Course Overview

CMPUT 654:
Modelling Human Strategic Behaviour

Strategic Modelling

This course is about **modelling human strategic behaviour**:

- **Modelling:** Constructing formal, predictive models of action
- **Strategic:** Outcomes that an agent cares about depend on:
    1. Agent's **own** actions
    2. Actions of **other** agents, with **independent** goals and priorities
- **Human:** Primarily concerned with modelling behaviour by **people**, not by algorithms (e.g., border gateway protocol)
    - Actual, **empirical** behaviour, not **ideal** behaviour

Part 1: Game Theory

- **Mathematical** framework for modelling interactions between **rational agents**
- **Format:**
  - First six weeks
  - Lecture format
  - Two assignments

Part 2:
Behavioural Game Theory

- **Inductive** models, not just implications of assumptions
- Models are typically cognitively inspired
- Less conceptually unified than standard game theory
- **Format**
  - Second four weeks
  - Student presentations of readings
  - Summaries of readings

Part 3: Research Survey

- Survey of literature of sub-area we did not cover in class
  - Could be an **application** area, **subset** of an area we covered
  - Ideally: Propose direction for **new research** (especially if you are considering working with me)
  - Novel research results **NOT REQUIRED** for full marks
- Presentations in final three weeks

Prerequisites

- Prior knowledge of game theory is **NOT REQUIRED**
- Need to be able to follow/construct formal **proofs** and **mathematical arguments**
- Basic knowledge of **probability** (random variables, expectations, conditional probability, Bayes' rule)

Lecture Outline

1. Overview
2. Course Topics
3. Logistics

Utility Theory:
Reward Hypothesis

Reward hypothesis [Sutton & Barto 2018]:
That all of what we mean by goals and purposes can be well thought of as the maximization of the **expected value** of the cumulative sum of a received **scalar signal** (called reward).

1. Why should we believe that an agent's preferences can be adequately represented by a **single number**?

2. Why should agents maximize **expected value** rather than some other criterion?

Utility Theory: Representation Theorem

- Utility theory deals with **preference relations** $\succeq$ over final outcomes $o \in O$
    - i.e.. $a \succeq b$ means "a is (weakly) preferred to b"
- von Neuman & Morgenstern's **representation theorem** says that if a preference relation $\succeq$ satisfies certain axioms, then there exists a utility function $u : O \to \mathbb{R}$ such that:

1. $o_1 \succeq o_2 \iff u(o_1) \geq u(o_2)$, and
2. $u([p_1 : o_1, \ldots, p_k : o_k]) = \sum_{i=1}^{k} p_i u(o_i) = \mathbb{E}[u(o)]$

Game Theory: Normal Form Games

- In a multiagent setting, what are the consequences of assuming that agents are **expected utility maximizers**?
- Normal form games:
  - Each agent picks an action simultaneously
  - **Profile** of utilities specified for each profile of actions
- **Question:** What *strategy* maximizes utility for the row agent?
  - **Solution concepts**: Outcomes that are consistent with the expected-utility maximization assumption

	L	R
T	4, 3	0, 0
B	1, -1	2, 8

Game Theory: Special Cases

- **Repeated games**: What happens when the same game is played between the **same agents multiple times**?
- **Extensive form games**: Explicitly represent **sequential action**
- **Bayesian games**: Explicitly represent **private information**

Game Theory:
Social Choice & Mechanism Design

- **Social choice:** Combining the preferences of multiple agents
- **Mechanism design:** "Game theory in reverse"
  - Design the **game itself** such that expected utility maximizers will reach the **socially optimal outcome**
  - ... even if you don't know their utilities
  - *Example:* allocating a valuable item

Behavioural Game Theory

- People aren't actually expected utility maximizers!
- **Behavioural game theory:** Accurate models of **human behaviour** in game theoretic settings
  - Demonstrate failures of standard game theory
  - Relaxing assumptions: expected utility maximization, common knowledge
  - Heuristic rules for interactions
  - Cognitive bounds

Survey Topics Examples

The ideal project is a **proposal** for novel work and a survey of the relevant **related work**.

1. Predictive Models

- Feedback and Dynamic Behaviour
- Interpretability
- Nonstrategic Factors in Behaviour
- Discovering Predictive Features

2. Agent Design

- Game Play
- Optimal Behaviour Discovery / Learning
- Behavioural Finance

3. Mechanism Design

- Spectrum Auctions
- Ad Auctions
- Peer Grading Platforms
- Misinformation in Social Networks
- Topic Selection in Election Coverage

Course Essentials

jrwright.info/bgtcourse/

- This is the **main source** for information about the class
- Slides, readings, deadlines

eClass

- This is where **assignments** are posted and handed in
- There is also a **class forum** for questions and discussions about course material

Contacting Me

- **Discussion board:** eClass forums
  for **public** questions about assignments, lecture material, etc.
- **Email:** james.wright@ualberta.ca
  for **private** questions (health problems, inquiries about grades)
- **Office hours:** After every lecture, or by appointment

Evaluation

- Assignments: 30%
- Reading presentation: 15%
- Reading summaries: 15%
- Research survey
  - Outline: 5%
  - Presentation: 15%
  - Writeup: 20%

Missed / Late Assignments

Late assignments

- Can be handed in up to 2 days after deadline
- Flat 20% penalty
- On time or late but please **not both**

Missed assignments

- Weight of assignments missed due to excused absences will be proportionally distributed among other assessments

Assignments

There will be **two** assignments (not necessarily weighted equally)

You are **encouraged to discuss** assignment questions with other students:

1. You **may not** share or look at each other's **written work**
2. You must **write up** your solutions individually
3. You must **list** everyone you talked with about the assignment.

Academic Conduct

- Submitting someone else's work as your own is **plagiarism**.
- So is helping someone else to submit your work as their own.
- I report **all cases** of academic misconduct to the university.
- The university takes academic misconduct **very seriously**.
  Possible consequences:
  - Zero on the assignment (virtually guaranteed)
  - Zero for the course
  - Permanent notation on transcript
  - Suspension or expulsion from the university

Readings

For Part 1 (Game theory)

- Yoav Shoham and Kevin Leyton-Brown, Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations

For Part 2 (Behavioural game theory):

- Original papers from the literature

For Part 3 (Research surveys):

- Self-directed readings from the literature
  - But feel free to ask me for pointers!

Enrollment

How many people present today are:

- Enrolled?
- Auditing with the hope of enrolling?
- Auditing without intending to enrol?

ABGT Reading Group

What: Topics related to algorithmic and behavioural game theory

When: Mondays at 3:15pm - 4:40pm
Where: ATH 3-32

Next meeting: January 15, 2024

Webpage: jrwright.info/abgt.html

Announcements: abgt slack channel (see website for link)

AI Seminar

What: Great talks on cutting-edge AI research
External (e.g., other universities, DeepMind, local startups) and internal speakers

When: Fridays at noon

Website: sites.google.com/ualberta.ca/ai-seminar/

Announcements: Sign up for mailing list (bottom of webpage)

Summary

- **Course webpage:** jrwright.info/bgtcourse/
- Data-driven behavioural modelling using lens of **game theory**
- Grading:
  - Two assignments
  - One reading presentation
  - Research survey
- Reading group: jrwright.info/abgt.html