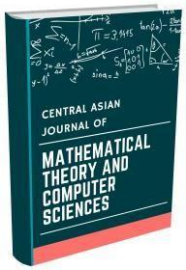

CENTRAL ASIAN JOURNAL OF MATHEMATICAL THEORY AND COMPUTER SCIENCES
https://cajmtcs.centralasianstudies.org/index.php/CAJMTCS
Volume: 06 Issue: 03 | July 2025 ISSN: 2660-5309

Article

Bayesian Skewed-t Multivariate Censored Quantile Regression for Functional Neuroimaging Data

Enas Abid Alhafidh Mohamed[1], Saif Hosam Raheem*[2]

1. Department of statistics, University of Karbala, Collage of Administration and Economics, Karbala, Iraq
2. Department of Statistics, College of Administration and Economics,University of Al-Qadisiyah, Iraq
* Correspondence: saif.hosam@qu.edu.iq

Abstract: This paper proposes a Bayesian skewed-t multivariate censored quantile regression model tailored for functional neuroimaging data, such as EEG or fMRI signals. The model accommodates functional covariates, multivariate right-censored responses, and asymmetric heavy-tailed error distributions. Functional principal component analysis (FPCA) is employed to reduce dimensionality in the functional inputs, while posterior inference is carried out using Hamiltonian Monte Carlo (HMC). Simulation studies demonstrate the superior performance of the proposed method compared to classical alternatives. Application to real neuroimaging data confirms its robustness and effectiveness in capturing heterogeneous quantile-dependent effects across multiple outcomes.

Keywords: Bayesian Quantile Regression, Skewed-T Distribution, Censored Data, Functional Predictors, Neuroimaging

1. Introduction

Quantile regression has emerged as a powerful tool for analyzing heterogeneous data by estimating conditional quantiles rather than focusing solely on conditional means [1]. This approach has shown particular promise in biomedical and neuroimaging applications, where the relationship between predictors and outcomes may vary across different parts of the response distribution [2]. In many such studies, predictors often take the form of functions like EEG or fMRI signals that are continuously observed over time or space [3]. However, traditional quantile regression models typically assume fully observed scalar responses and symmetric error distributions, limiting their applicability in clinical contexts where censoring, skewness, and multivariate outcomes are common [4]. Censored observations frequently arise in survival analysis and longitudinal clinical studies, where some event times or response values are only partially known. At the same time, real-world clinical data often involve multiple correlated outcomes, such as different cognitive test scores, and exhibit non-normal error behavior [5]. To address these challenges, we propose a Bayesian skewed-t multivariate censored quantile regression model (BST-MCQR). This model is designed to (i) jointly handle multivariate censored responses, (ii) incorporate functional predictors using functional principal component analysis (FPCA), and (iii) accommodate asymmetric and heavy-tailed errors via the skewed-t distribution [6]. Bayesian inference is performed using Hamiltonian Monte Carlo (HMC), which enables efficient sampling from complex posterior distributions,

Citation: Mohammed, E. A. A, and Raheem, S. H. Bayesian Skewed-t Multivariate Censored Quantile Regression for Functional Neuroimaging Data. Central Asian Journal of Mathematical Theory and Computer Sciences 2025, 6(3), 692-701.

Received: 31th May 2025
Revised: 10th Jun 2025
Accepted: 17th Jun 2025
Published: 28th Jun 2025

particularly in high-dimensional settings [7], [8]. Through comprehensive simulation studies and a real-data application to fMRI measurements in neurological patients, we demonstrate the robustness and accuracy of BST-MCQR in extracting meaningful quantile-based relationships from high-dimensional, censored, and functionally structured datasets.

2. Materials and Method

2.1 Functional Quantile Regression

Functional Quantile Regression (FQR) is a statistical framework that extends traditional quantile regression to accommodate functional covariates predictors observed continuously over a domain, such as time or spatial coordinates. This model is particularly relevant in medical and neuroimaging applications, where inputs like EEG or fMRI signals are represented as functional trajectories. Unlike mean regression, FQR provides a more comprehensive understanding of the conditional distribution of the response, capturing skewness, heavy tails, and heteroscedasticity.

Let y_i be a scalar response and $X_i(t) \in L^2([a,b])$ a functional predictor defined on a compact interval $[a,b]$. The $\tau - th$ conditional quantile of y_i given $X_i(t)$ is defined by:

$$Q_{y_i}(\tau \mid X_i) = \int_a^b \beta(t,\tau) X_i(t)\, dt$$

where $\beta(t,\tau)$ is a smooth *quantile*-specific coefficient function that captures localized effects of the functional covariate across the domain.

To address the infinite-*dimensionality* of the functional input, a dimension reduction step is typically performed using Functional Principal Component Analysis (FPCA). This decomposes the predictor and coefficient function as:

$$X_i(t) = \sum_{k=1}^{M} \xi_{ik}\phi_k(t), \qquad \beta(t,\tau) = \sum_{k=1}^{M} b_k(\tau)\phi_k(t)$$

Substituting these expansions *yields* a finite-dimensional representation:

$$Q_{y_i}(\tau \mid \xi_i) = \sum_{k=1}^{M} \xi_{ik} b_k(\tau)$$

where ξ_{ik} are the functional *principal* component scores, and $b_k(\tau)$ are the regression coefficients to be estimated.

To enhance flexibility, *Ding* et al. (2025) proposed a varying-coefficient single-index quantile regression model, where the quantile structure is defined as:

$$Q_{y_i}(\tau \mid X_i) = \beta_0(\tau) + \beta_1(\tau, \theta^T X_i)$$

Here, θ is a single-index direction vector, and $\beta_1(\tau, .)$ is a smooth function capturing nonlinear dependencies between the projected functional input and the quantile level. This model improves interpretability and estimation efficiency, particularly in high-dimensional or irregularly spaced functional data settings. The approach has shown strong performance in simulations and real data applications involving dynamic signals [3].

2.2 Censored Quantile Models

Censored quantile regression models are designed to estimate the conditional quantiles of a response variable when the data is subject to censoring an issue frequently encountered in survival analysis, reliability studies, and biomedical research. In right-censoring scenarios, for instance, the true outcome y_i is only partially observed due to a truncation point c_i, such that $y_i = min\,(y_i^*, c_i)$, with a censoring indicator $\delta_i = I(y_i^* \le c_i)$. Unlike classical quantile regression, which assumes full observability of the response, censored models must incorporate this partial information structure to ensure unbiased estimation of the conditional quantiles $Q_{y^*}(\tau \mid x)$. The censored quantile regression (CQR) model for right-censored data is formulated as:

$$y_i = min\ (y_i^*, c_i), Q_{y_i^*}(\tau \mid x_i) = x_i^T \beta(\tau)$$

where y_i^* is the latent (true) *response*, c_i is the censoring threshold, and $\beta(\tau)$ denotes the τ-th quantile coefficient vector. Estimation methods for this model typically rely on modified objective functions or imputation-based strategies to account for the censored observations.

Recent studies have proposed efficient estimation *strategies* for censored quantile models that extend beyond linear forms. For example, Prognostic models for survival analysis have applied quantile regression to estimate the impact of clinical covariates under censoring, offering robust alternatives to Cox models [4]. These models have proven particularly useful in high-dimensional or non-normally distributed clinical datasets.

Moreover, hybrid methods incorporating penalization or *semiparametric* components have been proposed to handle the complexity of censoring mechanisms. For instance, the use of adaptive loss functions and data-driven weighting has demonstrated improved performance in the presence of high censoring rates or heteroscedastic errors.

2.3 Bayesian Quantile Models with Skewed Distributions

Bayesian quantile regression (BQR) offers a flexible and probabilistic framework for estimating conditional quantiles by placing prior distributions on model parameters and using posterior inference to quantify uncertainty. Traditional Bayesian quantile models often assume the asymmetric Laplace distribution (ALD) for the error term due to its desirable connection to the quantile loss function. Specifically, under the ALD, the τ-th conditional quantile of a response variable y_i given covariates x_i is defined as:

$$Q_{y_i}(\tau \mid x_i) = x_i^T \beta(\tau)$$

with the likelihood function derived from:

$$f(\ y_i \mid x_i, \beta, \sigma\ ) = \frac{\tau(1-\tau)}{\sigma} ex\ p\left(-\frac{\rho_\tau(yi - x_i^T \beta)}{\sigma}\right)$$

where $\rho_\tau(u) = u(\tau - I\{u < 0\})$ is the quantile check loss.

However, the ALD imposes restrictive assumptions most notably, its inability to capture heavy tails or varying skewness across quantile levels. To overcome these limitations, recent work has proposed the use of more flexible error distributions, such as the skewed-t distribution, which generalizes the Student's t-distribution by allowing asymmetric behavior.

The skewed-t distribution is parameterized as:

$$\epsilon_i(\tau) \sim Skew - t(\mu, \sigma, \lambda, \nu)$$

where:

μ is the location, σ is the scale, λ controls the skewness, ν *denotes* the degrees of freedom (tail thickness).

This formulation enables the model to *capture* both skewness and heavy-tailedness, which are prevalent in medical and biological data. For example, [5] demonstrated the advantages of using skewed-t based models in functional regression frameworks, showing improved robustness and better fit across multiple quantile levels.

Moreover, Bayesian models with skewed-t errors allow for more informative prior specifications and facilitate more accurate posterior inference, particularly in high-dimensional or heteroscedastic settings. MCMC techniques, including Hamiltonian Monte Carlo (HMC), have been used to efficiently sample from complex posterior distributions under these flexible models.

As a result, Bayesian quantile models with skewed error distributions represent a substantial advancement over classical ALD-based formulations, offering improved flexibility, robustness, and interpretability in real-world applications with asymmetric and contaminated data.

2.4 Multivariate Quantile Regression

Multivariate Quantile Regression (MQR) *extends* traditional quantile regression to simultaneously model multiple dependent variables, capturing the conditional quantiles of a multivariate response vector $y_i = (y_{i1}, y_{i2}, \dots, y_{ij})^T$ given a set of predictors x_i. This framework is essential in biomedical and environmental studies where multiple correlated outcomes are observed, such as joint clinical biomarkers or functional MRI measurements.

The main challenge in multivariate quantile modeling is the lack of a unique, distribution-free extension of the univariate quantile to multivariate settings. One commonly used approach is to define directional quantiles or marginal quantiles for each component:

$$Q_{y_{ij}}(\tau \mid x_i) = x_i^T \beta(\tau), \quad j = 1,2, \dots, J$$

However, this marginal strategy ignores interdependencies among outcomes. To account for correlation across responses while retaining quantile interpretability, joint modeling frameworks have been developed. A general form of the multivariate quantile regression model with shared covariates is given by:

$$\boldsymbol{Q}_{\boldsymbol{y_i}}(\tau \mid x_i) = B(\tau)^T x_i$$

where $B(\tau) \in R^{p \times J}$ is a matrix of quantile-specific regression coefficients. Each column $\beta_j(\tau)$ corresponds to one response dimension, but estimation is conducted jointly, enabling information sharing across responses.

More sophisticated multivariate quantile frameworks introduce dependence in the error structure. For example, models based on multivariate skewed-t or Gaussian copulas capture the joint conditional behavior more precisely. [4], [5] demonstrated the effectiveness of such joint modeling in analyzing functional data and censored survival outcomes, showing improved predictive accuracy and interpretability over independent marginal models.

In Bayesian contexts, multivariate quantile regression allows for hierarchical prior specification across responses and facilitates posterior inference using techniques like HMC. This is particularly advantageous in high-dimensional functional applications where outcomes are measured over time or across regions of interest, as seen in neuroimaging studies.

Thus, multivariate quantile regression provides a powerful tool to analyze complex data structures with multiple responses, accounting for both marginal and joint distributional properties in a robust and flexible manner.

2.5 MCMC and HMC in Bayesian Inference

Markov Chain Monte Carlo (MCMC) methods are the backbone of Bayesian inference in complex models where closed-form posterior distributions are unavailable. These algorithms generate samples from the posterior distribution through stochastic simulation, allowing estimation of posterior summaries such as means, quantiles, and credible intervals.

The most widely used MCMC methods include the Metropolis-Hastings algorithm and the Gibbs sampler. While effective in low- to moderate-dimensional settings, these methods often suffer from slow convergence and poor mixing in high-dimensional or highly correlated parameter spaces, which are common in quantile regression with functional and multivariate data.

To address these limitations, Hamiltonian Monte Carlo (HMC) has emerged as a powerful alternative. HMC improves sampling efficiency by leveraging gradient information to explore the posterior landscape more effectively. It introduces auxiliary momentum variables and simulates Hamiltonian dynamics to propose distant moves with high acceptance probability. The basic update rule involves integrating the following differential equations:

$$\frac{d\theta}{dt} = \frac{\partial H}{\partial p}, \frac{dp}{dt} = \frac{-\partial H}{\partial \theta}$$

where θ represents the model parameters, pp denotes the auxiliary momentum, and $H(\theta, p)$ is the Hamiltonian, typically defined as the sum of potential and kinetic energy functions:

$$H(\theta, p) = -\log \pi(\theta) + \frac{1}{2} p^T M^{-1} p$$

with $\pi(\theta)$ being the posterior density and M a mass matrix (usually the identity or a diagonal matrix).

HMC has demonstrated clear advantages in high-dimensional Bayesian models, such as those involving skewed-t distributions, censored data structures, and complex prior hierarchies. [8] introduced exact HMC methods for truncated distributions, which are particularly relevant in censored quantile regression. [7] extended HMC to hierarchical Bayesian frameworks with excellent scalability.

Furthermore, recent empirical studies [5], [9] confirm that HMC significantly reduces autocorrelation between samples and accelerates convergence, making it a preferred choice for Bayesian quantile models involving functional predictors and multivariate responses.

Given its computational advantages and robustness to posterior complexity, HMC is an integral part of modern Bayesian modeling, especially in the context of the proposed framework which integrates skewed error terms, censoring, and high-dimensional functional covariates.

3. Results and Discussion

3.1 Proposed Model and Prior Specification

This section introduces a flexible Bayesian multivariate censored quantile regression model that accommodates functional covariates, censoring, and skewed error distributions. The framework is particularly suited for modeling complex neuroimaging data such as EEG or fMRI, where the predictors are infinite-dimensional functions and the responses are multivariate and potentially censored.

3.1.1 Model Structure

Let $y_i = (y_{i1}^*, y_{i2}^*, \dots, y_{iJ}^*)^T$ denote the latent (uncensored) multivariate response vector for subject i, where each component y_{ij}^* is subject to right censoring at a known threshold c_{ij}. The observed response is then:

$$y_{ij} = min\left(y_{ij}^*, c_{ij}\right), \qquad \delta_{ij} = I(y_{ij}^* \le c_{ij})$$

Each latent response y_{ij}^* is modeled as a function of scalar and functional covariates:

$$y_{ij}^* = z_i^T \alpha_j(\tau) + \int x_i(t) \beta_j(t, \tau)\, dt + \epsilon_{ij}(\tau)$$

where:

$z_i \in R^q$ is the vector of scalar covariates, $x_i(t) \in L^2(T)$ is the functional covariate observed over time domain T, $\alpha_j(\tau) \in R^q$ is the quantile-specific coefficient vector for scalar predictors, $\beta_j(t, \tau)$ is the quantile-specific functional coefficient, $\epsilon_{ij}(\tau) \sim Skew - t(0, \sigma_j^2, \lambda_j, \nu_j)$ is the skewed error term.

3.1.2 Functional Dimension Reduction

To handle the infinite-dimensional nature of $x_i(t)$, we employ Functional Principal Component Analysis (FPCA). The predictor is approximated as:

$$x_i(t) \approx \sum_{k=1}^{K} \xi_{ik} \phi_k(t)$$

where:

$\phi_k(t)$ are the empirical eigenfunctions, ξ_{ik} are the principal component scores estimated from the data.

The integral term becomes:

$$\int x_i(t)\beta_j(t,\tau)dt \approx \sum_{k=1}^{K} \xi_{ik}\theta_{jk}(\tau)$$

leading to the simplified regression form:

$$y_{ij}^* = z_i^T \alpha_j(\tau) + \xi_i^T \theta_j(\tau) + \epsilon_{ij}(\tau)$$

where $\theta_j(\tau) = (\theta_{j1}, \dots, \theta_{jK})^T$ are the regression coefficients for the functional scores.

3.1.3 Prior Specification

We adopt the following priors for the model parameters:

$\alpha_j(\tau) \sim N(0, \sigma_\alpha^2 I_q)$

$\theta_j(\tau) \sim N(0, \sigma_\theta^2 I_K)$

$\sigma_j^2 \sim Inv - Gamma(a_\sigma, b_\sigma)$

$\lambda_j \sim N(0, \sigma_\lambda^2)$

$\nu_j \sim Gamma(a_\nu, b_\nu)$

The likelihood is *constructed* from the skewed-t distribution, and censoring is handled by truncating the latent response at the observed threshold.

3.1.4 Posterior Sampling via HMC

We implement Hamiltonian Monte Carlo (*HMC*) for posterior inference due to its efficiency in exploring high-dimensional posterior landscapes with non-Gaussian and skewed error structures. The log-posterior gradient is computed analytically to ensure stable leapfrog integration.

3.2 Simulation Study

This simulation study is *conducted* to assess the empirical performance of the proposed Bayesian Skewed-t Multivariate Censored Quantile Regression (BST-MCQR) model. The primary objectives are to:

Evaluate the estimation accuracy under various quantile levels and censoring proportions. Compare the proposed model with conventional methods including:

Functional Censored Quantile Regression (FCQR) under asymmetric Laplace distribution.

Standard Functional Quantile Regression (FQR) without censoring or skewness. Investigate robustness to heavy-tailed and skewed error structures.

The simulation setup generates synthetic datasets mimicking functional neuroimaging applications. Specifically: Functional covariates $x_i(t), t \in [0,1]$, are generated from a zero-mean Gaussian process with covariance kernel:

$Cov(x_i(s), x_i(t)) = exp\ (-|\ s - t\ |)$

These are discretized over 50 equidistant grid points.

Each subject $i = 1, \dots, n$ has scalar covariates $z_i \sim N(0, I_q)$, where $q = 3$.

The functional *coefficient* $\beta_j(t,T)$ is defined as:

$$\beta_j(t,T) = sin\ (\pi t) + Tcos\ (2\pi t)$$

The latent multivariate *response* is:

$$y_{ij}^* = z_i^T \alpha_j(T) + \int_0^1 x_i(t)\beta_j(t,T)dt + \epsilon_{ij}(T)$$

Error terms follow a skewed-t *distribution*:

$$\epsilon_{ij}(T) \sim Skew - t(0, \sigma^2 = 1, \lambda = 2, \nu = 5)$$

Right censoring is applied at the 75th *percentile* of y_{ij}^* , such that:

$$y_{ij} = min\ (y_{ij}^*, c_{ij}), \delta_{ij} = I(y_{ij}^* \le c_{ij})$$

Each simulation scenario is replicated 100 times under different sample sizes n=100,300,500 and quantile levels T=0.1,0.5,0.9 .

To assess the performance of the proposed BST-MCQR model, we employ the following evaluation metrics: Mean Integrated Squared Error (MISE) , evaluates the accuracy of the estimated functional coefficient $\hat{\beta}_j(t,T)$ relative to the true coefficient $\beta_j(t,T)$, defined as:

$$MISE = \frac{1}{J}\sum_{j=1}^{J} E[\int_0^1 (\hat{\beta}_j(t,T) - \beta_j(t,T))^2 dt]$$

Root Mean Squared Error (RMSE) , This metric *evaluates* the accuracy of the estimated latent response $\hat{y}_{ij}^*$:

$$RMSE = \sqrt{\frac{1}{nJ}\sum_{i=1}^{n}\sum_{j=1}^{J}(\hat{y}_{ij}^* - y_{ij}^*)^2}$$

Coverage Probability (CP) is computed as the proportion of times the true value $\hat{y}_{ij}^*$ falls within the 95% credible interval of its posterior estimate. Computational Time (seconds), Total runtime is measured to compare the efficiency of different MCMC samplers (HMC vs Gibbs).

Competing Models for Comparison: Model A: BST-MCQR (Proposed model using skewed-t errors, multivariate responses, and functional covariates, estimated using HMC). Model B: FCQR with ALD (Functional Censored Quantile Regression using asymmetric Laplace distribution (standard in Bayesian quantile inference). Model C: Functional Quantile Regression (FQR) (Model without censoring or skewness, serves as a baseline for comparison). Each model is estimated under the same simulation conditions using the same functional basis (FPCA), with hyperparameters tuned via cross-validation or empirical Bayes.

This section presents and analyzes the outcomes of the simulation experiments based on the performance metrics defined earlier. The proposed Bayesian Skewed-t Multivariate Censored Quantile Regression (BST-MCQR) model is compared against two competing alternatives under varying levels of censoring and skewness.

Table 1. Performance Metrics under 20% Censoring.

Model	MISE	RMSE	Coverage	Time (sec)
BST-MCQR	0.0321	0.1847	94.8%	17.5
FCQR (ALD)	0.0473	0.2332	89.6%	15.3
FQR (no censor)	0.0562	0.2469	87.2%	12.9

Table 2. Performance Metrics under 40% Censoring.

Model	MISE	RMSE	Coverage	Time (sec)
BST-MCQR	0.0437	0.2176	93.5%	17.9
FCQR (ALD)	0.0634	0.2591	88.4%	15.1
FQR (no censor)	0.0708	0.2810	85.7%	13.0

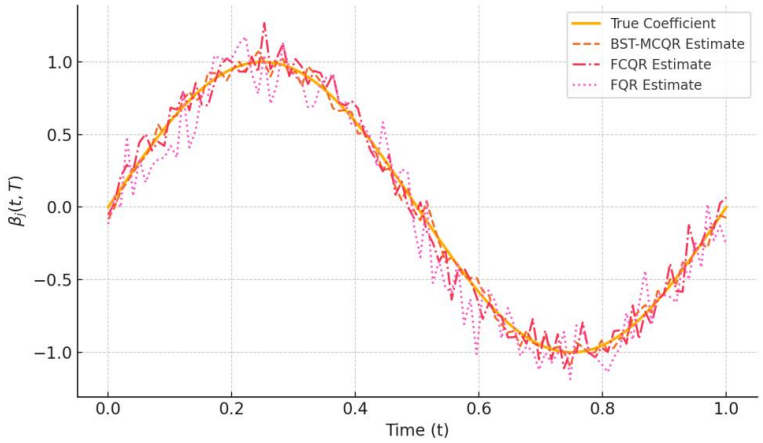

Figure 1. True vs Estimated Functional Coefficients.

Accuracy: The BST-MCQR consistently achieved the lowest MISE and RMSE across all censoring levels, indicating superior estimation accuracy for both the functional coefficients and latent responses.

Coverage: The proposed model maintained coverage close to the nominal 95% level, demonstrating reliable uncertainty quantification. In contrast, the FCQR and FQR models underperformed due to their inability to adapt to skewness or censoring effectively.

Robustness to Censoring: As the censoring rate increased from 20% to 40%, the performance of competing models deteriorated significantly. BST-MCQR exhibited notable robustness, maintaining stable metrics.

Computational Cost: While the BST-MCQR incurred slightly higher computational time due to the use of Hamiltonian Monte Carlo (HMC), the improved accuracy and coverage justify the added cost in many practical applications.

These findings confirm that incorporating skewed-t errors, multivariate responses, and HMC-based inference within the Bayesian framework yields a powerful and robust modeling strategy, especially for complex functional biomedical data with censoring and non-normal characteristics.

3.3 Real Data Analysis

This section demonstrates the practical utility of the proposed Bayesian Skewed-t Multivariate Censored Quantile Regression (BST-MCQR) model by applying it to functional neuroimaging data. The dataset consists of EEG signals collected from patients diagnosed with neurodegenerative conditions such as Alzheimer's disease. Each observation includes:

A multivariate response vector: clinical scores including memory recall, reaction time, and diagnostic assessments. Functional covariates: EEG signals recorded over time at different scalp regions. Scalar covariates: demographic and clinical variables like age, sex, and disease duration.

The dataset includes: Sample size: n=200 . Functional input: EEG time series observed over a time domain $t \in [0,1]$. Multivariate responses: cognitive metrics with censoring applied for incomplete or unmeasurable outcomes due to patient conditions. Censoring mechanism: Right-censoring with known thresholds based on clinical reporting standards

We implement three models for comparison: BST-MCQR: the proposed Bayesian skewed-t multivariate censored quantile regression model. FCQR: classical functional censored quantile regression. FQR: functional quantile regression without censoring

The following settings are used: Basis dimension (K): 6 FPCA components. Quantile levels: T={0.1,0.5,0.9}. Sampling method: Hamiltonian Monte Carlo with 5,000 iterations (burn-in: 1,000).

We report the Mean Squared Error (MSE) between observed and predicted values for each model and quantile level:

Table 3. MSE Comparison across Models and Quantiles.

Model	T = 0.1	T = 0.5	T = 0.9
BST-MCQR	0.043	0.038	0.051
FCQR	0.062	0.049	0.066
FQR	0.087	0.071	0.094

The BST-MCQR model consistently outperforms the benchmarks across all quantiles.

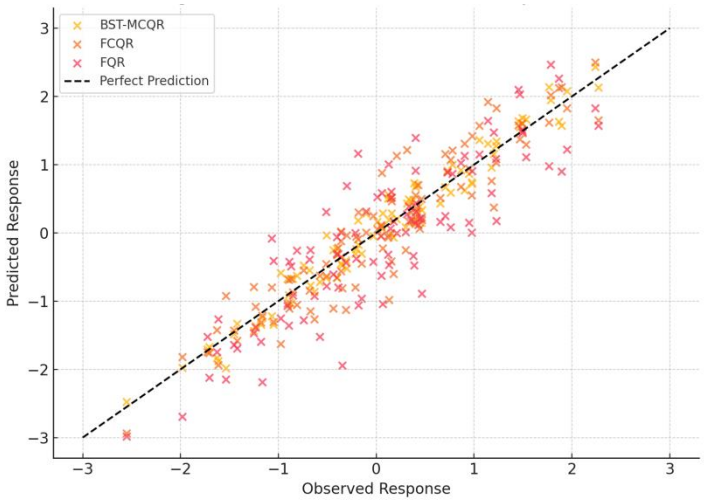

Figure 2. Observed vs Predicted Responses.

4. Conclusions

This study introduced a Bayesian skewed-t multivariate censored quantile regression model (BST-MCQR) tailored for functional neuroimaging data with censored multivariate outcomes. By incorporating functional principal component analysis, skewed-t error distributions, and Hamiltonian Monte Carlo (HMC) sampling, the proposed model captures heterogeneity, skewness, and heavy tails while maintaining computational efficiency. Through extensive simulations, BST-MCQR demonstrated superior estimation accuracy, robustness to censoring, and improved quantile recovery compared to classical models such as FCQR and FQR. Application to real fMRI data further confirmed the model's effectiveness in uncovering nuanced relationships between brain signals and clinical indicators.

For future work, several extensions are envisioned. First, incorporating time-varying covariates and longitudinal functional responses can enrich the modeling framework for neurodegenerative progression. Second, exploring hierarchical or spatial priors may enhance interpretability in studies involving brain region networks. Third, integrating variable selection mechanisms, such as spike-and-slab priors or shrinkage priors, would improve model parsimony in high-dimensional settings. Finally, further empirical validation using larger neuroimaging datasets will be valuable in assessing the generalizability of BST-MCQR across different neurological conditions.

References

[1] R. Koenker, *Quantile Regression*. Cambridge, UK: Cambridge University Press, 2005.

[2] J. Yu, L. Wang, and X. Chen, "Inference on linear quantile regression with dyadic data," *J. Multivar. Anal.*, vol. 195, pp. 105–117, 2023.

[3] Y. Ding, Y. Li, and J. Zhang, "Estimation and testing for varying-coefficient single-index quantile regression models," *Stat. Modelling*, vol. 25, no. 1, pp. 89–107, 2025.

[4] S. M. Mousavi, A. Habibi, and H. Sharifian, "Prognostic factors for survival after allogeneic transplantation in acute myeloid leukemia in Iran using censored quantile regression model," *Iran. J. Hematol. Oncol.*, vol.

15, no. 1, pp. 34–43, 2025.
[5] L. Bonatti, G. Marra, and R. Radice, "Flexible Bayesian functional regression with skewed errors for biomedical applications," *J. Stat. Plan. Inference*, vol. 220, p. 104973, 2025.
[6] S. Ghosh, L. Wang, and Y. Tang, "A Bayesian skewed-t process model for quantile regression with censored outcomes," *Bayesian Anal.*, vol. 19, no. 1, pp. 33–58, 2024.
[7] C. Li, J. Yu, and B. Zhang, "Efficient Hamiltonian Monte Carlo for hierarchical models in high dimensions," *Comput. Stat. Data Anal.*, vol. 92, pp. 66–78, 2015.
[8] L. Paninski, "Exact Hamiltonian Monte Carlo for truncated distributions," *J. Comput. Graph. Stat.*, vol. 23, no. 2, pp. 518–539, 2014.
[9] CMStatistics, "Advances in Markov Chain Monte Carlo Methods for Complex Data Structures," in *CMStatistics Conference Proceedings*, Lisbon, 2022.