SBI4GALEV 2026: Summary White Paper

alan (automated rapporteur)

June 2026

Abstract

The Simulation-Based Inference for Galaxy Evolution (SBI4GALEV) 2026 conference convened researchers to address the computational and statistical challenges of connecting complex astrophysical simulations with next-generation observational data. The meeting highlighted a transition from the introductory phase of Simulation-Based Inference (SBI) toward high-performance implementation, emphasizing differentiable programming, generative AI, and the mitigation of the "simulation bottleneck." Discussions spanned scales from the interstellar medium to cosmological large-scale structure, establishing a roadmap for scalable, robust, and amortized inference pipelines.

1 Overview of Topics

The conference was characterized by a multidisciplinary approach to the inverse problem in astrophysics. The primary focus remained on **Simulation-Based Inference (SBI) and Likelihood-Free Inference (LFI)**, with Tom providing a foundational overview of neural estimation techniques and model validation. This methodological core supported extensive research into **Galaxy Evolution and Formation**, where Christian Kragh Jespersen, Subhrata Dey, and Jonah Rose explored the physical drivers of galaxy growth and morphology.

A significant portion of the program was dedicated to **Cosmological Inference** and **Bayesian Methods**, with Hiranya Peiris and Alex Saoulis discussing large-scale structure and weak lensing, while Johannes Buchner compared the efficiencies of HMC, Nested Sampling, and SBI. The integration of **Deep Learning and Generative AI** was pervasive, featuring work on diffusion models by Stephen Thorp, Lawrence Faria, and Yichen Zang, as well as GNNs and foundation models presented by Ming-Shau Liu and Eduardo A. Hartmann.

Further technical discussions focused on **Astrophysical Simulations** and **SED Modeling**, with Patricia Iglesias-Navarro, Amanda Stoffers, and Stephen Ramnichal presenting advanced frameworks for spectral analysis. The meeting also covered **Observational Constraints** (James Kostas Ray), the **Interstellar and Circumgalactic Medium** (Caterina Bracci, Tanmay Singh), **Multi-Fidelity Modeling** (Richard Stiskalek, M. von Wietersheim), and **Differentiable Programming** (Suchetha Cooray, Andrew Green). Niche applications included **Gravitational Waves** (Philippa Cole), **Galactic Archaeology** (Shichen Su), and **AI Infrastructure** (Grant Stevens, Sotiria Fotopoulou, Chris Lovell).

2 Hot Subjects and Emerging Trends

The dominant theme of the 2026 meeting was the shift toward **Generative AI and Diffusion Models**. Rather than simple regression, researchers like Grégoire Aufort and Yichen Zang are utilizing score-based diffusion for priors and inpainting to handle missing data in SED fitting and light curves.

Differentiable Forward Modeling emerged as a critical efficiency driver. The adoption of JAX-based frameworks, such as Ceridwen (Amanda Stoffers) and Tengri (Suchetha Cooray), allows for automatic differentiation (AD) and gradient-based sampling, significantly reducing the wall-clock time for high-dimensional inference.

Amortized Inference Scaling has moved from proof-of-concept to production. The development of tools like Synference (Thomas Harvey) and SBIPIX (Patricia Iglesias-Navarro) demonstrates the ability to perform pixel-level or population-scale inference on billions of galaxies, bypassing the per-object cost of traditional MCMC.

Finally, **Robustness and Model Misspecification** became a central concern. Sherman Khoo and Arya Farahi presented methods using Maximum Mean Discrepancy (MMD) and generalized Bayes frameworks to ensure that neural estimators remain calibrated even when simulations deviate from real-world observations.

3 Main Topics and Tools Occurrence

The following table summarizes the frequency of key concepts and software mentioned across the conference transcripts and slides.

4 Comparison with Previous Year

Comparing the 2026 program to the 2025 abstract booklet reveals a clear evolution from *introduction* to *implementation*. While the 2025 meeting focused on basic tutorials (e.g., "Introduction to SBI," "A Brief Introduction to JAX"), the 2026 meeting centered on specific, high-performance frameworks like Ceridwen, Synference, and OASIS.

Methodologically, there is a marked shift in generative modeling. 2025 emphasized Invertible Neural Networks (cINNs) and basic Normalizing Flows; 2026 is dominated by Diffusion Models and Flow-Matching. Furthermore, the scope of SED fitting has expanded from integrated properties to 3D resolved modeling and pixel-level inference. Finally, the 2026 meeting introduced a stronger focus on robustness and the "simulation bottleneck," moving beyond the question of *how* to run SBI to *how to make it reliable* under model misspecification.

5 Outlook for the Future

The community is moving toward a unified, differentiable pipeline where the simulator and the inference engine are tightly integrated. The proposed community data challenge led by Luca and Chris Lovell will be pivotal in benchmarking these tools against realistic, misspecified data. Future efforts will likely focus on "foundation models" for galaxy spectra and the deployment of ambient, self-hosted AI (as discussed by Sotiria Fotopoulou and Chris Lovell) to accelerate academic

Table 1: Main topics and tools and their occurrence counts

Topic/Tool	Transcript mentions	Slide mentions
Topics/Concepts		
Spectral Energy Distribution Fitting	0	18
Neural Posterior Estimation	3	10
Diffusion Models	6	6
Normalizing Flows	7	0
Hamiltonian Monte Carlo	5	5
Maximum Mean Discrepancy	3	6
Neural Likelihood Estimation	3	5
Nested Sampling	4	2
Automatic Differentiation	1	4
Neural Ratio Estimation	1	2
Simulation-Based Calibration	1	2
Tools/Software		
sbi	24	41
Synthesizer	18	2
Synference	0	16
DREAMS	5	11
Tengri	0	14
IllustrisTNG	0	8
Bagpipes	7	0
Cloudy	6	0
OASIS	2	4
UNiverse Machine	2	3
L-Galaxies	0	3
Ceridwen	0	3
JAX	0	2
Open WebUI	0	1
Other (SIMBA, pop-cosmos, etc.)	$\sim$1	$\sim$10

synthesis and education. The ultimate goal remains the scaling of these methods to Stage-IV surveys, enabling the transition from fitting individual galaxies to performing field-level cosmological inference.