

Interactive Hyperparameter Optimization in Multi-Objective Problems via Preference Learning

Joseph Giovanelli¹, Alexander Tornede², Tanja Tornede², Marius Lindauer²

¹Alma Mater Studiorum — University of Bologna

²Institute of Artificial Intelligence, L3S Research Center, Leibniz University Hannover
j.giovanelli@unibo.it, {a.tornede, t.tornede, m.lindauer}@ai.uni-hannover.de

Abstract

Hyperparameter optimization (HPO) is important to leverage the full potential of machine learning (ML). In practice, users are often interested in multi-objective (MO) problems, i.e., optimizing potentially conflicting objectives, like accuracy and energy consumption. To tackle this, the vast majority of MO-ML algorithms return a Pareto front of non-dominated machine learning models to the user. Optimizing the hyperparameters of such algorithms is non-trivial as evaluating a hyperparameter configuration entails evaluating the quality of the resulting Pareto front. In literature, there are known indicators that assess the quality of a Pareto front (e.g., hypervolume, R2) by quantifying different properties (e.g., volume, proximity to a reference point). However, choosing the indicator that leads to the desired Pareto front might be a hard task for a user. In this paper, we propose a human-centered interactive HPO approach tailored towards multi-objective ML leveraging preference learning to extract desiderata from users that guide the optimization. Instead of relying on the user guessing the most suitable indicator for their needs, our approach automatically learns an appropriate indicator. Concretely, we leverage pairwise comparisons of distinct Pareto fronts to learn such an appropriate quality indicator. Then, we optimize the hyperparameters of the underlying MO-ML algorithm towards this learned indicator using a state-of-the-art HPO approach. In an experimental study targeting the environmental impact of ML, we demonstrate that our approach leads to substantially better Pareto fronts compared to optimizing based on a wrong indicator pre-selected by the user, and performs comparable in the case of an advanced user knowing which indicator to pick.

Introduction

Practical applications of machine learning often call for the optimization of more than one loss or objective function, called multi-objective machine learning (MO-ML) (Jin 2006) borrowing from the notion of multi-objective optimization (MO) (Deb, Sindhya, and Hakanen 2016; Gunantara 2018). For example, instead of only focusing on the performance of a model, its energy consumption is becoming more and more important in various domains such as edge computing, but also in general, sparked by efforts in the area of green artificial intelligence (Green AI) (Schwartz et al.

2019; van Wynsberghe 2021). Comparing two ML models with respect to several objectives is non-trivial and most approaches solve this problem by returning a set of solutions that cannot be improved further without going to the expense of at least one of the objectives, called the Pareto front.

Naturally, as standard ML algorithms, MO-ML algorithms expose hyperparameters controlling their learning behavior and thus, also the resulting Pareto front. To unleash the full potential of the MO-ML algorithms, these hyperparameters should be optimized with adequate methods. Unfortunately, optimizing the hyperparameters of such MO-ML algorithms is challenging for a user with standard hyperparameter optimization (HPO) (Feurer and Hutter 2019; Bischl et al. 2023) approaches, such as SMAC (Hutter, Hoos, and Leyton-Brown 2011; Lindauer et al. 2022), Optuna (Akiba et al. 2019), Hyperopt (Komer, Bergstra, and Eliasmith 2014), HpBandSter (Falkner, Klein, and Hutter 2018) or SyneTune (Salinas et al. 2022), which iteratively evaluate configurations. This is the case, as evaluating a hyperparameter configuration of an MO-ML algorithm involves evaluating the quality of the Pareto front of models returned by the corresponding algorithm. Several so-called (quality) indicators, such as hypervolume (Zitzler and Thiele 1999), are used to assess different properties of the shape of the Pareto front and in principle can also be used for the purpose of rating a configuration in the context of HPO. Nevertheless, although a user might have a clear idea about what kind of Pareto front shape they would like to choose their final model from, choosing the quality indicator leading to this Pareto front shape is a challenging task in practice. At the same time, however, correctly configuring the HPO tool with the right loss function, i.e. quality indicator in our case, is crucial to achieve the desired result.

With this paper, we propose an interactive human-centered HPO (Pfisterer, Thomas, and Bischl 2019; Souza et al. 2021; Moosbauer et al. 2021; Hvarfner et al. 2022; Moosbauer et al. 2022; Francia, Giovanelli, and Pisano 2023; Segel et al. 2023; Mallik et al. 2023) approach for MO-ML algorithms that frees users from choosing a predefined quality indicator suitable for their needs by learning one tailored towards them based on feedback. To achieve this, it first learns the desired Pareto front shape from the user in a short interactive session and then starts a corresponding HPO process optimizing towards the previously

learned Pareto front shapes. Instead of requiring the user to present us with a concrete Pareto front they would favor, we interactively and iteratively present them a few pairs of Pareto fronts asking them for their preferences. Based on these pairwise comparisons, we leverage methods from the field of preference learning (PL) (Fürnkranz and Hüllermeier 2010) to learn a latent utility function serving as a Pareto front quality indicator customized to the user. In the subsequent stage, we run a state-of-the-art HPO tool instantiated with the learned Pareto front quality indicator to evaluate configurations and the corresponding Pareto fronts.

In summary, we make the following contributions:

1. We propose an interactive approach to learn a Pareto front quality indicator from pairwise comparisons based on a latent utility function with methods from the field of preference learning. This quality indicator is customized to the user and can be learned from a small number of pairwise comparisons.
2. We combine the aforementioned learned quality indicator with an HPO tool to provide a full-fledged HPO approach for multi-objective MO-ML algorithms, which frees the user to choose an appropriate Pareto front quality indicator offering less opportunity for a mistake and thus, bad optimization results.
3. In an experimental case study, we demonstrate that our approach leads to substantially better Pareto fronts compared to optimizing based on a wrong indicator pre-selected by the user, and performs comparable in case of an advanced user knowing which indicator to pick. Thus, our approach makes HPO for MO-ML algorithms substantially more easily and robustly applicable in practice.

Background

Since our work is spanned along the dimensions of hyperparameter optimization (HPO), multi-objective optimization (MO), and preference learning (PL), we introduce in the following all of the aforementioned concepts.

Hyperparameter Optimization

HPO formalizes the task of finding a hyperparameter configuration for a machine learning algorithm leading to a well-performing model on a given dataset

$$\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N \in \mathbb{D} \subset \mathcal{X} \times \mathcal{Y} \quad (1)$$

with an instance space \mathcal{X} and a target space \mathcal{Y} . In addition to the dataset, we are provided with a hyperparameter configuration space $\Lambda = \Lambda_1 \times \dots \times \Lambda_K$ with K hyperparameters, where Λ_k is the domain of the k^{th} hyperparameter, and an algorithm $A : \mathbb{D} \times \Lambda \rightarrow \mathcal{H}$ which trains a model from the model space \mathcal{H} given a dataset and a hyperparameter configuration. Furthermore, we are provided with a loss function $\mathcal{L} : \mathcal{H} \times \mathbb{D} \rightarrow \mathbb{R}$ quantifying how well a given model performs on a given dataset. The loss function can be used to assess the quality of a hyperparameter configuration by splitting the original dataset \mathcal{D} into two disjoint datasets \mathcal{D}_{train} and \mathcal{D}_{test} , where the model is trained only based on \mathcal{D}_{train}

but evaluated with \mathcal{L} on \mathcal{D}_{test} . Overall, we seek to find the optimal hyperparameter configuration $\lambda^* \in \Lambda$ defined as

$$\lambda^* \in \arg \min_{\lambda \in \Lambda} \mathcal{L}(A(\mathcal{D}_{train}, \lambda), \mathcal{D}_{test}) . \quad (2)$$

There exist several approaches to automatically solving the optimization problem defined in (2), many of which are powered by Bayesian optimization (BO) (Frazier 2018) or evolutionary approaches (Bischl et al. 2023). Most of these techniques internally iteratively evaluate a large set of configurations based on their true estimated loss. To this end, they split up a so-called validation dataset \mathcal{D}_{val} from the training dataset \mathcal{D}_{train} to estimate the loss $\mathcal{L}(A(\mathcal{D}_{train}, \lambda), \mathcal{D}_{test})$ of a configuration λ by $\mathcal{L}(A(\mathcal{D}_{train}, \lambda), \mathcal{D}_{val})$ and avoid a biased overfit to \mathcal{D}_{test} .

Multi-Objective Machine Learning

When it comes to learning a machine learning model with more than one (possibly conflicting) objective or loss function $\mathcal{L}_1, \dots, \mathcal{L}_M$ in mind, comparing the quality of models becomes difficult. For instance, in the context of Green AI, searching for a model with higher performance usually leads to one with higher power consumption, and vice versa. Thus, in multi-objective ML, learning algorithms leverage the concept of dominance formalizing that one model $h_1 \in \mathcal{H}$ is better than another $h_2 \in \mathcal{H}$, if h_1 performs better in at least one of the objectives while not performing worse than h_2 in the remaining ones. Yet, this leaves us with a set of non-dominated models, which are indistinguishable with respect to this dominance idea. These models form a so-called Pareto front $P_{\mathcal{D}_{val}}(H)$ evaluated on dataset \mathcal{D}_{val} for a given set of models $H \subset \mathcal{H}$. Formally, a Pareto front is defined as

$$P_{\mathcal{D}_{val}}(H) = \left\{ h \mid \begin{array}{l} h \in H, \nexists h' \in H : \\ \forall m \in \{1, \dots, M\} : \\ \mathcal{L}_m(h', \mathcal{D}_{val}) \leq \mathcal{L}_m(h, \mathcal{D}_{val}), \\ \exists j \in \{1, \dots, M\} : \\ \mathcal{L}_j(h', \mathcal{D}_{val}) < \mathcal{L}_j(h, \mathcal{D}_{val}) \end{array} \right\} . \quad (3)$$

Due to the problem of further distinguishing non-dominated models among each other, MO-ML algorithms usually resolve to returning the complete Pareto front of models instead of a single model. Thus, formally, the signature of an algorithm changes to $A : \mathbb{D} \times \Lambda \rightarrow 2^{\mathcal{H}}$. As a consequence, the evaluation of each hyperparameter configuration of such an MO-ML algorithm also involves quantifying the quality of a Pareto front of models instead of a single model returned by the algorithm, yielding a much more difficult problem. As such, the signature of our HPO loss function used in Eq. 2 also needs to change to $\mathcal{L} : 2^{\mathcal{H}} \times \mathbb{D} \rightarrow \mathbb{R}$ and thus can no longer be instantiated with simple loss functions such as accuracy.

One way of assessing the quality of a Pareto front and thus instantiating the HPO loss function \mathcal{L} in this case, are so-called Pareto front quality indicators (Audet et al. 2021). These can be categorized into so-called external and internal indicators where external indicators measure the convergence of a Pareto front to the optimal one, i.e., the one that the user prefers. Yet, in real-case problems, computing the entire Pareto front space is unfeasible and there is

no way to describe the desiderata beforehand. In contrast, internal indicators assess the Pareto front quality by measuring specific characteristics. A very common indicator is called hypervolume (Zitzler and Thiele 1999) quantifying the volume occupied by the Pareto front w.r.t. a reference point. Other indicators evaluate factors such as uniformity of solution distribution (e.g., spacing indicator (Schott 1995)), range of values covered by the Pareto front (e.g., maximum spread indicator (Zitzler, Deb, and Thiele 2000)), or proximity to specific threshold points (e.g., R2 indicator (Hansen and Jaszkiewicz 1994)). Even though internal indicators can be effectively used as a loss function in HPO of MO-ML algorithms to quantify the quality of Pareto front and thus of a configuration, choosing the measure leading to a Pareto front which has a desired shape requires deep expert knowledge and thus, is a hard task for the user. In particular, a user has to map their implicit desiderata for the Pareto front shape to properties of the quality indicators without a concrete way of tailoring the HPO approach to these desiderata.

Preference Learning

Preference learning (PL) (Fürnkranz and Hüllermeier 2010) is a subfield of machine learning dealing with the problem of learning from different forms of preferences. Although, PL in general comprises a large set of learning problems, we focus on the object ranking problem (Fürnkranz and Hüllermeier 2010) and in particular of learning to rank objects based on a given set of pairwise preferences.

In the object ranking problem, we consider a space of objects \mathcal{O} , where each object $o \in \mathcal{O}$ is represented as a feature vector f_o . Such an object can be an item, a document, or (in our case) a Pareto front. Preferences among two objects $o_i, o_j \in \mathcal{O}$ are denoted as $o_i > o_j$ indicating that object o_i is preferred over o_j . The corresponding space of pairwise rankings over \mathcal{O} is denoted as $\mathcal{R}(\mathcal{O})$. Then, given a dataset of the form $\mathcal{U} = \{o_{i,1} > o_{i,2}\}_{i=1}^U$, the goal is to learn a function $f : \mathcal{O} \times \mathcal{O} \rightarrow \mathcal{R}(\mathcal{O})$ that, given two objects, returns the correct pairwise ranking of these two objects.

Most approaches to object ranking work by learning a utility function $u : \mathcal{O} \rightarrow \mathbb{R}$ returning a utility score for an object that can be used to create a ranking of two objects by sorting them according to their utility score.

Related Work

In the following, we give a short overview of related work in the areas of human-centered HPO/AutoML and the usage of preference learning in AutoML-related fields.

The field of human-centered HPO/AutoML has gained increasing traction in the last years with approaches targeted at explaining the hyperparameter optimization process to increase the trust in automated tools (Pfisterer, Thomas, and Bischl 2019; Moosbauer et al. 2021, 2022; Segel et al. 2023). This also includes concrete tools developed to help a user interpret the results such as XAutoML (Zöller et al. 2022) or DeepCave (Sass et al. 2022).

Motivated from a similar stance that the user should be put back into the loop of the AutoML process to a certain extent, Francia, Giovanelli, and Pisano (2023) propose

to leverage structured argumentation to semi-automatically constrain search spaces based on user input. Further going into this direction and most similar to our approach is the work by Kulbach, Philipp, and Thoma (2020) building upon the observation that users are often unable to concretely configure a loss function in an AutoML tool fitting for their problem at hand. They suggest to learn a loss function customized to the user as a scalarization over several frequently used standard loss functions via PL methods. Our work differs from theirs in two main aspects: (1) First, we consider tuning the hyperparameters of MO-ML algorithms, while they try to find complete AutoML pipelines with standard classification/regression algorithms. Hence, there is not only a difference in the considered meta-problem (AutoML vs. HPO), but also in the types of machine learning algorithms to be composed / tuned (pipelines of standard ML algorithms vs MO-ML algorithms). (2) Second, we demonstrate pairwise comparisons of Pareto fronts to the user compared to pairwise comparisons of feature vectors, targets and predictions presented by Kulbach, Philipp, and Thoma (2020). We believe that our comparisons are much easier to make for a user. (3) Moreover, the Pareto front quality indicator learned by us is not constrained to be a scalarization of existing loss functions or quality indicators. Instead, we can, in principle, leverage any object ranking approach that allows for the extraction of a utility function.

Lastly, our approach is not to be confused with multi-objective HPO (Morales-Hernández, Nieuwenhuysse, and Gonzalez 2021; Karl et al. 2022), where HPO problem itself features multiple loss functions to be optimized, but the ML algorithm itself only returns a single solution.

Interactive Hyperparameter Optimization in Multi-Objective Problems

Our approach tackles the HPO problem (cf. Eq. 2) for MO-ML algorithms, i.e., learning algorithms that return a Pareto front of models instead of a single model as a result of the learning process, works in three phases as depicted in Fig. 1:

1. **Preliminary Sampling:** In the preliminary sampling phase we sample a fixed but small amount of hyperparameter configurations, evaluate these configurations and store the corresponding Pareto fronts of models returned by the MO-ML algorithm.
2. **Interactive Preference Learning:** In the interactive preference learning phase, we construct pairs of Pareto fronts from the ones obtained in the preliminary sampling phase and show these pairs to the user, who rates which of two shown Pareto fronts they prefer. Based on the pairwise preferences obtained from the user and feature representations of the Pareto fronts underlying these pairwise preferences, we learn a latent utility function, which, given a Pareto front, outputs a utility score.
3. **Utility-driven HPO:** In this HPO phase, we instantiate an HPO tool (e.g. SMAC (Hutter, Hoos, and Leyton-Brown 2011; Lindauer et al. 2022) as in our experiments) with the previously learned utility function as a loss function and perform standard HPO.

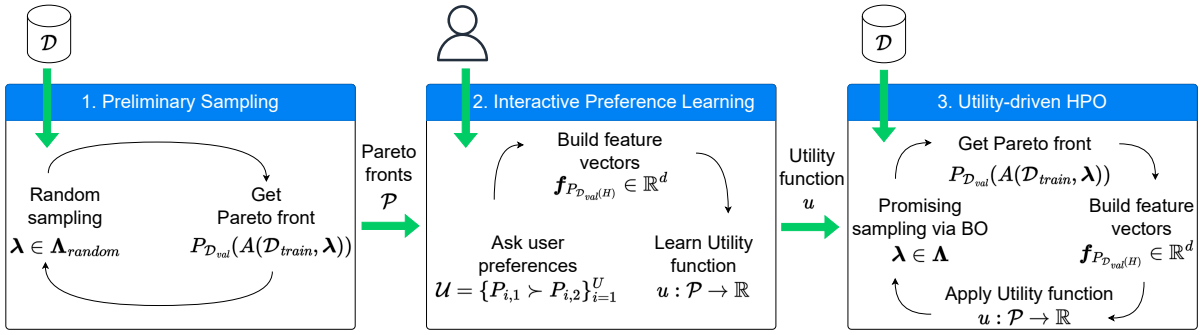


Figure 1: Overview of the three phases of our approach: Preliminary Sampling provides the user with different Pareto fronts, Interactive Preference Learning allows the user to express their preferences, finally, Utility-driven HPO guides the optimization to the user desiderata.

Preliminary Sampling

The underlying goal of the preliminary sampling phase is to obtain a set of Pareto fronts which we can use to construct pairs of Pareto fronts, which the user can rate in the subsequent stage. Keeping in mind, that we want to learn a utility function from the pairwise comparisons provided by the user, our set of Pareto fronts should ideally reasonably cover the space of possible Pareto fronts. This avoids potential generalization problems of the learned utility function, if it is provided with a Pareto front from a part of the Pareto front space which is very far off from any training data.

Recall that we perform HPO and our MO-ML algorithm $A : \mathbb{D} \times \Lambda \rightarrow 2^{\mathcal{H}}$ returns different Pareto fronts of models for a given dataset $\mathcal{D}_{train} \in \mathbb{D}$ based on different hyperparameter configurations $\lambda \in \Lambda$. As such, we can obtain a set of Pareto fronts of models by sampling a fixed number of hyperparameter configurations $\Lambda_{random} \subset \Lambda$ at random, training the algorithm instantiated with the corresponding hyperparameter configuration on the training data \mathcal{D}_{train} and evaluating it according to our loss functions \mathcal{L}_m on \mathcal{D}_{val} . This leads to a set of Pareto fronts defined as

$$\mathcal{P} = \{P_{\mathcal{D}_{val}}(A(\mathcal{D}_{train}, \lambda)) | \lambda \in \Lambda_{random}\}. \quad (4)$$

As the rich literature on AutoML and HPO shows, random search leads to a good coverage of the hyperparameter configuration space (Bischl et al. 2023). However, this does not automatically entail a reasonable coverage of the corresponding Pareto front space.

Interactive Preference Learning

The goal underlying this phase of our approach is construct pairs of Pareto fronts to show to the user and learn a utility function of the form $u : \mathcal{P} \rightarrow \mathbb{R}$, which, given a Pareto front, returns a utility score of the Pareto front based on the preferences obtained from the user.

Acquiring User Preferences More formally, we start by constructing a set of pairs of Pareto fronts as

$$\Omega = \{(P_1, P_2) | P_1, P_2 \in \mathcal{P}, P_1 \neq P_2\}, \quad (5)$$

which leads to a number of pairs quadratic in the number of samples Pareto fronts. Since we do not want to overwhelm

the user by showing them too many of such pairs, we can fallback to subsampling Ω to decrease the number of pairs to show to the user. However, as we show in the experimental evaluation later, we can also simply have a rather short preliminary sampling phase leading to a small number of Pareto fronts and thus also to a reasonably sized set Ω of pairs of Pareto fronts without compromising too much of the performance of our approach. Once the user has rated for each of such pairs $(P_1, P_2) \in \Omega$ whether he prefers the Pareto front P_1 over P_2 , we have an object ranking dataset of the form

$$\mathcal{U} = \{P_{i,1} > P_{i,2}\}_{i=1}^U, \quad (6)$$

where, without loss of generality, we assume that the user prefers Pareto front $P_{i,1}$ over Pareto front $P_{i,2}$.

Feature Representation of Pareto Fronts To learn a utility function based on the object ranking dataset defined in Eq. 6, we additionally require a feature representation of a Pareto front such that the object ranking learning algorithm we employ can generalize over unseen Pareto fronts. More precisely, we aim to encode the Pareto front $P_{\mathcal{D}_{val}(H)}$ returned by A in a d -dimensional feature representation $\mathbf{f}_{P_{\mathcal{D}_{val}(H)}} \in \mathbb{R}^d$ as depicted in Fig. 2. To this end, we assume to have access to all models returned by A , i.e. H including all dominated models internally learned by A and that there is some order imposed on the model space \mathcal{H} and hence on H . We evaluate all of these models $h_b \in H$ with each loss function \mathcal{L}_i on \mathcal{D}_{val} resulting in the following matrix

$$\mathbf{L} = \begin{pmatrix} \mathcal{L}_1(h_1, \mathcal{D}_{val}) & \dots & \mathcal{L}_M(h_1, \mathcal{D}_{val}) \\ \vdots & \ddots & \vdots \\ \mathcal{L}_1(h_B, \mathcal{D}_{val}) & \dots & \mathcal{L}_M(h_B, \mathcal{D}_{val}) \end{pmatrix}, \quad (7)$$

where $B \geq |H|$ is the maximal number of models returned by A . If A returns $B' < B$ values, we forward impute missing values for $B' + 1 \leq b \leq B' + (B - B')$ as

$$\mathbf{L}_{m,b} \leftarrow \mathcal{L}_m(h_{b-1}, \mathcal{D}_{val}) \quad \forall 1 \leq m \leq M. \quad (8)$$

For each model h_b we check if it is dominated by the previous model h_{b-1} as defined in Eq. 3. If it is dominated, its loss values in the matrix are replaced by the ones from the

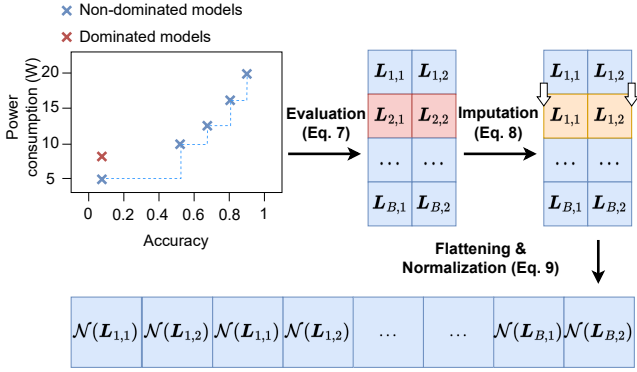


Figure 2: Visualization of the feature representation of a Pareto front based on two loss functions.

previous non-dominated model following Eq. 8. At the end of this process, this matrix only contains loss values of models contained in the Pareto front. Last, the matrix is flattened and standardized ($\mathcal{N}(\cdot)$) across Ω :

$$\mathbf{f}_{P_{\mathcal{D}_{\text{val}}(H)}} = [\mathcal{N}(L_{1,1}), \dots, \mathcal{N}(L_{B,M})]. \quad (9)$$

Note that we assume the order on our models to, on one hand ensure that the replacement and imputation strategy works as described above, and on the other hand works as a positional encoding which, assuming a suitable order relation is defined, in itself contains valuable domain knowledge. For our experimental evaluation, we define an ordering based on loss values of one of the considered loss functions.

Learning A Utility Function for Pareto Fronts With the object ranking dataset as defined in Eq. 6 and the feature representation previously described in Eq. 9, we can now learn the utility function. To this end, we leverage the RankSVM approach by Joachims (2002) which generalizes a standard SVM to the case of object ranking. This approach is very appealing for our use case as it allows to easily extract the latent utility function learned as part of the object ranker, which we want to use in the subsequent stage.

The idea underlying the (linear) RankSVM is that for every pairwise comparison $P_1 > P_2$ in our object ranking dataset \mathcal{U} we want that, without loss of generality, the hyperplane defined by the learned weight vector $\mathbf{w} \in \mathbb{R}^d$ separates P_1 and P_2 . Formally, it should hold for every pair:

$$P_1 > P_2 \Leftrightarrow \mathbf{w}^\top \mathbf{f}_{P_1} > \mathbf{w}^\top \mathbf{f}_{P_2} \quad (10)$$

$$\Leftrightarrow \mathbf{w}^\top (\mathbf{f}_{P_1} - \mathbf{f}_{P_2}) > 0 \quad (11)$$

$$\Leftrightarrow \mathbf{w}^\top (\mathbf{f}_{P_2} - \mathbf{f}_{P_1}) < 0. \quad (12)$$

As in the standard case of SVMs this problem is NP-hard as noted by (Joachims 2002), which can be solved by the standard problem relaxation leveraging slack variables as for normal SVMs. We refer to (Joachims 2002) for details.

With the observations formalized in Eq. 10-12 in mind, one can implement a RankSVM with a standard classification SVM trained with a dataset of the form

$$\mathcal{D}_{SVM} = \left\{ \left((\mathbf{f}_{P_1} - \mathbf{f}_{P_2}), 1 \right) \mid P_1 > P_2 \in \mathcal{U} \right\} \cup \left\{ \left((\mathbf{f}_{P_2} - \mathbf{f}_{P_1}), 0 \right) \mid P_1 > P_2 \in \mathcal{U} \right\}. \quad (13)$$

We add a positive example to encourage Eq. 11 and at the same time a negative example to encourage Eq. 12, which both enforce a balanced dataset.

After training a standard classification SVM on \mathcal{D}_{SVM} , we define the utility function $u : \mathcal{P} \rightarrow \mathbb{R}$ via the SVM weight vector \mathbf{w} leveraging the feature representation of a Pareto front $P \in \mathcal{P}$ defined in Eq. 9 as $u(P) = \mathbf{w}^\top \mathbf{f}_P$.

Although we only explained the linear RankSVM, the kernel trick (Schölkopf and Smola 2002) can be applied as usual in the case of SVMs leading to non-linear versions.

All in all, preference learning and in particular object ranking offers a way to quantify the quality of a Pareto front in terms of a single scalar value w.r.t. the user desired data without asking the user for this concrete scalar value. The RankSVM idea leverages the robustness and effectiveness of support vector machines to operationalize this idea while offering a clear and interpretable ranking mechanism.

Utility-driven HPO

As we learned the utility function $u : \mathcal{P} \rightarrow \mathbb{R}$ from the user preferences, we want to leverage it as a Pareto front quality indicator to optimize with a standard HPO tool. To this end, we leverage the well-known HPO tool SMAC (Hutter, Hoos, and Leyton-Brown 2011; Lindauer et al. 2022) instantiated with the learned utility function u as a cost function.

Evaluation

We evaluate our approach in a case-study related to decreasing the energy consumption of ML models and in particular consider an MO-ML algorithm optimizing for both accuracy and energy consumption showing how our approach can be used in the context of Green AutoML (Tornede et al. 2021). In the following, we first detail the general experimental setup and then present two experiments performed.

All code including detailed documentation and the technical appendix can be found on Github¹.

General Experimental Setup

The MO-ML algorithm, whose hyperparameters we want to tune, is a wrapper for a funnel-shaped deep neural network (DNN) learning algorithm that exposes the same hyperparameters as the underlying DNN learning algorithm, but performs a grid-search over the number of epochs to train the DNN to return a Pareto front of models (cf. Appendix D.1).

These models are all identical, but trained for different number of epochs and can thus be seen as snapshots of the DNN learning curve (Mohr and van Rijn 2022) at different epochs. A lower number of epochs usually leads to a lower energy consumption, but also to a lower accuracy indicating that the two objectives are potentially conflicting. The concrete hyperparameters our DNN algorithms exposes are defined by LCBench (Zimmer, Lindauer, and Hutter 2021), a well-known multi-fidelity deep learning benchmark, and given in Table 1 in the appendix. LCBench comprises evaluations of over 2000 funnel-shaped MLP neural networks with varying hyperparameters on 35 datasets of the OpenML

¹<https://github.com/automl/interactive-mo-ml>




Figure 3: Kendall’s Tau of the preference learning models.

CC-18 suite (Bischl et al. 2019). In the spirit of Green AutoML (Tornede et al. 2021), we leverage the benchmark surrogate for LCBench provided by YAHPO-GYM (Pfisterer et al. 2021).² We measure the accuracy of a model as the validation accuracy on 33% of the corresponding OpenML CC-18 dataset. We estimate the power consumption in watts (W) for training a model by assuming the maximum consumption for the whole provided training time. The evaluations present in LCBench were performed on an Intel Xeon Gold 6242 with a maximum consumption of 150 Wh .

As noted earlier, many users struggle to choose a Pareto front quality indicator as a loss function for an HPO tool that yields Pareto fronts with their desired shape. Our evaluation focuses on users that are likely to make a wrong decision wrt. choosing the correct quality indicator, but can very well label pairwise comparisons of Pareto fronts according to an indicator without knowing the indicator itself. As such, we simulate users by labeling pairwise comparisons according to hypervolume (HV), spacing (SP), maximum spread (MS) and R2 as Pareto front quality indicators. HV quantifies the volume of the front by merging the hypercubes determined by each of its models $h \in H$ and a reference point r (i.e., the worst possible model). HV values range from 0 to 1, where 1 is the optimal value. SP is one of the most popular uniformity indicators and gauges the variation of the distance between models in Pareto front. MS is a widely used spread indicator, which measures the range of a Pareto front by considering the maximum extent on each objective. R2 measures the proximity to a specific reference point r via the Chebyshev norm (formal definition in Appendix B).

The order over the models returned by A , which we assume for the feature representation, is given by the energy consumption loss function, which means that models are ordered based on their energy consumption.

Experiment: Object Ranking Performance

In the following, we evaluate our object ranker.

²We use the surrogate model to estimate the performance of those configurations for which no result is available in LCBench.

Additional Experimental Setup We tune the hyperparameters of our preference learning models, one for each quality indicator, on top of the first 3 datasets of LCBench: KDDCup09_appetency, covertype, Amazon_employee_access. Each configuration has been evaluated by averaging the performance achieved in those datasets over 3 different seeds. As to the evaluation within each dataset, we performed a cross-validation: we split the sampled Pareto fronts \mathcal{P} in 5 folds and compute all possible pairwise comparisons within each fold. Specifically, we set the number of sampled Pareto fronts $|\mathcal{P}|$ at 40, so that we can create 5 folds of 8 elements, which translates into $\binom{8}{2} = 28$ pairwise comparisons. Concretely, at each cross-validation evaluation, we use the pairwise comparisons of the 4 training folds to predict the global ranking of the 8 Pareto fronts in the test fold, and compare it with the ranking of the Pareto fronts given by the quality indicator at hand. These rankings are compared in terms of a ranking loss function called Kendall’s Tau correlation (Kendall 1948). Roughly speaking, this measures how correlated two rankings are. A correlation score of -1 can be interpreted as an inverse correlation, a score of 0 as no correlation and a score of 1 as perfect correlation.

Result Discussion In Fig. 3 we plot the performance of our models in terms of the Kendall’s Tau averaged across the folds with the error bars defined by the corresponding standard deviation. In particular, we show how this performance varies with the number of available pairwise comparisons, i.e. the size of the object ranking training data, highlighting that, depending on the quality indicator, we obtain a reliable ranking model for Pareto fronts with rather few training data.

Generally, as to be expected, the correlation scores increase for each utility function as the number of comparisons increases independent of the quality indicator based on which they are learned. However, both the score at the lowest number of comparisons and the degree to which it increases with a growing number of examples varies substantially between the different utility functions. This indicates that it is easier for our object ranker to learn a behavior similar to HV or R2 than to SP or MS which also coincides with the fact that HV and R2 are external quality indicators whereas SP and MS are internal indicators. A potential reason for these differences in modeling performance of our object ranker might be grounded in the feature representation of Pareto fronts we chose and in the linearity of the RankSVM used as an object ranker. In particular, both SP and MS require computations over pairs of models such that a quadratic kernel or feature transformation might be better suited for these indicators. Nevertheless, the ranking performance seems to be sufficient to guide the HPO process.

Experiment: HPO Approach Performance

In the following, we leverage the remaining 32 datasets of LCBench to evaluate our complete approach from phases 1 to 3. We will demonstrate that our HPO approach performs much better than SMAC assuming a user that chooses the wrong indicator and that it performs comparable assuming an advanced user knowing which indicator to pick.

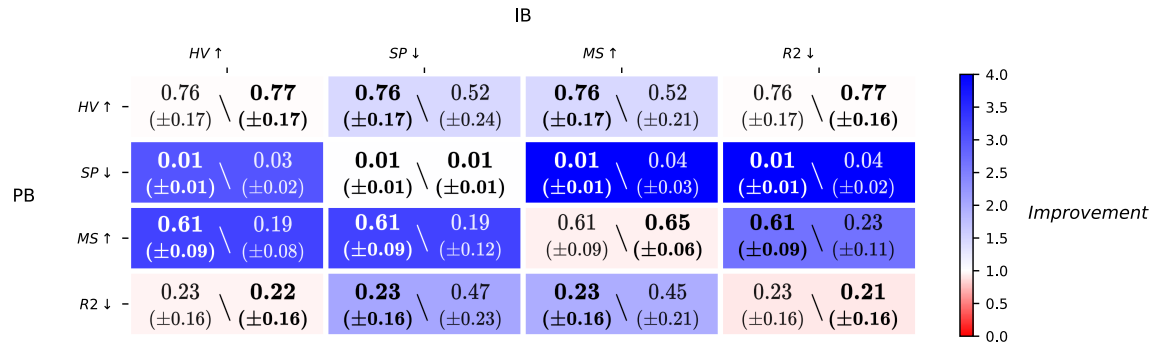


Figure 4: Comparison between PB approach, trained with 28 pairwise comparisons, and IB approach is denoted as $PB \setminus IB$. *Improvement* is computed as $\frac{PB}{IB}$ when the indicator in the row has to be maximized (\uparrow), and as $\frac{IB}{PB}$ otherwise (\downarrow).

Additional Experimental Setup In order to quantify how well our approach works, we compare how well SMAC instantiated with each of the Pareto front indicators above as a loss function (IB) works compared to our approach with the user simulated as mentioned above (PB). In particular, for each Pareto front indicator, we run our approach with a simulated user that behaves according to that indicator and compare against the HPO tool instantiated with each of the indicators as a loss function. That way, we cannot only quantify how much better our approach works wrt. to each of the quality indicators, under the assumption that a user chooses a wrong quality indicator, but also that our approach does not perform substantially worse in cases where the user picks the correct indicator. We run both IB and PB for a budget of 30 evaluations on each of the datasets for 3 seeds and report the mean and standard deviation over the seeds and datasets.

Result Discussion Figure 4 visualizes the comparison of SMAC initialized with different quality indicators as a loss function (IB) and our approach based on the learned utility as a quality indicator (PB). The indicators in the rows represent the ones used for labeling the user preferences, and hence, to train the preference learning models. The indicators in the columns represent the quality indicators chosen for optimization in SMAC. As a consequence, in each cell, we find the performance (averaged over seeds and datasets) and the respective standard deviation of the preference-based SMAC (PB) and the indicator-based SMAC (IB). This is expressed in terms of the quality indicator leveraged to rank the Pareto front (i.e. the one given in the row), hence, providing us with an estimation of how compliant the final Pareto front is with the user preferences. The cells in the diagonal correspond to situations where our approach is compared to IB initialized with the “correct” quality indicator function, whereas the off-diagonal cells correspond to scenarios where a user chooses the “wrong” quality indicator for IB. For better visual interpretability, we color cells with a blue tone depending on how much better our approach (PB) is relative to IB given in the respective column and red in cases where we are worse. Moreover, we highlight the better performance value for each cell in bold. If our learned utility function perfectly resembled the ground truth quality indicator, we would expect that the two values in each diag-

onal cell are identical and the coloring would be white as a consequence. Moreover, the better our approach is in one of the settings, the darker is the blue of the corresponding cell. At the same time, the worse our approach is in one of the settings, the darker will be the red in the corresponding cell.

As the figure shows, our approach (PB) behaves comparable to IB in cases where a user picks the correct Pareto front quality indicator (diagonal) and almost always better in cases where a user picks the wrong Pareto front quality indicator. In particular, we perform better or equal in 11/16 cases whereas IB performs slightly better only in 5/10 cases. In cases where our approach performs better, the improvements are often substantial, whereas in cases of degradation our approach is often only slightly worse.

Overall, our evaluation demonstrates that our approach successfully frees users from selecting the correct quality indicator aligned with their desiderata at the slight cost of visually comparing a low number of Pareto fronts.

Conclusion

In this paper, we propose a human-centered interactive HPO approach tailored towards MO ML leveraging preference learning to extract desiderata from users that guide the optimization. In particular, we learn a utility function for Pareto fronts based on pairwise preferences given by the user which we use as a loss function in a subsequent standard HPO process. In an experimental study, we demonstrate that our end-to-end approach performs much better than off-the-shelf HPO with a user that chooses an indicator not aligned with their desiderata and that it performs comparable to off-the-shelf HPO operated by an advanced user knowing which indicator to pick. As such, our approach successfully frees users from selecting the correct Pareto front quality indicator aligned with their desiderata at the slight cost of visually comparing a low number of Pareto fronts upfront.

As described in more detail in the limitations section in Appendix A, our approach naturally also offers room for future work. For example, we deem it interesting to design other feature representations with less assumptions and to generalize our approach to a larger number of loss functions as it is currently practically limited to two and users might have a hard time rating higher dimensional Pareto fronts.

Acknowledgements

Alexander Tornede and Marius Lindauer acknowledge funding by the European Union (ERC, “ixAutoML”, grant no.101041029). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Tanja Tornede was supported by the German Federal Ministry of the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (GreenAutoML4FAS project no. 67KI32007A).

References

- Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; and Koyama, M. 2019. Optuna: A Next-Generation Hyperparameter Optimization Framework. In *Proc. of KDD'19*, 2623–2631.
- Audet, C.; Bigeon, J.; Cartier, D.; Digabel, S. L.; and Salomon, L. 2021. Performance Indicators in Multiobjective Optimization. *EJOR*, 292(2): 397–422.
- Bischl, B.; Binder, M.; Lang, M.; Pielok, T.; Richter, J.; Coors, S.; Thomas, J.; Ullmann, T.; Becker, M.; Boulesteix, A.; Deng, D.; and Lindauer, M. 2023. Hyperparameter Optimization: Foundations, Algorithms, Best Practices, and Open Challenges. *Wiley IRDMKD*, e1484.
- Bischl, B.; Casalicchio, G.; Feurer, M.; Hutter, F.; Lang, M.; Mantovani, R.; van Rijn, J. N.; and Vanschoren, J. 2019. OpenML Benchmarking Suites and the OpenML100. *arXiv:1708.03731v1 [stat.ML]*.
- Deb, K.; Sindhya, K.; and Hakanen, J. 2016. Multi-objective Optimization. In *DSJ*, 161–200. Springer.
- Falkner, S.; Klein, A.; and Hutter, F. 2018. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In *Proc. of ICML'18*, 1437–1446.
- Feurer, M.; and Hutter, F. 2019. Hyperparameter Optimization. In Hutter, F.; Kotthoff, L.; and Vanschoren, J., eds., *Automated Machine Learning: Methods, Systems, Challenges*, chapter 1, 3 – 38. Springer. Available for free at <http://automl.org/book>.
- Francia, M.; Giovanelli, J.; and Pisano, G. 2023. HAMLET: A Framework for Human-centered AutoML via Structured Argumentation. *FGCS*, 142: 182–194.
- Frazier, P. 2018. A Tutorial on Bayesian Optimization. *arXiv:1807.02811 [stat.ML]*.
- Fürnkranz, J.; and Hüllermeier, E., eds. 2010. *Preference Learning*. Springer.
- Gunantara, N. 2018. A Review of Multi-objective Optimization: Methods and its Applications. *CE*, 5(1): 1502242.
- Hansen, M.; and Jaszkiewicz, A. 1994. *Evaluating the Quality of Approximations to the Non-dominated set*. IMM Technical Reports.
- Hutter, F.; Hoos, H.; and Leyton-Brown, K. 2011. Sequential Model-Based Optimization for General Algorithm Configuration. In *Proc. of LION'11*, 507–523.
- Hvarfner, C.; Stoll, D.; Souza, A.; Nardi, L.; Lindauer, M.; and Hutter, F. 2022. π BO: Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization. In *Proc. of ICLR'22*.
- Jin, Y., ed. 2006. *Multi-objective Machine Learning*. Springer.
- Joachims, T. 2002. Optimizing Search Engines Using Click-through Data. In *Proc. of KDD'02*, 133–142.
- Karl, F.; Pielok, T.; Moosbauer, J.; Pfisterer, F.; Coors, S.; Binder, M.; Schneider, L.; Thomas, J.; Richter, J.; Lang, M.; Garrido-Merchán, E.; Branke, J.; and Bischl, B. 2022. Multi-Objective Hyperparameter Optimization – An Overview. *arXiv:2206.07438 [cs.LG]*.
- Kendall, M. 1948. Rank Correlation Methods. *Hafner Publishing Co*.
- Komer, B.; Bergstra, J.; and Eliasmith, C. 2014. Hyperopt-Sklearn: Automatic Hyperparameter Configuration for Scikit-Learn. In *Proc. of SciPy'14*, 33–39.
- Kulbach, C.; Philipp, P.; and Thoma, S. 2020. Personalized Automated Machine Learning. In Giacomo, G. D.; Catalá, A.; Dilkina, B.; Milano, M.; Barro, S.; Bugarín, A.; and Lang, J., eds., *Proc. of ECAI'20*, 1246–1253.
- Lindauer, M.; Eggenberger, K.; Feurer, M.; Biedenkapp, A.; Deng, D.; Benjamins, C.; Ruhkopf, T.; Sass, R.; and Hutter, F. 2022. SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization. *JMLR*, 23(54): 1–9.
- Mallik, N.; Bergman, E.; Hvarfner, C.; Stoll, D.; Janowski, M.; Lindauer, M.; Nardi, L.; and Hutter, F. 2023. PriorBand: Practical Hyperparameter Optimization in the Age of Deep Learning. *arXiv:2306.12370 [cs.LG]*.
- Mohr, F.; and van Rijn, J. 2022. Learning Curves for Decision Making in Supervised Machine Learning—A Survey. *arXiv:2201.12150 [cs.LG]*.
- Moosbauer, J.; Casalicchio, G.; Lindauer, M.; and Bischl, B. 2022. Enhancing Explainability of Hyperparameter Optimization via Bayesian Algorithm Execution. *arXiv:2206.05447 [cs.LG]*.
- Moosbauer, J.; Herbinger, J.; Casalicchio, G.; Lindauer, M.; and Bischl, B. 2021. Explaining Hyperparameter Optimization via Partial Dependence Plots. In *Proc. of NeurIPS'21*.
- Morales-Hernández, A.; Nieuwenhuys, I. V.; and Gonzalez, S. 2021. A survey on multi-objective hyperparameter optimization algorithms for Machine Learning. *arXiv:2111.13755 [cs.LG]*.
- Pfisterer, F.; Schneider, L.; Moosbauer, J.; Binder, M.; and Bischl, B. 2021. YAHPO Gym – Design Criteria and a new Multifidelity Benchmark for Hyperparameter Optimization. *arXiv:2109.03670 [cs.LG]*.
- Pfisterer, F.; Thomas, J.; and Bischl, B. 2019. Towards Human Centered AutoML. *arXiv:1911.02391 [cs.HC]*.
- Salinas, D.; Seeger, M.; Klein, A.; Perrone, V.; Wistuba, M.; and Archambeau, C. 2022. Syne Tune: A Library for Large Scale Hyperparameter Tuning and Reproducible Research. In *Proc. of AutoML Conf'22*. PMLR.

- Sass, R.; Bergman, E.; Biedenkapp, A.; Hutter, F.; and Lindauer, M. 2022. DeepCAVE: An Interactive Analysis Tool for Automated Machine Learning. In *ICML ReALML Workshop*.
- Schölkopf, B.; and Smola, A. 2002. *Learning With Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press.
- Schott, J. 1995. Fault Tolerant Design Using Single and Multicriteria Genetic Algorithm Optimization. *MIT Libraries*.
- Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2019. Green AI. *arXiv:1907.10597v3 [cs.CY]*.
- Segel, S.; Graf, H.; Tornede, A.; Bischl, B.; and Lindauer, M. 2023. Symbolic Explanations for Hyperparameter Optimization. In *Proc. of AutoML Conf'23*. PMLR.
- Souza, A.; Nardi, L.; Oliveira, L.; Olukotun, K.; Lindauer, M.; and Hutter, F. 2021. Bayesian Optimization with a Prior for the Optimum. In *Proc. of ECML/PKDD'21*, 265–296.
- Tornede, T.; Tornede, A.; Hanselle, J.; Mohr, F.; Wever, M.; and Hüllermeier, E. 2021. Towards Green Automated Machine Learning: Status Quo and Future Directions. *JAIR*, 77: 427–457.
- van Wynsberghe, A. 2021. Sustainable AI: AI for Sustainability and the Sustainability of AI. *AIEJ*, 1(3): 213–218.
- Zimmer, L.; Lindauer, M.; and Hutter, F. 2021. AutoPytorch: Multi-Fidelity MetaLearning for Efficient and Robust AutoDL. *TPAMI*, 43: 3079–3090.
- Zitzler, E.; Deb, K.; and Thiele, L. 2000. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *EvoComp*, 8(2): 173–195.
- Zitzler, E.; and Thiele, L. 1999. Multiobjective Evolutionary Algorithms: a Comparative Case Study and the Strength Pareto Approach. *IEEE TEVC*, 3(4): 257–271.
- Zöller, M.; Titov, W.; Schlegel, T.; and Huber, M. F. 2022. XAutoML: A Visual Analytics Tool for Establishing Trust in Automated Machine Learning. *arXiv:2202.11954 [cs.LG]*.