View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by Massey Research Online

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

Massey University

Action-Selection in RoboCup Keepaway Soccer: Experimenting with Player Confidence

A thesis presented in partial fulfilment of the requirements for the degree of

Masters of Science

In

Computer Science

at Massey University, Albany,

New Zealand.

Samara Ann Neilson

2006

Copyright © 2006
Samara Ann Neilson
All Rights Reserved

Abstract

Through the investigation of collaborative multi-agent domains, in particular those of robot soccer and robot rescue, and the examination of many popular action-selection methodologies, this study identifies some of the issues surrounding entropy, action-selection and performance analysis.

In order to address these issues, a meaningful method of on-field player evaluation, the confidence model, was first proposed then implemented as an action-selection policy. This model represented player skill through the use of percentages signifying relative strength and weakness and was implemented using a combination of ideas taken from Bayesian Theory, Neural Networks, Reinforcement Learning, Q-Learning and Potential Fields.

Through the course of this study, the proposed confidence model action-selection methodology was thoroughly tested using the Keepaway Soccer Framework developed by Stone, Kuhlmann, Taylor and Liu and compared with the performance of its peers.

Empirical test results were also presented, demonstrating both the viability and flexibility of this approach as a sound, homogeneous solution, for a team wishing to implement a quickly trainable performance analysis solution.

Acknowledgements

I would like to gratefully acknowledge the enthusiastic supervision of Dr. Chris Messom during this work. I would also like to thank Stephen Ford for the assistance provided dealing with Helix's technical problems. And, finally, I am forever indebted to my friends and family for their understanding, endless patience, encouragement and support when it was most required.

Table of Contents

Abstract i

Acknowledgements i

List of Tables v

List of Figures vi

Chapter 1: Introduction & Study Background

1.1 Introduction 1

1.2 Statement of the Problem 4

1.3 General Research Hypotheses 4

1) Collaborative Multi-Agent Domains 5

2) Entropy and Action-Selection 5

3) Assessing Individual Performance 6

4) Implementation 6

5) Performance Assessment 7

1.4 Significance of the Study 7

1.5 Definitions & Operational Terms 7

1.6 Delimitations 9

Chapter 2: Collaborative Multi-Agent Domains

2.1 Introduction 10

2.2 Robot Soccer 16

2.2.1 Robot Soccer Initiatives 17

2.2.1.1 RoboCup Soccer 17

2.2.1.2 FIRA 19

2.2.2 Robot Soccer Simulators 21

2.2.2.1 RoboCup Soccer Server 22

2.2.2.2 SimuroSot 3D Soccer Simulator 23

2.2.2.3 TeamBots 25

2.2.2.4 Webots 26

2.2.3 Real Robot Soccer 27

2.2.4 Soccer Keepaway 30

2.3 Robot Search and Rescue 31

2.3.1 Robot Rescue Initiatives 32

2.3.1.1 RoboCup Rescue 32

2.3.1.2 NIST 33

2.3.2 Robot Rescue Simulation 34
2.3.2.1 RoboCup Rescue Simulator 34
2.3.3 Real Robot Rescue 35
2.4 Summary 36

Chapter 3: Action Selection Methodologies and Implementations
3.1 Action Selection Methodologies 38
3.1.1 Bayesian Theory 38
3.1.2 Neural Networks 39
3.1.3 Potential Field Methodologies 41
3.1.4 Reinforcement Learning 42
3.1.5 Q-Learning 44
3.2 Implementations
3.2.1 Simulation Environment 45
3.2.1.1 RoboCup Soccer Server 45
3.2.1.2 Keepaway Soccer Framework 47
3.2.2 Action-Selection Implementations 50
3.2.2.1 Confidence Model 51
3.2.2.2 Direct Implementation 52
3.2.2.3 Potential Field Confidence Model Implementation 54
3.3 Empirical Testing Regime 55
3.3.1 Benchmark Testing 55
3.3.2 Feasibility Testing 56
3.3.3 Scalability Testing 56
3.3.4 Entropy Testing 56
3.4 Summary 57

Chapter 4: Results
4.1 Benchmark Testing 59
4.2 Feasibility Testing 60
4.2.1 2 vs. 1 Keepaway 61
4.2.2 3 vs. 2 Keepaway 64
4.2.3 4 vs. 3 Keepaway 67
4.2.4 5 vs. 4 Keepaway 71
4.2.5 Summary 74
4.3 Scalability Testing 75
4.3.1 3 vs. 2 Keepaway on Differing Field Sizes 76

4.3.1.1 15 Metre Field 76
4.3.1.2 40 Metre Field 79
4.3.2 11 vs. 11 Keepaway 82
4.3.3 Summary 85
4.4 Entropy Testing 87

Chapter 5: Discussions and Conclusion
5.1 Discussion 90
5.1.1 Keepaway Soccer Framework 90
5.1.2 Confidence Model Observations & Results 91
5.2 Conclusion 94
5.3 Future Work 98

Chapter 6: References 99

Appendices:
Appendix A1: User Guide & Code Printout 104
A1.1 User Guide 104
A1.1.1 Installation 104
A1.1.2 Configuration Settings 104
A1.1.3 Sample Startup Script 105
A1.2 Printout of `LearningAgent.h` Code 106
A1.3 Printout of `LearningAgent.cc` Code 108
Appendix A2: Result Graphs 119
A2.1 Feasibility Testing Graphs 120
A2.2 Scalability Testing Graphs 132
A2.3 Entropy Testing Graphs 141

List of Tables

Table 2.1:	Differences between the RoboCup Rescue Simulator and the RoboCup Soccer Server.	35
Table 3.1:	Feasibility testing regime.	56
Table 4.1:	Keepaway framework results provided by Stone et al (2005).	59
Table 4.2:	Results recorded from testing conducted over 10,000 episodes on the Helix computer system.	59
Table 4.3:	Observed difference between results benchmarked by Stone et al and collected from Helix.	60
Table 4.4:	Summary of results for 2 vs. 1 keepaway soccer.	63
Table 4.5:	Summary of results for 3 vs. 2 keepaway soccer.	66
Table 4.6:	Summary of results for 4 vs. 3 keepaway soccer.	70
Table 4.7:	Summary of results for 5 vs. 4 keepaway soccer.	73
Table 4.8:	Policy performance summary for results collected using Helix.	74
Table 4.9:	Summary of results for 3 vs. 2 keepaway soccer on a 15m x 15m field.	78
Table 4.10:	Summary of results for 3 vs. 2 keepaway soccer on a 40m x 40m field.	81
Table 4.11:	Summary of results for 11 vs. 11 keepaway soccer.	84
Table 4.12	Policy performance results collected during 3 vs. 2 feasibility testing and 3 vs.2 scalability testing.	85
Table 4.13:	Performance results collected during feasibility testing and 11 vs. 11 scalability testing.	86
Table 5.1:	Average performance comparison between version 0.4 and 0.6 of the Framework.	93
Table 5.2:	Performance comparison between the 'Potential Fields' policy and 'Reinforcement Learning' policy for 3 vs. 2 keepaway over differing field sizes.	94
Table A1.1:	Keepaway settings.	104

List of Figures

Figure 2.1:	A taxonomy of autonomous agents.	15
Figure 2.2:	FIRA logo and examples of participant classes.	19
Figure 2.3:	RoboCup Soccer Server.	22
Figure 2.4:	Screenshot of the SimuroSot Simulator in action.	24
Figure 2.5:	Screenshot of SoccerBots showing set kick-off positioning.	25
Figure 2.6:	Simulated soccer game between two Sony AIBO robots.	26
Figure 2.7:	Small-sized team at RoboCup 2005.	28
Figure 2.8:	Medium-sized team at RoboCup 2005.	28
Figure 2.9:	Four-legged team at RoboCup 2005.	29
Figure 2.10:	Humanoid robot from team Osaka getting ready to compete at RoboCup 2005.	29
Figure 2.11:	Example of an episode with 3 keepers and 2 takers playing in a 20m x 20m region.	30
Figure 2.12:	RoboCup Rescue simulator in action at RoboCup 2005.	34
Figure 3.1:	Artificial neuron with input vectors (x), weighted input lines (w) and a thresholding function f()that determines the neurons output value.	40
Figure 3.2:	Example network topologies for the PassEvaluate() and GetOpen() functions.	41
Figure 3.3:	Uniform, attractive, repulsive and tangential potential fields.	42
Figure 3.4:	Model of a player's visual sensor showing its view cone and visible distance where any agents seen can be identified with a confidence of 100%. This confidence decreases by given proportions as players move farther outside the visible distance into the marked zones.	47
Figure 3.5:	Starting positions for 3 vs. 2 keepaway soccer.	48
Figure 3.6:	Representations of keepaway state variables.	49
Figure 3.7:	Human soccer player decision scenario.	50
Figure 3.8:	A direct confidence level approach to action-selection.	53
Figure 3.9:	A non-optimal passing scenario is displayed to the left with a visualisation of the beginning stages of the potential fields hybrid solution. The results of the solution, once optimised, are shown to the right along with the emergent behaviour.	55

Figure 4.1:	Policy Performance Graph for 2 vs. 1 keepway.	61
Figure 4.2:	Box and whisker analysis of 2 vs. 1 policy performance.	62
Figure 4.3:	Graph depicting ratio of episode outcomes for 2 vs. 1 keepaway.	63
Figure 4.4:	Policy performances for 3 vs. 2 keepaway.	64
Figure 4.5:	Box and whisker analysis of 3 vs. 2 policy performance.	65
Figure 4.6:	Graph depicting the ratio of outcomes of 3 vs. 2 keepaway.	66
Figure 4.7:	Policy performance graph for 4 vs. 3 keepaway.	68
Figure 4.8:	Box and whisker analysis of 4 vs. 3 keepaway soccer performance.	69
Figure 4.9:	Outcome ratios for 4 vs. 3 policy performance.	70
Figure 4:10:	Policy performance graph for 5 vs. 4 keepaway.	71
Figure 4.11:	Box and whisker analysis of 5 vs. 4 keepaway soccer performance.	72
Figure 4:12:	Outcome ratios for 5 vs. 4 policy performance.	73
Figure 4:13:	Policy performance summary for results collected using Helix.	75
Figure 4:14:	Policy performance graph for 3 vs.2 keepaway played on a 15m x 15m field.	76
Figure 4.15:	Box and whisker analysis of 3 vs.2 keepaway performance on a 15m x 15m field.	77
Figure 4:16:	Outcome ratios for 3 vs. 2 policy performance on a 15m x 15m field.	78
Figure 4:17	Policy performance graph for 3 vs.2 keepaway played on a 40m x 40m field.	79
Figure 4:18:	Box and whisker analysis of 3 vs.2 keepaway performance on a 15m x 15m field.	80
Figure 4:19:	Outcome ratios for 3 vs. 2 policy performance on a 40m x 40m field	81
Figure 4:20:	Policy performance graph for 11 vs.11 keepaway.	82
Figure 4:21:	Box and whisker analysis of 11 vs. 11 keepaway performance.	83
Figure 4:22	Outcome ratios for 11 vs. 11 policy performance.	84
Figure 4:23	3 vs. 2 policy performance comparison for differing field sizes.	86
Figure 4:24:	Policy performance comparison for differing team configurations.	87
Figure 4:25:	Keeper passing confidences observed for a handicapped 3 vs. 2 team.	88
Figure 4:26:	Keeper passing confidences observed for a 3 vs. 2 team.	89
Figure 5.1:	"Piggy in the middle" passing exhibited by players using the Potential Fields policy implementation.	92