

PHI1CTO - Critical Thinking

JAKE CHANDLER



1 Overview

Psychological studies have repeatedly shown that our thinking can be biased, distorted, and misinformed. This subject introduces you to some of the most common mistakes that people make in reasoning and equips you with essential tools to recognise and sidestep these errors. Among other things, you'll master the use of argument diagrams and basic logical principles to dissect and assess arguments as they occur in everyday life.

We explore both deductive and probabilistic reasoning, covering key concepts in formal logic alongside common cognitive biases that affect human judgment. In the final weeks, we examine how modern technology impacts our reasoning, with special attention to Large Language Models (LLMs) like ChatGPT and Claude. You'll gain insight into how these AI systems work, their reasoning capabilities compared to humans, and how to effectively interact with them.

Throughout this subject, you will improve your ability to think with clarity and autonomy, and to participate constructively in discussions and debates. Though critical to the practise of philosophy, these competencies are also vital across all disciplines in academia and are highly sought after in numerous professional fields.

The La Trobe Handbook entry for the subject can be found [here](#).

2 Student Testimonials

"The subject itself was very interesting and helped us question our day to day thinking processes. The teacher was also very helpful and made sure that we did not proceed to any part of the subject without providing a detailed explanation of the weekly content. Jake's teaching style was also very engaging and I looked forward to attending each class."

"I found this subject to be intellectually stimulating and participating in it has been a profoundly rewarding experience for me. This subject gave me hope that there are tools available to me to make sense of this complicated world - and also made me feel like I am not helpless, hopeless or stupid when it comes to making sense of things."

"I believe that the power point slides were extremely well constructed. Jake did an amazing job synthesising the information so that it was clear and easy to understand. The videos were kept short and engaging."

"Great lectures, explained very clearly, and just enough content to learn an incredible amount, but not enough to become overwhelmed. Tutorial sessions were huge for this class."

"The critical thinking skills learnt can be applied in many areas. The introduction to AI was very interesting and something that will be useful in the near future as AI becomes more prevalent."

“The lectures were very informative and the live tutorials were invaluable for consolidating my learning and maintaining focus as an online student.”

3 Assessment and teaching activities

Assessment:

- Two short essays (40%)
- Weekly quizzes (30%)
- Argument diagram (20%)
- Early assessment quiz (10%)

Teaching Activities:

- Weekly pre-recorded video lectures (0.5 hours)
- Weekly tutorials, face-to-face or online (1.5 hours)

4 Weekly breakdown

Week 1: Arguments: Part 1

In the first two weeks, we'll introduce the basics of critical thinking. We'll examine arguments, why they are important, why you should care about them, and how to identify premises and conclusions. The first two weeks will lay the foundation for the rest of the subject, so it's really important that you carefully do the readings and make sure you understand the key concepts and techniques.

In Week 1, we introduce some extremely basic concepts and distinctions, including the concepts of a reason, an argument, premises, conclusions, and joint versus independent reasons.

- Slides: [Week 1 Slides](#)
- *Lecture recordings are available on request.*

Week 2: Arguments: Part 2

This week, we first look at a number of different ways in which an argument could be objected to. We then learn how to present arguments and counterarguments in a clear and effective way, using 'standard form' as well as argument diagrams. Finally, we review some common difficulties that can be faced when trying to reconstruct an argument found 'in the wild' (e.g., in an article, a podcast or a speech).

- Slides: [Week 2 Slides](#)
- *Lecture recordings are available on request.*

Week 3: Deductive Reasoning: Part 1

After having seen how to reconstruct and present arguments, we now take a closer look at how to assess whether or not they are any good and, in particular, the extent to which the premises can be said to support the conclusion.

In Weeks 3–5 we'll look at deductive reasoning, which is one of the most important forms of reasoning used in all kinds of contexts, including law, politics, media, science, engineering, and of course, philosophy. In Week 3, we notably introduce the notion of a deductively valid argument, an argument such that the truth of its premises would guarantee the truth of its conclusion. We also discuss the idea that some arguments are valid merely because of their 'form', distinguishing two levels of description of argument form: sentential and subsentential.

- Slides: [Week 3 Slides](#)
- *Lecture recordings are available on request.*

Week 4: Deductive Reasoning: Part 2

This week, we continue our discussion of formal validity. We learn how to use special notation (the 'language of sentential logic') to express the sentential form of a sentence or argument and use it to present a number of well-known valid sentential argument forms. We then briefly look at subsentential form and the notation (the 'language of predicate logic') that is used to express it.

- Slides: [Week 4 Slides](#)
- *Lecture recordings are available on request.*

Week 5: Deductive Reasoning: Part 3

We've seen so far that one important way in which an argument can be good is by being deductively valid. We've also seen that deductive validity is often due to the mere form of the argument. But how do we recognise when a form is deductively valid or not?

This week, we learn an easy method, the method of truth tables, to help us determine whether any given sentential argument form is valid or not. We also use this method to prove the invalidity of three common 'deductive fallacies', before looking at a possible explanation of why these problematic patterns of reasoning might be attractive to some people.

- Slides: [Week 5 Slides](#)
- *Lecture recordings are available on request.*

Week 6: Probabilistic Reasoning: Part 1

In previous weeks, we have looked at one way in which the premises of an argument can be said to support its conclusion by being such that their truth would guarantee the truth of the conclusion. But in many apparently reasonable arguments, the relation between premises and conclusion falls short of this high standard; the truth of the premises would merely render the truth of the conclusion probable. This week we look at some basic concepts of probability, which are introduced informally, in a visual manner. We notably learn how to calculate the probability that a given hypothesis is true, given the evidence that has been gathered.

- Slides: [Week 6 Slides](#)
- *Lecture recordings are available on request.*

Week 7: Probabilistic Reasoning: Part 2

This week, we put to work the concept of conditional probability that I introduced last week to define the notion of an 'inductively valid' argument. After doing this, I show how inductively valid arguments have some important features that deductively valid arguments lack; they come in different strengths and are 'defeasible'.

We then start on a new topic, common errors in probabilistic reasoning. While assessing probabilities is a crucial skill in assessing arguments, it turns out that we are not particularly good at it. We make systematic mistakes that are generally attributed to our using mental 'heuristics'; cognitive rules of thumb that are easy to use but can lead us astray. This week, we take a look at one such common error, the Prosecutor's fallacy.

- Slides: [Week 7 Slides](#)
- *Lecture recordings are available on request.*

Week 8: Probabilistic Reasoning: Part 3

This week, we continue our discussion of so-called 'probabilistic fallacies', which have captivated the attention of psychologists, economists, marketing specialists and philosophers since the 1970s. We discuss four further fallacies and attempt to explain their causes: the Conjunction fallacy, the Gambler's fallacy, Availability fallacies and Anchoring fallacies.

- Slides: [Week 8 Slides](#)
- *Lecture recordings are available on request.*

Week 9: Confirmation Bias and Overconfidence

After having reviewed some of the most notorious ways in which our probabilistic reasoning is led astray, now take a look at a problematic type of behaviour that affects both deductive and probabilistic reasoning, confirmation bias. This bias involves a tendency to selectively disregard arguments that undermine our prior beliefs. It has also been linked to another well-known phenomenon, overconfidence. After looking at some striking experiments that expose both biases, we consider possible evolutionary explanations for their existence and strategies that we might follow in order to avoid them.

- Slides: [Week 9 Slides](#)
- *Lecture recordings are available on request.*

Week 10: Search Engines and Bias

After having spent a number of weeks on how to analyse and evaluate arguments for or against a certain claim, we now turn to the issue of how we come across such arguments in the first place. This week, we discuss the use of search engines in finding information online and the various forms of bias that can result from it. Among other things, we take a look at two well-known factors thought to influence the ranking of webpages returned by Google Search; PageRank and personalisation. Personalisation, in particular, has been associated with popular fears that our reliance on Google Search is creating 'filter bubbles' that amplify our existing confirmation biases and contribute to a more polarised society. We consider the extent to which these fears are well founded.

- Slides: [Week 10 Slides](#)
- *Lecture recordings are available on request.*

Week 11: Large Language Models: Part 1

This week introduces Large Language Models (LLMs) like ChatGPT and Claude, explaining the basic principles that underly their design and use. The lecture covers the basic architecture of LLMs, including embedding (encoding the meaning of user inputs into vectors), transformation (refining the vector representations by taking into account linguistic context and training data), and unembedding (using the transformed vectors to generate new words). We also take a look at the way in which LLMs are trained, using massive datasets, and stress the importance of effective prompting in order to obtain the best results. Finally, we touch on ongoing attempts to reverse engineer these complex models so as to understand exactly how they manage to perform the tasks we assign to them. Overall, this introduction aims to demystify these powerful AI tools, showing how they process language and generate human-like responses. It's an exciting field that's rapidly evolving, with new developments taking place constantly.

- Slides: [Week 11 Slides](#)
- *Lecture recordings are available on request.*

Week 12: Large Language Models: Part 2

Building on our introduction to Large Language Models from Week 11, this week we examine how LLMs perform on reasoning tasks compared to humans. We explore fascinating research showing that LLMs often exhibit the same reasoning biases and shortcomings as humans, such as struggling with the conjunction fallacy and cognitive reflection test items. We discuss how newer LLM versions like GPT-4 show improved reasoning capabilities, and how techniques like Chain of Thought (CoT) prompting can help overcome these limitations by encouraging step-by-step reasoning processes. The lecture addresses the parallels between human dual-process thinking (System 1 vs. System 2) and LLM reasoning, with implications for how we might better use these AI tools.

- Slides: [Large Language Models Pt 2](#)