# Kernel Tricks, Means and Ends 

Bernhard Schölkopf<br>Max Planck Institute for Biological Cybernetics<br>Tübingen, Germany

Empirical Inference Department<br>http://www.kyb.tuebingen.mpg.de/bs

## Learning theory in a nutshell

Learn $\quad f: \mathcal{X} \rightarrow\{ \pm 1\} \quad$ from examples
$\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right) \in \mathcal{X} \times\{ \pm 1\} \quad$ generated from $\mathrm{P}(x, y)$
Goal: minimize expected error

$$
\left.\left.R[f]=\int \frac{1}{2} \right\rvert\, f(x)-y\right) \mid d \mathrm{P}(x, y)
$$

Problem: P is unknown.
Induction principle: "empirical risk minimization"


$$
R_{\mathrm{emp}}[f]=\frac{1}{m} \sum_{i=1}^{m} \frac{1}{2}\left|f\left(x_{i}\right)-y_{i}\right|
$$

Vapnik \& Chervonenkis: this is consistent* iff the "capacity" of the function class is asymptotically well-behaved (e.g., finite VC dim).
Computing the capacity is nontrivial...


## Example of a Pattern Recognition Algorithm

Idea: classify points $x$ according to which of the two class means is closer.

$$
\mu_{+}:=\frac{1}{m_{+}} \sum_{y_{i}=1} x_{i}, \quad \mu_{-}:=\frac{1}{m_{-}} \sum_{y_{i}=-1} x_{i}
$$



- Decision function: hyperplane with normal vector $w:=\mu_{+}-\mu_{-}$
- How about problems that are not linearly separable?



## Feature Spaces

Preprocess the inputs with

$$
\begin{aligned}
\Phi: \mathcal{X} & \rightarrow \mathcal{H} \\
x & \mapsto \Phi(x)
\end{aligned}
$$

where $\mathcal{H}$ is a dot product space, and learn the mapping from $\Phi(x)$ to $y$.

## Example: All Degree 2 Monomials




## The Kernel Trick

$$
\begin{aligned}
\left\langle\Phi(x), \Phi\left(x^{\prime}\right)\right\rangle & =\left(x_{1}^{2}, \sqrt{2} x_{1} x_{2}, x_{2}^{2}\right)\left(x_{1}^{\prime 2}, \sqrt{2} x_{1}^{\prime} x_{2}^{\prime}, x_{2}^{\prime 2}\right)^{\top} \\
& =\left(x_{1} x_{1}^{\prime}+x_{2} x_{2}^{\prime}\right)^{2} \\
& =\left\langle x, x^{\prime}\right\rangle^{2} \\
& =: k\left(x, x^{\prime}\right)
\end{aligned}
$$

$\longrightarrow$ the dot product in $\mathcal{H}$ can be computed from the dot product in $\mathbb{R}^{2}$
More generally: for $x, x^{\prime} \in \mathbb{R}^{N}, d \in \mathbb{N}$,
$\left\langle x, x^{\prime}\right\rangle^{d}=\left(\sum_{j=1}^{N} x_{j} \cdot x_{j}^{\prime}\right)^{d}=\sum_{j_{1}, \cdots, j_{d}=1}^{N} x_{j_{1}} \cdots \cdots x_{j_{d}} \cdot x_{j_{1}}^{\prime} \cdots \cdots x_{j_{d}}^{\prime}=\left\langle\Phi(x), \Phi\left(x^{\prime}\right)\right\rangle$

More generally: works for positive definite kernels


## Positive Definite Kernels

Let $\mathcal{X}$ be a nonempty set. The following two are equivalent:

- $k$ is positive definite ( $p d$ ), i.e., $k$ is symmetric, and for
- any set of training points $x_{1}, \ldots, x_{m} \in \mathcal{X}$ and
- any $a_{1}, \ldots, a_{m} \in \mathbb{R}$
we have

$$
\sum_{i, j} a_{i} a_{j} K_{i j} \geq 0, \text { where } K_{i j}:=k\left(x_{i}, x_{j}\right)
$$

- there exists a map $\Phi$ into a dot product space $\mathcal{H}$ such that

$$
k\left(x, x^{\prime}\right)=\left\langle\Phi(x), \Phi\left(x^{\prime}\right)\right\rangle
$$

(RKHS)
$\mathcal{H}$ is a so-called reproducing kernel Hilbert space.
If for pairwise distinct points, $\Sigma=0$ iff all $a_{i}=0$, call $k$ strictly p.d.


## Construction of $\Phi$


$\Phi(x):=k(x,$.$) \quad (Aronszajn 1950), take linear hull ->$ vector space
$<\Phi(x), \Phi\left(x^{\prime}\right)>:=k\left(x, x^{\prime}\right)$, linear extension, can prove this is a dot product
Point evaluation: $f(x)=\langle f, k(x,)$.$\rangle . "Reproducing kernel Hilbert space"$


## The Kernel Trick - Main Points

- any algorithm that only depends on dot products can benefit from the kernel trick
- $\mathcal{X}$ need not be a vector space
- think of the kernel as a (nonlinear) similarity measure
- examples of common kernels:

Polynomial $k\left(x, x^{\prime}\right)=\left(\left\langle x, x^{\prime}\right\rangle+c\right)^{d}$
Gaussian $k\left(x, x^{\prime}\right)=\exp \left(-\left\|x-x^{\prime}\right\|^{2} /\left(2 \sigma^{2}\right)\right)$

## An Example of a Kernel Algorithm (Schölkopf \& Smola 2002)

Classify points $\mathbf{x}:=\Phi(x)$ in feature space according to which of the two class means is closer.

$$
\mu_{+}:=\frac{1}{m_{+}} \sum_{\left\{i: y_{i}=1\right\}} \Phi\left(x_{i}\right), \quad \mu_{-}:=\frac{1}{m_{-}} \sum_{\left\{i: y_{i}=-1\right\}} \Phi\left(x_{i}\right)
$$



Compute the sign of the dot product between $\mathrm{w}:=\mu_{+}-\mu_{\text {- }}$ and $\mathbf{x}-\mathbf{c}$.

ctd.

$$
\begin{aligned}
f(x) & =\operatorname{sgn}\left(\frac{1}{m_{+}} \sum_{\left\{i: y_{i}=1\right\}}\left\langle\Phi(x), \Phi\left(x_{i}\right)\right\rangle-\frac{1}{m_{-}} \sum_{\left\{i: y_{i}=-1\right\}}\left\langle\Phi(x), \Phi\left(x_{i}\right)\right\rangle+b\right) \\
& =\operatorname{sgn}\left(\frac{1}{m_{+}} \sum_{\left\{i: y_{i}=1\right\}} k\left(x, x_{i}\right)-\frac{1}{m_{-}} \sum_{\left\{i: y_{i}=-1\right\}} k\left(x, x_{i}\right)+b\right)
\end{aligned}
$$

with the constant offset

$$
b=\frac{1}{2}\left(\frac{1}{m_{-}^{2}} \sum_{\left\{(i, j): y_{i}=y_{j}=-1\right\}} k\left(x_{i}, x_{j}\right)-\frac{1}{m_{+}^{2}} \sum_{\left\{(i, j): y_{i}=y_{j}=1\right\}} k\left(x_{i}, x_{j}\right)\right) .
$$

If $k$ is a density, this is a classifier based on Parzen windows plug-in estimates of the two classes.



$$
f(x)=\operatorname{sgn}\left(\sum_{i} \lambda_{i} k\left(x_{i}, x\right)+b\right)
$$

representer theorem (Kimeldorf \& Wahba 1971, Schölkopf et al. 2000)

- unique solution found by convex QP


## Kernel PCA



Contains LLE, Laplacian Eigenmap, and (in the limit) Isomap as special cases with data dependent kernels (Ham et al. 2004)


Schölkopf, Smola \& Müller, 1998


## Application examples

## MNIST Benchmark

handwritten character benchmark ( 60000 training \& 10000 test examples, $28 \times 28$ )

| 5 |  | 4 |  |  | 92 |  |  | 1 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 3 |  | 3 | 6 | 1 | 17 | 72 | 24 | 8 |  | 9 |
| 4 | 0 | 9 |  | 1 | 12 | 24 | 43 | 2 | 2 |  |
| 3 | 8 | 6 | 9 | 0 | 0 | 6 | 60 | 0 | 7 | 6 |
| 1 | 8 | 1 | 19 | 3 | 3 | 98 | 8 | 9 | 9 | 3 |
| 3 | 0 | 7 | 4 | 9 | 48 | 80 | 0.9 | 94 | 4 |  |
| 4 | 4 | 6 | 0 | 4 | 45 | 56 | 61 | 1 | 0 | 0 |
| 1 | 7 | 1 | 6 | 3 | 30 | 02 | 21 | 1 | , | 7 |
| 9 | 0 | 2 | 6 | 7 | 78 | 83 | 39 | $9^{\circ} 0$ | 0 | 4 |
| 6 | 7 | 4 | 6 |  | 80 | 01 | 78 | 83 |  |  |

## MNIST Error Rates

| Classifier | test error | reference |
| :--- | :--- | :--- |
| linear classifier | $8.4 \%$ | Bottou et al. (1994) |
| 3-nearest-neighbour | $2.4 \%$ | Bottou et al. (1994) |
| SVM | $1.4 \%$ | Burges and Schölkopf (1997) |
| Tangent distance | $1.1 \%$ | Simard et al. (1993) |
| LeNet4 | $1.1 \%$ | LeCun et al. (1998) |
| Boosted LeNet4 | $0.7 \%$ | LeCun et al. (1998) |
| Translation invariant SVM | $0.56 \%$ | DeCoste and Schölkopf (2002) |

## PET attenuation correction



## Visual Impression of PET $_{\text {MRAC }}$ and $\mathrm{PET}_{\text {CTAC }}$ almost identical Quantification Error below 1\%

With M. Hofmann, B. Pichler, Radiologische Klinik, Tübingen Tracer: ${ }^{68} \mathrm{Ga}$ DOTA-TOC
Reconstruction performed on PET/CT Scanner, using Image Size $128 \times 128$, OSEM Reconstruction with 4 Iterations, 8 Subsets; Gaussian Filter FWHM 5 mm



## Learning of a Motor Primitive (Work in Progress)




## Kernel Means

Joint work with: K. Borgwardt, K. Fukumizu, A. Gretton, J. Huang, D. Janzing, Q. Le, M. Rasch, A. Smola, L. Song, B. Sriperumbudur, X. Sun

$X$ compact subset of a separable metric space, $m, n \in \mathbb{N}$.
Positive class $X:=\left\{x_{1}, \ldots, x_{m}\right\} \subset \mathcal{X}$
Negative class $Y:=\left\{y_{1}, \ldots, y_{n}\right\} \subset \mathcal{X}$
RKHS means $\mu(X)=\frac{1}{m} \sum_{i=1}^{m} k\left(x_{i}, \cdot\right), \mu(Y)=\frac{1}{n} \sum_{i=1}^{n} k\left(y_{i}, \cdot\right)$.
Get a problem if $\mu(X)=\mu(Y)$.
Schölkopf \& Smola, 2002


When do the means coincide?
$k\left(x, x^{\prime}\right)=\left\langle x, x^{\prime}\right\rangle: \quad$ the means coincide
$k\left(x, x^{\prime}\right)=\left(\left\langle x, x^{\prime}\right\rangle+1\right)^{d}$ : all empirical moments up to order $d$ coincide
$k$ strictly pd: $\quad X=Y$.

The mean "remembers" each point that contributed to it.

Proposition 1 Assume that $k$ is strictly pd, and for all $i, j$, $x_{i} \neq x_{j}$, and $y_{i} \neq y_{j}$. If for some $\alpha_{i}, \beta_{j} \in \mathbb{R}-\{0\}$, we have

$$
\begin{equation*}
\sum_{i=1}^{m} \alpha_{i} k\left(x_{i}, .\right)=\sum_{j=1}^{n} \beta_{j} k\left(y_{j}, .\right), \tag{1}
\end{equation*}
$$

then $X=Y$.
Proof (by contradiction): W.l.o.g., assume that $x_{1} \notin Y$. Subtract $\sum_{j=1}^{n} \beta_{j} k\left(y_{j},.\right)$ from (1), and make it a sum over distinct points, to get

$$
0=\sum_{i} \gamma_{i} k\left(z_{i}, .\right),
$$

where $z_{1}=x_{1}, \gamma_{1}=\alpha_{1} \neq 0$, and $z_{2}, \cdots \in X \cup Y-\left\{x_{1}\right\}, \gamma_{2}, \cdots \in \mathbb{R}$.
Take the dot product with $\sum_{j} \gamma_{j} k\left(z_{j},.\right)$, using $\left\langle k\left(z_{i},.\right), k\left(z_{j},.\right)\right\rangle=k\left(z_{i}, z_{j}\right)$, to get

$$
0=\sum_{i j} \gamma_{i} \gamma_{j} k\left(z_{i}, z_{j}\right),
$$

with $\gamma \neq 0$, hence $k$ cannot be strictly pd.


## The mean map

$$
\mu: X=\left(x_{1}, \ldots, x_{m}\right) \mapsto \frac{1}{m} \sum_{i=1}^{m} k\left(x_{i}, \cdot\right)
$$

satisfies

$$
\langle\mu(X), f\rangle=\left\langle\frac{1}{m} \sum_{i=1}^{m} k\left(x_{i}, \cdot\right), f\right\rangle=\frac{1}{m} \sum_{i=1}^{m} f\left(x_{i}\right)
$$

and
$\|\mu(X)-\mu(Y)\|=\sup _{\|f\| \leq 1}|\langle\mu(X)-\mu(Y), f\rangle|=\sup _{\|f\| \leq 1}\left|\frac{1}{m} \sum_{i=1}^{m} f\left(x_{i}\right)-\frac{1}{n} \sum_{i=1}^{n} f\left(y_{i}\right)\right|$.
Large distance $\Leftrightarrow$ can find a function distinguishing the two samples


## Witness function

$f=\frac{\mu(X)-\mu(Y)}{\|\mu(X)-\mu(Y)\|}$, thus $\left.f(x) \propto\langle\mu(X)-\mu(Y), k(x,)\rangle.\right):$


This function is in the RKHS of a Gaussian kernel, but not in the RKHS of the linear kernel.

The mean map for measures
$p, q$ Borel probability measures,
$\mathrm{E}_{x, x^{\prime} \sim p}\left[k\left(x, x^{\prime}\right)\right], \mathbf{E}_{x, x^{\prime} \sim q}\left[k\left(x, x^{\prime}\right)\right]<\infty(\|k(x,)\| \leq M<.\infty$ is sufficient $)$
Define

$$
\mu: p \mapsto \mathbf{E}_{x \sim p}[k(x, \cdot)]
$$

Note

$$
\langle\mu(p), f\rangle=\mathbf{E}_{x \sim p}[f(x)]
$$

and

$$
\|\mu(p)-\mu(q)\|=\sup _{\|f\| \leq 1}\left|\mathbf{E}_{x \sim p}[f(x)]-\mathbf{E}_{x \sim q}[f(x)]\right|
$$

Recall that in the finite sample case, for strictly p.d. kernels, $\mu$ was injective - how about now?

> Smola et al., ALT'07, Fukumizu et al., NIPS'07

Theorem 2 [Fortet and Mourier (1953); Dudley (2002)]

$$
p=q \Longleftrightarrow \sup _{f \in C(X)}\left|\mathbf{E}_{x \sim p}(f(x))-\mathbf{E}_{x \sim q}(f(x))\right|=0,
$$

where $C(X)$ is the space of continuous bounded functions on $x$.

Theorem 3 [Gretton et al. (2007)] If $k$ is universal, then

$$
p=q \Longleftrightarrow\|\mu(p)-\mu(q)\|=0 .
$$

Proof Idea: combine Theorem 2 with

$$
\|\mu(p)-\mu(q)\|=\sup _{\|f\| \leq 1}\left|\mathbf{E}_{x \sim p}[f(x)]-\mathbf{E}_{x \sim q}[f(x)]\right|
$$

Replace $C(X)$ by the unit ball in an RKHS that is dense in $C(X)$ — universal kernel [51], e.g., Gaussian.

Discussion: solves a high-dim. optimization problem...


- $\mu$ is invertible on its image
$\mathcal{M}=\{\mu(p) \mid p$ is a probability distribution $\}$ (the "marginal polytope", Wainwright and Jordan (2003))
- generalization of the moment generating function of a RV $x$ with distribution $p$ :

$$
M_{p}(.)=\mathbf{E}_{x \sim p}\left[e^{\langle x, \cdot\rangle}\right]
$$

- assume we have densities, the kernel is shift invariant, $k(x, y)=\phi(x-y)$, and all Fourier transforms exist. Note that $\mu$ is invertible iff

$$
\begin{aligned}
\int k(x-y) p(y) d x=\int k(x-y) q(y) d x & \Rightarrow p=q \\
\text { i.e., } & \hat{\phi}(\hat{p}-\hat{q})=0
\end{aligned} \begin{aligned}
& \Rightarrow p=q
\end{aligned}
$$

(Sriperumbudur et al., 2008)


## Application 1: Two-sample problem (Gretton et al., 2007)

$X, Y$ i.i.d. $m$-samples from $p, q$, respectively.

$$
\begin{aligned}
\|\mu(p)-\mu(q)\|^{2} & =\mathrm{E}_{x, x^{\prime} \sim p}\left[k\left(x, x^{\prime}\right)\right]-2 \mathbf{E}_{x \sim p, y \sim q}[k(x, y)]+\mathbf{E}_{y, y^{\prime} \sim q}\left[k\left(y, y^{\prime}\right)\right] \\
& =\mathrm{E}_{x, x^{\prime} \sim p, y, y^{\prime} \sim q}\left[h\left((x, y),\left(x^{\prime}, y^{\prime}\right)\right)\right]
\end{aligned}
$$

with

$$
h\left((x, y),\left(x^{\prime}, y^{\prime}\right)\right):=k\left(x, x^{\prime}\right)-k\left(x, y^{\prime}\right)-k\left(y, x^{\prime}\right)+k\left(y, y^{\prime}\right) .
$$

Define

$$
\begin{aligned}
D(p, q)^{2} & :=\mathbf{E}_{x, x^{\prime} \sim p, y, y^{\prime} \sim q} h\left((x, y),\left(x^{\prime}, y^{\prime}\right)\right) \\
\hat{D}(X, Y)^{2} & :=\frac{1}{m(m-1)} \sum_{i \neq j} h\left(\left(x_{i}, y_{i}\right),\left(x_{j}, y_{j}\right)\right) .
\end{aligned}
$$

$\hat{D}(X, Y)^{2}$ is an unbiased estimator of $D(p, q)^{2}$.
It's easy to compute, and works on structured data.

Theorem 4 Assume $k$ is bounded.
$\hat{D}(X, Y)^{2}$ converges to $D(p, q)^{2}$ in probability with rate $\mathcal{O}\left(m^{-\frac{1}{2}}\right)$.
This could be used as a basis for a test, but uniform convergence bounds are often loose..
Theorem 5 We assume $\mathbf{E}\left(h^{2}\right)<\infty$. When $p \neq q$, then $\sqrt{m}\left(\hat{D}(X, Y)^{2}-D(p, q)^{2}\right)$ converges in distribution to a zero mean Gaussian with variance

$$
\sigma_{u}^{2}=4\left(\mathbf{E}_{z}\left[\left(\mathbf{E}_{z} h\left(z, z^{\prime}\right)\right)^{2}\right]-\left[\mathbf{E}_{z, z}\left(h\left(z, z^{\prime}\right)\right)\right]^{2}\right)
$$

When $p=q$, then $m\left(\hat{D}(X, Y)^{2}-D(p, q)^{2}\right)=m \hat{D}(X, Y)^{2}$ converges in distribution to

$$
\begin{equation*}
\sum_{l=1}^{\infty} \lambda_{l}\left[q_{l}^{2}-2\right] \tag{2}
\end{equation*}
$$

where $q_{l} \sim \mathcal{N}(0,2)$ i.i.d., $\lambda_{i}$ are the solutions to the eigenvalue equation

$$
\int_{x} \tilde{k}\left(x, x^{\prime}\right) \psi_{i}(x) d p(x)=\lambda_{i} \psi_{i}\left(x^{\prime}\right)
$$

and $\tilde{k}\left(x_{i}, x_{j}\right):=k\left(x_{i}, x_{j}\right)-\mathbf{E}_{x} k\left(x_{i}, x\right)-\mathbf{E}_{x} k\left(x, x_{j}\right)+\mathbf{E}_{x, x^{\prime}} k\left(x, x^{\prime}\right)$ is the centred RKHS kernel.

## Application 2: Dependence Measures

Assume that $(x, y)$ are drawn from $p_{x y}$, with marginals $p_{x}, p_{y}$. Want to know whether $p_{x y}$ factorizes into its marginals.

Bach and Jordan (2002); Fukumizu et al. (2004): kernel generalized variance
Gretton et al. (2005a,b): kernel constrained covariance, HSIC
Main idea (Rényi, 1959; Jacod and Protter, 2000):
$x$ and $y$ independent $\Longleftrightarrow$

$$
\sup _{f, g \text { bounded \& continuous }} \operatorname{Cov}(f(x), g(y))=0
$$

Kernel version:

$$
\sup \quad \operatorname{Cov}(f(x), g(y))=0
$$

$f, g \in$ unit balls in RKHS

$$
\operatorname{cov}(\mathrm{x}, \mathrm{y}):=\mathbf{E}_{\mathrm{x}, \mathrm{y}}[\mathrm{xy}]-\mathbf{E}_{\mathrm{x}}[\mathrm{x}] \mathbf{E}_{\mathrm{y}}[\mathrm{y}]
$$


$k$ kernel on $x \times y$.

$$
\begin{aligned}
\mu\left(p_{x y}\right) & :=\mathbf{E}_{(x, y) \sim p_{x y}}[k((x, y), \cdot)] \\
\mu\left(p_{x} \times p_{y}\right) & :=\mathbf{E}_{x \sim p_{x}, y \sim p_{y}}[k((x, y), \cdot)] .
\end{aligned}
$$

Use $\Delta:=\left\|\mu\left(p_{x y}\right)-\mu\left(p_{x} \times p_{y}\right)\right\|$ as a measure of dependence.
For $k\left((x, y),\left(x^{\prime}, y^{\prime}\right)\right)=k_{x}\left(x, x^{\prime}\right) k_{y}\left(y, y^{\prime}\right)$ :
$\Delta^{2}$ equals the Hilbert-Schmidt norm of the covariance operator between the two RKHSs (HSIC), with empirical estimate $m^{-2} \operatorname{tr} H K_{x} H K_{y}$, where $H=I-1 / m$

Gretton et al. (2005a); Smola et al. (2007).


Witness function of the equivalent optimisation problem:


Application: learning causal structures (Sun, Janzing, Schölkopf, Fukumizu, ICML 2007; Fukumizu, Gretton, Sun, Schölkopf, NIPS 2007))

## Causal Inference

Forward model: $y=f(x)+n$ with $x, n$ independent. Question: when is there a corresponding backward model?

$$
f(x)=x
$$





Theorem 1 Let the joint probability density of $x$ and $y$ be given by

$$
\begin{equation*}
p(x, y)=p_{n}(y-f(x)) p_{x}(x) \tag{2}
\end{equation*}
$$

where $p_{n}, p_{x}$ are probability densities on $\mathbb{R}$. If there is a backward model of the same form, i.e.,

$$
\begin{equation*}
p(x, y)=p_{n}(x-g(y)) p_{y}(y), \tag{3}
\end{equation*}
$$

then, denoting $\nu:=\log p_{n}$ and $\xi:=\log p_{z}$, the triple $\left(f, p_{x}, p_{n}\right)$ must satisfy the following differential equation for all $x, y$ with $\nu^{\prime \prime}(y-f(x)) f^{\prime}(x) \neq 0$ :

$$
\begin{equation*}
\xi^{\prime \prime \prime}=\xi^{\prime \prime}\left(-\frac{\nu^{\prime \prime \prime} f^{\prime}}{\nu^{\prime \prime}}+\frac{f^{\prime \prime}}{f^{\prime}}\right)-2 \nu^{\prime \prime} f^{\prime \prime} f^{\prime}+\nu f^{\prime \prime \prime}+\frac{\nu^{\prime} \nu^{\prime \prime \prime} f^{\prime \prime} f^{\prime}}{\nu^{\prime \prime}}-\frac{\nu\left(f^{\prime \prime}\right)^{2}}{f^{\prime}}, \tag{4}
\end{equation*}
$$

where we have skipped the arguments $y-f(x), x$, and $x$ for $v, \xi$, and $f$, respectively. Moreover, if for a fuxed pair $(f, \nu)$ there exists $y \in \mathbb{R}$ such that $\nu^{\prime \prime}(y-f(x)) f^{\prime}(x) \neq 0$ for almost all $x \in \mathbb{R}$, the set of all $p_{2}$ for which $p$ has a backward model is contained in a 3-dimensional affine space.

A simple corollary is that if both the marginal density $p_{x}(x)$ and the noise density $p_{n}(y-f(x))$ are Gaussian then the existence of a backward model implies linearity of $f$ :

Corollary 1 Assume that $\nu^{\prime \prime \prime}=\xi^{\prime \prime \prime}=0$ everywhere. If a backward model exists, then $f$ is linear.
(Hoyer, Janzing, Mooij, Peters, Schölkopf, 2008)

## Application 3: Covariate Shift Correction and Local Learning

training set $X=\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)\right\}$ drawn from $p$,
test set $X^{\prime}=\left\{\left(x_{1}^{\prime}, y_{1}^{\prime}\right), \ldots,\left(x_{n}^{\prime}, y_{n}^{\prime}\right)\right\}$ from $p^{\prime} \neq p$.
Assume $p_{y \mid x}=p_{y \mid x}^{\prime}$.
Shimodaira (2000): reweight training set

Minimize

$$
\left\|\sum_{i=1}^{m} \beta_{i} k\left(x_{i}, \cdot\right)-\mu\left(X^{\prime}\right)\right\|^{2}+\lambda\|\beta\|_{2}^{2} \text { subject to } \beta_{i} \geq 0, \sum_{i} \beta_{i}=1 .
$$

Equivalent QP:

$$
\begin{aligned}
& \underset{\beta}{\operatorname{minimize}} \frac{1}{2} \beta^{\top}(K+\lambda \mathbf{1}) \beta-\beta^{\top} l \\
& \text { subject to } \beta_{i} \geq 0 \text { and } \sum_{i} \beta_{i}=1,
\end{aligned}
$$

where $K_{i j}:=k\left(x_{i}, x_{j}\right), l_{i}=\left\langle k\left(x_{i}, \cdot\right), \mu\left(X^{\prime}\right)\right\rangle$.
Experiments show that in underspecified situations (e.g., large kernel widths), this helps (Huang et al., 2007b).
$X^{\prime}=\left\{x^{\prime}\right\}$ leads to a local sample weighting scheme.


Application 4: Measure estimation and dataset squashing

Given a sample $X$, minimize

$$
\|\mu(X)-\mu(p)\|^{2}
$$

over a convex combination of measures $p_{i}$,

$$
p=\sum_{i} \alpha_{i} p_{i}, \quad \alpha_{i} \geq 0, \quad \sum_{i} \alpha_{i}=1
$$

This can be written as a convex QP with objective function

$$
\|\mu(X)-\mu(p)\|^{2}=\alpha^{\top} Q \alpha+1_{m}^{\top} K 1_{m}-2 \alpha^{\top} L 1_{m}
$$

where

$$
\begin{aligned}
L_{i j} & :=\mathbf{E}_{x \sim p_{i}}\left[k\left(x, x_{j}\right)\right] \\
Q_{i j} & :=\mathbf{E}_{x \sim p_{i}, x^{\prime} \sim p_{j}}\left[k\left(x, x^{\prime}\right)\right] \\
K_{i j} & =k\left(x_{i}, x_{j}\right) \\
1_{m} & :=(1 / m, \ldots, 1 / m)^{\top} \in \mathbb{R}^{m} .
\end{aligned}
$$



In practice, use

$$
\alpha^{\top}[Q+\lambda I] \alpha-2 \alpha^{\top} L 1_{m}
$$

Some cases where $Q$ and $L$ can be computed in closed form (Smola et al., 2007):

- Gaussian $p_{i}$ and $k$ (cf. Balakrishnan and Schonfeld (2006); Walder et al. (2007))
- $X$ training set, Dirac measures $p_{i}=\delta_{x_{i}}$ : dataset squashing, DuMouchel et al. (1999)
- $X$ test set, Dirac measures $p_{i}=\delta_{y_{i}}$ centered on the training points $Y$ : covariate shift correction Huang et al. (2007a)


## Implicit Surface Fitting

Given a sampling of a surface

$$
\boldsymbol{x}_{1}, \boldsymbol{x}_{2}, \ldots \boldsymbol{x}_{m} \subset \mathbb{R}^{d}
$$

possibly with corresponding surface normals

$$
\boldsymbol{n}_{1}, \boldsymbol{n}_{2}, \ldots \boldsymbol{n}_{m} \subset \mathbb{R}^{d}
$$

Construct a function $f$ whose zero level approximates the surface


## SVM Implicit Surface Approximation



## Large Scale Example (Walder et al. 2006)



Left: Rendered model of Lucy, constructed from 14 million points with normals.
Middle: Each of the 364,982 basis function centres
Right: A planar slice that cuts the nose.


## More Examples

## "

Dragon 1: 440K points - decreasing regularisation


## Interpolation in 4D

4D implicit. No data during red interval.


## The Morphing Problem




## Correspondence



Given a dense correspondence field (or warp), we can interpolate (and extrapolate) images, almost as in a linear space
(cf. Blanz \& Vetter, 1999)

## Correspondence via Machine Learning (Schölkopf, Steinke, Blanz, 2005)

- Objects $\boldsymbol{O}_{\boldsymbol{1}}$ and $\boldsymbol{O}_{2}$ living in X. The warp is a mapping

$$
\tau: X->X .
$$

- Given surface points $x_{i}$ of the $O_{1}$ and $z_{i}$ of $\boldsymbol{O}_{2}$.
- If they are in correspondence, we have a training set $\left(x_{1}, z_{1}\right), \ldots$, $\left(x_{m}, z_{m}\right)$ of "landmark points" and can do regression.
- What if they are not in correspondence?
- Main idea: $\tau$ should be such that


## $O_{1}$ relative to $x$ "looks like" $O_{2}$ relative to $\tau(x)$

- Formalize this as a locational cost

$$
c\left(O_{1}, x, O_{2}, \tau(x)\right)
$$

## Locational Cost Functions

feature functions $f_{1}, f_{2}: \mathcal{X} \rightarrow \mathbb{R}$ think of $f_{1}, f_{2}$ as the signed distance functions of $O_{1}, O_{2}$.

1. $d\left(f_{1}(x), f_{2}(\tau(x))\right)^{2}$
2. $\sum_{i=0}^{\infty} \alpha_{i} d\left(\nabla^{i} f_{1}(x), \nabla^{i} f_{2}(\tau(x))\right)^{2}$
3. If $\Psi$ is the feature map associated with a p.d. kernel on $(\mathcal{O} \times \mathcal{X}) \times(\mathcal{O} \times \mathcal{X})$.

$$
\left\|\Psi\left(O_{1}, x\right)-\Psi\left(O_{2}, \tau(x)\right)\right\|^{2}
$$

## Optimization Problem

- Component functions: for $d=1, \ldots, D$,

$$
\tau_{d}(x)=x_{d}+\left\langle\mathbf{w}_{d}, \Phi(x)\right\rangle
$$

- Minimize

$$
\begin{aligned}
\frac{1}{2} \sum_{d=1}^{D}\left\|\mathbf{w}_{d}\right\|^{2} & +\lambda_{p} \sum_{i=1}^{m}\left\|\tau\left(x_{i}\right)-z_{i}\right\|^{2} \\
& +\lambda_{l o c} \int_{\mathcal{X}} c_{l o c}\left(O_{1}, x, O_{2}, \tau(x)\right) d \mu(x)
\end{aligned}
$$

- For $\lambda_{l o c}=0: D$ SVR problems with quadratic loss
- in the generic case, nonconvex



## Toy Example

Signed distance


## Object Morphing


(signed distance and normals, no landmark points, no color information)

## Head Morphing




Steinke et al., NIPS 2006


with Dept. of Physiology, MPI for Biological Cybernetics

Bernhard Schölkopf, September 11, 2008



Walder et al., 2008


thank you for your attention

