Privacy-preserving decentralized detection in sensor networks

Sun, Meng

2019

Sun, M. (2019). Privacy-preserving decentralized detection in sensor networks. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/105947>

<https://doi.org/10.32657/10220/48831>

Privacy-preserving Decentralized Detection in Sensor Networks

Sun Meng

School of Electrical and Electronic Engineering

A thesis submitted to the Nanyang Technological University
in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

2018

Statement of Originality

I hereby certify that the work embodied in this thesis is the result of original research, is free of plagiarised materials, and has not been submitted for a higher degree to any other University or Institution.

2019-04-03
.....
Date

Sun Meng
.....
Sun Meng

Supervisor Declaration Statement

I have reviewed the content and presentation style of this thesis and declare it is free of plagiarism and of sufficient grammatical clarity to be examined. To the best of my knowledge, the research and writing are those of the candidate except as acknowledged in the Author Attribution Statement. I confirm that the investigations were conducted in accord with the ethics policies and integrity standards of Nanyang Technological University and that the research data are presented honestly and without prejudice.

2019-04-03
.....
Date

.....
Tay, Wee Peng

Authorship Attribution Statement

This thesis contains material from 3 paper(s) published in the following peer-reviewed journal(s) / from papers accepted at conferences in which I am listed as an author.

Chapter 3 is published as M. Sun, W. P. Tay, and X. He, "Toward information privacy for the internet of things: A nonparametric learning approach," IEEE Trans. Signal Process., vol. 66, no. 7, pp. 1734-1747, April 2018.

The contributions of the co-authors are as follows:

- I prepared the manuscript drafts. The manuscript was revised by Prof. Tay Wee Peng.
- Prof. Tay Wee Peng and Dr. He Xin discussed the problem formulation with me.
- Prof. Tay Wee Peng assisted in the proofs for theorems and propositions in this paper.

Chapter 4 is submitted as M. Sun and W. P. Tay, "Decentralized Detection with Robust Information Privacy Protection," submitted to IEEE Trans. Inf. Forensics Security.

The contributions of the co-authors are as follows:

- I prepared the manuscript drafts. The manuscript was revised by Prof. Tay Wee Peng.
- Prof. Tay Wee Peng discussed the problem formulation and with me.

- Prof. Tay Wee Peng assisted in the proofs for theorems and propositions in this paper.

Chapter 3 is published as M. Sun and W. P. Tay, "On the Relationship Between Inference and Data Privacy in Decentralized IoT Networks," submitted to IEEE Trans. Inf. Forensics Security.

The contributions of the co-authors are as follows:

- I prepared the manuscript drafts. The manuscript was revised by Prof. Tay Wee Peng.
- Prof. Tay Wee Peng discussed the problem formulation and with me.
- Prof. Tay Wee Peng assisted in the proofs for theorems and propositions in this paper.

2019-04-03
Date

.....
Sun Meng

Acknowledgments

Firstly, I want to express my sincere gratitude to my supervisor Prof. Tay Wee Peng, who patiently assisted me to learn how to do research from scratch, taught me to think and analyze problems systematically. With great patience, he always gave me useful suggestions and hints when I got stuck in my research. With immense knowledge and passion, he taught me what are the interesting and meaningful problems to consider. I feel so honored and grateful for working with him, and the serious work attitude I learnt from him will benefit me all my life.

Secondly, I would like to thank my former supervisor, Prof. Li Yubai at University of Electronic Science and Technology of China (UESTC) for his invaluable guidance and long lasting encouragement. The graduate studies with him stimulated my interests in exploring the unknowns and paved the way for my PhD research. His hard working attitude and profound knowledge have influenced significantly on my future career.

I would also like to thank each member in Distributed and Adaptive algorithms in Sensor Networks (DASN) group. My sincere thanks goes to Dr. Luo Wuqiong, Dr. Tang Jianhua, Dr. Hu Wuhua, Dr. Ji Feng and Dr. Zhangyi. The discussions with them really benefit me and inspire me to rethink my problems from different perspectives. I would also like to thank Dr. Leng Mei, Dr. Wang Chengcheng, Wang Yuan, Dr. Wang Junshan, Tang Wenchang, Yang Jielong, Kang qiyu, Wang Chongxiao, Dr. Song Yang for the insightful discussions and valuable comments to my work.

I thank my fellow labmates in INFINITUS for the stimulating discussions and

all the fun we have had in the last years, including Dr. Luo Sheng, Dr. Yang Rong, Dr. Zhao Ming, Dr. Ma Lijia, Li Beibei, Wang Guoming, Shao Hongxin, Wang Yang, Zhang Heng, Wang Wei, Gao Yumeng, Liu Yunxiang, Li Haoliang. The days with you in Singapore were so wonderful.

Last but not the least, I would like to thank my parents who have always been giving me unconditional love and support. In particular, I must thank my husband, Dr. Hu Hao, who sacrificed his time for delicious food and travel, and came to my place to accompany me every weekend during my Ph.D study.

Abstract

With a burgeoning number of Internet of Things (IoT) devices penetrating into all aspects of our lives, privacy-related issues are attracting increasing interest. Sensor network, one of the most important enabling technology of IoT, should be designed with privacy considerations. By giving users more control over what information can be shared from the sensors to the service providers, we will encourage the adoption of IoT technologies and put users at ease.

We model the sensor network with decentralized detection framework, where each sensor make a local decision based on its observation of hypotheses, and transmit the local decision to the fusion center. The fusion center make inferences based on the received sensor decisions. In this thesis, we aim to find privacy mapping at each sensor to distort sensor observations before send to the fusion center, such that privacy is protected, while still enabling the fusion center to make accurate detection of the public hypothesis.

Firstly, we consider protecting information privacy of private hypothesis without assuming knowledge of joint distribution of the sensor observations and hypotheses. In a sensor 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. 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.

Secondly, we consider protecting information privacy of a set of private hypotheses with known joint distribution of the sensor observations and hypotheses. We consider the fact that privacy concern is usually not about a single hypothesis, small deviation from a nominal private hypothesis should also be kept private from the fusion center. We find a representative private hypothesis, which is the easiest to detect among the set of private hypotheses. We propose an algorithm, by protecting the information privacy of the representative private hypothesis, the information privacy of the set of private hypotheses is protected. We consider the two cases where the number of sensors is finite and infinite respectively.

Finally, we discuss the relationship between various privacy metrics proposed in the literature. We divide the privacy metrics into inference privacy and data privacy metrics. Here, data privacy refers to the concealment of sensors' raw observation from the fusion center, while reducing the disclosure of private states of the object to the fusion center is inference privacy. We show that inference and data privacy are in general not equivalent. We propose methods to protect both inference and data privacy in decentralized detection, by incorporating local differential privacy (data privacy) and information privacy (inference privacy) metrics. We consider both case with and without the known prior knowledge of the sensor observations' distribution.

Contents

Statement of Originality	iii
Supervisor Declaration Statement	iv
Authorship Attribution Statement	v
Acknowledgment	vii
Abstract	ix
Table of Contents	xi
List of Figures	xv
List of Tables	xvii
List of Symbols	xviii
List of Abbreviations	xix
1 Introduction	1
1.1 Background	1
1.2 Research Motivation	3
1.2.1 Preserving Inference Privacy in Sensor Network	3
1.2.2 Parametric Robust Decentralized Detection with Information Privacy	4

1.2.3	Decentralized Detection with Inference Privacy and Data Privacy	5
1.3	Related Work	6
1.3.1	Privacy Preservation in Fully Distributed, Centralized and Decentralized Networks	6
1.3.2	Data Privacy and Inference Privacy	7
1.3.3	Differential Privacy	9
1.3.4	Decentralized Detection and Robust Detection	10
1.3.5	Theoretical Bounds	11
1.4	System Model	11
1.5	Thesis Contributions and Organization	13
2	Fundamental concepts	17
2.1	Introduction	17
2.2	Hypothesis Detection	17
2.2.1	Bayesian Hypothesis Testing	17
2.2.2	Decentralized Detection	18
2.3	Kernel Methods	20
2.4	Empirical Risk Minimization	22
2.5	Information Theory	23
2.6	Privacy Metrics	24
3	Nonparametric Decentralized Detection with Information Privacy	27
3.1	Problem Formulation and Information Privacy	27
3.1.1	System Model	27
3.1.2	Information Privacy	28
3.2	Privacy Mapping Design	30
3.2.1	Nonparametric Privacy-aware Optimization	31
3.2.2	Iterative Optimization	35
3.2.3	Extension to m -ary Private Hypothesis	38

3.3	Simulations and Experiments	41
3.3.1	Synthetic Data Set	42
3.3.2	UCI Repository Data Sets	47
Appendix 3.A	Proof of Proposition 3.1	49
Appendix 3.B	Proof of Theorem 3.1	51
Appendix 3.C	Proof of Proposition 3.2	54
Appendix 3.D	Proof of Proposition 3.3	55
Appendix 3.E	Proof of Theorem 3.2	55
4	Parametric Decentralized Detection with Robust Information Privacy Protection	56
4.1	System Model and Error Bounds	56
4.2	Error Bounds and Optimization Relaxation	61
4.2.1	Error Bounds	61
4.2.2	Optimization Relaxation	63
4.2.3	Most Favorable Hypothesis (MFH)	65
4.3	Privacy Mapping Design	68
4.3.1	Fixed number of sensors	68
4.3.2	Large number of sensors	71
4.4	Simulations	74
4.4.1	Asymptotically Large Number of Sensors	78
4.4.2	Comparisons	80
Appendix 4.A	Proof of Proposition 4.1	82
Appendix 4.B	Proof of Theorem 4.1	84
Appendix 4.C	Proof of Theorem 4.2	88
5	Decentralized Detection with both Inference and Data Privacy	90
5.1	System Model	90
5.2	Relationships Between Privacy Metrics	91
5.3	Parametric Case: Concatenated Privacy Mappings	94

5.3.1	Data Privacy using local differential privacy	95
5.3.2	Joint Inference and Data Privacy	97
5.4	Nonparametric Case: Empirical Risk Optimization	103
5.5	Numerical Results	106
5.5.1	Parametric Case Study	106
5.5.2	Nonparametric Case Study: Simulations	110
5.5.3	Nonparametric Case Study: OPPORTUNITY Data Set and Adult Data Set	113
Appendix 5.A	Proof of Theorem 5.1	116
6	Conclusion and Future Work	120
7	Author's Publications	122
	Bibliography	132

List of Figures

1.1	An IoT network with public hypothesis H and private hypothesis G .	13
2.1	Decentralized Detection	19
2.2	On the left, two variables are in two dimensional space, and they are not separable with a linear boundary. On the right, two variables are mapped to a three dimensional space, where there's a linear boundary to separate them.	21
3.1	Error rate of detecting H and G with different correlation coefficient between H and G	44
3.2	Error rate of detecting H and G with varying privacy threshold ratio p	45
4.1	Upper and lower bounds in Theorem 4.1 and $\mathbb{P}(\gamma_H(Z \neq H))$ calculated by Algorithm 4.1 with varying $I(X; H)$	75
4.2	Joint distribution p_{X^1, H, G° of sensor observation X^1 , public hypothesis H , and private hypothesis G° . The correlation coefficient between H and G° is 0.2.	76
4.3	Conditional probability distribution of $Z = (Z^1, \dots, Z^4)$ given different hypotheses. The horizontal axis shows the decimal form of the binary word $Z^1 Z^2 Z^3 Z^4$	76
4.4	Privacy mapping at sensor 1, with $r = 0.7$ and $\delta = 0.54$	77
4.5	Bayes error probability of detecting H , G° and G_{MF} with different privacy threshold ratio r and uncertainty parameter δ	78

4.6	Bayes error of detecting H , G° and G_{MF} with varying correlation coefficient between H and G°	79
4.7	Error exponent for detecting H and G° with varying correlation coefficient between H and G°	79
4.8	Error exponent for detecting H and G° with varying $ \mathcal{Z} $	80
4.9	Bayes error for detecting H and G° with varying conditional mutual information $I(X; H \mid G^\circ)$	82
5.1	Relationships between different privacy metrics for an IoT network with fixed q, s . An arrow $\rightarrow$ means “implies” while $\nrightarrow$ means “does not guarantee”.	92
5.2	Each sensor t 's privacy mapping $p_{Z^t X^t}^t = p_{Y^t X^t}^t \cdot p_{Z^t Y^t}^t$ consists of two privacy mappings concatenated together.	98
5.3	First stage privacy mapping at sensor 1 with Information-Local differential privacy (ILL).	107
5.4	Second stage privacy mapping at sensor 1 with ILL.	107
5.5	Bayes error for detecting H and G under Local differential-Information Privacy (LIP) and ILL for fixed privacy threshold ratio r and varying local differential privacy budget ϵ_{LD}	108
5.6	Bayes error for detecting H and G under LIP and ILL for fixed local differential privacy budget ϵ_{LD} and varying ϵ_I	108
5.7	Mutual information with $\epsilon_I = 0.15$ and varying ϵ_{LD} for ILL, LIP, InP and Local Differential Privacy (LDP).	109
5.8	Bayes error for detecting H and G , and mutual information between X^1 and Z^1 with different local differential privacy budget ϵ_{LD} . . .	111
5.9	Bayes error probability of detecting H and G with varying correlation coefficient between H and G	112

List of Tables

3.1	Joint distribution of Z and G	30
3.2	Loss functions, their conjugate duals, and parameters for Assumption 3.3	34
3.3	Sensor observations X for different realizations of (H, G)	43
3.4	Comparison of NPO with (3.21)	46
3.5	Bayes Error with m -ary Private Hypothesis	46
3.6	Detection errors using the OPPORTUNITY Activity Recognition Data Set.	48
3.7	Detection errors using the Census (Adult) Data Set.	49
5.1	Sensor observation for different realizations of (H, G)	111
5.2	Detection errors using the OPPORTUNITY Activity Recognition Data Set.	115
5.3	Detection errors using the Adult Data Set.	115
5.4	The joint distribution of V and U , where $\alpha \in [0, 1]$	116

List of Symbols

Throughout this thesis, capital letters like X denotes random variable or vector, lowercase letters like x denotes deterministic scalars, and boldface lowercase letters like $\mathbf{x}$ denote deterministic vectors. 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.

$\mathbb{R}$	Real numbers
$(\cdot)^T$	Transpose operation
Γ^c	Complement of set Γ
$\mathbf{1}_{\{A\}}$	Indicator function: takes value 1, if and only if clause A is true, otherwise takes value 0
$(x)_+$	$= \max\{x, 0\}$
$\mathbb{E}_X$	The expectation with respect to (w.r.t.) X
$p_{X Y}(\cdot \cdot)$	The conditional probability mass function of X given Y
$p_{A,B,C}$	The joint probability mass function of (A, B, C)
$I(A; B)$	Mutual information between A and B
$X \perp\!\!\!\perp Y$	X and Y independent of each other
$\ \cdot\ _{TV}$	Total variation
$-\mathbf{1}$	A vector of -1
$\mathbf{0}$	A vector of 0
$\log$	natural logarithm
$\downarrow$	approach a point from above
$\inf$	the greatest lower bound of a set

$x \sim x'$ two vectors x and x' are neighbors (they differ in only one of their vector components)

List of Abbreviations

ERM	Empirical Risk Minimization.
GDPR	General Data Protection Regulation.
i.i.d.	independently indentially distributed.
IoT	Internet of Things.
LFD	Least Favorable Distributions.
MFD	Most Favorable Distribution.
MFH	Most Favorable Hypothesis.
LP	Linear Program.
GS	block nonlinear Gauss-Siedel method.
RKHS	Reproducing Kernel Hilbert Space.
w.r.t.	with respect to.
NPO	Nonparamatric Privacy-aware Optimization.
E-LDP	A non-parametric method with Empirical Local Differential Privacy.
EPIC	A non-parametric method with Empirical information and local differential PrIvaCy.
InP	A parametric detection method with only Information Privacy metric.
LDP	A parametric detection method with only Local Differential Privacy metric.
ILL	A parametric detection method with Information-Local differential privacy.

LIP A parametric detection method with Local differential-
Information Privacy.

Chapter 1

Introduction

1.1 Background

IoT 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 IoT paradigm, large numbers of sensors are deployed to see, hear and share information, collaboratively enabling sense-making and intelligent analytics based on the sensors' observations. A large proportion of IoT network 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. The fusion center, furnished with all the sensor summaries will then makes the final inference on a phenomenon of interest.

IoT networks provide smarter, and more efficient solutions, and the number of connected “things” in sensor networks is ever increasing. By 2020, the number is expected to reach 50 billion [15]. Since huge amount of data is transmitted among these devices, security and privacy issues are becoming more severe. Previous researches focus more on the security issue in sensor networks. Extensive solutions have been proposed to tackle with silent attacks on service integrity, and attacks on network availability [16].

Attacks on secrecy and authentication is also studied to guarantee the confidentiality. Measures should be taken to make sure the right people can in fact get the important information, while preventing sensitive information from reaching the wrong people. However, a secure IoT network can still invade privacy, since even the right person can be curiously peeking at our private information. In IoT network, the fusion center collects our data wherever and whenever we are to provide better services for us. However, the fusion center can also use the data to make inference about our private information. People are not clear about what data is collected nor how the data is used, thus privacy concerns are surfacing.

Laws and regulations are put in place to protect privacy. On 25 May 2018, General Data Protection Regulation (GDPR) goes into effect, giving European Union citizens the right to clearly and explicitly opt into having their data collected and used by a company on the web. Technical companies also started to take privacy into concern. Apple Inc., for example, claimed that their architecture incorporated local differential privacy where data is randomized before sent from the device. These efforts aims at protecting privacy for online data, however the privacy preservation is not that strong, and utility might be severely sabotaged due to privacy preservation. This thesis studies privacy preservation in decentralized detection, and proposes better solution with more rigorous privacy controls and better privacy utility tradeoff.

In this thesis, both inference privacy and data privacy are considered. Here, inference privacy refers to preventing the fusion center from performing certain statistical inferences with no authorization. Data privacy, on the other hand, refers to preventing the fusion center from knowing the raw observation of sensors. Privacy-preserving decentralized detection methods is proposed so that utility is maximized while privacy is protected. Method is proposed for both cases where the joint distribution of sensor observation and hypothesis is known apriori or unknown. Furthermore, robust privacy preservation is studied, and multiple sensitive private information can be protected when performing inference on public

hypothesis of interest.

1.2 Research Motivation

In this section, the research motivations of different privacy-preserving decentralized detection methods are highlighted.

1.2.1 Preserving Inference Privacy in Sensor Network

The fusion center's role is to perform inference on a particular hypothesis of interest based on the received sensor information. However, there is nothing stopping it from using the same information to infer another correlated hypotheses. As a result, sensitive user information may be inferred by the fusion center without the user's consent. 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 thesis, 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.

In light of this motivation, part of this thesis focuses on finding a rule to distort sensor observations before send to the fusion center. With the distorted sensor information, the fusion center cannot perform accurate inference on the private hypothesis, while still be able to accurately infer the public hypothesis. It has been shown in [17–19] 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.

1.2.2 Parametric Robust Decentralized Detection with Information Privacy

In most applications, the sensor network is designed for a single public hypothesis. By contrary, there might exists multiple private aspects that people find sensitive. People who consider income as a sensitive private hypothesis, for example, might also find related aspects like occupation, monthly expenditure, housing condition sensitive. People who don’t hope to be tracked, might find their GPS data, restaurants they visit, or their daily step count private. Moreover, unlike a clearly specified public hypothesis, people usually don’t have a clear image of all the private hypotheses, thus enumerating all the private hypotheses would be impossible. Therefore, instead of considering a single private hypothesis, a set of sensitive hypotheses should be defined.

There are works studying detection method which prevent the fusion center from performing statistical inference of private hypothesis. However, to the best of our knowledge, there is no work studying the protection of a set of private

hypotheses. Motivated by this reason, we propose a framework that allows the user to specify a nominal private hypothesis, and our goal is to design privacy mappings at each sensor to protect the information privacy [22] of a set of private hypotheses that are within a certain distance of the nominal hypothesis (in a specific sense as defined in 4).

1.2.3 Decentralized Detection with Inference Privacy and Data Privacy

The data collected by the IoT devices enables the services brought to the customers. However, the data itself can be sensitive. Web search records and everything you type on a keyboard, for example, tell much more about you than you wish to share. Therefore, data privacy should be protected, so that people are aware of and have control over what data collected by IoT devices is shared to third parties. Laws and regulations are put in place to protect data privacy.

However, protecting data privacy alone is not reassuring, inference privacy should also be ensured. The key idea of a data-privacy-preserving system is to obfuscate the raw data while still being able to extract statistical information out of the data. So even if opted in data-privacy-preserving system, customers' sensitive private information may still be exposed to the service provider, since statistical information can be extracted out of large number of privatized data. Inference privacy, different from data privacy, refers to preventing the service provider from performing certain statistical inferences with no authorization.

In view of this motivation, part of this thesis is focused on designing a privacy-preserving decentralized detection architecture where customers can control the level of their data privacy and inference privacy. We first study the relationships between different privacy metrics proposed in the literature. We claim that both information privacy and local differential privacy metrics should be incorporated in system to provide a strong privacy guarantee. Based on this claim, we propose a parametric privacy-preserving decentralized detection architecture. In order

to protect data privacy and inference privacy, we propose an algorithm to design privacy mappings whereby each sensor distort its observations locally before sending them to the fusion center. We aim to achieve a good tradeoff among data privacy, inference privacy and the accuracy of detection at the fusion center for the public hypothesis.

1.3 Related Work

1.3.1 Privacy Preservation in Fully Distributed, Centralized and Decentralized Networks

Various works have focused on protecting data privacy while providing utility. In a privacy-preserving consensus network, each node share obfuscated information with each other. The papers [20–27] proposed methods that allow the nodes to obtain the correct information collaboratively without sharing their private observations. These works consider data privacy preserving methods for a fully distributed network where there is no fusion center. This is different from the IoT model that we study in this thesis and is out of our current scope. Moreover, the issue of inference privacy has also not been addressed.

In [28–30], the authors analysed distributed privacy-preserving data mining. They considered privacy preservation over both horizontally partitioned data sets where different parties hold different sets of record with the same attributes and vertically partitioned data sets where different parties hold same sets of record with different attributes. Various classical data mining applications such as Naive Bayes, SVM classifier have been extended to protect privacy of both kinds separately. Others also investigated privacy preservation from an information-theoretic point of view. In [31], the authors introduced information bottleneck problem. It aims at finding a short code $\tilde{X}$ to represent the information Y contained in the original signal X . It protects the privacy of information other than Y since ideally $\tilde{X}$ can squeeze other information out of X . In [32], the authors

considers the situation where we can get information about both Y^+ and Y^- from variable X . They proposed a method to extract information $\tilde{X}$ from X such that it contains relevant information of Y^+ while preserving little information about Y^- . All these works, however, develop privacy techniques for centralized platforms where a trusted data curator is assumed to be available. In this thesis, our aim is to develop privacy mechanisms for sensors in a decentralized network.

In a decentralized platform, without a trusted data curator, each party sanitizes its own data before sending the data to a fusion center. The sanitized data still enables the fusion center to perform certain functions, while not disclosing private information to the fusion center. The papers [33–37] propose methods to achieve (local) differential privacy, while the works [38–41] investigate privacy-preserving distributed data analytics. Other works resort to cryptographic techniques like fully homomorphic encryption (FHE) [42] and CryptoNets [43, 44] to encrypt information from each sensor while allowing the fusion center to perform computations on the sensors' ciphertext. All the abovementioned methods aims at protecting data privacy, so that the raw data is not accessible by the fusion center. However, either the fusion center may still be able to make statistical inference of a private hypothesis based on data that is sanitized for data privacy or inference of the public hypothesis is severely degraded [45].

1.3.2 Data Privacy and Inference Privacy

The privacy to be preserved can be mainly categorized into data privacy and inference privacy.

Data privacy refers to the ability an individual has to determine what data about him/her can be shared with third parties. In cloud services, homomorphic encryption [42–44, 46] is used to allow cloud computing on encrypted data without decryption. This serves as a good way to utilize cloud computing without disclosing the data and the result to the cloud. However, decentralized detection is a different scenario where the fusion center should perform the computation on the

data and get the result without knowing the data. For methods more suitable to decentralized detection, some protect data privacy with respect to a certain metric. The paper [47] analyzed the tradeoff between utility and data privacy, and compared the performance of different data privacy metrics: local differential privacy, identifiability, and mutual information. One example of data privacy metric is local differential privacy [41]. It guarantees the distributions of data sent from the sensors are similar, if sensors' observations are similar.

Inference privacy, on the other hand, refers to the ability an individual has to determine what statistical inference can be made about him/her. 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 its true state is not authorized to be detected or inferred by the fusion center. The authors of [48] proposed three inference privacy metrics to measure the exposure of private hypothesis: information privacy, differential privacy, and average information leakage, and showed that information privacy is the strongest among the three, while differential privacy does not guarantee information privacy. They define information privacy to ensure that the prior probability and posterior probability of a private hypothesis are similar. The paper [49] 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 [50] extends the result of [49], and designed a framework to quantify the privacy-utility tradeoffs. An information-theoretic scheme for information privacy was proposed by [48], 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 [51], 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.

In this thesis, data privacy and inference privacy refer to the ability to avoid information disclosure, while information privacy, differential privacy and so on, are defined as metrics measuring the leakage of information.

Inference privacy and data privacy generally don't imply each other. In [36] maximum leakage is used as the privacy metric to limit inference privacy leakage. With type II error exponent as the utility metric, the authors conclude that it will invade data privacy. On the other hand, data privacy constraints will not prevent the fusion center from making statistical inference, otherwise, accurate detection about the public hypothesis cannot be made.

1.3.3 Differential Privacy

Differential privacy was proposed by Dwork [40, 52, 53]. It considers the scenario where statistical information is to be learnt from a statistical database. With differential privacy, it guarantees that the release of statistical information does not compromise the privacy of each individual. Typical method to protect differential privacy is to add Laplacian noise to the released statistical information so that if the change of one data entry does not result in changes of the released information. It has wide application to various fields. In privacy-preserving data mining [54], deep learning [55], sensitive data is protected while mining for useful information. In [56], the authors proposed differentially private method to get accurate range-count queries via wavelet transform.

Differential privacy is proposed as a concept in database. In IoT sensor network, moreover, local differential privacy [35, 41, 57, 58] is proposed, which is closely related to differential privacy. It is a relatively stringent requirement, which guarantees for each sensor, different sensor observations result in the same sensor output with similar probability.

1.3.4 Decentralized Detection and Robust Detection

Decentralized detection was first introduced by [59], and [8] elaborates the formulation and solution. A large proportion of real life applications can be formulated with decentralized detection. In some applications related to IoT, it's hard to know the joint distribution of hypotheses and sensor observations, thus nonparametric decentralized detection is needed in these scenarios. A nonparametric decentralized detection method was introduced by [60], 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, [61] extended this method using a weighted kernel to allow sensor selection in the decentralized detection procedure. These works however do not address the privacy issue described above. Following these works, we design sensor decision rules in decentralized detection which protects data and inference privacy.

We propose privacy-preserving decentralized method that is robust against a small deviation from the nominal private hypothesis. Robust detection was first proposed by [62]. The paper [63] finds a general solution for the minimax test between two composite hypotheses. In [64], robust decentralized detection is considered in an asymptotic setting, where independently identically distributed (i.i.d.) sensor observations is assumed. In all these works, a minimax test between a pair of representative simple hypotheses is conducted. In this thesis, we study decentralized detection, in both non-asymptotic [8] and asymptotic [17] settings. Similar to the previous works, we use a representative pair of hypotheses when seeking robust privacy preservation for private hypotheses. However, different from the previous works, where the main focus is to find a pair of distributions that is the hardest to distinguish, the focus of our paper is to find a private hypothesis in the uncertainty set that is the easiest to detect.

1.3.5 Theoretical Bounds

There are works analyzing the theoretical bounds with privacy consideration. For stochastic source outputs, capacity bounds are proved with asymptotic equipartition property (AEP). In [65], a broadcast channel with one common and two private source outputs is studied, where mutual information between the private output and channel symbol string is used to quantify leakage rate. Inner and outer bounds on the leakage-capacity region are derived when the eavesdropper knows the codebook. In [49], two correlated stochastic outputs are transmitted through a one-way communication system, where one output should be kept as secret as possible, and the other output is supposed to be received within a distortion tolerance. Equivocation, the conditional entropy of private source output given the transmitted variable, is used to analyse rate-utility-privacy tradeoff in the paper.

Other works consider one-shot source outputs. [41] enables statisticians to make statistical inference from a dataset, without knowing exactly the data entries. Bounds for accuracy of statistical inference is analyzed with local differential privacy of the data. In [66], statistical inference is supposed to be made out of a measurement variable, while private variable correlated to that measurement variable is to be protected. minimum achievable privacy leakage is characterized when the amount of useful information is non trivial. In this thesis, we also study non-stochastic source outputs, and form a hypotheses detection problem. We analyze the tradeoff between the accuracy of detecting the public hypothesis and the information privacy of the private hypothesis.

1.4 System Model

In this thesis, we consider a decentralized detection network as shown in Fig.1.1. 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 $Z = \gamma^t(X) \in \mathcal{Z} = \{1, 2, \dots, |\mathcal{Z}|\}$, and transmits Z^t to a fusion center. Each sensor's local decision rule is allowed to be a probabilistic mapping from $\mathcal{X}$ to $\mathcal{Z}$. We interpret the privacy mapping p^t as a conditional probability $p^t(z^t | x^t) = p_{Z^t|X^t}(z^t | x^t)$. We constrain the transmission from each sensor to the fusion center to a limited number of bits, so that $|\mathcal{X}| \geq |\mathcal{Z}|$.

We denote $X = (X^t)_{t=1}^s$, $Z = (Z^t)_{t=1}^s$. Let $p_{Z|X}(\mathbf{z} | \mathbf{x}) = \mathbb{P}(Z = \mathbf{z} | X = \mathbf{x})$ denote the overall sensor mapping, and $\mathcal{Q}$ be the set of $p_{Z|X}(\cdot | \cdot)$ such that

$$p_{Z|X}(\mathbf{z} | \mathbf{x}) = \prod_{t=1}^s p^t(z^t | x^t), \quad (1.1a)$$

$$\sum_{z^t \in \mathcal{Z}} p^t(z^t | x^t) = 1, \quad (1.1b)$$

$$p^t(z^t | x^t) \geq 0, \quad (1.1c)$$

$$\sum_{x^t \in \mathcal{X}} p^t(z^t | x^t) \geq \Delta, \quad (1.1d)$$

$$\forall x^t \in \mathcal{X}, z^t \in \mathcal{Z}, t = 1, \dots, s,$$

where $\Delta \leq |\mathcal{X}|/|\mathcal{Z}|$ is a small positive constant. We impose the constraint in (1.1d) to ensure that every $z^t \in \mathcal{Z}$ has a positive probability mass because otherwise a smaller $|\mathcal{Z}|$ could have been used. Note that Δ has to be chosen no bigger than $|\mathcal{X}|/|\mathcal{Z}|$ since otherwise $|\mathcal{X}| = \sum_{z^t, x^t} p^t(z^t | x^t) > |\mathcal{X}|$, a contradiction.

This models an ad hoc IoT network with low-power devices that may be battery operated. For example, the NB-IoT standard [67] is developed for low-cost and low-power devices with limited communication bandwidths. In this thesis, we do not require that the sensor observations X^t , $t = 1, \dots, s$ are independent. 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 .

Figure 1.1: An IoT network with public hypothesis H and private hypothesis G .

1.5 Thesis Contributions and Organization

The main contributions of this thesis are summarized as follows:

- In Chapter 3, We adopt the concept of ϵ -information privacy from [48], 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.

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.

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.

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.

We verify the performance of our algorithm on both simulated data and real data. Our experiments suggest that our approach can achieve high error rates for detecting the private hypothesis, but low error rates for inferring the public hypothesis.

- In Chapter 4, We introduce uncertainty set to define a set of private hypotheses. We prove that if the set of private hypotheses are similar to the nominal hypothesis, it can be covered by a small uncertainty set. We define the most favorable hypothesis to be the hypothesis easiest to detect in the uncertainty set. We claim that by protecting the most favorable hypothesis, the set of private hypotheses are all protected. We define the sensor observations of the most favorable hypothesis as the most favorable distribution, and we find a probability distribution satisfies the definition of the most favorable distribution.

We consider the situation where a public hypothesis is correlated with a set of closely related private hypotheses. With known joint distribution of sensor observations and hypotheses, we propose a privacy-preserving decentralized detection algorithm to get the privacy mapping at each sensor. With the privacy mapping, the information privacy of the private hypotheses is protected,

while we can still have a high accuracy when detecting the public hypothesis. The proposed algorithm is proved to converge to a critical point. We also analyzed the theoretical bounds for utility given the privacy constraints.

We also study the case where the number of sensors is infinite. We formulate the problem with error exponent and prove that for optimal solution in this asymptotic setting, the cardinality of sensor decisions can be no larger than the cardinality of sensor observation added by two. We also prove that the optimal sensor decision rule is a randomization between two deterministic sensor decision rules.

- In Chapter 5, despite the many privacy metrics proposed in the literature, to the best of our knowledge, the interplay between inference privacy and data privacy and the relationship between different privacy metrics have not been adequately investigated.

In this thesis, we introduce the concept of privacy implication and non-guarantee, and show how one privacy metric is related to another in this framework. Moreover, we prove that information privacy metric is the strongest inference privacy metric amongst those considered, while local differential privacy metric is the strongest data privacy metric if sensor observations are uniformly distributed. Subsequently, we argue that in a practical IoT network, both information privacy and local differential privacy metrics should be incorporated when designing each sensor's privacy mapping.

We propose a parametric decentralized detection algorithm, with known joint distribution of public, private hypotheses and sensor observations. Supplementary to traditional decentralized detection [8], the proposed algorithm protect both data and inference privacy with local differential privacy and information privacy metrics. Taking into account that both local differential privacy and information privacy are proven to be immune to pre-processing

and post-processing, the proposed algorithm consists of two concatenate privacy mappings. One mapping protects local differential privacy, and the other protects information privacy with a comparatively low computational complexity. Simulations demonstrate that the proposed algorithm can protect both information privacy and local differential privacy, while maximizing the detection accuracy of public hypothesis.

The thesis is organized as follows.

Chapter 2 reviews the fundamental concept in hypothesis detection, including the principals of Bayesian detection, robust detection, kernel methods and empirical risk minimization. It also reviews information theory as an important tool utilized by this thesis. Chapter 3 proposes a nonparametric decentralized detection method to protect the information privacy of the private hypothesis. Chapter 4 allows the user to specify a nominal private hypothesis, and the goal is to design privacy mappings at each sensor to protect the information privacy [48] of an uncountable set of private hypotheses that are “close” (in a specific sense as defined in Section 4.1) to the nominal hypothesis. Chapter 5 compares different inference and data privacy metrics, and proposes decentralized detection method that prevent the fusion center from making statistical inference as well as knowing the sensor observation.

Chapter 2

Fundamental concepts

2.1 Introduction

In this section, we introduced the fundamental concepts necessary to understand this thesis. We will introduce the basics of hypothesis detection, kernel methods, empirical risk minimization, and information theory.

2.2 Hypothesis Detection

2.2.1 Bayesian Hypothesis Testing

Bayesian hypothesis testing is a commonly used tool to solve detection problems, where we wish to make decision among M possible statistical conditions based on observations [68]. Here we introduce some of the conclusion in binary hypothesis testing.

Assume there is a binary hypothesis, $H \in \{0, 1\}$. The observation related to H is X on a observation space $(\Gamma_0, \mathcal{G})$. Let $\Gamma \in \mathcal{G}$ and Γ^c be a partition of Γ_0 . For fixed partition of Γ , the decision rule γ is a function satisfying

$$\gamma(x) = \begin{cases} 0, & x \in \Gamma \\ 1, & x \in \Gamma^c \end{cases}$$

A correct decision is made if $\gamma(x)$ is the same as H . To minimize the error of detecting H , an optimum Γ should be chosen.

The error of detecting H is as the follows

$$\begin{aligned}
& \mathbb{P}(\gamma(x) \neq H) \\
&= \mathbb{P}(H = 1) \mathbb{P}(\gamma(x) = 0 | H = 1) + \mathbb{P}(H = 0) \mathbb{P}(\gamma(x) = 1 | H = 0) \\
&= \mathbb{P}(H = 1) \mathbb{P}(X \in \Gamma | H = 1) + \mathbb{P}(H = 0) \mathbb{P}(X \in \Gamma^c | H = 0) \\
&= \mathbb{P}(H = 0) + \mathbb{P}(H = 1) \mathbb{P}(X \in \Gamma | H = 1) - \mathbb{P}(H = 0) \mathbb{P}(X \in \Gamma | H = 0)
\end{aligned}$$

Therefore, the Bayesian error probability is minimized if we choose

$$\Gamma = \left\{ x \mid \frac{p(x|H=1)}{p(x|H=0)} < \frac{\mathbb{P}(H=0)}{\mathbb{P}(H=1)} \right\}.$$

2.2.2 Decentralized Detection

This thesis proposed privacy-preserving detection methods in decentralized sensor network. In this section, we will briefly introduce decentralized detection. Detailed explanation can be found in [8].

In a sensor network, each sensor has some observations about the state of the environment. In a centralized sensor network, each sensor sends all its observations to the fusion center, where a standard hypothesis detection is performed and the true state of the environment is decided. In a decentralized scheme, however, each sensor performs a local decision based on its observation and only sends its own local decision to the fusion center. Figure 2.1 shows a decentralized detection sensor network, where the state of the environment is denoted by random variable H , sensor observations are denoted by $X \in \{0, 1, \dots, |\mathcal{X}|\}$, and sensor local decisions are denoted by $Z \in \{0, 1, \dots, |\mathcal{Z}|\}$. The fusion center performs hypothesis testing about $H \in \{0, 1, \dots, M\}$ based on the received sensor local decisions. While the fusion center will do the same job, how should the sensors make the best local decisions becomes a problem.

Figure 2.1: Decentralized Detection

Compared to centralized scheme, the decentralized detection will yield weaker result since less information is received by the fusion center. However, the communication requirements of decentralized detection is dramatically reduced, since instead of sending the raw observation to the fusion center, the sensors only sends the local decisions (a few bits) to the fusion center.

Let γ_i denote the strategy at sensor i , where the local decision is made according to $Z^i = \gamma_i(X^i)$, $i \neq 0$. Let γ_0 denote the decision rule at the fusion center, $C(H, d, Z^1, \dots, Z^s)$ denote the cost of making decision d , when the true hypothesis is H and the received sensor decisions are $Z^1, \dots, Z^s$. We have the following proposition.

Proposition 2.1. *Assume $X^1, X^2, \dots, X^s$ are conditionally independent of each other given the hypothesis H .*

(a) *Fix some $i \in \{0, 1, \dots, s\}$, suppose for all $j \neq i$, γ_j is fixed. Then the γ_i to minimize the overall detection error is*

$$\gamma_i(x) = \arg \min_{z=1,2,\dots,|Z|} \sum_{j=1}^M p_{H|X}(j|x)$$

$$\mathbb{E} [C(H, \gamma_0(Z^1, \dots, Z^{i-1}, d, Z^{i+1}, \dots, Z^s), Z^1, \dots, Z^{i-1}, d, Z^{i+1}, \dots, Z^s) | H = j]$$

(b) Suppose that $\gamma_1, \gamma_2, \dots, \gamma_s$ are fixed. Then γ_0 is determined by

$$\gamma_0(Z^1, \dots, Z^s) = \arg \min_{d=1,2,\dots,M} \sum_{j=1}^M p(j|Z^1, \dots, Z^s) \times C(j, d, Z^1, \dots, Z^s)$$

2.3 Kernel Methods

Kernel methods provides a way to analyze data with higher dimensional feature space. We introduce kernel methods, since kernel methods are used as an important tool to form utility and constraint functions in this thesis. More detailed explanation can be found in [69].

We begin with Example 2.1 to show the necessity of mapping data to higher dimensional space in order to analyze more complex structure of data.

Example 2.1. *We have two variables in $\mathbb{R}^2$, marked with red and blue points in Figure 2.2 respectively. In this two dimensional space, it is impossible to separate the two variables with a linear boundary. However, if we map the variables to a three dimensional space, where the new variables becomes*

$$\phi(x) = [x_1, x_2, x_1x_2] \in \mathbb{R}^3,$$

the two variables become linearly separable as shown in Figure 2.2.

Instead of actually calculating the data, kernel methods only calculate the inner product of the high dimensional data, which dramatically reduce the computational complexity. Kernel trick has wide application in sequence data, text and image processing. To get a clearer picture of kernel method, we first introduce reproducing kernel.

Definition 2.1 (Hilbert space). *Hilbert space is a complete inner product space.*

Definition 2.2 (Reproducing kernel). *Let $\mathcal{H}$ be a hilbert space of real-valued functions. The function $\mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a reproducing kernel of $\mathcal{H}$ if it satisfies*

Figure 2.2: On the left, two variables are in two dimensional space, and they are not separable with a linear boundary. On the right, two variables are mapped to a three dimensional space, where there's a linear boundary to separate them.

- $\forall x \in \mathcal{X}$: it satisfies $k(\cdot, x) \in \mathcal{H}$,
- $\forall x \in \mathcal{X}$: $\langle f, k(\cdot, x) \rangle_{\mathcal{H}} = f(x)$.

Then we introduce the Reproducing Kernel Hilbert Space.

Definition 2.3 (Evaluation functional). *Let $\mathcal{H}$ be a Hilbert space of functions $f : \mathcal{X} \rightarrow \mathcal{H}$, defined on a non-empty set $\mathcal{X}$. For a fixed x , a evaluation functional $\delta_x : f \rightarrow f(x)$ is a function that evaluate each function in $\mathcal{H}$ at the point x .*

Definition 2.4 (Reproducing Kernel Hilbert Space). *Let $\mathcal{H}$ be a Hilbert space of functions $f : \mathcal{X} \rightarrow \mathcal{H}$, defined on a non-empty set $\mathcal{X}$. $\mathcal{H}$ is said to be a Reproducing Kernel Hilbert Space (RKHS) if for $\forall x \in \mathcal{X}$, δ_x is continuous.*

The reproducing kernel and RKHS is related with the following theorem.

Theorem 2.1. *$\mathcal{H}$ is a reproducing kernel if and only if $\mathcal{H}$ has a reproducing kernel.*

2.4 Empirical Risk Minimization

In supervised learning, given a set of training data $(x_i, y_i) \in (\mathcal{X}, \mathcal{Y}), i = 1, 2, \dots, n$, we would like to learn a function $f : \mathcal{X} \rightarrow \mathcal{Y}$. We aim at finding the optimal f to approximate y given x . Empirical Risk Minimization (ERM) is a statistical learning theory to achieve this goal. In this section, we briefly introduce ERM, which will be used in this thesis.

We first introduce loss function. A loss function is a function that maps an event or values of one or more variables onto a real number intuitively representing some “cost” associated with the event. Here, we use $L(f(x), y)$ to represent the loss function, which will yield a higher penalty cost if $f(x)$ is more different from y .

Here we list some commonly known loss functions.

- Classification loss:

$$L(f(x), y) = \begin{cases} 1, & f(x) \neq y \\ 0, & f(x) = y. \end{cases}$$

- Square loss:

$$L(f(x), y) = \frac{1}{2}(f(x) - y)^2.$$

- L_1 loss:

$$L(f(x), y) = |f(x) - y|^2.$$

The other loss functions with $\mathcal{Y} = \{-1, 1\}$, will be introduced in Chapter 3.

We denote the underlying joint distribution of (x, y) as $p_{X,Y}$, the true loss with

loss function L is

$$L(f(X), Y) = \int p_{X,Y}(x, y) L(f(x), y) dx dy.$$

However, in most applications, it is impossible to know $p_{X,Y}$, instead we have a series of training data, where we can calculate the empirical risk to approximate the true risk.

$$\hat{L}(f(X), Y) = \sum_{i=1}^n L(f(x_i), y_i).$$

ERM is a method to find the optimal f to minimize $\hat{L}(f(X), Y)$.

2.5 Information Theory

Information theory is used in hypothesis detection as an important tool to measure the correlation between random variables. Since the entropy and mutual information of random variables affects the performance of our proposed methods in this thesis dramatically, we briefly summarize the basic concepts in information theory. More detailed descriptions and proofs can be found in [70].

The entropy is a measure of uncertainty defined as follows.

Definition 2.5 (Entropy). *The entropy of $X \in \mathcal{X} \sim p$, $H(X)$ is defined as*

$$H(X) = - \int_{x \in \mathcal{X}} p(x) \log(p(x)) dx.$$

The relative entropy is defined as follows.

Definition 2.6 (KL-divergence). *The relative entropy or KL-divergence between probability distribution p and q is*

$$D(p||q) = \int_{x \in \mathcal{X}} p(x) \log\left(\frac{p(x)}{q(x)}\right).$$

KL-divergence measures the closeness of two probability distributions, it follows the following property.

Proposition 2.2. $D(p||q) \geq 0$. $D(p||q) = 0$ if and only if $p = q$.

Proof.

$$D(p||q) = \int_{x \in \mathcal{X}} p(x) \log\left(\frac{p(x)}{q(x)}\right) = - \int_{x \in \mathcal{X}} p(x) \log\left(\frac{q(x)}{p(x)}\right).$$

by Jensen's inequality

$$-\mathbb{E} \left[\log\left(\frac{q(x)}{p(x)}\right) \right] \geq -\log\left(\mathbb{E} \left[\frac{q(x)}{p(x)} \right]\right) = 0$$

with equality, if and only if $q(x) = cp(x), \forall x$, which is equivalent to $p = q$. $\square$

Finally, we introduce the mutual information.

Definition 2.7. *Two random variables, X and Y follow a joint probability distribution of $p_{X,Y}$, the marginal probability distribution is p_X and p_Y respectively. The mutual information $I(X;Y)$ is*

$$I(X;Y) = D(p_{X,Y}||p(X)p(Y)) = \mathbb{E}_{p_{X,Y}} \left[\log \frac{p_{X,Y}(X,Y)}{p_X(X)p_Y(Y)} \right]$$

2.6 Privacy Metrics

In this section, we consider different privacy metrics proposed in the literature. A privacy budget $\epsilon \geq 0$ is associated with each type of privacy metric, with a smaller ϵ corresponding to a more stringent privacy guarantee. We consider the following inference and data privacy metrics. Note that we use the joint distribution $p_{G,X,Z}$ in the definition of inference privacy and data metrics although they depend only on $p_{G,Z}$ while data privacy metrics depend only on $p_{X,Z}$. This is done to make it easier to present Definition 5.1, which allow us to relate inference and data privacy metrics.

Definition 2.8 (Inference privacy metrics). *Let $\epsilon \geq 0$, G, X, Z be random variables. We say that $p_{G,X,Z}$ satisfies each of the following types of inference privacy if the corresponding conditions hold.*

- ϵ -inference differential privacy [48]: for all $\mathbf{g}, \mathbf{g}' \in \mathcal{G}$ such that $\mathbf{g} \sim \mathbf{g}'$, and $\mathbf{z} \in \mathcal{Z}^s$,

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

- ϵ -average information leakage [48]: $I(G; Z) \leq \epsilon$.
- ϵ -information privacy [48]: for all $\mathbf{g} \in \mathcal{G}$ and $\mathbf{z} \in \mathcal{Z}^s$,

$$e^{-\epsilon} \leq \frac{p_{G|Z}(\mathbf{g}|\mathbf{z})}{p_G(\mathbf{g})} \leq e^\epsilon. \quad (2.1)$$

Note that the term “inference differential privacy” in Definition 2.8 is used to avoid confusion with “differential privacy”, which is usually associated with protecting the privacy of the data X . In Definition 2.8, the differential privacy refers to that for the private hypothesis G .

Definition 2.9 (Data privacy metrics). *Let $\epsilon \geq 0$, G, X, Z be random variables. We say that $p_{G,X,Z}$ satisfies each of the following types of data privacy if the corresponding conditions hold.*

- ϵ -local differential privacy [41]: for each sensor $t \in \{1, 2, \dots, s\}$, and all $x, x' \in \mathcal{X}$, and $z \in \mathcal{Z}$,

$$\frac{p^t(z|x)}{p^t(z|x')} \leq e^\epsilon.$$

- ϵ -mutual information privacy [47]: $I(X; Z) \leq \epsilon$.

- ϵ -identifiability [47]: for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^s$ such that $\mathbf{x} \sim \mathbf{x}'$, and $\mathbf{z} \in \mathcal{Z}^s$,

$$\frac{p_{X|Z}(\mathbf{x}|\mathbf{z})}{p_{X|Z}(\mathbf{x}'|\mathbf{z})} \leq e^\epsilon.$$

Chapter 3

Nonparametric Decentralized Detection with Information Privacy

3.1 Problem Formulation and Information Privacy

In this section, we first describe our system model, introduce 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 Internet of Things (IoT) network.

3.1.1 System Model

Consider the system model in Figure 1.1, our goal is to find a local decision rule or privacy mapping p^t for each sensor $t = 1, \dots, s$, 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 3.1.2, we show that by making the average of the Type I error $\mathbb{P}(\gamma_G(Z) = -1 \mid G = 1)$ and Type II error $\mathbb{P}(\gamma_G(Z) = 1 \mid G = -1)$ sufficiently large, we achieve *information privacy* [48] 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 independently identically distributed (i.i.d.) training data $(\mathbf{x}_i, h_i, g_i)_{i=1}^n$ sampled from $p_{X,H,G}$.¹ The framework of empirical risk minimization is adopted as in [60] to design the local sensor decision rules, which we discuss in detail in Section 3.2.1.

3.1.2 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 [48], which we recalled in Definition 2.8. Note that this differs from the more commonly used differential privacy concept [52, 71], which [48] shows is weaker than information privacy in general.

In an IoT network, both data and information privacy are important design aspects. Achieving data privacy has been comprehensively addressed in [30, 41, 72–74]. In this thesis, we focus on the information privacy aspect. In [48], 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 3.1. *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 3.1 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 ,

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.

let

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

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

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

and

$$c = \min \left\{ \mathbb{P} \left(\ell(\mathbf{z}) = \min_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) \mid G = -1 \right), \mathbb{P} \left(\ell(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) \mid G = 1 \right) \right\}. \quad (3.1)$$

From Assumption 3.1, 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 (2.1) 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 3.2.

Proposition 3.1. *Suppose that Assumption 3.1 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*

$$\mathcal{H}(\theta) \leq \mathcal{H}(G) - \epsilon, \quad (3.2)$$

with $\mathcal{H}(\cdot)$ being the binary entropy function [70].

- (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 3.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 2.8. An example is shown in Example 3.1, and a simulation to demonstrate this is given in Section 3.3.1. 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 3.1. Suppose $p_G(-1) < 1/2$, and G and $Z \in \{\mathbf{z}_1, \mathbf{z}_2\}$ have the joint distribution shown in Table 3.1. Then, for all sufficiently large d , the optimal Bayesian error probability detector declares $G^* = 1$ for all Z . The Bayes error $\mathbb{P}(G^* \neq G) = p_G(-1)$, 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 (2.1) for any $\epsilon \in [0, \infty)$. This shows that a large Bayesian error probability does not guarantee information privacy of G .

 Table 3.1: 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	$p_G(1) - p_G(-1)$	$p_G(-1)$

3.2 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.

3.2.1 Nonparametric Privacy-aware Optimization

From Proposition 3.1(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 :

$$\begin{aligned} \min_{\gamma^1, \dots, \gamma^s, \gamma_H} \quad & \mathbb{P}(\gamma_H(Z) \neq H) \\ \text{s.t.} \quad & \min_{\gamma_G} R(\gamma_G) \geq \theta, \end{aligned} \tag{3.3}$$

where θ is a privacy threshold that determines the information privacy budget for G . However, since (3.3) 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 [60], 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 thesis.

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

$$\min_{w_H \in \mathcal{H}, p_{Z|X} \in \mathcal{Q}} \frac{1}{n} \sum_{i=1}^n \sum_{\mathbf{z} \in \mathcal{Z}^s} \phi(h_i \langle w_H, \Phi(\mathbf{z}) \rangle) p_{Z|X}(\mathbf{z} | \mathbf{x}_i) + \frac{\lambda}{2} \|w_H\|^2, \quad (3.4a)$$

$$\text{s.t. } \min_{w_G \in \mathcal{H}} \frac{1}{2} \sum_{g \in \{-1, 1\}} \sum_{i \in \mathcal{S}_{g,n}} \sum_{\mathbf{z} \in \mathcal{Z}^s} \frac{\phi(g_i \langle w_G, \Phi(\mathbf{z}) \rangle) p_{Z|X}(\mathbf{z} | \mathbf{x}_i)}{|\mathcal{S}_{g,n}|} + \frac{\lambda_n}{2} \|w_G\|^2 \geq \theta. \quad (3.4b)$$

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, we can conclude that the generalization error can become large [69]. Note also that the empirical normalized risk in (3.4b) is different from the traditional empirical risk used in the machine learning literature (cf. (3.4a) and [69]), and serves as a surrogate for $R(\cdot)$ in Proposition 3.1.

The optimization problem (3.4) is intractable to compute because of the summation over $\mathcal{Z}^s$, therefore we employ the same lower bound relaxation as in (15) of [60], and let $\Phi_Q(\mathbf{x}) = \sum_{\mathbf{z} \in \mathcal{Z}^s} p_{Z|X}(\mathbf{z} | \mathbf{x}) \Phi(\mathbf{z})$. The corresponding kernel is $\kappa_Q(\mathbf{x}, \mathbf{x}') = \langle \Phi_Q(\mathbf{x}), \Phi_Q(\mathbf{x}') \rangle = \sum_{\mathbf{z}} \sum_{\mathbf{z}'} p_{Z|X}(\mathbf{z} | \mathbf{x}) p_{Z|X}(\mathbf{z}' | \mathbf{x}') \kappa(\mathbf{z}, \mathbf{z}')$. With this relaxation, the convex surrogate of the risk $R(\gamma_G)$ becomes

$$\begin{aligned} R_\phi(w_G, p_{Z|X}) &= \frac{1}{2} \mathbb{E}_X [\phi(-\langle w_G, \Phi_Q(X) \rangle) | G = -1] \\ &\quad + \frac{1}{2} \mathbb{E}_X [\phi(\langle w_G, \Phi_Q(X) \rangle) | G = 1], \end{aligned} \quad (3.5)$$

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

$$F(w_H, p_{Z|X}) = \frac{1}{n} \sum_{i=1}^n \phi(h_i \langle w_H, \Phi_Q(\mathbf{x}_i) \rangle) + \frac{\lambda}{2} \|w_H\|^2, \quad (3.6)$$

$$\hat{R}_n(w_G, p_{Z|X}) = \frac{1}{2} \sum_{g \in \{-1, 1\}} \sum_{i \in \mathcal{S}_{g,n}} \frac{\phi(g_i \langle w_G, \Phi_Q(\mathbf{x}_i) \rangle)}{|\mathcal{S}_{g,n}|} + \frac{\lambda_n}{2} \|w_G\|^2. \quad (3.7)$$

Our optimization problem (3.4) is then relaxed to

$$\min_{w_H \in \mathcal{H}, p_{Z|X} \in \mathcal{Q}} F(w_H, p_{Z|X}) \quad (3.8a)$$

$$\text{s.t.} \quad \min_{w_G \in \mathcal{H}} \hat{R}_n(w_G, p_{Z|X}) \geq \theta. \quad (3.8b)$$

For convenience, we call (3.8) the Nonparametric Privacy-aware Optimization (NPO). We note that (3.4b) holds if the constraint (3.8b) holds, since (3.8b) is a lower bound relaxation. In general, (3.8b) cannot guarantee information privacy with probability one since (3.8) 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 3.1 ((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 (3.8b). 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 [75].

Assumption 3.2 (Convex loss function). *The loss function ϕ is a convex function and $\phi(0) < \infty$.*

Assumption 3.3. *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 $\eta \in [0, 1]$,*

$$a^r (\phi(0) - R_\phi^*(\eta)) \geq \left| \frac{1}{2} - \eta \right|^r. \quad (3.9)$$

Examples of loss functions, with their corresponding a and r values in Assumption 3.3, are shown in Table 3.2. These include commonly used loss functions like the

Table 3.2: Loss functions, their conjugate duals, and parameters for Assumption 3.3

Loss function $\phi(u)$	Conjugate dual $\phi^*(-x^*)$	Assumption 3.3
e^{-u}	$\begin{cases} x^* \ln x^* - x^*, & x^* \in (0, \infty) \\ 0, & x^* = 0 \end{cases}$	$a = \frac{1}{\sqrt{2}}, r = 2$
$\log(1 + e^{-u})$	$\begin{cases} x^* \log x^* + (1 - x^*) \log(1 - x^*), & x^* \in (0, 1) \\ 0, & x^* = 0, 1 \end{cases}$	$a = \frac{1}{\sqrt{2}}, r = 2$
$\max\{1 - u, 0\}$	$-x^*, x^* \in [0, 1]$	$a = 1/2, r = 1$
$(1 - u)^2$	$(x^*)^2/4 - x^*, x^* \in \mathbb{R}$	$a = 1/2, r = 2$

logistic loss function [76], exponential loss function [77], hinge loss function [78], and quadratic loss function, which have been shown to be computationally efficient and have bounded approximation and estimation errors [75, 79]. In Table 3.2 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 $\eta(\mathbf{z}) = \mathbb{P}(\tilde{G} = 1 \mid Z = \mathbf{z})$, 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. (3.33)). Therefore, Assumption 3.3 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 (3.8b), we need the following assumption.

Assumption 3.4.

(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 [80]. With the above assumptions, the following result gives a theoretical guarantee of weak information privacy using the constraint (3.8b).

Theorem 3.1. *Suppose that Assumptions 3.1–3.4 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 $p_{Z|X} \in \mathcal{Q}$, if

$$\min_{w_G \in \mathcal{H}} \hat{R}_n(w_G, p_{Z|X}) \geq \theta,$$

where $\epsilon = \log \frac{c}{(c - 2a(\phi(0) - \theta + \delta)^{1/r})_+}$, with c as defined in (3.1).

Proof. See Appendix 3.B. □

From Theorem 3.1, we see that our NPO (3.8) 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.

3.2.2 Iterative Optimization

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

$$\begin{aligned} \hat{R}^*(\boldsymbol{\beta}, p_{Z|X}) = & - \sum_{g \in \{-1, 1\}} \sum_{i \in \mathcal{S}_{g,n}} \frac{\phi^*(-2|\mathcal{S}_{g,n}|\beta_i)}{2|\mathcal{S}_{g,n}|} \\ & - \frac{1}{2\lambda_n} \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j g_i g_j \kappa_Q(\mathbf{x}_i, \mathbf{x}_j), \end{aligned} \quad (3.10)$$

where ϕ^* is the conjugate dual of ϕ (see Table 3.2 for some examples) [81]. From Proposition 2 in [60], for any $p_{Z|X}$, we have

$$\min_{w_G \in \mathcal{H}} \hat{R}_n(w_G, p_{Z|X}) = \max_{\boldsymbol{\beta} \in \mathbb{R}^n} \hat{R}^*(\boldsymbol{\beta}, p_{Z|X}). \quad (3.11)$$

The optimal primal variable w_G and the optimal dual variable $\boldsymbol{\beta}$ in (3.11) are related by $w_G = \sum_{i=1}^n \beta_i g_i \Phi_Q(\mathbf{x}_i)$. The constraint (3.8b) can now be rewritten as $\max_{\boldsymbol{\beta} \in \mathbb{R}^n} \hat{R}^*(\boldsymbol{\beta}, p_{Z|X}) \geq \theta$. By using the interior-point method with log barrier [82],

we transform the optimization problem (3.8) into the following:

$$\min_{w_H \in \mathcal{H}, \beta \in \mathbb{R}^n, p_{Z|X} \in \mathcal{Q}} F_0(w_H, \beta, p_{Z|X}), \quad (3.12)$$

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

We propose a two-step algorithm to solve (3.12). 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 $p_{Z|X}$. 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 (3.12). We call p the privacy threshold ratio. In both steps, we apply a block nonlinear Gauss-Siedel method (GS) defined as follows.

Definition 3.2 (block nonlinear Gauss-Siedel method [83]). *For optimization*

$$\begin{aligned} & \min_x f(x) \\ & \text{s.t. } x \in \mathcal{X} = \mathcal{X}^1 \times \mathcal{X}^2 \times \dots \times \mathcal{X}^s, \end{aligned}$$

where $f(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuous differentiable function. $\mathcal{X}$ is the Cartesian product of closed and convex set $\mathcal{X}^t = \mathbb{R}^{m_t}$, $t = 1, 2, \dots, s$, where $\sum_{t=1}^s m_t = n$.

Block nonlinear Gauss-Siedel method is defined by the iteration

$$x_i^{k+1} = \arg \min_{y \in \mathcal{X}^i} f(x_1^{k+1}, x_2^{k+1}, \dots, y, \dots, x_s^k)$$

Our algorithm will be described in detail as follows.

Finding the privacy threshold

In Algorithm 3.1, we apply GS to solve

$$\max_{\beta \in \mathbb{R}^n, p_{Z|X} \in \mathcal{Q}'} \hat{R}^*(\beta, p_{Z|X}), \quad (3.13)$$

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

$$\sum_{x \in \mathcal{X}} p^t(z | x) \geq \Delta_1, \text{ for all } z \in \mathcal{Z}, \quad (3.14)$$

$$\left| p^t(z | x) - \frac{1}{|\mathcal{Z}|} \right| \geq \Delta_2, \text{ for all } z \in \mathcal{Z} \text{ and } x \in \mathcal{X}, \quad (3.15)$$

with Δ_1 and Δ_2 being small positive constants. The constraint (3.14) ensures that no $z \in \mathcal{Z}$ has small probability (otherwise we could have reduced the size of $\mathcal{Z}$), and (3.15) prevents equal probabilities from being assigned to all possible z . Since the optimization problem in (3.13) is non-convex, there is no guarantee that GS converges to the global optimum [83–85]. However, we can show the following convergence result.

Proposition 3.2. *Algorithm 3.1 converges to a critical point.*

Proof. See Appendix 3.C □

The objective value output from Algorithm 3.1 is denoted as θ^* , and we use its solution $(\tilde{\beta}, \widetilde{p_{Z|X}})$ as the initialization point in the second step.

Algorithm 3.1 Finding θ^*

- 1: **input:** $\{g_i, x_i^1, \dots, x_i^s\}_{i=1}^n$
- 2: **initialization:** $\beta[0] \in \mathbb{R}^n, p_{Z|X}[0] \in \mathcal{Q}', k = 0$,
- 3: **repeat**
- 4: • $k = k + 1$,

- fix $Q[k-1]$, solve the following convex optimization problem,

$$\beta[k] = \arg \max_{\beta \in \mathbb{R}^n} \hat{R}^*(\beta, p_{Z|X}[k-1]),$$

- fix $\beta[k]$, $p^j[k]$, for $j < t$, and $p^j[k-1]$, for $j > t$, $t = 1, 2, \dots, s$ update

$$p^t[k] = \arg \max_{p^t \in \mathcal{Q}'} \hat{R}^*(\beta[k], p^1[k], \dots, p^{t-1}[k], p^t, p^{t+1}[k-1], \dots, p^s[k-1]),$$

- 5: **until** $\frac{\hat{R}^*(\beta[k], p_{Z|X}[k]) - \hat{R}^*(\beta[k-1], p_{Z|X}[k-1])}{\hat{R}^*(\beta[k-1], p_{Z|X}[k-1])} \leq \epsilon$,

- 6: **return** $\tilde{\beta} = \beta[k], \widetilde{p_{Z|X}} = p_{Z|X}[k], \theta^* = \hat{R}^*(\beta[k], p_{Z|X}[k])$.
-

Finding sensor privacy mappings

With θ^* from Algorithm 3.1, we set $\theta = p\theta^*$, where $p \in (0, 1)$ is a constant typically chosen to be close to 1 (see Section 3.3.1). We minimize F_0 over $(w_H, \beta, p_{Z|X})$ using GS, as shown in Algorithm 3.2.

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

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

We have

$$\min_{w_H \in \mathcal{H}} F(w_H, p_{Z|X}) = \max_{\alpha \in \mathbb{R}^n} F^*(\alpha, p_{Z|X}), \quad (3.16)$$

where the optimal primal variable w_H and the optimal dual variable α in (3.16) are related by

$$w_H = \sum_{i=1}^n \alpha_i h_i \Phi_Q(\mathbf{x}_i). \quad (3.17)$$

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 (3.17), we turn the feature maps into the kernel in the primal space. Therefore, when it comes to updating β and p^t , we do the minimization in the primal space.

Proposition 3.3. *Algorithm 3.2 converges to a critical point.*

Proof. See Appendix 3.D. □

3.2.3 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

Algorithm 3.2 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:** $w[0] \in \mathbb{R}^n, \beta[0] \leftarrow \tilde{\beta}, p_{Z|X}[0] \leftarrow p_{\tilde{Z}|X}, k = 0,$
- 3: **repeat**
- 4: • $k = k + 1,$
- Fix $\beta[k-1]$ and $p_{Z|X}[k-1]$, and solve the following convex optimization problem,

$$\alpha[k] = \arg \max_{\alpha \in \mathbb{R}^n} \{F^*(\alpha, p_{Z|X}[k-1])\},$$

and we obtain $w_H[k] = \sum_{i=1}^n \alpha_i[k] h_i \Phi_Q(\mathbf{x})$,

- Fix $w_H[k]$ and $p_{Z|X}[k-1]$, update

$$\beta[k] = \arg \max_{\beta \in \mathbb{R}^n} \frac{1}{\mu} \log \left(\hat{R}^*(\beta, p_{Z|X}[k-1]) - \theta \right),$$

- Fix $w_H[k], \beta[k]$, and $p^1[k], \dots, p^{t-1}[k], p^{t+1}[k-1], p^s[k-1]$, update

$$p^t[k] = \arg \min_{p^t \in \mathcal{Q}} F_0(w_H[k], p^1[k], \dots, p^{t-1}[k], p^t, p^{t+1}[k-1], p^s[k-1], \beta[k]),$$

- 5: **until** $\frac{F_0(w_H[k-1], p_{Z|X}[k-1], \beta[k-1]) - F_0(w_H[k], p_{Z|X}[k], \beta[k])}{F_0(w_H[k-1], p_{Z|X}[k-1], \beta[k-1])} \leq \epsilon,$
 - 6: **output:** $\alpha[k], \beta[k], p_{Z|X}[k].$
-

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 [86] 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

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

and

$$\ell_g(\mathbf{z}) = \frac{p_{Z|G}(\mathbf{z}|g)}{p_{Z|G}(\mathbf{z}|0)}, \quad (3.18)$$

for $\mathbf{z} \in \mathcal{Z}^s$. The following result follows from Proposition 3.1(ii).

Theorem 3.2. *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*

$$c' = \min_{1 \leq g < m} \left\{ \mathbb{P} \left(\ell_g(\mathbf{z}) = \min_{\mathbf{z} \in \mathcal{D}} \ell_g(\mathbf{z}) \mid G = 0 \right), \mathbb{P} \left(\ell_g(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell_g(\mathbf{z}) \mid G = g \right) \right\} > 0.$$

Proof. See Appendix 3.E. □

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

$$\min_{w_g \in \mathcal{H}} \hat{R}_{g,n}(w_g, p_{Z|X}) \geq \theta, \text{ for } g = 1, \dots, m-1, \quad (3.19)$$

where

$$\hat{R}_{g,n}(w, p_{Z|X}) = \frac{1}{2} \sum_{g' \in \{0, g\}} \sum_{i \in \mathcal{S}_{g',n}} \frac{\phi(g'_i \langle w, \Phi_Q(\mathbf{x}_i) \rangle)}{|\mathcal{S}_{g',n}|} + \frac{\lambda_n}{2} \|w\|^2, \quad (3.20)$$

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 (3.18) are sufficient statistics.

Indeed, it can be shown, using the same arguments as in Theorem 3.1 that the following holds (we omit the proof here):

Theorem 3.3. *Suppose the corresponding assumptions as in Theorem 3.1 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 $p_{Z|X} \in \mathcal{Q}$ satisfying (3.19), where $\epsilon = \frac{1}{2} \log \frac{c'}{(c' - 2a(\phi(0) - \theta + \delta)^{1/r})_+}$, with c' as defined in Theorem 3.2.*

We apply Algorithm 3.1 to each of the $m - 1$ constraints in (3.19) 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 3.2, the objective function is replaced by

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

where $\hat{R}_g^*(\beta_g, p_{Z|X})$ is the dual form of $\hat{R}_{g,n}(w_g, p_{Z|X})$ (see (3.10)). It can be shown that Proposition 3.3 still holds.

3.3 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 [87], and compare its performance with RUCA [88] and MDR [89].

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 $p_{Z|X} \in \mathcal{Q}$, we have

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

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 3.4(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 [90], 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 3.1 and 3.2.

The complexity of both Algorithms 3.1 and 3.2 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 3.1 throughout.

3.3.1 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.

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 3.3, where n_i^t is distributed uniformly over $\{-1, 0, +1\}$. In this set 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.

Table 3.3: Sensor observations X for different realizations of (H, G)

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

In Fig. 3.1, 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 3.1. 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 (3.3), 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. 3.2, 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

Figure 3.1: Error rate of detecting H and G with different correlation coefficient between H and G .

Algorithm 3.2 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 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 .

Comparison with the empirical Bayes error privacy metric

Example 3.1 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 [91] 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

Figure 3.2: Error rate of detecting H and G with varying privacy threshold ratio p .

G . With the empirical risk corresponding to the Bayes error, our optimization problem becomes:

$$\begin{aligned}
 & \min_{w_H \in \mathcal{H}, p_{Z|X} \in \mathcal{Q}} F(w_H, p_{Z|X}) \\
 \text{s.t. } & \min_{w_G \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \phi(g_i \langle w_G, \Phi_Q(\mathbf{x}_i) \rangle) + \frac{\lambda}{2} \|w_G\|^2 \geq \theta.
 \end{aligned} \tag{3.21}$$

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 3.1 and 3.2 to find a privacy threshold and the sensor privacy mappings for (3.21). The results are shown in Table 3.4. We see that the Bayes errors of detecting H and G respectively are similar for both the NPO approach and (3.21). However, NPO achieves a much smaller privacy budget ϵ (i.e., stronger information privacy) than (3.21).

Table 3.4: Comparison of NPO with (3.21)

	Bayes Error H	Bayes Error G	$R(\gamma_G)$	Information privacy budget, ϵ
NPO	0.03	0.05	0.38	1.08
(3.21)	0.02	0.05	0.24	3.91

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 [60].

We see from Table 3.5 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.

Table 3.5: 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

3.3.2 UCI Repository Data Sets

We test our algorithm on the OPPORTUNITY Activity Recognition Data Set [92] and Census (Adult) Data Set (both available at UCI Repository [87]), and compare its performance with RUCA [88] and MDR [89]. In our NPO, we set the local decision space of each sensor to be $\mathcal{Z} = \{1, 2\}$.

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 [93] 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 [88], we group ‘Married-civ-spouse’, ‘Married-spouse-absent’ and ‘Married-AF-spouse’ 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.

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 [88] and MDR [89], which are nonparametric approaches that require a trusted data curator to first aggregate data from all sensors before applying an overall privacy mapping. Parameter settings for RUCA and MDR are as proposed in [88] and [89], 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 3.6 and Table 3.7, 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.

Table 3.6: 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%	1.02
DCA	10.88%	42.62%	0.88
RUCA ($\rho_p = 1$)	10.23%	45.73%	0.67
RUCA ($\rho_p = 100$)	10.10%	43.01%	0.69
RUCA ($\rho_p = 1000$)	10.10%	43.78%	0.69
$Z = X$	10.05%	5.57%	9.14

Table 3.7: 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%	1.68
DCA	38.38%	56.33%	2.47
RUCA ($\rho_p = 1$)	41.24%	65.25%	1.61
RUCA ($\rho_p = 100$)	41.10%	64.14%	1.66
RUCA ($\rho_p = 1000$)	40.67%	65.86%	1.61
$Z = X$	34.05%	30.48%	15.33

3.A Proof of Proposition 3.1

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 [70], we have

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

where the last inequality follows from Definition 2.8. 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 (3.2). 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

$$R(\gamma_G) = \frac{1}{2} - \frac{1}{2} \sum_{\mathbf{z} \in \Gamma} (p_{Z|G}(\mathbf{z} \mid 1) - p_{Z|G}(\mathbf{z} \mid -1)) \quad (3.22)$$

$$= \frac{1}{2} + \frac{1}{2} \sum_{\mathbf{z} \in \Gamma^c} (p_{Z|G}(\mathbf{z} \mid 1) - p_{Z|G}(\mathbf{z} \mid -1)). \quad (3.23)$$

For any $\mathbf{z}' \in \Gamma$, if $R(\gamma_G) \geq \theta$, we have from (3.22),

$$\begin{aligned}
 1 - 2\theta &\geq \sum_{\mathbf{z} \in \Gamma} (p_{Z|G}(\mathbf{z} | 1) - p_{Z|G}(\mathbf{z} | -1)) \\
 &= \sum_{\mathbf{z} \in \Gamma} p_{Z|G}(\mathbf{z} | 1) \left(1 - \frac{1}{\ell(\mathbf{z})}\right) \\
 &\geq \left(1 - \frac{1}{\ell(\mathbf{z}')} \right) \sum_{\{\mathbf{z}: \ell(\mathbf{z}) \geq \ell(\mathbf{z}')\}} p_{Z|G}(\mathbf{z} | 1) \\
 &\geq \left(1 - \frac{1}{\ell(\mathbf{z}')} \right) \mathbb{P} \left(\ell(\mathbf{z}) = \max_{\mathbf{z} \in \mathcal{D}} \ell(\mathbf{z}) \mid G = 1 \right) \\
 &\geq c \left(1 - \frac{1}{\ell(\mathbf{z}')} \right). \tag{3.24}
 \end{aligned}$$

From (3.24), we obtain

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

and

$$1 \leq \ell(\mathbf{z}') \leq \frac{1}{(1 - (1 - 2\theta)/c)_+} = e^\epsilon. \tag{3.25}$$

Similarly, for any $\mathbf{z}' \in \Gamma^c$, a similar argument using (3.23) yields

$$e^{-\epsilon} \leq \ell(\mathbf{z}') \leq 1. \tag{3.26}$$

Combining (3.25) and (3.26), we have for any $\mathbf{z} \in \mathcal{D}$,

$$e^{-\epsilon} \leq \ell(\mathbf{z}) \leq e^\epsilon, \tag{3.27}$$

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

$$\begin{aligned}
 e^{-\epsilon} p_{Z|G}(\mathbf{z} | g) &\leq p_Z(\mathbf{z}) \\
 &= \mathbb{E}_G [p_{Z|G}(\mathbf{z} | G)] \leq e^\epsilon p_{Z|G}(\mathbf{z} | g). \tag{3.28}
 \end{aligned}$$

We then obtain (2.1) 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.

3.B Proof of Theorem 3.1

We start with the following elementary lemma.

Lemma 3.B.1. *For any $w \in \mathcal{H}$,*

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

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 (3.7) and (3.5). □

We now prove Theorem 3.1. Let $\hat{w}_{G,n} = \arg \min_{w \in \mathcal{H}} \hat{R}_n(w, p_{Z|X})$. From Markov's inequality, we have

$$\begin{aligned}
 & \mathbb{P} \left(\hat{R}_n(\hat{w}_{G,n}, p_{Z|X}) \geq \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) + \delta \right) \\
 & \leq \frac{1}{\delta} \left(\mathbb{E} \left[\hat{R}_n(\hat{w}_{G,n}, p_{Z|X}) \right] - \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) \right) \\
 & = \frac{1}{\delta} \left(\mathbb{E} \left[\min_{w \in \mathcal{H}} \hat{R}_n(w, p_{Z|X}) \right] - \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) \right) \\
 & \leq \frac{1}{\delta} \left(\inf_{w \in \mathcal{H}} \mathbb{E} \left[\hat{R}_n(w, p_{Z|X}) \right] - \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) \right) \\
 & = \frac{1}{\delta} \left(\inf_{w \in \mathcal{H}} \{ R_\phi(w, p_{Z|X}) + \frac{\lambda_n}{2} \|w\|^2 \} - \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) \right), \tag{3.29}
 \end{aligned}$$

where the last equality follows from Lemma 3.B.1. If $n \rightarrow \infty$ in (3.29), its right hand side approaches to 0 uniformly in $p_{Z|X}$ because of Assumption 3.4(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 $1 - \delta$,

$$\hat{R}_n(\hat{w}_{G,n}, p_{Z|X}) \leq \inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) + \delta. \tag{3.30}$$

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

$$\begin{aligned}
 R_\phi(w, p_{Z|X}) &= \mathbb{E} \left[\phi \left(\tilde{G} \langle w, \Phi_Q(X) \rangle \right) \right] \\
 &= \mathbb{E} \left[\phi \left(\tilde{G} \langle w, \mathbb{E}_Z [\Phi(Z) | X] \rangle \right) \right] \\
 &\leq \mathbb{E} \left[\mathbb{E}_Z \left[\phi \left(\tilde{G} \langle w, \Phi(Z) \rangle \right) \mid X \right] \right] \\
 &= \mathbb{E} \left[\phi \left(\tilde{G} \cdot \gamma(Z) \right) \right] \\
 &= \mathbb{E} [(\eta(Z)\phi(\gamma(Z)) + (1 - \eta(Z))\phi(-\gamma(Z)))], \tag{3.31}
 \end{aligned}$$

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 $\eta(\mathbf{z}) = \mathbb{P}(\tilde{G} = 1 \mid Z = \mathbf{z})$. 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 3.4(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 (3.31), we obtain

$$\inf_{w \in \mathcal{H}} R_\phi(w, p_{Z|X}) \leq \mathbb{E}_Z [R_\phi^*(\eta(Z))]. \quad (3.32)$$

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

$$\begin{aligned} & \min_{\gamma} R(\gamma) \\ &= \min_{\gamma} \mathbb{E} [(1 - \eta(Z)) \mathbf{1}_{\{\gamma(Z) \geq 0\}} + \eta(Z) \mathbf{1}_{\{\gamma(Z) < 0\}}] \\ &= \frac{1}{2} - \mathbb{E} \left[\left| \frac{1}{2} - \eta(Z) \right| \right] \\ &= \mathbb{E} [R^*(\eta(Z))]. \end{aligned} \quad (3.33)$$

From Assumption 3.3, we obtain

$$\begin{aligned} \mathbb{E} [R_\phi^*(\eta(Z))] &\leq \phi(0) - \frac{1}{a^r} \mathbb{E} \left[\left(\frac{1}{2} - R^*(\eta(Z)) \right)^r \right] \\ &\leq \phi(0) - \frac{1}{a^r} \left(\frac{1}{2} - \mathbb{E} [R^*(\eta(Z))] \right)^r, \end{aligned} \quad (3.34)$$

where the last inequality follows from Jensen's inequality. Substituting (3.33) into (3.34), we have

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

where the last inequality follows from (3.30) and (3.32), and holds with probability $1 - \delta$. The theorem then follows from Proposition 3.1, and the proof is complete.

3.C Proof of Proposition 3.2

We first show that $-\hat{R}^*(\boldsymbol{\beta}, p_{Z|X})$ is a strictly quasi-convex function of $p_{Z|X}$ if $\boldsymbol{\beta} \neq \mathbf{0}$.

The definition of strictly quasi-convexity is borrowed from [83] as follows.

Definition 3.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 3.C.1. *Suppose that $\boldsymbol{\beta} \neq \mathbf{0}$. Then, $-\hat{R}^*(\boldsymbol{\beta}, p_{Z|X})$ is a strictly quasi-convex function of p^t , $t = 1, 2, \dots, s$, for $p_{Z|X} \in \mathcal{Q}'$.*

Proof. For each $p_{Z|X} \in \mathcal{Q}'$, consider

$$\begin{aligned} & \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j g_i g_j \kappa_Q(\mathbf{x}_i, \mathbf{x}_j) \\ &= \sum_{i,k} \sum_{j,k'} p_{Z|X}(\mathbf{z}_k | \mathbf{x}_i) p_{Z|X}(\mathbf{z}_{k'} | \mathbf{x}_j) \beta_i \beta_j g_i g_j \kappa(\mathbf{z}_k, \mathbf{z}_{k'}) \\ &= \left\| \sum_{i,k} \beta_i g_i p_{Z|X}(\mathbf{z}_k | \mathbf{x}_i) \Phi(\mathbf{z}_k) \right\|^2 > 0, \end{aligned} \tag{3.35}$$

where the last inequality holds due to Assumption 3.4(ii) since $\boldsymbol{\beta} \neq \mathbf{0}$, and $p^t \neq 0$ for all $t = 1, \dots, s$ because of (3.14). From (3.10) and (3.35), $-\hat{R}^*(\boldsymbol{\beta}, p_{Z|X})$ is a positive definite quadratic form of p^t , for each $t = 1, 2, \dots, s$, and is thus a strictly convex function of p^t . This implies strict quasi-convexity, and the proof of the lemma is complete. $\square$

We now show that $\{\boldsymbol{\beta}[k]\}$ and $\{p_{Z|X}[k]\}$ in Algorithm 3.1 have limit points. Since $\phi^*(\cdot) \geq -\phi(0) > -\infty$, $-\hat{R}^*(\boldsymbol{\beta}, p_{Z|X})$ is lower bounded. Therefore, the decreasing sequence $F[k] = -\hat{R}^*(\boldsymbol{\beta}[k], p_{Z|X}[k])$ converges [94]. From Proposition 4 of [83], convergence of $\{F[k]\}$ implies convergence of $\{\boldsymbol{\beta}[k]\}$ and $\{p_{Z|X}[k]\}$ to

limit points. By Proposition 5 of [83], this shows that Algorithm 3.1 converges to a critical point. The proof of the proposition is now complete.

3.D Proof of Proposition 3.3

From [60], $F(w_H, p_{Z|X})$ is a convex function of p^t , for each $t = 1, 2, \dots, s$. Since $-\log(\cdot)$ is a non-increasing convex function, and $\hat{R}^*(\boldsymbol{\beta}, p_{Z|X}) - \theta$ is a strictly concave function of p^t , $-\frac{1}{\mu} \log(\hat{R}^*(\boldsymbol{\beta}, p_{Z|X}) - \theta)$ is strictly convex with respect to p^t [82]. Therefore, $F_0(w_H, \boldsymbol{\beta}, p_{Z|X})$ is a strictly quasi-convex function of p^t , $t = 1, \dots, s$ (see Definition 3.C.1). The rest of the proof is similar to that in Appendix 3.C.

3.E Proof of Theorem 3.2

Following the same argument as that in the proof of Proposition 3.1(ii), for any $1 \leq g < m$ and $\mathbf{z} \in \mathcal{D}'$, we have from (3.27),

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

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 3.1(ii).

Chapter 4

Parametric Decentralized Detection with Robust Information Privacy Protection

4.1 System Model and Error Bounds

We consider a decentralized detection network as shown in Figure 1.1. The hypothesis $H \in \{-1, 1\}$ is the public hypothesis that the fusion center is authorized to infer, while $G^\circ \in \{-1, 1\}$ represents a private hypothesis that we wish to protect.

Given p_{X,H,G° , we define the utility to be the negative of the Bayes error $\mathbb{P}(\gamma_H(Z) \neq H)$ in inferring H so that maximizing the utility over $\mathcal{Q}$ means minimizing the Bayes error. Our goal is to find a $p_{Z|X} \in \mathcal{Q}$ to achieve an optimal trade-off between the utility of inferring H and protecting the privacy of G° . However, in practice, it may be hard for a user to specify a particular G° to protect. Furthermore, there could be a model mismatch. For example, G° may correspond to a particular sensitive phenomenon but during the actual operation of the sensor network, a related but different phenomenon is realized instead. Therefore, instead of considering only G° , we consider the privacy protection of all binary hypotheses

G that are close to G° in the following sense. We call G° the *nominal* private hypothesis.

Definition 4.1. For a given $p_{X|G^\circ}$ and an uncertainty parameter $\delta \in [0, 1)$, the uncertainty set $\mathcal{G}(p_{X|G^\circ}, \delta)$ of private hypotheses with respect to (w.r.t.) the nominal hypothesis G° is given by

$$\begin{aligned} \mathcal{G}(p_{X|G^\circ}, \delta) = \left\{ G : \min_{g \in \{-1, 1\}} p_G(g) > 0, \right. \\ p_{X|G}(\mathbf{x} | g) = (1 - \delta)p_{X|G^\circ}(\mathbf{x} | g) + \delta f_g(\mathbf{x}), \\ \left. f_g \in \mathcal{S}_{\mathcal{X}^s}, g \in \{-1, 1\}, \mathbf{x} \in \mathcal{X}^s \right\}, \end{aligned} \quad (4.1)$$

where $\mathcal{S}_{\mathcal{X}^s}$ is the set of probability mass functions over $\mathcal{X}^s$.

We assume that $G^\circ \in \mathcal{G}(p_{X|G^\circ}, \delta)$. We note that $\mathcal{G}(p_{X|G^\circ}, \delta)$ is an uncountable set if $\delta > 0$. Abusing notation, to avoid clutter, we sometimes denote $\mathcal{G}(p_{X|G^\circ}, \delta)$ as $\mathcal{G}_X$ if δ and G° are clear from the context.

We also assume for some $\alpha > 0$,

$$p_{X|G^\circ}(\mathbf{x}|g) \geq \alpha, \forall \mathbf{x} \in \mathcal{X}^s \text{ and } g \in \{-1, 1\}. \quad (4.2)$$

The assumption in (4.2) is to avoid the trivial case where the nominal private hypothesis G° is perfectly detectable for some observations $\mathbf{x}$.

In the example of fall detection using on-body wearables described in Section 1.2.1, the public hypothesis H is whether the person has fallen or not. The raw sensor information X from the on-body wearables may be able to let the fusion center infer whether the person is performing a specific exercise like running, doing sit-ups, weight training, climbing stairs, etc. Suppose we define G° to be the hypothesis that the person is running at a particular constant speed v . Then, the uncertainty set $\mathcal{G}(p_{X|G^\circ}, \delta)$ consists of exercise hypotheses that are related to running from the sensors' perspective. To be more specific, depending on the value of δ chosen, the uncertainty set $\mathcal{G}(p_{X|G^\circ}, \delta)$ can include hypotheses like the person is running at an average speed close to v , the person is running up a flight of stairs

at a speed similar to v , or the person is performing jumping jacks at a rate that generates similar but “noisier” sensor readings X than running at speed v .

The motivation for our Definition 4.1 is to protect all hypotheses that generate similar sensor observations as the nominal hypothesis G° . In this chapter, for simplicity, we consider only a single private hypothesis uncertainty set, although our framework can be extended to include multiple uncertainty sets defined around multiple or M -ary nominal private hypotheses (see Remark 4.5 later).

Our goal is to prevent statistical inference of all $G \in \mathcal{G}_X$. We choose information privacy [48] as the privacy metric because [45] has shown that the information privacy metric is one of the strongest privacy metrics that protects against statistical inference. The definition of information privacy is presented in Definition 2.8.

The definition of information privacy essentially says that the received information Z at the fusion center does not provide much information about G in addition to what the fusion center already knows about it. Therefore, it is difficult to statistically infer G given only Z . Definition 2.8 is similar to but different from the maximal realizable leakage defined in [95, Definition 4]. The maximal realizable leakage is shown to be equivalent to $\max_{\mathbf{x}, \mathbf{z}} p_{Z|X}(\mathbf{z}|\mathbf{x})/p_Z(\mathbf{z})$, which is related to the local differential privacy of the sensor data X [95] and protects all possible hypotheses including H . This privacy notion is too strong for our purposes as we are interested in only protecting the privacy of those hypotheses closely related to the nominal private hypothesis G° , while allowing the inference of H .

Our optimization problem is formulated as follows: for a privacy budget $\epsilon > 0$, we aim to

$$\min_{\gamma_H, p_{Z|X} \in \mathcal{Q}} \mathbb{P}(\gamma_H(Z) \neq H) \quad (4.3a)$$

$$\text{s.t. } e^{-\epsilon} \leq \frac{p_{Z|G}(\mathbf{z} | g)}{p_Z(\mathbf{z})} \leq e^\epsilon, \quad (4.3b)$$

$$\forall G \in \mathcal{G}(p_{X|G^\circ}, \delta), g \in \{-1, 1\}, \mathbf{z} \in \mathcal{Z}^s,$$

where we have made use of the equality $p_{G|Z}(g | \mathbf{z})/p_G(g) = p_{Z|G}(\mathbf{z} | g)/p_Z(\mathbf{z})$.

Note that here, $\mathbb{P}(\gamma_H(Z) \neq H)$ is the Bayes error w.r.t. the probability distributions p_{X,H,G° and $p_{Z|X}$.

Remark 1. Definition 4.1 utilizes the concept of a contamination uncertainty class from the robust detection literature [62, 64, 96]. A more general definition of an uncertainty class with desirable properties (e.g., the existence of a pair of least favorable distributions (LFDs) for the two uncertainty classes $\{p_{X|G}(\cdot | 0)\}$ and $\{p_{X|G}(\cdot | 1)\}$) is the set of distributions upper bounded by a 2-alternating capacity [63], which includes as special cases the contamination uncertainty set used in Definition 4.1, and the uncertainty class consisting of G such that $p_{X|G}$ is within a fixed total variation distance of $p_{X|G^\circ}$. Generalizing Definition 4.1 to utilize 2-alternating capacities is out of the scope of this chapter, and is part of our future work. Using contamination uncertainty class is also easier to interpret in practice and leads to the following natural result in Proposition 4.1.

Remark 2. The objective function in (4.3a) can be replaced by

$$E(\gamma_H, p_{Z|X})$$

where E is a general loss function, examples of which are the following. Let $\mathbb{P}(\gamma(Z) \neq H; p_{X,H,G})$ denote the Bayes error for inferring H w.r.t. a given $p_{X,H,G}$, $p_{Z|X}$ and γ .

- (a) The formulation in (4.3a) uses $E = \mathbb{P}(\gamma(Z) \neq H; p_{X,H,G^\circ})$ where the Bayes error is w.r.t. a given p_{X,H,G° .
- (b) If a prior distribution over $G \in \mathcal{G}(p_{X|G^\circ}, \delta)$ is known, a Bayesian formulation with $E = \mathbb{E}_G[\mathbb{P}(\gamma(Z) \neq H; p_{X,H,G})]$, where the expectation is taken over all the private hypotheses $G \in \mathcal{G}(p_{X|G^\circ}, \delta)$ can be used.
- (c) A minimax or robust formulation involves assuming that each $p_{X,H,G}$ belongs

to an uncertainty class

$$\mathcal{P} = \left\{ p_{X,H,G} = p_{G^\circ} \cdot p_{X|G} \cdot p_{H|X,G^\circ} : \right. \\ \left. G \in \mathcal{G}(p_{X|G^\circ}, \delta) \right\}$$

and letting

$$E = \max_{p_{X,H,G} \in \mathcal{P}} \mathbb{P}(\gamma_H(Z) \neq H; p_{X,H,G}).$$

It is easy to see that for any $p_{X,H,G} \in \mathcal{P}$, $\mathbf{x} \in \mathcal{X}^s$, $h \in \{-1, 1\}$,

$$p_{X|H}(\mathbf{x}|h) \\ \propto (1 - \delta) \sum_g p_{X,H,G^\circ}(\mathbf{x}, h, g) + \delta f(\mathbf{x}, h), \quad (4.4)$$

where $f(\mathbf{x}, h)$ is some joint probability mass function for (X, H) . Then minimizing E is equivalent to the robust detection of H in [62] over the contamination uncertainty classes defined by the class of distributions $\{p_{X|H}(\cdot | -1)\}$ and $\{p_{X|H}(\cdot | 1)\}$ of the form given in (4.4) as $f(\mathbf{x}, h)$ varies over all possible distributions. The loss function E can then be simplified by finding the pair of Least Favorable Distributions (LFD)s of these two classes, which correspond to a G' so that $E = \mathbb{P}(\gamma(Z) \neq H; p_{X,H,G'})$, which is similar to case (a).

Since the focus of this chapter is on robust information privacy, we use case (a) for illustrative purposes throughout this chapter.

Proposition 4.1. *Let $g_* = \arg \min_{g=-1,1} p_{G^\circ}(g)$ and*

$$\mathcal{F}_\pi = \{G : p_{G|G^\circ}(-1 | 1) \leq \pi, p_{G|G^\circ}(1 | -1) \leq \pi\},$$

where

$$\pi < \min\{p_{G^\circ}(g_*), \min_{\mathbf{x}, g} p_{X|G^\circ}(\mathbf{x} | g)\}.$$

Then $\mathcal{F}_\pi \subset \mathcal{G}(p_{X|G^\circ}, \delta)$, where

$$\delta = 1 - \frac{p_{G^\circ}(g_*) - \pi}{(1 - \pi)p_{G^\circ}(g_*)} \left(1 - \frac{\pi}{\min_{\mathbf{x}, g} p_{X|G^\circ}(\mathbf{x} | g)} \right). \quad (4.5)$$

Proof. See Appendix 4.A. □

In Proposition 4.1, we can interpret $\mathcal{F}_\pi$ to the set of outputs after passing G° through binary symmetric channels [70, Chapter 7.1] each of whose crossover probability is not more than π . Therefore, $G \in \mathcal{F}_\pi \subset \mathcal{G}(p_{X|G^\circ}, \delta)$ is a “noisy” version of G° .

4.2 Error Bounds and Optimization Relaxation

In this section, we first present bounds on the utility in the optimization problem (4.3). We then propose a series of steps to relax (4.3), and introduce the concept of a MFH, which simplifies our final optimization formulation significantly.

4.2.1 Error Bounds

Theorem 4.1. *Suppose that H has a uniform prior probability distribution under p_{X,H,G° . Then, the objective function in (4.3) satisfies*

$$\begin{aligned} & \frac{1}{2} - \frac{(e^\epsilon - 1) \|p_{X|H}(\cdot | -1) - p_{X|H}(\cdot | 1)\|_{TV}}{2 \min_{G \in \mathcal{G}_X} \{\max\{m_G, e^\epsilon - 1\}\}} \\ & \geq \min_{\gamma_H, p_{Z|X} \in \mathcal{Q}} \mathbb{P}(\gamma_H(Z) \neq H) \\ & \geq \frac{1}{2} - \sqrt{\frac{I(H; X | G^\circ) + \epsilon}{2}}, \end{aligned} \quad (4.6)$$

where $p_{X|H}(\mathbf{x} \mid h) = \sum_{g=-1,1} p_{X,G^\circ|H}(\mathbf{x}, g \mid h)/2$, and for each $G \in \mathcal{G}_X$ with corresponding joint distribution $p_{X,H,G}$,

$$\begin{aligned} m_G &= \max \left\{ |a_G(g) - b_G(g)|, |c_G(g) - d_G(g)| : g \in \{-1, 1\} \right\}, \\ a_G(g) &= \sum_{\mathbf{x} \in I^+} (p_{X|G}(\mathbf{x} \mid g) - e^\epsilon p_X(\mathbf{x}; G)), \\ b_G(g) &= \sum_{\mathbf{x} \in I^-} (p_{X|G}(\mathbf{x} \mid g) - e^\epsilon p_X(\mathbf{x}; G)), \\ c_G(g) &= \sum_{\mathbf{x} \in I^+} (e^{-\epsilon} p_X(\mathbf{x}; G) - p_{X|G}(\mathbf{x} \mid g)), \\ d_G(g) &= \sum_{\mathbf{x} \in I^-} (e^{-\epsilon} p_X(\mathbf{x}; G) - p_{X|G}(\mathbf{x} \mid g)), \end{aligned}$$

for $g = -1, 1$, and

$$\begin{aligned} p_X(\mathbf{x}; G) &= \sum_{g=-1,1} p_{X|G}(\mathbf{x} \mid g) p_G(g), \\ I^+ &= \{\mathbf{x} \in \mathcal{X}^s : p_{X|H}(\mathbf{x} \mid -1) - p_{X|H}(\mathbf{x} \mid 1) \geq 0\}, \\ I^- &= \{\mathbf{x} \in \mathcal{X}^s : p_{X|H}(\mathbf{x} \mid -1) - p_{X|H}(\mathbf{x} \mid 1) < 0\}. \end{aligned}$$

Proof. See Appendix 4.B. □

Remark 3. From the proof of Theorem 4.1, the lower bound in (4.6) can be generalized to case (b) in Remark 2 by replacing $I(H; X \mid G^\circ)$ (defined using p_{X,H,G°) with $I(H; X \mid G)$ (defined using $p_{X,H,G}$) and taking expectation w.r.t. G . Similarly, it can be generalized to case (c) in Remark 2 by replacing $I(H; X \mid G^\circ)$ with $I(H, X \mid G')$ (see Remark 2 for definition of G').

Both the bounds in Theorem 4.1 approach to $1/2$ as $\epsilon \rightarrow 0$, which is expected as a budget $\epsilon = 0$ means perfect information privacy. We present numerical results in Section 4.4 to show the behavior of these bounds for different $I(X; H)$.

4.2.2 Optimization Relaxation

The optimization problem (4.3) is a semi-infinite programming [97, (1)]. Without the constraints (4.3b), it is already a NP-complete problem [98]. In the following, we present a series of relaxations so that a simple block nonlinear Gauss-Siedel method (GS) approach can be developed (see Section 4.3).

For any fusion rule γ_G used to infer G at the fusion center, let

$$R_G(p_{Z|X}, \gamma_G) = \frac{1}{2} \left(\mathbb{P}(\gamma_G(Z) = -1 \mid G = 1) + \mathbb{P}(\gamma_G(Z) = 1 \mid G = -1) \right), \quad (4.7)$$

be the *average of its Type I and Type II detection error probabilities* for G . Note that we are not assuming $p_G(-1) = p_G(1) = 1/2$. As will be clear from the following discussion and (4.3b) itself, the prior distribution of G does not affect the information privacy constraint in (4.3b). For any $\mathbf{z} \in \mathcal{Z}^s$ and hypothesis G , let

$$\ell_{Z|G}(\mathbf{z}) = \frac{p_{Z|G}(\mathbf{z} \mid 1)}{p_{Z|G}(\mathbf{z} \mid -1)}.$$

From Proposition 3.1, to achieve ϵ -information privacy for $G \in \mathcal{G}_X$, it suffices to ensure that

$$\min_{\gamma_G} R_G(p_{Z|X}, \gamma_G) \geq \theta_G, \quad (4.8)$$

where $\theta_G = (1 - c_G(1 - e^{-\epsilon}))/2$ with

$$c_G = \min \left\{ \mathbb{P} \left(Z \in \arg \min_{\mathbf{z} \in \mathcal{Z}^s} \ell_{Z|G}(\mathbf{z}) \mid G = -1 \right), \mathbb{P} \left(Z \in \arg \max_{\mathbf{z} \in \mathcal{Z}^s} \ell_{Z|G}(\mathbf{z}) \mid G = 1 \right) \right\}.$$

The value c_G corresponds to the minimum probability of the fusion center making the correct inference about G over all realizations of Z . For example, if $Z = \mathbf{z}$

minimizes $\ell_{Z|G}(\cdot)$, then the fusion center makes the correct inference that $G = 0$.

We have

$$\begin{aligned}
\min_{G \in \mathcal{G}_X} c_G &\geq \min_{G \in \mathcal{G}_X, \mathbf{z} \in \mathcal{Z}^s, g \in \{-1, 1\}} p_{Z|G}(\mathbf{z}|g) \\
&= \min_{G, \mathbf{z}, g} \sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G}(\mathbf{x}|g) \\
&\geq (1 - \delta) \min_{\mathbf{z}, g} \sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G^\circ}(\mathbf{x}|g) \\
&\geq (1 - \delta) \alpha \Delta,
\end{aligned}$$

where the penultimate inequality follows from the definition (4.1) and the last inequality is due to the assumptions (1.1d), (4.2). Let

$$\theta = \frac{1}{2} (1 - (1 - e^{-\epsilon})(1 - \delta) \alpha \Delta). \quad (4.9)$$

Then, $\theta \geq \theta_G$ for all $G \in \mathcal{G}_X$. Therefore, by requiring that

$$\min_{G \in \mathcal{G}_X, \gamma_G} R_G(p_{Z|X}, \gamma_G) \geq \theta, \quad (4.10)$$

the information privacy constraints (4.3b) for all $G \in \mathcal{G}_X$ are satisfied. To further simplify the constraint (4.10), we note that for any privacy mapping $p_{Z|X} \in \mathcal{Q}$,

$$\mathcal{G}(p_{X|G^\circ}, \delta) \subset \mathcal{G}_Z \triangleq \mathcal{G}(p_{Z|G^\circ}, \delta),$$

since $p_{X|G}(\mathbf{x} | g) = (1 - \delta) p_{X|G^\circ}(\mathbf{x} | g) + \delta f_g(\mathbf{x})$ implies that

$$\begin{aligned}
p_{Z|G}(\mathbf{z}|g) &= \sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G}(\mathbf{x}|g) \\
&= (1 - \delta) p_{Z|G^\circ}(\mathbf{z} | g) + \delta f'_g(\mathbf{z}),
\end{aligned}$$

where $f'_g(\mathbf{z}) = \sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) f_g(\mathbf{x}) \in \mathcal{S}_{\mathcal{Z}^s}$.

Therefore, our optimization problem (4.3) is relaxed to:

$$\min_{\gamma_H, p_{Z|X} \in \mathcal{Q}} \mathbb{P}(\gamma_H(Z) \neq H) \quad (4.11a)$$

$$\text{s.t.} \quad \min_{G \in \mathcal{G}_Z, \gamma_G} R_G(p_{Z|X}, \gamma_G) \geq \theta. \quad (4.11b)$$

It turns out that there exists a “worst-case” hypothesis in the set $\mathcal{G}_Z$ so that achieving (4.11b) for this hypothesis implies information privacy for all other hypotheses in $\mathcal{G}_Z$ with the same privacy budget. We discuss how to find this “worst-case” hypothesis in the next subsection.

4.2.3 Most Favorable Hypothesis (MFH)

In this subsection, we introduce the notion of a Most Favorable Hypothesis (MFH) w.r.t. the privacy mapping $p_{Z|X}$. It is well known [68, Section II.B] that the decision rule γ_G^* minimizing (4.7) over all γ_G is given by

$$\gamma_G^*(\mathbf{z}) = \begin{cases} 1, & \text{if } \ell_{Z|G}(\mathbf{z}) \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Definition 4.2. For a privacy mapping $p_{Z|X} \in \mathcal{Q}$, if there exists a hypothesis $G \in \mathcal{G}_Z$, such that

$$R_G(p_{Z|X}, \gamma_G^*) = \min_{G' \in \mathcal{G}_Z} R_{G'}(p_{Z|X}, \gamma_{G'}^*),$$

then we call G a MFH w.r.t. $p_{Z|X}$, and $p_{Z|G}$ a Most Favorable Distribution (MFD) w.r.t. $\mathcal{G}_Z$.

The definitions of MFH and MFD mirror Huber’s definition of a pair of LFD [62]. While the LFD is defined for a pair of uncertainty classes within a hypothesis, our MFH and MFD are defined for an uncertainty set of hypotheses. Similar to [62], we show below that a MFH always exists and give an explicit construction of a MFD.

Theorem 4.2. *For any privacy mapping $p_{Z|X} \in \mathcal{Q}$ and $\mathcal{G}_Z = \mathcal{G}(p_{Z|G^\circ}, \delta)$, we have*

$$\min_{G \in \mathcal{G}_Z} R_G(p_{Z|X}, \gamma_G^*) = (1 - \delta) R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}^*), \quad (4.12)$$

and a MFD $p_{Z|G_{MF}}$ is given by

$$\begin{aligned} & p_{Z|G_{MF}}(\mathbf{z} \mid -1) \\ &= \begin{cases} A_1(1 - \delta)p_{Z|G^\circ}(\mathbf{z} \mid 1), & \mathbf{z} = \underline{\mathbf{z}} \\ (1 - \delta)p_{Z|G^\circ}(\mathbf{z} \mid -1), & \text{otherwise,} \end{cases} \end{aligned} \quad (4.13a)$$

$$\begin{aligned} & p_{Z|G_{MF}}(\mathbf{z} \mid 1) \\ &= \begin{cases} A_2(1 - \delta)p_{Z|G^\circ}(\mathbf{z} \mid -1), & \mathbf{z} = \bar{\mathbf{z}} \\ (1 - \delta)p_{Z|G^\circ}(\mathbf{z} \mid 1), & \text{otherwise,} \end{cases} \end{aligned} \quad (4.13b)$$

where

$$\begin{aligned} \underline{\mathbf{z}} &= \arg \min_{\mathbf{z}} \ell_{Z|G^\circ}(\mathbf{z}), \\ \bar{\mathbf{z}} &= \arg \max_{\mathbf{z}} \ell_{Z|G^\circ}(\mathbf{z}), \\ A_1 &= \frac{\frac{\delta}{1-\delta} + p_{Z|G^\circ}(\underline{\mathbf{z}} \mid -1)}{p_{Z|G^\circ}(\underline{\mathbf{z}} \mid 1)}, \\ A_2 &= \frac{\frac{\delta}{1-\delta} + p_{Z|G^\circ}(\bar{\mathbf{z}} \mid 1)}{p_{Z|G^\circ}(\bar{\mathbf{z}} \mid -1)}, \end{aligned}$$

with $\underline{\mathbf{z}}$ and $\bar{\mathbf{z}}$ randomly chosen from the set of minimizers and maximizers, respectively, if they are not unique.

Proof. See Appendix 4.C. □

Theorem 4.2 provides a MFD w.r.t. $\mathcal{G}_Z$, but the MFD may not be unique (consider for example the case where there are more than one $\mathbf{z}$ that minimize $\ell_{Z|G^\circ}(\mathbf{z})$). Given a MFD $p_{Z|G_{MF}}$, the corresponding MFH G_{MF} is also not unique (the MFH

distribution $p_{G_{\text{MF}}}$ is any binary distribution satisfying the conditions in the uncertainty set definition in Definition 4.1 and $p_{X|G_{\text{MF}}}$ may not be uniquely determined by $p_{Z|G_{\text{MF}}}$). However, (4.12) always holds, which allows us to replace the constraint (4.11b) with a simpler constraint to obtain the following optimization problem:

$$\min_{\gamma_H, p_{Z|X} \in \mathcal{Q}} \mathbb{P}(\gamma_H(Z) \neq H) \quad (4.14a)$$

$$\text{s.t. } \min_{\gamma_{G^\circ}} R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}) \geq \frac{\theta}{1 - \delta}. \quad (4.14b)$$

Remark 4. In this chapter, we have assumed binary private hypotheses for simplicity. To generalize to M -ary private hypotheses, suppose that every $G \in \mathcal{G}_X$ has range $\{0, 1, \dots, M - 1\}$. For a detector $\gamma_{G,g}$ at the fusion center that distinguishes between the hypothesis pair $G = 0$ and $G = g$, where $1 \leq g \leq M - 1$, let

$$R_{G,g}(p_{Z|X}, \gamma_{G,g}) = \frac{1}{2} \left(\mathbb{P}(\gamma_{G,g}(Z) = g \mid G = 0) + \mathbb{P}(\gamma_{G,g}(Z) = 0 \mid G = g) \right).$$

Then, from Theorem 3.2 and the discussions above, to achieve information privacy for all $G \in \mathcal{G}_X$, it is sufficient to replace the privacy constraint (4.14b) by

$$\min_{1 \leq g \leq M-1, \gamma_{G^\circ,g}} R_{G^\circ,g}(p_{Z|X}, \gamma_{G^\circ,g}) \geq \frac{\theta'}{1 - \delta}, \quad (4.15)$$

where $\theta' = \frac{1}{2} (1 - (1 - e^{-\epsilon/2})(1 - \delta)\alpha\Delta)$. In other words, the optimization problem (4.14) now has $M - 1$ constraints. The privacy mapping design and results in the next Section 4.3 can then be generalized accordingly.

Remark 5. Similar to the case of M -ary private hypotheses, suppose that we have multiple nominal binary private hypotheses $(G^\circ_k)_{k=1}^K$, with corresponding uncertainty sets $(\mathcal{G}(p_{X|G^\circ_k}, \delta_k))_{k=1}^K$. Then, (4.14b) can be generalized to the K constraints $\min_{\gamma} R_{G^\circ_k}(p_{Z|X}, \gamma) \geq \theta/(1 - \delta_k)$, for $k = 1, \dots, K$.

Algorithm 4.1 GS of local privacy mappings.

- 1: **input:** $p_{X,H,G^\circ}, \delta, r, \xi$
 - 2: **initialization:**
 - Let $p^{t,(0)}(z | x) = \frac{1}{|\mathcal{Z}|} + n^t(z | x), \forall z \in \mathcal{Z}$, with $\sum_{z \in \mathcal{Z}} n^t(z | x) = 0, \forall x \in \mathcal{X}, t = 1, 2, \dots, s$.
 - For fixed $p_{Z|X}^{(0)} = \prod_{t=1}^s p^{t,(0)}$, optimize θ_0 in (4.18) with likelihood ratio test.
 - Set $\theta = r\theta_0$.
 - Set $k = -1$ and $E^{(0)} = 1$.
 - 3: **repeat**
 - 4: (i) $k = k + 1$.
 - (ii) Fix $p_{Z|X}^{(k)} = \prod_{t=1}^s p^{t,(k)}$, optimize $\gamma_H^{(k+1)}$ to be the Bayes detector of hypothesis H .
 - (iii) Fix $\gamma_H^{(k+1)}$ and $p^{\setminus t,(k)} = \prod_{i=1}^{t-1} p^{i,(k+1)} \cdot \prod_{i=t+1}^s p^{i,(k)}$, for $t = 1, 2, \dots, s$, solve (4.17) to find the optimal $p^{t,(k+1)}$ and the optimal solution $E^{(k)}$. The linear solver used is [99, Algorithm 2].
 - 5: **until** $\frac{E^{(k)} - E^{(k+1)}}{E^{(k)}} \leq \xi$.
 - 6: **return** $p^t = p^{t,(k+1)}, t = 1, 2, \dots, s$.
-

4.3 Privacy Mapping Design

Since (4.14) is NP-complete (note that optimizing the objective is NP-complete [98]), we propose a suboptimal approach in this section to obtain the privacy mapping $p_{Z|X}$. We also derive some properties of this privacy mapping. We consider both the non-asymptotic case and asymptotic case where the number of sensors $s \rightarrow \infty$.

4.3.1 Fixed number of sensors

We optimize (4.14) using GS as follows. For a fixed $p_{Z|X}$, we first find the Bayesian detector γ_H that minimizes the error in detecting H . Let $p^{\setminus t} = \prod_{j \neq t} p^j$. Then for each sensor $t = 1, \dots, s$, we fix $\gamma_H, p^{\setminus t}$ and optimize (4.14) over all privacy mappings p^t satisfying (1.1b),(1.1c),(1.1d). This process is iterated until a convergence criterion is satisfied. Let Φ be the set of *deterministic* local privacy

mappings (a conditional probability distribution $p_{Z^t|X^t}(\cdot | x^t)$ with all its mass at a particular $z^t \in \mathcal{Z}$ for each $x^t \in \mathcal{X}$). The set Φ is finite since $|\mathcal{X}| \cdot |\mathcal{Z}| < \infty$. Then, for a fixed $p^{\setminus t}$, we can write

$$p_{Z|X} = \sum_{\phi \in \Phi} \nu_{\phi} \phi \cdot p^{\setminus t}, \quad (4.16)$$

where $\sum_{\phi \in \Phi} \nu_{\phi} = 1$, with $\nu_{\phi} \geq 0$, for all $\phi \in \Phi$. Let $L_H(p_{Z|X}, \gamma_H)$ be $\mathbb{P}(\gamma_H(Z) \neq H)$ when the fusion rule for H is γ_H and the privacy mapping $p_{Z|X}$ is used. From (4.16), we have

$$L_H(p_{Z|X}, \gamma_H) = \sum_{\phi \in \Phi} \nu_{\phi} L_H(\phi \cdot p^{\setminus t}, \gamma_H).$$

Similarly,

$$\begin{aligned} \min_{\gamma_{G^{\circ}}} R_{G^{\circ}}(p_{Z|X}, \gamma_{G^{\circ}}) &= \min_{\gamma_{G^{\circ}}} \sum_{\phi \in \Phi} \nu_{\phi} R_{G^{\circ}}(\phi \cdot p^{\setminus t}, \gamma_{G^{\circ}}) \\ &\geq \sum_{\phi \in \Phi} \nu_{\phi} \min_{\gamma_{G^{\circ}}} R_{G^{\circ}}(\phi \cdot p^{\setminus t}, \gamma_{G^{\circ}}). \end{aligned}$$

In our GS procedure, when optimizing for each sensor t with a fixed γ_H and $p^{\setminus t}$, we solve the following relaxed optimization problem:

$$\begin{aligned} \min_{\nu_{\phi}} \quad & \sum_{\phi \in \Phi} \nu_{\phi} L_H(\phi \cdot p^{\setminus t}, \gamma_H) \\ \text{s.t.} \quad & \sum_{\phi \in \Phi} \nu_{\phi} \min_{\gamma_{G^{\circ}}} R_{G^{\circ}}(\phi \cdot p^{\setminus t}, \gamma_{G^{\circ}}) \geq \frac{\theta}{1 - \delta}, \\ & \sum_{\phi \in \Phi} \nu_{\phi} = 1, \quad \nu_{\phi} \geq 0, \quad \forall \phi \in \Phi. \end{aligned} \quad (4.17)$$

Proposition 4.2. *The optimal solution of (4.17) is a randomization between at most two deterministic privacy mappings.*

Proof. The optimization problem (4.17) is a linear program, and has $|\mathcal{X}| \cdot |\mathcal{Z}|$ variables with $|\mathcal{X}| \cdot |\mathcal{Z}| + 2$ constraints. From linear programming theory [17], the

number of constraints in the optimal solution for which equality holds is no smaller than the number of variables. Therefore, at most two of the weights $\{\nu_\phi : \phi \in \Phi\}$ are nonzero. $\square$

Remark 6. In practice, it may not be easy to know *a priori* what privacy budget ϵ to set. Therefore, instead of fixing a privacy budget ϵ and choosing θ as in (4.9), we let $\theta = r\theta_0$, where $r \in (0, 1)$ is called the *privacy threshold ratio*, and

$$\theta_0 = (1 - \delta) \max_{p_{Z|X} \in \mathcal{Q}} \min_{\gamma_{G^\circ}} R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}) \quad (4.18)$$

corresponds to the maximal information privacy achievable for G° under the constraints imposed by $\mathcal{Q}$. The maximization on the right hand side of (4.18) is achieved by letting $p^t(z^t | x^t) = 1/|\mathcal{Z}|$, for all $z^t \in \mathcal{Z}$, $x^t \in \mathcal{X}$, and $t = 1, \dots, s$. With this privacy mapping, it can be shown that Z and G° are independent, and

$$\theta_0 = \frac{1 - \delta}{2}. \quad (4.19)$$

In some applications, we may include additional constraints in the privacy mapping set $\mathcal{Q}$ in order to improve the utility. For example, we may wish to constrain $|p^t(z^t | x^t) - 1/|\mathcal{Z}|| \geq \Delta'$ for some $\Delta' > 0$ to avoid a mapping that results in Z^t being independent of H . In this case, θ_0 can be found from (4.18) using a GS approach.

The above discussion on our proposed GS approach is summarized in (4.1). We initialize the privacy mapping for each sensor t to be $p^{t,(0)}(z | x) = 1/|\mathcal{Z}| + n^t(z | x)$, where $n^t(z | x)$ is a small random noise, satisfying $\sum_{z \in \mathcal{Z}} n^t(z | x) = 0$, for all $x \in \mathcal{X}$. This is done to help the GS procedure avoid local optimum points. For each iteration k , we perform the GS procedure for each sensor $t = 1, \dots, s$ in order. This is iterated until a convergence criterion is reached.

The optimization problem (4.17) at each sensor t is a Linear Program (LP), which can be solved efficiently in polynomial time using standard LP solvers. The storage complexity and time to encode this LP is $O(|\mathcal{Z}|^{|\mathcal{X}|})$ since each term

in the sums needs to be computed. Its solution time complexity is at most $O((|\mathcal{X}||\mathcal{Z}|)^{2.38} \log \xi_0^{-1})$, where ξ_0 is the relative accuracy of the LP solver (see [99] for a tighter bound), for each sensor at each iteration of (4.1).

Proposition 4.3. *Algorithm 4.1 converges to a critical point.*

Proof. In Algorithm 4.1, since the Bayes error probability of detecting H is non-increasing in steps 4(ii) and 4(iii), $E^{(k)}$ is non-increasing in each iteration. Since $E^{(k)}$ is lower bounded, E^k converges. From [83, Proposition 4], the convergence of $E^{(k)}$ implies the convergence of γ_H and $p^t, t = 1, 2, \dots, s$ to limit points. By [83, Proposition 5], every limit point of γ_H and $p^t, t = 1, 2, \dots, s$ is a critical point for Algorithm 4.1. The proposition is now proved. $\square$

4.3.2 Large number of sensors

In this subsection, we consider the case where the number of sensors goes to infinity. We assume that the sensor observations are conditionally independently identically distributed (i.i.d.) given the hypotheses, and every sensor t uses an identical local privacy mapping $p^t = q$. We have $p_{Z|X}(\mathbf{z} | \mathbf{x}) = \prod_{t=1}^s q(z^t | x^t)$.

Let $\tilde{\mathcal{Q}}$ be the set of privacy mappings $q(z|x)$. For any $q \in \tilde{\mathcal{Q}}$ such that Z , H and G° are not independent, the error probabilities $\min_{\gamma_H} \mathbb{P}(\gamma_H(Z) \neq H) \rightarrow 0$ and $\min_{\gamma_{G^\circ}} R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}) \rightarrow 0$ exponentially fast as $s \rightarrow \infty$ [17, 70]. In particular, we observe that other than using a trivial mapping that results in the worst case error probability for detecting H , there is no way to avoid violating the privacy constraint (4.3b) for a fixed ϵ budget when s is sufficiently large. This result follows from Proposition 3.1(i) since $\min_{\gamma_{G^\circ}} R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}) \rightarrow 0$ as $s \rightarrow \infty$. Therefore, we consider replacing the objective and constraint in the optimization problem (4.14) with their respective error exponents. The error exponent of (4.14a) is given by

$$\begin{aligned} & \lim_{s \rightarrow \infty} -\frac{1}{s} \log \min_{\gamma_H} \mathbb{P}(\gamma_H(Z) \neq H) \\ &= C_H(p) := -\min_{\lambda \in [0,1]} \log \sum_{z \in \mathcal{Z}} p_{Z^1|H}^\lambda(z | 1) p_{Z^1|H}^{1-\lambda}(z | -1), \end{aligned} \quad (4.20)$$

while that for (4.14b) is

$$\begin{aligned} & \lim_{s \rightarrow \infty} -\frac{1}{s} \log \min_{\gamma_{G^\circ}} R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}) \\ &= C_{G^\circ}(p) := - \min_{\lambda \in [0,1]} \log \sum_{z \in \mathcal{Z}} p_{Z^1|G^\circ}^\lambda(z | 1) p_{Z^1|G^\circ}^{1-\lambda}(z | -1). \end{aligned} \quad (4.21)$$

We then obtain the optimization problem

$$\max_{p \in \tilde{\mathcal{Q}}} C_H(p) \quad (4.22a)$$

$$\text{s.t. } C_{G^\circ}(p) \leq \beta, \quad (4.22b)$$

for some $\beta > 0$. The constraint (4.22b) is a privacy *rate* constraint that prevents the average error probability of inferring any $G \in \mathcal{G}_X$ from decaying faster than a fixed rate β as $s \rightarrow \infty$. This corresponds to preventing the information privacy leakage of any $G \in \mathcal{G}_X$ from worsening faster than a predefined rate.

Proposition 4.4. *There is no loss in optimality in (4.22) if we restrict $|\mathcal{Z}| \leq |\mathcal{X}| + 1$.*

Proof. For any $p \in \tilde{\mathcal{Q}}$, we have

$$\begin{aligned} e^{-C_H(p)} &= \sum_{z \in \mathcal{Z}} p_{Z^1}(z) d_H(z), \\ e^{-C_{G^\circ}(p)} &= \sum_{z \in \mathcal{Z}} p_{Z^1}(z) d_{G^\circ}(z), \end{aligned}$$

where

$$\begin{aligned} d_H(z) &= \frac{p_{H|Z^1}(-1 | z)}{p_H(-1)} \left(\frac{p_{H|Z^1}(1 | z) p_H(-1)}{p_{H|Z^1}(-1 | z) p_H(1)} \right)^{\lambda_1}, \\ d_{G^\circ}(z) &= \frac{p_{G^\circ|Z^1}(-1 | z)}{p_{G^\circ}(-1)} \left(\frac{p_{G^\circ|Z^1}(1 | z) p_{G^\circ}(-1)}{p_{G^\circ|Z^1}(-1 | z) p_{G^\circ}(1)} \right)^{\lambda_2}, \\ \lambda_1 &= \arg \min_{\lambda \in (0,1)} \log \sum_{z \in \mathcal{Z}} p_{Z^1|H}^\lambda(z | 1) p_{Z^1|H}^{1-\lambda}(z | -1), \\ \lambda_2 &= \arg \min_{\lambda \in (0,1)} \log \sum_{z \in \mathcal{Z}} p_{Z^1|G^\circ}^\lambda(z | 1) p_{Z^1|G^\circ}^{1-\lambda}(z | -1). \end{aligned}$$

For each $z \in \mathcal{Z}$, let w_z be the vector $(p_{X^1|Z^1}(x | z))_{x \in \mathcal{X}}$. We can then write

$$\sum_{z \in \mathcal{Z}} p_{Z^1}(z) [w_z, d_H(z), d_{G^\circ}(z)] = [p_{X^1}, e^{-C_H(p)}, e^{-C_{G^\circ}(p)}].$$

The vector $[p_{X^1}, e^{-C_H(p)}, e^{-C_{G^\circ}(p)}]$ belongs to a space with dimension $|\mathcal{X}| + 1$. From Carathéodory's theorem [100, Appendix C], since $w_z, d_H(z), d_{G^\circ}(z)$ are real-valued continuous functions of $p_{X^1|Z^1}(\cdot | z)$, we can find Z'^1 such that

$$|\{z : p_{Z'^1}(z) \neq 0\}| \leq |\mathcal{X}| + 1,$$

and

$$\sum_{z \in \mathcal{Z}} p_{Z'^1}(z) [w_z, d_H(z), d_{G^\circ}(z)] = [p_{X^1}, e^{-C_H(p)}, e^{-C_{G^\circ}(p)}].$$

The proposition now follows. □

Recall that Φ is the set of deterministic local privacy mappings. Any $p \in \tilde{\mathcal{Q}}$ can be written as a randomization of deterministic mappings or the convex combination $\sum_{\phi \in \Phi} \nu_\phi \phi$, where $\nu_\phi \geq 0$ for all $\phi \in \Phi$ and $\sum_{\phi \in \Phi} \nu_\phi = 1$. We now consider the case where each sensor performs independent randomization over Φ [8], and reports its choice of privacy mapping to the fusion center.

Proposition 4.5. *Suppose that each sensor performs independent randomization over Φ . Then, the optimal solution of (4.22) is a randomization between at most two deterministic decision rules.*

Proof. From [17, Lemma 1], for $J \in \{H, G^\circ\}$, we have

$$\lim_{s \rightarrow \infty} \frac{1}{s} \log \min \mathbb{P}(\gamma_J(Z) \neq J) = - \min_{\lambda \in [0,1]} \sum_{\phi \in \Phi} \nu_\phi \mu_J(\phi, \lambda),$$

where

$$\begin{aligned}\mu_J(\phi, \lambda) &= \log \sum_{z \in \mathcal{Z}} u_J^\lambda(\phi, z, 1) u_J^{1-\lambda}(\phi, z, -1), \\ u_J(\phi, z, j) &= \sum_{x \in \phi^{-1}(z)} p_{X|J}(x | j).\end{aligned}$$

Therefore, (4.22) can be reformulated as

$$\begin{aligned}\min_{\nu_\phi, \lambda_1 \in [0,1]} & \sum_{\phi \in \Phi} \nu_\phi \mu_H(\phi, \lambda_1) \\ \text{s.t.} \quad & \min_{\lambda_2 \in [0,1]} \sum_{\phi \in \Phi} \nu_\phi \mu_{G^\circ}(\phi, \lambda_2) \geq -\beta, \\ & \nu_\phi \geq 0, \forall \phi \in \Phi, \sum_{\phi \in \Phi} \nu_\phi = 1.\end{aligned}\tag{4.23}$$

For any fixed λ_1 and λ_2 , the linear program (4.23) has $|\mathcal{Z}| \cdot |\mathcal{X}|$ variables with $|\mathcal{Z}| \cdot |\mathcal{X}| + 2$ constraints. From linear programming theory [17], the number of constraints in the optimal solution for which equality holds is no smaller than the number of variables. Therefore, at most two of the ν_ϕ 's are nonzero, and the proposition is proved. $\square$

4.4 Simulations

In this section, we carry out simulation to provide insights into how different parameters impact the performance of our proposed approaches. We first consider a network of 4 sensors and a fusion center, with $\mathcal{X} = \{1, 2, \dots, 16\}$, and $\mathcal{Z} = \{1, 2\}$. We generate different p_{X,H,G° and calculate numerically the upper and lower bounds in Theorem 4.1. We set $I(X; H|G^\circ) = 4 \times 10^{-6}$ and impose an information privacy budget $\epsilon = 0.01$. Different p_{X,H,G° with different mutual information $I(X; H)$ are generated to test the performance of the bound. From Figure 4.1, we see that the difference between the upper and lower bounds become tighter when $I(X; H)$ is large. Furthermore, the upper bound is tighter than the

lower bound.

Figure 4.1: Upper and lower bounds in Theorem 4.1 and $\mathbb{P}(\gamma_H(Z) \neq H)$ calculated by Algorithm 4.1 with varying $I(X; H)$.

We set the correlation coefficient between the public hypothesis H and nominal private hypothesis G° to be 0.2. The joint distribution p_{X^1, H, G° at sensor 1 is shown in Figure 4.2. We assume that all sensors have the same joint distribution and are independent of each other conditioned on (H, G°) .

In our first simulation, we set the privacy threshold ratio $r = 0.7$ and $\delta = 0.54$ in (4.1) to obtain the privacy mapping $p_{Z|X}$. In Figure 4.3, we visualize the distribution of Z conditioned on different hypotheses, including the nominal private hypothesis G° , the MFH G_{MF} (constructed using Theorem 4.2), and the public hypothesis H . We observe that $p_{Z|G^\circ}(\cdot | 0)$ and $p_{Z|G^\circ}(\cdot | 1)$ are relatively similar to each other, which implies that Z does not provide much statistical information about G° . For the MFH G_{MF} , we have a similar observation that $p_{Z|G_{\text{MF}}}(\cdot | 0)$ and $p_{Z|G_{\text{MF}}}(\cdot | 1)$ are relatively similar to each other except at $Z = (0, 0, 0, 0)$ and $Z = (1, 1, 1, 1)$, which makes the MFH the easiest to distinguish amongst all hypotheses in $\mathcal{G}_Z$. In this case, we find that the privacy budget $\epsilon = 2.4$. We also observe that $p_{Z|H}(\cdot | -1)$ and $p_{Z|H}(\cdot | 1)$ are significantly different from

Figure 4.2: Joint distribution p_{X^1, H, G° of sensor observation X^1 , public hypothesis H , and private hypothesis G° . The correlation coefficient between H and G° is 0.2.

Figure 4.3: Conditional probability distribution of $Z = (Z^1, \dots, Z^4)$ given different hypotheses. The horizontal axis shows the decimal form of the binary word $Z^1 Z^2 Z^3 Z^4$.

each other, which implies that H can still be inferred from Z with good accuracy.

We also show the privacy mapping at sensor 1 in Figure 4.4 as an example.

In Figure 4.5, we show the Bayes error rates of detecting H and G_{MF} with

Figure 4.4: Privacy mapping at sensor 1, with $r = 0.7$ and $\delta = 0.54$.

different privacy threshold ratios r and uncertainty parameters δ . We let the privacy threshold $\theta = (1 - \delta)r/2$ as given by (4.19). We see that the our proposed (4.1) yields privacy mappings that allow the fusion center to detect the public hypothesis H with low error, while keeping the error for any private hypothesis $G \in \mathcal{G}_X$ high. As expected, the Bayes error of detecting all hypotheses increases as the threshold ratio r increases. Furthermore, the performance of our proposed (4.1) improves with smaller uncertainty parameter δ . This is intuitive since a larger δ implies a stricter privacy constraint.

Let ρ_{H,G° denote the correlation coefficient between H and G° . In Figure 4.6, we fix $\theta = 0.7$ and vary ρ_{H,G° . We see that if ρ_{H,G° is small, the Bayes error of detecting H is much smaller than the Bayes error of detecting any $G \in \mathcal{G}_Z$. However, if H and G° are highly correlated, the Bayes error of detecting H is close to that of G° , and greater than that of detecting G_{MF} . This is because G_{MF} is the MFH in $\mathcal{G}_Z$, and yields the smallest detection error amongst all $G \in \mathcal{G}_Z$.

Figure 4.5: Bayes error probability of detecting H , G^o and G_{MF} with different privacy threshold ratio r and uncertainty parameter δ .

4.4.1 Asymptotically Large Number of Sensors

Next, we plot the error exponents of detecting H and G^o as the number of sensors $s \rightarrow \infty$ in Figure 4.7. We use the same sensor observation distribution as in

Figure 4.6: Bayes error of detecting H , G° and G_{MF} with varying correlation coefficient between H and G° .

Figure 4.2 and set $\beta = 0.04$. We solve the optimization problem (4.22) using GS and Proposition 4.5 to obtain the optimal privacy mapping. We see that as the correlation coefficient between H and G° increases, the error exponent of detecting H decreases and approaches the error exponent of detecting G° .

Figure 4.7: Error exponent for detecting H and G° with varying correlation coefficient between H and G° .

In Figure 4.8, we study the accuracy-privacy tradeoff with varying $|\mathcal{Z}|$. We set the sensor observation set cardinality $|\mathcal{X}| = 6$, and $\beta = 0.05$. We consider the low-correlation case where $\rho_{H,G^\circ} = 0.2$, and the high-correlation case where $\rho_{H,G^\circ} = 0.8$. We see that as $|\mathcal{Z}|$ increases, the error exponent of detecting H increases and stays constant for $|\mathcal{Z}|$ sufficiently large. The point where the error exponent becomes constant is larger for smaller correlation coefficient ρ_{H,G° . The results in Figure 4.8 also verify Proposition 4.4.

Figure 4.8: Error exponent for detecting H and G° with varying $|\mathcal{Z}|$.

4.4.2 Comparisons

Finally, we compare our proposed approach with approaches using other types of privacy metrics. We set $\rho_{H,G^\circ} = 0.2$, and vary the distribution of X so that the conditional mutual information $I(X; H|G^\circ)$ varies accordingly. We consider the following three privacy metrics for comparison.

Definition 4.3 (Average information leakage [48]). *Let $\epsilon_A \geq 0$. We say that ϵ_A -average information leakage is achieved if $I(G^\circ; Z) \leq \epsilon_A$.*

Definition 4.4 (Local differential privacy [41]). *Let $\epsilon_L \geq 0$. We say that ϵ_L -local differential privacy is achieved if for all neighboring $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^s$, and $\mathbf{z} \in \mathcal{Z}^s$,*

$$\frac{p_{Z|X}(\mathbf{z}|\mathbf{x})}{p_{Z|X}(\mathbf{z}|\mathbf{x}')} \leq e^{\epsilon_L}.$$

Definition 4.5 (Maximal leakage [36]). *Let $\epsilon_{ML} \geq 0$. We say that ϵ_{ML} -maximal leakage privacy is achieved if $\sup_G \log \frac{\max_{\gamma} \mathbb{P}(\gamma(Z)=G)}{\max_g p_G(g)} \leq \epsilon_{ML}$.*

The approaches we compare against when $\delta = 0$ (i.e., $\mathcal{G}_X = \{G^\circ\}$) are optimization formulations to achieve the minimum detection error while guaranteeing ϵ_A -average information privacy and ϵ_L -local differential privacy respectively. This ensures fair comparison as these approaches deal only with a single private hypothesis. When $\delta > 0$, we compare against the optimization formulation (3) in [36], which guarantees ϵ_{ML} -maximal leakage privacy in the worst case private hypothesis.

In Figure 4.9, we choose the privacy thresholds and budgets for the different metrics so that the Bayes error for detecting the private hypothesis G° is the same for all metrics when $\delta = 0$. Similarly, when $\delta > 0$, we set the Bayes error for detecting G_{MF} to be the same across the metrics under comparison. We see that the Bayes error for detecting H decreases as $I(X; H \mid G^\circ)$ increases. As expected, our proposed approach using information privacy as the privacy metric yields the minimum Bayes error for detecting H . By comparison, the approach using average information leakage has a slightly higher Bayes error for detecting H . The approach using local differential privacy has a significantly larger Bayes error for detecting H , since local differential privacy protects the *data privacy* of the sensor observations X and does not distinguish between statistical inferences for H and G° . For the $\delta > 0$ comparison, the maximal leakage privacy formulation (3) in [36] achieves the worst performance because it considers the strictest privacy criterion using the worst-case private hypothesis G . This serves as a benchmark upper bound for our utility performance if we let $\delta \rightarrow 1$ in our uncertainty set $\mathcal{G}_X$.

Figure 4.9: Bayes error for detecting H and G^o with varying conditional mutual information $I(X; H | G^o)$.

4.A Proof of Proposition 4.1

If $\pi = 0$, the proposition trivially holds. We now assume that $\pi > 0$ so that $0 < \delta < 1$ from (4.5). Consider a $G \in \mathcal{F}_\pi$. For any $g \in \{0, 1\}$, let $\bar{g} = 1 - g$ and for any $\mathbf{x} \in \mathcal{X}^s$, let

$$f_g(\mathbf{x}) = \frac{1}{\delta} (p_{X|G}(x | g) - (1 - \delta)p_{X|G^o}(x | g)).$$

To prove the proposition, it suffices to show that $f_g(\mathbf{x}) \geq 0$ since $\sum_{\mathbf{x}} f_g(\mathbf{x}) = 1$.

We have

$$\begin{aligned} p_{X|G^o}(\mathbf{x} | g) &= \frac{p_{X,G^o,G}(\mathbf{x}, g, g)}{\sum_{g'=0,1} p_{G^o,G}(g, g')} + \frac{p_{X,G^o,G}(\mathbf{x}, g, \bar{g})}{p_{G^o}(g)} \\ &\leq p_{X|G^o,G}(\mathbf{x} | g, g) + \frac{p_{X,G^o,G}(\mathbf{x}, g, \bar{g})}{p_{G^o}(g)} \\ &\leq p_{X|G^o,G}(\mathbf{x} | g, g) + p_{G|G^o}(\bar{g} | g), \end{aligned}$$

and

$$\begin{aligned}
 p_{X|G}(\mathbf{x} | g) &= \sum_{g'=0,1} p_{X|G^\circ, G}(\mathbf{x} | g', g) p_{G^\circ|G}(g' | g) \\
 &\geq p_{G^\circ|G}(g | g)(p_{X|G^\circ}(\mathbf{x} | g) - p_{G|G^\circ}(\bar{g} | g)). \tag{4.24}
 \end{aligned}$$

We also have

$$\begin{aligned}
 p_{G^\circ|G}(\bar{g} | g) &= \frac{p_{G|G^\circ}(g | \bar{g})p_{G^\circ}(\bar{g})}{p_G(g)} \\
 &\leq \frac{\pi p_{G^\circ}(\bar{g})}{\sum_{g'=0,1} p_{G|G^\circ}(g | g')p_{G^\circ}(g')} \\
 &\leq \frac{\pi p_{G^\circ}(\bar{g})}{p_{G|G^\circ}(g | g)p_{G^\circ}(g)} \\
 &\leq \frac{\pi p_{G^\circ}(\bar{g})}{(1 - \pi)p_{G^\circ}(g)}. \tag{4.25}
 \end{aligned}$$

From (4.24),(4.25), we obtain

$$\begin{aligned}
 f_g(\mathbf{x}) &= \frac{1}{\delta}(p_{X|G}(\mathbf{x} | g) - (1 - \delta)p_{X|G^\circ}(\mathbf{x} | g)) \\
 &\geq \frac{1}{\delta} \left(\frac{p_{G^\circ}(g) - \pi}{(1 - \pi)p_{G^\circ}(g)} (p_{X|G^\circ}(\mathbf{x} | g) - \pi) \right. \\
 &\quad \left. - (1 - \delta)p_{X|G^\circ}(\mathbf{x} | g) \right) \\
 &= \frac{p_{X|G^\circ}(\mathbf{x} | g)}{\delta} \left(\frac{p_{G^\circ}(g) - \pi}{(1 - \pi)p_{G^\circ}(g)} \left(1 - \frac{\pi}{p_{X|G^\circ}(\mathbf{x} | g)} \right) \right. \\
 &\quad \left. - 1 + \delta \right) \\
 &\geq 0,
 \end{aligned}$$

where the last inequality follows from (4.5). The proposition is now proved.

4.B Proof of Theorem 4.1

It can be shown that

$$\begin{aligned} & \min_{\gamma_H, p_{Z|X}} \mathbb{P}(\gamma_H(Z) \neq H) \\ &= \frac{1}{2} - \frac{1}{2} \max_{p_{Z|X}} \|p_{Z|H}(\cdot | 0) - p_{Z|H}(\cdot | 1)\|_{TV}. \end{aligned} \quad (4.26)$$

We prove the theorem by bounding the total variation term in (4.26). We first prove the lower bound. From Pinsker's inequality [101], we have

$$\begin{aligned} & \frac{1}{2} \left(\sum_{\mathbf{z}} |p_{Z|H}(\mathbf{z} | 1) - p_Z(\mathbf{z})| \right)^2 \\ & \leq \sum_{\mathbf{z}} p_{Z|H}(\mathbf{z} | 1) \log \frac{p_{Z|H}(\mathbf{z} | 1)}{p_Z(\mathbf{z})} \end{aligned} \quad (4.27)$$

and

$$\begin{aligned} & \frac{1}{2} \left(\sum_{\mathbf{z}} |p_{Z|H}(\mathbf{z} | 0) - p_Z(\mathbf{z})| \right)^2 \\ & \leq \sum_{\mathbf{z}} p_{Z|H}(\mathbf{z} | 0) \log \frac{p_{Z|H}(\mathbf{z} | 0)}{p_Z(\mathbf{z})}. \end{aligned} \quad (4.28)$$

From (4.27),(4.28), we obtain

$$\begin{aligned} & I(H; Z) \\ & \geq \frac{1}{4} \left(\sum_{\mathbf{z}} \frac{1}{2} |p_{Z|H}(\mathbf{z} | 1) - p_{Z|H}(\mathbf{z} | 0)| \right)^2 \\ & \quad + \frac{1}{4} \left(\sum_{\mathbf{z}} \frac{1}{2} |p_{Z|H}(\mathbf{z} | 1) - p_{Z|H}(\mathbf{z} | 0)| \right)^2 \\ & = \frac{1}{2} \|p_{Z|H}(\cdot | 1) - p_{Z|H}(\cdot | 0)\|_{TV}^2. \end{aligned} \quad (4.29)$$

We have

$$I(H; Z) = I(H; Z | G^\circ) + I(G^\circ; Z) - I(G^\circ; Z | H)$$

$$\begin{aligned}
&\leq I(H; X \mid G^\circ) + I(G^\circ; Z) \\
&\leq I(H; X \mid G^\circ) + \epsilon,
\end{aligned} \tag{4.30}$$

where the last inequality follows from $I(G^\circ; Z) = \sum_{\mathbf{z}, g} p_{Z, G^\circ}(\mathbf{z}, g) \log(p_{G^\circ|Z}(g \mid \mathbf{z})/p_{G^\circ}(g)) \leq \epsilon$ and (4.3b). From (4.29), (4.30), we obtain

$$\begin{aligned}
&\|p_{Z|H}(\cdot \mid 1) - p_{Z|H}(\cdot \mid 0)\|_{TV} \\
&\leq \sqrt{2I(H; Z)} \\
&\leq \sqrt{2I(H; X \mid G^\circ) + \epsilon},
\end{aligned}$$

and the lower bound in (4.6) follows from (4.26).

We next prove the upper bound. Fix a subset of $\Gamma \in \mathcal{Z}^s$, and consider a $p_{Z|X}$ such that

$$\begin{aligned}
p_{Z|X}(\mathbf{z} \mid \mathbf{x}) &= \begin{cases} A, & \mathbf{x} \in I^+, \mathbf{z} \in \Gamma \\ B, & \mathbf{x} \in I^-, \mathbf{z} \in \Gamma \end{cases}, \\
p_{Z|X}(\mathbf{z} \mid \mathbf{x}) &= \begin{cases} B, & \mathbf{x} \in I^+, \mathbf{z} \in \Gamma^c \\ A, & \mathbf{x} \in I^-, \mathbf{z} \in \Gamma^c \end{cases}
\end{aligned}$$

for some $A, B \geq 0$. We then have

$$\begin{aligned}
&\|p_{Z|H}(\cdot \mid 0) - p_{Z|H}(\cdot \mid 1)\|_{TV} \\
&= \frac{1}{2} \sum_{\mathbf{z}} |p_{Z|H}(\mathbf{z} \mid 0) - p_{Z|H}(\mathbf{z} \mid 1)| \\
&= \frac{1}{2} \sum_{\mathbf{z}} \left| \sum_{\mathbf{x}} (p_{X|H}(\mathbf{x} \mid 0) - p_{X|H}(\mathbf{x} \mid 1)) p_{Z|X}(\mathbf{z} \mid \mathbf{x}) \right| \\
&= \frac{1}{2} \sum_{\mathbf{z} \in \Gamma} \left| A \sum_{\mathbf{x} \in I^+} (p_{X|H}(\mathbf{x} \mid 0) - p_{X|H}(\mathbf{x} \mid 1)) \right. \\
&\quad \left. + B \sum_{\mathbf{x} \in I^-} (p_{X|H}(\mathbf{x} \mid 0) - p_{X|H}(\mathbf{x} \mid 1)) \right|
\end{aligned} \tag{4.31}$$

$$\begin{aligned}
& + \frac{1}{2} \sum_{\mathbf{z} \in \Gamma^c} \left| B \sum_{\mathbf{x} \in I^+} (p_{X|H}(\mathbf{x} | 0) - p_{X|H}(\mathbf{x} | 1)) \right. \\
& \quad \left. + A \sum_{\mathbf{x} \in I^-} (p_{X|H}(\mathbf{x} | 0) - p_{X|H}(\mathbf{x} | 1)) \right| \\
& = \frac{1}{2} |\mathcal{Z}|^s |A - B| \|p_{X|H}(\cdot | 0) - p_{X|H}(\cdot | 1)\|_{TV}, \tag{4.32}
\end{aligned}$$

where (4.32) holds since

$$\begin{aligned}
& \sum_{\mathbf{x} \in I^+} (p_{X|H}(\mathbf{x} | 0) - p_{X|H}(\mathbf{x} | 1)) \\
& = - \sum_{\mathbf{x} \in I^-} (p_{X|H}(\mathbf{x} | 0) - p_{X|H}(\mathbf{x} | 1)) \\
& = \|p_{X|H}(\cdot | 0) - p_{X|H}(\cdot | 1)\|_{TV}.
\end{aligned}$$

To obtain the upper bound in (4.6), we find $A \geq B$ that

$$\begin{aligned}
& \max_{A, B} A - B \\
& \text{s.t.} \\
& e^{-\epsilon} \leq \frac{A \sum_{\mathbf{x} \in I^+} p_{X|G}(\mathbf{x} | g) + B \sum_{\mathbf{x} \in I^-} p_{X|G}(\mathbf{x} | g)}{A \sum_{\mathbf{x} \in I^+} p_X(\mathbf{x}; G) + B \sum_{\mathbf{x} \in I^-} p_X(\mathbf{x}; G)} \leq e^\epsilon, \\
& e^{-\epsilon} \leq \frac{B \sum_{\mathbf{x} \in I^+} p_{X|G}(\mathbf{x} | g) + A \sum_{\mathbf{x} \in I^-} p_{X|G}(\mathbf{x} | g)}{B \sum_{\mathbf{x} \in I^+} p_X(\mathbf{x}; G) + A \sum_{\mathbf{x} \in I^-} p_X(\mathbf{x}; G)} \leq e^\epsilon, \\
& \text{for } g = 0, 1, \text{ and all } G \in \mathcal{G}_X, \\
& A \geq 0, B \geq 0, A + B = \frac{2}{|\mathcal{Z}|^s}.
\end{aligned}$$

If we fix a $G \in \mathcal{G}_X$, and let A_G and B_G be the corresponding A and B of the above optimization problem, we obtain the linear program:

$$\begin{aligned}
& \max_{A_G, B_G} A_G - B_G \\
& \text{s.t. } a_G(g)A_G + b_G(g)B_G \leq 0, \\
& \quad b_G(g)A_G + a_G(g)B_G \leq 0, \\
& \quad c_G(g)A_G + d_G(g)B_G \leq 0,
\end{aligned}$$

$$d_G(g)A_G + c_G(g)B_G \leq 0,$$

$$\text{for } g = 0, 1,$$

$$A_G \geq 0, B_G \geq 0,$$

$$A_G + B_G = \frac{2}{|\mathcal{Z}|^s},$$

whose solution can be shown to be the following: Let

$$F = \{(a_G(g), b_G(g)), (b_G(g), a_G(g)), \\ (c_G(g), d_G(g)), (d_G(g), c_G(g)), g = 0, 1\}.$$

1. If there exists a pair $(f_1, f_2) \in F$ such that $f_1 > 0$ and $f_2 < 0$, then we find a pair (f_1^*, f_2^*) from F such that $f_1^* > 0$ and $f_2^* < 0$ and that minimizes f_1^*/f_2^* .

We obtain

$$A_G = \frac{-2f_2^*}{|\mathcal{Z}|^s(f_1^* - f_2^*)}, \\ B_G = \frac{2f_1^*}{|\mathcal{Z}|^s(f_1^* - f_2^*)},$$

and therefore $A_G - B_G = \frac{2(e^\epsilon - 1)}{|\mathcal{Z}|^s(f_1^* - f_2^*)}$, which is equivalent to

$$A_G - B_G \\ = \frac{2(e^\epsilon - 1)}{|\mathcal{Z}|^s \max\{|a_G(g) - b_G(g)|, |c_G(g) - d_G(g)| : g \in \{0, 1\}\}}.$$

2. If there does not exist any pairs $(f_1, f_2) \in F$ such that $f_1 > 0$ and $f_2 < 0$, we let $A_G = \frac{2}{|\mathcal{Z}|^s}$ and $B_G = 0$ to obtain $A_G - B_G = \frac{2}{|\mathcal{Z}|^s}$.

From (4.32) and combining the two cases above, we obtain

$$\|p_{Z|H}(\cdot | 0) - p_{Z|H}(\cdot | 1)\|_{TV} \\ \geq \frac{(e^\epsilon - 1)\|p_{X|H}(\cdot | 0) - p_{X|H}(\cdot | 1)\|_{TV}}{\max\{m_G, e^\epsilon - 1\}},$$

for any $G \in \mathcal{G}_X$. Maximizing over all $G \in \mathcal{G}_X$, we obtain the upper bound in (4.6). The proof of the theorem is now complete.

4.C Proof of Theorem 4.2

If $\delta = 0$, the theorem trivially holds. Suppose that $0 < \delta < 1$. It is easy to check that (4.13) defines a probability distribution in $\mathcal{G}_Z$. For any $G \in \mathcal{G}_Z$, let

$$\Gamma_G = \{\mathbf{z} : \ell_{Z|G}(\mathbf{z}) \geq 1\}.$$

We have $\ell_{Z|G^\circ}(\underline{\mathbf{z}}) \leq 1 \leq \ell_{Z|G^\circ}(\bar{\mathbf{z}})$, and

$$\begin{aligned} A_1 &> \frac{p_{Z|G^\circ}(\underline{\mathbf{z}} | 0)}{p_{Z|G^\circ}(\underline{\mathbf{z}} | 1)} = \frac{1}{\ell_{Z|G^\circ}(\underline{\mathbf{z}})} \geq 1, \\ A_2 &> \frac{p_{Z|G^\circ}(\bar{\mathbf{z}} | 1)}{p_{Z|G^\circ}(\bar{\mathbf{z}} | 0)} = \ell_{Z|G^\circ}(\bar{\mathbf{z}}) \geq 1, \end{aligned}$$

which yields the following observations:

- $\ell_{Z|G_{\text{MF}}}(\bar{\mathbf{z}}) = A_2 > 1$ and $\ell_{Z|G^\circ}(\bar{\mathbf{z}}) > 1$.
- $\ell_{Z|G_{\text{MF}}}(\underline{\mathbf{z}}) = 1/A_1 < 1$ and $\ell_{Z|G^\circ}(\underline{\mathbf{z}}) < 1$.
- For $\mathbf{z} \neq \underline{\mathbf{z}}, \bar{\mathbf{z}}$, we have $\ell_{Z|G_{\text{MF}}}(\mathbf{z}) = \ell_{Z|G^\circ}(\mathbf{z})$.

Therefore we have

$$\Gamma_{G^\circ} = \Gamma_{G_{\text{MF}}}. \quad (4.33)$$

For any $G \in \mathcal{G}_Z$, we have

$$\begin{aligned} &R_G(p_{Z|X}, \gamma_G^*) \\ &= \frac{1}{2} \left(\mathbb{P}(Z \in \Gamma_G | G = 0) + \mathbb{P}(Z \in \Gamma_G^c | G = 1) \right) \\ &\geq \frac{(1 - \delta)}{2} \left(\mathbb{P}(Z \in \Gamma_G | G^\circ = 0) \right) \end{aligned}$$

$$+ \mathbb{P}(Z \in \Gamma_G^c \mid G^\circ = 1) \Big) \quad (4.34)$$

$$\geq \frac{(1-\delta)}{2} \Big(\mathbb{P}(Z \in \Gamma_{G^\circ} \mid G^\circ = 0) + \mathbb{P}(Z \in \Gamma_{G^\circ}^c \mid G^\circ = 1) \Big) \quad (4.35)$$

$$= \frac{(1-\delta)}{2} \Big(\mathbb{P}(Z \in \Gamma_{G_{\text{MF}}} \mid G^\circ = 0) + \mathbb{P}(Z \in \Gamma_{G_{\text{MF}}}^c \mid G^\circ = 1) \Big) \quad (4.36)$$

$$= \frac{1}{2} \Big(\mathbb{P}(Z \in \Gamma_{G_{\text{MF}}} \mid G_{\text{MF}} = 0) + \mathbb{P}(Z \in \Gamma_{G_{\text{MF}}}^c \mid G_{\text{MF}} = 1) \Big) \quad (4.37)$$

$$= R_{G_{\text{MF}}}(p_{Z|X}, \gamma_{G_{\text{MF}}}^*),$$

where

- (4.34) follows from the definition of $\mathcal{G}_Z$;
- (4.35) is because Γ_{G° is the optimal critical region that minimizes $R_{G^\circ}(p_{Z|X}, \cdot)$;
- (4.36) follows from (4.33); and
- (4.37) follows from (4.13).

Therefore, $R_{G_{\text{MF}}}(p_{Z|X}, \gamma_{G_{\text{MF}}}^*) = \min_{G \in \mathcal{G}_Z} R_G(p_{Z|X}, \gamma_G^*)$. Furthermore, from (4.35), we obtain $R_{G_{\text{MF}}}(p_{Z|X}, \gamma_{G_{\text{MF}}}^*) = (1-\delta)R_{G^\circ}(p_{Z|X}, \gamma_{G^\circ}^*)$, and the theorem is proved.

Chapter 5

Decentralized Detection with both Inference and Data Privacy

5.1 System Model

Consider the system model in Figure 1.1, the fusion center infers the public hypothesis H from Z . However, it can also use Z to infer G , even though it has not been authorized to do so. At the same time, it may also try to recover X from Z . In this chapter, for simplicity, we consider the case where $H \in \{-1, 1\}$ is a binary hypothesis (our work is easily extended to the multiple hypothesis case), and $G = (G_1, \dots, G_q) \in \mathcal{G} = \{-1, 1\}^q$ is a random vector where each component is binary, i.e., G is a 2^q -ary hypothesis. Our goal is to design privacy mappings at sensors in order to make it difficult for the fusion center to both infer G (inference privacy) and to recover X (data privacy), while allowing it to infer H with reasonable accuracy. In this paper, we do not make any assumptions regarding the conditional independence of sensor observations, which is common in many of the works in decentralized detection [10, 11, 13, 14, 18, 102, 103].

In the example of fall detection, whether a fall happens is the public hypothesis H . Each binary G_i , $i = 1, \dots, q$, in the private hypothesis G can correspond to detecting if the person is performing different activities like running, climbing

stairs, squatting, and so on.

The utility of the network is the probability of inferring H correctly by the fusion center. Inference privacy is measured by the “difficulty” of inferring G . One of our objectives is to determine which inference privacy metric is most suitable for the Internet of Things (IoT) network in Figure 1.1. Furthermore, since some sensors’ observations may be uncorrelated with G , the raw observations from these sensors are transmitted to the fusion center to maximize the utility. There is then leakage of data privacy for these sensors. Therefore, we also require that the local privacy mappings at each sensor incorporate a data privacy mechanism.

5.2 Relationships Between Privacy Metrics

In this section, we study the relationships between privacy metrics to provide insights into the best inference and data privacy metrics for a decentralized IoT network. To relate one privacy metric to another, we introduce the concept of privacy implication and non-guarantee in the following definitions.

Definition 5.1 (Privacy implication). *We say that Type A privacy implies Type B privacy, if for all sequences of probability distributions $(p_{G,X,Z}^i)_{i \geq 1}$ such that $p_{G,X,Z}^i$ satisfies ϵ_i -Type A privacy with $\epsilon_i \downarrow 0$, then $p_{G,X,Z}^i$ satisfies ϵ'_i -Type B privacy with $\epsilon'_i \downarrow 0$.*

In non-technical terms, Definition 5.1 says that arbitrarily strong Type A privacy implies arbitrarily strong Type B privacy. Therefore, to achieve a desired level of Type B privacy, it suffices to ensure that Type A privacy with sufficiently small budget is satisfied. Conversely, we say Type A privacy does not guarantee Type B privacy if the condition in Definition 5.1 does not hold, i.e., there exists a sequence of probability distributions $(p_{G,X,Z}^i)_{i \geq 1}$, such that $p_{G,X,Z}^i$ satisfies ϵ_i -Type A privacy with $\epsilon_i \downarrow 0$, and ϵ'_i -Type B privacy with $\inf_{i \geq 1} \epsilon'_i > 0$.

The following theorem elucidates the relationships between different privacy metrics, which are summarized in Fig. 5.1. Some of these relationships are results

Figure 5.1: Relationships between different privacy metrics for an IoT network with fixed q, s . An arrow $\rightarrow$ means “implies” while $\times \rightarrow$ means “does not guarantee”.

proven in [47, 48], and are reproduced here for completeness.

Theorem 5.1. Consider the decentralized IoT network in Figure 1.1 with $s \geq 1$ sensors and $G = (G_1, \dots, G_q)$. Let $\epsilon \geq 0$. Then, the following holds for $p_{G,X,Z}$.

- (i) [48, Theorem 3] ϵ -information privacy implies 2ϵ -inference differential privacy for all $s \geq 1$.
- (ii) [48, Theorem 3] ϵ -information privacy implies $\frac{\epsilon}{\log 2}$ -average information leakage for all $s \geq 1$.
- (iii) ϵ -inference differential privacy implies $q\epsilon$ -information privacy. If $q \rightarrow \infty$, then inference differential privacy does not guarantee information privacy.
- (iv) ϵ -inference differential privacy implies $\frac{q\epsilon}{\log 2}$ -average information leakage. If $q \rightarrow \infty$, then inference differential privacy does not guarantee average information leakage.

- (v) *Average information leakage does not guarantee information privacy and inference differential privacy.*
- (vi) *ϵ -local differential privacy implies $2s\epsilon$ -information privacy.*
- (vii) *Information privacy does not guarantee local differential privacy.*
- (viii) *Information privacy does not guarantee mutual information privacy.*
- (ix) *Mutual information privacy does not guarantee information privacy.*
- (x) *ϵ -mutual information privacy implies ϵ -average information leakage.*
- (xi) *ϵ -local differential privacy implies $\frac{s\epsilon}{\log 2}$ -mutual information privacy.*
- (xii) *Mutual information privacy does not guarantee local differential privacy.*
- (xiii) *ϵ -local differential privacy yields $(\epsilon + \delta_X)$ -identifiability, where*

$$\delta_X = \max \log p_X(\mathbf{x})/p_X(\mathbf{x}')$$

with the maximum taken over all neighboring $\mathbf{x}, \mathbf{x}' \in \mathcal{X}^s$. Therefore, ϵ -local differential privacy implies ϵ -identifiability if X is restricted to have uniform distribution on $\mathcal{X}^s$. Otherwise, local differential privacy does not guarantee identifiability.

- (xiv) *ϵ -identifiability yields $(\epsilon + \delta_X)$ -local differential privacy. Therefore, ϵ -identifiability implies ϵ -identifiability if X is restricted to have uniform distribution on $\mathcal{X}^s$. Otherwise, identifiability does not guarantee local differential privacy.*

Proof. See Proof 5.A. □

From Theorem 5.1, we see that information privacy implies the other types of inference privacy metrics in (2.8). Although for a fixed number of components q of the private hypothesis $G = (G_1, \dots, G_q)$, inference differential privacy also implies other types of inference privacy metrics including information privacy, it does not guarantee information privacy when $q \rightarrow \infty$.

For data privacy, Theorem 5.1 shows that local differential privacy implies mutual information privacy. As the identifiability metric is essentially the same as local differential privacy up to a fixed constant, we consider only the local differential privacy metric in this paper.

Although local differential privacy implies information privacy for a fixed s number of sensors, this is no longer true if s is not fixed or known in advance. Furthermore, even if s is known *a priori*, Theorem 5.1 suggests that to achieve ϵ -information privacy based solely on preserving local differential privacy, the order of magnitude of the local differential privacy budget has to be not more than ϵ/s . Note that since the definition of local differential privacy does not distinguish between the public hypothesis H or the private hypothesis G , this implies that $p_{H,X,Z}$ also satisfies ϵ -information privacy. If s is large, Proposition 3.1(i) then implies that the Type I and II errors for detecting the public hypothesis H also become large, which is therefore undesirable. Hence, we propose to design the sensors' privacy mappings using both information privacy and local differential privacy constraints, where the local differential privacy budget can be chosen to be sufficiently large to achieve a reasonable utility for H while maintaining strong information privacy for G .

In the subsequent sections, we propose frameworks for designing the local privacy mappings for sensors in a decentralized IoT network under both the parametric and nonparametric cases. These privacy mappings are designed to achieve both information privacy and local differential privacy at the fusion center.

5.3 Parametric Case: Concatenated Privacy Mappings

In this section, we consider the parametric case where $p_{X,H,G}$ is known *a priori*. We first study decentralized detection that preserves only data privacy using the local differential privacy metric. Then we include information privacy as an additional constraint to achieve inference privacy. We prove that information

privacy is immune to post-processing, and local differential privacy is immune to both post-processing and pre-processing. Finally, we propose a local privacy mapping consisting of two concatenated privacy mappings that implement information privacy and local differential privacy mechanisms separately.

5.3.1 Data Privacy using local differential privacy

We know from Theorem 5.1 that local differential privacy is the strongest privacy metric amongst those we studied. Therefore we first consider the case where local differential privacy is adopted as the privacy metric for the IoT network in Figure 1.1 Let $\mathcal{Q}$ denote the set of $p_{Z|X}$ such that

$$p_{Z|X}(\mathbf{z} \mid \mathbf{x}) = \prod_{t=1}^s p^t(z^t \mid x^t), \quad (5.1a)$$

$$\sum_{z^t \in \mathcal{Z}} p^t(z^t \mid x^t) = 1, \quad (5.1b)$$

$$p^t(z^t \mid x^t) \geq 0, \forall x^t \in \mathcal{X}, z^t \in \mathcal{Z}, t = 1, \dots, s. \quad (5.1c)$$

Let $\gamma_H(Z)$ denote the decision rule used by the fusion center to infer the public hypothesis H from the received sensor information Z . Our goal is to

$$\begin{aligned} & \min_{\gamma_H, p_{Z|X} \in \mathcal{Q}} \mathbb{P}(\gamma_H(Z) \neq H) \\ \text{s.t. } & \frac{p^t(z|x)}{p^t(z|x')} \leq e^{\epsilon_{LD}}, \forall z \in \mathcal{Z}, x, x' \in \mathcal{X}, t = 1, 2, \dots, s, \end{aligned} \quad (5.2)$$

where $\epsilon_{LD} \geq 0$ is the local differential privacy budget.

To apply block nonlinear Gauss-Siedel method (GS), we iteratively optimize over the random variables.. For fixed $p^t, t = 1, 2, \dots, s$, (5.2) is a convex optimization over γ_H [82], which can be solved with standard approaches. Then for each $t = 1, \dots, s$, we fix γ_H and p^i where $i \neq t$ and optimize for p^t . This procedure is then repeated until a convergence criterion is met.

Theorem 5.2. Suppose $|\mathcal{Z}| = 2$. Consider optimizing (5.2) over p^t with γ_H and

$p^i, i \neq t$ fixed. The optimal solution is

$$\begin{aligned} p^t(1|x) &= \begin{cases} \frac{1}{1+e^{\epsilon_{LD}}}, & \text{if } f^t(1, x) \geq f^t(2, x), \\ \frac{e^{\epsilon_{LD}}}{1+e^{\epsilon_{LD}}}, & \text{if } f^t(1, x) < f^t(2, x), \end{cases} \\ p^t(2|x) &= \begin{cases} \frac{e^{\epsilon_{LD}}}{1+e^{\epsilon_{LD}}}, & \text{if } f^t(1, x) \geq f^t(2, x), \\ \frac{1}{1+e^{\epsilon_{LD}}}, & \text{if } f^t(1, x) < f^t(2, x), \end{cases} \end{aligned} \quad (5.3)$$

where

$$\begin{aligned} f^t(z, x) &= \sum_{\substack{\mathbf{z} \in \Psi(z) \\ \mathbf{x} \in \{\mathbf{x}: x^t = x\}}} \prod_{i \neq t} p^i(\mathbf{z}|\mathbf{x}) (p_{X,H}(\mathbf{x}, 0) - p_{X,H}(\mathbf{x}, 1)), \end{aligned}$$

with $\Psi(z) = \{\mathbf{z} : \gamma_H(\mathbf{z}) = 1, z^t = z\}$.

Proof. Let $\Gamma = \{\mathbf{z} : \gamma_H(\mathbf{z}) = 1\}$. We have

$$\begin{aligned} &\mathbb{P}(\gamma_H(Z) \neq H) \\ &= p_H(1) + \sum_{\mathbf{z} \in \Gamma} (p_{Z|H}(\mathbf{z}|-1)p_H(-1) - p_{Z|H}(\mathbf{z}|1)p_H(1)) \\ &= p_H(1) + \sum_{\mathbf{z} \in \Gamma, \mathbf{x}} p_{Z|X}(\mathbf{z}|\mathbf{x}) (p_{X,H}(\mathbf{x}, -1) - p_{X,H}(\mathbf{x}, 1)) \\ &= p_H(1) + \sum_{z \in \mathcal{Z}, x \in \mathcal{X}} p^t(z|x) f^t(z, x) \\ &= p_H(1) + \sum_{x \in \mathcal{X}} p^t(1|x) (f^t(1, x) - f^t(2, x)) + \sum_{x \in \mathcal{X}} f^t(2, x). \end{aligned} \quad (5.4)$$

We rewrite (5.2) as the following linear programming problem:

$$\begin{aligned} &\min_{p^t} \sum_{x \in \mathcal{X}} p^t(1|x) (f^t(1, x) - f^t(2, x)) \\ &\text{s.t. } p^t(z|x) - e^{\epsilon_{LD}} p^t(z|x') \leq 0 \\ &\quad p^t(z|x) \geq 0, \sum_z p^t(z|x) = 1, z \in \mathcal{Z}, x, x' \in \mathcal{X}. \end{aligned} \quad (5.5)$$

Without loss of generality, assume $a = p^t(1|1) \geq p^t(1|2) \geq \dots \geq p^t(1||\mathcal{X}) = b$ satisfy the constraints of (5.5). From (5.4), to minimize $\mathbb{P}(\gamma_H(Z) \neq H)$, we have $p^t(1|x) = a$ for $x \in \mathcal{X}_1 = \{x : f^t(2, x) > f^t(1, x)\}$ and $p^t(1|x) = b$ for $x \in \mathcal{X}_2 = \{x : f^t(2, x) \leq f^t(1, x)\}$. Thus, we can simplify (5.5) to

$$\begin{aligned}
& \min_{a,b} \sum_{x \in \mathcal{X}_1} a(f^t(1, x) - f^t(2, x)) \\
& \quad + \sum_{x \in \mathcal{X}_2} b(f^t(1, x) - f^t(2, x)) \\
& \text{s.t. } a - e^{-\epsilon_{LD}} b \geq 0 \\
& \quad a - e^{\epsilon_{LD}} b \leq 0 \\
& \quad (1 - a) - e^{-\epsilon_{LD}} (1 - b) \geq 0 \\
& \quad (1 - a) - e^{\epsilon_{LD}} (1 - b) \leq 0 \\
& \quad a \geq 0, b \geq 0.
\end{aligned}$$

It can be shown that the solution to the above linear program is

$$a = \frac{e^{\epsilon_{LD}}}{1 + e^{\epsilon_{LD}}}, \quad b = \frac{1}{1 + e^{\epsilon_{LD}}},$$

which proves the theorem. □

Theorem 5.2 provides a closed form solution for the local differential privacy mapping at each sensor t when the sensor is constrained to be binary. This is typically the case when the sensor is low-cost and has limited computational resources. The result in Theorem 5.2 thus allows efficient implementation in practice.

5.3.2 Joint Inference and Data Privacy

From Theorem 5.1, as information privacy is one of the strongest inference privacy metrics, we adopt the information privacy metric when designing our privacy

Figure 5.2: Each sensor t 's privacy mapping $p_{Z^t|X^t}^t = p_{Y^t|X^t}^t \cdot p_{Z^t|Y^t}^t$ consists of two privacy mappings concatenated together.

mechanism. To achieve joint inference and data privacy, we consider

$$\begin{aligned}
 & \min_{\gamma_H, p_{Z|X}} \mathbb{P}(\gamma_H(Z) \neq H), \\
 & e^{-\epsilon_I} \leq \frac{p_{Z|G}(\mathbf{z}|g_i)}{p_Z(\mathbf{z})} \leq e^{\epsilon_I}, i = 1, \dots, q, \\
 & \forall \mathbf{z} \in \mathcal{Z}^s, \forall \mathbf{g} = \{g_1, g_2, \dots, g_q\} \in \mathcal{G}, \\
 & \frac{p^t(z|x)}{p^t(z|x')} \leq e^{\epsilon_{LD}}, \forall z \in \mathcal{Z}, x, x' \in \mathcal{X}, t = 1, 2, \dots, s,
 \end{aligned} \tag{P0}$$

where $\epsilon_I > 0$ and ϵ_{LD} are the information privacy budget and local differential privacy budget, respectively.

Since (P0) is a NP-complete problem [98], we seek to find suboptimal solutions rather than solving (P0) directly. Similar to the work in [57], we break the privacy mapping $p_{Z|X}$ in (P0) into two concatenated stages as shown in Figure 5.2, where sensor observations $X \in \mathcal{X}^s$ are first mapped to $Y \in \mathcal{Y}^s$, which is then mapped to

$Z \in \mathcal{Z}^s$, i.e., the mappings $p_{Y|X}$ and $p_{Z|Y}$ satisfy

$$p_{Y|X}(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^s p_1^t(y^t|x^t), \quad p_{Z|Y}(\mathbf{z}|\mathbf{y}) = \prod_{t=1}^s p_2^t(z^t|y^t),$$

and for all $t = 1, \dots, s$,

$$p_1^t(y^t|x^t) \geq 0, \quad \sum_{y^t} p_1^t(y^t|x^t) = 1, \quad \forall y^t \in \mathcal{Y}, x^t \in \mathcal{X},$$

$$p_2^t(z^t|y^t) \geq 0, \quad \sum_{z^t} p_2^t(z^t|y^t) = 1, \quad \forall z^t \in \mathcal{Z}, y^t \in \mathcal{Y}.$$
(5.6)

The local privacy mapping for each sensor t is given by

$$p^t(z | x) = \sum_{y \in \mathcal{Y}} p_2^t(z|y) p_1^t(y|x). \quad (5.7)$$

We propose the following two architectures:

1. Information-Local differential privacy (ILL): the mapping from $X \in \mathcal{X}^s$ to $Y \in \mathcal{Y}^s$ preserves information privacy, while the mapping from $Y \in \mathcal{Y}^s$ to $Z \in \mathcal{Z}^s$ preserves local differential privacy.
2. Local differential-Information Privacy (LIP): the mapping from $X \in \mathcal{X}^s$ to $Y \in \mathcal{Y}^s$ preserves local differential privacy, while the mapping from $Y \in \mathcal{Y}^s$ to $Z \in \mathcal{Z}^s$ preserves information privacy.

In the following Proposition 5.1,5.2, we show that this two-stage approach achieves joint inference and data privacy. But first, we discuss how to optimize for the privacy mappings in practice.

In the ILL architecture, we find mappings $p_{Y|X}(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^s p_1^t(y^t|x^t)$ and

$p_{Z|Y}(\mathbf{z}|\mathbf{y}) = \prod_{t=1}^s p_2^t(z^t|y^t)$ satisfying

$$\begin{aligned}
 & \min_{\gamma_H, p_{Y|X}, p_{Z|Y}} \mathbb{P}(\gamma_H(Z) \neq H) \\
 & e^{-\epsilon_I} \leq \frac{p_{Y|G}(\mathbf{y}|\mathbf{g})}{p_Y(\mathbf{y})} \leq e^{\epsilon_I}, \forall \mathbf{g} \in \mathcal{G}, \mathbf{y} \in \mathcal{Y}^s, \\
 & \frac{p_2^t(z|y)}{p_2^t(z|y')} \leq e^{\epsilon_{LD}/2}, \forall z \in \mathcal{Z}, y, y' \in \mathcal{Y}, t = 1, 2, \dots, s, \\
 & p_{Y|X}, p_{Z|Y} \text{ satisfy (5.6)}.
 \end{aligned} \tag{P1}$$

To solve the problem (P1), we first consider the information privacy subproblem:

$$\min_{\gamma_H, p_{Y|X}} \mathbb{P}(\gamma_H(Y) \neq H), \tag{5.8a}$$

$$e^{-\epsilon_I} \leq \frac{p_{Y|G}(\mathbf{y}|\mathbf{g})}{p_Y(\mathbf{y})} \leq e^{\epsilon_I}, \forall \mathbf{g} \in \mathcal{G}, \mathbf{y} \in \mathcal{Y}^s, \tag{5.8b}$$

From Theorem 3.2, to meet the constraint (5.8b), it suffices to ensure that

$$\begin{aligned}
 & \min_{\mathbf{g} \in \mathcal{G} \setminus \{-\mathbf{1}\}, \gamma_G} \frac{1}{2} \left(\mathbb{P}(\gamma_G(Y) = \mathbf{g} \mid G = -\mathbf{1}) \right. \\
 & \quad \left. + \mathbb{P}(\gamma_G(Y) = -\mathbf{1} \mid G = \mathbf{g}) \right) \geq \theta,
 \end{aligned} \tag{5.9}$$

where $\theta = (1 - c_G(1 - e^{-\epsilon_I/2}))/2$ with

$$\begin{aligned}
 c_G &= \min_{\mathbf{g} \neq -\mathbf{1}} \left\{ \mathbb{P} \left(Y \in \arg \min_{\mathbf{y} \in \mathcal{Y}^s} \ell_{\mathbf{g}}(\mathbf{y}) \mid G = -\mathbf{1} \right), \right. \\
 & \quad \left. \mathbb{P} \left(Y \in \arg \max_{\mathbf{y} \in \mathcal{Y}^s} \ell_{\mathbf{g}}(\mathbf{y}) \mid G = \mathbf{g} \right) \right\}, \\
 \ell_{\mathbf{g}}(\mathbf{y}) &= \frac{p_{Y|G}(\mathbf{y} \mid \mathbf{g})}{p_{Y|G}(\mathbf{y} \mid -\mathbf{1})}.
 \end{aligned}$$

By using the constraint (5.9), we reduce the $2|\mathcal{G}| \times |\mathcal{Y}|^s$ constraints in (5.8b) to a single (but weaker) constraint, which is easier to optimize in practice. A GS variant of (5.8) similar to that used for solving (5.2) can then be used to find the privacy mapping $p_{Y|X}$ as follows.

- (i) For a fixed privacy mapping $p_{Y|X}$, we first find the optimal fusion center decision rule γ_H .
- (ii) For each sensor $t = 1, \dots, s$ in sequential order, we optimize for sensor t 's information privacy mapping $p_t^1(y_t|x_t)$, with γ_H and the privacy mappings of all other sensors $p_{\setminus t}^1 = \prod_{j \neq t} p_j^1$ fixed. Let the set of sensor t 's information privacy mapping be Φ . The optimization is done by solving the following linear program:

$$\begin{aligned}
 & \min_{\nu_\phi} \sum_{\phi \in \Phi} \nu_\phi L_H(\phi) \\
 & \text{s.t.} \sum_{\phi \in \Phi} \nu_\phi \min_{\gamma_G} R_{\mathbf{g}}(\phi \cdot p_{\setminus t}^1, \gamma_G) \geq \theta, \quad \forall \mathbf{g} \in \mathcal{G} \setminus \{\mathbf{0}\}, \\
 & \sum_{\phi \in \Phi} \nu_\phi = 1, \quad \nu_\phi \geq 0, \quad \forall \phi \in \Phi.
 \end{aligned}$$

where $L_H(\phi)$ is $\mathbb{P}(\gamma_H(Y) \neq H)$ when the privacy mapping $p_{Y|X} = \phi \cdot p_{\setminus t}^1$. Note that from [68, Section II.B], the decision rule $\gamma_G = \arg \min_{\gamma} R_{\mathbf{g}}(\phi \cdot p_{\setminus t}^1, \gamma)$ is given by

$$\gamma_G(\mathbf{y}) = \begin{cases} 1, & \text{if } \ell_{\mathbf{g}}(\mathbf{y}) \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

The above two steps are iterated until a convergence criterion (e.g., when the L_1 norm of the difference in the mapping $p_{Y|X}$ between two successive iterations is less than a small constant) is met.

In the second stage, we consider the local differential privacy subproblem:

$$\begin{aligned}
 & \min_{\gamma_H, p_{Z|Y}} \mathbb{P}(\gamma_H(Z) \neq H), \\
 & \frac{p_2^t(z|y)}{p_2^t(z|y')} \leq e^{\epsilon_{LD}/2}, \quad \forall z \in \mathcal{Z}, \quad y, y' \in \mathcal{Y}, \quad t = 1, 2, \dots, s.
 \end{aligned} \tag{5.10}$$

If $|\mathcal{Z}| = 2$, the GS solution follows from Theorem 5.2. If $|\mathcal{Z}| > 2$, we can use a

standard linear program solver [104] for (5.10) (see the discussion leading to (5.5) on how to formulate this linear program).

Similarly, for the LIP architecture, we consider the following optimization problem:

$$\begin{aligned}
 & \min_{\gamma_H, p_{Y|X}, p_{Z|Y}} \mathbb{P}(\gamma_H(Z) \neq H), \\
 & \frac{p_1^t(y|x)}{p_1^t(y|x')} \leq e^{\epsilon_{LD}}, \forall y \in \mathcal{Y}, x, x' \in \mathcal{X}, t = 1, 2, \dots, s, \\
 & e^{-\epsilon_I} \leq \frac{p_{Z|G}(\mathbf{z}|\mathbf{g})}{p_Z(\mathbf{z})} \leq e^{\epsilon_I}, \forall \mathbf{g} \in \mathcal{G}, \mathbf{z} \in \mathcal{Z}^s, \\
 & p_{Y|X}, p_{Z|Y} \text{ satisfy (5.6)}.
 \end{aligned} \tag{P2}$$

Solving (P2) can be done in a similar fashion as (P1).

We next show that the concatenation of information privacy mapping with local differential privacy mapping achieves joint information and local privacy in both the ILL and LIP architectures.

Proposition 5.1. *Let $\epsilon_I, \epsilon_{LD} \geq 0$. Suppose that $p_{G,X,Y}$ satisfies ϵ_I -information privacy and $p_{G,Y,Z}$ satisfies $\epsilon_{LD}/2$ -local differential privacy. Then, the following holds.*

- (a) *For any randomized mapping $p_{Z|Y}$, $p_{G,X,Z}$ satisfies ϵ_I -information privacy.*
- (b) *For any randomized mapping $p_{Y|X}$, $p_{G,X,Z}$ satisfies ϵ_{LD} -local differential privacy.*

Proof. (a) For any $\mathbf{z} \in \mathcal{Z}^s$ and $\mathbf{g} \in \mathcal{G}$, we have

$$\frac{p_{Z|G}(\mathbf{z}|\mathbf{g})}{p_Z(\mathbf{z})} = \frac{\sum_{\mathbf{y}} p_{Z|Y}(\mathbf{z}|\mathbf{y}) p_{Y|G}(\mathbf{y}|\mathbf{g})}{\sum_{\mathbf{y}} p_{Z|Y}(\mathbf{z}|\mathbf{y}) p_Y(\mathbf{y})}.$$

Since $e^{-\epsilon_I} \leq \frac{p_{Y|G}(\mathbf{y}|\mathbf{g})}{p_Y(\mathbf{y})} \leq e^{\epsilon_I}$ for all $\mathbf{y} \in \mathcal{Y}^s$, we obtain $e^{-\epsilon_I} \leq \frac{p_{Z|G}(\mathbf{z}|\mathbf{g})}{p_Z(\mathbf{z})} \leq e^{\epsilon_I}$.

- (b) Consider any sensor t . For any $y, y' \in \mathcal{Y}$ and $z \in \mathcal{Z}$, we have $e^{-\epsilon_{LD}/2} \leq$

$\frac{p_2^t(z|y)}{p_2^t(z|y')} \leq e^{\epsilon_{LD}/2}$. Therefore, for any $x, x' \in \mathcal{X}$, we then have

$$\begin{aligned} & \frac{p^t(z|x)}{p^t(z|x')} \\ &= \frac{\sum_y p_2^t(z|y)p_1^t(y|x)}{\sum_y p_2^t(z|y)p_1^t(y|x')} \\ &\leq \frac{\sum_y e^{\epsilon_{LD}/2} p_2^t(z|y')p_1^t(y|x)}{\sum_y e^{-\epsilon_{LD}/2} p_2^t(z|y')p_1^t(y|x')} \\ &= e^{\epsilon_{LD}}, \end{aligned}$$

for a fixed $y' \in \mathcal{Y}$.

The proposition is now proved. $\square$

Proposition 5.1 shows that joint information privacy for G and local differential privacy for X are preserved in the ILL architecture. In the LIP architecture, it is clear that information privacy for G is preserved since this is an explicit constraint in (P2). local differential privacy preservation follows from [53, Proposition 2.1], which is reproduced below for completeness.

Proposition 5.2. *Let $\epsilon_{LD} \geq 0$. Suppose that $p_{G,X,Y}$ satisfies ϵ_{LD} -local differential privacy. Then for any randomized mapping $p_{Z|Y}$, $p_{G,X,Z}$ satisfies ϵ_{LD} -local differential privacy.*

Proof. For any sensor t , $z \in \mathcal{Z}$, $x, x' \in \mathcal{X}$, we have

$$\frac{p^t(z|x)}{p^t(z|x')} = \frac{\sum_y p_2^t(z|y)p_1^t(y|x)}{\sum_y p_2^t(z|y)p_1^t(y|x')} \leq e^{\epsilon_{LD}},$$

since $\frac{p_1^t(y|x)}{p_1^t(y|x')} \leq e^{\epsilon_{LD}}$. The proposition is now proved. $\square$

5.4 Nonparametric Case: Empirical Risk Optimization

In many IoT applications, knowing the joint distribution of (H, G) and the sensor observations is impractical due to difficulties in accurately modeling this distribution. Therefore, we adopt a nonparametric approach to design the privacy mapping $p_{Z|X}$

based on a given set of n independently indentially distributed (i.i.d.) training data points $(X = x_i, H = h_i, G = \mathbf{g}_i)_{i=1}^n$.

Following Chapter 3, let ϕ be a loss function, $\mathcal{H}$ be a reproducing kernel Hilbert space with kernel $\kappa(\cdot, \cdot)$, kernel inner product $\langle \cdot, \cdot \rangle$, and associated norm $\|\cdot\|$. We restrict the rule used by the fusion center to infer H and G based on $Z = \mathbf{z}$ to be of the form $\langle w, \Phi(\mathbf{z}) \rangle$, where $\Phi(\mathbf{z}) = \kappa(\cdot, \mathbf{z})$ is the feature map. We seek to minimize the empirical ϕ -risk of deciding H while preserving information privacy.

We consider the following optimization problem:

$$\min_{w \in \mathcal{H}, p_{Z|X}} F(w, p_{Z|X}), \quad (5.11a)$$

$$\text{s.t. } \min_{v \in \mathcal{H}} \hat{R}_{\mathbf{g}}(v, p_{Z|X}) \geq \theta, \quad \forall \mathbf{g} \in \mathcal{G} \setminus \{-\mathbf{1}\}, \quad (5.11b)$$

$$\frac{p^t(z|x)}{p^t(z|x')} \leq e^{\epsilon_{LD}},$$

$$\forall z \in \mathcal{Z}, x, x' \in \mathcal{X}, t = 1, \dots, s, \quad (5.11c)$$

where

$$F(w, p_{Z|X}) = \frac{1}{n} \sum_{i=1}^n \phi(h_i \langle w, \Phi_Q(\mathbf{x}_i) \rangle) + \frac{\lambda}{2} \|w\|^2,$$

$$\hat{R}_{\mathbf{g}}(v, p_{Z|X}) = \frac{1}{2} \sum_{\mathbf{g}' \in \{-\mathbf{1}, \mathbf{g}\}} \sum_{i \in \mathcal{S}_{\mathbf{g}'}} \frac{\phi(g'_i \langle v, \Phi_Q(\mathbf{x}_i) \rangle)}{|\mathcal{S}_{\mathbf{g}'}|} + \frac{\lambda}{2} \|v\|^2,$$

$$\Phi_Q(\mathbf{x}) = \sum_{\mathbf{z} \in \mathcal{Z}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) \Phi(\mathbf{z}),$$

$\lambda > 0, \theta > 0$ is called the information privacy threshold,

$$g'_i = \begin{cases} -1, & \text{if } \mathbf{g}_i = -\mathbf{1}, \\ 1, & \text{otherwise,} \end{cases}$$

and

$$\mathcal{S}_{\mathbf{g}'} = \{i \in \{1, \dots, n\} : \mathbf{g}_i = \mathbf{g}'\}.$$

Note that $F(\cdot, \cdot)$ is the empirical ϕ -risk of detecting H while $\hat{R}_{\mathbf{g}}(\cdot, \cdot)$ is the empirical (normalized) ϕ -risk of distinguishing between $G = -\mathbf{1}$ and $G = \mathbf{g}$. For convenience, we call (5.11) the Empirical information and local differential PrIvaCy (EPIC) optimization.

For a detailed explanation of how the above optimization framework is derived, we refer the reader to Chapter 3. Briefly, we seek to find $p_{Z|X}$ such that the empirical risk for detecting G under *any decision rule* adopted by the fusion center is above the information privacy threshold θ . The mapping $p_{Z|X}$ is also required to satisfy ϵ_{LD} -local differential privacy in the constraint (5.11c).

From Theorem 3.1, for each ϵ_{LD} , by choosing θ appropriately, we can achieve ϵ_I -information privacy for any $\epsilon_I > 0$ under mild technical assumptions. However, this trades off the detection error rate for H . Therefore, we adopt the same two-step procedure in Chapter 3:

- (i) Determine the largest information privacy threshold θ^* achievable under additional constraints on $p_{Z|X}$ to ensure that the error rate of inferring H remains reasonable. This is achieved through an iterative block Gauss-Seidel method.
- (ii) Set a $r \in (0, 1)$, which we call the *information privacy threshold ratio*, set $\theta = r\theta^*$ in (5.11b) and use an iterative block Gauss-Seidel method to solve (5.11).

For the details of this two-step procedure, we again refer to Chapter 3. The only difference with the procedure in Chapter 3 is that now we have the additional linear inequality constraints (5.11c), which can be easily handled since each step in the block Gauss-Seidel method remains as a convex optimization problem.

5.5 Numerical Results

In this section, we carry out simulations to verify the performance of our proposed approaches. We also perform experiments and comparisons on the OPPORTUNITY Activity Recognition Data Set [92]. In our simulations, we consider binary public hypothesis H and private hypothesis G . To evaluate the performance, we compute the Bayes probability 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. The Bayes error of detecting H reflects the utility of our method, while the Bayes error of detecting G reflects the inference privacy of the private hypothesis G . Data privacy of the sensor t 's observation X^t is quantified by the mutual information $I(X^t; Z^t)$.

5.5.1 Parametric Case Study

Consider a network of 6 sensors and a fusion center. Suppose that $\mathcal{X} = \{1, 2, \dots, 16\}$ and $\mathcal{Z} = \{1, 2\}$. We set the correlation coefficient between the public hypothesis H and private hypothesis G to be 0.2. We assume that each sensor has identical joint distribution as shown in Figure 4.2.

In Figure 5.5, we let the information privacy budget be fixed at $\epsilon_I = 0.01$ and 0.18, and vary the local differential privacy budget ϵ_{LD} . We see that if ϵ_I is small, ILL is better at inferring the public hypothesis H while achieving a similar detection error for the private hypothesis G when compared to LIP. This is because ILL first sanitizes the sensor observations X for information privacy before applying a local differential privacy mapping, which allows it better control over sanitization of statistical information needed to infer G but keeping information for inferring H . On the other hand, if ϵ_I is large, LIP infers H with better accuracy. We also compare with the approach that uses only a local differential privacy constraint (i.e., the information privacy constraint in (P0) is removed), which we call Local Differential Privacy (LDP) in the left drawing in Figure 5.5. Without

Figure 5.3: First stage privacy mapping at sensor 1 with ILL.

Figure 5.4: Second stage privacy mapping at sensor 1 with ILL.

any constraint on ϵ_I , we see that LDP gives poor information privacy protection for G .

In Figure 5.6, we fix $\epsilon_{LD} = 0.07$, while varying ϵ_I . We see that when ϵ_{LD} is small, the Bayes error of detecting H is large regardless of the value of ϵ_I . This aligns with our discussion after Theorem 5.1 that we should not use local

Figure 5.5: Bayes error for detecting H and G under LIP and ILL for fixed privacy threshold ratio r and varying local differential privacy budget ϵ_{LD} .

Figure 5.6: Bayes error for detecting H and G under LIP and ILL for fixed local differential privacy budget ϵ_{LD} and varying ϵ_I .

differential privacy to achieve inference privacy for the private hypothesis G as this approach also leads to a poor inference performance for the public hypothesis H .

We next consider the case where sensor 1's observations are independent of

 (a) Mutual information $I(H; Z)$ and $(G; Z)$.

 (b) Mutual information $I(X^1; Z^1)$ and $(X^2; Z^2)$.

 Figure 5.7: Mutual information with $\epsilon_I = 0.15$ and varying ϵ_{LD} for ILL, LIP, InP and LDP.

G with marginal conditional distribution under H same as the joint distribution shown in Figure 4.2. All other sensors follow the distribution in Figure 4.2. In Figure 5.7, we fix $\epsilon_I = 0.15$ and vary ϵ_{LD} to illustrate the mutual information between different quantities. We also compare with the approach that uses only an Information Privacy (InP) (i.e., the local differential privacy constraint in (P0) is removed). From Fig. 5.7a, we observe that both ILL and LIP yield sanitized information Z that have a high mutual information with the public hypothesis H , and low mutual information with the private hypothesis G . However, with LDP the mutual information $I(H; Z)$ and $I(G; Z)$ are both much higher compared to other methods, since it does not protect the information privacy of G .

In Fig. 5.7b, we compare the mutual informations $I(X^1; Z^1)$ and $I(X^2, Z^2)$ under different privacy architectures. We see that $I(X^1; Z^1)$ under ILL and LIP are much lower than that under InP. In particular, InP does not achieve good data privacy for X^1 since the information privacy constraint only removes statistical information in X^1 related to G , which in this case is none as X^1 is independent of G . This example illustrates the need to include both inference and data privacy constraints in our privacy mapping design. We also see that $I(X^2; Z^2)$ under both ILL and LIP is lower than that under InP, but converges to that of InP as ϵ_{LD} becomes bigger.

5.5.2 Nonparametric Case Study: Simulations

In this subsection, we consider the nonparametric case where the underlying sensor distributions are unknown. We perform simulations to provide insights into the performance of our proposed EPIC approach in (5.11).

For simplicity, we use the count kernel in our simulations, which can be computed with a time complexity of $\mathcal{O}(s|\mathcal{Y}|)$. We choose the logistic loss function as the loss function ϕ in our simulations.

Consider a network of 4 sensors and a fusion center. Each sensor observation x_i^t is generated according to Table 5.1, where n_i^t is uniformly distributed over

$\{-2, -1, 0, +1, +2\}$. The sensor observation space is $\mathcal{X} = \{-5, -4, \dots, 5\}$, and the local decision space is chosen to be $\mathcal{Z} = \{1, 2\}$. Conditioned on (H, G) , sensor observations are independent of each other. We generate 40 i.i.d. training samples, and apply our proposed approach on the training data to learn the privacy mapping $p_{Z|X}$.

(h_i, g_i)	$(-1, -1)$	$(-1, 1)$	$(1, -1)$	$(1, 1)$
x_i^t	$-3 + n_i^t$	$-1 + n_i^t$	$1 + n_i^t$	$3 + n_i^t$

Table 5.1: Sensor observation for different realizations of (H, G) . Figure 5.8: Bayes error for detecting H and G , and mutual information between X^1 and Z^1 with different local differential privacy budget ϵ_{LD} .

We first compare the performance of three methods: our proposed EPIC, Nonparametric Privacy-aware Optimization (NPO) a detection method with only information privacy metric Chapter 3, and Local Differential Privacy (E-LDP) a detection method with only local differential privacy metric (i.e., solving (5.11a) without (5.11b)).

Fig. 5.8 demonstrates how ϵ_{LD} , the local differential privacy budget, affects the inference privacy, data privacy and utility of these methods. In the simulation, we fix the information privacy threshold ratio $r = 0.999$ when setting $\theta = r\theta^*$ in

(5.11b), and the correlation coefficient between H and G is 0.2. We observe that when ϵ_{LD} is small, the performance of EPIC is close to the performance of E-LDP, where the Bayes error rates of both hypotheses are close to 0.5. This is in line with Theorem 5.1(vi): a small local differential privacy budget implies information privacy for both hypotheses. With the increase of ϵ_{LD} , the performance of EPIC approaches the performance of NPO, where the error rate of H is low, while that for G is high. However, with E-LDP, the error rate of G also decreases with increasing ϵ_{LD} , which leads to inference privacy leakage. When analyzing the data privacy leakage, we find that $I(X^1; Z^1)$ stays high with NPO, whereas EPIC achieves a reasonable $I(X^1; Z^1)$ by choosing ϵ_{LD} to be around 5.

Figure 5.9: Bayes error probability of detecting H and G with varying correlation coefficient between H and G .

Fig. 5.9 shows how the correlation between H and G affects their Bayes error detection rate. For EPIC, we set $\epsilon_{LD} = 5$, and for LDP, we find a local differential privacy budget for each correlation coefficient tested that achieves the same error rate for G as in EPIC. We observe that for the same correlation coefficient, the error rate for H is higher in LDP compared to that in EPIC. This demonstrates our claim that local differential privacy should not be used to imply information privacy, as it can severely impact the detection error rate for H as well.

5.5.3 Nonparametric Case Study: OPPORTUNITY Data Set and Adult Data Set

We test our nonparametric EPIC framework on the OPPORTUNITY Activity Recognition Data Set [92] and Adult Data Set [105] available at UCI Repository [87], and compare its performance with RUCA [88], DCA [106] and MDR [89]. In EPIC, we set the local decision space of each sensor to be $\mathcal{Z} = \{1, 2\}$.

I

Data Preprocessing

In the OPPORTUNITY Activity Recognition Data Set, measurements from motion sensors including on-body sensors, sensors attached to objects, and ambient sensors like switches, are recorded while a person performs a series of typical daily activities. In this experiment, our public hypothesis H is whether the person is standing or walking, while the private hypothesis G is whether the person is touch a drawer or dishwasher. We used data from the ‘S2-Drill’ dataset, and sklearn [107] to select $s = 15$ sensors that are the 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 $n = 80$ instances of training data, and 3427 instances of testing data.

In the Adult Data Set, basic information of a certain population such as age, work class, education, income, marriage status was collected. In this thesis, the public hypothesis is $\text{income} > 50K$ and $\text{income} \leq 50K$ and the private hypotheses are married (‘Married-civ-spouse’, ‘Married-spouse-absent’ and ‘Married-AFspouse’), used to be married (‘Separated’, ‘Divorced’ and ‘Widowed’) and Never married (Never-married). We selected age, workclass, education-num, race, sex as the features, which represents sensor observation X in our problem formulation. Although the data is not collected from sensor network, we can apply our method similarly to demonstrate the performance of our proposed method. We discretized the continuous data to 5 bins and perform one-hot encoding to categorical data. We

select $n = 120$ instances of training data where both public and private hypotheses are evenly distributed and 15050 testing data.

Comparison Benchmarks

As comparison benchmarks, we compare our method to the above-mentioned NPO and E-LDP and the following additional methods:

1. NPO Chapter 3, which includes only information privacy.
2. The centralized approaches RUCA [88], DCA [106] and MDR [89], which require that all sensors send their observations to a central data curator that then applies an overall privacy mapping. Note that since the mapping in RUCA, DCA and MDR are deterministic, they do not provide any local differential privacy protection.
3. Sensors do not apply any privacy mapping and sends their raw observations, i.e., $Z = X$. In this case, no local differential privacy protection is available, while some information privacy maybe possible depending on the underlying distribution $p_{X|G}$. This serves as a benchmark to show the intrinsic error probability achievable.

Similar to Chapter 3, to estimate the privacy budgets achieved by each method, we compute

$$\hat{\epsilon}_I = \max_{g \in \mathcal{G}, \mathbf{z} \in \mathcal{Z}^s} \left| \log \frac{\hat{p}_{G,Z}(g, \mathbf{z})}{\hat{p}_G(g) \hat{p}_Z(\mathbf{z})} \right|, \quad (5.12)$$

$$\hat{\epsilon}_{LD} = \max_{z \in \mathcal{Z}, x, x' \in \mathcal{X}, t \in \{1, \dots, s\}} \log \frac{p^t(z | x)}{p^t(z | x')} \quad (5.13)$$

as estimates for the information privacy and local differential privacy budgets respectively. Here, $\hat{p}_A(a)$ is the empirical probability of the event $\{A = a\}$. Note that a smaller $\hat{\epsilon}$ implies stronger information privacy and a smaller $\hat{\epsilon}_{LD}$ implies stronger local differential privacy. We see that $\hat{\epsilon}_{LD} = \infty$ for RUCA, MDR, and the case $Z = X$.

Result and Discussion

From Table 5.2 and 5.3, we observe that EPIC achieves the lowest information privacy and local differential privacy budgets compared to all the other benchmarks. Compared to NPO, it has similar information privacy budget but significantly lower local differential privacy budget since NPO does not consider any data privacy constraints. Compared to E-LDP, it has similar local differential privacy budget, but a significantly lower information privacy constraint. Due to the local differential privacy constraint, we see that EPIC has the highest error rate for detecting H amongst all the methods, which is the price it pays for having the least privacy leakage.

Table 5.2: Detection errors using the OPPORTUNITY Activity Recognition Data Set.

Detection Method	H	G	$\hat{\epsilon}_I$	$\hat{\epsilon}_{LD}$
EPIC ($r = 0.99, \epsilon_{LD} = 1$)	10.91%	43.65%	0.46	0.81
NPO ($r = 0.99$)	10.53%	43.17%	0.47	2.22
E-LDP ($\epsilon_{LD} = 1$)	10.09%	7.31%	8.58	0.91
MDR	12.56%	40.16%	1.02	∞
DCA	10.88%	42.62%	0.88	
RUCA ($\rho_p = 1$)	10.23%	45.73%	0.67	
RUCA ($\rho_p = 100$)	10.10%	43.01%	0.69	
RUCA ($\rho_p = 1000$)	10.10%	43.78%	0.69	
$Z = X$	10.05%	5.57%	9.14	

Table 5.3: Detection errors using the Adult Data Set.

Detection Method	H	G	$\hat{\epsilon}_I$	$\hat{\epsilon}_{LD}$
EPIC ($r = 0.99, \epsilon_{LD} = 1$)	37.69%	62.25%	0.79	0.93
NPO ($r = 0.99$)	37.67%	62.14%	0.82	2.36
E-LDP ($\epsilon_{LD} = 1$)	36.11%	32.91%	9.48	0.97
MDR	37.57%	64.02%	1.68	∞
DCA	38.38%	56.33%	2.47	
RUCA ($\rho_p = 1$)	41.24%	65.25%	1.61	
RUCA ($\rho_p = 100$)	41.10%	64.14%	1.66	
RUCA ($\rho_p = 1000$)	40.67%	65.86%	1.61	
$Z = X$	34.05%	30.48%	15.33	

5.A Proof of Theorem 5.1

To show privacy non-guarantee, it suffices to provide an example of a sequence of joint distributions that do not satisfy Definition 5.1. We first present such an example that parts of the proof of Theorem 5.1 utilize.

Example 5.1. *If the random variables $U \in \mathcal{U}$ and $V \in \mathcal{V}$ satisfy the joint distribution as shown in Table 5.4, then we have*

$$\lim_{\alpha \rightarrow 0} I(V; U) = \lim_{\alpha \rightarrow 0} \left\{ p_{V,U}(0, 0) \log \frac{p_{V,U}(0, 0)}{p_V(0)p_U(0)} \right. \quad (5.14)$$

$$\left. + \sum_{i,j \neq 0} p_{V,U}(i, j) \log \frac{p_{V,U}(i, j)}{p_V(i)p_U(j)} \right\}$$

$$= \lim_{\alpha \rightarrow 0} \left\{ \alpha \log \frac{1}{\alpha} + (1 - \alpha) \log \frac{1}{1 - \alpha} \right\} \quad (5.15)$$

$$= 0,$$

and

$$\lim_{\alpha \rightarrow 0} \frac{p_{V,U}(0, 0)}{p_V(0)p_U(0)} = \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} = \infty, \quad (5.16)$$

$$\max_{v,u,u'} \frac{p_{V|U}(v|u)}{p_{V|U}(v|u')} = \frac{p_{V|U}(0|0)}{p_{V|U}(0|1)} = \infty. \quad (5.17)$$

$p_{V,U}$		V			
		0	1	$\dots$	$ \mathcal{V} - 1$
U	0	α	0	0	0
	1	0	$\frac{1-\alpha}{(\mathcal{U} -1)(\mathcal{V} -1)}$		
	$\vdots$	0			
	$ \mathcal{U} - 1$	0			

Table 5.4: The joint distribution of V and U , where $\alpha \in [0, 1]$.

We now proceed with the proof of Theorem 5.1.

(i-ii) These claims follow from [48, Theorem 3].

(iii-iv) For a fixed q , since $p_{G,X,Z}$ satisfies ϵ -differential privacy, for any $\mathbf{z} \in$

$\mathcal{Z}^s, \mathbf{g} \in \mathcal{G}$, we have

$$\frac{p_{Z|G}(\mathbf{z}|\mathbf{g}_1)}{p_{Z|G}(\mathbf{z}|\mathbf{g}_2)} \leq e^{q\epsilon},$$

for any $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}$. Therefore, we have

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

Thus $p_{G,X,Z}$ satisfies $q\epsilon$ -information privacy. Together with (ii), we obtain that $p_{G,X,Z}$ satisfies $q\epsilon/\log 2$ -average information leakage.

If $q \rightarrow \infty$, [48, Theorem 4] gives an example that shows differential privacy does not guarantee average information leakage. Together with (ii), it implies that differential privacy does not guarantee information privacy.

(v) Substitute G for U and Z for V in Example 5.1, then we get from (5.15) and (5.16), that average information leakage does not guarantee information privacy. From (iii), we also obtain that average information leakage does not guarantee differential privacy.

(vi) Since $p_{G,X,Z}$ satisfies ϵ -local differential privacy, for any $\mathbf{x}_0, \mathbf{x} \in \mathcal{X}^s$, and $\mathbf{z} \in \mathcal{Z}^s$, we have

$$e^{-s\epsilon} p_{Z|X}(\mathbf{z}|\mathbf{x}_0) \leq p_{Z|X}(\mathbf{z}|\mathbf{x}) \leq e^{s\epsilon} p_{Z|X}(\mathbf{z}|\mathbf{x}_0).$$

Then for any $\mathbf{g}_1, \mathbf{g}_2 \in \mathcal{G}, \mathbf{z} \in \mathcal{Z}^s$, we have

$$\begin{aligned} & \frac{p_{Z|G}(\mathbf{z}|\mathbf{g}_1)}{p_{Z|G}(\mathbf{z}|\mathbf{g}_2)} \\ &= \frac{\sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G}(\mathbf{x}|\mathbf{g}_1)}{\sum_{\mathbf{x} \in \mathcal{X}^s} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G}(\mathbf{x}|\mathbf{g}_2)} \\ &\leq \frac{\sum_{\mathbf{x} \in \mathcal{X}^s} e^{s\epsilon} p_{Z|X}(\mathbf{z}|\mathbf{x}_0) p_{X|G}(\mathbf{x}|\mathbf{g}_1)}{\sum_{\mathbf{x} \in \mathcal{X}^s} e^{-s\epsilon} p_{Z|X}(\mathbf{z}|\mathbf{x}_0) p_{X|G}(\mathbf{x}|\mathbf{g}_2)} \end{aligned}$$

$$= e^{2s\epsilon},$$

from which we obtain

$$e^{-2s\epsilon} \leq \frac{p_{G|Z}(\mathbf{g}|\mathbf{z})}{p_G(\mathbf{g})} = \frac{p_{Z|G}(\mathbf{z}|\mathbf{g})}{\sum_{\mathbf{g}'} p_{Z|G}(\mathbf{z}|\mathbf{g}') p_G(\mathbf{g}')} \leq e^{2s\epsilon},$$

any $\mathbf{g} \in \mathcal{G}$ and $\mathbf{z} \in \mathcal{Z}$.

(vii-viii) Suppose for any $\mathbf{x} \in \mathcal{X}^s, \mathbf{g} \in \mathcal{G}, p_{X|G}(\mathbf{x}|\mathbf{g}) = \frac{1}{|\mathcal{X}|^s}$. Then, $p_{X|G}(\mathbf{x}|\mathbf{g})/p_X(\mathbf{x}) = 1$, and

$$\frac{p_{Z|G}(\mathbf{z}|\mathbf{g})}{p_Z(\mathbf{z})} = \frac{\sum_{\mathbf{x}} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_{X|G}(\mathbf{x}|\mathbf{g})}{\sum_{\mathbf{x}} p_{Z|X}(\mathbf{z}|\mathbf{x}) p_X(\mathbf{x})} = 1,$$

for all privacy mappings $p_{Z|X}$. Therefore, $p_{G,X,Z}$ satisfies 0-information privacy but does not guarantee local differential privacy and mutual information privacy as $p_{Z|X}$ can be chosen arbitrarily.

(ix) Substitute X for U and Z for V in Example 5.1. From (5.15), there is a sequence of distributions $(p_{G,X,Z}^\alpha)_{\alpha \geq 0}$ satisfying ϵ_α -mutual information privacy with $\epsilon_\alpha \rightarrow 0$ as $\alpha \rightarrow 0$. Choose a $\mathbf{g}_0 \in \mathcal{G}$, and let

$$p_{X|G}(\mathbf{x}|\mathbf{g}_0) = \begin{cases} \alpha, & \text{if } \mathbf{x} = \mathbf{0} \\ (1 - \alpha)/(|\mathcal{X}|^s - 1), & \text{otherwise.} \end{cases}$$

For other $\mathbf{g} \in \mathcal{G}$, with $\mathbf{g} \neq \mathbf{g}_0$, we let $p_{X|G}(\mathbf{x}|\mathbf{g}) = \frac{1}{|\mathcal{X}|^s}$ for all $\mathbf{x} \in \mathcal{X}^s$. We also let G to be uniformly distributed. Then, we have

$$\begin{aligned} p_{Z|G}(\mathbf{0}|\mathbf{g}_0) &= \sum_{\mathbf{x}} p_{Z|X}(\mathbf{0}|\mathbf{x}) p_{X|G}(\mathbf{x}|\mathbf{g}_0) \\ &= p_{Z|X}(\mathbf{0}|\mathbf{0}) p_{X|G}(\mathbf{0}|\mathbf{g}_0) \\ &= \alpha, \end{aligned}$$

since $p_{Z|X}(\mathbf{0}|\mathbf{0}) = 1$ from Example 5.1. For $\mathbf{g} \neq \mathbf{g}_0$, we also have

$$\begin{aligned} p_{Z|G}(\mathbf{0}|\mathbf{g}) &= \sum_{\mathbf{x}} p_{Z|X}(\mathbf{0}|\mathbf{x}) p_{X|G}(\mathbf{x}|\mathbf{g}) \\ &= \frac{1}{|\mathcal{X}|^s}, \end{aligned}$$

and

$$p_Z(\mathbf{0}) = \frac{\alpha + \frac{|\mathcal{G}|-1}{|\mathcal{X}|^s}}{|\mathcal{G}|}.$$

Therefore, we have

$$\lim_{\alpha \rightarrow 0} \frac{p_{Z|G}(\mathbf{0}|\mathbf{g}_0)}{p_Z(\mathbf{0})} = \frac{\alpha|\mathcal{G}|}{\alpha + \frac{|\mathcal{G}|-1}{|\mathcal{X}|^s}} = 0,$$

which indicates that $\epsilon_I \rightarrow \infty$ as $\alpha \rightarrow 0$. This proves that mutual information privacy does not guarantee information privacy.

- (x) Since $I(G; Z|X) = 0$, we have $0 \leq I(Z; G) = I(X; Z) - I(X; Z|G) \leq I(X; Z)$, and the claim follows immediately.
- (xi) If $p_{G,X,Z}$ satisfies ϵ -local differential privacy, then $\frac{p_{Z|X}(\mathbf{z}|\mathbf{x}_1)}{p_{Z|X}(\mathbf{z}|\mathbf{x}_2)} \leq e^{s\epsilon}$, for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}^s$. The proof then proceeds similarly as that for (iv).
- (xii) Substitute X for U and Z for V in Example 5.1. From (5.15) and (5.17), we conclude that mutual information privacy does not guarantee local differential privacy.
- (xiii-xiv) These claims follow since for any $\mathbf{z} \in \mathcal{Z}^s$, $\mathbf{x} \sim \mathbf{x}' \in \mathcal{X}^s$, we have

$$\frac{p_{Z|X}(\mathbf{z}|\mathbf{x})p_X(\mathbf{x})}{p_{Z|X}(\mathbf{z}|\mathbf{x}')p_X(\mathbf{x}')} = \frac{p_{X|Z}(\mathbf{x}|\mathbf{z})}{p_{X|Z}(\mathbf{x}'|\mathbf{z})}.$$

The proof of the theorem is now complete.

Chapter 6

Conclusion and Future Work

In this thesis, we have investigated privacy preservation in wireless sensor network. We aim to prevent the fusion center from disclosing the private information while still enabling it to perform accurate detection of the public hypothesis of interest.

In Chapter 3, information privacy is the protection against statistical inference of a private phenomenon. 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 Internet of Things (IoT) networks in which every sensor operates independently of the others.

In Chapter 4, we have developed an approach to achieve information privacy for an uncertainty set of private hypotheses in a decentralized detection framework. We provided bounds for the utility of detecting the public hypothesis. We introduced the concept of a Most Favorable Hypothesis (MFH) of the uncertainty set of private hypotheses, and showed how to find a MFH, which allows us to transform the information privacy constraint into a more tractable metric that uses the average

of the Type I and II error probabilities of detecting the MFH. We proposed a block nonlinear Gauss-Siedel method (GS) algorithm to find the local privacy mappings at each sensor. Simulations indicate that our approach produces higher error in detecting the private hypotheses, while maintaining a reasonably low error for detecting the public hypothesis.

In Chapter 5, we have introduced the concept of privacy implication and non-guarantee to study the relationships between different inference and data privacy metrics. We showed that information privacy and local differential privacy are some of the strongest inference privacy and data privacy metrics, respectively. We considered the problem of preserving both information privacy of a private hypothesis and data privacy of the sensor observations in a decentralized network consisting of multiple sensors and a fusion center, whose task is to infer a public hypothesis of interest. In the parametric case, we proposed two different privacy mapping architectures, and showed that both achieve information privacy and local differential privacy to within the predefined budgets. In the nonparametric case, we proposed an empirical privacy optimization approach to learn the privacy mappings from a given training set. Simulations and tests on real data suggest that our proposed approaches achieve a good utility while protecting both inference and data privacy.

In this thesis, we have considered only sensor observations from a single time instance. An interesting future research direction is to generalize our approach to sensor observations over multiple time instances in a dynamic system model.

Chapter 7

Author's Publications

The following is a list of publications in refereed journals and conference proceedings produced during my Ph.D. candidature. In some cases, the journal papers contain partial material overlapping with the conference publications.

Journals:

1. **M. Sun**, W. P. Tay, and X. He, "Toward information privacy for the internet of things: A nonparametric learning approach," *IEEE Trans. Signal Process.*, vol. 66, no. 7, pp. 1734–1747, April 2018.

Conference:

1. **M. Sun**, and W. P. Tay, "Privacy-preserving Nonparametric Decentralized Detection," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Shanghai, 2016, pp. 6270–6274.
2. **M. Sun**, and W. P. Tay, "Inference and data privacy in IoT networks," *Proc. IEEE Workshop on Signal Processing Advances in Wireless Commun.*, Sapporo, 2017, pp. 1–5.
3. X. He, W. P. Tay and **M. Sun**, "Privacy-aware decentralized detection using linear precoding," in *Proc. IEEE Sensor Array and Multichannel Signal Processing Workshop*, Rio de Janeiro, 2016, pp. 1–5.

4. X. He, **M. Sun**, W. P. Tay and Y. Gong, "Multilayer nonlinear processing for information privacy in sensor networks", arXiv preprint arXiv:1711.04459, 2018.

Journals under Submission/ Revision

1. **M. Sun** and W. P. Tay, "Decentralized Detection with Robust Information Privacy Protection," submitted to *IEEE Trans. Inf. Forensics Security*.
2. **M. Sun** and W. P. Tay, "On the Relationship Between Inference and Data Privacy in Decentralized IoT Networks," submitted to *IEEE Trans. Inf. Forensics Security*.

Bibliography

- [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. on Inform. Process. in Sensor Networks*, 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. of the IEEE Eng. in Medicine and Biology 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. Alemdar and C. Ersoy, “Wireless sensor networks for healthcare: A survey,” *Computer Networks*, 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,” *Advances in Statistical Signal Processing*, 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] (2015) Gartner says 6.4 billion connected "things" will be in use in 2016, up 30 percent from 2015. [Online]. Available: <http://www.gartner.com/newsroom/id/3165317>
- [16] D. M. Mendez, I. Papapanagiotou, and B. Yang, "Internet of Things: Survey on security and privacy," *arXiv preprint arXiv:1707.01879*, 2017.
- [17] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals, Syst.*, vol. 1, pp. 167–182, 1988.
- [18] 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.
- [19] S. A. Aldosari and J. M. F. Moura, "Fusion in sensor networks with communication constraints," in *Proc. Int. Symp. on Inform. Process. in Sensor Networks*, Berkeley, California, USA, 2004, pp. 108–115.
- [20] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *Proc. ACM workshop on Privacy in the Electronic Society*, Raleigh, NC, 2012, pp. 81–90.
- [21] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private average consensus with optimal noise selection," in *Proc. IFAC Workshop Distrib. Estimation Control Networked Syst.*, vol. 48, no. 22, Philadelphia, PA, 2015, pp. 203–208.
- [22] N. E. Manitaras and C. N. Hadjicostis, "Privacy-preserving asymptotic average consensus," in *Proc. Eur. Control Conf.*, Zurich, Switzerland, 2013, pp. 760–765.
- [23] P. Braca, R. Lazzeretti, S. Marano, and V. Matta, "Learning with privacy in consensus + obfuscation," *IEEE Trans. Signal Process.*, vol. 23, no. 9, pp. 1174–1178, 2016.
- [24] J. He, L. Cai, C. Zhao, P. Cheng, and X. Guan, "Privacy-preserving average consensus: privacy analysis and algorithm design," *IEEE Trans. Signal and Information Processing over Networks*, vol. 5, no. 1, pp. 127–138, 2019.

-
- [25] R. Lazzeretti, S. Horn, P. Braca, and P. Willett, “Secure multi-party consensus gossip algorithms,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Florence, 2014, pp. 7406–7410.
 - [26] M. Ambrosin, P. Braca, M. Conti, and R. Lazzeretti, “Odin: Oblivious-based privacy-preserving consensus algorithm for decentralized information fusion in smart device networks,” *ACM Trans. on Internet Technology*, vol. 18, no. 1, p. 6, 2017.
 - [27] P. Hallgren, C. Orlandi, and A. Sabelfeld, “Privatepool: privacy-preserving ridesharing,” in *Proc. IEEE Computer Security Found. Symp.*, Santa Barbara, CA, 2017, pp. 276–291.
 - [28] C. C. Aggarwal and S. Y. Philip, “A general survey of privacy-preserving data mining models and algorithms,” in *Privacy-preserving data mining*. Springer, 2008, pp. 11–52.
 - [29] Z. Jun, “Differentially private deep learning with tight privacy-cost computation,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Calgary, 2018.
 - [30] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, “Differentially private empirical risk minimization,” *J. Machine Learning Research*, vol. 12, pp. 1069–1109, Mar. 2011.
 - [31] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” *arXiv preprint physics/0004057*, 2000.
 - [32] G. Chechik and N. Tishby, “Extracting relevant structures with side information,” in *Advances in Neural Information Processing Systems*, vol. 15. MIT Press, 2003, pp. 881–888.
 - [33] Y. Wang, X. Wu, and H. Donghui, “Using randomized response for differential privacy preserving data collection,” in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, Washington, D.C., 2003, pp. 505–510.
 - [34] S. Xiong, A. D. Sarwate, and N. B. Mandayam, “Randomized requantization with local differential privacy,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Shanghai, 2016, pp. 2189–2193.
 - [35] A. D. Sarwate and L. Sankar, “A rate-distortion perspective on local differential privacy,” in *Proc. Allerton Conf. on Commun., Control and Computing*, Monticello, IL, 2014, pp. 903–908.
 - [36] J. Liao, L. Sankar, F. P. Calmon, and V. Y. Tan, “Hypothesis testing under maximal leakage privacy constraints,” in *Proc. IEEE Int. Symp. on Inform. Theory*, Aachen, Germany, 2017, pp. 779–783.
 - [37] H. Imtiaz and A. D. Sarwate, “Differentially private distributed principal component analysis,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Calgary, 2018.

- [38] B. Razeghi and S. Voloshynovskiy, “Privacy-preserving outsourced media search using secure sparse ternary codes,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Calgary, 2018.
- [39] F. D. McSherry, “Privacy integrated queries: an extensible platform for privacy-preserving data analysis,” in *Proc. ACM SIGMOD Inter. Conf. Management of data*, Providence, Rhode Island, 2009, pp. 19–30.
- [40] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, “Our data, ourselves: Privacy via distributed noise generation,” in *Annual Inter. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2006, pp. 486–503.
- [41] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Local privacy and statistical minimax rates,” in *Proc. IEEE Symp. on Foundations of Computer Science*, Berkeley, 2013, pp. 429–438.
- [42] C. Gentry, “Fully homomorphic encryption using ideal lattices,” in *Proc. ACM Symp. Theory of Comput.*, Bethesda, MD, 2009, pp. 169–178.
- [43] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(Leveled) fully homomorphic encryption without bootstrapping,” *ACM Trans. Computation Theory*, vol. 6, no. 3, p. 13, 2014.
- [44] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, “Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy,” in *Proc. Int. Conf. on Machine Learning*, New York, 2016, pp. 201–210.
- [45] M. Sun and W. P. Tay, “Inference and data privacy in IoT networks,” in *Proc. IEEE Workshop on Signal Processing Advances in Wireless Commun.*, 2017.
- [46] A. Alabdulatif, I. Khalil, M. Reynolds, H. Kumarage, and X. Yi, “Privacy-preserving data clustering in cloud computing based on fully homomorphic encryption,” in *Proc. Pacific-Asia Conf. on Information Systems*, 2017, pp. 1–13.
- [47] W. Wang, L. Ying, and J. Zhang, “On the relation between identifiability, differential privacy, and mutual-information privacy,” *IEEE Trans. Inf. Theory*, vol. 62, no. 9, pp. 5018–5029, Jun. 2016.
- [48] F. du Pin Calmon and N. Fawaz, “Privacy against statistical inference,” in *Proc. Allerton Conf. on Commun., Control and Computing*, Monticello, IL, 2012, pp. 1401–1408.
- [49] H. Yamamoto, “A source coding problem for sources with additional outputs to keep secret from the receiver or wiretappers,” *IEEE Trans. Inf. Theory*, vol. 29, no. 6, pp. 918–923, 1983.
- [50] L. Sankar, S. R. Rajagopalan, and H. V. Poor, “A theory of utility and privacy of data sources,” in *Proc. IEEE Int. Symp. on Inform. Theory*, 2010, pp. 2642–2646.

-
- [51] Z. Li and T. J. Oechtering, “Differential privacy in parallel distributed Bayesian detections,” in *Proc. Int. Conf. on Information Fusion*, Salamanca, 2014, pp. 1–7.
 - [52] C. Dwork, “Differential privacy: A survey of results,” in *Proc. Int. Conf. Theory and Applications of Models of Computation*, Xi’an, 2008, pp. 1–19.
 - [53] C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
 - [54] A. Friedman and A. Schuster, “Data mining with differential privacy,” in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, Washington, DC, 2010, pp. 493–502.
 - [55] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proc. ACM SIGSAC Conf. on Computer & Commun. Security*. ACM, 2016, pp. 308–318.
 - [56] X. Xiao, G. Wang, and J. Gehrke, “Differential privacy via wavelet transforms,” *IEEE Trans. Knowledge and Data Engineering*, vol. 23, no. 8, pp. 1200–1214, 2011.
 - [57] J. Hamm, “Enhancing utility and privacy with noisy minimax filters,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, New Orleans, LA, Mar. 2017, pp. 6389–6393.
 - [58] Z. He, Z. Cai, Y. Sun, Y. Li, and X. Cheng, “Customized privacy preserving for inherent data and latent data,” *Personal and Ubiquitous Computing*, vol. 21, no. 1, pp. 43–54, 2017.
 - [59] R. R. Tenney and N. R. Sandell, “Detection with distributed sensors,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 17, no. 4, pp. 501–510, 1981.
 - [60] X. Nguyen, W. M. J., and M. I. Jordan, “Nonparametric decentralized detection using kernel methods,” *IEEE Trans. Signal Process.*, vol. 53, no. 11, pp. 4053–4066, Nov. 2005.
 - [61] 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.
 - [62] P. J. Huber, “A robust version of the probability ratio test,” *Ann. Math. Statist.*, vol. 36, no. 6, pp. 1753–1758, Dec. 1965.
 - [63] P. J. Huber and V. Strassen, “Minimax tests and the neyman-pearson lemma for capacities,” *The Annals of Statistics*, pp. 251–263, 1973.
 - [64] G. Moustakides, “Robust detection of signals: A large deviations approach,” *IEEE Trans. Inf. Theory*, vol. 31, no. 6, pp. 822–825, 1985.

- [65] Z. Goldfeld, G. Kramer, and H. H. Permuter, “Broadcast channels with privacy leakage constraints,” *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 5138–5161, 2017.
- [66] F. P. Calmon, A. Makhdoumi, and M. Médard, “Fundamental limits of perfect privacy,” in *Proc. IEEE Int. Symp. on Inform. Theory*, Hong Kong, 2015, pp. 1796–1800.
- [67] (2016) NB-IOT—Enabling new business opportunities. [Online]. Available: <http://www.huawei.com/minisite/4-5g/img/NB-IOT.pdf>
- [68] H. V. Poor, *An introduction to signal detection and estimation*. Springer Science & Business Media, 2013.
- [69] B. Scholkopf and A. J. Smola, *Learning With Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT press, 2001.
- [70] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience, 2006.
- [71] E. Shi, H. Chan, E. Rieffel, R. Chow, and D. Song, “Privacy-preserving aggregation of time-series data,” in *Proc. Network & Distributed System Security Symp.*, San Diego, 2011.
- [72] J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle, “Privacy in the Internet of Things: Threats and challenges,” *Security and Communication Networks*, vol. 7, no. 12, pp. 2728–2742, Jun. 2014.
- [73] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Privacy aware learning,” *J. ACM*, vol. 61, no. 6, p. 38, Dec. 2014.
- [74] R. Hall, A. Rinaldo, and L. Wasserman, “Differential privacy for functions and functional data,” *J. Machine Learning Research*, vol. 14, pp. 703–727, Feb. 2013.
- [75] T. Zhang, “Statistical behavior and consistency of classification methods based on convex risk minimization,” *Annals of Statistics*, vol. 32, no. 1, pp. 56–85, Feb. 2004.
- [76] J. Friedman, T. Hastie, and R. Tibshirani, “Additive logistic regression: a statistical view of boosting,” *Annals of Statistics*, vol. 28, no. 2, pp. 337–407, 2000.
- [77] Y. Freund and R. E. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” in *Proc. European Conf. on Computational Learning Theory*, 1995, pp. 23–37.
- [78] L. Rosasco, E. De Vito, A. Caponnetto, M. Piana, and A. Verri, “Are loss functions all the same?” *Neural Computation*, vol. 16, no. 5, pp. 1063–1076, May 2004.

-
- [79] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe, “Convexity, classification, and risk bounds,” *J. of the American Statistical Association*, vol. 101, no. 473, pp. 138–156, Jan. 2006.
- [80] H. Wendland, *Scattered Data Approximation*. Göttingen, Germany: Cambridge University Press, 2004.
- [81] R. T. Rockafellar, *Convex Analysis*. Princeton University Press, 2015.
- [82] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [83] L. Grippo and M. Sciandrone, “On the convergence of the block nonlinear Gauss-Seidel method under convex constraints,” *Operations Research Letters*, vol. 26, no. 3, pp. 127–136, 2000.
- [84] Y. Xu and W. Yin, “A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion,” *SIAM J. on Imaging Sciences*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [85] P. Tseng, “Convergence of a block coordinate descent method for nondifferentiable minimization,” *J. of Optimization Theory and Applicat.*, vol. 109, no. 3, pp. 475–494, 2001.
- [86] A. C. Lorena, A. C. De Carvalho, and J. M. Gama, “A review on the combination of binary classifiers in multiclass problems,” *Artificial Intelligence Review*, vol. 30, no. 1, pp. 19–37, 2008.
- [87] M. Lichman, “UCI machine learning repository,” 2013. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [88] M. Al, S. Wan, and S. Kung, “Ratio utility and cost analysis for privacy preserving subspace projection,” *arXiv preprint arXiv:1702.07976*, 2017.
- [89] K. Diamantaras and S. Kung, “Data privacy protection by kernel subspace projection and generalized eigenvalue decomposition,” in *IEEE Int. Workshop Machine Learning for Signal Processing*, Salerno, 2016, pp. 1–6.
- [90] S.-L. Sun, “Distributed optimal component fusion weighted by scalars for fixed-lag kalman smoother,” *Automatica*, vol. 41, no. 12, pp. 2153–2159, 2005.
- [91] M. Sun and W. P. Tay, “Privacy-preserving nonparametric decentralized detection,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Shanghai, 2016, pp. 6270–6274.
- [92] R. Chavarriaga, H. Sagha, A. Calatroni, S. T. Digumarti, G. Tröster, J. del R. Millán, and D. Roggen, “The opportunity challenge: A benchmark database for on-body sensor-based activity recognition,” *Pattern Recognition Lett.*, vol. 34, no. 15, pp. 2033–2042, 2013.

- [93] M. Hall *et al.*, “The WEKA data mining software: an update,” *SIGKDD Explorations*, vol. 11, no. 1, pp. 10–18, 2009.
- [94] W. Rudin, *Principles of Mathematical Analysis*. McGraw-Hill New York, 1964.
- [95] I. Issa and A. B. Wagner, “Operational definitions for some common information leakage metrics,” in *Proc. IEEE Int. Symp. on Inform. Theory*. IEEE, 2017, pp. 769–773.
- [96] R. Martin and S. Schwartz, “Robust detection of a known signal in nearly gaussian noise,” *IEEE Trans. Inf. Theory*, vol. 17, no. 1, pp. 50–56, 1971.
- [97] M. López and G. Still, “Semi-infinite programming,” *European Journal of Operational Research*, vol. 180, no. 2, pp. 491–518, 2007.
- [98] J. N. Tsitsiklis and M. Athans, “On the complexity of decentralized decision making and detection problems,” *IEEE Trans. Autom. Control*, vol. 30, pp. 440–446, 1985.
- [99] M. B. Cohen, Y. T. Lee, and Z. Song, “Solving linear programs in the current matrix multiplication time,” *CoRR*, vol. abs/1810.07896, 2018. [Online]. Available: <http://arxiv.org/abs/1810.07896>
- [100] A. El Gamal and Y.-H. Kim, *Network information theory*. Cambridge University Press, 2011.
- [101] I. Csiszar and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.
- [102] O. P. Kreidl, J. N. Tsitsiklis, and S. I. Zoumpoulis, “On decentralized detection with partial information sharing among sensors,” *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1759–1765, Apr. 2011.
- [103] Z. Zhang, E. Chong, A. Pezeshki, W. Moran, and S. Howard, “Learning in hierarchical social networks,” *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 2, pp. 305–317, Apr. 2013.
- [104] J. Lofberg, “Yalmip : A toolbox for modeling and optimization in matlab,” in *Proc. Int. Symp. Computer-Aided Control System Design*, Taipei, 2004, pp. 284–289.
- [105] R. Kohavi, “Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid,” in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, vol. 96, Portland, Oregon, 1996, pp. 202–207.
- [106] S.-Y. Kung, “Discriminant component analysis for privacy protection and visualization of big data,” *Multimedia Tools and Applications*, vol. 76, no. 3, pp. 3999–4034, 2017.

- [107] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.