

Frequentist Approximation of the Bayesian Posterior Inclusion Probability by Stochastic Subsampling

Vi Quang Ly[†], Ernest Fokoué^{*,‡}

[†]*M&T Bank
Buffalo, New York, USA
e-mail: viqly@aol.com*

[‡]*School of Mathematical Sciences
Rochester Institute of Technology, Rochester, New York, USA
e-mail: epfeqa@rit.edu*

Abstract: This paper uses stochastic subsampling of the dataset to provide a frequentist approximation to what is known in the Bayesian framework as the posterior inclusion probability (PIP). The distinct merit of this contribution lies in the fact that it makes it easier for typically non-Bayesian-minded practitioners, of which there are many, to relate to the way the Bayesian paradigm allows the computation of the nicely interpretable variable importance. Despite its computationally intensive nature, due to the need to fitting a very large number of models, the proposed approach is readily applicable to both classification and regression tasks, and can be done in comparatively competitive computational times thanks to the availability of parallel computing facilities through cloud and cluster computing. Finally, the scheme proposed is very general and can therefore be easily adapted to all kinds of statistical prediction tasks. Application of the proposed method to some very famous benchmark datasets shows that it mimics the Bayesian counterpart quite well in the important context of variable selection.

AMS 2000 subject classifications: Primary 62H30; secondary 62H25.

Keywords and phrases: Generalized Linear Model, Bayesian Variable Selection, Posterior Inclusion Probability, Stochastic Subsampling, Binary Classification, Prediction, Logistic, Logit, Probit, Cauchit, Spam data.

1. Introduction

Modern statistical machine learning is replete with thousands of studies where the main statistical task revolves around estimations and predictions based on the traditional generalized linear model (GLM) given by

$$g(\mathbb{E}[Y|X, \beta, \alpha]) = \alpha \mathbf{1}_n + \mathbf{X}\beta, \quad (1.1)$$

where $g(\cdot)$ is the so-called link function, $\beta = (\beta_1, \dots, \beta_p)^\top$, $Y = (y_1, \dots, y_n)^\top$, the design matrix $\mathbf{X}$ is an $n \times p$ matrix, and $\mathbf{1}_n = (1, 1, \dots, 1)^\top$ is a $n \times 1$ dimensional vector of 1's. We shall refer to (1.1) as the *full model*. We assume

*Corresponding author

that many of the β_j 's are essentially zero, so that the intrinsic rank of the design matrix $\mathbf{X}$ is a number $q \in \mathbb{N}$ with $q \ll p$. Many data mining problems do exhibit such a characteristic of rank deficiency, mainly because variables are typically picked up as they are available, and therefore will turn out to be either noise variable (no relationship with the response) or redundant variables. Let $\mathbf{I}_n$ denote the $n \times n$ identity matrix. A basic result in GLM analysis shows that when $\mathbf{X}$ is rank deficient, and $g(\cdot)$ is the identity function, and the density of $\mathbf{Y}$ is gaussian, ie, $(\mathbf{Y}|\mathbf{X}; \boldsymbol{\beta}, \alpha) \sim N_n(\alpha \mathbf{1}_n + \mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}_n)$, then the ordinary least squares estimator

$$\hat{\boldsymbol{\beta}}^{(\text{OLS})} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y} \quad (1.2)$$

of $\boldsymbol{\beta}$ will tend to exhibit a high (inflated) variance, thereby corrupting all predictions and inferences with the computed model. It is therefore crucial to determine (if possible) the intrinsic model that generated the data, i.e. the model made up of only the most significant and non redundant variables. For many decades, both frequentist and Bayesian statisticians have contributed substantially to this topic of *variable selection*. In elementary statistical regression analysis courses, the method of choice for variable selection has been overwhelmingly frequentist with *stepwise regression heuristic* occupying a prominent place, and *best subsets selection* occasionally used whenever possible. While a heuristic like stepwise regression does provide a workable approach to variable selection, it is not a principled method, and does have the extra limitation of not providing any measure of variable importance. In recent years, both Bayesians and non-Bayesians have developed new methods for handling some of the most formidable variable selection tasks, many of which arose from the statistical learning and data mining community.

2. Bayesian Approach to Variable Selection

The vast majority of Bayesian contributions to variable selection of late have concentrated on the use of conjugate prior, with the typical choice of prior on $\boldsymbol{\beta}$ being a Gaussian prior of the form

$$(\boldsymbol{\beta}|\sigma^2, \mathbf{W}) \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{W}^{-1}), \quad (2.1)$$

where $\mathbf{W}$ is the prior precision matrix. In other words, we have

$$p(\boldsymbol{\beta}|\sigma^2, \mathbf{W}) = \frac{1}{\sqrt{((2\pi)\sigma^2)^p |\mathbf{W}|}} \exp \left\{ -\frac{1}{2\sigma^2} \boldsymbol{\beta}^\top \mathbf{W} \boldsymbol{\beta} \right\}. \quad (2.2)$$

Of course, the use of a zero mean prior expresses the assumption of many insignificant coefficients. However, even more important is the use of a vector of indicator variables that ultimately provides a mechanism (device) for performing variable selection. One of the key building blocks of the Bayesian variable selection machinery is the use of a vector of indicator variables. With the p original predictor variables, there are $2^p - 1$ non empty models corresponding each

to a subset of the provided variables. We shall use a vector $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots, \gamma_p)^\top$ to denote the index of a given model, with each γ_j being an indicator of the variable's presence in the model under consideration, namely

$$\gamma_j = \begin{cases} 1 & \text{If variable } X_j \text{ appears in the model} \\ 0 & \text{Otherwise} \end{cases} \quad (2.3)$$

Clearly, $\boldsymbol{\gamma} = (1, \dots, 1)^\top$ corresponds to the *full model* M_f , while $\boldsymbol{\gamma} = (0, \dots, 0)^\top$ corresponds to the empty model also referred to as the *null model*, and given by

$$M_n : \quad g(\mathbb{E}[\mathbf{Y}|\alpha]) = \alpha \mathbf{1}_n. \quad (2.4)$$

Equipped with this index, $p_\gamma = \sum_{j=1}^p \gamma_j$ is the number of predictor variables in model M_γ , and $\boldsymbol{\beta}_\gamma$ is the subset of $\boldsymbol{\beta}$ made up of only the β_j 's picked up by $\boldsymbol{\gamma}$. Finally, $\mathbf{X}_\gamma$ is the submatrix of $\mathbf{X}$ whose columns are only those p_γ columns of $\mathbf{X}$ picked up by $\boldsymbol{\gamma}$, so that $\mathbf{X}_\gamma$ is really an $n \times p_\gamma$ matrix, and the corresponding model M_γ is given by

$$M_\gamma : \quad g(\mathbb{E}[\mathbf{Y}|\mathbf{X}_\gamma, \boldsymbol{\beta}_\gamma, \alpha]) = \alpha \mathbf{1}_n + \mathbf{X}_\gamma \boldsymbol{\beta}_\gamma. \quad (2.5)$$

For the normal linear model, we have $[\mathbf{Y}|\alpha, \boldsymbol{\beta}_\gamma, \sigma^2, M_\gamma] \sim N_n(\alpha \mathbf{1}_n + \mathbf{X}_\gamma \boldsymbol{\beta}_\gamma, \sigma^2 \mathbf{I}_n)$, which means that

$$p(\mathbf{Y}|\boldsymbol{\theta}_\gamma, M_\gamma) = \frac{1}{\sqrt{((2\pi)\sigma^2)^n}} \exp \left\{ -\frac{1}{2}(\mathbf{y} - \alpha \mathbf{1}_n - \mathbf{X}_\gamma \boldsymbol{\beta}_\gamma)^\top (\mathbf{y} - \alpha \mathbf{1}_n - \mathbf{X}_\gamma \boldsymbol{\beta}_\gamma) \right\}$$

where $\boldsymbol{\theta}_\gamma = \{\alpha, \boldsymbol{\beta}_\gamma, \sigma^2\}$. When it comes to Bayesian variable selection, arguably the most crucial ingredient is the posterior density of a given model, given by

$$p(M_\gamma|\mathbf{y}) = \frac{p(\mathbf{y}|M_\gamma)p(M_\gamma)}{\sum_{\boldsymbol{\gamma} \in \boldsymbol{\Gamma}} p(\mathbf{y}|M_\gamma)p(M_\gamma)}$$

where $\boldsymbol{\Gamma} = \{0, 1\}^p$ and $p(\mathbf{y}|M_\gamma)$ is the marginal density of the data, also referred to as the marginal likelihood of model M_γ , and defined by

$$p(\mathbf{y}|M_\gamma) = \int_{\Theta} p(\mathbf{y}|\boldsymbol{\theta}_\gamma, M_\gamma) p(\boldsymbol{\theta}_\gamma|M_\gamma) d\boldsymbol{\theta}_\gamma.$$

In some special cases, it is possible to derive closed-form (analytical) expressions for $p(\mathbf{y}|M_\gamma)$, but in general, it must be approximated using a variety of schemes. The posterior probability $p(M_\gamma|\mathbf{y})$ of model M_γ , plays a central role in Bayesian learning.

$$p(\mathbf{z}|\mathbf{y}) = \sum_{\boldsymbol{\gamma} \in \boldsymbol{\Gamma}} p(\mathbf{z}|M_\gamma, \mathbf{y}) p(M_\gamma|\mathbf{y})$$

and also

$$\mathbb{E}(\mathbf{z}|\mathbf{y}) = \sum_{\boldsymbol{\gamma} \in \boldsymbol{\Gamma}} \mathbb{E}(\mathbf{z}|M_\gamma, \mathbf{y}) p(M_\gamma|\mathbf{y})$$

Among Bayesian statisticians, there are those who suggest that when it comes to model selection, one must choose the model with the highest posterior density model, i.e.,

$$\gamma_{\text{HPM}} = \underset{\gamma \in \mathbf{P}}{\operatorname{argmax}} \{p(M_\gamma | \mathbf{y})\}$$

[1] have suggested selecting instead the so-called median probability model (MPM) given γ_{MPM} , such that

$$[\gamma_j]_{\text{MPM}} = \begin{cases} 1 & \text{if } \pi_j \equiv \Pr[\gamma_j = 1 | \mathbf{y}] \geq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases} \quad (2.6)$$

where the posterior inclusion probability is given by

$$\text{PIP}_j = \pi_j = \Pr[\gamma_j = 1 | \mathbf{y}] = \sum_{\gamma \in \mathbf{P}} 1(\gamma_j = 1) p(M_\gamma | \mathbf{y}), \quad (2.7)$$

In practice, the estimated posterior inclusion probability is given by

$$\widehat{\text{PIP}}_j = \hat{\pi}_j = \sum_{\gamma: \gamma_j=1} p(M_\gamma | \mathbf{y}), \quad (2.8)$$

which means that the importance of a variable is measured in terms of its relative frequency of appearance in models. In Equation (2.8), it is crucial to be able to compute the posterior density of a given model. Empirically, this means that it is able to generate, at least a representative subset (sample) of all the models, and then to compute estimates of the posterior density from them, [2].

3. Frequentist Approximation of the Bayesian PIP

It turns out that the theoretical derivation of the Bayesian Posterior Inclusion Probability defined in (2.6) is technically very complicated. In fact, the estimation indicated in (2.8) often requires sophisticated Markov Chain Monte Carlo (MCMC) simulations. Fortunately, from its definition, it turns out that the Bayesian PIP can be approximated quite well by frequentist stochastic subsampling, an alternative that is both viable, desirable and feasible thanks to the availability of computing resources. Besides, we demonstrate and argue later that the advantage of this approximation lies in the fact that any practitioner who understands GLM can readily implement the idea of this paper and gain great benefits. Our frequentist's alternative to the Bayesian PIP proceeds by building M different models based on random subsamples of the data. For each one of the M models, the corresponding indicator vector is estimated by simply identifying the significant variables based on their P-values or T-values, or (expensively) on heuristic like stepwise regression when the computing resources

allow it. After the M models are built, one should consider a matrix of the form

$$\hat{\mathbf{\Gamma}} = \begin{bmatrix} \hat{\gamma}_{11} & \hat{\gamma}_{12} & \cdots & \hat{\gamma}_{1j} & \cdots & \hat{\gamma}_{1p} \\ \hat{\gamma}_{21} & \hat{\gamma}_{22} & \cdots & \hat{\gamma}_{2j} & \cdots & \hat{\gamma}_{2p} \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \hat{\gamma}_{m1} & \hat{\gamma}_{m2} & \cdots & \hat{\gamma}_{mj} & \cdots & \hat{\gamma}_{mp} \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ \hat{\gamma}_{M1} & \hat{\gamma}_{M2} & \cdots & \hat{\gamma}_{Mj} & \cdots & \hat{\gamma}_{Mp} \end{bmatrix}, \quad (3.1)$$

where $\hat{\gamma}_{mj}$ is set to 1 if variable X_j appeared in model m or was significant in model m , and zero otherwise. In other words, $\hat{\gamma}_{mj} \in \{0, 1\}$. Once the matrix $\hat{\mathbf{\Gamma}} \in \{0, 1\}^{M \times p}$ is built, the posterior inclusion probability $\text{PIP}_j = \pi_j$ of variable X_j can be approximated using

$$\hat{\pi}_j = \frac{1}{M} \sum_{m=1}^M \hat{\gamma}_{mj}.$$

Finally, if $\hat{\boldsymbol{\gamma}}^{\text{FREQ}} = (\hat{\gamma}_1^{\text{FREQ}}, \dots, \hat{\gamma}_p^{\text{FREQ}})$ denotes the indicator vector, our final approximating model, where

$$\hat{\gamma}_j^{\text{FREQ}} = 1 \text{ if } \hat{\pi}_j > 0.5.$$


4. Computational demonstrations

4.1. Computational demonstration on pattern recognition

According to Microsoft, spam is a term used to classify unwanted email. Spam may contain viruses or other malicious programs that can harm a computer. Furthermore, spam may be used as scams to acquire vital personal information such as credit card accounts, bank accounts, social security numbers, etc. Spam filters have been developed as preventative measures to protect the end user from ever opening these emails. Text categorization is one of the several techniques used to create spam filters. A number of terms are identified as indicators of spam/non-spam from a training set of emails. In the simplest spam filters, the frequencies of these terms in an email are determined and used to flag emails as spam/non-spam, [4]. The purpose of this subsection is threefold:

1. Build generalized linear models with a variety of link functions in the binary (Bernoulli) family, and use the built model to accurately and precisely classify emails as spam or nonspam.
2. Present a frequentist alternative to Bayesian Posterior Inclusion Probability (PIP) for variable selection.
3. Compare logistic regression classification accuracy with accuracies of newer machine learning algorithms.

The figure below shows the linear pairwise correlations between all of the explanatory variables. Based on the correlation plot, there is potential for multicollinearity to affect regression results. Due to the multicollinearity, variable selection will be required during the model building phase. Even with variable selection, there are still some questions that remain unanswered. What is the optimal model size to achieve model complexity-testing accuracy tradeoff? How much confidence should be placed on the variables identified by variable selection as significant variables? Furthermore, how does one characterize the importance of a variable's contribution to the model? We will provide some insight into these questions through a frequentist approach to variable selection.



6

The frequentist's alternative approach to PIP will now be applied to the Spam-base dataset. The entire dataset consisted of 4,601 observations. For step 1, 70% of the observations (3,221 observations) were randomly sampled from the entire dataset to form the training set. The remaining 30% of the observations (1,380 observations) formed the test set. This process was repeated 100 times to form 100 replicates of 70/30 training/test split. Since the main goal will be binary classification of spam / non-spam, the logistic regression model was selected for step 2. A model was built for each of the 100 training replicates using logistic regression with Logit link function on all 57 explanatory variables. The 100 models were then applied on their corresponding test sets to calculate the out-of-sample accuracies. The left-most figure below is a comparative boxplot between the training and test set accuracies of the Logit link function on all 57 explanatory variables. The same process was repeated using the Probit (center plot below) and Cauchit (right plot below) link functions. Based on the plots and number summaries below, the Cauchit link function provided the best family of models with respect to training and test set accuracies.




FIG 2. Comparative Boxplots of Training/Testing Accuracies - Logit (Left); Probit (Center); Cauchit (Right)

TABLE 1
Accuracies of Training Sets for Different Link Functions

Model	Min	1st Qu	Median	Mean	3rd Qu	Max
Logit	83.64%	92.92%	93.14%	92.83%	92.83%	93.32%
Probit	81.50%	88.66%	92.50%	90.50%	92.90%	93.73%
Cauchit	69.95%	94.65%	94.88%	94.28%	95.16%	95.87%

TABLE 2
Accuracies of Test Sets for Different Link Functions

Link function	Min	1st Qu	Median	Mean	3rd Qu	Max
Logit	80.87%	92.39%	92.97%	92.69%	93.48%	94.42%
Probit	81.38%	88.73%	92.17%	90.43%	92.90%	94.28%
Cauchit	69.86%	93.99%	94.42%	93.84%	94.78%	95.65%

With our link function in hand, we will now approach the daunting task of variable selection in step 3. Stepwise regression via Akaike Information Criterion (AIC) was applied to each of the 100 replicates. This process was repeated using stepwise regression via Bayesian Information Criterion (BIC), forward selection via BIC, and the full model. The following figure provides boxplots for the number of variables deemed significant in a replicate. For example, in AIC stepwise regression method, 1 of the 100 replicates deemed 30 of the 57 variables significant while another replicate deemed 18 of the 57 variables significant. The figure provides insight that the optimal model size should be in the low 20's. However, there is now a dilemma. The plot also shows the variability in identifying significant variables due to the variability of training set data. By choosing the results from a single replicate, it is possible that a noise variable was deemed significant or a significant variable was missed due simply to the sample training data. This potential error is further demonstrated by the next two examples.




FIG 3. Number of Variables Deemed Significant in Replications

The first example below shows 2 of the 100 replicates after performing AIC stepwise regression. In the replicate on the left, 18 variables were deemed significant, while 30 variables were deemed significant in the replicate on the right. The second example below shows 2 of the 100 replicates after performing BIC stepwise regression. In the replicate on the left, 17 variables were deemed significant, while 26 variables were deemed significant in the replicate on the right. In the next section, we will examine how to find a more robust method that helps truly identify significant variables.

Coefficients:	Estimate	Std. Error	z value	Pr(> z)	(Intercept)	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-4.058e+00	5.171e-01	-7.849	4.21e-15 ***	wf.address	-2.196e-01	1.423e-01	-0.008	0.364784
wf.address	-3.165e-01	2.322e-01	-1.443	0.149060	wf.3d	8.523e+00	1.055e+01	0.808	0.419087
wf.all	-3.696e-01	2.859e-01	-1.293	0.198107	wf.our	1.707e+01	9.925e-01	5.835	5.29e-09 ***
wf.our	1.102e+00	3.105e-01	3.873	0.000107 ***	wf.over	3.123e+00	7.097e-01	4.400	1.08e-05 ***
wf.remove	3.870e+01	9.080e+00	4.262	2.03e-05 ***	wf.remove	7.263e+00	1.554e+00	4.673	2.96e-06 ***
wf.internet	1.117e+00	4.823e-01	2.288	0.023304 *	wf.order	1.634e+00	6.385e-01	2.481	0.013107 *
wf.order	9.449e-01	7.058e-01	1.339	0.180618	wf.will	-7.380e-01	2.069e-01	-3.567	0.000361 ***
wf.mail	1.736e-01	1.105e-01	1.570	0.116346	wf.free	1.920e+00	1.129e-01	4.805	1.55e-06 ***
wf.receive	-2.763e+00	8.873e-01	-3.114	0.001847 **	wf.business	4.444e+00	9.030e-01	4.921	8.61e-07 ***
wf.will	-1.860e-01	1.640e-01	-1.114	0.259687	wf.email	-6.050e-01	1.057e-01	-1.979	0.047843 *
wf.free	2.075e+00	5.221e-01	3.974	7.06e-05 ***	wf.you	-2.697e-01	1.028e-01	-2.624	0.008613 **
wf.business	1.442e+00	6.486e-01	2.056	4.99e-05 ***	wf.credit	5.886e+00	1.933e+00	3.046	0.002322 **
wf.you	1.874e-01	7.099e-02	2.639	0.008304 **	wf.your	3.046e-01	1.121e-01	2.718	0.005052 **
wf.credit	2.113e+00	1.673e+00	1.262	0.206779	wf.000	4.887e+00	1.815e+00	2.692	0.007104 **
wf.your	2.082e-01	1.082e-01	1.924	0.054407 *	wf.money	4.007e-01	1.648e-01	2.431	0.015000 *
wf.000	1.407e+00	9.058e-01	1.553	0.120398	wf.hp	-2.093e-01	4.716e-01	-0.438	0.67e-06 ***
wf.hp	-1.603e+01	4.463e+00	-4.039	5.36e-05 ***	wf.hpl	-3.625e+00	3.608e+00	-1.005	0.314943
wf.hpl	-6.207e+00	4.011e+00	-1.547	0.121745	wf.george	-2.358e+00	4.864e+00	-0.484	1.25e-06 ***
wf.george	-1.179e+02	2.515e+01	-4.680	3.39e-10 ***	wf.650	3.630e+00	8.377e-01	4.333	1.47e-05 ***
wf.tab	-1.926e-01	3.411e-01	-0.585	0.572241	wf.lab	-1.508e-01	1.528e-01	-0.987	0.323670
wf.data	-1.597e+00	1.172e+00	-1.363	0.172840	wf.data	-1.894e+00	6.934e-01	-2.732	0.006301 **
wf.650	-2.956e+00	3.248e+00	-0.910	0.361674	wf.85	-4.750e+00	1.086e+00	-1.539	0.123737
wf.technology	2.697e+00	6.051e-01	3.346	0.000821 ***	wf.technology	7.723e+00	4.478e-01	2.882	0.003967 **
wf.parts	-5.438e+00	4.424e+00	-1.229	0.219023	wf.pm	-1.749e+00	8.861e-01	-1.974	0.048418 *
wf.pm	-1.597e+00	1.114e+00	-1.433	0.151939	wf.direct	-1.793e+00	1.541e+00	-1.164	0.244463
wf.cs	-1.501e+02	1.443e+02	-1.040	0.298309	wf.cs	-1.229e+02	9.936e+01	-1.237	0.219304
wf.meeting	-6.747e+00	4.402e+00	-1.533	0.121339	wf.meeting	-2.607e+01	7.828e+00	-3.330	0.000869 ***
wf.project	-2.965e+00	1.345e+00	-2.204	0.027549 *	wf.project	-4.634e+00	1.556e+00	-2.978	0.002897 **
wf.re	-1.377e+00	4.040e-01	-3.408	0.000654 ***	wf.re	-1.814e+00	1.434e-01	-5.282	1.28e-07 ***
wf.edu	-2.816e-01	5.274e-02	-5.340	9.28e-08 ***	wf.edu	-5.762e+00	1.421e+00	-4.070	4.71e-05 ***
wf.table	-2.467e+01	4.616e+01	-0.514	0.591007	wf.table	-6.934e+00	8.322e+00	-0.833	0.404711
wf.confidence	-1.041e+01	8.982e+00	-1.161	0.245452	wf.confidence	-1.450e+00	1.436e+00	-1.719	0.085558 *
cf.semicon	-2.009e+00	1.051e+00	-1.911	0.056008 *	cf.semicon	-3.021e+00	6.809e-01	-0.968	0.329998 *
cf.parenthesis	-1.184e+00	2.767e-01	-2.523	0.011621 *	cf.parenthesis	-3.710e+00	1.176e+00	-3.155	0.001606 **
cf.exclamation	1.886e+00	3.303e-01	5.709	1.14e-08 ***	cf.exclamation	3.640e+00	1.241e-01	6.943	1.89e-12 ***
cf.dollar	1.807e+01	3.418e+00	5.288	1.24e-07 ***	cf.dollar	2.456e+01	1.320e+00	7.395	1.41e-13 ***
cap.run.length.avg	8.937e-01	1.404e-01	6.368	1.92e-10 ***	cf.pound	1.038e-01	1.472e+00	3.613	0.000393 ***
cap.run.length.total	3.499e-03	9.004e-04	3.886	0.000102 ***	cap.run.length.max	1.954e-02	4.749e-03	4.115	3.87e-05 ***
					cap.run.length.total	2.743e-03	6.881e-04	3.987	6.68e-05 ***

FIG 4. 2 Replicates After AIC Stepwise Regression

Coefficients:	Estimate	Std. Error	z value	Pr(> z)	Coefficients:	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-4.197e+00	4.633e-01	-9.060	< 2e-16 ***	(Intercept)	-4.628e+00	4.893e-01	-9.457	< 2e-16 ***
wf.our	4.208e-01	1.076e-01	3.909	9.26e-05 ***	wf.our	2.675e+00	4.090e-01	6.541	6.12e-11 ***
wf.remove	1.289e+01	2.499e+00	5.157	2.50e-07 ***	wf.remove	3.031e+00	7.866e-01	3.833	0.000117 ***
wf.internet	1.759e+00	6.568e-01	2.678	0.007402 **	wf.remove	9.283e+00	2.084e+00	4.433	9.29e-06 ***
wf.receive	-2.693e+00	7.370e-01	-3.654	0.000259 ***	wf.internet	1.732e+00	4.416e-01	3.905	8.41e-05 ***
wf.free	1.684e+00	3.626e-01	4.645	3.40e-06 ***	wf.order	3.114e+00	1.120e+00	2.781	0.005424 **
wf.business	5.114e+00	1.028e+00	4.976	6.49e-07 ***	wf.receive	-1.420e+00	5.867e-01	-2.420	0.015524 *
wf.you	2.397e-01	5.909e-02	4.056	5.00e-05 ***	wf.free	7.416e-01	2.289e-01	3.239	0.001198 **
wf.hp	-2.469e+01	5.373e+00	-4.595	4.33e-06 ***	wf.business	2.981e+00	6.893e-01	4.331	1.49e-05 ***
wf.hpl	-6.203e+00	5.081e+00	-1.221	0.222160	wf.you	2.986e-01	5.775e-02	5.171	2.32e-07 ***
wf.george	-2.907e+01	3.860e+00	-7.530	5.07e-14 ***	wf.credit	3.949e+00	1.936e+00	2.040	0.041350 *
wf.lab	-6.730e+01	3.088e+01	-2.179	0.029319 *	wf.hp	-1.349e+01	3.768e+00	-3.581	0.000342 ***
wf.pm	-1.106e+00	6.194e-01	-1.785	0.074187 *	wf.hpl	-9.718e+00	4.123e+00	-2.357	0.018431 *
wf.cs	-1.101e+02	8.226e+01	-1.338	0.180791	wf.george	-1.278e-01	6.011e-01	-0.453	4.95e-08 ***
wf.meeting	-8.776e+00	3.676e+00	-2.387	0.016964 *	wf.650	7.750e+00	1.214e+00	6.382	1.74e-10 ***
wf.project	-2.311e+00	1.248e+00	-1.852	0.064015 *	wf.labs	-3.435e+00	1.038e+00	-3.246	0.001172 **
wf.re	-1.269e+00	4.065e-01	-3.123	0.001790 **	wf.85	-9.342e+00	2.287e+00	-4.084	4.42e-05 ***
wf.edu	-3.347e+01	5.907e+00	-6.005	1.91e-09 ***	wf.pm	-1.565e+00	1.380e+00	-2.583	0.009782 **
wf.confidence	-6.943e+00	4.649e+00	-1.493	0.135340	wf.meeting	-7.209e+00	2.369e+00	-3.042	0.002348 **
cf.semicon	-1.964e+00	1.061e+00	-1.852	0.064049 *	wf.project	-7.105e+00	2.447e+00	-2.903	0.003690 **
cf.exclamation	1.184e+00	3.236e-01	3.715	1.10e-08 ***	wf.re	-2.332e+00	4.058e-01	-5.746	9.12e-09 ***
cf.dollar	2.533e+01	3.887e+00	6.517	1.91e-11 ***	wf.edu	-8.190e+00	1.414e+00	-5.793	6.50e-09 ***
cap.run.length.avg	9.439e-01	1.273e-01	7.417	1.20e-13 ***	wf.confidence	-5.339e+00	1.232e+00	-1.713	0.086641
cap.run.length.total	3.394e-03	7.376e-04	4.601	4.20e-06 ***	cf.semicon	-1.267e+00	6.267e-01	-2.022	0.043180 *
					cf.exclamation	2.072e+00	3.489e-01	5.938	2.88e-09 ***
					cf.dollar	2.423e+01	1.340e+00	7.254	4.06e-13 ***
					cap.run.length.avg	8.933e-01	1.183e-01	7.553	4.15e-14 ***
					cap.run.length.total	1.079e-03	3.716e-04	2.905	0.003674 **

FIG 5. 2 Replicates After BIC Stepwise Regression

The percentage in which each explanatory variable was deemed significant (p -value ≤ 0.05) out of the 100 training replicates was calculated and plotted in the figure below (step 4). There are several advantages to this approach. Firstly, we have identified, with high confidence, the 18 significant variables to comprise our core model. These variables are wf.remove, wf.hp, wf.re, wf.our, wf.free, wf.edu, cf.exclamation, cf.dollar, wf.business, wf.george, wf.project, cap.run.length.total, wf.your, wf.000, wf.internet, wf.receive, wf.over, and wf.money (note: wf is acronym for word frequency, cf is acronym for character frequency, and cap.run.length.total is the total number of capital letters in an email); for the remainder of this analysis, they will be referred to as the 18 core variables. There is high confidence that these 18 variables are significant variables due to their robustness. They were consistently identified as significant in $\geq 50\%$ of the replicates in all 4 of the variable selection methods despite the variability in the training sets' observations.




FIG 6. *Percentage in Which Each Explanatory Variable Deemed Significant*

Secondly, this approach also affords flexibility to the end user in modeling. There were 7 variables (marked by red dotted vertical lines) which were deemed significant in $\geq 50\%$ of the replicates by at least 1 of the 4 variable selection methods but not by all 4. These variables are `wf.technology`, `cap.run.length.avg`, `wf.meeting`, `wf.order`, `wf.you`, `wf.credit`, and `cap.run.length.max`. Depending on the end user's threshold for model complexity-accuracy tradeoff, the user can experiment building models with any combination of these 7 variables in addition to the 18 core variables. Furthermore, each variable's importance can now be characterized by the percent of replicates in which they are deemed significant. For example, it may not be cost-effective for a business to measure all 18 core variables or may require too much computing power; as a result, the user may be constrained to using only 10 variables. How would the end user decide which 10 variables to use? Based on the plot, the user should select the first 10 variables (`wf.remove`, `wf.hp`, `wf.re`, `wf.our`, `wf.free`, `wf.edu`, `cf.exclamation`, `cf.dollar`, `wf.business`, and `wf.george`) because they were statistically significant in $> 90\%$ of the replicates in all 4 variable selection techniques. One can view this approach as an alternative to Mallow's Cp.

The remaining 32 variables (the variables to the right of the rightmost red dotted line) can be discarded because they never exceeded 50% in the replicate sets for any of the 4 selection methods. These 32 variables are more susceptible to the randomness of the observations in the training replicates. When one compares multiple replicates, these variables will not be consistently of deemed significance. The potential error in using only one variable selection method becomes

evident. Suppose a modeler selected the model with the highest test set accuracy after performing only BIC stepwise regression (represented by green line in the plot above). There is $\sim 20\%$ chance the chosen model would not include `wf.project` (the percentage of an email in which the word "project" appears). However, when looking at variable selection from an ensemble point of view, there is high confidence that `wf.project` is a significant variable despite not being deemed significant in a replicate. Alternatively, there is a 30% chance that the chosen model would include `cf.semicolon` (the percentage of an email in which the character ";" appears). However, when looking at variable selection from an ensemble point of view, there is low confidence that `cf.semicolon` is a significant variable because it never reached $\geq 50\%$ in any of the variable selection methods.

Five families of models were built in increasing complexity (step 5). For the first family, a model was built for each training replicate using the 18 core variables. The accuracies were calculated for both training and test sets. This process was repeated 4 more times, in which different variables were added onto the 18 core variables (the variables used are listed below).

Let M_c denote the core model containing the 18 variables that always appear in every replication.

TABLE 3
Variable Selection For Various Model Complexities

Model Complexity	Variables Used
18 Variables	M_c
19 Variables	$M_{19} = M_c \cup \{\text{wf.technology}\}$
20 Variables	$M_{20} = M_{19} \cup \{\text{cap.run.length.avg}\}$
21 Variables	$M_{21} = M_{20} \cup \{\text{wf.meeting}\}$
23 Variables	$M_{23} = M_{21} \cup \{\text{wf.you}, \text{wf.credit}\}$

TABLE 4
Test Set Accuracies For Various Model Complexities

Model Complexity	Min	1st Qu	Median	Mean	3rd Qu	Max
18	91.09%	93.04%	93.51%	93.50%	93.93%	94.93%
19	91.09%	93.04%	93.48%	93.46%	93.91%	94.86%
20	91.67%	93.32%	93.94%	93.78%	94.22%	95.29%
21	91.09%	93.62%	94.13%	94.05%	94.49%	95.65%
23	90.94%	93.84%	94.13%	94.12%	94.42%	93.51%

For nonlinear models like the ones that arise in the generalized linear model framework, it is often the case that the Fisher scoring algorithm used to estimate the parameters of the model does not converge. When that happens, the solution delivered is at best suboptimal, and may lead to misleading conclusions. For instance, the gist of the method proposed in this paper lies on scanning the variables and choosing the model made up of variables whose p -values are less than 0.05. With a suboptimal solution, it is unwise and misleading to consider that p -values are meaningful. For that reason, we systematically tract all the

estimations throughout the totality of our random replications, and we provide an estimate of the percentage of times the estimates of the model are meaningful. It makes sense to us that only the cases where convergence is achieved should be used for inference, because -in a sense- that measures an aspect of the quality of the model space search. Below is a partial table of the percentages:

TABLE 5
Percentage of Convergence/Non-Convergence For Various Model Complexities

Model Complexity	% Converge	% Non-Converge
18	73	27
19	80	20
20	68	32
21	37	63
23	28	72

The table above shows the test set accuracies for all 5 models built. The last two columns in the table above list the percentage of 100 models which converged and the percentage of 100 models which did not converge. All 5 of the models have very respectable accuracies; furthermore, there is relatively little variation in test set accuracies throughout the 100 replicates. The end user now has several models to choose from. If the end user sought the most parsimonious model and is willing to accept a slight loss in accuracy, he/she can select the 18 variable model. If, on the other hand, the end user sought the highest prediction rate, he/she would select the 23 variable model. If the modeler sought the most computationally stable model, he/she would select the 19 variable model; this model had 80% of its replicates converge. For comparison with current machine learning methods, we selected the 20 variable model. The 20 variable model achieved the best tradeoff between model accuracy and convergence rate; this model complexity is in agreement with the predicted optimal model size. The following two figures plot the ROC curves for all 5 models. In the first figure, all 5 ROC curves achieve "right angle" shape. The second figure provides a close-up of the upper-left corner of the first figure. The 20 variable model (green line) tracks well with the 23 variable model (purple dotted line); both of these have more area under the curve compared to the remaining 3 models.

The 20 explanatory variables used to generate the model were wf.remove, wf.hp, wf.re, wf.our, wf.free, wf.edu, cf.exclamation, cf.dollar, wf.business, wf.george, wf.project, cap.run.length.total, wf.your, wf.000, wf.internet, wf.receive, wf.over, wf.money, wf.technology, and cap.run.length.avg. 6 out of the 20 explanatory variables are indicators of non-spam (as their frequency in an email increases, the more likely an email is not spam); these variables are wf.hp, wf.re, wf.george, wf.edu, wf.project, and wf.technology. 14 out of the 20 explanatory variables are indicators of spam (as their frequency in an email increases, the email is more likely to be spam); these variables are wf.000, wf.money, cf.dollar, wf.free, wf.remove, wf.business, wf.your, cf.exclamation, cap.run.length.avg, cap.run.length.total, wf.internet, wf.over, wf.our, and wf.receive.




FIG 7. *ROC Curves for 5 Models*




FIG 8. *Close Up of ROC Curves for 5 Models*

When one looks at both classes of indicators, certain patterns arise which can provide further explanation. The variables `wf.hp`, `wf.re`, `wf.george`, `wf.project`, and `wf.technology` in the non-spam indicators, suggest a personal or professional relationship with the recipient. Since the dataset was donated by [3] at Hewlett Packard labs, it makes sense that emails containing "hp" and "George" indicate that the sender either knew the recipient and/or the email was work-related. The characters "re" are often used in emails as replies. Therefore, email replies are flagged as non-spam because the recipient is receiving a reply to an earlier email sent out by the recipient. Additionally, the words meeting and project are usually work-related terms and hence why they are also indicators of non-spam. The variables `wf.000`, `wf.money`, `cf.dollar`, `wf.free`, and `wf.business` are associated with money. This makes sense since most spam emails are attempts to get money from the recipient.

In [6], different machine learning techniques were applied to the Spambase dataset with the goal of optimizing correct classification rate. The following table lists the testing set accuracy for the different machine learning techniques used in the paper. Our 20 variable model, with a median test accuracy of 93.94%, was outperformed by Ensemble Decision Tree and Adaboost methods.

TABLE 6
Test Set Accuracies For Various Machine Learning Techniques

Classifier	Accuracy
Ensemble Decision Tree (Nb of trees = 25)	96.40%
Adaboost	95.00%
Stacking	93.80%
SVM	93.40%
Bagging	92.80%
Decision Tree	92.58%
Neural Network	90.80%
Naive Bayes	89.57%
Nearest Neighbor (k=5)	89.40%

Independently, [5] applied Generalized Additive Neural Networks (GANNs) to the Spambase dataset to classify email as spam/non-spam. The spambase dataset was broken in 67% training and 33% test. The AutoGANN method used by [5] attained a 95.8% accuracy. While this methods accuracy beat the 20 variable model ($95.8\% > 93.94\%$), it required higher model complexity with 41 explanatory variables used.

[9] used a neural network method called Multi-Layer Perceptron (MLP). In this paper, the Spambase dataset was broken down into training set with 4,025 observations ($\sim 87.48\%$) and test set with 576 observations ($\sim 12.59\%$). The author applied this technique on 3 different scenarios. In the first scenario, [9] used MLP on the training set using all 57 exploratory variables. Due to the high dimensionality, the author noted that the MLP method had trouble converging and consequently, generated poor classification rates. A similar effect was observed in our approach. For our 5 models, the percentage of models that converged drastically decreased when model complexity increased above

20 variables. The following table lists the accuracies for different models built by the author using all 57 variables.

TABLE 7
Accuracies for MLP Using 57 Variables

MLP Architecture	Learning Rate	Momentum	Average Classification Rate
20 – 10 – 10 – 10 – 5	0.1	0.80	62.40%
20 – 10 – 10 – 10 – 5	0.1	0.95	63.20%
20 – 10 – 10 – 5	0.1	0.85	60.61%
15 – 15 – 15 – 5	0.1	0.85	60.59%

The author used the same MLP method after reducing dimensionality to 21 variables. The test accuracies listed below are very comparable to the test set accuracies for our 20 variable model.

TABLE 8
Accuracy for MLP Using 21 Variables

MLP Architecture	Learning Rate	Momentum	Average Classification Rate
20 – 10 – 10 – 10 – 5	0.1	0.80	93.50%
20 – 10 – 10 – 10 – 5	0.1	0.95	90.30%
20 – 10 – 10 – 10 – 7	0.1	0.80	93.80%

Lastly, in the third attempt, MLP was applied after reducing dimensionality to 9 variables. The test accuracies are listed below.

TABLE 9
Accuracy for MLP Using 9 Variables

MLP Architecture	Learning Rate	Momentum	Average Classification Rate
20 – 10 – 10 – 8	0.1	0.80	92.40%
20 – 10 – 10 – 5	0.1	0.80	91.70%
20 – 10 – 10 – 5	0.1	0.95	91.80%

In a fourth comparison, [8] utilized 9 different machine learning algorithms for their paper *Adaptive Approach for Spam Detection*. In their approach, the data was transformed into 1's and 0's. If a certain word appeared, that exploratory variable was a 1; if the certain word did not appear in an email, the exploratory variable was a 0. This was done for 55 of the 57 exploratory variables. The following table lists the performance of different algorithms after ten-fold cross validation. Our 20 variable model, with a median test accuracy of 93.94%, was out-performed by Random Committee and essentially tied with Random Forest technique.

There are several advantages to performing frequentist approach. First, this method provides a more robust variable selection by examining how often a variable is deemed significant by multiple traditional variable selection methods given random samples of observation data. Second, it also provides an approximation to the optimal model size. Third, it allows the modeler to characterize the importance of a variable to the model through the frequency in which a

TABLE 10
Accuracy for Various Machine Learning Algorithms

Algorithm	Accuracy
Bayes Network	88.56%
Logic Boost	89.70%
Random Tree	91.54%
JRip	92.32%
J48	92.34%
Multilayer Perceptron	93.28%
Kstar	93.56%
Random Forest	93.89%
Random Committee	94.28%

variable is deemed significant. Lastly, it affords the modeler flexibility in choosing certain variables to retain or discard depending on the modeler’s threshold for model-complexity accuracy tradeoff. The main downside to this method is computational intensity. On a Windows 7 64-Bit Laptop with Intel i7 2.7GHz processor and 16GB RAM, this process required ~ 12 hours to just apply BIC stepwise regression to 100 replicates and another 12 hours to apply AIC stepwise regression to the 100 replicates. Running the full model and forward selection via BIC on the 100 replicates was markedly faster and completed within minutes. By incorporating more variable selection methods to the ensemble, the modeler will have a serious tradeoff in computing time. However, this dilemma may be alleviated through the use of parallel processing in which multiple tasks are dispersed over multiple workstations rather than running the tasks sequentially on one computer. With further advances in parallel processing and increases in computing power, the ensemble variable selection method’s advantages will significantly dominate over its main weakness.

4.2. Computational demonstration on regression analysis

In the Spambase dataset, the frequentist approach was applied to classification. In the next example, the frequentist approach was applied to Multiple Linear Regression (MLR) on the Bodyfat dataset. The dataset, which was originally donated by [7], attempts to estimate body fat percentage by underwater weighing and various body circumference measurements for 252 men; this dataset may be found in the R package `mfp`¹. The dataset contained 2 response variables: `brozek` and `siri`. The `brozek` response variable calculated body fat percentage through the equation:

$$\text{brozek} = \frac{457}{\text{density}} - 414.2$$

The `siri` response variable calculated body fat percentage through the equation:

$$\text{siri} = \frac{495}{\text{density}} - 450$$

¹This bodyfat data set can also be found at <http://lib.stat.cmu.edu/datasets/bodyfat>

There were 14 explanatory variables. The first 3 variables are **density** (density determined from underwater weighing), **age**, and **weight**. The remaining 11 explanatory variables are body circumference measurements for **neck**, **chest**, **abdomen**, **hip**, **thigh**, **knee**, **ankle**, **biceps**, **forearm**, and **wrist**. The linear pairwise correlation plot below indicates a significant amount of multicollinearity and redundant variables.




FIG 9. Linear Correlation Plot for Bodyfat Dataset

The purpose of this section will be to compare the results from the frequentist approach for variable selection against results from Bayesian PIP. There were 5 cases identified in the dataset as erroneous observations/outliers, and as a result, were excluded during analyses. Cases 48, 76, and 96 were identified as having errors in body fat values. In case 42, the man weighed 200 lbs with a height of 3 ft. For case 182, the body fat percentage was rounded to 0 after having a negative body fat percentage. Additionally, the **density** variable was excluded from MLR because formulas to calculate **brozek** and **siri** response variables require **density** variable; consequently, the inclusion of **density** variable in MLR would dominate and bias the selection of other significant variables. The dataset was broken into 500 replicates of 70% training / 30% test sets (step 1 of frequentist approach). The linear model (MLR) was selected for the class of models (step 2). AIC stepwise regression, BIC stepwise regression BIC forward selection, and full model were applied to the 500 training sets (step 3). This section will concentrate solely on the **brozek** response variable. After scanning through the 500 replicates, the percentage in which each variable was deemed significant is listed below for each of the variable selection methods (step 4).

TABLE 11
Estimated Percentage of Inclusion of each variable

Variable	Full Model	Stepwise BIC	Forward BIC	Stepwise AIC
abdomen	1	1	1	1
wrist	0.8	0.698	0.8	0.842
weight	0.13	0.576	0.13	0.53
forearm	0.386	0.36	0.386	0.516
neck	0.21	0.222	0.21	0.304
hip	0.082	0.22	0.082	0.226
height	0.01	0.166	0.01	0.156
age	0.124	0.08	0.0124	0.2
biceps	0.036	0.08	0.036	0.1
thigh	0.07	0.038	0.07	0.208
chest	0.012	0.028	0.012	0.052
ankle	0.014	0.016	0.014	0.034
knee	0	0.002	0	0.006




FIG 10. Percentage of Inclusion of Explanatory Variables for **brozek** Response Variable - Frequentist Approach

The following section will compare the results from our frequentist approach against the results from Bayesian PIP. In order for a more objective comparison, only the stepwise BIC portion will be compared against the Bayesian PIP. The Bayesian PIP results were acquired using the Bayesian Model Selection (BMS) package in R and shown below.

The frequentist approach (looking at stepwise BIC only) and the Bayesian PIP both identified **abdomen** as the most significant explanatory variable; in both

	PIP	Post Mean	Post SD	Cond.Pos.Sign	Idx
abdomen	1.00000000	0.8994546675	0.05909395	1.00000000	6
weight	0.94254054	-0.1171392909	0.04085303	0.00000000	2
wrist	0.59224949	-0.7619333211	0.72426259	0.00000000	13
forearm	0.28606611	0.1181661995	0.20943063	0.99999662	12
neck	0.14135764	-0.0570056583	0.16396585	0.00000000	4
biceps	0.13742721	0.0382317610	0.11145263	0.99984420	11
hip	0.09782832	-0.0222768348	0.08206601	0.00000000	7
thigh	0.07307283	0.0110396174	0.05035270	0.99731043	8
age	0.05365942	0.0015489301	0.01020422	0.86215779	1
height	0.04282675	-0.0046580449	0.04223018	0.08660439	3
ankle	0.03962332	0.0056336034	0.05104796	0.98806057	10
knee	0.03664630	0.0043967238	0.04736250	0.96027152	9
chest	0.03268106	-0.0005889387	0.01716922	0.28242364	5

Mean no. regressors	Draws	Burnins	Time
"3.4760"	"8192"	"0"	"1.356216 secs"
No. models visited	Modelspace 2^K	% visited	% Topmodels
"8192"	"8192"	"100"	"6.1"
Corr PMP	No. Obs.	Model Prior	g-Prior
"NA"	"247"	"random / 6.5"	"UIP"
Shrinkage-Stats			
"AV=0.996"			

Time difference of 1.356216 secs

FIG 11. R Output - Bayesian PIP of Explanatory Variables for *brozek* Response Variable

methods, the **abdomen** variable was deemed significant in 100% of the 500 training sets. Additionally, in both methods, **weight** and **wrist** variables exceeded the median probability model (deemed significant in $\geq 50\%$ of 500 replicates) and will be retained as significant variables. The main difference between the frequentist approach and PIP is evident in the **weight** variable. The **weight** variable was deemed significant in $\sim 94\%$ of the 500 models by Bayesian PIP but only $\sim 58\%$ by the frequentist method. The following plot shows the distribution of model size across the 500 replicates for the frequentist approach. The optimal model size should include 3 to 4 variables. The average model size across 500 replicates was 3.49.

The following plot shows the distribution of model size across 500 replicates for the Bayesian PIP approach. The PIP approach also indicated an optimal model size between 3 to 4 variables with an average model size of 3.48 across 500 replicates. The results between the two methods are very similar.

The frequentist approach (looking at stepwise BIC only) and the Bayesian PIP both identified **abdomen**, **weight** and **wrist** as significant variables. Building 500 linear models using only the 3 variables, we obtain the following results shown below. As expected, the out of sample prediction results for the frequentist and Bayesian PIP because we are applying the same variables to the training sets. We will now build new sets of models to compare the frequentist and Bayesian PIP approaches. From a meta-analysis standpoint, **abdomen** and **wrist** variables were deemed significant in $\geq 50\%$ of the replicates for all 4 variable selection techniques in the frequentist approach. As a result, 500 linear models were




FIG 12. *Optimal Model Size - Frequentist Approach*




FIG 13. *Optimal Model Size - Bayesian PIP Approach*




FIG 14. Comparison of Relative PRESS - Both Approaches Use abdomen, weight and wrist Variables




FIG 15. Comparison of MSE - Both Approaches Use abdomen, weight and wrist Variables

built using only the **abdomen** and **wrist** explanatory variables. From the PIP standpoint, **abdomen**, **weight**, and **wrist** variables exceeded the median probability model. As a result, 500 linear models were built using **abdomen**, **weight**, and **wrist** explanatory variables. We could have justifiably included **weight** as an additional third variable in our frequentist approach because it was deemed significant by at least one variable selection technique; however, in doing so, we would get the same end results as the PIP linear models since both methods would now use the same variables. By comparing a two variable model (**abdomen** and **wrist**) against a three variable model (**abdomen**, **weight**, and **wrist**), we hope to achieve a more distinction between the two model types. In the next section, we will compare the out of sample predictive performance of the frequentist and the PIP approaches. The following table and plot will compare the relative PRESS values between the two methods.

TABLE 12
Comparison of PRESS Between Frequentist Approximation and Bayesian PIP

Model	Min	1st Qu	Median	Mean	3rd Qu	Max
Frequentist	0.6025	0.6638	0.6801	0.6812	0.6983	0.7585
BMS PIP	0.6243	0.6921	0.7065	0.7061	0.7208	0.7708




FIG 16. Comparison of Relative PRESS

The 500 linear models built for each of the approaches were applied on their corresponding test sets (step 5). The following table and plot will compare the Mean Squared Error (MSE) across the 500 test sets for both methods.

Model	Min	1st Qu	Median	Mean	3rd Qu	Max
Frequentist	11.31	15.84	17.52	17.68	19.27	25.31
BMS PIP	11.25	14.71	16.14	16.18	17.51	23.34




FIG 17. Figure: Comparison of MSE

While `weight` achieved a PIP of 0.94, its addition into the three variable model did not provide a practical improvement. The results between the PIP and frequentist approaches are very comparable and provides validity of using the frequentist approach as an alternative to the Bayesian PIP.

The results when using `siri` as the response variable are almost identical to the results attained above using `brozek` as the response variable. As a result, the comparison between the frequentist and PIP approaches for `siri` response variable will not be provided.

Applying the frequentist approach on linear models was significantly faster than on generalized linear models. Running stepwise regression via BIC and AIC, forward selection via BIC, and the full model on the 500 replicates for both `brozek` and `siri` response variables only took minutes.

5. Conclusion and discussion

We have used a straightforward, quite general and easily interpretable subsampling scheme to provide a frequentist approximation of the celebrated Bayesian posterior inclusion probability. Despite the relatively higher computational burden arising in the use of the proposed method on high dimensional classification

tasks, it is fair to say that the present method mimics the Bayesian framework quite well. All the scores, judging from the numerical values and the corresponding plots appear identical or at least very similar in shape and form. One would be particularly be excited to use this approach because it is easier to understand since it builds up on the widely used framework of variable selection by the step-wise regression heuristic. Even more importantly, it does not run into the some of the challenges of the Bayesian framework like the difficulty in computing the marginal density of the data. As we said earlier, the great challenge for this method is the heavy computational burden. However, with the availability of distributed and high performance parallel computing resources, this method becomes even more attractive for high dimensional data mining problems since one can perform the independent random split on different CPUs. Indeed, our future work will focus on substantially reducing the computing time by a careful use of the parallel computing resources.

Acknowledgement

Ernest Fokoué wishes to express his heartfelt gratitude and infinite thanks to Our Lady of Perpetual Help for Her ever-present support and guidance, especially for the uninterrupted flow of inspiration received through Her most powerful intercession.

Competing Interest

The authors declare that no competing interests exist.

References

- [1] M. M. Barbieri and J. O. Berger. Optimal predictive model selection. *Annals of Statistics*, 32:870–897, 2004.
- [2] E Fokoué. Estimation of Atom Prevalence for Optimal Prediction. *Contemporary Mathematics (The American Mathematical Society)*, 443:103–129, 2007.
- [3] G. Forman, M. Hopkins, and E. Reeber. UCI machine learning repository, 1999.
- [4] G. Fumera, I. Pillai, and F. Roli. Spam filtering based on the analysis of text information embedded into images. *Journal of Machine Learning Research*, 7:2699–2720, 2006.
- [5] J. C. Goosen and J. V. Du Toit. Spam detection with generalised additive neural networks. In *Southern Africa Telecommunication Networks and Applications Conference*, August 2009.
- [6] R. Kiran and I. Atmosukarto. Spam or not spam that is the question. Technical report, University of Washington, 2005.

- [7] K.W. Penrose, A.G. Nelson, and A.G. Fisher. Generalized body composition prediction equation for men using simple measurement techniques. *Medicine and Science in Sports and Exercise*, 17(2):189, 1985.
- [8] S. Sharma and A. Arora. Adaptive approach for spam detection. *International Journal of Computer Science Issues*, 10(1):23–26, July 2013.
- [9] T. Sivanadyan. Spam? not any more! detecting spam emails using neural networks. Technical report, University of Wisconsin, 2003.