LLM-Driven Applications with R and Python

Level: Intermediate
Duration: 6 hours

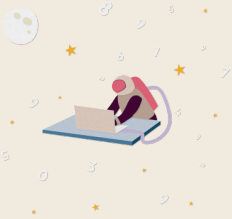

Learn how to work with large language models (LLMs) using R and Python. This course will start with basic concepts like sending user prompts and receiving a structured output, before moving onto move advanced topics like building LLM-powered web applications and configuring a knowledge store for retrieval-augmented generation (RAG). Throughout, we will emphasise important considerations for security, safety and responsible use of AI.

Course Outline

- **Generative AI basics**: A conceptual introduction to LLMs including tokens, pricing, strengths and weaknesses of LLMs, and a comparison of the most popular LLM providers.
- **Prompts**: Using R and Python to send prompts to an LLM and receiving a structured output, as well as good practices for writing user and system prompts.
- **LLM-powered Applications**: Building web applications that harness the power and flexibility of LLMs, including inserting a chatbox in a Shiny app and user-friendly data exploration with LLMs.
- **Retrieval-augmented Generation**: An introduction to RAG workflows followed by a hands-on demo in building a RAG knowledge store from scratch using web-based documents.
- **Security, Privacy & Responsible AI**: Exploring the main risks posed by LLM-powered applications including hallucination, prompt injection and data poisoning, as well as some techniques for mitigation.

Learning Outcomes

Session 1:

By the end of session 1 participants will...

- understand how LLMs work at a conceptual level.
- be familiar with the strengths and weaknesses of LLMs.
- know how to set up a connection with an LLM in R or Python.
- be able to write good user and system prompts.
- know how to format prompts to receive a structured output.
- have learned how to include images and PDFs in the user prompt.

Session 2:

By the end of session 2 participants will...

- be able to insert an LLM chatbox in a web application.
- understand how to use LLMs for intuitive data exploration.
- understand the steps involved in building a RAG knowledge store.
- be able to recognise the key risks when building an LLM-powered application and how to mitigate against these.
- have familiarity with the most popular web platforms for developing and maintaining LLM workflows.

This course does not include:

- AI-powered copilots for code generation.
- Introduction to programming with R and Python. See our **Introduction to R** and **Introduction to Python** courses.
- Building Shiny web applications from the ground up. See our **Introduction to Shiny** course for a primer.

Prior Knowledge

Participants must have a basic knowledge of either R or Python. No prior experience of working with LLMs is required for the course. A basic understanding of building web applications using Shiny (or a similar framework) is useful but not essential.

Attendee Feedback

- "Myles was very informative and friendly. I particularly liked the well-prepared hands-on demos and exercises to consolidate our learning straight away!"
- "I thought it struck the balance between conceptual and practical really well and certainly has given me a lot of confidence and de-mystified elements like RAG which I had heard of but never had the chance or confidence to experiment with."
- "I liked how it was structured. The theory was easy to understand."

Contact

hello@jumpingrivers.com