Co-Training as a Human Collaboration Policy

Xiaojin Zhu, **Bryan R. Gibson**, Timothy T. Rogers

University of Wisconsin-Madison

August 11, 2011

A Binary Classification Task

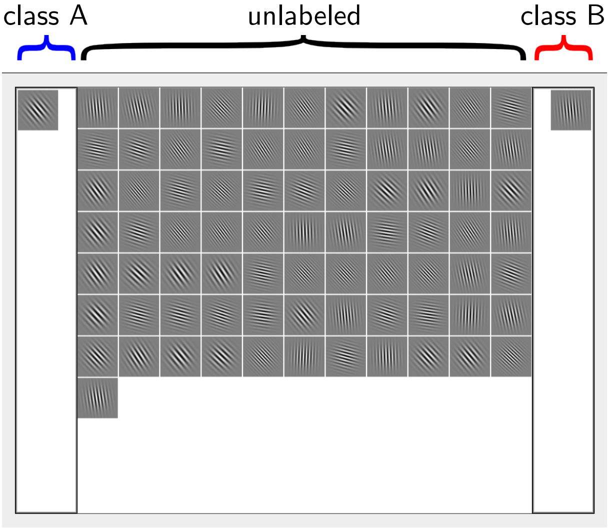

A Binary Classification Task

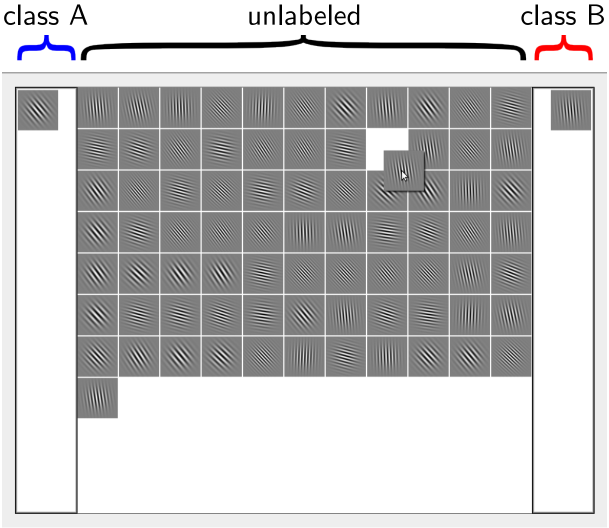

A Binary Classification Task

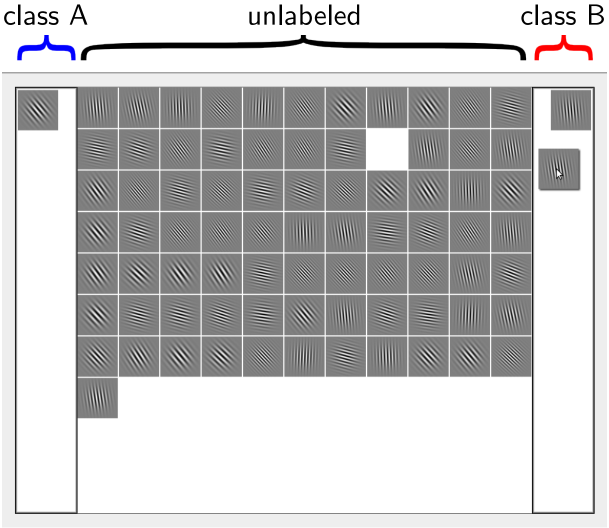

A Binary Classification Task

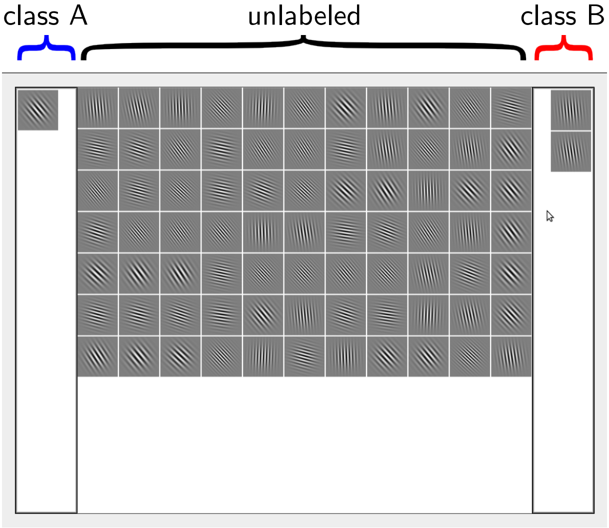

Three Different Collaboration Policies

no collaboration

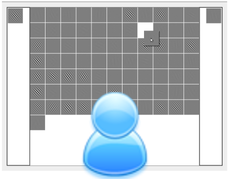

Three Different Collaboration Policies

no collaboration

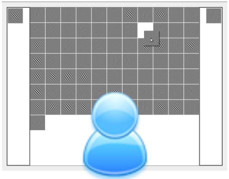

full collaboration

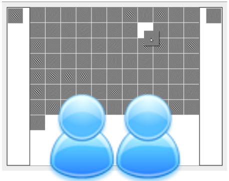

Three Different Collaboration Policies

no collaboration

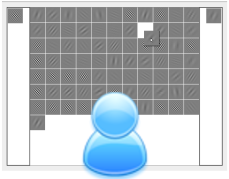

Co-Training

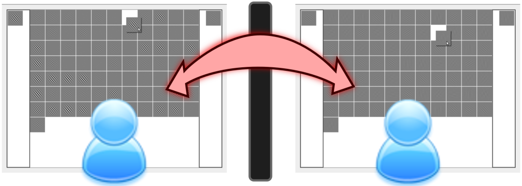

full collaboration

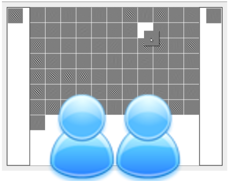

Stimuli vs. Feature Space

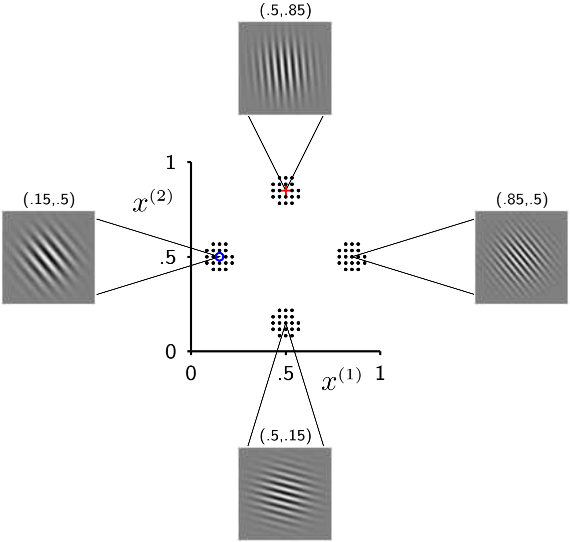

Possible Outcomes

Given:

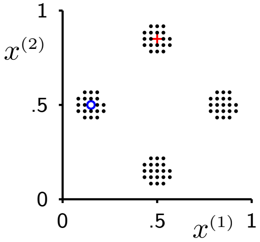

Possible Outcomes

Given:

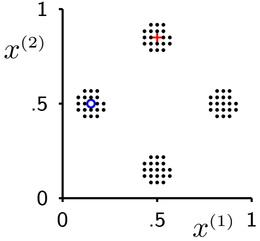

Majority Vote Per Cluster

Possible Outcomes

Given:

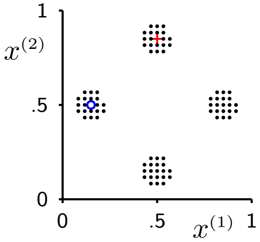

Majority Vote Per Cluster

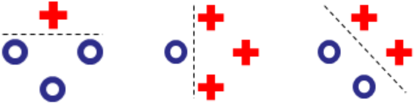

horz vert diag

Possible Outcomes

Given:

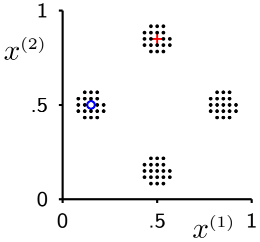

Majority Vote Per Cluster

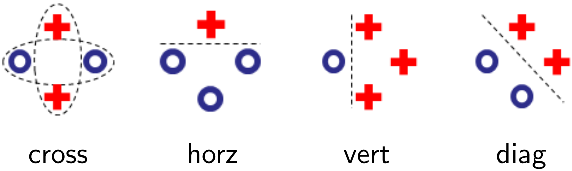

Possible Outcomes

Given:

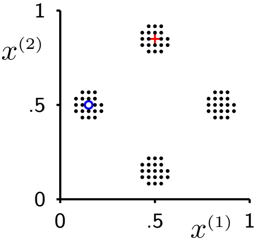

Majority Vote Per Cluster

12 more . . .

cross horz vert diag other

Review: Co-Training for Computers

Review: Co-Training for Computers [BM98]

- Given
  - ℓ labeled $(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_\ell, y_\ell)$
  - u unlabeled points $\mathbf{x}_{\ell+1}, \ldots, \mathbf{x}_{\ell+u}$

  Learn $f : \mathbf{x} \mapsto y$
- Assume feature vector $\mathbf{x}$ can be split into two "views":

$$\mathbf{x} = \begin{pmatrix} x^{(1)} \\ x^{(2)} \end{pmatrix}$$

- Train two base learners $f^{(1)} : x^{(1)} \mapsto y$ and $f^{(2)} : x^{(2)} \mapsto y$

The Co-Training Algorithm [BM98]

- First, learn from labeled data:

  $f^{(1)}$ on $(x_1^{(1)}, y_1) \ldots (x_\ell^{(1)}, y_\ell)$

  $f^{(2)}$ on $(x_1^{(2)}, y_1) \ldots (x_\ell^{(2)}, y_\ell)$

- Then, use unlabeled data in an iterative fashion:
  - $f^{(1)}$ classifies the unlabeled point that it is most confident in
  - $f^{(1)}$ adds this point to $f^{(2)}$'s labeled set
  - $f^{(2)}$ reciprocates
  - until the data is exhausted

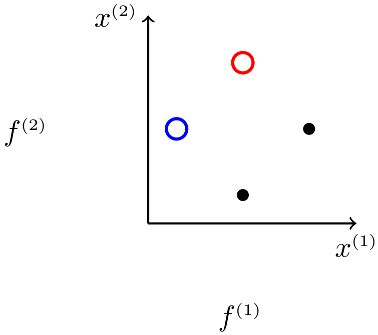

Sufficient Conditions for Co-Training [BM98]

compatibility: w.p. 1, $\mathbf{x} \sim p(\mathbf{x})$ satisfies $f^{(1)}(x^{(1)}) = f^{(2)}(x^{(2)})$

sufficiency: each base learner is able to learn the target concept under its view, *given enough labeled data*

conditional independence: $p(x^{(1)}, x^{(2)}|y) = p(x^{(1)}|y)p(x^{(2)}|y)$

Let Humans Enact Co-Training

Let Humans Enact Co-Training

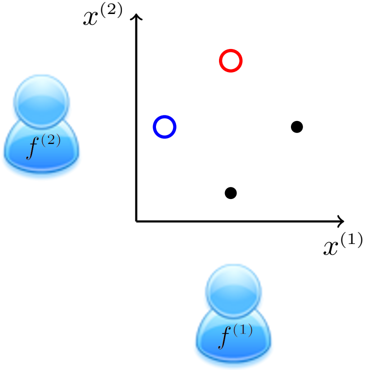

- Use humans as base learners
- Co-training as a policy for restricting communication

Collaboration Under Co-Training Policy

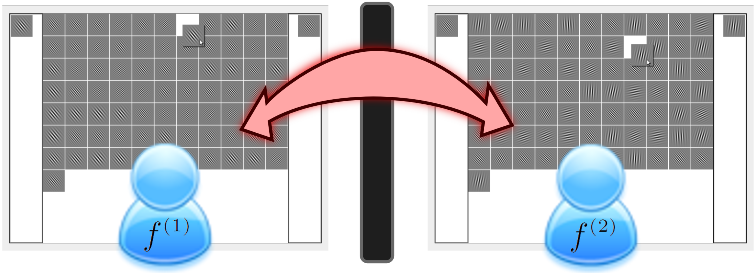

Collaboration Under Co-Training Policy

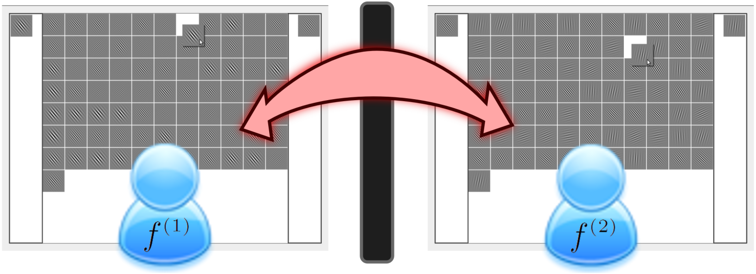

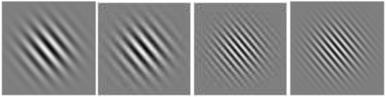

vary $x^{(1)}$, set $x^{(2)} = 0.5$

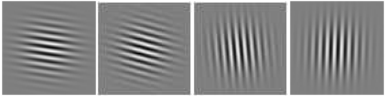

set $x^{(1)} = 0.5$, vary $x^{(2)}$

Collaboration Under Co-Training Policy

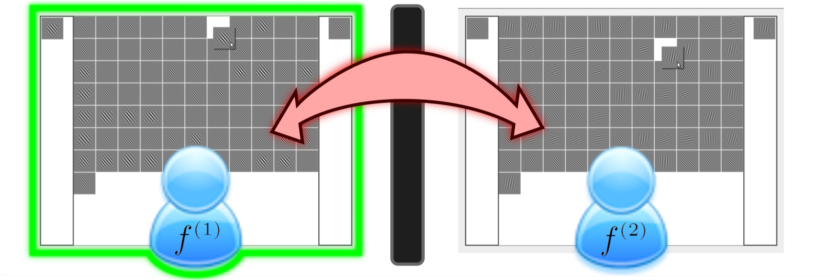

Collaboration Under Co-Training Policy

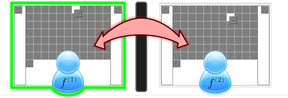

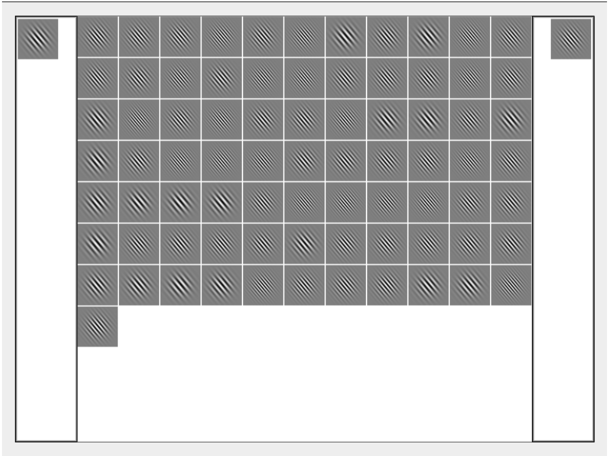

Collaboration Under Co-Training Policy

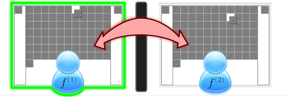

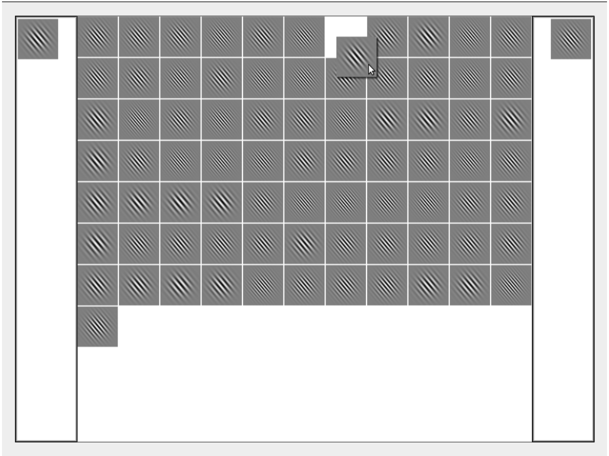

Collaboration Under Co-Training Policy

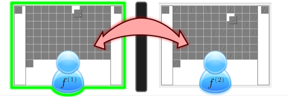

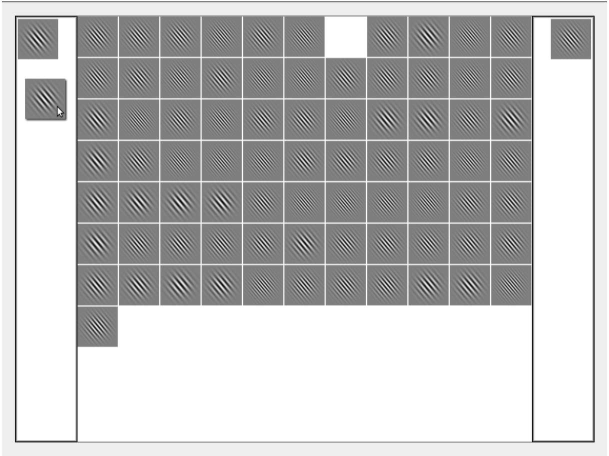

Collaboration Under Co-Training Policy

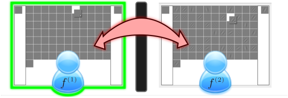

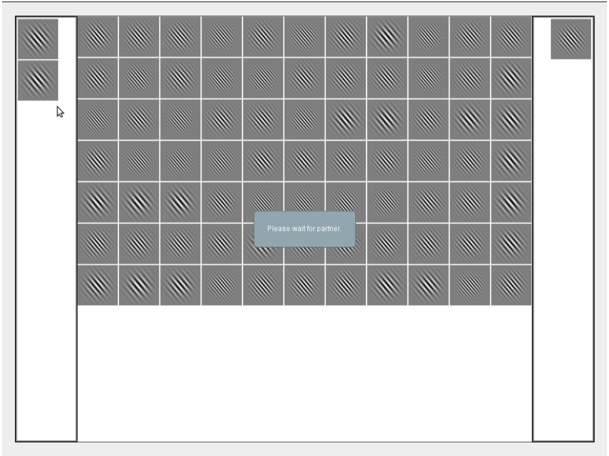

Collaboration Under Co-Training Policy

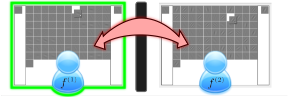

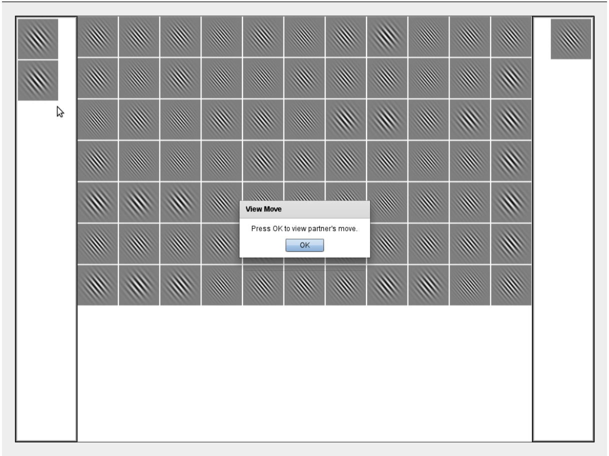

Collaboration Under Co-Training Policy

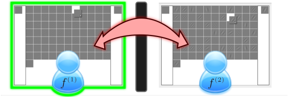

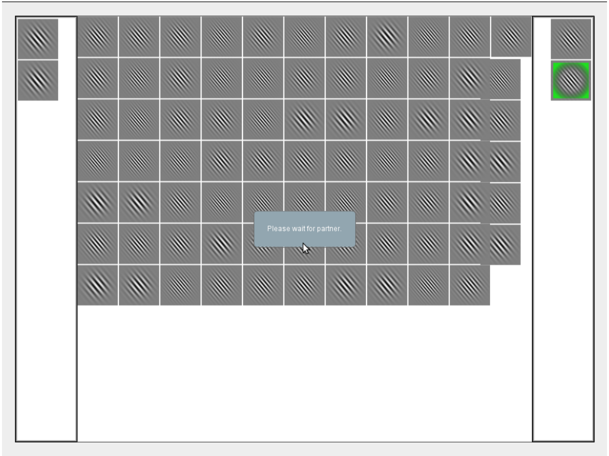

Collaboration Under Co-Training Policy

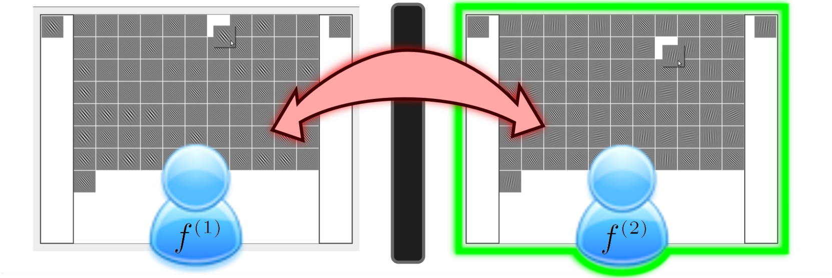

Collaboration Under Co-Training Policy

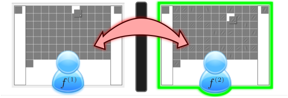

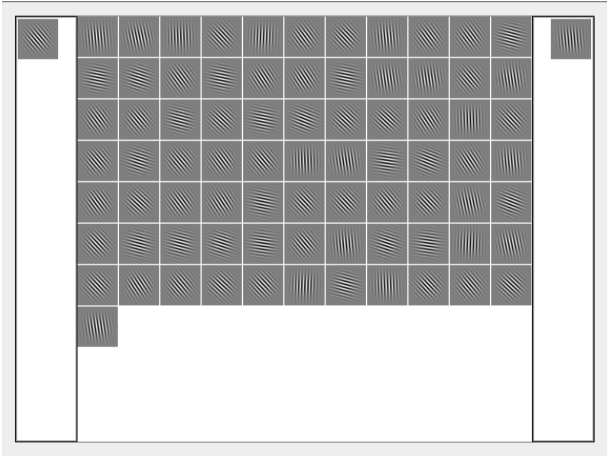

Collaboration Under Co-Training Policy

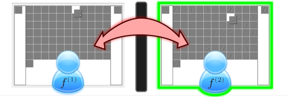

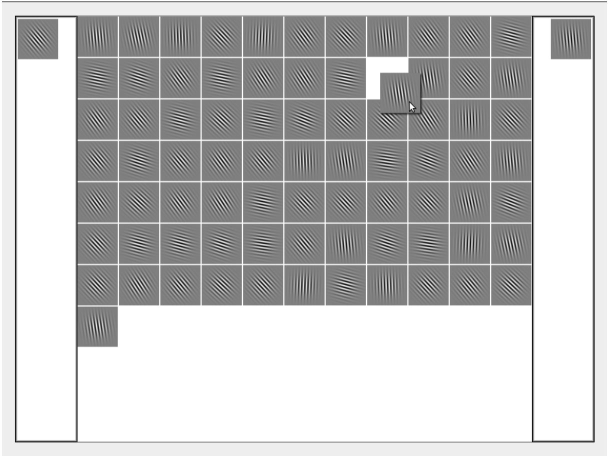

Collaboration Under Co-Training Policy

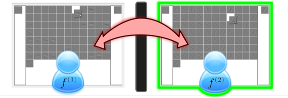

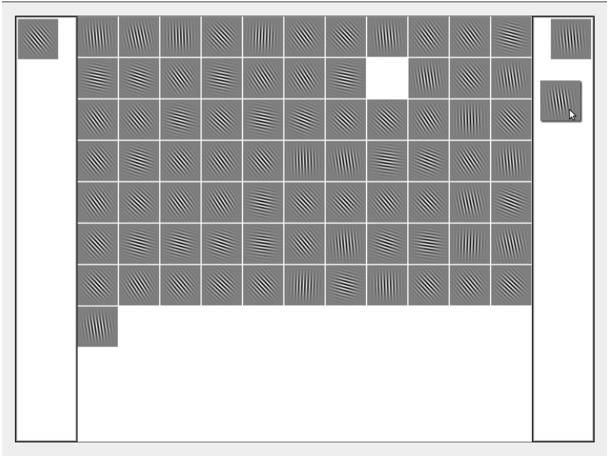

Collaboration Under Co-Training Policy

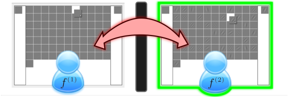

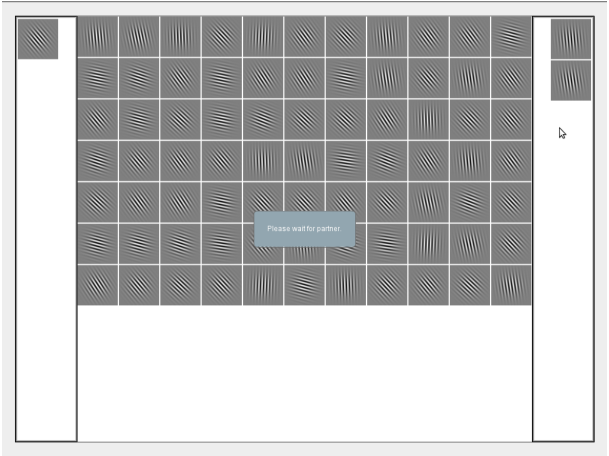

Collaboration Under Co-Training Policy

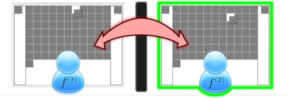

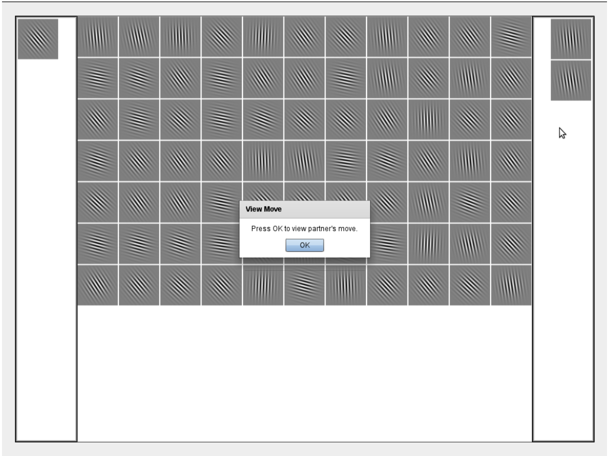

Collaboration Under Co-Training Policy

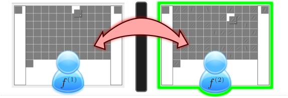

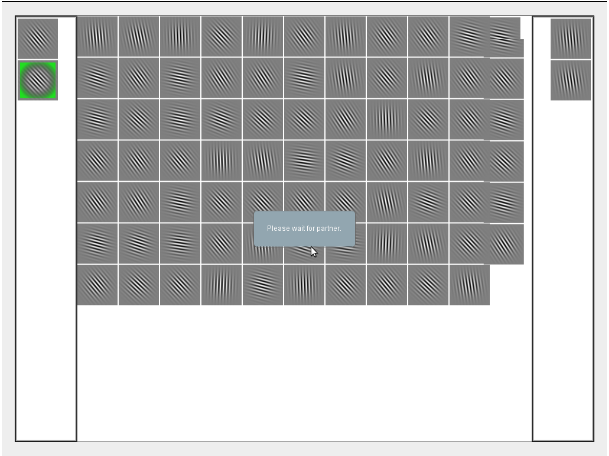

Collaboration Under Co-Training Policy

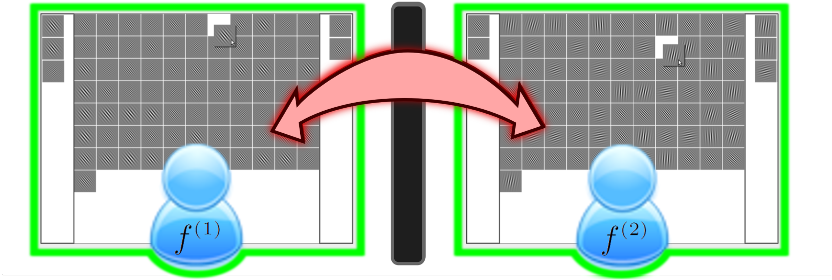

Review of Collaboration Policies

no collaboration

Review of Collaboration Policies

no collaboration

full collaboration

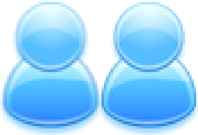

Review of Collaboration Policies

no collaboration

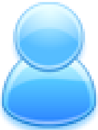

Co-Training

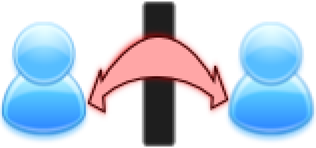

full collaboration

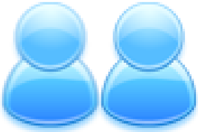

Results under Different Policies

	policy	(#subj.)	outcome				
			cross	horz	vert	diag	other
	no	(45x1)					
	Co-Training	(21x2)					
	full	(20x2)					

Results under Different Policies

policy		(#subj.)	cross	horz	vert	diag	other
	no	(45x1)	0.07	0.42	0.18	0.31	0.02
	Co-Training	(21x2)					
	full	(20x2)					

Results under Different Policies

policy	(#subj.)	outcome				
		cross	horz	vert	diag	other
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)					
full	(20x2)	0.05	0.25	0.35	0.30	0.05

Results under Different Policies

policy		(#subj.)	cross	horz	vert	diag	other
	no	(45x1)	0.07	0.42	0.18	0.31	0.02
	Co-Training	(21x2)					
	full	(20x2)	0.05	0.25	0.35	0.30	0.05

- Full and no-collaboration are not statistically significantly different

Results under Different Policies

policy	(#subj.)	cross	horz	vert	diag	other
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)	**0.80**	0.10	0.00	0.10	0.00
full	(20x2)	0.05	0.25	0.35	0.30	0.05

- Full and no-collaboration are not statistically significantly different

A Different Stimuli Set

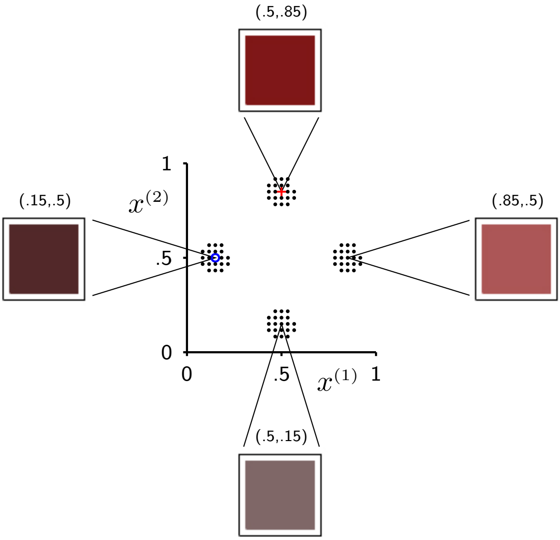

Results Using Color Stimuli

policy	(#subj.)	outcome				
		cross	horz	vert	diag	other
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)	**0.80**	0.10	0.00	0.10	0.00
full	(20x2)	0.05	0.25	0.35	0.30	0.05
no	(34x1)					
Co-Training	(25x2)					
full	(26x2)					

Results Using Color Stimuli

policy	(#subj.)	cross	horz	vert	diag	other
		outcome				
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)	**0.80**	0.10	0.00	0.10	0.00
full	(20x2)	0.05	0.25	0.35	0.30	0.05
no	(34x1)	0.00	0.00	0.00	**1.00**	0.00
Co-Training	(25x2)					
full	(26x2)					

Results Using Color Stimuli

policy	(#subj.)	cross	horz	vert	diag	other
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)	**0.80**	0.10	0.00	0.10	0.00
full	(20x2)	0.05	0.25	0.35	0.30	0.05
no	(34x1)	0.00	0.00	0.00	**1.00**	0.00
Co-Training	(25x2)					
full	(26x2)	0.00	0.08	0.00	**0.92**	0.00

Results Using Color Stimuli

policy	(#subj.)	cross	horz	vert	diag	other
no	(45x1)	0.07	0.42	0.18	0.31	0.02
Co-Training	(21x2)	**0.80**	0.10	0.00	0.10	0.00
full	(20x2)	0.05	0.25	0.35	0.30	0.05
no	(34x1)	0.00	0.00	0.00	**1.00**	0.00
Co-Training	(25x2)	**0.68**	0.04	0.04	0.20	0.04
full	(26x2)	0.00	0.08	0.00	**0.92**	0.00

Let's Break Co-Training

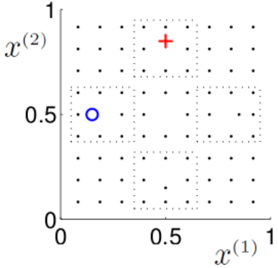

The only change:
u is now on a uniform grid over the entire space

This *violates* the Co-Training conditions (not *compatible*).

Co-Training Results (Human Behavior):
cross **0.00**, horz 0.21, vert 0.17, diag 0.33, other 0.29

Per-Item Average Labels

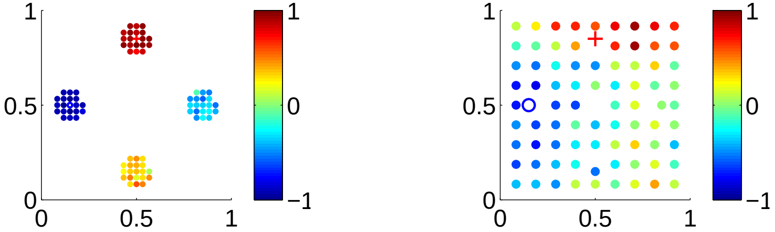

Conclusion and Future Work

- Co-Training collaboration can lead to non-trivial decision boundaries not found under no or full collaboration.
- When could this be useful?
  - data security is an issue
  - conflict is an issue
  - communication is costly
  - number of relevant features is overwhelming
- **Open:** Given $(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_n, y_n)$, is the problem Co-Trainable and if so, how do you optimally split $\mathbf{x}$?

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

Thank you!

References

F. G. Ashby and W. T. Maddox.
Integrating information from separable psychological dimensions.
Journal of Experimental Psychology: Human Perception and Performance, 16(3):598–612, 1990.

Avrim Blum and Tom Mitchell.
Combining labeled and unlabeled data with co-training.
In *COLT*, 1998.

B. C. Love.
Comparing supervised and unsupervised category learning.
Psychonomic Bulletin and Review, 9:829–835, 2002.

R. M. Nosofsky and T. J. Palmeri.
Learning to classify integral-dimension stimuli.
Psychonomic Bulletin and Review, 3(2):222–226, 1996.

The Co-Training Algorithm [BM98]

Input: labeled and unlabeled data, the 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 $f^{(1)}$ and $f^{(2)}$ 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 data is exhausted

Human Co-Training Algorithm

Input: labeled and unlabeled data, learning speed s.

Repeat

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

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

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

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

until unlabeled data is exhausted

Experiment 1 - Sufficiency Verification

Card sorting task:

- Integral stimuli
- 4 labeled items
- Subjects work alone
- $x^{(1)}$-view or $x^{(2)}$-view

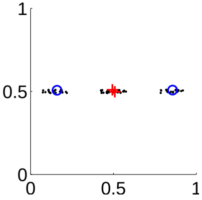

(a) $x^{(1)}$ view

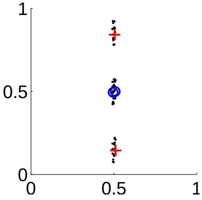

(b) $x^{(2)}$ view

Purpose: To verify that, when provided with enough labeled data, human learners are capable of learning the target concept projected in one view.

Results (classification accuracy):

$x^{(1)}$ view = 98.9%, $x^{(2)}$ view = 94.7% ✓

First Unlabeled Items Chosen

Exp. 2: `CS` and `CI`

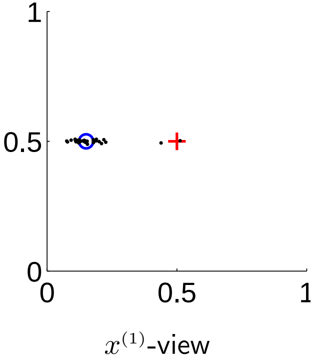

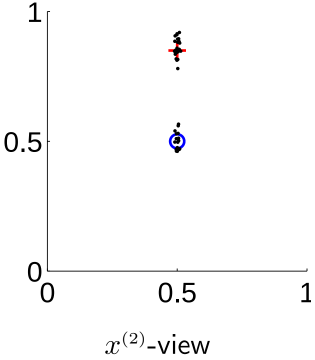

Human Learning Constraints

- difficulty learning nonlinear boundaries [Lov02]
- sensitive to stimulus representation [AM90, NP96]

	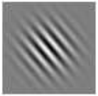	
separability	separable	non-separable
non-axis-parallel boundary	hard	easy

Question: Can we alter these constraints by leveraging insights from a machine learning algorithm?