A Batch Bayesian Approach for Bilevel Multi-Objective Decision Making Under Uncertainty

Vedat Dogan, Steven Prestwich
Confirm Centre for Smart Manufacturing
School of Computer Science and Information Technology,
University College Cork, Ireland
{vedat.dogan,steven.prestwich}@cs.ucc.ie

Abstract

Bilevel multiobjective optimization is a field of mathematical programming representing a nested hierarchical decision making process, with one or more decision makers at each level. These problems appear in many practical applications, solving tasks such as optimal control, process optimization, governmental and game playing strategy development, and transportation. Uncertainty cannot be ignored in these practical problems. We present a hybrid algorithm called BAMBINO, based on batch Bayesian approach via expected hypervolume improvement, that can handle uncertainty in the upper level. Three popular modified benchmark problems with multiple dimensions are used to evaluate its performance under objective noise compared to two popular algorithms in the literature. The results show that BAMBINO is computationally efficient and able to handle upper level uncertainty. We also evaluate the effect of batch size on performance.

Introduction

Hierarchical decision making has an extensive history, in Game Theory as first realized by von Stackelberg (Stackelberg 1952) and in the subfield of mathematical programming called bilevel optimization (Bracken and McGill 1973). A bilevel optimization problem contains a nested *inner* optimization problem which is a constraint of an *outer* optimization problem. The outer optimization task is referred to as the *upper level* or *leader* while the inner optimization problem is referred as the *lower level* or *follower*. Existing bilevel research has mainly focused on single-objective leaders and followers. *Multiobjective* bilevel optimization is relatively neglected, but there is work in the fields of classical optimisation (Eichfelder 2010) and evolutionary computation (Islam, Singh, and Ray 2016).

Our work focuses on the special case of multiobjective bilevel problems in which the leader has noisy objectives. We assume that the follower is free to choose any feasible solution from a Pareto-optimal set. We use batch Bayesian optimization to improve efficiency, approximating the leader's Pareto-front using fewer function evaluations than existing works. We also present a black-box approach to the noisy leader's objectives for handling the uncertainty during decision making.

Hierarchical decision making under uncertainty with noisy objectives becomes more interesting in a bilevel structure. The follower can observe the leader's decisions but the leader may have no idea how the follower is going to respond. Previously observed decisions are therefore important. Uncertainty in the objective also complicates the leader's decision making, and our algorithm uses a specifically designed acquisition function called qNEHVI to maximize expected hypervolume improvement under noisy objectives. We call our algorithm BAMBINO (**B**ayesian **A**pproach for **M**ultiobjective **Bi**level Problems with **N**oisy **O**bjectives). To evaluate its performance we consider three multi-dimensional test problems from two different suites of multiobjective bilevel optimization problems. Both examples illustrate the importance of taking uncertainty into account.

Motivation

Most studies in the multiobjective bilevel optimization literature focus on solving the optimization problem without addressing the impact of uncertainty. In practical problems, noise in the leader's objectives might represent environmental uncertainty, for example in a meta-learning regime (Al-Shedivat et al. 2017) that can be mathematically formulated as bilevel programming (Franceschi et al. 2018). As another example, a government might need to prevent terrorist attacks using information from unreliable sources. Yet another example occurs in computing optimal recovery policies for financial markets (Mannino, Bernt, and Dahl 2012), using bilevel optimization with objective uncertainties caused by several uncontrollable parameters.

Bilevel optimization problems are computationally expensive to solve because of their nested structure, and they become even more complex when there are multiple objectives and uncertainty (possibly at both levels). The main purpose of our work is to improve the efficiency of solving multiobjective bilevel optimization while handling leader objective uncertainties.

Background

We now provide some necessary background.

Bilevel Multiobjective Optimization Problems (BMOP). Because of the nature of multiobjective optimization prob-

Algorithm 1: BAMBINO

Inputs: $\mathbf{F}_u(\mathbf{x}_u, \mathbf{x}_l) : \mathbf{x_u} \in \mathbb{X}_u, \mathbf{x_l} \in \mathbb{X}_l$,
Batch points per epoch Q,
Total epoch N,
Reference point

$\mathbf{x_l}$: Find the Best Lower Level response as parameters with NSGA-II algorithm,
Initial decision data set with the objective noise $D = \{\mathbf{x}_{u_i}, \mathbf{F}_u(\mathbf{x}_{u_i}, \mathbf{x}_{l_i}), (\Sigma_i)\}_{i=1}^{n}$ with size of n,
Initialize the $\mathcal{GP}$ model with the observations and the objective noise
for i = 0 : N **do**
    Suggest new q-batch points by optimizing *qNEHVI*
    for j = 0 : q **do**
        For each upper-level decision $\mathbf{x}_u$, find optimal $\mathbf{x}_l^*$ by applying the NSGA-II
        Calculate fitness scores with noise $\mathbf{F}_u(\mathbf{x}_u, \mathbf{x}_l^*) + \xi$
    end for
    Update the data set D with new observations
end for
Update $\mathcal{GP}$ model with new observations
Return Pareto-front $\mathbf{F}_u^*$ and $(\mathbf{x}_u, \mathbf{x}_l^*)$

lems, only Pareto-optimal solutions at the lower level can be considered as feasible solutions for the upper level problem. We denote the upper level decision variables by $x_u \in X_u \subset \mathbb{R}^n$ and the lower level decision variables by $x_l \in X_l \subset \mathbb{R}^m$. The lower level problem is solved with respect to x_l while the upper level problem is solved with respect to both decisions $x = (x_u, x_l)$. Each x_u corresponds to a different lower level optimization problem with a different Pareto-front decision set. The lower level Pareto-front is defined as $P^* = \{f(\mathbf{x}_u, \mathbf{x}_l) : \mathbf{x}_l \in \mathbf{X}_l, \nexists \mathbf{x}'_l \in X_l \text{ s.t. } f(x') \succ f(x)\}$ where $f(x') \succ f(x)$ denotes $f(x')$ *dominates* $f(x)$. The Pareto-optimal decision set is $X_l^* = \{\mathbf{x}_l^* : f(\mathbf{x}_u, \mathbf{x}_l^*) \in P^*\}$. The definition of bilevel multiobjective problem with vector valued decision variables $\mathbf{x_u}$ and $\mathbf{x_l}$ is given by

$$\begin{aligned} &\underset{\mathbf{x}_u,\mathbf{x}_l}{\text{minimize}} \{F_1(\mathbf{x}_u, \mathbf{x}_l), ..., F_p(\mathbf{x}_u, \mathbf{x}_l)\} \\ &\text{subject to} \\ &\mathbf{x}_l \in \underset{\mathbf{x}_l}{\text{argmin}} \{f_1(\mathbf{x}_u, \mathbf{x}_l), ..., f_q(\mathbf{x}_u, \mathbf{x}_l); \\ &\qquad\qquad g_j(\mathbf{x}_u, \mathbf{x}_l) \leq 0,\ j = 1, 2, \ldots, J\} \\ &G_k(\mathbf{x}_u, \mathbf{x}_l) \leq 0,\ k = 1, 2, \ldots, K \end{aligned} \tag{1}$$

where $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^p$ represents the upper level function and $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^q$ represents the lower level function of the bilevel problem. Upper level and lower level constraints are defined by $G_k : X_u \times X_l \to \mathbb{R}$ and $g_j : X_u \times X_l \to \mathbb{R}$ for $k = 1, \ldots, K$ and $j = 1, \ldots, J$.

Bayesian Optimization (BO). BO is a sample-efficient approach that has demonstrated great potential in approximating a global optimum with a relatively small number of function evaluations. It uses a probabilistic surrogate model to make decisions by balancing exploration and exploitation (Shahriari et al. 2016). Gaussian process ($\mathcal{GP}$) is a common surrogate model with a flexible and non-parametric form. $\mathcal{GP}$ provides a posterior distribution for a decision point x in the search space by capturing the prior belief about the performances of unknown objective function, using a mean function $\mu(\mathbf{x})$ and a kernel function $k(\mathbf{x}_i, \mathbf{x}_j)$. BO uses an acquisition function to decide which point to choose next. The acquisition function specifies the value of the next point by using the surrogate's predictive distribution at the current point. We assume that the black-box function f is expensive to evaluate, but that optimizing the acquisition function is relatively cheap and fast.

Multiobjective Bayesian optimization (MOBO) combines the Bayesian surrogate model and an acquisition function specifically designed for multiobjective optimization problems such as qNEHVI (Daulton, Balandat, and Bakshy 2021). This is a hypervolume improvement based acquisition function that works well for noisy multiobjective optimization problems.

Method

We consider a case that is crucial in practice, in which the leader must make decisions under uncertainty based on noisy observations $\mathbf{F}_i = \mathbf{f}(\mathbf{x}_{u_i}, \mathbf{x}_{l_i}) + \xi_i$ where $\xi_i \sim \mathcal{N}(0, \Sigma_i)$ and Σ_i is the noise covariance and $\mathbf{x}_u, \mathbf{x}_l$ are upper and lower decision variables respectively. We reformulate leader's objective with noisy observations as

$$\underset{\mathbf{x}_u,\mathbf{x}_l}{\text{minimize}} \{F_1(\mathbf{x}_u, \mathbf{x}_l) + \xi_1, ..., F_p(\mathbf{x}_u, \mathbf{x}_l) + \xi_p\} \tag{2}$$

where $\xi_p \sim \mathcal{N}(0, \Sigma_p)$. The hypervolume indicator measures the volume of space between the non-dominated front and a reference point, which we assume is known by the upper level decision maker. The selection of reference point is tricky. In this work it is chosen to be an extreme point of the Pareto front, because reference points should be dominated by all Pareto-optimal solutions. Hypervolume improvement of a set of points $\mathcal{P}'$ is defined as $HVI(\mathcal{P}'|\mathcal{P}, \mathbf{r}) = HV(\mathcal{P} \cup \mathcal{P}'|\mathbf{r}) - HV(\mathcal{P}|\mathbf{r})$ where $\mathcal{P}$ represents the Pareto front and $\mathbf{r}$ the reference point. Given observations of the upper level decision making process, the $\mathcal{GP}$ surrogate model provides us with a posterior distribution over the upper level function values for each observation. These values can be used to compute the expected hypervolume improvement acquisition function defined by

$$\alpha_{ehvi}(\mathbf{x}_u)|\mathcal{P} = \mathbb{E}[HVI(\mathbf{F}_u|\mathcal{P})] \tag{3}$$

So the expected hypervolume improvement iterates over the posterior distribution, an approach that worked well in (Daulton, Balandat, and Bakshy 2021).

After n observations of the leader's decisions and the follower's response, the posterior distribution can be defined by the conditional probability $p(\mathbf{F}(\mathbf{x}_{u_n}, \mathbf{x}_{l_n})|\mathcal{D}_n)$ of the leader's objective values given decision variables $(\mathbf{x}_{u_n}, \mathbf{x}_{l_n})$ based on noisy observations $\mathcal{D}_n = \{\mathbf{x}_{u_i}, \mathbf{F}_i(\mathbf{x}_{u_i}, \mathbf{x}_{l_i}), (\Sigma_i)\}_{i=1}^{n}$. NEHVI is defined as

$$\alpha_{NEHVI}(\mathbf{x}_u) = \int \alpha_{ehvi}(\mathbf{x}_u|P_n) p(\mathbf{F}|\mathcal{D}_n) d\mathbf{F} \tag{4}$$

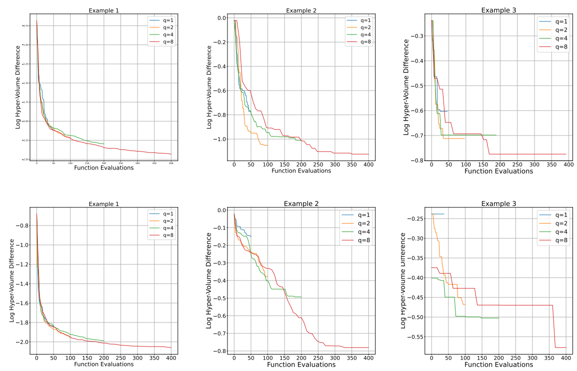

Figure 1: Hypervolume difference graph (log scale) with different batch sizes ($q = 1, q = 2, q = 4, q = 8$) for Example 1 with 15 (*top-left*) and 20 (*bottom-left*) dimensions, for Example 2 with 10 (*top-middle*) and 20 (*bottom-middle*) dimensions, for Example 3 with 10 (*top-right*) and 20 (*bottom-right*) dimensions.

where P_n denotes the Pareto-front optimal decision set over the leader's objectives $\mathbf{F}_n$. The aim is to improve the efficiency of the optimization, and the handling of noise in the leader's objective, by using the approach above and reformulating the bilevel multiobjective optimization problem. The algorithm details can be found in Algorithm 1.

Numerical Experiments

The test problems are selected from the literature (Deb and Sinha 2009b), with the aim of testing scalability in terms of decision variable dimensionality. The results are compared with state-of-art evolutionary algorithms m-BLEAQ (Sinha et al. 2016) and H-BLEMO (Deb and Sinha 2010). The Pareto-front is independent of the parameters.

Performance Metrics. We compare our results in terms of upper level function evaluations (FE) for the efficiency of the algorithm. Hypervolume improvement (HV) (Fonseca, Paquete, and Lopez-Ibanez 2006) and inverted generational distance (IGD) are also used to evaluate the success of approximation to Pareto-fronts, in terms of convergence and diversity. HV measures the volume of the space between the non-dominated front obtained and a reference point. IGD calculates the sum of the distances from each point of the true Pareto-front to the nearest point of the non-dominated set found by the algorithm. Therefore, smaller IGD value means approximated points are closer to the Pareto-front of the problem.

Parameters. In our experiments we use a number of iterations $N = 50$, and compute median FE from 21 runs. We use the independent $\mathcal{GP}$ model with *Matern52* kernel and fit the $\mathcal{GP}$ by maximizing the marginal log-likelihood. The method initialized with $2 \times (d+1)$ Sobol points where d represents the dimension of the problem to construct the initial $\mathcal{GP}$ model. All experiments are conducted using BoTorch (Balandat et al. 2019) library. We solved the follower's problem with the popular *non-dominated sorted genetic algorithm* (NSGA-II) (Deb et al. 2002) and choose the population size 100 and number of generations 200. We choose the follower's decisions from the obtained Pareto-front at random, as all solutions in the Pareto-front are feasible.

Example 1. The first example is a bi-objective problem that is scalable in terms of the number of follower decision variables. We choose $K = 14$ and $K = 19$, giving 15 and 20 follower variables respectively, with 1 leader decision variable. We choose the reference point required for measuring hypervolume improvement to be $(1.0, 0.5)$. The formulation of the problem is given by

$$\begin{aligned}
&\underset{(x_u,\mathbf{x_l})}{\text{Min}}\ \mathbf{F}(x_u,\mathbf{x_l}) = \begin{pmatrix} (x_{l_1}-1)^2 + \sum_{i=2}^{K} x_{l_i} + \\ (x_u)^2 + \xi \\ (x_{l_1}-1)^2 + \sum_{i=2}^{K} x_{l_i} + \\ (x_u-1)^2 + \xi \end{pmatrix} \\
&\text{subject to} \\
&\mathbf{x}_l \in \underset{\mathbf{x_l}}{\text{argmin}}\ \mathbf{f}(x_u,\mathbf{x_l}) = \begin{pmatrix} x_{l_1} + \sum_{i=2}^{K} x_{l_i}^2 \\ x_{l_1} - x_u + \sum_{i=2}^{K} x_{l_i}^2 \end{pmatrix} \\
&-1 \leq (x_u, x_{l_1}, x_{l_2}, ..., x_{l_K}) \leq 2 \\
&\xi \sim \mathcal{N}(0,\Sigma_\xi),\ \Sigma_\xi = \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}
\end{aligned} \tag{5}$$

The Pareto-optimal decision sets for this specific bilevel decision-making problem can be found in (Sinha et al. 2016).

Example 2. The second test problem is the modified test problem with 10 and 20 variable instances. We choose the required reference point to be $(1.1, 1.1)$. The formulation of the problem is given by:

$$\begin{aligned}
&\underset{(\mathbf{x_u},\mathbf{x_l})}{\text{Min}}\ \mathbf{F}(\mathbf{x_u},\mathbf{x_l}) = \\
&\begin{pmatrix} (1 + r - \cos(\alpha\pi x_{u_1})) + \sum_{j=2}^{K}(x_{u_j} - \frac{j-1}{2})^2 + \\ \tau\sum_{i=2}^{K}(x_{l_i} - x_{u_i})^2 - r\cos(\gamma\frac{\pi}{2}\frac{x_{l_1}}{x_{u_1}}) + \xi \\ (1 + r - \sin(\alpha\pi x_{u_1})) + \sum_{j=2}^{K}(x_{u_j} - \frac{j-1}{2})^2 + \\ \tau\sum_{i=2} K(x_{l_i} - x_{u_i})^2 - r\sin(\gamma\frac{\pi}{2}\frac{x_{l_1}}{x_{u_1}} + \xi) \end{pmatrix} \\
&\text{subject to} \\
&\mathbf{x}_l \in \underset{\mathbf{x_l}}{\text{argmin}} \\
&\mathbf{f}(\mathbf{x_u},\mathbf{x_l}) = \begin{pmatrix} x_{l_1}^2 + \sum_{i=2}^{K}(x_{l_i} - x_{u_i})^2 + \\ \sum_{i=2}^{K} 10(1 - \cos(\frac{\pi}{K}(x_{l_i} - x_{u_i})) + \xi \\ \sum_{i=2}^{K}(x_{l_i} - x_{u_i})^2 + \\ \sum_{i=2}^{K} 10|\sin(\frac{\pi}{K}(x_{l_i} - x_{u_i}))| + \xi \end{pmatrix} \\
&x_{l_i} \in [-K, K], i = 1, \ldots, K \\
&x_{u_1} \in [1, 4], x_{u_j} \in [-K, K], j = 2, \ldots, K \\
&\xi \sim \mathcal{N}(0,\Sigma_\xi),\ \Sigma_\xi = \begin{bmatrix} 0.25 & 0 \\ 0 & 0.16 \end{bmatrix}
\end{aligned} \tag{6}$$

The Pareto-front for a given leader is defined as a circle of radius $(1 + r)$ with centre $((1 + r), (1 + r))$. We choose

$K = 5$ for our experiments with parameters $r = 0.1, \tau = 1$ and $\alpha = 1$, following (Sinha et al. 2016) so that our results can be compared with those for m-BLEAQ and H-BLEMO.

Example 3. The third test problem is the modified test problem with 10 and 20 variable instances. We choose the required reference point $(0.8, 0.0)$ for measuring the hypervolume improvement during the optimization. The formulation of the problem is given by:

$$\underset{(\mathbf{x_u},\mathbf{x_l})}{\text{Min}} \ \mathbf{F}(\mathbf{x_u},\mathbf{x_l}) =$$

$$\begin{pmatrix} v_1(x_{u_1}) + \sum_{j=2}^{K}(x_{u_j}^2 + 10(1-\cos(\frac{\pi}{K}x_{u_i}))+ \\ \tau\sum_{i=2}^{K}(x_{l_i}-x_{u_i})^2 - r\cos(\gamma\frac{\pi}{2}\frac{x_{l_1}}{x_{u_1}}) + \xi \\ v_2(x_{u_1}) + \sum_{j=2}^{K}(x_{u_j}^2 + 10(1-\cos(\frac{\pi}{K}x_{u_i}))+ \\ \tau\sum_{i=2}^{K}(x_{l_i}-x_{u_i})^2 - r\sin(\gamma\frac{\pi}{2}\frac{x_{l_1}}{x_{u_1}}) + \xi \end{pmatrix}$$

where

$$v_1(x_{u_1}) = \begin{cases} \cos(0.2\pi)x_{u_1} + \\ \quad \sin(0.2\pi)\sqrt{|0.02\sin(5\pi x_{u_1})|}, \\ \text{for } 0 \le x_{u_1} \le 1. \\ x_{u_1} - (1-\cos(0.2\pi)), \quad \text{for } x_{u_1} \ge 1. \end{cases}$$

$$v_2(x_{u_1}) = \begin{cases} -\sin(0.2\pi)x_{u_1} + \\ \quad \cos(0.2\pi)\sqrt{|0.02\sin(5\pi x_{u_1})|}, \\ \text{for } 0 \le x_{u_1} \le 1. \\ 0.01(x_{u_1}-1) - \sin(0.2\pi), \quad \text{for } x_{u_1} \ge 1. \end{cases}$$

subject to

$$\mathbf{x}_l \in \underset{\mathbf{x_l}}{\text{argmin}}$$

$$\mathbf{f}(\mathbf{x_u},\mathbf{x_l}) = \begin{pmatrix} x_{l_1}^2 + \sum_{i=2}^{K}(x_{l_i}-x_{u_i})^2 + \xi \\ \sum_{i=2}^{K} i(x_{l_i}-x_{u_i})^2 + \xi \end{pmatrix}$$

$$x_{l_i} \in [-K, K], i = 1, \ldots, K$$

$$x_{u_1} \in [0.001, K], x_{u_j} \in [-K, K], j = 2, \ldots, K$$

$$\xi \sim \mathcal{N}(0, \Sigma_\xi), \ \Sigma_\xi = \begin{bmatrix} 0.09 & 0 \\ 0 & 0.09 \end{bmatrix} \tag{7}$$

Details on the Pareto-optimal solutions are given in (Deb and Sinha 2009a).

Results and Discussion

The performance of BAMBINO is compared with that of m-BLEAQ and H-BLEMO in Table 1, showing computational expense and convergence. The FE is calculated by $N_{initial} + (N_{batch} \times N_{iter} \times N_{restarts})$ for the leader problem, where $N_{initial}$ is the number of initial decisions for starting the algorithm. We choose it to be $2 \times (d+1)$ where d is dimensions of the decision variable. We run the experiment for different batch numbers $q = 1, q = 2, q = 4, q = 8$ to test the effect on performance. The HV difference is shown in Figure 1 for 15 and 20 variables, and 10 variables for Examples 1 and 2 respectively. Because of lack of information in the reference paper, we could not obtain the FE results for Example 1 with 20 variables.

Example 1. We can see from the Table 1 that the required upper level FE is significantly lower, while the algorithm approximates successfully to the Pareto-front while handling the uncertainty at leader's objective. For 15 variables BAMBINO achieves $\approx$ %38 improvement in terms of FE compared to m-BLEAQ and $\approx$ %89 compared to H-BLEMO. The IGD values in Table 1 for 15 and 20 variables show that BAMBINO successfully approximates to the Pareto-front of the problem while handling the uncertainty of leader's objective for both. We show the HV difference between Pareto-front solutions and approximated BAMBINO decisions algorithm in Figure 1. Again we tried different batch sizes for the experiment, and it can be seen that batch number of 8 is best for this specific example at both dimensions. We could not compare the 20 dimensional version of the problem with the selected algorithms because of the lack of information in (Sinha et al. 2016).

Example 2. Table 1 shows that BAMBINO obtains the best IGD results compared to the other algorithms. In terms of FE it significantly improves the state of the art, with $\approx$ %81 improvement for 10 variables and $\approx$ %88 improvement for 20 variables compared to m-BLEAQ. We also show the HV difference in Figure 1 and we can observe that, for this specific example, the batch number of 8 is the best selection for both 10 and 20 dimensional version.

Example 3. Table 1 shows that BAMBINO obtains the best IGD value compared to m-BLEAQ and H-BLEMO while improving efficiency in terms of FE: $\approx$ 84% and $\approx$ 97% with 10 variables, and $\approx$ 89% and $\approx$ 98% with 20 variables. Figure 1shows that batch size $q = 8$ gives best results. In summary, BAMBINO is successful on fairly high-dimensional problems, handling noisy objectives with less computational cost.

Noise in leader objective makes the problem harder to solve but more realistic for modelling practical problems, because of real-world uncertainty. We show the proposed BAMBINO algorithm works well on these test benchmark problems. We believe that BAMBINO can be applied to several practical bilevel problems applied successfully in machine learning community such as image classification (Mounsaveng et al. 2020), deep learning (Han et al. 2022), neural networks (Li and Zhang 2021), neural achitecture search and hyperparameter optimization (Liu et al. 2021).

Conclusions

In this paper we discussed bilevel multiobjective optimization under upper level uncertainty, and presented a hybrid algorithm called BAMBINO, based on batch Bayesian optimization with hypervolume improvement. We ran experiments using three benchmark problems, and BAMBINO performed very competitively in terms of computational efficiency and convergence. We also showed how batch size selection affects performance in terms of hypervolume improvement.

	Number of Variables	BAMBINO		m-BLEAQ		H-BLEMO	
		IGD	FE	IGD	FE	IGD	FE
Example 1	15	0.0044	4032	0.0013	6,464	0.0046	39,818
Example 2	10	0.0051	4022	0.0069	22,223	0.0134	106,003
Example 3	10	0.0076	4022	0.0079	25,364	0.0134	132,907
Example 1	20	0.0105	4042	-	-	-	-
Example 2	20	0.1032	4042	0.0435	34,110	0.1106	191,357
Example 3	20	0.0924	4042	0.0623	36,439	0.1321	216,083

Table 1: FE and IGD values for the examples with the number of variable dimensions.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number 16/RC/3918 which is co-funded under the European Regional Development Fund. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

Al-Shedivat, M.; Bansal, T.; Burda, Y.; Sutskever, I.; Mordatch, I.; and Abbeel, P. 2017. Continuous Adaptation via Meta-Learning in Nonstationary and Competitive Environments. *CoRR*, abs/1710.03641.

Balandat, M.; Karrer, B.; Jiang, D. R.; Daulton, S.; Letham, B.; Wilson, A. G.; and Bakshy, E. 2019. BoTorch: Programmable Bayesian Optimization in PyTorch. *CoRR*, abs/1910.06403.

Bracken, J.; and McGill, J. T. 1973. Mathematical Programs with Optimization Problems in the Constraints. *Operations Research*, 21(1): 37–44.

Daulton, S.; Balandat, M.; and Bakshy, E. 2021. Parallel Bayesian Optimization of Multiple Noisy Objectives with Expected Hypervolume Improvement. In *NeurIPS*.

Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2): 182–197.

Deb, K.; and Sinha, A. 2009a. Constructing test problems for bilevel evolutionary multi-objective optimization. In *2009 IEEE Congress on Evolutionary Computation*, 1153–1160.

Deb, K.; and Sinha, A. 2009b. Solving Bilevel Multi-Objective Optimization Problems Using Evolutionary Algorithms. In *EMO*.

Deb, K.; and Sinha, A. 2010. An Efficient and Accurate Solution Methodology for Bilevel Multi-Objective Programming Problems Using a Hybrid Evolutionary-Local-Search Algorithm. *Evolutionary computation*, 18: 403–49.

Eichfelder, G. 2010. Multiobjective bilevel optimization. *Math. Program.*, 123: 419–449.

Fonseca, C.; Paquete, L.; and Lopez-Ibanez, M. 2006. An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator. In *2006 IEEE International Conference on Evolutionary Computation*, 1157–1163.

Franceschi, L.; Frasconi, P.; Salzo, S.; Grazzi, R.; and Pontil, M. 2018. Bilevel Programming for Hyperparameter Optimization and Meta-Learning. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1568–1577. PMLR.

Han, Y.; Liu, J.; Xiao, B.; Wang, X.; and Luo, X. 2022. Bilevel Online Deep Learning in Non-stationary Environment. *CoRR*, abs/2201.10236.

Islam, M. M.; Singh, H. K.; and Ray, T. 2016. A Nested Differential Evolution Based Algorithm for Solving Multi-objective Bilevel Optimization Problems. In *ACALCI*.

Li, H.; and Zhang, L. 2021. A Bilevel Learning Model and Algorithm for Self-Organizing Feed-Forward Neural Networks for Pattern Classification. *IEEE Transactions on Neural Networks and Learning Systems*, 32: 4901–4915.

Liu, R.; Gao, J.; Zhang, J.; Meng, D.; and Lin, Z. 2021. Investigating Bi-Level Optimization for Learning and Vision from a Unified Perspective: A Survey and Beyond. *CoRR*, abs/2101.11517.

Mannino, C.; Bernt, F.; and Dahl, G. 2012. Computing Optimal Recovery Policies for Financial Markets. *Operations Research*, 60.

Mounsaveng, S.; Laradji, I. H.; Ayed, I. B.; Vázquez, D.; and Pedersoli, M. 2020. Learning Data Augmentation with Online Bilevel Optimization for Image Classification. *CoRR*, abs/2006.14699.

Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and de Freitas, N. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1): 148–175.

Sinha, A.; Malo, P.; Deb, K.; Korhonen, P.; and Wallenius, J. 2016. Solving Bilevel Multicriterion Optimization Problems With Lower Level Decision Uncertainty. *IEEE Transactions on Evolutionary Computation*, 20(2): 199–217.

Stackelberg, H. v. 1952. *The theory of the market economy*. London: William Hodge.