

EigenNet: A Bayesian hybrid of generative and conditional models for sparse learning

Yuan Qi
Departments of CS and Statistics
Purdue University

Feng Yan
Department of CS
Purdue University

May 28, 2018

Abstract

It is a challenging task to select correlated variables in a high dimensional space. To address this challenge, the elastic net has been developed and successfully applied to many applications. Despite its great success, the elastic net does not explicitly use correlation information embedded in data to select correlated variables. To overcome this limitation, we present a novel Bayesian hybrid model, the EigenNet, that uses the eigenstructures of data to guide variable selection. Specifically, it integrates a sparse conditional classification model with a generative model capturing variable correlations in a principled Bayesian framework. We reparameterize the hybrid model in the eigenspace to avoid overfitting and to increase the computational efficiency of its MCMC sampler. Furthermore, we provide an alternative view to the EigenNet from a regularization perspective: the EigenNet has an adaptive eigenspace-based composite regularizer, which naturally generalizes the $l_{1/2}$ regularizer used by the elastic net. Experiments on synthetic and real data show that the EigenNet significantly outperforms the lasso, the elastic net, and the Bayesian lasso in terms of prediction accuracy, especially when the number of training samples is smaller than the number of variables.

1 Introduction

In this paper we consider the problem of selecting correlated variables in a high dimensional space. Among many variable selection methods, the lasso and the elastic net are two popular choices [Tibshirani, 1994, Zou & Hastie, 2005]. The lasso uses a l_1 regularizer on model parameters. This regularizer shrinks the parameters towards zero, removing irrelevant variables and yielding a sparse model [Tibshirani, 1994]. However, the l_1 penalty may lead to over-sparsification: given many correlated variables, the lasso often only select a few of them. This not only degenerates its prediction accuracy but also affects the interpretability of the estimated model. For example, based on high-throughput

biological data such as gene expression and RNA-seq data, it is highly desirable to select multiple correlated genes specific to a phenotype since it may reveal underlying biological pathways. Due to its over-sparsification, lasso may not be suitable for this task.

To address this issue, the elastic net has been developed to encourage a grouping effect, where strongly correlated variables tend to be in or out of the model together [Zou & Hastie, 2005]. However, the grouping effect is just the result of its composite l_1 and l_2 regularizer; the elastic net does not explicitly incorporate correlation information among variables in its model.

In this paper, we propose a new sparse Bayesian hybrid model, called the *EigenNet*. Unlike the previous sparse models, it uses the eigen information from the data covariance matrix to guide the selection of correlated variables. Specifically, it integrates a sparse *conditional* classification model with a *generative* model capturing variable correlation in a principle Bayesian framework [Lasserre et al., 2006]. The hybrid model enables identification of groups of correlated variables guided by the eigenstructures. Also, it passes the information from the conditional model to the generative model, selecting informative eigenvectors for the classification task. Unlike frequentist approaches, the Bayesian hybrid model can reveal correlations between classifier weights via their joint posterior distribution.

We reparameterize the model in the eigenspace of the data. When the number of predictor variables (i.e., input features), (p), is bigger than the number of training samples (n), this reparameterization restricts the model in the data subspace, which not only reduces overfitting, but also allows us to develop efficient Markov Chain Monte Carlo sampler.

From the regularization perspective, the EigenNet naturally generalizes the elastic net by using a composite regularizer adaptive to the data eigenstructures. It contains a l_1 sparsity regularizer and a directional regularizer that encourages selecting variables associated with eigenvectors chosen by the model. When the variables are independent of each other, the eigenvectors are parallel to the axes and this composite regularizer reduces to the $l_{1/2}$ regularizer used by the elastic net; when some of the input variables are strongly correlated, the regularizer will encourage the classifier aligned with eigenvectors selected by the model. On one hand, our model is like the elastic net to retain ‘all the big fish’. On the other hand, our model is different from the elastic net by using the eigenstructure. Hence the name EigenNet.

Experiments on synthetic and real data are presented in Section 7. They demonstrate that the EigenNet significantly outperforms the lasso, the elastic net, and the Bayesian lasso [Park et al., 2008, Hans, 2009] in terms of prediction accuracy, especially when the number of training samples is smaller than the number of features.

2 Background: lasso and elastic net

We denote n independent and identically distributed samples as

$$\mathcal{D} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$$

, where \mathbf{x}_i is a p dimensional input features (i.e., explanatory variables) and y_i is a scalar label (i.e., response). Also, we denote $[\mathbf{x}_1, \dots, \mathbf{x}_n]$ by \mathbf{X} and $(\mathbf{y}_1, \dots, \mathbf{y}_n)$ by \mathbf{y} . In this paper, we consider the binary classification problem ($y_i \in \{-1, 1\}$), but our analysis and the proposed models can be extended to regression and other problems.

For classification, we use a logistic function as the data likelihood function:

$$p(\mathbf{y}|\mathbf{X}, \mathbf{w}, b) = \prod_i \sigma(y_i(\mathbf{w}^T \mathbf{x}_i + b)) \quad (1)$$

where $\sigma(z) = \frac{1}{1+\exp(-z)}$, and \mathbf{w} and b define the classifier.

To identify relevant variables for high dimensional problems, the lasso [Tibshirani, 1994] uses a l_1 penalty, effectively shrinking \mathbf{w} and b towards zero and pruning irrelevant variables. In a probabilistic framework this penalty corresponds to a Laplace prior distribution:

$$p(\mathbf{w}) = \prod_j \lambda \exp(-\lambda|w_j|) \quad (2)$$

where λ is a hyperparameter that controls the sparsity of the estimated model. The larger the hyperparameter λ , the sparser the model.

As described in Section 1, the lasso may over-penalize relevant variables and hurt its predictive performance, especially when there are strongly correlated variables. To address this issue, the elastic net [Zou & Hastie, 2005] combines l_1 and l_2 regularizers to avoid the over-penalization. The combined regularizer corresponds to the following prior distribution:

$$p(\mathbf{w}) \propto \prod_j \exp(-\lambda_1|w_j| - \lambda_2 w_j^2) \quad (3)$$

where λ_1 and λ_2 are hyperparameters. While it is well known that the elastic net tends to select strongly correlated variables together, it does not use correlation information embedded in the data. The selection of correlated variables is merely the result of a less aggressive regularizer for sparsity.

Besides the elastic net, there are many variants (and extensions) to the lasso, such as the bridge [Frank & Friedman, 1993] and smoothly clipped absolute deviation [Fan & Li, 2001]. These variants modify the l_1 penalty to choose variables, but again do not explicitly use correlation information in data.

3 EigenNet: eigenstructure-guided variable selection

In this section, we propose to use covariance structures in data to guide the sparse estimation of model parameters.

First, let us consider the following toy examples.

3.1 Toy examples

Figure 1(a) shows samples from two classes. Clearly the variables x^1 and x^2 are not correlated. The lasso or the elastic net can successfully select the relevant variable x^1 to classify the data. For the samples in Figure 1(b), the variables x^1 and x^2 are strongly correlated. Despite the strong correlation, the lasso would select only x^1 and ignore x^2 . The elastic net may select both x^1 and x^2 if the regularization weight λ_1 is small and λ_2 is big, so that the elastic net behaves like l_2 regularized classifier. The elastic net, however, does not explore the fact that x^1 and x^2 are correlated.

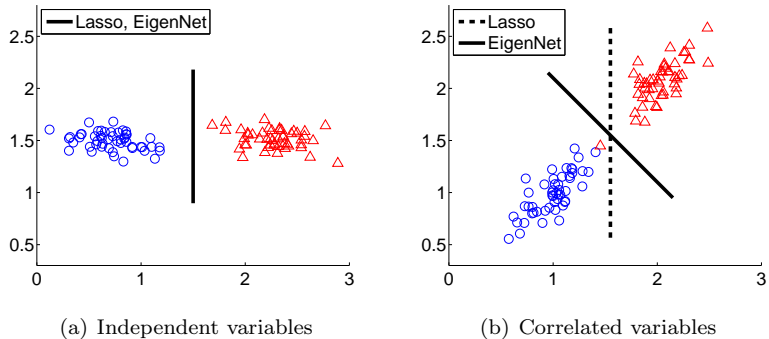


Figure 1: Toy examples. (a) When the variables x^1 and x^2 are independent of each other, both the lasso and the EigenNet select only x^1 . (b) When the variables x^1 and x^2 are correlated, the lasso selects only one variable. By contrast, guided by the major eigenvector of the data, the EigenNet selects both variables.

Since the eigenstructure of the data covariance matrix captures correlation information between variables, we propose to not only regularize the classifier to be sparse, but also encourage it to be aligned with certain eigenvector(s) that are helpful for the classification task. Since our new model uses the eigen information, we name it the EigenNet.

For the data in Figure 1(a), since the two eigenvectors are parallel with the horizontal and vertical axes, the EigenNet essentially reduces to the elastic net and selects x^1 . For the data in Figure 1(b), however, the eigenvectors (in particular, the principle eigenvector) will guide the EigenNet to select both x^1 and x^2 .

We use a Bayesian framework to materialize the above ideas in the EigenNet, as shown in the following section.

3.2 Bayesian hybrid of conditional and generative models

The EigenNet is a hybrid of conditional and generative models. The conditional component allows us to learn the classifier via "discriminative" training; the generative component captures the correlations between variables; and these two models are glued together via a joint prior distribution, so that the correlation information is used to guide the estimation of the classifier and the classification task is used to choose or scale relevant eigenvectors. Our approach is based on the general Bayesian framework proposed by Lasserre et al. [2006]), which allows one to combine conditional and generative models in an elegant principled way.

Specifically, for the conditional model we have the same likelihood as (1), $p(\mathbf{y}|\mathbf{X}, \mathbf{w}, b) = \prod_i \sigma(y_i(\mathbf{w}^T \mathbf{x}_i + b))$. To sparsify the classifier, we can use a Laplace prior on \mathbf{w} ,

$$p(\mathbf{w}) = \prod_j \lambda_1 \exp\{-\lambda_1 |w_j|\}. \quad (4)$$

To encourage the classifier aligned with certain eigenvectors, we use the following generative model:

$$p(\mathbf{V}\mathbf{s}|\tilde{\mathbf{w}}) \propto \exp\left(-\frac{\lambda_2}{2} \sum \eta_j \|\tilde{\mathbf{w}} - s_i \mathbf{v}_i\|_+^2\right) \quad (5)$$

where

$$\begin{aligned} & \|\tilde{\mathbf{w}} - s_i \mathbf{v}_i\|_+^2 \\ & \equiv -\frac{1}{2} \lambda_2 \sum_j \eta_j (\|\tilde{\mathbf{w}}\|^2 - 2s_j |\tilde{\mathbf{w}}^T \mathbf{v}_j| + s_j^2 \|\mathbf{v}_j\|^2) \\ & = -\frac{1}{2} \lambda_2 \sum_j \eta_j (\|\tilde{\mathbf{w}}\|^2 - 2s_j |\tilde{\mathbf{w}}^T \mathbf{v}_j| + s_j^2), \end{aligned} \quad (6)$$

\mathbf{s} are nonnegative continuous variables, \mathbf{v}_i and η_i are the i -th eigenvector and eigenvalue of the data covariance matrix, respectively. The reason we use absolute values of $\tilde{\mathbf{w}}^T \mathbf{v}_j$ in (6) is because we only care about the alignment of $\tilde{\mathbf{w}}$ and \mathbf{v}_i , not the sign of their product. Overall, the above model encourages the classifier to more aligned with the major eigenvectors with bigger eigenvalues. But the variables \mathbf{s} allow us to scale or select individual eigenvectors to remove irrelevant ones.

To integrate the conditional and generative models, we use a joint prior on \mathbf{w} and $\tilde{\mathbf{w}}$:

$$p(\mathbf{w}, \tilde{\mathbf{w}}) \propto \exp(-\lambda_1 |\mathbf{w}|_1) \exp\left(-\frac{\lambda_3}{2} \|\mathbf{w} - \tilde{\mathbf{w}}\|^2\right). \quad (7)$$




Figure 2: The graphical model of the EigenNet.

i.e., we have

$$p(\mathbf{w}, \tilde{\mathbf{w}}) = \lambda_1 \exp(-\lambda_1 |\mathbf{w}|_1) \mathcal{N}(\tilde{\mathbf{w}} | \mathbf{w}, \lambda_3^{-1}). \quad (8)$$

Finally we can assign Gamma priors on all the hyperparameters, λ_1 , λ_2 , and λ_3 . The whole model is depicted in the graphical model in Figure 2.

3.3 Reparameterization and constraint in Eigenspace

In this section we reparameterize the model in the eigenspace:

$$\mathbf{w} = \mathbf{V}\boldsymbol{\alpha} \quad \tilde{\mathbf{w}} = \mathbf{V}\boldsymbol{\beta} \quad (9)$$

where $\mathbf{V} \equiv [\mathbf{v}_1, \dots, \mathbf{v}_m]$ ($m = \min\{n, p\}$), and $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are the projections of \mathbf{w} and $\tilde{\mathbf{w}}$ on the eigenvectors, respectively.

The reparameterization restricts \mathbf{w} in the vector space spanned by $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$, which is equivalent to the data space $\mathcal{C}(\mathbf{X})$, spanned by the data points $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. When the number of features is bigger than the number of training points, i.e., $p > n$, it effectively reduces the number of free parameters in the model, helping avoid overfitting. Furthermore, it provides significant computational advantage when $p \gg n$.

Given $p(\mathbf{w}, \tilde{\mathbf{w}})$ and the relationship between $(\mathbf{w}, \tilde{\mathbf{w}})$ and $(\boldsymbol{\alpha}, \boldsymbol{\beta})$, we obtain $p(\boldsymbol{\alpha}, \boldsymbol{\beta})$ (Please see Appendix for the details):

$$p(\boldsymbol{\alpha}, \boldsymbol{\beta}) \propto \exp(-\lambda_1 |\mathbf{V}\boldsymbol{\alpha}|_1) \exp(-\frac{\lambda_3}{2} \|\boldsymbol{\alpha} - \boldsymbol{\beta}\|^2) \quad (10)$$

Based on the new reparameterization, the likelihood for the conditional model becomes

$$p(\mathbf{y} | \mathbf{X}, \boldsymbol{\alpha}, b) = \prod_i \sigma(y_i(\mathbf{x}_i^T \mathbf{V}\boldsymbol{\alpha} + b)). \quad (11)$$




Figure 3: Adaptive regularization of the EigenNet. The ellipses are the contours of a likelihood function. While the lasso draws the estimates towards the l_1 ball, the EigenNet’s estimate is guided by an eigenvector \mathbf{v} .

Similarly, the likelihood for the generative model becomes

$$\begin{aligned}
 p(\mathbf{V}, \mathbf{s} | \boldsymbol{\beta}) &\propto \exp\left(-\frac{1}{2}\lambda_2 \sum_j \eta_j (\|\mathbf{V}\boldsymbol{\beta}\|^2 \right. \\
 &\quad \left. - 2s_j |\mathbf{V}\boldsymbol{\beta}^T \mathbf{v}_j| + \|\mathbf{v}_j\|^2)\right) \\
 &\propto \exp\left(-\frac{1}{2}\lambda_2 \sum_j (\beta_j^2 - 2\eta_j s_j |\beta_j| + \eta_j s_j^2)\right) \quad (12)
 \end{aligned}$$

The second equation holds since \mathbf{V} is an orthonormal matrix.

Combining (10), (11) and (12), we obtain a complete model. We use Markov Chain Monte Carlo with a random walk proposal to estimate the model parameters \mathbf{s} , \mathbf{w} , and $\tilde{\mathbf{w}}$.

4 Alternative view: composite regularization

In this section, we provide an alternative view to the EigenNet by considering the limiting case of $\lambda_3 \rightarrow 0$. For such as case the prior $p(\boldsymbol{\alpha}, \boldsymbol{\beta})$ becomes

$$p(\boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\boldsymbol{\alpha})\delta(\boldsymbol{\alpha} - \boldsymbol{\beta})$$

This forces $\boldsymbol{\alpha} = \boldsymbol{\beta}$. From a regularization perspective, this prior is equivalent to a composite regularizer:

$$\lambda_1 |\mathbf{w}| + \frac{\lambda_2}{2} \sum \eta_j \|\mathbf{w} - s_j \mathbf{v}_j\|_+^2 \quad (13)$$

$$= \lambda_1 |\mathbf{w}| + \frac{\lambda_2}{2} \sum \eta_j (\|\mathbf{w}\|^2 - 2s_j |\mathbf{w}^T \mathbf{v}_j| + s_j^2) \quad (14)$$

Clearly, when $s_i = 0$ for all i ’s, the above regularizer reduces to the $l_{1/2}$ regularizer used by the elastic net¹. When $s_i \neq 0$ then the regularizer is *adaptive*

¹A subtle difference is that we also constrain \mathbf{w} in the data space for our model.

based on the eigenvector \mathbf{v}_i : First, if the elements of \mathbf{v}_i all have reasonably large values, then all the variables in \mathbf{w} will very likely to be selected. This effect is visualized in Figure 3(b). Second, if this eigenvector has only several large elements, the corresponding variables in $\tilde{\mathbf{w}}$ and \mathbf{w} are likely to be selected jointly. Unlike the $l_{1/2}$ regularizer that encourages the selection of groups of variables from all the variables, our regularizer directly targets at *specific* groups of variables corresponding to the sparse eigenvector. Third, if all the variables are independent of each other, then the eigenvectors are parallel to the axes and each of them contains only one nonzero element. In this case $|\mathbf{w}^T \mathbf{v}_j|$ reduces $|w_j|$, a l_1 regularizer. Figure 3(a) visualizes the eigen regularizer when variables are independent of each other.

In summary, the EigenNet can be viewed as an adaptive generalization of the elastic net by selecting groups of correlated variables based on eigenvectors of the data covariance matrix.

5 Related work

The EigenNet can be viewed as an extension of the classical eigenface approaches [Turk & Pentland, 1991, Sirovich & Kirby, 1987]. The eigenface approach uses PCA coefficients of samples to train a classifier. Naturally the major eigenvectors are often associated with large PCA coefficients and the classifier is constrained in the data subspace when the number of features is smaller than the number of training samples. The EigenNet essentially extends the eigenface approach by combining generative and conditional models in a Bayesian framework and performs sparse learning in an adaptive eigenspace (since the model selects or scales relevant eigenvectors based on s_j).

There are Bayesian versions of the lasso and the elastic net. Bayesian lasso [Park et al., 2008] puts a hyper-prior on the regularization coefficient and use a Gibbs sampler to jointly sample both regression weights and the regularization coefficient. Using a similar treatment to Bayesian lasso, Bayesian elastic net [Li & Lin, 2010] samples the two regularization coefficients simultaneously, potentially avoiding the “double shrinkage” problem described in the original elastic net paper [Zou & Hastie, 2005]. As the EigenNet, these methods are grounded in a Bayesian framework, sharing the benefits of obtaining posterior distributions for handling estimation uncertainty. However, Bayesian lasso and Bayesian elastic net are presented to handle regression problems (though certainly they can be generalized for classification problems) and sample in the original parameter space, not using the eigen information embedded in data. The EigenNet, by contrast, works in the eigenspace and uses eigen information to guide classification.

6 Experimental results

We evaluate the new sparse Bayesian model, the EigenNet, on both synthetic and real data and compare it with three representative state-of-the-art variable selection methods, including the lasso, the elastic net, and the Bayesian lasso modified for classification problems. For the lasso and the elastic net we use the Glmnet software package that uses cyclical coordinate descent in a pathwise fashion². The original Bayesian lasso was developed for regression and uses Gibbs sampling. For the classification tasks we consider, we change its Gaussian regression likelihood to the logistic likelihood (1) while keeping its Laplace prior distributions. We used Markov Chain Monte Carlo, instead of Gibbs sampler, to estimate the classifier for the Bayesian lasso. Bayesian approaches are capable of estimating all the hyperparameters from data. However, for easy and objective comparisons, we simply use cross-validation to tune the hyperparameters, λ_i , for all methods. For the Bayesian lasso and the EigenNet, we draw the 300,000 MCMC samples and use the last 150,000 samples to estimate the posterior mean of the classifiers, which are used for predicting the labels of test samples. We measure the prediction performance of all methods on test samples in terms of their average test error rate (e.g., the 0.2 error rate indicates 20% errors) and report the standard error of the error rates (except for the following visualization example).

6.1 Visualization of estimated classifiers

First, we test these methods on synthetic data that contain correlated features. We sample 40 dimensional data points, each of which contains two groups of correlated variables. The correlation coefficient between variables in each group is 0.81 and there are 4 variables in each group. We set the values of the classifier weights in one group as 5 and in the other group as -5. We also generate the bias term randomly from a standard Gaussian distribution. We set the number of training points to 80. Figure 4 shows the estimated classifiers and the true classifier. It is not surprising that the elastic net identifies more features than the lasso. What is interesting is that EigenNet does not suppress many the irrelevant features to be exactly 0, but it clearly identifies *all* the relevant one, which dominate the irrelevant ones. To save space, we did not show the estimated classifier by the Bayesian lasso. Similar to the EigenNet, its classifier also contains many small, but nonzero weights. On this dataset, the test error rates of the lasso, the elastic net, the Bayesian lasso, and the EigenNet are 0.297, 0.245, 0.251, and 0.137.

An advantage of the Bayesian treatment for feature selection over frequentist approaches is to possibly uncover the correlations between the classifier weights. These correlations can be revealed by the covariance matrices of the joint posterior distribution over the classifier weights. In Figure 5, we visualize the quantized covariance matrices estimated by the Bayesian lasso and the

²<http://www-stat.stanford.edu/~tibs/glmnet-matlab/>

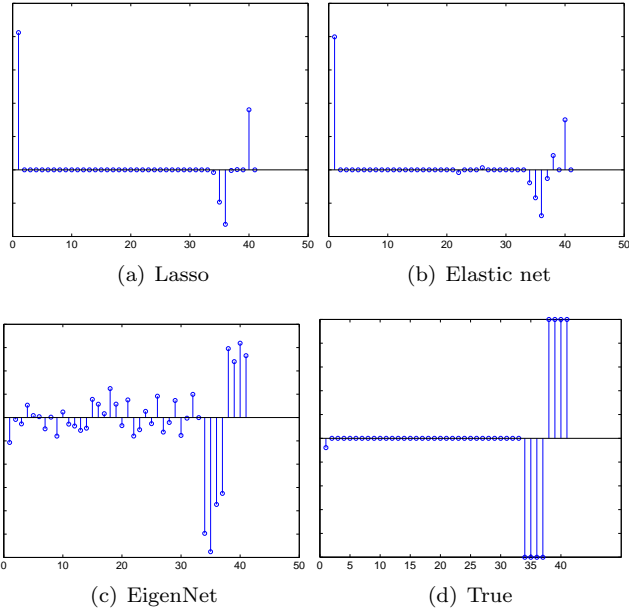


Figure 4: Visualization of the lasso, the elastic net, the EigenNet and the true classifier weights. These classifiers are estimated on 80 training samples with 40 features. Among the 40 features, 8 of them (as well as the bias) are relevant for the classification task. On this dataset the test error rates of the lasso, the elastic net, and the Bayesian lasso, the EigenNet are 0.297, 0.245, 0.251, and 0.137.

EigenNet. As shown in 5(a) and 5(b), while the Bayesian lasso suggests some correlation structures among features, they are fairly noisy. By contrast, the EigenNet shows the two groups of correlated features much more clearly.

6.2 Classification of synthetic data

Now we systematically compare these methods on synthetic datasets containing correlated features and datasets containing independent features. For this first case, we use a similar procedure as in the visualization example: we sample 40 dimensional data points, each of which contains two groups of correlated variables. The correlation coefficient between variables in each group is 0.81 and there are 4 variables in each group. However, unlike for the previous example where the classifier weights are the same for the correlated variables, now we set the weights within the same group to have the same sign, but with different random values. We vary the number of training points, ranging from 10 to 80, and test all these methods. For the datasets with independent features, we follow the same procedure except that the features are independently sampled.

We run the experiments 10 times. Figure 6 shows the error rates averaged over 10 runs. We do not plot the standard errors of the test error rates, since




Figure 5: Covariance matrices of the Bayesian lasso and the EigenNet classifiers. The covariance matrices are estimated based on the MCMC samples for these two models. We use 80 training samples with 40 features per sample. The covariance matrix of the EigenNet classifier correctly suggests the last few features are correlated. In particular, it clearly identifies a group of four correlated features.

they have very small values: the biggest one is less than 0.0183 for the results on data with correlated features, and for the results on data with independent features, the biggest one is less than 0.030. We report the numerical values of both the averaged error rates and the standard errors in the supplemental materials.

For the datasets with independent features, the EigenNet outperforms the alternative methods when the number of training samples are smaller than 40, the number of features (i.e., $p > n$). Since in this case the eigenstructures of the datasets are uninformative, we expect the improved prediction accuracy is the result of the subspace constraint used by the EigenNet. And once the number of training samples are not bigger than the data dimension, all these methods perform quite similarly.

For the datasets with correlated features, the EigenNet significantly outperforms the alternative methods consistently, not only when the number of training samples are smaller than 40 ($p > n$) but also when it is not. We believe this is because the EigenNet uses the valuable eigen information revealing the feature correlations to train its classifiers. Note that although the result of the elastic net appear to overlaps with those of the lasso. Actually for the data with correlated features, the elastic net often slightly outperforms the lasso (Please their numerical values in the supplemental materials).

6.3 Classification of real data

Besides the synthetic data, we also test all these methods on UCI benchmark datasets, two high-dimensional gene expression datasets, leukaemia and colon cancer, and a spambase dataset with relatively lower dimension but a lot more training samples.

For the leukaemia dataset, the task is to distinguish acute myeloid leukaemia

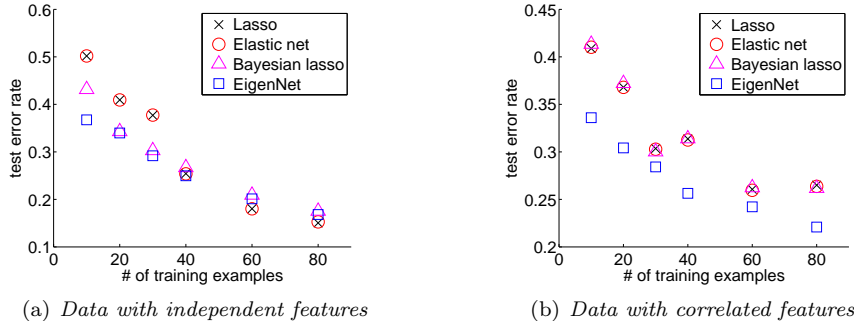


Figure 6: Test error rates on synthetic datasets with independent features and with correlated features. Each training sample has 40 features, 8 of which are relevant features. We increase the number of training samples from 10 to 80 and use 2000 test samples each time. The results are averaged over 10 runs. For the data with independent features, the EigenNet outperforms the alternative methods at beginning when the number of training samples are fewer than 40, the number of the features. With more training samples containing independent features, all these methods perform comparably. For data with correlated features, the EigenNet outperforms the alternative methods consistently.

(AML) from acute lymphoblastic leukaemia (ALL). The whole dataset has 47 and 25 samples of type ALL and AML respectively with 7129 features per sample. The dataset was randomly split 20 times into 37 training and 35 test samples.

For the colon cancer dataset, the task is to discriminate tumor from normal tissues using microarray data. The dataset has 22 normal and 40 cancer samples with 2000 features per sample. We randomly split the dataset into 31 training and 31 test samples 10 times.

For the spambase dataset, the task is to detect spam emails, i.e., unsolicited commercial emails. We use 57 features indicating whether a particular word or character was frequently occurring in the emails. We randomly split the dataset into 1533 training and 3066 test samples 10 times. Note that we do not use any kernel here and the results on this dataset are meant to examine how the performance of these methods compares to each other when there are more samples than features. Using a nonlinear basis function, e.g., a radial basis function, is expected to boost the predictive performance of all these methods.

Figure 7 summarizes the average test error rates and the standard errors of these methods on the three datasets. Again, the EigenNet significantly outperforms the alternative methods on three datasets. Note that for the leukaemia and colon cancer datasets Bayesian lasso does not perform much worse than the other methods. The reason, we believe, is that these two high dimensional datasets contain thousands of features and Bayesian lasso directly draws sam-

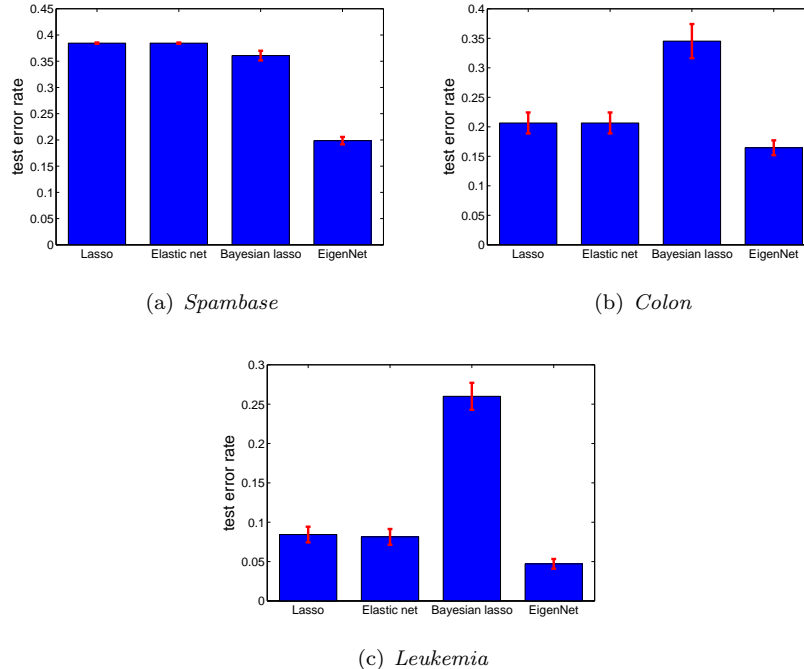


Figure 7: Test error rates on spambase, leukemia and colon cancer datasets. The error bars represent the standard errors of the error rates. The results on the spambase and colon cancer datasets are averaged over 10 random partitions and the results on the leukemia dataset are averaged over 20 partitions.

ples in such high dimensional spaces, leading to very slow mixing rates. By contrast, the EigenNet draws samples efficiently in a much smaller eigenspace, not only leading to faster mixing rates but also greatly saving the computing cost for obtaining each sample.

7 Conclusions

In this paper, we have presented a novel sparse Bayesian hybrid model, the EigenNet. It integrates a sparse conditional classification model with a generative model capturing the feature correlations. It also generalizes the elastic net by explicitly exploring correlations between features. Compared with several state-of-the-art methods, the EigenNet achieves significantly improved prediction accuracy on several benchmark datasets.

We plan to extend our hybrid model by utilizing other probabilistic generative models, such as sparse principle component analysis and related projection methods [Guan & Dy, Archambeau & Bach, 2009] and independent component

analysis models. Compared to the classical PCA models, these models could be used to better guide the selection of interdependent sparse features.

Acknowledgement

Thanks to Jyotishka Datta for his help on software implementation and to Tommi Jaakkola for stimulating discussion.

References

- Archambeau, Cédric and Bach, Francis. Sparse probabilistic projections. In *Advances in Neural Information Processing Systems 21*. 2009.
- Fan, Jianqing and Li, Runze. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- Frank, Ildiko E. and Friedman, Jerome H. A Statistical View of Some Chemometrics Regression Tools. *Technometrics*, 35(2):109–135, 1993.
- Guan, Yue and Dy, Jennifer. Sparse probabilistic principal component analysis. *JMLR W&CP: AISTATS*, 5.
- Hans, Chris. Bayesian lasso regression. *Biometrika*, 96(4):835–845, 2009.
- Lasserre, Julia A., Bishop, Christopher M., and Minka, Thomas P. Principled hybrids of generative and discriminative models. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 87–94, 2006.
- Li, Qing and Lin, Nan. The Bayesian Elastic Net. *Bayesian Analysis*, 5(1): 151–170, 2010.
- Park, Trevor, Casella, and George. The Bayesian Lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.
- Petersen, Kaare Brandt and Pedersen, Michael Syskind. The matrix cookbook, 2008. URL <http://matrixcookbook.com>.
- Sirovich, L. and Kirby, M. Low-dimensional procedure for the characterization of human faces. *J. Opt. Soc. Am. A*, 4(3):519–524, 1987.
- Tibshirani, Robert. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.
- Turk, Matthew and Pentland, Alex. Eigenfaces for recognition. *J. Cognitive Neuroscience*, 3:71–86, 1991.
- Zou, Hui and Hastie, Trevor. Regularization and variable selection via the Elastic Net. *Journal of the Royal Statistical Society B*, 67:301–320, 2005.

Appendix

Given the linear relationship between $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\mathbf{w}, \tilde{\mathbf{w}})$, the prior $p(\mathbf{w}, \tilde{\mathbf{w}})$ defined in (8) is equivalent to $p(\boldsymbol{\alpha}, \boldsymbol{\beta})$ defined in (10).

First, when $n \geq p$, we can easily obtain the $p(\boldsymbol{\alpha}, \boldsymbol{\beta})$ from $p(\mathbf{w}, \tilde{\mathbf{w}})$. In this case, the number of eigenvectors is p and the Jacobian matrix is the $p \times p$ full rank matrix \mathbf{V} . Furthermore, the determinant of \mathbf{V} is 1 since \mathbf{V} is an orthonormal matrix. Therefore, with $[\mathbf{w}, \tilde{\mathbf{w}}] = \mathbf{V}[\boldsymbol{\alpha}, \boldsymbol{\beta}]$ we have $p(\boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\mathbf{w}, \tilde{\mathbf{w}})$.

When $p > n$, $\mathbf{V}_{p \times n}$ is a tall matrix and therefore we cannot compute its determinant to transform the prior distribution $p(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Now $p(\mathbf{w}, \tilde{\mathbf{w}})$ is essentially a distribution on the data subspace embedded in the high dimensional space \mathcal{R}^p . To obtain the equivalence between these two priors, we consider the following theorem [Petersen & Pedersen, 2008]:

Theorem 1 *If \mathbf{A} is “tall”, i.e., “under-determined”, then $p(\mathbf{x}) = \int p(\mathbf{s})\delta(\mathbf{x} - \mathbf{A}\mathbf{s})d\mathbf{s} = \begin{cases} \frac{1}{\sqrt{|\mathbf{A}^T\mathbf{A}|}}p(\mathbf{A}^+\mathbf{x}) & \text{if } \mathbf{x} = \mathbf{A}\mathbf{A}^+\mathbf{x} \\ \mathbf{0} & \text{otherwise} \end{cases}$*

Using this theorem and the fact $|\mathbf{V}\mathbf{V}^+| = 1$, we see that with the simple linear relationship between the variables, $p(\boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\mathbf{w}, \tilde{\mathbf{w}})$.