

Table of Contents

1	Introduction	1
1.1	Motivation and Problem Statement	1
1.2	Objectives	3
1.2.1	General Objective	3
1.2.2	Specific Objectives	3
1.3	Hypothesis	4
1.4	Outline	4
2	Theoretical Framework	5
2.1	Sequential Decision Making	5
2.1.1	Markov Decision Process	5
2.2	Reinforcement Learning	6
2.2.1	Learning Policies	8
2.3	Function Approximation	11
2.3.1	Linear Model of Basis Functions	12
2.3.2	Artificial Neural Networks	12
2.3.3	Parameter Tuning	19
2.3.4	Deep RL Era	19
2.4	Interactive Machine Learning	20
2.4.1	Learning from Demonstration	20
2.4.2	Learning from Feedback	22
3	Introducing Deep Neural Networks into the COACH Framework	25
3.1	Overview	25
3.2	Methodology	26
3.2.1	Low-Dimensional vs High-dimensional States	27
3.2.2	The Algorithm	28
3.3	Experiments and Results	29
3.3.1	Validation of replay buffer with simulated teachers	32
3.3.2	Comparison of DRL and D-COACH OFF using real human teachers .	34
3.3.3	Validation in a real system	36
3.4	Discussion	37
4	Enhancing D-COACH for High-Dimensional State Problems	38
4.1	Overview	38
4.2	Online State Representation Learning	38

4.3	The Algorithm	39
4.4	Experiments and Results	41
4.4.1	Study with Simulated teacher	42
4.4.2	Experiments with real human teachers	44
4.4.3	Additional validation with real systems	46
4.5	Discussion	47
5	Adding Memory to the Agents	48
5.1	Overview	48
5.2	Methodology	49
5.2.1	Learning to Remember	49
5.2.2	Low-dimensional State with Memory	51
5.2.3	High-dimensional State with Memory	53
5.2.4	The Algorithm	55
5.3	Experiments and Results	56
5.3.1	Validation Low-Dimensional State	59
5.3.2	Validation High-Dimensional State	60
5.4	Discussion	62
6	Conclusion and Future Work	63
	Bibliography	66
7	Appendix	72
7.1	Publications	72