DAMG 7374-03 DATA ENGINEERING: Impact of Generative AI with LLM's

Course Information

Course Title: Data Engineering: Impact of Generative AI with LLM's
Course Number: DAMG 7374-03
Term and Year: Spring 2025
Credit Hour: 4
CRN: 18842
Course Format: Onsite with Virtual components

Instructor Information

Full Name: Kishore Aradhya
Email Address: k.aradhya@northeastern.edu

Instructor Biography

Kishore Aradhya is a senior technology executive with over twenty-five years of experience developing and scaling diverse technology organizations, spanning startups to Fortune 500 companies. Throughout his career, he has held leadership positions at organizations such as Stanley Black & Decker, Bose, Adobe, edX/MIT, Staples, and Monster Search. In his most recent role at Stanley Black & Decker, he led the AI and Enterprise Data Platform Engineering teams as a Senior Director of Advanced Analytics and Data. His work focused on bringing GenAI and LLM technologies to the forefront and advancing the use of critical Data Insights by working with key executive leadership.

His expertise is crafting production-ready Enterprise Data Platforms for Analytics, Machine Learning, and BI applications. This spans the design, development, and architecture of highly scalable enterprise SaaS cloud services, including Customer and Analytics Data Platforms, Data Engineering, Search, mobile, and e-commerce solutions. He also spearheaded an NLP and Computer-Vision-driven document extraction research initiative, introducing key innovations in Adobe Document product features.

Kishore Aradhya has an MBA from the Isenberg School, UMass Amherst, an MS in Computer Science, and multiple executive and academic certifications from MIT, Stanford, and Harvard professional programs. Outside his primary roles, he is an active contributor to the industry as a CDO Magazine Editorial Board Member, an Industry Advisory Board Member at DVSum (an AI Data insights startup), and a Product Advisory Council Member at Kensu (an Enterprise Data Observability startup). Furthermore, Kishore offers strategic advisories to various startups and founders, focusing on enterprise market fit and other technical and Go-to-market (GTM) strategies.

Teaching Assistant Information

Full Name: TBD
Email: TBD
Office Hours: TBD

Course Prerequisites

- Essential Python, Streamlit, and SQL programming skills.
- Essential understanding of data concepts like Data Quality, Data Transformation, and Data Integration.
- Basic understanding of AI concepts - Machine Learning, Deep Learning & NLP

Course Description

This seminar-style course explores the evolving relationship between Data and Artificial Intelligence (AI), explicitly emphasizing Generative AI with Large Language Models (LLM) in the evolving data platform architecture landscape. You will hear from data leaders and working practitioners in the industry share their insights, experiences, and challenges as they navigate this highly dynamic world. Students will engage in dynamic discussions and critical analyses of data challenges in LLM Application development with a specific impact on the Enterprise Data Platform and the ever-evolving role of Data Engineers.

Course Learning Outcomes

1. Master basic and advanced data engineering principles, focusing on AI with LLM applications.
2. Data lifecycle implementation from Business problem formulation to solution delivery.
3. Learn about current trends and challenges in modern data platform architectures through invited speakers.
4. Study real-world case studies to grasp practical tradeoffs, challenges, and limitations.
5. Use and learn about cutting-edge SOTA models (Llama 2, GPT4) in LLM applications while addressing data challenges inherent in a data engineer role.
6. Learn to read, analyze, and present NLP/LLM research through a practical implementation lens.
7. Build professional connections with field experts for future collaboration and growth.

Standard Learning Outcomes

Learning outcomes expected for all College of Engineering Graduate programs:

1. *An ability to identify, formulate, and solve complex engineering problems.*
2. *An ability to explain and apply engineering design principles appropriate to the program's educational objectives.*
3. *An ability to produce solutions that meet specified end-user needs with consideration of public health, safety, and welfare, as well as global, cultural, social, environmental, and economic factors.*

Required Tools:

We will be using the following services and tools as needed:

1. VS Code IDE (free app)
2. AWS Cloud Services
3. Snowflake
4. Data Orchestration tool: Airflow (both free and managed service)

5. Data Integration Tool: LlamaIndex, LangChain, Airbyte (free; we will use this as needed)
6. Data Transformation Tool: dbt (free, but will probably use a cloud service, dbtCloud)
7. API Services from OpenAI and/or Hugging Face.
8. Time permitting, governance, and Data Observability tools like BigID, Monte Carlo, or Kensu.

Reference Books:

1. **Fundamentals of Data Engineering**
   Authors: Joe Reis, Matt Housley Published by O'Reilly Media, Inc.
2. **Data Quality Fundamentals**
   Barr Moses, Lior Gavish, Molly Vorwerck
   Published by O'Reilly Media, Inc.

Topics Covered:

- Data Engineering
- Data Integration
- Data Orchestration
- Data Visualization
- Data Transformation
- Data Monitoring
- Data Observability
- Data Ingestion
- Data Platform Architecture
- Large Language Models (LLM)
- Generative Artificial Intelligence (AI)
- Data Governance
- Knowledge Graphs
- Data Quality

Course Activities

1. Activity #1 (Data Problem Formulation)
   Data Driven Decision Making: What Business Question are we trying to solve with Data?
2. Activity #2 (Data Platform Design & Architecture)
   Deliver an overall data platform design, the data model, and its associated implementation architecture.
3. Activity #3 (Phased Implementation Deliverable #1)
   Develop a 1st iteration of this deliverable and present this for feedback from the class.
4. Activity #4 (Phased Implementation Deliverable #2)
   Develop a 2nd iteration of this deliverable and present this for feedback from the class.
5. Activity #5 (Final Project Implementation Deliverable)
   Present the final project deliverable to the class for feedback and grading.
6. Group Research Paper (discussed throughout the class):

Selection, Analysis, and Presentation for grading.

Course Schedule (will change due to the dynamic nature of the class & speaker availability)

Week 1	- Class Introductions - Course overview, objectives, and expectations - Introduction to Data Engineering – Data Platform, GenAI, and LLM
Week 2	*Modern Data Engineering: Evolving approaches to support modern AI-driven Product Development* - Data Engineering principles - The role of a Data Engineer in the world of GenAI - The Evolving Role of Data Engineers *-- Break –* *Modern Data Platform Architectures* - Understanding data architectures - Trends in data platform architectures - How AI and LLM are shaping data platform architectures
Week 3	*AI and Data: A strong interdependency between them* - The role of data in machine learning and AI - Understanding data needs for AI - Challenges in data preparation for AI *-- Break –* *Large Language Models and Data Engineering* - Understanding LLM - The importance of data engineering in LLM - Data needs and challenges for LLM

	-- Break – Invited Speaker #1: TBD (Topic & Speaker) - Q&A session
Week 4	Prep and Discussion for: **Activity #1** (Data Problem Formulation) Data Driven Decision Making: What Business Question are we trying to solve with Data? Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. Invited Speaker #2: TBD (Topic & Speaker) - Q&A session

Week 5	Class Presentation & Feedback: **Activity #1** (Data Problem Formulation) Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Invited Speaker #3: TBD (Topic & Speaker) - Q&A session
Week 6	Prep and Discussion for: **Activity #2** (Data Platform Design & Architecture) Deliver an overall data platform design and the data model and its associated implementation architecture. Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Data Engineering in different Industries: - Manufacturing - Life Sciences - Finance - Retail (B2B & B2C)
Week 7	Class Presentation and Feedback: **Activity #2** (Data Platform Design & Architecture) Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Invited Speaker #4: TBD (Topic & Speaker) - Q&A session
Week 8	Prep and Discussion for: **Activity #3** (Phased Implementation Deliverable #1) Develop a 1st iteration of this deliverable and present this for feedback from the class. Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Invited Speaker #5: TBD (Topic & Speaker) - Q&A session
Week 9	Class Presentation and Feedback: **Activity #3** (Phased Implementation Deliverable #1) Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Invited Speaker #6: TBD (Topic & Speaker)

	- Q&A session

Week 10	Ethics, Privacy, and Security in Data Engineering - Data ethics in AI and LLM - Data privacy concerns - Security practices in data engineering -- *Break* – Prep and Discussion for: **Activity #4** (Phased Implementation Deliverable #2) Develop a 2nd iteration of this deliverable and present this for feedback from the class.
Week 11	Class Presentation & Feedback: **Activity #4** (Phased Implementation Deliverable #2) Develop a 2nd iteration of this deliverable and present this for feedback from the class. Miscellaneous Topics: Deep Dive into different Data Engineering areas as needed. -- *Break* – Invited Speaker #7: TBD (Topic & Speaker) Managing ethics, privacy, and security in Data Platforms
Week 12	What's with all the *Ops and how is that related to Data Engineering: - DataOps - MLOps - DevOps What is Data Observability: - Why is this so critical and the glue that holds everything together. -- *Break* – Invited Speaker #8: TBD (Topic & Speaker) - Q&A session
Week 13	Future of Data Engineering: Emerging Trends and Challenges - Impact of AI advancements on Data Engineering - The future of data platforms with AI and LLM - Open problems and challenges in Data Engineering -- *Break* – Invited Speaker #9: TBD (Topic & Speaker) - Q&A session
Week 14	Final Class Presentation & Feedback: **Activity #5** (Final Project Implementation Deliverable) Present the final project deliverable to the class for feedback and final refinement.
Week 15	Final Class Presentation & Feedback: **Activity #5** - Final Class Presentations (contd., from Week 14) Course Review and Wrap-up - Recap of major course themes - Where to go from here and discuss the challenges and do a class retrospective.

Grade Breakdown:

Class Participation: 10%
Project Completion: 35%
Final Project Presentation: 30%
Group Research Paper Analysis & Presentation: 25%

Please Note: If you miss more than two classes without clear, explicit permission from the Instructor/TA, you will automatically be dropped 1 grade level.

Grading Scale: 95-100% A	87-89.9% B+	77-79.9% C+	
	84-86.9% B	74-76.9% C	
90-94.9% A-	80-83.9%B-	70-73.9% C-	69.9% or below F

Attendance/Late Work Policy.

Attendance Policy

Students are expected to complete course readings, participate in class discussions or other learning activities during the unit, and complete written assignments for each unit during the time of that unit. It is understood that there might be one week when active participation in ongoing class conversations and learning activities might be delayed. Beyond one week, if there is an absence or lateness in participation (1) faculty must be notified in advance; (2) grades will be adjusted accordingly.

Late Work Policy

Students must submit assignments by the deadline in the time zone noted in the syllabus. Students must communicate with the faculty before the deadline if they anticipate work will be submitted late. Work submitted late without prior communication with faculty will not be graded.

End-of-Course Evaluation Surveys

Your feedback regarding your educational experience in this class is critical to the College of Engineering. Your comments will make a difference in our curriculum's future planning and presentation.

At the end of this course, please take the time to complete the evaluation survey at T**https://neu.evaluationkit.com**. Your survey responses are entirely **anonymous and confidential**. For courses six weeks in length or shorter, surveys will be open one week before the end of the courses; for courses greater than six weeks in length, surveys will be open for two weeks. An email will notify you when surveys are available to your Husky Mail account.

Academic Integrity

A commitment to the principles of academic integrity is essential to the mission of Northeastern University. Promoting independent and original scholarship ensures that students derive the most from their educational experience and pursuit of knowledge. Academic dishonesty violates the most fundamental values of an intellectual community and undermines the achievements of the entire University.

As academic community members, students must become familiar with their rights and responsibilities. In each course, they are responsible for knowing the requirements and restrictions regarding research and writing, examinations of whatever kind, collaborative work, using study aids, the appropriateness of assistance, and other issues. Students are responsible for learning documentation conventions and acknowledging sources in their fields. Northeastern University expects students to complete all examinations, tests, papers, creative projects, and assignments according to the highest ethical standards, as set forth either explicitly or implicitly in this Code or by the direction of instructors.

Go to http://www.northeastern.edu/osccr/academic-integrity-policy/ to access the full academic integrity policy.

University Health and Counseling Services

As a student enrolled in this course, you are fully responsible for assignments, work, and course materials as outlined in this syllabus and in the classroom. Over the semester, if you experience any health issues, please get in touch with UHCS.

For more information, visit https://www.northeastern.edu/uhcs.

Student Accommodations

Northeastern University and the Disability Resource Center (DRC) are committed to providing disability services that enable students who qualify under Section 504 of the Rehabilitation Act and the Americans with Disabilities Act Amendments Act (ADAAA) to participate fully in the university's activities. To receive accommodations through the DRC, students must provide appropriate documentation demonstrating a substantially limiting disability.

For more information, visit https://drc.sites.northeastern.edu.

Library Services

The Northeastern University Library is at the hub of campus intellectual life. Resources include over 900,000 print volumes, 206,500 e-books, and 70,225 electronic journals.

For more information and education-specific resources, visit https://library.northeastern.edu.

24/7 Canvas Technical Help

For immediate technical support for Canvas, call 617-373-4357 or email help@northeastern.edu
Canvas Faculty Resources: https://canvas.northeastern.edu/faculty-resources/
Canvas Student Resources: https://canvas.northeastern.edu/student-resources/ For assistance with my Northeastern e-mail, and basic technical support:
Visit ITS at https://its.northeastern.edu
Email: help@northeastern.edu
ITS Customer Service Desk: 617-373-4357

Diversity and Inclusion

Northeastern University is committed to equal opportunity, affirmative action, diversity, and social justice while building a climate of inclusion on and beyond campus. In the classroom, members of the University community work to cultivate an inclusive environment that denounces discrimination through innovation, collaboration, and an awareness of global perspectives on social justice.

Please visit http://www.northeastern.edu/oidi/ for complete information on Diversity and Inclusion

Title IX

Title IX of the Education Amendments of 1972 protects individuals from sex or gender-based discrimination, including discrimination based on gender-identity, in educational programs and activities that receive federal financial assistance.

Northeastern's Title IX Policy prohibits Prohibited Offenses, which are defined as sexual harassment, sexual assault, relationship or domestic violence, and stalking. The Title IX Policy applies to the entire community, including male, female, transgender students, faculty, and staff.

In case of an emergency, please call 911.

*Please visit **https://www.northeastern.edu/ouec** for a complete list of reporting options and resources both on- and off-campus.*