

2006 NSF WORKSHOP AND OUTREACH TUTORIALS ON APPROXIMATE DYNAMIC PROGRAMMING

April 3-6, 2006
Hotel Hacienda
Cocoyoc, Mexico

At the charming Hotel Hacienda in Mexico, we will advance the field of approximate dynamic programming by bringing together researchers and students from the fields of neural networks, operations research and artificial intelligence. A major theme of the workshop will be communication – bridging the inevitable language barriers that arise when different communities work on similar problems. All of the sessions have representatives from two or all three of the disciplines represented by the workshop, allowing participants to contrast how different communities address similar issues. An elegant setting and relaxed pace will allow workshop participants to discuss ideas informally.



Joining 40 leading researchers will be over 50 students and young professionals who will benefit from a series of tutorials that will help them learn the languages of approximate dynamic programming from the perspectives of different fields. In addition to an invaluable introduction to the field, the workshop will be an opportunity to meet people who are likely to form the basis of a professional peer group.

General Co-Chairs:

Warren Powell, Princeton University, Powell@princeton.edu

Jennie Si, Arizona State University, si@asu.edu

Tutorial and Local Co-Chairs

Francisco Cervantes Pérez, UNAM, Francisco_cervantes@cuaed.unam.mx

Nydia Lara, UNAM, nydia@aleph.cinstrum.unam.mx

CONFERENCE PROGRAM

The general organization of the workshop is as follows:

Sunday, April 2

Arrival, reception and opening remarks

Monday – Wednesday, April 3-5

Each day consists of four sessions with three speakers each. The middle of each day provides for a three-hour break for lunch and two tutorials. Students are expected to attend each tutorial, but these are open to the entire workshop. For those not attending a tutorial on a given day, this is an opportunity to meet informally with other workshop attendees to exchange thoughts and ideas.

Thursday, April 6

Panel – What did we learn? Panelists will be named on Wednesday afternoon. This will be an opportunity to summarize what we learned and discuss future directions for the field.

Each session consists of three speakers chosen from two and often three of the different disciplines represented at the workshop (“neural networks”, typically representing engineering, operations research and artificial intelligence). In addition, each session has a theme, where the goal is to illustrate how the different disciplines approach the theme. Each person will have 25 minutes to speak, providing the session chair with 10-15 minutes to synthesize the perspectives offered by each of the speakers.

2006 NSF Workshop and Outreach Tutorials on Approximate Dynamic Programming

Time	Sunday April 2
	Travel
17:30 - 19:30	Reception
19:30 - 20:30	Opening remarks: Goals, themes and the language of ADP

2006 NSF WORKSHOP AND OUTREACH TUTORIALS ON APPROXIMATE
DYNAMIC PROGRAMMING

Time	Monday April 3
	Session: A survey of applications Chair: Silvia Ferrari
8:30 - 8:55	Dan Bienstock: Risk, dynamic programming, and robust optimization: a study of two applications
8:55 - 9:20	Ben Van Roy: Internet Switches, Revenue Management, and a Question on Algorithm Performance
9:20 - 9:45	S.N. Balakrishnan: SNAC based Nonlinear Controllers and Observers
9:45 - 10:00	Chair's comments:
10:00 - 10:30	Break
	Session: Perspectives on learning Chair: Ron Parr
10:30 - 10:55	Eduardo Morales: Abstraction and Guidance for Scaling-up Reinforcement Learning
10:55 - 11:20	Warren Powell: Merging machine learning and math programming for solving high-dimensional resource allocation problems
11:20 - 11:45	Donald Wunsch: Integrating Supervised, Unsupervised, and Reinforcement Learning
11:45 - 12:00	Chair's comments:
12:00 - 13:30	Lunch
13:30 - 14:15	Tutorial: George Lendaris - Basics of DHP type Adaptive Critics / Approximate Dynamic Programming and some application issues
14:15 - 15:00	Tutorial: Mike Fu - Introduction to simulation
	Session: State variables in ADP Chair: Warren Powell
15:00 - 15:25	Dimitri Bertsekas: Recent Results on the Approximate Solution of POMDP
15:25 - 15:50	Edgar Sanchez: Discrete-time recurrent neural control: Nonlinear System Identification, State Estimation and Control
15:50 - 16:15	Satinder Singh: Rethinking State in Reinforcement Learning
16:15 - 16:30	Chair's comments:
16:30 - 17:00	Break
	Session: Stability, sensitivity and convergence Chair: John Tsitsiklis
17:00 - 17:25	Charles Anderson: Robust Reinforcement Learning Control with Guaranteed Stability While Learning
17:25 - 17:50	Julie Higle: Approximations and Convergence in Stochastic Linear Programming
17:50 - 18:15	Bob Narendra: Issues in the Stability and Optimization of Complex Adaptive Systems
18:15 - 18:30	Chair's comments:
18:30 - 19:30	Drinks
19:30 - 21:30	Dinner

2006 NSF WORKSHOP AND OUTREACH TUTORIALS ON APPROXIMATE
DYNAMIC PROGRAMMING

Time	Tuesday April 4
	Session: ADP and biology Chair: Jennie Si
8:30 - 8:55	Alfredo Weitzenfeld: Mobile Robotic Systems: From Biology to Soccer
8:55 - 9:20	Litsa Micheli-Tzanakou: Does the brain show any approximate dynamic programming capabilities in learning?
9:20 - 9:45	Emanuel Todorov: Biologically-inspired approximations to optimal control
9:45 - 10:00	Chair's comments:
10:00 - 10:30	Break
	Session: Bridging the gap: from DP, to SP to control Chair: Andrzej Ruszczynski
10:30 - 10:55	Xiren Cao: A Sensitivity View of Learning and Optimization
10:55 - 11:20	Terry Rockafellar: Dynamic Programming versus Stochastic Programming
11:20 - 11:45	Suvrajeet Sen: On Connections Between Differential Dynamic Programming and Nested Benders' Decomposition
11:45 - 12:00	Chair's comments:
12:00 - 13:30	Lunch
13:30 - 14:15	Tutorial: S.N. Balakrishnan - Implementation of ADP based Neurocontrollers-- Control of Vibration and Heat Transfer and Autopilot
14:15 - 15:00	Tutorial: Dimitri Bertsekas - Dynamic Programming and Suboptimal Control: A Survey from ADP to MPC
	Session: Approximation strategies Chair: Ben Van Roy
15:00 - 15:25	Ron Parr: Linear Value Function Approximation and Linear Models
15:25 - 15:50	Diego Klabjan: Ridge Functions in Approximate Dynamic Programming for Resource Allocation Problems
15:50 - 16:15	David Morton: Assessing Solution Quality in Stochastic Optimization
16:15 - 16:30	Chair's comments:
16:30 - 17:00	Break
	Session: Actor-critic methods Chair: Paul Werbos
17:00 - 17:25	Derong Liu: Adaptive Critic Learning Techniques for Engine Torque and Air-Fuel Ratio Control
17:25 - 17:50	Silvia Ferrari: Reconfigurable Neural Network Control by Adaptive Critics and Linear Matrix Inequalities
17:50 - 18:15	Jagannathan Sarangapani: Supervised Actor-Critic Reinforcement Learning Schemes for Control
18:15 - 18:30	Chair's comments:
18:30 - 19:30	Drinks
19:30 - 21:30	Dinner

2006 NSF WORKSHOP AND OUTREACH TUTORIALS ON APPROXIMATE
DYNAMIC PROGRAMMING

Time	Wednesday April 5
	Session: Learning for complex systems Chair: George Lendaris
8:30 - 8:55	Frank Lewis and Murad Abu Khalaf: Model-Free Adaptive Dynamic Programming for Uncertain Linear Continuous and Discrete-Time Systems
8:55 - 9:20	Sridhar Mahadevan: Learning Representation and Behavior: Manifold and Spectral Methods for Markov Decision Processes and Reinforcement Learning
9:20 - 9:45	Stuart Russell: Reinforcement learning with partial programs
9:45 - 10:00	Chair's comments:
10:00 - 10:30	Break
	Session: Perspectives on handling uncertainty Chair: Danil Prokhorov
10:30 - 10:55	Michael Fu: Simulation-Based and Evolutionary Algorithms for Markov Decision Processes
10:55 - 11:20	Shie Mannor: Reinforcement Learning with Gaussian Processes
11:20 - 11:45	Ganesh Venayagamoorthy: ADP for Dynamic Optimal Control of Power Systems
11:45 - 12:00	Chair's comments:
12:00 - 13:30	Lunch
13:30 - 14:15	Tutorial: Ron Parr - Speaking the language of artificial intelligence
14:15 - 15:00	Tutorial: Warren Powell - Modeling dynamic programs: an operations research perspective
	Session: Risk and robustness Chair: Shie Mannor
15:00 - 15:25	John Moody: Direct Reinforcement with Recurrence: Application to Games and Finance
15:25 - 15:50	Andrzej Ruszczyński: Risk-Averse Optimization
15:50 - 16:15	Jennie Si: Gradient algorithms, Robustness, and Partial Observability
16:15 - 16:30	Chair's comments:
16:30 - 17:00	Break
	Session: Managing complexity Chair: Litsa Micheli-Tzanakou
17:00 - 17:25	Daniela de Farias: Linear Programming for Approximate Dynamic Programming: New Formulations and Directions
17:25 - 17:50	George Lendaris: Autonomous Experience Based Identification and Control via Higher Level Learning Algorithms and Neural Networks, based on Context Discernment
17:50 - 18:15	Enrique Sucar: Approximate Solution of Complex Decision Problems based on Abstraction and Decomposition
18:15 - 18:30	Chair's comments:
18:30 - 19:30	Drinks
19:30 - 21:30	Dinner

2006 NSF WORKSHOP AND OUTREACH TUTORIALS ON APPROXIMATE
DYNAMIC PROGRAMMING

Time	Thursday April 6
9:00 - 10:30	Panel: What have we learned? Panelists to be named during workshop. Panel will focus on drawing lessons from the workshop.

Travel
