

Specialization

AI Engineer

This specialization prepares you to build and deploy production-ready AI solutions using LLMs, RAG and AI agents, giving you the practical skills to integrate AI into workflows and drive real business impact. The AI Engineer specialization consists of three required courses.

Total CEUs: 6

Learning Outcomes:

- Explain the architecture and functionality of LLMs, including transformer architecture, embeddings, pre-training and next token generation.
- Develop, deploy, and integrate LLM-powered applications using APIs and frameworks like LangChain for tasks such as customer support, productivity, and feedback analysis.
- Design and implement RAG-based applications to retrieve, summarize, and contextualize internal knowledge while evaluating performance with standard metrics.
- Build AI agents capable of reasoning, planning, automating workflows, and integrating with tools, APIs, and data systems for business applications.
- Evaluate, monitor, and improve generative AI systems using best practices, ensuring ethical, secure, and responsible deployment.



*Courses may have prerequisites; review the course page before enrolling. A checkmark indicates the course is typically offered during that term.**

Required Courses

6 CEUs | 3 Courses

COURSE NAME & NUMBER	CEUS	FALL	WINTER	SPRING	SUMMER
LLM Fundamentals and Practical Applications AISOV.819	2.0		✓		✓
Developing RAG Based Applications AISOV.820	2.0	✓		✓	
Developing AI Agent Applications AISOV.821	2.0	✓		✓	

Completion Review

Once all specialization requirements have been met and your final grades are posted, please access your Student Portal to enroll in the "[Specialization: AI Engineer Completion Fee](#)" to begin the review process. Please allow 4-6 weeks to receive your certificate.

Note: This specialization is intended for students with prior technical experience, and completing a foundational course such as [Generative AI Fundamentals](#) is recommended before enrolling.