Simulation Based Inference for Galaxy Evolution

Bristol, UK

27 - 30 May, 2025

Contents

Rationale	5
Invited Contributions	1
Introduction to SBI (Maximilian von Wietersheim-Kramsta)	1
Introduction to the LTU-ili Inference Pipeline (<i>Matthew Ho</i>)	1
SBI in the wild: the challenges of inferring galaxy properties from real data	
(Kartheik Iyer)	1
Pixel-based SED fitting of galaxies using SBI (<i>Patricia Iglesias Navarro</i>) .	2
Synthesizer: a Software Package for Synthetic Astronomical Observables	
$(Christopher Lovell) \dots \dots \dots \dots \dots \dots \dots \dots \dots $	2
Breaking astrophysical degeneracies using galaxy imaging at the pixel scale	
(John Wu)	3
Teaching Dark Matter Simulations to Speak the Galactic Language (Shivam	
Pandey)	3
Constraining Dark Matter through Stellar Kinematics and Simulation-	
based Inference $(Tri Nguyen)$	4
Tools for SBI ($Will Handley$)	4
A Brief Introduction to JAX (<i>Viraj Pandya</i>)	4
Contributed Abstracts	5
DREAMS of Dark Matter Annihilation Signals (Alex Garcia)	5
Unveiling Galaxy Evolution through UV Luminosity Functions with Simulatio Based Inference: Insights from CAMELS and Illustris-TNG (<i>Alice</i>	n-
Matthews)	5
GalactiKit: reconstructing mergers from $z=0$ debris using simulation-based	_
inference in Auriga (Andrea Sante)	6
Simulation-based inference of the AGN population based on the radiation-	0
regulated unification model (<i>Dimitra Gerolymatou</i>)	6
Panchromatic modelling of resolved high-z galaxy observations using Bayesian $N_{\rm exact}$ Network $N_{\rm exact}$ and $N_{\rm exact}$	-
Neural Networks (Steven Ramnichal)	(
Physics-Informed, AI-Accelerated Dynamical Modeling and Bayesian In-	7
Transfer learning for simulation based inference $(Alex Science)$	(0
DIPIN: Fast and differentiable forward modeling of IFS observations (Arms	0
Lena Schaible)	8
	0

	The assembly histories of HSC galaxy images via invertible neural net- works, contrastive learning, and cosmological simulations (Annalisa	
	Pillepich)	9
	The optical and infrared are connected (Christian Kragh Jespersen)	9
	Characterizing Resolved Stars and Dust Extinction in Nearby Galaxies	
	with Simulation-Based Inference (Christina Lindberg)	10
	Empirically predicting emission lines and deriving galaxy properties in	
	DESI (Eduardo A. Hartmann)	10
	Feedback constraints from old and new probes (<i>Giulio Fabbian</i>)	11
	CASBI (Giuseppe Viterbo)	11
	Recovering Dark Matter Halo Properties with SBI (Haley Bowden)	11
	Constraining the stochasticity of star formation across galaxy populations	
	with simulation based inference (<i>Harry Bevins</i>)	12
	Inferring the position and characteristic size of individual bubbles using	
	NIRSpec observations (Ivan Nikolić)	12
	CLONES: Digital Twins of the Local Universe for Bias-Free Inference	
	(Jenny Sorce)	13
	Direct inference of the full thermodynamic state of galaxy groups and	10
	clusters (Joeu Brasnenning)	13
	Inferring Dark Matter Properties from the Local Universe (<i>Jonah Rose</i>)	13
	Galaxy evolution and cosmological measurements through the forward	10
	modelling of galaxy surveys (Luca Tortorelli)	14
	Interpretable Scalable Galaxy Formation Parameter Inference at the Pop-	11
	ulation Level (Lucas Makinen)	14
	Forward-modelling galaxy surveys: A parametric model for non-local stochas-	11
	tic galaxy hiss (Marimilian von Wietersheim-Kramsta)	15
	Enhancing Galaxy Mass Estimation with Domain Adaptation using HaloFlow	10
	(Nikhil Garuda)	15
	Multi-tracer Inference of Astrophysics for the First Billion years of the	10
	Universe (Nikolaos Triantafullou)	16
	Decoding Calaxy Evolution Models with cINNs: Insights from SDSS-like	10
	Spectrophotometric Data (Nils Candebat)	16
	The unknown unknowns: Bopresentation Learning for simulation based in	10
	forence in galaxy formation and cosmology (Schastian Truille Comer)	17
	Approximate Maximum Likelihood Estimation with Local Score Matching	11
	(<i>Charman Khoo</i>)	17
	(Sherman Knoo)	11
	Decer)	10
	A Development Consistent Forward Model for Colory SED Development	10
	(Such sthe Coorner)	10
	(Suchetha Cooray)	10
	Towards combined constraints on galaxy astrophysics, dust attenuation,	
	and cosmology from galaxy luminosity functions and colors (1jitske	10
	Statkenbury)	19
Aı	thor Index	21

Rationale

Simulation Based Inference (SBI; see Cranmer et al. 2020) is currently being applied to many astrophysical problems which face similar computing barriers. Following the success of our 2024 meeting, we aim to bring together both experts in SBI as well as galaxy evolution experts with SBI experience, to discuss problems and solutions faced when applying SBI to problems in galaxy evolution. This will be a small, focused meeting to enable plenty of time for discussion and practical sessions.

Invited Contributions

Introduction to SBI

Maximilian von Wietersheim-Kramsta Durham University

Introduction to the LTU-ili Inference Pipeline

Matthew Ho Columbia University

The Learning the Universe Implicit Likelihood Inference (LtU-ILI) pipeline is a toolkit rapid, user-friendly, and cutting-edge machine learning inference for astrophysics and cosmology. In this tutorial, I will present an introduction to simulation-based inference with LtU-ILI, including how to adapt prepackaged models to complex data modalities, such as image, time-series, and graph data. I will also showcase new aspects of the LtU-ILI pipeline, including evidence computation and hyperparameter tuning.

SBI in the wild: the challenges of inferring galaxy properties from real data

Kartheik Iyer Columbia University

Simulation Based Inference opens up the possibility of recasting a wide range of problems into inference problems where we constrain key parameters of physical interest using carefully crafted training sets based on simulations + forward models. I will introduce three such problems where SBI is applied to real data: (i) inferring individual galaxy properties from multi-band photometry, (ii) inferring individual galaxy properties from few-band images, and (iii) constraining populationlevel burstiness in galaxy SFRs using spectral features. In each of these cases, I'll highlight the challenges in setting up and running the SBI pipeline, ranging from creating a representative forward model, to dealing with realistic noise and missing data. Promising as SBI is, the astronomy community would benefit from developing a common set of checks to ensure that the predicted posteriors are unbiased and interpretable.

Pixel-based SED fitting of galaxies using SBI

Patricia Iglesias Navarro Instituto Astrofísico de Canarias

We introduce an efficient Bayesian SED fitting framework tailored to multiwavelength, pixel-level data from the JADES and JEMS surveys based on Normalizing Flows. We combine the unprecedented spatial resolution, the wide wavelength coverage and the depth provided by HST and JWST. Our method is trained using simulated photometry generated from the MILES stellar population models, with non-parametric star formation histories, realistic noise, and filter sensitivity thresholds emulating ACS and NIRCam observations. This amortized, simulation-based inference is validated on mock datasets, demonstrating robust and well-calibrated posterior distributions for stellar properties, achieving a high R^2 score of 95% for the stellar mass and an inference speed of 10^{-4} seconds per pixel. We apply this pipeline to JADES and JEMS observations, producing spatially resolved maps for stellar population parameters across galaxy regions. We present six example galaxies with $0.6 \leq z \leq 5.6$. Our model provides a scalable solution for extracting high-fidelitity stellar population properties from HST+JWST datasets, opening the way for statistical studies of galaxy evolution at sub-galactic scales.

Synthesizer: a Software Package for Synthetic Astronomical Observables

Christopher Lovell University of Portsmouth

I will present Synthesizer, a fast, flexible, modular and extensible platform for modelling synthetic astrophysical observables. Synthesizer can be used for a number of applications, but is predominantly designed for generating mock observables from analytical and numerical simulations. These use cases include (but are not limited to) analytical modelling of the star formation and metal enrichment histories of galaxies, the creation of mock images and integral field unit observations from particle based simulations, detailed photoionisation modelling of the central regions of active galactic nuclei, and spectro-photometric fitting. We provide a number of stellar population synthesis models, photoionisation code configurations, dust models, and imaging configurations that can be used 'out-of-the-box' interactively. Synthesizer is unique in its speed and flexibility, which makes it ideal for Simulation Based Inference approaches. I will present early results using Synthesizer on mock photometric distribution functions from the CAMELS suite.

Breaking astrophysical degeneracies using galaxy imaging at the pixel scale

John Wu

STScI/JHU

Machine learning (ML) is revolutionizing our ability to study galaxy evolution and large-scale structure. Deep neural networks can now reliably predict galaxies' physical properties, such as cold gas content, metallicity, and even their entire optical spectra from imaging alone. Moreover, new generative models can produce synthetic galaxy images that span the space of observed galaxies conditioned on their physical properties. Thus, we can exploit simulation based inference (SBI) to constrain galaxy spectral energy distribution (SED) models using a combination of photometry and morphology. Galaxies' information-rich morphologies at the pixel scale enable us to break astrophysical degeneracies, such as the dust-age degeneracy, which cannot be resolved using broad-band color information.

Teaching Dark Matter Simulations to Speak the Galactic Language

Shivam Pandey

Johns Hopkins University

Developing fast and efficient methods for simulating our observable Universe is a key bottleneck in maximizing information extraction from cosmological datasets. I will discuss a novel approach to obtaining galaxies and their properties from fast dark matter simulations by treating this as a language translation problem. Using a transformer-based multi-modal architecture, I will first demonstrate how we can obtain high-resolution halo catalogs from very low-resolution (and approximate) dark matter simulations. Then, I will show how the same architecture can be augmented to directly generate hydrodynamical simulation-based galaxy catalogs and their observable photometry, conditioned on dark matter-only simulations. Finally, I will discuss the potential of this approach in the context of ongoing large-scale surveys such as DESI and Euclid.

Constraining Dark Matter through Stellar Kinematics and Simulation-based Inference

Tri Nguyen

Northwestern University

The particle nature of dark matter (DM) remains one of the most profound unsolved problems in modern astrophysics and cosmology. While the Lambda-Cold Dark Matter cosmology provides excellent agreement with observations on cosmological scales, significant tensions emerge on galactic and sub-galactic scales, where theoretical predictions systematically deviate from observation. The phase-space distribution of DM in the Milky Way and nearby galaxies offers a unique way to probe fundamental DM properties, including particle mass and self-interaction crosssection. Dwarf galaxies and stellar streams, in particular, retain subtle signatures of DM physics, making them exceptional testbeds for probing its nature. In this talk, I will demonstrate how precision stellar kinematics in these systems can be exploited to reconstruct the three-dimensional DM distribution. I will also highlight the transformative role of simulation-based inference in addressing the complexities of this problem and enabling more precise constraints on the particle properties of DM.

Tools for SBI

Will Handley University of Cambridge

A Brief Introduction to JAX

Viraj Pandya Columbia University

I will lead an introductory tutorial to the JAX python package showing how it can accelerate your code by serving as a drop-in replacement for numpy/scipy. Topics will include: (1) just-in-time compilation, (2) automatic vectorization and parallelization over multiple cores of a single CPU/GPU node as well as across multiple GPUs, (3) automatic differentiation for model sensitivity analysis and gradientbased optimization/inference, (4) reproducible random number generation, and (5) popular libraries in the JAX ecosystem for streamlining common tasks. This tutorial will be made available in a standalone Jupyter notebook.

Contributed Abstracts

DREAMS of Dark Matter Annihilation Signals

Alex Garcia University of Virginia

The detection of dark matter presents one of the greatest challenges of modern astronomy. Possibly one of the most promising avenues is via high energy particles created by dark matter self-annihilation events. These annihilation events, naturally, depend strongly on the particle physics of the dark matter, but also on the astrophysics of the Universe. The small changes to your assumptions of the density of dark matter within the Milky Way can have a significant impact on the predicted annihilation signal. We use DREAMS simulation suite of Milky Way galaxies, with systematic variations in astrophysical and dark matter particles, to make theory predictions for these annihilation signals. We present a systematic study on the impact of astrophysics variations on your predicted dark matter annihilation signal. We find: (i) increased supernova feedback can lower the expected signal, and that (iii) warm dark matter does not significantly change your expected signal from CDM. Ultimately, we provide a tool for the particle physics community to bound the astrophysical uncertainties in the search for dark matter annihilation signals.

Unveiling Galaxy Evolution through UV Luminosity Functions with Simulation-Based Inference: Insights from CAMELS and Illustris-TNG

Alice Matthews

UCL

Ultraviolet Luminosity Functions (UVLFs) provide key insights into galaxy formation and evolution and are a fundamental tool in understanding and tracing star formation histories and galaxy evolution processes. This work leverages simulationbased inference (SBI) using the CAMELS (Cosmology and Astrophysics with MachinE Learning Simulations) IllustrisTNG Sobol sequence dataset - consisting of 2048 simulations spanning 28 parameters, to generate synthetic photometry and ultraviolet luminosity functions and colours via the SYNTHESIZER module. We train an ensemble of neural networks within the LTU-ILI framework to infer posterior distributions for these 28 cosmological and astrophysical parameters. Initial results demonstrate that supernova and large-scale structure parameters generalise well, while AGN-related parameters remain challenging to constrain due to the limited simulation volume from CAMELS. Extending this analysis, we aim to apply our final trained models to real UVLF observations at z=0.5, derived from XMM-OM and Swift-UVOT data. This will allow us to test the ability of SBI methods to generalise to real data, bridging simulations and observations, ultimately improving constraints on galaxy evolution models.

GalactiKit: reconstructing mergers from z=0 debris using simulation-based inference in Auriga

Andrea Sante Liverpool John Moores University

GalactiKit is a methodology for estimating the lookback infall time, stellar mass, halo mass and mass ratio of the disrupted progenitors of Milky Way-like galaxies at the time of infall. GalactiKit uses simulation-based inference to extract the information on galaxy formation processes encoded in the Auriga cosmological MHD simulations of Milky Way-mass halos to create a model that relates the properties of mergers to those of the corresponding merger debris at z=0. I will discuss on three implementation of GalactiKit considering the dynamical, chemical, and the combined chemo-dynamical information of debris. The kinematics of the debris is found to trace the lookback time at which the progenitor was first accreted into the main halo. However, chemical information is necessary for inferring the stellar and halo masses of the progenitors. Stellar masses are predicted more accurately than the halo masses, which could be related to the scatter in the stellar mass-halo mass relation. The model trained combining chemical and dynamical data of debris provides the most accurate predictions for the merger parameters.

Simulation-based inference of the AGN population based on the radiation-regulated unification model

Dimitra Gerolymatou University of Geneva

X-ray surveys of active galactic nuclei (AGN) provide direct constraints on the properties of individual AGN, such as their accretion, reflection, and obscuration. Previous AGN population synthesis models have not addressed such constraints self-consistently. Here we use SBI to constrain the geometrical and physical properties of the AGN population. We perform numerical simulations with our ray-tracing code, RefleX, which allows the self-consistent modelling of the X-ray emission of AGN with flexible circumnuclear and source geometries. We create our synthetic population by sampling the black hole mass function and Eddington ratio distribution function of local AGN and we construct a geometry based on the radiation-regulated model. Using the RefleX-simulated emission of the AGN population, we simultaneously

reproduce several observed properties of Swift/BAT detected AGN, such as the lineof-sight NH distribution, the fraction of obscured AGN as a function of Eddington ratio, and their differential number counts. With this approach, we can test the consistency of the radiation-regulated model in the local Universe with the most comprehensive set of X-ray observables, while constraining the size and density of the dusty torus.

Panchromatic modelling of resolved high-z galaxy observations using Bayesian Neural Networks

Steven Ramnichal University of Bath

Galaxy growth has traditionally been studied from two complementary perspectives: the build-up of stellar mass as inferred from galaxy-integrated SED modelling, and the growth in size as captured by high-resolution imaging. Both approaches tend to gloss over the fact that stellar populations and dust attenuation may vary spatially. Increasingly, deep multi-wavelength observations of high-redshift galaxies are available from UV to mm wavelengths, at varying resolution. We developed a new modelling tool aimed at maximally exploiting all of this information without sacrificing unresolved or poorly resolved bands, and without sacrificing resolution where available. Our modelling is fundamentally 3D in nature, and extracts information on stellar populations and their radial variation as well as star-dust geometries, without imposing an attenuation law a priori. We jointly model observational constraints on wavelength dependent fluxes, sizes and light profile shapes using a computationally efficient machine learning emulator trained on dust radiative transfer calculations of an extensive library of toy model galaxies. Applications to observed and simulated galaxies will be discussed.

Physics-Informed, AI-Accelerated Dynamical Modeling and Bayesian Inference for Galaxy Populations

Viraj Pandya

Columbia University

The existence of morphological classes and "fundamental planes" for galaxies implies that their evolution must be deeply regulated by physical laws. However, deciphering the relevant astrophysics and robustly testing our modern paradigm for galaxy formation in a cosmological context requires a new generation of physics-informed, data-driven and AI-accelerated techniques for Bayesian inference and causal discovery. I will introduce my next-generation dynamical model ("sapphire") for evolving galaxy populations using a system of non-linearly coupled differential equations that track mass, metal and (newly) energy flows between galaxies and their complex dynamical ecosystems. It is the first and only galaxy formation model of its kind to be end-to-end differentiable and multi-GPU capable. I will highlight early applications of stochastic gradient descent, Hamiltonian Monte Carlo and implicit likelihood inference that reveal the key astrophysical parameters governing the galaxy-halo connection. By elucidating how and why different observables are sensitive to these parameters, I will summarize how sapphire can help definitively answer 50-year-old questions about how galaxies self-regulate and what they can teach us about the fundamental physics of dark matter, dark energy and inflation.

Transfer learning for simulation-based inference

Alex Saoulis

University College London

Many SBI methods are data-hungry, relying on large numbers of simulations to train neural density estimators and neural compression techniques that often utilise deep learning models. This poses a challenge when working with expensive, highfidelity simulations, where generating large datasets is infeasible. Transfer learning leverages knowledge from related datasets—often larger or of varying fidelity—to improve performance on more limited target datasets. In this work, we explore transfer learning strategies to reduce the number of required simulations while maintaining—or even improving—the accuracy of inference. Transfer learning is widely used in deep learning but remains underexplored in SBI, where domain adaptation between simulations of different fidelities could significantly improve efficiency. We apply established transfer learning techniques to CNNs used for neural compression, and present a novel approach for normalizing flows. By reducing dependence on large simulation suites, transfer learning could allow for fewer, more accurate, simulations, improving both the efficiency and precision of SBI.

RUBIX: Fast and differentiable forward modeling of IFS observations

Anna Lena Schaible IWR, Heidelberg University

Understanding galaxy formation and evolution requires comparing cosmological hydrodynamical simulations with observational data. Forward modeling provides a powerful solution by transforming simulation outputs into synthetic observations, enabling direct comparisons with real data. Integral Field Spectroscopy (IFS) generates spatially resolved spectra. RUBIX (https://astro-rubix.web.app/) is a cuttingedge, JAX-based framework for forward modeling IFS data cubes from cosmological simulations such as IllustrisTNG and NIHAO. By leveraging GPU acceleration, RU-BIX achieves a 600x speed improvement over traditional CPU-based tools, reducing computation times from hours to seconds. Beyond speed, RUBIX incorporates automatic differentiation, enabling gradient-based optimization and simulation-based inference, which extends its capabilities for both forward modeling and inverse analysis. The code is open source. RUBIX's modular design ensures compatibility with diverse simulations, spectral libraries, and telescope configurations and with that a broad range of applications.

The assembly histories of HSC galaxy images via invertible neural networks, contrastive learning, and cosmological simulations

Annalisa Pillepich MPIA Heidelberg

We constrain the merger and assembly history of hundreds of thousands galaxies observed with the Hyper Suprime Camera from their optical images by combining conditional invertible neural networks, contrastive learning, and cosmological simulations. In particular, we infer the posterior distributions of three summary statistics of the galaxies histories, namely the stellar mass and time of the last major merger and the fraction of ex-situ stars, by training on the output of the IllustrisTNG cosmological galaxy simulations. I will discuss details and choices of our algorithm and results, including comparisons across different cosmological simulation models, the importance of properly mocking the training datasets into the target observational space, and our attempts to quantitatively evaluate – and hence use in the inference process! – the level of realism of the simulated galaxies at the image level.

The optical and infrared are connected

Christian Kragh Jespersen Princeton University

Galaxies are often modelled as composites of separable components with distinct spectral signatures, implying that different wavelength ranges are only weakly correlated. They are not. We present a data-driven summarizer model to predict infrared emission from optical spectra, achieving accuracies of $\chi_N^2 \approx 1$ for all photometric bands in WISE, as well as good colors. Our empirically based summarizer exploits subtle correlations between physical processes and is thus able to predict typically IR-derived properties, e.g., the bolometric luminosity of AGN and dust parameters such as q_{PAH} . We find that current SED-fitting methods are incapable of making comparable predictions, and that model misspecification often leads to correlated biases in star-formation rates and AGN luminosities. To help improve SED models, we determine what features of the optical spectrum are responsible for our improved IR prediction accuracy, and identify several lines (CaII, SrII, FeI, [OII] and H α), which point to the complex chronology of star formation and chemical enrichment being incorrectly modelled.

Characterizing Resolved Stars and Dust Extinction in Nearby Galaxies with Simulation-Based Inference

Christina Lindberg

Johns Hopkins University

Using broadband photometry of resolved stars, we should in principle be able to characterize and infer their stellar parameters using just their spectral energy distributions (SED). These measurements are broadly useful for a wide range of science cases - estimating masses of supernova progenitors, measuring the initial mass function (IMF), characterizing planet hosts, etc. However, this seemingly simple inference is far more complicated in practice: a large fraction of stars are unresolved binaries, and photometry can be heavily extinguished by dust at shorter wavelengths. Additionally, assuming inference for a single star in isolation throws away valuable information about the surrounding stellar populations, which can be used to inform priors on the inference. To address these issues, we utilize simulation-based inference (SBI) to construct an adaptive Bayesian SED modeling tool to characterize resolved stellar populations. SBI provides a flexible framework, allowing us to easily incorporate new stellar models, adjust priors like the IMF or star-formation history (SFH), and recognize outliers in the data, all of which would require significant computational resources to rerun with traditional grid-based SED-fitting techniques.

Empirically predicting emission lines and deriving galaxy properties in DESI

Eduardo A. Hartmann

IAC

The spectra of galaxies contain a wealth of information about the physical properties of their stellar populations. Using stellar population synthesis techniques, we can derive key parameters such as ages, metallicities, masses, and star formation histories (SFHs). However, traditional algorithms for this analysis are computationally intensive, making them impractical for large-scale surveys. A simulation-based inference framework (Iglesias-Navarro et al. 2024) has demonstrated strong performance in deriving posterior distributions of metallicities and SFHs from absorption spectra, though it is limited to quiescent galaxies without emission lines. In this work, we present how we trained the AstroCLIP foundation model to empirically infer the equivalent widths of emission lines from the continuum. By integrating this approach, we can efficiently determine stellar population properties across a large, diverse sample of galaxies from the DESI survey.

Feedback constraints from old and new probes

Giulio Fabbian CNRS Institut d'Astrophysique Spatiale

The monopole of the tSZ effect provides an integrated measurement of the electron pression in collapsed halos. As such, it is a powerful probe of the thermal history of the universe and of galaxy formation models. In this talk I will show how a new reanalysis of the FIRAS data, the only data set suitable to carry out a tSZ monopole measurement, combined with modern hydrodynamical simulations and SBI can be used to rule out the feedback parameter space. I will show how this approach give results competitive with more recent data and give a prospect on the use of these techniques in the coming years in light of the new planned experiments.

CASBI

Giuseppe Viterbo IWR, Heidelberg University

Galaxies evolve hierarchically through merging with lower-mass systems and the remnants of destroyed galaxies are a key indicator of the past assembly history of our Galaxy. However, accurately measuring the properties of the accreted galaxies and hence unraveling the Milky Way's (MW) formation history is a challenging task. Here we introduce CASBI (Chemical Abundance Simulation Based Inference), a novel inference pipeline for Galactic Archeology based on Simulation-based Inference methods. CASBI leverages on the fact that there is a well defined mass-metallicity relation for galaxies and performs inference of key galaxy properties based on multi-dimensional chemical abundances of stars in the stellar halo.

Recovering Dark Matter Halo Properties with SBI

Haley Bowden University of Arizona

The properties of a galaxy are closely linked to those of its host dark matter halo. While numerous studies have demonstrated that halo mass—and assembly-related properties such as concentration and spin—play key roles in shaping galaxy evolution within simulations, directly inferring these halo characteristics for observed galaxy populations remains challenging. To enable this, we employ simulation-based inference to leverage the information encoded in the observable properties of galaxies and their local environments, thereby constraining the properties of their host halos. Our model, trained on mock galaxy catalogs from UNIVERSEMACHINE, decodes the connection between observables and the underlying halo mass and assembly history of dark matter halos. When applied to data from the Galaxy and Mass Assembly survey, this technique enables a direct investigation of the galaxy—halo connection across a large, low-redshift sample.

Constraining the stochasticity of star formation across galaxy populations with simulation based inference

Harry Bevins

University of Cambridge

Star formation is inherently stochastic, governed by processes such as gas inflows, feedback, and equilibrium regulation, each operating on distinct timescales. These processes imprint statistical fluctuations on star formation histories (SFHs), which can be characterised through a power spectral density (PSD) and its corresponding autocorrelation function (ACF). By adopting a physically motivated PSD, we can define a Gaussian process model for SFHs, capturing variations around a mean star formation rate. Observationally, several spectral features trace the burstiness or smoothness of star formation, including Dn(4000) and UV flux ratios. Building on previous work, we show that the distribution of these tracers across a galaxy population encodes information about the underlying PSD governing star formation variability. We forward-model galaxy SEDs across a range of PSD parameters and employ neural posterior estimation to infer the PSD parameters from observed star formation tracers. We validate our method on simulated galaxy populations and discuss its application to large surveys such as SDSS. This approach provides a novel, data-driven framework for constraining the physical drivers of star formation across galaxy populations.

Inferring the position and characteristic size of individual bubbles using NIRSpec observations

Ivan Nikolić

Scuola Normale Superiore Pisa

The Lyman- α line is a sensitive probe of the neutral hydrogen along the line-ofsight of a galaxy and can thus be used to constrain the topology of the Epoch of Reionization (EoR). Probing the topology is important to connect the sources that reionized the Universe to the growth of cosmic HII regions. The usual approaches include a number of simplifying assumptions including treating galaxies as independent probes of their environments. In this work we present a new framework to infer the properties of the local HII region around a group of galaxies. We forward-model Lyman-alpha spectra as would be observed by JWST, dealing with all relevant sources of uncertainty. We use likelihood-free inference (simulation-based inference, SBI) in the flux space to avoid biases from non-zero flux covariances. We take into account each galaxy's relative location with respect to the local HII region, using the full potential of the Lyman- α line. We find that the observed galaxy number densities of $n_{\rm gal} \sim 0.005 - 0.0085$ cMpc⁻³ should suffice to estimate the size and location of the local HII region at a percent level accuracy. Such number densities are well within reach of future JWST surveys.

CLONES: Digital Twins of the Local Universe for Bias-Free Inference

Jenny Sorce

CNRS CRIStAL / IAS

Understanding dark matter and dark energy, which make up 95

Direct inference of the full thermodynamic state of galaxy groups and clusters

Joey Braspenning Max Planck Institute for Astronomy

Current state-of-the art inference of physical quantities from X-ray spectroscopy gives single point estimates for the temperature, density and abundances. Often, individual abundances cannot even be separated. In this talk I will show how, using neural density estimators, direct inference from the spectrum can surpass all current methods and yield not only accurate point predictions, but even the full thermodynamic distribution of the gas. Using X-ray spectra generated from the TNG-Cluster and TNG-300 simulations, I will show how we uncover the cold gas phase, where traditional 2-temperature models fail, and are able to measure the gas temperature distribution over 3 orders of magnitude. Using the SHAP framework we find that the neural network extracts information not just from strong emission lines, but also from far more subtle features. In the context of large all-sky surveys such as eRosita, our model offers the opportunity to measure the full thermodynamic state of the gas for essentially all observed galaxy groups and clusters in the sky. This in turn gives powerful insights into the formation and evolution of cosmic structures.

Inferring Dark Matter Properties from the Local Universe

Jonah Rose Princeton University

We introduce the DREAMS project, an innovative approach to understanding the astrophysical implications of alternative dark matter models and their effects on galaxy formation and evolution. The DREAMS project will ultimately comprise thousands of cosmological hydrodynamic simulations that simultaneously vary over dark matter physics, astrophysics, and cosmology in modeling a range of systems—from galaxy clusters to ultra-faint satellites. Such extensive simulation suites can provide diverse training sets for machine-learning-based analyses. I will cover the analysis of two simulations. One suite has a unique Warm Dark Matter particle mass for each simulation, the other contains Cold Dark Matter, both vary the initial density field and several parameters controlling the strength of baryonic feedback within the IllustrisTNG model. With these simulations, we examine how astrophysical probes can be used to predict properties of both dark matter and astrophysics.

Galaxy evolution and cosmological measurements through the forward modelling of galaxy surveys

Luca Tortorelli LMU Munich

Stage IV experiments are set to provide the stringent constraints ever on cosmological parameters, provided that we carefully control systematics such as galaxy redshift distributions. The forward-modelling (FM) of photometric and spectroscopic galaxy surveys, a method that bridges cosmology with galaxy evolution, arose as one of the most promising approaches to solve this problem. This method requires as main pillar a realistic model for the galaxy population, which if properly calibrated allows us to do very precise galaxy evolution measurements. In this talk I will discuss the past (LT+18,20,21) and on-going efforts (LT+24) in FM galaxy surveys, from the modelling of the galaxy population to the simulation of images and spectra using simulators I developed, as well as the data required to calibrate this model. I'll show that FM already provides precise redshift distribution estimates for cosmology surveys and solid constraints on the galaxy luminosity function. I will also show a new galaxy population model based on stellar population synthesis that I'm currently developing (GalSBI, LT+25), from which I am able to draw physical galaxy quantities, such as stellar masses, SFHs, metallicities, gas, dust and AGN properties.

Interpretable, Scalable Galaxy Formation Parameter Inference at the Population Level

Lucas Makinen Imperial College London

The sensitivity of phenomenological galaxy formation models to variations in their free parameters remains largely uncharted territory. I will discuss a next-generation physical model of galaxy formation (sapphire; Pandya et al. 2023, in prep.) written in JAX that is automatically differentiable, capable of running on multiple GPUs and ripe for interpretable implicit likelihood inference. The sapphire model is both descriptive of the data and interpretable. Paired with simulation-based inference we can for the first time obtain constraints on the physical attributes of the model on data. Using the novel fishnets architecture (Makinen et al. 2023), I will show that a neural network can robustly and scalably learn the feedback parameters assumed in different model realizations if it is provided specific populations of noisy galaxy observables and/or data-derived empirical correlations. Fishnets inference can be scaled to arbitrary numbers of galaxies once trained. I will start by showing that the stellar-mass-halo-mass relation does not contain information about the mass loading of supernova-driven winds but it is sensitive to the energy loading. Then I will discuss the information content of the gas and stellar mass-metallicity relations, interstellar gas fractions and various observables of circumgalactic gas including warm/hot ion column densities, X-ray emission, the Sunyaev-Zel'dovich yparameter and dispersion measures to fast radio bursts as well as pulsars in the Large Magellanic Cloud. The combination of fishnets and sapphire appears promising for both constraining the physics of galaxy formation from populations in both astrophysical and cosmological contexts

Forward-modelling galaxy surveys: A parametric model for non-local stochastic galaxy bias

Maximilian von Wietersheim-Kramsta Durham University

Galaxy positions and shapes as tracers of the large-scale structure of the Universe are key observables to test cosmological models in the late Universe down to the non-linear regime. Accurate cosmological constraints require a good understanding of the galaxy-halo connection, i.e. galaxy bias, which can be influenced by complex multi-scale baryonic dynamics and evolve over time. I will present a novel semi-analytic model of galaxy bias which naturally incorporates stochasticity and non-locality while ensuring the physicality and isotropy of the galaxy field modelled from the underlying matter field. I will show how this model can be calibrated from large-scale hydrodynamical simulations, such as FLAMINGO, to accurately sample galaxy populations from simulated matter fields while conditioning on various galaxy properties. I will also present how this model can be used to efficiently forward model galaxy positions consistent with hydrodynamical simulations while varying over cosmological and astrophysical parameters. Therefore, this model may enable an accurate simulation-based inference (SBI) analysis of galaxy clustering and galaxy-galaxy lensing observables from galaxy surveys such as Euclid, Rubin LSST and DESI.

Enhancing Galaxy Mass Estimation with Domain Adaptation using HaloFlow

Nikhil Garuda University of Arizona

We present an extension of HaloFlow (Hahn et al. 2023), a machine learning approach that infers the dark matter and stellar mass of galaxies from their photometry and morphology. HaloFlow uses simulation-based inference with normalizing flows to conduct rigorous Bayesian inference. In this study, we use state-of-the-art synthetic galaxy images from Bottrell et al. (2023) that are constructed from the IllustrisTNG, Simba and Eagle hydrodynamic simulations and include realistic effects of the Hyper Suprime-Cam Subaru Strategy Program (HSC-SSP) observations. While HaloFlow improves mass estimation, it isn't robust on its own when applied to this dataset. In this talk, I'll discuss the domain adaptation techniques we've explored to enhance performance. I'll present the results from these methods, as well as future directions for improving model robustness and generalizing across different datasets.

Multi-tracer Inference of Astrophysics for the First Billion years of the Universe

Nikolaos Triantafyllou Scuola Normale Superiore (SNS)

A hundred million years following the big bang, the first stars and galaxies formed in the densest regions and shone light into the Universe. Their radiation heated and ionized their environs, driving the Universe's transition from a fully-neutral to a fully-ionized state. Modelling of these astrophysical processes typically requires the use of semi-analytical models, such as in our case, whose parameters must be tuned using observational tracers. In this work we aim, to assess the constraining power of a multi-tracer inference of these parameters, which is expected to improve upon singletracer approaches by leveraging the information encoded in correlated observables. Tracers such as the cosmic 21cm-signal, UV luminosity functions and interstellar medium transition lines are considered. To this end, we employ Simulation-Based Inference (SBI), specifically Neural Ratio Estimation (NRE), a method well-suited for inference in high-dimensional spaces. Our results demonstrate more accurate constraints on the astrophysical parameters, highlighting the method's potential.

Decoding Galaxy Evolution Models with cINNs: Insights from SDSS-like Spectrophotometric Data

Nils Candebat

INAF- Arcetri Observatory

Galaxy parameters extraction remains challenging despite rich spectroscopic data. Here, we present a conditional invertible neural networks(cINN) for the recovery of stellar population parameters from spectrophotometric data. Our analysis utilizes a synthetic spectral library comprising 500,000 SDSS-like spectra, which represent composite stellar populations with a diversity of star-formation and chemical enrichment histories as well as dust attenuation. At the core of the study, lies the application of the cINN, an elegant framework able to estimate the Bayesian overall uncertainty derived during the inference process. The cinn processes both spectral and photometric data simultaneously, aiming to recover key physical parameters including metallicity, age, dust attenuation, mass-to-light ratio, and star formation history indicators (age10

The unknown unknowns: Representation Learning for simulation-based inference in galaxy formation and cosmology

Sebastian Trujillo-Gomez Heidelberg Institute for Theoretical Studies (HITS)

Despite the great increase in the level of detail of modern galaxy observations, a hundred years after Hubble we continue to describe galaxies using his famous diagram. With the arrival of deep learning, we have focused on training machines to predict galaxy morphology efficiently across enormous datasets but still using variants of Hubble's century-old classification. Despite its importance in the big data era, this approach is limited by our own biased intuitions and expectations. To address this, we are developing a new framework to let the machine learn unbiased information-maximizing representations of large complex datasets. As an example, I will show how these tools find low-dimensional representations of galaxy structure from next-generation galaxy surveys and cosmological simulations, effectively learning their own Hubble diagram from the data. These representations automatically approximate the likelihood for use in simulation-based inference and model selection without resorting to lossy summary statistics. Our framework is designed to be interpretable and provides a path towards breakthroughs in galaxy formation theory and cosmology.

Approximate Maximum Likelihood Estimation with Local Score Matching

Sherman Khoo

University of Bristol

We consider simulation-based inference for maximum likelihood estimation, i.e., when the model can be simulated but the likelihood function is not available in closed form. We propose a sequential, gradient-based optimisation method based on a novel local score matching technique, which uses simulator outputs around a localised region of the parameter iterates. This produces a fast, efficient, and flexible approximation to the gradient, and in particular, recovers a smoothed version of the likelihood objective, mitigating issues arising in the optimisation of complex likelihood landscapes. We establish theoretical guarantees for our gradient estimator and provide bounds on the bias introduced by the approximation. Finally, we demonstrate the empirical performance of our method applied to various simulator-based models, including a simplified Supernova data analysis problem.

The Star Formation History of Galaxy Populations with pop-cosmos

Sinan Deger

Cambridge University, Institute of Astronomy

The next generation observatories such as Vera C. Rubin and Euclid are posing a massive data challenge. An obstacle needed to overcome is the inference of accurate redshifts from photometric observations that can be limited to a handful of bands. We addressed this challenge with a forward modeling framework, pop-cosmos, calibrated by fitting a population model to observations on the photometry space. We perform a high-dimensional fitting to the 26-band photometry from COSMOS2020, complete with data-driven noise modeling and flexible selection effects, which is achieved via a novel use of simulation-based inference. Having been fitted to a deep representative survey, samples from our trained model become realizations of the galaxy population within the survey limiting magnitude and up to z=4, capturing 90% of cosmic time. Therefore pop-cosmos unlocks a medium for the study of galaxy evolution that was not possible before. I will be presenting population-level results, specifically a comprehensive look at the star formation histories of galaxies inferred from our model. This will entail a detailed picture of the stellar mass assembly of galaxies together with the evolution of their metallicity and colors.

A Physically Self-Consistent Forward Model for Galaxy SED Populations

Suchetha Cooray Stanford University

Inferring galaxy physical parameters from observed spectral energy distributions (SEDs) for calibrating galaxy formation models remains challenging due to wellknown degeneracies between dust attenuation, metallicity, and stellar population age. We present a novel hierarchical forward model that combines cosmological structure formation with galaxy-halo connection models to generate physically selfconsistent galaxy SEDs. This model builds on the UniverseMachine (Behroozi et al. 2019), where we model dark matter halo property-dependent star formation rates, stellar masses, dust, metallicity (stellar and gas-phase), and nebular properties to generate self-consistent galaxy SED populations through stellar population synthesis. We calibrate our model using the SDSS Main Galaxy Sample, and show that it can simultaneously reproduce observed one-point joint magnitude and color distributions and two-point clustering statistics. I will also present recent progress in calibrating UniverseMachine with the latest observational constraints, including JWST data, using SBI. These works provide a foundation for robust population-level inference of galaxy properties while maintaining cosmological consistency through the galaxy-halo connection.

Towards combined constraints on galaxy astrophysics, dust attenuation, and cosmology from galaxy luminosity functions and colors

Tjitske Starkenburg Northwestern University

Simulation-based inference is often used to perform cosmological parameter inference. However, observable tracers depend on detailed galaxy physics that is challenging to model to sufficient detail and includes approximations leading to differences between model predictions. Moreover, he properties of the observable population can be strongly affected by the dust distribution in galaxies. We present forward modeled photometry for a suite of large-scale dark matter only simulation with different cosmological parameters as well as variations in galaxy formation physics and a flexible model for dust attenuation. Each dark matter only simulation is combined with a number of variations of the Santa Cruz semi-analytic galaxy formation model. We incorporate additional flexibility in the forward modeling pipeline through including variations in the dust attenuation that are based on radiative transfer simulations. We compare the resulting photometry to the SDSS Main Galaxy Sample and perform implicit likelihood inference to find combined constraints on the galaxy formation physics, dust attenuation and cosmology. Our work highlights the importance of building on small-scale physics to infer large-scale parameters.

Author Index

Bevins Harry, 12 Bowden Haley, 11 Braspenning Joey, 13 Candebat Nils, 16 Cooray Suchetha, 18 Deger Sinan, 18 Fabbian Giulio, 11 Garcia Alex, 5 Garuda Nikhil, 15 Gerolymatou Dimitra, 6 Hartmann Eduardo A., 10 Iglesias Navarro Patricia, 2 Jespersen Christian Kragh, 9 Kartheik Iyer, 1 Khoo Sherman, 17 Lindberg

Christina, 10 Lovell Christopher, 2 Makinen Lucas, 14 Matthew Ho, 1 Matthews Alice, 5 Nikolić Ivan, 12 Pandey Shivam, 3 Pandya Viraj, 7 Pillepich Annalisa, 9 Ramnichal Steven, 7 Rose Jonah, 13 Sante Andrea, 6 Saoulis Alex, 8 Schaible Anna Lena, 8 Sorce Jenny, 13 Starkenburg Tjitske, 19 Tortorelli Luca, 14

Tri Nguyen, 4 Triantafyllou Nikolaos, 16 Trujillo-Gomez Sebastian, 17 Viraj

Pandya, 4

Viterbo Giuseppe, 11 von Wietersheim-Kramsta Maximilian, 1, 15 Will

Handley, 4 Wu John, 3