ON THE EQUIVALENCE OF *f*-DIVERGENCE BALLS AND DENSITY BANDS IN ROBUST DETECTION

Michael Fau β^* , *Abdelhak M. Zoubir*^{*}, *and H. Vincent Poor*[†]

*Signal Processing Group, Technische Universität Darmstadt, 64283 Darmstadt, Germany [†]Dept. of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA

ABSTRACT

The paper deals with minimax optimal statistical tests for two composite hypotheses, where each hypothesis is defined by a nonparametric uncertainty set of feasible distributions. It is shown that for every pair of uncertainty sets of the f-divergence-ball type, a pair of uncertainty sets of the density-band type can be constructed, which is equivalent in the sense that it admits the same pair of least favorable distributions. This result implies that robust tests under f-divergence-ball uncertainty, which are typically only minimax optimal for the single sample case, are also fixed sample size minimax optimal with respect to the equivalent density-band uncertainty sets.

Index Terms— Hypothesis testing, robust detection, distributional uncertainty, *f*-divergence, density bands

1. INTRODUCTION

A statistical test for two composite hypotheses is called minimax optimal if it minimizes the maximum risk over the two corresponding sets of feasible distributions. In the context of robust statistics, these sets are referred to as *uncertainty sets*. In contrast to adaptive procedures [1], the minimax approach provides strict guarantees on the error probabilities for all feasible distributions. Moreover, minimax tests are often easy to implement since they typically reduce to an optimal test for two simple hypotheses, where each hypothesis is represented by a *least favorable distribution*.

A common way of specifying uncertainty sets is via a neighborhood around a nominal distribution, which represents an ideal system state or model [2]. In many works on robust detection, the use of *f*-divergence-balls has been proposed as a useful and versatile model to construct such neighborhoods [3–10]. In contrast to outlier models, such as ε -contamination [11], *f*-divergence balls do not allow for arbitrarily large deviations from the nominals and, therefore, have been argued to better represent scenarios where the shape of a distribution is subject to uncertainty, but there are no gross outliers in the data [5].

In order to present the result in this paper, the concept of single sample and fixed sample size tests needs to be introduced. A single sample test is based on the observation of a single, but possibly vector-valued, random variable X_1 . Consequently, the uncertainty sets are defined in terms of all possible joint distributions of the elements of X. Such an uncertainty model is suitable in some cases, but more often the observations are obtained by repeatedly performing independent random experiments so that the test is based on a sequence of independent random variables $X_1, \ldots, X_N, N > 1$. By definition, this independence constraint cannot be incorporated into a single sample test. Hence, tests whose observations are realizations of multiple independent random variables need to be considered separately. In order to highlight the difference to tests whose sample size is random [12], they are referred to as fixed sample size tests in what follows.

For most commonly used uncertainty models-including the density-band model, which will be discussed in detail later on-it can be shown that a single sample minimax optimal test is also fixed sample size minimax optimal. More precisely, the least favorable distributions for X_1 in the single sample case are also least favorable for all X_1, \ldots, X_N in the fixed sample size case. However, in general, this does not hold true for uncertainty sets of the f-divergence-ball type, where the fixed sample size minimax optimal solution is typically intractable. However, a commonly applied heuristic is to use the single sample least favorable distributions for the fixed sample size case anyhow, regardless of the fact that this extension does not hold in theory; compare [10, Sec. V.A]. Tests constructed this way are referred to as single sample minimax optimal tests with repeated observations. Evidently, such tests are no longer minimax optimal. However, in practice, it can be observed that they are still robust, meaning that they meet the specified error probabilities for most if not all feasible distributions.

In this paper, the favorable robustness properties of single sample minimax optimal tests with repeated observations are explained in a rigorous manner. It is shown that they are indeed fixed sample size minimax optimal, but for a density-band uncertainty model instead of the original f-divergence-ball model. That is, singlesample minimax optimal tests can be applied to repeated observations without sacrificing minimax optimality, if one is willing to accept a change in the uncertainty model. This result is proved by showing that for every f-divergence-ball model, there exists and equivalent density-band model that admits the same single sample minimax optimal solution. However, for the density-band model, this solution is known to be fixed sample size optimal as well so that it automatically extends to the case of repeated observations.

The paper is organized as follows: a brief review of minimaxoptimal detection is given in Section 2. The two uncertainty models of interest, i.e., f-divergence balls and density bands, are introduced in Section 3. The main result is stated and proved in Section 4, followed by a brief discussion in Section 5. An illustrative example is shown in Section 6, which also concludes the paper.

2. MINIMAX OPTIMAL DETECTION

The single sample case is considered first. Let $(\mathcal{X}, \mathcal{F})$ be a measurable space and let X_1 be a $(\mathcal{X}, \mathcal{F})$ -valued random variable that is distributed according to a probability measure (distribution) P. Throughout the paper it is assumed that all distributions on $(\mathcal{X}, \mathcal{F})$

This work was supported in part by the U.S. National Science Foundation under Grants CCF-1420575 and ECCS-1549881.

have a continuous density function with respect to some σ -finite reference measure μ . The set of all distributions on $(\mathcal{X}, \mathcal{F})$ that admit this property is denoted by \mathcal{M}_{μ} .

The goal of a simple, non-robust binary hypothesis test is to decide between the two hypotheses

$$\mathcal{H}_0: P = P_0, \qquad \mathcal{H}_1: P = P_1,$$

where $P_0, P_1 \in \mathcal{M}_{\mu}$ are two given distributions. The test is defined by a decision $d \in \{0, 1\}$ and a possibly randomized decision rule $\delta: \mathcal{X} \to [0, 1]$, where $\delta(x)$ denotes the conditional probability to decide for \mathcal{H}_1 , given the observation $X_1 = x$. The set of all decision rules is denoted by Δ . The type I and type II error probabilities are given by

$$P_0[d=1] = E_{P_0}[\delta(X)],$$

$$P_1[d=0] = E_{P_1}[1-\delta(X)].$$

The optimal decision rule δ^* for the simple binary hypothesis test is a threshold comparison of the likelihood ratio, i.e.,

$$\delta^*(x) = \begin{cases} 1, & l(x) > \lambda \\ \kappa, & l(x) = \lambda \\ 0, & l(x) < \lambda \end{cases}$$

where $\lambda>0$ is the threshold value, $\kappa\in[0,1]$ can be chosen arbitrarily, and l(x) denotes the likelihood ratio

$$l(x) = \frac{p_1(x)}{p_0(x)}.$$

The likelihood ratio test is optimal in a very general sense [13]. In particular, it minimizes the weighted sum error probability, i.e., it solves

$$\min_{\delta \in \Delta} E_{P_1}[\delta(X)] + \lambda E_{P_0}[1 - \delta(X)].$$
(1)

In robust detection, the distribution under each hypothesis is assumed not to be known exactly. The distributional uncertainty is modeled by two disjoint sets $\mathcal{P}_0, \mathcal{P}_1 \subset \mathcal{M}_{\mu}$ so that the hypotheses become

$$\mathcal{H}_0: P \in \mathcal{P}_0, \qquad \mathcal{H}_1: P \in \mathcal{P}_1.$$

The minimax problem corresponding to (1) is thus given by

$$\min_{\delta \in \Delta} \max_{\substack{H_0 \in \mathcal{P}_0 \\ H_1 \in \mathcal{P}_1}} E_{H_1}[\delta(X)] + \lambda E_{H_0}[1 - \delta(X)].$$
(2)

Problem (2) is central to robust detection. By definition, its solution is minimax optimal with respect to the weighted sum of error probabilities, but it can also be shown to be minimax optimal in the sense of Neyman–Pearson and the Baysian sense. This property is fixed in the following definition.

Definition 1. A triplet (δ^*, Q_0, Q_1) that solves (2) for a given $\lambda > 0$ is called single sample minimax optimal with respect to the threshold λ and the uncertainty sets \mathcal{P}_0 , \mathcal{P}_1 . This is written as

$$(\delta^*, Q_0, Q_1) \in \{\mathcal{P}_0, \mathcal{P}_1\}^*_{\lambda}.$$

In [11] and [14] it is shown that if the solution of (2) is independent of the threshold λ , it is also minimax optimal for fixed sample size tests with arbitrary thresholds and arbitrary sample sizes. This property is fixed in the next definition.

Definition 2. A triplet (δ^*, Q_0, Q_1) that jointly solves (2) for all $\lambda > 0$ is called fixed sample size minimax optimal with respect to $\mathcal{P}_0, \mathcal{P}_1$. This is written as

$$(\delta^*, Q_0, Q_1) \in \{\mathcal{P}_0, \mathcal{P}_1\}^*$$

3. UNCERTAINTY SETS

Two types of uncertainty sets are introduced in this section. The first one is the f-divergence-ball model, which specifies uncertainty sets in terms of a maximum feasible distance from a nominal distribution and allows the use of arbitrary f-divergences to define this distance. Formally, f-divergence-ball uncertainty sets are of the form

$$\mathcal{P}_f(P,\varepsilon) = \{ H \in \mathcal{M}_\mu : D_f(H \| P) \le \varepsilon \},\tag{3}$$

where P denotes the nominal distribution and D_f denotes the f-divergence induced by the function f, i.e.,

$$D_f(H||P) = \int_{\mathcal{X}} f\left(\frac{\mathrm{d}H}{\mathrm{d}P}(x)\right) \mathrm{d}P(x)$$
$$= \int_{\mathcal{X}} f\left(\frac{h(x)}{p(x)}\right) p(x) \mathrm{d}\mu(x)$$

where $f: \mathbb{R}_{\geq 0} \to \mathbb{R}$ is convex and satisfies f(1) = 0. The definition of the *f*-divergence ball in terms of $D_f(H||P)$ instead of $D_f(P||H)$ is arbitrary since for every feasible function *f* it holds that $D_f(P||H) = D_{\tilde{f}}(H||P)$, with $\tilde{f}(x) = f(\frac{1}{x})x$.

Owing to the mild constraints on f, uncertainty sets of the form (3) offer a great amount of flexibility and have attracted increased attention in recent years. In [5] and [10], minimax optimal tests based on the Kullback–Leibler divergence were derived under varying assumptions. Minimax optimal tests have also been derived for the squared Hellinger distance [6,7], the total variation distance [8], and α -divergence-ball uncertainty is that no fixed sample size minimax optimal solution is guaranteed to exist. In fact, most of the works cited above only consider the single sample case.

The second type of uncertainty model is the density-band model. In a robust detection context, it was first proposed by Kassam [15] and covers sets of the form

$$\mathcal{P}_{\mathsf{b}}(P',P'') = \{ H \in \mathcal{M}_{\mu} : p' \le h \le p'' \},\tag{4}$$

where P', P'' are nonnegative measures on $(\mathcal{X}, \mathcal{F})$ that admit densities p', p'' with respect to μ and satisfy

$$P'(\mathcal{X}) \le 1, \quad P''(\mathcal{X}) \ge 1, \quad \text{and} \quad 0 \le p' \le p''.$$

In words, the density-band model restricts the true density to lie within a band specified by p' and p''. Similar to the choice of f in the f-divergence-ball model, the choice of P' and P'' allows for a*priori* knowledge about the type of contamination to be incorporated into the model. Therefore, although it is still based on the concept of outliers, the band model can capture a much larger variety of contamination types and mismatches than the standard ε -contamination model. Another useful property of the density-band model is that for every pair of uncertainty sets of the form (4), a fixed sample size minimax optimal test is guaranteed to exist. Moreover, the corresponding least favorable densities can be calculated in a generic manner using a simple and efficient algorithm. See [14] for a more detailed discussion of the band model and the calculation of its least favorable densities.

In the next section, it is shown in how far the two uncertainty models can be considered equivalent in the single sample case and how the density-band model can be used to construct fixed sample size minimax optimal tests from tests that are only single sample size minimax optimal under f-divergence-ball uncertainty.

4. MAIN RESULT

In this section, the main result of the paper is stated and proved. A more detailed discussion is deferred to Section 5.

Theorem. Let $\mathcal{P}_{f_0}(P_0, \varepsilon_0)$ and $\mathcal{P}_{f_1}(P_1, \varepsilon_1)$ be two uncertainty sets of the form (3). If it holds that

$$(\delta^*, Q_0, Q_1) \in \{\mathcal{P}_{f_0}(P_0, \varepsilon_0), \mathcal{P}_{f_1}(P_1, \varepsilon_1)\}^*_{\lambda},\$$

then there exist nonnegative scalars $a_0 \leq b_0$ and $a_1 \leq b_1$ such that

$$(\delta^*, Q_0, Q_1) \in \{\mathcal{P}_b(a_0P_0, b_0P_0), \mathcal{P}_b(a_1P_1, b_1P_1)\}^*.$$

In words, the theorem states that if (δ^*, Q_0, Q_1) is single sample minimax optimal for an *f*-divergence-ball uncertainty model, a density-band model can be constructed from scaled versions of the nominal densities such that (δ^*, Q_0, Q_1) is fixed sample size minimax optimal with respect to this band model. A proof is detailed below.

Proof. Let P_0 , P_1 , f_0 , f_1 , ε_0 , ε_1 , and λ be given. Rewriting the optimization problem (2) in terms of the densities and with explicit constraints yields

$$\max_{\substack{h_0>0\\h_1>0}} \min_{\delta \in \Delta} \int_{\mathcal{X}} h_1 \,\delta + \lambda h_0 (1-\delta) \,\mathrm{d}\mu \quad s.t.$$
(5)

$$\int_{\mathcal{X}} f_0\left(\frac{h_0}{p_0}\right) p_0 \,\mathrm{d}\mu \le \varepsilon_0, \quad \int_{\mathcal{X}} f_1\left(\frac{h_1}{p_1}\right) p_1 \,\mathrm{d}\mu \le \varepsilon_1 \qquad (6)$$

$$\int_{\mathcal{X}} h_0 \,\mathrm{d}\mu = 1, \quad \int_{\mathcal{X}} h_1 \,\mathrm{d}\mu = 1. \tag{7}$$

By assumption, (δ^*, Q_0, Q_1) solves this minimax problem, which implies that (δ^*, Q_0, Q_1) is a saddle point of (5). This, in turn, implies that (δ^*, Q_0, Q_1) satisfies the corresponding Karush–Kuhn– Tucker conditions, which are first order necessary conditions for optimality [16]. In particular, stationarity of the saddle point solution implies that scalars η_0, η_1 and nonnegative scalars ν_0, ν_1 exists such that

$$\lambda(1-\delta^*) = \nu_0 f_0'\left(\frac{q_0}{p_0}\right) - \eta_0,\tag{8}$$

$$\delta^* = \nu_1 f_1' \left(\frac{q_1}{p_1} \right) - \eta_1, \tag{9}$$

$$\delta^* = \begin{cases} 1, & q_1 > \lambda q_0 \\ \kappa \in (0, 1), & q_1 = \lambda q_0 \\ 0, & q_1 < \lambda q_0 \end{cases}$$
(10)

where f'_0 , f'_1 denote the (sub)derivatives of f_0 and f_1 , η_0 , η_1 denote the Lagrange multipliers corresponding to the constraints (7), and ν_0 , ν_1 denote the Lagrange multipliers corresponding to the constraints (6). Since f_0 and f_1 are convex, their (sub)derivatives are nondecreasing. Moreover, the inverse functions g_0 and g_1 , which are implicitly defined by

$$g_0(f'_0(x)) = x$$
 and $g_1(f'_1(x)) = x$ $\forall x \in \mathbb{R}_{\geq 0},$

exist and are nondecreasing as well. Solving (8) and (9) for q_0 and q_1 yields

$$q_0 = g_0 \left(\frac{\lambda (1 - \delta^*) + \eta_0}{\nu_0} \right) p_0, \tag{11}$$

$$q_1 = g_1 \left(\frac{\delta^* + \eta_1}{\nu_1}\right) p_1.$$
 (12)

Combining (10), (11) *and* (12), *it follows that the least favorable densities are of the form*

$$q_0 = \begin{cases} b_0 p_0, & \delta^* = 0\\ \frac{1}{\lambda} q_1, & \delta^* \in (0, 1)\\ a_0 p_0, & \delta^* = 1 \end{cases}$$
(13)

and

$$q_1 = \begin{cases} a_1 p_1, & \delta^* = 0\\ \lambda q_0, & \delta^* \in (0, 1) \\ b_1 p_1, & \delta^* = 1 \end{cases}$$
(14)

where

and

$$a_0 = g_0\left(\frac{\eta_0}{\nu_0}\right) \le b_0 = g_0\left(\frac{\lambda + \eta_0}{\nu_0}\right) \tag{15}$$

$$y_1 = g_1\left(\frac{\eta_1}{\nu_1}\right) \le b_1 = g_1\left(\frac{1+\eta_1}{\nu_1}\right).$$
 (16)

Note that since q_0 and q_1 are valid densities, a_0, b_0 and a_1, b_1 are nonnegative. The next step of the proof is to show that the least favorable densities in (13) and (14) can be written as

$$q_0 = \min\{b_0 p_0, \max\{\frac{1}{\lambda}q_1, a_0 p_0\}\},$$
(17)

$$q_1 = \min\{b_1 p_1, \max\{\lambda q_0, a_1 p_1\}\}.$$
(18)

Only (17) is shown here since the proof for (18) can be given analogously. From (10) and (13) it follows that on $\{x \in \mathcal{X} : \delta(x) = 0\}$

$$q_1 < \lambda q_0 \quad \Rightarrow \quad \frac{1}{\lambda} q_1 < q_0 = b_0 p_0,$$
 (19)

on $\{x \in \mathcal{X} : \delta(x) \in (0,1)\}$

$$q_1 = \lambda q_0 \quad \Rightarrow \quad \frac{1}{\lambda} q_1 = q_0,$$
 (20)

and on $\{x \in \mathcal{X} : \delta(x) = 1\}$

$$q_1 > \lambda q_0 \quad \Rightarrow \quad \frac{1}{\lambda} q_1 > q_0 = a_0 p_0.$$
 (21)

Combining (19), (20), and (21) yields (17). From [14, Theorem 4], it follows immediately that (17) and (18) are fixed sample size least favorable for a density-band model with bounds

$$p'_0 = a_0 p_0 \le b_0 p_0 = p''_0,$$

 $p'_1 = a_1 p_1 \le b_1 p_1 = p''_1.$

5. DISCUSSION

The result presented in the previous section states that every single sample minimax optimal test under f-divergence-ball uncertainty is fixed sample size minimax optimal under the equivalent densityband uncertainty. This not only makes it possible to use single sample results for fixed sample size tests without sacrificing minimax optimality, but also to specify the exact sets of distributions for which the minimax property holds. In this sense, the theorem lifts the fdivergence-ball model to the same level of usefulness as the classic outlier models, whose single sample results automatically carry over to the fixed sample case. In addition to this generalization, the fact that for every f-divergence-ball model an equivalent densityband model can be constructed offers some deeper insights and also suggests an alternative approach to the design of robust tests under f-divergence-ball uncertainty.

One useful aspect of the equivalent band model is that it simplifies comparing the amount and type of uncertainty that is allowed for by different f-divergence-ball models. Such comparisons are nontrivial since the ε -tolerances in (3) do not directly translate to contamination ratios and might be of different scales altogether. While, for example, the Kullback-Leibler divergence can take on any nonnegative value, the Hellinger distance is bounded between zero and one. In such cases, one cannot simply compare the ε -tolerances in order to compare the maximum amount of uncertainty in the distributions. The corresponding band model, however, offers a way to make such comparisons possible. The lower bounds on the density functions, a_0p_0 and a_1p_1 , determine how much probability mass the nominal distributions contribute at least, namely a_0 and a_1 . Consequently, the outliers can at most contribute the remaining probability masses $1 - a_0$ and $1 - a_1$, which can hence be interpreted as contamination ratios. The larger they are, the more uncertainty a model allows. On the other hand, the upper bounds, b_0p_0 and b_1p_1 , offer an insight into what type of uncertainty is allowed. For $b_0, b_1 \gg 1$ the contamination is almost unconstrained, which corresponds to gross outliers. For $b_0, b_1 \approx 1$, the outlier distributions are close to the nominals, which corresponds to more subtle model mismatches. In general, the outlier distributions under \mathcal{H}_i , $i \in \{0, 1\}$, are constrained to lie within the set

$$\left\{ H \in \mathcal{M}_{\mu} : h \le \frac{b_i - a_i}{1 - a_i} p_i \right\}$$

This interpretation of a density-band model as a constrained ε contamination model often offers a useful intuition for the amount and type of contamination that cannot be obtained by inspection of the *f*-divergence-ball model.

There are several ways to determine the coefficients a_0 , b_0 and a_1 , b_1 in practice. If expressions for the least favorable distributions can already be found in the literature, the coefficients can be determined by a simple comparison with the expressions in (13) and (14). If the least favorable densities are unknown, the KKT conditions (8)–(10) can be solved for ν_0 , ν_1 and η_0 , η_1 . The scaling coefficients can then be calculated according to (15) and (16). In practice, however, this approach might be prohibitively complex.

An alternative to solving the KKT conditions for ν_0, ν_1 and η_0, η_1 is to solve them directly for a_0, b_0 and a_1, b_1 . From the result in the previous section, it follows that the least favorable densities are of the form (17) and (18). For given scaling coefficients a_0, b_0 and a_1, b_1 , these equations can be solved for q_0 and q_1 by finding a threshold λ so that the right hand sides of (17) and (18) are valid densities, i.e., they integrate to one. Finally, an outer search over a_0, b_0 and a_1, b_1 can be performed such that the primal constraints are fulfilled, i.e.,

$$\int_{\mathcal{X}} f_0\left(\frac{q_0}{p_0}\right) p_0 \,\mathrm{d}\mu = \varepsilon_0 \quad \text{and} \quad \int_{\mathcal{X}} f_1\left(\frac{q_1}{p_1}\right) p_1 \,\mathrm{d}\mu = \varepsilon_1.$$

This approach can be expected to be less efficient than a solution that exploits properties of a given function f, but is applicable in general and does not require prior analysis of the problem.

Yet another option to determine the scaling coefficients is to directly solve the primal problem (5) using a suitable convex optimization algorithm. Even if the latter does not calculate the dual variables explicitly, a_0 , b_0 and a_1 , b_1 can be obtained from the ratio of the least favorable and the nominal densities

$$\frac{q_0}{p_0} = \min\left\{b_0, \max\left\{\frac{1}{\lambda}\frac{q_1}{p_0}, a_0\right\}\right\},
\frac{q_1}{p_1} = \min\left\{b_1, \max\left\{\lambda\frac{q_0}{p_1}, a_1\right\}\right\}.$$
(22)



Fig. 1. Least favorable densities and equivalent density bands for uncertainty sets $\mathcal{P}_{-\log}(\mathcal{N}(-1,1), 0.03)$ and $\mathcal{P}_{-\log}(\mathcal{N}(1,2), 0.02)$.

By inspection of (22), the scaling coefficients can be identified form the regions where the likelihood ratio is constant.

The result in Section 4 also motivates further research into how exactly the two uncertainty models are related. Given equivalent fdivergence balls and density bands, does one uncertainty set contain the other, i.e., is one model a relaxation of the other one? Does every band model whose bounds are constructed by scaling a nominal density admit an equivalent f-divergence-ball model? If so, how can the corresponding function f be determined? Another question that might be asked is, whether similar equivalences exist for other types of uncertainty models as well and in how far there is a hierarchy between these models, i.e., whether the set of all possible least favorable distributions under one model is a sub- or super-set of all possible least favorable distributions under another model.

6. EXAMPLE

In order to highlight the connection to existing results, we consider the example from [10, Sec. VI.A], where the nominal distributions under \mathcal{H}_0 and \mathcal{H}_1 are chosen as $P_0 = \mathcal{N}(-1, 1)$ and $P_1 = \mathcal{N}(1, 2)$, respectively, and $\mathcal{N}(m, \sigma^2)$ denotes a Gaussian distribution with mean m and variance σ^2 . The Kullback–Leibler divergence is used as a distance measure, i.e., $f_0(x) = f_1(x) = -\log(x)$, the tolerances are chosen as $\varepsilon_0 = 0.03$, $\varepsilon_1 = 0.02$, and the likelihood ratio threshold is set to $\lambda = 1$. Using Theorem 2 in [10], the least favorable densities for this model can be calculated efficiently by solving two integral equations. The coefficients for the corresponding band model can be identified by comparing (6) in [10] to (13) and (14) in this paper. For the numerical values given above, they calculate to $a_0 \approx 0.9047$, $a_1 \approx 2.2519$, $b_0 \approx 0.8319$, and $b_1 \approx 1.3009$. The least favorable densities and the equivalent density bands are depicted in Fig. 1. Interestingly, the difference in the radii of the fdivergence balls is not reflected in the contamination ratio, which is smaller under $\mathcal{H}_0 \approx 10\%$) than under $\mathcal{H}_1 \approx 17\%$). However, the band under \mathcal{H}_0 is wider, meaning that it allows for larger deviations from the nominal distribution. Under \mathcal{H}_1 , the contamination ratio is higher, but the outlier distribution is much more restricted.

This example illustrates how existing results on robust tests under f-divergence-ball uncertainty can be used to construct fixed sample size minimax optimal tests for the equivalent density-band uncertainty sets and how the latter provide additional insight into the amount and type of contamination induced by the uncertainty model.

7. REFERENCES

- O. Zeitouni, J. Ziv, and N. Merhav, "When is the generalized likelihood ratio test optimal?," *IEEE Transactions on Information Theory*, vol. 38, no. 5, pp. 1597–1602, 1992.
- [2] S. A. Kassam and H. V. Poor, "Robust techniques for signal processing: A survey," *Proceedings of the IEEE*, vol. 73, no. 3, pp. 433–481, 1985.
- [3] A. L. McKellips and S. Verdu, "Worst case additive noise for binary-input channels and zero-threshold detection under constraints of power and divergence," *IEEE Transactions on Information Theory*, vol. 43, no. 4, pp. 1256–1264, 1997.
- [4] A. L. McKellips and S. Verdu, "Maximin performance of binary-input channels with uncertain noise distributions," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 947–972, 1998.
- [5] B. C. Levy, "Robust Hypothesis Testing With a Relative Entropy Tolerance," *IEEE Transactions on Information Theory*, vol. 55, no. 1, pp. 413–421, 2009.
- [6] G. Gül and A. M. Zoubir, "Robust hypothesis testing for modeling errors," in 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, 2013, pp. 5514– 5518.
- [7] G. Gül and A. M. Zoubir, "Robust hypothesis testing with squared Hellinger distance," in *Proc. of the 22nd European Signal Processing Conference (EUSIPCO)*, g, Ed., 2014, pp. 1083–1087.
- [8] G. Gül and A. M. Zoubir, "Robust hypothesis testing with composite distances," in 2014 IEEE Workshop on Statistical Signal Processing (SSP), 2014, pp. 432–435.
- [9] G. Gül and A. M. Zoubir, "Robust Hypothesis Testing with α -Divergence," *IEEE Transactions on Signal Processing*, vol. 64, no. 18, pp. 4737–4750, Sept 2016.
- [10] G. Gül and A. M. Zoubir, "Minimax Robust Hypothesis Testing," *IEEE Transactions on Information Theory*, vol. 63, no. 9, pp. 5572–5587, 2017.
- [11] P. J. Huber, "A Robust Version of the Probability Ratio Test," Annals of Mathematical Statistics, vol. 36, no. 6, pp. 1753– 1758, 1965.
- [12] A. Wald, *Sequential Analysis*, Wiley, Hoboken, New Jersey, USA, 1947.
- [13] R. Christensen, "Testing Fisher, Neyman, Pearson, and Bayes," *The American Statistician*, vol. 59, no. 2, pp. 121–126, 2005.
- [14] M. Fauß and A. M. Zoubir, "Old Bands, New Tracks— Revisiting the Band Model for Robust Hypothesis Testing," *IEEE Transactions on Signal Processing*, vol. 64, no. 22, pp. 5875–5886, 2016.
- [15] S. A. Kassam, "Robust hypothesis testing for bounded classes of probability densities (Corresp.)," *IEEE Transactions on Information Theory*, vol. 27, no. 2, pp. 242–247, 1981.
- [16] M. Guignard, "Generalized Kuhn–Tucker Conditions for Mathematical Programming Problems in a Banach Space," *SIAM Journal on Control*, vol. 7, no. 2, pp. 232–241, 1969.