Relative Goodness-of-Fit Tests for Models with Latent Variables

Arthur Gretton


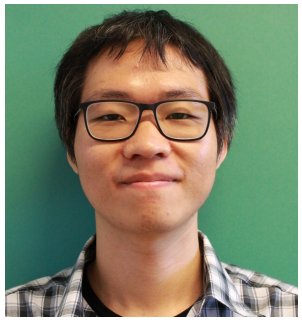

Gatsby Computational Neuroscience Unit,
University College London

Department of Statistics, Harvard, 2023

Model Criticism

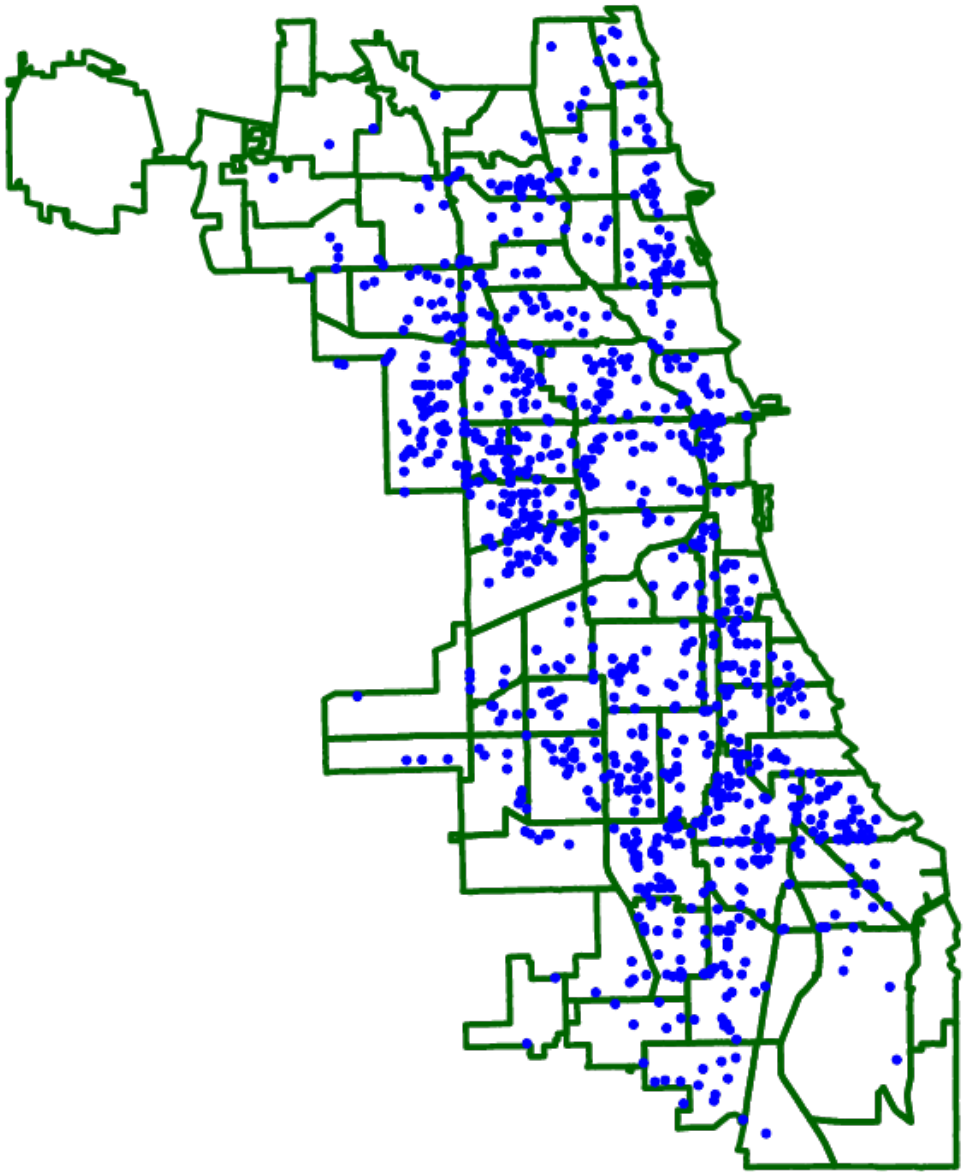

Data = robbery events in Chicago in 2016.

Model Criticism

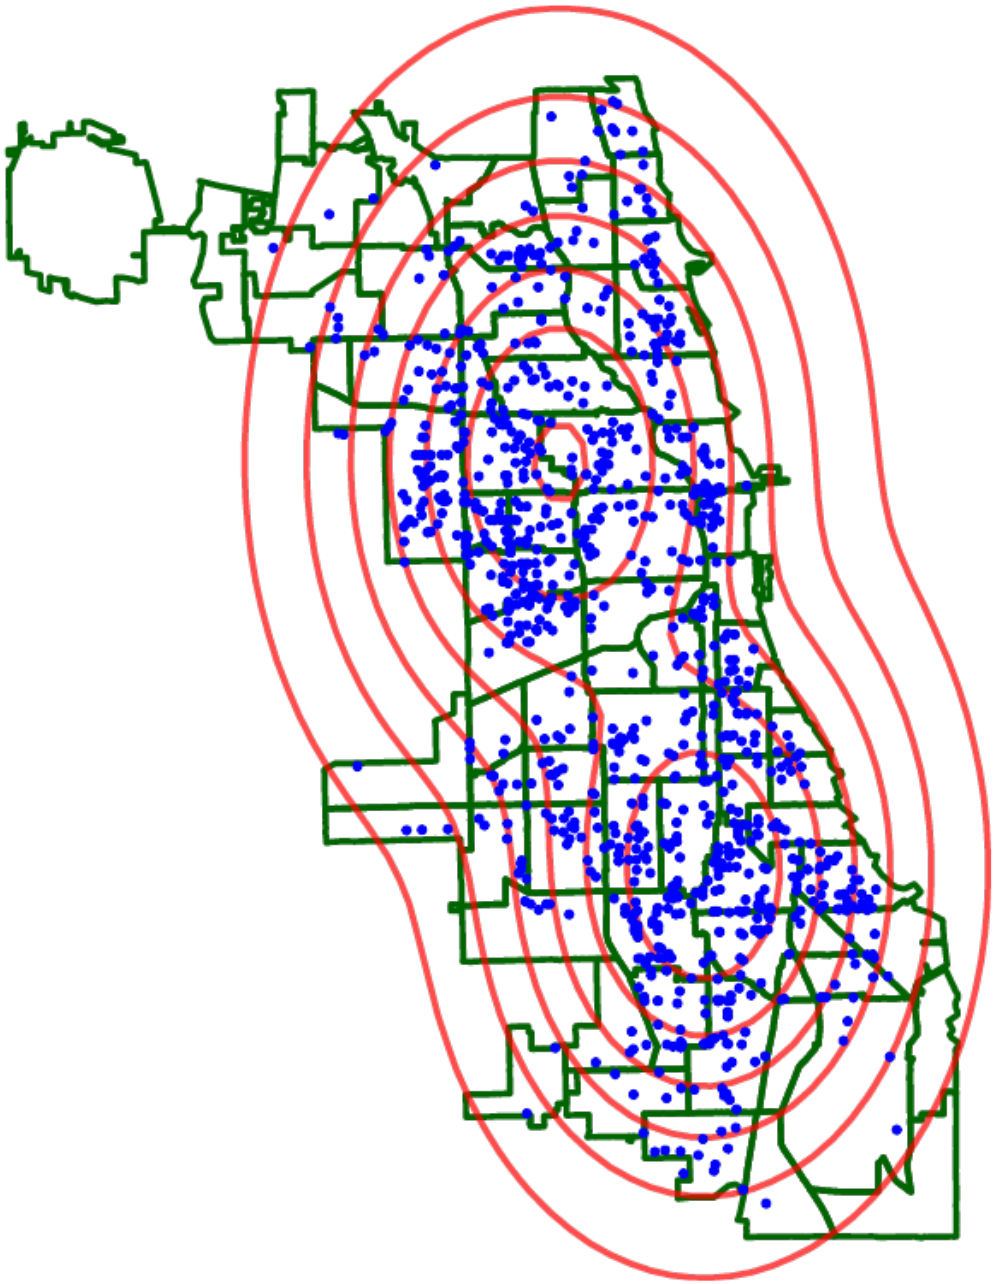

Is this a good model?

Model Criticism

"All models are wrong."

G. Box (1976)

Model comparison

- Have: two candidate models P and Q, and samples $\{x_i\}_{i=1}^n$ from reference distribution R
- Goal: which of P and Q is better?

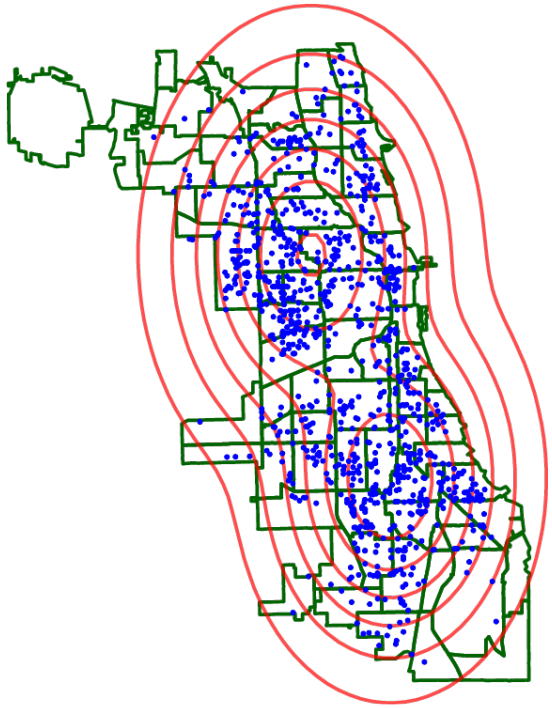

P : two components

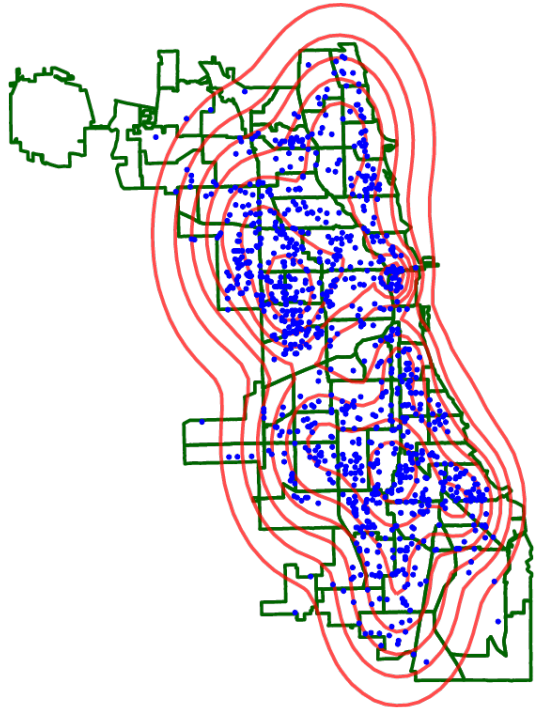

Q : ten components

Most interesting models have latent structure

Graphical model representation of hierarchical LDA with a nested CRP prior, Blei et al. (2003)

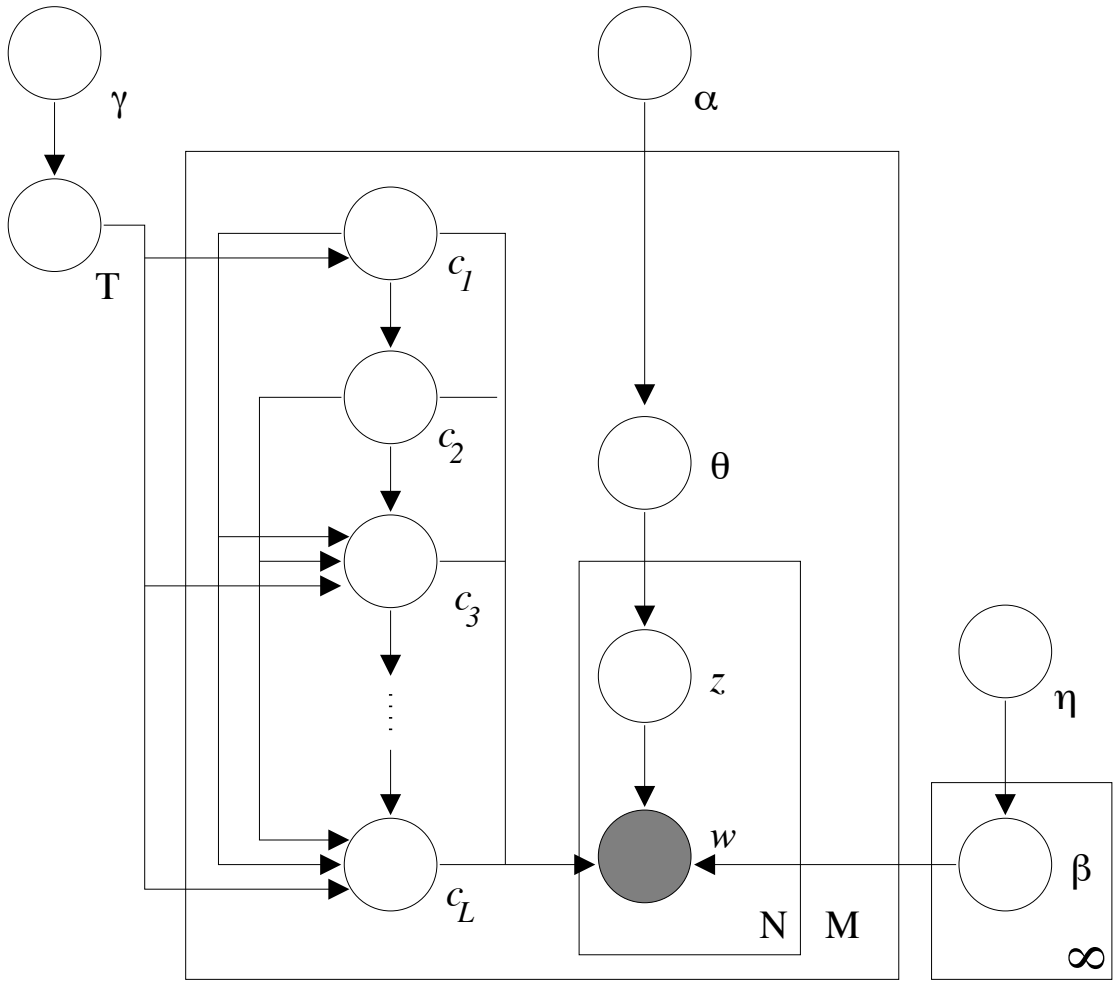

Outline

Relative goodness-of-fit tests for Models with Latent Variables

- The Maximum Mean Discrepancy: an integral probability metric
  - maximize difference in expectations using an RKHS witness class
- The kernel Stein discrepancy
  - Comparing a sample and a model: Stein modification of the witness class
- Constructing a relative hypothesis test using the KSD
- Relative hypothesis tests with latent variables

Kernel Stein Discrepancy

- Model P, data $\{x_i\}_{i=1}^n \sim Q$.
- "All models are wrong" ($P \neq Q$).

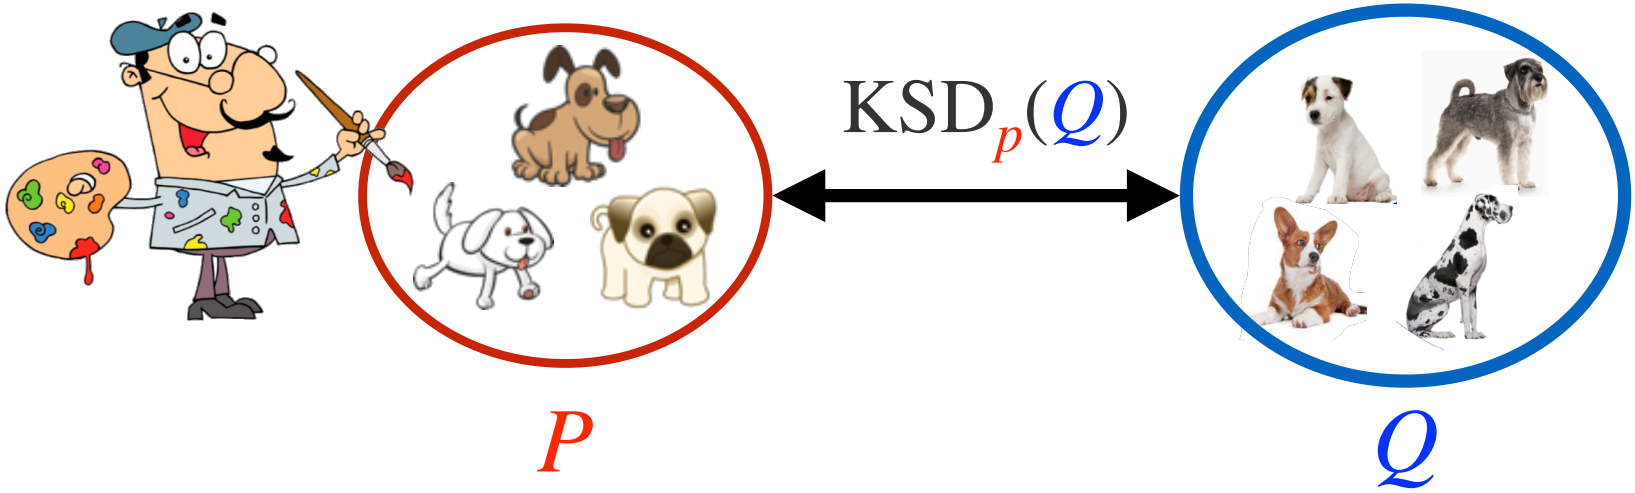

MMD: an integral probability metric

Maximum mean discrepancy: for P vs Q and "smooth" (Reproducing Kernel Hilbert Space) witness class $\mathcal{F}$

$$\mathrm{MMD}(P, Q; \mathcal{F}) := \sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y)]$$

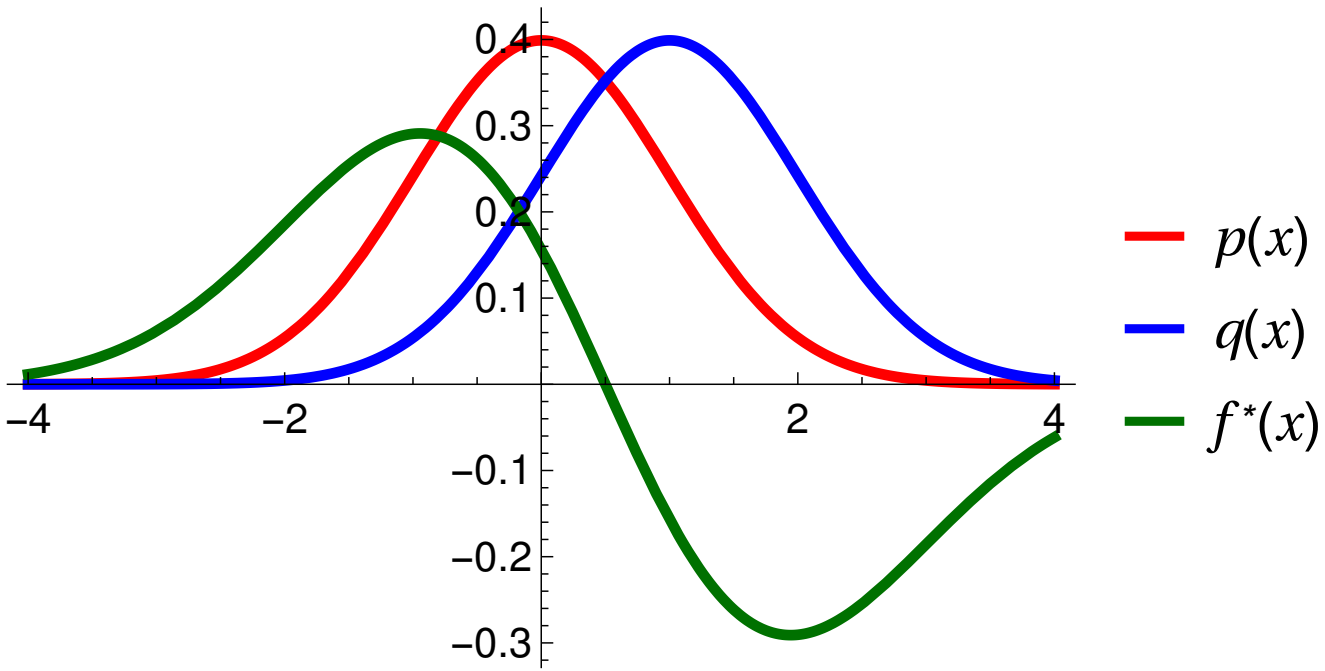

All of kernel methods

Kernels: dot products of features

Feature map $\varphi(x) \in \mathcal{F}$,

$$\varphi(x) = [\dots \varphi_\ell(x) \dots] \in \ell_2$$

For positive definite k,

$$k(x, x') = \langle \varphi(x), \varphi(x') \rangle_{\mathcal{F}}$$

Infinitely many features $\varphi(x)$, dot product in closed form!

Features are solutions to kernel eigenvalue equation

$$\lambda_\ell e_\ell(x) = \int k(x, x') e_\ell(x') dp(x') dx'$$

$$\varphi_\ell(x) = \sqrt{\lambda_\ell} e_\ell(x)$$

where $p(x)$ finite Borel measure satisfying Mercer (e.g. supported on $\mathcal{X}$ where $\mathcal{X}$ compact).

All of kernel methods

Kernels: dot products of features

Feature map $\varphi(x) \in \mathcal{F}$,

$$\varphi(x) = [\dots \varphi_\ell(x) \dots] \in \ell_2$$

For positive definite k,

$$k(x, x') = \langle \varphi(x), \varphi(x') \rangle_{\mathcal{F}}$$

Infinitely many features $\varphi(x)$, dot product in closed form!

Exponentiated quadratic kernel

$$k(x, x') = \exp\left(-\gamma \left\| x - x' \right\|^2\right)$$

$$p(x) = \mathcal{N}(0, 1)$$

$$\varphi(x) = \left[\begin{array}{c} \varphi_1(x) \\ \hline \varphi_2(x) \\ \hline \varphi_3(x) \\ \hline \vdots \end{array}\right]$$

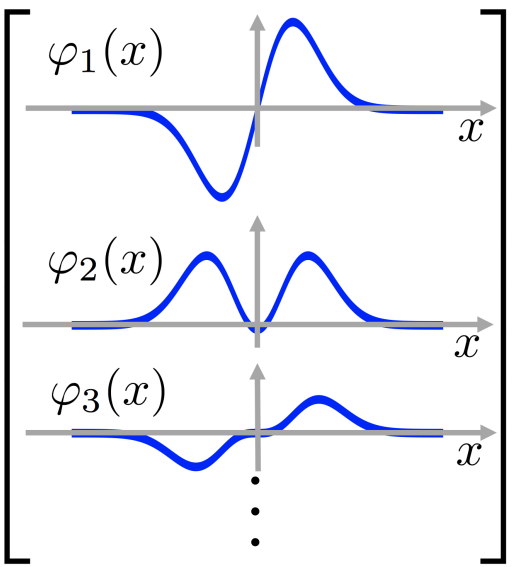

Features: Gaussian Processes for Machine learning, Rasmussen and Williams, Ch. 4.

All of kernel methods

Functions are linear combinations of features:

$$f(x) = \langle f, \varphi(x) \rangle_{\mathcal{F}} = \sum_{\ell=1}^{\infty} f_\ell \varphi_\ell(x) = \begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ \vdots \end{bmatrix}^\top$$

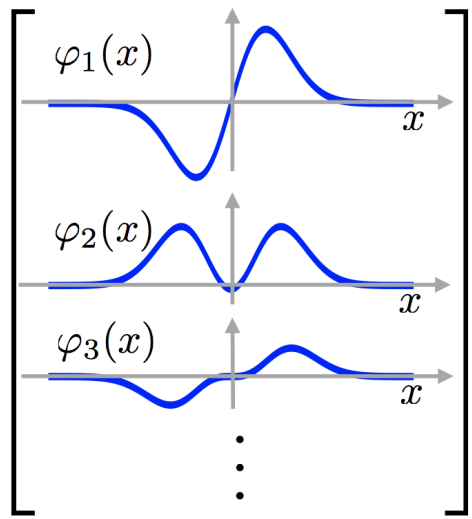

$$\|f\|_{\mathcal{F}}^2 := \sum_{i=1}^{\infty} f_i^2$$

All of kernel methods

"The kernel trick"

$$f(x) = \sum_{\ell=1}^{\infty} f_\ell \varphi_\ell(x)$$

$$= \sum_{i=1}^{m} \alpha_i \underbrace{k(x_i, x)}_{\langle \varphi(x_i), \varphi(x) \rangle_{\mathcal{F}}}$$

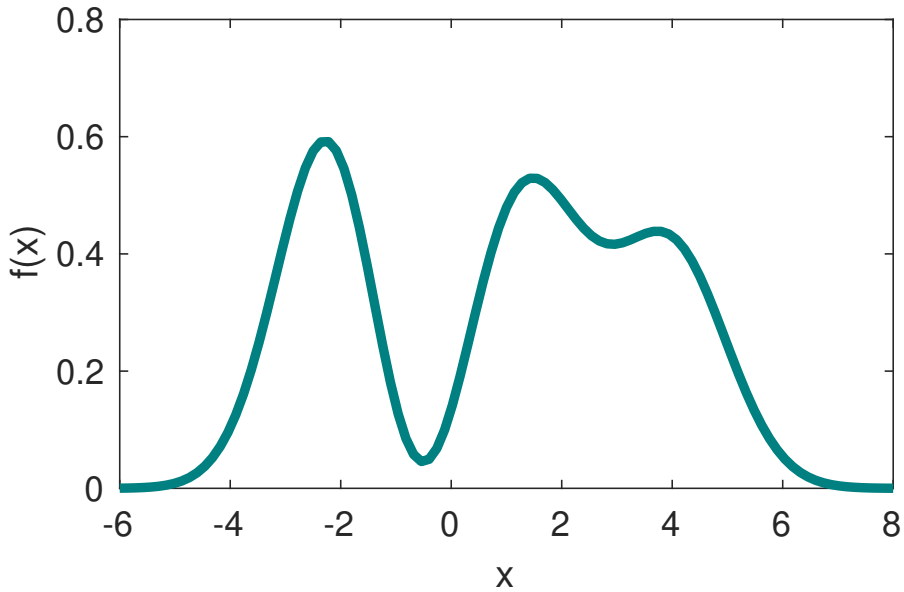

All of kernel methods

"The kernel trick"

$$f(x) = \sum_{\ell=1}^{\infty} f_\ell \varphi_\ell(x)$$
$$= \sum_{i=1}^{m} \alpha_i \underbrace{k(x_i, x)}_{\langle \varphi(x_i), \varphi(x) \rangle_{\mathcal{F}}}$$

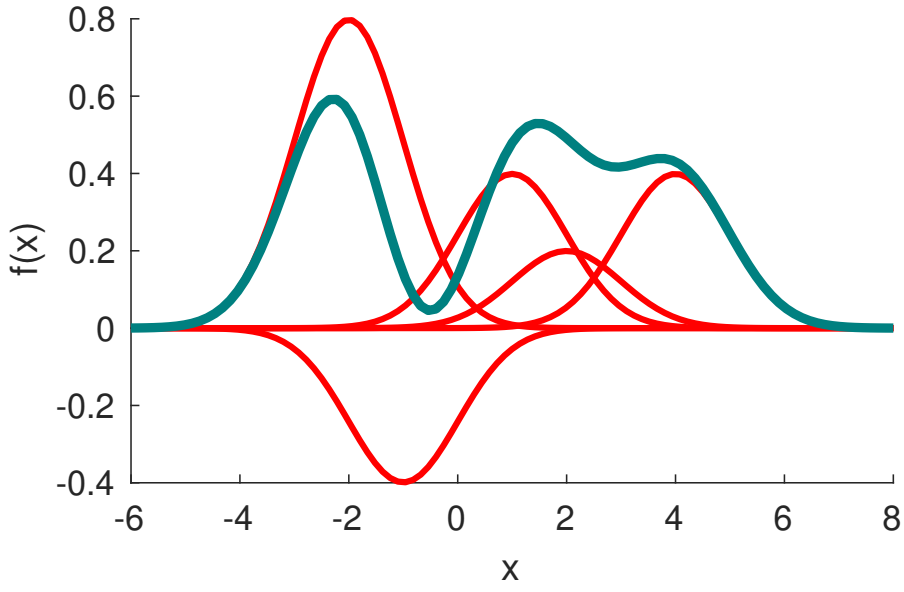

$f_\ell := \sum_{i=1}^{m} \alpha_i \varphi_\ell(x_i)$

Function of infinitely many features expressed using m coefficients.

MMD: an integral probability metric

Maximum mean discrepancy: smooth function for P vs Q

$$\mathrm{MMD}(P, Q; \mathcal{F}) := \sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y)]$$

For characteristic RKHS $\mathcal{F}$, $\mathrm{MMD}(P, Q; \mathcal{F}) = 0$ iff $P = Q$

Other choices for witness function class:

- Bounded continuous [Dudley, 2002]
- Bounded variation 1 (Kolmogorov metric) [Müller, 1997]
- 1-Lipschitz (Wasserstein distances) [Dudley, 2002]

The MMD in closed form

The MMD:

$$MMD(P, Q; F) = \sup_{||f||_{\mathcal{F}} \leq 1} [\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y)]$$

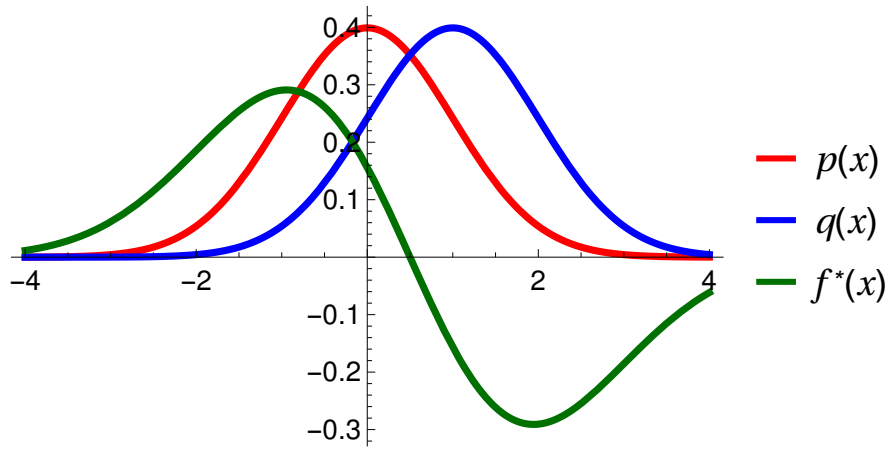

The MMD in closed form

The MMD:

$$
\begin{aligned}
&MMD(P, Q; F) \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} \left[\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y) \right] \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} \left\langle f, \mu_P - \mu_Q \right\rangle_{\mathcal{F}}
\end{aligned}
$$

use

$$
\begin{aligned}
\mathrm{E}_P f(X) &= \mathrm{E}_P \left\langle \varphi(X), f \right\rangle_{\mathcal{F}} \\
&= \left\langle \mathrm{E}_P \left[\varphi(X) \right], f \right\rangle_{\mathcal{F}} \\
&= \left\langle \mu_P, f \right\rangle_{\mathcal{F}}
\end{aligned}
$$

The MMD in closed form

The MMD:

$$
\begin{aligned}
&MMD(P, Q; F) \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y)] \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} \langle f, \mu_P - \mu_Q \rangle_{\mathcal{F}} \\
&= \|\mu_P - \mu_Q\|
\end{aligned}
$$

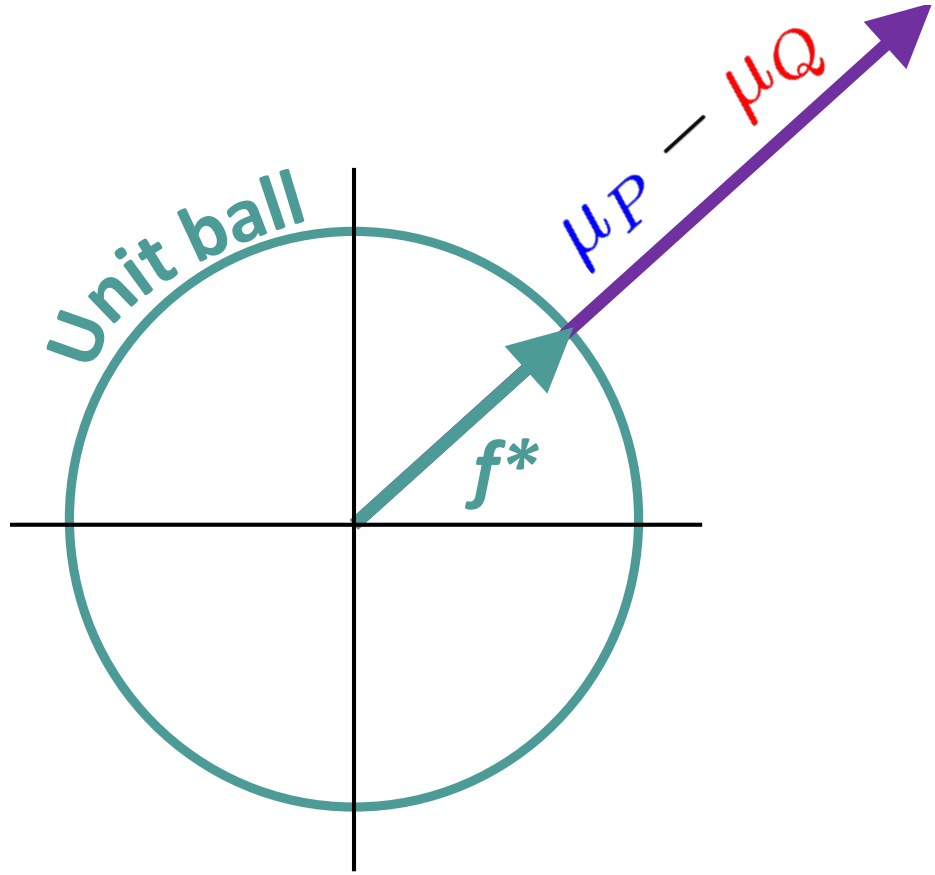

$$f^* = \frac{\mu_P - \mu_Q}{\|\mu_P - \mu_Q\|}$$

The MMD in closed form

The MMD:

$$
\begin{aligned}
& MMD(P, Q; F) \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_P f(X) - \mathrm{E}_Q f(Y)] \\
&= \sup_{\|f\|_{\mathcal{F}} \leq 1} \langle f, \mu_P - \mu_Q \rangle_{\mathcal{F}} \\
&= \|\mu_P - \mu_Q\|
\end{aligned}
$$

In terms of kernels:

$$
\begin{aligned}
MMD^2(P, Q) &= \|\mu_P - \mu_Q\|_{\mathcal{F}}^2 \\
&= \underbrace{\mathrm{E}_P k(x, x')}_{\text{(a)}} + \underbrace{\mathrm{E}_Q k(y, y')}_{\text{(a)}} - 2\underbrace{\mathrm{E}_{P,Q} k(x, y)}_{\text{(b)}}
\end{aligned}
$$

Comparing a sample and model

Can we compute MMD with samples from Q and a model P?

Problem: usualy can't compute $\mathrm{E}_p f$ in closed form.

$$\mathrm{MMD}(P, Q) = \sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_q f - \mathrm{E}_p f]$$

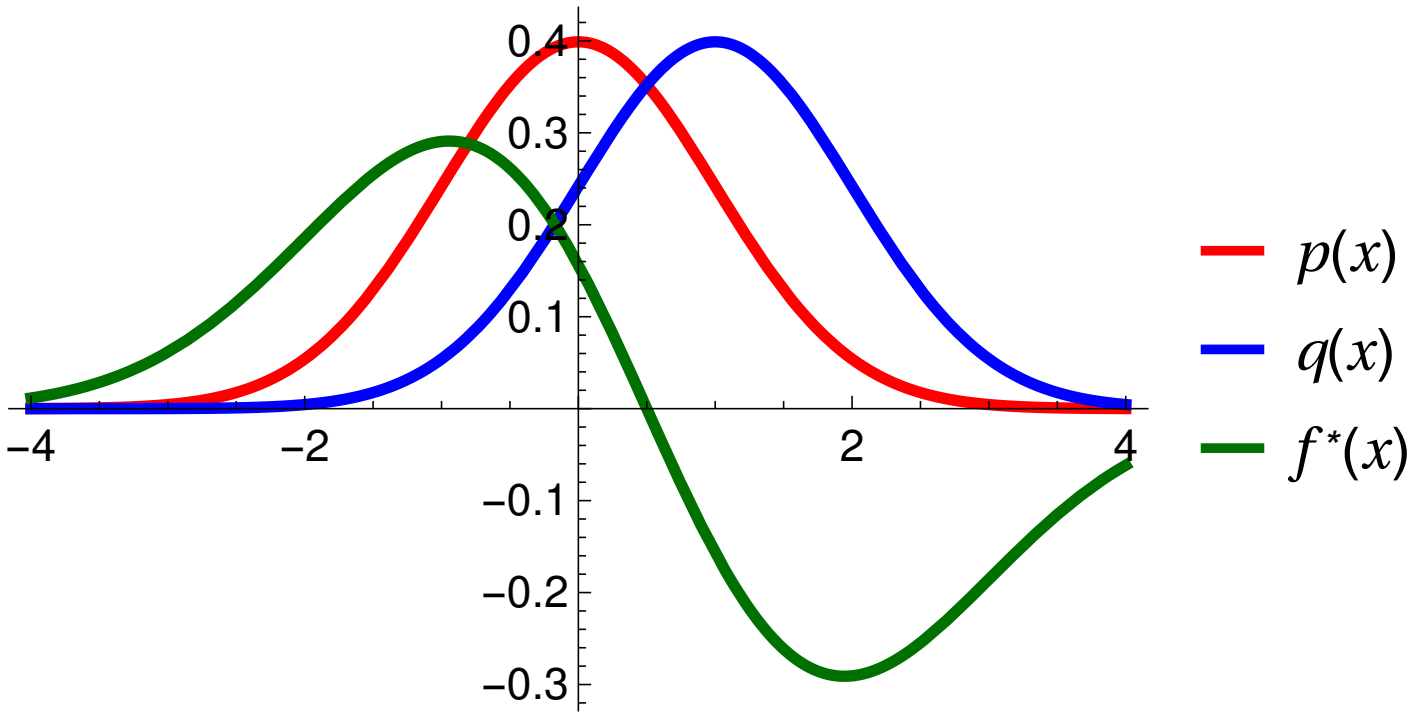

Stein idea

To get rid of $\mathrm{E}_p f$ in

$$\sup_{||f||_{\mathcal{F}} \leq 1} \left[\mathrm{E}_q f - \mathrm{E}_p f\right]$$

we use the (1-D) Langevin Stein operator

$$\begin{aligned}\left[\mathcal{A}_p f\right](x) &= \frac{1}{p(x)} \frac{d}{dx} \left(f(x) p(x)\right) \\ &= f(x) \frac{d}{dx} \log p(x) + \frac{d}{dx} f(x)\end{aligned}$$

Then

$$\mathrm{E}_p \mathcal{A}_p f = 0$$

subject to appropriate boundary conditions.

Gorham and Mackey (NeurIPS 15), Oates, Girolami, Chopin (JRSS B 2016)

Stein idea

To get rid of $\mathrm{E}_p f$ in

$$\sup_{\|f\|_{\mathcal{F}} \leq 1} \left[\mathrm{E}_q f - \mathrm{E}_p f\right]$$

we use the (1-D) Langevin Stein operator

$$\begin{aligned}\left[\mathcal{A}_p f\right](x) &= \frac{1}{p(x)} \frac{d}{dx} \left(f(x) p(x)\right) \\ &= f(x) \frac{d}{dx} \log p(x) + \frac{d}{dx} f(x)\end{aligned}$$

Then

$$\mathrm{E}_p \mathcal{A}_p f = 0$$

subject to appropriate boundary conditions.

$$\mathrm{E}_p \left[\mathcal{A}_p f\right] = \int \left[\frac{1}{\cancel{p(x)}} \frac{d}{dx} \left(f(x) p(x)\right)\right] \cancel{p(x)} dx = \left[f(x) p(x)\right]_{-\infty}^{\infty}$$

Gorham and Mackey (NeurIPS 15), Oates, Girolami, Chopin (JRSS B 2016)

Stein idea

To get rid of $\mathrm{E}_p f$ in

$$\sup_{\|f\|_{\mathcal{F}} \leq 1} [\mathrm{E}_q f - \mathrm{E}_p f]$$

we use the (1-D) Langevin Stein operator

$$[\mathcal{A}_p f](x) = \frac{1}{p(x)} \frac{d}{dx} (f(x)p(x))$$
$$= f(x) \frac{d}{dx} \log p(x) + \frac{d}{dx} f(x)$$

Then

$$\mathrm{E}_p \mathcal{A}_p f = 0$$

subject to appropriate boundary conditions.

Do not need to normalize p, or sample from it.

Gorham and Mackey (NeurIPS 15), Oates, Girolami, Chopin (JRSS B 2016)

Kernel Stein Discrepancy

Stein operator

$$\mathcal{A}_p f = f(x)\frac{d}{dx}\log p(x) + \frac{d}{dx}f(x)$$

Kernel Stein Discrepancy (KSD)

$$\mathrm{KSD}_p(Q) = \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_q \mathcal{A}_p g - \mathrm{E}_p \mathcal{A}_p g$$

Kernel Stein Discrepancy

Stein operator

$$\mathcal{A}_p f = f(x)\frac{d}{dx}\log p(x) + \frac{d}{dx}f(x)$$

Kernel Stein Discrepancy (KSD)

$$\mathrm{KSD}_p(Q) = \sup_{\|g\|_{\mathcal{F}} \le 1} \mathrm{E}_q \mathcal{A}_p g - \cancel{\mathrm{E}_p \mathcal{A}_p g} = \sup_{\|g\|_{\mathcal{F}} \le 1} \mathrm{E}_q \mathcal{A}_p g$$

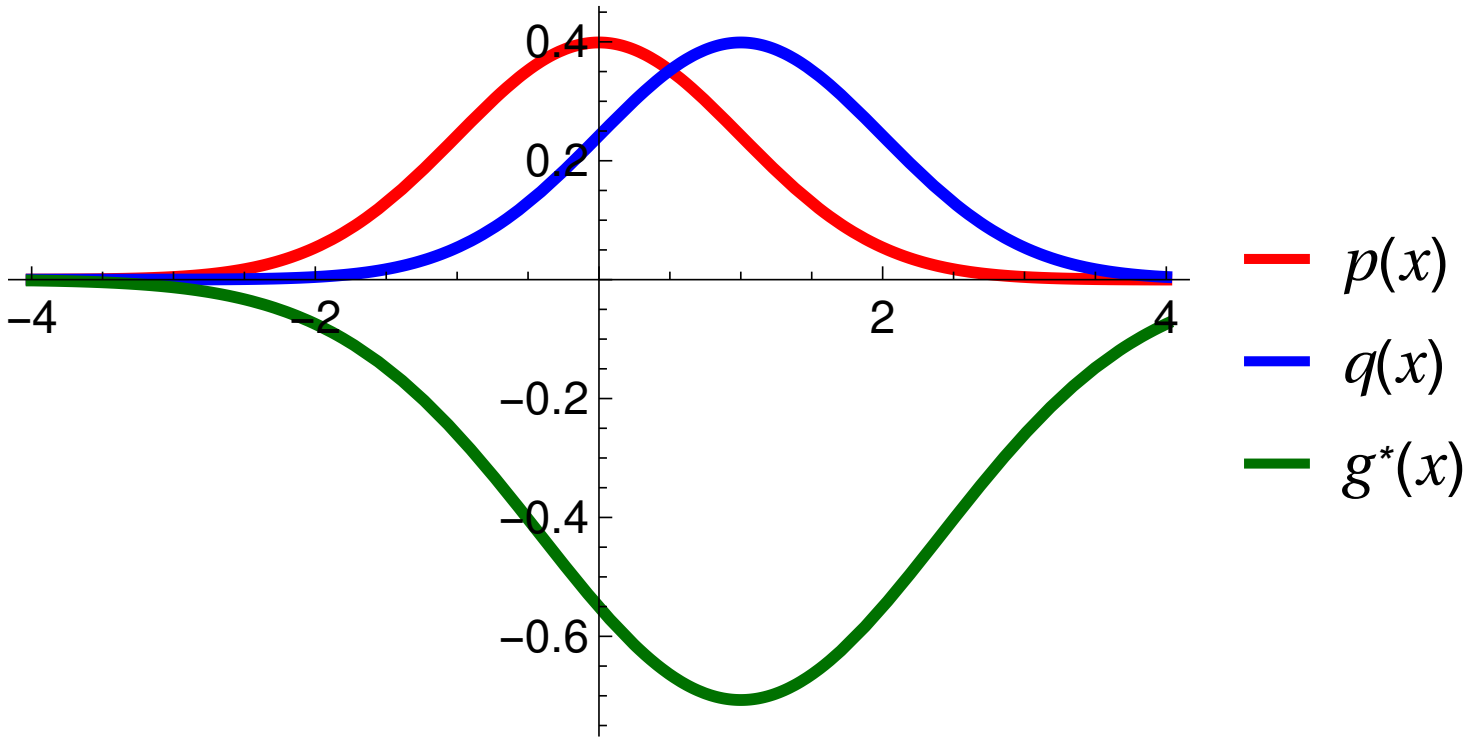

Kernel Stein Discrepancy

Stein operator

$$\mathcal{A}_p f = f(x)\frac{d}{dx}\log p(x) + \frac{d}{dx}f(x)$$

Kernel Stein Discrepancy (KSD)

$$\mathrm{KSD}_p(Q) = \sup_{\|g\|_{\mathcal{F}}\leq 1} \mathrm{E}_q \mathcal{A}_p g - \cancel{\mathrm{E}_p \mathcal{A}_p g} = \sup_{\|g\|_{\mathcal{F}}\leq 1} \mathrm{E}_q \mathcal{A}_p g$$

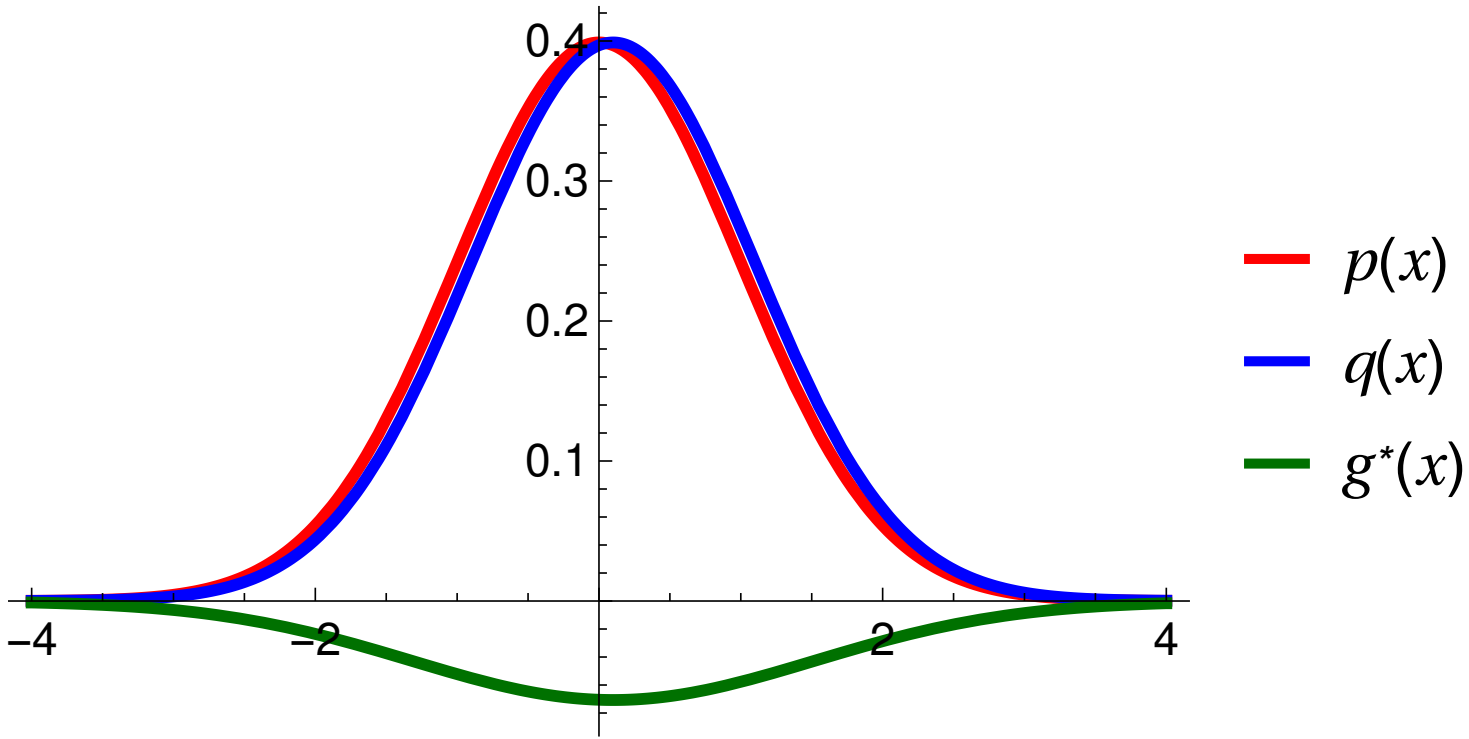

Computing the kernel Stein discrepancy

How do we get the KSD in closed form (with kernels)?

Can we define "Stein features"?

$$[\mathcal{A}_p f](x) = f(x)\frac{d}{dx}\log p(x) + \frac{d}{dx}f(x)$$
$$\stackrel{?}{=} \langle f, \underbrace{\xi(x)}_{\text{stein features}} \rangle_{\mathcal{F}}$$

where $\mathrm{E}_{x\sim p}\xi(x) = 0$.

Stein RKHS features

Reproducing property for the derivative: for differentiable $k(x, x')$,

$$\frac{d}{dx} f(x) = \left\langle f, \frac{d}{dx} \varphi(x) \right\rangle_{\mathcal{F}} \qquad \left\langle \frac{d}{dx} \varphi(x), \varphi(x') \right\rangle_{\mathcal{F}} = \frac{d}{dx} k(x, x')$$

Steinwart, Christmann, Support Vector Machines (2008), Lemma 4.3.4

Stein RKHS features

Reproducing property for the derivative: for differentiable $k(x, x')$,

$$\frac{d}{dx} f(x) = \left\langle f, \frac{d}{dx} \varphi(x) \right\rangle_{\mathcal{F}} \qquad \left\langle \frac{d}{dx} \varphi(x), \varphi(x') \right\rangle_{\mathcal{F}} = \frac{d}{dx} k(x, x')$$

Using kernel derivative trick in (a),

$$\begin{aligned}
\left[\mathcal{A}_p f\right](x) &= \left(\frac{d}{dx} \log p(x) \right) f(x) + \frac{d}{dx} f(x) \\
&= \left\langle f, \left(\frac{d}{dx} \log p(x) \right) \varphi(x) + \underbrace{\frac{d}{dx} \varphi(x)}_{(a)} \right\rangle_{\mathcal{F}} \\
&=: \langle f, \xi(x) \rangle_{\mathcal{F}} .
\end{aligned}$$

Steinwart, Christmann, Support Vector Machines (2008), Lemma 4.3.4

Kernel Stein discrepancy: derivation

Closed-form expression for KSD: given independent $x, x' \sim Q$, then

$$\mathrm{KSD}_p(Q) = \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \left([\mathcal{A}_p g](x) \right)$$

$$= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \langle g, \xi_x \rangle_{\mathcal{F}}$$

$$\underset{(a)}{=} \sup_{\|g\|_{\mathcal{F}} \leq 1} \langle g, \mathrm{E}_{x \sim q} \xi_x \rangle_{\mathcal{F}} = \| \mathrm{E}_{x \sim q} \xi_x \|_{\mathcal{F}}$$

Kernel Stein discrepancy: derivation

Closed-form expression for KSD: given independent $x, x' \sim Q$, then

$$\begin{aligned}
\mathrm{KSD}_p(Q) &= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \left([\mathcal{A}_p g](x) \right) \\
&= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \langle g, \xi_x \rangle_{\mathcal{F}} \\
&\underset{(a)}{=} \sup_{\|g\|_{\mathcal{F}} \leq 1} \langle g, \mathrm{E}_{x \sim q} \xi_x \rangle_{\mathcal{F}} = \| \mathrm{E}_{x \sim q} \xi_x \|_{\mathcal{F}}
\end{aligned}$$

Kernel Stein discrepancy: derivation

Closed-form expression for KSD: given independent $x, x' \sim Q$, then

$$\begin{aligned}\mathrm{KSD}_p(Q) &= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \left([\mathcal{A}_p g](x) \right) \\ &= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \langle g, \xi_x \rangle_{\mathcal{F}} \\ &\underset{(a)}{=} \sup_{\|g\|_{\mathcal{F}} \leq 1} \langle g, \mathrm{E}_{x \sim q} \xi_x \rangle_{\mathcal{F}} = \|\mathrm{E}_{x \sim q} \xi_x\|_{\mathcal{F}}\end{aligned}$$

Caution: (a) requires a condition for Bochner integrability,

$$\mathrm{E}_{x \sim q} \left(\frac{d}{dx} \log p(x) \right)^2 < \infty.$$

Kernel Stein discrepancy: derivation

Closed-form expression for KSD: given independent $x, x' \sim Q$, then

$$\begin{aligned}
\mathrm{KSD}_p(Q) &= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \left([\mathcal{A}_p g](x) \right) \\
&= \sup_{\|g\|_{\mathcal{F}} \leq 1} \mathrm{E}_{x \sim q} \langle g, \xi_x \rangle_{\mathcal{F}} \\
&\underset{(a)}{=} \sup_{\|g\|_{\mathcal{F}} \leq 1} \langle g, \mathrm{E}_{x \sim q} \xi_x \rangle_{\mathcal{F}} = \|\mathrm{E}_{x \sim q} \xi_x\|_{\mathcal{F}}
\end{aligned}$$

Kernel expression:

$$\begin{aligned}
&\|\mathrm{E}_{x \sim q} \xi_x\|_{\mathcal{F}}^2 \\
&= \left\| \mathrm{E}_{x \sim q} \left(\varphi(x) \frac{d}{dx} \log p(x) + \frac{d}{dx} \varphi(x) \right) \right\|_{\mathcal{F}}^2 \\
&= \mathrm{E}_{x, x' \sim Q} \Bigg(k(x, x') \frac{\partial p(x)}{p(x)} \frac{\partial p(x')}{p(x')} + \partial_1 k(x, x') \frac{\partial p(x')}{p(x')} \\
&\qquad + \partial_2 k(x, x') \frac{\partial p(x)}{p(x)} + \partial_{12} k(x, x') \Bigg)
\end{aligned}$$

Does the Bochner condition matter?

Consider the standard normal,

$$p(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-x^2/2\right).$$

Then

$$\frac{d}{dx} \log p(x) = -x.$$

If q is a Cauchy distribution, then the integral

$$\mathrm{E}_{x\sim q}\left(\frac{d}{dx}\log p(x)\right)^2 = \int_{-\infty}^{\infty} x^2 q(x)\, dx$$

is undefined.

Does the Bochner condition matter?

Consider the standard normal,

$$p(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-x^2/2\right).$$

Then

$$\frac{d}{dx} \log p(x) = -x.$$

If q is a Cauchy distribution, then the integral

$$\mathrm{E}_{x\sim q}\left(\frac{d}{dx} \log p(x)\right)^2 = \int_{-\infty}^{\infty} x^2 q(x)\, dx$$

is undefined.

Kernel Stein discrepancy: population expression

Population kernel Stein discrepancy (in $\mathbb{R}^D$):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') + \mathrm{s}_p(x)^\top k_2(x, x') + \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

- $\mathrm{s}_p(x) \in \mathbb{R}^D = \frac{\nabla p(x)}{p(x)}$
- $k_1(a, b) := \nabla_x k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
  $k_2(a, b) := \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
- $k_{12}(a, b) := \nabla_x \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^{D \times D}$

Kernel Stein discrepancy: population expression

Population kernel Stein discrepancy (in $\mathbb{R}^D$):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') + \mathrm{s}_p(x)^\top k_2(x, x') + \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

- $\mathrm{s}_p(x) \in \mathbb{R}^D = \frac{\nabla p(x)}{p(x)}$
- $k_1(a, b) := \nabla_x k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
  $k_2(a, b) := \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
- $k_{12}(a, b) := \nabla_x \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^{D \times D}$

Kernel Stein discrepancy: population expression

Population kernel Stein discrepancy (in $\mathbb{R}^D$):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') + \mathrm{s}_p(x)^\top k_2(x, x') + \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

- $\mathrm{s}_p(x) \in \mathbb{R}^D = \frac{\nabla p(x)}{p(x)}$
- $k_1(a, b) := \nabla_x k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
  $k_2(a, b) := \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^D$,
- $k_{12}(a, b) := \nabla_x \nabla_{x'} k(x, x')|_{x=a, x'=b} \in \mathbb{R}^{D \times D}$

If kernel is C_0-universal and Q satisfies $\mathrm{E}_{x \sim Q} \left\| \nabla \left(\log \frac{p(x)}{q(x)} \right) \right\|^2 < \infty$, then $\mathrm{KSD}_p^2(Q) = 0$ iff $P = Q$.

KSD for discrete-valued variables

Discrete domains: $\mathcal{X} = \{1, \ldots, L\}^D$ with $L \in \mathbb{N}$.

The population KSD (discrete):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') - \mathrm{s}_p(x)^\top k_2(x, x') - \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

$k_1(x, x') = \Delta_x^{-1} k(x, x')$, Δ_x^{-1} is difference on x, $\mathrm{s}_p(x) = \frac{\Delta p(x)}{p(x)}$

Ranganath et al. (NeurIPS 2016), Yang et al. (ICML 2018)

KSD for discrete-valued variables

Discrete domains: $\mathcal{X} = \{1, \ldots, L\}^D$ with $L \in \mathbb{N}$.

The population KSD (discrete):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') - \mathrm{s}_p(x)^\top k_2(x, x') - \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

$k_1(x, x') = \Delta_x^{-1} k(x, x')$, Δ_x^{-1} is difference on x, $\mathrm{s}_p(x) = \frac{\Delta p(x)}{p(x)}$

A discrete kernel: $k(x, x') = \exp\left(-d_H(x, x')\right)$, where $d_H(x, x') = D^{-1} \sum_{d=1}^D \mathbb{I}(x_d \neq x'_d)$.

Ranganath et al. (NeurIPS 2016), Yang et al. (ICML 2018)

KSD for discrete-valued variables

Discrete domains: $\mathcal{X} = \{1, \ldots, L\}^D$ with $L \in \mathbb{N}$.

The population KSD (discrete):

$$\mathrm{KSD}_p^2(Q) = \mathrm{E}_{x,x' \sim Q} h_p(x, x')$$

where

$$h_p(x, x') = \mathrm{s}_p(x)^\top \mathrm{s}_p(x') k(x, x') - \mathrm{s}_p(x)^\top k_2(x, x') - \mathrm{s}_p(x')^\top k_1(x, x') + \mathrm{tr}\left[k_{12}(x, x')\right]$$

$k_1(x, x') = \Delta_x^{-1} k(x, x')$, Δ_x^{-1} is difference on x, $\mathrm{s}_p(x) = \frac{\Delta p(x)}{p(x)}$

A discrete kernel: $k(x, x') = \exp\left(-d_H(x, x')\right)$, where $d_H(x, x') = D^{-1} \sum_{d=1}^{D} \mathbb{I}(x_d \neq x'_d)$.

$\mathrm{KSD}_p^2(Q) = 0$ iff $P = Q$ if

- Gram matrix over all the configurations in $\mathcal{X}$ is strictly positive definite,
- $P > 0$ and $Q > 0$.

Ranganath et al. (NeurIPS 2016), Yang et al. (ICML 2018)

Empirical statistic, asymptotic normality for $P \neq Q$

The empirical statistic:

$$\widehat{\mathrm{KSD}}_p^2(Q) := \frac{1}{n(n-1)} \sum_{i \neq j} h_p(x_i, x_j).$$

Empirical statistic, asymptotic normality for $P \neq Q$

The empirical statistic:

$$\widehat{\mathrm{KSD}}_p^2(Q) := \frac{1}{n(n-1)} \sum_{i \neq j} h_p(x_i, x_j).$$

Asymptotic distribution when $P \neq Q$:

$$\sqrt{n}\left(\widehat{\mathrm{KSD}}_p^2(Q) - \mathrm{KSD}_p(Q)\right) \xrightarrow{d} \mathcal{N}(0, \sigma_{h_p}^2) \qquad \sigma_{h_p}^2 = 4\mathrm{Var}[\mathbb{E}_{x'}[h_p(x, x')]].$$

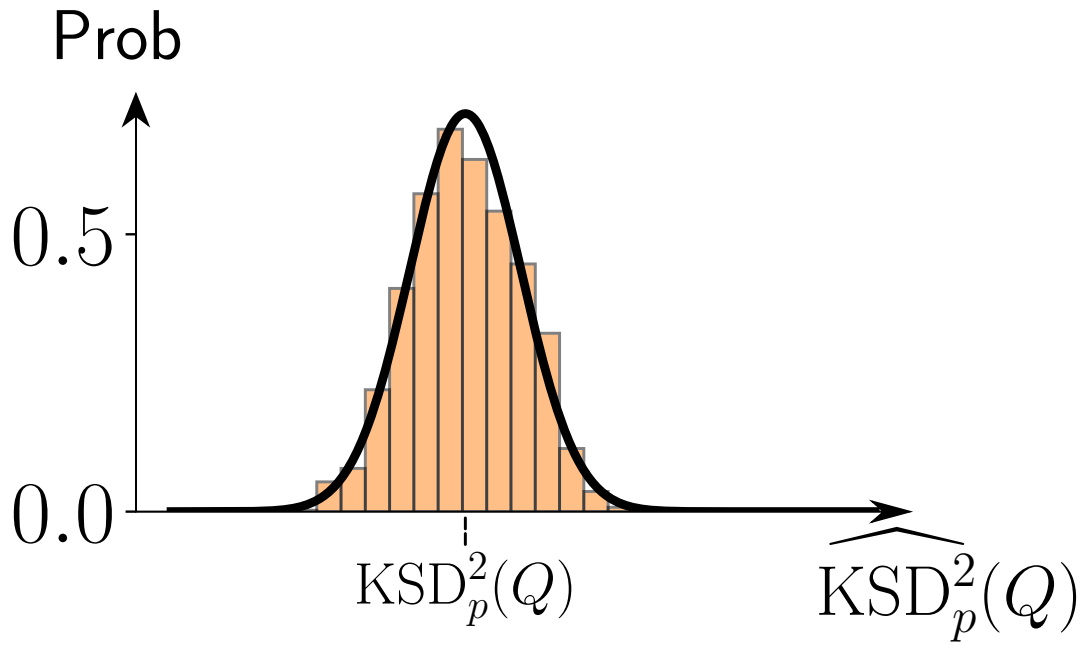

Relative goodness-of-fit testing

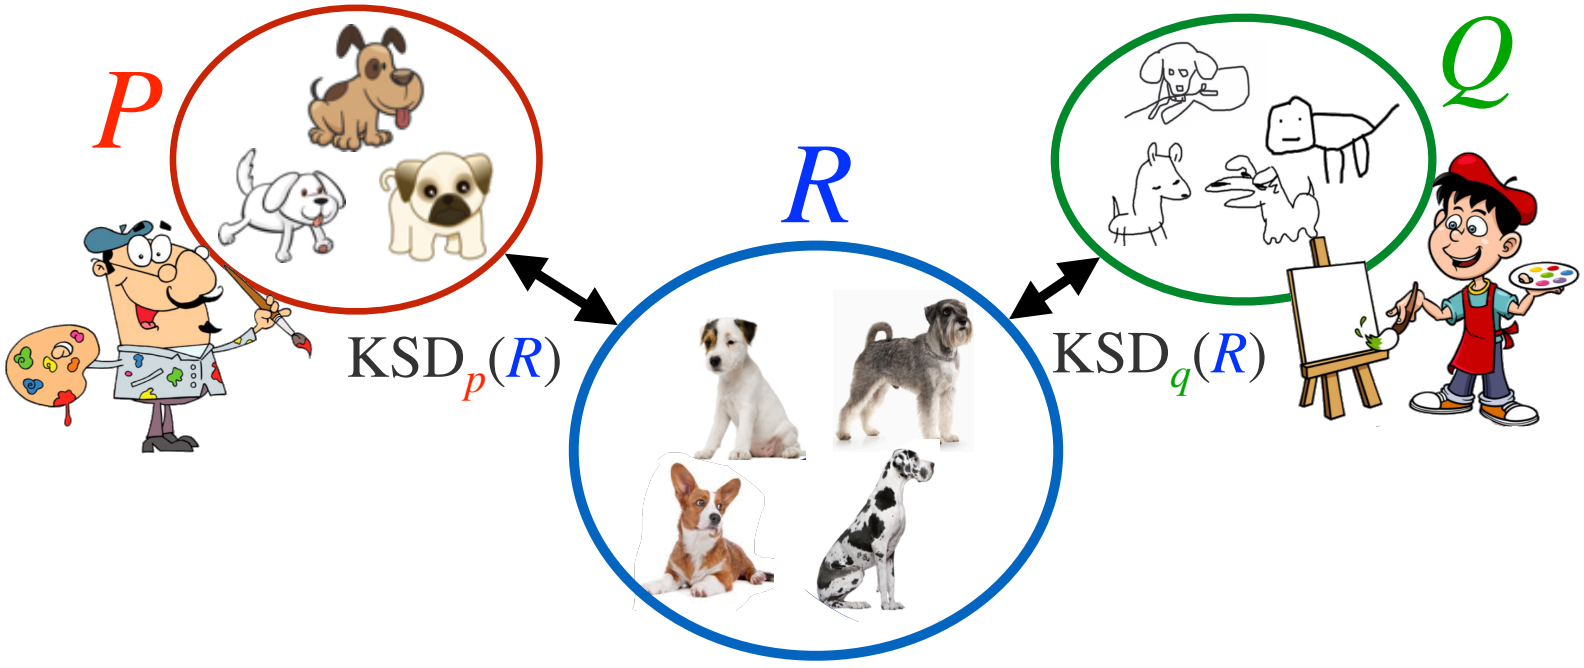

- Two latent variable models P and Q, data $\{x_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} R$.
- Distinct models $p \neq q$

Hypotheses:

$$H_0 : \mathrm{KSD}_p(R) \leq \mathrm{KSD}_q(R) \text{ vs. } H_1 : \mathrm{KSD}_p(R) > \mathrm{KSD}_q(R)$$

(H_0 : 'P is as good as Q, or better' vs. H_1 : 'Q is better')

Relative GOF testing: joint asymptotic normality

Joint asymptotic normality when $P \neq R$ and $Q \neq R$

$$\sqrt{n}\begin{bmatrix} \widehat{\mathrm{KSD}}_p^2(R) - \mathrm{KSD}_p(R) \\ \widehat{\mathrm{KSD}}_q^2(R) - \mathrm{KSD}_q(R) \end{bmatrix} \xrightarrow{d} \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{h_p}^2 & \sigma_{h_p h_q} \\ \sigma_{h_p h_q} & \sigma_{h_q}^2 \end{bmatrix}\right)$$

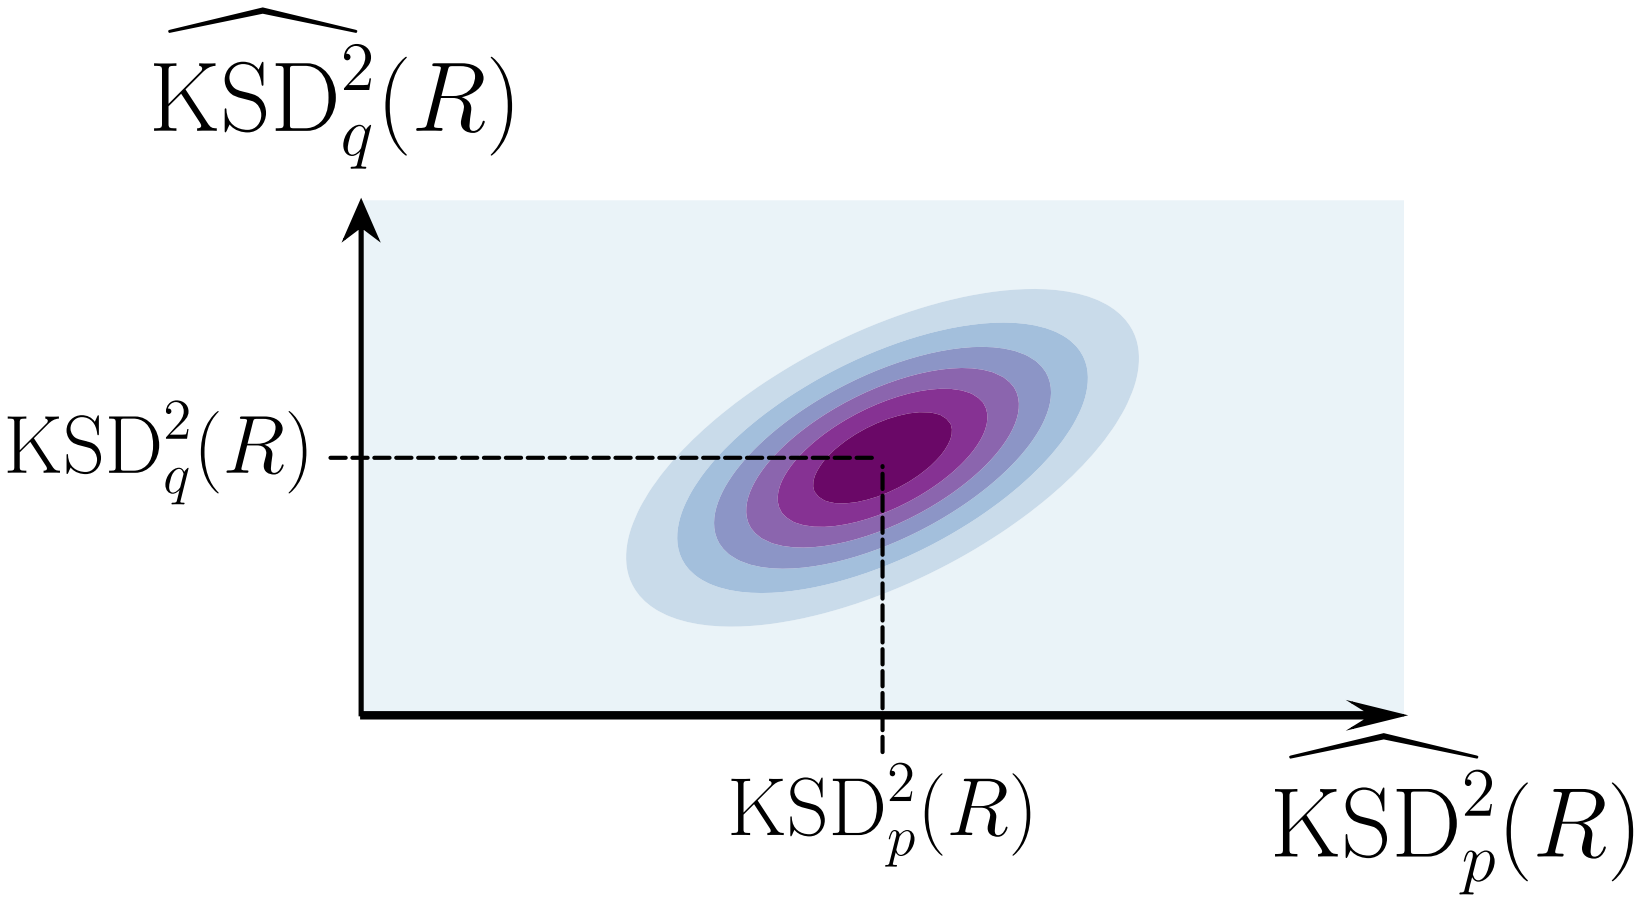

Relative GOF testing: joint asymptotic normality

Joint asymptotic normality when $P \neq R$ and $Q \neq R$

$$\sqrt{n}\begin{bmatrix} \widehat{\mathrm{KSD}_p^2}(R) - \mathrm{KSD}_p(R) \\ \widehat{\mathrm{KSD}_q^2}(R) - \mathrm{KSD}_q(R) \end{bmatrix} \xrightarrow{d} \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{h_p}^2 & \sigma_{h_p h_q} \\ \sigma_{h_p h_q} & \sigma_{h_q}^2 \end{bmatrix}\right)$$

Difference in statistics is asymptotically normal:

$$\sqrt{n}\left[\widehat{\mathrm{KSD}_p^2}(R) - \widehat{\mathrm{KSD}_q^2}(R) - (\mathrm{KSD}_p(R) - \mathrm{KSD}_q(R))\right]$$
$$\xrightarrow{d} \mathcal{N}\left(0, \sigma_{h_p}^2 + \sigma_{h_q}^2 - 2\sigma_{h_p h_q}\right)$$

$\implies$ a statistical test with null hypothesis $\mathrm{KSD}_p(R) - \mathrm{KSD}_q(R) \leq 0$ is straightforward.

Latent variable models

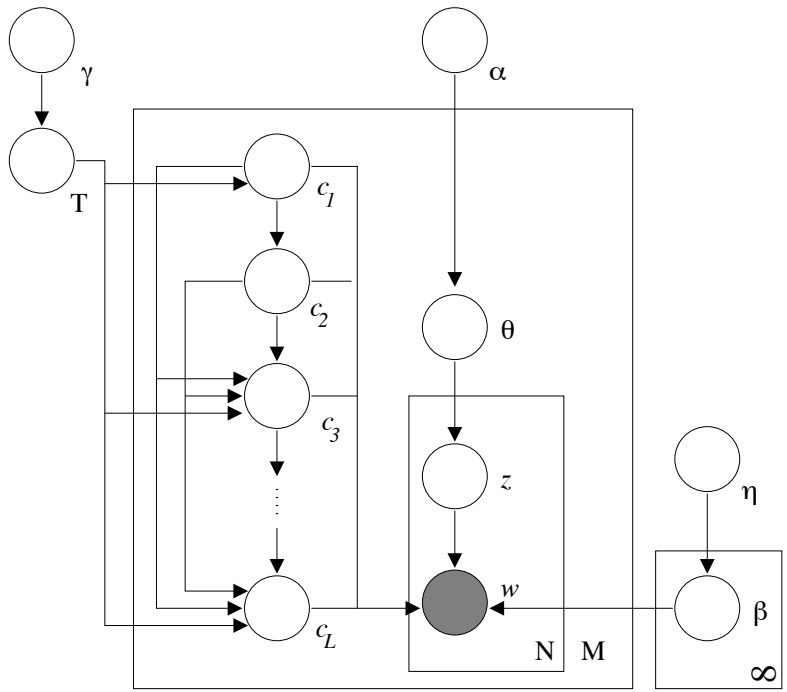

Latent variable models

Can we compare latent variable models with KSD?

$$p(x) = \int p(x|z)p(z)\,dz$$

$$q(x) = \int q(x|w)p(w)\,dw$$

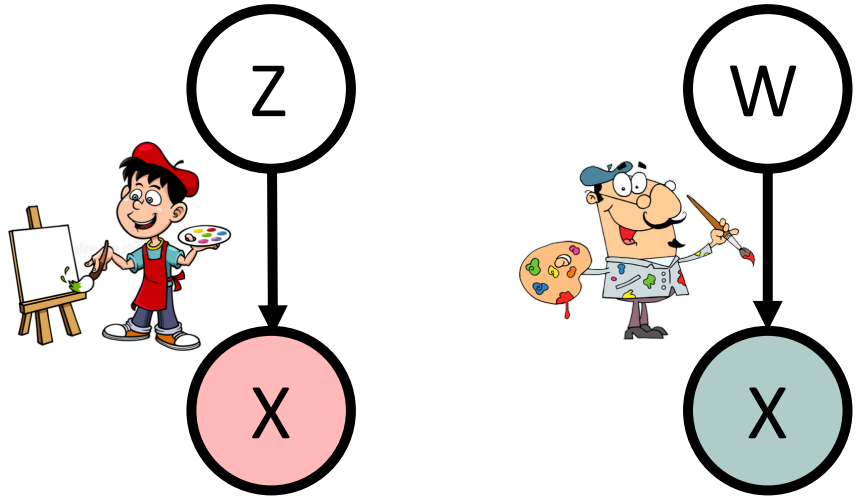

Multi-dimensional Stein operator:

$$[T_p f](x) = \left\langle f(x), \underbrace{\frac{\nabla p(x)}{p(x)}}_{(a)} \right\rangle + \langle \nabla, f(x) \rangle \,.$$

Expression (a) requires marginal $p(x)$, often intractable...

What not to do

Approximate the integral using $\{z_j\}_{j=1}^m \sim p(z)$:

$$p(x) = \int p(x|z)p(z)\,dz$$
$$\approx p_m(x) = \frac{1}{m}\sum_{j=1}^{m} p(x|z_j)$$

Estimate KSD with approximate density:

$$\widehat{\mathrm{KSD}^2_p}(R) \approx \widehat{\mathrm{KSD}^2_{p_m}}(R)$$

What not to do

Approximate the integral using $\{z_j\}_{j=1}^m \sim p(z)$:

$$p(x) = \int p(x|z)p(z)dz$$
$$\approx p_m(x) = \frac{1}{m}\sum_{j=1}^{m} p(x|z_j)$$

Estimate KSD with approximate density:

$$\widehat{\mathrm{KSD}^2_p}(R) \approx \widehat{\mathrm{KSD}^2_{p_m}}(R)$$

Problem: $\widehat{\mathrm{KSD}^2_{p_m}}(R)$ asymptotically normal but slow bias decay.

MCMC approximation of score function

Result we use:

$$\mathrm{s}_p(x) = \mathbb{E}_{z|x}[\mathrm{s}_p(x|z)]$$

Proof:

$$\begin{aligned}\mathrm{s}_p(x) = \frac{\nabla p(x)}{p(x)} &= \frac{1}{p(x)} \int \nabla p(x|z) \mathrm{d}p(z) \\ &= \int \frac{\nabla p(x|z)}{p(x|z)} \cdot \frac{p(x|z)\,dp(z)}{p(x)} = \mathbb{E}_{z|x}[\mathrm{s}_p(x|z)],\end{aligned}$$

Friel, N., Mira, A. and Oates, C. J. (2016) Exploiting multi-core architectures for reduced-variance estimation with intractable likelihoods. Bayesian Analysis, 11, 215–245.

MCMC approximation of score function

Result we use:

$$\mathrm{s}_p(x) = \mathbb{E}_{z|x}[\mathrm{s}_p(x|z)]$$

Proof:

$$\mathrm{s}_p(x) = \frac{\nabla p(x)}{p(x)} = \frac{1}{p(x)} \int \nabla p(x|z) \mathrm{d}p(z)$$
$$= \int \frac{\nabla p(x|z)}{p(x|z)} \cdot \frac{p(x|z)\,dp(z)}{p(x)} = \mathbb{E}_{z|x}[\mathrm{s}_p(x|z)],$$

Approximate intractable posterior $\mathbb{E}_{z|x_i}[\mathrm{s}_p(x_i|z)]$

$$\bar{\mathrm{s}}_p(x_i; z_i^{(t)}) := \frac{1}{m} \sum_{j=1}^{m} \mathrm{s}_p(x_i|z_{i,j}^{(t)}) \approx \mathrm{s}_p(x_i)$$

with $z_i^{(t)} = (z_{i,1}^{(t)}, \ldots, z_{i,m}^{(t)})$ via MCMC (after t burn-in steps)

Friel, N., Mira, A. and Oates, C. J. (2016) Exploiting multi-core architectures for reduced-variance estimation with intractable likelihoods. Bayesian Analysis, 11, 215–245.

KSD for latent variable models

Recall earlier KSD estimate:

$$U_n(P) = \frac{1}{n(n-1)} \sum_{i \neq j} h_p(x_i, x_j) \ (\approx \mathrm{KSD}_p^2(R))$$

KSD for latent variable models

Recall earlier KSD estimate:

$$U_n(P) = \frac{1}{n(n-1)} \sum_{i \neq j} h_p(x_i, x_j) \ (\approx \mathrm{KSD}_p^2(R))$$

KSD estimate for latent variable models:

$$U_n^{(t)}(P) := \frac{1}{n(n-1)} \sum_{i \neq j} \bar{H}_p\left[(x_i, z_i^{(t)}), (x_j, z_j^{(t)})\right] \ (\approx \mathrm{KSD}_p^2(R))$$

where $\bar{H}_p$ is the Stein kernel h_p with $\mathrm{s}_p(x_i)$ replaced with $\bar{\mathrm{s}}_p(x_i; z_i^{(t)})$.

Return to relative GOF test, latent variable models

Hypotheses:

$$H_0 : \mathrm{KSD}_p(R) \leq \mathrm{KSD}_q(R) \text{ vs. } H_1 : \mathrm{KSD}_p(R) > \mathrm{KSD}_q(R)$$

(H_0 : 'P is as good as Q, or better' vs. H_1 : 'Q is better')

Return to relative GOF test, latent variable models

Hypotheses:

$$H_0 : \mathrm{KSD}_p(R) \leq \mathrm{KSD}_q(R) \text{ vs. } H_1 : \mathrm{KSD}_p(R) > \mathrm{KSD}_q(R)$$

(H_0 : 'P is as good as Q, or better' vs. H_1 : 'Q is better')

Strategy:

- Estimate the difference $\mathrm{KSD}_p^2(R) - \mathrm{KSD}_q^2(R)$ by

$$D_n^{(t)}(P, Q) = U_n^{(t)}(P) - U_n^{(t)}(Q).$$

- If $D_n^{(t)}(P, Q)$ is sufficiently large, reject H_0.
  - "Sufficient": control type-I error (falsely rejecting H_0)
  - Requires the (asymptotic) behaviour of $D_n^{(t)}(P, Q)$

Asymptotic distribution for relative KSD test

Asymptotic distribution of approximate KSD estimate $n, t \to \infty$:

$$\sqrt{n}\left[D_n^{(t)}(P, Q) - \mu_{PQ}\right] \xrightarrow{d} \mathcal{N}(0, \sigma_{PQ}^2)$$

where

$$\mu_{PQ} = \mathrm{KSD}_p^2(R) - \mathrm{KSD}_q^2(R),$$
$$\sigma_{PQ}^2 = \lim_{n,t\to\infty} n \cdot \mathrm{Var}\left[D_n^{(t)}(P, Q)\right].$$

Fine print:

- The double limit requires fast bias decay $\sqrt{n}\left[\mathbb{E}\{D_n^{(t)}(P, Q)\} - \mu_{PQ}\right] \to 0$
- The fourth moment of $\bar{H}_p^{(t)} - \bar{H}_q^{(t)}$ has finite limit sup. $(t \to \infty)$.

Asymptotic distribution for relative KSD test

Asymptotic distribution of approximate KSD estimate $n, t \to \infty$:

$$\sqrt{n}\left[D_n^{(t)}(P, Q) - \mu_{PQ}\right] \xrightarrow{d} \mathcal{N}(0, \sigma_{PQ}^2)$$

where

$$\mu_{PQ} = \mathrm{KSD}_p^2(R) - \mathrm{KSD}_q^2(R),$$
$$\sigma_{PQ}^2 = \lim_{n,t\to\infty} n \cdot \mathrm{Var}\left[D_n^{(t)}(P, Q)\right].$$

Level-α test:

$$\text{Reject } H_0 \text{ if } D_n^{(t)}(P, Q) \geq \frac{\hat{\sigma}_{PQ}}{\sqrt{n}} c_{1-\alpha}$$

- $c_{1-\alpha}$ is $(1-\alpha)$-quantile of $\mathcal{N}(0, 1)$.
- $\hat{\sigma}_{PQ}$ estimated via jackknife

Experiments

Experiment 1: sensitivity to model difference

- Data R : Probabilistic Principal Component Analysis PPCA(A):

$$x_i \in \mathbb{R}^{100} \sim \mathcal{N}(Az_i, I), \;\; z_i \in \mathbb{R}^{10} \sim \mathcal{N}(0, I_z)$$

- Generate P, Q : perturb $(1,1)$-entry : $A_\delta = A + \delta E_{1,1}$

Experiment 1: sensitivity to model difference

- Data R : Probabilistic Principal Component Analysis PPCA(A):

$$x_i \in \mathbb{R}^{100} \sim \mathcal{N}(Az_i, I), \; z_i \in \mathbb{R}^{10} \sim \mathcal{N}(0, I_z)$$

- Generate P, Q : perturb $(1,1)$-entry : $A_\delta = A + \delta E_{1,1}$

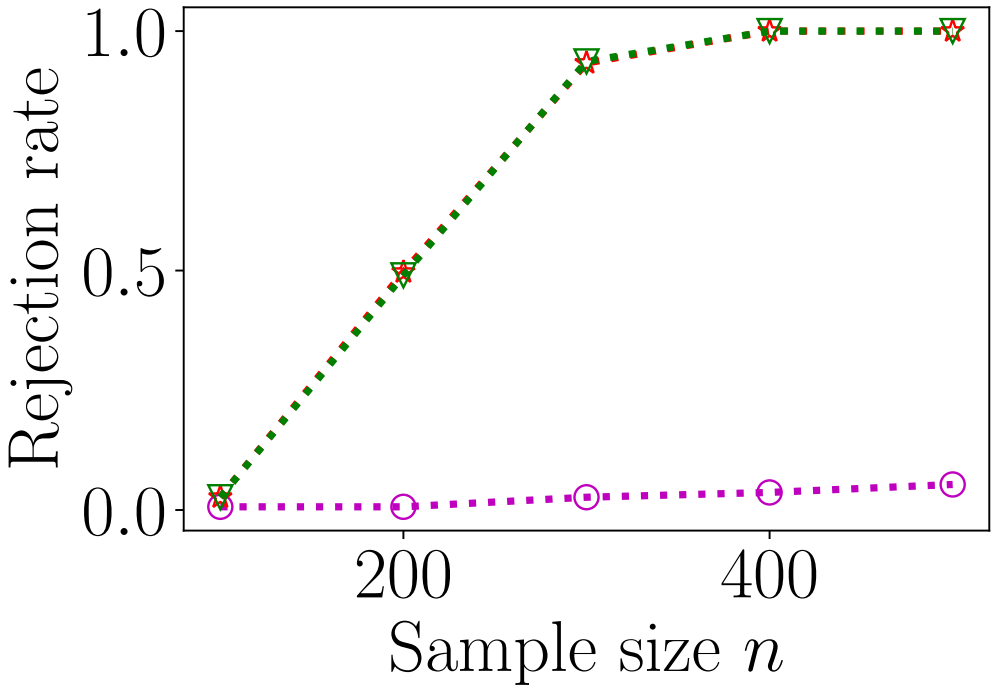

- Alt. H_1 (Q is better):
  - P's perturbation $\delta_P = 2$
  - Q's perturbation $\delta_Q = 1$
- IMQ kernel: $k(x, x') = \left(1 + \|x - x'\|_2^2 / \sigma_{\text{med}}^2\right)^{-1/2}$
- NUTS-HMC with sample size $m = 500$ (after $t = 200$ steps).

MMD KSD LKSD

Hoffman and Gelman (JMLR 2014)

Experiment 1: sensitivity to model difference

- Data R : Probabilistic Principal Component Analysis PPCA(A):

$$x_i \in \mathbb{R}^{100} \sim \mathcal{N}(Az_i, I), \; z_i \in \mathbb{R}^{10} \sim \mathcal{N}(0, I_z)$$

- Generate P, Q : perturb $(1,1)$-entry : $A_\delta = A + \delta E_{1,1}$

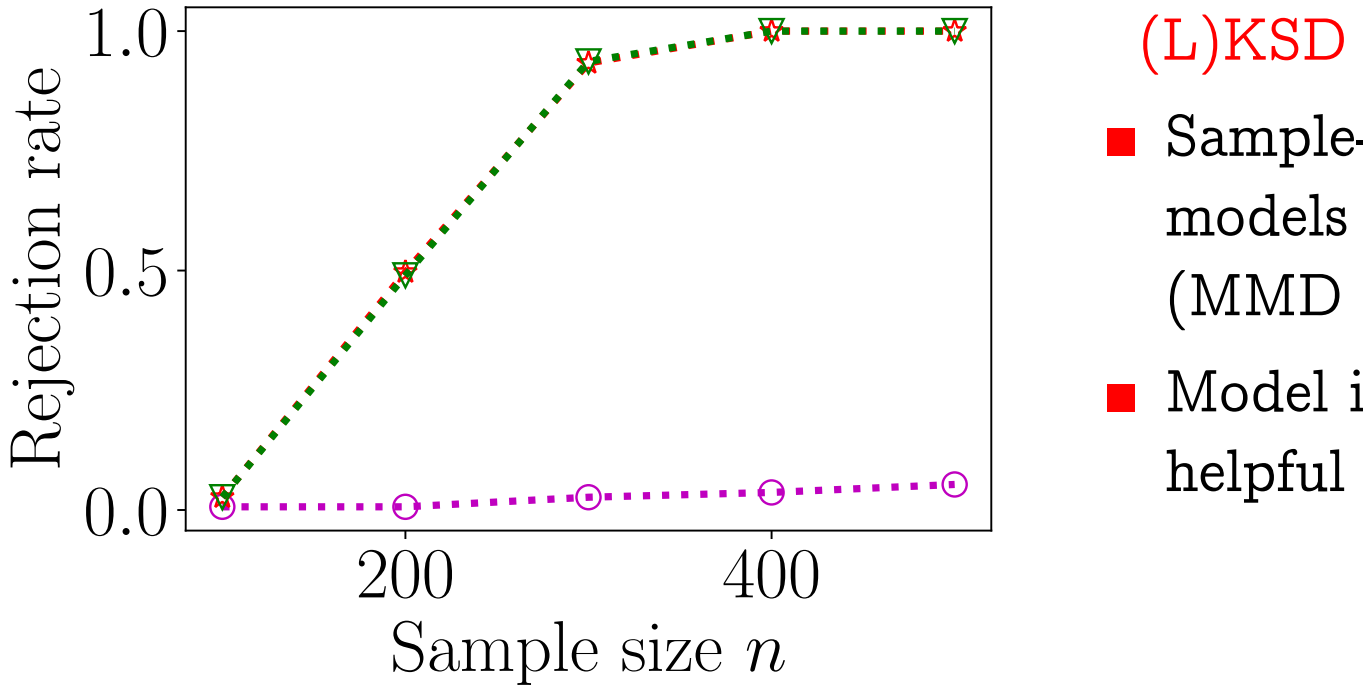

(L)KSD = higher power

- Sample-wise difference in models = subtle (MMD fails)
- Model information is helpful

Hoffman and Gelman (JMLR 2014)

Experiment 2: topic models for arXiv articles

- Data R : arXiv articles from category stat.TH (stat theory) :
- Models P, Q : LDAs trained on articles from different categories
  - P : math.PR (math probability theory)
  - Q : stat.ME (stat methodology). H_1: Q is better

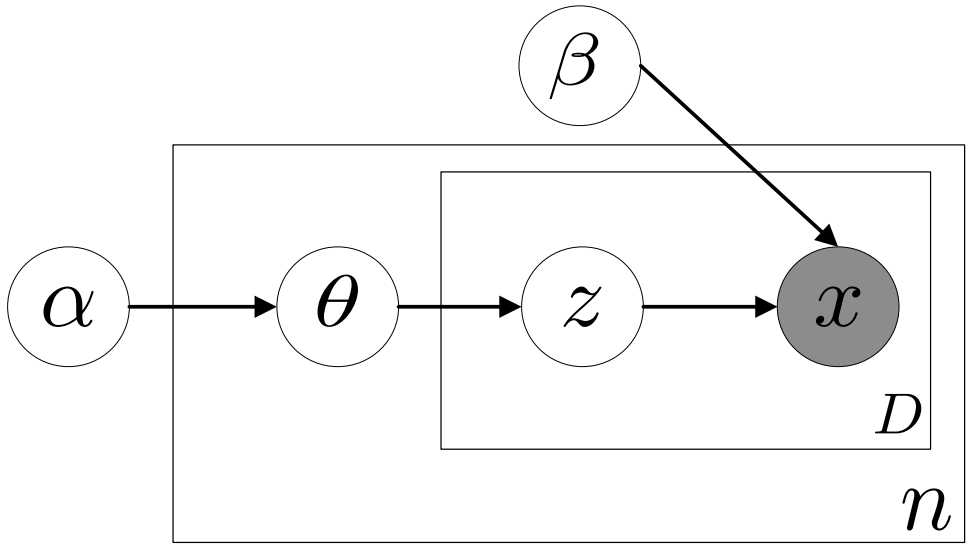

Graphical model of LDA

Blei, Ng, Jordan (JMLR 2003)

Experiment 2: topic models for arXiv articles

- Data R : arXiv articles from category stat.TH (stat theory) :
- Models P, Q : LDAs trained on articles from different categories (100 topics)
  - P : math.PR (math probability theory)
  - Q : stat.ME (stat methodology). H_1: Q is better

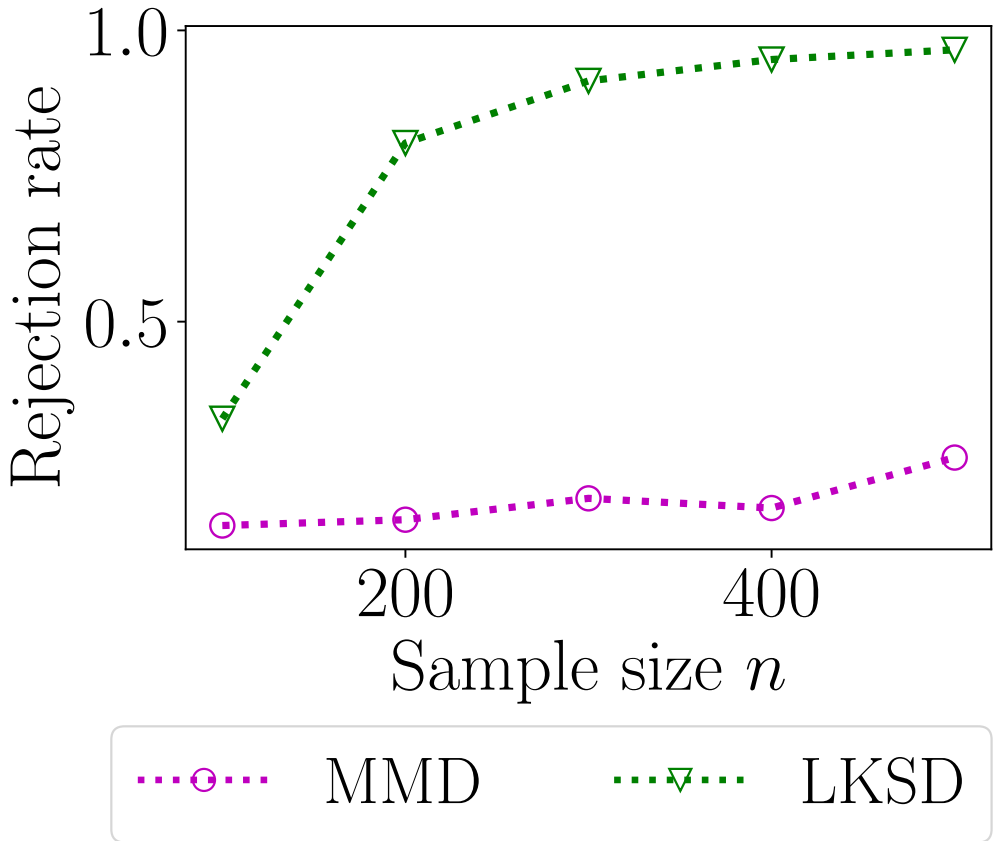

- $\mathcal{X} = \{1, \ldots, L\}^D$, $D = 100$, $L = 126,190$.
- IMQ kernel in BoW rep.: $k(x, x') = (1 + \|B(x) - B(x')\|_2^2)^{-1/2}$
- MCMC size $m = 5000$ (after $t = 500$ steps).

A failure mode

- Data R : arXiv articles from category stat.TH (stat theory) :
- Models P, Q : LDAs trained on articles from different categories (100 topics)
  - P : cs.LG (CS machine learning)
  - Q : stat.ME (stat methodology). H_1: Q is better

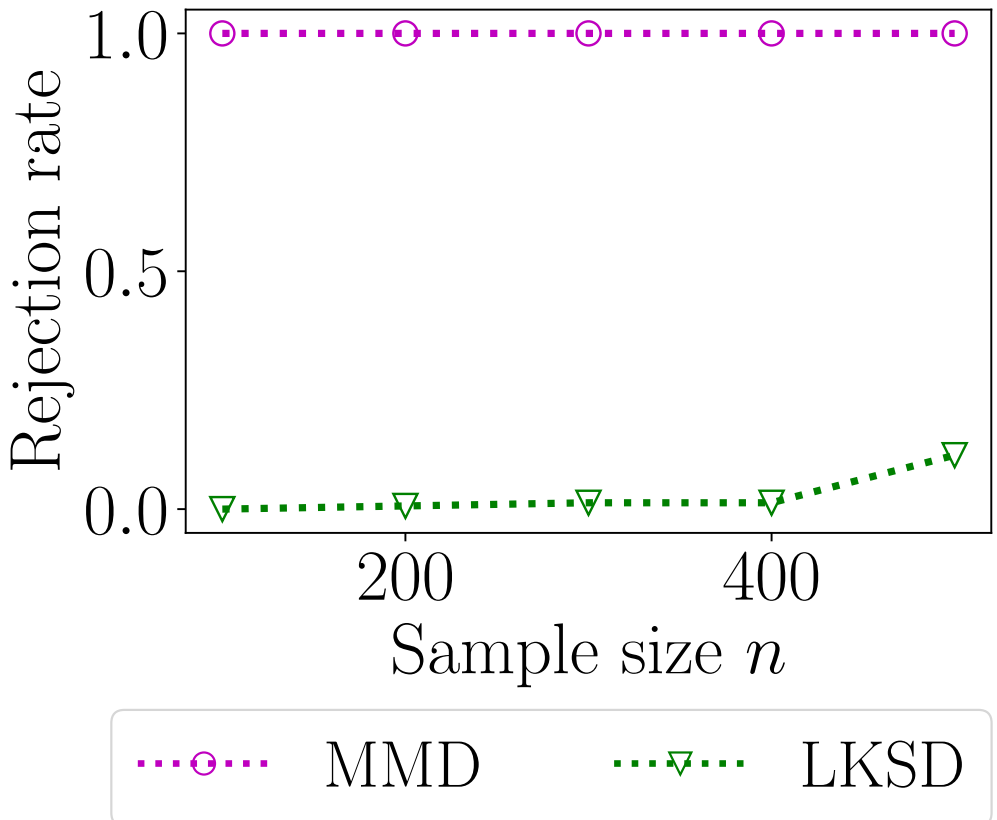

- $\mathcal{X} = \{1, \ldots, L\}^D$, $D = 100$, $L = 208,671$.
- IMQ kernel in BoW rep.: $k(x, x') = (1 + \|B(x) - B(x')\|_2^2)^{-1/2}$
- MCMC size $m = 5000$ (after $t = 500$ steps).

What went wrong?

Recall (one-dimension, informally)

$$\mathrm{s}_p(x) = \frac{p(x+1)}{p(x)} - 1$$

Numerical instability arises when

- Observed word x has low probability
- Word next to x in vocabulary has non-negligible probability

Zanella-Barker Stein operator

Zanella-Barker Stein operator (1-D):

$$\mathcal{A}_p^{\mathrm{ZB}} f(x) = \sum_{\tilde{x} \in \{x+1, x-1\}} \frac{p(\tilde{x})}{p(\tilde{x}) + p(x)} \cdot \{f(\tilde{x}) - f(x)\}$$

- More stable: the ratio $p(\tilde{x})/\{p(\tilde{x}) + p(x)\}$ is always between 0 and 1.
- Similarly applies to latent variable models.

Hodgkinson, Salomone, and Roosta (2020); Shi, Zhou, Hwang, Titsias, and Mackey. (2022)

A resolution to the failure mode

- Data R : arXiv articles from category stat.TH (stat theory) :
- Models P, Q : LDAs trained on articles from different categories (100 topics)
  - P : cs.LG (CS machine learning)
  - Q : stat.ME (stat methodology). H_1: Q is better

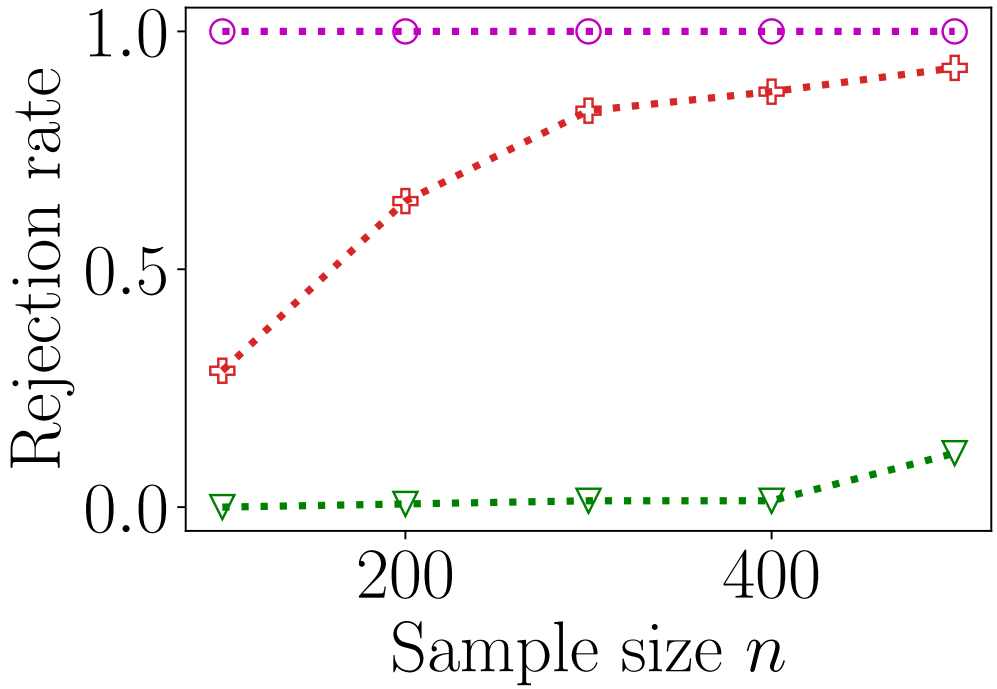

- Improved performance by an alternative Stein operator

Conclusion

Relative goodness-of-fit tests for Models with Latent Variables

- The kernel Stein discrepancy
  - Comparing two models via samples: MMD and the witness function.
  - Comparing a sample and a model: Stein modification of the witness class
- Constructing a relative hypothesis test using the KSD
- Relative hypothesis tests with latent variables

References

A Kernel Test of Goodness of Fit
Kacper Chwialkowski, Heiko Strathmann, Arthur Gretton
https://arxiv.org/abs/1602.02964

A Kernel Stein Test for Comparing Latent Variable Models
Heishiro Kanagawa, Wittawat Jitkrittum, Lester Mackey,
Kenji Fukumizu, Arthur Gretton
https://arxiv.org/abs/1907.00586

Questions?


THE KERNEL
BREWERY
LONDON
BALTIC PORTER
7.3% ABV
THE KERNEL
EXPORT STOUT

Can sampler influence test power?

How important is the quality of $\frac{1}{m}\sum_{j=1}^{m} \mathrm{s}_p(x|z_j^{(t)})$?

Experiment with PPCA:

- P : MALA with a bad step size (poor sampler)
- Q : NUTS-HMC (good sampler)

Expectation:

> If poor, the test would reject even if P and Q are equally good

Can sampler influence test power?

How important is the quality of $\frac{1}{m}\sum_{j=1}^{m} \mathrm{s}_p(x|z_j^{(t)})$?

Experiment with PPCA:

- P : MALA with a bad step size (poor sampler)
- Q : NUTS-HMC (good sampler)

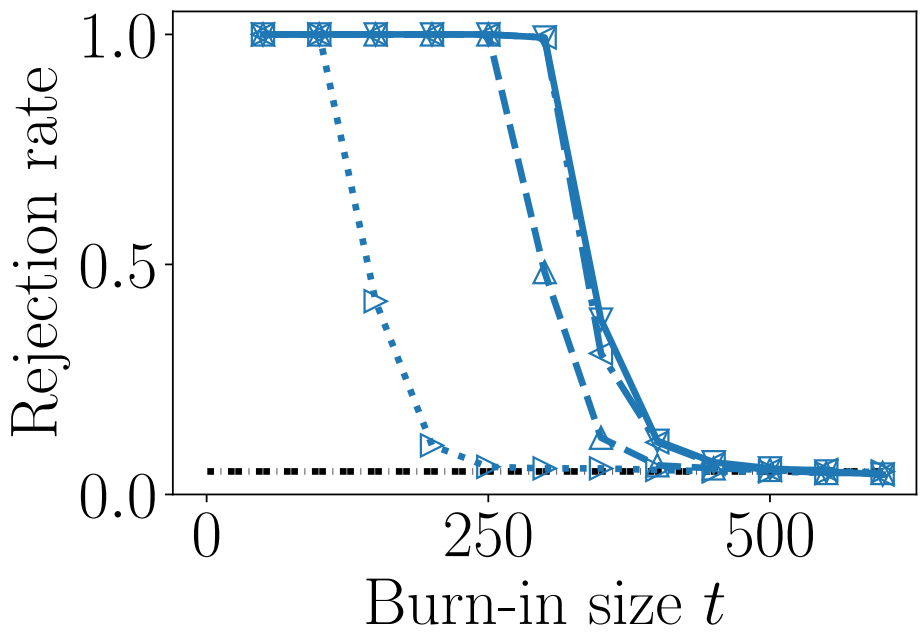

- Null H_0 (should not reject)
- Significance level $\alpha = 0.05$
- Sample size $n = 100$

Can sampler influence test power?

How important is the quality of $\frac{1}{m}\sum_{j=1}^{m} \mathrm{s}_p(x|z_j^{(t)})$?

Experiment with PPCA:

- P : MALA with a bad step size (poor sampler)
- Q : NUTS-HMC (good sampler)

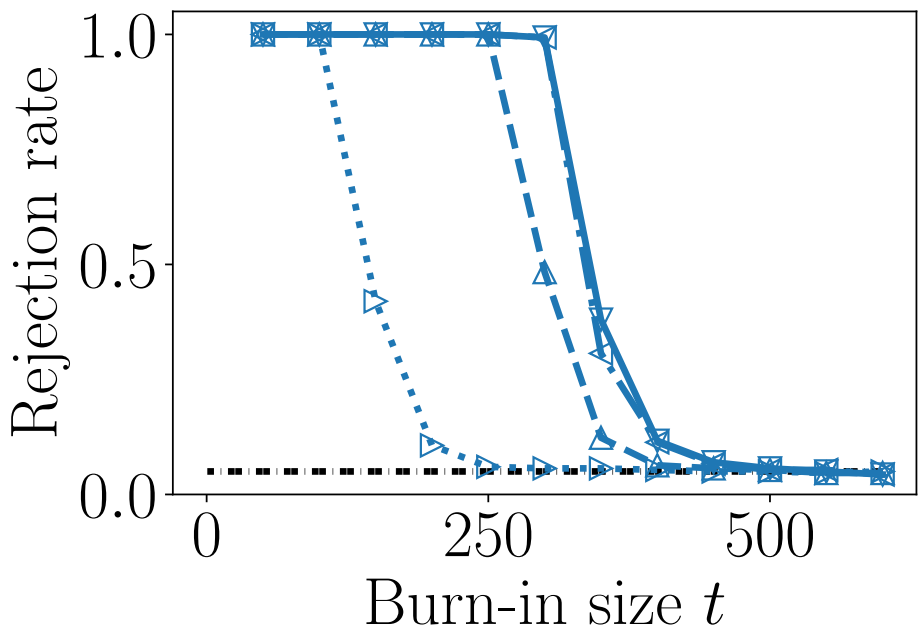

- Null H_0 (should not reject)
- Significance level $\alpha = 0.05$
- Sample size $n = 100$

Sufficient burn-in
$\rightarrow$ correct type-I error

Can sampler influence test power?

How important is the quality of $\frac{1}{m}\sum_{j=1}^{m} \mathrm{s}_p(x|z_j^{(t)})$?

Experiment with PPCA:

- P : MALA with a bad step size (poor sampler)
- Q : NUTS-HMC (good sampler)

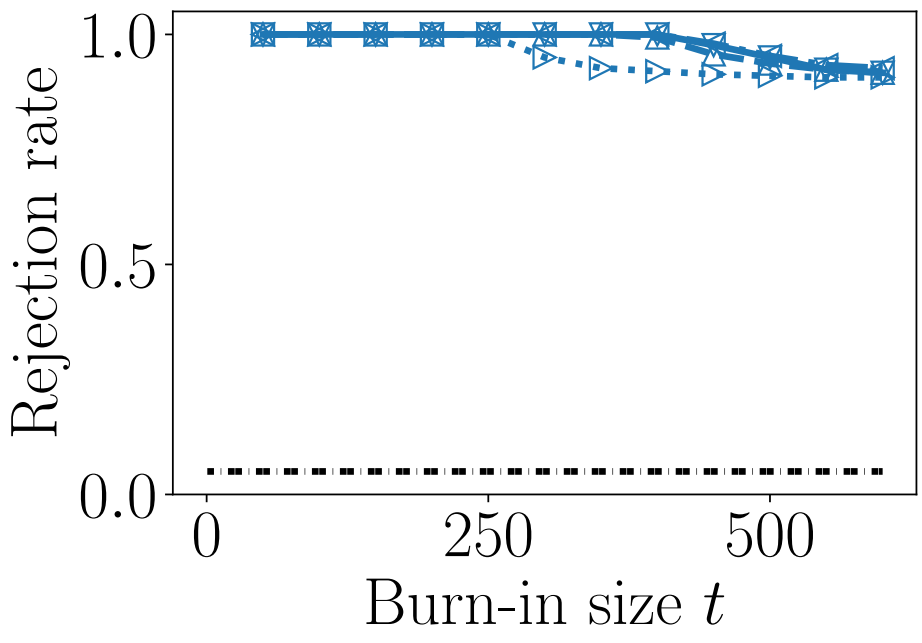

- Null H_0 (should not reject)
- Significance level $\alpha = 0.05$
- Sample size $n = 300$

Can sampler influence test power?

How important is the quality of $\frac{1}{m}\sum_{j=1}^{m} \mathrm{s}_p(x|z_j^{(t)})$?

Experiment with PPCA:

- P : MALA with a bad step size (poor sampler)
- Q : NUTS-HMC (good sampler)

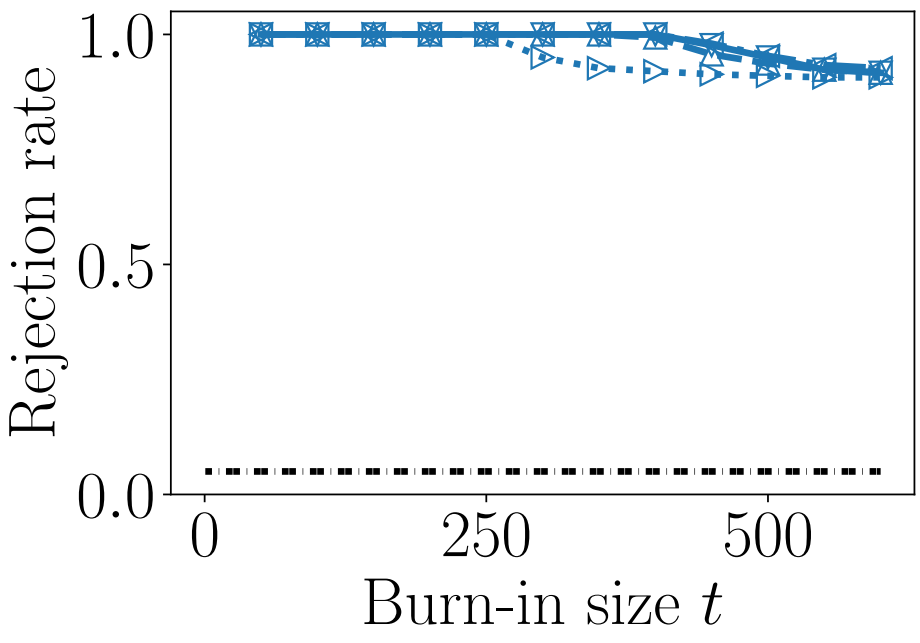

- Null H_0 (should not reject)
- Significance level $\alpha = 0.05$
- Sample size $n = 300$

Larger $n \implies$ more sensitive to mismatch