category learning, *Psychological Bulletin & Review*, in press.
- Love, B. C., Markman, A. B., and Yamauchi, T., 2000, Modeling classification and inference learning, *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pp. 136–141.
- Love, B. C., and Medin, D. L., 1998a, Modeling item and category learning, In *Proceedings of the 20th Annual Conference of the Cognitive Science Society*, pp. 639–644, Mahwah, NJ, Lawrence Erlbaum Associates.

- Love, B. C., and Medin, D. L., 1998b, SUSTAIN: A model of human category learning, In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pp. 671–676, Cambridge, MA, MIT Press.
- Love, B. C., Medin, D. L., and Gureckis, T., 2002, SUSTAIN: A network model of human category learning, Under Review, *Psychological Review*.
- Luce, R. D., 1959, *Individual choice behavior: A theoretical analysis*, Greenwood Press, Westport, Conn.
- Medin, D. L., and Schwanenflugel, P. J., 1981, Linear separability in classification learning, *Journal of Experimental Psychology: Human Learning & Memory*, 7:355–368.
- Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., McKinley, S. C., and Glauthier, P., 1994, Comparing models of rule based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961), *Memory & Cognition*, 22:352–369.
- Nosofsky, R. M., Palmeri, T. J., and McKinley, S. C., 1994, Rule-plus-exception model of classification learning, *Psychological Review*, 101(1):53–79.
- Ratcliff, R., 1990, Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions, *Psychological Review*, 97:285–308.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J., 1986, Learning representations by back-propagating errors, *Nature*, 323:533–536.
- Shepard, R. N., 1987, Toward a universal law of generalization for psychological science, *Science*, 237:1317–1323.
- Shepard, R. N., Hovland, C. L., and Jenkins, H. M., 1961, Learning and memorization of classifications, *Psychological Monographs*, 75(13, Whole No. 517).
- Sloman, S. A., 1997, Explanatory coherence and the induction of properties, *Thinking & Reasoning*, 3:81–110.
- Widrow, B., and Hoff, M. E., 1960, Adaptive switching circuits, In *IRE WESCON Convention Record*, New York, pp. 96–104.