



Bayesian Tensor Modeling for Image-based Classification of Alzheimer's Disease

Rongke Lyu¹ · Marina Vannucci¹ · Suprateek Kundu² · and for the Alzheimer's Disease Neuroimaging Initiative

Accepted: 20 May 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Tensor-based representations are being increasingly used to represent complex data types such as imaging data, due to their appealing properties such as dimension reduction and the preservation of spatial information. Recently, there is a growing literature on using Bayesian scalar-on-tensor regression techniques that use tensor-based representations for high-dimensional and spatially distributed covariates to predict continuous outcomes. However surprisingly, there is limited development on corresponding Bayesian classification methods relying on tensor-valued covariates. Standard approaches that vectorize the image are not desirable due to the loss of spatial structure, and alternate methods that use extracted features from the image in the predictive model may suffer from information loss. We propose a novel data augmentation-based Bayesian classification approach relying on tensor-valued covariates, with a focus on imaging predictors. We propose two data augmentation schemes, one resulting in a support vector machine (SVM) type of classifier, and another yielding a logistic regression classifier. While both types of classifiers have been proposed independently in literature, our contribution is to extend such existing methodology to accommodate high-dimensional tensor valued predictors that involve low rank decompositions of the coefficient matrix while preserving the spatial information in the image. An efficient Markov chain Monte Carlo (MCMC) algorithm is developed for implementing these methods. Simulation studies show significant improvements in classification accuracy and parameter estimation compared to routinely used classification methods. We further illustrate our method in a neuroimaging application using cortical thickness MRI data from Alzheimer's Disease Neuroimaging Initiative, with results displaying better classification accuracy throughout several classification tasks, including classification on pairs of the three diagnostic groups: normal control, AD patients, and MCI patients; gender classification (males vs females); and cognitive performance based on high and low levels of MMSE scores.

Keywords Alzheimer's disease · Bayesian tensor modeling · Logistic regression · Support vector machines · Neuroimaging analysis

Alzheimer's Disease Neuroimaging Initiative (ADNI) is a Group/Institutional Author.

Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.ucla.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.ucla.edu/wpcontent/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

✉ Rongke Lyu
rl77@rice.edu

¹ Department of Statistics, Rice University, Houston, TX, United States

² Department of Biostatistics, MD Anderson Cancer Center, Houston, TX, United States

Introduction

Neuroimaging studies stand as a cornerstone in contemporary neuroscience, fundamentally transforming our comprehension of the intricate structure and function of the brain. These non-invasive visualization techniques have not only enriched our understanding of neurological disorders but have also pioneered new frontiers in mental health research. Within the realm of risk prediction, neuroimaging studies have emerged as an invaluable tool for identifying individuals susceptible to neurological and psychiatric conditions. By discerning subtle abnormalities in brain structure and connectivity, researchers can now predict the risk of mental disorders (such as Alzheimer's disease or dementia) with greater precision. This early identification

facilitates timely intervention and treatment, potentially offering an opportunity to provide improved health outcomes and quality of life related to these disorders.

For example, neuroimaging studies have transformed the field of Alzheimer's disease (AD) research, providing invaluable insights into the pathological mechanisms underlying this devastating neurodegenerative disorder (Chouliaras & O'Brien, 2023). By visualizing the brain's structural and functional changes, neuroimaging techniques have enabled researchers to track the progression of AD, identify early signs of the disorder, and differentiate it from other causes of dementia. In particular, structural neuroimaging studies involving magnetic resonance imaging (MRI), have revealed the intricate patterns of atrophy in spatially distributed brain regions involved in memory and cognition that is a hallmark of AD (Frenzel et al., 2020). Neuroimaging features such as brain volume or cortical thickness that are derived from MRI scans can be used to provide quantitative assessments of disease severity and monitor disease progression over time. Importantly, such neuroimaging features can be embedded in machine learning algorithms in order to perform risk prediction or early detection in AD. However, several major challenges are encountered when analyzing neuroimaging data. For example, the brain imaging data is spatially dependent, high-dimensional and noisy, and it is often unclear how to identify suitable neurobiological markers for the mental disorder in the presence of heterogeneity.

In order to model such complex types of imaging data emerging at a rapid pace, several statistical and machine learning approaches have been proposed. Among them, classification models using neuroimaging features have seen a rapid development (Rathore et al., 2017; Arbabshirani et al., 2017; Falahati et al., 2014). These approaches typically either vectorize the image, or extract informative summary features from the image, to be used as covariates. For example, Plant et al. (2010) extracted the low-level-feature extraction algorithm with feature selection criterion to select most discriminating features, that are then coupled with a clustering algorithm to group spatially coherent voxels to predict Alzheimer's disease status. Ben Ahmed et al. (2015) proposed a multi-feature fusion algorithm used both extracted visual features from the hippocampal region of interest (ROI) and the quantity of cerebrospinal fluid (CSF) in the hippocampal region, and then applied a late fusion scheme to perform the binary classification of Alzheimer's disease subjects using the MRI images. Going beyond AD classification, Griffis et al. (2016) implemented a voxel-based Gaussian Naive Bayes Classification of ischemic stroke lesions in individual T1-weighted MRI scans, where the authors separately created two feature maps as predictor variables for missing and abnormal tissue to avoid including highly redundant information. Alternate types of shape-based image analysis that go beyond

voxel-level analysis have also been proposed for prediction (Wu et al., 2022).

The above approaches, while useful, did not explicitly account for the spatial configuration of imaging voxels. Some exceptions include Markov random field (MRF) based methods that have been proposed in the prediction context (Smith & Fahrmeir, 2007; Lee et al., 2014). However, given that these are not equipped to perform dimension reduction, they may not be fully scalable to high-dimensional images with tens of thousands of voxels, and their performance in classification problems is unclear. In order to tackle the spatial information in the image in the context of multi-class classification, Pan et al. (2018) proposed a penalized linear discriminant analysis (LDA) model using scalar and tensor covariates. Unfortunately, there is limited, if any, literature on Bayesian classification approaches based on imaging features that account for the spatial information in the image. This is surprising, given the utility of Bayesian methods that can predict the probability of an observation belonging to two or more classes, which can be useful in the presence of measurement error or uncertainty regarding class labels in medical imaging studies (Morales et al., 2013; Behler et al., 2022). Existing Bayesian classification approaches that use vectorized features can be not readily adapted to our problem of interest involving Bayesian image-based classification, since it ignores the spatial structure of the image resulting in information loss and potentially poor model performance. Additionally, simply vectorizing the imaging features without an appropriate lower dimensional representation also introduces the curse of dimensionality since the number of voxels in the image is typically tens of thousands. Alternate approaches that rely on first extracting lower dimensional features from the image and subsequently using these features for classification, may involve an additional layer of information loss resulting from the feature extraction step, resulting in potential loss in accuracy.

Recently, there has been a growing literature on tensor analysis in statistical modeling for imaging data that addresses some of the above concerns. Guhaniyogi et al. (2017) proposed a Bayesian tensor regression with a scalar response on scalar and tensor covariates. Other tensor models include Bayesian response regression models that model the image outcome as a tensor object. Guhaniyogi and Spencer (2021) implemented a Bayesian Tensor response on scalar regression with an application of neuronal activation detection in fMRI experiments using both tensor-valued brain images and scalar covariates. Kundu et al. (2023) proposed a longitudinal Bayesian tensor response regression model for mapping neuroplasticity across longitudinal visits. Billio et al. (2023) proposed a novel linear autoregressive tensor process model that introduces dynamics in linear tensor regression and allows for both tensor-valued covariates and outcomes. Under the frequentist approach, Lock (2018)

proposed a penalized tensor on tensor regression using a (L2) Ridge penalty. Zhou et al. (2013) proposed a tensor regression model by extending generalized linear model to include tensor-structured covariates, where the rank- R PARAFAC decomposition is assumed on the tensor parameters, with adaptive lasso penalties applied on the tensor margins. With the exception of the penalized GLM approach in Zhou et al. (2013), most existing tensor-based approaches in literature have mainly focused on linear regression models that can not be readily used for Bayesian classification.

In this article, we propose a data augmentation-based Bayesian classification approach that models binary outcomes based on imaging covariates using a tensor-based representation. We consider two different data augmentation schemes resulting in two distinct Bayesian classifiers: a support vector machine (SVM) and a logistic regression model. While these classifiers have been extensively used in literature, the focus has been on using non-structured covariates that ignore the spatial structure embedded in the image. Our specific interest is in Bayesian classification based on imaging predictors, where the images are registered across samples. Such a set-up is routinely used in neuroimaging studies. Our main contribution is to develop a Bayesian classification methodology based on high-dimensional tensor-valued predictors via low-rank decompositions of the coefficient matrix and using data augmentation. The low-rank PARAFAC decomposition assumed by the tensor model is able to preserve the spatial configuration of imaging voxels, while overcoming the challenges arising from the high dimensionality of the image that can often contain tens of thousands of voxels. This results in considerable improvements in classification accuracy, as illustrated via rigorous numerical examples. In contrast to existing feature extraction approaches that first use a tensor decomposition or alternate schemes to obtain low level features to be subsequently used in modeling (Sen & Parhi, 2021), the proposed approach uses the full image as is in the classification model, but employs a low rank PARAFAC decomposition to model the high-dimensional tensor model coefficients. This ensures no information loss due to feature extraction, while simultaneously allowing for dimension reduction. We adopt the multiway shrinkage prior from Guhaniyogi et al. (2017) to model the tensor margins of the assumed rank- R PARAFAC decomposition, which shrinks non-significant parameters to near zero while inducing a minimal shrinkage effect on the significant parameters. We develop efficient Markov chain Monte Carlo (MCMC) algorithms for posterior inference that use data augmentation techniques. Simulation studies show significant improvements in classification accuracy, parameter estimation and feature selection compared to routinely used classification methods that use vectorized images. We further illustrate the advantages under our method via a detailed neuroimaging application

using voxel-wise cortical thickness features data from Alzheimer's Disease Neuroimaging Initiative (ADNI) study, with the proposed Bayesian classifiers displaying better out-of-sample accuracy consistently throughout several classification tasks, including classification on pairs of the three diagnostic groups: normal control, AD patients, and MCI patients; gender classification (males vs females); and cognitive performance based on high and low levels of MMSE scores. We leverage cortical thickness as our neuroimaging feature of choice, since it is known to be a highly sensitive imaging biomarker for modeling neurodegeneration in AD (Weston et al., 2016; Fjell et al., 2015).

The rest of the paper is organized as follows: In Section 2, we propose the framework of the Bayesian tensor classification model with two types of data augmentation, specify the prior and hyperparameter choices, and list the posterior computation steps. In Section 3 we study the model performances through comprehensive simulation studies. In Section 4 we provide the results from the data analysis using the ADNI dataset. We conclude the manuscript with a discussion.

Methods

Brief Introduction to Tensors

Tensor-based models have gained recognition as a promising way to model neuroimaging data, due to their multifold advantages. Tensors naturally inherit a multidimensional structure to represent complex data structures such as spatial features of a brain region. Additionally, tensor-based techniques achieve dimension reduction, which is particularly useful with neuroimaging data to tackle the challenges of $p \gg n$ in statistical modeling. A tensor is a multi-dimensional array, with the order being the number of its dimensions. For example, a one-way or first-order tensor is a vector, and a second-order tensor is a matrix. A *fiber* resembles the idea of matrix rows and columns in higher dimensionality, which is obtained by fixing every dimension of a tensor except one. Similarly, a *slice* is defined by fixing every order of the tensor except two. Tensor decomposition is a mathematical technique that expresses a high dimensional tensor into the combination of lower dimensional factors. One type of tensor decomposition is the Tucker decomposition (Kolda & Bader, 2009), which decomposes a tensor into a core tensor and a set of matrices, one along each mode. It can be denoted as follows:

$$\begin{aligned} B &= \Lambda \times_1 A \times_2 B \times_3 \cdots \times_D D \\ &= \sum_{r_1=1}^{R_1} \cdots \sum_{r_D=1}^{R_D} \lambda_{r_1, \dots, r_D} a^{(r_1)} \circ \cdots \circ d^{(r_D)}, \end{aligned} \quad (1)$$

where Λ is the core tensor, and A, B, D are the factor matrices. The PARAFAC decomposition is a special case of the Tucker decomposition, where the core tensor Λ is restricted to be diagonal, and $R_1 = R_2 = \dots = R_D = R$. The rank- R PARAFAC model can then be expressed as

$$B = \sum_{r=1}^R \beta_1^{(r)} \circ \dots \circ \beta_D^{(r)}, \quad (2)$$

where $\beta_1, \dots, \beta_D$, known as tensor margins, are vectors of length $p_1, \dots, p_D$, and where $\beta_1 \circ \dots \circ \beta_D$ is a D -way outer product of dimension $p_1 \times p_2 \times \dots \times p_D$. It is essential to recognize that tensor margins can only be uniquely identified up to a permutation and a multiplicative constant unless we introduce additional constraints. However, the lack of identifiability in tensor margins does not create any complications for our scenario. This is because the tensor product is fully identifiable, which suffices for our primary objective of estimating coefficients. Consequently, we refrain from imposing extra identifiability conditions on the tensor margins, aligning with the principles in the Bayesian tensor modeling literature (Guhaniyogi, 2020). Furthermore, the PARAFAC decomposition dramatically reduces the number of coefficients from $p_1 \times \dots \times p_D$ to $R(p_1 + \dots + p_D)$, which grows linearly with the tensor rank R and results in significant dimension reduction. The appropriate tensor rank can vary depending on the specific application context and can be chosen using a goodness-of-fit approach.

Prior to applying the tensor model, the image's voxels are transformed onto a regularly spaced grid, making them more suitable for a tensor-based approach. This mapping conserves the spatial arrangements of the voxels, offering notable advantages over a univariate voxel-wise analysis or a multivariable analysis that vectorizes the voxels without considering their spatial arrangements. While the grid mapping might not preserve exact spatial distances between voxels, this has limited impact, as it can still capture correlations between neighboring elements in the tensor margins. Moreover, the tensor construction has the advantageous ability to estimate voxel-specific coefficients by leveraging information from neighboring voxels through the estimation of tensor margins with their inherent low-rank structure. This feature results in brain maps which are more consistent and robust to missing voxels and image noise. Additionally, it can be conveniently used for reliable imputation of imaging features related to missing voxels. In contrast, voxel-wise analysis lacks the capacity to share information among neighboring voxels, treating voxel

coefficients as independent entities regardless of their spatial arrangement. For more details on the advantages and characteristics of Bayesian tensor models, we refer readers to Kundu et al. (2023).

Loss-function Based Classification

A loss function is a mathematical tool to quantify the difference between predicted values under a model and the actual observed data values. Loss functions are often used for finding optimal estimates of model parameters in machine learning and statistical modeling tasks for regression, classification, and more. In the Bayesian paradigm, loss functions translate to different types of likelihood that are combined with additional priors on the model parameters embedded in the loss function to obtain posterior distributions that are subsequently used for estimation and uncertainty quantification. Although we motivate our approach by drawing connections with loss functions, we note that our work is distinct compared to decision theoretic Bayesian approaches that use loss functions as a post-processing step after Markov chain Monte Carlo (MCMC) to derive optimal estimates. See, for example, Hahn and Carvalho (2015) and Kundu et al. (2019). In this article we focus, in particular, on two types of commonly used loss functions: the hinge loss represented as the support vector machine classifier, and the logistic regression loss, both of which employ high-dimensional images as covariates for classification.

Support vector machine (SVM) SVMs play a pivotal role in classification tasks, and their significance stems from their ability to handle complex decision boundaries with remarkable efficiency. SVMs work by finding an optimal hyperplane that separates the data points into two classes. This hyperplane is chosen to maximize the margin, which is the distance between the hyperplane and the closest data points. This helps to ensure that the SVM model is generalizable to new data, and protects against overfitting. SVMs are well-suited for both linear and non-linear classification, via using suitable kernel functions. This versatility makes SVMs applicable across various domains, from image recognition and natural language processing to bioinformatics - see Cervantes et al. (2020) for a review. Their robust performance, ability to manage high-dimensional data, and capacity to handle intricate relationships between features underscore their importance in tackling diverse and challenging classification problems in machine learning.

Most SVM-based classifiers rely on point estimates with penalized approaches to tackle high-dimensional covariates (Peng et al., 2016; Dedieu, 2019). The SVM classifier uses the hinge loss function that takes the form

$$\mathcal{L}(y | \beta) = \frac{1}{\sigma^2} \max(1 - yf(\mathbf{x}; \beta), 0) + R, \quad (3)$$

where $y \in \{-1, 1\}$ is the binary outcome, $f(\cdot)$ is a linear or non-linear function of (potentially high-dimensional) covariates $\mathbf{x}$, with corresponding unknown parameters β that need to be estimated from the data, and additional tuning parameters σ^2 . See Fig. 1 for a visualization of the hinge loss function. Recently, Ma and Kundu (2022) generalized the hinge loss to a smooth hinge loss and derived provably flexible estimators in the presence of noisy high-dimensional covariates in the SVM framework. Bayesian approaches for SVM based on a pseudo-likelihood approach were proposed originally by Polson and Scott (2011) and subsequently adopted for biomedical applications in Sun et al. (2018). In particular, the pseudo-likelihood can be represented as a location-scale mixture of normals with latent variable ρ as

$$\begin{aligned} L &= \prod_{i=1}^n L_i(y_i | \mathbf{x}_i, \beta, \sigma^2) = \prod_{i=1}^n \left\{ \frac{1}{\sigma^2} \exp\left\{-\frac{2}{\sigma^2} \max(1 - y_i f(\mathbf{x}_i; \beta), 0)\right\} \right\} \\ &= \int_0^\infty \prod_{i=1}^n \frac{1}{\sigma \sqrt{2\pi} \rho_i} \exp\left(-\frac{(1 + \rho_i - y_i f(\mathbf{x}_i; \beta))^2}{2\rho_i \sigma^2}\right) d\rho_i, \end{aligned} \quad (4)$$

where L_i represents the contribution corresponding to the i th sample. The above representation essentially uses data augmentation techniques, introducing a latent variable ρ which, when marginalized over, gives back the hinge loss function. Such a latent variable representation enables an efficient Gibbs sampler for posterior inference, which will be described in Section 2.4 below.

Logistic regression classifier Logistic regression is a powerful and versatile machine learning algorithm that is widely used in diverse applications, including medical diagnosis, fraud detection, customer segmentation, marketing campaigns, and recommendation systems. It is characterized by simplicity in the implementation and interpretability of the results, as the logistic regression coefficients can be understood in terms of odd-ratios. In order to make this model scalable for high dimensional biomedical applications, penalized versions of logistic regression models have been proposed (Doerken et al., 2019; Devika et al., 2016).


The logistic loss function specifies a sigmoid type loss that takes the analytical form

$$\mathcal{L}(y = 1 | \beta) = \exp\{f(\mathbf{x}; \beta)\} / (1 + \exp\{f(\mathbf{x}; \beta)\}), \quad (5)$$

where $f(\mathbf{x}; \beta)$ represents the contribution of the covariates to the logistic loss that is quantified via the unknown parameters β , to be estimated from the data. See Fig. 1 for a visualization of the sigmoid loss function. The binary outcome variables $y \in \{0, 1\}$ are assumed to follow a Bernoulli distribution, with the corresponding probability function resembling the form $\mathcal{L}(y = 1 | \beta)$. Typically $f(\cdot)$ represents linear functions of covariates that facilitate straightforward interpretations for the model parameters, although non-linear logistic regression models have also been proposed (Tokdar & Ghosh, 2007). In the Bayesian paradigm, model fitting and inference for logistic regression is often achieved using Polya-Gamma latent variables (Nicholas et al., 2013). A random variable X has a Polya-Gamma distribution with parameters $b > 0$ and $c \in \mathcal{R}$, denoted as $X \sim PG(b, c)$ if

$$X \stackrel{D}{=} \frac{1}{2\pi^2} \sum_{k=1}^{\infty} \frac{g_k}{(k - 1/2)^2 + c^2 / (4\pi^2)}, \quad (6)$$

Fig. 1 From left to right: hinge loss and sigmoid loss function



where g_k follows a Gamma distribution $Ga(b, 1)$. The introduction of the Polya-Gamma latent variable allows one to represent the binomial likelihood as mixtures of Gaussians. It can be shown that the logistic loss function can be recovered by marginalizing out the latent Polya-Gamma latent variable using the following relationship

$$\frac{(e^{f(\cdot)})^y}{(1 + e^{f(\cdot)})^b} = 2^{-b} e^{\kappa \psi} \int_0^\infty e^{-\omega \psi^2/2} p(\omega) d\omega, b > 0, \kappa = y - b/2, \quad (7)$$

where $\omega \sim PG(b, 0)$, and $f(\mathbf{x}; \boldsymbol{\beta})$ is a linear predictor in most cases. The above identity facilitates conjugate updates under a Gaussian prior, conditional upon the latent Polya-Gamma variable, as detailed in Nicholas et al. (2013). The full data augmented likelihood is given by the following expression:

$$L = \prod_{i=1}^n \frac{(e^{f_i})^{y_i}}{(1 + e^{f_i})^b} = \prod_{i=1}^n 2^{-1} e^{\kappa_i \psi_i} \int_0^\infty e^{-\omega_i f_i^2/2} p(\omega_i) d\omega_i, \quad (8)$$

where $\kappa_i = y_i - 1/2$, $b = 1$, $\omega_i \sim PG(1, 0)$.

Priors

For our implementation, we consider the widely used linear predictor using the relationship

$$f_i = \langle \mathbf{X}_i, \mathbf{B} \rangle + \mathbf{z}_i' \boldsymbol{\gamma}, \quad (9)$$

where $\mathbf{X}_i$ and $\mathbf{z}_i$ denote the imaging predictors and the supplemental (eg: demographic/clinical) features, respectively, for the i th sample, $\langle \cdot, \cdot \rangle$ denotes the inner product operator, $\mathbf{B}$ denotes the tensor-valued coefficient matrix that quantifies the effect of the image on the classification model, and $\boldsymbol{\gamma}$ is a vector of dimension $p_z + 1$ capturing the effects of the supplemental covariates. Furthermore, we assume $\mathbf{B} \in \otimes_{j=1}^D \mathbb{R}^{p_j}$, that is modeled under a PARAFAC decomposition as in (2).

For the prior choice on the tensor margins, we adopt the multiway Dirichlet generalized double Pareto (M-DGDP) prior from Guhaniyogi et al. (2017), which shrinks small

coefficients towards zero while minimizing shrinkage of large coefficients. The prior can be expressed in hierarchical form on the tensor coefficient margins $\boldsymbol{\beta}_j^{(r)}$, $j = 1, \dots, D$ and $r = 1, \dots, R$ as:

$$\boldsymbol{\beta}_j^{(r)} \sim N(0, (\phi_r \tau) \mathbf{W}_{jr}), w_{jr,k} \sim \text{Exp}(\lambda_{jr}^2/2), \quad (10)$$

where $\tau \sim Ga(a_\tau, b_\tau)$ is a global scale parameter, $\Phi = (\phi_1, \dots, \phi_R)$ follows a Dirichlet distribution that encourages shrinkage to lower rank in the assumed PARAFAC decomposition with $(\phi_1, \dots, \phi_R) \sim \text{Dirichlet}(\alpha_1, \dots, \alpha_R)$, and where $\mathbf{W}_{jr} = \text{Diag}(w_{jr,1}, \dots, w_{jr,p_j})$ are scale parameters that are margin-specific for each element and modeled under an Exponential distribution as $w_{jr,k} \sim \text{Exp}(\lambda_{jr}^2/2)$ with λ_{jr} unknown and modeled as $\lambda_{jr} \sim Ga(a_\lambda, b_\lambda)$. Additionally, one can obtain

$$\boldsymbol{\beta}_{j,k}^{(r)} \mid \lambda_{jr}, \phi_r, \tau \stackrel{\text{iid}}{\sim} \text{DE}(\lambda_{jr}/\sqrt{\phi_r \tau}), 1 \leq k \leq p_j, \quad (11)$$

after marginalizing out the scale parameters $w_{jr,k}$, that is, prior (10) induces a GDP prior on the individual margin coefficients which in turn has the form of an adaptive Lasso penalty as in Armagan et al. (2013). Overall, the flexibility in estimating $\mathbf{B}_r = \{\boldsymbol{\beta}_j^{(r)}; 1 \leq j \leq D\}$ is accommodated by component-specific scaling parameter $w_{jr,k}$ and common rate parameter λ_{jr} , which shares information between margin elements and encourages shrinkage at the local scale. We complete the prior specification by assuming a $N(0, \Sigma_{0\gamma})$ as the prior on $\boldsymbol{\gamma}$.

MCMC Algorithms for Posterior Inference

For posterior inference, we implemented efficient MCMC algorithms that take advantage of data augmentation techniques. Algorithm 1 outlines the MCMC updates for the proposed Bayesian Tensor SVM model (BT-SVM), and Algorithm 2 illustrates the updating scheme for the the proposed Bayesian Tensor logistic regression model (BT-LR).

Algorithm 1 MCMC steps for BT-SVM

-
- 1: Update ρ_i from inverse Gaussian: $\rho_i^{-1} \sim IN(\mu_i, \lambda_i)$, where $\mu_i = |1 - y_i(< X_i, B > + z'\gamma)|^{-1}$ and $\lambda_i = 1/\sigma^2$
 - 2: Update hyperparameters $[\alpha, \Phi, \tau \mid \mathbf{B}, \mathbf{W}]$ compositionally as $[\alpha \mid \mathbf{B}, \mathbf{W}][\Phi, \tau \mid \alpha, \mathbf{B}, \mathbf{W}]$, as described in Guhaniyogi et al. (2017).
 - 3: Sample $\{\beta_j^{(r)}, \omega_{jr}, \lambda_{jr}\}$ using a back-fitting procedure to produce a sequence of draws from the margin-level conditional distributions across components.
 - (a) Draw $[w_{jr}, \lambda_{jr} \mid \beta_j^{(r)}, \phi_r, \tau] = [w_{jr} \mid \lambda_{jr}, \beta_j^{(r)}, \phi_r, \tau] [\lambda_{jr} \mid \beta_j^{(r)}, \phi_r, \tau]$.
 1. Draw $\lambda_{jr} \sim \text{Ga}(a_\lambda + p_j, b_\lambda + \|\beta_j^{(r)}\|_1 / \sqrt{\phi_r \tau})$;
 2. Draw $w_{jr,k} \sim \text{giG}(\frac{1}{2}, \lambda_{jr}^2, \beta_{j,k}^{2(r)} / (\phi_r \tau))$ independently for $1 \leq k \leq p_j$
 - (b) Draw $\beta_j^{(r)}$ from multivariate normal distribution: $\beta_j^{(r)} \sim N(\mu_{jr}, \Sigma_{jr})$, where $\mu_{jr} = \frac{\Sigma_{jr} H_j^{(r)} \tilde{\mathbf{y}}}{\sigma^2}$ and $\Sigma_{jr} = (\frac{H_j^{(r)T} H_j^{(r)}}{\sigma^2} + \frac{W_{jr}^{-1}}{\phi_r \tau})^{-1}$, with

$$h_{i,j,k}^{(r)} = \sum_{d_1=1, \dots, d_D=1}^{p_1, \dots, p_D} I(d_j = k) x_{d_1, \dots, d_D} \left(\prod_{l \neq j} \beta_{l, i_l}^{(r)} \right),$$

$$\mathbf{H}_{i,j}^{(r)} = \left(h_{i,j,1}^{(r)} / \sqrt{\rho_i}, \dots, h_{i,j,p_j}^{(r)} / \sqrt{\rho_i} \right)',$$

$$\tilde{y}_i = \frac{y_i}{\sqrt{\rho_i}} (\rho_i + 1 - y_i (z'_i \gamma + \sum_{l \neq r} < X_i, B_l >))$$

- 4: Update γ from a conjugate normal conditional distribution $\gamma \sim N(\mu_\gamma, \Sigma_\gamma)$ where $\mu_\gamma = \Sigma_\gamma Z^T (\tilde{y} y / \rho)$ and $\Sigma_\gamma = (G^T G / \sigma^2 + \Sigma_{0\gamma})^{-1}$, with $\mathbf{G}_{i,p_z} = Z_{i,p_z} / \sqrt{\rho_i}$ and $\tilde{y}_i = \rho_i + 1 - y_i < X_i, B >$.
-

Algorithm 2 MCMC steps for BT-LR

-
- 1: Update ω_i from Polya-gamma distribution: $\omega_i \sim \text{PG}(1, \langle X_i, \mathbf{B} \rangle + z_i^T \gamma)$
 - 2: Update hyperparameters $[\alpha, \Phi, \tau \mid \mathbf{B}, \mathbf{W}]$ compositionally as $[\alpha \mid \mathbf{B}, \mathbf{W}][\Phi, \tau \mid \alpha, \mathbf{B}, \mathbf{W}]$, as described in Guhaniyogi et al. (2017).
 - 3: Sample $\{\beta_j^{(r)}, \omega_{jr}, \lambda_{jr}\}$ using a back-fitting procedure to produce a sequence of draws from the margin-level conditional distributions across components.
 - (a) Draw $[w_{jr}, \lambda_{jr} \mid \beta_j^{(r)}, \phi_r, \tau] = [w_{jr} \mid \lambda_{jr}, \beta_j^{(r)}, \phi_r, \tau] [\lambda_{jr} \mid \beta_j^{(r)}, \phi_r, \tau]$.
 1. Draw $\lambda_{jr} \sim \text{Ga}(a_\lambda + p_j, b_\lambda + \|\beta_j^{(r)}\|_1 / \sqrt{\phi_r \tau})$;
 2. Draw $w_{jr,k} \sim \text{giG}(\frac{1}{2}, \lambda_{jr}^2, \beta_{j,k}^{2(r)} / (\phi_r \tau))$ independently for $1 \leq k \leq p_j$
 - (b) Draw $\beta_j^{(r)}$ from multivariate normal distribution: $\beta_j^{(r)} \sim N(\mu_{jr}, \Sigma_{jr})$, with $\mu_{jr} = \Sigma_{jr}(\Omega(H_j^{(r)})^T \tilde{y})$ and $\Sigma_{jr} = ((H_j^{(r)})^T \Omega H_j^{(r)} + W_{jr}^{-1} / (\phi_r \tau))^{-1}$, where $\tilde{y} = \kappa / \omega$, $\kappa = (y_1 - N_1/2, \dots, y_n - N_n/2)$, $N_1 = \dots = N_n = 1$, and Ω a diagonal matrix with diagonal elements ω_i 's.
 - 4: Update γ from a conjugate normal conditional distribution $\gamma \sim N(\mu_\gamma, \Sigma_\gamma)$, where $\mu_\gamma = \Sigma_\gamma Z^T (\tilde{y} \omega)$ and $\Sigma_\gamma = (G^T G + \Sigma_{0\gamma})^{-1}$, with $G_{i,p_z} = Z_{i,p_z} * \sqrt{\omega_i}$ and $\tilde{y}_i = \kappa_i / \omega_i - \langle X_i, \mathbf{B} \rangle$.
-

Simulation Study**Data Generation**

We illustrate the performance of our methods using several simulation settings and perform comparisons with competitive approaches, based on various types of generated data sets involving several types of functional signals and with data generated from SVM and logistic loss functions. We considered four different types of signals for the tensor coefficient $\mathbf{B}$ to generate the binary outcome, as defined below.

Scenario 1 In this setting, the tensor $\mathbf{B}$ is constructed from rank- R PARAFAC decomposition with rank $R_0 = 3$ and dimension $p = c(48, 48)$. Each beta margin $\beta_j^{(r)}$ is generated from the independent binomial distribution $\text{Binomial}(2, 0.2)$. After the construction of tensor, we set the maximal value of the tensor $\mathbf{B}$ cells to be 1.

Scenario 2 The tensor image is simulated by a rank- R PARAFAC decomposition with rank $R_0 = 3$. Here, instead of generating the tensor margin from a known distribution, we manually set up each value of $\beta_j^{(r)}$.

Scenario 3 Instead of generating the 2D tensor images from a PARAFAC decomposition, the tensor coefficient $\mathbf{B}$ is set to be 1 for a rectangular area and 0 otherwise. The non-zero elements cover approximately 30 percent of the area.

Scenario 4 The tensor coefficient $\mathbf{B}$ is set to be 1 for the circular area and 0 otherwise. The non-zero elements cover approximately 10 percent of the area.

Scenario 5 (using real 2D brain image) We considered a simulated scenario that uses 2D cortical thickness images from AD and NC patients in ADNI-1 study (see Section 4). In addition to these brain images, we use MMSE scores (used to measure cognitive impairment) as covariates to generate synthetic binary responses under both SVM and logistic link functions and simulated tensor coefficients. We selected three 2D cortical thickness slices (derived from the T1 weighted MRI scans and denoted as slices 20, 21, and 22) from the normal control group and the Alzheimer's disease group, with MMSE scores available. MMSE scores are treated as scalar covariates while the 2D cortical thickness slices are tensor covariates. The binary response is constructed under both the SVM and logistic loss functions. The true tensor coefficient $\mathbf{B}$ is constructed from a rank- R PARAFAC decomposition with rank $R_0 = 2$, where each tensor margin $\beta_j^{(r)}$ is being pre-specified.

The top panel in Fig. 2 shows the true 2D tensor images, for the different scenarios. For each setting, we generated the tensor covariates $\mathbf{X}$ from standard normal distribution $N(0, 1)$. For simplicity, we did not include non-tensor covariates in our simulation settings, i.e. we assumed the true $\gamma = (0, \dots, 0)'$. Finally for each scenario, the binary




Fig. 2 Row 1 from left: Simulated data with 48×48 2D tensor images from Scenarios 1-4. Row 2: Recovered images for the 48×48 2D tensor images using BT-SVM corresponding to 4 scenarios in row 1

1. Row 3: Recovered images for the 48×48 2D tensor images using BT-LR corresponding to 4 scenarios in row 1

outcome Y was generated according to both the SVM and logistic loss function as follows. Denoting the linear prediction as $\psi = \langle X_i, \mathbf{B} \rangle$, the binary outcome is generated as $Y_i = 1$ if $\psi > 0$ and $Y_i = -1$ otherwise, under the SVM loss, and from a Bernoulli distribution with probability $p = 1/(1 + \exp(-\psi))$ under a logistic loss. Therefore a total of 8 scenarios are considered in our simulation set-up, and for each of these scenarios 10 replicates were generated.

Parameter Settings

We choose suitable values of hyperparameters in the prior distributions that yield good overall performance. For example, we set the parameters of the hyperprior on the global scale τ to $a_\tau = 1$ and $b_\tau = \alpha R^{(1/D)}$, where R is the rank in the assumed PARAFAC decomposition, and set $\alpha_1 = \dots = \alpha_R = 1/R$. For the common rate parameter λ_{jr} , we set $a_\lambda = 3$ and $b_\lambda = \sqrt[3]{a_\lambda}$. Note that under the SVM loss, the scaling parameter σ^2 is a fixed parameter that can be manually adjusted for maximal model performance. Several values of the tuning parameter σ^2 from 0.1 through 10 are tested and we choose $\sigma^2 = 6$. In order to decide the rank of the fitted model, we fit the proposed model using ranks 2-5 and choose the rank that minimizes the Deviance

Information Criterion (DIC) scores. DIC measures the goodness-of-fit of a set of Bayesian hierarchical models adjusting for model complexity in a manner that penalizes more complex models.

Performance Evaluation

We report estimation accuracy in terms of the Root Mean Squared Error (RMSE) and correlation coefficient for point estimation of cell-level tensor coefficients. We also illustrate the classification accuracy by calculating misclassification error and the F1 score. Feature selection performances are evaluated by Sensitivity, Specificity, F1 score, and Matthews correlation coefficient (MCC). These metrics are defined as follows. Let θ_j , $j = 1, \dots, J$, be the vectorized tensor coefficients, with $J = \prod_{k=1}^D p_k$ the total number of cells in the tensor coefficient $\mathbf{B}$. Further, define the following terms related to classification performance under the SVM classifier: (a) TP is the true positive, i.e. the number of predictions where the classifier correctly predicts the positive class as positive; (b) FP is the false positive, i.e. the number of predictions where the classifier incorrectly predicts the negative class as positive; (c) TN is the true negative, i.e. the number of predictions where the classifier correctly predicts a negative class as negative;

and (d) FN is a false negative, i.e. the number of predictions where the classifier incorrectly classifies a positive class as negative. For logistic classification, the above definitions also hold, but after replacing the negative class with zero class. The above definitions of TP/FP/TN/FN can also be adopted for feature selection performance such that the positive class corresponds to non-zero coefficients, and the negative/zero class refers to absent or zero coefficients.

Metrics for evaluating coefficient estimation performance These metrics include: (i) Root-mean-square error

of θ , denoted as $RMSE(\theta) = \sqrt{\frac{\sum_{i=1}^J (\hat{\theta}_i - \theta_i)^2}{J}}$, that provides another measure of estimation accuracy; (ii) correlation coefficient between the true and estimated coefficients.

Metrics for evaluating classification performance These metrics include: (i) Mis-classification rate, defined as $\frac{FP+FN}{TP+TN+FP+FN}$; and (ii) F1-score, defined as the harmonic mean between precision (i.e. $TP/(TP+FP)$) and recall or sensitivity ($TP/(TP+FN)$). The expression of F1-score is given as $\frac{TP}{TP+(FP+FN)/2}$.

Metrics for evaluating feature selection performance These metrics include: (i) Sensitivity = $\frac{TP}{TP+FN}$; (ii) Specificity = $\frac{TN}{TN+FP}$; and (iii) MCC = $\frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$.

Results reported below were obtained by randomly splitting the data into training and test sets in the ratio 70:30. The metrics for point estimation and feature selection performances were calculated using the training set data, while the metrics for classification performance were calculated based on the test set. We report below averaged values of the selected metrics across 10 replicates. Two state-of-the-art classification methods are used as competing methods. The first is a penalized logistic regression model with lasso penalty, which is available in the R package `glmnet` (Friedman et al., 2010). The second competing method is the L1-norm SVM model from the R package `penalizedSVM` (Becker et al., 2009; Bradley & Mangasarian, 1998). Both methods use a vectorization approach, where the tensor covariates are vectorized into scalar variables and then fit into the statistical models. Therefore they do not respect the spatial information in the image. Additionally, we use a grid search algorithm and cross-validation to select the best tuning parameters prior to model fitting.

Results

We ran MCMC chains for 3,000 iterations, with 1,000 burn-in iterations. The computation time varies depending on the

rank and is expected to increase with higher ranks. It took around 19 minutes to run a single MCMC chain with rank 2, and around 36 minutes with rank 4. A Geweke diagnostic (Geweke, 1991) was applied to examine for signs of non-convergence parameters. We obtain the z-score from Geweke for each element of the coefficient matrix B . For the proposed Bayesian tensor SVM model (BT-SVM), we observe the z-scores lie in the range $(-1.96, 1.96)$ for 91 percent of the coefficient matrix elements. For the proposed Bayesian tensor logistic regression model (BT-LR), the z-scores lie in the range $(-1.96, 1.96)$ for 79 percent, indicating that most chains have reached ergodicity.

Scenarios 1-4 We report results for estimation and classification accuracy, and feature selection in Tables 1, 2, 3, and 4 corresponding to Scenarios 1-4. Specifically, Tables 1 and 2 reflect results across all 4 scenarios when the binary outcome Y is generated from a SVM loss, while Tables 3 and 4 reflect the logistic type of loss. These results demonstrate that both the proposed methods (BT-SVM and BT-LR) consistently outperform competing penalized methods in terms of coefficient estimation, feature selection, and classification performance across all scenarios. When the binary outcome data is generated from SVM loss, the BT-SVM approach has superior coefficient estimation (as evident from lower RMSE and higher correlation coefficient in Table 1) and improved classification accuracy (as evident from lower misclassification rate and higher F1-score in Table 1). Even when the data is generated from a logistic loss, the same trends hold, with the exception of Scenario 1 where the proposed BT-LR approach reports improved classification accuracy and comparable coefficient estimation, as evident from the results in Table 3.

Moreover, the proposed model with SVM loss (BT-SVM) generally performs worse than the corresponding model with logistic loss (BT-LR) in terms of feature selection, as evident from the results in Tables 2 and 4. In particular, the BT-LR approach almost always has considerably higher sensitivity compared to the BT-SVM model while having comparable or slightly lower specificity, even when the outcome data is generated under the SVM loss. This results in the BT-LR approach consistently having higher F1-score and higher or comparable MCC values for feature selection, when compared with the BT-SVM approach, even when the outcome data is generated under the SVM loss.

Combining the above discussions, the BT-SVM model appears to generally show improved coefficient estimation and classification performance over its counterpart with logistic loss, regardless of which type of loss function was used to generate the outcome data. On the contrary, the BT-LR model appears to have improved sensitivity and comparable specificity compared to its counterpart with SVM

Table 1 Point Estimation and Out-of-sample Classification results for the four 2D tensor images portrayed in Fig. 2 top panel; Y generated from SVM loss

Scenarios	Methods	RMSE	Corr.Coeff.	Mis. Class.	F1-score
Scenario 1	LR w/ lasso	0.558	0.071	0.54	0.147
	L1norm-SVM	0.687	0.027	0.513	0.522
	BT-SVM	0.489	0.442	0.260	0.764
	BT-LR	0.504	0.383	0.342	0.689
Scenario 2	LR w/ Lasso	0.382	0.055	0.46	0
	L1norm-SVM	0.533	0.092	0.48	0.502
	BT-SVM	0.246	0.874	0.149	0.837
	BT-LR	0.277	0.794	0.204	0.778
Scenario 3	LR w/ Lasso	0.541	0.074	0.507	0.191
	L1norm-SVM	0.537	0.109	0.533	0.512
	BT-SVM	0.412	0.810	0.197	0.823
	BT-LR	0.439	0.733	0.238	0.781
Scenario 4	LR w/ Lasso	0.330	0.122	0.533	0.2
	L1norm-SVM	0.535	0.053	0.44	0.565
	BT-SVM	0.225	0.773	0.200	0.819
	BT-LR	0.262	0.614	0.272	0.752

loss that translates to improved feature selection, regardless of which loss function is used to generate the outcome data. Figure 2 presents the coefficient matrix cell estimation

Table 2 Feature selection results for the four 2D tensor images portrayed in Fig. 2 top panel; Y generated from SVM loss

Scenarios	Methods	Sens.	Spec.	F1-score	MCC
Scenario 1	LR w/ lasso	0.031	0.984	0.057	0.046
	L1norm-SVM	0.058	0.950	0.1	0.017
	BT-SVM	0.070	0.999	0.130	0.213
	BT-LR	0.214	0.938	0.313	0.225
Scenario 2	LR w/ Lasso	0.003	1	0.007	0.055
	L1norm-SVM	0.021	0.985	0.037	0.016
	BT-SVM	0.628	1	0.771	0.772
	BT-LR	0.821	0.978	0.835	0.814
Scenario 3	LR w/ Lasso	0.025	0.981	0.047	0.019
	L1norm-SVM	0.019	0.988	0.037	0.027
	BT-SVM	0.192	1	0.32	0.377
	BT-LR	0.517	0.987	0.668	0.626
Scenario 4	LR w/ Lasso	0.094	0.974	0.145	0.122
	L1norm-SVM	0.020	0.984	0.034	0.009
	BT-SVM	0.368	0.998	0.530	0.573
	BT-LR	0.519	0.934	0.552	0.522

Table 3 Point Estimation and Out-of-sample Classification results for the four 2D tensor images portrayed in Fig. 2 top panel; Y generated from logistic regression loss

Scenarios	Methods	RMSE	Corr.Coeff.	Mis. Class.	F1-score
Scenario 1	LR w/ lasso	0.557	0.076	0.56	0.125
	L1norm-SVM	0.562	0.054	0.48	0.55
	BT-SVM	0.493	0.495	0.274	0.745
	BT-LR	0.490	0.445	0.247	0.767
Scenario 2	LR w/ Lasso	0.382	0.085	0.46	0
	L1norm-SVM	0.383	0.053	0.413	0.537
	BT-SVM	0.245	0.859	0.179	0.804
	BT-LR	0.260	0.732	0.252	0.734
Scenario 3	LR w/ Lasso	0.542	0.071	0.506	0.380
	L1norm-SVM	0.587	0.041	0.433	0.586
	BT-SVM	0.412	0.804	0.178	0.842
	BT-LR	0.421	0.618	0.253	0.770
Scenario 4	LR w/ Lasso	0.331	0.131	0.467	0.557
	L1norm-SVM	0.542	0.058	0.467	0.557
	BT-SVM	0.226	0.766	0.203	0.813
	BT-LR	0.255	0.658	0.227	0.788

using the proposed methods BT-SVM and BT-LR. From this Figure, it is evident that the proposed method is able to broadly recover the shapes of the 2D tensor $\mathbf{B}$ regardless of

Table 4 Feature selection results for the four 2D tensor images portrayed in Fig. 2 top panel; Y generated from logistic regression loss

Scenarios	Methods	Sens.	Spec.	F1-score	MCC
Scenario 1	LR w/ lasso	0.036	0.979	0.067	0.045
	L1norm-SVM	0.023	0.982	0.044	0.020
	BT-SVM	0.007	0.999	0.014	0.061
	BT-LR	0.343	0.913	0.446	0.318
Scenario 2	LR w/ Lasso	0.017	0.999	0.034	0.098
	L1norm-SVM	0.042	0.978	0.070	0.042
	BT-SVM	0.590	0.999	0.736	0.742
	BT-LR	0.801	0.948	0.764	0.736
Scenario 3	LR w/ Lasso	0.024	0.984	0.044	0.026
	L1norm-SVM	0.010	0.981	0.020	-0.031
	BT-SVM	0.180	1	0.300	0.361
	BT-LR	0.475	0.926	0.573	0.483
Scenario 4	LR w/ Lasso	0.086	0.972	0.133	0.102
	L1norm-SVM	0.040	0.983	0.067	0.051
	BT-SVM	0.409	0.999	0.577	0.612
	BT-LR	0.609	0.965	0.645	0.610

whether the underlying signal is generated using the PARAFAC decomposition or not.

In contrast, the penalized competing methods (logistic regression with LASSO penalty and L1norm-SVM) illustrate poor estimation, feature selection, and out-of-sample classification performance. This is evident from the results in Tables 2, 4, where the competing approaches report notably low sensitivity, almost approaching zero. This low sensitivity reflects their inability to detect the true signals, resulting in poor coefficient estimation performance as evident from Tables 1 and 3, where the correlation coefficients between the true and estimated coefficients are often close to zero. For a more detailed perspective, Fig. 3 visualizes the estimated coefficients under two competing methods. The Figure reveals spatially disparate non-zero coefficients, indicating that the absence of spatial smoothing hampers the accurate estimation of true signals by these competing methods. Ultimately, the substandard feature selection performance contributes to considerably inferior classification results, as demonstrated in Tables 1 and 3.

Scenario 5 For scenario 5 involving real brain cortical thickness images as covariates, Table 5 provides

parameter estimation for the tensor coefficients $\mathbf{B}$ and out-of-sample classification results under different methods when the outcome Y is generated from SVM loss, and Table 6 provides the corresponding results when Y is generated from the logistic loss. Overall, BT-SVM and BT-LR outperform the competing approaches consistently across all slices, with BT-SVM achieving a misclassification rate < 0.2 and an F1-score > 0.8 across all slices with both types of responses. The coefficient estimation under the spatially informed tensor approach is also improved compared to penalized approaches, as evident from lower RMSE. This is reasonable, as the penalized methods fail to detect most of the true signals, resulting in zero estimates for most cells of the coefficient matrix $\mathbf{B}$. On the other hand, BT-SVM and BT-LR recover the true signals adequately well (as evident from the lower RMSE and higher correlation coefficient values). However, they also result in false positives (they estimate voxels with zero effect sizes to be non-zero), due to the structure of tensor decomposition that results in spatial smoothing. This may occasionally result in slightly inflated coefficient estimation errors, especially under the BT-LR approach.




Fig. 3 Row 1 from left: Simulated data with 48x48 2D tensor images from Scenario 1, Scenario 2, Scenario 3, and Scenario 4. Row 2: Recovered images for the 48x48 2D tensor images using compet-

ing method L1norm-SVM for Scenario 1, Scenario 2, Scenario 3, and Scenario 4. Row 3: Recovered images for the 48x48 2D tensor images using logistic regression with LASSO penalty

Table 5 Simulated scenario 5: Cell Estimation of **B** and Out-of-sample Classification results for the three 2D cortical thickness slices; Y generated from SVM loss

Slices	Methods	RMSE	Corr.Coeff.	Mis. Class.	F1-score
slice 20	LR w/ lasso	0.253	0.141	0.229	0.686
	L1norm-SVM	0.261	0.061	0.409	0.542
	BT-SVM	0.222	0.550	0.156	0.823
	BT-LR	0.253	0.429	0.182	0.784
slice 21	LR w/ lasso	0.252	0.159	0.278	0.636
	L1norm-SVM	0.386	0.068	0.433	0.540
	BT-SVM	0.202	0.626	0.140	0.843
	BT-LR	0.236	0.566	0.196	0.782
slice 22	LR w/ Lasso	0.252	0.176	0.262	0.726
	L1norm-SVM	0.276	0.055	0.448	0.526
	BT-SVM	0.215	0.568	0.150	0.851
	BT-LR	0.245	0.454	0.200	0.806

Sensitivity Analysis

Guhaniyogi et al. (2017) set up a series of default values for the prior hyper-parameters for the Bayesian tensor regression model. Specifically, hyperparameter of the Dirichlet component $(\alpha_1, \dots, \alpha_R) = \alpha = (1/R, \dots, 1/R)$, $a_\lambda = 3$, $b_\lambda = \sqrt[3]{a_\lambda}$, $a_\tau = \sum_{i=1}^R \alpha_i$, and $b_\tau = \alpha R^{(1/D)}$ are set as defaulted values to control the cell-level variance on tensor **B**. We conducted a prior hyperparameter sensitivity analysis by tuning each hyperparameter while fixing other hyperparameters. Table 7 displays the cell-level RMSE values of the tensor coefficient matrix **B** of dimension $D = 2$ and parafac rank- R , where for BT-SVM model, **B** is generated from Scenario 4; and for BT-LR model, **B** is generated from Scenario 3. The overall results reveal no strong sensitivity to the hyper-parameter choices under the selected ranges that point to a robust performance.

Table 6 Simulated scenario 5: Cell Estimation of **B** and Out-of-sample Classification results for the three 2D cortical thickness slices; Y generated from logistic loss

Slices	Methods	RMSE	Corr.Coeff.	Mis. Class.	F1-score
slice 20	LR w/ lasso	0.253	0.155	0.201	0.729
	L1norm-SVM	0.268	0.058	0.411	0.557
	BT-SVM	0.227	0.509	0.155	0.821
	BT-LR	0.244	0.500	0.182	0.777
slice 21	LR w/ lasso	0.253	0.147	0.267	0.650
	L1norm-SVM	0.259	0.021	0.456	0.511
	BT-SVM	0.203	0.620	0.146	0.835
	BT-LR	0.216	0.587	0.196	0.784
slice 22	LR w/ Lasso	0.252	0.161	0.275	0.728
	L1norm-SVM	0.263	0.069	0.410	0.611
	BT-SVM	0.213	0.582	0.149	0.858
	BT-LR	0.231	0.445	0.205	0.806

ADNI Data Analysis

Data Source and Pre-processing

This study utilized the data obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI), which is a longitudinal multicenter study launched in 2004 for the early detection and tracking of Alzheimer's disease (AD). ADNI researchers collect multiple data types such as clinical, behavioral and genetic data, along with neuroimaging measurements such as magnetic resonance imaging (MRI), positron emission tomography (PET), and biospecimens. We use data from the ADNI 1 study collected at baseline, consisting of T1-weighted MRI scans, cognitive measurements in terms of Mini Mental State examination (MMSE), and basic demographic data (age, gender, year of education, APOE status) of 818 subjects. A more detailed description of the ADNI data is provided in Table 8.

The T1-weighted MRI images underwent processing through the Advance Normalization Tools (ANTs) registration pipeline (Tustison et al., 2014), where all images were registered to a template image to ensure consistent normalization of brain locations across participants. The population-based template was constructed based on data from 52 normal control participants in ADNI 1, originally from the ANTs group (Tustison et al., 2019). Notably, the ANTs pipeline includes the N4 bias correction step, addressing intensity discordance to inherently standardize intensity across samples (Tustison et al., 2010). It also employs a symmetric diffeomorphic image registration algorithm for spatial normalization, aligning each T1 image with a brain image template to facilitate spatial comparability (Avants et al., 2008). Subsequently, the pipeline utilized the processed brain images, estimated brain masks, and template tissue labels for 6-tissue Atropos segmentation, generating tissue masks for cerebrospinal fluid (CSF), gray matter (GM), white matter (WM), deep gray matter (DGM), brain stem, and cerebellum. Finally, cortical thickness measurements were derived using the DiReCT algorithm. The 3-D cortical thickness image was further downsampled to a dimension of $48 \times 48 \times 48$ and divided into 48 2-D axial slices of dimension 48×48 , and a subset of these 2D slices were used for our analysis. The downsampling step reduces the dimension of the image, and also somewhat alleviates the sparsity of the cortical thickness maps by consolidating adjacent voxels and presenting the average cortical thickness. A reduction of sparsity proves beneficial to fitting our Bayesian tensor model, and the same should hold of other commonly used statistical models.

Analysis Outline

We apply the proposed approaches to perform various classification tasks using data from the Alzheimer's Disease

Table 7 Prior hyperparameter sensitivity analysis

	Hyperparameters	RMSE	Hyperparameters	RMSE	Hyperparameters	RMSE
BT-SVM	$\alpha = 1/9$	0.224	$a_\lambda = 3$	0.225	$a_\tau = 1/3$	0.225
	$\alpha = 1/6$	0.228	$a_\lambda = 5$	0.235	$a_\tau = 1/2$	0.225
	$\alpha = 1/3$	0.226	$a_\lambda = 7$	0.227	$a_\tau = 1$	0.226
	$\alpha = 3^{(-0.1)}$	0.227	$a_\lambda = 10$	0.230	$a_\tau = 2$	0.231
BT-LR	$\alpha = 1/9$	0.418	$a_\lambda = 3$	0.413	$a_\tau = 1/3$	0.428
	$\alpha = 1/6$	0.422	$a_\lambda = 5$	0.421	$a_\tau = 1/2$	0.416
	$\alpha = 1/3$	0.419	$a_\lambda = 7$	0.413	$a_\tau = 1$	0.416
	$\alpha = 3^{(-0.1)}$	0.421	$a_\lambda = 10$	0.415	$a_\tau = 2$	0.420

Neuroimaging Initiative (ADNI) study. Demographic information of age, gender, years of education, APOE4 allele (0,1,2) and intracranial volume (ICV) are incorporated as scalar covariates, while 2D cortical thickness slices (derived from the T1 weighted MRI scans) are used as tensor covariates. Since the different 2D brain slices are expected to contain varying amount of the brain cortical regions, it is important to choose these slices carefully. In particular, we would like to choose 2D slices such that they contain at least a certain portion of the brain cortex, in order to contain enough information to perform classification. Therefore, we present the analysis results from 7 different axial slices (we will denote them slices 19 to 25), each of which has cortical brain regions that cover at least 65% of the slice. We evaluate the classification performance of the proposed approaches and compare with the penalized logistic regression with lasso and L1-norm SVM as described in Section 3.

In particular, we assess the ability of the proposed classifiers to differentiate between different types of disease phenotypes using the 2D imaging slices along with demographic information. In particular, we perform the following types of

classification tasks corresponding to disease phenotypes: (i) normal control vs. AD patients; (ii) normal control vs. MCI patients; (iii) MCI vs. AD patients. In addition to the above tasks, we also perform (iv) gender classification (males vs females); and (v) cognitive performance classification based on high and low levels of MMSE scores. MMSE is commonly used clinically for checking cognitive impairment with a low value indicative of a cognitive decline. We stratify individuals into high vs low MMSE categories, depending on whether their MMSE scores are above the 70th percentile or below the 30th percentile of the MMSE distribution. The corresponding sample sizes in the high and low MMSE categories were 280 and 249 respectively. For each classification task based on a given 2D slice, the data is randomly split into training and test splits in the ratio 70:30, and 10 such splits are considered. The results are averaged over these 10 replicates and reported in Table 9.

Results

Table 9 reports the misclassification rate and f1-score for slices 19 through 25. The proposed approach under both

Table 8 Summary of demographic variables and cognitive measurements under study

		NC (N=229)	AD (N=188)	MCI (N=401)	Overall (N = 818)
Age	Mean(SD)	75.86(5.02)	75.25(7.53)	74.74(7.35)	75.17(6.83)
	Median	75.6	75.65	75.10	75.45
	[Min, Max]	[59.90, 89.60]	[55.10, 90.90]	[54.40, 89.30]	[54.40, 90.90]
Yrs Ed^a	Mean(SD)	16.06(2.85)	14.65(3.13)	15.64(3.03)	15.53(3.05)
	Median	16	15	16	16
	[Min, Max]	[6, 20]	[4, 20]	[4, 20]	[4, 20]
Gender	Female	110(48.0%)	89(47.3%)	143(35.7%)	342(41.8%)
	Male	119(52.0%)	99(52.7%)	258(64.3%)	476(58.2%)
APOE4	0	168(73.4%)	64(34.0%)	186(46.4%)	418(51.1%)
	1	56(24.4%)	88(46.8%)	168(41.9%)	312(31.2%)
	2	5(2.2%)	36(19.2%)	47(11.7%)	88(10.7%)
MMSE	Mean(SD)	29.11(1.00)	23.27(2.03)	27.01(1.77)	26.74(2.67)
	Median	29	23	27	27
	[Min, Max]	[25, 30]	[18, 27]	[23, 30]	[18, 30]
log(ICV)	Mean(SD)	14.24(0.10)	14.24(0.12)	14.26(0.11)	14.25(0.11)
	Median	14.24	14.25	14.26	14.25
	[Min, Max]	[14.00, 14.53]	[13.92, 14.56]	[13.97, 14.56]	[13.92, 14.56]


^aYears of Education

Table 9 Classification on Gender, Disease phenotype, and High/Low MMSE scores with Demographic covariates

	BT-SVM		BT-LR		LR-lasso		L1norm-SVM	
	Mis. C. ^a	F1-score	Mis. C.	F1-score	Mis. C.	F1-score	Mis. C.	F1-score
Female vs. Male								
slice 19	0.288	0.764	0.317	0.728	0.337	0.750	0.422	0.636
slice 20	0.280	0.762	0.292	0.748	0.361	0.729	0.398	0.654
slice 21	0.296	0.752	0.3	0.756	0.357	0.75	0.447	0.595
slice 22	0.284	0.779	0.333	0.723	0.361	0.729	0.504	0.550
slice 23	0.256	0.789	0.325	0.733	0.373	0.657	0.390	0.684
slice 24	0.280	0.771	0.325	0.720	0.398	0.611	0.447	0.618
slice 25	0.256	0.801	0.345	0.719	0.369	0.674	0.455	0.591
NC. vs. AD								
slice 19	0.327	0.669	0.277	0.715	0.388	0.595	0.452	0.564
slice 20	0.301	0.712	0.325	0.682	0.341	0.644	0.365	0.671
slice 21	0.341	0.656	0.309	0.697	0.341	0.638	0.404	0.564
slice 22	0.278	0.724	0.325	0.655	0.333	0.655	0.444	0.582
slice 23	0.277	0.720	0.309	0.698	0.293	0.654	0.436	0.444
slice 24	0.293	0.689	0.333	0.681	0.325	0.601	0.341	0.626
slice 25	0.269	0.746	0.309	0.677	0.301	0.660	0.396	0.510
NC. vs MCI								
slice 19	0.280	0.8	0.312	0.768	0.344	0.732	0.414	0.686
slice 20	0.291	0.791	0.380	0.707	0.302	0.751	0.418	0.691
slice 21	0.338	0.761	0.365	0.718	0.333	0.731	0.440	0.688
slice 22	0.317	0.771	0.349	0.736	0.333	0.722	0.458	0.664
slice 23	0.322	0.776	0.349	0.766	0.365	0.696	0.455	0.586
slice 24	0.328	0.760	0.349	0.750	0.360	0.704	0.433	0.620
slice 25	0.291	0.792	0.354	0.729	0.333	0.720	0.465	0.582
AD vs. MCI								
slice 19	0.296	0.805	0.355	0.746	0.334	0.763	0.395	0.690
slice 20	0.282	0.816	0.305	0.784	0.325	0.577	0.420	0.495
slice 21	0.287	0.815	0.338	0.781	0.350	0.759	0.412	0.691
slice 22	0.282	0.810	0.322	0.753	0.322	0.778	0.356	0.770
slice 23	0.271	0.825	0.344	0.751	0.384	0.723	0.401	0.702
slice 24	0.282	0.812	0.338	0.758	0.350	0.747	0.446	0.663
slice 25	0.291	0.808	0.389	0.721	0.339	0.765	0.463	0.613
High vs Low MMSE scores								
slice 19	0.283	0.723	0.327	0.675	0.333	0.693	0.496	0.606
slice 20	0.295	0.715	0.358	0.655	0.333	0.674	0.471	0.534
slice 21	0.289	0.722	0.352	0.654	0.345	0.667	0.484	0.549
slice 22	0.295	0.710	0.339	0.686	0.377	0.634	0.471	0.460
slice 23	0.295	0.712	0.345	0.645	0.327	0.653	0.421	0.645
slice 24	0.289	0.720	0.333	0.675	0.321	0.622	0.459	0.587
slice 25	0.301	0.707	0.327	0.662	0.327	0.653	0.471	0.590

^aMisclassification Rate

Fig. 4 Top panel: Estimated effects of the brain cortical regions for classification task of AD vs. normal control subjects from BT-SVM, portrayed as a set of 2D brain slices overlaid with estimated points estimates of the model coefficients > 0.1 . Bottom panel: Estimated effects of the brain cortical regions for classification task of AD vs. normal control subjects from BT-LR, portrayed as a set of 2D brain slices overlaid with estimated points estimates of the model coefficients > 0.1



loss functions almost always results in better classification accuracy consistently across slices and classification tasks, compared to the penalized approaches. While both the two competing approaches have inferior performance compared to their Bayesian counterparts, the LR-Lasso generally performs slightly better than the L1norm-SVM. Among the three disease phenotype classification tasks, the highest accuracy is achieved for AD vs MCI classification under the BT-SVM approach, with f1-scores greater than 0.8 across all 2D slices. The proposed approaches are also able to perform well for the NC vs MCI classification task with F1-scores consistently greater than 0.75, while the classification performance for NC vs AD is slightly less impressive under all methods. We note that slice 23 provides the highest accuracy for both AD vs MCI and NC vs MCI classification tasks. With regards to the classification of non-imaging phenotypes, the gender classification performance is generally improved compared to cognitive performance classification, which is not surprising given the considerably lower training sample size for classifying the high versus low cohorts.

Clearly, there exist fluctuations in classification performances across the different 2-D slices. This is not surprising because cortical thickness contained in each 2-D slice varies depending on the sections of the slice. Further, different slices represent different brain regions that may have differential effects on classifying phenotypic classes. Overall,

these results demonstrate that our proposed Bayesian tensor approach offers consistent improvements for classifying Alzheimer's disease and subjects' demographic information, by leveraging the spatial information in the images and incorporating dimension reduction that potentially avoids overfitting. Moreover, the unsatisfactory performance of the competing methods clearly illustrates the perils of ignoring the spatial correlations in high-dimensional images in classification tasks. This is not surprising, given that Lasso-based approaches are known to be affected by multicollinearity. While it is possible to explore alternate approaches such as principal component regression (PCR), such approaches lead to loss of interpretability that is not desirable in imaging studies due to various reasons including the inability to perform feature selection.

Figure 4 shows brain maps that visualize the estimated coefficients of the brain cortical regions for the classification task of normal control vs. AD subjects based on BT-SVM and BT-LR, respectively. We up-sampled the 7 axial slices to full-dimension 3D images for visualization purposes. The 3D point estimate is then portrayed as a set of eight 2D brain slices, overlaid with the estimated effects of the brain cortical regions > 0.1 . The brain maps show regions with the strongest coefficient estimates that are directly responsible for differentiating the phenotypic class labels. Multiple Regions of Interests (ROIs) are identified under

the proposed method that are also known to be implicated by the progression of AD. Across the 7 axial slices used for the classification analysis, the proportion of important voxels with estimated coefficients greater than 0.1 under BT-SVM and BT-LR was reported to be over 10% of the right hippocampus, the left and right caudate, the right superior parietal gyrus, and the right putamen. Additionally, over 20% of the right insular cortex and right superior occipital gyrus regions report estimated coefficients > 0.1 under BT-SVM and BT-LR.

These brain regions are well-known to be associated with cortical atrophy and/or cognitive impairment due to healthy aging as well as AD. Hippocampal atrophy due to aging as well as AD is widely established in literature (Bettio et al., 2017; Salat et al., 2011), and hippocampal degeneration is known to be associated with cognitive impairment (Xiao et al., 2023). Further, both left and right caudate are known to exhibit volume loss due to AD and they are also associated with cognitive functioning as measured by MMSE (Madsen et al., 2010). Similarly, parietal regions are known to have differential atrophy patterns based on cognitive status for individuals with early dementia (Jacobs et al., 2011). Additionally, the putamen is known to have a strong reduction in volume for AD individuals that also correlated with reduced cognitive performance (de Jong et al., 2008). Further, recent findings suggest that amyloidosis in putamen is a valuable imaging marker for AD (Yang et al., 2023). Similarly, the insula has been shown to carry a considerable pathological burden that may affect behavioral traits (Bonthuis et al., 2005). Finally, significant cortical thickness and surface area atrophy was noticed in occipital lobes of the brain in AD individuals compared to normal controls and early MCI subjects (Yang et al., 2019).

Discussion

In this study, we proposed a Bayesian tensor classification approach based on high-dimensional neuroimaging data and scalar predictors, via data augmentation. The proposed approach essentially extends the literature on Bayesian tensor regression to classification problems that has the ability to perform inference and uncertainty quantification. The tensor structure is particularly suitable for neuroimaging data since it respects the spatial information of imaging voxels while keeping the number of parameters to be estimated at a manageable level via a PARAFAC decomposition. With two data augmentation schemes implemented, we demonstrated the superiority of the proposed method via simulation and data application. In particular, comprehensive simulation studies concretely illustrate the advantages under the proposed

approach, with each type of data augmentation having distinct advantages in classification, coefficient estimation and/or feature selection. We applied the proposed method to the ADNI dataset to classify disease phenotype, as well as gender and cognitive performance. Both data augmentation schemes showed consistently higher classification accuracy and improved feature selection than other penalized methods that vectorized the image. This showed the benefits of incorporating the spatial information in the image.

A potential limitation of the study is that we use 2D brain slices instead of using the whole 3D brain image. While using the full image may potentially result in improvements in accuracy, it typically requires higher rank tensors that may create computational bottlenecks due to a massive number of model parameters. In future work, we intend to propose more scalable versions of the proposed approach that can be used to incorporate high-dimensional 3D images for classification. We also propose to explore various prior choices for the tensor coefficients in order to enable stronger shrinkage and thicker tails, with a view to improving the feature selection performance. Another potential limitation is that we did not consider potential heterogeneity across samples in our modeling scheme. In future work, we plan to account for the heterogeneity across samples via a mixture of tensors approach that is particularly relevant for AD studies. While we plan to investigate the above issues in future research, we believe that the proposed approach in this article provides a valuable Bayesian classification tool based on imaging covariates that fills an important gap in literature.

Information Sharing Statement

The sample data from ADNI can be accessed at <https://adni.loni.usc.edu> by accepting the Data Use Agreement and submitting an online application form at <https://adni.loni.usc.edu/data-samples/access-data/>. The code for implementing the proposed approaches is available here: <https://github.com/rongke79/Bayesian-tensor-modeling>.

Author Contributions S.K. designed the research and conceived the analysis plan; R.L. performed model implementation and data analyses; S.K. and M.V. provided overall supervision and direction of the work. R.L., S.K. and M.V. wrote the paper.

Data Availability The sample data used in this study from ADNI can be accessed by accepting the Data Use Agreement and submitting an online application form at <https://adni.loni.usc.edu/data-samples/access-data/>. The code for implementing the proposed approaches is available here: <https://github.com/rongke79/Bayesian-tensor-modeling>.

Declarations

Competing Interests The authors declare no competing interests.

References

- Arbabshirani, M. R., Plis, S., Sui, J., & Calhoun, V. D. (2017). Single subject prediction of brain disorders in neuroimaging: Promises and pitfalls. *Neuroimage*, 145, 137–165.
- Armagan, A., Dunson, D. B., & Lee, J. (2013). Generalized double Pareto shrinkage. *Statistica Sinica*, 23(1), 119.
- Avants, B. B., Epstein, C. L., Grossman, M., & Gee, J. C. (2008). Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1), 26–41.
- Becker, N., Werft, W., Toedt, G., Lichter, P., & Benner, A. (2009). penalizedSVM: a R-package for feature selection SVM classification. *Bioinformatics*, 25(13), 1711–1712.
- Behler, A., Müller H. -P., Ludolph, A. C., Lulé, D., & Kassubek, J. (2022). A multivariate Bayesian classification algorithm for cerebral stage prediction by diffusion tensor imaging in amyotrophic lateral sclerosis. *NeuroImage: Clinical*, 35, 103094. <https://doi.org/10.1016/j.nicl.2022.103094>
- Ben Ahmed, O., Benois-Pineau, J., Allard, M., Ben Amar, C., Catheline, G., & Initiative, A. D. N. (2015). Classification of Alzheimer's disease subjects from MRI using hippocampal visual features. *Multi-media Tools and Applications*, 74, 1249–1266.
- Bettio, L. E., Rajendran, L., & Gil-Mohapel, J. (2017). The effects of aging in the hippocampus and cognitive decline. *Neuroscience & Biobehavioral Reviews*, 79, 66–86.
- Billio, M., Casarin, R., Iacopini, M., & Kaufmann, S. (2023). Bayesian dynamic tensor regression. *Journal of Business & Economic Statistics*, 41(2), 429–439.
- Bonthuis, D. J., Solodkin, A., & Van Hoesen, G. W. (2005). Pathology of the insular cortex in Alzheimer disease depends on cortical architecture. *Journal of Neuropathology & Experimental Neurology*, 64(10), 910–922.
- Bradley, P. S., & Mangasarian, O. L. (1998). Feature selection via concave minimization and support vector machines. In *ICML*, 98, 82–90.
- Cervantes, J., Garcia-Lamont, F., Rodríguez-Mazahua, L., & Lopez, A. (2020). A comprehensive survey on support vector machine classification: Applications, challenges and trends. *Neurocomputing*, 408, 189–215.
- Chouliaras, L., & O'Brien, J. T. (2023). The use of neuroimaging techniques in the early and differential diagnosis of dementia. *Molecular Psychiatry*, pages 1–14.
- de Jong, L. W., van der Hiele, K., Veer, I. M., Houwing, J., Westendorp, R., Bollen, E., de Bruin, P. W., Middelkoop, H., van Buchem, M. A., & van der Grond, J. (2008). Strongly reduced volumes of putamen and thalamus in Alzheimer's disease: an MRI study. *Brain*, 131(12), 3277–3285.
- Dedieu, A. (2019). Error bounds for sparse classifiers in high-dimensions. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 48–56. PMLR.
- Devika, S., Jeyaseelan, L., & Sebastian, G. (2016). Analysis of sparse data in logistic regression in medical research: A newer approach. *Journal of postgraduate medicine*, 62(1), 26. <https://doi.org/10.4103/0022-3859.173193>
- Doerken, S., Avalos, M., Lagarde, E., & Schumacher, M. (2019). Penalized logistic regression with low prevalence exposures beyond high dimensional settings. *PLoS One*, 14(5), e0217057. <https://doi.org/10.1371/journal.pone.0217057>
- Falahati, F., Westman, E., & Simmons, A. (2014). Multivariate data analysis and machine learning in Alzheimer's disease with a focus on structural magnetic resonance imaging. *Journal of Alzheimer's disease*, 41(3), 685–708.
- Fjell, A. M., Grydeland, H., Krogsrud, S. K., Amlie, I., Rohani, D. A., Ferschmann, L., Storsve, A. B., Tamnes, C. K., Sala-Llanch, R., Due-Tønnessen, P., et al. (2015). Development and aging of cortical thickness correspond to genetic organization patterns. *Proceedings of the National Academy of Sciences*, 112(50), 15462–15467.
- Frenzel, S., Wittfeld, K., Habes, M., Klinger-Koenig, J., Buelow, R., Voelzke, H., & Grabe, H. J. (2020). A biomarker for Alzheimer's disease based on patterns of regional brain atrophy. *Frontiers in psychiatry*, 10, 953.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1), 1.
- Geweke, J. (1991). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. Staff Report 148, Federal Reserve Bank of Minneapolis. <https://ideas.repec.org/p/fip/fedmsr/148.html>
- Griffis, J. C., Allendorfer, J. B., & Szaflarski, J. P. (2016). Voxel-based Gaussian naïve Bayes classification of ischemic stroke lesions in individual T1-weighted MRI scans. *Journal of neuroscience methods*, 257, 97–108.
- Guhaniyogi, R. (2020). *Bayesian methods for tensor regression* (pp. 1–18). Wiley StatsRef: Statistics Reference Online.
- Guhaniyogi, R., & Spencer, D. (2021). Bayesian tensor response regression with an application to brain activation studies. *Bayesian Analysis*, 16(4), 1221–1249.
- Guhaniyogi, R., Qamar, S., & Dunson, D. B. (2017). Bayesian tensor regression. *The Journal of Machine Learning Research*, 18(1), 2733–2763.
- Hahn, P. R., & Carvalho, C. M. (2015). Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective. *Journal of the American Statistical Association*, 110(509), 435–448.
- Jacobs, H. I., Van Boxtel, M. P., Uylings, H. B., Gronenschild, E. H., Verhey, F. R., & Jolles, J. (2011). Atrophy of the parietal lobe in preclinical dementia. *Brain and Cognition*, 75(2), 154–163.
- Kolda, T. G., & Bader, B. W. (2009). *Tensor decompositions and applications*. *SIAM review*, 51(3), 455–500.
- Kundu, S., Mallick, B. K., & Baladandayuthapan, V. (2019). Efficient Bayesian regularization for graphical model selection. *Bayesian Analysis*, 14(2), 449.
- Kundu, S., Reinhardt, A., Song, S., Han, J., Meadows, M. L., Crosson, B., Krishnamurthy, V. (2023). Bayesian longitudinal tensor response regression for modeling neuroplasticity. *Human Brain Mapping*.
- Lee, K.-J., Jones, G. L., Caffo, B. S., & Bassett, S. S. (2014). Spatial Bayesian variable selection models on functional magnetic resonance imaging time-series data. *Bayesian Analysis (Online)*, 9(3), 699.
- Lock, E. F. (2018). Tensor-on-tensor regression. *Journal of Computational and Graphical Statistics*, 27(3), 638–647.
- Ma, X., & Kundu, S. (2022). High-dimensional measurement error models for Lipschitz loss. Preprint at <https://arxiv.org/abs/2210.15008>
- Madsen, S. K., Ho, A. J., Hua, X., Saharan, P. S., Toga, A. W., Jack, C. R., Jr., Weiner, M. W., Thompson, P. M., Initiative, A. D. N., et al. (2010). 3d maps localize caudate nucleus atrophy in 400 Alzheimer's disease, mild cognitive impairment, and healthy elderly subjects. *Neurobiology of aging*, 31(8), 1312–1325.
- Morales, D. A., Vives-Gilabert, Y., Gómez-Ansón, B., Bengoetxea, E., Larrañaga, P., Bielza, C., Pagonabarraga, J., Kulisevsky, J., Corcuera-Solano, I., & Delfino, M. (2013). Predicting dementia development in Parkinson's disease using Bayesian network classifiers. *Psychiatry Research: Neuroimaging*, 213(2), 92–98.
- Nicholas, J. G. S., Polson, G., & Windle, J. (2013). Bayesian inference for logistic models using Pólya-Gamma latent variables. *Journal of the American Statistical Association*, 108(504), 1339–1349.
- Pan, Y., Mai, Q., & Zhang, X. (2018). Covariate-adjusted tensor classification in high dimensions. *Journal of the American statistical association*.

- Peng, P., Lekadir, K., Gooya, A., Shao, L., Petersen, S. E., & Frangi, A. F. (2016). A review of heart chamber segmentation for structural and functional analysis using cardiac magnetic resonance imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 29, 155–195.
- Plant, C., Teipel, S. J., Oswald, A., Böhm, C., Meindl, T., Mourao-Miranda, J., Bokde, A. W., Hampel, H., & Ewers, M. (2010). Automated detection of brain atrophy patterns based on MRI for the prediction of Alzheimer's disease. *Neuroimage*, 50(1), 162–174. <https://doi.org/10.1016/j.neuroimage.2009.11.046>
- Polson, N. G., & Scott, S. L. (2011). Data augmentation for support vector machines. *Bayesian Analysis*, 6, 1–24.
- Rathore, S., Habes, M., Iftikhar, M. A., Shacklett, A., & Davatzikos, C. (2017). A review on neuroimaging-based classification studies and associated feature extraction methods for Alzheimer's disease and its prodromal stages. *NeuroImage*, 155, 530–548.
- Salat, D. H., Chen, J. J., Van Der Kouwe, A., Greve, D. N., Fischl, B., & Rosas, H. D. (2011). Hippocampal degeneration is associated with temporal and limbic gray matter/white matter tissue contrast in Alzheimer's disease. *Neuroimage*, 54(3), 1795–1802.
- Sen, B., & Parhi, K. K. (2021). Predicting biological gender and intelligence from fMRI via dynamic functional connectivity. *IEEE Transactions on Biomedical Engineering*, 68(3), 815–825. <https://doi.org/10.1109/TBME.2020.3011363>
- Smith, M., & Fahrmeir, L. (2007). Spatial Bayesian variable selection with application to functional magnetic resonance imaging. *Journal of the American Statistical Association*, 102(478), 417–431.
- Sun, W., Chang, C. Zhao, Y., & Long, Q. (2018). Knowledge-guided Bayesian support vector machine for high-dimensional data with application to analysis of genomics data. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 1484–1493. IEEE.
- Tokdar, S. T., & Ghosh, J. K. (2007). Posterior consistency of logistic Gaussian process priors in density estimation. *Journal of statistical planning and inference*, 137(1), 34–42.
- Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., & Gee, J. C. (2010). N4ITK: improved N3 bias correction. *IEEE transactions on medical imaging*, 29(6), 1310–1320.
- Tustison, N. J., Cook, P. A., Klein, A., Song, G., Das, S. R., Duda, J. T., Kandel, B. M., van Strien, N., Stone, J. R., Gee, J. C., et al. (2014). Large-scale evaluation of ANTs and FreeSurfer cortical thickness measurements. *Neuroimage*, 99, 166–179.
- Tustison, N. J., Holbrook, A. J., Avants, B. B., Roberts, J. M., Cook, P. A., Reagh, Z. M., Duda, J. T., Stone, J. R., Gillen, D. L., Yassa, M. A., et al. (2019). Longitudinal mapping of cortical thickness measurements: An Alzheimer's Disease Neuroimaging Initiative-based evaluation study. *Journal of Alzheimer's Disease*, 71(1), 165–183.
- Weston, P. S., Nicholas, J. M., Lehmann, M., Ryan, N. S., Liang, Y., Macpherson, K., Modat, M., Rossor, M. N., Schott, J. M., Ourselin, S., et al. (2016). Presymptomatic cortical thinning in familial Alzheimer disease: A longitudinal MRI study. *Neurology*, 87(19), 2050–2057.
- Wu, Y., Kundu, S., Stevens, J. S., Fani, N., & Srivastava, A. (2022). Elastic shape analysis of brain structures for predictive modeling of PTSD. *Frontiers in Neuroscience*, 16, 954055. <https://doi.org/10.3389/fnins.2022.954055>
- Xiao, Y., Hu, Y., Huang, K., Initiative, A. D. N., et al. (2023). Atrophy of hippocampal subfields relates to memory decline during the pathological progression of Alzheimer's disease. *Frontiers in Aging Neuroscience*, 15.
- Yang, H., Xu, H., Li, Q., Jin, Y., Jiang, W., Wang, J., Wu, Y., Li, W., Yang, C., Li, X., et al. (2019). Study of brain morphology change in Alzheimer's disease and amnesic mild cognitive impairment compared with normal controls. *General psychiatry*, 32(2).
- Yang, Z., Cummings, J. L., Cordes, D., & Initiative, A. D. N. (2023). Amyloidosis at putamen predicts vulnerability to Alzheimer's disease. *Alzheimer's & Dementia*, 19, e079343.
- Zhou, H., Li, L., & Zhu, H. (2013). Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108(502), 540–552.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.