

A Framework for Secure Speech Recognition

Paris Smaragdis *, *Senior Member, IEEE* and Madhusudana Shashanka, *Student Member, IEEE*

Abstract—In this paper we present a process which enables privacy-preserving speech recognition transactions between two parties. We assume one party with private speech data and one party with private speech recognition models. Our goal is to enable these parties to perform a speech recognition task using their data, but without exposing their private information to each other. We will demonstrate how using secure multiparty computation principles we can construct a system where this transaction is possible, and how this system is computationally and securely correct. The protocols described herein can be used to construct a rudimentary speech recognition system and can easily be extended for arbitrary audio and speech processing.

Index Terms—Secure Multiparty Computation, Gaussian Mixture Models, HMM, Speech Recognition.

EDICS: SPE-GASR

I. INTRODUCTION

The widespread use of networking technology today has spawned an industry of online services. Business models based on client-server interactions are common place and increasingly more prominent. Speech recognition could easily be a part of this trend where servers can provide speech recognition services for remote clients. The private nature of speech data however is a stumbling block for such a development. Individuals, corporations and governments are understandably reluctant to send private speech data through a network to another party that cannot be trusted. In this paper we address this issue and show how such a cooperative model can be realized with no privacy leaks from any involved party.

We will specifically focus on the realization of a hidden Markov model (HMM) in a secure framework that allows training and classification between multiple parties, some owning speech data and some owning HMM models for speech recognition. Our formulation is shaped in such a way that the providers of the speech will not have to share any information about their data and the providers of the HMM will not share any information on their model. After evaluation the results will only be revealed to the parties that have provided the data, and not to the parties that provide the HMM models, thereby providing privacy at both the data and the semantic levels.

We will demonstrate the use of this idea using two scenarios. One scenario will involve the training of HMMs from data provided by multiple parties, and the other will deal with evaluating already trained HMMs on data from another party. The utility of these scenarios in collaborative speech recognition projects is easy to see. The first scenario can enable model training on multiple speech databases without requiring the disclosure of actual speech data, whereas the other scenario can enable speech recognition as a service model to remote customers who wish to maintain privacy on their data and their

transcription from both the service provider and malicious network intruders.

Strange as these constraints might seem, they can be satisfied using secure multiparty computation (SMC) protocols. SMC is a field of cryptography that provides means to perform arbitrary computations between multiple parties who are concerned with protecting their data. The field of SMC originated from the work of Yao [1] who gave a solution to the millionaire problem: two millionaires want to find which one has a larger fortune, without revealing any specific numbers to each other. Recently this concept has been employed for simple machine learning tasks such as multiple parties performing k-means [2], computation of means and related statistics from distributed databases [3] and rudimentary computer vision applications [4]. See [5] for a detailed treatment of the topic. In this paper we present a SMC formulation of training and evaluating HMMs as applied on speech data. To our knowledge this is the first application of SMC concepts for privacy constrained speech technology. We will consider HMMs where the observations are modeled by mixtures of Gaussians as is common in speech recognition applications. The main contributions in this paper are the creation of privacy-preserving protocols that support Gaussian mixture model and HMM learning and evaluation, as well as a secure method to combine these protocols so as to ensure data privacy.

The remainder of this paper is organized as follows. In section II we formally introduce the problem at hand and in section III we introduce the secure computation primitives that are employed for this task. Using these primitives, we deal with the problem of secure classification using Gaussian mixtures in section IV and in section V we extend that to present protocols for secure HMMs. We provide a brief discussion about security and efficiency of our protocols in section VI. Finally in section VII we provide conclusions and directions for future extensions.

II. PROBLEM FORMULATION

Hidden Markov models find use in a wide range of applications, and have been successfully used in speech recognition. There are three fundamental problems for HMM design, namely: the evaluation of the probability (or likelihood) of a sequence of observations given a specific HMM; the determination of a best sequence of model states; and the adjustment of model parameters so as to best account for the observed signal. The first problem is one of scoring how well a given model matches a given observation sequence. The second problem is one in which we attempt to uncover the hidden part of the model. The third problem is the problem of “training”. Algorithms for the above three problems are well known and described in detail in [6].

We will consider these problems using a transaction between two parties named Alice and Bob. Suppose Bob has a trained HMM with all the model parameters learned. Let the HMM be characterized as follows:

- N states $\{S_1, \dots, S_N\}$. Let the state at time t be q_t .
- The state transition probability distribution $\mathbf{A} = \{a_{ij}\}$ where

$$a_{ij} = P[q_{t+1} = S_j | q_t = S_i], \quad 1 \leq i, j \leq N. \quad (1)$$

- The observation symbol probability distribution in state j given by a mixture of Gaussians

$$b_j(\mathbf{x}) = \sum_{m=1}^M c_{jm} \mathcal{N}(\boldsymbol{\mu}_{jm}, \boldsymbol{\Sigma}_{jm}), \quad 1 \leq j \leq N, \quad (2)$$

where \mathbf{x} is the variable, c_{jm} is the mixture coefficient for the m -th mixture in state j , and $\mathcal{N}(\boldsymbol{\mu}_{jm}, \boldsymbol{\Sigma}_{jm})$ is a gaussian with mean vector $\boldsymbol{\mu}_{jm}$ and covariance matrix $\boldsymbol{\Sigma}_{jm}$.

- The initial state distribution $\pi = \{\pi_i\}$ where

$$\pi_i = P[q_1 = S_i] \quad 1 \leq i \leq N. \quad (3)$$

We use λ to denote the entire parameter set of the model.

Consider the first two problems where Bob has a trained HMM with all the model parameters learned. Let Alice have an observation sequence $\mathbf{X} = \mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_T$. We will show how Alice can securely compute $P(\mathbf{X} | \lambda)$, the probability of the observation sequence given the model, using the *forward-backward* procedure. We will also show how one can securely learn the best sequence of model states using the *viterbi* algorithm.

Once there is a secure way of computing likelihoods, it is easy to see how it can be extended to applications like speech recognition. Suppose Bob has trained several HMMs which characterize various speech sounds. Each HMM will correspond to a speech recognition unit. Let Alice's observation vector correspond to a small snippet of speech sound (we assume that Alice knows the features that Bob has used to train his HMMs on and has represented her sound sample in terms of those features). We then show how Alice and Bob can obtain additive shares of the likelihood of Alice's observation sequence for every speech recognition unit of Bob and use them to find out the unit that corresponds to Alice's sound snippet.

Now consider the third problem of training. The problem of security arises when Bob wants to train a HMM (or do data mining in general) on combined data from private databases owned by Alice and Charlie. We show how Bob can securely reestimate parameters of his HMM without gaining knowledge about the private data. The implicit assumption, of course, is that Alice and Charlie are willing to let Bob learn about distributions of their data.

III. SECURE TWO-PARTY COMPUTATIONS; BACKGROUND

The speech-recognition example that we will present is a specific example of a *secure two-party computation*. Consider the case where Alice and Bob have private data \mathbf{a} and \mathbf{b} respectively and they want to compute the result of a function




Fig. 1. Implementing an algorithm securely. The algorithm takes in private inputs \mathbf{a} and \mathbf{b} . Algorithm is split into steps that can be implemented as secure primitives (shown as grey boxes). Intermediate results are distributed as random additive shares and feed into the following steps. Final result \mathbf{c} is obtained by both parties (or the designated receiver).

$f(\mathbf{a}, \mathbf{b})$. Consider a trusted third-party who can take the private data, compute the result $\mathbf{c} = f(\mathbf{a}, \mathbf{b})$, and intimate the result to the parties. Any protocol that implements an algorithm to calculate $f(\mathbf{a}, \mathbf{b})$ is said to be *secure* only if it leaks no more information about \mathbf{a} and \mathbf{b} than what one can gain from learning the result \mathbf{c} from the trusted third-party. We assume a *semi-honest* model for the parties where they follow the protocol but could be saving messages and intermediate results to learn more about other's private data. In other words, the parties are *honest but curious* and will follow the agreed-upon protocol but will try to learn as much as possible from the data flow between the parties.¹

To implement an algorithm securely, we will have to implement each step of the algorithm securely. If one of the steps is insecurely implemented, either party could utilize the information to work their way backwards to gain knowledge about the other's private data. In addition, one must also consider the results of intermediate steps. If such results of intermediate steps are available, there is a possibility that one could also get back to the original private inputs. To prevent this:

- we express every step of the algorithm in terms of a handful of basic operations (henceforth called as *primitives*) for which secure implementations are already known, and
- we distribute intermediate results randomly between the two parties such that neither party has access to the entire result. For example, instead of obtaining the result z of a certain step, the parties receive *random additive shares* z_1 and z_2 ($z_1 + z_2 = z$). See figure 1 for a schematic illustration.

Secure protocols are often analyzed for correctness, security and complexity. Correctness is measured by comparing the proposed protocol with the ideal protocol using a third party. If the results are indistinguishable, the protocol is correct (note that one can use secure approximation to an ideal algorithm). All the protocols we present are exact protocols. For security, one needs to show what can and cannot be learned from

¹In a *malicious model*, no such assumptions are made about the parties' behavior. If both parties are malicious, security can be enforced by accompanying the protocols with zero-knowledge proofs that protocols are being followed. If only one of the parties is malicious, the other party can use *conditional disclosure of secrets* protocols [7] to make sure he/she receives valid inputs from the malicious party.

the data exchange. For complexity, one shows the computational and communication complexity of the secure algorithm. Based on the choice of primitives used and how they are implemented, one can achieve different levels of security and computational/communication efficiency. To evaluate the efficiency of protocols we propose in later sections, we provide measures in terms of efficiency of the primitives instead of absolute measures. Below, we describe the primitives that we use and briefly discuss about their implementations.

A. Secure Inner Products (SIP)

The primitive which we use most often is for computing secure inner products. If Alice has vector \mathbf{x} and Bob has vector \mathbf{y} , a secure inner product protocol produces two numbers a and b such that $a + b = \mathbf{x}^T \mathbf{y}$. Alice will get the result a and Bob will get the result b . To simplify notation, we shall denote a secure inner product computation $\mathbf{x}^T \mathbf{y}$ as $SIP(\mathbf{x}, \mathbf{y})$.

Many protocols have been proposed and they can be categorized as cryptographic protocols (eg. [8], [9]) and algebraic protocols (eg. [10], [11], [12], [13], [14]). They provide different levels of security and efficiency. Most of the algebraic protocols leak some information but are more straightforward and efficient than their cryptographic counterparts. The properties and weaknesses of some of these algebraic protocols have been analyzed in detail ([3], [9]). In this paper we will be using cryptographic protocols as they tend to be more secure. In appendices I and II we provide a description of two of the cryptographic protocols we have used. We have also included an algebraic protocol in appendix III for comparison.

B. Secure Maximum Index (SMAX)

Let Alice have a vector $\mathbf{x} = [x_1 \dots x_d]$ and Bob have the vector $\mathbf{y} = [y_1 \dots y_d]$, they would like to compute the index of the maximum of $\mathbf{x} + \mathbf{y} = [(x_1 + y_1) \dots (x_d + y_d)]$. At the end of the protocol, Alice (and/or Bob) will receive the result but neither party will know the actual value of the maximum. Notice that the same protocol can be used to compute the index of the minimum. We denote this as $j = SMAX(\mathbf{x}, \mathbf{y})$.

For this primitive, we use the permute protocol proposed by [15] (see appendix IV). The protocol enables Alice and Bob to obtain additive shares, \mathbf{q} and \mathbf{s} , of a permutation of the vector $\mathbf{x} + \mathbf{y}$, $\pi(\mathbf{x} + \mathbf{y})$, where π is chosen by Alice and Bob has no knowledge of π . The idea is for Alice to send $\mathbf{q} - r$, where r is a random number chosen by her, to Bob. Bob sends back the index of the maximum element of $\mathbf{q} + \mathbf{s} - r$ to Alice who then computes the real index using the inverse of the permutation π . Neither party learns the value of the maximum element and Bob does not learn the index of the maximum element.

If the security requirements of Alice are more strict, she can encrypt elements of \mathbf{q} using a homomorphic encryption scheme and send them to Bob along with the public key. Bob can make comparisons among the encrypted values of $\pi(\mathbf{x} + \mathbf{y})$ using protocols for Yao's millionaire problem (eg. [16], [17]) and send back the index of the maximum element.

C. Secure Maximum Value (SVAL)

Let Alice have a vector $\mathbf{x} = [x_1 \dots x_d]$ and Bob have the vector $\mathbf{y} = [y_1 \dots y_d]$, they would like to compute the value of the maximum element in $\mathbf{z} = \mathbf{x} + \mathbf{y}$. After the protocol, Alice and Bob receive additive shares of the result, a and b , but neither party will know the index of the maximum element. Notice that the same protocol can be used to compute the value of the minimum. Let us denote this as $a + b = SVAL(\mathbf{x}, \mathbf{y})$.

For this protocol, we can use the idea presented in [15]. Let us first consider a naive approach. Notice that $z_i \geq z_j \iff (x_i - x_j) \geq (y_j - y_i)$. Alice and Bob can do such pairwise comparisons and mimic any standard maximum finding algorithm to learn the value of the maximum. To perform the comparisons securely, they can use a protocol for Yao's millionaire problem [1].

However, if Alice and Bob follow the above naive approach, both will be able to also find the index of the maximum. Hence, the idea is for Alice and Bob to obtain two vectors whose sum is a random permutation of \mathbf{z} . Neither Alice nor Bob should know the permutation. They can then follow the above naive approach on their newly obtained vectors to compute additive shares of the maximum element. See appendix IV for a description of the permutation protocol.

D. Secure Logsum (SLOG)

This primitive, unlike the other three which we have introduced above, is not a cryptographic primitive. The main reason we introduce it is because it simplifies the presentation of many of the protocols we propose in future sections.

Let Alice have a vector $\mathbf{x} = [x_1 \dots x_d]$ and Bob have the vector $\mathbf{y} = [y_1 \dots y_d]$ such that $\mathbf{x} + \mathbf{y} = \ln \mathbf{z} = [\ln z_1 \dots \ln z_d]$. They would like to compute additive shares q and s such that $q + s = \ln(\sum_{i=1}^d z_i)$. Let us denote this secure computation as $q + s = SLOG(\mathbf{x}, \mathbf{y})$.

One can compute logarithm of a sum from the logarithms of individual terms as follows:

$$\ln \left(\sum_{i=1}^d z_i \right) = \ln \left(\sum_{i=1}^d e^{x_i + y_i} \right) \quad (4)$$

This suggests the following protocol.

- 1) Alice and Bob compute the dot product between vectors $e^{\mathbf{x}-q}$ and $e^{\mathbf{y}}$ using $SIP(e^{\mathbf{x}-q}, e^{\mathbf{y}})$ where q is a random number chosen by Alice. Let Bob obtain ϕ , the result of the dot product.
- 2) Notice that Bob has $s = \ln \phi = -q + \ln(\sum_{j=1}^d e^{x_j + y_j})$ and Alice has q .

In step (3), Bob receives the entire result of a dot product. However, this does not reveal any information to him about \mathbf{x} due to the presence of the random number q . In no other step does either party receive the complete result of an operation. Thus, the protocol is secure. In terms of efficiency, this primitive is equivalent to using the SIP primitive once.

IV. SECURE CLASSIFICATION: GAUSSIAN MIXTURE MODELS

Alice has a d -component data vector \mathbf{x} and Bob knows multivariate gaussian distributions of N classes $\omega_i, i =$

$\{1, \dots, N\}$ that the vector could belong to. They would like to engage in a protocol that lets Bob classify Alice's data but neither of them wants to disclose data to the other person. We propose protocols which enable such computations.

The idea is to evaluate the value of the *discriminant function*

$$g_i(\mathbf{x}) = \ln p(\mathbf{x}|\omega_i) + \ln P(\omega_i) \quad (5)$$

for all classes ω_i and assign \mathbf{x} to class ω_i if $g_i(\mathbf{x}) > g_j(\mathbf{x})$ for all $j \neq i$. Here, $p(\mathbf{x}|\omega_i)$ is the class-conditional probability density function and $P(\omega_i)$ is the *a priori* probability of class ω_i . We consider two cases where: (1) each class is modeled as a single multivariate gaussian, and (2) each class modeled as a mixture of gaussians.

A. Case 1: Single Multivariate Gaussian

We assume that the distribution of data is multivariate gaussian i.e. $p(\mathbf{x}|\omega_i) \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$, where $\boldsymbol{\mu}_i$ is the mean vector and $\boldsymbol{\Sigma}_i$ is the covariance matrix of class ω_i . Hence, the log-likelihood is given by:

$$\ln p(\mathbf{x}|\omega_i) = -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^t \boldsymbol{\Sigma}_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i) - \frac{d}{2} \ln 2\pi - \frac{1}{2} \ln |\boldsymbol{\Sigma}_i| \quad (6)$$

Ignoring the constant term $(d/2) \ln 2\pi$, we can write equation (5) as:

$$g_i(\mathbf{x}) = -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^t \boldsymbol{\Sigma}_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i) - \frac{1}{2} \ln |\boldsymbol{\Sigma}_i| + \ln P(\omega_i) \quad (7)$$

Simplifying, we have:

$$g_i(\mathbf{x}) = \mathbf{x}^T \bar{\mathbf{W}}_i \mathbf{x} + \bar{\mathbf{w}}_i^T \mathbf{x} + w_{i0} \quad (8)$$

where

$$\begin{aligned} \bar{\mathbf{W}}_i &= -\frac{1}{2} \boldsymbol{\Sigma}_i^{-1}, \quad \bar{\mathbf{w}}_i = \boldsymbol{\Sigma}_i^{-1} \boldsymbol{\mu}_i, \quad \text{and} \\ w_{i0} &= -\frac{1}{2} \boldsymbol{\mu}_i^T \boldsymbol{\Sigma}_i^{-1} \boldsymbol{\mu}_i - \frac{1}{2} \ln |\boldsymbol{\Sigma}_i| + \ln P(\omega_i) \end{aligned} \quad (9)$$

Let us create the $(d+1)$ -dimensional vectors $\bar{\mathbf{x}}$ and \mathbf{w}_i by appending the value 1 to \mathbf{x} and appending w_{i0} to $\bar{\mathbf{w}}_i$. By changing $\bar{\mathbf{W}}_i$ into a $(d+1) \times (d+1)$ matrix \mathbf{W}_i where the first d components of the last row are zeros and the last column is equal to $\bar{\mathbf{w}}_i^T$, we can express equation (8) in a simplified form:

$$g_i(\mathbf{x}) = \bar{\mathbf{x}}^T \mathbf{W}_i \bar{\mathbf{x}}$$

Expressing $\bar{\mathbf{x}}$ as \mathbf{x} for simplicity, we can write the above equation as:

$$g_i(\mathbf{x}) = \mathbf{x}^T \mathbf{W}_i \mathbf{x} \quad (10)$$

Henceforth, we shall use \mathbf{x} to denote a $(d+1)$ -dimensional vector with the last component equal to 1 unless otherwise mentioned.

Protocol SMG: Single Multivariate Gaussian

Input: Alice has vector \mathbf{x} , Bob has \mathbf{W}_i for $i = 1, 2, \dots, N$. We express the matrix \mathbf{W}_i as $[\mathbf{W}_i^1 \mathbf{W}_i^2 \dots \mathbf{W}_i^{d+1}]$, where \mathbf{W}_i^j is the j -th column of \mathbf{W}_i .

Output: Alice learns I such that $g_I(\mathbf{x}) > g_j(\mathbf{x})$ for all $j \neq I$. Bob learns nothing about \mathbf{x} .

1) For $i = 1, 2, \dots, N$

a) For $j = 1, \dots, d+1$, Alice and Bob perform *SIP*($\mathbf{x}, \mathbf{W}_i^j$) to obtain the vectors $\mathbf{a}_i = [a_i^1 \dots a_i^{d+1}]$ and $\mathbf{b}_i = [b_i^1 \dots b_i^{d+1}]$ respectively. Alice then computes $\mathbf{a}_i \mathbf{x}$.

b) Alice and Bob perform *SIP*(\mathbf{b}_i, \mathbf{x}) to obtain q_i and r_i respectively.

2) Alice has vector $\mathbf{A} = [(\mathbf{a}_1 \mathbf{x} + q_1) \dots (\mathbf{a}_N \mathbf{x} + q_N)]$ and Bob has vector $\mathbf{B} = [r_1 \dots r_N]$.

3) Alice and Bob perform the secure maximum index protocol between the vectors \mathbf{A} and \mathbf{B} and Alice obtains $I = \text{SMAX}(\mathbf{A}, \mathbf{B})$.

Correctness: In step 1, \mathbf{a}_i and \mathbf{b}_i are vectors such that $\mathbf{a}_i + \mathbf{b}_i = \mathbf{x}^T \mathbf{W}_i$. Also, $\mathbf{b}_i \mathbf{x} = q_i + r_i$. Hence, $\mathbf{x}^T \mathbf{W}_i \mathbf{x}$ is given by $\mathbf{a}_i \mathbf{x} + q_i + r_i$. I is the value of i for which $\mathbf{x}^T \mathbf{W}_i \mathbf{x}$ is maximum.

Efficiency: For a given $i = I$, the above protocol has $(d+2)$ *SIP* calls. Hence, it would take $N(d+2)$ *SIP* calls and one call of *SMAX*.

Security: If Bob gets to know the dot products of d different vectors with \mathbf{x} , he can learn \mathbf{x} completely. However, we see that neither Bob nor Alice ever learn the complete result of any dot product. Hence, if the protocols for *SIP* and *SMAX* are secure, the above protocol is secure.

B. Case 2: Mixture of Gaussians

Let us now consider the case where each class is modeled as a mixture of gaussians. Let the mean vector and covariance matrix of the j -th gaussian in class ω_i be $\boldsymbol{\mu}_{ij}$ and $\boldsymbol{\Sigma}_{ij}$ respectively. Hence we have $p(\mathbf{x}|\omega_i) = \sum_{j=1}^{J_i} \alpha_{ij} \mathcal{N}(\boldsymbol{\mu}_{ij}, \boldsymbol{\Sigma}_{ij})$ where J_i is the number of gaussians describing class ω_i and α_{ij} are the mixture coefficients. The log likelihood for the j -th gaussian in the i -th class is given by:

$$l_{ij}(\mathbf{x}) = \mathbf{x}^T \bar{\mathbf{W}}_{ij} \mathbf{x} + \bar{\mathbf{w}}_{ij}^T \mathbf{x} + w_{ij} \quad (11)$$

where

$$\bar{\mathbf{W}}_{ij} = -\frac{1}{2} \boldsymbol{\Sigma}_{ij}^{-1}, \quad \bar{\mathbf{w}}_{ij} = \boldsymbol{\Sigma}_{ij}^{-1} \boldsymbol{\mu}_{ij}, \quad \text{and}$$

$$w_{ij} = -\frac{1}{2} \boldsymbol{\mu}_{ij}^T \boldsymbol{\Sigma}_{ij}^{-1} \boldsymbol{\mu}_{ij} - \frac{1}{2} \ln |\boldsymbol{\Sigma}_{ij}| + \ln \alpha_{ij}$$

Expressing \mathbf{x} as a $(d+1)$ -dimensional vector and $\bar{\mathbf{W}}_{ij}, \bar{\mathbf{w}}_{ij}, w_{ij}$ together as the $(d+1) \times (d+1)$ matrix \mathbf{W}_{ij} as done in the previous case, we can simplify equation (11) as:

$$l_{ij}(\mathbf{x}) = \mathbf{x}^T \mathbf{W}_{ij} \mathbf{x} \quad (12)$$

Hence, the discriminant function for the i -th class can be written as

$$\begin{aligned} g_i(\mathbf{x}) &= \text{logsum}(l_{i1}(\mathbf{x}), \dots, l_{iJ_i}(\mathbf{x})) + \ln P(\omega_i) \\ &= \ln \left(\sum_{j=1}^{J_i} e^{l_{ij}(\mathbf{x})} \right) + \ln P(\omega_i) \end{aligned} \quad (13)$$

Input: Alice has vector \mathbf{x} , Bob has \mathbf{W}_{ij} and $P(\omega_i)$ for $i = 1, 2, \dots, N$, and $j = 1, 2, \dots, J_i$.

Output: Alice learns I such that $g_I(\mathbf{x}) > g_j(\mathbf{x})$ for all $j \neq I$. Bob learns nothing about \mathbf{x} .

- 1) For $i = 1, 2, \dots, N$
 - a) Alice and Bob engage in steps 1 and 2 of Protocol SMG for the J_i gaussians in the i -th mixture to obtain vectors $\mathbf{A}_i = [A_{i1} \dots A_{iJ_i}]$ and $\mathbf{B}_i = [B_{i1} \dots B_{iJ_i}]$. Notice that $A_{ij} + B_{ij} = l_{ij}(\mathbf{x})$.
 - b) Alice and Bob engage in the secure logsum protocol with vectors \mathbf{A}_i and \mathbf{B}_i to obtain u_i and z_i i.e. $u_i + z_i = SLOG(\mathbf{A}_i, \mathbf{B}_i)$.
- 2) Bob computes the vector $\mathbf{v} = [v_1 \dots v_N]$ where $v_i = z_i + \ln P(\omega_i)$. Alice forms the vector $\mathbf{u} = [u_1 \dots u_N]$.
- 3) Alice and Bob perform the secure maximum index protocol between vectors \mathbf{u} and \mathbf{v} and Alice obtains $I = SMAX(\mathbf{u}, \mathbf{v})$.

Correctness: If one follows the protocol carefully, it is easy to see that $u_i + v_i$ is equal to $g_i(\mathbf{x})$.

Efficiency: For a given i , there are $(J_i(d+2) + 1)$ SIP calls. Hence, in all, there are $(d+2) \sum_{i=1}^N J_i + N$ SIP calls and 1 SMAX call.

Security: If Protocol SMG and the protocols for SIP, SVAL and SMAX are secure, the above protocol is secure. In case Alice and Bob want to compute additive shares of the likelihood instead of the class label, they can use the SVAL protocol instead of SMAX in the last step.

C. Training Gaussian Mixture from Data

We now focus on a related problem. Let us suppose Alice has K d -component vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$. And she wants to learn a mixture of c gaussians from the data. She can use the iterative *expectation-maximization* (EM) algorithm to estimate the parameters of the c gaussians and the mixture weights. Now, consider the scenario when Bob wants to learn the parameters but Alice does not want to disclose her data to Bob. One solution is for Alice to estimate the parameters herself and then give them to Bob. Another approach is for Alice and Bob to engage in a secure protocol which lets Bob learn the parameters while Alice's data remain private. The latter approach becomes necessary in a scenario where Bob wants to do data mining on combined data from private databases owned by Alice and Charlie. Below, we describe a secure protocol which lets two parties perform such computations.

EM algorithm

We denote the estimate of a particular parameter after the r -th iteration by using a superscript. Thus, $\boldsymbol{\mu}_i^r$, $\boldsymbol{\Sigma}_i^r$ and $P^r(\omega_i)$ denote the mean vector, covariance matrix and the mixture weight for the i -th gaussian after the r -th iteration. For convenience, let us denote the entire parameter set after the r -th iteration by λ^r . At any given iteration, Alice has access to her data and Bob has access to the parameter $\boldsymbol{\Sigma}_i^r$. Alice and Bob have additive shares $\boldsymbol{\mu}_{iA}^r$, $\boldsymbol{\mu}_{iB}^r$ and ℓ_{iA} , ℓ_{iB} such that $\boldsymbol{\mu}_{iA}^r + \boldsymbol{\mu}_{iB}^r = \boldsymbol{\mu}_i^r$ and $\ell_{iA} + \ell_{iB} = \ln P^r(\omega_i)$. We

E Step:

$$P(\omega_i | \mathbf{x}_k, \lambda^r) = \frac{p(\mathbf{x}_k | \omega_i, \boldsymbol{\mu}_i^r, \boldsymbol{\Sigma}_i^r) P^r(\omega_i)}{\sum_{j=1}^c p(\mathbf{x}_k | \omega_j, \boldsymbol{\mu}_j^r, \boldsymbol{\Sigma}_j^r) P^r(\omega_j)} \quad (14)$$

Input: Alice has \mathbf{x}_k , $\boldsymbol{\mu}_{iA}^r$ and ℓ_{iA} ; Bob has $\boldsymbol{\mu}_{iB}^r$, $\boldsymbol{\Sigma}_i^r$ and ℓ_{iB} , $i = 1, 2, \dots, c$.

Output: Alice and Bob obtain u_{ik} and v_{ik} such that $u_{ik} + v_{ik} = \ln P(\omega_i | \mathbf{x}_k, \lambda^r)$.

- 1) Bob forms matrices \mathbf{W}_i for $i = 1, \dots, c$ with $\boldsymbol{\mu}_{iB}^r$, $\boldsymbol{\Sigma}_i^r$ as described in section IV-A and equation (9) (using $(d/2) \ln 2\pi$ instead of $\ln P(\omega_i)$ to compute w_{i0}). With $(\mathbf{x}_k - \boldsymbol{\mu}_{iA}^r)$ as Alice's input and \mathbf{W}_i for $i = 1, \dots, c$ as Bob's input and $N = c$, Alice and Bob engage in steps 1 and 2 of Protocol SMG (section IV-A) to obtain vectors \mathbf{A}'_k and \mathbf{B}'_k .
 - Log-likelihood $\ln p(\mathbf{x}_k | \omega_i, \boldsymbol{\mu}_i^r, \boldsymbol{\Sigma}_i^r)$ is given by equation 6. Notice that using $(\mathbf{x}_k - \boldsymbol{\mu}_{iA}^r)$ in place of \mathbf{x}_k and $\boldsymbol{\mu}_{iB}^r$ in place of $\boldsymbol{\mu}_i^r$ in equation 6 yields the same result as using \mathbf{x}_k and $\boldsymbol{\mu}_i^r$.
 - The sum of the i -th elements, $A'_{ik} + B'_{ik}$, is equal to $\ln p(\mathbf{x}_k | \omega_i, \boldsymbol{\mu}_i^r, \boldsymbol{\Sigma}_i^r)$.
- 2) Alice and Bob obtain vectors \mathbf{A}_k and \mathbf{B}_k , where for each i , $A_{ik} = A'_{ik} + \ell_{iA}$ and $B_{ik} = B'_{ik} + \ell_{iB}$.
 - Notice that $A_{ik} + B_{ik}$ is the logarithm of the numerator in equation 14.
- 3) Alice and Bob engage in the secure logsum protocol with the vectors \mathbf{A}_k and \mathbf{B}_k to obtain y_k and z_k i.e. $y_k + z_k = SLOG(\mathbf{A}_k, \mathbf{B}_k)$.
 - Notice that $y_k + z_k$ is the logarithm of the denominator of equation 14 (follows from equation 4).
- 4) Alice forms vector \mathbf{u}_k , where $u_{ik} = (A_{ik} - y_k)$. Bob forms the vector \mathbf{v}_k , where $v_{ik} = (B_{ik} - z_k)$.
 - $u_{ik} + v_{ik} = \ln P(\omega_i | \mathbf{x}_k, \lambda^r)$.

M Step:

$$\begin{aligned} \boldsymbol{\mu}_i^{r+1} &= \frac{\sum_{k=1}^K P(\omega_i | \mathbf{x}_k, \lambda^r) \mathbf{x}_k}{\sum_{k=1}^K P(\omega_i | \mathbf{x}_k, \lambda^r)} \\ P^{r+1}(\omega_i) &= \frac{\sum_{k=1}^K P(\omega_i | \mathbf{x}_k, \lambda^r)}{K} \\ \boldsymbol{\Sigma}_i^{r+1} &= \frac{\sum_{k=1}^K P(\omega_i | \mathbf{x}_k, \lambda^r) (\mathbf{x}_k - \boldsymbol{\mu}_i^{r+1})(\mathbf{x}_k - \boldsymbol{\mu}_i^{r+1})^T}{\sum_{k=1}^K P(\omega_i | \mathbf{x}_k, \lambda^r)} \end{aligned} \quad (15)$$

Input: Alice has $\mathbf{x}_k, k = 1, \dots, K$. Alice and Bob have K -vectors \mathbf{E} and \mathbf{F} such that $E_k + F_k = \ln P(\omega_i | \mathbf{x}_k, \lambda^r)$.

Output: Alice obtains $\boldsymbol{\mu}_{iA}^{r+1}$, ℓ_{iA} ; Bob obtains $\boldsymbol{\mu}_{iB}^{r+1}$, $\boldsymbol{\Sigma}_i^{r+1}$ and ℓ_{iB} . ($\boldsymbol{\mu}_{iA}^{r+1} + \boldsymbol{\mu}_{iB}^{r+1} = \boldsymbol{\mu}_i^{r+1}$ and $\ell_{iA} + \ell_{iB} = \ln P^{r+1}(\omega_i)$).

- 1) Alice and Bob engage in the secure logsum protocol with vectors \mathbf{E} and \mathbf{F} to obtain e and f i.e. $e + f = SLOG(\mathbf{E}, \mathbf{F})$.
- 2) Alice computes $\ell_{iA} = e - \ln K$, and Bob computes $\ell_{iB} = f$.

3) For $j = 1, 2, \dots, d$:

Let \mathbf{h}_j be the K -vector formed by the j -th elements of $\mathbf{x}_1, \dots, \mathbf{x}_K$. Alice and Bob engage in the secure logsum protocol with vectors $\mathbf{E} + \ln \mathbf{h}_j$ and \mathbf{F} to obtain e' and f' i.e. $e' + f' = SLOG(\mathbf{E} + \ln \mathbf{h}_j, \mathbf{F})$.

- Notice that $(e' - e) + (f' - f) = \ln \mu_{ij}^{r+1}$, the j -th element of $\boldsymbol{\mu}_i^{r+1}$.

Alice and Bob obtain the j -th elements of $\boldsymbol{\mu}_{iA}^{r+1}$ and $\boldsymbol{\mu}_{iB}^{r+1}$ respectively as a result of $SIP(\exp(e' - e), \exp(f' - f))$.

4) Consider the evaluation of σ_{mn} , the mn -th element of the matrix $\boldsymbol{\Sigma}_i^{r+1}$. We first consider evaluating the mn -th element of $(\mathbf{x}_k - \boldsymbol{\mu}_i^{r+1})(\mathbf{x}_k - \boldsymbol{\mu}_i^{r+1})^T$. As mentioned earlier, this is equivalent to evaluating the mn -th term of $(\bar{\mathbf{x}}_k - \bar{\boldsymbol{\mu}}_i)(\bar{\mathbf{x}}_k - \bar{\boldsymbol{\mu}}_i)^T$, where $\bar{\mathbf{x}}_k = (\mathbf{x}_k - \boldsymbol{\mu}_{iA}^{r+1})$ and $\bar{\boldsymbol{\mu}}_i = \boldsymbol{\mu}_{iB}^{r+1}$. Let the j -th elements of $\bar{\mathbf{x}}_k$ and $\bar{\boldsymbol{\mu}}_i$ be \bar{x}_{kj} and $\bar{\mu}_{ij}$ respectively. Notice that Alice has access to $\bar{\mathbf{x}}_k$ and Bob had access to $\bar{\boldsymbol{\mu}}_i$.

- For $k = 1, \dots, K$, Alice and Bob engage in the secure inner product protocol with vectors $\exp(\gamma_k)[\bar{x}_{km}\bar{x}_{kn}, -\bar{x}_{km}, \bar{x}_{kn}, 1]$ and $[1, \bar{\mu}_{im}, -\bar{\mu}_{im}, \bar{\mu}_{im}\bar{\mu}_{in}]$, where γ_k is a random scalar chosen by Alice. Let Bob obtain the result ϕ_k .
- Alice forms the K -vector $\boldsymbol{\gamma} = [\gamma_1, \dots, \gamma_K]$ and Bob forms the vector $\boldsymbol{\phi} = [\phi_1, \dots, \phi_K]$.

Alice and Bob engage in the secure logsum protocol with vectors $(\mathbf{E} - \boldsymbol{\gamma})$ and $(\mathbf{F} + \ln \boldsymbol{\phi})$ to obtain \bar{e} and \bar{f} i.e. $\bar{e} + \bar{f} = SLOG((\mathbf{E} - \boldsymbol{\gamma}), (\mathbf{F} + \ln \boldsymbol{\phi}))$.

- Notice that $(\bar{e} - e) + (\bar{f} - f) = \ln \sigma_{mn}$, the mn -th element of $\boldsymbol{\Sigma}_i^{r+1}$.

Alice sends $(\bar{e} - e)$ to Bob so that he can calculate σ_{mn} .

At the end of all iterations, Alice sends her shares $\boldsymbol{\mu}_{iA}$ and ℓ_{iA} to Bob so that he can calculate the mean $\boldsymbol{\mu}_i$ and the mixture weight $P(\omega_i)$ for $i = 1, 2, \dots, c$.

Efficiency: We only consider the cost of computations that occur between Alice and Bob. In the E-step, for a given \mathbf{x}_k and for all classes ω_i , there are $c(d+2)$ SIP calls with $(d+1)$ -dimensional vectors and one SIP call involving a c -vector. In the M-step, to compute a mixture weight, there is an SIP call involving a K -vector. To calculate a single mean vector, there are d SIP calls involving K -vectors and d SIP calls with scalars. To calculate each element of the covariance matrix for a given class, there are K SIP calls involving 4-dimensional vectors and one SIP call with a K -vector.

Security: We assume that $k \gg d$ and $d > c$. Until the end of the last iteration, Bob does not learn values of the means or the mixture weights. He does not learn the values of likelihoods or posterior probabilities during the iterations. He does learn the value of the covariance matrix with every iteration. This does leak some information about the distribution of Alice's data vectors but Bob's aim is to learn the distributions. The goal of Alice is to prevent Bob from knowing her individual data vectors and without the mean, Bob cannot gain any knowledge about the data vectors. Another important constraint is that Alice does not learn the values of the parameters and following the protocol closely shows that this holds true.

A. The Forward-Backward Procedure

Consider the forward variable $\alpha_t(i)$ defined as

$$\alpha_t(i) = P(\mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_t, q_t = S_i | \lambda) \quad (16)$$

We can solve for $\alpha_t(i)$ inductively and calculate $P(\mathbf{X} | \lambda)$ as follows:

1) Initialization:

$$\alpha_1(i) = \pi_i b_i(\mathbf{x}_1), \quad 1 \leq i \leq N$$

Input: Bob has the gaussian mixture distribution that defines $b_i(\mathbf{x})$ and the initial state distribution $\pi = \{\pi_i\}$; Alice has an observation \mathbf{x}_1 .

Output: Alice and Bob obtain vectors \mathbf{Q} and \mathbf{R} such that $Q_i + R_i = \ln \alpha_1(i)$.

- Bob forms the matrices \mathbf{W}_{ij} as mentioned in section IV-B. With matrices \mathbf{W}_{ij} and mixture weights c_{jm} as Bob's inputs and \mathbf{x}_1 as Alice's input, they perform steps 1 and 2 of the protocol MoG of section IV-B. Alice and Bob obtain vectors \mathbf{U} and \mathbf{V} . Notice that $U_i + V_i = \ln b_i(\mathbf{x}_1)$.
- Alice forms the vector $\mathbf{Q} = \mathbf{U}$. Bob forms vector \mathbf{R} , where for each i , $R_i = V_i + \ln \pi_i$. Thus, $Q_i + R_i = \ln b_i(\mathbf{x}_1) + \ln \pi_i = \ln \alpha_1(i)$.

2) Induction:

$$\alpha_{t+1}(j) = \left(\sum_{i=1}^N \alpha_t(i) a_{ij} \right) b_j(\mathbf{x}_{t+1})$$

$$\text{where } 1 \leq t \leq T-1, 1 \leq j \leq N$$

Input: Alice and Bob have vectors \mathbf{Q} and \mathbf{R} such that $Q_i + R_i = \ln \alpha_t(i)$. Alice and Bob have U_j and V_j such that $U_j + V_j = \ln b_j(\mathbf{x}_{t+1})$. Bob has the vector $\mathbf{a}_j = [a_{1j}, a_{2j}, \dots, a_{Nj}]$.

Output: Alice and Bob obtain \bar{Q} and \bar{R} such that $\bar{Q} + \bar{R} = \ln \alpha_{t+1}(j)$.

- Alice and Bob engage in the secure logsum protocol with vectors \mathbf{Q} and $(\mathbf{R} + \ln \mathbf{a}_j)$ to obtain y' and z' i.e. $y' + z' = SLOG(\mathbf{Q}, \mathbf{R} + \ln \mathbf{a}_j)$.
- Alice obtains $\bar{Q} = y' + U_j$ and Bob obtains $\bar{R} = z' + V_j$.

3) Termination:

$$P(\mathbf{X} | \lambda) = \sum_{i=1}^N \alpha_T(i).$$

Input: Alice and Bob have vectors \mathbf{Q} and \mathbf{R} such that $Q_i + R_i = \ln \alpha_T(i)$.

Output: Alice and Bob obtain y and z such that $y + z = \ln P(\mathbf{X} | \lambda)$.

- Alice and Bob engage in the secure logsum protocol with vectors \mathbf{Q} and \mathbf{R} to obtain y and z i.e. $y + z = SLOG(\mathbf{Q}, \mathbf{R})$.

Efficiency: In the initialization step, there are $(d+2)MN + N$ SIP calls and N $SMAX/SVAL$ calls involving d -dimensional vectors. In the induction step, for every j and for every t , there is one SIP call with an N -vector. In the

termination step, there is one *SIP* call with an N -vector.

Security: Bob does not learn any \mathbf{x}_k and Alice does not learn any of Bob's parameters. Hence, if the primitives *SIP*, *SMAX* and *SVAL* are secure, the protocol is secure.

We can obtain a similar procedure for a backward variable $\beta_t(i)$ defined as

$$\beta_t(i) = P(\mathbf{x}_{t+1}\mathbf{x}_{t+2}\dots\mathbf{x}_T | q_t = S_i, \lambda) \quad (17)$$

We can solve for $\beta_t(i)$ inductively as follows:

1) Initialization:

$$\beta_T(i) = 1, \quad 1 \leq i \leq N$$

2) Induction:

$$\beta_t(i) = \sum_{j=1}^N a_{ij} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j),$$

where $t = T-1, T-2, \dots, 1, \quad 1 \leq j \leq N$

Input: Alice and Bob have vectors \mathbf{Y} and \mathbf{Z} such that $Y_j + Z_j = \ln \beta_{t+1}(j)$. Alice and Bob have \mathbf{U} and \mathbf{V} such that $U_j + V_j = \ln b_j(\mathbf{x}_{t+1})$. Bob has the vector $\mathbf{a}'_i = [a_{i1}, a_{i2}, \dots, a_{iN}]$.

Output: Alice and Bob obtain \bar{Y} and \bar{Z} such that $\bar{Y} + \bar{Z} = \ln \beta_t(i)$.

a) Alice and Bob engage in the secure logsum protocol with vectors $\mathbf{Y} + \mathbf{U}$ and $(\mathbf{Z} + \mathbf{V} + \ln \mathbf{a}'_i)$ to obtain \bar{Y} and \bar{Z} i.e. $\bar{Y} + \bar{Z} = SLOG(\mathbf{Y} + \mathbf{U}, \mathbf{Z} + \mathbf{V} + \ln \mathbf{a}'_i)$.

B. Viterbi Algorithm

Consider the quantity

$$\delta_t(i) = \max_{q_1, q_2, \dots, q_{t-1}} P[q_1 q_2 \dots q_t = S_i, \mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_t | \lambda] \quad (18)$$

$\delta_t(i)$ is the best score (highest probability) along a single path, at time t , which accounts for the first t observations and ends in state S_i . The procedure for finding the best state sequence can be stated as follows:

1) Initialization:

$$\delta_1(i) = \pi_i b_i(\mathbf{x}_1), \quad \psi_1(i) = 0 \quad 1 \leq i \leq N$$

The procedure is evaluating $\delta_1(i)$ is analogous to the initialization step of the forward backward procedure. After this step, Alice and Bob will have additive shares of $\ln \delta_1(i)$.

2) Recursion:

$$\begin{aligned} \delta_t(j) &= \left(\max_{1 \leq i \leq N} [\delta_{t-1}(i) a_{ij}] \right) b_j(\mathbf{x}_t) \\ \psi_t(j) &= \operatorname{argmax}_{1 \leq i \leq N} [\delta_{t-1}(i) a_{ij}] \\ \text{where} \quad & 2 \leq t \leq T, 1 \leq j \leq N \end{aligned}$$

Input: Alice and Bob have vectors \mathbf{Q} and \mathbf{R} such that $Q_i + R_i = \ln \delta_{t-1}(i)$. Alice and Bob have U and V such that $U + V = \ln b_j(\mathbf{x}_t)$. Bob has the vector $\mathbf{a}_j = [a_{1j}, a_{2j}, \dots, a_{Nj}]$.

Output: Alice and Bob obtain \bar{Q} and \bar{R} such that $\bar{Q} + \bar{R} = \ln \delta_t(j)$. Alice obtains $\psi_t(j)$.

a) Alice and Bob engage in the secure maximum value protocol with vectors \mathbf{Q} and $(\mathbf{R} + \ln \mathbf{a}_j)$ to obtain y and z . They also perform *SMAX* on the same vectors and Alice obtains the result which is equal to $\psi_t(j)$.

b) Alice computes $\bar{Q} = y + U$ and Bob computes $\bar{R} = z + V$.

3) Termination:

$$P^* = \max_{1 \leq i \leq N} [\delta_T(i)] \quad q_T^* = \operatorname{argmax}_{1 \leq i \leq N} \delta_T(i).$$

Alice and Bob will use *SVAL* on their additive shares of $\ln \delta_T(i)$ for all i to evaluate $\ln P^*$. Similarly, they engage in *SMAX* on their shares and Alice obtains the result q_T^* .

4) Path backtracking:

$$q_t^* = \psi_{t+1}(q_{t+1}^*) \quad t = T-1, T-2, \dots, 1.$$

Alice, who has access to q_t and ψ_t , can evaluate the path sequence. Notice that Bob could be made to get this result instead of Alice if we let Bob learn the values of ψ_t and q_t in steps 2 and 3 instead of Alice.

Security and efficiency considerations for this protocol are similar to what was discussed with regard to the Forward Backward procedure (section V-A).

C. HMM Training

In the above formulation, we assumed that Bob had already trained his HMMs. Let us consider the case when Alice has all the training data and Bob wants to train a HMM using her data. Below, we show how Bob can securely reestimate parameters of his HMM.

Consider the variables

$$\begin{aligned} \gamma_t(i) &= P(q_t = S_i | \mathbf{X}, \lambda) = \frac{(\alpha_t(i) \beta_t(i))}{P(\mathbf{X} | \lambda)} \\ \xi_t(i, j) &= P(q_t = S_i, q_{t+1} = S_j | \mathbf{X}, \lambda) \\ &= \frac{(\alpha_t(i) a_{ij} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j))}{P(\mathbf{X} | \lambda)} \end{aligned}$$

In the previous subsections, we have shown how Alice and Bob can obtain additive shares of $\ln \alpha_t(i)$ (Q_i and R_i), $\ln \beta_t(i)$ (\bar{Y} and \bar{Z}), $\ln b_j(\mathbf{x}_{t+1})$ (U_j and V_j), $\ln \beta_{t+1}(j)$ (Y_j and Z_j) and $\ln P(\mathbf{X} | \lambda)$ (y and z). It is easy to see that using these shares, Alice and Bob can compute additive shares e_t, g_t and f_t, h_t such that $e_t + f_t = \ln \xi_t(i, j)$ and $g_t + h_t = \ln \gamma_t(i)$. Alice computes $g_t = Q_i + \bar{Y} - y$ and $e_t = Q_i + U_j + Y_j - y$. Bob computes $h_t = R_i + \bar{Z} - z$ and $f_t = R_i + \ln a_{ij} + V_j + Z_j - z$.

The variables π_i and a_{ij} can then be re-estimated as follows:

$$\begin{aligned} \bar{\pi}_i &= \gamma_1(i) \\ \bar{a}_{ij} &= \frac{\sum_{t=1}^{T-1} \xi_t(i, j)}{\sum_{t=1}^{T-1} \gamma_t(i)}. \end{aligned}$$

Input: Alice and Bob have $(T-1)$ -vectors \mathbf{e} and \mathbf{f} such that $e_t + f_t = \ln \xi_t(i, j)$. They also have vectors \mathbf{g} and \mathbf{h} such that

$$g_t + h_t = \ln \gamma_t(i).$$

Output: Bob obtains $\ln \bar{a}_{ij}$.

- 1) Alice and Bob engage in secure logsum protocol with vectors \mathbf{e} and \mathbf{f} to obtain \bar{e} and \bar{f} . They also engage in the secure logsum protocol with vectors \mathbf{g} and \mathbf{h} to obtain \bar{g} and \bar{h} respectively.
- 2) Alice sends $(\bar{e}-\bar{g})$ to Bob. Bob computes $(\bar{e}-\bar{g})+(\bar{f}-\bar{h})$ to obtain $\ln \bar{a}_{ij}$.

Notice that instead of Bob obtaining the final result, Alice and Bob can have additive shares of $\ln \bar{a}_{ij}$. Protocols for forward-backward and viterbi algorithms will then have to be modified so that Alice and Bob have additive shares of the vector $\ln \mathbf{a}_j$.

As for the gaussian mixture distributions $b_i(\mathbf{x})$, Bob can learn them from Alice's data securely as we have shown in section IV-C. We emphasize here that Bob does not learn all the parameters in every iteration. He learns the mean vector for every component gaussian only after the last iteration. He does learn the covariance matrix in every iteration but quantities used to calculate the covariance matrix are additive shares which does not help him in inferring Alice's data. The example shown in section IV-C uses two parties but it can be generalized to the case where Bob learns from multiple parties. In that case, learned statistics are averaged and provide an additional layer of security for the data providers.

VI. DISCUSSION

In this section, we discuss the computational efficiency considerations of protocols presented above. As mentioned earlier, efficiency of the protocols was evaluated in terms of primitives and absolute measures were not provided. This is due to the fact that efficiency of the primitives themselves varies widely and depends on how the primitives are implemented.

If one follows all the protocols carefully, efficiency mainly depends on the computational complexity of the *SIP* primitive. We shall focus on one particular implementation of this primitive: secure inner product using homomorphic encryption proposed by [9] (see appendix I, the reference provides proof that the protocol is correct and secure).

To validate the secure model we ran experiments performing learning and classification. The experiments were run using a MATLAB implementation and tested both the Gaussian mixture models and the hidden Markov models. Simulations were performed twice using the secure and the non-secure (traditional) methods. In all cases the results from both the secure and non-secure simulations were numerically identical as we have predicted. The secure versions were obviously less efficient due to the increased computational cost of the cryptographic operations and the increased network traffic. We did not study the communications complexity in these experiments and rather focused on the computational load, which is the primary bottleneck [18]. One simulation used the a generalized version [19] of the Paillier public-key scheme [20]. We used cryptographic keys of 1024 bits and the cryptosystem was implemented in Java. The computational load of this algorithm coupled with a non-optimal implementation resulted into a processing time per input vector in the order of a few seconds. An alternative implementation using algebraic primitives, which leak some information but are more

computationally efficient, resulted into a significant speedup of less than a second's time processing per input vector. As shown by the last experiment, the choice of implementation for a primitive (for example, *SIP* using algorithm in appendix I instead of *SIP* using algorithm in appendix II) significantly impacts performance, communication complexity and security. A wise choice will have to balance tradeoffs such as computational efficiency and network bandwidth as opposed to security/privacy. A discussion of these issues is out of the scope of this paper since it is a lengthy research project of its own and a moving target given the continuous discoveries of increasingly efficient protocols by the cryptography community.

The figures we have obtained were using non-optimized implementations, as noted by [18] careful implementation can produce significant speedups in computation. For practical implementations it is also possible (and recommended) that specialized hardware is used for the cryptography layer which can result into dramatic performance improvements.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we have presented an implementation of privacy preserving hidden Markov model and Gaussian mixtures computations. We first proposed a simple privacy-preserving protocol for computing logsums. Using primitives for computing scalar products and maxima, we proposed secure protocols for classification using Gaussian mixture models. We then proposed secure protocols for the *forward-backward* algorithm, the *viterbi* algorithm and HMM training. The protocols are defined modularly in terms of primitives so that future advances in cryptography, which will hopefully provide more robust and efficient protocols, can be readily employed in our framework by straightforward replacement. The approach we have taken also illustrates the process required to transform a signal processing algorithm to its privacy preserving version. Other data processing and classification algorithms can also be described in terms of secure primitives and easily reformulated for secure multiparty computations.

This paper is intended to be a starting point for secure audio frameworks and because of that it exposes a lot of new research directions which warrant more attention. One of these directions includes the design of alternative classifiers and algorithms using this process, and there is still ongoing work on the building block primitives (*SIP*, *SMAx*, *SVAl*, etc) themselves. These are all topics that present plenty of opportunities to explore efficiency and security and their tradeoffs. We expect these to be fruitful areas of research in the near future. It is our hope that a migration towards secure algorithms can help promote a more open collaboration setting where parties can freely exchange data and algorithms without legal and privacy issues.

The authors would like to acknowledge the help and influence of Shai Avidan in the making of this work. The authors also wish to thank anonymous reviewers for their comments and suggestions.

The following protocol is based on homomorphic encryption and was proposed by [9]. Let the triple $(\text{Ge}, \text{En}, \text{De})$ denote a public-key homomorphic cryptosystem (probabilistic polynomial time algorithms for key-generation, encryption and decryption). The key generation algorithm generates a valid pair (sk, pk) of private and public keys for a security parameter k . The encryption algorithm En takes as an input a plaintext m , a random value r and a public key pk and outputs the corresponding ciphertext $\text{En}(\text{pk}; m, r)$. The decryption algorithm De takes as an input a ciphertext c and a private key sk (corresponding to the public key pk) and outputs a plaintext $\text{De}(\text{sk}; c)$. It is required that $\text{De}(\text{sk}; \text{En}(\text{pk}; m, r)) = m$. A public-key cryptosystem is *homomorphic* if $\text{En}(\text{pk}; m_1, r_1) \cdot \text{En}(\text{pk}; m_2, r_2) = \text{En}(\text{pk}; m_1 + m_2, r_1 + r_2)$, where $+$ is a group operation and \cdot is a groupoid operation.

Inputs: Private vectors \mathbf{x} and \mathbf{y} with Bob and Alice respectively.

Outputs: Shares a and b such that $a + b = \mathbf{x}^T \mathbf{y}$.

- 1) Setup phase. Bob:
 - generates a private and public key pair (sk, pk) .
 - sends pk to Alice.
- 2) For $i \in \{1, \dots, d\}$, Bob:
 - generates a random new string r_i .
 - sends $c_i = \text{En}(\text{pk}; x_i, r_i)$ to Alice.
- 3) Alice:
 - sets $z \leftarrow \prod_{i=1}^d c_i^{y_i}$.
 - generates a random plaintext b and a random nonce r' .
 - sends $z' = z \cdot \text{En}(\text{pk}; -b, r')$ to Bob.
- 4) Bob computes $a = \text{De}(\text{sk}; z') = \mathbf{x}^T \mathbf{y} - b$.

See [9] for a proof that the protocol is correct and secure.

APPENDIX II
SECURE INNER PRODUCT FROM OBLIVIOUS POLYNOMIAL
EVALUATION

[8] proposes an elegant protocol for oblivious evaluation of multivariate polynomials using oblivious transfer [21] as a cryptographic primitive. It can be easily modified to securely evaluate dot products. Let Alice represent each x_i as $x_i = \sum_j a_{ij} 2^{j-1}$ with $a_{ij} \in \{0, 1\}$. Let $v_{ij} = 2^{j-1} y_i$. Notice that for each i , $1 \leq i \leq d$, $\sum_j a_{ij} v_{ij} = x_i y_i$. The idea is to have Bob prepare v_{ij} and have Alice get those v_{ij} with $a_{ij} = 1$ in some secret way. This is achieved as follows: Bob prepares the pair $(r_{ij}, v_{ij} + r_{ij})$ for randomly chosen r_{ij} and Alice runs independent Oblivious Transfer with Bob to get r_{ij} if $a_{ij} = 0$ and $v_{ij} + r_{ij}$ otherwise. At the end of the protocol, Alice will obtain $\sum_i \sum_j (a_{ij} v_{ij} + r_{ij}) = \sum_i x_i y_i + \sum_{i,j} r_{ij}$. Bob will have $-\sum_{i,j} r_{ij}$. Thus, Alice and Bob will have additive shares of the desired dot product.

[8] proves that this protocol is secure when the parties are semi-honest. The efficiency of the protocol depends on the implementation of oblivious transfer.

[11] proposes an algebraic approach which assumes that the dimensionality is even. Let us define \mathbf{x}_1 as the $d/2$ dimensional vector consisting of the first $d/2$ elements of \mathbf{x} and \mathbf{x}_2 as the vector consisting of the last $d/2$ elements of \mathbf{x} . We observe that $\mathbf{x}^T \mathbf{y} = \mathbf{x}_1^T \mathbf{y}_1 + \mathbf{x}_2^T \mathbf{y}_2$. Alice and Bob jointly generate a random invertible $d \times d$ matrix M . Alice computes $\mathbf{x}' = \mathbf{x}^T M$, splits it as \mathbf{x}'_1 and \mathbf{x}'_2 and sends \mathbf{x}'_2 to Bob. Bob computes $\mathbf{y}' = M^{-1} \mathbf{y}$, splits it as \mathbf{y}'_1 and \mathbf{y}'_2 and sends \mathbf{y}'_1 to Alice. Alice computes $\mathbf{x}'_1 \mathbf{y}'_1$ and Bob computes $\mathbf{x}'_2 \mathbf{y}'_2$ so that their sum is equal to the desired result.

This protocol has little communication and computational overhead compared to the cryptographic protocols but it comes at the cost of security. Alice and Bob learn $d/2$ linear equations for the d unknowns that constitute the other party's vector which leaks a lot of information. Hence, it is important that the same matrix M should be used when this protocol is used multiple times with the same vector \mathbf{x} (or \mathbf{y}). [3] has analyzed this protocol which showed serious security flaws and hence this is not practical when security is crucially important.

APPENDIX IV
PERMUTE PROTOCOL

This protocol was proposed in [15].

Input: Alice and Bob have d -component vectors \mathbf{x} and \mathbf{y} . Bob has a random permutation π .

Output: Alice and Bob obtain \mathbf{q} and \mathbf{s} such that $\mathbf{q} + \mathbf{s} = \pi(\mathbf{x}) + \pi(\mathbf{y})$.

- 1) Alice generates public and private keys for a homomorphic cryptosystem and sends the public key to Bob. Let $E(\cdot)$ denote encryption with Alice's public key.
- 2) Alice encrypts each element of \mathbf{x} and sends the resulting vector $\bar{\mathbf{x}}$ to Bob.
- 3) Bob generates a random vector \mathbf{r} and computes a new vector θ where $\theta_i = \bar{x}_i E(r_i) = E(x_i + r_i)$, for $i = 1, \dots, d$.
- 4) Bob permutes θ and sends $\pi(\theta)$ to Alice. Alice decrypts the vector to obtain \mathbf{q} .
- 5) Bob computes $\mathbf{y} - \mathbf{r}$ and then permutes it using π to obtain $\mathbf{s} = \pi(\mathbf{y} - \mathbf{r})$.

Alice and Bob engage in the above permute protocol twice, the second time with their roles interchanged. After this is done, Alice and Bob will have two vectors whose sum will be a random permutation of the original sum but neither of them will know what the permutation is.

REFERENCES

- [1] A. C.-C. Yao, "Protocols for secure computation," in *Proc. of the 23rd IEEE Symposium on Foundations of Computer Science*, 1982, pp. 160–164.
- [2] J. Vaidya and C. Clifton, "Privacy preserving k-means clustering over vertically partitioned data," in *ACM SIGKDD Conf on Knowledge Discovery and Data Mining*, 2003. [Online]. Available: citeseer.ist.psu.edu/vaidya03privacypreserving.html

- [3] E. Kiltz, G. Leander, and J. Malone-Lee, "Secure computation of the mean and related statistics," in *Proceedings of the Theory of Cryptography Conference*, ser. Lecture Notes in Computer Science, vol. 3378, 2005, pp. 283–302. [Online]. Available: <http://eprint.iacr.org/2004/359.pdf>
- [4] S. Avidan and M. Butman, "Blind vision," in *ECCV*, 2006. [Online]. Available: <http://www.merl.com/reports/docs/TR2006-006.pdf>
- [5] O. Goldreich, "Secure multi-party computation," Working Draft, 2000. [Online]. Available: citeseer.ist.psu.edu/goldreich98secure.html
- [6] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, Feb 1989.
- [7] S. Laur and H. Lipmaa, "Additive conditional disclosure of secrets and applications," *Cryptology ePrint Archive*, Report 2005/378, 2005, <http://eprint.iacr.org/>. [Online]. Available: citeseer.ist.psu.edu/laur05additive.html
- [8] Y.-C. Chang and C.-J. Lu, "Oblivious polynomial evaluation and oblivious neural learning," in *Advances in Cryptology, Asiacrypt '01*, ser. Lecture Notes in Computer Science, vol. 2248, 2001, pp. 369–384. [Online]. Available: citeseer.ist.psu.edu/chang01oblivious.html
- [9] B. Goethals, S. Laur, H. Lipmaa, and T. Mielikainen, "On private scalar product computation for privacy-preserving data mining," in *Intl. Conference on Information Security and Cryptology*, ser. Lecture Notes in Computer Science, C. Park and S. Chee, Eds., vol. 2506, 2004, pp. 104–120. [Online]. Available: <http://citeseer.ist.psu.edu/goethals04private.html>
- [10] W. Du and M. J. Atallah, "Privacy-preserving cooperative statistical analysis," in *Proceedings of the 17th Annual Computer Security Applications Conference*, New Orleans, Louisiana, December 2001. [Online]. Available: <http://citeseer.ist.psu.edu/549869.html>
- [11] W. Du and Z. Zhan, "A practical approach to solve secure multi-party computation problems," in *Proceedings of New Security Paradigms Workshop*, Virginia Beach, Virginia, USA, September 23-26 2002. [Online]. Available: citeseer.ist.psu.edu/du02practical.html
- [12] I. Ioannidis, A. Grama, and M. Atallah, "A secure protocol for computing dot-products in clustered and distributed environments," in *Proceedings of the Intl. Conf. on Parallel Processing*, Vancouver, Canada, 2002. [Online]. Available: <http://ieeexplore.ieee.org/iel5/8068/22315/01040894.pdf>
- [13] P. Ravikumar, W. W. Cohen, and S. E. Fienberg, "A secure protocol for computing string distance metrics," in *Proceedings of the Workshop on Privacy and Security Aspects of Data Mining*, Brighton, UK, 2004, pp. 40–46. [Online]. Available: <http://citeseer.ist.psu.edu/719096.html>
- [14] J. Vaidya and C. Clifton, "Privacy preserving association rule mining in vertically partitioned data," in *Proceedings of the Intl. Conf. on Knowledge Discovery and Data Mining*. Edmonton, Canada: ACM SIGKDD, 2002. [Online]. Available: <http://citeseer.ist.psu.edu/vaidya02privacy.html>
- [15] M. J. Atallah, F. Kerschbaum, and W. Du, "Secure and private sequence comparisons," in *Proceedings of Workshop on Privacy in the Electronic Society*, Washington, DC, USA, October 2003. [Online]. Available: <http://citeseer.ist.psu.edu/atallah03secure.html>
- [16] I. F. Blake and V. Kolesnikov, "Strong conditional oblivious transfer and computing on intervals," in *Proceedings of Advances in Cryptology - ASIACRYPT'04*, ser. LNCS, P. J. Lee, Ed., vol. 3329. Springer-Verlag, 2004, pp. 515–529. [Online]. Available: <http://citeseer.ist.psu.edu/blake04strong.html>
- [17] H.-Y. Lin and W.-G. Tzeng, "An efficient solution to the millionaires' problem based on homomorphic encryption," in *Proc of Intl. Conf. on Applied Cryptography and Network Security*, ser. LNCS, vol. 3531. Springer-Verlag, 2005, pp. 456–466. [Online]. Available: <http://citeseer.ist.psu.edu/lin05efficient.html>
- [18] Z. Yang, R. Wright, and H. Subramaniam, "Experimental analysis of a privacy-preserving scalar product protocol," *Intl Jnl of Computer Systems Science and Engineering*, vol. 21, no. 1, pp. 47–52, 2006. [Online]. Available: <http://www.cs.stevens.edu/~rwright/Publications/csse06.pdf>
- [19] I. Damgard and M. Jurik, "A generalisation, simplification and some applications of paillier's probabilistic public-key system," in *Proceedings of the Intl. Workshop on Practice and Theory in Public Key Cryptography*, ser. Lecture Notes in Computer Science, vol. 1992, 2001, pp. 119–136. [Online]. Available: <http://citeseer.ist.psu.edu/383099.html>
- [20] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Proceedings of Advances in Cryptology - EUROCRYPT '99*, ser. Lecture Notes in Computer Science, J. Stern, Ed., vol. 1592, 1999, pp. 223–238. [Online]. Available: citeseer.ist.psu.edu/paillier99publickey.html
- [21] M. Naor and B. Pinkas, "Oblivious transfer and polynomial evaluation," in *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, 1999, pp. 245–254. [Online]. Available: <http://portal.acm.org/citation.cfm?id=301312>

PLACE
PHOTO
HERE

Madhusudana Shashanka received BE (Hons) in computer science from the Birla Institute of Technology and Science, Pilani, India in June 2003. He is currently a PhD student with the Department of Cognitive & Neural Systems and the Hearing Research Center at Boston University.

PLACE
PHOTO
HERE

Paris Smaragdis is a member of the research staff in Mitsubishi Electric Research Laboratories in Cambridge MA, USA. Prior to that position he was at the Massachusetts Institute of Technology where he completed his graduate and postdoctoral training. His interests are computational audition, scene analysis and the intersection of machine learning with signal processing.