

Discussion of “On Nearly Assumption-Free Tests of Nominal Confidence Interval Coverage for Causal Parameters Estimated by Machine Learning”

Edward H. Kennedy, Sivaraman Balakrishnan and Larry Wasserman

1. INTRODUCTION

We congratulate the authors on their exciting paper, which introduces a novel idea for assessing the estimation bias in causal estimates. Doubly robust estimators are now part of the standard set of tools in causal inference, but a typical analysis stops with an estimate and a confidence interval. The authors give an approach for a unique type of model-checking that allows the user to check whether the bias is sufficiently small with respect to the standard error, which is generally required for confidence intervals to be reliable.

We begin our comments by looking at an example of a simple functional.

2. EXPECTED DENSITY EXAMPLE

In this section, we illustrate the main ideas in the paper by applying them to a simpler functional. This allows us to understand better some of the critical insights of Liu, Mukherjee and Robins (2020). In particular, we consider the classic expected density functional

$$\psi = \mathbb{E}\{p(X)\} = \int p(x)^2 dx.$$

This functional has been studied extensively, with estimation and inference by now well understood (Bickel and Ritov, 1988, Birgé and Massart, 1995). Further, although it is simple, it has many of the nice properties of more complicated functionals like the expected conditional covariance or average treatment effect.

An analog of a doubly robust estimator of ψ is the one-step or first-order corrected estimator given by

$$\hat{\psi} = \frac{2}{n} \sum_{i=1}^n \hat{p}(X_i) - \int \hat{p}(x)^2 dx,$$

Edward H. Kennedy is Assistant Professor, Department of Statistics & Data Science, Carnegie Mellon University, USA (e-mail: edward@stat.cmu.edu). Sivaraman Balakrishnan is Assistant Professor, Department of Statistics & Data Science, Carnegie Mellon University, USA (e-mail: siva@stat.cmu.edu). Larry Wasserman is University Professor, Department of Statistics & Data Science, Carnegie Mellon University, USA (e-mail: larry@stat.cmu.edu).

where \hat{p} is an initial pilot estimator of the density p , which for simplicity is based on an independent auxiliary sample of size n . The rest of this analysis is conditioned on this auxiliary sample.

To fix ideas, we briefly summarize some results regarding the estimation of ψ and the estimator $\hat{\psi}$. A simple calculation shows that we may write

$$\psi = \hat{\psi} - 2 \left(\frac{1}{n} \sum_{i=1}^n \hat{p}(X_i) - \mathbb{E}[\hat{p}] \right) + \int (\hat{p} - p)^2.$$

In rough terms, if $\int (\hat{p} - p)^2$ is $o_p(1/\sqrt{n})$ then the first-order estimator achieves parametric rates (and is semiparametrically efficient). As an example, over classical Sobolev or Hölder smoothness classes, the first-order estimator is efficient if $s > d/2$, where s denotes the smoothness parameter, and d the dimension of the data. On the other hand, it is well known that a second-order U-statistic estimator (Laurent, 1996) is semiparametrically efficient if $s > d/4$ and otherwise achieves the minimax rate of $n^{-4s/(4s+d)}$.

To understand the work of Liu, Mukherjee and Robins (2020), suppose we write our initial estimate as

$$\hat{p} = \sum_{j=1}^{\infty} \hat{\theta}_j \phi_j,$$

where the ϕ_j form an orthonormal basis with respect to the Lebesgue measure. Then for $p = \sum_j \theta_j \phi_j$ a straightforward calculation shows that the conditional bias (given the auxiliary sample) is

$$\begin{aligned} \text{Bias} &= \mathbb{E}(\hat{\psi} - \psi) = - \int \{\hat{p}(x) - p(x)\}^2 dx \\ &= - \sum_{j=1}^{\infty} (\hat{\theta}_j - \theta_j)^2. \end{aligned}$$

Note that this is the bias, not the squared bias; therefore for this functional, the standard first-order estimator has the monotone bias property. We can decompose this bias as

$$-\text{Bias} = \underbrace{\sum_{j=1}^k (\hat{\theta}_j - \theta_j)^2}_{\text{Bias}_k} + \underbrace{\sum_{j>k} (\hat{\theta}_j - \theta_j)^2}_{\text{Truncation Bias}}.$$

Then, in order to test if the estimator $\widehat{\psi}$ is biased, one can ignore the truncation bias and simply test if the first projected bias term is large. Of course, the projected bias is a quadratic function, and so can be estimated with a second-order estimator (in this case a simple U-statistic). For example, one could use a modified version of the estimator proposed by Laurent (1996)

$$\widehat{\text{Bias}}_k = \sum_{j=1}^k \widehat{\theta}_j^2 - \frac{2}{n} \sum_{j=1}^k \widehat{\theta}_j \left[\sum_{i=1}^n \phi_j(X_i) \right] + \frac{1}{n(n-1)} \sum_{j=1}^k \sum_{u=1}^n \sum_{v=1}^n \phi_j(X_u) \phi_j(X_v).$$

Using standard properties of U-statistics, or following the analysis of Laurent (1996) we find that if $k = o(n)$ then $\widehat{\text{Bias}}_k$ is a \sqrt{n} -consistent estimator of Bias_k , and the standard Wald interval can be used to test hypotheses regarding the magnitude of Bias_k .

To summarize, in this very simple example we are able to see two key insights of Liu, Mukherjee and Robins (2020):

1. For a wide class of functionals, the bias of the first-order estimator has a relatively simple form. Often a second-order U-statistic based estimator can be designed to estimate the bias (or a portion of the bias).

2. The statistician has considerable freedom in the setting of testing hypotheses regarding the bias of an estimator. In particular, in the above setting the statistician has license to both select a choice of an orthonormal basis, as well as to ignore the resulting “truncation bias”, and focus more narrowly on testing the magnitude of Bias_k . While these choices might affect the power of the resulting test, they do not affect the validity (Type I error control) of the test. We return to this point in Sections 6 and 7.

The work of Liu, Mukherjee and Robins (2020) presents an elegant extension of these ideas to a more complex class of functionals that arise in causal inference.

3. THE MONOTONE BIAS PROPERTY

In this section, we simply illustrate how the monotone bias property is a property of not only the functional but also a given estimator. In particular, given an estimator/functional pair satisfying the monotone bias property, we can construct a similar first-order estimator that does *not* have the monotone bias property, using the ideas of Newey and Robins (2018), and also discussed by Liu, Mukherjee and Robins (2020) in Supplementary Materials Section S.1.1.

For example, consider the first-order estimator of the expected density from Section 2

$$\widehat{\psi} = 2\mathbb{P}_n(\widehat{p}) - \int \widehat{p}^2,$$

which, as noted previously, has conditional bias $\mathbb{E}(\widehat{\psi} - \psi) = -\int (\widehat{p} - p)^2$ and so satisfies the monotone bias property. An alternative first-order estimator is given by

$$\widehat{\psi}_2 = \mathbb{P}_n(\widehat{p}_1 + \widehat{p}_2) - \int \widehat{p}_1 \widehat{p}_2,$$

where \widehat{p}_j are two different density estimators, built from separate independent samples (this can always be achieved by splitting the sample into thirds rather than halves). Importantly, this alternative estimator has conditional bias equal to

$$\mathbb{E}(\widehat{\psi}_2 - \psi) = -\int (\widehat{p}_1 - p)(\widehat{p}_2 - p)$$

and so does *not* satisfy the monotone bias property. The estimator $\widehat{\psi}_2$ has some nice properties: it is doubly-robust in the sense that it is consistent if either \widehat{p}_1 or \widehat{p}_2 is; it is root- n consistent, asymptotically normal, and efficient under the same conditions as the classic first-order estimator $\widehat{\psi}$; and further it can be shown that $\widehat{\psi}_2$ is minimax optimal when p is Hölder-smooth if \widehat{p} is a particular undersmoothed estimator.

The same trick can be used to demonotonize the expected conditional covariance: instead of using the standard first-order estimator $\mathbb{P}_n\{(A - \widehat{\pi})^2\}$ we would instead use

$$\mathbb{P}_n\{(A - \widehat{\pi}_1)(A - \widehat{\pi}_2)\},$$

which has essentially the same properties as $\widehat{\psi}_2$ (except that minimax optimality with undersmoothing in the Hölder setup would only be achieved when the Hölder smoothness of π is greater than $d/4$, unless the density of X was known).

Of course with these variants one could still use the approach detailed in Section 4 of Liu, Mukherjee and Robins (2020), for example, testing a null that the Cauchy–Schwarz bound on the bias is small. However this bound would no longer be as relevant for justifying small bias, since one would only need the typically weaker condition that, for example, for the expected density,

$$-\int (\widehat{p}_1 - p)(\widehat{p}_2 - p) = o_{\mathbb{P}}(n^{-1/2}).$$

Note for example that the mean of the left-hand side above is the squared bias, which can be smaller than the mean squared error with undersmoothing.

4. THE ROLE OF COVARIATE STRUCTURE

In this section, we raise some questions regarding the role covariate structure should play in estimation and inference for causal effects.

Covariate structure plays a relatively important role in Liu, Mukherjee and Robins (2020), as well as in the general higher-order influence function literature (Robins

et al., 2008, 2009, 2017). For example, minimax lower bounds for the average treatment effect and expected conditional covariance are known in models where the covariate density is known (Robins et al., 2009); these bounds can be attained even when the density is unknown, but this requires sufficient smoothness of the density, at least whenever the propensity scores and regression functions are smooth, but not so smooth to admit root- n rates (Robins et al., 2008, 2017). The complete minimax story with unknown density therefore appears to be an open problem. Based on upper bounds given in Robins et al. (2008), one might expect the rate to interpolate between the classic functional estimation rate $n^{-4s/(4s+d)}$ (where s is the average smoothness of the propensity score and regression functions), and the slower fixed design rate $n^{-2s/d}$, for example, an analog of which appears in Wang et al. (2008) in a related conditional variance estimation problem. The connection between nonrandomness (e.g., fixed design) and nonsmoothness in the covariate density would be interesting to explore further; at the nonsmooth extreme the aforementioned interpolation would suggest that these two structures align.

An interesting question is how and whether randomness of the covariates, and/or smoothness of the covariate density, should be relied upon in practice. In many problems one might expect convenience samples to be the norm, rather than true random samples, in which case it could be argued that the covariates should be conditioned on in any inferential procedures. In this kind of setup, one might wonder if estimators that exploit randomness/smoothness or other structure of the covariates might be less robust when such structure is not present; ideally our estimators would be able to adapt to such structure when present, but would otherwise be unaffected in the absence of randomness/smoothness.

More specifically, it would be interesting to explore whether higher-order influence function-based estimators have a robustness to misspecification of density estimates, for example, relative to a standard first-order estimator, which does not require modeling the density. Similarly, it would be useful to know how higher-order estimators and/or the bias test proposed by Liu, Mukherjee and Robins (2020) perform in a fixed design setup, insofar as this is the right context for conditional-on-covariate inference. Note however that there are at least four possible choices of setups, based on the design and target parameter; for example, for (half) the average treatment effect $\mathbb{E}(Y^1)$ one might consider:

1. the parameter $\int \mu_1(x) d\mathbb{P}(x)$ in a usual random design setup;
2. the parameter $\int \mu_1(x) dx$ in a fixed design setup with X equally spaced on a grid;
- 3/4. the parameter $\frac{1}{n} \sum_i \mu_1(x_i)$ in either a random or fixed design setup.

Since the fixed design setup leads to slower rates, for example, in variance estimation, one might expect this to affect the power but not validity of the bias test of Liu, Mukherjee and Robins (2020). However, the details are not obvious to us.

5. EXACT PARAMETRIC INFERENCE

The authors focus on nonparametric inference, and we agree that this is usually the best way to proceed. Still, it is interesting to ask what happens in a parametric model. Typically such models are too inflexible to yield reliable inference. But suppose we use a highly flexible parametric family. Examples include mixture models and exponential families based on truncated series expansions. The extra flexibility these models provide come at a cost: constructing confidence intervals can be even harder than in the nonparametric case. Indeed, models such as mixture models are nonregular and inference is challenging. However, a recent method called *universal inference* (Wasserman, Ramdas and Balakrishnan, 2020) provides a potential solution. The method is simple. Denote the model by $\mathcal{P} = \{p_\theta : \theta \in \Theta\}$. Split the data into two: \mathcal{D}_0 and \mathcal{D}_1 . Let

$$T(\theta) = \frac{\mathcal{L}_0(\hat{\theta}_1)}{\mathcal{L}_0(\theta)},$$

where \mathcal{L}_0 is the likelihood constructed from \mathcal{D}_0 and $\hat{\theta}_1$ is any estimate of θ based on \mathcal{D}_1 . This could be the maximum likelihood estimator, a robust estimator or a minimum distance estimator for example. Now repeat this process B times yielding $T_1(\theta), \dots, T_B(\theta)$ and define $\bar{T} = B^{-1} \sum_b T_b(\theta)$. Finally, set $C_n = \{\theta : \bar{T}(\theta) \leq 1/\alpha\}$. Then it was shown that

$$P_{\theta^*}(\theta^* \in C_n) \geq 1 - \alpha,$$

for all $\theta^* \in \Theta$. The result extends to functionals. Let $\psi = f(\theta^*)$ and replace the term $\mathcal{L}_0(\theta)$ with the profile likelihood for ψ . Then we get a finite sample valid confidence set for ψ . It would be interesting to compare exact, flexible parametric methods to approximate asymptotic methods in terms of coverage, robustness, and efficiency.

6. STRUCTURE-DRIVEN VS. METHODS-DRIVEN ESTIMATION & INFERENCE

In this section, we draw a distinction between structure-driven and methods-driven estimation and inference, and relate this to the work of Liu, Mukherjee and Robins (2020) and Robins et al. (2008, 2009, 2017).

Many successful tools in statistics and machine learning are algorithmic or methods-driven, that is, were not derived with a clear inferential goal (for example, optimality or minimaxity in a particular statistical model). The most successful of these tools are often “robust”, in

the sense of having strong guarantees under many different statistical models. Prominent examples include methods like spectral clustering, random forests, and deep neural networks. On the other hand, other methods are decidedly more model-based or structure-driven: for example, a truncated series estimator for regression or density estimation, where one assumes some structure (e.g., decay of coefficients in a particular basis, Hölder smoothness, etc.) and then tailors the estimator precisely to exploit this structure. These structure-driven estimators are often crucial for setting theoretical benchmarks, but their theoretical guarantees may not be as “robust”. This is reminiscent of the discussion in [Donoho et al. \(1995\)](#) (Section 3) distinguishing between estimators that are exactly minimax in a single model, versus approximately minimax across a large collection of models. The distinctions between structure and methods-driven methods are sometimes murky, but we find them nonetheless useful. One might attempt to similarly categorize classical first-order methods, as well as their higher-order counterparts ([Liu, Mukherjee and Robins, 2020](#), [Robins et al., 2008, 2009, 2017](#)), both in the context of functional estimation, as well as for the bias tests of [Liu, Mukherjee and Robins \(2020\)](#).

In the context of estimation, standard first-order estimators, such as the classic doubly robust estimator of the average treatment effect, seem to us more methods-driven. They are a “black-box” correction that can be applied to any initial nuisance function estimates, and in general yield improved guarantees over the plug-in estimator. This at least partially explains the strong empirical success these methods have found across various application domains. On the other hand, higher-order influence function methods typically rely on carefully constructed series estimates and achieve better performance, for example, over appropriate Hölder spaces, potentially at the expense of being more structure-driven. Higher-order estimators are useful in understanding the fundamental limits for functional estimation, but may be less agnostic about the underlying structure than corresponding first-order estimators, and may require more significant domain-knowledge in their practical application.

The previous discussion was focused on first and higher-order methods for estimation. However, [Liu, Mukherjee and Robins \(2020\)](#) address the slightly different problem of testing. To simplify the discussion, we focus on the simplest cases, where the monotone-bias property holds, and the claims of [Liu, Mukherjee and Robins \(2020\)](#) are the strongest. As we alluded to earlier, the bias testing problems that [Liu, Mukherjee and Robins \(2020\)](#) consider have some interesting features:

1. The statistician has flexibility to specify the (portion of the) bias she wishes to test for.

2. As is typical in hypothesis testing problems the primary focus is on controlling the Type I error (ensuring validity). In reasoning about the power of the test, the statistician also has the flexibility to specify the alternatives against which she desires power.

3. The authors work in a “falsification” setup, where in essence any rejection is useful from a practical standpoint in tempering the statistician’s claims about the first-order estimator and its associated Wald interval, and any failure to reject comes with few added harms.

From this standpoint, we can revisit the methods-driven versus structure-driven distinction, and ask: are there downsides to using structure-driven methods for inference? We believe the answer may be yes, but more work is required. In the next section, we consider a simple extension that could help move the [Liu, Mukherjee and Robins \(2020\)](#) approach to be more methods-driven.

7. A SIMPLE ENSEMBLED BIAS TEST

In this section, we propose a simple modification of the bias test that may provide some practical advantages in terms of power, by incorporating generic flexible regression tools, and thus weakening ties to classical series methods.

In the usual series setup, one specifies a basis that is a priori expected to capture underlying structure well (typically Hölder smoothness), such as Daubechies wavelets or a Fourier series. We instead suggest an alternative based on trying to estimate a good basis from the data, by constructing a dictionary of appropriate predictors, and then orthogonalizing them. The motivation behind this is that we want the truncation bias (the second term in Equation (2.3)) to be small in order for the test to have power, and this depends on the basis terms capturing enough important structure in the projected bias.

For simplicity consider the expected conditional variance functional $\mathbb{E}\{\text{var}(A | X)\}$. Suppose we have access to an auxiliary sample of observations (note we could always split the sample accordingly so as to have one more fold). In the auxiliary sample, we regress the residuals $(A - \hat{b}(X))$ on X , using m different methods. For example we could use random forests, lasso, boosting, etc. as well as variants of these and other methods with different tuning parameter choices. Importantly, these methods could capture different kinds of structure, for example, smoothness, sparsity, additivity, etc. This yields m different predictors $(\hat{f}_1, \dots, \hat{f}_m)$. Next, we evaluate these predictors on the sample used for testing the bias, to obtain m different n -vectors of predicted values $\hat{f}_j(X_i)$ for $j = 1, \dots, m$ and $i = 1, \dots, n$. Finally, we orthogonalize (the span of) these vectors and use them as an estimated basis; namely using the notation in the paper we let $\bar{z}_k(X) = (\hat{f}_1(X), \dots, \hat{f}_m(X))$, possibly concatenated with a typical fixed basis.

Intuitively since the m different predictors can potentially capture diverse kinds of structure, if one (or a weighted combination) of predictors leads to small truncation bias then the projection onto their span should as well (Tsybakov, 2003). It seems that the theory from Liu, Mukherjee and Robins (2020) would all go through as before, since conditional on the auxiliary sample the estimated basis functions are all fixed; the advantage is there should be some gain from tailoring the basis functions to the specific bias structure that is trying to be captured.

REMARK 1. Lastly, we would like to pick a small nit. The authors have taken the now common approach of referring to statistical prediction methods as “machine learning.” Perhaps the field of statistics should instead reclaim regression. Further, the field of machine learning is large and diverse, with many successes outside of pure prediction problems; so equating the term machine learning with prediction tools diminishes these other successes.

REFERENCES

- BICKEL, P. J. and RITOV, Y. (1988). Estimating integrated squared density derivatives: Sharp best order of convergence estimates. *Sankhyā Ser. A* **50** 381–393. [MR1065550](#)
- BIRGÉ, L. and MASSART, P. (1995). Estimation of integral functionals of a density. *Ann. Statist.* **23** 11–29. [MR1331653](#) <https://doi.org/10.1214/aos/1176324452>
- DONOHO, D. L., JOHNSTONE, I. M., KERKYACHARIAN, G. and PICARD, D. (1995). Wavelet shrinkage: Asymptopia? *J. Roy. Statist. Soc. Ser. B* **57** 301–369. [MR1323344](#)
- LAURENT, B. (1996). Efficient estimation of integral functionals of a density. *Ann. Statist.* **24** 659–681. [MR1394981](#) <https://doi.org/10.1214/aos/1032894458>
- LIU, L., MUKHERJEE, R. and ROBINS, J. M. (2020). On nearly assumption-free tests of nominal confidence interval coverage for causal parameters estimated by machine learning. *Statist. Sci.* **35** 518–539.
- NEWBY, W. K. and ROBINS, J. M. (2018). Cross-fitting and fast remainder rates for semiparametric estimation. Preprint. Available at [arXiv:1801.09138](https://arxiv.org/abs/1801.09138).
- ROBINS, J., LI, L., TCHETGEN, E. and VAN DER VAART, A. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman. Inst. Math. Stat. (IMS) Collect.* **2** 335–421. IMS, Beachwood, OH. [MR2459958](#) <https://doi.org/10.1214/193940307000000527>
- ROBINS, J., TCHETGEN, E., LI, L. and VAN DER VAART, A. (2009). Semiparametric minimax rates. *Electron. J. Stat.* **3** 1305–1321. [MR2566189](#) <https://doi.org/10.1214/09-EJS479>
- ROBINS, J. M., LI, L., MUKHERJEE, R., TCHETGEN, E. T. and VAN DER VAART, A. (2017). Minimax estimation of a functional on a structured high-dimensional model. *Ann. Statist.* **45** 1951–1987. [MR3718158](#) <https://doi.org/10.1214/16-AOS1515>
- TSYBAKOV, A. B. (2003). Optimal rates of aggregation. *Learning Theory and Kernel Machines* 303–313.
- WANG, L., BROWN, L. D., CAI, T. T. and LEVINE, M. (2008). Effect of mean on variance function estimation in nonparametric regression. *Ann. Statist.* **36** 646–664. [MR2396810](#) <https://doi.org/10.1214/009053607000000901>
- WASSERMAN, L., RAMDAS, A. and BALAKRISHNAN, S. (2020). Universal inference. *Proc. Natl. Acad. Sci. USA.* **117** 16880–16890.