Active Task Selection for Lifelong Machine Learning

Paul Ruvolo[1,2] Eric Eaton[1,3]

[1] Bryn Mawr College, [2] Olin College, [3] University of Pennsylvania

Abstract

In a lifelong learning framework, an agent acquires knowledge incrementally over consecutive learning tasks, continually building upon its experience. Recent lifelong learning algorithms have nearly identical accuracy to batch multi-task learning methods while learning tasks sequentially in over 1,000x less time. In this work, we further improve the scalability of lifelong learning by developing curriculum selection methods that enable an agent to actively select the next task to learn in order to maximize performance on future learning tasks. We demonstrate that active task selection is highly reliable and effective, allowing an agent to learn high performance models using up to 50% fewer tasks than when the agent has no control over the task order. We also explore a variant of transfer learning in the lifelong learning setting in which the agent can focus knowledge acquisition toward a particular target task.

Introduction

Goal: Develop intelligent agents that

1. Quickly learn new tasks
2. Learn continually with experience
3. Exhibit versatility over multiple tasks
4. Direct their own learning

	Transfer Learning	Batch Multi-Task Learning
Optimizes performance over	Target task	All tasks
Learns tasks consecutively	Yes, efficiently	Very inefficiently
Computational cost	Low	High

Lifelong learning includes elements of both transfer and multi-task learning

Contributions:

1. Active task selection methods that enable a lifelong learner to choose the next task to learn in order to maximize performance on future tasks
2. Targeted task selection that enables the lifelong learning agent to focus knowledge acquisition toward particular target tasks

Lifelong Learning Framework

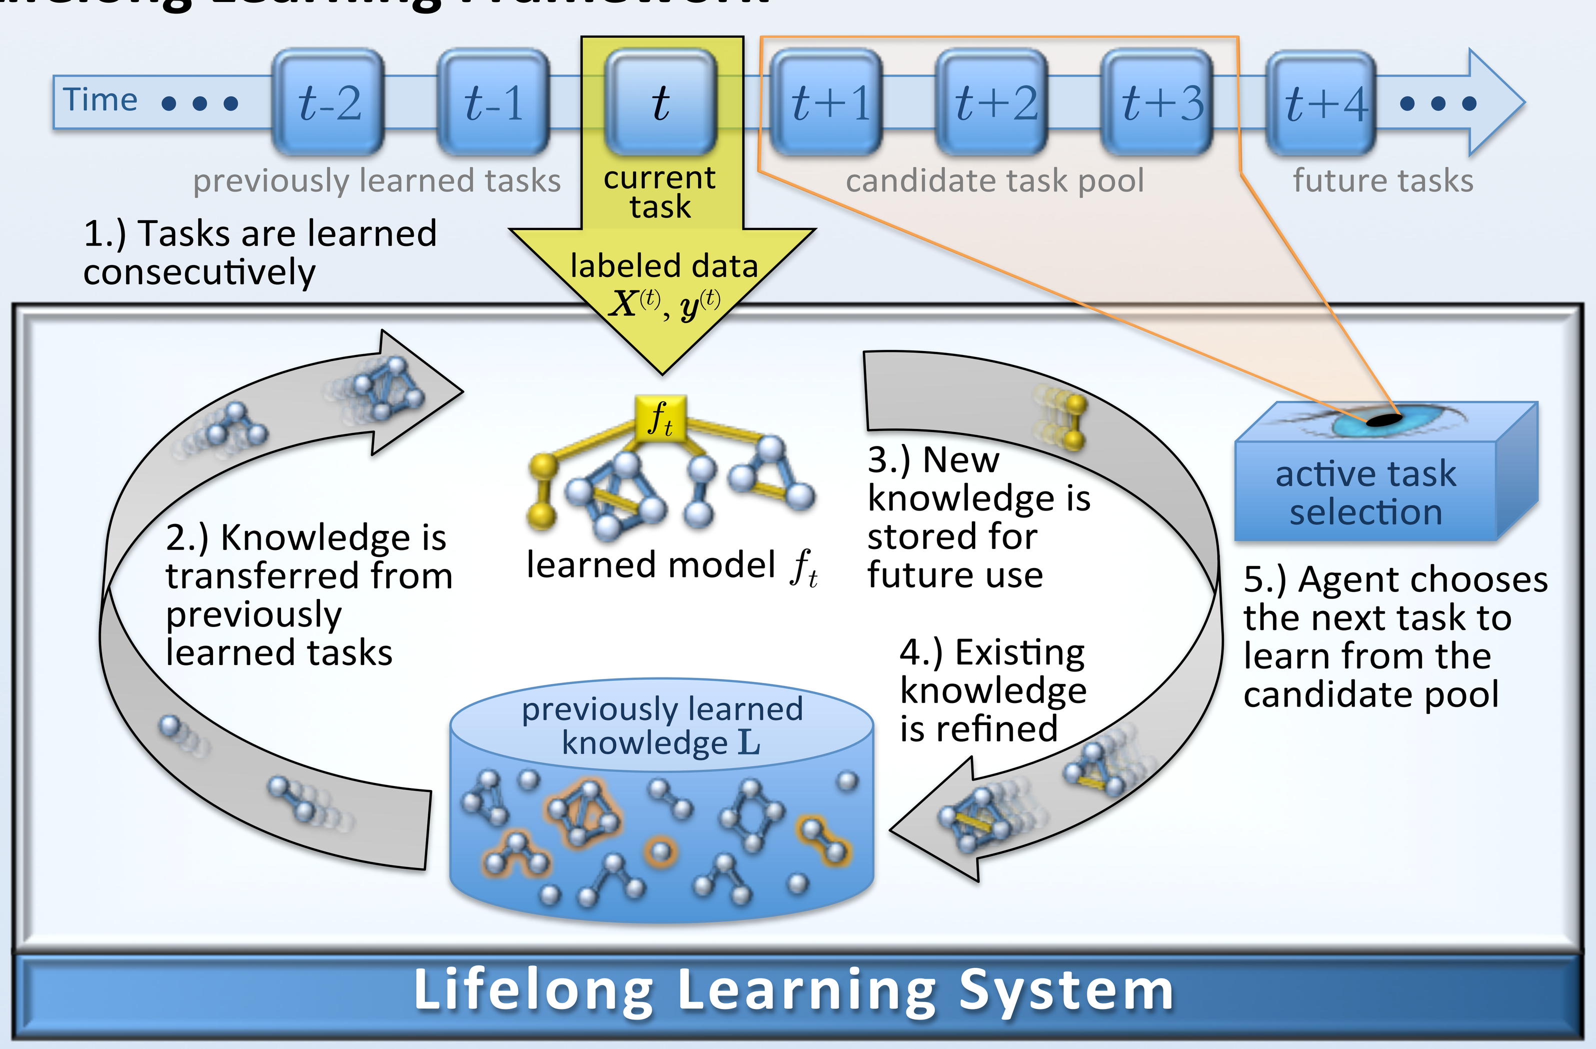

We learn a parametric model for each task t

$$f^{(t)}(\mathbf{x}) = f(\mathbf{x}; \boldsymbol{\theta}^{(t)}) \qquad \boldsymbol{\theta}^{(t)} \in \mathbb{R}^d$$

The parameter vectors for each model are linear combinations of a shared latent basis $\mathbf{L}$

$$\boldsymbol{\theta}^{(t)} = \mathbf{L}\boldsymbol{s}^{(t)} \qquad \mathbf{L} \in \mathbb{R}^{d \times k},\ \boldsymbol{s}^{(t)} \in \mathbb{R}^k$$

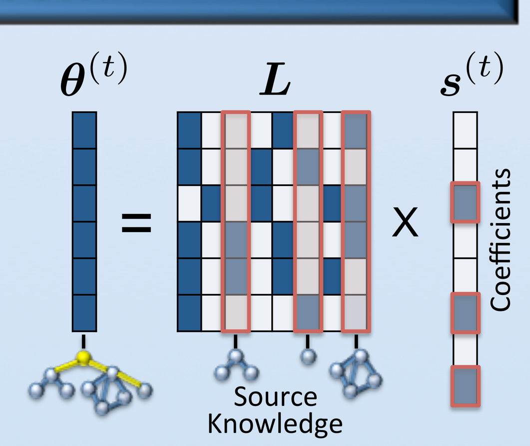

Overview of the Efficient Lifelong Learning Algorithm

Our active task selection is built on top of ELLA [Ruvolo & Eaton, ICML'13], an efficient online multi-task learner with the following properties:

1. Optimized performance over all tasks
2. Efficient learning of each new consecutive task via transfer
3. Computational complexity independent of: (1) the number of tasks learned, and (2) the amount of training data for all previous tasks
4. Close connections to online dictionary learning for sparse coding
5. Equivalent accuracy to batch MTL with over 1,000x speedup

ELLA minimizes an objective that encourages transfer between models:

$$e_T(\mathbf{L}) = \underset{\#\text{tasks seen so far}}{\frac{1}{T}}\sum_{t=1}^{T} \min_{\boldsymbol{s}^{(t)}} \left\{ \underbrace{\frac{1}{n_t}\sum_{i=1}^{n_t} \mathcal{L}\left(f\left(\boldsymbol{x}_i^{(t)}; \mathbf{L}\boldsymbol{s}^{(t)}\right), y_i^{(t)}\right)}_{\text{model fit to data}} + \underbrace{\mu \|\boldsymbol{s}^{(t)}\|_1}_{\text{sparsity}} \right\} + \underbrace{\lambda \|\mathbf{L}\|_{\mathsf{F}}^2}_{\text{complexity}}$$

To ensure scalability, ELLA makes the following simplifications:

1. Replace the inner sum with the 2nd-order Taylor expansion around the optimal task-specific model: $\boldsymbol{\theta}^{(t)} = \arg\min_{\boldsymbol{\theta}} \frac{1}{n_t}\sum_{i=1}^{n_t} \mathcal{L}\left(f\left(\boldsymbol{x}_i^{(t)}; \boldsymbol{\theta}\right), y_i^{(t)}\right)$
2. Eliminate the outer sum by optimizing $\boldsymbol{s}^{(t)}$ only when training on task t

These simplifications yield the following updates to learn given $(\mathbf{X}^{(t)}, \boldsymbol{y}^{(t)})$:

$$\boldsymbol{s}^{(t)} \leftarrow \arg\min_{\boldsymbol{s}^{(t)}} \ell(\mathbf{L}_m, \boldsymbol{s}^{(t)}, \boldsymbol{\theta}^{(t)}, \boldsymbol{D}^{(t)})$$

$$\mathbf{L}_{m+1} \leftarrow \arg\min_{\mathbf{L}} \lambda \|\mathbf{L}\|_{\mathsf{F}}^2 + \frac{1}{T}\sum_{t=1}^{T} \ell\left(\mathbf{L}, \boldsymbol{s}^{(t)}, \boldsymbol{\theta}^{(t)}, \boldsymbol{D}^{(t)}\right)$$

where

$$\ell(\mathbf{L}, \mathbf{s}, \boldsymbol{\theta}, \mathbf{D}) = \mu \|\mathbf{s}\|_1 + \|\boldsymbol{\theta} - \mathbf{L}\mathbf{s}\|_{\mathbf{D}}^2$$

$\boldsymbol{D}^{(t)}$ is ½ the Hessian of the single-task loss evaluated at $\boldsymbol{\theta}^{(t)}$

Active Task Selection

Goal: Choose the next task to learn from the candidate pool to best learn $\mathbf{L}$

- The agent can access a small set of labeled data for each candidate task

Information Maximization Approach

Selects the candidate task that maximizes the information gain on $\mathbf{L}$

$$t_{\text{next}} = \arg\min_{t} \iint p(\boldsymbol{\theta}^{(t)} = \mathbf{u}, \mathbf{D}^{(t)} = \mathbf{V} | \mathcal{I}_m) \times H\left[\mathbf{L} | \boldsymbol{\theta}^{(t)} = \mathbf{u}, \mathbf{D}^{(t)} = \mathbf{V}, \mathcal{I}_m\right] d\mathbf{u}\, d\mathbf{V}$$

To approximate this efficiently, we (1) use the optimal single task model $(\hat{\boldsymbol{\theta}}^{(t)}, \hat{\mathbf{D}}^{(t)})$, and (2) use a Laplace approximation of $\mathbf{L}$'s density as a multivariate Gaussian for the differential entropy term $H[]$, yielding:

$$t_{\text{next}} = \arg\min_{t \in \{T+1, \ldots, T_{\text{pool}}\}} \ln \left| \operatorname{Cov}\left[\operatorname{vec}(\mathbf{L}) | \boldsymbol{\theta}^{(t)} = \hat{\boldsymbol{\theta}}^{(t)}, \mathbf{D}^{(t)} = \hat{\mathbf{D}}^{(t)}, \mathcal{I}_m\right] \right|$$

Diversity Approach

Selects the candidate task that the current $\mathbf{L}$ is doing the worst job solving:

$$t_{\text{next}} = \arg\max_{t \in \{T+1, \ldots, T_{\text{pool}}\}} \min_{\mathbf{s}} \ell\left(\mathbf{L}_m, \mathbf{s}, \hat{\boldsymbol{\theta}}^{(t)}, \hat{\mathbf{D}}^{(t)}\right)$$

We also explore a probabilistic version, Diversity++, that chooses a candidate task proportionally to its inverse performance

Targeted Knowledge Acquisition with InfoMax

Idea: Instead of acquiring a general-purpose basis $\mathbf{L}$, focus on the knowledge needed for a specific target task, $t^{(\text{target})} = (\mathbf{X}^{(\text{target})}, \boldsymbol{y}^{(\text{target})})$

The targeted InfoMax objective is:

$$t_{\text{next}} = \arg\min_{t} \iint p(\boldsymbol{\theta}^{(t)} = \mathbf{u}, \mathbf{D}^{(t)} = \mathbf{V} | \mathcal{I}_m) \times H\left[\mathbf{L}\boldsymbol{s}^{(\text{target})} | \boldsymbol{\theta}^{(t)} = \mathbf{u}, \mathbf{D}^{(t)} = \mathbf{V}, \mathcal{I}_m\right] d\mathbf{u}\, d\mathbf{V}$$

which can be approximated efficiently as:

$$t_{\text{next}} = \arg\min_{t \in \{T+1, \ldots, T_{\text{pool}}\}} \ln \left| \boldsymbol{\Psi}^{\top} \boldsymbol{\Sigma}^{(t)} \boldsymbol{\Psi} \right| \qquad \boldsymbol{\Psi} = \mathbf{s}^{(\text{target})} \otimes \mathbf{I}_d$$

Applications

Facial Expression Recognition: identify presence of facial action units (#5 upper lid raiser, #10 upper lip raiser, #12 lip corner pull)

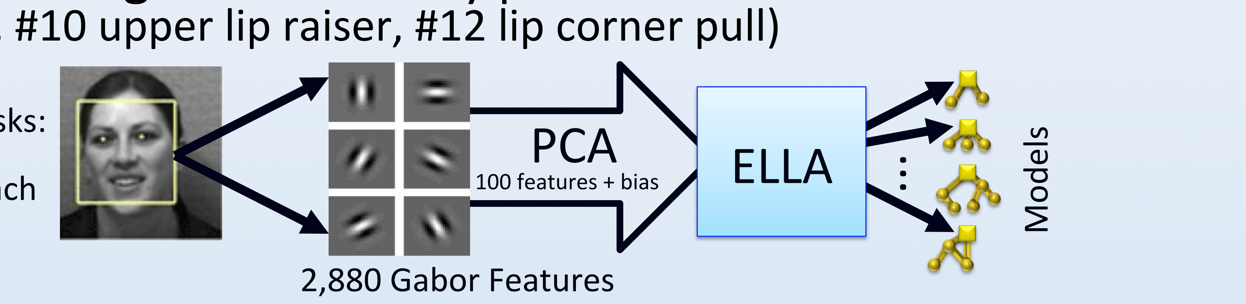

Land Mine Detection from radar

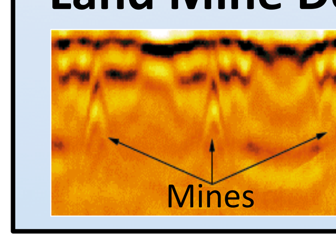

29 Classification Tasks:
- 29 regions
- 2 terrain types
- 14,820 instances total

Student Exam Score Prediction

139 Regression Tasks:
- 139 schools
- 15,362 students total
- 4 school-specific features
- 3 student-specific features

Results

Active task selection requires less tasks than random selection

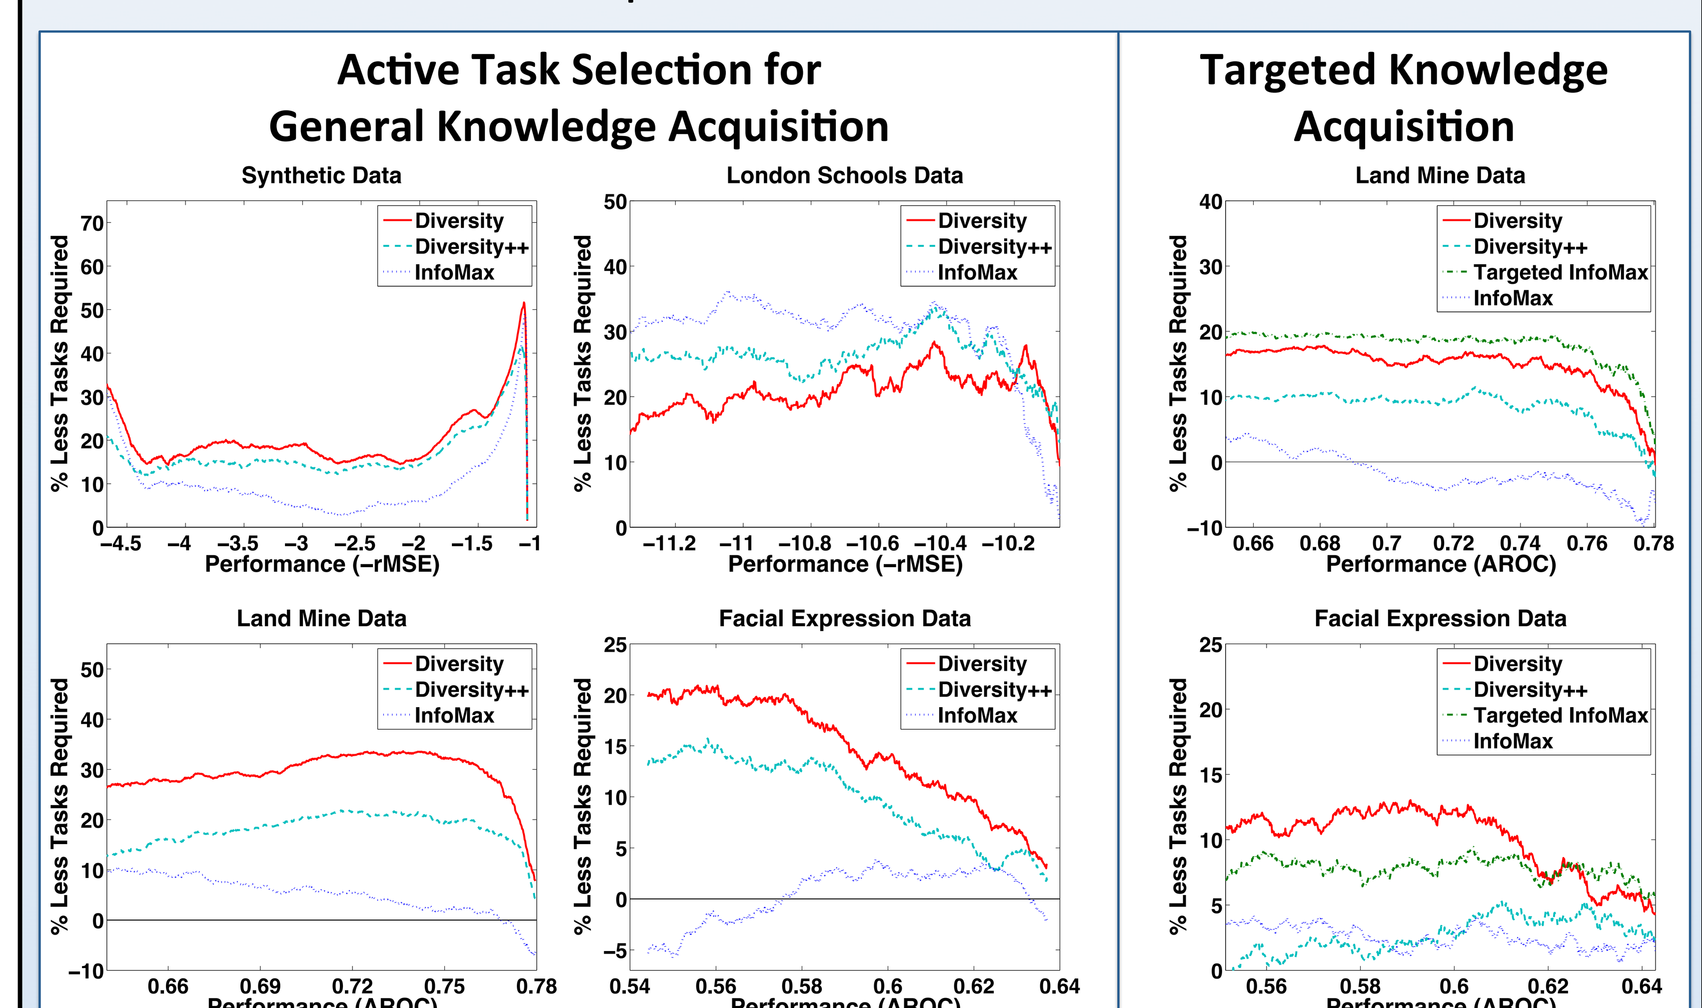

Plots shows the accuracy achieved versus the relative efficiency (in #tasks) as compared to random task selection

Average Task Reduction (%) for General Knowledge Acquisition

Data Set	InfoMax	Diversity	Diversity++
Land Mine	5.1±3.7	**29.4±4.1**	18.1±3.0
Facial Expr.	0.5±2.6	**14.6±5.1**	9.9±4.0
Syn. Data	10.2±7.9	**20.2±6.7**	17.0±5.9
London Sch.	**29.8±6.8**	21.0±3.1	26.2±3.1

Average Task Reduction (%) for Targeted Knowledge Acquisition

Data Set	Targeted InfoMax	InfoMax	Diversity	Diversity++
Land Mine	**17.9±2.7**	-1.7±3.0	14.9±3.2	8.5±2.5
Facial Expr.	7.8±0.7	2.6±0.8	**10.0±2.5**	2.7±1.3
Syn. Data	**38.4±7.5**	11.4±5.6	19.9±4.9	16.6±5.0
London Sch.	**26.9±1.8**	20.1±2.8	22.3±1.1	16.4±2.7

Acknowledgement: This research was supported by ONR grant #N00014-11-1-0139

Active task selection enables a lifelong learner to choose the next task to learn in order to maximize performance on future tasks