M3

Utility that matches the use

Which fidelity measures actually predict success on which downstream clinical tasks.

Master's · 30 credits | Statistics · Machine learning | Suits: KTH (ML, Data Science) · SU (Statistics, CS)

The question

Synthetic data is scored with dozens of fidelity metrics. Which of them predict how well a downstream task performs, and does the answer change between software testing, cohort sizing and model training?

Why it matters

A single utility score invites the wrong reading. Buyers and access bodies need to know which evidence counts for which purpose. Mapping metrics to tasks turns fitness for purpose from a slogan into something measurable.

What you would do

01 Catalogue fidelity metrics in common use: marginal divergences, correlation structure, propensity scores, train-synthetic-test-real performance.

02 Define a set of downstream tasks of different difficulty on public seed data.

03 Measure metric scores and task performance across many synthetic releases and look for which metrics are predictive, and where.

04 Recommend a minimal, purpose-specific evidence set.

What you would produce

- A metric-to-task predictiveness study
- A recommended evidence set per purpose
- Open evaluation code

Good fit if you have

Python, statistics, experimental design. Comfortable running many experiments and reasoning about them.

Data and access

Public seed data (CMS DE-SynPUF) and synthetic releases from day one. Work on real patient records only inside a partner's approved secure environment, with ethics approval. We hold no Swedish patient data.

Timing and place

Spring or autumn 2027, 20 weeks / 30 credits. Stockholm (Solna) or hybrid.

Supervision

Gaurav Raj Chattarki, Founder, Synthetix Health, together with your academic supervisor.

Questions or a proposal of your own? gaurav@synthetixhealth.com · synthetixhealth.com/careers#thesis
All projects run under a confidentiality and IP agreement agreed before work begins. We review drafts only for confidential information, never for results. Negative results are publishable.