> My main concern with the analysis is the choice of benchmark test functions and to what extent the conclusions are applicable in a wider context. While the scalar box integral with arbitrary masses is one of the most complicated integrals appearing at NLO, its numerical evaluation is by no means a bottleneck in modern algorithms. The two-loop self-energy is slightly more interesting since the analytic expressions for integrals of this type with arbitrary masses lead to the appearance of elliptic integrals which are difficult to evaluate numerically.

A: The reviewer rightly points out that the scalar box integral is not expensive to evaluate, and we have also mentioned this in the analysis. However, we chose it because of its complexity in higher dimensions while also being relatively quick to evaluate and generate the training data. A better test of speed does indeed come from the two-loop self-energy where we have stated the gain in speed of evaluation.

> Nevertheless many of the two-loop self energy integrals themselves are by now well understood (in four dimensions) in terms of elliptic polylogarithms which can be efficiently evaluated through q-expansions of iterated Eisenstein integrals (see for example 1912.00077). I suspect that neither of the test integral topologies would present any difficulties for state-of-the-art precision phenomenology.

A: We thank the reviewer for pointing out this reference. Examining it, we find that the iterated Eisenstein integrals are an efficient method for particular cases, such as the sunrise diagram, but not for general two-loop diagrams. The reviewer is interested in applications to state-of-the-art precision phenomenology as a better benchmark, but we believe that considering more general cases, albeit of less phenomenological need, provide a better discriminant of the approximation techniques, otherwise one might object to the generalizability of the analysis.

> The fact that these test cases are considered in odd dimensions is strange to me and unmotivated. Applications in high energy physics rarely involve quantum field theories in odd dimensions as far as I am aware, although I would be happy to hear otherwise.

A: We do not see an issue with odd dimensions since one can invoke the property in eq. (17) to go between even and odd dimensions (in these particular cases). The dimensions were chosen to be spread out between low dimensions and a maximum of $d = 9$ for the scalar box integral. In a more general case, we do not see a reason for the parity of the dimension to influence the results, but we understand the point of view of the reviewer in connecting the analysis more straightforwardly to applications in quantum field theory.

> The issue lies in the fact that the best method will clearly depend on the particular problem at hand. It is well known that Machine Learning methods scale better than conventional techniques for high dimensional problems and so the conclusions of the study do not appear to go against conventional wisdom. I am therefore unconvinced whether the additional information obtained by this systematic study would be of use to someone considering the application of such techniques to a more technically demanding computation.

A: We express our concerns about the first statement in the introduction and therefore specify to the case where the expensive function will be evaluated a large number of times. Although

this analysis does not contradict conventional wisdom, it provides an anchor for it, and further provides discrimination between various Machine Learning techniques. As far as we are aware, the current literature does not offer a reason to choose one ML technique over the other, whereas in this analysis we see strong evidence for the use of neural networks, especially in high dimensions.

> 1. The introduction summarises some previous attempts to use regression algorithms for high energy physics in the literature. The speed improvement declared in those articles is listed but I believe a suitable explanation of how the times should be compared should be added. The speed up declared refs. [7,8] for the call times of the approximate function is misleading and irrelevant when the generation of the input data dominates the overall evaluation time. The smaller speed up declared in refs [10-12] includes training, testing and interpolation and so represents a realistic measure of the optimisation. The comments related to refs. [10-12] in this article are perfectly well summarised but an explicit statement regarding the comparison with refs. [7,8] would be welcome.
>
> 2 The authors may consider referring to the application of Gaussian processes to the approximation of extremely complex two-loop five particle scattering amplitudes in the NNLO computation of pp->3 photons. (1911.00479) and Zaharid/GPTree (https://zenodo.org/record/3571309#.Ye6p_FjMKqA)
>
> 3. Recent work on the evaluation of multi-loop integrals provides useful context for this work (2112.09145). The authors could consider adding a reference.

We have modified the introduction to more accurately reflect the gain in speed for refs [9, 10] in the current version ([7, 8] in the previous version). We have also found the suggested articles useful to reference in the introduction for better context.