

1 We thank the reviewers for their thorough comments. One common comment from the reviewers is that they would like
2 to see experiments on more than one data set. If the reviewers find this point to be critical, we would be happy to add
3 another experiment to the supplement in the camera-ready version (we have already completed additional experiments,
4 but left them out because the results were very similar to those on the airfoil data—the theoretical guarantees we provide
5 are exact in finite sample without any distributions assumptions, so in the absence of estimating the likelihood ratio for
6 the weights, there should be no surprises when running our weighted conformal prediction on a new data set).
7 Below we respond to some of the specific comments and questions in the reviews. Square brackets [...] indicate the
8 reviewer’s comments and our replies follow.

9 **Reviewer #2**

Thank you for your detailed and positive feedback.

[...analysis of the effects of approximating the likelihood ratio...] Great question; if we have a guarantee on estimating $w(x) = dP_X(x)/dP_X(x)$ in the ℓ_∞ (essential supremum) norm, i.e., if we have bound on

$$\|\hat{w} - w\|_\infty = \text{ess sup}_{x \in \mathbb{R}^d} |\hat{w}(x) - w(x)|,$$

10 then we can show how to translate this into an approximate coverage guarantee on the conformal prediction intervals.
11 We can include a corollary to this effect in the main paper, with a proof in the supplement. This would of course be a
12 first step: an ℓ_∞ accuracy guarantee may be stringent, but a more careful analysis should be possible in more specialized
13 situations, out of the scope of the current paper.

14 [...classification...] This is a great point. The same methodology extends seamlessly to classification (only the notation
15 changes) and we can point this out in the camera-ready version.

16 [...code...] We apologize for not including an anonymized link. Our code is publicly available on github and we will
17 add the link as soon as the paper can be de-anonymized.

18 **Reviewer #4**

19 Thank you for your detailed and helpful review.

20 [...additional experiments...] Please see our comment at the top of the page.

21 [...analysis of the effects of approximating the likelihood ratio...] Please see our response to Reviewer #2.

22 [...computational tractability...] Yes, in general the complexity may be combinatorial. However, as mentioned in the
23 paper, fortunately for covariate shift (and a few other scenarios mentioned in the discussion), the calculation is easier.

24 **Reviewer #5**

25 Thank you for your detailed and astute review.

26 [...transfer learning...] As far as we can tell, this is not directly related to the goal of our paper, but as transfer learning
27 is often discussed alongside covariate shift, it would be interesting to consider this connection; space permitting we can
28 include a comment in the camera-ready version.

29 [...importance sampling...] These ideas are intimately connected to our work, and form the basis for much of the
30 methodology in covariate shift, though the modern covariate shift literature simply tends to use different language
31 than the older importance sampling literature in statistics. One could use importance sampling to approximately fix
32 the incompatibility of test and training distributions in conformal prediction (sample from a batch of test points with
33 probabilities that are proportional to their likelihood in the training sample). Our methodology is in essence a *direct* way
34 of doing this, without requiring subsampling, and is preferable since subsampling would introduce additional variance.
35 We can include this interpretation in the camera-ready if accepted.

36 [...experiments...] Beyond what we mentioned at the top of this page, the experiment is indeed proof-of-concept. But
37 we do believe it shows the method to be practically applicable; we have investigated several angles of the problem,
38 including what happens when we still estimate a likelihood ratio and there is actually no covariate shift (results given
39 in the supplement). Regarding testing whether our method is “competitive”, as far as we know, there are no other
40 methods that produce provably valid prediction intervals (with distribution-free, finite-sample coverage guarantees)
41 under covariate shift, so we do not know of competing methods that we might compare against.

42 [...detect drift incrementally...] This is an interesting question. There is orthogonal work on using conformal prediction
43 to detect a change in distribution: see “Inductive Conformal Martingales for Change-Point Detection”. Future work
44 could consider sequential versions of a problem of both detecting covariate shift and producing appropriately corrected
45 prediction intervals, combining the insights of our paper with theirs.

46 [...generalization to each point coming from different distribution...] While we do not know of any examples or
47 applications that need this full generality, it is very plausible that the test points lie in (multiple) unknown clusters, with
48 each cluster coming from a different distribution shift. This type of distribution could possibly be detected by mixture
49 methods, correcting for each cluster individually. We are still exploring the full power of the general technique and will
50 continue to do so in future work, but it is outside the scope of this short paper.