

# Administrative

## Reinforcement Learning and Planning under Uncertainty

ANU COMP6460/4640, Sem 2, 2008



Scott Sanner, Marcus Hutter

NICTA / RSISE

First.Last@nicta.com.au



# Course Goals

- 1) To realize that the main task of AI – *making optimal sequential decisions in an uncertain world* – can be formalized!
- 2) To understand formal models for decision-making under uncertainty and their properties
  - Unknown models (reinforcement learning)
  - Known models (planning under uncertainty)
- 3) To understand how to efficiently implement solution algorithms for some of these models

# Why *you* should Enrol / Audit

- **You're working or interested in graduate work in AI**
  - These models unify 50+ years of AI research
  - Core material you should know *and* cite
    - Don't reinvent the wheel (or decision-theory, MDP, etc.)
  - Broaden your view of your own topic area
- **You're interested in AI for fun (and profit)**
  - Perfect, Marcus + Scott = fun
  - (We'd like to share your profits)
- **You need credits to graduate**
  - Great, hopefully along the way we'll be able to impart our excitement for this work and pique your curiosity

# Enrolment

- **Enrolment for credit**
  - Email Scott if you need a permission code
- **Auditors are welcome**
  - Please realize material is cumulative
  - Hard to keep up if miss too many classes
  - Encourage you to attend tutorial / lab
    - Best way to understand is by solving problems

# Assessment

- **Grade breakdown**

- 3 written assignments (10% each)
- Take-home final exam (70%)

- **Important dates**

- Aug. 27, 23:59: [Assignment 1](#) due (email to Scott)
- Sep. 24, 23:59: [Assignment 2](#) due (email to Marcus)
- Oct. 29, 23:59: [Assignment 3](#) due (email to Scott)

- **Take-home final exam**

- After Oct. 29, TBD

Administrative:  
Any Questions?