

Learned Models of Performance for Many Planners

Mark Roberts, Adele Howe and Landon Flom
Computer Science Department
Colorado State University

The Study


- 28 publically available planners (from 1995 to 2006)
- 4726 publically available STRIPS planning problems
- Same hardware, default parameters, 30 minutes cut-off
- Extracted solutions to determine success and monitored time to solution

The Learned Models

- Features
 - 32 easily extracted from PDDL domain and problem descriptions
 - Examples: # of goals, # actions with negative effects, whether required fluents, max # of preconditions in actions...
- Restricted data to only 1215 “challenging problems”
- 32 Algorithms (from WEKA)
- Two models per planner:
 - success
 - time to successfully solve

Colorado State University


Model Accuracy: Success



Mean %fail: 86.61, Mean Accuracy: 96.59

Colorado State University

Model Accuracy: Time




Mean time: 238.38, Median time: 2.66

Colorado State University

Model Generalization

How well does a model trained on “old” problems do on “new” problems?

- Test on IPC5, train on others from challenge set



Mean %fail: 86.15, Mean Acc: 94.47

Mean time: 111.21, Mean RMSE 139.11

Colorado State University

Model Utility: Portfolio of Planners

- Constructed planner portfolio with subset of planners using models to drive planner selection and time allocation strategies.
 - Both time and success models worked well for informing selection.
 - Time model did not work well for allocation, beaten by just using mean time.
- Best learned portfolio out-performs best single planner on time and %fail.

Colorado State University