Running Head: ALIGNMENT AND MULTIPLE CATEGORY LEARNING

Alignment Effects on Learning Multiple, Use-Relevant Classification Systems

Cynthia M. Sifonis Brian H. Ross

Oakland University University of Illinois-Urbana Champaign

Please address all correspondence to:

Cynthia M. Sifonis
Department of Psychology
222 Pryale Hall
Oakland University
Rochester, MI, 48309- 4489

phone: (248) 370-4680 (office)
(248) 370-4612 (fax)
email: sifonis@oakland.edu

Abstract

People often learn multiple, classification systems that are relevant to some goal or use. We compared conditions in which subclassification within a category hierarchy was predicted by values on either the same (alignable) or different (nonalignable) dimensions between category hierarchies. The results indicated that learning in alignable conditions occurred in fewer blocks and with fewer errors than learning in nonalignable conditions. This facilitation was not the result of differences between conditions in the representations learned by the participants, the number of dimensions needed for subclassification (Exp. 1), or the objective complexity of the learning task (Exp. 2). The facilitated learning in the alignable conditions appears to reflect a commitment on the part of the learner to alignment - the belief that the structure relevant to the use of one category system will also be relevant to the use of a comparable system.

Alignment Effects on Learning Multiple, Use-Relevant Classification Systems

People often learn about and interact with instances in an unfamiliar domain. In the process, knowledge is built up about the instances, their typical features, and the application of the domain knowledge in a variety of tasks. The result is often an organization of domain knowledge into a category hierarchy that reflects the uses towards which the knowledge is applied (Boster & Johnson, 1989; Medin, Lynch, Coley, & Atran, 1997; Ross & Murphy, 1999). In this paper we explore the manner in which people learn such hierarchies. Of specific interest are factors affecting how people learn multiple hierarchies simultaneously.

When interacting with category instances, people frequently organize domain knowledge into several related hierarchies. For example, a person learning to train working dogs for a variety of tasks might form distinct subcategories of dogs such as “tracking dogs” and “service dogs”. In working with animals from each category, this person might further subdivide the categories into more specific subcategories. For example, they might subdivide the "service dogs" category into “vision dogs” and “hearing dogs.” At the same time they might also be subdividing the ”tracking dogs” category into “people-tracking dogs” and “narcotics dogs.”

Research supports the proposal that task-related subdivisions of domain knowledge occur in a wide range of domains. For instance, Ross and Murphy (1999) identified salient organizations of the food domain knowledge of college students resulting from the need to plan and eat several meals a day (e.g., breakfast foods). Tree experts and fish experts have been shown to possess subdivisions of domain knowledge reflecting goal-directed interactions with members of those domains (e.g., “weed trees”, “game fish”) (Boster & Johnson, 1989; Medin, et al., 1997). In short, the use of category knowledge frequently results in information relevant for that use becoming organized and grouped into goal-relevant subcategories.

Unfortunately, we know little about how people form multiple, goal-relevant category hierarchies. The research mentioned above examines the organization of domain knowledge in well-established domains. Consequently, it does not tell us how people initially learn to classify instances at multiple levels within a domain based on their interactions with category members. The research that actually examines the formation of class-inclusion hierarchies typically does so for a single superordinate. The focus of these studies is often on the factors affecting the ease with which category subdivisions are determined (e.g., types of dogs, types of butterfly) (Waxman, Lynch, Casey & Baer, 1997). One consequence of examining subclassification within a single superordinate is that the focus is on within-category factors that affect learning. As a result, we know little about possible between-category factors that could affect the ease with which multiple classification systems are learned.

It is important to investigate how people learn multiple classification hierarchies because people rarely learn category hierarchies in isolation. Sometimes we learn categories that bring to mind related categories and learning situations. At other times we need to learn to distinguish between multiple categories simultaneously and make further goal-driven distinctions within those categories. In the example described previously, the dog trainer is learning multiple use-relevant hierarchies. Distinctions at each hierarchical level depend on decisions concerning the types of tasks to which each dog is best suited. If a dog is too aggressive to be a service dog then perhaps it will make a better tracking dog. A dog docile enough to be a service dog might be too small to be an effective vision dog but would make a very good hearing dog.

As the example illustrates, an important part of forming a category hierarchy within a domain is identifying the dimensions important for classification at each level of the hierarchy. In the preceding example, values on the temperament dimension (i.e., aggression) allow the trainer to distinguish between service dogs and tracking dogs. For service dogs, values on the dimension of size determined the most effective use of the dog and allowed the trainer to make distinctions between types of service dogs (vision dogs versus hearing dogs).

Identifying the dimensions on which divisions within a hierarchy can be made is facilitated when those divisions can be made on the basis of contrastive information (Billman, 1996; Waxman et al., 1997). Contrastive information refers to an aspect of category structure in which the same dimensions matter across the category. When values on the same dimension allow a person to distinguish between categories at the same level of a hierarchy, learning is facilitated. This is true whether learning to distinguish between basic level categories (Kaplan, 2000; Lassaline & Murphy, 1998) or when learning to make subordinate level distinctions within a category (Billman, 1996; Waxman et al., 1997). It has been assumed that this facilitation reflects assumptions made on the part of the learner about the structure of the categories being learned. However, as discussed shortly, such facilitation may be a by-product of the designs used to study the phenomenon. Given the evidence that similarities in category structure facilitate classification learning at several hierarchical levels it is extremely important to investigate the source of this facilitation.

The present study extends the work reported in Sifonis and Ross (1999) demonstrating that between-category similarities in the dimensions important for the use of category knowledge facilitate learning multiple classification systems. The source of this facilitation will be explored in two experiments. The first experiment examines whether facilitation is due to the number of dimensions needing attention in order to make subordinate-level category distinctions. The second experiment examines whether the facilitation is a product of the psychological complexity of the learning task[1].

In the current work, we chose the domain of bank loan applications for two reasons. It is a complex domain that is similar to some real-world category learning and it also allows the learning of multiple, goal-driven category distinctions. When presented with a member of the domain (a loan application form), participants decide whether or not the applicant receives the loan. This initial decision establishes the superordinate categories in which participants will learn to make further goal-relevant category distinctions. After the loan application is categorized into one of the two superordinate categories, participants decide what type of financial advice to give the applicant and/or what consultant to assign to the applicant. The participants learn two hierarchies, each containing subordinate categories derived from the uses towards which category knowledge is applied (deciding financial advice and/or consultant).

Structural Alignment

Some of the work examining the effects of contrastive information on classification learning has used the concept of structural alignment to predict and explain classification performance (Goldstone, 1994; Kaplan, 2000; Lassaline & Murphy 1998). Structural alignment theory assumes that category representations encode information such as the dimensions associated with the category, the features typically instantiated by those dimensions, and the relationships between dimensions. Between-category comparisons involve aligning these representations to find the most structurally consistent match between them (Gentner, 1983). The alignment process highlights structural consistencies and inconsistencies between representations. Structural consistencies known as commonalities involve between-category matches on both dimensions and features. For example, if both loan-suitable and loan-unsuitable companies advertise in magazines this would be a commonality in their representations. There is a between-category match on dimension (method of advertising) and feature (magazine). Differences in representational structure can be either alignable or nonalignable. Alignable differences involve a match on dimension but a mismatch on features. Loan-suitable and loan-unsuitable companies would have alignable differences if the loan-suitable company advertised by television whereas the loan-unsuitable company advertised by newspaper. Nonalignable differences, on the other hand, are independent of the common structure. Differences of this type involve dimensions that have no corresponding dimension for the other entity. For instance, perhaps included on the loan application of a grocery company are dimensions representing shipping costs and distribution. These dimensions would not be applicable to companies specializing in web development and would not be included on their loan applications.

Alignable Differences and Learning Use-relevant Classification Systems

The distinction between alignable and nonalignable differences has proven useful for predicting classification learning at both basic (Kaplan, 2000, Exps. 2 & 3; Lassaline & Murphy, 1998) and subordinate levels in a hierarchy (Waxman,et al., 1997). It also predicts the ease with which multiple classification systems are learned (Billman, 1996; Sifonis & Ross, 1999). Of particular interest to the present study is the demonstration that between-category consistencies in the use of category knowledge affect the formation of use-relevant subordinate categories[2] (Sifonis & Ross, 1999).

As in the present study, Sifonis and Ross (1999) had participants make several classification distinctions based on information in bank loan application forms. Participants initially classified the loan applications into the superordinate categories of applications that received loans (loan category) and those that did not receive loans (no-loan category). Participants then had to make financial advice decisions based on superordinate category membership (loan/no-loan) and information found in the applications. The between-category alignability of the dimensions necessary for predicting financial advice was varied between conditions. In the alignable-use condition, the same dimension (i.e., “Type of Sales”) predicted financial advice for both superordinate categories of loan applications. In the nonalignable-use condition, the dimensions predicting financial advice differed between superordinates (i.e., “Type of Sales” for applications receiving a loan and “Method of Advertising” for applications that did not receive a loan). It was found that when the relationship between a dimension and its use were consistent between categories, learning the classification hierarchies was facilitated.

The Sifonis and Ross (1999) study demonstrated that between-category consistencies in the dimensions predictive of subclassification within a hierarchy affect the ease with which multiple classification hierarchies are learned. It also avoided an important confound of previous studies examining alignment effects on classification. Specifically, in the Sifonis and Ross (1999) study there were no differences between conditions in the category representations learned by participants.

Typically, experiments examining alignment effects on classification learning manipulate alignability in terms of whether contrastive categories possess alignable or nonalignable dimensions predictive of category membership (Kaplan, 2000 exps 2 & 3; Wisniewski & Markman, 1997 exps 1 & 2). By necessity, the contrast category representations and the exemplars experienced by participants during learning differ between conditions. By examining how people learn use-relevant subclassifications within multiple hierarchies, it becomes possible to manipulate the alignability of the dimensions important for category use between conditions while holding the alignability of the category representations constant.

In the Sifonis and Ross (1999) study, the category representations learned by participants were identical between conditions. Both categories of loan applications possessed the same dimensions in both the alignable-use and nonalignable-use conditions. Participants in both conditions experienced exactly the same category exemplars. As a result, the facilitated learning seen in the alignable-use condition of that study cannot be attributed to differences between conditions in the category representations learned by the participants. However, it is still possible the facilitation arises from a source other than the alignability of the dimensions predictive of subclassification within each hierarchy.

To date, all the studies examining multiple classification system learning have ensured that the number of features needed to make the classification and subclassification decisions are equated across conditions. However, the same cannot be said for the number of dimensions associated with those features (Billman, 1996; Sifonis & Ross, 1999). Consequently, it is possible the alignability effects exhibited in previous studies result from differences between conditions in the number of dimensions needed to make classification decisions. By definition, a single classification decision made on the basis of alignable differences between categories requires attention to a single dimension (e.g., Type of Sales: retail / Type of Sales: wholesale). In contrast, decisions made on the basis of nonalignable differences between categories require attention to several dimensions (Type of Sales: retail / Method of Advertising: magazine). In order to determine the independent contribution of alignability on learning multiple classification systems it is necessary to separate alignability from the number of dimensions necessary for making category distinctions.

Experiment 1

Experiment 1eliminates the confound between alignability and the number of attended dimensions by requiring that participants learn to make two use-relevant distinctions within each superordinate category. As in Sifonis and Ross (1999), participants reviewed bank loan application forms and classified the applications at both the superordinate and subordinate levels. However, in the present experiment, participants in both the alignable-use and nonalignable-use conditions were required to pay attention to two dimensions in order to classify at the subordinate level. The alignability of these dimensions was manipulated by varying the uses predicted by the dimensions in each of the two superordinate categories. In the alignable-use condition, the same dimension was used to predict the same use in both superordinate categories of applications (i.e., Dimension 1 predicts Use 1 and Dimension 2 predicts Use 2 for both the loan and no-loan categories). In the nonalignable-use condition, the uses predicted by the dimensions changed for each superordinate (i.e., Use 1 is predicted by Dimension 1 and Use 2 by Dimension 2 for the loan category; Use 1 is predicted by Dimension 2 and Use 2 by Dimension 1 for the no-loan category). Consequently, and in contrast to previous studies (Billman, 1996), "alignability" refers to the alignability of the use-relevant dimensions between superordinates rather than between the categories sharing a superordinate in the present study.

To the extent that between-category regularities in the use of category knowledge are a factor in learning multiple classification hierarchies, the alignability of use-relevant dimensions should affect how quickly these hierarchies are learned. When alignable dimensions predict subclassification in both hierarchies, we predict that learning those hierarchies will proceed more quickly and with fewer errors than when the dimensions are nonalignable. If the alignability effects seen in previous experiments were due simply to the number of dimensions needing attention in order to learn how to use the category, there should be no effects of alignability in the present experiment.

Method

Participants

Participants consisted of forty-eight University of Illinois students who participated for experimental credit or pay. The sessions lasted from 1 h. to 1 h. and 50 min.

Materials and Design

The materials were constructed to look like bank loan application forms (Appendix). Included on the forms were four, four-featured dimensions: Sales (wholesale, retail, mail order, internet), Advertising (television, radio, magazine, newspaper), Product (electronics, drugs, food, home furnishings), and International Market (Asia, South America, Europe, Australia).

Thirty-two loan applications were constructed, 16 in each category. Table 1 shows the abstract structure for the two categories of loan applications for each condition. Each dimension on the application (D1-D4) contained one of four possible features, denoted in the columns as values ranging from one to four. The features of two dimensions (D1 and D2) were perfectly predictive of superordinate category membership (superordinate-relevant dimensions).

Both of the superordinate-relevant dimensions were also predictive of the use of the category (use-relevant dimensions). One of the two use-relevant dimensions (indicated in bold type in Table 1) was predictive of the financial advice given to the applicant. The other dimension (indicated in italics in Table 1) was predictive of the consultant decision. The two remaining dimensions did not predict membership in either the superordinate or subordinate level categories.

Take, for example, the alignable-use condition in Table 1. Values on both the Sales dimension (D1) and the Advertising dimension (D2) predict membership in the superordinate category. If the feature of Sales was “wholesale” (indicated by a 1) or “retail” (indicated by a 2) or the feature of Advertising was “television” (indicated by a 1) or “radio” (indicated by a 2), the applicant received the loan. If the feature of Sales was “mail order” (indicated by a 3) or “internet” (indicated by a 4) or the feature of Advertising was “magazine” (indicated by a 3) or “newspaper” (indicated by a 4) the applicant did not receive the loan.

In this example, Sales and Advertising were also use-relevant dimensions, predicting subclassification within each superordinate. In the loan category, if the value of Sales was “wholesale,” the financial advice given was “product development” (indicated by 1 in the Use1 column). If it was “retail,” the financial advice was “expand line of merchandise” (as indicated by a 2 in the Use1 column). Similarly, if the value of Advertising was “television,” the consultant was “Ward” (indicated by 1 in the Use2 column). If the value of Advertising was “radio,” the consultant was “Smith” (indicated by 2 in the Use2 column). In the no-loan category, Sales (“mail order” or “internet”) also predicted financial advice (“sell some stock” or “focus on specialized market”) and Advertising (“magazine” or “newspaper”) predicted the consultant decision (“Jones” or “Brown”).

As can be seen in Table 1, the exemplars for the loan and no-loan categories are identical in both the alignable-use and nonalignable-use conditions. The superordinate-relevant and use-relevant dimensions are also identical. The difference between conditions is in the alignability of dimensions predictive of each use of the category. In the alignable-use condition, one superordinate-relevant dimension predicts financial advice for both categories and the other predicts the consultant for both categories. In the nonalignable-use condition, the dimension predictive of financial advice for one category predicts the consultant for the contrast category. For example, for the nonalignable-use condition, the loan category might be as above but the no-loan category would switch dimensions predictive of the uses, with Advertising predicting financial advice and Sales predicting consultant.

To insure that any learning differences between conditions was not due to the particular dimensions and their features, the use-relevant dimensions predictive of each use were counterbalanced over participants.

Procedure

Participants were told that the experiment examined how people learn to classify items and then use that information to make decisions related to the category. They were instructed to imagine that they were junior bank loan executives learning how to classify loan applications into those that receive business expansion loans and those that do not. Classification would be based on the information in the company's loan application form. After classifying the loan application, they would then use the information in the application to determine what financial advice to give the applicant. For companies that received a loan, the advice referred to how the loan money was to be spent ("expand the line of merchandise" or "product development"). For companies that did not receive a loan, the advice referred to how the company could generate revenue ("sell some stock" or "focus on specialized market"). Following the financial advice decision, they would then use the information in the application to determine which consultant would further process the loan. As with financial advice, the consultant decision consisted of four choices (two for each superordinate category) predicted by the variables in the loan application. Two of the choices were appropriate for applications receiving a loan (Smith or Ward) and two choices were appropriate for applications that did not receive a loan (Jones or Brown).

A sheet with the loan categories (loan/no-loan) and the corresponding financial advice and consultant decisions was visible throughout the experiment. Participants were given an example of the loan application form with X's replacing the features for the four dimensions. Their attention was then directed to those dimensions and they were told that they would use the information contained in those dimensions to make the loan, financial advice, and consultant decisions.

On each study trial, participants received a loan application form containing values on the four dimensions and made three responses. First, they classified the item (received a loan/did not receive a loan). Second, using the correct category, the participants then responded with one of the financial advice decisions. Finally, the participants responded with one of the consultant decisions. Following each response, the experimenter provided feedback about the correctness of each classification. The participants were given as much time as they wished to study the application before the next exemplar was presented.

Each block consisted of the random presentation of sixteen applications (eight loan, eight no-loan). Consequently, it took two complete blocks for the participant to see all thirty-two loan applications. All participants engaged in learning trials until they made no more than two mistakes on a block (out of 48 responses) or until they had completed twelve blocks.

For the analysis of blocks taken to learn both the superordinate[3] and subordinate category distinctions, participants who had not learned the loan/no-loan distinction, the financial advice, and/or the consultant decisions were recorded as having learned them by the thirteenth block. This was done as a conservative estimate of the next block they could have reached criterion had they been allowed to continue. To examine the effects of alignability on subclassification learning we averaged the number of blocks to learn each use-relevant subordinate category distinction to produce a combined blocks-to-learn-use score. We did the same with the number of errors to learn both uses to produce an errors-to-learn-use score.

Results and Discussion

If between-category alignability facilitates learning subordinate level distinctions then learning should be facilitated in the alignable-use condition because the same dimensions predict subclassification in both hierarchies. Earlier research shows an advantage for this condition when it was confounded with fewer attended dimensions. Do we still see facilitated learning in the alignable-use condition when the conditions are equated on the number of attended dimensions?

The distributions of scores were positively skewed so log transformations were performed. There was also significant variability introduced within each condition by the counterbalancing manipulation. To remove this source of variance for between-condition analyses, the log-transformed scores in the two counterbalancings within each condition were adjusted to have the same mean as the overall condition mean. This adjustment reduces the variability within each condition without affecting the condition means. Back-transformed means and standard deviations are reported.

When the dimensions predictive of use-relevant subclassifications within a hierarchy are alignable between hierarchies, learning is facilitated. As can be seen in Table 2, participants in the alignable-use condition learned the use-relevant subcategories for each hierarchy in fewer blocks and with fewer errors than participants in the nonalignable-use condition [$t(46) = 2.67$, $p < .01$; $t(46) =$

2.12, $p < .05$]. So, even when the number of attended dimensions is equated, the alignable-use condition shows an advantage in learning.

This learning advantage also held for each individual subclassification. Participants in the alignable-use condition learned the financial advice decisions in fewer blocks (5.9 versus 8.7) and made fewer errors (28.3 versus 44.8) than those in the nonalignable-use condition [$t(46) = 2.77$, $p < .01$; $t(46) = 2.38$, $p < .05$]. The same was true when learning the consultant decisions (alignable: blocks = 5.2, errors = 26.8; nonalignable: blocks = 8.0, errors = 42.5) [$t(46) = 2.54$, $p < .05$; $t(46) = 1.90$, $p = .06$].[4]

To summarize, between-category consistencies in the dimensions predictive of subclassification within a category facilitate learning multiple, use-relevant classification systems. Even when the number of dimensions needed for subclassification was held constant between conditions, we still see alignability effects on learning. When the dimensions important for subclassification are alignable between superordinates, learning is enhanced.

Experiment 1 controlled variables left to vary in other studies examining the effect of alignment on classification learning. As in Sifonis and Ross (1999), the current study eliminates differences between conditions in the category representations. Most importantly, it separated alignability from the number of dimensions important for making distinctions within a category. Demonstrating alignability effects while controlling for differences between conditions in the category representations and the number of attended dimensions suggests that such effects do not appear to be solely an artifact of the experimental designs used to study them.

Even though the alignability effects observed in Experiment 1 can no longer be attributed to the number of attended dimensions, it is possible that such effects arise from something other than the between-category alignability of the dimensions predictive of subclassification. An alternative explanation is that the alignable-use condition requires less complex functions to be learned than the nonalignable-use condition.[5] There are different ways one can think about the complexity of the functions in Experiment 1. One possibility is to consider it at the level of the mapping between features (e.g., Sales: Wholesale) and the particular use values (e.g., Financial advice: Product development). The number of mappings between feature and use values was the same for each condition (eight). In addition, the difficulty of classification at the superordinate level (loan/no loan) was exactly the same for the two conditions. So, if people were simply learning the mappings between features and use values, there was no difference in complexity.

Another possibility, however, is to consider complexity in terms of how dimensions, not features, related to uses. In the alignable-use condition, the association between a dimension and its use remained the same in both category systems, whereas in the nonalignable-use condition, the association between dimension and use changed as a function of superordinate category membership. Thus, by this view of complexity, the alignable-use condition was less complex. Even though the number of dimensions needed for category use were held constant between conditions, the manner in which those dimensions were used between category systems varied by condition. The complexity argument maintains that this difference requires something additional to be learned in the nonalignable-use condition, so learning was more difficult.

It is difficult to separate this type of complexity argument from the alignability argument, because the difference in the relation between the dimensions and the uses is both a complexity and an alignability difference. The important point about the complexity argument is that it does not view any particular complexity as critical; just that overall complexity will determine learning. In contrast, the alignability account claims that overall complexity is not the primary determinant of ease of learning, but rather that one type of complexity, alignability, is especially important. In other words, the argument is one of how complexity relates to psychological complexity. The complexity view is that psychological complexity (e.g., as measured by difficulty in learning) is a function of overall objective complexity. The alignability view is that some types of complexity, in particular alignability, are especially influential. We do not dispute that, keeping everything else constant, variations in alignability will lead to differences in complexity. We believe that between-category consistencies in how dimensions are used are a component of what it means for categories or classification systems to be alignable. Our claim is that the large differences in learning in Experiment 1 were due to this relatively subtle difference in a particular complexity, alignability. Despite the fact that alignability is usually confounded with overall complexity, this does not mean that the two views are indistinguishable. The purpose of Experiment 2 is to investigate whether overall complexity is sufficient for understanding the learning facilitation in the alignable-use condition.

Experiment 2

Experiment 2 compares learning in a complex alignable condition to learning in a less complex, nonalignable condition. The nonalignable condition was made less complex than the alignable condition by requiring only half as many category use decisions compared to the alignable condition. Participants in the alignable-use condition were treated exactly as in Experiment 1 (see top half of Table 3). They had to learn the loan/no-loan superordinate category distinctions and both use-relevant subordinate-level distinctions. Participants in the nonalignable-use condition learned only a subset of what was required in Experiment 1 (See bottom half of Table 3). They had to learn the same loan/no-loan superordinate category distinctions as participants in the alignable-use condition. However, they only had to learn one of the two use-relevant subordinate-level distinctions (e.g., financial advice but not consultant). Because subclassification in the nonalignable-use condition was a subset of subclassification in the alignable-use condition, the complexity was reduced relative to the alignable-use condition. Participants had fewer distinctions to learn at one time and four fewer feature-to-use mappings to establish. Note that the conditions remain equivalent in terms of the category representations and the number of dimensions needed to make use-relevant distinctions within each category. The conditions differ only in the alignability of the dimensions predictive of category use and the complexity of the learning task.

Although one can argue about how to equate complexity of alignability with the complexity of feature-value mappings, our approach was to make the nonalignable-use condition less complex in terms of the feature-value mappings (four instead of eight), as can be seen from comparing the two conditions in Table 3. If overall complexity was the reason behind the results of Experiment 1, then one would expect facilitated learning in the less complex, nonalignable condition of Experiment 2. Participants should learn to make use-relevant category distinctions in fewer blocks and with fewer errors in the nonalignable-use condition compared to the alignable-use condition. If, however, differences between conditions in alignability produced the results seen in Experiment 1, then even this large reduction of objective complexity in the nonalignable-use condition relative to that in the alignable-use condition may not be enough to overcome the alignability differences. Thus, even with only one use to learn (and far fewer mappings), participants in the nonalignable-use condition may not learn the classification systems in fewer blocks or make fewer errors than participants in the alignable-use condition.

Methods

Participants

Participants consisted of 64 Oakland University students who participated as volunteers or for experimental credit. The data from eight participants were excluded because they failed to learn the classification distinction (loan/no-loan) within 16 blocks or because time ran out before they were able to finish all the learning blocks.[6] The data from the remaining 56 participants were analyzed. The sessions lasted from 1 h. to 1 h. and 50 min.

Materials and Design

The study materials were as in Experiment 1. The design was similar to that used in Experiment 1 (See Table 3). Complexity was manipulated in terms of the number of category use decisions participants needed to learn in each condition. In the alignable-use condition, participants had to learn both the financial advice and consultant decisions. In design, the alignable-use condition of the present experiment is identical to the alignable-use condition of Experiment 1. As in Experiment 1, category use decisions were alignable between categories in that the same dimensions predicted the same use in both superordinate categories.

Complexity was reduced in the nonalignable-use condition relative to the alignable-use condition by having participants learn only one category use decision. Half the participants in the nonalignable-use condition learned the financial advice decision whereas the other half learned the consultant decision. Category use decisions were nonalignable between categories in that different dimensions predicted category use within each superordinate. The design of the present experiment is such that the number of dimensions needing attention to learn the use of both categories remains the same between conditions.

Procedure

The procedure for the alignable-use condition was identical to that used in Experiment 1 except the criteria for completion was extended to 16 blocks. The procedures for participants in the nonalignable-use condition were similar to those for participants in the alignable-use condition with the exception that participants learned only one of the two category use decisions. The nonalignable-use condition study criteria were no more than one mistake on a block (out of 32 total responses) or completion of 16 blocks. When participants in either condition failed to learn the correct use(s) of the category by the sixteenth block, they were counted as having learned by the seventeenth block for analysis purposes. This was done as a conservative estimate of the next block they could have reached criterion had they been allowed to continue.

As in Experiment 1, we averaged the number of blocks and errors to learn each use in the alignable-use condition. For participants in the nonalignable-use condition there was only the one use (either the financial advice or the consultant decision).

Results and Discussion

The critical issue is determining whether the alignability effects exhibited in Experiment 1 resulted from between-category consistencies in structure predictive of category use or were due to reduced complexity in the alignable-use condition. Because the nonalignable-use condition was less complex than the alignable-use condition in the current experiment, a complexity view predicts that learning use-relevant distinctions within each category system should be facilitated compared to learning in the alignable-use condition. As in Experiment 1, transformations were necessary to correct for negatively skewed distributions and variability introduced by the counterbalancing manipulation. Back-transformed means and standard deviations are reported.

Learning to make use-relevant distinctions in a less complex, nonalignable-use condition was not facilitated relative to that seen in a more complex, alignable-use condition. There were no significant differences between conditions on either blocks to learn subclassification within both hierarchies [t(54) < 1, n.s.] or errors made during learning [t(54) < 1, n.s.] (See Table 4).[7] Even though participants in the alignable-use condition had to learn more than participants in the nonalignable-use condition, the means indicate a nonsignificant advantage for learning in the alignable-use condition on all variables.

Although it is always difficult to interpret null results, the trends in the category learning means suggest it is unlikely the alignability effects seen in Experiment 1 were due to reduced complexity in the alignable-use condition. In the current experiment, participants in the alignable-use condition had to overcome a strong disadvantage in complexity compared to participants in the nonalignable-use condition yet performed equivalently. Effect sizes estimates of the differences between conditions (.27 for blocks, .12 for errors) (Cohen, 1988) and the superior performance in the alignable case suggests that overall complexity is not the main determinant of performance difficulty.

General Discussion

Experiments 1 and 2 demonstrate the importance of between-category structural regularities in category use on learning multiple, use-relevant category hierarchies. Experiment 1 found that between-category consistencies in the dimensions predictive of subclassification facilitate learning multiple category hierarchies. This facilitation was not due to differences between conditions in the category representations learned by participants or in the amount of information needed for subclassification in the hierarchies. Because we divorced alignability from the number of dimensions needed to use category knowledge in Experiment 1, the only difference between conditions lay in the structural regularities important for the use of the categories. When the values predicting use-relevant, subordinate-level distinctions within a category were associated with the same dimensions in both hierarchies, learning was facilitated.

Experiment 2 demonstrated that the alignability effects observed in Experiment 1 could not be explained by differences between conditions in the overall complexity of the learning task. Reducing the complexity of the relations to be learned in a nonalignable category hierarchy relative to an alignable hierarchy did not result in facilitated learning. In fact, rather than performing better than participants in the alignable-use condition of Experiment 2, participants in the less complex, nonalignable-use condition took somewhat longer and made slightly more errors learning the category hierarchies. This is especially noteworthy given the fewer associations they were required to learn relative to those in the alignable-use condition.

To summarize, when subordinate-level distinctions within category hierarchies can be made on the basis of alignable differences between those hierarchies, learning is facilitated. This is true even when the representations and the number of dimensions needed for subclassification are held constant between conditions. Even when the overall complexity of a nonalignable condition is less than that of an alignable condition, we fail to see facilitated classification learning.

The results of Experiments 1 and 2 suggest that alignment effects on classification learning are not simply the result of differences between conditions in the objective complexity of the learning task. Rather, they reflect a commitment on the part of the learner to a type of psychological complexity - that of alignment. Part of this commitment to alignment is the belief that the structure relevant to the use of one category system will also be relevant to the use of a comparable system. In terms of learning multiple classification hierarchies simultaneously, a commitment to alignability involves the assumption that dimensions important for making use-relevant distinctions within one hierarchy will also be important for making the same types of distinctions within the other hierarchy. In the present study, the similarity of the category representations in both hierarchies (differing only in the instantiation of shared dimensions) and the uses towards which they are applied probably encourage such assumptions. When this assumption is warranted, as it is in the alignable-use condition, learning is facilitated over situations in which the assumption is not justified (as in the nonalignable-use condition).

Learning Multiple Classification Systems

Examining the learning of multiple classification systems has both provided a better understanding of how multiple categories are learned together and extended our understanding of structural alignment in category learning. More specifically, the present study suggests that the commitment to alignability is not confined to classification distinctions made within a single superordinate. It extends to whole systems of relations such as those acquired when learning multiple, use-relevant category hierarchies.

Prior to the current study, research examining alignment effects on classification learning restricted its focus to the effect of between-category consistencies and inconsistencies in the dimensions predictive of classification in categories sharing a superordinate. This research demonstrated that classification predicted by alignable differences between categories, is easier to learn than classification based on nonalignable differences (Billman, 1996; Kaplan, 1999, Exps. 2 & 3; Lassaline & Murphy, 1998; Waxman,et al., 1997). Two constructs, consistent contrast and structural alignment, have been used to explain this alignment effect in classification learning.

Consistent contrast refers to a principle of category learning in which people are biased to learn sets of categories that contrast in the same way across the categories (Billman, 1996). In such categories, the same dimensions distinguish between categories within the same contrast set (mutually exclusive categories within a superordinate). Because the principle of consistent contrast is restricted to categories sharing a superordinate, whether or not the same dimension can be used for subclassification in multiple classification systems should not affect learning (Billman, 1996). Demonstrating that the learning of multiple classification systems is affected by between-category similarities in category structures suggests that the degree to which consistent contrast explains classification learning needs to be examined further. For example, the results of the current study can be explained by consistent contrast if both the loan and the no-loan categories are considered subordinates of the superordinate category "companies applying for a loan". An interesting avenue for future research would be to determine the level in a hierarchy in which categories need to share a common superordinate in order for consistent contrast to apply.

In comparison, the principle of structural alignment theory is such that it can easily be applied to predict and explain the formation of multiple classification systems. A structural alignment account of classification learning assumes that during learning category representations are compared in order to identify the elements of category structure important for distinguishing between category members. Part of this process involves the alignment of the representations to find the maximal structurally consistent match between them (Gentner, 1983). If the representations being compared share much of their structure then the principle of structural alignment causes the learner to infer that the structure important for the use of one representation will be so for the other representation as well. This is true whether classifying category members into categories sharing a superordinate or learning multiple, use-relevant classification hierarchies.[8]

The present study demonstrated that multiple classification hierarchies are not learned independently from each other. Learning was facilitated when the structure important for subclassification within one hierarchy was also important for subclassification within the other. If a principle such as structural alignment is operating during classification learning, it suggests that we must extend the scope of the commitment to alignability to include whole systems of relations such as those acquired when learning multiple, use-relevant category hierarchies.

Learning the Use of Multiple Categories

The results of the present study also inform us about the manner in which people learn to use multiple categories simultaneously. Ample evidence suggests that there is a close relationship between category use and classification. The use of category knowledge has been shown to result in use-relevant subdivisions of domain knowledge (Boster & Johnson, 1989; Medin, et al., 1997; Ross & Murphy, 1999). Category use has also been shown to affect classification performance (Ross, 1996, 1997, 1999, 2000).

The category use literature frequently employs a feature prediction paradigm such as that used in the present study to examine the effects of category use on the category representation. These studies have demonstrated that the features important for satisfying category goals (use-relevant features) become more central to the category representation. These features are generated more frequently in listing tasks, are judged to have occurred more frequently, and support more accurate classification than features not important for the use of the category (Ross, 1996, 1997). Importantly, these category use effects have also been demonstrated in tasks that do not involve feature prediction (Ross, 1999). This suggests that the results of the current study might very well apply to situations in which people are required to learn the use of two categories simultaneously. If so, between-category structural regularities in the use of category knowledge will be important when learning to use multiple categories simultaneously. The implication is as that with classification learning - learning to apply category knowledge is not an entirely within-category phenomenon. When learning the use of multiple categories simultaneously we attend to between-category structural regularities in the use of category knowledge. When the same structural regularities are important for the use of similar contrasting categories, learning is facilitated over situations in which the structure important for category use differs between categories.

Conclusion

These experiments demonstrate the importance of between-category structural regularities in learning to make use-relevant distinctions within multiple classification systems. Because between-category consistencies in the use of category knowledge affect classification learning, this suggests that multiple classification systems are not learned independently of one another. The present studies have also demonstrated that the objective complexity of the learning task is not the sole explanation for the effects of such consistencies on classification learning. Rather, the results appear to reflect assumptions about the psychological complexity of the learning task. Specifically, they reflect the assumption that the category structure relevant for the use of one category system will also be relevant for the use of a similar system being learned at the same time.

References

Anderson, J.R. (1991). The adaptive nature of human categorization. Psychological Review, 98, 409-429.

Billman, D. (1996). Structural biases in concept learning: Influences from multiple functions. In D.L. Medin (Ed.) The psychology of learning and motivation (Vol. 35, pp. 283-321). San Diego, Academic Press.

Boster, J.S. & Johnson, J.C. (1989). Form or function: A comparison of expert and novice judgments of similarity among fish. American Anthropologist, 91, 866-889.

Cohen, J. (1988). Statistical power analysis for the behavioral sciences. Hillsdale, NJ: Erlbaum.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.

Goldstone, R. (1994). The role of similarity in categorization: Providing a groundwork. Cognition, 52, 125-157.

Kaplan, A. (2000). Alignability and prior knowledge in category learning. Unpublished doctoral dissertation, University of Illinois, Urbana-Champaign.

Lassaline, M.E., & Murphy, G. L. (1998). Alignment and category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 24, 144-160.

Medin, D.L., Lynch, E.B., Coley, J.D., & Atran, S. (1997) Categorization and reasoning among tree experts: Do all roads lead to Rome? Cognitive Psychology, 32, 49-96.

Ross, B.H. (1996). Category representations and the effects of interacting with instances. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 1249-1265.

Ross, B.H. (1997). The use of categories affects classification. Journal of Memory and Language, 37, 240-267.

Ross, B.H. (1999). Post-classification category use: The effects of learning to use categories after learning to classify. Journal of Experimental Psychology: Learning, Memory and Cognition, 25, 743-757.

Ross, B.H. (2000). The effects of category use on learned categories. Memory & Cognition, 28, 51-63.

Ross, B.H. & Murphy, G.L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. Cognitive Psychology, 38, 495-553.

Sifonis, C.M. & Ross, B.H. (1999, August). The effects of alignability on learning the use of categories. Poster presented at the annual meeting of the Cognitive Science Society, Vancouver, BC.

Waxman, S.R., Lynch, E.B., Casey, K.L., & Baer, L. (1997). Setters and samoyeds: The emergence of subordinate level categories as a basis for inductive inference in preschool-age children. Developmental Psychology, 33, 1074-1090.

Wisniewski, E.J. & Markman, A.B. (1997). Feature alignment in concept learning. Manuscript submitted for publication.

Yamauchi, T., & Markman, A.B. (2000). Learning categories composed of varying instances: The effect of classification, inference, and structural alignment.

Memory & Cognition, 28, 64-78.

Author Note

Cynthia M. Sifonis, Department of Psychology, Oakland University; Brian H. Ross, Department of Psychology, Beckman Institute.

This study was supported by NSF grant SBR 97-20304 . We thank Gregory L. Murphy, Arthur B. Markman, Bob E. Rehder and three anonymous reviewers for helpful comments on an earlier version of this manuscript. We also thank Devon Baxter, Angela Castellani, Stacey Delikta, Alex Fedorowich, and Kim Lewis for help in testing subjects.

Parts of the current manuscript were presented at the annual meeting of the Cognitive Science Society, Vancouver, BC.

Correspondence concerning this article should be addressed to Cynthia M. Sifonis, Department of Psychology, 222 Pryale Hall, Oakland University, Rochester, MI 48309 - 4489. Electronic mail may be sent to sifonis@oakland.edu.

Footnotes

[1] It is important to state that even though we are studying the formation of use-relevant category hierarchies, we are not making any claims as to the manner in which use-relevant category hierarchies are learned compared to taxonomic hierarchies. Examining the role of contrastive information in learning use-relevant hierarchies will extend our knowledge of the subject. Whether or not taxonomic category hierarchies are learned in the same fashion is not addressed in the current paper.

[2] Even though Sifonis and Ross (1999) was designed to examine category use learning, participants in that study could be said to be learning multiple classification hierarchies. Deciding which of two types of financial advice to provide the applicant can be viewed as a classification task in which the loan/no-loan categorization is further divided into one of two subcategories (Anderson, 1991). For each category, one subcategory receives one type of financial advice with the other receiving the other type of advice. Even though there is some evidence suggesting that the manner in which a task is presented (e.g., as either classification or feature prediction) affects task performance (Yamauchi & Markman, 2000), it is possible for participants to view the prediction as a subclassification. To facilitate discussion of the issues in the present study, we will continue to refer to the Sifonis and Ross (1999) study as one examining subclassification rather than category use

[3] The focus of the present study is on the effects of alignability in the dimensions predictive of subordinate level distinctions on learning multiple

classification systems. Although we report the means for classification learning we will not discuss alignability effects on learning these distinctions in the present paper. We present statistics in the footnotes for interested readers.

[4] There were no significant alignability effects on learning the superordinate level distinctions in either the number of blocks needed to learn the classification (t(46) = 1.40, n.s.) or in the number of errors made during learning (t(46) = .86, n.s.).

[5] We would like to thank an anonymous reviewer for this explanation of the results.

[6] Participants for Experiment 2 were drawn from a different subject pool than those from Experiment 1. Pretesting indicated that the number of blocks to reach criterion would have to be increased to allow participants from this pool time enough to learn the classification hierarchies.

[7] There were no significant alignability effects on learning the superordinate level distinctions in either the number of blocks needed to learn the classification [t(54) = .89, n.s.] or in the number of errors made during learning [t(54) = .86, n.s.].

[8] Structural alignment theory does not specifically predict whether or not multiple classification systems will be learned independently. However, if comparisons are made between systems during learning, then structural alignment theory allows for structural consistencies and inconsistencies between systems to affect learning.

Appendix

Loan Application Form

Application for Business Expansion Loan

OMB No. 1840-0717 Form Approved Exp. Date 03/31/99

Warning: Any person who knowingly makes a false statement or misrepresentation on this form is to penalties which may include fines or imprisonment under the United States Criminal Code and 20 U.S.C. 1097

Guarantor Identification

First National City Bank

1st

Borrower Section **Please print neatly or type. Read the instructions carefully**

1. Company Name [redacted]	**2.** Company Address [redacted]
3. Industry Drug	**4.** Type of Sales Retail
5. International Market Europe	**6.** Method of Advertising Radio
7. Loan Period (Month/Year) 9/98 - 9/03	**8.** Loan Assistance Requested **$** 10,000,000 **.00**

Promissory Note

Promise to Pay: I promise to pay to the lender, or a subsequent holder of this Promissory Note, all sums dispersed (hear after "loan" or "loans") under the terms of this note, plus interest and other fees which may become due as provided by this Note. If I fail to make payments on this Note when due, I will also pay reasonable collection costs, including attorney's fees, court costs, and collection fees. I understand that I may cancel or reduce the size of any loan by refusing to accept any disbursement that is issued.

THIS IS A LOAN(S) THAT MUST BE REPAID

9. Borrower's Signature ____________________ Today's Date (Month/Day/Year)_________________

Bank Section **To be completed by an authorized bank official**

10. Lender Name First National City Bank	**11.** Lender Code 100XX1X2
12. Telephone Number (217) 333-2012	**13.** Amount(s) Approved **$** [redacted] **.00**

Table 1

Experiment 1 Design

Exemplars for the Alignable-use Condition

Loan Category Exemplars

Item	**D1**	*D2*	D3	D4	**Use1**	*Use2*
1	**1**	*1*	1	4	**1**	*1*
2	**1**	*2*	2	3	**1**	*2*
3	**1**	*1*	3	2	**1**	*1*
4	**1**	*2*	4	1	**1**	*2*
5	**1**	*1*	2	3	**1**	*1*
6	**1**	*2*	3	2	**1**	*2*
7	**1**	*1*	4	1	**1**	*1*
8	**1**	*2*	1	4	**1**	*2*
9	**2**	*1*	3	2	**2**	*1*
10	**2**	*2*	4	1	**2**	*2*
11	**2**	*1*	1	4	**2**	*1*
12	**2**	*2*	2	3	**2**	*2*
13	**2**	*1*	4	1	**2**	*1*
14	**2**	*2*	1	4	**2**	*2*
15	**2**	*1*	2	3	**2**	*1*
16	**2**	*2*	3	2	**2**	*2*

No-Loan Category Exemplars

Item	**D1**	*D2*	D3	D4	**Use1**	*Use2*
1	**3**	*3*	1	4	**3**	*3*
2	**3**	*4*	2	3	**3**	*4*
3	**3**	*3*	3	2	**3**	*3*
4	**3**	*4*	4	1	**3**	*4*
5	**3**	*3*	2	3	**3**	*3*
6	**3**	*4*	3	2	**3**	*4*
7	**3**	*3*	4	1	**3**	*3*
8	**3**	*4*	1	4	**3**	*4*
9	**4**	*3*	3	2	**4**	*3*
10	**4**	*4*	4	1	**4**	*4*
11	**4**	*3*	1	4	**4**	*3*
12	**4**	*4*	2	3	**4**	*4*
13	**4**	*3*	4	1	**4**	*3*
14	**4**	*4*	1	4	**4**	*4*
15	**4**	*3*	2	3	**4**	*3*
16	**4**	*4*	3	2	**4**	*4*

Exemplars for the Nonalignable-use Condition

Loan Category Exemplars

Item	**D1**	*D2*	D3	D4	**Use1**	*Use2*
1	**1**	*1*	1	4	**1**	*1*
2	**1**	*2*	2	3	**1**	*2*
3	**1**	*1*	3	2	**1**	*1*
4	**1**	*2*	4	1	**1**	*2*
5	**1**	*1*	2	3	**1**	*1*
6	**1**	*2*	3	2	**1**	*2*
7	**1**	*1*	4	1	**1**	*1*
8	**1**	*2*	1	4	**1**	*2*
9	**2**	*1*	3	2	**2**	*1*
10	**2**	*2*	4	1	**2**	*2*
11	**2**	*1*	1	4	**2**	*1*
12	**2**	*2*	2	3	**2**	*2*
13	**2**	*1*	4	1	**2**	*1*
14	**2**	*2*	1	4	**2**	*2*
15	**2**	*1*	2	3	**2**	*1*
16	**2**	*2*	3	2	**2**	*2*

No-Loan Category Exemplars

Item	*D1*	**D2**	D3	D4	**Use1**	*Use2*
1	*3*	**3**	1	4	**3**	*3*
2	*3*	**4**	2	3	**4**	*3*
3	*3*	**3**	3	2	**3**	*3*
4	*3*	**4**	4	1	**4**	*3*
5	*3*	**3**	2	3	**3**	*3*
6	*3*	**4**	3	2	**4**	*3*
7	*3*	**3**	4	1	**3**	*3*
8	*3*	**4**	1	4	**4**	*3*
9	*4*	**3**	3	2	**3**	*4*
10	*4*	**4**	4	1	**4**	*4*
11	*4*	**3**	1	4	**3**	*4*
12	*4*	**4**	2	3	**4**	*4*
13	*4*	**3**	4	1	**3**	*4*
14	*4*	**4**	1	4	**4**	*4*
15	*4*	**3**	2	3	**3**	*4*
16	*4*	**4**	3	2	**4**	*4*

Note. D1-D4 indicates exemplar dimensions and their features. Feature values of 1 or 2 for dimensions D1 or D2 are predictive of the loan category. Values of 3 or 4 for dimensions D1 or D2 are predictive of the no-loan category. Use1 indicates the financial advice decision. Use2 indicates the consultant decision. The dimensions in bold correspond to and predict the financial advice in bold. The dimensions in italics correspond to and predict the consultant advice in italics.

Table 2

Experiment 1: The Effects of Alignable Uses on Learning Multiple Classification Systems

	Use Blocks		Use Errors		Loan Blocks		Loan Errors	
	Align	Nonalign	Align	Nonalign	Align	Nonalign	Align	Nonalign
Mean	5.7	8.3	29.5	45.0	3.1	4.2	8.3	11.2
St. Dev	3.1	3.7	21.8	29.9	2.4	3.0	10.4	14.7

Note. **Align** indicates the scores for the alignable-use condition. **Nonalign** indicates the scores for the nonalignable-use condition. **Use Blocks/Errors** indicates the mean number of blocks/errors to learn both the financial advice and consultant decisions (averaged across both uses). **Loan Blocks/Errors** indicates the mean number of blocks/errors to learn the loan/no-loan decision.

Table 3

Experiment 2 Design

Exemplars for the Alignable-use Condition

Loan Category Exemplars

Item	**D1**	*D2*	D3	D4	**Use1**	*Use2*
1	**1**	*1*	1	4	**1**	*1*
2	**1**	*2*	2	3	**1**	*2*
3	**1**	*1*	3	2	**1**	*1*
4	**1**	*2*	4	1	**1**	*2*
5	**1**	*1*	2	3	**1**	*1*
6	**1**	*2*	3	2	**1**	*2*
7	**1**	*1*	4	1	**1**	*1*
8	**1**	*2*	1	4	**1**	*2*
9	**2**	*1*	3	2	**2**	*1*
10	**2**	*2*	4	1	**2**	*2*
11	**2**	*1*	1	4	**2**	*1*
12	**2**	*2*	2	3	**2**	*2*
13	**2**	*1*	4	1	**2**	*1*
14	**2**	*2*	1	4	**2**	*2*
15	**2**	*1*	2	3	**2**	*1*
16	**2**	*2*	3	2	**2**	*2*

No-Loan Category Exemplars

Item	**D1**	*D2*	D3	D4	**Use1**	*Use2*
1	**3**	*3*	1	4	**3**	*3*
2	**3**	*4*	2	3	**3**	*4*
3	**3**	*3*	3	2	**3**	*3*
4	**3**	*4*	4	1	**3**	*4*
5	**3**	*3*	2	3	**3**	*3*
6	**3**	*4*	3	2	**3**	*4*
7	**3**	*3*	4	1	**3**	*3*
8	**3**	*4*	1	4	**3**	*4*
9	**4**	*3*	3	2	**4**	*3*
10	**4**	*4*	4	1	**4**	*4*
11	**4**	*3*	1	4	**4**	*3*
12	**4**	*4*	2	3	**4**	*4*
13	**4**	*3*	4	1	**4**	*3*
14	**4**	*4*	1	4	**4**	*4*
15	**4**	*3*	2	3	**4**	*3*
16	**4**	*4*	3	2	**4**	*4*

Exemplars for the Nonalignable-use Condition

Loan Category Exemplars

Item	**D1**	D2	D3	D4	**Use**
1	**1**	1	1	4	**1**
2	**1**	2	2	3	**1**
3	**1**	1	3	2	**1**
4	**1**	2	4	1	**1**
5	**1**	1	2	3	**1**
6	**1**	2	3	2	**1**
7	**1**	1	4	1	**1**
8	**1**	2	1	4	**1**
9	**2**	1	3	2	**2**
10	**2**	2	4	1	**2**
11	**2**	1	1	4	**2**
12	**2**	2	2	3	**2**
13	**2**	1	4	1	**2**
14	**2**	2	1	4	**2**
15	**2**	1	2	3	**2**
16	**2**	2	3	2	**2**

No-Loan Category Exemplars

Item	D1	**D2**	D3	D4	**Use**
1	3	**3**	1	4	**3**
2	3	**4**	2	3	**4**
3	3	**3**	3	2	**3**
4	3	**4**	4	1	**4**
5	3	**3**	2	3	**3**
6	3	**4**	3	2	**4**
7	3	**3**	4	1	**3**
8	3	**4**	1	4	**4**
9	4	**3**	3	2	**3**
10	4	**4**	4	1	**4**
11	4	**3**	1	4	**3**
12	4	**4**	2	3	**4**
13	4	**3**	4	1	**3**
14	4	**4**	1	4	**4**
15	4	**3**	2	3	**3**
16	4	**4**	3	2	**4**

Note. D1-D4 indicates exemplar dimensions and their features. Feature values of 1 or 2 for dimension D1 or D2 are predictive of the loan category. Values of 3 or 4 for dimensions D1 or D2 are predictive of the no-loan category. For the alignable-use condition, Use1 indicates the financial advice decision and Use2 indicates the consultant decision. For the nonalignable-use condition, Use indicates either the financial advice or the consultant decision. The dimensions in bold correspond to and predict the financial advice in bold. For the Alignable-use condition, the dimensions in italics correspond to and predict the consultant advice in italics.

Table 4

Experiment 2: The Effects of Complexity on Learning Multiple Classification Systems

	Use Blocks		Use Errors		Loan Blocks		Loan Errors	
	Align	Nonalign	Align	Nonalign	Align	Nonalign	Align	Nonalign
Mean	7.3	8.4	44.8	49.0	4.2	5.7	15.0	25.5
St. Dev	3.2	4.9	25.6	42.3	2.7	4.6	14.8	29.3

Note. **Align** indicates the scores for the alignable-use condition. **Nonalign** indicates the scores for the nonalignable-use condition. **Use Blocks/Errors** indicates the mean number of blocks/errors to learn both the financial advice and consultant decisions (averaged across both uses). **Loan Blocks/Errors** indicates the mean number of blocks/errors to learn loan/no-loan decision.