

OPTIMAL CROSS-VALIDATION IN DENSITY ESTIMATION WITH THE L^2 -LOSS¹

BY ALAIN CELISSE

UMR 8524 CNRS–Université Lille 1

We analyze the performance of cross-validation (CV) in the density estimation framework with two purposes: (i) *risk estimation* and (ii) *model selection*. The main focus is given to the so-called leave- p -out CV procedure (Lpo), where p denotes the cardinality of the test set. Closed-form expressions are settled for the Lpo estimator of the risk of projection estimators. These expressions provide a great improvement upon V -fold cross-validation in terms of variability and computational complexity.

From a theoretical point of view, closed-form expressions also enable to study the Lpo performance in terms of risk estimation. The optimality of leave-one-out (Loo), that is Lpo with $p = 1$, is proved among CV procedures used for risk estimation. Two model selection frameworks are also considered: *estimation*, as opposed to *identification*. For estimation with finite sample size n , optimality is achieved for p large enough [with $p/n = o(1)$] to balance the overfitting resulting from the structure of the model collection. For identification, model selection consistency is settled for Lpo as long as p/n is conveniently related to the rate of convergence of the best estimator in the collection: (i) $p/n \rightarrow 1$ as $n \rightarrow +\infty$ with a parametric rate, and (ii) $p/n = o(1)$ with some nonparametric estimators. These theoretical results are validated by simulation experiments.

1. Introduction. For estimating a target quantity denoted by s , let $\{S_m\}_{m \in \mathcal{M}}$ denote a collection of sets of candidate parameters indexed by \mathcal{M} . From each S_m called a *model*, an estimator \hat{s}_m of s is computed. The goal of model selection is to design a criterion $\text{crit} : \mathcal{M} \rightarrow \mathbb{R}^+$ such that minimizing $\text{crit}(\cdot)$ over \mathcal{M} provides a final estimator $\hat{s}_{\hat{m}}$ that is “optimal.” Among various strategies of model selection, *model selection via penalization* has been introduced in the seminal papers by Akaike (1973), Mallows (1973), Schwarz (1978) on, respectively, AIC, C_p and BIC criteria. However, since AIC and BIC are derived from asymptotic arguments, their performances crucially depend on model collection and sample size [see Baraud, Giraud and Huet (2009)].

Received March 2012; revised May 2014.

¹Supported by the French Agence Nationale de la Recherche (ANR) under reference ANR-09-JCJC-0027-01 and ANR-11-BS01-0010.

MSC2010 subject classifications. Primary 62G09; secondary 62G07, 62E17.

Key words and phrases. Cross-validation, leave- p -out, resampling, risk estimation, model selection, density estimation, oracle inequality, projection estimators, concentration inequalities.

More recently, [Birgé and Massart \(1997, 2001, 2007\)](#) have developed a nonasymptotic approach inspired from the pioneering work of [Barron and Cover \(1991\)](#). It relies on concentration inequalities [[Ledoux \(2001\)](#), [Talagrand \(1996\)](#)] and aims at deriving *oracle inequalities* such as

$$(1) \quad \ell(s, \widehat{s}_m) \leq C \inf_{m \in \mathcal{M}} \{\ell(s, \widehat{s}_m)\} + r_n$$

with probability larger than $1 - c/n^2$, where $c > 0$ is a constant, $\ell(s, t)$ is a measure of the gap between parameters s and t , r_n is a remainder term with respect to $\inf_m \ell(s, \widehat{s}_m)$, and $C \geq 1$ denotes a constant independent of s . The closer C to 1 and the smaller r_n , the better the model selection procedure. If $C = C_n \rightarrow 1$ as $n \rightarrow +\infty$, the model selection procedure is said *asymptotically optimal* (or efficient) [see, e.g., [Arlot and Celisse \(2010\)](#)]. Note that other asymptotic optimality properties have been studied in the literature. For instance, a model selection procedure satisfying

$$\mathbb{P}[\widehat{m} = m_0] \xrightarrow{n \rightarrow +\infty} 1,$$

where m_0 denotes a fixed given model is said *model selection consistent* [see [Shao \(1997\)](#) for a study of various model selection procedures in terms of model selection consistency].

In the density estimation framework, model selection with deterministic penalties has been addressed: (i) for Kullback–Leibler divergence by [Barron, Birgé and Massart \(1999\)](#), [Castellan \(1999, 2003\)](#), [Yang and Barron \(1998\)](#) and further studied in [Birgé and Rozenholc \(2006\)](#), and (ii) for quadratic risk and projection estimators by [Birgé and Massart \(1997\)](#) and [Barron, Birgé and Massart \(1999\)](#).

The aforementioned approaches rely on some deterministic penalties such as AIC or BIC. These penalties are derived in some specific settings [e.g., a Gaussian noise is assumed by [Birgé and Massart \(2007\)](#)] and remain unjustified and even sometimes misleading in more general settings.

Conversely, cross-validation (CV) is a *resampling* procedure based on a *universal heuristics* which makes it applicable in a wide range of settings. CV procedures have been first studied in a regression context by [Stone \(1974, 1977\)](#) for the leave-one-out (Loo) and [Geisser \(1974, 1975\)](#) for the V -fold cross-validation (VFCV), and in the density estimation framework by [Rudemo \(1982\)](#), [Stone \(1984\)](#). Since these procedures can be computationally demanding or even intractable, [Bowman \(1984\)](#), [Rudemo \(1982\)](#) derived closed-form formulas for the Loo estimator of the risk of histograms or kernel estimators. These results have been recently extended to the leave- p -out cross-validation (Lpo) by [Celisse and Robin \(2008\)](#).

Although CV procedures are extensively used in practice, only few theoretical results exist on their performances, most of them being of asymptotic nature. For instance, in the regression framework, [Burman \(1989, 1990\)](#) proves Loo is asymptotically the best CV procedure in terms of risk estimation. Several papers are dedicated to show the equivalence between some CV procedures and penalized criteria

in terms of asymptotic optimality properties: (i) *efficiency* in Li (1987), Zhang (1993), and (ii) *model selection consistency* in Shao (1997), Yang (2006, 2007). Let us notice that in the parametric setting, Yang (2007) proved that efficiency and model selection consistency are contradictory objectives that cannot be achieved simultaneously. We refer interested readers to Shao (1997) for an extensive review about asymptotic optimality properties in terms of efficiency and model selection consistency of some penalized criteria as well as CV procedures.

As for nonasymptotic results in the density framework, Birgé and Massart (1997) have settled an oracle inequality that relies on a conjecture and may be applied to Loo. However, to the best of our knowledge, no such result has already been proved for Lpo in the density estimation framework. Recently, in the regression setting, Arlot (2007) established oracle inequalities for V -fold penalties, while Arlot and Celisse (2011) have carried out an extensive simulation study in the change-point detection problem with heteroscedastic observations.

In the present paper, we derive closed-form expressions for the Lpo risk estimator of the broad class of projection estimators (Section 2). Such closed-form expressions considerably improve upon V -FCV in terms of (i) variability [Celisse and Robin (2008)], and (ii) computational complexity (Section 2.3). A second improvement allowed by these formulas is the deep new understanding of the theoretical performance of CV in two respects: first for risk estimation (Section 2.4), and second for model selection (Section 3). For instance, it is proved that Loo is the best CV procedure for risk estimation (Theorem 2.1), while the story can be different for model selection (Corollary 3.1 and Theorems 3.3 and 3.4).

In Section 3, two aspects of model selection *via* CV have been explored. The *estimation* point of view is described in Section 3.1. It is shown that Lpo is optimal as long as $p/n = o(1)$ and p is large enough to balance the influence of the model collection structure. This phenomenon is supported by simulation experiments detailed in Section 3.1.4. Finally, Section 3.2 deals with the *identification* point of view. CV is proved to be model selection consistent in various settings where the choice of p is related to the convergence rate (parametric and nonparametric) of the best estimator one tries to recover. Simulation results illustrate these different behaviors in Section 3.2.2. Finally, a discussion is provided in Section 4 to give some guidelines toward a better understanding of CV procedures. The main proofs have been postponed to the Appendix. For reasons owing to space constraints, more technical ones are provided in the supplementary material [Celisse (2014)].

2. Cross-validation and risk estimation.

2.1. Statistical framework.

2.1.1. *Notation.* Throughout the paper, $X_1, \dots, X_n \in [0, 1]$ are independent and identically distributed (i.i.d.) random variables drawn from a probability distribution P of density $s \in L^2([0, 1])$ with respect to Lebesgue's measure on $[0, 1]$, and $X_{1,n} = (X_1, \dots, X_n)$.

Let \mathcal{S}^* denote the set of measurable functions on $[0, 1]$. The distance between s and any $u \in \mathcal{S}^*$ is measured by the quadratic *loss* denoted by

$$\ell : (s, u) \mapsto \ell(s, u) := \|s - u\|^2 = \int_{[0,1]} [s(t) - u(t)]^2 dt.$$

It is related to the *contrast* function

$$(2) \quad \gamma : (u, x) \mapsto \gamma(u; x) := \|u\|^2 - 2u(x) \quad \text{with } \ell(s, u) = P\gamma(u) - P\gamma(s),$$

where $P\gamma(u) = P(\gamma(u; \cdot))$ and $Pf := \mathbb{E}[f(X_1)]$ for every $f \in \mathcal{S}^*$. The performance of an estimator $\widehat{s} = \widehat{s}(X_1, \dots, X_n)$ of s is assessed by the *quadratic risk*

$$R_n(\widehat{s}) := \mathbb{E}[\ell(s, \widehat{s})] = \mathbb{E}[\|s - \widehat{s}\|^2].$$

Estimating $P\gamma(u)$ is made through the *empirical contrast* defined by

$$(3) \quad P_n\gamma(u) := \frac{1}{n} \sum_{i=1}^n \gamma(u; X_i) \quad \text{where } P_n = 1/n \sum_{i=1}^n \delta_{X_i}$$

denotes the empirical measure and $P_n f := 1/n \sum_{i=1}^n f(X_i)$ for every $f \in \mathcal{S}^*$.

For a collection of models $\{S_m\}_{m \in \mathcal{M}_n}$ indexed by a countable set \mathcal{M}_n , the *empirical contrast minimizer* is defined by

$$(4) \quad \widehat{s}_m := \underset{u \in S_m}{\text{Argmin}} P_n\gamma(u).$$

It results a collection $\{\widehat{s}_m\}_{m \in \mathcal{M}_n}$ of estimators of s depending on the choice of models S_m s. Instances of such models and estimators are described in Section 2.1.2.

2.1.2. Projection estimators. Let Λ_n be a set of countable indices and $\{\varphi_\lambda\}_{\lambda \in \Lambda_n}$ a family of vectors in $L^2([0, 1])$ such that for every $m \in \mathcal{M}_n$, $\{\varphi_\lambda\}_{\lambda \in \Lambda(m)}$ denotes an orthonormal family of $L^2([0, 1])$ with $\Lambda(m) \subset \Lambda_n$. For every $m \in \mathcal{M}_n$, S_m denotes the linear space spanned by $\{\varphi_\lambda\}_{\lambda \in \Lambda(m)}$, $D_m = \dim(S_m)$, and s_m is the *orthogonal projection* of s onto S_m

$$s_m := \underset{u \in S_m}{\text{Argmin}} P\gamma(u) = \sum_{\lambda \in \Lambda(m)} P\varphi_\lambda \varphi_\lambda \quad \text{with } P\varphi_\lambda = \mathbb{E}[\varphi_\lambda(X)].$$

DEFINITION 2.1. An estimator $\widehat{s} \in L^2([0, 1])$ is a projection estimator if there exists a family $\{\varphi_\lambda\}_{\lambda \in \Lambda}$ of orthonormal vectors of $L^2([0, 1])$ such that

$$\widehat{s} = \sum_{\lambda \in \Lambda} \theta_\lambda \varphi_\lambda \quad \text{with } \theta_\lambda = \frac{1}{n} \sum_{i=1}^n H_\lambda(X_i),$$

where $\{H_\lambda(\cdot)\}_{\lambda \in \Lambda}$ depends on the family $\{\varphi_\lambda\}_{\lambda \in \Lambda}$.

It is straightforward to check that the empirical contrast minimizer over $S_m = \text{Span}(\varphi_\lambda, \lambda \in \Lambda(m))$, defined by equation (4), is a projection estimator since

$$(5) \quad \widehat{s}_m = \sum_{\lambda \in \Lambda(m)} P_n \varphi_\lambda \varphi_\lambda \quad \text{with } P_n \varphi_\lambda = \frac{1}{n} \sum_{i=1}^n \varphi_\lambda(X_i).$$

Here are a few examples of projection estimators [see DeVore and Lorentz (1993)]:

- *Histograms*: For every $m \in \mathcal{M}_n$, let $\{I_\lambda\}_{\lambda \in \Lambda(m)}$ be a partition of $[0, 1]$ in $D_m = \text{Card}(\Lambda(m))$ intervals. Set $\varphi_\lambda = \mathbb{1}_{I_\lambda} / \sqrt{|I_\lambda|}$ for every $\lambda \in \Lambda(m)$, with $|I_\lambda|$ the Lebesgue measure of I_λ , and $\mathbb{1}_{I_\lambda}(x) = 1$ if $x \in I_\lambda$ and 0 otherwise. Then

$$(6) \quad \widehat{s}_m = \sum_{\lambda \in \Lambda(m)} P_n \mathbb{1}_{I_\lambda} \frac{\mathbb{1}_{I_\lambda}}{|I_\lambda|}.$$

- *Trigonometric polynomials*: For every $\lambda \in \mathcal{Z}$, let $\varphi_\lambda : t \mapsto \varphi_\lambda(t) = e^{2\pi i \lambda t}$. Then for any finite $\Lambda(m) \subset \mathbb{Z}$,

$$(7) \quad \widehat{s}_m(t) = \sum_{\lambda \in \Lambda(m)} P_n \varphi_\lambda e^{2\pi i \lambda t} \quad \forall t \in [0, 1]$$

is a trigonometric polynomial.

- *Wavelet basis*: Let $\{\varphi_\lambda\}_{\lambda \in \Lambda_n}$ be an orthonormal basis of $L^2([0, 1])$ made of compact supported wavelets, where $\Lambda_n = \{(j, k) \mid j \in \mathbb{N}^* \text{ and } 1 \leq k \leq 2^j\}$. Then for every subset $\Lambda(m)$ of Λ_n ,

$$(8) \quad \widehat{s}_m = \sum_{\lambda \in \Lambda(m)} P_n \varphi_\lambda \varphi_\lambda.$$

Some of these estimators can take negative values. A possible solution is truncating and normalizing the preliminary projection estimator

$$\widetilde{s}_m = \widehat{s}_m \mathbb{1}_{\widehat{s}_m \geq 0} \left(\int_{[0, 1]} \mathbb{1}_{\widehat{s}_m \geq 0}(t) \widehat{s}_m(t) dt \right)^{-1}.$$

However, the closed-form expressions provided in Section 2.3 are not available for these truncated and normalized estimators.

2.2. Leave- p -out cross-validation. In the literature, several cross-validation (CV) procedures have been successively introduced to overcome the defects of already existing ones. Let us describe the main CV procedures with some emphasis to computational aspects.

2.2.1. Cross-validation. For $1 \leq p \leq n-1$, let us define $\mathcal{E}_p = \{e \subset \{1, \dots, n\}, \text{Card}(e) = p\}$ and for $e \in \mathcal{E}_p$, set $X^e = \{X_i, i \in e\}$ (test set) and $X^{(e)} = \{X_i, i \in \{1, \dots, n\} \setminus e\}$ (training set). Let also $P_n^e := 1/p \sum_{i \in e} \delta_{X_i}$ and $P_n^{(e)} := 1/(n-p) \sum_{i \in (e)} \delta_{X_i}$ denote the empirical measures, respectively, defined from the test set X^e and the training set $X^{(e)}$.

Hold-out. *Simple validation* also called *Hold-out* was introduced in the early 1930s [Larson (1931)]. For every $1 \leq p \leq n - 1$, it consists in randomly splitting observations into a training set $X^{(e)}$ of cardinality $n - p$ and a test set X^e of cardinality p . Random data splitting is only made once and introduces additional variability. For every $e \in \mathcal{E}_p$ (randomly chosen), the hold-out estimator of $R_n(\hat{s})$ is

$$(9) \quad \widehat{R}_{\text{Ho},p}(\hat{s}) := P_n^e \gamma(\hat{s}(X^{(e)})) = \frac{1}{p} \sum_{i \in e} \gamma(\hat{s}(X^{(e)}); X_i).$$

Hold-out has been studied, for instance, by Bartlett, Boucheron and Lugosi (2002), Blanchard and Massart (2006) in classification and by Lugosi and Nobel (1999), Wegkamp (2003) in regression.

Leave-p-out. Unlike equation (9) where a single split e of the data is randomly chosen, which introduces additional unwanted variability, *leave-p-out* (Lpo) considers all the $\binom{n}{p} = \text{Card}(\mathcal{E}_p)$ splits. The Lpo estimator of $R_n(\hat{s})$ is defined by

$$(10) \quad \widehat{R}_p(\hat{s}) = \left(\binom{n}{p} \right)^{-1} \sum_{e \in \mathcal{E}_p} P_n^e \gamma(\hat{s}(X^{(e)})).$$

For instance, it has been studied by Shao (1993), Zhang (1993), and Arlot and Celisse (2011) in the regression framework. With $p = 1$, Lpo reduces to the celebrated *leave-one-out* (Loo) cross-validation introduced by Mosteller and Tukey (1968) and further studied by Stone (1974). Note that computing the Lpo estimator requires a computational complexity of order $\binom{n}{p}$ times that of computing \hat{s} , which becomes intractable as n grows.

V-fold cross-validation. To overcome the high computational burden of Lpo [equation (10)], Geisser (1974, 1975) introduced the *V-fold cross-validation* (V-FCV). Instead of considering all the $\binom{n}{p}$ possible splits, one (randomly or not) chooses a partition of X_1, \dots, X_n into V subsets X^{e_1}, \dots, X^{e_V} of approximately equal size $p = n/V = \text{Card}(e_i)$, $i = 1, \dots, V$. Every X^{e_i} , $i = 1, \dots, V$ is successively used as a test set leading to the *V-fold risk estimator* of $R_n(\hat{s})$

$$(11) \quad \widehat{R}_{V\text{-FCV}}(\hat{s}) = \frac{1}{V} \sum_{v=1}^V P_n^{e_v} \gamma(\hat{s}(X^{(e_v)})).$$

V-FCV has been studied in the regression framework by Burman (1989, 1990) who suggests a correction to remove its bias.

2.2.2. Lpo versus V-FCV. As explained in Section 2.2.1, the Lpo computational complexity is roughly $\binom{n}{p}$ times that of computing \hat{s} , which can be highly

time-consuming. Several surrogates of Lpo have been proposed such as V -FCV and the *repeated learning-testing cross-validation* [Breiman et al. (1984), Burman (1989), Zhang (1993)]. Unlike Lpo (and even Loo when $p = 1$), V -FCV involves only V such computations, which is less demanding as long as $V \ll n$. Note that usual values for V are 3, 5, and 10 (except $V = n$ where V -FCV and Loo coincide).

However, V -FCV relies on a preliminary (possibly random) partitioning of X_1, \dots, X_n into V subsets. This preliminary partitioning induces some additional variability which could be misleading. For instance, Celisse and Robin (2008) have theoretically quantified the amount of additional variability induced by V -FCV with respect to Lpo.

2.3. Closed-form expressions for the Lpo risk estimator. Closed-form formulas for the Lpo estimator are proved in the present section, which makes Lpo fully effective in practice and better than V -FCV. Such formulas also enable a more accurate theoretical analysis of CV procedures both in terms of risk estimation (Section 2.4) and model selection (Section 3).

With the notation introduced at the beginning of Section 2.2.1, let us consider projection estimators \hat{s}_m defined by equation (5). Closed-form formulas for the Lpo risk estimator are derived exploiting the “linearity” of projection estimators. Sums over \mathcal{E}_p (which cannot be computed in general) then reduce to binomial coefficients. In the present section, proofs have been deferred to Appendix A [Supplementary material Celisse (2014)]. Recalling the expression of the contrast $\gamma(\cdot; \cdot)$ [equation (2)], one has to compute both quadratic and linear terms.

LEMMA 2.1. *For every $m \in \mathcal{M}_n$, let $\hat{s}_m = \hat{s}_m(X_{1,n})$ denote a projection estimator defined by equation (5) and set $X^e = \{X_i, i \in e\}$ for every $e \in \mathcal{E}_p$. Then for every $p \in \{1, \dots, n - 1\}$,*

$$\begin{aligned} \sum_{e \in \mathcal{E}_p} \|\hat{s}_m(X^{(e)})\|^2 &= \frac{1}{(n-p)^2} \sum_{\lambda \in \Lambda(m)} \left[\binom{n-1}{p} \sum_{k=1}^n \varphi_\lambda^2(X_k) \right. \\ &\quad \left. + \binom{n-2}{p} \sum_{k \neq \ell} \varphi_\lambda(X_k) \varphi_\lambda(X_\ell) \right], \\ \sum_{e \in \mathcal{E}_p} \sum_{i \in e} \hat{s}(X^{(e)})(X_i) &= \frac{1}{n-p} \sum_{\lambda \in \Lambda(m)} \binom{n-2}{p-1} \sum_{i \neq j} \varphi_\lambda(X_i) \varphi_\lambda(X_j). \end{aligned}$$

Lemma 2.1 enables to derive closed-form formulas for the Lpo risk estimator, which makes Lpo procedure fully efficient in practice.

PROPOSITION 2.1. For every $m \in \mathcal{M}_n$, let $\widehat{s}_m = \widehat{s}_m(X_{1,n})$ denote a projection estimator defined by equation (5). Then for every $p \in \{1, \dots, n - 1\}$,

$$(12) \quad \begin{aligned} \widehat{R}_p(m) &= \widehat{R}_p(\widehat{s}_m) \\ &= \frac{1}{n(n-p)} \sum_{\lambda \in \Lambda(m)} \left[\sum_{j=1}^n \varphi_\lambda^2(X_j) - \frac{n-p+1}{n-1} \sum_{j \neq k} \varphi_\lambda(X_j) \varphi_\lambda(X_k) \right]. \end{aligned}$$

Proposition 2.1 enjoys a great interest. First it applies to the broad family of projection estimators. Second, it allows one to reduce the computation time from an exponential to a linear complexity since computing (12) is of order $\mathcal{O}(n)$. Note that in the more specific setting of histograms and kernel estimators, such closed-form formulas have been derived by Celisse and Robin (2008).

Let us now specify the Lpo estimator expressions for the three examples of projection estimators given in Section 2.1.2.

COROLLARY 2.1 (Histograms). For \widehat{s}_m given by equation (6) and for $p \in \{1, \dots, n - 1\}$,

$$\widehat{R}_p(m) = \frac{1}{(n-1)(n-p)} \sum_{\lambda=1}^{D_m} \frac{1}{|I_\lambda|} \left[(2n-p) \frac{n_\lambda}{n} - n(n-p+1) \left(\frac{n_\lambda}{n} \right)^2 \right],$$

where $n_\lambda = \text{Card}(\{i | X_i \in I_\lambda\})$.

COROLLARY 2.2 (Trigonometric polynomials). For every $k \in \mathbb{N}$, let φ_λ denote either $t \mapsto \cos(2\pi kt)$, if $\lambda = 2k$ or $t \mapsto \sin(2\pi kt)$, if $\lambda = 2k + 1$. Let us further assume $\Lambda(m) = \{0, \dots, 2K\}$ for $K \in \mathbb{N}^*$. Then for every $p \in \{1, \dots, n - 1\}$,

$$\widehat{R}_p(m) = \alpha(n, p) - \beta(n, p) \sum_{k=0}^K \left[\left\{ \sum_{j=1}^n \cos(2\pi k X_j) \right\}^2 + \left\{ \sum_{j=1}^n \sin(2\pi k X_j) \right\}^2 \right],$$

where $\alpha(n, p) = (p - 2)(K + 1)[(n - 1)(n - p)]^{-1}$ and $\beta(n, p) = (n - p + 1)[n(n - 1)(n - p)]^{-1}$.

COROLLARY 2.3 (Haar basis). Let us define $\varphi(\cdot) = \mathbb{1}_{[0,1]}(\cdot)$ and $\varphi_{j,k}(\cdot) = 2^{j/2} \varphi(2^j \cdot - k)$, where $j \in \mathbb{N}$ and $0 \leq k \leq 2^j - 1$, and assume $\Lambda(m) \subset \{(j, k) | j \in \mathbb{N}, 0 \leq k \leq 2^j - 1\}$ for every $m \in \mathcal{M}_n$. Then,

$$\widehat{R}_p(m) = \frac{1}{(n-1)(n-p)} \sum_{(j,k) \in \Lambda(m)} 2^j \left[(2n-p) \frac{n_{j,k}}{n} - n(n-p+1) \left(\frac{n_{j,k}}{n} \right)^2 \right],$$

where $n_{j,k} = \text{Card}(\{i | X_i \in [k/2^j, (k+1)/2^j]\})$.

2.4. *Risk estimation: Leave-one-out optimality.* From the general formula (12), one derives closed-form expressions for the expectation and variance of the Lpo risk. These expressions allow to analyze the theoretical behavior of CV in terms of risk estimation and model selection (see Section 3). In the present section, we prove the optimality of Loo for estimating the risk of any projection estimator (Theorem 2.1).

PROPOSITION 2.2. *For every $m \in \mathcal{M}_n$, let $\widehat{s}_m = \widehat{s}_m(X_{1,n})$ denote a projection estimator defined by equation (5). Then for every $1 \leq p \leq n - 1$,*

$$\mathbb{E}[\widehat{R}_p(m)] = \frac{1}{n-p} \sum_{\lambda \in \Lambda(m)} [\mathbb{E}\varphi_\lambda^2(X) - (\mathbb{E}\varphi_\lambda(X))^2] - \sum_{\lambda \in \Lambda(m)} (\mathbb{E}\varphi_\lambda(X))^2$$

and

$$(13) \quad \text{Var}[\widehat{R}_p(m)] = \frac{1}{(n-1)^2} \left[a_n + \frac{b_n}{(n-p)} + \frac{c_n}{(n-p)^2} \right],$$

where $a_n = \text{Var}[\sum_{\lambda \in \Lambda(m)} (n(P_n\varphi_\lambda)^2 - P_n\varphi_\lambda^2)]$, $c_n = \text{Var}[n \sum_{\lambda \in \Lambda(m)} (P_n\varphi_\lambda^2 - (P_n\varphi_\lambda)^2)]$, and $b_n = -2 \text{Cov}[\sum_{\lambda \in \Lambda(m)} (n(P_n\varphi_\lambda)^2 - P_n\varphi_\lambda^2), \sum_{\lambda \in \Lambda(m)} n(P_n\varphi_\lambda^2 - (P_n\varphi_\lambda)^2)]$.

The proof is a straightforward application of Proposition 2.1 and has been omitted. Note that the above quantities do exist as long as $P|\varphi_\lambda|^3 < +\infty$ for any $\lambda \in \Lambda(m)$, which holds true if s is bounded for instance and $\int |\varphi_\lambda|^3 < +\infty$ (φ_λ continuous and compact supported, e.g.). In (13), a_n , b_n and c_n do not depend on p . Then knowing the behavior of the variance with respect to p only depends on the magnitude of a_n , b_n and c_n , which is clarified by Corollary 2.5.

Let us first focus on the bias $\mathbb{B}[\widehat{R}_p(m)] := \mathbb{E}\widehat{R}_p(m) - \mathbb{E}[\|\widehat{s}_m\|^2 - 2 \int_{[0,1]} s \widehat{s}_m]$ of the Lpo estimator.

COROLLARY 2.4 (Bias). *For every $m \in \mathcal{M}_n$, let $\widehat{s}_m = \widehat{s}_m(X_{1,n})$ denote a projection estimator defined by equation (5). Then for every $m \in \mathcal{M}_n$ and $1 \leq p \leq n - 1$,*

$$\mathbb{B}[\widehat{R}_p(m)] = \frac{p}{n(n-p)} \sum_{\lambda \in \Lambda(m)} \text{Var}[\varphi_\lambda(X_1)] \geq 0.$$

The bias is nonnegative and increases with p , which means Loo ($p = 1$) has the smallest bias among CV procedures. If $p = p_n$ satisfies $p_n/n \xrightarrow{n \rightarrow +\infty} q \in [0, 1)$, then $\mathbb{B}[\widehat{R}_{p_n}(m)] \xrightarrow{n \rightarrow +\infty} 0$, and Loo is asymptotically unbiased.

Let us now describe the behavior of the variance with respect to p .

COROLLARY 2.5 (Variance). *With the same notation as Proposition 2.2, for every $m \in \mathcal{M}_n$ and $1 \leq p \leq n - 1$,*

$$\text{Var}[\widehat{R}_p(m)] = \frac{n}{(n - 1)^2} \left[A + \frac{B}{n - p} + \frac{C}{(n - p)^2} + O\left(\frac{1}{n}\right) \right],$$

where the big $O(\cdot)$ does not depend on p , but depends on S_m and P , and

$$\begin{aligned} A &= 4 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_2), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right] \geq 0, \\ B &= 8 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_2), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right] \\ &\quad - 4 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}^2(X_1), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right], \\ C &= 4 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_2), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right] \\ &\quad - 4 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}^2(X_1), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right] + \text{Var} \left[\sum_{\lambda} \varphi_{\lambda}^2(X_1) \right] \\ &\geq 0. \end{aligned}$$

In the more specific case of histogram and kernel density estimators, Celisse and Robin (2008) derived a similar (nonasymptotic) result for the variance.

The monotonicity of the variance with respect to p depends on the sign of B since $x \mapsto f(x) = Ax^2 + Bx + C$ has for derivative $x \mapsto f'(x) = 2Ax + B$ and $A \geq 0$. However, in full generality, the sign of B is unknown. The following proposition relates the monotonicity of $p \mapsto \text{Var}[\widehat{R}_p(m)]$ to this sign.

PROPOSITION 2.3. *Let us define $p_{0,n} = \text{Argmin}_{1 \leq p \leq n-1} \text{Var}[\widehat{R}_p(m)]$ in equation (13). Then,*

$$p_{0,n} = n + \left(1 - \frac{\text{Cov}[\sum_{\lambda} \varphi_{\lambda}^2(X_1), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3)]}{2 \text{Cov}[\sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_2), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3)]} \right) (1 + o(1)),$$

where the little $o(\cdot)$ only depends on S_m and P . Furthermore, if

$$\begin{aligned} &2 \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_2), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right] \\ (14) \quad &\geq \text{Cov} \left[\sum_{\lambda} \varphi_{\lambda}^2(X_1), \sum_{\lambda} \varphi_{\lambda}(X_1) \varphi_{\lambda}(X_3) \right], \end{aligned}$$

$p \in \{1, \dots, n - 1\} \mapsto \text{Var}[\widehat{R}_p(m)]$ is increasing. Otherwise, $p \mapsto \text{Var}[\widehat{R}_p(m)]$ is decreasing on $[1, p_{0,n}]$ and increasing on $[p_{0,n}, n - 1]$.

Equation (14) is related to the sign of B (Corollary 2.5) and to the minimum location value $p_{0,n}$. If it holds true, then $p_{0,n} \notin \{2, \dots, n-1\}$, which means Loo has the smallest variance among CV procedures. In particular, let us notice (14) holds true with any density estimated by regular histograms since $\sum_{\lambda} \varphi_{\lambda}^2(X_1)$ is then a constant and the covariance in the left-hand side is a variance by independence of X_1, X_2 , and X_3 . On the contrary, (14) is not fulfilled when using histograms based on a partition such that $\mathbb{P}[X \in I_{\lambda}] = C < 1/2$ for every λ , where C denotes a constant.

We are now in position to provide the main result of this section, which describes the behavior of \widehat{R}_p as a risk estimator in terms of mean-square error (MSE).

THEOREM 2.1. *For every $m \in \mathcal{M}_n$, let us define the MSE of \widehat{s}_m by $\text{MSE}(m; p) = (\mathbb{E}[\widehat{R}_p(m)])^2 + \text{Var}[\widehat{R}_p(m)]$, for every $p \in \{1, \dots, n-1\}$.*

1. *If (14) holds true, then for every $m \in \mathcal{M}_n$, $p \mapsto \text{MSE}(m; p)$ is minimum for $p = 1$.*
2. *Otherwise, for every $p = p_n \in \{1, \dots, n-1\}$ such that $\limsup_{n \rightarrow +\infty} p_n/n < 1$, then*

$$\text{MSE}(m; p) = \frac{A}{n} + O\left(\frac{1}{n^2}\right) \quad \text{as } n \rightarrow +\infty,$$

where A is given in Corollary 2.5 and the big $O(\cdot)$ depends on S_m and P .

If (14) holds true, Loo is the best CV procedure in terms of MSE when estimating the risk of an estimator. Otherwise as long as $\limsup_{n \rightarrow +\infty} p_n/n < 1$, choosing a value of $p \neq 1$ is useless since any value in $\{1, \dots, n-1\}$ asymptotically leads to the same performance in terms of MSE. Therefore, since Loo has the smallest bias (Corollary 2.4), *Loo is optimal among CV procedures for estimating the risk of an estimator*. This result confirms what has been previously stated by Burman (1989) in the regression framework.

3. Optimal cross-validation for model selection. From Section 2.4, Loo is proved to be the best CV procedures in the context of risk estimation. However, the best procedure for risk estimation is not necessarily the best one for *model selection*. Although the empirical risk $P_n \gamma(\widehat{s}_m)$ (4) is a reliable estimator of $\mathbb{E}[P_n \gamma(\widehat{s}_m)]$, using empirical risk minimization to choose one $\widehat{m} \in \mathcal{M}_n$ (without penalizing) would systematically lead to overfitting. The purpose of the present section is to study the performance of CV for model selection with respect to the cardinality p of the test set.

3.1. *Optimal cross-validation for estimation.* The performance of CV with respect to p is first characterized through a *sharp oracle inequality* (Theorem 3.1). A leading constant converging to 1 as $n \rightarrow +\infty$ is achieved for some values of p ,

highlighting the asymptotic optimality of corresponding CV procedures. From a theoretical point of view, Corollary 3.1 explores the link between (a proxy to) the optimal p and influential quantities related to the difficulty of the estimation problem for finite sample size. These results are further validated by simulation experiments (Section 3.1.4).

3.1.1. *Estimation point of view.* With the notation of Section 2.1, let us consider a family of projection estimators $\{\widehat{s}_m\}_{m \in \mathcal{M}_n}$, where \mathcal{M}_n denotes an (at most countable) index set allowed to depend on n . The best possible model, called *the oracle model*, is denoted by S_{m^*} , where

$$\begin{aligned} m^* &:= \underset{m \in \mathcal{M}_n}{\operatorname{Argmin}} P\gamma(\widehat{s}_m) - P\gamma(s) = \underset{m \in \mathcal{M}_n}{\operatorname{Argmin}} \|s - \widehat{s}_m\|^2 \\ &= \underset{m \in \mathcal{M}_n}{\operatorname{Argmin}} P\gamma(\widehat{s}_m). \end{aligned}$$

Since $P\gamma(\widehat{s}_m)$ has to be estimated, one uses CV (Lpo) to choose a candidate model for every $1 \leq p \leq n - 1$,

$$(15) \quad \widehat{m}(p) := \underset{m \in \mathcal{M}_n}{\operatorname{Argmin}} \widehat{R}_p(m),$$

and the final candidate model is denoted by $S_{\widehat{m}(p)}$. The purpose is now to infer the properties of $\widehat{s}_{\widehat{m}(p)}$ with respect to $p \in \{1, \dots, n - 1\}$ in terms of an oracle inequality such as (1).

3.1.2. *Main oracle inequality.* Let us introduce some notation and detail the main assumptions used along the following sections.

Square-integrable density.

$$(SqI) \quad s \in L^2([0, 1]).$$

Unlike [Castellan \(2003\)](#), for instance, it is not assumed that $s \geq \rho$ for a constant $\rho > 0$.

Polynomial collection. There exists $a_{\mathcal{M}} \geq 0$ such that

$$(Pol) \quad \operatorname{Card}(\mathcal{M}_n) \leq n^{a_{\mathcal{M}}}.$$

In particular, this holds true if there exists $\alpha \geq 0$ such that $\operatorname{Card}(\{m \in \mathcal{M}_n, D_m = D\}) \leq D^\alpha$, for every $1 \leq D \leq n$.

Model regularity.

$$(RegD) \quad \exists \Phi > 0, \quad \sup_{m \in \mathcal{M}_n} \frac{\|\phi_m\|_\infty}{D_m} \leq \Phi \quad \text{with } \phi_m = \sum_{\lambda \in \Lambda(m)} \varphi_\lambda^2.$$

It relates the regularity of the orthonormal basis (measured in terms of sup-norm) to the dimension of the model. For instance, using (6), (RegD) requires $|I_\lambda| \geq (\Phi D_m)^{-1}$ for every $\lambda \in \Lambda(m)$. The length of intervals I_λ cannot be too different from one another to some extent.

Maximal dimension.

$$(D_{\max}) \quad \exists \Gamma > 0, \quad \sup_{m \in \mathcal{M}_n} D_m \leq \Gamma \frac{n}{(\log n)^2}.$$

In the sequel, $\Gamma = 1$ is always considered to simplify expressions. Note that proofs and conclusions remain unchanged with this particular choice.

Estimation error and dimension.

$$(LoEx) \quad \exists \xi > 0, \quad \inf_{m \in \mathcal{M}_n} \frac{\sqrt{n} \mathbb{E}(\|s_m - \widehat{s}_m\|)}{\sqrt{D_m}} \geq \sqrt{\xi}.$$

This assumption makes the estimation error $\mathbb{E}(\|s_m - \widehat{s}_m\|^2)$ and D_m comparable. For instance, Lemma B.3 [supplementary material [Celisse \(2014\)](#)] proves (LoEx) is fulfilled with Hölder densities estimated by regular histograms defined by (6) such that $|\{x \in [0, 1] \mid s(x) \geq \eta\}| \geq \ell$ for some $\eta \in (0, 1)$, where $0 < \ell < 1$ satisfies

$$\ell > \left(\inf_{m \in \mathcal{M}_n} D_m \right)^{-1}.$$

Note the latter inequality amounts to exclude too small models for which the support of s is included in one single interval I_λ .

Richness of the collection. There exist $m_0 \in \mathcal{M}_n$ and $c_{\text{rich}} \geq 1$ such that,

$$(Rich) \quad \sqrt{n} \leq D_{m_0} \leq c_{\text{rich}} \sqrt{n}.$$

This requirement is rather mild since one can add such a model in our collection.

Approximation property. There exist $c_\ell, c_u > 0$ and $\ell > u > 0$ such that, for every $m \in \mathcal{M}_n$,

$$(Bias) \quad c_\ell D_m^{-\ell} \leq \|s - s_m\|^2 \leq c_u D_m^{-u}.$$

This assumption quantifies the bias (approximation error) incurred by model S_m in estimating s . It therefore relies on a smoothness assumption on s . Such an upper bound is classical for α -Hölderian functions with $\alpha \in (0, 1]$ and regular histograms (6), for instance. Note that [Stone \(1985\)](#) uses the same assumption (lower bound), which is the *finite sample counterpart* of the classical assumption $\|s - s_m\| > 0$ for every $m \in \mathcal{M}_n$ usually made to prove *asymptotic optimality* for a model selection procedure [see [Birgé and Massart \(2007\)](#)].

Rate of convergence for the oracle model.

$$(OrSp) \quad n R_n^* (\log n)^{-2} \xrightarrow{n \rightarrow \infty} +\infty \quad \text{with } R_n^* = \inf_{m \in \mathcal{M}_n} R_n(\widehat{s}_m).$$

The risk of the oracle model R_n^* does not decrease to 0 faster than $(\log n)^2/n$. In particular, this holds true for densities in $\mathcal{H}(L, \alpha)$ with $L > 0$ and $\alpha \in (0, 1]$ estimated by regular histograms [see Section B.5 in [Celisse \(2014\)](#)].

The performance of the Lpo estimator with respect to p is described by the following oracle inequality from which the CV optimality is deduced for some values of p . The proof is given in Appendix A.1.

THEOREM 3.1 (Optimal CV). *Let s denote a density on $[0, 1]$ such that (Sql) holds true, set $\{S_m\}_{m \in \mathcal{M}_n}$ a collection of models defined in Section 2.1.2, and assume (Pol), (RegD), (Dmax), (Rich), (LoEx), (Bias) and (OrSp). Let $\hat{m} = \hat{m}(p)$ denote the model minimizing $\hat{R}_p(m)$ over \mathcal{M}_n for every $p \in \{1, \dots, n - 1\}$. Then there exist a sequence $(\delta_n)_{\mathbb{N}}$ such that $\delta_n \rightarrow 0$, and $n\delta_n \rightarrow +\infty$ as $n \rightarrow +\infty$, and an event $\tilde{\Omega}$ with $\mathbb{P}(\tilde{\Omega}) \geq 1 - 6/n^2$ on which, for large enough values of n ,*

$$\|s - \hat{s}_{\hat{m}(p)}\|^2 \leq C_n(p) \inf_{m \in \mathcal{M}_n} \{ \|s - \hat{s}_m\|^2 \} \quad \text{with } C_n(p) = \frac{T_B^+ \vee T_V^+}{T_B^- \wedge T_V^-} \geq 1,$$

where

$$T_B^- = 1 - \delta_n K(n, p),$$

$$T_V^- = \frac{1}{1 - p/n} (1 - \delta_n)[1 - 4\delta_n] - 2\delta_n K(n, p)[3 - 4\delta_n],$$

$$T_B^+ = 1 + \delta_n K(n, p),$$

$$T_V^+ = \frac{1}{1 - p/n} (1 + \delta_n)[1 + 4\delta_n] + 2\delta_n K(n, p)[3 + 4\delta_n],$$

and $K(n, p) = 1 + \frac{2}{n-1} + \frac{p}{n-p} \frac{1}{n-1}$.

If $p/n \rightarrow 0$ then $C_n(p) \rightarrow 1$ as $n \rightarrow +\infty$, which leads to *efficient* (asymptotically optimal) model selection procedures. In particular, this holds true for $p = 1$ that is, *Loo* is asymptotically optimal since

$$\frac{\|s - \hat{s}_{\hat{m}(1)}\|^2}{\inf_{m \in \mathcal{M}_n} \{ \|s - \hat{s}_m\|^2 \}} \xrightarrow[n \rightarrow +\infty]{\text{a.s.}} 1.$$

From the proof, it also arises that the slowly decreasing sequence δ_n is related to the model collection structure. An increase of \mathcal{M}_n makes the model selection problem more difficult and δ_n larger.

While *asymptotic optimality* is deduced from Theorem 3.1 for any CV procedure as long as $p = o(n)$, it is also desirable to analyze the performance of CV as p depends on the finite sample size. From a theoretical point of view, this will provide the rate at which p/n has to decrease to 0 to reach efficiency. Based on Figure 1 [panel (c)] where $C_n(p)$ appears as a reliable proxy to the optimal $C_{\text{or},n}(p)$ [given by equation (19)], we suggest to optimize $C_n(p)$ with respect to p to get a surrogate optimal p depending on influential parameters such as n and δ_n . This strategy has been validated by simulation experiments of Section 3.1.4. The following Corollary 3.1 proves the best (surrogate) p/n slowly decreases to 0.

COROLLARY 3.1 (Optimizing upper bound). *With the notation and assumptions of Theorem 3.1, the constant $C_n(p)$ is minimized over $p \in \{1, \dots, n - 1\}$*

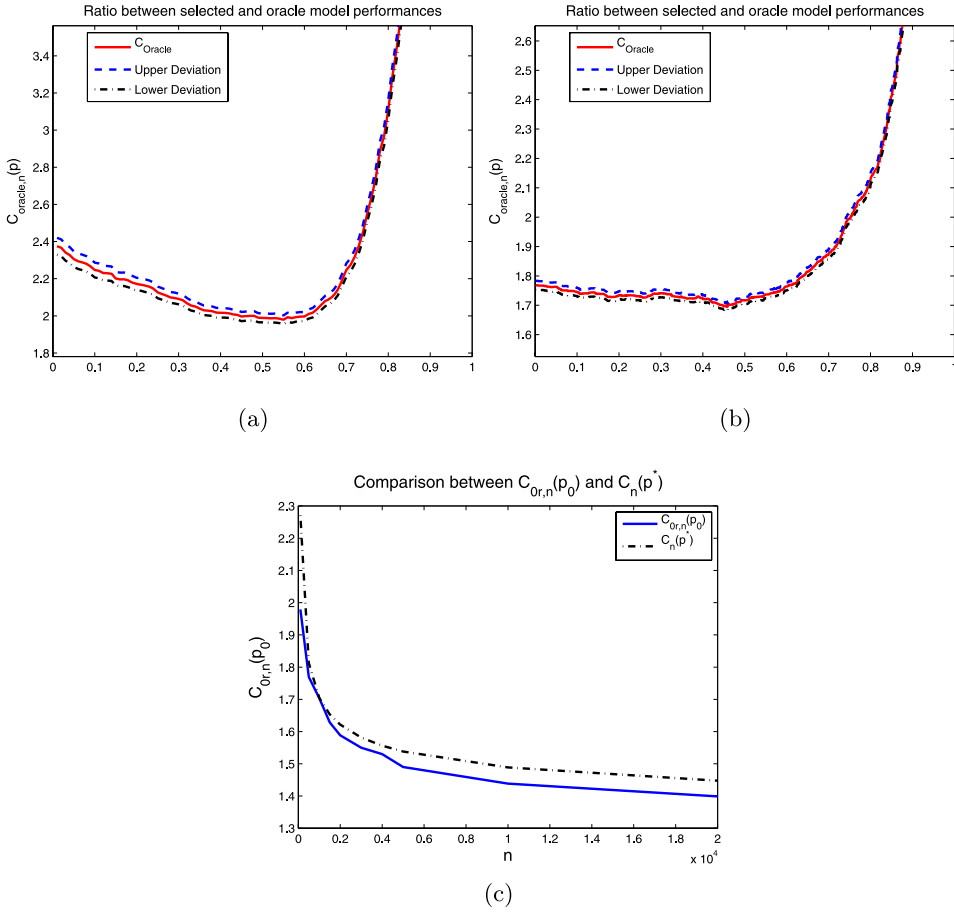


FIG. 1. Panels (a) and (b): $p/n \mapsto C_{\text{or},n}(p)$ (plain red line) is plotted for $\Gamma = 1$ [see (Dmax)] and different values of n : (a) $n = 100$, and (b) $n = 1000$. $p/n \mapsto C_{\text{oracle},n}^+(p)$ (blue dashed line) and $p/n \mapsto C_{\text{oracle},n}^-(p)$ (black dot-dashed line) have been plotted on the same graph as well [see (20)]. Panel (c): $n \mapsto C_{\text{or},n}(p_0)$ (plain blue line) and $n \mapsto C_n(p^*)$ (black dot-dashed line) are displayed. $N = 1000$ samples have been drawn from the mixture of Beta distributions (18).

for

$$0 < \frac{p_n^*}{n} = 1 - \frac{1 - 5\delta_n + 4\delta_n^2 - (2/(n-1))(3\delta_n - 4\delta_n^2) + \delta_n/(n-1)}{1 + 2(1 + 1/(n-1))(3\delta_n - 4\delta_n^2) - \delta_n(1 + 1/(n-1))} < 1.$$

Furthermore, the optimal ratio p^*/n is slowly decreasing to 0 as n tends to $+\infty$

$$(16) \quad p_n^* \underset{n \rightarrow +\infty}{\sim} 10n\delta_n \longrightarrow +\infty.$$

The proof has ben deferred to Appendix A.1. Corollary 3.1 describes the rate (up to constant) at which $p = p_n^*$ has to grow with n to achieve finite-sample

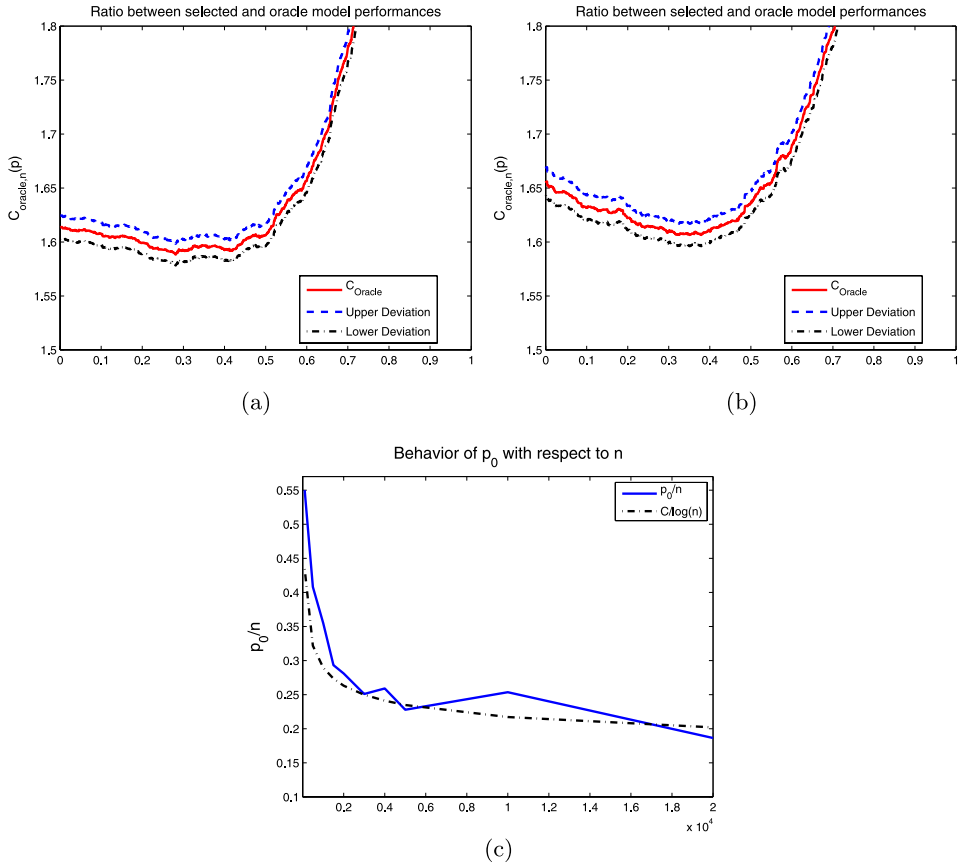


FIG. 2. For (a) and (b), $p/n \mapsto C_{\text{Or},n}(p)$ (plain red line) is plotted for $n = 2000$ and different values of Γ [see (Dmax)]: (a) $\Gamma = 1$, (b) $\Gamma = 2$. $p/n \mapsto C_{\text{oracle},n}^+(p)$ (blue dashed line) and $p/n \mapsto C_{\text{oracle},n}^-(p)$ (black dot-dashed line) have been plotted on the same graph as well [see (20)]. $N = 1000$ samples have been drawn from the mixture of Beta distributions (18). For (c), $n \mapsto p_0/n$ (blue plain line) and $n \mapsto C/(\log n)$ (black dot-dashed line) are displayed, where p_0 denotes the minimizer of $C_{\text{Or},n}(p)$ as a function of p and C is a constant.

optimality. In particular p_n^*/n in (16) is related to δ_n which is strongly connected to the structure of the model collection as explained following Theorem 3.1. A more complex collection leads to a larger δ_n and then to a larger optimal p_n^* . In other words, p must be chosen large enough to balance the overfitting induced by the structure of the model collection. This phenomenon is observed in practice in the simulation experiments of Section 3.1.4 (Figure 2).

3.1.3. *Adaptivity in the minimax sense.* *Adaptivity in the minimax sense* is a desirable property for model selection procedures. It means the considered proce-

ture automatically adapts to the unknown smoothness of the target function s to estimate [see Barron, Birgé and Massart (1999) for an extensive presentation].

Several adaptivity in the minimax sense results are provided in the present section. Deriving such results from oracle inequalities (1) is somewhat classical. Here, the novelty is first that CV enjoys such a desirable property as a model selection procedure, second that the leading constant $C_n(p)$ in Theorem 3.1 when converging to 1 as n tends to $+\infty$ provides accurate results.

Let us start providing a general theorem from which any adaptivity result will be immediate corollary. The proof is given in Appendix A.1.

THEOREM 3.2. *Let s denote a density on $[0, 1]$ such that (SqI) holds true, set $\{S_m\}_{m \in \mathcal{M}_n}$ a collection of models defined in Section 2.1.2, and assume (Pol), (RegD), (Dmax), (Rich), (LoEx), (Bias), and (OrSp). Let $\hat{m} = \hat{m}(p)$ denote the model minimizing $\hat{R}_p(m)$ over \mathcal{M}_n for every $p \in \{1, \dots, n-1\}$. Then for every $1 \leq p \leq n-1$,*

$$(17) \quad \mathbb{E}[\|s - \hat{s}_{\hat{m}(p)}\|^2] \leq C_n(p) \mathbb{E}\left[\inf_{m \in \mathcal{M}_n} \|s - \hat{s}_m\|^2\right] + (\Phi + \|s\|^2) \frac{12}{n(\log n)^2} + \frac{6c_u}{n^2},$$

where $C_n(p) = \frac{T_B^+ \vee T_V^+}{T_B^- \wedge T_V^-}$, with

$$T_B^- = 1 - \delta_n K(n, p),$$

$$T_V^- = \frac{1}{1 - p/n} (1 - \delta_n)[1 - 4\delta_n] - 2\delta_n K(n, p)[3 - 4\delta_n],$$

$$T_B^+ = 1 + \delta_n K(n, p),$$

$$T_V^+ = \frac{1}{1 - p/n} (1 + \delta_n)[1 + 4\delta_n] + 2\delta_n K(n, p)[3 + 4\delta_n],$$

and $K(n, p) = 1 + \frac{2}{n-1} + \frac{p}{n-p} \frac{1}{n-1}$.

The last two terms in the right-hand side of (17) are remainder terms by Assumptions (RegD), (Dmax), and (Bias).

Applying Theorem 3.2 to the collection of regular histograms defined by (6), the following corollary settles an adaptivity property with respect to Hölder balls [see DeVore and Lorentz (1993)].

COROLLARY 3.2. *Let us consider the model collection of Section 2.1.2 made of piecewise constant functions and the associated histograms defined by (6) such that, for every $m \in \mathcal{M}_n$ and $\lambda \in \Lambda(m)$, $|I_\lambda| = D_m^{-1}$ (regular histograms). Let us also assume (Dmax) and (LoEx) hold true.*

If the target density s belongs to the Hölder ball $\mathcal{H}(L, \alpha)$ for some $L > 0$ and $\alpha \in (0, 1]$, then there exist constants $0 < K_\alpha^- \leq K_\alpha^+$ such that for every $p = o(n)$,

$$\begin{aligned} & K_\alpha^- L^{2/(2\alpha+1)} n^{-2\alpha/(2\alpha+1)} \\ & \leq \sup_{s \in \mathcal{H}(L, \alpha)} \mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2] \\ & \leq (1 + o(1)) K_\alpha^+ L^{2/(2\alpha+1)} n^{-2\alpha/(2\alpha+1)} + O\left(\frac{1}{n(\log n)^2}\right), \end{aligned}$$

K_α^- and K_α^+ only depend on α (not on n or s).

Furthermore, since this property holds for every $L > 0$ and $\alpha \in (0, 1]$, then $\{\widehat{s}_{\widehat{m}(p)}\}_{n \in \mathbb{N}^*}$ is adaptive in the minimax sense with respect to $\{\mathcal{H}(L, \alpha)\}_{L > 0, \alpha \in (0, 1]}$ for every $p = o(n)$.

The proof has been deferred to Section B.5 in Celisse (2014). The upper bound is tight since the rate $n^{-2\alpha/(2\alpha+1)}$ and the dependence on the radius $L^{2/(2\alpha+1)}$ are the same as in the lower bound, which has been stated by Ibragimov and Has’minskiĭ (1981). The main contribution of this result is to prove $p = o(n)$ leads to adaptivity. Note that similar results can also be proved for Besov balls $\mathcal{B}_{\infty, 2}^\alpha(L)$, with $\alpha, L > 0$, for instance [see DeVore and Lorentz (1993)], by using an appropriate collection of models such as trigonometric polynomials defined by (7).

3.1.4. *Simulation experiments.* Results of simulation experiments are provided to check the conclusions drawn (from theory) in Section 3.1.2. A mixture of Beta distributions

$$(18) \quad \forall x \in [0, 1] \quad s(x) = \frac{\beta(3, 7; x) + \beta(10, 5; x)}{2}$$

has been used to generate samples of size $n = 100, 500, 1000, 2000, 3000, 4000, 5000, 10,000, 20,000$. Note that (18) defines a Hölder density on $[0, 1]$. For each n , every $p \in \{1, \dots, n - 1\}$ have been considered and (Dmax) is fulfilled with $\Gamma = 1$ (Figure 1) and $\Gamma = 2$ (Figure 2).

The model collection we used is made of piecewise constant functions described in Section 2.1.2 leading to regular histogram estimators defined by (6). Only regular histograms with dimension $D_m \geq 2$ are used so that (LoEx) holds true (Lemma B.3). For every $1 \leq p \leq n - 1$, $\widehat{m}(p)$ is defined by (15).

Let us also introduce

$$(19) \quad C_{\text{or}, n}(p) := \mathbb{E}\left[\frac{\|s - \widehat{s}_{\widehat{m}(p)}\|^2}{\inf_{m \in \mathcal{M}_n} \{\|s - \widehat{s}_m\|^2\}}\right] \quad \text{and} \quad p_0 := \underset{1 \leq p \leq n-1}{\text{Argmin}} C_{\text{or}, n}(p),$$

which measures the average performance of $\widehat{s}_{\widehat{m}(p)}$ with respect to that of \widehat{s}_{m^*} (oracle estimator). The closer $C_{\text{or}, n}(p)$ to 1, the better $\widehat{s}_{\widehat{m}(p)}$. Minimizing $C_{\text{or}, n}(p)$ as a function of p for various values of n enables to check whether the conclusions

drawn from minimizing $C_n(p)$ with respect to p (Theorem 3.1 and Corollary 3.1) hold true or not, that is whether $C_n(p)$ is an accurate approximation to $C_{\text{or},n}(p)$. For each curve $p \mapsto C_{\text{or},n}(p)$, a confidence band has been displayed. It is delimited by $p \mapsto C_{\text{or},n}^-(p)$ and $p \mapsto C_{\text{or},n}^+(p)$, respectively, defined by

$$(20) \quad C_{\text{or},n}^-(p) = C_{\text{or},n}(p) - \frac{\hat{\sigma}}{\sqrt{N}} \quad \text{and} \quad C_{\text{or},n}^+(p) = C_{\text{or},n}(p) + \frac{\hat{\sigma}}{\sqrt{N}},$$

where $\hat{\sigma}$ denotes the empirical standard deviation.

First, from panels (a) and (b) of Figure 1, $C_{\text{or},n}(p)$ (plain red lines) decreases pointwise as n grows. This is confirmed by panel (c) of Figure 1 at the particular value $p = p_0$ as n grows. This is in accordance with Theorem 3.1 and $C_n(p) \rightarrow 0$ as n increases. Second, the optimization strategy at the basis of Corollary 3.1 is empirically validated by panel (c) of Figure 1 where $C_{\text{or},n}(p_0)$ and its proxy $C_n(p_0)$ remain very close to each other. Furthermore, the optimal rate derived in Corollary 3.1 is supported up to constant by simulation results displayed in panel (c) of Figure 2 where p_0/n is almost equal to the predicted $\delta_n \approx C/(\log n)$ ($C > 0$) from the proof of Theorem 3.1.

The conclusion of Corollary 3.1 about the dependence of the optimal p_0 on the complexity of the model collection (through δ_n) is also illustrated by panels (a) and (b) in Figure 2 where Γ (Dmax), respectively, equals 1 and 2. As Γ grows the model collection becomes more complex, leading to a worse performance and a larger p_0 in panel (b). The need for a larger p_0 is all the more strong as the curve in panel (b) is less flat than in panel (a), indicating the problem becomes more difficult as Γ increases and any misspecification of p_0 leads to a stronger loss in accuracy. One concludes the more complex the model collection, the larger the optimal p .

Note that this conclusion does not apply to Loo ($p = 1$) [see (16)], suggesting Loo may be suboptimal for finite sample size. This is supported by Figure 1 [panels (a) and (b)] and Figure 2 [panels (a) and (b)] where the minimum of each curve is not reached at $p = 1$.

3.2. Optimal cross-validation for identification. With the notation of Section 2.1, $\{\hat{s}_m\}_{m \in \mathcal{M}_n}$ denotes a collection of projection estimators (Section 2.1.2) which is allowed to depend on n . The purpose is now to recover the best model denoted by $S_{\bar{m}}$ and defined by

$$(21) \quad \bar{m} := \underset{m \in \mathcal{M}_n}{\text{Argmin}} \mathbb{E}[\|s - \hat{s}_m\|^2],$$

where \bar{m} is a deterministic quantity unlike m^* from Section 3.1. Since this goal cannot be reached if other models can perform as well as $S_{\bar{m}}$ (even asymptotically), one also requires there exist $\mu > 0$ and $n_0 \in \mathbb{N}^*$ such that for every integer $n > n_0$,

$$(BeMo) \quad (1 + \mu) \mathbb{E}[\|s - \hat{s}_{\bar{m}}\|^2] \leq \inf_{m \in \mathcal{M}_n \setminus \{\bar{m}\}} \mathbb{E}[\|s - \hat{s}_m\|^2].$$

A similar assumption (in probability rather than in expectation) has been made by Yang (2007). Let us further assume the collection $\{S_m\}_{m \in \mathcal{M}_n}$ can be split into:

- *parametric models* indexed by $\mathcal{M}_{n,P}$ for which there exist constants $\pi, \tau > 0$ (independent of n) such that

$$(22) \quad \sup_{m \in \mathcal{M}_{n,P}} \{n\mathbb{E}[\|s_m - \widehat{s}_m\|^2]\} \leq \pi \quad \text{and} \quad \inf_{m \in \mathcal{M}_{n,P}, s \notin S_m} \{\|s - s_m\|^2\} \geq \tau.$$

- *nonparametric models* indexed by $\mathcal{M}_{n,NP}$ such that

$$(23) \quad n(\log n)^{-2} \inf_{m \in \mathcal{M}_{n,NP}} \mathbb{E}[\|s_m - \widehat{s}_m\|^2] \xrightarrow{n \rightarrow +\infty} +\infty.$$

Then

$$(P-NP) \quad \{S_m\}_{m \in \mathcal{M}_n} = \{S_m\}_{m \in \mathcal{M}_{n,P}} \cup \{S_m\}_{m \in \mathcal{M}_{n,NP}}.$$

Parametric models are models with convergence rate of order $1/n$. Since $\mathbb{E}[\|s - \widehat{s}_m\|^2] \approx \|s - s_m\|^2 + C \cdot D_m/n$, allowing D_m to depend on n makes the rate of the corresponding model slower than $1/n$ (nonparametric model). Consistently with this remark, (22) requires the largest dimension over parametric models is bounded by a constant independent of n , and that the bias of parametric models such that $s \notin S_m$ cannot decrease with n toward 0. Otherwise, such a model would be nonparametric. Conversely, (23) only requires that the dimension of nonparametric models must be larger than $(\log n)^2$. In particular, this does not prevent nonparametric models from containing s or having their bias decreasing to 0 as n grows.

3.2.1. *Main results.* Depending on whether s belongs or not to $\bigcup_{m \in \mathcal{M}_n} S_m$, the two following results prove model selection consistency for CV. Their main contribution is to relate the cardinality p of the test set to the rate of convergence of $\widehat{s}_{\widehat{m}}$ and the model collection complexity. Note that, in addition, the model consistency property is settled with a collection of models allowed to grow with n , which contrasts with earlier results [see, e.g., Yang (2007)].

Let us start with the setting where s belongs to $\bigcup_{m \in \mathcal{M}_n} S_m$, which implies the best estimator $\widehat{s}_{\widehat{m}}$ achieves the parametric rate $1/n$.

THEOREM 3.3 (Model consistency with $s \in \bigcup_m S_m$). *Let $\bigcup_{m \in \mathcal{M}_n} S_m$ denote a collection of models satisfying (Pol) and (P-NP), $\bar{m} \in \mathcal{M}_n$ given by (21) be such that (BeMo) holds true, and assume (Sql), (RegD), (Dmax), and (LoEx). For every $1 \leq p \leq n - 1$, let us also define $\widehat{m} = \widehat{m}(p) = \text{Argmin}_{m \in \mathcal{M}_n} \widehat{R}_p(m)$. If the target $s \in \bigcup_{m \in \mathcal{M}_n} S_m$, then every $1 \leq p = p_n \leq n - 1$ such that*

$$(24) \quad \log(n) \left(1 - \frac{p}{n}\right) \xrightarrow{n \rightarrow +\infty} 0 \quad \text{and} \quad n \left(1 - \frac{p}{n}\right) \xrightarrow{n \rightarrow +\infty} +\infty,$$

leads to

$$\mathbb{P}[\widehat{m} = \bar{m}] \xrightarrow{n \rightarrow +\infty} 1.$$

The proof has been deferred to Appendix B.1. When s belongs to $\bigcup_{m \in \mathcal{M}_n} S_m$, the best estimator $\widehat{s}_{\widehat{m}}$ in a polynomial collection can be recovered by CV provided p/n converges to 1 as n tends to $+\infty$. The proof establishes this rate (i) cannot exceed $1/n$ to allow distinguishing between parametric estimators (with convergence rate of order $1/n$), and (ii) has to be faster than $(\log n)^{-1}$ to allow dealing with the polynomial complexity of the model collection. For instance, a finite collection would lead to replace the $(\log n)^{-1}$ rate by a slower one determined by the control level of $\mathbb{P}[\widehat{m} = \bar{m}]$. In the regression setting, [Yang (2007)] already proved requiring $p/n \rightarrow 1$ enables to recover the best parametric estimator among parametric ones (see Corollary 1), while this requirement is no longer necessary when comparing parametric and nonparametric estimators. Our result is consistent with Yang's one, although our setting is somewhat different since we compare the best parametric estimator with both parametric and nonparametric ones in the same time.

Conversely, when s does not belong to $\bigcup_m S_m$, every parametric model is biased according to (22) and $\widehat{s}_{\widehat{m}}$ reaches a nonparametric rate, that is $nR_n(\bar{m}) \rightarrow +\infty$ as n tends to $+\infty$.

THEOREM 3.4 (Model consistency with $s \notin \bigcup_m S_m$). *Let $\bigcup_{m \in \mathcal{M}_n} S_m$ denote a collection of models satisfying (Pol) and (P-NP), $\bar{m} \in \mathcal{M}_n$ given by (21) be such that (BeMo) holds true, and assume (Sql), (RegD), (Dmax) and (LoEx). For every $1 \leq p \leq n-1$, let us also define $\widehat{m} = \widehat{m}(p) = \text{Argmin}_{m \in \mathcal{M}_n} \widehat{R}_p(m)$. Let us assume the target $s \notin \bigcup_{m \in \mathcal{M}_n} S_m$ and $R_n(\bar{m}) \rightarrow 0$ as n tends to $+\infty$.*

1. *If for large enough values of n $D_{\bar{m}} \leq (\log n)^4$, then every $1 \leq p = p_n \leq n-1$ such that*

$$(25) \quad \log(n) \left(1 - \frac{p}{n}\right) \xrightarrow{n \rightarrow +\infty} 0 \quad \text{and} \quad n\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2] = o(n-p)$$

leads to

$$\mathbb{P}[\widehat{m} = \bar{m}] \xrightarrow{n \rightarrow +\infty} 1.$$

(i) *If for large enough values of n $D_{\bar{m}} > (\log n)^4$, then every $1 \leq p = p_n \leq n-1$ such that*

$$(26) \quad \frac{(\log n)^5}{n\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2]} = o\left(\frac{p/n}{1-p/n}\right) \quad \text{and} \\ \frac{p/n}{1-p/n} = o\left(1 \vee \frac{\|s - s_{\bar{m}}\|^2}{\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2]}\right)$$

leads to

$$\mathbb{P}[\widehat{m} = \bar{m}] \xrightarrow{n \rightarrow +\infty} 1.$$

The proof is similar to that of Theorem 3.3 and has been postponed to Section C.1 [supplementary material Celisse (2014)]. The constraints on p strongly depend on the rate of convergence of $\widehat{s}_{\bar{m}}$ (nonparametric here). When $S_{\bar{m}}$ is a small nonparametric model ($D_{\bar{m}} \leq (\log n)^4$), (25) is very similar to (24) in the parametric setting. In particular, $n\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2] \rightarrow +\infty$ as n tends to $+\infty$ implies $n(1 - p/n) \rightarrow +\infty$ as well. For large nonparametric models ($D_{\bar{m}} > (\log n)^4$), the constraints on p are related to the ratio $\|s - s_{\bar{m}}\|^2/\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2]$. For instance, when estimating $s \in \mathcal{H}(L, \alpha)$ by regular histograms, this ratio remains bounded while $n\mathbb{E}[\|s_{\bar{m}} - \widehat{s}_{\bar{m}}\|^2]$ grows polynomially in n . Then p/n has to converge to 0 as n increases, but not too fast. In particular, Loo ($p = 1$) is suboptimal in that setting [see Figure 4 panel (b)]. Note that Theorem 3.4 has the same flavor as Corollary 1 in [Yang (2007)], except the density estimation setting allows to relate p to the features of the best estimator more closely.

3.2.2. *Simulation experiments.* Simulation experiments have been performed in the settings of Theorems 3.3 and 3.4, respectively, when s belongs to (resp., does not belong to) the model collection. We used a polynomial model collection made of regular piecewise constant functions described in Section 2.1.2 for which Assumptions (P-NP) and (BeMo) are fulfilled with $\mu = 5.10^{-1}$. In each setting, $N = 1000$ samples have been drawn. Results are given in Figures 3 and 4 where $\mathbb{P}[\widehat{m} = \bar{m}]$ is displayed with respect to the ratio p/n . Let us also mention that Lemma B.3 in the supplement Celisse (2014) clearly shows (LoEx) holds true for all densities defined in the following as long as $D_m \geq 2$ for every model in the collection.

When s belongs to the model collection (Figure 3), the following densities have been used:

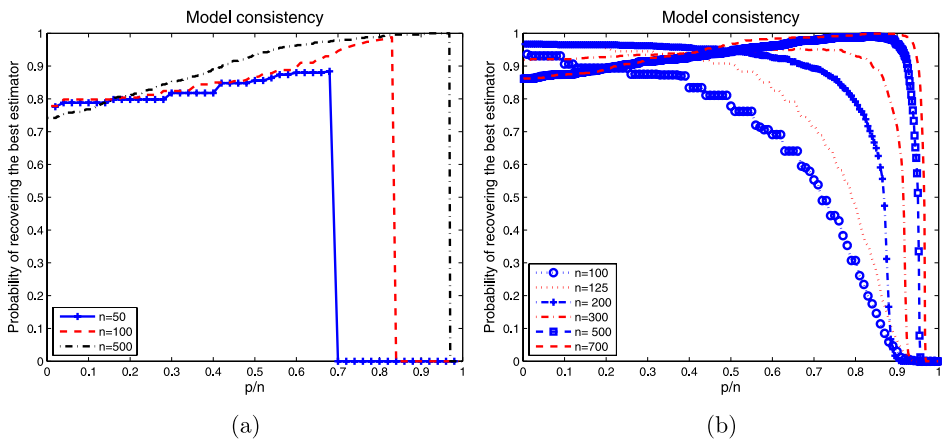


FIG. 3. $p/n \mapsto \mathbb{P}[\widehat{m} = \bar{m}]$ for density s_1 [panel (a)] and s_2 [panel (b)]. $N = 1000$ samples have been drawn.

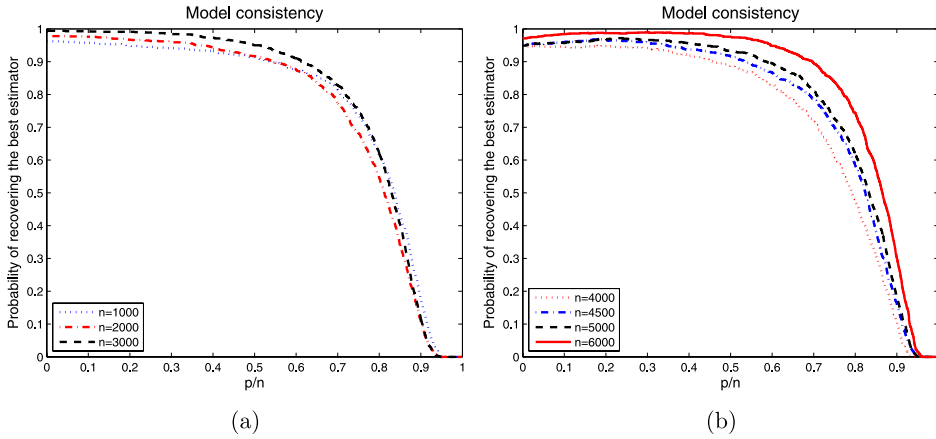


FIG. 4. $p/n \mapsto \mathbb{P}[\hat{m} = \bar{m}]$ for density s_3 [panel (a)] and s_4 [panel (b)]. $N = 1000$ samples have been drawn.

1. $s_1(t) = \frac{6}{8}\mathbb{1}_{[0,1/2]}(t) + \frac{10}{8}\mathbb{1}_{[1/2,1]}(t)$, $t \in [0, 1]$ [panel (a)],
2. $s_2(t) = \frac{135}{112}\mathbb{1}_{[0,1/3]}(t) + \frac{135}{56}\mathbb{1}_{[1/3,1/2]}(t) + \frac{1}{4}\mathbb{1}_{[1/2,5/7]}(t) + \frac{1}{2}\mathbb{1}_{[5/7,1]}(t)$, $t \in [0, 1]$ [panel (b)].

As predicted by Theorem 3.3, CV reaches model selection consistency for recovering the best parametric estimator $\hat{s}_{\bar{m}}$ on condition p/n increases to 1 as n grows to $+\infty$. Comparing (a) and (b), the convergence rate is slower in (b). Unlike (a) where \bar{m} remains almost unchanged as n increases, the best parametric estimator in (b) changes with n as allowed by (21). Therefore, the slower convergence rate in (b) results from the higher dimension of the space of piecewise constant functions s_2 belongs to.

When s does not belong to the model collection (Figure 4), densities with different smoothness assumptions have been considered:

1. $s_3(x) = \beta(10, 7; x)$, for every $t \in [0, 1]$ [panel (a)],
2. $s_4(x) = \frac{6}{5}x^{1/5}$, for every $t \in [0, 1]$ [panel (b)].

The converse situation arises since CV reaches model selection consistency as long as p/n decreases to 0 as n tends to $+\infty$. Consistently with Theorem 3.4, this rate strongly depends on the risk of the best estimator, that is on the smoothness of the target. While model selection consistency is illustrated by both panels (a) and (b), it is faster for the smoothest density s_3 than for s_4 . In Figure 4 panel (b), the highest probability is achieved for $p/n \approx 0.18$ with $n = 6000$.

4. Discussion. From the present analysis of CV procedures in the density estimation framework, we were able to prove the optimality of leave-one-out cross-validation for risk estimation, which is consistent with earlier results in the regression setting Burman (1989).

However, when CV is used as model selection procedure, the optimal p strongly depends on the structure of the model collection and on our goal (estimation or identification).

Estimation. When the best model has dimension growing with n [faster than $(\log n)^a$ for some $a > 0$] and the model collection has a polynomial complexity (Pol), Theorem 3.1 proves any p such that $p/n \rightarrow 0$ leads to an asymptotically optimal model selection procedure. This is consistent with the asymptotic equivalence between Lpo (as long as $p/n \rightarrow 0$) and Mallows' C_p previously settled in the regression setting [Shao (1997)].

From a nonasymptotic point of view, Corollary 3.1 suggests choosing $p > 1$ (for finite sample) could balance the overfitting phenomenon arising from selecting a model from a large collection. This overfitting phenomenon is already well known with penalized criteria such as Mallows' ones, inducing the need for heavier constants in front of the penalty Arlot and Massart (2009). Therefore, increasing p amounts to penalize more strongly complex models (with large dimension).

Identification. As settled by Yang (2007) for regression, Section 3.2 highlights the optimal p depends on the rate of convergence of the estimator one tries to recover (and on the structure of the model collection).

When the target estimator has a parametric rate with a polynomial collection, Theorem 3.3 proves $p/n \rightarrow 1$ leads to model selection consistency. This fact has been already noticed by Shao (1993) in the regression setting who proved leave-one-out is not model selection consistent. Remembering the asymptotic equivalence between Lpo and BIC-like criteria [Shao (1997)] established with the linear regression model, this confirms the somewhat paradoxical requirement [Yang (2006)] to devote most of available data ($p > n/2$) to the test set when trying to recover a parametric estimator.

Drawing such a simple conclusion is harder when the best estimator has a nonparametric rate as detailed by Theorem 3.4. If the best estimator has a rate close to parametric, then $p/n \rightarrow 1$ provides model selection consistency. Conversely, if the rate is slower (e.g., polynomial of order n^{-a} , for some $a > 0$), then requiring $p/n \rightarrow 0$ enables to recover the target estimator. Relating that way the optimal p to the rate of convergence of the best estimator has been already done in the regression context by [Yang (2007), see Corollary 1].

Note that when the best estimator is nonparametric (e.g., with polynomial rate), Theorem 3.1 and Theorem 3.4 imply $p/n \rightarrow 0$ leads to both efficiency and, respectively, model selection consistency. However, there is no contradiction with the earlier paper by Yang (2005) where it was proved no model selection criterion can share both efficiency and model selection consistency in a parametric setting. For instance, Li (1987) has established model selection consistency for leave-one-out with nonparametric estimators in regression.

APPENDIX A: ESTIMATION POINT OF VIEW

A.1. Main proofs.

PROOF OF THEOREM 3.1. First, let us use Proposition A.2 from [Celisse \(2014\)](#) applied with $m, m' \in \mathcal{M}_n$ such that $\widehat{R}_p(m') \leq \widehat{R}_p(m)$. Then it comes

$$\begin{aligned} & \frac{n}{n-p} \mathbb{E}[Z_{m'}^2] + \|s - s_{m'}\|^2 - K(n, p)[Z_{m'}^2 - \mathbb{E}[Z_{m'}^2]] \\ & \leq \frac{n}{n-p} \mathbb{E}[Z_m^2] + \|s - s_m\|^2 - K(n, p)[Z_m^2 - \mathbb{E}[Z_m^2]] \\ & \quad - 2K(n, p)v_n(s_{m'} - s_m) + \frac{1}{n} \left(K(n, p) + \frac{n}{n-p} \right) v_n(\phi_{m'} - \phi_m), \end{aligned}$$

where $K(n, p) = 1 + \frac{2}{n-1} + \frac{p}{n-p} \frac{1}{n-1}$.

Then, combining Propositions B.4 and B.5 from [Celisse \(2014\)](#) to control the remainder terms, there exist a sequence $(\delta_n)_{\mathbb{N}}$ with $\delta_n \rightarrow 0$ and $n\delta_n \rightarrow +\infty$ as $n \rightarrow +\infty$ and an event $\Omega = \Omega_{\text{rem},1} \cap \Omega_{\text{rem},2}$ of probability $1 - 4/n^2$ on which

$$\begin{aligned} & \frac{n}{n-p} \mathbb{E}[Z_{m'}^2] + \|s - s_{m'}\|^2 - K(n, p)[Z_{m'}^2 - \mathbb{E}[Z_{m'}^2]] \\ & \leq \frac{n}{n-p} \mathbb{E}[Z_m^2] + \|s - s_m\|^2 - K(n, p)[Z_m^2 - \mathbb{E}[Z_m^2]] \\ & \quad + \delta_n K(n, p)(\|s - s_{m'}\|^2 + \mathbb{E}[Z_{m'}^2] + \|s - s_m\|^2 + \mathbb{E}[Z_m^2]) \\ & \quad + \delta_n \left(K(n, p) + \frac{n}{n-p} \right) [\mathbb{E}[Z_{m'}^2] + \mathbb{E}[Z_m^2]]. \end{aligned}$$

In the following, δ_n always denotes such a sequence even if the precise expression of δ_n can differ from line to line.

Let us now use concentration results stated in Corollaries B.1 and B.2 from [Celisse \(2014\)](#) on the events Ω_{left} and Ω_{right} . The important point in this proof is given by Lemmas B.1 and B [[Celisse \(2014\)](#)], where it is proved that on the event $\Omega = \Omega_{\text{left}} \cap \Omega_{\text{right}} \cap \Omega_{\text{rem},1} \cap \Omega_{\text{rem},1}$, $\min\{D_{m^*}, D_{\widehat{m}(p)}\} \geq (\log n)^4$ for large enough values of n . Therefore, one can apply Lemma B.6 and Corollaries B.1 and B.2 from [Celisse \(2014\)](#) with $L_m = 0 = r_n(m)$ to get

$$\begin{aligned} & Z_{m'}^2 \left[\left(\frac{n}{n-p} (1 - \delta_n) - 2\delta_n K(n, p) \right) (1 - 4\delta_n) - 4K(n, p)\delta_n \right] \\ & \quad + [1 - \delta_n K(n, p)] \|s - s_{m'}\|^2 \\ & \leq Z_m^2 \left[\left(\frac{n}{n-p} (1 + \delta_n) + 2\delta_n K(n, p) \right) (1 + 4\delta_n) + 4K(n, p)\delta_n \right] \\ & \quad + [1 + \delta_n K(n, p)] \|s - s_m\|^2. \end{aligned}$$

Choosing $m' = \widehat{m}$, it comes

$$T_V^- Z_{\widehat{m}}^2 + T_B^- \|s - s_{\widehat{m}}\|^2 \leq T_V^+ Z_m^2 + T_B^+ \|s - s_m\|^2,$$

where

$$T_B^- = 1 - \delta_n K(n, p),$$

$$T_V^- = \frac{n}{n-p} (1 - \delta_n)[1 - 4\delta_n] - 2K(n, p)[3\delta_n - 4\delta_n^2],$$

$$T_B^+ = 1 + \delta_n K(n, p),$$

$$T_V^+ = \frac{n}{n-p} (1 + \delta_n)[1 + 4\delta_n] + 2K(n, p)[3\delta_n + 4\delta_n^2].$$

Finally on the event Ω , the following oracle inequality holds true for every $p \in \{1, n - 1\}$:

$$\|s - \widehat{s}_{\widehat{m}(p)}\|^2 \leq C_n(p) \inf_{m \in \mathcal{M}_n} \{\|s - \widehat{s}_m\|^2\} \quad \text{with } C_n(p) = \frac{T_B^+ \vee T_V^+}{T_B^- \wedge T_V^-}.$$

Moreover, on the event Ω , Lemmas B.1 and B.2 [Celisse (2014)] show $\min\{D_{m^*}, D_{\widehat{m}(p)}\} \geq (\log n)^4$. Then, it is enough to apply Propositions B.1 and B.2 from Celisse (2014) to models satisfying this constraint, which leads to the new event $\widetilde{\Omega}$ [where models with dimension smaller than $(\log n)^4$ have been omitted] of probability at least $1 - 6/n^2$. \square

PROOF OF COROLLARY 3.1. Let us recall the expression of the leading constant

$$C_n(p) = \frac{T_B^+ \vee T_V^+}{T_B^- \wedge T_V^-},$$

with

$$T_B^- = 1 - \delta_n K(n, p),$$

$$T_V^- = \frac{1}{1-p/n} (1 - \delta_n)[1 - 4\delta_n] - 2\delta_n K(n, p)[3 - 4\delta_n],$$

$$T_B^+ = 1 + \delta_n K(n, p),$$

$$T_V^+ = \frac{1}{1-p/n} (1 + \delta_n)[1 + 4\delta_n] + 2\delta_n K(n, p)[3 + 4\delta_n],$$

and $K(n, p) = 1 + \frac{2}{n-1} + \frac{p}{n-p} \frac{1}{n-1}$.

First, as long as n is large enough, simple calculations when $p = 1$ show $T_V^-(1) \leq T_B^-(1)$. Noticing moreover that $T_V^+(p) \geq T_B^+(p)$ for every p , it comes for p close to 1

$$C_n(p) = \frac{T_V^+}{T_V^-} = \frac{(1 + \delta_n)[1 + 4\delta_n] + 2(1 - p/n)\delta_n K(n, p)[3 + 4\delta_n]}{(1 - \delta_n)[1 - 4\delta_n] - 2(1 - p/n)\delta_n K(n, p)[3 - 4\delta_n]}.$$

It is then easy to show that $p \mapsto C_n(p)$ is decreasing on $\{1, \dots, p^*\}$, where p^* denotes the value of p such that $T_V^-(p) = T_B^-(p)$. Hence,

$$\frac{p_n^*}{n} = 1 - \frac{1 - 5\delta_n + 4\delta_n^2 - (2/(n - 1))(3\delta_n - 4\delta_n^2) + \delta_n/(n - 1)}{1 + 2(1 + 1/(n - 1))(3\delta_n - 4\delta_n^2) - \delta_n(1 + 1/(n - 1))}.$$

It results that for every $p \geq p^*$

$$C_n(p) = \frac{T_V^+}{T_B^-},$$

which is increasing with respect to p .

In the same way, it is easy to check that $p_n^*/(10n\delta_n) \xrightarrow{n \rightarrow +\infty} 1$, which enables us to conclude the proof. \square

PROOF OF THEOREM 3.2. Introducing the event $\tilde{\Omega}$ of Theorem 3.1, we get

$$\mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2] = \mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2 \mathbb{1}_{\tilde{\Omega}}] + \mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2 \mathbb{1}_{\tilde{\Omega}^c}].$$

Then Theorem 3.1 applied to the first expectation in the right-hand side leads to

$$\mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2] \leq C_n(p) \mathbb{E}\left[\inf_{m \in \mathcal{M}_n} \|s - \widehat{s}_m\|^2\right] + \mathbb{E}[\|s - \widehat{s}_{\widehat{m}(p)}\|^2 \mathbb{1}_{\tilde{\Omega}^c}].$$

Applying (Bias), one gets

$$\mathbb{E}[\|s - s_{\widehat{m}(p)}\|^2 \mathbb{1}_{\tilde{\Omega}^c}] \leq \mathbb{E}\left[\frac{c_u}{D_{\widehat{m}(p)}^u} \mathbb{1}_{\tilde{\Omega}^c}\right] \leq c_u \mathbb{P}(\tilde{\Omega}^c) \leq \frac{6c_u}{n^2},$$

while (RegD) and (Dmax) provide

$$\begin{aligned} & \mathbb{E}[\|s_{\widehat{m}(p)} - \widehat{s}_{\widehat{m}(p)}\|^2 \mathbb{1}_{\tilde{\Omega}^c}] \\ &= \mathbb{E}\left[\sum_{\lambda \in \Lambda(\widehat{m}(p))} (P_n \varphi_\lambda - P \varphi_\lambda)^2 \mathbb{1}_{\tilde{\Omega}^c}\right] \\ &\leq 2\mathbb{E}\left[\sum_{\lambda \in \Lambda(\widehat{m}(p))} (P_n \varphi_\lambda)^2 \mathbb{1}_{\tilde{\Omega}^c}\right] + 2\mathbb{E}\left[\sum_{\lambda \in \Lambda(\widehat{m}(p))} (P \varphi_\lambda)^2 \mathbb{1}_{\tilde{\Omega}^c}\right] \\ &\leq 2\mathbb{E}\left[\sum_{\lambda \in \Lambda(\widehat{m}(p))} \frac{1}{n^2} \sum_{i,j=1}^n \varphi_\lambda(X_i) \varphi_\lambda(X_j) \mathbb{1}_{\tilde{\Omega}^c}\right] + 2\|s\|^2 \mathbb{E}[D_{\widehat{m}(p)} \mathbb{1}_{\tilde{\Omega}^c}] \\ &\leq 2(\Phi + \|s\|^2) \frac{n}{(\log n)^2} \mathbb{P}(\tilde{\Omega}^c) \leq (\Phi + \|s\|^2) \frac{12}{n(\log n)^2}. \end{aligned} \quad \square$$

APPENDIX B: IDENTIFICATION POINT OF VIEW

B.1. Proof of Theorem 3.3. The general purpose is to prove there exist an event Ω_n with $\mathbb{P}(\Omega_n) \rightarrow 1$ as n tends to $+\infty$ and a positive integer N such that on Ω_n , for every $n \geq N$, every $m \neq \bar{m}$ satisfies

$$(27) \quad \widehat{R}_p(\bar{m}) - \widehat{R}_p(m) \leq -u_n(m)(1 + o(1)),$$

where $u_n(m) > 0$ denotes a real number for every n and m , and $o(1)$ does not depend on m . In particular, this implies

$$(28) \quad \mathbb{P}(\widehat{m} = \bar{m}) = \mathbb{P}(\forall m \neq \bar{m}, \widehat{R}_p(\bar{m}) - \widehat{R}_p(m) < 0) \xrightarrow{n \rightarrow +\infty} 1,$$

which would complete the proof.

Let us consider the event $\Omega_{\text{left}} \cap \Omega_{\text{right}}$ in Proposition C.2 from Celisse (2014) with $\beta_1 = \beta_2 = 1/n^2$, and the events $\Omega_{\text{rem},1}$ [Proposition B.4 from Celisse (2014)] and $\Omega_{\text{rem},3}$ (Proposition C.1 [Celisse (2014)]). Then with $\Omega_n = \Omega_{\text{left}} \cap \Omega_{\text{right}} \cap \Omega_{\text{rem},1} \cap \Omega_{\text{rem},3}$ and $\mathbb{P}[\Omega_n^c] \leq 8/n^2$, showing (28) amounts to prove

$$\mathbb{P}(\Omega_n \cap \{\forall m \neq \bar{m}, \widehat{R}_p(\bar{m}) - \widehat{R}_p(m) < 0\}) \xrightarrow{n \rightarrow +\infty} 1.$$

Let us now focus on the event Ω_n . The two main steps correspond to distinguishing between *parametric* and *nonparametric* models S_m [see (P-NP)]. For every m , let us define $B(m) = \|s - s_m\|^2$ and $V(m) = \mathbb{E}[\|s_m - \widehat{s}_m\|^2]$, where $s_m = \text{Argmin}_{t \in S_m} \|s - t\|^2$. From line to line, the value of δ_n may change, but it always denotes a sequence decreasing to 0 and such that $n\delta_n \rightarrow +\infty$ as n grows.

If \widehat{s}_m has a parametric rate.

- If $s \in S_m$:

Let us first notice $s_m = s = s_{\bar{m}}$, which implies $R_n(m) = V(m)$ and $R_n(\bar{m}) = V(\bar{m})$. Then Proposition A.2 from Celisse (2014), and Propositions B.4 and C.2 from Celisse (2014) lead to

$$\begin{aligned} & \left| [\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] - \frac{n}{n-p} [R_n(\bar{m}) - R_n(m)] \right| \\ & \leq \left(36L_n^2 + 3\delta_n + \delta_n \frac{n}{n-p} \right) [R_n(m) + R_n(\bar{m})] \\ & = o\left(\frac{n}{n-p}\right) [R_n(m) + R_n(\bar{m})], \end{aligned}$$

by requiring $L_n^2 = o((1 - p/n)^{-1})$, which provides (27) by use of (BeMo). Note that in the previous inequality, $r_n(\bar{m})$ and $r_n(m)$ [from Proposition C.2 from Celisse (2014)] have been omitted since they are negligible with respect to the other terms.

- If $s \notin S_m$:

Similarly, Proposition A.2, Proposition B.4, and Propositions C.1 and C.2 from Celisse (2014) lead to

$$\begin{aligned} & \left| [\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] - [B(\bar{m}) - B(m)] - \frac{n}{n-p}[V(\bar{m}) - V(m)] \right| \\ & \leq 6\delta_n[B(\bar{m}) + B(m)] \\ & \quad + \left(36L_n^2 + 3\delta_n \frac{\Phi}{\xi} + \delta_n \left(3 + \frac{n}{n-p} \right) \right) [V(m) + V(\bar{m})] \\ & \quad + 3\delta_n \|s\| \sqrt{\frac{\Phi}{\xi} \sqrt{V(\bar{m}) + V(m)}} + r_n(\bar{m}). \end{aligned}$$

With $s \in \bigcup_{m'} S_{m'}$ and $s \notin S_m$, it comes $B(\bar{m}) = 0$ and $B(m) \geq \tau > 0$ by (P-NP). Since both \widehat{s}_m and $\widehat{s}_{\bar{m}}$ have parametric rates, requiring $n(1 - p/n) \rightarrow +\infty$ as n grows implies (27), that is,

$$[\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] \leq -B(m)(1 + o(1)).$$

If \widehat{s}_m has a nonparametric rate.

- If $s \in S_m$:

Proposition A.2 from Celisse (2014), $s_m = s = s_{\bar{m}}$, and Propositions C.1 and C.2 from Celisse (2014) combined with $L_n^2 = o(n/(n-p))$ provide

$$\begin{aligned} & \left| [\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] - \frac{n}{n-p}[V(\bar{m}) - V(m)] \right| \\ & \leq o\left(\frac{n}{n-p}\right)V(m) + o\left(\frac{n}{n-p}\right)V(\bar{m}), \end{aligned}$$

where $o(n/(n-p))$ does not depend on m . Since S_m is nonparametric, (P-NP) gives

$$\frac{V(\bar{m})}{V(m)} \leq \frac{\pi}{(\log n)^2} \left(n / (\log n)^2 \inf_m V(m) \right)^{-1} \xrightarrow{n \rightarrow +\infty} 0,$$

which implies (27) with

$$(n-p)[\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] \leq -nV(m)(1 + o(1)).$$

- If $s \notin S_m$:

Both $L_n^2 = o(n/(n-p))$ and the same argument as above lead to

$$\begin{aligned} & \left| [\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] + B(m) + \frac{n}{n-p}V(m)[1 + o(1)] \right| \\ & \leq V(m)o\left(\frac{n}{n-p}\right) + B(m)o(1). \end{aligned}$$

Then (27) holds true with

$$[\widehat{R}_p(\bar{m}) - \widehat{R}_p(m)] \leq -B(m)(1 + o(1)) - \frac{n}{n-p} V(m)(1 + o(1)).$$

Then there exists an integer N such that for $n \geq N$, on the event Ω_n , (27) holds true for every $m \in \mathcal{M}_n$, which completes the proof.

Acknowledgements. We thank Sylvain Arlot and Stéphane Robin for helpful discussions, and also the Associate Editor as well as two anonymous referees for their comments and suggestions to improve the earlier version of the paper.

SUPPLEMENTARY MATERIAL

Supplement to “Optimal cross-validation in density estimation with the L^2 -loss”: Technical proofs and details (DOI: [10.1214/14-AOS1240SUPP](https://doi.org/10.1214/14-AOS1240SUPP); .pdf). Owing to space constraints, we have moved technical proofs to a supplementary document [[Celisse \(2014\)](#)].

REFERENCES

- AKAIKE, H. (1973). Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory (Tsahkadsor, 1971)* 267–281. Akadémiai Kiadó, Budapest. [MR0483125](#)
- ARLOT, S. (2007). V -fold penalization: An alternative to v -fold cross-validation. In *Oberwolfach Reports, Volume 4 of Mathematisches Forschungsinstitut*. EMS, Zürich.
- ARLOT, S. and CELISSE, A. (2010). A survey of cross-validation procedures for model selection. *Stat. Surv.* **4** 40–79. [MR2602303](#)
- ARLOT, S. and CELISSE, A. (2011). Segmentation of the mean of heteroscedastic data via cross-validation. *Stat. Comput.* **21** 613–632. [MR2826696](#)
- ARLOT, S. and MASSART, P. (2009). Data-driven calibration of penalties for least-squares regression. *Journal of Machine Learning* **10** 245–279.
- BARAUD, Y., GIRAUD, C. and HUET, S. (2009). Gaussian model selection with an unknown variance. *Ann. Statist.* **37** 630–672. [MR2502646](#)
- BARRON, A., BIRGÉ, L. and MASSART, P. (1999). Risk bounds for model selection via penalization. *Probab. Theory Related Fields* **113** 301–413. [MR1679028](#)
- BARRON, A. R. and COVER, T. M. (1991). Minimum complexity density estimation. *IEEE Trans. Inform. Theory* **37** 1034–1054. [MR1111806](#)
- BARTLETT, P., BOUCHERON, S. and LUGOSI, G. (2002). Model selection and error estimation. *Machine Learning* **48** 85–113.
- BIRGÉ, L. and MASSART, P. (1997). From model selection to adaptive estimation. In *Festschrift for Lucien Le Cam* (D. Pollard, E. Torgensen and G. Yang, eds.) 55–87. Springer, New York. [MR1462939](#)
- BIRGÉ, L. and MASSART, P. (2001). Gaussian model selection. *J. Eur. Math. Soc. (JEMS)* **3** 203–268. [MR1848946](#)
- BIRGÉ, L. and MASSART, P. (2007). Minimal penalties for Gaussian model selection. *Probab. Theory Related Fields* **138** 33–73. [MR2288064](#)
- BIRGÉ, L. and ROZENHOLC, Y. (2006). How many bins should be put in a regular histogram. *ESAIM Probab. Stat.* **10** 24–45 (electronic). [MR2197101](#)

- BLANCHARD, G. and MASSART, P. (2006). Discussion: “Local Rademacher complexities and oracle inequalities in risk minimization” [*Ann. Statist.* **34** (2006) 2593–2656] by V. Koltchinskii. *Ann. Statist.* **34** 2664–2671. [MR2329460](#)
- BOWMAN, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika* **71** 353–360. [MR0767163](#)
- BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A. and STONE, C. J. (1984). *Classification and Regression Trees*. Wadsworth, Belmont, CA. [MR0726392](#)
- BURMAN, P. (1989). A comparative study of ordinary cross-validation, v -fold cross-validation and the repeated learning-testing methods. *Biometrika* **76** 503–514. [MR1040644](#)
- BURMAN, P. (1990). Estimation of optimal transformations using v -fold cross validation and repeated learning-testing methods. *Sankhyā Ser. A* **52** 314–345. [MR1178041](#)
- CASTELLAN, G. (1999). Modified Akaike’s criterion for histogram density estimation. Technical Report 99.61, Univ. Paris-Sud.
- CASTELLAN, G. (2003). Density estimation via exponential model selection. *IEEE Trans. Inform. Theory* **49** 2052–2060. [MR2004713](#)
- CELISSE, A. (2014). Supplement to “Optimal cross-validation in density estimation with the L^2 -loss.” DOI:10.1214/14-AOS1240SUPP.
- CELISSE, A. and ROBIN, S. (2008). Nonparametric density estimation by exact leave- p -out cross-validation. *Comput. Statist. Data Anal.* **52** 2350–2368. [MR2411944](#)
- DEVORE, R. A. and LORENTZ, G. G. (1993). *Constructive Approximation. Grundlehren der Mathematischen Wissenschaften* **303**. Springer, Berlin. [MR1261635](#)
- GEISSER, S. (1974). A predictive approach to the random effect model. *Biometrika* **61** 101–107. [MR0418322](#)
- GEISSER, S. (1975). The predictive sample reuse method with applications. *J. Amer. Statist. Assoc.* **70** 320–328.
- IBRAGIMOV, I. A. and HAS’MINSKIĪ, R. Z. (1981). *Statistical Estimation: Asymptotic Theory. Applications of Mathematics* **16**. Springer, Berlin. [MR0620321](#)
- LARSON, S. C. (1931). The shrinkage of the coefficient of multiple correlation. *J. Educ. Psychol.* **22** 45–55.
- LEDoux, M. (2001). *The Concentration of Measure Phenomenon. Mathematical Surveys and Monographs* **89**. Amer. Math. Soc., Providence, RI. [MR1849347](#)
- LI, K.-C. (1987). Asymptotic optimality for C_p , C_L , cross-validation and generalized cross-validation: Discrete index set. *Ann. Statist.* **15** 958–975. [MR0902239](#)
- LUGOSI, G. and NOBEL, A. B. (1999). Adaptive model selection using empirical complexities. *Ann. Statist.* **27** 1830–1864. [MR1765619](#)
- MALLOWS, C. L. (1973). Some comments on C_p . *Technometrics* **15** 661–675.
- MOSTELLER, F. and TUKEY, J. W. (1968). Data analysis, including statistics. In *Handbook of Social Psychology, Vol. 2* (G. Lindzey and E. Aronson, eds.). Addison-Wesley, New York.
- RUDEMO, M. (1982). Empirical choice of histograms and kernel density estimators. *Scand. J. Stat.* **9** 65–78. [MR0668683](#)
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. [MR0468014](#)
- SHAO, J. (1993). Linear model selection by cross-validation. *J. Amer. Statist. Assoc.* **88** 486–494. [MR1224373](#)
- SHAO, J. (1997). An asymptotic theory for linear model selection. *Statist. Sinica* **7** 221–264. [MR1466682](#)
- STONE, M. (1974). Cross-validators choice and assessment of statistical predictions. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **36** 111–147. [MR0356377](#)
- STONE, M. (1977). An asymptotic equivalence of choice of model by cross-validation and Akaike’s criterion. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **39** 44–47. [MR0501454](#)
- STONE, C. J. (1984). An asymptotically optimal window selection rule for kernel density estimates. *Ann. Statist.* **12** 1285–1297. [MR0760688](#)

- STONE, C. J. (1985). An asymptotically optimal histogram selection rule. In *Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer, Vol. II* (Berkeley, CA, 1983). *Wadsworth Statist./Probab. Ser.* 513–520. Wadsworth, Belmont, CA. [MR0822050](#)
- TALAGRAND, M. (1996). New concentration inequalities in product spaces. *Invent. Math.* **126** 505–563. [MR1419006](#)
- WEGKAMP, M. (2003). Model selection in nonparametric regression. *Ann. Statist.* **31** 252–273. [MR1962506](#)
- YANG, Y. (2005). Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation. *Biometrika* **92** 937–950. [MR2234196](#)
- YANG, Y. (2006). Comparing learning methods for classification. *Statist. Sinica* **16** 635–657. [MR2267253](#)
- YANG, Y. (2007). Consistency of cross validation for comparing regression procedures. *Ann. Statist.* **35** 2450–2473. [MR2382654](#)
- YANG, Y. and BARRON, A. R. (1998). An asymptotic property of model selection criteria. *IEEE Trans. Inform. Theory* **44** 95–116. [MR1486651](#)
- ZHANG, P. (1993). Model selection via multifold cross validation. *Ann. Statist.* **21** 299–313. [MR1212178](#)

LABORATOIRE DE MATHÉMATIQUES PAINLEVÉ
UMR 8524 CNRS-UNIVERSITÉ LILLE 1
MODAL PROJECT-TEAM INRIA
F-59655, VILLENEUVE D'ASCQ CEDEX
FRANCE
E-MAIL: celisse@math.univ-lille1.fr