Representing and comparing probabilities: Part 2

Arthur Gretton

Gatsby Computational Neuroscience Unit,
University College London

Paris, 2018

Testing against a probabilistic model

Statistical model criticism

$$MMD(P, Q) = \|f^*\|^2 = \sup_{\|f\|_{\mathcal{F}} \leq 1} [E_Q f - E_p f]$$

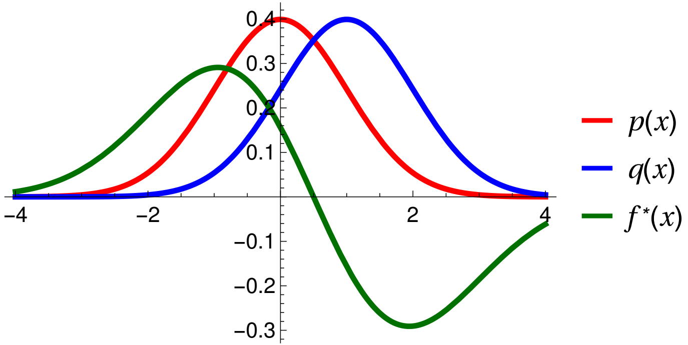

$f^*(x)$ is the witness function

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

Problem: usualy can't compute $E_p f$ in closed form.

Stein idea

To get rid of $E_p f$ in

$$\sup_{\|f\|_{\mathcal{F}} \leq 1} \left[E_q f - E_p f \right]$$

we define the **Stein operator**

$$\left[T_p f \right](x) = \frac{1}{p(x)} \frac{d}{dx} \left(f(x) p(x) \right)$$

Then

$$E_P \, T_P f = 0$$

subject to appropriate boundary conditions. (Oates, Girolami, Chopin, 2016)

Stein idea: proof

$$E_p[T_p f] = \int \left[\frac{1}{p(x)}\frac{d}{dx}(f(x)p(x))\right] p(x)dx$$

$$\int \left[\frac{d}{dx}(f(x)p(x))\right] dx$$

$$= [f(x)p(x)]_{-\infty}^{\infty}$$

$$= 0$$

Stein idea: proof

$$E_p[T_p f] = \int \left[\frac{1}{\cancel{p(x)}}\frac{d}{dx}\left(f(x)p(x)\right)\right]\cancel{p(x)}dx$$

$$\int \left[\frac{d}{dx}\left(f(x)p(x)\right)\right] dx$$

$$= [f(x)p(x)]_{-\infty}^{\infty}$$

$$= 0$$

Stein idea: proof

$$E_p\left[T_p f\right] = \int \left[\frac{1}{\cancel{p(x)}}\frac{d}{dx}\left(f(x)p(x)\right)\right]\cancel{p(x)}dx$$

$$\int \left[\frac{d}{dx}\left(f(x)p(x)\right)\right] dx$$

$$= \left[f(x)p(x)\right]_{-\infty}^{\infty}$$

$$= 0$$

Stein idea: proof

$$E_p\left[T_p f\right] = \int \left[\frac{1}{\cancel{p(x)}}\frac{d}{dx}\left(f(x)p(x)\right)\right]\cancel{p(x)}dx$$

$$\int \left[\frac{d}{dx}\left(f(x)p(x)\right)\right] dx$$

$$= \left[f(x)p(x)\right]_{-\infty}^{\infty}$$

$$= 0$$

Stein idea: proof

$$
\begin{aligned}
E_p\left[T_p f\right] &= \int \left[\frac{1}{\cancel{p(x)}} \frac{d}{dx}\left(f(x)p(x)\right)\right] \cancel{p(x)} dx \\
&\int \left[\frac{d}{dx}\left(f(x)p(x)\right)\right] dx \\
&= \left[f(x)p(x)\right]_{-\infty}^{\infty} \\
&= 0
\end{aligned}
$$

Kernel Stein Discrepancy

Stein operator

$$T_p f = \frac{1}{p(x)} \frac{d}{dx} (f(x) p(x))$$

Kernel Stein Discrepancy (KSD)

$$KSD(p, q, \mathcal{F}) = \sup_{\|g\|_{\mathcal{F}} \le 1} E_q T_p g - E_p T_p g$$

Kernel Stein Discrepancy

Stein operator

$$T_p f = \frac{1}{p(x)} \frac{d}{dx}\left(f(x)p(x)\right)$$

Kernel Stein Discrepancy (KSD)

$$KSD(p, q, \mathcal{F}) = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g - \cancel{E_p T_p g} = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g$$

Kernel Stein Discrepancy

Stein operator

$$T_p f = \frac{1}{p(x)} \frac{d}{dx}(f(x)p(x))$$

Kernel Stein Discrepancy (KSD)

$$KSD(p, q, \mathcal{F}) = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g - \cancel{E_p T_p g} = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g$$

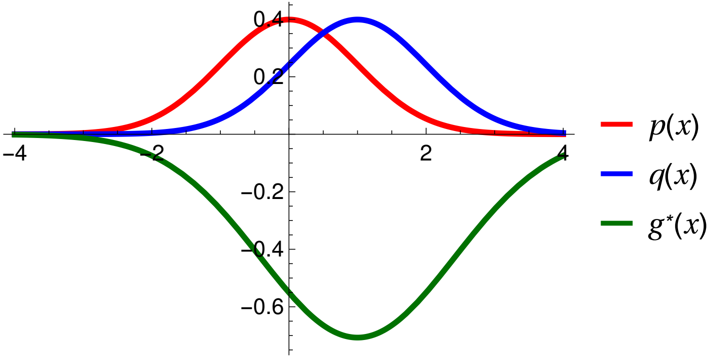

Kernel Stein Discrepancy

Stein operator

$$T_p f = \frac{1}{p(x)} \frac{d}{dx}\left(f(x)p(x)\right)$$

Kernel Stein Discrepancy (KSD)

$$KSD(p, q, \mathcal{F}) = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g - \cancel{E_p T_p g} = \sup_{\|g\|_{\mathcal{F}} \leq 1} E_q T_p g$$

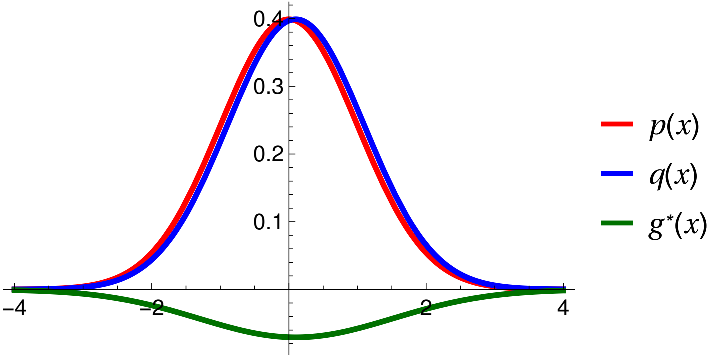

Kernel stein discrepancy

Closed-form expression for KSD: given $Z, Z' \sim q$, then
(Chwialkowski, Strathmann, G., ICML 2016) (Liu, Lee, Jordan ICML 2016)

$$\mathrm{KSD}(p, q, \mathcal{F}) = E_q h_p(Z, Z')$$

where

$$\begin{aligned} h_p(x, y) := {} & \partial_x \log p(x) \partial_x \log p(y) k(x, y) \\ & + \partial_y \log p(y) \partial_x k(x, y) \\ & + \partial_x \log p(x) \partial_y k(x, y) \\ & + \partial_x \partial_y k(x, y) \end{aligned}$$

and k is RKHS kernel for $\mathcal{F}$

Only depends on kernel and $\partial_x \log p(x)$. Do not need to normalize p, or sample from it.

Statistical model criticism

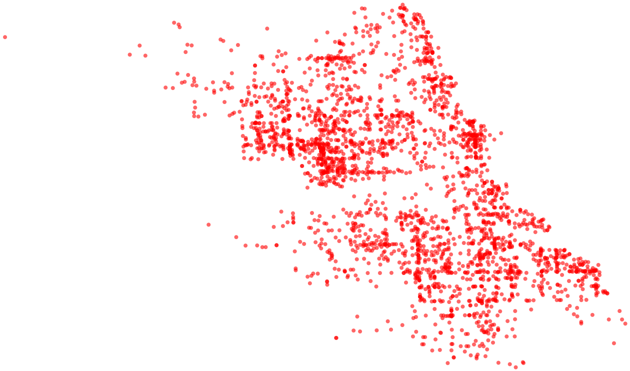

Chicago crime data

Statistical model criticism

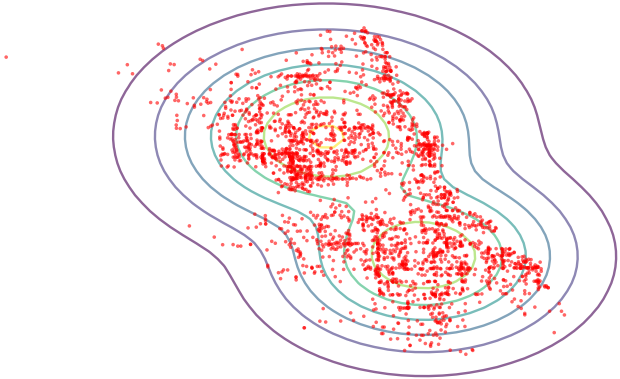

Chicago crime data

Model is Gaussian mixture with two components.

Statistical model criticism

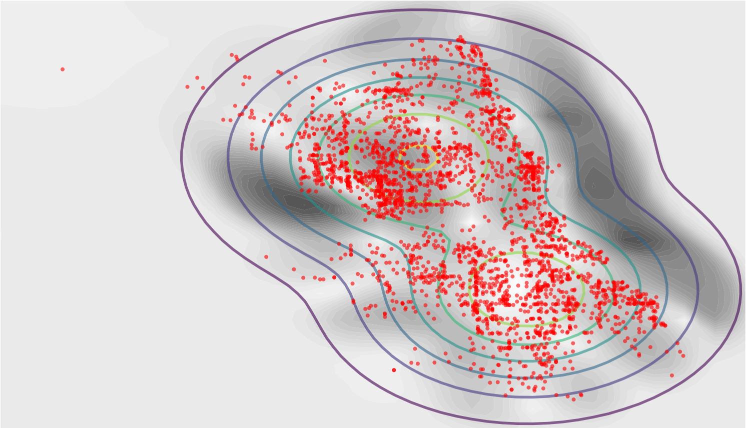

Chicago crime data

Model is Gaussian mixture with two components

Stein **witness function**

Statistical model criticism

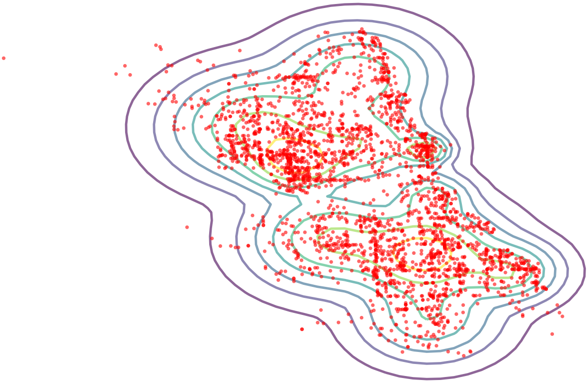

Chicago crime data

Model is Gaussian mixture with ten components.

Statistical model criticism

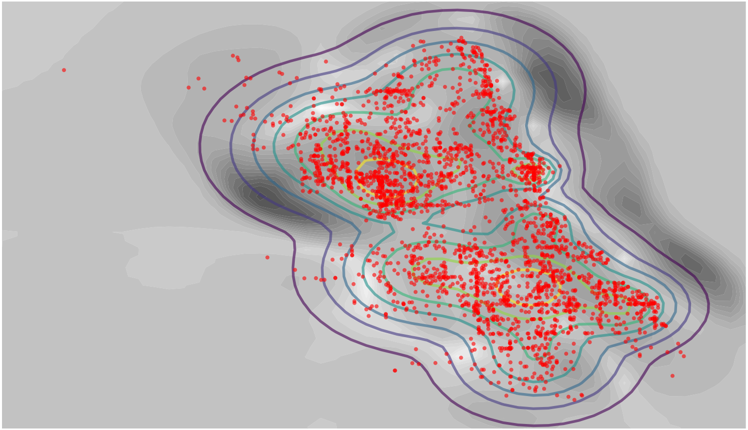

Chicago crime data

Model is Gaussian mixture with ten components

Stein witness function

Code: https://github.com/karlnapf/kernel_goodness_of_fit

Kernel stein discrepancy

Further applications:

- Evaluation of approximate MCMC methods.
  (Chwialkowski, Strathmann, G., ICML 2016; Gorham, Mackey, ICML 2017)

What kernel to use?

- The inverse multiquadric kernel,

$$k(x, y) = \left(c + \|x - y\|_2^2\right)^{\beta}$$

for $\beta \in (-1, 0)$.

arXiv.org > stat > arXiv:1703.01717

Statistics > Machine Learning

Measuring Sample Quality with Kernels

Jackson Gorham, Lester Mackey ICML 2017

(Submitted on 6 Mar 2017 (v1), last revised 3 Aug 2017 (this version, v6))

Testing statistical dependence

Dependence testing

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y independent?

X	Y
	A large animal who slings slobber, exudes a distinctive houndy odor, and wants nothing more than to follow his nose.
	Their noses guide them through life, and they're never happier than when following an interesting scent.
	A responsive, interactive pet, one that will blow in your ear and follow you everywhere.

Text from dogtime.com and petfinder.com

MMD as a dependence measure?

Could we use MMD?

$$MMD(\underbrace{P_{XY}}_{P}, \underbrace{P_X P_Y}_{Q}, \mathcal{H}_\kappa)$$

- We don't have samples from $Q := P_X P_Y$, only pairs $\{(x_i, y_i)\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$
  - Solution: simulate Q with pairs (x_i, y_j) for $j \neq i$.
- What kernel κ to use for the RKHS $\mathcal{H}_\kappa$?

MMD as a dependence measure?

Could we use MMD?

$$MMD(\underbrace{P_{XY}}_{P}, \underbrace{P_X P_Y}_{Q}, \mathcal{H}_\kappa)$$

- We don't have samples from $Q := P_X P_Y$, only pairs $\{(x_i, y_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$
  - Solution: simulate Q with pairs (x_i, y_j) for $j \neq i$.
- What kernel κ to use for the RKHS $\mathcal{H}_\kappa$?

MMD as a dependence measure?

Could we use MMD?

$$MMD(\underbrace{P_{XY}}_{P}, \underbrace{P_X P_Y}_{Q}, \mathcal{H}_\kappa)$$

- We don't have samples from $Q := P_X P_Y$, only pairs $\{(x_i, y_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$
  - Solution: simulate Q with pairs (x_i, y_j) for $j \neq i$.
- What kernel κ to use for the RKHS $\mathcal{H}_\kappa$?

MMD as a dependence measure

Kernel k on images with feature space $\mathcal{F}$,

$k($  ,  $)$

Kernel l on captions with feature space $\mathcal{G}$,

$l($ A large animal who slings slobber, ... , A responsive, interactive pet ... $)$

MMD as a dependence measure

Kernel k on images with feature space $\mathcal{F}$,

$k($  ,  $)$

Kernel l on captions with feature space $\mathcal{G}$,

$l($ A large animal who slings slobber, … , A responsive, interactive pet … $)$

Kernel κ on image-text pairs: are images and captions similar?

$\kappa($  A large animal who slings slobber, … ,  A responsive, interactive pet, … $)$

$= k($  ,  $) \times l($ A large animal who slings slobber, … , A responsive, interactive pet, … $)$

MMD as a dependence measure

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y independent?

$$MMD^2(\widehat{P}_{XY}, \widehat{P}_X \widehat{P}_Y, \mathcal{H}_\kappa) := \frac{1}{n^2}\text{trace}(K L)$$

(K, L column centered)

MMD as a dependence measure

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y independent?

$$MMD^2(\widehat{P}_{XY}, \widehat{P}_X \widehat{P}_Y, \mathcal{H}_\kappa) := \frac{1}{n^2}\text{trace}(KL)$$

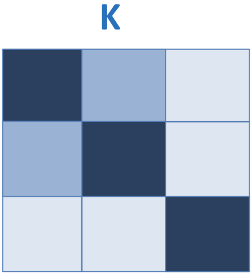

A large animal who slings slobber, exudes a distinctive houndy odor, ...

Their noses guide them through li
and they're never happier than wh
following an interesting scent.

A responsive, interactive pet, one that will blow in your ear and follow you everywhere.

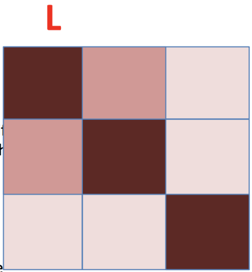

Text from dogtime.com and petfinder.com

MMD as a dependence measure

Two questions:

- Why the product kernel? Many ways to combine kernels - why not eg a sum?
- Is there a more interpretable way of defining this dependence measure?

Illustration: dependence $\neq$ correlation

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y dependent?

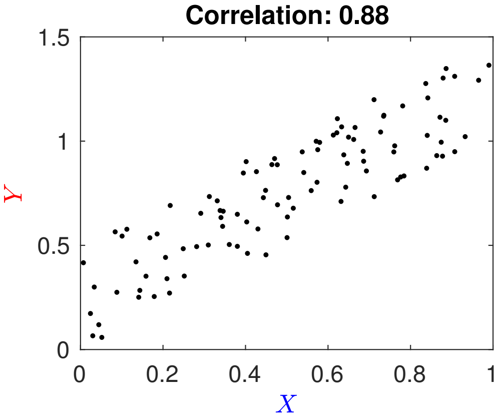

Illustration: dependence $\neq$correlation

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y dependent?

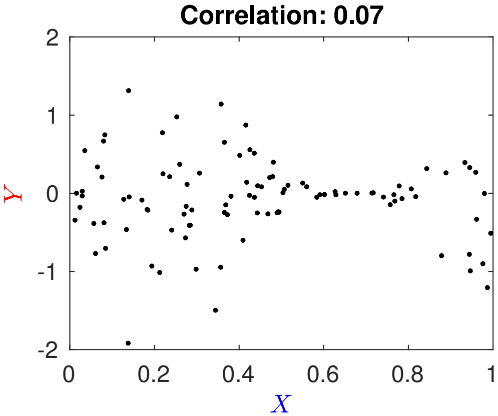

Illustration: dependence $\neq$correlation

- **Given:** Samples from a distribution P_{XY}
- **Goal:** Are X and Y dependent?

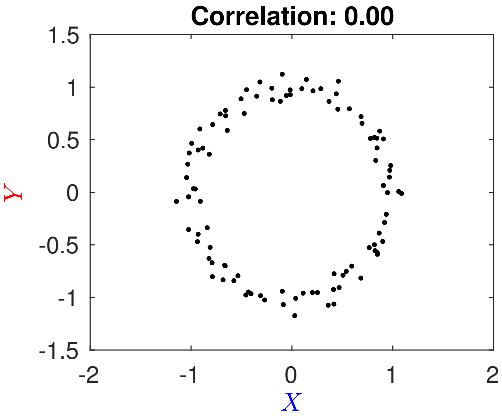

Finding covariance with smooth transformations

Illustration: two variables with no correlation but strong dependence.

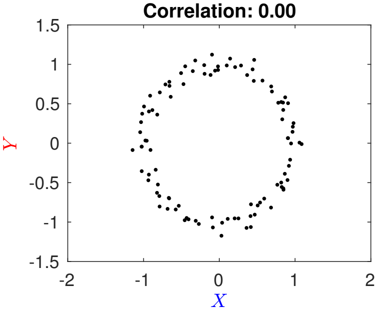

Finding covariance with smooth transformations

Illustration: two variables with no correlation but strong dependence.

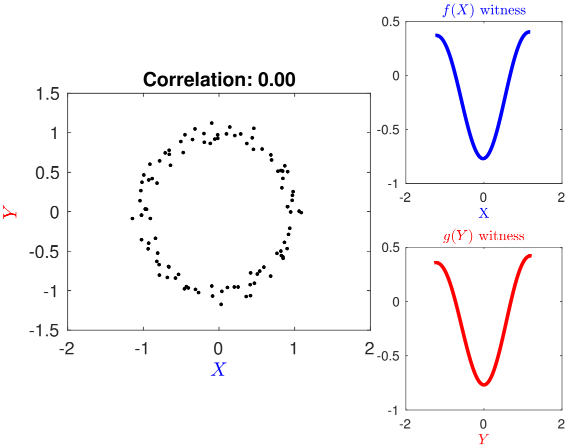

Finding covariance with smooth transformations

Illustration: two variables with no correlation but strong dependence.

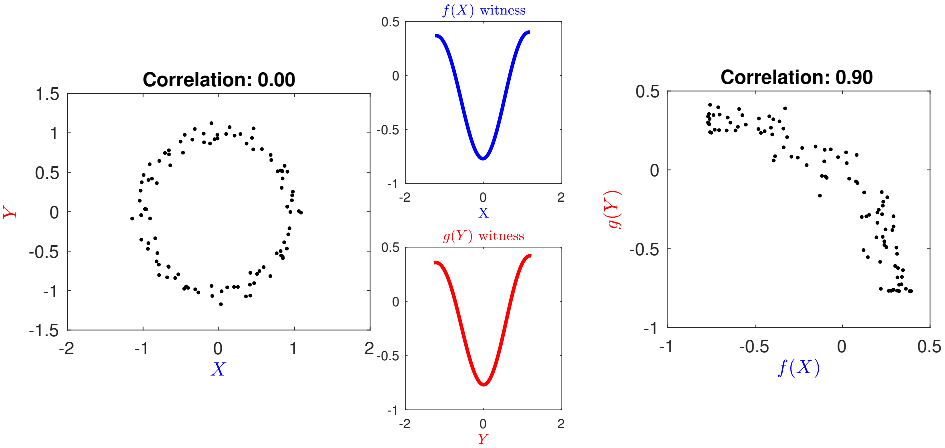

Define two spaces, one for each witness

Function in $\mathcal{F}$

$$f(x) = \sum_{j=1}^{\infty} f_j \varphi_j(x)$$

Feature map

$$\varphi(x) =$$

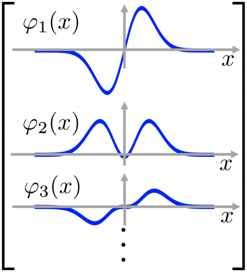

Kernel for RKHS $\mathcal{F}$ on $\mathcal{X}$:

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

Function in $\mathcal{G}$

$$g(y) = \sum_{j=1}^{\infty} g_j \phi_j(y)$$

Feature map

$$\phi(y) =$$

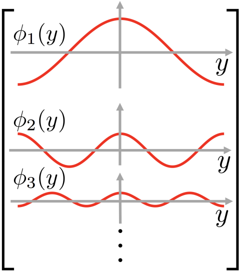

Kernel for RKHS $\mathcal{G}$ on $\mathcal{Y}$:

$$l(x, x') = \langle \phi(y), \phi(y') \rangle_{\mathcal{G}}$$

The constrained covariance

The constrained covariance is

$$\mathrm{COCO}(P_{XY}) = \sup_{\substack{\|f\|_{\mathcal{F}} \leq 1 \\ \|g\|_{\mathcal{G}} \leq 1}} \mathrm{cov}[f(x)g(y)]$$

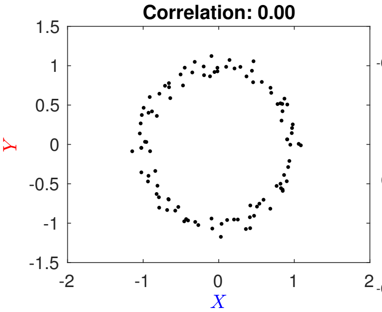

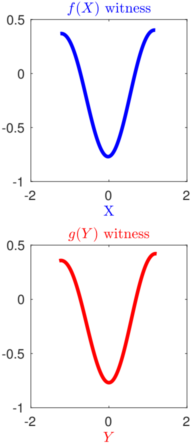

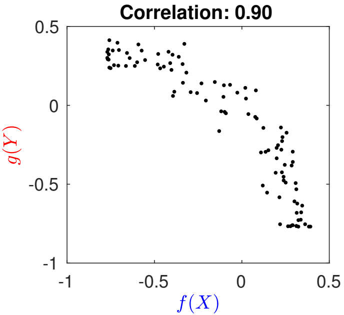

The constrained covariance

The constrained covariance is

$$\mathrm{COCO}(P_{XY}) = \sup_{\substack{\|f\|_{\mathcal{F}} \leq 1 \\ \|g\|_{\mathcal{G}} \leq 1}} \mathrm{cov}\left[\left(\sum_{j=1}^{\infty} f_j \varphi_j(x)\right)\left(\sum_{j=1}^{\infty} g_j \phi_j(y)\right)\right]$$

The constrained covariance

The constrained covariance is

$$\mathrm{COCO}(P_{XY}) = \sup_{\substack{\|f\|_{\mathcal{F}} \leq 1 \\ \|g\|_{\mathcal{G}} \leq 1}} E_{xy}\left[\left(\sum_{j=1}^{\infty} f_j \varphi_j(x)\right)\left(\sum_{j=1}^{\infty} g_j \phi_j(y)\right)\right]$$

Fine print: feature mappings $\varphi(x)$ and $\phi(y)$ assumed to have zero mean.

The constrained covariance

The constrained covariance is

$$\text{COCO}(P_{XY}) = \sup_{\substack{\|f\|_{\mathcal{F}} \leq 1 \\ \|g\|_{\mathcal{G}} \leq 1}} E_{xy}\left[\left(\sum_{j=1}^{\infty} f_j \varphi_j(x)\right)\left(\sum_{j=1}^{\infty} g_j \phi_j(y)\right)\right]$$

Fine print: feature mappings $\varphi(x)$ and $\phi(y)$ assumed to have zero mean.

Rewriting:

$$E_{xy}[f(x)g(y)]$$

$$= \begin{bmatrix} f_1 \\ f_2 \\ \vdots \end{bmatrix}^\top \underbrace{\mathbf{E}_{xy}\left(\begin{bmatrix} \varphi_1(x) \\ \varphi_2(x) \\ \vdots \end{bmatrix}\begin{bmatrix} \phi_1(y) & \phi_2(y) & \dots \end{bmatrix}\right)}_{C_{\varphi(x)\phi(y)}} \begin{bmatrix} g_1 \\ g_2 \\ \vdots \end{bmatrix}$$

The constrained covariance

The constrained covariance is

$$\mathrm{COCO}(P_{XY}) = \sup_{\substack{\|f\|_{\mathcal{F}} \leq 1 \\ \|g\|_{\mathcal{G}} \leq 1}} E_{xy}\left[\left(\sum_{j=1}^{\infty} f_j \varphi_j(x)\right)\left(\sum_{j=1}^{\infty} g_j \phi_j(y)\right)\right]$$

Fine print: feature mappings $\varphi(x)$ and $\phi(y)$ assumed to have zero mean.

Rewriting:

$$E_{xy}[f(x)g(y)]$$

$$= \begin{bmatrix} f_1 \\ f_2 \\ \vdots \end{bmatrix}^\top \underbrace{\mathbf{E}_{xy}\left(\begin{bmatrix} \varphi_1(x) \\ \varphi_2(x) \\ \vdots \end{bmatrix}\begin{bmatrix} \phi_1(y) & \phi_2(y) & \dots \end{bmatrix}\right)}_{C_{\varphi(x)\phi(y)}} \begin{bmatrix} g_1 \\ g_2 \\ \vdots \end{bmatrix}$$

COCO: max singular value of feature covariance $C_{\varphi(x)\phi(y)}$

Computing COCO in practice

Given sample $\{(x_i, y_i)\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{COCO}$?

Computing COCO in practice

Given sample $\{(x_i, y_i)\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{COCO}$?

$\widehat{COCO}$ is largest eigenvalue $\gamma_{\max}$ of

$$\begin{bmatrix} 0 & \frac{1}{n}KL \\ \frac{1}{n}LK & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \gamma \begin{bmatrix} K & 0 \\ 0 & L \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i, y_j)$.

Fine print: kernels are computed with empirically centered features $\varphi(x) - \frac{1}{n}\sum_{i=1}^n \varphi(x_i)$ and $\phi(y) - \frac{1}{n}\sum_{i=1}^n \phi(y_i)$.

G., Smola., Bousquet, Herbrich, Belitski, Augath, Murayama, Pauls, Schoelkopf, and Logothetis, AISTATS'05

Computing COCO in practice

Given sample $\{(x_i, y_i)\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{COCO}$?

$\widehat{COCO}$ is largest eigenvalue $\gamma_{\max}$ of

$$\begin{bmatrix} 0 & \frac{1}{n}KL \\ \frac{1}{n}LK & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \gamma \begin{bmatrix} K & 0 \\ 0 & L \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix}.$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i, y_j)$.

Witness functions (singular vectors):

$$f(x) \propto \sum_{i=1}^n \alpha_i k(x_i, x) \qquad g(y) \propto \sum_{i=1}^n \beta_i l(y_i, y)$$

Fine print: kernels are computed with empirically centered features $\varphi(x) - \frac{1}{n}\sum_{i=1}^n \varphi(x_i)$ and $\phi(y) - \frac{1}{n}\sum_{i=1}^n \phi(y_i)$.

G., Smola., Bousquet, Herbrich, Belitski, Augath, Murayama, Pauls, Schoelkopf, and Logothetis, AISTATS'05

Empirical COCO: proof (1)

The Lagrangian is

$$\mathcal{L}(f, g, \lambda, \gamma) = \underbrace{\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)]}_{\text{covariance}} - \underbrace{\frac{\lambda}{2}\left(\|f\|_{\mathcal{F}}^2 - 1\right) - \frac{\gamma}{2}\left(\|g\|_{\mathcal{G}}^2 - 1\right)}_{\text{smoothness constraints}}.$$

Fine print: $f(x_i)g(y_i)$ centered to have zero empirical mean.

Empirical COCO: proof (1)

The Lagrangian is

$$\mathcal{L}(f, g, \lambda, \gamma) = \underbrace{\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)]}_{\text{covariance}} - \underbrace{\frac{\lambda}{2}\left(\|f\|_{\mathcal{F}}^2 - 1\right) - \frac{\gamma}{2}\left(\|g\|_{\mathcal{G}}^2 - 1\right)}_{\text{smoothness constraints}}.$$

Fine print: $f(x_i)g(y_i)$ centered to have zero empirical mean.

Assume (cf representer theorem):

$$f = \sum_{i=1}^{n} \alpha_i \varphi(x_i) \qquad g = \sum_{i=1}^{n} \beta_i \psi(y_i)$$

for centered $\varphi(x_i)$, $\phi(y_i)$.

Empirical COCO: proof (1)

The Lagrangian is

$$\mathcal{L}(f, g, \lambda, \gamma) = \underbrace{\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)]}_{\text{covariance}} - \underbrace{\frac{\lambda}{2}\left(\|f\|_{\mathcal{F}}^2 - 1\right) - \frac{\gamma}{2}\left(\|g\|_{\mathcal{G}}^2 - 1\right)}_{\text{smoothness constraints}}.$$

Fine print: $f(x_i)g(y_i)$ centered to have zero empirical mean.

Assume (cf representer theorem):

$$f = \sum_{i=1}^{n} \alpha_i \varphi(x_i) \qquad g = \sum_{i=1}^{n} \beta_i \psi(y_i)$$

for centered $\varphi(x_i)$, $\phi(y_i)$.

First step is smoothness constraint:

$$\|f\|_{\mathcal{F}}^2 - 1 = \langle f, f\rangle_{\mathcal{F}} - 1$$

Empirical COCO: proof (1)

The Lagrangian is

$$\mathcal{L}(f,g,\lambda,\gamma) = \underbrace{\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)]}_{\text{covariance}} - \underbrace{\frac{\lambda}{2}\left(\|f\|_{\mathcal{F}}^2 - 1\right) - \frac{\gamma}{2}\left(\|g\|_{\mathcal{G}}^2 - 1\right)}_{\text{smoothness constraints}}.$$

Fine print: $f(x_i)g(y_i)$ centered to have zero empirical mean.

Assume (cf representer theorem):

$$f = \sum_{i=1}^{n}\alpha_i\varphi(x_i) \qquad g = \sum_{i=1}^{n}\beta_i\psi(y_i)$$

for centered $\varphi(x_i)$, $\phi(y_i)$.

First step is smoothness constraint:

$$\begin{aligned}\|f\|_{\mathcal{F}}^2 - 1 &= \langle f, f\rangle_{\mathcal{F}} - 1 \\ &= \left\langle \sum_{i=1}^{n}\alpha_i\varphi(x_i), \sum_{i=1}^{n}\alpha_i\varphi(x_i)\right\rangle_{\mathcal{F}} - 1\end{aligned}$$

Empirical COCO: proof (1)

The Lagrangian is

$$\mathcal{L}(f,g,\lambda,\gamma) = \underbrace{\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)]}_{\text{covariance}} - \underbrace{\frac{\lambda}{2}\left(\|f\|^2_{\mathcal{F}} - 1\right) - \frac{\gamma}{2}\left(\|g\|^2_{\mathcal{G}} - 1\right)}_{\text{smoothness constraints}}.$$

Fine print: $f(x_i)g(y_i)$ centered to have zero empirical mean.

Assume (cf representer theorem):

$$f = \sum_{i=1}^{n} \alpha_i \varphi(x_i) \qquad g = \sum_{i=1}^{n} \beta_i \psi(y_i)$$

for centered $\varphi(x_i)$, $\phi(y_i)$.

First step is smoothness constraint:

$$\begin{aligned}\|f\|^2_{\mathcal{F}} - 1 &= \langle f, f\rangle_{\mathcal{F}} - 1 \\ &= \left\langle \sum_{i=1}^{n} \alpha_i \varphi(x_i), \sum_{i=1}^{n} \alpha_i \varphi(x_i) \right\rangle_{\mathcal{F}} - 1 \\ &= \alpha^\top K \alpha - 1\end{aligned}$$

Proof sketch (2)

Second step is covariance:

$$\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)] = \frac{1}{n}\sum_{i=1}^{n}\langle f,\varphi(x_i)\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}}$$

$$= \frac{1}{n}\sum_{i=1}^{n}\left\langle \sum_{\ell=1}^{n}\alpha_\ell\varphi(x_\ell),\varphi(x_i)\right\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}}$$

$$= \frac{1}{n}\alpha^\top KL\beta$$

Proof sketch (2)

Second step is covariance:

$$
\begin{aligned}
\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)] &= \frac{1}{n}\sum_{i=1}^{n}\langle f,\varphi(x_i)\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}} \\
&= \frac{1}{n}\sum_{i=1}^{n}\left\langle \sum_{\ell=1}^{n}\alpha_\ell\varphi(x_\ell),\varphi(x_i)\right\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}} \\
&= \frac{1}{n}\alpha^\top KL\beta
\end{aligned}
$$

Proof sketch (2)

Second step is covariance:

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)] &= \frac{1}{n}\sum_{i=1}^{n}\langle f,\varphi(x_i)\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}} \\ &= \frac{1}{n}\sum_{i=1}^{n}\left\langle \sum_{\ell=1}^{n}\alpha_\ell\varphi(x_\ell),\varphi(x_i)\right\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}} \\ &= \frac{1}{n}\alpha^\top KL\beta\end{aligned}$$

where $K_{ij} = k(x_i, x_j) = \langle\varphi(x_i),\varphi(x_j)\rangle_{\mathcal{F}}$ $\qquad L_{ij} = l(y_i, y_j)$.

Proof sketch (2)

Second step is covariance:

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n}[f(x_i)g(y_i)] &= \frac{1}{n}\sum_{i=1}^{n}\langle f,\varphi(x_i)\rangle_{\mathcal{F}}\,\langle g,\varphi(y_i)\rangle_{\mathcal{G}}\\ &= \frac{1}{n}\sum_{i=1}^{n}\left\langle\sum_{\ell=1}^{n}\alpha_\ell\varphi(x_\ell),\varphi(x_i)\right\rangle_{\mathcal{F}}\langle g,\varphi(y_i)\rangle_{\mathcal{G}}\\ &= \frac{1}{n}\alpha^\top KL\beta\end{aligned}$$

where $K_{ij} = k(x_i, x_j) = \langle\varphi(x_i),\varphi(x_j)\rangle_{\mathcal{F}}$ $\qquad L_{ij} = l(y_i, y_j)$.

The Lagranian is now:

$$\mathcal{L}(f, g, \lambda, \gamma) = \frac{1}{n}\alpha^\top KL\beta - \frac{\lambda}{2}\left(\alpha^\top K\alpha - 1\right) - \frac{\gamma}{2}\left(\beta^\top L\beta - 1\right)$$

What is a large dependence with COCO?

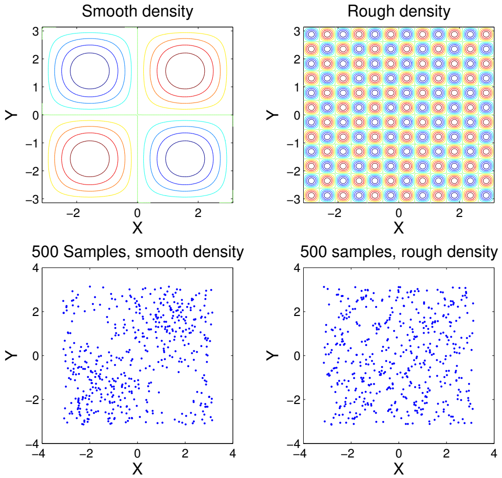

Which of these is the more "dependent"?

Density takes the form:

$$P_{XY} \propto 1 + \sin(\omega x)\sin(\omega y)$$

Finding covariance with smooth transformations

Case of $\omega = 1$:

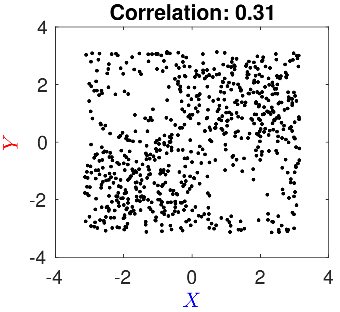

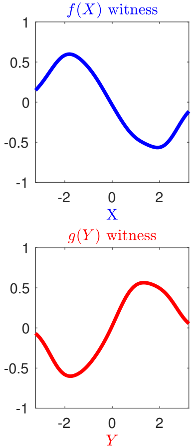

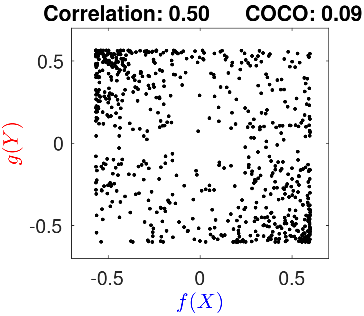

Finding covariance with smooth transformations

Case of $\omega = 2$:

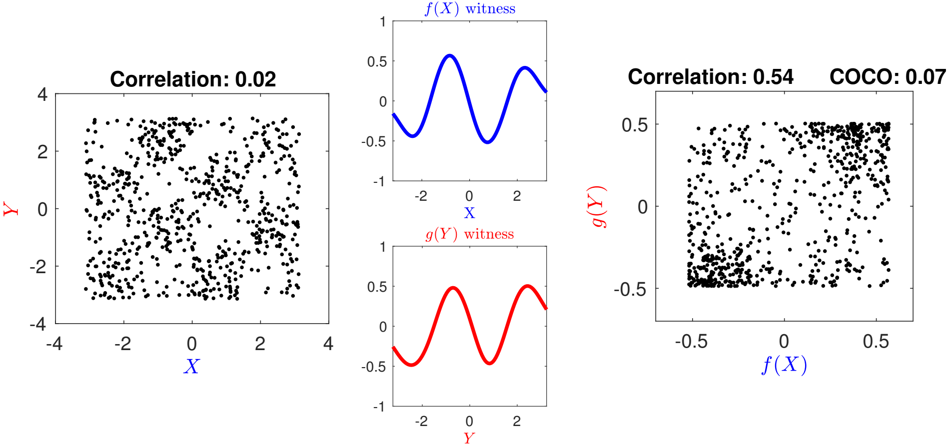

Finding covariance with smooth transformations

Case of $\omega = 3$:

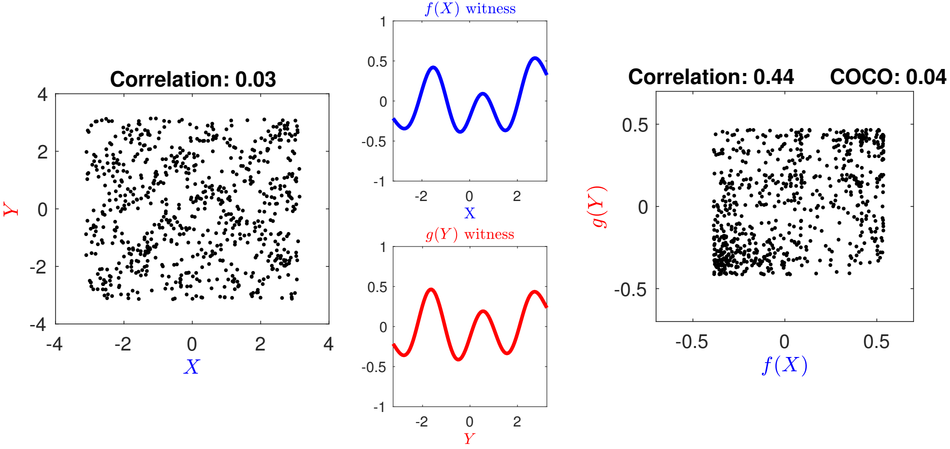

Finding covariance with smooth transformations

Case of $\omega = 4$:

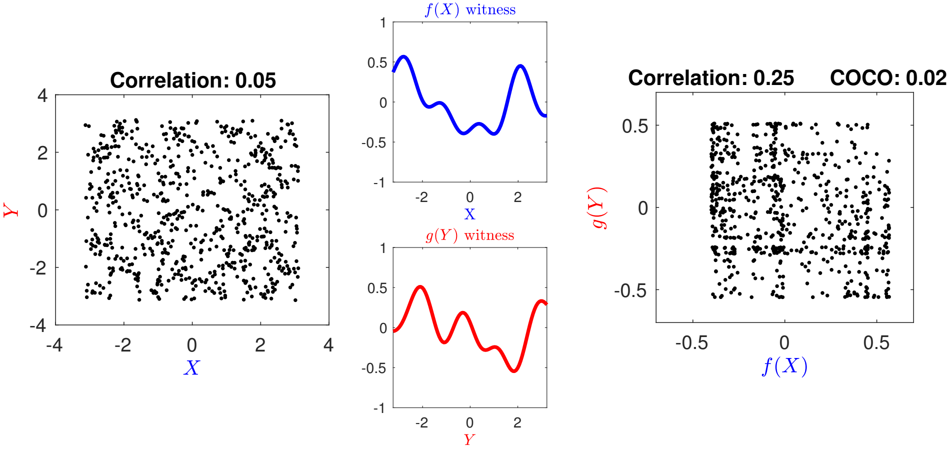

Finding covariance with smooth transformations

Case of $\omega = ??$:

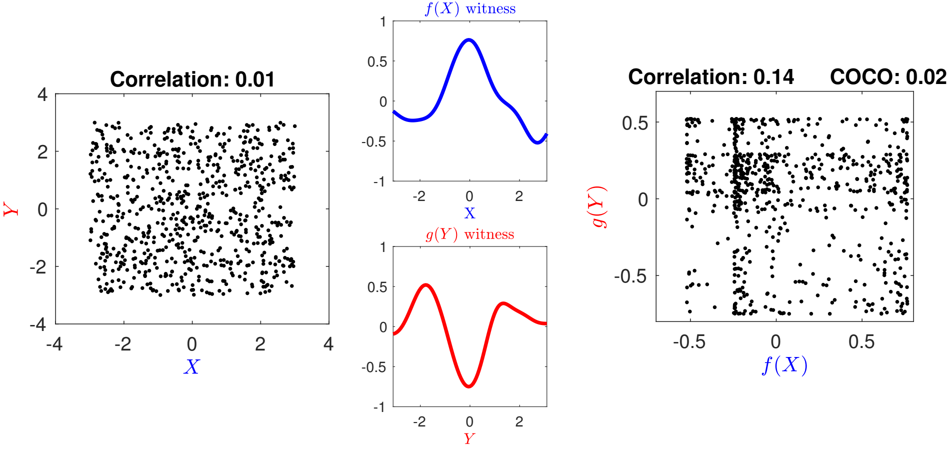

Finding covariance with smooth transformations

Case of $\omega = 0$: uniform noise! (shows bias)

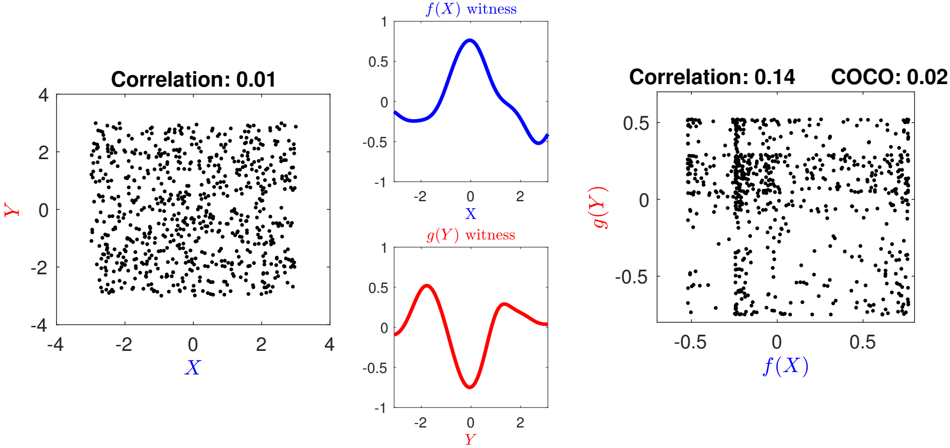

Dependence largest when at “low” frequencies

- As dependence is encoded at higher frequencies, the smooth mappings f, g achieve lower linear dependence.
- Even for independent variables, COCO will not be zero at finite sample sizes, since some mild linear dependence will be found by f,g (bias)
- This bias will decrease with increasing sample size.

Can we do better than COCO?

A second example with zero correlation.

First singular value of feature covariance $C_{\varphi(x)\phi(y)}$:

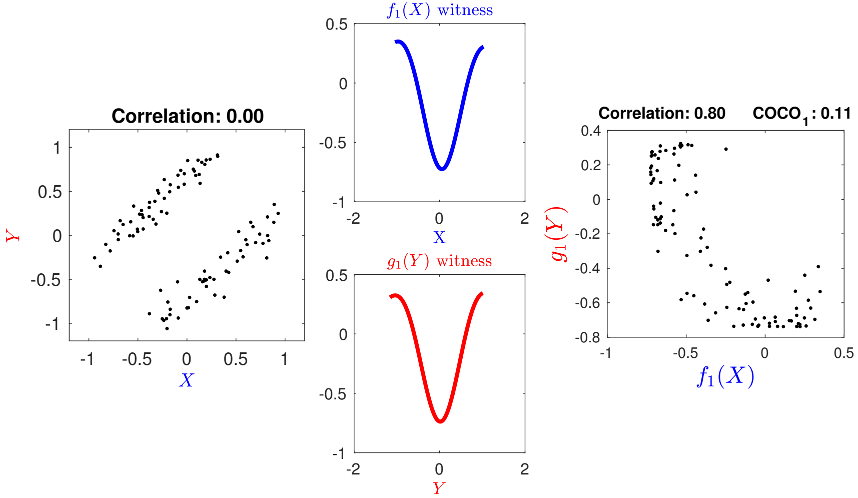

Can we do better than COCO?

A second example with zero correlation.

Second singular value of feature covariance $C_{\varphi(x)\phi(y)}$:

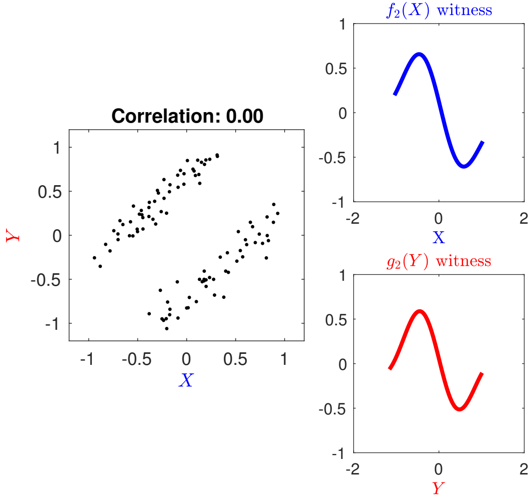

Can we do better than COCO?

A second example with zero correlation.

Second singular value of feature covariance $C_{\varphi(x)\phi(y)}$:

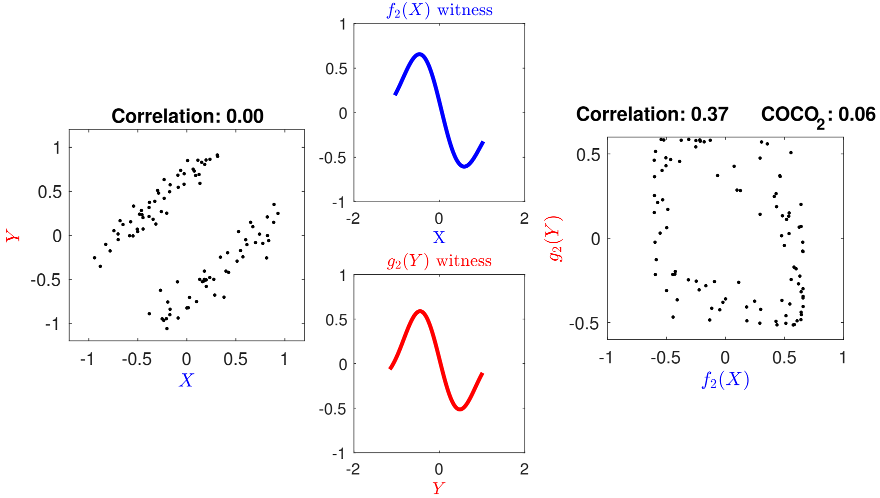

The Hilbert-Schmidt Independence Criterion

Writing the ith singular value of the feature covariance $C_{\varphi(x)\phi(y)}$ as

$$\gamma_i := COCO_i(P_{XY}; \mathcal{F}, \mathcal{G}),$$

define Hilbert-Schmidt Independence Criterion (HSIC)

$$HSIC^2(P_{XY}; \mathcal{F}, \mathcal{G}) = \sum_{i=1}^{\infty} \gamma_i^2.$$

G, Bousquet , Smola., and Schoelkopf, ALT05; G,., Fukumizu, Teo., Song., Schoelkopf., and Smola, NIPS 2007,.

The Hilbert-Schmidt Independence Criterion

Writing the ith singular value of the feature covariance $C_{\varphi(x)\phi(y)}$ as

$$\gamma_i := COCO_i(P_{XY}; \mathcal{F}, \mathcal{G}),$$

define Hilbert-Schmidt Independence Criterion (HSIC)

$$HSIC^2(P_{XY}; \mathcal{F}, \mathcal{G}) = \sum_{i=1}^{\infty} \gamma_i^2.$$

G, Bousquet , Smola., and Schoelkopf, ALT05; G,., Fukumizu, Teo., Song., Schoelkopf., and Smola, NIPS 2007,.

HSIC is MMD with product kernel!

$$HSIC^2(P_{XY}; \mathcal{F}, \mathcal{G}) = MMD^2(P_{XY}, P_X P_Y; \mathcal{H}_\kappa)$$

where $\kappa((x, y), (x', y')) = k(x, x')l(y, y')$.

Asymptotics of HSIC under independence

- Given sample $\{(x_i, y_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{HSIC}$?
- Empirical HSIC (biased)

$$\widehat{HSIC} = \frac{1}{n^2}\text{trace}(KL)$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i y_j)$ (K and L computed with empirically centered features)

- Statistical testing: given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α?
- Asymptotics of $\widehat{HSIC}$ when $P_{XY} = P_X P_Y$:

$$n\widehat{HSIC} \xrightarrow{D} \sum_{l=1}^{\infty} \lambda_l z_l^2, \qquad z_l \sim \mathcal{N}(0,1)\text{i.i.d.}$$

where $\lambda_l \psi_l(z_j) = \int h_{ijqr} \psi_l(z_i) dF_{i,q,r}, \quad h_{ijqr} = \frac{1}{4!}\sum_{(t,u,v,w)}^{(i,j,q,r)} k_{tu} l_{tu} + k_{tu} l_{vw} - 2k_{tu} l_{tv}$

Asymptotics of HSIC under independence

- Given sample $\{(x_i, y_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{HSIC}$?
- Empirical HSIC (biased)

$$\widehat{HSIC} = \frac{1}{n^2}\text{trace}(KL)$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i y_j)$ (K and L computed with empirically centered features)

- Statistical testing: given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α?
- Asymptotics of $\widehat{HSIC}$ when $P_{XY} = P_X P_Y$:

$$n\widehat{HSIC} \xrightarrow{D} \sum_{l=1}^{\infty} \lambda_l z_l^2, \qquad z_l \sim \mathcal{N}(0,1)\text{i.i.d.}$$

where $\lambda_l \psi_l(z_j) = \int h_{ijqr} \psi_l(z_i) dF_{i,q,r}$, $\quad h_{ijqr} = \frac{1}{4!}\sum_{(t,u,v,w)}^{(i,j,q,r)} k_{tu} l_{tu} + k_{tu} l_{vw} - 2k_{tu} l_{tv}$

Asymptotics of HSIC under independence

- Given sample $\{(x_i, y_i\}_{i=1}^n \overset{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{HSIC}$?
- Empirical HSIC (biased)

$$\widehat{HSIC} = \frac{1}{n^2}\text{trace}(KL)$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i y_j)$ (K and L computed with empirically centered features)

- Statistical testing: given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α?
- Asymptotics of $\widehat{HSIC}$ when $P_{XY} = P_X P_Y$:

$$n\widehat{HSIC} \overset{D}{\to} \sum_{l=1}^{\infty} \lambda_l z_l^2, \qquad z_l \sim \mathcal{N}(0,1)\text{i.i.d.}$$

where $\lambda_l \psi_l(z_j) = \int h_{ijqr} \psi_l(z_i) dF_{i,q,r}, \quad h_{ijqr} = \frac{1}{4!} \sum_{(t,u,v,w)}^{(i,j,q,r)} k_{tu} l_{tu} + k_{tu} l_{vw} - 2k_{tu} l_{tv}$

Asymptotics of HSIC under independence

- Given sample $\{(x_i, y_i\}_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} P_{XY}$, what is empirical $\widehat{HSIC}$?
- Empirical HSIC (biased)

$$\widehat{HSIC} = \frac{1}{n^2}\text{trace}(KL)$$

$K_{ij} = k(x_i, x_j)$ and $L_{ij} = l(y_i y_j)$ (K and L computed with empirically centered features)

- Statistical testing: given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α?
- Asymptotics of $\widehat{HSIC}$ when $P_{XY} = P_X P_Y$:

$$n\widehat{HSIC} \xrightarrow{D} \sum_{l=1}^{\infty} \lambda_l z_l^2, \qquad z_l \sim \mathcal{N}(0,1)\text{i.i.d.}$$

where $\lambda_l \psi_l(z_j) = \int h_{ijqr} \psi_l(z_i) dF_{i,q,r}, \quad h_{ijqr} = \frac{1}{4!}\sum_{(t,u,v,w)}^{(i,j,q,r)} k_{tu} l_{tu} + k_{tu} l_{vw} - 2k_{tu} l_{tv}$

A statistical test

- Given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α (prob. of false positive)?

- **Original time series:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_1 \; Y_2 \; Y_3 \; Y_4 \; Y_5 \; Y_6 \; Y_7 \; Y_8 \; Y_9 \; Y_{10}$$

- **Permutation:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_7 \; Y_3 \; Y_9 \; Y_2 \; Y_4 \; Y_8 \; Y_5 \; Y_1 \; Y_6 \; Y_{10}$$

- Null distribution via permutation
  - Compute HSIC for $\{x_i, y_{\pi(i)}\}_{i=1}^n$ for random permutation π of indices $\{1, \ldots, n\}$. This gives HSIC for independent variables.
  - Repeat for many different permutations, get empirical CDF
  - Threshold c_α is $1 - \alpha$ quantile of empirical CDF

A statistical test

- Given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α (prob. of false positive)?

- **Original time series:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_1 \; Y_2 \; Y_3 \; Y_4 \; Y_5 \; Y_6 \; Y_7 \; Y_8 \; Y_9 \; Y_{10}$$

- **Permutation:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_7 \; Y_3 \; Y_9 \; Y_2 \; Y_4 \; Y_8 \; Y_5 \; Y_1 \; Y_6 \; Y_{10}$$

- Null distribution via permutation
  - Compute HSIC for $\{x_i, y_{\pi(i)}\}_{i=1}^n$ for random permutation π of indices $\{1, \ldots, n\}$. This gives HSIC for independent variables.
  - Repeat for many different permutations, get empirical CDF
  - Threshold c_α is $1 - \alpha$ quantile of empirical CDF

A statistical test

- Given $P_{XY} = P_X P_Y$, what is the threshold c_α such that $P(\widehat{HSIC} > c_\alpha) < \alpha$ for small α (prob. of false positive)?
- **Original time series:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_1 \; Y_2 \; Y_3 \; Y_4 \; Y_5 \; Y_6 \; Y_7 \; Y_8 \; Y_9 \; Y_{10}$$

- **Permutation:**

$$X_1 \; X_2 \; X_3 \; X_4 \; X_5 \; X_6 \; X_7 \; X_8 \; X_9 \; X_{10}$$
$$Y_7 \; Y_3 \; Y_9 \; Y_2 \; Y_4 \; Y_8 \; Y_5 \; Y_1 \; Y_6 \; Y_{10}$$

- Null distribution via permutation
  - Compute HSIC for $\{x_i, y_{\pi(i)}\}_{i=1}^n$ for random permutation π of indices $\{1, \ldots, n\}$. This gives HSIC for independent variables.
  - Repeat for many different permutations, get empirical CDF
  - Threshold c_α is $1 - \alpha$ quantile of empirical CDF

Application: dependence detection across languages

Testing task: detect dependence between English and French text

X	Y
Honourable senators, I have a question for the Leader of the Government in the Senate	Honorables sénateurs, ma question s'adresse au leader du gouvernement au Sénat
No doubt there is great pressure on provincial and municipal governments	Les ordres de gouvernements provinciaux et municipaux subissent de fortes pressions
In fact, we have increased federal investments for early childhood development.	Au contraire, nous avons augmenté le financement fédéral pour le développement des jeunes
⋮	⋮

Text from the aligned hansards of the 36th parliament of canada, https://www.isi.edu/natural-language/download/hansard/

Application: dependence detection across languages

Testing task: detect dependence between English and French text

k-spectrum kernel, $k = 10$, sample size $n = 10$

X

Honourable senators, I have a question for the Leader of the Government in the Senate

No doubt there is great pressure on provincial and municipal governments

In fact, we have increased federal investments for early childhood development.

⋮

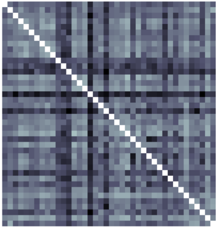

K

Y

Honorables sénateurs, ma question s'adresse au leader du gouvernement au Sénat

Les ordres de gouvernements provinciaux et municipaux subissent de fortes pressions

Au contraire, nous avons augmenté
le financement fédéral pour le développement des jeunes

⋮

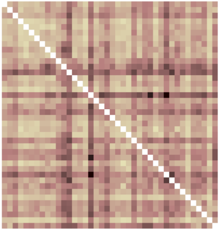

L

$$\widehat{HSIC} = \frac{1}{n^2}\operatorname{trace}(K\,L)$$

(K and L column centered)

Application:Dependence detection across languages

Results (for $\alpha = 0.05$)

- k-spectrum kernel: average Type II error 0
- Bag of words kernel: average Type II error 0.18

Settings: Five line extracts, averaged over 300 repetitions, for "Agriculture" transcripts. Similar results for Fisheries and Immigration transcripts.

Testing higher order interactions

Detecting higher order interaction

How to detect V-structures with pairwise weak individual dependence?

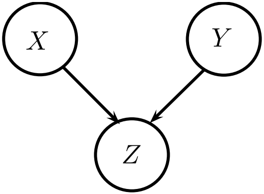

Detecting higher order interaction

How to detect V-structures with pairwise weak individual dependence?

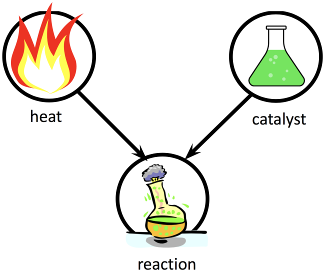

Detecting higher order interaction

How to detect V-structures with pairwise weak individual dependence?

$X \perp\!\!\!\perp Y, Y \perp\!\!\!\perp Z, X \perp\!\!\!\perp Z$

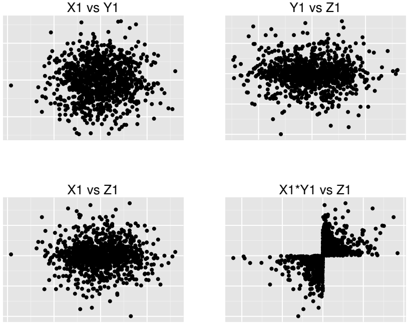

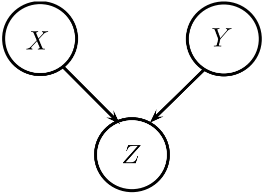

- $X, Y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$
- $Z \mid X, Y \sim \text{sign}(XY) Exp(\frac{1}{\sqrt{2}})$

Fine print: Faithfulness violated here!

V-structure discovery

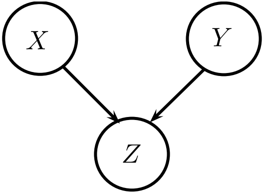

Assume $X \perp\!\!\!\perp Y$ has been established.

V-structure can then be detected by:

- Consistent CI test: $\mathbf{H_0}: X \perp\!\!\!\perp Y | Z$ [Fukumizu et al. 2008, Zhang et al. 2011]
- Factorisation test: $\mathbf{H_0}: (X, Y) \perp\!\!\!\perp Z \vee (X, Z) \perp\!\!\!\perp Y \vee (Y, Z) \perp\!\!\!\perp X$ (multiple standard two-variable tests)

How well do these work?

Detecting higher order interaction

Generalise earlier example to p dimensions

$X \perp\!\!\!\perp Y, Y \perp\!\!\!\perp Z, X \perp\!\!\!\perp Z$

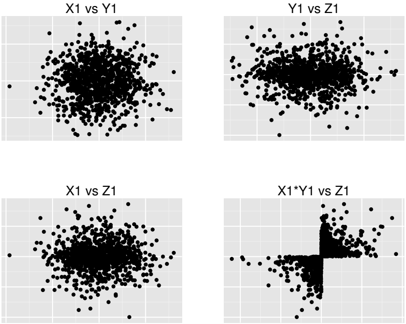

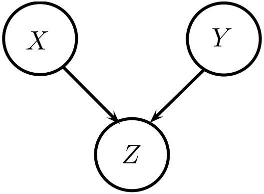

- $X, Y \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$
- $Z | \, X, Y \sim \text{sign}(XY) Exp(\frac{1}{\sqrt{2}})$
- $X_{2:p}, Y_{2:p}, Z_{2:p} \overset{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{I}_{p-1})$

Fine print: Faithfulness violated here!

V-structure discovery

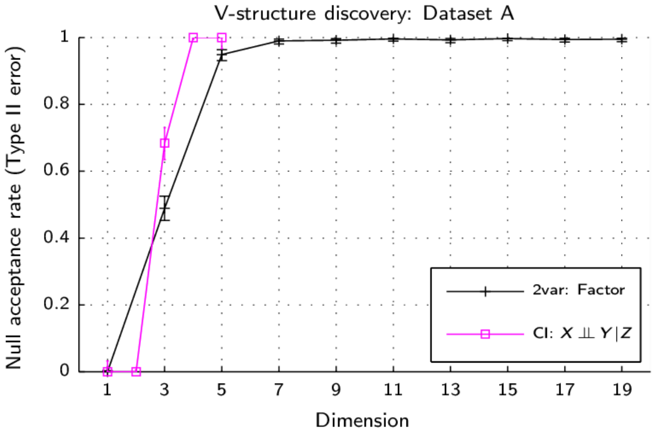

CI test for $X \perp\!\!\!\perp Y | Z$ from Zhang et al. (2011), and a factorisation test, $n = 500$

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \ldots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \dots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$$

$$D = 3: \qquad \Delta_L P = P_{XYZ} - P_X P_{YZ} - P_Y P_{XZ} - P_Z P_{XY} + 2 P_X P_Y P_Z$$

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \ldots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$$

$$D = 3: \qquad \Delta_L P = P_{XYZ} - P_X P_{YZ} - P_Y P_{XZ} - P_Z P_{XY} + 2 P_X P_Y P_Z$$

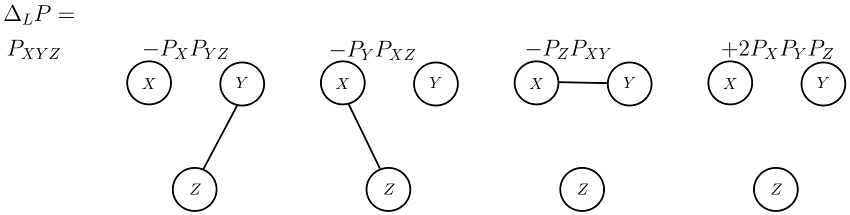

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \ldots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$$

$$D = 3: \qquad \Delta_L P = P_{XYZ} - P_X P_{YZ} - P_Y P_{XZ} - P_Z P_{XY} + 2 P_X P_Y P_Z$$

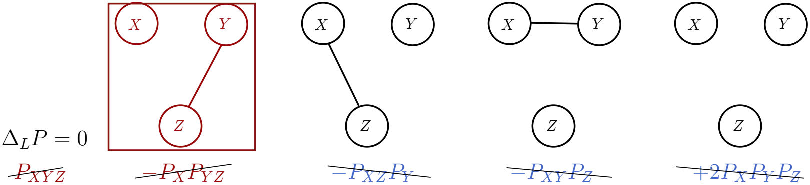

Case of $P_X \perp\!\!\!\perp P_{YZ}$

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \dots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$$

$$D = 3: \qquad \Delta_L P = P_{XYZ} - P_X P_{YZ} - P_Y P_{XZ} - P_Z P_{XY} + 2 P_X P_Y P_Z$$

$$(X, Y) \perp\!\!\!\perp Z \vee (X, Z) \perp\!\!\!\perp Y \vee (Y, Z) \perp\!\!\!\perp X \;\Rightarrow\; \Delta_L P = 0.$$

...so what might be missed?

Lancaster interaction measure

Lancaster interaction measure of $(X_1, \ldots, X_D) \sim P$ is a signed measure ΔP that vanishes whenever P can be factorised non-trivially.

$$D = 2: \qquad \Delta_L P = P_{XY} - P_X P_Y$$

$$D = 3: \qquad \Delta_L P = P_{XYZ} - P_X P_{YZ} - P_Y P_{XZ} - P_Z P_{XY} + 2 P_X P_Y P_Z$$

$$\Delta_L P = 0 \not\Rightarrow (X, Y) \perp\!\!\!\perp Z \vee (X, Z) \perp\!\!\!\perp Y \vee (Y, Z) \perp\!\!\!\perp X$$

Example:

$P(0,0,0) = 0.2$	$P(0,0,1) = 0.1$	$P(1,0,0) = 0.1$	$P(1,0,1) = 0.1$
$P(0,1,0) = 0.1$	$P(0,1,1) = 0.1$	$P(1,1,0) = 0.1$	$P(1,1,1) = 0.2$

A kernel test statistic using Lancaster Measure

Construct a test by estimating $\|\mu_\kappa(\Delta_L P)\|^2_{\mathcal{H}_\kappa}$, where $\kappa = k \otimes l \otimes m$:

$$\|\mu_\kappa(P_{XYZ} - P_{XY}P_Z - \cdots)\|^2_{\mathcal{H}_\kappa} =$$
$$\langle \mu_\kappa P_{XYZ}, \mu_\kappa P_{XYZ} \rangle_{\mathcal{H}_\kappa} - 2 \langle \mu_\kappa P_{XYZ}, \mu_\kappa P_{XY}P_Z \rangle_{\mathcal{H}_\kappa} \cdots$$

A kernel test statistic using Lancaster Measure

$\nu \backslash \nu'$	P_{XYZ}	$P_{XY}P_Z$	$P_{XZ}P_Y$	$P_{YZ}P_X$	$P_X P_Y P_Z$
P_{XYZ}	$(K \circ L \circ M)_{++}$	$((K \circ L)\,M)_{++}$	$((K \circ M)\,L)_{++}$	$((M \circ L)\,K)_{++}$	$tr(K_+ \circ L_+ \circ M_+)$
$P_{XY}P_Z$		$(K \circ L)_{++}\,M_{++}$	$(MKL)_{++}$	$(KLM)_{++}$	$(KL)_{++}M_{++}$
$P_{XZ}P_Y$			$(K \circ M)_{++}\,L_{++}$	$(KML)_{++}$	$(KM)_{++}L_{++}$
$P_{YZ}P_X$				$(L \circ M)_{++}\,K_{++}$	$(LM)_{++}K_{++}$
$P_X P_Y P_Z$					$K_{++}L_{++}M_{++}$

Table: V-statistic estimators of $\langle \mu_\kappa \nu, \mu_\kappa \nu' \rangle_{\mathcal{H}_\kappa}$ (without terms $P_X P_Y P_Z$). H is centering matrix $I - n^{-1}$

Lancaster interaction statistic: Sejdinovic, G, Bergsma, NIPS13

$$\|\mu_\kappa(\Delta_L P)\|^2_{\mathcal{H}_\kappa} = \frac{1}{n^2}\boxed{(HKH \circ HLH \circ HMH)_{++}}.$$

Empirical joint central moment in the feature space

A kernel test statistic using Lancaster Measure

$\nu \backslash \nu'$	P_{XYZ}	$P_{XY}P_Z$	$P_{XZ}P_Y$	$P_{YZ}P_X$	$P_X P_Y P_Z$
P_{XYZ}	$(K \circ L \circ M)_{++}$	$((K \circ L) M)_{++}$	$((K \circ M) L)_{++}$	$((M \circ L) K)_{++}$	$tr(K_+ \circ L_+ \circ M_+)$
$P_{XY}P_Z$		$(K \circ L)_{++} M_{++}$	$(MKL)_{++}$	$(KLM)_{++}$	$(KL)_{++}M_{++}$
$P_{XZ}P_Y$			$(K \circ M)_{++} L_{++}$	$(KML)_{++}$	$(KM)_{++}L_{++}$
$P_{YZ}P_X$				$(L \circ M)_{++} K_{++}$	$(LM)_{++}K_{++}$
$P_X P_Y P_Z$					$K_{++}L_{++}M_{++}$

Table: V-statistic estimators of $\langle \mu_\kappa \nu, \mu_\kappa \nu' \rangle_{\mathcal{H}_\kappa}$ (without terms $P_X P_Y P_Z$). H is centering matrix $I - n^{-1}$

Lancaster interaction statistic: Sejdinovic, G, Bergsma, NIPS13

$$\|\mu_\kappa(\Delta_L P)\|^2_{\mathcal{H}_\kappa} = \frac{1}{n^2} \boxed{(HKH \circ HLH \circ HMH)_{++}}.$$

Empirical joint central moment in the feature space

V-structure discovery

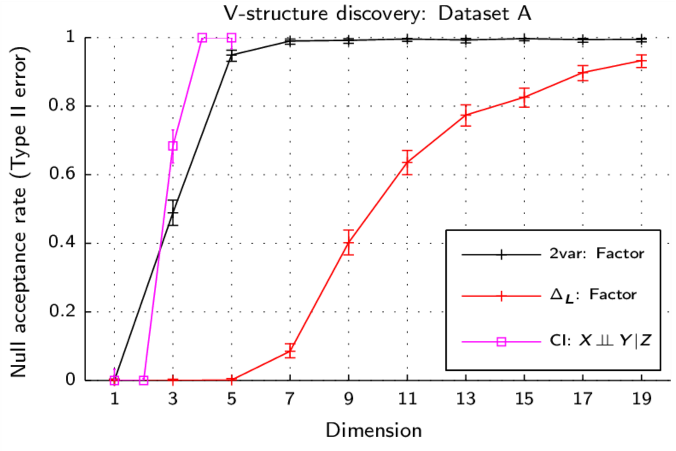

Lancaster test, CI test for $X \perp\!\!\!\perp Y|Z$ from Zhang et al. (2011), and a factorisation test, $n = 500$

Interaction for $D \geq 4$

- Interaction measure valid for all D:

(Streitberg, 1990)

$$\Delta_S P = \sum_{\pi} (-1)^{|\pi|-1} (|\pi| - 1)! \, J_\pi P$$

  - For a partition π, J_π associates to the joint the corresponding factorisation, e.g., $J_{13|2|4} P = P_{X_1 X_3} P_{X_2} P_{X_4}$.

Interaction for $D \geq 4$

- Interaction measure valid for all D:

(Streitberg, 1990)

$$\Delta_S P = \sum_{\pi} (-1)^{|\pi|-1} (|\pi| - 1)! J_\pi P$$

  - For a partition π, J_π associates to the joint the corresponding factorisation, e.g., $J_{13|2|4} P = P_{X_1 X_3} P_{X_2} P_{X_4}$.

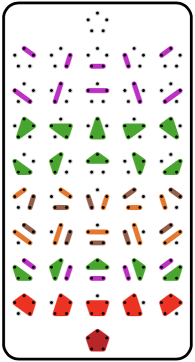

Interaction for $D \geq 4$

- Interaction measure valid for all D:

(Streitberg, 1990)

$$\Delta_S P = \sum_{\pi} (-1)^{|\pi|-1} (|\pi| - 1)! J_\pi P$$

  - For a partition π, J_π associates to the joint the corresponding factorisation, e.g., $J_{13|2|4} P = P_{X_1 X_3} P_{X_2} P_{X_4}$.

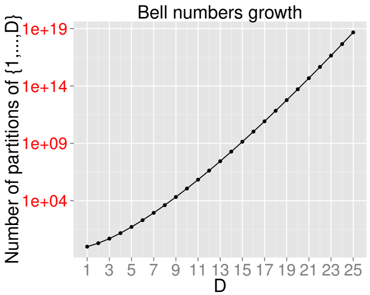

Co-authors

From Gatsby:

- Mikolaj Binkowski
- Kacper Chwialkowski
- Wittawat Jitkrittum
- Heiko Strathmann
- Dougal Sutherland
- Wenkai Xu

External collaborators:

- Kenji Fukumizu
- Bernhard Schoelkopf
- Bharath Sriperumbudur
- Alex Smola
- Zoltan Szabo

Questions?