A171 · Bridging Time-to-Event and Generative Deep Learning for Longitudinal Cardiovascular Digital Twins

Presenter: Siying (Avon) Yang — Department of Epidemiology & Biostatistics , Harvard T.H. Chan School of Public Health

Authors: Siying (Avon) Yang

Session: Day 2 · Fri Oct 2, 2026 · 2:15–4:15 PM

Longitudinal digital twins should both predict when disease occurs and simulate how health evolves, yet these objectives are usually studied separately. We develop a deep learning framework for cardiovascular digital twins using the Health Professionals Follow-up Study (HPFS), with repeated questionnaires from 1986 onward.

We represent each questionnaire wave as a structured health state containing demographics, anthropometrics, smoking, physical activity, diet, and blood-pressure measures. Among 51,387 participants, we constructed 547,699 prediction origins and supported cohorts of 531,277, 506,863, and 434,822 origins for 2-, 5-, and 10-year cardiovascular disease (CVD) prediction. Participant-level splitting, training-only preprocessing, time-support restrictions, and explicit competing-risk labels prevent temporal and cross-participant leakage.

On this common benchmark, we compare three Transformer-based strategies: (1) a MOTOR-style time-to-event model that learns task-conditioned survival representations across multiple endpoints; (2) a Delphi-style autoregressive model that predicts subsequent health states and event timing, enabling stochastic rollouts of future trajectories; and (3) a joint model optimized for both time-to-event and next-state prediction. The joint objective tests whether simulation-aware representation learning improves risk prediction and whether survival supervision improves the clinical validity of generated futures.

Models will be evaluated at 2-, 5-, and 10-year horizons using time-dependent discrimination, calibration, and Brier scores. For the generative model, CVD risk will be estimated from Monte Carlo trajectory rollouts and compared with direct survival estimates. This framework tests whether deep generative models can produce calibrated epidemiologic risks, rather than merely plausible trajectories, and establishes a foundation for longitudinal population-health digital twins.