

# Variational Information Maximization in Gaussian Channels

**Felix V. Agakov**

School of Informatics, University of Edinburgh, EH1 2QL, UK  
felixa@inf.ed.ac.uk, <http://anc.ed.ac.uk>

**David Barber**

IDIAP, Rue du Simplon 4, CH-1920 Martigny Switzerland  
[www.idiap.ch](http://www.idiap.ch)

## Abstract

Recently, we introduced a simple variational bound on mutual information, that resolves some of the difficulties in the application of information theory to machine learning. Here we study a specific application to Gaussian channels. It is well known that PCA may be viewed as the solution to maximizing information transmission between a high dimensional vector  $\mathbf{x}$  and its low dimensional representation  $\mathbf{y}$ . However, such results are based on assumptions of Gaussianity of the sources  $\mathbf{x}$ . In this paper, we show how our mutual information bound, when applied to this arena, gives PCA solutions, without the need for the Gaussian assumption. Furthermore, it naturally generalizes to providing an objective function for Kernel PCA, enabling the principled selection of kernel parameters.

## 1 Introduction

Maximization of information transmission in noisy channels is a common problem, ranging from the construction of good error-correcting codes [12] and feature extraction [17] to neural sensory processing [11], [6].

The key idea of information maximization is to choose a mapping from source variables (inputs)  $\mathbf{x}$  to response variables (outputs)  $\mathbf{y}$  such that the outputs contain as much information about which of the inputs was transmitted as possible. In a stochastic context, we have a source distribution  $p(\mathbf{x})$ , and a mapping  $p(\mathbf{y}|\mathbf{x})$ . The general aim will be to set any adjustable parameters of  $p(\mathbf{y}|\mathbf{x})$  in order to maximize information transfer. The principal measure of information transfer in this context is the mutual information defined as

$$I(\mathbf{x}, \mathbf{y}) \equiv H(\mathbf{x}) - H(\mathbf{x}|\mathbf{y}). \quad (1)$$

Equivalently, we may write  $I(\mathbf{x}, \mathbf{y}) \equiv H(\mathbf{y}) - H(\mathbf{y}|\mathbf{x})$ . Here  $H(\mathbf{y}) \equiv -\langle \log p(\mathbf{y}) \rangle_{p(\mathbf{y})}$  and  $H(\mathbf{y}|\mathbf{x}) \equiv -\langle \log p(\mathbf{y}|\mathbf{x}) \rangle_{p(\mathbf{x}, \mathbf{y})}$  are marginal and conditional entropies respectively, and angled brackets represent averages. The objective (1) is maximized with respect to parameters of the encoder  $p(\mathbf{y}|\mathbf{x})$ .

Our specific interest in this paper is when we have a finite set of training points  $\{\mathbf{x}_m | m = 1, \dots, M\}$ , which forms the empirical distribution  $p(\mathbf{x}) = (1/M) \sum_{m=1}^M \delta(\mathbf{x} - \mathbf{x}_m)$ . Our aim will be to form constrained representations  $p(\mathbf{y}|\mathbf{x})$  that maximize  $I(\mathbf{x}, \mathbf{y})$ . For a continuous encoder  $p(\mathbf{y}|\mathbf{x})$ , the theoretically optimal unconstrained decoder is given by Bayes' rule:

$$p(\mathbf{x}|\mathbf{y}) = \sum_{k=1}^M \delta(\mathbf{x} - \mathbf{x}_k) \omega_k(\mathbf{y}), \quad \omega_k(\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x}_k)}{\sum_{m=1}^M p(\mathbf{y}|\mathbf{x}_m)},$$

which is a mixture of Dirac delta functions. This reduces to a single delta peak only for the case when stochastic images  $\mathcal{I}$  of distinct source vectors  $\{\mathbf{x}\}$  are non-intersecting under the encoder  $p(\mathbf{y}|\mathbf{x})$ , i.e.

$\forall i, j \in \{1, \dots, M\}, i \neq j. \mathcal{I}(x_i) \cap \mathcal{I}(x_j) = \emptyset$ . Formally, this is never the case if  $\forall y. \forall x. p(y|x) \neq 0$ , which leads to a mixture form of  $p(x|y)$ .

Despite the conceptual simplicity of the statement, computation of the mutual information is generally intractable for all but special cases. For large-scale systems the key difficulty lies in the computation of the conditional entropy  $H(x|y)$  of the mixture distribution  $p(x|y)$ , which is in general NP hard.

Standard approaches address the problem of optimizing (1) by assuming that  $p(x, y)$  is jointly Gaussian [10], the output spaces are very low-dimensional [11], or the channels are deterministic and invertible [2]. Other popular methods suggest alternative objective functions (e.g. the Fisher information criterion [4]), which, however, do not preserve proper bounds on  $I(x, y)$ .

## 1.1 Variational Lower Bound on Mutual Information

In order to derive a simple lower bound on the mutual information, we consider the Kullback-Leibler divergence  $KL(p(x|y)||q(x|y))$  between the posterior  $p(x|y)$  and its variational approximation  $q(x|y)$ . Non-negativity of the divergence gives

$$\langle \log p(x|y) \rangle_{p(x|y)} - \langle \log q(x|y) \rangle_{p(x|y)} \geq 0 \Rightarrow \underbrace{\langle \log p(x|y) \rangle_{p(x|y)p(y)}}_{-H(x|y)} \geq \langle \log q(x|y) \rangle_{p(x|y)p(y)}. \quad (2)$$

This leads to

$$I(x, y) \geq H(x) + \langle \log q(x|y) \rangle_{p(x, y)} \stackrel{\text{def}}{=} \tilde{I}(x, y), \quad (3)$$

where  $q(x|y)$  is an arbitrary variational distribution which saturates the bound for  $q(x|y) \equiv p(x|y)$ . Note that the objective (3) explicitly includes both the encoder  $p(y|x)$  and decoder  $q(x|y)$ .

Other well known lower bounds on the mutual information could be considered [7]. However, our current experience suggests that for certain choices of the decoder  $q(x|y)$  the variational bound considered above is particularly computationally convenient. Moreover, since (3) is based on the KL divergence between the true and the approximating posteriors, it is equivalent to a moment matching approximation of  $p(x|y)$  by  $q(x|y)$ . This fact is beneficial in terms of decoding, since the more successful decoding algorithms approximate the mean of the posterior  $p(x|y)$  [13], whilst standard mode matching approaches (such as mean-field theory) typically get trapped in the one of many sub-optimal local minima.

Recently in [1] we discussed several applications of the bound (3) and outlined its basic properties. The principal objective of this paper is to investigate properties of the bound for the case when the decoder  $q(x|y)$  is Gaussian. We will also assume that the encoder  $p(y|x)$  is Gaussian, although this assumption is less critical to the tractable optimization of the objective.

## 1.2 Gaussian Decoders

In this paper we focus on a simple assumption that the approximate decoder  $q(x|y)$  is a Gaussian. Although the optimal decoder in the considered channels should be a mixture distribution  $q(x|y) = p(x|y)$ , one may hope that Gaussian decoders  $q(x|y)$  perform well if the codes are not strongly overlapping. We show that optimal Gaussian decoders have a strong relation to popular dimensionality reduction techniques, inducing PCA and kernel PCA solutions as special cases.

In the following sections we will consider increasingly complex encoders and decoders. Initially, however, we demonstrate the application of the bound to the simple case of a linear Gaussian encoder and decoder.

## 2 Linear Gaussian Decoder: $p(y|x) \sim \mathcal{N}_y(Wx, s^2\mathbf{I})$ , $q(x|y) \sim \mathcal{N}_x(Uy, \sigma^2\mathbf{I})$

Let the encoder and decoder be given by  $p(y|x) \sim \mathcal{N}_y(Wx, s^2\mathbf{I})$  and  $q(x|y) \sim \mathcal{N}_x(Uy, \sigma^2\mathbf{I})$  respectively. Our goal is to learn optimal settings of  $W \in \mathbb{R}^{|y| \times |x|}$  and  $U \in \mathbb{R}^{|x| \times |y|}$  (for fixed  $\sigma^2$  and  $s^2$ ) by maximizing the variational bound  $\tilde{I}(x, y)$  on the mutual information, which in this case is expressed as

$$\tilde{I}(x, y) \propto 2\text{tr}\{UWS\} - \text{tr}\{U\Sigma U^T\} + c. \quad (4)$$

Here  $c$  is a constant,  $S = \langle xx^T \rangle = \sum_m x_m(x_m)^T/M$  is the sample covariance of the zero-mean data, and

$$\Sigma = \mathbf{I}s^2 + WSW^T \in \mathbb{R}^{|y| \times |y|} \quad (5)$$

is the covariance of the distribution of the responses  $p(y)$ . In the following we assume that the weights  $W$ ,  $U$  and the sample covariance  $S$  are non-singular. Note that we make no assumptions about the distribution of the sources  $p(x)$ . Unsurprisingly, this objective is closely related to the least squares reconstruction error in a linear autoencoder. What is particularly interesting in this context is that it provides a lower bound on the mutual information.

### Nature of optimal solutions

Unconstrained optimization of (4) for the encoder's weights  $W$  leads to the extremum condition

$$U^T S = U^T U W S. \quad (6)$$

By assuming that  $y$  is a compressed representation of the source  $x$  (i.e.  $|x| > |y|$ ), we obtain  $W = (U^T U)^{-1} U^T$ . This transforms the objective function (4) into

$$\tilde{I}(x, y) = \text{tr}\{UWSW^T U^T\} - s^2 \text{tr}\{UU^T\} = \text{tr}\{U(U^T U)^{-1} U^T S\} - s^2 \text{tr}\{UU^T\}, \quad (7)$$

where we ignored the irrelevant constant  $c$ .

Let  $U = VLR^T$  be the singular value decomposition of the decoder weights  $U$  where, by definition,  $L \in \mathbb{R}^{|y| \times |y|}$  is diagonal and  $V^T V = R^T R = RR^T = \mathbf{I}_{|y|}$ . From (6) it is clear that  $W = RL^{-1}V^T$ , i.e.  $WU = \mathbf{I}_{|y|}$ . Substitution into (7) results in

$$\tilde{I}(x, y) = \text{tr}\{V^T S V\} - s^2 \text{tr}\{L^2\}. \quad (8)$$

Optimizing (8) for  $V$  under the orthonormality constraints on  $V$ , we readily obtain the PCA solution (and its rigid rotations). In the limit  $s^2 \rightarrow 0$ , the solutions are invariant with respect to the scalings  $L$ .

For noisy channels  $s^2 > 0$ , maximization of (8) leads to divergence of the Frobenius norm  $\|W\|_F$ , so that the contribution of the finite isotropic noise in (5) is negligible. One way to handle the problem of divergence for  $s^2 > 0$  is by constraining the singular values of  $U$ . In this case optimal weights  $U$  correspond to principal eigenvectors of  $S$ . Note that solutions are invariant with respect to complementary rotations of  $W$  and  $U$ .

This demonstrates the simple result that PCA provides a lower bound on the mutual information between  $x$  and  $y$ . Again, we stress that this conclusion is reached without the need for a Gaussian assumption about the source distribution. In order to go beyond the PCA solution, more complex encoders/decoders are required. In the following section we consider the effect of increasing the complexity of the encoder, whilst still using a simple linear decoder as above.

## 3 Nonlinear Gaussian Channels: $p(y|x) \sim \mathcal{N}_y(W\phi(x), \Sigma_y)$ , $q(x|y) \sim \mathcal{N}_x(Uy, \Sigma_x)$

Here we consider the case of a non-linear Gaussian channel with the encoder  $p(y|x) \sim \mathcal{N}_y(W\phi(x), \Sigma_y)$  and decoder  $q(x|y) \sim \mathcal{N}_x(Uy, \Sigma_x)$ . Note that for all the data points  $\{x_m | m = 1, \dots, M\}$ , the set of encodings

$\{y_m\}$  is given by a noisy linear projection from the (potentially high-dimensional) feature space  $\{\phi(x_m)\}$ . In what follows we assume that  $|\phi| > M$ , i.e. dimensionality of the feature space exceeds the number of training points.

It is easy to see that for the considered case the lower bound on the mutual information  $I(x, y)$  is given by

$$\tilde{I}(x, y) = -\frac{1}{2} \text{tr} \{ \Sigma_x^{-1} S \} + \text{tr} \{ \Sigma_x^{-1} U W \langle \phi(x) x^T \rangle \} - \frac{1}{2} \text{tr} \{ U^T \Sigma_x^{-1} U (\Sigma_y + W \langle \phi(x) \phi(x)^T \rangle W^T) \}$$

where  $S = \langle x x^T \rangle = \sum_m x_m (x_m)^T / M$  is the sample covariance of the zero-mean data, and the averages are computed with respect to the empirical distribution  $p(x) = (1/M) \sum_{m=1}^M \delta(x - x_m)$ . Note, however, that for high-dimensional feature spaces direct evaluation of the averages is implausible. It is therefore desirable to avoid explicit computations in  $\{\phi\}$ .

### 3.1 Kernelized Representation

Since each row  $\tilde{w}_i^T \in \mathbb{R}^{1 \times |\phi|}$  of the weight matrix  $W \in \mathbb{R}^{|\phi| \times |\phi|}$  has the same dimensionality as the feature vectors  $\phi(x_i)^T$ , it is representable as

$$\tilde{w}_i = \sum_{m=1}^M \alpha_{im} \phi(x_m) + \tilde{w}_i^\perp, \quad (9)$$

where  $\tilde{w}_i^\perp$  is orthogonal to the span of  $\phi(x_1), \dots, \phi(x_M)$ . Then

$$W = A F^T + W^\perp, \quad F \stackrel{\text{def}}{=} [\phi(x_1), \dots, \phi(x_M)] \in \mathbb{R}^{|\phi| \times M}, \quad (10)$$

where  $A = \{\alpha_{ij}\} \in \mathbb{R}^{|\phi| \times M}$  is the matrix of coefficients, and the transposed rows of  $W^\perp$  are given by  $\tilde{w}_i^\perp$ . In kernel literature  $F$  is often referred to as the *design* matrix (e.g. [18]).

From (10), we obtain expressions for the averages

$$\begin{aligned} W \langle \phi(x) x^T \rangle &= A \left\langle [\phi(x_1)^T \phi(x), \dots, \phi(x_M)^T \phi(x)]^T x^T \right\rangle \\ &= \frac{A}{M} \left[ \sum_m x_m \phi(x_m)^T \phi(x_1), \dots \right]^T = \frac{A B^T}{M} \end{aligned}$$

where we defined

$$B \stackrel{\text{def}}{=} \sum_{m=1}^M x_m k(x_m)^T \in \mathbb{R}^{|\phi| \times M}, \quad k(x_m) = F^T \phi(x_m) \in \mathbb{R}^M. \quad (11)$$

Here  $k(x_m)$  corresponds to the  $m^{\text{th}}$  column (or row) of the Gram matrix  $K = \{K_{ij}\} \stackrel{\text{def}}{=} \{\phi(x_i)^T \phi(x_j)\} \in \mathbb{R}^{M \times M}$ . Finally, note that for a fixed  $K$  the computed expectation  $W \langle \phi(x) x^T \rangle$  is a function of  $A$ .

Analogously, we can express

$$W \langle \phi(x) \phi(x)^T \rangle W^T = \frac{1}{M} A \sum_{m=1}^M k(x_m) k(x_m)^T A^T. \quad (12)$$

Again, for the fixed Gram matrix the term is a quadratic function of coefficients  $A$  alone.

By substitution, we may re-express the bound (9) as

$$\tilde{I}(x, y) = \frac{1}{M} \text{tr} \{ \Sigma_x^{-1} U A B^T \} - \frac{1}{2} \text{tr} \{ U^T \Sigma_x^{-1} U \Sigma_y \} - \frac{1}{2M} \text{tr} \{ U^T \Sigma_x^{-1} U A K^2 A^T \} - \frac{1}{2} \text{tr} \{ \Sigma_x^{-1} S \}. \quad (13)$$

In the simplest case when  $\phi(\mathbf{x}) \equiv \mathbf{x} \in \mathbb{R}^{|\mathbf{x}|}$ , we obtain  $\mathbf{K}^2 \propto \mathbf{X}^T \mathbf{S} \mathbf{X} \in \mathbb{R}^{M \times M}$ ,  $\mathbf{B} \propto \mathbf{S} \mathbf{X} \in \mathbb{R}^{|\mathbf{x}| \times M}$  where  $\mathbf{X} \stackrel{\text{def}}{=} [\mathbf{x}_1, \dots, \mathbf{x}_M] \in \mathbb{R}^{|\mathbf{x}| \times M}$  contains the training data. As expected, this transforms the bound (13) to the corresponding expression (7) for the linear Gaussian channel, thus resulting in PCA on the sample covariance  $\mathbf{S}$  for both the encoder  $\mathbf{W}^T$  and the decoder  $\mathbf{U}$  as the optimal choice.

The objective (13) may be used to learn optimal decoder weights  $\mathbf{U}$  and optimal coordinates  $\mathbf{A}$  in the space spanned by the feature vectors  $\{\phi(\mathbf{x}_i) | i \in [1, M] \cap \mathbb{N}\}$ . Moreover, note that if  $\mathcal{K}_{\Theta} : |\mathbf{x}| \times |\mathbf{x}| \rightarrow \mathbb{R}$  defines a symmetric positive-definite *kernel*, by Mercer's theorem the elements  $K_{ij}$  of the Gram matrix may be replaced by  $\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j; \Theta)$ . In this case, the bound  $\tilde{I}(\mathbf{x}, \mathbf{y})$  may be used to learn the optimal parameters<sup>1</sup> of the kernel function. Thus, in our framework the procedure for learning kernels may be viewed as a special case of the variational IM algorithm [1].

### 3.2 Nature of optimal solutions

In the following we assume for simplicity that  $\Sigma_y = s^2 \mathbf{I}$  and  $\Sigma_x = \sigma^2 \mathbf{I}$ . We also assume that  $|\mathbf{y}| \leq |\mathbf{x}| \leq |\phi|$  and  $|\mathbf{x}| \leq M$ , so that  $\mathbf{y}$  is a compressed representation of  $\phi(\mathbf{x})$ , and the number of training points is sufficient to ensure invertibility of the sample covariance.

#### Optimal Decoder

Optimization of  $\tilde{I}(\mathbf{x}, \mathbf{y})$  for the matrix of coefficients  $\mathbf{A}$  leads to the fixed point condition

$$\partial \tilde{I}(\mathbf{x}, \mathbf{y}) / \partial \mathbf{A} = 0 \Rightarrow \mathbf{U}^T \Sigma_x^{-1} \mathbf{B} = \mathbf{U}^T \Sigma_x^{-1} \mathbf{U} \mathbf{A} \mathbf{K}^2. \quad (14)$$

For a non-singular Gram matrix  $\mathbf{K}$  we obtain

$$\tilde{I}(\mathbf{x}, \mathbf{y}) \propto \text{tr} \{ \mathbf{U} \mathbf{A} \mathbf{K}^2 \mathbf{A}^T \mathbf{U}^T \} - s^2 \text{tr} \{ \mathbf{U} \mathbf{U}^T \} = \text{tr} \{ \mathbf{U} (\mathbf{U}^T \mathbf{U})^{-1} \mathbf{U}^T \mathbf{B} \mathbf{K}^{-2} \mathbf{B}^T \} - s^2 \text{tr} \{ \mathbf{U} \mathbf{U}^T \}. \quad (15)$$

Let  $\tilde{\mathbf{S}}_F = \sum_{m=1}^M \mathbf{x}_m \phi(\mathbf{x}_m)^T / M \in \mathbb{R}^{|\mathbf{x}| \times |\phi|}$ . By noticing that

$$\mathbf{B} \mathbf{K}^{-2} \mathbf{B}^T \propto \tilde{\mathbf{S}}_F \mathbf{K}^{-2} \mathbf{F}^T \tilde{\mathbf{S}}_F^T \propto \mathbf{S}, \quad (16)$$

the bound (15) reduces to the objective of the linear Gaussian channel (6). Note that  $\|\mathbf{W}\|_F \rightarrow \infty$  as  $\|\mathbf{U}\|_F \rightarrow 0$ .

Thus, under the specific assumption of isotropic Gaussian noise, optimal weights  $\mathbf{U}$  of the linear Gaussian decoder correspond to principal components (and their rotations) of the sample covariance  $\mathbf{S}$ . Similarly to the case of a linear channel, for  $s^2 > 0$  it is necessary to impose norm constraints on  $\mathbf{U}$  to ensure convergence of  $\|\mathbf{W}\|_F$ . Fundamentally, therefore, this simple decoder will restrict severely the power of the approach, as we will see below.

#### Optimal Encoder

From (14) it is clear that optimal solutions for the encoder are given by

$$\mathbf{W} \mathbf{F} = \mathbf{A} \mathbf{K} \propto \mathbf{U}^+ \mathbf{X}, \quad (17)$$

where  $\mathbf{U}^+$  denotes the pseudo-inverse, and left singular values of  $\mathbf{U}$  correspond to principal eigenvectors of  $\mathbf{S}$ . In the case when  $\phi(\mathbf{x}) \equiv \mathbf{x} \in \mathbb{R}^{|\mathbf{x}|}$  and  $\mathbf{X} \mathbf{X}^T$  is non-singular, condition (17) results in  $\mathbf{W} = \mathbf{U}^+ \in \mathbb{R}^{|\mathbf{y}| \times |\mathbf{x}|}$ ,

<sup>1</sup>It is also possible to learn  $\mathbf{K}$  directly by constraining it to satisfy properties of inner products; this may be useful, for example, when the source alphabet is exhausted by  $M$  training points.

which is the PCA solution of the linear Gaussian channel. However, for general non-linear mappings, optimal encoder weights  $W^T$  do not necessarily give rise to the non-linear PCA solution.

Finally, note that if the channel noise is isotropic and there are no constraints preventing the weights from taking optimal solutions according to (15) and (17), then the bound is given by the summation of  $|y|$  principal eigenvalues of the sample covariance  $\mathbf{S}$ . In this case the objective is invariant under the choice of non-linearity. Hence, we reach an important conclusion: *for isotropic channel noise, if the decoder is linear, nothing is gained by using a non-linear encoder* in the proposed variational settings. In other cases, for example when the channel noise is correlated or the encoder and decoder weights have structural constraints, optimal parameters of  $\mathcal{K}_\Theta$  may be obtained by maximizing (13).

These results are somewhat disappointing. In order to improve the power of the method, we need to consider both non-linear encoders and decoders. However, the naive method of using a non-linear decoder would typically result in intractable averages over  $y$ . In order to avoid this difficulty, we consider how to form decoding in the feature space.

## 4 Nonlinear Decoders and KPCA

The non-linear Gaussian channel discussed in Section 3 may be represented by the Markov chain  $\mathbf{x} \rightarrow \mathbf{f} \rightarrow \mathbf{y}$ , where  $\mathbf{f} \in \mathbb{R}^{|\phi|}$  and  $p(\mathbf{f}|\mathbf{x}) \sim \delta(\mathbf{f} - \phi(\mathbf{x}))$ ,  $p(\mathbf{y}|\mathbf{f}) \sim \mathcal{N}_y(\mathbf{W}\mathbf{f}, \Sigma_y)$ . Indeed, by marginalizing the feature variables  $\mathbf{f}$  it is clear that the encoder is given<sup>2</sup> by  $p(\mathbf{y}|\mathbf{x}) = \int_{\mathbf{f}} \delta(\mathbf{f} - \phi(\mathbf{x})) \mathcal{N}_y(\mathbf{W}\mathbf{f}, \Sigma_y) = \mathcal{N}_y(\mathbf{W}\phi(\mathbf{x}), \Sigma_y)$ .

**Proposition 1:** *Let  $\mathbf{s} \rightarrow \mathbf{t} \rightarrow \mathbf{r}$  define a Markov chain, such that  $p(\mathbf{t}|\mathbf{s}) = \delta(\mathbf{t} - \mathbf{f}(\mathbf{s}))$ , and  $p(\mathbf{r}|\mathbf{t})$  is a continuous differentiable density function satisfying  $\forall \mathbf{r}, \forall \mathbf{t}, p(\mathbf{r}|\mathbf{t}) \neq 0$ . Then  $I(\mathbf{s}, \mathbf{r}) = I(\mathbf{t}, \mathbf{r})$ .*

From proposition 1, the mutual information  $I(\mathbf{x}, \mathbf{y})$  may be bounded as

$$I(\mathbf{x}, \mathbf{y}) = I(\mathbf{f}, \mathbf{y}) \geq \tilde{I}(\mathbf{f}, \mathbf{y}), \text{ where } \tilde{I}(\mathbf{f}, \mathbf{y}) \stackrel{\text{def}}{=} \langle \log q(\mathbf{f}|\mathbf{y}) \rangle_{p(\mathbf{x})p(\mathbf{f}|\mathbf{x})p(\mathbf{y}|\mathbf{f})} + H(\mathbf{f}). \quad (18)$$

We make the simple assumption that the feature decoder is Gaussian,  $q(\mathbf{f}|\mathbf{y}) \sim \mathcal{N}_f(\mathbf{U}\mathbf{y}, \Sigma_f)$ , for which

$$\tilde{I}(\mathbf{x}, \mathbf{y}) = -\frac{1}{2} \text{tr} \left\{ \mathbf{U}^T \Sigma_f^{-1} \mathbf{U} (\Sigma_y + \mathbf{W} \mathbf{S}_F \mathbf{W}^T) \right\} + \text{tr} \left\{ \Sigma_f^{-1} \mathbf{U} \mathbf{W} \mathbf{S}_F \right\} + H(\mathbf{f}) + c \quad (19)$$

where  $\mathbf{S}_F \stackrel{\text{def}}{=} \langle \mathbf{f}\mathbf{f}^T \rangle_{p(\mathbf{f})}$ . If  $\langle \mathbf{f} \rangle = 0$  then clearly  $\mathbf{S}_F$  corresponds to the sample covariance in the feature space (see [14] for a discussion of centering of the data in feature spaces). This covariance is readily computable from the training set as

$$\mathbf{S}_F = \frac{1}{M} \sum_{m=1}^M \phi(\mathbf{x}_m) \phi(\mathbf{x}_m)^T, \quad (20)$$

where we assumed that  $\mathbf{x} = \phi^{-1}(\mathbf{f})$ . Note that if  $\phi : \mathbf{x} \rightarrow \mathbf{f}$  is deterministic and  $p(\mathbf{x}) = \sum_{i=1}^M \delta(\mathbf{x} - \mathbf{x}_i) / M$  then  $\phi(\mathbf{x}_i)$  corresponds to a re-labeling of the source pattern, leading to  $H(\mathbf{f}) = H(\mathbf{x})$ . As expected, linear mappings  $\phi(\mathbf{x}) \equiv \mathbf{x}$  result in  $\mathbf{S}_F = \langle \mathbf{x}\mathbf{x}^T \rangle \equiv \mathbf{S}$ , which reduces the objective (19) to (4).

It is easy to see that in general this case gives rise to non-linear source decoders  $q(\mathbf{x}|\mathbf{y}) = \int_{\mathbf{f}} p(\mathbf{x}|\mathbf{f}) \mathcal{N}_f(\mathbf{U}\mathbf{y}, \sigma_f^2 \mathbf{1})$ ; however, they may be difficult to compute. This is an important limitation of the approach, since for general feature mappings  $\phi$  it may be difficult to reconstruct  $\mathbf{x}$  from its encoded representation  $\mathbf{y}$ .

### 4.1 Kernelized Representation

In what follows we assume that  $\Sigma_y = s^2 \mathbf{1} \in \mathbb{R}^{|\mathbf{y}|}$ ,  $\Sigma_f = \sigma_f^2 \mathbf{1} \in \mathbb{R}^{|\phi|}$ , and  $|\mathbf{y}| < M < |\phi|$ . By analogy with Section 3 we notice that rows of  $\mathbf{W}$  and columns of  $\mathbf{U}$  have dimension  $|\phi|$ . Then they may be represented

<sup>2</sup>We assume Cartesian coordinates, i.e.  $\delta(\mathbf{x} - \mathbf{a}) = \prod_i \delta(x_i - a_i)$ .

in the basis defined by the span of  $\{\phi(x_m)|m = 1, \dots, M\}$  and its orthogonal complement as

$$W = AF^T + W^\perp \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{X}|}, \quad U = FC + U^\perp \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|}. \quad (21)$$

Here  $F \in \mathbb{R}^{|\mathcal{X}| \times M}$  is the design matrix,  $U^\perp, W^\perp$  are orthogonal to  $F$ , and  $A \in \mathbb{R}^{|\mathcal{Y}| \times M}, C \in \mathbb{R}^{M \times |\mathcal{Y}|}$  are matrices of coefficients to be learned. Substitution into expression (19) results in

$$\tilde{I}(x, y) \propto 2\text{tr}\{CAK^2\} - s^2\text{tr}\{C^TKC\} - \text{tr}\{AK^2A^T[C^TKC + (U^\perp)^TU^\perp]\}, \quad (22)$$

where we ignored the terms independent of  $A, C$  and  $K$ . It is clear that since we are interested in maximizing the bound, we may assume  $U^\perp = \mathbf{0} \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|}$ .

The objective (22) may be readily used for learning coefficients  $A, C$ . In the considered case it may also be applied to learning parameters  $\Theta$  of the kernel function  $\mathcal{K}(x_i, x_j; \Theta)$  which gives rise to the Gram matrix  $K$ .

## 4.2 Nature of Optimal Solutions

Optimization of the bound (22) for the coefficients  $A$  results in

$$A = (C^TKC)^{-1}C^T, \quad (23)$$

which transforms the objective to

$$\tilde{I}(x, y) = \text{tr}\{K^2C(C^TKC)^{-1}C^T\} - s^2\text{tr}\{C^TKC\} = \text{tr}\{U(U^TU)^{-1}U^TS_F\} - s^2\text{tr}\{UU^T\} \quad (24)$$

By analogy with Section 2, maximization of (24) gives rise to the non-linear PCA solution (and rotations) for the left singular vectors of  $U$  and  $W^T$ . (As before, we have ignored  $W^\perp$  and  $U^\perp$  in the definitions (21) since they have no influence on the bound).

Just as in the linear case, in order to prevent divergence of  $\|W\|_F$  for  $s^2 \neq 0$ , it is useful to constrain the singular values of  $U$ . In the special case when  $U^TU = I$ , expressions (21) and (24) lead to

$$K^2CR = KCR\Lambda_{S_F}. \quad (25)$$

Here  $\Lambda_{S_F} \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$  is a diagonal matrix of eigenvalues of  $S_F$  and  $R \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$  is a rotation matrix. From (23) and (25) it is clear that optimal  $C$  and  $A^T$  correspond to rotations of principal eigenvectors of the Gram matrix  $K$ , which is the kernel PCA solution. Hence, *kernel PCA can be viewed as a lower bound on the mutual information between  $x$  and  $y$  in nonlinear Gaussian channels*. Indeed, the bound enables us, in a principled way, to choose between different parameters of the kernel function, or to choose between competing kernels.

## 5 Optimal Kernel Functions

Optimal parameters  $\Theta$  of the kernel function  $\mathcal{K}_\Theta : |\mathcal{X}| \times |\mathcal{X}| \rightarrow \mathbb{R}$  may be obtained by maximizing the general objectives (13) or (19) or their kernelized versions. The optimization procedure may be viewed as a special case of the IM algorithm:

1. For the fixed Gram matrix  $K$ , optimize the bound w.r.t  $U, W$  (or the dual parameters  $C$  and  $A$ ).
2. For the fixed  $U, W$  (or  $C$  and  $A$ ), optimize the bound w.r.t. the parameters  $\Theta$  of  $\mathcal{K}_\Theta$ .

For non-isotropic channels (or constrained encoder-decoder pairs) this may in general result in non-trivial settings of the parameters  $\Theta$ . In what follows we describe a few special cases of optimal kernel functions for the simplest KPCA channel of Section 4.2. Our motivation here is gaining an intuition into choosing kernel parameters which maximize the bound on the mutual information (see [15], [16] for a detailed discussion of concentration properties of Gram matrices).

## 5.1 Kernel PCA Channels

As before, we assume that  $\mathbf{K} \in \mathbb{R}^{M \times M}$  is non-singular. From (24) and (25) it is clear that an alternative formulation of the information maximization problem for the KPCA case is given by

$$\max_{\Theta} \max_{\mathbf{C}, \mathbf{M}} [\text{tr} \{ \mathbf{C}^T \mathbf{K} \mathbf{C} \} - \text{tr} \{ \mathbf{M} (\mathbf{C}^T \mathbf{C} - \mathbf{I}) \}] \equiv \max_{\Theta} \sum_{i=1}^{|\mathcal{Y}|} \lambda_i(\mathcal{K}_{\Theta}) \leq \text{tr} \{ \mathbf{K} \}. \quad (26)$$

Here  $\mathbf{M} \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$  is a matrix of Lagrange multipliers imposing orthonormality constraints on  $\mathbf{C} \in \mathbb{R}^{M \times |\mathcal{Y}|}$ , and  $\lambda_i(\mathcal{K}_{\Theta})$  is the  $i^{\text{th}}$  principal component of the Gram matrix  $\mathbf{K} \in \mathbb{R}^{M \times M}$  corresponding to  $\mathcal{K}_{\Theta}$ .

In general, in order to avoid divergence it is necessary to impose norm constraints on  $\|\mathbf{K}\|_F$ . In this case it is intuitive that the worst kernel has a flat spectrum (which is the case for  $\mathbf{K} = c\mathbf{I}$ ), while the optimal kernel function results in the Gram matrix  $\mathbf{K}$  with an eigenspectrum concentrated at  $|\mathcal{Y}|$  principal components.

## Squared Exponential (RBF) Kernels

Let  $\mathcal{K}_{\gamma}(x_i, x_j; \gamma^2) \stackrel{\text{def}}{=} \exp\{-\|x_i - x_j\|^2 / (2\gamma^2)\}$  be a kernel function with the variance  $\gamma^2$ . It is clear that  $K_{ij} \leq K_{ii} = 1$  and  $\text{tr} \{ \mathbf{K} \} = M$ .

**Proposition 2:** *For Squared Exponential kernels  $\mathcal{K}_{\gamma}(x, \tilde{x}; \gamma^2)$  the optimal solution of (26) is attained in the limit of diverging variance  $\gamma^2$ .*

## Mixture Kernels

If  $\mathcal{K}_1$  and  $\mathcal{K}_2$  are positive semi-definite kernel functions then so is  $\mathcal{K}_{\alpha} \stackrel{\text{def}}{=} \alpha \mathcal{K}_1 + (1 - \alpha) \mathcal{K}_2$ , where  $\alpha \in [0, 1]$ . Here we show that  $\mathcal{K}_{\alpha}$  optimized for the mixing coefficient  $\alpha$  converges to the single best kernel component. This results from the following propositions:

**Proposition 3:** *Let  $\mathbf{L} = \text{diag}(\lambda_1, \dots, \lambda_M) \in \mathbb{R}^{M \times M}$ , where  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_M$ , and  $\mathbf{R} \in \mathbb{R}^{M \times |\mathcal{Y}|}$  such that  $\mathbf{R}^T \mathbf{R} = \mathbf{I}_{|\mathcal{Y}|}$ . Then  $\text{tr} \{ \mathbf{R}^T \mathbf{L} \mathbf{R} \} \leq \sum_{i=1}^{|\mathcal{Y}|} \lambda_i$ .*

**Proposition 4:** *Let  $\mathcal{K}_{\alpha} \stackrel{\text{def}}{=} \alpha \mathcal{K}_1 + (1 - \alpha) \mathcal{K}_2$ , where  $\alpha \in [0, 1]$ . Then the optimal solution of (26) is obtained for  $\alpha \in \{0, 1\}$ .*

The obtained result is quite intuitive. Just as a mixture distribution is generally vaguer than a distribution defined by a single mixture component, eigenspectra of mixtures of positive semi-definite matrices are generally flat. This makes them sub-optimal in the KPCA channels under the objective (26). Finally, note that the results of proposition 4 may be easily generalized to any convex mixture of fixed kernel functions.

## 6 Experiments

**Visualizing a Low-Dimensional Structure:** Let  $\mathcal{K}_{\gamma}(x_i, x_j; \gamma^2)$  be a Squared Exponential kernel with a fixed variance parameter  $\gamma^2$ . Here we show the influence of  $\gamma^2$  on the bound (22) for a simple KPCA channel (see Section 4.2). We also discuss performance of kernels with various parameter settings for visualizing an intrinsically low-dimensional structure of the data.

Figure 1 (a) shows  $M = 50$  training points  $x_i \in \mathbb{R}^2$ . The points were sampled uniformly from a semi-circular unit arc, and perturbed by a small noise  $\epsilon \sim \mathcal{N}_{\epsilon}(\mathbf{0}, 0.05\mathbf{I}_2)$ . From the construction it is clear






Figure 1: (a): training data,  $M = 50$ ,  $|\mathbf{x}| = 2$ . (b): change of  $\tilde{I}(x, y; \gamma^2)$  with an increase of the *inverse* variance  $\gamma^{-2}$  for Squared Exponential kernels,  $|y| = 1$ . *Insert*: Residuals  $\text{res}(\gamma^{-2}) \stackrel{\text{def}}{=} (M - \lambda_1(\mathcal{K}_\gamma))/M$  as a function of  $\gamma^{-2}$ .

that the (approximate) intrinsic dimension of each data point  $\mathbf{x}_i$  is defined by an angle  $\psi(\mathbf{x}_i) \in \mathbb{R}$ . One may hope that maximization of  $I(x, y)$  may result in the codes which capture useful 1d information about the angles. Since the points were sampled uniformly and the noise was small, the phase change  $\Delta\psi(\mathbf{x}_{i+1}, \mathbf{x}_i) \stackrel{\text{def}}{=} \psi(\mathbf{x}_{i+1}) - \psi(\mathbf{x}_i)$  between the neighboring points should be approximately constant. If  $y(\mathbf{x}_i)$  indeed corresponds to a (scaled) phase  $\psi(\mathbf{x}_i)$ , we can expect  $y(\mathbf{x}_i)$  to be roughly linear as a function of  $i$ .

Figure 1 (b) shows the change in  $\log \tilde{I}(x, y; \gamma^2)$  with a decrease in the variance. We see that the maximum of the bound is attained in the limit of  $\gamma^2 \rightarrow \infty$  (cf proposition 2). The inserted graph shows the relative weight of the  $M - 1$  minor eigenvalues  $\sum_{i=2}^M \lambda_i(\mathcal{K}_\gamma)/M$ , which decreases as  $\gamma^2 \rightarrow \infty$ . We therefore expect that in the specified channel the information transmission improves with the growth of  $\gamma^2$ .

To confirm that the choice of  $\gamma^2$  influences the visualization performance, we sampled  $M_1 = 75$  testing points at uniform from the same process. Figure 2 plots projections  $y_i$  of the testing points  $\mathbf{x}_i$  as a function of  $i$  (for illustration purposes, higher values of  $y_i$  are plotted in lighter colors). For a better visualization, the data was centered in the feature space (see [14]). Parameters  $\gamma^2$  of the kernel functions  $\mathcal{K}_\gamma$  were fixed at  $\gamma^2 = 0.1, 1$ , and 100, which resulted in the residuals decreasing as  $\text{res}(\gamma^2) \approx 0.809, 0.446, 0.170$ . From figure 2 (a), (b), (c) we see that as  $\gamma^2$  increases,  $y(\mathbf{x}_i)$  becomes approximately linear in  $i$ . For a high  $\gamma^2$  this indicates an approximately linear change of projection coordinates, which is characteristic of the intrinsic parameterization of the data. Finally, figure 2 (d) shows a linear projection of the testing data onto the principal eigenvector of the sample covariance  $\mathbf{S}$ . One can see that even though a part of the plot is approximately linear,  $\psi(\mathbf{x}_i)$  has an extraneous region of the constant phase towards the darker end of the line (which is easily explained by the form of the training data shown on figure 1 (a)).

**Digit Clustering:** Figure 3 shows projections of the real-world digits onto the 3 principal components of the feature covariance  $\mathbf{S}_F$  for the squared exponential kernel function  $\mathcal{K}_\gamma$ ,  $\gamma^2 = 100$ . The training data consisted of  $M = 90$  instances of digits 4, 5, and 7 (30 of each kind) with  $|\mathbf{x}| = 196$ . Prior to clustering, the patterns were centered and normalized. It is intuitive that if  $\mathbf{y}$  contains information about distances to cluster centers, it may be useful for reconstructing  $\mathbf{x}$ , so clustering is easily explainable. However, this is mainly an effect of constraints on the channel parameters (as will be discussed elsewhere).

## 7 Discussion

We have shown that optimal Gaussian decoders have a strong relation to popular dimensionality reduction techniques. One of the main contributions of this paper is a principled objective function which can be




Figure 2: Projection  $y(x_i)$  as a function of  $i$  for  $M_1 = 75$ ,  $|y| = 1$ . (a), (b), (c): RBF kernels,  $\gamma^2 = 0.1, 1, 100$ . (d): A linear projection onto the  $1^{st}$  principal eigenvector of  $S$ .




Figure 3: Clustering of 180 instances of 3 digits with RBF kernels,  $|x| = 196$ ,  $|y| = 3$ ,  $\gamma^2 = 100$ .

used to learn and compare kernel parameters. One can also envisage a similar approach based on a kernelized version of autoencoders. In particular, the bound here for the Gaussian decoders is strongly related to using a minimal least squares linear reconstruction error. However, our method is arguably more general, since it naturally generalizes to different channels, and is well founded from an information theoretic viewpoint. The result that using a linear decoder does not improve the situation much is analogous to the well-known result in autoencoders that one needs more than one hidden unit to improve on PCA [3].

Recently, Lawrence has introduced a nice way to use Gaussian processes for dimension reduction [9]. Our work is related to this, in the sense that both methods produce PCA as a limiting special case. However, the work of Lawrence is not directly related to kernel PCA, although it does indeed enable the evaluation of different kernel functions. From a practical viewpoint, our method may be computationally rather more convenient since, unlike [9], it does not require non-linear optimization to perform dimensionality reduction.

## Appendix

**Proposition 1 : Proof:** From basic properties of the mutual information (see e.g. [5]) it is easy to see that

$$I(\mathbf{s}, \mathbf{t}; r) = H(r) - H(r|\mathbf{t}, \mathbf{s}) = I(\mathbf{t}, r), \quad (27)$$

$$\begin{aligned} I(\mathbf{s}, \mathbf{t}; r) &= H(\mathbf{s}) + H(\mathbf{t}|\mathbf{s}) - H(\mathbf{s}|r) - H(\mathbf{t}|\mathbf{s}, r) \\ &= I(\mathbf{s}, r) + H(\mathbf{t}|\mathbf{s}) - H(\mathbf{t}|\mathbf{s}, r). \end{aligned} \quad (28)$$

Utilizing the chain structure and the deterministic mapping  $p(\mathbf{t}|\mathbf{s})$ , we obtain

$$p(\mathbf{t}|\mathbf{s}, r) = \frac{\delta(\mathbf{t} - \mathbf{f}(\mathbf{s}))p(r|\mathbf{t})}{\int_{\mathbf{t}} \delta(\mathbf{t} - \mathbf{f}(\mathbf{s}))p(r|\mathbf{t})} = \frac{\delta(\mathbf{t} - \mathbf{f}(\mathbf{s}))p(r|\mathbf{f}(\mathbf{s}))}{p(r|\mathbf{f}(\mathbf{s}))}, \quad (29)$$

i.e.  $p(\mathbf{t}|\mathbf{s}, r) = p(\mathbf{t}|\mathbf{s})$ . Here we used  $f(x)\delta(x - a) = \lim_{\epsilon \rightarrow 0} [f(a - \epsilon) + f(a + \epsilon)]\delta(x - a)/2$  (see e.g. [8]). Then  $H(\mathbf{t}|\mathbf{s}, r) = H(\mathbf{t}|\mathbf{s})$ , and from (27), (28) we obtain  $I(\mathbf{s}, r) = I(\mathbf{t}, r)$ . ■

**Proposition 2 : Proof:** From the definition of  $\mathcal{K}_\gamma(x, \tilde{x}; \gamma^2)$ , we get  $\forall i, \forall j. \lim_{\gamma^2 \rightarrow \infty} K_{ij} = 1$ , i.e.  $\text{rank}(\mathbf{K}) \rightarrow 1$  as  $\gamma^2 \rightarrow \infty$ . Then  $\lim_{\gamma^2 \rightarrow \infty} \sum_{i=1}^{|\mathcal{Y}|} \lambda_i(\mathcal{K}_\gamma) = \lambda_1(\mathcal{K}_\gamma) = M = \text{tr}\{\mathbf{K}\}$ , which is a global optimum of the objective in (26) independently of the size of  $|\mathcal{Y}|$ . ■

**Proposition 3 : Proof:** The proof follows straight away from the solution of the constrained optimization problem  $\tilde{\mathbf{R}} = \text{argmax}_{\mathbf{R}} \hat{I}_{\mathbf{R}}$ , where

$$\hat{I}_{\mathbf{R}} = \text{tr}\{\mathbf{R}^T \mathbf{L} \mathbf{R}\} - \text{tr}\{\mathbf{M}(\mathbf{R}^T \mathbf{R} - \mathbf{I}_{|\mathcal{Y}|})\}. \quad (30)$$

The optimization results in  $\tilde{\mathbf{R}} \in \mathbb{R}^{M \times |\mathcal{Y}|}$  corresponding to  $|\mathcal{Y}|$  principal components of  $\mathbf{L} \in \mathbb{R}^{M \times M}$ . Since  $\mathbf{L}$  is a diagonal matrix with a sorted eigenspectrum, we get  $\tilde{\mathbf{R}}^T = [\mathbf{I}_{|\mathcal{Y}|} \mathbf{0}]$ , where  $\mathbf{0} \in \mathbb{R}^{|\mathcal{Y}| \times (M - |\mathcal{Y}|)}$  is a matrix of zeros. Then  $\hat{I}_{\mathbf{R}} \leq \hat{I}_{\tilde{\mathbf{R}}} = \sum_{i=1}^{|\mathcal{Y}|} \lambda_i$ . ■

**Proposition 4 : Proof:**

Let  $\mathbf{K}_\alpha = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$ ,  $\mathbf{K}_1 = \mathbf{U}_1 \mathbf{\Lambda}_1 \mathbf{U}_1^T$ , and  $\mathbf{K}_2 = \mathbf{U}_2 \mathbf{\Lambda}_2 \mathbf{U}_2^T$  be eigenvalue decompositions of the Gram matrices corresponding to the kernel functions  $\mathcal{K}_\alpha$ ,  $\mathcal{K}_1$ , and  $\mathcal{K}_2$ . We are interested in maximizing  $\sum_{i=1}^{|\mathcal{Y}|} \lambda_i(\mathcal{K}_\alpha)$ , which may be written as

$$\begin{aligned} \sum_{i=1}^{|\mathcal{Y}|} \lambda_i(\mathcal{K}_\alpha) &= \text{tr}\{\mathbf{V}^T \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T \mathbf{V}\} \\ &= \alpha \text{tr}\{\mathbf{R}_1^T \mathbf{\Lambda}_1 \mathbf{R}_1\} + (1 - \alpha) \text{tr}\{\mathbf{R}_2^T \mathbf{\Lambda}_2 \mathbf{R}_2\}. \end{aligned} \quad (31)$$

Here columns of  $\mathbf{V} \in \mathbb{R}^{M \times |\mathcal{Y}|}$  correspond to the principal eigenvectors of  $\mathbf{K}_\alpha$ ,  $\mathbf{R}_1^T \stackrel{\text{def}}{=} \mathbf{V}^T \mathbf{U}_1 \in \mathbb{R}^{|\mathcal{Y}| \times M}$ , and  $\mathbf{R}_2^T \stackrel{\text{def}}{=} \mathbf{V}^T \mathbf{U}_2 \in \mathbb{R}^{|\mathcal{Y}| \times M}$ . From (31) and proposition 3 we get

$$\sum_{i=1}^{|\mathcal{Y}|} \lambda_i(\mathcal{K}_\alpha) \leq \alpha \sum_{i=1}^{|\mathcal{Y}|} \lambda_i^{(1)} + (1 - \alpha) \sum_{i=1}^{|\mathcal{Y}|} \lambda_i^{(2)} \leq \max_j \sum_{i=1}^{|\mathcal{Y}|} \lambda_i^{(j)}, \quad (32)$$

where  $\lambda_i^{(j)}$  corresponds to the  $i^{\text{th}}$  diagonal element of  $\mathbf{\Lambda}_j$ ,  $j \in \{1, 2\}$ . Apart from the invariant cases, the maximum is achieved for  $\alpha = 2 - \text{argmax}_{j \in \{1, 2\}} \sum_{i=1}^{|\mathcal{Y}|} \lambda_i^{(j)} \in \{0, 1\}$ . ■

## References

- [1] D. Barber and F. V. Agakov. The IM Algorithm: A Variational Approach to Information Maximization. In *NIPS*. MIT Press, 2003.

- [2] A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.
- [3] C.M. Bishop. *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford, 1995.
- [4] N. Brunel and J.-P. Nadal. Mutual Information, Fisher Information and Population Coding. *Neural Computation*, 10:1731–1757, 1998.
- [5] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- [6] N. Brunel J.-P. Nadal and N. Parga. Nonlinear feedforward networks with stochastic outputs: infomax implies redundancy reduction. *Network: Computation in Neural Systems*, 9(2):207–217, 1998.
- [7] T. S. Jaakkola and M. I. Jordan. Improving the Mean Field Approximation via the Use of Mixture Distributions. In M. I. Jordan, editor, *Learning in Graphical Models*. Kluwer Academic Publishers, 1998.
- [8] G. A. Korn and T. M. Korn. *Mathematical Handbook for Scientists and Engineers: Definitions, Theorems, and Formulas for Reference and Review*. McGraw-Hill, New York, 1968.
- [9] N. D. Lawrence. Gaussian Process Latent Variable Models for Visualization of High Dimensional Data. In *NIPS*, 2003.
- [10] R. Linsker. An Application of the Principle of Maximum Information Preservation to Linear Systems. In David Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 1. Morgan-Kaufmann, 1989.
- [11] R. Linsker. How to generate ordered maps by maximizing the mutual information between input and output signals. *Neural Computation*, 1, 1989.
- [12] D. J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [13] D. Saad and M. Opper. *Advanced Mean Field Methods Theory and Practice*. MIT Press, 2001.
- [14] B. Schoelkopf, A. Smola, and K.R. Mueller. Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation*, 10, 1998.
- [15] J. Shawe-Taylor, N. Cristianini, and J. Kandola. On the Concentration of Spectral Properties. In *NIPS*. MIT Press, 2002.
- [16] J. Shawe-Taylor and C. K. I. Williams. The Stability of Kernel Principal Components Analysis and its Relation to the Process Eigenspectrum. In *NIPS*. MIT Press, 2003.
- [17] K. Torkkola and W. M. Campbell. Mutual Information in Learning Feature Transformations. *Proc. 17th International Conf. on Machine Learning*, 2000.
- [18] C. K. I. Williams. Prediction with Gaussian Processes: From Linear Regression to Linear Prediction and Beyond. In M. I. Jordan, editor, *Learning in Graphical Models*. Kluwer Academic Publishers, 1998.