Lab: AI Fundamentals for Developers

- Launch the lab
- What you'll learn
- Modules

Get hands-on with the four core pillars of AI application development: models, prompt engineering, tool calling, and RAG. This lab runs entirely on your machine using Docker Model Runner — no API key or cloud account required.

Launch the lab

{{< labspace-launch image="dockersamples/labspace-ai-fundamentals" model-download="true" >}}

What you'll learn

By the end of this Labspace, you will have completed the following:

- Understand the Chat Completions API and how to structure messages for a model
- Use prompt engineering techniques including system prompts, few-shot examples, and structured output
- Implement tool calling and the agentic loop in code
- Build a RAG pipeline that grounds model responses in your own data

Modules

#	Module	Description

1	Welcome & Setup	Introduction to the lab and verifying your environment
2	Talking to Models	Chat Completions API, message roles, and stateless model behavior
3	Prompt Engineering	System prompts, few-shot examples, and structured output
4	Tool Calling	Tool definitions, the agentic loop, and executing tools in code
5	Retrieval Augmented Generation (RAG)	Retrieve, augment, and generate with your own knowledge base
6	Wrap-up	Summary of concepts and next steps

□□□ #2
□ □□□ □□□ 2026-10-31 18:46:13 UTC
□ □□□ □□□ 2026-10-31 18:46:14 UTC