Hybrid Diffusion-Flow-Matching Models for Mixed Numerical-Categorical Data Generation

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 9 peer-reviewed papers addressing the following research question: Do hybrid diffusion-flow-matching models outperform single-method baselines in terms of log-likelihood on mixed numerical-categorical datasets (e.g., California Housing or FICO) while maintaining. 16 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.4/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Flow Matching for Scalable Simulation-Based Inference. Research question: Do hybrid diffusion-flow-matching models outperform single-method baselines in terms of log-likelihood on mixed numerical-categorical datasets (e.g., California Housing or FICO) while maintaining computational efficiency at scale?.

2 Methodology

Systematic literature search across multiple databases yielded 9 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.4/10.

3 Results

9 papers retrieved. 16 claims extracted; 0 independently verified. Quality review score: 2.4/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The NPE baseline uses an embedding network to compress input x to a 128-dimensional feature vector.	×	0.02
The NPE embedding network consists of a learnable linear layer initialized with principal components of GW simulations f	×	0.01
FMPE modifies the baseline architecture by providing conditioning on (t, θ) via gated linear units in each hidden layer.	×	0.02
FMPE changes the dimension of the final feature vector to the dimension of θ to parameterize the conditional vector fiel	×	0.02
Both NPE and FMPE networks were trained with 5 $\cdot$ 10^6 simulations for 400 epochs using a batch size of 4096 on an A100 G	×	0.03
The FMPE network contains 1.9 $\cdot$ 10^8 learnable parameters.	×	0.02
Training the FMPE network takes approximately 2 days.	×	0.03
The NPE network contains 1.3 $\cdot$ 10^8 learnable parameters.	×	0.01
Training the NPE network takes approximately 3 days.	×	0.01
FMPE achieves a mean Jensen-Shannon divergence (JSD) of 0.5 mnat on the GW150914 evaluation.	×	0.04
NPE achieves a mean Jensen-Shannon divergence (JSD) of 3.6 mnat on the GW150914 evaluation.	×	0.03
FMPE achieves a maximum JSD of less than 2.0 mnat.	×	0.06
A maximum JSD of less than 2.0 mnat is an indistinguishability criterion for GW posteriors.	×	0.00
The evaluation omits the parameters c, JL, and θJN due to phase marginalization in importance sampling.	×	0.02
The GNPE results reported are from reference [7] and were generated with slightly different data conditioning.	×	0.03
The JSD for the tc parameter is not reported in reference [7] due to a tc marginalized reference.	×	0.02

References

- http://arxiv.org/abs/2305.17161v2
- http://arxiv.org/abs/2307.03672v3
- http://arxiv.org/abs/2502.17119v2