GP-Rush: Using Genetic Programming to Evolve Solvers for the Rush Hour Puzzle

A. Hauptman, A. Elyasaf, M. Sipper, A. Karmon

Ben-Gurion University

2009 "HUMIES" AWARDS FOR HUMAN-COMPETITIVE RESULTS

Friday, July 10, 2009

The Rush Hour Puzzle

- Sliding-blocks game played on 6x6 board
- Simple rules:
  - Car can move horizontally OR vertically
  - No hopping, no turning
- Purpose:
  move cars such that red car can exit

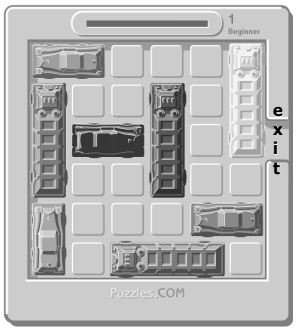

2

EASY TO LEARN

HARD TO PLAY

HARD FOR AIer

3

Previous Work

- n x n Rush Hour is PSPACE-complete [Flake & Baum, 2002]
  (hard to play, hard for AIer…)
- Discovery of all 6x6 solvable boards [Servais, 2006]
  Approach not scalable to larger boards

4

No Work on Solvers

- One free simple program
- BUT: No heuristic function currently exists
- Very difficult to estimate distance to goal
- WHY?
  1. Relaxing constraints spoils the game
     e.g., deleting cars, allowing cars to move freely
  2. Very difficult to find patterns / schemata
     one cell or car can totally alter play

5

Our Solution: 1. Heuristics

- We designed "human-like" heuristics for use with standard methods (e.g., IDA*)
- Example: BlockersLowerBound
  Lower bound on number of steps to goal,
  by counting moves needed to free blocking cars
- Goal distance, Hybrid, IsMoveToSecluded, …
- All proved limited (variable utility)

6

Our Solution: 2. Evolution

- Basic heuristics serve as building blocks
- Evolution may be used to:
  1) build new heuristics from existing building blocks
  2) Find weights for each heuristic
  3) Find conditions for applying each heuristic

7

Our Solution: 3. Policies

In the field of automated planning:
Policy = ordered set of deductive rules

Conditions	Results
Condition 1	Result 1
Condition 2	Result 2
…	…
Condition N	Result N
	Default Result

8

Two Goals

1. Evolve Solvers (GP-Evolved Policies)
2. Finding hard problems is hard:

   Evolve difficult 8x8 boards

Second goal arose because GP proved so successful at solving hard boards (and beating humans) we had to evolve new hard cases…

9

Results: 1. GP vs. Human AIer

our heuristics

	Without Heuristics	Blockers Estimation	Goal Distance	Hybrid	Hand-Crafted Policy	GP Policy
6 x 6	100%	72%	94%	102%	70%	**40%**
8 x 8	100%	69%	75%	70%	50%	**10%**

% of boards used in search compared to Iterative Deepening A*

Evolution drastically cuts amount of search

10

Results: 2. GP vs. Human Player

Time to solve (seconds)

	GP	Humans
Jam01…Jam08	0.03	2.6
Jam09…Jam16	0.6	8.15
Jam17…Jam24	0.83	10.32
Jam25…Jam32	1.17	14.1
Jam33…Jam40	2.65	20.00
Average	1.04	11.03

- Humans:
  - best of thousands at www.trafficjamgame.com
  - probably time to play (not solve), so gap much wider
- More than mere raw computing power

11

Why is Result Best? (1)

PUSHING EVOLUTION FURTHER

- Most difficult single-player search (i.e., planning) problem solved (so successfully) with evolution to date
- 6x6 Rush Hour more difficult than all other planning problems solved evolutionarily (difficult to design representation + huge, hard-to-navigate search space)
- Moreover, we evolved (& solved) yet harder 8x8 boards, never tackled before

12

Why is Result Best? (2)

SEVERAL DEGREES (& MODALITIES) OF IMPROVEMENT

- Popular Enhanced Iterative Deepening algorithm surpassed by our hand-crafted heuristics and policies, all of which were beaten by GP-evolved strategies
- Evolution managed to take our best designed ingredients of limited performance and transform them into highly successful strategies
- GP not only beat human AI researchers but also all human players of Rush Hour on record

13

Why is Result Best? (3)

SOLVE DIFFICULT PROBLEM WITH LONG HISTORY

- Difficult puzzles (involving search and planning) have a longstanding tradition in AI
- Rush Hour considered open problem until very recently [Kendall et al. 2008]

  No efficient solvers designed, despite fertility of field of automated planning

  (Note not only Rush Hour's open status but also its complexity, PSPACE-complete, superseding 23 other games described in 2008 Kendall survey paper, which are "only" NP-Complete)

14

Why is Result Best? (4)

- Our evolutionary algorithm "closed" Rush Hour's open status, in addition exhibiting the ability to scale up to new, more difficult problems — themselves discovered through evolution
- We used evolution to generate the most difficult Rush Hour problems known
- Thus, we evolved both the best known solvers

  and

  the most difficult existing boards

15

Result is Human-Competitive

(B) equal to / better than new scientific result

(D) publishable in its own right as new scientific result

(F) equal to / better than achievement in its field

(G) solves problem of indisputable difficulty in its field

(H) holds its own / wins competition vs. human

16