

GPU-Accelerated Bayesian Longitudinal Item Response Theory

Author:

Jonas Kupschus

Co-author:

Andreas Schuppert

Frequently, collections of binary or ordinal responses can be assumed to be noisy measurements of one or more underlying latent traits. These data arise, for example, in patient reported outcome measurements, educational testing, and employee or customer feedback questionnaires. Item response theory (IRT) is routinely used in these settings to estimate two complementary quantities: person-specific latent traits, which capture the underlying processes informing downstream interventions, and item-specific parameters, which help refine instruments and flag problematic items.

Increasingly, item response data are collected in an intensive longitudinal fashion, e.g., in political surveys or mHealth scenarios. This brings with it additional modeling requirements such as robust uncertainty quantification, out-of-sample prediction, and the ability to handle missing data.

We introduce dynamirt, a GPU-accelerated Python package for building and fitting Bayesian multidimensional item response theory models for repeated-measurement scenarios. We present a real-world case study using mHealth data from neuropathic pain patients.