Optimal Inference of Galaxy Properties from SEDs in the Era of Next Generation Surveys

〇Suchetha Cooray[1] (1.Stanford University, USA)

Galaxy physical properties inferred from spectral energy distributions (SEDs) have systematic uncertainties of up to ~0.3 dex in stellar masses and star formation rates, driven by modeling assumptions rather than data quality. As next-generation surveys (PFS, LSST, DESI, Euclid, SPHEREx, Roman) deliver 10^8–10^{10} galaxies, this systematic floor will hinder the science return unless inference codes can both marginalize over assumptions and exploit population-level information. We present Tengri, a GPU-accelerated, fully differentiable JAX framework built on DSPS (Hearin et al. 2023) that implements the full combinatorial space of SED modeling choices. We have implemented multiple stellar population synthesis libraries, dust attenuation/emission laws, SFH and metallicity history parameterizations, AGN component, and multiple nebular emission backends, all within a single unified pipeline. At ~50 µs per forward evaluation and ~10 s per posterior, Tengri makes Bayesian model averaging tractable at billion-galaxy scale, delivering the first honest, model-marginalized catalogs with fully propagated uncertainties. Tengri also leverages Information Field Theory (Enßlin 2019) for high-dimensional hierarchical inference: as a demonstration, we sample a >1.1-million-dimension posterior on a laptop CPU in ~10 minutes to recover star-formation burstiness parameters from many galaxy spectra jointly. Tengri provides a robust and scalable approach to maximizing the science return from next-generation multi-wavelength surveys.