

Co-Training as a Human Collaboration Policy

Xiaojin Zhu and Bryan R. Gibson and Timothy T. Rogers[†]

Department of Computer Sciences and [†]Psychology
University of Wisconsin-Madison, Madison, WI, USA. 53706
{jerryzhu, bgibson}@cs.wisc.edu, ttrogers@wisc.edu

Abstract

We consider the task of human collaborative category learning, where two people work together to classify test items into appropriate categories based on what they learn from a training set. We propose a novel collaboration policy based on the Co-Training algorithm in machine learning, in which the two people play the role of the base learners. The policy restricts each learner’s view of the data and limits their communication to only the exchange of their labelings on test items. In a series of empirical studies, we show that the Co-Training policy leads collaborators to jointly produce unique and potentially valuable classification outcomes that are not generated under other collaboration policies. We further demonstrate that these observations can be explained with appropriate machine learning models.

Though human learning abilities are remarkable in many respects, they are also constrained in ways that may seem puzzling to machine learning. As one example, people can have difficulty learning nonlinear decision boundaries without extensive supervision (Love 2002). As another example, psychologists often distinguish between feature dimensions that are “separable” versus “integral”. For separable features (e.g. color and shape), people can selectively attend to one dimension without processing the other. For integral dimensions (e.g. color saturation and brightness) they cannot. In learning problems that are identical from a machine-learning point of view, humans can show quite different patterns of behavior depending on whether the dimensions are integral or separable. For instance, people have difficulty learning non-axis-parallel boundaries for separable but not for integral feature dimensions (Nosofsky and Palmeri 1996; Ashby and Maddox 1990).

This paper considers whether these characteristics of human learning can be altered by leveraging insights from a machine learning algorithm, namely Co-Training. Co-Training uses unlabeled data to improve learning by encouraging agreement among multiple “base” machine learners, each exposed to a different “view” of the data (see below). The classic Co-Training algorithm (Blum and Mitchell 1998) and its extensions such as Co-EM (Nigam and Ghani 2000), Tri-Training (Zhou and Li 2005), and multiview learning (Brefeld et al. 2006) have enjoyed considerable

empirical success and theoretical justification (Johnson and Zhang 2007; Balcan and Blum 2010) in machine learning.

One often under-appreciated fact about Co-Training is that it has a different inductive bias, and so can produce quite different classification results from supervised learning. Figure 1(a) shows a “diamond” dataset with four clusters, with just one labeled item from each class (blue and red points). The task is to classify the unlabeled items (black dots). Supervised learning with the 1-nearest-neighbor (1NN) algorithm¹ learns a diagonal decision boundary in Figure 1(b). In contrast, with the same 1NN as base learners the Co-Training algorithm learns a very different solution (Figure 1(c)), grouping the top and bottom clusters together in the red class, and the left and right clusters in the blue class.




Figure 1: On this “diamond” dataset, supervised learning and Co-Training, both with 1NN classifiers, produce drastically different outcomes.

The linearly non-separable classification achieved by Co-Training is just the kind of solution that human beings have difficulty learning without extensive supervision (Love 2002). In the current work we consider whether the Co-Training algorithm can be used to design a *collaboration policy* for human participants that will promote learning of such “difficult” classifications over the linearly separable outcomes that individuals are prone to acquire on their own. Under this policy, each individual in the collaboration is treated as a “base” learner; each is exposed to a different “view” of the data; and the learning set-up is designed to promote agreement among the collaborators. We empirically assessed behavior in such teams for learning problems with

¹1NN classifiers are closely related to the generalized context model (Nosofsky 1986) in cognitive psychology, which assumes that people store training items in memory and categorize new items by computing similarities to these stored training items.

both psychologically-separable and integral stimulus dimensions, and compared performance to individual learners and to teams collaborating under an alternative policy. In simple learning problems like that shown in Figure 1 we will see that our Co-Training collaboration policy leads participants to learn classifications typically thought to be very difficult for humans, and also to show more homogeneous behavior for stimuli defined along separable versus integral dimensions. Though we do not extend the approach to a real-world learning problem here, we will consider how the approach might be used to design collaboration policies for such problems in cases where individuals have difficulty learning the appropriate classifications.

Review of the Co-Training Algorithm

We first review the classic Co-Training algorithm of (Blum and Mitchell 1998) as it is closely related to our policy. Assume that each item is parametrized by a feature vector \mathbf{x} and has a corresponding class label y . The input consists of l labeled items $(\mathbf{x}_1, y_1) \dots (\mathbf{x}_l, y_l)$ and u unlabeled items $\mathbf{x}_{l+1} \dots \mathbf{x}_{l+u}$. The goal is to learn a classifier $f : \mathbf{x} \mapsto y$ using both the labeled and unlabeled data.

Further assume that the feature vector can be split into two parts (called “views”): $\mathbf{x} = \begin{pmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{pmatrix}$. The Co-Training algorithm trains two base learners $f^{(1)} : \mathbf{x}^{(1)} \mapsto y$ and $f^{(2)} : \mathbf{x}^{(2)} \mapsto y$, each working exclusively on one view. In the beginning, these two base learners are trained on the labeled data. More specifically, $f^{(1)}$ is trained with the first view of the labeled data $(\mathbf{x}_1^{(1)}, y_1) \dots (\mathbf{x}_l^{(1)}, y_l)$. Subsequently, whenever $f^{(1)}$ encounters an item \mathbf{x} during training or prediction, it always works with the first view $\mathbf{x}^{(1)}$ of the item only and disregards the second view $\mathbf{x}^{(2)}$. $f^{(2)}$ operates similarly, working only with the second view. The ingenuity is in how the unlabeled data is utilized in an iterative fashion: At each iteration, $f^{(1)}$ classifies a few unlabeled items that it is most confident about and passes these and their predicted labels as *additional training data* to $f^{(2)}$. Simultaneously, $f^{(2)}$ reciprocates. Co-Training then updates both base learners with this additional “pseudo-labeled” data. This repeats until the unlabeled data is exhausted. A slightly simplified version of Blum and Mitchell’s Co-Training algorithm is given in Algorithm 1. To classify a new test item $\tilde{\mathbf{x}}$, one can compare the predictions $f^{(1)}(\tilde{\mathbf{x}}^{(1)})$ and $f^{(2)}(\tilde{\mathbf{x}}^{(2)})$ and pick the one with higher confidence.

Co-Training is a “wrapper” method in that the two base learners $f^{(1)}$ and $f^{(2)}$ can be any learning systems. The only requirement is that each base learner has a notion of confidence, which is used to select which unlabeled items to turn into pseudo labeled data for the other view. Importantly for this work, being a wrapper method enables Co-Training to treat two human collaborators as the base learners.

It is important to understand the conditions under which Co-Training will succeed. We present the sufficient conditions in the original analysis (Blum and Mitchell 1998), but with new interpretations geared toward our collaboration policy for human learning.

Input: labeled and unlabeled data where each item has two views; learning speed s .

Initialize $L_1 = L_2 =$ labeled data

repeat

Train $f^{(1)}$ from L_1 , $f^{(2)}$ from L_2 .

Classify unlabeled items with $f^{(1)}$, $f^{(2)}$ separately.

Add $f^{(1)}$ ’s top s most confident predictions $(\mathbf{x}, f^{(1)}(\mathbf{x}))$ to L_2 , and vice versa.

Remove these items from the unlabeled data.

until unlabeled data is exhausted ;

Algorithm 1: The Co-Training algorithm

The conditions are:

1. The unlabeled data distribution and the target concept f are *compatible* under the two views. In particular, let $p(\mathbf{x})$ be the marginal distribution of items. We require that with probability one, $\mathbf{x} \sim p(\mathbf{x})$ satisfies $f^{(1)}(\mathbf{x}^{(1)}) = f^{(2)}(\mathbf{x}^{(2)})$. That is, no item shall have conflicting labels between the two views.
2. Each base learner is able to learn the target concept under its view, given enough labeled data. This refers to standard supervised learning, where the amount of labeled data required may be much larger than in Co-Training.
3. The two views are conditionally independent given the class label: $p(\mathbf{x}^{(2)} | \mathbf{x}^{(1)}, y) = p(\mathbf{x}^{(2)} | y)$. If one knows the class y , then knowing the features in one view $\mathbf{x}^{(1)}$ does not help one guess the other view $\mathbf{x}^{(2)}$. This condition ensures that the most confident items from $f^{(1)}$ ’s perspective do not “pile up on top of each other” from $f^{(2)}$ ’s perspective. Rather, they spread out and provide representative (pseudo) training data for the second view.

In subsequent sections, we will see how consideration of these conditions shape our collaboration policy.

The reader might wonder why Co-Training keeps the two views separate. Why not stack the two views back into $\mathbf{x} = \begin{pmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{pmatrix}$, and train a supervised learner on \mathbf{x} ? One reason is their inductive biases leading to *different* classifiers for the same data, as shown in Figure 1. To see why this happens, consider how the base learners respond to the bottom and right clusters. For the bottom cluster, the x -axis view will be highly confident that the items belong to the red class because from this view they are nearly identical to the labeled red item. In contrast, the class of the right cluster will be uncertain from this view, since these items are not particularly similar to either labeled item. So, the x -view learner may choose to label some bottom cluster items and pass these to the y -view learner. For the y -view learner, the reverse pattern occurs: the right cluster items very likely belong to the blue class, whereas the class of the bottom cluster items is uncertain. Each view is confident about the items for which the opposing view is uncertain. Thus the two views, working together, converge on the solution shown in Figure 1(c). Such difference between supervised learning and Co-Training is general and can be observed with other datasets and choices of base learners. Another example is given in the last section.

Human Collaboration Policies

We now consider how these ideas from Co-Training can be used to shape a policy for human collaboration. The task we consider is category learning: Two human collaborators are given a number of labeled training items $(\mathbf{x}_1, y_1) \dots (\mathbf{x}_l, y_l)$ and together must label the unlabeled items $\mathbf{x}_{l+1} \dots \mathbf{x}_{l+u}$. One may view the labeled training items as *teaching experiences* given to the collaborators, e.g., by a teacher or a senior worker. It is reasonable to assume that in many cases the availability of teaching is limited. Therefore, the goal is for the dyad to grasp the target concept using as little teaching experience as possible. We assume that the collaborators can see all of the unlabeled items upfront, which is known as transduction in machine learning².

Our main interest is in exploring different *collaboration policies* between the two learners and how these policies affect the learning outcomes. One obvious policy is to allow the two collaborators full access to the data $(\mathbf{x}_1, y_1) \dots (\mathbf{x}_l, y_l), \mathbf{x}_{l+1} \dots \mathbf{x}_{l+u}$, and to allow them to fully interact with each other (in terms of discussions, gesturing, etc.). We call this the “full-collaboration” policy. Another policy might be to isolate the learners so that they each have full independent access to the data but cannot communicate or interact. We call this the “no-collaboration” policy.

We introduce a third policy, described in Algorithm 2 and explained below, that is inspired by and closely follows the Co-Training machine learning algorithm. This policy splits each item’s feature vector into two views: $\mathbf{x} = \begin{pmatrix} \mathbf{x}^{(1)} \\ \mathbf{x}^{(2)} \end{pmatrix}$. The intention is to allow each collaborator only one of the views. In contrast to machine learning, however, it is sometimes impossible to create artificial stimuli with a single view. For instance, the often used Gabor patches (Vandist, Schryver, and Rosseel 2009) vary in frequency and orientation, and it is impossible to depict an orientation without any information about frequency or vice versa. In this case, our policy constructs artificial stimuli that vary along the “viewed” dimension while holding a constant value on the “hidden” dimension (specifically the mean μ of the values on the missing view). So if Alice and Bob are the two collaborators, Alice might see the stimuli as $\mathbf{x}^{(1)}$ or $\begin{pmatrix} \mathbf{x}^{(1)} \\ \mu^{(2)} \end{pmatrix}$, while Bob sees them as $\mathbf{x}^{(2)}$ or $\begin{pmatrix} \mu^{(1)} \\ \mathbf{x}^{(2)} \end{pmatrix}$. Both Alice and Bob also see the labels for the labeled data.

Alice and Bob cannot directly communicate. Instead, at each iteration the policy requires both Alice and Bob to label the s unlabeled items that each is most confident about. After they have both finished, the policy shows Bob’s chosen items and labelings $(\mathbf{x}_{B1}, y_{B1}) \dots (\mathbf{x}_{Bs}, y_{Bs})$ to Alice. Note that, although Bob labeled these item from his view, Alice sees them from her own view. Alice understands that the labels come from Bob, but – in contrast to machine learning – it is up to her whether to believe Bob’s labelings (i.e., whether to use them as pseudo labeled data). At the same time, Alice’s labelings are shown to Bob. The policy then removes any unlabeled item that has been labeled by *either*

²However, the dyad is also capable of making *inductive* inferences when faced with new test items later on.

Input: labeled and unlabeled data, learning speed s .

Present the first-view data to Alice, second-view to Bob.

repeat

Let Alice label her s most confident unlabeled items; same for Bob.

Show Bob’s labelings $(\mathbf{x}_{B1}, y_{B1}) \dots (\mathbf{x}_{Bs}, y_{Bs})$ to Alice, and vice versa.

Remove $\{\mathbf{x}_{A1} \dots \mathbf{x}_{As}\} \cup \{\mathbf{x}_{B1} \dots \mathbf{x}_{Bs}\}$ from the unlabeled data.

until unlabeled data is exhausted ;

Algorithm 2: The Co-Training collaboration policy

Alice or Bob, and proceeds to the next iteration. This repeats until the unlabeled data is exhausted. In the end, each unlabeled item has received a label from Alice or Bob. In the rare cases when both Alice and Bob label the same item differently, the policy breaks the tie arbitrarily.

The only communication that is allowed in the Co-Training policy is label exchange³. In this sense, Co-Training falls between the no-collaboration and full-collaboration policies. Our main question is whether the Co-Training policy leads learners toward different classification outcomes than the no-collaboration and full-collaboration policies. We hypothesize that human behavior in the Co-Training policy will be well-predicted by the behavior of the Co-Training algorithm in machine learning, whereas participants will primarily learn linear category boundaries in the other two collaboration conditions. This is not a trivial hypothesis given the differences between human and machine learning discussed above, and the general difficulty human beings have in learning nonlinear decision boundaries without extensive supervision.

Experiments

We designed and conducted a series of experiments to compare human category learning behaviors under the three collaboration policies introduced in the previous section.

Participants and Materials

Across three separate experiments a total of 324 undergraduate students participated for course credit under IRB approval. We programmed networked software to run on a pair of computers so that two participants in separate rooms could collaborate according to the Co-Training policy, preventing any communication between them except that explicitly allowed by the software. The software also runs on a single computer for the full-collaboration and no-collaboration policies. The software was implemented in the ActionScript programming language and runs in Flash Player.

The category learning task was implemented as a card sorting game, see Figure 2. Each item \mathbf{x} is represented as a card. The user interface contains a central bin holding the unlabeled cards as well as a bin to the left and to the right

³In theory, Alice and Bob could agree on a *coding scheme* a priori and encode further information with their choices of items and labelings. We do not consider that possibility here.




Figure 2: The interface




Figure 3: Sample stimuli

into which labeled cards are placed. In the beginning, only the initially-labeled cards are shown in the left or right bins. The participants’ task is to sort all cards in the central bin into the left or right bins. Before starting the experiments, participants were told whether or not they would be working with a partner, and were instructed to begin with the card they were most confident about.

We assessed learning behavior in all collaboration conditions with two stimulus sets. Both included items defined over two continuous perceptual features, but differed in the psychological separability of the dimensions. The “separable” set contained Gabor patches varying in spatial frequency and orientation of the grating (Vandist, Schryver, and Rosseel 2009; Ashby and Maddox 1990). These dimensions are considered separable because it is possible for people to attend to one dimension to the exclusion of the other (Shepard 1964). The “integral” set contained colored squares of a fixed hue but varying in saturation and brightness. These dimensions are considered to be integral because it is difficult for people to attend to one dimension without also processing the other (Lockhead 1966). Extensive research has shown that people respond differently to stimuli defined on separable versus integral dimensions in supervised and unsupervised learning tasks (Ashby and Maddox 1990; Love 2002; Nosofsky and Palmeri 1996).

In both cases a stimulus is parametrized by $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in [0, 1]^2$. The range of values on each was determined in extensive pilot testing to ensure that participants could discriminate important distances along all dimensions. For Gabor patches, the frequencies were calculated using $\lambda = (x_1 * 5/34) + 2/17$, and the orientations were calculated using $\theta = x_2 * 100$, varying from 0 to 100 degrees clockwise from horizontal. For colored squares, the brightness was calculated using $b = x_1 * 0.5 + 0.25$ and the saturation was calculated using $c = x_2 * 0.9 + 0.5$. Figure 3 shows four stimuli corresponding to the cluster centers in Figure 1(a), in both the separable and integral stimulus spaces.

The Diamond Dataset and Co-Training Conditions

Most of our experiments employ the diamond dataset shown in Figure 1(a). It consists of $n = 80$ items evenly divided into 4 clusters. All clusters have radius 0.1. Items within a cluster lie on a regular grid. The two views are the x -axis and y -axis coordinates paired with the mean value of 0.5 on the hidden dimension as previously discussed.

We constructed this dataset with the aim of satisfying the three technical conditions for the Co-Training algorithm. Condition 1 is easy to verify: there exists at least one target concept f , shown in Figure 1(c), that is consistent with the




Figure 4: In Experiment 1 each participant worked with only one view of the dataset. There were four labeled items. Points dithered to show overlap.

marginal $p(\mathbf{x})$. In other words, no item receives contradictory labels across the two views (note this is not true for the concept in Figure 1(b)). From the Figure we can also verify that Condition 3 is approximately true:⁴ For both classes, knowing an item’s x -axis position tells us little about its y -axis position and vice versa.

Condition 2 cannot be verified by consideration of the stimulus set alone. It stipulates that each base learner in Co-Training must, with full supervision and sufficient labeled data, be capable of learning the target concept from only one view. Because the base learners in our study are human beings, we need to determine empirically whether this condition holds. Our first experiment addresses this question.

[Experiment 1] 13 participants were divided into two groups: 7 in the first-view group and 6 in the second-view group. Each worked alone as a base learner, and viewed stimuli from the “integral” stimulus set. Participants in the first-view condition saw items varying in the x dimension but fixed at 0.5 on the y dimension, whereas those in the second-view condition saw items varying along the y dimension and fixed at 0.5 in the x dimension, effectively collapsing the dataset into one dimension as shown in Figure 4. Participants viewed four labeled items corresponding to the four cluster centers in Figure 1(a), and were asked to classify the remaining 76 items. Note that this labeled data is twice what is provided in Co-Training. The purpose of the study is to verify that, when provided with this supervised experience, human learners are capable of learning the target concept as it is projected in one view.

Result: The average classification accuracy on the unlabeled items was quite high: 98.9% in the first-view group and 94.7% in the second-view group. These results suggest that people were able to learn the target concept f using only one view in a supervised learning setting given *four* labeled training items, thus verifying the final technical condition of Co-Training. Another pilot study also showed that in Experiment 1, humans cannot learn the concept in Figure 1(c) if they saw only the *two* labeled items in Figure 1(a) instead of the four. However, as we show next, they will be able to learn it from two labeled items if they perform Co-Training label exchange.

Results under Different Policies

[Experiment 2] Our second experiment compares human learning on the diamond dataset under the Co-Training, full-collaboration, and no-collaboration policies, now using just

⁴It would be exactly true if the clusters were squares, not circles.

Table 1: The fraction of patterns in cluster-level majority classification. “Other” includes the remaining $16 - 4 = 12$ possible patterns. Boldface indicates the largest fraction within a condition.

condition	pattern				other
	cross	horz	vert	diag	
CS	0.80	0.10	0.00	0.10	0.00
CI	0.68	0.04	0.04	0.20	0.04
FS	0.05	0.25	0.35	0.30	0.05
FI	0.00	0.08	0.00	0.92	0.00
NS	0.07	0.42	0.18	0.31	0.02
NI	0.00	0.00	0.00	1.00	0.00

two labeled items as in Figure 1(a). These three policies were implemented as follows:

Co-Training (C): Two partners sit in separate rooms working on a shared categorization task. Each partner sees one of the views and no communication is permitted except through the labeling of cards. Each partner labels one card ($s = 1$) and is then asked to wait for the other partner. The card labeled by the other partner is highlighted and automatically moves from the unlabeled bin to the appropriate labeled bin. If the partners have by chance labeled the same card, that card is automatically moved from the labeled bin, across the unlabeled bin, into the other labeled bin. This process of labeling followed by viewing is repeated until all cards are labeled.

Full-collaboration (F): Two partners sit side-by-side before a single computer working on the same categorization task. They are able to view both features on each card simultaneously. No restriction is made on their communication.

No-collaboration (N): A single participant categorizes all cards while viewing both features simultaneously.

Each collaboration policy was paired with the separable (S) or integral (I) stimuli, resulting in 6 conditions. Participants were assigned randomly to conditions as follows: 21 dyads for CS, 25 dyads for CI; 20 dyads for FS, 26 dyads for FI; 45 singles for NS, and 34 singles for NI.

To summarize the results of a given dyad or individual, we classified each cluster in the diamond dataset as either “red” or “blue” based on a simple majority vote (i.e. the cluster was designated red if more than 50% of the items in it were classified as red, and blue otherwise). Thus there were $2^4 = 16$ different possible patterns for the four clusters. Table 1 shows the proportion of participants whose behavior matched each of these patterns across the different conditions. For example, in the CS condition, $17/21 \approx 0.8$ fraction of dyads produced the “cross” pattern.

Several observations can be made from Table 1. First, the Co-Training policy robustly produces the nonlinear “cross” pattern in about three quarters of the dyads. This pattern was rarely observed in the full-collaboration and the no-collaboration policies (χ^2 test, $p \ll 0.01$), which both mainly produce linear decision boundaries. This is the main finding of our work: the Co-Training human collaboration policy leads to outcomes dramatically different from no-




Figure 5: Differences between humans and machines (aggregated over CS and CI). (a) The first unlabeled items (black dots) chosen by the first-view partners. (b) Same, but for the second-view. (c) Per-item average labels.

collaboration and full-collaboration policies, and consistent with that predicted by the machine learning algorithm.

Second, in the full-collaboration and no-collaboration policies, participants showed quite different behaviors for stimuli defined over separable versus integral dimensions, producing axis-parallel boundaries with separable dimensions and “integrated” oblique boundaries with integral dimensions. This pattern has been previously documented in a variety of work in cognition. In Co-Training, however, the separability of the stimulus dimensions does not affect behavior (CS vs. CI, χ^2 test, $p = 0.5$). This is not surprising given that each person sees only one view, but it suggests an interesting application of the policy: Co-Training can enforce consistent classification regardless of the separability of the stimulus dimensions.

Additionally, there was no significant difference between full- and no-collaboration (χ^2 test, $p = 0.7$). Thus the differences observed under the Co-Training policy were not simply the result of having two individuals working together.

Although the Co-Training human collaboration outcome fits machine learning model predictions at the cluster level, we observed some subtle differences suggesting that machine learning algorithms like 1NN may not be the ideal models for human base learners. One difference concerns the unlabeled items that humans label first. Machine base learners would label the items they are most confident about, which will likely be an item that *overlaps* with a labeled item under that view. Participants in our experiments did not always pick such overlapping items, but seemed to settle for items loosely similar to labeled ones, see Figure 5(a) and (b). Another difference is in how sure the humans are. For each unlabeled item, we may average its classification across all dyads in the Co-Training conditions where, if the average is close to -1 (blue) or 1 (red), all dyads label it consistently; 0 if they are quite unsure. Figure 5(c) shows this per-item average using a color coding. Items in the top and left clusters (with labeled items) are very certain, while those in the bottom and right clusters are relatively uncertain (though they do have the correct per-item majority vote label). Typical machine Co-Training learners will have higher certainty on these clusters.

A Counter-Example

Finally, we investigated human behavior under the Co-Training policy in a learning problem that *violates* Co-Training’s technical conditions. The new dataset was iden-




Figure 6: The counter-example

tical to the diamond dataset except that the unlabeled items were distributed on a grid, see Figure 6(a). The dataset therefore violates Condition 1: items near the four corners receive conflicting labels between the two views.

[Experiment 3] 24 dyads worked on this counter-example under the Co-Training policy with the separable stimuli. Apart from the distribution of the unlabeled items, all aspects of the study were identical to Experiment 2. Figure 6(b) shows the per-item average labels in Experiment 3. Classification decisions in this study were clearly less certain than those observed in Experiment 2 (see corresponding items in Figure 5(c)). To compare with the CS row in Table 1, we also computed the majority vote pattern for every dyad on each of the four rectangular “clusters” in Figure 6(a). The proportion of dyads showing each pattern were: cross 0.00, horz 0.21, vert 0.17, diag 0.33, other 0.29. No dyad produced the cross pattern on this dataset. Thus human Co-Training outcomes depends critically upon the distribution of the unlabeled items.

Discussions and Future Work

We showed that, when collaborating according to a novel policy inspired by Co-Training, two human learners behave differently than individual learners or learning pairs collaborating in an unconstrained manner. Specifically, they jointly acquire a nonlinear labeling on the diamond dataset that is highly consistent with the behavior of the machine learning algorithm, yet unusual for human category learning generally. We have also shown that the behaviors elicited by the policy depend upon the distribution of the unlabeled data.

The current work employs very simple stimuli constructed to highlight the differences between Co-Training and other learning models. The question thus arises, what relevance do these results have for real-world learning tasks? We believe there are several potentially important implications. First, under the Co-Training policy each participant need view only a subset of an item’s features. For problems where the number of relevant features are overwhelming, the policy may provide an efficient way of dividing the problem up so as to make best use of costly human effort. Second, in Co-Training each learner is satisfied with the final result (meaning there is little conflict between the labels given by one partner and the other), even though jointly the team arrives at a solution that would seem unlikely had they both viewed the full features. Co-Training thus provides a means of promoting agreement among team members for classification solutions that otherwise might cause disagreement. Third, the only communication required is label-exchange, which might be useful in situations where communication

is costly. Fourth, each learner is “blind” to some of the feature dimensions. The policy might therefore prove useful in sensitive classification tasks where data security is an issue.

Of course, all of these applications depend upon there being real-world tasks of interest that meet the technical conditions that allow Co-Training to work. In this vein, it is worth noting that Co-Training does apply to other datasets beyond the “diamond” set used here. For example, here is a 2D dataset with 8 clusters, two of them initially labeled:

$\begin{pmatrix} \bullet & \bullet & + \\ \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet \end{pmatrix}$. The outcome $\begin{pmatrix} o & + & o \\ o & + & o \\ o & + & o \end{pmatrix}$ is predicted by the Co-Training machine algorithm, and we have observed this behavior in preliminary human studies. To determine whether Co-Training has application for a real dataset, the task organizer must be able to assess whether the problem meets Co-Training’s technical conditions, and must also be able to find views of the data that exploit Co-Training’s properties. These constitute interesting problems for machine learning in their own right, and are a focus for future research.

Acknowledgments We thank Tom Mitchell for helpful discussions. Research supported by AFOSR FA9550-09-1-0313, NSF IIS-0953219, IIS-0916038.

References

- Ashby, F. G., and Maddox, W. T. 1990. Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance* 16(3):598–612.
- Balcan, M.-F., and Blum, A. 2010. A discriminative model for semi-supervised learning. *Journal of the ACM* 57(3).
- Blum, A., and Mitchell, T. 1998. Combining labeled and unlabeled data with co-training. In *COLT*.
- Brefeld, U.; Gaertner, T.; Scheffer, T.; and Wrobel, S. 2006. Efficient co-regularized least squares regression. In *ICML*.
- Johnson, R., and Zhang, T. 2007. Two-view feature generation model for semi-supervised learning. In *ICML*.
- Lockhead, G. R. 1966. Effects of dimensional redundancy on visual discrimination. *Journal of Experimental Psychology: Human Perception and Performance* 3:436–443.
- Love, B. C. 2002. Comparing supervised and unsupervised category learning. *Psychonomic Bulletin and Review* 9:829–835.
- Nigam, K., and Ghani, R. 2000. Analyzing the effectiveness and applicability of co-training. In *CIKM*.
- Nosofsky, R. M., and Palmeri, T. J. 1996. Learning to classify integral-dimension stimuli. *Psychonomic Bulletin and Review* 3(2):222–226.
- Nosofsky, R. M. 1986. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* 115(1):39–57.
- Shepard, R. N. 1964. Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology* 1:54–87.
- Vandist, K.; Schryver, M. D.; and Rosseel, Y. 2009. Semisupervised category learning: The impact of feedback in learning the information-integration task. *Attention, Perception, & Psychophysics* 71(2):328–341.
- Zhou, Z.-H., and Li, M. 2005. Tri-training: exploiting unlabeled data using three classifiers. *IEEE Transactions on Knowledge and Data Engineering* 17(11):1529–1541.