

# Maximally Informative Statistics

**David R. Wolf**

PO 8308, Austin, TX 78713-8308, USA, E-mail: drwolf@realtime.net  
Dr. Wolf is the corresponding author.

**Edward I. George**

Department of MSIS, University of Texas, Austin, TX 78712-1175, Email:  
egeorge@mail.utexas.edu

*Revision history: April 1996. Presented Bayesian Statistics 6, Valencia, 1998. Invited paper Monograph on Bayesian Methods in the Sciences, Rev. R. Acad. Sci. Exacta. Fisica. Nat. Vol. 93, No. 3, pp. 381–386, 1999. Arxiv version asserts bold vectors dropped in print.*

---

**Abstract:** In this paper we propose a Bayesian, information theoretic approach to dimensionality reduction. The approach is formulated as a variational principle on mutual information, and seamlessly addresses the notions of sufficiency, relevance, and representation. Maximally informative statistics are shown to minimize a Kullback-Leibler distance between posterior distributions. Illustrating the approach, we derive the maximally informative one dimensional statistic for a random sample from the Cauchy distribution.

**Keywords:** Bayesian Inference; Kullback-Leibler distance; Maximally informative statistics; Sufficient statistics; Mutual information; Calculus of variation.

---

## 1 Introduction

Dimensionality reduction is a fundamental goal of statistical science. In a modeling context, this is often facilitated by estimating a low dimensional quantity of interest. For example, suppose the quantities of interest are the labels of a classification of photographs of objects; of trees, children, etc. The data are the photographs, and the goal is to infer which of the several classes have been presented. In this case the data space often has dimension on the order of  $> 10^6$ , while the parameter space is a small discrete set of labels each having much lower dimension. A low dimensional summary of the photograph is then obtained as the estimate of the classification of the photograph.

In this paper, we propose a novel fully Bayesian information theoretic approach to dimensionality reduction, based on maximizing the mutual information between a statistic and a quantity of interest. The approach is formulated as a variational principle on mutual information, and it seamlessly addresses the notions of sufficiency, relevance, and representation. We refer to statistics which maximize this mutual information as *maximally informative* (MI) statistics. Such statistics are shown to minimize a Kullback-Leibler distance between posterior distributions.

The mutual information between a statistic and a quantity of interest is defined in section 2. The mutual information based variational principle for MI statistics is utilized in section 3 to derive non-variational derivative forms of the principle. In section 4 several properties of MI statistics are derived. The important result of this section is that MI statistics provide a generalization of the notion of sufficiency, because they are sensible both when they are not sufficient statistics, and when lower-than-data-dimension sufficient statistics do not exist. In section 5 we present the result that in inference the Kullback-Leibler (KL) distance is properly a functional of posterior distributions. There we find MI statistics at functional minima of a KL distance based on posterior distributions of the parameter of interest. The arguments made here suggest that the KL distance derived here is to be preferred to a maximum relative entropy distance, a fact which is not discussed in, for example, Kullback [3] or Shor [5], and numerous others. In section 6 the MI static for the location parameter of the univariate Gaussian distribution is derived, and shown to be the expected result, since in this case a one-dimensional sufficient statistic exists. In section 7 we find a one-

dimensional MI statistic for the Cauchy distribution, where a sufficiency reduction does not exist. In section 8 we discuss approximating the posterior distribution as a Gaussian and apply this technique to show that the MI statistics are then Bayes' estimators of the mean and standard deviation. In that section a contrast of the approximate MI inference approach with the Maximum Entropy method is made, and it is shown that although they agree for Gaussian likelihoods, they disagree for other distributions, with simplicity arguing in favor of the MI statistics.

## 2 The Mutual Information Between a Statistic and a Quantity of Interest

Let the data  $\mathbf{x} \in X$  be drawn according to a parameterized distribution  $P(\mathbf{x} | \boldsymbol{\theta})$ , with  $\boldsymbol{\theta} \in \Theta$ , the parameter space.  $\boldsymbol{\theta}$  itself is distributed according to the prior  $P(\boldsymbol{\theta})$ . The marginal distribution of  $\mathbf{x}$  is obtained from  $P(\mathbf{x}) = \int P(\mathbf{x} | \boldsymbol{\theta}) P(\boldsymbol{\theta}) d\boldsymbol{\theta}$ , and the posterior of  $\boldsymbol{\theta}$  given  $\mathbf{x}$  is obtained from Bayes Theorem as

$$P(\boldsymbol{\theta} | \mathbf{x}) = \frac{P(\mathbf{x} | \boldsymbol{\theta}) P(\boldsymbol{\theta})}{P(\mathbf{x})} \quad (1)$$

The quantity of interest  $\mathbf{q} = \boldsymbol{\xi}_Q(\boldsymbol{\theta})$  will be a function of  $\boldsymbol{\theta}$ , a mapping from the parameter space  $\Theta$  into some  $Q$ ,  $\boldsymbol{\xi}_Q(\cdot) : \Theta \rightarrow Q$ . It will be useful to use the Dirac delta-function  $\delta(\cdot)$  to represent the distribution of  $\mathbf{q}$  as

$$\begin{aligned} P(\mathbf{q} | \boldsymbol{\theta}) &:= P(\{\boldsymbol{\theta} : \mathbf{q} = \boldsymbol{\xi}_Q(\boldsymbol{\theta})\} | \boldsymbol{\theta}) \\ &= \delta(\mathbf{q} - \boldsymbol{\xi}_Q(\boldsymbol{\theta})) \end{aligned} \quad (2)$$

$$= \prod_{i=1}^{k_q} \delta(q_i - \xi_{Q,i}(\boldsymbol{\theta})), \quad (3)$$

where  $\delta(\mathbf{z}(\cdot)) = \prod_i \delta(z_i(\cdot))$ . Note that (2) may be seen directly by using Bayes' theorem to expand  $P(\mathbf{q}, \boldsymbol{\theta})$  as  $P(\mathbf{q} | \boldsymbol{\theta}) P(\boldsymbol{\theta})$ , integrating that over  $\mathbf{q}$ , which must produce  $P(\boldsymbol{\theta})$ , and noting that because the support of  $P(\mathbf{q} | \boldsymbol{\theta})$  is the unique  $\mathbf{q}$  such that  $\mathbf{q} = \boldsymbol{\xi}_Q(\boldsymbol{\theta})$  ( $\boldsymbol{\theta}$  is specified),  $P(\mathbf{q} | \boldsymbol{\theta})$  must therefore be the Dirac delta function. The distribution of  $\mathbf{q}$  given the data  $\mathbf{x}$ , may be written using (1) and (3) as

$$P(\mathbf{q} | \mathbf{x}) = \int P(\mathbf{q} | \boldsymbol{\theta}) P(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta} \quad (4)$$

A statistic  $\mathbf{r} = \boldsymbol{\xi}_R(\mathbf{x})$  will be a function of  $\mathbf{x}$ , a mapping from the data space  $X$  into some  $R$ ,  $\boldsymbol{\xi}_R(\cdot) : X \rightarrow R$ . Again using the delta notation, the distribution of the statistic given data is

$$P(\mathbf{r} | \mathbf{x}) = \delta(\mathbf{r} - \boldsymbol{\xi}_R(\mathbf{x})) \quad (5)$$

$$= \prod_{i=1}^{k_r} \delta(r_i - \xi_{R,i}(\mathbf{x})) \quad (6)$$

The joint distribution of the statistic  $\mathbf{r}$  and the quantity of interest  $\mathbf{q}$ , conditioned on the data  $\mathbf{x}$  is

$$P(\mathbf{r}, \mathbf{q} | \mathbf{x}) = P(\mathbf{r} | \mathbf{x}) P(\mathbf{q} | \mathbf{x}) \quad (7)$$

(since  $\mathbf{r} = \boldsymbol{\xi}_R(\mathbf{x})$  is specified once  $\mathbf{x}$  is known, making  $P(\mathbf{r} | \mathbf{x}, \mathbf{q}) = P(\mathbf{r} | \mathbf{x})$ ), and the unconditional joint distribution is

$$P(\mathbf{r}, \mathbf{q}) = \int P(\mathbf{r} | \mathbf{x}) P(\mathbf{q} | \mathbf{x}) P(\mathbf{x}) d\mathbf{x}. \quad (8)$$

Finally, we define the mutual information between a statistic  $\boldsymbol{\xi}_R(\cdot)$  and a quantity of interest  $\boldsymbol{\xi}_Q(\cdot)$  as

$$M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot)) = \int \int P(\mathbf{r}, \mathbf{q}) \log \left( \frac{P(\mathbf{r}, \mathbf{q})}{P(\mathbf{r}) P(\mathbf{q})} \right) d\mathbf{q} d\mathbf{r} \quad (9)$$

This mutual information is the Kullback-Leibler distance between the joint distribution  $P(\mathbf{r}, \mathbf{q})$  and the marginal product  $P(\mathbf{r})P(\mathbf{q})$  corresponding to independence between  $\mathbf{r}$  and  $\mathbf{q}$ . Note that this Kullback-Leibler distance is different from the Kullback-Leibler distance mentioned in the introduction (and seen later in section 5). A major contribution of this paper is the demonstration of how these two Kullback-Leibler distances are related.

### 3 MI Statistics and the Variational Principle

We now define the maximally informative (MI) statistic.

*Let  $S = \{\boldsymbol{\xi}_R(\cdot)\}$  be a set of statistics under consideration. A MI statistic for a quantity of interest  $\boldsymbol{\xi}_Q(\cdot)$  is any statistic  $\boldsymbol{\xi}_R(\cdot)$  from  $S$  maximizing the mutual information  $M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot))$  between the statistic and the quantity of interest.*

The following variational principle can be used to obtain an MI statistic. Let  $\frac{\delta}{\delta f(\cdot)}$  denote the functional derivative with respect to  $f(\cdot)$ .

*Choose  $\boldsymbol{\xi}_R(\cdot)$  from  $S$  such that  $\frac{\delta M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot))}{\delta \boldsymbol{\xi}_R(\cdot)} = 0$  and  $\frac{\delta^2 M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot))}{\delta \boldsymbol{\xi}_R(\cdot)^2}$  is negative semidefinite, i.e. so that  $\boldsymbol{\xi}_R(\cdot)$  maximizes the information between itself and  $\boldsymbol{\xi}_Q(\cdot)$ , the quantity of interest. If possible, choose the global maximum.*

Note that MI statistics in  $S$  may occur on the boundary of  $S$ . This may be a case of interest, which occurs when constraints are imposed on the statistics, and may be handled with a trivial modification. Note also that the space  $S$  of statistics may be constrained to contain only low-dimensional statistics, in order to force a dimensionality reduction of the data.

We now demonstrate the variational principle for MI statistics. The argument proceeds by varying (see, for example [1] for the variational calculus) the mutual information of (9) with respect to the statistic function  $\boldsymbol{\xi}_R(\cdot)$  of dimension  $k_r$ , i.e.  $\boldsymbol{\xi}_R(\cdot) = (\xi_{r,1}(\cdot), \dots, \xi_{r,k_r}(\cdot))$ . We now proceed to substitute  $\boldsymbol{\xi}_R(\mathbf{x}) = \boldsymbol{\xi}_R^0(\mathbf{x}) + \epsilon \boldsymbol{\eta}(\mathbf{x})$  in (9), and take the derivative with respect to  $\epsilon$ .

Assuming appropriate regularity conditions, we have

$$\begin{aligned} \partial_\epsilon M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot)) &= \int \int \left[ \partial_\epsilon P(\mathbf{r}, \mathbf{q}) \log \left( \frac{P(\mathbf{r}, \mathbf{q})}{P(\mathbf{r})P(\mathbf{q})} \right) \right. \\ &\quad \left. + P(\mathbf{r}) \partial_\epsilon P(\mathbf{q} | \mathbf{r}) \right] d\mathbf{q} d\mathbf{r} \end{aligned} \quad (10)$$

$$= \int \int \partial_\epsilon P(\mathbf{r}, \mathbf{q}) \log \left( \frac{P(\mathbf{r}, \mathbf{q})}{P(\mathbf{r})P(\mathbf{q})} \right) d\mathbf{q} d\mathbf{r}, \quad (11)$$

where simplification from (10) to (11) occurs because probability is conserved. Utilizing (7) we find

$$P(\mathbf{r}, \mathbf{q}) = \int \delta(\mathbf{r} - \boldsymbol{\xi}_R(\mathbf{x})) P(\mathbf{q} | \mathbf{x}) P(\mathbf{x}) d\mathbf{x} \quad (12)$$

Taking the derivative of (12) with respect to  $\epsilon$  yields

$$\partial_\epsilon P(\mathbf{r}, \mathbf{q}) = \sum_{j=1}^{k_r} \delta'(r_j - \xi_{R,j}(\mathbf{x})) \eta_j(\mathbf{x}) \prod_{i \neq j} \delta(r_i - \xi_{R,i}(\mathbf{x})) P(\mathbf{q} | \mathbf{x}) P(\mathbf{x}) d\mathbf{x} \quad (13)$$

Note that because  $\boldsymbol{\eta}$  is arbitrary, we may choose it to simplify as needed.

We proceed by considering  $k_r$  choices of  $\boldsymbol{\eta}$ . Label the choices by  $m \in \{1, \dots, k_r\}$ , and on choice  $m$  take the components of  $\boldsymbol{\eta}$  as follows:

$$\eta_\ell(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_c), \quad (\ell = m) \quad (14)$$

$$\eta_\ell(\mathbf{x}) = 0, \quad (\ell \neq m) \quad (15)$$

where  $\mathbf{x}_c$  is any data point we may choose. The condition that the mutual information is extremal then becomes the statement that for all  $\mathbf{x}_c$  and  $i \in \{1, \dots, k_r\}$ .

$$\partial_\epsilon M(\boldsymbol{\xi}_R(\cdot), \boldsymbol{\xi}_Q(\cdot)) |_{\epsilon=0} = 0 \quad (16)$$

$$= \int \int \delta'(r_i - \xi_{R,i}^0(\mathbf{x}_c)) \Pi_{i \neq j} \delta(r_i - \xi_{R,i}^0(\mathbf{x}_c)) \\ \times P(\mathbf{q} | \mathbf{x}_c) \log \left( \frac{P(\mathbf{r}, \mathbf{q})}{P(\mathbf{r})P(\mathbf{q})} \right) d\mathbf{q} d\mathbf{r} \quad (17)$$

Integrating (17) by parts with respect to  $\mathbf{r}$  (dropping both the “0” superscript and subscript “c”, since there is no distinction to be made at this point) yields the condition that for all  $\mathbf{x}$

$$\int P(\mathbf{q} | \mathbf{x}) \partial_{\mathbf{r}} \log \left( \frac{P(\mathbf{r}, \mathbf{q})}{P(\mathbf{r})P(\mathbf{q})} \right) |_{\mathbf{r}=\boldsymbol{\xi}_R(\mathbf{x})} d\mathbf{q} = 0 \quad (18)$$

where derivatives with respect to vectors are gradients (vectors of derivatives). The form from which the theorems of the next section are proven, is found by rewriting (18) as

$$\int \frac{P(\mathbf{q} | \mathbf{x})}{P(\mathbf{q} | \mathbf{r})} \partial_{\mathbf{r}} P(\mathbf{q} | \mathbf{r}) |_{\mathbf{r}=\boldsymbol{\xi}_R(\mathbf{x})} d\mathbf{q} = 0 \quad (19)$$

#### 4 MI Statistics and Sufficiency

Now we prove several important properties concerning MI statistics. The first property is the intuitively obvious property that *data is a MI statistic*. The second property is that *any sufficient statistic is a MI statistic*. Finally, we note that *MI statistics are not necessarily sufficient statistics*.

**Theorem 1** . *Any 1-1 function of data is a MI statistic of the quantity of interest*

Proof: Let  $\xi_R(\cdot)$  be the identity so that  $\xi_R(\mathbf{x}) = \mathbf{x}$  in (19). The fraction in that equation is then 1, and the derivative integrates to zero because probability is conserved. Having  $\xi_R(\mathbf{x})$  any invertible function changes nothing as any value of it determines  $\mathbf{x}$ .

**Theorem 2** . *Any sufficient statistic for the quantity of interest is a MI statistic of the quantity of interest*

Proof: Note that, using the definition of  $\xi_R(\mathbf{x})$  being a sufficient statistic, the ratio in (19) is one - the posterior distribution of the quantity of interest given the data  $\mathbf{x}$  is the same as the posterior distribution of the quantity of interest given the sufficient statistic  $\xi_R(\mathbf{x})$ . The derivative then integrates to zero because probability is conserved.

(Note that in both Theorems 1 and 2 the Hessian condition of the MI inference variational principle is easily established since then the extremum of the mutual information is easily seen to be a local maximum. Otherwise, one must check convexity.)

Although it is true that any sufficient statistic is a MI statistic, the converse is false. In problems (of data dimension greater than one) where a lower-than-data dimension sufficient statistic does not exist, there will exist a lower-than-data dimension statistic which is MI but not sufficient. Thus, the class of maximally informative statistics contains the sufficient statistics, but is broader. MI statistics need not provide all of the available information about the underlying quantity of interest. For example, as we show in Section 7, such a lower-than-data dimension MI statistic can be obtained for the Cauchy distribution where a lower-than-data dimension sufficient statistic is a-priori unavailable. In this manner, MI statistics seamlessly address *relevance* to the consumer of the information because it is about some *relevant quantity of interest* that MI statistics are maximally informative.

## 5 MI Statistics and the KL Distance

Equation (19) may be rewritten as

$$\partial_{\mathbf{r}} \left[ \int P(\mathbf{q} | \mathbf{x}) \log \left( \frac{P(\mathbf{q} | \mathbf{x})}{P(\mathbf{q} | \mathbf{r})} \right) d\mathbf{q} \right] \Big|_{\mathbf{r}=\xi_R(\mathbf{x})} = 0 \quad (20)$$

which, along with the curvature condition, states that

**Theorem 3.** *The Kullback-Leibler distance between the posterior distribution conditioned on the statistic and the posterior distribution conditioned on the data is minimized by a MI statistic.*

Again, note that MI statistics for the quantity of interest are generally not sufficient statistics for the quantity of interest. Indeed, rather than making the Kullback-Leibler distance zero, as in the case of sufficient statistics, MI statistics are found at local minima of the Kullback-Liebler distance - viewed as a functional of the statistic. This demonstrates how the approach of this paper generalizes that performed by Lindley [4].

## 6 MI Statistics for the Gaussian distribution

This section details the inference of the one-dimensional MI statistic for the one-dimensional Gaussian distribution. We take the position parameter of the Gaussian to be  $q$ , and the the goal is to find  $\xi_R(\mathbf{x})$  so that (19) holds. From there note that the calculation of  $P(q | r)$  and  $P(q | \mathbf{x})$  is necessary, and by Bayes' theorem therefore it is necessary to find  $P(r | q)$ , which may be written as

$$\begin{aligned} P(r | q) &= \int P(r|q, \mathbf{x}) P(\mathbf{x} | q) d\mathbf{x} \\ &= \int P(r|\mathbf{x}) P(\mathbf{x} | q) d\mathbf{x} \\ &= \int \delta(r - \xi_R(\mathbf{x})) \prod_{i=1}^N \frac{e^{-(x_i - q)^2 / 2\sigma^2}}{\sqrt{2\pi}\sigma} d\mathbf{x} \end{aligned} \quad (21)$$

The ansatz  $\xi_R(\mathbf{x}) = \sum_{i=1}^N \lambda_i x_i$  is useful (and not restrictive since the  $\lambda_i$ 's are implicitly only restricted to be functions of  $\mathbf{x}$ ), and making the changes of



variables  $y_i = \lambda_i x_i$  followed by  $u_i = \lambda_i q$  in (21) yields a form which may immediately be recognized as the convolution of  $N$  Gaussians with means  $\mu_i = \lambda_i q$  and standard deviations  $\sigma_i = \lambda_i \sigma$  respectively,

$$P(r | q) = \int \delta\left(r - \sum_{i=1}^N \lambda_i q - \sum_{i=1}^N u_i\right) \prod_{i=1}^N \frac{e^{-u_i^2/2(\sigma\lambda_i)^2}}{\sqrt{2\pi}(\lambda_i\sigma)} d\mathbf{u}. \quad (22)$$

This has the solution

$$P(r | q) = \phi(0, \sigma')(r - \sum_{i=1}^N \lambda_i q) \quad (23)$$

where  $\sigma' = \sigma \sqrt{\sum_{i=1}^N \lambda_i^2}$  and  $\phi(\cdot, \cdot)(\cdot)$  is the Gaussian density

$$\phi(\mu, \sigma)(z) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(z-\mu)^2/2\sigma^2}. \quad (24)$$

Finally, inserting this result into Bayes' theorem with uniform prior to find the posterior distribution of  $q$  conditioned on  $r$  yields

$$P(q | r) = S \phi(0, \sigma')(r - \sum_{i=1}^N \lambda_i q) \quad (25)$$

where  $S := \sum_{i=1}^N \lambda_i$ .

The calculation for  $P(q | \mathbf{x})$  is similar with the result is that

$$P(q | \mathbf{x}) = \phi(\bar{\mathbf{x}}, \frac{\sigma}{\sqrt{N}})(q) \quad (26)$$

where  $\bar{\mathbf{x}} := \frac{1}{N} \sum_{i=1}^N x_i$ . From the forms of (25) and (26) it is clear that not only will the integrand of (20) (that equation equivalent to (19)) be minimized, but that it will be zero, if all  $\lambda_i = 1/N$  is chosen. This of course is the expected result since  $\xi_R(x) = \sum_{i=1}^N x_i/N$  is a sufficient statistic for  $q$  when  $\sigma$  is known.

Alternatively, to satisfy that the calculation indicated in (19) is successful at finding the expected result, continue by taking (25) and (26) and substituting them into (19) to find after some simplification the equation which must be satisfied by  $\xi_R$

$$0 = \left[ \int (r - q \sum_{i=1}^N \lambda_i) e^{-(\bar{\mathbf{x}} - q)^2/2(\sigma/\sqrt{N})^2} dq \right] \Big|_{r=\xi_R(\mathbf{x})}. \quad (27)$$

This has the unique solution  $\xi_R(\mathbf{x}) = \bar{\mathbf{x}}$  when the arbitrary scale of the inferred statistic is fixed by setting  $1 = \sum_{i=1}^N \lambda_i$ . To conclude this section, the procedure culminating in (20) or (19) of finding MI statistics has been shown to produce the expected known result for the Gaussian case. The next section approaches the Cauchy distribution case for lower than data dimension statistics, where there is no sufficient statistic available and the result is novel.

## 7 MI Statistics for the Cauchy Distribution

This section outlines the inference of the one-dimensional MI statistic for the one dimensional Cauchy distribution. The detailed steps may be taken similarly to those of the last section but taking the Cauchy distribution instead of the Gaussian distribution. Take the position parameter of the Cauchy to be  $q$ , and the the goal is to find  $\xi_R(\mathbf{x})$  so that (19) holds. As in the last section it is necessary to determine both  $P(q | r)$  and  $P(q | \mathbf{x})$ . Assuming the same ansatz that  $\xi_R(\mathbf{x}) = \sum_i \lambda_i x_i$ , the necessary convolutions may be carried out with the use of the Fourier convolution theorem, with the results that

$$P(r | q) = \frac{S}{\pi(S^2 + (r - qS)^2)}, \quad (28)$$

$$P(q | r) = \frac{S}{\pi(S^2 + (r - qS)^2)}, \quad (29)$$

and

$$P(q | \mathbf{x}) \propto \prod_{i=1}^n \frac{1}{\pi(1 + (x_i - q)^2)} \quad (30)$$

where  $S := \sum_{i=1}^N \lambda_i$ . Substituting (28), (29), and (30) into (19) yields the equation that must be solved for  $\xi_R(\mathbf{x})$

$$0 = \left[ \int \left( \prod_{i=1}^n \frac{1}{\pi(1 + (x_i - q)^2)} \right) \frac{r/S - q}{1 + (r/S - q)^2} dq \right] \Big|_{r=\xi_R(\mathbf{x})}. \quad (31)$$

Rewriting this equation in more suggestive terms, while taking the scale  $S = 1$ , gives the result as an implicit equation for  $\xi_R(\mathbf{x})$ ,

$$\xi_R(\mathbf{x}) = \int q P(q | (\mathbf{x}, \xi_R(\mathbf{x}))) dq. \quad (32)$$

The form of the result (32) says that  $\xi_R(\mathbf{x})$  is the posterior mean of  $q$  given the data *and itself as an additional observation*. This form also suggests that  $\xi_R(\mathbf{x})$  could be the posterior mean of  $q$  given the data. However, this is not the case, as a check using the posterior moment forms derived in [6] immediately shows. Further, assuming a value for  $\xi_R(\mathbf{x})$  on the right-hand side of (32) allows that to be computed in closed form using the results of [6]. This finally yields that the left-hand side is a rational function of the right-hand side, a fixed point equation which may be solved by standard iterative methods. Other checks immediately show that the solution is not the maximum likelihood solution, nor the median.

To conclude this section, the one-dimensional MI statistic for the Cauchy distribution position parameter has been found as the posterior mean of the position parameter of the Cauchy distribution given the data and the MI statistic, and this statistic is different from the Bayes' estimator which is the posterior mean given the data only.

## 8 Approximate MI Inference and Bayes Estimators

In many cases of interest, if not in all cases of relevance with high dimensional data, the convolutions that appear similarly to those in (28) etc. will be quite impossible to do in closed form, and probably in a practical sense will even be numerically intractable. However, there is an approach that may be taken which does some harm to a fully rigorous Bayesian approach, but which may be necessary. The idea that is applicable in these cases of difficulty is to directly take  $P(\mathbf{q} | \mathbf{r})$  in (20) to be Gaussian with parameters  $r = \boldsymbol{\xi}_R(\mathbf{x}) = (\boldsymbol{\mu}(\mathbf{x}), \sigma(\mathbf{x}))$ . The approximate MI approach just outlined is applied below to finding the approximate MI statistics  $(\boldsymbol{\mu}(\mathbf{x}), \sigma(\mathbf{x}))$ . The approximate MI approach is then contrasted with an alternative approach using the KL distance inverted from that of (20), one that resembles Maximum Entropy inference. The results of this section hold for any likelihood, as will become apparent.

Take an arbitrary one-dimensional parameterized likelihood parameterized by  $q$  (i.e. with  $q$  the parameter of interest). Parameterize the inferred distribution  $P(q | \mathbf{r})$  of (20) as (see (24))

$$P(q | \mathbf{r} = (\mu, \sigma)) = \phi(\mu, \sigma)(q). \quad (33)$$

Equations (20) and (33) imply that the MI statistic is

$$\begin{aligned}\mu &= \int q P(q | \mathbf{x}) dq \\ \sigma^2 &= \int (q - \mu)^2 P(q | \mathbf{x}) dq\end{aligned}\tag{34}$$

These quantities are the Bayes' estimators for the mean and standard deviation of the distribution.

If, on the other hand, the inverted form of the KL distance is taken, as it often is in many of the cases we have observed, the statistic  $\mu$  is

$$\mu = \frac{\int q \phi(\mu, \sigma)(q) \log(P(q | \mathbf{x})) dq}{\int \phi(\mu, \sigma)(q) \log(P(q | \mathbf{x})) dq}\tag{35}$$

which, along with another non-linear equation for  $\sigma$ , is a complicated non-linear system to be solved for  $\mathbf{r} = (\mu, \sigma)$ .

Note that when the likelihood  $P(\mathbf{x} | \mathbf{q})$  is Gaussian these two approximate approaches produce the same statistic, the posterior mean and standard deviation; but for the Cauchy likelihood, for example, this is not the case, with necessity to solve the complicated nonlinear system. In contrast, the approximate MI inference technique always produces the posterior Bayes' moment estimators.

The difference between the forms of the approximate MI statistics and the inverted-KL statistics appearing in (34) and (35) respectively makes it clear that one needs a good first-principles approach to the KL distance.

## 9 Conclusion

We have formulated the mutual information based variational principle for statistical inference, a fully Bayesian approach to inference, defined MI statistics for a quantity of interest, shown how the principle may be reformulated as a minimal KL distance principle based on posterior distributions, and demonstrated how inference proceeds, when lower-than-data dimension sufficient statistics are absent, using the Cauchy distribution. Finally, an approximate approach to the inference of MI statistics was discussed, and the relationship of the resulting statistics to Bayes' estimators and the Maximum Entropy version of the same approximation was noted.

## 10 Acknowledgements

Thanks go to the Data Understanding Group at NASA Ames for their lively and interactive critique, friendship, mentoring, and support. This paper was improved by comments from Dr. Jeremy Frank and Hal Duncan, both of NASA. Much thanks to Tony O'Hagan for detailed comments. This work was supported by NASA Center for Excellence in Information Technology contract NAS-214217. This work was supported by NSF grant DMS-98.03756 and Texas ARP grants 003658.452 and 003658.690.

## References

- [1] George Arfken. *Mathematical Methods for Physicists*. Academic Press, Inc., London, 1985.
- [2] George E. P. Box and George C. Tiao. *Bayesian Inference in Statistical Analysis*. Wiley, NY, 1973.
- [3] Solomon Kullback. *Information Theory and Statistics*. John Wiley and Sons, Inc., New York, 1959.
- [4] D. V. Lindley. On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27:986–1005, 1961.
- [5] John E. Shore and Rodney V. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, January 1979.
- [6] David R. Wolf. Posterior moments of the cauchy distribution. In *Maximum Entropy and Bayesian Methods*, eds. W. van der Linden et. al., Kluwer Academic, Dordrecht, Netherlands, 1998.