NIPS*95 Post-Conference Workshop

"Learning to Learn: Knowledge Consolidation and Transfer in Inductive Systems"

[Motivation and Goals], [Submissions], [Organizers], [Schedule], [Talk Abstracts], [For More Information]

Length: 2 days

Organizers: Jonathan Baxter, Rich Caruana, Tom Mitchell, Lorien Y. Pratt, Daniel L. Silver, Sebastian Thrun

Invited Talks:

- Leo Breiman (Berkeley)
- Tom Mitchell (CMU)
- Tomaso Poggio (MIT)
- Noel Sharkey (Sheffield)
- Jude Shavlik (Wisconsin)

Motivation:

The power of tabula rasa learning is limited. As these limits become apparent, interest has increased in developing methods that capitalize on previously acquired domain knowledge. Examples of these methods include:

- using symbolic domain theories to bias connectionist networks
- using unsupervised learning on a large corpus of unlabelled data to learn features useful for subsequent supervised learning on a smaller labelled corpus
- using models previously learned for other problems as a bias when learning new, but related, problems
- using extra outputs on a connectionist network to bias the hidden layer representation towards more predictive features
- updating belief(s) from a set of priors with Bayes rule

The methods used go by many names: hints, knowledge-based artificial neural nets

(KBANN), explanation-based neural nets (EBNN), multitask learning (MTL), lifelong learning, knowledge consolidation, etc. What they all have in common is the attempt to transfer knowledge from other sources to benefit the current inductive task. Potential benefits include better generalization, faster learning, and a bias towards representations that are more robust or more broadly applicable.

Goals:

To provide an opportunity for researchers and practitioners to get together in one place to discuss progress in knowledge consolidation and transfer. More specifically:

- To identify the current directions and boundaries of research
- To share/debate current theories and empirical results
- To determine the unifying principles underlying different approaches
- To identify open problems, future directions, joint research opportunities, and key contact persons
- To try to start answering questions like:
  - what do we mean by "related" tasks and how can we identify them?
  - how do we predict when transfer will help (or hurt)?
  - what are the benefits: speed, generalization, intelligibility,...?
  - what should be transferred: internal representations, parameter settings, features,...?
  - how should it be transferred: weight initialization, biasing the error metric,...?
  - how do we look inside to see what has been transferred?

Workshop Schedule:

From the 21 submissions we received, 11 were selected for short presentations. The current workshop schedule is as follows (full titles and asbtracts available below):

```
Friday, Dec 1
=============

AM
7:30-7:35  Welcome
7:35-8:05  Tom Mitchell (invited talk) - "Situated Learning"
8:05-8:25  Lorien Pratt    - "Neural Transfer For Hazardous Waste"
8:25-8:45  Nathan Intrator - "Learning Internal Reps From Multiple Tasks"
8:45-9:05  Rich Caruana    - "Where is Multitask Learning Useful?"
9:05-9:30  Panel Debate and Discussion -
           Topics include:  serial vs. parallel transfer,
                            what should be transfered?
                            what domains are ripe for transfer?
                            what are the goals of transfer?
           (Baxter, Caruana, Intrator, Mitchell, Silver, Pratt, ...)

9:30-4:30  Extracurricular Recreation

PM
4:30-5:00  Jude Shavlik (invited talk) - "Talking to Your Neural Net"
5:00-5:20  Leo Breiman  (invited talk) - "Curds & Whey"
5:20-5:40  Jonathan Baxter - "Bayesian Model of Learning to Learn"
```

```
5:40-6:00  Sebastian Thrun - "Identifying Relevant Tasks"
6:00-6:30  Panel Debate and Discussion -
           Topics include:  transfer human to machine vs. machine to machine,
                            is practice meeting theory?
                            is theory meeting practice?
           (Baxter, Caruana, Breiman, Thrun, Mitchell, Shavlik, ...)

Saturday, Dec 2
===============

AM
7:30-8:00  Noel Sharkey (invited talk) - "Adaptive Generalisation"
8:00-8:20  Anthony Robbins  - "Rehearsal and Catastrophic Interference"
8:20-8:40  J. Schmidhuber   - "A Theoretical Model of Learning to Learn"
8:40-9:00  Bairaktaris/Levy - "Dual-weight ANNs: Short/Long Term Learning"
9:00-9:30  Panel Debate and Discussion -
           Topics include:  catastrophic interference,
                            is there evidence for transfer in cognition?
                            what can nature/cogsci tell us about transfer?
           (Bairaktaris, de Sa, Levy, Robbins, Sharkey, Silver, ...)

9:30-4:30  More Extracurricular Recreation

PM
4:30-5:00  Tomaso Poggio (invited talk) - "Virtual Examples"
5:00-5:20  Virginia de Sa - "On Segregating Input Dimensions"
5:20-5:40  Chris Thornton - "Learning to be Brave: A Constructive Approach"
5:40-6:00  Mark Ring      - "Continual Learning"
6:00-6:25  Panel Debate and Discussion -
           Topics include:  combining supervised and unsupervised learning,
                            where do we go from here?
                            *this space intentionally left flexible*
           (de Sa, Mitchell, Poggio, Ring, Thornton, ...)
6:25-6:30  Farewell
```

20 minute talks are 12 minutes presentation and 8 minutes questions and discussion. 30 minute invited talks are 20 minutes presentation and 10 minutes questions and discussion. There are four 30-minute panels, one for each session. Although topics are listed for each panel, these are intended merely as points of departure. Everyone attending the workshop should feel free to raise any issues during the panels that seem appropriate. We encourage speakers and members of the audience to prepare a terse list (preferably using inflammatory language) of your favorite transfer issues and questions.

There are 16 talks, but this is not a conference! If speakers don't abuse their question/discussion time too much, more than 50% of the workshop will be spent on questions and discussion. To promote this, talks will use few slides and will focus on one or two key issues. It's a workshop. Come preapred to speak up, be controversial, and have fun.

Talks: Full Titles and Abstracts (click here)

Submissions:

We solicited presentations from anyone working in (or near) the following areas:

- Fundamental problems/issues in learning to learn
- Sequential/incremental, compositional (learning by parts), and parallel learning
- Task representation learning, retention, and use
- Task knowledge transfer (symbolic-neural, neural-neural)
- Life-long learning
- Adaptation of learning algorithms based on prior learning
- Learning (domain-specific) inductive bias
- Combining predictions made for different tasks from one domain
- Combining supervised learning (where the goal is to learn one feature from the other fea unsupervised learning (where the goal is to learn every feature from all the other f
- Combining symbolic and connectionist methods via transfer

- Theoretical models of learning to learn

If you would like to make a presentation at the workshop, please send a short (1 page or less) description of what you want to talk about to one of the following contacts by October 15. Email is preferred.

Rich Caruana
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
email: caruana@cs.cmu.edu
ph: (412) 268-3043
fax: (412) 268-5576

Daniel L. Silver
Department of Computer Science
Middlesex College
University of Western Ontario
London, Ontario, Canada N6A 3K7
email: dsilver@csd.uwo.ca
ph: (519) 473-6168
fax: (519) 661-3515

Preference will be given to proposals likely to generate debate and that go beyond summarizing prior work by raising important general issues or suggesting directions for future work. Proposals for moderator or panel-led discussions (e.g., sequential vs. parallel transfer) are also encouraged. We plan to run the workshop as a workshop, not as a mini conference, so be daring.

We look forward to your submission.

For more information:

- contact Rich Caruana, Daniel L. Silver, or one of the other organizers
- Bibliography on Transfer and Knowledge Consolidation
- Lori Pratt's **Transfer Web Page**
- **NIPS Home Page**
- **Other NIPS*95 Workshops**

See you in Colorado!