Published in final edited form as:

Bernoulli (Andover). 2021 November ; 27(4): 2300–2336. doi:10.3150/20-BEJ1309.

Universal sieve-based strategies for efficient estimation using machine learning tools

HONGXIANG QIU^{1,*}, ALEX LUEDTKE², MARCO CARONE^{1,*,*}

¹Department of Biostatistics, University of Washington, Seattle, WA, USA.

²Department of Statistics, University of Washington, Seattle, WA, USA.

Abstract

Suppose that we wish to estimate a finite-dimensional summary of one or more function-valued features of an underlying data-generating mechanism under a nonparametric model. One approach to estimation is by plugging in flexible estimates of these features. Unfortunately, in general, such estimators may not be asymptotically efficient, which often makes these estimators difficult to use as a basis for inference. Though there are several existing methods to construct asymptotically efficient plug-in estimators, each such method either can only be derived using knowledge of efficiency theory or is only valid under stringent smoothness assumptions. Among existing methods, sieve estimators stand out as particularly convenient because efficiency theory is not required in their construction, their tuning parameters can be selected data adaptively, and they are universal in the sense that the same fits lead to efficient plug-in estimators for a rich class of estimands. Inspired by these desirable properties, we propose two novel universal approaches for estimating function-valued features that can be analyzed using sieve estimation theory. Compared to traditional sieve estimators, these approaches are valid under more general conditions on the smoothness of the function-valued features by utilizing flexible estimates that can be obtained, for example, using machine learning.

Keywords

nonparametric inference; asymptotic efficiency; sieve estimation

1. Introduction

1.1. Motivation

A common statistical problem consists of using available data in order to learn about a summary of the underlying data-generating mechanism. In many cases, this summary involves function-valued features of the distribution that cannot be estimated at a parametric rate under a nonparametric model — for example, a regression function or the density function of the distribution. Examples of useful summaries involving such features include average treatment effects [23], average derivatives [11], moments of the conditional mean function [24], variable importance measures [33] and treatment effect heterogeneity

* qiuhx@uw.edu; ** mcarone@uw.edu.

measures [12]. For ease of implementation and interpretation, in traditional approaches to estimation, these features have typically been restricted to have simple forms encoded by parametric or restrictive semiparametric models. However, when these models are misspecified, both the interpretation and validity of subsequent inferences can be compromised. To circumvent this difficulty, investigators have increasingly relied on machine learning (ML) methods to flexibly estimate these function-valued features.

Once estimates of the function-valued features are obtained, it is natural to consider plug-in estimators of the summary of interest. However, in general, such estimators are not root- n -consistent and asymptotically normal, and hence not asymptotically efficient (referred to as *efficient* henceforth). Lacking this property is problematic since it often forms the basis for constructing valid confidence intervals and hypothesis tests [3, 20]. When the function-valued features are estimated by ML methods, in order for the plug-in estimator to be CAN, the ML methods must not only estimate the involved function-valued features well, but must also satisfy a small-bias property with respect to the summary of interest [20, 31].

Unfortunately, because ML methods generally seek to optimize out-of-sample performance, they seldom satisfy the latter property.

1.2. Existing methodological frameworks

The targeted minimum loss-based estimation (TMLE) framework provides a means of constructing efficient plug-in estimators [27, 31]. Given an (almost arbitrary) initial ML fit that provides a good estimate of the function-valued features involved, TMLE produces an adjusted fit such that the resulting plug-in estimator has reduced bias and is efficient. This adjustment process is referred to as targeting since a generic estimate of the function-valued features is modified to better suit the goal of estimating the summary of interest. Though TMLE provides a general template for constructing efficient estimators, its implementation requires specialized expertise, namely knowledge of the analytic expression for an influence function of the summary of interest. Influence functions arise in semiparametric efficiency theory and are key to establishing efficiency, but can be difficult to derive. Furthermore, even when an influence function is known analytically, additional expertise is needed to construct a TMLE for a given problem.

Alternative approaches for constructing efficient plug-in estimators have been proposed in the literature, including the use of undersmoothing [18], twicing kernels [20], and sieves [4, 19, 24]. These methods neither require knowing an influence function nor performing any targeting of the function-valued feature estimates. Hence, the same fits can be used to simultaneously estimate different summaries of the data-generating distribution, even if these summaries were not pre-specified when obtaining the fit. These approaches also circumvent the difficulties in obtaining an influence function. However, these methods all rely on smoothness conditions on derivatives of the functional features that may be overly stringent. In addition, undersmoothing provides limited guidance on the choice of the tuning parameter; the twicing kernel method requires the use of a higher-order kernel, which may lead to poor performance in small to moderate samples [14].

In contrast, under some conditions, sieve estimation can produce a flexible fit with the optimal out-of-sample performance while also yielding an efficient — and therefore root- n -

consistent and asymptotically normal — plug-in estimator [24]. In this paper, we focus on extensions of this approach. In sieve estimation, we first assume that the unknown function falls in a rich function space, and construct a sequence of approximating subspaces indexed by sample size that increase in complexity as sample size grows. We require that, in the limit, the functions in the subspaces can approximate any function in the rich function space arbitrarily well. These approximating subspaces are referred to as *sieves*. By using an ordinary fitting procedure that optimizes the estimation of the function-valued feature within the sieve, the bias of the plug-in estimator can decrease sufficiently fast as the sieve grows in order for that estimator to be efficient. Thus sieve estimation requires no explicit targeting for the summary of interest.

The series estimator is one of the best known and most widely used sieve techniques. These sieves are taken as the span of the first finitely many terms in a basis that is chosen by the user to approximate the true function well. Common choices of the basis include polynomials, splines, trigonometric series and wavelets, among others. However, series estimators usually require strong smoothness assumptions on derivative of the unknown function in order for the flexible fit to converge at a sufficient rate to ensure the resulting plug-in estimator is efficient. As the dimension of the problem increases, the smoothness requirement may become prohibitive. Moreover, even if the smoothness assumption is satisfied, a prohibitively large sample size may be needed for some series estimators to produce a good fit. For example, if the unknown function is smooth but is a constant over a region, estimation based on a polynomial series can perform poorly in small to moderate samples.

Series estimators may also require the user to choose the number of terms in the series in such a way that results in a sufficient convergence rate. The rates at which the number of terms should grow with sample size have been thoroughly studied (e.g. [4, 19, 24]). However, these results only provide minimal guidance for applications because there is no indication on how to select the actual number of terms for a given sample size. In practice, the number of terms in the series is often chosen by CV. Upper bounds on the convergence rate of the series estimator as a function of sample size and the number of terms have been derived, and it has been shown that the optimal number of terms that minimizes the bound can also lead to an efficient plug-in estimator [24]. However, CV tends to select the number of terms that optimizes the actual convergence rate [29], which may differ from the number of terms minimizing the derived bound on the convergence rate. Even though the use of CV-tuned sieve estimators has achieved good numerical performance, to the best of our knowledge, there is no theoretical guarantee that they lead to an efficient plug-in estimator.

Two variants of traditional series estimators were proposed in [3]. These methods can use two bases to approximate the unknown function-valued features and the corresponding gradient separately, whereas in traditional series estimators, only one basis is used for both approximations. Consequently, these variants may be applied to more general cases than traditional series estimators. However, like traditional series estimators, they also suffer from the inflexibility of the pre-specified bases.

1.3. Contributions and organization of this article

In this paper we present two approaches that can partially overcome these shortcomings.

1. *Estimating the unknown function with Highly Adaptive Lasso (HAL)* [1, 26].

If we are willing to assume the unknown functions have a finite variation norm, then they may be estimated via HAL. If the tuning parameter is chosen carefully, then we may obtain an efficient plug-in estimator. This method can help overcome the stringent smoothness assumptions on derivatives that are required by existing series estimators, as we discussed earlier.

2. *Using data-adaptive series based on an initial ML fit.*

As long as the initial ML algorithm converges to the unknown function at a sufficient rate, we show that, for certain types of summaries, it is possible to obtain an efficient plug-in estimator with a particular data-adaptive series. The smoothness assumption on the unknown function can be greatly relaxed due to the introduction of the ML algorithm into the procedure. Moreover, for summaries that are highly smooth, we show that the number of terms in the series can be selected by CV.

Although the first approach is not an example of sieve estimation, both approaches are motivated by the sieve literature and can be shown to lead to asymptotically efficient plug-in estimators using the sieve estimation theory derived in [24]. The flexible fits of the functional features from both approaches can be plugged in for a rich class of estimands.

We remark that, although we do not have to restrict ourselves to the plug-in approach in order to construct an asymptotically efficient estimator, other estimators do not overcome the shortcomings described in Section 1.2 and can have other undesirable properties. For example, the popular one-step correction approach (also called debiasing in the recent literature on high-dimensional statistics) [22] constructs efficient estimators by adding a bias reduction term to the plug-in estimator. Thus, it is not a plug-in estimator itself, and as a consequence, one-step estimators may not respect known constraints on the estimand — for example, bounds on a scalar-valued estimand (e.g., the estimand is a probability and must lie in $[0, 1]$) or shape constraints on a vector-valued estimand (e.g., monotonicity constraints). This drawback is also typical for other non-plug-in estimators, such as those derived via estimating equations [30] and double machine learning [5, 6]. Additionally, as with the other procedures described above, the one-step correction approach requires the analytic expression of an influence function.

Our paper is organized as follows. We introduce the problem setup and notation in Section 2. We consider plug-in estimators based on HAL in Section 3, data-adaptive series in Section 4, and its generalized version that is applicable to more general summaries in Section 5. Section 6 concludes with a discussion. Technical proofs of lemmas and theorems (Appendix D), simulation details (Appendix E) and other additional details are provided in the Appendix.

2. Problem setup and traditional sieve estimation review

Suppose we have independent and identically distributed observations $V_1, \dots, V_n$ drawn from P_0 . Let Θ be a class of functions, and denote by $\theta_0 \in \Theta$ a (possibly vector-valued) functional feature of P_0 — for example, θ_0 may be a regression function. Throughout this paper we assume that the generic data unit is $V = (X, Z) \sim P_0$, where X is a (possibly vector-valued) random variable corresponding to the argument of θ_0 , and Z may also be a vector-valued random variable. In some cases $V = X$ and Z is trivial. We use $\mathcal{X}$ to denote the support of X . The estimand of interest is a finite-dimensional summary $\Psi(\theta_0)$ of θ_0 . We consider a plug-in estimator $\Psi(\hat{\theta}_n)$, where $\hat{\theta}_n$ is an estimator of θ_0 , and aim for this plug-in estimator to be asymptotically linear, in the sense that

$\Psi(\hat{\theta}_n) = \Psi(\theta_0) + n^{-1} \sum_{i=1}^n \text{IF}(V_i) + o_p(n^{-1/2})$ with IF an influence function satisfying $\mathbb{E}_{P_0}[\text{IF}(V)] = 0$ and $\mathbb{E}_{P_0}[\text{IF}(V)^2] < \infty$. This estimator is efficient under a nonparametric model if the estimator is also regular. By the central limit theorem and Slutsky's theorem, it follows that $\Psi(\hat{\theta}_n)$ is a CAN estimator of $\Psi(\theta_0)$, and therefore, $\sqrt{n}[\Psi(\hat{\theta}_n) - \Psi(\theta_0)] \xrightarrow{d} N(0, \mathbb{E}_{P_0}[\text{IF}(V)^2])$. This provides a basis for constructing valid confidence intervals for $\Psi(\theta_0)$.

We now list some examples of such problems.

Example 1. Moments of the conditional mean function [24]: Let $\theta_0: x \mapsto \mathbb{E}_{P_0}[Z | X = x]$ be the conditional mean function. The κ -th moment of $\theta_0(X)$, $X \sim P_0$, namely $\Psi_\kappa(\theta_0) = \mathbb{E}_{P_0}[\theta_0^\kappa(X)]$, can be a summary of interest. The values of $\Psi_1(\theta_0)$ and $\Psi_2(\theta_0)$ are useful for defining the proportion of $\text{Var}_{P_0}(Z)$ that is explained by X , which may be written as $\text{Var}_{P_0}(\theta_0(X))/\text{Var}_{P_0}(Z)$. This proportion is a measure of variable importance [33]. Generally, we may consider $\Psi(\theta_0) = \mathbb{E}_{P_0}[f(\theta_0(X))]$ for a fixed function f .

Example 2. Average derivative [11]: Let X follow a continuous distribution on $\mathbb{R}^d$ and $\theta_0: x \mapsto \mathbb{E}_{P_0}[Z | X = x]$ be the conditional mean function. Let θ'_0 denote the vector of partial derivatives of θ_0 . Then $\Psi(\theta_0) = \mathbb{E}_{P_0}[\theta'_0(X)]$ summarizes the overall (adjusted) effect of each component of X on Y . Under certain conditions, we can rewrite $\Psi(\theta_0) = \mathbb{E}_{P_0}[\theta_0(X)p'_0(X)/p_0(X)]$, where p_0 is the Lebesgue density of X and p'_0 is the vector of partial derivatives of p_0 . This expression clearly shows the important role of the Lebesgue density of X in this summary.

Example 3. Mean counterfactual outcome [23]: Suppose that $Z = (A, Y)$ where A is a binary treatment indicator and Y is the outcome of interest. Let $\theta_0: x \mapsto \mathbb{E}_{P_0}[Y | A = 1, X = x]$ be the outcome regression function under treatment value 1. Under causal assumptions, the mean counterfactual outcome corresponding to the intervention that assigns treatment 1 to the

entire population can be nonparametrically identified by the G -computation formula

$$\Psi(\theta_0) = \mathbb{E}_{P_0}[\theta_0(X)].$$

Example 4. Treatment effect heterogeneity measures [12]: Similarly to Example 3, suppose that A is a binary treatment indicator and Z is the outcome of interest. Let $\theta_0 = (\mu_{00}, \mu_{01})^T$, where $\mu_{0a}: x \mapsto \mathbb{E}_{P_0}[Z \mid A = a, X = x]$ is the outcome regression function for treatment arm $a \in \{0, 1\}$. Then, $\Psi(\theta_0) = \text{Var}_{P_0}(\mu_{01}(X) - \mu_{00}(X))$ is an overall summary of treatment effect heterogeneity.

To obtain an asymptotically linear plug-in estimator, $\hat{\theta}_n$ must converge to θ_0 at a sufficiently fast rate and approximately solve an estimating equation to achieve the small bias property with respect to the summary of interest [20, 26, 31]. For simplicity, we assume the estimand to be scalar-valued — when the estimand is vector-valued, we can treat each entry as a separate estimand, and the plug-in estimators of all entries are jointly asymptotically linear if each estimator is asymptotically linear. Therefore, this leads to no loss in generality if the same fits are used for all entries in the summary of interest.

Sieve estimation allows us to obtain an estimator $\Psi(\hat{\theta}_n)$ with the small bias property with respect to $\Psi(\theta_0)$ while maintaining the optimal convergence rate of $\hat{\theta}_n$ [4, 24]. The construction of sieve estimators is based on a sequence of approximating spaces Θ_n to Θ . These approximating spaces are referred to as *sieves*. Usually Θ_n is much simpler than Θ to avoid over-fitting but complex enough to avoid under-fitting. For example, Θ_n can be the space of all polynomials with degree K or splines with K knots with $K = K(n) \rightarrow \infty$ as $n \rightarrow \infty$. In this paper, with a loss function l such that $\theta_0 \in \arg\min_{\theta \in \Theta} \mathbb{E}_{P_0}[l(\theta)(V)]$, we consider estimating θ_0 by minimizing an empirical risk based on l , i.e., $\hat{\theta}_n \in \arg\min_{\theta \in \Theta_n} n^{-1} \sum_{i=1}^n l(\theta)(V_i)$. Under some conditions, the growth rate of Θ_n can be carefully chosen so that $\Psi(\hat{\theta}_n)$ is an asymptotically linear estimator of $\Psi(\theta_0)$ while $\hat{\theta}_n$ converges to θ_0 at the optimal rate.

Throughout this paper, for a probability distribution P and an integrable function f with respect to P , we define $Pf := \int f(v)dP(v) = \mathbb{E}_P[f(V)]$. We use P_n to denote the empirical distribution. We take $\langle \cdot, \cdot \rangle$ to be the $L^2(P_0)$ -inner product, i.e., $\langle \theta_1, \theta_2 \rangle = P_0(\theta_1^\top \theta_2)$, where $L^2(P_0)$ is the set of real-valued P_0 -squared-integrable functions defined on the support of P_0 . When the functions are vector-valued, we take $\langle \theta_1, \theta_2 \rangle = P_0(\theta_1^\top \theta_2)$. We use $\|\cdot\|$ to denote the induced norm of $\langle \cdot, \cdot \rangle$. We assume that $\Theta \subseteq L^2(P_0)$. We remark that we have committed to a specific choice of inner product and norm to fix ideas; other inner products can also be adopted, and our results will remain valid upon adaptation of our upcoming conditions. We discuss this explicitly via a case study in Appendix A.

For the methods we propose in this article, we assume that Θ is convex. Throughout this paper, we will further require a set of conditions similar to those in [24]. For any $\theta \in \Theta$, let

$\ell'_0[\theta - \theta_0](v) := \lim_{\delta \rightarrow 0} [\ell(\theta_0 + \delta(\theta - \theta_0))(v) - \ell(\theta_0)(v)]/\delta$ be the Gâteaux derivative of ℓ at θ_0 in the direction $\theta - \theta_0$ and $r[\theta - \theta_0](v) := \ell(\theta)(v) - \ell(\theta_0)(v) - \ell'_0[\theta - \theta_0](v)$ be the corresponding remainder.

Condition A1 (Linearity and boundedness of Gâteaux derivative operator of loss function). *For all $\theta \in \Theta$, $\theta \in \Theta$, $\ell'_0[\theta - \theta_0]$ exists and $\ell'_0[\theta - \theta_0](v) - P_0 \ell'_0[\theta - \theta_0]$ is linear and bounded in $\theta - \theta_0$.*

Condition A2 (Local quadratic behavior of loss function). *There exists a constant $\alpha_{0,1} \in (0, \infty)$ such that, for all $\theta \in \Theta$ such that $P_0 \{ \ell(\theta) - \ell(\theta_0) \}$ or $\|\theta - \theta_0\|$ is sufficiently small, it holds that $P_0 \{ \ell(\theta) - \ell(\theta_0) \} = \alpha_{0,1} \|\theta - \theta_0\|^2 / 2 + o(\|\theta - \theta_0\|^2)$.*

Remark 1. We now present an equivalent form of A2 that may be easier to verify in practice. For all $\theta \in \Theta \setminus \{\theta_0\}$, define $h_\theta := (\theta - \theta_0) / \|\theta - \theta_0\|$ and $a_\theta := \left. \frac{d^2}{d\delta^2} P_0 \ell(\theta_0 + \delta h_\theta) \right|_{\delta=0}$. Requiring Condition A2 is equivalent to requiring that $a_{\theta_1} = a_{\theta_2}$ for all $\theta_1, \theta_2 \in \Theta \setminus \{\theta_0\}$ and that

$$\sup_{\theta \in \Theta} \left| P_0 \ell(\theta_0 + \delta h_\theta) - P_0 \ell(\theta_0) - \frac{a_\theta \delta^2}{2} \right| = o(\delta^2).$$

Moreover, if A2 holds, then, for any $\theta \in \Theta \setminus \{\theta_0\}$, it is true that $\alpha_{0,1} = a_\theta$.

A large class of loss functions satisfy Conditions A1 and A2. For example, in the regression setting where Z is the outcome, the squared-error loss $\ell(\theta) : v \mapsto [z - \theta(x)]^2$ and the logistic loss $\ell(\theta) : v \mapsto -z\theta(x) + \log\{1 + \exp(\theta(x))\}$ both satisfy these conditions; a negative working log-likelihood usually also satisfies these conditions. In Examples 1–4, the unknown functions are all conditional mean functions, which can be estimated with the above loss functions. Thus, Conditions A1 and A2 hold. Examples 3 and 4 require a slight modification discussed in more details in Appendix A. We also note that Condition A2 is sufficient for Condition B in [24].

Condition A3 (Differentiability of summary of interest). $\Psi'_{\theta_0}[\theta - \theta_0] := \lim_{\delta \rightarrow 0} [\Psi(\theta_0 + \delta(\theta - \theta_0)) - \Psi(\theta_0)]/\delta$ exists for all $\theta \in \Theta$ and is a linear bounded operator.

If Condition A3 holds, then, by the Riesz representation theorem, $\Psi'_{\theta_0}[\theta - \theta_0] = \langle \theta - \theta_0, \dot{\Psi} \rangle$ for a gradient function $\dot{\Psi} = \dot{\Psi}_{\theta_0}$ in the completion of the space spanned by $\Theta - \theta_0 := \{x \mapsto \theta(x) - \theta_0(x) : \theta \in \Theta\}$.

Condition A4 (Locally quadratic remainder). *There exists a constant $C > 0$ so that, for all θ with sufficiently small $\|\theta - \theta_0\|$, it holds that*

$$\left| \Psi(\theta) - \Psi(\theta_0) - \Psi'_{\theta_0}[\theta - \theta_0] \right| \leq C \|\theta - \theta_0\|^2.$$

The above condition states that the remainder of the linear approximation to Ψ is locally bounded by a quadratic function.

Conditions A3 and A4 hold for Examples 1–4. For the generalized moment of the conditional mean function in Example 1, it holds that $\dot{\Psi} = f' \circ \theta_0$. For the average derivative of the conditional mean function in Example 2, it holds that $\dot{\Psi} = p'_0/p_0$. For the average treatment effect and the treatment effect heterogeneity measure in Examples 3 and 4, as we show in Appendix A, $\dot{\Psi}$ also exists and depends on the propensity score function $x \mapsto P_0(A = 1 | X = x)$.

3. Estimation with Highly Adaptive Lasso

3.1. Brief review of Highly Adaptive Lasso

Recently, the Highly Adaptive Lasso (HAL) was proposed as a flexible ML algorithm that only requires a mild smoothness condition on the unknown function and has a well-described implementation [1, 26]. In this subsection, we briefly review HAL. We first heuristically introduce its definition and desirable properties, and then introduce the definition and implementation more formally. For ease of presentation, for the moment, we assume that θ_0 is real-valued.

In HAL, θ_0 is assumed to fall in the class of càdlàg functions (right-continuous with left limits) defined on $\mathcal{X} \subseteq \mathbb{R}^d$ with variation norm bounded by a finite constant M . In this section, we denote this function class by $\Theta_{v,M}$. The variation norm of a càdlàg function θ , denoted by $\|\theta\|_v$, characterizes the total variability of θ as its argument ranges over the domain, so $\|\cdot\|_v$ is a global smoothness measure and $\Theta_{v,M}$ is a large function class that even contains functions with discontinuities. Fig. 1 presents some examples of univariate càdlàg functions with finite variation norms for illustration. Because $\Theta_{v,M}$ is a rich class, it can be plausible that $\theta_0 \in \Theta_{v,M}$ for some $M < \infty$. The HAL estimator of θ_0 is then $\hat{\theta}_n = \hat{\theta}_{n,M} \in \argmin_{\theta \in \Theta_{v,M}} n^{-1} \sum_{i=1}^n \ell(\theta)(V_i)$. Under this assumption, it has been shown that $\|\hat{\theta}_n - \theta_0\| = o_p(n^{-1/4})$ regardless of the dimension of X under additional mild conditions [26]. Thus, estimation with HAL replaces the usual smoothness requirement on derivatives of traditional series estimators by a requirement on global smoothness, namely $\theta_0 \in \Theta_{v,M}$ for some M .

We next formally present the definition of variation norm of a càdlàg function θ : $[x^{(\ell)}, x^{(u)}] \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$. Here, $x^{(\ell)}$ and $x^{(u)}$ are vectors in $\mathbb{R}^d$, with $\cdot$ being entrywise, $[x^{(\ell)}, x^{(u)}] := \{x \in \mathbb{R}^d : x^{(\ell)} \leq x \leq x^{(u)}\}$.

For any nonempty index set $s \subseteq \{1, 2, \dots, d\}$ and any $x = (x_1, x_2, \dots, x_d) \in [x^{(l)}, x^{(u)}]$, we define $x_s := \{x_j: j \in s\}$ and $x_{-s} := \{x_j: j \in \{1, 2, \dots, d\} \setminus s\}$ to be entries of x with indices in and not in s respectively. We defined the s -section of θ as $\theta_s := \theta(x_1 1(1 \in s), x_2 1(2 \in s), \dots, x_d 1(d \in s))$. We can subsequently obtain the following representation of θ at any $x \in [x^{(l)}, x^{(u)}]$ in terms of sums and integrals of the variation of s -sections of θ [9]:

$$\theta(x) = \theta(x^{(\ell)}) + \sum_{s \in \{1, \dots, d\}, s \neq \emptyset} \int_{x^{(\ell)}, x} \theta_s(d\tilde{x}).$$

The variation norm is then subsequently defined as

$$\|\theta\|_v := \|\theta(x^{(\ell)})\| + \sum_{s \in \{1, \dots, d\}, s \neq \emptyset} \int_{x^{(\ell)}, x^{(u)}} |\theta_s(d\tilde{x})|.$$

We refer to [1] and [26] for more details on variation norm. Notably, this notion of variation norm coincides with that of Hardy and Krause [21].

We finally briefly introduce the algorithm to compute a HAL estimator. It can be shown that an empirical risk minimizer in $\Theta_{v,M}$ is a step function that only jumps at sample points, namely

$$x \mapsto \beta_0 + \sum_{s \subseteq \{1, \dots, d\}, s \neq \emptyset} \sum_{j=1}^n \mathbb{1}(X_{j,s} \leq x_s) \beta_{s,j}.$$

Here, β_0 and all $\beta_{s,j}$ are real numbers. To find an empirical risk minimizer in $\Theta_{v,M}$ in the above form, we may solve the following optimization problem:

$$\min_{\theta} \sum_{i=1}^n \ell(\theta)(V_i)$$

$$\text{subject to } \theta: x \mapsto \beta_0 + \sum_{s \subseteq \{1, \dots, d\}, s \neq \emptyset} \sum_{j=1}^n \mathbb{1}(X_{j,s} \leq x_s) \beta_{s,j}$$

$$|\beta_0| + \sum_{s \subseteq \{1, \dots, d\}, s \neq \emptyset} \sum_{j=1}^n |\beta_{s,j}| \leq M.$$

The constraint imposes an upper bound on the l_1 norm of a vector. Therefore, for common loss functions, we may use software for LASSO regression [25]. For example, if the loss function is the squared-error loss, then we may run a LASSO linear regression to obtain a HAL estimate.

3.2. Estimation with an oracle tuning parameter

In this section, we consider plug-in estimators based on HAL. For ease of illustration, for the rest of this section, we consider scalar-valued Ψ , and will discuss vector-valued Ψ only at the end of this subsection. We further introduce the following conditions needed to establish that the HAL-based plug-in estimator is efficient.

Condition B1 (Càdlàg functions). θ_0 and $\dot{\Psi}$ are càdlàg.

Condition B2 (Bound on variation norm). For some $M < \infty$, $\|\theta_0\|_v + \|\dot{\Psi}\|_v \leq M$.

Condition B2 ensures that certain perturbations of θ_0 still lie in $\Theta_{v,M}$, a crucial requirement for proving the asymptotic linearity of our proposed plug-in estimator. In addition, since $\dot{\Psi}$ may depend on components of P_0 other than θ_0 as in Examples 2–4, Conditions B1–B2 may also impose conditions on these components.

In this section, we fix an M that satisfies Condition B2. Additional technical conditions can be found in Appendix B.1. Let $\hat{\theta}_n = \hat{\theta}_{n,M} \in \operatorname{argmin}_{\theta \in \Theta_{v,M}} n^{-1} \sum_{i=1}^n \ell(\theta)(V_i)$ denote the HAL fit obtained using the bound M in Condition B2.

We note that $\hat{\theta}_n$ is not a typical sieve estimator because M is fixed and there is no explicit sequence of growing approximating spaces Θ_n . Nevertheless, we may view this method as a special case of sieve estimation with degenerate sieves $\Theta_n = \Theta_{v,M}$ for all n . This allows us to utilize existing results [24] to show the asymptotic linearity and efficiency of the plug-in estimator based on $\hat{\theta}_n$. We next formally present this result.

Theorem 1 (Efficiency of plug-in estimator). *Under Conditions A1–A4 and B1–B4, $\Psi(\hat{\theta}_n)$ is an asymptotically linear estimator of $\Psi(\theta_0)$ with the influence function being*

$$v \mapsto \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\Psi}](v) + \mathbb{E}_{P_0}[\ell'_0[\dot{\Psi}](V)] \right\}, \text{ that is,}$$

$$\Psi(\hat{\theta}_n) = \Psi(\theta_0) + \frac{1}{n} \sum_{i=1}^n \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\Psi}](V_i) + \mathbb{E}_{P_0}[\ell'_0[\dot{\Psi}](V)] \right\} + o_p(n^{-1/2}).$$

As a consequence, $\sqrt{n}[\Psi(\hat{\theta}_n) - \Psi(\theta_0)] \xrightarrow{d} N(0, \xi^2)$ with $\xi^2 := \operatorname{Var}_{P_0}(\ell'_0[\dot{\Psi}](V)) / \alpha_{0,\ell}^2$. In addition, under Conditions E1 and E2 in Appendix B.4, $\Psi(\hat{\theta}_n)$ is efficient under a nonparametric model.

We note that, for HAL to achieve the optimal convergence rate, we only need that $M \geq \|\theta_0\|_v$ [1, 26]. The requirement of a larger M imposed by Condition B2 resembles undersmoothing [18], as using a larger M would result in a fit that is less smooth than that based on the CV-selected bound. The $L^2(P_0)$ -convergence rate of the flexible fit using the larger bound remains the same, but the leading constant may be larger. This is in contrast to traditional undersmoothing, which leads to a fit with a suboptimal rate of convergence.

Under some conditions, the following lemma provides a loose bound on $\|\dot{\Psi}\|_v$ in the case that $\dot{\Psi}$ has a particular structure. Such a bound can be used to select an appropriate bound on variation norm that satisfies Condition B2.

Lemma 1. *Suppose that $\dot{\Psi} = \dot{\psi} \circ \theta_0$, where $\dot{\psi}: \mathbb{R} \rightarrow \mathbb{R}$ is differentiable. Let $x^{(l)} = \sup\{x : P_0(X$
 $x) = 1\}$ where $\sup$ and $\inf$ are entrywise. Assume that θ_0 is differentiable. If each of
 $\|\theta_0\|_v, \|\dot{\psi}(x^{(l)})\|$ and $B := \sup_{z: |z| \leq \|\theta_0\|_v} |\dot{\psi}'(z)|$ is finite, then $\|\dot{\Psi}\|_v \leq B\|\theta_0\|_v + \|\dot{\psi}(x^{(l)})\|$. Hence,
 $\|\theta_0\|_v + \|\dot{\Psi}\|_v \leq (B + 1)\|\theta_0\|_v + \|\dot{\psi}(x^{(l)})\| < \infty$.*

As we discussed at the end of Section 2, such structures as $\dot{\Psi} = \dot{\psi} \circ \theta_0$ are common, especially if we augment θ_0 to include other implicitly relevant components of P_0 . For example, in Example 2, we may augment θ_0 with p_0 and p'_0 ; in Examples 3 and 4, we may augment θ_0 with the propensity score function.

When θ_0 is p'_0 -valued, θ_0 can often be viewed as a collection of q real-valued variation-independent functions $\eta_{10}, \dots, \eta_{q0}$. In this case, we can define $\Theta_{v,M} = \{(\eta_1, \dots, \eta_q) : \eta_j \text{ is c\`adl\`ag, } \|\eta_j\|_v \leq M_j, j = 1, \dots, q\}$ for a positive vector $M = (M_1, \dots, M_q)$. The subsequent arguments follow analogously, where now each η_j is treated as a separate function.

We remark that an undersmoothing condition such as B2 appears to be necessary for a HAL-based plug-in estimator to be efficient. We illustrate this numerically in Section 3.3. The choice of a sufficiently large bound M required by Theorem 1 is by no means trivial, since this choice requires knowledge that the user may not have. Nevertheless, this result forms the basis of the data-driven method that we propose in Section 3.3 for choosing M . We also remark that, if we wish to plug in the same $\hat{\theta}_n$ based on HAL for a rich estimands, the chosen bound M needs to be sufficiently large for all estimands of interest.

Another method to construct efficient plug-in estimators based on HAL has been independently developed [28]. Unlike our approach based on sieve theory, in this work, the authors directly analyzed the first-order bias of the plug-in estimator using influence functions. In terms of ease of implementation, their method requires specifying a constant involved in a threshold of the empirical mean of the basis functions, which may be difficult to specify in applications. Our approach in Section 3.3 may also require specifying an unknown constant to obtain a valid upper bound on $\|\dot{\Psi}\|_v$, but in some cases the constant may be set to zero, and our simulation suggests that the performance is not sensitive to the choice of the constant.

3.3. Data-adaptive selection of the tuning parameter

Since it is hard to prespecify a bound M on the variation norm that is sufficiently large to satisfy Condition B2 but also sufficiently small to avoid overfitting for a given data set, it is desirable to select M in a data-adaptive manner. A seemingly natural approach makes use of k -fold CV. In particular, for each candidate bound M , partition the data into k folds of

approximately equal size (k is fixed and does not depend on n), in each fold evaluate the performance of the HAL estimator fitted on all other folds based on this candidate M , and use the candidate bound M_n with the best average performance across all folds to obtain the final fit. It has been shown that $\hat{\theta}_{n, M_n}$ can achieve the optimal convergence rate under mild conditions [29], but M_n appears not to satisfy Condition B2 in general. In particular, the derived bound on $\|\hat{\theta}_n - \theta_0\|$ relies on an empirical process term, namely $\sup_{\theta \in \Theta_{v, M}} \left| (P_n - P_0) \{ \ell(\theta) - \ell(\theta_0) \} \right|$, and a larger M implies a larger space $\Theta_{v, M}$. Therefore, the bound on $\|\hat{\theta}_n - \theta_0\|$ grows with M . Because k -fold CV seeks to optimize out-of-sample performance, M_n generally appears to be close to $\|\theta_0\|_v$ and not sufficiently large to obtain an efficient plug-in estimator.

To avoid this issue with the CV-selected bound, we propose a method that takes inspiration from k -fold CV, but modifies the bound so that it is guaranteed to yield an efficient plug-in estimator for $\Psi(\theta_0)$. This method may require the analytic expression for $\dot{\Psi}$. In Sections 4 and 5, we present methods that do not require this knowledge.

1. Derive an upper bound on $\|\dot{\Psi}\|_v$. This bound is a non-decreasing function of the variation norms of functions that can be learned from data (e.g., using Lemma 1). In other words, find a non-decreasing function F such that $\|\dot{\Psi}\|_v \leq F(\|\eta_{10}\|_v, \dots, \|\eta_{q0}\|_v)$ for unknown functions $\eta_{10}, \dots, \eta_{q0}$ that can be assumed to be càdlàg with finite variation norm and can be estimated with HAL.
2. Estimate $\theta_0, \eta_{10}, \dots, \eta_{q0}$ by HAL with k -fold CV, and denote the CV-selected bounds for these functions by $M_n, M_{1n}, \dots, M_{qn}$.
3. For a small $\epsilon > 0$, use the bound $M_n + \epsilon + F(M_{1n} + \epsilon, \dots, M_{qn} + \epsilon)$ to estimate θ_0 with HAL and plug in the fit. We refer to this step of slightly increasing the bounds as ϵ -relaxation.

It follows from Lemma 2 in the Appendix that this method would yield a sufficiently large bound with probability tending to one. In practice, it is desirable for the bound derived on $\|\dot{\Psi}\|_v$ to be relatively tight to avoid choosing an overly large bound that leads to overfitting in small to moderate samples. We remark that multiplying by $1 + \epsilon$ rather than adding ϵ to each argument also leads to a valid choice for the bound; that is, the bound $M_n(1 + \epsilon) + F(M_{1n}(1 + \epsilon), \dots, M_{qn}(1 + \epsilon))$ is also sufficiently large with probability tending to one. In practice, the user may increase each CV-selected bound by, for example, 5% or 10%. Although it is more natural and convenient to directly use $M_n + F(M_{1n}, \dots, M_{qn})$ as the bound, we have only been able to prove the result with a small ϵ -relaxation. However, if the bound is loose and F is continuous, we can show that ϵ -relaxation is unnecessary. The formal argument can be found after Lemma 2 in the Appendix.

As for methods based on knowledge of an influence function, deriving $\dot{\Psi}$ and a bound for its variation norm requires some expertise, but in some cases this task can be straightforward. The derivation of an influence function is typically based on a fluctuation in the space of distributions, but in many cases, the relation between such fluctuations and the summary of

interest is implicit and difficult to handle. In contrast, the derivation of $\dot{\Psi}$ is based on a fluctuation of θ_0 , and the summary of interest explicitly depends on θ_0 . As a consequence, it can be simpler to derive $\dot{\Psi}$ than to derive an influence function. For example, for the summary $\Psi_\kappa(\theta_0) = P_0\theta_0^\kappa$ in Example 1, we find that $\dot{\Psi}_\kappa = \kappa\theta_0^{\kappa-1}$ by straightforward calculation, whereas the influence function given in Theorem 1 is more difficult to directly derive analytically.

We illustrate the fact that M_n may not be sufficiently large and show that our proposed method resolves this issue via a simulation study in which $\theta_0: x \mapsto \mathbb{E}_{P_0}[Y | X = x]$ and

$\Psi: \theta_0 \mapsto P_0\theta_0^2$. We compare the performance of the plug-in estimators based on the 10-fold CV-selected bound on variation norm (M.cv), the bound derived from the analytic expression of $\dot{\Psi}$ with and without ϵ -relaxation (M.gcv+ and M.gcv respectively), and a sufficiently large oracle choice satisfying Condition B2 (M.oracle). We According to Lemma 1, M.oracle is $3\|\theta_0\|_V$ and M.gcv is $3 \times \text{M.cv}$. We also investigate the performance of 95% Wald CIs based on the influence function. For each resulting plug-in estimator, we investigate the following quantities: $n \cdot \text{MSE}$, $\sqrt{n} \cdot |\text{bias}|$ and CI coverage. More details of this simulation are provided in Appendix E. In theory, for an efficient estimator, we should find that $n \cdot \text{MSE}$ tends to a constant (the variance of the influence function $\xi^2 := P_0\text{IF}^2$), $\sqrt{n} \cdot |\text{bias}|$ tends to 0, and 95% Wald CIs have approximately 95% coverage.

We report performance summaries in Fig 2 and Table 1 with this criterion, from which it appears that the plug-in estimators with M.oracle and M.gcv+ achieve efficiency, while the plug-in estimator based on M.cv does not. The desirable performance of M.oracle and M.gcv+ agrees with the available theory, whereas the poor performance of M.cv suggests that cross-validation may not yield a valid choice of variation norm in general. Interestingly, M.gcv performs similarly to M.oracle and M.gcv+. We conjecture that using an ϵ -relaxation is unnecessary in this setting. In Fig 3, we can also see that M.cv tends to $\|\theta_0\|_V$ and has a high probability of being less than M.oracle. Therefore, this simulation suggests that using a sufficiently large bound — in particular, a bound larger than the CV-selected bound — may be necessary and sufficient for the plug-in estimator to achieve efficiency.

4. Data-adaptive series

4.1. Proposed method

For ease of illustration, we consider the case that Ψ is scalar-valued in this section. As we will describe next, our proposed estimation procedure for function-valued features does not rely on Ψ and hence can be used for a class of summaries.

Suppose that Θ is a vector space of $\mathbb{R}^q$ -valued functions equipped with the $L_2(P_0)$ -inner product. Further, suppose that $\Psi = \psi \circ \theta_0$ for some function: $\psi: \mathbb{R}^q \rightarrow \mathbb{R}^q$. This holds, for example, when $\Psi: \theta \mapsto P_0(f \circ \theta)$ for a fixed differentiable function f in Example 1. In this case, $\dot{\Psi} = f' \circ \theta_0$ and hence $\dot{\psi} = f'$. Particularly useful examples include Examples 1 and 4.

For now we assume that the marginal distribution of X is known so that we only need to

estimate θ_0 for this summary. We will address the more difficult case in which the marginal distribution of X is unknown in Section 4.3.

Let θ_n^0 be a given initial flexible ML fit of θ_0 and consider the data-adaptive sieve-like subspaces based on $\theta_n^0, \Theta_n := \Theta_{n, \theta_n^0} := \text{Span}\{\phi_1, \phi_2, \dots, \phi_K\} \circ \theta_n^0$, where $\phi_1, \phi_2, \dots$ are $\mathbb{R}^q$ -valued basis functions in a series defined on $\mathbb{R}^q$ and $K = K(n)$ is a deterministic number of terms in the series — we will consider selecting K via CV in Section 4.4. Let $\theta_n^* = \theta_n^*(\theta_n^0) \in \arg\min_{\theta \in \Theta_n} n^{-1} \sum_{i=1}^b \ell(\theta)(V_i)$ denote the series estimator within this data-adaptive sieve-like subspace that minimizes the empirical risk. We propose to use $\Psi(\theta_n^*)$ to estimate $\Psi(\theta_0)$.

4.2. Results for a deterministic number of terms

Following [4, 24], our proofs of the validity of our data-adaptive series approach make heavy use of projection operators. We use $\pi_n := \pi_{n, \theta_n^0}$ to denote the projection operator for functions in Θ onto $\Theta_n = \Theta_{n, \theta_n^0}$ with respect to $\langle \cdot, \cdot \rangle$. For any function $\theta \in \Theta$, let $\Pi_{n, \theta}$ denote the operator that takes as input a function $g: \mathbb{R}^q \rightarrow \mathbb{R}^q$ for which $g \circ \theta \in L^2(P_0)$ and outputs a function $\Pi_{n, \theta}(g): \mathbb{R}^q \rightarrow \mathbb{R}^q$ such that $\Pi_{n, \theta}(g) \circ \theta = \pi_{n, \theta}(g \circ \theta)$. In other words, letting β_j be the quantity that depends on g and θ such that $\pi_{n, \theta}(g \circ \theta) = \left(\sum_{j=1}^K \beta_j \phi_j\right) \circ \theta$, we define $\Pi_{n, \theta}(g) := \sum_{j=1}^K \beta_j \phi_j$. The operator $\Pi_{n, \theta}$ may also be interpreted as follows: letting P_θ be the distribution of $\mathcal{A}(X)$ with $V = (X, Z) \sim P_0$, then $\Pi_{n, \theta}$ is the projection operator of functions $\mathbb{R}^q \rightarrow \mathbb{R}^q$ with respect to the $L^2(P_\theta)$ -inner product. We use $\mathcal{I}$ to denote the identity function in $\mathbb{R}^q$.

We now present additional conditions we will require to ensure that $\Psi(\theta_n^*)$ is an efficient estimator of $\Psi(\theta_0)$.

Condition C1 (Sufficient convergence rate of initial ML fit). $\|\theta_n^0 - \theta_0\| = o_p(n^{-1/4})$.

Condition C2 (Sufficiently small estimation error). $\|\theta_n^* - \pi_n(\theta_0)\| = o_p(n^{-1/4})$.

Condition C3 (Sufficiently small approximation error to $\mathcal{I}$ for Θ_{n, θ_0^0}).

$$\|\theta_0 - \Pi_{n, \theta_0^0}(\mathcal{I}) \circ \theta_0\| = o(n^{-1/4}).$$

Condition C4 (Sufficiently small approximation error to ψ for Θ_{n, θ_0} and convergence rate of

θ_n^*). $\left\| \left[\psi - \Pi_{n, \theta_0}(\psi) \right] \circ \theta_0 \right\| \cdot \|\theta_n^* - \theta_0\| = o_p(n^{-1/2})$ Appendix B.2 contains further technical conditions and Appendix C discusses their plausibility. As discussed in Appendix C,

Conditions C2–C4 typically imply restrictions on the growth rate of K : if K grows too fast with n , then Condition C2 may be violated; if K instead grows too slow, then Conditions C3 and C4 may be violated. For the generalized moment $\Psi : \theta \mapsto P_0(f \circ \theta)$ with a fixed known function f in Example 1, Condition C4 typically also imposes a smoothness condition f so that f can be approximated by the series well. Our conditions are closely related to the conditions in Theorem 1 of [24]. Conditions C1–C3 and C6 serve as sufficient conditions for the condition on the smoothness of Ψ and the convergence rate of θ_n^* in Theorem 1 of [24].

Together with Conditions C4 and C7, we can derive Lemma 4, which is similar to the first part of Condition C of [24]. The empirical process condition C8 is sufficient for Conditions A, D and the second part of C in Theorem 1 in [24].

We now present a theorem ensuring the asymptotic linearity and efficiency of the plug-in estimator based on θ_n^* .

Theorem 2 (Efficiency of plug-in estimator). *Under Conditions A1–A4 and C1–C9, $\Psi(\theta_n^*)$ is an asymptotically linear estimator of $\Psi(\theta_0)$ with the influence function being*

$$v \mapsto \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\Psi}](v) + \mathbb{E}_{P_0}[\ell'_0[\dot{\Psi}](V)] \right\}, \text{ that is,}$$

$$\Psi(\theta_n^*) = \Psi(\theta_0) + \frac{1}{n} \sum_{i=1}^n \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\Psi} \circ \theta_0](V_i) + \mathbb{E}_{P_0}[\ell'_0[\dot{\Psi} \circ \theta_0](V)] \right\} + o_p(n^{-1/2}).$$

As a consequence, $\sqrt{n}[\Psi(\hat{\theta}_n) - \Psi(\theta_0)] \xrightarrow{d} N(0, \xi^2)$ with $\xi^2 := \text{Var}_{P_0}(\ell'_0[\dot{\Psi}](V))/\alpha_{0,\ell}^2$. In addition, under Conditions E1 and E2 in Appendix B.4, $\Psi(\hat{\theta}_n)$ is efficient under a nonparametric model.

Remark 2. Consider the general case in which it may not be true that $\dot{\Psi}$ can be represented as $\dot{\psi} \circ \theta_0$ for some $\dot{\psi}: \mathbb{R}^q \rightarrow \mathbb{R}^q$. If the analytic expression of $\dot{\Psi}$ can be derived and $\dot{\Psi}$ can be estimated by $\dot{\Psi}_n$ such that $\|\dot{\Psi}_n - \dot{\Psi}\| \cdot \|\theta_n^0 - \theta_0\| = o_p(n^{-1/2})$, then our data-adaptive series can take a special form that is targeted towards Ψ . Specifically, letting $\vartheta_0 := (\theta_0, \dot{\Psi})^\top$ and $\Psi(\vartheta_0) := \Psi(\theta_0)$, it is straightforward to show that the gradient of Ψ is $\dot{\Psi} = (\dot{\Psi}, 0)^\top = (e_2, 0)^\top \vartheta_0$ with $\mathbf{0} = (0, 0)^\top$ and $e_2 = (0, 1)^\top$, which is a function composed with ϑ_0 . We can set $\vartheta_n^0 = (\theta_n^0, \dot{\Psi}_n)$ and $\Theta_n = \text{Span}\{\theta_n^0, \dot{\Psi}_n\}$ in our data-adaptive series. This approach does not have a growing number of terms in Θ_n and is not similar to sieve estimation, but can be treated as a special case of data-adaptive series. It can be shown that Conditions C1–C4 are still satisfied for ϑ and Ψ with this choice of Θ_n , and hence our data-adaptive series estimator leads to an efficient plug-in estimator. We remark that the introduction of ϑ and Ψ is a purely theoretical device, and this targeted approach to estimation is quite similar to that used in the context of TMLE [27, 31].

4.3. Summaries involving the marginal distribution of X

We now generalize the setting considered thus far by allowing the parameter to depend both on θ_0 and on P_0 , i.e., estimating $\Psi(\theta_0, P_0)$. The example given at the beginning of Section 4.1, namely that of estimating $\Psi(\theta_0) = P_0(f \circ \theta_0)$, is a special case of this more general setting. In what follows, we will make use of the following conditions:

Condition D1 (Conditions with P_0 fixed). *When we regard $\Psi(\theta_0, P_0)$ as the mapping $\theta \mapsto \Psi(\theta, P_0)$ evaluated at θ_0 , Conditions A1–A4, C1–C4 and C6–C9 are satisfied for estimating $\Psi(\theta_0, P_0)$.*

Condition D2 (Hadamard differentiability with θ_0 fixed). *The mapping $P \mapsto \Psi(\theta_0, P)$ is Hadamard differentiable at P_0 .*

By the functional delta method, it follows that $\Psi(\theta_0, P_n) = \Psi(\theta_0, P_0) + P_n \text{IF}_0 + o_p(n^{-1/2})$ for a function IF_0 satisfying $P_0 \text{IF}_0 = 0$ and $P_0 \text{IF}_0^2 < \infty$.

Condition D3 (Negligible second-order difference).

$$[\Psi(\theta_n^*, P_n) - \Psi(\theta_0, P_n)] - [\Psi(\theta_n^*, P_0) - \Psi(\theta_0, P_0)] = o_p(n^{-1/2}).$$

This condition usually holds, for example, when $\Psi(\theta_0, P_0) = P_0(f \circ \theta_0)$, as in this case the left-hand side is equal to $(P_n - P_0)(f \circ \theta_n^* - f \circ \theta_0)$, which is $o_p(n^{-1/2})$ under empirical process conditions.

Theorem 3 (Asymptotic linearity of plug-in estimator). *Under Conditions D1–D3, $\Psi(\theta_n^*, P_n)$ is an asymptotically linear estimator of $\Psi(\theta_0, P_0)$ with influence function*

$$v \mapsto \alpha_{0, \ell}^{-1} \left\{ -\ell'_0[\dot{\psi} \circ \theta_0](v) + \mathbb{E}_{P_0}[\ell'_0[\dot{\psi} \circ \theta_0](V)] \right\} + \text{IF}_0(V),$$

that is,

$$\Psi(\theta_n^*, P_n) = \Psi(\theta_0, P_0) + \frac{1}{n} \sum_{i=1}^n \left\{ -\alpha_{0, \ell}^{-1} \ell'_0[\dot{\psi} \circ \theta_0](V_i) + \alpha_{0, \ell}^{-1} \mathbb{E}_{P_0}[\ell'_0[\dot{\psi} \circ \theta_0](V)] + \text{IF}(V_i) \right\} + o_p(n^{-1/2}).$$

As a consequence, $\sqrt{n}[\Psi(\theta_n^*, P_n) - \Psi(\theta_0, P_0)] \xrightarrow{d} N(0, \xi^2)$ with

$$\xi^2 := \text{Var}_{P_0}(\alpha_{0, \ell}^{-1} \ell'_0[\dot{\psi} \circ \theta_0](V) + \text{IF}(V)).$$

This result is easy to verify by decomposing $\Psi(\theta_n^*, P_n) - \Psi(\theta_0, P_0)$ as

$$[\Psi(\theta_n^*, P_0) - \Psi(\theta_0, P_0)] + [\Psi(\theta_0, P_n) - \Psi(\theta_0, P_0)] + \{[\Psi(\theta_n^*, P_n) - \Psi(\theta_0, P_n)] - [\Psi(\theta_n^*, P_0) - \Psi(\theta_0, P_0)]\}$$

Moreover, under conditions similar to the conditions E1 and E2 given in Appendix B.4, we can show that $\Psi(\theta_n^*, P_n)$ is efficient under a nonparametric model.

Remark 3. Conditions D2 and D3 can be relaxed. Specifically, if $\hat{P}_n$ is an estimator of P_0 that satisfies that $\Psi(\theta_0, \hat{P}_n) = \Psi(\theta_0, P_0) + P_n \text{IF}_0 + o_p(n^{-1/2})$ for an influence function IF_0 and Condition D3 holds with P_n replaced by $\hat{P}_n$, then $\Psi(\theta_n^*, \hat{P}_n)$ is an asymptotically linear estimator of $\Psi(\theta_0, P_0)$.

4.4. CV selection of the number of terms in data-adaptive series

In the preceding subsections, we established the efficiency of the plug-in estimator based on suitable rates of growth for K relative to the sample size n . In this subsection, we show that, under some conditions, such a K can be selected by k -fold CV: after obtaining θ_n^0 for each K in a range of candidates, we can calculate the cross-validated risk from k folds and choose the value of K with the smallest CV risk. We denote the number of terms in the series that CV selects by K^* . In this section, we use K in the subscripts for notation related to data-adaptive sieves-like spaces and projections; this represents a slight abuse of notation because, in Sections 4.1 and 4.2, these subscripts were instead used for sample size n . That is, we use $\Theta_{K,\theta}$ to denote $\text{Span}\{\varphi_1, \varphi_2, \dots, \varphi_K\} \circ \theta$, $\pi_{K,\theta}$ to denote the projection onto $\Theta_{K,\theta}$, $\Pi_{K,\theta}$ to denote the operator such that $\Pi_{K,\theta}(g) \circ \theta = \pi_{K,\theta}(g \circ \theta)$ for all $g: \mathbb{R}^q \rightarrow \mathbb{R}^q$ with $g \circ \theta \in L^2(P_0)$, and $\theta_n^\# := \theta_{K^*}^*(\theta_n^0)$ to be the data-adaptive series estimator based on θ_n^0 and K^* .

Condition C5 (Bounded approximation error of ψ relative to $\mathcal{F}$). *There exists a constant $C > 0$ such that, with probability tending to one,*

$$\left\| \psi \circ \theta_n^0 - \Pi_{K, \theta_n^0}(\psi) \circ \theta_n^0 \right\| \leq C \left\| \theta_n^0 - \Pi_{K, \theta_n^0(\mathcal{F})}(\theta_n^0) \right\|.$$

This condition is equivalent to

$$\left\| \psi - \Pi_{K, \theta_n^0}(\psi) \right\|_{L^2(P_{\theta_n^0})} \leq C \left\| \mathcal{F} - \Pi_{K, \theta_n^0(\mathcal{F})} \right\|_{L^2(P_{\theta_n^0})}$$

for all K with probability tending to one, which may be interpreted in terms of two simultaneous requirements. The first requirement is that the identity function $\mathcal{F}$ is not exactly contained in the span of $\varphi_1, \dots, \varphi_K$ for any K , since otherwise, the right-hand side would be zero for all sufficiently large K . Therefore, common series such as polynomial and spline series are not permitted for general summaries. In contrast, other series such as trigonometric series and wavelets satisfy this requirement. The second requirement is that the approximation error of the chosen series for the identity function $\mathcal{F}$ is not much larger than ψ . If a trigonometric or wavelet series is used, then this condition imposes a strong smoothness condition on derivatives of ψ . Nonetheless, this may not be stringent in some interesting examples. For example, if $\Psi(\theta) = P_0(f \circ \theta)$ for a fixed function f in Example 1,

then $\dot{\psi}$ equals f' and hence can be expected to satisfy this strong smoothness condition provided that f is infinitely differentiable with bounded derivatives. The estimands encountered in many applications involve f satisfying this smoothness condition.

The following theorem justifies the use of k -fold CV to select K under appropriate conditions.

Theorem 4 (Efficiency of CV-based plug-in estimator). *Assume that Conditions A1–A4, C1–C3, C5, C8 and C9 hold for a deterministic $K = K(n)$. Suppose part (a) of Condition C7 holds, then, with $\theta_n^\# := \theta_{K^*}(\theta_n^0)$, $\Psi(\theta_n^\#)$ is an asymptotically linear estimator of $\Psi(\theta_0)$ with influence function $v \mapsto \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\psi} \circ \theta_0](v) + \mathbb{E}_{P_0}[\ell'_0[\dot{\psi} \circ \theta_0](V)] \right\}$, that is,*

$$\Psi(\theta_n^\#) = \Psi(\theta_0) + \frac{1}{n} \sum_{i=1}^n \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\dot{\psi} \circ \theta_0](V_i) + \mathbb{E}_{P_0}[\ell'_0[\dot{\psi} \circ \theta_0](V)] \right\} + o_p(n^{-1/2}).$$

As a consequence, $\sqrt{n}[\Psi(\theta_n^\#) - \Psi(\theta_0)] \xrightarrow{d} N(0, \xi^2)$ with $\xi^2 := \text{Var}_{P_0}(\ell'_0[\dot{\psi} \circ \theta_0](V)) / \alpha_{0,\ell}^2$. In addition, under Conditions E1 and E2 in Appendix B.4, $\Psi(\hat{\theta}_n)$ is efficient under a nonparametric model.

4.5. Simulation

4.5.1. Demonstration of Theorem 4—We illustrate our method in a simulation in which we take $\theta_0(x) = \mathbb{E}_{P_0}[Z | X = x]$ and $\Psi(\theta_0) = P_0\theta_0^2$. This is a special case of Example 1.

The true function θ_0 is chosen to be discontinuous, which violates the smoothness assumptions commonly required in traditional series estimation. In this case, $\dot{\psi} = 2\mathcal{J}$ and so the constant in Condition C5 is 2. We compare the performance of plug-in estimators based on three different nonparametric regressions: (i) polynomial regression with degree selected by 10-fold CV (poly), which results in a traditional sieve estimator, (ii) gradient boosting (xgb) [7, 8, 15, 16], and (iii) data-adaptive trigonometric series estimation with gradient boosting as the initial ML fit and 10-fold CV to select the number of terms in the series (xgb.trig). We also compare these plug-in estimators with the one-step correction estimator [22] based on gradient boosting (xgb.1step). Further details of this simulation can be found in Appendix E.

Fig 4 presents $n \cdot \text{MSE}$ and $\sqrt{n} \cdot |\text{bias}|$ for each estimator, whereas Table 2 presents the coverage probability of 95% Wald CIs based on these estimators. We find that xgb.trig and xgb.1step estimators perform well, while poly and xgb plug-in estimators do not appear to be efficient. Since polynomial series estimators only work well when estimating smooth functions, in this simulation, we would not expect the fit from the polynomial series estimator to converge sufficiently fast, and consequently, we would not expect the resulting plug-in estimator to be efficient. In contrast, gradient boosting is a flexible ML method that can learn discontinuous functions, so we can expect an efficient plug-in estimator based on this ML method. However, gradient boosting is not designed to approximately solve the

estimating equation that achieves the small-bias property for this particular summary, so we would not expect its naïve plug-in estimator to be efficient. Based on gradient boosting, our estimator and the one-step corrected estimator both appear to be efficient, but our method has the advantage of being a plug-in estimator. Moreover, the construction of our estimator does not require knowledge of the analytic expression of an influence function.

We also investigate the effect of the choice of K on the performance of our method. Fig 5 presents $n \cdot \text{MSE}$ for the data-adaptive series estimator with different choices of K . We can see that our method is insensitive to the choice of K in this simulation setting. Although a relatively small K performs better, choosing a much larger K does not appear to substantially harm the behavior of the estimator. This insensitivity to the selected tuning parameter suggests that in some applications, without using CV, an almost arbitrary choice of K that is sufficiently large might perform well.

4.5.2. Violation of Condition C5—For the k -fold CV selection of K in our method to yield an efficient plug-in estimator, Ψ must be highly smooth in the sense that ψ can be approximated by the series about as well as can the identity function (see Condition C5). Although we have argued that this condition is reasonable, in this section, we explore via simulation the behavior of our method based on CV when ψ is rough. We again take $\theta_0: x \mapsto \mathbb{E}_{P_0}[Z | X = x]$ and an artificial summary $\Psi(\theta_0) = P_0(f \circ \theta_0)$ where f is an element of $C^1[-1,1]$ but not of $C^2[-1,1]$. In this case, $\psi = f'$ is very rough, so we do not expect it to be approximated by a trigonometric series as well as the identity function. However, it is sufficiently smooth to allow for the existence of a deterministic K that achieves efficiency. Further simulation details are provided in Appendix E.

Table 3 presents the performance of our estimator based on 10-fold CV. We note that it performs reasonably well in terms of $n \cdot \text{MSE}$ criterion. However, it is unclear whether its scaled bias converges to zero for large n , so our method may be too biased. The coverage of 95% Wald CIs is close to the nominal level, suggesting that the bias is fairly small relative to the standard error of the estimator at the sample sizes considered. One possible explanation for the good performance observed is that the $L^2(P_0)$ -convergence rate of θ_n^* is much faster than $n^{-1/14}$, which allows for a slower convergence rate of the approximation error $\|\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0\|$ (see Appendix C). This simulation shows that our proposed method may still perform well even if Condition C5 is violated, especially when the initial ML fit is close to the unknown function.

5. Generalized data-adaptive series

5.1. Proposed method

As in Section 4, we consider the case that Ψ is scalar-valued in this section. The assumption that $\dot{\Psi} = \dot{\psi} \circ \theta_0$ may be too restrictive for general summaries as in Examples 2–4, especially if $\dot{\Psi}$ is not derived analytically (see Remark 2). In this section, we generalize the method in Section 4 to deal with these summaries. Letting $\mathcal{I}_x$ be the identity function defined on $\mathcal{X}$, we

can readily generalize the above method to the case where $\dot{\Psi}$ can be represented as $\dot{\psi} \circ (\theta_0, \mathcal{F}_x)$ for a function $\dot{\psi}: \mathbb{R}^q \times \mathcal{X} \rightarrow \mathbb{R}^q$; that is, $\dot{\Psi}(x) = \dot{\psi}(\theta_0(x), x)$. This form holds trivially if we set $\dot{\psi}(t, x) = \dot{\Psi}(x)$, i.e., $\dot{\psi}$ is independent of its first argument, but we can utilize flexible ML methods if $\dot{\psi}$ is nontrivial. Again, we assume Θ is a vector space of $\mathbb{R}^q$ -valued function equipped with the $L^2(P_0)$ -inner product. We assume $\dot{\psi}$ can be approximated well by a basis $\phi_1, \phi_2, \dots: \mathbb{R}^q \times \mathcal{X} \rightarrow \mathbb{R}^q$, and consider the data-adaptive sieve-like subspace $\Theta_n := \Theta_{n, \theta_n^0} := \text{Span}\{\phi_1, \dots, \phi_K\} \circ (\theta_n^0, \mathcal{F}_x)$. We propose to use $\Psi(\theta_n^*)$ to estimate $\Psi(\theta_0)$, where $\theta_n^* = \theta_n^{*(\theta_n^0)} \in \arg\min_{\theta \in \Theta_n} n^{-1} \sum_{i=1}^n \ell(\theta)(V_i)$ denotes the series estimator within Θ_n minimizing the empirical risk.

5.2. Results for proposed method

With a slight abuse of notation, in this section we use $\mathcal{F}$ to denote the function $(t, x) \mapsto t$ where $t \in \mathbb{R}^q$ and $x \in \mathcal{X}$. Again, we use $\pi_n := \pi_{n, \theta_n^0}$ to denote the projection operator onto Θ_{n, θ_n^0} .

Let $\Pi_{n, \theta}$ be defined such that, for any function $g: \mathbb{R}^q \times \mathcal{X} \rightarrow \mathbb{R}^q$ with $g \circ (\theta, \mathcal{F}_x) \in L^2(P_0)$, it holds that $\Pi_{n, \theta}(g) \circ (\theta, \mathcal{F}_x) = \pi_{n, \theta}(g \circ (\theta, \mathcal{F}_x))$; that is, letting β_j be the quantity that depends on g and θ such that $\pi_{n, \theta}(g \circ (\theta, \mathcal{F}_x)) = \left(\sum_{j=1}^K \beta_j \phi_j\right) \circ (\theta, \mathcal{F}_x)$, we define $\Pi_{n, \theta}(g) := \sum_{j=1}^K \beta_j \phi_j$.

We introduce conditions and derive theoretical results that are parallel to those in Section 4.

Condition C3* (Sufficiently small approximation error to $\mathcal{F}$ for Θ_{n, θ_0}).

$$\theta_0 - \Pi_{n, \theta_0}(\mathcal{F}) \circ (\theta_0, \mathcal{F}_x) = o_p(n^{-1/4}).$$

Condition C4* (Sufficiently small approximation error to $\dot{\psi}$ for Θ_{n, θ_0} and convergence rate of θ_n^*).

$\left[\dot{\psi} - \Pi_{n, \theta_0}(\dot{\psi})\right] \circ (\theta_0, \mathcal{F}_x) \cdot \theta_n^* - \theta_0 = o_p(n^{-1/2})$ Additional regularity conditions can be found in Appendix B.3. Note that $\dot{\Psi}$ may depend on components of P_0 other than θ_0 , Condition C4* may impose smoothness conditions on these components so that $\dot{\psi}$ can be well approximated by the chosen series. For example, in Example 2, Condition C4* requires that p'_0/p_0 and the propensity score can be approximated by the series well; in Examples 3 and 4, Condition C4* imposes the same requirement on the propensity score. We now present a theorem that establishes the efficiency of the plug-in estimator based on θ_n^* .

Theorem 5 (Efficiency of plug-in estimator). *Under Conditions A1–A4, C1, C2, C3*, C4*, C6*, C7*, C8 and C9, $\Psi(\theta_n^*)$ is an asymptotically linear estimator of $\Psi(\theta_0)$ with influence function $v \mapsto \alpha_{0, \ell}^{-1} \left\{ -\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](v) + \mathbb{E}_{P_0}[\ell'_0[\dot{\psi} \circ (\theta_0, \mathcal{F}_x)](V)] \right\}$, that is,*

$$\Psi(\theta_n^*) = \Psi(\theta_0) + \frac{1}{n} \sum_{i=1}^n \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](V_i) + \mathbb{E}_{P_0}[\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](V)] \right\} + o_p(n^{-1/2}).$$

As a consequence, $\sqrt{n}[\Psi(\theta_n^*) - \Psi(\theta_0)] \xrightarrow{d} N(0, \xi^2)$ with $\xi^2 := \text{Var}_{P_0}(\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](V)) / \alpha_{0,\ell}^2$. In addition, under Conditions E1 and E2 in Appendix B.4, $\Psi(\hat{\theta}_n)$ is efficient under nonparametric models.

Remark 4. When Ψ depends on both θ_0 and P_0 , we can readily adapt this method as in Section 4.3.

We now present a condition for selecting K via k -fold CV, in parallel with Condition C5 from Section 4.4.

Condition C5* (Bounded approximation error of ψ relative to $\mathcal{F}$). *There exists a constant $C > 0$ such that, with probability tending to one,*

$$\psi \circ (\theta_n^0, \mathcal{F}_x) - \Pi_{K, \theta_n^0(\psi)}(\psi) \circ (\theta_n^0, \mathcal{F}_x) \leq C \theta_n^0 - \Pi_{K, \theta_n^0(\mathcal{F})}(\theta_n^0) \circ (\theta_n^0, \mathcal{F}_x).$$

Remark 5. Similarly to Condition C5, Condition C5* requires that the identity function $\mathcal{F}$ is not contained in the span of finitely many terms of the chosen series and that ψ is sufficiently smooth so that ψ can be approximated well by the chosen series. However, Condition C5* may be far more stringent than Condition C5. In fact, it may be overly stringent in practice. Since Ψ may depend on components of P_0 other than θ_0 , Condition C5* may require these components to be sufficiently smooth. When a common candidate series such as the trigonometric series is used, a sufficient condition for Condition C5* is that ψ is infinitely differentiable with bounded derivatives, which further imposes assumptions on the smoothness of other components of P_0 . For example, in Example 2, a sufficient condition for Condition C5* is that p_0'/p_0 is infinitely differentiable with bounded derivatives; in Examples 3 and 4, a sufficient condition for Condition C5* is that Condition C5* is that the propensity score function satisfies the same requirement. Due to the stringency of Condition C5*, we conduct a simulation in Section 5.3.2 to understand the performance of our proposed method when this condition is violated. The simulation appears to indicate that our method may be robust against violation of Condition C5*.

The following theorem shows that k -fold CV can be used to select K under certain conditions.

Theorem 6 (Efficiency of CV-based plug-in estimator). *Assume Conditions A1–A4, C1, C2, C3*, C6*, C7*, C8 and C9 hold for a deterministic $K = K(n)$. Suppose that part (a) of Condition C7* holds. With $\theta_n^\# := \theta_{K^*}^0(\theta_n^0)$, $\Psi(\theta_n^\#)$ is an asymptotically linear estimator of $\Psi(\theta_0)$*

with influence function $v \mapsto \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](v) + \mathbb{E}_{P_0}[\ell'_0[\psi \circ (\theta_0, \mathcal{F}_x)](V)] \right\}$, that is,

$$\Psi(\theta_n^\#) = \Psi(\theta_0) + \frac{1}{n} \sum_{i=1}^n \alpha_{0,\ell}^{-1} \left\{ -\ell'_0[\psi \circ (\theta_0 \cdot \mathcal{F}_x)](V_i) + \mathbb{E}_{P_0}[\ell'_0[\psi \circ (\theta_0 \cdot \mathcal{F}_x)](V)] \right\} + o_p(n^{-1/2})$$

Therefore, $\sqrt{n}[\Psi(\theta_n^\#) - \Psi(\theta_0)] \xrightarrow{d} N(0, \xi^2)$ with $\xi^2 := \text{Var}_{P_0}(\ell'_0[\psi \circ \theta_0](V)) / \alpha_{0,\ell}^2$. In addition, under Conditions E1 and E2 in Appendix B.4, $\Psi(\hat{\theta}_n)$ is efficient under a nonparametric model.

5.3. Simulation

In the following simulations, we consider the problem in Example 4. As we show in Appendix A, letting $g_0 : x \mapsto P_0(A = 1 | X = x)$ be the propensity score and setting $\theta = (\mu_0, \mu_1)$, with $\ell(\theta) : v \mapsto a[z - \mu_1(x)]^2 + (1-a)[z - \mu_0(x)]^2$, the generalized data-adaptive series methodology may be used to obtain an efficient estimator. As in Section 4.5, we conduct two simulation studies, the first demonstrating Theorem 6 and the other exploring the robustness of CV against violation of Condition C5*.

5.3.1. Demonstration of Theorem 6—We choose θ_0 to be a discontinuous function while g_0 is highly smooth. We compare the performance of plug-in estimators based on three different nonparametric regressions: (i) polynomial regression with the degree selected by 5-fold CV (poly), which results in a traditional sieve estimator, (ii) gradient boosting (xgb) [7, 8, 15, 16], and (iii) generalized data-adaptive trigonometric series estimation with gradient boosting as the initial ML fit and 5-fold CV to select the number of terms in the series (xgb.trig). Further details of the simulation setting are provided in Appendix E.

Fig 6 presents $n \cdot \text{MSE}$ and $\sqrt{n} \cdot |\text{bias}|$ for each estimator, whereas Table 4 presents the coverage probability of 95% Wald CIs based on these estimators. There are a few runs in the simulation with noticeably poor behavior, so we trimmed the most extreme values when computing MSE and bias in Fig 6 (1% of all Monte Carlo runs). The outliers may be caused by the performance of gradient boosting and the instability of 5-fold CV. In practice, the user may ensemble more ML methods and use 10-fold CV to mitigate such behavior. We note that xgb.trig and xgb.1step estimators perform well, while poly and xgb plug-in estimators do not appear to be efficient. Based on gradient boosting, our estimator and the one-step corrected estimator both appear to be efficient, but the construction of our estimator has the advantage of not requiring the analytic expression of an influence function.

5.3.2. Violation of Condition C5*—We also study via simulation the behavior of our estimator when Condition C5* is violated. We note that whether Condition C5* holds depends on the smoothness of g_0 . We choose g_0 to be rougher than $\mathcal{F}$ with g_0 being an element of $\mathcal{C}^2[-1, 1]$ but not of $\mathcal{C}^3[-1, 1]$. Consequently, ψ cannot be approximated by our generalized data-adaptive series as well as $\mathcal{F}$, but its smoothness is sufficient for the existence of a deterministic K to achieve efficiency. Appendix E describes further details of this simulation setting.

Table 5 presents the performance of our estimator based on 5-fold CV. We observe that its scaled MSE appears to converge to one, but it is unclear whether its scaled bias converges to zero for large n , and so our method may be overly biased.. The coverage of 95% Wald CIs is close to the nominal level, suggesting that the bias may be fairly small relative to the standard error of the estimator at the sample sizes considered. Therefore, according to this simulation, our generalized data-adaptive series methodology appears to be robust against violation of Condition C5*.

6. Discussion

Numerous methods have been proposed to construct efficient estimators for statistical parameters under a nonparametric model, but each of them has one or more of the following undesirable limitations: (i) their construction may require specialized expertise that is not accessible to most statisticians; (ii) for any given data set, there may be little guidance, if any, on how to select a key tuning parameter; and (iii) they may require stringent smoothness conditions, especially on derivatives. In this paper, we propose two sieve-like methods that can partially overcome these difficulties.

Our first approach, namely that based on HAL, can be further generalized to the case in which the flexible fit is an empirical risk minimizer over a function class assumed to contain the unknown function. The key condition B2 may be modified in that case as long as it ensures that certain perturbations of the unknown function still lie in that function class. We note that our methods may also be applied under semiparametric models.

A major direction for future work is to construct valid CIs without the knowledge of the influence function of the resulting plug-in estimator. The nonparametric bootstrap is in general invalid when the overall summary is not Hadamard differentiable and especially when the method relies on CV [2, 10], but a model-based bootstrap is a possible solution (Chapter 28 of [31]). In many cases only certain components of the true data-generating distribution must be estimated to obtain a plug-in estimator, while its variance may depend on other components that are not explicitly estimated. Therefore, generating valid model-based bootstrap samples is generally difficult.

Our proposed sieve-like methods may be used to construct efficient plug-in estimators for new applications in which the relevant theoretical results are difficult to derive. They may also inspire new methods to construct such estimators under weaker conditions.

Acknowledgements

This work was partially supported by the National Institutes of Health under award number DP2-LM013340 and R01HL137808. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Appendix A:: Modification of chosen norm for evaluating the conditions: case study of mean counterfactual outcome

In this appendix, we consider a parameter that requires a modification in the chosen norm for evaluating the conditions. In particular, we discuss estimating counterfactual mean outcome in Example 3.

Let $g_0 : x \mapsto P_0(A = 1|X = x)$ be the propensity score function. A natural choice of the loss function is $\ell(\theta) : v \mapsto a[y - \theta(x)]^2$. Indeed, learning a function with this loss function is equivalent to fitting a function within the stratum of observations that received treatment 1. Unfortunately, this loss function does not satisfy Condition A2 with $L^2(P_0)$ -norm, because $P_0\{\ell(\theta) - \ell(\theta_0)\} = P_0\{g_0 \cdot (\theta - \theta_0)^2\}$ cannot be well approximated by $\alpha_0 P_0\{(\theta - \theta_0)^2\}/2$ for any constant $\alpha_0, \rho > 0$ unless g_0 is a constant. One way to overcome this challenge is to choose the alternative inner product $\langle \theta_1, \theta_2 \rangle_{g_0} := P_0\{g_0 \theta_1 \theta_2\}$ and its induced norm $\|\cdot\|_{g_0}$.

In this case, Condition A2 is satisfied once k is replaced by $\|\cdot\|_{g_0}$ in the condition statement. Under this choice, $\Psi'_{\theta_0} = P_0(\theta - \theta_0) = \langle 1/g_0, \theta - \theta_0 \rangle_{g_0}$.

We may redefine the corresponding $\dot{\Psi}$ similarly as the function that satisfies

$$\Psi'_{\theta_0} = \langle \dot{\Psi}, \theta - \theta_0 \rangle_{g_0},$$

and it immediately follows that $\dot{\Psi} = 1/g_0$. Moreover, under a strong positivity condition, namely $g_0(X) \geq \delta_g > 0$ a.s. for some δ_g , which is a typical condition in causal inference literature [31, 34], then it is straightforward to show that $\delta_g \|\cdot\| \leq \|\cdot\|_{g_0} \leq \|\cdot\|$; that is, $\|\cdot\|_{g_0}$ is equivalent to $L^2(P_0)$ -norm. Using this fact, it can be shown that all other conditions with respect to the $L^2(P_0)$ -inner product are equivalent to the corresponding conditions with respect to $\langle \cdot, \cdot \rangle_{g_0}$.

Therefore, the data-adaptive series can be applied to estimation of the counterfactual mean outcome under our conditions for $L^2(P_0)$ -inner product. If we use the targeted form in Remark 2, then we need a flexible estimator of g_0 and the procedure is almost identical to a TMLE [31]. If we use the generalized data-adaptive series, we would require sufficient amount of smoothness for $g_0(\cdot)$. In the latter case, the change in norm when evaluating the conditions is a purely technical device and the estimation procedure is the same as would have been used if we had used the $L^2(P_0)$ -norm. We also note that the same argument may be used to show that in Example 4, with $\ell(\theta) : v \mapsto a[z - \mu_1(x)]^2 + (1 - a)[z - \mu_0(x)]^2$ being the usual squared-error loss, we may choose the alternative inner product

$$\langle \theta_1, \theta_2 \rangle_{g_0} := P_0\{\theta_1^\top \cdot \text{diag}(1 - g_0, g_0) \cdot \theta_2\} \text{ and find that ,}$$

$\dot{\Psi} = (-2/(1 - g_0) \cdot [(\mu_{01} - \mu_{00}) - P_0(\mu_{01} - \mu_{00})], 2/g_0 \cdot [(\mu_{01} - \mu_{00}) - P_0(\mu_{01} - \mu_{00})])^\top$ as we did in Section 5.3.

Appendix B:: Additional conditions

Throughout the rest of this appendix, we use C to denote a general absolute positive constant that can vary line by line.

B.1. HAL

Condition B3 (Empirical processes conditions). For any fixed $\vartheta \in \Theta_{v,M}$ and some $\delta > 0$, it holds that $\ell(\theta)$, $\ell'_0[\theta - \theta_0]$ and $\{r[\theta - \theta_0] - r[\theta + \delta(\vartheta - \theta) - \theta_0]\}/\delta$ are càdlàg for all $\theta \in \Theta_{v,M}$ and all $\delta \in [0, \Delta]$. Moreover, the following terms are all finite:

$$\sup_{\theta \in \Theta_{v,M}} \left\| \ell(\theta) \right\|_v, \sup_{\theta \in \Theta_{v,M}} \left\| \ell'_0[\theta - \theta_0] \right\|_v, \sup_{\theta \in \Theta_{v,M}, \delta \in [0, \Delta]} \left\| \frac{r[\theta - \theta_0] - r[\theta - \theta_0 + \delta(\vartheta - \theta)]}{\delta} \right\|_v.$$

In addition, $\left\| \ell'_0[\hat{\theta}_n - \theta_0] \right\|$ and $\sup_{\delta \in [0, \Delta]} \left\| \{r[\hat{\theta}_n - \theta_0] - r[\hat{\theta}_n - \theta_0 + \delta(\vartheta - \hat{\theta}_n)]\}/\delta \right\|$ converge to 0 in probability.

Condition B4 (Finite variance of influence function) $\xi^2 := \text{Var}_{P_0}(\ell'_0[\dot{\Psi}](V))/\alpha_{0,\ell}^2 < \infty$.

B.2. Data-adaptive series

Condition C6 (Local Lipschitz continuity of $\Pi_{n,\theta_0}(\mathcal{J})$). For sufficiently large n ,

$$\left\| \Pi_{n,\theta_0}(\mathcal{J}) \circ \theta - \Pi_{n,\theta_0}(\mathcal{J}) \circ \theta_0 \right\| \leq C \left\| \theta - \theta_0 \right\|$$

for all $\theta \in \Theta$ with $\left\| \theta - \theta_0 \right\| \leq n^{-1/4}$.

Condition C7 (Local Lipschitz continuity of ψ and $\Pi_{n,\theta_0}(\psi)$). For sufficiently large n , for all $\theta \in \Theta$ with $\left\| \theta - \theta_0 \right\| \leq n^{-1/4}$,

- a. $\left\| \psi \circ \theta - \psi \circ \theta_0 \right\| \leq C \left\| \theta - \theta_0 \right\|;$
- b. $\left\| \Pi_{n,\theta_0}(\psi) \circ \theta - \Pi_{n,\theta_0}(\psi) \circ \theta_0 \right\| \leq C \left\| \theta - \theta_0 \right\|.$

Condition C8 (Empirical process conditions). There exists some constant $\delta > 0$ such that

$$\sup_{\delta \in [0, \Delta]} \left\| (P_n - P_0) \left\{ \frac{r[\theta_n^* - \theta_0] - r[\pi_n((1 - \delta)\theta_n^* + \delta(\pm \dot{\Psi} + \theta_0)) - \theta_0]}{\delta} \right\} \right\| = o_p(n^{-1/2}),$$

$$(P_n - P_0) \ell'_0[(\pm \dot{\Psi} + \theta_0) - \pi_n(\pm \dot{\Psi} + \theta_0)] = o_p(n^{-1/2}).$$

$$(P_n - P_0)\ell'_0[\theta_n^* - \theta_0] = o_p(n^{-1/2}).$$

Condition C9 (Finite variance of influence function) $\xi^2 := \text{Var}_{P_0}(\ell'_0[\dot{\Psi}](V))/\alpha_{0,\ell}^2 < \infty$.

B.3. Generalized data-adaptive series

Condition C6* (Local Lipschitz continuity of projected $\mathcal{J}$ for Θ_{n,θ_0}). For sufficiently large

$$n, \left\| \Pi_{n,\theta_0}(\mathcal{J}) \circ (\theta, \mathcal{J}_x) - \Pi_{n,\theta_0}(\mathcal{J}) \circ (\theta_0, \mathcal{J}_x) \right\| \leq C \|\theta - \theta_0\| \text{ for all } \|\theta - \theta_0\| \leq n^{-1/4}.$$

Condition C7* (Local Lipschitz continuity of $\dot{\psi}$ and its projection for Θ_{n,θ_0}). For sufficiently large n , for all $\|\theta - \theta_0\| \leq n^{-1/4}$,

- a. $\left\| \dot{\psi} \circ (\theta, \mathcal{J}_x) - \dot{\psi} \circ (\theta_0, \mathcal{J}_x) \right\| \leq C \|\theta - \theta_0\|;$
- b. $\left\| \Pi_{n,\theta_0}(\dot{\psi}) \circ (\theta, \mathcal{J}_x) - \Pi_{n,\theta_0}(\dot{\psi}) \circ (\theta_0, \mathcal{J}_x) \right\| \leq C \|\theta - \theta_0\|.$

B.4. Conditions for efficiency of the plug-in estimator

Define a collection of submodels

$$\left\{ \left\{ P_{H,\delta} : \delta \in B_H \subseteq \mathbb{R} \right\} : H \in \mathcal{H} \right\}$$

for which: (i) $\mathcal{H}$ is a subset of $L_0^2(P_0)$ and the $L_0^2(P_0)$ of its linear span is $L_0^2(P_0)$; and (ii) each $\{P_{H,\delta} : \delta \in B_H \subseteq \mathbb{R}\}$ is a regular univariate parametric submodel that passes through P_0 and has score H for δ at $\delta = 0$. For each $H \in \mathcal{H}$ and $\delta \in B_H$, we define $\theta_{H,\delta} \in \arg\min_{\theta \in \Theta} P_{H,\delta} \ell(\theta)$. In this appendix, for all small o and big O notations, we let $\delta \rightarrow 0$ with H fixed.

Condition E1 (Sufficiently close risk minimizer). For any given $H \in \mathcal{H}$.

$$\|\theta_{H,\delta} - \theta_0\| = o(\delta^{1/2}).$$

Condition E2 (Quadratic behavior of loss function remainder near 0). For any given $H \in \mathcal{H}$ and ϑ , there exists positive $\delta' = \alpha(\delta)$ such that

$$(P_{H,\delta} - P_0) \left\{ r[(1 - \delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}] - r[\theta_{H,\delta} - \theta_0] \right\} / \delta' = o(\delta).$$

Appendix C:: Discussion of technical conditions for data-adaptive series and its generalization

C.1. Theorem 2

Condition C2 usually imposes an upper bound on the growth rate of K . To see this, we show that Condition C2 is equivalent to a term being $o_p(n^{-1/4})$, and an upper bound of this term is controlled by K . Let $\theta_n^* \in \operatorname{argmin}_{\theta \in \Theta_n} P_0 \ell(\theta)$ be the true-risk minimizer in Θ_n . Under Conditions A2, C1, C3 and C6, by Lemma 5, it follows that Condition C2 is equivalent to requiring that $\|\theta_n^* - \theta_n^\dagger\| = o_p(n^{-1/4})$. Note that θ_n^* minimizes the empirical risk in Θ_n and M-estimation theory [32] can show that $\|\theta_n^* - \theta_n^\dagger\|$ can be upper bounded by an empirical process term, whose upper bound is related to the complexity of Θ_n namely how fast K grows with sample size. To ensure this bound is $o_p(n^{-1/4})$, K must not grow too quickly.

Condition C3 assumes that the identity function can be well approximated by the series ϕ_k with the specified number of terms K in the $L^2(P_{\theta_0})$ sense. If $\operatorname{Span}\{\phi_1, \dots, \phi_K\}$ does not contain $\mathcal{F}$ for any K , then sufficiently many terms must be included to satisfy this condition; that is, this condition imposes a lower bound on the rate at which K should grow with n . Even if $\operatorname{Span}\{\phi_1, \dots, \phi_K\}$ does contain $\mathcal{F}$ for some finite K , this condition still requires that K is not too small.

Condition C4 is implied by the following condition in view of Lemma 3:

$$\textbf{Condition C4s.} \quad \left\| \left[\dot{\psi} - \Pi_{n, \theta_0}(\dot{\psi}) \right] \circ \theta_0 \right\| = o(n^{-1/4}).$$

This condition is similar to Condition C3. However, in general, we do not expect $\dot{\psi}$ to be contained in $\operatorname{Span}\{\phi_1, \dots, \phi_K\}$ for any K , and hence this condition generally imposes a lower bound on the rate of K . Note that Condition C4s is stronger than Condition C4, and there are interesting examples where C4 holds but C4s fails to hold. Indeed, if θ_n^* converges to θ_0 at a rate much faster than $n^{-1/4}$, then C4 can be satisfied even if $\left\| \left[\dot{\psi} - \Pi_{n, \theta_0}(\dot{\psi}) \right] \circ \theta_0 \right\|$ decays to zero in probability relatively slowly — that is, the convergence rate of θ_n^* can compensate for the approximation error of $\dot{\psi}$. This is one way in which we can benefit from using flexible ML algorithms to estimate θ_0 : if θ_n^0 converges to θ_0 at a fast rate, then we can expect θ_n^* to also have a fast convergence rate.

Conditions C2, C3 and C4 are not stringent provided sufficient smoothness on derivatives of $\dot{\psi}$ and a reasonable series. For example, as noted in [4], when $\dot{\psi}$ has a bounded p -th order derivative and the polynomial, trigonometric series or spline with degree at least $p+1$ is used, then if $K^2/n \rightarrow 0$ ($K^3/n \rightarrow 0$ for polynomial series), the term in Condition C2 is $O_p(\sqrt{K/n})$; the terms in Condition C3 and the sufficient Condition C4s are $O(K^{-p/q})$.

Therefore, we can select K to grow at a rate faster than $n^{q/(4p)}$ and slower than $n^{1/2}$ ($n^{1/3}$ for polynomial series). If p is large, then this allows for a wide range of rates for K . Typically $\dot{\Psi}$ (and hence $\dot{\psi}$) is only related to the summary of interest Ψ but not the true function θ_0 . For example, for the summary $\Psi(\theta) = P_0(f \circ \theta)$ at the beginning of Section 4.1, $\dot{\psi} = f'$ is variation independent of θ_0 . It is often the case that Ψ is smooth and so is $\dot{\psi}$, so p is often sufficiently large for this window to be wide.

Condition C6 is usually easy to satisfy. Since $\Pi_{n, \theta_0}(\mathcal{J})$ is a linear combination of $\{\phi_k : k \in \{1, \dots, K\}\}$ and is an approximation of a highly smooth function $\mathcal{J}$, if the series ϕ_k is smooth, then we can expect that $\Pi_{n, \theta_0}(\mathcal{J})$ will be Lipschitz uniformly over n , that is, that Condition C6 holds. For example, using polynomial series, cubic splines or trigonometric series imply that this condition holds.

Condition C7 imposes Lipschitz continuity conditions on $\dot{\psi}$ and $\Pi_{n, \theta_0}(\dot{\psi})$ uniformly over n . The Lipschitz continuity of $\dot{\psi}$ has been discussed above. As for $\Pi_{n, \theta_0}(\dot{\psi})$, similarly to Condition C6, as long as the series ϕ_k being used is smooth, $\Pi_{n, \theta_0}(\dot{\psi})$ would be Lipschitz continuous uniformly over n .

C.2. Theorem 5

The conditions are similar to those in Theorem 2. However, Condition C4* can be more stringent than Condition C4. For generalized data-adaptive series, the dimension of the argument of the series is larger. Hence, as noted in [4], C4* may require more smoothness of $\dot{\psi}$ in order that $\dot{\psi}$ can be well approximated by $\Pi_{n, \theta_0}(\dot{\psi})$. However, in general, we do not expect the smoothness of $\dot{\psi}$ to depend on Ψ alone but no components of P_0 , so the amount of smoothness of $\dot{\psi}$ may be more limited in practice.

It is also worth noting that, similarly to Theorem 2, a sufficient condition for Condition C4* is the following:

$$\text{Condition C4*s. } \left[\dot{\psi} - \Pi_{n, \theta_0}(\dot{\psi}) \right] \circ (\theta_0, \mathcal{J}_x) = o(n^{-1/4}).$$

Appendix D:: Lemmas and technical proofs

D.1. Highly Adaptive Lasso (HAL)

Proof of Theorem 1. Under Conditions A2 and B1–B3, Lemma 1 and its corollary in [26] show that $\|\hat{\theta}_n - \theta_0\| = o_p(n^{-1/4})$.

We show that the small perturbations of $\hat{\theta}_n$ in certain directions are contained in $\Theta_{v, M}$. Let $\vartheta_\delta = \hat{\theta}_n + \delta(\dot{\Psi} + \theta_0 - \hat{\theta}_n)$ be a path indexed by δ ($0 < \delta < 1$) that is a perturbation of $\hat{\theta}_n$. Note that for all δ , ϑ_δ is càdlàg by Condition B1 and we have that

$$\|\theta_\delta\|_v = \|(1-\delta)\hat{\theta}_n + \delta(\dot{\Psi} + \theta_0)\|_v \leq (1-\delta)\|\hat{\theta}_n\|_v + \delta(\|\dot{\Psi}\|_v + \|\theta_0\|_v) \leq (1-\delta)M + \delta M = M$$

by Condition B2. Hence $\theta_\delta \in \Theta_{v,M}$. The same result holds for the path

$$\tilde{\theta}_\delta := \hat{\theta}_n + \delta(-\dot{\Psi} + \theta_0 - \hat{\theta}_n).$$

Combining this observation with the P_0 -Donkser property of $\Theta_{v,M'}$ for any fixed $M' > 0$ [9] and Conditions A1–A2, B4, we have that all of the conditions of Theorem 1 in [24] are satisfied with all sieves being $\Theta_{v,M}$. The desired asymptotic linearity result follows. The efficiency result is shown in Appendix D.3. $\square$

Proof of Lemma 1. Recall that $\mathcal{X} \subseteq \mathbb{R}^d$. Similar to $x^{(j)}$, let $x^{(u)} = \inf\{x : P_0(X \leq x) = 1\}$ where $\inf$ and $\leq$ are entrywise. To avoid clumsy notations, in this proof we drop the subscript in θ_0 and use θ instead. This should not introduce confusion because other functions (e.g., an estimator of θ_0) are not involved in the statement or proof. Using the results reviewed in Section 3.1,

$$\begin{aligned} \|\dot{\Psi}\|_v &= \left| \dot{\Psi}(x^{(\mathcal{L})}) \right| + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} |\dot{\Psi}_s(du)| \\ &= \left| \dot{\Psi}(x^{(\mathcal{L})}) \right| + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} |\dot{\Psi}'(z)| \Big|_{z=\theta_s(u)} |\theta_s(du)|. \end{aligned}$$

Since

$$\begin{aligned} |\theta(x)| &= \left| \theta(x^{(\mathcal{L})}) + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} \theta_s(du) \right| \\ &\leq \left| \theta(x^{(\mathcal{L})}) \right| + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} |\theta_s(du)| \\ &\leq \left| \theta(x^{(\mathcal{L})}) \right| + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} |\theta_s(du)| = \|\theta\|_v, \end{aligned}$$

we have $|\dot{\Psi}'(z)|_{z=\theta_s(u)} \leq \sup_{z': |z'| \leq \|\theta_0\|_v} |\dot{\Psi}'(z')| = B$ for all $x^{(\mathcal{L})} \leq u \leq x^{(u)}$, so

$$\|\dot{\Psi}\|_v \leq \left| \dot{\Psi}(x^{(\mathcal{L})}) \right| + \sum_{s \subseteq \{1, 2, \dots, d\}, s \neq \emptyset} \int_{x_s^{(\mathcal{L})}}^{x_s^{(u)}} B |\theta_s(du)| \leq \left| \dot{\Psi}(x^{(\mathcal{L})}) \right| + B \|\theta_0\|_v.$$

$\square$

Lemma 2 (CV-selected bound not much smaller than the bound of the true function's variation norm). *Suppose that Condition B1 holds, θ_0 is càdlàg, $\|\theta_0\|_v < \infty$ and for any M , $\sup_{\theta \in \Theta_{v,M}} \|\ell(\theta)\| < \infty$. Let M_n be a (possibly random) sequence such that $P_0\left\{\ell(\hat{\theta}_{n,M_n}) - \ell(\theta_0)\right\} = o_p(1)$. Then for any $\epsilon > 0$, with probability tending to one, $M_n - \|\theta_0\|_v - \epsilon$. Therefore, for any fixed $\epsilon > 0$, with probability tending to one, $M_n + \epsilon - (\|\theta_0\|_v - \epsilon) + \epsilon = \|\theta_0\|_v$*

Proof of Lemma 2. We prove by contradiction. Suppose the claim is not true, i.e. there exists $\epsilon, \delta > 0$ such that $P(M_n < \|\theta_0\|_v - \epsilon)$ for all $n \in \mathcal{N}$, where $\mathcal{N}$ is an infinite set. Let $\theta_{0,M} \in \operatorname{argmin}_{\theta \in \Theta_{v,M}} P_0 \ell(\theta)$. Then for all $n \in \mathcal{N}$, with probability at least δ ,

$$\begin{aligned} P_0\left\{\ell(\hat{\theta}_{n,M_n}) - \ell(\theta_0)\right\} &= P_0\left\{\ell(\hat{\theta}_{n,M_n}) - \ell(\theta_{0,M_n})\right\} + P_0\left\{\ell(\theta_{0,M_n}) - \ell(\theta_0)\right\} \\ &\geq P_0\left\{\ell(\theta_{0,M_n}) - \ell(\theta_0)\right\} \\ &\geq P_0\left\{\ell(\theta_0, \|\theta_0\|_v - \epsilon) - \ell(\theta_0)\right\}, \end{aligned}$$

which is a positive constant since the function class $\Theta_{\|\theta_0\|_v - \epsilon}$ does not contain θ_0 and this

term is non-negligible bias. This contradicts the assumption that

$$P_0\left\{\ell(\hat{\theta}_{n,M_n}) - \ell(\theta_0)\right\} = o_p(1) \text{ and hence the desired follows. } \square$$

Therefore, if $\|\dot{\Psi}\|_v \leq F(\|\theta_0\|_v)$ for a known increasing function F , then with probability tending to one, $M_n + \epsilon + F(M_n + \epsilon)$ is a valid bound on $\|\hat{\theta}_n\|_v$ that can be used to obtain an efficient plug-in estimator. Moreover, if the bound is loose, i.e. $\|\dot{\Psi}\|_v \leq F(\|\theta_0\|_v)$, and F is continuous, then there exists some $\epsilon > 0$ such that $\|\dot{\Psi}\|_v \leq F(\|\theta_0\|_v - \epsilon) - \epsilon$ and hence $\|\theta_0\|_v + \|\dot{\Psi}\|_v \leq M_n + F(M_n)$ with probability tending to one.

Note that this lemma only concerns learning a function-valued feature but not estimating $\Psi(\theta_0)$. There are examples where $\dot{\Psi}$ depends on components of P_0 , say η_0 , other than θ_0 . However, if η_0 can be learned via HAL, then Lemma 2 can be applied. Therefore, if it is known that $\|\dot{\Psi}\|_v \leq F(\|\theta_0\|_v, \|\eta_0\|_v)$ for a known increasing function F , then we can use a bound on $\|\hat{\theta}_n\|_v$ obtained in a similar fashion as above from the sequence M_n to construct an efficient plug-in estimator $\Psi(\hat{\theta}_n)$.

Now consider obtaining M_n by k -fold CV from a set of candidate bounds. Then, under Conditions B1–B3, by (i) Lemma 1 and its corollary of [26], and (ii) the oracle inequality

for k -fold CV in [29], $P_0\left\{\ell\left(\hat{\theta}_{n,M_n}\right)-\ell\left(\theta_0\right)\right\}=o_p\left(n^{-1/4}\right)$ if (i) one candidate bound is no smaller than $\|\theta_0\|_v$, and (ii) the number of candidate bounds is fixed. Therefore, the above results apply to this case.

D.2. Data-adaptive series estimation

We first present and prove two lemmas that lead to Theorems 2 and 5.

Lemma 3 (Convergence rate of the sieve estimator). *Under Conditions C1, C3 and C6, $\|\pi_n(\theta_0)-\theta_0\|=o_p(n^{-1/4})$. Under an additional condition C2, $\|\theta_n^*-\theta_0\|=o_p(n^{-1/4})$.*

Proof of Lemma 3. By triangle inequality,

$\|\pi_n(\theta_0)-\theta_0\|\leq\left\|\theta_0-\theta_n^0\right\|+\left\|\theta_n^0-\pi_n\left(\theta_n^0\right)\right\|+\left\|\pi_n\left(\theta_n^0\right)-\pi_n\left(\theta_0\right)\right\|$. We bound these three terms separately.

Term 1: By Condition C1, $\left\|\theta_0-\theta_n^0\right\|=o_p\left(n^{-1/4}\right)$.

Term 2: By the definition of projection operator,

$$\left\|\theta_n^0-\pi_n\left(\theta_n^0\right)\right\|=\left\|\theta_n^0-\Pi_{n,\theta_n^0(\mathcal{F})}\circ\theta_n^0\right\|\leq\left\|\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0\right\|.$$

We bound the right-hand side by showing this term is close to $\left\|\theta_0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\|$ up to an $o_p(n^{-1/4})$ term. By the reverse triangle inequality and the triangle inequality,

$$\begin{aligned} & \left|\left\|\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0\right\|-\left\|\theta_0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\|\right| \\ & \leq\left|\left\|\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0\right\|-\left\|\theta_0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\|\right| \\ & =\left|\left\|\theta_n^0-\theta_0\right\|-\left\|\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\|\right| \\ & \leq\left\|\theta_n^0-\theta_0\right\|+\left\|\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\| \\ & \leq\left\|\theta_n^0-\theta_0\right\|+C\left\|\theta_n^0-\theta_0\right\| \end{aligned}$$

which is $o_p(n^{-1/4})$ by Condition C1. Therefore, by Condition C3,

$$\left\|\theta_n^0-\pi_n\left(\theta_n^0\right)\right\|\leq\left\|\theta_n^0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_n^0\right\|\leq\left\|\theta_0-\Pi_{n,\theta_0(\mathcal{F})}\circ\theta_0\right\|+o_p\left(n^{-1/4}\right)=o_p\left(n^{-1/4}\right)$$

Term 3: By the definition of projection and Condition C1,

$$\left\|\pi_n\left(\theta_n^0\right)-\pi_n\left(\theta_0\right)\right\|\leq\left\|\theta_n^0-\theta_0\right\|=o_p\left(n^{-1/4}\right)$$

Conclusion from the three bounds: $\left\|\pi_n\left(\theta_0\right)-\theta_0\right\|=o_p\left(n^{-1/4}\right)$.

If, in addition, Condition C2 also holds, then

$$\|\theta_n^* - \theta_0\| \leq \|\pi_n(\theta_0) - \theta_0\| + \|\theta_n^* - \pi_n(\theta_0)\| = o_p(n^{-1/4}). \quad \square$$

The same result holds for the generalized data-adaptive series under Conditions C1, C6*, C3* and C2 (if relevant). The proof is almost identical and is therefore omitted.

Lemma 4 (Approximation error to ψ). *Under Condition C7,*

$$\|\psi \circ \theta_0 - \pi_n(\psi \circ \theta_0)\| \leq C \|\theta_n^0 - \theta_0\| + \|\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0\|. \text{ Therefore, under Conditions C1--C4, } \|\psi \circ \theta_0 - \pi_n(\psi \circ \theta_0)\| \cdot \|\theta_n^* - \theta_0\| = o_p(n^{-1/2}).$$

Proof of Lemma 4. By the definition of the projection operator and triangle inequality,

$$\|\psi \circ \theta_0 - \pi_n(\psi \circ \theta_0)\| \leq \|\psi \circ \theta_0 - \pi_n(\psi \circ \theta_n^0)\| \leq \|\psi \circ \theta_0 - \psi \circ \theta_n^0\| + \|\psi \circ \theta_n^0 - \pi_n(\psi \circ \theta_n^0)\|.$$

We bound the two terms on the right-hand side separately.

Term 1: By Condition C7, $\|\psi \circ \theta_0 - \psi \circ \theta_n^0\| \leq C \|\theta_0 - \theta_n^0\|$.

Term 2: This term can be bounded similarly as in Lemma 3. By the reverse triangle inequality and the triangle inequality,

$$\begin{aligned} & \left\| \psi \circ \theta_n^0 - \Pi_{n, \theta_0}(\psi) \circ \theta_n^0 - \left[\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0 \right] \right\| \\ & \leq \left\| \left[\psi \circ \theta_n^0 - \Pi_{n, \theta_0}(\psi) \circ \theta_n^0 \right] - \left[\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0 \right] \right\| \\ & = \left\| \left[\psi \circ \theta_n^0 - \psi \circ \theta_0 \right] - \left[\Pi_{n, \theta_0}(\psi) \circ \theta_n^0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0 \right] \right\| \\ & \leq \|\psi \circ \theta_n^0 - \psi \circ \theta_0\| + \|\Pi_{n, \theta_0}(\psi) \circ \theta_n^0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0\| \\ & \leq C \|\theta_n^0 - \theta_0\| + C \|\theta_n^0 - \theta_0\| \quad (\text{Condition C7}) \\ & = C \|\theta_n^0 - \theta_0\|. \end{aligned}$$

Therefore, by the definition of the projection operator and Condition C7,

$$\begin{aligned} \|\psi \circ \theta_n^0 - \pi_n(\psi \circ \theta_n^0)\| & \leq \|\psi \circ \theta_n^0 - \Pi_{n, \theta_0}(\psi) \circ \theta_n^0\| \\ & \leq \|\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0\| + C \|\theta_n^0 - \theta_0\|. \end{aligned}$$

Conclusion from the two bounds:

$$\|\psi \circ \theta_0 - \pi_n(\psi \circ \theta_0)\| \leq C \|\theta_n^0 - \theta_0\| + \|\psi \circ \theta_0 - \Pi_{n, \theta_0}(\psi) \circ \theta_0\|.$$

Under Conditions C1–C4, using Lemma 3, it follows that

$$\|\dot{\psi} \circ \theta_0 - \pi_n(\dot{\psi} \circ \theta_0)\| \cdot \|\theta_n^* - \theta_0\| = o_p(n^{-1/2}). \quad \square$$

Note that π_n is a linear operator. Lemma 3 and 4 along with other conditions essentially satisfy the assumptions in Corollary 2 in [24]. We can prove the asymptotic linearity result of Theorem 2 similarly to this result as follows.

Proof of Theorem 2. We note that

$$\begin{aligned} P_n \ell(\theta_n^*) &= P_n \ell(\theta_0) + P_0[\ell(\theta_n^*) - \ell(\theta_0)] + (P_n - P_0)[\ell(\theta_n^*) - \ell(\theta_0)] \\ &= P_n \ell(\theta_0) + P_0[\ell(\theta_n^*) - \ell(\theta_0)] + (P_n - P_0)\ell'_0[\theta_n^* - \theta_0] + (P_n - P_0)r[\theta_n^* - \theta_0]. \end{aligned}$$

Let ϵ_n be an arbitrary sequence of positive real numbers that is $\alpha(n^{-1/2})$. We may replace θ_n^* with $\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))$ in the above equation. We first consider $\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))$:

$$\begin{aligned} &P_n \ell(\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) \\ &= P_n \ell(\theta_0) + P_0[\ell(\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) - \ell(\theta_0)] \\ &\quad + (P_n - P_0)\ell'_0[\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0] \\ &\quad + (P_n - P_0)r[\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0]. \end{aligned} \quad (1)$$

Take the difference between the above two equations. By the linearity of ℓ'_0 and π_n , we have that

$$\begin{aligned} &P_n \ell(\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) - P_n \ell(\theta_n^*) \\ &= P_0[\ell(\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) - \ell(\theta_0)] - P_0[\ell(\theta_n^*) - \ell(\theta_0)] \\ &\quad + (P_n - P_0)\ell'_0[\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_n^*] + (P_n - P_0)\{r[\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0] - r[\theta_n^* - \theta_0]\} \\ &\quad \}. \end{aligned}$$

We next analyze the three lines on the right-hand side of the above equation separately.

Line 1: Under Condition A2,

$$\begin{aligned} &P_0[\ell(\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) - \ell(\theta_0)] - P_0[\ell(\theta_n^*) - \ell(\theta_0)] \\ &= \frac{\alpha_{0,\ell}}{2} \|\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0\|^2 - \frac{\alpha_{0,\ell}}{2} \|\theta_n^* - \theta_0\|^2 \\ &\quad + o_p(\|\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0\|^2 + \|\theta_n^* - \theta_0\|^2) \end{aligned}$$

We subtract and add $(1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})$ in the first term. By the fact that π_n is linear and $\pi_n(\theta_n^*) = \theta_n^*$, the display continues as

$$\begin{aligned}
&= \frac{\alpha_{0,\ell}}{2} \left\| \left\{ \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - ((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) \right\} + \left\{ (1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}) - \theta_0 \right\} \right\|^2 \\
&\quad - \frac{\alpha_{0,\ell}}{2} \left\| \theta_n^* - \theta_0 \right\|^2 + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right) \\
&= \frac{\alpha_{0,\ell}}{2} \left\| \epsilon_n \left\{ \pi_n(\theta_0 + \dot{\Psi}) - (\theta_0 + \dot{\Psi}) \right\} + (\theta_n^* - \theta_0) + \epsilon_n(\dot{\Psi} + \theta_0 - \theta_n^*) \right\|^2 - \frac{\alpha_{0,\ell}}{2} \left\| \theta_n^* - \theta_0 \right\|^2 \\
&\quad + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right) \\
&= \epsilon_n \alpha_{0,\ell} \langle \theta_n^* - \theta_0, \dot{\Psi} \rangle + \epsilon_n^2 \frac{\alpha_{0,\ell}}{2} \left\| \pi_n(\theta_0 + \dot{\Psi}) - \theta_n^* \right\|^2 + \epsilon_n \alpha_{0,\ell} \langle \pi_n(\theta_0) - \theta_0, \theta_n^* - \theta_0 \rangle \\
&\quad + \epsilon_n \alpha_{0,\ell} \langle \pi_n(\dot{\Psi}) - \dot{\Psi}, \theta_n^* - \theta_0 \rangle - \epsilon_n \alpha_{0,\ell} \left\| \theta_n^* - \theta_0 \right\|^2 + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right)
\end{aligned}$$

By Cauchy-Schwards inequality, the display continues as

$$\begin{aligned}
&\leq \epsilon_n \alpha_{0,\ell} \langle \theta_n^* - \theta_0, \dot{\Psi} \rangle + \epsilon_n^2 \frac{\alpha_{0,\ell}}{2} \left\| \pi_n(\theta_0 + \dot{\Psi}) - \theta_n^* \right\|^2 + \epsilon_n \alpha_{0,\ell} \left\| \pi_n(\theta_0) - \theta_0 \right\| \left\| \theta_n^* - \theta_0 \right\| \\
&\quad + \epsilon_n \alpha_{0,\ell} \left\| \pi_n(\dot{\Psi}) - \dot{\Psi} \right\| \left\| \theta_n^* - \theta_0 \right\| - \epsilon_n \alpha_{0,\ell} \left\| \theta_n^* - \theta_0 \right\|^2 + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right)
\end{aligned}$$

By Lemmas 3–4 and the assumption that $\epsilon_n = \alpha(n^{-1/2})$, the display continues as

$$= \epsilon_n \alpha_{0,\ell} \langle \theta_n^* - \theta_0, \dot{\Psi} \rangle + \epsilon_n o_p(n^{-1/2}) + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right).$$

Line 2: We subtract and add $(1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})$. By linearity of ℓ'_0 , Condition C8, and the fact that $\pi_n(\theta_n^*) = \theta_n^*$, we have that

$$\begin{aligned}
&(P_n - P_0)\ell'_0[\pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_n^*] \\
&= (P_n - P_0)\ell'_0[(1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}) - \theta_n^*] + \epsilon_n(P_n - P_0)\ell'_0[\pi_n(\theta_0 + \dot{\Psi}) - (\theta_0 + \dot{\Psi})] \\
&= \epsilon_n(P_n - P_0)\ell'_0[\dot{\Psi}] - \epsilon_n(P_n - P_0)\ell'_0[\theta_n^* - \theta_0] + \epsilon_n o_p(n^{-1/2}) \\
&= \epsilon_n(P_n - P_0)\ell'_0[\dot{\Psi}] + \epsilon_n o_p(n^{-1/2}).
\end{aligned}$$

Line 3: By Condition C8, this term is $\epsilon_n o_p(n^{-1/2})$.

Conclusion of the three lines: It holds that

$$\begin{aligned}
&P_n \ell(\pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))) - P_n \ell(\theta_n^*) \\
&\leq \epsilon_n \alpha_{0,\ell} \langle \theta_n^* - \theta_0, \dot{\Psi} \rangle + \epsilon_n(P_n - P_0)\ell'_0[\dot{\Psi}] + \epsilon_n o_p(n^{-1/2}) \\
&\quad + o_p \left(\left\| \pi_n((1-\epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi})) - \theta_0 \right\|^2 + \left\| \theta_n^* - \theta_0 \right\|^2 \right).
\end{aligned}$$

Since θ_n^* is an empirical risk minimizer, the left-hand side is non-negative. Thus,

$$0 \leq \theta_n^* - \theta_0, \dot{\Psi} + (P_n - P_0) \alpha_{0, \ell'0}^{-1} [\dot{\Psi}] + o_p(n^{-1/2}).$$

Similarly, by replacing $\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))$ with $\pi_n((1 - \epsilon_n)\theta_n^* + \epsilon_n(\theta_0 + \dot{\Psi}))$ in (1), we derive that

$$0 \leq -\langle \theta_n^* - \theta_0, \dot{\Psi} \rangle - (P_n - P_0) \alpha_{0, \ell'0}^{-1} [\dot{\Psi}] + o_p(n^{-1/2}).$$

Therefore, $|\langle \theta_n^* - \theta_0, \dot{\Psi} \rangle - (P_n - P_0) \alpha_{0, \ell'0}^{-1} [\dot{\Psi}]|$. By Conditions A3–A4 and Lemma 3,

$$\begin{aligned} \Psi(\theta_n^*) - \Psi(\theta_0) &= \langle \theta_n^* - \theta_0, \dot{\Psi} \rangle + o_p(n^{-1/2}) \\ &= - (P_n - P_0) \alpha_{0, \ell'0}^{-1} [\dot{\Psi}] + o_p(n^{-1/2}). \end{aligned}$$

The asymptotic linearity of $\Psi(\theta_n^*)$ follows. We prove the efficiency in Appendix D.3. $\square$

The proof of Theorem 5 is almost identical.

Nest we present and prove a lemma allows us to interpret Condition C2 as an upper bound on the rate of K .

Lemma 5. *Under Conditions A2, C1, C3 (C3* resp.) and C6 (C6* resp.),*

$$\|\pi_n(\theta_0) - \theta_n^\dagger\| = o_p(n^{-1/4}).$$

Proof of Lemma 5. By definition of $\theta_n^\dagger$ and Condition A2, we have

$$\|\theta_n^\dagger - \theta_0\|^2 \leq CP_0\{\ell(\theta_n^\dagger) - \ell(\theta_0)\} \leq CP_0\{\ell(\pi_n(\theta_0)) - \ell(\theta_0)\} \leq C\|\pi_n(\theta_0) - \theta_0\|^2,$$

the right-hand side of which is $o_p(n^{-1/2})$ by Lemma 3 (or its corresponding version under Conditions C6* and C3*). Therefore, $\|\theta_n^\dagger - \theta_0\| = o_p(n^{-1/4})$ and hence

$$\|\pi_n(\theta_0) - \theta_n^\dagger\| \leq \|\pi_n(\theta_0) - \theta_0\| + \|\theta_n^\dagger - \theta_0\| = o_p(n^{-1/4}).$$

We finally prove the efficiency of the data-adaptive series estimator with K selected by CV.

Proof of Theorem 4. By Lemma 3 and Condition A2, for that existing deterministic K ,

$$P_0\{\ell(\theta_K^*(\theta_n^0)) - \ell(\theta_0)\} \leq C\|\theta_K^*(\theta_n^0) - \theta_0\|^2 = o_p(n^{-1/2}).$$

By the oracle inequality for CV in [29], $P_0\{\ell(\theta_n^\#) - \ell(\theta_0)\} = o_p(n^{-1/2})$. By Condition A2, $\|\theta_n^\# - \theta_0\|^2 \leq CP_0\{\ell(\theta_n^\#) - \ell(\theta_0)\} = o_p(n^{-1/2})$ and hence $\|\theta_n^\# - \theta_0\| = o_p(n^{-1/2})$. So with probability tending to one,

$$\begin{aligned}
& \left\| \psi \circ \theta_n^0 - \pi_{K^*, \theta_n^0}(\psi \circ \theta_n^0) \right\| = \left\| \psi \circ \theta_n^0 - \Pi_{K^*, \theta_n^0}(\psi) \circ \theta_n^0 \right\| \\
& \leq C \left\| \theta_n^0 - \Pi_{K^*, \theta_n^0}(\psi) \circ \theta_n^0 \right\| \quad (\text{Condition C5}) \\
& \leq C \left\| \theta_n^0 - \theta_n^\# \right\| \quad (\text{definition of the projection operator}) \\
& \leq C \left(\left\| \theta_n^0 - \theta_0 \right\| + \left\| \theta_n^\# - \theta_0 \right\| \right), \quad (\text{triangle inequality})
\end{aligned}$$

which is $o_p(n^{-1/4})$ by Condition C1. Hence,

$$\begin{aligned}
& \left\| \psi \circ \theta_0 - \pi_{K^*, \theta_n^0}(\psi \circ \theta_0) \right\| \leq \left\| \psi \circ \theta_0 - \pi_{K^*, \theta_n^0}(\psi \circ \theta_n^0) \right\| \\
& \leq \left\| \psi \circ \theta_0 - \psi \circ \theta_n^0 \right\| + \left\| \psi \circ \theta_n^0 - \pi_{K^*, \theta_n^0}(\psi \circ \theta_n^0) \right\| \\
& \leq C \left\| \theta_n^0 - \theta_0 \right\| + o_p(n^{-1/4}), \quad (\text{Condition C7})
\end{aligned}$$

which is $o_p(n^{-1/4})$ by Condition C1.

This bounds the approximation error $\left\| \psi \circ \theta_0 - \pi_{K^*, \theta_n^0}(\psi \circ \theta_0) \right\|$ for ψ , a result that is similar to

Lemma 4 combined with Conditions C1 and C4*s. Similarly to Theorem 2, along with other conditions, the assumptions in Corollary 2 in [24] are essentially satisfied and hence an almost identical argument shows that $\Psi(\theta_n^\#)$ is an asymptotically linear estimator of $\Psi(\theta_0)$.

We prove the efficiency in Appendix D.3. $\square$

D.3. Efficiency

Proof of efficiency of the proposed estimators. It is sufficient to show that the influence function of our proposed estimators is the canonical gradient under a nonparametric model. Let $H \in \mathcal{H}$ be fixed. In the rest of this proof, for all small o and big O notations, we let $\delta \rightarrow 0$. The proof is similar to the proof of asymptotic linearity in [24] except that the estimator of θ_0 and the empirical distribution P_n are replaced by $\theta_{H,\delta}$ and $P_{H,\delta}$ respectively.

Let δ' satisfy Condition E2. We note that

$$\begin{aligned}
P_{H,\delta} \ell(\theta_{H,\delta}) &= P_{H,\delta} \ell(\theta_0) + P_0 [\ell(\theta_{H,\delta}) - \ell(\theta_0)] + (P_{H,\delta} - P_0) [\ell(\theta_{H,\delta}) - \ell(\theta_0)] \\
&= P_{H,\delta} \ell(\theta_0) + P_0 [\ell(\theta_{H,\delta}) - \ell(\theta_0)] + (P_{H,\delta} - P_0) \ell'_0 [\theta_{H,\delta} - \theta_0] + (P_{H,\delta} - P_0) r [\theta_{H,\delta} - \theta_0].
\end{aligned}$$

We also note that $(1 - \delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi}) \in \Theta$ if $|\delta|$ is sufficiently small. Then, similarly, by replacing $\theta_{H,\delta}$ with $(1 - \delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})$ in the above equation, we have that

$$\begin{aligned}
 & P_{H,\delta} \ell((1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})) \\
 &= P_{H,\delta} \ell(\theta_0) + P_0 [\ell((1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})) - \ell(\theta_0)] \\
 &+ (P_{H,\delta} - P_0) \ell'_0 [(1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi}) - \theta_0] \\
 &+ (P_{H,\delta} - P_0) r [(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}].
 \end{aligned} \tag{2}$$

Take the difference between the above two equations. By the linearity of ℓ'_0 , we have that

$$\begin{aligned}
 & P_{H,\delta} \ell((1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})) - P_{H,\delta} \ell(\theta_{H,\delta}) \\
 &= P_0 [\ell((1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})) - \ell(\theta_0)] - P_0 [\ell(\theta_{H,\delta}) - \ell(\theta_0)] + \delta' (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi} - \theta_{H,\delta} + \theta_0] \\
 &+ (P_{H,\delta} - P_0) \{ r [(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}] - r [\theta_{H,\delta} - \theta_0] \} \\
 &= \frac{\alpha_{0,\ell}}{2} \|(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}\|^2 - \frac{\alpha_{0,\ell}}{2} \|\theta_{H,\delta} - \theta_0\|^2 \\
 &+ o(\|\theta_{H,\delta} - \theta_0\|^2 + \|(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}\|^2) \quad (\text{Condition A2}) \\
 &+ \delta' (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] - \delta' (P_{H,\delta} - P_0) \ell'_0 [\theta_{H,\delta} - \theta_0] + \delta' o(\delta) \quad (\text{Condition E2}) \\
 &= \delta' \alpha_{0,\ell} \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle - \delta' \alpha_{0,\ell} \|\theta_{H,\delta} - \theta_0\|^2 + \delta'^2 \frac{\alpha_{0,\ell}}{2} \|\theta_{H,\delta} - \theta_0 + \dot{\Psi}\|^2 \\
 &+ o(\|\theta_{H,\delta} - \theta_0\|^2 + \|(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}\|^2) + \delta' (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + \delta' o(\delta) \\
 &\leq \delta' \alpha_{0,\ell} \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + \delta'^2 \frac{\alpha_{0,\ell}}{2} \|\theta_{H,\delta} - \theta_0 + \dot{\Psi}\|^2 + o(\|\theta_{H,\delta} - \theta_0\|^2 + \|(1-\delta')(\theta_{H,\delta} - \theta_0) + \delta' \dot{\Psi}\|^2) \\
 &+ \delta' (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + \delta' o(\delta).
 \end{aligned}$$

Since the left-hand side of the above display is nonnegative, by Condition E1, we have that

$$\begin{aligned}
 0 &\leq \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + \alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + o(\delta') + o(\delta) \\
 &= \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + \alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + o(\delta).
 \end{aligned}$$

Similarly, by replacing $(1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})$ with $(1-\delta')\theta_{H,\delta} + \delta'(\theta_0 + \dot{\Psi})$ in (2), we show that $0 \leq -\langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + \alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + o(\delta)$. Therefore,

$$\left| \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + \alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] \right| = o(\delta) \text{ and}$$

$$\begin{aligned}
 \Psi(\theta_{H,\delta}) - \Psi(\theta_0) &= \langle \theta_{H,\delta} - \theta_0, \dot{\Psi} \rangle + o(\|\theta_{H,\delta} - \theta_0\|^2) \\
 &= -\alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + o(\delta) + o(\|\theta_{H,\delta} - \theta_0\|^2) \\
 &= -\alpha_{0,\ell}^{-1} (P_{H,\delta} - P_0) \ell'_0 [\dot{\Psi}] + o(\delta). \quad (\text{Condition E1})
 \end{aligned}$$

Consequently, $\lim_{\delta \rightarrow 0} [\Psi(\theta_{H,\delta}) - \Psi(\theta_0)]/\delta = P_0 \{-\alpha_{0,\ell}^{-1} \ell'_0 [\dot{\Psi}] \cdot H\}$ and hence the canonical gradient of Ψ under a nonparametric model is $\alpha_{0,\ell}^{-1} \{-\ell'_0 [\dot{\Psi}] + P_0 \ell'_0 [\dot{\Psi}]\}$. Since the influence

functions of our asymptotically linear estimators are equal to this canonical gradient, our proposed estimators are efficient under a nonparametric model. $\square$

Appendix E:: Simulation setting details

In all simulations, since $\theta_0(x) = \mathbb{E}_{P_0}[Z|X = x]$ is the conditional mean function, the loss function was chosen to be the square loss $\ell(\theta) : v \mapsto (z - \theta(x))^2$.

E.1. HAL

In the simulation, we generate data from the distribution defined by

$$X \sim N(0, 1), \theta_0(x) = \exp\left\{-\left(-1 + 2x + 2x^2\right)/2\right\}, Z|X = x \sim \text{Exponential}\left(\text{rate} = 1/\theta_0(x)\right).$$

The sample sizes being considered are 500, 1000, 2000, 5000 and 10000. For each scenario we run 1000 replicates. We chose $M_{\text{gcv+}}$ to be 3.1 times M_{cv} .

E.2. Data-adaptive series

E.2.1. Demonstration of Theorem 4

In the simulation, we generate data from the distribution defined by $X \sim \text{Unif}(-1, 1)$, $Z|X = x \sim N(\theta_0(x), 0.25^2)$ where

$$\theta_0 : x \mapsto I(-1 \leq x < -3/4) + \pi I(-3/4 \leq x < -1/2) + 10x^2 I(-1/4 \leq x < 1/4) + \sqrt{2} I(1/4 \leq x < 1/2) + \exp(-1) I(1/2 \leq x < 3/4) + \sqrt[3]{3} I(3/4 \leq x \leq 1),$$

When using the trigonometric series, we first shift and scale the initial function range to be $[-1/2, 1/2]$ and then use the basis for the interval $[-1, 1]$ (i.e. $\sin(j\pi z)$, $\cos(j\pi z)$) in sieve estimation to avoid the poor behavior of trigonometric series near the boundary. We consider sample sizes 500, 1000, 2000, 5000, 10000 and 20000. For each sample size, we run 1000 simulations.

E.2.2. Violation of Condition C5

In the simulation, we generate data from the distribution defined by $X \sim \text{Unif}(-1, 1)$, $Z|X = x \sim N(\theta_0(x), 1)$ where $\theta_0 : x \mapsto \cos(10x)$. The estimand is $\Psi(\theta_0) = P_0(f \circ \theta_0)$ where

$$f : z \mapsto \left[\frac{3}{10\pi} \cos(5\pi z) - \frac{3}{8} \right] I\left(z < -\frac{1}{2}\right) - \frac{3}{2} z^2 I\left(-\frac{1}{2} \leq z < 0\right) + 3z^2 I\left(0 \leq z < \frac{1}{2}\right) + \left[-\frac{3}{2} \exp(2 - 4z) - 3z + \frac{15}{4} \right] I\left(z \geq \frac{1}{2}\right).$$

We consider sample sizes 500, 1000, 2000, 5000, 10000 and 20000; for each sample size, we run 1000 simulations. Our goal is to explore the behavior of the plug-in estimator when f , instead of θ_0 , is rough, so we use kernel regression [17] to estimate θ_0 for convenience.

E.3. Generalized data-adaptive series

E.3.1. Demonstration of Theorem 6

In the simulation, we generate data from the distribution defined by $X \sim \text{Unif}(-1, 1)$, $A|X = x \sim \text{Bern}(\text{expit}(-x))$, $Y|A = a, X = x \sim \mathcal{N}(\mu_{0,a}(x), 0.25^2)$ where

$$\mu_{00}: x \mapsto I(-1 \leq x < -3/4) + \pi I(-3/4 \leq x < -1/2) + 10x^2 I(-1/4 \leq x < 1/4) + \sqrt{2} I(1/4 \leq x < 1/2) + \exp(-1) I(1/2 \leq x < 3/4) + \sqrt[3]{3} I(3/4 \leq x \leq 1),$$

$$\mu_{01}: x \mapsto x^2 I(x < -1/3) + \exp(x) I(-1/3 \leq x < 1/3) + I(x) 1/3$$

The series is the tensor product [4] of univariate trigonometric series in E.2.1. The sample sizes are the same as in E.2.1.

E.3.2. Violation of Condition C5*

In the simulation, we generate data from the distribution defined by $X \sim \text{Unif}(-1, 1)$, $A|X = x \sim \text{Bern}(g_0(x))$, $Y|A = a, X = x \sim \mathcal{N}(\mu_{0,a}(x), 0.25^2)$ where $\mu_{0,a}: x \mapsto \exp(-x^2 + 0.8ax + 0.5a)$ ($a \in \{0, 1\}$) and

$$g_0: x \mapsto \text{expit}\left\{\left(-\frac{5}{3}x^3 - \frac{15}{4}x^2 - \frac{5}{3}x - \frac{25}{96}\right)I\left(x \leq -\frac{1}{2}\right) + \left(\frac{5}{6}x^4 + \frac{5}{3}x^3\right)I\left(-\frac{1}{2} < x \leq 0\right) + \frac{5}{3}x^3 I\left(0 < x \leq \frac{1}{2}\right) + \left(5x^2 - \frac{15}{4}x + \frac{5}{6}\right)I\left(x > \frac{1}{2}\right)\right\}.$$

We consider sample sizes 500, 1000, 2000, 5000, 10000 and 20000; for each sample size, we run 1000 simulations. Our goal is to explore the behavior of the plug-in estimator when $\dot{\Psi}$, instead of θ_0 , is rough, so we use kernel regression [17] to estimate θ_0 for convenience.

References

- [1]. Benkeser D and van Der Laan M (2016). The Highly Adaptive Lasso Estimator. In Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on 689–696. IEEE.
- [2]. Bickel P, Götze F and vanZwet W (1997). Resampling fewer than n observations: Gains, losses, and remedies for losses. STATISTICA SINICA 7.
- [3]. Bickel PJ and Ritov Y (2003). Nonparametric estimators which can be “plugged-in”. Annals of Statistics 31 1033–1053.
- [4]. Chen X (2007). Chapter 76 Large Sample Sieve Estimation of Semi-Nonparametric Models. Handbook of Econometrics 6 5549–5632.
- [5]. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C and Newey W (2017). Double/debiased/Neyman machine learning of treatment effects. American Economic Review 107 261–265.
- [6]. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W and Robins J (2018). Double/debiased machine learning for treatment and structural parameters. Econometrics Journal 21 C1–C68.
- [7]. Friedman JH (2001). Greedy Function Approximation: A Gradient Boosting Machine Technical Report No. 5.

- [8]. Friedman JH (2002). Stochastic gradient boosting. Computational Statistics and Data Analysis 38 367–378.
- [9]. Gill RD, van der Laan MJ and Wellner JA (1993). Inefficient estimators of the bivariate survival function for three models. Rijksuniversiteit Utrecht. Mathematisch Instituut.
- [10]. Hall P (2013). The bootstrap and Edgeworth expansion. Springer Science & Business Media.
- [11]. Härdle W and Stoker TM (1989). Investigating smooth multiple regression by the method of average derivatives. Journal of the American Statistical Association 84 986–995.
- [12]. Levy J, van der Laan M, Hubbard A and Pirracchio R (2018). A Fundamental Measure of Treatment Effect Heterogeneity.
- [13]. Mallat S (2009). A Wavelet Tour of Signal Processing. Elsevier.
- [14]. Marron JS (1994). Visual understanding of higher-order kernels. Journal of Computational and Graphical Statistics 3 447–458.
- [15]. Mason L, Baxter J, Bartlett P and Frean M (1999). Boosting Algorithms as Gradient Descent in Function Space Technical Report.
- [16]. Mason L, Baxter J, Bartlett PL and Frean M (2000). Boosting Algorithms as Gradient Descent Technical Report.
- [17]. Nadaraya EA (1964). On estimating regression. Theory of Probability & Its Applications 9 141–142.
- [18]. Newey W, Hsieh F and Robins J (1998). Undersmoothing and Bias Corrected Functional Estimation. Working papers.
- [19]. Newey WK (1997). Convergence rates and asymptotic normality for series estimators. Journal of Econometrics 79 147–168.
- [20]. Newey WK, Hsieh F and Robins JM (2004). Twicing Kernels and a Small Bias Property of Semiparametric Estimators. Econometrica 72 947–962.
- [21]. Owen AB (2005). Multidimensional variation for quasi-Monte Carlo. In Contemporary Multivariate Analysis And Design Of Experiments: In Celebration of Professor Kai-Tai Fang's 65th Birthday 49–74. World Scientific.
- [22]. Pfanzagl J (1982). Contributions to a General Asymptotic Statistical Theory. Lecture Notes in Statistics 13. Springer New York, New York, NY.
- [23]. Rubin DB (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology 66 688–701.
- [24]. Shen X (1997). On methods of sieves and penalization. Annals of Statistics 25 2555–2591.
- [25]. Tibshirani R (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological) 267–288.
- [26]. van Der Laan M (2017). A Generally Efficient Targeted Minimum Loss Based Estimator based on the Highly Adaptive Lasso. International Journal of Biostatistics 13.
- [27]. van der Laan M and Rubin D (2006). Targeted Maximum Likelihood Learning. U.C. Berkeley Division of Biostatistics Working Paper Series.
- [28]. van der Laan MJ, Benkeser D and Cai W (2019). Efficient Estimation of Pathwise Differentiable Target Parameters with the Undersmoothed Highly Adaptive Lasso. arXiv preprint arXiv: 1908.05607v1.
- [29]. van der laan MJ and Dudoit S (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. U.C. Berkeley Division of Biostatistics Working Paper.
- [30]. van der Laan MJ and Robins JM (2003). Unified Methods for Censored Longitudinal Data and Causality. Springer Series in Statistics. Springer New York, New York, NY.
- [31]. van der Laan MJ and Rose S (2018). Targeted Learning in Data Science.
- [32]. van der Vaart A and Wellner J (2000). Weak Convergence and Empirical Processes: With Applications to Statistics. Springer Series in Statistics. Springer.
- [33]. Williamson B, Gilbert P, Simon N and Carone M (2017). Nonparametric variable importance assessment using machine learning techniques. UW Biostatistics Working Paper Series.
- [34]. Yang S and Ding P (2018). Asymptotic inference of causal effects with observational studies trimmed by the estimated propensity scores. Biometrika 105 487–493.

- [35]. Zhang CH (2010). Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics* 38 894–942.

Figure 1. Examples of univariate càdlàg functions with finite variation norms. The top-left, top-right, bottom-left and bottom-right plots present the standard normal density function, a minimax concave penalty function [35], a step function and the real part of a Morlet wavelet [13] respectively.

Figure 2.

The relative MSE, $n \cdot \text{MSE} / \xi^2$, and the relative absolute bias, $\sqrt{n} \cdot \left| \text{bias} / \Psi(\theta_0) \right|$, of the plug-in estimator of $\Psi(\theta_0) = P_0 \theta_0^2$ based on HAL for an oracle choice of the bound on variation norm (M.oracle), the 10-fold CV-selected bound (M.cv), a bound based on M.cv and analytic expression of $\dot{\Psi}$ without and with ϵ -relaxation (M.gcv and M.gcv+ respectively). $\xi^2 := P_0 \text{IF}^2$ is the asymptotic variance that the $n \cdot \text{MSE}$ of an AL estimator should converge to. Note that then MSE for M.oracle, M.gcv and M.gcv+ tends to 2 but that for M.cv does not.

Figure 3.

A boxplot of the ratio of bounds based on 10-fold CV and M.oracle. The horizontal gray thick dashed lines are 1 and 1/3. The y-axis is scaled based on logarithm for readability. There is a high probability that M.cv is much smaller than M.oracle; M.cv tends to the variation norm of the function being estimated, $\|\theta_0\|_v$, corresponding to 1/3 of M.oracle. Enlarging M.cv according to the analytic expression of $\hat{\Psi}$ with ϵ -relaxation results in sufficiently large bounds. The enlargement without ϵ -relaxation appears to have similar performance.

Figure 4.

The relative MSE, $n \cdot \text{MSE} / \xi^2$, and the relative absolute bias, $\sqrt{n} \cdot \left| \text{bias} / \Psi(\theta_0) \right|$, of estimators of $\Psi(\theta_0) = P_0 \theta_0^2$ is the asymptotic variance that then $n \cdot \text{MSE}$ of an AL estimator should converge to. poly: plug-in estimator based on polynomial sieve estimation. xgb: plug-in estimator based on gradient boosting. xgb.1step: one-step correction (debiasing) of the plug-in estimator based on gradient boosting. xgb.trig: data-adaptive series with trigonometric series composed with gradient boosting. All tuning parameters are CV-selected. The y-axis for relative MSE is scaled based on logarithm for readability. Note that then $n \cdot \text{MSE}$ for xgb.trig and xgb.1step tend to ξ^2 , but those for poly and xgb do not.

Figure 5.

$n \cdot \text{MSE}$ of estimators of $\Psi(\theta_0) = P_0 \theta_0^2$ based on data-adaptive series with different choices of K . The horizontal gray thick dashed line is the asymptotic variance that then $n \cdot \text{MSE}$ of an AL estimator should converge to, $\xi^2 := P_0 \text{IF}^2$. Note that $n \cdot \text{MSE}$ is not sensitive to the choice of K over a wide range of K .

Figure 6.

The relative MSE, $n \cdot \text{MSE}/\xi^2$, and the relative absolute bias, $\sqrt{n} \cdot |\text{bias}/\Psi(\theta_0)|$, of estimators of $\Psi(\theta_0) = \text{Var}_{P_0}(\mu_{0,1}(X) - \mu_{0,0}(X))$ where $\mu_{0,a}: x \mapsto \mathbb{E}_{P_0}[Y|A = a, X = x]$. $\xi^2 := P_0\text{IF}^2$ is the asymptotic variance that the $n \cdot \text{MSE}$ of an AL estimator should converge to. poly: plug-in estimator based on polynomial sieve estimation. xgb: plug-in estimator based on gradient boosting. xgb.1step: one-step correction (debiasing) of the plug-in estimator based on gradient boosting. xgb.trig: data-adaptive series with trigonometric series composed with gradient boosting. All tuning parameters are CV-selected. The y-axis for relative MSE is scaled based on logarithm for readability. Note that the $n \cdot \text{MSE}$ for xgb.trig and xgb.1step tend to ξ^2 , but those for poly and xgb do not.

Table 1.

Coverage probability of 95% Wald CI of the plug-in estimator of $\Psi(\theta_0) = P_0\theta_0^2$ based on HAL for an oracle choice of the bound on variation norm (M.oracle), the 10-fold CV-selected bound (M.cv), a bound based on M.cv and analytic expression of $\dot{\Psi}$ without and with ϵ -relaxation (M.gcv and M.gcv+ respectively). The CI is constructed based on the influence function. The coverage for M.oracle, M.gcv and M.gcv+ is approximately 95%, but that for M.cv is not.

n	M.cv	M.gcv	M.gcv+	M. oracle
500	0.87	0.96	0.96	0.97
1000	0.87	0.97	0.97	0.97
2000	0.90	0.95	0.95	0.96
5000	0.93	0.95	0.95	0.95
10000	0.89	0.95	0.95	0.95

Table 2.

Coverage probability of 95% Wald CI based on estimators of $\Psi(\theta_0) = P_0\theta_0^2$. poly: plug-in estimator based on polynomial sieve estimation. xgb: plug-in estimator based on gradient boosting. xgb.lstep: one-step correction (debiasing) of the plug-in estimator based on gradient boosting. xgb.trig: data-adaptive series with trigonometric series composed with gradient boosting. All tuning parameters are CV-selected. The CI is constructed based on the influence function. The coverage probabilities for xgb.trig and xgb.lstep are approximately 95%, but those for poly and xgb are not.

n	poly	xgb	xgb.lstep	xgb.trig
500	0.90	0.90	0.95	0.95
1000	0.86	0.89	0.95	0.95
2000	0.74	0.88	0.96	0.96
5000	0.47	0.88	0.94	0.94
10000	0.16	0.87	0.95	0.96
20000	0.02	0.86	0.96	0.96

Table 3.

Performance of the plug-in estimator of $\Psi(\theta_0) = P_0(f \circ \theta_0)$ based on data-adaptive series. Here f is not infinitely differentiable. The relative MSE is $n \cdot \text{MSE}/\xi^2$ where $\xi^2 := P_0 \text{IF}^2$ is the asymptotic variance that the $n \cdot \text{MSE}$ of an AL estimator should converge to; the root- n abs relative bias is $\sqrt{n}|\text{bias}/\Psi(\theta_0)|$. The performance appears to be acceptable in view of the small MSE and reasonable CI coverage.

n	relative MSE	root- n absolute relative bias	95% Wald CI coverage
500	0.88	3.95	0.97
1000	0.89	3.73	0.96
2000	0.79	3.15	0.97
5000	0.78	2.02	0.97
10000	0.88	2.57	0.97
20000	0.88	1.75	0.96

Table 4.

Coverage probability of 95% Wald CI based on estimators of $\Psi(\theta_0) = \text{Var}_{P_0}(\mu_{0,1}(X) - \mu_{0,0}(X))$ where $\mu_{0,a}: x \mapsto \mathbb{E}_{P_0}[Y|A = a, X = x]$. poly: plug-in estimator based on polynomial sieve estimation. xgb: plug-in estimator based on gradient boosting. xgb.1step: one-step correction (debiasing) of the plug-in estimator based on gradient boosting. xgb.trig: data-adaptive series with trigonometric series composed with gradient boosting. All tuning parameters are CV-selected. The CI is constructed based on the influence function. The coverage probabilities for xgb.trig and xgb.1step are relatively close to 95%, but those for poly and xgb are not.

n	poly	xgb	xgb.1step	xgb.trig
500	0.85	0.76	0.89	0.90
1000	0.68	0.78	0.93	0.93
2000	0.44	0.81	0.93	0.92
5000	0.11	0.80	0.89	0.87
10000	0.00	0.79	0.92	0.90
20000	0.00	0.67	0.91	0.88

Table 5.

Performance of the plug-in estimator of $\Psi(\theta_0) = \text{Var}_{P_0}(\mu_{0,1}(X) - \mu_{0,0}(X))$ where $\mu_{0,a}: x \mapsto \mathbb{E}_{P_0}[Y|A=a, X=x]$ based on data-adaptive series. Here the propensity score $g_0: x \mapsto \mathbb{E}_{P_0}[A|X=x]$ is rough. The relative MSE is $n \cdot \text{MSE}/\xi^2$ where $\xi^2 := P_0\text{IF}^2$ is the asymptotic variance that the $n \cdot \text{MSE}$ of an AL estimator should converge to; the root- n abs relative bias is $\sqrt{n}|\text{bias}/\Psi(\theta_0)|$. The performance appears to be acceptable in view of small MSE and reasonable CI coverage.

n	relative MSE	root- n absolute relative bias	95% Wald CI coverage
500	1.02	0.28	0.92
1000	1.13	0.26	0.91
2000	1.10	0.19	0.94
5000	1.03	0.02	0.93
10000	0.96	0.23	0.95
20000	0.99	0.24	0.94