

# Model Selection Techniques

*An overview*

©ISTOCKPHOTO.COM/GREMLIN

---

Jie Ding, Vahid Tarokh, and Yuhong Yang

In the era of big data, analysts usually explore various statistical models or machine-learning methods for observed data to facilitate scientific discoveries or gain predictive power. Whatever data and fitting procedures are employed, a crucial step is to select the most appropriate model or method from a set of candidates. Model selection is a key ingredient in data analysis for reliable and reproducible statistical inference or prediction, and thus it is central to scientific studies in such fields as ecology, economics, engineering, finance, political science, biology, and epidemiology. There has been a long history of model selection techniques that arise from researches in statistics, information theory, and signal processing. A considerable number of methods has been proposed,

following different philosophies and exhibiting varying performances. The purpose of this article is to provide a comprehensive overview of them, in terms of their motivation, large sample performance, and applicability. We provide integrated and practically relevant discussions on theoretical properties of state-of-the-art model selection approaches. We also share our thoughts on some controversial views on the practice of model selection.

## **Why model selection**

Vast developments in hardware storage, precision instrument manufacturing, economic globalization, and so forth have generated huge volumes of data that can be analyzed to extract useful information. Typical statistical inference or machine-learning procedures learn from and make predictions on data by fitting parametric or nonparametric models (in a broad

---

*Digital Object Identifier 10.1109/MSP.2018.2867638*  
*Date of publication: 13 November 2018*

sense). However, there exists no model that is universally suitable for any data and goal. An improper choice of model or method can lead to purely noisy discoveries, severely misleading conclusions, or disappointing predictive performances. Therefore, a crucial step in a typical data analysis is to consider a set of candidate models (referred to as the *model class*) and then select the most appropriate one. In other words, model selection is the task of selecting a statistical model from a model class, given a set of data. We may be interested, e.g., in the selection of

- variables for linear regression
- basis terms, such as polynomials, splines, or wavelets in function estimation
- order of an autoregressive (AR) process
- number of components in a mixture model
- most appropriate parametric family among a number of alternatives
- number of change points in time series models
- number of neurons and layers in neural networks
- best choice among logistic regression, support vector machine, and neural networks
- best machine-learning techniques for solving real-world data challenges on an online competition platform.

There have been many overview papers on model selection scattered in the communities of signal processing [1], statistics [2], machine learning [3], epidemiology [4], chemometrics [5], and ecology and evolution [6]. Despite the abundant literature on model selection, existing overviews usually focus on derivations, descriptions, or applications of particular model selection principles. In this article, we aim to provide an integrated understanding of the properties and practical performances of various approaches by reviewing their theoretical and practical advantages, disadvantages, and relations.

## Some basic concepts

### Notation

We use  $\mathcal{M}_m = \{p_{\theta_m} : \theta_m \in \mathcal{H}_m\}$  to denote a model (in the formal probabilistic sense), which is a set of probability density functions to describe the data  $z_1, \dots, z_n$ . Here,  $\mathcal{H}_m$  is the parameter space associated with  $\mathcal{M}_m$ . A model class,  $\{\mathcal{M}_m\}_{m \in \mathbb{M}}$ , is a collection of models indexed by  $m \in \mathbb{M}$ . The number of models (or the cardinality of  $\mathbb{M}$ ) can be fixed or depend on the sample size  $n$ . For each model  $\mathcal{M}_m$ , we denote by  $d_m$  the dimension of the parameter in model  $\mathcal{M}_m$ . Its log-likelihood function is written as  $\theta_m \mapsto \ell_{n,m}(\theta_m) = \log p_{\theta_m}(z_1, \dots, z_n)$ , and the maximized log-likelihood value is

$$\ell_{n,m}(\hat{\theta}_m), \text{ with } \hat{\theta}_m = \operatorname{argmax}_{\theta_m \in \mathcal{H}_m} p_{\theta_m}(z_1, \dots, z_n), \quad (1)$$

the maximum likelihood estimator (MLE) under model  $\mathcal{M}_m$ . We will write  $\ell_{n,m}(\hat{\theta}_m)$  as  $\hat{\ell}_{n,m}$  for simplicity. We use  $p_*$  and  $E_*$  to

**Vast developments in hardware storage, precision instrument manufacturing, economic globalization, and so forth have generated huge volumes of data that can be analyzed to extract useful information.**

denote the true data-generating distribution and expectation with respect to the true data-generating distribution, respectively. In the parametric framework, there exists some  $m \in \mathbb{M}$  and some  $\theta_* \in \mathcal{H}_m$  such that  $p_*$  is exactly  $p_{\theta_*}$ . In the nonparametric framework,  $p_*$  is excluded in the model class. We sometimes call a model class  $\{\mathcal{M}_m\}_{m \in \mathbb{M}}$  *well-specified* (respectively, *misspecified*) if the data-generating process is in a parametric (respectively nonparametric) framework. We use  $\rightarrow_p$  and  $\rightarrow_d$  to denote convergence in

probability and in distribution (under  $p_*$ ), respectively. We use  $\mathcal{N}(\mu, V)$  to denote a Gaussian distribution of mean  $\mu$  and covariance  $V$ ,  $\chi_d^2$  to denote a chi-squared distribution with  $d$  degrees of freedom, and  $\|\cdot\|_2$  to denote the Euclidean norm. The word *covariate* is often referred to as the *covariate* in a regression setting.

A typical data analysis can be thought of as consisting of two steps.

- *Step 1:* For each candidate model  $\mathcal{M}_m = \{p_{\theta_m}, \theta_m \in \mathcal{H}_m\}$ , fit all of the observed data to that model by estimating its parameter  $\theta_m \in \mathcal{H}_m$ .
- *Step 2:* Once we have a set of estimated candidate models  $p_{\hat{\theta}_m}(m \in \mathbb{M})$ , select the most appropriate one for either interpretation or prediction.

We note that not every data analysis and its associated model selection procedure formally rely on probability distributions. Examples of model-free methods are nearest-neighbor learning, certain reinforcement learning, and expert learning. Before we proceed, it is helpful to first introduce the following two concepts: the model fitting and the best model.

### The model fitting

The fitting procedure (also called *parameter estimation*) given a certain candidate model  $\mathcal{M}_m$  is usually achieved by minimizing the following (cumulative) loss:

$$\tilde{\theta}_m = \operatorname{arg min}_{\theta_m \in \mathcal{H}_m} \sum_{t=1}^n s(p_{\theta_m}, z_t). \quad (2)$$

In (2), each  $p_{\theta_m}$  represents a distribution for the data, and  $s(\cdot, \cdot)$ , referred to as the *loss function* (or *scoring function*), is used to evaluate the goodness of fit between a distribution and the observation. A commonly used loss function is the logarithmic loss

$$s(p, z_t) = -\log p(z_t), \quad (3)$$

the negative logarithm of the distribution of  $z_t$ . Then, (2) produces the MLE for a parametric model. For time series data, (3) is written as  $-\log p(z_t | z_1, \dots, z_{t-1})$ , and the quadratic loss  $s(p, z_t) = \{z_t - E_p(z_t | z_1, \dots, z_{t-1})\}^2$  is often used, where the expectation is taken over the joint distribution  $p$  of  $z_1, \dots, z_t$ .

### The best model

Let  $\hat{p}_m = p_{\hat{\theta}_m}$  denote the estimated distribution under model  $\mathcal{M}_m$ . The predictive performance can be assessed via the out-sample prediction loss, defined as

$$E_*(s(\hat{p}_m, Z)) = \int s(\hat{p}_m(z), z)p_*(z) dz, \quad (4)$$

where  $Z$  is independent with and identically distributed as the data used to obtain  $\hat{p}_m$ . Here,  $Z$  does not have the subscript  $t$  as it is the out-sample data used to evaluate the predictive performance. There can be a number of variations to this in terms of the prediction loss function [8] and time dependency. In view of this definition, the best model can be naturally defined as the candidate model with the smallest out-sample prediction loss, i.e.,

$$\hat{m}_0 = \arg \min_{m \in \mathbb{M}} E_*(s(\hat{p}_m, Z)).$$

In other words,  $\mathcal{M}_{\hat{m}_0}$  is the model whose predictive power is the best offered by the candidate models. We note that the best is in the scope of the available data, the class of models, and the loss function.

In a parametric framework, typically the true data-generating model, if not too complicated, is the best model. In this vein, if the true density function  $p_*$  belongs to some model  $\mathcal{M}_m$  or, equivalently,  $p_* = p_{\theta_*}$  for some  $\theta_* \in \mathcal{H}_m$  and  $m \in \mathbb{M}$ , then we seek to select such  $\mathcal{M}_m$  (from  $\{\mathcal{M}_m\}_{m \in \mathbb{M}}$ ) with probability going to one as the sample size increases, which is called *consistency in model selection*. In addition, the MLE of  $p_{\theta_m}$  for  $\theta_m \in \mathcal{H}_m$  is known to attain Cramer–Rao lower bound asymptotically. In a nonparametric framework, the best model depends on the sample size—typically the larger the sample size, the larger the dimension of the best model because more observations can help reveal weak variables (whose effects are relatively small) that are out of reach at a small sample size. As a result, the selected model is sensitive to the sample size, and selection consistency becomes statistically unachievable. We revisit this point in the “Because All Models Are Wrong, Why Pursue Consistency in Selection?” section.

We note that the aforementioned equivalence between the best model and the true model may not hold for regression settings where the number of independent variables is large relative to the sample size. Here, even if the true model is included as a candidate, its dimension may be too high to be appropriately identified based on relatively small data. Then the parametric framework becomes practically nonparametric. We will emphasize this point in the “An Illustration on Fitting and the Best Model” section.

### Goals of data analysis and model selection

There are two main objectives in learning from data. One is for scientific discovery, understanding of the data-generation process, and interpretation of the nature of the data. A scientist, e.g., may use the data to support a physical model or identify genes that clearly promote early onset of a disease. Another objective of learning from data is for prediction, i.e., to quantitatively describe future observations. Here the data

scientist does not necessarily care about obtaining an accurate probabilistic description of the data. Of course, one may also be interested in both directions.

In tune with the two different objectives, model selection can also have two directions: model selection for inference and model selection for prediction. The first one is intended to identify the best model for the data, which hopefully provides a reliable characterization of the sources of uncertainty for scientific insight and interpretation. And the second is to choose a model as a vehicle to arrive at a model or method that offers top performance. For the former goal, it is crucially important that the selected model is not too sensitive to the sample size. For the latter, however, the selected model may simply be the lucky winner among a few close competitors, yet the predictive performance can still be (nearly) the best possible. If so, the model selection is perfectly fine for the second goal (prediction), but the use of the selected model for insight and interpretation may be severely unreliable and misleading. Associated with the first goal of model selection for inference or identifying the best candidate is the following concept of selection consistency.

#### Definition 1

A model selection procedure is consistent if the best model is selected with probability going to one as  $n \rightarrow \infty$ . In the context of variable selection, in practical terms, model selection consistency is intended to mean that the important variables are identified and their statistical significance can be ascertained in a follow-up study of a similar sample size but the rest of the variables cannot. In many applications, prediction accuracy is the dominating consideration. Even when the best model as defined earlier cannot be selected with high probability, other models may provide asymptotically equivalent predictive performance. The following asymptotic efficiency property demands that the loss of the selected model or method is asymptotically equivalent to the smallest among all of the candidates.

#### Definition 2

A model selection procedure is asymptotically efficient if

$$\frac{\min_{m \in \mathbb{M}} \mathcal{L}_m}{\mathcal{L}_{\hat{m}}} \rightarrow_p 1 \quad \text{as } n \rightarrow \infty, \quad (5)$$

where  $\hat{m}$  is the selected model,  $\mathcal{L}_m = E_*(s(\hat{p}_m, Z)) - E_*(s(p_*, Z))$  is the adjusted prediction loss, and  $\hat{p}_m$  denotes the estimated density function under model  $m$ .

The subtraction of  $E_*(s(p_*, Z))$  allows for better comparison of competing model selection methods. Another property often used to describe model selection is minimax-rate optimality, which will be elaborated on in the “Theoretical Properties of the Model Selection Criteria” section. A related but different school of thought is the structural risk minimization in the literature of statistical learning theory. In that context,

**In this article, we aim to provide an integrated understanding of the properties and practical performances of various approaches by reviewing their theoretical and practical advantages, disadvantages, and relations.**

a common practice is to bound the out-sample prediction loss using in-sample loss plus a positive term (e.g., a function of the Vapnik–Chervonenkis dimension [9] for a classification model). The major difference of the current setting compared with that in statistical learning is the (stronger) requirement that the selected model should exhibit prediction loss comparable to the best offered by the candidates. In other words, the positive term plus the in-sample loss should asymptotically approach the true out-sample loss (as sample size goes to infinity).

The goals of inference and prediction as assessed in terms of asymptotic efficiency of model selection can often be well aligned in a parametric framework, although there exists an unbridgeable conflict when a minimax view is taken to assess the prediction performance. We will elaborate on this and related issues in the “War and Peace—Conflicts Between AIC and BIC and Their Integration” section.

In light of all of the preceding discussions, we note that the task of model selection is primarily concerned with the selection of  $\mathcal{M}_m$  ( $m \in \mathbb{M}$ ), because once  $m$  is identified, the model fitting part is straightforward. Thus, the model selection procedure can also be regarded as a joint estimation of both the distribution family ( $\mathcal{M}_m$ ) and the parameters in each family ( $\theta_m \in \mathcal{H}_m$ ).

A model class  $\{\mathcal{M}_m\}_{m \in \mathbb{M}}$  is nested if smaller models are always special cases of larger models. For a nested model class, the model selection is sometimes referred to as the *order selection problem*. The task of model selection in its broad sense can also refer to *method* (or *modeling procedure*) *selection*, which we shall revisit in the “Modeling Procedure Selection” section.

### An illustration on fitting and the best model

We provide a synthetic experiment to illustrate the general rules that 1) better fitting does not imply better predictive per-

formance, and 2) the predictive performance is optimal at a candidate model that typically depends on both the sample size and the unknown data-generating process. As a result, an appropriate model selection technique is important to single out the best model for inference and prediction in a strong, practically parametric framework or to strike a good balance between the goodness of fit and model complexity on the observed data to facilitate optimal prediction in a practically nonparametric framework.

#### Example 1


Suppose that a set of time series data  $\{z_t : t = 1, \dots, n\}$  is observed, and we specify an AR model class with order at most  $d_n$ . Each model of dimension (or order)  $k$  ( $k = 1, \dots, d_n$ ) is in the form of

$$z_t = \sum_{i=1}^k \psi_{k,i} z_{t-i} + \varepsilon_t, \quad (6)$$

referred to as the AR( $k$ ), where  $\psi_{k,i} \in \mathbb{R}$  ( $i = 1, \dots, k$ ),  $\psi_{k,k} \neq 0$ , and  $\varepsilon_t$ s are independent Gaussian noises with zero mean and variance  $\sigma^2$ . Adopting quadratic loss, the parameters  $\psi_{k,1}, \dots, \psi_{k,k}$  can be estimated by the method of least squares. When the data-generating model is unknown, one critical problem is the identification of the (unknown) order of the AR model. We need to first estimate parameters with different orders  $1, \dots, d_n$  and then select one of them based on a certain principle.

#### Experiment

In this experiment, we first generate time series data using each of the following three true data-generating processes, with the sample sizes  $n = 100, 500, 2,000, 3,000$ . We then fit the data using the model class in Example 1, with maximal order  $d_n = 15$ .



**FIGURE 1.** The parametric framework: the best predictive performance is achieved at the true order three. (a) The in-sample loss for each sample size  $n$  monotonically decreases as the order (model complexity) increases. (b) The predictive performance is only optimal at the true order (circled). (c) The most efficient model (circled) is therefore the true model.

1) *Parametric framework*: The data are generated in the way described by (6) with true order  $k_0 = 3$  and parameters  $\psi_{3,\ell} = 0.7^\ell$  ( $\ell = 1, 2, 3$ ).

Suppose that we adopt the quadratic loss in Example 1. Then we obtain the average in-sample loss

$$\hat{e}_k = (n - k)^{-1} \sum_{t=k+1}^n (z_t - \sum_{i=1}^k \hat{\psi}_{k,i} z_{t-i})^2.$$

In Figure 1(a), we plot  $\hat{e}_k$  against  $k$  for  $k = 1, \dots, d_n$ , averaged over 50 independent replications. The curve for each sample size  $n$  is monotonically decreasing, because larger models fit the same data better. We compute and plot in Figure 1(b) the out-sample prediction loss in (4), which is equivalent to  $E_*(s(\hat{p}_k, Z_t)) = E_*(Z_t - \sum_{i=1}^k \hat{\psi}_{k,i} Z_{t-i})^2$  in this example. The above expectation is taken over the true stationary distribution of an independent process of  $Z_t$ . (An alternative definition is based on the same-realization expectation that calculates the loss of the future of an observed time series [10].) The curves in Figure 1(b) show that the predictive performance is only optimal at the true order.

Under the quadratic loss, we have  $E_*(s(p_*, Z_t)) = \sigma^2$ , and the asymptotic efficiency (Definition 2) requires that

$$\frac{\min_{k=1, \dots, d_n} E_*(Z_t - \sum_{i=1}^k \hat{\psi}_{k,i} Z_{t-i})^2 - \sigma^2}{E_*(Z_t - \sum_{j=1}^{\hat{k}} \hat{\psi}_{\hat{k},j} Z_{t-j})^2 - \sigma^2} \quad (7)$$

converges to one in probability. To describe how the predictive performance of each model deviates from the best possible, we define the efficiency of each model of order  $k'$  to be the quantity in (7) with  $\hat{k}$  being replaced with  $k'$  ( $k' = 1, \dots, d_n$ ). Note that the concepts of efficiency and asymptotic efficiency in model selection are reminiscent of their counterparts in

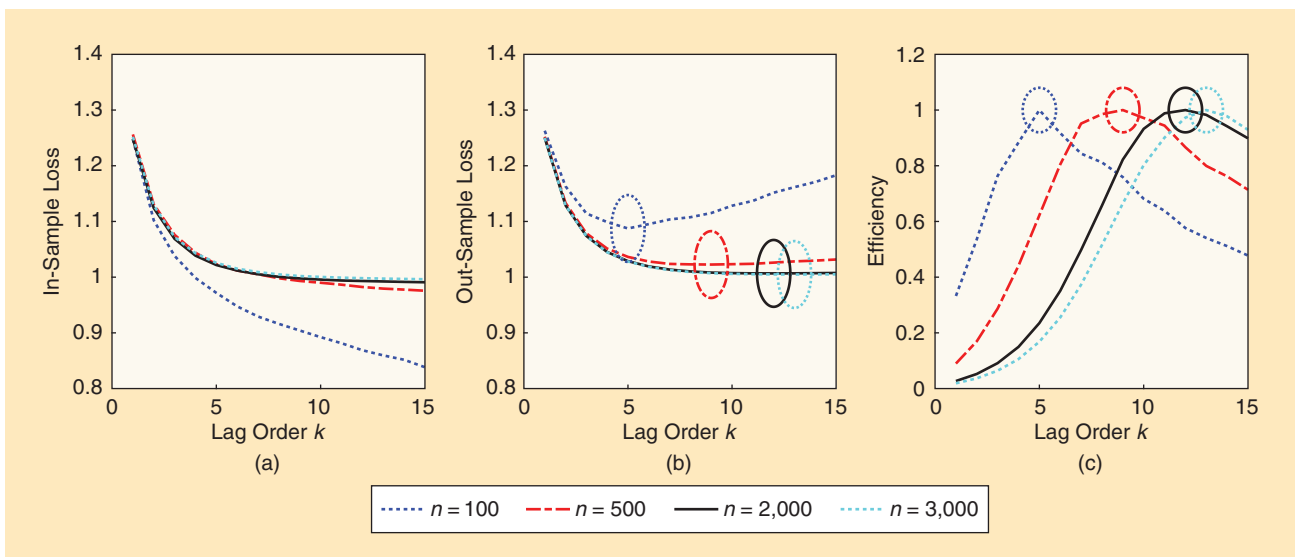
parameter estimation. We plot the efficiency of each candidate model in Figure 1(c). Similarly to Figure 1(b), the curves here show that the true model is the most efficient model. We note that the minus- $\sigma^2$  adjustment of out-sample prediction loss in the above definition makes the property highly nontrivial to achieve (see, e.g., [11]–[13]). Consider, e.g., the comparison between AR(2) and AR(3) models, with the AR(2) being the true data-generating model. It can be proved that without subtracting  $\sigma^2$ , the ratio (of the mean square prediction errors) for each of the two candidate models approaches one; by subtracting  $\sigma^2$ , the ratio for AR(2) still approaches one, whereas the ratio for AR(3) approaches 2/3.

2) *Nonparametric framework*: The data are generated by the moving average (MA) model  $z_t = \varepsilon_t - 0.8\varepsilon_{t-1}$ , with  $\varepsilon_t$  being independent standard Gaussian.

Similarly to case 1, we plot the results in Figure 2. Different from case 1, the predictive performance is optimal at increasing model dimensions as  $n$  increases. In such a nonparametric framework, the best model is sensitive to the sample size, so that pursuing an inference of a fixed good model becomes unrealistic. The model selection task aims to select a model that is asymptotically efficient [see Figure 2(c)]. Note that Figure 2(b) and (c) is drawn based on the information of the underlying true model, which is unavailable in practice; hence, we need a model selection method to achieve the asymptotic efficiency.

3) *Practically nonparametric framework*: The data are generated in the same way as in case 1, except that  $k_0 = 10$ .

We plot the results in Figure 3. For  $n = 2,000, 3,000$ , the sample sizes are large enough to support the evidence of a true model with a relatively small model dimension. Similarly to experiment 1, this is a parametric framework in which the optimal predictive performance is achieved at the true model. For  $n = 100, 500$ , where the sample sizes are not large enough



**FIGURE 2.** The nonparametric framework: the best predictive performance is achieved at an order that depends on the sample size. (a) The in-sample loss for each sample size  $n$  monotonically decreases as the order (model complexity) increases. (b) The predictive performance is optimal at increasing orders (circled) as  $n$  increases. (c) The order of the most efficient model (circled) therefore increases as  $n$  increases.

compared to the true model dimension, however, fitting too many parameters actually causes an increased variance that diminishes the predictive power. In such a scenario, even though the true model is included as a candidate, the best model is not the true model, and it is unstable for small or moderate sample sizes as if in a nonparametric setting. In other words, the parametric framework can turn into a practically nonparametric framework in the small data regime. It can also work the other way around, i.e., for a true nonparametric framework, for a large range of sample sizes (e.g., 100–2,000), a relatively small parametric model among the candidates continues to be the best model [14].

### Principles and approaches from various philosophies or motivations

A wide variety of model selection methods have been proposed in the past few decades, motivated by different viewpoints and justified under various circumstances. Many of them originally aimed to select either the order in an AR model or a subset of variables in a regression model. We review some of the representative approaches in these contexts in this section.

#### Information criteria based on likelihood functions

Information criteria generally refer to model selection methods that are based on likelihood functions and applicable to parametric model-based problems. Here we introduce some information criteria whose asymptotic performances are well understood.

Akaike information criterion (AIC) is a model selection principle proposed by Akaike [15]. A detailed derivation of it from an information theoretic perspective can be found in [1]. Briefly speaking, the idea is to approximate the out-sample prediction loss by the sum of the in-sample loss and a correc-

tion term. We refer to [1] for a detailed derivation of this correction term. In the typical setting where the loss is logarithmic, the AIC procedure is to select the model  $\mathcal{M}_m$  that minimizes

$$\text{AIC}_m = -2\hat{\ell}_{n,m} + 2d_m, \quad (8)$$

where  $\hat{\ell}_{n,m}$  is the maximized log likelihood of model  $\mathcal{M}_m$  given  $n$  observations as defined in (1), and  $d_m$  is the dimension of model  $\mathcal{M}_m$ . It is clear that more complex models (with larger  $d_m$ ) will suffer from larger penalties.

In the task of AR order selection, it is also common to use

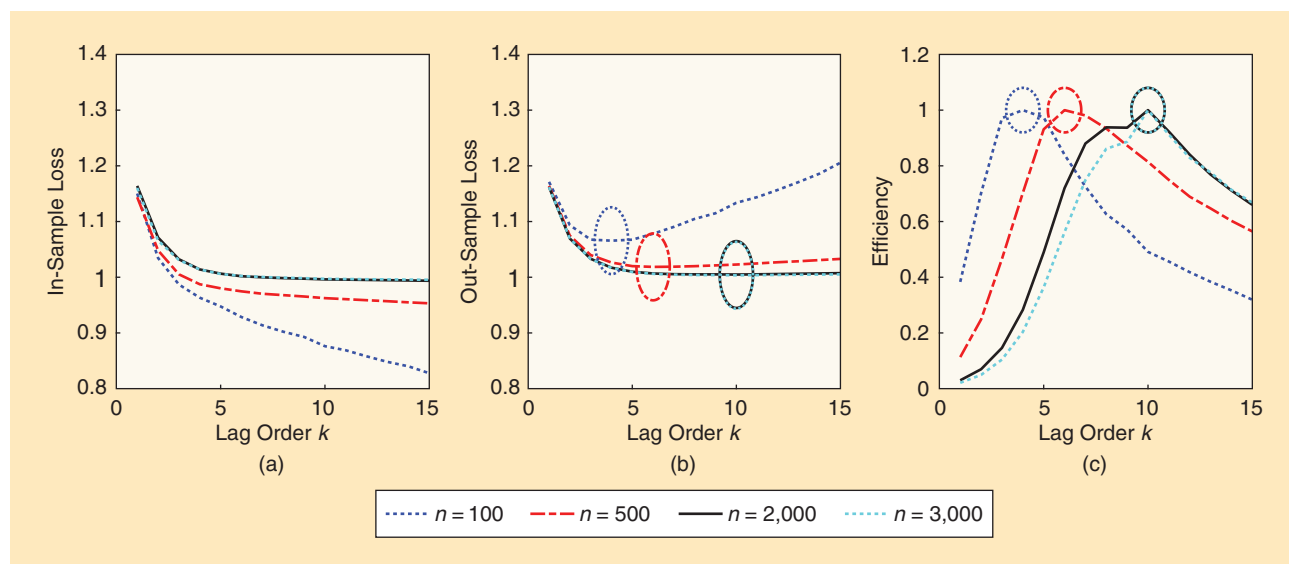
$$\text{AIC}_k = n \log \hat{e}_k + 2k \quad (9)$$

for the model of order  $k$ , where  $\hat{e}_k$  is the average in-sample loss based on the quadratic loss. In fact, (9) can be derived from (8) by assuming that AR noises are Gaussian and by regarding ARs of different orders as  $\{\mathcal{M}_m\}_{m \in \mathbb{M}}$ . A predecessor of AIC is the final prediction error criterion (FPE) [16] (also by Akaike). An extension of AIC is the Takeuchi's information criterion [17], derived in a way that allows model misspecification, but it is rarely used in practice due to its computational complexity. In the context of generalized estimating equations for correlated response data, a variant of AIC based on quasi-likelihood is derived in [18].

Finite-sample corrected AIC (AICc) [19] was proposed as a corrected version of the AIC for small-sample study. It selects the model that minimizes

$$\text{AICc}_m = \text{AIC}_m + \frac{2\{d_m + 1\}\{d_m + 2\}}{n - d_m - 2}.$$

Unless the sample size  $n$  is small compared with model dimension  $d_m$ , there is little difference between AICc and



**FIGURE 3.** The practically nonparametric framework: the best predictive performance is achieved at an order that depends on the sample size in the small data regime. (a) The in-sample loss for each sample size  $n$  monotonically decreases as the order (model complexity) increases. (b) The predictive performance is optimal at increasing orders (circled) as  $n$  increases in a certain range. (c) The order of the most efficient model (circled) therefore depends on  $n$  in a certain range.

AIC. Another modified AIC that replaces the constant two with a different positive number has also been studied in [20].

Bayesian information criterion (BIC) [21] is another popular model selection principle. It selects the model  $m$  that minimizes

$$\text{BIC}_m = -2\hat{\ell}_{n,m} + d_m \log n. \quad (10)$$

The only difference with AIC is that the constant two in the penalty term is replaced with the logarithm of the sample size. The original derivation of BIC by Schwarz turned out to have a nice Bayesian interpretation, as its current name suggests.

To see the interpretation, we assume that  $z_1, \dots, z_n$  are the realizations of independent, identically distributed random variables, and  $\pi(\cdot)$  is any prior distribution on  $\theta$  that has dimension  $d$ . We let  $\ell_n(\theta) = \sum_{i=1}^n \log p_\theta(z_i)$  be the log-likelihood function and  $\hat{\theta}_n$  the MLE of  $\theta$ . Based on classical Bayesian asymptotics we have (see, e.g., [22, eq. (1.5)]) under regularity conditions

$$\begin{aligned} \int_{\mathbb{R}^d} \pi(\hat{\theta}_n + n^{-\frac{1}{2}}r) \exp(\ell_n(\hat{\theta}_n + n^{-\frac{1}{2}}r) - \ell_n(\hat{\theta}_n)) dr \\ \rightarrow p_\pi(\theta_*) (2\pi)^{d/2} \{\det E_*(-\nabla_\theta^2 \log p_{\theta_*}(z))\}^{-1/2} \end{aligned} \quad (11)$$

as  $n \rightarrow \infty$ , for some constant  $\theta_*$ . Note that the right-hand side of (11) is a constant that does not depend on  $n$  and the left-hand side of (11) equals

$$p(z_1, \dots, z_n) \exp(-\ell_n(\hat{\theta}_n) + \frac{d}{2} \log n). \quad (12)$$

Therefore, selecting a model with the largest marginal likelihood  $p(z_1, \dots, z_n)$  (as advocated by Bayesian model comparison) is asymptotically equivalent to selecting a model with the smallest BIC in (10). It is interesting to see that the marginal likelihood of a model does not depend on the imposed prior at all in the large sample limit. Intuitively speaking, this is because, in the integration of  $p(z_1, \dots, z_n) = \int_{\mathcal{H}_m} \pi(\theta) p_\theta(z_1, \dots, z_n) d\theta$ , the mass is concentrated around  $\hat{\theta}_n$  with radius  $O(n^{-1/2})$  and dimension  $d$ , so its value is proportional to the maximized likelihood value multiplied by the volume approximately at the order of  $n^{-d/2}$ , which is in line with (12).

Hannan and Quinn (HQ) criterion [23] was proposed as an information criterion that achieves strong consistency in AR order selection. In other words, if the data are truly generated by an AR model of fixed order  $k_0$ , then the selected order  $k$  converges almost surely to  $k_0$  as the sample size goes to infinity. We note that strong consistency implies (the usual) consistency. In general, this method selects a model by minimizing  $\text{HQ}_m = -2\hat{\ell}_{n,m} + 2c d_m \log \log n$  (for any constant  $c > 1$ ). It can be proved under some conditions that any penalty no larger than  $2d_m \log \log n$  is not strongly consistent [23]; therefore, HQ employs the smallest possible penalty to guarantee strong consistency.

Bridge criterion (BC) [24], [25] is a recently proposed information criterion that aims to bridge the advantages of both AIC and BIC in the asymptotic regime. It selects the model

$\mathcal{M}_m$  that minimizes  $\text{BC}_m = -2\hat{\ell}_{n,m} + c_n(1 + 2^{-1} + \dots + d_m^{-1})$  (with the suggested  $c_n = n^{2/3}$ ) over all of the candidate models whose dimensions are no larger than  $d_{m_{\text{AIC}}}$ , the dimension of the model selected by AIC. Note that the penalty is approximately  $c_n \log d_m$ , but it is written as a harmonic number to highlight some of its nice interpretations. Its original derivation was motivated by a recent finding that the information loss of underfitting a model of dimension  $d$  using dimension  $d-1$  is asymptotically  $\chi_1^2/d$  for large  $d$ , assuming that nature generates the model from a noninformative uniform distribution over its model space (in particular the coefficient space of all stationary autoregressions) [24, Appendix A]. BC was proved to perform similarly to AIC in a nonparametric framework and similarly to BIC in a parametric framework. We further discuss BC in the “War and Peace—Conflicts Between AIC and BIC and Their Integration” section.

### Methods from other perspectives

In addition to information criteria, some other model selection approaches have been motivated from either Bayesian, information-theoretic, or decision-theoretic perspectives.

Bayesian posterior probability is commonly used in Bayesian data analysis. Suppose that each model  $m \in \mathbb{M}$  is assigned a prior probability  $p(\mathcal{M}_m) > 0$  (such that  $\sum_{m \in \mathbb{M}} p(\mathcal{M}_m) = 1$ ), interpreted as the probability that model  $\mathcal{M}_m$  contains the true data-generating distribution  $p_*$ . Such a prior may be obtained from scientific reasoning or knowledge from historical data. For each  $m \in \mathbb{M}$ , we also introduce a prior, with density  $\theta_m \mapsto p_m(\theta_m)$  ( $\theta_m \in \mathcal{H}_m$ ), and a likelihood of data  $p_m(z | \theta_m)$ , where  $z = [z_1, \dots, z_n]$ . A joint distribution on  $(z, \theta_m, \mathcal{M}_m)$  is therefore well defined, based on which various quantities of interest can be calculated. We first define the marginal likelihood of model  $\mathcal{M}_m$  by

$$p(z | \mathcal{M}_m) = \int_{\mathcal{H}_m} p_m(z | \theta_m) p_m(\theta_m) d\theta_m. \quad (13)$$

Based on (13), we obtain the following posterior probabilities on models by Bayes formula:

$$p(\mathcal{M}_m | z) = \frac{p(z | \mathcal{M}_m) p(\mathcal{M}_m)}{\sum_{m' \in \mathbb{M}} p(z | \mathcal{M}_{m'}) p(\mathcal{M}_{m'})}. \quad (14)$$

The maximum a posteriori approach [26] would select the model with the largest posterior probability.

Bayes factors are also popularly adopted for Bayesian model comparison, defined for a pair of models  $(\mathcal{M}_m, \mathcal{M}_{m'})$  by

$$B_{m,m'} = \frac{p(\mathcal{M}_m | z)}{p(\mathcal{M}_{m'} | z)} = \frac{p(\mathcal{M}_m)}{p(\mathcal{M}_{m'})} \frac{p(z | \mathcal{M}_m)}{p(z | \mathcal{M}_{m'})}.$$

The model  $\mathcal{M}_m$  is favored over  $\mathcal{M}_{m'}$  if  $B_{m,m'} > 1$ . Bayes factors remove the impact of prior probabilities on the models from the selection process to focus on the ratio of marginal likelihoods. Compared with the Bayesian posterior probability, Bayes factors are appealing when it is difficult to formulate prior probabilities on models.

Bayesian marginal likelihood, defined in (13), also referred to as the *evidence* or *model evidence*, is a quantity naturally motivated by Bayes factors. In the presence of multiple models, the one with the largest Bayesian marginal likelihood is favored over all other models in terms of the Bayes factor. Moreover, it can be seen that the model with the highest marginal likelihood is the model with the highest posterior probability given that the Bayesian prior probabilities on models are all equal. Interestingly, this Bayesian principle using marginal likelihood is asymptotically equivalent to the BIC (as we have introduced in the “Information Criteria Based on Likelihood Function” section). In practice, the preceding Bayesian model selection methods can be computationally challenging. Calculation of the quantities in (13) and (14) are usually implemented using Monte Carlo methods, especially sequential Monte Carlo (for online data) and Markov chain Monte Carlo (for batch data) (see, e.g., [27]). It is worth noting that improper or vague priors on the parameters of any candidate model can have a nonnegligible impact on the interpretability of marginal likelihood and Bayes factors in the nonasymptotic regime, and that has motivated some recent research on Bayesian model selection (see, e.g., [28] and the references therein).

The minimum message length (MML) principle [29] was proposed from an information-theoretic perspective. It favors the model that generates the shortest overall message, which consists of a statement of the model and a statement of the data concisely encoded with that model. Specifically, this criterion aims to select the model that minimizes

$$-\log p(\theta) - \log p(x|\theta) + \frac{1}{2} \log |I(\theta)| + \frac{d}{2}(1 + \log \kappa_d),$$

where  $p(\theta)$  is a prior,  $p(x|\theta)$  is the likelihood function,  $I(\theta) = \int \{\partial \log p(x|\theta) / \partial \theta\}^2 p(x|\theta) dx$  is the Fisher information,  $d$  is the dimension of  $\theta$ , and  $\kappa_d$  is the so-called optimal quantizing lattice constant that is usually approximated by  $\kappa_1 = 1/12$ . A detailed derivation and application of MML can be found in [30].

The minimum description length (MDL) principle [31]–[34] describes the best model as the one that leads to the best compression of a given set of data. It was also motivated by an information-theoretic perspective (which is similar to MML). Different from MML, which is in a fully Bayesian setting, MDL avoids assumptions on prior distribution. Its predictive extension, referred to as the *predictive minimum description length criterion (PMDL)*, is proposed in [35]. One formulation of the principle is to select the model by minimizing the stochastic complexity  $-\log p_{\theta_t}(z_1) - \sum_{t=2}^n \log p_{\theta_t}(z_t | z_1, \dots, z_{t-1})$ , in which  $\theta_t$ 's are restricted to the same parameter space (with the same dimension). Here, each  $\theta_t$  ( $t > 1$ ) is the MLE calculated using  $z_1, \dots, z_{t-1}$ , and  $p_{\theta_t}(\cdot)$  can be an arbitrarily chosen prior distribution. The above PMDL criterion is also closely related to the prequential (or predictive sequential) rule [36] from a decision-theoretic perspective.

**Interestingly, the LOO was shown to be asymptotically equivalent to either AIC/Takeuchi's information criterion under some regularity conditions.**

Deviance information criterion (DIC) [37] was derived as a measure of Bayesian model complexity. Instead of being derived from a frequentist perspective, DIC can be thought of as a Bayesian counterpart of AIC. To define DIC, a relevant concept is the deviance under model  $m$ :  $D_m(\theta) = -2 \log p_m(y|\theta) + C$ , where  $C$  does not depend on the model being compared. Also, we define the effective number of parameters of the model to be  $p_D = E_{\theta|z} D_m(\theta) - D_m(E_{\theta|z}(\theta))$ , where  $E_{\theta|z}(\cdot)$  is the expectation taken over  $\theta$  conditional on all of the observed data  $z$  under model  $\mathcal{M}_m$ . Then the DIC selects the model  $\mathcal{M}_m$  that minimizes

$$\text{DIC}_m = D_m(E_{\theta|z}(\theta)) + 2p_D. \quad (15)$$

The conceptual connection between DIC and AIC can be readily observed from (15). The MLE and model dimension in AIC are replaced with the posterior mean and effective number of parameters, respectively, in DIC. Compared with AIC, DIC enjoys some computational advantage for comparing complex models whose likelihood functions may not even be in analytic forms. In Bayesian settings, Markov chain Monte Carlo tools can be utilized to simulate posterior distributions of each candidate model, which can be further used to efficiently compute DIC in (15).

### Methods that do not require parametric assumptions

Cross validation (CV) [38], [39] is a class of model selection methods widely used in machine-learning practice. CV does not require the candidate models to be parametric, and it works as long as the data are permutable and one can assess the predictive performance based on some measure. A specific type of CV is the delete-1 CV method [40] [or leave-one-out (LOO)]. The idea is as follows. For brevity, let us consider a parametric model class as before. Recall that we wish to select a model  $\mathcal{M}_m$  with as small out-sample loss  $E_*(s(p_{\hat{\theta}_m}, Z))$  as possible. Its computation involves an unknown true data-generating process, but we may approximate it by  $n^{-1} \sum_{i=1}^n s(p_{\hat{\theta}_{m,-i}}, z_i)$ , where  $\hat{\theta}_{m,-i}$  is the MLE under model  $\mathcal{M}_m$  using all of the observations except  $z_i$ . In other words, given  $n$  observations, we leave each one observation out in turn and attempt to predict that data point by using the  $n-1$  remaining observations, and we record the average prediction loss over  $n$  rounds. Interestingly, the LOO was shown to be asymptotically equivalent to either AIC/Takeuchi's information criterion under some regularity conditions [40].

In general, CV works in the following way. It first randomly splits the original data into a training set of  $n_t$  data  $1 \leq n_t \leq n-1$  and a validation set of  $n_v = n - n_t$  data; each candidate model is then trained from the  $n_t$  data and validated on the remaining data (i.e., to record the average validation loss). This procedure is independently replicated a few times (each with a different validation set) to reduce the variability caused by splitting. Finally, the model with the



smallest average validation loss is selected, and it is retrained using the complete data for future prediction.

A special type of CV is the so-called  $k$ -fold CV (with  $k$  being a positive integer). It randomly partitions data into  $k$  subsets of (approximately) equal size; each model is trained on  $k-1$  folds and validated on the remaining one fold. The procedure is repeated  $k$  times, and the model with the smallest average validation loss is selected. The  $k$ -fold CV is perhaps more commonly used than LOO, partly due to the large computational complexity involved in LOO. The holdout method, as often used in data competitions (e.g., Kaggle competition), is also a special case of CV. It does data splitting only once, one part as the training set and the remaining part as the validation set. We note that there exist fast methods, such as generalized cross validation (GCV) [90], as surrogates to LOO to reduce the computational cost. Some additional discussion on CV will be provided in the ‘‘Clarification of Some Misconceptions’’ section.

### Methods proposed for specific types of applications

There have been some other criteria proposed for specific types of applications, mostly for time series or linear regression models.

The predictive least-squares (PLS) principle proposed by Rissanen [41] is a model selection criterion based on his PMDL principle. PLS aims to select the stochastic regression model by minimizing the accumulated squares of prediction errors (in a time-series setting), defined as

$$\text{PLS}_m = \sum_{t=t_0+1}^n (y_t - x_{m,t}^T \beta_{m,t-1})^2,$$

where  $y_t$  is each response variable,  $x_{m,t}$  is the vector of covariates corresponding to model  $m$ , and  $\beta_{m,t-1}$  is the least-squares estimate of model  $\mathcal{M}_m$  based on data before time  $t$ . The time index  $t_0$  is the first index such that  $\beta_t$  is uniquely defined. Conceptually, PLS is not like AIC and BIC, which select the model that minimizes a loss plus a penalty. It seems more like the counterpart of LOO in sequential contexts. Interestingly, it has been proved that PLS and BIC are asymptotically close, both strongly consistent in selecting the data-generating model (in a parametric framework) [42]. Extensions of PLS where the first index  $t_0$  is a chosen sequence indexed by  $n$  have also been studied. It has been shown, e.g., that PLS with  $t_0/n \rightarrow 1$  shares the same asymptotic property of AIC under some conditions (see, e.g., [43, Example 8]).

Generalized information criterion ( $\text{GIC}_{\lambda_n}$ ) [12], [44] represents a wide class of criteria whose penalties are linear in model dimension. It aims to select the regression model  $\mathcal{M}_m$  that minimizes

$$\text{GIC}_{\lambda_n, m} = \hat{e}_m + \frac{\lambda_n \hat{\sigma}_n^2 d_m}{n}.$$

Here,  $\hat{\sigma}_n^2$  is an estimator of  $\sigma^2$ , the variance of the noise, and  $\hat{e}_m = n^{-1} \|y - \hat{y}_m\|_2^2$  is the mean square error between the observations and least-squares estimates under regression model  $\mathcal{M}_m$ .  $\lambda_n$  is a deterministic sequence of  $n$  that controls the tradeoff between the model fitting and model complexity.

If we replace  $\hat{\sigma}_n^2$  with  $(n - d_m)^{-1} n \hat{e}_m$ , it can be shown under mild conditions that minimizing  $\text{GIC}_{\lambda_n}$  is equivalent to minimizing [12, p. 232]

$$\log \hat{e}_m + \frac{\lambda_n d_m}{n}. \quad (16)$$

In this case,  $\lambda_n = 2$  corresponds to AIC, and  $\lambda_n = \log n$  corresponds to BIC. Mallows’s  $C_p$  method [45] is a special case of GIC with  $\hat{\sigma}_n^2 \triangleq (n - d_m)^{-1} n \hat{e}_m$  and  $\lambda_n = 2$ , where  $\hat{m}$  indexes the largest model that includes all of the covariates.

### Theoretical properties of the model selection criteria

Theoretical examinations of model selection criteria have centered on several properties: consistency in selection, asymptotic efficiency, and minimax-rate optimality. Selection consistency targets the goal of identifying the best model or method on its own for scientific understanding, statistical inference, insight, or interpretation. Asymptotic efficiency and minimax-rate optimality (defined in Definition 3, which follows) are in tune with the goal of prediction. Before we introduce the theoretical properties, it is worth mentioning that many model selection methods can also be categorized into two classes according to their large-sample performances, represented by AIC and BIC. In fact, it has been known that AICc, FPE, and GCV are asymptotically close to AIC, whereas Bayes factors, HQ, and the original PLS are asymptotically close to BIC. For some other methods, such as CV and GIC, their asymptotic behavior usually depends on the tuning parameters.  $\text{GIC}_{\lambda_n}$  is asymptotically equivalent to AIC when  $\lambda_n = 2$  and to BIC when  $\lambda_n = \log n$ . In general, any sequence of  $\lambda_n$  satisfying  $\lambda_n \rightarrow \infty$  would exhibit the consistency property shared by BIC. As a corollary, the  $C_p$  method (as a special case of  $\text{GIC}_2$ ) is asymptotically equivalent to AIC. For CV with  $n_t$  training data and  $n_v$  validation data, it is asymptotically similar to AIC when  $n_v/n_t \rightarrow 0$  (including the LOO as a special case) and to BIC when  $n_v/n_t \rightarrow \infty$  [12, eq. (4.5)].

In general, AIC and BIC have served as the golden rules for model selection in statistical theory during their existence. Though cross validations or Bayesian procedures have also been widely used, their asymptotic justifications are still rooted in frequentist approaches in the form of AIC, BIC, and so forth. Therefore, understanding the asymptotic behavior of AIC and BIC is crucial in both theory and practice. We thus focus on the properties of AIC and BIC in the rest of this section and the ‘‘War and Peace—Conflicts Between AIC and BIC and Their Integration’’ section. It is remarkable that the asymptotic watershed of AIC and BIC (and their closely related methods) simply lies in whether the penalty is a fixed, well-chosen constant or goes to infinity as a function of  $n$ .

First of all, AIC is proved to be minimax-rate optimal for a range of variable selection tasks, including the usual subset selection and order selection problems in linear regression and nonparametric regression based on series expansion with such bases as polynomials, splines, or wavelets (see, e.g., [46] and the references therein). Consider, e.g., the minimax risk of estimating the regression function  $f \in \mathcal{F}$  under the squared error

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}} n^{-1} \sum_{i=1}^n E_* (\hat{f}(x_i) - f(x_i))^2, \quad (17)$$

where  $\hat{f}$  is over all estimators based on the observations and  $f(x_i)$  is the expectation of the  $i$ th response variable (or the  $i$ th value of the regression function) conditional on the  $i$ th vector of variables  $x_i$ . Each  $x_i$  can refer to a vector of explanatory variables, or polynomial basis terms, and so on. For a model selection method  $\nu$ , its worst-case risk is  $\sup_{f \in \mathcal{F}} R(f, \nu, n) = n^{-1} \sum_{i=1}^n E_* \{\hat{f}_\nu(x_i) - f(x_i)\}^2$ , with  $\hat{f}_\nu$  being the least-squares estimate of  $f$  under the variables selected by  $\nu$ .

### Definition 3

A method  $\nu$  is said to be minimax-rate optimal over  $\mathcal{F}$  if  $\sup_{f \in \mathcal{F}} R(f, \nu, n)$  converges at the same rate as the minimax risk in (17).

Another good property of AIC is that it is asymptotically efficient [as defined in (5)] in a nonparametric framework (see, e.g., [11] and [47]). In other words, the predictive performance of its selected model is asymptotically equivalent to the best offered by the candidate models (even though it is sensitive to the sample size).

BIC, on the other hand, is known to be consistent in selecting the smallest true data-generating model in a parametric framework (see, e.g., [12] and [23]). Suppose, e.g., that the data are truly generated by an AR(2) and the candidate models are AR(2), AR(3), and an MA model that is essentially AR( $\infty$ ). Then AR(2) is selected with probability going to one as the sample size tends to infinity. MA(1) is not selected because it is a wrong model, and AR(3) is not selected because it overfits [even though it nests AR(2) as its special case]. Moreover, it can be proved that the consistency of BIC also implies that it is asymptotically efficient in a parametric framework [12], [24]. We will elaborate more on the theoretical properties of AIC and BIC in the next section.

## War and peace—Conflicts between AIC and BIC and their integration

In this section, we review some research advances in the understanding of AIC, BIC, and related criteria. The choice of AIC and BIC to focus on here because they represent two cornerstones of model selection principles and theories. We are only concerned with the settings where the sample size is larger than the model dimension. Details of the following discussions can be found in such original papers as [11], [12], [24], [47]–[49], and the references therein.

Recall that AIC is asymptotically efficient for the nonparametric framework and is also minimax optimal [46]. In contrast, BIC is consistent and asymptotically efficient for the parametric framework. Despite the good properties of AIC and BIC, they have their own drawbacks. AIC is known to be inconsistent in a parametric framework where there are at least two correct candidate models. As a result, AIC is not asymptotically efficient in such a framework.

If data are truly generated by an AR(2), e.g., and the candidate models are AR(2), AR(3), and so forth, then AR(2) cannot be selected with probability going to one by AIC as the sample size increases. The asymptotic probability of it being selected can actually be analytically computed [48]. BIC, on the other hand, does not enjoy the properties of minimax-rate optimality and asymptotic efficiency in a nonparametric framework [12], [50].

Why do AIC and BIC work in those ways? Theoretical arguments in those aspects are highly nontrivial and have motivated a vast literature since the formulations of AIC and BIC. Here we provide some heuristic explanations. For AIC, its formulation in (8) was originally motivated by searching the candidate model  $p$  that is the closest in Kullback–Leibler (KL) divergence (denoted by  $D_{KL}$ ) from  $p$  to the data-generating model  $p_*$ . Because  $\min_p D_{KL}(p_*, p)$  is equivalent to  $\min_p E_*(-\log p)$  for a fixed  $p_*$ , AIC is expected to perform well in minimizing the prediction loss. But AIC is not consistent for a model class containing a true model and at least one oversized model, because fitting the oversized model would only reduce the first term  $-2\hat{\ell}_{n,m}$  in (8) by a random quantity that is approximately chi-square distributed (by, e.g., Wilks’s theorem [51]), whereas the increased penalty on the second item  $2d_m$  is at a constant level, which is not large enough to suppress the overfitting gain in fitness with high probability. Selection consistency of BIC in a parametric framework is not surprising due to its nice Bayesian interpretation (see the “Principles and Approaches from Various Philosophies or Motivations” section). However, its penalty  $d_m \log n$  in (10) is much larger than the  $2d_m$  in AIC, so it cannot enjoy the predictive optimality in a typical nonparametric framework (if AIC already does so).

To briefly summarize, for asymptotic efficiency, AIC (respectively, BIC) is only suitable in nonparametric (respectively, parametric) settings. Figure 4 illustrates the two situations. There has been a debate between AIC and BIC in model selection practice, centering on whether the data-generating process is in a parametric framework or not. The same debate was sometimes raised under other terminology. In a parametric (respectively, nonparametric) framework, the true data-generating model is often said to be well specified (respectively, misspecified) or finite (respectively, infinite) dimensional. (To see a reason for such terminology, consider, e.g., the regression analysis using polynomial basis function as covariates. If the true regression function is indeed a polynomial, then it can be parameterized with a finite number of parameters; if it is an exponential function, then it cannot be parameterized with any finite dimensional parameter.) Without prior knowledge on how the observations were generated, determining which method to use becomes very challenging. It naturally motivates the following fundamental question: Is it possible to have a method that combines the strengths of AIC and BIC?

**There has been a debate between AIC and BIC in model selection practice, centering on whether the data-generating process is in a parametric framework or not.**

The combining of strengths can be defined in two ways. First, can the properties of minimax-rate optimality and consistency be shared? Unfortunately, it has been theoretically shown under rather general settings that there exists no model selection method that achieves both optimality and consistency simultaneously [49]. For any model selection procedure to be consistent, i.e., it must behave suboptimally in terms of minimax rate of convergence in the prediction loss. Second, can the properties of asymptotic efficiency and consistency be shared? In contrast to minimax-rate optimality, which allows the true data-generating model to vary, asymptotic efficiency is in a pointwise sense, meaning that the data are already generated by some fixed (unknown) data-generating model. Therefore, the asymptotic efficiency is a requirement from a more optimistic view and thus weaker in some sense than the minimaxity. Recall that consistency in a parametric framework is typically equivalent to asymptotic efficiency [12], [24]. Clearly, if an ideal method can combine asymptotic efficiency and consistency, it achieves asymptotic efficiency in both parametric and nonparametric frameworks. That motivated an active line of recent advances in reconciling the two classes of model selection methods [24], [43], [52].

In particular, the new model selection method BC was recently proposed (see the “Principles and Approaches from Various Philosophies or Motivations” section) to simultaneously achieve consistency in a parametric framework and asymptotic efficiency in both (parametric and nonparametric) frameworks. The key idea of BC is to impose a BIC-like heavy penalty for a range of small models but to alleviate the penalty for larger models if more evidence is supporting an infinite dimensional true model. In that way, the selection procedure is automatically adaptive to the appropriate setting (either parametric or nonparametric). A detailed statistical interpretation of how BC works in both theory and practice and how it relates to AIC and BIC is elaborated in [24].

Moreover, in many applications, data analysts would like to quantify to what extent the framework under consideration can be practically treated as parametric, or, in other words, how likely the postulated model class is well specified. This motivated the concept of the parametricness index (PI) [14], [24], which assigns a confidence score to model selection. One definition of PI, which we shall use in the following experiment, is this quantity on  $[0, 1]$ :

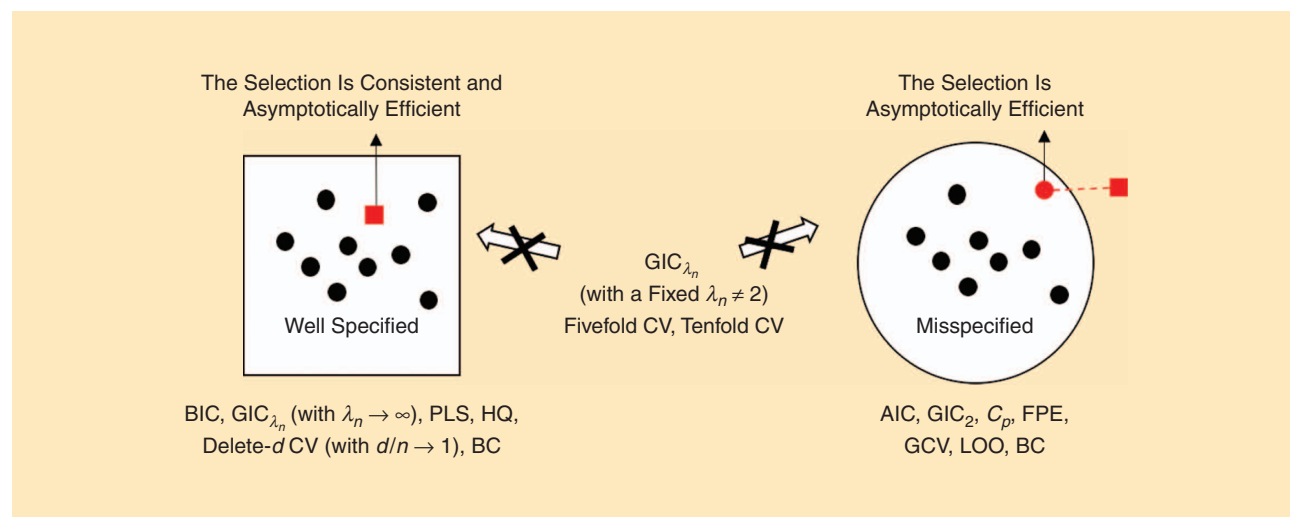
$$PI_n = \frac{|d_{m_{BC}} - d_{m_{AIC}}|}{(|d_{m_{BC}} - d_{m_{AIC}}| + |d_{m_{BC}} - d_{m_{BIC}}|)}$$

if the denominator is not zero, and  $PI_n = 1$  otherwise. Here,  $d_{m_\nu}$  is the dimension of the model selected by the method  $\nu$ . Under some conditions, it can be proved that  $PI_n \rightarrow_p 1$  in a parametric framework and  $PI_n \rightarrow_p 0$  otherwise.

### Experiments

We now revisit Example 1 in the “An Illustration on Fitting and the Best Model” section and numerically demonstrate the performances of different methods based on 100 replications and  $n = 500$ . For each of the three cases, we compute the means and standard errors of the efficiency [defined in (7)], dimension of the selected model, and PI and summarize them in Table 1. In case 1, BIC and BC perform much better than AIC in terms of efficiency, and PI is close to 1. This is expected from theory, as we are in a parametric setting. In cases 2 and 3, which are (practically) nonparametric, BC performs similarly to AIC, much better than BIC, and PI is closer to zero.

In practice, AIC seems more widely used compared with BIC, perhaps mainly due to the thinking that all models are wrong and minimax-rate optimality of AIC offers more robustness in adversarial settings than BIC. Nevertheless, the parametric setting is still of vital importance. First of all, being consistent in selecting the true model if it is really among the candidates is certainly mathematically appealing, and a nonparametric



**FIGURE 4.** A graph illustrating a parametric setting where the model class (by large square) includes the true data-generating model (by small red square) and a nonparametric setting where the model class (by large circle) excludes the true data-generating model, along with an asymptotically efficient model (by red circle) in the second case. It also lists some popular methods suitable for either situation and a class of GIC and CV that are asymptotically suboptimal for regression models.

framework can be a practically parametric framework. More importantly, when decisions need to be made on the use of certain variables, the concept of consistency that avoids over-selection of variables is practically very important. If medical researchers need to decide if certain genes should be further studied in costly experiments, e.g., the protection of overfitting of BIC avoids recommending variables that are hard to be justified statistically in a follow-up study, whereas AIC may recommend quite a few variables that may have some limited predictive values but their effects are too small to be certain with the limited information in the data for decision making and inference purposes.

The war between AIC and BIC originates from two fundamentally different goals: one to minimize certain loss for prediction purpose and the other to select the best model for inference purpose. A unified view on reconciling two such different goals wherever possible is a fundamental issue in model selection, and it remains an active line of research. We have witnessed some recent advances in that direction, and we expect more discoveries to flourish in the future.

### High-dimensional variable selection

The methods introduced in the “Principles and Approaches from Various Philosophies or Motivations” section were designed for small models, where the dimension  $d_n$  is often required to be  $o(\sqrt{n})$  in technical proofs. In this section, we elaborate on high-dimensional regression variable selection, an important type of model selection problems in which  $d_n$  can be comparable with or even much larger than  $n$ . To alleviate the difficulties, the data-generating model is often assumed to be a well-specified linear model, i.e., one of the following candidate models.

Each model  $\mathcal{M}$  assumes that  $y = \sum_{i \in \mathcal{M}} \beta_i x_i + \varepsilon$ , with  $\varepsilon$  being random noises. Here, with a slight abuse of notation, we have also used  $\mathcal{M}$  to denote a subset of  $\{1, \dots, d_n\}$ , and each data point is written as  $z = [y, x_1, \dots, x_{d_n}]$ , with  $y$  being the observed response and  $x_i$  being the (either fixed or random) covariates. Here,  $d_n$  instead of  $d$  is used to highlight that the number of candidate variables may depend on the sample size  $n$ .

The variable selection problem is also known as *support recovery* or *feature selection* in different literature. The mainstream idea to select the subset of variables is to either solve a penalized regression problem or iteratively pick up significant variables. The proposed methods differ from each other in terms of how they incorporate unique domain knowledge (e.g., sparsity, multicollinearity, group behavior) or what desired properties (e.g., consistency in coefficient estimation, consistency in variable selection) to achieve. The list of methods we will introduce is far from complete. Wavelet shrinkage, iterative thresholding, Dantzig selector,  $\ell_q$ -regularization with  $q \in (0, 1)$  (see, e.g., [53]–[57]), e.g., will not be covered.

### Penalized regression for variable selection

In a classical setting, a model class is first prescribed to data analysts (either from scientific reasoning or from exhaustive

search over  $d_n$  candidate variables), and then a criterion is used to select the final model (by applying any properly chosen method explained in the “Principles and Approaches from Various Philosophies or Motivations” section). When there is no ordering of variables known in advance and the number of variables  $d_n$  is small, one may simply search over  $2^{d_n}$  possible subsets and perform model selection. But it is usually computationally prohibitive to enumerate all possible subsets for large  $d_n$ , especially when  $d_n$  is comparable with or even larger than the sample size  $n$ . Note also that the problem of obtaining a sparse representation of signal  $y$  through some chosen basis  $x_i$  (say polynomial, spline, or wavelet basis) usually falls under the framework of variable subset selection as well (but with a different motivation). Such a representation can be practically useful in, e.g., compressing image signals, locating radar sources, or understanding principal components.

Suppose that we have response  $Y_n$  and design matrix  $X_n$  whose entries are  $n$  observations of  $[y, x_1, \dots, x_{d_n}]$ . For high-dimensional regression, a popular solution is to consider the following penalized regression that amalgamates variable selection and prediction simultaneously in operation. Solve

$$\hat{\beta} = \arg \min_{\beta} \left\{ \|Y_n - X_n \beta\|_2^2 + \sum_{j=1}^{d_n} p(|\beta_j|; \lambda, \gamma) \right\}, \quad (18)$$

and let  $\{i: \hat{\beta}_i \neq 0\}$  be the selected subset of variables. Here, the  $p(\beta; \lambda, \gamma)$  is a penalty function of  $\beta$  with tuning parameters  $\lambda, \gamma$  (which are usually determined by cross validation). It is crucial that the penalty function is not differentiable at  $\beta = 0$  so that the resulting solution becomes sparse when  $\lambda$  gets large.

Least absolute shrinkage and selection operator (LASSO) [58] in the form of  $p(\beta; \lambda) = \lambda |\beta|$  is perhaps the most commonly used penalty function. Here,  $\lambda$  is a tuning parameter that controls the strength of the penalty term. Increasing  $\lambda$  leads to fewer variables selected. In practice, data analysts can either 1) numerically sweep over a range of  $\lambda$  or 2) use the least-angle regression method [59] to find all of the possible candidate models (also called the *solution paths*) and then

**Table 1. The AR order selection: The average efficiency, dimension, and PI (along with standard errors).**

		AIC	BC	BIC
<b>Case 1</b>	Efficiency	0.78 (0.04)	0.93 (0.02)	0.99 (0.01)
	Dimension	3.95 (0.20)	3.29 (0.13)	3.01 (0.01)
	PI		0.93 (0.03)	
<b>Case 2</b>	Efficiency	0.77 (0.02)	0.76 (0.02)	0.56 (0.02)
	Dimension	9.34 (0.25)	9.29 (0.26)	5.39 (0.13)
	PI		0.13 (0.03)	
<b>Case 3</b>	Efficiency	0.71 (0.02)	0.67 (0.02)	0.55 (0.02)
	Size	6.99 (0.23)	6.61 (0.26)	4.02 (0.10)
	PI		0.35 (0.05)	

select the model with the best cross-validation performance. In a time series setting where LASSO solutions need to be continuously updated, fast online algorithms have been proposed (e.g., in [60]). Given that the data are truly generated by a linear model, tight prediction error bounds have been established for LASSO. Though originally designed for linear regression, LASSO has been also extended to a wide range of statistical models, such as generalized linear models (see [61] and the references therein).

Smoothly clipped absolute deviation (SCAD) [62] is another penalized regression that can correct the bias in LASSO estimates that comes from the  $\ell_1$ -penalty function being unbounded. It was also shown to exhibit oracle property, meaning that, as the sample size and model dimension go to infinity, all and only the true variables will be identified with probability going to one, the estimated parameters converge in probability to the true parameters, and the usual asymptotic normality holds as if all of the irrelevant variables have already been excluded. More discussions on such an oracle property will be included in the “Clarification of Some Misconceptions” section. The penalty of SCAD is in the form of

$$p(\beta; \lambda, \gamma) = \begin{cases} \lambda |t| & \text{if } |t| \leq \lambda \\ \frac{2\gamma\lambda |t| - t^2 - \lambda^2}{2(\gamma - 1)} & \text{if } \lambda < |t| \leq \gamma\lambda \\ \frac{\lambda^2(\gamma + 1)}{2} & \text{if } |t| > \gamma\lambda \end{cases}$$

In choosing a parsimonious set of variables, LASSO tends to overshrink the retained variables. In the SCAD penalty, the idea is to let  $\lambda$  and  $\gamma$  jointly control that the penalty first suppresses insignificant variables as LASSO does and then tapers off to achieve bias reduction. The tuning parameters in SCAD can be chosen by sweeping over a range of them and then applying cross validation.

Minimax concave penalty (MCP) [63] in the form of

$$p(\beta; \lambda, \gamma) = \begin{cases} \lambda |t| - \frac{t^2}{2\gamma} & \text{if } |t| \leq \gamma\lambda \\ \frac{\gamma\lambda^2}{2} & \text{if } |t| > \gamma\lambda \end{cases}$$

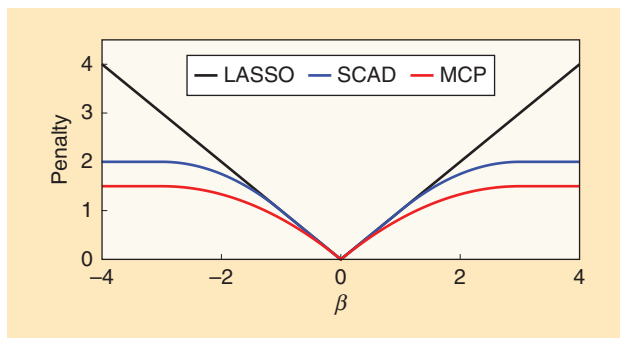


FIGURE 5. The penalties in LASSO, SCAD, and MCP.

is a penalized regression that works in a similar way as SCAD. Under some conditions, MCP attains minimax convergence rates in probability for the estimation of regression coefficients. Figure 5 illustrates the penalties in LASSO, SCAD, and MCP for  $\lambda = 1$  and  $\gamma = 3$ .

Elastic net [64] in the form of  $p(\beta; \lambda) = \lambda_1 |t| + \lambda_2 t^2$  is proposed to address several shortcomings of LASSO when

the covariates are highly correlated. The solution  $\hat{\beta}$  of the elastic net penalty exhibits mixed effects of the LASSO and ridge penalties. Recall that ridge regression in the form of  $p(\beta; \lambda) = \lambda t^2$  introduces bias to the regression estimates to reduce the large variances of ordinary least-squares estimates in the case of multicollinearity, and that LASSO tends to select a sparse subset. Interestingly, under elastic net, highly correlated covariates will tend to have similar

regression coefficients. This property, distinct from LASSO, is appealing in many applications when data analysts would like to find all of the associated covariates rather than selecting only one from each set of strongly correlated covariates.

Group LASSO [65] is another penalty introduced to restrict that all of the members of each predefined group of covariates are selected together. Different from (18), the penalty of the regression is not a sum of  $n$  terms but is replaced with  $\lambda \sum_{j=1}^r \|\beta_{I_j}\|_2$ , where  $\beta_{I_j}$  is a subvector of  $\beta$  indexed by  $I_j$  (the  $j$ th group), and  $I_1, \dots, I_r$  form a partition of  $\{1, \dots, n\}$ . It can be proved that  $\hat{\beta}_{I_j}$  is restricted to be vanishing together for each  $j$  [65]. The groups are often predefined using prior knowledge.

Adaptive LASSO [66] has been introduced to overcome the inconsistency in variable selection of LASSO. It replaces the penalty in (18) with  $\lambda \sum_{j=1}^{d_n} |\hat{\beta}_j|^{-u} |\beta_j|$ , where  $\hat{\beta}_j$  is referred to as a *pilot estimate* that can be obtained in various ways (e.g., by least squares for  $d_n < n$  or univariate regressions for  $d_n \geq n$ ). Adaptive LASSO was shown to exhibit the aforementioned oracle property. The adaptive LASSO can be solved by the same efficient algorithm for solving the LASSO, and it can be easily extended for generalized linear models as well.

In addition to the preceding penalized regression, a class of alternative solutions is known as *greedy algorithms* (or *stepwise algorithms*), which select a set of variables by making locally optimal decisions in each iteration.

Orthogonal matching pursuit (OMP) [67], [68], also referred to as the *forward stepwise regression algorithm*, is a very popular greedy algorithm that also inspired many other greedy algorithms. The general idea of OMP is to iteratively build a set of variables that are the most relevant to the response. It works in the following way. In each iteration, the variable most correlated with the current residual (in absolute value) is added to the subset (which is initialized as the empty set). Here, the residual represents the component of the observation vector  $y$  not in the linear span of the selected variables. Stopping criteria that guarantee good asymptotic properties, such as consistency in variable selection, remain an active line of research.

**A unified view on reconciling two such different goals wherever possible is a fundamental issue in model selection, and it remains an active line of research.**

The OMP algorithm can sequentially identify all of the significant variables with high probability under some conditions, such as weak dependences of the candidate variables (see, e.g., [69] and [70] and the references therein).

Least-angle regression (LARS) [59] is a greedy algorithm for stepwise variable selection. It can also be used for computing the solution paths of LASSO. Different from OMP, it does not permanently maintain a variable once it is selected into the model. Instead, it only adjusts the coefficient of the most correlated variable until that variable is no longer the most correlated with the recent residual. Briefly speaking, LARS works in the following way. It starts with all coefficients  $\beta_i$  being zeros. In each iteration, it looks for the variable  $x_i$  most correlated with the current residual  $r$  and increases its coefficient  $\beta_i$  in the direction of the sign of its correlation with  $y$ . Once some other variable  $x_j$  has the same correlation with  $r$  as  $x_i$  has, it increases  $\beta_i$  and  $\beta_j$  in the direction of their joint least squares until another variable has the same correlation with the residual. The procedure is repeated until all of the variables are in the model or the residuals have become zero.

### Properties of the penalized regression methods

Theoretical examinations of the penalized regression methods have mainly focused on the properties of tight prediction error bounds and consistency in selection. These asymptotic properties are mostly studied by assuming a parametric framework, i.e., data are truly generated by a linear regression model. Analysis for nonparametric, high-dimensional regression models has been also investigated in terms of oracle inequalities for prediction loss [71] and nonlinear additive models [72], [73].

The goal for prediction in high-dimensional regression focuses the control of the prediction loss (usually squared loss) bound so that it eventually vanishes even for a very large number of variables  $d_n$  (compared with the sample size  $n$ ). Suppose that data are generated, e.g., by  $Y_n = X_n\beta_* + \varepsilon$ , where  $Y_n \in \mathbb{R}^n$ ,  $\beta_* \in \mathbb{R}^{d_n}$ , and  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_n)$ . Let  $\|\beta_*\|_0$  denote the number of nonzero entries in  $\beta_*$ . Then, under certain restricted eigenvalue assumptions [71], there exist some constants  $c_1 > 2\sqrt{2}$  and  $c_2 > 0$  such that the LASSO solution satisfies  $n^{-1} \|X_n\beta_* - X_n\hat{\beta}\|_2^2 \leq c_2\sigma^2 \|\beta_*\|_0 n^{-1} \log d_n$  with probability at least  $1 - d_n^{1-c_1/8}$ , if we choose  $\lambda = c_1\sigma\sqrt{n \log d_n}$ . Note that the choice of  $\lambda$  depends on an unknown  $c_1\sigma$  that, though it does not scale with  $n$ , can have an effect for small sample size. Notably, the number of variables  $d_n$  is allowed to be much larger than  $n$  to admit a good predictive performance, as long as  $\log d_n$  is small compared with  $n$ . Similar tight bounds can be obtained by making other assumptions on  $\beta_*$  and  $X_n$ .

Selection consistency, as before, targets the goal of identifying the significant variables for scientific interpretation. The property of asymptotic efficiency we introduced before is rarely considered in high-dimensional regressions, because

it is implied by selection consistency in the parametric setting. For any vector  $\beta \in \mathbb{R}^{d_n}$ , let  $r(\beta)$  denote the indicator vector of  $\beta$  such that for any  $j = 1, \dots, d_n$ ,  $r_j(\beta) = 0$  if  $\beta_j = 0$ , and  $r_j(\beta) = 1$  otherwise. Selection consistency requires that the probability of  $r(\hat{\beta}) = r(\beta)$  converges in probability to one (as  $n \rightarrow \infty$ ). Under various conditions, such as fixed design or random design matrices, consistency of LASSO in estimating the significant variables has been widely studied under such

various technical conditions as sparsity, restricted isometry [74], mutual coherence [75], irrepresentable condition [76], and restricted eigenvalue [71], which create theoretical possibilities to distinguish the true subset of variables from all of the remaining subsets for large  $n$ .

At the same time, it has been known that LASSO is not generally consistent in parameter/coefficient estimation. This motivates the methods, such as SCAD, MCP, adaptive LASSO, and so forth, that correct the esti-

mation bias of LASSO. These three methods are also known to enjoy the so-called oracle property. The oracle property is perhaps more widely considered than selection consistency for high-dimensional regression analysis, because the penalized regression methods target simultaneous parameter estimation and prediction loss control. An oracle estimator [62] must be consistent in variable selection and parameter estimation, and satisfy 1) the sparsity condition, meaning that  $P_*\{r(\hat{\beta}) = r(\beta)\} \rightarrow 1$  as  $n \rightarrow \infty$ , where the inequality is componentwise; and 2) the asymptotic normality  $\sqrt{n}(\hat{\beta}_S - \beta_S) \rightarrow_d \mathcal{N}(0, I^{-1}(\beta_S))$ , where  $S$  is the support set of  $\beta$ ,  $\beta_S$  is the subvector of  $\beta_*$  indexed by  $S$ , and  $I(\beta_S)$  is the Fisher information knowing  $S$  in advance. Intuitively speaking, an oracle estimator enjoys the properties achieved by the MLE knowing the true support. We will revisit the oracle property in the ‘‘Controversy over the Oracle Property’’ section.

### Practical performance of penalized regression methods

With the huge influx of high-dimensional regression data, the penalized regression methods have been widely applied for sparse regression where a relatively small (or tiny) number of variables are selected out of a large number of candidates. In applications with a gene expression type of data, e.g., although the number of subjects may be only tens or hundreds, a sparse set of genes is typically selected out of thousands of choices. This has created a lot of excitement, with thousands of publications of such research and applications. This celebrated sparsity feature of penalized regression methods has generated an optimistic view that, even with, e.g., fewer than a hundred observations, the modern variable selection tool can identify a sparse subset out of thousands or even many more variables as the set of the most important ones for the regression problem. The estimated model is often readily used for data-driven discoveries.

There is little doubt that penalized regression methods have produced many successful results for the goal of prediction

**The war between AIC and BIC originates from two fundamentally different goals: one to minimize certain loss for prediction purpose and the other to select the best model for inference purpose.**

(see, e.g., [77]). As long as a proper cross validation is done for tuning parameter selection, the methods can often yield good predictive performance. However, given the challenge of high dimension and diverse data sources, the different penalized regression methods may have drastically different relative performance for various data sets. Therefore, proper choice of a method is important, to which end cross validation may be used, as will be presented in the next section.

For the goal of model selection for inference, however, the picture is much less promising. Indeed, many real applications strongly suggest that the practice of using the selected model for understanding and inference may be far from reliable. It has been reported that the selected variables from these penalized regression methods are often severely unstable, in the sense that the selection results can be drastically different under a tiny perturbation of data (see [78] and the references therein). Such high uncertainty damages reproducibility of the statistical findings [79]. Overall, being overly optimistic about the interpretability of high-dimensional regression methods can lead to spurious scientific discoveries.

The fundamental issue still lies in the potential discrepancy between inference and prediction, which is also elaborated in the “War and Peace—Conflicts Between AIC and BIC and Their Integration” and “Controversy over the Oracle Property” sections. If data analysts know in advance that the true model is exactly (or close to) a stable low-dimensional linear model, then the high-dimensional methods with the aforementioned oracle property may produce stable selection results not only good for prediction but also for inference purposes. Otherwise, the produced selection is so unstable that analysts can only focus on prediction alone. In practice, data analysts may need to utilize data-driven tools, such as model averaging [80], resampling [81], and confidence set for models [82], or model selection diagnostic, such as the parametricness index introduced in the “War and Peace—Conflicts Between AIC and BIC and Their Integration” section, to make sure the selected variables are stable and properly interpretable. Considerations along these lines also lead to stabilized variable selection methods [81], [83], [84]. The instability of penalized regression also motivated some recent research on postselection inference [85], [86]. Their interesting results in specific settings call for more research for more general applications.

## Modeling procedure selection

The discussions in the previous sections have focused on model selection in the narrow sense, where the candidates are models. In this section, we review the use of CV as a general tool for modeling procedure selection, which aims to select one from a finite set of modeling procedures [87]. Multiple modeling procedures such as AIC, BIC, and CV could be used for variable selection, and one of those procedures (together with the model selected by the procedure) is selected using an appropriately designed CV (which is at the second level). Another example is the emerging online competition platforms, such as Kaggle, that compare new problem-solving procedures and award prizes using cross validation.

The best procedure is defined in the sense that it outperforms, with high probability, the other procedures in terms of out-sample prediction loss for sufficiently large  $n$  (see, e.g., [13, Definition 1]).

There are two main goals of modeling procedure selection. The first is to identify with high probability the best procedure among the candidates. The property of selection consistency is of interest here. The second goal of modeling procedure selection is to approach the best performance (in terms of out-sample prediction loss) offered by the candidates, instead of pinpointing which candidate procedure is the best. Note again that, in case there are procedures that have similar best performances, we do not need to single out the best candidate to achieve the asymptotically optimal performance.

Similarly to model selection, for the task of modeling procedure selection, CV randomly splits  $n$  data into  $n_t$  training data and  $n_v$  validation data (so  $n = n_t + n_v$ ). The first  $n_t$  data are used to run different modeling procedures, and the remaining  $n_v$  data are used to assess the predictive performance. We will see that, for the first goal, the evaluation portion of CV should be large enough. For the second goal, a smaller portion of the evaluation may be enough to achieve optimal predictive performance.

In the literature, much attention has been focused on choosing whether to use the AIC procedure or BIC procedure for data analysis. For regression variable selection, it has been proved that the CV method is consistent in choosing between AIC and BIC given  $n_t \rightarrow \infty$ ,  $n_v/n_t \rightarrow \infty$ , and some other regularity assumptions [87, Th. 1]. In other words, the probability of BIC being selected goes to one in a parametric framework, and the probability of AIC being selected goes to one otherwise. In this way, the modeling procedure selection using CV naturally leads to a hybrid model selection criterion that builds upon strengths of AIC and BIC. Such a hybrid selection combines some theoretical advantages of both AIC and BIC. This aspect is seen in the context of the “War and Peace—Conflicts Between AIC and BIC and Their Integration” section. The task of classification is somewhat more relaxed compared with the task of regression. To achieve consistency in selecting the better classifier, the splitting ratio may be allowed to converge to infinity or any positive constant, depending on the situation [13]. In general, it is safe to let  $n_t \rightarrow \infty$  and  $n_v/n_t \rightarrow \infty$  for consistency in modeling procedure selection.

Closely related to this discussion is the following paradox. Suppose that a set of newly available data is given to an analyst. The analyst would naturally add some of the new data in the training phase and some in the validation phase. Clearly, with more data added to the training set, each candidate modeling procedure is improved in accuracy; with more data added to the validation set, the evaluation is also more reliable. It is tempting to think that improving the accuracy on both training and validation would lead to a sharper comparison between procedures. However, this is not the case. The prediction error estimation and procedure comparison are two different targets.

The cross-validation paradox says that better training and better estimation (e.g., in both bias and variance) of the prediction error by CV together do not imply better modeling procedure selection [13], [87]. Intuitively speaking, when comparing two procedures that are naturally close to each other, the improved estimation accuracy achieved by adopting more observations in the training part only makes the procedures more difficult to be distinguishable. The consistency in identifying the better procedure cannot be achieved unless the validation size diverges fast enough.

### Experiments

We illustrate the cross-validation paradox using the synthetic data generated from the linear regression model  $y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \varepsilon$ , where  $\beta = [1, 2, 0]^T$ , and the covariates  $X_j$  ( $j = 1, 2, 3$ ) and noise  $\varepsilon$  are independent standard Gaussian. Given  $n$  observations  $(y_i, x_{1,i}, x_{2,i}, x_{3,i})_{i=1, \dots, n}$ , we compare the following two different uses of linear regression. The first is based on  $X_1$  and  $X_2$ , and the second is based on all three covariates. Note that, in this experiment, selecting the better procedure is equivalent to selecting a better model. The data-generating model indicates that  $x_3$  is irrelevant for predicting  $y$ , so that the first procedure should be better than the second. Suppose that we start with 100 observations. We randomly split the data 100 times, each with 20 training data and 80 validation data, and record which procedure gives the smaller average quadratic loss during validation. We then add 50 new data to the training set and 50 to the validation set and record again which procedure is favored. We continue doing this until the sample size reaches 500. By running 1,000 independent replications, we summarize the frequency of the first procedure being favored in Table 2. As the paradox suggests, the accuracy of identifying the better procedure does not necessarily increase when more observations are added to both the estimation phase and the validation phase.

### Clarification of some misconceptions

#### *Pitfall of one-size-fits-all recommendation of data splitting ratio of cross validation*

There are widespread general recommendations on how to apply cross validation for model selection. It is stated in the literature, e.g., that tenfold CV is the best for model selection. Such guidelines seem to be unwarranted. First, it mistakenly disregards the goal of model selection. For prediction purposes, LOO is actually preferred in tuning parameter selection for traditional nonparametric regression. In contrast, for selection consistency, tenfold often leaves too few observations in evaluation to be stable. Indeed, fivefold often produces more stable selection results for high-dimensional regression. Second,  $k$ -fold CV, regardless of  $k$ , in general, is often unstable in the sense that a different dividing of data can produce a very different selection result. A common way to improve performance is to randomly divide the data into  $k$  folds several times and use the average validation loss for selection.

**Table 2. The cross-validation paradox: More observations in training and evaluation do not lead to higher selection accuracy in selecting the better procedure.**

Sample size $n$	100	200	300	400	500
Training size $n_t$	20	70	120	170	220
Accuracy	98.3%	94.9%	93.7%	92.3%	92.5%

For model selection, CV randomly splits  $n$  data into  $n_t$  training data and  $n_v$  validation data. Common practices using fivefold, tenfold, or 30% for validation do not exhibit asymptotic optimality (either consistency or asymptotic efficiency) in simple regression models, and their performances can be very different depending on the goal of applying CV. In fact, it is known that the delete- $n_v$  CV is asymptotically equivalent to  $\text{GIC}_{\lambda_n}$  with  $\lambda_n = n/(n - n_v) + 1$  for linear regression models under some assumptions [12]. It is also known that  $\text{GIC}_{\lambda_n}$  achieves asymptotic efficiency in a nonparametric framework only with  $\lambda_n = 2$ , and asymptotic efficiency in a parametric framework only with  $\lambda_n \rightarrow \infty$  (as  $n \rightarrow \infty$ ). In this context, from a theoretical perspective, the optimal splitting ratio  $n_v/n_t$  of CV should either converge to zero or diverge to infinity to achieve asymptotic efficiency, depending on whether the setting is nonparametric or parametric.

For modeling procedure selection, it is often necessary to let the validation size take a large proportion (e.g., half) to achieve good selection accuracy. In particular, the use of LOO for the goal of comparing procedures is the least trustworthy (see the “Modeling Procedure Selection” section).

### Experiment

We show how the splitting ratio can affect CV for model selection using the Modified National Institute of Standards and Technology database [88], which consists of 70,000 images of handwritten digits (from 0 to 9) with  $28 \times 28$  pixels. We implement six candidate feed-forward neural network models for classification. The first four models have one hidden layer, and the number of hidden nodes are 17, 18, 19, and 20. The fifth model has two hidden layers with 20 and four nodes; the sixth model has three hidden layers with 20, two, and two nodes. Because the true data-generating model for the real data is unavailable, we take 35,000 data (often referred to as the *test data*) out for approximating the true prediction loss and use the remaining data to train and validate. For model selection, we run CV with different  $n_v/n_t$ . For each ratio, we compute the average validation loss of each candidate model based on ten random partitions. We then select the model with the smallest average loss and calculate its true predictive loss using the remaining 35,000 data. The results recorded in Table 3 indicate that a smaller splitting ratio  $n_v/n_t$  leads to better classification accuracy. This is in line with the existing theory, because the neural network modeling is likely to be of nonparametric nature. This example also provides a complementing message to the cross-validation paradox. At ratio 0.95, the training sample size is too small to represent the full



**Table 3. The classification for handwritten digits: Smaller tends to give better predictive performance.**

<b>Ratio</b>	0.95	0.9	0.5	0.1	0.05
<b>Accuracy</b>	72.24%	90.28%	91.47%	91.47%	92.99%

sample size, so the ranking of the candidate models estimated from training data can be unstable and deviate from the ranking of models estimated from the full data set.

### *Because all models are wrong, why pursue consistency in selection?*

Because the reality is usually more complicated than a parametric model, perhaps everyone agrees that all models are wrong and the consistency concept of selecting the true model in a parametric framework cannot hold in the rigid sense. One view on such selection consistency is that, in many situations, a stable parametric model can be identified, and it can be treated as the true model. Such an idealization for theoretical investigation with practical implications is no more sinful than deriving theories under nonparametric assumptions. The true judge should be the performance in real applications. The notion of consistency in a nonparametric framework, however, is rarely used in the literature. In fact, it was shown that there does not exist any model selection method that can guarantee consistency in nonparametric regression settings (see, e.g., [25]). This partly explains why the concept of asymptotic efficiency (which is a weaker requirement) is more widely used in nonparametric frameworks.

### *Controversy over the oracle property*

The popular oracle property (as mentioned in the “Properties of the Penalized Regression Methods” section) for high-dimensional variable selection has been a focus in many research publications. However, it has been criticized by some researchers (see, e.g., [89]). At first glance, the oracle property may look very stringent. But we note that its requirement is fundamentally only as stringent as consistency in variable selection. In fact, if all of the true variables can be selected with probability tending to one by any method, then one can obtain MLE or the like restricted to the relevant variables for optimal estimation of the unknown parameters in the model. To our knowledge, there is neither claim nor reason to believe that the original estimator should be better than the refitted one by MLE based on the selected model. Though the oracle property is not theoretically surprising beyond consistency, it is still interesting and nontrivial to obtain such a property with only one stage of regression (as SCAD and MCP do). These methods, when armed with efficient algorithms, may save the computational cost in practice.

It was emphasized in [89] that the oracle estimator does not perform well in a uniform sense for point or interval estimation of the parameters. A price paid for the oracle property is that the risk of any oracle estimator (see [62]) has a supremum that diverges to infinity, i.e.,

$$\sup_{\beta \in \mathbb{R}^p} E_{\beta} \{n(\hat{\beta} - \beta)^T (\hat{\beta} - \beta)\} \rightarrow \infty,$$

as sample size  $n \rightarrow \infty$  (see, e.g., [89]). Here, we let  $E_{\beta}$  denote expectation with respect to the true linear model with coefficients  $\beta$ . In fact, for any consistent model selection method, we can always find a parameter value that is small enough so that the selection method tends to not include it (because it has to avoid overselection), yet the parameter value is big enough so that dropping it has a detrimental effect in rate of convergence (see, e.g., [49] and [90]). Although uniformity and robustness are valid and important considerations, we do not need to overly emphasize such properties. Otherwise, we are unduly burdened to retain not very useful variables in the final model and have to lose the ability in choosing a practically satisfying parsimonious model for interpretation and inference.

### **Some general recommendations**

Model selection, no matter how it is done, is exploratory in nature and cannot be confirmatory. Confirmatory conclusions can only be drawn based on well-designed follow-up studies. Nevertheless, good model selection tools can provide valuable and reliable information regarding explanation and prediction. Obviously there are many specific aspects of the data, nature of the models and practical considerations of the variables in the models, and so on that make each model selection problem unique to some degree. In spite of that, based on the literature and our own experiences, we give some general recommendations.

- 1) Keep in mind the main objective of model selection. First, if one needs to declare a model for inference, model selection consistency is the right concept to think about. Model selection diagnostic measures need to be used to assess the reliability of the selected model. In a high-dimensional setting, penalized regression methods are typically highly uncertain. For selection stability, when choosing a tuning parameter by cross validation, e.g., fivefold tends to work better than tenfold (see, e.g., [78]). Second, if one’s main goal is prediction, model selection instability is less of a concern, and any choice among the best performing models may give a satisfying prediction accuracy. In a parametric framework, consistent selection leads to asymptotic efficiency. In a nonparametric framework, selection methods based on the optimal tradeoff between estimation error and approximation error lead to asymptotic efficiency. When it is not clear if a (practically) parametric framework is suitable, we recommend the use of an adaptively asymptotic efficient method (e.g., the BC criterion).
- 2) When model selection is for prediction, the minimax consideration gives more protection in the worst case. If one postulates that the nature is adversary, the use of a minimax optimal criterion (e.g., AIC) is safer (than, e.g., BIC).
- 3) When prediction is the goal, one may consider different types of models and methods and then apply cross validation

to choose one for final prediction. If one needs to know which model or method is really the best, a large enough proportion (e.g., one-third or even half) for validation is necessary. If one just cares about the prediction accuracy and has little interest in declaring the chosen one being the best, the demand on the validation size may be much lessened.

## Acknowledgments

This research was funded in part by the Defense Advanced Research Projects Agency under grant W911NF-18-1-0134. We thank Dr. Shuguang Cui and eight anonymous reviewers for giving feedback on the initial submission of the manuscript. We are also grateful to Dr. Matthew McKay and Dr. Osvaldo Simeone for handling the full submission of the manuscript, and to three anonymous reviewers for their comprehensive comments that have greatly improved the article.

## Authors

**Jie Ding** (dingj@umn.edu) received his Ph.D. degree in engineering sciences from Harvard University, Cambridge, in 2017, where he worked as a postdoctoral researcher from January 2017 to December 2018. He was a postdoctoral researcher with the Information Initiative at Duke University, Durham, North Carolina, from January to August 2018. He is currently an assistant professor with the School of Statistics, University of Minnesota, Minneapolis. His research topics include signal processing, statistical inference, data prediction, and machine learning with applications to multimedia processing, human-computer interface, finance, and so on.

**Vahid Tarokh** (vahid.tarokh@duke.edu) received his Ph.D. degree in electrical engineering from the University of Waterloo, Ontario, Canada, in 1995. He worked at AT&T Labs-Research and AT&T Wireless Services until August 2000 as a member, principal member of technical staff and, finally, as the head of the Department of Wireless Communications and Signal Processing. In September 2000, he joined the Massachusetts Institute of Technology, Cambridge, as an associate professor of electrical engineering and computer science. In June 2002, he joined Harvard University, Cambridge, as a Gordon McKay Professor of Electrical Engineering and Hammond Vinton Hayes Senior Research Fellow. He was named Perkins Professor of Applied Mathematics and Hammond Vinton Hayes Senior Research Fellow of Electrical Engineering in 2005. In January 2018, he joined Duke University as the Rhodes Family Professor of Electrical and Computer Engineering, Computer Science, and Mathematics. From January 2018 to May 2018, he was also a Gordon Moore Distinguished Scholar with the California Institute of Technology, Pasadena. He is a Fellow of the IEEE.

**Yuhong Yang** (yyang@stat.umn.edu) received his Ph.D. degree in statistics from Yale University, New Haven, Connecticut, in 1996. He first joined the Department of Statistics at Iowa State University, Ames, then moved to the University of Minnesota, Minneapolis, in 2004, where he has

been a full professor since 2007. His research interests include model selection, multiarmed bandit problems, forecasting, high-dimensional data analysis, and machine learning. He has published in journals in several fields, including *Annals of Statistics*, *IEEE Transactions on Information Theory*, *Journal of Econometrics*, *Journal of Approximation Theory*, *Journal of Machine Learning Research*, and *International Journal of Forecasting*. He is a fellow of the Institute of Mathematical Statistics.

## References

- [1] P. Stoica and Y. Selen, "Model-order selection: A review of information criterion rules," *IEEE Signal Process. Mag.*, vol. 21, no. 4, pp. 36–47, 2004.
- [2] J. B. Kadane and N. A. Lazar, "Methods and criteria for model selection," *J. Amer. Statist. Assoc.*, vol. 99, no. 465, pp. 279–290, 2004.
- [3] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
- [4] S. Greenland, "Modeling and variable selection in epidemiologic analysis," *Am. J. Public Health*, vol. 79, no. 3, pp. 340–349, 1989.
- [5] C. M. Andersen and R. Bro, "Variable selection in regression: A tutorial," *J. Chemometrics*, vol. 24, nos. 11–12, pp. 728–737, 2010.
- [6] J. B. Johnson and K. S. Omland, "Model selection in ecology and evolution," *Trends Ecol. Evolut.*, vol. 19, no. 2, pp. 101–108, 2004.
- [7] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. New York: Springer-Verlag, 2003.
- [8] M. Parry, A. P. Dawid, and S. Lauritzen, "Proper local scoring rules," *Ann. Statist.*, vol. 40, no. 1, pp. 561–592, 2012.
- [9] V. Vapnik and A. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," in *Measures of Complexity*, V. Vovk, H. Papadopoulos, and A. Gammerman, Eds. New York: Springer-Verlag, 1971, pp. 11–30.
- [10] C.-K. Ing and C.-Z. Wei, "Order selection for same-realization predictions in autoregressive processes," *Ann. Statist.*, vol. 33, no. 5, pp. 2423–2474, 2005.
- [11] R. Shibata, "Asymptotically efficient selection of the order of the model for estimating parameters of a linear process," *Ann. Statist.*, vol. 8, no. 1, pp. 147–164, 1980.
- [12] J. Shao, "An asymptotic theory for linear model selection," *Stat. Sinca*, vol. 7, no. 2, pp. 221–242, 1997.
- [13] Y. Yang, "Comparing learning methods for classification," *Stat. Sinca*, vol. 16, pp. 635–657, 2006.
- [14] W. Liu and Y. Yang, "Parametric or nonparametric? A parametricness index for model selection," *Ann. Statist.*, vol. 39, no. 4, pp. 2074–2102, 2011.
- [15] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control*, vol. 19, no. 6, pp. 716–723, 1974.
- [16] H. Akaike, "Fitting autoregressive models for prediction," *Ann. Inst. Statist. Math.*, vol. 21, no. 1, pp. 243–247, 1969.
- [17] K. Takeuchi, "Distribution of informational statistics and a criterion of model fitting," *Suri-Kagaku (Mathematical Sciences)*, vol. 153, pp. 12–18, 1976.
- [18] W. Pan, "Akaike's information criterion in generalized estimating equations," *Biometrics*, vol. 57, no. 1, pp. 120–125, 2001.
- [19] C. M. Hurvich and C.-L. Tsai, "Regression and time series model selection in small samples," *Biometrika*, vol. 76, no. 2, pp. 297–307, 1989.
- [20] P. M. Broersen, "Finite sample criteria for autoregressive order selection," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3550–3558, 2000.
- [21] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 1978.
- [22] J. K. Ghosh and R. V. Ramamoorthi, *Bayesian Nonparametrics*. New York: Springer, 2003, p. 36.
- [23] E. J. Hannan and B. G. Quinn, "The determination of the order of an autoregression," *J. Roy. Statist. Soc. Ser. B*, vol. 41, no. 2, pp. 190–195, 1979.
- [24] J. Ding, V. Tarokh, and Y. Yang, "Bridging AIC and BIC: A new criterion for autoregression," *IEEE Trans. Inf. Theory*, vol. 64, no. 6, pp. 4024–4043, 2018.
- [25] J. Ding, V. Tarokh, and Y. Yang, (2016). Optimal variable selection in regression models. [Online]. Available: <http://jdjng.org/jie-uploads/2017/11/regression.pdf>
- [26] P. M. Djuric, "Asymptotic map criteria for model selection," *IEEE Trans. Signal Process.*, vol. 46, no. 10, pp. 2726–2735, 1998.
- [27] C. Andrieu, P. Djuric, and A. Doucet, "Model selection by MCMC computation," *Signal Process.*, vol. 81, no. 1, pp. 19–37, 2001.

- [28] S. Shao, P. E. Jacob, J. Ding, and V. Tarokh, 2018. "Bayesian model comparison with the Hyvarinen score: Computation and consistency," *J. Am. Statist. Assoc.* doi: 10.1080/01621459.2018.1518237.
- [29] C. S. Wallace and D. M. Boulton, "An information measure for classification," *Computer J.*, vol. 11, no. 2, pp. 185–194, 1968.
- [30] M. A. T. Figueiredo and A. K. Jain, "Unsupervised learning of finite mixture models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 3, pp. 381–396, 2002.
- [31] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, no. 5, pp. 465–471, 1978.
- [32] J. Rissanen, "Estimation of structure by minimum description length," *Circuits Syst. Signal Process.*, vol. 1, no. 3, pp. 395–406, 1982.
- [33] A. Barron, J. Rissanen, and B. Yu, "The minimum description length principle in coding and modeling," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2743–2760, 1998.
- [34] M. H. Hansen and B. Yu, "Model selection and the principle of minimum description length," *J. Amer. Statist. Assoc.*, vol. 96, no. 454, pp. 746–774, 2001.
- [35] J. Rissanen, "Stochastic complexity and modeling," *Ann. Statist.*, vol. 14, no. 3, pp. 1080–1100, 1986.
- [36] A. P. Dawid, "Present position and potential developments: Some personal views: Statistical theory: The prequential approach," *J. Roy. Statist. Soc. Ser. A*, vol. 147, no. 2, pp. 278–292, 1984.
- [37] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, "Bayesian measures of model complexity and fit," *J. Roy. Statist. Soc. Ser. B*, vol. 64, no. 4, pp. 583–639, 2002.
- [38] D. M. Allen, "The relationship between variable selection and data augmentation and a method for prediction," *Technometrics*, vol. 16, no. 1, pp. 125–127, 1974.
- [39] S. Geisser, "The predictive sample reuse method with applications," *J. Amer. Statist. Assoc.*, vol. 70, no. 350, pp. 320–328, 1975.
- [40] M. Stone, "An asymptotic equivalence of choice of model by cross-validation and Akaike's criterion," *J. Roy. Statist. Soc. Ser. B*, vol. 39, no. 1, pp. 44–47, 1977.
- [41] J. Rissanen, "A predictive least-squares principle," *IMA J. Math. Control Inform.*, vol. 3, no. 2–3, pp. 211–222, 1986.
- [42] C.-Z. Wei, "On predictive least squares principles," *Ann. Statist.*, vol. 20, no. 1, pp. 1–42, 1992.
- [43] C.-K. Ing, "Accumulated prediction errors, information criteria and optimal forecasting for autoregressive time series," *Ann. Statist.*, vol. 35, no. 3, pp. 1238–1277, 2007.
- [44] R. Nishii, "Asymptotic properties of criteria for selection of variables in multiple regression," *Ann. Statist.*, vol. 12, no. 2, pp. 758–765, 1984.
- [45] C. L. Mallows, "Some comments on CP," *Technometrics*, vol. 15, no. 4, pp. 661–675, 1973.
- [46] A. Barron, L. Birgé, and P. Massart, "Risk bounds for model selection via penalization," *Probab. Theory Relat. Fields*, vol. 113, no. 3, pp. 301–413, 1999.
- [47] R. Shibata, "An optimal selection of regression variables," *Biometrika*, vol. 68, no. 1, pp. 45–54, 1981.
- [48] R. Shibata, "Selection of the order of an autoregressive model by Akaike's information criterion," *Biometrika*, vol. 63, no. 1, pp. 117–126, 1976.
- [49] Y. Yang, "Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation," *Biometrika*, vol. 92, no. 4, pp. 937–950, 2005.
- [50] D. P. Foster and E. I. George, "The risk inflation criterion for multiple regression," *Ann. Statist.*, vol. 22, no. 4, pp. 1947–1975, 1994.
- [51] S. S. Wilks, "The large-sample distribution of the likelihood ratio for testing composite hypotheses," *Ann. Math. Statist.*, vol. 9, no. 1, pp. 60–62, 1938.
- [52] T. van Erven, P. Grünwald, and S. De Rooij, "Catching up faster by switching sooner: A predictive approach to adaptive estimation with an application to the AIC–BIC dilemma," *J. Roy. Statist. Soc. Ser. B*, vol. 74, no. 3, pp. 361–417, 2012.
- [53] D. L. Donoho, and I. M. Johnstone, "Minimax estimation via wavelet shrinkage," *Ann. Statist.*, vol. 26, no. 3, pp. 879–921, 1998.
- [54] I. Daubechies, M. Deffrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [55] E. Candès, and T. Tao, "The Dantzig selector: Statistical estimation when  $p$  is much larger than  $n$ ," *Ann. Statist.*, vol. 35, no. 6, pp. 2313–2351, 2007.
- [56] S. Foucart and M.-J. Lai, "Sparsest solutions of underdetermined linear systems via  $\ell_q$ -minimization for  $0 < q \leq 1$ ," *Appl. Comput. Harmon. A*, vol. 26, no. 3, pp. 395–407, 2009.
- [57] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [58] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," *Ann. Statist.*, vol. 32, no. 2, pp. 407–499, 2004.
- [59] B. Babadi, N. Kalouptsidis, and V. Tarokh, "SPARLS: The sparse RLS algorithm," *IEEE Trans. Signal Process.*, vol. 58, no. 8, pp. 4013–4025, 2010.
- [60] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity: The LASSO and Generalizations*. Boca Raton, FL: CRC, 2015.
- [61] J. Fan and R. Li, "Variable selection via nonconcave penalized likelihood and its oracle properties," *J. Amer. Statist. Assoc.*, vol. 96, no. 456, pp. 1348–1360, 2001.
- [62] C.-H. Zhang, "Nearly unbiased variable selection under minimax concave penalty," *Ann. Statist.*, vol. 38, no. 2, pp. 894–942, 2010.
- [63] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *J. Roy. Statist. Soc. Ser. B*, vol. 67, no. 2, pp. 301–320, 2005.
- [64] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Roy. Statist. Soc. Ser. B*, vol. 68, no. 1, pp. 49–67, 2006.
- [65] H. Zou, "The adaptive lasso and its oracle properties," *J. Amer. Statist. Assoc.*, vol. 101, no. 476, pp. 1418–1429, 2006.
- [66] S. Chen, S. A. Billings, and W. Luo, "Orthogonal least squares methods and their application to non-linear system identification," *Int. J. Control*, vol. 50, no. 5, pp. 1873–1896, 1989.
- [67] Y. C. Pati, R. Rezaifar, and P. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. IEEE 7th Asilomar Conf. Signals, Systems, Computers*, 1993, pp. 40–44.
- [68] C.-K. Ing and T. L. Lai, "A stepwise regression method and consistent model selection for high-dimensional sparse linear models," *Statist. Sinca*, vol. 21, pp. 1473–1513, Oct. 2011.
- [69] J. Ding, L. Chen, and Y. Gu, "Perturbation analysis of orthogonal matching pursuit," *IEEE Trans. Signal Process.*, vol. 61, no. 2, pp. 398–410, 2013.
- [70] P. J. Bickel, Y. Ritov, and A. B. Tsybakov, "Simultaneous analysis of LASSO and Dantzig selector," *Ann. Statist.*, vol. 37, no. 4, pp. 1705–1732, 2009.
- [71] J. Lafferty and L. Wasserman, "Rodeo: Sparse, greedy nonparametric regression," *Ann. Statist.*, vol. 36, no. 1, pp. 28–63, 2008.
- [72] Q. Han, J. Ding, E. M. Airoldi, and V. Tarokh, "SLANTS: Sequential adaptive nonlinear modeling of time series," *IEEE Trans. Signal Process.*, vol. 65, no. 19, pp. 4994–5005, 2017.
- [73] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [74] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.
- [75] P. Zhao and B. Yu, "On model selection consistency of LASSO," *J. Mach. Learn. Res.*, vol. 7, pp. 2541–2563, Nov. 2006.
- [76] S. Yang, M. Santillana, and S. Kou, "Accurate estimation of influenza epidemics using Google search data via ARGO," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 112, no. 47, pp. 14,473–14,478, 2015.
- [77] Y. Nan and Y. Yang, "Variable selection diagnostics measures for high-dimensional regression," *J. Comp. Graph. Statist.*, vol. 23, no. 3, pp. 636–656, 2014.
- [78] J. P. Ioannidis, "Why most published research findings are false," *PLoS Medicine*, vol. 2, no. 8, pp. 696–701, 2005.
- [79] Y. Yang, "Adaptive regression by mixing," *J. Amer. Statist. Assoc.*, vol. 96, no. 454, pp. 574–588, 2001.
- [80] N. Meinshausen and P. Bühlmann, "Stability selection," *J. Roy. Statist. Soc. Ser. B*, vol. 72, no. 4, pp. 417–473, 2010.
- [81] D. Ferrari and Y. Yang, "Confidence sets for model selection by  $f$ -testing," *Stat. Sinca*, vol. 25, no. 4, pp. 1637–1658, 2015.
- [82] C. Lim and B. Yu, "Estimation stability with cross-validation (ESCV)," *J. Comput. Graph. Statist.*, vol. 25, no. 2, pp. 464–492, 2016.
- [83] W. Yang and Y. Yang, "Toward an objective and reproducible model choice via variable selection deviation," *Biometrics*, vol. 73, no. 1, pp. 20–30, 2017.
- [84] R. Berk, L. Brown, A. Buja, K. Zhang, and L. Zhao, "Valid post-selection inference," *Ann. Statist.*, vol. 41, no. 2, pp. 802–837, 2013.
- [85] J. Taylor, R. Lockhart, R. J. Tibshirani, and R. Tibshirani. (2014). Post-selection adaptive inference for least angle regression and the LASSO. [Online]. Available: [https://www.researchgate.net/publication/259764783\\_Post-selection\\_adaptive\\_inference\\_for\\_Least\\_Angle\\_Regression\\_and\\_the\\_Lasso](https://www.researchgate.net/publication/259764783_Post-selection_adaptive_inference_for_Least_Angle_Regression_and_the_Lasso)
- [86] Y. Zhang and Y. Yang, "Cross-validation for selecting a model selection procedure," *J. Econometrics*, vol. 187, no. 1, pp. 95–112, 2015.
- [87] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [88] H. Leeb and B. M. Pötscher, "Sparse estimators and the oracle property, or the return of Hodges estimator," *J. Econom.*, vol. 142, no. 1, pp. 201–211, 2008.
- [89] Y. Yang, "Prediction/estimation with simple linear models: Is it really that simple?" *Econom. Theory*, vol. 23, no. 1, pp. 1–36, 2007.
- [90] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.