Introduction to binary and categorical outcomes with BART

> Rodney Sparapani
> Medical College of Wisconsin
> Copyright (c) 2023 Rodney Sparapani

June 7 \& 8: BART workshop<br>Medical College of Wisconsin, Milwaukee campus

Funding for this research was provided, in part, by the Advancing Healthier Wisconsin Research and Education Program under awards 9520277 and 9520364.

## Outline

Sparapani, Spanbauer \& McCulloch 2021
Journal of Statistical Software

- Motivation: chronic spine pain and obesity
- Dichotomous outcomes with probit BART
- Dichotomous outcomes with logistic BART
- Geweke convergence diagnostics for binary BART
- Categorical outcomes with logistic BART
- Categorical outcomes with probit BART


## Motivation: chronic spine pain and obesity

- Hypothesis a: obesity is a risk factor for chronic lower back/buttock pain
- Hypothesis b: obesity is NOT a risk factor for chronic neck pain
- Data available from the National Health and Nutrition Examination Survey (NHANES) 2009-2010 Arthritis Questionnaire
- 5106 subjects were surveyed
- Demographics: age and gender
- Anthropometrics available: weight (kg), height (cm), body mass index ( $\mathrm{kg} / \mathrm{m}^{2}$ ), waist circumference ( cm )
- Sampling weights to estimate for the US as a whole
- For obesity quantified by BMI, see demo/nhanes.pbart1.R and demo/nhanes.pbart2. R in the BART R package
- For obesity quantified by waist circumference, see demo/nhanes.pbart. R in the BART3 R package


## Probit BART for binary outcomes

Probit regression with latent variables: Albert \& Chib 1993 JASA

$$
\begin{aligned}
y_{i} \mid p_{i} & \stackrel{\text { ind }}{\sim} \mathbf{B}\left(p_{i}\right) \\
p_{i} \mid f & =\Phi\left(\mu+f\left(x_{i}\right)\right) \text { where } f \stackrel{\text { prior }}{\sim} \text { BART and } \mu=\Phi^{-1}(\bar{y}) \\
z_{i} \mid y_{i}, f & \sim \mathrm{~N}\left(\mu+f\left(x_{i}\right), \mathbf{1}\right) \begin{cases}\mathbf{I}(-\infty, 0) & \text { if } y_{i}=\mathbf{0} \\
\mathbf{I}(0, \infty) & \text { if } y_{i}=1\end{cases} \\
f \mid z_{i}, y_{i} & \stackrel{d}{=} f \mid z_{i} \\
{[y \mid f] } & =\prod_{i=1}^{N} p_{i}^{y_{i}}\left(1-p_{i}\right)^{1-y_{i}} \quad \text { Likelihood }
\end{aligned}
$$

Continuous BART with unit variance, $\boldsymbol{\sigma}^{\mathbf{2}}=\mathbf{1}$, and $z_{i}$ are the data

## Friedman's partial dependence function for probit BART

Friedman 2001 AnnStat

$$
\begin{aligned}
& p(x)=p\left(x_{S}, x_{C}\right) \quad \text { BART function where } x=\left[x_{S}, x_{C}\right] \\
& p\left(x_{S}\right)=\mathbf{E}_{x_{C}}\left[p\left(x_{S}, x_{C}\right) \mid x_{S}\right] \\
& \approx N^{-1} \sum_{i} p\left(x_{S}, x_{i C}\right) \equiv N^{-1} \sum_{i} \Phi\left(\mu+f\left(x_{S}, x_{i C}\right)\right) \\
& p_{m}\left(x_{S}\right) \equiv N^{-1} \sum_{i} p_{m}\left(x_{S}, x_{i C}\right) \\
& \hat{p}\left(x_{S}\right) \equiv M^{-1} \sum_{m} p_{m}\left(x_{S}\right)
\end{aligned}
$$

## gbart and mc. gbart input and output

```
post <- gbart(x.train, y.train, type="pbart", ...,
    ndpost=M, keepevery=10) or
post <- mc.gbart(x.train, y.train, type="pbart", ...,
    ndpost=M, keepevery=10, mc.cores=2, seed=99)
```

Input matrices: x.train and, optionally, x.test: $\boldsymbol{x}_{\boldsymbol{i}}$

$$
\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{N}
\end{array}\right]
$$

Output object, post, of type pbart (essentially a list)
Matrices: post\$prob.train and, optionally, post\$prob.test:

$$
\begin{gathered}
\hat{p}_{i m}=\Phi\left(\mu+f_{m}\left(x_{i}\right)\right) \\
{\left[\begin{array}{ccc}
\hat{p}_{11} & \cdots & \hat{p}_{N 1} \\
\vdots & \vdots & \vdots \\
\hat{p}_{1 M} & \cdots & \hat{p}_{N M}
\end{array}\right]}
\end{gathered}
$$

## predict. pbart input and output

```
pred <- predict(post, x.test, mc.cores=1, ...)
```

Input matrices: x.test: $\boldsymbol{x}_{\boldsymbol{i}}$

$$
\left[\begin{array}{c}
x_{1} \\
x_{2} \\
\vdots \\
x_{Q}
\end{array}\right]
$$

Output list with prob.test: $\hat{p}_{i m}=\Phi\left(\mu+f_{m}\left(x_{i}\right)\right)$

$$
\left[\begin{array}{ccc}
\hat{p}_{11} & \cdots & \hat{p}_{Q 1} \\
\vdots & \vdots & \vdots \\
\hat{p}_{1 M} & \cdots & \hat{p}_{Q M}
\end{array}\right]
$$

## Demo: chronic spine pain and obesity

- Hypothesis a: obesity is a risk factor for chronic lower back/buttock pain
- Hypothesis b: obesity is NOT a risk factor for chronic neck pain
- system.file('demo/nhanes.pbart1.R', package='BART')
- system.file('demo/nhanes.pbart2.R', package='BART')


## Friedman's partial dependence function:

 Probability of chronic pain vs. BMI

## Friedman's partial dependence function:

Probability of chronic pain vs. BMI


## Logistic BART for binary outcomes

Logistic regression with latent variables
Devroye 1986 Non-uniform random variate generation
Holmes \& Held 1993 Bayesian Analysis
Gramacy \& Polson 2012 Bayesian Analysis

$$
\begin{aligned}
y_{i} \mid p_{i} & \stackrel{\text { ind }}{\sim} \mathbf{B}\left(p_{i}\right) \\
p_{i} \mid f & =\Phi\left(\mu+f\left(x_{i}\right)\right) \text { where } f \stackrel{\text { prior }}{\sim} \text { BART }(\mu) \text { and } \mu=\Phi^{-1}(\bar{y}) \\
z_{i} \mid y_{i}, f, \sigma_{i} & \sim \mathrm{~N}\left(\mu+f\left(x_{i}\right), \sigma_{i}^{2}\right) \begin{cases}\mathrm{I}(-\infty, 0) & \text { if } y_{i}=\mathbf{0} \\
\mathrm{I}(0, \infty) & \text { if } y_{i}=1\end{cases} \\
\sigma_{i}^{2} & =\mathbf{4} \psi_{i}^{2} \text { where } \psi_{i} \sim \text { Kolmogorov-Smirnov (see Devroye) }
\end{aligned}
$$

Continuous BART with heteroskedastic variance and $z_{i}$ is the data

## Geweke convergence diagnostics for binary BART

Hastings 1970 Biometrika, Silverman 1986 Chapman and Hall

$$
\begin{aligned}
& \hat{\theta}_{M}=M^{-1} \sum_{m=1}^{M} \theta_{m} \\
& \sigma_{\hat{\theta}}^{2}=\lim _{M \rightarrow \infty} \mathrm{~V}\left[\hat{\theta}_{M}\right]
\end{aligned}
$$

Bayesian estimator

Asymptotic variance

Suppose $\boldsymbol{\theta}_{\boldsymbol{m}}$ is an ARMA $(\boldsymbol{p}, \boldsymbol{q})$

$$
\begin{aligned}
\gamma(w) & =(2 \pi)^{-1} \sum_{m=-\infty}^{\infty} \mathrm{V}\left[\theta_{0}, \theta_{m}\right] \mathrm{e}^{\mathrm{i} m w} \\
\hat{\sigma}_{\hat{\theta}}^{2} & =\hat{\gamma}^{2}(0)
\end{aligned}
$$

Variance estimator

## Geweke convergence diagnostics for binary BART

Geweke 1992 Bayesian Statistics

- Divide your chain into two segments: $\boldsymbol{A}$ and $\boldsymbol{B}$
- $m \in A=\left\{1, \ldots, M_{A}\right\}$ where $M_{A}=a M$
- $m \in B=\left\{M-M_{B}+1, \ldots, M\right\}$ where $M_{B}=b M$
- $a+b<1$, Geweke suggests $a=0.1$ and $b=0.5$

$$
\begin{array}{rlr}
\hat{\theta}_{A}=M_{A}^{-1} \sum_{m \in A} \theta_{m} & \hat{\theta}_{B}=M_{B}^{-1} \\
\hat{\sigma}_{\hat{\theta}_{A}}^{2}=\hat{\gamma}_{m \in A}^{2}(0) & \hat{\sigma}_{\hat{\theta}_{B}}^{2}=\hat{\gamma}_{m \in B}^{2} \\
z & =\frac{\sqrt{M}\left(\hat{\theta}_{A}-\hat{\theta}_{B}\right)}{\sqrt{a^{-1}} \hat{\sigma}_{\hat{\theta}_{A}}^{2}+b^{-1} \hat{\sigma}_{\hat{\theta}_{B}}^{2}} & \sim \mathbf{N}(0,1)
\end{array}
$$

## Geweke convergence diagnostics for binary BART

- We have a $z_{i}$ corresponding to each $\theta_{i}=\boldsymbol{h}\left(\boldsymbol{\mu}+\boldsymbol{f}\left(\boldsymbol{x}_{i}\right)\right)$
- In the BART R package, we created the gewekediag function which was adapted from the coda $R$ package Plummer, Best et al. 2006
system.file('demo/geweke.pbart2.R', package='BART')


## Geweke convergence diagnostics for binary BART: simulated data scenario

system.file('demo/geweke.pbart2.R', package='BART')

$$
\begin{aligned}
N & =\mathbf{2 0 0}, \mathbf{1 0 0 0}, \mathbf{1 0 0 0 0} \quad \text { sample sizes } \\
K & =\mathbf{5 0} \quad \text { number of covariates } \\
f\left(x_{i}\right) & =-\mathbf{1 . 5}+\sin \left(\pi x_{1 i} x_{2 i}\right)+\mathbf{2}\left(x_{3 i}-\mathbf{0 . 5}\right)^{2}+x_{4}+\mathbf{0 . 5} x_{5} \\
z_{i} & \sim \mathbf{N}\left(f\left(x_{i}\right), \mathbf{1}\right) \\
y_{i} & =\mathbf{I}\left(z_{i}>\mathbf{0}\right)
\end{aligned}
$$

## Geweke convergence diagnostics for binary BART:

$N=200$



m
N:200, k:50

$\mathrm{N}: 200, \mathrm{k}: 50$

## Geweke convergence diagnostics for binary BART:

$N=1000$





1000, k:50

## Geweke convergence diagnostics for binary BART:

$N=10000$


## Multinomial BART with logit link

 mbart2 function for a larger number of categoriesSparapani, Spanbauer and McCulloch 2021 JSS

- $y=\left[\begin{array}{c}y_{1} \\ \vdots \\ y_{k}\end{array}\right] \sim \operatorname{Multinomial}(n, p)$ where $p=\left[\begin{array}{c}p_{1} \\ \vdots \\ p_{k}\end{array}\right]$
- $n=\sum_{j} y_{j}$ and $\sum_{j} p_{j}=1$
- If $\boldsymbol{n}=\mathbf{1}$, computing Multinomial BART is facilitated by modeling the binary outcomes with $\boldsymbol{k}$ logistic BARTs
$y_{i j} \sim \mathbf{B}\left(p_{i j}\right)$ where $f_{j} \stackrel{\text { prior }}{\sim}$ BART $\left(\mu_{j}\right)$ and $p_{i j} \propto F\left(\mu_{j}+f_{j}\left(x_{i}\right)\right)$
- And then combining the inference as follows
$\boldsymbol{p}_{i j}=\frac{\exp \left(\mu_{j}+f_{j}\left(x_{i}\right)\right)}{\sum_{j^{\prime}} \exp \left(\mu_{j}+f_{j}\left(x_{i}\right)\right)}$ (but each fit is slow and we need $k$ of them)
- This would work with the probit link (and it would be much faster), but there is no theoretical basis for combining probits in this way
- Or another alternative (that also doesn't follow from theory)
$-\tilde{p}_{i j}=\frac{\Phi\left(\mu_{j}+f_{j}\left(x_{i}\right)\right)}{\sum_{j^{\prime}} \Phi\left(\mu_{j}+f_{j}\left(x_{i}\right)\right)}$


## Multinomial BART with probit link

 mbart function for a smaller number of categories Sparapani, Spanbauer and McCulloch 2021 JSS- If $\boldsymbol{n}=\mathbf{1}$, fit a sequence of binary probit models (this bears some resemblance to continuation-ratio logits)
- assume $\boldsymbol{k}$ categories where each are represented by mutually exclusive binary indicators: $y_{i 1}, \ldots, y_{i k}$
- the probability of these outcomes, $\boldsymbol{p}_{i j}$, where $\boldsymbol{j}=\mathbf{1}, \ldots, \boldsymbol{k}$

$$
\begin{aligned}
p_{i 1} & =\mathbf{P}\left[y_{i 1}=1\right] \\
p_{i 2} & =\mathbf{P}\left[y_{i 2}=1 \mid y_{i 1}=0\right] \\
p_{i 3} & =\mathbf{P}\left[y_{i 3}=\mathbf{1} \mid y_{i 1}=y_{i 2}=0\right] \\
\vdots & \\
p_{i, k-1} & =\mathbf{P}\left[y_{i, k-1}=\mathbf{1} \mid y_{i 1}=\cdots=y_{i, k-2}=0\right] \\
p_{i k} & =\mathbf{P}\left[y_{i, k-1}=\mathbf{0} \mid y_{i 1}=\cdots=y_{i, k-2}=0\right]
\end{aligned}
$$

Notice that $\boldsymbol{p}_{\boldsymbol{i k}}=\mathbf{1}-\boldsymbol{p}_{i, k-1}$ so we can specify the $\boldsymbol{k}$ conditional probabilities via $\boldsymbol{k}-\mathbf{1}$ parameters

## Multinomial BART with probit link

 mbart function for a smaller number of categories- these conditional probabilities are, by construction, defined for subsets of subjects: let $S_{1}=\{\mathbf{1}, \ldots, N\}$ and $S_{j}=\left\{i: y_{i 1}=\cdots=y_{i, j-1}=0\right\}$ where $j=2, \ldots, k-1$
- the unconditional probability of these outcomes, $\pi_{i j}$, can be defined in terms of the conditional probablities and their complements, $q_{i j}=\mathbf{1}-p_{i j}$, for all subjects

$$
\begin{aligned}
\pi_{i 1} & =\mathbf{P}\left[y_{i 1}=1\right]=p_{i 1} \\
\pi_{i 2} & =\mathbf{P}\left[y_{i 2}=1\right]=p_{i 2} q_{i 1} \\
\pi_{i 3} & =\mathbf{P}\left[y_{i 3}=1\right]=p_{i 3} q_{i 2} q_{i 1} \\
\vdots & \\
\pi_{i, k-1} & =\mathbf{P}\left[y_{i, k-1}=1\right]=p_{i, k-1} q_{i, k-2} \cdots q_{i 1} \\
\pi_{i k} & =\mathbf{P}\left[y_{i k}=1\right]=q_{i, k-1} q_{i, k-2} \cdots q_{i 1}
\end{aligned}
$$

N.B. the rules of probability ensure that $\sum_{j=1}^{k} \pi_{i j}=\mathbf{1}$

## Multinomial BART with probit link <br> Alligator food choice: demo/alligator.R

- 219 alligators were taken by hunters in 1985 from 4 Florida lakes
- From 1 to 4 meters long, their stomachs were removed for study
- Each gator's primary food choice was determined 5 categories: bird, fish, invertebrate, reptile or other
- Covariates: lake, sex, and size (small vs. large)

