ROBUST BINARY HYPOTHESIS TESTING UNDER CONTAMINATED LIKELIHOODS

Dennis Wei and Kush R. Varshney

Thomas J. Watson Research Center, IBM Research, Yorktown Heights, NY 10598, USA

ABSTRACT

In hypothesis testing, the phenomenon of label noise, in which hypothesis labels are switched at random, contaminates the likelihood functions. In this paper, we develop a new method to determine the decision rule when we do not have knowledge of the uncontaminated likelihoods and contamination probabilities, but only have knowledge of the contaminated likelihoods. In particular we pose a minimax optimization problem that finds a decision rule robust against this lack of knowledge. The method simplifies by application of linear programming theory. Motivation for this investigation is provided by problems encountered in workforce analytics.

Index Terms— label noise, linear programming, minimax, signal detection theory, workforce analytics

1. INTRODUCTION

Label noise in hypothesis testing problems results in the crosscontamination of the likelihood functions and possible degradation in detection performance if not accounted for when determining a decision rule. In this paper, we propose a linear programming framework for robustly dealing with contaminated likelihoods. Specifically, we propose an algorithm for obtaining a minimax optimal decision rule under label noise that is applicable under general likelihood models.

We are motivated by problems encountered in workforce analytics: data-driven decision making to manage the human capital of a corporation. For example, decision makers may want to use human resources data to predict whether or not an employee will voluntarily resign within the next 12 months [1], or to determine whether an employee from another division is a suitable candidate to fill an open position on a team in their division, based on skills and expertise data about the employee.¹ We face label noise and contamination of hypotheses in both examples. In the voluntary resignation example, we can take all employees that resigned in the recent past as samples from the alternative hypothesis and all employees that are currently active as samples from the null hypothesis. However, among currently active employees, some will resign in the coming months. Therefore, we are not in a position to observe an uncontaminated null distribution. In the suitable candidate example, we can take all employees in the decision maker's team as samples from the alternate distribution and all other employees as samples from the null distribution. However, not all team members may be suitable for the open position and not all other employees are unsuitable (which is why this problem is posed in the first place). Thus in this example, we observe contaminated versions of both likelihoods.

The problem of contaminated likelihoods in binary hypothesis testing was recently studied in considerable generality in [2,3]. The

theoretical framework in the present work is largely guided by [2, 3]. These previous works assume that the true likelihoods have an irreducibility property (described more fully in Section 3) that allows consistency results to be established. However, the assumption of irreducibility is restrictive. It is not satisfied for example by two Gaussian distributions with different variances, nor is it likely to be satisfied by real-world distributions such as may be encountered in workforce analytics. A contribution of the current paper in Section 3 is to remove the irreducibility assumption and extend the analysis to arbitrary true likelihoods. Furthermore, the approach taken herein, described in Section 4, differs fundamentally from [2, 3] in focusing not on consistent learning of a particular contamination model, but rather on designing hypothesis tests that are robust to uncertainty in the model. In Section 5, the utility of the robust viewpoint is demonstrated in two numerical examples.

More broadly, various types of label noise have been studied in the machine learning literature, including random, adversarial, and observation-dependent, and noise that affects different classes symmetrically and asymmetrically [4]. However, the vast majority of that work has been devoted to specific supervised classification algorithms operating on finite training data, see references in [2,3]. In contrast, our work deals with the regime encountered in signal detection theory and hypothesis testing, not the regime with finite training samples. Therefore, we work with likelihood ratio tests and true error probabilities rather than with specific classification algorithms and generalization bounds. Somewhat more related is the mixture modeling approach of [5,6], which attempts to learn the contamination model using the EM algorithm. This approach however requires parametric assumptions on the true likelihoods that we do not make. In the statistical signal processing literature, notions of robustness similar to that discussed in Section 4 have been explored [7, 8], for example to cope with uncertainty in prior probabilities [9, 10] and likelihoods [11, 12], but not for contamination/label noise.

2. PROBLEM STATEMENT

We consider the binary hypothesis testing problem of deciding between a null hypothesis $H = h_0$ and an alternative hypothesis $H = h_1$ based on observation of a random variable Y. Under hypothesis $H = h_0$, Y follows the probability distribution P_0 , while under $H = h_1$, Y follows distribution P_1 . A decision rule \hat{H} is desired that maps every possible observation Y = y to either h_0 or h_1 . For a rule \hat{H} , define $R_0(\hat{H}) = \Pr(\hat{H} = h_1 \mid H = h_0)$ and $R_1(\hat{H}) = \Pr(\hat{H} = h_0 \mid H = h_1)$ to be the Type I and Type II error probabilities. In this paper we focus on the Bayesian formulation in which the hypotheses have prior probabilities $\Pr(H = h_0) = q_0$, $\Pr(H = h_1) = 1 - q_0$, and the performance measure is the Bayes risk

$$R_B(\hat{H}) = c_{01}q_0R_0(\hat{H}) + c_{10}(1-q_0)R_1(\hat{H}), \qquad (1)$$

where c_{01} and c_{10} are the costs of Type I and Type II errors.

¹The types of data used for these purposes are generally necessary to the corporation's operations and part of the employment contract, and as such do not involve special privacy considerations.

Given knowledge of the conditional distributions P_0 and P_1 , it is straightforward to construct a likelihood ratio test that minimizes the Bayes risk [13]. However, in the contaminated version of the problem considered herein, P_0 and P_1 are not known. Instead, we have access to the contaminated distributions

$$\dot{P}_0 = (1 - \pi_0)P_0 + \pi_0 P_1, \qquad (2a)$$

$$\tilde{P}_1 = (1 - \pi_1)P_1 + \pi_1 P_0, \tag{2b}$$

where the contamination proportions $\pi_0, \pi_1 \in [0, 1]$ are also unknown. The following constraint is placed on π_0, π_1 ,

$$\pi_0 + \pi_1 < 1, \tag{3}$$

to resolve an interchange ambiguity and with essentially no loss of generality. Indeed, if $\pi_0 + \pi_1 > 1$, then as noted in [2], interchanging P_0 and P_1 yields complementary proportions $1 - \pi_0$, $1 - \pi_1$ satisfying $(1 - \pi_0) + (1 - \pi_1) < 1$. If $\pi_0 + \pi_1 = 1$, then (2) implies that $\tilde{P}_0 = \tilde{P}_1$ and discrimination is not possible.

As discussed in [2], it is not possible in general to design a test \hat{H} that minimizes the Bayes risk (1) given only the contaminated distributions \tilde{P}_0 , \tilde{P}_1 and no knowledge of P_0 , P_1 , π_0 , π_1 . Therefore in this paper we revise the objective to that of choosing \hat{H} to be robust to the uncertainty in the true distributions P_0 , P_1 , subject to limited additional input. We note that in the absence of further conditions, there is a large range of possible solutions to (2). In particular, it cannot be ruled out that there is no contamination, i.e. $\pi_0 = \pi_1 = 0$, $P_0 = \tilde{P}_0$, and $P_1 = \tilde{P}_1$. In the sequel, we seek to identify conditions that require minimal knowledge of or assumptions on P_0 , P_1 , π_0 , π_1 while also restricting uncertainty in a meaningful way in terms of Bayes risk.

We focus in this paper on the population setting where the distributions \tilde{P}_0 and \tilde{P}_1 are known exactly. Our results can be extended fairly straightforwardly to the finite-sample setting where \tilde{P}_0 and \tilde{P}_1 are approximated using training data, for example following the learning-theoretic approach of [2]. In the finite-sample case, the lack of knowledge of P_0 , P_1 translates into an inability to draw samples from P_0 , P_1 .

3. CONTAMINATION MODEL THEORY

In this section we summarize results that precisely characterize the possible solutions (P_0, P_1, π_0, π_1) to the contamination model (2). These results generalize parallels in [2] as discussed shortly. Due to space constraints, we refer the reader to [14] for full proofs.

First we recall some definitions from [2]. For probability distributions P and Q, define the maximal mixture proportion $\nu^*(P,Q)$ as

$$\nu^*(P,Q) = \max\{\alpha \in [0,1] : \\ \exists \text{ probability distribution } S : P = \alpha Q + (1-\alpha)S\}.$$
(4)

One way of interpreting $\nu^*(P,Q)$ is as the infimum of the ratio p(x)/q(x) if P and Q have probability densities p(x) and q(x) [3, Lem. 5]. From this it can be seen that $\nu^*(P,Q)$ is not necessarily symmetric. If $\nu^*(P,Q) = 0$, P is said to be irreducible with respect to Q, and if $\nu^*(Q,P) = 0$ also, then P and Q are mutually irreducible. Many of the results in [2] depend on the assumption that the true distributions P_0 and P_1 are mutually irreducible. This assumption is relaxed in the present paper.

The first result below relates maximal mixture proportions between P_0 and P_1 to mixed counterparts involving both pure and contaminated distributions. Lemma 1. Under condition (3),

$$\nu^*(P_0, \tilde{P}_1) = \frac{\nu^*(P_0, P_1)}{1 - \pi_1 + \pi_1 \nu^*(P_0, P_1)},$$

$$\nu^*(P_1, \tilde{P}_0) = \frac{\nu^*(P_1, P_0)}{1 - \pi_0 + \pi_0 \nu^*(P_1, P_0)}.$$

Proof outline. The proof involves showing that a decomposition of P_0 in terms of P_1 and a distribution Q implies a decomposition of P_0 in terms of \tilde{P}_1 and Q, and vice versa. Combining the implications yields the equalities in the lemma. Details can be found in [14]. \Box

This lemma generalizes [2, Lem. 3], which states that $\nu^*(P_0, \tilde{P}_1) = 0$ if and only if $\nu^*(P_0, P_1) = 0$, and similarly for the second equation.

Given condition (3), the contamination model (2) has an equivalent representation as specified by [2, Lem. 1]:

$$\tilde{P}_0 = (1 - \tilde{\pi}_0)P_0 + \tilde{\pi}_0\tilde{P}_1, \qquad \tilde{\pi}_0 = \frac{\pi_0}{1 - \pi_1} \in [0, 1),$$
 (5a)

$$\tilde{P}_1 = (1 - \tilde{\pi}_1)P_1 + \tilde{\pi}_1\tilde{P}_0, \qquad \tilde{\pi}_1 = \frac{\pi_1}{1 - \pi_0} \in [0, 1).$$
 (5b)

This alternative form makes clear that once $(\tilde{P}_0, \tilde{P}_1)$ and the modified parameters $(\tilde{\pi}_0, \tilde{\pi}_1)$ (or equivalently (π_0, π_1)) are fixed, (P_0, P_1) are also specified exactly. Using (5), [2, Cor. 1] shows that $\tilde{\pi}_0$ and $\tilde{\pi}_1$ are uniquely determined under the irreducibility conditions $\nu^*(P_0, \tilde{P}_1) = \nu^*(P_1, \tilde{P}_0) = 0$. The next lemma provides general expressions for $\tilde{\pi}_0, \tilde{\pi}_1$ that do not require irreducibility. The proof of this lemma in [14] is based on an extension of [2, Cor. 1].

Lemma 2. The contamination model (5) has a unique solution in $(\tilde{\pi}_0, \tilde{\pi}_1)$ in terms of maximal mixture proportions:

$$\begin{split} \tilde{\pi}_0 &= \frac{\nu^*(\tilde{P}_0, \tilde{P}_1) - \nu^*(P_0, \tilde{P}_1)}{1 - \nu^*(P_0, \tilde{P}_1)},\\ \tilde{\pi}_1 &= \frac{\nu^*(\tilde{P}_1, \tilde{P}_0) - \nu^*(P_1, \tilde{P}_0)}{1 - \nu^*(P_1, \tilde{P}_0)}. \end{split}$$

Combining Lemmas 1 and 2 as detailed in [14] yields a characterization of the contamination proportions π_0 , π_1 .

Theorem 1. Under condition (3), we have the relations

$$\pi_0 + \nu^* (\tilde{P}_0, \tilde{P}_1) \pi_1 = \frac{\nu^* (\tilde{P}_0, \tilde{P}_1) - \nu^* (P_0, P_1)}{1 - \nu^* (P_0, P_1)},$$

$$\nu^* (\tilde{P}_1, \tilde{P}_0) \pi_0 + \pi_1 = \frac{\nu^* (\tilde{P}_1, \tilde{P}_0) - \nu^* (P_1, P_0)}{1 - \nu^* (P_1, P_0)}.$$

Since \tilde{P}_0 , \tilde{P}_1 and hence $\nu^*(\tilde{P}_0, \tilde{P}_1)$, $\nu^*(\tilde{P}_1, \tilde{P}_0)$ are assumed to be known, Theorem 1 can be interpreted as a system of equations relating π_0 , π_1 to the maximal proportions $\nu^*(P_0, P_1)$, $\nu^*(P_1, P_0)$ for the pure distributions. If $\nu^*(\tilde{P}_0, \tilde{P}_1)$, $\nu^*(\tilde{P}_1, \tilde{P}_0) < 1$, i.e., if $\tilde{P}_0 \neq \tilde{P}_1$, then this system is invertible and Theorem 1 describes a bijection.

Fig. 1 depicts the set of feasible (π_0, π_1) values given the contaminated maximal proportions $\nu^*(\tilde{P}_0, \tilde{P}_1), \nu^*(\tilde{P}_1, \tilde{P}_0)$. The solid outer lines correspond to the mutually irreducible case, namely $\nu^*(P_0, P_1) = \nu^*(P_1, P_0) = 0$ in Theorem 1, and the intersection of the lines is the solution characterized in [2, Prop. 3]. Theorem 1 generalizes to the interior of the region by specifying solutions for nonzero values of $\nu^*(P_0, P_1), \nu^*(P_1, P_0)$. In particular, the dashed lines in Fig. 1 are lines of constant $\nu^*(P_0, P_1)$ or $\nu^*(P_1, P_0)$ and are parallel to the boundary lines. This geometry is used in the next section to describe uncertainty in π_0, π_1 .



Fig. 1. Region of feasible contamination proportions (π_0, π_1) given contaminated distributions \tilde{P}_0 and \tilde{P}_1 .

4. CONTAMINATION-ROBUST HYPOTHESIS TESTING

This section discusses the determination of decision rules that are robust to uncertainty in the contamination proportions π_0 and π_1 . Defining $\boldsymbol{\pi} = (\pi_0, \pi_1)$, we rewrite the Bayes risk (1) as follows,

$$R_{\rm B}(\hat{H}, \boldsymbol{\pi}) = c_{01} q_0 R_0(\hat{H}, \boldsymbol{\pi}) + c_{10} (1 - q_0) R_1(\hat{H}, \boldsymbol{\pi}), \quad (6)$$

to make explicit the dependence on the contamination proportions. From (5), the two error probabilities under the true distributions P_0 , P_1 can be expressed as

$$R_0(\hat{H}, \boldsymbol{\pi}) = \frac{(1 - \pi_1)\tilde{R}_0(\hat{H}) - \pi_0(1 - \tilde{R}_1(\hat{H}))}{1 - \pi_0 - \pi_1}, \qquad (7a)$$

$$R_1(\hat{H}, \boldsymbol{\pi}) = \frac{(1 - \pi_0)\tilde{R}_1(\hat{H}) - \pi_1(1 - \tilde{R}_0(\hat{H}))}{1 - \pi_0 - \pi_1}.$$
 (7b)

The performance thus depends on the error probabilities $R_0(\hat{H})$, $\tilde{R}_1(\hat{H})$ under the contaminated distributions, which can be determined for fixed decision rule \hat{H} , and π_0 , π_1 , which are only partially known.

The set of possible (π_0, π_1) values is constrained by knowledge of \tilde{P}_0 and \tilde{P}_1 as shown in Fig. 1. In addition to these initial constraints, we also consider lower and/or upper bounds on π_0, π_1 and the maximal mixture proportions $\nu^*(P_0, P_1), \nu^*(P_1, P_0)$ for the pure distributions. As seen from Theorem 1 and Fig. 1, bounds on $\nu^*(P_0, P_1), \nu^*(P_1, P_0)$ correspond to linear inequalities in π_0, π_1 . It follows that the feasible region for (π_0, π_1) is in general a convex polygon, which we may represent as a system of linear inequalities:

$$\Pi = \{ \boldsymbol{\pi} : \mathbf{a}_i^T \boldsymbol{\pi} \le b_i, \ i = 1, \dots, m \}$$

with appropriate choices of $\mathbf{a}_i \in \mathbb{R}^2$ and $b_i \in \mathbb{R}$.

The additional bounds on π_0 , π_1 , $\nu^*(P_0, P_1)$, $\nu^*(P_1, P_0)$ may be provided by application-specific knowledge and past experience. For example, with voluntary resignation, we can examine the resignation rate historically and use it to roughly characterize or bound π_0 . Moreover, examining data from more than a year in the past, we can observe P_0 and P_1 without contamination because any employee who was active then and has not resigned yet is by definition not a contaminated sample. Such historical P_0 and P_1 can be used to bound present values of $\nu^*(P_0, P_1)$ and $\nu^*(P_1, P_0)$.² In the case of finding suitable internal candidates for openings, similar openings filled in adjacent groups can provide bounds on $\pi_0, \pi_1, \nu^*(P_0, P_1), \nu^*(P_1, P_0)$.

In this paper, the decision rule \hat{H} is chosen to minimize the Bayes risk subject to worst-case uncertainty in (π_0, π_1) within the set Π :

$$\hat{H}_{\rm B} = \arg\min_{\hat{\pi}} \max_{\boldsymbol{\pi} \in \Pi} R_{\rm B}(\hat{H}, \boldsymbol{\pi}).$$
(8)

Alternative formulations include minimizing the worst-case deviation from the true Bayes risk (instead of the absolute Bayes risk in (8)) and minimizing the average Bayes risk over Π with respect to some distribution for π . We leave these alternatives for future work.

The inner maximization in (8) can be restricted to a subset of the vertices of Π . For a vertex $\pi \in \Pi$, define $I(\pi) \subseteq \{1, \ldots, m\}$ to be the set of constraints $\mathbf{a}_i^T \pi \leq b_i$ that are met with equality (active constraints), and cone ($\{\mathbf{a}_i, i \in I(\pi)\}$) to be the cone formed by non-negative combinations of the corresponding \mathbf{a}_i . We use \mathbb{R}^2_- as a shorthand for the non-positive quadrant of \mathbb{R}^2 .

Lemma 3. Assume that \hat{H} satisfies $\tilde{R}_0(\hat{H}) + \tilde{R}_1(\hat{H}) \leq 1$. Let π^k , $k = 1, \ldots, V$, be the vertices of Π such that

$$\operatorname{cone}\left(\{\mathbf{a}_i, i \in I(\boldsymbol{\pi}^k)\}\right) \cap \mathbb{R}^2_- \neq \emptyset.$$
(9)

Then

$$\max_{\boldsymbol{\pi}\in\Pi} R_{\mathrm{B}}(\hat{H},\boldsymbol{\pi}) = \max_{k=1,\ldots,V} R_{\mathrm{B}}(\hat{H},\boldsymbol{\pi}^k).$$

Abbreviated proof. The restriction to vertices of Π follows from the fact that $R_{\rm B}(\hat{H}, \pi)$ is a linear-fractional function of π for fixed \hat{H} . This property is seen by substituting (7) into (6) to obtain

$$R_{\rm B}(\hat{H}, \boldsymbol{\pi}) = \frac{\mathbf{c}^T \boldsymbol{\pi} + d}{1 - \pi_0 - \pi_1},\tag{10}$$

where $\mathbf{c} \in \mathbb{R}^2$ and $d \in \mathbb{R}$ do not depend on $\boldsymbol{\pi}$ (explicit expressions are omitted here). Given (10), the maximization of $R_{\mathrm{B}}(\hat{H}, \boldsymbol{\pi})$ may be carried out as a search for the largest $t \geq 0$ for which the linear program

$$\max_{\boldsymbol{\pi}\in\Pi} \mathbf{c}^T \boldsymbol{\pi} + d - t(1 - \pi_0 - \pi_1) \tag{11}$$

has a non-negative optimal value, implying that the superlevel set $\{\pi \in \Pi : R_{\rm B}(\hat{H}, \pi) \geq t\}$ is non-empty. Since (11) is a linear optimization over a bounded polygon, there exists a vertex of Π that is optimal [15, Thm. 2.8]. This holds in particular for $t = \max_{\pi \in \Pi} R_{\rm B}(\hat{H}, \pi)$ and hence it is sufficient to consider only the vertices of Π in maximizing $R_{\rm B}(\hat{H}, \pi)$.

The restriction to vertices satisfying (9) is due to the KKT optimality condition for the maximization of $R_{\rm B}(\hat{H}, \pi)$:

$$\nabla_{\boldsymbol{\pi}} R_{\mathrm{B}}(\hat{H}, \boldsymbol{\pi}) = \sum_{i \in I(\boldsymbol{\pi})} \mu_i \mathbf{a}_i, \quad \mu_i \ge 0, \quad (12)$$

which is a necessary condition because Π is defined by linear inequalities [16, Prop. 3.3.7]. It is shown in [14] that the gradient with respect to π on the left-hand side of (12) belongs to \mathbb{R}^2_- , while the right-hand side of (12) can range over cone({ $\mathbf{a}_i, i \in I(\pi)$ }). Therefore it suffices to consider vertices satisfying (9).

²One may ask why historical P_0 and P_1 cannot simply be used to determine the decision rule in the present; this is not possible in dynamic business environments where the resignation rate within job roles, skill sets, professions, and organizational units—which are all observations to predict resignation—changes rapidly due to technology trends and management changes. It is the level of differentiation between the classes that we assume does not change much over time, allowing us to bound $\nu^*(P_0, P_1)$ and $\nu^*(P_1, P_0)$.

Remark. The condition $\tilde{R}_0(\hat{H}) + \tilde{R}_1(\hat{H}) \leq 1$ is satisfied by any decision rule \hat{H} that is at least as good as random guessing. Hence no generality is lost.

Combining (8) and Lemma 3 yields

$$\hat{H}_{\rm B} = \operatorname*{arg\ min}_{\hat{H}} t \quad \text{s.t.} \quad R_{\rm B}(\hat{H}, \boldsymbol{\pi}^k) \le t, \quad k = 1, \dots, V.$$
 (13)

In the two-dimensional case considered here, the number V of vertices satisfying (9) is very small and π^1, \ldots, π^V are easily enumerated. Therefore (13) represents a significant simplification compared to (8). However, enumeration becomes increasingly difficult in higher dimensions that would arise in hypothesis testing with more than two hypotheses.

5. NUMERICAL EXAMPLES

In this section we illustrate the proposed minimax procedure via two examples with likelihoods that are not mutually irreducible: Gaussian distributions with different means and different variances, and exponential distributions with different inverse scale parameters. The Gaussians example provides a rough model for features that predict voluntary resignation, since features such as time since the last job promotion and annual performance rating tend to be approximately normal in many organizations. The exponentials example provides a rough model for abilities among a high-performing group, which arises when finding suitable candidates.

Consider $P_0 \sim \mathcal{N}(\mu_0, \sigma_0^2)$ and $P_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$ where $\mu_0 \neq \mu_1$ and, without loss of generality, $\sigma_0 < \sigma_1$. For this problem, the uncontaminated error probabilities for a likelihood ratio test with threshold value γ are:

$$R_0(\gamma) = Q\left(\frac{y^+ - \mu_0}{\sigma_0}\right) + Q\left(\frac{-y^- + \mu_0}{\sigma_0}\right)$$
$$R_1(\gamma) = 1 - Q\left(\frac{y^+ - \mu_1}{\sigma_1}\right) - Q\left(\frac{-y^- + \mu_1}{\sigma_1}\right),$$

where $Q(y) = \frac{1}{\sqrt{2\pi}} \int_y^\infty \exp(-y'^2/2) dy'$, and y^+ and y^- are the solutions to the quadratic equation:

$$(\sigma_1^2 - \sigma_0^2)y^2 + 2(\mu_1 \sigma_0^2 - \mu_0 \sigma_1^2)y + \mu_0^2 \sigma_1^2 - \mu_1^2 \sigma_0^2 - 2\sigma_0^2 \sigma_1^2 \ln\left(\gamma \frac{\sigma_1}{\sigma_0}\right) = 0.$$

We examine the situation in which $\mu_0 = 0$, $\mu_1 = 0.2$, $\sigma_0 = 1$, and $\sigma_1 = 2$. Additionally, for the Bayes risk, we consider the simple case when $q_0 = 0.5$ and $c_{01} = c_{10} = 1$. The true contamination proportions, unknown to an observer, are $\pi_0 = 0.2$ and $\pi_1 = 0.3$. These contamination proportions result in $\nu^*(\tilde{P}_0, \tilde{P}_1) = 0.2857$ and $\nu^*(\tilde{P}_1, \tilde{P}_0) = 0.7202$, which are observed. Additional information on the contamination gives us the constraints $\pi_0 \ge 0.05$ and $\pi_1 \ge 0.1$, as well as $\pi_0 + \nu^*(\tilde{P}_0, \tilde{P}_1)\pi_1 \ge 0.2$ and $\nu^*(\tilde{P}_1, \tilde{P}_0)\pi_0 +$ $\pi_1 \ge 0.25$. The last two inequalities follow from Theorem 1 and upper bounds on $\nu^*(P_0, P_1)$, $\nu^*(P_1, P_0)$. With these constraints, the polygon II has six vertices.

After performing the inner maximization of the minimax procedure, we find the vertex of Π that maximizes the Bayes risk to be (0.1619, 0.1334). This maximum Bayes risk is shown in Fig. 2(a) as a function of the threshold λ applied to the contaminated likelihood ratio (λ is related to γ through a transformation derived in [2, Prop. 1]). The minimum value of this function, i.e. the minimax Bayes risk we seek, is 0.3845.

The figure also shows the Bayes risk if we use the unknown true contamination proportions (which equals the uncontaminated Bayes risk) and the Bayes risk if we use the (0,0) point, i.e., we do not



Fig. 2. Bayes risk as a function of the threshold on the contaminated likelihood ratio for (a) Gaussian example and (b) exponential example: using unknown true contamination proportions (dashed), max solution (solid), and (0, 0) contamination proportions (dash-dot).

account for contamination. The minimum Bayes risk using the true contamination proportions is 0.3372 and the minimum when using (0,0) is 0.4186. The minimax solution is between these two values. Notably, it is less pessimistic than the default (0,0) solution. The solution under irreducibility [2] is not selected under the minimax criterion as it is too optimistic about the Bayes risk value.

As a second example, consider $P_0 \sim \mathcal{E}(\alpha_0)$ and $P_1 \sim \mathcal{E}(\alpha_1)$ where without loss of generality, $\alpha_0 < \alpha_1$. For this problem, the uncontaminated error probabilities for a likelihood ratio test threshold value γ are: $R_0(\gamma) = 1 - e^{-\alpha_0 y^*}$ and $R_1(\gamma) = e^{-\alpha_1 y^*}$, where $y^* = \ln\left(\frac{\alpha_0}{\alpha_1}\gamma\right)/(\alpha_0 - \alpha_1)$. We set $\alpha_0 = 1$ and $\alpha_1 = 2$ and (to economize on space for exposition) keep all other parameters the same as in the first example. With these exponential likelihoods and parameter settings, $\nu^*(\tilde{P}_0, \tilde{P}_1) = 0.7059$ and $\nu^*(\tilde{P}_1, \tilde{P}_0) = 0.3750$ and the resulting II has five vertices. The maximizing vertex is (0.1619, 0.1334) and the maximum Bayes risk is shown in Fig. 2. The minimax Bayes risk is 0.4130, which lies between the minimum Bayes risk with known contamination proportions, 0.3750, and the minimum Bayes risk using proportions (0, 0), 0.4375, in the same manner as the previous example.

6. CONCLUSION

In this paper, we have examined the problem of contaminated likelihood functions that arise due to label noise in hypothesis testing. In contrast to previous work on the subject which derived consistency results for the case when the likelihoods are mutually irreducible, we deal with arbitrary likelihoods and obtain decision rules robust to uncertainty in the contamination proportions. Toward this end, we have posed an optimization problem that is naturally subject to linear constraints and shown that its objective function is a linearfractional function. Therefore, the optimization problem reduces to linear programs that can be simplified using the KKT conditions into a search over certain vertices of the constraint set. We have shown the method on two numerical examples.

7. REFERENCES

- [1] M. Singh, K. R. Varshney, J. Wang, A. Mojsilović, A. R. Gill, P. I. Faur, and R. Ezry, "An analytics approach for proactively combating voluntary attrition of employees," in *Proc. IEEE Int. Conf. Data Min. Workshops*, Brussels, Belgium, Dec. 2012, pp. 317–323.
- [2] C. Scott, G. Blanchard, and G. Handy, "Classification with asymmetric label noise: Consistency and maximal denoising," in *JMLR W&CP, Conf. on Learning Theory (COLT)*, 2013, vol. 30, pp. 489–511.
- [3] C. Scott, G. Blanchard, G. Handy, S. Pozzi, and M. Flaska, "Classification with asymmetric label noise: Consistency and maximal denoising," arXiv:1303.1208v1, Mar. 2013.
- [4] B. Frénay and A. Kabán, Eds., *Special Issue on Advances in Learning with Label Noise*, Neurocomputing, in press.
- [5] N. Lawrence and B. Schölkopf, "Estimating a kernel Fisher discriminant in the presence of label noise," in *Proc. Int. Conf. Machine Learning (ICML)*, 2001.
- [6] C. Bouveyron and S. Girard, "Robust supervised classification with mixture models: Learning from data with uncertain labels," *J. Pattern Recogn.*, vol. 42, pp. 2649–2658, 2009.
- [7] S. A. Kassam and H. V. Poor, "Robust techniques for signal processing: A survey," *Proc. IEEE*, vol. 73, no. 3, pp. 433– 481, Mar. 1985.
- [8] A. M. Zoubir, V. Koivunen, Y. Chakhchoukh, and M. Muma, "Robust estimation in signal processing: A tutorial-style treatment of fundamental concepts," *IEEE Signal Process. Mag.*, vol. 29, no. 4, pp. 61–80, July 2012.
- [9] J. Jiao, L. Zhang, and R. D. Nowak, "Minimax-optimal bounds for detectors based on estimated prior probabilities," *IEEE Trans. Inf. Theory*, vol. 58, no. 9, pp. 6101–6109, Sept. 2012.
- [10] K. R. Varshney and L. R. Varshney, "Optimal grouping for group minimax hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 60, no. 10, pp. 6511–6521, Oct. 2014.
- [11] M. Feder and N. Merhav, "Universal composite hypothesis testing: A competitive minimax approach," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1504–1517, June 2002.
- [12] V. V. Veeravalli, T. Başar, and H. V. Poor, "Minimax robust decentralized detection," *IEEE Trans. Inf. Theory*, vol. 40, no. 1, pp. 35–40, Jan. 1994.
- [13] H. L. Van Trees, Detection, Estimation, and Modulation Theory, Wiley, New York, NY, 1968.
- [14] D. Wei and K. R. Varshney, "Robust binary hypothesis testing under contaminated likelihoods," arXiv:1410.0952, Oct. 2014.
- [15] D. Bertsimas and J. N. Tsitsiklis, *Introduction to Linear Opti*mization, Athena Scientific, Belmont, MA, 1997.
- [16] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, MA, 2nd edition, 1999.