TRANSFER LEARNING OF SEARCH HEURISTICS FOR CONSTRAINT SATISFACTION

(Introducing Constrained Heuristic Search to the Soar Cognitive Architecture)

(30th Soar Workshop, University of Michigan)

Sean A. Bittle
Mark S. Fox

May 27th, 2010

university of toronto
mechanical & industrial engineering

Agenda

- Introduction: The Problem
- Implementation
- Experiments and Results
- Summary and Future Work

Introduction

"What is the problem?"

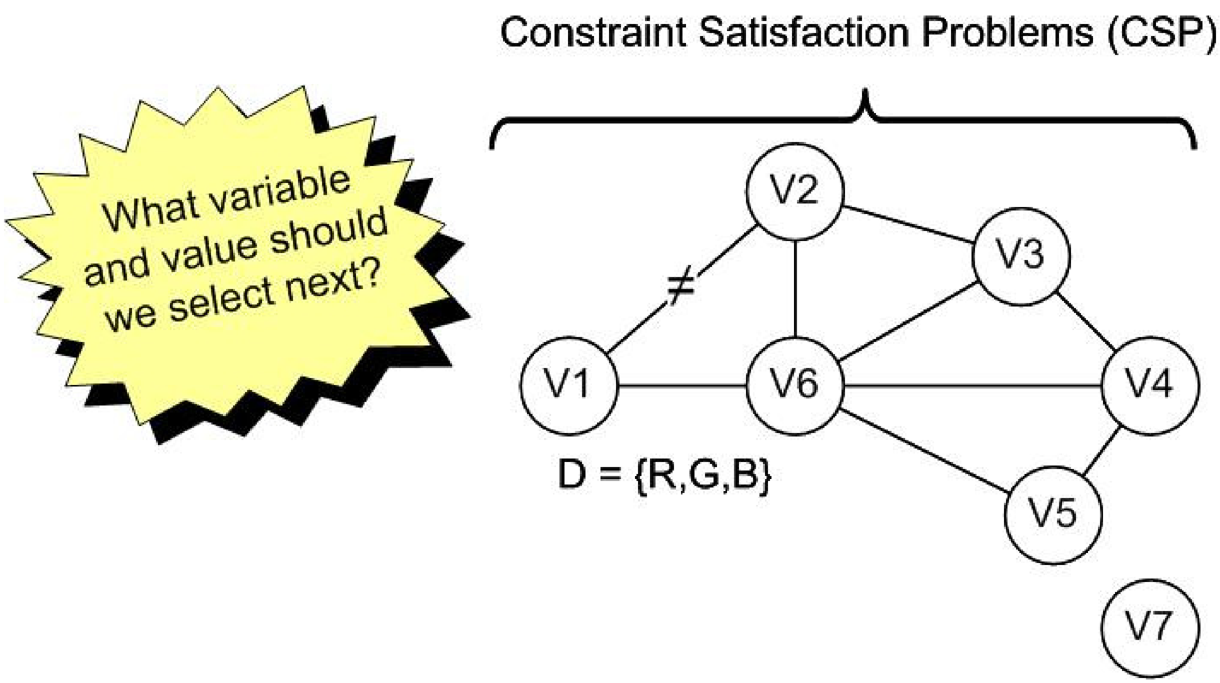

Introduction

"What is the problem?"

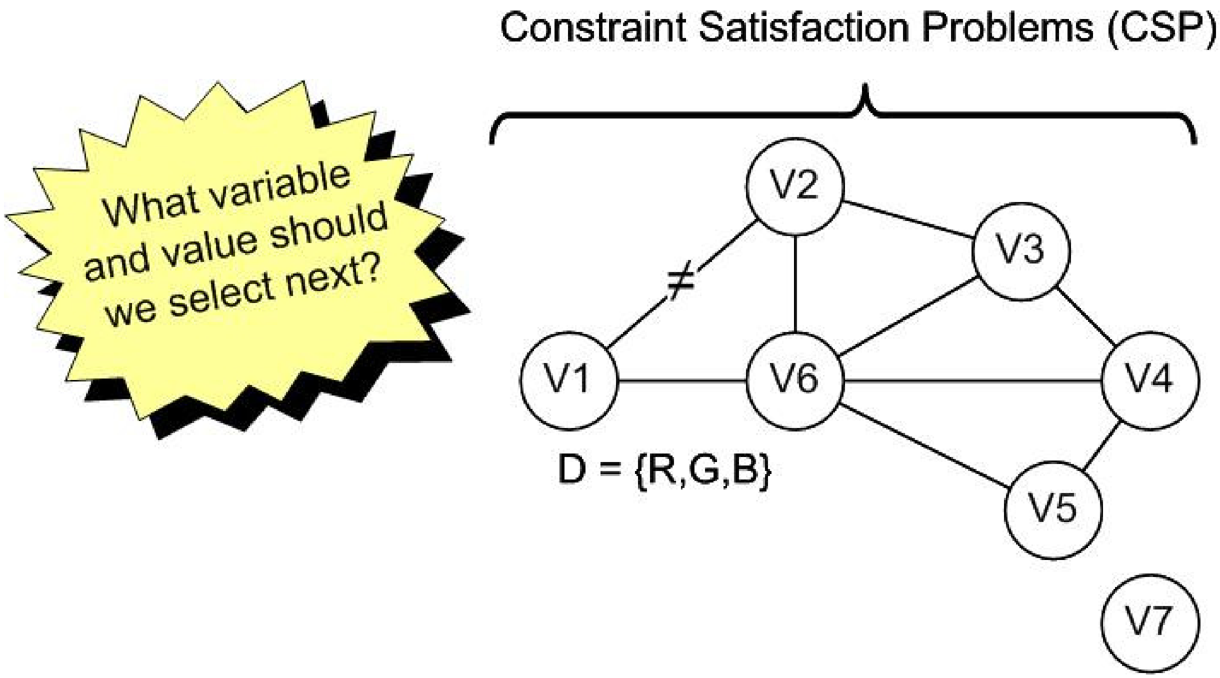

A central challenge in efficiently solving CSP's is variable and value ordering.

Current approaches use heuristics to guide ordering

Tradeoff between generality and effectiveness

Can we learn variable and value ordering heuristics that are effective and general?

Introduction

- Our goal is to demonstrate that computers can apply knowledge learned from one set task(s) to achieve superior performance on a different set of task(s).

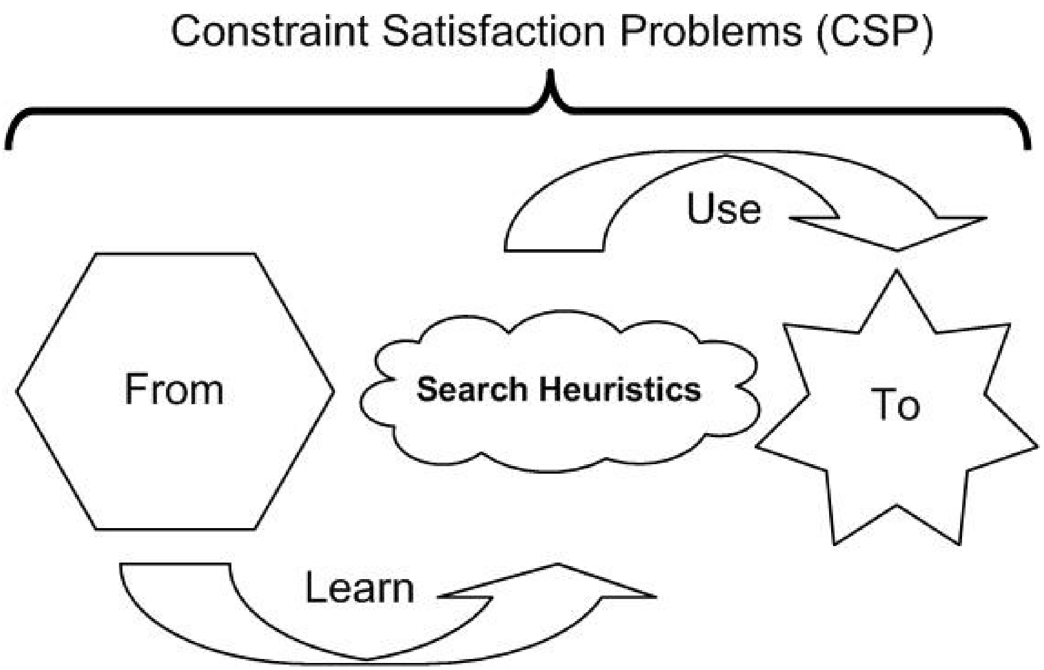

- Our thesis is that we can abstract and learn variable and value ordering search heuristics from a set of task(s) that can be used to secure improved problem solving performance on a different set of task(s).

Introduction

"What is transfer learning?"

Transfer learning (TL) is the ability to extract knowledge learned for an original (source) set of tasks in one domain to solve problems on a different (target) set of tasks in order to improve performance or enhance future learning [DARPA, 2005]

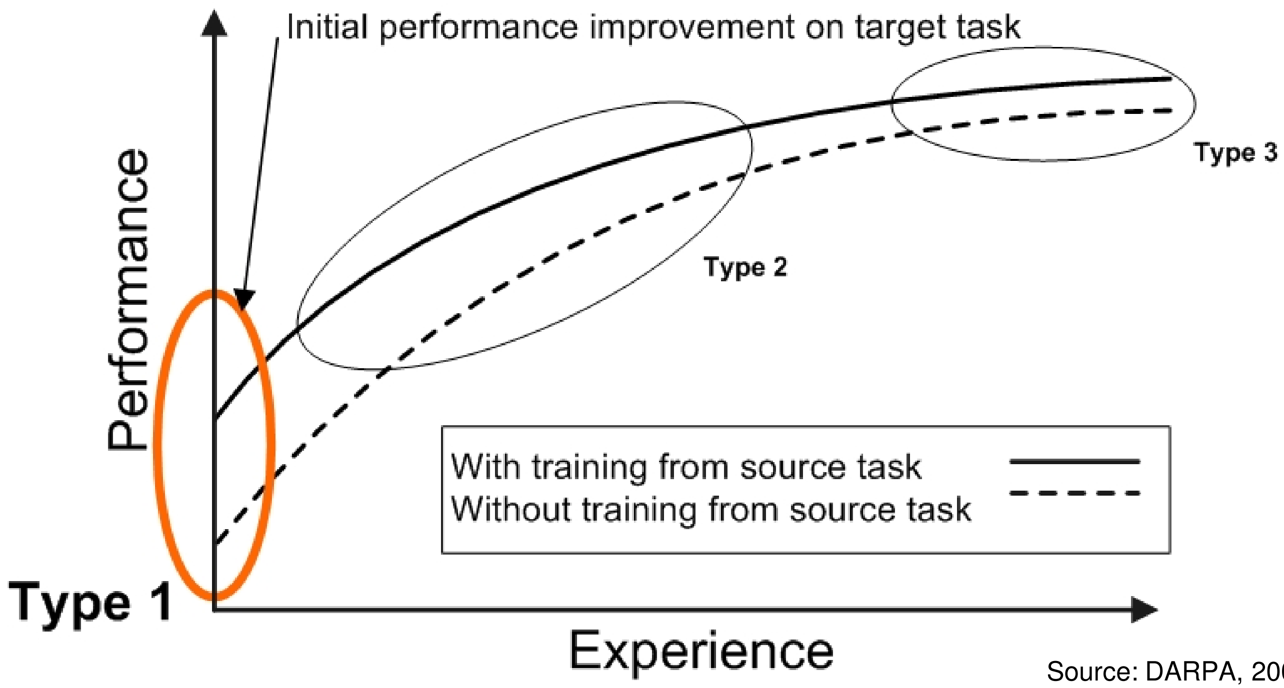

Source: DARPA, 2005

Implementation

"What is CHS-Soar-RL?"

- Combines constraint and rule based reasoning
- Soar: Symbolic cognitive architecture developed by Newell, Laird and Rosenbloom, 1983
  - Two forms of learning (Chunking and RL)
- Constrained Heuristic Search (CHS) developed by Fox, Sadeh and Baykan, 1989
  - Problem Topology
  - Problem Textures
  - Problem Objective

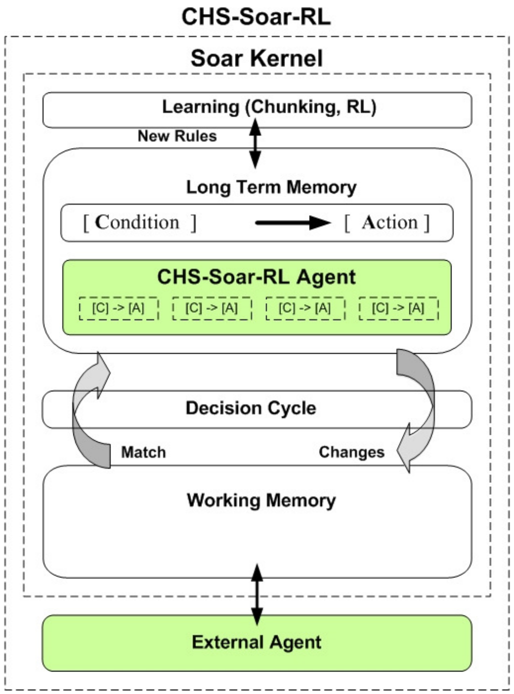

Implementation

A “texture” is not a heuristic itself, but can be considered the constituent parts of a heuristic

“What are texture measures?”

- A texture measurement is a technique for distilling information embedded in the constraint graph into a form that heuristics can use

Ordering	Name	Texture	Heuristic
Variable	Degree (DEG)	C_i, number of constraints to other unassigned variables linked to variable.	Select max texture.

Implementation

"What are texture measures?"

- A texture measurement is a technique for distilling information embedded in the constraint graph into a form that heuristics can use

Ordering	Name	Texture	Heuristic
Variable	Degree (DEG)	C_i, number of constraints to other unassigned variables linked to variable.	Select Max texture.

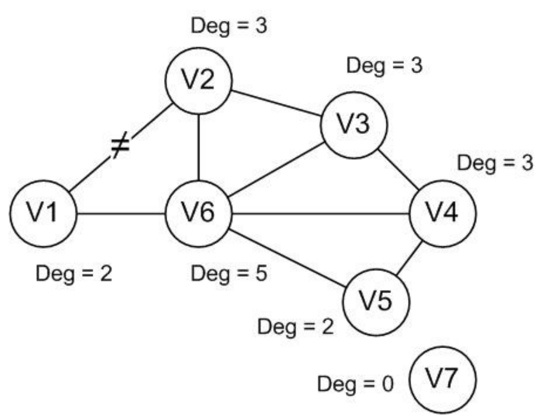

Variable	DEG	DEG
	Actual	*Normalized*
V_1	2	0.4
V_2	3	0.6
V_3	3	0.6
V_4	3	0.6
V_5	2	0.4
V_6	5	1.0
V_7	0	0.0

Range of Structural Detail

DEG Heuristic

Implementation

"How do we solve problems?"

CHS-SOAR PROBLEM SPACE

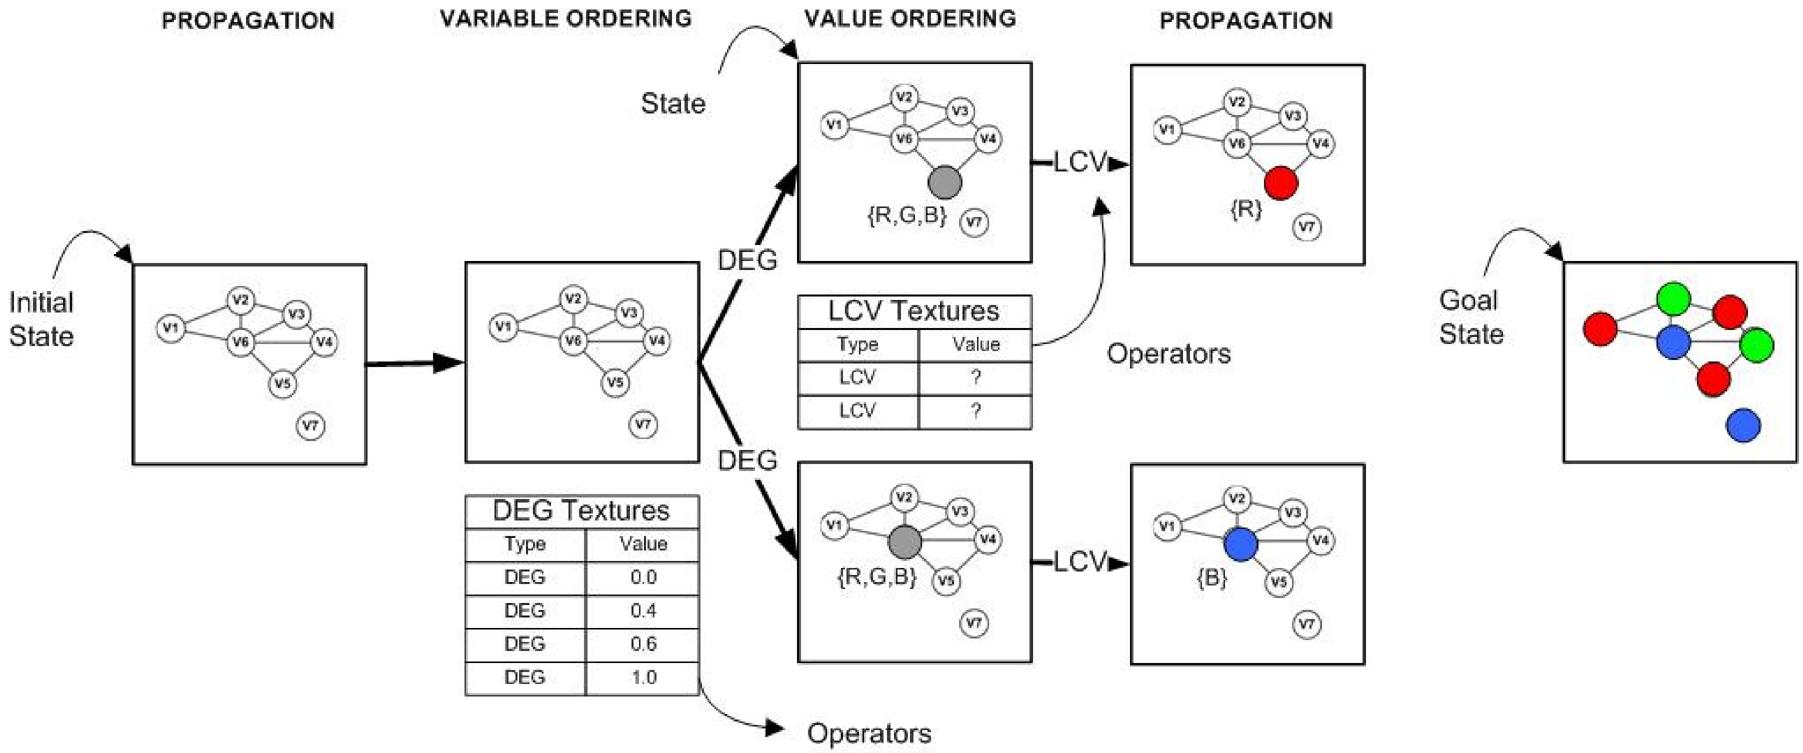

Implementation

"How do we learn?"

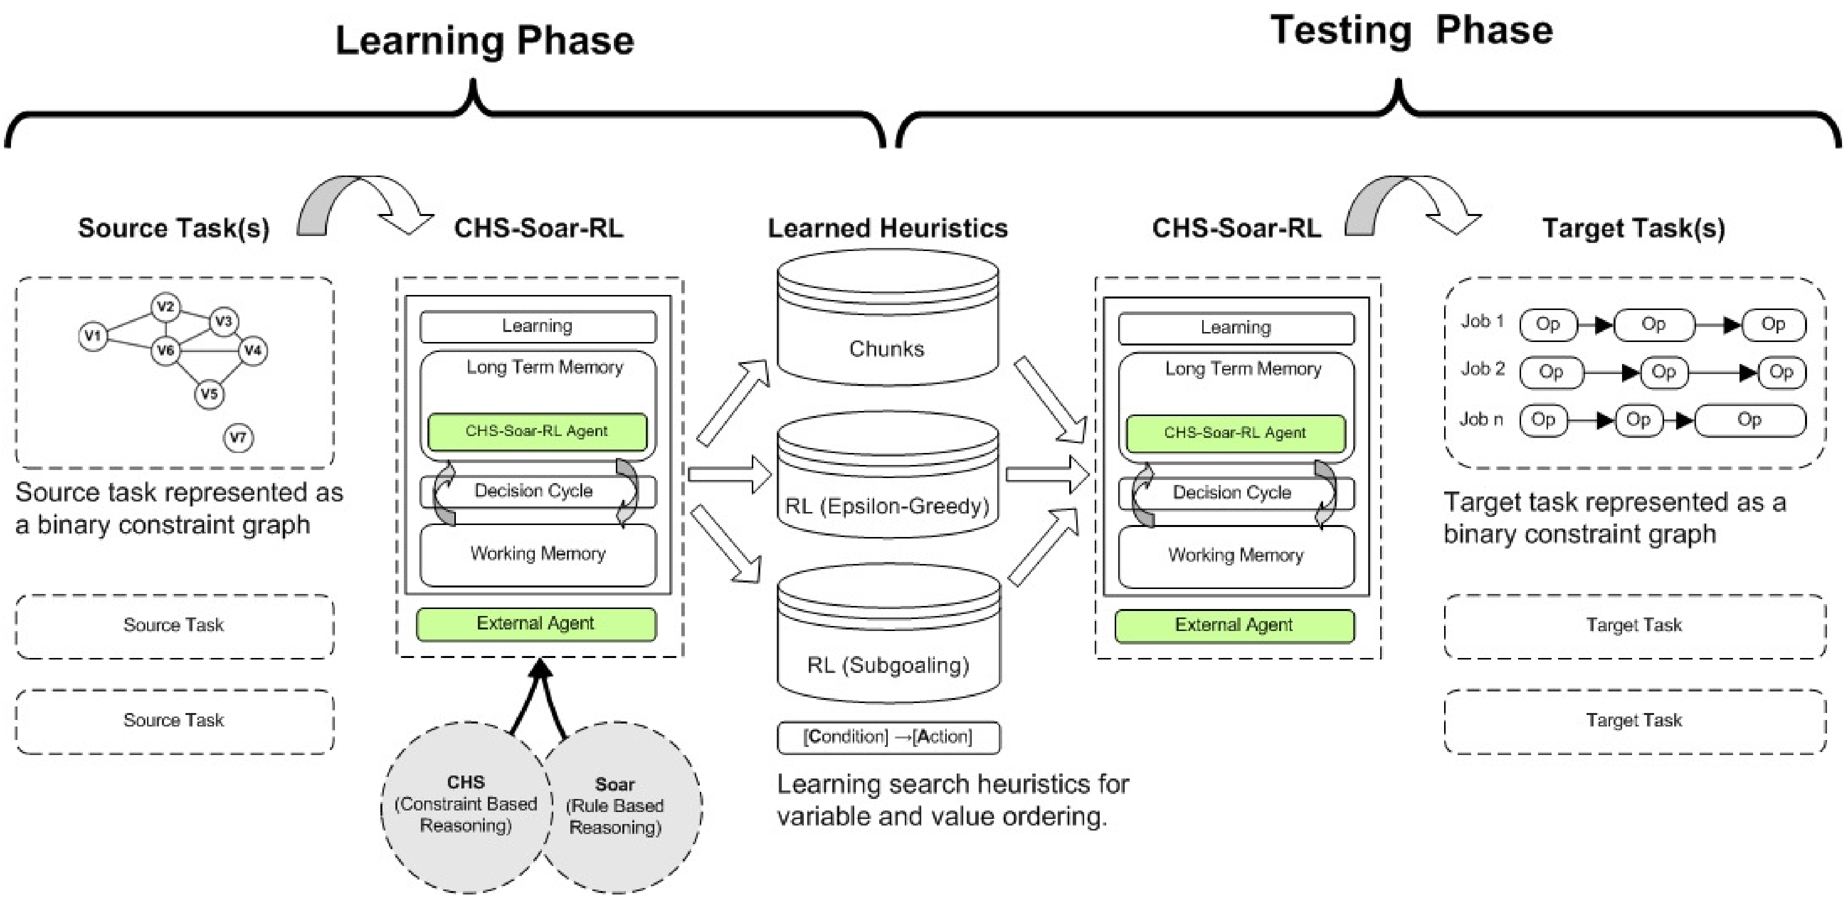

Experiments

Investigate target task initial performance improvement via transfer learning for 3 different learning approaches and 2 different sets of source tasks:

Experiment 1: Source Tasks: RGP (Random Generated CSP's)
Experiment 2: Source Tasks: Mix (set of CSP's comprised of: MCP, FAP & JSP)

Target Tasks	***Base***	***Instance Size (# of variables)***			
Size multiple of base		*1x*	*2x*	*3x*	*4x*
Map Coloring (MCP)	25	25	50	75	100
Job Shop Scheduling (JSP)					
Frequency Assignment (FAP)					
Traveling Salesman (TSP)					
Randomly Generated (RGP)					
N-Queens (NQP)	8	8	16	24	32

Experiments

Evaluation of target task performance based on:

- No Transfer (random selection)
- Transfer (with training from source tasks)
- Benchmark uses standard CSP heuristics (i.e. MRV, DEG, LCV)

Allows us to compare two metrics:

- Transfer Ratio: Performance change with transfer over no transfer
- Benchmark Ratio: Performance change with transfer over benchmark

Experiment 1: (RGP Source Tasks)

Transfer Ratio: How did we perform with transfer over the no transfer case?

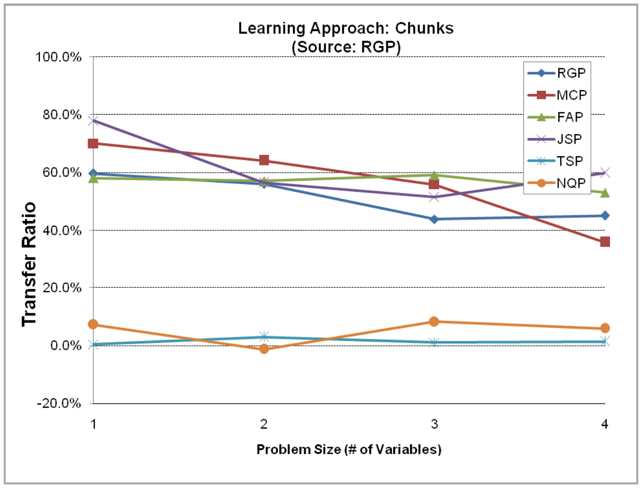

Chunks

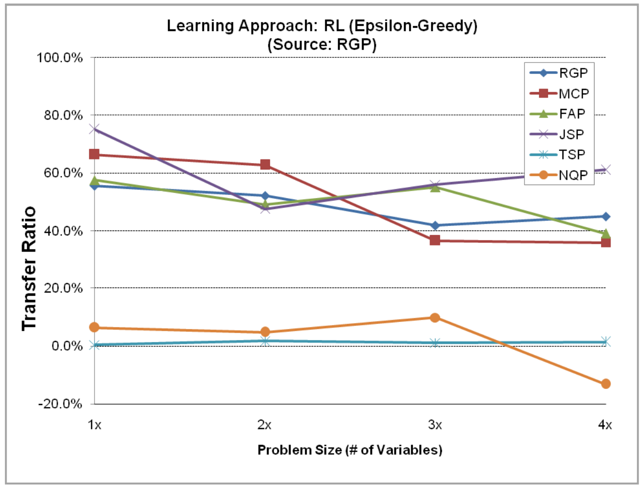

RL (Epsilon-Greedy)

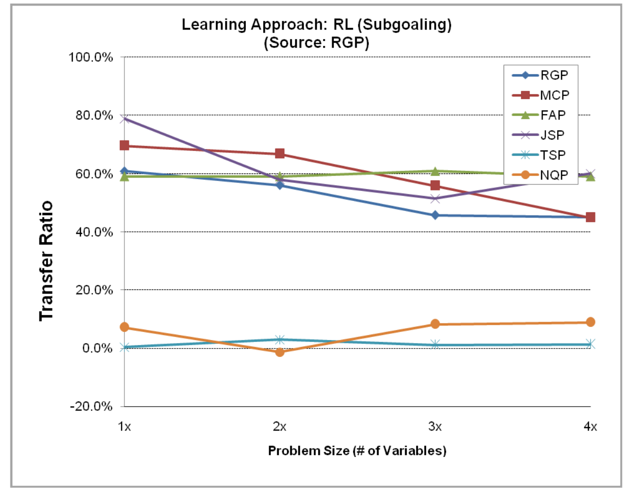

RL (Subgoaling)

For FAP, MCP and JSP we observe approximately 50% improvement for Chunks and RL (Subgoaling).

Experiment 1: (RGP Source Tasks)

Benchmark Ratio: How did we perform with transfer over the benchmark case?

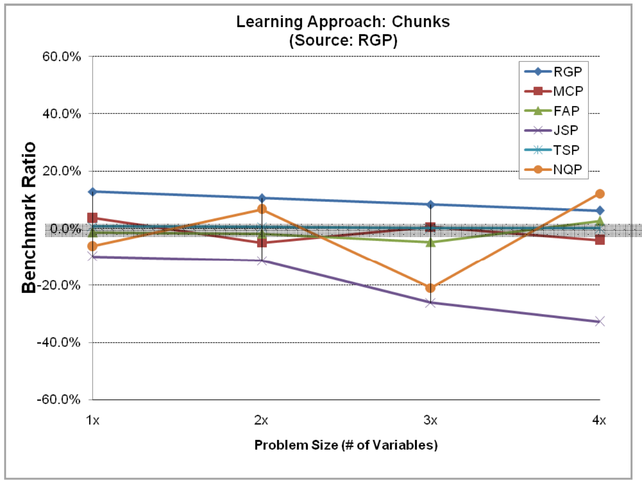

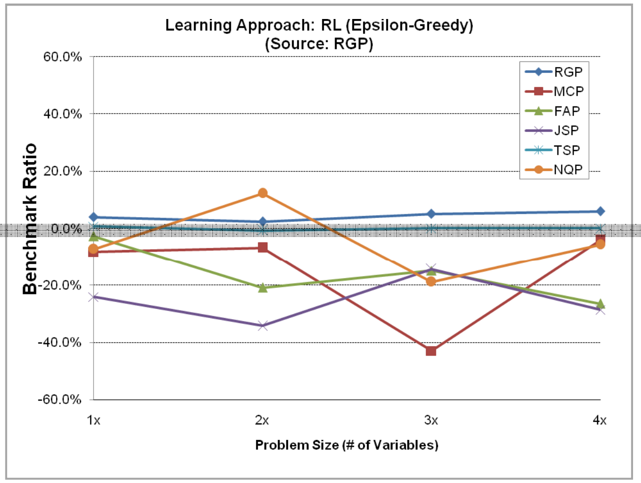

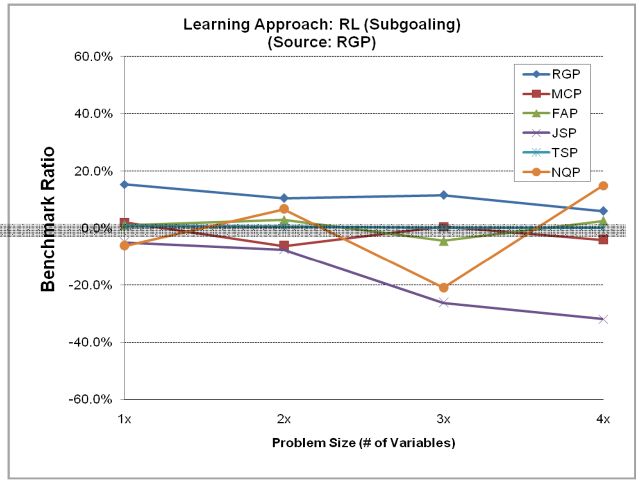

Chunks **RL (Epsilon-Greedy)** **RL (Subgoaling)**

Overall negative benchmark performance for all three learning approaches.

Experiment 2: (Mix Source Tasks)

Transfer Ratio: How did we perform with transfer over the no transfer case?

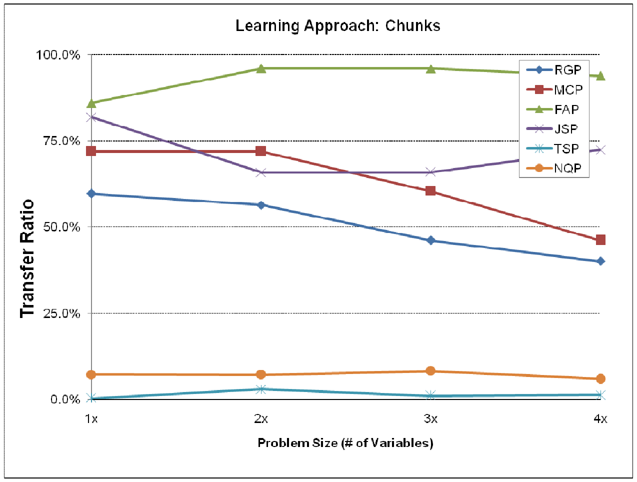

Chunks

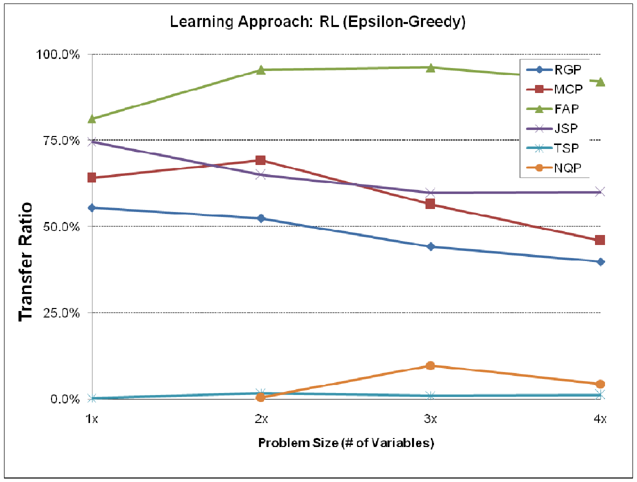

RL (Epsilon-Greedy)

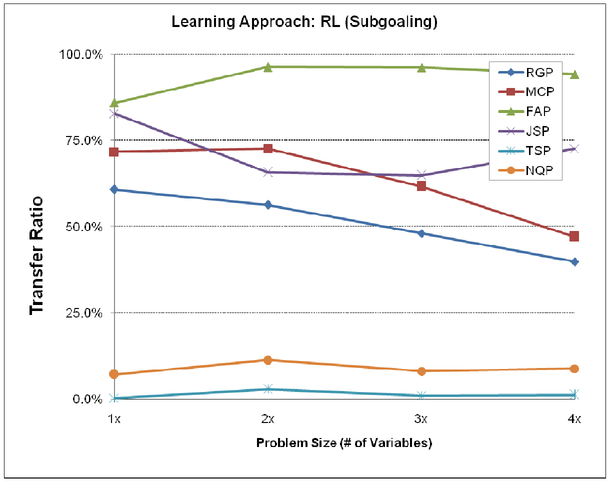

RL (Subgoaling)

For FAP, MCP and JSP we observe approximately 75% improvement for Chunks and RL (Subgoaling).

Experiment 2: (Mix Source Tasks)

Benchmark Ratio: How did we perform with transfer over the benchmark case?

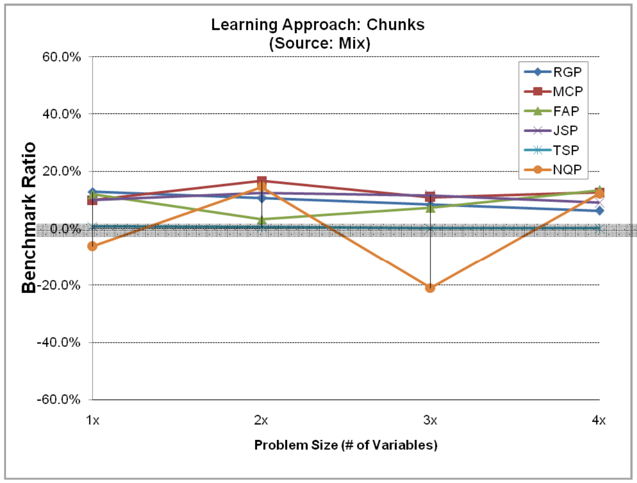

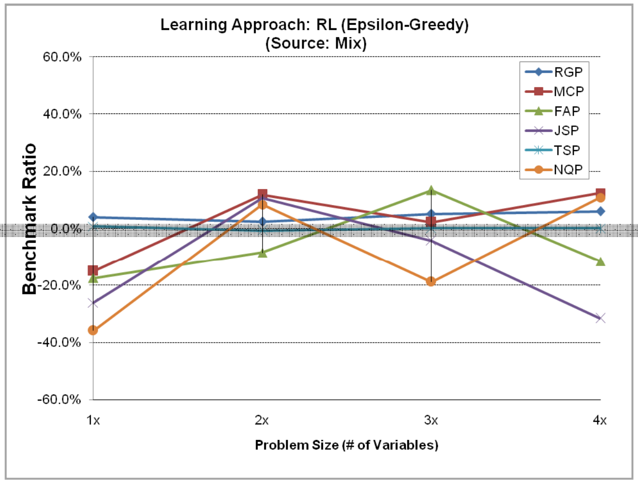

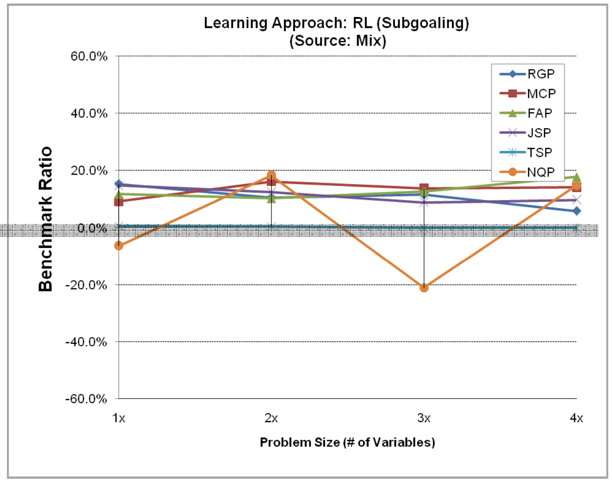

Chunks **RL (Epsilon-Greedy)** **RL (Subgoaling)**

For FAP, MCP and JSP we observe 9% improvement for Chunks and RL (Subgoaling) over benchmark case.

Experiment 2: (Mix Source Tasks)

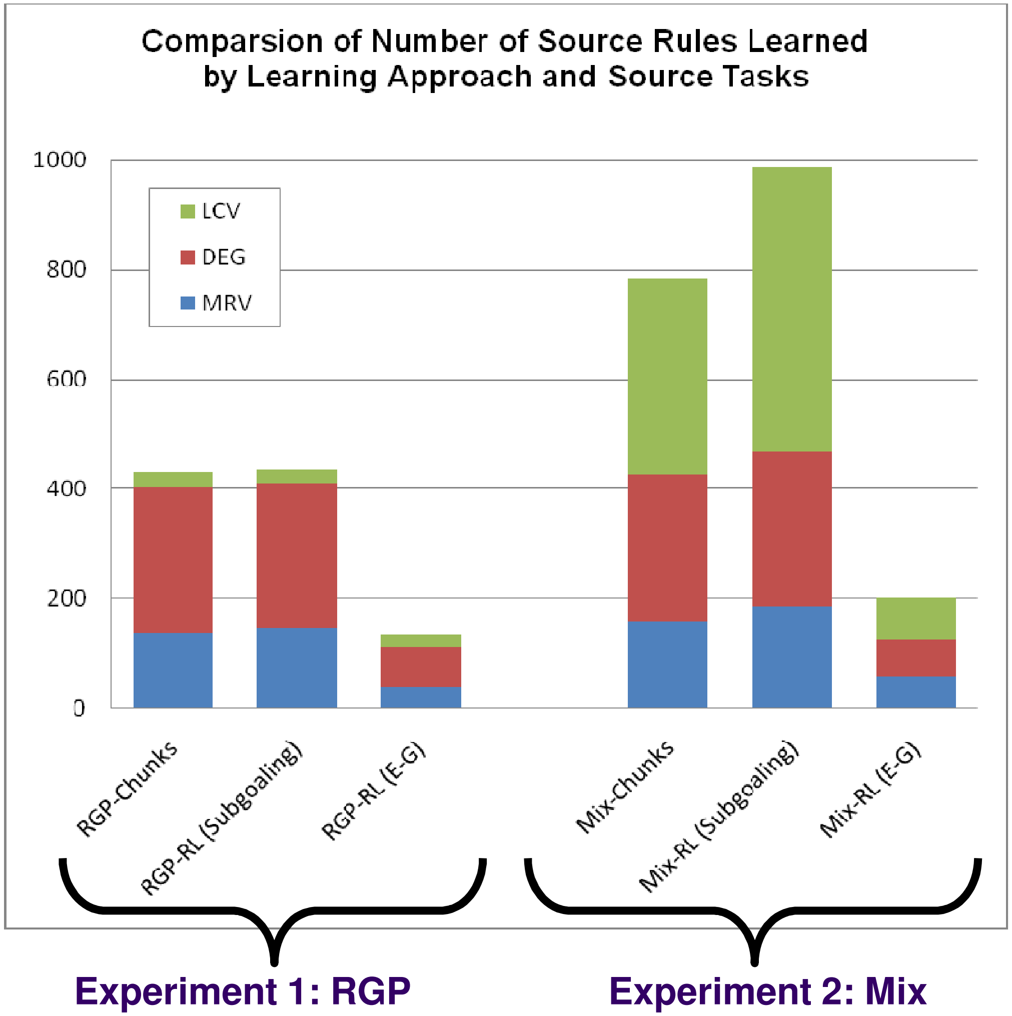

Results suggest "Mixed" source tasks provide more opportunities to encode search heuristic knowledge, particularly for value ordering over RGP.

Experiment 2: (Mix Source Tasks)

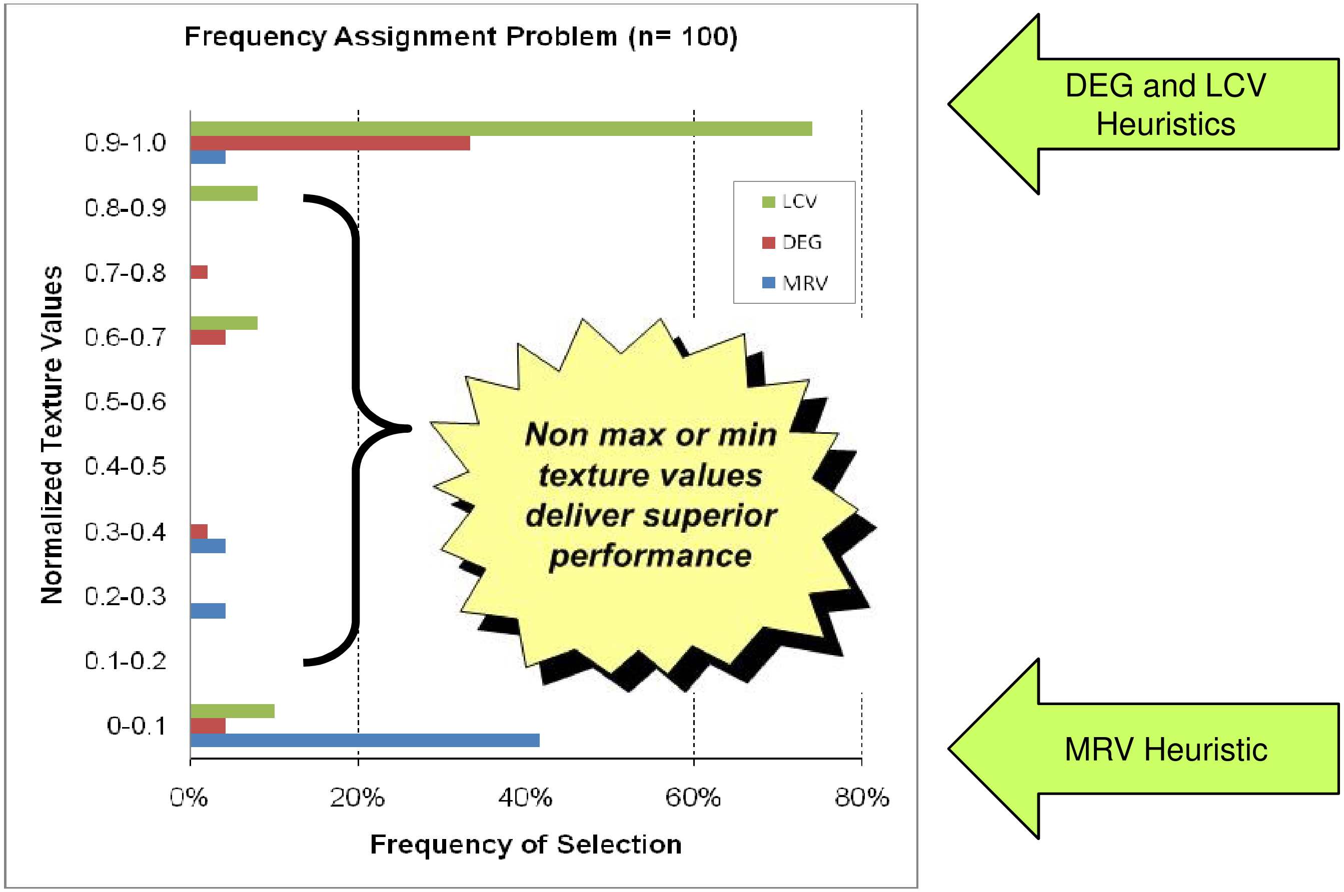

Summary

- Progress to date has demonstrated transfer learning for variable and value ordering in binary CSPs by combining constraint and rule based reasoning.
- Initial performance improvement (type 1) of:
  - 50% above the no transfer case
  - 9% above the benchmark case
  - Sensitive to the type of source training tasks.
- Expanding the expressiveness of the search heuristics by encoding intermediate (i.e. non min/max) texture details allows us to secure improved performance over benchmark heuristics for selected target tasks.
- Did not confirm the inherent benefit of reinforcement learning which allows us to dynamically revise the numerical preferences.

Current/Future Work

- Duration and diversity of training
- Improved texture evaluation functions (e.g. tree search)
- Additional (more insightful) texture measurements
- New structural features (i.e. density, tightness)
- Additional problem types (i.e. tasks)

Questions

TRANSFER LEARNING OF SEARCH HEURISTICS FOR CONSTRAINT SATISFACTION

(Introducing Constrained Heuristic Search to the Soar Cognitive Architecture)

Sean A. Bittle
Mark S. Fox