

CS 594 Special Topics Proposal

January 2026

Title: Vector Databases

Instructor: Abolfazl Asudeh (asudeh@uic.edu)

Course Description

The rapid rise of large language models (LLMs) and Retrieval-Augmented Generation (RAG) for tasks such as information retrieval, question answering, and knowledge-intensive reasoning, together with the increasing prevalence of large-scale unstructured data, has driven the emergence of vector database management systems. These systems are designed to extend traditional database architectures by efficiently storing, indexing, and querying vector representations such as embeddings, which are central to modern machine learning and AI-driven applications.

This course explores advanced topics in vector databases and Retrieval-Augmented Generation (RAG), with a focus on the underlying data structures, algorithms, and system designs. Students will learn about information retrieval fundamentals, embedding models, vector database architectures, indexing techniques for high-dimensional data, and modern approaches for efficient query processing in vector databases. The course also covers other data models, such as knowledge graphs, and a practical implementation of RAG pipelines using vector databases.

Prerequisites: Students should be familiar with the topics covered in CS 401 and CS 480.

Method of instruction

- The primary mode of instruction will be via lectures given by the instructor.
- Depending upon the size and background of the class, students will be asked to read new research papers, summarize the papers, and make a brief presentation in the class.
- Selected students will be provided with open research questions and will be encouraged to work on them. Should a publishable work develop,

the instructor will provide guidance towards its eventual publication in a suitable journal or conference.

Class meetings: The class will have two 75-minute lectures per week.

Course Topics

The course will cover the following topics:

1. Introduction / Fundamentals of Information Retrieval
2. Information Retrieval and Ranking
3. Introduction to Vector Embeddings and Contextual Embeddings
4. Vector Database Architectures and Challenges
5. Similarity Measures
6. Query Models
7. Vector indexing: Basics
8. Indexing in High Dimensions
9. Quantized / Table-Based Indexing
10. Tree and Graph Based Indexing
11. Filtered Search
12. Introduction to Knowledge Graphs
13. Introduction to RAGs
14. RAG Models and Challenges
15. Vector Databases for RAG

Textbook: No textbook.

Student Deliverables

Students are expected to come to lecture and are encouraged to participate.

1. **Scribe Notes.** Each student will prepare scribe notes for one or two lectures during the semester. All the lecture notes will be made available on the class web page.
2. **Programming Project.** Each student will implement a Vector Database and develop a RAG system on top of it.

3. **Research Project.** Students will work on a research project related to vector databases and unstructured data processing. Students should form teams of up to 2 people. Students are encouraged to propose their own projects. The instructor will also provide suggested topics and will highlight open problems during the lectures.

Grading

- Scribe Notes: 20%
- Programming Project: 30%
- Research Project: 40%
- Active Class Participation: 10%

Overlap with other courses

To the best of my knowledge, the course material has minimal overlap with any other CS or MCS courses, including CS 580, CS 581, and CS 582.