

A novel transfer learning model for predictive analytics using incomplete multimodality data

Xiaonan Liu^a, Kewei Chen^b, David Weidman^b, Teresa Wu^a, Fleming Lure^c, Jing Li^a, and for the Alzheimer's Disease Neuroimaging Initiative*

^aIndustrial Engineering, Arizona State University, Tempe, AZ, USA; ^bBanner Alzheimer's Institute, Phoenix, AZ, USA; ^cMS Technologies, Rockville, MD, USA

ABSTRACT

Multimodality datasets are becoming increasingly common in various domains to provide complementary information for predictive analytics. One significant challenge in fusing multimodality data is that the multiple modalities are not universally available for all samples due to cost and accessibility constraints. This situation results in a unique data structure called an Incomplete Multimodality Dataset. We propose a novel Incomplete-Multimodality Transfer Learning (IMTL) model that builds a predictive model for each sub-cohort of samples with the same missing modality pattern, and meanwhile couples the model estimation processes for different sub-cohorts to allow for transfer learning. We develop an Expectation-Maximization (EM) algorithm to estimate the parameters of IMTL and further extend it to a collaborative learning paradigm that is specifically valuable for patient privacy preservation in health care applications. We prove two advantageous properties of IMTL: the ability for out-of-sample prediction and a theoretical guarantee for a larger Fisher information compared with models without transfer learning. IMTL is applied to diagnosis and prognosis of Alzheimer's disease at an early stage called Mild Cognitive Impairment using incomplete multimodality imaging data. IMTL achieves higher accuracy than competing methods without transfer learning.

ARTICLE HISTORY

Received 17 July 2018
Accepted 14 June 2020

KEYWORDS

Incomplete multimodality data; transfer learning; predictive analytics; health care

1. Introduction

Multimodality datasets are becoming increasingly common in various domains and provide complementary information for predictive analytics. For example, in health care, images of different types such as structural Magnetic Resonance Imaging (MRI) and FluDeoxyGlucose Positron Emission Tomography (FDG-PET) provide complementary information about the organ of interest, which allows a predictive model to be built to accurately diagnosing a certain disease. In manufacturing, data collected from multiple different types of sensors provide complementary information about the process and product, allowing for more accurate assessment of process and product quality.

One important challenge for integration of multimodality datasets in building a predictive model is that the multiple different modalities are not universally available for all the samples. Take the diagnosis of Alzheimer's Disease (AD) – a fatal neurological disorder – as an example. Figure 1 shows the special “Incomplete Multimodality Dataset (IMD)” we are focusing on in this article, which includes three complementary image modalities, i.e., MRI, FDG-PET, and amyloid-PET for detection of AD at an early stage called Mild

Cognitive Impairment (MCI). In the recently published expert consensus criteria by the National Institute of Aging and Alzheimer's Association, the use of multimodality images for early detection of AD has been highly recommended (Albert *et al.*, 2011). In Figure 1, each sub-cohort consists of patients who have the same availability of modalities. Different sub-cohorts have different missing modality patterns. The reasons for IMD are multifold: some imaging equipment such as PET is costly and only available in limited clinics; some modalities are not accessible to patients due to insurance coverage; it is not safe to put patients with some pre-existing conditions through a certain imaging examination.

If we applied existing methods to model IMD, there would be three options:

1. Filling in missing data using some imputation algorithms (He *et al.*, 2017). As an IMD dataset misses the entire modality/modalities and not individual features within a modality, there are too many missing values to fill in. The resulting dataset may have poor quality.

CONTACT Jing Li  jinglz@asu.edu

*Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/24725854.2020.1798569>.

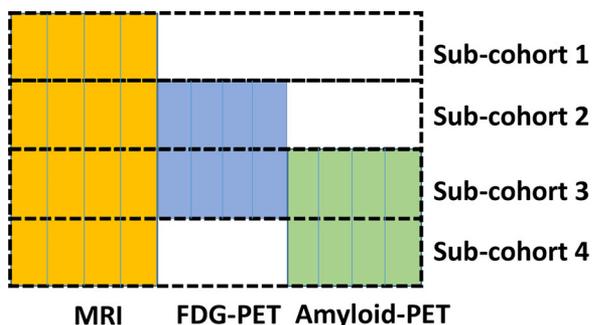


Figure 1. An example of the IMD, in which MRI, FDG-PET, and amyloid-PET are considered as three modalities. Columns within each modality represent features extracted from the image. Each sub-cohort consists of patients with the same availability of modalities.

2. **Separate Modeling (SM).** Since each sub-cohort has different availability for the modalities, SM builds a separate predictive model for each sub-cohort using the available modality/modalities within that sub-cohort. The limitation of SM is obvious: as each model can only use the data specific to the corresponding sub-cohort, sample size shortage may prevent building a robust model.
3. **All Available Data Modeling (AADM).** To build a model for each sub-cohort l , one can incorporate data from another sub-cohort l' whose available modalities include those in l . For example, to build a model for sub-cohort 2, one can combine the data of MRI and PDG-PET in sub-cohort 3. As all the available data is used regardless of the sub-cohort in which the data resides, this approach is called AADM. Compared with SM, AADM alleviates the sample size shortage. However, it requires data pooling from different sub-cohorts. In reality, different sub-cohorts likely correspond to different health institutions or hospitals, data pooling is not easy due to the concern of patient privacy. Furthermore, it is known that AADM may not produce statistical consistent estimators; when it is used to estimate a covariance matrix, the estimate may not be positive definite (Little and Rubin, 2002).

In this article, we propose a novel Incomplete-Multimodality Transfer Learning (IMTL) model to tackle the limitations of existing methods. IMTL models all the sub-cohorts simultaneously under a unified framework. In this way, knowledge obtained from the modeling of each sub-cohort can be “transferred” to help the modeling of other sub-cohorts. This makes IMTL a transfer learning model. Compared with SM, IMTL is not limited by the sample size of each sub-cohort. Compared with AADM, IMTL estimates model parameters in an integrated manner, which overcomes the limitations of AADM in the lack of positive definiteness and consistency guarantees. Furthermore, we propose two algorithms for parameter estimation of IMTL: with and without data pooling. The latter is a computational framework that includes iterative communication between a global learner and local learners residing within each sub-cohort/institution. This allows for between-institutional collaborative model estimation without the need for data

pooling. This is particularly important for patient privacy preservation in health care applications of IMTL.

This article is organized as follows: Section 2 provides a literature review. Section 3 presents the IMTL. Section 4 investigates unique properties of IMTL. Section 5 presents case studies. Section 6 is the conclusion.

2. Literature review

This article primarily intersects with the research area of statistical and machine learning models using data missed in modalities. To our best knowledge, this area only has limited work. In what follows, we review each related paper in detail.

Yuan *et al.* (2012) proposed an incomplete MultiSource Feature (iMSF) learning method, which used an l_{21} penalty to enforce the same features within each modality to be selected across different sub-cohorts. One limitation of iMSF is that it cannot do “out-of-sample prediction”. That is, if a modality-wise missing pattern is not included in the training data, iMSF cannot make a prediction on new samples with that missing pattern. Also, the l_{21} -enabled feature selection scheme is most effective if different modalities have little correlation. To overcome the limitations of iMSF, Xiang *et al.* (2014) proposed an Incomplete Source-Feature Selection (ISFS) model. The main idea was to estimate a set of common coefficients across different sub-cohorts and specific coefficients to account for the uniqueness of each sub-cohort. To gain this flexibility, ISFS needs to estimate many parameters.

Thung *et al.* (2014) developed a matrix completion method, which selected samples and features in the original dataset to produce a smaller dataset. This was done by using the group-lasso-based multitask learning algorithm twice on features and samples, respectively. Then, standard missing data imputation algorithms were applied to the reduced dataset and classifiers were built on the imputed data. Although the proposed idea of data reduction is novel, imputation would still have to be used.

Liu *et al.* (2017) proposed a View-Aligned Hypergraph Learning (VAHL) method. VAHL divided the dataset into several views based on the availability of different modalities. A hypergraph of subjects was constructed on each view. Then, the hypergraphs were fused by a view-aligned regularizer under a classification framework. VAHL has a novel perspective of exploiting subject relationships using hypergraphs to naturally get around the issue of missing modalities. Also because of this “subject” perspective, the model has to be retrained from scratch every time new data become available. Also, VAHL has many parameters to estimate.

Li *et al.* (2014) proposed a Deep Learning (DL) framework specifically for imaging data. The basic idea was to train a 3-D Convolutional Neural Network (CNN) to establish a voxel-wise mapping from an MRI image to an FDG-PET image based on a dataset with both images available. Then, the CNN could be used to create a “pseudo” FDG-PET from an MRI image for any patient with a missing

FDG-PET. To perform a diagnosis or prognosis for a patient, both MRI and FDG-PET (real or pseudo) would be used. This work represents one of the pioneering works that introduced DL into imaging-based AD research. On the other hand, as MRI measures brain structure and FDG-PET measures brain function, crafting one from the other may not be biologically valid, even though this is possible from a pure data-driven perspective. Furthermore, there is a concern of uncertainty propagation as the uncertainty in establishing the voxel-wise mapping between MRI and FDG-PET will propagate to the uncertainty of the pseudo FDG-PET, which further affects the diagnosis and prognosis based on the pseudo FDG-PET. Also, this approach was developed to model two image modalities, and is not directly applicable to datasets with more than two modalities and complicated missing patterns.

In summary, limited work has been done to develop statistical models for IMD data. All the above-reviewed models, despite their specific weakness, share some common limitations: (i) most models cannot do out-of-sample prediction, which limits broader utilization; (ii) model estimation needs data pooling from different sub-cohorts. If the sub-cohorts correspond to different health institutions, which is typically the case, protection of the patient's data privacy is a concern; (iii) although it shows an empirically good performance on specific datasets, there is a lack of theoretical study on why the performance is guaranteed.

3. Development of the IMTL model

For notational simplicity, we present our model development in the context of three modalities; however, the model is generalizable to more modalities. For example, the three modalities can be MRI, FDG-PET and amyloid-PET as shown in Figure 1. There are four sub-cohorts corresponding to different availabilities of the modalities: MRI alone; MRI & FDG-PET; MRI & amyloid-PET; all three modalities.

Let k be the index for modalities, $k = 1, 2, 3$; l be the index for sub-cohorts, $l = 1, 2, 3, 4$; and i be the index for samples/patients, $i = 1, \dots, n$. Denote the sample size of each sub-cohort by n_l . $\sum_{l=1}^4 n_l = n$. Furthermore, let $\mathbf{x}_i^{(kl)}$ contain features in modality k for patient i in sub-cohort l . Let $y_i^{(l)}$ be the response variable for patient i in sub-cohort l . We propose two IMTL models: a predictive model and a classification model. Both models are useful in disease diagnosis and prognosis.

3.1. IMTL predictive model

3.1.1. Formulation and estimation

Consider the joint distribution of $y_i^{(l)}$, $\mathbf{x}_i^{(2l)}$, and $\mathbf{x}_i^{(3l)}$ given $\mathbf{x}_i^{(1l)}$ to be multivariate normal, that is to say

$$\left(y_i^{(l)}, \mathbf{x}_i^{(2l)}, \mathbf{x}_i^{(3l)} \mid \mathbf{x}_i^{(1l)} \right) \sim \text{MVN} \left(\boldsymbol{\mu} \left(\mathbf{x}_i^{(1l)} \right), \boldsymbol{\Sigma} \right). \quad (1)$$

Here, we consider $\mathbf{x}_i^{(1l)}$ (e.g., features of MRI) to be fixed covariates instead of random variables based on the fact that MRI is in the standard clinical care of AD and thus available

to all the patients. However, $\mathbf{x}_i^{(1l)}$ could be considered as random in the most general formulation, doing so would need a joint distribution of $(y_i^{(l)}, \mathbf{x}_i^{(2l)}, \mathbf{x}_i^{(3l)}, \mathbf{x}_i^{(1l)})$, which requires more parameters to be estimated than the proposed approach.

In Equation (1), $\boldsymbol{\mu}(\cdot)$ is a vector function of covariates:

$$\boldsymbol{\mu} \left(\mathbf{x}_i^{(1l)} \right) = \left(\mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 + \beta_0, \mathbf{x}_i^{(1l)} \mathbf{A}_2 + \mathbf{b}_2, \mathbf{x}_i^{(1l)} \mathbf{A}_3 + \mathbf{b}_3 \right),$$

where $\boldsymbol{\beta}_1, \beta_0, \mathbf{A}_2, \mathbf{b}_2, \mathbf{A}_3, \mathbf{b}_3$ are coefficients. $\boldsymbol{\Sigma}$ in Equation (1) can be written in a more explicit format to include sub-matrices of covariances between the response and each modality and between the modalities:

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_y^2 & \Sigma_{y2} & \Sigma_{y3} \\ \Sigma_{2y} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{3y} & \Sigma_{32} & \Sigma_{33} \end{pmatrix}, \quad (2)$$

Let $\Theta = (\boldsymbol{\Sigma}, \boldsymbol{\beta}_1, \beta_0, \mathbf{A}_2, \mathbf{b}_2, \mathbf{A}_3, \mathbf{b}_3)$ contain all the unknown parameters for the model in Equation (1). Furthermore, let D^{mis} and D^{obs} contain the missing and observed data, respectively. That is,

$$D^{mis} = \left\{ \left\{ \mathbf{x}_i^{(21)}, \mathbf{x}_i^{(31)} \right\}_{i=1}^{n_1}, \left\{ \mathbf{x}_i^{(32)} \right\}_{i=1}^{n_2}, \left\{ \mathbf{x}_i^{(24)} \right\}_{i=1}^{n_4} \right\},$$

and

$$D^{obs} = \left\{ \left\{ \mathbf{x}_i^{(11)}, y_i^{(1)} \right\}_{i=1}^{n_1}, \left\{ \mathbf{x}_i^{(12)}, \mathbf{x}_i^{(22)}, y_i^{(2)} \right\}_{i=1}^{n_2}, \left\{ \mathbf{x}_i^{(13)}, \mathbf{x}_i^{(23)}, \mathbf{x}_i^{(33)}, y_i^{(3)} \right\}_{i=1}^{n_3}, \left\{ \mathbf{x}_i^{(14)}, \mathbf{x}_i^{(34)}, y_i^{(4)} \right\}_{i=1}^{n_4} \right\}.$$

Then, we can write down the complete-data log-likelihood function:

$$\begin{aligned} l(\Theta; D^{obs}, D^{mis}) &= -n \log |\boldsymbol{\Sigma}| - \sum_{l=1}^4 \sum_{i=1}^{n_l} \left(y_i^{(l)} - \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 - \beta_0, \right. \\ &\quad \left. \mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_2, \mathbf{x}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_3 \right) \boldsymbol{\Sigma}^{-1} \\ &\quad \left. \left(y_i^{(l)} - \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 - \beta_0, \mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_2, \mathbf{x}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_3 \right)^T \right). \end{aligned} \quad (3)$$

Since $l(\Theta; D^{obs}, D^{mis})$ includes missing data, we resort to the Expectation-Maximization (EM) algorithm. The general EM framework includes an E-step and an M-step. The E-step is to find the expectation of the complete-data log-likelihood function with respect to the missing data given the observed data and the current parameter estimates. In our case, the E-step is to find:

$$E_{D^{mis} \mid D^{obs}, \Theta^{(t)}} \left(l \left(\Theta; D^{obs}, D^{mis} \right) \right), \quad (4)$$

where $\Theta^{(t)}$ contains the parameter estimates obtained at the t th iteration. Then, the M-step is to update the parameter estimates by maximizing the expectation in the E-step:

$$\Theta^{(t+1)} = \arg \max_{\Theta} E_{D^{mis} \mid D^{obs}, \Theta^{(t)}} \left(l \left(\Theta; D^{obs}, D^{mis} \right) \right). \quad (5)$$

The two steps are iterated until convergence. The challenges in using the general EM framework are to derive the expectation and solve the maximization for a specific model (e.g.,

IMTL in our case). In what follows, we will develop the details of the E-step and M-step for IMTL.

E-step:

Inserting the complete-data likelihood function in Equation (3) into the expectation formula in Equation (4), we need to derive expectations for the statistics contained in a set S :

$$S = \left\{ \begin{array}{l} \mathbf{x}_i^{(21)}, \mathbf{x}_i^{(31)}, \mathbf{x}_i^{(32)}, \mathbf{x}_i^{(24)}, \left(\begin{array}{cc} (\mathbf{x}_i^{(21)})^T \mathbf{x}_i^{(21)} & (\mathbf{x}_i^{(31)})^T \mathbf{x}_i^{(21)} \\ (\mathbf{x}_i^{(21)})^T \mathbf{x}_i^{(31)} & (\mathbf{x}_i^{(31)})^T \mathbf{x}_i^{(31)} \end{array} \right), \\ (\mathbf{x}_i^{(32)})^T \mathbf{x}_i^{(32)}, (\mathbf{x}_i^{(24)})^T \mathbf{x}_i^{(24)} \end{array} \right\}.$$

Next, we need to derive the expectation of each element contained in S . For example, focus on $(\mathbf{x}_i^{(21)}, \mathbf{x}_i^{(31)})$ first. We can derive that:

$$\begin{aligned} (\tilde{\mathbf{x}}_i^{(21)}, \tilde{\mathbf{x}}_i^{(31)}) &\triangleq E \left[\left(\mathbf{x}_i^{(21)}, \mathbf{x}_i^{(31)} \right) \mathbf{x}_i^{(11)}, y_i^{(1)}, \boldsymbol{\Theta}^{(t)} \right] \\ &= \left(\mathbf{x}_i^{(11)} \mathbf{A}_2^{(t)} + \mathbf{b}_2^{(t)}, \mathbf{x}_i^{(11)} \mathbf{A}_3^{(t)} + \mathbf{b}_3^{(t)} \right) \\ &\quad + \left(\boldsymbol{\Sigma}_{2y}^{(t)}, \boldsymbol{\Sigma}_{3y}^{(t)} \right) \left(\sigma_y^{(t)2} \right)^{-1} \left(y_i^{(1)} - \mathbf{x}_i^{(11)} \boldsymbol{\beta}_1^{(t)} - \beta_0^{(t)} \right), \end{aligned} \quad (6)$$

Similarly, the expectations of $\mathbf{x}_i^{(32)}$ and $\mathbf{x}_i^{(24)}$ can be obtained as follows:

$$\begin{aligned} \tilde{\mathbf{x}}_i^{(32)} &\triangleq E \left[\mathbf{x}_i^{(32)} \mathbf{x}_i^{(12)}, \mathbf{x}_i^{(22)}, y_i^{(2)}, \boldsymbol{\Theta}^{(t)} \right] \\ &= \mathbf{x}_i^{(12)} \boldsymbol{\beta}_1^{(t)} + \beta_0^{(t)} + \left(\boldsymbol{\Sigma}_{3y}^{(t)}, \boldsymbol{\Sigma}_{32}^{(t)} \right) \left(\begin{array}{cc} \sigma_y^{(t)2} & \boldsymbol{\Sigma}_{y2}^{(t)} \\ \boldsymbol{\Sigma}_{2y}^{(t)} & \boldsymbol{\Sigma}_{22}^{(t)} \end{array} \right)^{-1} \\ &\quad \times \left(\begin{array}{c} y_i^{(2)} - \mathbf{x}_i^{(12)} \boldsymbol{\beta}_1^{(t)} - \beta_0^{(t)} \\ \mathbf{x}_i^{(22)} - \mathbf{x}_i^{(12)} \mathbf{A}_2^{(t)} - \mathbf{b}_2^{(t)} \end{array} \right), \end{aligned} \quad (7)$$

$$\begin{aligned} \tilde{\mathbf{x}}_i^{(24)} &\triangleq E \left[\mathbf{x}_i^{(24)} \mathbf{x}_i^{(14)}, \mathbf{x}_i^{(34)}, y_i^{(4)}, \boldsymbol{\Theta}^{(t)} \right] \\ &= \mathbf{x}_i^{(14)} \boldsymbol{\beta}_1^{(t)} + \beta_0^{(t)} + \left(\boldsymbol{\Sigma}_{2y}^{(t)}, \boldsymbol{\Sigma}_{23}^{(t)} \right) \left(\begin{array}{cc} \sigma_y^{(t)2} & \boldsymbol{\Sigma}_{y3}^{(t)} \\ \boldsymbol{\Sigma}_{3y}^{(t)} & \boldsymbol{\Sigma}_{33}^{(t)} \end{array} \right)^{-1} \\ &\quad \times \left(\begin{array}{c} y_i^{(4)} - \mathbf{x}_i^{(14)} \boldsymbol{\beta}_1^{(t)} - \beta_0^{(t)} \\ \mathbf{x}_i^{(34)} - \mathbf{x}_i^{(14)} \mathbf{A}_3^{(t)} - \mathbf{b}_3^{(t)} \end{array} \right). \end{aligned} \quad (8)$$

Using Equations (6)–(8), we can further derive the expectations of the second-order elements in S as:

$$\begin{aligned} E \left[\left(\begin{array}{cc} (\mathbf{x}_i^{(21)})^T \mathbf{x}_i^{(21)} & (\mathbf{x}_i^{(31)})^T \mathbf{x}_i^{(21)} \\ (\mathbf{x}_i^{(21)})^T \mathbf{x}_i^{(31)} & (\mathbf{x}_i^{(31)})^T \mathbf{x}_i^{(31)} \end{array} \right) \middle| \mathbf{x}_i^{(11)}, y_i^{(1)}, \boldsymbol{\Theta}^{(t)} \right] \\ = \left(\begin{array}{cc} (\tilde{\mathbf{x}}_i^{(21)})^T \tilde{\mathbf{x}}_i^{(21)} & (\tilde{\mathbf{x}}_i^{(31)})^T \tilde{\mathbf{x}}_i^{(21)} \\ (\tilde{\mathbf{x}}_i^{(21)})^T \tilde{\mathbf{x}}_i^{(31)} & (\tilde{\mathbf{x}}_i^{(31)})^T \tilde{\mathbf{x}}_i^{(31)} \end{array} \right) + \left(\begin{array}{cc} \boldsymbol{\Sigma}_{22|y}^{(t)} & \boldsymbol{\Sigma}_{23|y}^{(t)} \\ \boldsymbol{\Sigma}_{32|y}^{(t)} & \boldsymbol{\Sigma}_{33|y}^{(t)} \end{array} \right), \end{aligned} \quad (9)$$

$$E \left[\left(\mathbf{x}_i^{(24)} \right)^T \mathbf{x}_i^{(24)} \mathbf{x}_i^{(14)}, \mathbf{x}_i^{(34)}, y_i^{(4)}, \boldsymbol{\Theta}^{(t)} \right] = (\tilde{\mathbf{x}}_i^{(24)})^T \tilde{\mathbf{x}}_i^{(24)} + \boldsymbol{\Sigma}_{22|3y}^{(t)}, \quad (10)$$

$$E \left[\left(\mathbf{x}_i^{(32)} \right)^T \mathbf{x}_i^{(32)} \mathbf{x}_i^{(12)}, \mathbf{x}_i^{(22)}, y_i^{(2)}, \boldsymbol{\Theta}^{(t)} \right] = (\tilde{\mathbf{x}}_i^{(32)})^T \tilde{\mathbf{x}}_i^{(32)} + \boldsymbol{\Sigma}_{33|2y}^{(t)}, \quad (11)$$

where

$$\begin{aligned} \left(\begin{array}{cc} \boldsymbol{\Sigma}_{22|y}^{(t)} & \boldsymbol{\Sigma}_{23|y}^{(t)} \\ \boldsymbol{\Sigma}_{32|y}^{(t)} & \boldsymbol{\Sigma}_{33|y}^{(t)} \end{array} \right) &= \left(\begin{array}{cc} \boldsymbol{\Sigma}_{22}^{(t)} & \boldsymbol{\Sigma}_{23}^{(t)} \\ \boldsymbol{\Sigma}_{32}^{(t)} & \boldsymbol{\Sigma}_{33}^{(t)} \end{array} \right) - \left(\begin{array}{c} \boldsymbol{\Sigma}_{2y}^{(t)} \\ \boldsymbol{\Sigma}_{3y}^{(t)} \end{array} \right) \left(\sigma_y^{(t)2} \right)^{-1} \left(\boldsymbol{\Sigma}_{y2}^{(t)} \quad \boldsymbol{\Sigma}_{y3}^{(t)} \right), \\ \boldsymbol{\Sigma}_{22|3y}^{(t)} &= \boldsymbol{\Sigma}_{22}^{(t)} - \left(\boldsymbol{\Sigma}_{2y}^{(t)}, \boldsymbol{\Sigma}_{23}^{(t)} \right) \left(\begin{array}{cc} \sigma_y^{(t)2} & \boldsymbol{\Sigma}_{y3}^{(t)} \\ \boldsymbol{\Sigma}_{3y}^{(t)} & \boldsymbol{\Sigma}_{33}^{(t)} \end{array} \right)^{-1} \left(\begin{array}{c} \boldsymbol{\Sigma}_{y2}^{(t)} \\ \boldsymbol{\Sigma}_{y3}^{(t)} \end{array} \right), \\ \boldsymbol{\Sigma}_{33|2y}^{(t)} &= \boldsymbol{\Sigma}_{33}^{(t)} - \left(\boldsymbol{\Sigma}_{3y}^{(t)}, \boldsymbol{\Sigma}_{32}^{(t)} \right) \left(\begin{array}{cc} \sigma_y^{(t)2} & \boldsymbol{\Sigma}_{y2}^{(t)} \\ \boldsymbol{\Sigma}_{2y}^{(t)} & \boldsymbol{\Sigma}_{22}^{(t)} \end{array} \right)^{-1} \left(\begin{array}{c} \boldsymbol{\Sigma}_{y3}^{(t)} \\ \boldsymbol{\Sigma}_{y2}^{(t)} \end{array} \right). \end{aligned}$$

Next, we plug the derived expectations in Equations (6)–(11) into the expected complete-data log-likelihood function in Equation (4). Through some algebra calculations, Equation (4) can be written as

$$\begin{aligned} E_{\mathbf{D}^{mis} | \mathbf{D}^{obs}, \boldsymbol{\Theta}^{(t)}} (l(\boldsymbol{\Theta}; \mathbf{D}^{obs}, \mathbf{D}^{mis})) \\ = -n \log |\boldsymbol{\Sigma}| - \sum_{l=1}^4 \sum_{i=1}^{n_l} \left(y_i^{(l)} - \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 - \beta_0, \tilde{\mathbf{x}}_i^{(2l)} \right. \\ \left. - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_2, \tilde{\mathbf{x}}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_3 \right) \\ \times \boldsymbol{\Sigma}^{-1} \left(y_i^{(l)} - \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 - \beta_0, \tilde{\mathbf{x}}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_2, \tilde{\mathbf{x}}_i^{(3l)} \right. \\ \left. - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_3 \right)^T - \text{tr} \left(\boldsymbol{\Sigma}^{-1} \left(\begin{array}{ccc} 0 & 0 & 0 \\ 0 & \boldsymbol{\Sigma}_{22|3y}^{(t)} & 0 \\ 0 & 0 & 0 \end{array} \right) \right) \\ + n_2 \left(\begin{array}{ccc} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \boldsymbol{\Sigma}_{33|2y}^{(t)} \end{array} \right) + n_1 \left(\begin{array}{ccc} 0 & 0 & 0 \\ 0 & \boldsymbol{\Sigma}_{22|y}^{(t)} & \boldsymbol{\Sigma}_{23|y}^{(t)} \\ 0 & \boldsymbol{\Sigma}_{32|y}^{(t)} & \boldsymbol{\Sigma}_{33|y}^{(t)} \end{array} \right) \Bigg). \end{aligned} \quad (12)$$

In Equation (12), a notational trick is used, i.e., $\tilde{\mathbf{x}}_i^{(kl)}$ is used to represent the data $\mathbf{x}_i^{(kl)}$ no matter if the data is observed or missing, $k = 2, 3; l = 1, \dots, 4$. When the data is observed, e.g., $\mathbf{x}_i^{(22)}$, we make $\tilde{\mathbf{x}}_i^{(22)} = \mathbf{x}_i^{(22)}$. When the data is missing, e.g., $\mathbf{x}_i^{(21)}$, the expression of $\tilde{\mathbf{x}}_i^{(21)}$ is given in Equation (6). This notational trick facilitates the maximization in the M-step, which will become apparent in the following discussion.

M-step:

Split the parameter set $\boldsymbol{\Theta}$ into two subsets: $(\boldsymbol{\beta}_1, \beta_0, \mathbf{A}_2, \mathbf{b}_2, \mathbf{A}_3, \mathbf{b}_3)$ and $\boldsymbol{\Sigma}$. The maximization problem can be solved by taking the partial derivative of the expectation in Equation (12) with respect to each subset and equating the partial derivatives to zero:

$$\frac{\partial E_{\mathbf{D}^{mis} | \mathbf{D}^{obs}, \boldsymbol{\Theta}^{(t)}} (l(\boldsymbol{\Theta}; \mathbf{D}^{obs}, \mathbf{D}^{mis}))}{\partial (\boldsymbol{\beta}_1, \beta_0, \mathbf{A}_2, \mathbf{b}_2, \mathbf{A}_3, \mathbf{b}_3)} = 0, \text{ and } \frac{\partial E_{\mathbf{D}^{mis} | \mathbf{D}^{obs}, \boldsymbol{\Theta}^{(t)}} (l(\boldsymbol{\Theta}; \mathbf{D}^{obs}, \mathbf{D}^{mis}))}{\partial \boldsymbol{\Sigma}} = 0.$$

Instead of directly solving these equations, which is computationally involved, we take an indirect approach by first obtaining the least square estimators for the coefficients in the following three regressions:

$$\begin{cases} y_i^{(l)} \sim \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 + \beta_0 \\ \tilde{\mathbf{x}}_i^{(2l)} \sim \mathbf{x}_i^{(1l)} \mathbf{A}_2 + \mathbf{b}_2 \\ \tilde{\mathbf{x}}_i^{(3l)} \sim \mathbf{x}_i^{(1l)} \mathbf{A}_3 + \mathbf{b}_3 \end{cases}$$

The least square estimators are:

$$\begin{cases} \begin{pmatrix} \beta_0^{(t+1)} \\ \boldsymbol{\beta}_1^{(t+1)} \end{pmatrix} = \left(\sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T (\mathbf{1}, \mathbf{x}_i^{(1l)}) \right)^{-1} \sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T y_i^{(l)} \\ \begin{pmatrix} \mathbf{b}_2^{(t+1)} \\ \mathbf{A}_2^{(t+1)} \end{pmatrix} = \left(\sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T (\mathbf{1}, \mathbf{x}_i^{(1l)}) \right)^{-1} \sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T \tilde{\mathbf{x}}_i^{(2l)} \\ \begin{pmatrix} \mathbf{b}_3^{(t+1)} \\ \mathbf{A}_3^{(t+1)} \end{pmatrix} = \left(\sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T (\mathbf{1}, \mathbf{x}_i^{(1l)}) \right)^{-1} \sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{1}, \mathbf{x}_i^{(1l)})^T \tilde{\mathbf{x}}_i^{(3l)} \end{cases} \quad (13)$$

It is not hard to show that these estimators are equivalent to the optimal solutions for $(\boldsymbol{\beta}_1, \beta_0, \mathbf{A}_2, \mathbf{b}_2, \mathbf{A}_3, \mathbf{b}_3)$ in the M-step. Furthermore, let $\mathbf{z}_i^{(l)} = (y_i^{(l)} - \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1^{(t+1)} - \beta_0^{(t+1)}, \tilde{\mathbf{x}}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2^{(t+1)} - \mathbf{b}_2^{(t+1)}, \tilde{\mathbf{x}}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3^{(t+1)} - \mathbf{b}_3^{(t+1)})$. Then, we can obtain the optimal solution for $\boldsymbol{\Sigma}$ as

$$\boldsymbol{\Sigma}^{(t+1)} = \frac{1}{n} \left\{ \sum_{l=1}^4 \sum_{i=1}^{n_l} (\mathbf{z}_i^{(l)})^T \mathbf{z}_i^{(l)} + n_4 \begin{pmatrix} 0 & 0 & 0 \\ 0 & \boldsymbol{\Sigma}_{22|3y}^{(t)} & 0 \\ 0 & 0 & 0 \end{pmatrix} + n_2 \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \boldsymbol{\Sigma}_{33|2y}^{(t)} \end{pmatrix} + n_1 \begin{pmatrix} 0 & 0 & 0 \\ 0 & \boldsymbol{\Sigma}_{22|y}^{(t)} & \boldsymbol{\Sigma}_{23|y}^{(t)} \\ 0 & \boldsymbol{\Sigma}_{32|y}^{(t)} & \boldsymbol{\Sigma}_{33|y}^{(t)} \end{pmatrix} \right\}. \quad (14)$$

3.1.2. Prediction

Upon convergence of the above EM iterations, we can obtain the estimated parameters $\hat{\boldsymbol{\Theta}} = (\hat{\boldsymbol{\Sigma}}, \hat{\boldsymbol{\beta}}_1, \hat{\beta}_0, \hat{\mathbf{A}}_2, \hat{\mathbf{b}}_2, \hat{\mathbf{A}}_3, \hat{\mathbf{b}}_3)$. Using these estimated parameters, we can further obtain the estimated joint distribution of $y_i^{(l)}, \mathbf{x}_i^{(2l)}, \mathbf{x}_i^{(3l)} | \mathbf{x}_i^{(1l)}$ in Equation (1). Then, we can derive any conditional distribution allowing for prediction of new samples in each sub-cohort through conditional probability rules and probability marginalization. Specifically, let i^* denote a new sample. The formulas for predicting the response variable \hat{y}_{i^*} when the sample belongs to different sub-cohorts are as follows:

$$\begin{aligned} \hat{y}_{i^*} &= \mathbf{x}_{i^*}^{(11)} \hat{\boldsymbol{\beta}}_1 + \hat{\beta}_0, \text{ if } i^* \in \text{sub-cohort 1;} \\ \hat{y}_{i^*} &= \mathbf{x}_{i^*}^{(12)} (\hat{\boldsymbol{\beta}}_1 - \hat{\mathbf{A}}_2 \hat{\boldsymbol{\Sigma}}_{22}^{-1} \hat{\boldsymbol{\Sigma}}_{2y}) + \mathbf{x}_{i^*}^{(22)} \hat{\boldsymbol{\Sigma}}_{22}^{-1} \hat{\boldsymbol{\Sigma}}_{2y} + (\hat{\beta}_0 - \hat{\mathbf{b}}_2 \hat{\boldsymbol{\Sigma}}_{22}^{-1} \hat{\boldsymbol{\Sigma}}_{2y}), \text{ if } i^* \in \text{sub-cohort 2;} \end{aligned}$$

$$\begin{aligned} \hat{y}_{i^*} &= \mathbf{x}_{i^*}^{(13)} \left(\hat{\boldsymbol{\beta}}_1 - (\hat{\mathbf{A}}_2, \hat{\mathbf{A}}_3) \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{22} & \hat{\boldsymbol{\Sigma}}_{23} \\ \hat{\boldsymbol{\Sigma}}_{32} & \hat{\boldsymbol{\Sigma}}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{2y} \\ \hat{\boldsymbol{\Sigma}}_{3y} \end{pmatrix} \right) \\ &+ (\mathbf{x}_{i^*}^{(23)}, \mathbf{x}_{i^*}^{(33)}) \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{22} & \hat{\boldsymbol{\Sigma}}_{23} \\ \hat{\boldsymbol{\Sigma}}_{32} & \hat{\boldsymbol{\Sigma}}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{2y} \\ \hat{\boldsymbol{\Sigma}}_{3y} \end{pmatrix} \\ &+ \left(\hat{\beta}_0 - (\hat{\mathbf{b}}_2, \hat{\mathbf{b}}_3) \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{22} & \hat{\boldsymbol{\Sigma}}_{23} \\ \hat{\boldsymbol{\Sigma}}_{32} & \hat{\boldsymbol{\Sigma}}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\boldsymbol{\Sigma}}_{2y} \\ \hat{\boldsymbol{\Sigma}}_{3y} \end{pmatrix} \right), \end{aligned}$$

$$\text{if } i^* \in \text{sub-cohort 3;} \\ \hat{y}_{i^*} = \mathbf{x}_{i^*}^{(14)} (\hat{\boldsymbol{\beta}}_1 - \hat{\mathbf{A}}_3 \hat{\boldsymbol{\Sigma}}_{33}^{-1} \hat{\boldsymbol{\Sigma}}_{3y}) + \mathbf{x}_{i^*}^{(34)} \hat{\boldsymbol{\Sigma}}_{33}^{-1} \hat{\boldsymbol{\Sigma}}_{3y} + (\hat{\beta}_0 - \hat{\mathbf{b}}_3 \hat{\boldsymbol{\Sigma}}_{33}^{-1} \hat{\boldsymbol{\Sigma}}_{3y}), \text{ if } i^* \in \text{sub-cohort 4.}$$

3.2. IMTL classification model

In the classification model, $y_i^{(l)}$ can take a value of either zero or one, depending on the class it represents. Within each class, consider the joint distribution $\mathbf{x}_i^{(2l)}$ and $\mathbf{x}_i^{(3l)}$ given $\mathbf{x}_i^{(1l)}$ to be multivariate normal:

$$(\mathbf{x}_i^{(2l)}, \mathbf{x}_i^{(3l)}) | \mathbf{x}_i^{(1l)}, y_i^{(l)} = 1 \sim \text{MVN}(\boldsymbol{\mu}_1(\mathbf{x}_i^{(1l)}), \boldsymbol{\Sigma}), \quad (15)$$

$$(\mathbf{x}_i^{(2l)}, \mathbf{x}_i^{(3l)}) | \mathbf{x}_i^{(1l)}, y_i^{(l)} = 0 \sim \text{MVN}(\boldsymbol{\mu}_0(\mathbf{x}_i^{(1l)}), \boldsymbol{\Sigma}), \quad (16)$$

where the class-specific means are linear functions of $\mathbf{x}_i^{(1l)}$:

$$\begin{aligned} \boldsymbol{\mu}_1(\mathbf{x}_i^{(1l)}) &= (\mathbf{x}_i^{(1l)} \mathbf{A}_2 + \mathbf{b}_{21}, \mathbf{x}_i^{(1l)} \mathbf{A}_3 + \mathbf{b}_{31}), \\ \text{and } \boldsymbol{\mu}_0(\mathbf{x}_i^{(1l)}) &= (\mathbf{x}_i^{(1l)} \mathbf{A}_2 + \mathbf{b}_{20}, \mathbf{x}_i^{(1l)} \mathbf{A}_3 + \mathbf{b}_{30}). \end{aligned}$$

The same covariance matrix,

$$\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{22} & \boldsymbol{\Sigma}_{23} \\ \boldsymbol{\Sigma}_{32} & \boldsymbol{\Sigma}_{33} \end{pmatrix},$$

is assumed for the two classes. Furthermore, we assume the distribution of $y_i^{(l)}$ given $\mathbf{x}_i^{(1l)}$ to be Bernoulli, that is

$$y_i^{(l)} = 1 | \mathbf{x}_i^{(1l)} \sim \text{Bernoulli} \left(\frac{1}{1 + \exp\{-\mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 - \beta_0\}} \right). \quad (17)$$

Let $\tilde{\boldsymbol{\Theta}} = (\tilde{\boldsymbol{\Sigma}}, \tilde{\boldsymbol{\beta}}_1, \tilde{\beta}_0, \mathbf{A}_2, \mathbf{b}_{21}, \mathbf{b}_{20}, \mathbf{A}_3, \mathbf{b}_{31}, \mathbf{b}_{30})$ contain the unknown parameters for the model in (15)-(17). We can write down the complete-data log-likelihood function:

$$\begin{aligned} l(\tilde{\boldsymbol{\Theta}}; \mathbf{D}^{obs}, \mathbf{D}^{mis}) &= -n \log |\boldsymbol{\Sigma}| + \sum_{l=1}^4 \sum_{i=1}^{n_l} (y_i^{(l)} (\mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_{21}, \mathbf{x}_i^{(3l)} \\ &- \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_{31}) \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 - \mathbf{b}_{21}, \mathbf{x}_i^{(3l)} \\ &- \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_{31})^T + (1 - y_i^{(l)}) (\mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 \\ &- \mathbf{b}_{20}, \mathbf{x}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_{30}) \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i^{(2l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_2 \\ &- \mathbf{b}_{20}, \mathbf{x}_i^{(3l)} - \mathbf{x}_i^{(1l)} \mathbf{A}_3 - \mathbf{b}_{30})^T) + y_i^{(l)} (\mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 + \beta_0) \\ &- \log(1 + \exp(\mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 + \beta_0)). \end{aligned} \quad (18)$$

Equation (18) can be decomposed into a logistic regression and a conditional multivariate normal distribution. As a result, we can estimate $(\boldsymbol{\beta}_1, \beta_0)$ and the remaining parameters in $\tilde{\boldsymbol{\Theta}}$ separately. Specifically, $(\boldsymbol{\beta}_1, \beta_0)$ are coefficients of the logistic regression model:

$$\text{logit}(P(y_i^{(l)} = 1)) = \mathbf{x}_i^{(1l)} \boldsymbol{\beta}_1 + \beta_0.$$

This model does not involve missing data, which means that $(\boldsymbol{\beta}_1, \beta_0)$ can be estimated by iteratively reweighted least squares estimation.

Furthermore, let $\boldsymbol{\Theta}$ be the parameters in $\tilde{\boldsymbol{\Theta}}$ excluding $(\boldsymbol{\beta}_1, \beta_0)$. $\boldsymbol{\Theta}$ can be estimated by a similar EM algorithm to the predictive model in Section 3.1. Please see Online Supplement for the formula in the EM algorithm and in the classification models on new samples.

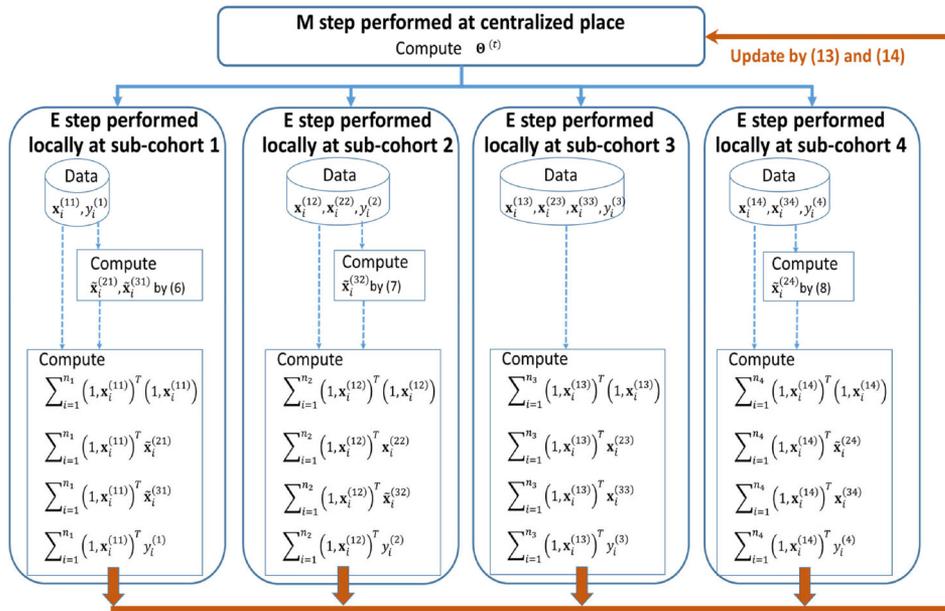


Figure 2. A computational framework for collaborative model estimation of IMTL without data pooling from different sub-cohorts.

3.3. Collaborative model estimation without data pooling

One reason leading to the generation of IMD data in health care applications is that each sub-cohort corresponds to a different institution. The availability of modalities varies across the different institutions due to accessibility and cost. In the IMTL models proposed in Sections 3.1 and 3.2, model estimation is assumed to happen at a centralized place into which the data from different institutions (i.e., sub-cohorts) have been deposited. This requires a multi-institutional data sharing agreement – a process known to be both time- and effort-intensive. A more commonly encountered scenario is that different institutions like to collaborate on model estimation without having to share their respective patients' data. In this section, we address the latter scenario by proposing a computational framework for model estimation. This framework uses the same equations as those in Sections 3.1 and 3.2, but it computes the equations at the E-step locally at each institution, as each equation only involves the data from a single institution. This is shown as the four vertical rectangular boxes in Figure 2. For example, Equation (6) in the E-step computes the expectations of two missing modalities in sub-cohort 1, $\mathbf{x}_i^{(21)}$ and $\mathbf{x}_i^{(31)}$, which only involve the data from sub-cohort 1, $\mathbf{x}_i^{(11)}$ and $y_i^{(1)}$. This nice “local” property also holds for other equations in the E-step. At the M-step, the proposed framework combines the locally computed results in a centralized place, which is shown as the horizontal rectangular box at the top of Figure 2. As these results do not reveal the raw data in each institution, patient privacy is preserved within each institution. The key idea of this computational framework is to consider the M- and E-steps as a global and a local learner, respectively. The global learner resides in a centralized place, whereas the local learners reside in each sub-cohort. A local learner can only “see” the data within the respective sub-cohort and performs calculations locally.

Results from the local computation, not the data, are sent to the global learner for combination.

Compared with the centralized model estimation in Sections 3.1 and 3.2, this computational framework involves communications among the global and local learners. As a result, there may be loss of efficiency due to limited communication bandwidth. On the other hand, this problem can be potentially mitigated because the computations of local learners can be performed in parallel.

3.4. Generalization to M modalities

The presentation of IMTL in Sections 3.1 to 3.3 is within the context of three modalities based on the consideration of notational simplicity. In this section, we provide the steps of extending IMTL to the general case of M modalities:

1. Given a multimodality dataset from an application, subjects (a.k.a. samples) are grouped into sub-cohorts with each sub-cohort having a different pattern of missing modalities.
2. Depending on the type of the response variable, one can decide if the problem to be tackled should be formulated as regression or classification. For a regression problem, a multivariate normal distribution can be assumed for the modalities and the response. For classification problems, a multivariate normal distribution can be assumed for the modalities and a Bernoulli distribution can be assumed for the response. For most applications, there is at least one modality available across all the sub-cohorts. If this is the case, the aforementioned distributions can be modified into conditional distributions given the available modality. Next, one can write down the complete-data log-likelihood function under the distribution assumption.

3. In the E-step of the EM algorithm, the key is to identify the sufficient statistics of the log-likelihood function, which include the missing modalities in each sub-cohort, the quadratic term of each missing modality, and pair-wise products between the missing modalities in that sub-cohort (if there is more than one missing modality). Then, one can derive expectations of the sufficient statistics given the observed data and parameter estimates from the previous iteration. Further, these expected sufficient statistics are plugged into the expected complete-data log-likelihood function, which will be maximized in the M-step.
4. In the M-step, an intuitive, but mathematically-involved, approach is to equate the first-order partial derivative of each parameter to zero and solve the parameter-wise simultaneous equations to update the parameter set. Alternative approaches can be developed to solve the maximization problem easier, depending on the form of the log-likelihood function. For example, in the three-modality case, we use a notational trick, which allowed us to convert the maximization into least square estimations.

Computational complexity: The proposed EM algorithm for IMTL estimation has analytical solutions in the E-step and M-step. Therefore, the computational complexity primarily comes from the iterations between the two steps until convergence. The complexity of the EM iterations has been well-studied in the literature. Furthermore, within the E-step, since there are no iterations but just arithmetic operations based on derived mathematical formula, the complexity primarily depends on how many expectations to compute. Given M modalities, the total number of sub-cohorts with missing modalities is $L = 2^{M-1} - 1$. Within each sub-cohort, there are two types of expectations to compute, including the mean vector and variance-covariance matrix of the missing features. Therefore, the complexity of the E-step can be considered as $2(2^{M-1} - 1)$. Within the M-step, the complexity primarily depends on how many parameters to estimate. Suppose each modality has p features. The total number of parameters is

$$\frac{1}{2}(M-1)^2 p^2 + (M-1)p^2 + \frac{5}{2}(M-1)p + p + 2.$$

4. Properties of IMTL

In this section, we discuss two unique properties of IMTL: (i) the ability for out-of-sample prediction; and (iii) a theoretical guarantee for a larger Fisher information compared with models without transfer learning, which explains the superiority of IMTL from a theoretical point of view.

4.1. Ability for out-of-sample prediction

Definition: Let D_{tr} denote a training set. Suppose the training samples can be divided into L sub-cohorts, where each sub-cohort corresponds to a different missing modality

pattern in $\{\mathcal{P}_1, \dots, \mathcal{P}_L\}$. Let i^* be a sample in a test set, whose missing modality pattern is $\mathcal{P}(i^*)$. Assume $\mathcal{P}(i^*) \notin \{\mathcal{P}_1, \dots, \mathcal{P}_L\}$. If a model trained on D_{tr} can be used predict i^* , the model is called capable of *out-of-sample prediction*.

For example, a training set can include only sub-cohorts 1, 2, and 4 in [Figure 1](#), whereas the test set includes sub-cohort 3. It is obvious that the two competing methods, i.e., SM and AADM, cannot do out-of-sample prediction. In contrast, IMTL is capable of out-of-sample prediction. Next, we provide an illustrative proof for this capability of IMTL. We focus on the predictive model in [Section 3.1](#). Also, for notational simplicity, each modality is assumed to contain one feature.

Consider a sample i^* in the test set which belongs to sub-cohort 3. To predict the response variable of this sample, [Equation \(19\)](#) will be used:

$$\begin{aligned} \hat{y}_{i^*}^{(3)} = & x_{i^*}^{(13)} \left(\hat{\beta}_1 - (\hat{A}_2, \hat{A}_3) \begin{pmatrix} \hat{\Sigma}_{22} & \hat{\Sigma}_{23} \\ \hat{\Sigma}_{32} & \hat{\Sigma}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\Sigma}_{2y} \\ \hat{\Sigma}_{3y} \end{pmatrix} \right) \\ & + \left(x_{i^*}^{(23)}, x_{i^*}^{(33)} \right) \begin{pmatrix} \hat{\Sigma}_{22} & \hat{\Sigma}_{23} \\ \hat{\Sigma}_{32} & \hat{\Sigma}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\Sigma}_{2y} \\ \hat{\Sigma}_{3y} \end{pmatrix} \\ & + \left(\hat{\beta}_0 - (\hat{b}_2, \hat{b}_3) \begin{pmatrix} \hat{\Sigma}_{22} & \hat{\Sigma}_{23} \\ \hat{\Sigma}_{32} & \hat{\Sigma}_{33} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\Sigma}_{2y} \\ \hat{\Sigma}_{3y} \end{pmatrix} \right). \end{aligned} \quad (19)$$

The parameters of the model in [Equation \(19\)](#) are estimated from a training set which only includes sub-cohorts 1, 2, and 4 but not 3. It is easy to understand why estimation of other parameters is possible except Σ_{23} . Intuitively, since Σ_{23} is the covariance between features in modalities 2 and 3, one would expect to have at least *some* training data from sub-cohort 3, which have both modalities 2 and 3 available, in order to estimate Σ_{23} . However, this is not our case. The key to demonstrating that IMTL can do out-of-sample prediction is to demonstrate that the estimation for Σ_{23} is possible by IMTL even without any data from sub-cohort 3 in the training set. To show this, consider the estimation for Σ_{23} by the EM algorithm. At convergence, it can be derived that Σ_{23} can be estimated by [Equation \(20\)](#). The detailed derivation is skipped due to space limitations:

$$\begin{aligned} \hat{\Sigma}_{23} = & \frac{1}{n - n_1 - \tilde{n}_2 - \tilde{n}_4} (\tilde{n}_1 - n_1) \hat{\Sigma}_{2y} (\hat{\sigma}_y^2)^{-1} \hat{\Sigma}_{y3} \\ & + \frac{1}{k_2} \sum_{i=1}^{n_2} (x_i^{(22)} - x_i^{(12)} \hat{A}_2 - \hat{b}_2) \\ & \times (y_i^{(2)} - x_i^{(12)} (\hat{\beta}_1 - \hat{A}_2 \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{2y}) - x_i^{(22)} \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{2y} \\ & - (\hat{\beta}_0 - \hat{b}_2 \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{2y})) \hat{\Sigma}_{y3} \\ & + \frac{1}{k_3} \sum_{i=1}^{n_4} (x_i^{(34)} - x_i^{(14)} \hat{A}_3 - \hat{b}_3) \\ & \times (y_i^{(4)} - x_i^{(14)} (\hat{\beta}_1 - \hat{A}_3 \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y}) \\ & - x_i^{(34)} \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y} - (\hat{\beta}_0 - \hat{b}_3 \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y})) \hat{\Sigma}_{y2}, \end{aligned} \quad (20)$$

where

$$\begin{aligned}
k_2 &= \hat{\sigma}_y^2 - \hat{\Sigma}_{y2} \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{2y}, \quad k_3 = \hat{\sigma}_y^2 - \hat{\Sigma}_{y3} \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y}, \\
\tilde{n}_1 &= \sum_{i=1}^{n_1} \frac{\left(y_i^{(1)} - x_i^{(11)} \hat{\beta}_1 - \hat{\beta}_0 \right)^2}{\hat{\sigma}_y^2}, \\
\tilde{n}_2 &= \sum_{i=1}^{n_2} \left(x_i^{(22)} - x_i^{(12)} \hat{A}_2 - \hat{b}_2 \right)^2 \\
&\quad \times \left(\hat{\Sigma}_{22}^{-1} + \frac{1}{k_2} \hat{\Sigma}_{22}^{-1} \hat{\Sigma}_{2y} \hat{\Sigma}_{y2} \hat{\Sigma}_{22}^{-1} \right) \\
&\quad - \frac{1}{k_2} \sum_{i=1}^{n_2} \left(x_i^{(22)} - x_i^{(12)} \hat{A}_2 - \hat{b}_2 \right)^T \\
&\quad \times \left(y_i^{(2)} - x_i^{(12)} \hat{\beta}_1 - \hat{\beta}_0 \right) \hat{\Sigma}_{y2} \hat{\Sigma}_{22}^{-1}, \\
\tilde{n}_4 &= \sum_{i=1}^{n_4} \left(x_i^{(34)} - x_i^{(14)} \hat{A}_3 - \hat{b}_3 \right)^2 \\
&\quad \times \left(\hat{\Sigma}_{33}^{-1} + \frac{1}{k_3} \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y} \hat{\Sigma}_{y3} \hat{\Sigma}_{33}^{-1} \right) \\
&\quad - \frac{1}{k_3} \sum_{i=1}^{n_2} \hat{\Sigma}_{33}^{-1} \hat{\Sigma}_{3y} \left(y_i^{(4)} - x_i^{(14)} \hat{\beta}_1 - \hat{\beta}_0 \right) \\
&\quad \times \left(x_i^{(34)} - x_i^{(14)} \hat{A}_3 - \hat{b}_3 \right).
\end{aligned}$$

Equation (20) indicates that although training data from sub-cohort 3 are not available, Σ_{23} can be estimated *indirectly* through a summation of three parts: $(\tilde{n}_1 - n_1) \hat{\Sigma}_{2y} (\hat{\sigma}_y^2)^{-1} \hat{\Sigma}_{y3}$ contributes to estimating the covariance between modalities 2 and 3 through exploiting their respective covariances with y . The second part leverages the training data in sub-cohort 2, and contributes to estimating Σ_{23} by exploring the covariance between residual modality 2 and residual modality 3. Here, residual modality 2 is modality 2 after factoring out modality 1; residual modality 3 is the residual of the response variable regressing on modalities 1 and 2. Both residual modalities are computed using the training data in sub-cohort 2. The third part leverages the training data in sub-cohort 4, and contributes to estimating Σ_{23} by exploring the covariance between residual modality 2 and 3 that are computed on the training data in sub-cohort 4.

Furthermore, we can explain why estimation of Σ_{23} without sub-cohort 3 is possible from another angle. Using the Law of Total Covariance, Σ_{23} can be decomposed as $\Sigma_{23} = \Sigma_{23|y} + \Sigma_{2y} \Sigma_{y3} / \sigma_y^2$. We cannot estimate $\Sigma_{23|y}$ due to the lack of sub-cohort 3. However, Σ_{2y} , Σ_{y3} , and σ_y^2 in the second term on the right-hand side can be estimated using the available sub-cohorts. This means that the estimator for Σ_{23} will be biased. Although it would be ideal to have data for sub-cohort 3 to mitigate the bias in estimating Σ_{23} , our simulation experiments in Section 5.1 show that the biased estimator performs well in prediction. In statistical models, biased estimators are used quite often and show good performance in prediction tasks.

4.2. Fisher information performance

The next section shows empirical evidence that IMTL outperforms SM and AADM, i.e., models without transfer learning. Here we would like to explain the performance improvement

from a theoretical standpoint, particularly through comparing the Fisher information of parameter estimates from IMTL and SM/AADM. Fisher information characterizes the variance of a maximum likelihood estimate, and larger Fisher information means smaller variance. It is our goal to find out if IMTL has larger Fisher information for some parameter estimates than SM/AADM, indicating more robust parameter estimation.

For clarity of presentation, we focus on a two-modality case in deriving the Fisher information for IMTL, SM, and AADM. Modality 1 is available for all patients, but modality 2 is missed for some patients. This divides the patients into two sub-cohorts: sub-cohort 1 has modality 1 available, but misses modality 2; sub-cohort 2 has both modalities available. Following the notational convention in Section 3, let $y_i^{(l)}$, $x_i^{(1l)}$, $x_i^{(2l)}$ denote the response variable, feature in modality 1, and feature in modality 2 for patient i in sub-cohort l , $l = 1, 2$. Assume one feature in each modality for notational simplicity. Let D^{mis} and D^{obs} contain the missing and observed data, i.e.,

$$D^{mis} = \left\{ x_i^{(21)} \right\}_{i=1}^{n_1}$$

and

$$D^{obs} = \left\{ \left\{ x_i^{(11)}, y_i^{(1)} \right\}_{i=1}^{n_1}, \left\{ x_i^{(12)}, x_i^{(22)}, y_i^{(2)} \right\}_{i=1}^{n_2} \right\}.$$

To model this dataset, IMTL assumes a multivariate normal distribution of $(y_i^{(l)}, x_i^{(2l)})$ given $x_i^{(1l)}$:

$$(y_i^{(l)}, x_i^{(2l)}) | x_i^{(1l)} \sim MVN(\boldsymbol{\mu}(x_i^{(1l)}), \boldsymbol{\Sigma}), \quad l = 1, 2, \quad (21)$$

where

$$\boldsymbol{\mu}(x_i^{(1l)}) = \left(x_i^{(1l)} \beta_1 + \beta_0, x_i^{(1l)} A_2 + b_2 \right),$$

and

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_{yy} & \sigma_{y2} \\ \sigma_{2y} & \sigma_{22} \end{pmatrix}.$$

To estimate the parameters of this IMTL model, a similar EM algorithm to that proposed in Section 3.1 can be used. To make predictions on new samples, we can derive the distributions of $y_i^{(1)} | x_i^{(11)}$ and $y_i^{(2)} | x_i^{(12)}, x_i^{(22)}$ from Equation (21) and use them for predicting samples from sub-cohort 1 and 2, respectively.

If SM is used to model this dataset, there will be two separate models for the two sub-cohorts. Since sub-cohort 2 has no missing modality, the model for sub-cohort 2 has the same form as Equation (21) but with $l = 2$ only. Sub-cohort 1 is separately modeled by a conditional distribution of the response variable given modality 1:

$$(y_i^{(1)} | x_i^{(11)}) \sim N(x_i^{(11)} \alpha_1 + \alpha_0, \zeta_y^2). \quad (22)$$

AADM is similar to SM in the sense that two separate models are built for the two sub-cohorts. These models take the same forms as those in SM. However, in estimating the model parameters for sub-cohort 1, AADM uses all available data which includes the data of modality 1 from both sub-cohort 1 and 2, since modality 1 is available for both sub-cohorts. Recall that in SM, only the data from sub-cohort 1

Table 1. Out-of-sample prediction accuracy on the test set with different training sample sizes (between-modality correlation is kept as 0.6 in both settings).

Training size	PMSE: ave(std)	PC: ave(std)
300	1.174 (0.175)	0.945 (0.010)
150	1.469 (0.289)	0.931 (0.017)

is used. To estimate the parameters of SM/AADM, maximum likelihood estimation can be used, as no missing data is involved in the model formulation. To make predictions on new samples, we can use Equation (22) directly if the sample is from sub-cohort 1, and use $y_i^{(2)} \mid x_i^{(12)}, x_i^{(22)}$ if the sample is from sub-cohort 2.

It can be seen from the above descriptions that IMTL, SM, and AADM share the same model for sub-cohort 2, but they estimate the model parameters in different ways. Theorem 1 compares the Fisher information of the parameter estimates for sub-cohort 2 by the three methods, specifically focusing on the estimates for the elements in the precision matrix:

$$\Omega \triangleq \Sigma^{-1} = \begin{pmatrix} \Omega_{yy} & \Omega_{y2} \\ \Omega_{2y} & \Omega_{22} \end{pmatrix}.$$

Theorem 1: Let $I_{IMTL}(\Omega_{22})$, $I_{IMTL}(\Omega_{2y})$, $I_{IMTL}(\Omega_{yy})$ be the Fisher information of the estimates for Ω_{22} , Ω_{2y} , and Ω_{yy} under IMTL, respectively. Let $I_{SM}(\cdot)$ and $I_{AADM}(\cdot)$ be the Fisher information of the estimates for the same parameters under SM and AADM, respectively. Then,

$$I_{IMTL}(\Omega_{2y}) > I_{SM}(\Omega_{2y}) = I_{AADM}(\Omega_{2y}), \text{ and } I_{IMTL}(\Omega_{yy}) > I_{SM}(\Omega_{yy}) = I_{AADM}(\Omega_{yy}).$$

Furthermore,

$$I_{IMTL}(\Omega_{22}) > I_{SM}(\Omega_{22}) = I_{AADM}(\Omega_{22})$$

under the condition that

$$\frac{-n_1 + 2p_1 + \sqrt{(n_1 - 2p_1)^2 + 4n_1p_1}}{4p_1} < \frac{\sigma_{2y}^2}{\sigma_{22}\sigma_{yy}}, \quad (23)$$

where n_1 is the sample size of sub-cohort 1 and p_1 is number of features in modality 1.

Please see the proof in the Online Supplement. Theorem 1 shows that the Fisher information for the estimates of Ω_{2y} and Ω_{yy} under IMTL is greater than SM/AADM unconditionally. This relationship holds for the estimate of Ω_{22} under the condition given in Equation (22). This condition is worthy of further discussion. Specifically, if considering n_1 and p_1 to be fixed (i.e., the left-hand side of Equation (22) is a constant), Theorem 1 indicates that the correlation between modality 2 and the response variable must be sufficiently large (i.e., larger than the constant) in order for IMTL to have a larger Fisher information for the estimate of Ω_{22} than SM/AADM. This means that IMTL will be most effective when the modality with missing data is a significant predictor for the response. Otherwise, IMTL may not perform as well as models without transfer learning because it runs the risk of transferring noise and thus hurting the model performance.

Table 2. Out-of-sample prediction accuracy on the test set with different between-modality correlations (training sample size is kept as 300 in both settings).

Between-modality correlation	PMSE: ave(std)	PC: ave(std)
0.6	1.174 (0.175)	0.945 (0.010)
0	1.300 (0.187)	0.866 (0.028)

5. Application case study

In this section, we apply IMTL to simulated and real-world datasets. Simulation experiments are presented in Section 5.1, with the purpose of demonstrating the out-of-sample prediction ability of IMTL, which the competing methods (i.e., SM and AADM) do not possess. Section 5.2 presents an application of AD diagnosis and prognosis of MCI patients using the Alzheimer’s Disease Neuroimaging Initiative (ADNI) dataset. Here, “diagnosis” means detection of the existence of AD pathology in the brain of an MCI patient. “Prognosis” means prediction if an MCI patient will progress to AD by a certain year of interest, e.g., 6 years. Both tasks are important for treatment and management of the patients.

5.1. Simulation experiments

5.1.1. Out-of-sample prediction by IMTL predictive model

We conduct simulation experiments for the IMTL predictive model and classification model. For the predictive model, we first generate data for three modalities, i.e., $\mathbf{x}_i^{(1)}$, $\mathbf{x}_i^{(2)}$, $\mathbf{x}_i^{(3)}$, from a zero-mean multivariate normal distribution $MVN(\mathbf{0}, \Sigma)$. The number of features in each modality is set to be $p_1 = 10$, $p_2 = p_3 = 5$, which are close to the size of features in the real-world data presented in Section 5.1. All diagonal elements of Σ are set to be one. Σ includes two parts: within-modality correlation and between-modality correlation. The former has been found to have little impact on the model performance and therefore is set to be 0.6. We investigate two settings for between-modality correlation: 0.6 and 0, which represent moderately strong correlation and no correlation. We investigate two training sample sizes: 300 and 150.

Once the data for features are generated, we generate the response variable $y^{(l)}$ by a linear model,

$$y_i^{(l)} = \mathbf{x}_i^{(1)}\boldsymbol{\beta}_1 + \mathbf{x}_i^{(2)}\boldsymbol{\beta}_2 + \mathbf{x}_i^{(3)}\boldsymbol{\beta}_3 + \beta_0 + \epsilon.$$

Here, $\beta_0 = 2$; elements in $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$, $\boldsymbol{\beta}_3$ are set to be 0.2; $\epsilon \sim \mathcal{N}(0, 1)$. Then, the simulated training data are equally separated into three sub-cohorts, $l = 1, 2, 4$, corresponding to sub-cohorts 1, 2, and 4 in Figure 1. To obtain the incomplete modality pattern in each sub-cohort, we remove the training data of modalities 2 and 3 for sub-cohort 1, remove modality 3 for sub-cohort 2, and remove modality 2 for sub-cohort 4. As our intention in this experiment is to assess the out-of-sample prediction capability of IMTL, we generate data in a test data that includes only sub-cohort 3, i.e., all the modalities are available. The sample size of the test set is 100.

Table 3. Out-of-sample classification accuracy on the test set with different training sample sizes (between-modality correlation is kept as 0.6 in both settings).

Training size	AUC: ave(std)
300	0.882 (0.037)
150	0.832 (0.05)

IMTL is trained on the training set that includes only data from sub-cohorts 1, 2, and 4. Then, the trained model is used to predict on the test set that only includes samples from sub-cohort 3. The predicted response variables of the test set are compared with the true responses to compute a Prediction Mean Square Error (PMSE) and a Pearson Correlation (PC). We repeat the entire experiment for 100 times. Tables 1 and 2 summarize the results. As expected, increasing the training sample size significantly improves PMSE and PC ($p < 0.001$). The correlation between modalities also helps improve PMSE and PC ($p < 0.001$), consistent with the theoretical discovery in Section 4.1, in which we found that the key to out-of-sample prediction was to be able to estimate Σ_{23} from the training data. From Equation (20), it is known that the estimation of Σ_{23} is affected by the correlation between modality 2 and 3. Even though the training data does not include samples with both modality 2 and 3 available, Σ_{23} can still be estimated indirectly by IMTL through exploiting the between-modality correlation and the relationship between modalities and the response variable.

5.1.2. Out-of-sample prediction by IMTL classification model

The data generation process of this experiment is the same as the previous section, except that we use a logistic regression model to link the response variable with predictors/features. Specifically, we first simulate a linear predictor

$$z_i^{(l)} = \mathbf{x}_i^{(1)} \boldsymbol{\beta}_1 + \mathbf{x}_i^{(2)} \boldsymbol{\beta}_2 + \mathbf{x}_i^{(3)} \boldsymbol{\beta}_3 + \beta_0 + \epsilon.$$

Then, $y_i^{(l)}$ is generated from a Bernoulli distribution with success probability equal to $1/(1 + e^{-z_i^{(l)}})$. Test accuracy is reported as the Area Under the Curve (AUC). Tables 3 and 4 summarize the results. Doubling the training sample size does not seem to dramatically improve the AUC, although this improvement is still statistically significant ($p < 0.001$). The correlation between modalities also helps improve the AUC significantly ($p < 0.001$).

5.2. Early diagnosis and prognosis of AD

5.2.1. Introduction to ADNI

ADNI (<http://adni.loni.ucla.edu>) was launched in 2003 by the NIH, FDA, private pharmaceutical companies, and non-profit organizations, as a \$60 000 000, 5-year public-private partnership. The primary goal of ADNI has been to test whether MRI, PET, other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and early AD. Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians

Table 4. Out-of-sample classification accuracy on the test set with different between-modality correlations (training sample size is kept as 300 in both settings).

Between-modality correlation	AUC: ave(std)
0.6	0.882 (0.037)
0	0.781 (0.046)

to develop new treatments and monitor their effectiveness, as well as lessen the time and cost of clinical trials. The Principal Investigator of this initiative is Michael W. Weiner, MD, VA Medical Center and University of California-San Francisco. ADNI is the result of efforts of many co-investigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the US and Canada. For up-to-date information, see <http://www.adni-info.org/>.

5.2.2. Patient inclusion and diagnostic/prognostic end points

Our study includes 214 MCI patients from ADNI through our collaborative intuition, Banner Alzheimer's Institute (BAI), with which two co-authors are affiliated. BAI is a member of ADNI PET core (PI, William Jagust UC Berkeley). Multimodality image data include MRI, FDG-PET, amyloid-PET, which follow the IMD structure in Figure 1. Each sub-cohort has the same sample size. For *diagnosis*, we use $A\beta$ positivity as an indicator for high-risk AD. We follow the recommendation by Fleisher *et al.* (2011) and use a threshold of mean standard uptake value ratios (SUVR) greater than or equal to 1.18 to define $A\beta$ positivity. According to this criterion, there are 87 and 127 patients in class 1 (high-risk) and 0 (otherwise). For *prognosis*, the purpose is to predict when an MCI patient will convert to AD. We searched the ADNI database for the 214 patients from the time when their imaging data were collected up to 6 years' follow up, and found that 46 converted to AD, i.e., there are 46 converters (class 1) and 168 non-converters (class 0).

5.2.3. Image processing and feature computation

For MRI images, we use the FreeSurfer (<http://surfer.nmr.mgh.harvard.edu/>) software to extract volumetric measurements for pre-defined Regions Of Interest (ROIs). We focus on three ROIs including hippocampal, ventricle, and entorhinal volumes relative to intracranial volume. All three have been widely reported to be related to AD (Thompson *et al.*, 2004; Devanand *et al.*, 2007). Both FDG-PET and amyloid-PET are PET images, so they share the same image processing step in which we use SPM8 (<http://www.fil.ion.ucl.ac.uk/spm/>) to spatially normalize each patient's PET images into the common Montreal Neurological Institute (MNI) atlas space. Then, we extract features from each type of PET image separately. From FDG-PET, the features include Hypometabolic Convergence Index (HCI) (Chen *et al.*, 2011), statistical Region Of Interest (sROI) (Chen *et al.*, 2010), and regional precuneus metabolism and posterior cingulate metabolism. All these features have been previously reported to be related to AD (Del Sole *et al.*, 2008; Bailly *et al.* 2015). From amyloid-PET, the features include

Table 5. Diagnostic and prognostic performance: ave (std) and p value for hypothesis testing that IMTL is better than a competing method.

	Diagnosis			Prognosis		
	IMTL	SM	AADM	IMTL	SM	AADM
AUC	0.93(0.03)	0.86(0.06)	0.90(0.04)	0.85(0.05)	0.72(0.09)	0.78(0.07)
		$p < 0.001$	$p < 0.001$		$p < 0.001$	$p < 0.001$
Sensitivity	0.91(0.06)	0.84(0.09)	0.88(0.06)	0.96(0.09)	0.58(0.18)	0.76(0.17)
		$p < 0.001$	$p = 0.03$		$p < 0.001$	$p < 0.001$
Specificity	0.87(0.06)	0.82(0.06)	0.86(0.05)	0.85(0.05)	0.86(0.05)	0.86(0.05)
		$p < 0.001$	$p = 0.27$		$p = 0.78$	$p = 0.34$

Note: The numbers in bold indicate that IMTL is better than competing methods.

SUVRs from six brain regions including orbital frontal cingulate, temporal cortex, anterior cingulate, posterior cingulate, parietal lobe, and precuneus. These regions are known to be associated with amyloid depositions and AD (Fleisher *et al.*, 2011). As the six SUVRs are highly correlated, we apply Principal Component Analysis (PCA) and include the first Principal Component (PC) as a feature for amyloid-PET. Note that IMTL assumes normal distributions of the features. In this application, this assumption naturally holds because each feature is an average of voxel-wise measurements within a brain region. Since many voxels are involved in generating the average, the Central Limit Theorem applies. Also, we generate normal Q-Q plots for the features and find that the normality assumption holds.

5.2.4. Inclusion of clinical variables and feature screening

We also include the following clinical variables that could potentially help the early diagnosis and prognosis of AD: age, gender, years of education, APOE e4 status, and cognitive test scores from several commonly used instruments such as Mini-Mental State Examination (MMSE), AD assessment scale-cognitive, Clinical Dementia Rating (CDR), and auditory verbal learning test. No patient has missing data for these clinical variables so they are used in the same way as MRI features in our model. Furthermore, we put all the features through a feature screening module using the approach by Fan and Lv (2008). Note that feature screening is only applied to the training set not the entire dataset to avoid overfitting.

5.2.5. Application of IMTL

Within each sub-cohort, the samples are divided into five folds. We combine four folds from each sub-cohort into a training set and use the remaining data as the test set. We apply IMTL to the training set and then use the trained model to predict on the test set. We exhaust all four-fold combinations in training, which produces a 5-fold cross-validation procedure for evaluating the accuracy of IMTL. This process is repeated for 50 times. For comparison, two competing methods are applied on the same data: SA and AADM. Table 5 summarizes the results. IMTL has significantly higher AUC and sensitivity than both competing methods in both diagnosis and prognosis. Notably, competing methods have low AUC and sensitivity in prognosis. This is greatly improved by IMTL. Prognosis is more challenging than diagnosis, due to the former having a heavily imbalanced dataset (46 converters vs. 168 non-converters). Clearly, IMTL is more robust to sample imbalance. All models achieve

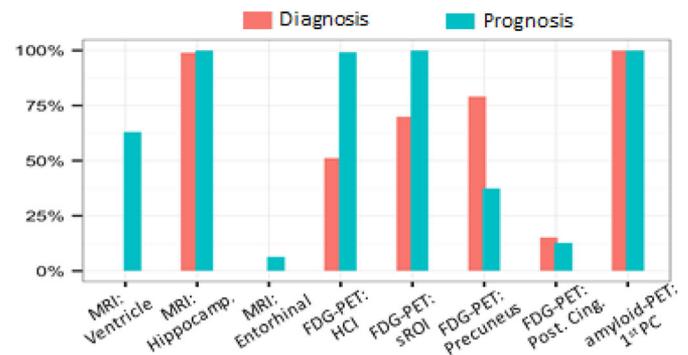


Figure 3. Percentage of times imaging features are included in IMTL over 5-fold cross-validation and 50 repeated experiments.

a similar level of specificity. Finally, we show the contribution of each imaging feature to diagnosis and prognosis by plotting the percentage of times a feature is included in the IMTL model. The result is shown in Figure 3. Hippocampal volume from MRI and the first PC of six SUVRs from amyloid-PET are almost always included in both diagnostic and prognostic models. This is consistent with findings in the literature that hippocampal atrophy and amyloid-PET SUVRs provide most important biomarkers for AD (Devanand *et al.*, 2007; Fleisher *et al.*, 2011). Other features that are selected for over 50% of the time include HCl, sROI and precuneus metabolism from FDG-PET for diagnosis; and ventricle volume from MRI and HCl and sROI from FDG-PET for prognosis. Clinical variables such as age, APOE e4 status, MMSE, and CDR are more frequently selected. These variables have been widely reported to be related to AD.

6. Conclusion

In this article, we proposed IMTL to build predictive and classification models for IMD data. We developed an EM algorithm for parameter estimation of IMTL and further extended it to achieve between-institutional collaborative model estimation without the need for data pooling. We demonstrated that IMTL was capable of out-of-sample prediction and proved that it had larger Fisher information than models without transfer learning under mild conditions. This explained the superiority of IMTL from a theoretical standpoint. Simulation experiments demonstrated high accuracy in using IMTL for out-of-sample prediction and classification. IMTL was applied to AD early diagnosis and prognosis, i.e., at the MCI stage, using incomplete multimodality imaging data. Significantly higher AUC and sensitivity were achieved in both diagnosis and prognosis compared with competing methods. Image features selected to include in the models were widely reported in the literature to be related to AD.

This research has several limitations: First, IMTL assumes normal distributions of features. To make IMTL an appropriate choice for an application, feature normality needs to be checked. If the features do not follow a normal distribution, Cox transformation may be used. Nevertheless, extension of IMTL to non-normal features provides a more general approach, and thus being an interesting future

research direction. Second, this article focuses on response variables that are either normal or binary. Extending the current modeling framework to other types of response variables will be valuable to address the need of various application domains. Third, IMTL assumes equal variance-covariance for the models in different sub-cohorts. This assumption can be relaxed to accommodate potential sub-cohort heterogeneity. Last but not least, although **Theorem 1** reveals one advantage of IMTL in terms of greater Fisher information (thus smaller variance asymptotically) for estimators of the elements in the precision matrix, more theoretical studies may be performed on the estimators of other parameters of IMTL to reveal their properties compared with competing methods.

Funding

This research is partly supported by NIH-NIA grant no. R41AG053149 and NSF CAREER 1149602. Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

Notes on contributors

Xiaonan Liu received his PhD in industrial engineering at Arizona State University in 2019. His research focuses on statistical modeling and machine learning with applications in health care and manufacturing.

Kewei Chen received his PhD in biomathematics from UCLA in 1993. He is a Sr. Scientist at Banner Alzheimer's Institute, an adjunct professor in the School of Math and Statistics, Arizona State University and a research professor in the College of Medicine-Phoenix, University of Arizona, USA. His current research interests are neuroimaging-based biomarkers for Alzheimer's disease studies, machine learning and statistics.

David Weidman, MD, is a neurologist and associate director of the trials division at Banner Alzheimer's Institute. His research interests include the utility of neuroimaging in Alzheimer's disease and allied disorders. He is a principal investigator of several interventional trials testing disease modifying therapies for Alzheimer's disease across its spectrum.

Teresa Wu is a professor in industrial engineering at Arizona State University. She received her PhD from the University of Iowa in 2007. Her research interests are health informatics and distributed decision supports. She is a recipient of an NSF CAREER award. She is a member of IISE and INFORMS.

Fleming Y. M. Lure is Chief Product Officer at MS Technologies Corp, Rockville, MD. He received his PhD in electrical engineering at Pennsylvania State University in 1990. His research interests are computer-assisted radiology and pathology with machine learning, image processing, artificial intelligence, and portable and touchless devices for monitoring and treating individuals.

Jing Li is a professor in industrial engineering at Arizona State University. She received her PhD in industrial and operations engineering from the University of Michigan in 2007. Her research interests are statistical modeling and machine learning for health care applications. She is a recipient of an NSF CAREER award. She is a member of IISE, INFORMS, and IEEE.

References

- Albert, M.S., Steven T.D., Dickson, D., Dubois, B., Feldman, H.H., Fox, N.C., Gamst, A., Holtzman, D.M., Jagust, W.J., Petersen, R.C., Snyder, P.J., Carrillo, M.C., Thine, B. and Phelps, C.H. (2011) The diagnosis of mild cognitive impairment due to Alzheimer's disease: Recommendations from the National Institute on Aging-Alzheimer's Association Workgroups on Diagnostic Guidelines for Alzheimer's Disease. *Alzheimer's and Dementia*, 7(3), 270–279.
- Bailly, M., Destrieux, C., Hommet, C., Mondon, K., Cottier, J.P., Beaufils, E., Vierron, E., Vercoullie, J., Ibazizene, M., Voisin, T., Payoux, P., Barre, L., Camus, V., Guilloteau, D. and Ribeiro M. (2015) Precuneus and cingulate cortex atrophy and hypometabolism in patients with Alzheimer's disease and mild cognitive impairment: MRI and 18F-FDG PET quantitative analysis using freesurfer. *BioMed Research International*.
- Chen, K., Ayutyanont, N., Langbaum, J.B.S., Fleisher, A.S., Reschke, C., Lee, W., Liu, X., Bandy, D., Alexander, P.M., Thompson, L.S., John, Q.T., Jack Jr, C.R., Landau, S.M., Foster, N.L., Harvey, D.J., Weiner, M.W., Koeppe, R.A., Jagust, W.J., Reiman, E.M. the Alzheimer's Disease Neuroimaging Initiative. (2011) Characterizing Alzheimer's disease using a hypometabolic convergence index. *NeuroImage*, 56(1), 52–60.
- Chen, K., Langbaum, J.B., Fleisher, A.S., Ayutyanont, N., Reschke, C., Lee, W., Liu, X., Bandy, D., Alexander, G.E., Thompson, P.M., Foster, N.L., Harvey, D.J., de Leon, M.J., Koeppe, R.A., Jagust, W.J., Weiner, M.W., Reiman, E.M. and the Alzheimer's Disease Neuroimaging Initiative. (2010) Twelve-month metabolic declines in probable Alzheimer's disease and amnesic mild cognitive impairment assessed using an empirically pre-defined statistical region-of-interest: Findings from the Alzheimer's Disease Neuroimaging Initiative. *NeuroImage*, 51(2), 654–664.
- Del Sole, A., Clerici, F., Chiti, A., Lecchi, M., Mariani, C., Maggiore, L., Mosconi, L. and Lucignani, G. (2008) Individual cerebral metabolic deficits in Alzheimer's disease and amnesic mild cognitive impairment: An FDG PET study. *European Journal of Nuclear Medicine and Molecular Imaging*, 35(7), 1357–1366.
- Devanand, D.P., Pradhaban, G., Liu, X., Khandji, A., De Santi, S., Segal, S., Rusinek, H., Pelton, G.H., Honig, L.S., Mayeux, R., Stern, Y., Tabert, M.H. and de Leon M.J. (2007) Hippocampal and entorhinal atrophy in mild cognitive impairment: Prediction of Alzheimer disease. *Neurology*, 68(11), 828–836.
- Fan, J. and Lv, J. (2008) Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 70(5), 849–911.
- Fleisher, A.S., Chen, K., Liu, X., Roontiva, A., Thiyyagura, P., Ayutyanont, N., Joshi, A.D., Clark, C.M., Mintun, M.A., Pontecorvo, M.J., Doraiswamy, P.M., Johnson, K.A., Skovronsky, D.M. and Reiman, E.M. (2011) Using positron emission tomography and

- florbetapir F 18 to image cortical amyloid in patients with mild cognitive impairment or dementia due to Alzheimer disease. *Archives of Neurology*, **68** (11), 1404–1411.
- He, D., Wang, Z., Yang, L. and Dai, W. (2017) Study on missing data imputation and modeling for the leaching process. *Chemical Engineering Research and Design*, **124**, 1–19.
- Li, R., Zhang, W., Suk, H., Wang, L., Li, J., Shen, D., and Ji, S. (2014) Deep learning based imaging data completion for improved brain disease diagnosis. *Medical Image Computing and Computer-Assisted Intervention*, **17**, 305–312.
- Little, R. and Rubin, D.B. (2002) *Statistical Analysis with Missing Data*, Second Edition, John Wiley & Sons, Hoboken, NJ.
- Liu, M., Zhang, J., Yap, P.T. and Shen, D. (2017) View-aligned hypergraph learning for Alzheimer’s disease diagnosis with incomplete multi-modality data. *Medical Image Analysis*, **36**, 123–134.
- Thompson, P.M., Hayashi, K.M., de Zubicaray, G.I., Janke, A.L., Rose, S.E., Semple, J., Hong, M.S., Herman, D.H., Gravano, D., Doddrell, D.M. and Toga, A.W. (2004) Mapping Hippocampal and ventricular change in Alzheimer disease. *NeuroImage*, **22**(4), 1754–1766.
- Thung, K.H., Wee, C.Y., Yap, P.T. and Shen, D. (2014) Neurodegenerative disease diagnosis using incomplete multi-modality data via matrix shrinkage and completion. *NeuroImage*, **91**, 386–400.
- Xiang, S., Yuan, L., Fan, W., Wang, Y., Thompson, P.M. and Ye, J. (2014) Bi-level multi-source learning for heterogeneous block-wise missing data. *NeuroImage*, **102**, 192–206.
- Yuan, L., Wang, Y., Thompson, P.M., Narayan, V.A. and Ye, J. (2012) Multi-source feature learning for joint analysis of incomplete multiple heterogeneous neuroimaging data. *NeuroImage*, **61**(3), 622–632.