

Model Selection Confidence Sets by Likelihood Ratio Testing

Chao Zheng¹, Davide Ferrari² and Yuhong Yang³

²*Lancaster University*, ²*University of Melbourne* and ³*University of Minnesota*

Abstract: The traditional activity of model selection aims at discovering a single model superior to other candidate models. In the presence of pronounced noise, however, multiple models are often found to explain the same data equally well. To resolve this model selection ambiguity, we introduce the general approach of model selection confidence sets (MSCSs) based on likelihood ratio testing. A MSCS is defined as a list of models statistically indistinguishable from the true model at a user-specified level of confidence, which extends the familiar notion of confidence intervals to the model-selection framework. Our approach guarantees asymptotically correct coverage probability of the true model when both sample size and model dimension increase. We derive conditions under which the MSCS contains all the relevant information about the true model structure. In addition, we propose natural statistics based on the MSCS to measure importance of variables in a principled way that accounts for the overall model uncertainty. When the space of feasible models is large, MSCS is implemented by an adaptive stochastic search algorithm which samples MSCS models with high probability. The MSCS methodology is illustrated through numerical experiments on synthetic data and real data examples.

Key words and phrases: Adaptive sampling; Likelihood ratio test; Model selection confidence set; Optimal detectability condition

1 Introduction

Likelihood inference is a centerpiece of statistical theory and plays an important role in many research fields. Numerous methods relying on likelihood objective functions have been de-

veloped in the literature of model selection, ranging from classic information criteria to more recent sparsity-inducing penalization methods; see McQuarrie and Tsai (1998), Claeskens and Hjort (2008) and Bühlmann and van de Geer (2011) for book-length expositions. In the presence of noise in data, however, it is typically difficult to declare a single model significantly superior to all possible competitors, due to the prevailing effect of the model selection uncertainty. In this situation, multiple or even a large number of models may be equally supported by data, so that any selection procedure is likely to pick at random a single model from a large set of more or less equivalent models. Clearly, this implies tossing away valuable information; for example, in regression analysis, alternative combinations of predictors may be discarded, whilst such combinations may contain scientifically valid explanations of the phenomenon under examination.

Motivated by the above issues, there has been a growing interest in developing statistical measures of model selection uncertainty. The approach followed in this paper proposes to construct a model selection confidence set (MSCS), defined as a set of models indistinguishable from the true model at a user-defined confidence level. Simply put, the MSCS extends the familiar frequentist notion of confidence intervals to the model-selection framework. Ferrari and Yang (2015) first introduced confidence sets for variable selection in the context of linear models by F-testing. They achieve the exact coverage probability for the globally optimal model from the model space. Thus, this is the first work introducing confidence sets in the frequentist sense for variable selection. Moreover, in their framework the number of predictors can grow with the sample size, so that the number of potentially useful models is allowed to be large.

Different from the MSCS of Ferrari and Yang (2015), Hansen, Lunde and Nason (2011)

studied the methodology of model confidence sets. Their approach builds on classic step-down procedures for multiple hypothesis testing (Lehmann and Romano, 2005; Romano and Wold, 2005). Starting from a fixed user-defined set of models, which has limited sizes, they carry out step-wise equivalence testing under a user-defined loss function, followed by an elimination rule to drop the worst performing models. Previously, Shimodaira (1998) constructed confidence sets containing models with AIC values near the smallest among the candidate models. We refer to the Section 6 of Ferrari and Yang (2015) for a detailed discussion.

In this paper, we introduce a general methodology to construct model selection confidence sets via likelihood inference. We begin by considering a full model with p variables to form a reference model or full model, where p is required to be less than n . This preliminary step can be achieved by any over-consistent model screening method, which selects the relevant variables plus a few other variables. We then test candidate sub-models against the full model, by a likelihood ratio test (LRT) at the significance level $0 < \alpha < 1$. The MSCS is formed by all the candidate models that survive the LRT screening. This way of construction guarantees that the globally optimal model is included in the MSCS with probability at least $1 - \alpha$ as the sample size increases (under appropriate regularity conditions). From a theoretical viewpoint, we investigate the condition for the MSCS to contain all the relevant information about the model structure when both p and n diverge. Since in practice the MSCS cannot be computed by exhaustive search unless p is very small, when the model space is moderate or large, we propose a stochastic algorithm (MSCS-AS) which samples MSCS models with high probability.

The proposed MSCS methodology can be used for various tasks in support of the model

selection activity. First, given a model selected from some rule, one can immediately use the MSCS to check if such a model is too parsimonious in terms of missing important variables. Second, the frequency of variables in the MSCS can be used to rank their usefulness in a principled way that accounts for the model selection uncertainty. Third, the MSCS and the associated importance measures may be used to narrow down the list of candidate models by considering the most important variables.

The rest of the paper is organized as follows. In Section 2, we describe the main MSCS methodology and study the condition needed to learn the true underlying model structure. In the same section, we propose a measure of importance for the individual variables. In Section 3, we give an adaptive sampling algorithm that implements the MSCS methodology. In Section 4, we study the finite sample properties of MSCS by Monte Carlo simulations for various models. In Section 5, we illustrate the MSCS procedure using the European E.coli outbreak data and the Australian breast cancer family study data. In Section 6, we conclude and give final remarks. Technical proofs are deferred to the Appendix.

2 Model selection confidence sets

Consider independent observations, $\mathbf{Y}_1, \dots, \mathbf{Y}_n$, from a family of models indexed by the parameter $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)^T \in \Theta$, with corresponding log-likelihood function $\ell_n(\mathbf{y}, \boldsymbol{\theta})$. Each parameter element $\theta_j \in \boldsymbol{\theta}$ describes a possibly relevant part of the overall model structure. We suppose that only a subset of $\boldsymbol{\theta}$ is useful for describing the data, while the others are regarded as unnecessary. A generic model index γ is defined as a subset of indexes in $\{1, \dots, p\}$ and we write the correspondent parameter space as Θ_γ . Denote $\boldsymbol{\theta}_\gamma \in \Theta_\gamma$ as a parameter with the model γ , and let $p_\gamma = \text{card}(\gamma)$ denotes the cardinality (number of

elements) of γ . The true parameter vector and the true model is denoted by $\boldsymbol{\theta}^*$ and γ^* , respectively, while the full model with p parameters is denoted by γ_f . The space of feasible candidate models is Γ which contains the true model γ^* . The cardinality of Γ may be as large as 2^p ; however it also may be restricted in some special problems.

In the rest of the paper, we assume $p < n$, but p is allowed to slowly grow with n , reflecting the notion that with more observations available, the statistician is tempted to introduce additional variables into the model. For simplicity, we omit the sub-index n when it is clear from the context. In what follows, we use “ \lesssim ” to denote that the left hand side is bounded by the right hand side up to some positive constant independent of n . We write $a \gtrsim b$ if $b \lesssim a$.

2.1 Construction by likelihood ratio testing

A MSCS is constructed from the known models space, Γ , and a criterion to assess models in Γ empirically. To screen out implausible models in the context of maximum likelihood estimation, it is natural to use the likelihood ratio test. Given a candidate model γ , we consider testing the null hypothesis $H_0 : \boldsymbol{\theta}^* \in \Theta_\gamma$ against the alternative hypothesis $H_1 : \boldsymbol{\theta}^* \notin \Theta_\gamma$. Then model γ is rejected if

$$\Lambda_\gamma \equiv 2 \left\{ \ell_n(\widehat{\boldsymbol{\theta}}_{\gamma_f}) - \ell_n(\widehat{\boldsymbol{\theta}}_\gamma) \right\} \geq q(\alpha; p - p_\gamma), \quad (2.1)$$

where: $\widehat{\boldsymbol{\theta}}_\gamma$ and $\widehat{\boldsymbol{\theta}}_{\gamma_f}$ denote, respectively, the MLEs for the candidate and full models; $\ell_n(\cdot)$ is the log-likelihood function; $q(\alpha; d)$ is the upper α -quantile for the central chi-squared distribution with d degree of freedom. The $(1 - \alpha)100\%$ -MSCS is defined by the set of all

models surviving the LRT screening:

$$\widehat{\Gamma}_\alpha \equiv \{\gamma \in \Gamma : \Lambda_\gamma \leq q(\alpha; p - p_\gamma)\}. \quad (2.2)$$

The LRT procedure is applied to all models $\gamma \in \Gamma$: If a model is rejected, then we have evidence that it is too parsimonious in the sense that it is likely to miss at least one important variable. By default, the full model γ_f is included in $\widehat{\Gamma}_\alpha$.

When p is fixed, and γ^* is a proper subset of γ , the limiting null distribution of the LRT statistics Λ_γ is a central chi-square distribution, which follows directly from Wilks theorem (e.g., see van der Vaart (2000)). By construction, this implies that the true model is in the MSCS with probability approximately $1 - \alpha$ in large samples. Specifically, if the true model is not the full model ($\gamma^* \neq \gamma_f$), we have:

$$\lim_{n \rightarrow \infty} P(\gamma^* \in \widehat{\Gamma}_\alpha) = 1 - \alpha. \quad (2.3)$$

If $\gamma^* = \gamma_f$, then $P(\gamma^* \in \widehat{\Gamma}_\alpha) = 1$.

When p increases with n , similar Wilks-type results are given by Portnoy (1988), Murphy (1993), Fan and Peng (2004) and Fan and Zhou (2016) for exponential family models, Cox regression, penalized likelihood and goodness of spurious fit for GLMs. These results yield asymptotic coverage probability as in (2.3). Spokoiny (2012, 2013) establish Wilks-type behaviours for rather general families of models, which quantify and explicitly describe the error term in the approximation of the likelihood ratio statistics under mild regularity conditions on the parametric family.

We remark that although the MSCS includes γ^* with at least probability $1 - \alpha$, one

cannot simply conclude that a variable is important just because it appears in some of the models in $\widehat{\Gamma}_\alpha$. Actually, unimportant variables tend to appear with a respectable frequency in the MSCS models since larger models containing the true model plus other irrelevant variables is likely to survive the LRT screening.

2.2 Asymptotic detectability

In this section, we study the conditions under which the variables in the true model appear with large frequency in MSCS. The results presented in this section extend the analysis given by Ferrari and Yang (2015) for linear models.

Definition 1 (Asymptotic detectability). The MSCS ($\widehat{\Gamma}_\alpha$) is said to asymptotically detect all the true variables, if all the variables in the true model γ^* are included in each of the models in $\widehat{\Gamma}_\alpha$, with probability going to 1.

The concept of detectability is closely related to the power of the LRT. In the fixed p scenario when a candidate model γ misses at least one important variable, under appropriate regularity conditions ensuring asymptotic normality of the MLE, the Λ_γ converges in distribution to a non-central chi-square random variable with degree of freedom $d_\gamma = p - p_\gamma$. Let $\boldsymbol{\theta}_\gamma^*$ denote the parameter value in the model γ that minimizes the Kullback-Leibler divergence from the true density (hence providing the best approximation to the true density). Then the non-centrality parameter of the asymptotic chi-square distribution is $\delta_\gamma = (\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*)^T \mathcal{F}(\boldsymbol{\theta}^*)(\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*)$, where the $\mathcal{F}(\boldsymbol{\theta}) = -E[\partial^2 \ell_n(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}^2]$ is the Fisher information matrix. In the normal regression case as in Ferrari and Yang (2015), the asymptotic distribution is exact. Clearly, when applying LRT to a model γ , a large value of δ_γ makes it easier to reject γ .

In the following theorem, we show that the limiting non-central chi-square alternative distribution is still valid under certain conditions for an exponential model where $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are i.i.d. observations from the pdf

$$f(\mathbf{y}; \boldsymbol{\theta}) = \exp [\boldsymbol{\theta}^T \mathbf{y} - A(\boldsymbol{\theta})] . \quad (2.4)$$

with respect to a sigma-finite dominating measure. For this model, the true model $\boldsymbol{\gamma}^*$ is defined as the indexes of non-zero component of $\boldsymbol{\theta}$. Appropriate generalizations of the following result may be derived for other models but they are not pursued in this paper.

Theorem 1. *Assume conditions (A1)–(A3) given in the Appendix. Let $\boldsymbol{\gamma}$ be a model missing at least one variable in the true model $\boldsymbol{\gamma}^*$ and $d_\boldsymbol{\gamma} \rightarrow \infty$ as n grows. Moreover, assume $\|\boldsymbol{\theta}^* - \boldsymbol{\theta}_\boldsymbol{\gamma}^*\| \gtrsim \sqrt{p/n}$ and $p = o(n^{2/3})$. Then for model (2.4) we have*

$$\frac{\Lambda_\boldsymbol{\gamma} - (d_\boldsymbol{\gamma} + \delta_\boldsymbol{\gamma})}{\sqrt{2d_\boldsymbol{\gamma} + 4\delta_\boldsymbol{\gamma}}} \xrightarrow{\mathcal{D}} \mathcal{N}_1(0, 1), \quad \text{as } n \rightarrow \infty. \quad (2.5)$$

Denote $X_{\delta,d}$ as a chi-square random variable with degree of freedom d and non-centrality parameter δ . Recall that $(X_{\delta,d} - d - \delta)/\sqrt{2(d + 2\delta)}$ converges to a standard normal distribution $\mathcal{N}_1(0, 1)$, when $d \rightarrow \infty$. This means that in view of (2.5), $\Lambda_\boldsymbol{\gamma}$ is approximately non-central chi-square variable with degree of freedom $d_\boldsymbol{\gamma}$ and non-centrality parameter $\delta_\boldsymbol{\gamma}$.

The non-centrality parameter $\delta_\boldsymbol{\gamma}$ may be interpreted as the discrepancy measure due to missing important variables in the true model. From this viewpoint, the relative magnitude of $\delta_\boldsymbol{\gamma}$ provides us with some insight on how informative is the data in relation to the feasibility of the model selection task. Let $K_n(s) = s \log(p/s)$. With $p \rightarrow \infty$, it is typically the case that the true model dimension is bounded away significantly from p . In the rest of the paper,

we assume that d_γ increases to ∞ (however slowly) uniformly for the candidate models. The following result gives explicit sufficient conditions involving δ_γ for detectability in the general parametric setting, which includes model (2.4) described in Theorem 2.2.

Theorem 2. *Let Γ_u denote the set of models missing at least one of the true variables.*

Suppose that it holds that for all $\gamma \in \Gamma_u$, we have

$$|P(\Lambda_\gamma \leq q(\alpha; k)) - P(X_{\delta_\gamma, k} \leq q(\alpha; k))| \leq c_1 \exp[-c_2 K_n(k)], \quad (2.6)$$

where c_1 and c_2 are positive constants. A sufficient condition for asymptotic detectability is

$$\min_{\gamma \in \Gamma_u} \frac{\delta_\gamma}{K_n(d_\gamma)} > B, \quad (2.7)$$

for some large enough positive constant B .

The additional assumption above requires an exponential probability bound for the chi-square approximation of LRT statistics with model misspecification. We refer to Theorem 3.10 and Proposition B.1 in Spokoiny (2013), for the chi-square approximation, where a similar bound can be achieved for certain i.i.d and regression models. Moreover, consider normal linear regression as in Ferrari and Yang (2015), $(RSS_\gamma - RSS_{\gamma_f})/d_\gamma$, the numerator of their F -test, follows an exactly non-central chi-square distribution where the assumption is trivially satisfied.

2.3 Sharpness of the sufficient condition for detectability

The above detectability condition theorem is a general extension of Theorem 2.3 in Ferrari and Yang (2015), where in the context of normal linear regression, a sufficient condition for detectability is given as $\min_{\gamma \in \Gamma_u} \delta_\gamma / \left\{ \xi_n + \sqrt{K_n(d_\gamma)} \right\}$ is greater than some large enough constant, where $\xi_n \rightarrow \infty$ is any arbitrarily slowly growing sequence. It turns out the condition is in fact not sufficient and an error occurred in their derivation. A correct sufficient condition is that $\min_{\gamma \in \Gamma_u} \delta_\gamma / K_n(d_\gamma)$ is larger than some constant, which matches (2.7) in this paper. In this subsection, we show the new sufficient condition cannot be generally improved. Due to space limitation and the need to correct Theorem 2.3 of Ferrari and Yang (2015), we focus on the normal regression case here. A generalization to other models, e.g. GLMs, can be done similarly with additional technical developments.

Clearly the detectability condition relates to the size of the coefficients. For the following results, we assume the sparse Riesz condition (SRC) (Zhang, 2010) holds and consider $0 < \alpha < 1/2$.

Let $r^* \leq p/2$ be a positive integer as the number of non-zero coefficients in the true model. Write $f_\beta(\mathbf{X}) = \mathbf{X}^T \beta = \sum_{j=1}^p \beta_j X_j$ and let $\mathcal{B} = \{\beta : \|\beta\|_0 = r^* \text{ and } \|f_\beta\|_n^2 \leq cK_n(r^*)\}$ for some small constant $c > 0$, where $\|\cdot\|_0$ denotes the ℓ_0 -norm, and $\|f_\beta\|_n^2 = \sum_{i=1}^n f_\beta^2(\mathbf{X}_i)$ with \mathbf{X}_i being the covariate vector for the i -th observation. It represents all linear regression models γ with only r^* non-zero coefficients.

Theorem 3. *Let \mathcal{D} denote the event that all the variables in γ^* are included in each of the models in the MSCS. Then when c is small enough, we must have*

$$\limsup_{n \rightarrow \infty} \inf_{\beta \in \mathcal{B}} P_\beta(\mathcal{D}) < 1.$$

From the theorem, for the true models of dimension r^* with $\|f_\beta\|_n^2 \leq cK_n(r^*)$, detection of the true terms is impossible in a proper minimax sense. Note that, for instance, for the model γ that contains all the wrong variables and none of the true variables, it results in $d_\gamma = r^*$ and the noncentrality parameter is of order $K_n(d_\gamma)$. This matches the lower bound requirement (2.7) in Theorem 2 in order. So from this aspect, the sufficient condition (2.7) for detectability cannot be generally weakened in order.

2.4 Inclusion importance

Under the detectability conditions in Theorem 2, the MSCS includes all the relevant information concerning the model selection variability. Thus, a natural measure for ranking the importance of each parameter element $\theta_k \in \boldsymbol{\theta}$ is its relative frequency over all the MSCS models. This suggests the following definition.

Definition 2 (Inclusion Importance). The inclusion importance (II) for any $\theta_k \in \boldsymbol{\theta}$ is defined as

$$II_k = \sum_{\gamma \in \widehat{\Gamma}_\alpha} I(k \in \gamma) / \text{card}(\widehat{\Gamma}_\alpha). \quad (2.8)$$

When θ_k appears in all MSCS models, its importance is $II_k = 1$, meaning that θ_k is most likely part of the true model.

As already mentioned, however, we note that a variable cannot be declared relevant just because it has a non-zero importance index. Actually, unimportant variable of $\boldsymbol{\theta}$ tend to appear in the MSCS with frequency near 1/2. The reason is that when a small model is included in $\widehat{\Gamma}_\alpha$, also larger models containing the same variables plus some others tend to be included via the LRT by construction. The following theorem describes an asymptotic behavior for the inclusion importance.

Theorem 4. *If the asymptotic detectability conditions in Theorem 2 are satisfied, we have:*

- (i) $\lim_{n \rightarrow \infty} P(II_k = 1) = 1$, for all $k \in \gamma^*$;
- (ii) $\lim_{n \rightarrow \infty} P\left(II_k > \frac{1}{2} + \Delta\right) \leq \frac{\alpha(1 + 2\Delta)}{4\Delta}$, for all $k \notin \gamma^*$, where $0 < \Delta < \frac{1}{2}$.

If we have sufficient information to learn all the the relevant variables of the true model, we expect that their importance to be close to 1, while the unimportant variables are not likely exceeding by much the value 0.5. The upper bound in Theorem 4 can be used as a guidance to control the error probability of over selection. For example, one can set the error probability $\epsilon = \alpha(1 + 2\Delta)/(4\Delta)$ to be some small number and then find the corresponding Δ so as to use II for an understanding if a variable is really important. For example, if the significance level is $\alpha = 0.05$, setting $\Delta = 1/6$ implies $\epsilon \leq 0.1$.

2.5 The multivariate normal location model

In this section we consider the special case of the multivariate normal distribution with unknown location. Let \mathbf{Y} follows the p -variate normal distribution $\mathcal{N}_p(\boldsymbol{\mu}, I_p)$. Then the pdf with form (2.4) can be obtained by setting the parameter vector as $\boldsymbol{\theta} = \boldsymbol{\mu}$ and the cumulant generating function is $A(\boldsymbol{\theta}) = \boldsymbol{\theta}^T \boldsymbol{\theta} / 2 + p \log(2\pi) / 2$. Assume the true parameter $\boldsymbol{\theta}^*$ is sparse with $p_{\gamma^*} = o(p)$. The model space Γ is then with cardinality 2^p .

For a misspecified model $\gamma \in \Gamma_u$, the corresponding non-centrality parameter is $\delta_\gamma = n \sum_{j \in \gamma^* \text{ and } j \notin \gamma} \boldsymbol{\theta}_j^{*2}$. A large value of δ_γ enables us to detect inadequacy of such models. For example, the asymptotic detectability conditions in (2.7) states that as long as the minimum signal is large enough, $\min\{|\theta_i|, i \in \gamma^*\} > B\sqrt{p/n}$ for a large enough positive constant B , then all the models in MSCS are expected to contain all the nonzero parameters with

probability going to 1 as $n \rightarrow \infty$. Otherwise if the size of some non-zero parameters in θ^* is too small, then the LRT has not enough power to screen some wrong models out.

3 Implementation by adaptive sampling

Testing all the models in Γ is computationally challenging unless p is small, since the cardinality of the model space may grow exponentially in p . Thus, in order to find models in the MSCS it seems natural to turn to sampling methods. Let $\mathbf{U} = (u_1, \dots, u_p)^T \sim p(\mathbf{u}; \boldsymbol{\omega})$ be a random binary vector representing a model sampled from Γ ($u_j = 1$ if the j th variable is included in the model), and $p(\cdot; \boldsymbol{\omega})$ is a user-defined pmf indexed by $\boldsymbol{\omega}$. Our main objective is to choose a value of the parameter $\boldsymbol{\omega}$ that maximizes the probability to sample MSCS models

$$P(\mathbf{U} \in \hat{\Gamma}_\alpha) = \sum_{\mathbf{u} \in \Gamma} p(\mathbf{u}; \boldsymbol{\omega}) I(\mathbf{u} \in \hat{\Gamma}_\alpha). \quad (3.1)$$

Note one is unlikely to find models in $\hat{\Gamma}_\alpha$ just by sampling from some arbitrary pmf $p(\cdot; \boldsymbol{\omega})$, unless α is sufficiently small. Thus, given a target significance level $\alpha = \alpha^*$ (e.g. 0.05) we propose to start from some small initial confidence level, say $\alpha^{(0)}$, and then construct a sequence of significance levels, $0 < \alpha^{(0)} \leq \alpha^{(1)} \leq \dots \leq \alpha^*$, corresponding to sampling distributions $p(\cdot; \boldsymbol{\omega}^{(0)}), p(\cdot; \boldsymbol{\omega}^{(1)}), \dots, p(\cdot; \boldsymbol{\omega}^{(*)})$ increasingly concentrated on the target subspace $\hat{\Gamma}_{\alpha^*}$.

At each step $t \geq 0$ of our algorithm, the parameter $\boldsymbol{\omega}$ of the sampling distribution is retrieved by the following weighted likelihood approach. We generate B models $\{\mathbf{u}_b^{(t)}, b =$

$1, \dots, B\}$ from $p(\cdot; \widehat{\boldsymbol{\omega}}^{(t-1)})$ and then compute

$$\widehat{\boldsymbol{\omega}}^{(t)} = \underset{\boldsymbol{\omega}}{\operatorname{argmax}} \sum_{b=1}^B I(\mathbf{u}_b^{(t)} \in \widehat{\Gamma}_{\alpha^{(t-1)}}) p(\mathbf{u}_b^{(t)}; \boldsymbol{\omega}). \quad (3.2)$$

This finds the pmf $p(\cdot; \widehat{\boldsymbol{\omega}}^{(t)})$ closest to the best subset of previously sampled models in terms of their resemblance to MSCS models. As t increases and $\alpha^{(t)}$ gradually gets closer to α^* , $p(\cdot; \widehat{\boldsymbol{\omega}}^{(t)})$ tends to assign larger probability to models in $\widehat{\Gamma}_{\alpha}$.

Since this procedure is useful only when the indicator $I(\mathbf{u}_b^{(t-1)} \in \widehat{\Gamma}_{\alpha^{(t-1)}}) = 1$ for a sufficiently large fraction of sampled models, $\mathbf{u}_1^{(t-1)}, \dots, \mathbf{u}_B^{(t-1)}$, we propose to increase adaptively the significance level as $\alpha^{(t)} = \min\{\text{p-val}_{\lfloor (1-\zeta)B \rfloor}^{(t)}, \alpha^*\}$, $0 < \zeta < 1$, where $\text{p-val}_{\lfloor (1-\zeta)B \rfloor}^{(t)}$ is the empirical $(1 - \zeta)$ -quantile computed from the distribution of p-values. This ensures that the event $\{\boldsymbol{\gamma}_b^{(t)} \in \widehat{\Gamma}_{\alpha^{(t)}}\}$ is not too rare and occurs with probability of approximately ζ . The proposed approach is closely related to cross-entropy (CE) sampling. See Rubinstein and Kroese (2004) for a book-length exposition on this topic and Costa, Jones and Kroese (2007) for convergence analysis. In our practical implementation, we use $p(\mathbf{u}; \boldsymbol{\omega}) = \prod_{j=1}^p \omega_j^{u_j} (1 - \omega_j)^{1-u_j}$, which gives a closed-form solution to (3.2) and leads to a fast algorithm; all our numerical experiments showed reliable results with relatively fast convergence. Other choices for $p(\mathbf{u}; \boldsymbol{\omega})$ may enhance the performance of the algorithm, but they are not pursued here. The following steps outline the stochastic procedure for MSCS construction.

First, note that Step 4 carries out smoothing at each iteration; if $\xi = 1$, the algorithm avoids smoothing. In our simulations, we found that $\xi < 1$ performs better than the non-smooth update with $\xi = 1$ since it prevents occurrences of too many zeros and ones in situations where p is moderate or large. Smoothing avoids local optima where some model variables do not have the chance to be selected, while others are always selected. In our

Algorithm 1 : MSCS construction by adaptive sampling (MSCS-AS)

0. Initialize $t = 0$ (iteration counter) and $\widehat{\boldsymbol{\omega}}^{(0)}$ (parameter vector for pmf $p(\mathbf{u}, \boldsymbol{\omega})$).
1. Set $t \leftarrow t + 1$. Generate $S^{(t)} = \{\mathbf{u}_1^{(t)}, \dots, \mathbf{u}_B^{(t)}\}$ from $p(\cdot; \widehat{\boldsymbol{\omega}}^{(t-1)})$, and compute the sorted p-values, $\text{p-val}_{(1)}^{(t)} \leq \dots \leq \text{p-val}_{(B)}^{(t)}$, by the LRT defined in (2.1).
2. Update $\alpha^{(t)} = \min\{\text{p-val}_{\lfloor (1-\zeta)B \rfloor}^{(t)}, \alpha^*\}$.

3. Use models sampled in Step 1, maximize the weighted likelihood as in(3.2) by computing:

$$c_j^{(t)} = \frac{\sum_{b=1}^B I\{\text{p-val}_{(b)}^{(t)} > \alpha^{(t)}, \widehat{\theta}_j \in \mathbf{u}_b^{(t)}\}}{\sum_{b=1}^B I\{\text{p-val}_{(b)}^{(t)} > \alpha^{(t)}\}}, \quad j = 1, \dots, p,$$

where $\{\theta_j \in \mathbf{u}\}$ denotes the event that the variable θ_j appears in model \mathbf{u} .

4. Update $\widehat{\omega}_j^{(t)} \leftarrow \xi c_j^{(t)} + (1 - \xi)\widehat{\omega}_j^{(t-1)}$ for some constant $0 < \xi < 1$.
 5. Repeat Steps 1–4 until $\alpha^{(t-d)} = \dots = \alpha^{(t)} = \alpha^*$, for some d (e.g. $d = 10$). The final MSCS is obtained by drawing $B^{(T)}$ models from $p(\mathbf{u}; \widehat{\boldsymbol{\omega}}^{(T)})$, where T denotes the last iteration.
-

experience, the MSCS-AS algorithm is robust to the choice of ξ , with $\xi = 0.2$ performing well across all our numerical examples.

Second, the MSCS-AS algorithm requires setting the initial weights $\widehat{\boldsymbol{\omega}}^{(0)}$, and the number of models sampled at each iteration B . We found that the procedure is quite robust to the choices of such parameters. When no prior information on inclusion importance is available, the initial probabilities $\widehat{\boldsymbol{\omega}}^{(0)}$ can be set as $\widehat{\boldsymbol{\omega}}^{(0)} = (0.5, \dots, 0.5)$. The performance of the method, however, can be improved by assigning larger weights to variables that are known to contain more information about the true model. The number of models B generated in each iteration should be decided based on affordable computational resources. However, if B is too small this will affect the accuracy of the weighted likelihood criterion (3.2). In all our numerical examples we set $B = 300$.

Finally, the constant ζ prevents overly small p-values in the first few iterations; thus,

it ensures a balanced growth of $\alpha^{(t)}$ and guides the sampling process towards the MSCS models. The parameter ζ governs the trade-off between exploration and exploitation of the model space Γ and it should be also fixed based on the available computational resources. In our simulations, $\zeta = 0.25$ is found to work well and is compatible with choices of analogous parameters often found in the CE literature.

4 Monte Carlo experiments

4.1 MSCS construction by exhaustive search

In this subsection, we study the finite sample properties of MSCSs constructed by exhaustive search on the model space Γ . We generate samples from the following four models:

Model 1: p -variate normal with unknown location, $\mathbf{Y} = (Y_1, \dots, Y_p)^T \sim \mathcal{N}_p(\boldsymbol{\theta}, I)$.

Model 2: p -variate normal $\mathbf{Y} = (Y_1, \dots, Y_p)^T \sim \mathcal{N}_p(0, \Sigma)$, with unknown covariance matrix Σ . Additionally we assume that (Y_1, \dots, Y_p) can be partitioned into independent subvectors ($\Sigma_{i,j} = 0$ if Y_i and Y_j belong to different subvectors).

Model 3: Logistic regression: $Y|\mathbf{x} \sim \text{Bernoulli}(\pi(\mathbf{x}))$, $\text{Logit}(\pi(x)) = -\mathbf{x}^T \boldsymbol{\theta}$, $\boldsymbol{\theta} \in \mathbb{R}^p$.

Model 4: Poisson regression: $Y|\mathbf{x} \sim \text{Poisson}(\lambda(\mathbf{x}))$, $\log(\lambda(\mathbf{x})) = -\mathbf{x}^T \boldsymbol{\theta}$, $\boldsymbol{\theta} \in \mathbb{R}^p$.

In Models 3 and 4, the vectors of covariates are sampled from multivariate normal distribution $\mathcal{N}_p(0, I)$ at each Monte Carlo run. For all the above models we consider the following two scenarios describing the relative size of parameters.

Setting 1 (constant parameter size): The first $p/2$ parameters have the same size and the others are equal to 0. Specifically, $\theta_j = \psi$, $j = 1, \dots, p/2$, and $\theta_j = 0$,

$j = p/2 + 1, \dots, p$, where the parameters size ψ is set to be 1 for Models 1 and 3 and 0.2 for Models 4, respectively. For Model 2, (Y_1, \dots, Y_p) is partitioned into $p/2 + 1$ groups as $\{Y_1, \dots, Y_{p/2}\}, \{Y_{p/2+1}\}, \dots, \{Y_p\}$. We set $\Sigma_{i,j} = 1$ if $i = j$; $\Sigma_{i,j} = 0.5$, $1 \leq i < j \leq p/2$, and $\Sigma_{i,j} = 0$ otherwise.

Setting 2 (decreasing parameter size): The first $p/2$ coefficients have decreasing size and the others are equal to 0. Specifically, $\theta_j = \psi/j$, $j = 1, \dots, p/2$, and $\theta_j = 0$, $j = p/2 + 1, \dots, p$, where ψ is set as 1, 2 and 0.4 for Models 1, 3 and 4, respectively. For Model 2, (Y_1, \dots, Y_p) is partitioned into $p/2 + 1$ groups as $\{Y_1, \dots, Y_{p/2}\}, \{Y_{p/2+1}\}, \dots, \{Y_p\}$. We set $\Sigma_{i,j} = 1$ if $i = j$; $\Sigma_{i,j} = 0.5/|i - j|$ if $1 \leq i < j \leq p/2$, and $\Sigma_{i,j} = 0$ otherwise.

The above settings are designed to achieve a small signal relative to the noise, so that the resulting data are affected by model-selection uncertainty. The model space consists of 2^p models for Models 1, 3 and 4. For Model 2, the model space is equivalent to all the possible partitions of the set $\{1, 2, \dots, p\}$ and its cardinality can be given by the Bell number B_p . In the following simulations, we use $p = 6, 8$ corresponding to $B_6 = 203$ and $B_8 = 4140$.

Tables 1 and 2 show Monte Carlo estimates for the coverage probability and cardinality of MSCS corresponding to different sample sizes, n , number of predictors, p , at the 90, 95 and 99% confidence levels. As one expects, the cardinality of the MSCS grows as α decreases, while it increases rapidly with p , especially when the sample size n is relatively small. This reflects the situation where the data contain too much noise and the subsequent model selection variability is pronounced. The cardinality of the MSCS drops quickly as n increases.

In most cases, the true coverage probability is quite close to the nominal confidence level.

		Setting 1				Setting 2				
		$n =$	100	250		100	250			
	α				Model 1					
	$p =$	8	12	8	12	8	12	8	12	
Coverage (%)	0.10	91.2	90.0	89.8	90.2	91.2	89.4	91.0	88.6	
	0.05	94.4	95.8	94.4	96.2	94.8	95.4	95.0	94.4	
	0.01	98.4	99.8	99.2	98.4	99.4	98.8	99.0	99.4	
Cardinality	0.10	14.5	58.1	14.4	58.3	22.7	241.3	15.5	101.3	
	0.05	15.3	61.0	15.2	61.2	27.2	319.0	16.8	126.5	
	0.01	15.8	63.5	15.9	63.4	37.1	497.8	19.4	184.6	
					Model 2					
	$p =$	6	8	6	8	6	8	6	8	
Coverage (%)	0.10	89.2	88.4	89.2	91.0	89.0	87.6	88.4	88.8	
	0.05	94.6	94.8	94.2	94.4	94.4	93.2	94.4	94.4	
	0.01	98.4	97.8	98.8	99.4	98.4	99.0	98.6	99.2	
Cardinality	0.10	13.5	47.1	13.5	47.0	13.8	57.1	13.4	46.4	
	0.05	14.3	50.7	14.2	47.3	15.1	72.7	14.2	49.2	
	0.01	15.6	57.9	14.8	46.9	18.8	126.6	14.8	51.5	

Table 1: Monte Carlo estimates of MSCS coverage probability and cardinality by exhaustive search for Model 1 and Model 2 under varying confidence level $1 - \alpha$, sample size n , and number of variables p . Results based on 500 MC runs.

And will be improved in general as the sample sizes increases. We note that the true coverage probability tends to be more off from the nominal level when the size of the true parameters is decreasing (Setting 2). Clearly, in such settings model selection is more challenging, which leads to a increased cardinality of the MSCS but maintains the same coverage probability.

4.2 MSCS construction by stochastic search

In this subsection, we study the performance of the MSCS-AS algorithm described in Section 3. We generate data from Models 3 and 4 (Poisson and Logistic regression models) using a setting similar to that in Fan and Lv (2011). For both models, we set $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^T, 0, 0, \dots, 0)^T$,

	$n =$	Setting 1				Setting 2			
		100		250		100		250	
		$p =$	8	12	8	12	8	12	8
	α	Model 3							
Coverage(%)	0.10	86.6	80.0	86.6	87.0	86.6	83.6	87.4	89.2
	0.05	92.4	89.6	94.6	91.2	93.0	91.8	93.3	94.6
	0.01	97.6	97.4	99.2	99.0	96.6	97.6	98.6	98.6
Cardinality	0.10	17.5	97.3	14.3	56.8	35.4	459.5	18.2	195.3
	0.05	20.8	147.8	15.2	60.9	43.6	612.6	21.5	252.4
	0.01	32.4	257.0	15.9	67.0	62.0	941.4	29.2	383.6
		Model 4							
Coverage(%)	0.10	89.0	90.6	90.4	89.8	90.6	89.2	89.0	90.4
	0.05	95.6	94.2	95.2	94.6	95.2	94.2	94.6	95.0
	0.01	99.4	99.4	99.0	99.0	99.8	99.4	98.6	99.0
Cardinality	0.10	85.2	871.6	24.6	148.7	75.6	1269.3	47.9	597.8
	0.05	109.6	1217.8	32.4	210.3	94.1	1633.6	51.0	796.7
	0.01	157.9	1987.0	54.8	402.1	124.9	2177.4	71.3	1167.7

Table 2: Monte Carlo estimates of MSCS coverage probability and cardinality by exhaustive search for Logistic regression (Model 3) and Poisson regression (Model 4) under varying confidence level $1 - \alpha$, sample size n , and number of variables p . Results based on 500 MC runs.

where $\boldsymbol{\theta}_1 = (2.5, -1.9, 2.8, -2.2, 3)^T$ in Model 3 and $\boldsymbol{\theta}_1 = (1.25, -0.95, 0.9, -1.1, 0.6)^T$ in Model 4. The vector of covariates are sampled from a multivariate normal distribution $\mathcal{N}_p(0, \Sigma)$ at each Monte Carlo run, where Σ has elements $\Sigma_{i,j} = 0.5^{|i-j|}$, $i, j = 1, 2, \dots, p$. For illustration of capability to handle large p , we show the results for (n, p) equal to $(200, 100)$ and $(1000, 500)$. The initial weights for the MSCS-AS algorithm are $\widehat{\boldsymbol{\omega}}^{(0)} = (0.5, \dots, 0.5)$, corresponding to lack of prior information about predictors' importance. The remaining tuning parameters are set as $\zeta = 0.25$, $\alpha^* = 0.05$, and $\xi = 0.2$.

Figure 1 shows the trajectories for the importance weights $\{\omega_j\}_{j=1}^p$ during the first 50 iterations of the algorithm at the 95% confidence level. In all the considered cases, the

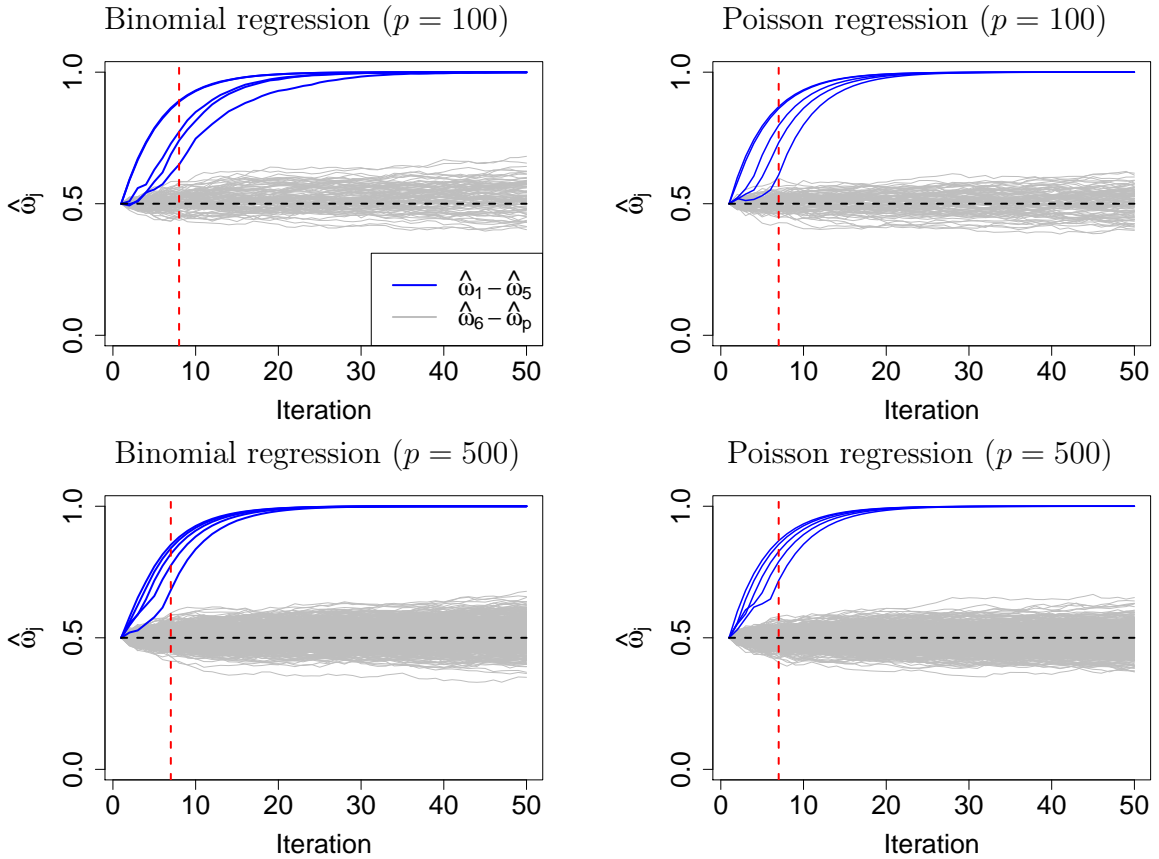


Figure 1: Sampling weights in $p(\cdot; \boldsymbol{\omega})$ for 50 iterations of the MSCS-AS algorithm at the 95% confidence level. The vertical dashed line corresponds to iteration t such that $\alpha^{(t)} = \alpha^*$. Left and right panels correspond to Binomial and Poisson regression models described in Section 4. Settings for the algorithm: $B = 300$, $\zeta = 0.25$, $\alpha^* = 0.05$, and $\xi = 0.2$.

trajectories corresponding to terms with non-zero coefficients are clearly distinguished from the others after a few iterations. The inclusion importance for the relevant terms increases to around 1, while the others only have importance weights near 0.5. The graphs show that the MSCS-AS algorithm samples with probability progressively concentrating on the true model terms, while unimportant terms are sampled quite randomly. This behaviour mimics the structure of the true MSCS, $\hat{\Gamma}_\alpha$, thus enabling us to detect MSCS models at a much cheaper computational cost than exhaustive search on Γ .

As suggested in Section 3 (Step 5 of the algorithm), we stop updating the importance weights shortly after the sequence of significance levels $\alpha^{(0)}, \alpha^{(1)}, \dots$ reaches the target sig-

nificance level α^* . However, after $\alpha^{(t)}$ reaches α^* , the weights of the irrelevant predictors eventually converge to 0 or 1 according to Kolmogorov’s zero-one law, thus one should stop before that happens. For example, stopping shortly after $\alpha^{(t)} = \alpha^* = 0.05$ – say around 15 iterations – already enables us to detect useful predictors from the rest.

To illustrate that the MSCS-AS algorithm generates MSCS models with large probability, we stop at iteration 15 and sample 10^6 models using $p(\cdot; \hat{\omega}^{(15)})$, where $\hat{\omega}^{(15)}$ is the importance weight at the 15th iteration. For the binomial regression model with $p = 100$ predictors, 80.5% of the models generated are included in MSCS. For the Poisson regression model with $p = 100$ predictors, 76.6% of the total models generated are in the MSCS. In comparison, if we generate 10^6 models using the uninformative weights $\omega^{(T)} = (0.5, \dots, 0.5)$, the proportion of MSCS models is basically 0, due to the largeness of the model space.

5 Real data examples

Example 1: European Escherichia coli(E.coli) O104:H4 outbreak data. In this example, we apply the MSCS methodology to the E.coli data as described in Edwards, Pope and Holt (2016). E.coli O104:H4 is a particularly aggressive pathogen and caused a serious outbreak in northern Germany in 2011 (Rasko et al., 2011). Both during and after the outbreak, scientists have examined the genome of E.coli to find genetic causes for the severity of the outbreak. The data set used here consists of 56 outbreak isolates. For each isolate, 10 genes (or hypothetical genes) in the O104:H4 pangenome (a full collection of genes in a species of bacteria) that have been identified as might have been associated with the outbreak are considered. The main goal of our analysis is to select a model which can explain the most meaningful interaction effects between those genes.

The presence of genes in E.coli is denoted by binary variables taking values 1 when the correspondent gene is present and 0 otherwise. Let $\mathbf{Y} = (Y_1, \dots, Y_{10})$ be a random variable with 10 binary variables each denoting the activity of a particular gene. The pmf of \mathbf{Y} is modelled by the Ising Model,

$$P(\mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}) = \exp \left(\sum_{1 \leq j < k \leq 10} \theta_{j,k} y_j y_k + \psi(\boldsymbol{\theta}) \right), \quad (5.1)$$

where $\boldsymbol{\theta} = (\theta_{j,k})_{1 \leq j < k \leq 10}$ is the parameter of interest with $p = 55$ and $\psi(\boldsymbol{\theta})$ is the normalizing constant. The variable $\theta_{j,j}$ is regarded as the main effect for gene j , whilst $\theta_{j,k}$ is interpreted as an interaction effect between genes j and k . Here, we wish to choose $\theta_{j,k} \neq 0$ if genes j and k have interaction (in the same group) and $\theta_{j,k} = 0$ otherwise. In our analysis, we assume $\theta_{j,j} \neq 0$ for $1 \leq j \leq 10$, meaning that the main effects are always included. The total number of possible models in Γ is $B_{10} = 115975$, corresponding to the 10-th Bell number, which counts the number of different ways to partition a set containing 10 elements.




Figure 2: Inclusion Importance (II) for $\theta_{j,k}$ in model (5.1). The II values are calculated from the 95%-MSCS.

Here we use the exhaustive search to construct the MSCS for θ . The MSCSs at the 90, 95 and 99% confidence levels contain 7, 12 and 38 models, respectively; these numbers are small compared to the model space size, meaning that most of the models in Γ are rejected by the LRT procedure and there is not too much model selection uncertainty here. The jk -th element of the matrix in Figure 2 represents the inclusion importance (II) for the variable $\theta_{j,k}$. Note that certain gene pairs have high II values; for example, pairwise interactions among genes 1141, 3614, 5245, 0051, 0904, 0901 are close to 1 in terms of II values, suggesting that such genes form a synergetic network associated with the outbreak occurrence. Other genes, such as 6065, 0233, 4619, 4633, show inclusion importance for interaction effects close or smaller than 0.5, which suggests that the corresponding interactions are small or irrelevant.

Example 2: Australian breast cancer family study data. In the second example, we apply the MSCS methodology to the ABCFS genotype data, consisting of 356 observations (284 breast cancer patients and 72 controls). Cases are obtained from the Australian Breast Cancer Family Study (ABCFS) (Dite et al., 2003), while controls are from the Australian Mammographic Density Twins and Sisters Study by Odefrey et al. (2010). Patients are genotyped using a Human610-Quad beadchip array. The response is the binary disease status (presence/absence of breast cancer), while the predictors are 50 SNPs, measured at different loci encoding a candidate susceptibility pathway (probe IDs are listed in Figure 3). To model the binary disease status, we use a logistic regression model. MSCS models are sampled using the MSCS-AS algorithm described in Section 3 with tuning parameters $\alpha^* = 0.05$, $\zeta = 0.25$, $B = 300$ and $B^{(T)} = 10^6$.

Figure 3 (top) shows the models selected by forward step-wise AIC and BIC (F-AIC and F-BIC), and penalized likelihood methods under Lasso (Tibshirani, 1996), SCAD (Fan




Figure 3: Analysis of the ABCFS case-control genotype data. Top: model selection by penalized likelihood methods with Lasso, SCAD and MCP penalties and step-wise forward AIC and forward BIC. Colored cells denote selected variables. Bottom: Estimate of II by MSCS-AS algorithm (red points) and the 95% bootstrap confidence intervals. The number of bootstrap replicates is 50. Settings for the MSCS-AS algorithm parameters: $B = 300$, $B^{(T)} = 10^6$, $\alpha^* = 0.05$, $\zeta = 0.25$, $\xi = 0.2$.

and Li, 2001) and MCP (Zhang, 2010) penalties where tuning parameters are all chosen by five-fold cross-validation. Figure 3 (bottom) shows Inclusion Importance, II values, for each SNP. The vertical bars represent the 95% bootstrap confidence intervals for II values. First note that SNPs with high II values show considerable overlap with those selected by the other methods. Particularly, most of the SNPs reported as important by more than one model selection method have large II values. Predictors with II values that are not significantly larger than 0.5 may not necessarily be important, and need to be considered more carefully.

The p-values of the F-test for the AIC, BIC, Lasso, SCAD and MCP models are 0.96, 0.09, 0.88, 0.86 and 0.80. respectively. Therefore, while the BIC model is included in the MSCS at the 95% and 99% levels, it is not accepted at the 90% confidence level. This suggests that F-BIC model may miss some important SNP predictors and is not as reliable as other models. Moreover, note that there are only 5 SNPs (rs0082248_A, rs12470143_T, rs2754530_T, rs8192207_T, rs2257157_G) that have relatively large II values and confidence interval significantly above 0.5. Many predictors chosen by some of the methods have II confidence intervals covering 0.5. This suggests that the sample is not sufficiently informative to declare such terms relevant so they should be further studied with particular care. To confirm this, we investigate the marginal significance of all the selected SNPs in each of those models. At the 0.05 level, we have 8 (Lasso, SCAD and MCP) to 12 (AIC) significant coefficients, while the 5 SNPs with II interval not including 0.5 are significant in all cases.

Finally, we show the instability of common model selection methods for this dataset. We consider the selected models as in Figure 3 (top). For each model we obtain the fitted values $\{\hat{Y}_i\}_{i=1}^n$, which are equal to $\{\hat{p}\}_{i=1}^n$, the estimated probability in the logistic regression.

	Lasso	SCAD	MCP	F-AIC	F-BIC
Size	21	16	15	15	5
AHD	14.74	11.132	10.44	11.58	5.59

Table 3: The number of SNPs in the original selected model and the average Hamming distance between the 500 bootstrapped models and the original selected model.

Next, parametric bootstrap is used to generate bootstrap replicates $\{X_i, Y_i^*\}_{i=1}^n$, and all the methods are applied again to corresponding bootstrap samples. We repeat this step for $S = 500$ times and compute the average hamming distances (AHD) of the S bootstrapped Lasso, SCAD, MCP, F-AIC and F-BIC models to the the respective models obtained from the original sample, see Table 3. The AHD is large compared the size of the models, which means for each bootstrapped sample, those model selection methods will choose quite different predictors.

6 Conclusion and final remarks

The MSCS methodology in this paper introduces new tools supporting the activity of model selection in the context of likelihood-based inference. Since the MSCS is asymptotically guaranteed to contain the true model at a pre-specified confidence level, it represents a natural extension of the familiar notion of confidence intervals to the model selection framework. Furthermore, Theorem 4 suggests that important variables tend to appear in the MSCS models with large probability as $n \rightarrow \infty$, while unimportant terms appear randomly with frequency not significantly larger than 0.5. By looking at the variables appearing frequently in the MSCS one can also choose a single central model representing the entire MSCS by taking predictors with inclusion importance significantly larger than 0.5. In the future, developing a theoretical understanding of the optimal way to combine MSCS models would be

very valuable as it can potentially lead to improved model combining and model selection strategies.

The main focus of the current MSCS approach is based on maximum likelihood estimation. We have shown that in exponential family models this requires $p = o(n^{2/3})$ and correct model specification for the MSCS to be meaningful. In the future, however, higher-dimensional problems may be pursued, by replacing the LRT statistics with other tools to construct the MSCS, e.g., using penalized likelihood methods. Computational methods to tackle the case where the model space and MSCS is large is also of great interest.

Appendix: Proofs

Notice that quantities such as p , $\boldsymbol{\theta}$, $\boldsymbol{\gamma}$ and Θ may depend on n , hence array asymptotics are considered in this section.

Proof of Theorem 2.2. Since within the considered exponential family affine mappings are preserved, without loss of generality we assume $E(\mathbf{Y}_i) = A'(\boldsymbol{\theta}^*) = 0$ and $\text{cov}(\mathbf{Y}_i) = A''(\boldsymbol{\theta}^*) = I_p$.

Denote $\mathbf{V}_\theta = \mathbf{Y}_\theta - E(\mathbf{Y}_\theta)$, where $\mathbf{Y}_\theta \sim f(\mathbf{y}; \boldsymbol{\theta})$. First, we need to assume following regularity conditions for the exponential family model $f(\cdot; \boldsymbol{\theta})$:

$$(A1) \quad EY_{ij}^6 < \infty, \quad (j = 1, \dots, p);$$

$$(A2) \quad \sup_{\substack{\|a\|=1 \\ \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|^2 \lesssim p/n}} |E(a^T \mathbf{V}_\theta)^3| \lesssim \sqrt{n/p}, \text{ and } \sup_{\|a\|=1} |E(a^T \mathbf{V}_\theta)^3| = O(1);$$

$$(A3) \quad \sup_{\substack{\|a\|=1 \\ \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|^2 \lesssim p/n}} E(a^T \mathbf{V}_\theta)^4 = O(1).$$

Suppose we have $\boldsymbol{\gamma}$ which is a model not containing all the elements in $\boldsymbol{\gamma}^*$. Let $\bar{\mathbf{Y}} =$

$\sum_{i=1}^n \mathbf{Y}_i/n$ and $\bar{\mathbf{Y}}_\gamma = \sum_{i=1}^n \mathbf{Y}_{i,\gamma}/n$, where $\mathbf{Y}_{i,\gamma}$ denote the vector with elements equal to \mathbf{Y}_i for indexes in γ and zero otherwise.

Let $\mathbf{G} = \boldsymbol{\theta}^* - \boldsymbol{\theta}_\gamma^*$ satisfying $\|\mathbf{G}\| \gtrsim \sqrt{p/n}$ and assume $p = o(n^{2/3})$, where $\boldsymbol{\theta}_\gamma^*$ is the p_γ -variate vector with components equal to $\boldsymbol{\theta}^*$ at $\gamma \cap \gamma^*$ and equal to zero otherwise. . To proceed, we give following auxiliary lemmas .

Lemma 1. *Suppose conditions (A2) and (A3) hold. Then, for the MLEs, $\hat{\boldsymbol{\theta}}_\gamma$, we have:*

$$\|\hat{\boldsymbol{\theta}}_\gamma - \boldsymbol{\theta}_\gamma^*\| = O_p\left(\sqrt{p_\gamma/n}\right), \text{ and } \|\hat{\boldsymbol{\theta}}_\gamma - \boldsymbol{\theta}_\gamma^* - \bar{\mathbf{Y}}_\gamma\| = O_p(p_\gamma/n).$$

Lemma 1 combines Theorems 2.1 and 3.1 in Portnoy (1988). The existence of MLE for $\boldsymbol{\theta}_\gamma$ in a $L-2$ neighbourhood of order $\sqrt{p_\gamma/n}$ is still valid by simply changing the true model $\boldsymbol{\theta}^*$ with the partial model $\boldsymbol{\theta}_\gamma^*$.

Lemma 2. *Assume (A2) and (A3) hold, for model γ , $|A(\boldsymbol{\theta}_\gamma^*) - A(\boldsymbol{\theta}^*)| = \|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*\|^2/2 + o(\|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*\|^2)$.*

Proof. Note that from (A3), by Taylor expansion we have:

$$|A(\boldsymbol{\theta}_\gamma^*) - A(\boldsymbol{\theta}_n^*)| = \left| \frac{1}{2} \|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*\|^2 + \frac{1}{6} E [(\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}_n^*)^T \mathbf{V}_{\boldsymbol{\theta}^*}]^3 \right| + O(\|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}_n^*\|^4).$$

Applying (A2) gives $E [(\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*)^T \mathbf{V}_{\boldsymbol{\theta}^*}]^3 \lesssim \|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*\|^3$, which completes the proof. \square

Lemma 3. *Suppose conditions (A1), (A2) and (A3) hold, when $p - p_\gamma \rightarrow \infty$ we have*

$$\frac{n (\|\bar{\mathbf{Y}}\|^2 - \|\bar{\mathbf{Y}}_\gamma\|^2) + 2n\mathbf{G}^T\bar{\mathbf{Y}} - (p - p_\gamma)}{\sqrt{2(p - p_\gamma + 2\delta_n)}} \xrightarrow{\mathcal{D}} \mathcal{N}_1(0, 1), \quad (6.1)$$

where $\delta_n = n\|\mathbf{G}\|^2$.

Proof. Define $\mathbf{T}_k = \sum_{i=1}^k \mathbf{Y}_i$, $\mathbf{T}_{k,\gamma} = \sum_{i=1}^k \mathbf{Y}_{i,\gamma}$ and $S_k = \|\mathbf{T}_k\|^2 - \|\mathbf{T}_{k,\gamma}\|^2 + 2k\mathbf{G}^T\mathbf{T}_k - k(p - p_\gamma)$. It is easy to see that $E(S_n) = 0$ and $\text{var}(S_n) = 2n^2(p - p_\gamma + 2n\|\mathbf{G}\|^2)$.

Let $D_k = S_k - S_{k-1}$, thus

$$D_k = 2\mathbf{Y}_k^T\mathbf{T}_{k-1} + \|\mathbf{Y}_k\|^2 - 2\mathbf{Y}_{k,\gamma}^T\mathbf{T}_{\gamma,k-1} - \|\mathbf{Y}_{k,\gamma}\|^2 + 2\mathbf{G}^T\mathbf{T}_{k-1} + 2k\mathbf{G}^T\mathbf{Y}_k - (p - p_\gamma).$$

Next, define $\sigma_k^2 = ED_k^2$ and $s_k^2 = \sum_{i=1}^k \sigma_i^2$. Note that \mathbf{Y}_k and \mathbf{T}_{k-1} are independent. A simple calculation shows that:

$$s_n^2 = \sum_{k=1}^n \sigma_k^2 \lesssim n^2p + n^3\|\mathbf{G}\|^2, \quad (6.2)$$

Next, let $\mathcal{F}_k = \mathcal{F}(\mathbf{Y}_1, \dots, \mathbf{Y}_k)$ denote the σ -field generated by $\mathbf{Y}_1, \dots, \mathbf{Y}_k$. Then $\{S_k\}$ are martingales on $\{\mathcal{F}_k\}$, and $\{D_k\}$ are the martingale differences. From Chow and Teicher (1978), by Martingale Central Limit Theorem, we have $S_n / (n\sqrt{p - p_\gamma + \delta_\gamma}) \xrightarrow{\mathcal{D}} \mathcal{N}_1(0, 1)$ if

$$\sum_{k=1}^n E|D_k|^3/s_n^3 \rightarrow 0, \quad \text{and} \quad \sum_{k=1}^n E|E(D_k^2|\mathcal{F}_{k-1}) - \sigma_k^2|/s_n^2 \rightarrow 0. \quad (6.3)$$

Since $E(Y_{ij}^6) < \infty$, by Proposition A.3 in Portnoy (1988), we have $E(\mathbf{Y}_k^T\mathbf{T}_{k-1})^6 \lesssim k^3p^3$ and therefore $\sum_{k=1}^n E|D_k|^3 \lesssim n^{5/2}p^{3/2} + n^2p^{5/2} + n^4\|\mathbf{G}\|^3$,

Together with Equation (6.2) implies $\sum_{k=1}^n E|D_k|^3/s_n^3 \rightarrow 0$.

Next, note that

$$\sum_{k=1}^n E|E(D_k^2|\mathcal{F}_{k-1}) - \sigma_k^2| \leq \left\{ E \left[E(D_k^2|\mathcal{F}_{k-1}) - \sigma_k^2 \right]^2 \right\}^{1/2} \quad (6.4)$$

$$\lesssim n^{3/2}p^{3/2} + n^2p^{1/2} + n^{3/2}\|\mathbf{G}\|^2 \quad (6.5)$$

Thus, we have $\sum_{k=1}^n E|E(D_k^2|\mathcal{F}_{i-1}) - \sigma_k^2|/s_n^2 \rightarrow 0$. \square

Now we are able to prove the main results.

Proof of Theorem 2.2. Under the alternatives, we have

$$\begin{aligned} \Lambda_\gamma &= 2n(\widehat{\boldsymbol{\theta}}_{\gamma_f} - \widehat{\boldsymbol{\theta}}_\gamma)^T \bar{\mathbf{Y}} - 2n \left[A(\widehat{\boldsymbol{\theta}}_{\gamma_f}) - A(\widehat{\boldsymbol{\theta}}_\gamma) \right] \\ &= 2n \left[(\widehat{\boldsymbol{\theta}}_{\gamma_f} - \boldsymbol{\theta}^*) - (\widehat{\boldsymbol{\theta}}_\gamma - \boldsymbol{\theta}_\gamma^*)^T \right] \bar{\mathbf{Y}} - 2n \left[A(\widehat{\boldsymbol{\theta}}_{\gamma_f}) - A(\boldsymbol{\theta}_\gamma^*) - A(\widehat{\boldsymbol{\theta}}_\gamma) + A(\boldsymbol{\theta}_\gamma^*) \right] \\ &\quad - 2n \left[(\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*)^T \bar{\mathbf{Y}} - A(\boldsymbol{\theta}_\gamma^*) + A(\boldsymbol{\theta}^*) \right]. \end{aligned} \quad (6.6)$$

The first two terms of the LRT statistic in (6.6) are approximately $n(\|\bar{\mathbf{Y}}\|^2 - \|\bar{\mathbf{Y}}_\gamma\|^2) + O_p(p^2/n)$. Additionally, Lemma 2 implies that the third term in (6.6) is approximately equals to $-2n(\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*)^T \bar{\mathbf{Y}} + n\|\boldsymbol{\theta}_\gamma^* - \boldsymbol{\theta}^*\|^2$. Hence, applying Lemma 3 completes the proof. \square

Proof of Theorem 2. We show that if the sufficient condition in Theorem 2 is satisfied, we have

$$P \left(\bigcup_{\gamma \in \Gamma_u} \{ \Lambda_\gamma \leq q(\alpha; d_\gamma) \} \right) \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

An union bound of the above probability is:

$$P \left(\bigcup_{\gamma \in \Gamma_u} \{ \Lambda_\gamma \leq q(\alpha; d_\gamma) \} \right) \leq \sum_{k=1}^{p-1} \sum_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} P(\Lambda_\gamma \leq q(\alpha; k)) \quad (6.7)$$

$$\leq \sum_{k=1}^{p-1} \exp[K_n(k)] \max_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} P(\Lambda_\gamma \leq q(\alpha; k)). \quad (6.8)$$

Note that we assume the exponential bound as in (2.6). Thus, combining Lemma 8.1 in

Birgé (2001) and Theorem A in Inglot (2010) gives the following probability upper bound:

$$\begin{aligned} \max_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} P(\Lambda_\gamma(d_\gamma, \delta_\gamma) \leq q(\alpha; k)) &\leq \max_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} P(X_{\delta_\gamma, k} \leq q(\alpha, k)) + c_1 \exp[-c_2 K_n(k)] \\ &\leq \exp \left[- \min_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} \frac{(\delta_\gamma + 2 \log(\alpha) - 2\sqrt{-k \log(\alpha)})^2}{2(k + 2\delta_\gamma)} \right] \\ &\quad + c_1 \exp[-c_2 K_n(k)] \end{aligned} \quad (6.9)$$

where c_1, c_2 are positive constants. When $\min_{\substack{d_\gamma=k \\ \gamma \in \Gamma_u}} \frac{\delta_\gamma}{K_n(k)} > B$ for some large enough positive constant B , the first term in the last upper bound (6.9) will be small than $\exp[-(c_3-1)K_n(k)]$ for some $c_3 > 1$. Therefore, we have

$$P \left(\bigcup_{\gamma \in \Gamma_u} \{\Lambda_\gamma \leq q(\alpha; d_\gamma)\} \right) \lesssim \sum_{k=1}^{p-1} \exp[-\min(c_2, c_3 - 1)K_n(k)] \rightarrow 0. \quad (6.10)$$

□

Proof of Theorem 3 Let $\widehat{\gamma}$ be the largest model that are nested in all the models in MSCS. It may be the intercept only model. Note that with probability at least $1 - \alpha$ the true model γ^* is included in the MSCS, i.e., $P(\gamma^* \in \widehat{\Gamma}_\alpha) \geq 1 - \alpha$. Therefore

$$\sup_{\beta \in \mathcal{B}} P_\beta(\widehat{\gamma} \neq \gamma^*) \leq \sup_{\beta \in \mathcal{B}} P_\beta(\gamma^* \notin \widehat{\Gamma}_\alpha) + \sup_{\beta \in \mathcal{B}} P_\beta(\mathcal{D}^C) \quad (6.11)$$

$$\leq \alpha + 1 - \inf_{\beta \in \mathcal{B}} P_\beta(\mathcal{D}). \quad (6.12)$$

It follows that if $\inf_{\beta \in \mathcal{B}} P_\beta(\mathcal{D}) \rightarrow 1$ for any subsequence of $\{n_j\} \subset \{1, 2, \dots\}$, we must have $\sup_{\beta \in \mathcal{B}} P_\beta(\widehat{\gamma} \neq \gamma^*) \leq \alpha'$ for some $\alpha < \alpha' < 1/2$ when n_j is large enough in the subsequence.

Hence if we can actually show $\sup_{\beta \in \mathcal{B}} P_\beta(\hat{\gamma} \neq \gamma^*) \geq 1/2$ for a small enough $c > 0$ in the definition of \mathcal{B} , the theorem is proved.

Note that this now becomes a traditional minimax framework where Fano's inequality can be applied. Without loss of generality, we assume $\sigma^2 = 1$ for the error variance. Consider a packing set N_{ϵ_n} in $\mathcal{B} = \{\beta : \|\beta\|_0 = r^* \text{ and } \|f_\beta\|_n^2 \leq cK_n(r^*)\}$ with packing distance ϵ_n being a small fraction of $\sqrt{cK_n(r^*)}$ under the $\|\cdot\|_n$ norm. Let β be randomly chosen from the uniform distribution on N_{ϵ_n} . Using similar arguments as in the proof of Theorem 11 in Wang, et al. (2014), under the SRC, by choosing the constant c small enough, the mutual information $\mathbf{I}(\beta; \{X_i, Y_i\}_{i=1}^n)$ between the random β and the observations is upper bounded by $\delta = cK_n(r^*)/2$, and the local packing ϵ_n -entropy $\log |N_{\epsilon_n}|$ is lower bounded by $2\delta + 2 \log 2$. Apply Fano's inequality (see, e.g., Yang and Barron, 1999) to this linear regression model gives:

$$\sup_{\beta \in \mathcal{B}} P_\beta(\hat{\gamma} \neq \gamma^*) \geq 1 - \frac{\mathbf{I}(\beta; \{X_i, Y_i\}_{i=1}^n) + \log 2}{\log |N_{\epsilon_n}|} \geq \frac{1}{2}, \quad (6.13)$$

this completes the proof. \square

Proof of Theorem 4 (i) Let γ be a model missing at least one variable in γ^* . From Theorem 2, $P(\gamma \in \hat{\Gamma}_\alpha) \rightarrow 0$, as $n \rightarrow \infty$. Hence, for all θ_j in θ^* , $\lim_{n \rightarrow \infty} P(II_j = 1) = 1$. This completes the first part of the theorem.

(ii) Let $\tilde{\Gamma} = \{\gamma_1, \gamma_2, \dots, \gamma_N\}$ be the set of models larger than γ^* . The construction of MSCS implies $\lim_{n \rightarrow \infty} P(\gamma_i \in \hat{\Gamma}_\alpha) \geq 1 - \alpha$ for any $1 \leq i \leq N$. Let X_i be the random variable taking value 1 when $\gamma_i \in \hat{\Gamma}_\alpha$ and 0 otherwise. Note that when $k \notin \gamma^*$, there are $N/2$ models in $\tilde{\Gamma}$ which contain θ_k . Without loss of generality, let $\gamma_1, \gamma_2, \dots, \gamma_{N/2}$ be models containing k . Let $Y_1 = \sum_{i=1}^{N/2} X_i$, denoting the number of models in $\{\gamma_1, \gamma_2, \dots, \gamma_{N/2}\}$ that are included

in $\widehat{\Gamma}_\alpha$. Similarly, let $Y_2 = \sum_{i=N/2+1}^N X_i$. Then $Y_1 \leq N/2$ and $E(Y_2) \geq N(1 - \alpha)/2$. The II_k defined in Section 2.4 equals to $Y_1/(Y_1 + Y_2)$ with probability going to 1 as $n \rightarrow \infty$. Note that for any $0 < \Delta \leq 1/2$ we have

$$\begin{aligned} P\left(\frac{Y_1}{Y_1 + Y_2} \geq \frac{1}{2} + \Delta\right) &\leq P\left(\frac{N/2}{N/2 + Y_2} \geq \frac{1}{2} + \Delta\right) \\ &= P\left(N/2 - Y_2 \geq \frac{2\Delta N}{1 + 2\Delta}\right) \\ &\leq \frac{(N/2 - E(Y_2))(1 + 2\Delta)}{2N\Delta} \\ &\leq \frac{\alpha(1 + 2\Delta)}{4\Delta}, \end{aligned}$$

where the second inequality follows from Markov's inequality, which completes the proof. \square

Acknowledgements

We sincerely thank the two reviewers and the AE for their very helpful comments and suggestions for improving our work.

References

- Birgé, L. (2001). In *State of the Art in Probability and Statistics. IMS Lecture Notes*. **37**, 113–133.
- Bühlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, New York.
- Chow, Y. S. and Teicher, H. (1978). *Probability Theory: Independence, Interchangeability, Martingales*. Springer, New York.
- Claeskens, G. and Hjort, N. L. (2008). *Model Selection and Model Averaging*. Cambridge University Press, Cambridge.
- Costa, A., Jones, O. D. and Kroese, D. (2007). Convergence properties of the cross-entropy method for discrete

-
- optimization. *Oper. Res. Lett.* **35**, 573–580.
- Dite, G.S., et al. (2003). Familial risks, early-onset breast cancer, and brca1 and brca2 germline mutations. *J. Natl. Cancer Inst.* **95**, 448–457.
- Edwards, D., Pope, B.J. and Holt, K.E. (2016). Reddog tutorial. *GitHub technical report*, Web: <https://github.com/katholt/RedDog/blob/master/docs/RedDogTutorial.pdf>.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96**, 1348–1360.
- Fan, J. and Lv, J. (2011). Nonconcave penalized likelihood with np-dimensionality. *IEEE Trans. Inform. Theory.* **57**, 5467–5484.
- Fan, J. and Peng, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *Ann. Statist.* **32**, 928–961.
- Fan, J. and Zhou, W.-X. (2015). Guarding from spurious discoveries in high dimension. *J. Mach. Learn. Res.* **17**, 1–34.
- Ferrari, D. and Yang, Y. (2015). Confidence sets for model selection by F -testing. *Statist. Sinica.* **25**, 1637–1658.
- Hansen, P.R., Lunde, A. and Nason, J.M. (2011). The model confidence set. *Econometrica.* **79**, 453–497.
- Heritier, S. and Ronchetti, E. (1994). Robust bounded-influence tests in general parametric models. *J. Amer. Statist. Assoc.* **89**, 897–904.
- Inglot, T. (2010). Inequalities for quantiles of the chi-square distribution. *Probab. Math. Statist.* **30**, 339–351.
- Lehmann, E.L. and Romano, J.P. (2005). *Testing Statistical Hypotheses*. 3rd edition. Springer, New York.
- Mcquarrie, A.D. and Tsai, C.-L. (1998). *Regression and Time Series Model Selection*. World Scientific, Singapore.
- Murphy, S.A. (1993). Testing for a time dependent coefficient in cox regression model. *Scand. J. Stat.* **20**, 35–50.
- Odefrey, F. et al. (2010). Common genetic variants associated with breast cancer and mammographic density measures

-
- that predict disease. *Cancer Res.* **70**, 1449–1458.
- Portnoy, S. (1988). Asymptotic behavior of likelihood methods for exponential families when the number of parameters tends to infinity. *Ann. Statist.* **16**, 356–366.
- Rasko, D. et al. (2011). Origins of the e. coli strain causing an outbreak of hemolytic-uremic syndrome in germany. *N. Engl. J. Med.* **365**, 709–717.
- Romano, J.P. and Wolf, M. (2005). Exact and approximate stepdown methods for multiple hypothesis testing. *J. Amer. Statist. Assoc.* **100**, 94–108.
- Rubinstein, R. Y. and Kroes, D.P. (2004). *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. Springer-Verlag, New York.
- Shimodaira, H. (1998). An application of multiple comparison techniques to model selection. *Ann. Inst. Statist. Math.* **50**, 1–13.
- Spokoiny, V. (2012). Parametric estimation. finite sample theory. *Ann. Statist.* **40**, 2877–2909.
- Spokoiny, V. (2013). Bernstein-von mises theorem for growing parameter dimension. Available at [arXiv:1302.3430](https://arxiv.org/abs/1302.3430).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58**, 267–288.
- Van der Vaart, A.W. (2000). *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- Wang, Z., Paterlini, S., Gao, F. and Yang, Y. (2014). Adaptive minimax regression estimation over sparse l_q -hulls. *J. Mach. Learn. Res.* **15**, 1675–1711.
- Yang, Y. and Barron, A. (1999). Information-theoretic determination of minimax rates of convergence. *Ann. Statist.* **27**, 1564–1599.
- Zhang, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38**, 894–942.

Chao Zheng

Department of Mathematics and Statistics, Lancaster University

E-mail: c.zheng5@lancaster.ac.uk

Davide Ferrari

School of Mathematics and Statistics, University of Melbourne

E-mail: dferrari@unimelb.edu.au

Yuhong Yang

School of Statistics, University of Minnesota

E-mail: yyang@stat.umn.edu