amarula solutions

Go Far, Go Together

TRAINING PROGRAMME

THREE-DAY HANDS-ON MASTER CLASS

OpenAI Master Class

Mastering OpenAI APIs and Beyond

From the text API and prompt engineering to retrieval, agents, multimedia and AI-powered developer tools.

DURATION	PARTICIPANTS	PRICE	TRAINER
3 days	**Minimum 3**	**€1,350 each**	**Patrizio Gelosi**

Build with AI, not just talk to it

A three-day master class for software engineers who want to move from using ChatGPT to building production applications with the OpenAI API and modern LLM tooling. The programme begins with the foundations of the text API, grows into prompt engineering, retrieval and agents, and finishes with multimedia APIs and practical AI-powered developer tools.

FORMAT	LABS	LEVEL	MINIMUM
Instructor-led, hands-on	**Python and Jupyter**	**Intermediate**	**3 participants**

WHO IT IS FOR

- Software engineers and technical leads building AI features
- Developers moving from ChatGPT usage to API integration
- Architects evaluating LLM providers, endpoints and frameworks
- Teams adopting AI in their software development workflow

PREREQUISITES

- Some experience using ChatGPT or another LLM assistant
- Basic Python programming
- Familiarity with REST APIs and JSON
- A laptop with Python 3.10+ and internet access for the labs

WHAT YOU WILL LEARN

- ✓ Integrate OpenAI's text, vision, audio and image APIs into Python applications.
- ✓ Choose the right endpoint, model and parameters for quality, latency and cost.
- ✓ Design prompts and structured outputs that are reliable enough for production.
- ✓ Build retrieval-augmented and tool-using assistants.
- ✓ Build single- and multi-agent systems with the OpenAI Agents SDK and provider-agnostic frameworks.
- ✓ Apply AI to code generation, documentation, debugging and code review.

> **No machine-learning background is required.** The course is about building with LLM APIs: the concepts are introduced from first principles and reinforced with working code.

The three-day programme

Each day combines short lectures, live API demonstrations and hands-on Python notebooks. The course moves from foundations to agents, then finishes with practical developer applications.

DAY 1

Foundations and the text API

- Module 1 - AI: from origins to OpenAI's breakthrough
- Module 2 - Base interaction endpoints
- Module 3 - Advanced text API techniques

DAY 2

Prompting and agents

- Module 4 - Prompt engineering in API interactions
- Module 5 - Agentic AI
- Module 6 - Provider-agnostic LLM frameworks

DAY 3

Multimedia and developer applications

- Module 7 - Multimedia API
- Module 8 - Applications for software development
- Module 9 - Applications for code review

> Selected slides from the course deck are reproduced on the module pages that follow. The complete deck and the lab notebooks are used in class.

HOW THE DAYS RUN

LEARN

- Short lectures and live API demonstrations
- Architecture and design trade-offs explained in context
- Provider and model comparisons grounded in real tests

BUILD

- Guided Python notebooks for every module
- Working examples you can adapt to your own use case
- Time for questions, debugging and discussion with the trainer

> *"Theory meets practice with real-world coding sessions, all brought to life in Jupyter Notebooks."*

MODULE 01 OF 09

AI: from origins to OpenAI's breakthrough

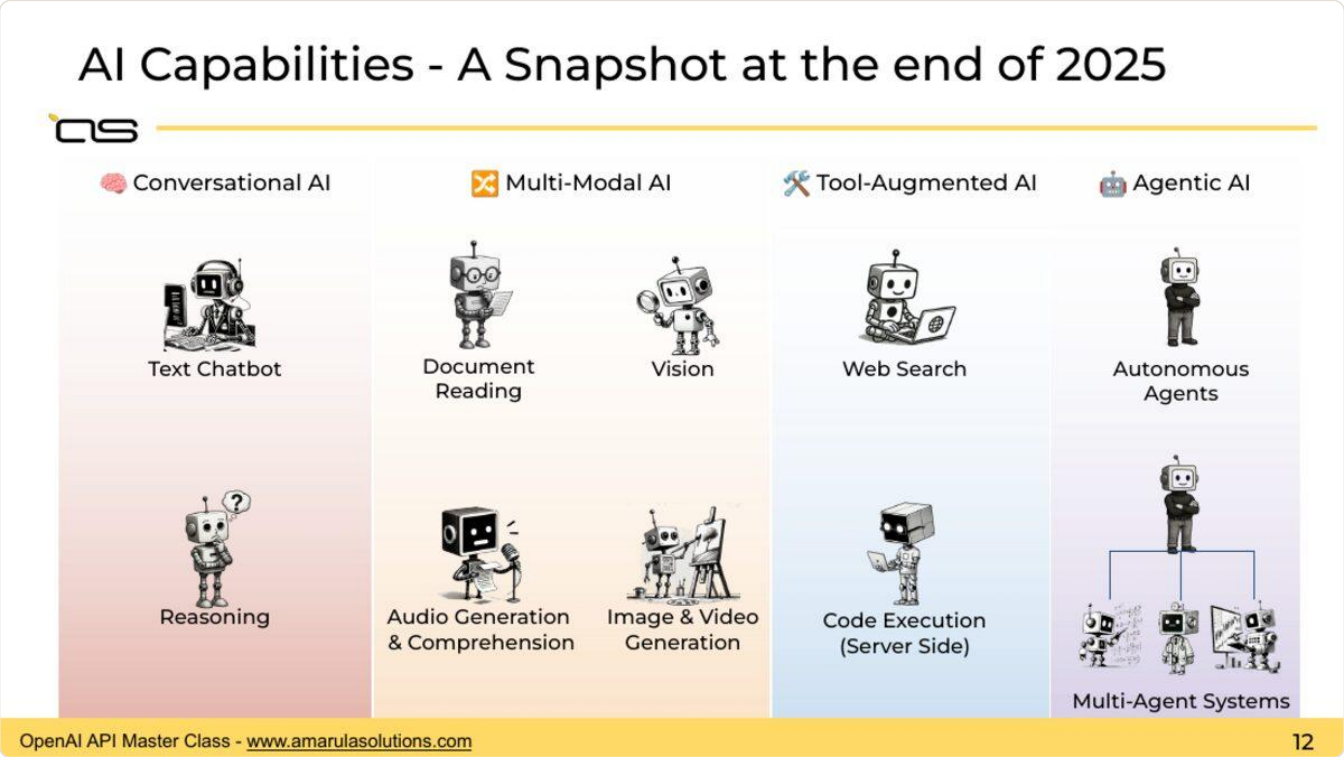

AI Capabilities - A snapshot at the end of 2025

Start with the big picture: how AI arrived at transformers and ChatGPT, where large language models fit in modern software, and how to choose a provider and model with confidence.

WHAT YOU WILL LEARN

- The evolution from early AI to transformers, embeddings and ChatGPT
- User-centric versus outcome-centric AI adoption, and where to start
- How to compare LLM providers on quality, cost, latency and privacy
- OpenAI models, SDKs, pricing and token limits

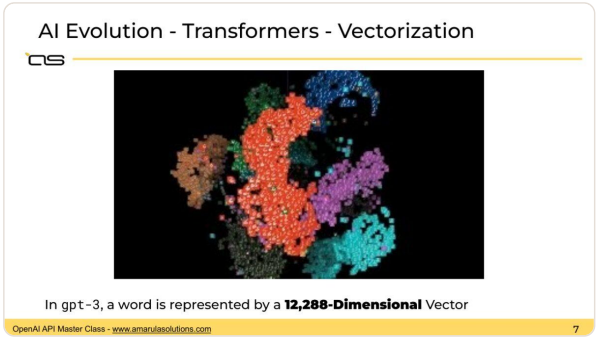

AI Evolution - Transformers - Vectorization

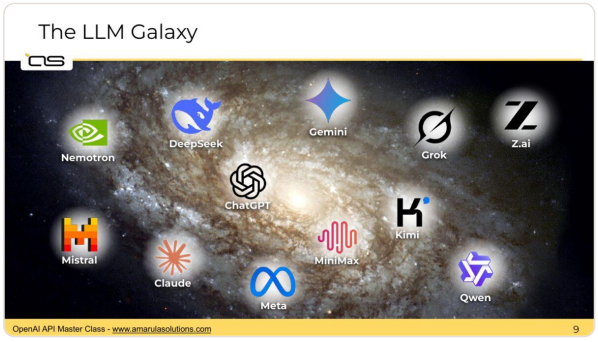

The LLM Galaxy

MODULE 02 OF 09

Base interaction endpoints

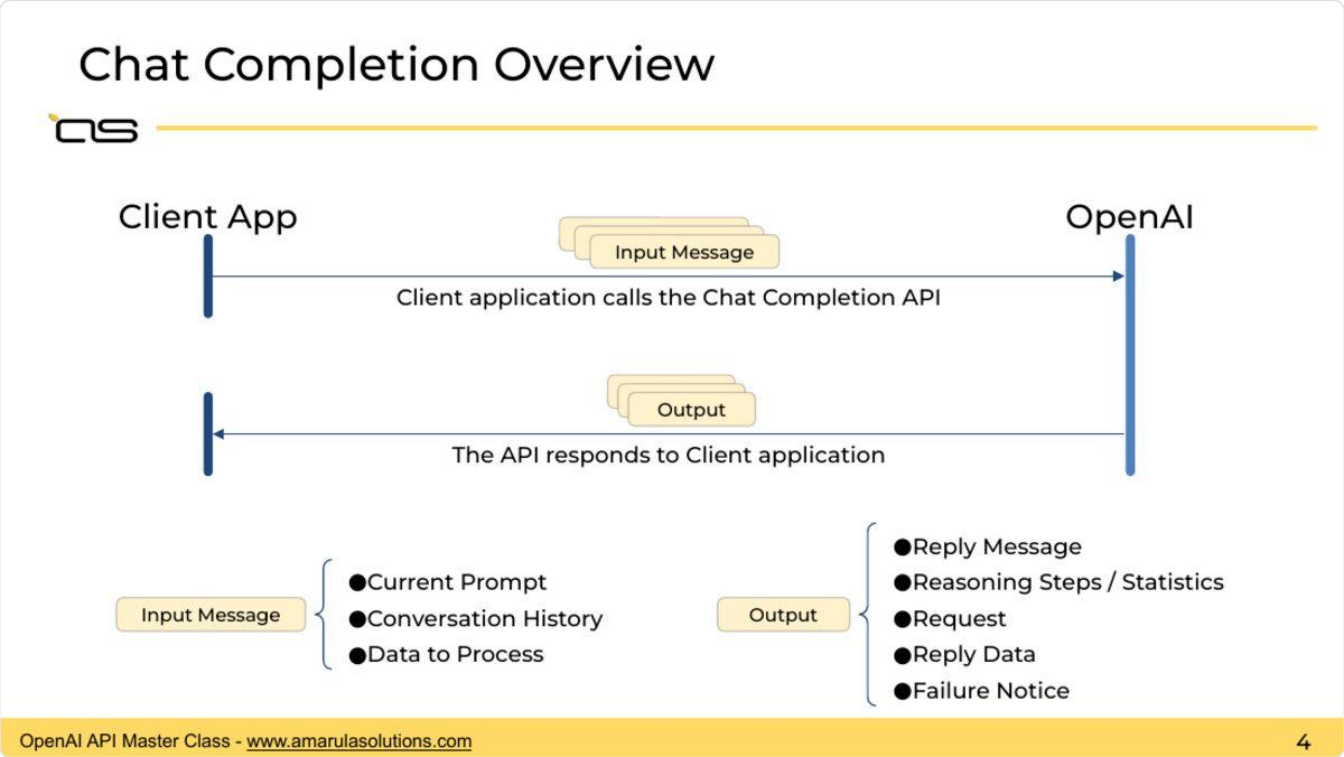

Chat Completion Overview

Learn the three core ways to talk to OpenAI models and when to use each: stateless Chat Completions, lightweight Responses and persistent Conversations.

WHAT YOU WILL LEARN

- Chat Completions: messages, roles, parameters and multi-turn history
- Responses: instructions, response chaining and temporary requests
- Conversations: server-managed state for longer interactions
- Model selection: reasoning versus GPT models and endpoint trade-offs

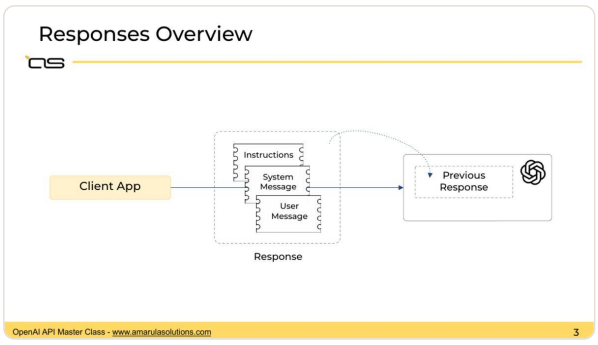

Responses Overview

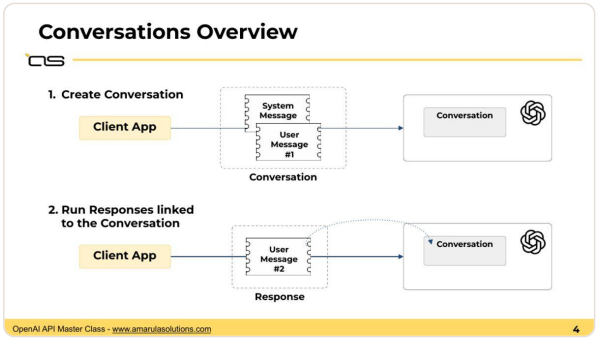

Conversations Overview

MODULE 03 OF 09

Advanced text API techniques

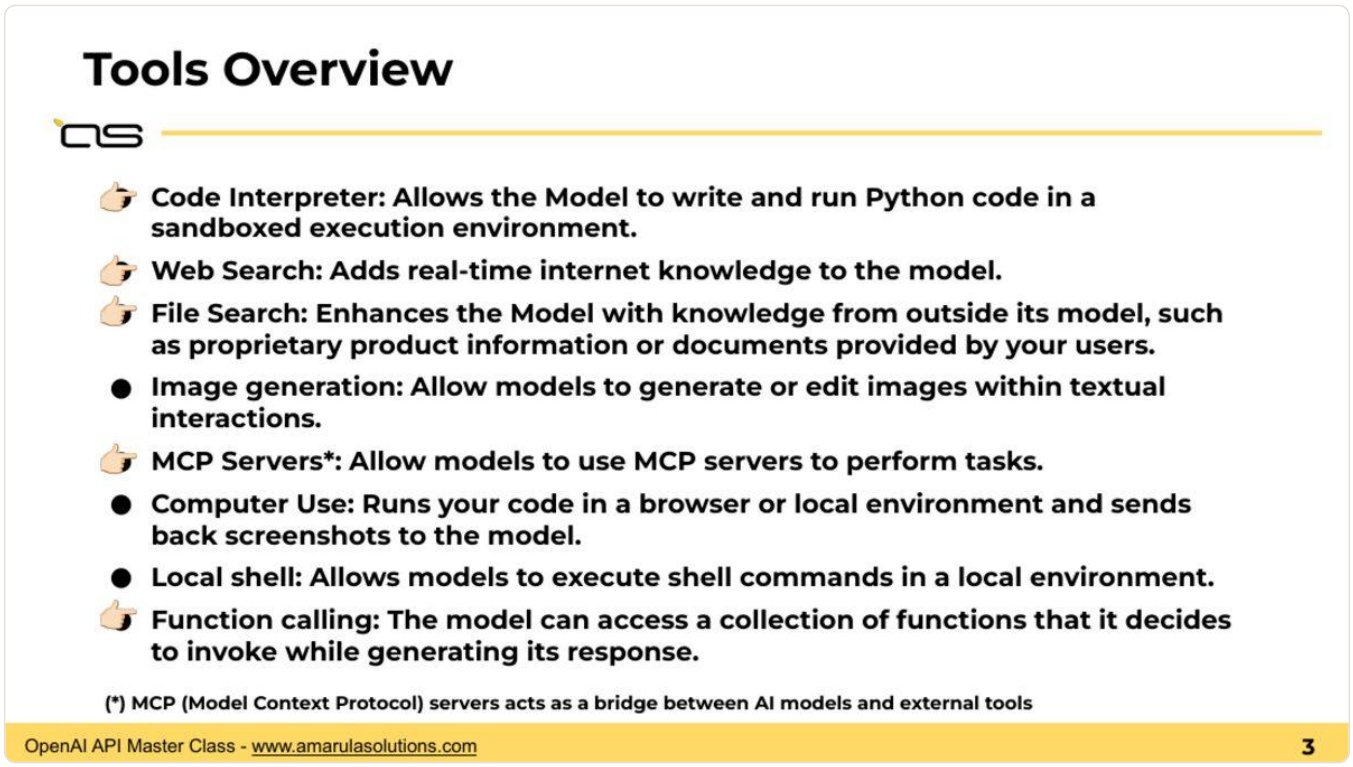

Tools Overview

Move beyond simple prompts. Attach files, call tools, return structured data and ground answers in your own documents with retrieval-augmented generation.

WHAT YOU WILL LEARN

- Built-in tools: Code Interpreter, web search, file search and MCP servers
- Function calling and structured JSON outputs
- Embeddings, vector search and retrieval strategies
- Retrieval-augmented generation (RAG): one-shot retrieval, reranking and hybrid search

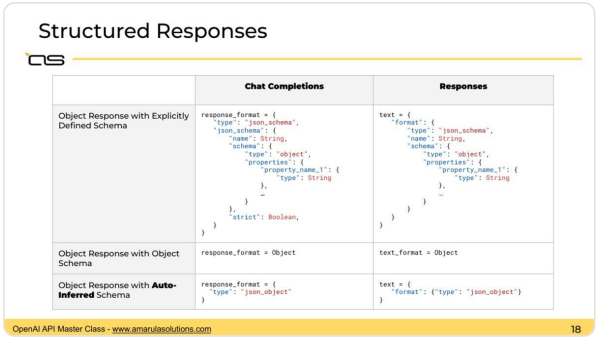

Structured Responses

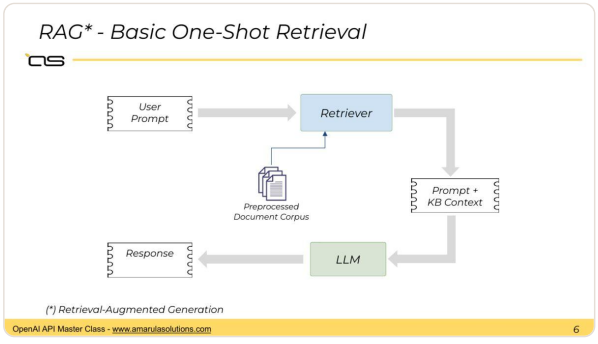

RAG - Basic One-Shot Retrieval

MODULE 04 OF 09

Prompt engineering in API interactions

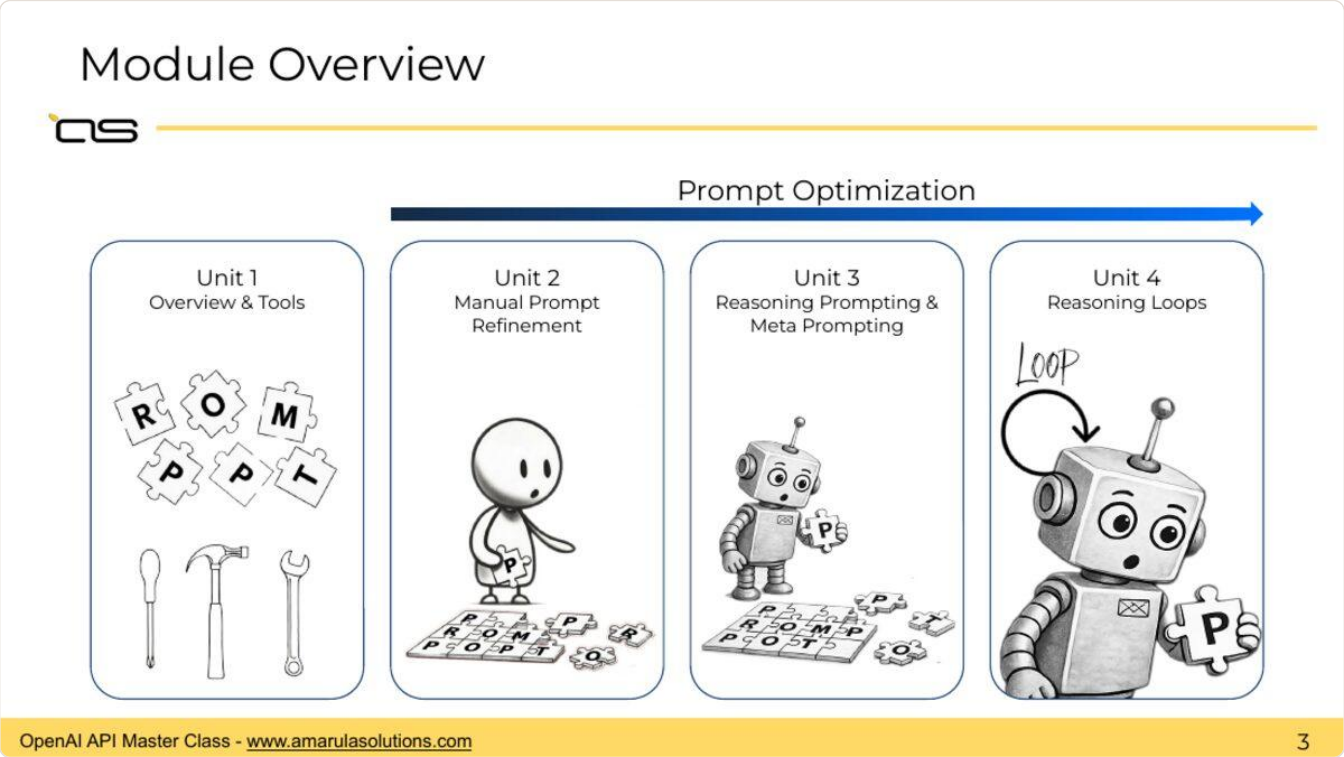

Module Overview

Turn prompting into an engineering discipline: design reusable prompts, control reasoning and improve results through structured iteration.

WHAT YOU WILL LEARN

- Prompt components: directives, context, examples and placeholders
- Manual and stepwise prompt refinement
- Chain of Thought, Tree of Thoughts and meta-prompting
- Iterative prompting, response length and variability control

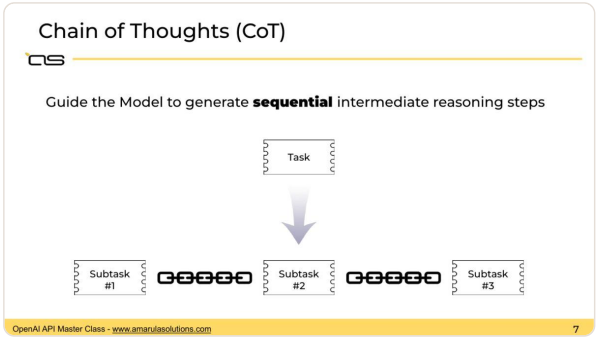

Chain of Thoughts (CoT)

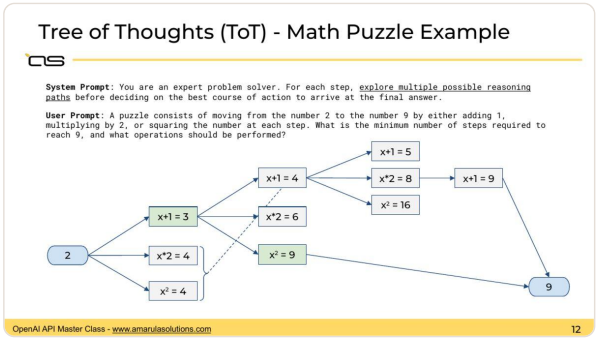

Tree of Thoughts (ToT) - Math Puzzle Example

MODULE 05 OF 09

Agentic AI

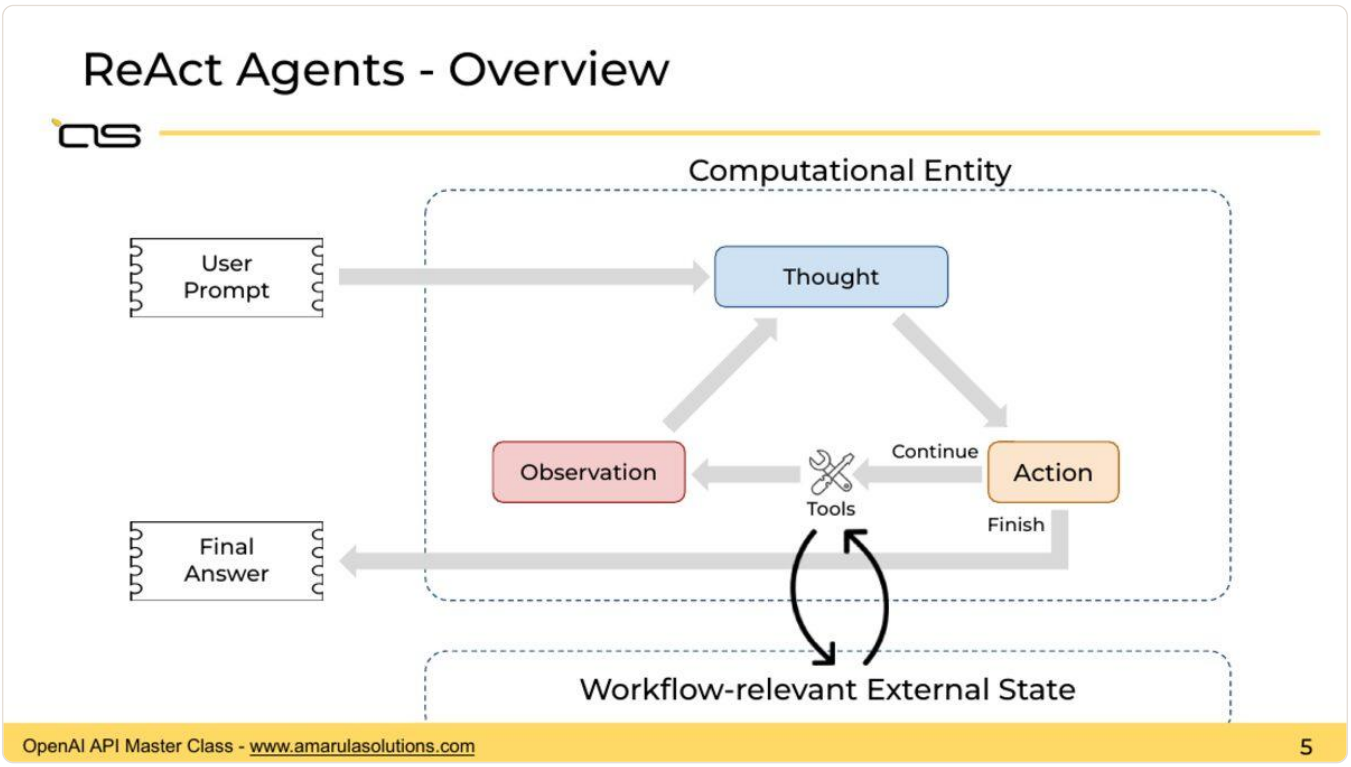

ReAct Agents - Overview

Build systems that can plan, use tools and act. Start from the ReAct loop, then move to the OpenAI Agents SDK, multi-agent orchestration and persistent agents.

WHAT YOU WILL LEARN

- What makes an AI system agentic; ReAct from first principles
- OpenAI Agents SDK: agents, runs, sessions and memory
- Handoffs, agents-as-tools, guardrails and orchestration
- Scheduled, reactive and persistent agent patterns

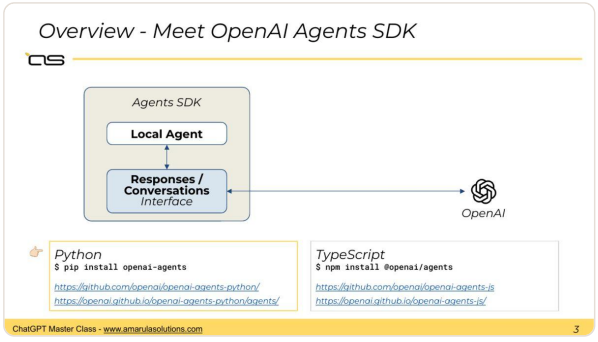

Meet OpenAI Agents SDK

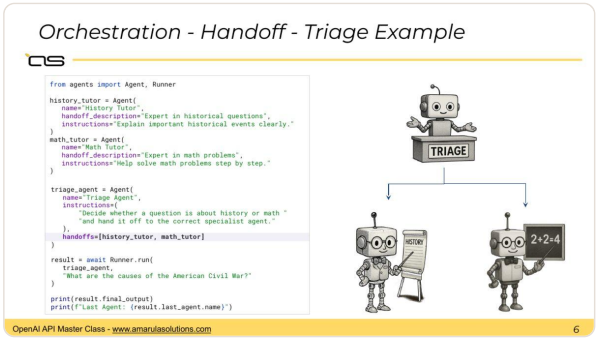

Orchestration - Handoff - Triage Example

MODULE 06 OF 09

Provider-agnostic LLM frameworks

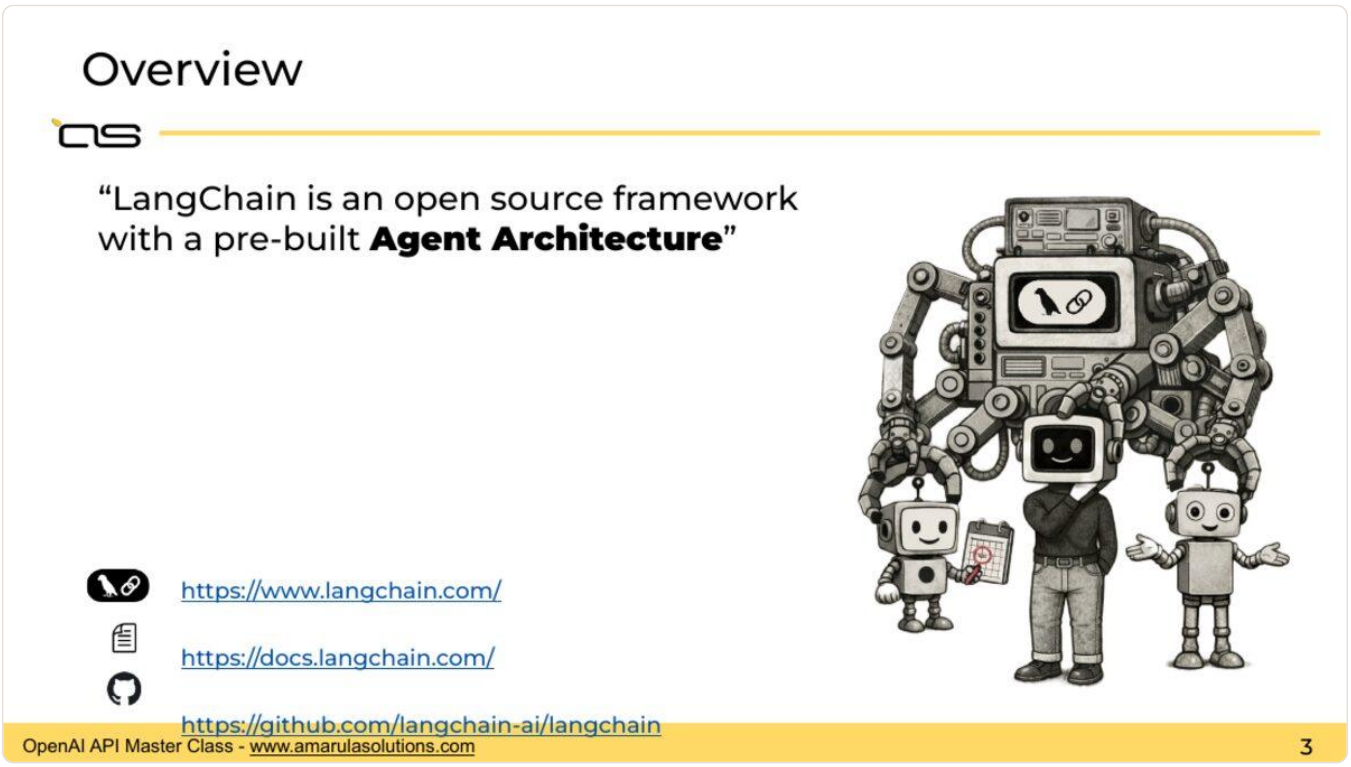

LangChain - Overview

Avoid lock-in. Use frameworks that work across providers and learn how to choose between native SDKs, LangChain and DSPy for real applications.

WHAT YOU WILL LEARN

- LiteLLM: one interface for hosted and local models
- LangChain: prompt templates, agents, middleware and LCEL
- LangGraph: stateful, graph-based orchestration
- DSPy: modules, signatures and automatic prompt optimisation

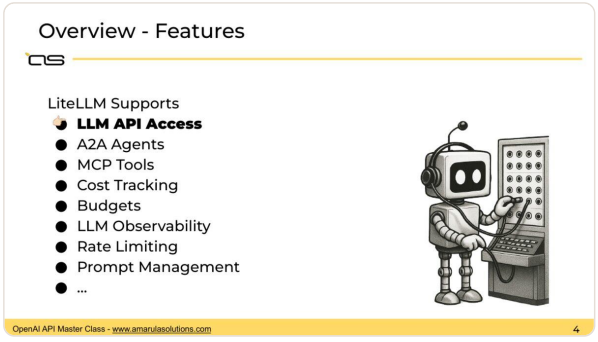

LiteLLM - Features

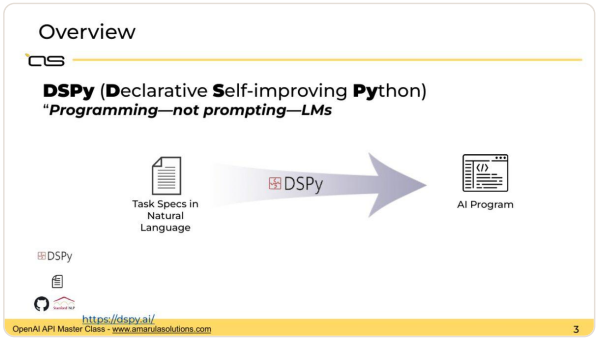

DSPy - Overview

MODULE 07 OF 09

Multimedia API

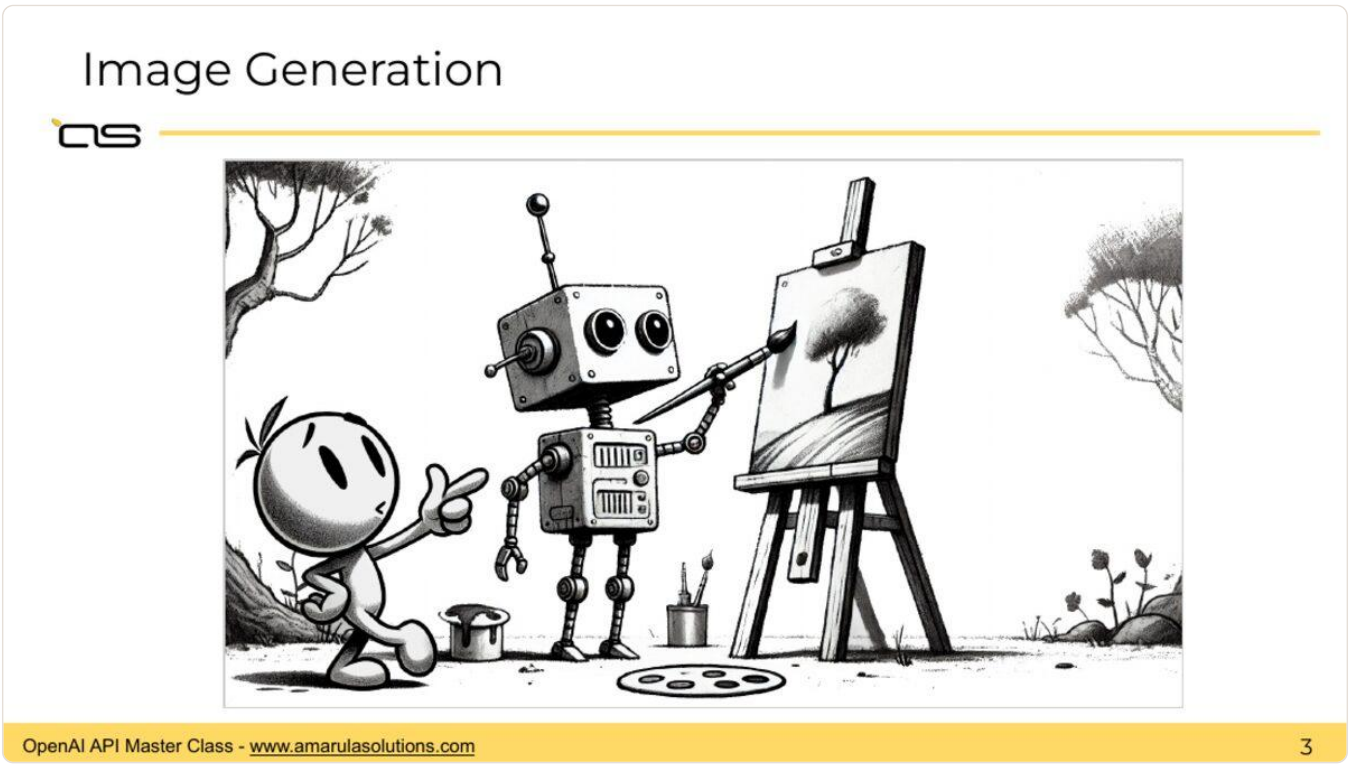

Image Generation

Add speech, vision and image generation to your applications, and understand the quality, latency and cost trade-offs of each modality.

WHAT YOU WILL LEARN

- Text-to-speech, transcription, translation and audio understanding
- Vision: image inputs, extraction, geolocation and defect analysis
- Image generation and editing with DALL·E and GPT Image
- Model and pricing trade-offs for multimodal workloads

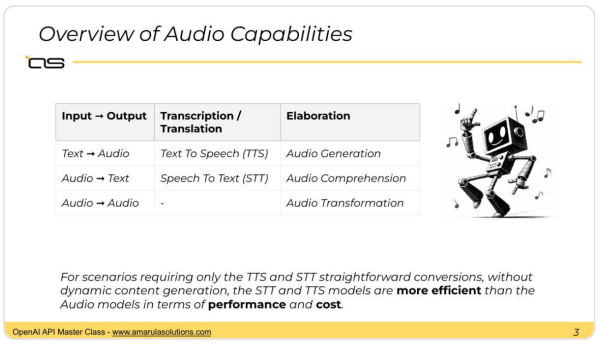

Overview of Audio Capabilities

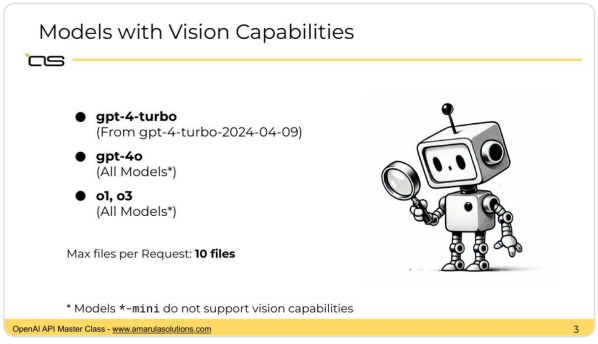

Models with Vision Capabilities

MODULE 08 OF 09

Applications for software development

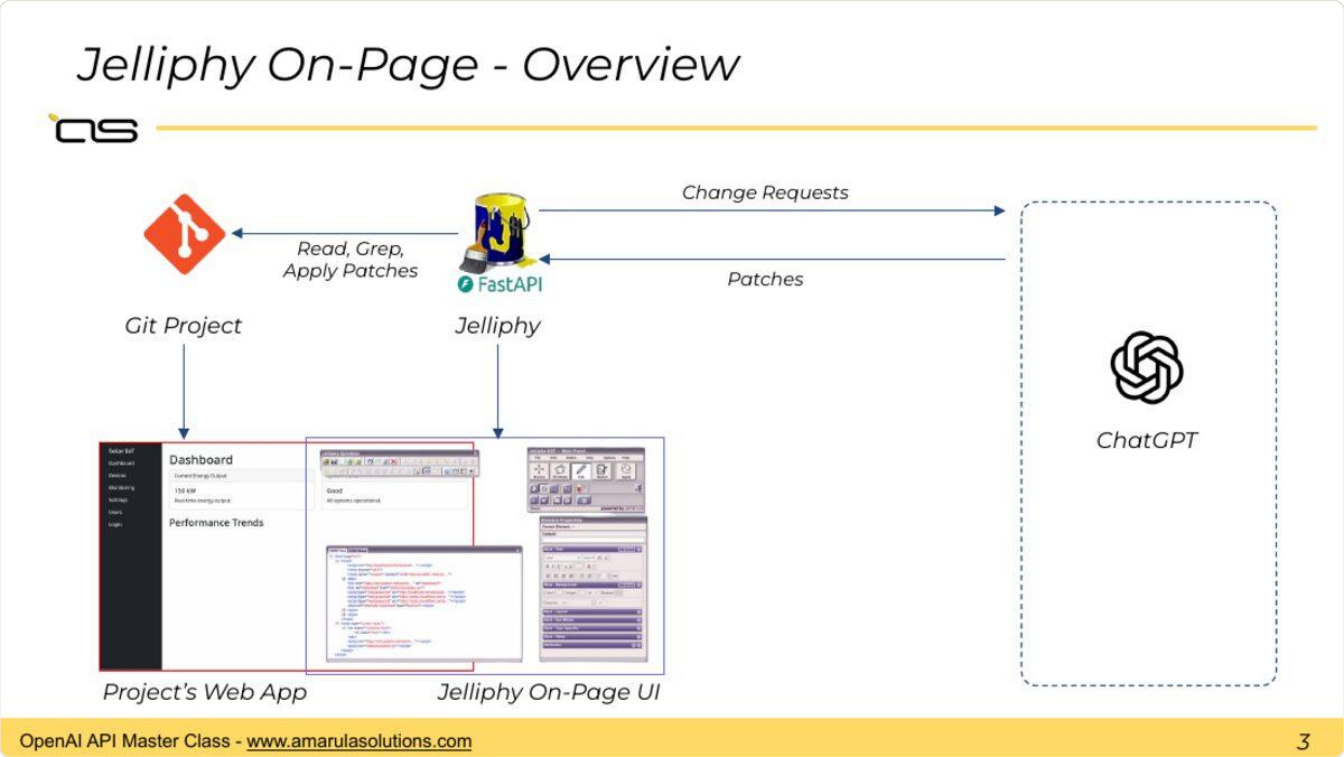

Jelliphy On-Page - Overview

Apply LLMs to the software lifecycle: generate and change code, write documentation, debug failures and build AI-powered developer tools.

WHAT YOU WILL LEARN

- Code generation, change requests and documentation assistance
- Context strategies: whole repository, diffs and on-demand retrieval
- Debugging and iterative fixing with reasoning models
- Jelliphy: building and editing web applications in the browser

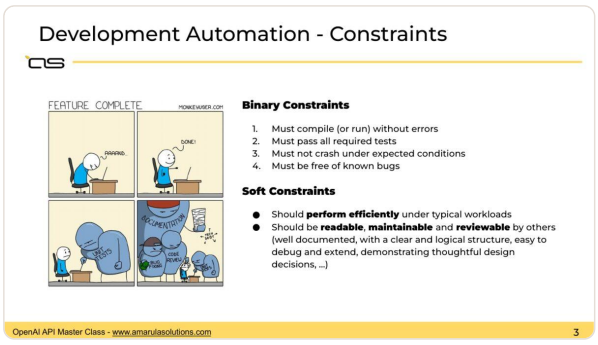

Development Automation - Constraints

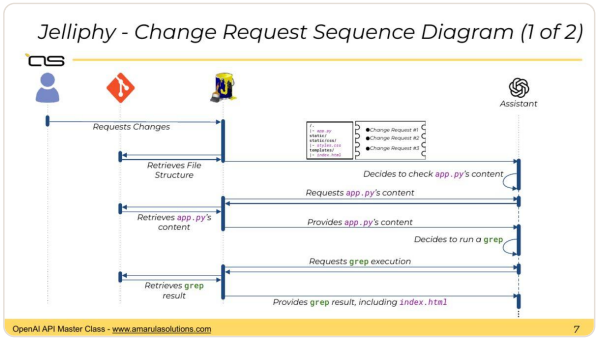

Jelliphy - Change Request Sequence

MODULE 09 OF 09

Applications for code review

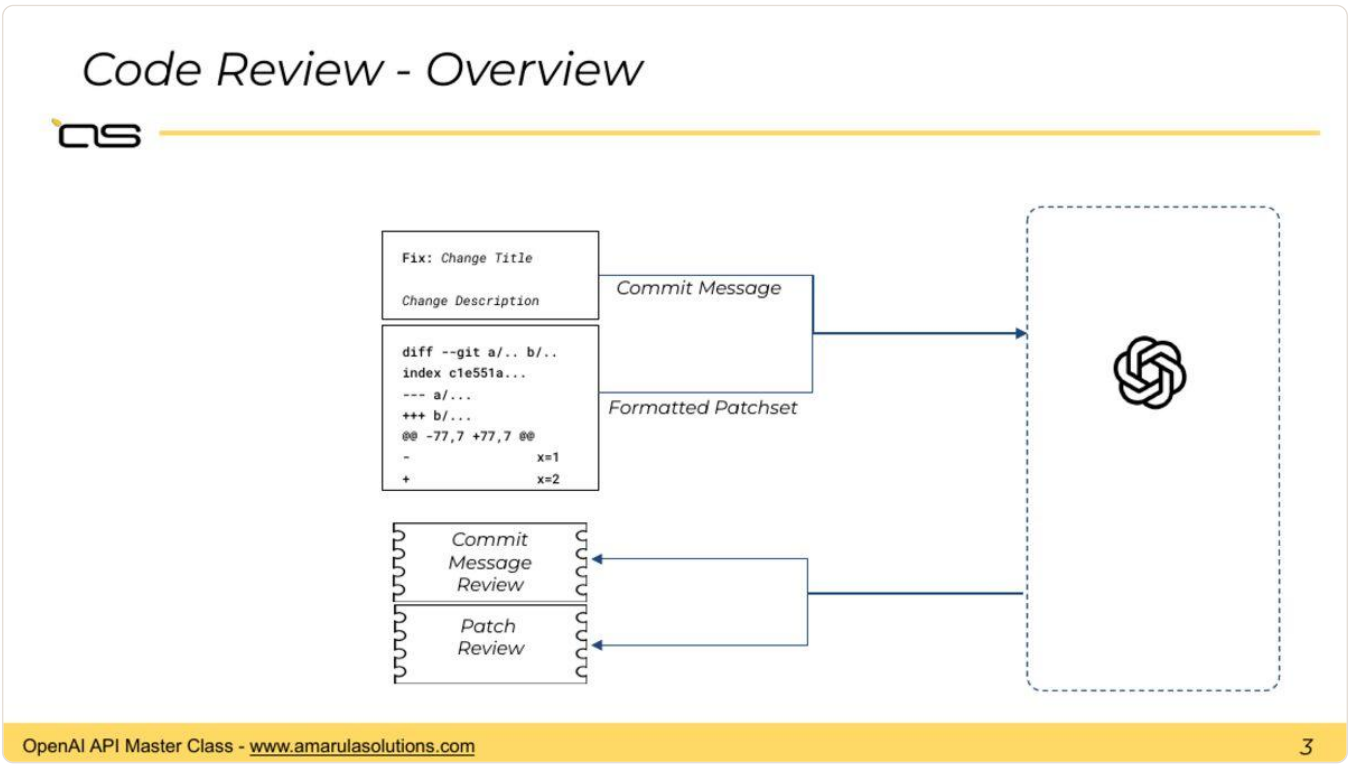

Code Review - Overview

See how AI can review code at scale, from diff analysis to a production Gerrit plugin, and how to customise review behaviour for your team.

WHAT YOU WILL LEARN

- Code review context strategies and review quality
- The ChatGPT Gerrit plugin architecture and workflow
- Reviewing change sets, votes, comments and commands
- Custom prompts, assistant instructions and directives

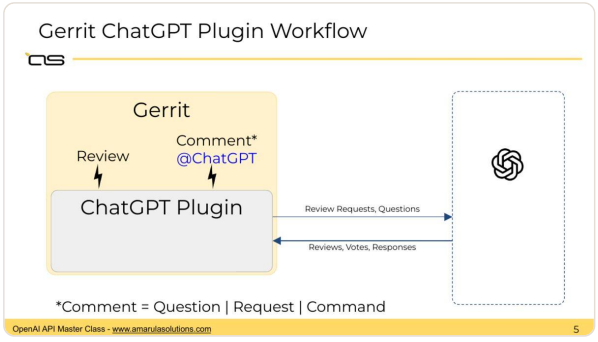

Gerrit ChatGPT Plugin Workflow

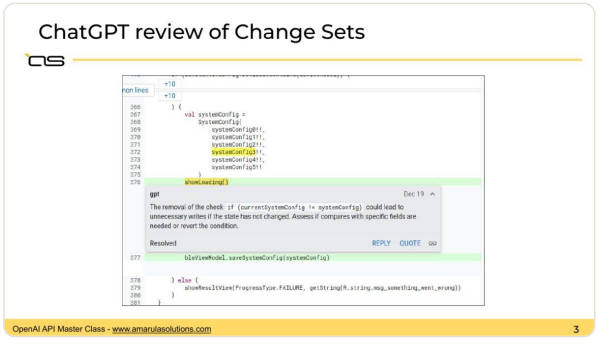

ChatGPT Review of Change Sets

Your trainer



Patrizio Gelosi

Senior Software Developer, Amarula Solutions

Patrizio holds a Master's degree in Software Engineering from the University of Pisa and has more than two decades of experience across natural-language processing, document and image processing, web editors and developer tooling.

He designed this master class around the way he builds AI-powered tools at Amarula Solutions: start with a working prototype, understand the model and the API, then iterate until the result is reliable enough for production.

Bring the master class to your team

DURATION	PARTICIPANTS	PRICE	TRAINER
3 days	**Minimum 3**	**€1,350 each**	**Patrizio Gelosi**

Sessions are arranged on request for groups of at least three participants. The programme can be adapted to the team's existing stack and experience.

Request a quotation

Use the quotation form at amarulasolutions.com/quotation/ to tell us about your team and preferred dates.

Talk to us

Contact the Amarula team at amarulasolutions.com/contact/ or info@amarulasolutions.com.

© 2026 Amarula Solutions. Selected slides are excerpts from the course deck and remain the property of Amarula Solutions.