SEC495™: Leveraging LLMs: Building and Securing RAG, Contextual RAG, and Agentic RAG™

1 Day Course | 7 CPEs | Laptop Required

You Will Be Able To

- Build an end-to-end RAG back-end solution
- Extend a RAG to implement Contextual RAG solutions
- Understand and implement AI agents in an LLM context for RAG supervision
- Implement security controls limiting information disclosure from an LLM
- Prevent and defend against prompt injection attacks

Who Should Attend

- Anyone seeking to implement a generative-AI solution for information retrieval
- Individuals who have successfully completed SEC595™ who wish to know more about LLMs
- Professionals seeking to understand how to leverage LLMs for internal and customer-facing-information-retrieval purposes

Business Takeaways

- Students will understand how to work with and leverage vector databases.
- Students will be able to implement chatbot (and similar) style solutions internally.
- Students will know how to build AI/LLM solutions without disclosing sensitive internal information to a third party or using a public or commercial API.
- Students will understand how to build cutting-edge contextual RAG solutions.
- Students will understand how to implement agent-based AI solutions related to LLMs.

The leadership teams of many organizations have directed that the enterprise be on the lookout for opportunities to leverage AI in the business process. The problem many face is that there is very little clear articulation of what the vision for AI in the enterprise is.

While SEC595 training teaches you everything you need to know to be able to build cutting edge machine learning and AI solutions for real world cybersecurity problems, SEC495 training has an entirely different goal. Our experience tells us that most management teams asking for AI are reacting to the Large Language Models (LLMs) that have dominated the mainstream news for the last few years. In this course, you will work along with the instructor to build a completely self-hosted Retrieval Augmented Generation (RAG) system that leverages an LLM. More than this, you will learn how to implement security controls to defend the LLM from prompt injection and how to implement information sensitivity controls to limit the answers the LLM can provide based on the rights of the user.

If you need to build an LLM-based solution for question answering, knowledge-base retrieval, policy creation, or any related task, this class will get you up and running quickly.

Course Author Statement

More and more, management is looking to us to leverage AI in useful ways in the enterprise. How can we do this? What does it look like? While there are many answers to these questions, and SEC595 training provides clear answers with regard to threat hunting and monitoring, this course teaches you everything you need to know to get started building solutions that leverage LLMs. While the SEC595 solutions are extremely useful and cutting edge, the SEC495 focus on building RAG solutions using LLMs is much more readily understood by management teams who can instantly see and understand how the solution is useful.

While we focus on leveraging (and securing) RAGs for information retrieval, there are some beautiful natural extensions, such as identifying standards compliance based on policies, automated report generation, and many more. Perhaps the best part of SEC495 training is that everything is completed using on-premises containers. Of course, you can host these in the cloud, scale them up, or even trade them out for commercial APIs, but you learn how to implement all these pieces without every needing to send sensitive information to a third party. That's a big win!

—Dave Hoelzer

Section Description

SECTION 1: Building Useful LLM Solutions

After introducing the fundamental components upon which LLMs are built, you will work through how to build a traditional RAG solution. Once this is functioning, you will learn how to implement access controls to limit which data the AI can return to specific users. Next, you will improve the quality of the information delivered by your AI by learning how to build a Contextual RAG. The class finishes by introducing the notion of Agentic RAG, implementing auditor and prompting agents, and discussing other possible extensions.

Something that sets this class apart is that we are focused on the security of your information. While all the techniques taught in the class can be translated directly to the use of commercial APIs, the entire course is taught leveraging locally hosted components. While it can seem that this makes the process more complex, it actually makes the process much more transparent—and demonstrates how to implement LLM-based solutions without disclosing sensitive information to a third-party API.

TOPICS: Tokenization; Word Embeddings; CBOW and Skipgrams; Vector Data Stores; Hosting LLMs Effectively and Efficiently; Interfacing with LLMs Programmatically; Document Preprocessing and Ingestion; RAG Prompt Engineering; Limiting LLM Hallucinations; Implementing ACLs within a RAG Effectively; Implementing Contextual RAG Efficiently; Building Agentic AI Solutions; Deployment Considerations for Cost Mitigation