BBCI Summer School, Berlin, Germany Sep. 25, 2012

Density Ratio Estimation in Machine Learning

Masashi Sugiyama
Tokyo Institute of Technology, Japan
sugi@cs.titech.ac.jp
http://sugiyama-www.cs.titech.ac.jp/~sugi/

Machine Learning (ML)

- Goal: Learn information hidden behind data
- There are many ML tasks:
  - Learning under non-stationarity, domain adaptation, multi-task learning, two-sample test, outlier detection, change detection in time series, independence test, feature selection, dimension reduction, independent component analysis, canonical dependency analysis, causal inference, clustering, object matching, conditional probability estimation, probabilistic classification, etc.

Universal Approach

- Learning data-generating probability distribution allows us to solve all ML tasks.

Knowing data generating distributions ⟷ Knowing all about data

- Ex: Pattern recognition is possible if data-generating probability distributions for each class are available.

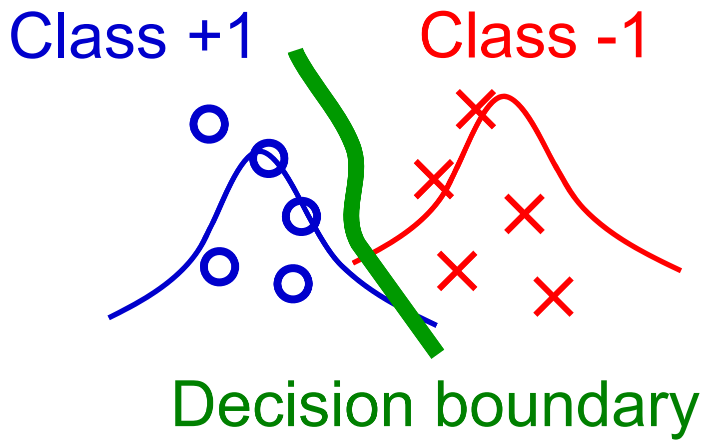

Task-Specific Approach

- However, estimation of probability distributions is known to be difficult.
- Avoid probability distribution estimation and solve the target task directly.
- Ex: Support vector machine

Cortes & Vapnik (ML1995)

  - Directly learn a decision boundary without estimating data-generating distributions.

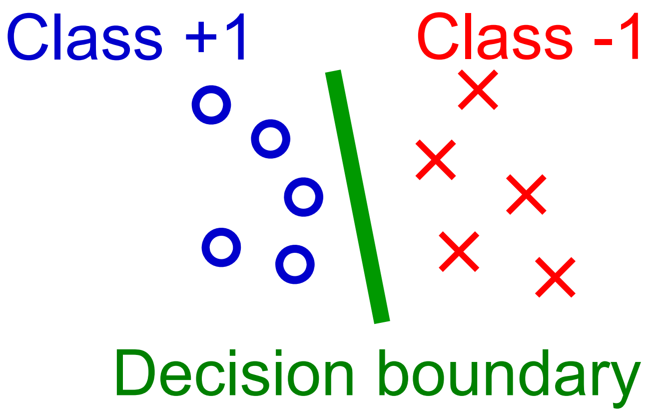

Task-Specific Approach (cont.)

- In principle, task-specific approaches can be more accurate than the universal approach.

- However, research and development for each ML task is highly costly and cumbersome:
  - Theory, algorithms, implementation, education, etc.

Intermediate Approach

- Our target: R&D for a group of tasks:
  - Learning under non-stationarity, domain adaptation, multi-task learning, two-sample test, outlier detection, change detection in time series, independence test, feature selection, dimension reduction, independent component analysis, canonical dependency analysis causal inference, clustering, object matching, conditional probability estimation, probabilistic classification, etc.

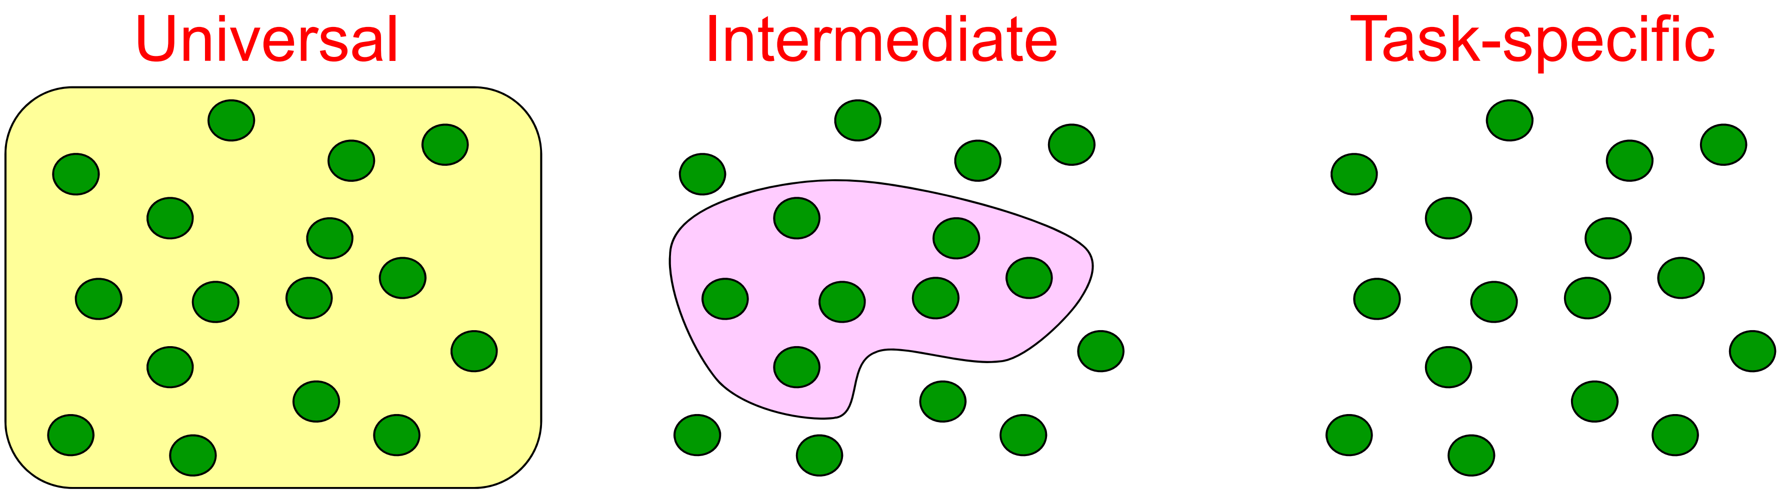

Density-Ratio Estimation

- All ML tasks listed in the previous page include multiple probability distributions.

$$p(\boldsymbol{x}), q(\boldsymbol{x})$$

- For solving these ML tasks, individual distributions are not necessary.
- But knowing the density ratio is enough.

$$r(\boldsymbol{x}) = \frac{p(\boldsymbol{x})}{q(\boldsymbol{x})}$$

- We directly estimate the density ratio without going through density estimation.

Intuitive Justification

Vapnik's principle: Vapnik (1998)

When solving a problem of interest, one should not solve a more general problem as an intermediate step

Knowing densities

$$p(\boldsymbol{x}), q(\boldsymbol{x})$$

Knowing ratio

$$r(\boldsymbol{x}) = \frac{p(\boldsymbol{x})}{q(\boldsymbol{x})}$$

- Estimating the density ratio is substantially easier than estimating densities!

Quick Conclusions

- Density ratios can be accurately and efficiently estimated by simple least-squares!
- Many ML tasks can be solved just by LS:
  - Importance sampling: $\sum_{i=1}^{n} \frac{p_{\mathrm{test}}(\boldsymbol{x}_i)}{p_{\mathrm{train}}(\boldsymbol{x}_i)} \mathrm{loss}(\boldsymbol{x}_i)$
  - Divergence estimation: $\int p(\boldsymbol{x}) \log \frac{p(\boldsymbol{x})}{q(\boldsymbol{x})} \mathrm{d}\boldsymbol{x}$
  - Mutual information estimation: $\iint p(\boldsymbol{x}, \boldsymbol{y}) \log \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x}) p(\boldsymbol{y})} \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{y}$
  - Conditional probability estimation: $p(\boldsymbol{y}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})}$

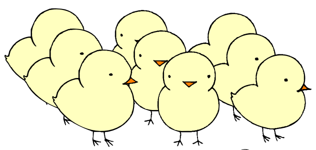

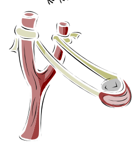

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
4. More on Density Ratio Estimation
5. Conclusions

Density Ratio Estimation: Problem Formulation

- Goal: Estimate the density ratio

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

from data

$$\{\boldsymbol{x}_i^{\mathrm{nu}}\}_{i=1}^{n_{\mathrm{nu}}} \overset{i.i.d.}{\sim} p_{\mathrm{nu}}(\boldsymbol{x})$$

$$\{\boldsymbol{x}_j^{\mathrm{de}}\}_{j=1}^{n_{\mathrm{de}}} \overset{i.i.d.}{\sim} p_{\mathrm{de}}(\boldsymbol{x})$$

Density Estimation Approach

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

$$\{\boldsymbol{x}_i^{\mathrm{nu}}\}_{i=1}^{n_{\mathrm{nu}}} \stackrel{i.i.d.}{\sim} p_{\mathrm{nu}}(\boldsymbol{x})$$

$$\{\boldsymbol{x}_j^{\mathrm{de}}\}_{j=1}^{n_{\mathrm{de}}} \stackrel{i.i.d.}{\sim} p_{\mathrm{de}}(\boldsymbol{x})$$

- Naïve 2-step approach:
  1. Perform density estimation:

     $$\widehat{p}_{\mathrm{nu}}(\boldsymbol{x}), \widehat{p}_{\mathrm{de}}(\boldsymbol{x})$$

  2. Compute the ratio of estimated densities:

     $$\widehat{r}(\boldsymbol{x}) = \frac{\widehat{p}_{\mathrm{nu}}(\boldsymbol{x})}{\widehat{p}_{\mathrm{de}}(\boldsymbol{x})}$$

- However, this works poorly because 1. is performed without regard to 2.

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
   - A) Density Fitting
   - B) Density-Ratio Fitting
3. Usage of Density Ratios
4. More on Density Ratio Estimation
5. Conclusions

Kullback-Leibler Importance Estimation Procedure (KLIEP)

Nguyen, Wainwright & Jordan (NIPS2007)
MS, Nakajima, Kashima, von Bünau & Kawanabe (NIPS2007)

- Minimize KL divergence from $p_{\mathrm{nu}}(\boldsymbol{x})$ to $\widehat{p}_{\mathrm{nu}}(\boldsymbol{x}) = \widehat{r}(\boldsymbol{x})p_{\mathrm{de}}(\boldsymbol{x})$:

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

$$\min_{\widehat{r}} \underbrace{\int p_{\mathrm{nu}}(\boldsymbol{x}) \log \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{\widehat{r}(\boldsymbol{x})p_{\mathrm{de}}(\boldsymbol{x})} d\boldsymbol{x}}_{=:\mathrm{KL}(\widehat{r})}$$

- Decomposition of KL:

$$\mathrm{KL}(\widehat{r}) = C - \int p_{\mathrm{nu}}(\boldsymbol{x}) \log \widehat{r}(\boldsymbol{x}) d\boldsymbol{x}$$

KLIEP: Formulation

- Objective function:

$$\max_{\widehat{r}} \int p_{\mathrm{nu}}(\boldsymbol{x}) \log \widehat{r}(\boldsymbol{x}) d\boldsymbol{x}$$

- Constraints:
  - $\widehat{p}_{\mathrm{nu}}(\boldsymbol{x}) = \widehat{r}(\boldsymbol{x}) p_{\mathrm{de}}(\boldsymbol{x})$ is a probability density:

$$\int \widehat{r}(\boldsymbol{x}) p_{\mathrm{de}}(\boldsymbol{x}) d\boldsymbol{x} = 1 \qquad \widehat{r}(\boldsymbol{x}) \geq 0$$

- Linear-in-parameter density-ratio model:

$$\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{b} \alpha_\ell \phi_\ell(\boldsymbol{x}) = \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x}) \qquad \phi_\ell(\boldsymbol{x}) \geq 0$$

(ex. Gauss kernel)

KLIEP: Algorithm

- Approximate expectations by sample averages:

$$\max_{\boldsymbol{\alpha}} \sum_{i=1}^{n_{\mathrm{nu}}} \log(\boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}})) \quad \text{subject to } \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}}) = 1 \text{ and } \boldsymbol{\alpha} \geq \mathbf{0}$$

- This is convex optimization, so repeating
  - Gradient ascent
  - Projection onto the feasible region

  leads to the global solution.
- The global solution is sparse!

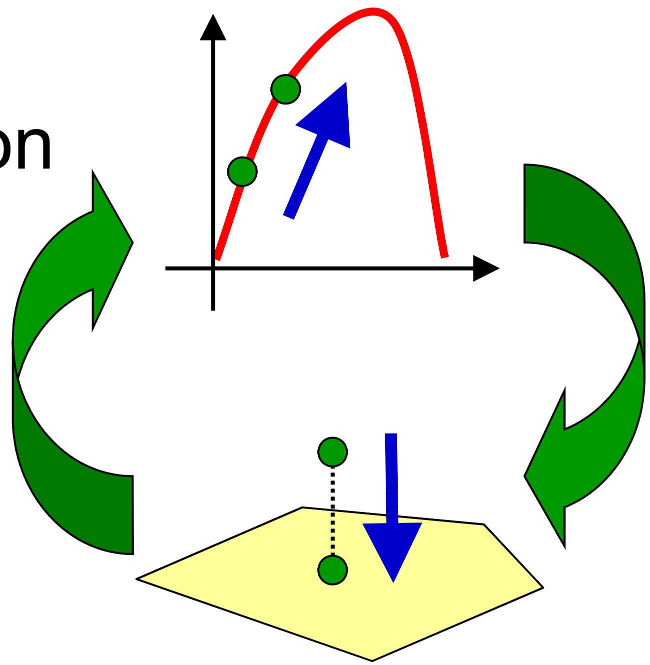

KLIEP: Convergence Properties

Nguyen, Wainwright & Jordan (IEEE-IT2010)
MS, Suzuki, Nakajima, Kashima, von Bünau & Kawanabe (AISM2008)

- Parametric case: $\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{b} \alpha_\ell \phi_\ell(\boldsymbol{x})$
  - Learned parameter converge to the optimal value with order $n^{-\frac{1}{2}}$, which is the optimal rate.

    $n = \min(n_{\mathrm{nu}}, n_{\mathrm{de}})$
- Non-parametric case: $\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{n_{\mathrm{nu}}} \alpha_\ell K(\boldsymbol{x}, \boldsymbol{x}_\ell^{\mathrm{nu}})$
  - Learned function converges to the optimal function with order $n^{-\frac{1}{2+\gamma}}$, which is the optimal rate.

    $0 < \gamma < 2$: Complexity of the function class related to the covering number or bracketing entropy

KLIEP: Numerical Example

$$n_{\mathrm{nu}} = n_{\mathrm{de}} = 200$$

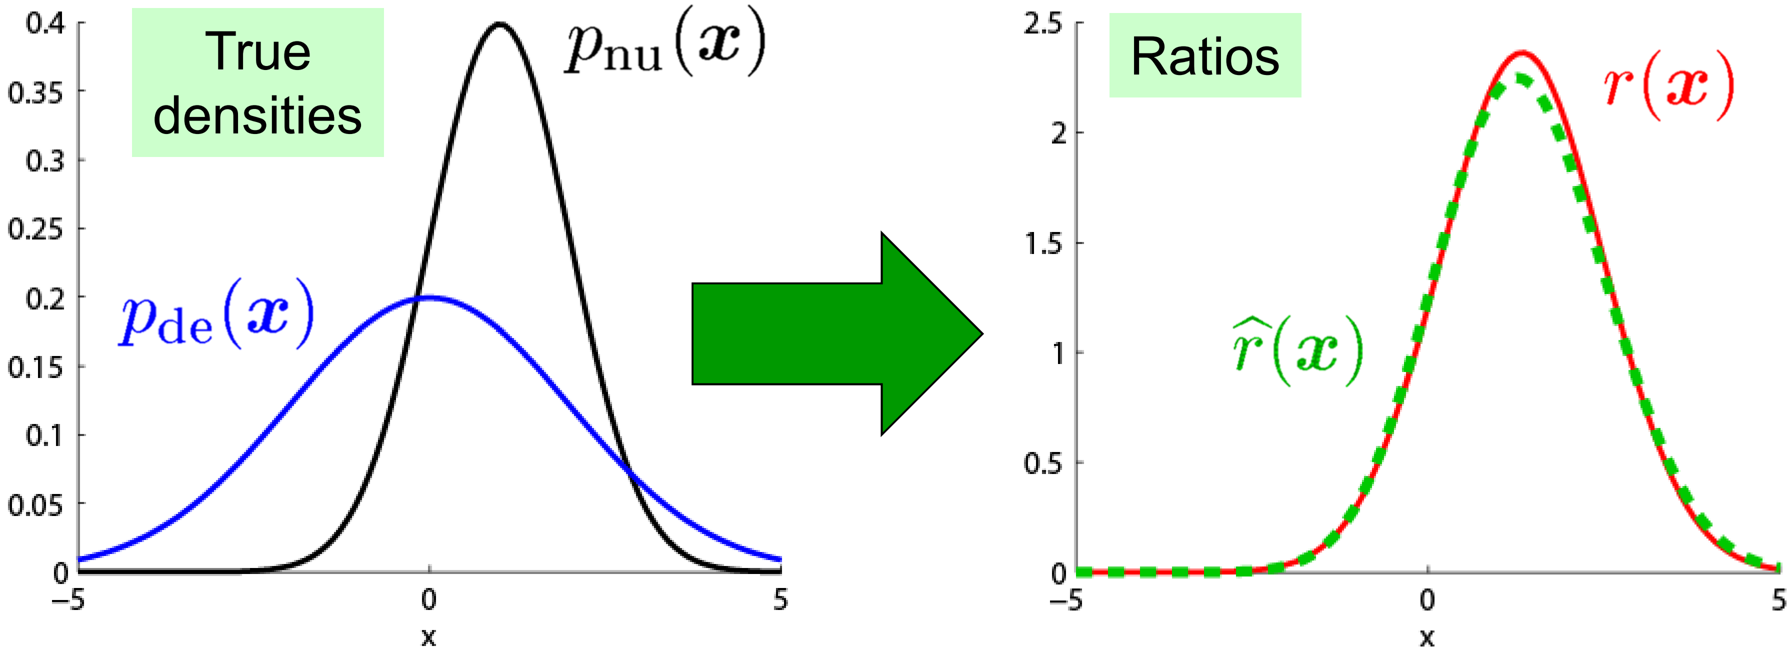

- Gaussian width can be determined by cross-validation with respect to KL.

$$\int p_{\mathrm{nu}}(\boldsymbol{x}) \log \widehat{r}(\boldsymbol{x}) d\boldsymbol{x}$$

KLIEP: Summary

- Density estimation is not involved.
- Cross-validation is available for kernel parameter selection.
- Variations for various models exist:
  - Log-linear, Gaussian mixture, PCA mixture, etc.
- An unconstrained variant corresponds to maximizing a lower-bound of KL divergence.

$$\int p_{\mathrm{nu}}(\boldsymbol{x}) \log \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})} \mathrm{d}\boldsymbol{x}$$

Nguyen, Wainwright & Jordan (NIPS2007)

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
   - A) Density Fitting
   - B) Density-Ratio Fitting
3. Usage of Density Ratios
4. More on Density Ratio Estimation
5. Conclusions

Least-Squares Importance Fitting (LSIF)

Kanamori, Hido & MS (NIPS2008)

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

- Minimize squared-loss (SQ):

$$\min_{\widehat{r}} \underbrace{\int \left(\widehat{r}(\boldsymbol{x}) - r(\boldsymbol{x})\right)^2 p_{\mathrm{de}}(\boldsymbol{x}) d\boldsymbol{x}}_{=:\mathrm{SQ}(\widehat{r})}$$

- Decomposition and approximation of SQ:

$$\mathrm{SQ}(\widehat{r}) = \int \left(\widehat{r}(\boldsymbol{x})\right)^2 p_{\mathrm{de}}(\boldsymbol{x}) d\boldsymbol{x} - 2 \int \widehat{r}(\boldsymbol{x}) p_{\mathrm{nu}}(\boldsymbol{x}) d\boldsymbol{x} + C$$

$$\approx \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \widehat{r}(\boldsymbol{x}_j^{\mathrm{de}})^2 - \frac{2}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \widehat{r}(\boldsymbol{x}_i^{\mathrm{nu}}) + C$$

Constrained LSIF Formulation

- Linear (or kernel) density-ratio model:

$$\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{b} \alpha_\ell \phi_\ell(\boldsymbol{x}) = \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x})$$

- Constrained LSIF (cLSIF):
  - Non-negativity constraint with ℓ_1-regularizer

$$\min_{\boldsymbol{\alpha}} \left[\frac{1}{2} \boldsymbol{\alpha}^\top \widehat{\boldsymbol{H}} \boldsymbol{\alpha} - \widehat{\boldsymbol{h}}^\top \boldsymbol{\alpha} + \lambda \boldsymbol{\alpha}^\top \boldsymbol{1} \right]$$

$$\text{subject to } \boldsymbol{\alpha} \geq \boldsymbol{0}$$

  - A convex quadratic program with sparse solution.

$$\widehat{\boldsymbol{H}} = \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}}) \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}})^\top \qquad \widehat{\boldsymbol{h}} = \frac{1}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}})$$

cLSIF: Regularization Path Tracking

$$\min_{\boldsymbol{\alpha}} \left[\frac{1}{2} \boldsymbol{\alpha}^\top \widehat{\boldsymbol{H}} \boldsymbol{\alpha} - \widehat{\boldsymbol{h}}^\top \boldsymbol{\alpha} + \lambda \boldsymbol{\alpha}^\top \mathbf{1} \right] \quad \text{subject to } \boldsymbol{\alpha} \geq \mathbf{0}$$

- The solution path is piece-wise linear with respect to the regularization parameter λ .

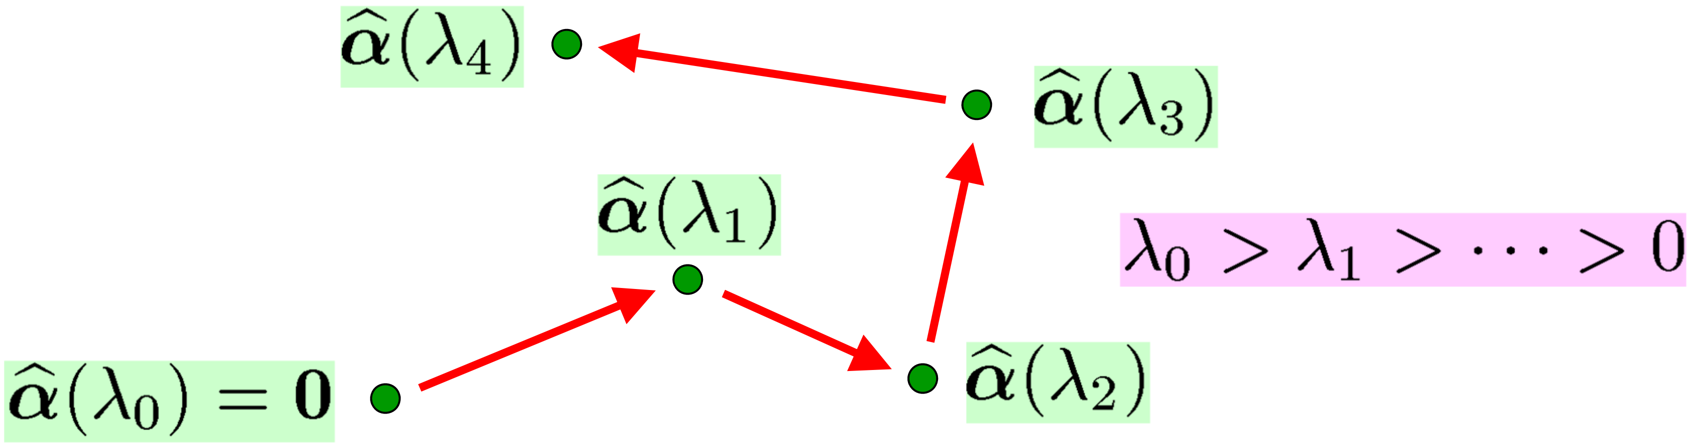

- Solutions for all λ can be computed efficiently without QP solvers!

Unconstrained LSIF Formulation

$$\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{b} \alpha_\ell \phi_\ell(\boldsymbol{x}) = \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x})$$

Unconstrained LSIF (uLSIF):

- uLSIF: No constraint with ℓ_2-regularizer

$$\min_{\boldsymbol{\alpha}} \left[\frac{1}{2} \boldsymbol{\alpha}^\top \widehat{\boldsymbol{H}} \boldsymbol{\alpha} - \widehat{\boldsymbol{h}}^\top \boldsymbol{\alpha} + \frac{\lambda}{2} \boldsymbol{\alpha}^\top \boldsymbol{\alpha} \right]$$

- Analytic solution is available: $(\widehat{\boldsymbol{H}} + \lambda \boldsymbol{I})^{-1} \widehat{\boldsymbol{h}}$

$$\widehat{\boldsymbol{H}} = \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}}) \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}})^\top \qquad \widehat{\boldsymbol{h}} = \frac{1}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}})$$

uLSIF: Analytic LOOCV Score

- Leave-one-out cross-validation (LOOCV):

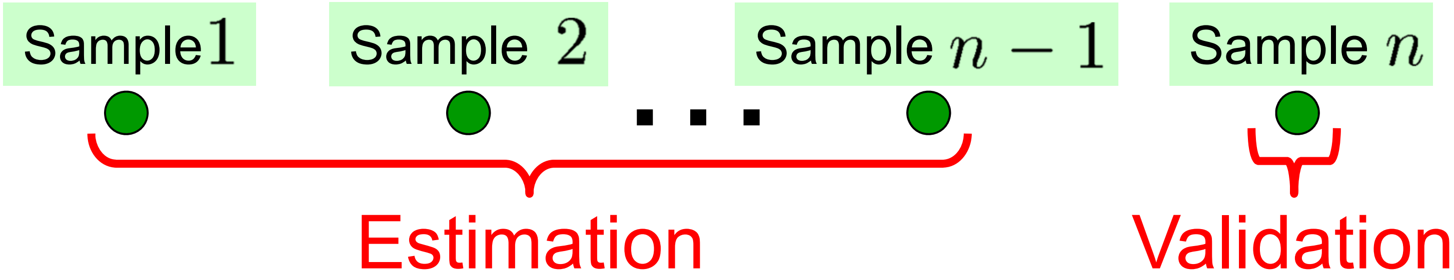

$\widehat{r}(\boldsymbol{x})$

$\frac{1}{2}\widehat{r}(\boldsymbol{x})^2 - \widehat{r}(\boldsymbol{x})$

- LOOCV generally requires n repetitions.
- However, it can be analytically computed for uLSIF (Sherman-Woodbury-Morrison formula).
- Computation time including model selection is significantly reduced.

uLSIF: Theoretical Properties

- **Parametric convergence:**
  - Learned parameter converge to the optimal value with order $n^{-\frac{1}{2}}$, which is the optimal rate.

    $n = \min(n_{\mathrm{nu}}, n_{\mathrm{de}})$ Kanamori, Hido & MS (JMLR2009)

- **Non-parametric convergence:**
  - Learned function converges to the optimal function with order $n^{-\frac{1}{2+\gamma}}$ (depending on the bracketing entropy), which is the optimal rate.

    $0 < \gamma < 2$ Kanamori, Suzuki & MS (MLJ2012)

- **Non-parametric numerical stability:**
  - uLSIF has the smallest condition number among a class of density ratio estimators.

    Kanamori, Suzuki & MS (MLJ2012)

uLSIF: Numerical Example

$$p_{\mathrm{nu}}(\boldsymbol{x}) = N(\boldsymbol{x}; (0,0,\ldots,0)^\top, \boldsymbol{I}_d)$$

$$p_{\mathrm{de}}(\boldsymbol{x}) = N(\boldsymbol{x}; (1,0,\ldots,0)^\top, \boldsymbol{I}_d)$$

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

$n_{\mathrm{nu}} = 1000, n_{\mathrm{de}} = 100$

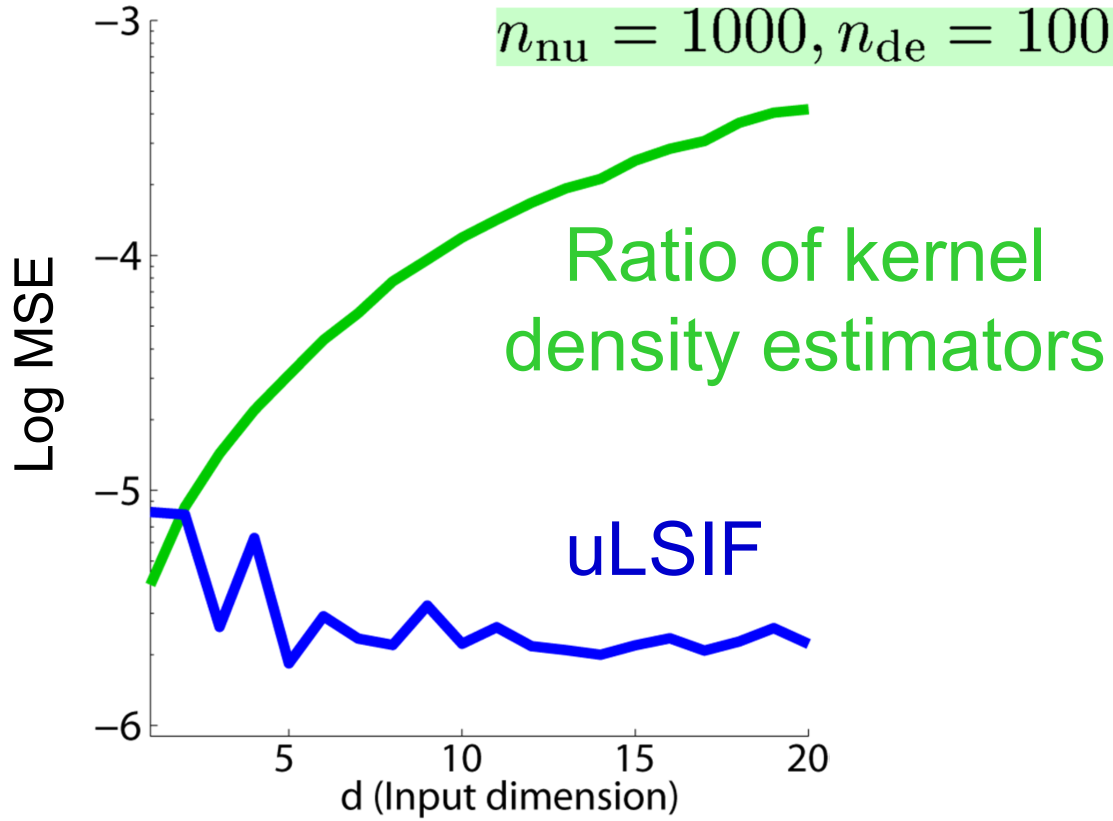

LSIF: Summary

- LS formulation is computationally efficient:
  - cLSIF: Regularization path tracking
  - uLSIF: Analytic solution and LOOCV
- Gives an accurate approximator of Pearson (PE) divergence (an *f*-divergence):

$$\int p_{\mathrm{de}}(\boldsymbol{x}) \left(\frac{p_{\mathrm{nu}}(\boldsymbol{x})}{q_{\mathrm{de}}(\boldsymbol{x})} - 1 \right)^2 \mathrm{d}\boldsymbol{x}$$

- Analytic solution of uLSIF allows us to compute the derivative of PE divergence approximator:
  - Useful in dimension reduction, independent component analysis, causal inference etc.

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
   A) Importance sampling
   B) Distribution comparison
   C) Mutual information estimation
   D) Conditional probability estimation
4. More on Density Ratio Estimation
5. Conclusions

Learning under Covariate Shift

- Covariate shift: Shimodaira (JSPI2000)
  - Training/test input distributions are different, but target function remains unchanged.
  - (Weak) extrapolation.

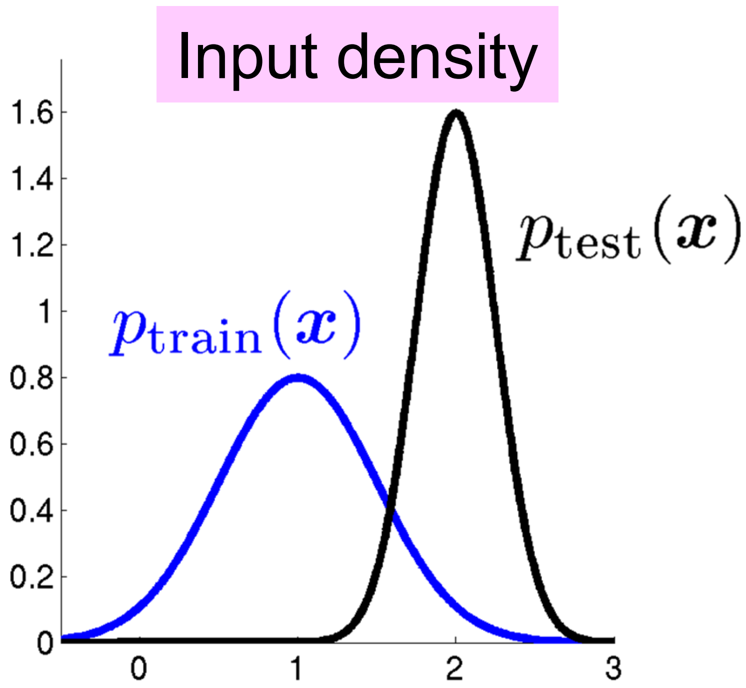

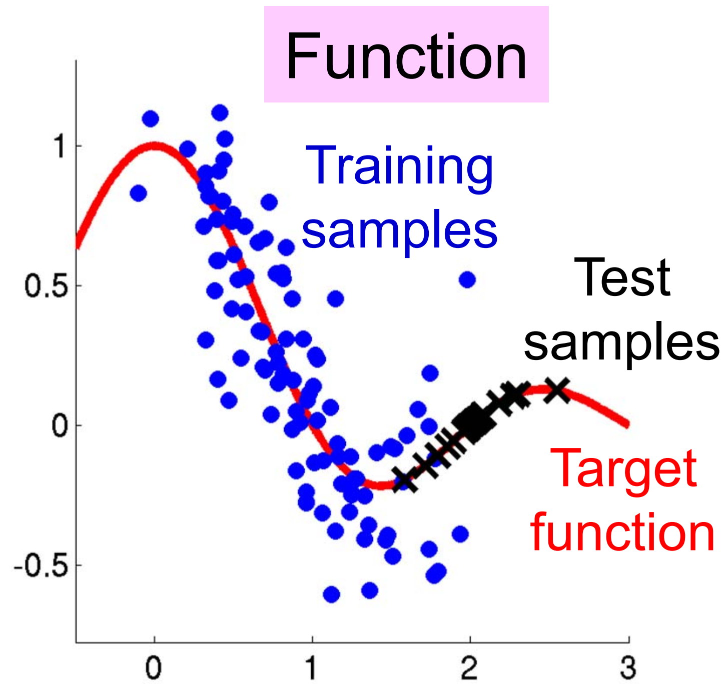

Ordinary Least-Squares (OLS)

$$\min_{\boldsymbol{\alpha}} \left[\sum_{i=1}^{n} \left(\widehat{f}(\boldsymbol{x}_i) - y_i \right)^2 \right]$$

- In standard setting, OLS is consistent, i.e., the learned function converges to the best solution when $n \to \infty$.
- Under covariate shift, OLS is no longer consistent.

$$\widehat{f}(x) = \alpha_1 + \alpha_2 x$$

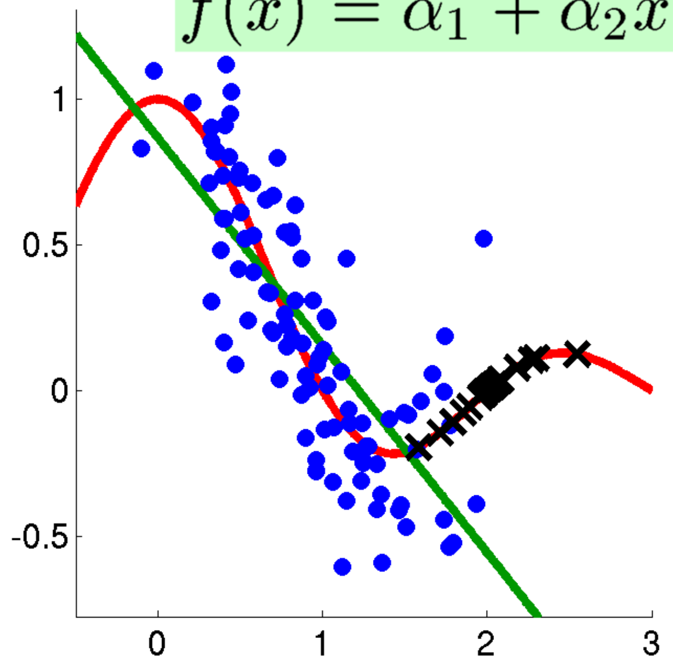

Law of Large Numbers

- Sample average converges to the population mean:

$$\frac{1}{n}\sum_{i=1}^{n}\mathrm{loss}(\boldsymbol{x}_i) \longrightarrow \int \mathrm{loss}(\boldsymbol{x})p_{train}(\boldsymbol{x})d\boldsymbol{x}$$

$$\boldsymbol{x}_i \overset{i.i.d.}{\sim} p_{train}(\boldsymbol{x})$$

- We want to estimate the expectation over test input points only using training input points $\{\boldsymbol{x}_i\}_{i=1}^{n}$.

$$\int \mathrm{loss}(\boldsymbol{x})p_{test}(\boldsymbol{x})d\boldsymbol{x}$$

Importance Weighting

- Importance : Ratio of test and training input densities

$$\frac{p_{test}(\boldsymbol{x})}{p_{train}(\boldsymbol{x})}$$

- Importance-weighted average:

$$\frac{1}{n}\sum_{i=1}^{n}\frac{p_{test}(\boldsymbol{x}_i)}{p_{train}(\boldsymbol{x}_i)}\mathrm{loss}(\boldsymbol{x}_i) \qquad \boldsymbol{x}_i \overset{i.i.d.}{\sim} p_{train}(\boldsymbol{x})$$

$$\longrightarrow \int \frac{p_{test}(\boldsymbol{x})}{p_{train}(\boldsymbol{x})}\mathrm{loss}(\boldsymbol{x})p_{train}(\boldsymbol{x})d\boldsymbol{x}$$

$$= \int \mathrm{loss}(\boldsymbol{x})p_{test}(\boldsymbol{x})d\boldsymbol{x}$$

Importance-Weighted Least-Squares

$$\min_{\boldsymbol{\alpha}} \left[\sum_{i=1}^{n} \frac{p_{test}(\boldsymbol{x}_i)}{p_{train}(\boldsymbol{x}_i)} \left(\widehat{f}(\boldsymbol{x}_i) - y_i \right)^2 \right]$$

- IWLS is consistent even under covariate shift.
- The idea is applicable to any likelihood-based methods!
  - Support vector machine, logistic regression, conditional random field, etc.

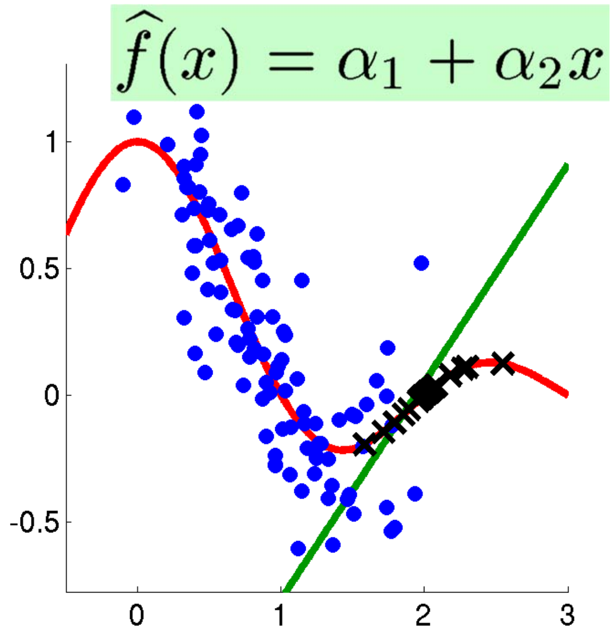

Model Selection

- Controlling bias-variance trade-off is important.
  - No weighting: low-variance but high-bias
  - Importance weighting: low-bias but high-variance

- "Flattened"-IWLS:

$$\min_{\boldsymbol{\alpha}} \left[\sum_{i=1}^{n} \left(\frac{p_{test}(\boldsymbol{x}_i)}{p_{train}(\boldsymbol{x}_i)} \right)^{\lambda} \left(\widehat{f}(\boldsymbol{x}_i) - y_i \right)^2 \right]$$

Shimodaira (JSPI2000)

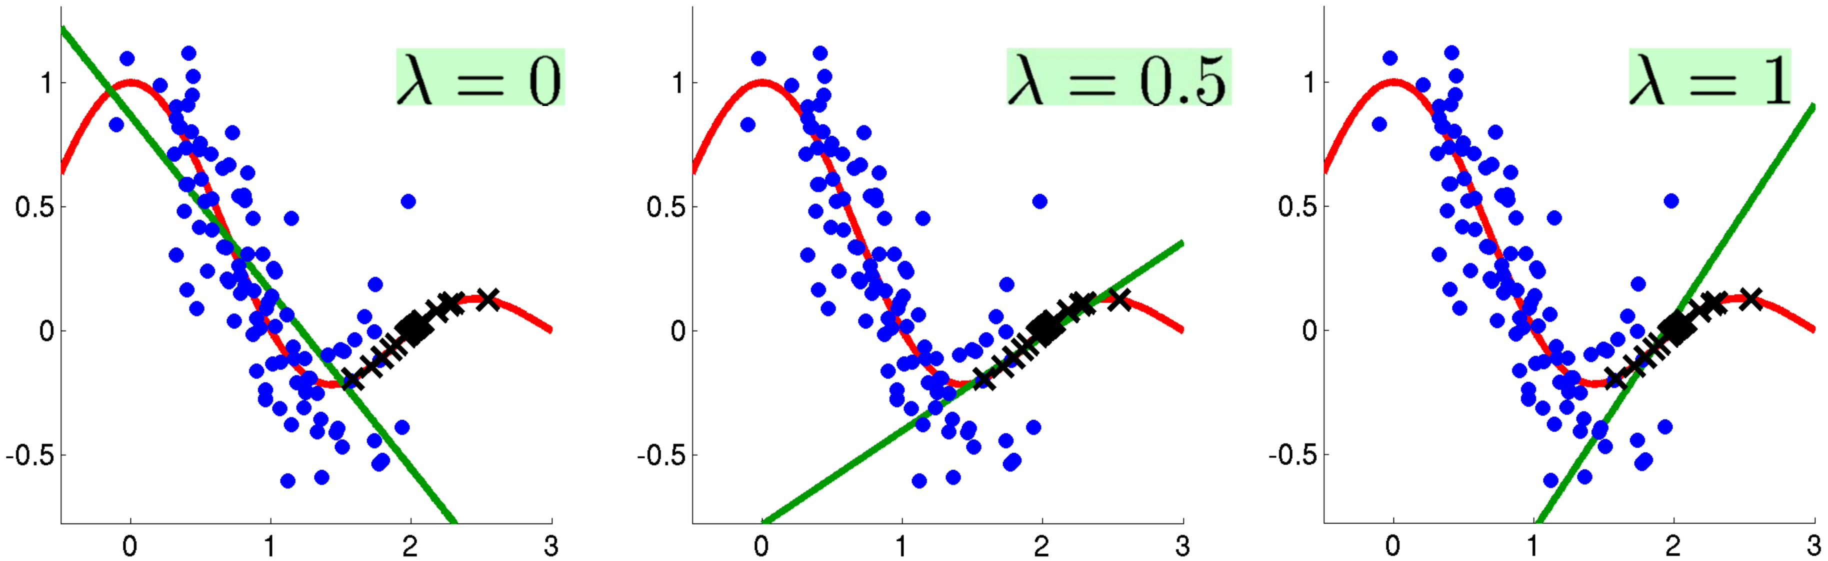

Model Selection

- Importance weighting also plays a central role for unbiased model selection:
  - Akaike information criterion (regular models)
    Shimodaira (JSPI2000)
  - Subspace information criterion (linear models)
    MS & Müller (Stat&Dec.2005)
  - Cross-validation (arbitrary models)
    MS, Krauledat & Müller (JMLR2007)

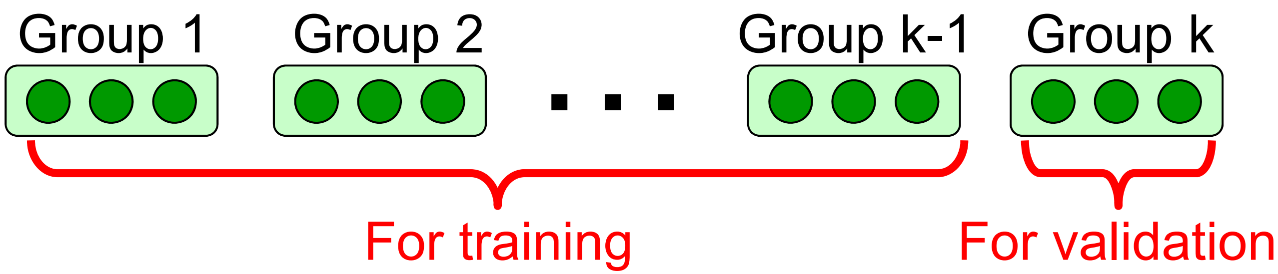

$$\widehat{f}(\boldsymbol{x})$$

$$\frac{p_{test}(\boldsymbol{x}')}{p_{train}(\boldsymbol{x}')}\left(\widehat{f}(\boldsymbol{x}')-y'\right)^2$$

Experiments: Speaker Identification

Yamada, MS & Matsui (SigPro2010)

- NTT Japanese speech dataset. Matsui & Furui (ICASSP1993)
- Text-independent speaker identification accuracy for 10 male speakers.
- Kernel logistic regression (KLR) with sequence kernel.

Training data	Speech length	IWKLR+IWCV+KLIEP	KLR+CV
9 months before	1.5 [sec]	91.0 %	88.2 %
	3.0 [sec]	95.0 %	92.9 %
	4.5 [sec]	97.7 %	96.1 %
6 months before	1.5 [sec]	91.0 %	87.7 %
	3.0 [sec]	95.3 %	91.1 %
	4.5 [sec]	97.4 %	93.4 %
3 months before	1.5 [sec]	94.8 %	91.7 %
	3.0 [sec]	97.9 %	96.3 %
	4.5 [sec]	98.8 %	98.3 %

Experiments: Text Segmentation

Tsuboi, Kashima, Hido, Bickel & MS (JIP2009)

こんな失敗はご愛敬だよ.
→ こんな／失敗／は／ご／愛敬／だ／よ／.

- Japanese word segmentation dataset.

  Tsuboi, Kashima, Mori, Oda & Matsumoto (COLING2008)

- Adaptation from daily conversation to medical domain.
- Segmentation by conditional random field (CRF).

	IWCRF+IWCV +KLIEP	CRF+CV	CRF+CV (use additional test labels)
F-measure (larger is better)	94.46	92.30	94.43

Semi-supervised adaptation with importance weighting is comparable to supervised adaptation!

Other Applications

- Age prediction from faces:
  - Illumination change

    Ueki, MS & Ihara (ICPR2010)

- Brain-computer interface:
  - Mental condition change

    MS, Krauledat & Müller (JMLR2007)
    Li, Kambara, Koike & MS (IEEE-TBME2010)

- Robot control:
  - Efficient sample reuse

    Hachiya, Akiyama, MS & Peters (NN2009)
    Hachiya, Peters & MS (NeCo2011)

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
   - A) Importance sampling
   - B) Distribution comparison
   - C) Mutual information estimation
   - D) Conditional probability estimation
4. More on Density Ratio Estimation
5. Conclusions

Inlier-Based Outlier Detection

Hido, Tsuboi, Kashima, MS & Kanamori (ICDM2008, KAIS2011)
Smola, Song & Teo (AISTATS2009)

- Goal: Given a set of inlier samples, find outliers in a test set (if exist)

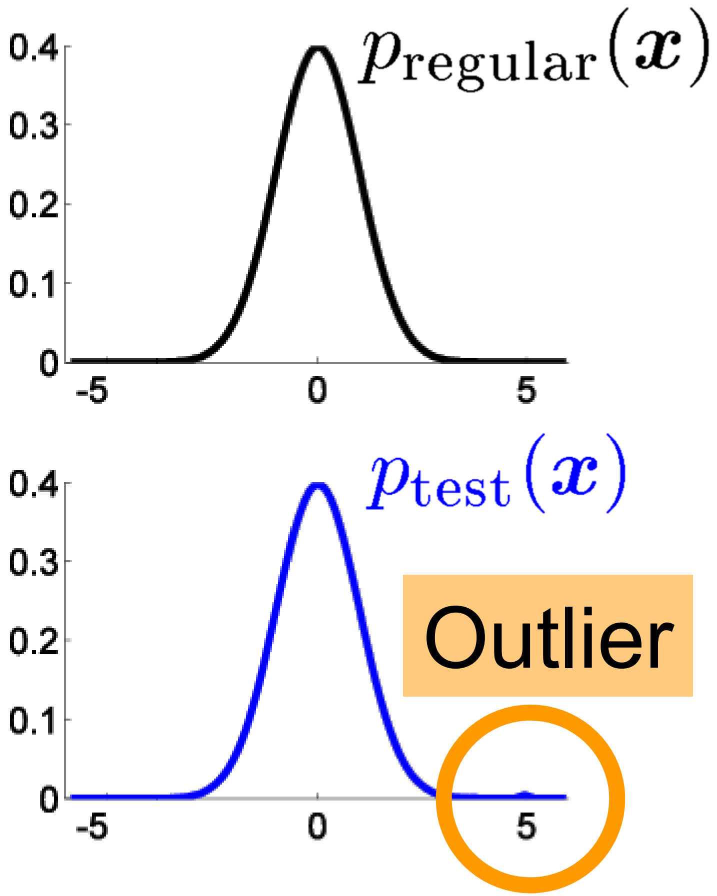

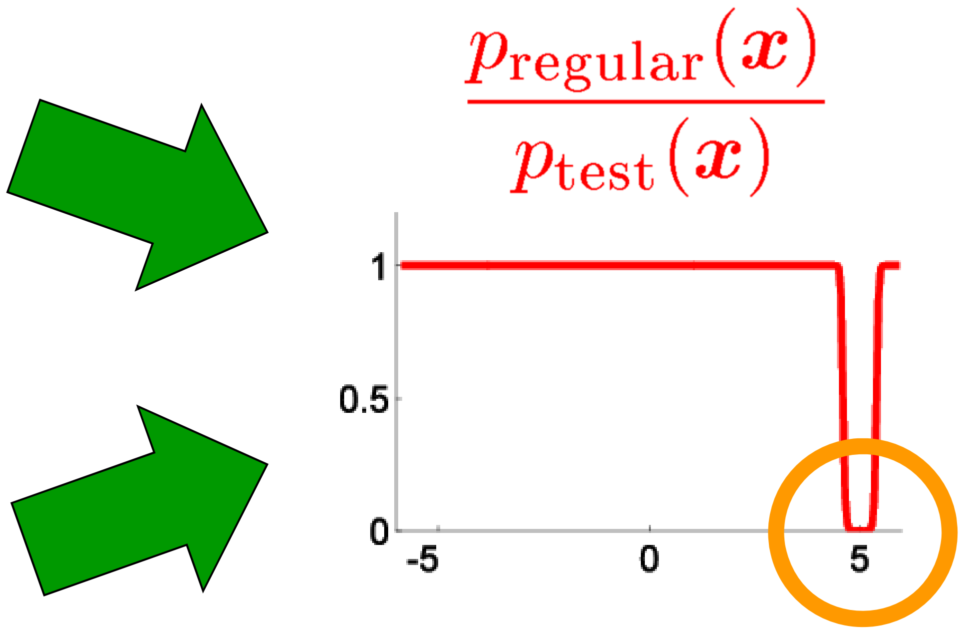

Tuning parameters can be optimized in terms of ratio approximation error

Experiments

Hido, Tsuboi, Kashima, MS & Kanamori (ICDM2008, KAIS2011)

- Top10 outliers in the USPS test dataset found based on the USPS training dataset.

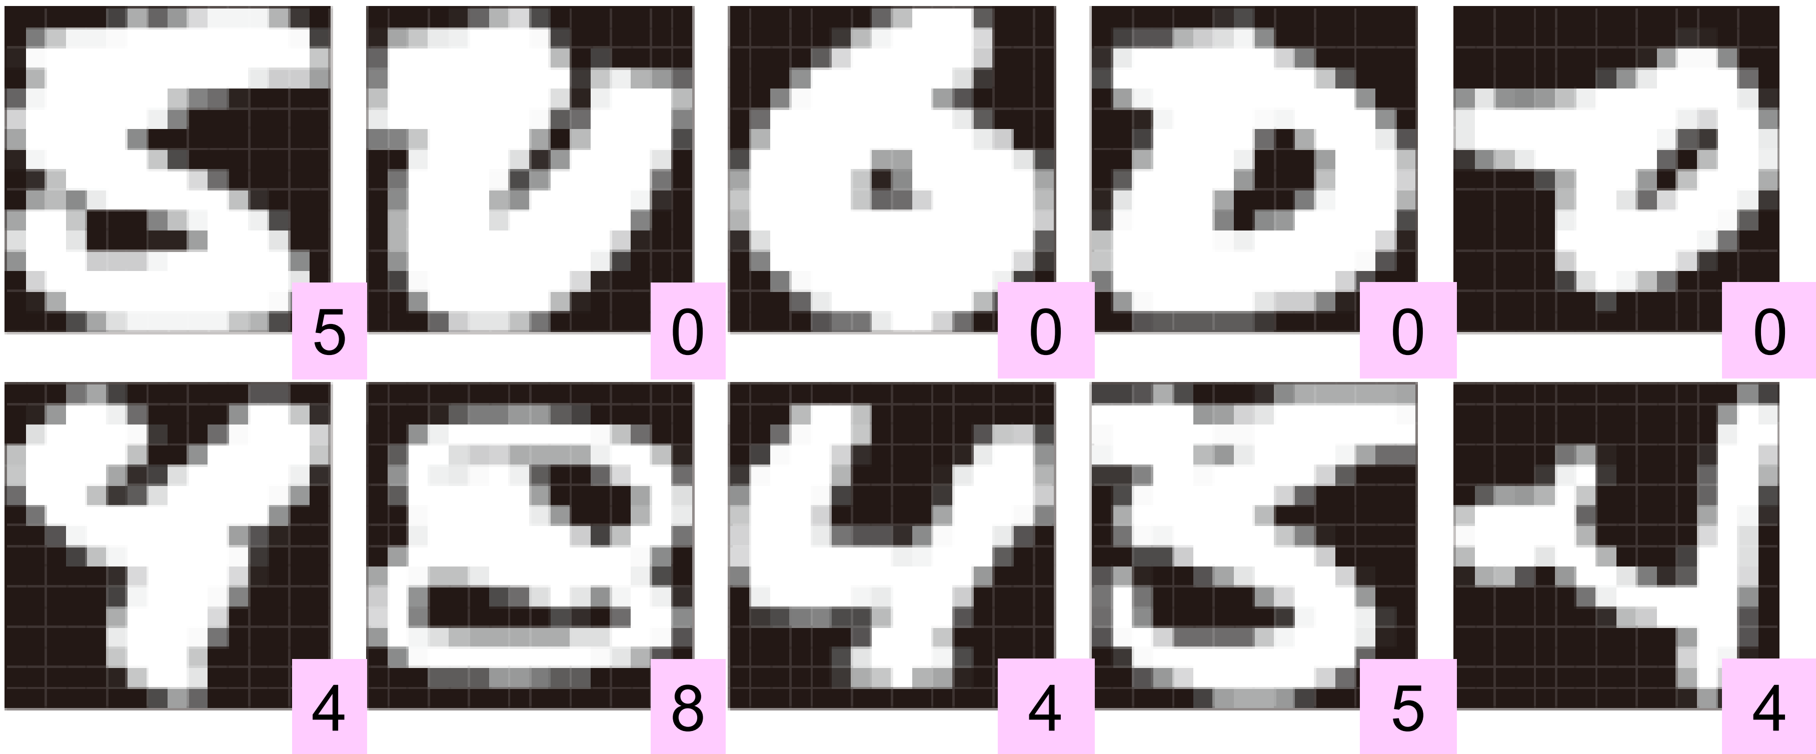

Most of them are not readable even by human.

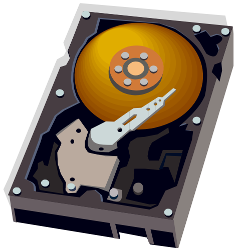

Failure Prediction in Hard-Disk Drives

- Self-Monitoring And Reporting Technology (SMART):

Murray, Hughes & Kreutz-Delgado (JMLR 2005)

	Least-squares density ratio	One-class SVM	Local outlier factor #NN=5	Local outlier factor #NN=30
AUC (larger is better)	0.881	0.843	0.847	0.924
Comp. time	1	26.98	65.31	

- LOF works well, given #NN is set appropriately. But there is no objective model selection method.
- Density ratio method can use cross-validation for model selection, and is computationally efficient.

OSVM: Schölkopf, Platt, Shawe-Taylor, Smola & Williamson (NeCo2001)
LOF: Breunig, Kriegel, Ng & Sander (SIGMOD2000)

Other Applications

- Steel plant diagnosis — Hirata, Kawahara & MS (Patent2011)
- Printer roller quality control — Takimoto, Matsugu & MS (DMSS2009)
- Loan customer inspection — Hido, Tsuboi, Kashima, MS & Kanamori (KAIS2011)
- Sleep therapy — Kawahara & MS (SADM2012)

Divergence Estimation

Nguyen, Wainwright & Jordan (IEEE-IT2010)
MS, Suzuki, Ito, Kanamori & Kimura (NN2011)

- **Goal**: Estimate a divergence functional from

$$\{\boldsymbol{x}_i\}_{i=1}^{n} \overset{i.i.d.}{\sim} p(\boldsymbol{x}) \qquad \{\boldsymbol{x}'_j\}_{j=1}^{n'} \overset{i.i.d.}{\sim} p'(\boldsymbol{x})$$

  - Kullback-Leibler divergence: $\displaystyle\int p(\boldsymbol{x}) \log \frac{p(\boldsymbol{x})}{p'(\boldsymbol{x})} \mathrm{d}\boldsymbol{x}$
  - Pearson divergence: $\displaystyle\int p'(\boldsymbol{x}) \left(\frac{p(\boldsymbol{x})}{p'(\boldsymbol{x})} - 1 \right)^2 \mathrm{d}\boldsymbol{x}$

    (an *f*-divergence)

- Use density ratio estimation: $\displaystyle r(\boldsymbol{x}) = \frac{p(\boldsymbol{x})}{p'(\boldsymbol{x})}$

Real-World Applications

- Regions-of-interest detection in images:

Yamanaka, Matsugu & MS (IEEJ2011)

- Event detection in movies:

Matsugu, Yamanaka & MS (VECTaR2011)

- Event detection from Twitter data:

Liu, Yamada, Collier & MS (arXiv2012)

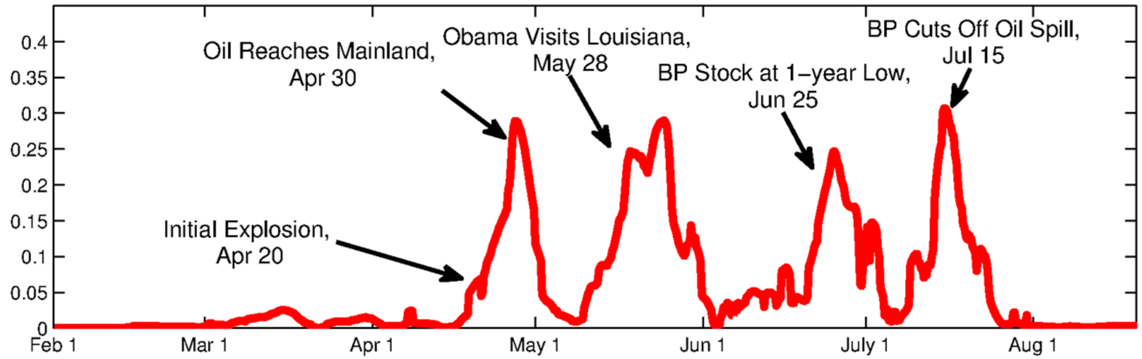

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
   - A) Importance sampling
   - B) Distribution comparison
   - C) Mutual information estimation
   - D) Conditional probability estimation
4. More on Density Ratio Estimation
5. Conclusions

Mutual Information Estimation

Suzuki, MS, Sese & Kanamori (FSDM2008)

- Mutual information (MI): Shannon (1948)

$$\mathrm{MI} = \iint p(\boldsymbol{x}, \boldsymbol{y}) \log \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})p(\boldsymbol{y})} \mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{y}$$

- MI works as an independence measure:

$\mathrm{MI} = 0$ $\Longleftrightarrow$ $\boldsymbol{x}$ and $\boldsymbol{y}$ are statistically independent

- Use KL-based density ratio estimation (KLIEP):

$$r(\boldsymbol{x}, \boldsymbol{y}) = \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})p(\boldsymbol{y})}$$

Experiments: Methods Compared

- KL-based density ratio method.
- Kernel density estimation (KDE).
- K-nearest neighbor density estimation (KNN). Kraskov, Stögbauer & Grassberger (PRE2004)
  - The number of NNs is a tuning parameter.
- Edgeworth expansion density estimation (EDGE). van Hulle (NeCo2005)

Datasets for Evaluation

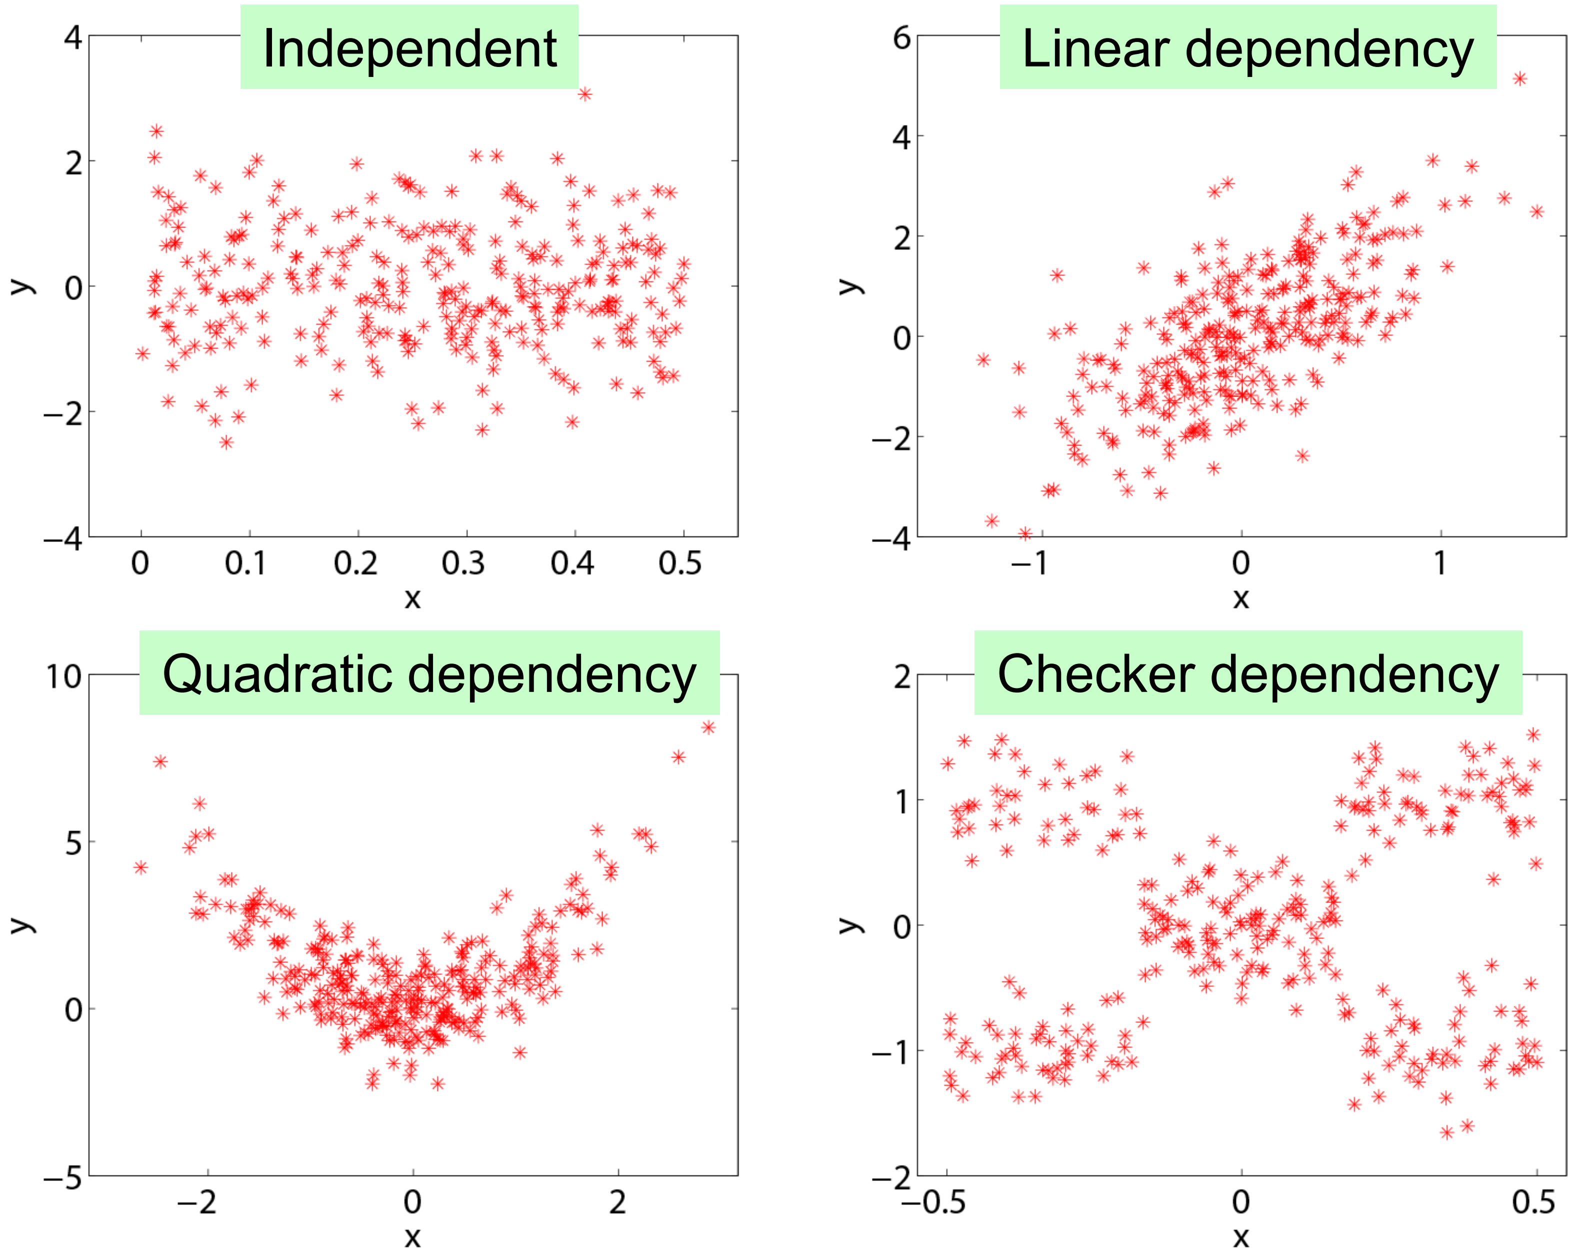

MI Approximation Error

$$\mathrm{Error} = |\widehat{\mathrm{MI}} - \mathrm{MI}|$$

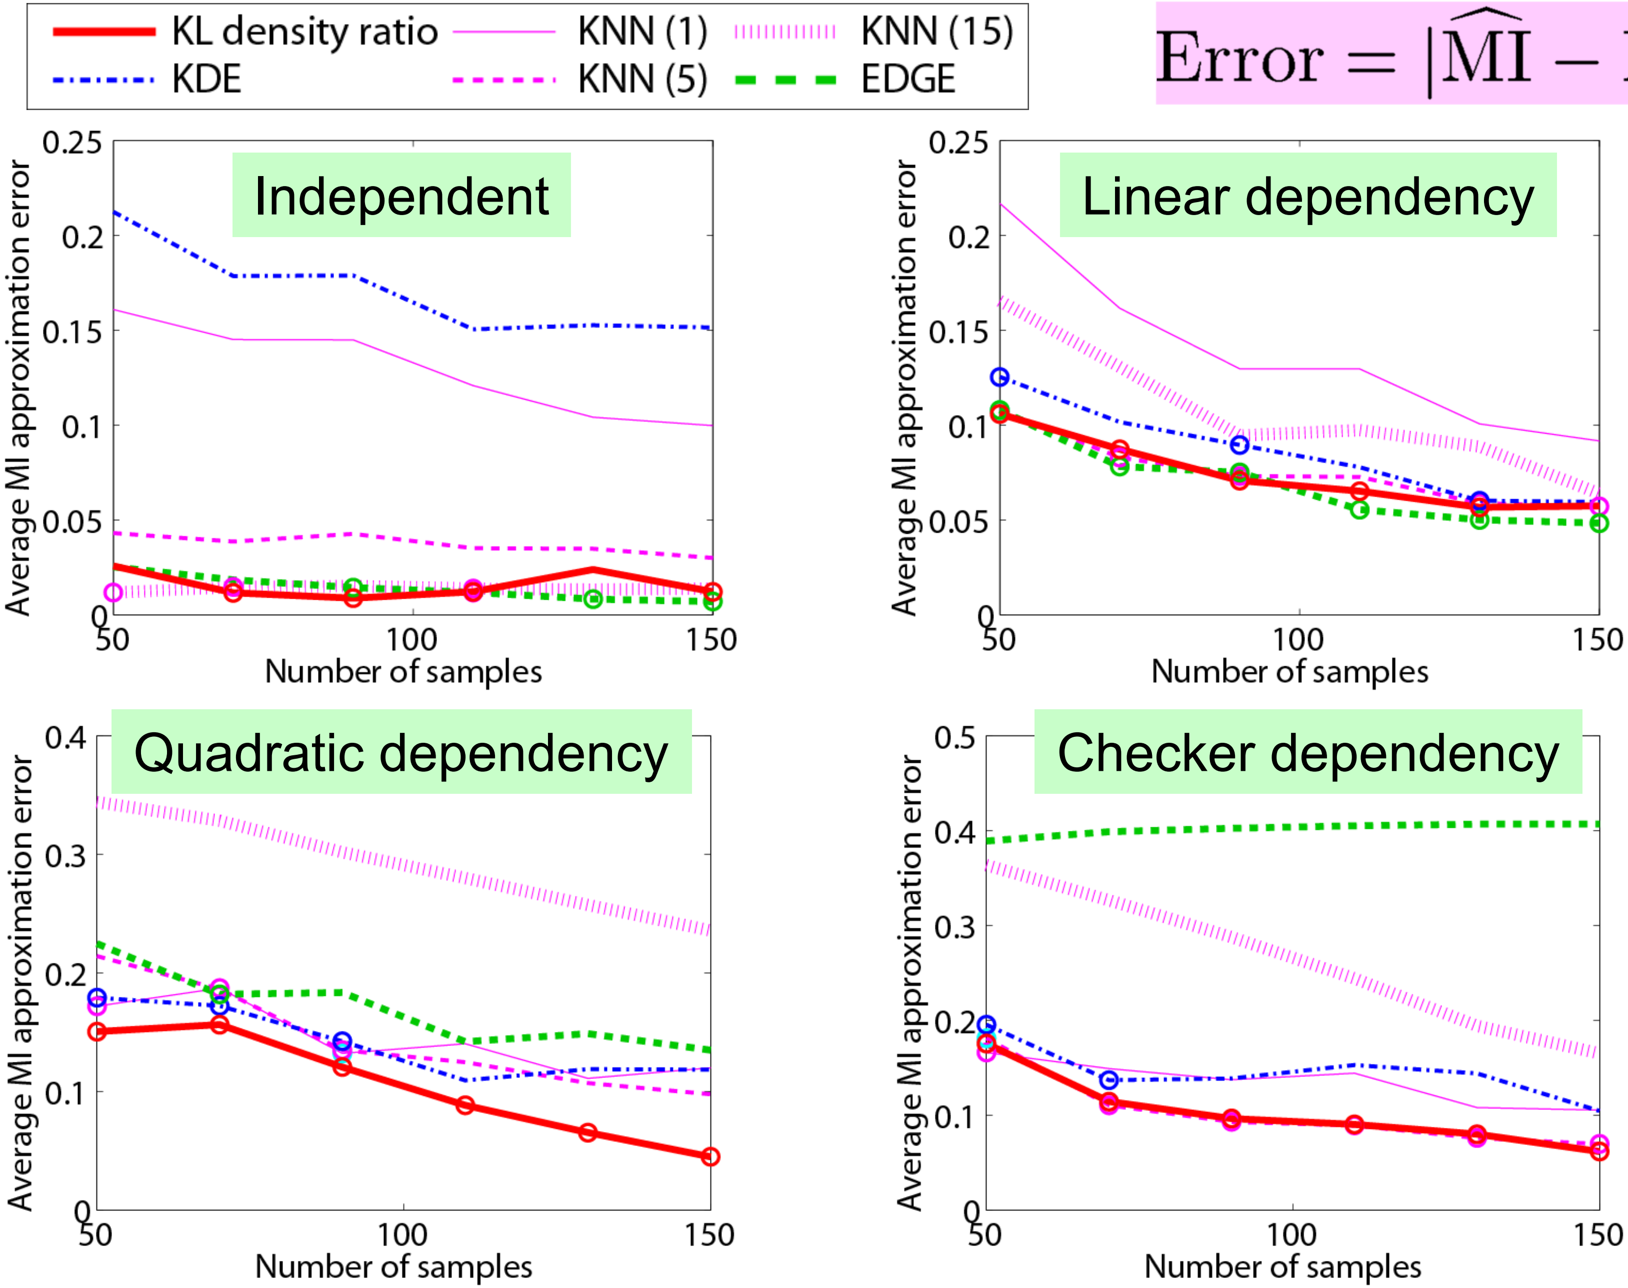

Estimation of Squared-Loss Mutual Information (SMI)

Suzuki, MS, Sese & Kanamori (BMC Bioinfo. 2009)

- Ordinary MI is based on the KL-divergence.
- SMI is based on the Pearson divergence:

$$\mathrm{SMI} = \iint p(\boldsymbol{x})p(\boldsymbol{y}) \left(\frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})p(\boldsymbol{y})} - 1 \right)^2 \mathrm{d}\boldsymbol{x}\mathrm{d}\boldsymbol{y}$$

  - Can also be used as an independence measure.
  - Can be approximated analytically and efficiently by least-squares density ratio estimation (uLSIF).

Usage of SMI Estimator

- Between input and output:
  - Feature ranking — Suzuki, MS, Sese & Kanamori (BMCBioinfo 2009)
  - Sufficient dimension reduction — Suzuki & MS (NeCo2012)
  - Clustering — MS, Yamada, Kimura & Hachiya (ICML2011) Kimura & MS (JACIII2011)
- Between inputs:
  - Independent component analysis — Suzuki & MS (NeCo2010)
  - Object matching — Yamada & MS (AISTATS2011)
  - Canonical dependency analysis — Karasuyama & MS (NN2012)
- Between input and residual:
  - Causal inference

    Yamada & MS (AAAI2010)

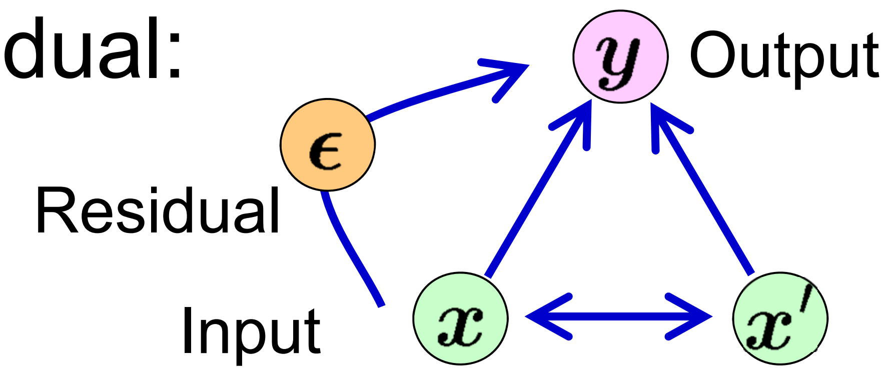

Sufficient Dimension Reduction

Li (JASA1991)

- Input: $\boldsymbol{x}$
- Output: $\boldsymbol{y}$
- Projected input: $\boldsymbol{z} = \boldsymbol{W}\boldsymbol{x}$

$$\boldsymbol{z} = \boldsymbol{W} \boldsymbol{x}$$

$$\boldsymbol{W}\boldsymbol{W}^\top = \boldsymbol{I}$$

- Goal: Find $\boldsymbol{W}$ so that $\boldsymbol{z}$ contains all information on $\boldsymbol{y}$, i.e., $\boldsymbol{y} \perp\!\!\!\perp \boldsymbol{x} \mid \boldsymbol{z}$
  - In terms of SMI: Suzuki & MS (NeCo2012)

$$\boldsymbol{y} \perp\!\!\!\perp \boldsymbol{x} \mid \boldsymbol{z} \quad \Longleftrightarrow \quad \max_{\boldsymbol{W}} \ \mathrm{SMI}(\boldsymbol{W}\boldsymbol{x}, \boldsymbol{y})$$

Sufficient Dimension Reduction via SMI Maximization

- Let's solve $\max_{\boldsymbol{W}} \ \widehat{\mathrm{SMI}}(\boldsymbol{W}) \text{ subject to } \boldsymbol{W}\boldsymbol{W}^\top = \boldsymbol{I}$.

$$\widehat{\mathrm{SMI}}(\boldsymbol{W}) = 2\widehat{\boldsymbol{h}}^\top \widehat{\boldsymbol{\alpha}} - \widehat{\boldsymbol{\alpha}}^\top \widehat{\boldsymbol{H}} \widehat{\boldsymbol{\alpha}} - 1$$

$\widehat{\boldsymbol{\alpha}}$: uLSIF solution

- Since $\boldsymbol{W}$ is on a Grassmann manifold, natural gradient gives the steepest direction:

Amari (NeCo1998)

$$\boldsymbol{W} \longleftarrow \boldsymbol{W} + \epsilon \frac{\partial \widehat{\mathrm{SMI}}}{\partial \boldsymbol{W}} \left(\boldsymbol{I} - \boldsymbol{W}^\top \boldsymbol{W} \right)$$

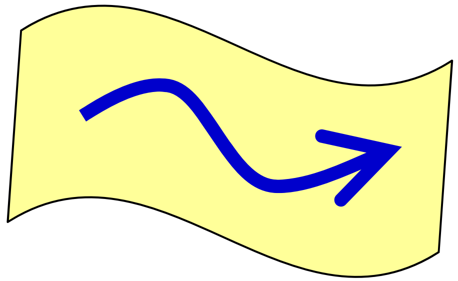

- A computationally efficient heuristic update is also available.

Yamada, Niu, Takagi & MS (ACML2011)

Experiments

Yamada, Niu, Takagi & MS (ACML2011)

- Dimension reduction for multi-label data:

Pascal VOC 2010 image classification

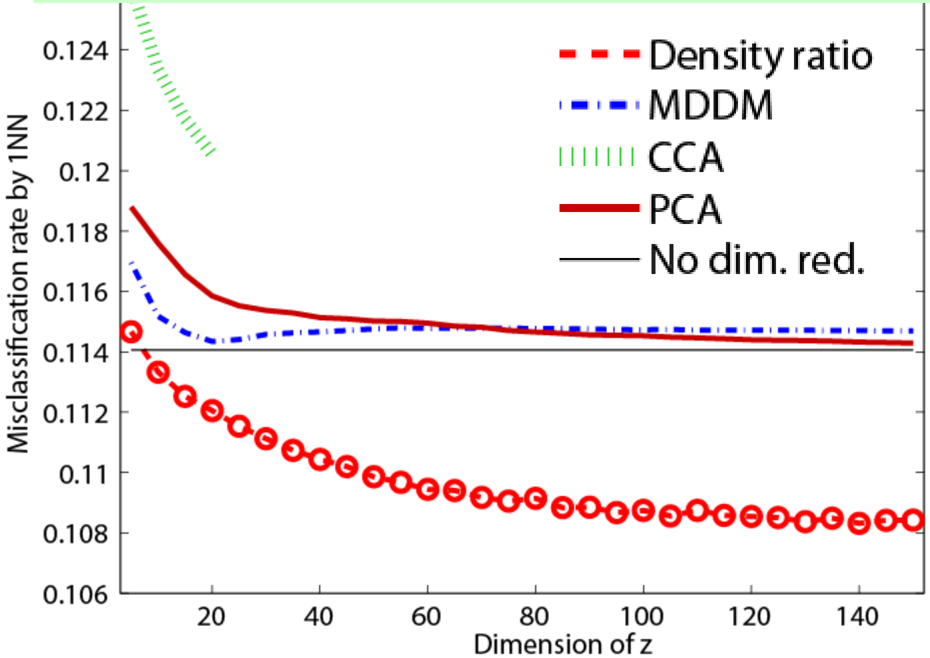

Freesound audio tagging

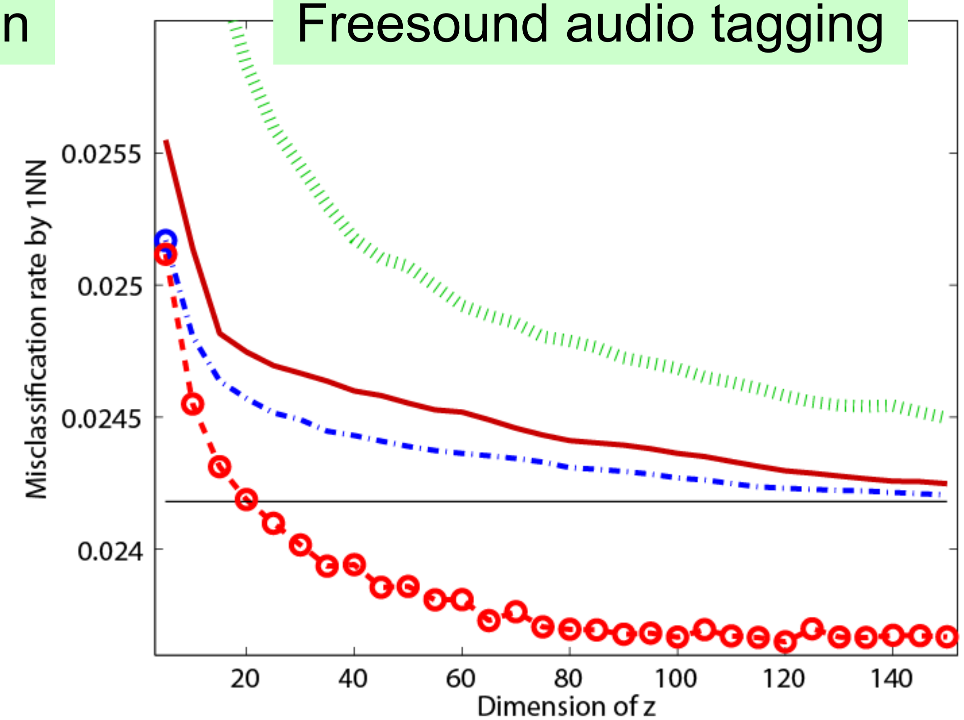

- MDDM: Multi-label dimensionality reduction via dependence maximization (MDDM) Zhang & Zhou (ACM-TKDD2010)
- CCA: Canonical correlation analysis
- PCA: Principal component analysis

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
   - A) Importance sampling
   - B) Distribution comparison
   - C) Mutual information estimation
   - D) Conditional probability estimation
4. More on Density Ratio Estimation
5. Conclusions

Conditional Density Estimation

$$p(\boldsymbol{y}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})}$$

MS, Takeuchi, Suzuki, Kanamori,
Hachiya & Okanohara (IEICE-ED2010)

- Regression = Conditional mean estimation
- However, regression is not informative enough for complex data analysis:
  - Multi-modality
  - Asymmetry
  - Hetero-scedasticity
- Directly estimation of conditional density via density-ratio estimation.

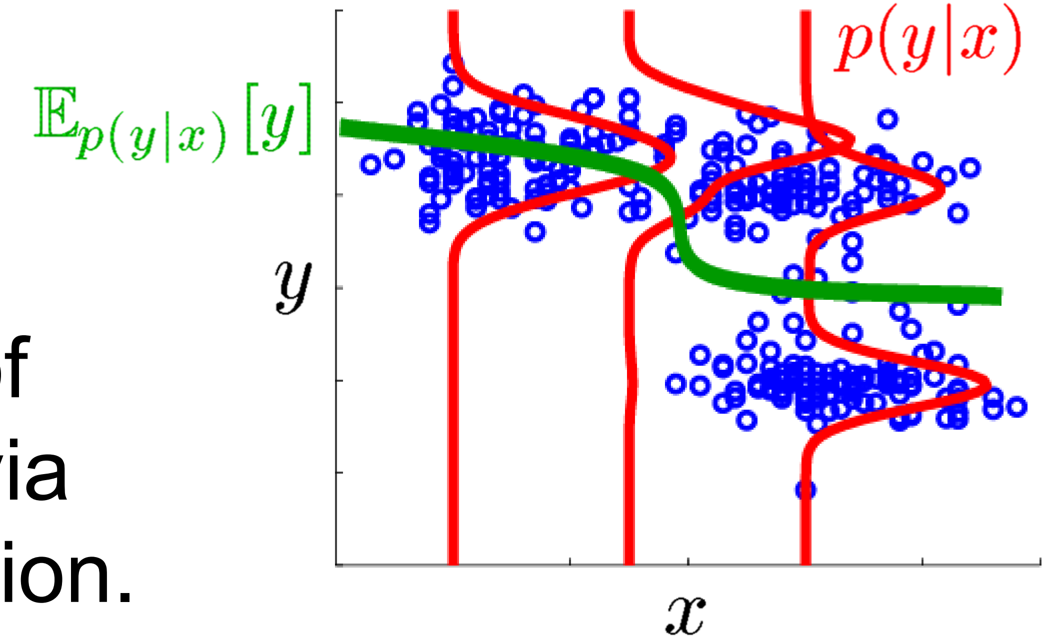

Experiments: Transition Estimation for Mobile Robot

- Transition probability $p(s'|s,a)$: Probability of being at state s' when action a is taken at s.

Khepera robot
- State: Infrared sensors
- Action: Wheel speed

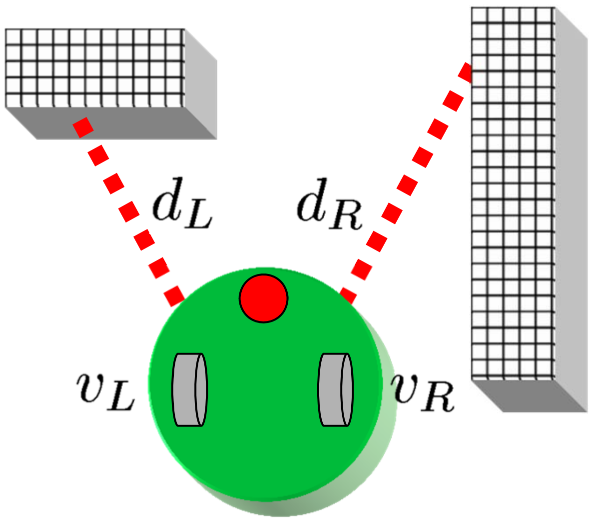

Mean (std.) test negative log-likelihood over 10 runs (smaller is better) (red: comparable by 5% t-test)

Data	uLSIF	ε-KDE	MDN
Khepera1	1.69(0.01)	2.07(0.02)	1.90(0.36)
Khepera2	1.86(0,01)	2.10(0.01)	1.92(0.26)
Pendulum1	1.27(0.05)	2.04(0.10)	1.44(0.67)
Pendulum2	1.38(0.05)	2.07(0.10)	1.43(0.58)
Comp. Time	**1**	**0.164**	**1134**

ε-KDE: ε-neighbor kernel density estimation
MDN: Mixture density network Bishop (Book2006)

Probabilistic Classification

$$p(\boldsymbol{y}|\boldsymbol{x}) = \frac{p(\boldsymbol{x}, \boldsymbol{y})}{p(\boldsymbol{x})}$$

MS (IEICE-ED2010)

- If y is categorical, conditional probability estimation corresponds to learning class-posterior probability.
- Least-squares density ratio estimation (uLSIF) provides an analytic estimator:
  - Computationally efficient alternative to kernel logistic regression.
  - No normalization term included.
  - Classwise training is possible.

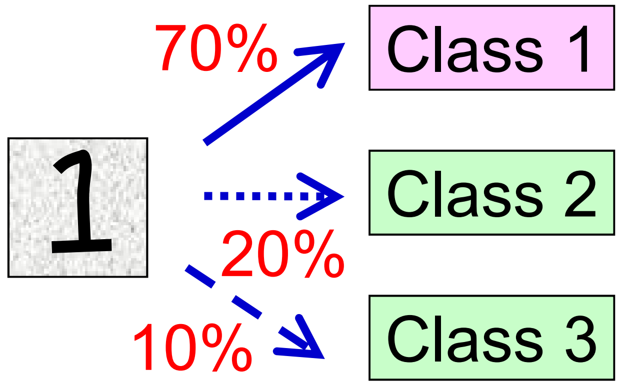

Numerical Example

- Letter dataset (26 classes):

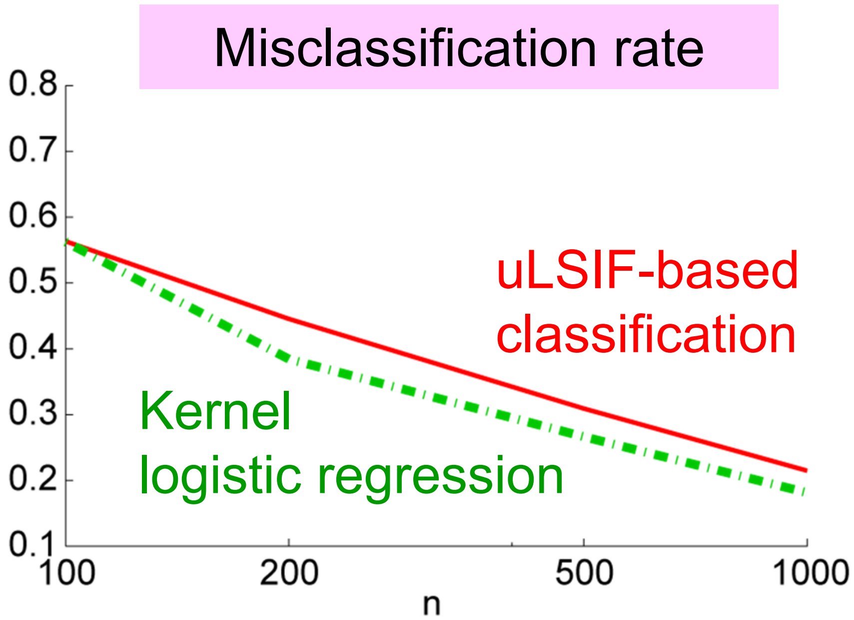

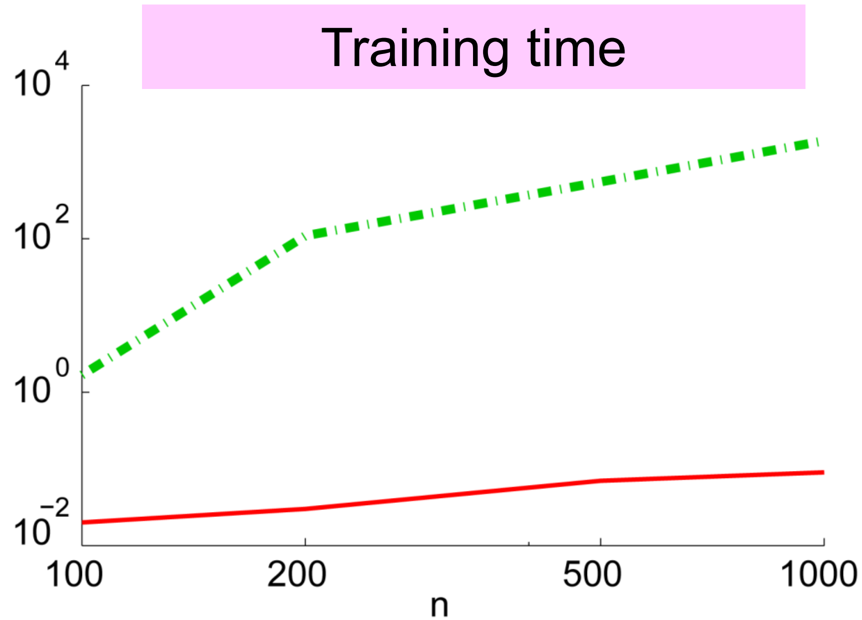

- uLSIF-based classification method:
  - Comparable accuracy with KLR.
  - Training is 1000 times faster!

More Experiments

Yamada, MS, Wichern & Simm (IEICE2011)

Dataset	uLSIF	KLR
Aeroplane	82.6(1.0)	83.0(1.3)
Bicycle	77.7(1.7)	76.6(3.4)
Bird	68.7(2.0)	70.8(2.2)
Boat	74.4(2.0)	72.8(2.6)
Bottle	65.4(1.8)	62.1(4.3)
Bus	85.4(1.4)	85.6(1.4)
Car	73.0(0.8)	72.1(1.2)
Cat	73.6(1.4)	74.1(1.7)
Chair	71.0(1.0)	70.5(1.0)
Cow	71.7(3.2)	69.3(3.6)
Diningtable	75.0(1.6)	71.4(2.7)
Dog	69.6(1.0)	69.4(1.8)
Horse	64.4(2.5)	61.2(3.2)
Motorbike	77.0(1.7)	75.9(3.3)
Person	67.6(0.9)	67.0(0.8)
Pottedplant	66.2(2.6)	61.9(3.2)
Sheep	77.8(1.6)	74.0(3.8)
Sofa	67.4(2.7)	65.4(4.6
Train	79.2(1.3)	78.4(3.0)
Tvmonitor	76.7(2.2)	76.6(2.3)
Training time [sec]	**0.7**	**24.6**

Pascal VOC 2010
image classification:
Mean AUC (std) over 50 runs
(red: comparable by 5% t-test)

Freesound audio tagging:
Mean AUC (std) over 50 runs

	uLSIF	KLR
AUC	70.1(9.6)	66.7(10.3)
Training time [sec]	0.005	0.612

Other Applications

- Action recognition from accelerometer

  Hachiya, MS & Ueda (Neurocomputing2011)

- Age prediction from faces

  Ueki, MS, Ihara & Fujita (ACPR2011)

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
4. More on Density Ratio Estimation
   - A) Unified Framework
   - B) Dimensionality Reduction
   - C) Relative Density Ratios
5. Conclusions

Bregman (BR) Divergence

Bregman (1967)

- f : Differentiable convex function
- BR divergence with function f :

$$\mathrm{BR}_f(t^* \| t) := f(t^*) - \underbrace{f(t) - \partial f(t)(t^* - t)}_{\text{Linear prediction from } f(t) \text{ to } f(t^*)}$$

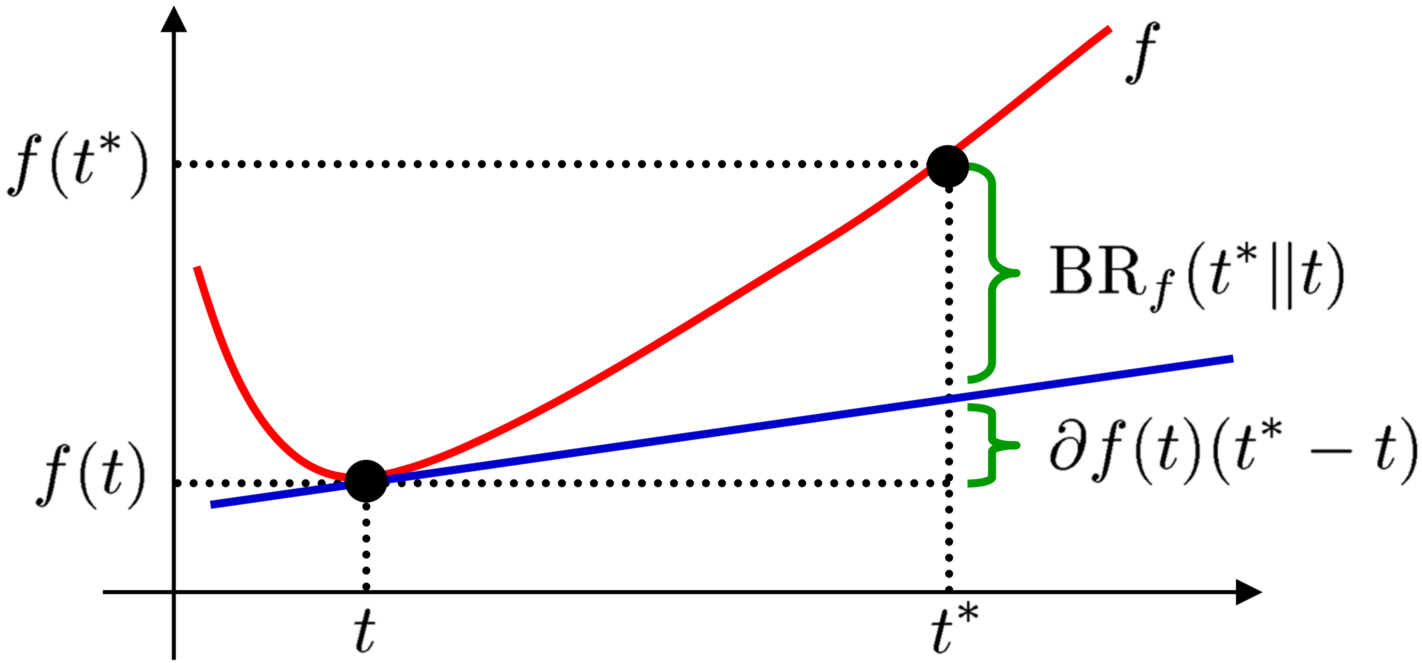

Density-Ratio Fitting under BR Divergence

MS, Suzuki & Kanamori (AISM2012)

- Fit a ratio model $\widehat{r}(\boldsymbol{x})$ to true ratio $r(\boldsymbol{x})$ under the BR divergence:

$$\min_{\widehat{r}} \mathrm{BR}_f\left(\widehat{r}\right)$$

$$\mathrm{BR}_f\left(\widehat{r}\right) = \int p_{\mathrm{de}}(\boldsymbol{x}) \nabla f(\widehat{r}(\boldsymbol{x}))\widehat{r}(\boldsymbol{x}) d\boldsymbol{x} - \int p_{\mathrm{de}}(\boldsymbol{x}) f(\widehat{r}(\boldsymbol{x})) d\boldsymbol{x}$$

$$- \int p_{\mathrm{nu}}(\boldsymbol{x}) \nabla f(\widehat{r}(\boldsymbol{x})) d\boldsymbol{x} + C$$

$$\approx \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \nabla f(\widehat{r}(\boldsymbol{x}_j^{\mathrm{de}}))\widehat{r}(\boldsymbol{x}_j^{\mathrm{de}}) - \frac{1}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} f(\widehat{r}(\boldsymbol{x}_j^{\mathrm{de}}))$$

$$- \frac{1}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \nabla f(\widehat{r}(\boldsymbol{x}_i^{\mathrm{nu}})) + C$$

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

Unified View

- Logistic regression:

$$f(t) = t\log t - (1+t)\log(1+t)$$

- (Extended) kernel mean matching:

$$f(t) = (t-1)^2/2 \qquad \min_{\widehat{r}} \|\nabla J(\widehat{r})\|^2$$

- KL-based method:

$$f(t) = t\log t - t$$

- uLSIF:

$$f(t) = (t-1)^2/2 \qquad \min_{\widehat{r}} J(\widehat{r})$$

- Robust estimator (power divergence):

$$f(t) = \alpha^{-1}(t^{1+\alpha} - t) \quad \alpha > 0$$

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
4. More on Density Ratio Estimation
   - A) Unified Framework
   - B) Dimensionality Reduction
   - C) Relative Density Ratios
5. Conclusions

Direct Density-Ratio Estimation with Dimensionality Reduction (D^3)

- Directly density-ratio estimation without density estimation is promising.
- However, for high-dimensional data, density-ratio estimation is still challenging.
- We combine direct density-ratio estimation with dimensionality reduction!

Hetero-distributional Subspace (HS)

MS, Kawanabe & Chui (NN2010)

- Key assumption: $p_{\mathrm{nu}}(\boldsymbol{x})$ and $p_{\mathrm{de}}(\boldsymbol{x})$ are different only in a subspace (called HS).

$$r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

$$\begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} = \begin{bmatrix} \boldsymbol{U} \\ \boldsymbol{V} \end{bmatrix} \boldsymbol{x}$$

$$= \frac{p(\boldsymbol{v}|\boldsymbol{u})p_{\mathrm{nu}}(\boldsymbol{u})}{p(\boldsymbol{v}|\boldsymbol{u})p_{\mathrm{de}}(\boldsymbol{u})} = \frac{p_{\mathrm{nu}}(\boldsymbol{u})}{p_{\mathrm{de}}(\boldsymbol{u})}$$

$\begin{bmatrix} \boldsymbol{U} \\ \boldsymbol{V} \end{bmatrix}$: Full-rank and orthogonal

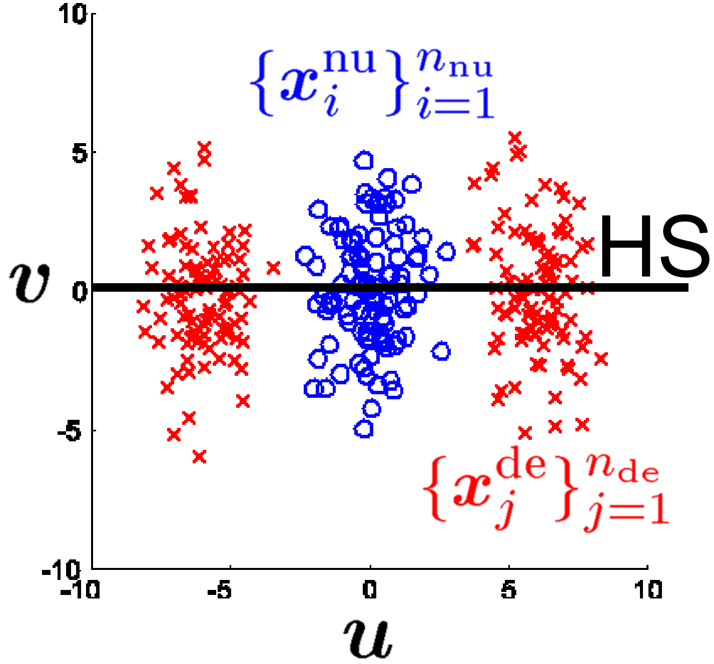

- This allows us to estimate the density ratio only within the low-dimensional HS!

Characterization of HS

MS, Yamada, von Bünau, Suzuki, Kanamori & Kawanabe (NN2011)

- HS is given as the maximizer of the Pearson divergence with respect to $\boldsymbol{U}$:

$$\mathrm{PE}[p_{\mathrm{nu}}(\boldsymbol{u}), p_{\mathrm{de}}(\boldsymbol{u})] = \int \left(\frac{p_{\mathrm{nu}}(\boldsymbol{u})}{p_{\mathrm{de}}(\boldsymbol{u})} - 1 \right)^2 p_{\mathrm{de}}(\boldsymbol{u}) \mathrm{d}\boldsymbol{u}$$

- PE can be analytically approximated by uLSIF (with good convergence property).
- HS search by
  - Natural gradient
  - A heuristic update   Yamada & MS (AAAI2011)

Numerical Example

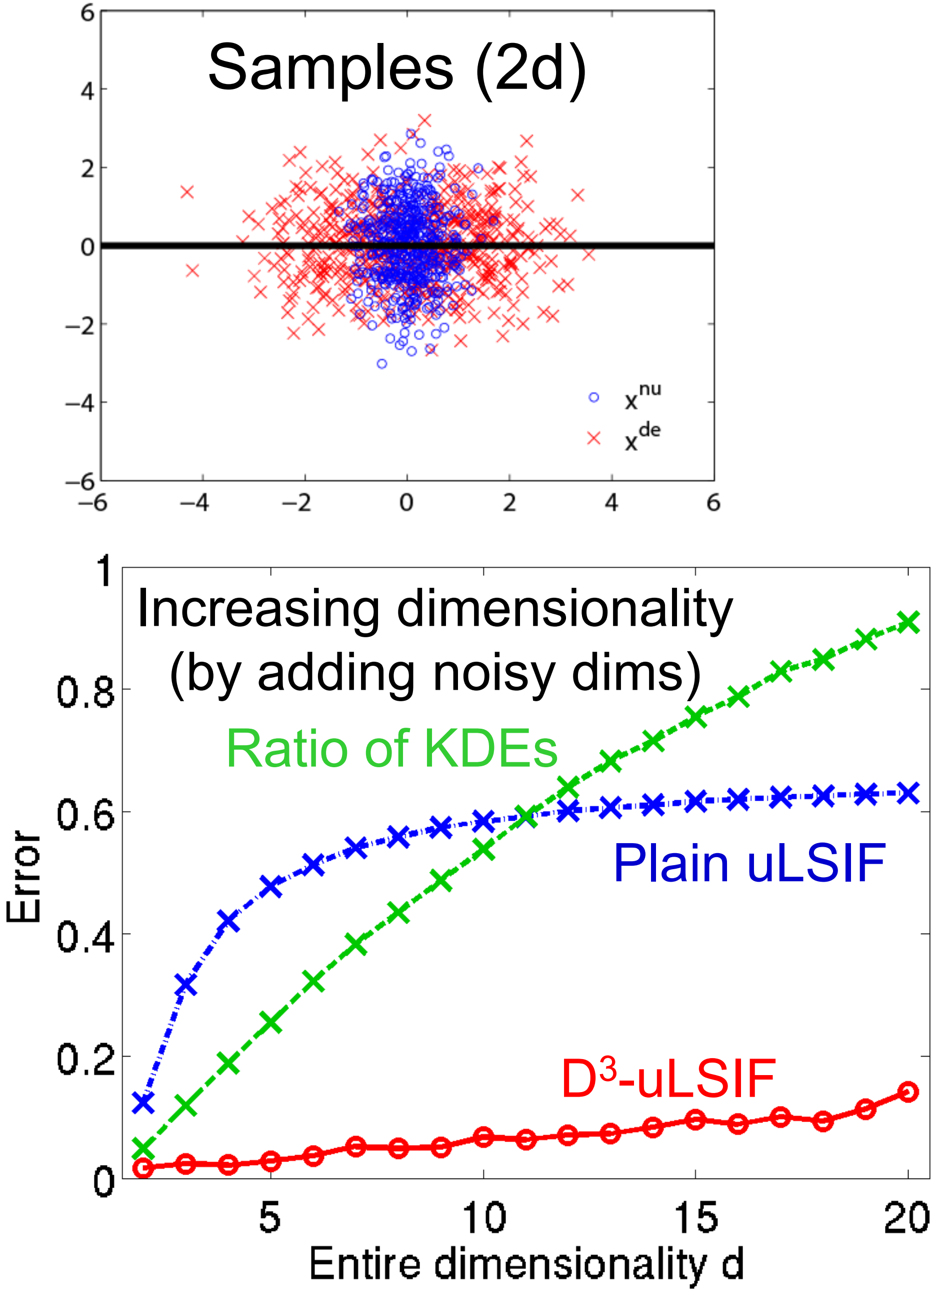

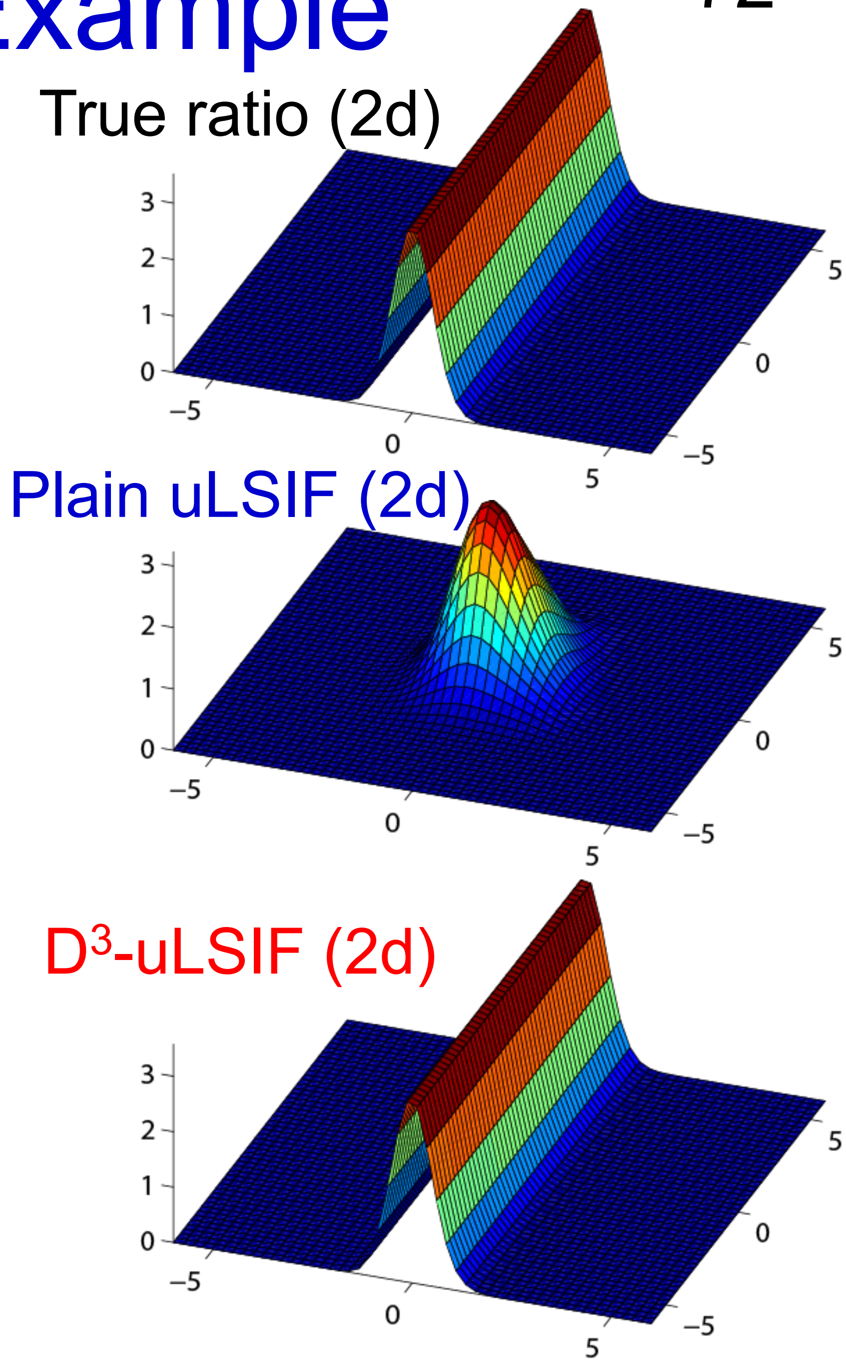

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
4. More on Density Ratio Estimation
   - A) Unified Framework
   - B) Dimensionality Reduction
   - C) Relative Density Ratios
5. Conclusions

Weakness of Density Ratios

- Density ratio can diverge to infinity:

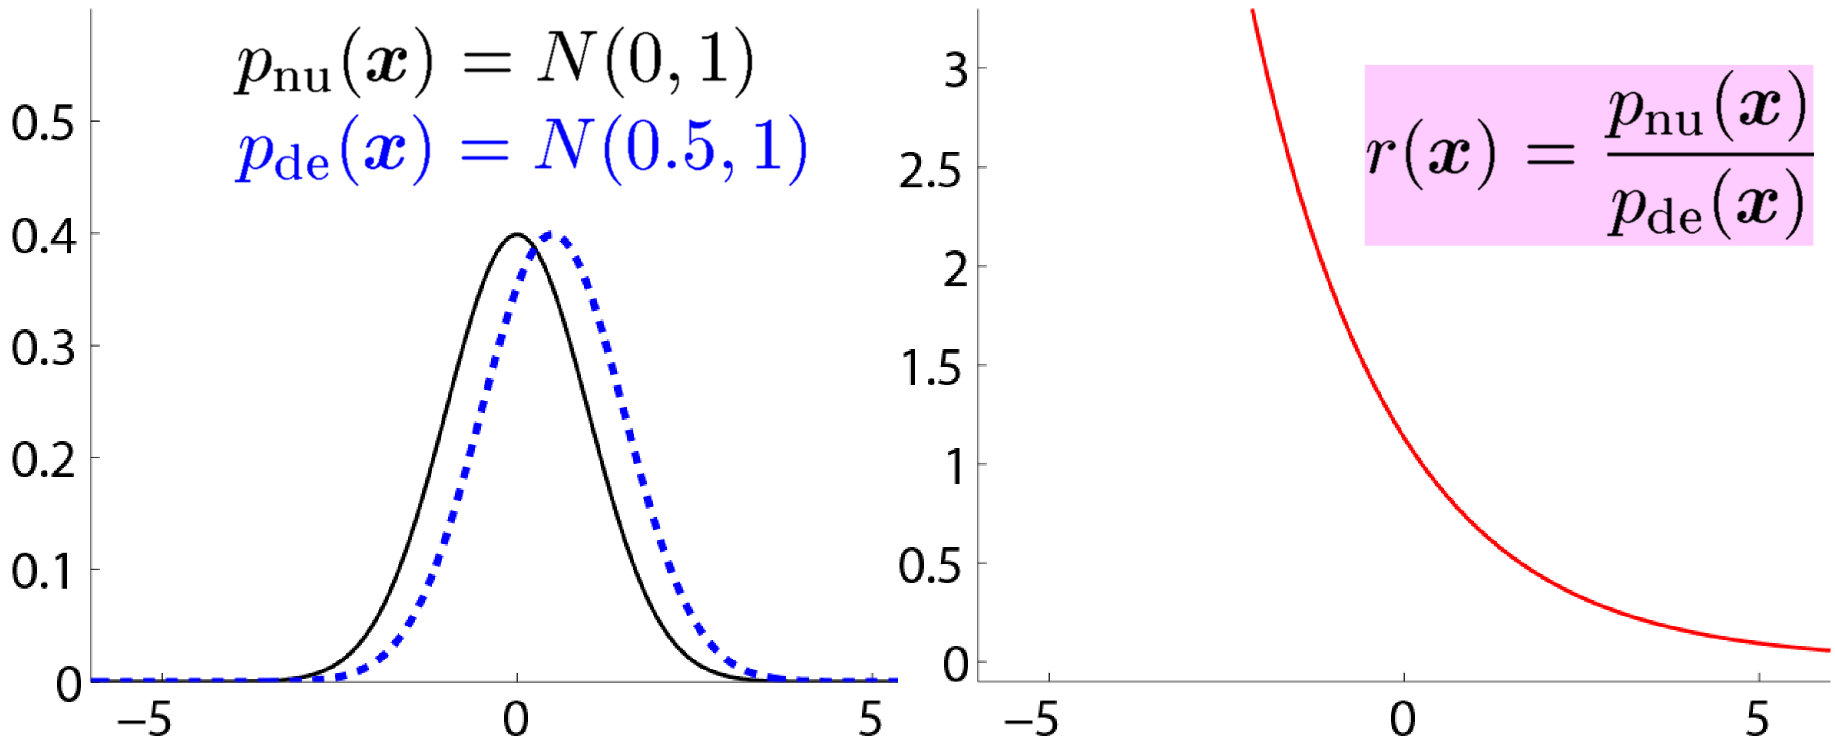

- Estimation becomes unreliable!

Relative Density Ratios

Yamada, Suzuki, Kanamori, Hachiya & MS (NIPS2011)

$$r_\beta(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{\beta p_{\mathrm{nu}}(\boldsymbol{x}) + (1-\beta) p_{\mathrm{de}}(\boldsymbol{x})} \qquad 0 \le \beta < 1$$

- Bounded for any $p_{\mathrm{nu}}(\boldsymbol{x}), p_{\mathrm{de}}(\boldsymbol{x})$: $r_\beta(\boldsymbol{x}) < \frac{1}{\beta}$

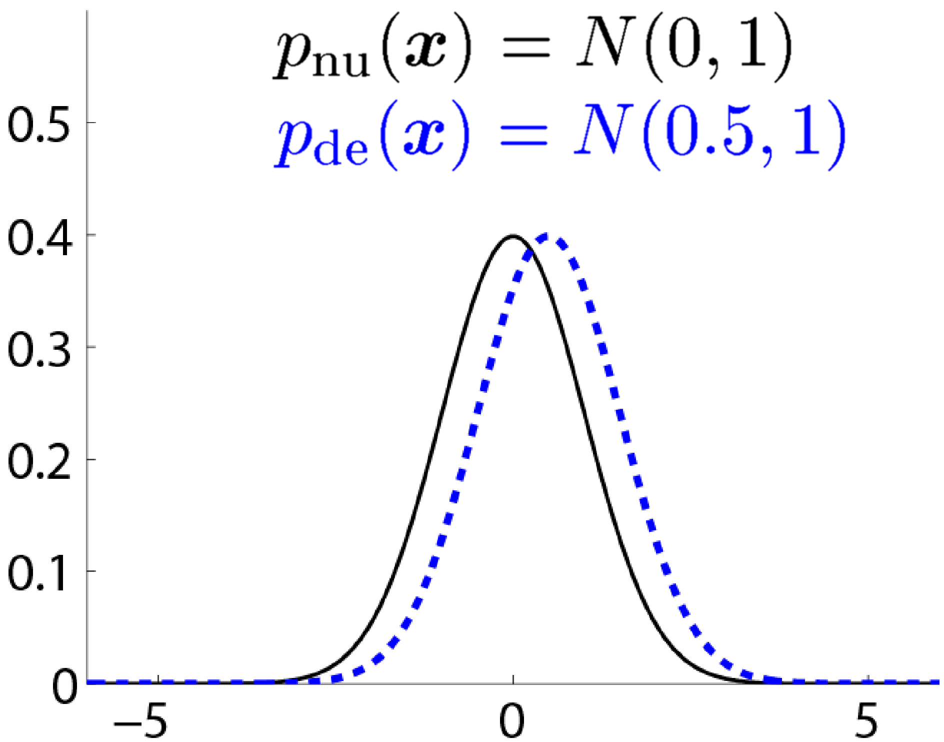

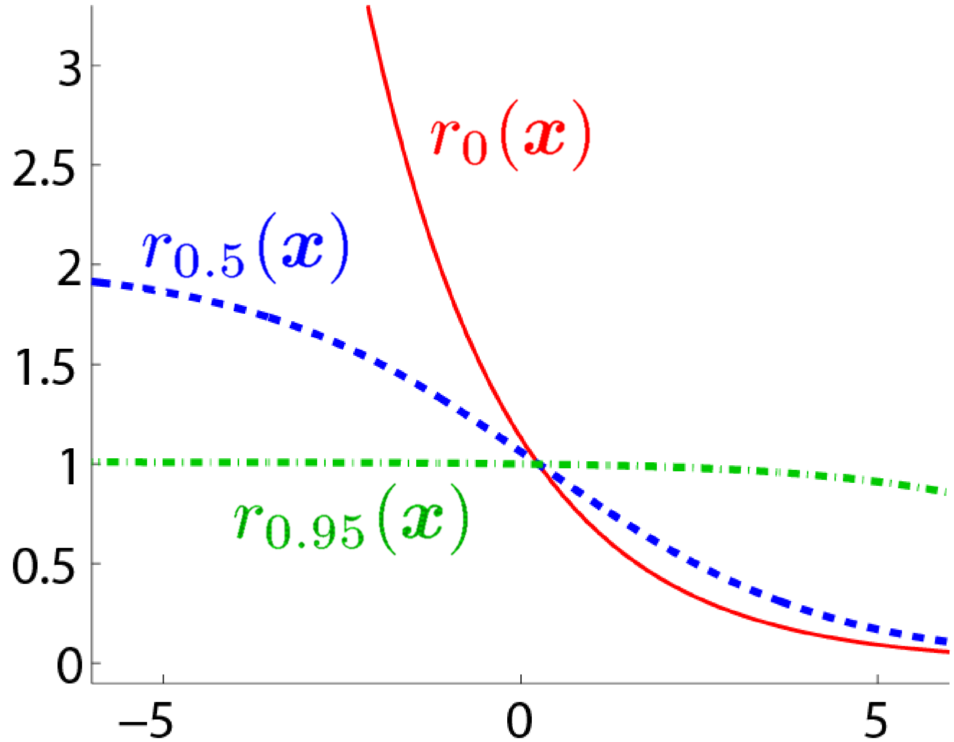

Estimation of Relative Ratios

- Linear model: $$\widehat{r}(\boldsymbol{x}) = \sum_{\ell=1}^{b} \alpha_\ell \phi_\ell(\boldsymbol{x}) = \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\boldsymbol{x})$$

- Relative unconstrained least-squares importance fitting (RuLSIF):

$$\min_{\widehat{r}} \int \Big(\widehat{r}(\boldsymbol{x}) - r_\beta(\boldsymbol{x})\Big)^2 q_\beta(\boldsymbol{x}) d\boldsymbol{x} \qquad r_\beta(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{q_\beta(\boldsymbol{x})}$$

$$q_\beta(\boldsymbol{x}) = \beta p_{\mathrm{nu}}(\boldsymbol{x}) + (1-\beta) p_{\mathrm{de}}(\boldsymbol{x})$$

- The solution can be computed analytically:

$$\mathop{\mathrm{argmin}}_{\boldsymbol{\alpha}} \left[\frac{1}{2} \boldsymbol{\alpha}^\top \widehat{\boldsymbol{H}} \boldsymbol{\alpha} - \widehat{\boldsymbol{h}}^\top \boldsymbol{\alpha} + \frac{\lambda}{2} \boldsymbol{\alpha}^\top \boldsymbol{\alpha} \right] = (\widehat{\boldsymbol{H}} + \lambda \boldsymbol{I})^{-1} \widehat{\boldsymbol{h}}$$

$$\widehat{\boldsymbol{H}} = \frac{\beta}{n_{\mathrm{de}}} \sum_{j=1}^{n_{\mathrm{de}}} \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}}) \boldsymbol{\phi}(\boldsymbol{x}_j^{\mathrm{de}})^\top + \frac{1-\beta}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}}) \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}})^\top \qquad \widehat{\boldsymbol{h}} = \frac{1}{n_{\mathrm{nu}}} \sum_{i=1}^{n_{\mathrm{nu}}} \boldsymbol{\phi}(\boldsymbol{x}_i^{\mathrm{nu}})$$

Relative Pearson Divergence

$$\mathrm{PE}_\beta[p_{\mathrm{nu}}(\boldsymbol{x}), p_{\mathrm{de}}(\boldsymbol{x})] = \frac{1}{2}\int \left(r_\beta(\boldsymbol{x}) - 1\right)^2 q_\beta(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$$

$$r_\beta(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{q_\beta(\boldsymbol{x})} \qquad q_\beta(\boldsymbol{x}) = \beta p_{\mathrm{nu}}(\boldsymbol{x}) + (1-\beta)p_{\mathrm{de}}(\boldsymbol{x})$$

- Relative Pearson divergence can be more reliably approximated:

$$\widehat{\mathrm{PE}}_\beta - \mathrm{PE}_\beta = \mathcal{O}_p(n^{-1/2}c\|r_\beta\|_\infty + \lambda_n \max(1, R(r_\beta)^2))$$

$$n = \min(n_{\mathrm{nu}}, n_{\mathrm{de}}) \qquad \lambda_n \to o(1) \ \text{ and } \ \lambda_n^{-1} = o(n^{2/(2+\gamma)}),\ 0 < \gamma < 2$$

$$\|r_\beta\|_\infty = \max_{\boldsymbol{x}} r_\beta(\boldsymbol{x}) = \left\|\left(\beta + (1-\beta)/r(\boldsymbol{x})\right)^{-1}\right\|_\infty < \frac{1}{\beta} \qquad r(\boldsymbol{x}) = \frac{p_{\mathrm{nu}}(\boldsymbol{x})}{p_{\mathrm{de}}(\boldsymbol{x})}$$

Organization of This Lecture

1. Introduction
2. Methods of Density Ratio Estimation
3. Usage of Density Ratios
4. More on Density Ratio Estimation
5. Conclusions

Conclusions

- Estimating data-generating probability distributions is universal, but inaccurate.
- Solving each task directly is ideal, but costly.
- Density ratio estimation is realistic compromise:
  - Systematically avoiding density estimation.
  - Applicable to a reasonably rich class of tasks.
  - Useful in many real-world problems.

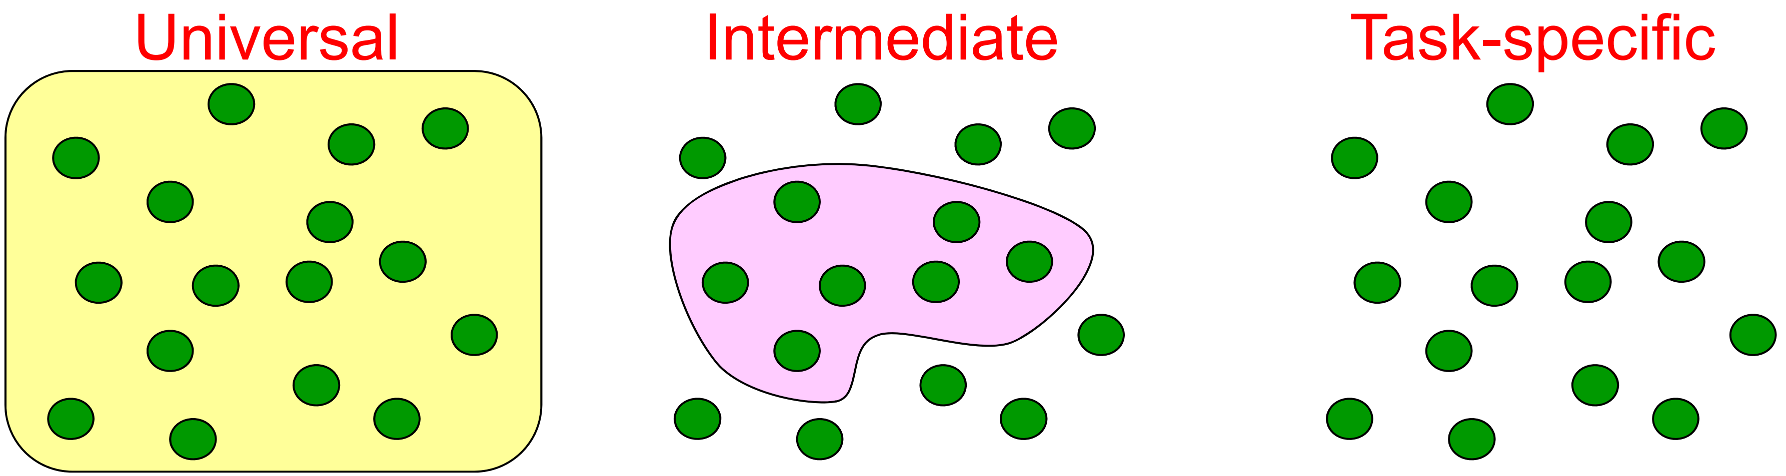

Books on Density Ratios

- Sugiyama, Suzuki & Kanamori,
  Density Ratio Estimation
  in Machine Learning,
  Cambridge University Press, 2012

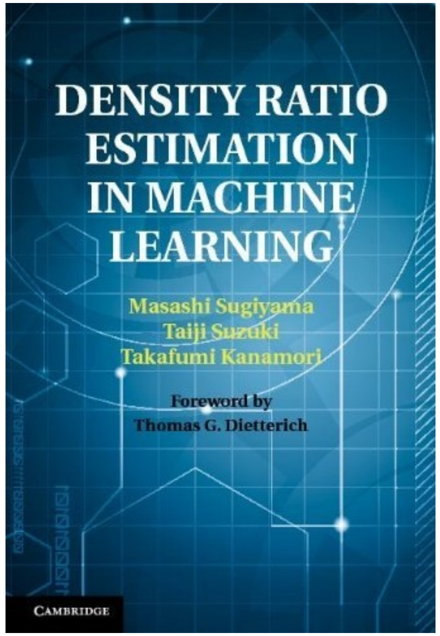

- Sugiyama & Kawanabe
  Machine Learning
  in Non-Stationary Environments,
  MIT Press, 2012

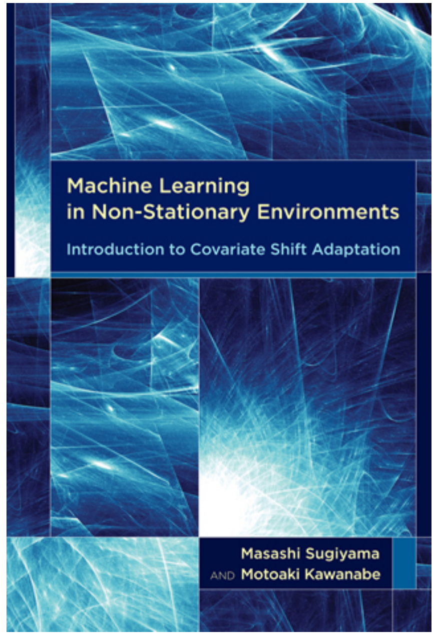

Acknowledgements

- Colleagues: Hirotaka Hachiya, Shohei Hido, Yasuyuki Ihara, Hisashi Kashima, Motoaki Kawanabe, Manabu Kimura, Masakazu Matsugu, Shin-ichi Nakajima, Klaus-Robert Müller, Jun Sese, Jaak Simm, Ichiro Takeuchi, Masafumi Takimoto, Yuta Tsuboi, Kazuya Ueki, Paul von Bünau, Gordon Wichern, Makoto Yamada.
- Funding Agencies: Ministry of Education, Culture, Sports, Science and Technology, Alexander von Humboldt Foundation, Okawa Foundation, Microsoft Institute for Japanese Academic Research Collaboration Collaborative Research Project, IBM Faculty Award, Mathematisches Forschungsinstitut Oberwolfach Research-in-Pairs Program, Asian Office of Aerospace Research and Development, Support Center for Advanced Telecommunications Technology Research Foundation, Japan Science and Technology Agency

- Papers, articles, and software of density ratio estimation are available from

  http://sugiyama-www.cs.titech.ac.jp/~sugi/