Novelty Assessment Report

Paper: Neural Posterior Estimation with Latent Basis Expansions

PDF URL: https://openreview.net/pdf?id=jsPQFNmnln

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Neural posterior estimation (NPE) is a likelihood-free amortized variational inference method that approximates projections of the posterior distribution. To date, NPE variational families have been either simple and interpretable (such as the Gaussian family) or highly flexible but black-box and potentially difficult to optimize (such as normalizing flows). In this work, we parameterize variational families via basis expansions of the latent variables. The log density of our variational distribution is a linear combination of latent basis functions (LBFs), which may be fixed a priori or adapted to the problem class of interest. Our training and inference procedures are computationally efficient even for problems with high-dimensional latent spaces, provided only a low-dimensional projection of the posterior is of interest, owing to NPE's automatic marginalization capabilities. In numerous inference problems, the proposed variational family exhibits better performance than existing variational families used with NPE, including mixtures of Gaussians (mixture density networks) and normalizing flows, as well as outperforming an existing basis expansion method for variational inference.

Core Task Landscape

This paper addresses: **Likelihood-Free Amortized Posterior Inference with Basis Expansions**

A total of **12 papers** were analyzed and organized into a taxonomy with **11 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Amortized Neural Posterior Estimation Methods**
- **Spectral and Basis Function Approximation for Likelihoods**
- **Gaussian Approximation and Moment Matching Methods**
- **Spatial Process Models with Basis Representations**
- **Unbiased Estimation for Intractable Models**

Complete Taxonomy Tree

- Likelihood-Free Amortized Posterior Inference with Basis Expansions Survey Taxonomy
- Amortized Neural Posterior Estimation Methods
  - Latent Basis Expansion Approaches ★ (2 papers)
  - [0] Neural Posterior Estimation with Latent Basis Expansions (Anon et al., 2026) View paper
  - [10] Neural Amortization of Bayesian Point Estimation (Y Li, 2024) View paper
  - Deep Learning Variational Inference (1 papers)
  - [4] Advances in Amortized Bayesian Inference, with Applications to Astronomy (McNamara, 2025) View paper
  - Simulation-Based Spatial Inference (1 papers)
  - [12] SpatFormer: simulation-based inference with transformers for spatial statistics (H Tesso, n.d.) View paper
- Spectral and Basis Function Approximation for Likelihoods
  - Spectral Likelihood Expansions (1 papers)
  - [11] Spectral likelihood expansions for Bayesian inference (Bruno, 2022) View paper
  - Wavelet-Based Approximate Bayesian Computation (1 papers)
  - [8] Estimating parameters in complex systems with functional outputs: A wavelet-based approximate Bayesian computation approach (Hongxiao Zhu, 2018) View paper
  - Radial Basis Function Approximation (1 papers)
  - [2] Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation (Nikolay Bliznyuk, 2008) View paper
- Gaussian Approximation and Moment Matching Methods
  - Variational Gaussian Approximation for Non-Gaussian Likelihoods (1 papers)
  - [5] Likelihood approximations via Gaussian approximate inference (Bui, 2024) View paper
  - Message Passing for Signal Tracking (1 papers)
  - [6] Clutter Tracking Using Variational Message Passing (Anders Malthe Westerkam, 2025) View paper
- Spatial Process Models with Basis Representations
  - Semiparametric Spatial Autoregressive Models (1 papers)
  - [3] A Semiparametric Bayesian Approach to Heterogeneous Spatial Autoregressive Models (Ting Liu, 2024) View paper
  - Point Process Intensity Modeling (2 papers)
  - [1] Approximation of Bayesian Hawkes process with inlabru (Francesco Serafini, 2022) View paper
  - [7] Modelling for Poisson process intensities over irregular spatial domains (Zhao, 2021) View paper
- Unbiased Estimation for Intractable Models (1 papers)
  - [9] Unbiased Monte Carlo: Posterior estimation for intractable/infinite-dimensional models (Sergios Agapiou, 2014) View paper

Narrative

Core task: likelihood-free amortized posterior inference with basis expansions. The field addresses settings where likelihood evaluations are intractable or prohibitively expensive, yet one wishes to perform Bayesian inference efficiently across many observed datasets. The taxonomy organizes work into several main branches. Amortized Neural Posterior Estimation Methods train neural networks to map observations directly to posterior distributions, enabling fast inference at test time without repeated MCMC runs. Spectral and Basis Function Approximation for Likelihoods construct tractable surrogates by expanding likelihoods or log-likelihoods in orthogonal or radial basis functions, as seen in Spectral Likelihood Expansions[11] and Bayesian Calibration RBF[2]. Gaussian Approximation and Moment Matching Methods simplify inference by fitting Gaussian or low-rank approximations to posteriors, often via variational or Laplace techniques. Spatial Process Models with Basis Representations handle geospatial or point-process data using basis decompositions to manage high-dimensional latent fields, exemplified by Bayesian Hawkes inlabru[1] and Poisson Irregular Domains[7]. Finally, Unbiased Estimation for Intractable Models provides Monte Carlo schemes that remove bias when likelihoods are accessible only through simulation, as in Unbiased Monte Carlo[9].

A particularly active line of work explores how neural amortization can be combined with latent basis expansions to achieve both flexibility and computational speed. Neural Posterior Latent Basis[0] sits squarely in this cluster, learning a low-dimensional basis representation within the amortized inference pipeline. This approach contrasts with purely spectral methods like Spectral Likelihood Expansions[11], which predefine basis functions analytically, and with standard neural estimators such as Neural Amortization Point[10], which may not explicitly exploit structured basis decompositions. Meanwhile, spatial models like SpatFormer[12] and Semiparametric Spatial Autoregressive[3] also leverage basis functions but focus on domain-specific priors rather than general amortization. A central open question is how to balance the expressiveness of learned bases against the interpretability and computational guarantees offered by fixed spectral or radial basis schemes, especially when scaling to high-dimensional parameter spaces or complex observation models.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Neural Amortization of Bayesian Point Estimation

Authors: Y Li, G ClartÃ© | **Year/Venue**: 2024 | **URL**: View paper

Abstract

â¦ to approximate the posterior distribution without requiring the evaluation of the likelihood â¦ function f over a set X can be decomposed into an element-wise transformation Ï followed by a â¦

Relationship Analysis

Both papers belong to the latent basis expansion approaches category, using neural networks to parameterize posterior approximations in likelihood-free settings. The original paper (Neural Posterior Estimation with Latent Basis Expansions) focuses on variational inference by parameterizing log-densities through basis function expansions in latent space, targeting full posterior distributions with exponential family representations. The candidate paper (Neural Amortization of Bayesian Point Estimation) differs fundamentally by targeting point estimates (mean, median, mode) rather than full posteriors, using an encoder-decoder architecture with a flexible loss function parameter α to control which statistic is estimated, rather than representing the entire posterior distribution.

Contributions Analysis

Overall novelty summary. The paper proposes a variational family for neural posterior estimation that parameterizes the log density as a linear combination of latent basis functions, either fixed or adapted to the problem class. Within the taxonomy, it resides in the 'Latent Basis Expansion Approaches' leaf under 'Amortized Neural Posterior Estimation Methods'. This leaf contains only two papers total, indicating a relatively sparse research direction. The sibling paper in this leaf represents the only other work explicitly combining neural amortization with latent basis expansions, suggesting the approach occupies a niche intersection between structured basis methods and flexible neural inference.

The taxonomy reveals several neighboring directions that contextualize this work. The sibling category 'Deep Learning Variational Inference' houses amortized methods without explicit basis expansions, while 'Spectral and Basis Function Approximation for Likelihoods' contains non-neural basis methods like orthogonal polynomial expansions and radial basis surrogates. The paper bridges these areas by embedding basis expansions within neural amortization, contrasting with purely spectral approaches that predefine bases analytically and with black-box neural flows that lack interpretable structure. The taxonomy's scope notes clarify that methods without explicit basis expansions belong elsewhere, positioning this work at a distinct methodological boundary.

Among 25 candidates examined, the contribution-level analysis shows mixed novelty signals. The core LBF-NPE variational family examined 10 candidates with zero refutable prior work, suggesting this specific parameterization is relatively unexplored. The computational efficiency claim examined 10 candidates and found one potentially overlapping result, indicating some prior work addresses efficient inference for low-dimensional projections. The convex optimization formulation examined 5 candidates with no refutations. Overall, the limited search scope (25 papers, not exhaustive) reveals that while the basis expansion parameterization appears novel, the efficiency advantages may have partial precedent in the examined literature.

Based on the top-25 semantic matches and taxonomy structure, the work appears to occupy a genuinely sparse research area where neural amortization meets structured basis representations. The single sibling paper and absence of refutations for the core variational family suggest meaningful novelty, though the computational efficiency contribution shows some overlap. The analysis does not cover the full literature landscape, and a broader search might reveal additional related work in adjacent fields like kernel methods or functional data analysis that were not captured by the semantic search.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Latent Basis Function NPE (LBF-NPE) variational family

Description: The authors introduce a new variational family for neural posterior estimation where the log density is expressed as a linear combination of basis functions over the latent space. This exponential family parameterization can use either fixed basis functions (such as B-splines or wavelets) or adaptively learned basis functions fitted jointly with the inference network.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Incremental variational sparse Gaussian process regression

URL: View paper

Brief Assessment

Incremental Sparse Gaussian[35] focuses on incremental variational inference for Gaussian process regression using basis functions in RKHS, not neural posterior estimation with basis expansions over latent spaces for likelihood-free inference.

---

2. Fast Bayesian Basis Selection for Functional Data Representation with Correlated Errors

URL: View paper

Brief Assessment

Fast Bayesian Basis[33] focuses on basis function selection for functional data representation with correlated errors using variational EM, not neural posterior estimation with basis expansions over latent spaces for likelihood-free inference.

3. Shape statistics in kernel space for variational image segmentation

URL: View paper

Brief Assessment

Shape Statistics Kernel[34] focuses on kernel-based shape statistics for image segmentation, not neural posterior estimation with basis function expansions for variational inference. The candidate addresses a completely different problem domain (computer vision/ segmentation) rather than likelihood-free inference.

4. Nonparametric variational inference

URL: View paper

Brief Assessment

Nonparametric Variational Inference[31] uses a mixture of Gaussians with kernel density estimation for variational inference, not basis function expansions over latent space for neural posterior estimation. The candidate focuses on general variational inference without the NPE framework or basis function parameterization of log densities.

5. Functional variational inference based on stochastic process generators

URL: View paper

Brief Assessment

Functional Variational Stochastic[32] focuses on stochastic process generators (SPGs) for function-space variational inference, not basis function expansions for neural posterior estimation in latent space.

6. A Scalable Variational Bayes Approach for Fitting Non-Conjugate Spatial Generalized Linear Mixed Models via Basis Expansions

URL: View paper

Brief Assessment

Scalable Variational Spatial[37] focuses on spatial generalized linear mixed models using basis expansions for spatial random effects in a Bayesian hierarchical framework, not on neural posterior estimation with basis function parameterizations of variational densities for likelihood-free inference.

7. Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming

URL: View paper

Brief Assessment

Hilbert Gaussian Processes[30] focuses on basis function approximations for Gaussian processes in probabilistic programming, not neural posterior estimation with basis expansions over latent spaces for simulation-based inference.

8. Clustering functional data via variational inference

URL: View paper

Brief Assessment

Clustering Functional Variational[29] focuses on clustering functional data using variational inference with B-spline basis expansions for mean curves, not on neural posterior estimation with basis function parameterizations for log densities in simulation-based inference.

9. Variational Phase Estimation with Variational Fast Forwarding

URL: View paper

Brief Assessment

Variational Phase Forwarding[36] focuses on quantum phase estimation using time-evolved states for molecular Hamiltonians, not neural posterior estimation with basis function expansions for Bayesian inference.

10. EigenVI: score-based variational inference with orthogonal function expansions

URL: View paper

Brief Assessment

EigenVI[28] uses orthogonal function expansions (e.g., Hermite polynomials) for non-amortized variational inference, while LBF-NPE uses basis expansions specifically for amortized neural posterior estimation with automatic marginalization capabilities.

Contribution 2: Computationally efficient training and inference for low-dimensional posterior projections

Description: The method exploits NPE's automatic marginalization to efficiently handle high-dimensional latent spaces when only low-dimensional posterior projections are needed. This allows the approach to avoid modeling nuisance variables explicitly while maintaining computational tractability through numerical integration in the low-dimensional space of interest.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Understanding posterior projection effects with normalizing flows

URL: View paper

Brief Assessment

Posterior Projection Flows[15] focuses on computing posterior profiles (maximization over nuisance parameters) rather than marginals for interpretability in cosmology applications. The original paper's contribution concerns automatic marginalization in NPE for efficient handling of high-dimensional latent spaces when only low-dimensional projections are needed, which is a different technical approach and use case.

2. Lens Modeling of STRIDES Strongly Lensed Quasars Using Neural Posterior Estimation

URL: View paper

Brief Assessment

STRIDES Neural Posterior[13] applies NPE to lens modeling for cosmology but does not discuss exploiting automatic marginalization for computational efficiency in low-dimensional projections. The candidate focuses on astrophysical applications rather than methodological contributions to efficient inference.

3. Neural networks as optimal estimators to marginalize over baryonic effects

URL: View paper

Brief Assessment

Neural Baryonic Marginalization[19] focuses on cosmological parameter estimation from contaminated observational data, not on general-purpose neural posterior estimation frameworks with automatic marginalization capabilities for arbitrary Bayesian models.

4. Automatic Forward Model Parameterization with Bayesian Inference of Conformational Populations

URL: View paper

Brief Assessment

Forward Model Conformational[17] focuses on Bayesian inference for conformational populations in molecular systems with forward model parameter optimization, not on neural posterior estimation with automatic marginalization for general low-dimensional projections.

5. Field-level simulation-based inference of galaxy clustering with convolutional neural networks

URL: View paper

Brief Assessment

Field Level Galaxy[14] focuses on galaxy clustering inference using convolutional neural networks for field-level simulation-based inference. The minimal context provided does not demonstrate that this work addresses efficient handling of high-dimensional latent spaces with automatic marginalization in neural posterior estimation, which is the core novelty claim of the original contribution.

6. Collapsed Inference for Bayesian Deep Learning

URL: View paper

Brief Assessment

Collapsed Inference Deep[18] focuses on Bayesian neural networks using collapsed sampling with weighted model integration for volume computation. The original paper addresses neural posterior estimation with automatic marginalization for simulation-based inference, which is a fundamentally different inference framework and problem domain.

7. The case for Bayesian deep learning

URL: View paper

Brief Assessment

Bayesian Deep Learning[22] focuses on marginalization over neural network parameters for epistemic uncertainty quantification, not on efficient inference methods for high-dimensional latent spaces with automatic marginalization capabilities as in the original paper's NPE framework.

8. Truncated marginal neural ratio estimation

URL: View paper

Brief Assessment

Truncated Marginal Ratio[16] focuses on marginal posterior estimation through truncated priors and ratio estimation, not on automatic marginalization in neural posterior estimation frameworks. The candidate's efficiency comes from targeted simulation in truncated regions, while the original exploits NPE's inherent marginalization capabilities.

9. swyft: Truncated marginal neural ratio estimation in python

URL: View paper

Brief Assessment

swyft[20] focuses on truncated marginal neural ratio estimation for likelihood-to-evidence ratios in simulation-based inference, not on neural posterior estimation with automatic marginalization for high-dimensional latent spaces as described in the original paper.

10. Forward amortized inference for likelihood-free variational marginalization

URL: View paper

Prior Art Analysis

Forward Amortized Marginalization[21] demonstrates that efficient handling of high-dimensional latent spaces when only low-dimensional posterior projections are needed was achieved prior to the original paper. The candidate explicitly describes a method that marginalizes over nuisance variables without explicit modeling, using forward KL divergence in a likelihood-free setting. The theoretical framework (Theorem 2) proves that marginalization can be performed by simply ignoring nuisance variables during sampling, avoiding the need to model their joint density. This directly addresses the same computational efficiency challenge claimed as novel by the original paper.

Evidence

Evidence 1 - **Rationale**: Both papers claim computational efficiency for high-dimensional problems when only low-dimensional projections are needed through automatic marginalization, suggesting Forward Amortized Marginalization[21] achieved this capability earlier. - **Original**: our training and inference procedures are computationally efficient even for problems with high-dimensional latent spaces, provided only a low-dimensional projection of the posterior is of interest, owing to npe's automatic marginalization capabilities - **Candidate**: one of the most important features of favi is that it can be used for marginalizing over a large space of nuisance variables without explicitly modeling their joint density. marginalization of nuisance variables is important in many real-world problems such as weather forecasting

Evidence 2 - **Rationale**: The candidate's Theorem 2 formally proves that marginalization over nuisance variables ξ can be achieved without explicit modeling, matching the original's claim about automatic marginalization through discarding nuisance variables. - **Original**: when the generative model contains both parameters of interest and nuisance variables, npe can automatically marginalize over the nuisance parameters during training: by simulating complete data and then discarding the nuisance variables to create training pairs, the method infers posterior projecti... - **Candidate**: theorem 2 (consistent marginalization) . consider a joint distribution p(z,ξ,x) and the (nonparametric) conditionally independent variational model q(z,ξ |x) = qz(z |x)qξ(ξ |x). the following equality holds: ∫ arginf q lfa[p(z,ξ,x),q(z,ξ |x)]dξ = arginf qz lfa[p(z,x),qz(z|x)]

Evidence 3 - **Rationale**: Both papers emphasize avoiding explicit modeling of nuisance variables when only low-dimensional parameters are of interest, with Forward Amortized Marginalization[21] providing the theoretical foundation earlier. - **Original**: unlike standard variational inference, where nuisance variables must be modeled in the variational distribution, npe applications typically require posteriors over just a few scientifically relevant parameters - **Candidate**: therefore, there is no need to explicitly model the conditional dependencies between zand ξwhen the aim is to estimate qz(z|x). in practice, it is straightforward to obtain monte carlo estimates of the marginalized variational loss

Contribution 3: Convex optimization formulation with fixed basis functions

Description: The authors establish that when basis functions are fixed a priori, the resulting optimization problem is convex in the inference network parameters. This convexity property ensures stable convergence to global optima and addresses optimization difficulties that plague more flexible variational families like normalizing flows.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Variational Quantum Optimization with Multi-Basis Encodings

URL: View paper

Brief Assessment

Variational Quantum Multi-Basis[25] addresses quantum optimization for classical combinatorial problems using multi-basis graph encodings, not variational inference with fixed basis functions for posterior approximation.

2. Sparse Template-Based Variational Image Segmentation

URL: View paper

Brief Assessment

Sparse Template Segmentation[26] addresses image segmentation using shape templates with convex relaxation, not variational inference for posterior estimation. The technical domains and problem formulations are fundamentally different.

3. Variational ground states of two-dimensional antiferromagnets in the valence bond basis

URL: View paper

Brief Assessment

Variational Antiferromagnets Valence[27] addresses quantum spin systems using valence bond basis functions for ground state approximation, not variational inference with neural networks. The convexity discussed relates to energy functionals in quantum physics, not inference network parameters in machine learning.

4. Variational Methods for Inference, Filtering, and Control

URL: View paper

Brief Assessment

Variational Inference Control[23] focuses on control and filtering problems using variational methods. The provided context fragments do not contain sufficient detail about basis function parameterization or convexity properties to challenge the original paper's novelty claim regarding convex optimization with fixed basis functions for neural posterior estimation.

5. Bayesian Electromagnetic Spatio-Temporal Imaging of Extended Sources Based on Matrix Factorization

URL: View paper

Brief Assessment

Bayesian Electromagnetic Imaging[24] focuses on electromagnetic source imaging for E/MEG data using matrix factorization with temporal basis functions, not on general variational inference with fixed basis functions for neural posterior estimation.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Neural Posterior Estimation with Latent Basis Expansions View paper
- [1] Approximation of Bayesian Hawkes process with inlabru View paper
- [2] Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation View paper
- [3] A Semiparametric Bayesian Approach to Heterogeneous Spatial Autoregressive Models View paper
- [4] Advances in Amortized Bayesian Inference, with Applications to Astronomy View paper
- [5] Likelihood approximations via Gaussian approximate inference View paper
- [6] Clutter Tracking Using Variational Message Passing View paper
- [7] Modelling for Poisson process intensities over irregular spatial domains View paper
- [8] Estimating parameters in complex systems with functional outputs: A wavelet-based approximate Bayesian computation approach View paper
- [9] Unbiased Monte Carlo: Posterior estimation for intractable/infinite-dimensional models View paper
- [10] Neural Amortization of Bayesian Point Estimation View paper
- [11] Spectral likelihood expansions for Bayesian inference View paper
- [12] SpatFormer: simulation-based inference with transformers for spatial statistics View paper
- [13] Lens Modeling of STRIDES Strongly Lensed Quasars Using Neural Posterior Estimation View paper
- [14] Field-level simulation-based inference of galaxy clustering with convolutional neural networks View paper
- [15] Understanding posterior projection effects with normalizing flows View paper
- [16] Truncated marginal neural ratio estimation View paper
- [17] Automatic Forward Model Parameterization with Bayesian Inference of Conformational Populations View paper
- [18] Collapsed Inference for Bayesian Deep Learning View paper
- [19] Neural networks as optimal estimators to marginalize over baryonic effects View paper
- [20] swyft: Truncated marginal neural ratio estimation in python View paper
- [21] Forward amortized inference for likelihood-free variational marginalization View paper
- [22] The case for Bayesian deep learning View paper
- [23] Variational Methods for Inference, Filtering, and Control View paper
- [24] Bayesian Electromagnetic Spatio-Temporal Imaging of Extended Sources Based on Matrix Factorization View paper
- [25] Variational Quantum Optimization with Multi-Basis Encodings View paper

- [26] Sparse Template-Based Variational Image Segmentation View paper
- [27] Variational ground states of two-dimensional antiferromagnets in the valence bond basis View paper
- [28] EigenVI: score-based variational inference with orthogonal function expansions View paper
- [29] Clustering functional data via variational inference View paper
- [30] Practical Hilbert space approximate Bayesian Gaussian processes for probabilistic programming View paper
- [31] Nonparametric variational inference View paper
- [32] Functional variational inference based on stochastic process generators View paper
- [33] Fast Bayesian Basis Selection for Functional Data Representation with Correlated Errors View paper
- [34] Shape statistics in kernel space for variational image segmentation View paper
- [35] Incremental variational sparse Gaussian process regression View paper
- [36] Variational Phase Estimation with Variational Fast Forwarding View paper
- [37] A Scalable Variational Bayes Approach for Fitting Non-Conjugate Spatial Generalized Linear Mixed Models via Basis Expansions View paper