Novelty Assessment Report

Paper: Frozen Priors, Fluid Forecasts: Prequential Uncertainty for Low-Data Deployment with Pretrained Generative Models

PDF URL: https://openreview.net/pdf?id=3FCHmUPmhe

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Deploying ML systems with only a few real samples makes operational metrics (such as alert rates or mean scores) highly unstable. Existing uncertainty quantification (UQ) methods fail here: frequentist intervals ignore the deployed predictive rule, Bayesian posteriors assume continual refitting, and conformal methods offer per-example rather than long-run guarantees. We introduce a forecast-first UQ framework that blends the empirical distribution with a frozen pretrained generator using a unique Dirichlet schedule, ensuring time-consistent forecasts. Uncertainty is quantified via martingale posteriors: a lightweight, likelihood-free resampling method that simulates future forecasts under the deployed rule, yielding sharp, well-calibrated intervals for both current and long-run metrics without retraining or density evaluation. A single hyperparameter, set by a small-N minimax criterion, balances sampling variance and model--data mismatch; for bounded scores, we provide finite-time drift guarantees. We also show how this framework informs optimal retraining decisions. Applicable off-the-shelf to frozen generators (flows, diffusion, autoregressive models, GANs) and linear metrics (means, tails, NLL), it outperforms bootstrap baselines across vision and language benchmarks (WikiText-2, CIFAR-10, and SVHN datasets); e.g., it achieves $\sim$90\% coverage on GPT-2 with 20 samples vs.\ 37\% for bootstrap. Importantly, our uncertainty estimates are operational under the deployed forecasting rule agnostic of the population parameters, affording practicable estimators for deployment in real world settings.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Uncertainty Quantification for Operational Metrics with Frozen Pretrained Generators**

A total of **35 papers** were analyzed and organized into a taxonomy with **18 categories**.

Taxonomy Overview

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

- **Uncertainty Quantification Frameworks for Frozen Pretrained Models**
- **Generative Model-Based Uncertainty Estimation**
- **Uncertainty-Aware Learning with Pretrained Vision-Language and Multimodal Models**
- **Uncertainty-Driven Active Learning and Data-Efficient Training**
- **Domain-Specific Applications of Uncertainty Quantification with Pretrained Models**
- **Uncertainty Quantification via Data Augmentation and Synthetic Sample Generation**

Complete Taxonomy Tree

- Uncertainty Quantification for Operational Metrics with Frozen Pretrained Generators Survey Taxonomy
- Uncertainty Quantification Frameworks for Frozen Pretrained Models
  - Post-Hoc Uncertainty Estimation via Auxiliary Models ★ (3 papers)
  - [0] Frozen Priors, Fluid Forecasts: Prequential Uncertainty for Low-Data Deployment with Pretrained Generative Models (Anon et al., 2026) View paper
  - [4] Principled Input-Output-Conditioned Post-Hoc Uncertainty Estimation for Regression Networks (Bramlage, 2025) View paper
  - [12] BayesCap: Bayesian Identity Cap for Calibrated Uncertainty in Frozen Neural Networks (Upadhyay, 2022) View paper
  - Evidential and Meta-Learning Approaches for Uncertainty (2 papers)
  - [13] Integrating large pre-trained models into multimodal named entity recognition with evidential fusion (Liu, 2023) View paper
  - [28] Guided Uncertainty Learning Using a Post-Hoc Evidential Meta-Model (Barker, 2025) View paper
  - Pretrained Uncertainty Modules and Transfer Learning (2 papers)
  - [17] UGrafting: Uncertainty Quantification in Self-Supervised Learning (EstevÃ£o, 2024) View paper
  - [32] Pretrained Visual Uncertainties (Kirchhof, 2024) View paper
- Generative Model-Based Uncertainty Estimation
  - Diffusion and Score-Based Generative Uncertainty (3 papers)
  - [1] Quantifying generative model uncertainty in posterior sampling methods for computational imaging (C Ekmekci, 2023) View paper
  - [7] Estimating uncertainty in diffusion MRI models using generative deep learning (Frank Zijlstra, 2024) View paper
  - [10] Generative Uncertainty in Diffusion Models (Jazbec, 2025) View paper
  - Adversarial and Conditional Generative Uncertainty (3 papers)
  - [9] Estimation with uncertainty via conditional generative adversarial networks (Minhyeok Lee, 2021) View paper
  - [21] Application of conditional generative adversarial network in ground motion modelling encompassing epistemic uncertainty (Ravi Kanth Sriwastav, 2025) View paper
  - [24] Evaluating Aleatoric Uncertainty via Conditional Generative Models (Huang ZiYi, 2022) View paper
  - Bayesian and Variational Generative Frameworks (2 papers)
  - [16] Uncertainty Quantification with Generative Models (BÃ¶hm, 2022) View paper

- [34] Deep Generative Models for Stochastic Seismic Imaging and Uncertainty Quantification (Y. Xie, 2024) View paper
- Generative Emulation for Scientific and Physical Systems (3 papers)
- [11] Generative emulation and uncertainty quantification of geological CO2 storage with conditional diffusion models (Zhongzheng Wang, 2025) View paper
- [23] Scalable crystal structure relaxation using an iteration-free deep generative model with uncertainty quantification (Ziduo Yang, 2024) View paper
- [25] Generative modelling for mass-mapping with fast uncertainty quantification (Jessica J Whitney, 2025) View paper

- Uncertainty-Aware Learning with Pretrained Vision-Language and Multimodal Models
  - Vision-Language Model Uncertainty and Calibration (3 papers)
  - [3] Probabilistic Prototype Calibration of Vision-Language Models for Generalized Few-shot Semantic Segmentation (Liu Jie, 2025) View paper
  - [26] Toward More Reliable Artificial Intelligence: Reducing Hallucinations in Vision-Language Models (Kassoum Sanogo, 2025) View paper
  - [31] Towards Understanding and Quantifying Uncertainty for Text-to-Image Generation (Gianni Franchi, 2024) View paper
  - Multimodal Evidential Fusion and Integration (1 papers)
  - [6] AVadCLIP: Audio-Visual Collaboration for Robust Video Anomaly Detection (Wu Peng, 2025) View paper
  - Uncertainty in Language Models and Text Generation (1 papers)
  - [5] Distinguishing the Knowable from the Unknowable with Language Models (Qin Tian, 2024) View paper
- Uncertainty-Driven Active Learning and Data-Efficient Training
  - Active Learning with Pretrained Models (2 papers)
  - [8] Active Learning of non-Semantic Speech Tasks with Pretrained models (Harlin Lee, 2022) View paper
  - [29] Uncertainty and Generalizability in Foundation Models for Earth Observation (Ramos-PollÃ¡n, 2024) View paper
  - Semi-Supervised and Co-Training with Uncertainty (2 papers)
  - [15] Harmonizing Generalization and Specialization: Uncertainty-Informed Collaborative Learning for Semi-supervised Medical Image Segmentation (Wenjing Lu, 2025) View paper
  - [30] TRiCo: Triadic Game-Theoretic Co-Training for Robust Semi-Supervised Learning (He Hongyang, 2025) View paper
  - Uncertainty-Guided Anomaly Detection and Incremental Learning (1 papers)
  - [22] I Detect What I Don't Know: Incremental Anomaly Learning with Stochastic Weight Averaging-Gaussian for Oracle-Free Medical Imaging (Nand Kumar Yadav, 2025) View paper
- Domain-Specific Applications of Uncertainty Quantification with Pretrained Models
  - Industrial and Manufacturing Process Uncertainty (1 papers)
  - [2] Large Pre-Trained Models and Few-Shot FineTuning for Virtual Metrology: A Framework for Uncertainty-Driven Adaptive Process Control in Semiconductor â¦ (CY Lin, 2025) View paper
  - Medical and Biological Imaging Uncertainty (1 papers)
  - [27] Deep Generative Models for Anomaly Detection and Uncertainty Quantification in Biology and Medicine (Unknown, 2025) View paper
  - AI-Generated Content Detection and Verification (1 papers)
  - [33] CINEMAE: Leveraging Frozen Masked Autoencoders for Cross-Generator AI Image Detection (Minsuk Jang, 2025) View paper
  - Uncertainty in Planning and Robotic Perception (2 papers)
  - [18] Into the Unknown: Towards using Generative Models for Sampling Priors of Environment Uncertainty for Planning in Configuration Spaces (Bhattacharjee, 2025) View paper
  - [35] Deep Generative Models for Competency Awareness and Uncertainty Quantification (Acharya, 2023) View paper
- Uncertainty Quantification via Data Augmentation and Synthetic Sample Generation (3 papers)
  - [14] Uncertainty-Aware Deep Classifiers using Generative Models (M. Sensoy, 2022) View paper
  - [19] Novel Uncertainty Quantification Through Perturbation-Assisted Sample Synthesis (Liu Yifei, 2024) View paper
  - [20] Stochastic latent feature distillation: Enhancing dataset distillation via structured uncertainty modeling (Zhe Li, 2025) View paper

Narrative

Core task: uncertainty quantification for operational metrics with frozen pretrained generators. The field addresses how to estimate predictive uncertainty when leveraging large pretrained models without retraining them, a practical constraint in many deployment scenarios. The taxonomy reveals several complementary directions: some branches focus on general frameworks for post-hoc uncertainty estimation via auxiliary models or ensemble-like methods, while others emphasize generative model-based approaches that exploit the stochastic nature of diffusion models or GANs. A third cluster explores uncertainty-aware learning with vision-language and multimodal models, adapting pretrained representations to downstream tasks while quantifying confidence. Additional branches address active learning strategies that use uncertainty to guide data collection, domain-specific applications ranging from medical imaging to geophysics, and data augmentation techniques that generate synthetic samples to probe model reliability. Representative works such as Generative Model Uncertainty Imaging[1] and Conditional GAN Uncertainty[9] illustrate how generative architectures naturally produce distributional outputs, while Pretrained Models Virtual Metrology[2] and BayesCap[12] demonstrate auxiliary modeling strategies.

A particularly active line of work centers on post-hoc methods that attach lightweight uncertainty estimators to frozen backbones, balancing computational efficiency with calibration quality. Frozen Priors Fluid Forecasts[0] sits within this branch, sharing methodological kinship with Input Output Conditioned Uncertainty[4] and Probabilistic Prototype Calibration[3], which similarly avoid fine-tuning the base model. Compared to Input Output Conditioned Uncertainty[4], which conditions auxiliary networks on both inputs and outputs, Frozen Priors Fluid Forecasts[0] emphasizes operational metrics—quantities directly tied to decision-making in deployment contexts. Meanwhile, Probabilistic Prototype Calibration[3] focuses on prototype-based representations, whereas Frozen Priors Fluid Forecasts[0] targets broader operational forecasting scenarios. These distinctions highlight ongoing trade-offs between generality, computational overhead, and the granularity of uncertainty estimates, with open questions around how to best propagate uncertainty from pretrained features to task-specific metrics without sacrificing the efficiency gains that motivate freezing the generator in the first place.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Principled Input-Output-Conditioned Post-Hoc Uncertainty Estimation for Regression Networks

Authors: Bramlage, Lennart, Curio, CristÃ³bal, Lennart Bramlage, et al. (6 authors total) | **Year/Venue**: 2025 • arXiv.org | **URL**: View paper

Abstract

Uncertainty quantification is critical in safety-sensitive applications but is often omitted from off-the-shelf neural networks due to adverse effects on predictive performance. Retrofitting uncertainty estimates post-hoc typically requires access to model parameters or gradients, limiting feasibility in practice. We propose a theoretically grounded framework for post-hoc uncertainty estimation in regression tasks by fitting an auxiliary model to both original inputs and frozen model outputs. Dr...

Relationship Analysis

Both papers belong to the Post-Hoc Uncertainty Estimation via Auxiliary Models category, fitting secondary models on frozen network outputs to quantify uncertainty without retraining. The original paper focuses on blending empirical data with frozen pretrained generators using a prequential Dirichlet schedule for operational metrics in low-data deployment scenarios, while the candidate paper proposes fitting auxiliary models to both inputs and frozen outputs for regression tasks, emphasizing input-output conditioning and OOD detection. The key difference is that the original paper addresses uncertainty for operational metrics (alert rates, mean scores) under a deployed forecasting rule with martingale posteriors, whereas the candidate paper targets input-dependent predictive uncertainty in regression with auxiliary model fitting on diverse data.

2. BayesCap: Bayesian Identity Cap for Calibrated Uncertainty in Frozen Neural Networks

Authors: Upadhyay, Uddeshya, Uddeshya Upadhyay, Karthik, Shyamgopal, et al. (14 authors total) | **Year/Venue**: 2022 • European Conference on Computer Vision | **URL**: View paper

Abstract

. High-quality calibrated uncertainty estimates are crucial for numerous real-world applications, especially for deep learning-based deployed ML systems. While Bayesian deep learning techniques allow uncertainty estimation, training them with large-scale datasets is an expensive process that does not always yield models competitive with non-Bayesian counterparts. Moreover, many of the high-performing deep learning models that are already trained and deployed are non-Bayesian in nature and do not...

Relationship Analysis

Both papers belong to the Post-Hoc Uncertainty Estimation via Auxiliary Models category, fitting secondary models on frozen network outputs to quantify uncertainty. They overlap in addressing uncertainty quantification for frozen pretrained models without retraining, but differ fundamentally in approach: the original paper uses prequential forecasting with Dirichlet-weighted blending of empirical and generator distributions to produce martingale-based uncertainty intervals for operational metrics, while the candidate paper (BayesCap) trains a Bayesian autoencoder as an identity mapping on frozen model outputs to estimate aleatoric uncertainty via heteroscedastic distributions for image-level predictions.

Contributions Analysis

Overall novelty summary. The paper introduces a forecast-first uncertainty quantification framework for operational metrics when deploying frozen pretrained generators with limited real samples. It resides in the 'Post-Hoc Uncertainty Estimation via Auxiliary Models' leaf, which contains only three papers total. This leaf sits within the broader 'Uncertainty Quantification Frameworks for Frozen Pretrained Models' branch, indicating a relatively sparse research direction focused on attaching auxiliary models to frozen networks. The sibling papers address input-output conditioned uncertainty and probabilistic prototype calibration, suggesting the leaf covers diverse post-hoc strategies but remains underpopulated compared to generative model-based branches.

The taxonomy reveals neighboring leaves in 'Evidential and Meta-Learning Approaches' and 'Pretrained Uncertainty Modules,' both emphasizing meta-learned or evidential reasoning over frozen representations. The paper diverges from these by focusing on operational metrics and martingale posteriors rather than evidential frameworks or transfer learning. Adjacent branches like 'Generative Model-Based Uncertainty Estimation' contain substantially more papers (diffusion, GAN, Bayesian methods), indicating that generative-centric uncertainty is a more crowded area. The paper's emphasis on operational forecasting and time-consistent guarantees distinguishes it from these generative-focused directions, which typically target per-example or reconstruction uncertainty.

Among nine candidates examined across three contributions, none were flagged as clearly refuting the work. The prequential forecasting framework with Dirichlet blending examined one candidate with no refutation. The martingale posterior method examined five candidates, all non-refutable or unclear. The minimax hyperparameter criterion examined three candidates, again with no refutations. This limited search scope—nine papers total—suggests the analysis captures a narrow semantic neighborhood rather than exhaustive prior work. The absence of refutations within this small sample indicates the specific combination of martingale posteriors, Dirichlet blending, and operational metric forecasting may be underexplored, though broader literature beyond these nine candidates remains unexamined.

Given the sparse taxonomy leaf and limited search scope, the work appears to occupy a relatively novel position within post-hoc uncertainty estimation for frozen models. However, the analysis explicitly covers only top-K semantic matches and does not claim exhaustive coverage. The framework's integration of prequential forecasting, martingale posteriors, and operational metrics may represent a distinctive synthesis, but definitive novelty assessment would require examining a larger candidate pool and exploring connections to adjacent fields like online learning or sequential decision-making under uncertainty.

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

Contribution 1: Prequential forecasting framework with Dirichlet blending for frozen generative models

Description: The authors propose a prequential forecasting approach that blends empirical data with a fixed pretrained generator using a Dirichlet-style schedule ($\lambda_i = \alpha/(i+\alpha)$). They prove this is the unique affine combination ensuring time-consistent forecasts, making the sequence of forecasted functionals form a martingale.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Human Activity Recognition with an HMM-Based Generative Model

URL: View paper

Brief Assessment

HMM Generative Model[41] applies Dirichlet distributions to emission probabilities in hidden Markov models for human activity recognition, not to prequential forecasting or blending empirical data with pretrained generators. The technical contexts are entirely different.

Contribution 2: Martingale posterior method for uncertainty quantification without retraining

Description: The authors develop a martingale posterior approach that quantifies uncertainty by simulating future forecasts under the deployed blending rule. This method provides calibrated predictive intervals for operational metrics without requiring model retraining or likelihood evaluation.

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. Test Time Scaling for Neural Processes

URL: View paper

Brief Assessment

Test Time Scaling[40] focuses on refining latent variable samples in neural processes at test time using sequential Monte Carlo samplers, not on martingale posteriors for blended forecasting rules. The candidate addresses uncertainty in meta-learning contexts with variational posteriors, while the original develops martingale posteriors for operational metrics under deployed predictive rules with frozen generators.

2. Moment Martingale Posteriors for Semiparametric Predictive Bayes

URL: View paper

Brief Assessment

Moment Martingale Posteriors[37] focuses on semiparametric predictive Bayes using moment-based martingale posteriors for mixture distributions, not on uncertainty quantification for deployed generative models without retraining. The candidate addresses a different problem domain (semiparametric inference) rather than operational metrics for frozen pretrained generators.

3. Alternative formats

URL: View paper

Brief Assessment

Alternative Formats[39] is a PhD thesis on drill-system parameter estimation and hazardous event detection, with no content related to martingale posteriors, uncertainty quantification, or machine learning deployment methods.

4. Asymptotics for parametric martingale posteriors

URL: View paper

Brief Assessment

Parametric Martingale Posteriors[36] focuses on asymptotic properties and acceleration of martingale posterior sampling through normal approximations and Bernstein-von Mises results. The original paper develops a specific application to deployment scenarios with frozen pretrained generators and dirichlet-weighted blending, which is a distinct methodological contribution not addressed in the candidate.

5. Towards the Uncertainty-aware Geospatial Artificial Intelligence

URL: View paper

Brief Assessment

Uncertainty Aware Geospatial AI[38] appears to focus on geospatial AI applications. The provided candidate context is too limited (only fragments) to assess whether it addresses martingale posteriors, likelihood-free resampling, or blending rules for deployment without retraining.

Contribution 3: Minimax criterion for hyperparameter selection in low-data regime

Description: The authors provide a principled method for selecting the hyperparameter α by formulating a small-sample minimax problem that explicitly trades off sampling variance against model-data mismatch. This yields a closed-form expression $\alpha^* = \sigma^2/\Delta^2$ that is independent of sample size.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Enhanced Balancing of Bias-Variance Tradeoff in Stochastic Estimation: A Minimax Perspective

URL: View paper

Brief Assessment

Bias Variance Tradeoff Minimax[44] focuses on asymptotic minimax risk ratios for stochastic estimators in general settings, not on small-sample hyperparameter selection for blending empirical and pretrained generative models in deployment scenarios.

2. Locally adaptive label smoothing improves predictive churn

URL: View paper

Brief Assessment

Locally Adaptive Label Smoothing[43] addresses minimax hyperparameter selection for label smoothing in classification tasks, not for blending pretrained generators with empirical data in deployment scenarios. The candidate's minimax approach optimizes k-nn smoothing parameters for prediction churn reduction, while the original paper's minimax criterion balances sampling variance against model-data mismatch for frozen generative models in low-data deployment.

3. Lookbehind-SAM: k steps back, 1 step forward

URL: View paper

Brief Assessment

Lookbehind SAM[42] addresses hyperparameter selection in sharpness-aware minimization for neural network optimization, not low-data regime uncertainty quantification with pretrained generative models. The minimax formulations serve entirely different purposes in distinct problem domains.

Appendix: Text Similarity Detection

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

References

- [0] Frozen Priors, Fluid Forecasts: Prequential Uncertainty for Low-Data Deployment with Pretrained Generative Models View paper
- [1] Quantifying generative model uncertainty in posterior sampling methods for computational imaging View paper
- [2] Large Pre-Trained Models and Few-Shot FineTuning for Virtual Metrology: A Framework for Uncertainty-Driven Adaptive Process Control in Semiconductor â¦ View paper
- [3] Probabilistic Prototype Calibration of Vision-Language Models for Generalized Few-shot Semantic Segmentation View paper
- [4] Principled Input-Output-Conditioned Post-Hoc Uncertainty Estimation for Regression Networks View paper

- [5] Distinguishing the Knowable from the Unknowable with Language Models View paper
- [6] AVadCLIP: Audio-Visual Collaboration for Robust Video Anomaly Detection View paper
- [7] Estimating uncertainty in diffusion MRI models using generative deep learning View paper
- [8] Active Learning of non-Semantic Speech Tasks with Pretrained models View paper
- [9] Estimation with uncertainty via conditional generative adversarial networks View paper
- [10] Generative Uncertainty in Diffusion Models View paper
- [11] Generative emulation and uncertainty quantification of geological CO2 storage with conditional diffusion models View paper
- [12] BayesCap: Bayesian Identity Cap for Calibrated Uncertainty in Frozen Neural Networks View paper
- [13] Integrating large pre-trained models into multimodal named entity recognition with evidential fusion View paper
- [14] Uncertainty-Aware Deep Classifiers using Generative Models View paper
- [15] Harmonizing Generalization and Specialization: Uncertainty-Informed Collaborative Learning for Semi-supervised Medical Image Segmentation View paper
- [16] Uncertainty Quantification with Generative Models View paper
- [17] UGrafting: Uncertainty Quantification in Self-Supervised Learning View paper
- [18] Into the Unknown: Towards using Generative Models for Sampling Priors of Environment Uncertainty for Planning in Configuration Spaces View paper
- [19] Novel Uncertainty Quantification Through Perturbation-Assisted Sample Synthesis View paper
- [20] Stochastic latent feature distillation: Enhancing dataset distillation via structured uncertainty modeling View paper
- [21] Application of conditional generative adversarial network in ground motion modelling encompassing epistemic uncertainty View paper
- [22] I Detect What I Don't Know: Incremental Anomaly Learning with Stochastic Weight Averaging-Gaussian for Oracle-Free Medical Imaging View paper
- [23] Scalable crystal structure relaxation using an iteration-free deep generative model with uncertainty quantification View paper
- [24] Evaluating Aleatoric Uncertainty via Conditional Generative Models View paper
- [25] Generative modelling for mass-mapping with fast uncertainty quantification View paper
- [26] Toward More Reliable Artificial Intelligence: Reducing Hallucinations in Vision-Language Models View paper
- [27] Deep Generative Models for Anomaly Detection and Uncertainty Quantification in Biology and Medicine View paper
- [28] Guided Uncertainty Learning Using a Post-Hoc Evidential Meta-Model View paper
- [29] Uncertainty and Generalizability in Foundation Models for Earth Observation View paper
- [30] TRiCo: Triadic Game-Theoretic Co-Training for Robust Semi-Supervised Learning View paper
- [31] Towards Understanding and Quantifying Uncertainty for Text-to-Image Generation View paper
- [32] Pretrained Visual Uncertainties View paper
- [33] CINEMAE: Leveraging Frozen Masked Autoencoders for Cross-Generator AI Image Detection View paper
- [34] Deep Generative Models for Stochastic Seismic Imaging and Uncertainty Quantification View paper
- [35] Deep Generative Models for Competency Awareness and Uncertainty Quantification View paper
- [36] Asymptotics for parametric martingale posteriors View paper
- [37] Moment Martingale Posteriors for Semiparametric Predictive Bayes View paper
- [38] Towards the Uncertainty-aware Geospatial Artificial Intelligence View paper
- [39] Alternative formats View paper
- [40] Test Time Scaling for Neural Processes View paper
- [41] Human Activity Recognition with an HMM-Based Generative Model View paper
- [42] Lookbehind-SAM: k steps back, 1 step forward View paper
- [43] Locally adaptive label smoothing improves predictive churn View paper
- [44] Enhanced Balancing of Bias-Variance Tradeoff in Stochastic Estimation: A Minimax Perspective View paper