

Agent Networks

Understanding, Designing and Mastering Networks of LLM Agents

4 parts · 16 chapters · 83 sections · 220 pages

About this book

Objectives. Agent networks are a new and fascinating frontier of artificial intelligence, and behind their apparent magic lies a complex technology. This book gives the reader the keys to mastering it: a deep understanding of the architecture, the mechanisms and the limits of these systems, so as to use the resulting tools with discernment and to design robust, controlled agent networks. It is also a clear, properly referenced exploration of the large language models that are their cornerstone.

Audience. Written for curious minds and hands-on builders: designers, project managers, developers and “no-coders” who want to grasp the ins and outs of agent networks in order to build more resilient, reliable and efficient systems — and, more broadly, anyone who would rather understand this technological wave than simply be subjected to it.

Approach. Two principles. First, ideas are made tangible through code: every concept is illustrated in Python, and every excerpt is directly executable, so the reader can watch the mechanism in action rather than only read about it. Second, direct access to the sources of knowledge, through a rigorous selection of references to foundational publications.

Outline

Foreword p. 1

PART 1

Foundations of Agent Networks

Chapter 1 **The Evolution of AI Systems** p. 5

Brute Force: The Dawn of Machine Computation · Abstraction: Languages and Data Structures · Formalism: the machine manipulates symbols · Expert Systems: Knowledge Encapsulated · Machine Learning: Knowledge Induced from Data · Deep Learning: The Quest for a Universal Predictor · LLMs: Words Become Vectors · Agents: AI Becomes an Actor

Chapter 2 **Measuring Intelligence and the Ability to Handle Uncertainty** p. 27

The Turing test and its limits · Knowledge tests for LLMs · Intelligence as skill-acquisition efficiency · Generalisation and handling uncertainty · Evaluating an agent: beyond the generic benchmark

Knowledge, information extraction, and tools · Overview

PART 2

Exploring Generative Models and LLMs

Generative AI · MoE: ensemble approach and efficiency · Dense architecture · Training LLMs · The randomness inherent in LLMs · Plasticity · Retrieval Augmented Generation

PART 3

Understanding the Reasoning Capabilities of LLMs

The Stochastic Parrot · What LLMs Lack

Zero-shot · Few-shots · Chain-of-Thought · Zero-shot Chain-of-Thought

Self-Consistency · Tree-of-Thoughts · Graph of Thoughts · Confronting Reality: ReAct

PART 4

Implementing Agents with LLMs

Implementation with Few-Shot Prompting · Implementation with Retraining · Model with a Grammar · Invoking the Tool

Graph Theory · Loops and Iterations · Human in the Loop · Obtaining Reliable Answers · Controlling the Infrastructure · The Importance of State · The Plan-Based Approach · Conclusion

Selection Criteria · Routing · Sizing · Hosting · Optimisation

The Limits of Bare Reasoning · A Typology of the Knowledge to Be Provided · Modes of Knowledge Injection · Building and Maintaining the Knowledge Base · Acculturation as a Human Process

Chapter 12	Securing Access to Tools and Data	p. 133
Introduction · What Makes Agents Different · Sources of Danger · Identity and Rights: Authentication vs. Authorisation · Technical Solutions for Identity and Access · Security Solutions · Towards a Proactive Defensive Approach		
Chapter 13	Designing the Software Architecture	p. 149
The Model Context Protocol (MCP) · Server · Distribution · Monitoring		
Chapter 14	Deploying and Monitoring Agent Networks	p. 155
Kubernetes as a Deployment Foundation · Supervision and Monitoring · Managing Jobs and Workflows		
Chapter 15	Open Source Tools for Development	p. 169
Programming Agents · Easy-to-Use Agentic Libraries · Guaranteeing Code Quality · Towards a Coherent Toolbox · Communities and Keeping Watch · OpenHands: An Open Source Developer Agent		
Chapter 16	Case Study: An Agent Network for Designing 3D Parts	p. 187
Why Parametric CAD Lends Itself Well to Agents · Breaking the Problem Down into Roles · The Verification Tool: Running CadQuery in a Sandbox · The Four Nodes in PocketFlow · An End-to-End Example · Natural Extensions · Variant: Delegating Orchestration to the Model · Explicit or Delegated Orchestration: Which Regime to Choose · What Generalises Beyond CAD		
	LLMs and the Future of AI	p. 205