

# **Course Summary: Dimensions of RL & Planning**

**Reinforcement Learning and  
Planning under Uncertainty**

ANU COMP6460/4640, Sem 2, 2008

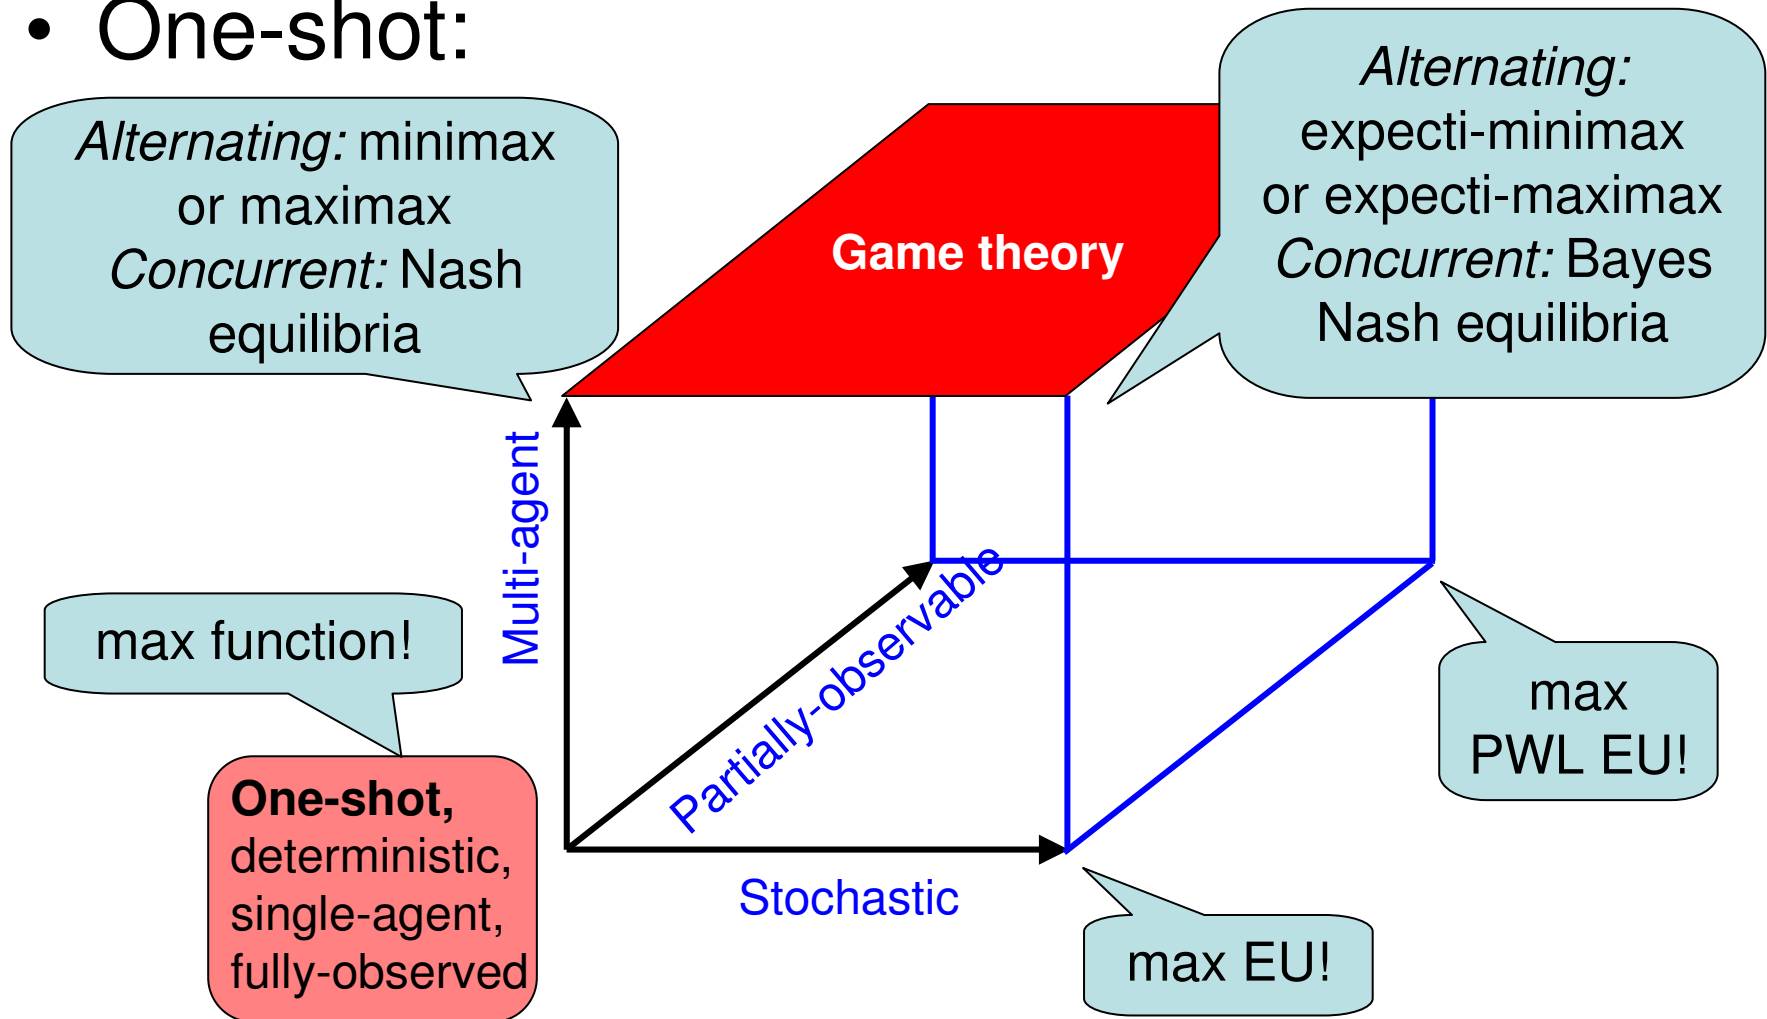
**Scott Sanner**

**NICTA / RISE**

**First.Last@nicta.com.au**

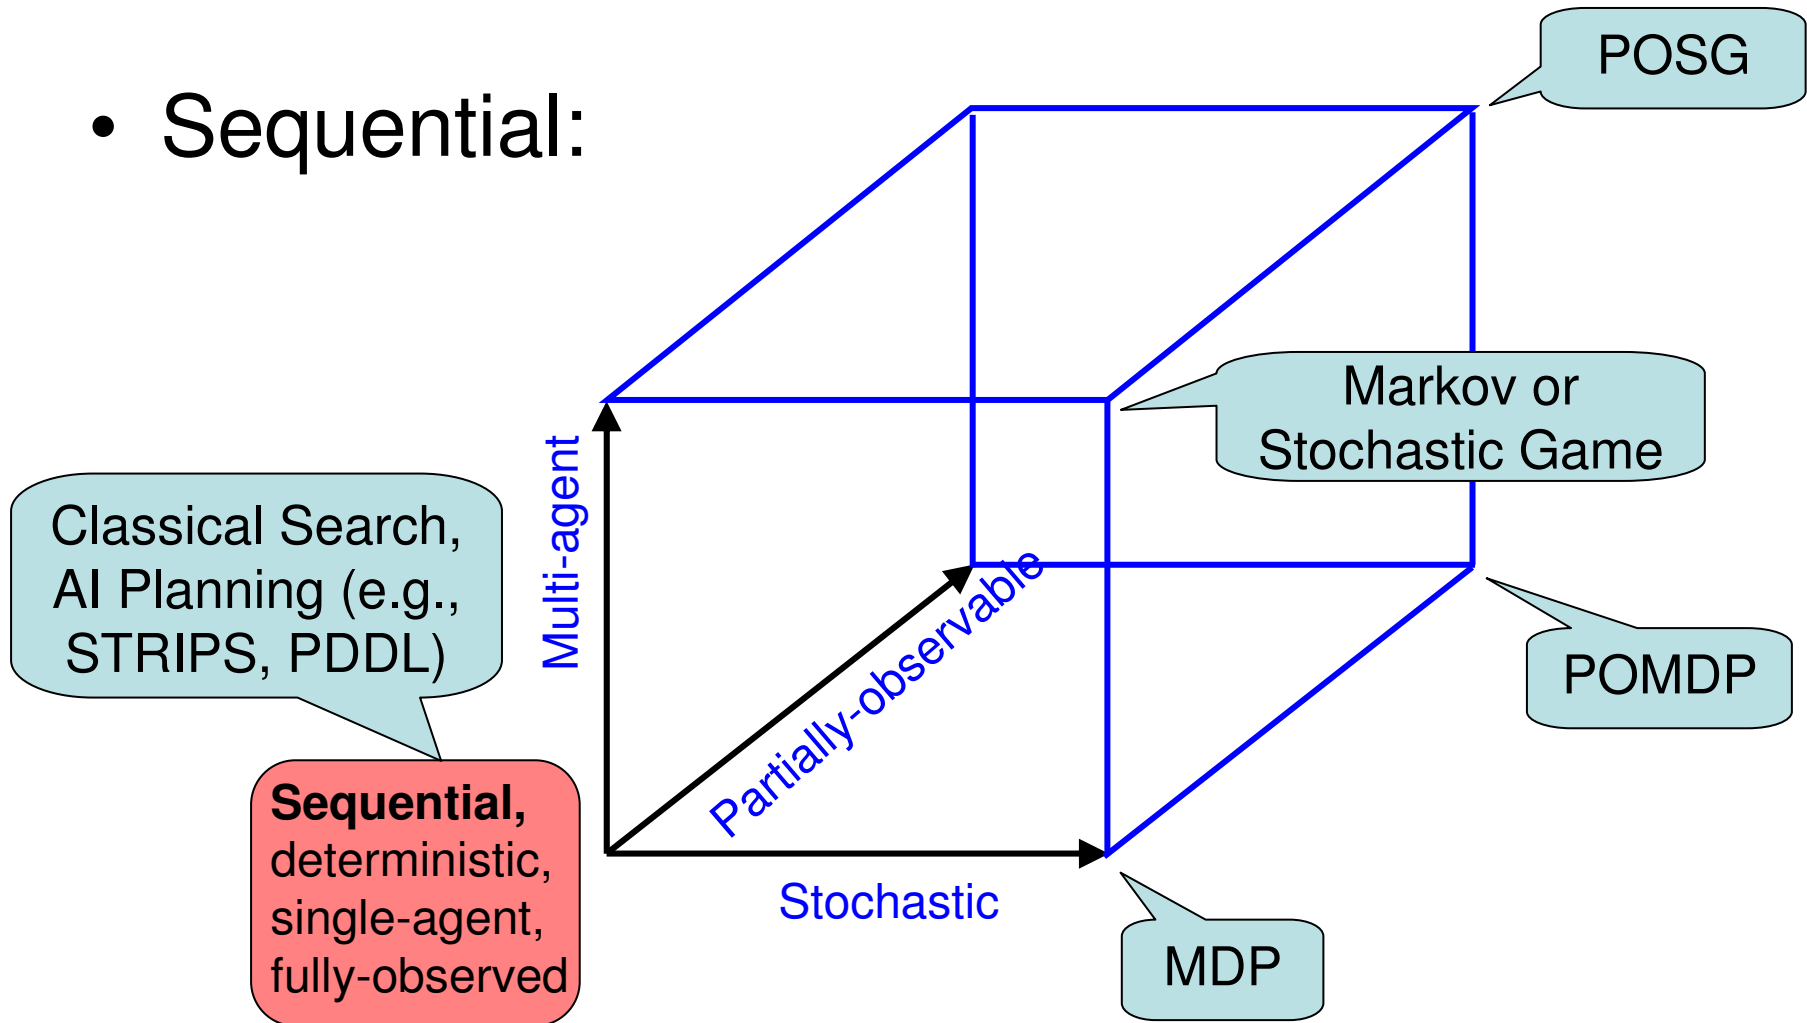
# Dimensions of RL & Planning

- One-shot:



# Dimensions of RL & Planning

- Sequential:



# More Dimensions

- State and action space
  - Continuous?
- Structure
  - Factored (i.e., state & action variables)
  - Relational?
- Actions
  - Multiple time steps per action? (Semi-MDP)
  - Time in state?
- Policy constraints (or hints as partial programs)

Remember: Just sample  
if model-free RL case.

Can save experience if  
learn model as well.

# Recall Course Goals

- 1) To realize that the main task of AI – *making optimal sequential decisions in an uncertain world* – can be formalized!

Can you formalize Intro lecture examples now?

- 2) To understand formal models for decision-making under uncertainty and their properties
  - Unknown models (reinforcement learning)
  - Known models (planning under uncertainty)

Optimal deterministic policies? Exploration required?

- 3) To understand how to efficiently implement solution algorithms for some of these models

Async. methods. TD & Fn Approx. Factored methods.

# Concluding Remarks

- Hopefully you know a lot more about how to formalize and solve sequential decision making problems than when you started...
- This course is just the tip of the iceberg
  - But a large chunk of the tip
- If you never formalize another (PO)(...)MDP
  - Hopefully this course changes the way you view decision-making in AI and everyday life
  - Many of the tools (e.g., factored probability models, decision diagrams) indispensable for non-decision making problems

