

over, when the response data is very noisy (i.e., low signal-to-noise ratio), `tgpr` can be expected to partition heavily under the `bprior="bflat"` prior. In such cases, one of the other proper priors like the full hierarchical `bprior="b0"` or `bprior="bmzt"` might be preferred.

3.5 Friedman data

This Friedman data set is the first one of a suite that was used to illustrate MARS (Multivariate Adaptive Regression Splines) [10]. There are 10 covariates in the data ($\mathbf{x} = \{x_1, x_2, \dots, x_{10}\}$). The function that describes the responses (Z), observed with standard Normal noise, has mean

$$E(Z|\mathbf{x}) = \mu = 10 \sin(\pi x_1 x_2) + 20(x_3 - 0.5)^2 + 10x_4 + 5x_5, \quad (18)$$

but depends only on $\{x_1, \dots, x_5\}$, thus combining nonlinear, linear, and irrelevant effects. Comparisons are made on this data to results provided for several other models in recent literature. Chipman et al. [5] used this data to compare their treed LM algorithm to four other methods of varying parameterization: linear regression, greedy tree, MARS, and neural networks. The statistic they use for comparison is root mean-square error (RMSE)

$$\text{MSE} = \sum_{i=1}^n (\mu_i - \hat{z}_i)^2 / n \quad \text{RMSE} = \sqrt{\text{MSE}}$$

where \hat{z}_i is the model-predicted response for input \mathbf{x}_i . The \mathbf{x} 's are randomly distributed on the unit interval.

Input data, responses, and predictive locations of size $N = 200$ and $N' = 1000$, respectively, can be obtained by a function included in the `tgpr` package.

```
> f <- friedman.1.data(200)
> ff <- friedman.1.data(1000)
> X <- f[, 1:10]
> Z <- f$Y
> XX <- ff[, 1:10]
```

This example compares Bayesian treed LMs with Bayesian GP LLM (not treed), following the RMSE experiments of Chipman et al. It helps to scale the responses so that they have a mean of zero and a range of one. First, fit the Bayesian treed LM, and obtain the RMSE.

```
> fr.btlm <- btlm(X = X, Z = Z, XX = XX, tree = c(0.95,
+      2), pred.n = FALSE, m0r1 = TRUE, verb = 0)
> fr.btlm.mse <- sqrt(mean((fr.btlm$ZZ.mean - ff$Ytrue)^2))
> fr.btlm.mse
```

```
[1] 1.939446
```

Next, fit the GP LLM, and obtain its RMSE.

```

> fr.bgp1lm <- bgp1lm(X = X, Z = Z, XX = XX, pred.n = FALSE,
+   m0r1 = TRUE, verb = 0)
> fr.bgp1lm.mse <- sqrt(mean((fr.bgp1lm$ZZ.mean - ff$Ytrue)^2))
> fr.bgp1lm.mse

[1] 0.4241515

```

So, the GP LLM is 4.573 times better than Bayesian treed LM on this data, in terms of RMSE (in terms of MSE the GP LLM is 2.138 times better).

Parameter traces need to be gathered in order to judge the ability of the GP LLM model to identify linear and irrelevant effects.

```

> XX1 <- matrix(rep(0, 10), nrow = 1)
> fr.bgp1lm.tr <- bgp1lm(X = X, Z = Z, XX = XX1, pred.n = FALSE,
+   trace = TRUE, verb = 0)

```

Notice that the `m0r1=TRUE` has been omitted so that the β estimates provided below will be on the original scale. A summary of the parameter traces show that the Markov chain had the following (average) configuration for the booleans.

```

> apply(fr.bgp1lm.tr$trace$XX[[1]][, 27:36], 2, mean)

b1 b2 b3 b4 b5 b6 b7 b8 b9 b10
1  1  1  0  0  0  0  0  0  0

```

Therefore the GP LLM model correctly identified that only the first three input variables interact only linearly with the response. This agrees with dimension-wise estimate of the total area of the input domain under the LLM (out of a total of 10 input variables).

```

> mean(fr.bgp1lm.tr$trace$linarea$ba)

[1] 7

```

A similar summary of the parameter traces for β shows that the GP LLM correctly identified the linear regression coefficients associated with the fourth and fifth input covariates (from (18))

```

> summary(fr.bgp1lm.tr$trace$XX[[1]][, 9:10])

      beta4      beta5
Min.   : 8.623  Min.   :4.309
1st Qu.: 9.370  1st Qu.:5.176
Median : 9.564  Median :5.376
Mean    : 9.550  Mean    :5.375
3rd Qu.: 9.735  3rd Qu.:5.582
Max.    :10.431  Max.    :6.313

```

and that the rest are much closer to zero.

```

> apply(fr.bgp1lm.tr$trace$XX[[1]][, 11:15], 2, mean)

      beta6      beta7      beta8      beta9      beta10
-0.23968561  0.37046946  0.13081722 -0.07842566  0.11911203

```