

Statistical Distances and Their Role in Robustness

Marianthi Markatou, Yang Chen, Georgios Afendras, and Bruce G. Lindsay

1 Introduction

Distance measures play a ubiquitous role in statistical theory and thinking. However, within the statistical literature this extensive role has too often been played out behind the scenes, with other aspects of the statistical problems being viewed as more central, more interesting, or more important.

The behind the scenes role of statistical distances shows up in estimation, where we often use estimators based on minimizing a distance, explicitly or implicitly, but rarely studying how the properties of the distance determine the properties of the estimators. Distances are also prominent in goodness-of-fit (GOF) but the usual question we ask is how powerful is our method against a set of interesting alternatives not what aspects of the difference between the hypothetical model and the alternative are we measuring?

How can we interpret a numerical value of a distance? In goodness-of-fit we learn about Kolmogorov-Smirnov and Cramér-von Mises distances but how do these compare with each other? How can we improve their properties by looking at what statistical properties are they measuring?

Past interest in distance functions between statistical populations had a two-fold purpose. The first purpose was to prove existence theorems regarding some

M. Markatou (✉) • G. Afendras
Department of Biostatistics, SPHHP and Jacobs School of Medicine and Biomedical Sciences,
University at Buffalo, Buffalo, NY, USA
e-mail: markatou@buffalo.edu; gafendra@buffalo.edu

Y. Chen
Department of Biostatistics, University at Buffalo, Buffalo, NY 14214, USA
e-mail: ychen57@buffalo.edu

B.G. Lindsay
Department of Statistics, The Pennsylvania State University, University Park, PA 16820, USA
e-mail: bgl@psu.edu

© Springer International Publishing AG 2017
D.-G. Chen et al. (eds.), *New Advances in Statistics and Data Science*,
ICSA Book Series in Statistics, https://doi.org/10.1007/978-3-319-69416-0_1

optimum solutions in the problem of statistical inference. Wald (1950) in his book on statistical decision functions gave numerous definitions of distance between two distributions which he primarily introduced for the purpose of creating decision functions. In this context, the choice of the distance function is not entirely arbitrary, but it is guided by the nature of the mathematical problem at hand.

Statistical distances are defined in a variety of ways, by comparing distribution functions, density functions or characteristic functions and moment generating functions. Further, there are discrete and continuous analogues of distances based on comparing density functions, where the word “density” is used to also indicate probability mass functions. Distances can also be constructed based on the divergence between a nonparametric probability density estimate and a parametric family of densities. Typical examples of distribution-based distances are the Kolmogorov-Smirnov and Cramér-von Mises distances. A separate class of distances is based upon comparing the empirical characteristic function with the theoretical characteristic function that corresponds, for example, to a family of models under study, or by comparing empirical and theoretical versions of moment generating functions.

In this paper we proceed to study in detail the properties of some statistical distances, and especially the properties of the class of chi-squared distances. We place emphasis on determining the sense in which we can offer meaningful interpretations of these distances as measures of statistical loss. Section 2 of the paper discusses the definition of a statistical distance in the discrete probability models context. Section 3 presents the class of chi-squared distances and their statistical interpretation again in the context of discrete probability models. Section 3.3 discusses metric and other properties of the symmetric chi-squared distance. One of the key issues in the construction of model misspecification measures is that allowance should be made for the scale difference between observed data and a hypothesized model continuous distribution. To account for this difference in scale we need the distance measure to exhibit discretization robustness, a concept that is discussed in Sect. 4.1. To achieve discretization robustness we need sensitive distances, and this requirement dictates a balance of sensitivity and statistical noise. Various strategies that deal with this issue are discussed in the literature and we briefly discuss them in Sect. 4.1. A flexible class of distances that allows the user to adjust the noise/sensitivity trade-off is the kernel smoothed distances upon which we briefly remark on in Sect. 4. Finally, Sect. 5 presents further discussion.

2 The Discrete Setting

Procedures based on minimizing the distance between two density functions express the idea that a fitted statistical model should summarize reasonably well the data and that assessment of the adequacy of the fitted model can be achieved by using the value of the distance between the data and the fitted model.

The essential idea of density-based minimum distance methods has been presented in the literature for quite some time as it is evidenced by the method of minimum chi-squared (Neyman 1949). An extensive list of minimum chi-squared

methods can be found in Berkson (1980). Matusita (1955) and Rao (1963) studied minimum Hellinger distance estimation in discrete models while Beran (1977) was the first to use the idea of minimum Hellinger distance in continuous models.

We begin within the discrete distribution framework so as to provide the clearest possible focus for our interpretations. Thus, let $\mathcal{T} = \{0, 1, 2, \dots, T\}$, where T is possibly infinite, be a discrete sample space. On this sample space we define a true probability density $\tau(t)$, as well as a family of densities $\mathcal{M} = \{m_\theta(t) : \theta \in \Theta\}$, where Θ is the parameter space. Assume we have independent and identically distributed random variables X_1, X_2, \dots, X_n producing the realizations x_1, x_2, \dots, x_n from $\tau(\cdot)$. We record the data as $d(t) = n(t)/n$, where $n(t)$ is the number of observations in the sample with value equal to t .

Definition 1 We will say that $\rho(\tau, m)$ is a statistical distance between two probability distributions with densities τ, m if $\rho(\tau, m) \geq 0$, with equality if and only if τ and m are the same for all statistical purposes.

Note that we do not require symmetry or the triangle inequality, so that $\rho(\tau, m)$ is not formally a metric. This is not a drawback as well known distances, such as Kullback-Leibler, are not symmetric and do not satisfy the triangle inequality.

We can extend the definition of a distance between two densities to that of a distance between a density and a class of densities as follows.

Definition 2 Let \mathcal{M} be a given model class and τ be a probability density that does not belong in the model class \mathcal{M} . Then, the distance between τ and \mathcal{M} is defined as

$$\rho(\tau, \mathcal{M}) = \inf_{m \in \mathcal{M}} \rho(\tau, m),$$

whenever the infimum exists. Let $m_{\text{best}} \in \mathcal{M}$ be the best fitting model, then

$$\rho(\tau, m_{\text{best}}) \triangleq \rho(\tau, \mathcal{M}).$$

We interpret $\rho(\tau, m)$ or $\rho(\tau, \mathcal{M})$ as measuring the “lack-of-fit” in the sense that larger values of $\rho(\tau, m)$ mean that the model element m is a worst fit to τ for our statistical purposes. Therefore, we will require $\rho(\tau, m)$ to indicate the worst mistake that we can make if we use m instead of τ . The precise meaning of this statement will be obvious in the case of the total variation distance, as we will see that the total variation distance measures the error, in probability, that is made when m is used instead of τ .

Lindsay (1994) studied the relationship between the concepts of efficiency and robustness for the class of f - or ϕ -divergences in the case of discrete probability models and defined the concept of Pearson residuals as follows.

Definition 3 For a pair of densities τ, m define the Pearson residual by

$$\delta(t) = \frac{\tau(t)}{m(t)} - 1, \tag{1}$$

with range the interval $[-1, \infty)$.

This residual has been used by Lindsay (1994), Basu and Lindsay (1994), Markatou (2000, 2001), and Markatou et al. (1997, 1998) in investigating the robustness of the minimum disparity and weighted likelihood estimators respectively. It also appears in the definition of the class of power divergence measures defined by

$$\begin{aligned}\rho(\tau, m) &= \frac{1}{\lambda(\lambda + 1)} \sum \tau(t) \left\{ \left(\frac{\tau(t)}{m(t)} \right)^\lambda - 1 \right\} \\ &= \frac{1}{\lambda(\lambda + 1)} \sum m(t) \{ (1 + \delta(t))^{\lambda+1} - 1 \}.\end{aligned}$$

For $\lambda = -2, -1, -1/2, 0$ and 1 one obtains the well-known Neyman's chi-squared (divided by 2) distance, Kullback-Leibler divergence, twice-squared Hellinger distance, likelihood disparity and Pearson's chi-squared (divided by 2) distance respectively. For additional details see Lindsay (1994) and Basu and Lindsay (1994).

A special class of distance measures we are particularly interested in is the class of chi-squared measures. In what follows we discuss in detail this class.

3 Chi-Squared Distance Measures

We present the class of chi-squared disparities and discuss their properties. We offer loss function interpretations of the chi-squared measures and show that Pearson's chi-squared is the supremum of squared Z-statistics while Neyman's chi-squared is the supremum of squared t -statistics. We also show that the symmetric chi-squared is a metric and offer a testing interpretation for it.

We start with the definition of a generalized chi-squared distance between two densities τ, m .

Definition 4 Let $\tau(t), m(t)$ be two discrete probability distributions. Then, define the class of generalized chi-squared distances as

$$\chi_a^2(\tau, m) = \sum \frac{[\tau(t) - m(t)]^2}{a(t)},$$

where $a(t)$ is a probability density function.

Notice that if we restrict ourselves to the multinomial setting and choose $\tau(t) = d(t)$ and $a(t) = m(t)$, the resulting chi-squared distance is Pearson's chi-squared statistic. Lindsay (1994) studied the robustness properties of a version of $\chi_a^2(\tau, m)$ by taking $a(t) = [\tau(t) + m(t)]/2$. The resulting distance is called *symmetric chi-squared*, and it is given as

$$S^2(\tau, m) = \sum \frac{2[\tau(t) - m(t)]^2}{\tau(t) + m(t)}.$$

The chi-squared distance is symmetric because $S^2(\tau, m) = S^2(m, \tau)$ and satisfies the triangle inequality. Thus, by definition it is a proper metric, and there is a strong dependence of the properties of the distance on the denominator $a(t)$. In general we can use as a denominator $a(t) = \alpha\tau(t) + \bar{\alpha}m(t)$, $\bar{\alpha} = 1 - \alpha$, $\alpha \in [0, 1]$. The so defined distance is called blended chi-squared (Lindsay 1994).

3.1 Loss Function Interpretation

We now discuss the loss function interpretation of the aforementioned class of distances.

Proposition 1 *Let τ, m be two discrete probabilities. Then*

$$\rho(\tau, m) = \sup_h \frac{\{\mathbb{E}_\tau(h(X)) - \mathbb{E}_m(h(X))\}^2}{\text{Var}_a(h(X))},$$

where $a(t)$ is a density function, and $h(X)$ has finite second moment.

Proof Let h be a function defined on the sample space. We can prove the above statement by looking at the equivalent problem

$$\sup_h \{\mathbb{E}_\tau(h(X)) - \mathbb{E}_m(h(X))\}^2, \quad \text{subject to } \text{Var}_a(h(X)) = 1.$$

Note that the transformation from the original problem to the simpler problem stated above is without loss of generality because the first problem is scale invariant, that is, the functions \hat{h} and $c\hat{h}$ where c is a constant give exactly the same values. In addition, we have location invariance in that $h(X)$ and $h(X) + c$ give again the same values, and symmetry requires us to solve

$$\sup_h \{\mathbb{E}_\tau(h(X)) - \mathbb{E}_m(h(X))\}, \quad \text{subject to } \sum h^2(t)a(t) = 1.$$

To solve this linear problem with its quadratic constraint we use Lagrange multipliers. The Lagrangian is given as

$$L(t) = \sum h(t)(\tau(t) - m(t)) - \lambda \left\{ \sum h^2(t)a(t) - 1 \right\}.$$

Then

$$\frac{\partial}{\partial h} L(t) = 0, \quad \text{for each value of } t,$$

is equivalent to

$$\tau(t) - m(t) - 2\lambda h(t)a(t) = 0, \quad \forall t,$$

or

$$\widehat{h}(t) = \frac{\tau(t) - m(t)}{2\lambda a(t)}.$$

Using the constraint we obtain

$$\sum \frac{[\tau(t) - m(t)]^2}{4\lambda^2 a(t)} = 1 \Rightarrow \widehat{\lambda} = \frac{1}{2} \left\{ \sum \frac{[\tau(t) - m(t)]^2}{a(t)} \right\}^{1/2}.$$

Therefore,

$$\widehat{h}(t) = \frac{\tau(t) - m(t)}{a(t) \sqrt{\sum \frac{[\tau(t) - m(t)]^2}{a(t)}}}.$$

If we substitute the above value of h in the original problem we obtain

$$\begin{aligned} \sup_h \{ \mathbb{E}_\tau(h(X)) - \mathbb{E}_m(h(X)) \}^2 &= \sup_h \left\{ \sum h(t) [\tau(t) - m(t)]^2 \right\} \\ &= \left\{ \sum \frac{[\tau(t) - m(t)]^2}{a(t) \sqrt{\sum \frac{[\tau(t) - m(t)]^2}{a(t)}}} \right\}^2 \\ &= \left\{ \frac{1}{\sqrt{\sum \frac{[\tau(t) - m(t)]^2}{a(t)}}} \left(\sum \frac{[\tau(t) - m(t)]^2}{a(t)} \right) \right\}^2 = \sum \frac{[\tau(t) - m(t)]^2}{a(t)}, \end{aligned}$$

as was claimed. \square

Remark 1 Note that $\widehat{h}(t)$ is the least favorable function for detecting differences between means of two distributions.

Corollary 1 *The standardized function which creates the largest difference in means is*

$$\widehat{h}(t) = \frac{\tau(t) - m(t)}{a(t) \sqrt{\chi_a^2}},$$

where $\chi_a^2 = \sum \frac{[\tau(t) - m(t)]^2}{a(t)}$, and the corresponding difference in means is

$$\mathbb{E}_\tau[\widehat{h}(t)] - \mathbb{E}_m[\widehat{h}(t)] = \sqrt{\chi_a^2}.$$

Remark 2 Are there any additional distances that can be obtained as solutions to an optimization problem? And what is the statistical interpretation of these optimization problems? To answer the aforementioned questions we first present the optimization problems associated with the Kullback-Leibler and Hellinger distances. In fact, the entire class of the blended weighted Hellinger distances can be obtained as a solution to an appropriately defined optimization problem. Secondly, we discuss the statistical interpretability of these problems by connecting them, by analogy, to the construction of confidence intervals via Scheffé’s method.

Definition 5 The Kullback-Leibler divergence or distance between two discrete probability density functions is defined as

$$KL(\tau, m_\beta) = \sum_x m_\beta(x) [\log m_\beta(x) - \log \tau(x)].$$

Proposition 2 *The Kullback-Leibler distance is obtained as a solution of the optimization problem*

$$\sup_h \sum_x h(x) m_\beta(x), \quad \text{subject to } \sum_x e^{h(x)} \tau(x) \leq 1,$$

where $h(\cdot)$ is a function defined on the same space as τ .

Proof It is straightforward if one writes the Lagrangian and differentiates with respect to h . □

Definition 6 The class of squared blended weighted Hellinger distances ($BWHD_\alpha$) is defined as

$$(BWHD_\alpha)^2 = \sum_x \frac{[\tau(x) - m_\beta(x)]^2}{2 \left[\alpha \sqrt{\tau(x)} + \bar{\alpha} \sqrt{m_\beta(x)} \right]^2},$$

where $0 < \alpha < 1$, $\bar{\alpha} = 1 - \alpha$ and $\tau(x)$, $m_\beta(x)$ are two probability densities.

Proposition 3 *The class of $BWHD_\alpha$ arises as a solution to the optimization problem*

$$\sup_h \sum_x h(x) [\tau(x) - m_\beta(x)], \quad \text{subject to } \sum_x h^2(x) \left[\alpha \sqrt{\tau(x)} + \bar{\alpha} \sqrt{m_\beta(x)} \right]^2 \leq 1.$$

When $\alpha = \bar{\alpha} = 1/2$, the $(BWHD_{1/2})^2$ gives twice the squared Hellinger distance.

Proof Straightforward. □

Although both Kullback-Leibler and blended weighted Hellinger distances are solutions of appropriate optimization problems, they do not arise from optimization problems in which the constraints can be interpreted as variances. To exemplify and illustrate further this point we first need to discuss the connection with Scheffé’s confidence intervals.

One of the methods of constructing confidence intervals is Scheffé's method. The method adjusts the significance levels of the confidence intervals for general contrasts to account for multiple comparisons. The procedure, therefore, controls the overall significance for any possible contrast or set of contrasts and can be stated as follows,

$$\sup_{\mathbf{c}} |\mathbf{c}^T (\mathbf{y} - \boldsymbol{\mu})| < K \widehat{\sigma}, \quad \text{subject to } \|\mathbf{c}\| = 1, \mathbf{c}^T \mathbf{1} = 0,$$

where $\widehat{\sigma}$ is an estimated contrast variance, $\|\cdot\|$ denotes the Euclidean distance and K is an appropriately defined constant.

The chi-squared distances extend this framework as follows. Assume that \mathcal{H} is a class of functions which are taken, without loss of generality, to have zero expectation. Then, we construct the optimization problem $\sup_h \int h(x)[\tau(x) - m_\beta(x)]dx$, subject to a constraint that can possibly be interpreted as a constraint on the variance of $h(x)$ either under the hypothesized model distribution or under the distribution of the data.

The chi-squared distances arise as solutions of optimization problems subject to variance constraints. As such, they are interpretable as tools that allow the construction of "Scheffé-type" confidence intervals for models. On the other hand, distances such as the Kullback-Leibler or blended weighted Hellinger distance do not arise as solutions of optimization problems subject to interpretable variance constraints. As such they cannot be used to construct confidence intervals for models.

3.2 Loss Analysis of Pearson and Neyman Chi-Squared Distances

We next offer interpretations of the Pearson chi-squared and Neyman chi-squared statistics. These interpretations are not well known; furthermore, they are useful in illustrating the robustness character of the Neyman statistic and the non-robustness character of the Pearson statistic.

Recall that the Pearson statistic is

$$\begin{aligned} \sum \frac{[d(t) - m(t)]^2}{m(t)} &= \sup_h \frac{[\mathbb{E}_d(h(X)) - \mathbb{E}_m(h(X))]^2}{\text{Var}_m(h(X))} \\ &= \frac{1}{n} \sup_h \frac{[\frac{1}{n} \sum h(X_i) - \mathbb{E}_m(h(X))]^2}{\frac{1}{n} \text{Var}_m(h(X))} = \frac{1}{n} \sup_h Z_h^2, \end{aligned}$$

that is, the Pearson statistic is the supremum of squared Z-statistics.

A similar argument shows that Neyman's chi-squared equals $\sup_h t_h^2$, the supremum of squared t -statistics.

This property shows that the chi-squared measures have a statistical interpretation in that a small chi-squared distance indicates that the means are close on the scale of standard deviation. Furthermore, an additional advantage of the above interpretations is that the robustness character of these statistics is exemplified. Neyman’s chi-squared, being the supremum of squared t -statistics, is robust, whereas Pearson’s chi-squared is non-robust, since it is the supremum of squared Z -statistics.

Signal-to-Noise There is an additional interpretation of the chi-squared statistic that rests on the definition of signal-to-noise ratio that comes from the engineering literature.

Consider the pair of hypotheses $H_0 : X_i \sim \tau$ versus the alternative $H_1 : X_i \sim m$, where X_i are independent and identically distributed random variables. If we consider the set of randomized test functions that depend on the “output” function h , the distance between H_0 and H_1 is

$$S^2(\tau, m) = \frac{[\mathbb{E}_m(h(X)) - \mathbb{E}_\tau(h(X))]^2}{\text{Var}_\tau(h(X))}.$$

This quantity is a generalization of one of the more common definitions of signal-to-noise ratio. If, instead of working with a given output function h , we take supremum over the output functions h , we obtain Neyman’s chi-squared distance, which has been used in the engineering literature for robust detection. Further, the quantity $S^2(\tau, m)$ has been used in the design of decision systems (Poor 1980).

3.3 Metric Properties of the Symmetric Chi-Squared Distance

The symmetric chi-squared distance, defined as

$$S^2(\tau, m) = \sum \frac{2[\tau(t) - m(t)]^2}{m(t) + \tau(t)},$$

can be viewed as a good compromise between the non-robust Pearson distance and the robust Neyman distance. In what follows, we prove that $S^2(\tau, m)$ is indeed a metric. The following series of lemmas will help us establish the triangle inequality for $S^2(\tau, m)$.

Lemma 1 *If a, b, c are numbers such that $0 \leq a \leq b \leq c$ then*

$$\frac{c - a}{\sqrt{c + a}} \leq \frac{b - a}{\sqrt{b + a}} + \frac{c - b}{\sqrt{c + b}}.$$

Proof First we work with the right-hand side of the above inequality. Write

$$\begin{aligned} \frac{b-a}{\sqrt{b+a}} + \frac{c-b}{\sqrt{c+b}} &= (c-a) \left\{ \frac{b-a}{c-a} \frac{1}{\sqrt{b+a}} + \frac{c-b}{c-a} \frac{1}{\sqrt{c+b}} \right\} \\ &= (c-a) \left\{ \frac{\alpha}{\sqrt{a+b}} + (1-\alpha) \frac{1}{\sqrt{c+b}} \right\}, \end{aligned}$$

where $\alpha = (b-a)/(c-a)$. Set $g(t) = 1/\sqrt{t}$, $t > 0$. Then $g''(t) = \frac{d^2}{dt^2}g(t) > 0$, hence the function $g(t)$ is convex. Therefore, the aforementioned relationship becomes

$$(c-a) \{ \alpha g(a+b) + (1-\alpha)g(c+b) \}.$$

But

$$\alpha g(a+b) + (1-\alpha)g(c+b) \geq g(\alpha(a+b) + (1-\alpha)(c+b)),$$

where

$$\alpha(a+b) + (1-\alpha)(c+b) = \frac{b-a}{c-a}(a+b) + \frac{c-b}{c-a}(b+c) = c+a.$$

Thus

$$\alpha g(a+b) + (1-\alpha)g(c+b) \geq g(c+a),$$

and hence

$$\frac{b-a}{\sqrt{b+a}} + \frac{c-b}{\sqrt{c+b}} \geq \frac{c-a}{\sqrt{c+a}},$$

as was stated. □

Note that because the function is strictly convex we do not obtain equality except when $a = b = c$.

Lemma 2 *If a, b, c are numbers such that $a \geq 0, b \geq 0, c \geq 0$ then*

$$\left| \frac{c-a}{\sqrt{c+a}} \right| \leq \left| \frac{b-a}{\sqrt{b+a}} \right| + \left| \frac{c-b}{\sqrt{c+b}} \right|.$$

Proof We will distinguish three different cases.

Case 1: $0 \leq a \leq b \leq c$ is already discussed in Lemma 1.

Case 2: $0 \leq c \leq b \leq a$ can be proved as in Lemma 1 by interchanging the role of a and c .

Case 3: In this case b is not between a and c , thus either $a \leq c \leq b$ or $b \leq a \leq c$.

Assume first that $a \leq c \leq b$. Then we need to show that

$$\frac{c-a}{\sqrt{c+a}} \leq \frac{b-a}{\sqrt{b+a}}.$$

We will prove this by showing that the above expressions are the values of an increasing function at two different points. Thus, consider

$$f_1(t) = \frac{t-a}{\sqrt{t+a}}.$$

It follows that

$$f_1(b) = \frac{b-a}{\sqrt{b+a}} \quad \text{and} \quad f_1(c) = \frac{c-a}{\sqrt{c+a}}.$$

The function $f_1(t)$ is increasing because $f_1' > 0$ (recall $a \geq 0$) and since $c \leq b$ this implies $f_1(c) \leq f_1(b)$. Similarly we prove the inequality for $b \leq a \leq c$. \square

Lemma 3 *The triangle inequality holds for the symmetric chi-squared distance $S^2(\tau, m)$, that is,*

$$\{S^2(\tau, m)\}^{1/2} \leq \{S^2(\tau, g)\}^{1/2} + \{S^2(g, m)\}^{1/2}.$$

Proof Set

$$\alpha_t = \frac{|\tau(t) - g(t)|}{\sqrt{\tau(t) + g(t)}}, \quad \beta_t = \frac{|g(t) - m(t)|}{\sqrt{g(t) + m(t)}}.$$

By Lemma 2

$$\left\{ \sum \alpha_t^2 \right\}^{1/2} \leq \left\{ \sum (\alpha_t + \beta_t)^2 \right\}^{1/2}.$$

But

$$\begin{aligned} \sum (\alpha_t + \beta_t)^2 &= \sum \alpha_t^2 + \sum \beta_t^2 + 2 \sum \alpha_t \beta_t \\ &\leq \sum \alpha_t^2 + \sum \beta_t^2 + 2 \left\{ \sum \alpha_t^2 \right\}^{1/2} \left\{ \sum \beta_t^2 \right\}^{1/2}. \end{aligned}$$

Therefore

$$\sum (\alpha_t + \beta_t)^2 \leq \left\{ \sqrt{\sum \alpha_t^2} + \sqrt{\sum \beta_t^2} \right\}^2,$$

and hence

$$\left\{ \sum (\alpha_i + \beta_i)^2 \right\}^{1/2} \leq \left\{ \sum \alpha_i^2 \right\}^{1/2} + \left\{ \sum \beta_i^2 \right\}^{1/2},$$

as was claimed. \square

Remark 3 The inequalities proved in Lemmas 1 and 2 imply that if $\tau \neq m$ there is no “straight line” connecting τ and m , in that there does not exist g between τ and m for which the triangle inequality is an equality.

Therefore, the following proposition holds.

Proposition 4 *The symmetric chi-squared distance $S^2(\tau, m)$ is indeed a metric.*

A testing interpretation of the symmetric chi-squared distance: let ϕ be a test function and consider the problem of testing the null hypothesis that the data come from a density f versus the alternative that the data come from g . Let θ be a random variable with value 1 if the alternative is true and 0 if the null hypothesis is true. Then

Proposition 5 *The solution ϕ_{opt} to the optimization problem*

$$\min_{\phi} \mathbb{E}_{\pi}[(\theta - \phi(x))^2],$$

where $\pi(\theta)$ is the prior probability on θ , given as

$$\pi(\theta) = \begin{cases} 1/2, & \text{if } \theta = 0 \\ 1/2, & \text{if } \theta = 1 \end{cases},$$

is not a 0 – 1 decision, but equals the posterior expectation of θ given X . That is

$$\phi(t) = \mathbb{E}(\theta | X = t) = \mathbb{P}(\theta = 1 | X = t) = \frac{\frac{1}{2}g(t)}{\frac{1}{2}f(t) + \frac{1}{2}g(t)} = \frac{g(t)}{f(t) + g(t)},$$

the posterior probability that the alternative is correct.

Proof We have

$$\mathbb{E}(\theta | X) = \frac{1}{2} \mathbb{E}_{H_1}[(1 - \phi)^2] + \frac{1}{2} \mathbb{E}_{H_0}(\phi^2).$$

But

$$\mathbb{E}_{H_1}[(1 - \phi(X))^2] = \sum_t (1 - \phi(t))^2 g(t),$$

and

$$\mathbb{E}_{H_0}(\phi^2(X)) = \sum_t \phi^2(t) f(t),$$

hence

$$\phi_{\text{opt}}(t) = \frac{g(t)}{f(t) + g(t)},$$

as was claimed. □

Corollary 2 *The minimum risk is given as*

$$\frac{1}{4} \left(1 - \frac{S^2}{4} \right),$$

where

$$S^2 = S^2(f, g) = \sum \frac{[f(t) - g(t)]^2}{\frac{1}{2}f(t) + \frac{1}{2}g(t)}.$$

Proof Substitute ϕ_{opt} in $\mathbb{E}_{\pi}[(\theta - \phi)^2]$ to obtain

$$\mathbb{E}_{\pi}[(\theta - \phi_{\text{opt}})^2] = \frac{1}{2} \sum_t \frac{f(t)g(t)}{f(t) + g(t)}.$$

Now set

$$A = \sum_t \frac{[f(t) + g(t)]^2}{f(t) + g(t)} = 2, \quad B = \sum_t \frac{[f(t) - g(t)]^2}{f(t) + g(t)} = \frac{1}{2}S^2.$$

Then

$$A - B = 4 \sum_t \frac{f(t)g(t)}{f(t) + g(t)} = 2 - \frac{1}{2}S^2,$$

or, equivalently,

$$\sum_t \frac{f(t)g(t)}{f(t) + g(t)} = \frac{1}{4} \left(2 - \frac{1}{2}S^2 \right).$$

Therefore

$$\mathbb{E}_{\pi}[(\theta - \phi_{\text{opt}})^2] = \frac{1}{2} \sum_t \frac{f(t)g(t)}{f(t) + g(t)} = \frac{1}{4} \left(1 - \frac{S^2}{4} \right),$$

as was claimed. □

Remark 4 Note that $S^2(f, g)$ is bounded above by 4 and equals 4 when f, g are mutually singular.

The Kullback-Leibler and Hellinger distances are extensively used in the literature. Yet, we argue that, because they are obtained as solutions to optimization problems with non-interpretable (statistically) constraints, are not appropriate for our purposes. However, we note here that the Hellinger distance is closely related to the symmetric chi-squared distance, although this is not immediately obvious. We elaborate on this statement below.

Definition 7 Let τ, m be two probability densities. The squared Hellinger distance is defined as

$$H^2(\tau, m) = \frac{1}{2} \sum_x \left[\sqrt{\tau(x)} - \sqrt{m(x)} \right]^2.$$

We can more readily see the relationship between the Hellinger and chi-squared distances if we rewrite $H^2(\tau, m)$ as

$$H^2(\tau, m) = \frac{1}{2} \sum_x \frac{[\tau(x) - m(x)]^2}{[\sqrt{\tau(x)} + \sqrt{m(x)}]^2}.$$

Lemma 4 *The Hellinger distance is bounded by the symmetric chi-squared distance, that is,*

$$\frac{1}{8}S^2 \leq H^2 \leq \frac{1}{4}S^2,$$

where S^2 denotes the symmetric chi-squared distance.

Proof Note that

$$\left(\sqrt{\tau(x)} + \sqrt{m(x)} \right)^2 = \tau(x) + m(x) + 2\sqrt{\tau(x)m(x)} \geq \tau(x) + m(x).$$

Also

$$\left(\sqrt{\tau(x)} + \sqrt{m(x)} \right)^2 \leq 2[\tau(x) + m(x)],$$

and putting these relationships together we obtain

$$\tau(x) + m(x) \leq \left(\sqrt{\tau(x)} + \sqrt{m(x)} \right)^2 \leq 2[\tau(x) + m(x)].$$

Therefore

$$H^2(\tau, m) \leq \frac{1}{2} \sum_x \frac{[\tau(x) - m(x)]^2}{\tau(x) + m(x)} = \frac{1}{4}S^2,$$

and

$$H^2(\tau, m) \geq \frac{1}{2} \sum \frac{[\tau(x) - m(x)]^2}{2[\tau(x) + m(x)]} = \frac{1}{8} S^2,$$

and so

$$\frac{1}{8} S^2(\tau, m) \leq H^2(\tau, m) \leq \frac{1}{4} S^2(\tau, m),$$

as was claimed. \square

3.4 Locally Quadratic Distances

A generalization of the chi-squared distances is offered by the locally quadratic distances. We have the following definition.

Definition 8 A locally quadratic distance between two densities τ, m has the form

$$\rho(\tau, m) = \sum K_m(x, y)[\tau(x) - m(x)][\tau(y) - m(y)],$$

where $K_m(x, y)$ is a nonnegative definite kernel, possibly dependent on m , and such that

$$\sum_{x,y} a(x)K_m(x, y)a(y) \geq 0,$$

for all functions $a(x)$.

Example 1 The Pearson distance can be written as

$$\begin{aligned} \sum \frac{(d(t) - m(t))^2}{m(t)} &= \sum \frac{\mathbb{1}[s = t]}{\sqrt{m(s)m(t)}} [d(s) - m(s)][d(t) - m(t)] \\ &= \sum K_m(s, t)[d(s) - m(s)][d(t) - m(t)], \end{aligned}$$

where $\mathbb{1}(\cdot)$ is the indicator function. It is a quadratic distance with kernel

$$K_m(s, t) = \frac{\mathbb{1}[s = t]}{\sqrt{m(s)m(t)}}.$$

Sensitivity and Robustness In the classical robustness literature one of the attributes that a method should exhibit so as to be characterized as robust is the attribute of being resistant, that is insensitive, to the presence of a moderate number of outliers and to inadequacies in the assumed model.

Similarly here, to characterize a statistical distance as robust it should be insensitive to small changes in the true density, that is, the value of the distance should not be greatly affected by small changes that occur in τ . Lindsay (1994), Markatou (2000, 2001), and Markatou et al. (1997, 1998) based the discussion of robustness of the distances under study on a mechanism that allows the identification of distributional errors, that is, on the Pearson residual. A different system of residuals is the set of symmetrized residuals defined as follows.

Definition 9 If τ, m are two densities the symmetrized residual is defined as

$$r_{\text{sym}}(t) = \frac{\tau(t) - m(t)}{\tau(t) + m(t)}.$$

The symmetrized residuals have range $[-1, 1]$, with value -1 when $\tau(t) = 0$ and value 1 when $m(t) = 0$. Symmetrized residuals are important because they allow us to understand the way different distances treat different distributions.

The symmetric chi-squared distance can be written as a function of the symmetrized residuals as follows

$$S^2(\tau, m) = 4 \sum \left(\frac{1}{2}\tau(t) + \frac{1}{2}m(t) \right) \left\{ \frac{\tau(t) - m(t)}{\tau(t) + m(t)} \right\}^2 = 4 \sum b(t)r_{\text{sym}}^2(t),$$

where $b(t) = [\tau(t) + m(t)]/2$.

The aforementioned expression of the symmetric chi-squared distance allows us to obtain inequalities between $S^2(\tau, m)$ and other distances.

A third residual system is the set of logarithmic residuals, defined as follows.

Definition 10 Let τ, m be two probability density. Define the logarithmic residuals as

$$\delta(t) = \log \left(\frac{\tau(t)}{m(t)} \right),$$

with $\delta \in (-\infty, \infty)$.

A value of this residual close to 0 indicates agreement between τ and m . Large positive or negative values indicate disagreement between the two models τ and m .

In an analysis of a given data set, there are two types of observations that cause concern: outliers and influential observations. In the literature, the concept of an outlier is defined as follows.

Definition 11 We define an outlier to be an observation (or a set of observations) which appears to be inconsistent with the remaining observations of the data set.

Therefore, the concept of an outlier may be viewed in relative terms. Suppose we think a sample arises from a standard normal distribution. An observation from this sample is an outlier if it is somehow different in relation to the remaining observations that were generated from the postulated standard normal model. This means that, an observation with value 4 may be surprising in a sample of size 10, but is less so if the sample size is 10,000. In our framework therefore, the extent to which

an observation is an outlier depends on both the sample size and the probability of occurrence of the observation under the specified model.

Remark 5 Davies and Gather (1993) state that although detection of outliers is a topic that has been extensively addressed in the literature, the word “outlier” was not given a precise definition. Davies and Gather (1993) formalized this concept by defining outliers in terms of their position relative to a central model, and in relationship to the sample size. Further details can be found in their paper.

On the other hand, the literature provides the following definition of an influential observation.

Definition 12 (Belsley 1980) An influential observation is one which, either individually or together with several other observations, has a demonstrably larger impact on the calculated values of various estimates than is the case for most of the other observations.

Chatterjee and Hadi (1986) use this definition to address questions about measuring influence and discuss the different measures of influence and their inter-relationships.

The aforementioned definition is subjective, but it implies that one can order observations in a sensible way according to some measure of influence. Outliers need not be influential observations and influential observations need not be outliers. Large Pearson residuals correspond to observations that are *surprising*, in the sense that they occur in locations with small model probability. This is different from influential observations, that is from observations for which their presence or absence greatly affects the value of the maximum likelihood estimator.

Outliers can be surprising observations as well as influential observations. In a normal location-scale model, an outlying observation is both surprising and influential on the maximum likelihood estimator of location. But in the double exponential location model, an outlying observation is possible to be surprising but never influential on the maximum likelihood estimator of location as it equals the median.

Lindsay (1994) shows that the robustness of these distances is expressed via a key function called *residual adjustment function* (RAF). Further, he studied the characteristics of this function and showed that an important class of RAFs is given by $A_\lambda(\delta) = \frac{(1+\delta)^\lambda - 1}{\lambda + 1}$, where δ is the Pearson residual (defined by Eq. (1)). From this class we obtain many RAFs; in particular, when $\lambda = -2$ we obtain the RAF corresponding to Neyman’s chi-squared distance. For details, see Lindsay (1994).

4 The Continuous Setting

Our goal is to use statistical distances to construct model misspecification measures. One of the key issues in the construction of misspecification measures in the case of data being realizations of a random variable that follows a continuous distribution is that allowances should be made for the scale difference between

observed data and hypothesized model. That is, data distributions are discrete while the hypothesized model is continuous. Hence, we require the distance to exhibit discretization robustness, so it can account for the difference in scale.

To achieve discretization robustness, we need a sensitive distance, which implies a need to balance sensitivity and statistical noise. We will briefly review available strategies to deal with the problem of balancing sensitivity of the distance and statistical noise.

In what follows, we discuss desirable characteristics we require our distance measures to satisfy.

4.1 *Desired Features*

Discretization Robustness Every real data distribution is discrete, and therefore is different from every continuous distribution. Thus, a reasonable distance measure must allow for discretization, by saying that the discretized version of a continuous distribution must get closer to the continuous distribution as the discretization gets finer.

A second reason for requiring discretization robustness is that we will want to use the empirical distribution to estimate the true distribution, but without this robustness, there is no hope that the discrete empirical distribution will be closed to any model point.

The Problem of Too Many Questions Thus, to achieve discretization robustness, we need to construct a sensitive distance. This requirement dictates us to carry out a delicate balancing act between sensitivity and statistical noise.

Lindsay (2004) discusses in detail the problem of too many questions. Here we only note that to illustrate the issue Lindsay (2004) uses the chi-squared distance and notes that the statistical implications of a refinement in partition are the widening of the sensitivity to model departures in new “directions” but, at the same time, this act increases the statistical noise and therefore decreases the power of the chi-squared test in every existing direction.

There are a number of ways to address this problem, but they all seem to involve a loss of statistical information. This means we cannot ask all model fit questions with optimal accuracy. Two immediate solutions are as follows. First, limit the investigation only to a finite list of questions, essentially boiling down to prioritizing the questions asked of the sample. A number of classical goodness-of-fit tests create exactly such a balance. A second approach to the problem of answering infinitely many questions with only a finite number of data points is through the construction of kernel smoothed density measures. Those measures provide a flexible class of distances that allows for adjusting the sensitivity/noise trade-off. Before we briefly comment on this strategy, we discuss statistical distances between continuous probability distributions.

4.2 The L_2 -Distance

The L_2 distance is very popular in density estimation. We show below that this distance is not invariant to one-to-one transformations.

Definition 13 The L_2 distance between two probability density functions τ , m is defined as

$$L_2^2(\tau, m) = \int [\tau(x) - m(x)]^2 dx.$$

Proposition 6 The L_2 distance between two probability density functions is not invariant to one-to-one transformations.

Proof Let $Y = a(X)$ be a transformation of X , which is one-to-one. Then $x = b(y)$, $b(\cdot)$ is the inverse transformation of $a(\cdot)$, and

$$\begin{aligned} L_2^2(\tau_Y, m_Y) &= \int [\tau_Y(y) - m_Y(y)]^2 dy \\ &= \int [\tau_X(b(y)) - m_X(b(y))]^2 (b'(y))^2 dy \\ &= \int [\tau_X(x) - m_X(x)]^2 (b'(a(x)))^2 a'(x) dx \\ &= \int [\tau_X(x) - m_X(x)]^2 b'(a(x)) dx \\ &\neq \int [\tau_X(x) - m_X(x)]^2 dx = L_2^2(\tau_X, m_X). \end{aligned}$$

Thus, the L_2 distance is not invariant under monotone transformations. □

Remark 6 It is easy to see that the L_2 distance is location invariant. Moreover, scale changes appear as a constant factor multiplying the L_2 distance.

4.3 The Kolmogorov-Smirnov Distance

We now discuss the Kolmogorov-Smirnov distance used extensively in goodness-of-fit problems, and present its properties.

Definition 14 The Kolmogorov-Smirnov distance between two cumulative distribution functions F , G is defined as

$$\rho_{KS}(F, G) = \sup_x |F(x) - G(x)|.$$

Proposition 7 (Testing Interpretation) *Let $H_0 : \tau = f$ versus $H_1 : \tau = g$ and that only test functions φ of the form $\mathbb{1}(x \leq x_0)$ or $\mathbb{1}(x > x_0)$ for arbitrary x_0 are allowed. Then*

$$\rho_{KS}(F, G) = \sup |\mathbb{E}_{H_1}[\varphi(X)] - \mathbb{E}_{H_0}[\varphi(X)]|.$$

Proof The difference between power and size of the test is $G(x_0) - F(x_0)$. Therefore,

$$\sup_{x_0} |G(x_0) - F(x_0)| = \sup_{x_0} |F(x_0) - G(x_0)| = \rho_{KS}(F, G),$$

as was claimed. □

Proposition 8 *The Kolmogorov-Smirnov distance is invariant under monotone transformations.*

Proof Write

$$F(x_0) - G(x_0) = \int \mathbb{1}(x \leq x_0)[f(x) - g(x)]dx.$$

Let $Y = a(X)$ be a one-to-one transformation and $b(\cdot)$ be the corresponding inverse transformation. Then $x = b(y)$ and $dy = a'(x)dx$, so

$$\begin{aligned} F_Y(y_0) - G_Y(y_0) &= \int \mathbb{1}(y \leq y_0)[f_Y(y) - g_Y(y)]dy \\ &= \int \mathbb{1}(y \leq y_0)[f_X(b(y))b'(y) - g_X(b(y))b'(y)]dy \\ &= \int \mathbb{1}(x \leq b(y_0))[f_X(b(y))b'(y) - g_X(b(y))b'(y)]dy \\ &= \int \mathbb{1}(x \leq x_0)[f_X(x) - g_X(x)]dx. \end{aligned}$$

Therefore,

$$\sup_{y_0} |F_Y(y_0) - G_Y(y_0)| = \sup_{x_0} |F_X(x_0) - G_X(x_0)|,$$

and the Kolmogorov-Smirnov distance is invariant under one-to-one transformations. □

Proposition 9 *The Kolmogorov-Smirnov distance is discretization robust.*

Proof Notice that we can write

$$|F(x_0) - G(x_0)| = \left| \int \mathbb{1}(x \leq x_0)d[F(x) - G(x)] \right|,$$

with $\mathbb{1}(x \leq x_0)$ being thought of as a “smoothing kernel”. Hence, comparisons between discrete and continuous distributions are allowed and the distance is discretization robust. \square

The Kolmogorov-Smirnov distance is a distance based on the probability integral transform. As such, it is invariant under monotone transformations (see Proposition 8). A drawback of distances based on probability integral transforms is the fact that there is no obvious extension in the multivariate case. Furthermore, there is not a direct loss function interpretation of these distances when the model used is incorrect. In what follows, we discuss chi-squared and quadratic distances that avoid the issues listed above.

4.4 Exactly Quadratic Distances

In this section we briefly discuss exactly quadratic distances. Rao (1982) introduced the concept of an exact quadratic distance for discrete population distributions and he called it *quadratic entropy*. Lindsay et al. (2008) gave the following definition of an exactly quadratic distance.

Definition 15 (Lindsay et al. 2008) Let F, G be two probability distributions, and K is a nonnegative definite kernel. A quadratic distance between F, G has the form

$$\rho_K(F, G) = \iint K_G(x, y)d(F - G)(x)d(F - G)(y).$$

Quadratic distances are of interest for a variety of reasons. These include the fact that the empirical distance $\rho_K(\widehat{F}, G)$ has a fairly simple asymptotic distribution theory when G identifies with the true model τ , and that several important distances are exactly quadratic (see, for example, Cramér-von Mises and Pearson’s chi-squared distances). Furthermore, other distances are asymptotically locally quadratic around $G = \tau$. Quadratic distances can be thought of as extensions of the chi-squared distance class.

We can construct an exactly quadratic distance as follows. Let F, G be two probability measures that a random variable X may follow. Let ε be an independent error variable with known density $k_h(\varepsilon)$, where h is a parameter. Then, the random variable $Y = X + \varepsilon$ has an absolutely continuous distribution such that

$$f_h^*(y) = \int k_h(y - x)dF(x), \quad \text{if } X \sim F,$$

or

$$g_h^*(y) = \int k_h(y - x)dG(x), \quad \text{if } X \sim G.$$

Let

$$P^{*2}(F, G) = \int \frac{[f^*(y) - g^*(y)]^2}{g^*(y)} dy,$$

be the kernel-smoothed Pearson's chi-squared statistic. In what follows, we prove that $P^{*2}(F, G)$ is an exactly quadratic distance.

Proposition 10 *The distance $P^{*2}(F, G)$ is an exactly quadratic distance provided that $\iint K(s, t)d(F - G)(s)d(F - G)(t) < \infty$, where $K(s, t) = \int \frac{k_h(y-s)k_h(y-t)}{g^*(y)} dy$.*

Proof Write

$$\begin{aligned} P^{*2}(F, G) &= \int \frac{[f^*(y) - g^*(y)]^2}{g^*(y)} dy \\ &= \int \frac{[\int k_h(y-x)dF(x) - \int k_h(y-x)dG(x)]^2}{g^*(y)} dy \\ &= \int \frac{[\int k_h(y-x)d(F-G)(x)]^2}{g^*(y)} dy \\ &= \int \frac{[\int k_h(y-s)d(F-G)(s)][\int k_h(y-t)d(F-G)(t)]}{g^*(y)} dy. \end{aligned}$$

Now using Fubini's theorem, the above relationship can be written as

$$\begin{aligned} &\iint \left\{ \int \frac{k_h(y-s)k_h(y-t)}{g^*(y)} dy \right\} d(F-G)(s)d(F-G)(t) \\ &= \iint K(s, t)d(F-G)(s)d(F-G)(t), \end{aligned}$$

with $K(s, t)$ given above. □

Remark 7

- (a) The issue with many classical measures of goodness-of-fit is that the balance between sensitivity and statistical noise is fixed. On the other hand, one might wish to have a flexible class of distances that allows for adjusting the sensitivity/noise trade-off. Lindsay (1994) and Basu and Lindsay (1994) introduced the idea of smoothing and investigated numerically the blended weighted Hellinger distance, defined as

$$BWHD_\alpha(\tau^*, m_\theta^*) = \int \frac{(\tau^*(x) - m_\theta^*(x))^2}{(\alpha \sqrt{\tau^*(x)} + \bar{\alpha} \sqrt{m_\theta^*(x)})^2} dx,$$

where $\bar{\alpha} = 1 - \alpha$, $\alpha \in [1/3, 1]$. When $\alpha = 1/2$, the $BWHD_{1/2}$ equals the Hellinger distance.

- (b) Distances based on kernel smoothing are natural extensions of the discrete distances. These distances are not invariant under one-to-one transformations, but they can be easily generalized to higher dimensions. Furthermore, numerical integration is required for the practical implementation and use of these distances.

5 Discussion

In this paper we study statistical distances with a special emphasis on the chi-squared distance measures. We also introduce an extension of the chi-squared distance, the quadratic distance, introduced by Lindsay et al. (2008). We offered statistical interpretations of these distances and showed how they can be obtained as solutions of certain optimization problems. Of particular interest are distances with statistically interpretable constraints such as the class of chi-squared distances. These allow the construction of confidence intervals for models. We further discussed robustness properties of these distances, including discretization robustness, a property that allows discrete and continuous distributions to be arbitrarily close. Lindsay et al. (2014) study the use of quadratic distances in problems of goodness-of-fit with particular focus on creating tools for studying the power of distance-based tests. Lindsay et al. (2014) discuss one-sample testing and connect their methodology with the problem of kernel selection and the requirements that are appropriate in order to select optimal kernels. Here, we outlined the foundations that led to the aforementioned work and showed how these elucidate the performance of statistical distances as inferential functions.

Acknowledgements The first author dedicates this paper to the memory of Professor Bruce G. Lindsay, a long time collaborator and friend, with much respect and appreciation for his mentoring and friendship. She also acknowledges the Department of Biostatistics, School of Public Health and Health Professions and the Jacobs School of Medicine and Biomedical Sciences, University at Buffalo, for supporting this work. The second author acknowledges the Troup Fund, Kaleida Foundation for supporting this work.

References

- Basu, A., & Lindsay, B. G. (1994). Minimum disparity estimation for continuous models: Efficiency, distributions and robustness. *Annals of the Institute of Statistical Mathematics*, 46, 683–705.
- Belsley, D. A., Kuh, E., & Welsch, R. E. (1980). *Regression diagnostics: Identifying influential data and sources of collinearity*. New York: Wiley.
- Beran, R. (1977). Minimum Hellinger distance estimates for parametric models. *Annals of Statistics*, 5, 445–463.
- Berkson, J. (1980). Minimum chi-square, not maximum likelihood! *Annals of Statistics*, 8, 457–487.

- Chatterjee, S., & Hadi, A. S. (1986). Influential observations, high leverage points, and outliers in linear regression. *Statistical Science*, *1*, 379–393.
- Davies, L., & Gather, U. (1993). The identification of multiple outliers. *Journal of the American Statistical Association*, *88*, 782–792.
- Lindsay, B. G. (1994). Efficiency versus robustness: The case for minimum Hellinger distance and related methods. *Annals of Statistics*, *22*, 1081–1114.
- Lindsay, B. G. (2004). Statistical distances as loss functions in assessing model adequacy. In *The nature of scientific evidence: Statistical, philosophical and empirical considerations* (pp. 439–488). Chicago: The University of Chicago Press.
- Lindsay, B. G., Markatou, M., & Ray, S. (2014). Kernels, degrees of freedom, and power properties of quadratic distance goodness-of-fit tests. *Journal of the American Statistical Association*, *109*, 395–410.
- Lindsay, B. G., Markatou, M., Ray, S., Yang, K., & Chen, S. C. (2008). Quadratic distances on probabilities: A unified foundation. *Annals of Statistics*, *36*, 983–1006.
- Markatou, M. (2000). Mixture models, robustness, and the weighted likelihood methodology. *Biometrics*, *56*, 483–486.
- Markatou, M. (2001). A closer look at weighted likelihood in the context of mixture. In C. A. Charalambides, M. V. Koutras, & N. Balakrishnan (Eds.), *Probability and statistical models with applications* (pp. 447–468). Boca Raton: Chapman & Hall/CRC.
- Markatou, M., Basu, A., & Lindsay, B. G. (1997). Weighted likelihood estimating equations: The discrete case with applications to logistic regression. *Journal of Statistical Planning and Inference*, *57*, 215–232.
- Markatou, M., Basu, A., & Lindsay, B. G. (1998). Weighted likelihood equations with bootstrap root search. *Journal of the American Statistical Association*, *93*, 740–750.
- Matusita, K. (1955). Decision rules, based on the distance, for problems of fit, two samples, and estimation. *Annals of Mathematical Statistics*, *26*, 613–640.
- Neyman, J. (1949). Contribution to the theory of the χ^2 test. In *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability* (pp. 239–273). Berkeley: The University of California Press.
- Poor, H. (1980). Robust decision design using a distance criterion. *IEEE Transactions on Information Theory*, *26*, 575–587.
- Rao, C. R. (1963). Criteria of estimation in large samples. *Sankhya Series A*, *25*, 189–206.
- Rao, C. R. (1982). Diversity: Its measurement, decomposition, apportionment and analysis. *Sankhya Series A*, *44*, 1–21.
- Wald, A. (1950). *Statistical decision functions*. New York: Wiley.