Intelligent Medical Engineering

(智慧醫療工程) Fall 2026 | College of Intelligence Computing

Instructor: Professor Kuan-Fu Chen, MD, PhD

"Bridging the gap between clinical practice and cutting-edge artificial intelligence: From raw medical data to causal decision-making and real-world deployment."

Welcome to **Intelligent Medical Engineering**! This course provides a comprehensive, hands-on approach to applying Artificial Intelligence in healthcare. Rather than just learning predictive algorithms in isolation, students will master the entire lifecycle of next-generation medical AI—transitioning from basic data integration to advanced causal inference, dynamic treatment recommendations, and rigorous regulatory deployment.

🌟 Course Features & Highlights

- **Real-World Clinical Problem-Solving:** Engage in project-based learning using complex clinical datasets. Tackle advanced challenges such as developing algorithms for an **Artificial Pancreas** or designing **Reinforcement Learning (RL) for clinical decision-making** in critical care (e.g., personalized sepsis management).
- **Next-Generation Architectures:** Move beyond standard predictive models to prescriptive AI. Learn to implement **Causal Machine Learning**, **Reinforcement Learning (RL)**, **Variational Autoencoders (VAEs)**, and **Causal Distillation Models** to handle the messy reality of unmeasured clinical confounders.
- **Healthcare Interoperability:** Master real-world deployment standards, specifically focusing on **SMART on FHIR**, ensuring the AI solutions you build can seamlessly integrate into modern Electronic Medical Records (EMR).
- **Innovating Education: Personalized Learning:** We will test-drive agent-based learning into the overhauled IME course this year, offering students a tailored, interactive educational experience.

🧠 Key Learning Objectives & Core Concepts

By the end of this course, students will be proficient in:

1. **Causal Inference & Reinforcement Learning:** Moving beyond correlation by utilizing VAEs to disentangle confounding factors and designing robust RL policies for dynamic, multi-timescale treatment regimes.
2. **Advanced Deep Learning in Medicine:** Implementing advanced architectures, including Causal Distillation Models and state-space models (e.g., Decision Mamba), to process irregular clinical time-series and waveforms.

3. **Data Preprocessing & EDA:** Handling severe missingness, selection bias, and integrating large-scale clinical databases using HL7/FHIR standards.
4. **Clinical Rigor & Ethics:** Uniquely focusing on the medical validity and safety of models. Learn to apply **QUADAS-AI** for diagnostic bias evaluation and the CONSORT-AI guidelines for designing robust medical AI clinical trials.

🚀 Prerequisites & Pathways Forward

- **(Suggested) Prerequisites:** Machine Learning, Deep Learning, and Basic Programming.
- **Your Future in Medical AI:** Completing this course builds a robust foundation for high-impact careers and advanced research, including developing AI-driven Clinical Decision Support Systems (CDSS), healthcare analytics, and leading physician-scientist collaborations.

For any questions regarding course material, please reach out to the department office or contact Prof. Kuan-Fu Chen directly. We look forward to innovating the future of healthcare with you!

Figure 1: example of SCOPE-Sepsis Architecture. Demonstrating the causal bridge from "blind" to "sighted" AI using AutoIV and EDVAE, leading to Counterfactual Experience Augmentation (CEA) and safe Causal RL Policy Learning.

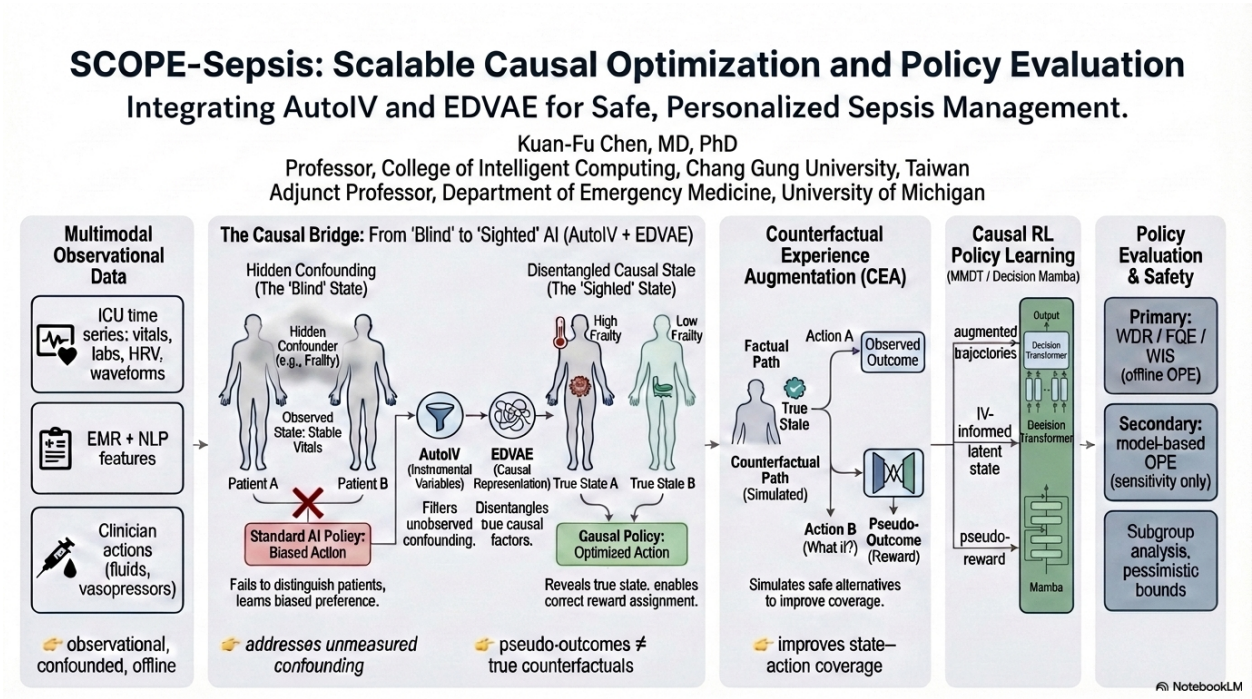

Figure 2: The Messy Reality of Clinical RL. Highlighting the clinical bottlenecks of incomplete information, practice variation, and selection bias, alongside the Out-Of-Distribution (OOD) problem in traditional algorithmic baselines.