Toward information privacy for the internet of things : a nonparametric learning approach

Sun, Meng; Tay, Wee Peng; He, Xin

2018

Sun, M., Tay, W. P., & He, X. (2018). Toward information privacy for the internet of things : a nonparametric learning approach. IEEE Transactions on Signal Processing, 66(7), 1734-1747. doi:10.1109/TSP.2018.2793871

<https://hdl.handle.net/10356/107063>

<https://doi.org/10.1109/TSP.2018.2793871>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org.remotexs.ntu.edu.sg/10.1109/TSP.2018.2793871>

Downloaded on 26 Apr 2023 19:30:19 SGT

Toward Information Privacy for the Internet of Things: A Nonparametric Learning Approach

Meng Sun^{ID}, *Student Member, IEEE*, Wee Peng Tay^{ID}, *Senior Member, IEEE*, and Xin He

Abstract—In an Internet of things network, multiple sensors send information to a fusion center for it to infer a public hypothesis of interest. However, the same sensor information may be used by the fusion center to make inferences of a private nature that the sensors wish to protect. To model this, we adopt a decentralized hypothesis testing framework with binary public and private hypotheses. Each sensor makes a private observation and utilizes a local sensor decision rule or privacy mapping to summarize that observation independently of the other sensors. The local decision made by a sensor is then sent to the fusion center. Without assuming knowledge of the joint distribution of the sensor observations and hypotheses, we adopt a nonparametric learning approach to design local privacy mappings. We introduce the concept of an empirical normalized risk, which provides a theoretical guarantee for the network to achieve information privacy for the private hypothesis with high probability when the number of training samples is large. We develop iterative optimization algorithms to determine an appropriate privacy threshold and the best sensor privacy mappings, and show that they converge. Finally, we extend our approach to the case of a private multiple hypothesis. Numerical results on both synthetic and real data sets suggest that our proposed approach yields low error rates for inferring the public hypothesis, but high error rates for detecting the private hypothesis.

Index Terms—Information privacy, decentralized hypothesis testing, decentralized detection, nonparametric, empirical risk, Internet of things.

I. INTRODUCTION

SENSOR networks have seen widespread applications in industrial, military, and civilian monitoring applications like intrusion detection, target tracking, leakage detection and fall detection [1]–[7]. In the emerging Internet of Things (IoT) paradigm, large numbers of sensors are deployed to enable sense-making and intelligent analytics based on the sensors' observations. This can be modeled using the decentralized detection framework [8]–[14], where each sensor makes an observation, summarizes this observation using a local decision

rule, and sends the summary to a fusion center. Based on the received sensor summaries, the fusion center then makes the final inference on a phenomenon of interest.

While the fusion center's role is to perform inference on a particular hypothesis of interest, there is nothing stopping it from using the received sensor information to infer another correlated hypothesis. An example is the deployment of home-monitoring video cameras in old folks' homes for fall detection. If the cameras transmit the raw video feed to a fusion center, the fusion center can not only use these video feeds for fall detection, but also has the potential to intrude on the privacy of the home inhabitants. The camera sensors therefore need to perform intelligent observation summary with a suitable sensor rule in order to limit the amount and quality of information they send to the fusion center. Another example is when an insurance company wishes to determine if a person has a particular pre-existing medical condition using medical records from hospitals the person has been treated at. However, these medical records may reveal more than the particular condition that the insurance company is investigating. The hospitals will need to decide what to send to the insurance company to avoid disclosing the person's other medical conditions. Although this latter example is not in the context of sensor networks, we can see that it nevertheless falls in the framework of how to preserve privacy in decentralized detection while still enabling the fusion center to make an inference on a particular hypothesis. In this paper, we call a hypothesis a public hypothesis if its inference or detection is to be achieved by a sensor network specifically designed for this purpose. We call a hypothesis a private hypothesis if it can also be inferred based on the same sensor observations, but whose true state the sensor network wishes to protect. We call preventing the accurate inference of the private hypothesis *information privacy* (see Section II for a precise technical definition).

The main focus of this paper is to protect information privacy by making it difficult for the fusion center to perform inference on the private hypothesis. It has been shown in [15]–[17] that the error decay rate at the fusion center increases with the quality of information that the sensors convey to it. Therefore, by appropriately “degrading” the information sent by each sensor to the fusion center, we aim to achieve a good tradeoff between privacy leakage and the ability of the fusion center to infer the public hypothesis.

A. Related Work

With ubiquitous IoT devices monitoring every aspect of a user's life, privacy is a main consideration for users when

Manuscript received June 29, 2017; revised November 7, 2017; accepted December 29, 2017. Date of publication January 15, 2018; date of current version February 13, 2018. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Pierre Borgnat. This work was supported in part by the Singapore Ministry of Education Academic Research Fund Tier 2 Grant MOE2014-T2-1-028 and the NTU-NXP Intelligent Transport System Test-Bed Living Lab Fund S15-1105-RF-LLF from the Economic Development Board, Singapore. (Corresponding author: Meng Sun.)

The authors are with the Department of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798 (e-mail: MSUN002@e.ntu.edu.sg; wptay@ntu.edu.sg; hexin@ntu.edu.sg).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2018.2793871

adopting IoT technologies. Sensitive personal data like lifestyle preferences and location information may be abused for unwanted advertisement purposes or for more nefarious objectives like unauthorized surveillance. Privacy in IoT networks can be classified as data privacy and inference privacy.

Data privacy refers to the protection of sensor data from unauthorized parties. Achieving data privacy has been comprehensively addressed by methods typically involving encryption or perturbation [18]–[25].

Although data privacy ensures that each sensor datum is protected, it does not stop a fusion center from inferring about a private hypothesis if statistical information about it is still present in the aggregated data.

The focus of this paper is inference privacy, whose aim is to prevent the fusion center from using its received information to accurately infer a private hypothesis. The paper [26] utilized source coding so that a receiver can decode one source within a prescribed distortion tolerance, while ensuring that the mutual information between the decoded sequence and another correlated private source is lower than a threshold. The author also analyzed the privacy-utility tradeoff under this formulation. The reference [27] extends the result of [26], and designed a framework to quantify the privacy-utility tradeoffs. An information-theoretic scheme for information privacy was proposed by [28], which formulated a convex program to find the privacy-preserving mapping to minimize the mutual information between the private hypothesis and information received at the fusion center, while satisfying certain utility constraints. In [29], the minimum Bayesian error probability for a fusion center to infer a private phenomenon is used as the privacy metric, and a person-by-person optimization approach is proposed to find sensor decision rules, but inference privacy is not guaranteed in general.

All the aforementioned approaches are designed to achieve a good utility-privacy tradeoff, but assume knowledge of the joint distribution of sensor observations, public and private hypotheses. If a mismatched joint distribution is utilized, the utility-privacy tradeoff is impacted [30]. Knowing the underlying joint distribution may not be practical in an IoT network. Therefore, in this paper, we propose a nonparametric hypothesis testing approach with information privacy constraints.

A nonparametric approach to inference privacy was proposed by [31], [32], in which sensor data is mapped to a subspace before being made available to the fusion center. These methods can be adopted in a centralized platform like cloud computing, but impractical for a sensor network because it requires the use of a *trusted data curator* to first aggregate observations from all sensors, and then performing the privacy mapping on the aggregated data. Furthermore, no theoretical guarantees of the level of privacy achievable are provided in [31], [32]. In a decentralized architecture like an IoT network, each sensor makes its local decision based solely on its own observation, and independently of the other sensors.

A nonparametric decentralized detection method was introduced by [33], which proposed the use of a kernel-based method to learn the optimal sensor decision rules from a given set of labeled training data. Subsequently, [34] extended this method

using a weighted kernel to allow sensor selection in the decentralized detection procedure. These works however do not address the inference privacy issue described above. Following these works, we assume that a set of labeled training data is available and employ a kernel-based approach to learn the sensor rules subject to a privacy metric called information privacy [28].

B. Our Contributions

In this paper, we develop an information privacy preserving framework for nonparametric hypothesis testing and algorithms to realize our framework. Our main contributions are as follows:

- 1) We adopt the concept of ϵ -information privacy from [28], and show that a sufficient condition to achieve ϵ -information privacy when the private hypothesis is binary, is to ensure that the average of the Type I and II detection error probabilities are large. We show that contrary to intuition, a large Bayes detection error for the private hypothesis does not ensure information privacy.
- 2) Since we do not assume knowledge of the underlying joint distribution and we adopt a nonparametric optimization framework, we introduce the concept of (ϵ, δ) -information privacy, which is a weak form of ϵ -information privacy. When the private hypothesis is binary, we propose a privacy metric constraint, which we call the *empirical normalized risk*, and show that under some mild technical assumptions, this achieves (ϵ, δ) -information privacy for any $\delta > 0$ when the training sample size becomes large.
- 3) When both public and private hypotheses are binary, we propose a nonparametric privacy-aware optimization framework and iterative algorithms to learn an appropriate privacy threshold for our empirical normalized risk, and the optimal sensor privacy mappings. We show that both algorithms converge to the critical points of their respective objective functions.
- 4) We extend our optimization framework to the case where the private hypothesis is m -ary with $m > 2$, and provide a sufficient condition to achieve information privacy. This condition translates into $m - 1$ empirical normalized risk constraints in our optimization framework, which achieves (ϵ, δ) -information privacy for any $\delta > 0$ when the training sample size becomes large.
- 5) We verify the performance of our algorithm on both simulated data and real data. Our experiments suggest that our approach can achieve low error rates for inferring the public hypothesis, but high error rates for detecting the private hypothesis.

This paper is an extension of our conference paper [35], which utilized a nonparametric approach to learn sensor decision rules under a Bayesian error probability privacy constraint. As shown in this paper, that approach does not guarantee information privacy in general.

The rest of this paper is organized as follows. In Section II, we present our system model and assumptions. We also define information privacy and present a sufficient condition to achieve it. In Section III, we propose a nonparametric privacy-aware

Fig. 1. An IoT network with public hypothesis H and private hypothesis G .

optimization framework based on an empirical normalized risk constraint, to obtain the sensor privacy mappings. We show that our approach achieves information privacy with high probability as the training sample size becomes large. We then propose an iterative algorithm for the optimization problem. We present simulation results to verify the effectiveness of the proposed algorithms in Section IV, and we conclude in Section V.

Notations: We use capital letters like X to denote random variables or vectors, lowercase letters like x for deterministic scalars, and boldface lowercase letters like $\mathbf{x}$ for deterministic vectors. We use $\mathbb{R}$ to denote the set of real numbers, and Γ^c to be the complement of the set Γ . The indicator function $\mathbf{1}_{\{A\}}$ takes value 1 iff the clause A is true. We let $(x)_+ = \max\{x, 0\}$. We assume that all random variables are defined on the same underlying probability measure space with probability measure $\mathbb{P}$, and $\mathbb{E}$ is the associated expectation operator. In some cases, for clarity, we use $\mathbb{E}_X$ to emphasize that the expectation is with respect to (w.r.t.) X . We use $p_X(\cdot)$ to denote the probability mass function of X , and $p_{X|Y}(\cdot | \cdot)$ to denote the conditional probability mass function of X given Y .

II. PROBLEM FORMULATION AND INFORMATION PRIVACY

In this section, we first describe our system model and assumptions, and then discuss the connection of our problem setup with information privacy and differential privacy. We provide a sufficient condition to achieve information privacy for an IoT network.

A. System Model

We consider a decentralized detection network as shown in Fig. 1. Suppose that two hypotheses H and G each takes binary values in $\{-1, +1\}$. Each sensor $t \in \{1, 2, \dots, s\}$, makes a noisy observation $X^t \in \mathcal{X}$ of (H, G) , where $\mathcal{X} = \{1, 2, \dots, |\mathcal{X}|\}$. It then summarizes its observation using a local decision rule or privacy mapping $\gamma^t: \mathcal{X} \mapsto \mathcal{Z}$, with $\mathcal{Z} = \{1, 2, \dots, |\mathcal{Z}|\}$, and transmits $Z^t = \gamma^t(X^t)$ to a fusion center. Each sensor's local decision rule is allowed to be a probabilistic mapping from $\mathcal{X}$ to $\mathcal{Z}$. The transmission from each sensor to the fusion center is constrained to a limited number of bits, so that $|\mathcal{X}| \geq |\mathcal{Z}|$. This models an ad hoc IoT

network with low-power devices that may be battery operated. For example, the NB-IoT standard [36] is developed for low-cost and low-power devices with limited communication bandwidths. In this paper, we do not require that the sensor observations $X^t, t = 1, \dots, s$ are independent.

Let $X = (X^1, \dots, X^s)$ and $Z = (Z^1, \dots, Z^s)$. Based on the received messages Z , the fusion center makes a decision $\hat{H} = \gamma_H(Z) \in \{-1, +1\}$ about the state of the hypothesis H . We consider H to be the public hypothesis that the sensors want the fusion center to infer correctly. On the other hand, G is a private hypothesis that the sensors wish to hide from the fusion center. The fusion center however is curious, and after receiving the local decisions from the sensors, implements a decision rule $\hat{G} = \gamma_G(Z) \in \{-1, +1\}$ to infer the private hypothesis G . Our goal is to find, for each sensor $t = 1, \dots, s$, a local decision rule or privacy mapping γ^t to minimize the error probability $\mathbb{P}(H \neq \gamma_H(Z))$, while making it difficult for any fusion rule γ_G the fusion center may employ to detect G .

In the following Section II-B, we show that by making the average of the Type I error $\mathbb{P}(\gamma_G(Z) = -1 | G = 1)$ and Type II error $\mathbb{P}(\gamma_G(Z) = 1 | G = -1)$ sufficiently large, we achieve *information privacy* [28] for the private hypothesis G . However, to find the optimal local sensor decision rules in order to facilitate the inference of H while ensuring that both the Type I and II errors are large, one needs to know the joint probability distribution $p_{X,H,G}$ a priori. This assumption is impractical for many IoT applications as the underlying joint distribution may be difficult to model accurately. Therefore, we assume that this distribution is unknown, but we are given a set of n independent and identically distributed (i.i.d.) training data $(\mathbf{x}_i, h_i, g_i)_{i=1}^n$ sampled from $p_{X,H,G}$.¹ We adopt the framework of empirical risk minimization as in [33] to design the local sensor decision rules, which we discuss in detail in Section III-A.

B. Information Privacy

To prevent the fusion center from inferring the true state of the private hypothesis G based on the received sensor messages Z , we need to design Z so that the prior and posterior distributions of G are close to each other. This is the definition of ϵ -information privacy given by [28], which we recall below. Note that this differs from the more commonly used differential privacy concept [21], [37], which [28] shows is weaker than information privacy in general.

Definition 1 (Information privacy): Let G and Z be random variables. For $\epsilon > 0$, G given Z or $p_{G|Z}$ has ϵ -information privacy, if for almost surely any $(G, Z) = (g, \mathbf{z})$, we have

$$e^{-\epsilon} \leq \frac{p_{G|Z}(g | \mathbf{z})}{p_G(g)} \leq e^\epsilon. \quad (1)$$

The value ϵ is called the information privacy budget.

In an IoT network, both data and information privacy are important design aspects. Achieving data privacy has been comprehensively addressed in [25], [38]–[41]. In this paper, we focus on

¹In the sequel, we adopt the following convention: a superscript index corresponds to a sensor index (e.g., X^t), while a subscript index corresponds to either a training data index (e.g., $\mathbf{x}_i$) or a hypothesis.

the information privacy aspect. In [28], an optimization framework based on minimizing the mutual information between G and Z is proposed. This leads to a weaker privacy guarantee than information privacy. In the following, we show that if G is a binary hypothesis, then constraining the total probability of error in detecting G leads to information privacy for G under some technical conditions. We make following assumption.

Assumption 1: The priors $p_G(1), p_G(-1) > 0$. The supports of the conditional distributions of Z given $G = 1$ and $G = -1$ are the same, and denoted as $\mathcal{D}$.

Assumption 1 leads to no loss of generality and is required to avoid trivial cases where G is perfectly detectable. For any detector $\gamma(Z)$ for G based on Z , let

$$R(\gamma) = \frac{1}{2}(\mathbb{P}(\gamma(Z) = -1 | G = 1) + \mathbb{P}(\gamma(Z) = 1 | G = -1))$$

be the average of the Type I and II error probabilities. For each $\mathbf{z} \in \mathcal{D}$, let

$$\ell(\mathbf{z}) = \frac{p_{Z|G}(\mathbf{z} | 1)}{p_{Z|G}(\mathbf{z} | -1)},$$

and

$$c = \min \left\{ \mathbb{P} \left(\ell(\mathbf{z}) = \min_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) \middle| G = -1 \right), \right. \\ \left. \mathbb{P} \left(\ell(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) \middle| G = 1 \right) \right\}. \quad (2)$$

From Assumption 1, since $0 < \ell(\mathbf{z}) < \infty$ for all $\mathbf{z} \in \mathcal{D}$, $c > 0$.

The following result shows that a small information privacy budget is essentially equivalent to a large $\min_{\gamma} R(\gamma)$. Since using (an empirical proxy of) $\min_{\gamma} R(\gamma)$ as a privacy constraint is more convenient than imposing (1) directly (which leads to $2|\mathcal{Z}|$ constraints when G is binary), the following proposition allows us to formulate a compact privacy constraint in Section III.

Proposition 1: Suppose that Assumption 1 holds.

- i) Let $\epsilon > 0$ be a sufficiently small positive constant. If G given Z has ϵ -information privacy, then $R(\gamma_G) \geq \theta/(2 \max_g p_G(g))$, for all $\theta \in [0, 1/2]$ that satisfies

$$\mathcal{H}(\theta) \leq \mathcal{H}(G) - \epsilon, \quad (3)$$

with $\mathcal{H}(\cdot)$ being the binary entropy function [42].

- ii) If $\min_{\gamma} R(\gamma) \geq \theta \in [0, 1/2]$, G given Z achieves ϵ -information privacy where $\epsilon = \log \frac{c}{(c+2\theta-1)_+}$.

Proof: See Appendix A. ■

If the prior probabilities $p_G(1) = p_G(-1) = 1/2$, $\min_{\gamma} R(\gamma)$ is the Bayesian error probability. When the prior for G is not uniform, the Bayes detector γ_G^* with a large Bayes error $\mathbb{P}(\gamma_G^*(Z) \neq G)$ do not guarantee information privacy as defined in Definition 1. An example is shown in Example 1, and a simulation to demonstrate this is given in Section IV-A2. To ensure information privacy, we require that both the Type I and II error probabilities are large, which is equivalent to having the risk $\min_{\gamma} R(\gamma)$ to be sufficiently large.

Example 1: Suppose $p_G(-1) < 1/2$, and G and $Z \in \{\mathbf{z}_1, \mathbf{z}_2\}$ have the joint distribution shown in Table I. Then, for all sufficiently large d , the optimal Bayesian error probability detector declares $G^* = 1$ for all Z . The Bayes error

TABLE I
JOINT DISTRIBUTION OF Z AND G

$P_{G,Z}(g, \mathbf{z})$		Z	
		$\mathbf{z}_1$	$\mathbf{z}_2$
G	-1	$p_G(-1)/d$	$(1 - 1/d)p_G(-1)$
	1	$p_G(1) - p_G(-1)$	$p_G(-1)$

$\mathbb{P}(G^* \neq G) = p_G(-1)$, which is the largest possible Bayes error for any distribution $p_{Z|G}$. However, we have $\frac{p_{G|Z}(-1|\mathbf{z}_1)}{p_G(-1)} = \frac{1}{p_G(-1) + d(1 - 2p_G(-1))} \rightarrow 0$ as $d \rightarrow \infty$, which invalidates (1) for any $\epsilon \in [0, \infty)$. This shows that a large Bayesian error probability does not guarantee information privacy of G .

III. PRIVACY MAPPING DESIGN

In this section, we formulate a nonparametric privacy-aware optimization framework to find privacy mappings or local decision rules for the sensors in order to achieve information privacy for the private hypothesis G . We show that our approach achieves a weak form of information privacy, and we propose iterative algorithms to solve the optimization problem.

A. Nonparametric Privacy-Aware Optimization

From Proposition 1(ii), we can solve the following optimization problem to obtain sensor privacy mappings $\{\gamma^t : t = 1, \dots, s\}$ in order to achieve a desired level of information privacy for G :

$$\begin{aligned} \min_{\gamma^1, \dots, \gamma^s, \gamma_H} \quad & \mathbb{P}(\gamma_H(Z) \neq H) \\ \text{s.t.} \quad & \min_{\gamma_G} R(\gamma_G) \geq \theta, \end{aligned} \quad (4)$$

where θ is a privacy threshold that determines the information privacy budget for G . However, since (4) requires knowledge of the joint distribution of sensor observations and the hypotheses H and G , such an approach may not be practical in IoT applications. Instead, we propose a nonparametric approach in which we use empirical regularized risks to approximate the Bayesian error probability for H and the risk $R(\gamma_G)$ for G .

Let $\mathcal{H}$ be a reproducing kernel Hilbert space (RKHS) associated with a kernel $\kappa(\cdot, \cdot)$. Let $\langle \cdot, \cdot \rangle$ be the kernel inner product associated with $\kappa(\cdot, \cdot)$, and $\|\cdot\|$ denote the norm induced by the kernel inner product. Following [33], we restrict γ_H and γ_G to be from $\mathcal{H}$, i.e., γ_H and γ_G are of the form

$$\begin{aligned} \gamma_H(\mathbf{z}) &= \langle w_H(\cdot), \Phi(\mathbf{z})(\cdot) \rangle, \\ \gamma_G(\mathbf{z}) &= \langle w_G(\cdot), \Phi(\mathbf{z})(\cdot) \rangle, \end{aligned}$$

where $w_H(\cdot), w_G(\cdot) \in \mathcal{H}$, and $\Phi(\mathbf{z})(\cdot) = \kappa(\cdot, \mathbf{z})$ is the feature map, which maps a point $\mathbf{z}$ to a function $\Phi(\mathbf{z})(\cdot) : \mathcal{Z}^s \mapsto \mathbb{R}$. To simplify notations, we use w_H, w_G , and $\Phi(\mathbf{z})$ respectively in this paper.

For each $t = 1, \dots, s$, $x^t \in \mathcal{X}$, and $z^t \in \mathcal{Z}$, let $Q^t(z^t | x^t) = \mathbb{P}(Z^t = z^t | X^t = x^t)$ represent the privacy mapping γ^t . Let $Q(\mathbf{z} | \mathbf{x}) = \mathbb{P}(Z = \mathbf{z} | X = \mathbf{x})$ denote the

overall sensor mapping, and $\mathcal{Q}$ be the set of $Q(\cdot | \cdot)$ such that

$$Q(\mathbf{z} | \mathbf{x}) = \prod_{t=1}^s Q^t(z^t | x^t), \quad (5)$$

$$\sum_{z^t \in \mathcal{Z}} Q^t(z^t | x^t) = 1, \quad (6)$$

$$Q^t(z^t | x^t) \geq 0. \quad (7)$$

Let ϕ be a loss function. We seek to minimize the regularized empirical ϕ -risk of deciding H , while ensuring that the *empirical normalized ϕ -risk* of deciding G , is higher than a given threshold θ . We assume that we are given n i.i.d. training data points $(x_i, h_i, g_i)_{i=1}^n$, and we seek to

$$\min_{w_H \in \mathcal{H}, Q \in \mathcal{Q}} \frac{1}{n} \sum_{i=1}^n \sum_{\mathbf{z} \in \mathcal{Z}^s} \phi(h_i \langle w_H, \Phi(\mathbf{z}) \rangle) Q(\mathbf{z} | \mathbf{x}_i) + \frac{\lambda}{2} \|w_H\|^2, \quad (8a)$$

$$\text{s.t. } \min_{w_G \in \mathcal{H}} \frac{1}{2} \sum_{g \in \{-1,1\}} \sum_{i \in \mathcal{S}_{g,n}} \sum_{\mathbf{z} \in \mathcal{Z}^s} \frac{\phi(g_i \langle w_G, \Phi(\mathbf{z}) \rangle) Q(\mathbf{z} | \mathbf{x}_i)}{|\mathcal{S}_{g,n}|} + \frac{\lambda_n}{2} \|w_G\|^2 \geq \theta. \quad (8b)$$

Here, $\mathcal{S}_{g,n} = \{i \in \{1, \dots, n\} : g_i = g\}$, for $g \in \{-1, 1\}$, and $\lambda > 0$ and $\lambda_n > 0$ are regularization weights. The regularized empirical risks are used, since without the regularization, it is known that the generalization error can become large [43]. Note also that the empirical normalized risk in (8b) is different from the traditional empirical risk used in the machine learning literature (cf. (8a) and [43]), and serves as a surrogate for $R(\cdot)$ in Proposition 1.

The optimization problem (8) is intractable to compute because of the summation over $\mathcal{Z}^s$, therefore we employ the same lower bound relaxation as in (15) of [33], and let $\Phi_Q(\mathbf{x}) = \sum_{\mathbf{z} \in \mathcal{Z}^s} Q(\mathbf{z} | \mathbf{x}) \Phi(\mathbf{z})$. The corresponding kernel is $\kappa_Q(\mathbf{x}, \mathbf{x}') = \langle \Phi_Q(\mathbf{x}), \Phi_Q(\mathbf{x}') \rangle = \sum_{\mathbf{z}} \sum_{\mathbf{z}'} Q(\mathbf{z} | \mathbf{x}) Q(\mathbf{z}' | \mathbf{x}') \kappa(\mathbf{z}, \mathbf{z}')$. With this relaxation, the convex surrogate of the risk $R(\gamma_G)$ becomes

$$R_\phi(w_G, Q) = \frac{1}{2} \mathbb{E}_X [\phi(-\langle w_G, \Phi_Q(X) \rangle) | G = -1] + \frac{1}{2} \mathbb{E}_X [\phi(\langle w_G, \Phi_Q(X) \rangle) | G = 1], \quad (9)$$

and the surrogates for the empirical risks for H and G are respectively,

$$F(w_H, Q) = \frac{1}{n} \sum_{i=1}^n \phi(h_i \langle w_H, \Phi_Q(\mathbf{x}_i) \rangle) + \frac{\lambda}{2} \|w_H\|^2, \quad (10)$$

$$\hat{R}_n(w_G, Q) = \frac{1}{2} \sum_{g \in \{-1,1\}} \sum_{i \in \mathcal{S}_{g,n}} \frac{\phi(g_i \langle w_G, \Phi_Q(\mathbf{x}_i) \rangle)}{|\mathcal{S}_{g,n}|} + \frac{\lambda_n}{2} \|w_G\|^2. \quad (11)$$

Our optimization problem (8) is then relaxed to

$$\min_{w_H \in \mathcal{H}, Q \in \mathcal{Q}} F(w_H, Q) \quad (12a)$$

$$\text{s.t. } \min_{w_G \in \mathcal{H}} \hat{R}_n(w_G, Q) \geq \theta. \quad (12b)$$

For convenience, we call (12) the nonparametric privacy-aware optimization (NPO) framework. We note that (8b) holds if the constraint (12b) holds, since (12b) is a lower bound relaxation. In general, (12b) cannot guarantee information privacy with probability one since (12) is solved using a finite number of training data points $(x_i, h_i, g_i)_{i=1}^n$. Therefore, we introduce the notion of *weak information privacy* below.

Definition 2 ((ϵ, δ) -Information privacy): Let G and Z be random variables. For $\epsilon > 0$, $\delta \in [0, 1]$, G given Z or $p_{G|Z}$ has (ϵ, δ) -information privacy, if

$$\mathbb{P} \left(e^{-\epsilon} \leq \frac{p_{G|Z}(G | Z)}{p_G(G)} \leq e^\epsilon \right) \geq 1 - \delta.$$

Our goal is to provide a theoretical guarantee for weak information privacy with some $\delta \in (0, 1)$ using (12b). However, since we are using a loss function ϕ as a surrogate for the 0–1 loss implied in $R(\cdot)$, we require the following assumptions, which have been used to bound generalization errors in [44].

Assumption 2 (Convex loss function): The loss function ϕ is a convex function and $\phi(0) < \infty$.

Assumption 3: For any $\eta \in [0, 1]$, let $R_\phi^*(\eta) = \inf_{\gamma \in \mathbb{R}} (\eta \phi(\gamma) + (1 - \eta) \phi(-\gamma))$. There exist $a > 0$ and $r \geq 1$, such that for all $\eta \in [0, 1]$,

$$a^r (\phi(0) - R_\phi^*(\eta)) \geq \left| \frac{1}{2} - \eta \right|^r. \quad (13)$$

Examples of loss functions, with their corresponding a and r values in Assumption 3, are shown in Table II. These include commonly used loss functions like the logistic loss function [45], exponential loss function [46], hinge loss function [47], and quadratic loss function, which have been shown to be computationally efficient and have bounded approximation and estimation errors [44], [48]. In Table II we also list the conjugate dual ϕ^* of these loss functions, which will be useful in the sequel.

For $\mathbf{z} \in \mathcal{Z}^s$, let $\eta(\mathbf{z}) = \mathbb{P}(\tilde{G} = 1 | Z = \mathbf{z})$, where $\tilde{G}$ induces the same conditional distribution for Z as G but has uniform prior. It can be shown that $\min_\gamma R(\gamma) = 1/2 - \mathbb{E}[|1/2 - \eta(Z)|]$ (cf. (37)). Therefore, Assumption 3 gives a lower bound for $\min_\gamma R(\gamma)$ in terms of $\mathbb{E}[R_\phi^*(\eta(Z))]$. To relate this to the left hand side of (12b), we need the following assumption.

Assumption 4:

i) $\lambda_n \rightarrow 0$, as $n \rightarrow \infty$.

ii) The kernel $\kappa(\cdot, \cdot)$ is a strictly positive kernel.

Examples of strictly positive kernels include the Gaussian kernel, and kernels defined by inverse multiquadrics [49]. With the above assumptions, the following result gives a theoretical guarantee of weak information privacy using the constraint (12b).

Theorem 1: Suppose that Assumptions 1–4 hold. Then for any $\delta \in (0, 1]$, there exists n_0 , such that for all $n \geq n_0$, G

TABLE II
LOSS FUNCTIONS, THEIR CONJUGATE DUALS, AND PARAMETERS FOR ASSUMPTION 3

Loss function ϕ	Conjugate dual $\phi^*(-x^*)$	Assumption 3
$\phi(u) = e^{-u}$	$\begin{cases} x^* \ln x^* - x^*, & x^* \in (0, \infty) \\ 0, & x^* = 0 \end{cases}$	$a = \frac{1}{\sqrt{2}}, r = 2$
$\phi(u) = \log(1 + e^{-u})$	$\begin{cases} x^* \log x^* + (1 - x^*) \log(1 - x^*), & x^* \in (0, 1) \\ 0, & x^* = 0, 1 \end{cases}$	$a = \frac{1}{\sqrt{2}}, r = 2$
$\phi(u) = \max\{1 - u, 0\}$	$-x^*, x^* \in [0, 1]$	$a = 1/2, r = 1$
$\phi(u) = (1 - u)^2$	$(x^*)^2/4 - x^*, x^* \in \mathbb{R}$	$a = 1/2, r = 2$

achieves (ϵ, δ) -information privacy for any $Q \in \mathcal{Q}$, if

$$\min_{w_G \in \mathcal{H}} \hat{R}_n(\hat{w}_G, Q) \geq \theta,$$

where $\epsilon = \log \frac{c}{(c - 2a(\phi(0) - \theta + \delta)^{1/r})_+}$, with c as defined in (2).

Proof: See Appendix B. ■

From Theorem 1, we see that our NPO (12) can achieve arbitrarily strong information privacy with $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$ by taking the sample size $n \rightarrow \infty$ and $\theta \rightarrow \phi(0)$. The latter condition will however result in a poor detection rate for H , since it leads to the case where all sensor observations are mapped to the same output. Therefore, in the following, we propose an iterative procedure to find a suitable threshold θ with appropriate constraints on the privacy mappings.

B. Iterative Optimization

Let $\beta = [\beta_1, \beta_2, \dots, \beta_n]^T \in \mathbb{R}^n$, and

$$\begin{aligned} \hat{R}^*(\beta, Q) = & - \sum_{g \in \{-1, 1\}} \sum_{i \in \mathcal{S}_{g,n}} \frac{\phi^*(-2|\mathcal{S}_{g,n}|\beta_i)}{2|\mathcal{S}_{g,n}|} \\ & - \frac{1}{2\lambda_n} \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j g_i g_j \kappa_Q(\mathbf{x}_i, \mathbf{x}_j), \end{aligned} \quad (14)$$

where ϕ^* is the conjugate dual of ϕ (see Table II for some examples) [50]. From Proposition 2 in [33], for any Q , we have

$$\min_{w_G \in \mathcal{H}} \hat{R}_n(w_G, Q) = \max_{\beta \in \mathbb{R}^n} \hat{R}^*(\beta, Q). \quad (15)$$

The optimal primal variable w_G and the optimal dual variable β in (15) are related by $w_G = \sum_{i=1}^n \beta_i g_i \Phi_Q(\mathbf{x}_i)$. The constraint (12b) can now be rewritten as $\max_{\beta \in \mathbb{R}^n} \hat{R}^*(\beta, Q) \geq \theta$. By using the interior-point method with log barrier [51], we transform the optimization problem (12) into the following:

$$\min_{w_H \in \mathcal{H}, \beta \in \mathbb{R}^n, Q \in \mathcal{Q}} F_0(w_H, \beta, Q), \quad (16)$$

where $F_0(w_H, \beta, Q) = F(w_H, Q) - \frac{1}{\mu} \log(\hat{R}^*(\beta, Q) - \theta)$ and $\mu > 0$ is the barrier parameter.

We propose a two-step algorithm to solve (16). Since its unclear how to choose the threshold θ a priori, in the first step, we use an iterative algorithm to find θ^* , which is the maximal possible θ , subject to some constraints on Q . The first step also provides the initial point for the iterative algorithm in the second step. In the second step, we set the threshold θ to be a fixed fraction p of θ^* , and iteratively solve (16). We call p the privacy threshold ratio. In both steps, we apply a *block Gauss-Seidel method* [52]. We describe our algorithm in detail as follows.

Algorithm 1: Finding θ^* .

- 1: **input:** $\{g_i, x_i^1, \dots, x_i^s\}_{i=1}^n$
- 2: **initialization:** $\beta[0] \in \mathbb{R}^n, Q[0] \in \mathcal{Q}, k = 0$,
- 3: **repeat**
- 4: • $k = k + 1$,
- fix $Q[k - 1]$, solve the following convex optimization problem,
$$\beta[k] = \arg \max_{\beta \in \mathbb{R}^n} \hat{R}^*(\beta, Q[k - 1]),$$
- fix $\beta[k]$, $Q^j[k]$, for $j < t$, and $Q^j[k - 1]$, for $j > t$, $t = 1, 2, \dots, s$ **update**

$$Q^t[k] = \arg \max_{Q^t \in \mathcal{Q}} \hat{R}^*(\beta[k], Q^1[k], \dots, Q^{t-1}[k],$$

$$Q^t, Q^{t+1}[k - 1], \dots, Q^s[k - 1]),$$
- 5: **until** $\frac{\hat{R}^*(\beta[k], Q[k]) - \hat{R}^*(\beta[k - 1], Q[k - 1])}{\hat{R}^*(\beta[k - 1], Q[k - 1])} \leq \epsilon$,
- 6: **return** $\tilde{\beta} = \beta[k], \tilde{Q} = Q[k], \theta^* = \hat{R}^*(\beta[k], Q[k])$.

1) *Finding the Privacy Threshold:* In Algorithm 1, we apply the block Gauss-Seidel method to solve

$$\max_{\beta \in \mathbb{R}^n, Q \in \mathcal{Q}'} \hat{R}^*(\beta, Q), \quad (17)$$

where $\mathcal{Q}'$ is the set of $Q \in \mathcal{Q}$, such that for all $t = 1, \dots, s$,

$$\sum_{x \in \mathcal{X}} Q^t(z | x) \geq \Delta_1, \text{ for all } z \in \mathcal{Z}, \quad (18)$$

$$\left| Q^t(z | x) - \frac{1}{|\mathcal{Z}|} \right| \geq \Delta_2, \text{ for all } z \in \mathcal{Z} \text{ and } x \in \mathcal{X}, \quad (19)$$

with Δ_1 and Δ_2 being small positive constants. The constraint (18) ensures that no $z \in \mathcal{Z}$ has small probability (otherwise we could have reduced the size of $\mathcal{Z}$), and (19) prevents equal probabilities from being assigned to all possible z . Since the optimization problem in (17) is non-convex, there is no guarantee that the block Gauss-Seidel method converges to the global optimum [52]–[54]. However, we can show the following convergence result.

Proposition 2: Algorithm 1 converges to a critical point.

Proof: See Appendix C. ■

The objective value output from Algorithm 1 is denoted as θ^* , and we use its solution $(\tilde{\beta}, \tilde{Q})$ as the initialization point in the second step.

2) *Finding Sensor Privacy Mappings:* With θ^* from Algorithm 1, we set $\theta = p\theta^*$, where $p \in (0, 1)$ is a constant

Algorithm 2: Optimizing sensor decision rules and fusion center rules.

1: **input:** $\{h_i, g_i, x_i^1, \dots, x_i^s\}_{i=1}^n$
2: **initialization:** $w[0] \in \mathbb{R}^n, \beta[0] \leftarrow \tilde{\beta}, Q[0] \leftarrow \tilde{Q}, k = 0$,
3: **repeat**
4: • $k = k + 1$,
• Fix $\beta[k-1]$ and $Q[k-1]$, and solve the following convex optimization problem,

$$\alpha[k] = \arg \max_{\alpha \in \mathbb{R}^n} \{F^*(\alpha, Q[k-1])\},$$
and we obtain $w_H[k] = \sum_{i=1}^n \alpha_i[k] h_i \Phi_Q(\mathbf{x}_i)$,
• Fix $w_H[k]$ and $Q[k-1]$, update

$$\beta[k] = \arg \max_{\beta \in \mathbb{R}^n} \frac{1}{\mu} \log \left(\hat{R}^*(\beta, Q[k-1]) - \theta \right),$$
• Fix $w_H[k]$, $\beta[k]$, and $Q^1[k], \dots, Q^{t-1}[k], Q^{t+1}[k-1], Q^s[k-1]$, update

$$Q^t[k] = \arg \min_{Q^t \in \mathcal{Q}} F_0(w_H[k], Q^1[k], \dots, Q^{t-1}[k], Q^t, Q^{t+1}[k-1], Q^s[k-1], \beta[k]),$$
5: **until** $\frac{F_0(w_H[k-1], Q[k-1], \beta[k-1]) - F_0(w_H[k], Q[k], \beta[k])}{F_0(w_H[k-1], Q[k-1], \beta[k-1])} \leq \epsilon$,
6: **output:** $\alpha[k], \beta[k], Q[k]$.

typically chosen to be close to 1 (see Section IV-A). We minimize F_0 over (w_H, β, Q) using a *block Gauss-Seidel method*, as shown in Algorithm 2.

Let $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_n]^T \in \mathbb{R}^n$, and

$$F^*(\alpha, Q) = -\frac{1}{n} \sum_{i=1}^n \phi^*(-n\alpha_i) - \frac{1}{2\lambda} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j h_i h_j \kappa_Q(\mathbf{x}_i, \mathbf{x}_j).$$

We have

$$\min_{w_H \in \mathcal{H}} F(w_H, Q) = \max_{\alpha \in \mathbb{R}^n} F^*(\alpha, Q), \quad (20)$$

where the optimal primal variable w_H and the optimal dual variable α in (20) are related by

$$w_H = \sum_{i=1}^n \alpha_i h_i \Phi_Q(\mathbf{x}_i). \quad (21)$$

Instead of recording the optimal w_H , which is associated with the feature map, we record the value of α , which is a vector of length n . Note that if we substitute the updated w_H with α according to (21), we turn the feature maps into the kernel in the primal space. Therefore, when it comes to updating β and Q^t , we do the minimization in the primal space.

Proposition 3: Algorithm 2 converges to a critical point.

Proof: See Appendix D. ■

C. Extension to m -ary Private Hypothesis

In this section, we extend our NPO framework to the case where G is an m -ary hypothesis, with $m > 2$. For simplicity, the public hypothesis H remains as a binary hypothesis since in many IoT applications like intrusion detection, detection of whether an event has occurred is the main phenomenon of interest. By using standard multi-class classification techniques, our framework can be generalized to the case where H is a multi hypothesis. We refer the reader to [55] for details.

Suppose that $G \in \mathcal{G} = \{0, 1, \dots, m-1\}$, where $m > 2$. For a detector γ_g that distinguishes between the hypothesis pair $(0, g)$, where $1 \leq g < m$, let

$$R_g(\gamma_g) = \frac{1}{2} (\mathbb{P}(\gamma_g(Z) = g | G = 0) + \mathbb{P}(\gamma_g(Z) = 0 | G = g)),$$

and

$$\ell_g(\mathbf{z}) = \frac{p_{Z|G}(\mathbf{z}|g)}{p_{Z|G}(\mathbf{z}|0)}, \quad (22)$$

for $\mathbf{z} \in \mathcal{Z}^s$. The following result follows from Proposition 1(ii).

Theorem 2: Suppose that $\min_{g \in \mathcal{G}} p_G(g) > 0$ and the support $\mathcal{D}$ of the conditional distributions of Z given $G = g$ for all $g \in \mathcal{G}$ are the same. If $\min_{1 \leq g < m, \gamma_g} R_g(\gamma_g) \geq \theta$, with $\theta \in [0, 1/2]$, then G given Z achieves ϵ -information privacy where $\epsilon = 2 \log \frac{c'}{(c' + 2\theta - 1)_+}$, with $c' = \min_{1 \leq g < m} \{\mathbb{P}(\ell_g(\mathbf{z}) = \min_{\mathbf{z} \in \mathcal{D}} \ell_g(\mathbf{z}) | G = 0), \mathbb{P}(\ell_g(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell_g(\mathbf{z}) | G = g)\} > 0$.

Proof: See Appendix E.

To achieve (ϵ, δ) -information privacy for G for $\delta \in (0, 1]$, we impose the following $m-1$ empirical normalized risk constraints:

$$\min_{w_g \in \mathcal{H}} \hat{R}_{g,n}(w_g, Q) \geq \theta, \text{ for } g = 1, \dots, m-1, \quad (23)$$

where

$$\hat{R}_{g,n}(w, Q) = \frac{1}{2} \sum_{g' \in \{0, g\}} \sum_{i \in \mathcal{S}_{g',n}} \frac{\phi(g'_i \langle w, \Phi_Q(\mathbf{x}_i) \rangle)}{|\mathcal{S}_{g',n}|} + \frac{\lambda_n}{2} \|w\|^2, \quad (24)$$

with $g'_i = -1$ if $g_i = 0$ and $g'_i = 1$ otherwise.

Note that here, we only need to consider the empirical normalized risk of confusing $G = g \geq 1$ with $G = 0$, instead of every pair of hypothesis values. The intuition is the same as in m -ary hypothesis testing, where (22) are sufficient statistics.

Indeed, it can be shown, using the same arguments as in Theorem 1 that the following holds (we omit the proof here):

Theorem 3: Suppose the corresponding assumptions as in Theorem 1 for m -ary hypothesis G hold. Then for any $\delta \in (0, 1]$, there exists n_0 , such that for all $n \geq n_0$, G achieves (ϵ, δ) -information privacy for any $Q \in \mathcal{Q}$ satisfying (23), where $\epsilon = \frac{1}{2} \log \frac{c'}{(c' - 2a(\phi(0) - \theta + \delta)^{1/r})_+}$, with c' as defined in Theorem 2.

We apply Algorithm 1 to each of the $m-1$ constraints in (23) to obtain a threshold θ_g^* for each $g = 1, \dots, m-1$. Then, we set $\theta = p \cdot \min_{1 \leq g < m} \theta_g^*$, for some $p \in (0, 1)$. Finally, in

Algorithm 2, the objective function is replaced by

$$\min_{w_H \in \mathcal{H}, \beta_g, Q \in \mathcal{Q}} F(w_H, Q) - \frac{1}{\mu} \sum_{g=1}^{m-1} \log(\hat{R}_g^*(\beta_g, Q) - \theta),$$

where $\hat{R}_g^*(\beta_g, Q)$ is the dual form of $\hat{R}_{g,n}(w_g, Q)$ (see (14)). It can be shown that Proposition 3 still holds.

IV. SIMULATIONS AND EXPERIMENTS

In this section, we first perform simulations to provide insights into how different parameters impact the performance of our NPO approach. We then test our algorithm on real datasets from the UCI Repository [56], and compare its performance with RUCA [32] and MDR [31].

For simplicity, we use the *count kernel* in our simulations, which is defined as $\kappa(\mathbf{z}, \tilde{\mathbf{z}}) = \sum_{t=1}^s \mathbf{1}_{\{z^t = \tilde{z}^t\}}$. Then, for any $Q \in \mathcal{Q}$, we have

$$\kappa_Q(\mathbf{x}, \tilde{\mathbf{x}}) = \sum_{t=1}^s \sum_{z \in \mathcal{Z}} Q^t(z | x^t) Q^t(z | \tilde{x}^t),$$

which can be computed with a time complexity of $\mathcal{O}(s|\mathcal{Z}|)$. On the other hand, for some kernels like the Gaussian kernel, the same computation incurs a time complexity of $\mathcal{O}(|\mathcal{Z}|^{2s})$. Note however that the count kernel does not satisfy Assumption 4(ii). Nevertheless, our simulations suggest that using the count kernel does not prevent our NPO approach from protecting the information privacy of G . Since many IoT devices and gateways that serve as fusion centers are embedded platforms with limited computation power [57], the count kernel allows practical implementation. Therefore, in our simulations, we evaluate the performance of our algorithms using the count kernel.

We choose the logistic loss function as ϕ in our simulations, and employ gradient descent in the optimization steps in Algorithms 1 and 2.

The complexity of both Algorithms 1 and 2 using gradient descent in each optimization step and the count kernel is $\mathcal{O}(n^2 s |\mathcal{Z}|)$. We use $\Delta_1 = \Delta_2 = 0.005$ in Algorithm 1 throughout.

A. Synthetic Data Set

In this section, we generate a synthetic data set to verify the performance of our proposed method. We first consider the case where G has uniform prior probability and then give an example where G has a skewed prior. Finally, we present simulation results for the case where G is a m -ary hypothesis, with $m > 2$. To evaluate the performance of our algorithm, we compute the Bayes errors for detecting H and G since these are the minimum detection errors any detector can achieve so that our results are oblivious to the choice of learning method adopted by the fusion center. We also show the Bayes errors of detecting H and G when $Z = X$ (i.e., the raw sensor observations are available at the fusion center) as a baseline for comparison.

1) *Performance of NPO Framework*: Consider a network of 4 sensors and a fusion center. Each sensor observation $x_i^t \in \mathcal{X} = \{1, 2, \dots, 9\}$ is generated according to Table III, where n_i^t is distributed uniformly over $\{-1, 0, +1\}$. In this set

TABLE III
SENSOR OBSERVATIONS X FOR DIFFERENT REALIZATIONS OF (H, G)

(h_i, g_i)	x_i^t
$(-1, -1)$	$2 + n_i^t$
$(-1, 1)$	$4 + n_i^t$
$(1, -1)$	$6 + n_i^t$
$(1, 1)$	$8 + n_i^t$

Fig. 2. Error rate of detecting H and G with different correlation coefficient between H and G .

of simulations, G has uniform prior. Conditioned on (H, G) , sensor observations are independent of each other. We generate 80 i.i.d. training samples and 1000 i.i.d. testing samples.

In Fig. 2, we show how the correlation coefficient between H and G affects the detection error rates. Recall that the privacy threshold is chosen to be $\theta = p\theta^*$, where θ^* is found using Algorithm 1. We use a privacy threshold ratio $p = 0.999$. We generated both H and G with zero mean, but varying correlation coefficient between them. As the correlation between H and G becomes larger, the Bayes error for both hypotheses converge to each other. As expected, to achieve a reasonable Bayes error for H requires that H is not too correlated with G .

We compare the performance of our approach with the optimal detectors found using (4), which assumes knowledge of the underlying joint distribution. We use the same θ as that in our NPO to achieve the same error rate for G . We see that our NPO achieves an error rate for H that is not too different from the optimal detectors, even though no prior knowledge of the underlying distribution is assumed in NPO.

In Fig. 3, we show the effect of the threshold ratio p on the error rates of detecting H and G , when their correlation coefficients are 0.2 and 0.8, respectively. The testing error is the empirical detection rate of the classifier found using Algorithm 2 applied to the testing samples we generated.

We see that when the correlation between H and G is small, p has no significant effect on their error rates over a large range. This is because the privacy mapping Q that minimizes the error rate of H does not contain much information about G . We also observe that with the NPO framework, the error rate for H is

Fig. 3. Error rate of detecting H and G with varying privacy threshold ratio p .

not significantly higher than using the raw sensor observations, whereas the error rate for G is increased significantly. However when the correlation between H and G is large, p has significant impact on the error rates. In this case, as expected, we cannot find a p that induces a low error rate for H and a high error rate for G .

2) *Comparison With the Empirical Bayes Error Privacy Metric:* Example 1 shows that even if the Bayes error is large, information privacy cannot be guaranteed. In this simulation, we demonstrate that using the empirical risk corresponding to the Bayes error [35] as a privacy metric can lead to privacy mappings that result in a $p_{G|Z}(g | \mathbf{z})$ that is significantly different from $p_G(g)$, i.e., some realization $Z = \mathbf{z}$ can leak significant information about the true state of G . With the empirical risk corresponding to the Bayes error, our optimization problem becomes:

$$\begin{aligned} & \min_{w_H \in \mathcal{H}, Q \in \mathcal{Q}} F(w_H, Q) \\ \text{s.t. } & \min_{w_G \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \phi(g_i \langle w_G, \Phi_Q(\mathbf{x}_i) \rangle) + \frac{\lambda}{2} \|w_G\|^2 \geq \theta. \end{aligned} \quad (25)$$

We let $\mathbb{P}(G = 1) = 0.95$ and the correlation coefficient between H and G to be 0.23. We use a similar procedure as Algorithms 1 and 2 to find a privacy threshold and the sensor privacy mappings for (25). The results are shown in Table IV. We see that the Bayes errors of detecting H and G respectively are similar for both the NPO approach and (25). However, NPO achieves a much smaller privacy budget ϵ (i.e., stronger information privacy) than (25).

3) *NPO With m -ary Private Hypothesis:* We next show the performance of NPO when G is a m -ary hypothesis with $m > 2$. We let $(H, G) \in \{-1, 1\} \times \{0, \dots, m-1\}$, H and G be uncorrelated, with both having uniform priors. We let $X^t = m(H + 1) + 2(G + 1) + N^t$, where $t = 1, 2, 3, 4$, and for sensors $t = 5, \dots, 4 + m$, we let

$$X^t = \begin{cases} H + N^t, & \text{if } G - H/2 - 0.5 \pmod{m} \equiv t - 5 \\ 0, & \text{otherwise.} \end{cases}$$

The noise N^t is chosen uniformly at random from $\{-1, 0, +1\}$. We perform training using 80 training samples, with $\mathcal{Z} = \{1, 2\}$.

We compare NPO with two other methods: (i) detecting H and G when $Z = X$; and (ii) finding the sensor local decision rules by minimizing $F(w_H, Q)$ without any privacy constraints. We call this the nonparametric decentralized detection (NDD) method [33].

We see from Table V that with NPO, the Bayes error of detecting H is small, while the Bayes error of detecting G is large, and increasing with m . Note that the maximum possible Bayes error for G (achieved by random guessing, i.e., Z is independent of G) is $1 - 1/m$. We see that the methods we compare with have significantly lower error rates for detecting G when compared to NPO. However, NPO has a slightly bigger Bayes error for detecting H when compared with NDD.

B. UCI Repository Data Sets

We test our algorithm on the OPPORTUNITY Activity Recognition Data Set [58] and Census (Adult) Data Set (both available at UCI Repository[56]), and compare its performance with RUCA [32] and MDR [31]. In our NPO, we set the local decision space of each sensor to be $\mathcal{Z} = \{1, 2\}$.

1) *Data and Implementation:* The OPPORTUNITY Activity Recognition Data Set consists of recorded readings of on-body, object and ambient sensors installed in a living space, while a person executes typical daily activities. We let the public hypothesis be locomotion detection (standing or walking), and the private hypothesis be the object the person is touching (a drawer or a dish washer). We used the data from ‘S2-Drill’, and Weka [59] to select $s = 15$ sensors that are most correlated with our chosen labels. Since the sensor reading is continuous, unsupervised discretization was applied to quantize each continuous sensor reading to 10 levels. We randomly sampled $n = 40$ instances of training data, and 3427 instances of testing data.

The Census (Adult) Data Set contains information about a population, such as income, marital status, gender, and nationality. We let the public hypothesis be whether a person has high income ($> 50K$) or low income ($\leq 50K$), and the private hypothesis be the marital status of the person. Following [32], we group ‘Married-civ-spouse’, ‘Married-spouse-absent’ and ‘Married-AFSpouse’ into a single category called ‘Married’, while ‘Divorced’, ‘Separated’ and ‘Widowed’ are grouped into a single category called ‘Used to be Married’. The ‘Never Married’ category was left as it is. Thus the private hypothesis is a 3-ary hypothesis. We used data from ‘adult.data’, and selected $s = 5$ attributes (age, education, occupation, relationship, sex) that are most correlated with our chosen labels. Unsupervised discretization is applied to quantize each numerical attributes to 15 levels. We then randomly sampled $n = 180$ instances of training data, and 2100 instances of testing data.

2) *Results and Discussion:* To the best of our knowledge, our work is the first to provide a nonparametric approach to design *decentralized* sensor privacy mappings. Therefore, we compare our method to RUCA [32] and MDR [31], which are nonparametric approaches that require a trusted data curator to

TABLE IV
COMPARISON OF NPO WITH (25)

	Bayes Error H	Bayes Error G	$R(\gamma_G)$	Information privacy budget, ϵ
NPO	0.03	0.05	0.38	1.08
(25)	0.02	0.05	0.24	3.91

TABLE V
BAYES ERROR WITH m -ARY PRIVATE HYPOTHESIS

m	Detect H with NPO	Detect G with NPO	Detect H with $Z = X$	Detect G with $Z = X$	Detect H with NDD	Detect G with NDD
3	3.62×10^{-2}	0.65	3.77×10^{-3}	1.21×10^{-2}	2.41×10^{-2}	0.13
4	2.72×10^{-2}	0.73	1.10×10^{-3}	3.92×10^{-3}	1.85×10^{-2}	0.19
5	2.13×10^{-2}	0.79	3.02×10^{-3}	4.22×10^{-3}	1.78×10^{-2}	0.23

TABLE VI
DETECTION ERRORS USING THE OPPORTUNITY ACTIVITY
RECOGNITION DATA SET

Detection Method	H Locomotion	G Object	$\hat{\epsilon}$
NPO	10.53%	43.17%	0.47
MDR	12.56%	40.16%	0.69
DCA	10.88%	42.62%	
RUCA ($\rho_p = 1$)	10.23%	45.73%	
RUCA ($\rho_p = 100$)	10.10%	43.01%	
RUCA ($\rho_p = 1000$)	10.10%	43.78%	
$Z = X$	10.05%	5.57%	

TABLE VII
DETECTION ERRORS USING THE CENSUS (ADULT) DATA SET

Detection Method	H Income	G Marital-status	$\hat{\epsilon}$
NPO	37.67%	62.14%	0.82
MDR	37.57%	64.02%	1.61
DCA	38.38%	56.33%	
RUCA ($\rho_p = 1$)	41.24%	65.25%	
RUCA ($\rho_p = 100$)	41.10%	64.14%	
RUCA ($\rho_p = 1000$)	40.67%	65.86%	
$Z = X$	34.05%	30.48%	

first aggregate data from all sensors before applying an overall privacy mapping. Parameter settings for RUCA and MDR are as proposed in [32] and [31], respectively. We also compare with the case where sensors send their observations directly to the fusion center, namely $Z = X$.

To estimate the information privacy budget achieved by each method, we compute

$$\hat{\epsilon} = \max_{g \in \mathcal{G}, \mathbf{z} \in \mathcal{Z}^s} \left| \log \frac{\hat{p}_{G,Z}(g, \mathbf{z})}{\hat{p}_G(g) \hat{p}_Z(\mathbf{z})} \right|,$$

where $\hat{p}_A(a)$ is the empirical probability of the event $\{A = a\}$. Note that a smaller $\hat{\epsilon}$ implies stronger information privacy.

From Table VI and Table VII, we observe that our proposed NPO is comparable to the other nonparametric methods that need a trusted data curator. We also observe that NPO achieves stronger information privacy than the other methods, since it explicitly optimizes for this criterion.

V. CONCLUSIONS

Information privacy is the protection against statistical inference of a private phenomenon. In this paper, we proved that for a binary hypothesis, ensuring that the average of its Type I and II errors is large, is sufficient to guarantee information privacy. We proposed a nonparametric privacy-aware optimization framework to find sensor privacy mappings that can limit the fusion center's detection rate of a private hypothesis, without significantly compromising the detecting ability of the public hypothesis. We proved a theoretical guarantee of weak information privacy for our proposed framework, and proposed iterative algorithms to solve the proposed optimization problem. Our approach is suitable for IoT networks in which every sensor operates independently of the others.

Future research directions include generalizing our privacy-aware optimization framework to networks with multiple layers so that nodes of different sensing or computation ability can apply local decision rules of different complexities to achieve a better overall utility-privacy tradeoff. It would also be of interest to design sensor privacy mappings that achieve both information and data privacy simultaneously.

APPENDIX A PROOF OF PROPOSITION 1

We first show (i). Let $P_e = \mathbb{P}(\gamma_G(Z) \neq G)$ and $\mathcal{I}(\cdot; \cdot)$ denote the mutual information operator. From Fano's inequality [42], we have

$$\begin{aligned} \mathcal{H}(P_e) &\geq \mathcal{H}(G | Z) \\ &= \mathcal{H}(G) - \mathcal{I}(G; Z) \\ &= \mathcal{H}(G) - \mathbb{E} \left[\log \frac{p_{G|Z}(G|Z)}{p_G(G)} \right] \\ &\geq \mathcal{H}(G) - \epsilon, \end{aligned}$$

where the last inequality follows from Definition 1. Since $\mathcal{H}(P_e)$ is a non-decreasing function of P_e for $0 \leq P_e \leq 1/2$, we obtain $P_e \geq \theta$, where θ satisfies (3). Part (i) now follows from $R(\gamma_G) \geq P_e / (2 \max_g p_G(g))$.

We next show (ii). Let $\gamma_G(Z) = \arg \min_{\gamma} R(\gamma)$. It can be shown that $\gamma_G(Z) = 1$ if $\ell(Z) \geq 1$ and $\gamma_G(Z) = -1$ otherwise.

Let $\Gamma = \{\mathbf{z} \in \mathcal{D} : \ell(\mathbf{z}) \geq 1\}$ and $\Gamma^c = \mathcal{D} \setminus \Gamma$. We have

$$R(\gamma_G) = \frac{1}{2} - \frac{1}{2} \sum_{\mathbf{z} \in \Gamma} (p_{Z|G}(\mathbf{z} | 1) - p_{Z|G}(\mathbf{z} | -1)) \quad (26)$$

$$= \frac{1}{2} + \frac{1}{2} \sum_{\mathbf{z} \in \Gamma^c} (p_{Z|G}(\mathbf{z} | 1) - p_{Z|G}(\mathbf{z} | -1)). \quad (27)$$

For any $\mathbf{z}' \in \Gamma$, if $R(\gamma_G) \geq \theta$, we have from (26),

$$\begin{aligned} 1 - 2\theta &\geq \sum_{\mathbf{z} \in \Gamma} (p_{Z|G}(\mathbf{z} | 1) - p_{Z|G}(\mathbf{z} | -1)) \\ &= \sum_{\mathbf{z} \in \Gamma} p_{Z|G}(\mathbf{z} | 1) \left(1 - \frac{1}{\ell(\mathbf{z})}\right) \\ &\geq \left(1 - \frac{1}{\ell(\mathbf{z}')} \right) \sum_{\{\mathbf{z} : \ell(\mathbf{z}) \geq \ell(\mathbf{z}')\}} p_{Z|G}(\mathbf{z} | 1) \\ &\geq \left(1 - \frac{1}{\ell(\mathbf{z}')} \right) \mathbb{P}(\ell(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) | G = 1) \\ &\geq c \left(1 - \frac{1}{\ell(\mathbf{z}')} \right). \end{aligned} \quad (28)$$

From (28), we obtain

$$\frac{1}{\ell(\mathbf{z}')} \geq 1 - \frac{1 - 2\theta}{c},$$

and

$$1 \leq \ell(\mathbf{z}') \leq \frac{1}{(1 - (1 - 2\theta)/c)_+} = e^\epsilon. \quad (29)$$

Similarly, for any $\mathbf{z}' \in \Gamma^c$, a similar argument using (27) yields

$$e^{-\epsilon} \leq \ell(\mathbf{z}') \leq 1. \quad (30)$$

Combining (29) and (30), we have for any $\mathbf{z} \in \mathcal{D}$,

$$e^{-\epsilon} \leq \ell(\mathbf{z}) \leq e^\epsilon, \quad (31)$$

and for any $g \in \{-1, 1\}$,

$$\begin{aligned} e^{-\epsilon} p_{Z|G}(\mathbf{z} | g) &\leq p_Z(\mathbf{z}) \\ &= \mathbb{E}_G[p_{Z|G}(\mathbf{z} | G)] \leq e^\epsilon p_{Z|G}(\mathbf{z} | g). \end{aligned} \quad (32)$$

We then obtain (1) by noting that

$$\frac{p_{G|Z}(g | \mathbf{z})}{p_G(g)} = \frac{p_{Z|G}(\mathbf{z} | g)}{p_Z(\mathbf{z})},$$

and the proof of the theorem is complete.

APPENDIX B PROOF OF THEOREM 1

We start with the following elementary lemma.

Lemma B.1: For any $w \in \mathcal{H}$,

$$\mathbb{E}[\hat{R}_n(w, Q)] = R_\phi(w, Q) + \frac{\lambda_n}{2} \|w\|^2,$$

where the expectation is taken with respect to the training samples $(X_i, G_i)_{i=1}^n$.

Proof: For each $g \in \{-1, 1\}$, we have

$$\begin{aligned} &\mathbb{E} \left[\sum_{i \in \mathcal{S}_{g,n}} \frac{\phi(G_i \langle w, \Phi_Q(X_i) \rangle)}{|\mathcal{S}_{g,n}|} \right] \\ &= \mathbb{E} \left[\mathbb{E}_X \left[\sum_{i \in \mathcal{S}_{g,n}} \frac{\phi(G_i \langle w, \Phi_Q(X_i) \rangle)}{|\mathcal{S}_{g,n}|} \middle| G_1, \dots, G_n \right] \right] \\ &= \mathbb{E} \left[\frac{1}{|\mathcal{S}_{g,n}|} \sum_{i \in \mathcal{S}_{g,n}} \mathbb{E}_X [\phi(G_i \langle w, \Phi_Q(X_i) \rangle) | G_i = g] \right] \\ &= \mathbb{E}_X [\phi(g \langle w, \Phi_Q(X) \rangle) | G = g]. \end{aligned}$$

The lemma now follows from (11) and (9). $\blacksquare$

We now prove Theorem 1. Let $\hat{w}_{G,n} = \arg \min_{w \in \mathcal{H}} \hat{R}_n(w, Q)$. From Markov's inequality, we have

$$\begin{aligned} &\mathbb{P} \left(\hat{R}_n(\hat{w}_{G,n}, Q) \geq \inf_{w \in \mathcal{H}} R_\phi(w, Q) + \delta \right) \\ &\leq \frac{1}{\delta} \left(\mathbb{E}[\hat{R}_n(\hat{w}_{G,n}, Q)] - \inf_{w \in \mathcal{H}} R_\phi(w, Q) \right) \\ &= \frac{1}{\delta} \left(\mathbb{E}[\min_{w \in \mathcal{H}} \hat{R}_n(w, Q)] - \inf_{w \in \mathcal{H}} R_\phi(w, Q) \right) \\ &\leq \frac{1}{\delta} \left(\inf_{w \in \mathcal{H}} \mathbb{E}[\hat{R}_n(w, Q)] - \inf_{w \in \mathcal{H}} R_\phi(w, Q) \right) \\ &= \frac{1}{\delta} \left(\inf_{w \in \mathcal{H}} \{R_\phi(w, Q) + \frac{\lambda_n}{2} \|w\|^2\} - \inf_{w \in \mathcal{H}} R_\phi(w, Q) \right), \end{aligned} \quad (33)$$

where the last equality follows from Lemma B.1. If $n \rightarrow 0$ in (33), its right hand side approaches to 0 uniformly in Q because of Assumption 4(i) and $\mathcal{Q}$ is a finite set. Therefore,

there exists n_0 , such that for all $n \geq n_0$, we have with probability $1 - \delta$,

$$\hat{R}_n(\hat{w}_{G,n}, Q) \leq \inf_{w \in \mathcal{H}} R_\phi(w, Q) + \delta. \quad (34)$$

Let $\tilde{G}$ be a binary hypothesis with uniform prior, and satisfies $p_{Z|\tilde{G}}(\mathbf{z}|g) = p_{Z|G}(\mathbf{z}|g)$ for all $\mathbf{z} \in \mathcal{Z}^s$ and $g \in \{-1, 1\}$. We then have

$$\begin{aligned} R_\phi(w, Q) &= \mathbb{E} \left[\phi \left(\tilde{G} \langle w, \Phi_Q(X) \rangle \right) \right] \\ &= \mathbb{E} \left[\phi \left(\tilde{G} \langle w, \mathbb{E}_Z [\Phi(Z) | X] \rangle \right) \right] \\ &\leq \mathbb{E} \left[\mathbb{E}_Z \left[\phi \left(\tilde{G} \langle w, \Phi(Z) \rangle \right) \middle| X \right] \right] \\ &= \mathbb{E} \left[\phi \left(\tilde{G} \cdot \gamma(Z) \right) \right] \\ &= \mathbb{E} [(\eta(Z)\phi(\gamma(Z)) + (1 - \eta(Z))\phi(-\gamma(Z)))], \end{aligned} \quad (35)$$

where the inequality follows from Jensen's inequality, and we let $\gamma(\mathbf{z}) = \langle w, \Phi(\mathbf{z}) \rangle$ in the penultimate equality. Recall that $\eta(\mathbf{z}) = \mathbb{P}(\tilde{G} = 1 | Z = \mathbf{z})$. For each $\mathbf{z} \in \mathcal{Z}^s$, let $\gamma^*(\mathbf{z}) = \arg \inf_{\gamma \in \mathbb{R}} (\eta(\mathbf{z})\phi(\gamma) + (1 - \eta(\mathbf{z}))\phi(-\gamma))$, and let K be the

Gram matrix of $\kappa(\cdot, \cdot)$ with respect to all elements in $\mathcal{Z}^s = \{\mathbf{z}_1, \dots, \mathbf{z}_{|\mathcal{Z}^s|}\}$. From Assumption 4(ii), since K is strictly positive definite, we can define

$$[\psi_1, \psi_2, \dots, \psi_{|\mathcal{Z}^s|}] = [\gamma^*(\mathbf{z}_1), \gamma^*(\mathbf{z}_2), \dots, \gamma^*(\mathbf{z}_{|\mathcal{Z}^s|})] \cdot K^{-1},$$

and take $\gamma(\mathbf{z}) = \sum_{i=1}^{|\mathcal{Z}^s|} \psi_i \kappa(\mathbf{z}_i, \mathbf{z}) = \gamma^*(\mathbf{z})$ for all $\mathbf{z} \in \mathcal{Z}^s$. Therefore, from (35), we obtain

$$\inf_{w \in \mathcal{H}} R_\phi(w, Q) \leq \mathbb{E}_Z[R_\phi^*(\eta(Z))]. \quad (36)$$

For $\eta \in [0, 1]$, let $R^*(\eta) = 1/2 - |1/2 - \eta|$. We then have

$$\begin{aligned} \min_{\gamma} R(\gamma) &= \min_{\gamma} \mathbb{E}[(1 - \eta(Z))\mathbf{1}_{\{\gamma(Z) \geq 0\}} + \eta(Z)\mathbf{1}_{\{\gamma(Z) < 0\}}] \\ &= \frac{1}{2} - \mathbb{E} \left[\left| \frac{1}{2} - \eta(Z) \right| \right] \\ &= \mathbb{E}[R^*(\eta(Z))]. \end{aligned} \quad (37)$$

From Assumption 3, we obtain

$$\begin{aligned} \mathbb{E}[R_\phi^*(\eta(Z))] &\leq \phi(0) - \frac{1}{a^r} \mathbb{E} \left[\left(\frac{1}{2} - R^*(\eta(Z)) \right)^r \right] \\ &\leq \phi(0) - \frac{1}{a^r} \left(\frac{1}{2} - \mathbb{E}[R^*(\eta(Z))] \right)^r, \end{aligned} \quad (38)$$

where the last inequality follows from Jensen's inequality. Substituting (37) into (38), we have

$$\begin{aligned} \min_{\gamma} R(\gamma) &\geq \frac{1}{2} - a(\phi(0) - \mathbb{E}[R_\phi^*(\eta(Z))])^{\frac{1}{r}} \\ &\geq \frac{1}{2} - a(\phi(0) - \theta + \delta)^{\frac{1}{r}}, \end{aligned}$$

where the last inequality follows from (34) and (36), and holds with probability $1 - \delta$. The theorem then follows from Proposition 1, and the proof is complete.

APPENDIX C PROOF OF PROPOSITION 2

We first show that $-\hat{R}^*(\beta, Q)$ is a strictly quasi-convex function of Q if $\beta \neq 0$. The definition of strictly quasi-convexity is borrowed from [52] as follows.

Definition C.1: A function $f(x_1, x_2, \dots, x_s)$ is called a strictly quasi-convex function of $x_1, \dots, x_s$, if for all $k \in \{1, \dots, s\}$,

$$\begin{aligned} f(x_1, \dots, px_k^1 + (1-p)x_k^2, \dots, x_s) \\ < \max(f(x_1, \dots, x_k^1, \dots, x_s), f(x_1, \dots, x_k^2, \dots, x_s)), \end{aligned}$$

for any $p \in (0, 1)$.

Lemma C.1: Suppose that $\beta \neq 0$. Then, $-\hat{R}^*(\beta, Q)$ is a strictly quasi-convex function of Q^t , $t = 1, 2, \dots, s$, for $Q \in \mathcal{Q}'$.

Proof: For each $Q \in \mathcal{Q}'$, consider

$$\begin{aligned} &\sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j g_i g_j \kappa_Q(\mathbf{x}_i, \mathbf{x}_j) \\ &= \sum_{i,k} \sum_{j,k'} Q(\mathbf{z}_k | \mathbf{x}_i) Q(\mathbf{z}_{k'} | \mathbf{x}_j) \beta_i \beta_j g_i g_j \kappa(\mathbf{z}_k, \mathbf{z}_{k'}) \quad (39) \\ &= \left\| \sum_{i,k} \beta_i g_i Q(\mathbf{z}_k | \mathbf{x}_i) \Phi(\mathbf{z}_k) \right\|^2 > 0, \end{aligned}$$

where the last inequality holds due to Assumption 4(ii) since $\beta \neq 0$, and $Q^t \neq 0$ for all $t = 1, \dots, s$ because of (18). From (14) and (39), $-\hat{R}^*(\beta, Q)$ is a positive definite quadratic form of Q^t , for each $t = 1, 2, \dots, s$, and is thus a strictly convex function of Q^t . This implies strict quasi-convexity, and the proof of the lemma is complete. ■

We now show that $\{\beta[k]\}$ and $\{Q[k]\}$ in Algorithm 1 have limit points. Since $\phi^*(\cdot) \geq -\phi(0) > -\infty$, $-\hat{R}^*(\beta, Q)$ is lower bounded. Therefore, the decreasing sequence $F[k] = -\hat{R}^*(\beta[k], Q[k])$ converges [60]. From Proposition 4 of [52], convergence of $\{F[k]\}$ implies convergence of $\{\beta[k]\}$ and $\{Q[k]\}$ to limit points. By Proposition 5 of [52], this shows that Algorithm 1 converges to a critical point. The proof of the proposition is now complete.

APPENDIX D PROOF OF PROPOSITION 3

From [33], $F(w_H, Q)$ is a convex function of Q^t , for each $t = 1, 2, \dots, s$. Since $-\log(\cdot)$ is a non-increasing convex function, and $\hat{R}^*(\beta, Q) - \theta$ is a strictly concave function of Q^t , $-\frac{1}{\mu} \log(\hat{R}^*(\beta, Q) - \theta)$ is strictly convex with respect to Q^t [51]. Therefore, $F_0(w_H, \beta, Q)$ is a strictly quasi-convex function of Q^t , $t = 1, \dots, s$ (see Definition C.1). The rest of the proof is similar to that in Appendix C.

APPENDIX E PROOF OF THEOREM 2

Following the same argument as that in the proof of Proposition 1(ii), for any $1 \leq g < m$ and $\mathbf{z} \in \mathcal{D}'$, we have from (31),

$$e^{-\epsilon/2} \leq \ell_g(\mathbf{z}) \leq e^{\epsilon/2},$$

and for any $g, g' \in \{0, \dots, m-1\}$,

$$e^{-\epsilon} \leq \frac{p_{Z|G}(\mathbf{z} | g)}{p_{Z|G}(\mathbf{z} | g')} \leq e^{\epsilon}.$$

The proof then proceeds similarly as that in Proposition 1(ii).

REFERENCES

- [1] I. Butun, S. D. Morgera, and R. Sankar, "A survey of intrusion detection systems in wireless sensor networks," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 1, pp. 266–282, Jan. 2014.
- [2] W.-P. Chen, J. C. Hou, and L. Sha, "Dynamic clustering for acoustic target tracking in wireless sensor networks," *IEEE Trans. Mobile Comput.*, vol. 3, no. 3, pp. 258–271, Jul. 2004.

- [3] I. Stoianov, L. Nachman, S. Madden, T. Tokmouline, and M. Csail, "Pipenet: A wireless sensor network for pipeline monitoring," in *Proc. Int. Symp. Inf. Process. Sensor Netw.*, 2007, pp. 264–273.
- [4] J. Chen, K. Kwong, D. Chang, J. Luk, and R. Bajcsy, "Wearable sensors for reliable fall detection," in *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Shanghai, 2006, pp. 3551–3554.
- [5] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "On the impact of node failures and unreliable communications in dense sensor networks," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2535–2546, Jun. 2008.
- [6] H. Alemдар and C. Ersoy, "Wireless sensor networks for healthcare: A survey," *Comput. Netw.*, vol. 54, no. 15, pp. 2688–2710, Oct. 2010.
- [7] W. Xu, F. Quitin, M. Leng, W. P. Tay, and S. G. Razul, "Distributed localization of a RF target in NLOS environments," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 7, pp. 1–14, Jul. 2015.
- [8] J. N. Tsitsiklis, "Decentralized detection," *Adv. Stat. Signal Process.*, vol. 2, pp. 297–344, 1993.
- [9] R. Viswanathan and P. K. Varshney, "Distributed detection with multiple sensors: Part I—Fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, Jan. 1997.
- [10] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "Data fusion trees for detection: Does architecture matter?" *IEEE Trans. Inf. Theory*, vol. 54, no. 9, pp. 4155–4168, Sep. 2008.
- [11] W. P. Tay, "The value of feedback in decentralized detection," *IEEE Trans. Inf. Theory*, vol. 58, no. 12, pp. 7226–7239, Dec. 2012.
- [12] H. Chen, B. Chen, and P. Varshney, "A new framework for distributed detection with conditionally dependent observations," *IEEE Trans. Signal Process.*, vol. 60, no. 3, pp. 1409–1419, Mar. 2012.
- [13] W. P. Tay, "Whose opinion to follow in multihypothesis social learning? A large deviations perspective," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 2, pp. 344–359, Mar. 2015.
- [14] J. Ho, W. P. Tay, T. Q. Quek, and E. K. Chong, "Robust decentralized detection and social learning in tandem networks," *IEEE Trans. Signal Process.*, vol. 63, no. 19, pp. 5019–5032, Oct. 2015.
- [15] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals, Syst.*, vol. 1, pp. 167–182, 1988.
- [16] J.-F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 407–416, Feb. 2003.
- [17] S. A. Aldosari and J. M. F. Moura, "Fusion in sensor networks with communication constraints," in *Proc. Int. Symp. Inf. Process. Sensor Netw.*, Berkeley, CA, USA, 2004, pp. 108–115.
- [18] D. Boneh, E.-J. Goh, and K. Nissim, "Evaluating 2-DNF formulas on ciphertexts," in *Proc. Int. Conf. Theory Cryptography*, Cambridge, MA, USA, 2005, pp. 325–341.
- [19] Y. Ishai and A. Paskin, "Evaluating branching programs on encrypted data," in *Proc. Int. Conf. Theory Cryptography*, Berlin, Germany, 2007, pp. 575–594.
- [20] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proc. ACM Symp. Theory Comput.*, Bethesda, MD, USA, 2009, pp. 169–178.
- [21] C. Dwork, "Differential privacy: A survey of results," in *Proc. Int. Conf. Theory Appl. Models Comput.*, Xi'an, China, 2008, pp. 1–19.
- [22] A. Friedman and A. Schuster, "Data mining with differential privacy," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Washington, DC, USA, 2010, pp. 493–502.
- [23] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security*, Berlin, Germany, 2013, pp. 901–914.
- [24] R. Chen, B. C. Fung, S. Y. Philip, and B. C. Desai, "Correlated network data publication via differential privacy," *Int. J. Very Large Data Bases*, vol. 23, no. 4, pp. 653–676, Nov. 2014.
- [25] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Proc. IEEE Symp. Found. Comput. Sci.*, Berkeley, CA, USA, 2013, pp. 429–438.
- [26] H. Yamamoto, "A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 6, pp. 918–923, Nov. 1983.
- [27] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "A theory of utility and privacy of data sources," in *Proc. IEEE Int. Symp. Inf. Theory*, 2010, pp. 2642–2646.
- [28] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, USA, 2012, pp. 1401–1408.
- [29] Z. Li and T. J. Oechtering, "Differential privacy in parallel distributed Bayesian detections," in *Proc. Int. Conf. Inf. Fusion*, Salamanca, 2014, pp. 1–7.
- [30] S. Salamatian et al., "Managing your private and public data: Bringing down inference attacks against your privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1240–1255, Oct. 2015.
- [31] K. Diamantaras and S.-Y. Kung, "Data privacy protection by kernel subspace projection and generalized eigenvalue decomposition," in *IEEE Int. Workshop Mach. Learn. Signal Process.*, Salerno, 2016, pp. 1–6.
- [32] M. Al, S. Wan, and S.-Y. Kung, "Ratio utility and cost analysis for privacy preserving subspace projection," arXiv:1702.07976.
- [33] X. Nguyen, M. J. Wainwright, and M. I. Jordan, "Nonparametric decentralized detection using kernel methods," *IEEE Trans. Signal Process.*, vol. 53, no. 11, pp. 4053–4066, Nov. 2005.
- [34] W. Wang, Y. Liang, E. P. Xing, and L. Shen, "Nonparametric decentralized detection and sparse sensor selection via weighted kernel," *IEEE Trans. Signal Process.*, vol. 64, no. 2, pp. 306–321, Jan. 2016.
- [35] M. Sun and W. P. Tay, "Privacy-preserving nonparametric decentralized detection," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Shanghai, 2016, pp. 6270–6274.
- [36] NB-IOT—Enabling new business opportunities. 2016. [Online]. Available: <http://www.huawei.com/minisite/4-5g/img/NB-IOT.pdf>
- [37] E. Shi, H. Chan, E. Rieffel, R. Chow, and D. Song, "Privacy-preserving aggregation of time-series data," in *Proc. Netw. Distrib. Syst. Security Symp.*, San Diego, 2011.
- [38] J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle, "Privacy in the Internet of things: Threats and challenges," *Security Commun. Netw.*, vol. 7, no. 12, pp. 2728–2742, Jun. 2014.
- [39] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, "Differentially private empirical risk minimization," *J. Mach. Learn. Res.*, vol. 12, pp. 1069–1109, Mar. 2011.
- [40] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Privacy aware learning," *J. ACM*, vol. 61, no. 6, Dec. 2014, Art. no. 38.
- [41] R. Hall, A. Rinaldo, and L. Wasserman, "Differential privacy for functions and functional data," *J. Mach. Learn. Res.*, vol. 14, pp. 703–727, Feb. 2013.
- [42] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 2006.
- [43] B. Scholkopf and A. J. Smola, *Learning With Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT Press, 2001.
- [44] T. Zhang, "Statistical behavior and consistency of classification methods based on convex risk minimization," *Ann. Statist.*, vol. 32, no. 1, pp. 56–85, Feb. 2004.
- [45] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *Ann. Statist.*, vol. 28, no. 2, pp. 337–407, 2000.
- [46] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Proc. Eur. Conf. Comput. Learn. Theory*, 1995, pp. 23–37.
- [47] L. Rosasco, E. De Vito, A. Caponnetto, M. Piana, and A. Verri, "Are loss functions all the same?" *Neural Comput.*, vol. 16, no. 5, pp. 1063–1076, May 2004.
- [48] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe, "Convexity, classification, and risk bounds," *J. Amer. Stat. Assoc.*, vol. 101, no. 473, pp. 138–156, Jan. 2006.
- [49] H. Wendland, *Scattered Data Approximation*. Göttingen, Germany: Cambridge Univ. Press, 2004.
- [50] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 2015.
- [51] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [52] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss-Seidel method under convex constraints," *Oper. Res. Lett.*, vol. 26, no. 3, pp. 127–136, 2000.
- [53] Y. Xu and W. Yin, "A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion," *SIAM J. Imag. Sci.*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [54] P. Tseng, "Convergence of a block coordinate descent method for non-differentiable minimization," *J. Optim. Theory Appl.*, vol. 109, no. 3, pp. 475–494, 2001.
- [55] A. C. Lorena, A. C. De Carvalho, and J. M. Gama, "A review on the combination of binary classifiers in multiclass problems," *Artif. Intell. Rev.*, vol. 30, no. 1, pp. 19–37, 2008.

- [56] M. Lichman, "UCI machine learning repository," 2013. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [57] S.-L. Sun, "Distributed optimal component fusion weighted by scalars for fixed-lag Kalman smoother," *Automatica*, vol. 41, no. 12, pp. 2153–2159, 2005.
- [58] R. Chavarriaga *et al.*, "The opportunity challenge: A benchmark database for on-body sensor-based activity recognition," *Pattern Recognit. Lett.*, vol. 34, no. 15, pp. 2033–2042, 2013.
- [59] M. Hall *et al.*, "The WEKA data mining software: An update," *SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, 2009.
- [60] W. Rudin, *Principles of Mathematical Analysis*. New York, NY, USA: McGraw-Hill, 1964.

Meng Sun (S'17) received the B.Eng. degree in communication engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2013. She is currently working toward the Ph.D. degree at the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. Her research interests include decentralized detection, cybersecurity, applied probability, and information theory.

Wee Peng Tay (S'06–M'08–SM'14) received the B.S. degree in electrical engineering and mathematics, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2002, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2008.

He is currently an Associate Professor with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. His research interests include distributed inference and signal processing, sensor networks, social networks, information theory, and applied probability. He was the recipient of the Singapore Technologies Scholarship in 1998, the Stanford University President's Award in 1999, the Frederick Emmons Terman Engineering Scholastic Award in 2002, and the Tan Chin Tuan Exchange Fellowship in 2015. He is a coauthor of the paper that received the Best Student Paper Award at the Asilomar Conference on Signals, Systems, and Computers in 2012, and the IEEE Signal Processing Society Young Author Best Paper Award in 2016. He is currently an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, an Editor for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, serves on the MLSP TC of the IEEE Signal Processing Society, and is the Chair of DSNIG in IEEE Multimedia Communications Technical Committee. He has also served as a Technical Program Committee member for various international conferences.

Xin He received the B.Eng. degree in electronic information from Sichuan University, Chengdu, China, in 2007, the M.Eng. degree in information and communication from the Beijing University of Posts and Telecommunications, Beijing, China, in 2010, and the Ph.D. degree in EEE from the University of Hong Kong, Hong Kong, in 2015. He is currently a Research Assistant Professor with the College of Innovation and Entrepreneurship, Southern University of Science and Technology, Shenzhen, China. His research interests include optimization and statistics for probabilistic inference and beamforming.