



Prof Roy Ka-Wei Lee

Roy Ka-Wei Lee is an Assistant Professor at the Information Systems Technology and Design Pillar, Singapore University of Technology and Design. He is appointed Deputy Sector Lead of SUTD AI Mega Center. His research lies in the intersection of data mining, machine learning, social computing, and natural language processing. He is leading the Social AI Studio, which focuses on designing next-generation social artificial intelligence systems. He has published in top-tier venues in data mining, computation linguistics, and artificial intelligence domains such as IEEE TKDE, IEEE TNNLS, IEEE ICME, IEEE ICDM, ACM Multimedia, SIGKDD, IJCAI, WWW, ACL, EMNLP, COLING, etc. He also serves on the program committees of these top conferences and is a reviewer for the top journals. He is a senior IEEE member.

Link: <https://istd.sutd.edu.sg/people/faculty/roy-lee>

Course Outline:

1. Introduction to Retrieval-Augmented Generation:

- Definition and core components of RAG systems.
- How RAG differs from traditional generation models.
- Overview of real-world applications for RAG in information retrieval and AI-driven content generation.

2. Understanding Retrieval Systems:

- Key principles behind retrieval models (BM25, dense retrieval with transformers).
- Techniques for integrating external knowledge sources into retrieval systems.
- Hands-on Activity 1: Implementing a basic retrieval system using text-based queries.

3. Generation Models and Integration with Retrieval:

- Overview of generation models (e.g., GPT, T5) and their application in AI.
- Strategies for combining retrieval and generation to improve task performance.

continue ...

4. Training and Fine-Tuning RAG Systems:

- How to fine-tune retrieval and generation models for specific use cases.
- Optimizing RAG systems for efficiency and accuracy.
- Hands-on Activity 3: Fine-tuning a RAG model for a specific task (e.g., answering user questions from a custom knowledge base).

5. Case Studies: RAG in Real-World Applications:

- **Online Trust & Safety:** Using RAG for misinformation detection or flagging inappropriate content.
- Hands-on Activity 4: Applying a RAG system to a case study, such as retrieving and generating responses for online content moderation.

6. Evaluation and Future Trends:

- Evaluation metrics for RAG systems (precision, recall, fluency).
- The future of retrieval-augmented systems in various AI applications.
- Ethical considerations in using RAG for tasks involving sensitive information retrieval.

Pre-requisite:

- **Basic Understanding of Machine Learning and Deep Learning:** Familiarity with transformers, language models, and information retrieval techniques.
- **Knowledge of Natural Language Processing (NLP):** Experience with text-based language models and search algorithms.

Programming Skills: Proficiency in Python, and experience with deep learning frameworks such as PyTorch or TensorFlow