

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) xx-xx-2012		2. REPORT TYPE Final		3. DATES COVERED (From - To) Apr 2009 - Dec 2011	
4. TITLE AND SUBTITLE A Cognitive Study of Learning with Labeled and Unlabeled Data				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-09-1-0313	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Xiaojin Zhu, Timothy T. Rogers				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Wisconsin-Madison				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2012-0535	
12. DISTRIBUTION / AVAILABILITY STATEMENT A					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In semi-supervised learning (SSL) the learner is presented with both labeled and unlabeled data. If the learner makes certain assumptions regarding the distribution of the unlabeled items $p(x)$ and the class conditional $p(y x)$, they can learn the concept faster and more accurately. We investigate how humans are affected by unlabeled data in a supervised categorization task. Our project lead to better understanding of human learning, improvements in human teaching strategy, improvements in human/machine cooperative learning and improvements in machine learning models. Our empirical evidence for human SSL includes several human behavioral studies that definitively show the influence of unlabeled data in human category learning. Our theoretical models produce plausible semi-supervised learning models for human learning and machine learning. We utilize these observations to enhance human learning.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

Final Report FA9550-09-1-0313

A Cognitive Study of Learning with Labeled and Unlabeled Data

PI: Xiaojin Zhu CoPI: Timothy T. Rogers

2009 – 2011

1 Human Semi-Supervised Learning

While both Supervised Learning (SL), in the form of classification, and Unsupervised Learning, in the form of clustering, have been well studied in Cognitive Science, it is only recently that the Machine Learning (ML) concept of Semi-Supervised Learning (SSL) has been applied to human learning.

In a SL setting, a learner is presented with a set of labeled items (x, y) and is asked to use these item/label pairs to learn the underlying classifier (a.k.a. concept) $f : X \mapsto y$. SSL differs in that, in addition to the labeled data, the learner is also presented with a (usually much larger) set of unlabeled data. If the learner makes certain assumptions regarding the distribution of the unlabeled items $p(x)$ and the class conditional $p(y | x)$, they may be able to learn the concept more accurately and potentially faster than with labeled items alone, given that the SSL assumptions made are appropriate.

Investigation of how humans are affected by *unlabeled* data in a *supervised* categorization task, and how the resulting behavior compares to the well-understood behavior of SSL ML models, can lead to further understanding of human learning, improvements in human teaching strategy, improvements in human/machine cooperative learning and, potentially, improvements in the ML models themselves.

2 Our Empirical Evidence for Human Semi-Supervised Learning

Our group is among the first to investigate the effect of unlabeled data on human category learning [13]. Our team of ML, Cognitive Science and Educational Psychology researchers showed that humans are affected by unlabeled data. Furthermore, the resulting behavior can be accurately modeled by ML (SSL) techniques. In this study human participants were first trained to learn on labeled items varying in one feature to learn a binary classification concept.

Each participant was then exposed to and asked to classify unlabeled data drawn from a bimodal Gaussian Mixture Model (GMM) distribution. The trough of this GMM was shifted away from the decision boundary indicated by the labeled data. One of the assumptions available in the SSL framework is the gap assumption: that a classification boundary will lie along a low-density region (gap) of the unlabeled distribution while a boundary which runs through a high density region is assumed unlikely. The unlabeled distribution is shifted so that the trough, or gap, in the distribution violates this assumption. The fact that the classification boundaries implied by participant behavior drifted towards this shifted trough showed that humans are affected by unlabeled data. Additionally, the behavior matches existing SSL model predictions. A second experiment by Kalish et al. resulted in similar findings [6].

To further understand human SSL, a third experiment was devised to explore how humans would behave when exposed to unlabeled drawn from a distribution designed to be explicitly contradictory [8]. The task was again binary classification, with participants asked to label unlabeled items interspersed with labeled items, where items varied in two dimensions. Participants were split into two conditions which varied in the underlying unlabeled distribution. In the “helpful” condition a gap in the unlabeled distribution existed overlapping and parallel to the labeled classification boundary. In the “harmful” condition the gap in the unlabeled distribution was orthogonal to the labeled boundary. Learners making use of the gap assumption should learn the concept faster with the helpful unlabeled distribution. It was found that, without time pressure, participants in both conditions performed equally well. However, when required to respond rapidly, participants performed substantially better in the helpful condition, indicating that they were affected by the underlying distribution of unlabeled data in a way that enhanced their performance.

While the use of gaps in the unlabeled distribution is a common method of achieving SSL, there are other properties of unlabeled data that can affect learning. A fourth experiment was designed to investigate the effect of presenting unlabeled items ordered in time [12]. Human participants were shown a sequence of labeled training items, varying in one dimension, and asked to learn a binary classification. They were then asked to label a separate set of unlabeled test items. Participants in each of two conditions were shown exactly the same set of labeled training and unlabeled test items, but the ordering of the unlabeled test items differed by condition: either in a sequence ordered from left to right in feature space or right to left. It was found that humans, shown the same labeled data, produced different labelings of the test items depending on the ordering. The classification boundary was found to shift in one direction or the other in feature space depending on the direction of sequence presentation order. Several SSL models were presented which produced behavior similar to that of the human participants.

Another SSL assumption which we investigated is the manifold assumption [2]. Under this assumption, items are assumed to lie along a lower dimensional manifold in a higher dimensional space. For instance, a set of items described in a two dimensional feature space may in fact all lie along a 1D line,

or manifold, in this 2D space. Items assumed to lie along a manifold can also be assumed to share any label information attached to items which fall on that manifold. A label for one point on the manifold can be allowed to propagate to any unlabeled items sharing that manifold. A fifth experiment was designed to test whether exposing humans to a mixture of labeled and unlabeled data following a 1D manifold in 2D space would lead to behavior similar to that of an ML making the manifold assumption. We found that participants were able to produce labelings similar to that of an SSL model using the manifold assumption, but that, for our chosen distribution and stimuli, two things were necessary: a number of labeled points which ruled out simple hypotheses, and hints that particular stimuli were similar to each other.

3 Our Theoretical Models for Semi-Supervised Learning

To account for human SSL behaviors, we developed several new ML models that are cognitively plausible [12]. Recall the empirical experiments showed that two people receiving exactly the same training experience will classify certain test items in opposite ways depending on the other items that appear in the test set. This test-item effect can be induced by either the order or the distribution of test items. We consider test-item effects as arising from online semi-supervised learning, and compared three novel computational models: (i) a non-parametric Bayesian model (Dirichlet Process Mixture model or DPMM) similar to Anderson’s Rational model of categorization but extended to online semi-supervised learning by marginalization; (ii) a non-parametric regression model (Nadaraya-Watson kernel estimator) similar to exemplar models of categorization but extended to online semi-supervised learning by a self-training procedure; and (iii) an online semi-supervised parametric mixture model (PMM) similar to prototype models of categorization. The empirical data are consistent with some parametrization of the DPMM and PMM approaches but are not well explained by the NKWE approach, suggesting that test-item effects can provide important empirical constraints on theories of human category learning.

Another SSL model we developed is for a learning setting of importance to large scale machine learning: potentially unlimited data arrives sequentially, but only a small fraction of it is labeled. The learner cannot store the data; it should learn from both labeled and unlabeled data, and it may also request labels for some of the unlabeled items. This setting is frequently encountered in real-world applications and has the characteristics of online, semi-supervised, and active learning. Yet previous learning models fail to consider these characteristics jointly. We present OASIS, a Bayesian model for this learning setting [4]. The main contributions of the model include the novel integration of a semi-supervised likelihood function, a sequential Monte Carlo scheme for efficient online Bayesian updating, and a posterior-reduction criterion for active learning.

Finally, we introduced sparsity into SSL. We pose transductive classification as a matrix completion problem. By assuming the underlying matrix has a low rank, our formulation is able to handle three problems simultaneously: i) multi-label learning, where each item has more than one label, ii) transduction, where most of these labels are unspecified, and iii) missing data, where a large number of features are missing. We obtained satisfactory results on several real-world tasks, suggesting that the low rank assumption may not be as restrictive as it seems. Our method allows for different loss functions to apply on the feature and label entries of the matrix. The resulting nuclear norm minimization problem is solved with a modified fixed-point continuation method that is guaranteed to find the global optimum [5].

4 Our Enhancement of Human Learning Based on Machine Learning Principles

We developed “human algorithms” informed by SSL which can affect the learning of cooperative groups of learners. One such algorithm is the Human Co-Training procedure [11]. Under Co-Training, two learners collaborate to label a set of unlabeled data according to a concept learned from a set of labeled data. This method is unique in that neither learner has a full view of the data. Instead, the features are split into two views such that each collaborator sees all of the data, but only represented by the features within their split or view. For example, if the data exists in two dimensions, then both learners would perceive the data as varying in only one dimension, that dimension being different for both collaborators. If the data and the classification concept follow certain constraints, the unlabeled data can be labeled correctly by the Co-Training pair using a smaller number of labeled examples than either learner could on their own. In an experiment where participant pairs collaborated on classification tasks under varying communication constraints, we were able to show that the Co-Training policy leads collaborators to jointly produce unique and potentially valuable classification outcomes. These outcomes are not generated under other collaboration policies and that these behaviors are expected by existing machine learning models.

We also investigated the reverse problem of human teaching in the presence of labeled and unlabeled data [7]. We study the empirical strategies that humans follow as they teach a target concept with a simple 1D threshold to a robot. Previous studies of computational teaching, particularly the teaching dimension model and the curriculum learning principle, offer contradictory predictions on what optimal strategy the teacher should follow in this teaching task. We show through behavioral studies that humans employ three distinct teaching strategies, one of which is consistent with the curriculum learning principle, and propose a novel theoretical framework as a potential explanation for this strategy. This framework, which assumes a teaching goal of minimizing the learners expected generalization error at each iteration, extends the standard

teaching dimension model and offers a theoretical justification for curriculum learning.

5 Other Related Work

We list some other work not directly on human SSL, but is in service of (or has the potential to) the project.

An important problem in cognitive psychology is to quantify the perceived similarities between stimuli. This is of great importance to the study of human SSL. Previous work attempted to address this problem with multi-dimensional scaling (MDS) and its variants. However, there are several shortcomings of the MDS approaches. We propose Yada, a novel general metric learning procedure based on two-alternative forced-choice behavioral experiments [9]. Our method learns forward and backward nonlinear mappings between an objective space in which the stimuli are defined by the standard feature vector representation, and a subjective space in which the distance between a pair of stimuli corresponds to their perceived similarity. Yada outperforms several standard embedding and metric learning algorithms, both in terms of likelihood and recovery error.

How does one know if a human learner has truly learned a concept, or is he simply overfitting? We offer a measure that combines computational learning theory and cognitive psychology to gauge human generalization abilities [14]. We propose to use Rademacher complexity, originally developed in computational learning theory, as a measure of human learning capacity. Rademacher complexity measures a learner's ability to fit random labels, and can be used to bound the learner's true error based on the observed training sample error.

Other work includes human multi-arm bandit tasks [3], sensorimotor child-parent interaction for word learning [10], and human expert knowledge in latent topic models [1].

References

- [1] David Andrzejewski, Xiaojin Zhu, Mark Craven, and Ben Recht. A framework for incorporating general domain knowledge into Latent Dirichlet Allocation using First-Order Logic. In *The Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, 2011.
- [2] Bryan Gibson, Xiaojin Zhu, Tim Rogers, Chuck Kalish, and Joseph Harrison. Humans learn using manifolds, reluctantly. In *Advances in Neural Information Processing Systems (NIPS) 24*. 2010.
- [3] Bryan R Gibson, Kwang-Sung Jun, and Xiaojin Zhu. With a little help from the computer: Hybrid human-machine systems on bandit problems. In *NIPS 2010 Workshop on Computational Social Science and the Wisdom of Crowds*, 2010.

- [4] Andrew Goldberg, Xiaojin Zhu, Alex Furger, and Jun-Ming Xu. OASIS: Online active semi-supervised learning. In *The Twenty-Fifth Conference on Artificial Intelligence (AAAI-11)*, 2011.
- [5] Andrew Goldberg, Xiaojin Zhu, Benjamin Recht, Jun-Ming Xu, and Robert Nowak. Transduction with matrix completion: Three birds with one stone. In *Advances in Neural Information Processing Systems (NIPS) 24*. 2010.
- [6] Charles W. Kalish, Timothy T. Rogers, Jonathan Lang, and Xiaojin Zhu. Can semi-supervised learning explain incorrect beliefs about categories? *Cognition*, 2011.
- [7] Faisal Khan, Xiaojin Zhu, and Bilge Mutlu. How do humans teach: On curriculum learning and teaching dimension. In *Advances in Neural Information Processing Systems (NIPS) 25*. 2011.
- [8] Timothy Rogers, Charles Kalish, Bryan Gibson, Joseph Harrison, and Xiaojin Zhu. Semi-supervised learning is observed in a speeded but not an unspeeded 2D categorization task. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society (CogSci)*, 2010.
- [9] Jun-Ming Xu, Xiaojin Zhu, and Timothy T. Rogers. Metric learning for estimating psychological similarities. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*, 2011.
- [10] Chen Yu, Jun-Ming Xu, and Xiaojin Zhu. Word learning through sensorimotor child-parent interaction: A feature selection approach. In *The 33rd Annual Conference of the Cognitive Science Society (CogSci 2011)*, 2011. (oral, acceptance rate 32%).
- [11] Xiaojin Zhu, Bryan Gibson, and Timothy Rogers. Co-training as a human collaboration policy. In *The Twenty-Fifth Conference on Artificial Intelligence (AAAI-11)*, 2011.
- [12] Xiaojin Zhu, Bryan R. Gibson, Kwang-Sung Jun, Timothy T. Rogers, Joseph Harrison, and Chuck Kalish. Cognitive models of test-item effects in human category learning. In *The 27th International Conference on Machine Learning (ICML)*, 2010.
- [13] Xiaojin Zhu, Timothy Rogers, Ruichen Qian, and Chuck Kalish. Humans perform semi-supervised classification too. In *Twenty-Second AAAI Conference on Artificial Intelligence (AAAI-07)*, 2007.
- [14] Xiaojin Zhu, Timothy T. Rogers, and Bryan Gibson. Human Rademacher complexity. In *Advances in Neural Information Processing Systems (NIPS) 23*. 2009.