



Bayesian Numerical Integration with Neural Networks

Katharina Ott
Cologne
10 April 2024

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN



Wind farm simulation



★ f : energy output of wind farm

★ π : Distribution over wind direction α

Expected energy output:

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Numerical integration: Monte Carlo sampling



π and f are cheap

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Assumption: We can sample from π

$$\Pi[f] \approx \frac{1}{N} \sum_{n=1}^N f(\alpha_i), \quad \alpha_i \sim \pi(\alpha)$$

- * extremely simple
- * extensions: MCMC sampling



Numerical integration: Monte Carlo sampling



π and f are cheap

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Assumption: We can sample from π

$$\Pi[f] \approx \frac{1}{N} \sum_{n=1}^N f(\alpha_i), \quad \alpha_i \sim \pi(\alpha)$$

- * extremely simple
- * extensions: MCMC sampling



Numerical integration: Monte Carlo sampling



π and f are cheap

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Assumption: We can sample from π

$$\Pi[f] \approx \frac{1}{N} \sum_{n=1}^N f(\alpha_n), \quad \alpha_n \sim \pi(\alpha)$$

- * extremely simple
- * extensions: MCMC sampling



Numerical integration: Monte Carlo sampling



π and f are cheap

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Assumption: We can sample from π

$$\Pi[f] \approx \frac{1}{N} \sum_{n=1}^N f(\alpha_n), \quad \alpha_n \sim \pi(\alpha)$$

- * extremely simple
- * extensions: MCMC sampling



Numerical integration: Monte Carlo sampling



π and f are cheap

$$\Pi[f] = \int f(\alpha)\pi(\alpha)d\alpha$$

Assumption: We can sample from π

$$\Pi[f] \approx \frac{1}{N} \sum_{n=1}^N f(\alpha_i), \quad \alpha_i \sim \pi(\alpha)$$

- * extremely simple
- * extensions: MCMC sampling



Numerical integration: Bayesian Quadrature

What if π or f are **expensive**?

- ✦ When is our approximation good enough?
- ✦ Use knowledge about function $f \rightarrow$ Use GP as a prior for f

Numerical integration: Bayesian Quadrature

Quick recap (+ my notation):

Gaussian process $\mathbf{g} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$

Data $\mathcal{D} = \{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$

Compute predictive posterior $\mathbf{g} \mid \mathcal{D} = \mathcal{GP}(\mathbf{m}_p, \mathbf{k}_p)$,

$$\mathbf{m}_p(\mathbf{x}) = \mathbf{m}(\mathbf{x}) + \mathbf{k}_X(\mathbf{x})^T (\mathbf{k}_{XX} + \sigma^2 \mathbf{I})^{-1} (\mathbf{Y} - \mathbf{m}_X)$$

$$\mathbf{k}_p(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_X(\mathbf{x})^T (\mathbf{k}_{XX} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_X(\mathbf{x})$$

Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$

$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$



$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Conditioning the GP on observations $\Lambda | \mathcal{D} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$

$$\mu_{\mathcal{D}} = \int_{\mathcal{X}} m_p(x)\pi(x)dx$$

$$\sigma_{\mathcal{D}} = \int_{\mathcal{X}} \int_{\mathcal{X}} k_p(x, x')\pi(x)\pi(x')dx dx'$$



Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$



$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Conditioning the GP on observations $\Lambda \mid \mathcal{D} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$

$$\mu_{\mathcal{D}} = \int_{\mathcal{X}} m_p(x)\pi(x)dx$$

$$\sigma_{\mathcal{D}} = \int_{\mathcal{X}} \int_{\mathcal{X}} k_p(x, x')\pi(x)\pi(x')dx dx'$$



Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$



$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Conditioning the GP on observations $\Lambda \mid \mathcal{D} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$

$$\mu_{\mathcal{D}} = \int_{\mathcal{X}} m_p(x)\pi(x)dx$$

$$\sigma_{\mathcal{D}} = \int_{\mathcal{X}} \int_{\mathcal{X}} k_p(x, x')\pi(x)\pi(x')dx dx'$$



Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$



$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Conditioning the GP on observations $\Lambda \mid \mathcal{D} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$

$$\mu_{\mathcal{D}} = \int_{\mathcal{X}} m_p(x)\pi(x)dx$$

$$\sigma_{\mathcal{D}} = \int_{\mathcal{X}} \int_{\mathcal{X}} k_p(x, x')\pi(x)\pi(x')dx dx'$$



Numerical integration: Bayesian Quadrature

Prior $g \sim \mathcal{GP}(m, k)$



$$\Lambda = \int_{\mathcal{X}} g(x)\pi(x)dx$$

Note $\Lambda \sim \mathcal{N}(\mu, \sigma)$

Conditioning the GP on observations $\Lambda \mid \mathcal{D} \sim \mathcal{N}(\mu_{\mathcal{D}}, \sigma_{\mathcal{D}})$

$$\mu_{\mathcal{D}} = \int_{\mathcal{X}} m_p(x)\pi(x)dx$$

$$\sigma_{\mathcal{D}} = \int_{\mathcal{X}} \int_{\mathcal{X}} k_p(x, x')\pi(x)\pi(x')dxdx'$$



Numerical integration: ?

- ✦ f and π cheap \rightarrow MC-sampling
- ✦ f or π expensive \rightarrow Bayesian quadrature

Wind farm example: f takes 2 min

- ✦ sampling expensive
- ✦ But: we can obtain many samples \rightarrow Bayesian quadrature expensive

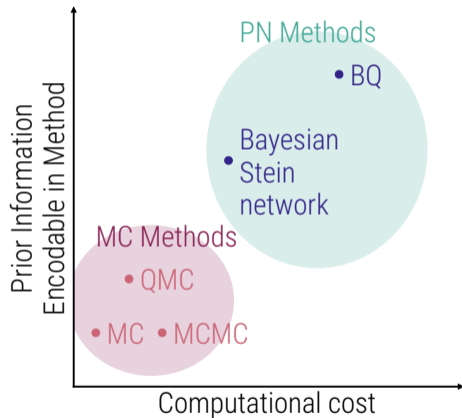
How can we handle large amounts of data + higher dimensions?

Neural networks for numerical integration

$$\Pi[f] = \int_{\mathcal{X}} f(x)\pi(x)dx.$$

Why Bayesian neural networks for integration tasks?

- ✦ Scaling to many data points
- ✦ Scaling to higher dimensions
- ✦ Memory efficient
- ✦ Neural networks interpolate well
- ✦ Uncertainty estimates



Numerical integration with a neural networks

1. Approximate f with a Bayesian neural networks $f_\theta \approx f$
2. Compute the integral (how? \rightarrow later)

$$\Pi[f_\theta] = \int f_\theta(\alpha)\pi(\alpha)d\alpha \approx \Pi[f]$$

3. Compute posterior for $\Pi[f_\theta]$

Numerical integration with a neural networks

1. Approximate f with a **Bayesian neural networks** $f_\theta \approx f$
2. Compute the integral (how? \rightarrow later)

$$\Pi[f_\theta] = \int f_\theta(\alpha)\pi(\alpha)d\alpha \approx \Pi[f]$$

3. Compute posterior for $\Pi[f_\theta]$

Bayesian neural networks

Neural networks

Neural network $f_\theta: \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$

Network parameters: $\theta \in \Theta \subseteq \mathbb{R}^{d_\theta}$

Dataset: $\mathcal{D} = \{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$

Define cost function $\ell: \Theta \rightarrow \mathbb{R}$, e.g.:


★ Regression:

$$\ell(\theta) = \frac{1}{N} \sum_{n=1}^N \|f_\theta(\mathbf{x}_n) - \mathbf{y}_n\|_2^2$$

Goal:


$$\theta^* = \arg \min_{\theta} \ell(\theta)$$

Regression




Overfitting and weight decay

Overfitting



Weight decay



Reduce overfitting \rightarrow regularize loss
e.g., via weight decay $\ell_{\text{decay}}(\theta) = \ell(\theta) + \lambda \|\theta\|_2^2$.

Bayesian neural networks

Prior: No knowledge about network parameters

→ Gaussian prior $p(\theta) = \mathcal{N}(\mathbf{0}, \sigma_0^2 \mathbb{I})$

Likelihood: $\log p(\mathcal{D} | \theta) \propto \ell(\theta)$


Posterior:

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta)p(\theta)}{\int p(\mathcal{D} | \theta)p(\theta)d\theta}$$

Use MCMC to compute posterior and predictive posterior


- ✦ computationally expensive
- ✦ cheap alternative → Laplace approximation

BNN



Laplace approximation

Prior $p(\theta)$




Mackay, A practical Bayesian framework for backpropagation networks,
Neural Computation (1992)


$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$

Laplace approximation

Prior $p(\theta)$



Likelihood $p(y | \theta, x)$




$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$


Mackay, A practical Bayesian framework for backpropagation networks,
Neural Computation (1992)

Laplace approximation


Prior $p(\theta)$



Likelihood $p(y | \theta, x)$



Posterior $p(\theta | x, y)$




$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$


Mackay, A practical Bayesian framework for backpropagation networks,
Neural Computation (1992)

Laplace approximation


Prior $p(\theta)$



Likelihood $p(y | \theta, x)$



Posterior $p(\theta | x, y)$




$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$

Mackay, A practical Bayesian framework for backpropagation networks,
Neural Computation (1992)

Laplace approximation

Prior $p(\theta)$




θ

$\log p(\theta)$

+

$$\sim \lambda \|\theta\|_2^2$$

Likelihood $p(y | \theta, x)$





x

$\log p(y | \theta, x)$

$$\sim \frac{1}{N} \sum_{i=1}^N (f_{\theta}(x_i) - y_i)^2$$

Posterior $p(\theta | x, y)$



θ


$= \log p(\theta | x, y) + \log Z$

$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$

Mackay, A practical Bayesian framework for backpropagation networks, *Neural Computation* (1992)


Laplace approximation

Prior $p(\theta)$





θ

Likelihood $p(y | \theta, x)$



x

Posterior $p(\theta | x, y)$



θ

$\log p(\theta)$

+

$\log p(y | \theta, x)$

= $\log p(\theta | x, y) + \log Z$

$$\sim \lambda \|\theta\|_2^2$$

$$\sim \frac{1}{N} \sum_{i=1}^N (f_{\theta}(x_i) - y_i)^2$$

Regularized loss

$$p(\theta | x, y) = \frac{p(y | \theta, x)p(\theta)}{\int p(y | \theta, x)p(\theta)d\theta}$$

Mackay, A practical Bayesian framework for backpropagation networks, *Neural Computation* (1992)

Linearization and GGN


Taylor series approximation requires:

$$H = \nabla_{\theta}^2 \log p(\theta | x, y)$$

- ✦ Not necessarily positive definite
- ✦ Computationally expensive

Use additional approximation of Hessian:

- ✦ Generalized Gauss Newton (GGN) approximation



Schraudolph, Fast Curvature Matrix-Vector Products for Second-Order Gradient Descent, *Neural Computation* (2002)

Ritter et al., A scalable Laplace approximation for neural networks *ICLR* (2018)

Daxberger et al., Laplace Redux - Effortless Bayesian Deep Learning, *NeurIPS* (2021)

Integrating neural networks

Integrating neural networks

1. Approximate f with a Bayesian neural networks $f_\theta \approx f$
2. **Compute the integral**

$$\Pi[f_\theta] = \int f_\theta(\alpha)\pi(\alpha)d\alpha \approx \Pi[f]$$

3. Compute posterior for $\Pi[f_\theta]$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

where $u : \mathbb{R}^n \rightarrow \mathbb{R}^n$

and $\nabla_x u(x) = \sum_i \partial_{x_i} u(x)_i$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

where $u : \mathbb{R}^n \rightarrow \mathbb{R}^n$

and $\nabla_x u(x) = \sum_i \partial_{x_i} u(x)_i$

$$\Pi [\mathcal{S}[u]] = \int_{\mathbb{R}^n} [(\nabla_x \log \pi(x)) \cdot u(x) \pi(x) + \nabla_x u(x) \cdot \pi(x)] dx$$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

where $u : \mathbb{R}^n \rightarrow \mathbb{R}^n$

and $\nabla_x u(x) = \sum_i \partial_{x_i} u(x)_i$

$$\begin{aligned} \Pi [\mathcal{S}[u]] &= \int_{\mathbb{R}^n} [(\nabla_x \log \pi(x)) \cdot u(x)\pi(x) + \nabla_x u(x) \cdot \pi(x)] dx \\ &= \int_{\mathbb{R}^n} \left[\frac{\nabla_x \pi(x)}{\pi(x)} \cdot u(x)\pi(x) + \nabla_x u(x) \cdot \pi(x) \right] dx \end{aligned}$$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

where $u : \mathbb{R}^n \rightarrow \mathbb{R}^n$

and $\nabla_x u(x) = \sum_i \partial_{x_i} u(x)_i$

$$\begin{aligned} \Pi [\mathcal{S}[u]] &= \int_{\mathbb{R}^n} [(\nabla_x \log \pi(x)) \cdot u(x)\pi(x) + \nabla_x u(x) \cdot \pi(x)] dx \\ &= \int_{\mathbb{R}^n} \left[\frac{\nabla_x \pi(x)}{\pi(x)} \cdot u(x)\pi(x) + \nabla_x u(x) \cdot \pi(x) \right] dx \\ &= \underbrace{[\pi(x)u(x)]_{\mathbb{R}^n}}_{=0} + \int_{\mathbb{R}^n} \underbrace{[-\pi(x) \cdot \nabla_x u(x) + \nabla_x u(x) \cdot \pi(x)]}_{=0} dx \\ &= 0. \end{aligned}$$

Stein operator see: Anastasiou et al., *Statistical Science* (2023)

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

$$\Pi [\mathcal{S}[u]] = 0$$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

$$\Pi [\mathcal{S}[u]] = 0$$

Neural network:

$$f_\theta(x) := \mathcal{S}[u_{\theta_u}](x) + \theta_0$$

Stein network

Consider:

$$\mathcal{S}[u](x) := (\nabla_x \log \pi(x))^\top u(x) + \nabla_x u(x)$$

$$\Pi[\mathcal{S}[u]] = 0$$

Neural network:

$$f_\theta(x) := \mathcal{S}[u_{\theta_u}](x) + \theta_0$$


Then

$$\Pi[f_\theta] = \Pi[\mathcal{S}[u_{\theta_u}]] + \Pi[\theta_0] = \theta_0$$

If $f_\theta \approx f$

$$\Pi[f] \approx \Pi[f_\theta] = \theta_0$$

Neural network regression for f_θ
 $\{x_i, f(x_i), \nabla \log \pi(x_i)\}$



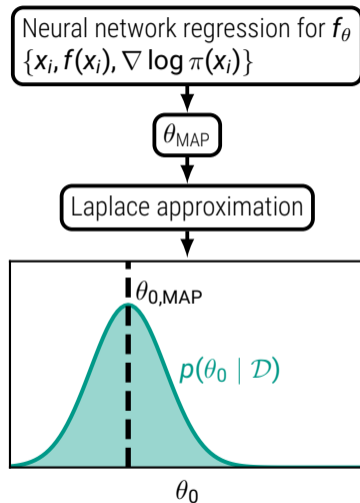
Laplace Approximation for Stein network

Via Laplace approximation:

$$p(\theta | \mathcal{D}) \approx \mathcal{N}(\theta | \theta_{\text{MAP}}, \Sigma_{\text{GGN}})$$

Then:

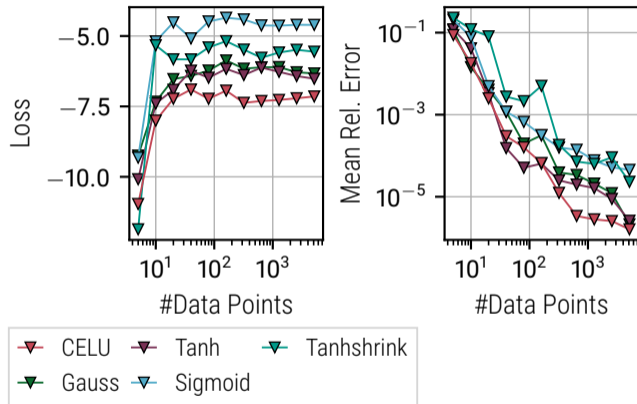
$$p(\theta_0 | \mathcal{D}) \approx \mathcal{N}(\theta_0 | \theta_{0,\text{MAP}}, (\Sigma_{\text{GGN}})_{0,0})$$



Architectural choices

Stein network requires special choices:


- ✦ **Choice of optimizer:** L-BFGS
- ✦ **Choice of activation function:**
Needs to be differentiable
 - ✦ ReLU does not work
 - ✦ CELU works best



Does it work?

Continuous Genz dataset




- ✦ Standard integration benchmark dataset.
- ✦ Integration against a standard normal distribution $\pi(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{0}, \mathbf{1})$.
- ✦ Dimension: $d = 20$.



Does it work?

Continuous Genz dataset



- ★ Standard integration benchmark dataset.
- ★ Integration against a standard normal distribution $\pi(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{0}, \mathbf{1})$.
- ★ Dimension: $d = 20$.



Does it work?

Continuous Genz dataset

- ★ Standard integration benchmark dataset.
- ★ Integration against a standard normal distribution $\pi(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \mathbf{0}, \mathbf{1})$.
- ★ Dimension: $d = 20$.



Calibration:

$$\gamma = \mathbf{e}_{\text{abs}} / \sigma_{\theta_0}$$

\mathbf{e}_{abs} : integration error

σ_{θ_0} : standard deviation from Laplace/BQ

Parameter inference for ODEs

Task

- ✦ Data + ODE model for data

$$\frac{dz}{dt} = f(t, z, x)$$

- ✦ Prior for parameters $p(x)$
- ✦ Goal: compute posterior mean

$$\Pi[x] = \int x \pi(x) dx$$

here $\pi(x) = p(x | \mathcal{D})$.

Issues

- ✦ Sampling expensive (expensive $\nabla \log \pi$).
- ✦ π not available in normalized form \rightarrow BQ impractical.

Parameter inference for ODEs

Task

- ★ Data + ODE model for data

$$\frac{dz}{dt} = f(t, z, x)$$


- ★ Prior for parameters $p(x)$
- ★ Goal: compute posterior mean

$$\Pi[x] = \int x \pi(x) dx$$

here $\pi(x) = p(x | \mathcal{D})$.

Issues

- ★ Sampling expensive (expensive $\nabla \log \pi$).
- ★ π not available in normalized form \rightarrow BQ impractical.



Parameter inference for ODEs

Task

- ★ Data + ODE model for data

$$\frac{dz}{dt} = f(t, z, x)$$

- ★ Prior for parameters $p(x)$
- ★ Goal: compute posterior mean

$$\Pi[x] = \int x \pi(x) dx$$

here $\pi(x) = p(x | \mathcal{D})$.

Issues

- ★ Sampling expensive (expensive $\nabla \log \pi$).
- ★ π not available in normalized form \rightarrow BQ impractical.

Observation

$\nabla_x \log \pi(x)$ can take very large values

Solution

Diffusion Stein operator

$$\mathcal{S}_m[u](x) := (m(x)^\top \nabla_x \log \pi(x))^\top u(x) + \nabla_x \cdot (m(x)u(x)),$$

Choice: Constant $m(x) = C^{-1}$ with

$$C = \max_{i=1, \dots, n} \nabla_x \log \pi(x_i)$$

Parameter inference for ODEs

Task

- ★ Data + ODE model for data

$$\frac{dz}{dt} = f(t, z, x)$$

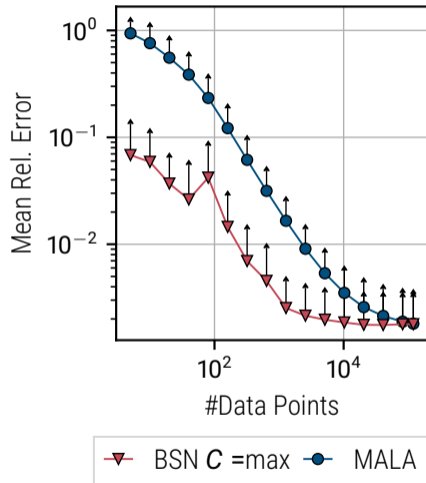
- ★ Prior for parameters $p(x)$
- ★ Goal: compute posterior mean

$$\Pi[x] = \int x \pi(x) dx$$

here $\pi(x) = p(x | \mathcal{D})$.

Issues

- ★ Sampling expensive (expensive $\nabla \log \pi$).
- ★ π not available in normalized form \rightarrow BQ impractical.



Wind farm simulation

- ✦ Expected energy production of a wind farm
- ✦ Distributions over wind direction, turbulence intensity, ...
- ✦ Sampling slow (~ 2 min)

Wind farm simulation

- ✦ Expected energy production of a wind farm
- ✦ Distributions over wind direction, turbulence intensity, ...
- ✦ Sampling slow (~ 2 min)

Bounded domains

For the wind direction $\mathcal{X} = [a, b]$

But we need

$$\Pi [\mathcal{S}[u]] = 0$$

Choose

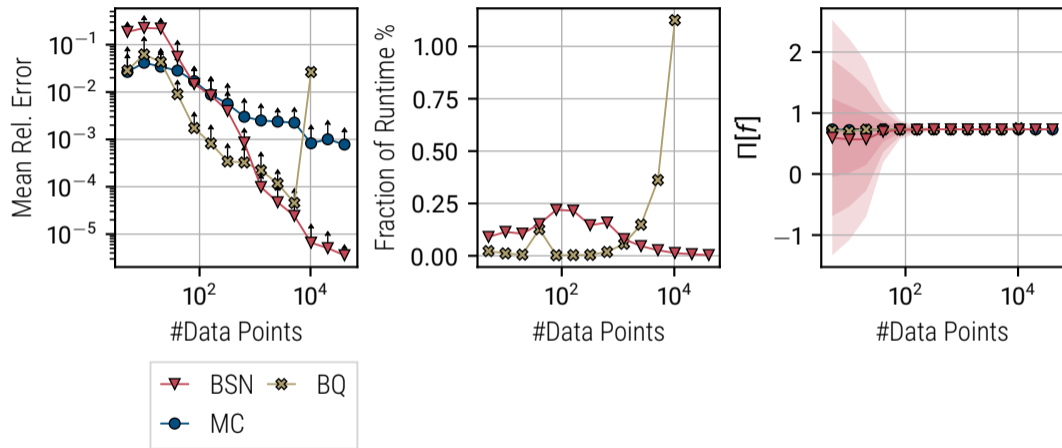
$$\tilde{u}(x) = u(x)\delta(x)$$

$$\delta(a) = 0, \quad \delta(b) = 0$$

For example

$$\delta(x) = (x - a)(x - b)$$

Wind farm simulation



Ott et al., Bayesian numerical integration with neural networks *UAI* (2023)

Limitations and discussion

Choice of prior

- ✦ Encoding more specific information hard
- ✦ Choice of Gaussian prior might be limiting

Scaling the architecture

- ✦ L-BFGS only feasible for small architectures
- ✦ To scale to high dimensions/ more datapoints
 - ✦ Batching
 - ✦ First order optimization (e.g. Adam, SGD)

Summary

Approach

- ✦ Neural networks + Stein operator + Laplace approximation

Advantages

- ✦ Provides uncertainty estimates of good quality
- ✦ Scales well to large number of data points and higher dimensions
- ✦ Handles wide range of integration densities
 - ✦ e.g. if density only available in unnormalized form

Thank You!



Francois-Xavier Briol
University College London



Philipp Hennig
University of Tübingen



Michael Tiemann
Bosch Center for Artificial
Intelligence