
HEBO: Pushing The Limits of Sample-Efficient Hyperparameter Optimisation

1 Reproducibility Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
 - (d) Have you read the ethics author’s and review guidelines and ensured that your paper conforms to them? <https://automl.cc/ethics-accessibility/> [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results, including all requirements (e.g., requirements.txt with explicit version), an instructive README with installation, and execution commands (either in the supplemental material or as a URL)? [Yes]
 - (b) Did you include the raw results of running the given instructions on the given code and data? [Yes]
 - (c) Did you include scripts and commands that can be used to generate the figures and tables in your paper based on the raw results of the code, data, and instructions given? [Yes]
 - (d) Did you ensure sufficient code quality such that your code can be safely executed and the code is properly documented? [Yes]
 - (e) Did you specify all the training details (e.g., data splits, pre-processing, search spaces, fixed hyperparameter settings, and how they were chosen)? [Yes]
 - (f) Did you ensure that you compared different methods (including your own) exactly on the same benchmarks, including the same datasets, search space, code for training and hyperparameters for that code? [Yes]
 - (g) Did you run ablation studies to assess the impact of different components of your approach? [Yes]
 - (h) Did you use the same evaluation protocol for the methods being compared? [Yes]
 - (i) Did you compare performance over time? [No]
 - (j) Did you perform multiple runs of your experiments and report random seeds? [Yes]

(k) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]	36 37
(l) Did you use tabular or surrogate benchmarks for in-depth evaluations? [Yes]	38
(m) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]	39 40
(n) Did you report how you tuned hyperparameters, and what time and resources this required (if they were not automatically tuned by your AutoML method, e.g. in a NAS approach; and also hyperparameters of your own method)? [Yes]	41 42 43
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...	44
(a) If your work uses existing assets, did you cite the creators? [Yes]	45
(b) Did you mention the license of the assets? [N/A]	46
(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]	47
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]	48 49
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]	50 51
5. If you used crowdsourcing or conducted research with human subjects...	52
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]	53 54
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]	55 56
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]	57 58