PLANT: Conditional Generative Models for Fast Gravitational-Wave Population Synthesis

Rhea Senthil Kumar, Floor S. Broekgaarden (PI)
Gravitational Wave Paleontology Group · University of California, San Diego · rhsenthilkumar@ucsd.edu
The Gravitational Wave Paleontology Lab

Introduction

PLANT (**P**opulation synthesis with **L**earned **A**strophysical ge**N**erative models for gravi**T**ational-wave populations) uses neural generative models to emulate GW merger populations orders of magnitude faster than traditional simulation.

- LIGO-Virgo-KAGRA have detected 390 compact-binary mergers in the gravitational-wave transient catalog (GWTC-5.0; LVK Collaboration 2026).
- Linking mergers to star formation requires costly population-synthesis runs governed by the cosmic star formation rate (SFR) (Madau & Dickinson 2014):

$$\psi(z) \;=\; a\,\frac{(1+z)^{b}}{1+\left(\frac{1+z}{c}\right)^{d}} \quad [\,M_\odot\ \mathrm{yr}^{-1}\ \mathrm{Gpc}^{-3}\,]$$

- **Forward model:** COMPAS binary evolution + cosmic integration (SSPC; Neijssel et al. 2019) maps astrophysical hyperparameters $\mathbf{\Lambda}$ to merger catalogs $(m_{\mathrm{chirp}}, q, z)$.
- **Bottleneck:** each $\mathbf{\Lambda}$ needs $\sim 10^6$–10^7 binaries (many CPU-hours per point; Barrett et al. 2017). Hierarchical inference or dense grids require $\sim 10^3$–10^4 computationally intractable forward evaluations (Team COMPAS et al. 2022).

Question: Can conditional generative models learn to produce realistic GW merger catalogs directly from astrophysical parameters, bypassing expensive population-synthesis runs?

Training Data

The SSPC (Synthetic Stellar Population Convolve) framework convolves COMPAS binaries with the Madau-Dickinson SFR and Neijssel+19 metallicity model. We train on COMPAS (Riley+ 2022): 3.35×10^6 DCO systems across SMT, CE, CHE channels. We grid over a_{SFR} and μ_0 (two parameters that strongly drive predictions) as a feasibility test, sampling the other 7 nuisance parameters stochastically. For each grid point we compute the intrinsic merger rate:

$$\frac{dN}{dz_m} = \int_0^{z_{\max}} \psi(z_f;\mathbf{\Lambda})\, p(Z \mid z_f;\mathbf{\Lambda})\, p(\tau)\, \frac{dV_c}{dz_f}\, dz_f$$

weighting each binary by its formation rate and delay time τ.

- $a_{\mathrm{SFR}} \in [10^{-6},\ 0.092]$, $\mu_0 \in [10^{-6},\ 0.061]$ (van Son+ 2022, TNG100 $\pm 1\sigma$)
- 7 nuisance params $(b, c, d,\ \mu_z,\ \sigma_0,\ \sigma_z,\ \alpha_{\mathrm{skew}})$ sampled per cell
- 50×50 × 3 channels = 7,500 cells; 50K events each → 375M mergers
- Observables: $\mathcal{M}_c$, q, z; $\mathbf{\Lambda}$ (12-D) conditioning

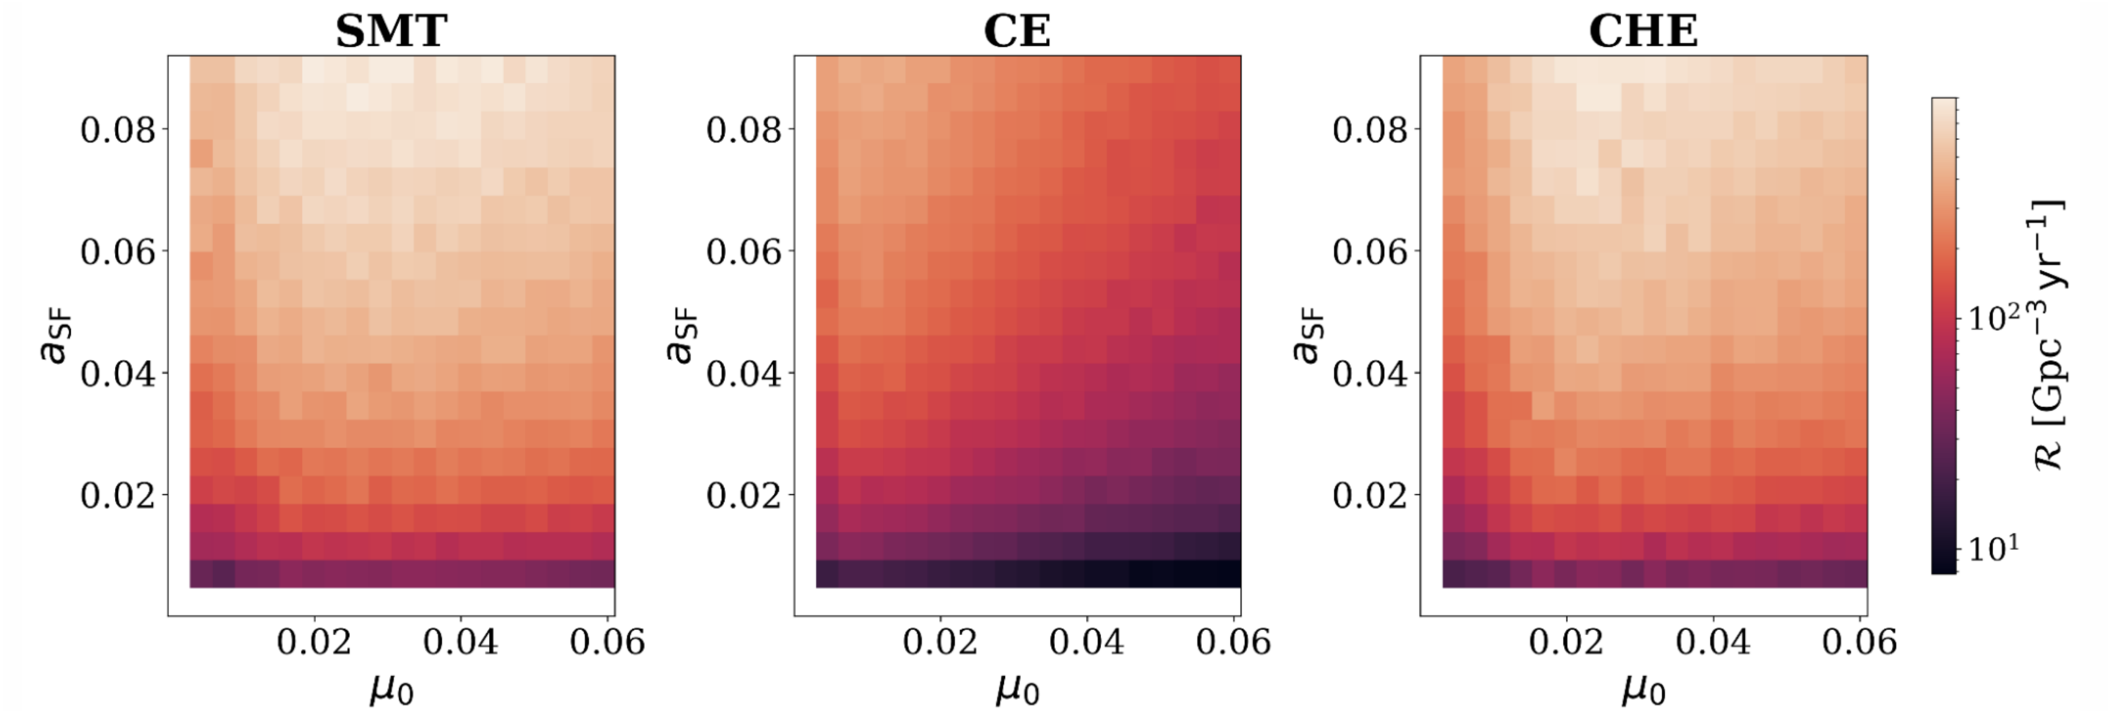

Figure 1: Merger-rate density across the $(a_{\mathrm{SFR}},\, \mu_0)$ grid.

Methods: Generative Models

Both models learn $p(\mathcal{M}_c, q, z \mid \mathbf{\Lambda})$ with a shared conditioning encoder.

Conditional Flow Matching (CFM)

- Deterministic ODE velocity field, optimal-transport paths (Lipman+ 2023)
- $\mathcal{L} = \mathbb{E}\big\|v_\theta - (\mathbf{x}_1 - \mathbf{x}_0)\big\|^2$
- Single forward pass at inference

Score-Based Diffusion

- Iterative denoising, linear schedule, $T\!=\!100$ steps (Ho+ 2020)
- $\mathcal{L} = \mathbb{E}_{t,\epsilon}\big\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{\Lambda}) - \epsilon\big\|^2$
- Strong mode coverage for multimodal targets

Baseline: Naive Bayes

- Per-cell Gaussians, kernel-weighted mix: $\pi_g \propto \exp(-\|\Lambda - \Lambda_g\|^2/2\tau^2)$
- No training, statistical estimate. Lower bound on fidelity.

Methods: Pipeline

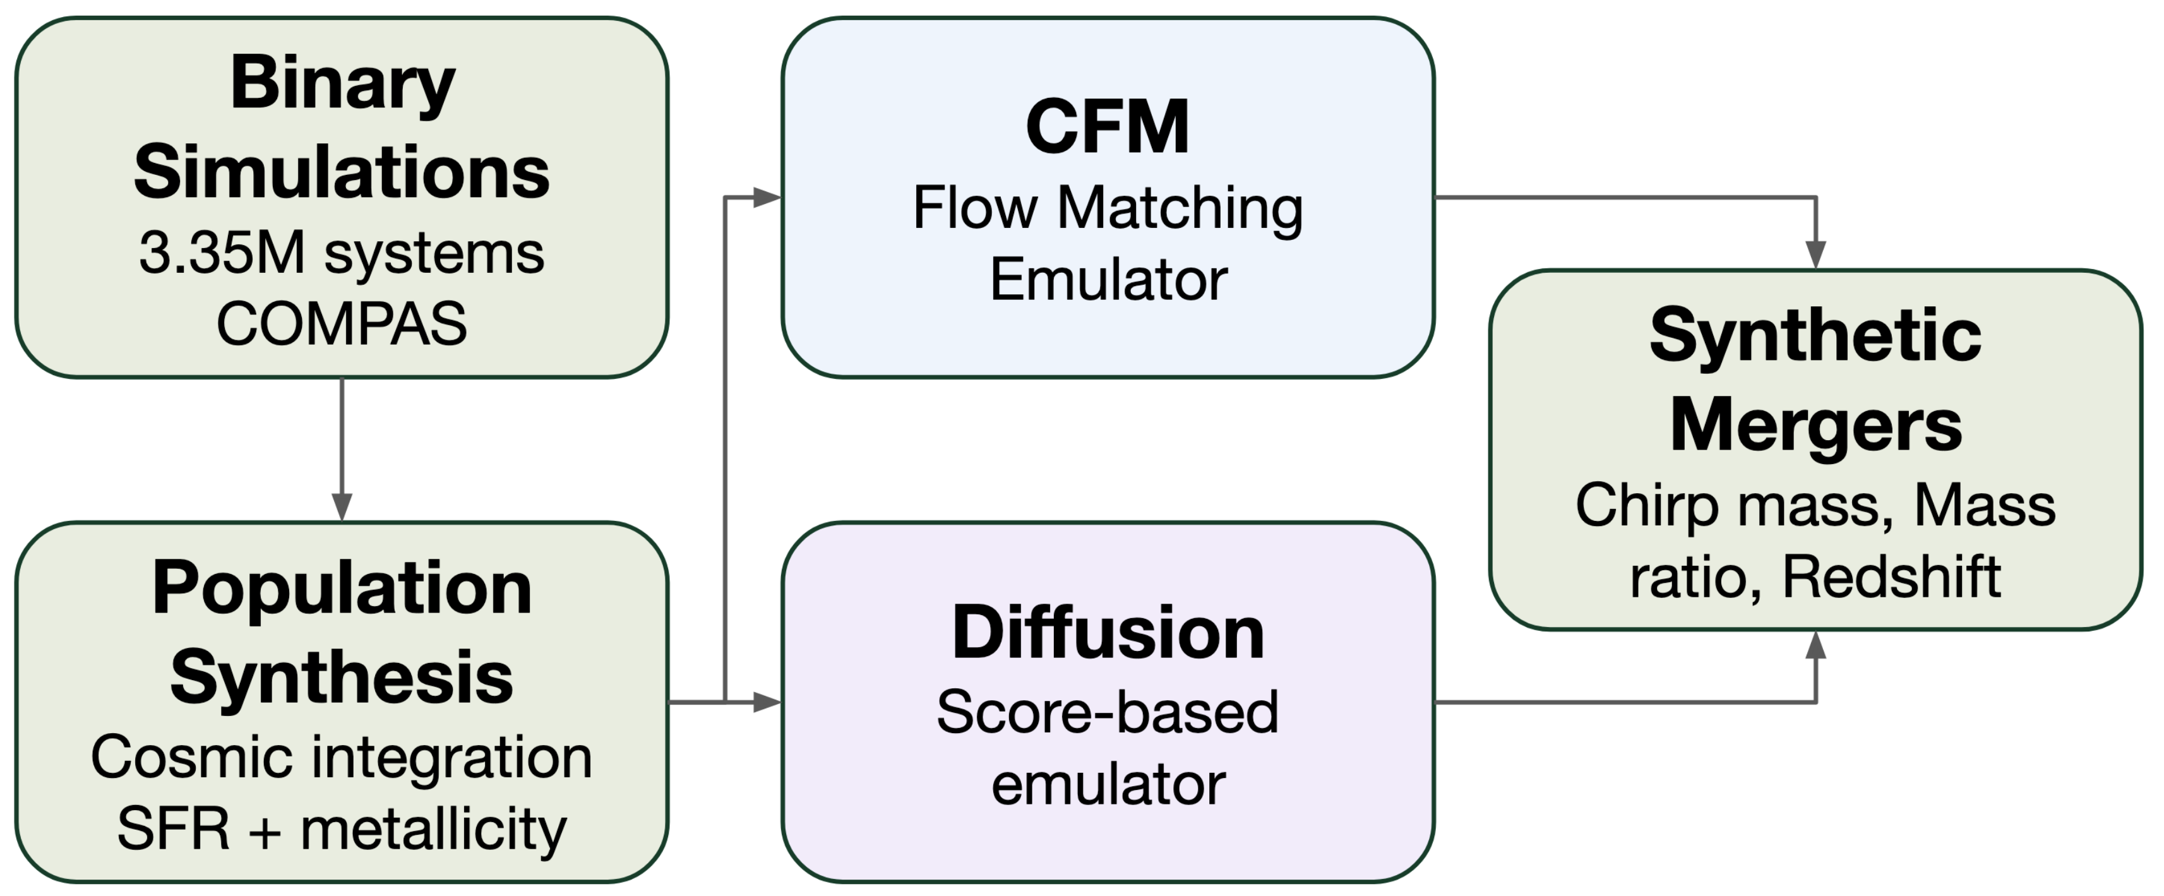

Binary Simulations
3.35M systems
COMPAS

Population Synthesis
Cosmic integration
SFR + metallicity

CFM
Flow Matching
Emulator

Diffusion
Score-based
emulator

Synthetic Mergers
Chirp mass, Mass ratio, Redshift

Figure 2: PLANT pipeline, from binary simulations to fast catalog generation.

Methods: Architecture

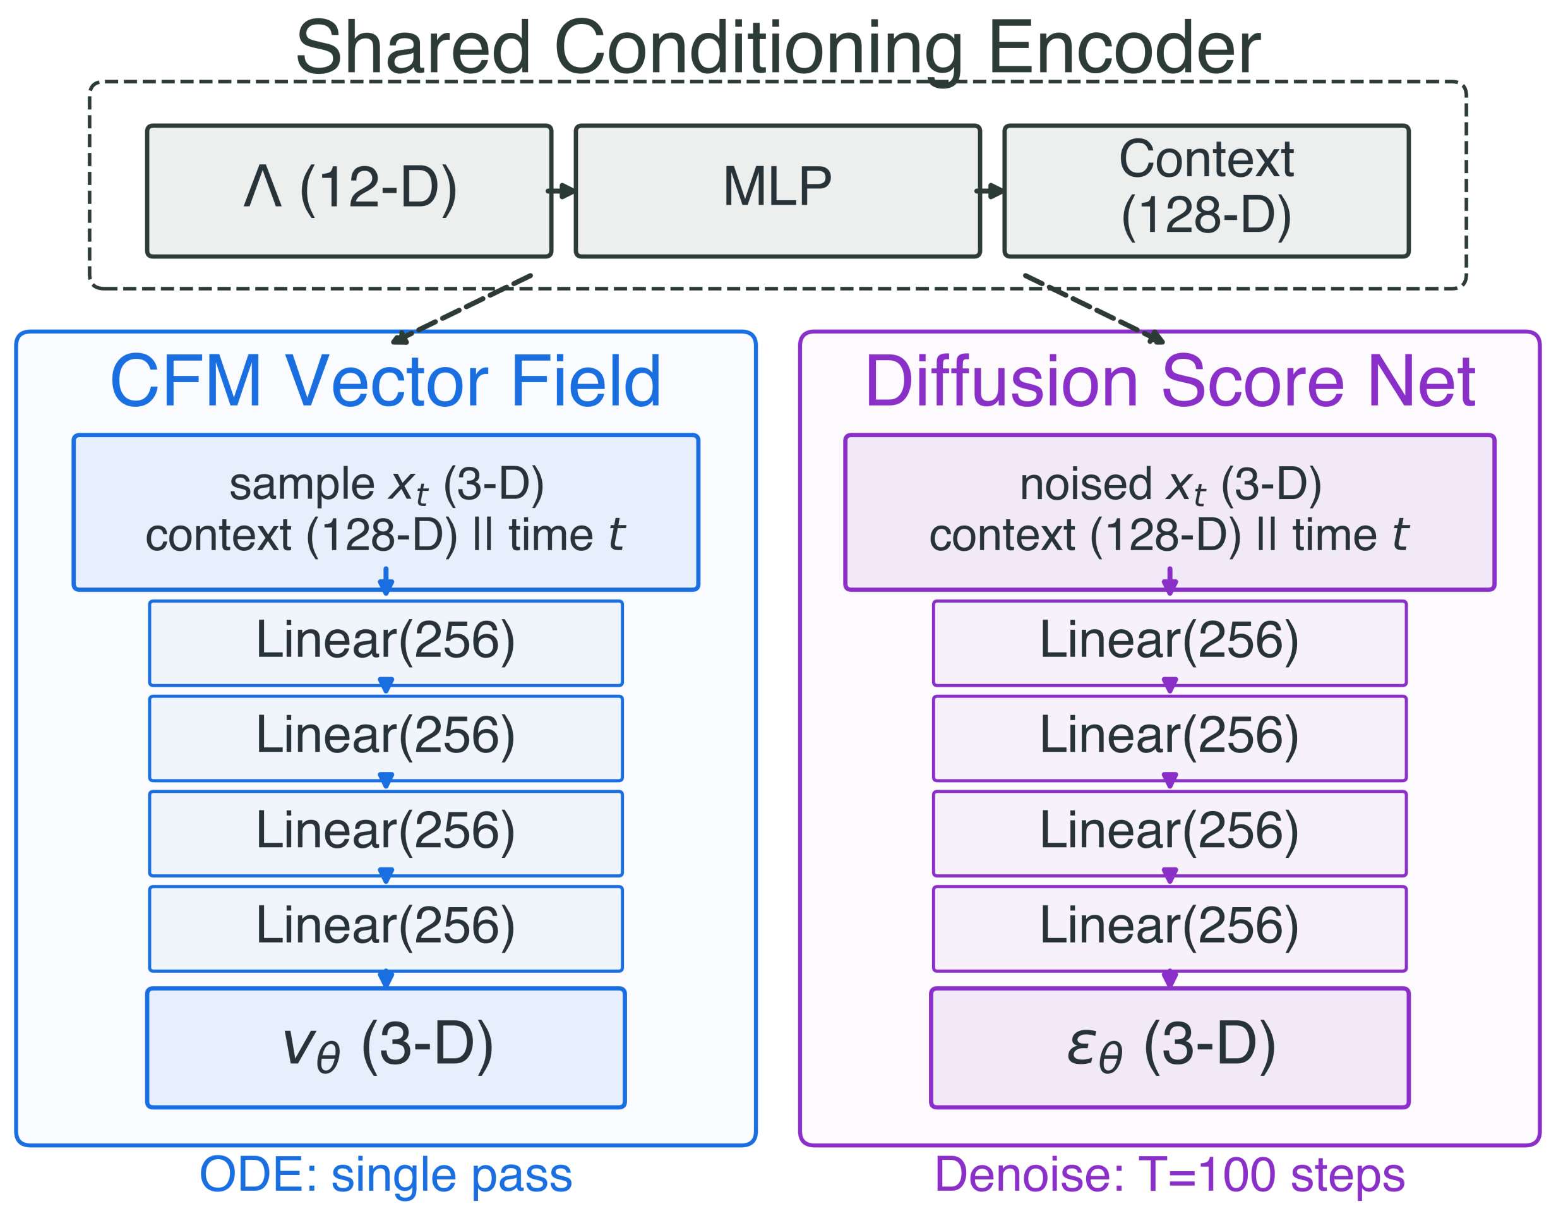

Shared Conditioning Encoder
Λ (12-D)
MLP
Context (128-D)

CFM Vector Field
sample x_t (3-D)
context (128-D) ll time t
Linear(256)
Linear(256)
Linear(256)
Linear(256)
v_θ (3-D)
ODE: single pass

Diffusion Score Net
noised x_t (3-D)
context (128-D) ll time t
Linear(256)
Linear(256)
Linear(256)
Linear(256)
ε_θ (3-D)
Denoise: T=100 steps

Figure 3: Shared encoder feeds both generation heads.

Results

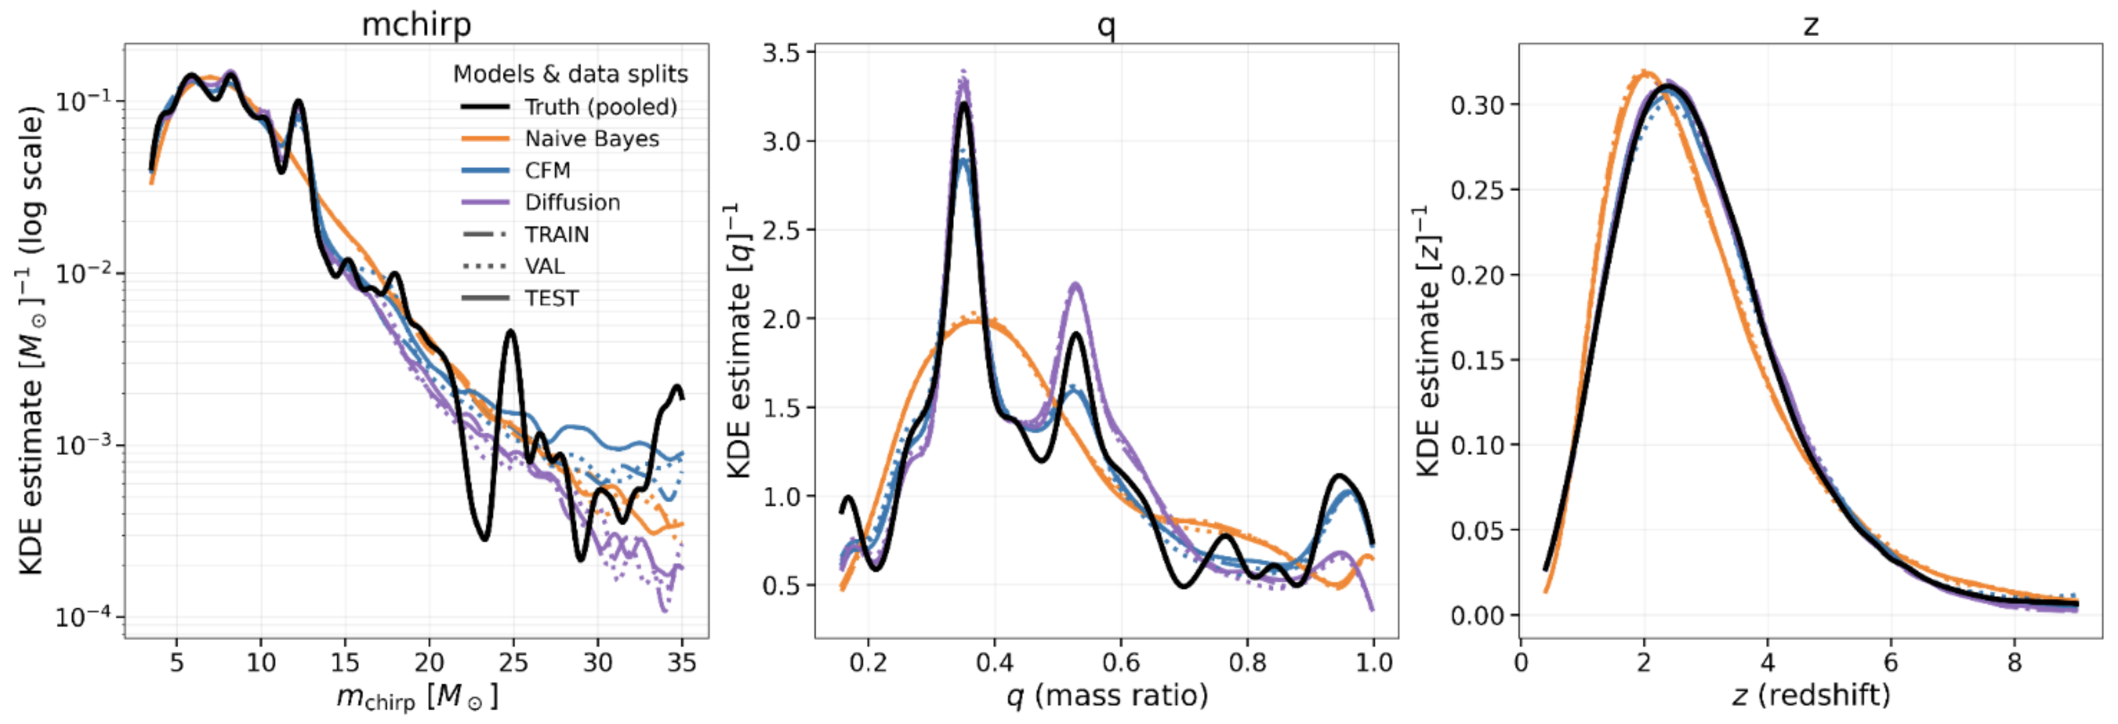

Figure 4: Emulator vs. ground truth marginals $(\mathcal{M}_c,\, q,\, z)$ on held-out grid points.

On held-out settings, CFM and Diffusion reproduce the bulk of intrinsic merger distributions; all emulators struggle on the high-$\mathcal{M}_c$ tail, which is sparse.

W_1	$\mathcal{M}_c\ [M_\odot]$	q	z
CFM	**0.12**	**0.005**	**0.04**
Diffusion	0.40	0.021	0.07
Naive Bayes	0.33	0.016	0.15

- CFM outperforms other models on Wasserstein metric across all three observables
- Naive Bayes is fast (CPU, no backprop) but factorized Gaussians miss q structure
- All emulators 300–500× faster than SSPC per catalog

Conclusions & Next Steps

- Demonstrated feasibility of generative models on selected star formation parameter space
- Both models reproduce multimodal merger distributions across all three formation channels
- Opens up real-time parameter exploration previously bottlenecked by simulation cost

Next steps:

- Finetune hyperparameters of generative models
- Add χ_{eff}; scale to full 9-D SSPC grid
- Selection effects and measurement uncertainty
- Ensemble of emulators for epistemic uncertainty
- Apply to real LVK O4/O5 data

References

[1] Madau & Dickinson 2014, ARAA 52, 415
[2] Neijssel+ 2019, MNRAS 490, 3740
[3] Broekgaarden+ 2022, MNRAS 516, 5737
[4] Team COMPAS (Riley+) 2022, ApJS 258, 34
[5] Lipman+ 2023, ICLR; Ho+ 2020, NeurIPS
[6] Taylor & Gerosa 2018, PRD 98, 083017
[7] COMPAS data: zenodo.org/records/4947741
[8] SSPC: davidhendriks.com/code/sspc.html
[9] SDSC Expanse: sdsc.edu/systems/expanse

Acknowledgements: I thank Prof. Floor Broekgaarden for her mentorship and guidance throughout this project, and Sandip Roy and Aobo Li for valuable discussions and collaboration. This work used the SSPC framework and publicly available COMPAS simulation data. Computational resources were provided by SDSC Expanse through ACCESS. Grateful to the AstroAI Workshop organizers for the opportunity to present.